



A MITEL
PRODUCT
GUIDE

Virtual Appliance Deployment Guide

March 2026

Notices

The information contained in this document is believed to be accurate in all respects but is not warranted by **Mitel Networks Corporation (MITEL®)**. The information is subject to change without notice and should not be construed in any way as a commitment by Mitel or any of its affiliates or subsidiaries. Mitel and its affiliates and subsidiaries assume no responsibility for any errors or omissions in this document. Revisions of this document or new editions of it may be issued to incorporate such changes. No part of this document can be reproduced or transmitted in any form or by any means - electronic or mechanical - for any purpose without written permission from Mitel Networks Corporation.

Trademarks

The trademarks, service marks, logos and graphics (collectively "Trademarks") appearing on Mitel's Internet sites or in its publications are registered and unregistered trademarks of Mitel Networks Corporation (MNC), its affiliates, parents, or subsidiaries (collectively "Mitel") or others. Use of the Trademarks is prohibited without the express consent from Mitel. Please contact our legal department at legal@mitel.com for additional information. For a list of the worldwide Mitel Networks Corporation registered trademarks, please refer to the website: <http://www.mitel.com/trademarks>.

®,™ Trademark of Mitel Networks Corporation

© Copyright 2026, Mitel Networks Corporation

All rights reserved

Contents

1 Introduction.....	1
1.1 Overview.....	1
1.1.1 Benefits of Virtualization.....	3
1.2 Software Requirements.....	5
1.2.1 VMware Software Requirements.....	5
1.2.2 Hyper-V Software Requirements.....	5
1.2.3 Mitel Product Version Support for VMware, Hyper-V, Nutanix, and Proxmox.....	6
1.2.4 Hardware Requirements.....	9
1.2.5 Documentation and Resources.....	10
2 Deploying On A VMware Infrastructure.....	13
2.1 Deploying Mitel Virtual Applications on a VMware Virtual Infrastructure.....	13
2.2 Architecture and Topology of Mitel Virtual Solutions.....	13
2.3 Deployment considerations.....	15
2.3.1 Deployment types.....	15
2.3.2 Preparing the VMware Infrastructure.....	16
2.4 VMware Support by Mitel Product Release.....	19
2.5 Resource Dimensioning.....	21
2.5.1 Using Resource Pools.....	21
2.6 Host Server Sizing and Resource Management Considerations.....	22
2.6.1 Host Server Sizing for Mitel Virtual Appliances.....	22
2.6.2 Mitel OEM resourcing information for VMware.....	22
2.6.3 Increasing Resources for Larger Dimensions.....	26
2.6.4 CPU and Memory Resource Reservation.....	27
2.6.5 Resource Sharing Guidelines.....	27
2.6.6 Storage Provisioning.....	29
2.6.7 Other Virtual Appliance Management Considerations.....	29
2.7 VMware Features.....	31
2.8 Resiliency, High Availability, and Disaster Recovery.....	35
2.8.1 System Resiliency without VMware High Availability Features.....	35
2.8.2 System Resiliency with VMware High Availability Features.....	36
2.8.3 VMware High Availability (HA).....	40
2.8.4 vSphere Replication.....	44
2.8.5 VMware Fault Tolerance.....	44
2.8.6 VMware Distributed Resource Scheduler (DRS) and vMotion.....	45
3 Deploying on a Nutanix Cluster Running AHV Hypervisor.....	46
3.1 Deploying on a Nutanix Cluster Running AHV Hypervisor.....	46
3.2 Deployment Considerations for Nutanix Running AHV Hypervisor.....	46
3.2.1 Deployment types for Nutanix Cluster.....	46
3.2.2 Preparing the Nutanix Infrastructure.....	47
3.3 Prism Central Support for Nutanix by Mitel Product Release.....	47
3.4 Host Server Sizing for Mitel Virtual Appliances.....	48

3.4.1 Increasing Resources for Larger Dimensions.....	48
3.4.2 Resource Dimensioning.....	49
3.4.3 Other Virtual Appliance Management Considerations.....	50
3.5 Resource Sharing Guidelines.....	51
3.5.1 Sharing with Third-Party Virtual Appliances.....	52
3.5.2 Enhanced Instance Density on Virtualization Platforms.....	52
3.5.3 Sharing with Other Mitel Virtual Appliances.....	52

4 Deploying on a Nutanix Cluster Running ESXi Hypervisor..... 54

4.1 Deploying on a Nutanix Cluster Running ESXi Hypervisor.....	54
4.2 Architecture and Topology of Mitel Virtual Solutions.....	54
4.3 Deployment Considerations for Nutanix Running ESXi Hypervisor.....	55
4.3.1 Deployment types for Nutanix Cluster.....	55
4.3.2 Preparing the Nutanix Infrastructure.....	56
4.4 vSphere ESXi Support for Nutanix by Mitel Product Release.....	56
4.5 Host Server Sizing for Mitel Virtual Appliances.....	57
4.5.1 Increasing Resources for Larger Dimensions.....	57
4.5.2 Resource Dimensioning.....	57
4.5.3 Resource Sharing Guidelines.....	59
4.5.4 Other Virtual Appliance Management Considerations.....	61

5 Deploying on Microsoft Hyper-V Infrastructure..... 63

5.1 Deploying Mitel Virtual Applications in a Hyper-V Virtual Infrastructure.....	63
5.2 Preparing the Hyper-V infrastructure.....	64
5.2.1 Adding NICs.....	64
5.2.2 Storage.....	65
5.2.3 Host Server Sizing and Resource Management Considerations.....	65
5.2.4 Resource Dimensioning.....	67
5.2.5 Resource Sharing Guidelines.....	68
5.2.6 Clock Sync.....	68
5.2.7 Keep Host NIC Driver up-to-date.....	68
5.2.8 Product-specific Guidelines.....	68
5.3 Installing and upgrading on Hyper-V.....	69
5.3.1 Installing Mitel Virtual Applications on Hyper-V Virtual Machines.....	70
5.4 Upgrading Mitel Applications on Hyper-V.....	72
5.4.1 Upgrade Backup/restore Variants.....	72
5.4.2 Supported Upgrade Methods.....	72
5.4.3 Hyper-V software and versions.....	73

6 Deploying on a Proxmox Infrastructure..... 75

6.1 Deployment Considerations for Proxmox.....	75
6.1.1 Deployment types for Proxmox.....	75
6.1.2 Preparing the Proxmox Infrastructure.....	76
6.1.3 Support for Proxmox by Mitel Product Release.....	76
6.2 Host Server Sizing for Mitel Virtual Appliances.....	76
6.2.1 Resource Sharing Guidelines.....	77
6.2.2 Increasing Resources for Larger Dimensions.....	78
6.2.3 Resource Dimensioning.....	79
6.2.4 Other Virtual Appliance Management Considerations.....	80

7	Deployment on Microsoft Azure	82
7.1	Deployment Considerations for Microsoft Azure	82
7.2	Host Server Sizing and Resource Management for Azure	83
7.2.1	Host Server Sizing for Mitel Virtual Appliances	83
7.2.2	Mitel OEM resourcing information for Azure	83
7.3	Resource and Infrastructure considerations (Azure)	87
8	Deployment on Amazon Web Services (AWS)	89
8.1	Deployment Considerations for AWS	89
8.2	Host Server Sizing and Resource Management for AWS	90
8.2.1	Host Server Sizing for Mitel Virtual Appliances	90
8.2.2	Mitel OEM resourcing information for AWS	90
8.3	Resource and infrastructure Considerations (AWS)	94
9	Install and Upgrade	96
9.1	Deploying Mitel Virtual Appliances in Hyper-V	96
9.2	Deploying Mitel Virtual Appliances in VMware	96
9.2.1	Disk Provisioning	96
9.2.2	Network driver	97
9.3	Deploying Mitel virtual appliances in Nutanix Cluster Running AHV Hypervisor	97
9.4	Deploying Mitel OVAs	98
9.5	Upgrading Mitel Virtual Appliances and Applications	112
9.5.1	Upgrading on Hyper-V	112
9.5.2	Upgrading on VMware	112
9.5.3	Software Upgrades	113
9.5.4	OVA Upgrade Advantages	114
9.5.5	MSL Application Blade Upgrade Advantages	114
9.5.6	VMware Virtual Hardware Versions	115
9.5.7	VMware software and versions	116
9.5.8	VMware Tools	121
9.5.9	Additional VMware Features	123
10	Maintenance and Troubleshooting	125
10.1	Maintenance	125
10.1.1	Backups	125
10.1.2	Restore	128
10.1.3	Migrate a Virtual Machine Between Hypervisors	129
10.1.4	Troubleshooting	130
10.1.5	Virtualization Diagnostic Tool (VDT)	134
10.1.6	Known Issues (VMware infrastructure)	145
11	Appendix A	153
11.1	Glossary	153
11.2	Deployment Types	166
12	Appendix B	167

12.1 VMware vSphere Storage Setup Guidelines.....	167
12.1.1 Storage Virtualization.....	167
12.1.2 Hard Disk Drive Types.....	168
12.1.3 Types of storage.....	168
12.1.4 RAID levels.....	170
12.1.5 Virtual disk modes.....	171
12.1.6 Virtual disk types.....	171
12.1.7 Partition alignment.....	172
12.1.8 SAN multipathing.....	172
12.1.9 File systems.....	172
12.1.10 Recommendations.....	173
12.1.11 Other Considerations.....	174
12.2 Reference.....	176

13 Appendix C..... 177

13.1 VMWARE vCloud Director.....	177
13.1.1 vCloud Director Overview.....	177
13.1.2 How vCloud Director capabilities are used in Mitel products.....	180
13.2 Creating and Publishing Virtual Application Catalog Templates.....	180
13.2.1 Unpack and Publish OVA in Linux.....	181
13.2.2 Unpack and publish OVA in Microsoft® Windows®.....	182
13.2.3 Import VM from vSphere.....	183
13.2.4 General Settings for Catalog Template Items.....	185
13.3 Deploying and Configuring Virtual Applications from Catalogs.....	185
13.3.1 Deployment Differences Relative to vSphere.....	187
13.3.2 Post-deployment Configuration.....	187
13.4 Run-time Resource Behavior of Deployed Virtual Applications.....	188
13.4.1 Pay-As-You-Go Allocation Model.....	189
13.4.2 Allocation Pool Allocation Model.....	190
13.4.3 Reservation Pool Allocation Model.....	191
13.4.4 Effects of Other Underlying Resources.....	192
13.4.5 Modifying vApp Resourcing for Deployment Configuration.....	193
13.5 vCloud Director Lease and Lease Expiry.....	197
13.5.1 Lease Expiry and Mission-critical Applications.....	198
13.6 Un-deploying Virtual Applications.....	198
13.7 Networking and Security Considerations.....	198
13.7.1 Deployment in a Direct-connected, Flat Organization Network.....	199
13.7.2 Deployment in Direct-connected, Layered Organization Network.....	200
13.7.3 Deployment in NAT-routed Organization Network.....	202

14 Appendix D..... 204

14.1 Security in Mitel/VMware Virtual Networks.....	204
14.1.1 vCloud Networking and Security (vCNS) Overview.....	204
14.2 vCNS Security Topologies.....	206
14.2.1 General Best Practices.....	207
14.2.2 Topology 1: vCNS Based Private Cloud - MBG Virtual Parallel with vCNS.....	207
14.2.3 Topology 2: vCNS Based Private Cloud - MBG Virtual Behind vCNS.....	209
14.2.4 Topology 3: vCNS Based Private Cloud - MiVoice Business Express.....	211
14.2.5 Topology 4: vCNS Based Public Cloud.....	212
14.2.6 Topology 5: vCNS Based Public Cloud - MiVoice Business Express.....	215
14.2.7 Topology Advantages, Disadvantages, and Recommended Applications.....	216
14.2.8 Network Component Configuration by Topology Type.....	223

This chapter contains the following sections:

- [Overview](#)
- [Software Requirements](#)

1.1 Overview

This guide describes the equipment, network, and configuration considerations that apply when setting up Mitel virtual appliances on servers enabled with VMware® vSphere™ virtualization or Microsoft Hyper-V.

Note:

This guide provides information about VMware and Hyper-V applications and requirements that were valid for previous releases, and which may not be up-to-date. Always refer to your current VMware or Hyper-V documentation for the latest information.

Mitel currently offers the following virtual appliances.

Supported in VMware:

- MiVoice Business Virtual
- MiCollab Virtual
- MiVoice Border Gateway Virtual
- MiContact Center Office
- MiContact Center Business
- MiContact Center Enterprise
- MiVoice Call Recording
- Virtual Open Integration Gateway (vOIG)
- Virtual Contact Center (vCC), Virtual MiVoice Call Accounting, Virtual IVR, Virtual Contact Center Bundle
- MiVoice Business Express
- MiVoice Business Solution Virtual Instance
- MiCloud Management Portal (formerly Oria)
- MiCloud Management Gateway
- MiVoice Office 250 (PS-1 only)
- MiVoice MX-ONE
- MiVoice 5000
- MiVoice Office 400
- MiVoice Connect
- MiVoice Office Mobile Application

- Mitel Mass Notification
- Mitel CX

Note:

- MiCollab Client Virtual has been discontinued as a standalone product. The server function for all of the clients is now included in MiCollab.
- MiCollab Outlook Plugin and MiCollab Audio, Web and Video Conference sharing are not supported in the virtual environment.

Supported in Hyper-V:

- MiVoice Business Virtual
- MiCollab Virtual
- MiVoice Border Gateway Virtual
- MiContact Center Business
- MiVoice Business Express
- MiVoice 5000
- MiVoice Office 400
- MiVoice Connect
- MiVoice MX-ONE
- MiContact Center Enterprise
- MiVoice Business Solution Virtual Instance
- Mitel CX

Note:

MiCollab Client Virtual has been discontinued as a standalone product. The server function for all of the clients is now included in MiCollab.

Supported in Nutanix:

- MiVoice Business Virtual
- MiCollab Virtual
- MiVoice Border Gateway Virtual
- MiVoice MX-ONE
- MiContact Center Business
- MiVoice 5000
- MiVoice Business Solution Virtual Instance
- Mitel CX

Supported in Proxmox:

- MiVoice Business Virtual
- MiCollab Virtual

Introduction

- MiVoice Border Gateway Virtual
- MiVoice MX-ONE
- MiVoice 5000
- MiVoice Business Solution Virtual Instance

The following Mitel applications are available as VMware-Ready® applications that you can install in a virtual machine:

- MiContact Center (also available as a virtual appliance, as shown in the list above)
- MiVoice Business Reporting (also available as a virtual appliance, as shown in the list above)
- MiVoice Enterprise Manager

Related Documents

The deployment procedures for each product on the supported hypervisor are detailed in the product-specific documents available on the [Document Center](#). The table below provides the product names and corresponding documents that outline the deployment procedures

Product	Document
MiVoice Business	MiVoice Business Installation and Administration Guide for Industry Standard Servers (ISS) and vMiVB
MiVoice Border Gateway	MiVoice Border Gateway Installation and Maintenance Guide
MiCollab	MiCollab Installation and Maintenance Guide
MX-ONE	MiVoice MX-ONE CPI library
MiVoice 5000	MiVoice 5000 - Implementation Manual
MiVoice Business Solution Virtual Instance	MiVoice Business Solution Virtual Instance
Mitel CX	Mitel CX Deployment and BluePrint Guide

1.1.1 Benefits of Virtualization

Virtualization can offer significant benefits to a data center; they are described on the VMware and Hyper-V Web sites, and summarized here:

Reduced Capital Expenditure on Physical Servers

Dedicated physical servers or appliances are not required for each application, resulting in server consolidation and hardware savings when using a virtual infrastructure. In addition, multiple operating systems can run on a single server.

Reduced Operations and Maintenance Costs

Having fewer servers reduces time, effort, and cost for server management. Placing Mitel virtual appliances in the virtual infrastructure with other virtualized business applications provides further cost savings through integrated IT processes.

Reduced Power Consumption

Virtual data centers have inherent power savings due to the reduced number of servers.

In VMware deployments, VMware Distributed Power Management (DPM) and the vMotion feature reduce the number of servers running at any given time.

Improved Application Availability

VMware offers several high availability features. One example is the VMware High Availability feature, which can monitor server and virtual appliance health. In the case of a host server failure, the virtual appliances can be restarted on a different server, minimizing down time. When server maintenance is necessary, Mitel virtual appliances can be moved to another server using vMotion, avoiding a service outage.

Hyper-V availability features are not available for Mitel products at this time.

Integrated Business Continuity

By consolidating Mitel virtual appliances within the customer's virtual infrastructure, the disaster recovery features provided by VMware will be applied consistently to all of the critical services.

Hyper-V availability features are not available for Mitel products at this time.

Seamless Telephony Integration when using Mitel Virtual Appliances

Mitel's adoption of virtualization has allowed continued compatibility and the ability to integrate with existing customer network and telephony infrastructure, just as Mitel physical server-based applications such as the 3300 ICP do.

Faster System Deployment

Mitel virtual appliances are inherently simpler to deploy into an existing data center. Increasing communication capacity or adjusting overall workload is easier and faster than the equivalent operations in traditional hardware and dedicated servers.

Update Roll-back

VMware features provide methods to quickly and easily restore a virtual appliance to its exact state before an upgrade, should a problem arise during an application or virtual appliance upgrade. On a physical system, a failed upgrade usually requires rebuilding the system and performing a data restore.

Hyper-V features are not available for Mitel products at this time.

This Solutions Guide discusses configuration and performance at the general solution level. The planning and configuration guidelines discussed here apply to the deployment of all Mitel virtual products. Product installation and configuration procedures and engineering guidelines are described in the product documentation for each individual product. See Mitel [Document Center](#).

For more information about designing multi-node Mitel networks, both physical MiVoice Business-based and MiVoice Business Virtual-based, refer to the *Multi-Node Networking Solutions Guide*.

1.2 Software Requirements

Refer to Mitel's [Product Compatibility Matrix](#) (PCM) to verify software product versions supported for use in a Mitel virtualized data center.

To check the supported versions, complete the following steps:

1. Go to the [Mitel MiAccess](#) portal.
2. Under the left-menu, go to **Compatibility Matrix > View Product Compatibility Details**.
3. On the **Product Compatibility Matrix** page, select a product category from the **Product Type** drop-down.
4. Select the product and release number from the respective drop-down options.
5. Select **Third-party** from the **Compatible with** drop-down, and the specific product from the **All compatible products** drop-down to check Microsoft HyperV, Microsoft Windows, VMware, vCenter, vSphere, and EXS/EXSi.

1.2.1 VMware Software Requirements

Refer to Mitel's [Product Compatibility Matrix](#) (PCM) to check VMware software product versions supported for use in a Mitel virtualized data center.

Note:

1. VMware also provides a bundled product collection known as vCloud Suite which contains several of the specific products listed in the table (vSphere, vCloud Networking and Security, vCloud Director, Site Recovery Manager). vCloud Suite contains several products such as, vCenter Operations Manager, vCloud Automation Center, and so on. See <http://www.vmware.com/products/vCloud-suite>. Availability of features and products in the vCloud Suite lineup for use with Mitel products are as listed in the PCM.
2. VMware vSphere version 5.5 is no longer supported.

1.2.2 Hyper-V Software Requirements

Refer to Mitel's [Product Compatibility Matrix](#) (PCM) to check Hyper-V software product versions required for use in a Mitel virtualized data center.

1.2.3 Mitel Product Version Support for VMware, Hyper-V, Nutanix, and Proxmox

Following table shows the Mitel virtualized product versions supported in VMware, Hyper-V, Nutanix, and Proxmox.

Table 1: Mitel Product Version Support

Software	VMware supported	Hyper-V supported	Nutanix AOS supported	Proxmox VE
Current Releases				
MiVoice Business Virtual	10.5	10.5	10.5	10.5
MiVoice Connect	19.3 SP3	19.3 SP3	not supported	
MiVoice Border Gateway Virtual	12.3	12.3	12.3	12.3
MiVoice MX-ONE	7.4 or later	7.1 or later	7.8 or later	7.8 or later
MiVoice Office 400	7.1 SP1	7.1 SP1	not supported	
MiVoice 5000	8.3	8.3	8.3	8.3
MiCollab Virtual	10.2 SP1	10.2 SP1	10.2 SP1	10.2 SP1
Mitel Mass Notification	6.5	not supported	not supported	
MiVoice Office Mobile Application	6.1	not supported	not supported	
MiContact Center Office Virtual	6.2 SP1	6.2 SP1	not supported	

Software	VMware supported	Hyper-V supported	Nutanix AOS supported	Proxmox VE
MiContact Center Business Virtual	10.1	10.1	not supported	
Mitel CX for MiVB	1.1, 1.0	1.1, 1.0	1.1, 1.0	
Mitel CX for SIP	1.1, 1.0	1.1, 1.0	1.1, 1.0	
MiContact Center Enterprise	9.7 SP1	9.7 SP1	not supported	
MiVoice Call Recording	9.2 SP7	9.2 SP7	not supported	
Virtual Open Integration Gateway (vOIG)	4.2 SP4	4.2 SP4	not supported	
MiVoice Business Solution Virtual Instance	SVI R 1.0 (2312)	SVI R 1.0 (2312)	SVI R 1.0 (2312)	SVI R 2.0
MiVoice Business Express	8.0	8.0	not supported	
MiCloud Management Portal (Oria)	6.0	not supported	not supported	
MiCloud Management Gateway	5.0	not supported	not supported	
Previous Releases				
MiVoice Business Virtual / vMCD	all releases since 6.0 SP2	all releases since 7.02	all releases since 10.1	all releases 10.4

Software	VMware supported	Hyper-V supported	Nutanix AOS supported	Proxmox VE
MiVoice Connect	all releases since 19.1	all releases since 19.1	not supported	
MiVoice Border Gateway Virtual	all releases since 7.1	all releases since 8.0 SP1	all releases since 11.6	all releases since 12.2
MiVoice MX-ONE	all releases since 6.3	all releases since 7.1	all releases since 9.8	
MiVoice Office 400	all releases since 5.0	all releases since 6.0	not supported	
MiCollab Virtual / vMAS	all releases since 9.3	all releases since 9.3	all releases since 10.0	all releases since 10.1
MiCollab Client Virtual / vUCA	all releases since 9.3	all releases since 9.3	all releases since 10.0	all releases since 10.1
NuPoint Unified Messaging Virtual / vNPM	all releases since 6.0	all releases since 7.0	not supported	
MiContact Center Virtual (MiCC Virtual)	all releases since 7.1 SP1	all releases since 7.1 SP1	not supported	
MiContact Center Enterprise	all releases since 7.0	all releases since 9.2	not supported	
Virtual Customer Service Manager (vCSM)	all releases since 6.0	not supported	not supported	
Virtual Open Integration Gateway (vOIG)	all releases since 2.0	all releases since 4.1	not supported	

Software	VMware supported	Hyper-V supported	Nutanix AOS supported	Proxmox VE
MiContact Center Office Virtual	all releases since 6.1 SP1	not supported	not supported	
Oria	all releases since 5.0	not supported	not supported	
MiVoice Call Recording	all releases since 8.1 SP1	9.2 SP4	not supported	
MiVoice Business Express	all releases since 6.0	all releases since 7.1	not supported	

Note:

Refer to the Engineering Guidelines for each product; see [Table 7](#).

1.2.4 Hardware Requirements

VMware Hardware Requirements:

For the list of processors supported by VMware, refer to <http://www.vmware.com/resources/compatibility/search.php>.

Nutanix Hardware Requirements:

For the list of processors supported by Nutanix, refer to [Recommended Community Edition Hardware](#).

Proxmox Hardware Requirements:

For the list of processors supported by Proxmox, refer to [System Requirements](#).

Hyper-V Hardware Requirements:

Refer to the Hyper-V 2025 documentation for the hardware requirements [here](#).

At minimum, Mitel virtual appliances running on both VMware and Hyper-V virtual machines require:

- Currently shipping Intel Xeon E3v2 / E5 / E7 Server Class series processors with a minimum of 4 cores, 2 GHz, Hyper-threading enabled, and Extended Page Table support (<http://ark.intel.com/>)
- Legacy server technology may also be used but must meet the minimum requirements shown in the following table:

Table 2: Minimum Hardware Requirements

Hardware Supplier	Processor Series Required	Additional Requirements
Intel	Xeon® 55xx/56xx/65xx/75xx Series	<ul style="list-style-type: none"> • 2.26 GHz or better, with a minimum of 4 cores • Hyper-Threading must be enabled in the BIOS (in some cases, this may be called Logical Processor) • Intel VT¹ must be enabled in the BIOS • must use EPT (Extended Page Tables)
AMD	Opteron™ 2400 Series	<ul style="list-style-type: none"> • 2.4 GHz or better with 6 cores • including Rapid Virtualization Indexing technology
	Opteron 3200, 4200, 6200 Series	<ul style="list-style-type: none"> • Supported for vSphere 5.0+

Note:

- Intel VT consists of Intel VT-x in the processor, Intel VT-d in the I/O chip set, and Intel VT-c in the network adapter chip set. The Intel Xeon 5500 Series incorporates Intel VT as do many newer Intel processors. Refer to <http://ark.intel.com/VTLList.aspx>.
- Refer to the Release Notes for the Mitel product version you are using for an up-to-date list of supported processors.
- Some processors allow selection of power management modes in BIOS settings (Low Power vs. High Performance, for example). It is recommended that hosts used to support real-time sensitive applications such as MiVoice Business and MiVoice Border Gateway, have their BIOS setting configured to maximize performance. This reduces the risk of voice quality issues and missed interrupts, for example.

1.2.5 Documentation and Resources

The following table lists additional resources and documentation. Mitel documentation is available on Mitel OnLine.

Table 3: Documentation and Resources

Content	Document Title or Location
Detailed MiVoice Business Virtual engineering information	MiVoice Business Engineering Guidelines for Industry Standard Servers and MiVoice Business Virtual
MiVoice Business Virtual installation and administration	MiVoice Business Installation and Administration Guide for MiVoice Business Virtual
MiVoice Business Virtual known issues	Mitel Virtual Appliance Quick Reference Guide
MiVoice Business Solution Virtual Instance	MiVoice Business Solution Virtual Instance Solution Guide
MiVoice Business Express	MiVoice Business Express Deployment Guide
MiCollab	MiCollab Installation and Maintenance Guide
	MiCollab Client Administrator Guide
	MiCollab Audio, Web, and Video Configuration and Maintenance Manual
	MiCollab Engineering Guidelines
	MiCollab Client Engineering Guidelines
MiVoice Border Gateway Virtual installation and maintenance	MiVoice Border Gateway Installation and Maintenance Guide
MiContact Center Office	MiContact Center Office Technician's Guide
MiContact Center Business	MiContact Center Business Deployment Guide
Mitel CX	Mitel CX Deployment Guide

Content	Document Title or Location
MiVoice MX-ONE	MiVoice MX-ONE Virtualization, Description
MiCloud Management Portal (Oria) deployment	MiCloud Management Portal Installation and Maintenance Guide MiCloud Management Portal Engineering Guidelines
The main VMware documentation Web page	http://www.vmware.com/support/pubs/
The main VMware vSphere Web page, with links to documentation, videos, and release notes	http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html
The VMware Hardware Compatibility guide	http://www.vmware.com/resources/compatibility/search.php
vSphere resource management	VMware vSphere Resource Management Guide
vSphere performance best practices	Performance Best Practices for VMware vSphere
Microsoft Hyper-V documentation	Microsoft Technet Library
Mitel Mass Notification	Mitel Mass Notification Administration Guide
MiContact Center Enterprise	Virtualization Description, MiCC Enterprise

Deploying On A VMware Infrastructure

2

This chapter contains the following sections:

- [Deploying Mitel Virtual Applications on a VMware Virtual Infrastructure](#)
- [Architecture and Topology of Mitel Virtual Solutions](#)
- [Deployment considerations](#)
- [VMware Support by Mitel Product Release](#)
- [Resource Dimensioning](#)
- [Host Server Sizing and Resource Management Considerations](#)
- [VMware Features](#)
- [Resiliency, High Availability, and Disaster Recovery](#)

2.1 Deploying Mitel Virtual Applications on a VMware Virtual Infrastructure

This chapter describes VMware deployment considerations specific to introducing Mitel virtual appliances and virtualized applications into your IT infrastructure, including:

- [Architecture and topology of Mitel virtual solutions](#)
- [Deployment considerations](#)
- [VMware features](#)
- [Resiliency, high availability, and disaster recovery](#)

2.2 Architecture and Topology of Mitel Virtual Solutions

The following figure shows the Mitel virtual appliances on the VMware infrastructure, distributed on hardware servers. Mitel virtual applications deliver capital cost savings associated with the reduction in server hardware and real estate; operational savings related to the reduction in power and server provisioning costs; and productivity improvements and resource efficiency in data center management and risk mitigation. Additionally, use of VMware availability features enable even higher levels of solution resiliency, business continuity and disaster recovery.

With virtualized voice, rather than having to handle voice communications with a separate budget and separate sets of hardware, processes, tools, and often staff, voice can be treated like any other business application in the data center. Rather than managing boxes, IT managers can manage the overall services that IT provides to the business, and in the process, reap the benefits and cost savings of a simplified test, development, and production cycle, streamlined administration, and a single disaster recovery and business continuity plan that applies to the whole data center.




IP1595

Figure 1: Mitel virtual applications on VMware Infrastructure

Mitel virtual appliances are bundled with Mitel Standard Linux (MSL) and packaged in Open Virtualization Format (OVF) for deployment in a VMware environment. Mitel virtual applications are delivered as OVA (.ova) files, ready to import into a VMware ESXi Hypervisor.

Given the potential complexity of your solution, and the time-sensitivity of the functions, it is important that the systems be configured to:

- Ensure that the Mitel products work optimally with each other and with your existing infrastructure.
- Ensure that the overall solution is optimized for use with the virtual (VMware) infrastructure.
- Ensure that the VMware servers have the capacity to deliver the standard of service you need.
- Increase the availability of your services and applications.



IP1594

Figure 2: MiVoice Business Virtual application layer resiliency with vSphere High Availability

This scenario consists of a Mitel Voice Over IP (VOIP) and Unified Communications (UC) virtual solution in a single data center, similar to the previous scenario. This scenario adds MiVoice Business Virtual controllers for application layer resiliency that ensures immediate voice recovery, and VMware High Availability (HA) to provide rapid recovery and return to the primary MiVoice Business Virtual, in the event the host or virtual machine fails.

MiVoice Border Gateway Virtual application layer resiliency (clustering) ensures immediate availability of voice connectivity to the service provider. A failed MiVoice Border Gateway Virtual recovers using HA, restoring full service quickly.

The resumption of voice connectivity occurs in preference to data network connectivity: SIP and PRI connectivity. In this example, PRI is the front-end to the virtual appliances with a physical Mitel 3300 ICP acting as the media gateway and connection to each of the data centers.

2.3 Deployment considerations

The following considerations will affect the design of your deployment:

- [Deployment types](#)
- [Preparing the VMware infrastructure](#)
- [VMware support by Mitel product release](#)
- [Host server sizing and resource management considerations](#)

2.3.1 Deployment types

There are three primary deployment types used for the installation of Mitel virtual applications:

Small and Medium-size Business (SMB)

SMB deployments of virtual applications are usually used in smaller businesses, for server consolidation. Management is generally not the primary issue, but running Mitel Unified Communication and Collaboration (UCC) applications on VMware virtual servers offers cost savings because fewer servers are required.

An SMB deployment is likely to have a small number of host servers, and possibly just a single stand-alone host. vCenter Server-related functions such as vMotion and HA virtualized environment availability features may not be present. Mitel applications generally coexist with non-Mitel applications in SMB environments.

Enterprise

Enterprise-level deployments are used either to fit the Mitel system into the existing virtual IT infrastructure, or to create a new virtual IT infrastructure specifically dedicated to Enterprise UCC functions.

An Enterprise-level deployment allows, and generally requires, the addition of VMware vCenter Server-based management, along with related functions such as vMotion and HA. Mitel virtual applications usually run on a subset of the available servers, to allow space for vMotion, HA, DRS, and DPM operations, for example, as required for the installation.

Enterprise deployments may be characterized as Private Cloud, effectively a hosted model, operated by Enterprise IT, using one of the deployment models discussed below.

Hosted

Rather than setting up a Mitel network on your premises with Mitel hardware and software, you can have your UCC and other applications hosted in a Service Provider data center, in a Public Cloud, for a monthly hosting charge or similar billing arrangement, using one of the deployment models discussed below.

This environment is likely to be a large-scale cloud cluster in which you lease capacity to run your virtual applications. Sophisticated management, including vCenter Server and related operations, are usually present, but not directly accessible by the customer.

The Mitel virtual UCC solutions are well-suited to the following deployment models:

- **Software as a Service (SaaS)** - Service Providers or Enterprise IT host the Mitel UCC solution as a set of software applications within a VMware vSphere shared server infrastructure, and offer it to customers as a service. There may also be other non-Mitel applications offered as part of the overall service. Customers, in this context, may be individual end-customers, Mitel resellers managing the UCC and other applications on behalf of the end-customer, or Enterprise IT (Private Cloud) providing service to organizations within the enterprise.
- **Unified Communications as a Service (UCaaS)** - Similar to SaaS, Service Providers or the Enterprise IT host the Mitel UCC solution as a set of applications within a VMware vSphere shared server infrastructure, and offer UCC service to customers. Customers may be individual end-customers, Mitel resellers managing the UCC applications on behalf of the end-customers, or Enterprise IT (Private Cloud) providing service to organizations within the enterprise.
- **Infrastructure as a Service (IaaS)** - Infrastructure providers lease out computing and network resources (for example: vCPU, GHz, RAM, HDD, and network access) required to host the Mitel UCC solution on their VMware vSphere shared infrastructure. The end customer, or a Mitel reseller acting on their behalf, must handle all aspects of installation, administration, management and maintenance of the UCC applications.
- **Customer Premise Equipment** - Mitel certified dealers or customers install and configure Mitel UCC virtual applications in the VMware environment on the customer's premise. The virtual infrastructure is directly managed by the customer. This is primarily suited to Enterprise and SMB type deployments.

2.3.2 Preparing the VMware Infrastructure

Regardless of the deployment environment, it is important that the overall VMware infrastructure is correctly configured to guarantee the highest level of availability and performance.

2.3.2.1 Host Servers and Clustering

When deciding on the host server computing environment, and planning how to set up one or more clusters of virtual machines, consider the following:

- Plan for adequate overall computing capacity across the cluster of host servers to accommodate immediate and foreseeable future needs for overall CPU and memory. This should include availability plans, if one or more hosts fail. See [Host server sizing and resource management considerations](#).
- Plan to deploy on multiple hosts of similar capacity and compatible processor type within each cluster so that you can enable Distributed Resource Scheduler (DRS), vMotion, High Availability (HA), Distributed Power Management (DPM), and other VMware availability and optimization features.

Note:

Mitel expects vMotion and Storage vMotion to function properly with Mitel virtual appliances in accordance with supported configuration from VMware. Similarly, SAN, VMware Virtual SAN (vSAN), and other storage connectivity is expected to work in VMware-supported configurations, subject to the I/O requirements listed in [Table 6](#). Refer to Mitel application-specific *Engineering Guidelines* documentation ([Table 7](#)) for vMotion host and storage performance requirements. If you want to use vMotion in a configuration that may be outside these specifications, contact Mitel Professional Services for assistance.

- Configure hosts into clusters, with VMware Distributed Resource Scheduler (DRS) enabled. DRS continuously monitors utilization across a resource pool, and allocates available resources among virtual machines according to business needs. See [Resiliency, high availability, and disaster recovery](#) for more information about DRS.
- Ensure that all hosts that will be used to run real-time sensitive applications, such as MiVoice Business Virtual, NuPoint UM Virtual, and MBG Virtual, are configured for maximum performance. This includes processor and BIOS settings requirements as listed in [Minimum hardware requirements](#).

2.3.2.2 Networking

Follow VMware best practices for configuration of networking capacity and fault tolerance. Mitel recommends the following minimum guidelines:

- Supply plenty of bandwidth: 1 Gbit/s across each of at least four physical ports on each host server is considered the minimum for most deployments.
- Separate the Virtual Local Area Network (VLAN) and IP interfaces for storage traffic (assuming Storage Area Network (SAN), VMware Virtual SAN (vSAN), Network File System (NFS), or similar IP storage is used) on a separate VLAN or subnet. This provides adequate throughput for storage operations, while isolating traffic from other flows. Preferably, as per VMware best practices, you should configure storage networking on two or more physical ports on the host servers, directed at different network switches to provide resiliency. You should also consider configuring these ports for NIC bonding (using Link Aggregation Control Protocol (LACP) or similar) to improve overall throughput and/or resiliency.
- If VMware vMotion features such as Distributed Resource Scheduler (DRS) and High Availability (HA) will be used, separate VLAN, IP interface, and multiple physical ports should be used for vMotion traffic.
- Voice traffic to and from the Mitel virtual appliances should be separate from data traffic. Implement a dedicated voice VLAN, and support it through multiple physical ports. The same recommendation applies for physical implementations.

2.3.2.3 Storage

In planning connections to, and stability of, the network storage, consider the following guidelines:

- For networked storage, ensure a robust Storage Network to support storage network traffic with minimal network delays.
- Ensure adequate throughput and IOPS to support all hosted applications in the storage device. Refer to the following documentation for detailed requirements.
- VMware vSphere storage setup guidelines for Mitel virtual appliances
- Product-specific documentation, for detailed requirements
- VMware documentation for storage set-up. VMware Publication: vSphere Storage (for version number)

- Deploy with support for replication, including remote replication if you plan to implement Site Recovery Manager (SRM), or a similar disaster recovery scheme.
- Deploy storage with multi-pathing, with a minimum of two physical ports on each host server and on each of the storage units, directed at different network switches. This provides resiliency.
- RAID protection of all storage is strongly recommended.

Mitel virtual appliances support for various storage architectures matches that of VMware vSphere, unless specifically stated otherwise. This includes iSCSI, NFS, Fibre Channel, VSAN, VSA (VMware and other vendors), and host-local storage. Host-local storage has a major limitation in that the virtual application cannot be moved to a different host server using vMotion for load balancing or maintenance, or protected by HA. For details, refer to the Mitel documentation for each virtual appliance, available at Mitel OnLine.

Note:

Regardless of the storage technology used, ensure that the storage used by Mitel applications meets or exceeds the specific IOPS requirements of the applications involved. See [Virtual Appliance Deployment hosting server sizing](#) website, and application-specific Engineering Guidelines.

In addition to Network Storage, Mitel supports on-board server storage that can be enabled with VMware vSphere Storage Appliance (VSA), and is ideal for smaller businesses that have a small data center and do not want to invest in expensive off-board network storage solutions. VSA allows for support of HA, DRS, and DPM with on-board server storage for up to three servers. (The three-host limit is imposed by the VMware implementation of VSA, as of vSphere 5.1.) vMCD Release 6.0 and MiVoice Business Virtual Release 7.0 support VSA.

Note:

VMware VSA may be lower performance than hardware-based high performance shared storage. If VMware VSA is to be used, ensure that VSA can meet the performance requirements of the applications using this as storage.

2.3.2.3.1 Virtual SAN (vSAN)

The VMware vSAN storage solution is fully integrated with vSphere. It automatically aggregates server disks in a cluster to create shared storage that can be rapidly provisioned from VMware vCenter during VM creation. It is an object-based storage system and a platform for VM Storage Policies designed to simplify virtual machine storage placement decisions for vSphere administrators. It is fully integrated with core vSphere features such as VMware vSphere High Availability (vSphere HA), VMware vSphere Distributed Resource Scheduler™ (vSphere DRS) and VMware vSphere vMotion®.

Its goal is to provide both high availability and scale-out storage functionality. It can also be considered in the context of quality of service (QoS) because VM Storage Policies can be created that define the level of performance and availability required on a per-virtual machine basis. See VMware vSphere storage setup guidelines for Mitel virtual appliances.

VMware recommends a 10 GbE network for this feature.

Table 4: Comparing VSA and vSAN

Attributes	vSphere Storage Appliance (VSA)	Virtual SAN (vSAN)
Description	Low cost, simple shared storage for small deployments.	Scalable, distributed storage designed for virtualized and cloud environments.
Form Factor	Virtual appliance	Built-in vSphere kernel
Ideal targets	<ul style="list-style-type: none"> • Small SMB • Remote Office/Branch Office (ROBO) deployments 	<ul style="list-style-type: none"> • Enterprise • Commercial
Scalability	<ul style="list-style-type: none"> • 2-3 vSphere servers • Does not scale beyond three hosts maximum 	<ul style="list-style-type: none"> • Minimum of three hosts in deployment • Scalable to vSphere cluster size
Performance	No Solid State Drive (SSD) - low performance	Solid State Drive (SSD) caching - high performance
Functionality	<ul style="list-style-type: none"> • Simple to install and configure • Scales up to about 16 TB of usable storage 	<ul style="list-style-type: none"> • vCenter-integrated management • SSD caching and intelligent data placement • Rapid storage provisioning • Scaling for large deployments • Granular scaling • Policy-based management

2.4 VMware Support by Mitel Product Release

Table 1, “VMware version support,” shows the vSphere releases that are supported by Mitel virtual appliance releases. The following table shows compatibility with the VMware tools and applications.

Table 5: VMware feature compatibility by Mitel product - latest Mitel releases

	Site Recovery Manager ^{1,2}	VMware Horizon	VMware Cloud Director	VMware NSX
Mitel Mass Notification	not compatible	not compatible	not compatible	not compatible
MiVoice Business 10.5	7.0, 6.7, 6.5	6.1, 5.3	5.5	5.1
MiVoice MX-ONE	not compatible	not compatible	not compatible	not compatible
MiCollab Client 10.2 SP1	7.0, 6.7	8.12, 8.11, 8.8, 7.11	5.5	5.1
MiVoice Border Gateway 12.3	8.0, 7.0, 6.7, 6.5	6.1 5.3	5.5 5.1	5.1
MiContact Center Office 6.2 SP1	not compatible	not compatible	not compatible	not compatible
MiCloud Management Portal (Oria) 6.0 Platform Manager	not compatible	not compatible	not compatible	not compatible
MiCloud Management Gateway 5.0	not compatible	6.1 5.3, 5.2	5.5, 5.1	5.5, 5.1

Notes:

1. MCDs and vMCDs using MCD clustering and IP Trunks are not compatible in SRM for versions earlier than MiVoice Business Release 7.0. Only SIP trunks and standalone vMCDs inter-operate with SRM. A deployment configuration has been tested with MiVoice Business Virtual that uses resiliency along with SRM for MiVoice Business Virtual Release 7.x.
2. After SRM recovery—that is, after the system becomes active on the recovery site—you must configure and test SRM at the recovery site to ensure the ability to move back to the original site or other alternate site. Then you can use SRM to move the system back to the original site.

2.5 Resource Dimensioning

In keeping with VMware best practices and Mitel requirements, you must have a minimum of 2 GB of memory allocated for the VMware Hypervisor (ESXi).

It is common practice to use resource over-subscription when provisioning your virtual machines. This is allowed, but you must observe the guidelines in the VMware best practices, and the Mitel maximum of 70% loading. This limit includes all of the Mitel appliances and all of the third-party virtual applications; the peak load for the total of all applications running on an individual host server must not exceed 70%.

Refer to the [VMware vSphere Resource Management Guide](#) and the [Performance Best Practices for VMware vSphere](#) on the VMware Web site.

CAUTION:

You must maintain CPU Headroom of 30% over the peak load to guarantee service level agreement (SLA) performance.

2.5.1 Using Resource Pools

In an advanced environment, you might choose to set up one or more resource pools. Resource pool level resource reservations can be used to set upper and lower limits on resources available to all contained VMs (or further child resource pools), so if you choose to set up a resource pool, you must still allocate enough CPU and RAM to meet the needs of all resource reservations of the contained VMs. The resource pool reservations must be greater than or equal to the sum of all reservations of the virtual appliances that will be deployed in that pool.

- Resource pools only apply to the reservation of CPU GHz. The published resource requirements for vCPU, RAM, IOPS and storage requirements remain at 100% utilization and cannot be part of the burst model.
- It must always be possible, during any burst period, to reach the full 100% CPU GHz determined for the reservations model.
- The Pool CPU GHz resources shall be 67%, or higher, of the CPU GHz determined for the reservations model¹
- The virtual machines (VM) in the pool shall be monitored for CPU Ready. The CPU Ready measurement indicates the wait time of an application to get resources, which can impact voice quality. The VM CPU Ready shall not exceed 40ms per vCPU. CPU Ready numbers that are regularly in this region indicate potential resource contention, and may be a precursor indicator of voice quality issues.
- Mission Critical deployments shall continue to use the reservation model and shall not be included in the reservation pool. Examples of a Mission Critical business include but are not limited to: Emergency Call or Dispatch centre (112, 911, 999, etc.), Hospital Emergency, life impacting activities or services (airport control tower, fire/ambulance/police, mountain rescue, etc.), etc.
- Contact Centre deployments shall continue to use the reservation model and shall not be included in the reservation pool. Contact centres are by their very nature traffic intensive, and may be considered business critical. Dedicated and reserved resources for contact centres and associated components should be considered separate from other less business critical components, such as back office staff that use components with pooled resources.

Note:

If an application requires 2vCPU and 2GHz CPU under the reservation model, under the pooled reservation/burst model, the settings will become 2vCPU and 1.33GHz. The number of vCPU does not change, only the pooled GHz resources. If, after deploying a reservation/burst model with resource pools, there are continuous voice quality issues, then this deployment should be transitioned to a reservation only model in order to guarantee peak load resource allocation.

2.6 Host Server Sizing and Resource Management Considerations

When setting up your environment in preparation for deployment of Mitel virtual appliances, you must consider host sizing and resource management.

The following section shows the resource allocation and resource reservation requirements for:

1. Mitel Virtual Appliances
2. Mitel OEM's

2.6.1 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

2.6.2 Mitel OEM resourcing information for VMware

This section lists additional third-party applications such as *Mitel Interaction Recording (MIR)*, *MiContact Center Outbound (MICC - O)*, *MiCC Speech*, and *MiCAM*. The following table provides the defined resource profiles and deployment details for these applications.

Note: *MiVoice Business* and *MiVoice Business with MPA* are available as defined resource profiles for VMware deployments. However, deployment with an integrated iCall Suite (also known as MBA) is currently not supported. For VMware environments, it is recommended to use an external server for iCall/ MBA integration.

