




N4866A PCIe Interconnect

Installation and User's Guide



Manual Part Number: N4866-97000, April 2010

© Copyright Agilent Technologies 2010
All Rights Reserved.



Agilent Technologies

Contents

This Guide ... 3

Section 1: Introduction and Overview 4

Section 2: Quick Installation Guide 6

Section 3. Requirements and Planning 7

3.1 Supported Platforms 7

3.2 Interconnect Planning 8

Section 4. Initial Installation 9

4.1 N4866A PCIe Interconnect Adapter Card/Riser Card Installation 9

4.2 Software and Cable Installation 11

4.3 Installation Overview 31

Section 5. Interconnect Maintenance 32

5.1 Verifying Functionality and Performance 32

5.2 Replacing Interconnect Cables 36

5.3 Replacing an Adapter 37

Section 6. SuperSockets 38

6.1 Make Generic Applications use SuperSockets 38

6.2 SuperSockets Functionality and Performance 38

Appendix A. Configuration Files 40

Appendix B. DX Firmware upgrade 43

This Guide ...

This document describes the installation of the Agilent N4866A PCIe Interconnect hardware and the Dolphin software stack (including SuperSockets) on single machines or on a cluster of machines. This software stack is needed to use Agilent's high-performance interconnect and consists of drivers (kernel modules), user space libraries / applications, documentation, and more. SuperSockets drastically accelerate generic socket communication as used by clustered applications.

This document is for clusters built using the Dolphin Express DX software.

Section 1: Introduction and Overview

Who needs the Agilent N4866A / Dolphin Express Software and SuperSockets?

Clustered applications running on multiple machines communicating via an Ethernet-based network often suffer from the delays that occur when data needs to be exchanged between processes running on different machines. These delays caused by the communication time make processes wait for data when they otherwise could perform useful work. The Agilent N4866A combined with the Dolphin Express software is a combination of a high-performance interconnect (to replace the Ethernet network) together with a highly optimized software stack. One part of this software stack is SuperSockets which implements a bypass of the TCP/UDP/IP protocol stack for standard socket-based inter-process communication. This bypass moves data directly via the high-performance interconnect and thereby reduces the minimal latency typically by a factor of 10 or more with 100% binary application compatibility.

How does the Agilent N4866A / Dolphin Express and SuperSockets work?

The Agilent N4866A PCIe Interconnect hardware provides means for a process on one machine to write data directly into the address space of a process running on a remote machine. This can be done using either direct store operations of the CPU (for lowest latency) or using the DMA engine of the N4866A Interconnect adapter (for lowest CPU utilization). WinSock2 applications are linked to several dynamic link libraries that are loaded at runtime. With the exception of kernel32.dll, the operating system looks for import symbols first in the directory from which the application loaded. If the Dolphin Sockets DLL resides in the same directory and the configuration file dishosts.conf is valid, then all calls are redirected to the low-latency SCI transport. Otherwise, the standard NDIS (TCP/IP protocol stack) transport is used.

What is Special about the Agilent N4866A / Dolphin Express and Super-Sockets?

The Agilent N4866A / Dolphin Express and SuperSockets offer a number of unique and outstanding features for high-performance communication via the socket interface:

- **Lowest Latency SuperSockets:** Offers the lowest effective latency.
- **Highest Message Rate:** The effective application performance is not only determined by the low latency, but also by the rate at which messages can be processed. This is different from raw latency. SuperSockets has an excellent message rate as well.
- **High Scalability:** The low protocol overhead of SuperSockets for both CPU and host adapter ensures low CPU utilization and a good scalability with the number of cores in each machine.
- **Interruptless Mode:** For real-time environments, SuperSockets can be operated in a mode which does not trigger interrupts during communication. This allows for the lowest overhead and high timing predictability.
- **WinSock2 Semantics:** SuperSockets offer a basic subset of the Winsock2 socket protocols. Please consult the release note for more information.
- **Free Open Source:** SuperSockets are licensed by the Gnu Public License v2 and thus are free open source with all the inherent benefits.

Contents of this Document

This document is the installation and reference guide for the Agilent N4866A which is to be used with the Dolphin Express and SuperSockets software. It is structured as follows:

- Section 1: *Introduction & Overview*** is what you are currently reading. (Page 4)
- Section 2: *Quick Installation Guide*** is for those with some experience in cluster management (these are the abbreviated installation instructions). The *Quick Installation Guide* might be sufficient if you have experience. Otherwise, the subsequent chapters cover all of the details. (Page 6)
- Section 3: *Requirements and Planning*** discusses the platform support and choice of the N4866A PCIe Interconnect hardware and is best read before you decide on using the Agilent N4866A / Dolphin Express and SuperSockets. (Page 7)
- Section 4: *Initial Installation*** should be followed if you are installing the Agilent N4866A PCIe Interconnect, the Dolphin Express software, and SuperSockets for the first time. (Page 9)
- Section 5: *Interconnect Maintenance*** lists procedures for typical maintenance tasks and for trouble shooting once the cluster is installed. (Page 32)
- Appendices:** The appendices provide the configuration files of the Dolphin Express software stack and a listing of known issues with the platform as well as default software limitations.

Terminology

We define some terms that will be used throughout this document:

- Adapter:** An Agilent N4866A host adapter. This is the Agilent N4866A hardware installed in the cluster nodes.
- Node:** A computer or oscilloscope that is part of the Agilent N4866A interconnect, which means it has an *adapter* installed. All nodes together constitute the *cluster*.
- CPU architecture:** The CPU architecture relevant in this guide is characterized by the addressing width of the CPU (32 or 64 bit) and the instruction set (x86, Sparc, etc.). If these two characteristics are identical, the CPU architecture is identical for the scope of this guide.
- Link:** The cable between two adapters or between an adapter and a switch.
- Cluster:** All *nodes* constitute the *cluster*.
- Node manager:** The node manager is a daemon process that is running on each node and provides remote access to the interconnect driver and other node status to the network manager. It reports status and performs actions like configuring the installed adapter or changing the interconnect routing table if necessary.

The service name of the network manager is `dis_nodemgr`.

•**Windows Installer (MSI):** The Windows Installer is an engine for the installation, maintenance, and removal of software on modern Microsoft Windows systems. The installation information, and often the files themselves, are packaged in installation packages, loosely relational databases structured as OLE Structured Storage Files and commonly known as "MSI files" from their default file extension.

•**x4:** PCI express slots and the Agilent N4866A combine multiple lanes (serial high-speed communication channels using few electrical connections) into communication paths with a higher bandwidth. With PCI Express Gen. 1, each lane carries 2.5Gbit/s of traffic. Combining 4 lanes into a single communication path is called x4 and thus delivers 10Gbit/s bandwidth.

Support

If that have any problems with the procedures described in this document, or have suggestions for improvement, please contact Agilent Technical Support. Visit www.agilent.com and select "Contact Us" at the top of the page to locate Agilent's contact information for your specific country.

Section 2: Quick Installation Guide

This section gives an overview of the steps necessary for a default installation using the Dolphin Express Windows Installer (MSI).

- 1 *Installation requirements:* The Windows Installer will work well on x86 and x86_64 platforms.

For details, see Section 3, *Requirements and Planning* for general platform requirements of the Dolphin Express software stack and Section 4, *Initial Installation* - “Installation Overview” for specific requirements for this installation procedure.

- 2 *Adapter installation:* Install the Agilent N4866A PCIe interconnect adapters/riser cards into the nodes (PC and scope(s)).

NOTE The cables should not be installed in this step!

For details, see Section 4, *Initial Installation* - “Agilent N4866A Hardware Installation”.

- 3 *Software download:* Download the MSI installers from http://www.dolphinics.com/support/dolphin_agilent_download.html (the oscilloscope driver package should be installed on the oscilloscope and the host driver package should be installed on the PC).

- 4 *Start software installation:* Ensure you are an Administrator on the PC/Scope and start the software installation by double clicking on the correct MSI file.

Repeat the procedure for all nodes in the cluster.

For details on this part of the installation, please see Section 4, *Initial Installation* - “Software and Cable Installation”.

- 5 *Interconnect configuration:* After the software has been installed the Dolphin Network Configurator is started to verify the configuration files for the interconnect.

For details on how to perform this task, see Section 4, *Initial Installation* - “Working with the Dolphin Network Configurator, dis_netconfig”.

- 6 *Install cables:* Using the Dolphin Network Configurator (dis_netconfig), you can create the cabling instructions for the Agilent N4866A PCIe interconnect. Connect the cables according to these instructions.

For details on cable installation, see Section 4, *Initial Installation* - “Cluster Cabling”.

- 7 Enable your applications to make use of SuperSockets: ws2_32.dll - Dolphin Sockets dynamic library. It is loaded instead of the original system socket switch (%SYSTEMROOT%/System32/ws2_32.dll). The ws2_32.dll has to be placed in the directory where the target application binary is located. (ATTENTION: Do not delete the original file %SYSTEMROOT%/System32/ws2_32.dll).

For details on this step, please see Section 4, *Initial Installation* - “Making Cluster Application use Dolphin Express”

Section 3. Requirements and Planning

Before you start using the N4866A PCIe Interconnect by either adding it to an existing system or by planning it into a new system, some considerations on selection of products and the physical setup are necessary.

3.1 Supported Platforms


The Agilent N4866A/Dolphin Express software stack is designed to run on the Windows XP software platform.

