

SANRAD Application Note: Multi Path I/O for iSCSI

APP-017-01

Important Notice

The configurations described or tested in this application note are not the only available solution. This application note is not intended nor may it be construed as an endorsement of any product(s) tested.

This application note provides no warranty of any kind. SANRAD limited warranties for SANRAD products are stated in separate documentation accompanying each product. No damages or remedy of any kind shall be recoverable by any party for any claim or loss in any way arising or alleged to arise from or as a result of this application note, whether considered alone or in addition to any other claim.

Copyright SANRAD 2006

All rights reserved. The copyright and all intellectual property rights in this article belong to SANRAD. It is strictly forbidden to copy, duplicate or otherwise use this article or any part thereof in any way shape or form without the prior written consent of SANRAD.

Table of Contents

Introduction	3
Multi Path I/O for iSCSI – Basic Concept	3
Multiple Sessions or Multiple Connections	4
Multi Path I/O Load Balance	4
Integrating Multi Path I/O into High Availability Environment.....	4
Configuring Multi Path I/O for Windows Servers	5
How to Configure Multi Sessions Connectivity	5
Requirements	5
Configuration	5
Verify.....	9
How to Configure Multi Connection Connectivity.....	12
Requirements	12
Configuration	12
Testing.....	19
Basic Troubleshooting	20
Summary	20

Figures

Figure 1. Basic Multi Path I/O for iSCSI	3
Figure 2. Multi-Path I/O for “No Single Point of Failure”	5
Figure 3. Configure Multi Sessions – iSCSI initiator Properties.....	6
Figure 4. Configure Multi Sessions – Log On to Target	6
Figure 5. Configure Multi Sessions – Log on Multi-Path Window.....	7
Figure 6. Configure Multi Sessions - Advance settings Window	8
Figure 7. Configure Multi Sessions - Advance settings Window	9
Figure 8. Verify - Advance settings Window.....	10
Figure 9. Verify – Target Properties Window.....	11
Figure 10. Verify – Device Details Window	12
Figure 11. Configure Multi Connections – iSCSI Initiator Properties Window	13
Figure 12. Configure Multi Connections – Log On Window.....	14
Figure 13. Configure Multi Connections - Advance settings Window	15
Figure 14. Configure Multi Connections – Target Properties Window.....	16
Figure 15. Configure Multi Connections – Session Connections Window.....	17
Figure 16. Configure Multi Connections – Add Connection Window	17
Figure 17. Configure Multi Connections - Advance settings Window	18
Figure 18. Configure Multi Connections – Session Connections Window.....	19

Introduction

When designing a system with high availability and no single point of failure it's necessary to consider several levels of redundancy. In order to protect the path all the way from the server to the storage the system should:

1. Have redundant paths from the server to the V-Switch Ethernet ports
2. Contain a cluster of V-Switches with failover capabilities
3. Have redundant paths from the V-Switch to the attached storage

This paper covers the redundancy between the host and the V-Switch (Item #1), using Multi Path I/O (MPIO) technologies. It explains the concept of MPIO and provides a "How to" guide for configuring MPIO in MS Windows environment.

For more information regarding High Availability solutions, refer to the following application notes:

1. <http://www.sanrad.com/objects/solutions/APP-15-01-NoSinglePointofFailure.pdf>
2. http://www.sanrad.com/objects/support/Windows_VSwitch_cluster-APP-008-02.pdf

Multi Path I/O for iSCSI – Basic Concept

Figure 1 shows a basic multi path I/O design.




Figure 1. Basic Multi Path I/O for iSCSI

A multi path I/O design for connectivity between the iSCSI initiator and the iSCSI portals (e.g. between the iSCSI initiator and portals on V-Switch) comprised of both:

- A physical layer (cables)
- A software layer.
The software layer is a “MPIO” driver which can be part of the iSCSI initiator or can be a separate add-on for an iSCSI HBA.

The first step in implementing any MPIO design is to connect the cables as required. The actual configuration of the MPIO is done either from the iSCSI initiator interface or from the MPIO interface (if it has one). Once configured, the MPIO driver is responsible for coordinating the IO traffic between the paths to the portals and to reroute it in case of a failure within one the paths. Most main stream operating systems (e.g. Microsoft, Versions of Linux and AIX) which have iSCSI initiator also support MPIO. Using SANRAD’s V-Switch, multi path I/O traffic management between the V-Switch and the disk array is performed automatically (including rerouting the I/O in case of a path failure) by the V-Switch.

Multiple Sessions or Multiple Connections

Usually, iSCSI targets are exposed from all iSCSI portals. In other words, all targets can be connected by logging in to any of the configured iSCSI portals.

- **Multi Sessions** connectivity occurs when you log in to the same target with different sessions as follows:
Login to a target via portal A and login again to target A via portal B.
- **Multi Connections** occurs when you have one iSCSI login session with multiple connections as follows:
Login to a target via portal A and then add a connection to target A via portal B.

The difference between the two is that with multi connections, both the target side and the initiator side are jointly managing the I/O traffic between the iSCSI initiator and the iSCSI target. Whereas with multi sessions, only the MPIO driver from the server side is responsible for managing the I/O traffic between the iSCSI initiator and the iSCSI target

SANRAD’s V-Switch supports both types of connectivity and it’s up to the user to decide which one to implement based on the iSCSI initiator capabilities (not all iSCSI initiators support both types of connectivity), type of HA design and performance tests.

Multi Path I/O Load Balance

Besides failover capabilities, an additional benefit of using Multi-Path I/O is load balancing. Multi-Path I/O routes traffics via the least busy path, enhancing I/O performance. Most MPIO drivers use the “Round Robin” policy for I/O load balancing.

Integrating Multi Path I/O into High Availability Environment

Multi Path I/O can be integrated into a high availability iSCSI storage environment ensuring “no single point of failure”. While there are many design options, Figure 2 shows one example integrating Multi Path I/O:

