

iSCSI Multi-Path I/O (MPIO) Example

This is an example of how to create an MPIO connection to an iSCSI target in a high performance environment. By high performance, we mean that we have an isolated VLAN for iSCSI traffic, and that **only** iSCSI traffic will travel on the VLAN. All non-essential protocols (Netbios, client for Microsoft networks, etc.) have been unbound from our iSCSI interfaces. We will be connecting to a target that supports active-active using the Microsoft DSM Round-Robin mode.

This document is provided as a convenience only. If you have access to the Microsoft iSCSI initiator Users Guide version 2.06 or later, you should reference that document, as it has a much more thorough treatment of this subject. Furthermore, that document should be regarded as authoritative if there are conflicts with these instructions, unless those conflicts are specific to the way Alacritech products function.

Before starting, you need to make sure that the target you plan to use supports MPIO, and that it supports active-active connections. If the target supports MPIO, but does not support active-active, you can still make an MPIO connection, but the only supported mode will be failover. There's nothing wrong with failover mode, it gives you network redundancy, but it does not provide the performance increase that the other MPIO modes do.

The other thing you should know is whether your target manufacturer has their own MPIO DSM (Device Specific Module). If they do, it may be preferable to use their DSM mode instead of using the round-robin mode that we use in this example.

Our target is an EqualLogic PS100 Storage Array with a portal address of 192.168.51.1/22. The target has 3 ethernet interfaces all on this subnet. The target has already had a volume set up, and has been configured so that our initiator host has access to it.

