

**Huawei MZ912 NIC  
V100R001**

# **White Paper**

**Issue**                    08  
**Date**                     2018-07-13



**Copyright © Huawei Technologies Co., Ltd. 2018. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI 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 MZ912 in terms of its functions, appearance, features, applications, and technical specifications. You can obtain comprehensive information about the MZ912 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 Conventions

Symbol	Description
 <b>DANGER</b>	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.
 <b>WARNING</b>	Indicates a potentially hazardous situation which, if not avoided, could result in death or serious injury.
 <b>CAUTION</b>	Indicates a potentially hazardous situation which, if not avoided, may result in minor or moderate injury.

Symbol	Description
 <b>NOTICE</b>	Indicates a potentially hazardous situation which, if not avoided, could result in equipment damage, data loss, performance deterioration, or unanticipated results. NOTICE is used to address practices not related to personal injury.
 <b>NOTE</b>	Calls attention to important information, best practices and tips. NOTE is used to address information not related to personal injury, equipment damage, and environment deterioration.

## Change History

Issue	Date	Description
08	2018-07-13	<ul style="list-style-type: none"><li>● This issue is the eighth official release.</li><li>● Modified the section in <a href="#">3 Applications</a>.</li><li>● Removed the section Compatible Compute Nodes and I/O Modules.</li></ul>
07	2016-11-21	This issue is the seventh official release.
06	2016-08-17	This issue is the sixth official release.
05	2016-05-15	This issue is the fifth official release.
04	2016-03-31	This issue is the fourth official release.
03	2015-10-19	This issue is the third official release.
02	2015-05-30	This issue is the second official release.
01	2015-03-20	This issue is the first official release.

---

# Contents

---

<b>About This Document.....</b>	<b>ii</b>
<b>1 Overview.....</b>	<b>1</b>
1.1 Functions.....	1
1.2 Appearance.....	1
<b>2 Features.....</b>	<b>4</b>
2.1 Feature List.....	4
2.2 Feature Description.....	4
2.3 Standards Compliance.....	5
<b>3 Applications.....</b>	<b>7</b>
3.1 Software and Hardware Compatibility.....	7
3.2 Networking.....	7
<b>4 Technical Specifications.....</b>	<b>8</b>
4.1 Technical Specifications.....	8
<b>A Acronyms and Abbreviations.....</b>	<b>10</b>

# 1 Overview

---

## 1.1 Functions

The MZ912 is a 4-port multi-purpose NIC. It is used for E9000 compute nodes, and provides two 10GE ports and two 8 Gbit/s Fibre Channel (FC) or 10 Gbit/s Fibre Channel over Ethernet (FCoE) ports for connecting to switch modules in the chassis.

The MZ912 uses the Intel 82599 chip and Emulex XE201 chip. The Intel 82599 chip provides two physical 10GE ports and supports NIC applications. The XE201 chip provides two 8 Gbit/s FC or 10 Gbit/s FCoE ports (the port mode is configurable), and supports 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 MZ912 provides independent physical ports to connect the service network and storage network, implementing physical isolation between the two networks.

## 1.2 Appearance

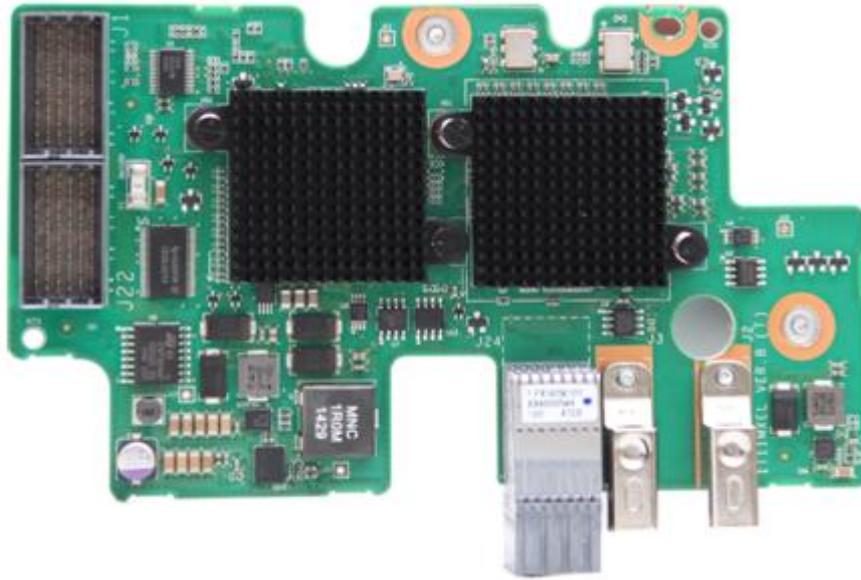
The MZ912 can be installed in slot Mezzanine1 (Mezz1 for short) or Mezzanine2 (Mezz2 for short) on a half-width compute node or in slot Mezz1, Mezz2, Mezzanine3 (Mezz3 for short), or Mezzanine4 (Mezz4 for short) on a full-width compute node.

