



# **Supported iSCSI Initiators**

## **Application Notes**

**July 2003**

All trademark names are the property of their respective companies. This publication contains opinions of StoneFly Networks, which are subject to change from time to time.

This publication is copyright © 2003 by StoneFly Networks, Inc. and is intended for use only by recipients authorized by StoneFly Networks. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of StoneFly Networks, Inc., is in violation of U.S. copyright law.

## Table of Contents

Introduction.....	<b>3</b>
Initiator Types .....	<b>3</b>
SNICs and HBAs .....	3
Software Driver and NIC .....	3
TOE.....	4
iSCSI Initiators Supported by StoneFly Networks.....	<b>4</b>
Intel PRO/1000 T IP Storage Adapter .....	4
Microsoft iSCSI Software Driver.....	5
Adaptec ASA-7211C .....	5
Linux Open Source (Source Forge).....	5
IBM Software Initiator.....	6
Novell NetWare iSCSI Initiator .....	6
Novell iSCSI Initiator Server Requirements .....	7
Installing the Novell iSCSI Initiator Software .....	7
Using Server Console Commands.....	7
Using NetWare Remote Manager.....	7
QLogic QLA4000 for Solaris (Prerelease Only).....	9
NIC FailOver with Adaptive Load Balancing .....	<b>10</b>
Installing Intel® PROSet Software .....	11

## Introduction

The StoneFly Networks™ Storage Concentrator™ is the interface between hosts and storage devices in an IP network. IP-based storage area networks (IP SANs) use the iSCSI protocol over an Ethernet and TCP/IP network. Each server, host, or user that wants to access the storage devices must also implement the iSCSI protocol over an Ethernet and TCP/IP network. The iSCSI initiator is the component of a server that “initiates” or originates storage commands.

This application note describes how to implement iSCSI initiators supported by StoneFly Networks Storage Concentrator Version 1.7. StoneFly Networks will actively certify and support additional iSCSI initiators as they come to market. For the current list of supported initiators, please check the StoneFly Networks Web site at:

[http://www.stonefly.com/Products/stonefly\\_compat\\_matrix.htm](http://www.stonefly.com/Products/stonefly_compat_matrix.htm)

The current implementation of the Storage Concentrator supports both iSCSI specification Draft Version 0.8 and Draft Version 0.19 & 0.20 (RFC Version 1.0).

## Initiator Types

There are three ways to set up iSCSI initiators to work with the StoneFly Networks Storage Concentrator.

Note: These are general guidelines; specific implementations may vary in configuration, performance, and functionality.

### SNICs and HBAs

A storage network interface card (SNIC) or iSCSI adapter interfaces with the host operating system (OS) in the same way as a storage Host Bus Adapter (HBA). Some vendors call their cards an **iSCSI HBA**.

Once a SNIC or iSCSI adapter is installed into a server, the operating system will present disk drives that appear to be additional SCSI devices. A SNIC has good throughput, low CPU utilization, and a gigabit Ethernet connection on board that offloads both TCP/IP and iSCSI processing. This solution adds some cost and requires an available PCI slot.

In order for a SNIC or iSCSI adapter to be operational, the card must have its IP address and other settings configured using a driver management software program.

### Software Driver and NIC

Software-based iSCSI initiators with a Network Interface Connector (NIC) connect to the OS storage stack and present disk drives in much the same way as a SNIC.

With an iSCSI software initiator and NIC adapter, the storage traffic is processed and redirected to the normal networking interface where TCP/IP processing is implemented.

This solution normally has higher CPU utilization than the SNIC approach. Benchmark tests show throughput similar to a SNIC when there is no other processing on the server and plenty of CPU cycles are available. In actual use, throughput will not be as good or the other applications or database programs will respond more slowly.

The advantage of the software iSCSI initiator is that the server hardware is not shut down to install an adapter card; only a software driver is added to the OS. This type of driver is normally used in conjunction with a server's Gigabit Ethernet connection. A 10/100 connection is sufficient for many applications.

## **TOE**

A TCP Offload Engine (TOE) adapter card offloads the compute-intensive processing associated with maintaining a TCP/IP connection.

The TOE card reduces CPU utilization for all network communications. When used with an iSCSI software initiator, the performance can be similar to a SNIC. Normally there is a small amount of extra overhead getting data from the storage stack to the communications stack.

The advantage is that the financial investment of implementing a TOE will improve performance on both the storage and network connections. An add-on card and PCI slot are required in addition to the iSCSI driver. The TOE card normally includes the iSCSI driver and the network adapter card driver.

