

# Virtualization Description

MIVoice MX-ONE SERVICE NODE 6.X AND MEDIA SERVER



## NOTICE

The information contained in this document is believed to be accurate in all respects but is not warranted by Mitel Networks™ Corporation (MITEL®). Mitel makes no warranty of any kind with regards to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. 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](mailto: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>.

© Copyright 2018, Mitel Networks Corporation

All rights reserved

---

# CONTENTS

<b>1</b>	<b>REVISION INFORMATION.....</b>	<b>1</b>
<b>2</b>	<b>INTRODUCTION .....</b>	<b>2</b>
2.1	SCOPE OF THIS DOCUMENT	2
2.2	MIVOICE MX-ONE VIRTUALIZATION	2
<b>3</b>	<b>BENEFITS OF VIRTUALIZATION .....</b>	<b>4</b>
<b>4</b>	<b>SOLUTION DESCRIPTION .....</b>	<b>5</b>
4.1	MX-ONE SERVICE NODE	5
4.2	MX-ONE MEDIA SERVER	6
4.3	INTRODUCTION OF THE SOLUTION	7
4.4	VALIDATED SCENARIOS – SERVICE NODE	9
4.4.1	CONSOLIDATED SETUP	9
4.4.2	AVAILABILITY SETUPS WITH VMWARE AVAILABILITY SOFTWARE OPTIONS	11
4.4.3	MAINTENANCE USING VMWARE VMOTION	17
4.5	VALIDATED SCENARIOS – MEDIA SERVER	19
4.5.1	CONSOLIDATED SETUP	19
4.5.2	VMWARE HIGH AVAILABILITY SETUP	20
<b>5</b>	<b>MX-ONE REQUIREMENTS.....</b>	<b>23</b>
5.1	MX-ONE REQUIREMENTS	23
5.2	MX-ONE SERVICE NODE RESOURCES REQUIREMENT	23
5.2.1	BASIC GUIDELINES	23
5.2.2	SERVICE NODE GUEST MACHINE TESTED CAPACITY	24
5.2.3	SERVICE NODE VIRTUAL MACHINE DIMENSIONING	26
5.2.4	MEDIA GATEWAY CHASSIS DISTRIBUTION PER MX-ONE SERVICE NODE VIRTUAL MACHINE	28
5.2.5	MEDIA SERVER	29
5.2.6	MX-ONE SERVICE NODE VIRTUALIZATION BENEFITS	30
5.3	UPGRADE PROCESS	31
5.4	MX-ONE SERVICE NODE SOFTWARE	31
5.5	CREATING SNAPSHOT ON VIRTUAL MACHINE RUNNING THE MX-ONE SERVICE NODE SOFTWARE	31
5.6	MX-ONE MEDIA SERVER	32
5.7	PROVISIONING MANAGER	32
5.8	RECOMMENDED TRAINING	32
5.9	PRODUCT BUSINESS APPROVAL (PBA)	32
<b>6</b>	<b>ACRONYMS, ABBREVIATIONS AND GLOSSARY .....</b>	<b>33</b>
<b>7</b>	<b>REFERENCE DOCUMENTS.....</b>	<b>34</b>

# 1 REVISION INFORMATION

2015-03-06 Version /02: This is the first official version of the document.

2015-05-12 Version /03: One figure updated with correct naming of Service Node.

2015-05-22 Version /04: New template

2016-02-10 Version /05: Re-naming done for some concepts and re-branding for Mitel.

2016-04-21 Version /06: Trademarks text updated.

2016-06-10 Version /07: Solidus renamed MiContact Center Enterprise.

2018-02-20 Version /08: Virtual Description updated under Chapter 5

2018-03-26 Version /09: Changed SLES 11 to SLES 12.

2018-07-12 Version /10: Changed Recommended number Virtual Machine and Fault Tolerance compatible.

## 2 INTRODUCTION

This guide describes the setup configurations that apply when setting up Mitel MiVoice MX-ONE 6.x as well as MX-ONE Media Server on servers enabled with VMware® vSphere™ virtualization.


### 2.1 SCOPE OF THIS DOCUMENT

The aim of this document is to provide a description of MX-ONE Service Node and MX-ONE Media Server 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 6.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.

### 2.2 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 6.x, the Service Node 6.x as 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 Media Server was validated in a virtualized environment being collocated 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 in order 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, 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 Service Node.

Currently, MiVoice MX-ONE relies on VMware software as part of its virtualization solution. Mitel has verified MX-ONE Service Node 6.x call manager software, Media Server as well as Mitel UC applications, such as Mitel CMG, MiContact Center Enterprise and Mitel Advanced Messaging, included as part of the MiVoice MX-ONE 6.x solution, running as virtual machines in a VMware 5.x infrastructure. As this document focuses on the MX-ONE Service Node and MX-ONE Media Server in a virtualized environment, please refer to the relevant UC application documentation for virtualization guidelines.

## 3 BENEFITS OF VIRTUALIZATION

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

MX-ONE Service Node virtualization benefits are described in the item 4.2.5.

## 4 SOLUTION DESCRIPTION

MX-ONE Service Node 5.0 and later, including MX-ONE Media Server can run in a hardware virtualization environment, formally named Virtualization. Currently, MX-ONE relies on VMware software as part of its virtualization solution.

Mitel has verified MX-ONE Service Node 6.x call manager as well as MX-ONE Media Server software running as virtual machines in a VMware 5.5 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 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.

It is assumed:

- That a VCP engineer 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.

### 4.1 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 at least the following VMware's reference documents:

VMware vSphere Basics Guide – ESXi 5.0

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-basics-guide.pdf>

vSphere Availability ESXi 5.5

<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-availability-guide.pdf>



vSphere High Availability Deployment Best Practices

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

Protecting Mission-Critical Workloads with VMware Fault Tolerance

<http://www.vmware.com/resources/techresources/1094>

VMware Fault Tolerance Recommendations and Considerations on VMware vSphere 4

<http://www.vmware.com/resources/techresources/10040>

Performance Best Practices for VMware vSphere™ 5.5

[http://www.vmware.se/pdf/Perf\\_Best\\_Practices\\_vSphere5.5.pdf](http://www.vmware.se/pdf/Perf_Best_Practices_vSphere5.5.pdf)

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

<http://www.vmware.se/files/pdf/techpaper/voip-perf-vsphere5.pdf>

VMware vSphere with Operations Management and VMware vSphere

Licensing, Pricing and Packaging

<http://www.vmware.com/files/pdf/products/vsphere/VMware-vSphere-Pricing-Whitepaper.pdf>

Please, always check the latest products' documentation.

## 4.2 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 collocated as well as standalone deployments. In a collocated 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.se/files/pdf/techpaper/voip-perf-vsphere5.pdf>

Please, always check the latest products documentation.

