



## **Dolphin Interconnect Solutions**

### **eXpressWare**

### **Software release notes**

### **for MX products**

**5.23.0**

**Date: January 30th, 2025.**

# Table of contents

Definitions.....	3
Release categories .....	3
Release version scheme .....	4
Types of releases .....	4
Release candidate .....	4
Release .....	4
Dolphin eXpressWare .....	5
SuperSockets.....	5
SuperSockets User-space Library for Windows and Linux .....	5
SuperSockets Kernel Space library for Linux.....	5
Dolphin technical support and documentation.....	6
Supported Hardware and configurations.....	7
Supported Dolphin products .....	7
Configurations supported by this release .....	7
Standard BMC firmware revisions, MXH83x .....	9
Standard BMC firmware revisions, MXH93x .....	9
Standard BMC firmware revisions, MXH94x .....	9
Standard BMC firmware revisions, MXH95x .....	10
Standard BMC firmware revisions, MXH53x .....	10
Standard configuration file revisions .....	10
Support for OEM hardware.....	10
Dolphin eXpressWare MX release 5.23.0 Linux .....	12
Dolphin eXpressWare MX release 5.23.0 Windows.....	20
Dolphin eXpressWare MX release 5.23.0 RTX64 .....	24
Dolphin eXpressWare MX release 5.21.0 VxWorks.....	26
Dolphin eXpressWare MX release 5.23.0 QNX .....	28

# Definitions

## Release categories

The following definitions are used for the eXpressWare software.

### Prototype:

Only parts of the full functionality are implemented. No or little error handling. The purpose of a prototype release is to demonstrate some limited important functionality to a selected number of customers. Should normally not be made generally available.

### Alpha:

All functionality implemented. Limited error handling. Can be released to collaborating partners.

### Beta:

Error handling and performance optimization completed. Very close to the final product. Can be made available on a general basis.

### GA:

General Availability. Well tested product that can be shipped to OEM customers and end users.

### Snapshot:

A snapshot release (or code drop) can be done at any of the above defined release categories. Code distributed by an individual, as a part of an agreed collaboration etc. for the purpose of testing a fix or new functionality requested by the other party. Snapshot releases are generally not supported by Dolphin, but used as an important part of the test and qualification of new code. The person doing the snapshot is responsible for making sure the letter «d» (for development) is added to the version string before distributing the software and to make sure the other party understands the terms (not supported) and conditions of the software release. The person doing the snapshot release should make sure to keep an exact copy of the release (by a private copy or source control tag), no other formal steps are required.

## **Release version scheme**

All eXpressWare components have a version string that easily can be retrieved by users to identify the product and the release version. The version string is constructed using decimal numbers formatted like MAJOR\_VER.MINOR\_VER.BUGFIX\_VER (E.g. DIS 5.13.1).

### **MAJOR\_VER**

Major changes to the software. Significant improvements or changes that may require changes to how the product is used.

### **MINOR\_VER**

Changes to the software that significantly improve functionality. Minor changes and improvements that do not affect general usage of the product. Fixes to support new versions of operating systems. Bug fixes.

### **BUGFIX\_VER**

Minor Bug fixes to a previously released software. No functionality or general improvements are allowed.

The letter “d” should always be added to the version string while the code is being implemented to specify that the running code is an internal development version. Snapshot releases should always display the letter “d”.

## **Types of releases**

Two types of releases are used by Dolphin to satisfy rapid development, concurrent external testing and stability and long-term test requirements. These ideas are collected from and widely used within the Linux community. These two types apply to Prototype, Alpha, Beta and GA releases.

## **Release candidate**

A release candidate is a release that is still undergoing testing and qualification. The purpose of distributing release candidates is to enable customers and partners to have access to updated drivers without having to wait for the final release. The only difference between a release candidate and a release is the amount of testing performed. The release note should specify if this is a release candidate or a completely tested release. Critical bugs found during release candidate testing may cause the release to be canceled. In such cases, the version date/string of the release must be changed before a new release candidate can be produced following the general release engineering steps.

## **Release**

A release is a product that have completed the full release procedure.

Normally a release candidate will be renamed to a release if all tests are passing, and no new critical bugs are found. This will be reflected in the release note.

## Dolphin eXpressWare

The eXpressWare software suite includes the following components. Please note that all combinations of components, PCIe cards and operating systems may not be supported. Please refer to the release note for each operating system and card for details.

IRM	Interconnect Resource Manager
SISCI	Software Infrastructure for Shared-Memory Cluster Interconnects. This includes binaries , sources, documentation and development tools.
SuperSockets	Socket accelerator for PCIe
IPoPCIe	TCP-IP driver for PCIe. Windows NDIS driver or Linux DISip.
SmartIO	Sharing and access to IO management and control. Includes Device Lending . and SISCI SmartIO management. Supports Hot-Add with Linux.

### SuperSockets

Dolphins SuperSockets is a family of Berkeley Sockets API compliant libraries that will accelerate embedded applications written to standard networking functionality. The functionality depends on the implementation approach.

### SuperSockets User-space Library for Windows and Linux

The Dolphin SuperSockets user space library is currently provided for the purpose of supporting embedded applications. The Windows Winsock2 environment contains a rich set of socket functions and options. SuperSockets version 5.2 and newer adds support for connectivity to the new Linux SuperSockets user space library. We have verified and support the following socket calls:

accept, bind, connect, getpeername, getsockname, getsockopt, listen, recv, select, send, setsockopt(TCP\_NODELAY), socket, shutdown.

For Windows OS, specific functions are implemented: closesocket, WSASendDisconnect, ioctlsocket, WSAAsyncSelect, WSAGetOverlappedResult, WSAIoctl, WSARecv, WSASelect and WSASendDisconnect.

For Linux OS, normal file operations are implemented: close, ioctl.

### SuperSockets Kernel Space library for Linux

The SuperSockets Kernel-space library is our standard recommendation for most standard applications. It implements a new socket address family AF\_SSOCKS, and support virtually all networked Linux applications. The software is also compliant with the Linux Kernel Sockets API and can also be used by kernel services that can be configured to use AF\_SSOCKS (decimal value 27). This version of SuperSockets includes automatic fail-over to Ethernet if there is a failure with the PCIe network.