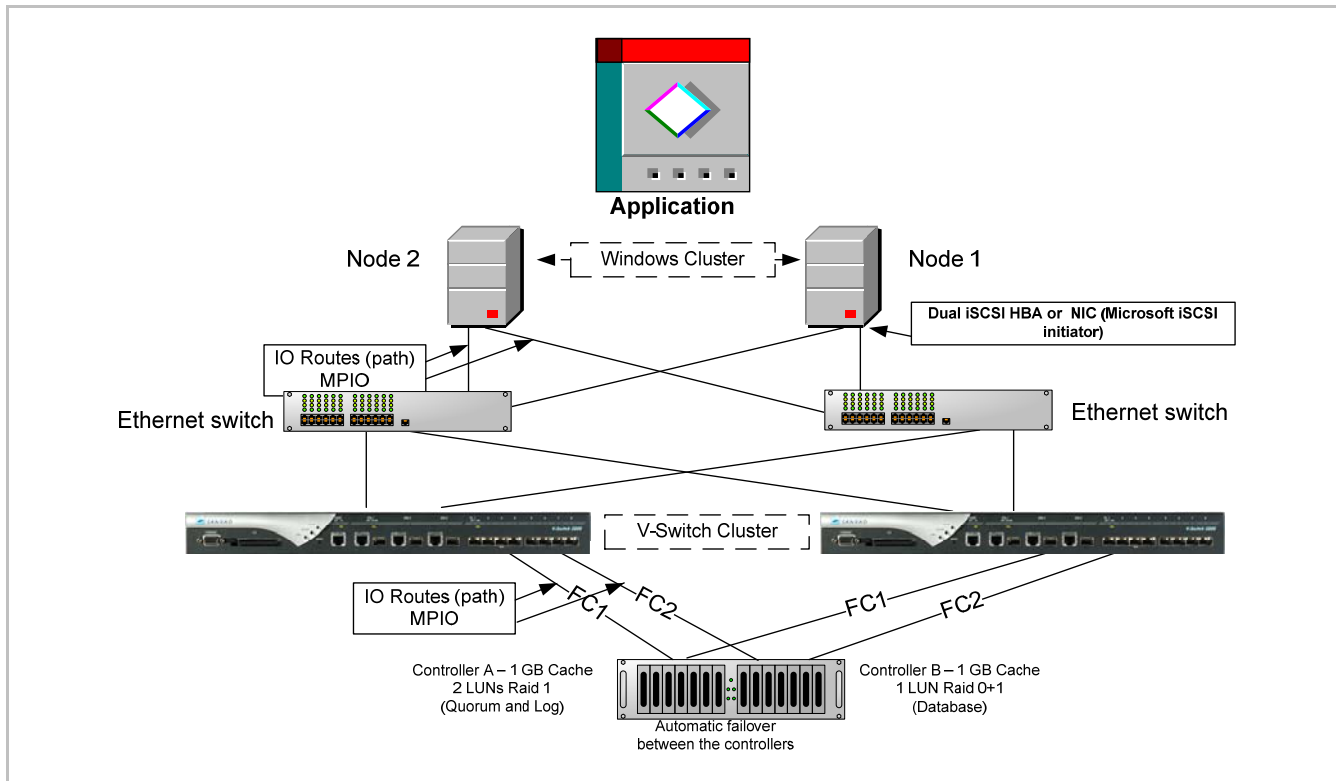


Figure 2. Multi-Path I/O for “No Single Point of Failure”

Configuring Multi Path I/O for Windows Servers

The following describes how to configure multi path I/O from Windows servers to SANRAD’s V-Switch.

How to Configure Multi Sessions Connectivity

Requirements

1. On the window server, must have two Network Cards each configured to a different subnet:
Subnet A for example 100.0.100.0
Subnet B for Example 100.0.102.0
2. The V-Switch must have two portals one on each of Ethernet port 1 (ETH1) and one on port 2 (ETH2 or ETH3) each one on a different subnet: One on subnet A and one subnet B (same as the above).
3. The Windows server must have the latest Microsoft iSCSI initiator (V2 or higher).
4. The V-Switch must have the latest firmware (V3.1 or higher).

Configuration

1. Open the Microsoft iSCSI initiator GUI and click the Discovery Tab and enter the IP’s of the two portals from subnet A and Subnet B.
2. Click the Target Tab (Figure 3).

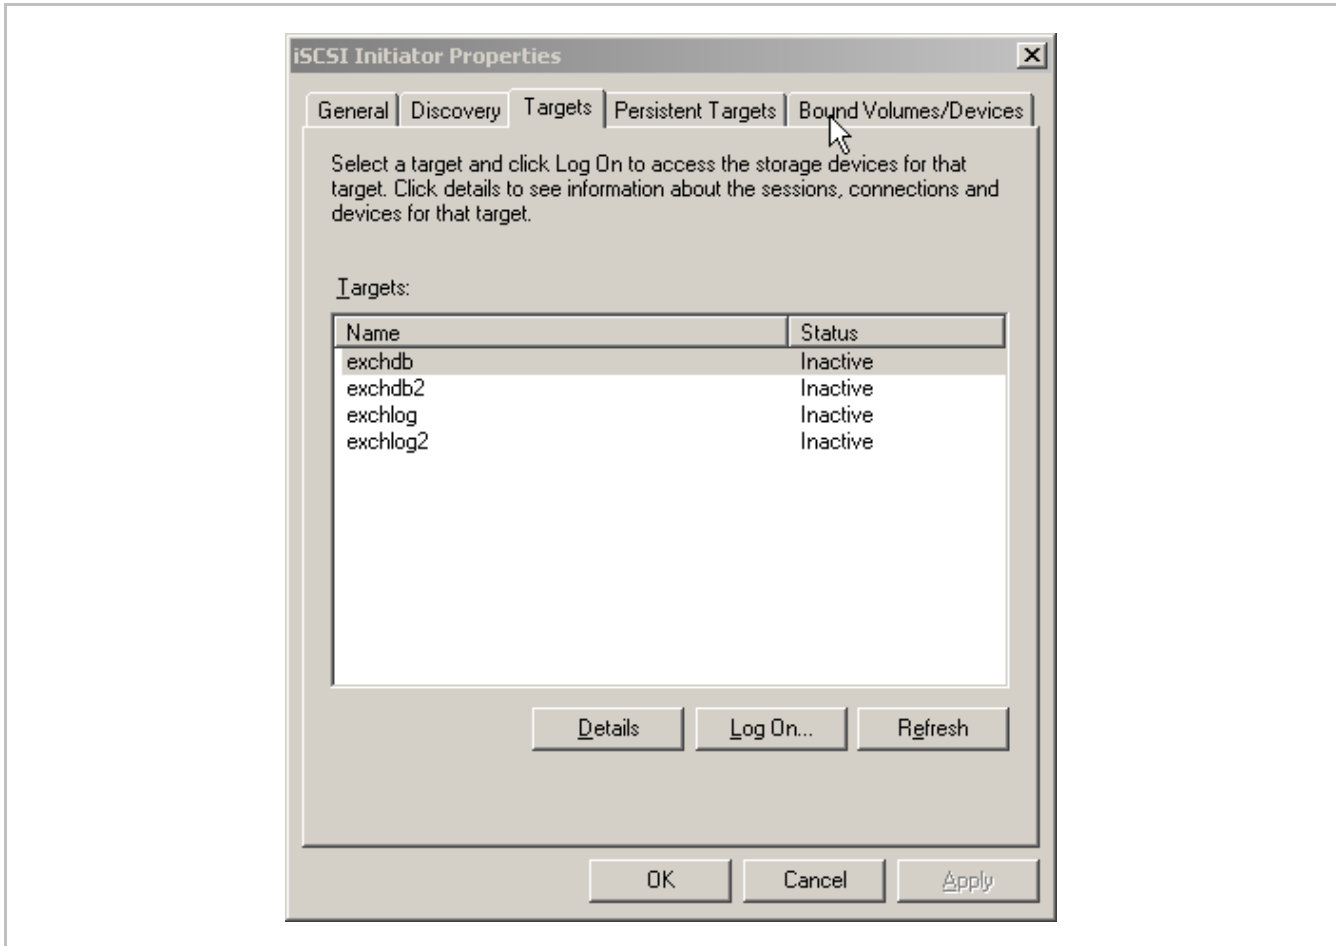


Figure 3. Configure Multi Sessions – iSCSI initiator Properties

3. Select the target and click [**Log On...**].
The Log On to Target Window opens (Figure 4).




Figure 4. Configure Multi Sessions – Log On to Target

4. Check **Enable multi-path** (Figure 4).




Figure 5. Configure Multi Sessions – Log on Multi-Path Window

5. Click [**Advanced...**].
In the Advanced setting window (Figure 6):
 - a. Change **Local Adapter** to **Microsoft iSCSI initiator**
 - b. Change **Target Portal** to the **portal IP in subnet A**.