Virtual Appliance	System Capacity	vCPU	Memory (GB)	Disk (GB)
Mitel Interaction Recording (MIR) Server A ¹	Small	4	16	210
	Medium	6	16	210
	Large	12	32	260
MIR Server B Audio ²	Small	4	4	120
	Medium	6	6	160
	Large	8	8	200
MIR Server B Screen ²	Small	4	4	120
	Medium	6	6	160
	Large	8	8	200
MIR Server C	Small	4	4	100
	Medium	6	6	100
	Large	8	8	100
MIR Server D ²	Small	4	4	120
	Medium	6	6	160
	Large	8	8	200
MIR Server E ³	Small	4	4	80

Virtual Appliance	System Capacity	vCPU	Memory (GB)	Disk (GB)
	Medium	6	6	80
	Large	8	8	80
MIR Server H ¹	Small	4	8	210
	Medium	6	16	210
	Large	12	32	260
MIR ⁴	MIR Archive	4	4	80
Speech Analysis (SA)	SA Transcription Core	2	8	120
	SA Transcription Decoder	8	52	120
MiContact Center-Outbound (MiCC-O) All-in-One	Small	8	12	190
	Medium	10	12	210
MICC-O Dialer	Small	4	8	120
	Mid-Small	6	8	150
	Medium	8	8	210
	Large	12	16	340
MICC-O NVP	Small	4	6	120

Virtual Appliance	System Capacity	vCPU	Memory (GB)	Disk (GB)
	Medium	6	6	150
	Large	8	8	230
MICC-O Web	Small	4	8	160
	Medium	6	8	190
	Large	8	12	320
MICC-O SQLSSRS	Small	4	14	190
MICC-O SQLPlus	Mid-Small	6	24	440
	Medium	6	30	640
	Large	12	54	1060
MICC Speech	ASR & Standard TTS 60	2	12	780
	ASR & Standard TTS 120	4	12	1240
	Enhanced TTS 60	6	32	780
	Enhanced TTS 120	12	64	780
MiCAM ⁵	Call Server Type B	6	16	160
	Call Server Type C	12	24	160

Virtual Appliance	System Capacity	vCPU	Memory (GB)	Disk (GB)
	System Server Type B	6	16	500
	System Server Type C	12	24	500
Management	MPA (Probe)	2	4	50
	Revolution	2	4	80

Key:**Note:**

1. HDD specification for database and storage needs to be added to this value. See application description for further details.
2. Default cache time included. Additional cache will require more HDD. See application description for further details.
3. Custom configuration. See application description for further details.
4. Recording storage in addition to server requirements. See application description for further details.
5. The HDD is split into partitions with the Operating System and ASR/TTS on one partition, and the MiCAM application and local cache/storage on the other partition.


2.6.3 Increasing Resources for Larger Dimensions

Note:

If you find that the deployed OVA is too small, you can use this procedure to increase the size for better performance.

To configure for larger dimensions:

1. From the VMware vSphere client, start the OVF deployment by specifying the location of the virtual OVA file.

2. Complete the steps in the deployment wizard and:
 - Select **Enterprise** as the **Configuration** in the **Deployment Configuration** section.
 - Ensure that the option to **Power on after deployment** is **NOT** selected.
 3. After the deployment is complete and the virtual machine is created, right-click the virtual machine and select **Edit Settings**.
 4. In the **Hardware** tab:
 - Select **CPUs** and change the **Number of virtual sockets** to the appropriate value.
 - Select **Memory** and change the **Memory Size** to the appropriate value.
 - Select **Hard disk 2** and change the **Provisioned Size** to the appropriate value.
-  **Note:** Refer [Virtual Appliance Deployment Server Sizing](#) website for the values.
5. In the **Resources** tab:
 - Select **CPU** and change the **Reservation** to the appropriate value.
 - Select the **Memory** and change the **Reservation** to the appropriate value.
 6. Apply the changes by clicking **OK**.
 7. Power on the virtual machine.

2.6.4 CPU and Memory Resource Reservation

Each Mitel virtual appliance is built with the resource reservation required for the performance Service Level Agreements already set. Do not change the resource reservation. If you choose to reduce the resource reservation, you do so at your own risk, and you must monitor performance carefully to ensure that you always retain some headroom.

CAUTION:

If you reduce your resource reservation, Mitel support will ask you to increase the resource level to the as-shipped value before providing support related to the service level agreements.

2.6.5 Resource Sharing Guidelines

One of the major advantages of virtualization is that you can co-locate multiple VMs on one host server.

Because most of Mitel's virtual applications provide real time voice, the customer will experience voice quality issues if insufficient CPU resources are allocated. Mitel builds in VMware resource reservation settings that will prevent resource limitation issues for typical customer installations.

VMware uses resource reservation to guarantee that the specified amount of the resource (CPU cycles and memory) is always available to the VM, regardless of the needs of other VMs sharing the resource.

The host CPU utilization must be maintained below 70% to prevent voice quality issues, as discussed in [Resource dimensioning](#).

Should performance or quality issues arise after reducing the resource reservation below the settings provided by Mitel, the first action Mitel Product Support will suggest is to restore the reservations to Mitel-mandated values.

2.6.5.1 Sharing with Third-Party Virtual Appliances

Mitel virtual appliances may be installed into the Virtual Infrastructure, where they will share resources with third-party virtual appliances and applications. A primary consideration in this scenario is ensuring that resource allocations for Mitel voice operations (MiVoice Business Virtual and other real-time voice applications) are adequate for good voice quality. Giving them too little resource reservation and priority may result in voice quality issues. These issues may be very transitory, based on what all the virtual appliances are doing at any given point in time. This can make troubleshooting very difficult.

There are two ways to ensure good performance of Mitel real-time applications; by setting a guaranteed resource reservation for the time-sensitive applications, or by giving each application a share-based percentage of the total CPU.

Setting a guaranteed resource reservation is the best way to ensure that enough CPU is allocated to the voice applications; a percentage-based allocation does not guarantee the required performance level in cases where the CPU is overloaded.

When co-locating Mitel and third-party applications, you must ensure that restart priorities are set correctly. MiVoice Business Virtual should be set to the highest priority. This sets the MiVoice Business Virtual VM ahead of other non-critical appliances in the event that a host fails. For example, you would set instances of MiVoice Business Virtual at the top priority, with MiCollab Virtual at a slightly lower priority, followed by your other business applications. Refer to [Resiliency, high availability, and disaster recovery](#) for more information.

Should performance or quality issues arise after you have reduced the resource reservation below the settings provided by Mitel, you will be asked to restore the resource reservations before support is offered for SLA issues. Maximum loading, including all Mitel and third-party virtual appliances, is 70%.

2.6.5.2 Increased Density with MiVoice Business Virtual Instances

In some cases where the MiVB instances are running less traffic and possibly less users than the standard deployments, it may be possible to tune and reduce the GHz reservations. This can improve instance density per server. For further details on settings, refer the MiVoice Business Engineering Guidelines (MiVB Release 8.0 onwards) and Engineering Guidelines for MiVoice Business for ISS and MiVoice Business Virtual for earlier releases.

If the installer is unfamiliar with determining and applying the resource settings, it is recommended to contact Mitel Professional Services.

2.6.5.3 Sharing with Other Mitel Virtual Appliances

Deploying multiple Mitel applications on one virtual server has advantages in that the various Mitel applications have similar requirements, and they are intended to work together. This generally simplifies resource considerations, and can result in lower maintenance costs.

When co-locating multiple instances of different Mitel virtual appliances, consider setting up a resource pool. In a resource pool, you can set priorities for resource allocation within the pool. For more information about resource pools, see [Using Resource Pools](#).

Minimum storage requirements are product specific, and different for every Mitel virtual appliance. For product-specific storage requirements, refer to the applicable Mitel Engineering Guide. For the I/O characterization specifics, see [Virtual Appliance Deployment Server Sizing Website](#).

2.6.6 Storage Provisioning

During deployment of the OVA file, you must choose the provisioning format. The choices are:

- **Thin provisioned format** - The VM storage is initially allocated based on the storage required at the time of deployment, and can expand as storage needs increase. This is more efficient in terms of total storage requirement, but can also be considerably slower as the storage grows and becomes more fragmented.
- **Thick provisioned format** - Causes immediate allocation of all storage required. This is less efficient in terms of total storage requirement (does not allow over-commit), but it guarantees optimum performance. Mitel virtual appliances that are sensitive to voice quality, including MiVoice Business Virtual, MiVoice Border Gateway, NuPoint Virtual, and MiVoice Business Express require thick provisioning to ensure good voice quality from the first call. For more information, refer to the product-specific documentation.

Thick provisioned format includes two additional options: Lazy-zeroed and Eager-zeroed. Mitel recommends using the Lazy-zeroed option. For more information about storage provisioning and these settings, see [Disk provisioning](#).

Note:

If you are using a storage technology with de-duplication capability, then the difference between thin and thick-provisioned formats is reduced, and you may find that Thin provisioning is sufficient. The Mitel recommendation is to use Thick provisioning for all Mitel virtual appliances.

Contact your storage provider for information about the over-commit capabilities of the storage option you are using.

2.6.7 Other Virtual Appliance Management Considerations

This section describes some additional items to consider when you deploy Mitel virtual appliances and applications.

2.6.7.1 Anti-affinity Rules

Some applications, especially when deployed in resiliency configurations, require that VMs for the applications do not reside on the same physical host server. Rules that ensure that certain VMs remain on different servers are called anti-affinity rules.

For example, a MiVoice Business Virtual primary controller VM should never be allowed to run on the same hardware as the corresponding MiVoice Business Virtual secondary controller VM. VMware vCenter anti-

affinity rules should be set to ensure that the primary and secondary are never moved to the same host server, including during automated actions such as vMotion or HA. See [Resiliency, high availability, and disaster recovery](#) and the VMware documentation for more information.

2.6.7.2 Hyper-Threading and Performance Impact

Intel's Hyper-Threading Technology makes one core appear as two logical processors, but the two logical processors share most of the core's resources. If one logical processor is running at high utilization, the other logical processor will have very little capacity that can be used before the performance of both logical processors suffers.

AMD Opteron processors use Rapid Virtualization Indexing, which helps accelerate the performance of many virtualized applications by enabling hardware-based virtual machine memory management.

ESXi assigns adjacent vCPU numbers to the logical processors on the same core. VMware attempts to keep the vCPUs of different VMs on separate cores. Use caution if you decide to use VMware advanced features to override VMware default behavior—assigning vCPUs to VMs. If two high-priority VMs have vCPUs on the same core, performance problems are likely.

CPU utilization shown in VMware management tools (Performance tab in vSphere Client, for example) can be misleading. If half the vCPUs are running at near 100% (on different cores) and the other half of the vCPUs are running near 0%, the system is actually fully loaded, although the average CPU may show 50% (the average of all the vCPUs).

Similarly, the GHz value for the system is really closer to the GHz sum of all the cores, not the sum of all the vCPUs.

With servers hosting Mitel virtual appliances that are less demanding—appliances that do not have to handle voice traffic—the assigned vCPU capacity may exceed the number of logical processors (twice the number of physical cores) by up to 50%. For example, a server with 8 physical cores provides 16 logical cores when hyperthreading is enabled. In this case, with non-voice traffic, you can provision 50% more vCPUs, for a total of 24 vCPUs.

Also refer the MiVoice Business Engineering Guidelines (MiVB Release 8.0) for MiVoice Business on ISS and MiVoice Business Virtual.

2.6.7.3 Using 3D Acceleration

By default, 3D acceleration is not enabled in ESXi 5.1 and later. You can enable 3D support on virtual machines that have the Windows 8+ guest operating system.

1. For the virtual machine you are enabling 3D acceleration on:
 - Verify that the virtual machine compatibility is ESXi 5.5+.
 - Verify that the virtual machine is powered off.
2. Right-click the virtual machine in the inventory.
3. Click **Edit Settings**.
4. In the left column of the **Virtual Hardware** tab, expand **Video Card**.
5. Select **Enable 3D Support**.

2.6.7.4 Product-specific Guidelines

For additional product-specific restrictions and requirements, see the *Mitel Engineering Guidelines* for each product.

Table 6: Engineering Guidelines documents for Mitel virtual applications

Mitel virtual appliance product	Engineering Guidelines
MiVoice Business Virtual	MiVoice Business Engineering Guidelines
MiCollab Virtual	MiCollab Engineering Guidelines
MiVoice Border Gateway Virtual	MiVoice Border Gateway Engineering Guidelines
MiCollab Client	MiCollab Client Engineering Guidelines
Virtual NuPoint Unified Messaging	NuPoint Engineering Guidelines; NuPoint Technician's Handbook.
MiContact Center Office	MiContact Center Office Technician's Handbook.
MiVoice Business Express	MiVoice Business Express Deployment Guide (includes Eng. Guidelines)
Mitel Open Integration Gateway	Mitel Open Integration Gateway Engineering Guidelines
MiCloud Management Portal (Oria)	MiCloud Management Portal Engineering Guidelines
MiVoice Office 400	MiVoice Office 400 Virtual Appliance System Manual

2.7 VMware Features

VMware offers many features and capabilities, many of which are compatible or applicable for Mitel virtual applications. The following table lists the basic VMware features, and specifies which are expected to work with Mitel virtual appliances. See also "[Mitel compatibility with VMware resiliency features](#),".

Note:

Mitel maintains support for current VMware vSphere/vCenter releases as outlined in this document. VMware also provides a number of infrastructure services that require the underlying VMware (vSphere / vCenter) such as SRM, vCNS, and vCloud Director that do not have any product integrations. As such, these VMware capabilities are supported by the customer and do not require direct product support.

Table 7: Basic VMware features

Feature	Compatible?	Notes
Virtual appliance deployment (Import)	Yes	Some virtual appliances may require a License Server hardware ID reset on re-deployment. Refer to the product-specific installation documentation and the Release Notes.
Export virtual appliance	Yes	
Power ON	Yes	The Power On starts a virtual application. To minimize downtime of Mitel virtual applications, set the Virtual Machine Startup and Shutdown setting to allow virtual machines to start and stop automatically with the system on a single server configuration. If several Mitel virtual appliances are configured to run on the same server, you may need to adjust the startup order based on the virtual appliance startup dependencies.
Shutdown Guest	Yes	Performs a graceful shutdown of the operating system.
Power OFF	Not recommended	This is the equivalent of cutting the power to a physical server. This can result in data corruption because the O/S is not shutdown cleanly.

Feature	Compatible?	Notes
Reset	Not recommended	This is the equivalent of performing a forced reset on a physical server. This can result in data corruption because the O/S is not shutdown cleanly.
Restart Guest	Yes	Restarts the virtual machine guest operating system gracefully.
Suspend/Resume	No	DO NOT suspend operation. This results in a loss of communication services and could cause database corruption. In addition, the remote end-points for Mitel virtual appliances that use VoIP connections will time out.
Cold Migration	Yes	Migration is the process of moving a powered-off virtual appliance from one host or storage location to another. Optionally, the user can relocate configuration and disk files to a new storage locations. Cold migration can be used to migrate virtual appliances from one data center to another. For MiContact Center Office, you must comply with Windows licensing restrictions.
Migration while suspended	No	Suspend/resume is not available.
Snapshot (Powered OFF)	Yes	Using a snapshot before an upgrade is recommended, as long as the snapshot is deleted after the upgrade is completed successfully.

Feature	Compatible?	Notes
Snapshot (Powered ON/Live)	No	The performance impact on Disk I/O is too large, making it impossible for Mitel virtual appliances to meet voice quality requirements.
Snapshot (Suspended)	Not recommended	Suspended snapshots are suggested only for limited maintenance periods.
Cloning	Yes	<p>The virtual appliance must be shut down before creating a clone.</p> <p>The clone must be registered with the License Server using a new Serial ID.</p> <p>Cloning is not available for configured MiVoice Business Express OVA files.</p>
Health Monitoring	Yes	This is the vSphere monitoring. (This not referring to the HA heartbeat mechanisms).
High Availability VM monitoring	Yes	VM Monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received within a set time. You can enable this feature and configure the sensitivity with which VMware HA monitors non-responsiveness.
High Availability Application monitoring	No	Application Monitoring is not available on Mitel virtual appliances.
Performance Reports	Yes	

Feature	Compatible?	Notes
Virtual SAN	Yes	vSAN automatically aggregates server disks in a cluster to create shared storage that can be rapidly provisioned from VMware vCenter during VM creation.
Virtual Flash	Yes	Virtual Flash allows you to accelerate virtual machine performance through the use of local SSD disks, which serve flash memory cache to chosen virtual machines running on the ESXi host.
vCloud Air	Yes	
Storage I/O Control	Yes	
Network I/O Control	Yes	

2.8 Resiliency, High Availability, and Disaster Recovery

Running your Mitel solution in a virtual environment allows even more options for high availability and fail-over. VMware availability-related features do not replace the underlying Mitel resiliency architecture (for those Mitel products that offer resiliency; MiVoice Business and MiVoice Border Gateway for example), but can supplement it to provide even higher availability levels, and improved ease of management within the virtualized environment.

For those Mitel products that do not offer application layer availability solutions, the VMware high availability (HA) features are highly recommended.

Whether or not you use VMware features for high availability, critical host servers must be arranged to ensure that one physical or electrical disruption cannot disable all servers in the cluster.

2.8.1 System Resiliency without VMware High Availability Features

Running your Mitel solution in a virtual environment, even without VMware high availability-related features enabled, allows options for high availability and fail-over because you can set up a physical primary MiVoice Business, with a virtual secondary for fail-over, or both primary and secondary MiVoice Business servers can be virtual (as long as you ensure that both primary and secondary do not end up on the same server; see [Anti-affinity rules](#) for more information).

MiVoice Border Gateway (MBG) also includes resiliency capabilities; supported MiNet sets can support a resiliency list of up to four IP addresses. If a set loses its connection and cannot re-establish it, it will try the next IP address on its list until it has exhausted all IP addresses on its resiliency list. For MiNet sets supporting persistent resiliency lists, this resiliency list is retained through power-cycling. The resiliency list can be manually configured with arbitrary IP addresses. Ideally, they will be configured to correspond to multiple MBG servers.

The following table shows the list of Mitel resiliency guides.

Table 8: Resiliency and high availability documentation

Product	Guide
MiVoice Business	MiVoice Resiliency Guidelines 3300 ICP Multi-Node Management Clustering
Mitel Automatic Call Distribution	ACD Resiliency Getting Started Guide
MiVoice Border Gateway (MBG)	MiVoice Border Gateway Engineering Guidelines

2.8.2 System Resiliency with VMware High Availability Features

Mitel products are compatible with many of the VMware high availability features, as shown in the following table. It is important to note that the resiliency capabilities afforded through the use of VMware high availability features are not a substitute for the Mitel resiliency features (depending on your requirements), but can be used in addition to the comprehensive capabilities designed into the Mitel products.

Note:

Mitel maintains support for current VMware vSphere/vCenter releases as outlined in this document. VMware also provides a number of infrastructure services that require the underlying VMware (vSphere / vCenter) such as SRM, vCNS, and vCloud Director that do not have any product integrations. As such, these VMware capabilities are supported by the customer and do not require direct product support.

Table 9: Mitel compatibility with VMware resiliency features

Function	Compatible?	Notes
vMotion ^{2,3}	Yes ¹	<p>While a vMotion operation is in progress, users may notice voice degradation, and some communication services may become unavailable for a brief time.</p> <p>For MiVoice Business Virtual, the voice degradation has been observed only in features provided by the Media Server. This includes conferencing, music-on--hold, paging, and RAD to external voice mail applications like NuPoint Unified Messaging.</p> <p>Ensure that anti-affinity rules are applied such that vMotion cannot move primary and secondary instances of a resilient Mitel application (MiVoice Business Virtual, MBG Virtual) to the same physical host server. For more information, see Anti-affinity rules.</p>

Function	Compatible?	Notes
Long Distance vMotion	Yes	<p>Movement of VMs from one vCenter server to another requires that both source and destination are running at least vCenter 6.0.</p> <ul style="list-style-type: none"> • When using the UI, source and destination vCenter servers must be in the same Single Sign-on (SSO) domain. • When using the API, source and destination vCenter servers can be in different Single Sign-on (SSO) domain. <p>To use vMotion across vSwitches, you should have L2 VM connectivity, and L3 connections between vCenter servers. 250 Mbps network bandwidth is required per vMotion operation.</p> <p>See also the limitations of vMotion.</p>
Storage vMotion	Yes ¹	<p>Storage vMotion (also called storage migration) may reduce storage performance, with the potential for degraded voice quality while the migration is in progress. Use of Storage vMotion is recommended only when the VA is powered off or during periods of low workload.</p>

Function	Compatible?	Notes
High Availability (HA)	Yes ^{1,3}	<p>Ensure that critical Mitel applications, such as MiVoice Business Virtual and MiVoice Border Gateway, are assigned high restart priorities.</p> <p>You must set anti-affinity rules such that clustered primary and secondary instances cannot run on the same host. See Anti-affinity rules.</p> <p>Note that MiVoice Business Express requires a large amount of CPU and memory reservation and it will not restart on a server unless the appropriate resources are available.</p>
Fault Tolerance (FT)	No	<p>FT resource requirements have a negative impact on performance of Mitel real-time vApps.</p> <p>FT also has strict requirements around set-up, the other VMware features that can be used with FT, and server BIOS settings, for example.</p>
Distributed Resource Scheduler (DRS)	Yes ¹	<p>vMotion is used for this service, so the vMotion cautions apply.</p> <p>Configuring this feature with the Aggressive setting may affect Mitel virtual appliance performance.</p>
Storage DRS	Yes ¹	
Distributed Power Management (DPM)	Yes ¹	vMotion is used for this service, so the vMotion cautions apply.

Function	Compatible?	Notes
vSphere Replication	Yes	vSphere Replication offers replication on a per-VM basis. vSphere Replication does not depend on Site Recovery Manager.
Site Recovery Manager (SRM)	Yes	Contact Mitel Professional Services.

i Note:

1. For MiContact Center Office, these features must comply with the Windows licensing restrictions.
2. Mitel expects vMotion and Storage vMotion to function properly with Mitel virtual appliances in accordance with supported configuration from VMware. Similarly, SAN or other storage connectivity is expected to work in VMware-supported configurations, subject to the I/O requirements listed. If you want to use vMotion in a configuration that may be outside these specifications, contact Mitel Professional Services for assistance.
3. VMware's vMotion is not supported with MiContact Center Business as it can degrade the performance of Ignite, IVR Routing, and other components.

i Note:

VMware features such as vMotion, Storage vMotion, HA, SRM, cloning, templates, and so on, do not require re-licensing of the application, because Mitel virtual application licensing Serial IDs are based on globally unique IDs (GUIDs) used by the Serial ID, and are not dependent on the specific hardware the application is running on.

2.8.3 VMware High Availability (HA)

The VMware High Availability (HA) feature enables monitoring of running VMs and host servers, and the ability to restart VMs elsewhere in the cluster in the case of host failure. HA provides protection for VMs at the cluster level, within a single data center.

HA can be a very effective supplement to Mitel built-in resiliency features. HA can protect against individual VMware host failures by restarting VMs that were running on the failed host, on a different host. While some down-time is involved for the VMs during the HA operation (as they are started on the new host), the inherent Mitel resiliency architecture provides zero downtime for the overall solution. Use of VMware HA does not take the place of Mitel resiliency operations, but supplements them by greatly reducing downtime of the nodes involved.

Note:

- In configuring any system supporting Mitel resilient applications (MiVoice Business Virtual and MiVoice Border Gateway, for example), it is extremely important to configure anti-affinity rules to ensure primary and corresponding secondary nodes are never allowed to run on the same physical host server. This would defeat the inherent Mitel resiliency operations. See [Anti-affinity rules](#).
- VMware's vMotion is not supported with MiContact Center Business as it can degrade the performance of Ignite, IVR Routing, and other components.




Figure 2: Resiliency using virtual MiVoice Business

In the example in the above figure, if MiVB Virtual A is a primary node deployed on host X, with HA protection to restart on some other host, and MiVB Virtual B is a secondary node deployed on host Y, also with HA protection, then the following failure scenario can be configured:

1. Host X fails. MiVB Virtual A is now out of service.
2. All IP Phones homed to MiVB Virtual A immediately re-home (fail over) to MiVB Virtual B on Host Y (their secondary) with near-zero loss of service (seconds to fail over, and no loss of active calls).
3. VMware HA then ensures that MiVB Virtual A is restarted on Host Z (anti-affinity rules apply so that MiVB Virtual A will not be restarted on the same host server as MiVB Virtual B), and during this time service is maintained by the secondary controller: MiVB Virtual B on Host Y.
4. When MiVB Virtual A is recovered a few minutes later, the IP Phones again re-home to MiVB Virtual A, now running on Host Z. In this case, HA protection of the primary node effectively reduces down time of the primary node to minutes.

In the physical world, it could take hours or even days to replace physical hardware running the primary node, and the devices would continue to run on the secondary for the duration. Not only is this non-optimal operation, it also adds the risk of additional down time, should node B also fail during the Node A repair time.

In the reverse scenario where failure happens on host Y, the secondary MiVB Virtual is restarted elsewhere, but with no loss of service since the IP Phones remain homed to their primary MiVB Virtual A.

HA protection can also be applied to Mitel virtual appliances that do not support inherent resiliency. While these applications are generally not considered mission critical, HA protection will allow the application to be monitored and automatically restarted elsewhere in the case of host failure. This will have the effect of reducing down time from hours or even days in the physical world to minutes in the virtual world.

When planning and configuring VMware HA protection, it is very important to consider the following:

- Anti-affinity rules must be put in place, such that primary and secondary nodes within a cluster cannot be run on the same physical host. This also applies in the event of HA operation restarting a node on a different host; the target host must also have the anti-affinity rules applied. See [Anti-affinity rules](#).
- Carefully consider HA restart priority settings for the protected virtual appliances. For example, it may be desirable to give primary nodes higher restart priority than secondary ones, or some application categories higher priority than others. For example, mission critical MiVoice Business Virtual controllers and MBG Virtual instances would be set for a higher restart priority than non-mission-critical applications.
- Plan the HA admission control policies to ensure adequate resources are reserved, in advance, to accommodate all high priority HA-protected applications, especially mission-critical Mitel virtual appliances. The admission control policies also affect the restart priority discussion above. See [Admission control for High Availability \(HA\) clusters](#) for recommendations for setting HA admission control policies.
- In circumstances where insufficient host resources exist after a failure, HA may override anti-affinity rules set at the Distributed Resource Scheduler (DRS) level. This may result, for example, in MiVoice Business Virtual primary and secondary nodes running on the same host until the failed host can be repaired. For a single pair of MiVoice Business Virtual controllers, a minimum of three host servers would be required to avoid this, each with sufficient guaranteed spare capacity to accept and run the failed-over MiVoice Business node.

2.8.3.1 Admission Control for High Availability (HA) Clusters

When configuring your high-availability VM cluster, you choose one of the following three admission policies:

- Host Failures Cluster Tolerates
- Percentage of Cluster Resources Reserved
- Specify a Failover Host

Alternatively, you can disable admission control if that makes sense for your implementation. If availability is more important than Service Level Agreement (SLA) performance upon failure, then, in VMware vSphere, setting Disable: Power on VMs that violate availability constraints may be worth considering. Keep in mind that there is a risk of voice quality degradation.

If you disable admission control, you must manually manage the number of VMs allowed to run on the hosts and/or reduce the VM reservations to ensure that the VMs will be able to restart on another host in the event of a host failure; this is the source of the voice quality risk. Enabling admission control is generally recommended.

The three available VMware admission control policies are described in the following sections.

2.8.3.1.1 Host Failures Cluster Tolerates

The Host Failures Cluster Tolerates admission control policy uses slot size calculations to determine whether enough capacity exists in the High Availability cluster to power on a virtual machine. This mode ensures HA for all the VMs in the cluster, and it is easy to configure, since VMware does all the resource calculations.

On the other hand, if a VM with a large resource reservation is introduced into the HA cluster, it can limit the number of VMs that can be powered on at one time. The large resource reservations set in the Mitel OVA files can change the slot size calculations for all the VMs in the cluster. Since most applications do not have an explicitly-set resource reservation, they work within the default 64 KB slot size, unless the slot size is changed by the introduction of an extra-large VM.

To avoid the potential problems associated with the cluster-wide slot size being changed to a large value by the introduction of Mitel appliances, you should set the maximum slot size manually, before adding the Mitel appliances. If you set a maximum, virtual machines larger than the set slot size are assigned multiple slots.

To reduce resource overhead due to the inefficiency in the slot size definition, you might consider creating separate clusters for different MiVoice Business Virtual and MBG Virtual deployments. For example, you could have a cluster with MiVoice Business instances dedicated and optimized for 150 users and another cluster optimized for 2500 users. The second cluster could use larger slot sizes.

2.8.3.1.2 Percentage of Cluster Resources Reserved

Using this percentage-based admission control policy is potentially more cost effective. To allow a trade-off between cost and availability SLA, for example, you can set the percentage to ensure availability only for critical VMs (defined using VM Restart Priority in Virtual Machine Options under the VMware HA configuration panel).

The flexibility in this option comes with some complexity, however, since you have to determine the appropriate percentage for the required SLA. In addition, the percentage value may require on-going adjustment depending on:

- The SLA level (how many VMs must be protected)
- Changes in available resources in the cluster
- How many simultaneous host failures must be accommodated

The complexity of determining, monitoring, and re-adjusting the percentage setting affects the ongoing cost of maintaining the system, and the risk of undershooting customer SLA requirements during failures (lowered voice quality or inability to support full MiVoice Business resiliency, for example).

Another risk in using the Percentage of Cluster Resources Reserved admission control policy is that, in a cluster running close to capacity, there might be enough aggregate resources for a virtual machine to be failed over, but the resources might be fragmented across different hosts. When that happens, they are not usable. A virtual machine can run on only one ESXi host at a time.

2.8.3.1.3 Specify a Failover Host

Using the Specify a Failover Host admission control policy, you specify one host to serve as a backup. It is kept idle so that it is always available to take over if another host fails. A couple of disadvantages to this admission strategy are:

- One dedicated host per cluster (the maximum allowed) is probably not sufficient in a large cluster.
- This resource is not used at all unless there is a failure, so under normal circumstances, the fail-over host resource is wasted.

For more information and detailed procedures, see the VMware vSphere Administration Guide and the VMware vSphere Availability Guide.

2.8.4 vSphere Replication

VMware vSphere Replication is a feature of the vSphere platform. It copies a virtual machine to another location, within or between clusters, and makes the new copy available for restoration through the vCenter Server Web-based user interface. It provides hypervisor-based virtual machine level replication and recovery from a primary site to a secondary site.

vSphere Replication continues to protect the virtual machine by replicating changes made to the virtual machine, to the copy. This ensures that the virtual machine remains protected and is available for recovery without requiring restore from backup. vSphere Replication is provided as a no-charge component of all eligible vSphere licenses.

Unified management of vSphere Replication is provided through the new vSphere Web Client. This provides a common and unified screen for all aspects of virtual data center management, including many aspects of protection, such as replication, backup and restore.

Note:

vSphere Replication is not intended to provide complex orchestration of disaster recovery in the manner of VMware Site Recovery Manager (SRM). vSphere Replication offers replication on a per-virtual machine basis to enable fine-grain recovery configurations for environments that need it.

There are two levels of replication, the storage level and the hypervisor level. vSphere (hypervisor level) replication is a part of SRM, and as a separately installed extension to vCenter Server.

2.8.5 VMware Fault Tolerance

The VMware Fault Tolerance (FT) feature provides real-time recovery of running virtual appliances, with no restart required, and very low interruption of service. FT provides run-time protection for VMs at the cluster level, within a single data center.

Mitel virtual appliances are not compatible with the Fault Tolerance feature (see [Table 9](#)). FT is very resource heavy, especially in the level of CPU resources consumed on multiple hosts, and the network load necessary to maintain synchronized state across the instances on those hosts negatively affects the performance of the Mitel real-time applications.

A combination of Mitel application level resiliency and VMware HA, as described previously, is recommended.

2.8.6 VMware Distributed Resource Scheduler (DRS) and vMotion

VMware Distributed Resource Scheduler (DRS) is a vMotion feature that enables dynamic load balancing across multiple hosts in a cluster, both at VM deployment time and, potentially, dynamically at run time.

While not directly related to resiliency and availability, DRS can play an important role in ensuring optimal performance during other operations. DRS ensures that VM resource reservations are honored, and can re-distribute loads to ensure overall optimal performance of all applicable VMs in the cluster. This can, in turn, help ensure optimal responsiveness and real-time behavior (voice quality) of Mitel virtual appliances.

vMotion can also be used to manually migrate running VMs to different hosts in the overall system, for example, in order to do maintenance of specific host servers (host software updates, trouble-shooting of host issues, hardware configuration, and so on). Migration to different hosts avoids downtime of the Mitel solution components during the host server maintenance operation, increasing overall availability. This is not feasible in the physical world.

Mitel strongly recommends enabling DRS at the cluster level where Mitel virtual appliances are deployed. To avoid the administrator having to monitor distribution behavior, the DRS automation level can be set to Fully Automated.

Note:

Since DRS uses vMotion to dynamically distribute loads at run time, and vMotion operation can cause momentary pauses in voice streaming, it is an administrator decision whether to enable full automation, and if enabled, what Migration threshold to set (a more aggressive migration threshold will result in higher chance of vMotion operating during active voice streaming).

When DRS is enabled for Mitel virtual appliances, such as MiVoice Business Virtual, MBG Virtual, and MiVoice Connect Virtual that implement their resiliency mechanisms, it is critical to set anti-affinity rules within DRS configuration to ensure that primary and secondary elements (for example, a primary MiVoice Business Virtual node and a secondary MiVoice Business Virtual node) are never allowed to run on the same physical host, since a single host failure could then defeat the inherent Mitel resiliency mechanisms. See also [Anti-affinity rules](#).

Deploying on a Nutanix Cluster Running AHV Hypervisor

3

This chapter contains the following sections:

- [Deploying on a Nutanix Cluster Running AHV Hypervisor](#)
- [Deployment Considerations for Nutanix Running AHV Hypervisor](#)
- [Prism Central Support for Nutanix by Mitel Product Release](#)
- [Host Server Sizing for Mitel Virtual Appliances](#)
- [Resource Sharing Guidelines](#)

3.1 Deploying on a Nutanix Cluster Running AHV Hypervisor

This chapter describes Nutanix deployment considerations specific to introducing Mitel virtual appliances and virtualized applications into your IT infrastructure.

3.2 Deployment Considerations for Nutanix Running AHV Hypervisor

The following considerations will affect the design of your deployment:

- [Deployment types for Nutanix Cluster](#)
- [Preparing the Nutanix infrastructure](#)
- [Prism Central support for Nutanix by Mitel product release](#)
- [Host server sizing and resource management considerations](#)

3.2.1 Deployment types for Nutanix Cluster

Following sections describe three primary deployment types used for the installation of Nutanix virtual application:

Small and Medium-size Business (SMB)

SMB deployments of virtual applications are usually used in smaller businesses, for server consolidation. Management is generally not the primary issue, but running Mitel Unified Communication and Collaboration (UCC) applications on Nutanix virtual servers offers cost savings because fewer servers are required.

An SMB deployment is likely to have a small number of host servers, and possibly just a single stand-alone host. Prism Central Server-related functions such as vMotion and HA virtualized environment availability

features may not be present. Mitel applications generally coexist with non-Mitel applications in SMB environments.

Enterprise

Enterprise-level deployments are used either to fit the Mitel system into the existing virtual IT infrastructure, or to create a new virtual IT infrastructure specifically dedicated to Enterprise UCC functions.

An Enterprise-level deployment allows, and generally requires, the addition of Prism Central Server-based management, along with related functions such as vMotion and HA. Mitel virtual applications usually run on a subset of the available servers, to allow space for vMotion, HA, DRS, and DPM operations, for example, as required for the installation.

Enterprise deployments may be characterized as Private Cloud, effectively a hosted model, operated by Enterprise IT, using one of the deployment models discussed below.

Hosted

The Mitel virtual UCC solutions are well-suited to the following deployment models:

- **Customer Premise Equipment** - Mitel certified dealers or customers install and configure Mitel UCC virtual applications in the Nutanix environment on the customer's premise. The virtual infrastructure is directly managed by the customer. This is primarily suited to Enterprise and SMB type deployments.

3.2.2 Preparing the Nutanix Infrastructure

Regardless of the deployment environment, it is important that the overall Nutanix infrastructure is correctly configured to guarantee the highest level of availability and performance. For more information about preparing the Nutanix cluster running Prism Central AHV Hypervisor, see the *Nutanix documentation*.

3.3 Prism Central Support for Nutanix by Mitel Product Release

[Prism Central AHV Hypervisor feature compatibility by Mitel product - latest Mitel releases](#) shows the Prism Central/Nutanix Web Console releases that are supported by Mitel virtual appliance releases. The following table shows compatibility with the Nutanix AHV Hypervisor tools and applications.

Table 10: Prism Central AHV Hypervisor feature compatibility by Mitel product - latest Mitel releases

Product	Nutanix AOS	Prism Central
MiVoice Business (10.5)	7.0	pc.2024.3.1
MiVoice Border Gateway (12.3)	7.0	pc.2024.3.1
MiCollab	6.5.4	pc.2022.6.0.9

Product	Nutanix AOS	Prism Central
MiContact Center Business 10.2	6.5.4	pc.2022.6.0.9
Mitel CX	7.0.1	pc.2024.3.1
MX-ONE	6.8.1	pc.2024.3.1
MiVoice 5000 (8.0 and later)	5.16 and later versions	5.16 and later versions
MiVoice Business Solution Virtual Instance SVI R1.0, R2.0	7.0	pc.2024.3.1

3.4 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

3.4.1 Increasing Resources for Larger Dimensions

Note:

If you find that the deployed ISO is too small, you can use this procedure to increase the size for better performance.

To configure for larger dimensions:

1. From the Prism Central client, start the ISO deployment by specifying the location of the virtual ISO file.
2. Complete the steps in the deployment wizard and:
 - Select **Enterprise** as the **Configuration** in the **Deployment Configuration** section.
 - Ensure that the option to **Power on after deployment** is **NOT** selected.
3. After the deployment is complete and the virtual machine is created, right-click the virtual machine and select **Edit Settings**.

4. In the **Hardware** tab:

- Select **CPUs** and change the **Number of virtual sockets** to the appropriate value in [Nutanix AOS/Prism Central required resource reservation and allocation](#).
- Select **Memory** and change the **Memory Size** to the appropriate value in [Nutanix AOS/Prism Central required resource reservation and allocation](#).
- Select **Hard disk 2** and change the **Provisioned Size** to the appropriate value in [Nutanix AOS/Prism Central required resource reservation and allocation](#).

5. In the **Resources** tab:

- Select **CPU** and change the **Reservation** to the appropriate value in [Nutanix AOS/Prism Central required resource reservation and allocation](#).
- Select the **Memory** and change the **Reservation** to the appropriate value in [Nutanix AOS/Prism Central required resource reservation and allocation](#).

6. Apply the changes by clicking **OK**.

7. Power on the virtual machine.

3.4.2 Resource Dimensioning

In keeping with Nutanix best practices and Mitel requirements, you must have a minimum of 2 GB of memory allocated for the Nutanix cluster running AHV Hypervisor.

It is common practice to use resource over-subscription when provisioning your virtual machines. This is allowed, but you must observe the guidelines in the Nutanix best practices, and the Mitel maximum of 70% loading. This limit includes all of the Mitel appliances and all of the third-party virtual applications; the peak load for the total of all applications running on an individual host server must not exceed 70%.

See the *Nutanix documentation* for more information.

CAUTION:

You must maintain CPU Headroom of 30% over the peak load to guarantee service level agreement (SLA) performance.

3.4.2.1 Using Resource Pools

In an advanced environment, you might choose to set up one or more resource pools. Resource pool level resource reservations can be used to set upper and lower limits on resources available to all contained VMs (or further child resource pools), so if you choose to set up a resource pool, you must still allocate enough CPU and RAM to meet the needs of all resource reservations of the contained VMs. The resource pool reservations must be greater than or equal to the sum of all reservations of the virtual appliances that will be deployed in that pool.

- Resource pools only apply to the reservation of CPU GHz. The published resource requirements for vCPU, RAM, IOPS and storage requirements remain at 100% utilization and cannot be part of the burst model.
- It must always be possible, during any burst period, to reach the full 100% CPU GHz determined for the reservations model.

- The Pool CPU GHz resources shall be 67%, or higher, of the CPU GHz determined for the reservations model¹
- The virtual machines (VM) in the pool shall be monitored for CPU Ready. The CPU Ready measurement indicates the wait time of an application to get resources, which can impact voice quality. The VM CPU Ready shall not exceed 40ms per vCPU. CPU Ready numbers that are regularly in this region indicate potential resource contention, and may be a precursor indicator of voice quality issues.
- Mission Critical deployments shall continue to use the reservation model and shall not be included in the reservation pool. Examples of a Mission Critical business include but are not limited to: Emergency Call or Dispatch center (112, 911, 999, etc.), Hospital Emergency, life impacting activities or services (airport control tower, fire/ambulance/police, mountain rescue, etc.), etc.
- Contact Center deployments shall continue to use the reservation model and shall not be included in the reservation pool. Contact centers are by their very nature traffic intensive, and may be considered business critical. Dedicated and reserved resources for contact centers and associated components should be considered separate from other less business critical components, such as back office staff that use components with pooled resources.

Note:

If an application requires 2vCPU and 2GHz CPU under the reservation model, under the pooled reservation/burst model, the settings will become 2vCPU and 1.33GHz. The number of vCPU does not change, only the pooled GHz resources. If, after deploying a reservation/burst model with resource pools, there are continuous voice quality issues, then this deployment should be transitioned to a reservation only model in order to guarantee peak load resource allocation.

3.4.3 Other Virtual Appliance Management Considerations

This section describes some additional items to consider when you deploy Mitel virtual appliances and applications.

3.4.3.1 Anti-affinity Rules

Some applications, especially when deployed in resiliency configurations, require that VMs for the applications do not reside on the same physical host server. Rules that ensure that certain VMs remain on different servers are called anti-affinity rules.

For example, a MiVoice Business Virtual primary controller VM should never be allowed to run on the same hardware as the corresponding MiVoice Business Virtual secondary controller VM. Prism Central anti-affinity rules should be set to ensure that the primary and secondary are never moved to the same host server, including during automated actions such as vMotion or HA.

3.4.3.2 Hyper-Threading and Performance Impact

Intel's Hyper-Threading Technology makes one core appear as two logical processors, but the two logical processors share most of the core's resources. If one logical processor is running at high utilization, the other logical processor will have very little capacity that can be used before the performance of both logical processors suffers.

AMD Opteron processors use Rapid Virtualization Indexing, which helps accelerate the performance of many virtualized applications by enabling hardware-based virtual machine memory management.

Prism Central assigns adjacent vCPU numbers to the logical processors on the same core. Nutanix attempts to keep the vCPUs of different VMs on separate cores. Use caution if you decide to use Nutanix advanced features to override Nutanix default behavior—assigning vCPUs to VMs. If two high-priority VMs have vCPUs on the same core, performance problems are likely.

CPU utilization shown in Nutanix management tools (Performance tab in Prism Central Client, for example) can be misleading. If half the vCPUs are running at near 100% (on different cores) and the other half of the vCPUs are running near 0%, the system is actually fully loaded, although the average CPU may show 50% (the average of all the vCPUs).

Similarly, the GHz value for the system is really closer to the GHz sum of all the cores, not the sum of all the vCPUs.

With servers hosting Mitel virtual appliances that are less demanding—appliances that do not have to handle voice traffic—the assigned vCPU capacity may exceed the number of logical processors (twice the number of physical cores) by up to 50%. For example, a server with 8 physical cores provides 16 logical cores when hyperthreading is enabled. In this case, with non-voice traffic, you can provision 50% more vCPUs, for a total of 24 vCPUs.

Also, see the *MiVoice Business Engineering Guidelines* (MiVoice Business Release 10.1) for MiVoice Business on ISS and MiVoice Business Virtual.

3.4.3.3 Product-specific Guidelines

For additional product-specific restrictions and requirements, see the *Mitel Engineering Guidelines* for each product.

Table 11: Engineering Guidelines documents for Mitel virtual applications

Mitel virtual appliance product	Engineering Guidelines
MiVoice Business Virtual	MiVoice Business Engineering Guidelines
MiVoice Border Gateway Virtual	MiVoice Border Gateway Engineering Guidelines
MiCollab Client	MiCollab Client Engineering Guidelines

3.5 Resource Sharing Guidelines

One of the major advantages of virtualization is that you can co-locate multiple VMs on one host server.

Because most of Mitel's virtual applications provide real time voice, the customer will experience voice quality issues if insufficient CPU resources are allocated. Mitel builds in Nutanix resource reservation settings that will prevent resource limitation issues for typical customer installations.

Nutanix uses resource reservation to guarantee that the specified amount of the resource (CPU cycles and memory) is always available to the VM, regardless of the needs of other VMs sharing the resource.

The host CPU utilization must be maintained below 70% to prevent voice quality issues, as discussed in [Resource dimensioning](#).

Should performance or quality issues arise after reducing the resource reservation below the settings provided by Mitel, the first action Mitel Product Support will suggest is to restore the reservations to Mitel-mandated values.

3.5.1 Sharing with Third-Party Virtual Appliances

Mitel virtual appliances may be installed into the Virtual Infrastructure, where they will share resources with third-party virtual appliances and applications. A primary consideration in this scenario is ensuring that resource allocations for Mitel voice operations (MiVoice Business Virtual and other real-time voice applications) are adequate for good voice quality. Giving them too little resource reservation and priority may result in voice quality issues. These issues may be very transitory, based on what all the virtual appliances are doing at any given point in time. This can make troubleshooting very difficult.

There are two ways to ensure good performance of Mitel real-time applications; by setting a guaranteed resource reservation for the time-sensitive applications, or by giving each application a share-based percentage of the total CPU.

Setting a guaranteed resource reservation is the best way to ensure that enough CPU is allocated to the voice applications; a percentage-based allocation does not guarantee the required performance level in cases where the CPU is overloaded.

When co-locating Mitel and third-party applications, you must ensure that restart priorities are set correctly. The virtualization platforms should be set to the highest priority. This sets the Virtual VM ahead of other non-critical appliances in the event that a host fails. For example, you would set instances of MiVoice Business Virtual at the top priority, with MiCollab Virtual at a slightly lower priority, followed by your other business applications.

Should performance or quality issues arise after you have reduced the resource reservation below the settings provided by Mitel, you will be asked to restore the resource reservations before support is offered for SLA issues. Maximum loading, including all Mitel and third-party virtual appliances, is 70%.

3.5.2 Enhanced Instance Density on Virtualization Platforms

In some cases, if the instances are running less traffic and possibly fewer users than the standard deployments, it might be possible to tune and reduce the GHz reservations. This can improve instance density per server.

If the installer is unfamiliar with determining and applying the resource settings, it is recommended to contact Mitel Professional Services.

3.5.3 Sharing with Other Mitel Virtual Appliances

Deploying multiple Mitel applications on one virtual server has advantages in that the various Mitel applications have similar requirements, and they are intended to work together. This generally simplifies resource considerations, and can result in lower maintenance costs.

When co-locating multiple instances of different Mitel virtual appliances, consider setting up a resource pool. In a resource pool, you can set priorities for resource allocation within the pool. For more information about resource pools, see [Using Resource Pools](#).

Minimum storage requirements are product-specific, and different for every Mitel virtual appliance. For product-specific storage requirements, refer to the applicable Mitel Engineering Guide. For the I/O characterization specifics, see [Nutanix AOS/Prism Central required resource reservation and allocation](#).