### NOTE

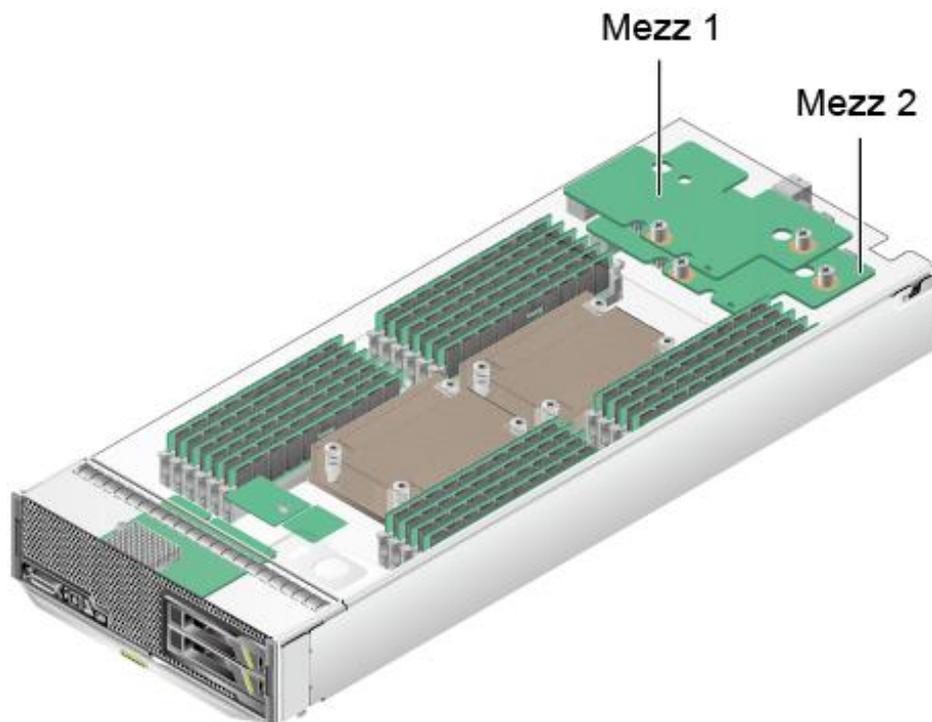
The MZ912 must be in slot Mezz1 when installed on the CH242 compute node.

- When the MZ912 is installed in slot Mezz1 or Mezz3, its two 10GE ports and two 8 Gbit/s FC (or 10 Gbit/s FCoE) ports connect to switch modules in slots 2X and 3X.
- When the MZ912 is installed in slot Mezz2 or Mezz4, its four ports connect to switch modules in slots 1E and 4E.