The initiator host is running Windows Server 2003 R2 SP2, and has the Windows iSCSI initiator version 2.07, including MPIO, installed. It has an Alacritech SES2002XT dual port iSCSI Accelerator that we will be using to connect to the target. The iSCSI interfaces are configured as follows:

```
C:\Temp>ipconfig
Windows IP Configuration
...
Ethernet adapter 3A_iSCSI:

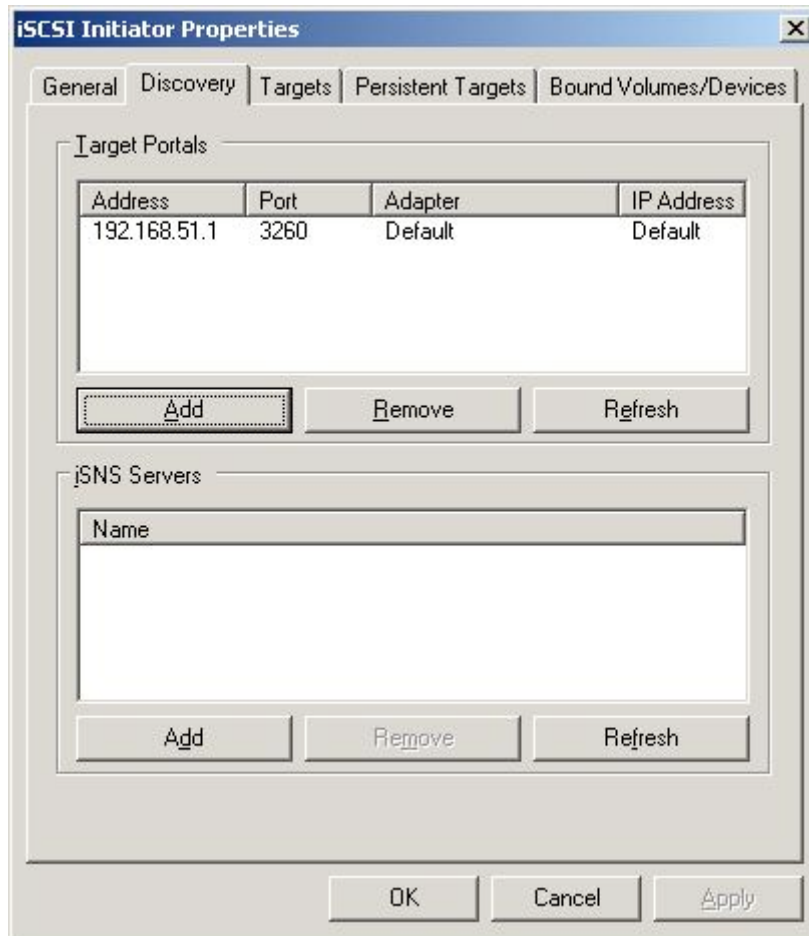
    Connection-specific DNS Suffix  . :
    IP Address. . . . . : 192.168.50.105
    Subnet Mask . . . . . : 255.255.252.0
    Default Gateway . . . . . :

Ethernet adapter 3B_iSCSI:

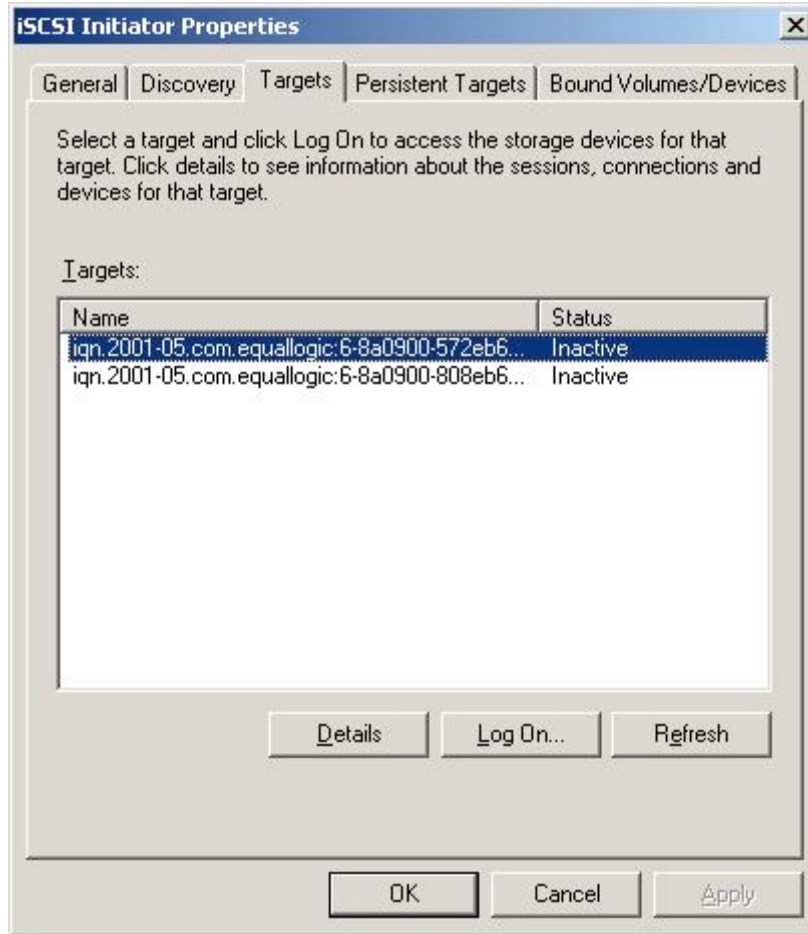
    Connection-specific DNS Suffix  . :
```

IP Address : 192.168.50.125
Subnet Mask : 255.255.252.0
Default Gateway :

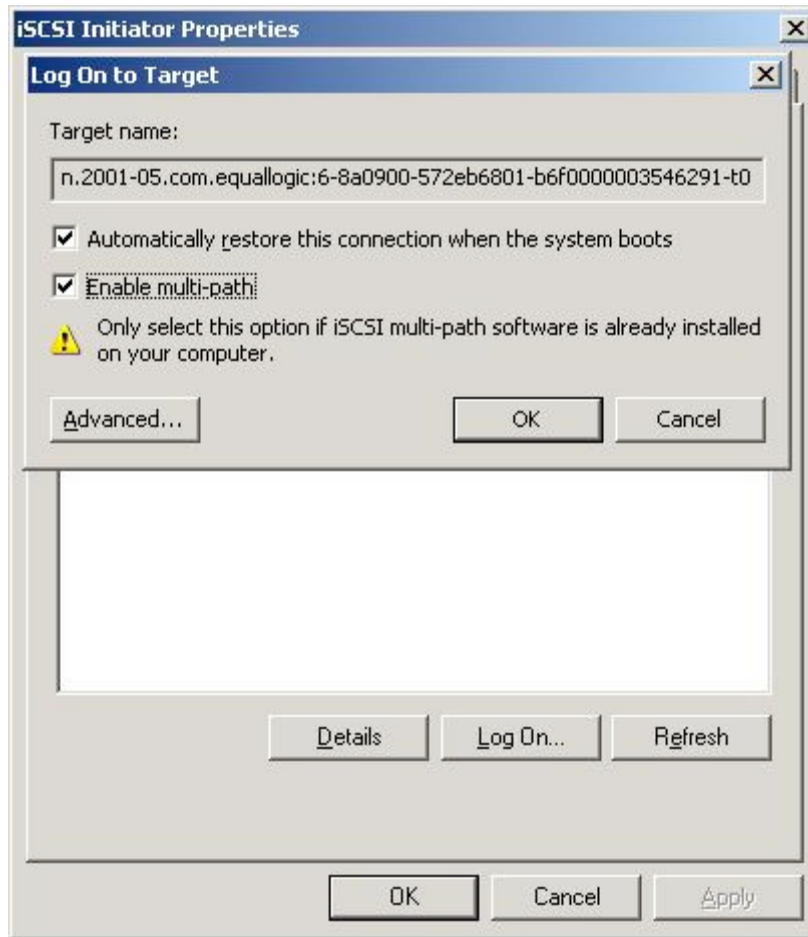
After opening the iSCSI initiator applet, the first thing we do is add the portal address of the PS100.



