A MITEL
PRODUCT
GUIDE

MiVoice MX-ONE

Virtualization and Public Cloud - Description

Release 7.5 SP1
74/1551-ANF 901 43 Uen H 2023-04-24

April 2023

Mitel
Powering connections

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) or its 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.

Contents

1 Introduction....1
1.1 Scope of this Document....1
1.2 MiVoice MX-ONE Virtualization....1

2 Benefits of Virtualization....3

3 Solution Description....4
3.1 MiVoice MX-ONE Support- VMware....5
3.1.1 MX-ONE Service Node....5
3.1.2 MX-ONE Media Server....5
3.1.3 Database....6
3.1.4 Management System....7
3.1.5 Introduction of the Solution....8
3.1.6 VMware Validated Scenarios....10
3.2 MiVoice MX-ONE Support - Hyper-V....21
3.2.1 Basic Hyper-V....21
3.2.2 Failover Clustering....22
3.3 MiVoice MX-ONE Support - KVM and Red Hat Enterprise....22

4 MX-ONE Requirements....23
4.1 System Requirements for Virtualized Infrastructure....23
4.2 Requirements - VMware....23
4.3 MX-ONE Service Node Resources Requirement....25
4.3.1 Basic Guidelines....25
4.3.2 MX-ONE Guest Machine Requirements - General....26
4.4 Requirements - VMware....27
4.4.1 MX-ONE Guest Machine Requirements - VMware....29
4.4.2 Media Gateway Chassis Distribution....37
4.4.3 Media Server....38
4.4.4 Provisioning Manager....40
4.4.5 Upgrade Process....40
4.4.6 MX-ONE Service Node Software....40
4.4.7 MX-ONE Media Server....40
4.4.8 Creating Snapshot on a VM....40
4.5 Requirements - Hyper-V and KVM....40
4.5.1 Service Node with MX-ONE Applications....41
4.5.2 Media Server Standalone....42
4.5.3 Provisioning Manager Standalone....43
4.5.4 Database Standalone or Large Co-Located Databases....44
4.6 Recommended Training....45
4.7 Product Business Approval (PBA)....45

4.8 Upgrading MX-ONE in KVM to Support Minimum Downtime....46

5 Public Cloud....49
5.1 MiVoice MX-ONE Support - Microsoft Azure....49
5.2 Requirements - Microsoft Azure....50
5.2.1 MX-ONE Server requirements....50
5.2.2 Connectivity between on-premises and Azure....51
5.2.3 Availability - MX-ONE Redundancy....52
5.2.4 Gateway Calls - Forced Gateway....53
5.2.5 Security Requirements....53
5.2.6 Mitel Applications....54

6 Acronyms, Abbreviations, and Glossary....55

7 Reference Documents....57

Introduction 1

This chapter contains the following sections:

- Scope of this Document
- MiVoice MX-ONE Virtualization

This guide describes the setup configurations that apply when setting up Mitel MiVoice MX-ONE 7.x on servers enabled with:

- VMware® vSphere™ virtualization
- Microsoft® Hyper-V™ virtualization
- KVM (Kernel-based Virtual Machine)

1.1 Scope of this Document

The aim of this document is to provide a description of MX-ONE virtualization solution explaining the supported scenarios and their requirements as well as benefits and limitations. Although Mitel UC applications, such as Mitel CMG, MiContact Center Enterprise and Mitel Advanced Messaging, included as part of the MiVoice MX-ONE 7.x solution, have also been validated in a virtualized environment, they are not covered in this document. Please refer to relevant documents for these applications where virtualization options are described.

1.2 MiVoice MX-ONE Virtualization

MiVoice MX-ONE is composed mainly by a piece of software called Service Node that is responsible for call control and a gateway that is responsible by media transcoding. The gateways can be software only called Media Server or a dedicated hardware/software called MGU – Media Gateway Unit.

MX-ONE Service Node software as well as Media Server was originally designed to run on top of a standard physical server (traditional x86 architecture), meaning server hardware was dedicated to each server in a MX-ONE.

Figure 1: Traditional server x86 architecture example

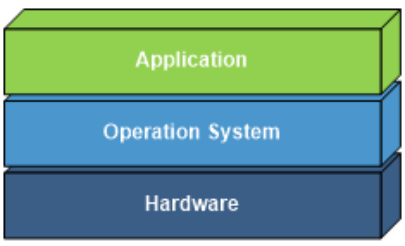

More and more virtualization is part of the companies ITs and since MX-ONE 5.0 Mitel has validated MX-ONE software to run on top of a virtualized environment.

With the introduction of MiVoice MX-ONE 7.x, the MX-ONE main components Service Node, Database, Management system (Provisioning Manager) well as Media Server software has been re-validated in a virtualized environment, where the physical server hardware can be shared between several Service Node call managers.

The Database, Management system and Media Server was validated in a virtualized environment being co-located in the same virtual machine as MX-ONE Service Node and in standalone virtual machine.

In a virtualization environment an abstraction layer, called a hypervisor, is installed between the physical hardware (host server/machine) and the operating system (guest machine). This abstraction layer allows several guest machines to co-exist on the same physical hardware to share resources like memory, CPU, etc. These guest machines are commonly referred to as Virtual Machines (VM).

Figure 2: Virtual architecture example

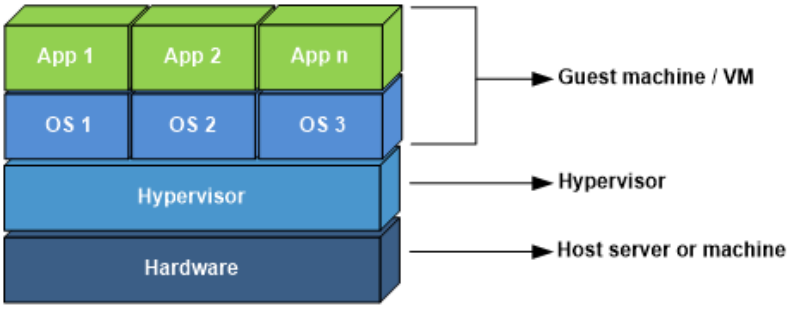

This means that multiple MX-ONE Service Node call managers, Database servers, Media Server, applications or guest machines can reside in the same physical host. The consolidation of server hardware through virtualization reduces the server footprint offering lower power consumption and cooling requirements as well as decreased physical space requirements. Virtualization also offers the possibility to take advantage of high availability options that can provide increased resiliency for real time applications, such as MX-ONE.

Currently, MiVoice MX-ONE relies on the following virtual machine software:

- VMware software as a part of its virtualization solution. Mitel has verified MX-ONE 7.x running as virtual machines in a VMware 6.5 or later infrastructure.
- Microsoft Hyper-V support based on Windows Server 2016 and Windows Server 2019 as part of its virtualization solution. Mitel has verified MX-ONE 7.1, and later, running as virtual machines in a Windows Server 2016 and Windows Server 2019 hypervisor infrastructure.
- KVM virtual machine based on SUSE Linux Enterprise Server (SLES) 12 SP5 or later.
- KVM virtual machine based on Red Hat Enterprise Linux 7.6 or later.

Benefits of Virtualization 2

IT departments are more and more looking for alternatives to reduce the total cost of ownership and increase productivity. However, in the last two decades, x86 architecture based servers have increased their numbers in the IT segment. Indeed, in a traditional x86 architecture, only one or few applications share the same hardware. This meant the multiplication of servers which were each dedicated to a different back office applications (Mail servers, Web servers, Databases, CRM, etc.) to ensure enough resources to handle peak traffic. In most of the cases these applications do not use all the power of the server hardware, which as a consequence causes waste of CPU, memory, etc. Additional servers require more power, cooling, physical space, maintenance, etc. that raises the TCO in data centers.

Virtualization software companies address such issues by offering possible ways to optimize server efficiency. For instance, this can be achieved by sharing hardware resources between virtual machines (guest machine).

In general advantages of virtualization are:

- **IT Hardware Consolidation**
  - Improve the efficiency and availability of IT resources
  - Reduce capital costs
    - Reduce expenditure on physical servers
    - Maintenance and Hardware savings
- **Infrastructure savings**
  - Lower power consumption
  - Reduced cooling requirements
  - Decrease physical space requirements
- **Increase application reliability**
  - Potentially more availability options available
  - Business continuity options

There are also some disadvantages, like:

- Network complexity
- Hardware costs
- Software costs

Solution Description 3

This chapter contains the following sections:

- MiVoice MX-ONE Support- VMware
- MiVoice MX-ONE Support - Hyper-V
- MiVoice MX-ONE Support - KVM and Red Hat Enterprise

MX-ONE Service Node 7.0 and later, including Database servers, Management Servers and MX-ONE Media Server can run in a hardware virtualization environment, formally named Virtualization.

Mitel has verified MX-ONE Service Node 7.x call manager as well as MX-ONE Media Server software running as virtual machines in a VMware 6.5 or later infrastructure.

Note:
Mitel does not supply any VMware software together with MiVoice MX-ONE. The VMware software related parts described in this document must be designed and implemented by a VCP (VMware Certified Professional) certified engineer via the VMware partner certification program.

The solution described in this document runs on top of the following infrastructures:

- A VMware infrastructure and therefore all VMware´s requirements for such a scenario must be in place before beginning the MX-ONE installation in a virtualized environment.
- A Microsoft Windows Server 2016 and Windows Server 2019 infrastructure and therefore all Microsoft´s requirements for such a scenario must be in place before beginning the MX-ONE installation in a virtualized environment.
- A KVM infrastructure and therefore all KVM´s requirements for such a scenario must be in place before beginning the MX-ONE installation in a virtualized environment.

It is assumed:

- For a VMware setup, that a VCP engineer and a network engineer do the design and implementation of the customer project.
- For a Hyper-V setup, that a Microsoft Hyper-V certified engineer and a network engineer do the design and implementation of the customer project.
- For KVM setup, an engineer with good knowledge of KVM architecture and a network engineer do the design and implementation of the customer project.

That the network is properly designed and configured to provide connectivity to the server infrastructure and data storage as well as handle the traffic generated by the applications that runs on it.

When availability is required, it is assumed:

That a high availability network is properly designed and configured to provide reliability, scalability and security.

3.1 MiVoice MX-ONE Support- VMware

This topic describes how you can setup the MX-ONE system in a VMware virtual environment.

3.1.1 MX-ONE Service Node

Mitel has validated three VMware software options to run MX-ONE Service Node in a virtualized infrastructure.

The three options are:

- Consolidated setup
- Availability setups with VMware vSphere Availability
  - VMware vSphere High Availability setup
  - VMware vSphere Fault Tolerance setup

Note:
VMware´s High Availability (HA) and Fault Tolerance (FT) options require a specific network as well as storage infrastructure to be in place.

Note:
The network and storage implementation are not covered in this description, as they are a pre-requisite for implementing VMware´s HA and FT infrastructure, which is described in standard VMware documentation.

Before offering MX-ONE Service Node virtualization solution, please read carefully the VMware´s documentation.

https://docs.vmware.com/en/VMware-vSphere/

Please, always check the latest products' documentation.

3.1.2 MX-ONE Media Server

MX-ONE Media Server is part of the MiVoice MX-ONE softswitch concept, which means that no hardware is required to run MiVoice MX-ONE. The typical scenario for softswitch is a pure SIP solution. MiVoice MX-ONE softswitch is composed by at least one MX-ONE Service Node and one MX-ONE Media Server.

Mitel has validated two VMware software options to run MX-ONE Media Server in a virtualized infrastructure.

These two options are:

- Consolidated setup
- Availability setups with VMware vSphere Availability
    - VMware vSphere High Availability setup

Media Server was validated in co-located as well as standalone deployments. In a co-located deployment MX-ONE Service Node and Media Server are installed in the same virtual machine. On the other hand, in a standalone deployment, the Media Server is installed in a separate virtual machine.

Note:
VMware´s High Availability (HA) option requires a specific network as well as storage infrastructure to be in place.

Note:
The network and storage implementation are not covered in this description, as they are a pre-requisite for implementing VMware´s HA infrastructure, which is described in standard VMware documentation.

Before offering MX-ONE Media Server virtualization solution, please read carefully at least the following VMware´s reference document:

Voice over IP (VoIP) Performance Evaluation on VMware vSphere® 5

http://www.vmware.com/files/pdf/techpaper/voip-perf-vsphere5.pdf

Please, always check the latest products documentation.

3.1.3 Database

From MX-ONE 7.0 onwards, a new type of server is introduced, the database server, which Cassandra database runs.

Mitel has validated two VMware software options to run the database in a virtualized infrastructure.

These two options are:

- Consolidated setup
- Availability setups with VMware vSphere Availability
    - VMware vSphere High Availability setup

Database was validated in co-located as well as standalone deployments. In a co-located deployment MX-ONE component e.g. Service Node, Database, Management and Media Server are installed in the same virtual machine. On the other hand, in a standalone deployment, the database is installed in a separate virtual machine.

Note:
VMware´s High Availability (HA) option requires a specific network as well as storage infrastructure to be in place.

Note:
The network and storage implementation are not covered in this description, as they are a pre-requisite for implementing VMware´s HA infrastructure, which is described in standard VMware documentation.

3.1.4 Management System

The MX-ONE Management system, Provisioning Manager and Service Node Manager can also run in a virtualized environment.

Mitel has validated two VMware software options to run the management system in a virtualized infrastructure.

These two options are:

- Consolidated setup
- Availability setups with VMware vSphere Availability
    - VMware vSphere High Availability setup

Provisioning Manager was validated in co-located as well as standalone deployments. In a co-located deployment MX-ONE component e.g. Service Node, Database, Management and Media Server are installed in the same virtual machine. On the other hand, in a standalone deployment, the Provisioning Manager is installed in a separate virtual machine.

Note:
VMware´s High Availability (HA) option requires a specific network as well as storage infrastructure to be in place.

Note:
The network and storage implementation are not covered in this description, as they are a pre-requisite for implementing VMware´s HA infrastructure, which is described in standard VMware documentation.

3.1.5 Introduction of the Solution

In the next sessions some basic concepts and requirements of a virtualized environment are introduced.

Virtual Machine

First of all let us define a virtual machine according to the VMware documentation: “A virtual machine is a tightly isolated software container that can run its own operating systems and applications as if it were a physical computer. A virtual machine behaves exactly like a physical computer and contains its own virtual (i.e. software-based) CPU, RAM, hard disk and network interface card (NIC). An operating system cannot tell the difference between a virtual machine and a physical machine, nor can applications or other computers on a network. Even the virtual machine thinks it is a “real” computer. Nevertheless, a virtual machine is composed entirely of software and contains no hardware components whatsoever. As a result, virtual machines offer a number of distinct advantages over physical hardware”.