## 4.3 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 6.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 6.x software running on top of a host machine with VMware's vSphere ESXi 5.x 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 Service Node 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**

## 4.4 VALIDATED SCENARIOS – SERVICE NODE

Mitel validated three different setups for virtualization on top of VMware's software. These setups are:

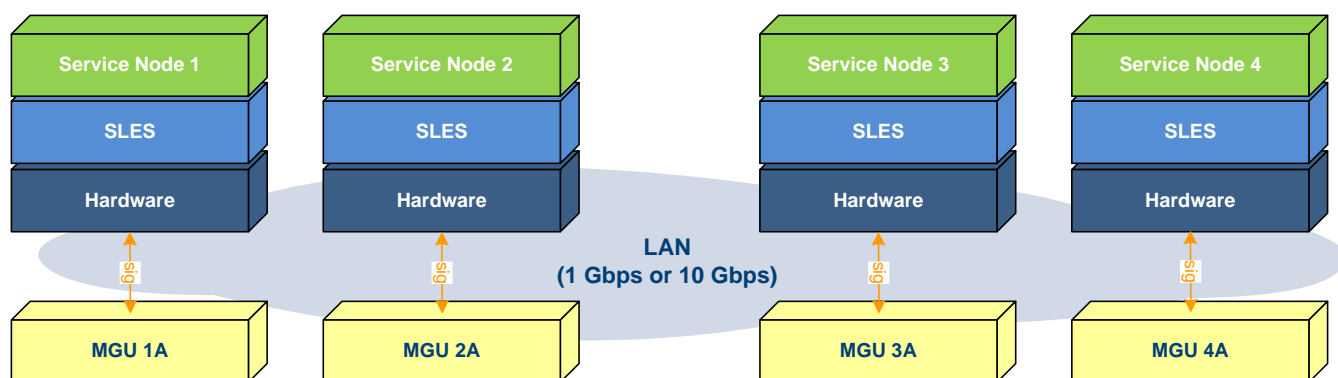
- Consolidated setup where several Service Nodes are consolidated in one or more physical servers. In this case no redundancy is provided by VMware's software.
- Availability setups where several physical servers are consolidated in one or more physical servers. In these solutions, redundancy is provided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

### 4.4.1 CONSOLIDATED SETUP

Mitel validated a consolidation setup where a guest machine running Linux SLES 12 SP3 and where MX-ONE Service Node 6.0 runs on top of a host machine running VMware's vSphere ESXi 5.x hypervisor.

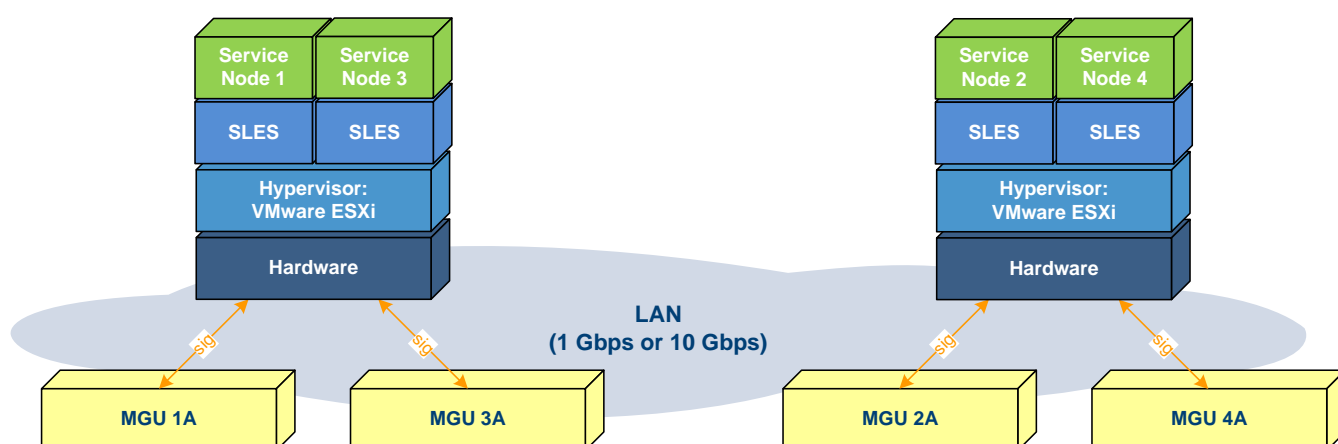
In such a scenario a single physical server or several servers with the appropriate hardware capacity and configuration can be used to reduce the number of physical servers that would be needed in traditional server architecture. In fact, it can reduce hardware footprint as well as maintenance costs.

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.

#### 4.4.1.1 Requirements for consolidation setup

The consolidation setup requires VMware vSphere Hypervisor: ESXi 5.0 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 -vSphere 5.5, ESXi 5.5, vCenter Server 5.5

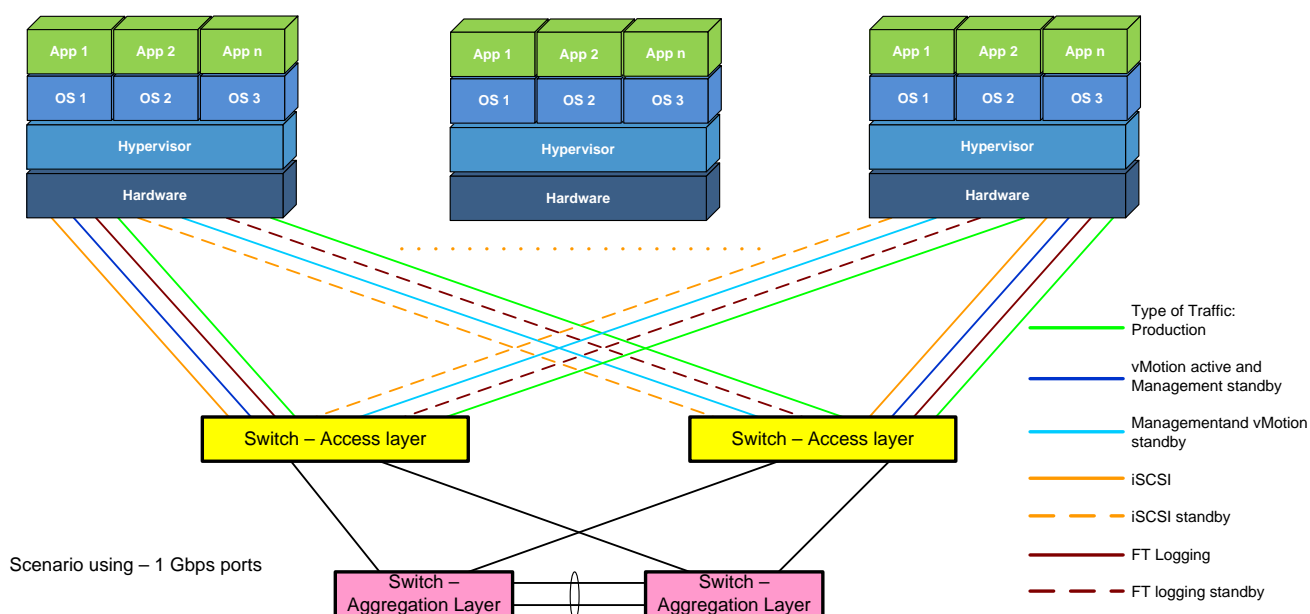
<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-networking-guide.pdf>

#### 4.4.2 AVAILABILITY SETUPS WITH VMWARE AVAILABILITY SOFTWARE OPTIONS

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 6.0, 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. Please check VMware white paper: "VMware Fault Tolerance Recommendations and Considerations on VMware vSphere 4" - the Figure 3 called "Diagram of networking and storage redundancy" shows VMware recommendation for redundant network.

