



# **eXpressWare release for PX products**

**Dolphin Interconnect Solutions**

**January 28, 2026**

# Contents

<b>1</b>	<b>Definitions</b>	<b>1</b>
	Product PCIe Chipset . . . . .	1
	Release categories . . . . .	1
	Release Version Scheme . . . . .	2
	Types of Releases . . . . .	2
<b>2</b>	<b>Dolphin eXpressWare Overview</b>	<b>3</b>
	Components . . . . .	3
	SuperSockets . . . . .	3
	Documentation & Support . . . . .	4
<b>3</b>	<b>Supported Hardware &amp; Configurations</b>	<b>5</b>
	Dolphin PCIe Adapter Cards supported by the NTB software . . . . .	5
	Dolphin PCIe Adapter Cards supported by the Dolphin Board Management Software . . . . .	5
	Third-Party Devices Supported by Dolphin Software . . . . .	6
	Dolphin PCIe Switches supported . . . . .	6
	Supported Configurations . . . . .	6
	Standard Firmware Revisions . . . . .	8
	Support for OEM hardware . . . . .	8
	General eXpressWare defaults and settings for PX . . . . .	8
	Known Limitations . . . . .	8
<b>4</b>	<b>Change log for Linux Software</b>	<b>10</b>
	New in DIS Release 5.25.1 . . . . .	10
	Changes in previous releases . . . . .	11
	DIS Release 5.25.0 . . . . .	11
	DIS Release 5.24.0 . . . . .	13
	DIS Release 5.23.0 . . . . .	15
	DIS Release 5.22.0 . . . . .	16
	DIS Release 5.21.4 . . . . .	17
	DIS Release 5.21.3 . . . . .	17
	DIS Release 5.21.2 . . . . .	17
	DIS Release 5.21.1 . . . . .	17
	DIS Release 5.21.0 . . . . .	17
	DIS Release 5.20.1 . . . . .	17
	DIS Release 5.20.0 . . . . .	18

DIS Release 5.19.2 . . . . .	19
DIS Release 5.18.1 . . . . .	20
DIS Release 5.17.0 . . . . .	20
DIS Release 5.16.0 . . . . .	20
DIS Release 5.15.2 . . . . .	21
DIS Release 5.15.1 . . . . .	21
DIS Release 5.15.0 . . . . .	21
Included software . . . . .	22
Supported architectures . . . . .	22
Tested and supported OS platforms . . . . .	23
Cluster installation requirements . . . . .	23
Installation and management . . . . .	24
High availability features . . . . .	24
Bundled (major) management tools . . . . .	24
Bundled (major) demo tools . . . . .	24
Firmware upgrade . . . . .	24
<b>5 Change log for Windows Software . . . . .</b>	<b>25</b>
New in DIS Release 5.25.1 . . . . .	25
Changes in previous releases . . . . .	25
DIS Release 5.25.0 . . . . .	25
DIS Release 5.24.0 . . . . .	27
DIS Release 5.23.0 . . . . .	27
DIS Release 5.22.0 . . . . .	28
DIS Release 5.21.4 . . . . .	28
DIS Release 5.21.3 . . . . .	28
DIS Release 5.21.2 . . . . .	28
DIS Release 5.21.0 . . . . .	28
DIS Release 5.20.1 . . . . .	29
DIS Release 5.20.0 . . . . .	29
DIS Release 5.19.2 . . . . .	29
DIS Release 5.18.1 . . . . .	29
DIS Release 5.16.0 . . . . .	29
DIS Release 5.15.2 . . . . .	29
DIS Release 5.15.1 . . . . .	30
DIS Release 5.15.0 . . . . .	30
DIS Release 5.14.0 . . . . .	30
Included software . . . . .	30
Supported OS platforms . . . . .	31
Cluster installation requirements . . . . .	31
Installation and management . . . . .	31
High availability features . . . . .	31
Bundled (major) management tools . . . . .	32
Bundled (major) demo tools . . . . .	32
Firmware upgrade . . . . .	32
<b>6 Change log for RTX Software . . . . .</b>	<b>33</b>
New in DIS Release 5.25.1 . . . . .	33
Changes in previous releases . . . . .	33
DIS Release 5.25.0 . . . . .	33
DIS Release 5.24.0 . . . . .	35
DIS Release 5.23.0 . . . . .	35
DIS Release 5.22.0 . . . . .	35
DIS Release 5.20.1 . . . . .	35

DIS Release 5.20.0 . . . . .	36
DIS Release 5.19.2 . . . . .	36
DIS Release 5.18.1 . . . . .	36
DIS Release 5.17.0 . . . . .	36
DIS Release 5.15.2 . . . . .	36
DIS Release 5.15.0 . . . . .	36
Included software . . . . .	36
Supported OS platforms . . . . .	37
Installation and management . . . . .	37
High availability features . . . . .	37
Bundled (major) management tools . . . . .	37
Bundled (major) demo tools . . . . .	37
Firmware upgrade . . . . .	37
<b>7 Change log for VxWorks Software</b>	<b>38</b>

## Definitions

### Product PCIe Chipset

PX products are based on the Broadcom PCIe 3.0 NTB enabled chipsets in various configurations and topologies.

