
# Dolphin Interconnect Solutions

# eXpressWare

# Software release notes

# for PX products

# 5.23.0

**Date: January 30th, 2025**

# Table of Contents

## 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 5.7.0).

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 has 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. |
| IO Monitoring | Transparent Board Management Software for transparent PCIe adapter cards and cable connections. |

**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), shutdown, socket.

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

**Dolphin technical support and documentation**

We do take software development and product testing seriously, please let us know your experience or any issue by contacting our support team at http://www.dolphinics.com/csp.

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

SISCI API resources can be found at
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**

- PXH810, PXH820, PXH830, PXH840, PXH841, PXH843, PXH844

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

- PXH812, PXH822, PXH832, PXH842, PXH845, PXH846, PXH847, PXH848, PXH860

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

### Supported Dolphin PCI Express switches

- IXS600 (PXH810 only)
- MXS824 (PXH830 and PXH820)
- MXS824 (PXH810 using a x8 iPass to dual x4 SFF-8644 cable)
- MXS924 (PXH830 and PXH820)

### Configurations supported by this release

The software release supports the following configurations.

**General use with PXH810:**

- 2 to 14 nodes using multiple IXS600-HN switches.
- 2 to 12 nodes using MXS824 (using a x8 iPass to dual x4 SFF-8644 cable)
- 2 nodes using a single PCIe x8 cable.

.

**General use with PXH820 and PXH830:**

- 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.
- 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.
- Up to 30 nodes x16 using multiple MXS824 switches.
- Up to 60 nodes x8 using multiple MXS824 switches.

**NOTE:** Please contact Dolphin for more information on large PCIe networks, some limitations may apply.

**General use with PXH840:**

- 2 nodes using 1 or 2 (x8) MPO 24 strand fiber cables.
- 3 nodes using 2 (x8) MPO 24 strand fiber cables.

**Reflective memory use with PXH810:**

- 2 to 14 nodes using one IXS600-HN switch.
- 2 nodes using a single PCIe x8 iPass cable.

**Reflective memory use with PXH820 and PXH830:**

- 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.

- Up to 30 nodes x16 using multiple MXS824 switches.
- Up to 60 nodes x8 using multiple MXS824 switches.

**Support for redundant fabrics**

- Up 5 Cards pr host connected directly or through separate switches are supported at the SISCI level. SuperSockets supports fail-over from single adapter card to Ethernet.

**PCI Express Hot Add with PXH820, PXH830 and PXH840 (Linux Only)**

- PXH820 or PXH830 connected to SFF8644 target devices.
- PXH820 or PXH830 connected to MXS824 in transparent mode, with attached SFF8644 target devices.
- NTB PXH84x connected to a FireFly target device.

**NOTE:** Please contact Dolphin if you are planning to connect directly to non-Dolphin target card or device.

**Standard firmware revisions**

| Adapter / DIS release | PXH82x | | PXH830 | | PXH832 | | PXH84x | |
|---|---|---|---|---|---|---|---|---|
| | MCU | EEPROM | MCU | EEPROM | MCU | EEPROM | MCU | EEPROM |
| 5.23.0 | 8.12 | 2 | 8.16 | 11 | 8.16 | 13 | 8.11 | 8 |
| 5.22.0 | 8.12 | 2 | 8.16 | 11 | 8.16 | 14 | 8.11 | 8 |
| 5.21.x | 8.12 | 2 | 8.14 | 11 | 8.14 | 13 | 8.11 | 8 |
| 5.20.0 / 5.20.1 | 8.12 | 2 | 8.14 | 11 | 8.14 | 13 | 8.11 | 8 |
| 5.19.2 | 8.12 | 2 | 8.12 | 11 | 8.12 | 11 | 8.11 | 8 |
| 5.18.1 | 8.12 | 2 | 8.12 | 11 | 8.12 | 11 | 8.11 | 8 |
| 5.17 | 8.12 | 2 | 8.12 | 11 | 8.12 | 11 | 8.11 | 8 |
| 5.16 | 8.11 | 2 | 8.11 | 11 | 8.11 | 11 | 8.11 | 8 |
| 5.15.2 | 8.11 | 2 | 8.11 | 11 | 8.11 | 11 | 8.11 | 8 |
| 5.15.1 | 8.11 | 2 | 8.10 | 11 | 8.10 | 11 | 8.10 | 8 |
| 5.14.0 | 8.10 | 2 | 8.10 | 11 | 8.10 | 11 | 8.6 | 8 |
| 5.13.1 | 8.9 | NA | 8.4 | 11 | 8.4 | 11 | 8.6 | 8 |

