

MiVoice MX-ONE

Engineering Guidelines

Release 7.3 SP1
January 27, 2021



Notice

The information contained in this document is believed to be accurate in all respects but is not warranted by **Mitel Networks™ Corporation (MTEL®)**. 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>.

®,™ Trademark of Mitel Networks Corporation
© Copyright 2021, Mitel Networks Corporation
All rights reserved

Contents

Introduction	1
Scope of this Document.....	1
Architecture	2
MiVoice MX-ONE Architecture	2
Database Node Design.....	6
IP Protocol	6
System Requirements	8
General MX-ONE Requirements	8
Hardware Requirements.....	8
Software Requirements	9
Network Services Requirements.....	10
Service Node Requirements	10
Database Node Requirements.....	10
Provisioning Manager Requirements.....	10
Media Server Requirements	11
Media Gateways Requirements.....	12
Deployment Scenarios	13
Scenario 1 – 5,000 SIP Extensions.....	13
Scenario 2 – 15,000 SIP Extensions.....	14
Scenario 3 – 81,000 SIP Extensions.....	15
Redundant MX-ONE Systems	16
Standby Server and Database Nodes	17
Media Server Redundancy and Load Balancing	17
Home Location Register (HLR) Redundancy.....	18
Performance MX-ONE Systems	19
Service Node.....	19
Media Server	19
Virtualization	20
VMware	20
KVM	20
Hyper-V.....	20
High Availability Networks	21
Design a High Availability Network	21
Security	23
Application Security.....	23

MX-ONE and MiCollab.....	25
Appendix A	26
Database Nodes - MX-ONE 6.3 SPX Migration Path to MX-ONE 7.1 and Later	26
General Database Guidance	26
One Site Scenario, One Datacenter	26
One Service Node	26
Two Service Nodes.....	27
Three Service Nodes	28
Four Service Nodes	30
Ten Service Nodes.....	31
Two Sites Scenario, Two Datacenters	35
Four Service Nodes	35
Ten Service Nodes.....	36
Six Sites Scenario, Six datacenters	38
Fourteen Service Nodes.....	38
Appendix B.....	41
Service Nodes and Database Node Distribution in a Virtual- ized Environment	41
Database in Network Storage	41
Acronyms, Abbreviations and Glossary	42
Reference Documents	43

Introduction

MiVoice MX-ONE unified communication system has been updated with a new version of the operating system SUSE Linux Enterprise Server 12 SP5 (SLES 12 SP5) and application server Jboss, now called WildFly.

Operating system is used by all MiVoice MX-ONE components, e.g. Service Node, Media Server, Provisioning Manager, Service Node Manager and system database node.

The application server is used by Provisioning Manager and Service Node Manager.

These two main components bring latest technology and more security options to the MX-ONE system. The changes done for these two main components are under the hood, so it is not so visible to a partner/customer, however the benefits that it brings to the solution can be perceived as the company start to use the system.

Additionally, a new database is introduced, it is called Cassandra and it replaces the Open-LDAP. The main differences between the databases is the fact that Cassandra is a multi-master database instead of master-slave as Open-LDAP and that there is no need to have database nodes per each Service Node.

Scope of this Document

The aim of this document is to provide guidelines on how to deploy the MiVoice MX-ONE. Detailed information regarding the functions mentioned in this document can be found in the MX-ONE CPI documentation.

Architecture

MiVoice MX-ONE Architecture

The MiVoice MX-ONE system is composed by the following major components:

- Service Node (SN)
- Media Server (MS)
- Media Gateway (MGU)
- System database node (Cassandra database)

MiVoice MX-ONE is managed by the following management system components:

- Provisioning Manager (PM)
- Service Node Manager (SNM)

The MiVoice MX-ONE components are fully described in the MiVoice MX-ONE System Description and they will not be described in this document.

The figures 1 and 2 below show high-level architecture views of MiVoice MX-ONE system.

- Service Node communicates with other Service Nodes via SCTP protocol (inter-lim signaling).
- Service Node communicates with Media Server and Media Gateway via SCTP protocol.
- Service Node communicates with system database server via an internal C++ driver.
- Provisioning Manager communicates with Service Node Manager via Web Services.
- Service Node Manager communicates with Service Node 1 via XML using ACS appServ (internal component), command line and direct calls to system database using a Java driver.
- System database nodes communicate with each other with via Cassandra internal protocol.

Figure 2.1: MX-ONE High level view - Single Service Node

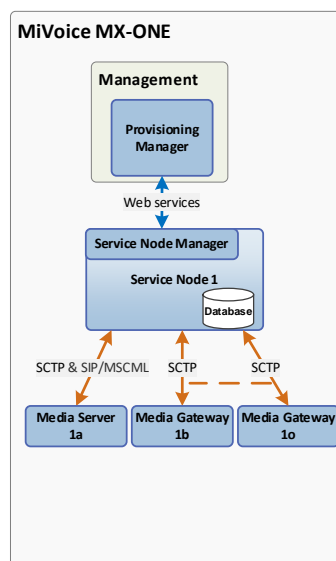
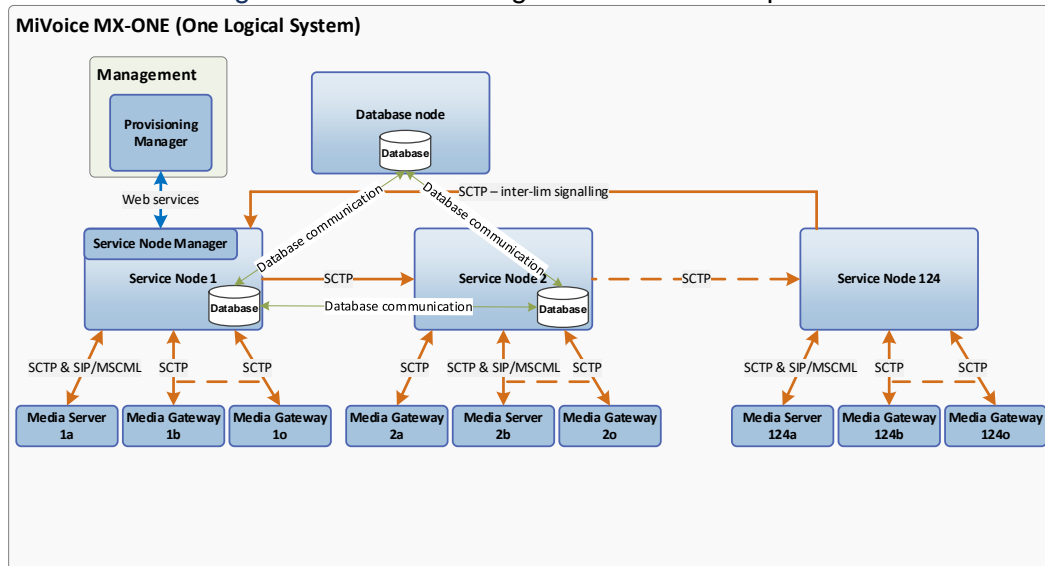


Figure 2.2: MX-ONE High level view - Multiple Service Nodes

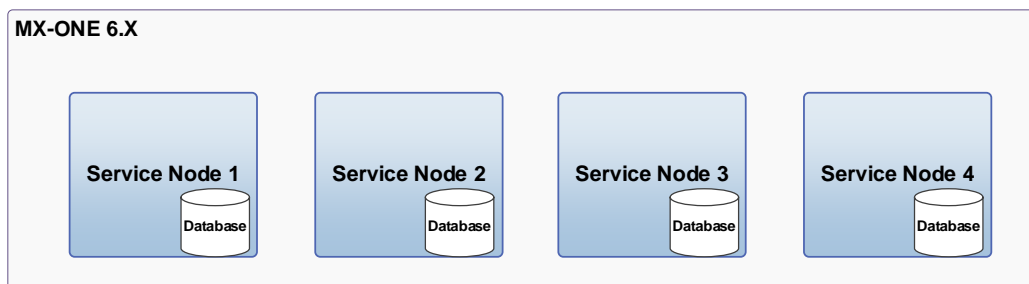


Architecture Changes in MX-ONE 7.0 Onwards

From MX-ONE 7.0 release onwards, a new database is introduced which change the architecture of MiVoice MX-ONE.

In MX-ONE 6.X, each Service Node had its own database and the system used to work in a master-slave relationship. The master database was installed in the Service Node 1, figure 3 shows an example of a MX-ONE 6.X system.

Figure 2.3: MX-ONE 6.X high-level view



The Cassandra database is distributed, decentralized and highly available, which are some of the most important requirements on MX-ONE, so changes were done to use such capabilities together with the MX-ONE system, improving the system availability and reliability.

From MX-ONE 7.0 release onwards, a new multi master concept is used, which means that each database can act as a master, because they might contain all the data required for the system to work.

The requirement to have one database per Service Node is no longer valid, indeed the number of database nodes in the system is reduced and the system setup is done differently when compared with previous versions of MX-ONE. However, in systems with 1 to 3 Service Nodes the number of databases is almost the same as in MX-ONE 6.X.

A new type of server is introduced, called system database (Cassandra) node.

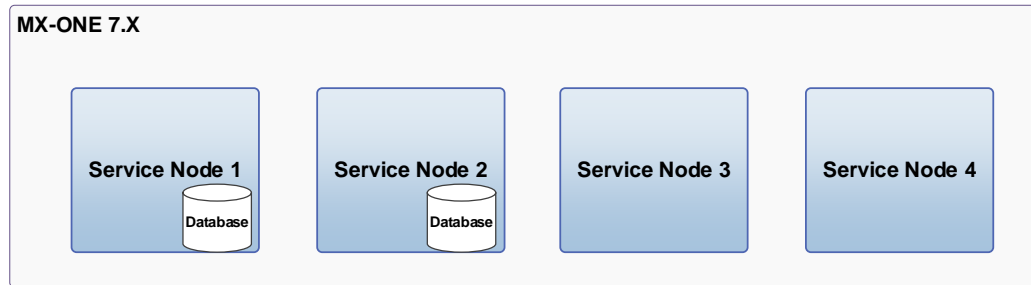
The database node can be deployed co-located with the Service Node or as a stand-alone server.

The stand-alone server runs the same SLES 12 as the Service Node and it is defined in the MX-ONE system installation as the other servers, e.g. Media Server and Provisioning Manager.

The choice of using co-located or stand-alone depend of the system size and the required level of redundancy for the MiVoice MX-ONE system.

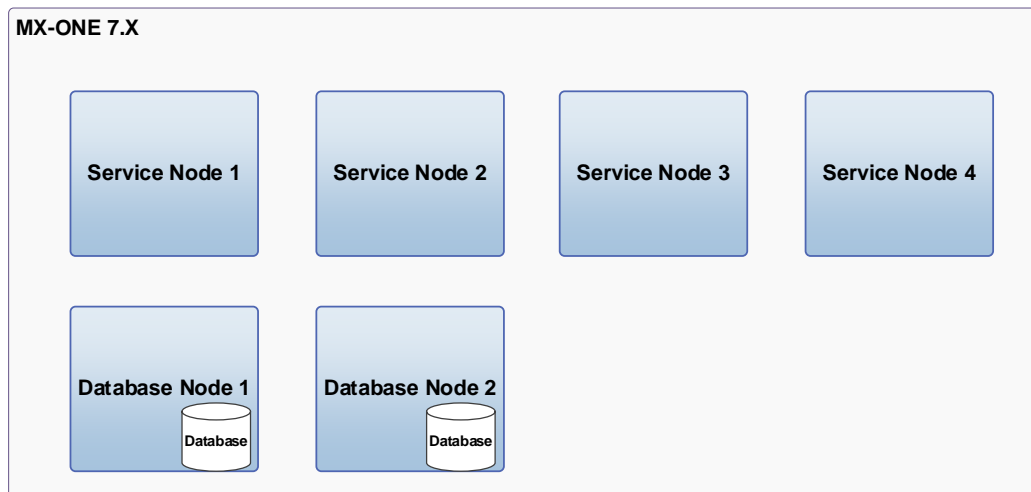
The figure 4 shows a high-level view of an MX-ONE with co-located database nodes.

Figure 2.4: MX-ONE with co-located database nodes



The figure 5 shows a high-level view of an MX-ONE with stand-alone database nodes.

Figure 2.5: MX-ONE with stand-alone database nodes



System Database Data Center and Rack Concept

Another important difference in MX-ONE 7.0 and later releases, is the fact that the database nodes are allocated within the concept of Data Centers (physical locations) and racks (physical racks) located inside of a Data Center. It is assumed that each data center has infrastructure that leverage redundant elements, avoiding single point of failures.

Each geographic location is a Data Center, so for remote sites, a Data Center is required.

During the MX-ONE installation you must answer questions about the quantity of Data Centers and Racks will be used in the system.

The figure 6 shows the Data Center and Rack concept used by the Cassandra database.