#### **Note**

Please note that Dolphin eXpressWare also supports Microchip PCIe chipsets. Please see <https://www.dolphinics.com> for more information.

### Release categories

The following definitions are used for the eXpressWare software.

#### **Prototype**

Only parts of the full functionality are implemented. Minimal error handling. Intended to demonstrate limited functionality to selected customers. Not for general availability.

#### **Alpha**

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

#### **Beta**

Complete with error handling and performance optimization. Nearly final product. Can be made generally available.

#### **GA**

General Availability. Fully tested and ready for shipment 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 5.23.1).

- **MAJOR\_VER:**

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

- **MINOR\_VER:**

Changes to the software that significantly improves functionality. Enhancements, OS support, non-breaking improvements.

- **BUGFIX\_VER:**

Minor fixes of previously released software. No functional changes.

## Types of Releases

Dolphin uses two types of releases to satisfy rapid development, long-term test requirements and concurrent external testing and stability.

### Release Candidate

A pre-release version that is still undergoing final testing and qualification. It is made available to customers and partners to provide early access to updated drivers and features, without requiring them to wait for the official release. The primary distinction between a release candidate and a final release is the level of testing completed. If critical issues are discovered during release candidate testing, the release may be withdrawn. In such cases, the version string and date must be updated before a new release candidate is issued, in accordance with standard release engineering procedures.

### Release

Has completed all test procedures. Typically a renamed release candidate that passes all tests, and with no new critical bugs found.

## Dolphin eXpressWare Overview

The eXpressWare software suite includes the following components. Not all combinations of software, hardware and OS are supported. Refer to release note section for compatibility.

### Components

- **DIS-PX:** Low-level PX device driver
- **IRM:** Interconnect Resource Manager
- **SISCI:** Shared-memory cluster interconnect infrastructure
- **SuperSockets:** Socket accelerator for PCIe
- **IPoPCIe:** TCP/IP over PCIe (Windows NDIS or Linux DISip)
- **SmartIO:** IO sharing and control. Supports Device Lending and Hot-Add (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.

#### User-space Library (Windows/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.

##### Supports:

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

##### Windows-specific:

closesocket, WSASendDisconnect, ioctlsocket, WSAAsyncSelect, WSAGetOverlappedResult, WSAIoctl, WSAREcv, WSASend

##### Linux-specific:

close, ioctl

---

## Kernel-space Library (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. 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 Info: <https://www.dolphinics.com/software>

## Documentation & Support

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

- **Additional info and manuals:**

- <https://www.dolphinics.com/px>
- <https://docs.dolphinics.com>

- **SISCI API resources:**

- <https://docs.dolphinics.com/sisci/guides/index.html>
- <https://www.dolphinics.com/products/embedded-sisci-developers-kit.html>

- **SmartIO resources:**

- <https://docs.dolphinics.com/smartio/introduction.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 & Configurations

### Dolphin PCIe Adapter Cards supported by the NTB software

- **PXH810:** PCIe 3.0 NTB Host Adapter, x8, iPass
- **PXH820:** PCIe 3.0 XMC NTB Adapter, x16, SFF-8644
- **PXH830:** PCIe 3.0 NTB Host Adapter, x16, SFF-8644
- **PXH840:** PCIe 3.0 NTB Host Adapter, 2 x8 FireFly modules, 2 x8, x16
- PXH841
- PXH843
- PXH844

### Dolphin PCIe Adapter Cards supported by the Dolphin Board Management Software

- **PXH812:** PCIe 3.0 Transparent Host / Target Adapter, x16, iPass
- **PXH822:** PCIe 3.0 Transparent XMC Adapter, x16, SFF-8644
- **PXH832:** PCIe 3.0 Transparent Host / Target adapter, x16, SFF-8644
- **PXH842:** PCIe 3.0 Transparent Host / Target adapter, 2 x8 FireFly modules, 2 x8, x16
- PXH845
- PXH846
- PXH847
- PXH848
- PXH860

---

## Third-Party Devices Supported by Dolphin Software

- PCH258
- PCH482
- PCH420
- PCH133
- ICH131
- PVH801
- PVH802
- PVH210
- PVH230
- PVH240
- PVH250
- PVH280

### Note

If you are not using a listed adapter card, please use another software download.

## Dolphin PCIe Switches supported

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

## Supported Configurations

### Note

See the *Known Limitations* section about support for Intel Xeon Scalable 3rd, 4th and 5th gen processors if you plan to use any of these in a switch configuration.

#### 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

**Support for redundant fabrics:**

- Up to 5 Cards per host connected directly or through separate switches are supported at the SIBICI 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

Table 1: Standard Firmware revisions for PXH8xx

Release	PXH82x		PXH830		PXH832		PXH84x	
	MCU	EEPROM	MCU	EEPROM	MCU	EEPROM	MCU	EEPROM
5.25.1	8.12	2	8.16	11	8.16	13	8.11	8
5.25.0	8.12	2	8.16	11	8.16	13	8.11	8
5.24.0	8.12	2	8.16	11	8.16	13	8.11	8
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
5.20.1	8.12	2	8.14	11	8.14	13	8.11	8
5.20.0	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.x	8.12	2	8.12	11	8.12	11	8.11	8
5.16.x	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	8.1	8
5.14.0	8.10	2	8.10	11	8.10	11	8.6	8
5.13.1	8.9	N/A	8.4	11	8.4	11	8.6	8