**Support for OEM hardware**

This version of Dolphin eXpressWare has general support for most Broadcom Gen3 NTB enabled chipset in various configurations and topologies. Please contact Dolphin for details and licensing information if you would like to run eXpressWare on y our 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.

Please note that Dolphin eXpressWare also supports IDT and Microchip PCIe chipset. Please contact Dolphin for more information.

**General eXpressWare defaults and settings for PX**

eXpressWare is tuned for general use and operations. In some cases, you need to tune some parameters to optimize resources for our use case.

- SISCI Multicast
  - Default setting 4 groups, 2 Megabyte per segment.
  - Max setting 64 groups, 128 Gigabytes 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 3rd, 4th and 5th gen Intel Xeon Scalable processors.**

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

Certain configurations of Intel Xeon Scalable 3rd, 4th and 5th 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 3rd, 4th and 5th gen processors in the following ways:

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

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

## Dolphin eXpressWare PX 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/px

**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 PX 5.23.0**

- Added support for PXH822-L1
- SISCI:
  - Added interrupt support for SISCI SmartIO
  - Enabled possibility for non-zero offset when calling SCIMapRemoteSegment for multicast segments
  - Bug fix: Fix for SCIStartDMATransferMem deadlock
- SmartIO:
  - Added interrupt support for SISCI 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: Lane state is now only printed for adapter with SFF-8644 connectors
- Improved dis_nodemgr:
  - Unresolvable hosts are skipped and no longer stop cluster configuration.
- 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
- Improved dis_services:
  - Bug fix: dis_services now properly handles services that end up in a failure state
- Improved dis_license_tool:
  - Bug fix: Fixed issue where applying a new license key failed
- 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:**

**DIS PX 5.22.0**

- Added support for Intel Xeon Scalable 5$^{th}$ gen (“Emerald Rapids”).
- Improved support for Intel Xeon Scalable 3$^{rd}$, 4$^{th}$ and 5$^{th}$ 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.
- Improved dis_diag.
  - Added support for json-output.
  - Added support for showing root port link capabilities.
  - Improved link error detection.
  - Now shows temperature information for Firefly.
  - Added lane state information to output.
- 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.

**DIS PX 5.21.4:**

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

**DIS PX 5.21.3:**

- No changes for PX.

**DIS PX 5.21.2:**

- For Linux users, this release is equivalent to the 5.21.1 release below.

**DIS PX 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 PX 5.21.0:**

- Added support for Linux kernel 5.19
- Added support for Linux kernel 6.0
- MXS924:
  - Eeprom version updated to VER5
  - Updated throttling values for DSWITCH topology
- Added preliminary support for Intel Ice Lake and Sapphire Rapids Xeon CPUs in DIRECT topology. Limited support for MXS824/MXS924 switches.
- 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 warnings when running `dis_firmwaretool` utility

**DIS PX 5.20.1:**

- Initial support for CentOS 9
- Added 'make' package as required dependency
- Added support for Ubuntu 22.04
- 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
- 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 dis_config reordering issue during configuration
  - Fixed an issue with cancelling callback threads that caused SISCI applications to deadlock

**DIS PX 5.20.0:**

- Added support for MXS924 switch topologies:
  - NTB
  - Transparent
  - D-SWITCH
- Added support for Linux 5.16
- Added support for Linux 5.18
- Added support for dis_firmwaretool on the MXS924 switch
- Added support for using MXS824 switches with PVH240
- Added preliminary support for Intel Ice Lake Xeon CPUs
- Added option "-userid" for example programs reflective_write and reflective_bench to allow multiple concurrent instances
- Bug fixes and improvements
  - 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
  - Fix OS deadlock due to mtbl_index_bitmap overflow for PLX_CAPELLA_1, when deleting a mapping table
  - Fixed regression with Asymmetric topology on PCH482
  - Fixed session problem after running multicast programs with MXS CASCADE topology
  - Fixed problem with unsupported asymmetric adapters confusing the transparent driver
  - 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
  - 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