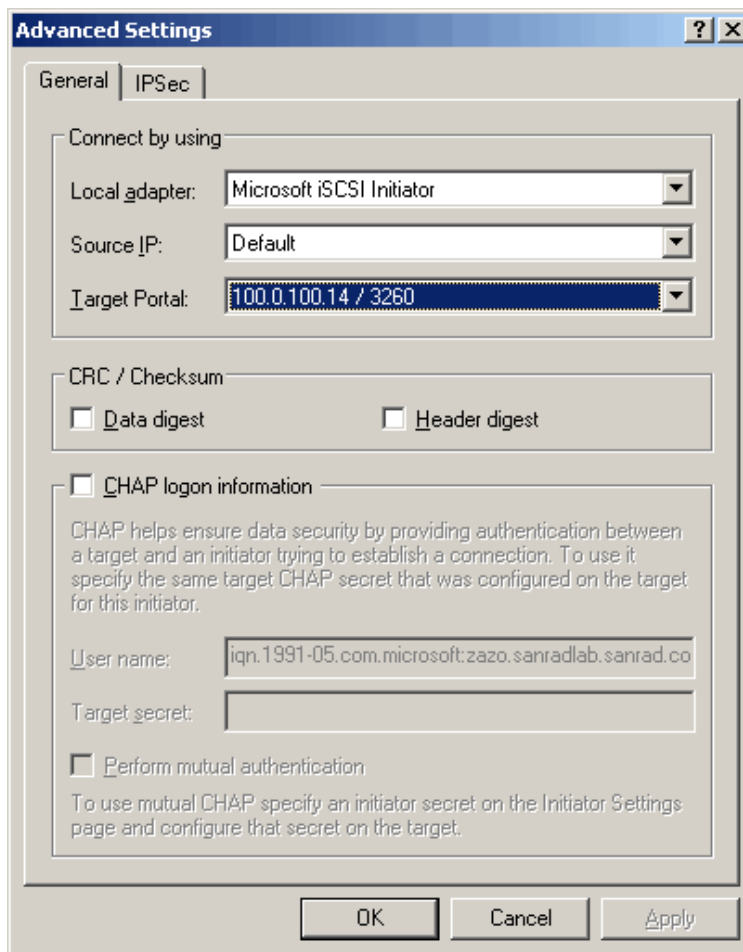


Figure 6. Configure Multi Sessions - Advance settings Window

6. Click OK (Figure 6).
7. Click OK (Figure 5).
8. In the Microsoft iSCSI initiator GUI (Target Tab - Figure 3) choose the same target again and click **[Log On...]**.
9. Check **Enable multi-path** (Figure 4).
10. Click **[Advanced...]**.
In the Advanced setting window (Figure 7):
 - a. Change Local Adapter to Microsoft iSCSI initiator
 - b. Change **Target Portal** to the **portal IP in subnet B**.

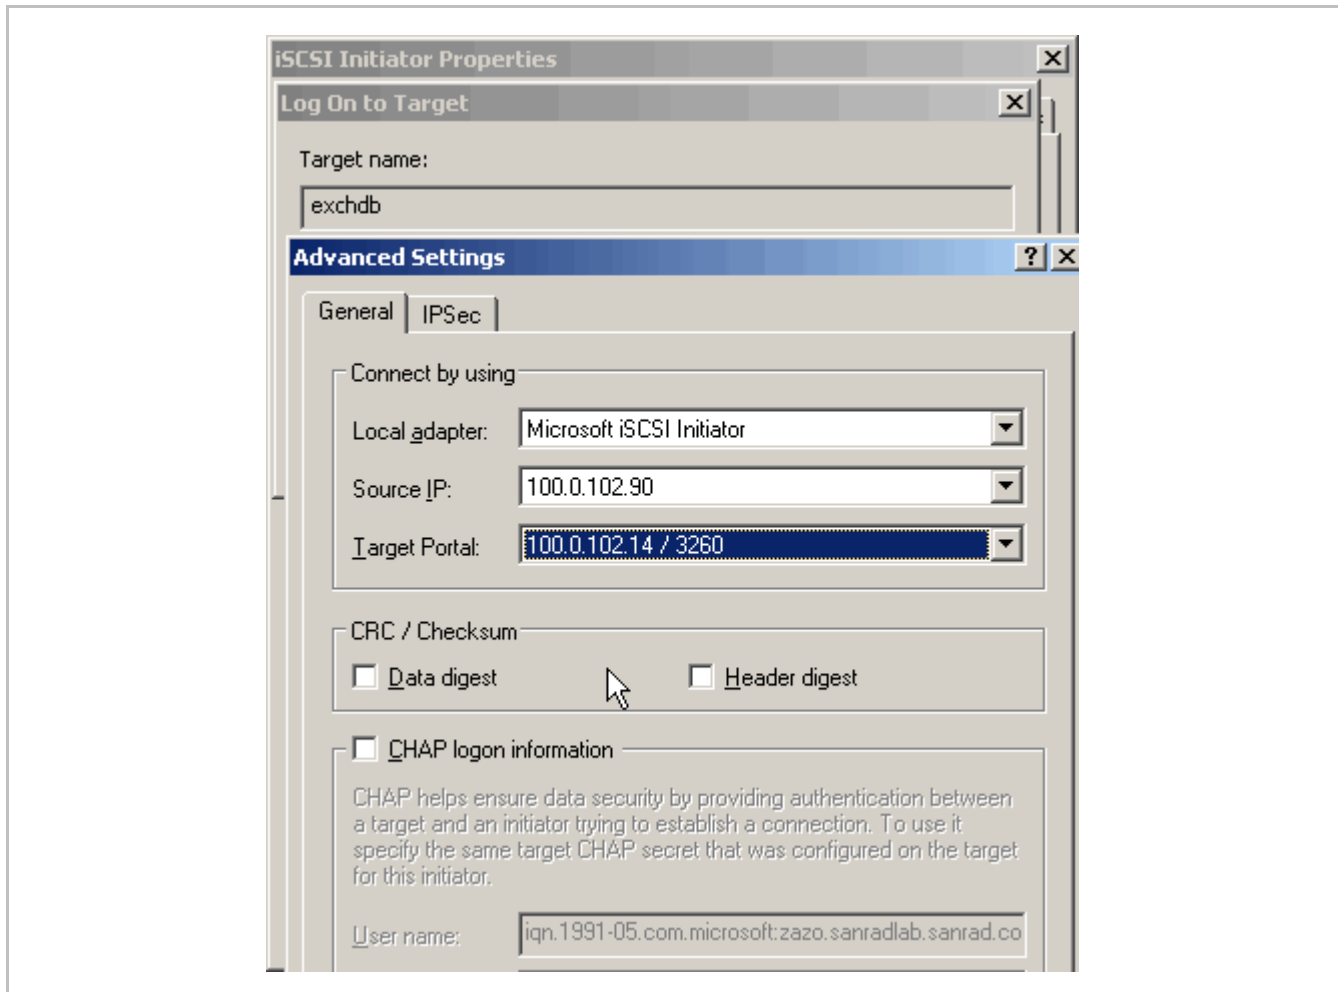


Figure 7. Configure Multi Sessions - Advance settings Window

11. Click OK (Figure 7).
12. Click OK (Figure 4).

Verify

1. Open the Microsoft iSCSI initiator GUI and click the Target Tab (Figure 3) and choose the target that was configured with MPIO.
2. Click [Details...]

