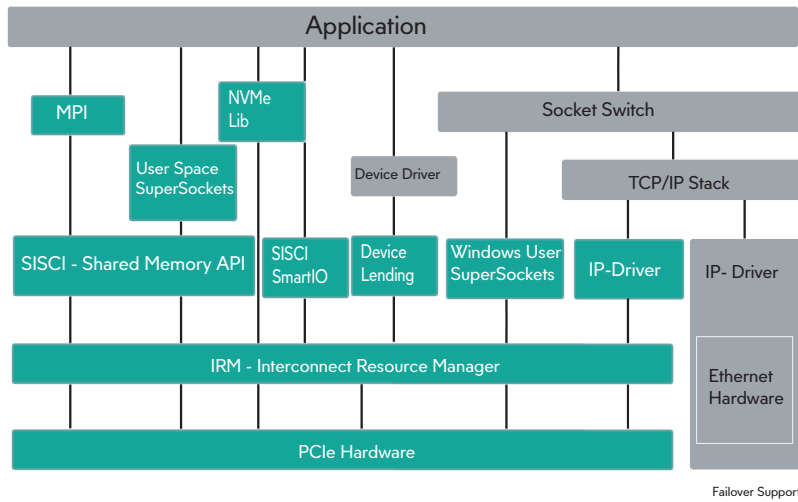


eXpressWare



Overview

eXpressWare is an advanced software suite for PCI Express hardware. It supports the non-transparent features in many PCIe switches from leading PCIe vendors.

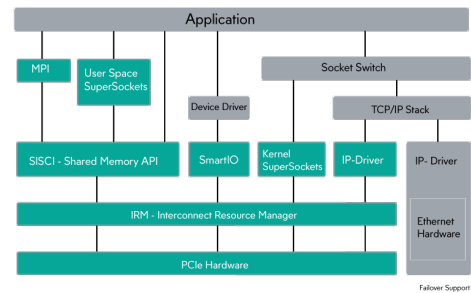
eXpressWare includes a low level API called SISCI, a sockets API - SuperSockets, an IPoPCIE driver, and advanced IO software - SmartIO for sharing PCIe devices.

eXpressWare comes with all Dolphin eXpressWare NTB adapters. It also can be licensed for third party PCIe hardware.

SISCI API

Dolphin’s Software Infrastructure Shared-Memory Cluster Interconnect (SISCI) API makes developing PCI Express® Network applications faster and easier. The SISCI API is a well-established API for shared memory environments and can be used for both intra-and inter-system communication.

In a multiprocessing architecture with PCI Express®, the SISCI API enables PCI Express® based applications to use distributed resources such as CPUs, I/O, and memory. The resulting application features reduced system latency and increased data throughput. The Dolphin SISCI Developers kit consists of driver and API software, tools, documentation and source files needed to develop your own embedded application. The development kit provides a C system call interface to ease customer integration to the PCI Express Network. SISCI enables customer application to easily bypass the limitations of traditional network solutions, avoiding time consuming operating system calls, and network protocol software overhead. The SISCI software supports clusters of hundreds of nodes.



Benefits

- Low Latency API designed for direct memory access between multiple systems and devices
- Supports Windows, Linux, VxWorks, and RTX
- Cross O/S support between Windows and Linux
- Example code and performance test suite for latency and throughput
- Peer-to-Peer support and memory writes into device memory of GPUs and FPGAs

Features


Shared memory API for clusters and SMP's	Create and trigger remote application level interrupts
Reflective Memory / PCIe Multicast	Direct remote CPU load and stores, memcpy()
PCI Express Peer-to-Peer support	Windows, RTX, VxWorks and Linux OS support
Distributed shared memory	Caching and error checking support
Easy deployment of DMA transfers	Events and callbacks

Device Lending

PCIe Device Lending offers a flexible way for NVMeS, FPGAs, GPUs, and other devices to be accessed within a PCIe Fabric. PCIe devices can be borrowed between nodes in a PCIe Fabric. This is done without any software overhead and at PCIe data rates.

Device Lending is a simple way to reconfigure systems and reallocate resources. GPUs, NVMe drives or FPGAs can be logically added from and to any system in the fabric. The result is a flexible simple method to maximize usage by creating a pool of devices. Since this solution uses standard PCIe, it doesn't add any software overhead to the communication path. Standard PCIe transactions are used between the systems.

Device lending also enables a SR-IOV device to be shared as a MR-IOV device. SR-IOV functions can be borrowed by any system in the PCIe Fabric. Thereby enabling the device to be shared by multiple systems. This maximizes the use of SR-IOV devices such as 100 Gbit Ethernet cards.