UDP multicast is supported if the underlying PCIe hardware and topology supports multicast. Currently, only a single receiver per node for a specific multicast group is supported. The number of available groups is hardware dependent.

More details on the eXpressWare software can be found at: <https://www.dolphinics.com/software>

**Dolphin technical support and documentation**

We do take software development and product testing seriously, please let us know your experience or any issue by sending an email to [pci-support@dolphinics.com](mailto:pci-support@dolphinics.com).

Additional information, installation manuals, adapter users guide etc. can be found at <http://www.dolphinics.com/mx>

SISCI API resources can be found at [www.dolphinics.com/products/embedded-sisci-developers-kit.html](http://www.dolphinics.com/products/embedded-sisci-developers-kit.html)

This release note contains a summary of the important changes made to eXpressWare. Please contact Dolphin for a complete list of changes.

## Supported Hardware and configurations

### Supported Dolphin products

This software release supports the following PCIe adapter cards.

#### The NTB software supports the following Dolphin adapter cards

- MXH530
- MXH930
- MXH94x / MXH95x
  - 0 (2 x8 FireFly modules, x16, 2 x8)
  - 1 (1 x8 FireFly module x8)
  - 3 (2 x4 FireFly modules x8, 2 x4)
  - 4 (1 x4 FireFly module x4)
  - 9 (4 x4 FireFly modules x16, 2 x8, 4 x4)
- MXH830
- MXC948
- MXH832, MXH833 (As transparent target only, PCIe hot add)
- MCH663
- MXH910
- MXH914
- MXH916
- MXH918
- MXP908
- MXC960

#### The Transparent Board Management software supports the following adapter cards

- MXH932
- MXH94x / MXH95x
  - 2 (2 x8 FireFly modules, x16, 2 x8)
  - 6 (1 x8 FireFly module x8)
  - 7 (2 x4 FireFly modules x8, 2 x4)
  - 8 (1 x4 FireFly module x4)
  - 5 (4 x4 FireFly modules x16, 2 x8, 4 x4)
- MXP924
- MXH832
- MXH912
- MXH915
- MXH917
- MXH919
- MXP909
- MXC962

Please choose another software download if you are using a not listed adapter card.

### Supported Dolphin PCI Express switches

- MXS824
- MXS924

### Configurations supported by this release

The software release supports the following configurations.

#### General use with MXH840:

- 2 nodes using 1, 2, 4 Samtec Gen3 FireFly / cables.

#### General use with MXH850:

- 2 nodes using 1, 2, 4 Samtec Gen3 FireFly / cables.

**General use with MXH940:**

- 2 nodes using 1, 2, 4 Samtec Gen4 FireFly / cables.

**General use with MXH950:**

- 2 nodes using 1, 2, 4 Samtec Gen4 FireFly / cables.

**General use with MXH910:**

- 2 nodes using 1 x8 iPass cable.

**General use with MXH830:**

- 2 nodes using 1, 2, 4 (x4, x8, x16) PCIe 3.0 cables.
- 3 nodes using 1, 2 (x4 or x8) PCIe 3.0 cables.
- 5 nodes using 1 (x4) PCIe 3.0 cables.

**General use with MXH930:**

- 2 nodes using 1, 2, 4 (x4, x8, x16) PCIe 4.0 cables.
- 3 nodes using 1, 2 (x4 or x8) PCIe 4.0 cables.
- 5 nodes using 1 (x4) PCIe 4.0 cables.

**General use with MXH530:**

- 2 nodes using 1, 2, 4 (x4, x8, x16) PCIe 5.0 cables.
- 3 nodes using 1, 2 (x4 or x8) PCIe 5.0 cables.
- 5 nodes using 1 (x4) PCIe 5.0 cables.

**Reflective memory use with MXH830 and MXS824:**

- 2 nodes using 1, 2, 4 (x4, x8, x16) PCIe 3.0 cables.
- Up to 6 nodes x16 using one MXS824 switch.
- Up to 12 nodes x8 using one MXS824 switch.
- Up to 24 nodes x4 using one MXS824 switch.

**Reflective memory use with MXH930 and MXS924:**

- 2 nodes using 1, 2, 4 (x4, x8, x16) PCIe 4.0 cables.
- Up to 6 nodes x16 using one MXS924 switch.
- Up to 12 nodes x8 using one MXS924 switch.
- Up to 24 nodes x4 using one MXS924 switch.
- Up to 30 nodes x16 using multiple MXS924 switches.

**Reflective memory use with MXH530 and MXS924:**

- 2 nodes using 1, 2, 4 (x4, x8, x16) PCIe 4.0 cables.
- Up to 6 nodes x16 using one MXS924 switch.
- Up to 12 nodes x8 using one MXS924 switch.
- Up to 24 nodes x4 using one MXS924 switch.
- Up to 30 nodes x16 using multiple MXS924 switches.

**PCI Express Hot Add with MXH830 (Linux Only)**

- MXH830 connected to MXH832 transparent target cards and endpoints.
- MXH830 connected to MXS824, fan out to max 5 MXH832 transparent cards and endpoints.

**PCI Express Hot Add with MXH530, MXH930, MXH940 and MXH950 (Linux only)**

- MXH530 connected to any transparent target cards and endpoints.
- MXH930 connected to any transparent target cards and endpoints.



- NTB MXH94x connected to any transparent target cards and endpoints.
- NTB MXH95x connected to any transparent target cards and endpoints.