The target properties window appears (Figure 8). You should see two sessions.

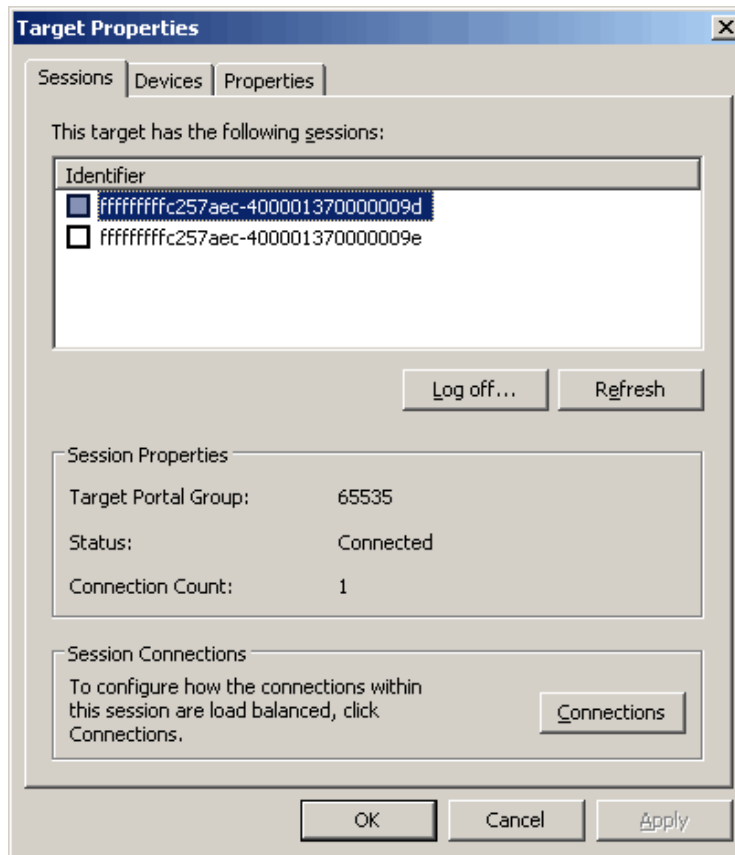


Figure 8. Verify - Advance settings Window

3. Click the Devices tab (Figure 9).
You should see two devices.

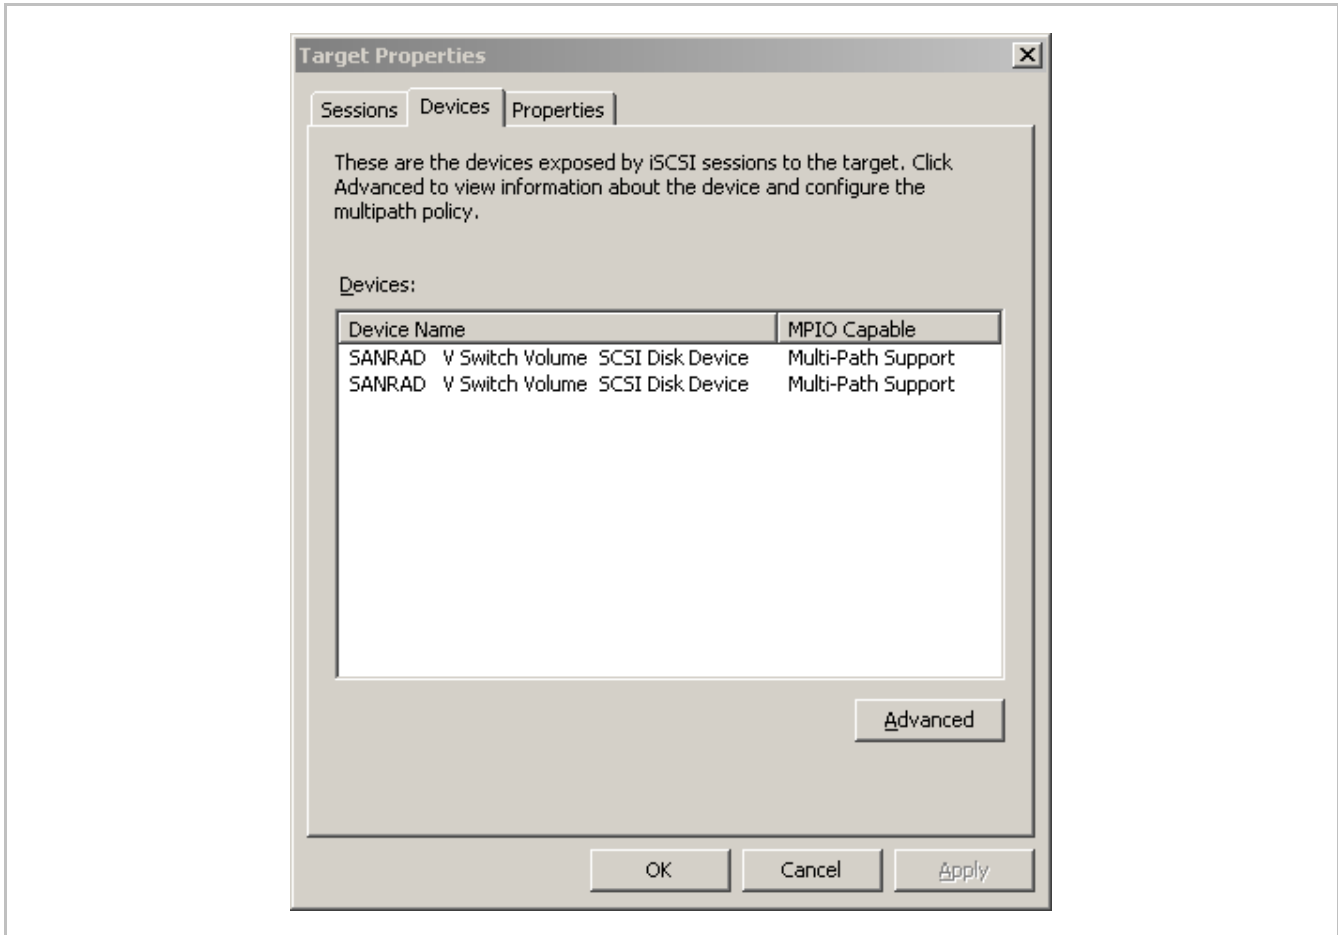


Figure 9. Verify – Target Properties Window

4. Choose one of the devices and click [**Advanced**].
The device details window will open (Figure 10).
5. Click the MPIO tab.
You should see two paths in the device details window.
6. Choose the **Load Balance Policy**:
“**Round Robin**” for I/O load balance traffic and fail over
or
“**Fail Over Only**”.