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



**Figure 8 – Example: Availability – traffic type distributed by 1 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.


##### 4.4.2.1 VMware High Availability setup

MX-ONE Service Node 6.0 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 6.0 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 5.0 or later
- 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. The VMware vSphere standard edition should include HA 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 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 5.5

<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-availability-guide.pdf>

vSphere High Availability Deployment Best Practices

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

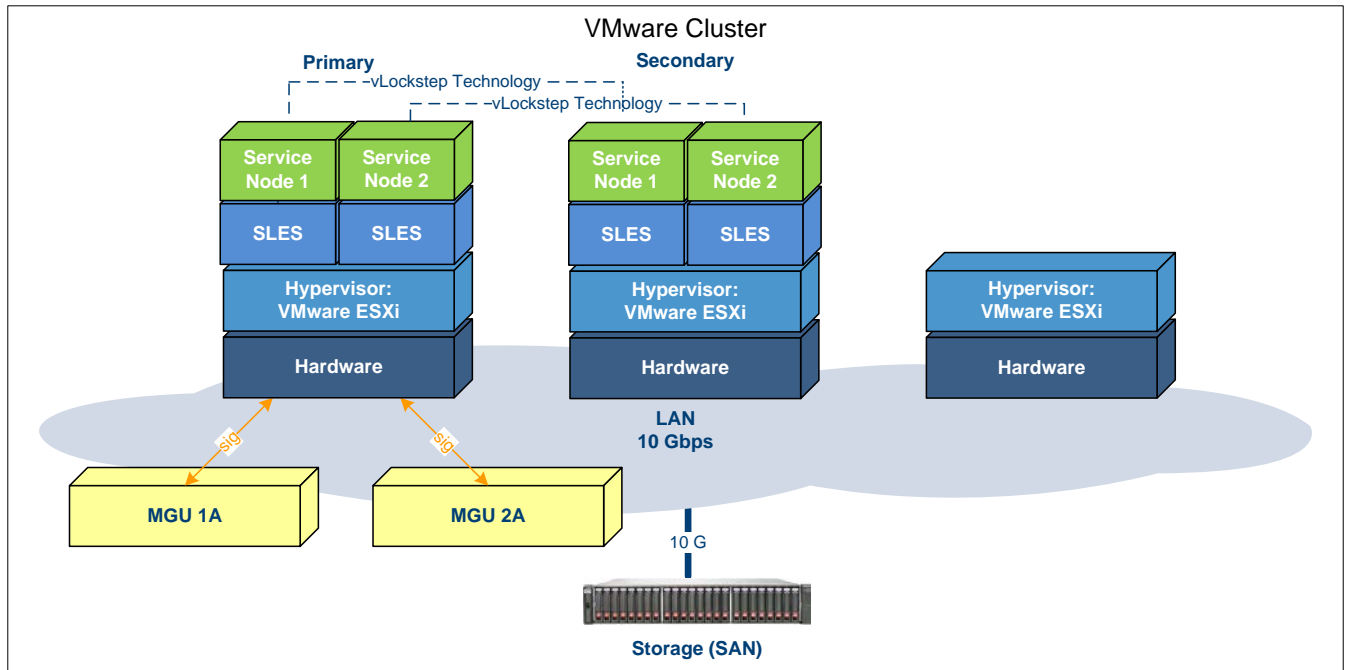
#### 4.4.2.2 *VMware Fault Tolerance setup*

MX-ONE Service Node 6.0 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.

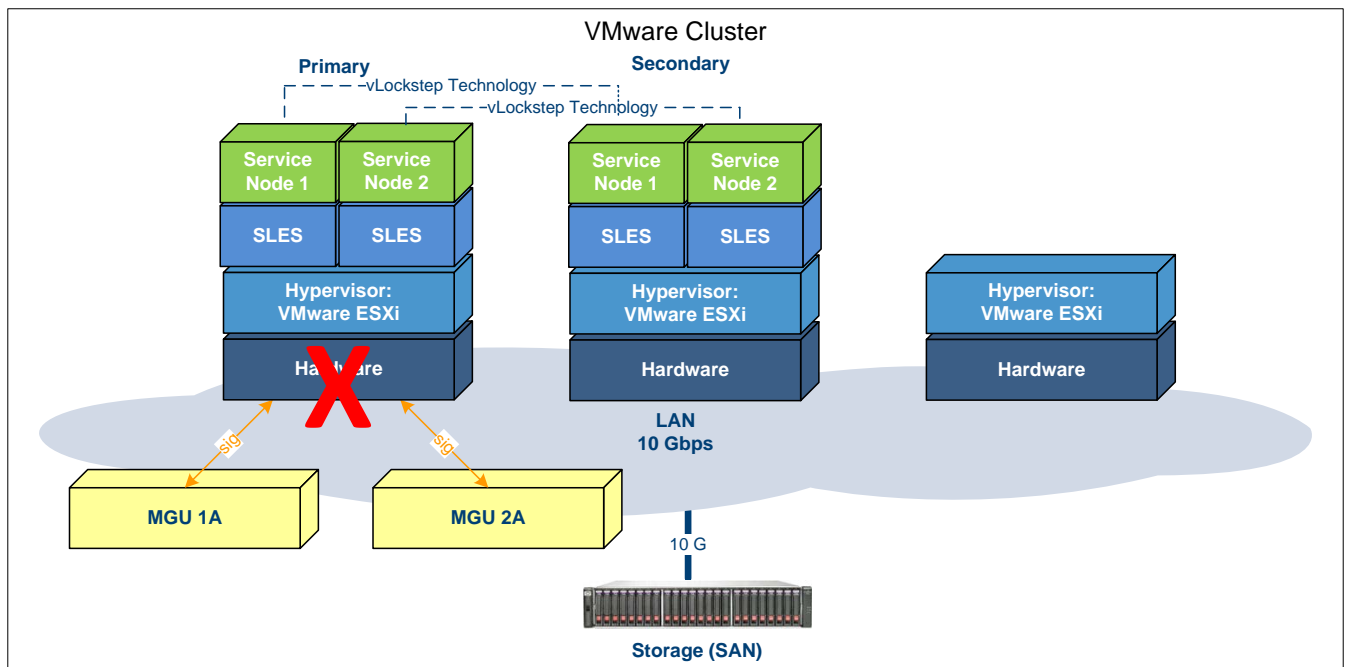
The figure below shows a standard MX-ONE composed of two Service Node guest machines running on top of VMware's FT infrastructure.