#### **Standard BMC firmware revisions, MXH83x**

DIS release / Adapter	MXH83x
5.23.0	8.12
5.22.0	8.12
5.21.x	8.12
5.20.0 / 5.20.1	8.12
5.19.2	8.12
5.18.0	8.12
5.17.0	8.12
5.16.0	8.10
5.15.0	8.10
5.14.0	8.10
5.13.1	8.8

#### **Standard BMC firmware revisions, MXH93x**

DIS release / Adapter	MXH93x
5.23.0	2.10
5.22.0	2.10
5.21.x	2.8
5.20.0	2.8
5.19.2	2.1/2.4
5.18.0	1.5
5.17.0	1.0

#### **Standard BMC firmware revisions, MXH94x**

DIS release / Adapter	MXH94x
5.23.0	2.11
5.22.0	2.10
5.21.x	2.8
5.20.0	2.8
5.19.2	2.1

### Standard BMC firmware revisions, MXH95x

DIS release / Adapter	MXH95x
5.23.0	2.11
5.22.0	2.10
5.21.x	2.8
5.20.0	2.8
5.19.2	2.1

### Standard BMC firmware revisions, MXH53x

DIS release / Adapter	MX530x
5.23.0	3.3

### Standard configuration file revisions

DIS release / Adapter	MXH930	MXH94x	MXH95x	MXP908	MXP924
5.23.0	VER13	VER12	VER9	VER3	VER9
5.22.0	VER13	VER11	VER7	VER3	VER8
5.21.x	VER12	VER11	VER7	VER3	VER7
5.20.0	VER11	VER10	VER6	VER3	VER7

DIS release / Adapter	MXH910	MXH912	MXH914 MXH915	MXH916 MXH917	MXH918 MXH919
5.23.0	VER6	VER3	VER4	VER4	VER4
5.22.0	VER6	VER3	VER4	VER4	VER4
5.21.x	VER1	VER1	VER2	VER2	VER1

DIS release / Adapter	MXP908	MXP909	MXC96x	MXH530
5.23.0	VER3	VER4	VER2	VER2

### Support for OEM hardware

This version of Dolphin eXpressWare has general support for the Microchip PFX Switchtec Gen3, Gen4 and Gen5 chipsets in various configurations and topologies. Please contact Dolphin for licensing information if you would like to run eXpressWare on your own hardware.

If you already are running eXpressWare on non-Dolphin hardware, please contact your hardware vendor for additional information on how to upgrade your software.

### General eXpressWare defaults and settings for MX

- SISCIMulticast
  - Default setting 4 groups, 2 Megabyte per segment.

- Max setting 16 groups, 128 Gigabyte per segment.

Please consult the “Dolphin eXpressWare Installation and Reference Guide, section “Managing PCIe and eXpressWare Resources” for information how to tune eXpressWare parameters.

### **Support for 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> gen Intel Xeon Scalable processors.**

- Intel Xeon Scalable 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> gen processors are supported fully in back-to-back configurations.

Certain configurations of Intel Xeon Scalable 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> gen processors require special configurations at the current point in time. Please contact Dolphin support at <http://www.dolphinics.com/csp> if you plan to use Intel Xeon Scalable 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> gen processors in the following ways:

- You plan on connecting more than six Intel Xeon Scalable 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> gen processors to a MXS824 or MXS924 switch in x8 or x4 mode.
- You plan on using any of the following Intel Xeon Scalable 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> gen Xeon Platinum processors in x8 or x4 mode:

Intel Xeon Platinum 8380	Intel Xeon Platinum 84xx	Intel Xeon Platinum 85xx.
--------------------------	--------------------------	---------------------------

## Dolphin eXpressWare MX release 5.23.0 Linux

Description of content: Clustering package for Linux.

Release category: GA

Target audience: Production systems.

Release date: January 30th, 2025.

**Current status:** Available for download from <https://www.dolphinics.com/mx>

**NOTE: See the section about support for Intel Xeon Scalable 3rd, 4th and 5th gen processors if you plan to use any of these.**

### New in DIS MX 5.23.0

- Added support for Gen 5 adapter MXH530 in the driver and tools
- Added support for MXC960
- Added driver support for MXS924 cascade
- Added support for license file
- SISI:
- Added interrupt support for SISI SmartIO
- Enabled possibility for non-zero offset when calling SCIMapRemoteSegment for multicast segments
- Bug fix: Fix for SCISStartDMATransferMem deadlock
- SmartIO:
- Added interrupt support for SISI SmartIO
- Added support for RHEL9-based Linux distributions
- Fixed regression for devices with IO BARs
- Improved support for device lending of GPUs
- Improved smartio\_tool:
  - Now makes all devices in the same IOMMU group available when a device in the group is made available
  - Prevent adding devices upstream from a Dolphin device
  - Bug fix: Match hexadecimal BDF
- Improved dis\_diag:
  - Bug fix: No longer prints invalid temperature
  - Bug fix: Fixed info readout for all firefly adapters
- Improved dis\_nodemgr:
  - Unresolvable hosts are skipped and no longer stop cluster configuration
- Improved dis\_services:
  - Bug fix: dis\_services now properly handles services that end up in a failure state
- Improved installer:
  - Better error messages are printed if something goes wrong
  - Bug fix: No longer loops if the --batch option is passed when no configuration file is present
- Bug fix: Fix issue with MXP924 reporting “Hot Plug Capability”
- Bug fix: IOMMU now works as expected on kernels >= 6.3
- Bug fix: Driver configuration now correctly exits if IOMMU initialization fails

### Changes in previous releases

#### MX 5.22.0

- Added support for Intel Xeon Scalable 5<sup>th</sup> gen (“Emerald Rapids”).
- Improved support for Intel Xeon Scalable 3<sup>rd</sup>, 4<sup>th</sup> and 5<sup>th</sup> generation CPUs.
  - Requester IDs for the CPU and link are now determined dynamically by the NTB driver.
  - Added support for requester ID scaling with hyperthreading disabled.

- Added support for requester ID scaling with different NUMA node configurations.
- Added support for Linux kernel 6.8
- Added support for Ubuntu 24.04
- Added support for Ubuntu 22.04 HWE
- Added support for pre-retbleed Centos 7.
- Added support for Centos 9 and Centos 9 RT.
- MXS924: New firmware version 5.
- MXS824: New firmware version 10.
- MXP924: Upgraded to NXP version 3.2
  - Added support for 4x4 configuration
- MXH91x: Applied newest Microchip patch to resolve boot issues in certain systems.
- Improved dis\_diag.
  - Added support for json-output.
  - Added support for showing root port link capabilities.
  - Improved reporting for central PFX switch being reset.
  - Throttling events are now separated into relevant categories.
  - Improved link error detection.
  - Now shows temperature information for Firefly.
- Improved installer.
  - Bugfix in transparent configurations for Rocky 8.
  - Various bugfixes.
- Improved dis\_status script, now retrieves more information.
- Improved dis\_networkmanager.
  - Added support for limiting maximum log file size.
  - Fixed CPU usage issue related to large log file sizes.
- Improved dis\_firmwaretool.
  - Added support for virtual configurations.
  - Various bugfixes.
- SmartIO: Increased default maximum read request size.
- Added support for sanitization in upgrade-eeprom.