**DIS PX 5.19.2:**

- Fixed unbinding of plx_dma driver from DMA functions of Dolphin adapters
- Added support for PXH832-H1 where Suppress Surprise down bit is enabled
- Added upgrade support for PXH860
- Added support for adapters with IPASS connectors connected to adapters with Mini SAS HD connectors. I.e., PXH810 can be connected to PXH830 or MXS824/MXS924 with special cable.
- Added support for Linux kernels up to 5.14
- Fixed handling of completion timeout on Root Port for transparent adapters.
- Fixed regression with Asymmetric topology introduced in 5.19.0
- Fixed driver crash in DSWITCH mode
- Fixed possible crash when aborting sysdma transfers on some systems
- Use more conservative PCIe MPS (max packet size) handling.

By default (ntb_set_mps=1) to prevent automatic increase.

- General bug fixes and improvements.
- Installer support EL7 for Fedora 22, 23 and 24
- Added link speed (--link-speed) option to the Linux installer
- Fixed crash when IOMMU was disabled on some systems
- Improved deallocation of segments for smart-IO
- Added support for Supersockets on Ubuntu 20.04 with hwe-kernel
- Improved interrupt handling for global DMA
- Resetting the link with causes LUTs to be marked as zombies
- Fixed DMA driver hang on some systems
- Applied copy function patch for the unaligned crash issue on some ARM platforms.
- Added support for Ubuntu 20.04.2 LTS

**DIS PX 5.18.1:**

- Improved detection and reporting of old firmware.
- General bug fixes and improvements.
- Added 2x8 support for MXH830 adapters
- Fixed DMA performance optimization on ARM64 platforms.
- Added initial support for monitoring transparent downstream links in D-Switch mode.
- Fixed DMA hang on some systems
- General DMA improvements
- Fixed regression in session handling with MXS824 switch
- Fixed hardware descriptor index bug when using MSI interrupts for Global DMA
- Fixed problem with Global DMA engine hanging due to wrong SSID register content on some systems.

**DIS PX 5.17.0:**

- SISCI
  - 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 issue upgrading PXH card firmware on some platforms.
- Fixed Ubuntu 18.04 installation issue.
- Do not report Unsupported Requests (UR) to the IRM.
- Add BDF information in dis_diag and dis_tool.
- General bug fixes and improvements.

**DIS PX 5.16.0:**

- Added support for system DMA with IOMMU ON (Off already supported)
- Added full support for all flavors of PXH84x with reduced number of FireFly ( PXH843 – PXH848).
- Added support for Linux kernel >= 5.5.3. Tested on Linux kernel 5.5.8
- SISCI API
    - Added support for using multiple DMA channels.
    - Added support for registering more than one PCIe requester ID per SISCI 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 PX 5.15.2:**

- Improved detection and reporting of old firmware
- SmartIO – fixed potential assert in device scanning process

**DIS PX 5.15.1:**

- Fixed bug in SISCI supporting larger than 4GB segments.
- Fixed multi endian issue with reflective memory test program
- PX Firmware 5.11 – fixes issue setting prefetch space size introduced in Firmware 5.10
- SmartIO 2.5.1
    - Improved scanning and error resilience, fixed dual adapter problem.

**DIS PX 5.15.0:**

- Increase max number of adapters from 4 to 5.
- SmartIO 2.5
    - SmartIO SISCI 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 MXS824 switch not being detected by the driver running on big endian host machines.
- 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).
- 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 systems running Windows SuperSockets.
- SISCI API 2.0.2
  - Shared memory (DMA, PIO, RDMA, Interrupts).
  - Reflective memory/multicast.
  - SmartIO Extension 1.0
  - PCIe peer to peer communication (FPGAs, GPUs, NVMe etc.).
  - Full connectivity to other systems running Windows, RTX or VxWorks.
- SmartIO 2.5
  - Device Lending
  - SmartIO Hot-Add
  - SISCI 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.
- Transparent Board Management
  - Transparent board management software for transparent cards PXH812, PXH832 and PXH84x.