Network Infrastructure

A reliable and secure network infrastructure supporting standard protocols is required to provide server and storage connectivity, refer to VMware´s documentation for such requirements.

Furthermore, avoid any single point of failure in the network when designing a high available system.

Virtual infrastructure

VMware Availability software requires a virtual infrastructure that is composed by virtual network and virtual switches. There are two types of virtual switches, vNetwork standard switches and vNetwork distributed switches. Which one to be used depends of customer environment, because every customer has different needs as well as customers’ infrastructures are different and customers/VMware partners have to take that into account while designing the virtual network.

In Mitel´s validation setup, five traffic cases were defined in order to test VMware´s High Availability and Fault Tolerance.

- Management – that is used by vCenter Server
- VMotion – that is used by VMotion
- FT – Logging that is used by Fault Tolerance
- iSCSI – that is used by the IP storage
- Production – that is used by MX-ONE Service Node Virtual Machines

Below is a virtual switch example with the different traffic types that are required by VMware´s software when running availability. In a customer network more traffic types can exist.

Figure 3: Virtual switch example

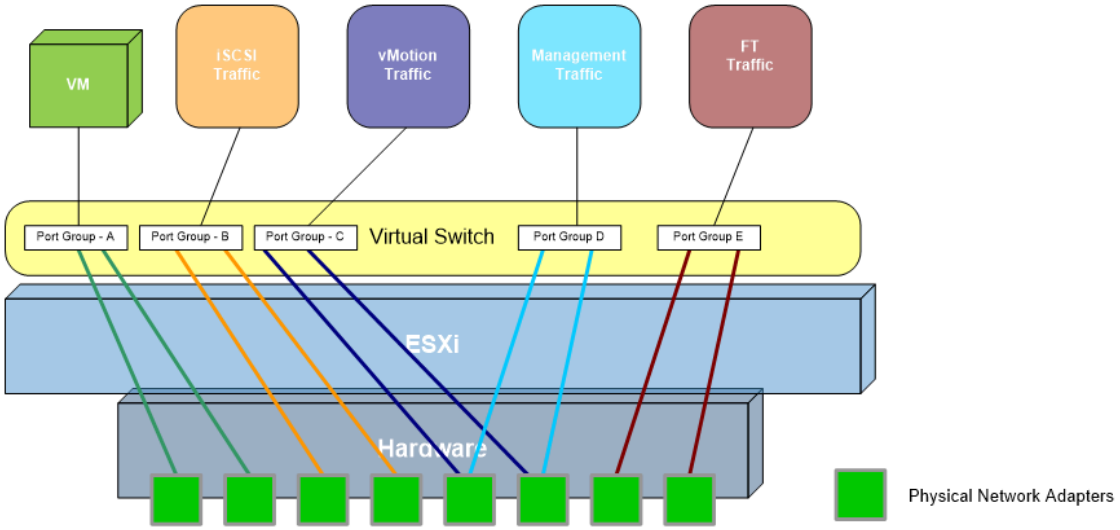

With that definition in mind, this document will now describe the MX-ONE Service Node 7.x virtualized solution that was validated by Mitel.

Figure 4: MX-ONE Service Node virtualized example

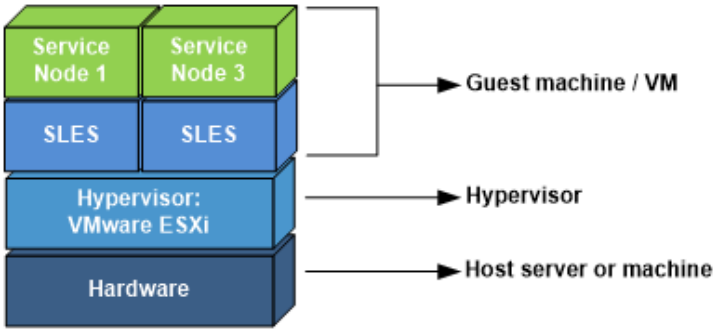

The MX-ONE Service Node virtualized solution is basically composed by a guest machine with Linux SLES and MX-ONE Service Node 7.x software running on top of a host machine with VMware´s vSphere ESXi 6.7 hypervisor.

In this approach, the main advantage of the virtualization from a customer perspective is the fact that it allows the possibility that more than one MX-ONE server runs on a single physical server with the appropriate hardware capacity and configuration that can reduce hardware footprint as well as maintenance costs. However, it is important to mention that some caution is needed when doing a virtualization project due to the fact that MX-ONE is a real-time communication system and it is always a good practice to have a certain level of redundancy.

Note:
For customers running multiple Service Nodes in a virtualized environment, Mitel strongly recommends that the Service Node virtual machines be spread evenly across at least two physical hosts in order to avoid that a host failure jeopardizes the whole communication system.

Figure 5: High level MiVoice MX-ONE virtualized deployment

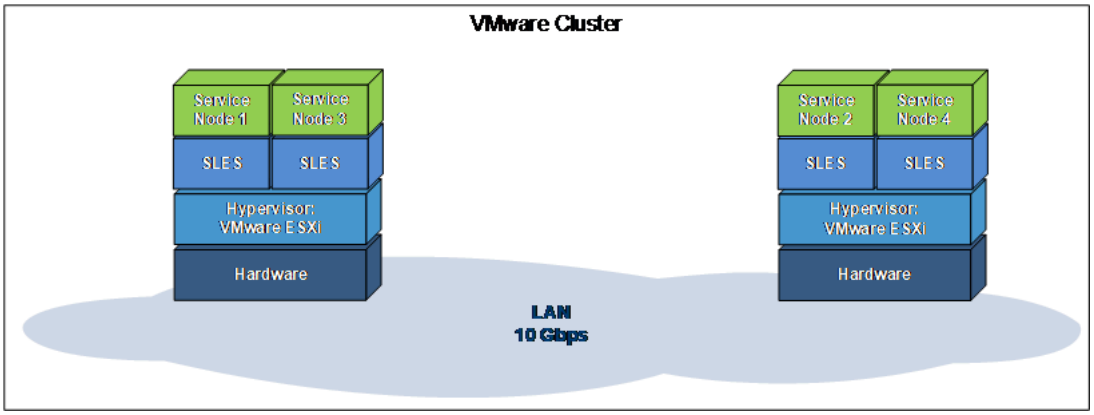

3.1.6 VMware Validated Scenarios

Mitel validated three different setups for virtualization on top of VMware's software for MX-ONE 7.0 and later releases. These setups are:

- Consolidated setup where several Service Nodes are consolidated in one or more physical servers. In this case no redundancy is provided by VMware's software.
- Availability setups where several physical servers are consolidated in one or more physical servers. In these solutions, redundancy is provided by VMware´s software. This option offers two possible setups:
  - VMware vSphere High Availability
  - VMware vSphere Fault Tolerance (only for Service Node)

3.1.6.1 Consolidated Setup

Mitel validated a consolidation setup where a guest machine running Linux SLES 12 SP5 and where MX-ONE Service Node runs on top of a host machine running VMware´s vSphere ESXi 6.7 hypervisor.

In such a scenario a single physical server or several servers with the appropriate hardware capacity and configuration can be used to reduce the number of physical servers that would be needed in traditional server architecture. In fact, it can reduce hardware footprint as well as maintenance costs.

MX-ONE Service Node, database, media server and management system can be deployed in a consolidated setup.

The figure below shows an example how it can be achieved. The standard MX-ONE is composed by four Service Nodes each one running on top of a physical server and four MX-ONE Lite media gateways using MGU.

Figure 6: MX-ONE with four Service Nodes

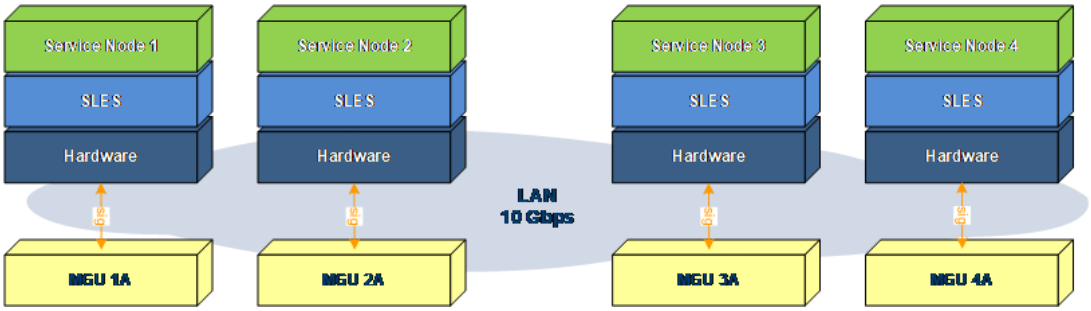

When virtualization is deployed in this setup it can, for example, be reduced to two physical servers, each with two Service Nodes.

Figure 7: MX-ONE consolidation example

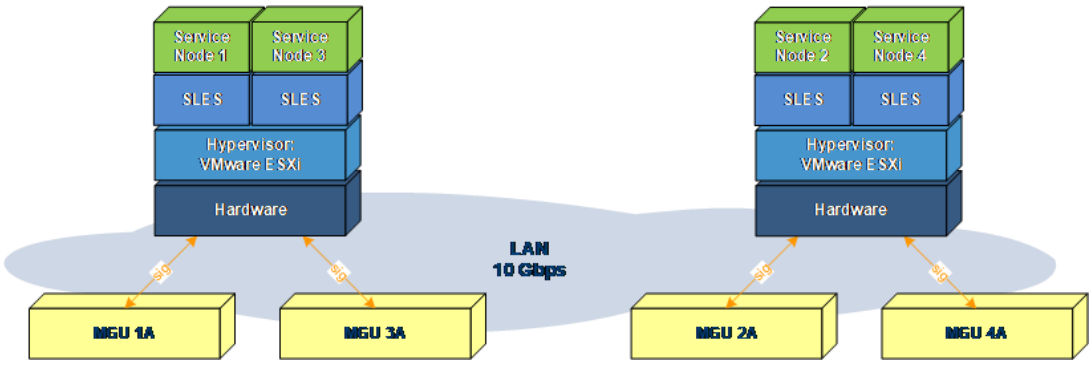

The disadvantage is that in the case one physical host fails, there will be two Service Node nodes out of order. Having multiple Service Nodes on one host defeats the purpose of the distributed architecture. If all Service Nodes would be hosted in one physical server, this would be a single point of failure, and the entire system would be compromised.

As this solution does not implement any type of redundancy, such a configuration is not recommended to be used for multi Service Node systems. However, if a resilient setup is established, e.g. hardware host (dual power supply, RAID, etc.) and IP network (switches, paths, etc.) where components are redundant, the risk can be minimized.

3.1.6.1.1 Requirements for Consolidation Setup

The consolidation setup requires VMware vSphere Hypervisor: ESXi 6.7 or later.

Although not mandatory in this setup, VMware´s vCenter management software should be used to manage this solution.

Hardware hosts and IP network should have a certain level of redundancy.

Since network setup needs to follow VMware´s recommendations, please check the document below as a reference:

vSphere Networking

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.networking.doc/GUID-35B40B0B-0C13-43B2-BC85-18C9C91BE2D4.html

3.1.6.2 Availability Setups with VMware

Mitel has also verified some more advanced VMware software options in order to give customers the possibility to choose the best availability option for their IT environment. The required software is VMware vSphere High Availability (HA) and VMware vSphere Fault Tolerance (FT). VMware vCenter management software is required to manage these solutions.

For the purposes of VMware HA and FT validation with MX-ONE Service Node, a minimum of four physical networks and a SAN infrastructure were required. Additionally, to ensure network redundancy in a production environment, secondary LANs should be added to the production, storage and management networks.

Note:
VMware Fault Tolerance requires a 10 Gb Ethernet network.

Please read carefully VMware´s documentation for more detailed information.

Best Practices for Networking

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-B1906BCD-E538-4FFF-AAE9-5403FE253F38.html

The picture below shows an example of traffic distribution per network adapter.

Figure 8: Example: Availability – traffic type distributed by 10 Gigabit Ethernet network adapters

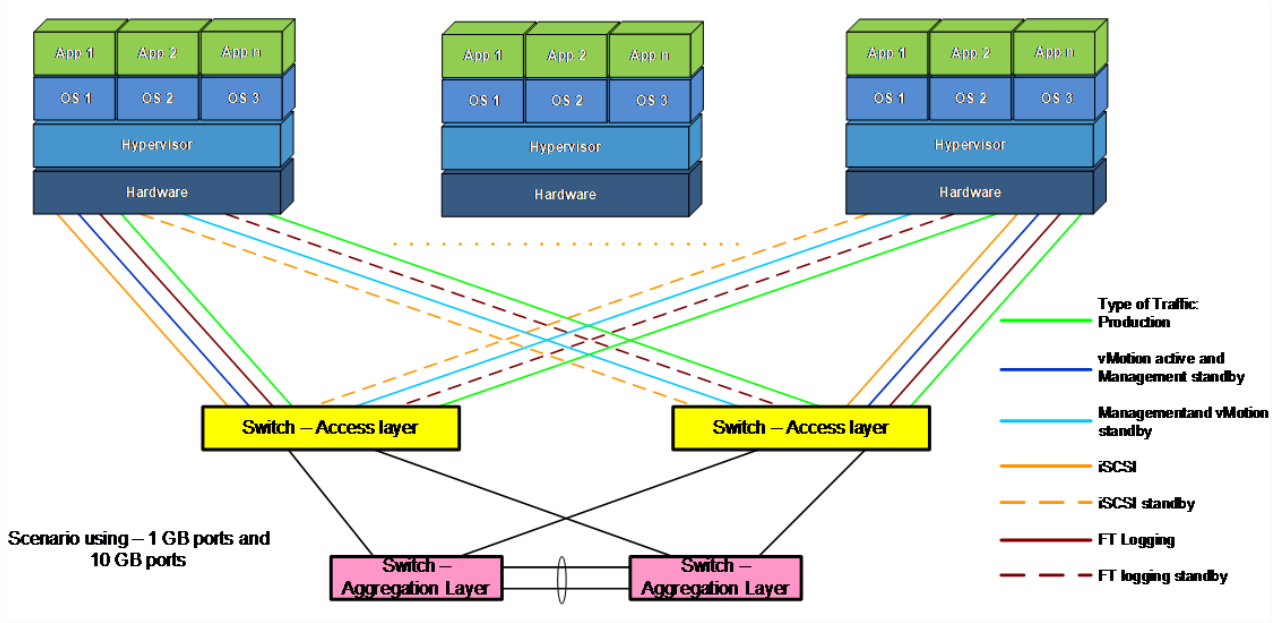

VMware High Availability or Fault Tolerance?