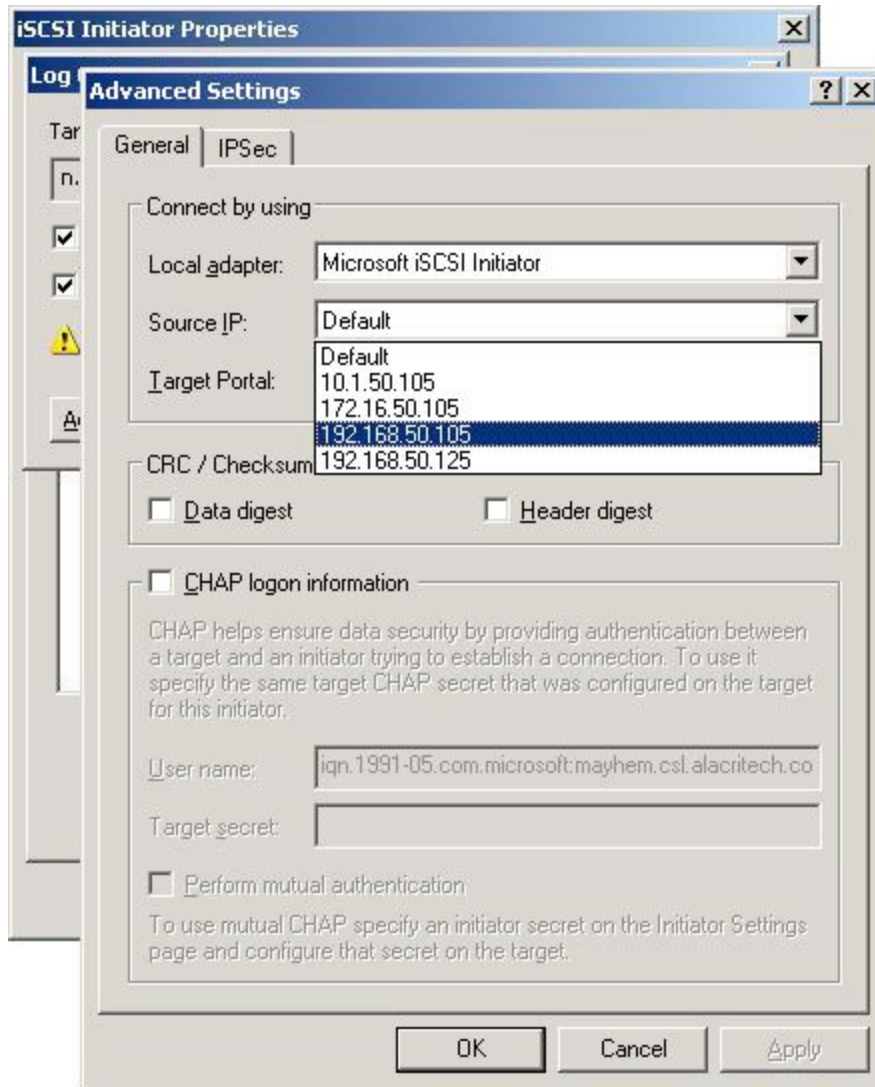
Once the portal address has been entered, the available targets are displayed. It is not obvious which target we want to log on to, so we need to look at the configuration of the target itself. Since that process is different for each manufacturer's target, we won't cover that here, but by default, we know the highlighted target is the one we want.