3.1.a Hardware

Agilent N4866A Hardware Kit

Kit contents

- One PCI Express PC Host Adapter
- One PCI Express SCOPE Host Adapter
- One riser card
- One 10 meter fiber optic cable
- This *User's Guide* on CD-ROM



Hardware Installation Procedure

- Please refer to Section 4.1 *N4866A PCIe Interconnect Adapter Card/Riser Card Installation* (Page 9) for information on installing the N4866A PCIe Interconnect.

Recommended Node Hardware

The Agilent N4866A host adapters are low profile cards with a full-height bracket. The power consumption of the N4866A adapters is between 5W and 15W (Consult the separate data sheets for detail). The Agilent N4866A PCIe Interconnect is fully inter-operable between all supported hardware platforms, also with different PCI or CPU architectures.

3.1.b Software Support

The Dolphin Express software stack that is to be used with the Agilent N4866A is currently available for Windows XP.

3.2 Interconnect Planning

This section discusses the decisions that are necessary when planning to install a N4866A PCIe Interconnect.

3.2.a Nodes to Equip with Agilent N4866A PCIe Interconnect

Depending on the application that will run on the cluster, the choice of N4866A PCIe Interconnect equipped machines differs.

3.2.b Agilent N4866A PCIe Interconnect Topology

Two nodes can be connected directly using one CX4 cables to connect port 0 of each adapter.

Three nodes can be connected directly with a link width x4 between each pair of nodes. This enables you to control two oscilloscopes with one PC. You can also connect up to 10 nodes by using a Dolphin DXS410 switch. This switch has 10 ports and they can be used flexibly to connect to nodes.

3.2.c Physical Node Placement

Generally, nodes that are to be equipped with Agilent N4866A PCIe Interconnect host adapters should be placed close to each other to keep cable lengths short. This reduces costs and allows for better routing of the cables. The N4866A supports both copper and fiber optic cables (a 10 m fiber optic cable is included in the N4866A). The maximum cable length is 10 meter for copper cables (CX4) and 100 meter for fiber.

Section 4. Initial Installation

This chapter guides you through the initial hardware and software installation of the N4866A PCIe Interconnect, Dolphin Express software, and SuperSockets. This means that no Dolphin software is installed on the nodes or the frontend prior to these instructions.

The recommended installation procedure, which is described in this chapter, uses the windows installer image *.msi distribution of the Dolphin Express software stack.

4.1 N4866A PCIe Interconnect Adapter Card/Riser Card Installation


The N4866A ships with a PC host adapter, an oscilloscope host adapter, a riser card, and a 10 m fiber optic cable. The adapters and riser card must be installed before downloading the Dolphin software.

WARNING: Make sure you are properly grounded to avoid static discharges that may destroy the hardware. When the card is properly inserted and fixed, close the oscilloscope chassis and power up the node again. The LED(s) on the adapter slot cover should have an orange color.

WARNING (RISK OF SHOCK): Disconnect the AC power cord from both the oscilloscope and the PC before starting the installation procedures described below.

4.1.a Installation of PC Host Adapter

To install the PC host adapter, open the PC tower chassis so you can see the motherboard and insert the PC host adapter into a free PCI Express slot as shown below. Please note that the PCIe slot may be larger than the PC host adapter. If this is the case then insert the host adapter into the slot such that it is as close to the rear of the PC as possible (in other words, allow any extra portion of the PCIe slot to be towards the front of the tower instead of the rear).




4.1.b Installation of Oscilloscope Host Adapter and Riser Card

WARNING (RISK OF SHOCK): Disconnect the AC power cord from both the oscilloscope before starting the installation procedures described below.

To install the oscilloscope host adapter, remove the top cover and top plate from the 90000A Series oscilloscope as shown in the *90000A Series Service Guide* (“Replacing Assemblies” chapter). Once the oscilloscope chassis is open, find the motherboard and the PCIe slot.

Remove the slot cover from the rear panel of the oscilloscope. **Be sure to keep the screw as you will need it to install the adapter card.** Now, insert the riser card into the PCIe slot. Then insert the oscilloscope host adapter into the riser card. Again, the host adapter may be smaller than the riser card so insert it such that any open portion is towards the front of the oscilloscope (the host adapter should be as close to the rear of the oscilloscope as possible as shown below). Secure the host adapter to the oscilloscope using the screw you removed when you took out the slot cover.



You do not need to connect the cables at this point. Cabling instructions customized for the specific cluster will be created during the software installation. Generally, the cables can safely be connected and disconnected with the nodes powered up as the N4866A is a fully hot-plug capable interconnect.

Note If you know how to connect the cables, you can connect them now. Please inspect and verify that the cabling is correct as described in the remainder of this chapter.

4.2 Software and Cable Installation

After the host adapters are installed, the Dolphin software needs to be installed on both the PC and the oscilloscope(s). On the nodes, the hardware driver and additional kernel modules, user space libraries, and the node manager will be installed. On the frontend, the network manager and the cluster administration tool will be installed.

4.2.a Running the Dolphin Installer

The software installation on Windows is plug and play with the Windows installer package (*.msi). The directions below will cover installing the Dolphin software on both your PC and oscilloscope(s).

On your oscilloscope, close the oscilloscope application and any other open applications. Close any open applications on your PC as well.

Connect your oscilloscope and PC to the Internet and visit the following website on both of them:
http://www.dolphinics.com/support/dolphin_agilent_download.html

If this is your first time downloading software from Dolphin, you may need to register for a valid user name and password (the website will inform you of this). If you are required to register, then once you receive your user name and password from Dolphin, visit the URL above and login. If you are not required to register then simply continue.

On your PC, you will want to select the Agilent-specific installer image called Driver Package for Host. Double-click on the **Download** link as shown below. On your oscilloscope, you will want to select the Driver package for Oscilloscope's **Download** link.

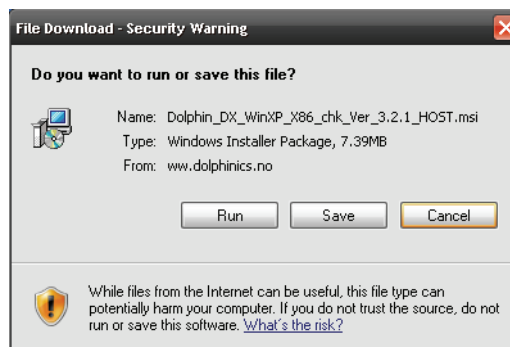
Description	OS Version	Driver Version	Software
Driver package for Oscilloscope	Windows XP 32 bit	3.2.1	Download
Driver package for Host	Windows XP 32 bit	3.2.1	Download

Click for software to install on oscilloscope(s) Click for software to install on PC

NOTE

Do not use the Windows installer for either the PC or oscilloscope installation (If the Windows installer dialog box appears, cancel it). You should instead use the File Download dialog box that appears when you click on the Download links.

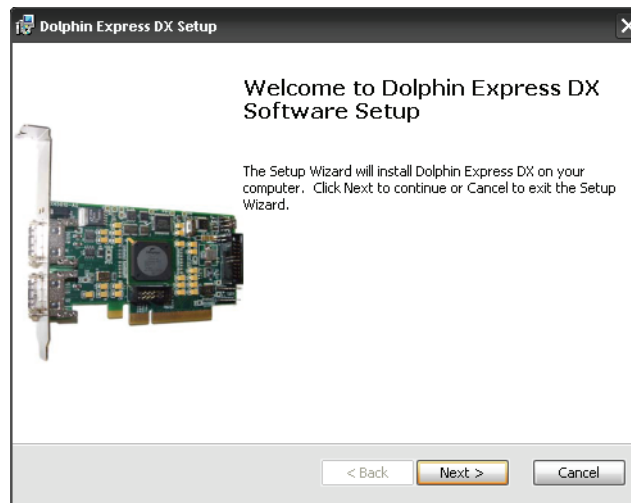
You should see the following **File Download** dialog box.



Select **Run**. Wait for the software to download and then you will be met with another dialog box asking you if you want to Run the software. Select **Run**.

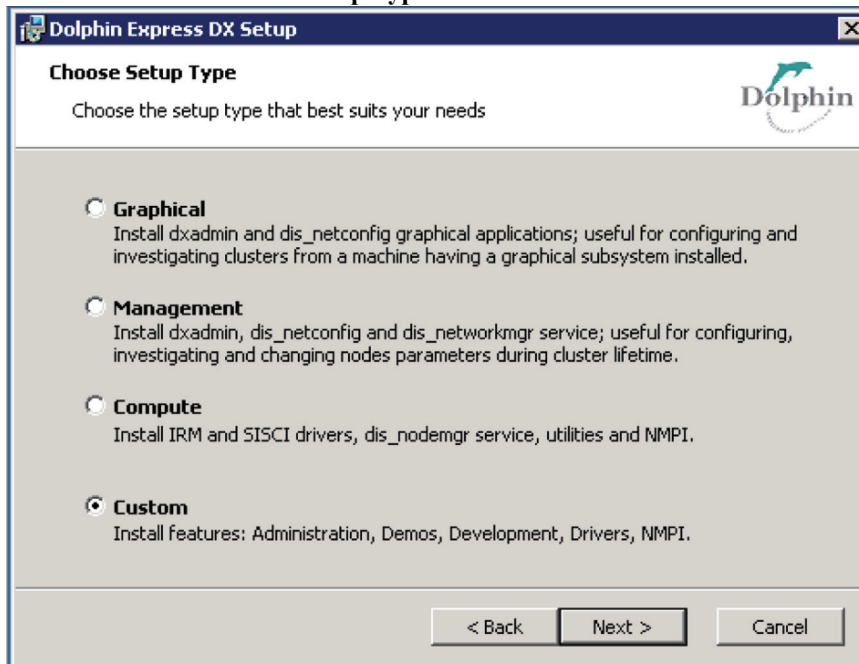
You will then be met by a welcome message.