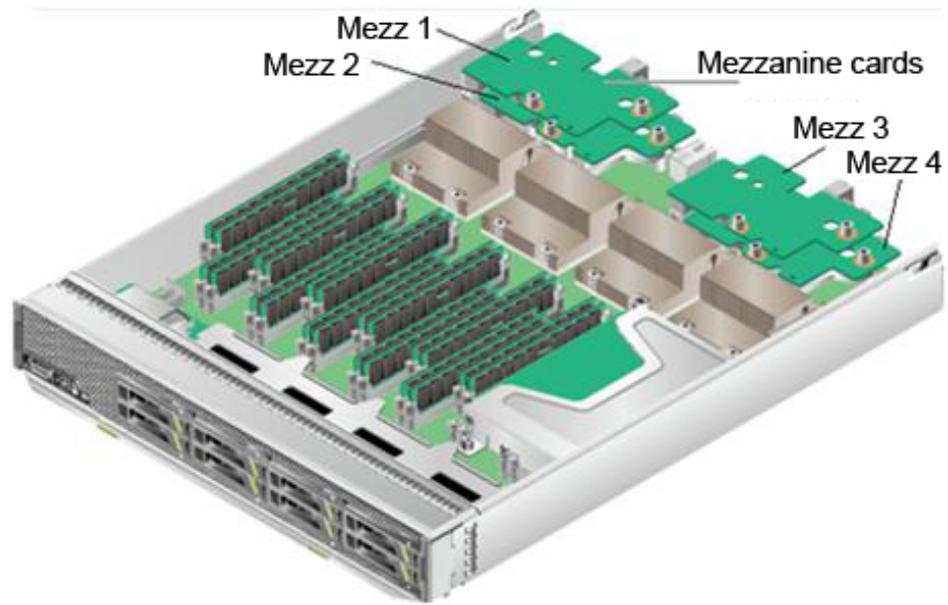
**Figure 1-1** MZ912 appearance



**Figure 1-2** MZ912 installation positions on a half-width compute node



**Figure 1-3** MZ912 installation positions on a full-width compute node



# 2 Features

## 2.1 Feature List

The MZ912 supports the following features and performance specifications:

- FC and FCoE
- N\_Port ID Virtualization (NPIV), supporting a maximum of 255 virtual N\_Port IDs per FC or FCoE port
- Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), Data Center Bridging (DCB), and Data Center Bridging Exchange (DCBX)
- Single-root I/O virtualization (SR-IOV), supporting a maximum of 128 virtual functions (VFs)
- Virtual Machine Device Queue (VMDq)
- Ethernet Preboot Execution Environment (PXE), FC/FCoE storage area network (SAN) boot
- 802.1Q VLAN, supporting a maximum of 4094 VLANs
- TCP Checksum Offload, UDP Checksum Offload, TCP segmentation offload (TSO), Receive Side Scaling (RSS), and Receive Side Coalescing (RSC)
- Jumbo frames of 9.5 KB
- Data Plane Development Kit (DPDK)

 **NOTE**

If the flow control function is configured for the NIC, this function must also be configured for the peer device (switch module or external switch).

## 2.2 Feature Description

### Priority Flow Control (PFC)

The 10GE ports on the MZ912 support the PFC feature. This feature supports flow control and back pressure on a per class-of-service (CoS) basis. The MZ912 supports a maximum of eight CoS queues. If both the MZ912 and its connected peer device support PFC, the link between them supports PFC. The MZ912 can identify PFC frames sent by the peer device and

adjusts the transmitting bandwidth. If the inbound traffic exceeds the threshold, the MZ912 sends PFC frames to instruct the peer device to control the traffic.

The MZ912 can negotiate with the peer device about PFC parameters over DCBX.

## Enhanced Transmission Selection (ETS)

The 10GE ports on the MZ912 support the ETS feature. This feature assigns port transmitting bandwidth based on priority groups (PGs). The MZ912 supports a maximum of eight CoS queues and eight PGs (BWGs). 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 MZ912 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 MZ912 does not tag or untag packets, but transparently transmits them. VLAN IDs are specified by the OS on an E9000 compute node.

## N\_Port\_ID Virtualization (NPIV)

The FC or FCoE ports on the MZ912 support the NPIV feature. 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.

## Single Root I/O Virtualization (SR-IOV)

The 10GE ports on the MZ912 support the SR-IOV feature. The NIC supports a maximum of 128 VFs or two physical functions (PFs) and 126 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.

## PXE and SAN Boot

The 10GE ports on the MZ912 support PXE, and the FC or FCoE ports on the MZ912 support SAN Boot.

- 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.
- This feature 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.