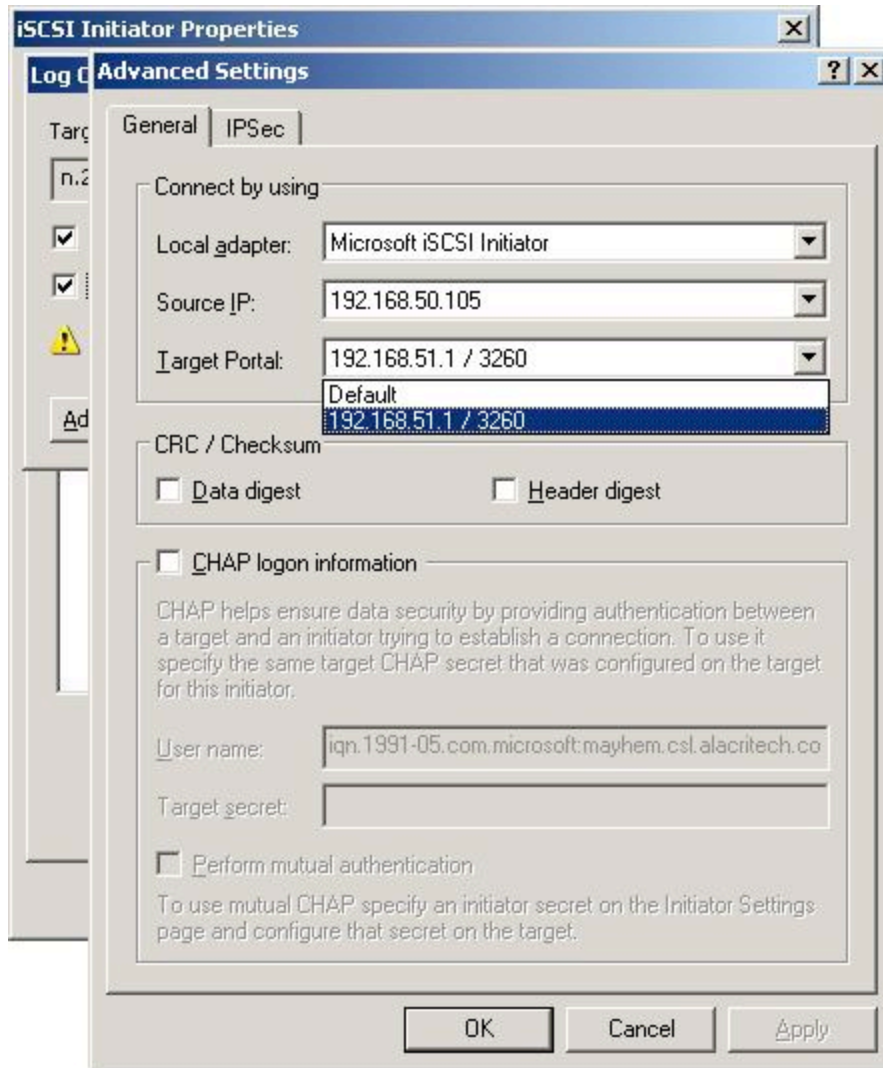
Click "Log On", and check both the box for persistent connections and the box for Multi-Path.




We then select "Advanced" so that we can configure the individual connects. Set "Local adapter" to "Microsoft iSCSI Initiator", and select one of the two IP addresses on our iSCSI network as the source IP.