Deploying on a Nutanix Cluster Running ESXi Hypervisor

4

This chapter contains the following sections:

- [Deploying on a Nutanix Cluster Running ESXi Hypervisor](#)
- [Architecture and Topology of Mitel Virtual Solutions](#)
- [Deployment Considerations for Nutanix Running ESXi Hypervisor](#)
- [vSphere ESXi Support for Nutanix by Mitel Product Release](#)
- [Host Server Sizing for Mitel Virtual Appliances](#)

4.1 Deploying on a Nutanix Cluster Running ESXi Hypervisor

This chapter describes Nutanix deployment considerations specific to introducing Mitel virtual appliances and virtualized applications into your IT infrastructure, including:

- [Architecture and topology of Mitel virtual solutions](#)
- [Deployment considerations for Nutanix](#)

4.2 Architecture and Topology of Mitel Virtual Solutions

The following figure shows the Mitel virtual appliances on the vSphere infrastructure, distributed on hardware servers. Mitel virtual applications deliver capital cost savings associated with the reduction in server hardware and real estate; operational savings related to the reduction in power and server provisioning costs; and productivity improvements and resource efficiencies in data center management and risk mitigation. Additionally, use of Nutanix availability features enable even higher levels of solution resiliency, business continuity and disaster recovery.

With virtualized voice, rather than having to handle voice communications with a separate budget and separate sets of hardware, processes, tools, and often staff, voice can be treated like any other business application in the data center. Rather than managing boxes, IT managers can manage the overall services that IT provides to the business, and in the process, reap the benefits and cost savings of a simplified test, development, and production cycle, streamlined administration, and a single disaster recovery and business continuity plan that applies to the whole data center.

Mitel virtual appliances are bundled with Mitel Standard Linux (MSL) and packaged in Open Virtualization Format (OVF) for deployment in a Nutanix running on ESXi Hypervisor environment. Mitel virtual applications are delivered as ISO (.iso) files, ready to import into a Nutanix running on ESXi Hypervisor.

Given the potential complexity of your solution, and the time-sensitivity of the functions, it is important that the systems be configured to:

- Ensure that the Mitel products work optimally with each other and with your existing infrastructure.

- Ensure that the overall solution is optimized for use with the virtual (Nutanix) infrastructure.
- Ensure that the Nutanix servers have the capacity to deliver the standard of service you need.
- Increase the availability of your services and applications.

This MiVoice Business Virtual application layer resiliency with vSphere High Availability scenario consists of a Mitel Voice Over IP (VOIP) and Unified Communications (UC) virtual solution in a single data center, similar to the previous scenario. This scenario adds MiVoice Business Virtual controllers for application layer resiliency that ensures immediate voice recovery, and Nutanix High Availability (HA) to provide rapid recovery and return to the primary MiVoice Business Virtual, in the event the host or virtual machine fails.

MiVoice Border Gateway Virtual application layer resiliency (clustering) ensures immediate availability of voice connectivity to the service provider. A failed MiVoice Border Gateway Virtual recovers using HA, restoring full service quickly.

The resumption of voice connectivity occurs in preference to data network connectivity: SIP and PRI connectivity. In this example, PRI is the front-end to the virtual appliances with a physical Mitel 3300 ICP acting as the media gateway and connection to each of the data centers.

4.3 Deployment Considerations for Nutanix Running ESXi Hypervisor

The following considerations will affect the design of your deployment:

- [Deployment types for Nutanix Cluster](#)
- [Preparing the Nutanix infrastructure](#)
- [vSphere ESXi support for Nutanix by Mitel product release](#)
- [Host server sizing and resource management considerations](#)

4.3.1 Deployment types for Nutanix Cluster

Following sections describe three primary deployment types used for the installation of Nutanix virtual application:

Small and Medium-size Business (SMB)

SMB deployments of virtual applications are usually used in smaller businesses, for server consolidation. Management is generally not the primary issue, but running Mitel Unified Communication and Collaboration (UCC) applications on Nutanix virtual servers offers cost savings because fewer servers are required.

An SMB deployment is likely to have a small number of host servers, and possibly just a single stand-alone host. vCenter Server-related functions such as vMotion and HA virtualized environment availability features may not be present. Mitel applications generally coexist with non-Mitel applications in SMB environments.

Enterprise

Enterprise-level deployments are used either to fit the Mitel system into the existing virtual IT infrastructure, or to create a new virtual IT infrastructure specifically dedicated to Enterprise UCC functions.

An Enterprise-level deployment allows, and generally requires, the addition of vSphere Server-based management, along with related functions such as vMotion and HA. Mitel virtual applications usually run on a subset of the available servers, to allow space for vMotion, HA, DRS, and DPM operations, for example, as required for the installation.

Enterprise deployments may be characterized as Private Cloud, effectively a hosted model, operated by Enterprise IT, using one of the deployment models discussed below.

Hosted

The Mitel virtual UCC solutions are well-suited to the following deployment models:

- **Customer Premise Equipment** - Mitel certified dealers or customers install and configure Mitel UCC virtual applications in the Nutanix environment on the customer's premise. The virtual infrastructure is directly managed by the customer. This is primarily suited to Enterprise and SMB type deployments.

4.3.2 Preparing the Nutanix Infrastructure

Regardless of the deployment environment, it is important that the overall Nutanix infrastructure is correctly configured to guarantee the highest level of availability and performance. For more information about preparing the Nutanix cluster running ESXi Hypervisor, see the *Nutanix documentation*.

4.4 vSphere ESXi Support for Nutanix by Mitel Product Release

"[VMware version support](#)," shows the vSphere releases that are supported by Mitel virtual appliance releases. The following table shows compatibility with the vSphere ESXi tools and applications.

Table 12: vSphere ESXi feature compatibility by Mitel product - latest Mitel releases

	Nutanix AOS	ESXi
MiVoice Business (10.5)	7.0	8.0
MiVoice Border Gateway (12.3)	7.0	8.0
MiCollab 10.2	6.5.4	8.0
MiVoice Business Solution Virtual Instance 1.0, 2.0	7.0	8.0

4.5 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

4.5.1 Increasing Resources for Larger Dimensions

Note:

If you find that the deployed ISO is too small, you can use this procedure to increase the size for better performance.

To configure for larger dimensions:

1. From the ESXi Hypervisor client, start the ISO deployment by specifying the location of the virtual ISO file.
2. Complete the steps in the deployment wizard and:
 - Select **Enterprise** as the **Configuration** in the **Deployment Configuration** section.
 - Ensure that the option to **Power on after deployment** is **NOT** selected.
3. After the deployment is complete and the virtual machine is created, right-click the virtual machine and select **Edit Settings**.
4. In the **Hardware** tab:
 - Select **CPUs** and change the **Number of virtual sockets** to the appropriate value in [Nutanix AOS/ESXi required resource reservation and allocation](#).
 - Select **Memory** and change the **Memory Size** to the appropriate value in [Nutanix AOS/ESXi required resource reservation and allocation](#).
 - Select **Hard disk 2** and change the **Provisioned Size** to the appropriate value in [Nutanix AOS/ESXi required resource reservation and allocation](#).
5. In the **Resources** tab:
 - Select **CPU** and change the **Reservation** to the appropriate value in [Nutanix AOS/ESXi required resource reservation and allocation](#).
 - Select the **Memory** and change the **Reservation** to the appropriate value in [Nutanix AOS/ESXi required resource reservation and allocation](#).
6. Apply the changes by clicking **OK**.
7. Power on the virtual machine.

4.5.2 Resource Dimensioning

In keeping with Nutanix best practices and Mitel requirements, you must have a minimum of 2 GB of memory allocated for the Nutanix cluster running Hypervisor (ESXi).

It is common practice to use resource over-subscription when provisioning your virtual machines. This is allowed, but you must observe the guidelines in the Nutanix best practices and the Mitel maximum of 70% loading. This limit includes all of the Mitel appliances and third-party virtual applications; the peak load for all applications running on an individual host server must not exceed 70%.

See the *Nutanix documentation* for more information.

CAUTION:

You must maintain CPU Headroom of 30% over the peak load to guarantee service level agreement (SLA) performance.

4.5.2.1 Using Resource Pools

In an advanced environment, you might choose to set up one or more resource pools. Resource pool level resource reservations can be used to set upper and lower limits on resources available to all contained VMs (or further child resource pools), so if you choose to set up a resource pool, you must still allocate enough CPU and RAM to meet the needs of all resource reservations of the contained VMs. The resource pool reservations must be greater than or equal to the sum of all reservations of the virtual appliances that will be deployed in that pool.

- Resource pools only apply to the reservation of CPU GHz. The published resource requirements for vCPU, RAM, IOPS and storage requirements remain at 100% utilization and cannot be part of the burst model.
- It must always be possible, during any burst period, to reach the full 100% CPU GHz determined for the reservations model.
- The Pool CPU GHz resources shall be 67%, or higher, of the CPU GHz determined for the reservations model¹
- The virtual machines (VM) in the pool shall be monitored for CPU Ready. The CPU Ready measurement indicates the wait time of an application to get resources, which can impact voice quality. The VM CPU Ready shall not exceed 40ms per vCPU. CPU Ready numbers that are regularly in this region indicate potential resource contention, and may be a precursor indicator of voice quality issues.
- Mission Critical deployments shall continue to use the reservation model and shall not be included in the reservation pool. Examples of a Mission Critical business include but are not limited to: Emergency Call or Dispatch centre (112, 911, 999, etc.), Hospital Emergency, life impacting activities or services (airport control tower, fire/ambulance/police, mountain rescue, etc.), etc.
- Contact Centre deployments shall continue to use the reservation model and shall not be included in the reservation pool. Contact centres are by their very nature traffic intensive, and may be considered business critical. Dedicated and reserved resources for contact centres and associated components should be considered separate from other less business critical components, such as back office staff that use components with pooled resources.

Note:

If an application requires 2vCPU and 2GHz CPU under the reservation model, under the pooled reservation/burst model, the settings will become 2vCPU and 1.33GHz. The number of vCPU does not change, only the pooled GHz resources. If, after deploying a reservation/burst model with resource pools, there are continuous voice quality issues, then this deployment should be transitioned to a reservation only model in order to guarantee peak load resource allocation.

4.5.2.2 Storage Provisioning

During deployment of the ISO file, you must choose either of the following provisioning format:

- **Thin provisioned format** - The VM storage is initially allocated based on the storage required at the time of deployment, and can expand as storage needs increase. This is more efficient in terms of total storage requirement, but can also be considerably slower as the storage grows and becomes more fragmented.
- **Thick provisioned format** - Causes immediate allocation of all storage required. This is less efficient in terms of total storage requirement (does not allow over-commit), but it guarantees optimum performance. Mitel virtual appliances that are sensitive to voice quality, including MiVoice Business Virtual, MiVoice Border Gateway, NuPoint Virtual, and MiVoice Business Express require thick provisioning to ensure good voice quality from the first call. For more information, refer to the product-specific documentation.

Thick provisioned format includes two additional options: Lazy-zeroed and Eager-zeroed. Mitel recommends using the Lazy-zeroed option. For more information about storage provisioning and these settings, see [Disk provisioning](#).

Note:

If you are using a storage technology with de-duplication capability, then the difference between thin and thick-provisioned formats is reduced, and you may find that Thin provisioning is sufficient. The Mitel recommendation is to use Thick provisioning for all Mitel virtual appliances.

Contact your storage provider for information about the over-commit capabilities of the storage option you are using.

4.5.3 Resource Sharing Guidelines

One of the major advantages of virtualization is that you can co-locate multiple VMs on one host server.

Because most of Mitel's virtual applications provide real time voice, the customer will experience voice quality issues if insufficient CPU resources are allocated. Mitel builds in Nutanix resource reservation settings that will prevent resource limitation issues for typical customer installations.

Nutanix uses resource reservation to guarantee that the specified amount of the resource (CPU cycles and memory) is always available to the VM, regardless of the needs of other VMs sharing the resource.

The host CPU utilization must be maintained below 70% to prevent voice quality issues, as discussed in [Resource dimensioning](#).

Should performance or quality issues arise after reducing the resource reservation below the settings provided by Mitel, the first action Mitel Product Support will suggest is to restore the reservations to Mitel-mandated values.

4.5.3.1 Sharing with Third-Party Virtual Appliances

Mitel virtual appliances may be installed into the Virtual Infrastructure, where they will share resources with third-party virtual appliances and applications. A primary consideration in this scenario is ensuring that resource allocations for Mitel voice operations (MiVoice Business Virtual and other real-time voice applications) are adequate for good voice quality. Giving them too few resource reservations and priorities may result in voice quality issues. These issues may be very transitory, based on what all the virtual appliances are doing at any given point in time. This can make troubleshooting very difficult.

There are two ways to ensure good performance of Mitel real-time applications; by setting a guaranteed resource reservation for the time-sensitive applications, or by giving each application a share-based percentage of the total CPU.

Setting a guaranteed resource reservation is the best way to ensure that enough CPU is allocated to the voice applications; a percentage-based allocation does not guarantee the required performance level in cases where the CPU is overloaded.

When co-locating Mitel and third-party applications, you must ensure that restart priorities are set correctly. The virtualization platforms should be set to the highest priority. This sets the Virtual VM ahead of other non-critical appliances in the event that a host fails. For example, you would set instances of MiVoice Business Virtual at the top priority, with MiCollab Virtual at a slightly lower priority, followed by your other business applications. See [Resiliency, high availability, and disaster recovery](#) for more information.

Should performance or quality issues arise after you have reduced the resource reservation below the settings provided by Mitel, you will be asked to restore the resource reservations before support is offered for SLA issues. Maximum loading, including all Mitel and third-party virtual appliances, is 70%.

4.5.3.2 Enhanced Instance Density on Virtualization Platforms

In some cases, if the instances are running less traffic and possibly less users than the standard deployments, it might be possible to tune and reduce the GHz reservations. This can improve instance density per server.

If the installer is unfamiliar with determining and applying the resource settings, it is recommended to contact Mitel Professional Services.

4.5.3.3 Sharing with Other Mitel Virtual Appliances

Deploying multiple Mitel applications on one virtual server has advantages in that the various Mitel applications have similar requirements, and they are intended to work together. This generally simplifies resource considerations, and can result in lower maintenance costs.

When co-locating multiple instances of different Mitel virtual appliances, consider setting up a resource pool. In a resource pool, you can set priorities for resource allocation within the pool. For more information about resource pools, see [Using Resource Pools](#).

Minimum storage requirements are product-specific, and different for every Mitel virtual appliance. For product-specific storage requirements, refer to the applicable Mitel Engineering Guide. For the I/O characterization specifics, see [Nutanix AOS/ESXi required resource reservation and allocation](#).

4.5.4 Other Virtual Appliance Management Considerations

This section describes some additional items to consider when you deploy Mitel virtual appliances and applications.

4.5.4.1 Anti-affinity Rules

Some applications, especially when deployed in resiliency configurations, require that VMs for the applications do not reside on the same physical host server. Rules that ensure that certain VMs remain on different servers are called anti-affinity rules.

For example, a MiVoice Business Virtual primary controller VM should never be allowed to run on the same hardware as the corresponding MiVoice Business Virtual secondary controller VM. vCenter anti-affinity rules should be set to ensure that the primary and secondary are never moved to the same host server, including during automated actions such as vMotion or HA. See [Resiliency, high availability, and disaster recovery](#) and the Nutanix documentation for more information.

4.5.4.2 Hyper-Threading and Performance Impact

Intel's Hyper-Threading Technology makes one core appear as two logical processors, but the two logical processors share most of the core's resources. If one logical processor is running at high utilization, the other logical processor will have very little capacity that can be used before the performance of both logical processors suffers.

AMD Opteron processors use Rapid Virtualization Indexing, which helps accelerate the performance of many virtualized applications by enabling hardware-based virtual machine memory management.

ESXi assigns adjacent vCPU numbers to the logical processors on the same core. Nutanix attempts to keep the vCPUs of different VMs on separate cores. Use caution if you decide to use Nutanix advanced features to override Nutanix default behavior—assigning vCPUs to VMs. If two high-priority VMs have vCPUs on the same core, performance problems are likely.

CPU utilization shown in Nutanix management tools (Performance tab in vSphere Client, for example) can be misleading. If half the vCPUs are running at near 100% (on different cores) and the other half of the vCPUs are running near 0%, the system is actually fully loaded, although the average CPU may show 50% (the average of all the vCPUs).

Similarly, the GHz value for the system is really closer to the GHz sum of all the cores, not the sum of all the vCPUs.

With servers hosting Mitel virtual appliances that are less demanding—appliances that do not have to handle voice traffic—the assigned vCPU capacity may exceed the number of logical processors (twice the number of physical cores) by up to 50%. For example, a server with 8 physical cores provides 16 logical cores when hyperthreading is enabled. In this case, with non-voice traffic, you can provision 50% more vCPUs, for a total of 24 vCPUs.

Also, see the *MiVoice Business Engineering Guidelines* (MiVoice Business Release 10.1) for MiVoice Business on ISS and MiVoice Business Virtual.

4.5.4.3 Product-specific Guidelines

For additional product-specific restrictions and requirements, see the *Mitel Engineering Guidelines* for each product.

Table 13: Engineering Guidelines documents for Mitel virtual applications

Mitel virtual appliance product	Engineering Guidelines
MiVoice Business Virtual	MiVoice Business Engineering Guidelines
MiVoice Border Gateway Virtual	MiVoice Border Gateway Engineering Guidelines
MiCollab Client	MiCollab Client Engineering Guidelines

Deploying on Microsoft Hyper-V Infrastructure

5

This chapter contains the following sections:

- [Deploying Mitel Virtual Applications in a Hyper-V Virtual Infrastructure](#)
- [Preparing the Hyper-V infrastructure](#)
- [Installing and upgrading on Hyper-V](#)
- [Upgrading Mitel Applications on Hyper-V](#)

5.1 Deploying Mitel Virtual Applications in a Hyper-V Virtual Infrastructure

Mitel currently supports Power On and Shutdown Guest functions in the Hyper-V environment. Advanced features and clustering have not been fully tested and are not supported

This chapter outlines key considerations for deploying Mitel virtualized applications in a Microsoft Hyper-V environment, including installation and upgrade guidelines.

Hyper-V is available in two deployment options:

- **Standalone Hyper-V Server**, versions 2016, 2019, 2022, and 2025
- **Hyper-V role installed on Windows Server** versions 2016, 2019, 2022, and 2025

 **Note:** Older versions of Hyper-V and earlier Mitel product releases are not supported.

The following Mitel virtual appliances are supported in Hyper-V environments:

- MiVoice Business
- MiVoice Business Express
- MiCollab
- MiCollab Client
- NuPoint Unified Messaging (*no longer available for purchase as a standalone product*)
- MiVoice Border Gateway
- MiContact Center
- Mitel CX
- MiVoice Business Solution Virtual Instance

Currently, Mitel supports **Power On** and **Shutdown Guest** operations in Hyper-V environments. However, advanced features and clustering configurations have not been fully tested and are therefore **not supported**.

5.2 Preparing the Hyper-V infrastructure

Regardless of the deployment environment, it is important that the overall Hyper-V infrastructure is correctly configured to guarantee the highest level of availability and performance.

For best practices for physical servers hosting Hyper-V roles, refer to the Microsoft documentation, and the following recommendations:

- Ensure that all hosts that will be used to support real-time sensitive applications, such as NuPoint UM Virtual, MiVoice Border Gateway Virtual, and so on, are configured for maximum performance. This includes processor and BIOS settings requirements as listed in [Minimum hardware requirements](#).
- Supply plenty of bandwidth: 1 GB across each of at least four physical ports on each host server is considered the minimum for most deployments.
- Separate the Virtual Local Area Network (VLAN) and IP interfaces for storage traffic (assuming Storage Area Network (SAN), Network File System (NFS), or similar IP storage is used). This provides adequate throughput for storage operations, while isolating traffic from other flows.
- Preferably, support storage networking on two or more physical ports on the host servers, directed at different network switches to provide resiliency. You should also consider configuring these ports for NIC bonding (using Link Aggregation Control Protocol (LACP) or similar) to improve overall throughput.
- Voice traffic to and from the Mitel applications should be separate from data traffic. Implement a dedicated voice VLAN, and support it through multiple physical ports. The same recommendation applies for physical implementations.

5.2.1 Adding NICs

The following shows how many NICs you need for various configurations of the Mitel virtual applications in the Hyper-V environment. Refer to the Microsoft documentation for instructions.

Table 14: NICs required by Mitel virtual application and configuration

Virtual machine	Deployment option	Number of NICs required
MiVoice Business	Server-only	1
MiVoice Business Express	Server-only	1
	Server gateway	2
MiCollab Virtual	Server-only	1
	Server gateway	2
NuPoint UM Virtual	Server-only	1

Virtual machine	Deployment option	Number of NICs required
MiVoice Border Gateway	Server-only	1
	Server gateway	2
Mitel CX	Server-only	1
	Server(IVR)	1
MiVoice Business Solution Virtual Instance	Server gateway	2

5.2.2 Storage

In planning connections to, and stability of, the network storage, consider the following guidelines:

- Ensure adequate throughput to support the intended applications. Refer to the product-specific documentation for detailed requirements.
- Support storage using multi-path networking, with a minimum of two physical ports on each host server and on each of the storage units, directed at different network switches. This provides resiliency.
- RAID protection of all storage is strongly recommended.

Mitel virtual applications support for various storage architectures matches that of Hyper-V , unless specifically stated otherwise. This includes iSCSI, NFS, Fibre Channel, and host-local storage. For details, refer to the Mitel documentation for each virtual application, available at Mitel Document Center.

Note:

Regardless of the storage technology used, ensure that the storage used by Mitel applications meets or exceeds the specific requirements of the applications involved, in terms of IOPS and latency.

5.2.3 Host Server Sizing and Resource Management Considerations

When setting up your environment in preparation for the deployment of Mitel virtual appliances, you must consider host sizing and resource management.

The following section shows the resource allocation and resource reservation requirements for:

1. Mitel Virtual Appliances
2. OEM suppliers

5.2.3.1 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

5.2.3.2 Mitel OEM resourcing information for Hyper-V

The following third party applications have been verified on Hyper-V and are recorded here, along with the resource requirements:

Virtual Appliance	System Capacity	vCPU	Memory (GB)	Disk (GB)
MIR Server A ¹	Small	4	16	210
	Medium	6	16	210
	Large	12	32	260
MIR Server B Audio ²	Small	4	4	120
	Medium	6	6	160
	Large	8	8	200
MIR Server B Screen ²	Small	4	4	120
	Medium	6	6	160
	Large	8	8	200
MIR Server C	Small	4	4	100
	Medium	6	6	100
	Large	8	8	100

Virtual Appliance	System Capacity	vCPU	Memory (GB)	Disk (GB)
MIR Server D ²	Small	4	4	120
	Medium	6	6	160
	Large	8	8	200
MIR Server E ³	Small	4	4	80
	Medium	6	6	80
	Large	8	8	80
MIR Server H ¹	Small	4	8	210
	Medium	6	16	210
	Large	8	32	260
MIR ⁴	MIR NAS/Archive	4	4	80

Note:

1. HDD specification for database and storage needs to be added to this value. See application description for further details.
2. Default cache time included. Additional cache will require more HDD. See application description for further details.
3. Custom configuration. See application description for further details.
4. Recording storage in addition to server requirements. See application description for further details.

5.2.4 Resource Dimensioning

You must have a minimum of 2 GB of memory allocated for the Hyper-V Hypervisor. Refer to the Hyper-V documentation on the Microsoft web site. Refer to <http://technet.microsoft.com/en-us/library/hh831531.aspx>.

 **CAUTION:**

You must maintain CPU Headroom of 30% over the peak load to guarantee service level agreement (SLA) performance.

5.2.5 Resource Sharing Guidelines

One of the major advantages of virtualization is that you can co-locate multiple VMs on one host server.

Because most of Mitel's virtual applications provide real-time voice, the customer will experience voice quality issues if insufficient CPU resources are allocated. Hyper-V uses priorities to guarantee that the specified amount of the resource (CPU cycles and memory) is always available to the VM, regardless of the needs of other VMs sharing the resource.

Set the CPU priority to High for the voice-sensitive Mitel applications. Lower settings can be used for the other applications.

5.2.5.1 Dynamic Memory Support

Dynamic memory is not supported for Mitel Virtual Appliances. Mitel Virtual Appliances are 32-bit machines and Dynamic memory works only on 64-bit machines.

Assign the amount of Start up Memory and make sure Dynamic memory is not selected.

5.2.6 Clock Sync

The VM guest must not do clock sync with both NTP and sync-to-host. Turn off one of them.

5.2.7 Keep Host NIC Driver up-to-date

The host NIC driver should be up to date. NIC Teaming should be used and Virtual Machine Queue (VMQ) should be enabled.

5.2.8 Product-specific Guidelines

For additional product-specific restrictions and requirements, see the Mitel Engineering Guidelines for each product.

The Engineering Guidelines also list limits, requirements, and recommendations related to:

- I/O considerations, including effects on performance
- Power conservation
- Security

5.3 Installing and upgrading on Hyper-V

Deploying Hyper-V involves creating a Virtual Machine (VM) with the correct resource allocation to support the installation of the particular Mitel virtual application. This section provides instructions for creating the Hyper-V virtual machine on which you can install the Mitel virtual applications. For detailed installation instructions, refer to the product documentation.

The supported Mitel virtual applications are:

- MiVoice Business Virtual
- MiCollab Virtual
- MiVoice Border Gateway Virtual
- NuPoint Unified Communications Virtual (no longer available for purchase as a stand-alone product)
- MiContact Center
- Mitel CX

Note:

When choosing the license type to purchase, select the virtual license, but installation is manual. As with installation on the Industry Standard Server (ISS), you install MSL, and then the virtual appliance.

The custom features to configure during the VM creation are:

- Name the Hyper-V virtual machine (VM).
- Set Virtual Machine Type to Generation 1.
- Assign a Memory allocation to match the Mitel vApp requirements.
- Configure an additional NIC if the Mitel application is being deployed in server-gateway mode.
- Configure appropriate networking.
- Connect Virtual Hard Disk (fixed and correct size).
- If power is restored, select Static MAC Address option.

After creating the Hyper-V VM, follow the physical installation procedures for the particular Mitel application. Refer to the documentation for the Mitel product you are installing for detailed instructions.

Note:

Hyper-V VMs running the MSL OS do NOT support the connection of USB drives directly to the OS. This means that in Mitel virtual application installations such as MiCollab Virtual, you cannot scan for attached ISOs on a USB drive.

CAUTION:

Make sure that, while installing Windows Server 2012-R2 on the host system, no USB drives are plugged in, either externally or internally. If there are USB drives present, the installation fails with a “partition error” as it tries to install the operating system on the USB device.

5.3.1 Installing Mitel Virtual Applications on Hyper-V Virtual Machines

Installing Mitel applications in the Hyper-V environment is similar to installing on physical server, with the additional preliminary step of creating a suitable virtual machine for the application. You must add the appropriate virtual devices and set the correct resource dimensioning. Refer [Virtual Appliance Deployment host server sizing](#) website.

Installing the Mitel application on Hyper-V is a three step process:

- Create the virtual environment
- Install and configure Mitel Standard Linux (MSL) on the virtual machine.
- Install the Mitel application by mounting the ISO images from a network drive, installing the software from CD/DVD, or for some applications, installing from the Blades panel. (Applies to MiCollab and NuPoint UM.)

Note:

When choosing the license type to purchase, select the virtual license, but installation is manual. As with installation on the Industry Standard Server (ISS), you install MSL, and then the virtual appliance.

Table 15: Installation parameters for Hyper-V

Settings	Possible choices	Supported values for Mitel virtual applications
Hyper-V version		Windows Server 2016, 2019, 2022, 2025
Hyper-V hardware version	<ul style="list-style-type: none"> • Generation 1 • Generation 2 	Generation 1 Generation 2
Storage controller type	<ul style="list-style-type: none"> • IDE • SCSI 	IDE

Settings	Possible choices	Supported values for Mitel virtual applications
Virtual hard disk format	<ul style="list-style-type: none"> VHD VHDX 	VHDX
Virtual hard disk type	<ul style="list-style-type: none"> Dynamically Expanding Fixed Size Differencing 	Fixed Size
Network driver	<ul style="list-style-type: none"> Network Adapter Network Adapter (Legacy) 	Network Adapter

5.3.1.1 Create the Virtual Environment

When configuring the virtual environment, the operating system should be set as CentOS Linux 6 (64-bit) for MSL version lower than 11.0 and CentOS Linux 7 for MSL 11.0 or higher.

The Guest Hardware should be configured with the number of processors, amount of memory and disk size as specified in and the Engineering Guidelines for the application being deployed. The disk should be added to the IDE Device.

The CPU priority choices are High, Medium, Low, and Custom. Set the CPU priority to High for the voice-sensitive Mitel applications. Lower settings can be used for the other applications.

Before starting the machine (Action > Start) and before MSL is installed, you must modify the Virtual Machine (in Settings) to match for the specific product you are installing. This is a mandatory step.

5.3.1.2 Install and Configure Mitel Standard Linux on the Virtual Machine

Refer to *Mitel Standard Linux Installation and Administration Guide*, available on Mitel OnLine.

5.3.1.3 Install the Mitel Application

Installing Mitel virtual applications in Hyper-V is very similar to installing on a physical server. For Hyper-V installations, you use the same MSL and application ISO images as for installing during the physical procedures, but you must still purchase the virtual version of the license for the Mitel products.

Refer to the installation documentation for the specific Mitel product on Mitel MiAccess portal for detailed instructions for installing the appliance on a physical server.

After installation is complete, the Mitel virtual solutions for Hyper-V must maintain online connectivity to the License Server. The license synchronization of the solution is similar to that of vSphere-based deployments.

5.4 Upgrading Mitel Applications on Hyper-V

Mitel applications running on Hyper-V are installed from the ISO disk image similar to those used when installing MSL and applications on physical servers. At this time, Mitel does not offer pre-installed disk images similar to the OVA files provided for VMware installations. The upgrade choices are similar to the choices available for physical servers.

Because there is no pre-installed disk image option as there is for VMware installations, any changes to the configuration of the virtual applications such as revised resource requirements must be made manually.

As is done for physical server deployments, you will typically need to upgrade MSL and the application together. Refer to the Installation and Maintenance Guides and Release Notes for product-specific guidance.

5.4.1 Upgrade Backup/restore Variants

There are a couple of upgrade variants that do not exist for physical servers, to ensure that you can revert to your starting place if problems occur in the upgrade process.

- Before beginning the upgrade, you clone the virtual machine (VM). Then you begin the upgrade on a copy of the VM. The old VM can be kept in a shut-down state until the new, upgraded virtual application is verified to be working properly. In the event of a problem with the new virtual application, the old virtual application can be powered up in seconds and be back in service, exactly as it was before the upgrade was started.
- Alternatively, you can take a snapshot before beginning the upgrade process. Use of the Hyper-V snapshot is recommended only for use during a system upgrade, and only when the VM is powered down. You can take a snapshot before starting the upgrade, when the system is out of service. If there is a problem with the upgrade, you can revert to the snapshot to restore the appliance to service. After the upgrade is complete and verified, you must remove the snapshot. Removing the snapshot ensures adequate performance for voice-intensive Mitel virtual applications.

5.4.2 Supported Upgrade Methods

The two supported upgrade methods are:

- Blade upgrade method: Use this upgrade method for minor release and Service Pack releases.
- Fresh Install/Restore method: Use this upgrade method for major release, or if indicated in the product release notes.

Blade upgrade method

1. Back up the existing version, using the product-specific backup method. Also see [Upgrade backup/restore variants](#).
2. For MiCollab, select Upgrade MiCollab from the server console.
3. For other applications, start the upgrade from the Blades panel in the server manager.
4. After the upgrade, ensure that your virtual machine has the resources required for the new release, as appropriate for your environment.

5. After the upgrade is complete, and the first License Server synchronization has occurred, verify that the upgraded licenses have been correctly applied.

Fresh install/Restore method

1. Back up the existing version, using the product-specific backup method. Also see [Upgrade backup/restore variants](#).
2. Create a new virtual machine.
3. Configure your virtual machine, ensuring that your virtual machine has the resources required for the new release, as appropriate for your environment.
4. Install MSL from an ISO image.
5. Restore the backup data. The installation process prompts you for the name of the backup file to restore.
6. Install the Mitel Application Blades either from ISO images or download from the SWDL as documented in the Mitel application-specific Installation and Maintenance Guide.
7. After the first License Server synchronization has occurred, verify that the upgraded licenses have been correctly applied.

5.4.3 Hyper-V software and versions

Hyper-V exists in two variants:

- As a standalone product called Hyper-V Server
- As an installable role in Windows Server

The latest releases of the following Mitel virtual appliances are supported on Hyper-V when deployed in the following environments:

- **Standalone Hyper-V Server:**
 - Version 2016
 - Version 2019
 - Version 2022
 - Version 2025
- **Hyper-V role on Windows Server:**
 - Windows Server 2016
 - Windows Server 2019
 - Windows Server 2022
 - Windows Server 2025

Support is not available for older Hyper-V versions or older Mitel product versions

Linux Integration Services allows Hyper-V to interact with the Mitel guest VM. This allows operations such as shutting down the VM from Hyper-V. Microsoft has provided the software for the Linux Integration Services to the open source communities.

The drivers necessary for Linux Integration Services are included in MSL (introduced in MSL 10.0). It is not necessary to add another Mitel Blade to get these components. Do not attempt to upgrade these components yourself. Please refer your questions to Mitel Support.

Deploying on a Proxmox Infrastructure 6

This chapter contains the following sections:

- [Deployment Considerations for Proxmox](#)
- [Host Server Sizing for Mitel Virtual Appliances](#)

This chapter describes Proxmox deployment considerations specific to introducing Mitel virtual appliances and virtualized applications into your IT infrastructure.

6.1 Deployment Considerations for Proxmox

The following considerations will affect the design of your deployment:

- [Deployment types for Proxmox](#)
- [Preparing the Proxmox infrastructure](#)
- [Support for Proxmox by Mitel product release](#)
- [Host server sizing and resource management considerations](#)

6.1.1 Deployment types for Proxmox

Following sections describe three primary deployment types used for the installation of Proxmox virtual application:

Small and Medium-size Business (SMB)

SMB deployments of virtual applications are usually used in smaller businesses, for server consolidation. Management is generally not the primary issue, but running Mitel Unified Communication and Collaboration (UCC) applications on Proxmox virtual servers offers cost savings because fewer servers are required.

An SMB deployment is likely to have a small number of host servers, and possibly just a single stand-alone host. Mitel applications generally coexist with non-Mitel applications in SMB environments.

Enterprise

Enterprise-level deployments are used either to fit the Mitel system into the existing virtual IT infrastructure, or to create a new virtual IT infrastructure specifically dedicated to Enterprise UCC functions.

Mitel virtual applications usually run on a subset of the available servers, to allow space for vMotion, HA, DRS, and DPM operations, for example, as required for the installation.

Enterprise deployments may be characterized as Private Cloud, effectively a hosted model, operated by Enterprise IT, using one of the deployment models discussed below.

Hosted

The Mitel virtual UCC solutions are well-suited to the following deployment models:

- **Customer Premise Equipment** - Mitel certified dealers or customers install and configure Mitel UCC virtual applications in the Proxmox environment on the customer's premise. The virtual infrastructure is directly managed by the customer. This is primarily suited to Enterprise and SMB type deployments.

6.1.2 Preparing the Proxmox Infrastructure

Regardless of the deployment environment, it is important that the overall Proxmox infrastructure is correctly configured to guarantee the highest level of availability and performance. For more information about preparing the Proxmox, see the *Proxmox documentation*.

6.1.3 Support for Proxmox by Mitel Product Release

[Feature compatibility by Mitel product - latest Mitel releases](#) shows the Proxmox Web Console releases that are supported by Mitel virtual appliance releases. The following table shows compatibility with the Proxmox tools and applications.

Table 16: Feature compatibility by Mitel product - latest Mitel releases

Product	Proxmox
MiVoice Business (10.5)	8.4.0
MiVoice Border Gateway (12.2 and later)	8.3.5 and 8.4.0
MiCollab 10.2 SP1	8.3.5 and 8.4.0
MX-ONE	8.3.1
MiVoice 5000 (8.0 and later)	8.1.0 and later versions
MiVoice Business Solution Virtual Instance	8.4.0

6.2 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

6.2.1 Resource Sharing Guidelines

One of the major advantages of virtualization is that you can co-locate multiple VMs on one host server.

Because most of Mitel's virtual applications provide real time voice, the customer will experience voice quality issues if insufficient CPU resources are allocated. Mitel builds in Proxmox resource reservation settings that will prevent resource limitation issues for typical customer installations.

Proxmox uses resource reservation to guarantee that the specified amount of the resource (CPU cycles and memory) is always available to the VM, regardless of the needs of other VMs sharing the resource.

The host CPU utilization must be maintained below 70% to prevent voice quality issues, as discussed in [Resource dimensioning](#).

Should performance or quality issues arise after reducing the resource reservation below the settings provided by Mitel, the first action Mitel Product Support will suggest is to restore the reservations to Mitel-mandated values.

6.2.1.1 Sharing with Third-Party Virtual Appliances

Mitel virtual appliances may be installed into the Virtual Infrastructure, where they will share resources with third-party virtual appliances and applications. A primary consideration in this scenario is ensuring that resource allocations for Mitel voice operations (MiVoice Business Virtual and other real-time voice applications) are adequate for good voice quality. Giving them too little resource reservation and priority may result in voice quality issues. These issues may be very transitory, based on what all the virtual appliances are doing at any given point in time. This can make troubleshooting very difficult.

There are two ways to ensure good performance of Mitel real-time applications; by setting a guaranteed resource reservation for the time-sensitive applications, or by giving each application a share-based percentage of the total CPU.

Setting a guaranteed resource reservation is the best way to ensure that enough CPU is allocated to the voice applications; a percentage-based allocation does not guarantee the required performance level in cases where the CPU is overloaded.

When co-locating Mitel and third-party applications, you must ensure that restart priorities are set correctly. The virtualization platforms should be set to the highest priority. This sets the Virtual VM ahead of other non-critical appliances in the event that a host fails. For example, you would set instances of MiVoice Business Virtual at the top priority, with MiCollab Virtual at a slightly lower priority, followed by your other business applications.

Should performance or quality issues arise after you have reduced the resource reservation below the settings provided by Mitel, you will be asked to restore the resource reservations before support is offered for SLA issues. Maximum loading, including all Mitel and third-party virtual appliances, is 70%.

6.2.1.2 Sharing with Other Mitel Virtual Appliances

Deploying multiple Mitel applications on one virtual server has advantages in that the various Mitel applications have similar requirements, and they are intended to work together. This generally simplifies resource considerations, and can result in lower maintenance costs.

When co-locating multiple instances of different Mitel virtual appliances, consider setting up a resource pool. In a resource pool, you can set priorities for resource allocation within the pool. For more information about resource pools, see [Using Resource Pools](#).

Minimum storage requirements are product-specific, and different for every Mitel virtual appliance. For product-specific storage requirements, refer to the applicable Mitel Engineering Guide. For the I/O characterization specifics, see [Proxmox required resource reservation and allocation](#).

6.2.1.3 Enhanced Instance Density on Virtualization Platforms

In some cases, if the instances are running less traffic and possibly less users than the standard deployments, it might be possible to tune and reduce the GHz reservations. This can improve instance density per server.

If the installer is unfamiliar with determining and applying the resource settings, it is recommended to contact Mitel Professional Services.


6.2.2 Increasing Resources for Larger Dimensions

Note:

If you find that the deployed ISO is too small, you can use this procedure to increase the size for better performance.

To configure for larger dimensions:

1. Start the ISO deployment by specifying the location of the virtual ISO file.

 **Note:** MiVoice Business Virtual on Proxmox supports both OVA and ISO deployment methods. However, the ISO-based deployment is the recommended procedure.

2. Complete the steps in the deployment wizard and:
 - Select **Enterprise** as the **Configuration** in the **Deployment Configuration** section.
 - Ensure that the option to **Power on after deployment** is **NOT** selected.
3. After the deployment is complete and the virtual machine is created, right-click the virtual machine and select **Edit Settings**.
4. In the **Hardware** tab:
 - Select **CPUs** and change the **Number of virtual sockets** to the appropriate value in [Proxmox required resource reservation and allocation](#).
 - Select **Memory** and change the **Memory Size** to the appropriate value in [Proxmox required resource reservation and allocation](#).
 - Select **Hard disk 2** and change the **Provisioned Size** to the appropriate value in [Proxmox required resource reservation and allocation](#).

5. In the **Resources** tab:

- Select **CPU** and change the **Reservation** to the appropriate value in [Proxmox required resource reservation and allocation](#).
- Select the **Memory** and change the **Reservation** to the appropriate value in [Proxmox required resource reservation and allocation](#).

6. Apply the changes by clicking **OK**.

7. Power on the virtual machine.

6.2.3 Resource Dimensioning

In keeping with Proxmox best practices and Mitel requirements, you must have a minimum of 2 GB of memory allocated for the Proxmox.

It is common practice to use resource over-subscription when provisioning your virtual machines. This is allowed, but you must observe the guidelines in the Proxmox best practices, and the Mitel maximum of 70% loading. This limit includes all of the Mitel appliances and all of the third-party virtual applications; the peak load for the total of all applications running on an individual host server must not exceed 70%.

See the *Proxmox documentation* for more information.

CAUTION:

You must maintain CPU Headroom of 30% over the peak load to guarantee service level agreement (SLA) performance.

6.2.3.1 Using Resource Pools

In an advanced environment, you might choose to set up one or more resource pools. Resource pool level resource reservations can be used to set upper and lower limits on resources available to all contained VMs (or further child resource pools), so if you choose to set up a resource pool, you must still allocate enough CPU and RAM to meet the needs of all resource reservations of the contained VMs. The resource pool reservations must be greater than or equal to the sum of all reservations of the virtual appliances that will be deployed in that pool.

- Resource pools only apply to the reservation of CPU GHz. The published resource requirements for vCPU, RAM, IOPS and storage requirements remain at 100% utilization and cannot be part of the burst model.
- It must always be possible, during any burst period, to reach the full 100% CPU GHz determined for the reservations model.
- The Pool CPU GHz resources shall be 67%, or higher, of the CPU GHz determined for the reservations model¹
- The virtual machines (VM) in the pool shall be monitored for CPU Ready. The CPU Ready measurement indicates the wait time of an application to get resources, which can impact voice quality. The VM CPU Ready shall not exceed 40ms per vCPU. CPU Ready numbers that are regularly in this region indicate potential resource contention, and may be a precursor indicator of voice quality issues.
- Mission Critical deployments shall continue to use the reservation model and shall not be included in the reservation pool. Examples of a Mission Critical business include but are not limited to: Emergency

Call or Dispatch center (112, 911, 999, etc.), Hospital Emergency, life impacting activities or services (airport control tower, fire/ambulance/police, mountain rescue, etc.), etc.

- Contact Center deployments shall continue to use the reservation model and shall not be included in the reservation pool. Contact centers are by their very nature traffic intensive, and may be considered business critical. Dedicated and reserved resources for contact centers and associated components should be considered separate from other less business critical components, such as back office staff that use components with pooled resources.

Note:

If an application requires 2vCPU and 2GHz CPU under the reservation model, under the pooled reservation/burst model, the settings will become 2vCPU and 1.33GHz. The number of vCPU does not change, only the pooled GHz resources. If, after deploying a reservation/burst model with resource pools, there are continuous voice quality issues, then this deployment should be transitioned to a reservation only model in order to guarantee peak load resource allocation.

6.2.4 Other Virtual Appliance Management Considerations

This section describes some additional items to consider when you deploy Mitel virtual appliances and applications.

6.2.4.1 Anti-affinity Rules

Some applications, especially when deployed in resiliency configurations, require that VMs for the applications do not reside on the same physical host server. Rules that ensure that certain VMs remain on different servers are called anti-affinity rules.

For example, a MiVoice Business Virtual primary controller VM should never be allowed to run on the same hardware as the corresponding MiVoice Business Virtual secondary controller VM.

6.2.4.2 Hyper-Threading and Performance Impact

Intel's Hyper-Threading Technology makes one core appear as two logical processors, but the two logical processors share most of the core's resources. If one logical processor is running at high utilization, the other logical processor will have very little capacity that can be used before the performance of both logical processors suffers.

AMD Opteron processors use Rapid Virtualization Indexing, which helps accelerate the performance of many virtualized applications by enabling hardware-based virtual machine memory management.

Similarly, the GHz value for the system is really closer to the GHz sum of all the cores, not the sum of all the vCPUs.

With servers hosting Mitel virtual appliances that are less demanding—appliances that do not have to handle voice traffic—the assigned vCPU capacity may exceed the number of logical processors (twice the number of physical cores) by up to 50%. For example, a server with 8 physical cores provides 16 logical cores when hyperthreading is enabled. In this case, with non-voice traffic, you can provision 50% more vCPUs, for a total of 24 vCPUs.

Also, see the *MiVoice Business Engineering Guidelines* (MiVoice Business Release 10.4) for MiVoice Business on ISS and MiVoice Business Virtual.

6.2.4.3 Product-specific Guidelines

For additional product-specific restrictions and requirements, see the *Mitel Engineering Guidelines* for each product.

Table 17: Engineering Guidelines documents for Mitel virtual applications

Mitel virtual appliance product	Engineering Guidelines
MiVoice Business Virtual	MiVoice Business Engineering Guidelines
MiVoice Border Gateway Virtual	MiVoice Border Gateway Engineering Guidelines
MiCollab Client	MiCollab Client Engineering Guidelines

Deployment on Microsoft Azure

7

This chapter contains the following sections:

- [Deployment Considerations for Microsoft Azure](#)
- [Host Server Sizing and Resource Management for Azure](#)
- [Resource and Infrastructure considerations \(Azure\)](#)

Microsoft Azure is a public cloud infrastructure platform that provides scalable, on-demand compute and networking resources for enterprise application hosting. Azure's virtualization environment, backed by availability zones, high-performance storage, and global reach, makes it a suitable platform for deploying Mitel Unified Communications solutions in both centralized and distributed configurations.

Mitel supports the deployment of several virtual applications on Azure, allowing customers and partners to benefit from cloud-based resilience, simplified infrastructure management, and faster service activation.

The following Mitel virtual appliances are validated and supported for deployment on Microsoft Azure:

- MiVoice Business Virtual (MiVB)
- MiCollab
- MiVoice Border Gateway (MBG)
- MiContact Center Business
- MiContact Center Outbound
- MiContact Center Speech
- Mitel Interaction Recording
- MiCAM
- Open Integrated Gateway (OIG)

These products can be deployed as standalone instances, high-availability pairs, or across geographically dispersed regions depending on customer requirements. This chapter provides an overview of supported deployment models, resource planning guidelines, and infrastructure considerations specific to Azure environments.