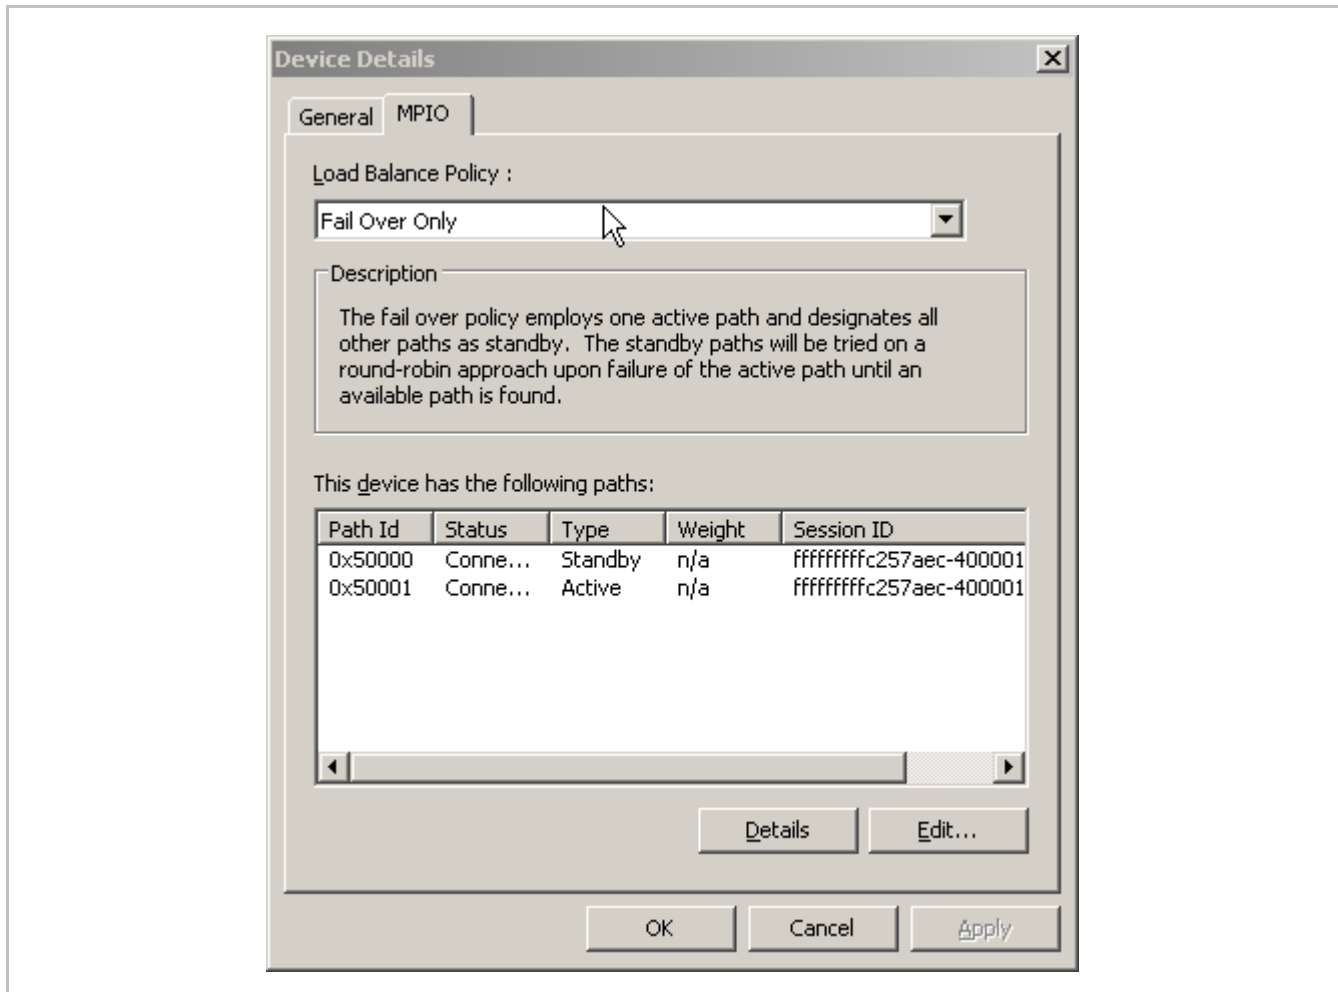


Figure 10. Verify – Device Details Window

How to Configure Multi Connection Connectivity

Requirements

1. The windows server must have two Network Cards each configured with a different subnet:
For example: Subnet A 100.0.100.0
For example: Subnet B 100.0.102.0
2. The V-Switch must have two portals one on each of Ethernet port 1 (ETH1) and one on port 2 (ETH2 or ETH3) each one on a different subnet: One on subnet A and one subnet B (same as the above).
3. The Windows server must have the latest Microsoft iSCSI initiator (V2 or higher).
4. The V-Switch must have the latest firmware (V3.1 or higher).

Configuration

1. Open the Microsoft iSCSI initiator GUI and Click the Discovery Tab. Enter the IP's of the two portals from subnet A and Subnet B.
2. Click the Target Tab (Figure 11).