## 2.3 Standards Compliance

**Table 2-1** lists the standards and protocols that the MZ912 complies with.

**Table 2-1** Standards and protocols that the MZ912 complies with

<b>Standard</b>	<b>Protocol</b>
IEEE 802.3x	Flow Control and Back Pressure
IEEE 802.3z	1000BASE-X
IEEE 802.3ap	10GBASE-KR
IEEE 802.1Q	VLAN Tagging
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)
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
DCBX	Data Center Bridging Exchange

# 3 Applications

## 3.1 Software and Hardware Compatibility

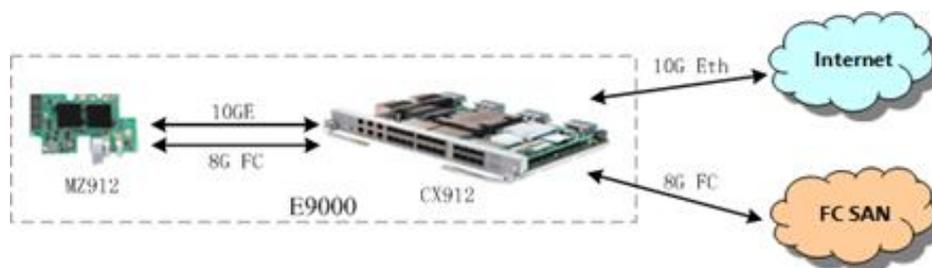
For details about the software and hardware that are compatible with the MZ912, see [Huawei Server Compatibility Checker](#).

## 3.2 Networking

The MZ912 can connect to I/O modules (switch modules or pass-through modules) to provide Ethernet and FC (or FCoE) services.

When the MZ912 works with the CX912 switch module, the MZ912 works in NIC+FC mode. The two 10 Gbit/s ports on the MZ912 provide 20 Gbit/s interface bandwidth to support the Ethernet service. The MZ912 connects to the Ethernet switch module on the CX912 and then to the Internet through 10GE ports on the CX912. The two 8 Gbit/s ports on the MZ912 provide 2 x 8 Gbit/s interface bandwidth to support the FC storage service. The MZ912 connects to the FC switch module MX210 on the CX912 and then to the external FC SAN through 8 Gbit/s FC ports on the CX912. See [Figure 3-1](#).

**Figure 3-1** Connection between the MZ912 and the CX912



# 4 Technical Specifications

## 4.1 Technical Specifications

**Table 4-1** lists the technical specifications for the MZ912.

**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	0.16 kg
Maximum power consumption	22 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)
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 900 m (2952 ft.) When the device is used at an altitude of 900 m to 5000 m (2952 ft. to 16404 ft.), the highest operating temperature decreases by 1°C (1.8°F) for every increase of 300 m (984 ft.).
Corrosive air pollutant	<ul style="list-style-type: none"><li>● Corrosion rate of the copper test piece: &lt; 300 Å/month (in compliance with the ANSI/ISA-71.04-2013 gaseous corrosion level G1)</li><li>● Corrosion rate of the silver test piece: &lt; 200 Å/month</li></ul>

Item	Specifications
Particulate pollutant	<ul style="list-style-type: none"><li>● The ISO14664-1 Class 8 requirements are met. You are advised to ask a professional organization to monitor particulate pollutants in the equipment room.</li><li>● There is no explosive, conductive, magnetic, or corrosive dust in the equipment room.</li></ul>
Number of ports	4
Port type	Ethernet and FC/FCoE
Chip model/ manufacturer	82599/Intel and XE201/Emulex

# A Acronyms and Abbreviations

<b>B</b>	
BWG	bandwidth group
<b>C</b>	
CoS	class of service
<b>D</b>	
DCB	Data Center Bridging
DCBX	Data Center Bridging Exchange
DPDK	Data Plane Development Kit
<b>E</b>	
ETS	Enhanced Transmission Selection
<b>F</b>	
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet
<b>L</b>	
LACP	Link Aggregation Control Protocol
LLDP	Link Layer Discovery Protocol
<b>N</b>	
NIC	network interface card
NPIV	N_Port ID Virtualization

<b>O</b>	
OS	operating system
<b>P</b>	
PCIe	Peripheral Component Interconnect Express
PF	physical function
PFC	Priority-based Flow Control
PG	priority group
PXE	Preboot Execution Environment
<b>R</b>	
RSC	Receive Side Coalescing
RSS	Receive Side Scaling
<b>S</b>	
SAN	storage access network
SR-IOV	single-root I/O virtualization
<b>T</b>	
TSO	TCP segmentation offload
<b>V</b>	
VF	virtual function
VLAN	virtual local access network
VM	virtual machine
VMDq	Virtual Machine Device Queue