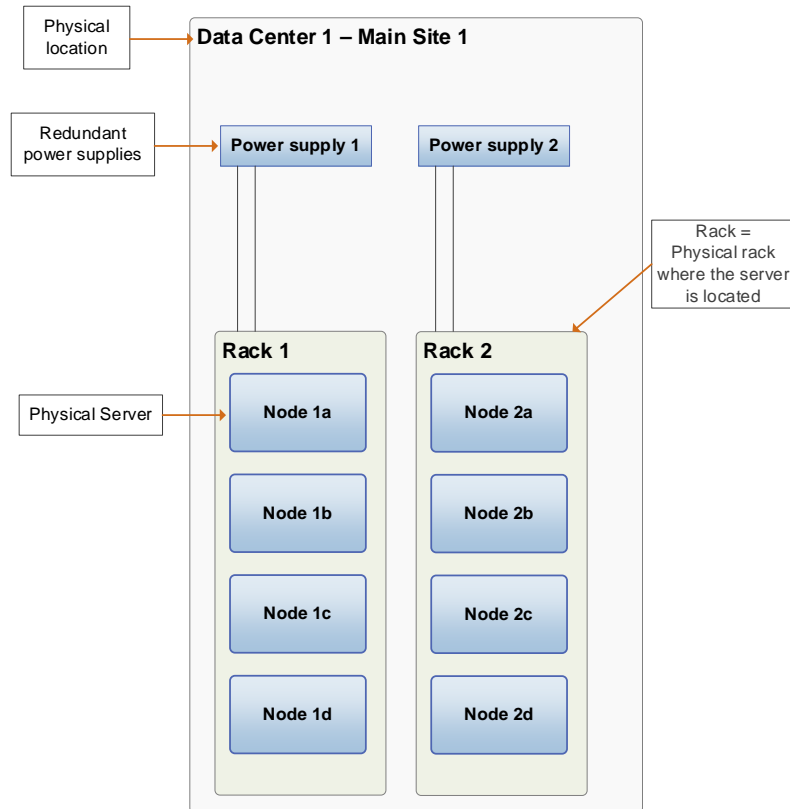


Figure 6 - Data Center and Rack concept

The figure 7 shows an example of a system deployed in 2 data centers, where 6 Service Nodes and 2 Service Nodes standby installed in 8 servers, distributed in 4 racks.

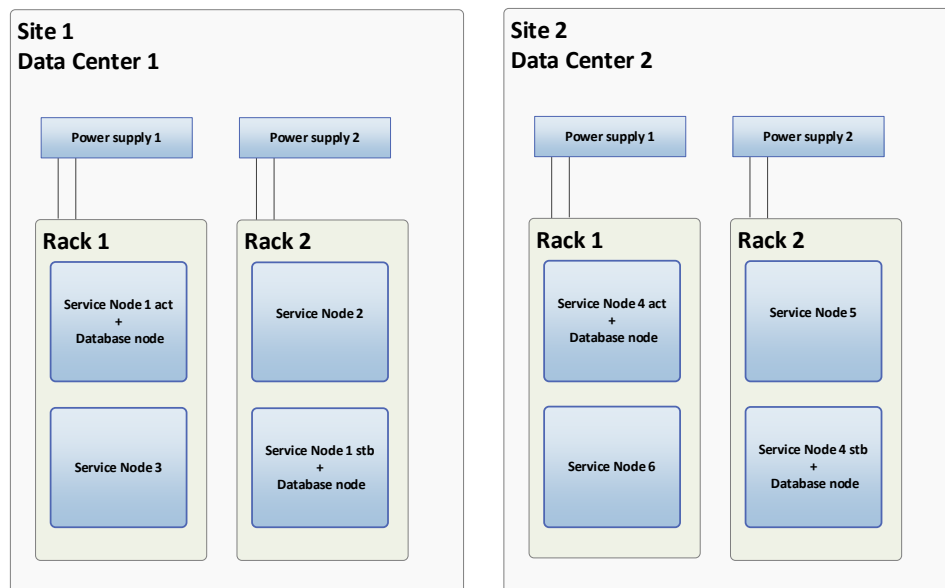


Figure 7 - Service Nodes and Database Nodes distribution

For a detailed description of the MX-ONE database, please read the MX-ONE CPI document System Database (Cassandra) Description.

Database Node Design

The database nodes in MX-ONE can be deployed in co-located or stand-alone mode. The choice of using one or the other depend of the system size and the desired level of availability of the MX-ONE system.

The MX-ONE Service Nodes communicate with the database nodes via a driver and the database nodes communicate with each other using it owns protocol.

The figure below shows a high-level view of the communication between MX-ONE and database nodes and databases nodes.



Note! Database node requires unique IP address.

Database node – High Level communication view – single site

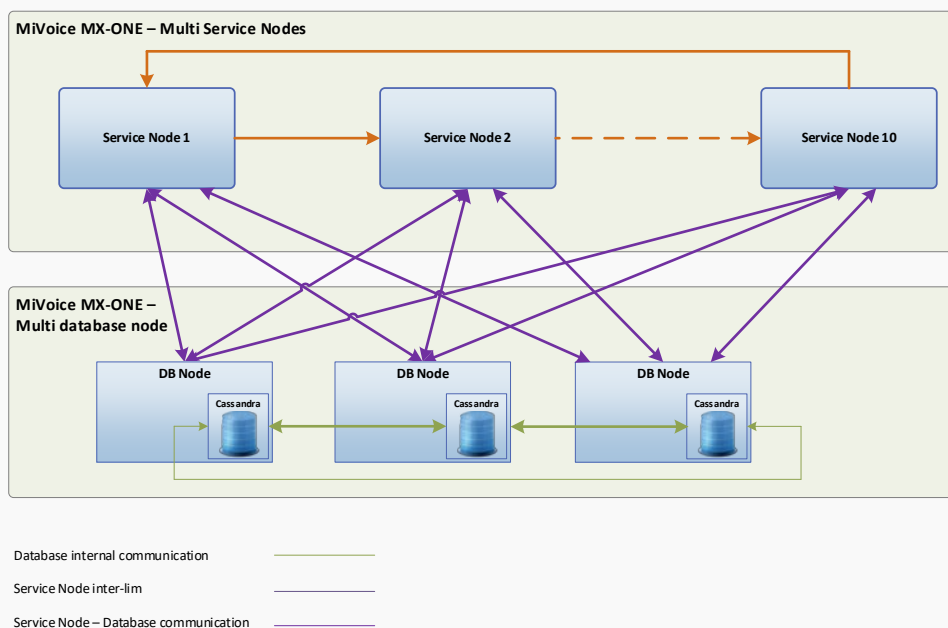


Figure 8 - High level view of the database node communications

IP Protocol

MiVoice MX-ONE 7.0 and later releases supports IPv4 or IPv4/IPv6 in dual mode.

Please note that when MX-ONE is configured with dual mode, the communication between the Service Node and Database nodes is done in IPV6 only. So, it is important to make the correct IP Protocol choice before setup the system.

The figure below shows how the different MX-ONE components interwork with IPv4-IPv6.

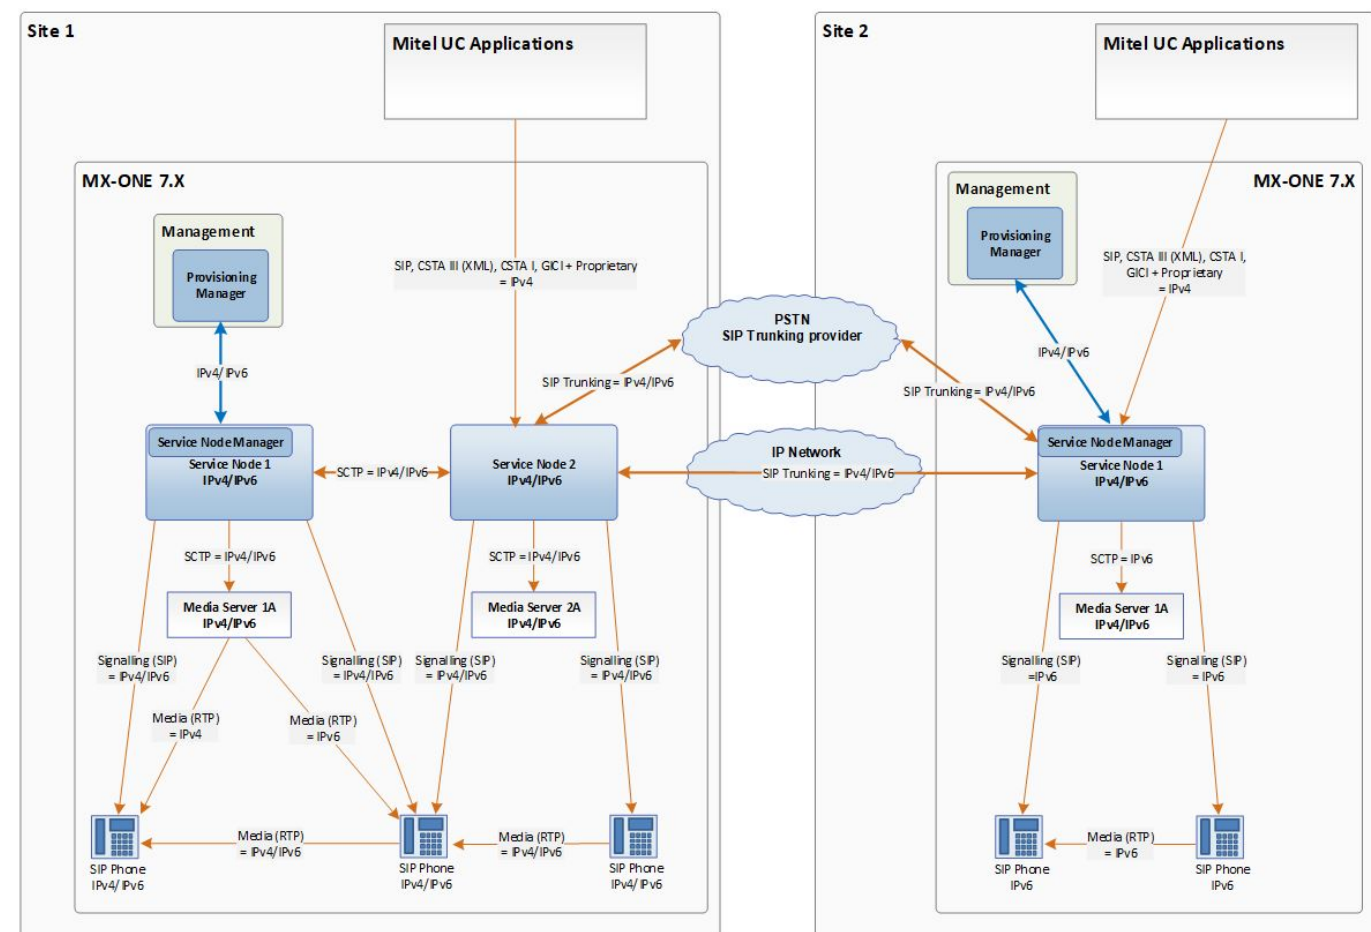


Figure 9 - MX-ONE IP Protocol view

System Requirements

General MX-ONE Requirements

MX-ONE system requires a good system planning, refer for *MX-ONE System Planning* document for further information.

The MX-ONE system requirements for servers is changed from MX-ONE 7.0 release onwards and it depends of the system size as well as the application running on the server. It is important that the MX-ONE Capacity and the Virtualization description documents are read carefully before deploying a new MX-ONE.

Note: When deploying an MX-ONE system with a large amount of extensions in one Service Node it is important to consider redundancy, it can be achieved via MX-ONE 1+1 redundancy or VMware High Availability.

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

Hardware Requirements

The Operating System software used by MX-ONE, SLES 12 SP5 is patched with the latest security software available to date. Additionally, the firmware used by the servers (physical or virtual machines) shall be updated according to the hardware or virtualization supplier. These security updates were performance tested in Mitel's labs and the test results concluded that there is performance degradation, which results in less calls per second depending of the hardware used. The number of calls described in capacity document takes the new values in consideration.

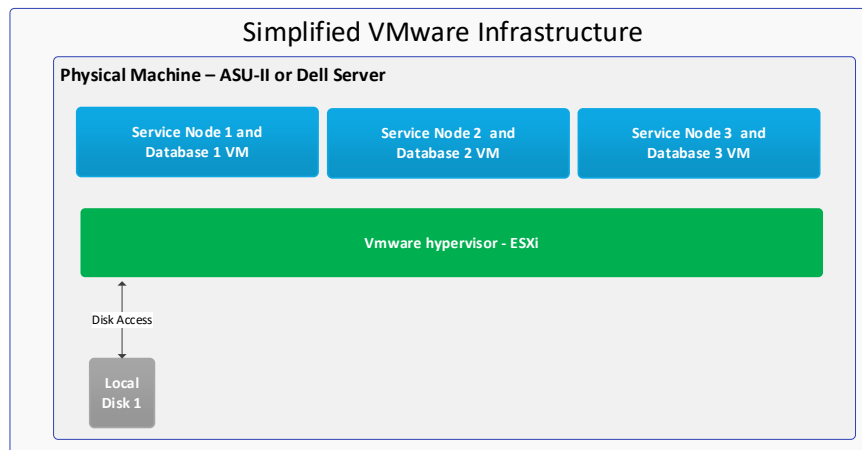
Please note that some servers are not supported from MX-ONE 7.0 release onwards, due to performance limitations, refer to the capacity document to get the list of not supported servers.

It is highly recommended to use Industrial Standard Servers using Intel processors, family Skylake or higher.

It is highly recommended to use Industrial Standard Servers for larger and very large systems for Service Nodes and database nodes (both co-located and stand-alone).

It is highly recommended to use SSD disk to MX-ONE database instances (Cassandra).

When MX-ONE is deployed in a hypervisor environment, e.g. VMware ESXi free version, and the MX-ONE data is stored in local disks, instead of Storage Area Networks (SAN), it is highly recommended to use SSD disks, at least in the VMware host that have a Cassandra database on it, this is to avoid any disk latency issues in a Multi-VM environment. Please see the scenario below:



Disk performance (latency). In general, MX-ONE 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 10 ms, 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. This is valid for systems running local disk and SAN.