7.1 Deployment Considerations for Microsoft Azure

Mitel virtual applications can be deployed on Microsoft Azure using different infrastructure configurations based on customer needs and availability requirements. The following deployment models are supported:

- **Standalone Deployment** A single Mitel virtual appliance is deployed in one Azure Availability Zone. This model is suitable for non-critical or development environments where redundancy is not required.
- **High Availability (HA) Pair Deployment** Two appliances are deployed in separate Availability Zones within the same Azure region. This configuration provides fault tolerance and service continuity in case of infrastructure failure.
- **Multi-Region Deployment** Mitel applications are deployed across two geographically separated Azure regions. This setup supports advanced disaster recovery scenarios and is ideal for enterprises with distributed operations or regulatory compliance needs.

Note: For detailed deployment considerations, refer to the [MiVoice Business Solution Engineering Guidelines](#). To learn about deploying on Microsoft Azure, refer to the [MiVoice Business Solution Azure Deployment Guide](#).

7.2 Host Server Sizing and Resource Management for Azure

When setting up your environment in preparation for the deployment of Mitel virtual appliances, you must consider host sizing and resource management.

The following section shows the resource allocation and resource reservation requirements for:

1. Mitel Virtual Appliances
2. Mitel OEM's

7.2.1 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

7.2.2 Mitel OEM resourcing information for Azure

Virtual Appliance	System Capacity	CPU	OS HDD	Data SSD	vCPU	Memory (GB)	OS HDD (GB)	Data SSD (GB)
MIR Server A ²	Small	D4sv4	S15	*Note 1	4	16	256	*Note 1
	Medium	D8sv4	S15	*Note 1	8	32	256	*Note 1
	Large	D16sv4	S20	*Note 1	16	64	512	*Note 1
MIR Server B Audio	Small	F4	S10		4	8	128	
	Medium	F8	S15		8	16	256	
	Large	F8	S15		8	16	256	

Virtual Appliance	System Capacity	CPU	OS HDD	Data SSD	vCPU	Memory (GB)	OS HDD (GB)	Data SSD (GB)
MIR Server B Screen	Small	F4	S10		4	8	128	
	Medium	F8	S15		8	16	256	
	Large	F8	S15		8	16	256	
MIR Server C	Small	F4	S10		4	8	128	
	Medium	F8	S10		8	16	128	
	Large	F8	S10		8	16	128	
MIR Server D	Small	F4	S10		4	8	128	
	Medium	F8	S15		8	16	256	
	Large	F8	S15		8	16	256	
MIR Server E	Small	F4	S10		4	8	128	
	Medium	F8	S10		8	16	256	
	Large	F8	S10		8	16	256	
MIR Server H	Small	D4sv4	S15	*Note 1	4	16	256	*Note 1
	Medium	D8sv4	S15	*Note 1	8	32	256	*Note 1
	Large	D8sv4	S20	*Note 1	8	32	256	*Note 1
MIR ²	MIR Archive	F4	S10		4	8	128	
Speech Analysis	Transcription Core	D2sv4	S10		2	8	128	

Virtual Appliance	System Capacity	CPU	OS HDD	Data SSD	vCPU	Memory (GB)	OS HDD (GB)	Data SSD (GB)
	Transcription Decoder	E8asv4	S10		8	64	128	
MICC-O All-in-One	Small	F8	S15		8	16	256	
	Medium	F16	S15		16	32	256	
MICC-O Dialer	Small	F4	S10		4	8	128	
	Mid-Small	F8	S15		8	16	256	
	Medium	F8	S15		8	16	256	
	Large	F16	S20		16	32	512	
MICC-O NVP	Small	F4	S10		4	8	128	
	Medium	F8	S15		8	16	256	
	Large	F8	S15		8	16	256	
MICC-O Web	Small	F4	S15		4	8	256	
	Medium	F8	S15		8	16	256	
	Large	F8	S20		8	16	512	
MICC-O SQL SSRS	Small	D4sv4	S15		4	16	256	
MICC-O SQL Plus	Mid-Small ³	D8sv4	S15	E20	8	32	256	512
	MICC-O SQL Plus Medium ³	D8sv4	S15	E20	8	32	256	512

Virtual Appliance	System Capacity	CPU	OS HDD	Data SSD	vCPU	Memory (GB)	OS HDD (GB)	Data SSD (GB)
	MICC-O SQLPlus Large ³	D16sv4	S20	E30	16	64	512	1024
MiContact Center Speech	ASR & Standard TTS 60	E2sv4	S30		2	16	1024	
	ASR & Standard TTS 120	D4sv4	S30		4	16	1024	
	Enhanced TTS 60	D8sv4	S30		8	32	1024	
	Enhanced TTS 120	D16sv4	S30		16	64	1024	
MiCAM	Call Server Type B	F8	S10	S10	8	16	128	128
	Call Server Type C	F16	S10	S10	16	32	128	128
	System Server Type B	F8	S10	S20	8	16	128	512
	System Server Type C	F16	S10	S20	16	32	128	512
Management	MPA (Probe)	D2sv4	S10		2	8	128	
	Revolution	D2sv4	S10		2	8	128	

Note:

1. The database needs to be provided on a separate partition, and the size of the partition is determined by customer requirements, such as the number of calls, the number of users, and the database retention period. Calculators are available to determine the size of the database in advance. Since this is a database with high level of access, the use of SSD rather than HDD is recommended for the Medium and Large server configurations.
2. Call and Screen Recording files are stored in Azure Cool Storage. This storage is expandable, depending on use and customer usage. Calculators are available to help determine the storage requirements in advance. This can be mapped directly for smaller systems. Access to the Cool Storage can be via the core server. However, where this storage exceeds 700G, it is recommended to use the Archive server as the front end to the Cool Storage, especially in a multi-server deployment.

There is no change in the Azure definitions

For MiContact Center Outbound, the preference is to split the single (OS) HDD drive into separate drive partitions with the operating system on drive C: and the application on drive D:, and may also include smaller data partitions. For the SQLPlus servers, the HDD specification is defined to include both the operating system and application partition, with the active data on the SSD storage for improved performance. With the Azure storage definition, there will likely be more capacity available than required for the minimum operation. Share out this capacity across the different partitions. Logs and database are especially storage hungry.

For MiCAM, two disk spaces are defined. The OS drive includes the operating system and ASR/TTS application, and the Data drive includes the MiCAM application and local cache/storage.

7.3 Resource and Infrastructure considerations (Azure)

This section outlines the infrastructure considerations when deploying Mitel virtual appliances in a Microsoft Azure environment. Proper selection of compute, storage, networking, and availability options is essential to ensure reliable operation and voice quality.

Storage

Azure deployments must use **Premium SSD managed disks** for all production workloads. These disks deliver consistent IOPS and low latency, which are required for real-time voice services. Disk sizes vary by application, ranging from **50 GB to 180 GB**. Azure Backup can be used to schedule regular snapshots and restore points for disaster recovery.

Networking

Each Mitel appliance should be deployed in an Azure Virtual Network (VNet) with separate subnets for signaling, media, and management traffic. Where supported, Accelerated Networking must be enabled to reduce jitter and improve media performance. Public IP addresses should be assigned only when necessary and secured using Network Security Groups (NSGs). For hybrid deployments, Azure ExpressRoute or a site-to-site VPN can be configured to extend on-premise connectivity.

High Availability and Redundancy

For improved availability, Mitel applications should be deployed in a **High Availability (HA) configuration** using Availability Zones. Appliances should be distributed across at least two zones within a single Azure region. Azure Load Balancer or DNS-based failover can be configured to ensure seamless redirection of signaling and user traffic. For disaster recovery scenarios, Azure Site Recovery (ASR) can be leveraged to replicate workloads to a secondary region.

Monitoring and Management

Use Azure Monitor to collect and analyze real-time performance metrics such as CPU, memory, disk I/O, and network throughput. Log Analytics can provide deeper diagnostics for troubleshooting and event analysis. For cost optimization in lab or test environments, auto-shutdown schedules should be configured. Resources should be logically grouped and tagged for simplified management and cost tracking.

Deployment on Amazon Web Services (AWS)

8

This chapter contains the following sections:

- [Deployment Considerations for AWS](#)
- [Host Server Sizing and Resource Management for AWS](#)
- [Resource and Infrastructure Considerations \(AWS\)](#)

Amazon Web Services (AWS) is a comprehensive cloud computing platform that provides infrastructure-as-a-service (IaaS) capabilities including virtual servers (EC2), scalable storage (EBS), and global network reach. AWS is well-suited for deploying Mitel communications solutions that require high availability, dynamic scalability, and rapid provisioning in enterprise and service provider environments.

Mitel has validated the deployment of the following virtual appliances on AWS:

- MiVoice Business Virtual (MiVB)
- MiCollab
- MiVoice Border Gateway (MBG)
- MiContact Center Business
- MiContact Center Outbound
- Mitel Interaction Recording
- MiCAM
- Open Integration Gateway

These solutions can be implemented in single-AZ or multi-AZ (Availability Zone) architectures, with full support for AWS-native features such as Elastic Load Balancing, Route 53 failover, CloudWatch monitoring, and EBS backup services. The sections that follow outline key AWS deployment models, instance selection guidelines, and infrastructure best practices for hosting Mitel applications in AWS environments.

8.1 Deployment Considerations for AWS

Mitel solutions deployed on AWS can leverage the platform's flexible architecture to achieve varying levels of redundancy and fault tolerance. The supported deployment models include:

- **Single-Instance Deployment** A Mitel virtual appliance is hosted in one Availability Zone. This setup is suitable for testing, small offices, or environments where uptime is not business-critical.
- **Multi-AZ (High Availability) Deployment** Appliances are deployed in two separate Availability Zones within the same AWS region. This design enhances system resilience and ensures continuity during zonal failures.

- **Multi-Region Deployment** Separate AWS regions host identical application stacks to enable full disaster recovery. This model is appropriate for large enterprises requiring geographic redundancy and low-latency service delivery across multiple regions.

Note: For detailed deployment considerations, refer to the [MiVoice Business Solution Engineering Guidelines](#). To learn about deploying on Amazon Web Services (AWS), refer to the [MiVoice Business Solution Amazon Web Services Deployment Guide](#).

8.2 Host Server Sizing and Resource Management for AWS

When setting up your environment in preparation for the deployment of Mitel virtual appliances, you must consider host sizing and resource management.

The following section shows the resource allocation and resource reservation requirements for:

1. Mitel Virtual Appliances
2. Mitel OEM's

8.2.1 Host Server Sizing for Mitel Virtual Appliances

Refer to the [Virtual Appliance Deployment hosting server sizing](#) webpage for detailed information on the recommended hardware resources, including CPU, memory, storage, and virtualization host requirements needed to support the deployment.

8.2.2 Mitel OEM resourcing information for AWS

Virtual Appliance	System Capacity	AWS Definition			AWS Resources			
		Definition	OS SSD	Data SSD	vCPU	Memory (GB)	OS SSD (GB)	Data SSD (GB)
MIR Server A ²	Small	m5.xlarge	Gp3: 210	*Note 1	4	16	210	*Note 1
	Medium	c5.2xlarge	Gp3: 210	*Note 1	8	16	210	*Note 1
	Large	c5.4xlarge	Gp3: 260	*Note 1	16	32	260	*Note 1
MIR Server B Audio	Small	c5.xlarge	Gp3: 120		4	8	120	
	Medium	c5.2xlarge	Gp3: 160		8	16	160	

Virtual Appliance	System Capacity	AWS Definition			AWS Resources			
		Definition	OS SSD	Data SSD	vCPU	Memory (GB)	OS SSD (GB)	Data SSD (GB)
	Large	c5.2xlarge	Gp3: 200		8	16	200	
MIR Server B Screen	Small	c5.xlarge	Gp3: 120		4	8	120	
	Medium	c5.2xlarge	Gp3: 160		8	16	160	
	Large	c5.2xlarge	Gp3: 200		8	16	200	
MIR Server C	Small	c5.xlarge	Gp3: 100		4	8	100	
	Medium	c5.2xlarge	Gp3: 100		8	16	100	
	Large	c5.2xlarge	Gp3: 100		8	16	100	
MIR Server D	Small	c5.xlarge	Gp3: 120		4	8	120	
	Medium	c5.2xlarge	Gp3: 160		8	16	160	
	Large	c5.2xlarge	Gp3: 200		8	16	200	
MIR Server E	Small	c5.xlarge	Gp3: 80		4	8	80	
	Medium	c5.2xlarge	Gp3: 80		8	16	80	
	Large	c5.2xlarge	Gp3: 80		8	16	80	
MIR Server H	Small	c5.xlarge	Gp3: 210	*Note 1	4	8	210	*Note 1
	Medium	c5.2xlarge	Gp3: 210	*Note 1	8	16	210	*Note 1
	Large	c5.4xlarge	Gp3: 260	*Note 1	16	32	260	*Note 1
MIR ²	MIR Archive	c5.xlarge	Gp3: 80		4	8	80	

Virtual Appliance	System Capacity	AWS Definition			AWS Resources			
		Definition	OS SSD	Data SSD	vCPU	Memory (GB)	OS SSD (GB)	Data SSD (GB)
Speech Analysis	Transcription Core	m5.large	Gp3: 120		2	8	120	
	Transcription Decoder	r5.2xlarge	Gp3: 120		8	64	120	
MiCC-O All-in-One	Small	c5.2xlarge	Gp3: 190		8	16	190	
	Medium	c5.4xlarge	Gp3: 210		16	32	210	
MiCC-O Dialer	Small	c5.xlarge	Gp3: 120		4	8	120	
	Mid-Small	c5.2xlarge	Gp3: 150		8	16	150	
	Medium	c5.2xlarge	Gp3: 180		8	16	180	
	Large	c5.4xlarge	Gp3: 340		16	32	340	
MiCC-O NVP	Small	c5.xlarge	Gp3: 120		4	8	120	
	Medium	c5.2xlarge	Gp3: 150		8	16	150	
	Large	c5.2xlarge	Gp3: 230		8	16	230	
MiCC-O Web	Small	c5.xlarge	Gp3: 160		4	8	160	
	Medium	c5.2xlarge	Gp3: 190		8	16	190	
	Large	c5.2xlarge	Gp3: 320		8	16	320	
MiCC-O SQLSSRS	Small	m5.xlarge	Gp3: 190		4	16	190	

Virtual Appliance	System Capacity	AWS Definition			AWS Resources			
		Definition	OS SSD	Data SSD	vCPU	Memory (GB)	OS SSD (GB)	Data SSD (GB)
MiCC-O SQLPlus	Mid-Small ³	m5.2xlarge	Gp3: 100	gp2: 400	8	32	100	400
	MiCC-O SQLPlus Medium ³	m5.2xlarge	Gp3: 100	gp2: 600	8	32	100	600
	MiCC-O SQLPlus Large ³	m5.4xlarge	Gp3: 100	gp2: 1000	16	64	100	1000
MiCC Speech	ASR & Standard TTS 60	r5.large	Gp3: 780		2	16	780	
	ASR & Standard TTS 120	m5.xlarge	Gp3: 1240		4	16	1240	
	Enhanced TTS 60	m5.2xlarge	Gp3: 780		8	32	780	
	Enhanced TTS 120	m5.4xlarge	Gp3: 780		16	64	780	
MiCAM	Call Server Type B	c5.2xlarge	Gp3: 80	gp2: 80	8	16	80	80
	Call Server Type C	c5.4xlarge	Gp3: 80	gp2: 80	16	32	80	80
	System Server Type B	c5.2xlarge	Gp3: 80	gp2: 420	8	16	80	420
	System Server Type C	c5.4xlarge	Gp3: 80	gp2: 420	16	32	80	420

Virtual Appliance	System Capacity	AWS Definition			AWS Resources			
		Definition	OS SSD	Data SSD	vCPU	Memory (GB)	OS SSD (GB)	Data SSD (GB)
Management	MPA (Probe)	c5.large	Gp3: 50		2	4	50	
	Revolution	c5.large	Gp3: 80		2	4	80	

Server Resource Definitions (AWS)

Note:

1. The database needs to be provided on a separate partition and the size of the partition is determined by customer requirements, such as number of calls, number of users and database retention period. Calculators are available to determine the size of the database in advance. Since this is a database with high level of access, the use of SSD rather than HDD is recommended for the Medium and Large server configurations.
2. Call and Screen Recording files are stored in Azure Cool Storage. This storage is expandable, depending on use and customer usage. Calculators are available to determine the size of the storage requirements in advance. This can be mapped directly for smaller systems. Access to the Cool Storage can be via the core server. However, where this storage exceeds 700G, it is recommended to use the Archive server as the front end to the Cool Storage, especially in a multi-server deployment

There is no change in the AWS definitions.

For MiContact Center Outbound, the preference is to split the single (OS) HDD drive into separate drive partitions with the operating system on drive C: and the application on drive D and may also include smaller data partitions. For the SQLPlus servers the HDD specification is defined to include both the operating system and application partition, with the active data on the SSD storage for improved performance. With the Azure storage definition there will likely be more capacity available than required for the minimum operation. Share out this capacity across the different partitions. Logs and database are especially storage hungry.

For MiCAM, two disk spaces are defined. The OS drive includes the operating system and ASR/TTS application, and the Data drive includes the MiCAM application and local cache/storage. See Engineering Guidelines.

8.3 Resource and infrastructure Considerations (AWS)

This section describes the key infrastructure components required for successful deployment of Mitel virtual applications in an Amazon Web Services (AWS) environment.

Storage

AWS deployments use Elastic Block Store (EBS) volumes to provide persistent storage for virtual appliances. The recommended volume type is **gp3**, which offers balanced performance for most workloads. For applications requiring higher throughput, such as MiCollab or MX-ONE, **io1** volumes may be selected. Disk sizes typically range from **50 GB to 180 GB**. AWS Backup or EBS Snapshots can be configured for regular backups and recovery.

Networking

Mitel workloads should be deployed within an Amazon Virtual Private Cloud (VPC) with appropriate subnet segmentation for media, signaling, and management traffic. Each EC2 instance must be associated with Security Groups that define allowed traffic. Elastic IPs can be used to assign static public addresses to internet-facing services. Site-to-site VPN or AWS Direct Connect can provide secure and low-latency connectivity to on-premise environments.

High Availability and Redundancy

To achieve high availability, appliances must be deployed across Multiple Availability Zones (AZs) within a single AWS region. Elastic Load Balancers (ELB) or Route 53 DNS failover should be configured for traffic redirection. Enable EC2 Auto Recovery to automatically restore instances in the event of hardware failure. For advanced recovery, a Multi-Region DR architecture may be implemented using replicated workloads.

Monitoring and Management

Use **Amazon CloudWatch** to monitor system health metrics including CPU, memory usage (via CloudWatch Agent), disk I/O, and network traffic. Alerts and thresholds can be configured to detect anomalies. AWS CloudTrail can be used for auditing API calls and changes made to the infrastructure. For administrative automation and secure shell-less management, enable AWS Systems Manager (SSM) on all instances.

This chapter contains the following sections:

- [Deploying Mitel Virtual Appliances in Hyper-V](#)
- [Deploying Mitel Virtual Appliances in VMware](#)
- [Deploying Mitel virtual appliances in Nutanix Cluster Running AHV Hypervisor](#)
- [Deploying Mitel OVAs](#)
- [Upgrading Mitel Virtual Appliances and Applications](#)

Mitel virtual products installation differs between VMware deployments and Hyper-V deployments.

9.1 Deploying Mitel Virtual Appliances in Hyper-V

See [Installing and upgrading on Hyper-V](#).

9.2 Deploying Mitel Virtual Appliances in VMware

Mitel virtual appliances that are delivered as an OVA file must be deployed as an OVA. The OVA packaging of the Mitel virtual appliance includes additional components (VMware Tools and drivers, for example), and settings such as resource dimensioning, that are required for proper operation.

Note:

- Manually installing Mitel Standard Linux (MSL) in a VM, and then manually installing the corresponding Mitel application is not supported.
- Manually changing the disk storage to an unsupported dimension is not allowed.

Note:

MiContact Center Office: On first configuration after OVA file deployment, configure the network settings. After new network settings have been applied, the Network Configuration screen will no longer be available. In the case of any network issues (IP address conflict, for example) close the installer window (click Cancel) and troubleshoot the network settings from the Windows Network and Sharing Center. Then the MiContact Center Office setup wizard can be run and the configuration can be resumed.

9.2.1 Disk Provisioning

Mitel virtual appliances that require low latency for voice streaming (such as MiVoice Business Virtual, NuPoint Virtual, and MiCollab Virtual) require the use of thick provisioning of the disks. Testing has found that the additional run-time overhead from using thin-provisioning noticeably affects voice quality.

Thick-provisioned disks may be either eager-zeroed or lazy-zeroed. Mitel recommends the use of lazy-zeroed disks.

Lazy-zeroed disks are zeroed on the first write to each block. The first write to each block is much slower than every successive write. Conversely, eager-zeroed disks have all their blocks zeroed during deployment when the disk is created for the virtual appliance.

Eager-zeroed disks take longer to deploy (on the order of a few minutes depending on the size of the disk), but do not risk slow disk performance at random times during run-time, as not-yet-used disk blocks are zeroed.

Refer to VMware Knowledge Base article 1011170, for instructions for determining whether the disk of an already-installed VM is lazy-zeroed or eager-zeroed.

9.2.2 Network driver

Mitel virtual appliances require low latency for voice streaming. Therefore, they must use the VMXNET3 network driver.

9.3 Deploying Mitel virtual appliances in Nutanix Cluster Running AHV Hypervisor

To deploy Virtual Machines on a Nutanix Cluster running on the AHV hypervisor, follow these steps:

1. Login to the Nutanix website.
2. Upload the .iso file for Windows Server 2022 to the Nutanix Prism Central.
3. Upload the Nutanix Virtio for Windows Server 2019 to the Nutanix Prism Central.
4. Click **Create VM** from Catalog Item and select Catalog type as image.
5. Select the .iso image files and click **Begin**.
6. In the **Configuration** page,
 - a. Enter a name for the new Virtual Machine.
 - b. Enter a description for the Virtual Machine.
 - c. In the VM properties section, set the CPU, Cores per CPU and Memory values.
 - d. Click Next.
7. In the **Resources** page,
 - a. Enter the Network details, select the Subnet from the list and retain the Network Connection site and Boot configuration to default.
 - b. Select the Vlans for internal, external and management ports to connect to this VM.
 - c. Select the Flash Mode checkbox for all disks.
 - d. Click Next.

8. In the **Management** page,
 - a. Select the category type and timezone.
 - b. Click Create VM.
9. In the **Review** page, review all the configurations you entered.
10. Click **Save**. The VM creation will begin.
11. Enter a username and password for the VM.
12. In the **VM > Action** section, toggle the **Power on button** to power on the VM.
13. Open the Console page of the VM.
14. Install the Windows Server 2022, enter the Windows key, and accept the terms and conditions.
15. Click the CD-Rom icon and select Nutanix Virtio driver.
16. Select Windows Server 2019/amd64 and select all the settings.

The installation will begin, and on successful installation, enter the Windows password.

9.4 Deploying Mitel OVAs

These are the basic instructions for deploying Mitel virtual appliances. Mitel virtual appliance OVAs (as opposed to VMware-ready products) contain Mitel Standard Linux, the VMware Tools, and the Mitel product files.

Note:

If you are installing multiple Mitel vApps, it is strongly recommended that you deployed them one at a time. This will avoid having the guest properties for all of the vApps displayed at the same time, which can cause confusion in completing the deployment.

There may be product-specific settings that apply to the OVA you are deploying. There may also be additional configuration steps to be performed after the OVA is deployed. Refer to the product-specific documentation for the details.

Note:

- See [Virtual Appliance Deployment host server sizing](#) website for detailed resource and capacity information.
- Refer to the product-specific installation guide for details and limitations particular to installing the product OVA, and for post-install and configuration instructions.

The following procedure is an example of how to deploy vMiCollab using the vSphere 5.5 web interface to an ESXi host via vCenter Manager. Note that the screens that are displayed during the deployment wizard are dependent on your network configuration and VMware environment. Details of the deployment for MiVoice Business, MiVoice Border Gateway, and MiVoice Business Express are added in.



Note:

It is strongly recommended that the deployment be done using the VMware Web Client, as all the new features used for Mitel OVA deployment may not be available in the older client.

1. Obtain the **Mitel OVA** for the product you are deploying.

CAUTION:

Do not install Mitel Blades from any CD ROM drive that may be attached to the virtual machine.

2. Log into the **vSphere Web Client** with your user name and password.
3. In the vSphere Client, in the upper right corner of the tool bar, click . Ensure that you have installed the latest **VMware Client Integration Plugin**.
4. In the left-hand menu, click **vCenter Inventory Lists**.
5. Click **Virtual Machines**.
6. On the **Objects** tab, click  **Deploy OVF Template**.
 - If you are using the desktop client, click **File > Deploy OVF Template**.
7. If prompted, click to allow “Client Integration Access Control”.
8. Select the Source of the OVF template file (OVA file extension):
 - **URL**: If the OVF template file is on the Internet or accessible through a web browser; enter the URL of the location of the file.
 - **Local file**: If the OVF template file was downloaded to the local computer or to a network share drive, click **Browse** to locate the file.
9. Click **Next**.
10. Review the OVF template details and click **Next**.
11. Click **Accept** to accept the End User License Agreements and click **Next**. The Deploy OVF Template Name and Location screen appears.
12. Enter a meaningful name for the product instance, or accept the default name. In **Inventory Location**, you must select a target folder if you are deploying to a vCenter Server.

Note:

When deploying in a vCenter environment, the wizard may prompt for a Data store and Network Mapping if several options are available. Contact your Data Center administrator for details about which Datastore or Network Mapping to use.

13. Click **Next**. The **Select configuration** screen appears.
14. Select the required deployment configuration, if applicable, for your site from the drop-down menu. Refer to the product engineering guidelines for the user capacities supported by each configuration.

See “[Engineering Guidelines documents for Mitel virtual applications](#),” for the name of the engineering guidelines document you need.

Table 18: Selecting deployment configurations - by product

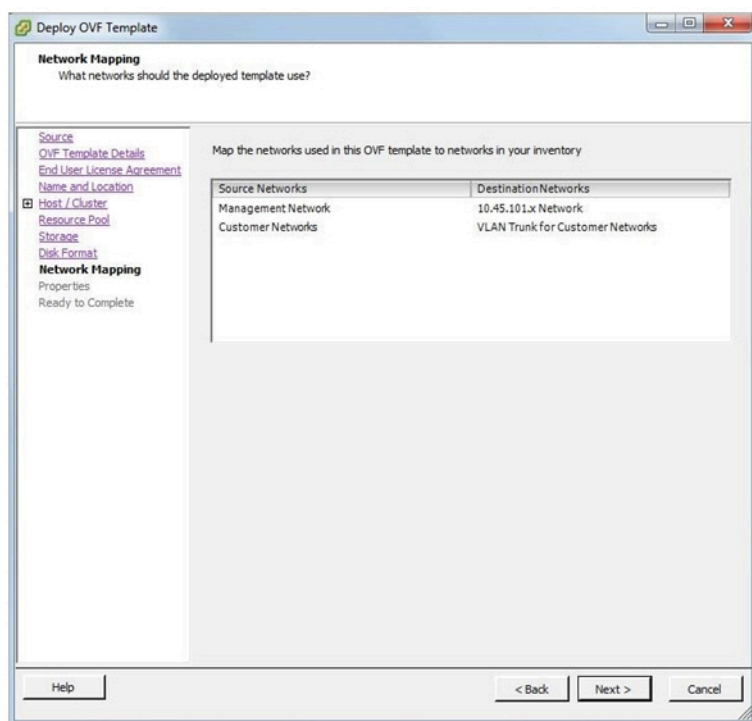
Product	Deployment Configuration
MiVoice Business	
	<p>Select Small business, Mid-Enterprise, or Enterprise as appropriate from the Configuration drop-down list.</p> <p>After you select a deployment configuration, user limits and required hardware resources are displayed.</p>
MiCollab	
	<p>In MAS Release 7.0 or above, the Enterprise configuration is supported up to 1500 users. If you are upgrading from MAS Release 5.0 to MiCollab Release 7.x, select the Mid-market configuration for a site up to 1500 users. To support a large Enterprise (5000-user multi-application capacity or 5000-user single application capacity) you must manually increase the VMware resources for the MiCollab virtual machine. Deploy the OVA using the Enterprise configuration. Before you power up the virtual machine, edit the virtual machine settings and increase the Virtual Hardware resources to the requirements, in Virtual Appliance Deployment Host Server Sizing website. Then power up the virtual machine and proceed with configuration.</p>
MiVoice Border Gateway	
	<p>Choose the required deployment configuration for your site from the drop-down menu: Small Business or Enterprise.</p> <p>After you select a deployment configuration, user limits and required hardware resources are displayed.</p>

Product	Deployment Configuration
MiVoice Business Express	
	<p>Select Small business or Mid Market as appropriate for your site from the Configuration drop-down menu.</p> <p>After you select a configuration, the user limits and required hardware resources are displayed on the screen.</p>
MiVoice Office 400	
	<p>Up to R5.0, deployments available were MBA, MBB and VA.</p>

15. Click **Next**. The Select a resource screen appears. When installing in a vCenter Server, select the Resource Pool for the virtual instance of the OVA. This page is displayed only if the cluster contains a resource pool.
16. If you are deploying the Mitel vApp in a vCenter Server, select the Host/Cluster for the vMiCollab instance. The Host/Cluster screen only appears if you are deploying on a vCenter Server.
17. Click **Next**. The Select storage screen appears.
18. Set the virtual disk format to **Thick Provisioned Lazy Zeroed**.
19. Select a destination **Datastore**.
20. Click **Next**. The Setup networks screen appears. This screen is different for each of the Mitel products that use the custom Setup networks screen. If the network defined in the OVF template does not match the name of the template on the host to which you are deploying, you are prompted to configure the Network Mapping. The required settings may be different, depending on your deployment configuration.
 - **Network Edge (Server-Gateway) Mode:** In this configuration mode, the server functions as a firewall/Internet gateway with two Ethernet interfaces. One interface is connected to the internal network (LAN) while the other is connected to the external network (Internet). Select the destination LAN and WAN networks for the OVF template. These are the “Associated Networks” that are

assigned in the LAN and WAN IP Pools. You must assign the LAN and WAN destinations to different networks.

- **LAN Only (Server-only) Mode:** In this configuration mode, the server is only connected to the internal network (LAN). For this mode, select only a destination LAN network for the OVF template. The WAN destination is ignored.
- **LAN (Optional):** This interface can be used to connect a management application or to route the SIP Proxy to an isolated SIP Proxy network.
- If you are deploying the MiCloud Management Gateway:
 - The MiCloud Management Gateway has two interfaces, Management Network and Customer Network. You must map Customer Network to a VLAN trunk.
 - If a dvSwitch is configured, ensure that it is selected.



21. Click **Next**. If you are installing MiCollab, MiVoice Business, MiVoice Business Express, MiVoice Border Gateway, or MiCloud Management Gateway, continue through the Properties screen instructions. If you are installing any other Mitel OVA, skip to Step 23.
22. For certain Mitel appliances, the **Properties** screen appears. The Properties screen only appears for deployments that use vCenter. These are the Mitel virtual appliances that use the additional Properties screen.
 - MiCollab
 - MiVoice Business
 - MiVoice Business Express
 - MiVoice Border Gateway
 - MiCloud Management Gateway

You use this screen to configure the MSL operating system parameters. Complete the fields in this screen using the information that you recorded in preparing for installation. Mandatory fields are

highlighted with a red border. [Properties page - Networking settings by OVA type](#) describes how to fill in the Properties screen for each product.

 **CAUTION:**

Make sure to enter the correct values for each of the properties.

- a. If any mandatory property values are missing, the wizard finishes, but the virtual appliance will not be able to boot. A corresponding error message appears, and you will have to re-start the OVF deployment from the beginning.
- b. If all mandatory property values are set, but one or more of the values are invalid, the Mitel Virtual Framework (MVF) blade attempts to retrieve the settings it can, and set them as defaults for the MSL console. You will have to set parameters using the MSL console configuration procedures. See the product-specific guide for details.
- c. You can use this screen only when setting up the LAN IP and WAN IP addresses for the initial deployment of the appliance. After initial boot-up, you must use the MSL server console for the product to modify the LAN IP or WAN IP addresses.

Table 19: Properties page - Common settings

Properties section	Field	Field description
Localization		
	Timezone	Select the MSL time zone setting. Default: North America.
	Keyboard	Select the keyboard used by the MSL console.
	Restore from backup	Select to open the MSL console and step through a database restore. See the MSL Installation and Administration Guide. Default: Disabled
Application		

Properties section	Field	Field description
	Password	This is the initial administrator password.
	Hostname	The host name to be used by the virtual machine. The default is MiVB. Note: This field can be left blank for template creation.
	Domain name	Specifies the domain name for this host. The default is mycompany.local. Note: This field can be left blank for template creation.
	License key	License key for the application.
	DNS server IP	The IP address of the DNS Server
	Remote network address	

Properties section	Field	Field description
	Remote network net mask	

Note:

To create a blank template for cloning, leave the following fields empty: Administrator Password, Hostname, Domain Name, LAN and WAN IP addresses. After you create the clone, you must complete these fields before you can proceed with deployment. You cannot clone an active (deployed) virtual machine.

Table 20: Properties page - Networking settings by OVA type

Mitel Product	Field name	Field Description
MiVoice Business Fields		
	IP Address	<p>The IP address to be used for the local (LAN) interface.</p> <p>This field can be left blank when creating a template, but must be set in the VM before powering on the VM. (Edit Settings > Option > vApp Options > Properties)</p> <p>No checking is done that this IP address is in the attached network or that it is available. so check your entry carefully.</p>
	LAN Netmask	<p>The netmask for the local subnet. Checking is done to make sure that this is a valid IP address, but not that this is a valid network, so check your entry carefully.</p>

Mitel Product	Field name	Field Description
	Default Gateway IP Address	The Gateway IP address. There is no check that this is a valid gateway or that the IP address exists in the subnet, so check your entry carefully.
MiVoice Border Gateway Fields		
	LAN IP Address	The IP address of the LAN side of the MiVoice Border Gateway.
	LAN Net mask	The net mask for the LAN.
	WAN IP Address	The IP address of the WAN side of the MiVoice Border Gateway.
	WAN Netmask	The net mask of the WAN.
	Optional LAN IP	
	Optional LAN Netmask	
	Default Gateway IP	
MiCollab Fields		

Mitel Product	Field name	Field Description
	<p>LAN IP Address</p> <p>(IP address of the MiCollab Virtual)</p>	<p>The IP address of the local (LAN) interface. This must be a valid IP address on the local LAN.</p> <div data-bbox="1086 449 1451 816" style="background-color: #e1f5fe; padding: 10px;"> <p>Note:</p> <p>You can leave this field blank if you are creating a blank template of the OVA file for cloning. However, you must set it before powering up the virtual appliance. You can set this IP address from vSphere Client.</p> </div> <p>Right-click MiCollab and click Edit Settings. Click the Options tab, click Properties and enter the LAN IP Address.</p>
	<p>LAN Netmask</p>	<p>The netmask of the LAN.</p>

Mitel Product	Field name	Field Description
	WAN IP Address (Optional)	<p>For Network Edge (Server-gateway) deployments, record the IP address of the external (WAN) interface. This must be a valid IP address on external WAN.</p> <p>For LAN only (Server-only) deployments, use an IP address of 0.0.0.0.</p> <div style="background-color: #e1f5fe; padding: 5px;"> <p>Note:</p> <p>You can leave this field blank if you are creating a blank template of the OVA file for cloning. However, you must set it before powering up the virtual appliance. You can set this address from vSphere Client. Right-click on the MiCollab and click Edit Settings. Click the Options tab, click Properties and enter the WAN IP Address.</p> </div>
	WAN Netmask (Optional)	Enter the Netmask of the WAN.
	2nd WAN Address	
	Default Gateway IP Address	<p>Record the Gateway IP address. For Server-gateway deployments this gateway typically points to the Internet.</p> <p>For Server-only deployments, this gateway typically points to a LAN router.</p>
MiVoice Business Express Fields		

Mitel Product	Field name	Field Description
	LAN IP Address (IP address of the vMiCollab)	<p>The IP address of the local (LAN) interface. This must be a valid IP address on the local LAN.</p> <div data-bbox="1086 449 1451 816" style="background-color: #e1f5fe; padding: 10px;"> <p>i Note:</p> <p>You can leave this field blank if you are creating a blank template of the OVA file for cloning. However, you must set it before powering up the virtual appliance. You can set this IP address from vSphere Client.</p> </div> <p>Right-click MiCollab and click Edit Settings. Click the Options tab, click Properties and enter the LAN IP Address.</p>
	LAN Netmask	The Netmask of the LAN.

Mitel Product	Field name	Field Description
	WAN IP Address (Optional)	<p>For Network Edge (Server-gateway) deployments, record the IP address of the external (WAN) interface. This must be a valid IP address on external WAN.</p> <p>For LAN only (Server-only) deployments, use an IP address of 0.0.0.0.</p> <div style="background-color: #e1f5fe; padding: 5px;"> <p>Note:</p> <p>You can leave this field blank if you are creating a blank template of the OVA file for cloning. However, you must set it before powering up the virtual appliance. You can set this address from vSphere Client. Right-click on the MiCollab and click Edit Settings. Click the Options tab, click Properties and enter the WAN IP Address.</p> </div>
	WAN Netmask (Optional)	Record the Netmask of the WAN.
	LAN (Optional)	Network interface that can be used to connect a management application or to route the SIP Proxy to an isolated SIP proxy network.
	Default Gateway IP Address	<p>Record the Gateway IP address. For Server-gateway deployments this gateway typically points to the Internet.</p> <p>For Server-only deployments, this gateway typically points to a LAN router.</p>
MiCloud Management Gateway Fields		

Mitel Product	Field name	Field Description
	Management IP Address	The Management IP Address for the Virtual Machine must be a valid unused address on the Management Network that is being used. For example, 10.45.101.185 is a valid address on 10.45.101.x network.
	Management Netmask	This is the netmask for the management interface.
	Default Gateway Address	This address typically points to the router.

Note:

- a. IP addresses must be specified for the LAN, WAN, and Optional LAN; otherwise, the virtual appliance will not power on. If you are deploying the virtual machine in LAN only (server-only) mode, set the WAN IP address to 0.0.0.0.
- b. For Network Edge (Server-Gateway) deployments, ensure that the LAN IP and WAN IP addresses are on different subnets and the Gateway IP address is on the subnet of the WAN IP address.
- c. You can use this screen only when setting up the LAN IP and WAN IP addresses for the initial deployment of the appliance. After initial boot-up, you must use the MiCollab server console interface to modify the LAN IP or WAN IP addresses.

23. Click **Next**. The Ready to Complete screen appears.
24. Review the information and click **Finish**. vSphere starts the deployment of the Mitel vApp on the server. Progress is indicated in the upper right panel of the screen.
25. After the OVA template file has been deployed, the new virtual machine appears in the inventory list in the left side navigation pane.
26. Do one of the following:
 - If you deployed the vApp on vSphere vCenter and used the Properties screen to configure the MSL Operating System parameters, continue with configuring the vApp. See the specific vApp deployment or installation guide for application configuration details.
 - If you did not use vSphere vCenter, you must configure the MSL Operating System parameters. See the specific vApp deployment or installation guide for MSL configuration details.

9.5 Upgrading Mitel Virtual Appliances and Applications

How to upgrade Mitel virtual appliances and applications depends on whether they are installed in a VMware environment or a Hyper-V environment, and whether a minor or major upgrade.

9.5.1 Upgrading on Hyper-V

Hyper-V installation and upgrades follow the physical model. Refer to the installation guides for the specific product or application to perform upgrades.

9.5.2 Upgrading on VMware

Mitel virtual appliances include three components, all of which must be considered when making upgrade decisions:

- the virtual application
- the Mitel Virtual Framework (MVF) blade
- the Mitel Standard Linux (MSL) OS

Note:

All products use MSL except for MiContact Center Office, which runs on Microsoft® Windows®.

When upgrading a virtual appliance, there are two options.

- The new OVA can be downloaded and deployed into the Virtual Infrastructure, with an MSL data backup/restore (for MSL-based virtual appliances) used to move the data from the old virtual appliance to the new virtual appliance.
- The application's Blades can be upgraded individually through the MSL Blades panel, in the Install Applications tab.

Note:

The Blade upgrade is not available if the OVA dimensions have changed. If the OVA dimensions have changed in the release, upgrade by installing the new OVA.

Refer to the Installation and Maintenance Guides and Release Notes for product-specific guidance.

Follow these upgrade guidelines:

- Use Blade upgrades of the MVF, MSL, and virtual application for minor releases and Service Pack (SP) releases.
- Use an OVA upgrade for major releases (the first two digits of the version number changes between old and new loads).

There may occasionally be releases where something in the OVA packaging itself is modified and an OVA upgrade is required. In such cases, this requirement will be highlighted in the Release Notes.

9.5.3 Software Upgrades

You upgrade each Mitel virtual appliance individually, using Mitel upgrade procedures. VMware vCenter Update Manager is not supported for Mitel virtual applications. Log and core files must be backed up separately if you want to keep them.

Do not update the virtual appliance manually. Manual updates miss changes that affect the virtualization-specific aspects of the product, such as resource reservation changes, VMware Tools updates, and hooks to the VMware environment. Following the procedures in this section will ensure that any virtualization-specific changes are applied.

Use these rules when you upgrade Mitel applications:

Blade upgrade method

For minor release and Service Pack release (within the same release)

1. Back up the existing version for disaster recovery purposes, using the product-specific backup method.
2. Power down the virtual machine.
3. After OVA deployment, ensure that your virtual machine has the resources required for the new release, as appropriate for your environment. The OVA deployment will set default values that may need adjustment for your specific needs.
4. After the first License Server sync has occurred, verify that the upgraded licenses have been correctly applied.

OVA Upgrade method

For major release upgrades, or as indicated in the product release notes:

1. Back up the existing version, using the MSL backup method.
2. Power down the virtual machine.
3. Deploy a new virtual appliance using the new OVA file for the upgrade version of the virtual appliance.
4. After OVA deployment, ensure that your virtual machine has the resources required for the new release, as appropriate for your environment. The OVA deployment will set default values that may need adjustment for your specific needs.
5. Restore the backup data. The deployment process prompts you for the backup file to restore.
6. After the restore is complete, and the first License Server sync has occurred, verify that the upgraded licenses have been correctly applied.

Note:

MCD Release 6.0+ and MiVoice Business Virtual: Each MiVoice Business Virtual requires two IP addresses, one for MSL and one for MCD. MCD Release 5.x and previous: Each vMCD requires five IP addresses, one for MSL and 4 contiguous IP addresses for the MCD space.

CAUTION:

Ensure that there are no IP address duplication in the network.

VMware snapshot is recommended only for use during a system upgrade, and only when the VM is powered down. You can take a snapshot before starting the upgrade, when the system is out of service. If there is a problem with the upgrade, you can revert to the snapshot to try again. After the upgrade is complete and verified, you must remove the snapshot. This ensures adequate performance for voice-intensive Mitel virtual appliances. It is not necessary to clear the hardware ID on the License Server.

9.5.4 OVA Upgrade Advantages

The advantages to upgrading by deploying a new OVA are:

- This is often the faster upgrade method for applications with many blades (such as NuPoint UM Virtual, for example).
- This is a well-understood upgrade method for Virtual Infrastructure administrators familiar with VMware.
- The new OVA will have a new file system, free from fragmentation, corruption, and obsolete files, that may have accumulated over a long run time in the old VM.
- The new OVA is guaranteed to have a tested lineup of software components that are known to work together.
- The old VM can be kept in a shut down state until the new upgraded virtual appliance is verified to be working properly. In the event of a problem with the new virtual appliance, the old virtual appliance can be powered up in seconds and be back in service, exactly as it was before the upgrade was started.
- Any changes to the configuration of the virtual appliance (in the OVA), such as revised resource reservations or updated VMware Tools, will be deployed in this approach, but not in the Blade-upgrade approach.

Note:

Use of the VMware vCenter Update Manager is not supported.

9.5.5 MSL Application Blade Upgrade Advantages

The advantages of using Blade upgrades are:

- This is often the faster upgrade method for small applications.
- This is a familiar upgrade method for experienced MSL Administrators.

- All logs, core files, and so on, on the existing system are retained in their current locations.

9.5.6 VMware Virtual Hardware Versions

VMware ESXi 6.0 introduces Virtual Hardware version 11, whereas earlier versions of ESXi supported earlier virtual hardware versions. Since VMware does not provide backward compatibility for virtual hardware versions, VMs running on hardware version 9 cannot run on previous versions of ESXi. In a mixed environment, this will affect which servers vMotion or HA can relocate each VM to, complicating operation of the virtual infrastructure.

Virtual hardware version 10 is not supported due to issues with some versions of vSphere client.

All Mitel virtual appliances can be deployed, and will run, on virtual hardware versions 7, 8, or 9. Current release Mitel virtual appliances are packaged to automatically deploy using the highest available virtual machine version available on the ESXi host they are initially deployed on.

If you are deploying on a mixed environment—a cluster containing multiple different ESXi versions—start by deploying on a host server running the lowest ESXi version available within the cluster. This will allow vMotion and HA operations on all hosts in the cluster.

Note:

Under some circumstances, newer versions of ESXi may prevent normal virtual machine startup due to interactive prompting and other issues that appear if they are not running the latest virtual hardware version available on the host. This can affect Power on, vMotion, HA, and other availability features. Virtual machine version upgrade to the latest available version fixes these issues. These are not Mitel-specific problems, but can affect any applications running on ESXi. It is strongly recommended that you upgrade the virtual hardware of deployed Mitel virtual machines to the latest available on the host ESXi. After the upgrade is done, you cannot downgrade to an earlier Virtual Machine Version that may be required for running on older ESXi version Refer to VMware documentation for performing the upgrades, and for more information about ESXi/Virtual Machine version compatibility.

For a definitive listing of VMware virtual hardware compatibility, refer to <http://kb.vmware.com/>.

