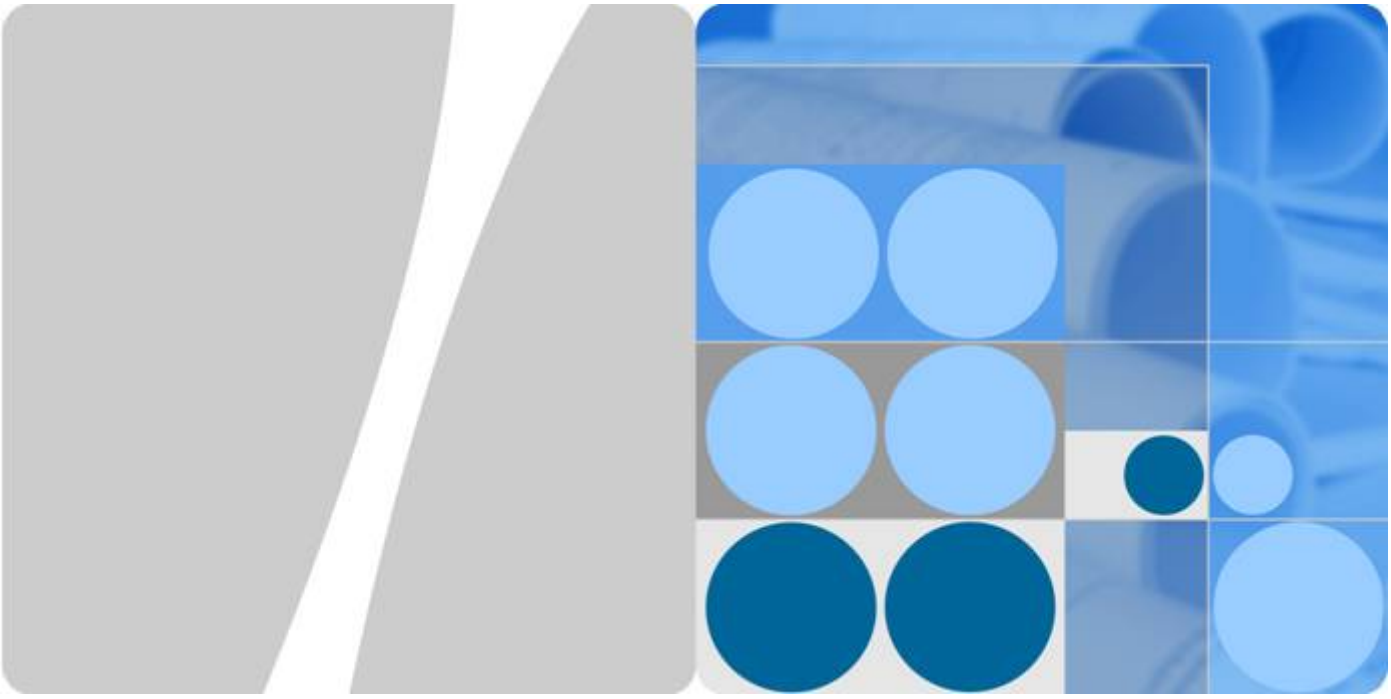

Huawei MZ910 NIC

V100R001

White Paper

Issue 09

Date 2016-11-21

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://e.huawei.com

About This Document

Purpose

This document describes the MZ910 in terms of its functions, appearance, features, applications, and technical specifications. You can obtain comprehensive information about the MZ910 by reading this document.

Intended Audience

This document is intended for:

- Huawei presales engineers
- Channel partner presales engineers
- Enterprise presales engineers

Symbol Conventions

The symbols that may be found in this document are defined as follows.

Symbol	Description
 DANGER	Alerts you to a high risk hazard that could, if not avoided, result in serious injury or death.
 WARNING	Alerts you to a medium or low risk hazard that could, if not avoided, result in moderate or minor injury.
 CAUTION	Alerts you to a potentially hazardous situation that could, if not avoided, result in equipment damage, data loss, performance deterioration, or unanticipated results.
 NOTE	Provides additional information to emphasize or supplement important points in the main text.

Change History

Issue 09 (2016-11-21)

This issue is the ninth official release.

Issue 08 (2016-08-17)

This issue is the eighth official release.

Issue 07 (2016-05-15)

This issue is the seventh official release.