**Figure 12 – Two Service Nodes protected by VMware FT**

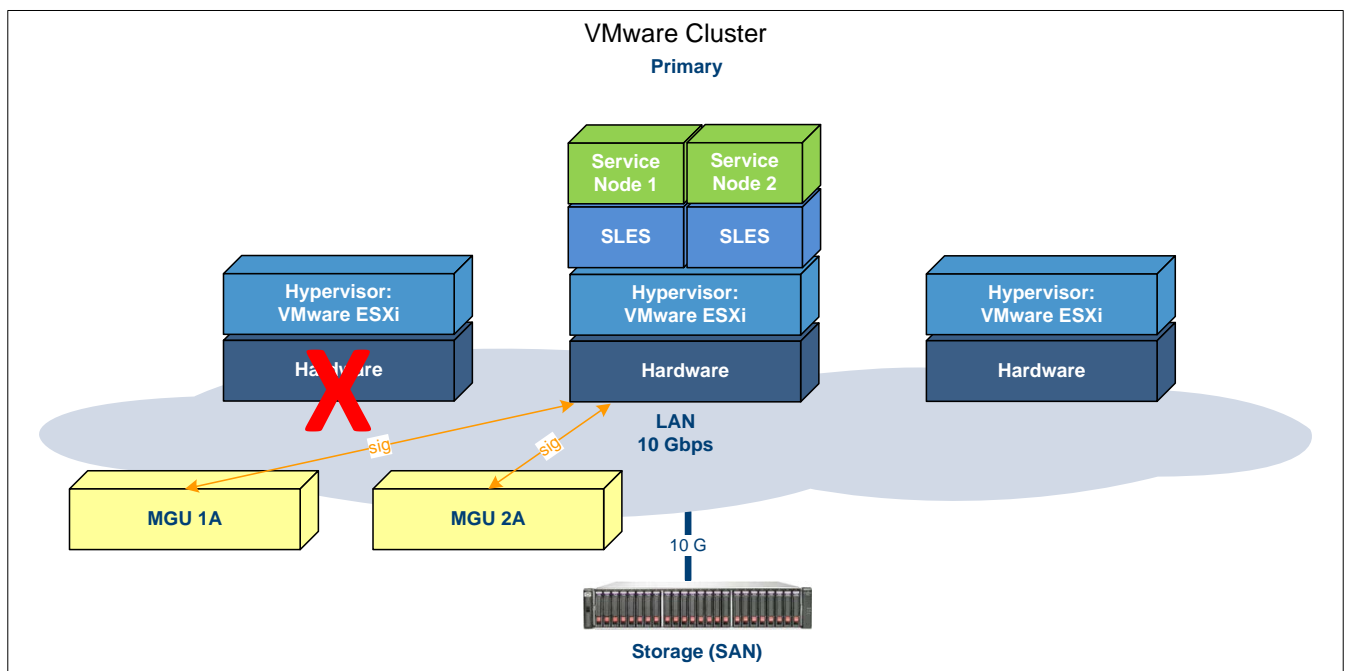
In normal operation, the signaling between Service Node and MGU is transmitted via the “primary physical server”.

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.



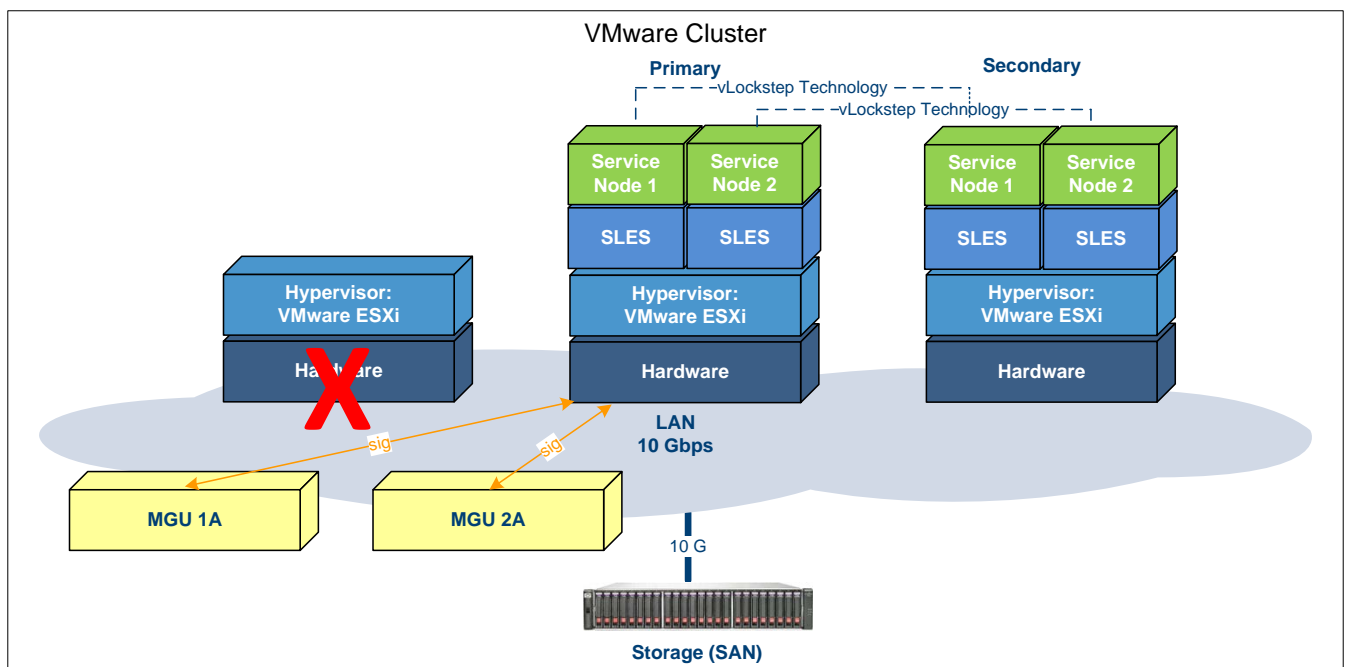
**Figure 13 - VMware FT physical server failure example**

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 a 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 5.5 or later
- 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 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.

Recommendations and limitations with VMware Fault Tolerance



**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:

vSphere Availability ESXi 5.5

<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-availability-guide.pdf>

VMware Fault Tolerance Recommendations and Considerations on VMware vSphere 4

<http://www.vmware.com/resources/techresources/10040>

Some of the VMware’s Fault Tolerance limitations/recommendations are listed below:

- It is recommended that no more than four Primary and Secondary VMs be placed onto the same ESXi host.
- USB devices are not supported with Fault Tolerance.


#### 4.4.3 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:

<http://www.vmware.com/files/pdf/VMware-VMotion-DS-EN.pdf>

## 4.5 VALIDATED SCENARIOS – MEDIA SERVER

In this section MiVoice MX-ONE softswitch is used to refer to a virtual machine that contains MX-ONE Service Node and MX-ONE Media Server.

For Media Server, Mitel has validated two different setups for virtualization on top of VMware's software.