Raid Controller with at least 1GB cache it is highly recommended.

Database stand-alone servers are recommended to be used to Large (above 3,000 SIP extensions) and Very Large (above 10,000 SIP extensions) systems.

Dell Servers or Industrial Standard Servers

The minimum recommended Dell server to run MX-ONE is a server equipped with Intel Xeon E series 3.x GHz CPU or Xeon Silver 4215 2.x GHz with 4-8 cores equivalent or higher (family Skylake or higher), 16-64GB of RAM memory, HW Raid with at least Raid 1 enabled (Raid Controller must have at least 1GB cache), Dual power and 2 x 300GB SSDs, or storage with equivalent performance, redundant 1 GB or 10 GB Ethernet cards.

The server shall support at least Linux SUSE Enterprise Server 12 SP5.

When deploying a SIP system with huge number of extensions, an Industrial Standard Server or a Dell server is recommended, because they offer the redundant fans units, and can be provided with additional hard disks in a RAID configuration, CPU, Memory, Ethernet cards and additional Power Supply Units (PSU).

It is strongly recommended to always include the additional hard disk (at least RAID1) and the additional PSU when choosing and configuring the MX-ONE servers.

ASU-II Servers

The ASU-II is a card designed to be used in the MX-ONE chassis where the total number of extensions does not exceed 7,500 extensions.

The ASU-II can be used in Multi-Service Nodes systems, however the total number of users in the system shall not exceed 25,000 extensions.

Ideally a large system using ASU-II shall not have more than 10 Service Nodes in a single logical system, however there is no limitation when the ASU-II card is used in a standalone Service Node using SIP tie-lines between Service Nodes.

Software Requirements



Note! Provisioning Manager (PM) and the MX-ONE database (Cassandra) uses Java technology, as they require large amount of memory and CPU, the recommendation is to use them separate for systems above 1000 users.

MX-ONE 7.3 release requires the SUSE Linux Enterprise Server 12 SP5 (SLES 12 SP5) operating system. All other software required by MX-ONE and MX-ONE applications are delivered by MX-ONE installation packages, e.g. .bin, .ova, .iso., qcow2 (KVM).

Network Services Requirements

NTP – Network Time Protocol

An accurate NTP setup is essential to MX-ONE. The MX-ONE Service Nodes and database nodes must run under milliseconds time precision to avoid data inconsistency.

SNTP (Simple Network Time Protocol) is not supported.

DNS – Domain Name System

MX-ONE requires a DNS setup for name resolution.

Service Node Requirements

MiVoice MX-ONE supports up to 124 Service Nodes as one single system, however in order to simplify the system maintenance, updates and upgrades, it is recommended to deploy more than one Very Large system when more than 100 K (100,000) SIP extensions is required. These systems can be connected by SIP trunkings forming a large network as there is no limit to connect systems via SIP trunking.

Generally, less Service Nodes are better for performance and maintenance, so if the system requires 30,000 SIP extensions, for example deploy a system with 4 Service Nodes with 7,500 SIP extensions in each, instead of 10 Service Nodes with 3,000 SIP extensions on each. Of course, that every customer has a different demand and the customer requirements must be taken in consideration without compromise the MX-ONE system deployment.

The MX-ONE Capacity and the Virtualization documents have the requirements for Service Node 1 with applications, Other Service Nodes with Provisioning Manager, Other Service Nodes and only Service Nodes, please read the documents before deploying a new MX-ONE solution.

Database Node Requirements

MX-ONE requires minimum 1 database node, however for large and very large system, a minimum of 3 database nodes are required to secure a high level of data consistency and availability.

The server requirements regarding RAM memory, number of CPU, disk size to a database node shall cover the full system capacity, it is per system and not per node.

In virtualized environment the recommendation is a maximum of two database nodes per physical servers.

For example, if the MX-ONE system has 4 Service Nodes with the total of 5,000 users, each Service Node with 1,250 users, the minimum database memory shall be 10 GB RAM with 4 CPUs (capacity/virtualization document), for both co-located and stand-alone cases. It means that Service Nodes running database nodes require more resources than a Service Node without database node.

The MX-ONE Capacity and the Virtualization documents have the requirements for database nodes both co-located and stand-alone, please read the documents before deploying a new MX-ONE solution.

Provisioning Manager Requirements

Provisioning Manager offers single point of entry for user management.

Provisioning Manager supports several Service Nodes systems, so in case the customer has a network with 3 MX-ONE systems, only one Provisioning Manager is required.

The MX-ONE systems are created as sub-systems in Provisioning Manager. So, for the administration perspective all the user management is done in Provisioning Manager.

Provisioning Manager will communicate with the Service Node Managers located in each MX-ONE system.

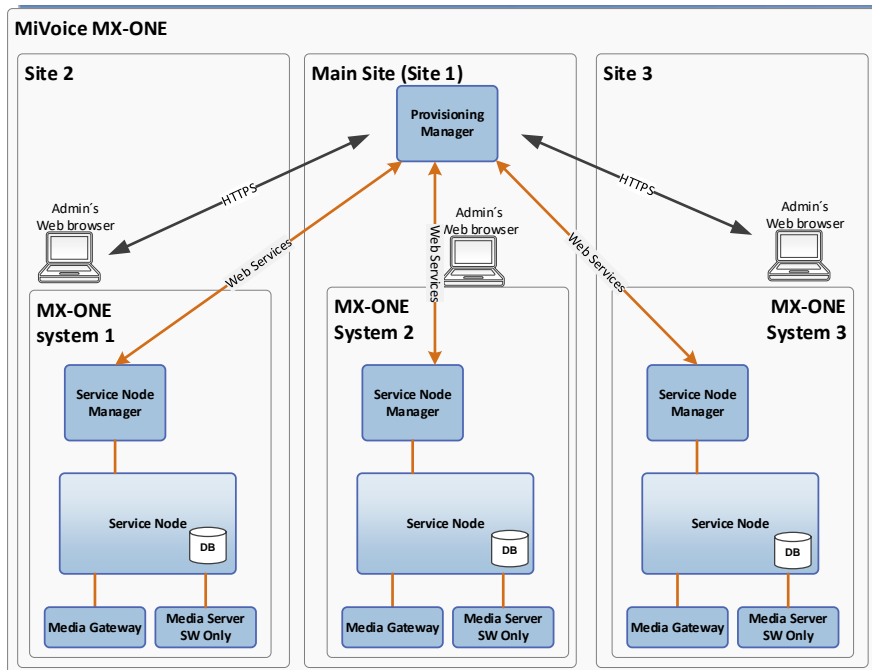


Figure 10 - Provisioning Manager



Note! The PM and SNM heap size need to be updated manually for large systems.

The MX-ONE Capacity and the Virtualization documents have the requirements for Provisioning Manager when running in the same server as Service Node 1 as well as when Provisioning Manager is running in a stand-alone server.

Media Server Requirements

Media Server can be deployed as co-located with a Service Node or as stand-alone.

Media Server requires more resources when SRTP is used to encrypt media streams and complex codecs like G.729 and G.722.

In some traffic cases transcoding is required and for those cases more RTP gateway resources are needed. Therefore, more CPU resources are necessary to handle those gateway calls in the MediaServer.

The MX-ONE Capacity and the Virtualization documents have the requirements for Media Server when running in the same server as Service Node as well as when Media Server Manager is running in a stand-alone server.

Media Gateways Requirements

MX-ONE 7.0 and later releases only supports Media Gateway Unit (MGU) as gateway for legacy devices. Encryption requires more DSP resources when a MX-ONE system has encryption enabled, the number of RTP resources is reduced.

Refer the *Media Gateway Unit (MGU)* documentation in MX-ONE CPI for more information.



Note! LSU-E is not supported.

Deployment Scenarios

Scenario 1 – 5,000 SIP Extensions

The figure below shows a typical MiVoice MX-ONE, containing Service Nodes with database co-located using 1+1 redundancy in the Service Node 1. The Service Node Manager is co-located with Service Node 1, Provisioning Manager and Media Server for SN1 runs on stand-alone servers. Service Node 2 has a MS co-located. The Service Nodes contain Media Gateway Unit (MGU) which allows ISDN connectivity to PSTN and legacy devices support.

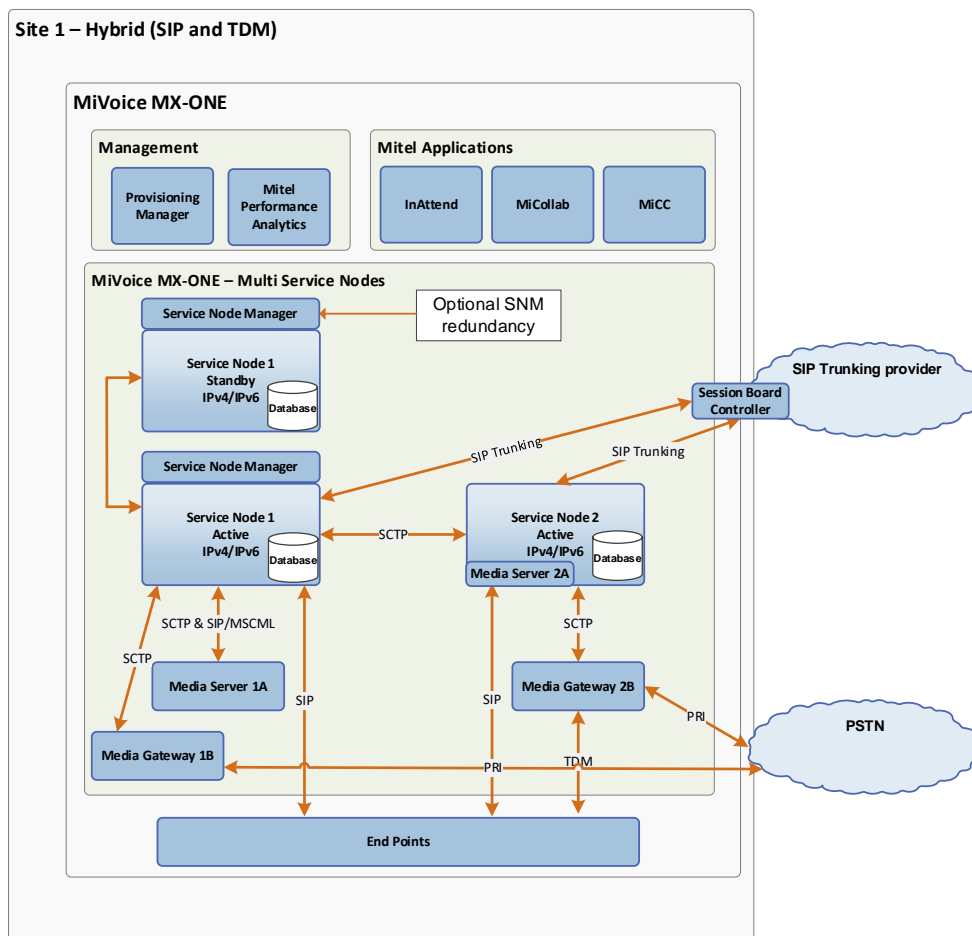


Figure 11 - MX-ONE deployed in one data center

In this scenario, there is only one Data Center and the database nodes are deployed in the Service Node 1 active, in the Service Node 1 standby and Service Node 2. The database nodes need to be installed in both active and standby, because in case Service Node 1 is not available, the Service Node 1 standby can take over the full functionality.

For example, if in this scenario the MX-ONE system has 5,000 users distributed in the 2 Service Nodes, the database requirement for 5,000 users is 10 GB of RAM memory and 4 CPUs/vCPUs which in this case is equal to the requirements for 2.5 K SIP users per Service Node.

From MX-ONE 7.0 release onwards Service Node Manager redundancy is introduced, which might be used in this case to avoid that the Management System (Provisioning Manager and Service Node Manager) do not work when the system is running in the Service Node standby.

Customers which require moves, adds and changes when a MX-ONE Service Node is down must use server redundancy, MX-ONE 7.0 and later releases still requires that all MX-ONE Service Nodes are up and running to be able to make changes in the system, because of the reload data used in some functionalities. So, in the scenario described in figure 6, the Service Node 2 must also support 1 + 1 server redundancy.

Scenario 2 – 15,000 SIP Extensions

The figure below shows a typical MiVoice MX-ONE with four sites, three data centers, containing Service Nodes, stand-alone database in the two main sites, the third site has a Service Node with co-located database node and the fourth no Service Nodes, only SIP phones.

Provisioning Manager runs on stand-alone servers.