Figure 4.3. Windows Installer: Welcome Message



Click **Next** to select which components to install.

Figure 4.4. Windows Installer: Setup Type



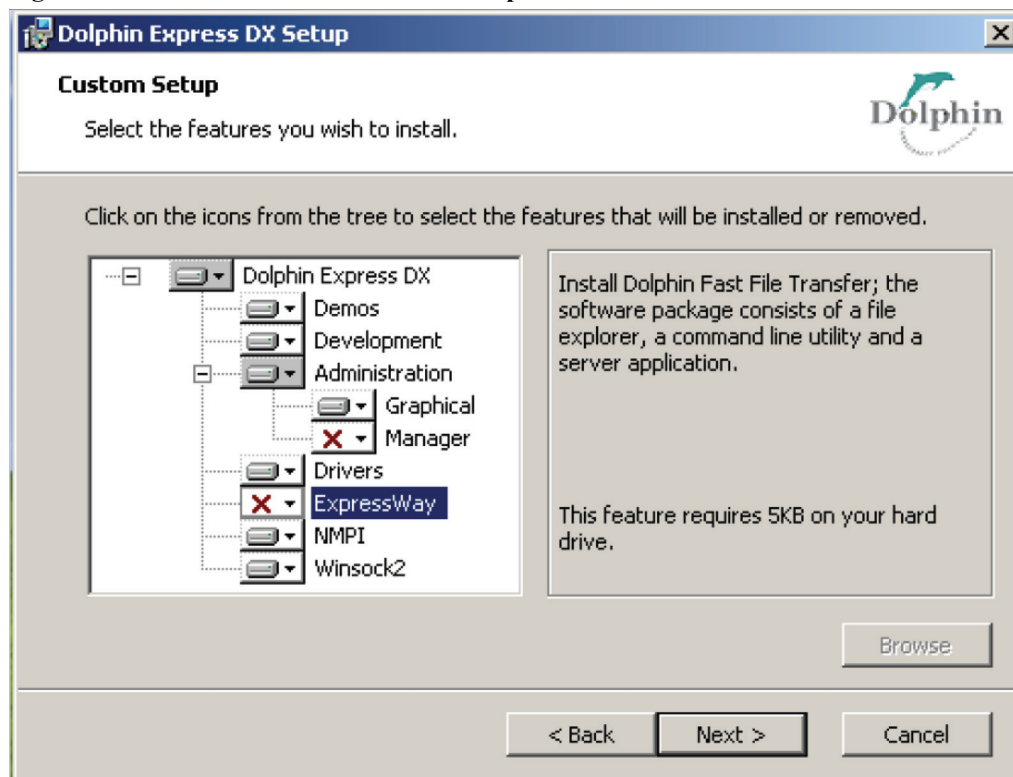
Choose **Custom** for your setup type.

The following is a general description of each of the choices in the previous dialog box (this is for your knowledge, but again, you should choose **Custom**):

- Custom** This setup type is typically used if you want to install more than one Dolphin Software module (e.g. both Compute and Management on a single node), if you want to de-select some components, or have other special requirements.
- Graphical** installation can be used to:
 - Generate configuration files on a machine away from the running cluster.
 - Monitoring the Dolphin Admin GUI on a loosely connected machine.
- Management** selection installs the centralized management tools. You may run the centralized management daemon, Dolphin Network Manager on one of the nodes in the cluster, or on a separate node that does not require Dolphin hardware.
- Compute** node installs drivers, the node manager daemon and other software that runs on the Dolphin Adapters.

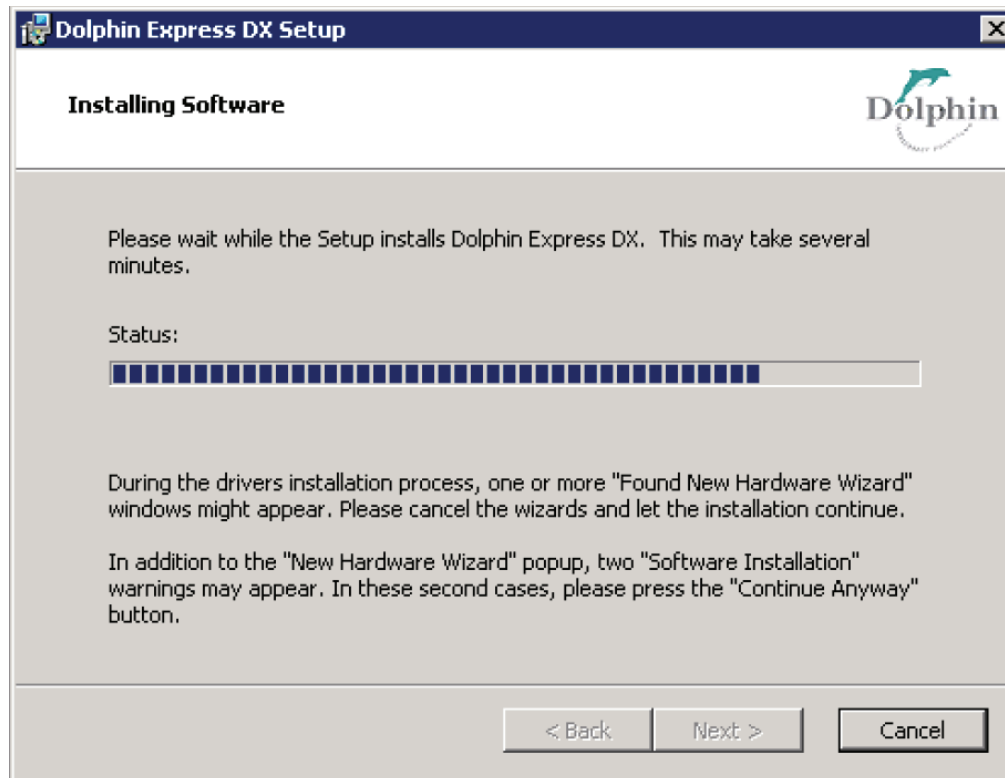
After selecting **Custom**, the **Custom Setup** dialog box will appear. De-select **ExpressWay** and **Manager**. To de-select these options, click on the down arrow located next to each one of these options and then choose “de-select” from the drop-down menu. A red “X” will then appear next to each of those options as shown in the following screen shot.

Figure 4.5. Windows Installer: Custom Setup



Click **Next**. Then agree to the two term agreements and select **Install**.

Figure 4.7. Windows Installer: Installing Software




Installing this new hardware will automatically cause some pop-ups to appear (Found New Hardware Wizard, for example). You do not need these in the Dolphin installation so please select **Cancel**.

Figure 4.8. Windows Installer: Cancel Wizard



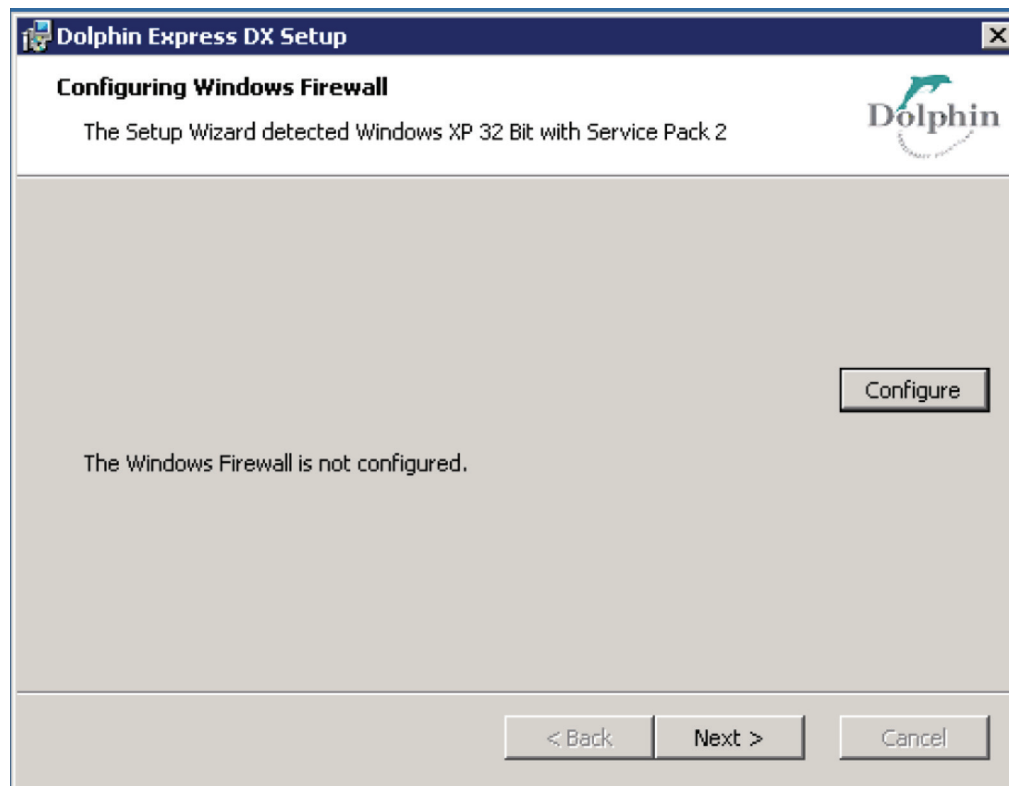
You may also see the following dialog box. We ask you to trust Dolphin's software even if they are not part of the Windows logo program. Select **Continue Anyway**

Figure 4.9. Windows Installer: Continue Anyway



If your machine has a firewall installed, then you need to open up connections between the modules in the Dolphin Admin (node manager, network manager and admin).