There are differences in performance and availability when running VMware vSphere High Availability and VMware vSphere Fault Tolerance. Indeed, a choice between performance and availability needs to be taken in consideration when doing a system project.

Note:
VMware vSphere High Availability provides better VM performance, but warm standby level of availability. However, VMware vSphere Fault Tolerance provides “five 9s” level of availability, but lower VM performance, due the fact that some VMware limitations are imposed for such a setup.

Please read carefully VMware´s documentation for more detailed information.

About vSphere Availability

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-63F459B7-8884-4818-8872-C9753B2E0215.html

3.1.6.2.1 VMware High Availability Setup

MX-ONE running on top of a VMware vSphere HA infrastructure can be protected from hardware failures as well as benefit of planned hardware maintenance. According to the VMware documentation: “vSphere HA provides high availability for applications running in virtual machines. In the event of a server failure, affected virtual machines are automatically restarted on other productions servers with enough capacity to run them.”

MX-ONE (Service Node, database, media server and management) guest machines running on top of VMware vSphere High Availability allows a cold standby solution, which means in the event that a physical server where the Service Node guest machine is running goes down, a short downtime will occur, due to the fact that Service Node guest machine needs to be initiated in another physical server.

The VMware High Availability setup can be compared with MX-ONE Server Redundancy. However, the VMware High Availability replaces the need for a Service Node redundancy configuration in MX-ONE. Once a customer has the VMware HA option in place, the MX-ONE will be setup as a standard system in the VMware HA cluster without any MX-ONE server redundancy option enabled.

The following figure shows a standard MX-ONE composed of four Service Nodes guest machines running on top of VMware´s HA infrastructure.

Figure 9: Four Service Nodes running on VMware HA

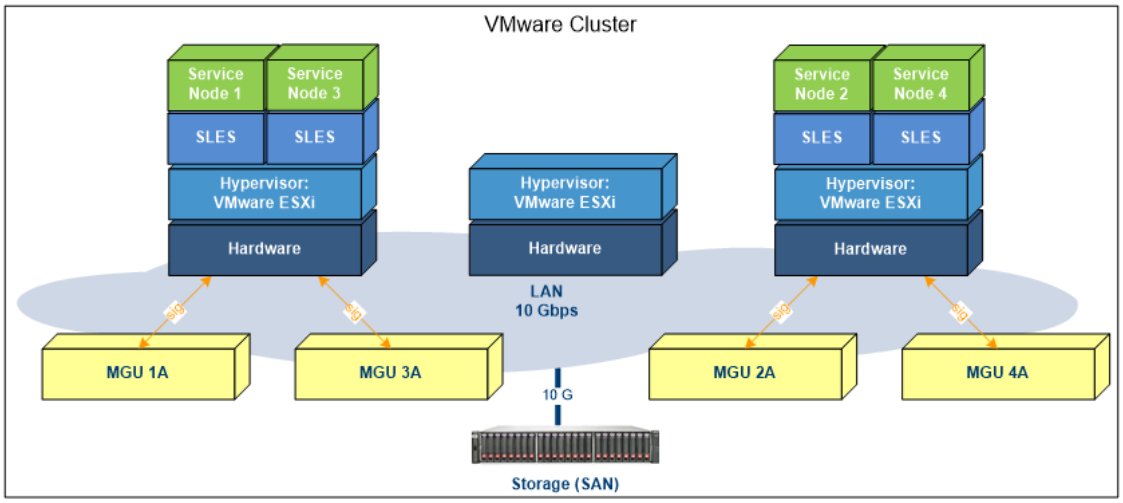

Note:
VMware High Availability is enabled in a cluster, spare host resources, such as memory and CPU, are required in order to fulfill the failover requirements. As VMware requires that vSphere HA operates in a single physical datacenter, Mitel's validation of HA has been carried out in this environment.

In case of hardware failure, the Service Nodes guest machines running on top of the failed machine goes down.

Figure 10: VMware HA physical server failure example

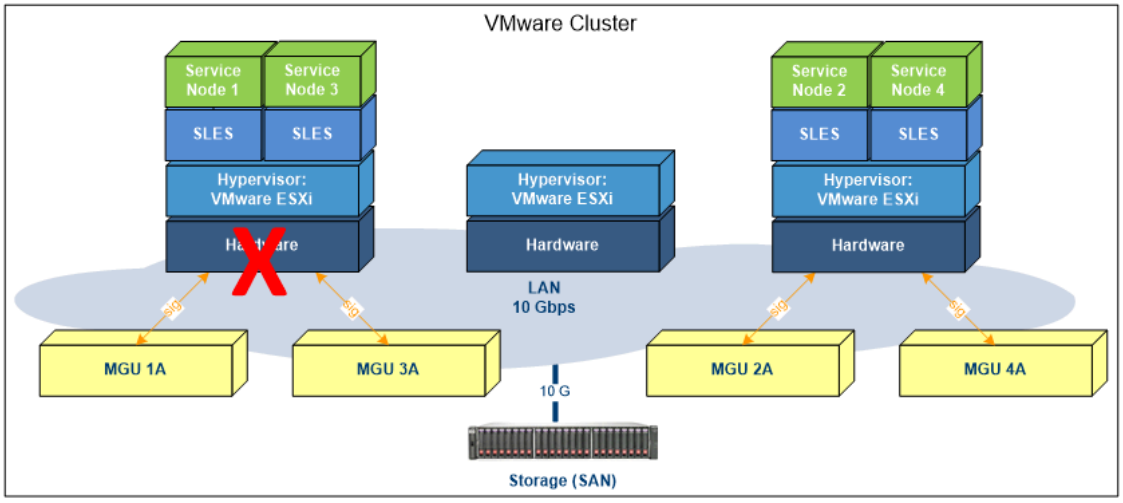

The VMware HA mechanism will initialize these Service Node guest machines in another available machine in the VMware cluster environment.

In the example below a spare machine in the cluster with enough spare resources (memory, CPU, etc.) was used to reallocate the two Service Node guest machines.

Figure 11: VMware HA recovery example

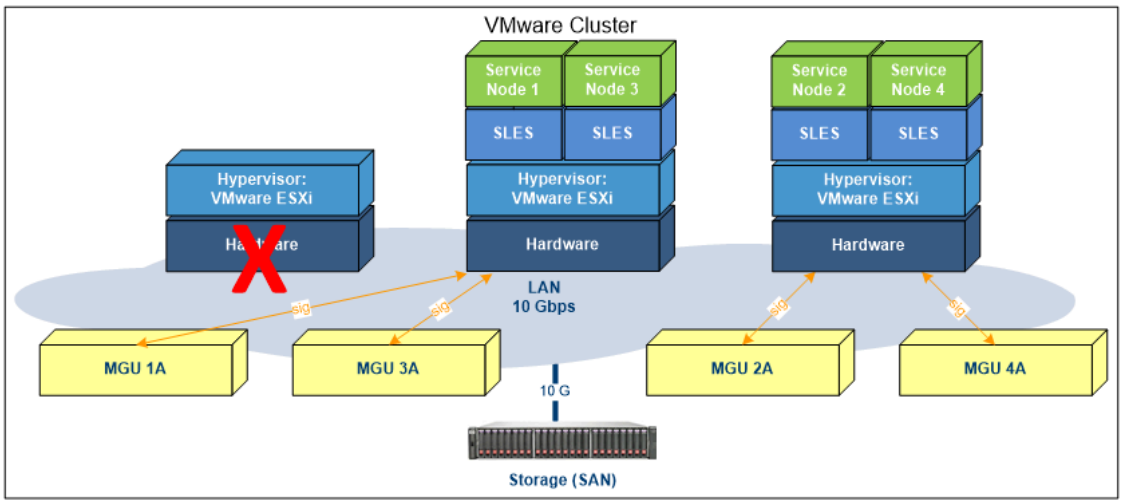

During the re-initialization process, ongoing media gateway calls will be dropped, whereas ongoing direct media calls will be maintained. This is similar to the server redundancy option offered with MX-ONE. Additionally, as the VMware HA cluster environment requires a SAN setup, the customer data and management continuity are maintained.

Requirements for VMware High Availability setup

- SAN (Storage Area Network) and Network requirements for High Availability according to VMware specifications
- VMware vSphere, Hypervisor: ESXi 6.7
- VMware vCenter
- VMware VMotion

- VMware High Availability

It should be noticed that a SAN environment and multiple LAN segments are required by VMware in order for this option to be deployed. Refer to the latest VMware packaging options to determine the VMware software editions that best fit the requirements.

Mitel strongly recommends that partners/customers always check the latest High Availability requirements with a qualified VMware technical representative.

It is also recommended that partner/customers read the latest versions of the following VMware´s documents that can be found on VMware website:

vSphere Availability ESXi 6.7

https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.avail.doc/GUID-63F459B7-8884-4818-8872-C9753B2E0215.html

vSphere High Availability Deployment Best Practices

http://www.vmware.com/resources/techresources/10232

3.1.6.2.2 VMware Fault Tolerance setup

MX-ONE Service Node running on top of a VMware vSphere HA and FT infrastructure can have a higher level of business continuity than VMware vSphere HA, according to the VMware documentation: “FT provides a higher level of availability, allowing users to protect any virtual machine from a host failure with no loss of data, transactions, or connections. FT provides zero downtime, zero data loss, and continuous availability for your applications”.

When MX-ONE Service Node guest machines are running in a VMware vSphere Fault Tolerance cluster, a transparent failover solution can be achieved. This means that in the event that a physical server where a Service Node guest machine is running goes down, no calls will be dropped during the failover process and continuity will be maintained. This transparent failover is possible, according to VMware, because Fault Tolerance uses VMware vLockstep technology, which guarantees the primary and secondary VMs execute exactly the same x86 instruction sequences. Fault Tolerance requires that the hosts CPUs are compatible with vLockstep technology, which requires additional physical processor extensions.

That is to say when a MX-ONE Service Node guest machine is running in VMware vSphere Fault Tolerance infrastructure, there is an additional Service Node guest machine running in parallel on a different physical server executing the same instructions. In this scenario, the two guest machines are synchronized or mirrored, so that in the case of a primary host server failure, the second parallel virtual machine takes over and becomes the new primary.

The figure below shows a standard MX-ONE composed of two Service Node guest machines running on top of VMware´s FT infrastructure.

Figure 12: Two Service Nodes protected by VMware FT

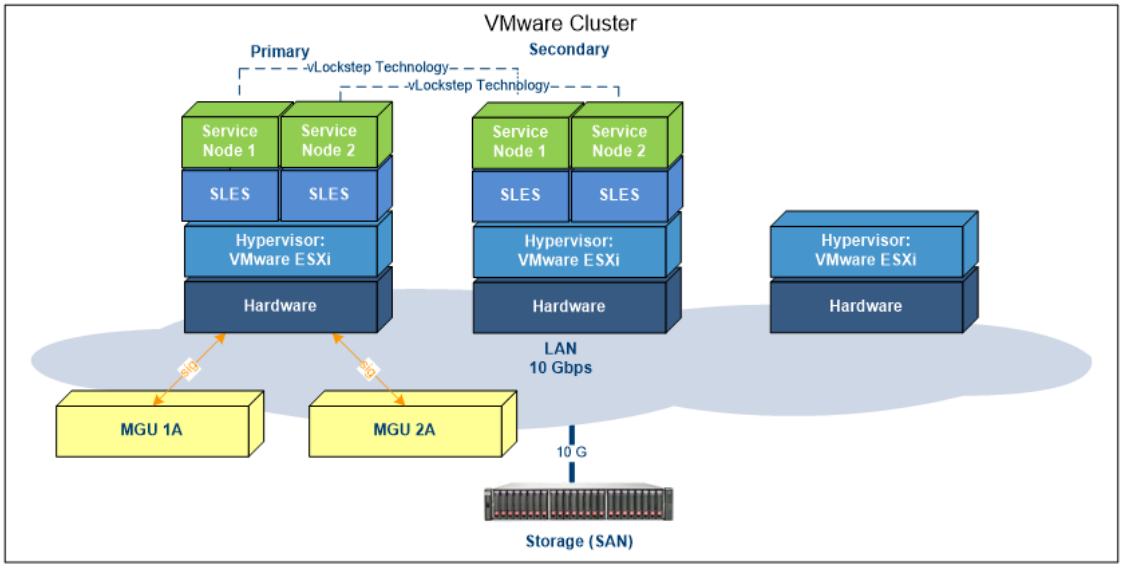

In normal operation, the signaling between Service Node and MGU is transmitted via the “primary physical server”.

To describe this process briefly, when a failure in the primary server hardware occurs, VMware´s mechanism in the secondary server will immediately detect it, take over and start to process pending I/O operations.

Figure 13: VMware FT physical server failure example

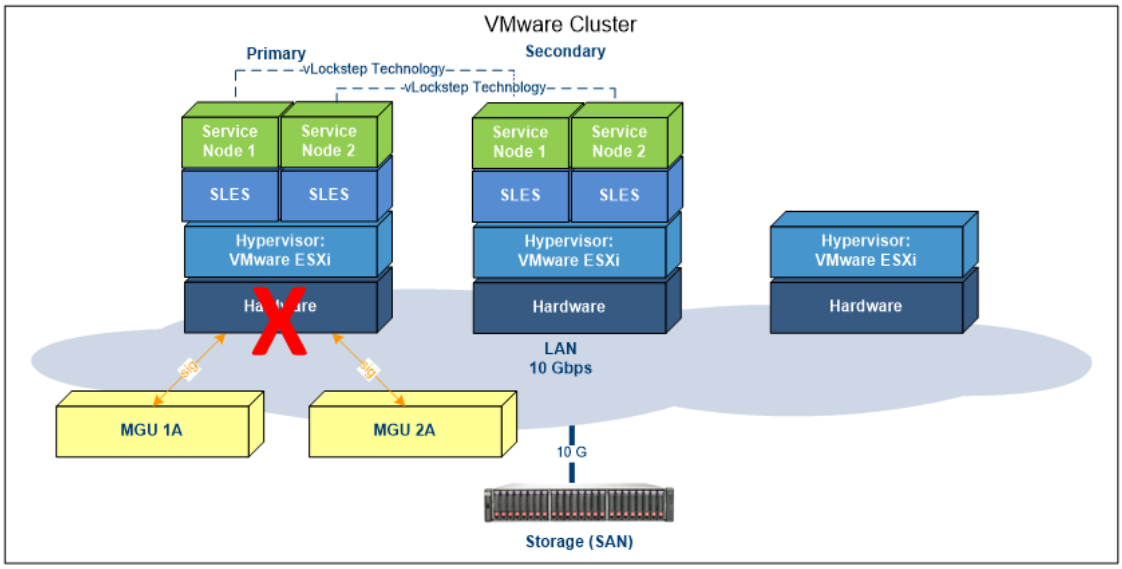

The secondary machine will then perform a “go live” operation and become the new primary server.