## Support for OEM hardware

This version of Dolphin eXpressWare has general support for most Broadcom PCIe 3.0 NTB enabled chipsets in various configurations and topologies. Please contact Dolphin for details and 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 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 Gigabyte per segment

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

## Known Limitations

### Support for 3rd, 4th and 5th gen Intel Xeon Scalable processors

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

**SmartIO is not yet supported on these processors.**

---

Certain configurations of Intel Xeon Scalable 3rd, 4th and 5th gen processors require special configurations. Please contact Dolphin support at <https://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 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

### **SmartIO**

Borrowing multiple NVIDIA GPUs using a PX adapter connected to a **MXS824** or **MXS924** switch is not supported.

## Change log for Linux Software

**Description of content:** Clustering package for Linux.

**Release category:** GA.

**Target audience:** Production systems.

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

### Note

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

## New in DIS Release 5.25.1

### General

- (#8274) *dis\_diag* now correctly prints the root port's BDF
- (#8329) Fix issue where the build date for certain programs on RPM-based systems would show *Jan 28 2008*
- (#8342) Fixed CSR value printouts

### Installer

- (#8266) Fix warning in installer output: 'grep: warning: stray \ before white space'

### Platform support

- (#8258) Add support for Debian 13
- (#8285) Add support for Redhawk 9.6

### SISCI

- (#8207) SCICreateMapSequence returns error when used on map of local segment
- (#8280) Improved performance for *memcpy\_nt\_avx2*
- (#8341) Fix arithmetic overflow inside SISCI testlib, on systems where CPU frequency is 4GHz+

### SmartIO

- 
- (#8172) Fix issue where *smartio\_tool unavailable* does not make devices unavailable
  - (#8249) Fixed issue where CUDA programs would fail with the assertion *Assertion failed: 0 == (physAddr & (RM\_PAGE\_SIZE\_HUGE - 1))*
  - (#8301) Add information to *smartio\_tool list* when session to borrower is disconnected
  - (#8302) Add information to *smartio\_tool list* when session to lender is disconnected
  - (#8302) Fix some bugs related to lender rebooting while borrowed device is in use
  - (#8303) Report DMA window size during borrow in a more consistent way
  - (#8325) *smartio\_tool connect* now persists after a borrower is rebooted

## Changes in previous releases

### DIS Release 5.25.0

#### Benchmarks

- (#7882) Adds check for div by zero when *walltime* is used for *dma\_bench* or *scibench*
- (#7997) Creates a benchmark analogous to *reflective\_bench* but which uses DMA for sending
- (#8063) *dma\_bench* regression fixed so *-veclen* parameter can be used again
- (#8075) Added *-vl* argument to *reflective\_dma\_bench* which can be used to increase DMA transfer vector length
- (#8094) The *sisci\_benchmarks.sh* script now also runs the vectorized version of *dma\_bench*

#### Demos

- (#7851) Added performance demo, video transfer demo, and rfm wrapper to DIS example programs

#### Dis\_diag

- (#7817) New json output when *-json* option is used
- (#7893) *dis\_diag* output reorganized to print info per adapter
- (#8123) Make the report of the upstream link more generic in *dis\_diag*

#### General

- (#6980) Fixed doorbell performance spikes by removing an unneeded loopback doorbell
- (#7317) The configuration files of DIS hosts no longer supports 'STRIPE' and 'REDUNDANT' socket adapters. A socket adapter specified with this mode will be ignored
- (#7359) Fixed issue where *\*\_setup* scripts would not exit with an error code on failure
- (#7737) Added support for VXP adapter board VPX6-1961
- (#7845) Improved stability with multi-socket Icelake configurations
- (#7887) Require that endpoint devices in same group as dolphin device is not attached by other drivers
- (#7891) Fixed problem with driver crash during *dis\_config* for adapters with invalid hardware configurations
- (#7940) Added support to detect and report the upstream hwport
- (#7968) Fixed issue where *dis\_services stop* would fail if there is no adapter installed in the system
- (#7977) Fixed a bug where a node would end up in a livelock with no session, spamming *getAliveCnt - mbxMsgQ\_localHeader == NULL* to the kernel buffer

- 
- (#8026) Fix issue where old versions of the *dis\_intel\_dma* kernel module would be loaded, even if the driver is not compiled with *-enable-intel-dma*
  - (#8027) Fixed a bug where multicast would fail when IOMMU is enabled and the system contains a transparent adapter
  - (#8031) Vastly improved PIO performance for the Zen3 architecture
  - (#8042) Added DUAL NTB Switch support for MXS924
  - (#8052) Added flag *-skip-distro-check* to eXpressWare installers
  - (#8156) Improved the detection of adapter types in *dis\_firmwaretool*
  - (#8173) Fixed bug in MXS824 NTB DUAL configs that prevented version 10 of the configs from working
  - (#8203) *dis\_intel\_dma* is now enabled by default if the driver is built with the flag *-enable-intel-dma*
  - (#8237) Improved support for Granite Rapids Xeon Scalable CPUs