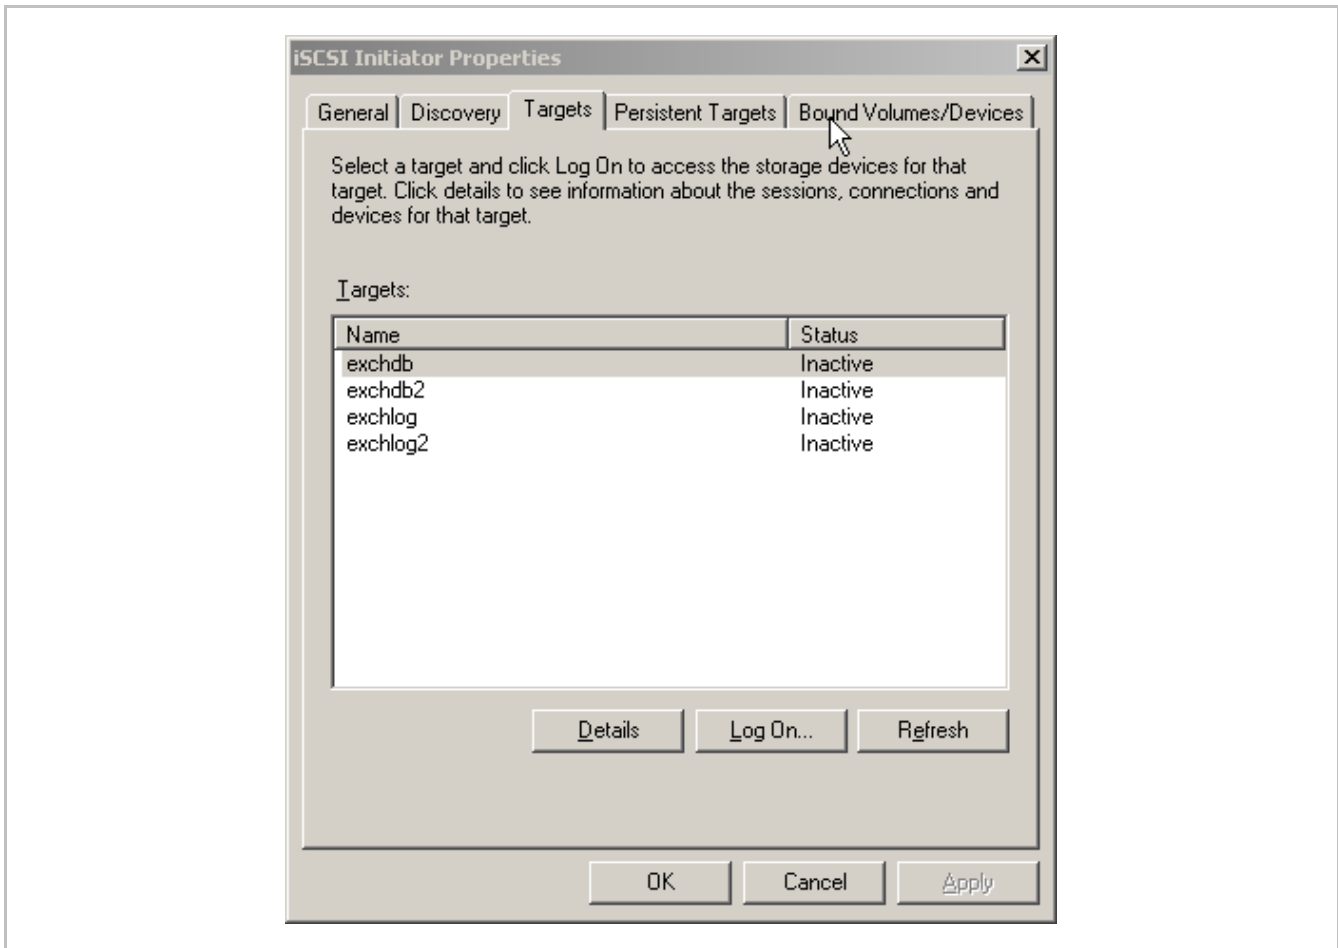


Figure 11. Configure Multi Connections – iSCSI Initiator Properties Window

3. Select the target and click [**Log On...**].
The Log On to Target Window appears (Figure 12).




Figure 12. Configure Multi Connections – Log On Window

4. Click [Advanced...].

In the Advanced setting window (Figure 13):

- a. Change Local Adapter to Microsoft iSCSI initiator**
- b. Change Target Portal to the portal IP in subnet A.**

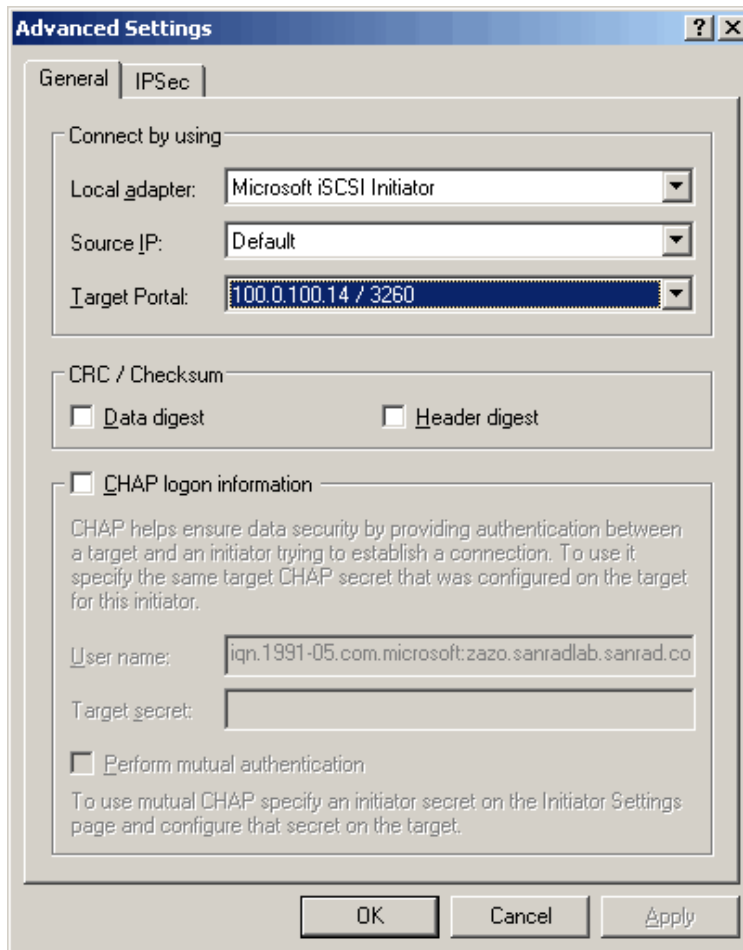


Figure 13. Configure Multi Connections - Advance settings Window

5. Click OK (Figure 13).
6. Click OK (Figure 12).
7. In the Microsoft iSCSI initiator GUI (Target Tab Figure 11) choose the same target and click **[Details]**. The target properties window appears (Figure 14).