Figure 14: VMware FT transparent failover example

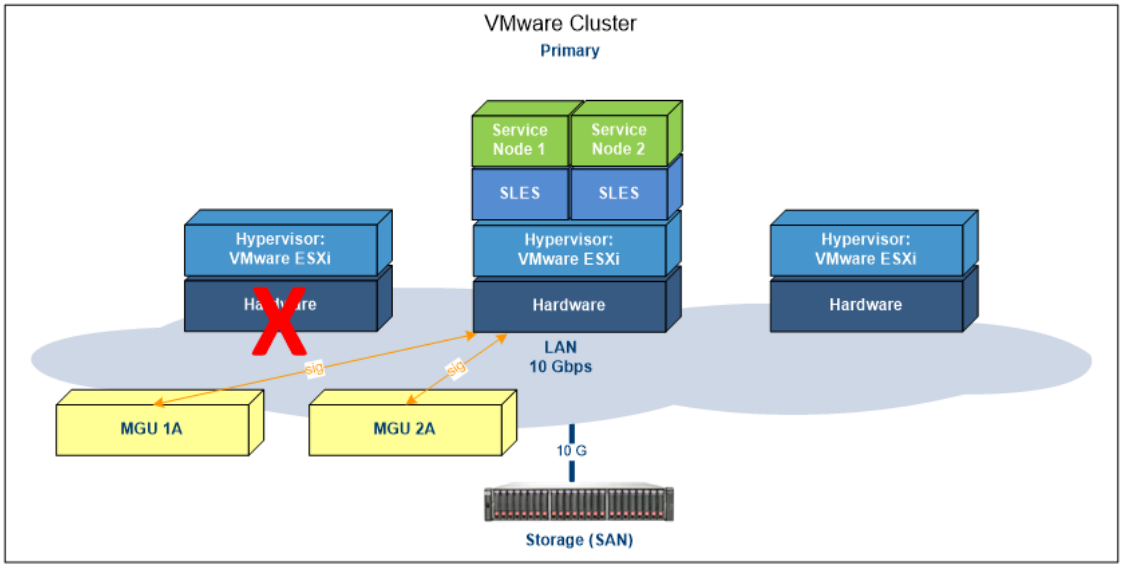

As part of the recovery process, after the secondary machine successfully takes over, VMware´s HA algorithm selects a new host machine in the cluster that is working properly and has available resources to create a new secondary host machine. This is to ensure that the system is protected again in case of a new hardware failure.

Figure 15: VMware FT transparent failover example

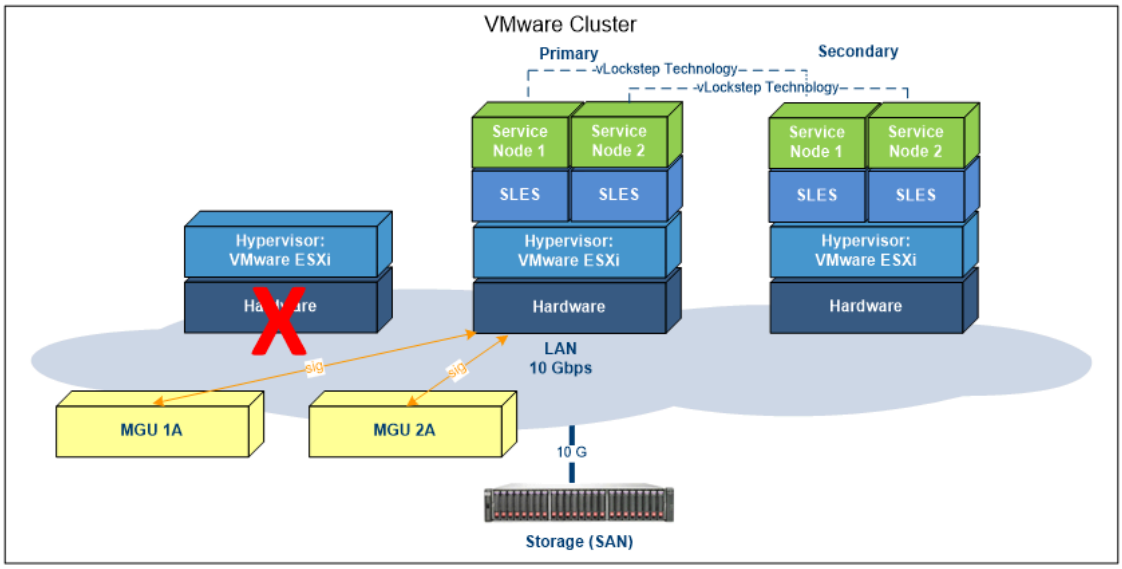

This last process is optional, although highly recommended. It is, of course, possible to limit the cluster to a primary and secondary host. Although, for added security it is always better to have three or more hosts in a cluster to re-establish a full active-active situation within minutes of the initial failure of a primary host.

From the MX-ONE Service Node call manager perspective no server failure is detected. The failover process is handled by VMware´s Fault Tolerance, where the signaling between MX-ONE Service Node and its associated MGUs continue to work normally, although the call processing is maintained by the “secondary physical server” instead. From an end user perspective, the MX-ONE will continue to work normally and ongoing calls and feature requests are maintained. Even from a management continuity perspective there is no loss of functionality, as the database is in a SAN environment, which is shared by

the primary and secondary server. This functionality is similar to a hot standby solution. Such a solution can be used by mission critical customers if they desire a more reliable system.

Requirements for VMware Fault Tolerance setup:

- SAN (Storage Area Networks) and Network requirements for High Availability and Fault Tolerance according to VMware specifications
- VMware vSphere, Hypervisor: ESXi 6.7
- VMware vCenter
- VMware VMotion
- VMware High Availability
- VMware Fault Tolerance

It should be noted that a SAN environment and multiple LAN segments are required by VMware in order for this option to be deployed. Refer to the latest VMware packaging options to determine the VMware software editions that best fit the requirements. The VMware vSphere enterprise editions should include HA/FT and VMotion. VMware vCenter is usually ordered separately and required to set this environment in place.

Mitel strongly recommends that partners/customers always check the latest High Availability/FT requirements with a qualified VMware technical representative.

Note:
VMware Fault Tolerance has some limitations/recommendations that need to be taken in consideration before implementation of such a solution. As VMware requires that vSphere HA/FT operates in a single physical datacenter, Mitel's validation of HA/FT has been carried out in this environment.

Mitel strongly recommends that the partner/customers read the following VMware´s documents:

About vSphere Availability

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-63F459B7-8884-4818-8872-C9753B2E0215.html

Providing Fault Tolerance for Virtual Machines

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-7525F8DD-9B8F-4089-B020-BAA4AC6509D2.html

Fault Tolerance Requirements, Limits, and Licensing

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.avail.doc/GUID-57929CF0-DA9B-407A-BF2E-E7B72708D825.html

3.1.6.3 Maintenance using VMware VMotion

VMware VMotion technology can be used to do maintenance in a physical server avoiding an unexpected failure. It allows that a virtual machine is moved from one server to another without interruption.

To be able to use VMotion, enough spare resources shall be available in the VMware cluster as well the required licenses.

The figures below show the VMotion process.

First, the Virtual Servers need to be moved to another physical machine with enough resources to handle the MX-ONE servers.

Figure 16: VMware VMotion example

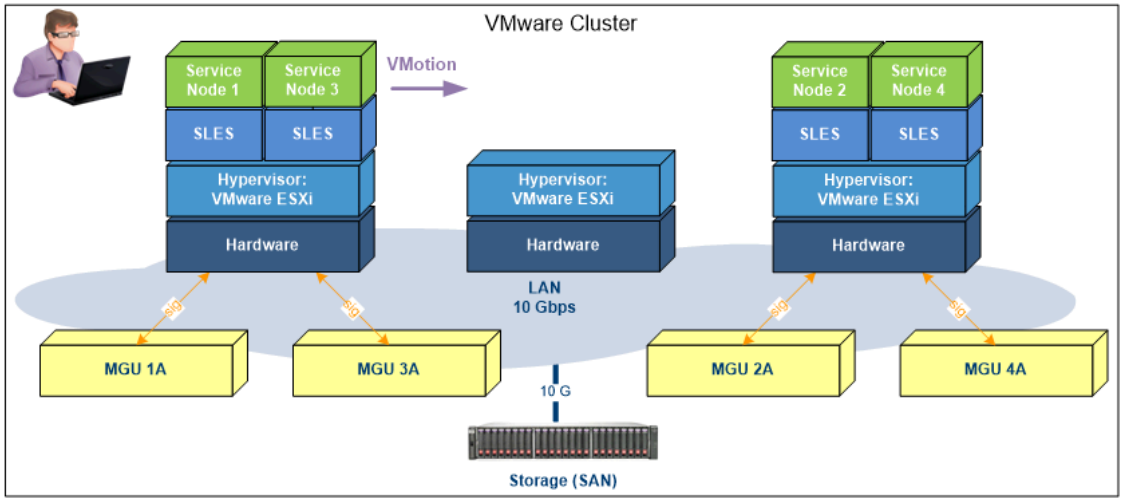

After the migration, the physical server can be repaired.

Figure 17: VMware VMotion example

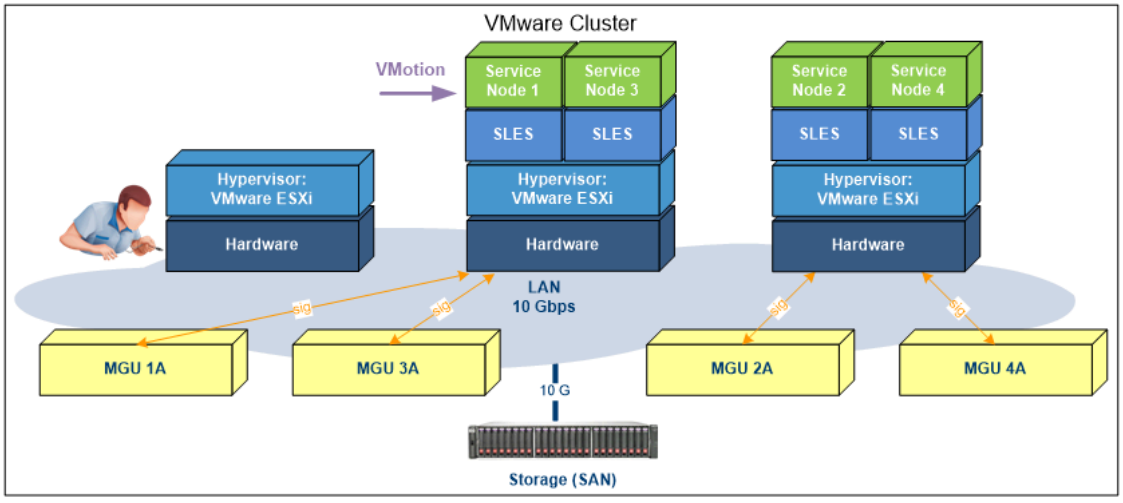

As soon as the physical server is repaired, the MX-ONE servers can be moved back to the original physical server.

Figure 18: VMware VMotion example

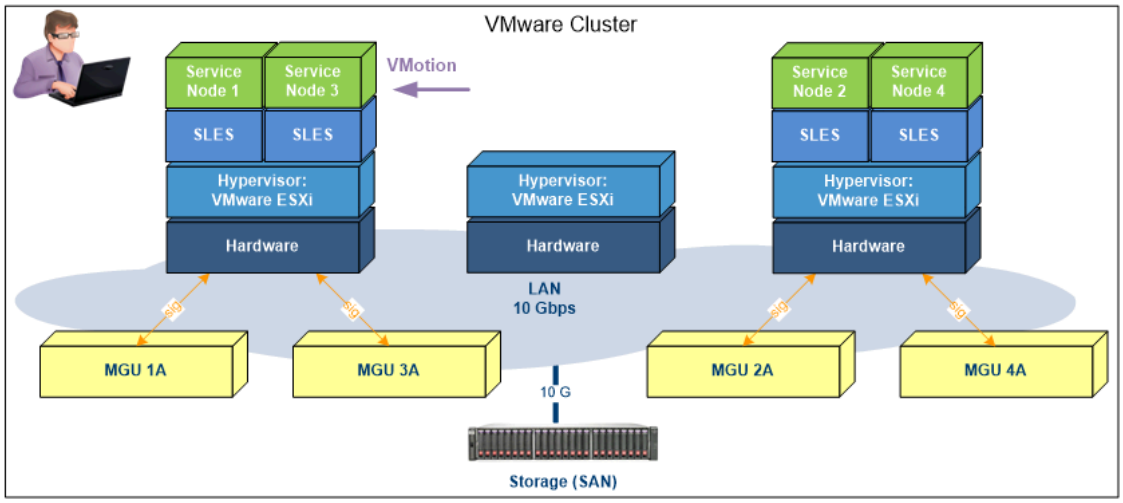

For more information regarding VMware VMotion, please read:

Migration with vMotion

https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.vcenterhost.doc/GUID-D19EA1CB-5222-49F9-A002-4F8692B92D63.html

3.2 MiVoice MX-ONE Support - Hyper-V

MiVoice MX-ONE 7.1 and later can be deployed in a Microsoft Hyper-V virtual machine based on Windows 2016 and Windows 2019 in two main scenarios, Basic Hyper-V setup and Failover Clustering.

3.2.1 Basic Hyper-V

In a basic Hyper-V setup, you can deploy multiple virtual instances of MX-ONE on a single physical network. The operating systems for each of the virtual machines are isolated from each other freeing up the unused hardware resources.