### Platform support

- (#7790) Added support for Linux kernel 6.16
- (#8001) Added proper support for Linux kernel 6.15
- (#8182) Added support for Debian 12
- (#8182) Added support for Linux kernel 6.17
- (#8182) Added support for Linux kernel 6.18

### SISCI

- (#7336) Multi-threading is now required to compile eXpressWare
- (#7501) Add “gpudirect” SISCI example using CUDA
- (#7879) Fixed off by one error in *reflective\_bench*
- (#8001) Fix issue where SISCI would fail to load kernel modules for Linux kernel  $\geq 6.15$
- (#8044) *\_REENTRANT* no longer has to be defined for SISCI-features that require multi-threading
- (#8230) Fixed a bug in *SCICreateDeviceSegment* where it would fail if called with either of the flags *SCI\_FLAG\_LOCAL\_ONLY* or *SCI\_FLAG\_DEVICE\_SIDE\_ONLY*

### SmartIO

- (#7522) Improved default DMA window size
- (#7547) Add *disable-p2p* command to *smartio\_tool*. Mappings set up by *enable-p2p* are now properly cleaned up
- (#7597) Add experimental support for virtualization. See SmartIO documentation for more details
- (#7616) Fix issue where borrowed GPUs would fail to initialize CUDA with nvidia-drivers  $\geq 575$
- (#7823) Fixed a small memory leak on initialization
- (#7877) Allow local borrow of other devices in an IOMMU group when a device in that group is already locally borrowed
- (#7911) Add virtualization support for Linux kernel 5.15-5.19. See the SmartIO documentation for details
- (#7941) Identified issue with AMD EPYC lenders. See the SmartIO documentation for details
- (#7953) Fix regression where *smartio\_tool enable-p2p* would fail for two remote devices located in the same lender

- 
- (#7984) Added support for Device Lending on Linux kernels  $\geq$  v6.12
  - (#7998) Fixed issue borrowing SR-IOV virtual functions
  - (#8113) Add support for a proposed Linux patch regarding AMD EPYC lenders
  - (#8175) Resolved issue with NT-Transparent causing repeated borrow/return to fail
  - (#8186) Fixed an issue where two lenders borrowing each others' devices would fail with IOMMU disabled
  - (#8226) Fixed issue with automatic DMA windows size calculation with NT-Transparent
  - (#8230) Fixed a bug in *SCICreateDeviceSegment* where it would fail if called with either of the flags *SCI\_FLAG\_LOCAL\_ONLY* or *SCI\_FLAG\_DEVICE\_SIDE\_ONLY*

### Supersockets

- (#7350) Systemd unit file for *dis\_supersocket.service* changed to make it start after *dis\_nodemgr.service*
- (#7629) Sockperf: 1 second walltime is standard measurement
- (#7629) Sockperf: Improved explanation of measurement options
- (#7878) Fixed issue where SuperSockets failed to build on RHEL/centos 10
- (#7980) Add a *-size* flag to *network\_benchmarks.sh*
- (#8005) Added UDP mode to Network Benchmarks
- (#8005) Improved usability for *network\_benchmarks.sh*
- (#8009) Fixed a memory leak related to the usage of DMA channels
- (#8071) Improvements to *network\_benchmarks.sh*
- (#8114) Solved issue with *dis\_ssocks\_diag* that could cause segfaults on newer kernels

## DIS Release 5.24.0

### Benchmarks

- (#7644) Added *-wt* (walltime) option for *dma\_bench* and *scibench2*
- (#7801) Added warmup loops to *scibench2*, *dma\_bench* and *dma\_user\_bench*
- (#7801) Fixed potential sending of too large sizes for *scibench* if offset is used
- (#7801) Made usage text more consistent for *scibench2* and *dma\_bench*

### General

- (#7446) Fixed reset option for MXS924, Port 5 in DSWITCH mode
- (#7516) Fix DMA transfer failure due to uninitialized variable
- (#7518) Added option to configure processor I/O requester IDs manually in the configuration files. Particularly useful for heterogenous Intel Xeon Scalable clusters
- (#7526) Fixed bug in *example\_to\_json* script
- (#7536) Check that any linux NTB drivers (e.g. *switchtec* for MX) are stopped before starting driver
- (#7539) Minor fixes in installer scripts
- (#7569) Fix installer counting the wrong number of adapters
- (#7581) Fixed a bug where *dis\_services* would yield undefined behavior in some environments
- (#7593) Fixed issue with make commands not always using the correct compiler if *-host* supplied

- 
- (#7726) Removed use of deprecated kernel function *flush\_scheduled\_work()*. Fixes kernel warning when stopping driver on kernels  $\geq 5.19$
  - (#7784) Stop building and loading *dis\_intel\_dma* kernel module by default
  - (#7797) Increased max number of downstream ports from 8 to 12
  - (#7831) Fix installer bug related to shell output redirection

### Platform support

- (#7451) Added support for Linux kernel 6.13
- (#7507) Dropped support for Debian 7
- (#7507) Dropped support for Ubuntu 14.04
- (#7556) Added support for debian9-rt and debian10-rt
- (#7588) Added support for GCC 15
- (#7679) Added support for Rocky Linux / RHEL 10
- (#7686) Added support for kernel 6.15
- (#7768) Added support for Concurrent RedHawk 9.2