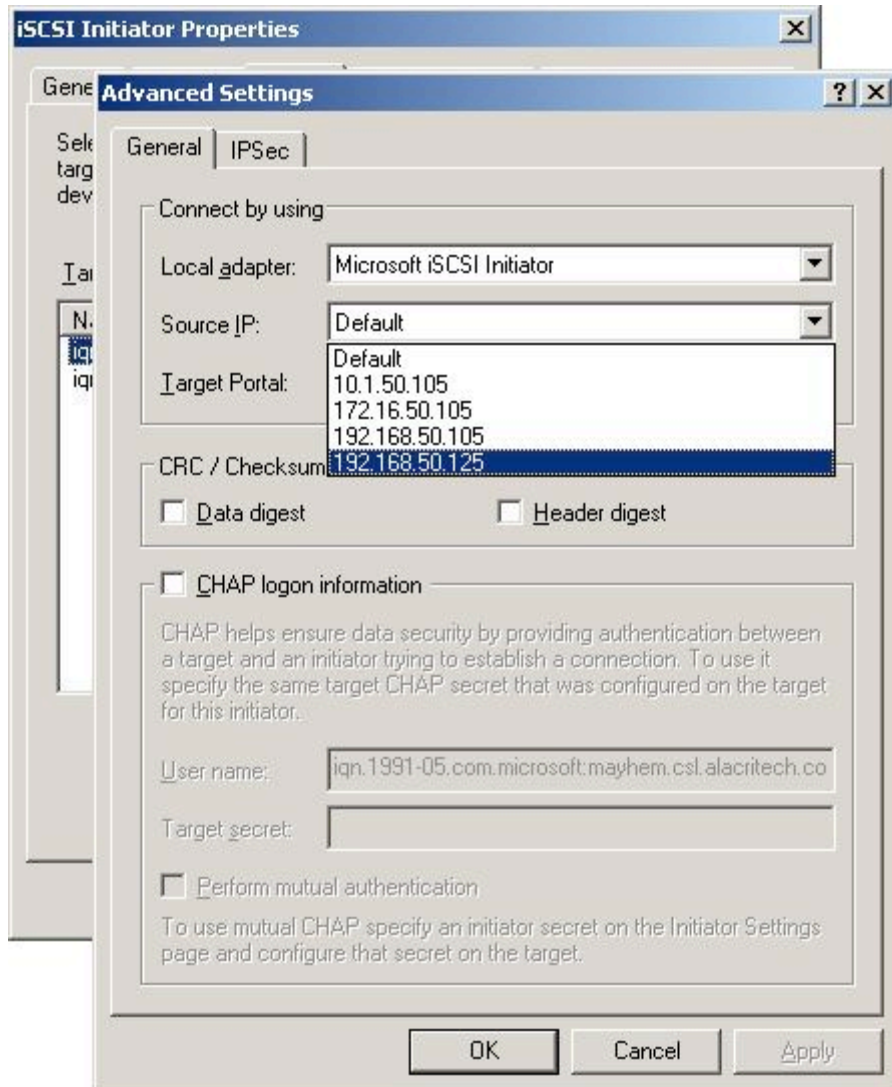
Next we select the destination address. On most targets we would see several destination addresses here, one for each ethernet interface, and we would choose one. The PS100 is different though. Because EqualLogic uses a technology called iSCSI redirect, we see only one address here. We select that address, and the PS100 will take care of redirecting the connection to the appropriate interface.



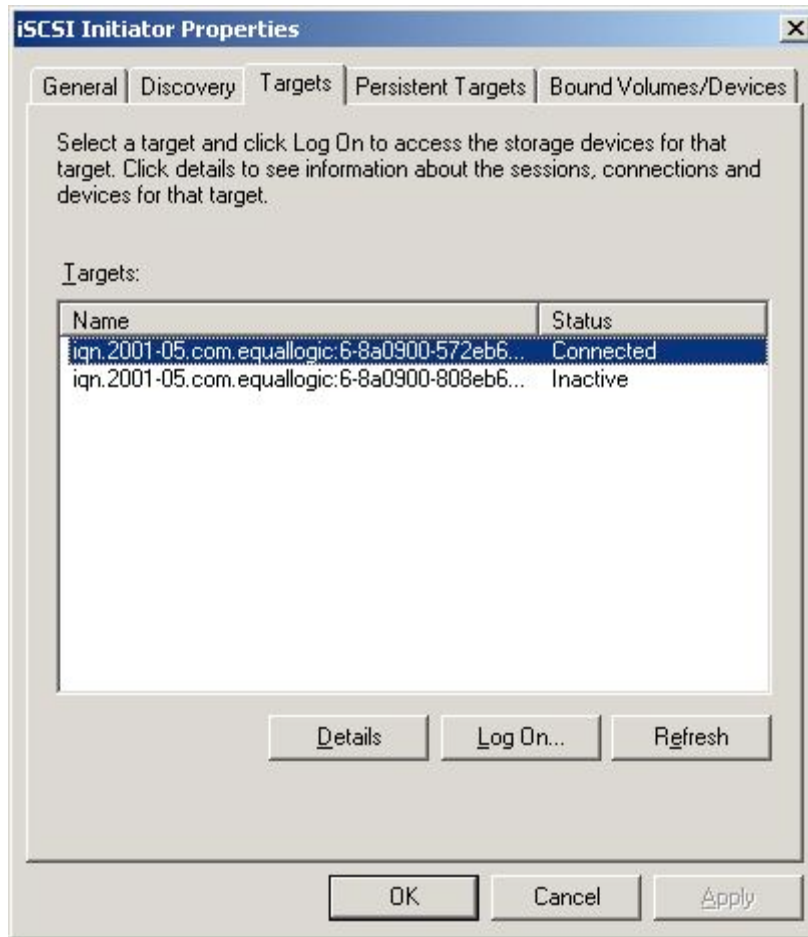
After hitting "OK" several times, we get back to the target window which now shows as "connected".