**Supported architectures:**

The software supports the following architectures and platforms:

- Intel x86 / x64
- ARM32, Jetson Tegra TK1
- ARM64, Jetson Tegra TX1, TX2. Xavier
- ARM64v8 Applied  Micro X-Gene

**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
- 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

- OpenSUSE Leap 15.1 x86_64
- Debian 7 x86_64
- Debian 8 x86_64
- Debian 9 x86_64
- Debian 10 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:**
- Nodes can join and leave cluster (node reboot, power cycled) any time without disturbing communication between other nodes.
- Power cycling the switch will cause SuperSockets to fail over to Ethernet while cluster is automatically reconfiguring.
- 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 PX release 5.23.0 Windows

Description of content: Clustering package for Windows.
Release category: GA
Release date: January 30th, 2025

**Current status: Available for download from**
https://www.dolphinics.com/px

**New in DIS PX 5.23.0:**

- SISCI:
  - Enabled possibility for non-zero offset when calling SCIMapRemoteSegment for multicast segments
  - Bug fix: Fix for SCIStartDMATransferMem deadlock
- Bug fix: Driver configuration now correctly exits if IOMMU initialization fails
- Improved dis_diag:
  - Bug fix: No longer prints invalid temperature
- Improved dis_nodemgr:
  - Unresolvable hosts are skipped and no longer stop cluster configuration
- Bug fix: Update BAR computation on transparent devices
- Bug fix: Set the DMA function's MPS equal to the one used by the root port

**Changes in previous releases:**

**DIS PX 5.22.0**

- Added support for Intel Xeon Scalable 5$^{th}$ gen ("Emerald Rapids").
- Added support for Windows 11, Server 2022
- Improved support for daisy chain configurations.
- 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.

**DIS PX 5.21.4:**

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

**DIS PX 5.21.3:**

- No changes for PX.

**DIS PX 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 PX 5.21.0:**

- Removed dis_coinstaller from the driver package

- General bug fixes and improvements

**DIS PX 5.20.1:**
- Bug fixes and improvements
  - Fixed initial adapterno value for IRM transparent inf files

**DIS PX 5.20.0:**
- Bug fixes and improvements
  - Fixed crash on PXH830 where NTB_B0 is loaded while NTB_B1 is unloaded
  - Fixed upgrade_eeprom.ps1 for PXH810

**DIS PX 5.19.2:**
- Improved detection of DMA hardware PEX ids to prevent conflicts with other PEX switches.
- Improved memory handling allocating large segments
- General improvements
- Improved Windows installer

**DIS PX 5.18.1:**
- Fixed crash on transparent adapter installation.
- Fixed completion timeout on Root Port for transparent adapters.
- Fixed problem with Global DMA engine hanging due to wrong SSID register content on some systems.
- Fixed hardware descriptor index bug when using MSI interrupts for Global DMA.
- Improved detection and reporting of old firmware.
- Improved driver handling of transparent adapters

**DIS PX 5.16.0:**
- SISCI API
  - Added support for using multiple DMA channels.
  - Added support for registering more than one PCIe requester ID per SISCI descriptor.
- Prevent Board Management Software installation on Compute nodes.

**DIS PX 5.15.2:**
- Added support for SISCI Hot Add support for transparent devices.
- Improved detection and reporting of old firmware.

**DIS PX 5.15.1:**
- Fixed bug in SISCI supporting larger than 4GB segments.
- Fixed multi endian issue with reflective memory test program.

**DIS PX 5.15.0:**
- Increase max number of adapters from 4 to 5.
- User-space SuperSockets: added a listen thread to accept incoming connections.

**DIS PX 5.14.0:**
- Support for PXH820/822 Revision B (XMC)
- Increased maximum multicast segment size to 512 MB
- Fixed problem with dis_diag reporting erroneous board temperature for PXH8xx
- Improved allocation of large memory areas
- Fixed IRM failure to load on certain runtime contexts
- Fixed detection of RootPort's BDF on some systems

- Fixed deadlock or crash due to improper computation of CompletionTimeout offset inside PCIe AER capability
- General bug fixes and improvements.