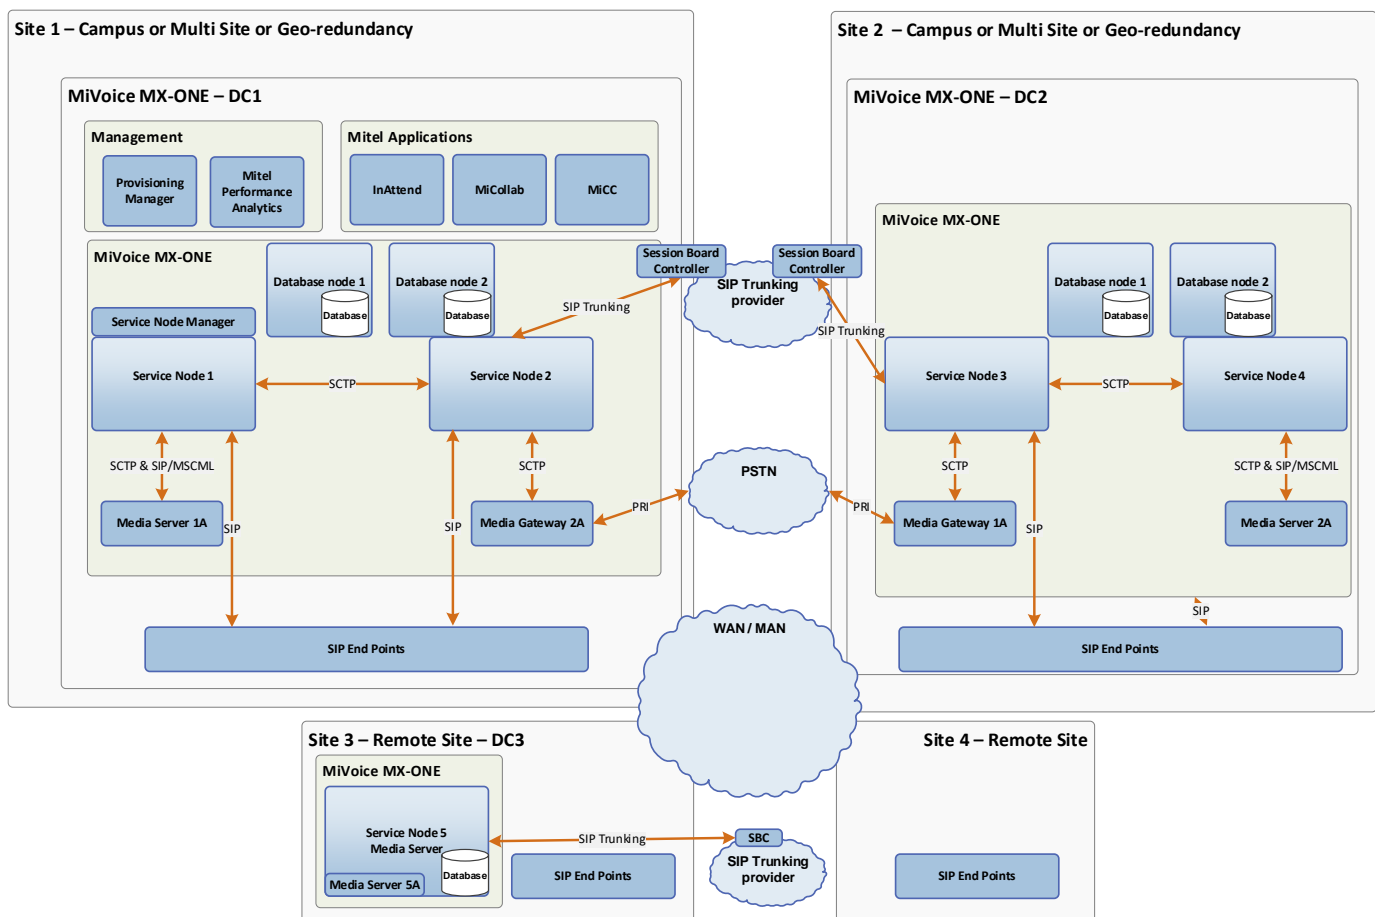


Figure 12 - MX-ONE deployed in three data centers

For example, if in this scenario the MX-ONE system has 15,000 users distributed in the 5 Service Nodes, SN1, SN2, SN3 and SN4 has 3,500 SIP users and SN5 has 1,000 SIP users. The database requirement for 15,000 users is 16 GB of RAM memory and 6 CPUs/vCPUs, so in this case all database nodes need to have 16 GB of RAM and 6 CPUs/vCPUs including the Service Node 5 which has a co-located system database node.

Scenario 3 – 81,000 SIP Extensions

The figure below shows a typical MiVoice MX-ONE with eight sites, eight data centers, containing Service Nodes, stand-alone database in the two main sites and Service Node with co-located database node in the remote sites. Provisioning Manager runs on stand-alone servers.

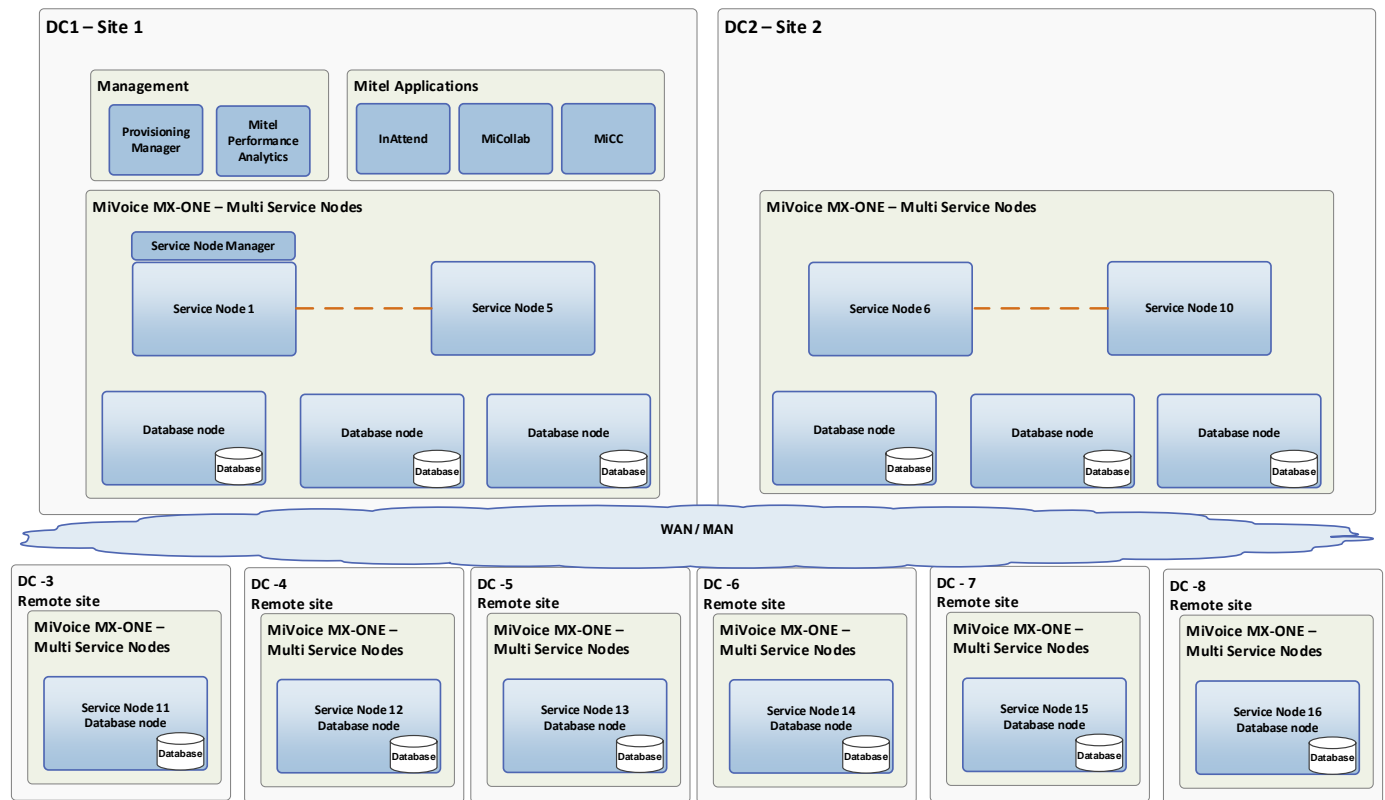


Figure 13 - MX-ONE deployed in eight data centers

For example, if in this scenario the MX-ONE system has 81,000 users distributed in the 16 Service Nodes, SN1 up to SN10 has 7,500 SIP users and SN11 to SN16 has 1,000 SIP users. The database requirement for 81,000 users is 32 GB of RAM memory and 8 CPUs/vCPUs (100,000 in the capacity table), so in this case all database nodes need to have 32 GB of RAM and 8 CPUs/vCPUs including the Service Nodes 11 to 16 which has a co-located system database node.

Redundant MX-ONE Systems

MX-ONE supports different types of server redundancy.

Server redundancy is achieved by adding one or more standby servers to the network with the ability to take over any failing Service Node in the cluster.

The following server redundancy are supported by MX-ONE.

- N+1
- 1+1
- 1+1 with Preload standby

NOTE: Customers which require moves, adds and changes when a MX-ONE Service Node is down must use server redundancy, MX-ONE requires that all MX-ONE Service Nodes are up and running to be able to make changes in the system, because of the reload data used in some functionalities, the Service Node can be a redundant server (1+1 or N+1).

Cluster

A cluster is composed by grouping together a number of regular servers and a standby server.

A cluster can have up to ten servers and only one standby server. A cluster with 1 active and 1 passive is named 1+1. A cluster with more than 2 actives servers and one standby is named N+1.

In a multi-server system, it is possible to have as many clusters as there are servers in the system (with a maximum of one standby server per regular server). A server and standby server can belong to only one cluster.

Server Redundancy

Using server redundancy, a standby server can take over the tasks of a regular server suffering from, for example, hardware failure. This way, a faulty server can be replaced with a minimum of disturbance. When using server redundancy, regular servers and an additional, standby server are grouped as a cluster. The standby server is prepared with data from the regular servers in the cluster and ready to start an instance of any of these servers in case of a server fault.

Each server in a cluster supervises the state of the other servers. In case of a server failure, the software and configuration running on the faulty server will be activated and started on the standby server. The standby server will also manage the media gateway or gateways of the faulty server. In a N+1 Server Redundancy scenario, if in the very unlikely event that there are more than one faulty servers at the same time in a cluster, the standby server will only replace one of the faulty servers. Other faulty servers will not operate until the fault is fixed and they are restarted.

When a regular server recovers from a failure, an automatic or a manual fallback will take place. During fallback, the Service Node software is stopped on the standby server and is then reloaded and restarted on the regular server.

At failover, the standby server will take over the identity of the failing server and the control of the media gateways in the failing server. For ongoing TDM traffic related to the media gateways controlled by that server, all connections will be lost when the server fails. All new traffic will be redirected to the standby server. In a distributed system connected over limited bandwidth (WAN), each remote domain must have its own standby server.



Note! All servers in one standby cluster must be on the same subnet.

Server Redundancy N + 1

In the server redundancy N+1, a cluster of up to 10 Service Nodes can have 1 standby server.

At failover, the standby server will take over the identity of the failing server and the control of the media gateways/servers in the failing server.

Server Redundancy 1 + 1

In the server redundancy 1+1, a cluster of 1 Service Node is created with 1 standby server.

At failover, the standby server will take over the identity of the failing server and the control of the media gateways/servers in the failing server.

Server Redundancy 1 + 1 with preloaded standby

In the server redundancy 1+1 with preloaded standby, a cluster of 1 Service Node is created with 1 standby server. In this case both Service Nodes are running, but the traffic is redirected to the active one, the traffic to the preloaded standby is blocked using IP tables.

At failover, the standby server will take over the identity of the failing server and the control of the media gateways/servers in the failing server. The failover time is faster in the 1+1 preloaded standby than the 1+1 server redundancy.

Standby Server and Database Nodes

In MX-ONE a server can have more than one role, for example a stand-alone database node can also be a standby server.

It is highly recommended to install the database nodes before promoting a server as standby. For example, install a new server, it will be a free server, after adding a stand-alone database node in this server, when the installation finishes go and add the standby role. As a last step create the cluster with the desired type of server redundancy.

Service Node Manager Server Redundancy

In MX-ONE Service Node Manager can be setup redundant. This option is important when customer requires Moves, Adds and Changes in all circumstances.

Refer to *Service Node Manager* documentation to read more about it.

Media Server Redundancy and Load Balancing

The Media Server component is enhanced in MX-ONE 7.0 and later releases and now it supports:

Server Redundancy

The Service Node using Media Server now supports redundancy, i.e. a stand-by Service Node can have one or several Media Servers, which are also on stand-by, and can start up and take over in case of a change-over to the backup server.

Load Balancing

The Service Node and Media Server now support load balancing, i.e. that media resources can be seized in alternative Media Servers (or in some cases also MGU), in case all involved end-points are IP-based, and thus not

“locked” to a certain media gateway (as for TDM end points). The load sharing is both between Service Nodes and within one Service Node.

Home Location Register (HLR) Redundancy

The Home Location Register (HLR) Redundancy is designed for customers that are deploying IP/SIP phones and have more than one server with load balancing enabled to share the load between servers. The user data is synchronized between the servers. Should one of the servers fail, the IP/SIP phones will re-register to one of the secondary servers (using a guest HLR) automatically and recover most of their features. When the primary server comes back, they will re-register to their primary HLR server.

HLR redundancy feature is intended for customers that do NOT use server redundancy (N+1 or 1+1), but still want a redundancy function for the SIP/IP phones (possibility to log in in another server).

Each server in an HLR redundancy setup can handle up to 15,000 HLR (registered phones) or guest HLR (“visitors”) IP/SIP end-points.

Performance MX-ONE Systems

Service Node

MX-ONE has been verified to handle the traffic described in the capacity and virtualization documents per server physical server or Virtual Machine.

The traffic values described in the documentation 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 = $1/\text{erlang} * \text{average call time} * \text{traffic (speed in calls per second)}$.

Number of users = $1/0.2 * 100.1$

Number of users = 500.

The performance tests were executed using Visitor Desktop (VDP), TLS 1.2 and SRTP enabled in the SIP phones.

Media Server

For customer running high traffic and contact center systems, standalone Media Servers are highly recommended.