#### **DIS MX 5.21.4**

- Driver now recalibrates to support both older and newest MXS824 and MXS924 switch configurations.

#### **DIS MX 5.21.3**

- Improved device detection for MX adapters.

#### **DIS MX 5.21.2**

- Identical to 5.21.1 for Linux

#### **DIS MX 5.21.1**

- Added support for all Intel Xeon Ice Lake and Sapphire Rapids Bronze and Silver CPUs.
- Added support for all Intel Xeon Ice Lake and Sapphire Rapids Gold and Platinum CPUs with up to two sockets, three UPI links and 36 cores.

#### **DIS MX 5.21.0:**

- Added support for MXH914, MXH915, MXH916, MXH917, MXH918 and MXH919
- Added support for MXH910 and MXH912 adapter with IPASS connector
- Added support for Linux kernel 5.19
- Added support for Linux kernel 6.0
- Added support for MXH833 in dis\_firmwaretool
- MXS924:
  - Eeprom version updated to VER5
  - Updated throttling values for DSWITCH topology
- Added support for AGX ORIN
- Added configurations for adapter DSWITCH topology for MXH930, MXH940 and MXH950

- Improved session handling after re-plugging cables.
- EEPROM firmware version updates:
  - Extended the requester ids table for Ice Lake and Sapphire Rapids.
  - Added 64GB NTB configuration
  - Added 128GB NTB configuration
  - MXH930 updated to VER12
  - MXH940-C updated to VER11
  - MXH950-C updated to VER7
- Added preliminary support for Intel Ice Lake and Sapphire Rapids Xeon CPUs in DIRECT topology. Limited support for MXS824/MXS924 switches.
- Added support for Secure Boot
- Fixed MSI-X support regression for device lending on newer kernels (4.15 and earlier)
- SmartIO:
  - Fixed SmartIO showing “Illegal Vendor ID” 0xFF for virtual functions
  - Fixed kernel warning ‘proc\_dir\_entry already registered’
- Improved installer functionality
- Fixed x86 Spectre Mitigation V2-2 breaking compilation for Centos/Ubuntu/Debian with symbol `__x86_return_thunk`. This will affect all **RHEL7** before **kernel-3.10.0-1160.83.1** using installer built with a kernel-3.10.0-1160.83.1 or newer.
- Fixed installer issue related to Cuda support on RHEL7
- Fixed installer issue for Fedora related to broken RUNPATH
- Improved Xavier/Ubuntu/Debian package platform-integration
- Fixed ‘cache line size’ detection in NT-Transparent topology

#### **DIS MX 5.20.1:**

- Initial support for CentOS 9
- Added 'make' package as required dependency
- Added support for Ubuntu 22.04
- Added support in the upgrade\_eeprom script to prevent an upgrade for customer specific adapters, e.g. (-0101-)
- Improved platform integration for CentOS 8 and CentOS 9 (SELinux)
- Added support for -Os optimization level
- Add support for IPv4-mapped IPv6 addresses in kernel SuperSockets
- SmartIO: Added implementation of dma\_mmap\_coherent
- Extended dis\_diag with register information
- Bug fixes and improvements
  - Fixed SuperSockets build\_RPM failure with certain distros due to missing Module.symvers from GX
  - Fixed GUI tools regression with QString .arg() usage
  - Fixed problem with SuperSockets getsockopt causing kernel error
  - Fixed regressions from 5.20.0 causing crash during driver load or link reset on some systems
  - Fixed dis\_config reordering issue during configuration
  - Fixed an issue with cancelling callback threads that caused SISI applications to deadlock
  - Fixed DISIP not working in D-Switch topology with transparent MX host card

#### **DIS MX 5.20.0:**

- Added support for MXH941, MXH943, MXH944, MXH945, MXH946, MXH947, MXH948 and MXH949
- Added support for MXH951, MXH953, MXH954, MXH955, MXH956, MXH957, MXH958 and MXH959

- Changed the adapter serial number format from MXH9xx-AA-123456 to MXH9xx-AA-0001-123456
- Updated configuration files for MXH930, MXH94x, MXH95x to MR4
- Added support for MXS924 switch topologies:
  - NTB
  - Transparent
  - D-SWITCH
- Added support for MXC960
- Added support for Linux 5.16
- Added support for Linux 5.18
- Added support for MXH330 adapter
- Add support for NXP 2.8 for MXH9xx, MXP9xx and MXC9xx adapters
- Added support for MXP908 and MXP909
- Added support for MVH222
- Added support for MVH242
- Added support for dis\_firmwaretool on the MXS924 switch
- Added support for proxy id changes in PFX Gen4 MR4 firmware release
- Added driver support for evaluation board from Microchip PM42100-KIT
- Added support for 3-slot and 5-slot backplane configuration for MXH930, MXH94x, MXH95x
- Added support for MXH330 in DIRECT topology
- Added support for 12 port x8 in DSWITCH mode for MXS924
- Added support for MCU PFX config format version 2
- Add SW lock synchronize partition operations
- Added option "-userid" for example programs reflective\_write and reflective\_bench to allow multiple concurrent instances
- dis\_firmwaretool: Added option to patch IDEEPROM files
- Bug fixes and improvements
  - Fixed incorrect BDF being used for outbound DMA RID entry when P2P is not enabled in USP
  - Fixed problem with transparent MXS924 switch being erroneously detected as adapter
  - Fixed GPIO reset pin assignment for MXH940 and MXH950 adapters for 2x8 NTB and 4x4 NTB configuration
  - Fixed problem with traffic interruption to downstream devices when the driver is being loaded by preventing BME from being cleared
  - Fixed problem with unbinding of DMA functions from plx\_dma driver not working on Ubuntu
  - Fixed problem with module parameter ntb\_mcast\_alloc\_groups not working as expected
  - Improved handling of memory allocation failures during driver startup
  - Fixed problem with MPS exceeding device capabilities under rare conditions
  - Made changes to the installer to install QT5 for Ubuntu 20.04 and newer
  - Made changes to the installer to install libgcc-s1 instead of libgcc1 for Debian 10 and newer
  - Reinsert outbound DMA RID during link setup to avoid Gen4 DMA proxyid collision
  - Created workaround to solve issue with IOMMU attachment for > Linux 5.8
  - Fixed regression with driver not loading properly on PowerPC
  - Added p2p\_keeper in sysfs
  - RTSS: implement osif\_kmem\_get\_contig2 on top of RtAllocateContiguousMemorySpecifyCache to improve large memory alignment and allocation
  - RTSS: update build files to Visual Studio 2019
  - Added handling of AER events for type GX\_EVENT\_LINK\_RECEIVER\_ERROR
  - Improved link status output for dis\_diag
  - Fixed problem with shared segments not being unexported when