**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.
- SISCI API V2.0
  - 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.
- 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 cards PXH812, PXH832 and PXH842.

**Supported OS platforms:**

Windows Vista - 32 bit
Windows Vista - 64 bit
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:**

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

PXH cards may need a firmware update to support the 5.x.x series of software. After installing the drivers, please run the command:

Start PowerShell with administrative capabilities (Press Windows+X, select Windows PowerShell (Admin))

```
PS > cd "${env:ProgramFiles}\Dolphin Express PX\Util"
PS > Set-ExecutionPolicy AllSigned -Scope Process
PS > .\upgrade_eeprom.ps1 --upgrade
```

Please carefully review the output from the upgrade utility.

A reboot is required after upgrading the firmware. Please file a support request at http://www.dolphinics.com/csp if you have any problems.

## Dolphin eXpressWare PX release 5.23.0 RTX64

Description of content: Clustering package for RTX (64-bit only)
Release category: GA
Release date: January 30th, 2025.

**Current status:** Available. Contact Dolphin.

**Currently supported versions:**
- IntervalZero RTX64 v4.2+
- IntervalZero RTX64 v3.7

**New in DIS PX 5.23.0**
- General bug fixes and improvements

**Changes in previous releases:**

**New in DIS PX 5.22.0**
- SISCI: General bug fixes and improvements
- Improved PX adapter function detection
- Added support for driver configuration when the monolith must be available on a limited account

**DIS PX 5.21.x:**
- General bug fixes and improvements

**DIS PX 5.20.1:**
- General bug fixes and improvements

**DIS PX 5.20.0:**
- General bug fixes and improvements

**DIS PX 5.19.2:**
- General bug fixes and improvements.

**DIS PX 5.18.1:**
- General bug fixes and improvements.

**DIS PX 5.17.0:**
- SISCI
  - 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 PX 5.15.2:**
- Improved detection and reporting of old firmware.

**DIS PX 5.15.0:**
- Increase max number of adapters from 4 to 5.

**Included software:**
- SISCI 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, Windows or VxWorks.

**Supported OS platforms:**
Windows 10 – 64 bit
Windows 11 – 64 bit

**Installation and management:**
- Windows MSI Installer package.

**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 (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:**
PXH cards may need a firmware update to support the 5.x.x series of software. 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 file a support request at http://www.dolphinics.com/csp if you have any problems.

## Dolphin eXpressWare PX release 5.21.0 VxWorks

Description of content: Clustering package for VxWorks
Release category: GA
Release date: May 23rd, 2023.

**Current status:** Contact Dolphin.

**New in DIS PX 5.21.0:**

- General bug fixes and improvements

**Changes in previous releases**

**DIS PX 5.20.1:**

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

**DIS PX 5.20.0:**

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

**DIS PX 5.19.2:**

- Added support for MCH663 D-Switch on Linux and VxWorks 7
- Fixed regression with Asymmetric topology on PCH482
- Added support for MCH663 D-Switch
- General bug fixes and improvements.
- Extended adapter installer support
- Fixed issue with cache allocator running out of memory
- Added support for multiple MSI-X vectors for VxWorks 7

**DIS PX 5.18.1:**

- Fixed timer tasks crash during init on some configurations
- Improved driver handling of transparent adapters

**DIS PX 5.17.0:**

- Added support for VxWorks 7 SR0640
- SISCI
  - 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 PX 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.
- SISCI API
  - Added support for using multiple DMA channels.
  - Added support for registering more than one PCIe requester ID per SISCI descriptor.

**DIS PX 5.15.2:**

- Improved detection and reporting of old firmware.

**DIS PX 5.15.1:**

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

**DIS PX 5.15.0:**

- Increase max number of adapters from 4 to 5.
- Verbosity control for driver output.
- Fixed problem with MXS824 switch not being detected by the driver running on big endian host machines.

**Included software:**

- SISCI 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 6.8 PPC32 SMP
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 (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:**
PXH cards may need a firmware update to support the 5.x.x series of software. 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 file a support request at http://www.dolphinics.com/csp if you have any problems.