Virtualization

MX-ONE can be virtualized with the following technologies:

- VMware ESXi 6.5 and later is the recommended, however VMware ESXi 5.5 and above are supported.
- KVM, which is delivered with SLES 12.
- Hyper-V based on Windows 2016 server.

NOTE: The virtualization environment shall be properly dimensioned to avoid resources starvation for the MX-ONE components. So, it is recommended to have the Memory reservation and CPU reservation equal to the Memory and CPU required values to avoid resources starvation.

VMware

MX-ONE 7.1 and later is tested with VMware 6.5 and later in Mitel's labs in the following setups:

- High Availability
- Fault Tolerance

For more information please read the Virtualization document.

For information on how to distribute the MiVoice MX-ONE VMs, please verify the appendix B in this document.

KVM

MX-ONE has been tested using KVM when MX-ONE is running in the EX gateway.

Hyper-V

MX-ONE 7.1 and later is tested with Hyper-V running on Windows 2016 server. For more information please read the Virtualization document.

Azure

MX-ONE 7.3 and later can be deployed in the Microsoft Azure public cloud, for more information refer to the *MiVoice MX-ONE Virtualization and Public Cloud – Description document (public cloud section)*.

High Availability Networks

Mitel MiVoice MX-ONE that is a communication system offering real time voice capability requires a high availability network. Mitel recommends that partners and customer follow the best practice in the market regarding high availability network design. A well-designed network reduces considerably the risk of failure and it increases the uptime of whole solution.

A robust network can be built using mechanisms/techniques that are very well known in the market. For instance, some mechanisms/techniques that need to be taken in consideration are: to build a layered Local Area Network (access, distribution and core layers), reduce the broadcast domains, enable QoS, avoid single point of failures and use redundant paths. Furthermore, basic techniques like Spanning Tree Protocol, Rapid Spanning Tree Protocol, Virtual Router Redundancy Protocol, Link aggregation (bonding) are highly recommended when building a redundant network for MiVoice MX-ONE. Additionally, use redundant network devices and make sure that switches and routers have duplicated power supply.

Moreover, when Service Node, Media Gateways and IP telephones are deployed over WAN (Wide Area Network) be sure to make it as much available as possible. Always try to avoid common WAN single point of failures, use a second link, preferable via a second provider and enable end to end QoS.

Design a High Availability Network

Redundant Network can be designed in several different ways, it will depend of the customer requirements, and normally a good network design includes a combination of hardware and software techniques.

The best way to increase the availability in any network is to ensure that network elements are always fitted with redundant components to avoid single point of failure.

Some examples of components to consider are:

- Power supplies
- Links
- Switches
- Routers
- Physical location of equipment
- Disks
- Storages

Network single point of failure reduction can be achieved using duplication of some hardware and the introduction of some network protocols that can help the availability of the network.

Equipment's can also be installed in different physical locations. For example, switches with redundant power supplies or switches using 48 VDC, using the same source than MiVoice MX-ONE. It means that switch 1 that Service Node 1 is connected to is powered by power supply 1 and switch 2 that Service Node 2 is connected to is powered by power supply 2.

Techniques likes as Virtual Local Area Networks (VLAN), Virtual Redundant Route Protocol (VRRP) and Layer 3 Protocols can be used to create a redundant environment that can increase the network availability. When using VRRP different VRID shall be used to avoid that the same MAC address is used by two VLANs.

Avoid the common mistake of using different VLANs in the same physical switch for network redundancy. If the switch fails, both VLANs will fail.

Avoid overbuilt (excessive network redundancy), because it is more difficult to manage.

Avoid under built (for example, having redundancy links between sites, but the links are fully loaded), because when needed, there will not be sufficient bandwidth to handle the extra traffic that will come from the link that presents the failure.

The figure below shows a high-level network redundant example.

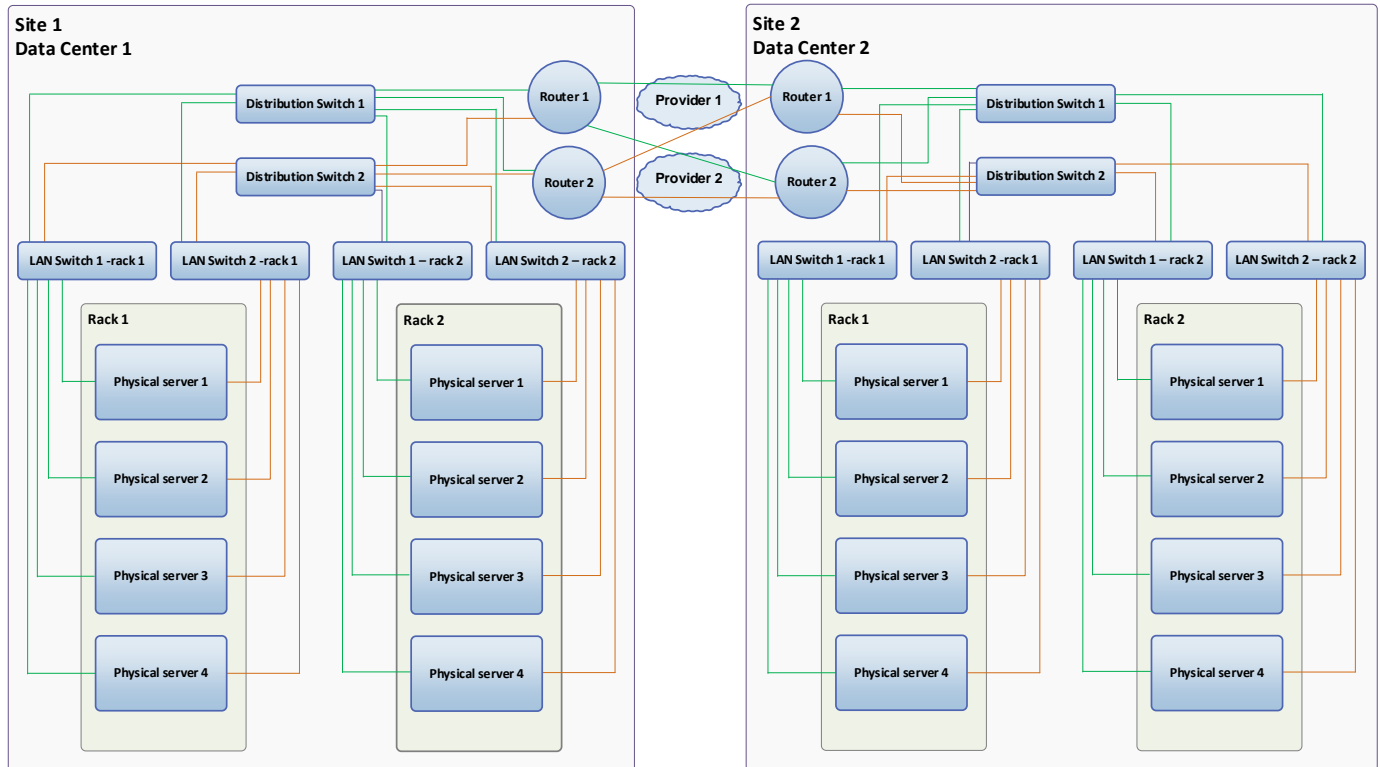


Figure 14 - Network diagram high level view

Security

The following MiVoice MX-ONE main components were updated regarding technology and security.

- Operating System – SLES 12 SP5.
The operation system updated the following main components used by MX-ONE.
 - OpenSSL
 - OpenSSH
 - Java
 - Postgres
 - Network
 - Security mechanisms
- Application Server – Wildfly 20

Security Mechanism in MX-ONE

By default, the audit log is enabled.

The communication with MX-ONE is done using SSH version 2.

Database server communication is encrypted using TLS 1.2.

Authorization codes can be hashed using SHA 256.

Authorization codes is expanded, and it supports a maximum of 64 characters.

Passwords typed in the SIP Phones are encrypted since the phone up to Service Node. This function requires SIP phones 68XX and 69XX running the latest 5.X firmware.

SRTP can be used to encrypt media, license is required. AES 256 is supported with the latest SIP 68XX and 69XX running the latest 5.X firmware.

TLS 1.2 can be used to encrypt signaling, license is required.

Provisioning Manager and Service Node Manager supports HTTPS with standard level of security or with high level of encryption (a Java license is required).

More information about MX-ONE security can be found in the security description.

Operating System Patches

Security flaws are discovered almost every day, in order to maintain the operating system used in MX-ONE up to date, Mitel deliveries once a month SLES patches update.

Application Security

MX-ONE software as well as the operating system are frequently verified against security flaws by application scanners. If security flaws are detected in Mitel's products, they will be fixed according to the Mitel product security policy.

The Mitel Product Security Policy can be found in the following link.

<https://www.mitel.com/support/security-advisories/mitel-product-security-policy>

Mitel Security Advisories can be found in the following link.

<https://www.mitel.com/support/security-advisories>

MX-ONE and MiCollab

MiCollab is the unified communication and collaboration used by MX-ONE.

MiCollab is connect in MX-ONE Service Node via SIP interfaces both SIP extensions and trunks depending of the service used and CSTA III (internal CSTA server in Service Node).

In the integration between MiVoice MX-ONE and MiCollab, Provisioning Manager is the only tool used to provision users in MiCollab.

It is important to note that one MiCollab system can be used by more than one MiVoice MX-ONE.

For example, if the customer has 5,000 SIP extensions distributed in 3 MiVoice MX-ONE connected by SIP trunks, 3,000 SIP extensions in the main site (site 1), 1,000 SIP extensions in the site 2 and 1,000 SIP extensions in the site 3, only one MiCollab is required.

For more information about MX-ONE and MiCollab, please read the MiVoice MX-ONE Solution Overview.

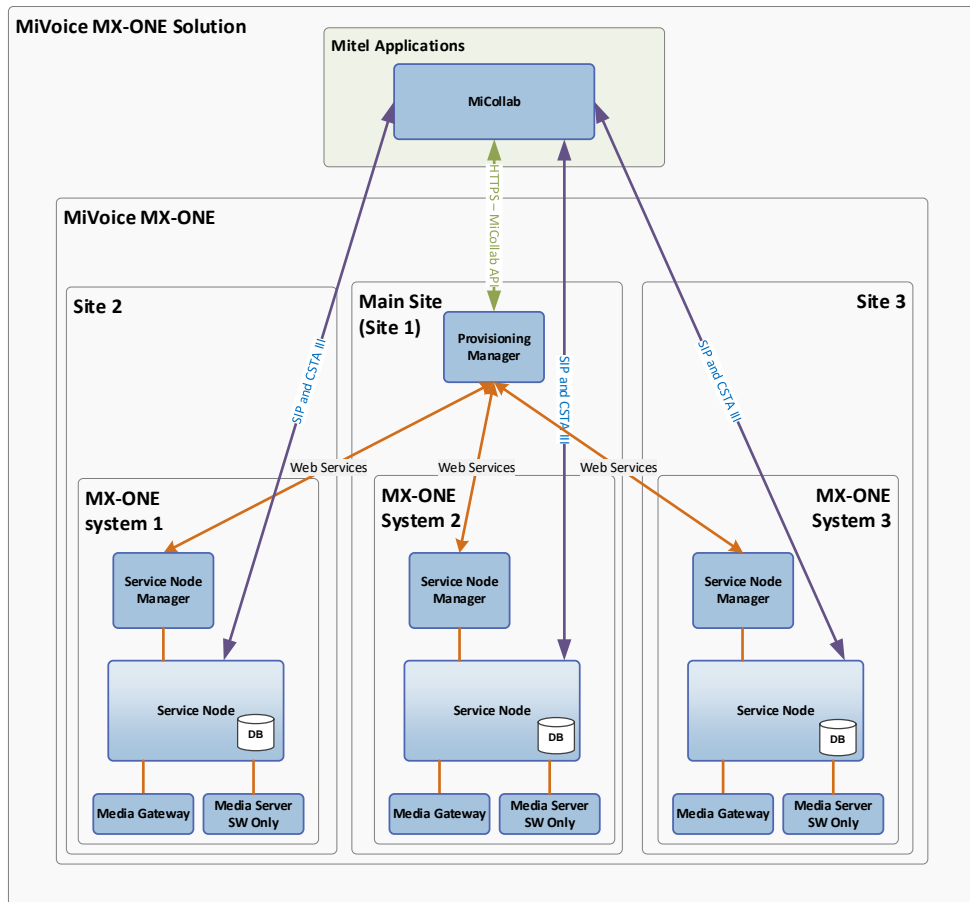


Figure 15 - MX-ONE and MiCollab

Appendix A

Database Nodes - MX-ONE 6.3 SPX Migration Path to MX-ONE 7.1 and Later

The database deployment in MX-ONE is different from the previous versions of MX-ONE. The aim with this guide is give a general overview on how to migrate a previous MX-ONE to MX-ONE 7.1 and later in the database context.

General Database Guidance

Stand-alone deployments are the primary choice for large (above 3,000 SIP users) and very large (above 10,000 SIP users) MiVoice MX-ONE systems.

The recommendation is to use three or four database nodes per site in large and very large systems.

In a system using co-located database, the memory and CPU requirements for the server shall cover the total number of users in the whole system. For example, if the MX-ONE system has 20,000 SIP extensions, the total memory and CPU per Service Node shall be the total memory and CPU required for the database nodes, 20,000 SIP extensions.

In virtualized environment the recommendation is a maximum of 2 database nodes per physical host servers in a typical entry Industrial Standard Server (e.g. Dell R4X0, ASU-II or equivalent), if the MX-ONE data is saved on the local disks. This recommendation does not apply when properly dimensioned SANs or High End Industrial Standard Servers are used in a Virtualization environment.