Type	Change Description
Modify	The document content is optimized.

Issue 06 (2016-03-31)

This issue is the sixth official release.

Issue 05（2015-11-09）

This issue is the fifth official release.

Type	Change Description
Modify	Modified the note in **3.3 MZ910 Networking**.

Issue 04 (2015-05-30)

This issue is the fourth official release.

Issue 03 (2015-03-20)

This issue is the third official release.

Issue 02 (2014-11-10)

This issue is the second official release.

Issue 01 (2014-09-15)

This issue is the first official release.

Contents

Figures

Tables

1 Overview

About This Chapter

1.1 Functions

The MZ910 is a multi-purpose NIC. It is used for E9000 compute nodes, and provides two 10GE ports and two 8G Fibre Channel (FC) or 10G Fibre Channel over Ethernet (FCoE) ports for compute nodes to connect to switch modules in the chassis.

The MZ910 uses the Emulex XE201 chip. It provides two 10GE ports to support NIC applications, and provides two 8G FC or 10G FCoE ports (the port mode is configurable) to support FC and FCoE network applications. Software can automatically set the port mode (FC or FCoE)depending on the type of the connected switch modules.

The MZ910 provides independent physical ports to connect the service network and storage network respectively, which implements physical isolation between the two networks.

1.2 Appearance

The MZ910 can be installed in slot Mezzanine1 (Mezz1 for short) or Mezzanine2 (Mezz2 for short) on a half-width E9000 compute node or in slot Mezz1, Mezz2, Mezzanine3 (Mezz3 for short), or Mezzanine4 (Mezz4 for short) on a full-width E9000 compute node.

The MZ910 provides network ports for connecting to switch modules:

- When the MZ910 is installed in slot Mezz1 or Mezz3, its four ports connect to switch modules in slots 2X and 3X.
- When the MZ910 is installed in slot Mezz2 or Mezz4, its four ports connect to switch modules in slots 1E and 4E.

