



Case Study: Third I/O Product: TargetConnect™ SDK

The TargetConnect SDK provides a time-to-market shortcut for developing feature-rich Linux SCSI target mode applications based on the Emulex-developed Linux driver

Third I/O

Third I/O Incorporated specializes in high-speed bandwidth and supercomputing technologies with expertise in the enterprise server, storage, and networking industries. The company, which focuses on highly efficient development and testing methodologies, was founded on the belief that intensive evaluation and testing is critical in delivering high quality products. The company has a two-fold mission: as the developer and distributor of Iris, a bandwidth-intensive data storage solution which was recently benchmarked at over 11 GB/s and 1.6 million I/O operations per second; and as a consultant for leading technology companies to create customized testing and evaluation studies and benchmarks. Third I/O can be found at www.thirdio.com.

Mark Lanteigne, founder of Third I/O, admits he has always had a fascination with coaxing as much performance out of a system as possible. "It's always fun to see how fast you can push a system and also watch a variety of bugs surface during high bandwidth testing," he says.

During his years of enterprise testing, Lanteigne discovered that a particular dynamic of system performance was routinely overlooked: the performance or bandwidth of embedded or slot-based peripherals such as storage and networking adapters. "Many times, these performance anomalies can be traced back to poor design or motherboard layout, but they also can be signs of deeper systemic problems," Lanteigne explained. "A true annoyance of mine is when a server only allows for a fraction of the advertised available bandwidth. And unfortunately, this is often the case."

So Lanteigne founded Third I/O in order to develop its own bandwidth-intensive product to resolve these performance issues. With years of experience in enterprise server development strategies, it didn't take long for Third I/O's "bandwidth brigade" to develop and deploy Iris, a hybrid data storage and computing platform that allows for extraordinary levels of bandwidth and I/O operations per second (IOPS) while still maintaining the ability to perform as a fully functional enterprise server platform. Recently, Third I/O showcased extraordinary performance at Supercomputing 2008 in a collaborative benchmark with Medusa Labs, Dell Computer, AMD and Emulex Corporation, documenting a sustained data rate of over 1.6 million I/O operations per second (IOPS), using Third I/O's Iris Storage Platform.

Software Developer Kit Enables Greater Performance, Design Customization and Proprietary Functionality

Storage solution developers like Third I/O experience intense deadline and budget pressure to design and deliver advanced features and functionality required by networked storage applications. The superior data throughput achieved by the Iris solution early on is in part due to a solid partnership between Emulex Corporation and Third I/O, in which Third I/O engineers leveraged the advanced development tools in the Emulex TargetConnect Software Developer Kit (SDK).

The TargetConnect SDK provides the flexibility to quickly develop and deploy feature-rich storage solutions based on Emulex SCSI target mode drivers. These solutions include sophisticated storage devices which provide advanced functionality, including test and diagnostics solutions (such as Iris from Third I/O), storage appliances, backup engine devices, deduplication, and more.



Case Study: Third I/O Product: TargetConnect™ SDK



The Emulex SDK includes an Emulex Lightpulse® driver (Linux initiator base driver) delivered in source code form with an Emulex target mode API added to the base driver. Developers like Third I/O write directly to the Emulex target mode API which interfaces to a custom SCSI target mode driver (a sample SCSI target mode driver with simulated disks is available as a reference implementation to guide the development of a custom driver).

“The Linux target mode API driver from Emulex gave us more control of our development efforts to create an infrastructure for a target mode application at a low level-maintained API,” Lanteigne said. “This feature provided us with greater design customization to achieve the performance and proprietary functionality we required in the first release of our I/O storage platform.”

The Iris solution does not require specialized or proprietary hardware as is the case with most enterprise storage devices. It was designed to transform any x86-64 Intel or AMD system into a high-speed Fibre Channel or FCoE solid state disk (SSD), and still allow the same platform to perform the duties of a single or multiple virtual OS platform. The Iris solution can replace the I/O and bandwidth of thousands of hard drives or dozens of enterprise storage platforms, runs entirely from a bootable CD, and was developed as an easy-to-operate storage application.

“The flexibility of our product is in part due to the flexibility of the Emulex API, which placed no material restrictions to our development,” Lanteigne said. “In fact, it comes with code that enables a functional sample target that allowed us to expedite our development by several weeks. The fact that this code is so comprehensive enabled a smoother architectural and coding process from start to finish.”