One Site Scenario, One Datacenter

One Service Node

When the MiVoice MX-ONE system is a single Service Node, there are two options:

- 1 - Keep the current database co-located in the Service Node as it was in MX-ONE 6.3 SPX.

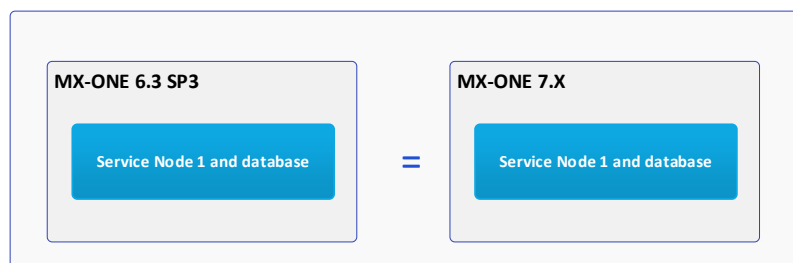


Figure 16 - Co-located scenario – One Service Node

- 2 - Add a database node as a stand-alone.

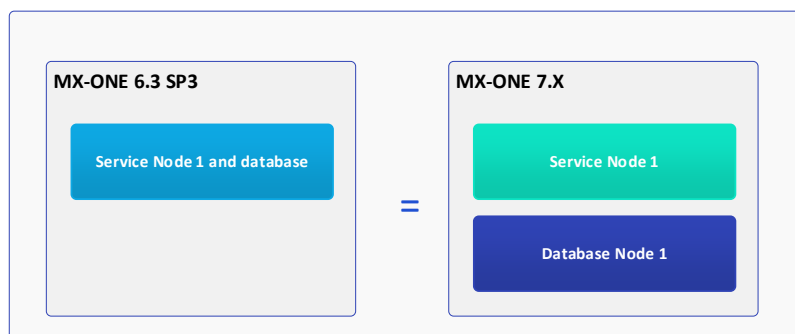


Figure 17- Stand-alone scenario – One Service Node

The second option shall be used if there are no available resources in the Service Node (Memory, CPU, Disk type, etc.) machine to accommodate the database node requirements.

Two Service Nodes

When the MiVoice MX-ONE is a two Service Nodes system, there are three options:

1 - Keep the current databases co-located in the Service Nodes as it was in MX-ONE 6.3 SPX.

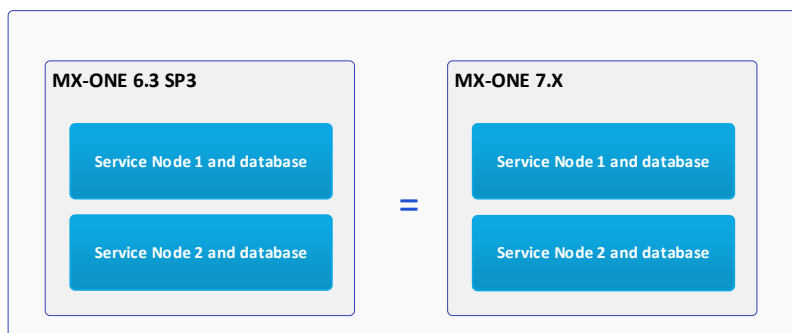


Figure 18 - Co-located scenario – Two Service Nodes

2 - Hybrid setup, keep one co-located database in one of the Service Nodes and add a second database as stand-alone.

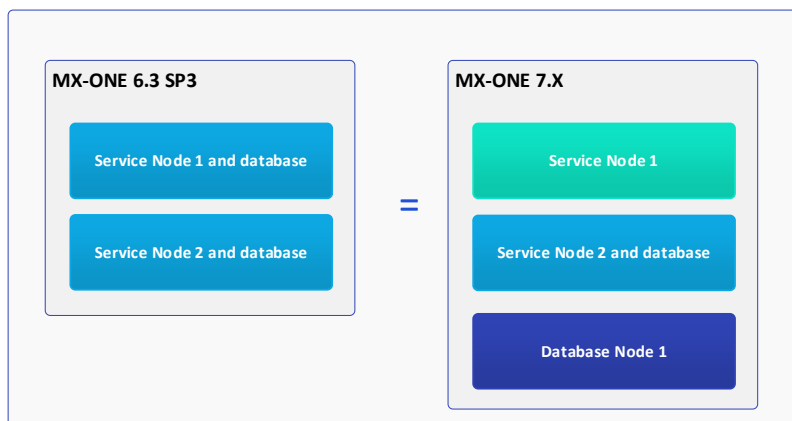


Figure 19 - Hybrid scenario - Two Service Nodes

3 - Add database nodes as stand-alone servers.

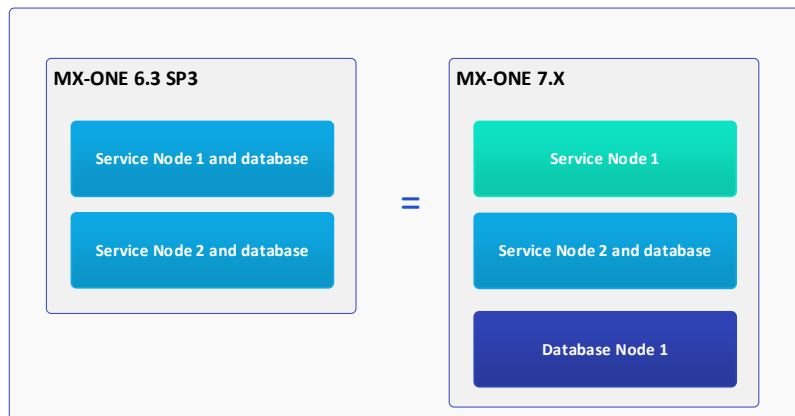


Figure 20 - Stand-alone scenario - Two Service Nodes

The second and the third options shall be used if there are no available resources in the Service Node (Memory, CPU, Disk type, etc.) machine to accommodate the database node requirements.

The second database is required for redundancy purpose.

Three Service Nodes

When the MiVoice MX-ONE is a three Service Nodes system, there are three options:

1 - Keep the current databases co-located in the Service Nodes as it was in MX-ONE 6.3 SPX.

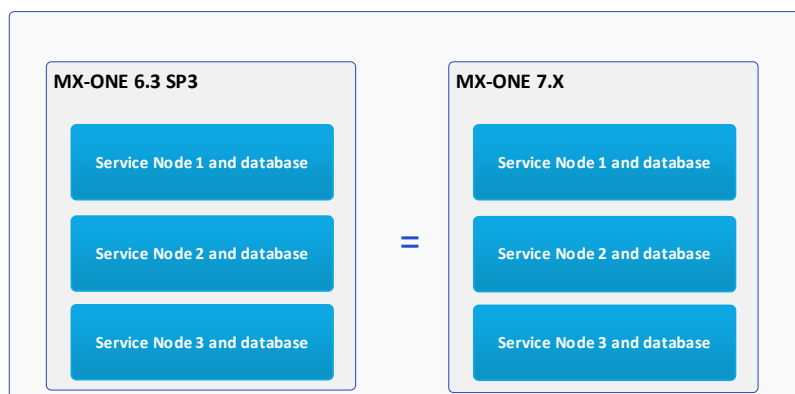


Figure 21 - Co-located scenario - Three Service Nodes

2 - Hybrid setup, keep one/two database/s co-located in one or two Service Nodes and add a second/third database as stand-alone servers.

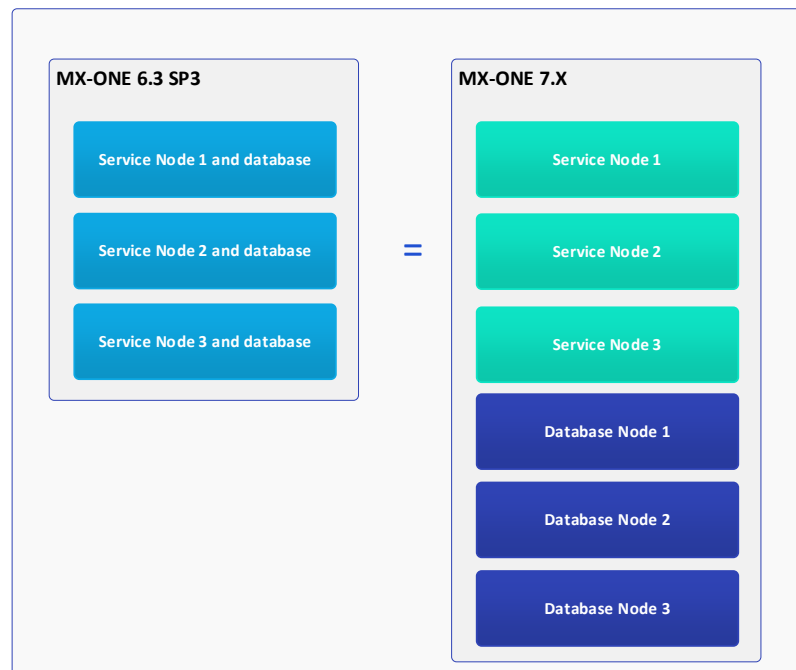


Figure 22 - Hybrid scenario - Three Service Nodes

3 - Add database nodes as stand-alone servers.

Figure 23 - Stand-alone scenario - Three Service Nodes

The second and the third options shall be used if there are no available resources in the Service Node (Memory, CPU, Disk type, etc.) machine to accommodate the database node requirements.



Four Service Nodes

When the MiVoice MX-ONE is four Service Nodes system, there are three options:

1 - Keep three databases co-located in the Service Nodes.

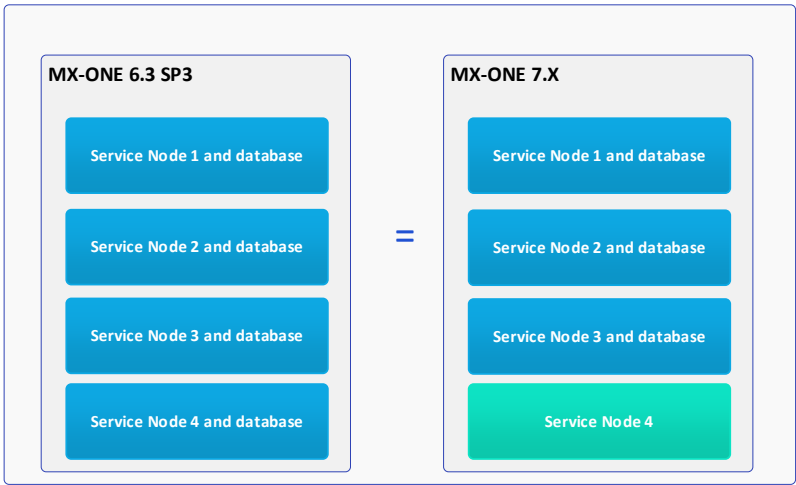


Figure 24 - Co-located scenario - Four Service Nodes

2 - Hybrid setup, keep one/two database/s co-located in one or two Service Nodes and add second/third databases as stand-alone servers.



Figure 25 - Hybrid scenario - Four Service Nodes

3 - Add databases node as stand-alone servers.



Figure 26 - Stand-alone scenario - Four Service Nodes

The second and the third options shall be used if there are no available resources in the Service Node (Memory, CPU, Disk type, etc.) machine to accommodate the database node requirements.

Ten Service Nodes

When the MiVoice MX-ONE is a ten Service Nodes system, there are three options:

1 - Keep at least three databases co-located in the Service Nodes.

Optionally, a fourth database can be added if the system is very large.

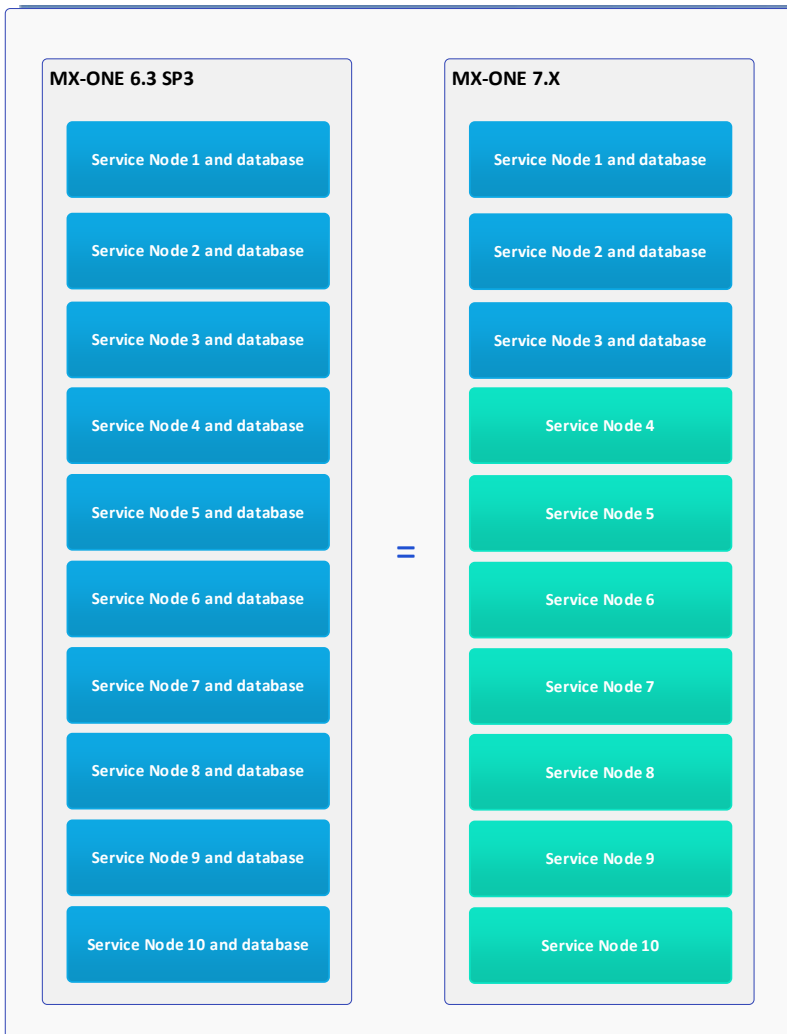


Figure 27 - Co-located scenario - Ten Service Nodes

2 - Hybrid setup, keep one/two database/s co-located in one or two Service Nodes and add second/third databases as stand-alone servers.