Figure 1-1 MZ910 appearance

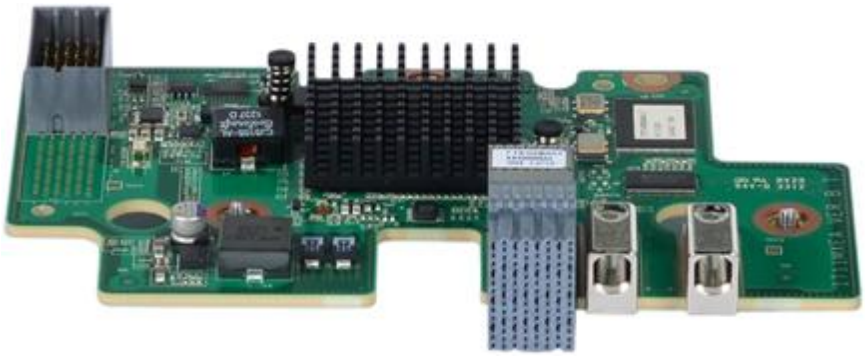

Figure 1-2 MZ910 installation positions on a half-width compute node

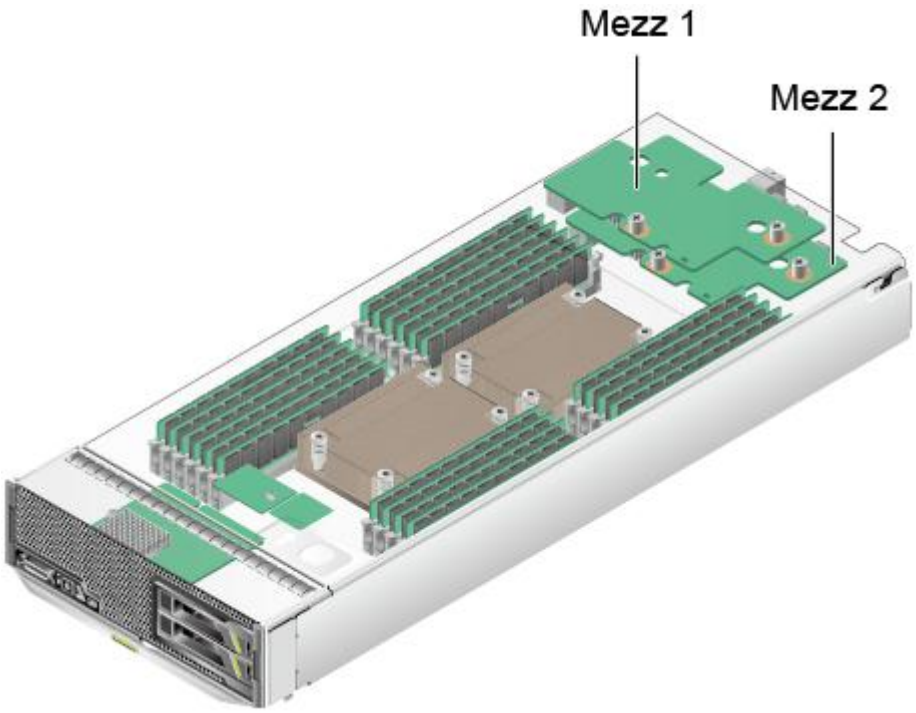

Figure 1-3 MZ910 installation positions on a full-width compute node

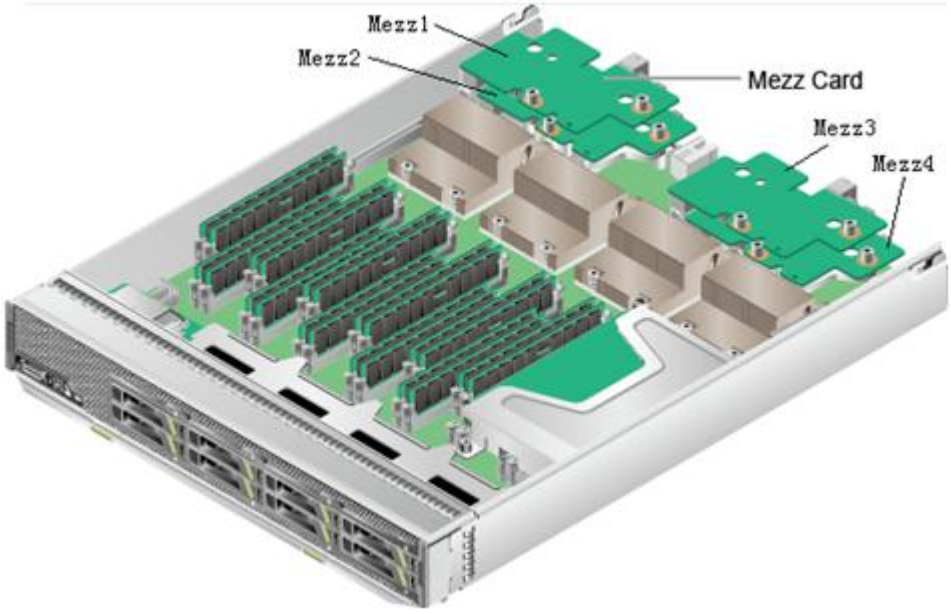

2 Features

About This Chapter

2.1 Feature List

The MZ910 supports the following features and performance specifications:

- Support FC and FCoE
- N_Port_ID Virtualization (NPIV), supporting a maximum of 255 virtual N_Port IDs per port
- Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), and Data Center Bridging Exchange (DCBX)
- Single Root I/O Virtualization (SR-IOV), supporting a maximum of 128 virtual functions (VFs)
- Virtual Machine Queue (VMQ) and NetQueue
- Virtual Ethernet Bridging (VEB)
- Ethernet PXE, FCoE storage area network (SAN) Boot, and iSCSI Boot (iBFT: iSCSI Boot Firmware Table)
- 802.1Q VLAN, supporting a maximum of 4094 VLANs
- Jumbo frames of 9 KB

2.2 Feature Description

PFC

The MZ910 supports the PFC feature. This feature supports flow control and back pressure on a per class-of-service (CoS) basis. The MZ910 supports a maximum of eight CoS queues. If both the MZ910 and its connected peer device support PFC, the link between them supports PFC. The MZ910 can identify PFC frames sent by the peer device and adjusts the transmitting bandwidth. If the inbound traffic exceeds the threshold, the MZ910 sends PFC frames to instruct the peer device to control the traffic.

The MZ910 can negotiate with the peer device about PFC parameters over DCBX. Typically, CoS3 is used for FCoE traffic, which is assigned by the FCoE Forwarder (FCF). When FCoE is enabled, only two CoS queues are available: one for FCoE traffic, and the other for native Ethernet traffic.