- removed.(SCIPrepareSegment failed with SCI\_ERR\_BUSY for multicast segments.)
- Fixed problem with "dis\_firmwaretool pci:" not working with Linux 5.15
- Added TTY access mode to dis\_firmwaretool
- Fixed problem with dis\_firmwaretool showing the wrong version for the running PFX IMG under certain conditions
- Fixed problem with STM firmware update and log dump not working for MXH83x adapters since release 5.19.1

#### **DIS MX 5.19.2:**

- Added support for direct 3 and 5 node NTB topology for MXH930, MXH940, MXH950.
- Added support for NT-only D-Switch ports.
- Added support for multicast receivers on D-Switch Downstream ports.
- Updated configuration files for MXH930, MXH940, MXH950:
  - MXH930 – VER10
  - MXH940-B – VER8
  - MXH940-C – VER8
  - MXH950-B – VER4
  - MXH950-C – VER4
- Fixed incorrect reporting of correctable errors.
- Added support for adapters MXP924, MXC948, MCH663
- General bug fixes and improvements.
  - Fixed crash due to missing check in DSWITCH mode.
  - Added –link-speed to Linux installer.
  - Made link watchdog ignore transparent and virtual (D-Switch) links.
  - Added optical support for MXS824 x4 configuration.
  - Fixed bug that caused local io address to be set incorrectly when using System DMA + IOMMU.
  - Fixed 32-bit compilation on 64-bit kernel.
  - Improved SISI callback status for DMA.
  - Fixed aborting GX/osif sysdma transfers triggering crashes.
  - More conservative MPS handling by default (ntb\_set\_mps=1) to prevent automatic increase.
  - Fixed problem with multicast not working in Gen3 D-Switch configurations.
  - Fixed AER Correctable Error Status offset definition.
  - Fixed problem on certain systems where BME in P2P is not set.
  - Fixed Multicast only working with group 0 on Gen3 D-Switch.
  - Fixed Problem with NMI on some systems when cable is unplugged.
  - Fixed potential small memory leak when using CUDA segments.
  - Linux installer for el7 accepts fedora22-24.
  - MSI: install Gen4 firmware files with Transparent setup option.
  - **Improved parameter checking for IRM tools to prevent crash.**

#### **DIS MX 5.18.0:**

- Improved support for Gen4 adapters MXH930, MXH932, MXH940, MXH942, MXH950, MXH952. New firmware bundle.
- Initial support for MXP924, MXC940 and MXC948
- dis\_diag: Added reporting of board firmware versions and temperatures.
- Applied copy function patch for an unaligned crash issue on some ARM platforms.
- SISI
  - Adjusted the number of dev nodes from 4096 to 255.
- General bug fixes and improvements.

#### **DIS MX 5.17.0:**

- Added support for Gen4 adapters MXH930, MXH932, MXH940, MXH942
- SISI
  - Added support for send-only multicast configurations.
  - Fixed bug in segment cleanup for shared segments that could cause segment to be



- unavailable for all users.
- Increased max multicast segment size to 128 GB. Fixed 32 bit size limitations with internal copy function. Tested up to 4GB.
- Fixed bug related to local DMA (memory to memory) with systems with IOMMU ON.
- Fixed Ubuntu 18.04 installation issue.
- Add BDF information in dis\_diag and dis\_tool.
- General bug fixes and improvements.

#### **DIS MX 5.16.0:**

- Added support for system DMA with IOMMU ON (Off already supported)
- Added support for Linux kernel >= 5.5.3. Tested on Linux kernel 5.5.8
- SISI API
  - Added support for using multiple DMA channels.
  - Added support for registering more than one PCIe requester ID per SISI descriptor.
  - Fixed a crash when an invalid fdid was given to SCIBorrowDevice()
  - Allow SCIBorrowDevice() of PCI-PCI bridges in NT-Transp
- SmartIO
  - Fixed issue when some devices are behind an IOMMU, but others are not.
  - Fixed issues with rescan when link goes up / down with multiple adapters.
  - More robust scanning.
  - Fixed crash during device lending when amd\_iommu was enabled.
  - Fixed issue with RedHat 8 / new kernels Assertion (vdev->orig\_dma\_ops != NULL).
  - Fixed issue in NT-Transp detecting directly connected endpoints.

#### **DIS MX 5.15.2:**

- SmartIO – fixed potential assert in device scanning process

#### **DIS MX 5.15.1:**

- Fixed bug in SISI supporting larger than 4GB segments.
- Fixed multi endian issue with reflective memory test program.

#### **DIS MX 5.15.0:**

- Transparent driver support for MXH832
- Increased maximum number of adapters from 4 to 5.
- Support for GUI on Centos/RHEL 8 (Qt5)
- SmartIO 2.5
  - SmartIO SISI functions stabilized
  - Add support for `smartio\_tool remove`
  - Show ethernet device in `smartio\_tool show`
  - Support for multi-link in NT-Transp.
  - Fixed bug disabling NVidia persistence mode.
  - Fixed issue affecting SR-IOV VFs
- Fixed problem with slot AER errors not being logged.
- SIA Installer: Fixed problem with the --install-all option
- User-space SuperSockets:
  - Fixed problem with maximum number of connections decreasing over time.

#### **Included software:**

- SuperSockets (kernel)
  - Ultra-fast, low latency Berkeley Sockets API (TCP, UDP, UDP multicast) for PCIe.
  - Local and remote socket communication acceleration (accelerated loop-back device support local IP and localhost address).
  - This version of SuperSockets only supports communication to other Linux systems.