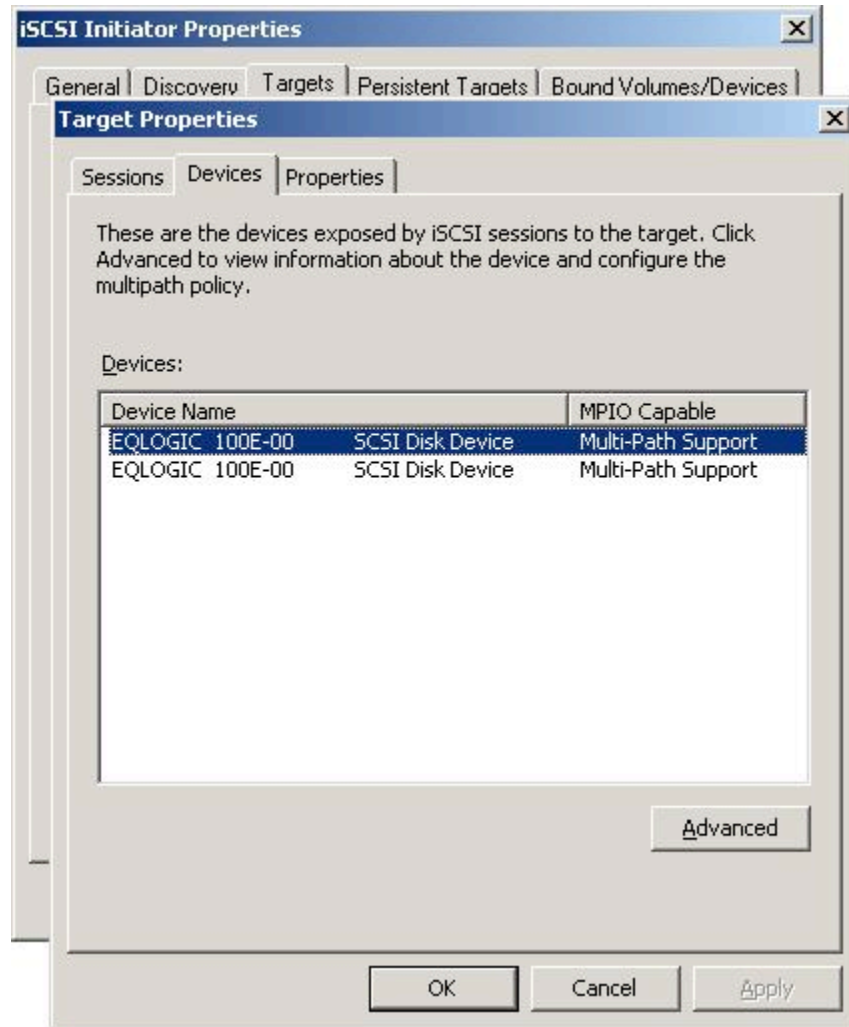
Now we simply repeat the login process. The only difference is that this time we select a different IP source address. Had this been a target that exports all target addresses, we would also select a different target address, but the iSCSI redirect implementation on the PS100 makes this unnecessary.