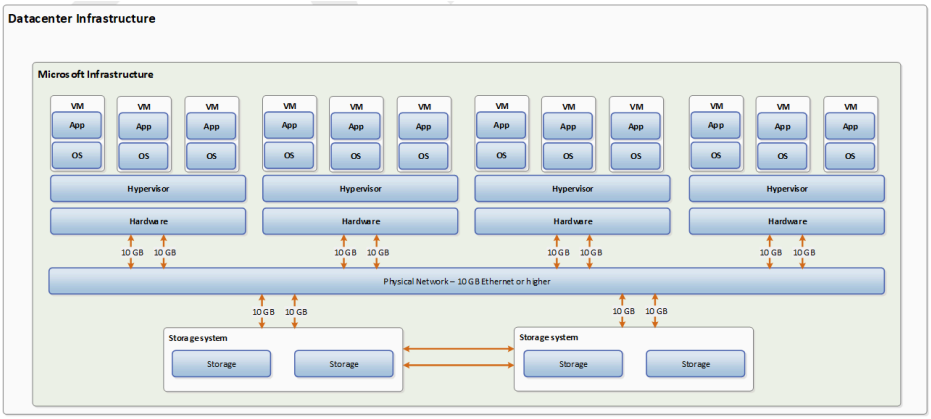

3.2.2 Failover Clustering

In a Failover clustering setup, if the physical machine on which the MX-ONE virtual machine is running goes down, the MX-ONE virtual machine can be re-initiated in another physical host. (Similar to VMware High Availability).

For more information, see https://docs.microsoft.com/en-us/windows-server/failover-clustering/clustering-requirements

3.3 MiVoice MX-ONE Support - KVM and Red Hat Enterprise

MiVoice MX-ONE 7.2 and later can be deployed in a KVM virtual machine based on SUSE Linux Enterprise Server (SLES) 12 SP5 or later as well as Red Hat Enterprise Linux 7.6 or later.

Note:
Mitel does not provide Red Hat Enterprise Linux 7.6 software.

MX-ONE Requirements 4

This chapter contains the following sections:

- System Requirements for Virtualized Infrastructure
- Requirements - VMware
- MX-ONE Service Node Resources Requirement
- Requirements - VMware
- Requirements - Hyper-V and KVM
- Recommended Training
- Product Business Approval (PBA)
- Upgrading MX-ONE in KVM to Support Minimum Downtime

The section discusses the MX-ONE requirements needed in a virtualized environment.

4.1 System Requirements for Virtualized Infrastructure

Mitel integrated MX-ONE to run on top of virtualized (VMware, Hyper-V based on Windows 2016 Server and Windows 2019 Server, KVM (SUSE Linux Enterprise Server or Red Hat Enterprise Server 7.6) infrastructure.

MiVoice MX-ONE virtualization solution supports the followings gateways/chassis:

- MGU and MGU2 (high capacity media gateway board for the 19” building practice – (MX-ONE Lite, MX-ONE Slim and MX-ONE Classic chassis)

MX-ONE Media Server

- MX-ONE Server software

Database Server

Management

- Provisioning Manager
- Service Node Manager

4.2 Requirements - VMware

Note:
MX-ONE Virtual Machines requires minimum ESXi 5.5 Virtual Machine version 10, but Mitel recommends ESXi 6.7.

Mitel recommends the use of thick provisioned lazy zeroed disks when deploying Consolidated or High Availability setups. Fault Tolerance setup requires thick provisioned eager zeroed.

When running VMware HA and FT, it is necessary that sufficient server resources and priority are configured for MX-ONE guest machines in a cluster in order to avoid that any failed virtual machine have standby resource precedence over the MX-ONE guest machines in a failover scenario. Such a situation would cause real-time traffic disturbances in a customer environment, which would defeat the purpose of the High Availability or Fault Tolerance.

The MX-ONE 7.x and later OVA default setup provided by Mitel is shown in the following table. .

Note:
VM requirement in the following table needs to be changed according to the system size (number of devices per server). See the tables MX-ONE Guest Machine Requirements for adjusting the requirements appropriately.

VMware ESXi version	ESXi 5.5 and later
vCPU	2
CPU Reservation	1000 MHz
Memory	6 GB
NIC driver	VMXNET3
VM Hard disk space	120 GB
Virtual Machine Version	10

In general, it is always recommended to follow VMware recommendations for host server configurations, but when running MX-ONE in a virtualized environment, the minimum configuration for the consolidated and high availability setup is described in the item MX-ONE Guest Machine requirements.

Mitel recommends two or more virtual CPUs. If the MHz need exceeds 2500, then three virtual CPUs should be configured. If the MHz need exceeds 3500, four or five virtual CPU should be configured and so on.

Note:
Partners/customers must change the default OVA setup according to the MX-ONE application that the VM runs and the system size as described in the MX-ONE Guest Machine requirement tables. The default setting cannot be changed from 120GB SSD disk size. However, you can use the template from the recovery image (MX-ONE_7.x_Server_Template.ovf) to create a new virtual machine and change the disk size to appropriate value and install the virtual machine from the recovery image.

Note:
The default setting cannot be changed since we have delivered 120GB ssd disk. A workaround is done to use the template from the recovery image (MX-ONE_7.x_Server_Template.ovf) to create a new virtual machine and to change the disk size to appropriate value. Install the virtual machine from the recovery image.

4.3 MX-ONE Service Node Resources Requirement

Mitel has performed several tests in its laboratories in order to determine the optimum setup for the host machines as well as the guest machines when running MX-ONE Service Node as well as MX-ONE Media Server in a virtualized environment. There are different aspects that need to be analyzed when deploying a real-time communication system in a virtual environment.

Mitel recommends the following to achieve an optimum performance:

4.3.1 Basic Guidelines

This section provides the basic guidelines to follow for a virtualized infrastructure (VMware, Hyper-V based on Windows 2016 Server and Windows 2019 server, KVM (SUSE Linux Enterprise Server or Red Hat Enterprise Server 7.6) setup.

Mitel strongly recommends dedicated host machines to MX-ONE Service Node (SN), Provisioning Manager (PM), Database (Cassandra), Media Server (MS) and guest machines.

If this is not possible, then it is strongly recommended to avoid mixing MX-ONE guest machines on the same host server with CPU intensive guest machines, such as Windows applications, mails servers, databases, etc., as they could affect performance with real-time applications such as MX-ONE Service Node and the MX-ONE database (Cassandra).

CPU over provisioning and Memory over commitment are not supported by MX-ONE system.

In general, MX-ONE communication system as with any real-time application, requires low I/O read and write disk average latency. Although 2 ms or less would be considered optimum and highly recommended, an average delay of less than 10ms, which can be expressed in either "await" in Linux terms or GAVG in VMware terminology should not induce any adverse effects in the system under normal voice traffic load conditions. A minimum of 1 Gb access to the SAN array in the same site is necessary to achieve this requirement and assumes that there is no mirroring or synchronization of SANs involved, especially with a remote site (for example, stretched Cluster).

However, in a redundant SAN scenario when there is mirroring/synchronization between network storage systems, the delay between the SAN systems must be less than 2ms with 10 Gb link between SANs. This is to ensure that if the site A fails, the storage in the site B can to take over within a couple of milliseconds (1 – 2 ms) and vice versa. Otherwise, this can compromise MX-ONE I/O read and write capability to the disks, causing the system to collapse. Therefore, it is essential that the storage system be designed to support real time applications like the MX-ONE communications system.

4.3.2 MX-ONE Guest Machine Requirements - General

Mitel has executed performance tests on the following virtualized infrastructure:

- MX-ONE Service Node 7.x call manager software running as virtual machines in a VMware 6.7 infrastructure.
- MX-ONE Service Node 7.x call manager software running as virtual machines in a Microsoft Hyper-V based on Windows 2016 and Windows 2019 infrastructure.
- MX-ONE Service Node 7.X call manager software running as virtual machines in a KVM infrastructure (SUSE Linux Enterprise Server or Red Hat Enterprise Server 7.6)

The tables below show the minimum recommended configuration of virtual CPU (vCPU) and MHz reservations for the virtual machine (VM) running MX-ONE Service Node software.

This data is from Dell servers running Intel Xeon processors, family Skylake or newer.

Note:
The performance of the virtual machine should be monitored in order to verify that there is no starvation of the Virtual Machine. This is really important if there is more than one VM on the same host.

MX-ONE 7.x has been verified to handle the traffic described in the tables per server VM when running in Virtualized solution.

Traffic results are based on 0.2 Erlang and approximately 6 calls per user per hour in average of 100 seconds per call.

The formula for the calculation is:

Number of user = 1/erlang * average call time * traffic (speed in calls per second).

Number of users = 1/0.2 * 100.1

Number of user = 500.

Note:
The information provided in this document does not cover Contact Center traffic.

The following table shows the number of extensions supported per server for 15000 devices based on the number of users per device.

Max Capacity	Devices Per User	Max Number of Extensions Per Server
15000	1	15000
15000	2	7500
15000	3	5000
15000	4	3750

For information about IOPS and Bandwidth requirements, see the document IOPS Disk and Network Bandwidth Requirements, 33/1551-ASP11301.

4.4 Requirements - VMware

Note:
MX-ONE Virtual Machines requires minimum ESXi 5.5 Virtual Machine version 10, but Mitel recommends ESXi 6.7.

Mitel recommends the use of thick provisioned lazy zeroed disks when deploying Consolidated or High Availability setups. Fault Tolerance setup requires thick provisioned eager zeroed.

When running VMware HA and FT, it is necessary that sufficient server resources and priority are configured for MX-ONE guest machines in a cluster in order to avoid that any failed virtual machine have standby resource precedence over the MX-ONE guest machines in a failover scenario. Such a situation would cause real-time traffic disturbances in a customer environment, which would defeat the purpose of the High Availability or Fault Tolerance.

The MX-ONE 7.x and later OVA default setup provided by Mitel is shown in the following table. .

Note:
VM requirement in the following table needs to be changed according to the system size (number of devices per server). See the tables MX-ONE Guest Machine Requirements for adjusting the requirements appropriately.

VMware ESXi version	ESXi 5.5 and later
vCPU	2
CPU Reservation	1000 MHz
Memory	6 GB
NIC driver	VMXNET3
VM Hard disk space	120 GB
Virtual Machine Version	10

In general, it is always recommended to follow VMware recommendations for host server configurations, but when running MX-ONE in a virtualized environment, the minimum configuration for the consolidated and high availability setup is described in the item MX-ONE Guest Machine requirements.

Mitel recommends two or more virtual CPUs. If the MHz need exceeds 2500, then three virtual CPUs should be configured. If the MHz need exceeds 3500, four or five virtual CPU should be configured and so on.

Note:
Partners/customers must change the default OVA setup according to the MX-ONE application that the VM runs and the system size as described in the MX-ONE Guest Machine requirement tables. The default setting cannot be changed from 120GB SSD disk size. However, you can use the template from the recovery image (MX-ONE_7.x_Server_Template.ovf) to create a new virtual machine and change the disk size to appropriate value and install the virtual machine from the recovery image.

Note:
The default setting cannot be changed since we have delivered 120GB ssd disk. A workaround is done to use the template from the recovery image (MX-ONE_7.x_Server_Template.ovf) to create a new virtual machine and to change the disk size to appropriate value. Install the virtual machine from the recovery image.

4.4.1 MX-ONE Guest Machine Requirements - VMware

Guest machine VMware requirements are listed in this section.

Note:
MX-ONE is a real-time application which requires resources available when it is needed. So, CPU over-provisioning or Memory over commitment of the Virtual Machine cannot be used by the MX-ONE system.

4.4.1.1 VM Consolidate and in a High Availability Setup

The tables below show the MX-ONE 7.x requirements for RAM Memory, CPU or vCPU, the clock speed and the required disk space when running combinations of the following MX-ONE applications: Service Node (SN), Service Node Manager (SNM), Provisioning Manager (PM), Database (Cassandra) and Media Server (MS). In some cases, Mitel strongly recommend running Provisioning Manager, Cassandra and/or the Media Server applications in different servers/VMs to avoid performance constraints.

Raid Controller must have at least 1 GB cache.

For the database disks, SSD disks are highly recommended to avoid performance issues both in bare metal or virtualized systems.

Service Node 1 with MX-ONE Applications

Minimum system requirements for Service Node 1 and Service Node standby, including other MX-ONE applications Service Node Manager, Provisioning Manager, Database (Cassandra) and Media Server.

Provisioning Manager must be installed in stand-alone VM for systems above of 2000 users.

The table below includes Service Node Manager, Database (Cassandra) and Media Server.

Service Node 1 and Service Node 1 Standby with Applications							
Standard VMware and High Availability							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic

Service Node 1 and Service Node 1 Standby with Applications							
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	8	2	2000	100	120	SN, SNM, Cassandra, and Media Server	1.00
1000	8	3	2500	100	120	SN, SNM, Cassandra, and Media Server	2.00
2500	10	4	3500	100	120	SN, SNM, Cassandra, and Media Server	5.00
5000	12	6	5500	100	120	SN, SNM, Cassandra	10.00
7500	14	6	6000	100	120	SN, SNM, Cassandra	15.00
10000	16	8	6000	150	180	SN, SNM, Cassandra	20.00
15000	20	8	8000	150	180	SN, SNM, Cassandra	30.00

The table below includes Service Node Manager, Provisioning Manager, Database (Cassandra) and Media Server.

Service Node 1 and Service Node 1 Standby with Provisioning Manager							
Standard VMware and High Availability							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	8	3	2500	100	120	SN, SNM, Cassandra, and Media Server	1.00

Service Node 1 and Service Node 1 Standby with Provisioning Manager							
1000	10	4	3000	100	120	SN, SNM, Cassandra, and Media Server	2.00

Other Service Nodes with MX-ONE applications

Minimum system requirements for Other Service Nodes and other Service Nodes standby, including other MX-ONE applications Provisioning Manager, Database (Cassandra) and Media Server.

Provisioning Manager must be installed in stand-alone VM for systems above of 2000 users.

The table below includes Service Node Manager, Database (Cassandra) and Media Server.

Other Service Node and Other Service Node Standby with Applications							
Standard VMware and High Availability							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	6	2	2000	100	100	SN, Cassandra, and Media Server	1.00
1000	8	3	2500	100	100	SN, Cassandra, and Media Server	2.00
2500	10	4	3500	100	100	SN, Cassandra, and Media Server	5.00
5000	12	6	5500	100	100	SN, Cassandra	10.00
7500	12	6	6000	100	100	SN, Cassandra	15.00
10000	14	8	6000	100	100	SN, Cassandra	20.00
15000	16	8	8000	100	100	SN, Cassandra	30.00

The table below includes Service Node Manager, Provisioning Manager, Database (Cassandra) and Media Server.

Other Service Node and Other Service Node Standby with applications and Provisioning Manager							
Standard VMware and High Availability							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	8	3	2500	100	100	SN, PM, Cassandra, and Media Server	1.00
1000	10	4	3500	100	100	SN, PM, Cassandra, and Media Server	2.00

The table below includes only Provisioning Manager.