- SuperSockets user space
  - User space version of SuperSockets. Provides lower latency than the kernel version but proves a limited number of socket functionality. Currently only TCP.
  - Support communication to Windows SuperSockets.
- SISI API 2.0.2
  - Shared memory (DMA, PIO, RDMA, Interrupts).
  - Reflective memory/multicast.
  - SmartIO Extension 2.0
  - PCIe peer to peer communication (FPGAs, GPUs, NVMe etc.).
  - Full connectivity to other systems running Windows, RTX or VxWorks.
- SmartIO 2.1
  - Device Lending
  - Device Hot Add
  - SISI API SmartIO Extension.
- TCP/IP driver
  - Full IP networking over PCIe to other nodes running Linux.
- Network Installer (SIA)
  - Installs the above software on a cluster of computers interconnected with Ethernet.

### **Supported architectures:**

The software supports the following architectures and platforms:

- Intel x86 / x64
- ARM64
- Nvidia Xavier (only supported by MXH93x, MXH94x cards)
- NXP BlueBox

### **Tested and supported OS platforms:**

Dolphin strive to support all major Linux distributions and kernels from 2.6.32 and up. This release has been tested on these platforms but are expected to work on many more. Please let us know if you run into problems or need another kernel:

Linux Kernel 2.6.32 - 6.8

- CentOS 9 x86\_64
- CentOS 8 x86\_64
- CentOS 7 x86\_64
- CentOS 6 x86\_64
- CentOS 5 x86
- OpenSUSE Leap 15.1
- Ubuntu 14.04 x86\_64
- Ubuntu 16.04 x86\_64
- Ubuntu 18.04 x86\_64
- Ubuntu 20.04 x86\_64
- Ubuntu 22.04 x86\_64
- Ubuntu 24.04 x86\_64
- Debian 7 x86\_64
- Debian 8 x86\_64
- Debian 9 x86\_64
- Debian 10 x86\_64
- Debian 11 x86\_64
- Fedora 25 x86\_64

RedHawk Linux from Concurrent Real-Time is supported. Please contact Concurrent for access to drivers.

We also provide OpenRC init scripts as used by Gentoo and other distributions, please contact Dolphin Support for more detailed information.

**Cluster installation requirements:**

- All nodes connected and properly configured with Ethernet.
  - Non-Ethernet configurations optional, contact Dolphin.

**Installation and management:**

- RPM based via Self Installing Archive (SIA).
- Automatic configuration via Interconnect manager.
- Centralized monitoring and diagnostic.
  - Both Linux and Windows supported as GUI platform.

**High availability features:**

- Unplugging a cable will cause SuperSockets to fail over to Ethernet for all connections affected. SuperSockets will automatically re-establish communication when cable is inserted.

**Bundled (major) management tools:**

- dis\_diag (diagnostic tool)
- dis\_admin (cluster monitor and manager GUI)
- networkmanager (cluster configuration and maintenance demon)
- dis\_netconfig (configuration editor GUI)

**Bundled (major) demo tools:**

- Scibench2 (SISCI based PIO throughput benchmark)
- scipp (SISCI based ping pong benchmark)
- dma\_bench (SISCI based DMA benchmark)
- interrupt\_bench (SISCI based remote system interrupt benchmark)
- reflective\_bench (SISCI based reflective memory benchmark)
- latency\_bench (TCP latency benchmark, runs on any IP network including SuperSockets)

**Firmware upgrade:**

The Driver installer will detect and suggest to automatically update the firmware during installation if an update is required (new firmware may be included in software distribution). Please upgrade by running /opt/DIS/sbin/upgrade\_eeprom.sh. A reboot is required after firmware upgrade.

## Dolphin eXpressWare MX release 5.23.0 Windows

Description of content: Clustering package for Windows.

Release category: Production systems

Release date: January 30th, 2025.

**Current status:** Available for download from <https://www.dolphinics.com/mx>

### New in DIS MX 5.23.0

- Added support for Gen 5 adapter MXH530 in the driver and tools
- Added support for MXC960
- Added support for MXS924 cascade
- SISI:
  - Enabled possibility for non-zero offset when calling SCIMapRemoteSegment for multicast segments
  - Bug fix: Fix for SCISStartDMATransferMem deadlock
- Bug fix: Driver configuration now correctly exits if IOMMU initialization fails
- Improved dis\_diag:
  - Show PCIe parent link capability
  - Bug fix: No longer prints invalid temperature
  - Bug fix: Fixed info readout for all firefly adapters
- Improved dis\_nodemgr:
  - Unresolvable hosts are skipped and no longer stop cluster configuration.
- Bug fix: Updated BAR computation on transparent devices
- Bug fix: Fix issue with MXP924 reporting “Hot Plug Capability”
- Bug fix: IOMMU now works as expected on kernels  $\geq 6.3$
- Bug fix: Driver configuration now correctly exits if IOMMU initialization fails

### Changes in previous releases:

#### DIS MX 5.22.0

- Added support for Intel Xeon Scalable 5<sup>th</sup> gen (“Emerald Rapids”).
- Added support for Windows 11, Server 2022
- Improved support for daisy chain configurations.
- MXH930: Fixed a bug where Windows 10 Enterprise LTSC occupies BAR0 for the DMA.
- Improved dis\_diag.
  - Improved link error detection.
  - Now shows temperature information for Firefly.
  - Added support for json-output.
- Improved installer.
  - Fixed bug for transparent install on GUI installer.
  - Various bugfixes.
- Improved dis\_status script, now retrieves more information.
- Fixed logging issue with dis\_networkmanager.
- Added support for virtual configuration through dis\_firmwaretool.
- Added support for sanitization in upgrade-eeprom.

#### DIS MX 5.21.4:

- The driver now recalibrates to support both older and newest MXS824 and MXS924 switch configurations.

#### DIS MX 5.21.3:

- Improved device detection for MX.

#### DIS MX 5.21.2:

- Added support for all Intel Xeon Ice Lake and Sapphire Rapids Bronze and Silver CPUs.
- Added support for all Intel Xeon Ice Lake and Sapphire Rapids Gold and Platinum CPUs with up to two sockets, three UPI links and 36 cores.