Benefits

- Dynamic reallocation of NVMe Drives and other resources
- Flexible use of GPUs
- Sharing of SR-IOV devices

Features

Linux OS support	Lending and Borrowing Software
Command line interface	API to enable scripting
Any system can be lender and borrower	Supports GPUs, FPGAs, Ethernet cards, NVMe drivers and other devices
Supports Hot add supported devices only	Enables sharing of devices with SR-IOV between multiple system.


PCIe Hot Add

PCIe NTB technology enables software based resource management and enumeration of standard PCIe devices. This process of adding or removing devices can be automatic or managed through some basic command line utilities.


The main difference between PCIe hot add and traditional expansion setups is that the PCIe host uplink card is replaced with a Dolphin PCIe NTB card. eXpressWare software is used to enumerate downstream devices vs the system BIOS. The solution works with directly attached PCIe devices as well as PCIe devices installed into PCI Express expansion systems.

The solution does not require any special BIOS version. Hot adding many devices or devices that require a large PCIe address space may require the BIOS to support 64bit encoding and large PCIe BARs.


Single device:



Multiple devices in an expansion box:



Using a switch to connect multiple expansion boxes:



Benefits


- Add multiple devices or new devices to a running system
- No power-on sequencing
- Dynamic Hot Add and replace
- Works with standard PCIe devices (NVMe drives)

Features

NTB host adapter	Dolphin eXpressWare software
Software based resource management and enumeration	Automatic management of device adding and removing
Scalability	Add or remove a number of devices
No prior installation required	No changes to device drivers

SuperSockets

Dolphin’s SuperSockets™ delivers maximum application performance without necessitating application changes. SuperSockets™ is a unique implementation of the Berkeley Sockets API that capitalizes on the PCI Express® transport. It transparently achieves performance gains for existing socket-based network applications. Dolphin Express hardware and the SuperSockets™ software layer create an ultra-low latency, high-bandwidth, low overhead, and high availability platform to support the most demanding sockets-based applications.



Benefits

- Increased application performance without changing your application
- Extremely low latency with under 1µs of latency
- Highly available network with failover to Ethernet
- High bandwidth and throughput

Features

Compliant with Linux Socket library and Berkeley Sockets	No OS patches or application modifications required
All popular Linux distributions supported	Supports both user space and kernel space clients
Both TCP and UDP support	Transparent failover to Ethernet if high speed connection is down.

IPoPCIE

Dolphin’s performance optimized TCP IP driver for PCIe (IPoPCIE) provides a fast and transparent way for any networked applications to dramatically improve network throughput. The software is highly optimized to reduce system load (e.g. system interrupts) and uses both PIO and RDMA operations to implement most efficient transfers for all message sizes. The major benefits are plug and play and much higher bandwidth than network technologies like 10G Ethernet.

Feature	SISCI	SuperSockets	IPoPCle Driver	SISCI SmartIODevice Lending	
Linux Platform Support	x86, x86_64	x86, x86_64	x86, x86_64	x86, x86_64	
Kernel 2.6 - 4.x	ARM 32/64. PPC	ARM 64 PPC	ARM 64 PPC	ARM 32/64. PPC	x86, x86_64
Windows Vista, 7, 10, Server 2008-2019 Platform Support	x86, x86_64	x86, x86_64	x86, x86_64	coming	2020
VxWorks 6.9 , 7.0	x86, x86_64 PPC	No	No	No	No
Latency	0.54 μ s Linux	0.98 μ s Linux	5.6 μ s Linux	Native PCIe	Native PCIe
Max Bandwidth	13 GB/s Linux	65 Gb/s	65 Gb/s	Native PCIe	Native PCIe
Reflective Memory Support	Yes	N/A	N/A	N/A	N/A
Address based Multi-cast support	Yes	N/A	N/A	N/A	N/A
TCP and UDP Support	N/A	TCP -Yes UDP Linux - Yes UDP Windows - No	Yes	N/A	N/A
UDP Multi-cast Support	N/A	Linux -Yes Windows -No	Linux - No Windows - Yes	N/A	N/A
Application modifications required or unmodified device driver	Yes	No	No	Yes	No
Cross O/S and Platform data transfer support	All	User space Linux / Windows	No	N/A	2020
Peer to Peer Transfers support	Yes	N/A	N/A	Yes	Yes
Sharing of NON-SR-IOV devices				Yes	No
Sharing SR-IOV devices to multiple hosts				Yes	Yes
Direct Transfers NVME to GPU / FPGA buffers				Yes	No
Flexible device memory management				Yes,	No
Support for GPU-Direct				Yes	Yes
Support for PCIe Multicast - device				Yes	No
PCIe DMA support				Yes	No
Device DMA support				Yes	Yes