Other Service Node and Other Service Node Standby with Provisioning Manager							
Standard VMware and High Availability							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	6	3	2500	100	100	SN, PM	1.00
1000	8	4	3500	100	100	SN, PM	2.00

Other Service Nodes

Minimum system requirements for Other Service Nodes (Service Node).

Other Service Node and Other Service Node Standby							
Standard VMware and High Availability							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic

Other Service Node and Other Service Node Standby							
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	6	2	2000	100	100	SN	1.00
1000	8	3	2500	100	100	SN	2.00
2500	8	4	3500	100	100	SN	5.00
5000	10	4	5500	100	100	SN	10.00
7500	12	4	6000	100	100	SN	15.00
10000	14	6	6000	100	100	SN	20.00
15000	16	6	8000	100	100	SN	30.00

Media Server Standalone

In larger networks with lots of SIP end-point traffic, it may be beneficial to have the Media Server (MS) reside in a separate server or VM to avoid putting load on the Service Node. Also, if there is a need to have more than one Media Server controlled by the same MX-ONE Service Node, the other MS instances would have to run in separate Linux server / VM, although managed by the same Service Node.

Minimum system requirements for Media Server standalone.

Media Server Standalone					
Standard VMware and High Availability					
RTP	Memory (GB)	CPU	CPU	Disk	Application
Number of RTP channels	Minimum Memory required/reserved in GB per VM	Minimum CPU/vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application
500	6	2	2000	100	Media Server
1000	8	3	2500	100	Media Server
2000	10	4	4000	100	Media Server

Provisioning Manager Standalone

This configuration with the Provisioning Manager (PM) on a separate standalone server/VM from Service Node 1 is optional and recommended for networks with more than 2000 users. This is to avoid creating

load on the SN itself, in particular when things like back-up or bulk provisioning are done. It is also the preferred configuration when a customer has more than one MiVoice MXONE sub-system (multiple networked MX-ONE systems) to manage, as it removes any dependency any given system. This is especially important when the customer's MX-ONE systems are running different SW versions, as provisioning manager can support connecting and managing different MX-ONE versions concurrently. Additionally, in this setup, the PM software version can be kept to the latest release without affecting any given MX-ONE system in the network.

Minimum system requirements for Provisioning Manager standalone.

Provisioning Manager Standalone					
Standard VMware and High Availability					
Number of Users	Memory (GB)	CPU	CPU	Disk	Application
Number of Users managed	Minimum Memory required/reserved in GB per VM	Minimum CPU/vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application
500	4	2	1000	100	PM Standalone
1000	4	2	1000	100	PM Standalone
2500	6	2	1000	100	PM Standalone
5000	6	4	3500	100	PM Standalone
7500	6	4	3500	100	PM Standalone
10000	8	4	3500	100	PM Standalone
15000	8	4	3500	100	PM Standalone
30000	12	6	4000	100	PM Standalone
50000	16	6	4000	100	PM Standalone
70000	20	6	4000	100	PM Standalone
100000	24	6	4000	100	PM Standalone

Database Standalone Or Large Co-located Databases

This configuration with the Cassandra DB on a standalone server from the SN is optional and typically only used for larger networks to avoid load on the SN itself or if there is a desire to connect more than one Service Node to the same database. Note that this will increase the IP interface load, so 10 Giga Ethernet cards are recommended in database standalone servers. For the database, SSD disks are highly recommended to reduce latency.

In virtualized systems, it is highly recommended to create database (Cassandra) servers in local SSD disks.

Minimum system requirements for Database standalone or large co-located databases.

Database (Cassandra) Standalone					
Standard VMware and High Availability					
Number of Users	Memory (GB)	CPU	CPU	Disk	Application
Number of SIP Extensions with 40 Keys	Minimum Memory required/reserved in GB per VM	Minimum vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application
500	4	2	2000	100	Cassandra Standalone
1000	4	3	2000	100	Cassandra Standalone
2500	4	3	2000	100	Cassandra Standalone
5000	7	3	3000	100	Cassandra Standalone
7500	7	3	3500	100	Cassandra Standalone
10000	9	4	4000	100	Cassandra Standalone
15000	12	4	5000	100	Cassandra Standalone
30000	15	4	5000	100	Cassandra Standalone
50000	24	8	5000	100	Cassandra Standalone
70000	28	8	5000	100	Cassandra Standalone
100000	32	8	6000	100	Cassandra Standalone

Note:

If a system includes integrated DECT phones, the calculations will need to be adjusted compared to systems with only SIP extensions due to roaming activities. For more information, refer to *Cordless Phone Installation Planning (5/1531-ANF90114)*.

The load from integrated DECT roaming is much heavier than a SIP reregistration, but one roaming is less heavy than a basic call. The heaviness of a roaming depends on the type of roaming and can range from 0.1 to 0.5 of a basic call. However, there can be many roamings for one user.

For instance, in a highly mobile DECT system, DECT users can roam 30-90 times per day or more, whereas really mobile users can do several thousand roamings per day. Suppose an example of 30 roamings per day for one user, where 50% of the roamings are heavier (0.5 call) and 50% are less heavy (0.1 call). In that case, the load from this integrated DECT terminal (call + roamings) should at least be equal to the load from nine SIP terminals (calls plus re-registrations).

In a standard DECT system, a terminal on average roams less frequently than in a highly mobile system, but assume, for example, 10 times per day. Suppose 50% of the roamings are heavier (0.5 call load) and 50% are less heavy (0.1 call load). Calculate with approximately five times more roaming events than call events (example based on Mobility Logging data). The load from a DECT terminal (call + roamings) should at least be equal to the load from a SIP terminal (call + reregistrations). For the roamings, add an equal amount to the load of three SIP users for one integrated DECT user, which equals the sum of four SIP users.

4.4.1.2 VM in a Fault Tolerance Setup

Service Node 1 and Service Node 1 Standby							
VMware Fault Tolerance							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	8	4 * **	3000	100	120	SN, SNM, Cassandra	1.00
1000	8	4 * **	4000	100	120	SN, SNM, Cassandra	2.00
2500	10	4 * **	5000	100	120	SN, SNM, Cassandra	5.00

Service Node 1 and Service Node 1 Standby							
5000	14	6 * **	6500	100	120	SN, SNM, Cassandra	10.00

Note:
* up to 8 vCPU are supported

Note:
** The capacity was validated using VMware ESXi 6.7, the network used between physical host is 10 Gigabit Ethernet. SIP traffic generators were used to simulate real traffic.

Note:
Backing up data and execute config_mirror might slow down traffic flow; hence, Mitel recommends that you back up data during low traffic periods.

4.4.1.3 Capacity Test Conclusion

Although it is technically possible to run more users and traffic in a VM, testing showed when running several smaller VMs with a lower number of users per VM in one physical host, gives better performance than one large VM with many users. Spreading the load over several VMs takes advantage of the VMware inherent resource sharing capabilities and offers better performance.

4.4.2 Media Gateway Chassis Distribution

The following table shows the Media Gateway chassis distribution per MX-ONE Service Node Virtual Machine.

Media Gateway - MGU			
	Consolidation	High Availability	Fault Tolerance
Number of vCPU	Minimum 2 *	Minimum 2 *	4
Maximum number of Media Gateways Chassis	up to 15*	up to 15*	Up to 2*
	*depends on the number of users in the Service Node	*depends on the number of users in the Service Node	*depends on the number of users in the Service Node

Consolidated and High Availability Setup

The following table shows the Media Gateway chassis distribution per number of users per MX-ONE Service Node Virtual Machine, Consolidated and High Availability setup.

Note:
The traffic limitation is in the Service Node and not in the Media Gateway, so the number of trunks should be proper dimensioned to handle the quantity of calls supported by the Service Node.

Maximum number of users per Service Node Virtual Machine	Minimum Recommended number of virtual CPU	Minimum recommended reservation, MHz	SIP traffic capacity with the recommended MHz reservation Traffic rate: 0,2 Erlang	Supported number of Media Gateway Chassis
500	2	1000 MHz	1 Call/second	2
1000	3	1500 MHz	2 Calls/second	2
2500	4	2000 MHz	5 Calls/second	5
5000	6	3500 MHz	10 Calls/second	15

Fault Tolerance Setup

The following table shows the Media Gateway chassis distribution per number of users per MX-ONE Service Node Virtual Machine, Fault Tolerance setup.

Maximum number of users per Service Node Virtual Machine	Minimum Recommended number of virtual CPU	Minimum recommended reservation, MHz	SIP traffic capacity with the recommended MHz reservation Traffic rate: 0,2 Erlang	Supported number of Media Gateway Chassis
500	4	3000 MHz	1 Call/second	Up to 2*
				*depends on the number of users
1000	4	4000 MHz	2 Calls/second	Up to 2*
				*depends on the number of users

4.4.3 Media Server

Consolidated and High Availability setup

The following tables show the Media Server requirements for Direct Media and Forced Gateway scenario. The recommended figures in below tables are done with the Media Server co-located in the same virtual machine as the Service Node, called MX-ONE Softswitch.

Direct Media Setup

It means that a full SIP solution (SIP trunk and SIP extensions) was deployed and the media server is used for conference calls, for example.

Other Service Node and Other Service Node Standby without applications						
Standard VMware and High Availability						
Number of Devices	Memory (GB)	CPU	CPU	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application	Calls per Second
500	6	2	2000	100	SN, SNM, Cassandra	1.00
1000	8	3	2500	100	SN	2.00
2500	8	4	3500	100	SN	5.00
5000	10	4	5500	100	SN	10.00
15000	16	6	8000	100	SN	15.00

* G.729/G.722 codec is not supported in this configuration.

Forced Gateway

In some traffic cases transcoding is required and for those cases more RTP gateway resources are needed. As a consequence, more CPU resources are necessary to handle those gateway calls in the media server.

Note:
The maximum number of users for the Forced Gateway traffic cases is limited by the number of RTP resources available in the Media Server.

Maximum number of users per Service Node Virtual Machine	Minimum Recommended number of virtual CPU	Minimum recommended reservation, MHz - Media Server collocated with Service Node using G.711 codec and Forced Gateway *	Minimum recommended reservation, MHz - Media Server including Service Node	SIP traffic capacity with the recommended MHz reservation - Traffic rate: 0,2 Erlang
500	2	1200 MHz	2200 MHz	1 Call/second
1000	3	1600 MHz	3100 MHz	2 Calls/second

* G.729/G.722 codec is not supported in this configuration.

4.4.4 Provisioning Manager

Mitel recommends that Provisioning Manager is installed in a separated Virtual Machine in the same environment. It can be in the same physical host or not.

4.4.5 Upgrade Process

The upgrade process in the virtualized environment is the normal MX-ONE upgrade process and the downtime will depend of the total system configuration.

4.4.6 MX-ONE Service Node Software

The Service Node is also available as a virtual appliance.

The virtual appliance is an ESXi 6.7 Virtual Machine Version 10 and the based appliance/image is not compatible with previous versions.

4.4.7 MX-ONE Media Server

MX-ONE Media Server can be installed in the same virtual machine as MX-ONE Service Node or on a separate machine. Please follow the installation instructions in the MX-ONE CPI.

4.4.8 Creating Snapshot on a VM

While you create a snapshot on a VM running the MX-ONE Service Node software, use the command `data_change` to prevent MX-ONE Service Node data changes while creating VM snapshot.

4.5 Requirements - Hyper-V and KVM

The following sections describes the MX-ONE 7.x requirements for RAM Memory, CPU or vCPU, the clock speed and the required disk space when running combinations of the following MX-ONE applications:

- Service Node (SN)
- Service Node Manager (SNM)
- Provisioning Manager (PM)
- Database (Cassandra)
- Media Server (MS)

In some cases, Mitel strongly recommend running Provisioning Manager, Cassandra and/or the Media Server applications in different servers/VMs to avoid performance constraints.

Raid Controller must have at least 1 GB cache.

If local disks are used for MX-ONE Service Node, Database Node, SSD disks are highly recommended to avoid performance issues.

Note:
For systems running Hyper-V and KVM images are delivered with a default setup, which is the minimum setup required to enable the Virtual Machine. The default values needs to be changed according to the system size (number of devices per server). See the tables MX-ONE Guest Machine Requirements for adjusting the requirements appropriately.

4.5.1 Service Node with MX-ONE Applications

The table below shows the minimum system requirements for Service Node and Service Node standby, including other MX-ONE applications Service Node Manager, Database (Cassandra) and Media Server.

Provisioning Manager must be installed in stand-alone VM for systems above of 2000 users.

Service Node and Service Node Standby with Applications							
Standard Hyper-V and Failover Clustering/ KVM							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	8	2	2000	100	120	SN, SNM, Cassandra, and Media Server	1.00
1000	8	3	2500	100	120	SN, SNM, Cassandra, and Media Server	2.00
2500	10	4	3500	100	120	SN, SNM, Cassandra, and Media Server	5.00
5000	12	6	5500	100	120	SN, SNM, Cassandra, and Media Server	10.00
7500	14	6	6000	100	120	SN, SNM, Cassandra, and Media Server	15.00

Service Node and Service Node Standby with Applications							
10000	16	8	6000	150	180	SN, SNM, Cassandra, and Media Server	20.00
15000	20	8	8000	150	180	SN, SNM, Cassandra, and Media Server	30.00

The table below includes Service Node Manager, Database (Cassandra), Media Server and Provisioning Manager running on the same machine.

Service Node and Service Node Standby with Provisioning Manager							
Standard Hyper-V and Failover Clustering/ KVM							
Number of Devices	Memory (GB)	CPU	CPU	Disk	Disk	Application	Max Traffic
Max SIP Devices Capacity	Minimum Memory required/ reserved in GB per VM	Minimum CPU/ vCPU required	Minimum MHz required/ reserved	Disk Size (GB)	Large System (more than 3 servers) (GB)	Type of Application	Calls per Second
500	8	3	1500	100	120	SN, SNM, PM Cassandra, and Media Server	1.00
1000	10	4	2000	100	120	SN, SNM, Cassandra, PM and Media Server	2.00

4.5.2 Media Server Standalone

In larger networks with many SIP end-point traffic, it may be beneficial to have the Media Server reside in a separate server or VM to avoid putting load on the Service Node. Also, if there is a need to have more than one Media Server controlled by the same MX-ONE Service Node, the other Media Server instance would have to run in separate Linux server / VM, although managed by the same Service Node.

Media Server Standalone					
Standard Hyper-V and Failover Clustering/ KVM					
RTP	Memory (GB)	CPU	CPU	Disk	Application

Media Server Standalone					
Number of RTP channels	Minimum Memory required/reserved in GB per VM	Minimum CPU/vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application
500	6	2	2000	100	Media Server
1000	8	3	2500	100	Media Server
2000	10	4	4000	100	Media Server

Note:
Force gateway calls require more CPU resources.