Optionally, a fourth database can be added if the system is very large.

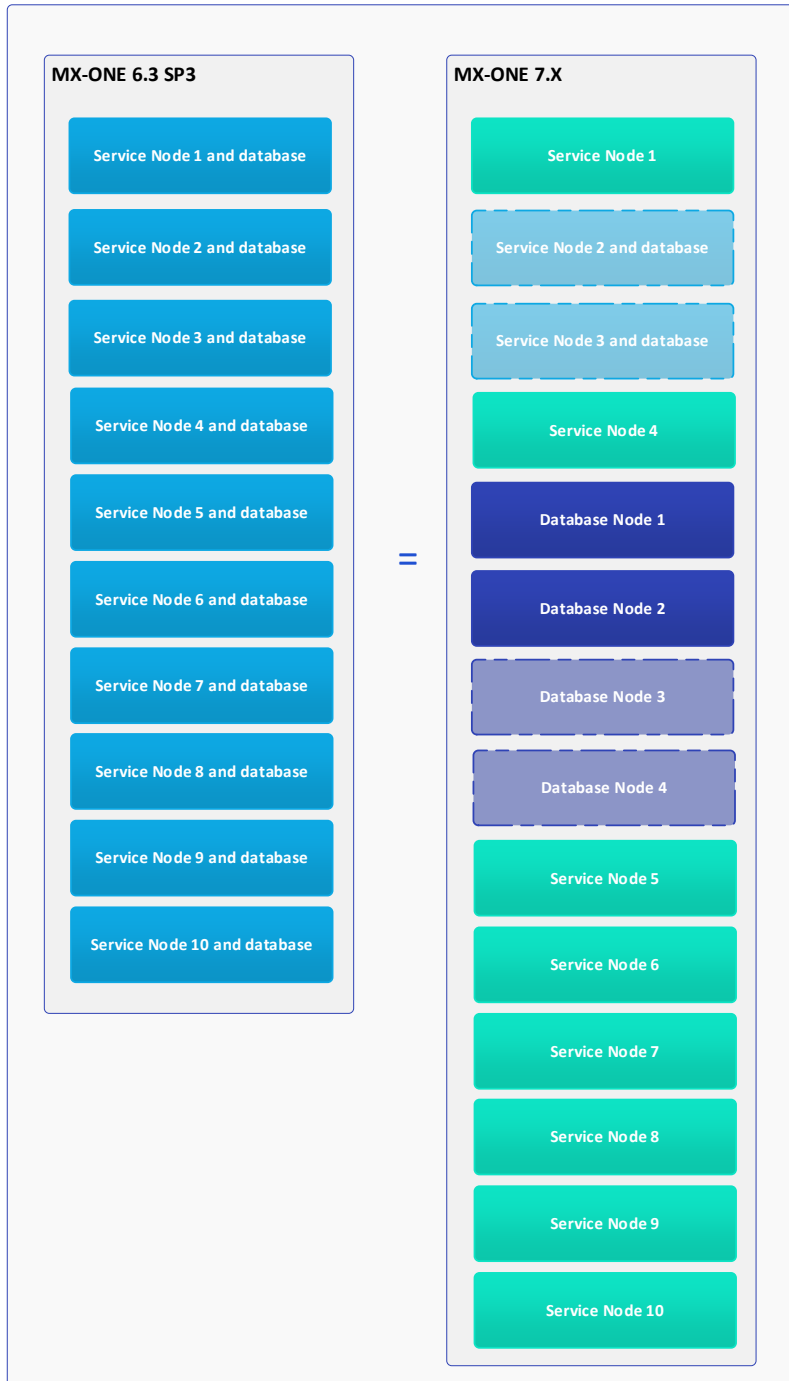


Figure 28 -Hybrid scenario - Ten Service Nodes

3 - Add database nodes as stand-alone servers.

Optionally, a fourth database can be added if the system is very large.

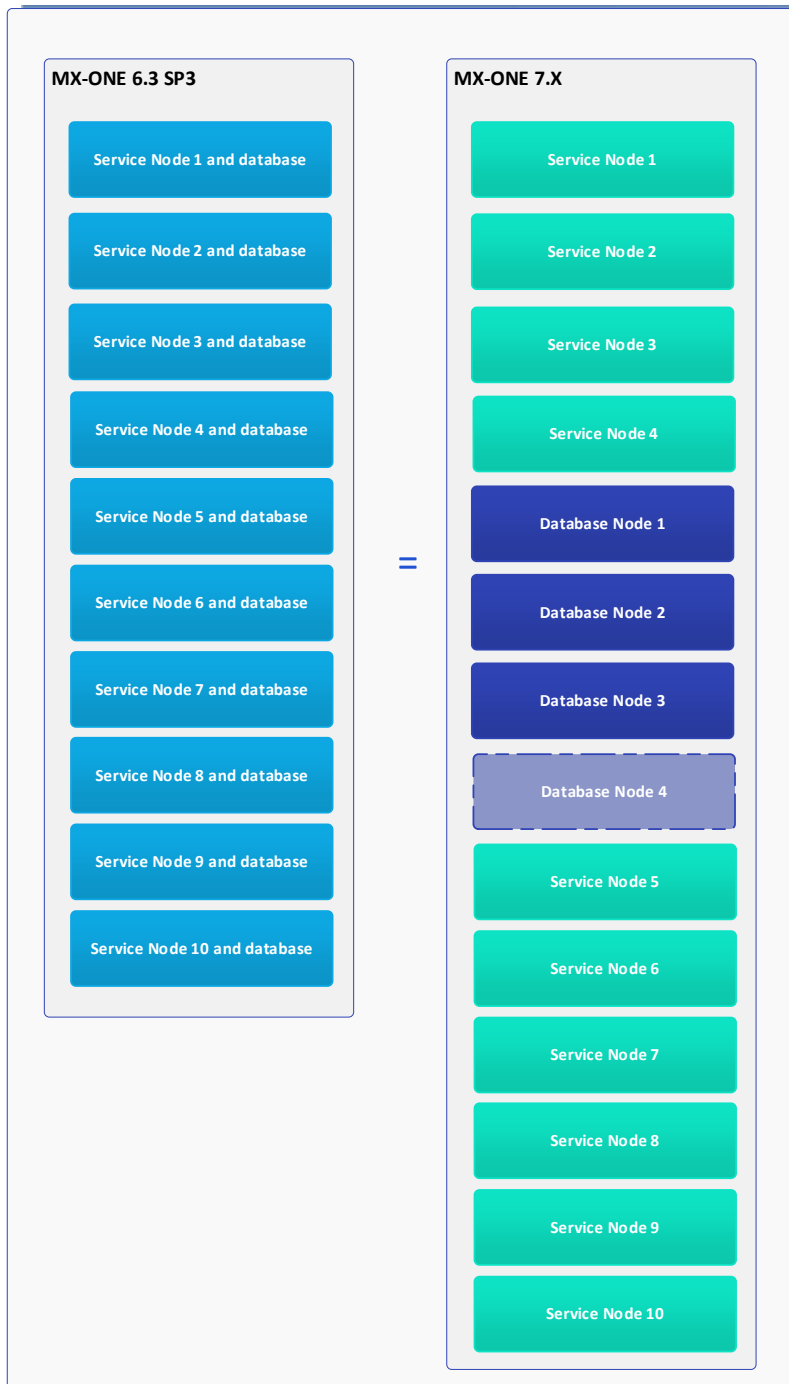


Figure 29 - Stand-alone scenario - Ten Service Nodes

The second and the third options shall be the preferred deployment for large (above 3,000 SIP users) and very large (above 10,000 SIP users) deployments.

Two Sites Scenario, Two Datacenters

Four Service Nodes

When the MiVoice MX-ONE is a four Service Nodes system distributed in two sites, there are three options:

1 - Keep the current databases co-located as it was in MX-ONE 6.3 SPX.

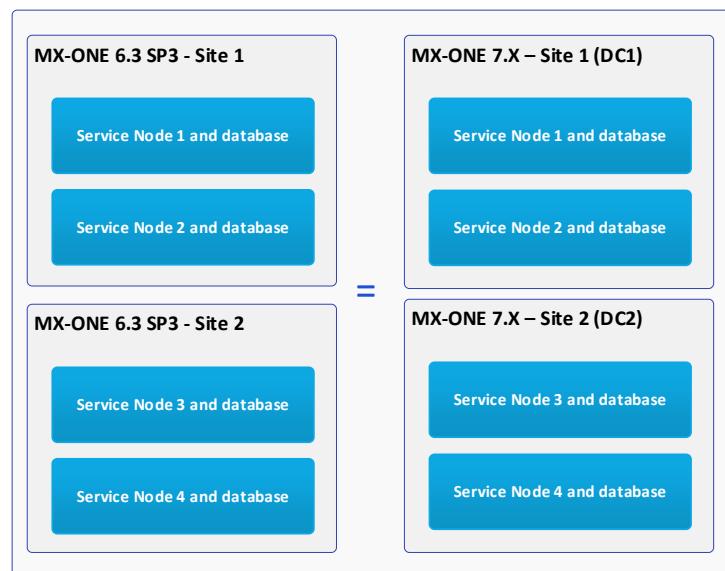


Figure 30 - Co-located scenario - Four Service Nodes distributed in two sites

2 - Hybrid setup, where one/two database/s are co-located in one or two Service Nodes and add second/third database nodes as stand-alone servers.

- Add databases nodes as stand-alone servers.



Figure 31 -Stand-alone scenario - Four Service Nodes distributed in two sites

The second and the third options shall be used if there are no available resources in the Service Node (Memory, CPU, Disk type, etc.) machine to accommodate the database node requirements.

Ten Service Nodes

When the MiVoice MX-ONE is a ten Service Nodes system distributed in two sites, there are three options:

- 1 - Keep at least two databases co-located in the Service Nodes of the site 1 and site 2.



Figure 32 - Co-located scenario - Ten Service Nodes distributed in two sites

2 - Hybrid setup, keep one/two database/s co-located in one or two Service Nodes and add second/third database nodes as stand-alone servers.

3 - Add database nodes as stand-alone servers.

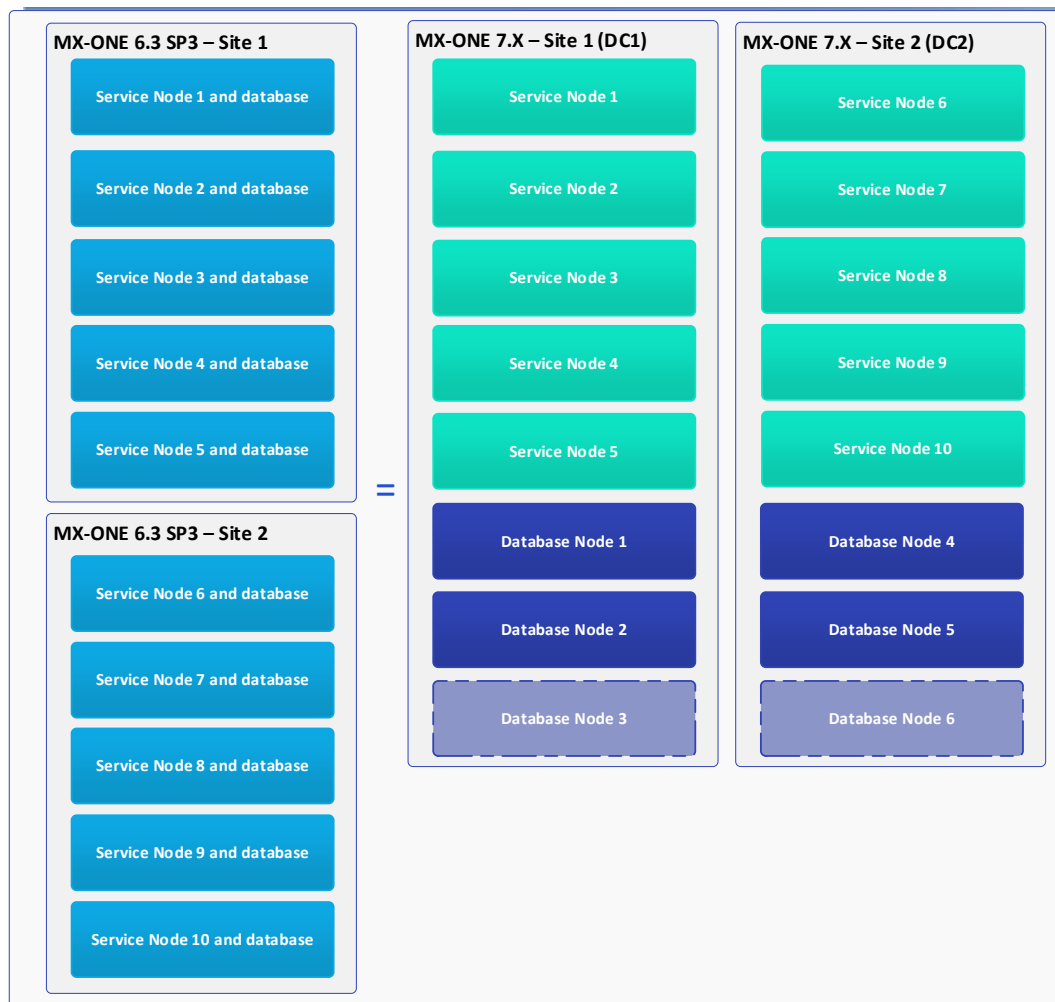


Figure 33 - Stand-alone scenario - Ten Service Nodes distributed in two sites

Six Sites Scenario, Six datacenters

Fourteen Service Nodes

When the MiVoice MX-ONE is a fourteen Service Nodes system distributed in six sites, there are three options:

1 - Keep at least two databases co-located in the Service Nodes of the site 1 and site 2 and keep the databases of the Service Nodes located in the remote sites.