Figure 4.10. Windows Installer: Configure Firewall



Click **Next** to complete the installation.

Figure 4.11. Windows Installer: Completed



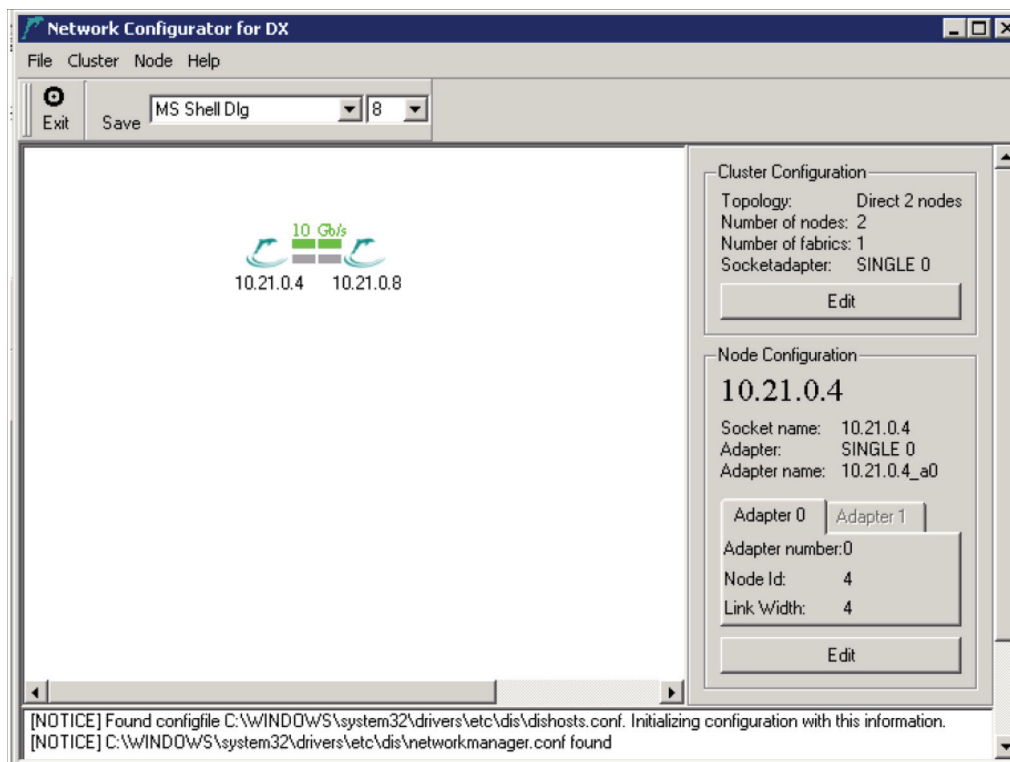
Click **Finish**.

After installation, navigate to the Dolphin Network Configurator via the Start menu on your Desktop (either your scope or your PC). The default path is **Start > All Programs > Dolphin Express DX > Administration > Configuration > dis_networkconfig**.

Your configuration parameters will already be filled in (and should look like the screen shot below). You will see a new ip address on each side of the graphic. To confirm, you can type "ipconfig /all" in a command prompt on your computer (most likely found at **Start > All Programs > Accessories > Command Prompt**) . On the **SCOPE** side, the address is 10.21.0.8. On the **HOST**, the ip address is 10.21.0.4. The next section in this *User's Guide* will discuss the settings in the Dolphin Network Configurator.

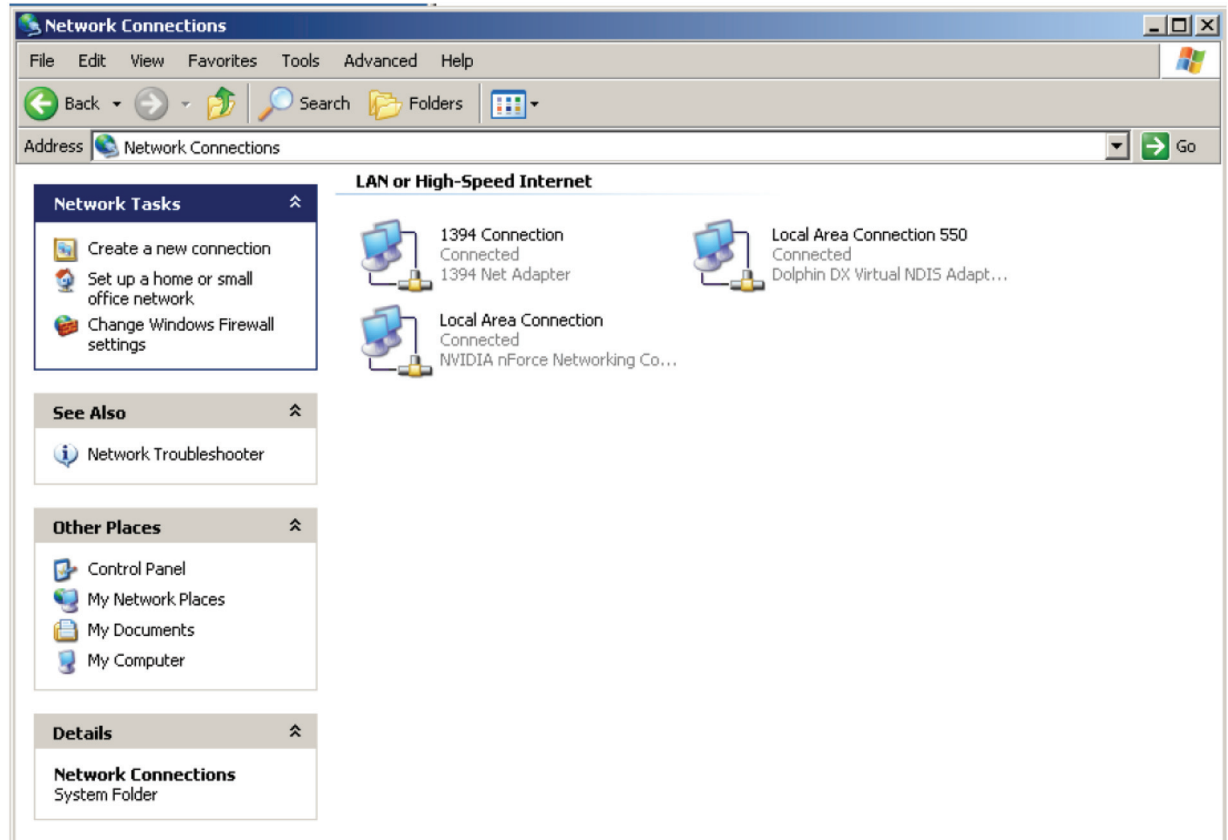
NOTE: You will not have to fill in the nodeids or launch the graphical configuration because this is done automatically.

Figure 4.12. Windows Installer: Configurator



If you navigate to **Start > Control Panel > Network Connections**, The Dolphin DX adapters will be shown (as in the following screen shot).

Figure 4.13. Windows Installer: Network Connection

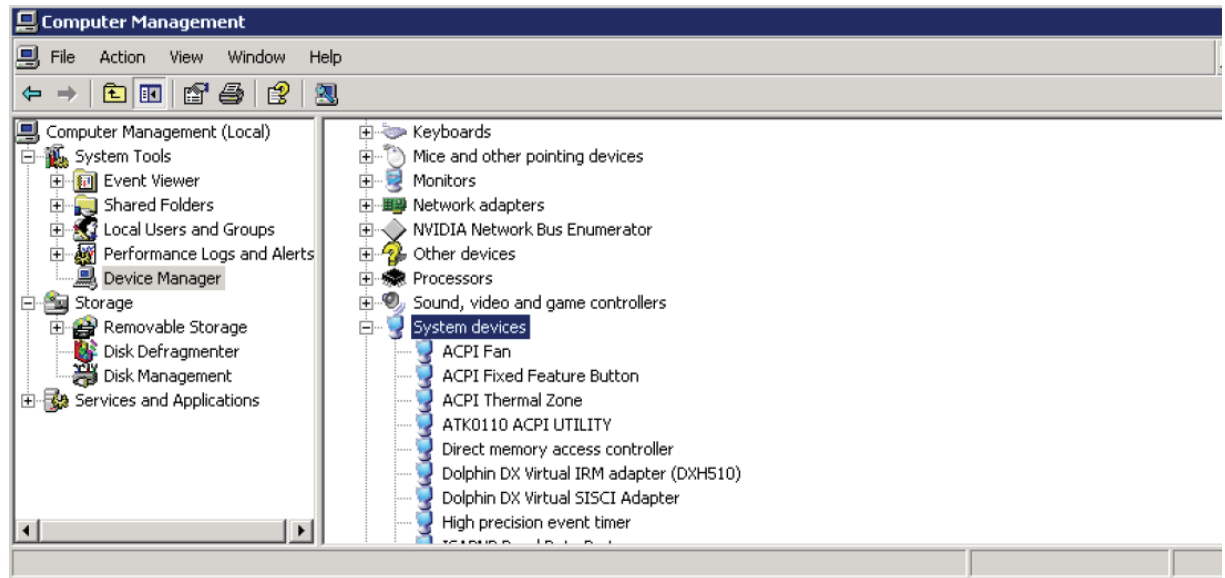


Warning

You should ignore all sections related to Dolphin Network Admin, node manager dis_nodemgr, network manager and dxadmin for the remainder of this document.