### SISCI

- (#6917) Added documentation for `DIS_BROADCAST_NODEID_GROUP_ALL` constant for `SCICConnectSegment`
- (#7558) Return code fixes for `SCIQuery`
- (#7578) Fixed issue where requester ID added to adapter was not added to NTB switch table
- (#7591) Simplified SISCI API header file

### SmartIO

- (#6092) Fixed issue with borrowing multiple NVIDIA GPUs
- (#7044) Add support for device lending NVIDIA 5000-series / Blackwell GPUs
- (#7044) Fixed issue with borrowing the Physical Function (PF) of an SR-IOV device. “Failed to export bar7: 22”
- (#7231) Added experimental support for peer-to-peer between borrowed- and local devices. Users of this experimental feature should be aware that P2P mappings set up with *smartio\_tool enable-p2p* are not properly cleaned up until all dolphin drivers are restarted
- (#7295) Fixed issue with NVIDIA GPUs and “Can’t find an IRQ for your NVIDIA card!”
- (#7394) Fixed assertion failure when disabling interrupts on some devices
- (#7394) Fixed issue where `lspci` showed `<broken chain>` for borrowed devices
- (#7485) SmartIO now works on Linux kernels  $\geq 5.19$  when IOMMU is enabled
- (#7530) Support drivers that use `dma_alloc_pages` including *xhci\_hcd*
- (#7542) Fixed a bug that could cause interrupts to be missed on some systems
- (#7553) Fixed bug that allowed users to lend out boot VGA devices
- (#7559) Added support for borrowing local devices
- (#7628) Fixed issue where the MPS was incorrectly downgraded after borrowing a device, causing a fatal AER error

- 
- (#7727) Added `-force` option to `smartio_tool` available
  - (#7815) Add support for NVIDIA A100 and NVIDIA RTX 4500 PRO Blackwell
  - (#7815) Add support for modern GPUs where all BAR regions are occupied

#### **Smartio known limitations**

- (#6092) Borrowing multiple NVIDIA GPUs on PX + MXS switch is not supported

#### **Supersockets**

- (#7645) Supersocket fix for kernel versions 6.0 and higher, fixes crash that occurred when incorrectly interpreting user memory as kernel memory

## **DIS Release 5.23.0**

#### **General**

- Added support for PXH822-L1
- Bug fix: IOMMU now works as expected on kernels  $\geq 6.3$
- Bug fix: Driver configuration now correctly exits if IOMMU initialization fails

#### **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_license_tool`**

- Bug fix: Fixed issue where applying a new license key failed

#### **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

#### **SISCI**

- Added interrupt support for SISCI SmartIO
- Enabled possibility for non-zero offset when calling `SCIMapRemoteSegment` for multicast segments
- Bug fix: Fix for `SCISStartDMATransferMem` 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

## DIS Release 5.22.0

### General

- Added support for Intel Xeon Scalable 5th gen (“Emerald Rapids”)
- 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\_status script, now retrieves more information
- SmartIO: Increased default maximum read request size

### 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 dis\_firmwaretool

- Added support for virtual configurations
- Various bugfixes

### Improved dis\_networkmanager

- Added support for limiting maximum log file size
- Fixed CPU usage issue related to large log file sizes

### Improved installer

- Bugfix in transparent configurations for Rocky 8
- Various bugfixes

### Improved support for intel xeon scalable 3rd

- 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

---

## DIS Release 5.21.4

### General

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

## DIS Release 5.21.3

### General

- No changes for PX

## DIS Release 5.21.2

### General

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

## DIS Release 5.21.1

### General

- 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 Release 5.21.0

### General

- Added support for Linux kernel 5.19
- Added support for Linux kernel 6.0
- 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

### Mxs924

- Eeprom version updated to VER5
- Updated throttling values for DSWITCH topology

## DIS Release 5.20.1

### 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

### **General**

- 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

## **DIS Release 5.20.0**

### **Added support for mxs924 switch topologies**

- NTB
- Transparent
- D-SWITCH

### **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

### General

- 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

## DIS Release 5.19.2

### General

- 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 Release 5.18.1

### General

- 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 Release 5.17.0

### General

- 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

### 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

## DIS Release 5.16.0

### General

- 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)
- = 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 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 Release 5.15.2**

### **General**

- Improved detection and reporting of old firmware

### **SmartIO**

- fixed potential assert in device scanning process

## **DIS Release 5.15.1**

### **General**

- Fixed bug in SISI 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 Release 5.15.0**

### **General**

- Increase max number of adapters from 4 to 5
- 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

### **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

## 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 supports local IP and local-host address)

### SuperSockets user space

- User space version of SuperSockets
- Provides lower latency than the kernel version but provides a limited number of socket functionalities.
- Currently only TCP supported
- Supports communication to Windows SuperSockets

### SISCI 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 2.5.1

- 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 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 strives to support all major Linux distributions and kernels from 2.6.32 and up. This release has been tested on the following 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 - 6.18

### Supported Distributions:

- Ubuntu 24.04 x86\_64
- Ubuntu 22.04 x86\_64
- Ubuntu 20.04 x86\_64
- Rocky / RHEL 10
- Rocky / RHEL 9
- Rocky / RHEL 8
- Debian 13 x86\_64
- Debian 12 x86\_64
- Debian 11 x86\_64