Lanteigne continued to explain that other vendors offer very little or nothing in terms of proactive support for their partners’ target mode development efforts. “Emulex is far superior by offering a full target mode SDK. Not only does Emulex provide a nice framework and API, it’s gone one step further by including a functional reference target to illustrate how to get started on a full-blown target mode development.”

I/O Intensive Applications Demand Bandwidth Intensive Storage

Third I/O launched the Iris Storage Platform at a time when bandwidth and I/O intensive applications are an essential component of high performance computing, business, and even daily life. High bandwidth storage refers to computer storage that is capable of extremely fast data transmission rates. A high bandwidth storage device is one that can read or write data more rapidly than standard storage devices such as hard drives or RAID arrays. Among the applications that have the highest need for data throughput are streaming media, data warehousing (required to analyze key business trends) and server virtualization (which greatly reduces IT capital expenses and data center space requirements). Because many of these applications are in their early deployment, the requirement for high speed storage will grow exponentially for many years.

“We created Iris as a hybrid solution because the current roadmap of virtualized computing dictates that significantly more work and operations need to be performed while using an equivalent—or even reduced—hardware infrastructure,” he said.

Lanteigne continued to explain that virtualization is going to be much more demanding of server hardware. “With multiple operating systems and applications on a single platform, the need for additional CPU, memory, and external bandwidth all increase dramatically. And Fibre Channel connectivity is playing an important role in meeting this need.”

The performance of a Fibre Channel HBA is often measured using two common metrics. One is the number of I/O (input/output) transactions that can be performed per second, referred to as IOPS. The second way to measure performance is the number of bytes of data that can be moved per second, which is measured in MB/s.

IOPS performance rates are a critical comprehensive system benchmarking metric. This is because high I/O transaction rates require several components of a system to work together efficiently in order to achieve respectable results. Small block I/O sizes can be very taxing on a system, both from a software and hardware perspective. IOPS testing has a higher overhead because every I/O is required to travel through the operating system’s SCSI layers and I/O scheduler. In general, there is also a direct correlation between number of IOPS and the number of times that a storage controller will physically interrupt the CPU subsystem. With small block I/O benchmarking, the generation of more IOPS leads to significantly more accesses on the SCSI layer, OS I/O scheduler, and the number of interrupts sent to the processors. This makes IOPS analysis a great benchmark for examining the OS, the storage controller, and the host processors simultaneously.

“In our opinion, Emulex has the best IOPS performance profile in the industry,” stated Lanteigne. “The ability to run at extremely high bandwidth and IOPS has been a collaborative and focused effort by both Emulex and Third I/O. One point of certainty, both companies have been and are still working very diligently to create the fastest and most reliable target in the industry.”

A Partnership Based on Leading Technology and Unparalleled Support

The TargetConnect SDK provides a time-to-market shortcut for developing feature-rich Linux SCSI target mode applications based on the Emulex-developed Linux driver. Additionally, the device driver and application architecture is fully backward and forward compatible, which means that a development on an Emulex 8Gb/s HBA solution will also run on a 10Gb/s FCoE CNA as well as on previous generations of the Emulex product line.

Case Study: Third I/O Product: TargetConnect™ SDK



“From a software development point of view, backwards driver compatibility is as critical as forward compatibility,” Lanteigne said. “It allows for a very streamlined process, enables us to develop one solution that will work on many products, and opens opportunities into new product mixes with relatively few software changes. I’ve been surprised to see that other manufacturers have released software and/or drivers that are quite specific to product generations or specific revisions of silicon, which inhibits our development process.”

A unique feature to the Emulex software developer kit is that partners who sign a TargetConnect™ software license agreement will benefit from the world-class support infrastructure at Emulex, including dedicated technical support and engineering training.

“Based on my many years of experience, I believe that Emulex has the best hardware and software engineers in the storage industry,” stated Lanteigne. “Emulex engineering has been very supportive of our development since day one and our communication has been strong, fast, efficient, and bi-directional. They assist us when we encounter issues, have implemented many of our requests, and are proactive in updating the latest specifications to the target mode kit.”