Figure 34 - Co-located scenario - Fourteen Service Nodes distributed in six sites

- 2 - Hybrid setup, keep one/two database/s co-located in one or two Service Nodes of the site 1 and site 2 and add a second/third database as stand-alone and keep the databases of the Service Nodes located in the remote sites.
- 3 - Add databases node as stand-alone servers of the site 1 and site 2 and keep the databases of the Service Nodes located in the remote sites.

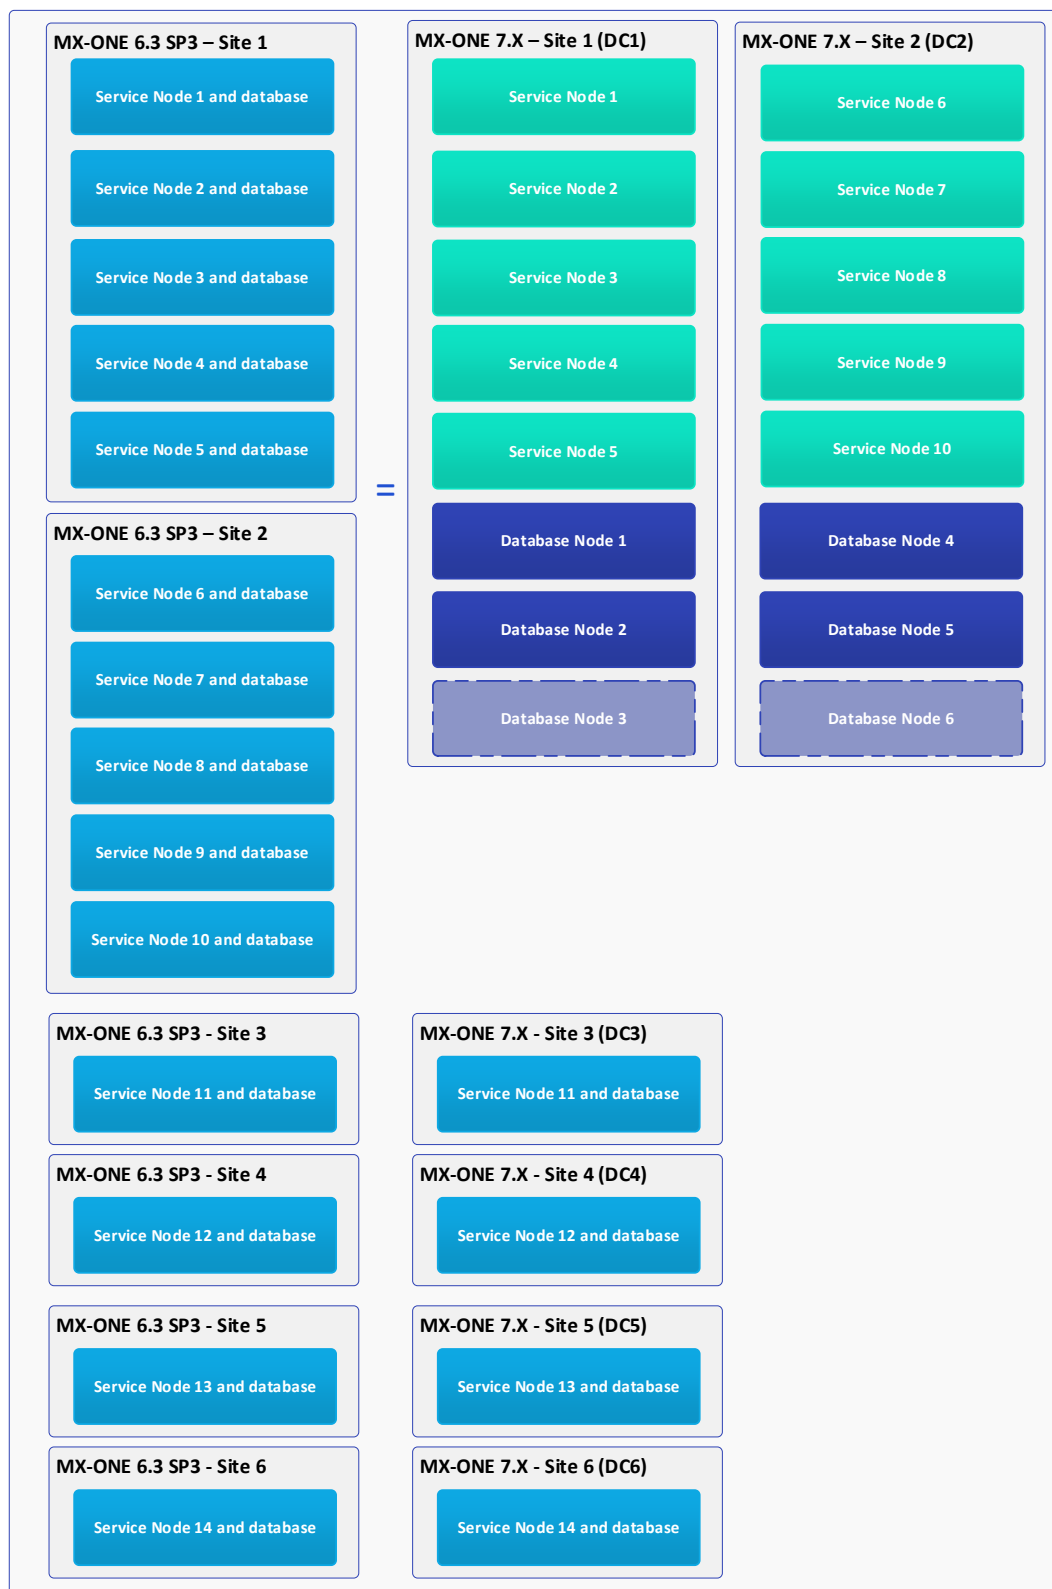


Figure 35 - Stand-alone scenario - Fourteen Service Nodes distributed in six sites

Appendix B

Service Nodes and Database Node Distribution in a Virtualized Environment

Database in Network Storage

The figure below shows an example of MX-ONE Virtual Machines (VM) distributed in four physical servers and two storages for redundancy purpose.

The advantage with this setup is the possibility to use the VMware High Availability, so the MX-ONE VMs can be moved for another server (if enough resources are available) in case of physical server failure.

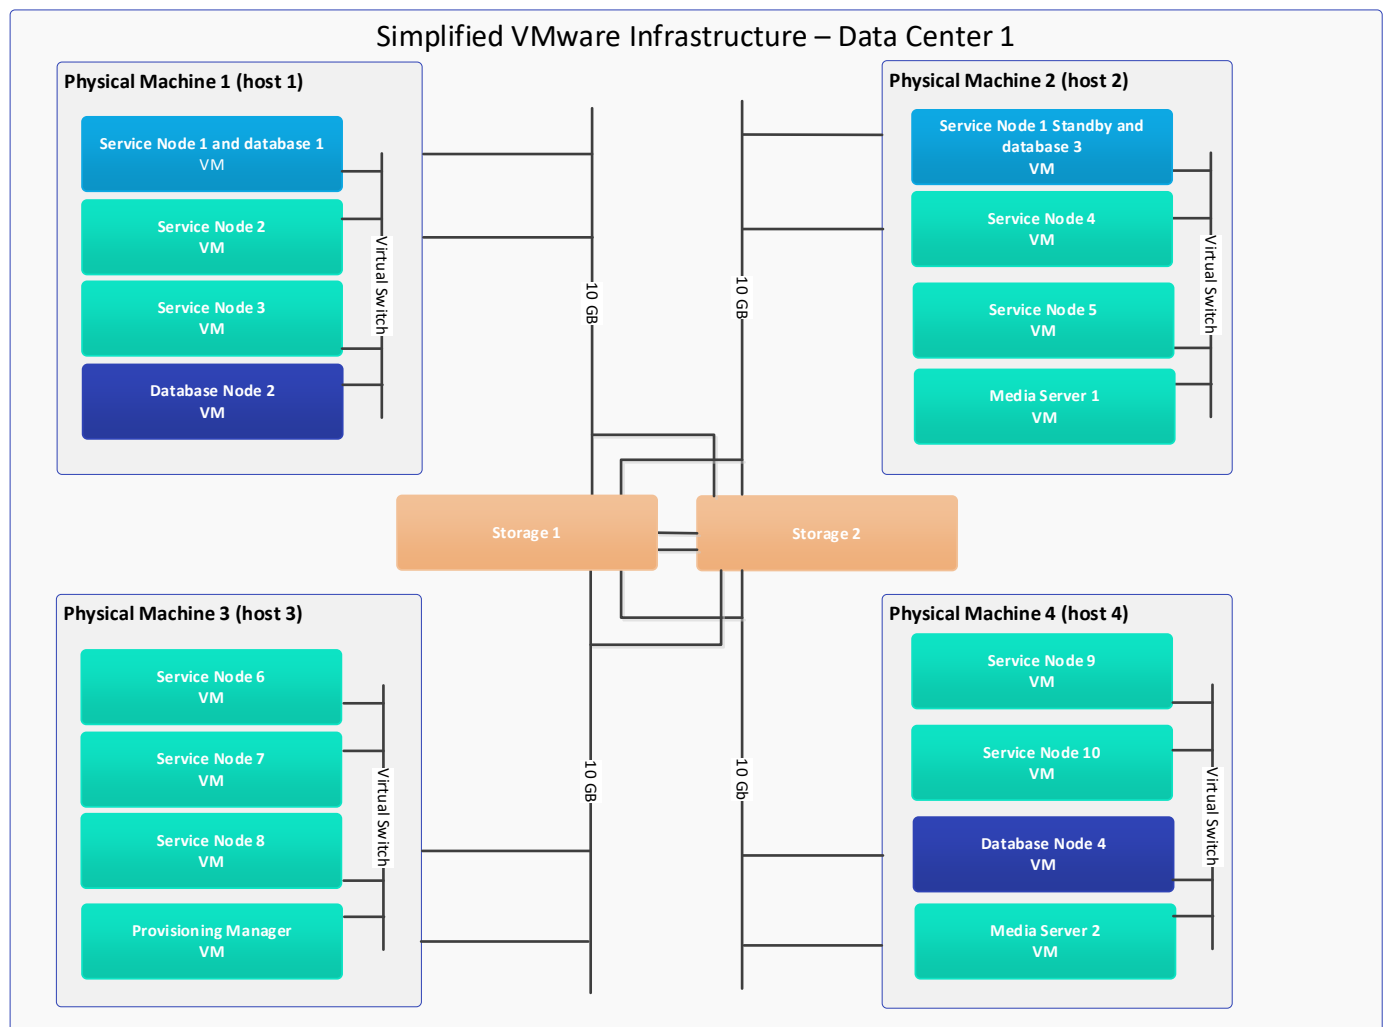


Figure 36 - MX-ONE VM distribution – storage

Acronyms, Abbreviations and Glossary

DB – Database
DNS – Domain Name System
HTTP – Hypertext Transfer Protocol
HTTPS – Hypertext Transfer Protocol Secure
IPv4 – Internet Protocol version 4
IPv6 – Internet Protocol version 6
KVM – Kernel-based Virtual Machine
LAN – Local Area Network
LIM – same as Service Node
MS – Media Server
MGW – Media Gateway
NTP – Network Time Protocol
PM – Provisioning Manager
QoS – Quality of Service
SCTP – Stream Control Transmission Protocol
SHA – Secure Hash Algorithms
SN – Service Node
SNM – Service Node Manager
SNTP – Simple Network Time Protocol
SRTP – Secure Real-time Transport Protocol
SSH – Secure Shell
RTP – Real-time Transport Protocol
TLS – Transport Layer Security
VLAN – Virtual Local Area Networks
VRRP – Virtual Redundant Route Protocol
WAN – Wide Area Network

Reference Documents

MiVoice MX-ONE Capacity

Description MiVoice MX-ONE

System Description MiVoice

MX-ONE Solution Overview

MiVoice MX-ONE

Virtualization Description

MiVoice MX-ONE Security

Description MiVoice MX-ONE

Security Guidelines MiVoice

MX-ONE System Planning

MiVoice MX -ONE Solution

IPv6 Support

Cassandra requirements:

<http://cassandra.apache.org/doc/latest/operating/hardware.html> SUSE SLES:

www.suse.com

Wildfly: www.wildfly.org