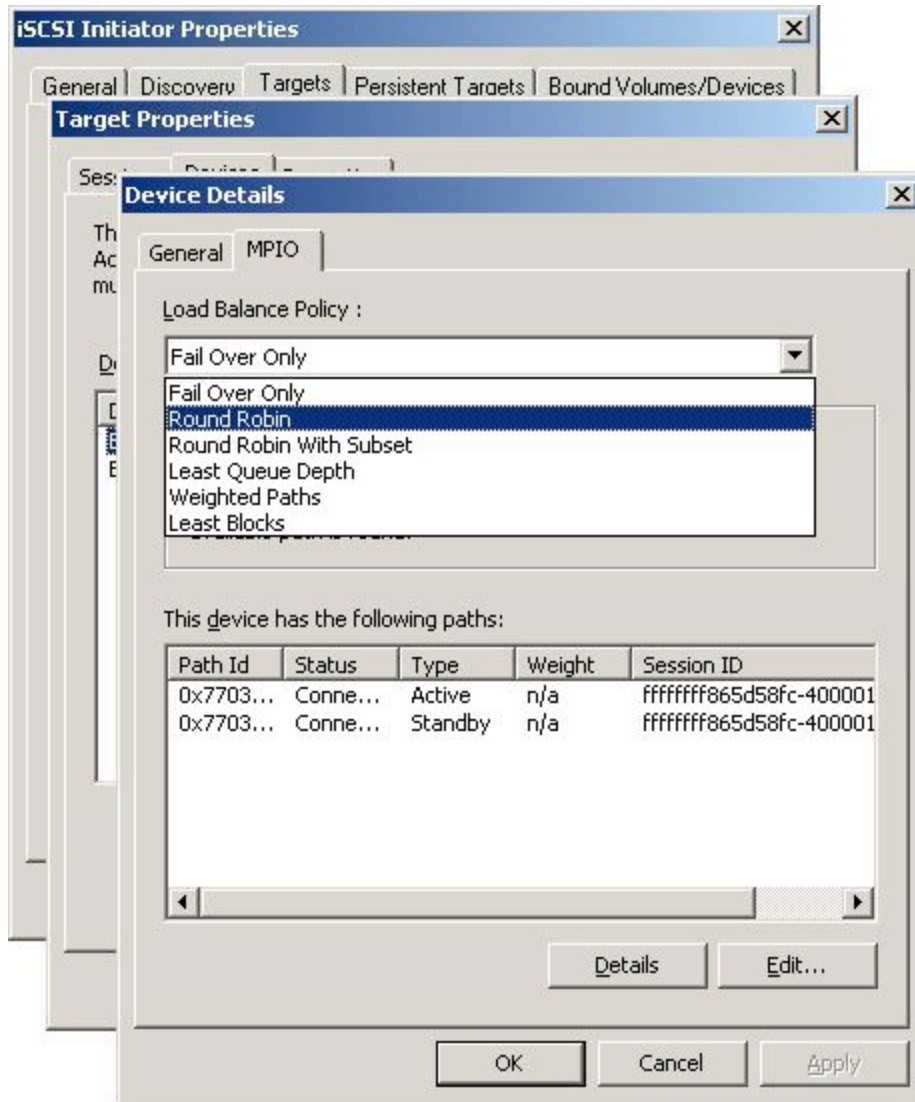
Now we are back at our target window, and it looks no different than it did after logging on the first time. Select the target again and hit the "Details" button.