ETS

The MZ910 supports the ETS feature. This feature assigns port transmitting bandwidth based on priority groups (PGs). The MZ910 supports a maximum of eight CoS queues and two PGs. The minimum bandwidth and maximum bandwidth are assigned to each PG. The minimum bandwidth means the committed bandwidth, which is also known as the committed information rate (CIR). The maximum bandwidth means the maximum shared bandwidth, which is also known as the peak information rate (PIR). If other PGs are underloaded, the local PG can share (borrow) their bandwidths and its maximum bandwidth can reach 10 Gbit/s.

802.1Q VLAN

The MZ910 supports a maximum of 4094 VLANs. Each 10GE port supports a maximum of 4094 VLANs. The VLAN IDs are integers ranging from 1 to 4094.

The MZ910 does not tag or untag packets, but transparently transmits them. VLAN IDs are specified by the operating system (OS) on an E9000 compute node.

SR-IOV

The MZ910 supports the SR-IOV feature. The NIC supports a maximum of two physical functions (PFs, that is, 10GE physical ports) and 128 VFs. Each PF supports a maximum of 64 VFs.

The VFs derived from each PF can be assigned to virtual machines (VMs). In this way, the mapping between VFs and VMs is set up.

In VMware, each physical port supports a maximum of 16 VFs.

NOTE

The NIC with firmware RT9.6.7 or earlier supports a maximum of two PFs and 14 VFs. Each PF supports seven VFs.

NPIV

The MZ910 supports the NPIV feature when it operates in FCoE mode. With this feature, each port supports a maximum of 255 virtual N_Port_IDs and its own physical N_Port_ID.

Virtual N_Port_IDs can be assigned to VMs. In this way, the mapping between virtual N_Port_IDs and VMs is set up.

PXE, SAN Boot, and iBFT

The MZ910 supports PXE, SAN Boot, and iBFT.

- PXE is used for remote boot over the Ethernet or IP network. It enables users to connect to the remote PXE server for loading an OS.
- SAN Boot is used for remote boot over the FC or FCoE SAN. It enables users to connect to the remote FC or FCoE array for loading an OS.
- iBFT is used for remote boot over the Ethernet or IP network. Like iSCSI Boot, iBFT enables users to connect to the remote iSCSI array for loading an OS.

NOTE

The MZ910 does not support iSCSI Offload. It has to use upper-layer software (such as the OS, Hypervisor, and application) to implement iSCSI. However, the MZ910 implements iSCSI Boot by using the iBFT feature and connects to a remote iSCSI array for loading an OS.

2.3 Standards Compliance

Table 2-1 lists the standards and protocols that the MZ910 complies with.

Table 2-1 Standards compliance

Standard	Protocol
IEEE 802.3x	Flow Control and Back Pressure
IEEE 802.3z	1000BASE-X

Standard	Protocol
IEEE 802.3ap	10GBASE-KR
IEEE 802.1Qbb	Priority-based Flow Control (PFC)
IEEE 802.1Qaz	Enhanced Transmission Selection (ETS)
IEEE 802.1ab	Station and Media Access Control Connectivity Discovery (LLDP)
IEEE 802.3ad	Link Aggregation Control Protocol (LACP)
IEEE 802.1Qbg	Edge Virtual Bridging (EVB)
FC-LS	FC Link Service
FC-FS-2	FC Framing and Signaling
FC-GS-5	FC Generic Service
FCP-3	Fibre Channel Protocol for SCSI
FC-AL-2	Fibre Channel Arbitrated Loop
FC-BB-5	Fibre Channel-Backbone-5 (FCoE)
DCBX	Data Center Bridging Exchange

3 Applications

About This Chapter

3.1 Compatible Compute Nodes

The MZ910 can be installed in slot Mezz1 or Mezz2 on a half-width compute node or in slot Mezz1, Mezz2, Mezz3, or Mezz4 on a full-width compute node. **Table 3-1** lists the compute nodes that support the MZ910 and its installation positions on them.