These setups are:

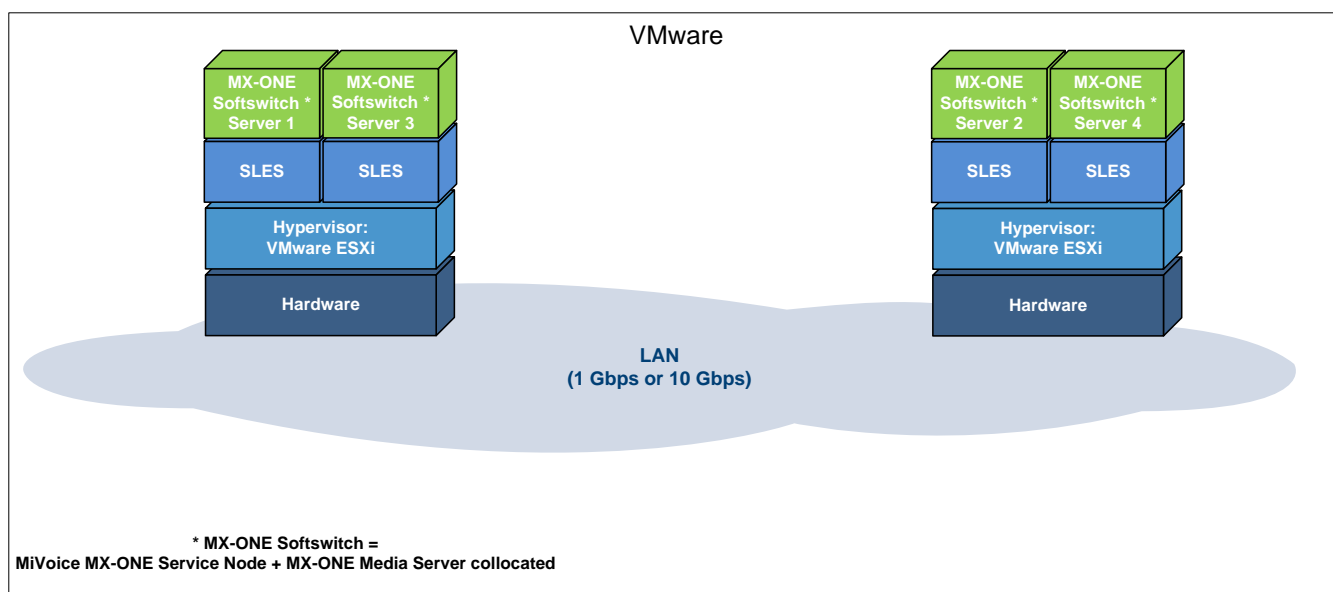
- Consolidated setup where several MX-ONE softswitches are consolidated in one or more physical servers. In this case no redundancy is provided by VMware's software.
- Availability setup, where MX-ONE softswitch physical servers are consolidated in one or more physical servers. In these solutions, redundancy is provided by VMware's software called VMware vSphere High Availability.

### 4.5.1 CONSOLIDATED SETUP

Mitel validated a consolidation setup where a guest machine running Linux SLES 12 and MX-ONE softswitches on top of a host machine running VMware's vSphere ESXi 5.5 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.

The MX-ONE softswitch shown in the picture is composed by four Service Nodes and four Media Servers distributed in two physical servers.



**Figure 19 - MX-ONE with four MX-ONE softswitches**

The disadvantage is that in case one physical host fails, there will be two MX-ONE softswitch nodes out of order. Having multiple MX-ONE softswitches in 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.) components are redundant, the risk can be minimized.

#### 4.5.1.1 Requirements for consolidation setup

The consolidation setup requires VMware vSphere Hypervisor: ESXi 5.0 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.

#### 4.5.2 VMWARE HIGH AVAILABILITY SETUP

MX-ONE softswitch 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 softswitch guest machines running on top of VMware vSphere High Availability allow a cold standby solution, which means that 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 following figure shows a MX-ONE softswitch composed by four Service Nodes and four Media Servers distributed in two physical servers running on top of VMware's HA infrastructure.



**Figure 20 - Four MX-ONE softswitches running on VMware HA**



**Note!** When 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 MX-ONE softswitch guest machines running on top of the failed machine goes down.

The VMware HA mechanism will initialize these MX-ONE softswitch guest machines in another available machine in the VMware cluster environment.

During the re-initialization process, ongoing forced media gateway calls will be dropped, whereas ongoing direct media calls will be maintained.

Additionally, as the VMware HA cluster environment requires a SAN setup, the customer data and management continuity are maintained.

#### 4.5.2.1 Requirements for VMware High Availability setup

- SAN (Storage Area Network) and Network requirements for High Availability according to VMware specifications
- VMware vSphere, Hypervisor: ESXi 5.0 or later
- 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. The VMware vSphere standard edition should include HA 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 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 5.5



<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-availability-guide.pdf>

vSphere High Availability Deployment Best Practices

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

## 5 MX-ONE REQUIREMENTS

### 5.1 MX-ONE REQUIREMENTS

Mitel integrated MX-ONE 6.0 to run on top of VMware 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)
- LSU-E (Media gateway board for older building practice - Stackable).

MX-ONE Media Server

- MX-ONE Server software

### 5.2 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:

#### 5.2.1 BASIC GUIDELINES

Mitel strongly recommends dedicated host machines to MX-ONE Service Node guest machines.

If this is not possible, then it is strongly recommended to avoid mixing MX-ONE Service Node 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.

In general, MX-ONE Service Node as with any real-time application, requires low I/O read and write disk average latency. Although 2ms 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 1Gb 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 (e.g. 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 10Gb link between SANs. This is to ensure that if the site A fails, the storage in the site B can 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.

MX-ONE Service Node Virtual Machines requires minimum ESXi 5.0 Virtual Machine version 8.

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.

It is strongly recommended the use of the Media Gateway Unit (MGU) board housed in the MX-ONE Slim, MX-ONE Lite and/or MX-ONE Classic chassis when running virtualization for several reasons. Firstly, this building practice reduces the hardware footprint of a system as well as power consumption and cooling requirements. Additionally, the MGU comes with dual NIC cards with built in failover capabilities that increase the overall robustness of the system.

MiVoice MX-ONE 6.x has enhanced the Media Server support in the virtualized environment. From this release the Media Server can be installed either on the same server as MX-ONE Service Node or on a separate server.

When running VMware HA and FT, it is necessary that sufficient server resources and priority are configured for MX-ONE Service Node guest machines in a cluster in order to avoid that any failed virtual machine have standby resource precedence over the MX-ONE Service Node 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.

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 setup is the following:

**Host machine hardware minimum for consolidation for 1 MX-ONE Service Node with maximum 1000 users and 2 calls/second using 0,2 Erlang as base:**

- Dell PowerEdge R620 with 1 x CPU @ 2,50 GHz Intel Xeon with 6 cores supporting Hyper-Threading (Intel® Xeon® E5-2640 2.50 GHz, 15 M Cache, 7.2GT/s QPI, Turbo, 6C, 95 W) with minimum 6 GB RAM or an HP DL360 G7 with 1 x CPU @2.4G Hz and 6 cores supporting Hyper-Threading (Intel® Xeon® E5645 (6 Core, 2.4 GHz, 80 W) with minimum 6 GB RAM memory or an equivalent server certified to operate with VMware).
- ESXi Hypervisor, minimum 2 GB RAM.
- Service Node Guest machine, minimum 2 vCPU 1500 MHz, 4 GB RAM and 60 GB HD for the Service Node 1.