The first-generation Iris solution ran exclusively on Emulex LightPulse® 4Gb/s Fibre Channel Host Bus Adapters (HBAs). In 2008 the company upgraded to support the LightPulse 8Gb/s HBAs, and in 2009 added Emulex 10 Gb/s CNA FCoE adapters into the mix. “Our engineers evaluated other Fibre Channel offerings, but ultimately chose Emulex for a variety of technical and support reasons,” Lanteigne remembers. “Among other factors, we found Emulex to have the best Fibre Channel performance of tested Fibre Channel vendors, excellent performance scaling across multiple ports, and the best error recovery at the Fibre Channel and system level. Additionally, nothing could compete with its flexible target mode development kit.”

The First Storage Platform to Achieve 1 Million IOPS

In July of 2008, Third I/O announced the results of a collaborative benchmark with AMD and Emulex which achieved a sustained data rate of over 1.1 million I/O operations per second (IOPS), using Third I/O’s Iris Storage Platform. These results mark the first time that over one million IOPS has been achieved on a single server or storage hardware platform. In addition, 100% of the benchmarked configuration was created using shipping and proven hardware and software.

The configuration consisted of two HP DL 585 G5 servers connected via Emulex LightPulse LPe12002 8Gb/s Fibre Channel adapters. Both servers were populated with four AMD Opteron

Quad Core 8360 processors, which allowed for an extraordinary level of IOPS and bandwidth performance. The first system was configured as a storage initiator, running Windows 2008 Enterprise Server x64. The second system was configured running Third I/O’s Iris Storage Platform, allowing it to operate as a high-speed Fibre Channel solid state disk device. The open source Iometer benchmark was used to verify I/O data rates and disk access times. In addition to the groundbreaking IOPS performance, Third I/O’s Iris access times averaged an extraordinary 470 microseconds.

This unsurpassed benchmark recorded 1,122,742 and 1,088,976 IOPS on 512 byte and 1 KB I/O sizes respectively. The complete report can be viewed at www.emulex.com/1million.pdf.

Fortunately, Emulex and Third I/O continue to pursue the highest levels of storage performance in the industry, and Lanteigne revealed that several of his customers bear witness to this fact daily. “In 2009, multiple customers have been reporting single platform Iris data rates in excess of 11 GB/s and 1.6 million IOPS. It is our belief that the Third I/O and Emulex combination will shatter the 2 million IOPS barrier in 2010.”

Conclusion

“Our early evaluation showed us that Emulex was the best choice for Third I/O to partner with to develop a target mode solution,” Lanteigne recalled. “The fact that it had a target mode SDK made the partnership even more attractive. To the best of my knowledge, Emulex is the only major storage silicon vendor that has a formalized target mode development team and project.”

Third I/O has embraced the Linux Open Source features of the Emulex target mode SDK and transformed a comprehensive engineering conception into an efficient, stable and easy-to-use proprietary platform with unprecedented performance results.

Lanteigne concluded by saying that the Iris storage solution plays an important role in external bandwidth testing and public results in revealing performance flaws in data center infrastructures.

“What good is a massively virtualized server if the data moving in and out of the server is being transferred at a snail’s pace?” he questioned. “Our solution was built to jump start performance today and enable it for more growth tomorrow. We can deliver this expandability because Iris is based on the open technology of Linux and supported across all generations of Emulex’s advanced, HBA technology for enterprise-class SANs.

“In the words of Ricky Bobby (*Talladega Nights*, 2006), ‘If you ain’t first, you’re last!’ So, let’s go looking for the fastest server shipping today. And let’s expose the pretenders to the throne along the way!”



www.emulex.com

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600
Wokingham, UK +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547
Tokyo, Japan +81 3 5322 1348 | **Bangalore, India** +91 80 40156789

Connect with Emulex

twitter.com/emulex [friendfeed.com/emulex](https://www.facebook.com/emulex) bit.ly/emulexlinks bit.ly/emulexfb

©2009 Emulex, Inc. All rights reserved. This document refers to various companies and products by their trade names. In most, if not all cases, their respective companies claim these designations as trademarks or registered trademarks. This information is provided for reference only. Although this information is believed to be accurate and reliable at the time of publication, Emulex assumes no responsibility for errors or omissions. Emulex reserves the right to make changes or corrections without notice. This report is the property of Emulex and may not be duplicated without permission from the Company.