Table 3-1 Compute nodes that support the MZ910

Compute Node	Number of Mezz Module Slots	MZ910 Installation Position
CH121	2	Mezz1 and Mezz2
CH121 V3	2	Mezz1 and Mezz2
CH220	1	Mezz1
CH220 V3	4	Mezz1, Mezz2, Mezz3, and Mezz4
CH221	1	Mezz1
CH222	2	Mezz1 and Mezz2
CH222 V3	2	Mezz1 and Mezz2
CH240	2	Mezz1 and Mezz2
CH242	4	Mezz1, Mezz2, Mezz3, and Mezz4
CH242 V3	4	Mezz1, Mezz2, Mezz3, and Mezz4

3.2 Connected I/O Modules

MZ910s can connect to I/O modules (switch modules or interface boards). **Figure 3-1** shows the connections between the MZ910s on a half-width compute node and the ports on I/O modules.

Figure 3-1 Connections between the MZ910s on a half-width compute node and the ports on I/O modules

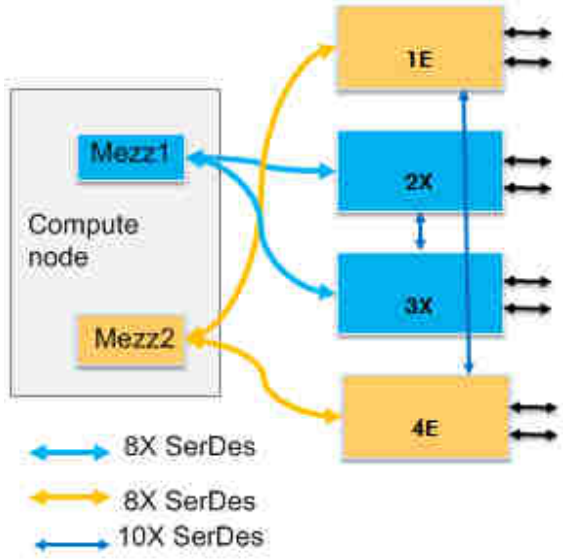

There are two or four groups of Serializer/Deserializer (SerDes, known as high-speed interconnect line) between each compute node and I/O module slots.

- Mezz1: 8X SerDes for connecting to I/O module slots 2X and 3X
- Mezz2: 8X SerDes for connecting to I/O module slots 1E and 4E
- Mezz3 (available only on a full-width compute node): 8X SerDes for connecting to I/O module slots 2X and 3X
- Mezz4 (available only on a full-width compute node): 8X SerDes for connecting to I/O module slots 1E and 4E

NOTE

The MZ910 provides four ports, and only 2X of each 8X SerDes is used.

Table 3-2 describes the I/O modules to which the MZ910 can connect.

Table 3-2 I/O modules to which the MZ910 can connect

I/O Module	I/O Module Slot	MZ910 (Mezz1)	MZ910 (Mezz2)	Typical Configuration	Remarks
CX911	2X/3X	√	X	Yes	-
	1E/4E	X	√	No	It is recommended that the CX911 not be installed in slot 1E or 4E.
CX912	2X/3X	√	X	Yes	-
	1E/4E	X	√	No	It is recommended that the CX912 not be installed in slot 1E or 4E.

I/O Module	I/O Module Slot	MZ910 (Mezz1)	MZ910 (Mezz2)	Typical Configuration	Remarks
CX913	2X/3X	√	X	Yes	The FC pass through module on the CX913 provides only eight 8G FC pass-through (directly attached) ports to connect to the MZ910s on eight half-width compute node slots in the chassis. Only one of every two neighboring half-width compute nodes on the left and right can be selected for the connection.
	1E/4E	X	√	No	The FC pass through module on the CX913 provides only eight 8G FC pass-through (directly attached) ports to connect to the MZ910s on eight half-width compute node slots in the chassis. Only one of every two neighboring half-width compute nodes on the left and right can be selected for the connection. It is recommended that the CX913 not be installed in slot 1E or 4E.
CX915	2X/3X	√	X	No	-
	1E/4E	X	√	No	-

3.3 MZ910 Networking

The MZ910 can connect to I/O modules (switch modules or interface boards) to provide Ethernet and FC/FCoE services.