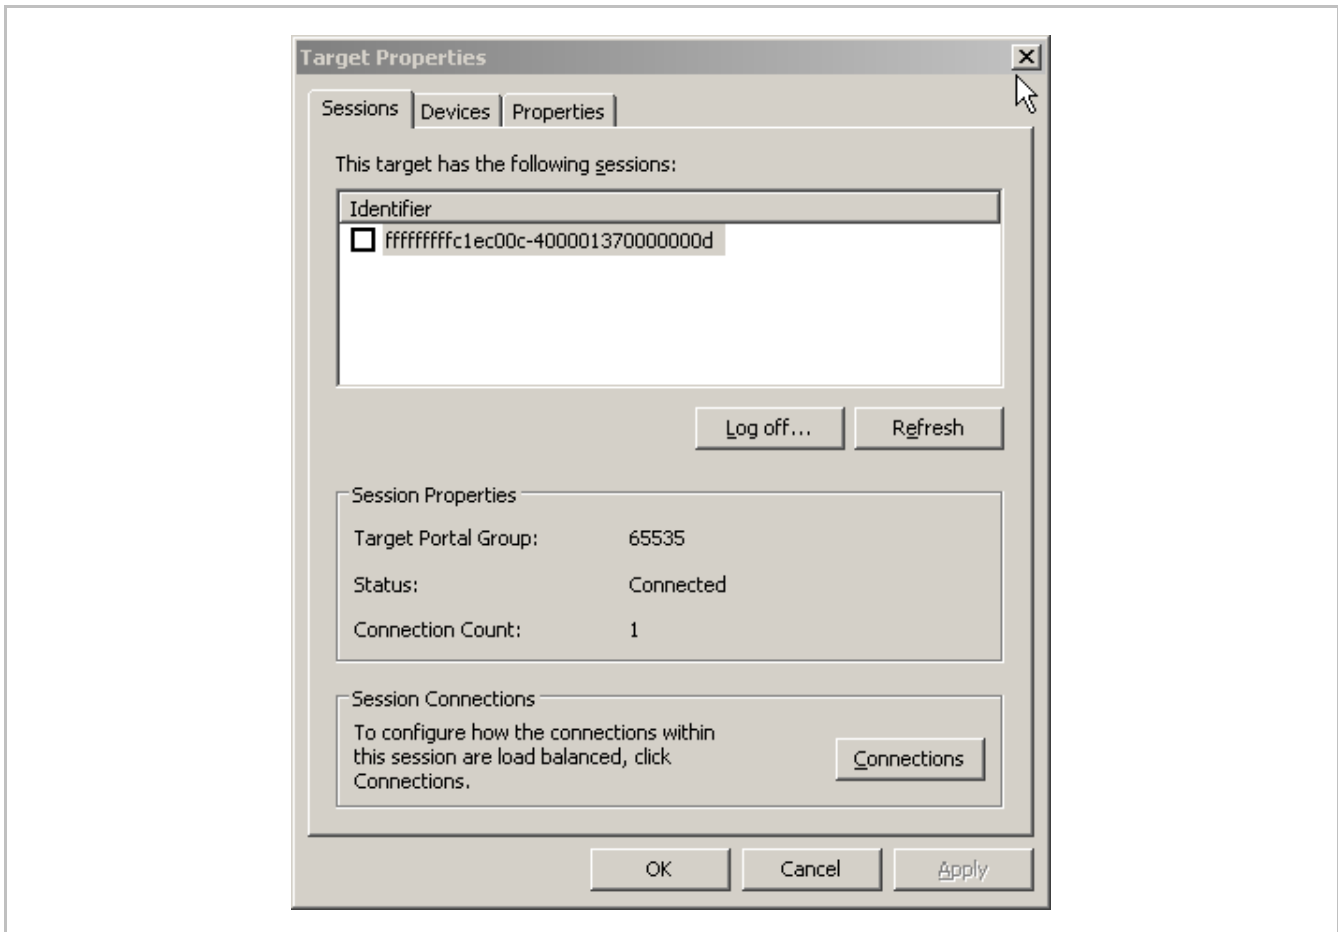


Figure 14. Configure Multi Connections – Target Properties Window

8. Click [Connections].
The session connection window appears (Figure 15).

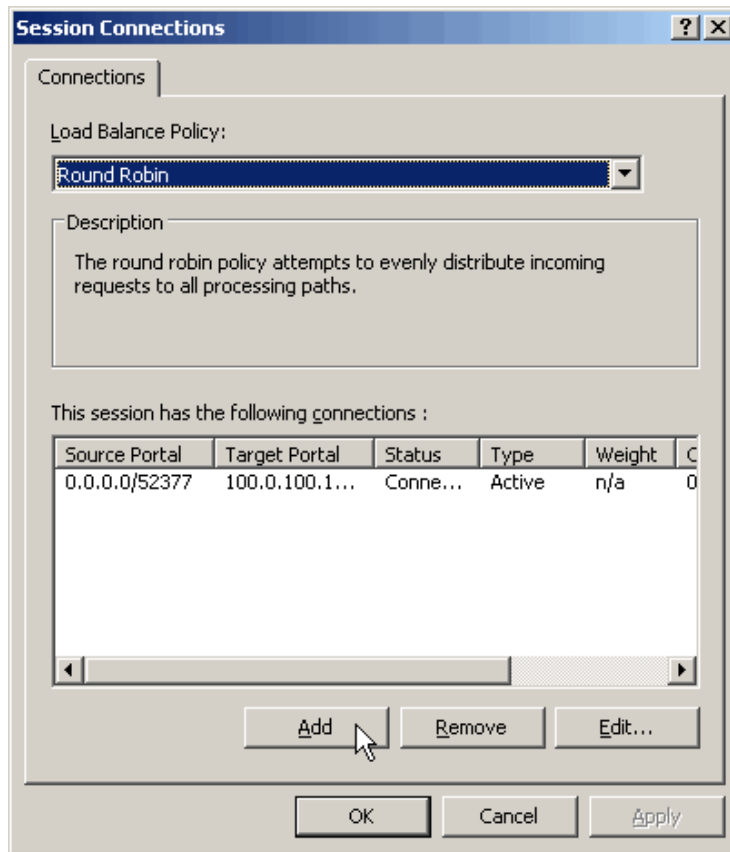


Figure 15. Configure Multi Connections – Session Connections Window

9. Click [Add].
The Add Connection window appears (Figure 16).




Figure 16. Configure Multi Connections – Add Connection Window

10. Click [Advanced...].
The Advance settings window appears (Figure 17).

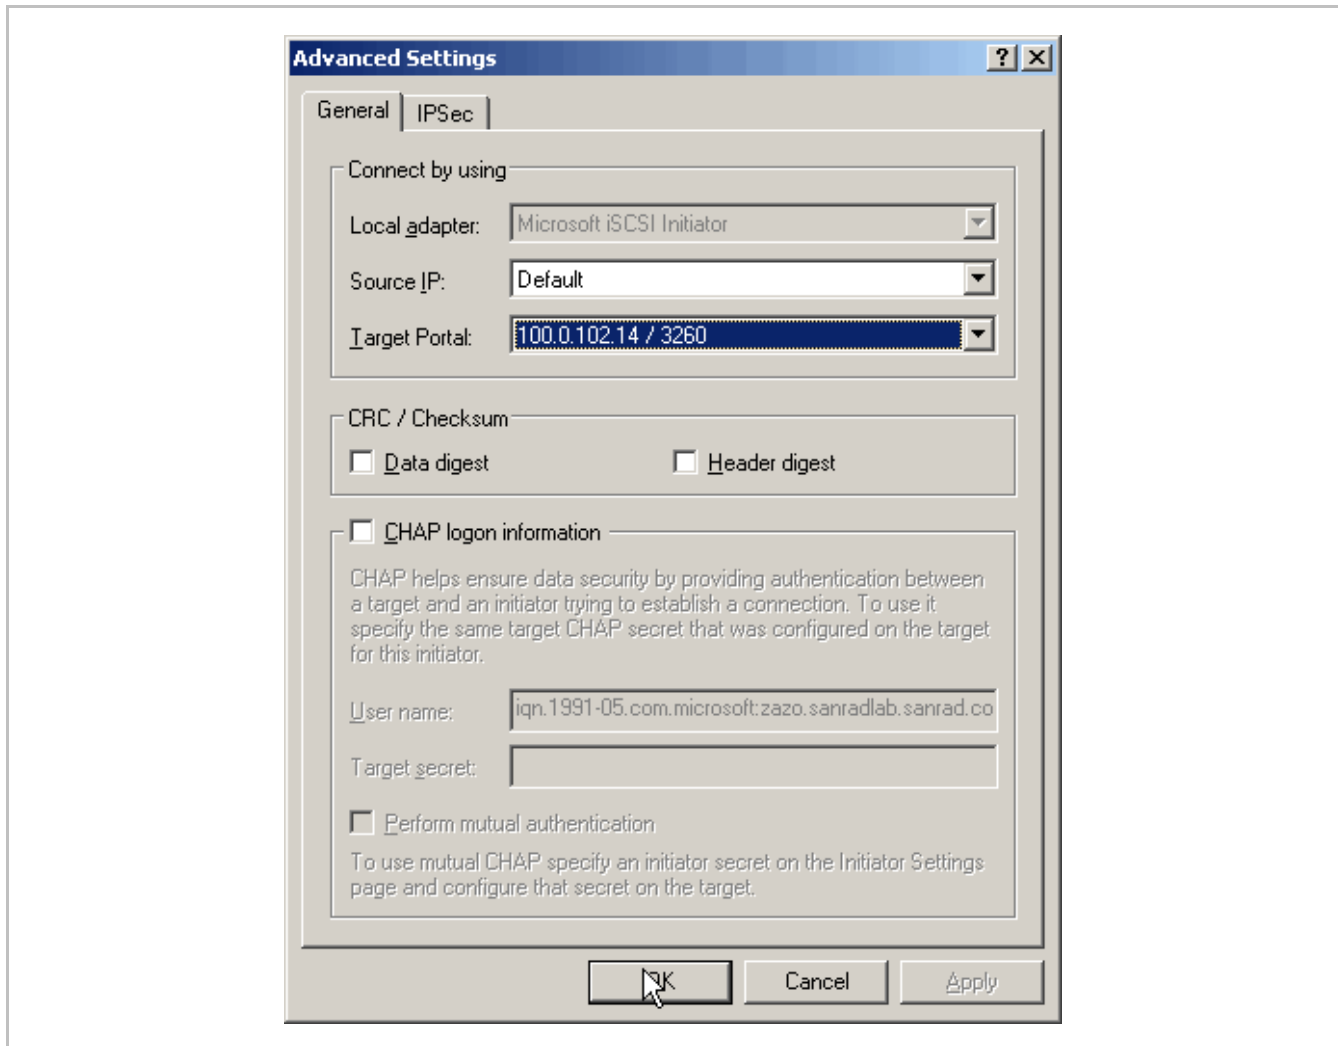


Figure 17. Configure Multi Connections - Advance settings Window

11. Choose the portal on subnet B and click **[OK]**.
12. Click **[OK]** (Figure 16).
The Session connections window appears (Figure 18). You should see two connections for this session.
13. Choose the **Load Balance Policy**:
 “**Round Robin**” for I/O load balance traffic and fail over
 or
 “**Fail Over Only**”.