### EOL Distributions With Available Builds:

- Ubuntu 18.04 x86\_64
- Ubuntu 16.04 x86\_64
- CentOS 8 x86\_64
- CentOS 7 x86\_64
- CentOS 6 x86\_64
- Debian 10 x86\_64
- Debian 9 x86\_64
- Debian 8 x86\_64

#### Note

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 for assistance

---

## Installation and management

- RPM-based installation via Self Installing Archive (SIA)
- Automatic configuration via Interconnect Manager
- Centralized monitoring and diagnostic - Both Linux and Windows supported as GUI platforms

## 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 affected connections. SuperSockets will automatically re-establish communication when the cable is reinserted

## Bundled (major) management tools

- **dis\_diag**: diagnostic tool
- **dis\_services**: driver service manager
- **dis\_config**: configuration tool
- **dis\_mkconf**: configuration file generator

## 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
- **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).

For further instructions, please see the “Firmware Release Notes” of your adapter: [https://www.dolphinics.com/download/PX/RELEASE/PW\\_DOC/download-px-doc.html](https://www.dolphinics.com/download/PX/RELEASE/PW_DOC/download-px-doc.html)

Please file a support request at <https://www.dolphinics.com/csp> if you have any problems.

## Change log for Windows Software

**Description of content:** Clustering package for Windows.

**Release category:** GA

**Target audience:** Production systems.

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

### New in DIS Release 5.25.1

#### General

- (#8274) *dis\_diag* now correctly prints the root port's BDF
- (#8342) Fixed CSR value printouts

#### SISCI

- (#8207) SCICreateMapSequence returns error when used on map of local segment
- (#8280) Improved performance for *memcpy\_nt\_avx2*
- (#8341) Fix arithmetic overflow inside SISCI testlib, on systems where CPU frequency is 4GHz+

### Changes in previous releases

#### DIS Release 5.25.0

##### Benchmarks

- (#7882) Adds check for div by zero when walltime is used for *dma\_bench* or *scibench*
- (#7997) Creates a benchmark analogous to *reflective\_bench* but which uses DMA for sending
- (#8063) *dma\_bench* regression fixed so *-veclen* parameter can be used again
- (#8075) Added *-vl* argument to *reflective\_dma\_bench* which can be used to increase DMA transfer vector length

##### Demos

- (#7851) Added performance demo, video transfer demo, and rfm wrapper to DIS example programs

---

## Dis\_diag

- (#7817) New json output when *-json* option is used
- (#7893) dis\_diag output reorganized to print info per adapter
- (#8123) Make the report of the upstream link more generic in dis\_diag

## General

- (#6980) Fixed doorbell performance spikes by removing an unneeded loopback doorbell
- (#7317) The configuration files of DIS hosts no longer supports 'STRIPE' and 'REDUNDANT' socket adapters. A socket adapter specified with this mode will be ignored
- (#7737) Added support for VXP adapter board VPX6-1961
- (#7845) Improved stability with multi-socket Icelake configurations
- (#7891) Fixed problem with driver crash during dis\_config for adapters with invalid hardware configurations
- (#7940) Added support to detect and report the upstream hwport
- (#7968) Fixed issue where *dis\_services stop* would fail if there is no adapter installed in the system
- (#7977) Fixed a bug where a node would end up in a livelock with no session, spamming *getAliveCnt - mbxMsgQ\_localHeader == NULL* to the kernel buffer
- (#8031) Vastly improved PIO performance for the Zen3 architecture
- (#8042) Added DUAL NTB Switch support for MXS924
- (#8156) Improved the detection of adapter types in dis\_firmwaretool
- (#8173) Fixed bug in MXS824 NTB DUAL configs that prevented version 10 of the configs from working
- (#8237) Improved support for Granite Rapids Xeon Scalable CPUs
- (#8240) During MSI upgrade, the internal configuration files must be preserved. *\_dishosts.conf\_* is kept, while **dis\_<adapter>.conf** and **dis\_irm.conf** are overwritten

## SISCI

- (#7336) Multi-threading is now required to compile eXpressWare
- (#7879) Fixed off by one error in *reflective\_bench*
- (#8230) Fixed a bug in *SCICreateDeviceSegment* where it would fail if called with either of the flags *SCI\_FLAG\_LOCAL\_ONLY* or *SCI\_FLAG\_DEVICE\_SIDE\_ONLY*

## SmartIO

- (#8175) Resolved issue with NT-Transparent causing repeated borrow/return to fail
- (#8186) Fixed an issue where two lenders borrowing each others' devices would fail with IOMMU disabled
- (#8226) Fixed issue with automatic DMA windows size calculation with NT-Transparent
- (#8230) Fixed a bug in *SCICreateDeviceSegment* where it would fail if called with either of the flags *SCI\_FLAG\_LOCAL\_ONLY* or *SCI\_FLAG\_DEVICE\_SIDE\_ONLY*

## Supersockets

- (#7629) Sockperf: 1 second walltime is standard measurement
- (#7629) Sockperf: Improved explanation of measurement options
- (#8009) Fixed a memory leak related to the usage of DMA channels