## **iSCSI Initiators Supported by StoneFly Networks**

For the most current list of initiators, please check the StoneFly Networks Web site at:

[http://www.stonefly.com/Products/stonefly\\_compat\\_matrix.htm](http://www.stonefly.com/Products/stonefly_compat_matrix.htm)

Refer to the release notes that came with your driver for critical information regarding patch levels that are required for your operating system.

The following are release notes for using the following iSCSI initiators with the StoneFly Networks Storage Concentrator:

### **Intel PRO/1000 T IP Storage Adapter**

The Intel PRO/1000 T IP Storage Adapter is available from StoneFly Networks, Intel or its distributors. The drivers and related software are available from Intel's Web site at:

<http://www.intel.com/support/network/adapter/iscsi/tip/index.htm>

Hints for using the Intel PRO/1000 T IP Storage Adapter iSCSI initiator with the StoneFly Networks Storage Concentrator:

1. **IMPORTANT:** The current version of the Intel PRO/1000 T storage adapter is PWSA8701T. This Version supports both 32- and 64-bit PCI slots. The prior version (PWSA8700T) only supported 64-bit PCI slots. For more information, see your Intel PRO/1000 T documentation.

2. The Intel Storage Adapter iSCSI initiator performs discovery for its targets by issuing a "send targets" command. If the Storage Concentrator has been preconfigured to assign volumes to the host, the Intel Storage Adapter iSCSI initiator will discover the volumes. If no volumes are discovered, the Intel Storage Adapter iSCSI initiator will not automatically reissue its discovery command. To generate another discovery request, after the volume access lists are configured on the Storage Concentrator, remove the iSCSI target from the Intel Management GUI and then add it again.

## Microsoft iSCSI Software Driver

The Microsoft iSCSI driver is now available directly from Microsoft. You can use the driver with a standard NIC. Some TOE cards also support the Microsoft driver. The driver can be found at:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=12cb3c1a-15d6-4585-b385-befd1319f825&DisplayLang=en>

## Adaptec ASA-7211C

The Adaptec ASA-7211C is a gigabit Ethernet iSCSI HBA. StoneFly has tested this card for both Windows 2000 and Windows NT. The card is currently available from Computer Discount Warehouse (CDW). Internet users can find it at:

<http://www.cdw.com/shop/products/default.asp?EDC=470875>

## Linux Open Source (Source Forge)

Sources for the Open Source iSCSI initiator (from Source Forge) can be found at <http://sourceforge.net/projects/linux-iscsi>.

StoneFly has fully tested the following Versions:

- For iSCSI Version 0.8, the correct version is 2.1.0.20 for Red Hat Versions 7.1, 7.2, or 7.3
- For iSCSI Version 0.20, the correct version is 3.1.04 for Red Hat Version 8.0.

Contact StoneFly Networks Customer Care for additional information on how to obtain or create executable versions of this iSCSI initiator. StoneFly Networks tests the Storage Concentrator with new revisions of the source as they become available from Source Forge. For the latest information on supported versions, please check StoneFly Networks' Web site at [www.stonefly.com/support](http://www.stonefly.com/support).

Hints for using the Open Source iSCSI initiator with the StoneFly Networks Storage Concentrator:

- The Open Source iSCSI initiator is available in source code form. It must be compiled with the specific version of the Linux kernel on the host system. Newer versions of the iSCSI initiator and the Linux kernel may be available at installation time. For information on the latest supported combinations, go to the Knowledge Base section of StoneFly Networks' Web site at <http://www.stonefly.com> or call StoneFly Networks Customer Care.

- The /etc/iscsi.conf file requests an entry for each target and LUN to be used by the iSCSI initiator. This provides easy control of the target/LUN list. However, the iSCSI initiator accepts all target/LUN combinations presented to it, so to gain ultimate flexibility with the Storage Concentrator do not specify any target or LUN combinations. The initiator will then allow its targets to be specified only at the Storage Concentrator Graphical User Interface.

## IBM Software Initiator

IBM provides a software initiator that works with standard NIC cards for Windows and Linux. StoneFly supports the Windows version. Currently, this software initiator supports iSCSI Draft Version 0.8 only.