4.5.3 Provisioning Manager Standalone

This configuration with the Provisioning Manager (PM) on a separate standalone server/VM from Service Node 1 is optional and recommended for networks with more than 2000 users. This is to avoid creating load on the Service Node itself, in particular when things such as back-up or bulk provisioning are done. It is also the preferred configuration when a customer has more than one MiVoice MX-ONE sub-system (multiple networked MX-ONE systems) to manage, as it removes any dependency in any given system. This is especially important when the customer's MX-ONE systems are running different software versions. This is because PM can support connecting and managing different MX-ONE versions concurrently. Additionally, in this setup, the PM software version can be kept to the latest release without affecting any given MX-ONE system in the network.

Minimum system requirements for Provisioning Manager stand-alone are shown in the following table.

Provisioning Manager Standalone					
Standard Hyper-V and the Failover Clustering/KVM					
Number of Users	Memory (GB)	CPU	CPU	Disk	Application
Number of Users managed	Minimum Memory required/reserved in GB per VM	Minimum CPU/vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application
500	4	2	1000	100	PM Standalone
1000	4	2	1000	100	PM Standalone
2500	6	2	1000	100	PM Standalone
5000	6	4	3500	100	PM Standalone

Provisioning Manager Standalone					
7500	6	4	3500	100	PM Standalone
10000	8	4	3500	100	PM Standalone
15000	8	4	3500	100	PM Standalone
30000	12	6	4000	100	PM Standalone
50000	16	6	4000	100	PM Standalone
70000	20	6	4000	100	PM Standalone
100000	24	6	4000	100	PM Standalone

4.5.4 Database Standalone or Large Co-Located Databases

This configuration with the Database Node (Cassandra) on a stand-alone server from the Service Node is optional and typically only used for larger networks to avoid load on the Service Node itself or if there is a desire to connect more than one Service Node to the same database.I

If local disks are used for MX-ONE Databases, SSD disks are highly recommended to reduce latency.

Minimum system requirements for Database standalone or large co-located databases are shown in the following table:

Database (Cassandra) Standalone					
Standard Hyper-V and the Failover Clustering/KVM					
Number of Users	Memory (GB)	CPU	CPU	Disk	Application
Number of SIP Extensions with 40 Keys	Minimum Memory required/reserved in GB per VM	Minimum vCPU required	Minimum MHz required/reserved	Disk Size (GB)	Type of Application
500	6	2	2000	100	Cassandra Standalone
1000	6	3	2000	100	Cassandra Standalone
2500	6	4	2000	100	Cassandra Standalone
5000	10	4	3000	100	Cassandra Standalone

Database (Cassandra) Standalone					
7500	10	4	3500	100	Cassandra Standalone
10000	12	6	4000	100	Cassandra Standalone
15000	16	6	5000	100	Cassandra Standalone
30000	20	8	5000	100	Cassandra Standalone
50000	24	8	5000	100	Cassandra Standalone
70000	28	8	5000	100	Cassandra Standalone
100000	32	8	6000	100	Cassandra Standalone

4.6 Recommended Training

Partners that will install the MiVoice MX-ONE 7.x virtualization solution must be trained in installing and maintaining a VMware infrastructure. Mitel requires engineers with VMware VCP training level and certification to do the implementation in a VMware environment.

The minimum recommended training is:

- VMware vSphere: Install, Configure, Manage [Version 6.5 or later]
- Hyper-V: Install, Configure, Manage

4.7 Product Business Approval (PBA)

Although, VMware certification is not obligatory to obtain PBA for the MiVoice MX-ONE 7.x solution, it is obligatory that the partner have at least one engineer that holds a VCP certification before they do the virtualization implementation themselves in a live customer environment. Alternatively, if the partner themselves are PBA certified, but do not have a VCP certified engineer, then they must ensure that the implementation be carried out with assistance of a VCP certified engineer, such as a customer IT engineer or their VMware certified partner.

4.8 Upgrading MX-ONE in KVM to Support Minimum Downtime

This section describes how to achieve minimum downtime when upgrading MX-ONE using a Kernel-based Virtual Machine (KVM) server.

Description

Using KVM, you can configure two systems in parallel in MX-ONE server. For example, each physical server can have two service nodes, named *SN current* and *SN upgrade to*. Therefore, during the switchover from the *SN current* to the *SN upgrade to* node, the phones (configured in MX-ONE server) must be moved to *SN upgrade to* without causing system interruption so that there would be minimum or no downtime for the end-users.

Note:

The *SN upgrade to* node must have its configuration completed, tested, and the OS-patched included to be ready to move to the next phase.

Assumptions

- During the switchover, ongoing calls will be dropped, but new calls will be connected in within a few seconds.

Requirements

- The physical server must have enough capacity enough to handle at least two SNs and the KVM server.
- The physical server must have at least 2 NICs.
- *SN current* is running on the same KVM Host Server as that of *SN upgrade to* service node.
- Configure a different VLAN for configuring and testing the *SN upgrade to* system.
- An extra license is needed for the *SN upgrade to* service node as this is a new installation.

Procedure

Prepare Network

Connect eth1 of the physical server to an isolated switch or special VLAN when the *SN upgrade to* MX-ONE is up, it will not make conflicts with the *SN upgrade to*. All the test and upgrade work are done in that isolated environment.

Create Network Bridge

1. Log in as root on the KVM Host Server.
2. Start **yast**> **System**> **Network Settings**.
3. Click the **Overview** tab and click **Add**.

   a. Select **Bridge** from the **Device Type** list and enter the bridge device interface number, that is **1**, in the **Configuration Name** entry field.
   b. Press **Next** to proceed.

> **Note**:
>
> Do not add more than one Ethernet interface to the bridge. Otherwise, there will be package loss.

4. Set IP address, network mask on the **bridge br1**.
5. Click **OK**> **Run**> **Quit yast**.
6. Check whether **bridge br1** is created on the KVM Host Server.

```
root # brctl show
bridge name bridge id           STP enabled   interfaces
br0         8000.28b9d9e339d4       no        eth0
                                              vnet0
br1         8000.00085d9ca893       no        eth1
```

Install the SN Upgrade to

1. Download the .qcow2 and .xml files from the package repository and save all under folder `/local/images`.
2. Create a virtual machine and enter the following details:

```
virt-clone --original-xml Virtual_Image-MX-
ONE_7.3.sp1.hf3.rc6.xml --name <your system name> --file
/local/images/<your system name>.gcow2
```

   For example:

```
cd /local/images/virt-clone --original-xml Virtual_Image-MX-
ONE_7.3.sp1.hf1.rc6.xml --name kvm73116 --file /local/images/kvm73116.gcow2
```

3. Install Mobaxterm (GUI or X server support) to start your virtual machine. Refer to https://mobaxterm.mobatek.net/download-home-edition.html.

4. Run the KVM machine and run the command `virt-manager` on the command line in **Mobaxterm**.

5. If you have the virtual manager access, you can start, stop, change, and install the MX-ONE deployed on KVM.

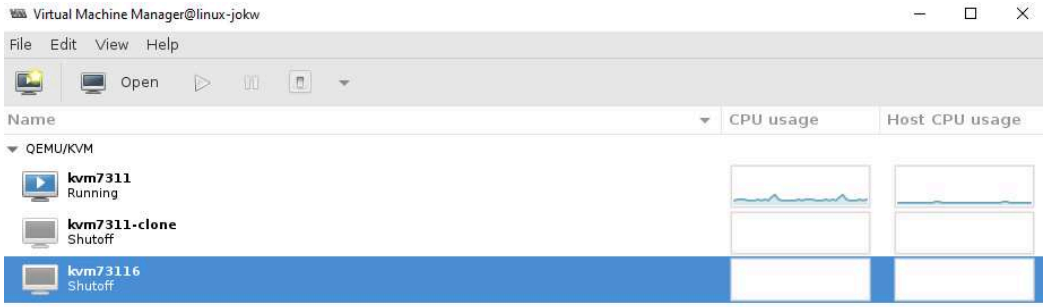

6. Open the new MXONE KVM image, find its NIC configuration, choose **bridge br1** as network source, and apply the change so that the *SN Upgrade to* node that can access the network through **bridge br1**.

7. Start the new MX-ONE KVM machine, then start configuring the new MX-ONE instance by following the instructions described in the *INSTALLING AND CONFIGURING MIVOICE MX-ONE* document.

Public Cloud 5

This chapter contains the following sections:

- MiVoice MX-ONE Support - Microsoft Azure
- Requirements - Microsoft Azure

This chapter describes the MX-ONE VM deployment in the public cloud environment, such as Microsoft Azure.

5.1 MiVoice MX-ONE Support - Microsoft Azure

A standard MiVoice MX-ONE system composed by Service Node, Cassandra DB, Media Server, Provisioning Manager and Service Node Manager, which can be deployed in Microsoft Azure on one or several Virtual Machines. The above components can all reside in the same VM or depending on size, traffic and load, can be spread out on more than one VM.

As an example, in larger deployments, it would be recommended to put the Provisioning Manager and/or the Media Server component(s) on separate VMs. It is always recommended to follow the server capacity and deployment guidelines as stated in the MiVoice MX-ONE Dimensioning and Capacity documents in the CPI. The main target for this solution is a centralized system.

For an MX-ONE system to be deployed in Azure, a specific MX-ONE image is required. The MX-ONE image includes SLES 12 SP5.

The figure below shows the high-level architecture of the solution. The MiVoice MX-ONE servers are installed in Azure, while the other Mitel devices/applications such as SIP Phones, SIP clients, Mitel Common Gateways EX and GX and Media Gateway Unit are installed on the customer site (on-premises).

In this setup, the PSTN trunks are terminated in the customer site while SIP trunks might be terminated in Azure.

Secure IP connections in the form of a dedicated network or Azure Express Connect VPN services are always required between the customer site (on-premises) and the customer's MX-ONE environment in Azure Cloud Services. Depending on customer needs and premises, an option with SBC and secure SD-WAN (Software-Defined Networking in a wide area network) over internet should also be considered.

Figure 19: High level architecture - MX-ONE deployed in Microsoft Azure

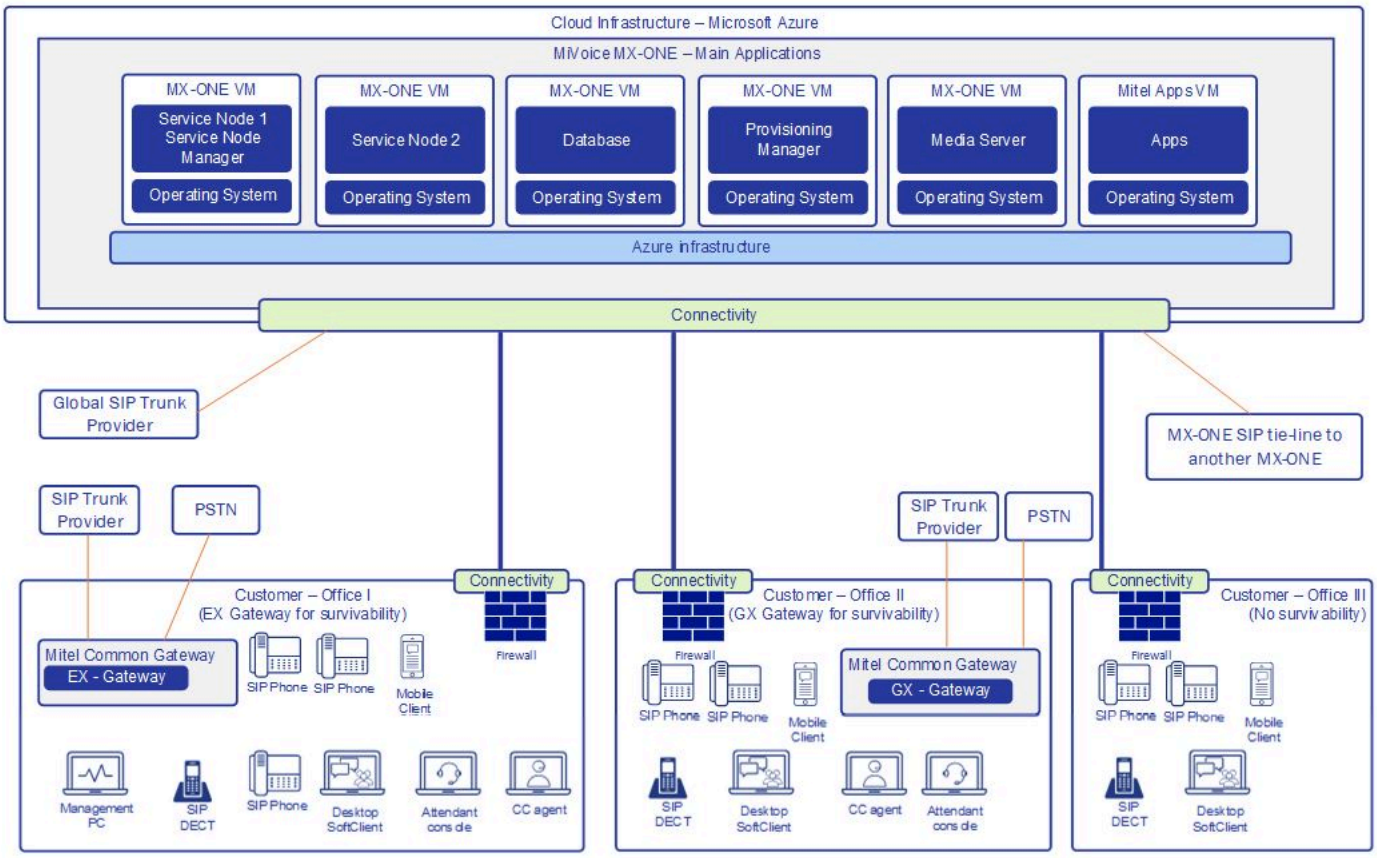

5.2 Requirements - Microsoft Azure

This section discusses the MX-ONE requirements and recommendations for Azure.

5.2.1 MX-ONE Server requirements

For MX-ONE server requirements (vCPU, Memory, etc) please check the Virtualization - Description, Hyper-V and Azure section and IOPS Disk and Network Bandwidth Requirements documents.

Azure VM sizes

Azure supports different sizes for Linux virtual machines.

MX-ONE was tested with Azure General Purpose Dsv3-series (D4s_v3 and D8s_v3) for system with low and medium traffic and Compute optimized Fsv2-series (F4s_v2 and F8s_v2) for system with high traffic requirements.

