



Dolphin Interconnect Solutions

eXpressWare

Software release notes

for IX products

5.16.0

Date: 8th May 2020

Table of Contents

Definitions.....	3
Release categories.....	3
Release version scheme.....	4
Types of releases.....	4
Release candidate.....	4
Release.....	4
Dolphin eXpressWare.....	5
SuperSockets.....	5
SuperSockets User-space Library for Windows and Linux.....	5
SuperSockets Kernel Space library for Linux.....	5
Dolphin technical support and documentation.....	6
Supported Hardware and configurations.....	7
Supported Dolphin hardware.....	7
Configurations supported by this release.....	7
Support for OEM hardware.....	7
General eXpressWare defaults and settings for IX.....	8
Dolphin eXpressWare IX release DIS 5.16.0 Linux.....	9
Changes in previous releases.....	9
Dolphin eXpressWare IX release DIS 5.16.0 Windows.....	12
Changes in previous releases.....	12
Dolphin eXpressWare IX release DIS 5.16.0 RTX.....	14
Changes in previous releases.....	14
Dolphin eXpressWare IX release DIS 5.16.0 VxWorks.....	16
Changes in previous releases.....	16

Definitions

Release categories

The following definitions are used for the eXpressWare software.

Prototype:

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

Alpha:

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

Beta:

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

GA:

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

Snapshot:

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

Release version scheme

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

MAJOR_VER

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

MINOR_VER

Changes to the software that significant improve functionality. Minor changes and improvements that does 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. The ideas are collected from and widely used within the Linux community. These two types applies to Prototype, Alpha, Beta and GA releases.

Release candidate

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

Release

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

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

Dolphin eXpressWare

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

IRM	Interconnect Resource Manager
SISCI	Software Infrastructure for Shared-Memory Cluster Interconnects. This includes binaries , sources, documentation and development tools.
SuperSockets	Socket accelerator for PCIe
IPoPCIe	TCP-IP driver for PCIe. Windows NDIS driver or Linux DISip.
SmartIO	Sharing and access to IO management and control. Includes Device Lending and SISCI SmartIO management. SmartiIO is only available with Dolphins PX and MX products.
IO Monitoring	Transparent Board Management Software for transparent PCIe adapter cards and cable connections. Transparent Board Management is currently only available with Dolphins PX products.

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 supports the following socket calls:

accept, bind, connect, getpeername, getsockname, getsockopt, listen, recv, select, send, setsockopt, shutdown.

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.

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 sending an email to pci-support@dolphinics.com.

Additional information, installation manuals etc can be found at

http://www.dolphinics.com/support/index_support_ix.html

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 hardware

This software release supports the following PCIe adapter cards based on IDT PCIe technology.

- IXH610, IXH611, IXH620, IXH650, IXH631

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

Supported Dolphin PCI Express switches

- IXS600

Configurations supported by this release

The software release support the following configurations.

General use with IXH610/IXH611/IXH620

- 2 to 20 nodes using multiple IXS600 switches
- 2 to 8 nodes using a single IXS600 switch
- 2 nodes using a single PCIe x8 iPass cable
- Support for redundant fabrics
 - Up to 4 Cards pr host connected directly or through separate switches are supported at the SISCI level.
 -

Reflective memory use with with IXH610/IXH611/IXH620

- 2 nodes using a single PCIe x8 cable
- 2 to 56 nodes using multiple IXS600 switches
 - 20 lower NodeIds may still be used for general clustering
 - 36 high NodeIds can only communicate using reflective memory. Please contact Dolphin for more information.
 -

General use with IXH631 (4 port MiniSAS-HD):

- 2 and 3 nodes x8 without switch
- 2 - 5 nodes x4 without switch
 - Connectivity is point to point, no re-routing is supported.

PCI Express Hot Add

The IXH610, 611, IXH620 cards does not support PCIe Hot add. Please contact Dolphin for more information.

Support for OEM hardware

This version of Dolphin eXpressWare has general support for all IDT PCIe NTB enabled chipset 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.

Please note that Dolphin eXpressWare also support Broadcom and Microsemi PCIe chipset. Please contact Dolphin for more information.

General eXpressWare defaults and settings for IX

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

- SISI Multicast
 - Default setting 4 groups, 2 Megabyte per segment.
 - Max setting 4 groups, 1 Gigabyte per segment.

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

Dolphin eXpressWare IX release DIS 5.16.0 Linux

Description of content: Clustering package for Linux

Release category: GA

Target audience: Production systems.

Release date: 21st April 2020.

Current status: Available for download from
https://www.dolphinics.com/support/index_support_ix.html

New in DIS IX 5.16.0:

- Added support for system DMA with IOMMU ON (Off already supported)
- Added support for Linux kernel \geq 5.5.3. Tested on Linux kernel 5.5.8
- SISI API
 - Added support for using multiple DMA channels.
 - Added support for registering more than one PCIe requester ID per SISI descriptor.

Changes in previous releases

DIS IX 5.15.2:

DIS IX 5.15.1:

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

DIS IX 5.15.0:

- Increase max number of adapters from 4 to 5.
- SIA Installer: Fixed problem with the --install-all option
- User-space SuperSockets: Fixed problem with maximum number of connections decreasing over time.