If the installation was successful the Dolphin Adapters will also appear in **Start > Control Panel > Administrative Tools > Computer Management > Device Manager** as shown in the screen shot below.

Figure 4.14. Windows Installer: Computer Management

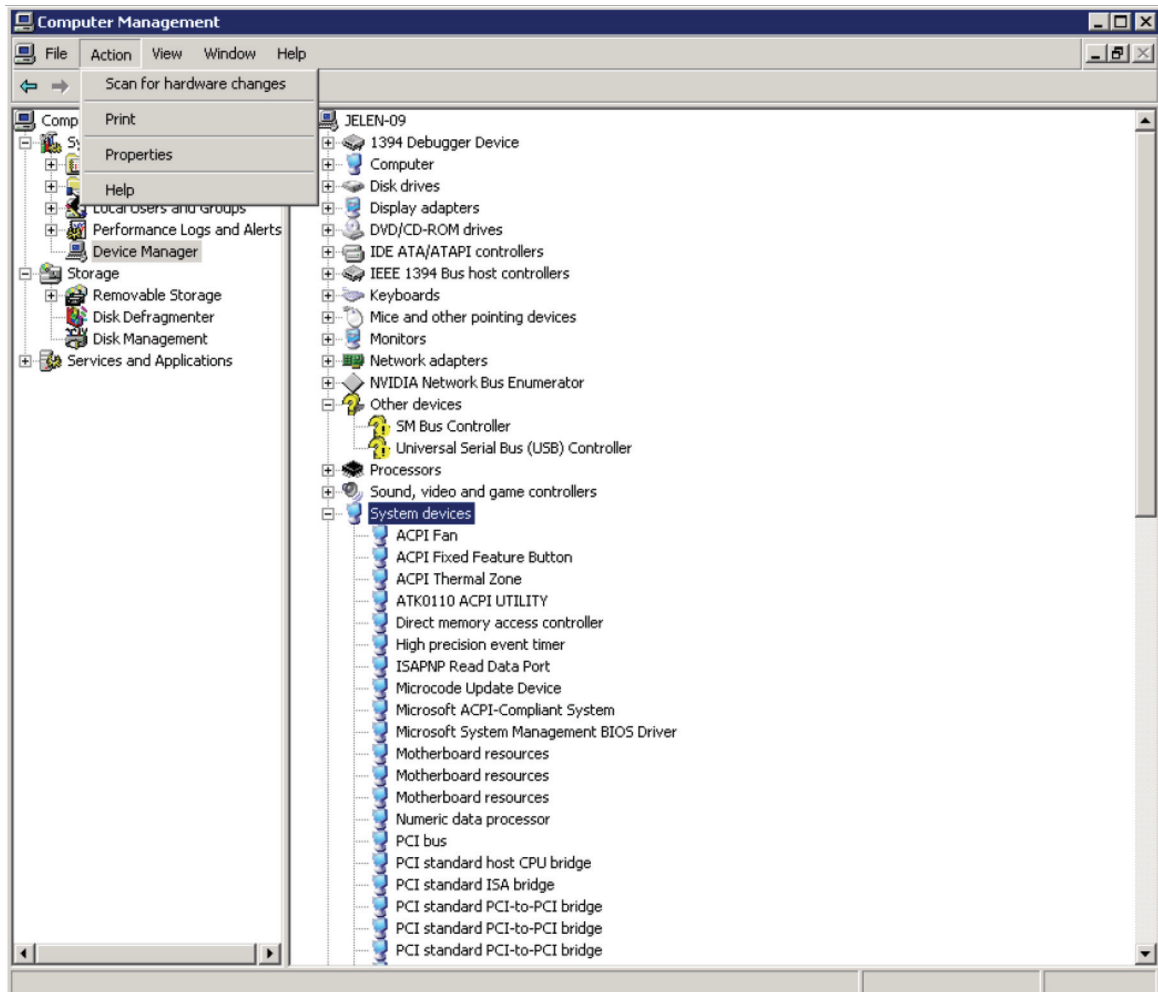


The Dolphin Network Configurator, `dis_netconfig` is only opened if the **Management** or **Graphical** tools are selected. Typically the Dolphin Network Manager, `dis_networkmgr`, is started when the you have exited the `dis_netconfig`.

You should now have the N4866A hardware and the Dolphin software installed on both your PC and oscilloscope. If you do not then repeat the previous steps until it is installed on both.

Whenever you uninstall and install again, the Device Manager needs a rescan of the devices before the new installation. You do not have to reboot the machine.

Figure 4.15. Rescan devices



4.2.b Commandline Installation

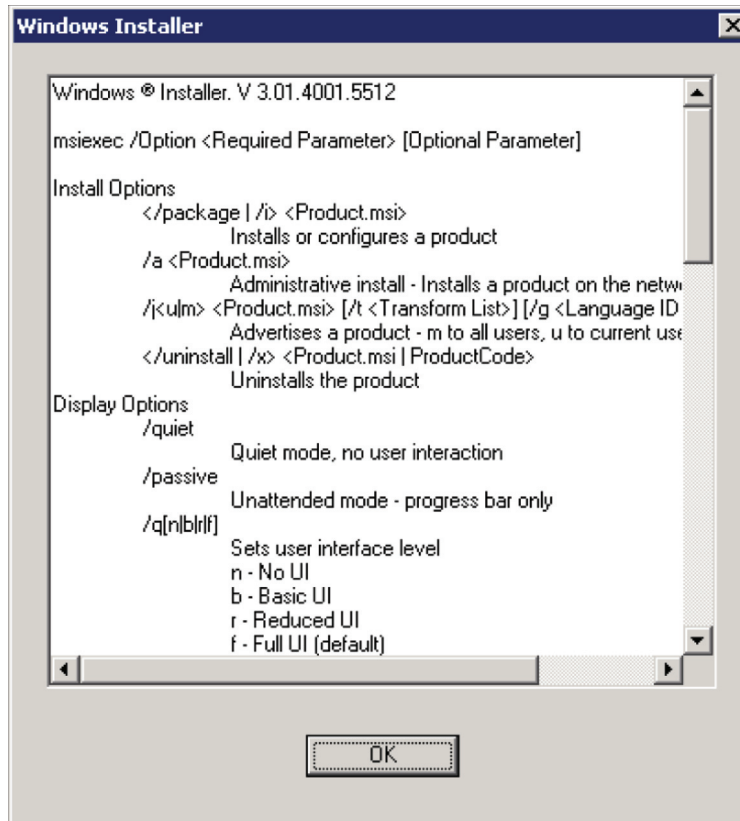
It is possible to run the Dolphin windows installer package from the commandline as well (only perform this operation if you are familiar with commandline installations in general. If you do not have experience with this then the previous installation method is recommended). Agilent SCOPE and HOST are two different MSI files. When you install from the commandline you use the **msiexec** command.

The commandline switch to **msiexec** is **/i** for install. If you also use the **/ADDLOCAL** switch you may select specific modules that you want to install. Running without the **/ADDLOCAL** switch installs all modules. These are the modules that you may preselect:

- Drivers
- Winsock2
- Manager
- Graphical
- Demos
- Development
- NMPI
- ExpressWay

Another useful switch combination to **msiexec** is **/qb** which lets you run without opening the windows installer in Section 3.1, “Running the Dolphin Installer”. Run **msiexec** without options for documentation.

Figure 4.16. msiexec help menu



Install Agilent SCOPE

On the SCOPE side users should typically select the components:

- Drivers
- Winsock2
- Graphical
- Demos
- Development

```
C:\Documents and Settings\tester\Desktop>msiexec /i Dolphin_DX_WinXP_X86_chk_Ver_3.2.1_SCOPE.msi ADDLOCAL
```

Dolphin_DX_WinXP_X86_chk_Ver_3.2.1_SCOPE.msi ADDLOCAL=The installation is not 100% silent. You need to perform the step described in Figure 4.9, “Windows Installer: Continue Anyway”.

On the HOST users would typically select:

- Drivers
- Winsock2
- Graphical
- Demos
- Development

```
C:\Documents and Settings\tester\Desktop>msiexec /i Dolphin_DX_WinXP_X86_chk_Ver_3.2.1_HOST.msi ADDLOCAL=Drivers
```

The installation is not 100% silent. You need to perform the step described in Figure 4.9, “Windows Installer: Continue Anyway” and Figure 4.8, “Windows Installer: Cancel Wizard”.

Uninstall

Whenever you uninstall and install again, the Device Manager needs a rescan of the devices before the new installation. You don't have to reboot the machine. If you want to uninstall you may use the /uninstall switch:

```
C:\Documents and Settings\tester\Desktop>msiexec /uninstall Dolphin_DX_WinXP_X86_chk_Ver_3.2.1_SCOPE.msi /
```

4.2.c Installation Results

After the installation, the nodes get installed and drivers and the node manager are started. Then, the basic packages are installed on the frontend, and the dis_netconfig application is launched to create the required configuration files c:\WINDOWS\system32\drivers\etc\dis\dishosts.conf and c:\WINDOWS\system32\drivers\etc\dis\networkmanager.conf if they do not already exist.