---

## DIS Release 5.24.0

### Benchmarks

- (#7644) Added *-wt* (walltime) option for *dma\_bench* and *scibench2*
- (#7801) Added warmup loops to *scibench2*, *dma\_bench* and *dma\_user\_bench*
- (#7801) Fixed potential sending of too large sizes for *scibench* if offset is used
- (#7801) Made usage text more consistent for *scibench2* and *dma\_bench*

### General

- (#7446) Fixed reset option for MXS924, Port 5 in DSWITCH mode
- (#7491) Verify presence of IOMMU OS mitigation for PCIe devices
- (#7505) Fix buffer overrun resulting in potential Requester ID corruption
- (#7516) Fix DMA transfer failure due to uninitialized variable
- (#7516) Provision extra entry to CAPELLA1 PEX link side to match SapphireRapids DRAC02 silent RID generation
- (#7518) Added option to configure processor I/O requester IDs manually in the configuration files. Particularly useful for heterogenous Intel Xeon Scalable clusters
- (#7797) Increased max number of downstream ports from 8 to 12

### SISCI

- (#6917) Added documentation for `DIS_BROADCAST_NODEID_GROUP_ALL` constant for `SCIDebugSegment`
- (#7558) Return code fixes for `SCIQuery`
- (#7578) Fixed issue where requester ID added to adapter was not added to NTB switch table
- (#7591) Simplified SISCI API header file

## DIS Release 5.23.0

### General

- Bug fix: Driver configuration now correctly exits if IOMMU initialization fails
- Bug fix: Update BAR computation on transparent devices
- Bug fix: Set the DMA function's MPS equal to the one used by the root portChanges in previous releases

### Improved `dis_diag`

- Bug fix: No longer prints invalid temperature

### Improved `dis_nodemgr`

- Unresolvable hosts are skipped and no longer stop cluster configuration

### SISCI

- Enabled possibility for non-zero offset when calling `SCIMapRemoteSegment` for multicast segments
- Bug fix: Fix for `SCIStartDMATransferMem` deadlock

---

## DIS Release 5.22.0

### General

- Added support for Intel Xeon Scalable 5th gen (“Emerald Rapids”)
- Added support for Windows 11, Server 2022
- Improved support for daisy chain configurations
- Improved dis\_status script, now retrieves more information
- Fixed logging issue with dis\_networkmanager
- Added support for virtual configuration through dis\_firmwaretool

### 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

## DIS Release 5.21.4

### General

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

## DIS Release 5.21.3

### General

- No changes for PX

## DIS Release 5.21.2

### General

- 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 Release 5.21.0

### General

- Removed dis\_coinstaller from the driver package
- General bug fixes and improvements

---

## DIS Release 5.20.1

### Bug fixes and improvements

- Fixed initial adapterno value for IRM transparent inf files

## DIS Release 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 Release 5.19.2

### General

- 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 Release 5.18.1

### General

- 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 Release 5.16.0

### General

- Prevent Board Management Software installation on Compute nodes

### Sisci api

- Added support for using multiple DMA channels
- Added support for registering more than one PCIe requester ID per SISI descriptor

## DIS Release 5.15.2

### General

- Added support for SISI Hot Add support for transparent devices

- 
- Improved detection and reporting of old firmware

## DIS Release 5.15.1

### General

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

## DIS Release 5.15.0

### General

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

## DIS Release 5.14.0

### General

- 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 local-host address)
- Connectivity to other systems running Linux user space SuperSockets

### SISI API 2.0.2

- 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

### 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)

## 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 for assistance

## 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 eXpressWare network is unavailable during application startup

---

## Bundled (major) management tools

- **dis\_diag**: diagnostic tool
- **dis\_services**: driver service manager
- **dis\_config**: configuration tool
- **dis\_mkconf**: configuration file generator

## 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
- **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).

For further instructions, please see the “Firmware Release Notes” of your adapter: [https://www.dolphinics.com/download/PX/RELEASE/PW\\_DOC/download-px-doc.html](https://www.dolphinics.com/download/PX/RELEASE/PW_DOC/download-px-doc.html)

Please file a support request at <https://www.dolphinics.com/csp> if you have any problems.

## Change log for RTX Software

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

**Release category:** GA

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

**Currently supported versions:**

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

### New in DIS Release 5.25.1

#### General

- (#8274) *dis\_diag* now correctly prints the root port's BDF
- (#8342) Fixed CSR value printouts

#### SISCI

- (#8207) SCICreateMapSequence returns error when used on map of local segment
- (#8280) Improved performance for *memcpy\_nt\_avx2*
- (#8341) Fix arithmetic overflow inside SISCI testlib, on systems where CPU frequency is 4GHz+

### Changes in previous releases

#### DIS Release 5.25.0

##### Benchmarks

- (#7882) Adds check for div by zero when walltime is used for *dma\_bench* or *scibench*
- (#7997) Creates a benchmark analogous to *reflective\_bench* but which uses DMA for sending
- (#8063) *dma\_bench* regression fixed so *-veclen* parameter can be used again
- (#8075) Added *-vl* argument to *reflective\_dma\_bench* which can be used to increase DMA transfer vector length

##### Demos

- 
- (#7851) Added performance demo, video transfer demo, and rfm wrapper to DIS example programs