Monitoring

Azure offers graphic tools to monitor the VM performances (figure below), so it is recommended that the partner/customer monitor the MX-ONE VMs to verify if the VM used is correct.

Figure 20: Data collected during the MX-ONE upgrade

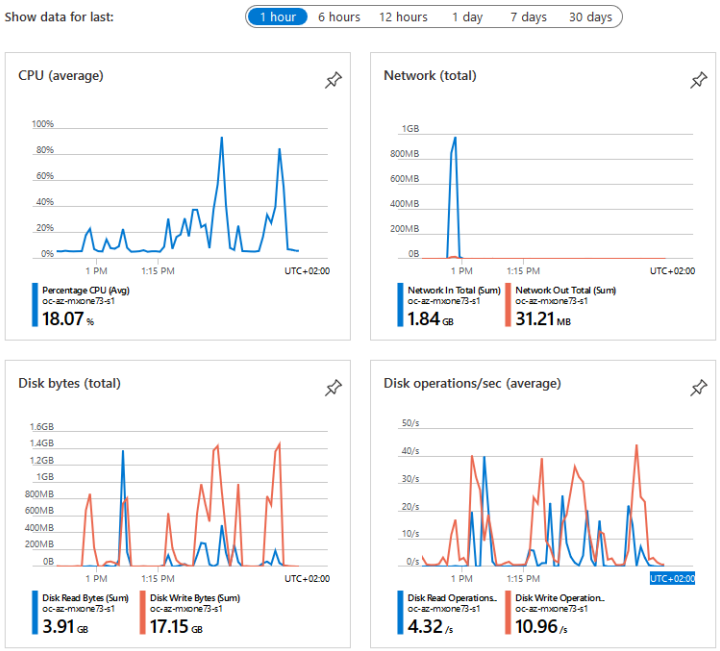

Disks

Azure Premium SSD disks are recommended as they are designed to production and performance sensitive workloads. For additional information about Azure disk types please check Premium SSD in Azure documentation site.

Another important Azure article regarding disks is Azure premium storage: design for high performance.

Scale up or down VMs.

Azure offers the possibility of scale the VM vertically (up and down) via VM resize, this function can be used by MX-ONE, however the MX-ONE VM will be restarted and the traffic will be disconnected, so if scale up or down is required, it is highly recommended, and possibly to be done during low or no traffic hours.

Dedicated host

Azure offers dedicated host, which is a service that provides dedicated physical servers. Azure Dedicated Hosts is highly recommended for customers that require control of their environment and run 24 x 7 x 365.

Mitel recommends the Azure Dedicated Hosts service for MX-ONE system, because the customer/partner can control the load in that specific physical server avoiding resource starvation (vCPU and memory).

5.2.2 Connectivity between on-premises and Azure

The figure below shows the typical on-premises and Azure network connectivity. In this case, the Azure Express Route service is used.

This option is recommended by Mitel, because it provides better network connectivity (latency, jitter, etc) than other types of connections available in Azure.

Figure 21: Azure connectivity example

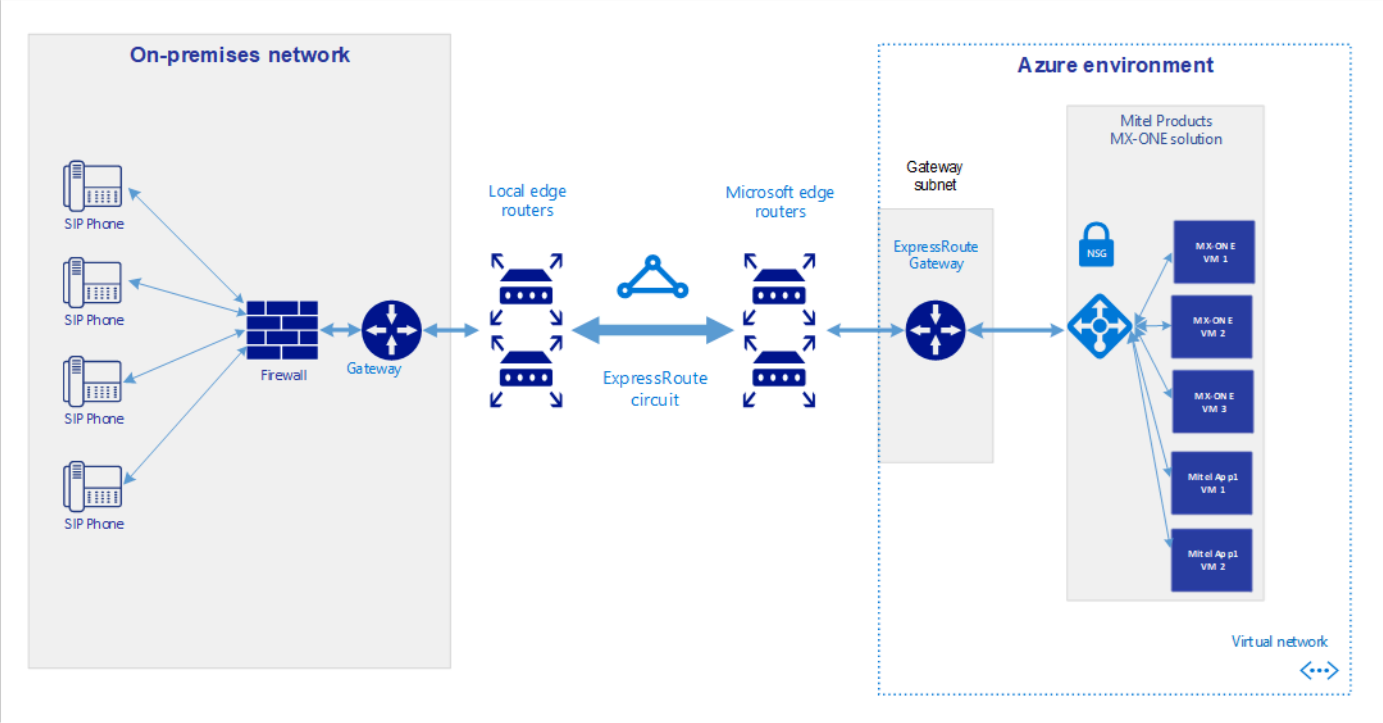

Latency

Voice over IP systems requires low latency, so the maximum supported Round Trip Delay is 150 ms between MX-ONE components. E.g. MX-ONE Service Node and a SIP Phone on-premises, MX-ONE Service Node and Mitel Common Gateway (EX and GX) on-premises, MX-ONE Service Node and Media Server on-premises. Therefore, it is recommended to deploy MX-ONE image in the nearest located Azure DC. Monitoring of latency can be achieved with MPA (Mitel Performance Analytics) and should be considered in lack of other network monitoring tools.

5.2.3 Availability - MX-ONE Redundancy

The MX-ONE redundancy options are not available in Azure. To have availability in MX-ONE system, the Azure built-in resources must be used. The resources provided by Azure are for example Availability and Azure Dedicated Hosts.

- **Availability** - https://docs.microsoft.com/en-us/azure/virtual-machines/linux/availability
- **Azure Dedicated Hosts** - https://docs.microsoft.com/en-us/azure/virtual-machines/windows/dedicated-hosts#groups-hosts-and-vms

Select the Availability options and Availability zone required by the customer that will use the MiVoice MX-ONE system.

Under Advanced selected Host group select Azure Dedicated Hosts in the customer Azure subscription, if available.

5.2.4 Gateway Calls - Forced Gateway

Media Servers using forced gateway calls require a dedicated VM in Azure. So, a stand-alone Media Server needs to be used.

Note:
This setup will require that the media pass through Azure, which is not an optimal setup as it will cause extra bandwidth costs to the customers (traffic leaving Azure is charged after certain amount of data). This scenario is not recommended by Mitel.

5.2.5 Security Requirements

In a cloud environment the security is a shared responsibility between the customer and the cloud provider, so Mitel highly recommends that the MX-ONE system is deployed using the available security mechanisms provided by MX-ONE as well as the Azure security best practices.

Connections

TLS setup is configurable in MX-ONE and it has four levels (low, medium, high, and modern security level).

Mitel highly recommend that the whole system is deployed using modern security level, which means TLS 1.3 (SIP signaling) and HTTPS with TLS 1.3 (VDP) for protecting data in transit.

It is also recommended that data at rest is encrypted when that is available.

Secure media (SRTP) is also highly recommended when this option is allowed.

It is highly recommended that Provisioning Manager and Service Node Manager are setup to use TLS 1.3 only.

Certificates

Mitel recommends the use of certificates Issued by well-known CA authorities. In case customer has its own CA (banks, authorities, etc.) a possibility to import customer's Root CA and issuer exists, but it's solely customer's responsibility.

Firewall

It is highly recommended the use of firewalls between the enterprise and Azure as well as Session Border Controller (SBC) for SIP phones used by remote workers (Tele-worker solution).

Mitel MBG was tested as part of this solution, however the MBG was sitting in the enterprise side and not in Azure.

Mitel does not recommend that MX-ONE services (SIP or management) are directly exposed in the Internet without the proper setup. All tests executed in Mitel´s laboratories/Azure used only private IP addresses and public IP addresses were deleted from the standard Azure setup.

The required ports for MX-ONE, SIP Phones and Mitel applications are described in the MX-ONE document System Planning, chapters "IP Protocols and Ports" and "Proprietary Protocols".

Phone Models

Mitel recommends the use of SIP phones family 68XX and 69XX for system that are connected in Azure, because they support TLS 1.3 for signaling encryption as well as 802.1x network authentication.

5.2.6 Mitel Applications

Please verify if the Mitel application that needs to be integrated with MX-ONE is Azure ready, before starting to deploy it.

Acronyms, Abbreviations, and Glossary

6

- **FT –** Fault Tolerance
- **Guest Machine –**The same as Virtual Machine
- **HA –**High Availability
- **HD –**Hard Disk
- **Host Machine –**Physical machine that executes the hypervisor
- **HTTPS –** Hypertext Transfer Protocol Secure
- **Hyper-V –** Microsoft® Hyper-V is one of the virtualization options available in the market. It provides a layer of abstraction between the physical hardware and the operating system. One physical hardware can run multiple operating systems, while keeping these operating systems isolated from each other.
- **IT –**Information Technology
- **KVM** - Kernel-based Virtual Machine.
- **MGU –**Media Gateway Unit
- **NIC –**Network Interface Card
- **TDM –**Time Division Multiplex
- **SAN –** Storage Area Network
- **VCP –**VMware Certified Professional
- **MX-ONE Lite –**MX-ONE Lite is a 3U high, 19-inch wide sub-rack with 4 board positions for different functions and interfaces. The media gateway it is based on the MGU board
- **MX-ONE Media Server –**The MX-ONE Media Server provides a software emulated version of the Media Gateway Unit board (MGU board) to bring media services like recorded voice announcements and conferencing to a SIP-only environment. In such environment the whole MiVoice MX-ONE runs in a Linux hosted server and no proprietary hardware is needed
- **MX-ONE Softswitch –**A version of MiVoice MX-ONE composed by software only. It contains at least one MX-ONE Service Node and one MX-ONE Media Server
- **MiVoice MX-ONE –** Unified Communication System that provides business class telephony features. It performs call control, call-signaling, and media transcoding and conversion functions. It is composed by at least one call manager software named Service Node and at least one media gateway
- **NTP –** Network Time Protocol
- **SBC –** Session Border Controller
- **SIP –**Session Initiation Protocol
- **SSD –** Solid-State drive
- **SSH –** Secure Shell
- **TCP –** Transport Control Protocol
- **TLS –** Transport Layer Security
- **UDP –** User Data Protocol
- **VM –**Virtual Machine according to VMware: “A virtual machine is a tightly isolated software container that can run its own operating systems and applications as if it were a physical computer. A virtual machine behaves exactly like a physical computer and contains it own virtual (i.e. software-based) CPU, RAM hard disk and network interface card (NIC). An operating system can’t tell the difference between a virtual machine and a physical machine, nor can applications or other computers on a network. Even the virtual machine thinks it is a “real” computer. Nevertheless, a virtual machine is composed entirely of software and contains no hardware

components whatsoever. As a result, virtual machines offer a number of distinct advantages over physical hardware”.

Reference Documents 7

MX-ONE CPI documentation

MX-ONE Media Server Description, part of the MX-ONE CPI documentation

VMware documentation

VMware provide online documentation, for more information please go to VMware site in order to get the latest documentation.

https://docs.vmware.com/en/VMware-vSphere/

Microsoft Hyper-V – Windows Server 2016 and Windows Server 2019 documentation

Microsoft provides online documentation, for more information please go to Microsoft Hyper-V Website to get the latest documentation.

The basic Hyper-V functionality.

https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/hyper-v-on-windows-server

The Failover Clustering functionality.

https://docs.microsoft.com/en-us/windows-server/failover-clustering/whats-new-in-failover-clustering

Networking

https://docs.microsoft.com/en-us/windows-server/networking/networking

Linux KVM Hypervisor Documentation

The Kernel-based-virtual-machine documentation is available in the web page https://www.linux-kvm.org/page/Main_Page

General Azure documentation page:

https://docs.microsoft.com/en-us/azure/?product=featured

Sizes for Linux virtual machines in Azure

https://docs.microsoft.com/en-us/azure/virtual-machines/linux/sizeshttps://docs.microsoft.com/en-us/azure/virtual-machines/linux/sizes

Availability

https://docs.microsoft.com/en-us/azure/virtual-machines/linux/availability

Azure Dedicated Hosts

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/dedicated-hosts#groups-hosts-and-vms

Azure networking services overview

https://docs.microsoft.com/en-us/azure/networking/networking-overview

Virtual Network documentation

https://docs.microsoft.com/en-us/azure/virtual-network/

ExpressRoute overview

https://docs.microsoft.com/en-us/azure/expressroute/expressroute-introduction

https://docs.microsoft.com/en-us/azure/expressroute/

Designing for high availability with ExpressRoute

https://docs.microsoft.com/en-us/azure/expressroute/designing-for-high-availability-with-expressroute

ExpressRoute FAQ

https://docs.microsoft.com/en-us/azure/expressroute/expressroute-faqs

Virtual Hard Disk

https://docs.microsoft.com/en-us/previous-versions/windows/desktop/legacy/dd323654(v%3Dvs.85)

Azure security best practices

https://azure.microsoft.com/en-us/resources/security-best-practices-for-azure-solutions/

Azure premium storage: design for high performance

https://docs.microsoft.com/en-us/azure/virtual-machines/linux/premium-storage-performance