Table 21: Virtual Hardware support by Mitel virtual appliance

Current Releases		
Product	Release	VMware Hardware Version in OVA
MiVoice Business Virtual	10.5	13, 14, 17 and 20
MiVoice Connect	19.3 SP3	not supported
MiVoice Office 400	7.1 SP1	13 and above (14, 15, 16, 17, 18, 19)

Current Releases		
Product	Release	VMware Hardware Version in OVA
MiCollab Virtual	10.2	7, 8, 9, 11
Mitel Mass Notification	6.5	not supported
MiVoice Border Gateway Virtual	12.3	14,15,17 and 20
MiContact Center Business Virtual	10.0	7, 8, 9
Mitel CX	1.0	7, 8, 9
MiVoice Business Express	8.0	7, 8, 9, 11
MiContact Center Enterprise (formerly Solidus)	9.7 SP1	10 and above
MiContact Center Office	6.2 SP1	7, 8, 9
MiCloud Management Portal (Oria)	6.0	7, 8, 9
MiCloud Management Gateway	5.0	7, 8, 9, 11
Enterprise Manager	9.0	not supported
MiVoice Business Solution Virtual Instance	2.0	14, 15, 17, 18, 19, 20 and 21

9.5.7 VMware software and versions

Not all versions of VMware components are compatible with each other. Refer to the “VMware Product Interoperability Matrixes” at the VMware Web site: <https://knowledge.broadcom.com/external/article/312100/esxi-hosts-and-compatible-virtual-machin.html>

Table 22: VMware software versions by Mitel virtual appliance versions

		ESXi						vSphere Web Client ¹					
		8.0	7.0	6.7	6.5	6.0	5.5	6.7	6.5	6.0	5.5	5.1	5.0
Current Releases and updates													
MiVB Virtual	10.5	Y	Y	N	N	N	N	Y	Y	N	N	N	N
MiVoice Connect	19.3 SP3			Y	Y	Y		Y	Y	Y			
MiVoice Office 400	7.1 SP1		Y	Y	Y			Y	Y				
MiCollab	10.2	Y	Y	Y	Y	Y	Y	Y	Y				
Mitel Mass Notification	6.5				Y						Y	Y	Y
MiVoice Border Gateway Virtual	12.3	Y	Y	N	N	N	N	N	N	N	N	N	
MiVoice Business Express	8.0				Y	Y	Y		Y	Y	Y	Y	
MiContact Center Office	6.2 SP1					Y	Y					Y	Y
Mitel CX	1.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	
MiContact Center Enterprise	9.7 SP1							Y	Y	Y	Y		
MiCloud Management Portal (Oria)	6.0				Y	Y	Y		Y	Y	Y	Y	
MiCloud Management Gateway	5.0				Y	Y			Y	Y	Y		
MiVoice Enterprise Manager	9.0							Y	Y	Y	Y		
MiVoice Business Solution Virtual Instance	2.0	Y	Y					Y	Y				
Previous Releases													
MiVoice Business Virtual	10.4, 10.3, 10.2, 10.1x, 10.0x, 9.4x, 9.3, 9.2, 9.1			Y	Y	Y	N	Y	Y	Y	N	N	N
	9.0 SP1, 9.0				Y	Y	Y		Y	Y	N	N	N
	8.0				Y	Y	Y			Y	Y	Y	
	7.3, 7.2, 7.1, 7.0					Y	Y				Y	Y	Y
MiVoice Connect	19.1 x, 19.2 x, 19.3 x			Y	Y	Y		Y	Y	Y			

		ESXi						vSphere Web Client ¹					
		8.0	7.0	6.7	6.5	6.0	5.5	6.7	6.5	6.0	5.5	5.1	5.0
MiVoice Express	7.3, 7.2, 7.1					Y	Y			Y	Y	Y	Y
MiCollab with Voice	7.0, 6.x					Y	Y			Y	Y	Y	Y
MiCollab	10.1, 10.0 SP1, 10.0	Y	Y	Y	Y	Y	Y	Y	Y				
	9.8 SP1, 9.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	9.7 SP1, 9.7, 9.6, 9.5, 9.4, 9.3+		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
MiVoice Border Gateway Virtual	12.2, 12.1	Y	Y	Y	Y	N	N	Y	Y	N	N	N	
	12.0, 11.x			Y	Y	N	N	Y	Y	N	N	N	
vUC Advanced	6.x, 5.x					Y	Y				Y	Y	Y
MiVoice Office 400	7.1		Y	Y	Y			Y	Y				
	7.0, 6.3			Y	Y			Y	Y				
	6.2, 6.1 HF3, 6.1, 6.0 SP1				Y				Y				
vNuPoint UM	8.3, 8.2, 8.1					Y	Y				Y	Y	Y
vNuPoint UM	7.0						Y				Y	Y	Y
vNuPoint UM	6.0, 5.0 SP2											Y	Y
vNuPoint UM	5.0												Y
MiContact Center Business	10.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	
MiContact Center Business	9.4.x, 9.3.x, 9.2.x, 9.1.x, 9.0.x				Y	Y	Y		Y	Y	Y	Y	
Oria Virtual	5.3, 5.2, 5.1, 5.0+					Y	Y				Y	Y	Y
Oria Virtual	4.0						Y				Y	Y	Y
vCSM	6.0												Y
vUIC	3.0						Y					Y	Y

Note:

The vSphere Client can be used to manage VMs with VMware virtual hardware versions up to version 9, but you will not be able to create new VMs. If you need to both create or manage VMs, you must stay at virtual hardware version 7 or 8, and you must manage your VMs with the vSphere Client.

Table 23: vCenter and vSphere compatibility with Mitel virtual appliances

		vSphere Web Client ¹						vCenter Server and vSphere Client ¹							
		8.0	7.0	6.7	6.5	6.0	5.5	8.0(UR2)	7.0	6.7	6.5	6.0	5.5	5.1	5.0
Current Releases															
MiVoice Business Virtual	10.5	Y	Y	N	N	N	N	Y	Y	N	N	N	N	N	N
MiVoice Connect	19.3 SP3			Y	Y	Y						Y			
MiVoice Office 400	7.1 SP1			Y	Y					Y	Y				
MiCollab Virtual	10.1	Y	Y	Y				Y	Y	Y					
MiCollab Client	10.1	Y	Y	Y				Y	Y	Y					
MiVoice Border Gateway Virtual	12.3	Y	Y	N	N	N	N	Y	Y	N	N	N	N		
Mitel Mass Notification	6.5														Y
MiContact Center Office	6.2 SP1				*	Y	Y					Y	Y	Y	Y
MiCloud Management Portal (Oria)	6.0				Y	Y	Y				Y	Y	Y		
MiCloud Management Gateway	5.0				Y	Y	Y				Y	Y	Y		
Mitel CX	1.0	Y							Y	Y	Y	Y	Y		
MiVoice Business Express	8.0					Y	Y				Y	Y	Y		
MiVoice Enterprise Manager	9.0														
MiVoice Business Solution Virtual Instance	2.0, 1.0														
Previous Releases															
MiVoice Business Virtual	10.4, 10.3, 10.2, 10.x, 10.0x, 9.2, 9.1x			Y	Y	Y	N			Y	Y	Y	N	N	N
	9.0+, 8.0				Y	Y	Y				Y	Y	Y		
	7.3, 7.2, 7.1, 7.0					Y	Y					Y	Y	Y	Y

		vSphere Web Client ¹						vCenter Server and vSphere Client ¹							
		8.0	7.0	6.7	6.5	6.0	5.5	8.0(UR2)	7.0	6.7	6.5	6.0	5.5	5.1	5.0
MiVoice Connect	19.1x, 19.2x, 19.3x			Y	Y	Y						Y			
MiVoice Business Express	7.3, 7.2, 7.1, 7.0					Y	Y					Y	Y	Y	Y
MiCollab Virtual	10.0 SP1, 10.0	Y	Y	Y				Y	Y	Y					
	9.8 SP1, 9.8			Y	Y			Y	Y	Y	Y				
	9.7 SP1, 9.7, 9.6, 9.5, 9.4, 9.3			Y	Y	Y	Y			Y	Y	Y	Y		
vMAS	7.3, 7.2, 7.1, 7.0, 6.0, 5.0					Y	Y					Y	Y	Y	Y
NuPoint UM Virtual	8.3, 8.2, 8.1					Y	Y					Y	Y		Y
NuPoint UM Virtual	8.0, 7.0					Y	Y					Y	Y	Y	Y
vNuPoint UM	6.0											Y	Y	Y	Y
MiCollab Client	9.5, 9.4, 9.3			Y	Y	Y	Y			Y	Y	Y	Y		Y
MiVoice Office 400	7.1			Y	Y	Y					Y	Y			
	7.0, 6.3, 6.2, 6.1 HF3				Y						Y				
MiVoice Border Gateway Virtual	12.2, 12.1	Y	Y	Y	Y	N	N				Y	Y	N		
	12.0, 11.x			Y	Y	N	N				Y	Y	N		
Oria Virtual	5.3, 5.2, 5.1					Y	Y					Y	Y		
	5.0 SP1, 5.0					Y	Y					Y	Y	Y	
MiContact Center Business	10.0									Y	Y	Y	Y		
MiContact Center Business	10.0.x, 9.5.x, 9.2.x, 9.1.x, 9.0.x											Y	Y		
vCSM	6.0														Y
vUIC	3.0					Y	Y					Y	Y	Y	Y

Note:

The vSphere Client can be used to manage VMs with VMware virtual hardware versions up to version 9, but you will not be able to create new VMs. If you need to both create or manage VMs, you must stay at virtual hardware version 7 or 8, and you must manage your VMs with the vSphere Client.

9.5.8 VMware Tools

VMware Tools is a software deliverable provided by VMware. It is installed and run inside the Guest O/S, to integrate the virtual and guest environments, and to improve VM performance.

VMware Tools packages are designed to integrate with particular, well known commercial operating systems like Red Hat Linux, Windows, and so on. In the case of Mitel Standard Linux (MSL), which is not considered well known by VMware, VMware Tools distributions are managed within the Mitel Virtualization Framework (MVF). Mitel virtual applications based on MSL include the MVF Blade as part of their OVA packaging, with the MVF containing the latest VMware Tools at time of product release. Do not attempt to install or upgrade the VMware Tools using VMware vSphere functions.

CAUTION:

Do not perform a manual install of VMware tools in MSL. VMware Tools upgrades must be done by upgrading the MVF BLADE. MiContact Center Office is the one exception. Refer to the MiContact Center Office documentation for details.

In the case of Mitel applications containing MVF in their OVA packaging, in the vSphere Client Summary Tab, the VMware Tools field displays “(3rd-party/Independent)”. This means that VMware cannot manage or upgrade the VMware Tools in that Guest O/S. The status will be displayed as Unmanaged in the vSphere 4.1 Client.

While it is desirable to always upgrade VMware Tools to the latest available for the version of vSphere being used, there are constraints on how this can be performed.

- For Mitel virtual appliances containing MSL 10.0 and above combined with MVF 2.0 and above, an MVF Blade upgrade can be applied in order to upgrade VMware Tools. From the MSL blades panel, it may be possible to upgrade MVF (and VMware tools) if a newer MVF version is shown. OVA upgrade to a newer version of the appliance containing the desired VMware Tools level is also applicable. See [OVA upgrade advantages](#).
- For Mitel virtual appliances that contain either MSL 9.x series or MVF 1.x series, MVF Blade upgrade is not applicable. OVA upgrade to a newer version of the appliance, containing the desired VMware Tools, is the only method available. See [OVA upgrade advantages](#).
- Those Mitel Applications that are not packaged as an OVA, but are installed into a pre-existing VMware Virtual Machine and Guest O/S, can have their VMware Tools software installed and upgraded in the Guest O/S using the normal VMware-supplied methods. This also applies to OVA-packaged virtual appliances based on Windows OS (MiContact Center Office Virtual, for example). Refer to the VMware documentation for instructions.

The following table summarizes VMware Tools upgrade methods applicable to each Mitel virtual appliance or application.

Table 24: VMware Tools packaging and upgrade by Mitel virtual appliance

Current Releases		VMware Tools included	VMware Tools Upgrade
MiVoice Business Virtual	10.5	Y	OVA upgrade
MiVoice Office 400	7.1 SP1	Y	OVA upgrade
MiCollab Virtual	10.2	Y	OVA to 6.0 and above
Mitel Mass Notification	6.5	Y	Firmware upgrade
MiVoice Border Gateway Virtual	12.3	Y	MVF blade upgrade or OVA to 8.1 and above
NuPoint UM Virtual (This is the final release of NuPoint VM.)	8.3	Y	MVF blade upgrade or OVA to 7.0 and above
MiVoice Business Express	8.0	Y	MVF blade upgrade or OVA to 6.0 and above
Open Integration Gateway	4.2 SP4	Y	MVF blade upgrade or OVA
MiContact Center Business	10.0	Y	Use VMware methods
Mitel CX	1.0	Y	Use VMware methods
MiContact Center Office	6.2 SP1	Y	

Current Releases		VMware Tools included	VMware Tools Upgrade
MiCloud Management Portal (Oria)	6.0	Y	
MiCloud Management Gateway	5.0	Y	Use VMware methods
MiVoice Enterprise Manager	7.0	N	Use VMware methods
MiVoice Call Recording	9.2 SP7	Y	Use VMware methods
MiVoice Business Solution Virtual Instance	1.0	Y	Use VMware methods
MiVoice Business Solution Virtual Instance	2.0	Y	MVF blade upgrade or OVA upgrade

9.5.9 Additional VMware Features

The following table shows the VMware features not already discussed elsewhere in this guide, and discusses their use with Mitel virtual solutions.

Table 25: Additional VMware features

Function	Available?	Notes
Update Manager and O/S Patching	Not for MSL virtual appliances	VMware does not have the necessary knowledge of MSL or its updates and patches.
Consolidated Backup	No	

Function	Available?	Notes
vStorage APIs (VADP)	Yes	<p>Except for MiContact Center Office Virtual</p> <p>The “vStorage APIs for Data Protection” feature enables backup software to protect system, application, and user data in their virtual appliances in a simple and scalable way. These APIs enable backup software to:</p> <ul style="list-style-type: none"> • Perform full, differential, and incremental image backup and restore of virtual appliances. • Perform file-level backup of virtual appliances using supported Windows and Linux operating systems.
vSphere Storage Appliance	Yes	Virtual storage running on VMware-enabled servers. Provides storage virtualization, with the ability to remove the danger of a single point of failure. Used as a cost-effective solution for smaller deployments that do not have SAN
Data Recovery	Yes	Except for MiContact Center Office Virtual
Hot Add	No	
VMsafe	No	
vCenter Converter	No	

This chapter contains the following sections:

- [Maintenance](#)

After your system is up and running, there are regular tasks to perform to keep the system running well:

- [Backups](#)
- [Restores](#)
- [Software upgrades](#)
- [Monitoring resource usage and performance](#)

10.1 Maintenance

10.1.1 Backups

Virtual machine backups can be done for the entire VM, as a whole, or at the product level, and there are different ways to do the various types of backup. In general, you should start with the Release Notes for each product for detailed information.

The following sections give background information on each backup and restore approach. See the product-specific backup and restore methods, as found in the corresponding Mitel product guides.

10.1.1.1 Backup Methods in Hyper-V

To perform backups of Mitel virtual applications running on Hyper-V virtual machines, follow the backup procedures in the Mitel product-specific documentation.

10.1.1.2 Backup Methods in VMware

The sections that follow describe several ways of backing up the virtual appliance or the whole VM, and the best ways to perform backups, depending on circumstances.

10.1.1.2.1 Backups Using Virtual Machine Snapshots and Replication

Snapshots are typically used to create a restoration point before installing a different version of an application, or before performing potentially destabilizing or error-prone re-configuration of the applications. After the new application is up and running, or re-configuration verified, the restoration point must be removed to avoid any performance degradation. Snapshots should not be used as a method of backing up the Mitel virtual appliance. Snapshots must be deleted after restore is complete.

To perform a complete backup of your system, you can use VMware management functions to back up the full virtual appliance VM, which includes all of the virtual applications running on it. When creating backups of the whole VM, you must ensure that you have plenty of storage space because the backups can be very large for a large system.

In the general case, you can take a snapshots of the whole VM while it is powered on (running), suspended, or shut down.

Note:

Live snapshots are not supported for Mitel virtual appliances. Snapshots of a suspended Virtual Machine are also not supported. Mitel virtual appliances must be powered down before a snapshot is saved.

Snapshots preserve the state of a VM such that it can be returned to the same state repeatedly. A snapshot captures the entire state of a virtual machine at a given moment in time. This includes:

- Memory state: the contents of the virtual machine memory
- Settings state: the virtual machine settings
- Disk state: the state of all virtual disks on the virtual machine
- Power state: whether the virtual machine is running, suspended, or shut down

When reverting to a snapshot, these items are returned to their state at the time the snapshot was taken. With snapshots, a user can create backup and restore positions in a linear process. Snapshots are typically used as restoration points during a linear or iterative process, such as installing update packages, or during a branching process, such as installing different versions of an application.

Because Mitel virtual appliances are communicating with other devices on the network and because the risk of a database corruption increases when live snapshots are taken, it is strongly recommended that the Virtual Machine hosting a Mitel virtual appliances be powered down prior to taking a snapshot.

Note:

Keeping a snapshot consumes run-time resources, because the virtual disk subsystem must keep track of changed blocks from the snapshot. This presents a load on the system that can cause performance degradation; snapshots should be kept only temporarily, and then deleted when they are no longer needed.

10.1.1.3 VMware Data Protection

VMware Data Protection (VDP) protects against data loss in your virtual environment by enabling fast backups to disk, and fast and complete recovery of data. It is built on the vStorage APIs for Data Protection (VADP), and is fully integrated with vCenter Server for centralized scheduling of backups. vSphere Data Protection is itself a virtual machine that runs on VMware ESX and ESXi hosts. vSphere Data Protection requires vCenter Server.

vSphere Data Protection also provides a centralized management interface to enable backup and recovery of your VMs directly through VMware vCenter Server, and automatically monitors VMs that are moved by HA, vMotion, and DRS, so that scheduled backups can continue uninterrupted.

During a backup, Data Recovery creates a quiesced snapshot of the virtual machine. Deduplication is automatically performed with every backup operation.

Data Recovery can concurrently back up a maximum of eight virtual machines. To start multiple backups, CPU utilization must be less than 90 percent. Due to memory constraints, Data Recover does not support using more than two backup destinations simultaneously. If more than two backup destinations must be used, configure them to be used at different times.

You can create backup jobs that include which virtual machines to backup, where to store the backups, and for how long. For details on installing and running Data Recovery, refer to the VMware documentation.

Note:

To avoid adverse effects on NuPoint Virtual voice quality, use vSphere Data Protection backups only at times when the call rate is less than 500 Calls Per Hour (CPH).

10.1.1.3.1 vStorage APIs for Data Protection (VADP)

The VADP allows third-party backup software providers to access VMFS volumes directly, off-loading the work from the main processors. A VADP-based backup can be expected to have no adverse effects on voice quality.

Note:

Third party backup applications have not been specifically qualified by Mitel, and are not fully supported.

10.1.1.4 Backup Routine Recommendations

You should rely primarily on the Mitel product-specific backup recommendations; see following table for the guides containing this information for the various products.

Table 26: Recommended backup methods

Backup/Restore requirement	Recommended method
After initial configuration or provisioning	Application or MSL backup*

Backup/Restore requirement	Recommended method
Day-to-day backup for application data recovery	Use one or more of these methods: <ul style="list-style-type: none"> Scheduled Application backup or MSL backup.* Scheduled vSphere Data Protection Third-party backup software using VADP (if available) Export/Import OVA
Application Blade software upgrade	Application or MSL backup.* Snapshot (remove after upgrade validated).
Deploying a new version of a Mitel virtual appliance with a new release	Application or MSL backup.* Import after OVA deployment of upgrade (major release).
Migration from physical application to virtual appliance	Application or MSL backup.* Import after OVA deployment.

Note:

1. See [Table 22](#) for the product-specific guides.
2. MiVoice Business Virtual requires both an MiVoice Business backup and an MSL backup.
3. vUIC requires only an application backup.

10.1.2 Restore

If you perform backups using the product-specific backup procedures, you should also use the product-specific restore routines described in the product documentation. See [Table 22](#) for the product-specific guides.

10.1.2.1 MiCollab Virtual System Recovery

You can recover a MiCollab Virtual system database on the same virtual machine by deploying the latest MiCollab Virtual OVF file, restoring your MiCollab Virtual database backup, and then installing the latest MiCollab application software.

For detailed instructions, refer to the *MiCollab Installation and Maintenance Guide*.

10.1.2.2 Monitoring resource usage and performance

To keep your Mitel virtual applications running at peak performance, it is recommended that you track ongoing CPU and memory utilization, network I/O performance, and storage I/O performance. This can be done manually, by regularly observing performance statistics for each host server in the solution, and for critical VMs. These statistics are available through vSphere Client/vCenter Performance tabs. Ensure that vCPU, memory, network and disk performance is within the limits stated for the specific applications, and the general guidelines given elsewhere in this document.

Tracking performance and resource use in an automatic way is especially useful when monitoring critical resources. VMware vCenter Operations integrates directly with third-party monitoring tools for a complete view of your infrastructure. You can set the upper and lower bounds of the “normal” range, and configure alerts that appear when performance criteria stray outside of the normal range. For more information, see the VMware White Paper, [Sophisticated Dynamic Thresholds with VMware vCenter Operations](#).

Note:

NuPoint has specific limits for minimum I/O storage that may require monitoring. See the Mitel NuPoint Unified Messaging Engineering Guidelines.

10.1.3 Migrate a Virtual Machine Between Hypervisors

To migrate a virtual machine to or from a different hypervisor, follow these steps:

1. Deploy the new virtual machine on the destination hypervisor.
 - To deploy a VMware virtual machine, see [Deploying On A VMware Infrastructure](#).
 - To deploy a Nutanix virtual machine, do either of the following:
 - Deploy a virtual machine on Nutanix Cluster running ESXi hypervisor, see [Deploying on a Nutanix Cluster Running ESXi Hypervisor](#).
 - Deploy a MiVoice Business virtual machine running Nutanix AHV, see [Deploying on a Nutanix Cluster Running AHV Hypervisor](#).
 - To deploy a virtual machine in a Hyper-V Virtual Infrastructure, see [Deploying on Microsoft Hyper-V Infrastructure](#).
2. Manually configure and apply the Serial ID to MSL.
3. Download all the licensed blades or application .iso files and install.

4. Restore an MSL backup (not recommended for MiVoice Business, see note below)

Note:

- Restore from running server is not a supported method to migrate between hypervisors.
- For MiVoice Business, restore an ESM backup to ensure any customisation of the database is retained. Do not restore using an MSL backup.
- For all other MSL based applications, an MSL backup is required to restore the database.

10.1.4 Troubleshooting

Each Mitel virtual appliance is installed with its own built-in diagnostics for use in troubleshooting problems local to the specific Mitel virtual application.

Refer to the following product guides for troubleshooting help.

Table 27: Troubleshooting guides for Mitel products

Product	Guide containing troubleshooting information
MiVoice Business	MiVoice Business Technician's Handbook
MiVoice Business Virtual	MiVoice Business Voice Quality Troubleshooting Guide Installation Guide for MiVoice Business Virtual Engineering Guidelines for Industry Standard Servers and MiVoice Business Virtual
MiCollab	MiCollab Installation and Maintenance Guide
NuPoint UM	NuPoint Unified Messaging Technician's Handbook
MiCollab Client	MiCollab Client Engineering Guidelines
MiVoice Border Gateway MBG and MBG Virtual	MiVoice Border Gateway (MBG) Installation and Maintenance Guide
MiVoice Business Express	MiVoice Business Express Deployment Guide

10.1.4.1 Troubleshooting Virtual Appliances in Hyper-V Environment

For Hyper-V specific troubleshooting information, refer to the Hyper-V documentation, including <http://technet.microsoft.com/en-us/library/cc742454.aspx>.

10.1.4.2 Troubleshooting Virtual Appliances in VMware Environment

For Enterprise-level troubleshooting in the VMware environment, you must install VMware vCenter Virtualization Management. For more information about VMware vCenter Virtualization Management, see <http://www.vmware.com/products/vcenter-server/overview.html>.

10.1.4.2.1 Mitel Virtualization Diagnostic Tool (VDT)

Troubleshooting is especially difficult in the service provider environment in which the customer is renting computing capacity from the service provider. In this deployment model, the Mitel customer has no access to the service provider's virtual infrastructure. Isolation of issues within the virtual infrastructure from problems within the Mitel virtual solution is almost impossible without some level of cooperation from the service provider.

For the full VDT reference, see [Virtualization Diagnostic Tool \(VDT\)](#). This reference is also maintained in the *Mitel Standard Linux (MSL) online help*.

10.1.4.2.2 OAuth2 Access Token Issue on VMware ESXi

Symptom: OAuth2 Access Token has expired or is not yet valid, resulting in service disruptions.

Sample Error Message: "Office365 Oauth2 Graph Error _code:InvalidAuthenticationToken,message: Access token has expired or is not yet valid".

Environment:

- Mitel virtualized application running on VMware ESXi
- OAuth2 authentication with external services (e.g., Office 365)

Cause: The issue is related to time synchronization problems on the VMware ESXi host.

Voice Mail to E-Mail stops working.

Corrective Action:

1. Make sure the ESXi host uses the same NTP source and the time is correct on the ESXi Host Server.

2. Check the time on the ESXi Host Server using the following command via SSH or Putty:

```
# /usr/bin/vmware-toolbox-cmd stat hosttime
```

```
12 Dec 2023 17:15:07
```

3. Verify that the displayed time is accurate.
4. If the displayed time is not accurate, disable the periodic time synchronization in the VMware Tools using the following steps:

- a. Start **vSphere Web Client**

- b. Navigate to **Edit Settings > VM Options > VMware Tools > Synchronize Time with Host.**

- c. Deselect the checkbox **Synchronize time periodically.**

5. Use the following command to check VMware Tools time synchronization status:

```
# /usr/bin/vmware-toolbox-cmd timesync status
```

```
Disabled
```

6. Confirm that it reports **Disabled**.
7. If necessary, restart VMware Tools on the virtual machine to apply changes.

10.1.4.3 Collecting VMware Logs

There may be times when the logs from a Mitel Virtual Machine do not have enough information to diagnose a problem. In some of these cases it may be necessary to get the logs from vCenter or the hypervisor. This is done using the vSphere client in either the Windows or the Web client variant. Although typically more information is better than less, as a minimum, Mitel Product Support will need at least the **vmware*.log** for your VM and the **hostd.log** for the host the VM was running on.

Although you can collect logs using either the vSphere Windows client or the new Web client, VMware plans to drop support for the Windows client in the future, so it may be more appropriate to use the Web client.

VMware has documented the process for collecting diagnostics using the Web client here <http://kb.vmware.com/>.

This is the definitive reference for this process but for your convenience, the process is presented below.

To collect ESX/ESXi and vCenter Server diagnostic data using the Web client:

1. Start the **vSphere Web Client** and log in to the **vCenter Server** system.
2. Under **Inventory Lists**, select **vCenter Servers**.
3. Click the **vCenter Server** that contains the ESX/ESXi hosts from which you want to export logs.
4. Click the **Monitor** tab and click **System Logs**.
5. Click **Export System Logs**.
6. Select the **ESX/ESXi hosts** from which you want to export logs.
7. Optional: Select the Include vCenter Server and vSphere Web Client logs option.
8. Click **Next**.

9. Select the system logs that are to be exported. Make sure that the **vmware*.log** for your VM is acquired by selecting the check box **VirtualMachines/logs** and the **host.d log** is acquired by selecting **Logs/System**. If you are not concerned about the size of the log archive it would be best to select more options.
10. Select **Gather performance data** to include performance data information in the log files.

Note:

You can update the duration and interval time for the data collection.

11. Click **Next**.
12. Click **Generate Log Bundle**. The **Download Log Bundles** dialog appears when the **Generating Diagnostic Bundle** task completes.
13. Click **Download Log Bundle** to save it to your local computer.

Note:

The host or vCenter Server generates ZIP bundles containing the log files. The **Recent Tasks** panel shows the **Generate diagnostic bundles task** in progress.

14. After the download is complete, click **Finish** or generate another log bundle.

Note:

You can use the events log to troubleshoot issues, particularly when trying to determine a time stamp.

To export the events log:

1. Click **File > Export > Export Events**.
2. Browse to a location to save the log bundle.
3. Choose a time frame to export and click **OK**.
4. Choose the format in which to export and click **Save**.

VMware has documented the process for collecting diagnostics using the Windows client [here](#).

This is the definitive reference for this process but as a convenience the process is presented below.

To collect ESX/ESXi and vCenter Server diagnostic data using the Windows client:

1. Open the **vSphere Client** and connect to **vCenter Server** or directly to an **ESXi 5.x host**.
2. Log in using an account with administrative privileges or with the **Global Diagnostics** permission.
3. Select an **ESXi host**, cluster, or data center in the inventory.
4. Click **File > Export > Export System Logs**.

5. Select the host or group of hosts from the Source list, if a group of ESXi hosts are available in the selected context,
6. Click **Next**.
7. In the **System Logs** pane, select the components for which the diagnostic information must be obtained. To collect diagnostic information for all the components, click **Select All**. Make sure the `vmware*.log` for your VM is acquired by selecting the check box **VirtualMachines/logs** and the `host.d` log is acquired by selecting **Logs/System**. If you are not concerned about the size of the log archive, it would be best to select more options.
8. Select the **Gather performance data** option if required and specify a duration and interval.
9. Click **Next**.
10. In the **Download Location** pane, click **Browse** and select a location on the client's disk where you want to save the support bundle.
11. Click **Next**.
12. In the **Ready to Complete** pane, review the summary and click **Finish**. The **Downloading System Logs Bundles** dialog box appears and provides progress status for the creation and downloading of the support bundle from each source. A Generate system logs bundles task is created.

10.1.5 Virtualization Diagnostic Tool (VDT)

To simplify the gathering of diagnostic information, the Virtualization Diagnostic Tool provides the ability to extract some statistical information from the service provider virtual infrastructure automatically, and make this infrastructure data available.

VDT is delivered to all Mitel virtual appliances through the MVF blade with the exception of the Windows-based MiContact Center Office, and MiVoice Business Virtual. As VDT collects virtual infrastructure data on a periodic basis, it may impose additional CPU and network loading on the ESXi host, virtual network and vCenter.

To use the VDT, you need:

- Operating System: MSL 10.0 or later
- VMware environment: vSphere 4.1 or later
- Mitel Virtual Framework: MVF 2.0 or later

By default, Mitel virtual appliances are delivered with the **Virtualization Diagnostic Tool** activated.

10.1.5.1 Using the VDT

To use the VDT web page:

1. Log in to the **MSL Server Manager** on a Mitel virtual application that s VDT.

Note: Full VDT functionality requires Administrator level access to either the ESXi host or the vCenter that manages that host.

2. Navigate to **Administration > Virtualization** (left-side menu). The Mitel Virtualization page is displayed. From this page, you can view the **Diagnostic Overview**, which displays status for hypervisor support, OVA dimensioning, License Server connectivity and Last Nightly Analysis.

3. In the **Credentials and Diagnostics** section of the page, enter the vCenter or ESXi credentials to gain access to all the information available. Without these credentials, a subset of the information is still available. See [Table 24](#) for details.
4. Click **Run diagnostics** to initiate the manual analysis of seven days of Mitel virtual appliance statistics.

Note: The VDT is packaged with the Mitel Virtualization Framework, and is supported on, and run from, Mitel Standard Linux (MSL). This introductory release of the VDT is not supported on Microsoft Windows, but is available as a separate deliverable, if required.

10.1.5.1.1 Required credentials

The following table shows the VDT features, and the credential level required for each one.

Table 28: Feature by required credential level

Feature	Required credentials			
	No credentials required	ESXi credentials	vCenter credentials (read-only access)	vCenter credentials (Admin access)
Collect reservation limits (Guest SDK)	Yes	Yes	Yes	Yes
Collect performance statistics of the VM and the host	No	Yes	Yes	Yes
Collect the VM vHW info (CPU count, speed, memory, and number of NICs)	Yes	Yes	Yes	Yes
Collect task and events dump of the VM host	No	No	Yes	Yes

Feature	Required credentials			
	No credentials required	ESXi credentials	vCenter credentials (read-only access)	vCenter credentials (Admin access)
Analysis of performance counters, VM and host configuration host	Minimal - no performance data analysis	Uses performance counters only	Performance counters plus infrastructure info (HA, DRS)	Performance counters plus infrastructure info and licensing info

10.1.5.1.2 Virtual Machine Properties

The Virtual Machine Properties table displays information about the virtual machine and the Mitel Virtual Framework (MVF). The information is presented in two columns:

- **Current Dimensions:** Lists the configuration at the time that the current Mitel Virtualization page was loaded. Refreshing the page resets the settings.
- **First Boot Dimensions:** Lists the configuration after the Mitel OVA package has been installed and the settings configured, but before the virtual machine has been powered on for the first time.

Setting	Description
MVF Version	The version number of the Mitel Virtualization Framework (MVF), a software package that enable Mitel applications to run in a virtual infrastructure. MVF has the capacity to support multiple operating systems and hypervisor products.
Virtualization Agent Version (VMware Tools)	The version number of VMware Tools, a suite of utilities that enhances the performance of the virtual machine's guest operating system and improves management of the virtual machine.
Hypervisor Version	The version number of the VMware ESXi hypervisor that hosts one or more virtual machines and their guest operating systems.
vCPU count	The number of virtual Central Processing Units (vCPUs) configured on this virtual machine.

Setting	Description
Memory (MB)	The amount of virtual physical memory available for use by the operating system on this virtual machine.
Disk size (GB)	The virtual disk size available for use by the operating system on this virtual machine.
NIC count	The number of virtual network interface cards configured on this virtual machine.
CPU Reservation (MHz)	The guaranteed minimum allocation of CPU resources for this virtual machine.
Memory Reservation (MB)	The guaranteed minimum allocation of memory resources for this virtual machine.
CPU Limit (MHz)	The upper limit of CPU resources that can be allocated to this virtual machine. This limit is expressed in concrete units (Megahertz) and cannot be exceeded.
Memory Limit (MB)	The upper limit of for memory resources that can be allocated to this virtual machine. This limit is expressed in concrete units (Megabytes) and cannot be exceeded.
vCPU Speed (MHz)	<p>The speed of the virtual CPU, which is dependant on the speed of your underlying processor. So if you have a 12 cores and a processor speed of 3.36GHz, That means a virtual machine with a single vCPU running a single threaded application can consume 3.36GHz.</p> <p>This setting defines what a single vCPU will consume, not the aggregated amount among multiple vCPUs on a single virtual machine. Accordingly, if you have two vCPUs this figure should be doubled.</p>

10.1.5.1.3 Storage Monitoring

Use the following settings to specify whether the tool performs file system monitoring and whether it reboots on read-only state.

Setting	Description
File System Monitoring	Use this setting to specify whether the system monitoring is enabled or disabled. If the main feature is enabled (the default), the system will check the file system for disk errors every five seconds. If any disk errors are discovered, the system will generate a warning notification and send it to the administrator e-mail address configured on the Email Settings screen.
Reboot on read-only state	<p>Select Yes or No to specify whether the system reboots when it enters a read-only state. A read-only state occurs when there are I/O errors on the virtual machine disk drives.</p> <p>The default setting is Yes.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #004a87; margin-top: 10px;"> <p>Note: File System Monitoring must be enabled for this feature to be effective.</p> </div>

10.1.5.1.4 Diagnostic Overview

The VDT constantly monitors the system in order to report on three alarm conditions and the state of the last nightly analysis.

Setting	Description
Hypervisor Version	<p>Indicates whether or not the version of VMware ESXi Hypervisor is supported. The Hypervisor is also known as the Virtual Machine Monitor (VMM).</p> <p>If your ESXi version is not supported and you must switch to a supported version in order to restore monitoring functionality. For example, if you are running ESXi 4.0 or earlier, you must upgrade to version 4.1 or later.</p>

Setting	Description
Current Dimensions	<p>Indicates whether the currently configured application resource dimensions are supported.</p> <p>If your configuration is not supported due to a setting (vCPU count, Memory, Disk size, or NIC count) being out of boundaries, you can resolve any performance issues, by doing the following:</p> <p>Revert to the default configuration for your deployment.</p> <p>Contact Mitel Product Support for assistance.</p>
License Server Connectivity	<p>Indicates whether the Virtual Machine can connect to the Mitel License Server for licensing purposes.</p> <p>The two states are Connected and Error. In the Error state, the VM cannot connect to the License Server. Check the networking configuration and Serial ID. See the Mitel Standard Linux Installation and Maintenance Guide for more information.</p>
Last Nightly Analysis	<p>Indicates the date and time that the last nightly analysis was completed, and whether any problems occurred while it was being run. Upon successful completion, the nightly analysis generates the following log file:</p> <p>NIGHTLY-REPORT-YYYY-MM-DD.txt</p>

10.1.5.1.5 Virtualization Diagnostics Credentials

To enable the VDT to collect statistics for the virtual machine and the host, and then use the statistics to generate log files, you must enter credentials for the vCenter server or ESXi hypervisor.

The information collected depends on the credentials entered:

- Admin login to vCenter - full range of features and statistics.
- Read-only login to vCenter - subset of features and statistics.
- Read-only login to ESXi - subset of features and statistics.
- No credentials - Allocation and Reservation & Limits information only.

Note:

For optimum results, enter credentials for the vCenter. Entering credentials for the ESXi may result in connectivity problems if settings are changed on the hypervisor.

10.1.5.1.5.1 Enter New Credentials

To enter the virtualization diagnostics credentials:

1. Under **Administration**, click **Virtualization**.
2. Under **Virtualization Diagnostics**, enter the following settings.
3. Click **Save**.

Setting	Description
FQDN or IP address	Enter the Fully Qualified Domain Name or IP address of the vCenter or ESXi hypervisor.
Username	Enter the user name required to access the vCenter or ESXi hypervisor.
Password	Enter the password required to access vCenter or ESXi hypervisor.
Nightly Analysis Time	Specify the one-hour period during which the nightly analysis will be run each day. Select hours between 0-1 and 23-24. Upon successful completion, the nightly analysis generates the following log file: NIGHTLY-REPORT-YYYY-MM-DD.txt.

After a connection is established, the system obtains performance statistics for the virtual machine and the host. You can click Run Diagnostics to manually generate log files and an on-line report.

Note:

For a newly installed system, wait for at least 15 minutes before clicking **Run Diagnostics**.

10.1.5.1.5.2 Remove Current Credentials

To remove the virtualization diagnostics credentials:

1. Under **Administration**, click **Virtualization**.
2. Under **Virtualization Diagnostics**, click **Remove**. You may now enter new credentials.

Note: Without credentials, the system will not collect statistics or generate log files for virtualization diagnostics.

10.1.5.1.6 Log Files

The system generates log files containing performance and configuration data, and statistical events.

10.1.5.1.6.1 Automatically Generated Log Files

The following log files are generated automatically by the system on a periodic basis.

Report Name	Description
NIGHTLY-REPORT-YYYY-MM-DD.txt	This report contains the previous day's detailed performance and configuration information, and is generated daily in the Nightly Analysis Time you have specified. The system retains seven reports, deleting the oldest file after seven days.
VM-STATS-YYYY-MM-DD.csv	This report contains virtual machine statistics for the previous week. The system polls for new data every 15 minutes and deletes existing data after seven days.
HOST-STATS-YYYY-MM-DD.csv	This report contains host system statistics for the previous week. The system polls for new data every 15 minutes and deletes existing data after seven days.
ALL-CONFIG-YYYY-MM-DD.csv	This report contains all CPU, performance and network configuration statistics concerning the host and virtual machine for the previous week. The system polls for new data every 15 minutes and deletes existing data after seven days.

10.1.5.1.6.2 Manually Generated Log Files

A number of log files are created when you request them. To manually generate the log files and an online report:

1. Under **Administration**, click **Virtualization**.
2. Under **Virtualization Diagnostics**, click **Run Diagnostics**.

Note:

- For a newly installed system, allow it to collect statistics for at least 15 minutes before you click Run Diagnostics.
- If you repeatedly click Run Diagnostics, you may exceed the storage capacity of the host server's hard drive.

Report Name	Description
USER-REPORT-YYYY-MM-DD-HH-MM-SS.txt	<p>This report is similar to the NIGHTLY-REPORT-YYYY-MM-DD.txt report, but it detailed performance and configuration information for the previous week (rather than a single day), collected from the moment you click the Run Diagnostics button.</p> <p>The report file is retained for seven days and then deleted.</p> <p>If you run into a problem you cannot resolve, Mitel Product Support is likely to ask you to sent this log file to them for analysis</p>
USER-SUMMARY.tmp	<p>This report is an abbreviated version of the USER-REPORT-YYYY-MM-DD-HH-MM-SS.txt report. It contains performance and configuration overviews for each day of the previous week.</p> <p>This report is presented in two formats:</p> <p>Displayed on the Mitel Virtualization screen. This report is retained until you navigate away from the screen.</p> <p>Recorded in the log files. This file is retained until the Run Diagnostics button is clicked again.</p>

Report Name	Description
VM-EVENTS-YYYY-MM-DD.csv	The report contains 15 days' activity regarding the operation of the Virtual Machine. This file is retained until Run Diagnostics is clicked again.

10.1.5.1.6.3 Log File Contents

Although the log files are primarily intended for use by Mitel Product Support, you may use them to troubleshoot basic issues with the following issues:

- **Performance issues:** The system analyzes performance data and if it detects five consecutive “out of bounds” events, an problem will be reported. For example, if the virtual machine waits longer than two seconds to be serviced by the host, five times in a row, the system will report a “CPU Ready” error. System events are registered every twenty seconds.
- **Configuration issues:** The system checks configuration data and statistical events on an ongoing basis. If a problem is found, and error is logged immediately.

See the MSL online help for detailed information concerning the system settings which control the generation of log file problems.

Performance issues	Description
CPU Ready (seconds)	The virtual machine has exceeded the maximum amount of time that it can wait to be run on the physical CPUs. The default is 2 seconds
CPU Usage (percent)	The virtual machine has exceeded its CPU capacity limit, which is expressed as a percentage of the total amount available. For example, with a limit of 50%, if the virtual machine has four CPUs with 2 GHz processors, and you are running an application that requires 6 GHz (75% of capacity), the limit has been exceeded by 25%. The default is 50%.
Disk Latency (seconds)	The virtual machine has exceeded the maximum amount of time permitted for a SCSI command to be issued by the guest operating system to the virtual machine hard disk. The default is 0.02.

Performance issues	Description
Network Usage (MB)	The virtual machine has exceeded the maximum network utilization (combined transmit and receive rates, in Megabytes per second). The default is 50.0 MB.
Memory Swapped (MB)	The virtual machine has exceeded the maximum amount of memory, in Megabytes, that can be swapped into memory from disk. The default is 0 MB.
Memory Use (MB)	The virtual machine has exceeded the maximum amount of memory capacity that it can use, expressed as a percentage of the total amount available. For example, if the virtual machine has 4 GHz of memory, and you are running an application that requires 3 GHz (75% of capacity), an event will be registered. The default value is 50%.
Number of Packets Dropped (average)	The virtual machine has exceeded the maximum number of received packets that can be dropped at the network interface. The default value is 0.
Disk Usage (MB)	The virtual machine has exceeded the maximum amount of data, in Megabytes per second, that can be read from the virtual machine hard disk. The default value is 30 MB.
Configuration Detection	Description (Yes/No)
High VM-to-host CPU ratio	If Yes is displayed, the ESXi host has exceeded the virtual CPU to host CPU ratio, which is 0.79 by default. For example, if five virtual machines with 4 GHz vCPUs are powered on, and the host has 8 physical/16 logical cores, then the ratio is $4 + 4 + 4 + 4 + 4 \div 16 = 1.25$. Since 1.25 exceeds 0.79, a potential configuration issue is detected.

Performance issues	Description
High VM-to-host Memory ratio	If Yes is displayed, the ESXi host has exceeded the virtual memory to host memory ratio, which is 1.20 by default. For example, if five VMs are powered on, each using 2 GHz of memory, and the host has 8 GHz of physical memory, then the ratio is $2 + 2 + 2 + 2 + 2 \div 8 = 1.25$, which will cause an event to be registered.
Snapshots Present	If Yes is displayed, the system checks to determine if snapshots are supported on the virtual machine. Because snapshots create considerable disk I/O load, use of this feature may degrade the voice quality of calls.
Low CPU Speed (MHz)	If Yes is displayed, the maximum speed of the virtual CPU, which is dependant on the speed of the underlying processor on the ESXi host, has been exceeded.
No Hyperthreading (Ignore if running on non-Intel processor)	If Yes is displayed, the system checks to determine if hyperthreading is enabled on the ESXi host. Note: This parameter can only report on Intel processors that support hyperthreading. It cannot report on AMD or other non-Intel processors.
vMotion occurred	If Yes is displayed, the system checks to determine if vMotion is enabled on the ESXi host.
Low CPU Reservation (MHz)	If Yes is displayed, the guaranteed minimum allocation of CPU resources for this virtual machine has been exceeded.
Low Memory Reservation (MB)	If Yes is displayed, the guaranteed minimum allocation of memory resources for this virtual machine has been exceeded.

10.1.6 Known Issues (VMware infrastructure)

This section describes known issues with Mitel virtual appliances running in a VMware infrastructure and describes fixes or work-arounds, where available.

The following known issues are discussed here:

- Mitel application performance issues when CPU and memory requirements are satisfied
- Mitel Virtual Appliance performance becomes sluggish in the presence of VMware snapshots
- File system becomes read-only, causing Mitel Applications to go out of service
- Virtual Machine is no longer responsive, and is at or close to 100% CPU usage
- Slow network throughput for streaming or large file transfers
- Mitel virtual machines are blocked during power up due to checking file systems un-expected inconsistency error
- Mitel virtual machines experience issues on power up after powering off
- Mitel resiliency solutions not working (such as MiVoice Business resiliency or MBG clustering)
- Mitel application message record/playback or voice prompt issues
- Mitel application repeatedly resets after startup, or voice quality issues

10.1.6.1 Performance issues after CPU and memory requirements are met

There can be cases where, even though sufficient CPU and memory resources are available for Mitel virtual machines, Mitel applications appear to be performing poorly, with problems such as poor quality voice streaming, message playback disruptions, or dropped calls.

Probable cause

This can often be attributed to one of the following causes:

1. I/O to and from the virtual machine disk (vmdk) is slow, high latency, or intermittent.
2. Host(s) may not meet minimum processor requirements or are not properly configured for high performance.
3. Share or reservation settings may not be correctly configured and applied.

Resolution

Use one or both of the following resolutions:

1. Examine the Virtual Machine Disk Configuration and Statistics and address any exceptions.
2. Examine host processor, BIOS, and related processor settings, to ensure that they meet the requirements listed in [Table 2, "Minimum hardware requirements,"](#). Correct any settings or move the Mitel applications to fully qualified hosts.
3. 'CPU Ready' values that consistently exceed 100ms (<0.5%) of delay indicate onset of CPU resource starvation which may impact the real time operation of voice and video streaming services. Lower values of 'CPU Ready' may still be impacted. Ensure that the CPU resource reservations are correctly configured and applied. Ensure that the resources are not oversubscribed.

10.1.6.2 VMware Snapshots Slowdown Performance

Mitel Virtual Appliance performance is sluggish, possibly after upgrade. Depending on the virtual appliance, this could appear as poor voice quality.

Probable Cause

Snapshots were created before upgrade, and were not deleted before restarting the Mitel Virtual Appliance. The more snapshots there are, the more the performance will be degraded.

Resolution

To restore performance, delete all of the VM snapshots.

10.1.6.3 Read-only Files Render Mitel Applications Inactive

Mitel Applications may cease to function if MSL switches the file system to “read-only” in an effort to prevent file system corruption. Before concluding that a read-only file system is the problem, confirm that MSL has switched to read-only by selecting the Mitel virtual machine, and clicking Launch Virtual Machine Console. Observe the output, which may indicate that the file system is now read-only. The console is the only place where you will be able to see the read-only trigger, since logs cannot be written to disk in the read-only state.

Probable Cause

Very high disk latency between the ESXi host and the virtual machine's data store over short time period can trigger file system protection and the read-only state in MSL.

Resolution

To restore service immediately, right-click on the Mitel virtual machine, select Power and Restart Guest. Use the Virtual Machine Console to confirm that system shut-down and start-up occur before confirming that service is restored.