For additional Service Node guest machines, add minimum 2 vCPU 1500 MHz, 4 GB RAM and 60 GB HD per extra Service Node guest machine.



**Note!** That is in addition to the 6 GB above.

The minimum configuration for the **High Availability setup** is the following:

**Host machine hardware minimum for High Availability for 1 MX-ONE Service Node:**

- Dell PowerEdge R620 with 2 x CPU @ 2,50 GHz Intel Xeon with 6 cores supporting Hyper-Threading (Intel® Xeon® E5-2640 2.50 GHz, 15 M Cache, 7.2GT/s QPI, Turbo, 6C, 95 W) with minimum 6 GB RAM or an HP DL360 G7 with 2 x CPU @2.50 GHz and 6 cores supporting Hyper Threading with minimum 6 GB RAM memory or an equivalent server certified to operate with VMware.
- ESXi Hypervisor, minimum 2 GB RAM.
- Service Node guest machine, minimum 4 GB RAM and 60 GB HD for the Service Node 1. For the number of vCPUs, please check the table test results in the item 4.2.2.

Additional Service Node guest machines, add minimum 4 GB RAM and 60 GB HD per extra Service Node guest machine.



**Note!** That is in addition to the 6 GB above.

Mitel strongly recommends that the partner/customers read the following VMware's document: Performance Best Practices for VMware vSphere® 5.5

[http://www.vmware.se/pdf/Perf\\_Best\\_Practices\\_vSphere5.5.pdf](http://www.vmware.se/pdf/Perf_Best_Practices_vSphere5.5.pdf)

## 5.2.2 SERVICE NODE GUEST MACHINE TESTED CAPACITY

Mitel has executed performance tests of MX-ONE Service Node 6.0 call manager software running as virtual machines in a VMware 5.5 infrastructure. The table below shows the recommended configuration of virtual CPU (vCPU) and MHz reservations for the virtual machine (VM) running MX-ONE Service Node software. This data is from a server with 2 x CPU

Intel Xeon E5620 2.4 GHz with a total of 8 cores (2 x 4 cores) supporting Hyper Threading. The configuration may be different for other CPU/servers.



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

Mitel recommends two or more virtual CPUs. If the MHz need exceeds 3000, then three virtual CPUs should be configured. If the MHz need exceeds 4500, four or five virtual CPU should be configured.

Virtual Machine configuration setup used in the performance test:

<b>VMware ESXi version</b>	<b>ESXI 5.0</b>
<b>Memory Reserved</b>	4 GB
<b>Memory Limit</b>	4 GB
<b>NIC driver</b>	VMXNET3
<b>VM Hard disk space</b>	60 GB
<b>Virtual Machine Version</b>	8

MX-ONE 6.0 has been verified to handle the following traffic per server VM when running in Virtualized solution. Traffic results are based on 0.2 Erlang.

#### VM Consolidate and in a High Availability setup

Tests results:

Maximum number of users per Service Node Virtual Machine	Recommended number of virtual CPU	Minimum recommended reservation, MHz	SIP traffic capacity with the recommended MHz reservation Traffic rate: 0,2 Erlang
500	2	1000 MHz	1 Call/second
1000	2	1500 MHz	2 Calls/second
2000	2	2000 MHz	4 Calls/second
3000	2	2500 MHz	5 Calls/second
4000	3	3000 MHz	7 Calls/second
5000	4	4000 MHz	9 Calls/second

#### VM in a Fault Tolerance setup

Tests results:

Maximum number of users per Service Node Virtual Machine	Imposed number of virtual CPU by VMware	Minimum recommended reservation, MHz	SIP traffic capacity with the recommended MHz reservation Traffic rate: 0,2 Erlang
500	4 (VMWare 6.5 is required)	2500 MHz	1 Call/second
1000	4 (VMWare 6.5 is required)	3000 MHz	2 Calls/second



**Note!** Backing up data might slow down traffic flow; hence, Mitel recommends that you back up data during low traffic periods.



**Note!** Fault Tolerance setup reading a large amount of system data using a management tool can cause traffic disturbances, so Mitel's recommendation is to read small amount of system data during high traffic hours.

### 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.

Although it is observed that up to 4 vCPUs attributed to an MX-ONE Service Node VM improved performance over 1 vCPU per VM, no significant difference in performance occurs with more than 4 vCPUs per VM.

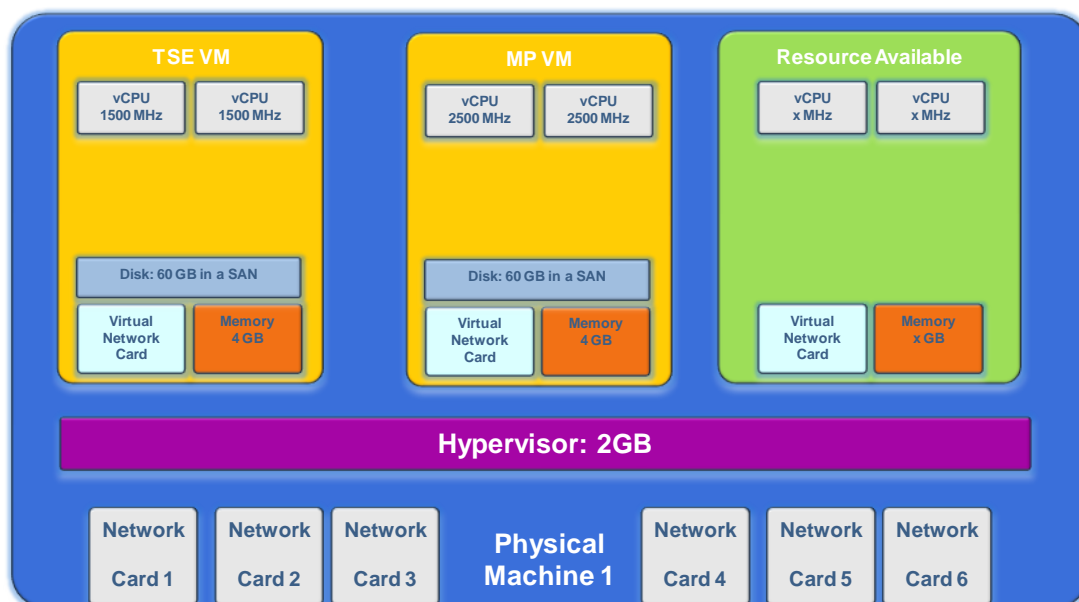
## 5.2.3 SERVICE NODE VIRTUAL MACHINE DIMENSIONING

### Consolidated setup

The following example shows how to dimension a Service Node to run in the VMware environment.

The following information is used as input:

- Number of Service Nodes: 1
- Number of users per Service Node: 1000 users
- Physical machine model: Dell PowerEdge R620 with 1 x CPU @ 2,50 GHz Intel Xeon with 6 cores supporting Hyper-Threading (Intel® Xeon® E5-2640 2.50 GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95 W) with minimum 10 GB RAM



**Figure 21 Example of a consolidated setup with 1000 users**



**Note!** CPU above has 6 physical cores, so the maximum recommended numbers of Service Nodes are 2, when the Manager Provisioning is used in the same physical host (another Virtual Machine).