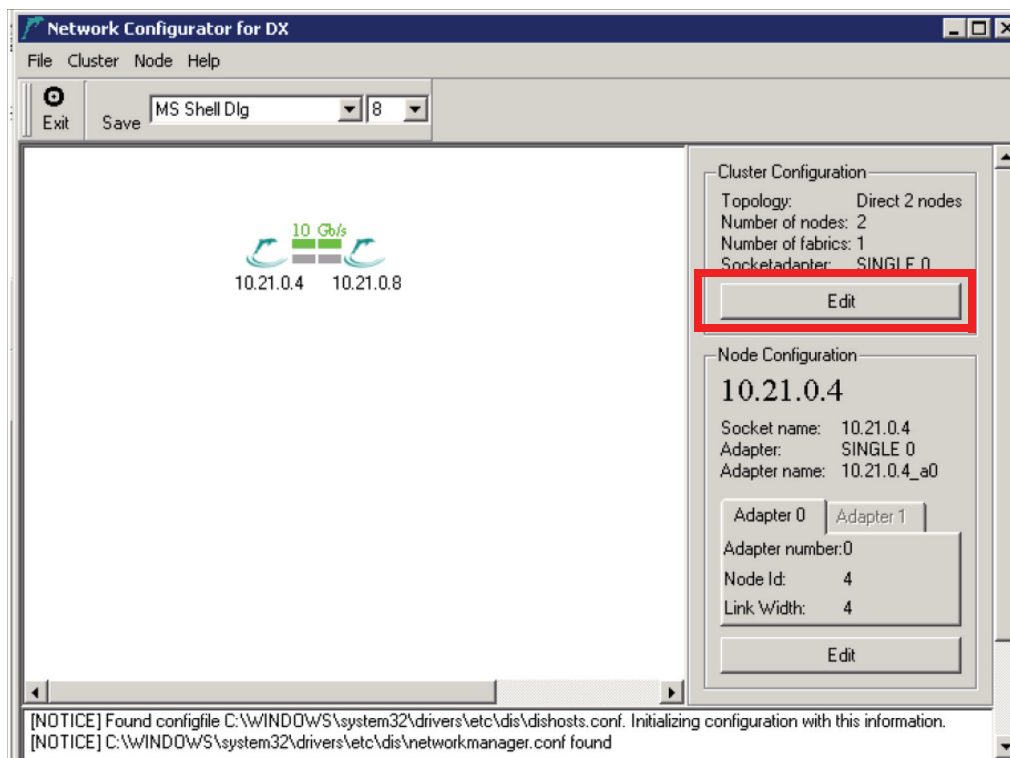
4.2.d Working with the Dolphin Network Configurator, dis_netconfig

The Dolphin Network Configurator, dis_netconfig is a GUI tool that helps gather the cluster configuration (and is used to create the cluster configuration file c:\WINDOWS\system32\drivers\etc\dis\dishosts.conf and the network manager configuration file c:\WINDOWS\system32\drivers\etc\dis\networkmanager.conf). A few global interconnect properties need to be set, and the position of each node within the interconnect topology needs to be specified.

To start the Dolphin Network Configurator, navigate to its default path of **Start > All Programs > Dolphin Express DX > Administration > Configuration > dis_networkconfig**.

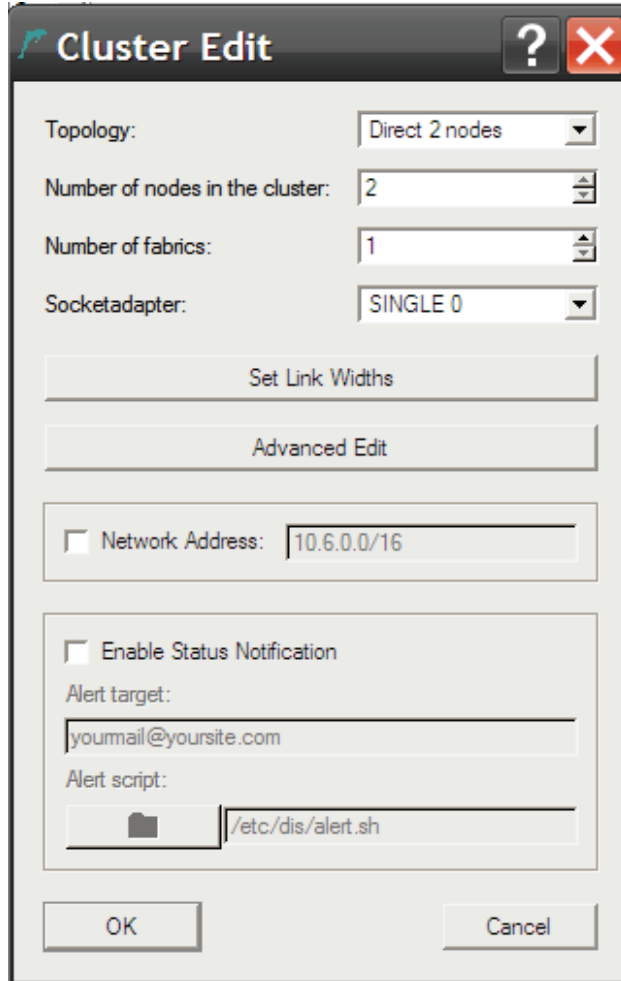
Cluster Edit

After the Dolphin Network Configurator (dis_netconfig) is launched, click the **Edit** button in the *Cluster Configuration* section.



The following dialog box will appear where the global interconnect properties need to be specified (see screen shot below).

Figure 4.17. Cluster Edit dialog of dis_netconfig



Interconnect Topology

In the upper half of the Cluster Edit dialog, you need to specify the interconnect topology that you will be using with your cluster. If `dis_netconfig` is launched by the installation script, the script tries to set these values correctly, but you need to verify the settings.

First, select the *Topology* of your cluster:

- Direct 2 Nodes*: This will be the correct setting for 95% of users. This is the instance where you have a single PC connected to a single oscilloscope.
- Switch (1-16 nodes)*: This setting enables you to use a single DXS switch for 2-10 nodes or two connected DXS switches for up to 16 nodes. This setting is for connecting up to 15 oscilloscopes to a single PC.
- Direct 3 Nodes*: This setting is for when you have two oscilloscopes connected to a single PCIe card in a PC (using both the 0 and 1 ports in the PCIe card).

Then, specify the *Number of nodes* in your cluster (in other words, the number of oscilloscopes + the one PC).

The *Number of fabrics* needs to be set to the minimum number of adapters in every node (typically, this value is 1).

The *Socketadapter* setting determines which of the available adapters are used for SuperSockets. If you look at the PCIe adapter, you will see that it actually has two ports (port 0 and port 1). This setting determines which one you are using.

- SINGLE 0: only adapter 0 is used
- SINGLE 1: only adapter 1 is used (only valid for more than one fabric)
- Channel Bonding: SuperSockets distributes the traffic across both adapters 0 and 1 (only valid for more than one fabric)
- NONE: SuperSockets should not be used.

You then need to *Set Link Widths* for each node. This must be set to x4 (connected with a single cable).

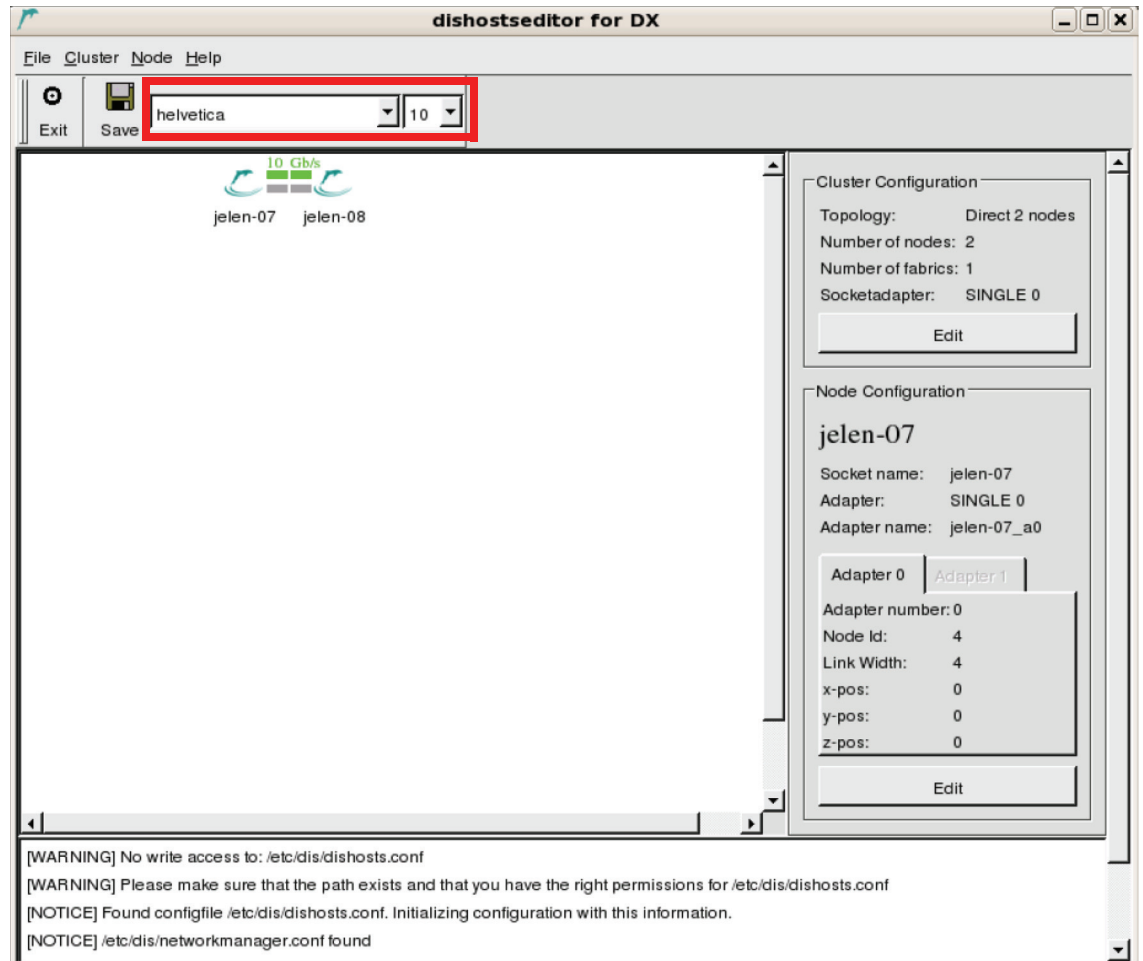
The *Advanced Edit* option does not need to be changed: the session between the nodes should typically always be set up automatically.