#### **DIS MX 5.21.0:**

- Removed dis\_coinstaller from the driver package
- Fixed internal DMA channels being improperly computed
- Fixed BAR2 mapping on systems where a bogus BAR1 is allocated
- Fixes to prevent potential leaks during driver failure
- Enable option to sign drivers with a self-sign certificate
- Remove dishosts.conf during uninstallation of Windows drivers

#### **DIS MX 5.20.1:**

- Bug fixes and improvements
  - Reset the Completion Timeout on root ports in cases where more than 1 transparent adapter are installed
  - Fixed initial adaptno value for IRM transparent inf files

#### **DIS MX 5.20.0:**

- Bug fixes and improvements
  - Generate .nupkg for MX Windows x64 builds
  - Added TTY access mode to dis\_firmwaretool

#### **DIS MX 5.19.2:**

- dis\_firmwaretool
  - Added support for 'pci:' communicator.
- Increased max number of transparent adapters from 5 to 20.
- General bug fixes and improvements.
  - Fixed crash on management EP if there is not enough contiguous memory.
  - Reset completion timeout on Root Port for transparent adapters.
  - **Firmware upgrade**

#### **DIS MX 5.18.0:**

- Improved support for Gen4 adapters MXH930, MXH932, MXH940, MXH942, MXH950, MXH952.
- General bug fixes and improvements.

#### **DIS MX 5.17.0:**

- 5.17.0 for Windows was not released due to a regression test issue.
- Added support for Gen4 adapters MXH930, MXH932, MXH940, MXH942
- SISI
  - Added support for send-only multicast configurations.
  - Fixed bug in segment cleanup for shared segments that could cause segment to be unavailable for all users.
  - Increased max multicast segment size to 128 GB. Fixed 32 bit size limitations with internal copy function.

#### **DIS MX 5.16.0:**

- SISI API
  - Added support for using multiple DMA channels.
  - Added support for registering more than one PCIe requester ID per SISI descriptor.
- Prevent Board Management Software installation on Compute nodes.

#### **DIS MX 5.15.1:**

- Fixed bug in SISI supporting larger than 4GB segments.
- Fixed multi endian issue with reflective memory test program.

#### **DIS MX 5.15.0:**

- Increase max number of adapters from 4 to 5.
- Fixed problem with slot AER errors not being logged.
- User-space SuperSockets: added a listen thread to accept incoming connections.

**Included software:**

- SuperSockets
  - Ultra-fast, low latency WinSock2 Sockets API (TCP) for PCIe
  - Local and remote socket communication acceleration (accelerated loop-back device support local IP and localhost address).
  - Connectivity to other systems running Linux user space SuperSockets.
- SISCIAPI V2.0
  - Shared memory (DMA, PIO, RDMA, Interrupts)
  - PCIe peer to peer communication (FPGAs, GPUs etc.)
  - Full connectivity to other systems running Linux, RTX or VxWorks.
- TCP/IP driver (IPoPCIe, only included in installers for Windows Vista and newer)
  - Private network to other nodes running Windows.
  - Routing to other network (Connect your PCIe cluster to a 10G Ethernet up-link)
- Transparent Board Management
  - Transparent board management software for transparent card MXH832.

**Dolphin SuperSockets**

The Dolphin SuperSockets library is currently provided for the purpose of supporting embedded applications. The Windows Winsock2 environment contains a rich set of socket functions and options. SuperSockets supports fail-over from single adapter card to Ethernet. SuperSockets version 5.2 and newer adds support for connectivity to the new Linux SuperSockets user space library. We have verified and support the following socket calls:

accept, bind, close, socket, connect, WSASendDisconnect, getpeername, getsockname, getsockopt, ioctlsocket, listen, recv, WSARecv, WSAAsyncSelect, send, WSASend, setsockopt, shutdown, closesocket, ioctlsocket, WSAGetOverlappedResult, WSAIoctl, WSASendDisconnect.

**Supported OS platforms:**

Windows Server 2008 - 32 bit  
Windows Server 2008 - 64 bit  
Windows Server 2008 R2 - 64 bit  
Windows Server 2012 R2 – 64 bit  
Windows Server 2016 – 64 bit  
Windows Server 2019 – 64 bit  
Windows Server 2022 - 64 bit  
Windows 7 - 32 bit  
Windows 7 - 64 bit  
Windows 8 - 32 bit  
Windows 8 - 64 bit  
Windows 8.1 – 64 bit  
Windows 10 – 64 bit  
Windows 11 - 64 bit

**Cluster installation requirements:**

- All nodes connected and properly configured with Ethernet.
  - Non-Ethernet configurations optional, contact Dolphin.

**Installation and management:**

- Windows MSI Installer package.
- Automatic configuration via Interconnect manager.
- Centralized monitoring and diagnostic.
  - Both Linux and Windows supported as GUI platform.

**High availability features:**

- SuperSockets will fail-over to Ethernet if Dolphin Express network is unavailable during application startup.

**Bundled (major) management tools:**

- dis\_diag (diagnostic tool)
- dis\_admin (cluster monitor and manager GUI)
- networkmanager (cluster configuration and maintenance demon)
- dis\_netconfig (configuration editor GUI)

**Bundled (major) demo tools:**

- Scibench2 (SISCI based PIO throughput benchmark)
- scipp (SISCI based ping pong benchmark)
- dma\_bench (SISCI based DMA benchmark)
- interrupt\_bench (SISCI based remote system interrupt benchmark)
- reflective\_bench (SISCI based reflective memory benchmark)
- latency\_bench (TCP latency benchmark, runs on any IP network including SuperSockets)

**Firmware upgrade:**

MXH cards may need a firmware update. After installing the drivers, please run the command:

```
C:\Program Files\Dolphin Express MX\Util\upgrad_eeprom.cmd --upgrade
```

to upgrade the firmware. A reboot is required after upgrading the firmware. Please contact pci-support@dolphinics.com if you have any problems.

## **Dolphin eXpressWare MX release 5.23.0 RTX64**

Description of content: Clustering package for RTX (64-bit only)

Release category: Beta

Release date: January 30th, 2025.