To prevent the problem from recurring, determine whether a SAN connection disruption occurred, and examine Virtual Machine Disk Configuration and Statistics. Address any failures or exceptions.

It is very important to make sure the system virtual disk can sustain the I/O throughput requirement and the disk latency requirement. Because the physical storage is shared, performance of storage could be momentarily reduced or degraded as the workload increases.

Other external factors can also affect storage performance. For example, congestion in the storage network induces network latency that, in turn, increases the I/O latency on storage devices. In general, storage device average latency should not be more than 30 ms. Performance of virtual machine will be affected as storage latency increases.

If SSD is used as part of the storage device, the maximum acceptable average latency should be less. Disk I/O latency can be monitored by VMware vCenter Server. vCenter Server can trigger actions (send notification to system administrators, for example) if it sees a degraded or abnormal latency on storage devices. For example, vCenter server can be set up to send warnings if disk latency exceeds 1000 ms for one minute, and send alerts if disk latency exceeds 3000 ms for one minute. The following procedure is required to enable Storage I/O monitoring in VMware vCenter Server from vSphere Client.

Note:

The threshold values are just examples. The administrator of the virtual infrastructure must determine the actual values suitable within the environment. The steps that follow are based on vSphere 5.5 and may be slightly different in other versions.

1. Login to **vCenter Server** from **VMware vSphere** (Windows) Client.
2. On the left pane, highlight the **virtual machine** to be monitored.
3. Select the **Alarm** tab on the right pane.
4. Under the row of tabs, select **Definitions** view.
5. Right-click on the empty space on the right pane, and select **New Alarm**.
6. In the **General** tab, fill in the **Alarm name** and **Description**. In **Alarm Type**, select:
 - a. Virtual Machine
 - b. Monitor for specific conditions or state
 - c. Enable this alarm
7. In the **Triggers** tab, right-click empty space to **Add Trigger**. A trigger entry is created. Set the following conditions:
 - a. Trigger Type to VM Max Total Disk Latency (ms).
 - b. Condition to **Is Above**
 - c. Warning to **1000**
 - d. Condition Length to **for 1 minute**
 - e. Alert to **3000**
 - f. Condition Length to **for 1 minute**
8. Select **Trigger** if any of the conditions are satisfied.
9. In the **Reporting** tab, enter **0** in both text boxes.
10. In the **Actions** tab, right-click on empty space and select **Add Action**.
11. Change Action to **Send a notification email**:
 - a. Enter **e-mail address** in the **Configuration** field.
 - b. Select **Once** for all condition transitions.
12. Right-click on the empty space to add another action, if needed.
13. Click **OK** to close the dialog box. There may be a **vCenter e-mail settings are not configured** popup. Click **OK** to close all of the dialog boxes.
14. If vCenter e-mail settings are not configured:
 - a. In vCenter, go to **Administration > vCenter Server settings > Mail**.
 - b. Enter the **SMTP Server name** and the **Sender Accountname**.

System administrators should be notified when storage performance degradation begins. This should allow enough time for them to correct the problem before storage performance falls to an unacceptable level.

10.1.6.4 Unresponsive Virtual Machine and CPU Usage Close to 100%

A Mitel Virtual Machine consumes CPU reaching near 100% continuously, and becomes unresponsive. Memory consumption may also increase to near 100%.

Probable Cause

The SAN volume hosting the Mitel virtual machine's vmdk is no longer available, possibly due to network connectivity or SAN system failure.

Resolution

Confirm that the SAN volume is no longer accessible from the ESXi host running the Mitel Virtual Machine. This should be reported next to the data store name in the vSphere client, such as:

datastore1 (inactive)

or

datastore1 (inaccessible)

Shut down (Shut Down Guest) the Mitel virtual machine, if it is responsive. If it is not responsive, plan and execute a restart of the ESXi host. Investigate and address the issue with the SAN volume.

10.1.6.5 Slow Network throughput for Large File Transfers and Streaming

Large file transfers to or from a Mitel virtual machine may have very slow throughput, or streaming media between multiple Mitel virtual machines may be poor quality or interrupted.

Probable Cause

Large Receive Offload (LRO) may be enabled for the VMware vmxnet3 kernel module within MSL. The vmxnet3 kernel module is a software driver used to support all VMxnet3 virtual network cards used within a virtual machine.

Some Linux-based operating systems such as MSL do not support LRO. MSL will ignore LRO packets and they will require re-transmission.

Resolution

For up-to-date information on this issue, check the VMware knowledge base at <http://kb.vmware.com/>.

10.1.6.6 Virtual machines blocked during power up for file system inconsistency error checks

After deploying a Mitel virtual appliance and powering on the Mitel virtual machine created, the operating system boot does not finish successfully. When examining the Virtual Machine Console, the Checking file systems command will indicate FAILED and report an error similar to the following output will be seen:

```
Checking filesystems
/dev/mapper/VolGroup-lv_root: Superblock last mount time (Thu Nov 22 06:48:26 20
12,
    now = Fri Oct 22 05:52:34 2010) is in the future.

/dev/mapper/VolGroup-lv_root: UNEXPECTED INCONSISTENCY: RUN fsck MANUALLY.
(i.e., without -a or -p options)
                                [FAILED]

*** An error occurred during the file system check.
*** Dropping you to a shell; the system will reboot
*** when you leave the shell.
Give root password for maintenance
(or type Control-D to continue): _
```

Probable Cause

The ESXi host that the Mitel virtual appliance was deployed to does not have the correct system time set. The incorrect time is pushed from the ESXi host to the virtual machine on power up. This time is earlier than the last one recorded in the MSL file system. This causes interactive file maintenance to be performed, however, this requires a root password which is not yet set if this is the first boot of the Mitel Virtual Machine.

Resolution

1. Shut down (**Shut Down Guest**) the Mitel virtual machine.
2. Change the system time of the ESXi server hosting the Mitel Virtual Machine to the current time.
3. Power on the Mitel virtual machine and observe MSL boot-up in the Virtual Machine Console.

10.1.6.7 Issues with powering on Virtual Machines after shut down

On **Power On** of a Mitel virtual machine that was shut down with a Power Off, startup delays, file system repairs, or startup failures occur.

Probable Cause

The Power Off request disrupted the running applications or system state in the Mitel virtual machine. The virtual machine disk may be corrupted or in an inconsistent state

Resolution

Always use **Shut Down Guest** to turn off the virtual machine. This is a safe, graceful shutdown. Do not use Power Off, since this effects an immediate shutdown, equivalent to sudden powering off of a physical server.

10.1.6.8 Resiliency Solutions Do Not Work

Mitel resiliency solution is not exhibiting the correct resiliency behavior when an ESXi host or its dependent resources become unavailable.

Probable Cause

Mitel virtual machines configured for resiliency must reside on different ESXi hosts to ensure that one ESXi-related outage does not affect all Mitel virtual machines in a resilient configuration.

Resolution

Ensure that Mitel virtual machines configured for resiliency never run on the same ESXi server. If deploying a resilient solution to one ESXi cluster, ensure that Distributed Resources Scheduler (DRS) anti-affinity rules are set up for the ESXi cluster to ensure that the Mitel Virtual Machines are distributed among different ESXi hosts in the cluster. Also see [Anti-affinity rules](#) and [Other virtual appliance management considerations](#).

10.1.6.9 Issues with Message Recordings, Voice Prompts, and Playback

Mitel applications appear to be performing poorly with respect to message record and playback, prompt playing, and so on, with problems such as poor quality voice record or playback, message playback disruptions, or similar.

Probable Cause

This can often be attributed to high disk latency or slow storage IOPs, or intermittent disk access disruptions. This could include issues with the overall storage environment, or issues in configuration of the VM's virtual disk.

Resolution

Examine and address issues in Virtual Machine storage configuration.

1. Select **Edit Settings** for the Mitel virtual machine, and under the **Hardware** tab, select each Hard disk used, and confirm that **Disk Provisioning Type** is set to set to **Thick Provision Lazy Zeroed** or **Thick Provision Eager Zeroed**. It should NOT be set to Thin Provisioned. If Thin Provisioned is

set, conversion to one of the Thick provisioning options is recommended. Convert to a Thick Provision method by moving the virtual machine to another data store:

- a. Select **Migrate, Change Datastore** option.
- b. Change **Select a virtual disk format** to **Thick Provision Lazy Zeroed**.
- c. Select a different data store.
- d. Complete the wizard to start the migration.
- e. Select the Mitel virtual machine, open Snapshot Manager and ensure that there are no snapshots. If there are snapshots listed, verify that all snapshots are no longer required; then click **Delete All** to remove them.
- f. Select the Mitel Virtual Machine, click the **Performance** tab, select the **Advanced** button and investigate the Datastore and Disk statistics for latency and read/write throughout. (Write Latency, Read Latency and/or Highest Latency).
- g. Examine these values over the time period when issues were experienced. Values should be consistently near 0 ms latency. Consistent latency responses above 30 ms or peak values over 80 ms can result in many performance problems. Read/Write throughput should meet or exceed requirements for the application if running at full load (see [Table 6](#)). If high latency or low throughput is observed, it is highly recommended that you locate and migrate the virtual machine to a data store that is performing better.

10.1.6.10 Repeated Application Resets after Startup, and Voice Quality Issues

Mitel logs files show missing interrupts, high clock skew, hardware watchdog warnings or similar.

Probable Cause

Host(s) may not meet minimum processor requirements or not properly configured for high performance.

Resolution

Examine host processor, BIOS and related processor settings, to ensure they meet requirements listed in [Table 2, "Minimum hardware requirements,"](#). Correct any settings or move to fully qualified hosts.

This chapter contains the following sections:

- [Glossary](#)
- [Deployment Types](#)

11.1 Glossary

Following is a glossary of terms and acronyms:

Acronym or term	Definition
5-tuple	<p>5-tuple is a term used in computer networks to refer to a set of five different values that make up a Transmission Control Protocol/Internet Protocol (TCP/IP) connection.</p> <p>The 5-tuple is made up of source IP address, destination IP address, source port number, destination port number and the protocol in use.</p>
ACD	Automatic Call Distribution
ARID	Application Record ID - A bundle of Mitel licenses managed as single set with a unique identifier; used for a single deployment of MSL.
CC vCC vCCM vCCS	Mitel Contact Center Virtual Contact Center Virtual Contact Center Management Virtual Contact Center Solutions (now called MiContact Center)

Acronym or term	Definition
CSM vCSM	Customer Service Manager (now called MiContact Center Office) - A contact center solution available exclusively on MiVoice Office (previously known as Mitel 5000 CP), a communications solution for small and medium-sized businesses. It enables basic contact centers or workgroups to efficiently monitor, manage, and route calls. It provides real-time business intelligence, including call performance and agent activity reporting, as well as agent productivity tools, including screen pop and Personal Information Manager (PIM) integration.
DLM	Designated License Manager Starting in MCD Release 6.0, vMCD and MiVoice Business 7.0+ can act as a DLM, able to host license sharing for a group of MCDs. In MCD Release 5.0 and previous releases, MCD can run a License Manager that can participate in group license sharing, but cannot act as a DLM.
DMZ	De-Militarized Zone: A computer host or small network inserted as a “neutral zone” between a company’s private network and the outside public network. It prevents outside users from getting direct access to a server that contains company data.
DRS Distributed Resource Scheduler	VMware Distributed Resource Scheduler - A VMware feature allowing dynamic load sharing across hosts within a vSphere cluster.
dvSwitch	Distributed vSwitch - VMware vCNS fully featured vSwitch, configured at the data center level across multiple hosts. See also vSwitch.
ESM	Embedded System Management; also called the System Administration Tool
ESX	Main Hypervisor from VMware. Phased out in favour of ESXi in release 4.1.

Acronym or term	Definition
ESXi	The latest variant of ESX. It has a smaller footprint because it does not require or include the ESX Service Console.
FT Fault Tolerance	VMware Fault Tolerance feature - A VMware feature that improves on High Availability (HA) by providing zero downtime and zero data loss in the event of a failure. A secondary virtual machine (VM) mirrors the execution state of the primary VM (a pair of VMs running in lock step, same IP, same MAC address, and so on). In the event of a failure, the secondary VM picks up execution without any loss of connectivity or transactions. Fault Tolerance introduces large latencies in the running virtual machines. Fault Tolerance is not suitable for Mitel applications.
guest	The guest, or guest O/S, is the operating system that runs inside a virtual machine. It is almost completely unaware it is running on a virtual platform.
HA High Availability	<p>VMware High Availability feature - In the event of physical server failure, affected virtual machines are automatically restarted on other production servers with spare capacity. In the case of operating system failure, VMware HA restarts the affected virtual machine on the same physical server.</p> <p>For a similar service, but without the downtime of HA see Fault Tolerance (FT). For fault tolerance at the data center scale, see Site Recovery Manager (SRM).</p>
host	The host is the hardware the hypervisor and its virtual machines reside on and run on.
hypervisor	A software entity that implements virtualization by presenting a virtual operating platform to one or more guest operating systems and monitors their execution. Also called virtual machine monitor (VMM). Examples are VMware ESXi and Microsoft Hyper-V Server (standalone).

Acronym or term	Definition
Hyper-V	<p>Hyper-V is a Microsoft Windows native hypervisor that enables platform virtualization on x86-64 systems.</p> <p>Hyper-V exists in two variants:</p> <p>As a stand-alone product called Hyper-V Server: Four major versions have so far been released: Hyper-V Server 2012 R2 (containing the current release of Hyper-V), Hyper-V Server 2012, Hyper-V Server 2008 R2 and Hyper-V Server 2008.</p> <p>As an install able role in Windows Server 2016, Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008 and the x64 edition of Windows 8 Pro.</p>
ISO image	<p>An ISO image is an archive file of an optical disc, a type of disk image composed of the data contents of every written sector of an optical disc, including the optical disc file system. ISO image files usually have a file extension of .iso.</p>
I/O	Input/Output
LBG	<p>Live Business Gateway - Enables enterprises to implement IP-based solutions. Working with Microsoft® Office Communications Server 2007 R2 or Microsoft Lync™ 2010 Server, a SIP-based presence and collaboration server, the Mitel / Microsoft partnership delivers advanced applications for global communication and collaboration.</p>
Licensing Mechanism	<p>Refers to the procedures applications use to acquire software licenses from the licensing authority, and the delivery and maintenance of the software licenses. At Mitel, the licensing mechanism involves ARID activation, synchronization and de-activation. Current Mitel licensing mechanisms allow software licenses to be obtained online or offline, but only the online method is supported for virtual appliances.</p>

Acronym or term	Definition
Licensing Model	Refers to the business logic and technologies about the scheme on how a valid software licenses are determined and granted to applications. The licensing model may change from time to time to reflect the current business need. At Mitel this intelligence is held in the License Server.
Linux Integration Tools	A set of drivers that enable synthetic device in supported Linux virtual machines under Hyper-V.
LM	License Manager
MAS vMAS	Mitel Application Suite (now called MiCollab)
MBG MBG Virtual	MiVoice Border Gateway - A highly scalable and optionally redundant solution that gives remote workers, road warriors, and day-extenders seamless access to the voice and data capabilities of the office, wherever they are.
MCA	Mitel Collaboration Advanced (now called MiCollab Audio, Web, and Video (AWV))
MCD vMCD	Mitel Communications Director (now called MiVoice Business) - Highly scalable, proven IP-PBX software; the foundation of the Mitel Freedom Architecture.
MICD	Multi-Instance Communications Director (now called MiVoice Business Multi-instance) - Delivers multiple instances of Mitel's award-winning MCD call control software on a single industry standard server.
MiCollab	Previously known as Mitel Application Suite (MAS)
MiCollab Audio, Web and Video Conferencing	Previously known as Mitel Collaboration Advanced (MCA)

Acronym or term	Definition
MiCollab Client	Previously known as Unified Communicator Advanced (UCA)
MiCollab with Voice	An appliance that includes MiVoice Business with MiCollab applications. Now called MiVoice Business Express.
MiVoice Business MiVoice Business Virtual MiVoice Business for ISS MiVoice Business Multi-instance	Previously known as: Mitel Communications Director (MCD) Virtual Mitel Communications Director (vMCD) Mitel Communications Director for Industry Standard Servers (MCD for ISS) Multi-instance Communications Director (MICD)
MiVoice Business Express	A virtual appliance that includes MiVoice Business with MiCollab applications.
MiContact Center Business MiContact Center Enterprise	Previously known as: Contact Center Business Edition MiContact Center Solidus
Mitel CX	Mitel Customer Experience
Mitel Licenses & Services	License Server used for Mitel license management - Mitel Internet service to which Mitel applications authenticate themselves and retrieve license information specific to their unique ARID, as entered by the Installer.
MiVB	MiVoice Business
MiVoice Office	Previously known as Mitel 5000 CP.

Acronym or term	Definition
MSL	Mitel Standard Linux - A Linux distribution based on Red Hat Linux, created and maintained by Mitel, and on which most Mitel applications are developed.
MVF	Mitel Virtual Framework - An install-able component that delivers all the generic Mitel components to support running in a virtual infrastructure.
NAT	Network Address Translation - A technique in which a router or firewall rewrites the source and/or destination Internet addresses in a packet as it passes through, typically to allow multiple hosts to connect to the Internet via a single external IP address. NAT keeps track of outbound connections and distributes incoming packets to the correct machine.
NFS	Network File System - A distributed file system protocol originally developed by Sun Microsystems, allowing a user on a client computer to access files over a network in a manner similar to the way local storage is accessed.
NP UM	<p>NuPoint Unified Messenger - the Mitel voice mail product. Sometimes written as NPM.</p> <p>NOTE: NuPoint UM is no longer available in stand-alone mode. NuPoint is still a part of MiCollab.</p>
Offline Licensing	<p>Offline licensing refers to the scenario where Mitel products obtain software licenses by generating and saving information required for another piece of software to send licensing requests to the License Server, and receive licensing responses from the License Server. Software licenses from license responses are copied back to the virtual appliances to enable software features.</p> <p>Offline licensing is not supported for virtual appliances.</p>

Acronym or term	Definition
Online Licensing	Online licensing refers to the scenario where Mitel virtual appliances obtain software licenses by sending licensing requests to the License Server and receiving licensing responses from the License Server through networking. License Server licensing traffic could travel directly between the nodes or through a licensing proxy.
Open virtual appliance	Defined by VMware as “an appliance that is accessible to end customers for modifications and edits. The O/S and the application, for example, can be patched individually and agents may be installed at the customer's site.”
OVA	Open Virtualization Archive format - a compact, single 'tar' file packaging of OVF.
OVF	Open Virtualization Format - A format for the packaging and distribution of one or more virtual machines. A collection of items in a single folder, most commonly a description file (*.ovf), a manifest file (*.mf), and virtual machine state files (*.vhd or *.vmdk)..
PCoIP®	PC-over-IP protocol: Teradici™ - VMware proprietary remote display virtual desktop presentation protocol.
pCPU	Physical CPU - Refers to the number of logical CPUs, or threads, that are available in a machine, rather than actual physical cores. For example, a CPU with hyperthreading enabled might have 4 real physical cores, but will report on 8 pCPU. Also see vCPU.
PIM	Personal Information Manager
PSTN	Public Switched Telephone Network -The network of the world's public circuit-switched telephone networks.

Acronym or term	Definition
RDN	Remote Directory Number - Remote portable directory numbers belong to devices that are connected to remote cluster elements. In relation to the network element that you are programming, all other elements in the cluster network are remote elements.
RDN Synchronization Mode	RDN Synchronization Mode is the data model used for post-4.0 releases of the MCD.
RDP	Remote Display Protocol: Microsoft proprietary remote display virtual desktop presentation protocol.
RTP	Real-time Transport Protocol version: The standard protocol for carrying media streams such as voice audio and video between end-points.
SAN	Storage Area Network - A high-speed special-purpose network (or subnetwork) that interconnects different kinds of data storage devices with associated data servers, on behalf of a larger network of users. See also vSAN.
SBC	Session Border Controller
SLA	Service Level Agreement
SMB	Small and Medium-size Businesses
Software upgrade	Software upgrade requires a valid software license, a copy of the new software and possibly some additional resources. This document focuses only on the licensing portion of the upgrade.
SRC	Secure Recording Connector - A feature of MiVoice Border Gateway that facilitates the recording of Mitel-encrypted voice streams by third-party call recording equipment.

Acronym or term	Definition
SRM	VMware Site Recovery Manager - A VMware product that allows the virtual machines at one site to fail over to a secondary site if the primary site experiences a catastrophic outage. This is not a hot failover; the secondary site will take 10's of minutes or more to become fully functional.
UCA vUCA	Mitel Unified Communicator Advanced (now called MiCollab Client)
VA	See Virtual appliance.
VADP	vStorage APIs for Data Protection
vApp	A VMware term that describes a solution that contains multiple Virtual Machines configured to work together for a cloud environment, and packaged as an OVF. Not to be confused with a VA.
vCA	MiVoice Call Accounting - Included in MiContact Center
vCC vIVR	Mitel Virtual Contact Center (now called MiContact Center) - Allows easy deployment and management of a virtual contact center to achieve superior customer service, effective resource utilization, and business continuity. Includes Virtual MiVoice Call Accounting (vCA) and Virtual Intelligent Voice Routing (vIVR).
vCD	vCloud Director - VMware product
vCNS	VMware vCloud Networking and Security - VMware security and networking components of vCloud Director Suite, including vShield Manager, Edge, App, and Endpoint components.

Acronym or term	Definition
vCPU	Virtual CPU - Refers to the number of virtual processor cores that a virtual machine needs to operate correctly. When a virtual machine is scheduled to run, the hypervisor maps its vCPUs to physical cores on the host.
vCSM	Virtual Customer Service Manager (now called MiContact Center Office) - A contact center solution, available exclusively on MiVoice Office. It enables basic contact centers or workgroups to efficiently monitor, manage, and route calls. It provides real-time business intelligence, including call performance and agent activity reporting, as well as agent productivity tools, including screen pop and PIM integration.
vDC	Virtual Data Center - Pool of resources dedicated to Provider (system level) or Organizations.
Virtual appliance	Defined by VMware as: “a pre-built software solution, comprised of one or more virtual machines that is packaged, maintained, updated, and managed as a unit.”
vApp	Virtual appliance - A virtualized application, or solution, consisting of one or more VMs, packaged in OVF/OVA format.
VDT	Virtualization Diagnostic Tool - Simplifies the gathering of diagnostic information by providing the ability to extract statistical information from the service provider virtual infrastructure automatically for use.

Acronym or term	Definition
VHD	Microsoft Virtual Hard Disk - The Virtual Hard Disk (VHD) format is a publicly-available image format specification that allows encapsulation of the hard disk into an individual file for use by the operating system as a virtual disk in all the same ways physical hard disks are used. These virtual disks are capable of hosting native file systems (NTFS, FAT, exFAT, and UDFS) while supporting standard disk and file operations. VHD API support allows management of the virtual disks. Virtual disks created with the VHD API can function as boot disks.
VHDX	<p>VHDX is a Hyper-V virtual hard disk (VHD) format found in Windows Server 2012.</p> <p>The main advantage of switching to VHDX is its increased storage capacity of 64 TB (terabytes), instead of VHD's standard storage limit of 2 TB. Other advantages of VHDX include file corruption protection and the ability to create differencing disks.</p>
VLAN	Virtual Local Area Network - A group of hosts with a common set of requirements that communicate as if they were attached to the same broadcast domain, regardless of their physical location.
VM	Virtual Machine - A completely isolated guest operating system installation within a normal host operating system. Modern virtual machines are implemented with either a software emulation or hardware virtualization or (in most cases) both together.
vMAS	Virtual Mitel Applications Suite (now called MiCollab) - Leverages VMware VSphere to enable businesses to consolidate Mitel's leading unified communications applications in the data center.
vMCA	Virtual Mitel Collaboration Advanced (now called MiCollab Audio, Web, and Video Conferencing)

Acronym or term	Definition
vMCD	Virtual Mitel Communications Director (now called MiVoice Business Virtual) - A virtual telephony services platform that provides IP-PBX features for small to large enterprises. It offers the same MiVoice Business software deployed on the Mitel 3300 IP Communications Platform (ICP) and industry standard servers.
VMDK	Virtual Machine Disk - File format for a type of virtual appliance developed from VMware products. It is proprietary to VMware, but the format is open and documented.
vMotion	VMware vMotion - A VMware feature allowing VMs to be moved automatically from host to host within a cluster. Several other VMware features make use of this capability (DRS and DPM, for example).
VMM	Virtual Machine Manager - See hypervisor.
vNIC	Virtual Network Interface Card
vOIG	Mitel Virtual Open Integration Gateway
vSAN	VMware Virtual SAN™ - Virtual SAN is a new software-defined storage solution that is fully integrated with vSphere. Virtual SAN aggregates locally attached disks in a vSphere cluster to create a storage solution that rapidly can be provisioned from VMware vCenter™ during virtual machine provisioning operations.
vSwitch	<p>Virtual Switch - A VMware vCNS virtual switch, an emulation of a layer 2 LAN switch resource available for connecting virtual machines to each other within the VMware infrastructure and/or to real physical network resources.</p> <p>The vSwitch is configured at the individual host level. It includes fewer features than the dvSwitch. See also dvSwitch.</p>

Acronym or term	Definition
vUCA Server	Mitel Virtual Unified Communicator Advanced Server (now called MiCollab Client Services) - Combines the call control capabilities of Mitel communications platforms with contact management, dynamic status, and collaboration applications to simplify and enhance real-time communications.
vUCC	Mitel Virtual Unified Communication and Collaboration virtual appliance (now called MiVoice Business Express)

11.2 Deployment Types

The table represents the maximum user count limits for different deployment types:

Deployment Type	Maximum User Count
Small	250
Mid-market	1500
Enterprise	2500
Large Enterprise	5000
Extra Large Enterprise	7500

This chapter contains the following sections:

- [VMware vSphere Storage Setup Guidelines](#)
- [Reference](#)

12.1 VMware vSphere Storage Setup Guidelines

CAUTION:

Recommendations in this document are intended to help build a better storage subsystem for hosting Mitel virtual appliances in VMware environment. All recommendations should be used at the discretion of your virtual infrastructure specialists. Always refer to official VMware documentations and recommendations for designing and implementing your VMware virtual infrastructure.

This chapter describes some ESXi storage options and recommends some storage configurations that can be optimized for running Mitel virtual appliances on ESXi hosts. Various options and considerations are discussed first, with recommendations to follow. Other practices and more advanced areas specific to VMware and real-time, latency sensitive workloads are also mentioned for considerations.

Note:

This chapter provides information about VMware and Hyper-V applications and requirements that were added for previous releases, and which may not be up-to-date. Always refer to your current VMware or Hyper-V documentation for the latest information.

12.1.1 Storage Virtualization

ESXi hosts provide storage virtualization (also called virtual disk), which logically abstracts the physical storage layer from virtual machines. ESXi virtual machines use virtual disks the same way a physical machine would use a local disk drive. However, a virtual disk can be mapped to one large physical file, a physical disk, or a set of files that are located either locally on the ESXi host or remotely.

A virtual machine uses a virtual SCSI controller to access virtual disks. The virtual SCSI controllers can be any of the following:

- BusLogic Parallel
- LSI Logic Parallel
- LSI Logic SAS
- VMware Paravirtual

The latest Windows operating systems use LSI Logic SAS by default, and most of the advanced Linux operating systems use LSI Logic Parallel by default. Whether the storage devices are being accessed through parallel SCSI, iSCSI, network, Fiber Channel, or FCoE, they are transparent to the guest operating system and applications running on the virtual machine.

12.1.2 Hard Disk Drive Types

Hard disk drives are used to construct virtual disks. There are two major types of hard disk drives: mechanical hard disk drives and solid state disk drives.

12.1.2.1 Mechanical Hard Disk Drive

Traditional hard disk drives use rotating mechanical disks as electromechanical devices to store data. Since it takes time to move a read/write head to the specific region on a rotating disc to access information, access time is slower when compared to a solid state disk drive. Hard disk drives have moving parts: the rotating disk and the moving read/write head. The overall average access latency is affected by the speed of spindle rotation and the read/write head access time. The data interface is one of SATA, SAS, or Fiber Channel.

12.1.2.2 Solid State Disk (SSD) Drive

Solid State disks use semiconductors as the storage medium rather than rotating disks. Since there are no moving parts, SSDs provide much faster access to data than conventional hard disk drives. Like mechanical hard disk drives, SSDs come with different interfaces: SATA, SAS, and Fiber Channel. One of the disadvantages of using SSD is its shorter service life. Storage cells inside a SSD do not survive as many write cycles as mechanical disks do. Operators must carefully monitor how intensively the SSD device is utilized to calculate its estimated lifetime.

12.1.3 Types of storage

There are two major types of physical storage.

- **Local Storage:** stores information on disks that are directly connected to the host, either internally or externally.
- **Networked Storage:** stores information on external disks or arrays that attached to the host through a high-speed network.

12.1.3.1 Networked storage

Networked Storage consists of external storage systems connected to ESXi hosts through a high-speed storage network. The networked storage can be shared among ESXi hosts; it can be accessed by multiple hosts concurrently, and is therefore eliminated as a single point of failure for local storage.

Networked Storage can be implemented by using the following network storage technologies:

- Fiber Channel (FC)
- iSCSI
- NFS
- Shared Serial Attached SCSI (SAS)

- Virtual SAN

12.1.3.2 Local Storage

Local storage is directly connected to ESXi hosts, either internally or externally. It can be accessed by only one host and cannot be shared across multiple hosts. Storage access is fast because it is accessible only by the attached host. It also creates a single point of failure because storage access is impossible when the host fails.

12.1.3.2.1 Fiber Channel (FC)

FC SAN is a very high performance type of storage device. The ESXi host must be equipped with Fiber Channel host bus adapters (HBA) and Fiber Channel switches to connect to the FC SAN. It uses Fiber Channel protocol and a Fiber Channel network to transport SCSI traffic.

A less expensive alternative is to use FCoE (Fiber Channel over Ethernet) adapters on an Ethernet network as the physical network to transport SCSI traffic using Fiber Channel protocol.

12.1.3.2.2 Virtual SAN

Virtual SAN is a distributed layer of software that runs natively as part of the ESXi Hypervisor. Virtual SAN aggregates local or direct-attached disks of a host cluster and creates a single storage pool that is shared across all hosts in the cluster. Virtual SAN offers some of the benefits of SAN without the need to purchase a real SAN.

Mitel virtual appliances support deployment on Virtual SAN. However, Virtual SAN has its own limitations and requires additional effort to setup and configure. Consult the VMware documentation to make sure that Virtual SAN will meet your needs before deploying it.

12.1.3.2.3 NFS

ESXi hosts can use Network File System (NFS) version 3 to communicate with NAS or NFS servers for accessing virtual disks. While SCSI is a block-level type of access, NFS is a file-level type of access. In theory, the performance is lower when compared to block-based access types.

12.1.3.2.4 iSCSI

iSCSI packages SCSI traffic into the TCP/IP protocol and it travels through a standard TCP/IP network instead of the specialized FC network. ESXi host serves as iSCSI initiator that communicates with iSCSI targets in remote iSCSI storage systems. ESXi host s two types of iSCSI connections:

- **Hardware iSCSI:** available through a third-party adapter capable of off-loading iSCSI and network processing.
- **Software iSCSI:** a software-based iSCSI initiator provided by VMware. It uses VMkernel to connect to the storage without any specialized third-party adapter.

12.1.3.2.5 Shared Serial Attached SCSI (SAS)

Shared SAS is a newer technology that depends on hardware vendors. It has the advantage of cost, performance, and simplicity over other storage technologies in localized environments.

12.1.4 RAID levels

RAID is a data storage virtualization technology that combines multiple disk drives into a single logical unit for the purposes of data redundancy and/or performance improvement. Data can be distributed across the drives in one of several ways, referred to as RAID levels. Each RAID level provides a different balance between reliability, availability, performance and capacity. Each RAID level inherits a specific write penalty. The write penalty is a measure of the extra waiting time for parity info or redundant info to be written onto the array of disks. A list of some popular RAID configurations is shown in the following sections.

12.1.4.1 RAID 0

RAID 0 provides no data redundancy or fault tolerance. It improves performance through the use of parallel read and write operations across multiple disk drives. RAID 0 has no error detection mechanism, so failure of one disk causes data loss on the array. Since there is no parity information to calculate, and only the data is written onto the disk, this type of RAID array has a write penalty of 1.

12.1.4.2 RAID 1

RAID 1 uses mirroring. Data is written identically to two or more drives, producing a mirrored set. Each read request is served by the disk with the lowest seek latency and rotational latency. Write performance is degraded, compared to RAID 0 because all drives being mirrored must be updated. However, RAID 1 provides redundancy and the disk array continues to operate as long as at least one drive is functioning. Because data is written to two disks (requires two write operations), RAID 1 has a write penalty of 2.

12.1.4.3 RAID 10

RAID 10 is a nested RAID level of RAID 1 and RAID 0. It creates a striped set from a series of mirrored drives. The array can sustain multiple drive losses, as long as no mirror loses all of its drives. Since RAID 0 carries a write penalty of 1 (no penalty), RAID 10 has the same write penalty as RAID 1, which is 2.

12.1.4.4 RAID 5

RAID 5 uses block-level striping with distributed parity. Parity information is distributed among the drives. It requires all drives but one to be present to operate. Upon a single drive failure subsequent reads can be calculated from the distributed parity. RAID 5 requires a minimum of three disks, and it is seriously affected by array rebuild time and the chance of failure during rebuild. RAID 5 has a write penalty of 4.

12.1.4.5 RAID 6

RAID 6 uses block-level striping with double distributed parity. Double parity provides fault tolerance for up to two failed drives. RAID 6 has a write penalty of 6.

12.1.5 Virtual disk modes

VMware Virtual Disks can be operated in a number of modes with different behaviors. After a virtual machine is created manually, without first startup, the mode of the virtual disk can be changed in Edit Settings on the virtual machine (in vSphere client).

12.1.5.1 Independent Persistent

In this mode, changes are written and remain on the disk (persistent), providing the best performance. However, snapshots cannot be taken on an independent persistent disk. Since most backup software needs to create a snapshot of a virtual disk to do backups, operators need to handle backups on independent persistent disks differently.

12.1.5.2 Independent Non-persistent

In this mode, changes are appended to a redo log. The redo log is erased when the virtual machine is powered off or reverts to a snapshot, causing changes to be discarded. Because the redo log tracks the changes in the file system of the virtual machine, performance is not as high as for disks in independent persistent mode.

12.1.5.3 Dependent

In this mode, changes are appended to a redo log that persists between power cycles. Similar to Independent non-persistent mode, disk performance of dependent disks are lower than for independent persistent mode. This is the Mitel default.

12.1.6 Virtual disk types

VMware virtual disks can be prepared in a few different ways.

12.1.6.1 Thick-Provisioned

Thick provisioned disks have all their space allocated at creation time. Because the space is pre-allocated it takes more time to create. There are two types of thick-provisioned disk.

12.1.6.1.1 Eager-zeroed

Eager-zeroed disks have all the space pre-allocated and zeroed out at the time of creation. It takes more time to create the disk, but the pre-zeroing action results in the best disk performance.

12.1.6.1.2 Lazy-zeroed

Lazy-zeroed disks have all the space pre-allocated at creation time, but blocks are zeroed only on first write. It reduces creation time but reduces performance the first time a block is written.

12.1.6.2 Thin-provisioned

Thin-provisioned disks have space allocated and zeroed upon the first write, as opposed to upon creation. For the first write on any unwritten file block this method carries a higher I/O cost. However, subsequent writes on the same blocks have the same performance as eager-zeroed thick disks.

12.1.7 Partition alignment

Earlier versions of ESX/ESXi create VMFS3 partitions with 64 KB alignment. The 64 KB alignment is retained even after the file system is upgraded to VMFS5. The vSphere Client automatically aligns VMFS3 and VMFS5 partitions along the 1 MB boundary at creation. To obtain the better 1 MB alignment, a 64 KB aligned partition must be deleted and recreated using vSphere Client and an ESXi 5.0 or later host.

12.1.8 SAN multipathing

SAN Multipathing is a fault-tolerance and performance enhancement technique. There are multiple physical and logical paths between the host and the storage device through different controllers and switches.

Paths are accessed based on the configured path policies. SAN path policies can have a significant effect on storage performance. For most Active/Passive storage arrays, ESXi uses the Most Recently Used (MRU) path policy by default. Fixed path policy is not recommended for Active/Passive arrays because it can cause frequent and rapid path switching resulting in slow LUN access.

For most Active/Active storage arrays, ESXi defaults to use of the Fixed path policy. For optimal performance, the Round Robin path policy can be considered.

Round Robin path policy can improve storage performance in some environments by cycling I/O requests through all active paths. However, not all storage arrays support the Round Robin path policy. Check the storage array documentation or the storage vendor to make sure the Round Robin path policy is supported and recommended for your storage array and configuration.

12.1.9 File systems

VMware ESXi host s two different file systems for virtual machines: VMFS and NFS.

12.1.9.1 VMFS

VMFS is the VMware native high-performance clustered file system, supported either locally or network-attached.

12.1.9.2 NFS

The NFS distributed file system has been in use for nearly 20 years. NFS is strictly network-attached and uses Remote Procedure Calls to access remote files.

12.1.9.3 Raw Disk

VMware vSphere supports storing virtual machine guest files on a raw disk. Raw device mapping (RDM) provides a mechanism for a virtual machine to have direct access to a LUN on the physical storage subsystem (iSCSI or Fiber Channel only). RDM is a mapping file in a separate VMFS volume that acts as a raw physical storage device. It has some benefits, but it also has limitations, and is not suitable in every situation. Consult the VMware documentation to thoroughly understand Raw Device Mapping before using it.

12.1.10 Recommendations

Storage network

Either iSCSI or NFS would be suitable. iSCSI is a block-based protocol operating at a lower level, while NFS is a file-based protocol operating at a higher level. However, iSCSI has the advantage of being slightly more efficient. The following implementation is recommended for both iSCSI and NFS.

- For better iSCSI performance, enable jumbo frames when possible. Jumbo frames reduce packet-processing overhead, thus improving CPU efficiency of storage I/O. Use of jumbo frames requires end-to-end hardware support from iSCSI adapters, storage arrays, and hardware network switches. The default Ethernet frame size is 1500.
- It is recommended that both iSCSI and NFS have their own dedicated physical network to minimize network interference from other packet sources. If this is not possible, logical separation using VLAN can be used.

12.1.10.1 RAID level

RAID 10 provides the most balanced performance, redundancy, cost, and write penalty, so it is the recommended RAID level. Other RAID levels may still be used to build virtual disks to better fit specific workload requirements.

12.1.10.2 Virtual Disk Mode

Mitel supports the default Dependent virtual disk mode. Dependent disk mode allows temporary disk snapshot, which Mitel uses during backup and upgrade. Virtual disks of Mitel virtual appliances must not contain snapshots during normal operation.

12.1.10.3 Virtual Disk Types

Mitel applications are real-time applications and require high system performance during operation. It is preferred that all disk spaces be allocated during system installation so that they are ready to roll during operation. Therefore, the thin-provisioned disk type is not recommended. Thick Provisioned with lazy-zeroed virtual disk type (the default) is recommended because it has a more balanced disk performance and OVA deployment time.

12.1.10.4 Partition Alignment

Use the latest version of VMware software to perform disk partitioning and formatting, for optimum disk drives preparation.

12.1.10.5 SAN Multipathing

Enabling SAN Multipathing is recommended for fault-tolerance and performance enhancements. Although Multipathing usually requires additional networking equipment such as HBAs and network switches, the benefits outweigh the costs.

12.1.10.6 File System

In general, VMware VMFS is recommended for use as the underlying file system because VFMS is native to VMware and it is designed specifically for the VMware environment.

12.1.11 Other Considerations

CPU Utilization

Since a higher level of CPU resources is needed to process higher levels of networking traffic, it is important to ensure sufficient CPU resources are available, without limiting the maximum network throughput. It is important to monitor CPU utilization of high-throughput workloads to ensure that the CPU resource is sufficient.

12.1.11.1 Separate Physical NIC

If a physical NIC is shared by multiple virtual machines, each virtual machine could affect the performance of other virtual machines. For best network performance, use separate physical NICs for heavy networking I/O virtual machines and latency-sensitive virtual machines. 10 Gbps NICs have multi-queue support, so this recommendation typically does not apply, as long as a sufficient number of queues is available.

12.1.11.2 Jumbo Frames

ESXi supports the use of Jumbo Frames with iSCSI. ESXi allows Jumbo Frames with MTU size up to 9000 bytes. When a Jumbo Frame is used for iSCSI traffic, the following must be true:

- The network must support Jumbo Frames of that MTU size end-to-end to be effective.
- Network equipment and storage vendors must be consulted to ensure that the physical NICs, iSCSI HBAs, and network switches support Jumbo Frames from end to end.
- Network switches have been set up and verified for Jumbo Frames, to ensure proper operation on the network.

12.1.11.3 vSphere Flash Read Cache

Flash Read Cache accelerates virtual machine performance through the use of host-resident flash-based storage devices as a cache. VMware Virtual Flash configured as a cache can be used for virtual machine

migration, read caching, and write caching. VMware Distributed Resource Scheduling (DRS), and vSphere High Availability support Virtual Flash.

12.1.11.4 Storage Hardware Acceleration

The functionality of hardware acceleration enables an ESXi host to integrate compliant storage arrays and off-load specific virtual machine and storage management operations to the storage hardware. Off-loading tasks to storage hardware means that the storage consumes less CPU, memory and storage fabric bandwidth on the ESXi hosts. Storage hardware acceleration works only when appropriate hosts and storage arrays are combined. Verify with the hardware vendors to make sure that storage hardware acceleration is supported.

12.1.11.5 VMware-specific optimizations

VMware-specific optimizations

When using ESXi with a SAN, use the following tips to avoid SAN problems.

- Put only one VMFS data store on each LUN. Multiple VMFS data store on one single LUN is a configuration that is not recommended.
- Do not change the path policy the system sets unless the implications of the change are fully understood.
- Ensure there is no single point of failure from the SAN to the host HBAs.
- I/O device latency should not be consistently greater than 20-30 ms (for more information, refer to this VMware vSphere Blog post: <http://blogs.vmware.com/vsphere/2012/06/troubleshooting-storage-performance-in-vsphere-part-2.html>).
- Tune the virtual network adapter (VMXNET3). Some of the default parameters may not suit your environment (retries and time out values, for example). In that case, fine tuning of parameters may be required.
- If the virtual machine is using an LSI Logic vHBA, the reqCallThreshold value could be adjusted.

The lower the reqCallThreshold value, the lower the time the I/O requests are likely to stay in the vHBA queue. If reqCallThreshold is set to 1, it means the I/O will be dispatched to the lower layer even there is only one I/O request in the LSI Logic vHBA queue. The default reqCallThreshold value is 8, and it is a system-wide configuration on a ESXi host.

12.1.11.6 vSphere pluggable storage architecture

vSphere pluggable storage architecture

VMware vSphere employs the vSphere Pluggable Storage Architecture (PSA) in its storage subsystem. PSA is essentially a collection of plugins inside the VMkernel layer of the ESXi host. The top-level plugin in the PSA is the Multipathing Plugin (MPP). MPP defines how vSphere manages and accesses storage, including load balancing, path selection, and fail over. VMware provides the Native Multipathing Plugin (NMP) as its MPP, but storage vendors may also provide their own MPP implementations.

Within each MPP, including VMware NMP are two sub-level plugins: the Storage Array Type Plugin (SATP) and the Path Selection Plugin (PSP). SATP handles path fail over and PSP handles load balancing.

Depending on the type of storage arrays used, the default PSP policy could be:

- Fixed: The default policy for most active-active arrays.
- Most Recently Used: The default policy for most active-passive arrays.
- Round Robin: The default policy for some active-active and active-passive arrays.

For best performance, check with your storage vendor to obtain the appropriate plugins, if available.

12.2 Reference

Reference

VMware Publication: vSphere Storage (for ESXi 5.5 and vCenter Server 5.5) <http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-storage-guide.pdf>

This chapter contains the following sections:

- [VMWARE vCloud Director](#)
- [Creating and Publishing Virtual Application Catalog Templates](#)
- [Deploying and Configuring Virtual Applications from Catalogs](#)
- [Run-time Resource Behavior of Deployed Virtual Applications](#)
- [vCloud Director Lease and Lease Expiry](#)
- [Un-deploying Virtual Applications](#)
- [Networking and Security Considerations](#)

13.1 VMWARE vCloud Director

Note:

This chapter provides information about VMware applications and requirements that were valid for previous releases, and which may not be up-to-date. Always refer to your current VMware documentation for the latest information.

This chapter describes the general behavior of all Mitel virtual applications when deployed in a VMware vCloud Director environment (vCD). This includes deployment methods, runtime behavior, and networking-related considerations.

vCloud Director requires vSphere (ESXi + vCenter Server), and vCloud Networking and Security (vCNS) as an underpinning, at a compatible version level.

Note:

vCD 5.1 and above only, are compatible for use with Mitel virtual appliances.

13.1.1 vCloud Director Overview

This section provides a brief overview of the VMware vCloud Director (vCD) product. For a more complete understanding of vCloud Director architecture, operation, and use, refer to the VMware documentation. The figure below shows the VMware vCloud Director system architecture at a very high level (from VMware documentation).



Figure 3: vCloud Director high level architecture

vCloud Director operates as an administration layer on top of VMware vSphere, using vSphere resources, and providing an abstracted view of vSphere resources to users. vCloud Director uses VMware vCNS to create secured organization networks and other networking constructs; this is largely invisible to the user.

The following table outlines key concepts for understanding how vCloud Director operates.

Table 29: vCloud Director Key Concepts

Concept	Definition
<p>Organization</p>	<p>The unit of multi-tenancy that represents a single logical security boundary, assigned sets of resources, and policies.</p> <p>An Organization contains users, one or more virtual data centers (vDCs), and one or more networks. Organizations may represent a customer, in the hosted service provider/public cloud case, or organizations within a corporation (private cloud case).</p> <p>Organization administrators and users control their resources through an Organization-specific web administration portal provided by vCloud Director (vCD), and are completely segregated from other Organizations and the vCD System layer.</p>

Concept	Definition
Virtual Data Center (vDC)	A grouping of computational and storage resources from a single vCenter Server. vDCs at the vCloud Director level map to resource pools at the vSphere level.
Provider vDC	<p>A provider vDC represents an overall set of resources from which sub-units (such as Organization vDCs) are allocated and managed by vCloud Director system administration.</p> <p>A vCloud Director installation may make use of more than one provider vDC.</p>
Organization vDC	<p>A sub-grouping of computational and storage resources allocated from a Provider vDC and assigned to a single Organization, under control of the Organization. An Organization may have more than one Organization vDC.</p> <p>An Organization vDC is a deployment environment in which VAs and vApps can be instantiated, deployed, and powered on. An Organization vDC allocates resources to VAs and vApps using one of the following allocation models:</p> <ul style="list-style-type: none"> • Pay-As-You-Go • Reservation Pool • Allocation Pool
Catalog	<p>A repository of vApp Templates and media available to users for deployment.</p> <p>Catalogs can be private to an Organization or published to all Organizations in the same vCloud Director environment.</p>

Concept	Definition
vApp	<p>A container for a software solution in the vCloud Director, and the standard unit of deployment for workloads in vCloud Director. A vApp contains one or more virtual machines, has power-on operations, and can be imported or exported as an OVF.</p> <p>A virtual appliance (VA) is effectively a single-VM vApp.</p>
Lease	<p>Both resources representing Catalog Templates and deployed vApps are subject to Leasing from the vCloud Director provider. A lease sets limits on the amount of time the resource may be used before the lease must be renewed.</p>

13.1.2 How vCloud Director capabilities are used in Mitel products

Mitel virtual appliances use vCloud Director in the following operations:

- Publishing Mitel virtual appliances or virtual appliances as Organization Catalog items
- Deploying Mitel virtual appliances from Catalog, within an Organization's vCD
- Ensuring Mitel virtual appliances receive adequate resource allocations for their run-time needs
- Understanding the implications of vCloud Director lease expiry, and planning appropriate notification of upcoming expiry to customers
- Un-deploying Mitel virtual appliances

Only information specific to Mitel virtual appliances is described in this guide; for full information and instructions, refer to the VMware documentation. For specifics related to using the vCloud Director Web Admin interface, refer to the VMware vCloud Director User Guide.