The MZ910 can work with the CX911 switch module. If this happens, the MZ910 operates in NIC+FCoE or NIC+FC mode. The two 10GE ports on the MZ910 provide 20 Gbit/s bandwidth to support the Ethernet service. The MZ910 connects to the Ethernet switch

module on the CX911 and then to the Internet through 10GE ports on the CX911. The two 8G FC or 10G FCoE ports on the MZ910 provide 2 x 8 Gbit/s or 2 x 10 Gbit/s bandwidth to support the FC or FCoE storage service. The MZ910 connects to the FC switch module MX510 on the CX911 and then to the external FC SAN through 8G FC ports on the CX911. See **Figure 3-2**.

Figure 3-2 Connection between the MZ910 and the CX911

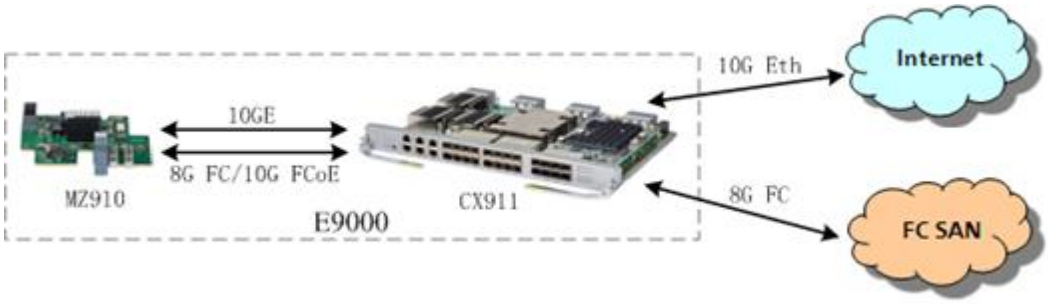

The MZ910 can work with the CX912 switch module. If this happens, the MZ910 operates in NIC+FC mode. The two 10GE ports on the MZ910 provide 20 Gbit/s bandwidth to support the Ethernet service. The MZ910 connects to the Ethernet switch module on the CX912 and then to the Internet through 10GE ports on the CX912. The two 8G FC ports on the MZ910 provide 2 x 8 Gbit/s bandwidth to support the FC storage service. The MZ910 connects to the FC switch module MX210 on the CX912 and then to the external FC SAN through 8G FC ports on the CX912. See **Figure 3-3**.

Figure 3-3 Connection between the MZ910 and the CX912

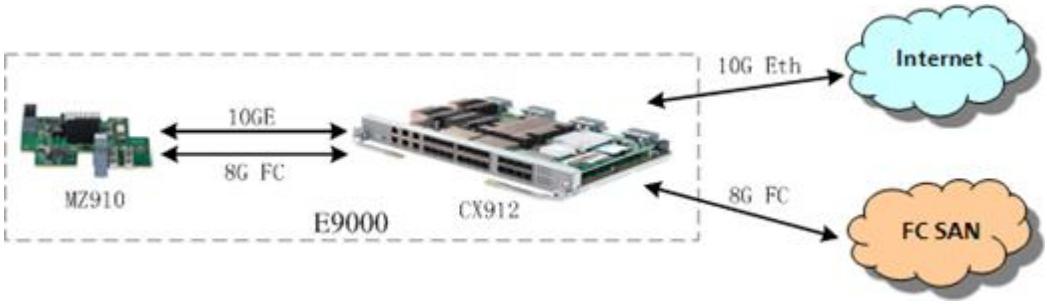

The MZ910 can work with the CX913 switch module. If this happens, the MZ910 operates in NIC+FC mode. The two 10GE ports on the MZ910 provide 20 Gbit/s bandwidth to support the Ethernet service. The MZ910 connects to the Ethernet switch module on the CX913 and then to the Internet through 10GE ports on the CX913. The two 8G FC ports on the MZ910 provide 2 x 8 Gbit/s bandwidth to support the FC storage service. The MZ910 connects to the FC pass through module on the CX913 and then to the external FC SAN through 8G FC ports on the CX913. See **Figure 3-4**.