Node Arrangement

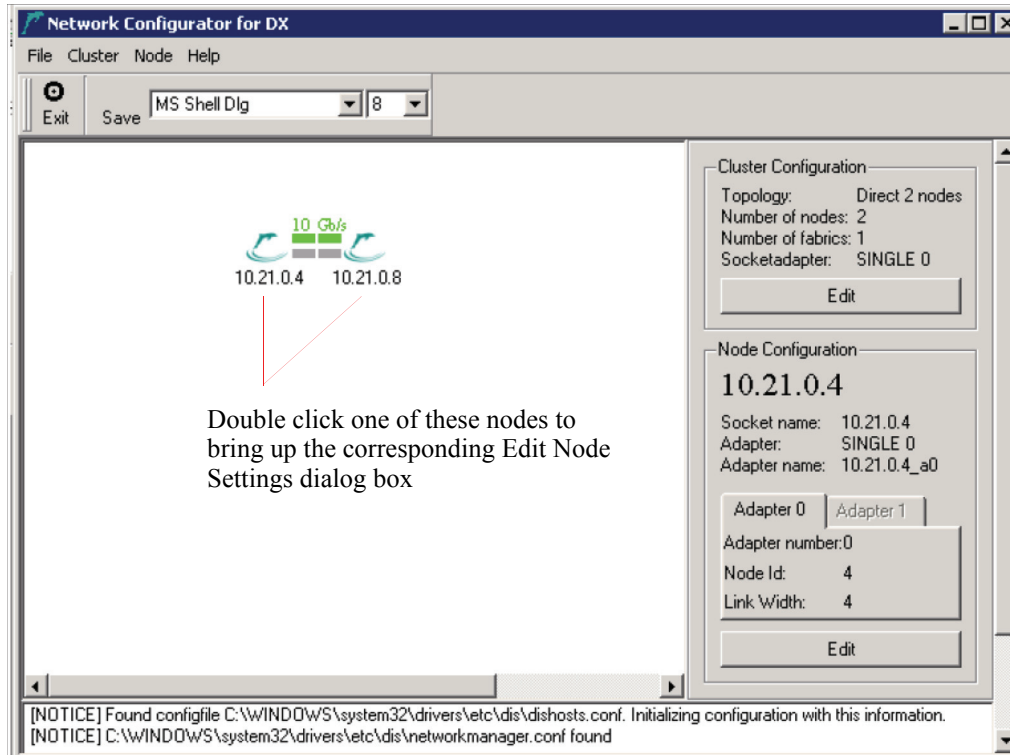
The nodes in the cluster will be arranged in the topology that was selected in the previous dialog (Cluster Configuration). To change this topology and other general interconnect settings, you can always go back and click **Edit** in the *Cluster Configuration* area which will bring up the **Cluster Edit** dialog again.

If the font settings of your X server cause `dis_netconfig` to print unreadable characters, you can change the font size and the type with the drop-down box at the top of the windows, next to the floppy disk icon.

Figure 4.18. Main dialog of `dis_netconfig`

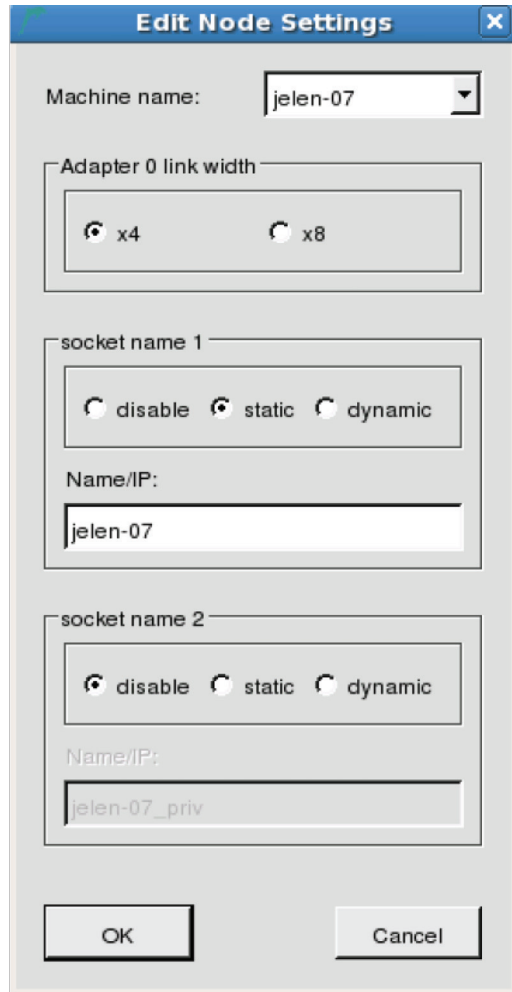


At this point, you need to arrange the nodes (marked by their hostnames) such that the placement of each node in the torus as shown by `dis_netconfig` matches its placement in the physical torus. You do this by assigning the correct hostname for each node by double-clicking its node icon which will open the configuration dialog of the node.



The following dialog box will appear:

Figure 4.19. Node dialog of dis_netconfig



In this dialog, select the correct **machine name** (which is the hostname as seen from the frontend) from the drop-down list. This will either be the ip address for the node or a name you assigned the node (using the **Name/IP** field in this dialog box).

In the node dialog, you need to specify that you want to use 4 lanes (by choosing **x4**). After you have assigned the correct hostname to this machine (should be set to **static**), you may need to configure SuperSockets on this node. If you selected the Network Address in the **Cluster Configuration** dialog, then SuperSockets will use this subnet address and will not allow for editing this property on the nodes. Otherwise, you can choose between 3 different options for each of the currently supported 2 SuperSockets-accelerated IP interfaces per node:

- **disable:** Do not use SuperSockets. If you set this option for both fields, SuperSockets can not be used with this node, although the related kernel modules will still be loaded.
- **static: (Recommended for the Agilent N4866A PCIe Interconnect)** Enter the hostname or IP address for which SuperSockets should be used. This hostname or IP address will be statically assigned to this physical node (its Agilent N4866A interconnect adapter). Choosing a *static* socket means that the mapping between the node (its adapters) and the specified hostname/IP address is static and will be specified within the configuration file `dishosts.conf`. All nodes will use this identical file (which is automatically distributed from the frontend to the nodes by the network manager) to perform this mapping. This option works fine if the nodes in your cluster do not change their IP addresses over

time and is recommend as it does not incur any name resolution overhead.

- **dynamic:** Enter the hostname or IP address for which SuperSockets should be used. This hostname or IP address will be dynamically resolved to the interconnect adapter that is installed in the machine with this hostname/IP address. SuperSockets will therefore resolve the mapping between adapters and hostnames/IP addresses dynamically. This incurs a certain initial overhead when the first connection between two nodes is set up and in some other specific cases. This option is similar to using a subnet but resolves only the explicitly specified IP addresses (for all nodes) and not all possible IP addresses of a subnet. Use this option if nodes change their IP addresses or node identities move between physical machines, i.e. in a fail-over setup.

Cabling Instructions

To create the cabling instruction, choose the menu item **File -> Create Cabling Instructions**. You can save and/or print the instructions. It is a good idea to print the instructions so you can take them with you to the cluster. It is important that you verify that the actual cable setup matches the topology you just specified.

4.2.e Cluster Cabling

If the cables are already connected, please proceed to the next section (*Verifying the Cabling*).

Note

In order to achieve the trouble-free operation of your cluster, it is critical that you set up the cables correctly. Please take your time to perform this task properly.

The cables can be installed while nodes are powered up.


Connecting the N4866A cable

Please proceed by connecting the nodes as described by the cabling instructions generated by the Dolphin Network Configurator (**File > Create Cabling Instructions**). Insert one or more cables into the connectors on the N4866A adapters. The connectors are labeled P0 or P1.

Note

Each of the ports on the N4866A adapters has a LED that should glow green if the port is connected.

Figure 4.20. Two cables plugged into the two ports on an adapter



Additional information can be found in the DXS410 - Dolphin DX Switch Quick Start Guide available from the Dolphin web site

Verifying the Cabling

NOTE

A green link LED indicates that the link between the output plug and input plug could be established and synchronized. It does *not* assure that the cable is actually placed correctly! It is therefore important to verify once more that the cables are plugged according to the cabling instructions generated by the Dolphin Network Configurator (`dis_netconfig`)!