13.2 Creating and Publishing Virtual Application Catalog Templates

vCloud Director 5.5 supports direct deployment of OVA packages into a Catalog or a Organization. Mitel virtual applications are available in OVA format. Deploying the OVA directly is the preferred method if your vCloud Director environment supports it.

This section describes how to create and publish Mitel virtual appliances as vCloud Director Catalog Template items from OVA packages. vCloud Director releases before 5.5 do not directly support deployment of VAs or vApps packaged in OVA format, which is Mitel standard packaging and distribution format.

If you are using vCloud Director 5.1 or earlier, which support uploading of VA and vApps only when packaged in OVF format, (the unpacked file hierarchy making up the VA or vApp definition), then you must use the following procedures. The OVA is essentially a TAR-packed, one-file packaging of the OVF directory tree, so Mitel OVAs must first be unpacked to OVF, or otherwise prepared to be uploaded to vCloud Director Catalogs. Multiple methods are available.

Note:

Catalog items can also be created as ISO images in vCloud Director. However, Mitel VAs and vApps do not need this capability, since they are packaged as OVAs. While it is possible, deploying from an ISO is not recommended practice for OVA-packaged applications, and is not covered in this guide.

The following subsections provide descriptions of three different procedures to create Catalog items from OVA distribution files. The procedures described here are:

- [Unpack and publish OVA in Linux](#)
- [Unpack and publish OVA in Microsoft® Windows®](#)
- [Import VM from vSphere](#)

13.2.1 Unpack and Publish OVA in Linux

If you are using Linux desktop-based administration, the Linux desktop may be the best way to unpack and publish the OVA. Both command line and GUI options are described. The GUI method described assumes CentOS/Red Hat.

The Linux method is preferred to the others discussed in this section because it may be somewhat easier to perform than the Windows desktop method.

To unpack and publish from OVA format file in a Linux environment, use the following steps.

Table 30: Linux method

Command line procedure	GUI procedure
Unpack the OVA	
1. Copy or download the application OVA file to a temporary location in Linux, or mount it as a share drive. # cp app.ova /Temp/	1. Drag and drop the application OVA into a temporary folder.
2. Navigate to the temporary directory. # cd /Temp	2. Right-click the OVA file, and select Open with Archive Manager .

Command line procedure	GUI procedure
<p>3. Create a destination directory for the unpacked OVF.</p> <p># mkdir app-ovf</p>	<p>3. In the Archive Manager:</p> <ul style="list-style-type: none"> • Click Extract. • In the folder drop down menu, select Extract > Other > New Folder. • Enter a name for the destination folder, app-ovf, for example. • Click Extract.
<p>4. Unpack the OVA using the TAR command.</p> <p># tar -xf app.ova --directory app-ovf/</p> <p>The files are unpacked into the specified destination directory.</p>	<p>4. Close the Archive Manager.</p>
<p>Upload the OVF to the Catalog</p>	
<ol style="list-style-type: none"> 1. Open the vCloud Director Web Admin for the Organization as Organization Administrator or Catalog Author. 2. Open the target Catalog. 3. Click the Upload icon. 4. Browse to the OVF folder you created in the previous step, and select app.ovf., or the name of the destination folder you created. 5. Click Upload. 6. Certificate exceptions may pop up during progress. Click Accept on each one. 7. Confirm progress and final success. 	

13.2.2 Unpack and publish OVA in Microsoft® Windows®

If you are using mostly Windows desktop-based administration, the Windows desktop approach may be the best way to unpack and publish the OVA.

This approach involves extra steps, compared to the Linux method, and the use of additional tools, so the Linux method is preferable, if it is available.

Unpack and publish in Windows:

1. Install and configure a TAR file unpacking utility.
 - Configure the unpacking tool so that it does not automatically adjust text files to DOS/Windows line breaks (keep line breaks in UNIX format). Recommended tool: 7-Zip (www.7-Zip.org)
 - If the unpacking tool you have chosen cannot be configured to so that it does not automatically adjust text file format, then download and configure a text editor that is capable of converting the files back to UNIX format. Recommended tool: IDM Computer Solutions UltraEdit
2. Create a destination directory for the unpacked OVF; for example, C:\Temp\App\App-OVF.
3. Unpack the OVA using the TAR utility.
 - a. Right-click the OVA file, select **Open with**, and select the **TAR utility** (or select the TAR utility directly if that is possible).
 - b. Click **Extract All** (or the equivalent command for the utility), select the destination folder you just created, and click **OK**.
 - c. Click **Extract**.
4. Close the unpacking utility.
5. If the unpacking utility does not support leaving the file format as-is (instead, automatically converting text files to DOS format), then open each text file (app.mf, app.ovf, and so on) with the text editor described above. For each text file:
 - a. Right-click the text file, select **Open with**, and select the editor utility (or select the utility directly if available).
 - b. For each file, use **Save As** to convert it to **UNIX** format.

Upload the OVF to the Catalog:

1. Open the vCloud Director Web Admin for the Organization (as Organization Administrator or Catalog Author).
2. Open the target **Catalog**, and click **Upload**.
 1. Click the **Upload** icon, browse to the OVF folder created previously, select the **OVF file**, and click **Upload**,
 2. Certificate exceptions may pop up during progress. **Accept** them.
 3. Confirm progress and final success.

As an alternative to the above procedure, VMware also provides a tool called OVF Tool, which has a command line interface. Since this tool is more difficult to install and use, is considerably slower, and can alter file names as provided in the original OVA packaging, it is not generally recommended. However, this tool may be useful where scripting is used to automate OVF conversion and importing.

13.2.3 Import VM from vSphere

This method requires that you have both vCloud Director system administrator level access and vCenter administrator access. This cannot be used by normal Organization level administrators Organization users.

This method may be most appropriate if you wish to set up and publish master Catalogs for your customer Organizations to use.

To accomplish this, you can create a separate special Organization specifically for the purpose of importing and publishing the OVA-based VAs and vApps with a Public Catalog. Users of this special Organization are given system-level administrator access to vCloud Director and to vCenter, allowing them to do the procedures required to populate the Catalog. The special Organization's Public Catalog can then be accessed by normal customer Organizations for the purpose of deploying the VAs and vApps.

To import from vSphere, use the following procedure.

Deploy OVA as a virtual appliance in the vSphere level:

1. Log in to vSphere vCenter as Administrator, or another administrative user with sufficient privileges for deploying virtual appliances.

Note:

You must be logged in as vCenter Administrator. Normally, vCloud Director Organization-level administrators do not have such access.

2. Create a new, separate resource pool and/or a folder at the vSphere level to manage the new VMs you will deploy. This step is optional, but recommended to make organizing your VMs easier.
3. Deploy the application OVA as a VM in vSphere vCenter using normal VMware procedures. Leave the VM powered off in vSphere.

Note:

You must select valid network connections for each vNIC on the VM in vSphere. If you do not, the Import operation (described in the next procedure) will fail in vCloud Director. The connections you enter do not have to be valid for vCloud Director use, but they must be valid in vSphere.

Import a virtual appliance to a vCloud Director Catalog

1. Log in to vCloud Director as a vCloud Director System Administrator (or equivalent).

Note:

You must be logged in as vCloud Director System Administrator. Normally, vCloud Director Organization-level administrators do not have such access.

2. Navigate to the target Organization and the specific public or private Catalog to import the VA or vApp into. Select the Catalog. The catalog opens in List view.
3. Select the **vApp Templates** tab.
4. In the **vApp Templates List** view, click Import from vSphere.
 - OR: Click the Actions drop-down menu and select Import from vSphere.
 - OR: Right-click in the Template list and select Import from vSphere. The Import vApp Template wizard opens.

5. In **Catalog**, enter the catalog item name.
6. In **Description**, enter a description of this catalog item. Adding a description is optional, but recommended.
7. If there is more than one vDC configured for this Organization, select the **Organizational vDC** to import to.
8. Select **Move VM**, to move the VM from its current vSphere location to the vCloud Director-generated resource pool corresponding to this Catalog.

Note: It is recommended that you not use Copy VM. This is the default setting; Copy VM copies the VM to vCloud Director space, leaving the original in place in vSphere. This can waste space and be more difficult to manage.

9. Optional: Select **Gold Master**. See [General settings for Catalog Template items](#) for details.
10. Click **OK** to exit the Import dialog box, and start the import operation.
11. In the **Catalog List** view, watch the progress of the import operation and confirm that it finishes successfully. This process can be quite long, depending on the size of the VM (or OVA), network performance, and the total load on the hosts involved.

13.2.4 General Settings for Catalog Template Items

After Catalog Template items are uploaded or imported, and after they appear in the vApp Templates List view, you can modify the Template settings.

To open a template, select the template. Right-click and select Properties, or select Properties from the Actions drop-down menu.

You can modify the properties as follows:

- Change the name or description of the Template.
- Set Gold Master to Yes or No (the default is No). Changing the Gold Master designation to Yes has no effect on the Template item, either when importing or when deployed, other than to highlight the Template. This could be useful, for example, to mark the Template items as ready for general deployment, as opposed to “under test” or “under construction”.
- Reset the Template lease. The Template Lease setting determines how long the Template remains in the Catalog. This is limited to the maximum Catalog item lease time set for the vDC in which the Template is deployed. See [vCloud Director Lease and Lease Expiry](#) for more information.
- Set Create vApp from Template Options. This is not applicable to Mitel VAs, and other single VM virtual appliances.
- Set VM startup order. This is not applicable to Mitel VAs, and other single VM virtual appliances.

13.3 Deploying and Configuring Virtual Applications from Catalogs

To deploy a Mitel VA or vApp from the vCloud Director Catalog, follow normal vCloud Director Web Admin procedures. Two different methods are described here.

i Note: If you are installing multiple Mitel vApps, it is strongly recommended that you deploy them one at a time. This will avoid having the guest properties for all of the vApps displayed at the same time, which can cause confusion in completing the deployment.

Method 1 - Deploying from Catalog View

1. Log in to the **Organization** as an **Administrator**, or another user authorized to deploy VMs.
2. Navigate to the **Catalogs** tab and open the Catalog containing the vApp Template to deploy. This can be an Organization private Catalog, or one published externally by a different Organization.
3. On the **vApp Templates** tabs, select a **vApp template**, right-click, and select **Add to My Cloud**.
4. Follow the steps in the Add to my Cloud wizard that is launched:

- a. When the Mitel EULA appears, click **Accept**.
- b. Enter a name and description for the vApp. The description is optional, but recommended.

i Note: You cannot modify the runtime and storage lease duration for the VA at this stage, but you can modify and renew the lease at a later time. See [vCloud Director Lease and Lease Expiry](#).

- c. Configure the network mappings: For each network interface on the VM, select the network to connect the interface to, and the IP Allocation type, For Mitel VAs, select Static - Manual IP Allocation. Static - Manual is the only IP Allocation type currently available for Mitel VAs. For VAs with more than one vNIC, the network and IP Assignment must be done for each vNIC separately.

i Note: For IP Assignment type, three options are available. Each of these enable address and related configuration to be pushed into the VA when it is deployed. Which one to use depends on desired outcome, and compatibility of the specific VA being deployed. However, current release Mitel VAs do not support either IP Pool or DHCP-based auto configuration of addresses. Always select Static – Manual allocation.

- d. Configure Properties. Not all Mitel product releases support the modification of custom Properties; these Mitel VAs will display the message, “There are no user configurable properties.”When the custom properties screen appears, fill in all fields, For more information about the Custom properties for each Mitel VA, see [Deploying Mitel OVAs](#).

i Note: Configure Networks and Custom Properties settings are not available for use with current MSL-based VAs.

- e. The Ready to Complete screen appears. Check the settings displayed to ensure that they are all correct, and click Finish.
- f. Navigate to My Cloud > vApps view to watch progress and ensure success.

Method 2 - Deploying from My Cloud view

1. Log in to the **Organization** as an **Administrator**, or another user authorized to deploy VMs.
2. Navigate to the **My Cloud** tab or the **vApps list** view.
3. Right-click in the **vApps list** and select **Add vApp From Catalog**, or click the **Actions** drop-down menu and select **Add vApp From Catalog**. The Add vApp from Catalog wizard opens.

4. Select the appropriate catalog (either Organization private or Public), then select the **Template to deploy**.
5. Click **Next**.
6. For the rest of the procedure, see [Method 1 - Deploying from Catalog View](#).

13.3.1 Deployment Differences Relative to vSphere

One important difference between vCloud Director-based deployment and deployment directly on vSphere is that the Configuration Options dialog is not presented in vCloud Director, as it is in vSphere. Therefore the administrator cannot choose the specific deployment profile to use; Small Business or Enterprise, for example. In vCloud Director, the vApp is deployed in its default configuration.

While this issue is created when the Virtual Appliance is uploaded or imported to a vApp Template in the Catalog, it manifests only when the Template is deployed.

Unless the vApp is appropriately modified to correspond to the desired deployment profile, its resourcing may be incorrect; it may be too small for the desired capacity, or unnecessarily large.

Refer to section [Modifying vApp Resourcing for Deployment Configuration](#) for details on how to work around this issue.

13.3.2 Post-deployment Configuration

For current Mitel VAs other than MiVoice Business Express, after the Mitel VA has been created, the VA is then configured as for deployment of the VA under vSphere vCenter. In the case of MiVoice Business Express, initial configuration is performed using the Mitel Initial Configuration Wizard at deployment time (see [Deploying and configuring virtual applications from catalogs](#)).

This procedure varies, depending on Mitel VA specifics; the basic steps are as follows:

1. Right-click the VA name, and select Start.
2. Double-click the Console icon to open the console on the VA.
3. Configure the administrator password, the host name, the IP address, IP mask, and DNS.
4. Select the network interfaces, and so on, as required.

Refer to application-specific administration and engineering guides, available on Mitel Document Center, for detailed procedures.

The VA is now ready to use.

Note:

With the exception of MiVoice Business Express, Mitel VAs are not pre-configured with IP address as set in the address auto-configuration step above. The MiVoice Business Express Initial Configuration Wizard guides you through the initial system setup at the time you first login to the MiVoice Business Express web administration. Refer to MiVoice Business Express documentation.

13.4 Run-time Resource Behavior of Deployed Virtual Applications

Virtual appliances and vApps deployed through vCloud Director are run at the vSphere/vCenter level on the same ESXi, so most application behavior and performance characteristics are exactly the same as if deployed directly through vSphere/vCenter.

There are, however, large potential differences in the way resources are allocated to the run-time VAs in vCloud Director environments. Many Mitel VAs have CPU and memory reservations built into their OVA packaging. vCloud Director makes modifications to internal resource reservations of the underlying OVA/OVF when the VA is deployed. vCloud Director does not modify other aspects, such as disk size and number of vCPUs, for example.

In the vCloud Director environment, the allocation model assigned to the Organization vDC into which virtual appliances are being deployed affects how resource reservations built into the virtual appliance OVA are applied. The following table summarizes these differences.

Table 31: Resource Settings vs. Allocation Model

Allocation Model	VA run-time Setting					
	CPU			Memory		
	Reservation	Limit	Shares	Reservation	Limit	Shares
Pay-As-You-Go	VA # vCPU x vCPU speed x guarantee %	VA # vCPU x vCPU speed	Normal (forced)	VA memory x guarantee %	VA Memory (forced)	Normal (forced)
Allocation Pool	0 (forced)	VA # vCPU	Normal (forced)	VA memory reservation x guarantee %	VA Memory	Normal (forced)

Allocation Model	VA run-time Setting					
	CPU			Memory		
	Reservation	Limit	Shares	Reservation	Limit	Shares
Reservation Pool	VA CPU reservation	VA CPU limit	VA CPU shares	VA memory reservation	VA memory limit	VA memory shares

i Note:

1. **VA # vCPU** - number of vCPUs assigned to the VA in OVA/OVF
2. **VA memory** - amount of memory (memory size) assigned to VA in OVA/OVF
3. **VA CPU reservation** - CPU resource reservation (GHz) applied to VA in OVA/OVF
4. **VA memory reservation** - memory resource reservation (GB) applied to VA in OVA/OVF
5. **VA limit** - resource limit (CPU or memory) applied to the VA in OVA/OVF
6. **VA shares** - resource shares (CPU or memory) applied to VA in OVA/OVF
7. **vCPU speed** - vCPU speed value assigned to Pay-As-You-Go, as set in Organization vDC allocation policies (see following subsections)
8. **Guarantee %** - CPU or Memory resource guarantee value, as set in the Organization vDC allocation

i Note:

In early testing, Mitel found discrepancies in the way Allocation Pool reservations are set, relative to the formulas in VMware documentation. It is important to check actual applied values in the customer vCloud Director environment by inspecting actual reservations of the deployed VMs at vSphere level while the VM is running.

In vCloud Director, each allocation model allows different controls for how resource over-commit is managed. For example, as can be seen in table above, both Pay-As-You-Go and Allocation Pool allocation models provide a “guarantee percent” control, that sets the reservation at run time, based on adjustment of the virtual appliances built-in reservations. Even these two models apply the allocation percent control somewhat differently, however. Only the Reservation Pool allocation model fully preserves the virtual appliances reservations.

Further details and discussion about each allocation model are provided in the following sub-sections.

13.4.1 Pay-As-You-Go Allocation Model

The Pay-As-You-Go allocation model is not recommended for real-time voice applications that provide their own strict built-in reservations, since it cannot provide the required CPU reservation guarantees.

Pay-As-You-Go allocation, with suitable tuning, could be a good choice in some circumstances, such as large hosting providers, or try-and-buy agreements.

The Pay-As-You-Go allocation model is designed to:

- Allow customer organizations to control the amount of resources being used (and therefore, what they will be billed by their provider).
- Allow the provider (the vCloud Director system administrator) to manage Service Level Agreement (SLA) assurances by managing their overall resource over-commit levels.
- Allocate resources and calculate reservations at the VA level.

As shown in [Table 27](#), this model effectively ignores any CPU GHz reservation built into the OVA file for the VA, instead basing reserved resources only on the number of vCPUs and the vCPU speed, as modified by the percent guarantee. Similarly, memory reservation is based on the VA's built-in memory size, as modified by the percent guarantee, and ignoring VA built-in reservations. Both CPU and memory shares are always forced to Normal, independent of any built-in VA setting.

Recommendations:

- The Pay-As-You-Go allocation model is not recommended for use with real-time-sensitive UCC applications (MiVoice Business Virtual, MBG Virtual, MiCollab Virtual), and is not recommended for deployments in which voice quality guarantees are part of the SLA.
- If you are using the Pay-As-You-Go allocation model, customers with a strict SLA can be provided with a different Organization vDC, specifically for their real time applications, using a different allocation model (Reservation Pool, for example). This is the preferred approach.
- The Pay-As-You-Go allocation model can be a viable choice in larger hosting providers, where there are plenty of resources, overall, and there are large numbers of Organizations using the same allocation approach.
- As guaranteed percent values approach 100%, this model becomes increasingly close to the Reservation Pool model, with the exception that it is based on set vCPU speed in the Organization vDC (rather than the VA's CPU reservations), and on the VA's set memory size (rather than memory reservations).
- If voice quality is part of the SLA provided to the Organization, then allocation guaranteed percent values should be set relatively high for the Organization vDC, especially when the CPU reservation is to be used for real time voice.
- If voice quality is not an important part of the SLA provided to the Organization, then allocation guaranteed percent values can be set somewhat lower for the CPU. However, memory percent guarantees should still be set fairly high, to avoid slow response time.

For detailed information on configuring the Pay-As-You-Go allocation model, refer to VMware documentation.

13.4.2 Allocation Pool Allocation Model

The Allocation Pool allocation model is not available for real-time voice applications that provide their own strict built-in reservations, since it cannot provide the required CPU or memory reservation guarantees.

Allocation Pool allocation, with adequate and guaranteed resourcing at the Organization vDC level, can be a good choice in some circumstances, such as small systems running alongside other low-demand workloads.

The Allocation Pool allocation model is designed to:

- Allow customer organizations limited control of the amount of resources being used (along with what they will be billed by their provider).
- Allow the provider (the vCloud Director system administrator) to maximize resource over-commit while maintaining SLAs at the Organization vDC level.
- Allocate resources and calculate reservations at the Organization vDC level, with individual VAs drawing on resources from the overall reservation of the Organization vDC.

In the Allocation Pool model, resources are committed to the Organization vDC (drawn from the overall Provider vDC) when the Organization vDC is created. In effect, this allows the customer Organization to control their own over-commit levels. If the CPU guaranteed percent value remains set at 0 (the default) larger numbers of virtual appliances and vApps may be run simultaneously, with the effect that they all slow down, but they are still able to run.

VAs with built-in reservations are then modified when deployed, as shown in [Table 27](#), and draw their resources from the overall Organization pool. This effectively reduces the Mitel built in reservations, which may affect voice quality.

Recommendations:

- The Allocation Pool allocation model is not recommended for use with real-time-sensitive UCC applications (MiVoice Business Virtual, MBG Virtual, MiCollab Virtual, etc), and is not recommended for deployments where voice quality guarantees are part of the SLA.
- Alternatively, the customer with the strict SLA could be provided with a different Organization vDC, specifically for their real time applications, using a different allocation model (Reservation Pool, for example). This is the preferred approach.
- The Allocation Pool allocation model can be a good choice, under some circumstances, especially when the customer Organization can afford lots of overall resourcing for their Organization vDC.
- In this model, it would primarily be up to the customer Organization to ensure they provide enough resources to ensure voice quality. This, in effect, provides customer-controlled SLA.
- If it is known that the customer will be deploying voice or other real-time sensitive virtual appliances or vApps, the hosting provider should set CPU guaranteed percent values relatively high. This will ensure that the Organization vDC overall gets adequate CPU and memory resources for their needs, including the real-time applications, even when there is high overall load across the host's Provider vDC.

For detailed information about configuring the Allocation Pool allocation model, refer to VMware documentation.

13.4.3 Reservation Pool Allocation Model

The Reservation Pool allocation model is the only allocation model that is fully recommended for real-time voice applications that provide their own strict built-in reservations (like Mitel voice applications).

Where the Organization is using a different model for general applications, you can configure a separate Organization vDC using the Reservation Pool allocation model to run only the real-time applications.

The Reservation Pool allocation model is designed to:

- Allow customer organizations full control of the resources being used within their Organization vDC, including specific reservations where required at the VA level, with strong guarantees on the total resources available to the Organization vDC as a whole.
- Allow the provider (the vCloud Director system administrator) to provide very tight SLA guarantees, with the cost being loss of the ability to configure overall resource over-commit. This cost will usually be charged back to the customer.
- Allow resources to be allocated and guaranteed at the Organization vDC level, as well as at the individual VA level, with individual VAs drawing on resources from the Organization vDC's overall reservation as they are run.
- Reservation Pool is the only allocation model that allows Administrators (at either Organization or System level) to access and modify resource reservations on VMs; the other allocation models do not allow this.

In the Reservation Pool model, resources are fully committed to the Organization vDC (drawn from the overall Provider vDC) when the Organization vDC is created. Built in OVA reservations are not modified, as shown in [Table 27](#). This allows the customer Organization to control their own over-commit levels, within their vDC. If resources are over-committed (that is, total reservations becomes higher than the total reservation for the Organization vDC), then VAs will fail to deploy.

Recommendations:

- The Reservation Pool allocation model is recommended for Mitel VAs, which have built in resource reservations.
- The Reservation Pool allocation model is ideal for real-time applications, where both CPU and memory reservations are important to guarantee voice quality SLAs. This comes at expense, however, of zero over-commit available to the provider.
- In this model, it may be primarily up to the customer Organization to ensure that they provide enough resources to ensure voice quality. In effect, SLA is customer-controlled, using reservations built into the VAs as they are deployed.
- Alternatively, if the hosting provider is also supplying the VAs the customer Organization is deploying, (they could be selected from a Catalog provided by the host Organization, for example), then specific reservations can be set in advance for customer use. This allows flexibility for setting strict reservations where important, and not setting them where it is not important.
- This allocation policy can be an ideal solution as a voice (or other real-time) application-specific Organization vDC, to supplement other vDCs used by the Organization for less critical needs. For example, the customer can run non-critical applications in a vDC using Allocation Pool or Pay-As-You-Go, and run only the real-time application in a Reservation Pool vDC.

For detailed information on configuring the Reservation Pool allocation model, refer to VMware documentation.

13.4.4 Effects of Other Underlying Resources

Aside from the allocation model, other underlying resources used to create the Organization vDC can also have large effects on the performance of the deployed virtual appliances.

13.4.4.1 Storage

Where storage performance is important to the applications to be run in a particular Organization vDC (for example, NuPoint Virtual), then the Organization vDC should be supported by high performance storage resources.

vCloud Director 5.1 adds the concept of “storage profiles”. These must be configured at the vSphere level for vCloud Director, which uses them to automatically select storage for vDCs, based on the specified profile.

The names you assign to the storage should reflect its performance and/or level of protection. For example, a high performance store, with full network RAID 10 and replication protection might be labeled “Platinum”, the same but with no replication might be “Gold”, mid-performance with RAID 10 protection, “Silver”, and low performance with no RAID 10, “Bronze”.

Note:

Due to vMotion requirements, storage for vCloud Director must be shared storage (SAN, vSAN, NFS, VMware VSA), and must be available on all hosts supporting the vCloud Director-provided vDC cluster(s). Host-local storage cannot be used.

13.4.4.2 Network

The same performance considerations apply within vCloud Director as in the underlying vSphere layer.

Note that vCloud Director requires the use of virtual distributed switches (dvSwitch) in the underlying vSphere networks. Regular vSwitches are not compatible.

13.4.5 Modifying vApp Resourcing for Deployment Configuration

As noted in [Deployment differences relative to vSphere](#), when vApps are deployed on vCloud Director, the Configuration Options dialog box is not presented as it is in direct vSphere deployments. Therefore, the administrator cannot choose the specific deployment profile to use (Small Business or Enterprise, for example), and resourcing for the vApp may be incorrect for the desired capacity (too small, or unnecessarily large).

This section describes ways to work around this issue.

Note:

Methods described here can only be fully implemented when the VA or VA Template (Catalog item) are deployed in an Organization vDC configured to use the Reservation Pool allocation model.

To modify the VA resourcing to match the appropriate profile, two methods are available, as follows:

13.4.5.1 Method 1 - Modify Resourcing in vCloud Director

In this approach, the VA is brought into a vApp Template. Then it is deployed in a Reservation Pool based vDC and modified. Finally, the VA is added back into the Catalog as a new vApp Template. This may be done by either Organization or System level administrators.

Follow these steps:

1. Upload the VA into the Catalog (or import from vSphere), as described in previous sections. See [Creating and publishing virtual application Catalog Templates](#). The Templates resourcing can now be inspected (but not modified), by selecting and opening the vApp, selecting the desired contained VM, and opening the Properties dialog box on the VM. The Hardware tab shows the number of vCPUs, memory, and disk assignments; the Resource Allocation tab shows CPU and memory reservations, shares, and limits.

Note:

The Resource Allocation tab is available only when the vApp Template is stored on a Reservation Pool vDC.

2. Deploy the **vApp Template** to **My Cloud**, as described earlier. See [Deploying and configuring virtual applications from catalogs](#).

Important:

During deployment, at the Network Mapping step, the VM's built-in network interface Source name will be modified to match whatever network is selected to connect the source to. For example, the Source name "LAN" or "WAN" will be overwritten to become whatever the name of the selected destination network is: "LAN" will become "My Organization-internal" (or whatever is selected). Therefore, it is important to connect to network names that are easily understood when the modified vApp Template is re-deployed at later stages. Preferably, Organization network names should be configured to match the Source names of the VMs to be modified (the network name is also "LAN" or "WAN," for example). Organization network names can be modified by either Organization or System-level administrators, under the Organization's Administration / Org vDC Networks tab.

Important:

Do not run the deployed vApp (or its VMs) at this stage or configure its operating system (MSL) or any attributes other than resource profile. Any configuration, such as OS-level host name, network configuration, and so on, will be captured in the modified vApp Template Catalog item in later steps, and are likely to be incorrect for the deployment of the modified vApp Template. In this case, the MSL initial configuration will not run (because MSL is already configured).

3. Open the deployed **vApp** (double-click or right-click Open), and open the desired contained VM Properties dialog box (double-click or right-click Properties).
4. In the **General** tab, rename the **VM** appropriately.

5. In the **Hardware** tab (see the figure below), modify the Number of vCPUs, Total memory, or increase Disk size as appropriate to the desired profile. If required, additional disks may be added in this step.

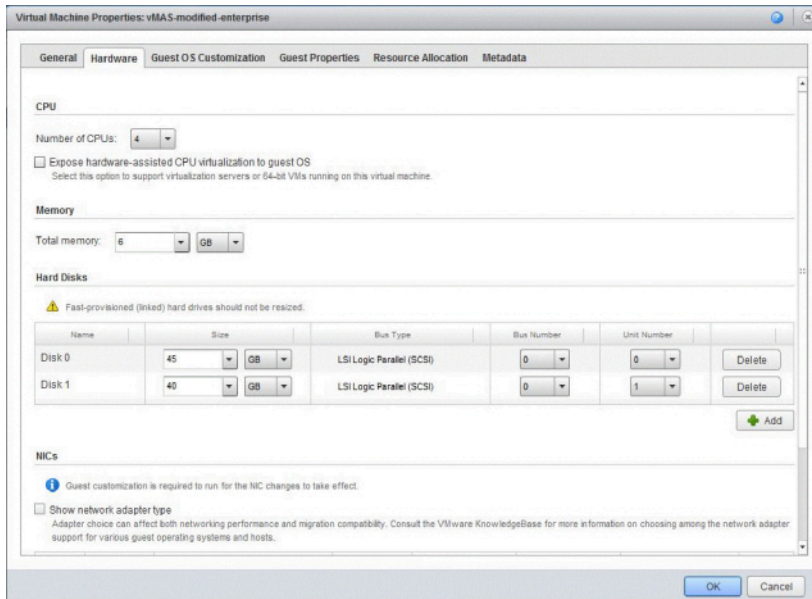


Figure 4: VM Properties - Hardware tab

Note:

- The disk size of an already deployed disk may only be increased, not decreased.
- If the Organization vDC is configured to enable Fast-Provisioning (linked clone), the existing disk(s) cannot be increased or decreased after the VM is deployed.

6. Where [Table 6](#) indicates two disks, a new disk must be added to the VM when it is modified. For example, when modifying from MiCollab Virtual or MiVoice Border Gateway Virtual Small Business profile to the corresponding Enterprise profile, an additional disk of the appropriate size must be configured at this step. If memory is increased, and the corresponding reservation is also to be increased, then the dialog box must be closed (Click OK), and then re-opened for the new value to be saved. Do this before moving to the next step (Resource Allocation).

7. In the **Resource Allocation** tab (see the figure below), modify the vCPU Reservation, Priority (shares) and Limit, as well as Memory Reservation, Priority (shares) and Limit, as appropriate for the desired profile. Disk shares and I/O limits cannot be modified. Click **OK** to save the changes.

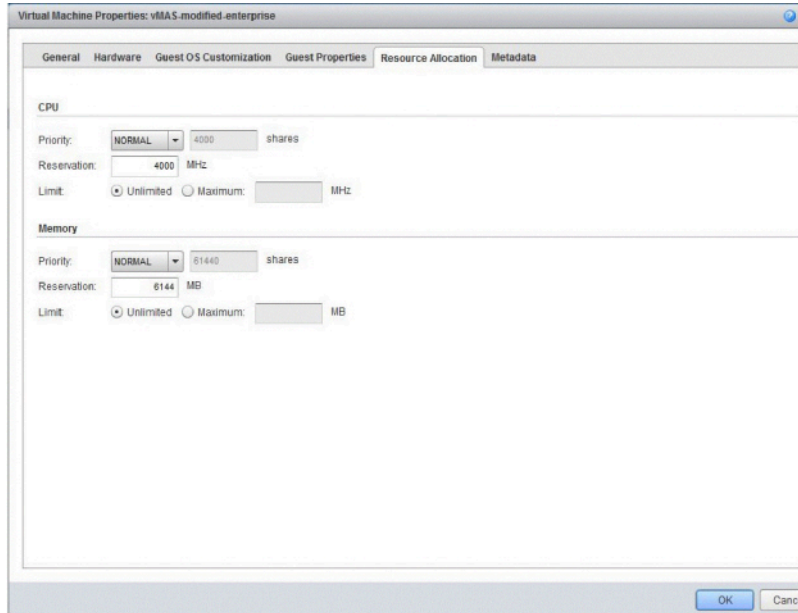


Figure 5: VM Properties - Resource Allocation Tab

Note:

The Resource Allocation tab is available only for vApps deployed in a vDC using the Reservation Pool allocation model.

8. Add the modified vApp into the Catalog as a vApp Template. Right-click Add to Catalog or select the vApp, click Tools, and select Add to Catalog. Enter a suitable name for the vApp Template.
9. Select **Customize VM** settings or Make identical copy as appropriate. Make identical copy will cause all initial configuration at deploy time to be identical to the settings captured in the Catalog Template; Customize VM settings presents Configure Resources, Network Mapping, and Custom Properties steps when vApps are deployed from the Template. Customize VM settings is usually the more appropriate. Click OK to complete the process.
10. After completing the vApp Template, inspect the new Catalog item to ensure that it is correct To inspect the Catalog item:
 - a. Select and open the **vApp Template**.
 - b. Select the desired contained **VM**.
 - c. Open the **Properties** dialog box on the VM.

13.4.5.2 Method 2 - Import from vSphere

In this approach, the VA is first deployed at vSphere level, then into vCloud Director as a vApp Template Catalog item. While the number of vCPUs, memory, and disk are preserved under all allocation models,

only Reservation Pool-based vDC fully preserves vCPU and memory reservations, as described in [Table 27, "Resource Settings vs. Allocation Model,"](#) above.

Note:

This method is available only to system level administrators. Organization level administrators do not have permissions for Import from vSphere.

Follow these steps:

1. Deploy the Virtual Appliance in vSphere:
 - a. In the vSphere Deploy OVF Template wizard, in the Deployment Configuration step, select the desired deployment profile for the VA to be created.
 - b. In the vSphere Deploy OVF Template wizard, in the Disk Format step, select the appropriate format for the VA to be created, as recommended for that application. Typically, Thick Provisioning is recommended for storage intensive real time applications like MiVoice Business Virtual, NuPoint UM Virtual, and MiCollab Virtual.
2. Import the created VA into a vApp Template Catalog item as previously described in [Import VM from vSphere](#).
 - a. **Optional:** In the Import VM as a vApp Template dialog, select Move VM. This has the effect of moving the VM to a new location controlled by vCloud Director. This removes the VM from vSphere control as part of the Import operation, saving space.
 - b. **Optional:** Mark the vApp Template as a Gold Master if appropriate.
3. **Optional:** If the VA was copied in Step2 (not moved), then remove the created VA in vSphere (Delete from Disk), unless it is needed for further Importing to other Catalogs. This saves space.
4. After completing the vApp Template, inspect the new Catalog item to ensure that it is correct. To inspect the Catalog item:
 - a. Select and open the **vApp Template**.
 - b. Select the desired contained VM.
 - c. Open the **Properties** dialog box on the **VM**.

13.5 vCloud Director Lease and Lease Expiry

vCloud Director allows control of several aspects of resource leasing policy when an Organization is created and configured. These policies can also be modified, after the Organization is created and is in use, by the system Administrator or the Organization Administrator.

These controls include:

- **Maximum vApp run-time lease** - how long a deployed VA or vApp will be allowed to run
- **Maximum vApp storage lease** - how long the VA or vApps storage will be maintained
- **vApp Template maximum storage lease** - how long a Catalog Template item will be maintained

- **vApp and vApp Template storage cleanup** - whether expired items will move to the Expired Items folder, or be completely removed

When the VA is deployed from the Catalog, the administrator or user can also modify the lease time assigned to the VA, up to the maximums configured for the Organization. After deployment, the administrator or user must periodically reset the leases for the VA to avoid expiry, unless the lease is set to Never Expires (recommended for mission-critical VAs). This also applies to Catalog Templates (vApp template lease).

13.5.1 Lease Expiry and Mission-critical Applications

When the vApp or VA run-time lease expires, the VA will be shut down, potentially interrupting service. The VA can still be re-started from its storage, as long as the vApp storage has not yet expired.

For mission critical VAs—and most Mitel VAs are mission-critical—you must either:

- Set the run-time and storage leases to Never Expires, or
- Set up automated processes to monitor and track expiry dates, and automate lease renewal to avoid interruptions due to expired leases.

Robust processes for preventing unplanned and unexpected lease expiry are strongly recommended.

13.6 Un-deploying Virtual Applications

To un-deploy a Mitel VA under vCloud Director, the administrator selects the specific VA, powers off, and deletes it, as follows:

1. Click **My Cloud**.
2. Click **vApps**.
3. Select the **VA**.
4. Right-click **Stop** and wait for this to finish.
5. Click **My Cloud**.
6. Click **vApps**.
7. Select the **VA**.
8. Right-click **Delete**.

Un-deploying a VA does not affect other VAs running in the same Organization. Resources are relinquished, allowing other VAs to be started, or granted a larger share of the resources allocated to the Organization vDC.

13.7 Networking and Security Considerations

vCloud Networking and Security (vCNS), and specifically vCNS, is tightly integrated with vCloud Director, and provides flexibility in the ways Organization networks can be configured. The options are:

- **Direct-connected external network:** Allows the application to connect directly to an external network, such as an enterprise LAN (private cloud), extended customer WAN (public cloud with MPLS or other VPN support), or public Internet (public cloud or enterprise private cloud).
- **Network Address Translation (NAT)-routed external network:** Allows connection to external networks only through NAT/firewall. The NAT/firewall is implemented by vCNS, and can be configured by the system or organization administrator for NAT address mapping, firewall rules, and so on.
- **Internal network:** Isolated, and private to the Organization, with no connection to external networks.

These options can be used in combination for any particular organization.

During Organization network creation, the vCloud Director Create Organization Network Wizard provides some typical setup options. Advanced options are also available, which allow you to configure the same set up steps in finer detail.

Organization networks are created when Organizations are initially configured or later, when additional networks are created for an existing Organization. These options are discussed in the following subsections.

13.7.1 Deployment in a Direct-connected, Flat Organization Network

A typical Organization network configuration uses a direct-connected external network, shown in the figure below. A second isolated internal network is not provided, so this is considered a “flat” topology.

Mitel applications can readily be deployed in this topology, by treating the direct-connected external network as if it were an enterprise LAN. In this case, it is assumed that no MBG-based protection of the Mitel applications is required. Mitel applications are deployed to directly connect to the external network, as is usual in a LAN environment.

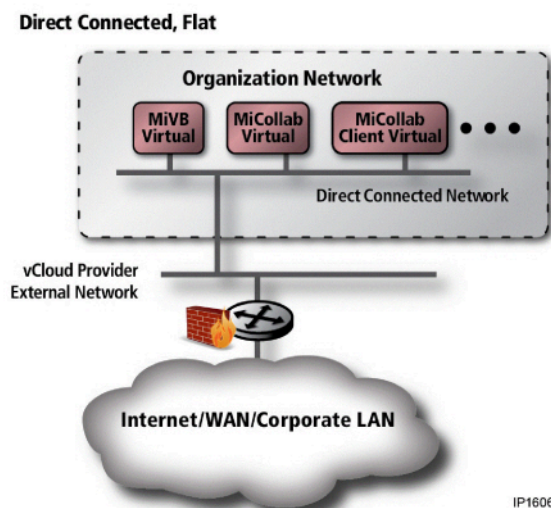


Figure 6: Direct connected, flat network topology

MiVoice Business Express does not fit this network topology. This network topology is appropriate only for stand-alone applications with no MBG protection, and where there is some form of firewall protection provided between the vCloud Director networks and the Internet (or other untrusted network).

Recommendations:

- The direct-connected, flat Organization network topology is recommended for Mitel virtual applications in cases where MBG protection of the applications is not required. Typically, this would be in an Enterprise or SMB private cloud deployment, where the enterprise has direct control over the entire network. In the hosting, public cloud case, the external network must be an extension of the customer WAN, through MPLS or other VPN, for example.
- This topology is not suitable for MiVoice Business Express.
- In many circumstances this topology is not available for deployment due to hosting provider restrictions, for example. When this is the case, another direct-connected (layered) topology or a NAT-routed topology should be used.

13.7.2 Deployment in Direct-connected, Layered Organization Network

In this typical Organization network configuration, applications can connect directly to the external network (where appropriate), with no NAT/firewall in the path, and a second isolated internal network is provided, as shown in the figure below.

Mitel applications can readily be deployed in this topology, by treating the direct-connected external network as if it were the Internet. It may be any of the following:

- Live Internet connection (private cloud DMZ, or public cloud, hosting-provided Internet)
- Enterprise LAN (private cloud)
- The direct connection may map to an external-facing network for a particular customer (public cloud/hosting). In this case, the external network may be an extension of the customer WAN, for example, through MPLS or other VPN.

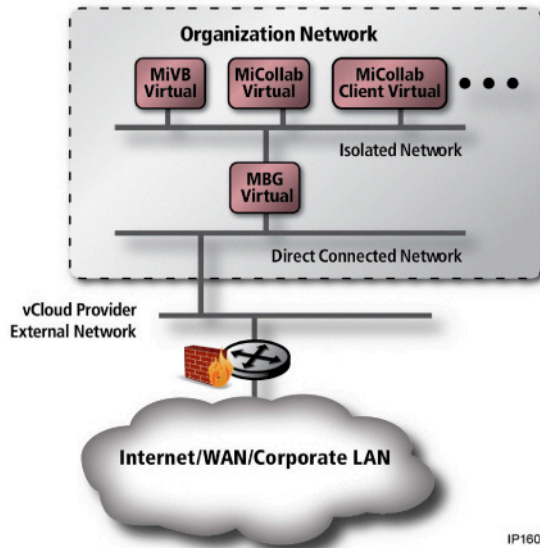


Figure 7: Direct connected, layered network topology

In this case, MBG Virtual is deployed in server-gateway mode, bridging between the direct-connected external network and the internal network. This provides the necessary security to the other Mitel applications, which are deployed on the isolated internal network. External devices such as IP Phones are configured in teleworker mode, following normal MBG deployment rules. Other MBG services, such as web proxy, may also be enabled following normal MBG deployment rules.

Alternatively, MiVoice Business Express may be deployed, with its WAN interface connected to the external network, and its LAN interfaces connected to the isolated internal network.

An example vCloud Director configuration for this is shown in the figure below. This is a screen capture from **vCloud Director Organization > My Cloud > VM** view.

Console	Name	Networks	IP Address
	vMAS	NIC 0 : Organization-01-internal-netw	192.168.1.80
	vMBG	NIC 2 : - NIC 1 : Organization-01-internal-netw NIC 0 : Organization-01-external-netw	- 192.168.1.60 10.45.110.60
	vMCD	NIC 0 : Organization-01-internal-netw	192.168.1.70

Figure 8: Deployment Configuration in Direct Connected Layered Topology

Recommendations:

- Direct-connected, layered Organization network topology is recommended approach for Mitel virtual applications for which MBG protection is required.
- This is preferred over the NAT-routed topology (described in the next section), because it requires reduced configuration overhead for the vCNS component (which is not required in this topology), and reduced risk of vCNS adding to media flow latency and jitter (and potentially reduced voice quality).
- In some circumstances, this topology may not be available for deployment, due to hosting provider restrictions, for example. In that case, the NAT-routed topology should be used.

13.7.3 Deployment in NAT-routed Organization Network

A typical Organization network configuration using NAT/firewall routed external network access provides external network access through a NAT/firewall, plus internal Organization communication through an isolated internal network, as shown in the figure below.

Note:

In a MiCloud environment, MiCloud Management Gateway can be used to replace vShield Edge/vCNS. MiCloud Management Gateway supports connections for up to 100 customers.

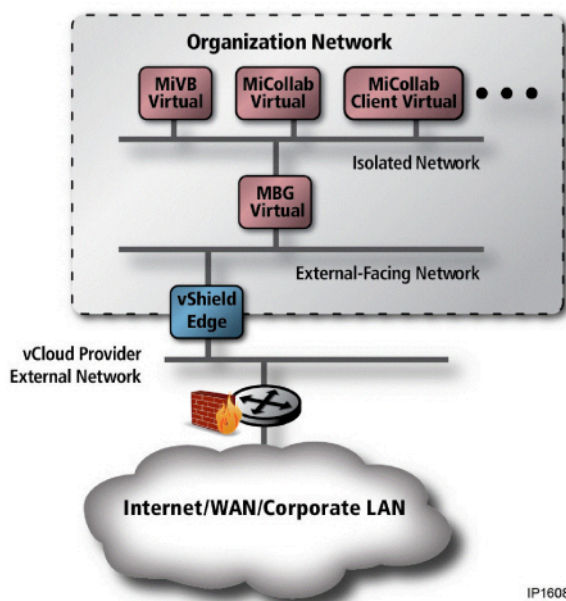


Figure 9: NAT routed, layered topology

In the NAT-routed topology, Mitel applications are deployed behind the vCNS NAT/Firewall, however, some may not operate properly due to the address translation. Operation of Mitel applications behind NAT is application-specific, and the same as the behavior of the application behind a physical NAT device. Refer to the application-specific documentation for details.

The following configurations are supported.

13.7.3.1 Approach 1: Supported by MiVoice Border Gateway Virtual

For this network topology, MiVoice Border Gateway Virtual is deployed between the external-facing network and the isolated internal network. In effect, MBG Virtual is in series with a vCNS NAT/Firewall.

MBG Virtual is configured in server-gateway mode. The vCNS NAT/Firewall is configured to provide a 1:1 mapping of an externally visible, public Internet IP address to the MBG Virtual WAN interface (no port mapping permitted), following normal MBG firewall configuration rules. If MCA is also required, then two address mappings are required. Other UCC applications are then deployed on the internal network only. MiVoice Border Gateway allows IP Phones in Teleworker mode to connect to the MCD, and other applications are connected to the internal network.