Figure 3-4 Connection between the MZ910 and the CX913

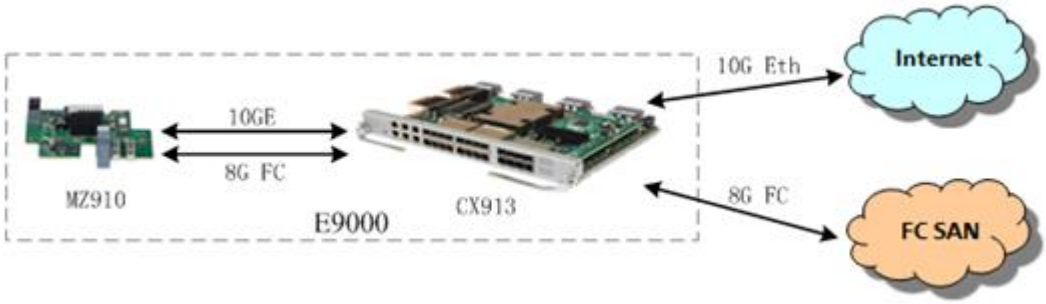

NOTE

The FC pass through module on the CX913 provides only eight 8G FC pass-through (directly attached) ports to connect to the MZ910s on eight half-width compute node slots in the chassis. Only one of every two neighboring half-width compute nodes on the left and right can be selected for the connection. There are eight pairs of 2-to-1 selection, that is, compute node slots 1 and 9, slots 2 and 10, and so on. For details, see the *CX913 Switch Module White Paper* and *MM910 Management Module Configuration Guide*.

The MZ910 can work with the CX915 switch module. If this happens, the MZ910 operates in NIC+FC mode. The two GE ports on the MZ910 provide 2 Gbit/s bandwidth (the rate of each 10GE port on the MZ910 decreases to GE) to support the Ethernet service. The MZ910 connects to the Ethernet switch module on the CX915 and then to the Internet through 10GE ports on the CX915. The two 8G FC or 10G FCoE ports on the MZ910 provide 2 x 8 Gbit/s or 2 x 10 Gbit/s bandwidth to support the FC or FCoE storage service. The MZ910 connects to the FC switch module MX510 on the CX915 and then to the external FC SAN through 8G FC ports on the CX915. See **Figure 3-5**.

Figure 3-5 Connection between the MZ910 and the CX915

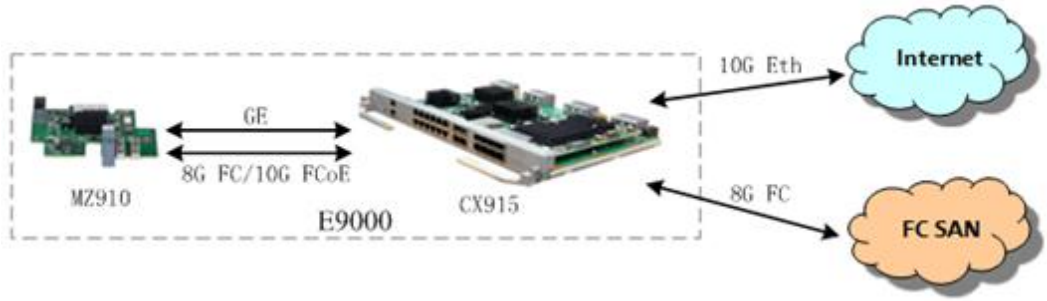

NOTE

When the MZ910 works with the CX911, CX912, or CX913, forcing the NIC ports on the MZ910 into GE mode has reliability risks and therefore is not supported. To force the NIC ports into GE mode, contact Huawei technical support.

3.4 Supported OSs

Table 3-3 lists the OSs supported by the MZ910.

Table 3-3 OSs supported by the MZ910

OS	Version	Remarks
Redhat	RHEL 6.5	-
	RHEL 6.6	-
	RHEL 6.7	-
	RHEL 7.0	-
	RHEL 7.1	-
	RHEL 7.2	-
Suse	SLES 11.3	-
	SLES 11.4	-
	SLES 12.0	-
	SLES 12.1	-
VMware	VMware ESXi 5.1.3	-
	VMware ESXi 5.5.2	-
	VMware ESXi 5.5.3	-
	VMware ESXi 6.0	-
	VMware ESXi 6.0.1	-
	VMware ESXi 6.0.2	-
Windows	Windows 2008 R2 SP1	-
	Windows 2012	-
	Windows 2012 R2	-

The preceding table is for reference only. Compatible OSs for the MZ910 vary with the compute node type. For details, see the ***Huawei Server Compatibility Checker***.

3.5 Connected FC Devices

Table 3-4 lists the FC arrays and FC switches to which the MZ910 can connect.