### Dis\_diag

- (#7817) New json output when *-json* option is used
- (#7893) dis\_diag output reorganized to print info per adapter
- (#8123) Make the report of the upstream link more generic in dis\_diag

### General

- (#6980) Fixed doorbell performance spikes by removing an unneeded loopback doorbell
- (#7317) The configuration files of DIS hosts no longer supports 'STRIPE' and 'REDUNDANT' socket adapters. A socket adapter specified with this mode will be ignored
- (#7737) Added support for VXP adapter board VPX6-1961
- (#7845) Improved stability with multi-socket Icelake configurations
- (#7891) Fixed problem with driver crash during dis\_config for adapters with invalid hardware configurations
- (#7940) Added support to detect and report the upstream hwport
- (#7968) Fixed issue where *dis\_services stop* would fail if there is no adapter installed in the system
- (#7977) Fixed a bug where a node would end up in a livelock with no session, spamming *getAliveCnt - mbxMsgQ\_localHeader == NULL* to the kernel buffer
- (#8031) Vastly improved PIO performance for the Zen3 architecture
- (#8042) Added DUAL NTB Switch support for MXS924
- (#8156) Improved the detection of adapter types in dis\_firmwaretool
- (#8173) Fixed bug in MXS824 NTB DUAL configs that prevented version 10 of the configs from working
- (#8237) Improved support for Granite Rapids Xeon Scalable CPUs

### SISCI

- (#7336) Multi-threading is now required to compile eXpressWare
- (#7879) Fixed off by one error in reflective\_bench
- (#8230) Fixed a bug in *SCICreateDeviceSegment* where it would fail if called with either of the flags *SCI\_FLAG\_LOCAL\_ONLY* or *SCI\_FLAG\_DEVICE\_SIDE\_ONLY*

### SmartIO

- (#8175) Resolved issue with NT-Transparent causing repeated borrow/return to fail
- (#8186) Fixed an issue where two lenders borrowing each others' devices would fail with IOMMU disabled
- (#8226) Fixed issue with automatic DMA windows size calculation with NT-Transparent
- (#8230) Fixed a bug in *SCICreateDeviceSegment* where it would fail if called with either of the flags *SCI\_FLAG\_LOCAL\_ONLY* or *SCI\_FLAG\_DEVICE\_SIDE\_ONLY*

### Supersockets

- (#7629) Sockperf: 1 second walltime is standard measurement
- (#7629) Sockperf: Improved explanation of measurement options
- (#8009) Fixed a memory leak related to the usage of DMA channels

---

## DIS Release 5.24.0

### Benchmarks

- (#7644) Added *-wt* (walltime) option for *dma\_bench* and *scibench2*
- (#7801) Added warmup loops to *scibench2*, *dma\_bench* and *dma\_user\_bench*
- (#7801) Fixed potential sending of too large sizes for *scibench* if offset is used
- (#7801) Made usage text more consistent for *scibench2* and *dma\_bench*

### General

- (#7446) Fixed reset option for MXS924, Port 5 in DSWITCH mode
- (#7491) Implement *iommu\_verify.exe* as precondition for driver start
- (#7505) Fix buffer overrun resulting in potential Requester ID corruption
- (#7516) Fix DMA transfer failure due to uninitialized variable
- (#7518) Added option to configure processor I/O requester IDs manually in the configuration files. Particularly useful for heterogenous Intel Xeon Scalable clusters
- (#7797) Increased max number of downstream ports from 8 to 12
- (#7810) Minor update on *RtssStack.ps1* to keep support for RTX64 v3.7

### SISCI

- (#6917) Added documentation for `DIS_BROADCAST_NODEID_GROUP_ALL` constant for `SCICConnectSegment`
- (#7558) Return code fixes for `SCIQuery`
- (#7578) Fixed issue where requester ID added to adapter was not added to NTB switch table
- (#7591) Simplified SISCI API header file

## DIS Release 5.23.0

### General

- General bug fixes and improvements

## DIS Release 5.22.0

### General

- Improved PX adapter function detection
- Added support for driver configuration when the monolith must be available on a limited account

### SISCI

- General bug fixes and improvements

## DIS Release 5.20.1

### General

- General bug fixes and improvements

---

## DIS Release 5.20.0

### General

- General bug fixes and improvements

## DIS Release 5.19.2

### General

- General bug fixes and improvements

## DIS Release 5.18.1

### General

- General bug fixes and improvements

## DIS Release 5.17.0

### General

- General bug fixes and improvements

### 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

## DIS Release 5.15.2

### General

- Improved detection and reporting of old firmware

## DIS Release 5.15.0

### General

- Increase max number of adapters from 4 to 5

## Included software

### SISCI API 2.0.2

- 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

## 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
- **dis\_config**: configuration 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

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).

For further instructions, please see the “Firmware Release Notes” of your adapter: [https://www.dolphinics.com/download/PX/RELEASE/PW\\_DOC/download-px-doc.html](https://www.dolphinics.com/download/PX/RELEASE/PW_DOC/download-px-doc.html)

Please file a support request at <https://www.dolphinics.com/csp> if you have any problems.

## Change log for VxWorks Software

**Description of content:** Clustering package for VxWorks

**Release category:** GA

**Current status:** Contact Dolphin.