The Windows user documentation can be found at:

<http://www-1.ibm.com/support/manager.wss?rs=513&rt=0&org=ssg&doc=S7000058>

The Windows executable can be found at:

<http://www-1.ibm.com/support/manager.wss?rs=513&rt=0&org=ssg&doc=S4000059>

Hints for using the IBM software iSCSI initiator with the StoneFly Networks Storage Concentrator:

- The IBM iSCSI initiator performs discovery by using the SLP protocol. However, the Storage Concentrator responds only to a "send targets" request. As a result, the iSCSI name of targets must be entered manually at the server. The release documents from IBM describe how to set the iSCSI target names in the iSCSI initiator. All driver configuration is manual.
- Since the Storage Concentrator does not connect with the IBM iSCSI initiator for a discovery login, the iSCSI identification of the IBM initiator must be added at the Storage Concentrator's management console. For information on adding new hosts, see the *Storage Concentrator User's Guide*.
- Port number 3261 must be used for the iSCSI listening port. For more information on configuring the iSCSI listening port, see the *Storage Concentrator User's Guide*. Typically, the user interface for the IBM iSCSI initiator is pre-set to use port 5003.

## Novell NetWare iSCSI Initiator

Novell's iSCSI initiator stack for NetWare 5.1 and NetWare 6 is now available. This is iSCSI Draft 20 compliant. The driver is available at:

<http://download.novell.com/filedist/pages/PublicSearch.jsp>

To locate the driver, first select Netware as the platform. Then select the radio button for Search by Keyword. Then enter "iSCSI" as the search parameter. The software driver, set-up instructions and a link to the documentation are all available at this site.

**Note:** Novell 6.5 will include the iSCSI driver as part of the OS distribution, which will be on CD. Novell 6.5 is targeted for General Availability on Aug. 15, 2003.

### Novell iSCSI Initiator Server Requirements

- NetWare 5.1 requires Service Pack 5 installed on all servers that will run iSCSI initiator software.
- NetWare 6 requires Service Pack 3 installed on all servers that will run iSCSI initiator software.
- Netware 6.5 is in Novell Beta Status---contact Novell for information.

### Installing the Novell iSCSI Initiator Software

Download the iSCSI initiator package (.zip and .ncf files) and copy it to the root of the sys: volume on each server that will be used to access shared volumes connected to a Storage Concentrator. Extract the files from the download by entering INSTINIT at the server console. When you extract the files from the iSCSI Initiator download, they are automatically copied to the appropriate directories on your NetWare server.

**IMPORTANT:** *If you intend to install Novell Cluster Services software, you must do so after installing and configuring the iSCSI initiator software and before partitioning the disks on the shared disk system*

### Configuring the Novell iSCSI Initiator

iSCSI initiator software can be configured either at the server console using server console commands or remotely using NetWare Remote Manager.

#### Using Server Console Commands

For each server that you want to function as an iSCSI initiator, do the following:

1. Enter **ion** at the server console to load iSCSI initiator software.  
You can also enter **ioff** at the server console to unload iSCSI initiator software.
2. Enter **icsinit connect a.b.c.d target\_name** at the server console.  
Replace *a.b.c.d* with the IP address of the iSCSI storage router that is connected to the shared storage system.  
Replace *target\_name* with the iSCSI target name that is displayed after running the **icsinit discover a.b.c.d** command. The iSCSI target name is case sensitive.

**Note:** At this point, because the Concentrator does not use an authentication procedure on login, you will have to go to the management console for the Storage Concentrator and give the host ACL rights to the volume. When you have done that, repeat Step 2.

3. Enter **list devices** at the server console to have the server scan for and find the shared storage disks.  
The disks on the shared storage device will be listed along with the other devices that are directly connected to the server.

If you want iSCSI initiator software to load automatically when servers start, you can add the commands in the above steps to the AUTOEXEC.NCF file of each initiator server.

#### Using NetWare Remote Manager

For each server that you want to function as an iSCSI initiator do the following:

1. Run **ion** on the server to load iSCSI initiator software.

You can do this either at the server console or remotely by using NetWare Remote Manager to access the server console.

2. On the NetWare Remote Manager main screen, click the iSCSI Services link at the bottom of the left column.
3. Click Add Target and enter the IP address of the iSCSI target device (storage router) that is connected to the shared storage system.

Each target device can have multiple targets.

If you want a list of possible target names for a given IP address, click Find Target and enter the IP address of the target device.

**Note:** At this point, because the Concentrator does not use an authentication procedure on login, you will have to go to the management console for the Storage Concentrator and give the host ACL rights to the volume. When you have done that, repeat Step 2.

4. Click Next, select the target name you want to establish a session with, then click Next.
5. Enter **list devices** at the server console to have the server scan for and find the shared storage disks.

This can also be done remotely by using NetWare Remote Manager to access the server console.

The SCSI disks on the shared storage system will be listed along with the other devices that are directly connected to the server.

### Viewing and Editing Initiator Properties

You can view iSCSI initiator properties at the initiator server console or by using the NetWare Remote Manager. To change iSCSI initiator and driver properties, you must use the NetWare Remote Manager.