DIS IX 5.14.0:

- Supports up to kernel 5.3
- Added support for Debian 10
- Added support for Ubuntu 18.04 with HWE kernel
- SuperSockets:
 - Fixed problem with crash during accept() if NO_NATIVE_CONN is use inconsistently across the cluster.
- General bug fixes and improvements.

DIS IX 5.13.1:

- SISI
 - SISI virtual devices can now hold up to 32 instances of each resource type.
 - Max multicast group size increased to 4 x1 GB. Default setting is 4 x 2MB.
- Improved SIA installers for all platforms. Added advanced configuration option allowing only required modules to be installed.
- Fixed problem compiling eXpressWare from tarball source.
- Support for new Linux kernels and distributions.
- General bug fixes and improvements.

Please note (for 4.x.x users only): All utilities previously having an “ix” prefix (eg. ixdiag) has now been renamed to use a “dis” prefix (e.g dis_diag).

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 Windows 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 Windows, RTX or VxWorks.
- 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 and Dolphin eXpressWare IX.

Tested and supported OS platforms and distributions:

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 support for another kernel:

Linux Kernel 2.6.32 - 5.5.8

- 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
- 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 interrupt benchmark)
- reflective_bench (SISCI based reflective memory benchmark)
- latency_bench (TCP latency benchmark, runs on any IP network including SuperSockets)

Firmware upgrade:

All older IXH cards needs a firmware update to support DIS 5.x drivers. The Driver installer will detect and suggest to automatically update the firmware during installation if an update is required (new firmware included in software distribution). Please upgrade by running `/opt/DIS/sbin/upgrade_eeprom.sh --upgrade`. A reboot is required after firmware upgrade.

Dolphin eXpressWare IX release DIS 5.16.0 Windows

Description of content: Clustering package for Windows

Release category: GA

Release date: 21st April 2020.

Current status: Available for download from

https://www.dolphinics.com/support/index_support_ix.html

New in DIS IX 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.

Changes in previous releases

DIS IX 5.15.1:

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

DIS IX 5.15.0:

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

DIS IX 5.14.0:

- 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
- Improved allocation of large memory areas
- General bug fixes and improvements.

DIS IX 5.13.1:

- SISCI
 - Max multicast group size increased to 4 x1 GB. Default setting is 4 x 2MB.
- Fixed crash that may occur during reboot of system.
- 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 V2.0.2
 - 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 7 - 32 bit
Windows 7 - 64 bit
Windows 8 - 32 bit
Windows 8 - 64 bit
Windows 8.1 – 64 bit
Windows 10 – 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 eXpressWare 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:

IXH 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 IX\Util\upgrade_eeprom.cmd --upgrade

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

Dolphin eXpressWare IX release DIS 5.16.0 RTX

Description of content: Clustering package for RTX

Release category: GA

Release date: 21st April 2020.

Current status: Available. Contact Dolphin.

New in DIS IX 5.16.0:

.

Changes in previous releases

DIS IX 5.15.0:

- Increase max number of adapters from 4 to 5.

DIS IX 5.14.0:

- Improved allocation of large memory areas
- General bug fixes and improvements.

DIS IX 5.13.1

- SISI
 - Max multicast group size increased to 4 x1 GB. Default setting is 4 x 2MB.
- General bug fixes and improvements.

Included software:

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

Supported OS platforms:

Windows Server 2008 R2 - 64 bit

Windows Server 2012 R2 – 64 bit

Windows Server 2016 – 64 bit

Windows 7 - 64 bit

Windows 8.1 – 64 bit

Windows 10 – 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 (SISI based PIO throughput benchmark)
- scipp (SISI 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:

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

```
C:\Program Files\Dolphin Express RTSS\Util> upgrade_eeprom.cmd --upgrade
```

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

Dolphin eXpressWare IX release DIS 5.16.0 VxWorks

Description of content: Clustering package for Windows

Release category: GA

Release date: 21st April 2020.

Current status: Available. Contact Dolphin.

New in DIS IX 5.16.0:

- Added support for VxWorks 7 SR0620 with LLVM/Clang compiler
- Fixed VxWorks initialization ordering
- Added CDF parameter EXPRESSWARE_SLOT_0_NODEID to explicitly define the auto-configuration NodeId if the geographical address (slot number) is reported as 0 by the BSP.
- SISI API
 - Added support for using multiple DMA channels.
 - Added support for registering more than one PCIe requester ID per SISI descriptor.

Changes in previous releases

DIS IX 5.15.1:

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

DIS IX 5.15.0:

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

DIS IX 5.14.0:

- Improved allocation of large memory areas
- General bug fixes and improvements.

DIS IX 5.13.1:

- SISI
 - Max multicast group size increased to 4 x1 GB. Default setting is 4 x 2MB.
- General bug fixes and improvements.

Included software:

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

Supported OS platforms:

VxWorks 6.9 Intel x64 SMP

VxWorks 6.9 Intel x86 SMP

VxWorks 7.0 Intel x64 SMP

VxWorks 7.0 Intel x86 SMP

Please note that AMP model and User-space communication is not supported with this release.

Please contact Dolphin for more information.

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:

IXH 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 contact pci-support@dolphinics.com if you have any problems.