**Current status:** Please contact Dolphin for updated status.

### **Currently supported versions:**

- IntervalZero RTX64 v4.2+
- IntervalZero RTX64 v3.7

### **New in DIS MX 5.23.0**

- General bug fixes and improvements

### **Changes in previous releases:**

#### **DIS MX 5.22.0**

- SISI: General bug fixes and improvements

#### **DIS MX 5.21.x:**

- General bug fixes and improvements

#### **DIS MX 5.20.1:**

- General bug fixes and improvements

#### **DIS MX 5.20.0:**

- General bug fixes and improvements

#### **DIS MX 5.19.2:**

- General bug fixes and improvements.

#### **DIS MX 5.18.0:**

- General bug fixes and improvements.

#### **DIS MX 5.17.0:**

- SISI
  - Added support for send-only multicast configurations.
  - Fixed bug in segment cleanup for shared segments that could cause segment to be unavailable for all users.
  - Increased max multicast segment size to 128 GB. Fixed 32 bit size limitations with internal copy function.
- General bug fixes and improvements.

#### **DIS MX 5.15.2:**

- Increase max number of adapters from 4 to 5.
- Fixed problem with slot AER errors not being logged.

### **Included software:**

- SISI API V2.0.2
  - Shared memory (DMA, PIO, RDMA, Interrupts)
  - PCIe peer to peer communication (FPGAs, GPUs etc.)
  - Full connectivity to other systems running Linux, Windows or VxWorks.

### **Supported OS platforms:**

Windows 10 – 64 bit

Windows 11 – 64 bit

### **Installation and management:**

- Windows MSI Installer package.



**Bundled (major) management tools:**

- dis\_diag (diagnostic tool)

**Bundled (major) demo tools:**

- Scibench2 (SISCI based PIO throughput benchmark)
- scipp (SISCI based ping pong benchmark)
- dma\_bench (SISCI based DMA benchmark)
- interrupt\_bench (SISCI based remote system interrupt benchmark)
- reflective\_bench (SISCI based reflective memory benchmark)

**Firmware upgrade:**

MXH cards may need a firmware update. After installing the drivers, please run the command:

C:\Program Files\Dolphin Express RTSS\Util\upgrade\_eeprom.cmd --upgrade

to upgrade the firmware. A reboot is required after upgrading the firmware. Please contact pci-support@dolphinics.com if you have any problems.

## Dolphin eXpressWare MX release 5.21.0 VxWorks

Description of content: Clustering package for VxWorks

Release category: GA

Release date: May 23<sup>rd</sup>, 2023.

**Current status:** Contact Dolphin.

### New in DIS MX 5.21.0:

- General bug fixes and improvements

### Changes in previous releases

#### DIS MX 5.21.0:

- Fixed internal DMA channels being improperly computed

#### DIS MX 5.20.1:

- Bug fixes and improvements
  - Fixed an issue with cancelling callback threads that caused SISI applications to deadlock

#### DIS MX 5.20.0:

- Bug fixes and improvements
  - Fixed problem with warning about unknown ioctl 0x52 on SR640+

#### DIS MX 5.19.2:

- Added support for adapter MCH663.
- Fixed issue with cache allocator running out of memory
- Added supports for multiple MSI-X vectors
- General bug fixes and improvements.

#### DIS MX 5.18.0:

- Fix potential timer tasks crash during init on some configurations
- SISI
  - Fixed issue with SCISstartDmaTransfer() reporting false errors.

#### DIS MX 5.17.0:

- Added support for VxWorks 7 SR0640
- SISI
  - Added support for send-only multicast configurations.
  - Fixed bug in segment cleanup for shared segments that could cause segment to be unavailable for all users.
  - Increased max multicast segment size to 128 GB. Fixed 32 bit size limitations with internal copy function. Tested up to 4GB.
- General bug fixes and improvements.

#### DIS MX 5.16.0:

- Added support for VxWorks 7 SR0620 with LLVM/Clang compiler
- Fixed VxWorks initialization ordering
- Added CDF parameter EXPRESSWARE\_SLOT\_0\_NODEID to explicitly define the auto-configuration NodeId if the geographical address (slot number) is reported as 0 by the BSP.
- SISI API
  - Added support for using multiple DMA channels.
  - Added support for registering more than one PCIe requester ID per SISI descriptor.

#### DIS MX 5.15.2:

- Increase max number of adapters from 4 to 5.
- Verbosity control for driver output.

- Fixed problem with slot AER errors not being logged.

#### **DIS MX 5.15.1:**

- Fixed bug in SISI supporting larger than 4GB segments.
- Fixed multi endian issue with reflective memory test program.

#### **Included software:**

- SISI API V2.0.2
  - Shared memory (DMA, PIO, RDMA, Interrupts)
  - Reflective memory/multicast
  - PCIe peer to peer communication (FPGAs, GPUs etc.)
  - Full connectivity to other systems running Linux, RTX or VxWorks.

#### **Supported OS platforms:**

VxWorks 6.9 Intel x64 SMP

VxWorks 6.9 Intel x86 SMP

VxWorks 7.0 Intel x64 SMP

VxWorks 7.0 Intel x86 SMP

Note: AMP and User Space communication not supported.

#### **Installation and management:**

- Windows MSI Installer package.
- Linux ZIP file.

#### **High availability features:**

- Nodes can join and leave cluster (node reboot, power cycled) any time without disturbing communication between other nodes.

#### **Bundled (major) management tools:**

- dis\_diag (diagnostic tool)

#### **Bundled (major) demo tools:**

- Scibench2 (SISI based PIO throughput benchmark)
- scipp (SISI based ping pong benchmark)
- dma\_bench (SISI based DMA benchmark)
- interrupt\_bench (SISI based remote system interrupt benchmark)
- reflective\_bench (SISI based reflective memory benchmark)

#### **Firmware upgrade:**

MXH cards may need a firmware update. After installing the drivers, please run the command:

```
upgrade_eeprom.cmd --upgrade
```

to upgrade the firmware. A reboot is required after upgrading the firmware. Please contact pci-support@dolphinics.com if you have any problems.

## **Dolphin eXpressWare MX release 5.23.0 QNX**

Description of content: Clustering package for QNX.

Release category: GA

Target audience: Production systems.

Release date: January 30th, 2025.

**Current status: Please contact Dolphin for details**