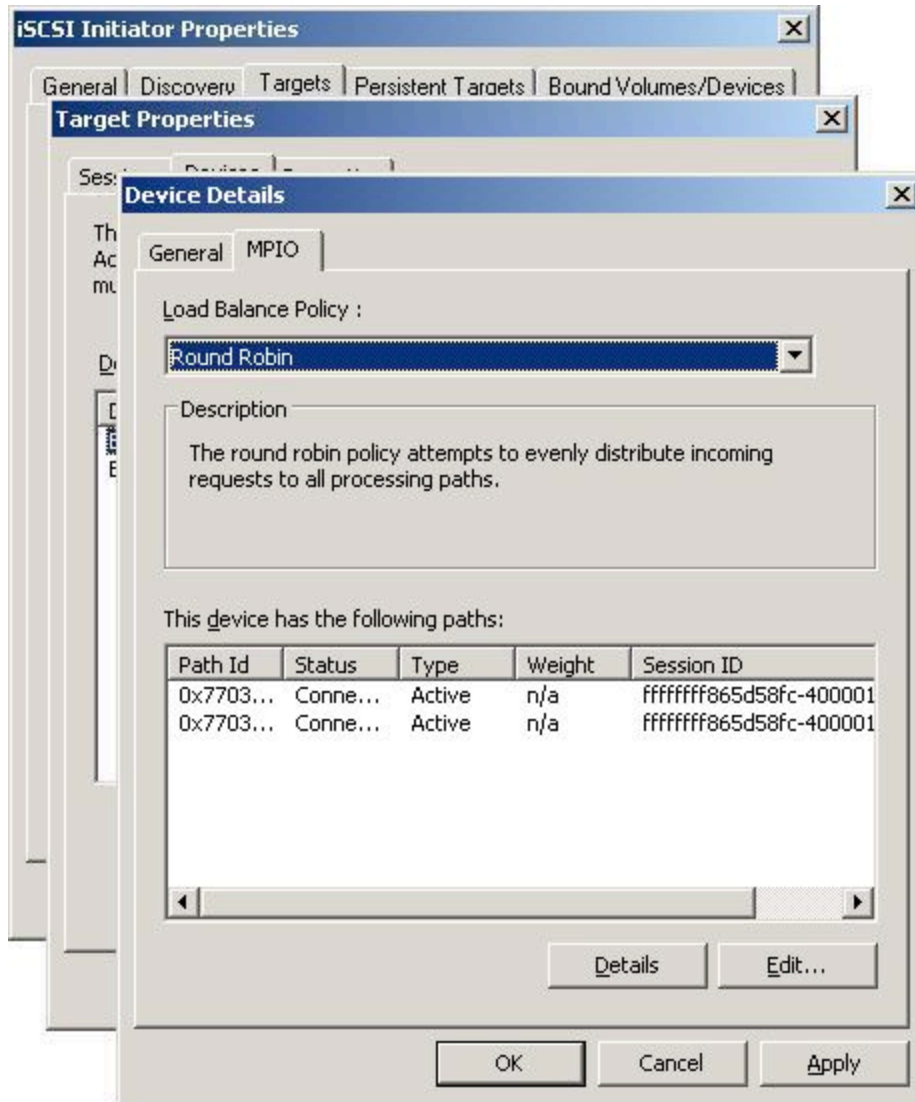
This opens the target properties pane. Now select the "Devices" tab and highlight one of the available devices.



With a device highlighted, activate the "Advanced" button to bring up the "Device Details" pane. This pane should show the "Load Balance Policy" as "Fail over only", and should show one connection as active, and one as standby.



Change the "Load Balance Policy" to "Round-Robin". All paths should immediately change to active, and your MPIO connection will now be fully functional. The bandwidth to target is now equal to the sum of the bandwidth of completely independent paths. In this example, we have two completely independent paths of one gigabit each, so the bandwidth to target is two gigabit. Had we more network interfaces available we could increase the bandwidth even further.



If we now run the command `netstat -nt` we see that we have a connection from 192.168.50.105:2817 to 192.168.51.2:3260, and from 192.168.50.125:2819 to 192.168.51.4:3260, and that both connections are offloaded, indicating that everything is working as it should.

```

c:\ Shortcut to cmd.exe
Subnet Mask . . . . . : 255.255.0.0
Default Gateway . . . . . : 10.1.1.1

Ethernet adapter 3A_iSCSI:

    Connection-specific DNS Suffix . :
    IP Address . . . . . : 192.168.50.105
    Subnet Mask . . . . . : 255.255.252.0
    Default Gateway . . . . . :

Ethernet adapter 3B_iSCSI:

    Connection-specific DNS Suffix . :
    IP Address . . . . . : 192.168.50.125
    Subnet Mask . . . . . : 255.255.252.0
    Default Gateway . . . . . :

C:\Temp>netstat -nt

Active Connections

   Proto Local Address          Foreign Address        State                   Offload S
tate
   TCP    172.16.50.105:1133     172.16.50.2:445       ESTABLISHED            Offloaded
   TCP    192.168.50.105:2817   192.168.51.2:3260     ESTABLISHED            Offloaded
   TCP    192.168.50.125:2819   192.168.51.4:3260     ESTABLISHED            Offloaded

C:\Temp>

```