**Note!** Additional 4 GB is required to run the second Service Node Text.

### High Availability setup

The following example shows how to dimension two Service Nodes to run in the VMware HA environment.

The following information is used as input:

- Number of Service Nodes: 2
- Number of users per Service Node: 5000 users
- 2 x Dell PowerEdge R620 with 2 x CPU @ 2,50 GHz Intel Xeon with 6 cores supporting Hyper-Threading (Intel® Xeon® E5-2640 2.50 GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95 W) with minimum 16 GB RAM

The MX-ONE Service Node Virtual Machine is divided in the two machines, 5000 users for each that requires 4 vCPU (cores).

Manager Provisioning that requires 2 vCPU (cores) is installed in the physical machine 1.

The figures below show the dimensioning result for an HA configuration:

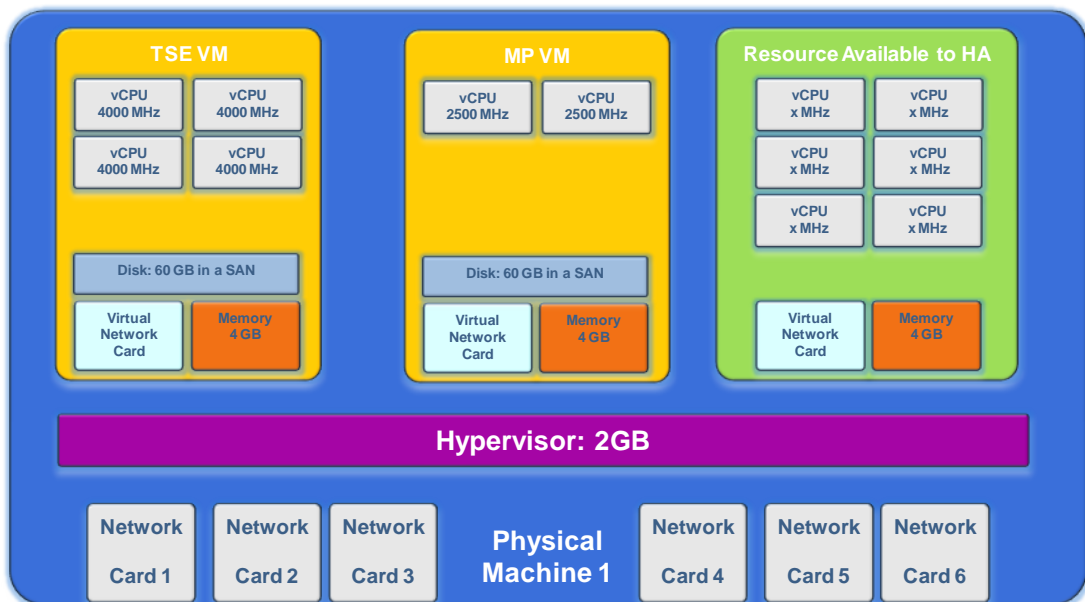
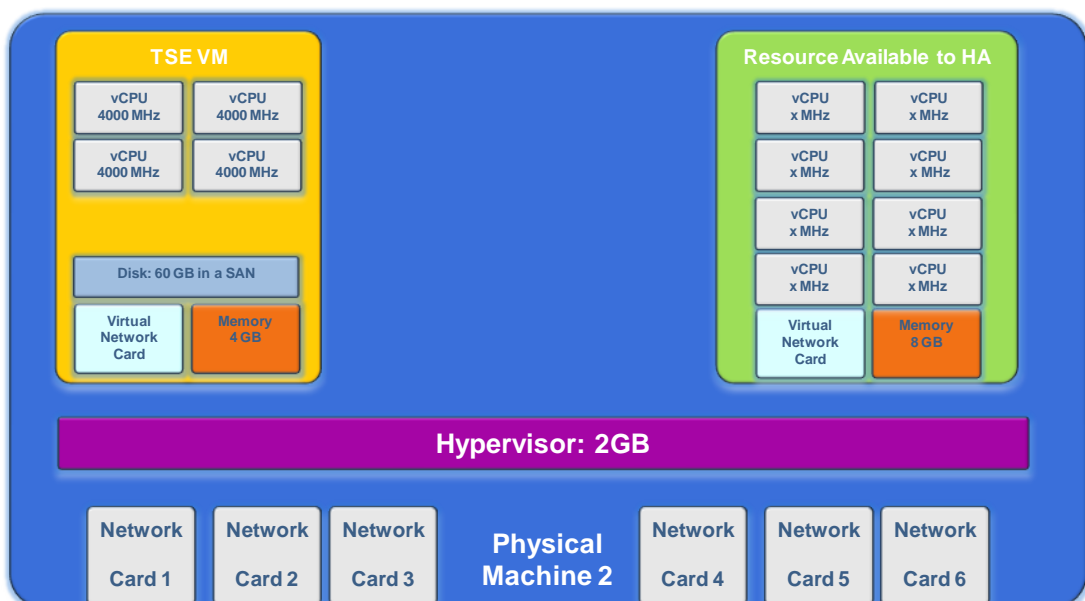


Figure 22 - HA setup 5000 users - physical machine 1



**Figure 23 - HA setup 5000 users - physical machine 2**

As can be seen in the Figures 22 and 23, the green box shows that machines 1 and 2 have enough resources (CPU, memory, etc.) to re-initiate the Service Node and the Manager Provisioning in case of hardware failure in one of the physical machines.

#### 5.2.4 MEDIA GATEWAY CHASSIS DISTRIBUTION PER MX-ONE SERVICE NODE VIRTUAL MACHINE

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

	Consolidation	High Availability	Fault Tolerance
<b>Number of vCPU</b>	minimum 2, maximum 4	minimum 2, maximum 5	4 (VMWare 6.5 is required)
<b>Maximum number of Media Gateways Chassis</b>	up to 15*  *depends on the number of users	up to 15*  *depends on the number of users	Up to 2*  *depends on the number of users

##### 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	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	2	1500 MHz	2 Calls/second	2
2000	2	2000 MHz	4 Calls/second	5
3000	2	2500 MHz	5 Calls/second	7
4000	3	3000 MHz	7 Calls/second	10
5000	4	4000 MHz	9 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.



**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	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 (VMWare 6.5 is required)	2500 MHz	1 Call/second	Up to 2*  *depends on the number of users

1000	4 (VMWare 6.5 is required)	3000 MHz	2 Calls/second	Up to 2* *depends on the number of users
------	----------------------------	----------	----------------	---

### 5.2.5 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 collocated 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.

Maximum number of users per Service Node Virtual Machine	Recommended number of virtual CPU	Minimum recommended reservation, MHz Media Server collocated with Service Node using G.711 codec and Direct Media *	Minimum recommended reservation, MHz Media Server including Service Node	SIP traffic capacity with the recommended MHz reservation Traffic rate: 0,2 Erlang
500	2	500 MHz	1500 MHz	1 Call/second
1000	2	500 MHz	2000 MHz	2 Calls/second
2000	2	500 MHz	2500 MHz	4 Calls/second
3000	2	500 MHz	3000 MHz	5 Calls/second
4000	3	500 MHz	3500 MHz	7 Calls/second
5000	4	500 MHz	4500 MHz	9 Calls/second

\* G.729 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	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	700 MHz	1700 MHz	1 Call/second
1000	2	1100 MHz	2600 MHz	2 Calls/second

\* G.729 codec is not supported in this configuration.