To view iSCSI initiator properties at the initiator server console, enter **iscsinit info**.

To view or change iSCSI initiator properties using NetWare Remote Manager:

1. On the NetWare Remote Manager main screen, click the iSCSI Services link at the bottom of the left column.
2. Click the Properties link to bring up a page that lets you view or change initiator and driver properties.
3. View or change the desired properties, then click Finish to save changes.

Current initiator and driver properties include the following:

- Authentication Method
- Front Page Display Controls
- Number of LUN Probes per Target
- Display Driver Statistics

### Release Notes for Novell NetWare iSCSI Initiator Beta 2

- 1) AUTH\_NONE is the only authorization method supported.
- 2) After executing the List Devices command at the server console, you might experience a delay before the device list is displayed.
- 3) A server failure can occur if the server is running a version of NetWare 6 earlier than NetWare 6 SP3 and you enter IOFF at the server console with active sessions running. To avoid the failure, first dismount iSCSI-based volumes, then deactivate iSCSI based pools.

4) Occasional slow performance and/or dropped iSCSI connection problems have been experienced with certain router, switch, and hub configurations.

These network configurations may require that SACK (Selective ACK) be turned off for the TCP/IP stack running on the initiator. This can be done at the NetWare server console by entering: set tcp sack option=off

## **QLogic QLA4000 for Solaris (Prerelease Only)**

StoneFly currently supports the Solaris version of the QLogic QLA4000 for iSCSI Version Draft 20. The QLogic QLA4000 is available from QLogic or its distributors. Drivers for the QLA4000 for Solaris can be found at:

[http://www.qlogic.com/support/os\\_detail.asp?productid=320&osid=22](http://www.qlogic.com/support/os_detail.asp?productid=320&osid=22)

**Configuring volumes:** Once the QLA4000 Solaris iSCSI driver has been installed on the Solaris host system run the "SANblade Control iX" utility to configure the HBA.

First, Select the HBA Options tab then Network to configure the HBA's iSCSI Name, IP Address, Subnet Mask and Gateway. Enter all the values.

Next select Target Settings tab to configure targets. Use the "+" button to add a new target then fill in target IP Address, Port and iSCSI Target Name. Note if Target Name is left blank the HBA will perform a Discovery Session then login to all targets returned for this HBA.

On the Storage Concentrator's GUI select the "Sessions" screen to see a separate line for each volume assigned to the Solaris host.

In order for the Solaris Host to use the iSCSI volumes from the Storage Concentrator several steps must be taken to format the volume properly.

The Solaris Host does not have a saved definition of a Storage Concentrator logical volume. With the current release of the Storage Concentrator software drive geometry for the iSCSI logical volume must be given to the Solaris operating system through the format command.

Therefore, you must enter the drive geometry information that is equivalent to the StoneFly logical volume. If the values given for "cylinders, heads, sectors" is less than the actual size of the logical volume you would not have full use of the volume. If the values for "cylinders, heads, sectors" is greater than actual, you would be attempting to create a larger volume than actual, which will fail.

Please enter the following information at the command prompt:

1. # devfsadm
  - This identifies the iSCSI volumes for the operating system.
2. # format
  - Lists the known drives (including iSCSI volumes). First, identify the drive to format by selecting the first iSCSI drive and entering it's number.
  - Next the system will ask you to identify the type of drive you have selected. Enter the number for "other."

- The format command will then request information on geometry for the "other" drive you have selected.
    - Enter the number of data cylinders =  $(512 * n)$  - where n = number of GbE for the volume
    - Enter the default for "alternate data cylinders" = 2
    - Enter the number of heads = 64
    - Enter the number of sectors = 64
    - Enter the drive rpm = slowest drive in volumes rpm (enter 7200)
    - Enter the defaults for all other values
  - Next partition the drive.
  - Next format the drive. (Note that this can't be interrupted and will take quite a while, depending on the size of the volume ---for example, a 6 GB drive will take approximately 40 minutes to format)
  - The format command changes to its command line functions at this point. Enter: > label. Answer "y" to confirm the label should be written now. Then type > quit
3. The volume will now be ready to mount and use.
  4. Repeat this process for each iSCSI volume. After all the volumes are labeled, a file system may be placed on the drives and used in applications.

## NIC FailOver with Adaptive Load Balancing

Designing a fault tolerant storage network requires that there is no single point of failure in the network. Many network storage experts recommend that host servers feature more than one path to the storage. To this end, StoneFly has tested the Intel PRO/1000 Network Interface Cards configured as a "team" running adaptive load balancing (ALB).