Table 3-4 FC devices to which the MZ910 can connect

Category	Vendor	Model	Remarks
FC array	EMC	Symmetrix DMX-4	-

Category	Vendor	Model	Remarks
		Symmetrix VMAX	-
		Symmetrix VMAXe	-
		VNX	-
	Huawei	HVS85T	-
		S2200T	-
		S2600T	-
		S5500T	-
		S3900	-
		S5600T	-
		S5800T	-
		S5900	-
		S6800T	-
		S6900	-
	Futijsu	DX80	-
	Synology	DS3611xs	-
		DS3612xs	-
		RS3411RPxs	-
		RS3411xs	-
		RS3412RPxs	-
		RS3412xs	-
FC switch	Brocade	Brocade 300	The MZ910 connects to Brocade 300 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		Brocade 5100	The MZ910 connects to Brocade 5100 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		Brocade 5300	The MZ910 connects to Brocade 5300 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		Brocade 8510	The MZ910 connects to Brocade 8510 in NPV mode through the CX911, CX912, or CX915 on the E9000.

Category	Vendor	Model	Remarks
	Huawei	CX911	The CX911 is a switch blade on the E9000, and its FC switch module is an ODM product from QLogic.
		CX912	The CX912 is a switch blade on the E9000, and its FC switch module is an ODM product from Brocade.
		CX915	The CX915 is a switch blade on the E9000, and its FC switch module is an ODM product from QLogic.
		SNS2124	OEM Brocade 300 The MZ910 connects to SNS2124 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		SNS2248	OEM Brocade 6510 The MZ910 connects to SNS2248 in NPV mode through the CX911, CX912, or CX915 on the E9000.
	Cisco	MDS 9148	The MZ910 connects to MDS 9148 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		MDS 9513	The MZ910 connects to MDS 9513 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		MDS 9505	The MZ910 connects to MDS 9505 in NPV mode through the CX911, CX912, or CX915 on the E9000.
		MDS 9222i	The MZ910 connects to MDS 9222i in NPV mode through the CX911, CX912, or CX915 on the E9000.

4 Technical Specifications

About This Chapter

4.1 Technical Specifications

Table 4-1 lists the technical specifications for the MZ910.

Table 4-1 Technical specifications

Item	Specifications
Dimensions (length x width)	148 mm x 85 mm (5.83 in. x 3.35 in.)
Power supply	12 V DC
Net weight	0.15 kg (0.33 lb)
Maximum power consumption	13 W
Temperature	Operating temperature: 5°C to 40°C (41°F to 104°F) (ASHRAE Class A3 compliant)
	Storage temperature: –40°C to +65°C (–40°F to +149°F)
Temperature change rate	15°C/h (27°F/h)
Humidity	Operating humidity: 5% RH to 85% RH (non-condensing)
	Storage humidity: 5% RH to 95% RH (non-condensing)
Altitude	● 40°C (104°F) at 1800 m (5905.44 ft) ● 30°C (86°F) at 3000 m (9842.40 ft) When the MZ910 is used at an altitude between 1800 m and 3000 m, the highest operating temperature decreases by 1°C (1.8°F) as the altitude increases by 120 m (393.70 ft).
PCIe port bandwidth	64 Gbit/s (PCIe 3.0 x8)
Port rate	10.3125 Gbit/s and 8.5 Gbit/s
Number of ports	4
Port type	Ethernet and FC/FCoE
Chip model/manufacturer	XE201/Emulex

A Acronyms and Abbreviations

C	
CoS	class of service
D	
DCB	Data Center Bridging
DCBX	Data Center Bridging Exchange
E	
ETS	Enhanced Transmission Selection
F	
FC	Fibre Channel
FCF	FCoE Forwarder
FCoE	Fibre Channel over Ethernet
I	
iSCSI	Internet Small Computer System Interface
iBFT	iSCSI Boot Firmware Table
L	
LACP	Link Aggregation Control Protocol
N	
NIC	network interface card
NPIV	N_Port_ID Virtualization

O	
OS	operating system
P	
PCIe	Peripheral Component Interconnect Express
PF	physical function
PFC	Priority-based Flow Control
PG	Priority Group
PXE	Preboot Execution Environment
S	
SAN	storage access network
SR-IOV	Single Root I/O Virtualization
V	
VEB	Virtual Ethernet Bridging
VF	virtual function
VLAN	virtual local access network
VM	virtual machine
VMQ	virtual machine queue