

# Technical Brief

**Infortrend**<sup>®</sup>

## Introduction to SANWatch – EonPath

### **Abstract**

This document introduces SANWatch's multipathing driver, EonPath and its data path failover and load-balancing mechanism.

### **Infortrend<sup>®</sup> Technology, Inc.**

Networked Storage Solution Provider

Revision 1.0

Aug, 2008

Infortrend Technology makes no representations or warranties with respect to the contents hereof and specifically disclaims, any implied warranties of merchantability or fitness for any particular purpose. Infortrend reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify of such changes.

Infortrend, Infortrend logo, and EonStor are registered trademarks of Infortrend Technology, Inc.

## What is EonPath?

EonPath is a multipathing driver developed by Infortrend. It can recognize and manage up to 32 redundant paths between a server and an individual RAID volume. In the event of cabling component failure, the path failover mechanism can prevent data flow from being interrupted, thus increasing the fault-tolerant capability of the system. EonPath can also redistribute I/O requests among the multiple paths to balance workloads and optimize throughput across data links. In addition, users can also monitor path throughput performance and flexibly enable or disable paths to meet application needs. Once users have EonPath installed onto the in-band hosts<sup>1</sup>, they can launch it through SANWatch and do the configuration settings with only a few clicks. The supported operating systems of EonPath include Windows, Linux and Solaris.

## Data Path Failover Mechanism

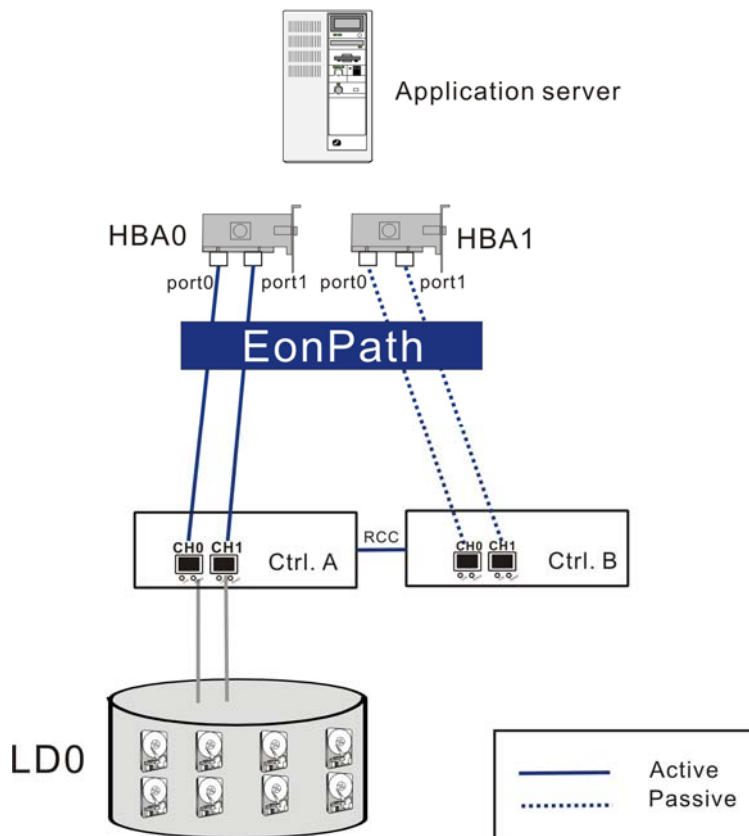
In the dual-controller configuration, when users configure a logical drive (LD), SANWatch will require them to assign the LD to either Controller A or Controller B<sup>2</sup>. According to the assignment, EonPath would recognize which data paths are active and which are passive. For example, in the figure below, the host has four data paths through which it can distribute I/O to LD0: HBA0/port0 → Controller A/CH0 → LD0; HBA0/port1 → Controller A/CH1 → LD0; HBA1/port0 → Controller B/CH0 → RCC<sup>3</sup> → Controller A → LD0; HBA1/port1 → Controller B/CH1 → RCC → Controller A → LD0.

---

<sup>1</sup> In-band hosts refer to the servers communicating with RAID subsystems over the Fibre, SAS or iSCSI data paths.

<sup>2</sup> User can still change the assignment after the LD is created. For the detailed instructions, please refer to SANWatch User's Manual.

<sup>3</sup> RCC means the communication path between the two controllers.



Since LD0 is assigned to Controller A, the two data paths connected to Controller A will be recognized as active, while the two connected to Controller B as passive. The passive paths will automatically take over the responsibility of handling the data flow when active paths are all broken. EonPath GUI represents the attributes of data paths in an easy-to-understand way, whether they are **Active**, **Passive**, **Not Used**, or **Failed**. With as few as two clicks, users can change the attributes of data paths (**Active to Passive**, **Passive to Active**, or **Not Used to Active**). After users make the changes, EonPath allows them to make sure that I/Os are really distributed through the desirable paths by checking the throughput graph of each channel.

### Load-balancing Between Data Paths

When multiple active paths are available between a host and a LD, EonPath can redistribute read/write requests among these paths to increase system throughput performance. This act is called load-balancing. EonPath provides three load-balancing policies: **Mini Queue**, **Round Robin** and **Failover**. Selecting **Mini Queue** means that EonPath would check the number of I/O requests in each path and dynamically distribute the new I/Os to the path with lower loading (fewer requests in the job queue). This policy can reduce the possibility that a single path is congested with an overload of I/O requests. If users select **Round Robin**, I/Os will be equally distributed among paths in a round-robin fashion, regardless of their loadings. When the data flow is not

so heavy that it exceeds the bandwidth of one data path, users can select **Failover** to prevent the system from wasting computing power on distributing I/Os among multiple paths. By selecting **Failover**, host I/Os will go through only one active data path and therefore system does not have to make any load-balancing efforts. When the active data path serving as I/O channel fails (cabling or HBA failure), another active data path will take over the role as the I/O channel.