Intel describes a team as a group of adapters configured to enhance network functionality, including failover, increased transmission throughput or fault tolerance. Intel uses its PROSet software to support teaming with Windows 2000, Windows NT 4.0, and Windows XP.

The multiple adapters in an ALB team provide emergency backup protection, as well as increased bandwidth. If one server link goes down – due to a broken cable, a bad switch port, or a failed adapter – the other adapters automatically accept the additional load. There's no interruption in server operation, and a network alert is generated to inform IT staff of the problem.

For iSCSI (IP) SANs, an iSCSI software initiator that supports ALB is also required. StoneFly has tested initiators from Microsoft and IBM that both work with two Intel NIC cards and ALB enabled with the Intel PROSet software. Intel also has *Advanced Network Services* (ANS) package, version 2.3.x, for Linux systems that supports teaming. However, StoneFly Networks has not tested other configurations and they are not certified to work with the StoneFly IP SAN.

To run ALB, each server adapter must be connected to a switch. The adapters can be connected to different switches, as long as all the switches are on the same network segment (they can't be separated by a router). All outgoing server traffic will be balanced across the adapter team. A single adapter carries incoming traffic. In most environments, this is a highly effective solution, since server traffic is primarily outbound – from the server to the clients. Adapter Fault Tolerance is a team of two or more adapters configured to

provide a backup network connection. In the case of hub or switch port, cable, or adapter failure, the computer maintains uninterrupted network performance.

## Installing Intel® PROSet Software

Intel PROSet enables you to test and configure Intel network adapters. You can install Intel PROSet on computers running Windows 2000, Windows XP, and Windows Server 2003. You must have administrative rights on the computer where you are installing the software.

When you run the Installer, it checks to see if a previous version of the drivers or Intel PROSet exists on your computer. If a previous version exists, it is upgraded to the latest version. Any settings specified in the previous version are applied to the version.

The default setting is for auto-negotiation to be enabled. Change this setting only to match a specific link partner.

After installing Intel PROSet and restarting the computer, two icons appear, as shown below. Double-click either icon to run Intel PROSet.

System Tray Icon:



Control Panel Icon:



**NOTE 1:** The adapter will not operate at gigabit speed (1000 Mbps) if you disable auto-negotiation.

**NOTE 2:** You cannot change the speed or duplex of Intel PRO/1000 adapters that use optical fiber cabling.

If Auto Negotiation is disabled, you can specify Forced Speed (to 10 or 100 Mbps, but not 1000 Mbps) and Duplex. In this situation, the Negotiable Speeds and Duplexes checkboxes are disabled, since these are set to advertise speeds during auto-negotiation.

If Auto Negotiation is *enabled*, the Negotiable Speeds and Duplexes checkboxes are enabled, allowing specific speed advertisement and disabling the Forced Speed and Duplex selection buttons.

If you need to manually change the settings:

1. Double-click the Intel PROSet icon from the Control Panel or System Tray.
2. Click the Speed tab and highlight the desired adapter in the left window.
3. Press the F1 key. More specific instructions are available in the Intel PROSet Help window.

Next, setup the team.

1. In the Network Component Tree, right-click a non-teamed adapter.
2. Click Add to Team, then Create New Team.  
The Teaming Wizard appears.
3. Select Adaptive Load Balancing and click Next.
4. If you are creating an Adaptive Load Balancing team, you can choose to either enable or disable Receive Load Balancing (RLB).  
Note: If you enable RLB, team members capable of the highest speed receive IP traffic, and the primary adapter receives non-IP traffic. With RLB enabled, Intel(R) PROSet II automatically selects a preferred primary adapter. Intel PROSet II's selection of a preferred primary adapter is not based on adapter model or speed. If you disable RLB, only the primary adapter receives traffic.
5. The list displays all adapters that are not currently part of a team. From the list, click the adapters you want to include in this team, then click Next.
6. Click Finish.
7. If you choose, you can specify an adapter in the team as the primary adapter. To specify a primary adapter, right-click the adapter you would like to use as the primary then click Preferred Primary.  
If you choose not to select a primary adapter, Intel PROSet II selects a preferred primary adapter.
8. Click Close to close Intel PROSet II and apply changes.

The team is now configured.

Because the Intel PRO/1000 Network Interface Card already supports high-bandwidth, using the ALB feature might cause reduced performance on some servers. This is because many servers available today do not have the CPU or Bus capacity to effectively handle teamed gigabit-class adapters. For best ALB performance, use newer servers that support multiple PCI buses and multiple processors running at a speed of 400 MHz or more. On such a server, install the teamed Intel PRO/1000 adapters on separate PCI buses for maximum ALB performance.