## 5.2.6 MX-ONE SERVICE NODE VIRTUALIZATION BENEFITS

New and current MX-ONE Service Node customers can benefit from a virtualization solution. The benefits for each type of system are described below.

### 5.2.6.1 New MX-ONE Service Node 6.0 with Media Server (MX-ONE softswitch)

A new customer deploying a MX-ONE softswitch with virtualization can benefit from the following:

- No need of MX-ONE hardware
  - Lower power consumption
  - Reduced cooling requirements

- Decreased physical space requirements
- Increased reliability adding VMware High Availability
- The customer data and management continuity is maintained

#### 5.2.6.2 *New MX-ONE Service Node 6.0 with MGU2*

A new customer deploying a Service Node with virtualization can benefit from the following:

- Chassis infrastructure savings (inherent of MGU2 board gateways)
  - Lower power consumption
  - Reduced cooling requirements
  - Decreased physical space requirements
  - Increased network redundancy options
- Increased reliability adding VMware High Availability and Fault Tolerance Solutions
- Decreased number of servers when running VMware HA if compared with MX-ONE Server Redundancy solution
- The customer data and management continuity is maintained

#### 5.2.6.3 *Upgrade from V.3.2, V.4.0, V.4.1, (5.0) or Migration from TSW*

A MX-ONE upgrading from V.3.2, V.4.0, V.4.1, (5.0) or migrating from TSW can benefit from the following:

- MX-ONE Service Node 6.x with MGU board (hardware replacement) using MX-ONE Lite or/and MX-ONE Classic chassis:
- Infrastructure enhancements provided with Media Gateway Unit (MGU2) board gateways
  - Lower power consumption
  - Reduced cooling requirements
  - Decreased physical space requirements
- Increased network redundancy options
- Reduced server and telephony hardware maintenance costs
- Increased resilience by adding VMware High Availability and Fault Tolerance Solutions

## 5.3 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.

## 5.4 MX-ONE SERVICE NODE SOFTWARE

The Service Node is also available as a virtual appliance.

The virtual appliance is an ESXi 5.0 Virtual Machine Version 8 and the based appliance/image is not compatible with previous versions.

## 5.5 CREATING SNAPSHOT ON VIRTUAL MACHINE RUNNING THE MX-ONE SERVICE NODE SOFTWARE

Use the command `data_change` to prevent MX-ONE Service Node data changes while creating VM snapshot.

## 5.6 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.

## 5.7 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.

The VM requirements for Provisioning Manager are:

- Minimum of 2 GB of RAM memory
- Minimum 2 Virtual CPUs with 2,5 GHz
- Minimum 20 GB of hard disk per 5,000 users

## 5.8 RECOMMENDED TRAINING

Partners that will install the MiVoice MX-ONE 6.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 VMware training is:

- VMware vSphere: Install, Configure, Manage [V5.5]

[http://mylearn.vmware.com/mgrreg/courses.cfm?ui=www\\_edu&a=one&id\\_subject=46424](http://mylearn.vmware.com/mgrreg/courses.cfm?ui=www_edu&a=one&id_subject=46424)

## 5.9 PRODUCT BUSINESS APPROVAL (PBA)

Although, VMware certification is not obligatory to obtain PBA for the MiVoice MX-ONE 6.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.

## 6 ACRONYMS, ABBREVIATIONS AND GLOSSARY

**FT** – Fault Tolerance

**Guest Machine** – The same as Virtual Machine

**HA** – High Availability

**HD** – Hard Disk

**Host Machine** – Physical machine that executes the hypervisor

**LSU-E** – LIM Switch Unit Ethernet

**MGU** – Media Gateway Unit

**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

**NIC** –Network Interface Card

**TDM** – Time Division Multiplex

**SAN** – Storage Area Network

**VCP** – VMware Certified Professional

**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 its 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”*

## 7 REFERENCE DOCUMENTS

### **MX-ONE 6.x CPI documentation**

MX-ONE Media Server Description, part of the MX-ONE 6.x CPI documentation

### **VMware documentation**

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

Below are some of the VMware documents that were used as reference in this description:

Protecting Mission-Critical Workloads with VMware Fault Tolerance

<http://www.vmware.com/resources/techresources/1094>

[http://www.vmware.com/files/pdf/resources/ft\\_virtualization\\_wp.pdf](http://www.vmware.com/files/pdf/resources/ft_virtualization_wp.pdf)

VMware vSphere Basics Guide – ESXi 6.0

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-basics-guide.pdf>

vSphere Availability ESXi 5.5

<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-availability-guide.pdf>

vSphere High Availability Deployment Best Practices

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

VMware Fault Tolerance Recommendations and Considerations on VMware vSphere 4

<http://www.vmware.com/resources/techresources/10040>

vSphere Networking -vSphere 5.5, ESXi 5.5, vCenter Server 5.5

<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-networking-guide.pdf>

Performance Best Practices for VMware vSphere™ 5.5

[http://www.vmware.se/pdf/Perf\\_Best\\_Practices\\_vSphere5.5.pdf](http://www.vmware.se/pdf/Perf_Best_Practices_vSphere5.5.pdf)

VMware vSphere: Install, Configure, Manage [V5.5]

[http://mylearn.vmware.com/mgrreg/courses.cfm?ui=www\\_edu&a=one&id\\_subject=46424](http://mylearn.vmware.com/mgrreg/courses.cfm?ui=www_edu&a=one&id_subject=46424)

VMware vSphere: Install, Configure, Manage

- Student Manual – Volume 1 and 2
- ESXi 5.0 and vCenter Server 5.0

VMware vSphere Metro Storage Cluster Case Study

<http://www.vmware.com/resources/techresources/10299>

Stretched Clusters and VMware vCenter™ Site Recovery Manager

<http://www.vmware.com/resources/techresources/10262>

Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs

<http://www.vmware.se/files/pdf/techpaper/VMW-Tuning-Latency-Sensitive-Workloads.pdf>

Deploying Extremely Latency – Sensitive Applications in VMware vSphere 5.5

<http://www.vmware.se/files/pdf/techpaper/latency-sensitive-perf-vsphere55.pdf>

Network I/O Latency on VMware vSphere® 5

<http://www.vmware.com/files/pdf/techpaper/network-io-latency-perf-vsphere5.pdf>

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

<http://www.vmware.se/files/pdf/techpaper/voip-perf-vmware5.pdf>

VMware VMotion

<http://www.vmware.com/files/pdf/VMware-VMotion-DS-EN.pdf>

VMware vSphere with Operations Management and VMware vSphere

Licensing, Pricing and Packaging

<http://www.vmware.com/files/pdf/products/vsphere/VMware-vSphere-Pricing-Whitepaper.pdf>