An example vCloud Director configuration for this is shown in the figure below. This screen capture is taken from the vCloud Director Organization > My Cloud > VM view.

Console	Name	Networks	IP Address
	vMAS	NIC 0 : Organization-01-internal-netw	192.168.1.80
	vMBG	NIC 2 : - NIC 1 : Organization-01-internal-netw NIC 0 : Organization-01-external-netw	- 192.168.1.60 192.168.100.60
	vMCD	NIC 0 : Organization-01-internal-netw	192.168.1.70

Figure 10: MBG in Series Configuration in NAT Routed Layered Topology

13.7.3.2 Approach 2: Supported by MiVoice Business Express

Deployment of MiVoice Business Express in this network topology is completely equivalent to the MBG Virtual in series with vCNS approach above. In this case, the MiVoice Business Express WAN interface is connected to the external facing network, and the LAN interfaces are connected to the isolated internal network. The vCNS NAT/Firewall is configured in the same way as for the MBG Virtual approach.

Recommendations:

- The NAT-routed/layered Organization network topology is recommended approach for Mitel virtual applications in which MBG protection is required.
- The direct connected/layered Organization network topology described in [Deployment in direct-connected, layered Organization network](#) would generally be preferred, if possible, due to added complexity in configuration imposed by adding NAT in the path and the risk of reduced voice quality due to latency and jitter possibly added by the Edge component (considered low risk).

This chapter contains the following sections:

- [Security in Mitel/VMware Virtual Networks](#)
- [vCNS Security Topologies](#)

14.1 Security in Mitel/VMware Virtual Networks

Note:

This chapter provides information about VMware and Hyper-V applications and requirements that were valid for previous releases, and which may not be up-to-date. Always refer to your current VMware or Hyper-V documentation for the latest information.

When setting up your network with virtual servers, security considerations are the same as for traditional servers, and all of the same security measures, including firewalls, DMZ, authentication, encryption, and SSL-protected connections, apply.

In addition, there are measures you can take specific to your virtualized environment; these are described in this chapter. For detailed VMware product information, see the VMware documentation.

14.1.1 vCloud Networking and Security (vCNS) Overview

This section provides a brief overview of VMware's vCloud Networking and Security (vCNS) family of security products.

Note:

vCNS 5.1 and above only, are supported for use with Mitel virtual appliances.

Security components of vCNS were formerly known as the VMware product suite.

Additional networking functionality is included in vCNS, but is not directly relevant to Mitel deployments, so it is not described or discussed here. Please refer to VMware documentation for details.

14.1.1.1 vShield Manager

vShield Manager provides centralized management for the other vCNS/vShield security products. One instance of vShield Manager is installed and configured for each vCenter in the overall environment.

vShield Manager is accessed through a plug-in to vCenter, which appears in the vSphere Client the administrator uses to manage the vCenter, or through a Web administration interface on the vShield Manager application.

14.1.1.2 vCNS

vCNS provides gateway services at the network edge. The basic architecture is similar to that of a NAT/Router in the physical world, as shown in the figure below.

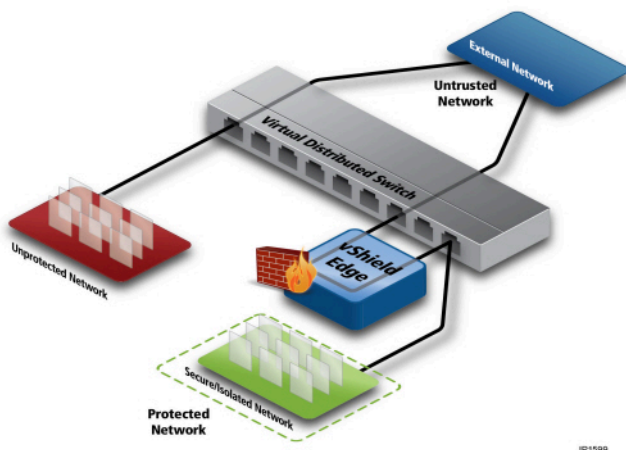


Figure 11: vCNS architecture

vCNS is deployed as one or more virtual appliances within the virtual infrastructure. vCNS instances may be deployed between sections of the data center network, at the external facing edge of the data center network, in a virtual DMZ.

Capabilities include:

- Network Address Translation (NAT) with external and internal address mapping capability, inbound and outbound access rules, and so on.
- Firewall, with 5-tuple inbound and outbound access rules.

- vCNS High Availability, with two vCNS instances in active-passive configuration, allowing stateful fail-over (vCNS Advanced only).

Note:

While the overall service will recover, there could be some delay to vCNS fail-over, so there would be loss of voice or other media flows for that period and signaling delays as well.

Note:

This fail-over of vCNS instances is not equal or equivalent to the vSphere HA feature, which provides automatic restart of VMs in case of VM failure, or failure of the host on which the VM resides.

- vCNS is available at different performance levels: Compact, Large and X-Large, allowing scaling up of the solution. In-place vCNS appliances can be transformed to the larger capacity configuration, but cannot be changed to a lower performance level.
- Site-to-site VPN.
- Load balancing across multiple external-facing interfaces.
- DHCP address configuration on the internal-facing network.
- DNS relay service for the internal network.

vCNS-based security has several benefits:

- Simple, familiar VLAN-based model, allowing administrators to easily understand their security topology, and configure it.
- Using NAT capabilities, different tenant “internal” networks can be made identical to each other (or nearly so), allowing for easy mass deployment of multiple customers in a hosted environment, with internal addressing pre-determined.
- Using vCNS High Availability feature, reduced service interruption in the case of failure of the (currently) active vCNS appliance.

vCNS, on its own, does not provide protection against network-internal threats (threats from inside the NAT boundary). Nor does it provide isolation of network-internal concerns, as may be required for regulatory compliance (like Sarbanes Oxley or Payment Card Industry (PCI) Data Security Standard, for example).

14.2 vCNS Security Topologies

The following subsections describe five vCNS-based security topologies, including configuration details, recommended usage, and pros and cons of each. Each of these topologies uses the vShield component of vCNS.

Many other approaches are feasible, including use of other vCNS components such as vShield App, depending on overall solution requirements and architecture. The topologies described in the following sections have been tested and verified, so they are supported and recommended.

14.2.1 General Best Practices

This section describes general considerations, applicable across all security topologies. The following best practices should be observed for all of the topologies described in this chapter.

- Use of VMware distributed virtual switching (dvSwitch) as the foundation of networking is strongly recommended. Use of dvSwitches greatly simplifies network configuration, and ensures consistency and correctness. dvSwitch is also a requirement for several vCNS features.
- Where UCC traffic flows through a vCNS, always deploy the Edge appliance in at least the Large profile. (Big UCC systems may require Extra Large.) This is to ensure adequate throughput to maintain high voice quality. This applies to topologies 2 - 4 described below, or any similar topology used.

Note:

Due to the possibility of reduced voice quality, use of vCNS Compact sizing is not recommended or supported, except where no UCC traffic will flow through the vCNS appliance.

- Deploy all security related components (vCNS, MBG Virtual, MiVoice Business Express) on clusters with vSphere DRS/vMotion and HA enabled. Security related components should be configured in HA for high restart priority.
- Enable vCNS High Availability on all vCNS appliances. This ensures maximum continuity should one of the active-passive pair of Edge appliances fail. Note that this feature is not the same as vSphere HA, and that there are some serious limitations for VoIP applications. See [vCNS](#) for more information.

14.2.2 Topology 1: vCNS Based Private Cloud - MBG Virtual Parallel with vCNS

The following figure shows a simple private cloud network topology, using vCNS as the general access control mechanism in the DMZ, between Internet (or other untrusted network) and the enterprise-internal network applications. In this topology, MBG Virtual is placed in parallel to vCNS, and is used for UCC-specific access and NAT handling.

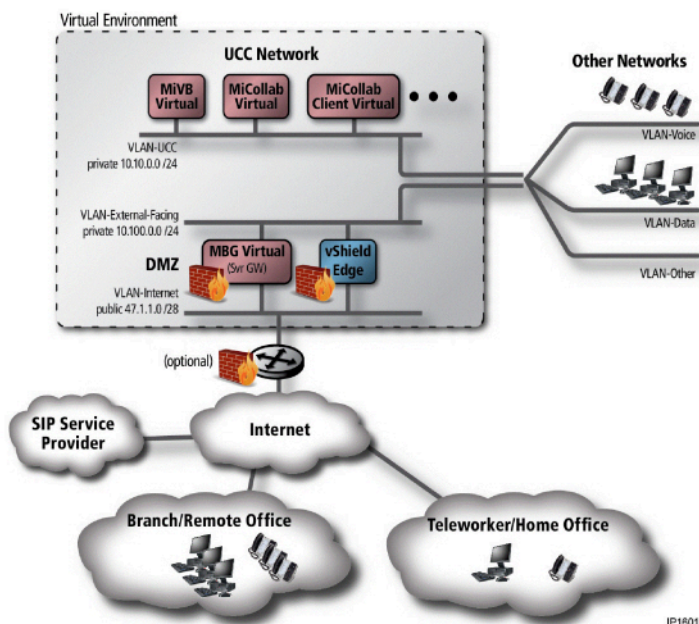


Figure 12: Private Cloud, vCNS Controlled - Parallel

In the diagrams, VMware elements are shown in blue shading; Mitel elements are shown in red shading.

This topology is suitable for enterprise private cloud type deployments, where the enterprise administrator has full control over their network and virtual infrastructure. It is assumed that network-internal threats and strong isolation between internal networks are not major concerns.

Normally, the virtualized infrastructure would be on premise (in the corporate data center). However, it may be feasible for the infrastructure to be hosted externally by a provider as an IaaS deployment, using a VLAN technology such as MPLS to extend the enterprise network to inside the host data center. vCNS site-to-site VPN capability could also be used. Such an extended enterprise network would connect on the “internal side” nominally at the same place MBG Virtual internal interface connects (VLAN-External-Facing in Figure 11).

14.2.2.1 Topology Description

The vCNS controlled Private Cloud - Parallel topology consists of the following elements:

- Some number of remote networks, such as branch or home offices and/or Teleworker users. These access the private cloud across Internet (or other non-trusted network). These are typically protected by a NAT/firewall device (not shown). All endpoints in these remote networks are assumed to be in Teleworker mode of operation.
- Internet public network.
- Access Router, providing access to and from the untrusted network, usually the Internet.
- (Optional) In practice, there may also be a physical firewall in the Internet path, behind or as part of the access router. This is quite likely in large enterprise case, for protection of non-virtualized resources inside the enterprise network.
- DMZ network layer, where machines meant to span between untrusted and trusted network domains reside. The external side of the DMZ attaches to the untrusted network (VLAN-Internet in the diagram), typically using public IP addressing. The internal side attaches to the trusted internal network (VLAN-

External-Facing in the figure). Machines in this layer are virtualized, and may be running on the same host servers as other non-DMZ VMs. Access to the DMZ VM vNICs is strictly controlled through the use of VLANs.

- MBG Virtual, residing in the DMZ network layer, and configured in server-gateway mode, provides firewall, NAT handling, and access, specifically to and from the Mitel UC applications (MiVoice Business Virtual, and so on). The MBG Virtual may also be used for the SIP Service Provider gateway function (SBC), if required. The external-facing “WAN” interface of the MBG Virtual attaches to the external untrusted network (VLAN-Internet), and the internal-facing LAN interface attaches to the internal trusted network (VLAN-External-Facing).
- vCNS, residing in the DMZ network layer, and providing NAT/firewall service to general purpose VMs in the overall solution. The external-facing interface of the vCNS instance attaches to the external untrusted network (VLAN-Internet), and the internal-facing interface attaches to the internal trusted network (VLAN-External-Facing).

Note:

In this topology, vCNS is not directly required to support Mitel UCC applications. vCNS and MBG do not directly interact, only co-exist. If only UCC applications are included in the virtual environment, the vCNS component is not required.

- UCC applications network containing the Mitel UCC application VAs, other than the MBG Virtual gateway (VLAN-UCC in the diagram). These are configured as they would be in a traditional LAN environment. Note that MBG Virtual may still be present, if required for internal services such as Secure Recording Connector (SRC).
- One or more “Other” internal networks. In the private cloud case, this is likely to include separate networks for end user devices such as PCs, a voice network for VoIP devices such as IP Phones, and possibly several more. Some of these other networks may exist inside the virtual infrastructure (rather than external to it, as shown here). Devices used to manage the virtual infrastructure (administrator PCs, for example) may reside on one or more of these networks.

For information about configuring vCloud Networking and Security (formerly vShield Edge), MBG Virtual, the IP Phones (and other VoIP end-points), and the external firewall (if one exists in the deployment), see [Table 29, “Network component configuration by Topology type,”](#).

For a comparison of the advantages and disadvantages, and the recommended applications of each topology, see [Table 28, “Topology comparisons and recommendations,”](#).

14.2.3 Topology 2: vCNS Based Private Cloud - MBG Virtual Behind vCNS

The following figure shows a simple private cloud network topology, using vCNS as the general access control mechanism in the DMZ, between Internet (or other untrusted network) and the enterprise-internal network applications. In this topology, MBG Virtual is placed behind vCNS, in series with the vCNS virtual firewall/NAT, and is used for UCC-specific access and NAT handling. In this topology, MBG Virtual provides an additional security layer controlling access to UCC applications.

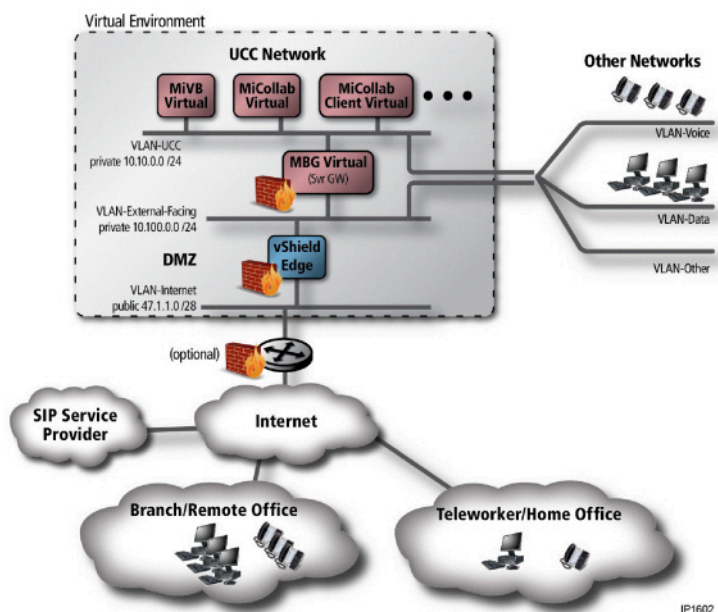


Figure 13: Private Cloud, vCNS Controlled - Series

This topology is similar to the topology described in the previous section, but moves the MBG Virtual function “inside”, behind the vCNS NAT/firewall.

This topology is appropriate for enterprise private cloud type deployments, where the enterprise administrator has full control over their network and virtual infrastructure. It is assumed that network-internal threats and strong isolation between internal networks other than UCC are not major concerns. Additional layering of security is provided for UCC applications, so security from internal threats on UCC is stronger.

14.2.3.1 Topology Description

The vCNS controlled Private Cloud - Series topology consists of the following elements. Only elements different from the previous topology are described here.

- MBG is moved to the internal network, behind vCNS and the (optional) physical firewall. MBG Virtual is in series with the vCNS NAT/firewall.
- MBG Virtual is configured in Server-Gateway mode, providing firewall, NAT handling and access to and from the Mitel UCC applications (MiVoice Business Virtual, and so on).
- The MBG Virtual may also be used for SIP Service Provider gateway function, if required.
- The external-facing WAN interface of the MBG Virtual attaches to an internal trusted network (VLAN-External-Facing in the diagram), and the internal-facing LAN interface attaches to a different internal trusted network (VLAN-UCC in the diagram).
- vCNS, residing in the DMZ network layer, provides address and port mapping for Mitel UCC applications via MBG Virtual, as well as provides NAT/firewall service to general purpose VMs in the overall solution.
- Specific vCNS configuration is required to support Mitel UCC/MBG Virtual interface.

For information about configuring vCloud Networking and Security (vCNS), MBG Virtual, the IP Phones (and other VoIP end-points), and the external firewall (if one exists in the deployment), see [Table 29](#), “Network component configuration by Topology type,”.

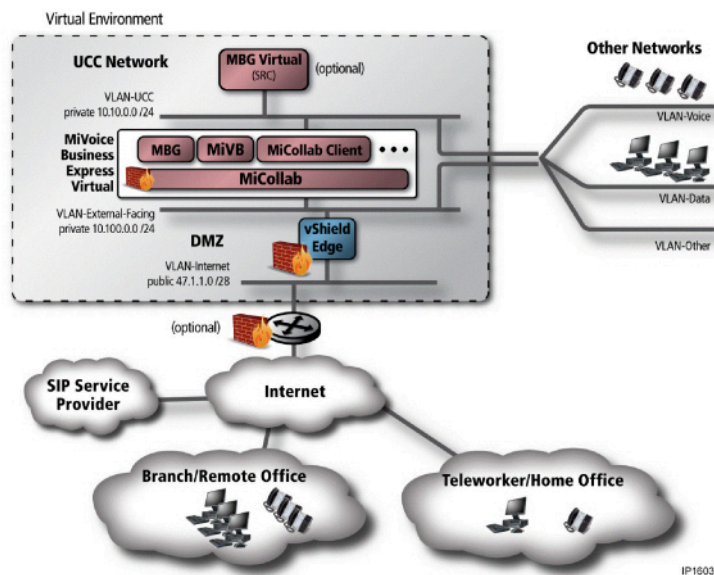
For a comparison of the advantages and disadvantages, and the recommended applications of each topology, see [Table 28](#), “Topology comparisons and recommendations,”.

14.2.4 Topology 3: vCNS Based Private Cloud - MiVoice Business Express

The following figure adds deployment of MiVoice Business Express, and from the security perspective is really a direct extension of the previous topology. vCNS is used as the general access control mechanism in the DMZ, between Internet (or other untrusted network) and the enterprise-internal network applications. In this topology, MiVoice Business Express is placed behind vCNS, in series with the vCNS virtual firewall/NAT. The MBG component of MiVoice Business Express, which is configured in Server-Gateway mode, provides for UCC-specific access and NAT handling. In this topology, the MBG component adds an additional security layer controlling access to UCC applications.

For MiVoice Business Express, two IP addresses are used on each of the external-facing connections (VLAN-External-Facing), and on the internal-facing connection (VLAN-UCC). This is required for MSL/MiCollab/MBG external addressing and for MCA proxy external address, as well as the MSL/MiCollab/MBG internal address, and the MiVoice Business internal address.

Figure 14: Private Cloud, vCNS Controlled - MiVoice Business Express



This topology is suitable for enterprise private cloud deployments, where the enterprise administrator has full control over the network and virtual infrastructure. It is assumed that network-internal threats, and strong isolation between internal networks other than UCC are not major concerns. Additional layering of security is provided for UCC applications by the MBG component of MiVoice Business Express, so security from internal threats on UCC is stronger than in the previous two topologies described. In addition, use of MiVoice Business Express simplifies deployment.

14.2.4.1 Topology Description

The vCNS controlled Private Cloud - MiVoice Business Express topology consists of the following elements. Only elements different from the previous topologies are described here.

- MiVoice Business Express is deployed between the trusted side of the DMZ and the UCC network. In effect, the MBG component of MiVoice Business Express is in series with the vCNS NAT/firewall.
- In MiVoice Business Express, the MBG component is configured in Server-Gateway mode, providing firewall, NAT handling and access to and from the other MiCollab applications.
- The MBG component may also be used for SIP Service Provider gateway function, if required.
- The external-facing WAN interface of MiVoice Business Express attaches to an internal trusted network (VLAN-External-Facing in the diagram), and the internal-facing LAN interface attaches to a different internal trusted network (VLAN-UCC in the diagram). Two fixed IP addresses are assigned to each interface.
- vCNS, residing in the DMZ network layer, provides address and port mapping for MiVoice Business Express, and provides NAT/firewall service to general purpose VMs in the overall solution. For MiVoice Business Express, two externally visible fixed IP addresses are mapped to MiVoice Business Express WAN addresses.

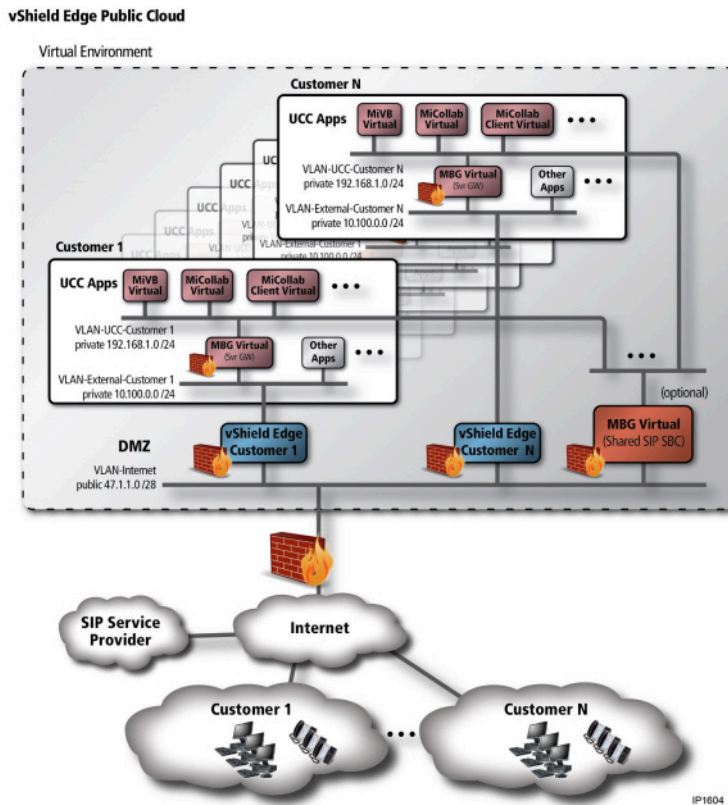
For information about configuring vCNS, MBG Virtual, the IP Phones (and other VoIP end-points), and the external firewall (if one exists in the deployment), see [Table 29, “Network component configuration by Topology type,”](#).

For a comparison of the advantages and disadvantages, and the recommended applications of each topology, see [Table 28, “Topology comparisons and recommendations,”](#).

14.2.5 Topology 4: vCNS Based Public Cloud

The following figure shows a public cloud network topology, using vCNS as the access control mechanism for each tenant/customer in the DMZ, between Internet (or other untrusted network) and the customer networks in the hosting environment. vCNS also provides isolation between customers. MBG Virtual, in series with each customer's vCNS, performs the NAT traversal proxy function for the Mitel UCC applications.

Figure 15: Public Cloud, vCNS Controlled



This topology is well-suited for hosted multi-tenant public cloud deployments, providing UCC as a service (UCaaS), possibly along with other services. The hosting provider has full control over their network and virtual infrastructure; however the customer is assumed to have limited control. It is assumed that network-internal threats within a particular customer are not a concern. Strong isolation between customers is a very high concern.

Note:

This topology maps very closely to VMware vCloud Director NAT routed Organization Networks, where deployment of vCNS appliances to isolate customer networks is automated by vCloud Director. See “vCloud Director” on page 97.

If only UCC applications are being hosted, then the vCNS layer is not required for UCC functionality. (It may, however, be required by the host provider for other reasons.) In this case the MBG Virtual could be connected directly to the Internet, as a normal Server Gateway deployment, without vCNS in the path. This variation is not discussed in detail in this guide; refer to the MiVoice Border Gateway documentation for more information.

14.2.5.1 Topology Description

The Public Cloud / vCNS topology is largely an expansion of the previous topology ([Topology 2: vCNS Based Private Cloud - MBG Virtual Behind vCNS](#)), by completely replicating the topology across multiple customers. vCNS is used to isolate customers from each other. Within each customer network, MBG

Virtual performs the NAT traversal proxy function, as well as providing an additional firewall layer in front of the UCC applications.

Optionally, a shared MBG Virtual (or cluster of MBG Virtual instances), acting as SIP Proxy/Session Border Controller (SBC) for multiple customers, is also introduced. Consolidation of capacity is provided for multiple customers. All trunks come in to one location, and calls are routed to the appropriate customer's UCC environment (MiVoice Business Virtual or MiVoice Business Express). This has potential advantages, but at the cost of forcing customer networks to be not completely identical to each other.

- May enable reduction of total resource reservations across the overall hosting solution to provide the SIP SBC services across multiple customers, due to statistical multiplexing (peak loads for multiple customers are unlikely to occur at the same time). It could also reduce the resource reservations required per customer MBG instance, while creating a reservation for the shared MBG function that is lower than the total—reduced across the customer base. This potential advantage can be expected to be higher where there are large numbers of smaller end customers involved.
- Potential cost savings if SIP services are purchased in bulk by the hosting provider.
- Increased simplicity for bandwidth management. All SIP trunking becomes a single bandwidth management zone.
- Increased resiliency of the SIP SBC functions, using MBG clustering.

In a public cloud environment, it can be assumed that an external firewall will usually also be in place, outside of the virtualized environment.

The Public Cloud/vCNS topology consists of the following elements.

- Each individual customer is provided one or more separate external network addresses (Internet addresses), separated from other customers, from within the public Internet address range operated by the hosting company (VLAN-Internet). Each customer's IP addresses are configured as the external facing interface(s) of the customer-specific vCNS instance.
- Since each customer network is isolated by vCNS, internal addressing is allowed to be identical, or overlapped, between customers. This allows Mitel (and other) components within the customer network to be pre-configured with respect to addressing when the customer is deployed.
- Optionally, a shared MBG Virtual (or cluster of MBG Virtual instances), acting as SIP Proxy/SBC for multiple customers, can also be deployed. This connects to each individual customer internal network, nominally at the same point as the MBG Virtual internal interface. This is the expected scenario in most hosting environments.

Note:

If a shared SIP SBC is used, then customer internal networks can not be identical to each other. Each customer network must have a different UCC network address range/subnet, so that the shared SIP SBC can route customer-specific SIP traffic correctly.

- If a shared MiVoice Border Gateway (MBG) is provided, then the customer MBG Virtual is used for Teleworker access and Secure Recording Connector only. If there is no shared MBG, then the customer MBG Virtual is also used as SIP Proxy/SBC for the individual customer, as required.
- The per-customer MBG Virtual can also be used for remote administration and other purposes, as per normal MBG operation. Refer to the MiVoice Border Gateway documentation for more information.

For information about configuring vCNS, MiVoice Border Gateway, the IP Phones (and other VoIP endpoints), and the external firewall (if one exists in the deployment), see [Table 32, “Network component configuration by Topology type,”](#).

For a comparison of the advantages and disadvantages, and the recommended applications of each topology, see [Table 31, “Topology comparisons and recommendations,”](#).

14.2.6 Topology 5: vCNS Based Public Cloud - MiVoice Business Express

The following figure shows a public cloud network topology, using vCNS as the access control mechanism for each tenant/customer in the DMZ, using MiVoice Business Express as the UCC application set. vCNS also provides isolation between customers. The MiVoice Business Express MBG component, in series with each customer's vCNS, performs the NAT traversal proxy function.

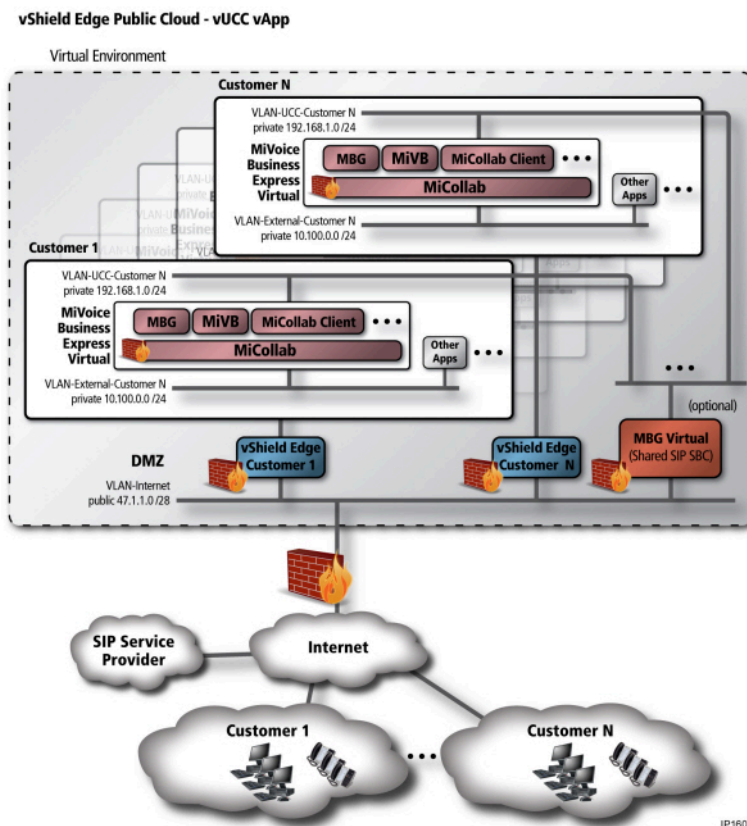


Figure 16: Public Cloud, vCNS Controlled - MiVoice Business Express

This topology is appropriate for hosted multi-tenant public cloud deployments, providing UCC as a service (UCaaS), possibly along with other services, if required. The hosting provider has full control of the network and virtual infrastructure; however, the customer is assumed to have limited control. It is assumed that network-internal threats within each customer are not a concern. Strong isolation between customers is a very high concern. Use of MiVoice Business Express simplifies deployment and initial configuration.

Optionally, a shared MiVoice Border Gateway (or cluster of MBG Virtual instances), acting as SIP Proxy/SBC for multiple customers, can also be introduced. A discussion of advantages and disadvantages of the shared MBG are discussed in Topology 4.

Note:

This topology maps very closely to the VMware vCloud Director NAT-routed Organization Networks. This deployment of vCNS appliances to isolate customer networks is automated by vCloud Director. For more information, see “vCloud Director” on page 97.

14.2.6.1 Topology Description

The Public Cloud/vCNS with MiVoice Business Express topology is largely an expansion of the previous [Topology 3: vCNS Based Private Cloud - MiVoice Business Express](#), completely replicating the topology across multiple customers. vCNS is used to isolate customers from each other. Within each customer network, MiVoice Business Express MBG component performs the NAT traversal proxy function, as well as providing an additional firewall layer in front of the UCC applications.

Optionally, a shared MBG Virtual (or cluster of MBG Virtual instances), acting as SIP Proxy/SBC for multiple customers, is also introduced.

In a public cloud environment, it can be assumed that an external firewall will also be in place, outside of the virtualized environment.

The Public Cloud/vCNS MiVoice Business Express topology consists of the following elements. Only elements different from the previous topology are described here.

- MiVoice Business Express is deployed between the vCNS internal interface and the UCC network. This is functionally equivalent to the previous scenario.

For information about configuring vCNS, MBG Virtual, the IP Phones (and other VoIP end-points), and the external firewall (if one exists in the deployment), see [Table 29, “Network component configuration by Topology type,”](#).

For a comparison of the advantages and disadvantages, and the recommended applications of each topology, see [Table 28, “Topology comparisons and recommendations,”](#).

14.2.7 Topology Advantages, Disadvantages, and Recommended Applications

The following table summarizes the advantages and disadvantages of each of the five topologies, it summarizes the situations best suited to each topology.

Table 32: Topology comparisons and recommendations

	Advantages	Disadvantages	Recommended application
Topology 1 vCNS Based Private Cloud - MBG Virtual Parallel with vCNS			
	<ul style="list-style-type: none"> • Very simple configuration and management for both vCNS and MBG Virtual (basic default configurations). • No coupling of MBG Virtual configuration requirements to vCNS, or visa-versa. • Lowest effort topology to implement • Similar to traditional VLAN-based approach. Easily understood and familiar to administrator. 	<ul style="list-style-type: none"> • No protection from network-internal threats • No isolation of concerns (separate trust zones, compliance zones, tenants) 	<ul style="list-style-type: none"> • Private cloud/ enterprise deployments, where admins have full control of the internal network and virtual environment, and network-internal threats are not a strong concern.
Topology 2 vCNS Based Private Cloud - MBG Virtual Behind vCNS			

	Advantages	Disadvantages	Recommended application
	<ul style="list-style-type: none"> • Sample configuration and management, adding only 1:1 address mapping rule at vCNS. • Additional security specifically for applications behind MBG Virtual (on MiVoice Business Express VLAN) • Similar to traditional VLAN-based approach. Easily understood and familiar to administrator. 	<ul style="list-style-type: none"> • Additional configuration overhead on vCNS (minimal). • Risk that vCNS could impact voice/video quality under load, due to added delay and jitter in media streams, as compared to Topology 1. See vCNS. • No protection from network-internal threats, other than MiVoice Business Express VLANs. • No isolation of concerns (separated trust zones, compliance zones, tenants) 	<ul style="list-style-type: none"> • Private cloud/enterprise deployments, where admins have full control of the internal network and virtual environment, and network-internal threats are not a strong concern. • Networks where vCNS is already deployed as a firewall, and customer is not willing (or able) to allow MBG Virtual in parallel with the vCNS appliance (cannot implement Topology 1). • Networks where additional protection from network-internal threats for UCC applications is desired. • Networks requiring alignment with MiVoice Business Express in future.
<p>Topology 3</p> <p>vCNS Based Private Cloud - MiVoice Business Express</p>			

	Advantages	Disadvantages	Recommended application
	<ul style="list-style-type: none"> • MiVoice Business Express simplified deployment and initial configuration. • Simple configuration and management adding only two 1:1 address mapping rules at vCNS. • Additional security specifically for applications behind the MBG component of MiVoice Business Express (on UCC VLAN). • Similar to traditional VLAN-based approach. Easily understood and familiar to administrator. 	<ul style="list-style-type: none"> • Additional configuration overhead on vCNS (minimal). • Risk that vCNS could impact voice/video quality under load, due to added delay and jitter in media streams. This is a relatively lower risk than in Topology 2, due to the relatively small scale of MiVoice Business Express, unless the vCNS Appliance is under high load for other reasons (non-UCC loads). • No protection from network-internal threats, other than MiVoice Business Express VLANs. • No isolation of concerns (separated trust zones, compliance zones, tenants) 	<ul style="list-style-type: none"> • Recommended for simplified deployment and configuration, using MiVoice Business Express. • Private cloud/enterprise deployments, where admins have full control of the internal network and virtual environment, and network-internal threats are not a strong concern. • Networks where additional protection from network-internal threats for UCC applications is desired.
<p>Topology 4</p> <p>vCNS Based Public Cloud</p>			

	Advantages	Disadvantages	Recommended application
	<ul style="list-style-type: none"> • Strong isolation between customers, provided by vCNS. • Simple configuration and management on a per-customer basis. • If shared MiVoice Border Gateway SIP SBC function is used, potential reduced overall resource reservation, cost and administration advantages. <div style="background-color: #e1f5fe; padding: 5px; margin: 5px 0;"> <p>i Note: This precludes having all customer networks completely identical.</p> </div> <ul style="list-style-type: none"> • Simplified mass deployment re: addressing within individual customers, if identical customer-internal network addressing, or easily mapped. Customer apps can be pre-configured for networking (other than MBG Virtual and vCNS external interfaces). • Similar to traditional VLAN-based approach, within each customer. • Maps directly to vCloud Director NAT-routed Organization networks. 	<ul style="list-style-type: none"> • Risk that vCNS could impact voice/video quality under load, due to added delay and jitter in media streams. Likely a lower risk in Public Cloud deployments, due to expected relatively smaller customer size. • If the shared MiVoice Border Gateway SIP SBC option is used, then per-customer networks cannot be completely identical. • No protection from customer network internal threats. Expected to be a low concern in hosted environments, where the internal network is highly controlled. • No isolation of concerns (separated trust zones, compliance zones). Low concern in hosted environments. 	<ul style="list-style-type: none"> • Public cloud UCaaS or IaaS deployments, where the host admins have full control of the internal network and virtual infrastructure, and customer network-internal threats are not a strong concern. • If shared MBG SIP SBC option does not provide significant advantage (cost saving, reduced admin overhead), then not recommended, in favor of advantages of making all customer internal networks identical.

	Advantages	Disadvantages	Recommended application
<p>Topology 5</p> <p>vCNS Based Public Cloud - MiVoice Business Express</p>			

	Advantages	Disadvantages	Recommended application
	<ul style="list-style-type: none"> • Strong isolation between customers, provided by vCNS. • Simple configuration and management on a per-customer basis. • Simplified deployment and initial configuration of the UCC applications, due to use of MiVoice Business Express. • If the shared MiVoice Border Gateway SIP SBC function is used, potential reduced overall resource reservation, cost and admin advantages. <p>Note: This precludes having all customer networks completely identical.</p> <ul style="list-style-type: none"> • Simplified mass deployment re: addressing within individual customers, if identical customer-internal network addressing, or easily mapped. Customer apps can be pre-configured for networking (other than MBG Virtual and vCNS external interfaces). 	<ul style="list-style-type: none"> • Risk that vCNS could impact voice/video quality under load, due to added delay & jitter in media streams. Low risk in Public Cloud MiVoice Business Express deployments, due to expected relatively smaller customer size as well as relatively smaller scale of MiVoice Business Express. • If the shared MiVoice Border Gateway SIP SBC option is used, then per-customer networks cannot be identical. • No protection from customer network internal threats. Likely a low concern in hosted environments, where the internal network is highly controlled. • No isolation of concerns (separated trust zones, compliance zones). Likely a low concern in hosted environments. 	<ul style="list-style-type: none"> • Recommended for simplified deployment and configuration, using MiVoice Business Express. • Public cloud UCaaS or IaaS deployments, where the host admins have full control of the internal network and virtual infrastructure, and customer network-internal threats are not a strong concern. • If shared MiVoice Border Gateway SIP SBC does not provide significant advantage (cost saving, reduced admin overhead), then it is recommended to not use it. Instead, take advantage of the benefits of making all customer internal networks identical.

	Advantages	Disadvantages	Recommended application
	<ul style="list-style-type: none"> • Similar to traditional VLAN-based approach, within each customer. • Maps directly to vCloud Director NAT-routed Organization networks. 		

14.2.8 Network Component Configuration by Topology Type

The following table provides configuration instructions for the five topology types described in this chapter. For detailed descriptions of the advantages, disadvantages and recommended uses for each topology, see [Table 28, "Topology comparisons and recommendations,"](#).

Table 33: Network component configuration by Topology type

	Network component	Configuration considerations
Topology 1		
vCNS Based Private Cloud - MBG Virtual Parallel with vCNS		
	vCNS	<ul style="list-style-type: none"> • No UCC-specific configuration required. • Configure for needs of any non-UCC applications (web access from user desktops, or other applications requiring Internet access, for example).

	Network component	Configuration considerations
	MiVoice Border Gateway Virtual	<ul style="list-style-type: none"> • Standard Server-Gateway configuration • If UCC applications are on a different LAN segment (subnet) than the MiVoice Border Gateway internal interface (VLAN-External-Facing in the diagram), then you must configure LAN gateway address (address of the internal router) as a routing rule for accessing these applications. Refer to the MiVoice Border Gateway documentation describing “Additional Local Networks”. • If there is an external firewall, configure the public WAN address as the IP address of the external firewall. • If MiCollab Audio, Web, and Video is part of the solution, then you must configure a second public WAN IP address for MiCollab AWW Connection Point HTTPS, and configure the LAN address of the MiCollab AWW Virtual application. MiVoice Border Gateway then forwards HTTPS (port 443) in the second WAN interface to MiCollab AWW. Refer to the MiVoice Border Gateway documentation for details.
	IP Phones, other VoIP endpoints	<ul style="list-style-type: none"> • All external endpoints are in Teleworker configuration. • If there is an external firewall, then internal end points must not be in the Teleworker configuration.

	Network component	Configuration considerations
	External firewall (if it exists)	<ul style="list-style-type: none"> • Configure static 1:1 external to internal address mapping for MBG Virtual external (WAN) interface. • The NAT/firewall function must preserve the TCP and UDP port numbers in packets exchanged between the MBG Virtual and the external network. In other words, only the address field may be changed. • The public address bound to the MiVoice Border Gateway Virtual external facing address must be a static IP address, visible from the external network (Internet). This should be a separate address from the external IP address of the vCNS (or other upstream firewall). <p>Note: If external access to MiCollab AWW services is required, then a second 1:1 external to internal address mapping, similarly configured, must be provided for the second MBG Virtual address corresponding to the virtual MiCollab AWW.</p>
<p>Topology 2</p> <p>vCNS Based Private Cloud - MBG Virtual Behind vCNS</p>		

	Network component	Configuration considerations
	vCNS	<ul style="list-style-type: none"> • Configure static 1:1 external to internal address mapping for MBG Virtual external (WAN) interface. • The NAT/firewall function must preserve the TCP and UDP port numbers in packets exchanged between the MBG Virtual and the external network. In other words, only the address field may be changed. • The public address bound to the MBG Virtual external facing address must be a static IP address, visible from the external network (Internet). This should be a separate address from the external IP address of the vCNS (or other upstream firewall). • If MiCollab AWV is part of the solution, then you must also configure a second 1:1 external to internal address mapping for the second MiVoice Border Gateway WAN address required for MiCollab AWV Connection Point HTTPS. • Configure for needs of any non-UCC applications (web access from user desktops, or other applications requiring Internet access).

	Network component	Configuration considerations
	MiVoice Border Gateway Virtual	<ul style="list-style-type: none"> • Standard Server-Gateway configuration. • Configure the public WAN address as the IP address of the vCNS public IP address (or as the public IP address of the external firewall, if present). • If MiCollab AWV is part of the solution, then you must also configure a second public WAN IP address for MCA Connection Point HTTPS and configure LAN address of the MiCollab AWV, as in previous topology.
	IP Phones, other VoIP endpoints	<ul style="list-style-type: none"> • All external endpoints are in Teleworker configuration. • Internal endpoints must not be in Teleworker configuration.
	External firewall (if it exists)	<ul style="list-style-type: none"> • Configure as per vCNS. It is basically the same function, in series.
Topology 3 vCNS Based Private Cloud - MiVoice Business Express		

	Network component	Configuration considerations
	vCNS	<ul style="list-style-type: none"> • Configure two static 1:1 external to internal address mappings for MiVoice Business Express external (WAN) interface • The NAT/firewall function must preserve the TCP and UDP port numbers in packets exchanged between MiVoice Business Express and the external network; only the address fields may be changed. • The public addresses bound to MiVoice Business Express external facing addresses must be static IP addresses, visible from the external network (Internet). These should be separate addresses from the external IP address of the vCNS (or other upstream firewall). • If MiCollab AWV is part of the solution, then you must also configure a second 1:1 external to internal address mapping for the second MiVoice Business Express/MBG WAN address required for the MiCollab AWV Connection Point HTTPS. • Configure for the needs of any non-UCC applications (web access from user desktops, and any other applications requiring Internet access).

	Network component	Configuration considerations
	MiVoice Business Express - MBG Virtual component	<ul style="list-style-type: none"> Standard configuration, as built into MiVoice Business Express. Configure the public WAN address as the IP address of the vCNS public IP address (or as the public IP address of the external firewall, if one exists). If MiCollab AWV is part of the solution, then you must also configure a second public WAN IP address for the MiCollab AWV Connection Point HTTPS, as in previous topologies. Configuration of the MiCollab AWV LAN address is not required.
	IP Phones, other VoIP endpoints	<ul style="list-style-type: none"> All external endpoints are in Teleworker configuration. Internal endpoints must not be in Teleworker configuration.
	External firewall (if it exists)	<ul style="list-style-type: none"> Configure as per vCNS. It is basically the same function, in series.
Topology 4 vCNS Based Public Cloud		

	Network component	Configuration considerations
	vCNS per customer	<ul style="list-style-type: none"> • Configure static 1:1 external to internal address mapping for the MBG Virtual external (WAN) interface. • The NAT/firewall function must preserve the TCP and UDP port numbers in packets exchanged between the MiVoice Border Gateway and the external network; only the address field may be changed. • The public address bound to the MiVoice Border Gateway external facing address must be a static IP address, visible from the external network (Internet). This should be a separate address from the external IP address of the vCNS (or other upstream firewall). • If MiCollab AWV is part of the customer solution—requires external (WAN/Internet) access—then you must also configure a second 1:1 external to internal address mapping for the second MBG WAN address needed for MiCollab AWV Connection Point HTTPS. • Configure for the needs of any non-UCC applications and any applications requiring Internet access.

	Network component	Configuration considerations
	MiVoice Border Gateway per customer	<ul style="list-style-type: none"> • Standard Server-Gateway configuration. • Configure the public WAN address as the IP address of the per-customer vCNS public IP address (or as the public IP address of the external firewall, if present). • If MiCollab AWV is part of the customer solution, then you must also configure a second public WAN IP address for MiCollab AWV Connection Point HTTPS and configure LAN address of the MiCollab AWV, as in topologies described earlier in this guide.
	IP Phones, other VoIP endpoints	<ul style="list-style-type: none"> • All endpoints are in Teleworker configuration. • There are no host site internal endpoints.
	External firewall (if it exists)	<ul style="list-style-type: none"> • Configure as per vCNS, as described for Topology 2, and replicated across all customers. It is basically the same function, in series, one for each customer.
Topology 5 vCNS Based Public Cloud - MiVoice Business Express		

	Network component	Configuration considerations
	vCNSper customer	<ul style="list-style-type: none"> • Configure static 1:1 external to internal address mapping for the MBG Virtual external (WAN) interface. • The NAT/firewall function must preserve the TCP and UDP port numbers in packets exchanged between the MBG Virtual and the external network; only the address field may be changed. • The public address bound to the MBG Virtual external facing address must be a static IP address, visible from the external network (Internet). This should be a separate address from the external IP address of the vCNS (or other upstream firewall). • If MiCollab AWV is part of the customer solution—requires external (WAN/Internet) access—then you must also configure a second 1:1 external to internal address mapping for the second MBG WAN address needed for MiCollab AWV Connection Point HTTPS. • Configure for the needs of any non-UCC applications and any applications requiring Internet access.

	Network component	Configuration considerations
	MiVoice Business Express - MBG Virtual component	<ul style="list-style-type: none"> • Standard Server-Gateway configuration, as built in to MiVoice Business Express. • Configure the public WAN address as the IP address of the per-customer vCNS public IP address (or as the public IP address of the external firewall, if present). • If MiCollab AWV is part of the customer solution, then you must also configure a second public WAN IP address for MiCollab AWV Connection Point HTTPS, as in previous topologies. Configuration of the MiCollab AWV LAN address is not required.
	IP Phones, other VoIP endpoints	<ul style="list-style-type: none"> • All endpoints are in Teleworker configuration. • There are no host site internal endpoints.
	External firewall (if it exists)	<ul style="list-style-type: none"> • Configure as per vCNS, as described for Topology 2, and replicated across all customers. It is basically the same function, in series, one for each customer.