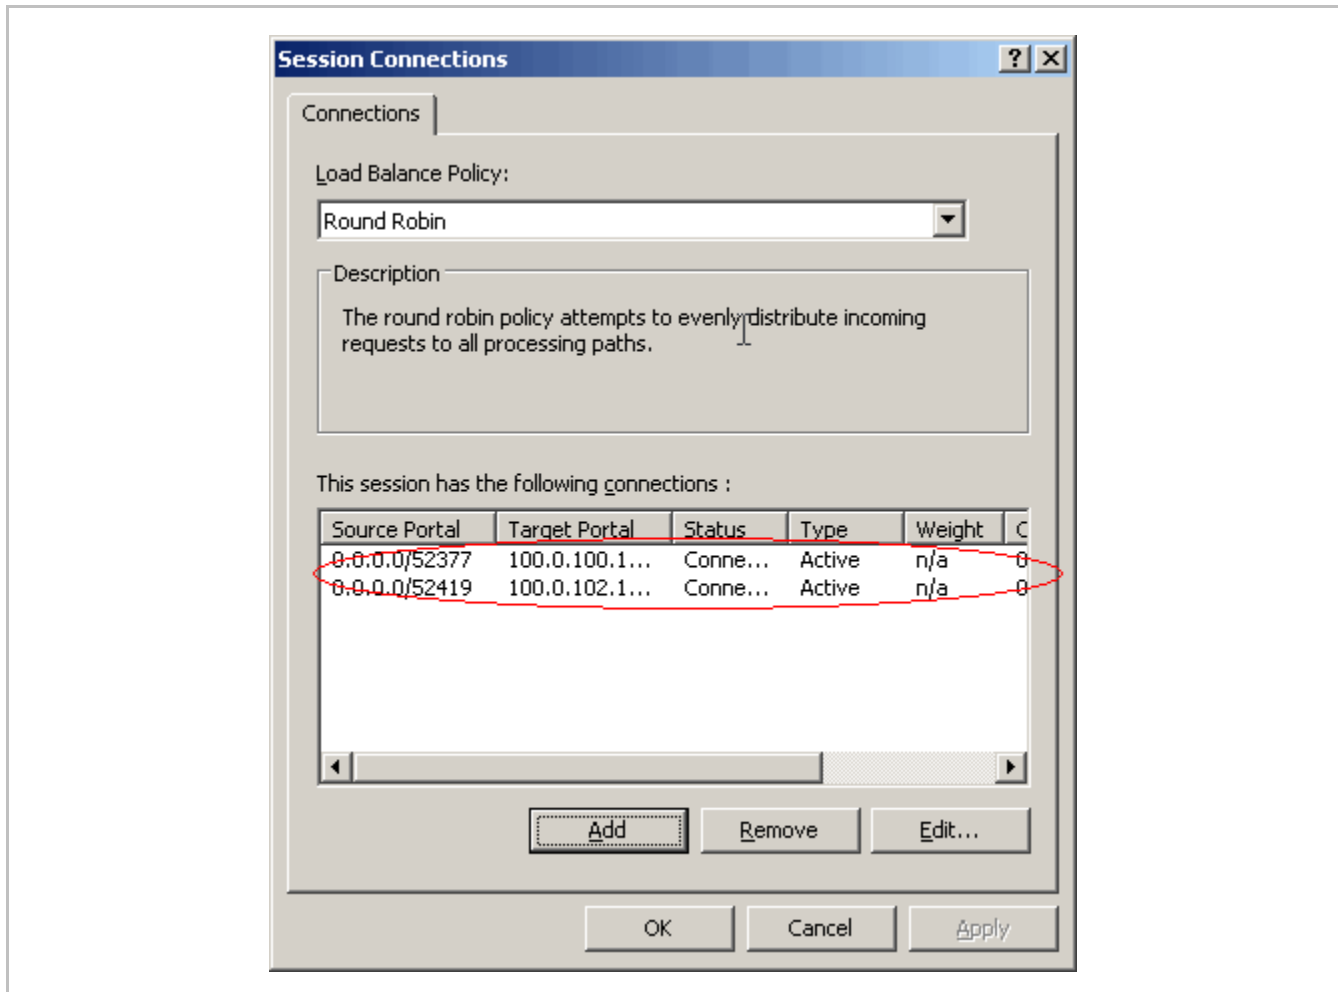


Figure 18. Configure Multi Connections – Session Connections Window

Testing

1. Disconnect the cable from one of the Ethernet ports on the V-Switch.
2. Wait approximately 2 minutes
3. Open Explorer and open a file inside the disk (represented by the target with the MPIO).
4. Try to read the file or write to it.

Basic Troubleshooting

After disconnecting the cable from one of the Ethernet ports on the V-Switch:

- Make sure you can ping the redundant portal IP from the server.
- Open the Microsoft iSCSI initiator GUI, choose the target and click [**Refresh**]. Make sure the target status is still "Connected".

Summary

Multi Path I/O for iSCSI provides redundancy for connectivity between the servers and the iSCSI targets exposed by SANRAD's iSCSI V-Switch. When designing a "No Single Point of System Failure" high availability application, it's important to include MPIO.