If a pair of LEDs do not turn green, please perform the following steps:

- If the LEDs still do not turn green, use a different cable.
- If the LEDs still do not turn green, swap the cable of the problematic connection with a working one and observe if the problem moves with the cable.
- Power-cycle the nodes with the orange LEDs
- Contact Agilent support, www.agilent.com, if you cannot make the LEDs turn green after trying all proposed measures.

When you are done connecting the cables, no more user interaction is required. Please confirm that all cables are connected and all LEDs are green, and the installation will proceed. The network manager will be started on the frontend, configuring all cluster nodes according to the configuration specified in `dishosts.conf`.

4.2.f Handling Installation Problems

If you still fail to install the software successfully, refer to Chapter 5, *Interconnect Maintenance*. If you still fail to install the software successfully, you should contact Agilent support. Please provide all installation logfiles. To get the installation logs, please run each MSI with /l*xv log_install.txt switch: Please also include the configuration files that can be found in c:\WINDOWS\system32\drivers\etc\dis on the frontend.

4.2.g Making Cluster Application use Dolphin Express

After the N4866A hardware and Dolphin software has been installed and tested, you will want your cluster application to make use of the increased performance.

4.3 Installation Overview

4.3.a Installation Result

The result of the installation is that all necessary drivers and services are installed on the nodes and that the whole PCIe interconnect is configured and functional. By default, the installation target of all Dolphin software is %SYSTEMROOT%\Program Files\Dolphin Express DX.

In detail, the installation results on each node are the following services:

- dis_dx: This service loads the dis_dx module which manages access to the DX adapter hardware.

You can check the status with

```
:\Program Files\Dolphin Express DX\util>sc query dis_dx
```

- dis_irm: This service loads the dis_irm module which performs the configuration, session, and resource management.

```
:\Program Files\Dolphin Express DX\util>sc query dis_irm
```

- dis_sisci: This service loads the dis_sisci module which provides abstractions and mappings for the SISCI user-level API.

You can check the status with

```
:\Program Files\Dolphin Express DX\util>sc query dis_sisci
```

Next to this, a number of administration tools and test programs are installed into %SYSTEMROOT%\Program Files\Dolphin Express DX\util and %SYSTEMROOT%\Program Files\Dolphin Express DX/demo.

Monitoring and test tools are installed to %SYSTEMROOT%\Program Files\Dolphin Express DX\util.

Section 5. Interconnect Maintenance

This chapter describes how to perform a number of typical tasks related to the Agilent N4866A PCIe interconnect and Dolphin software.

5.1 Verifying Functionality and Performance

The following sections describe how to verify that the interconnect is setup correctly, which means that all nodes can communicate with all other nodes via the N4866A PCIe interconnect by sending low-level control packets and performing remote memory access.

5.1.a Availability of Drivers and Services

Without the required drivers and services running on all nodes and the frontend, the cluster will fail to operate. On the nodes, the kernel services `dis_irm` (low level hardware driver) and `dis_sisci` (upper level hardware services) need to be running. Next to these kernel drivers, the user-space service `dis_nodemgr` (node manager, which talks to the central network manager) needs to be active for configuration and monitoring. On the frontend, only the user-space service `dis_networkmgr` (the central network manager) needs to be running.

Because the drivers also appear as services, you can query their status with the usual tools of the installed operating system distribution.

```
> sc query dis_irm
Dolphin IRM 3.3.0 ( November 13th 2007 ) is running.
```

If any of the required services are not running, you will find more information on the problem that may have occurred in the system log facilities. Call **Start > Control Panel > Administrative Tools > Event Viewer** to inspect the kernel messages or check for related messages.

5.1.b Static Interconnect Test

The static interconnect test makes sure that all adapters are working correctly by performing a self-test, and determines if the setup of the routing in the adapters is correct (matches the actual hardware topology). It will also check if all cables are plugged in to the adapters, but this has already been done in the *Cable Connection Test*. The tool to perform this test is **dxdiag** (default location `c:\Program Files\Dolphin Express DX\util\dxdiag`).

Running **dxdiag** on a node will perform a self test on the *local* adapter(s) and list all remote adapters that this adapter can see via the N4866A PCIe interconnect. This means, to perform the static interconnect test on a full cluster, you will basically need to run **dxdiag** on each node and see if any problems with the adapter are reported and if the adapters in each node can see all remote adapters installed in the other nodes.

An example output of **dxdiag** for a node which is part of a 2 node cluster and using one adapter per node looks like this:

```
=====g=====
Dolphin diagnostic tool -- Dxdiag version 3.3.1d ( Mon Apr 21 2008 )
=====

***** VARIOUS INFORMATION *****

Dxdiag compiled in 64 bit mode
Driver : Dolphin IRM (DX)3.3.1d Dec 5th 2007 (rev unknown)
Date : Mon Apr 21 13:30:38 CEST 2008
System : Linux jelen-07 2.6.9-55.0.9.ELsmp #1 SMP Thu Sep 27 18:28:00 EDT 2007 x86_64 x86_64 x86_64 GNU/Linux

Number of configured local adapters found: 1

Hostbridge : , 0xbfff870

Local adapter 0 > Type           : DXH510
                   NodeId(log)   : 4
                   Serial number  : 0xad015107061
                   DXH chipId     : 0x01019902
                   DXH chip revision : 0x00000002 Assembly part revision : 0x01b1
                   Firmware version : 0x0157
                   SROM version    : 0x01a0
                   Switch type     : unknown
                   Topology type   : DX direct 2 nodes
                   PCIe link width : x8

OK: Link alive in adapter 0.
==> Local adapter 0 ok.

***** TOPOLOGY SEEN FROM ADAPTER 0 *****

Adapters found: 2 Switch ports found: 0
----- List of all ranges (rings) found:
In range 0: 0004 0008
-----

dxdiag discovered 0 note(s).
dxdiag discovered 0 warning(s).
dxdiag discovered 0 error(s).
TEST RESULT: *PASSED*
```

5.1.d Interconnect Load Test

While the static interconnect test sends very a few packets over the links to probe remote adapters, the *Interconnect Load Test* puts significant stress on the interconnect and observes if any data transmissions have to be retried due to link errors. This can happen if cables are not correctly connected (for example, plugged in without screws being tightened). Before running this test, make sure your cluster is cabled and configured correctly by running the tests described in the previous sections.

Test Execution from Command Line

To run this test from the command line, simply invoke **sciconntest** (default location /opt/DIS/bin/sciconntestc:\Program Files\Dolphin Express DX\demo\sciconntest) on all nodes.

Note

It is recommended to run this test from the sciadmin GUI because it will perform a more controlled variant of this test and give more helpful results.

All instances of **sciconntest** will connect and start to exchange data, which can take up to 30 seconds. The output of sciconntest on one node which is part of a 9-node cluster looks like this:

```
/opt/DIS/bin/sciconntest compiled Oct 2 2007 : 22:29:09
-----
Local node-id : 76
Local adapter no. : 0
Segment size : 8192
MinSize : 4
Time to run (sec) : 10
Idelay : 0
No Write : 0
Loopdelay : 0
Delay : 0
Bad : 0
Check : 0
Mcheck : 0
Max nodes : 256
rnl : 0
Callbacks : Yes
-----
Probing all nodes
Response from remote node 4
Response from remote node 8
Response from remote node 12
Response from remote node 68
Response from remote node 72
Response from remote node 132
Response from remote node 136
Response from remote node 140
Local segment (id=4, size=8192) is created.
Local segment (id=4, size=8192) is shared.
Local segment (id=8, size=8192) is created.
Local segment (id=8, size=8192) is shared.
Local segment (id=12, size=8192) is created.
Local segment (id=12, size=8192) is shared.
Local segment (id=68, size=8192) is created.
```

Local segment (id=68, size=8192) is shared.
Local segment (id=72, size=8192) is created.
Local segment (id=72, size=8192) is shared.
Local segment (id=132, size=8192) is created.
Local segment (id=132, size=8192) is shared.
Local segment (id=136, size=8192) is created.
Local segment (id=136, size=8192) is shared.
Local segment (id=140, size=8192) is created.
Local segment (id=140, size=8192) is shared.
Connecting to 8 nodes
Connect to remote segment, node 4
Remote segment on node 4 is connected.
Connect to remote segment, node 8
Remote segment on node 8 is connected.
Connect to remote segment, node 12
Remote segment on node 12 is connected.
Connect to remote segment, node 68
Remote segment on node 68 is connected.
Connect to remote segment, node 72
Remote segment on node 72 is connected.
Connect to remote segment, node 132
Remote segment on node 132 is connected.
Connect to remote segment, node 136
Remote segment on node 136 is connected.
Connect to remote segment, node 140
Remote segment on node 140 is connected.
SCICONNTEST_REPORT
NUM_TESTLOOPS_EXECUTED 1
NUM_NODES_FOUND 8
NUM_ERRORS_DETECTED 0
node 4 : Found
node 4 : Number of failures : 0
node 4 : Longest failure : 0.00 (ms)
node 8 : Found
node 8 : Number of failures : 0
node 8 : Longest failure : 0.00 (ms)
node 12 : Found
node 12 : Number of failures : 0
node 12 : Longest failure : 0.00 (ms)

```
node 68 : Found
node 68 : Number of failures : 0
node 68 : Longest failure : 0.00 (ms)
node 72 : Found
node 72 : Number of failures : 0
node 72 : Longest failure : 0.00 (ms)
node 132 : Found
node 132 : Number of failures : 0
node 132 : Longest failure : 0.00 (ms)
node 136 : Found
node 136 : Number of failures : 0
node 136 : Longest failure : 0.00 (ms)
node 140 : Found
node 140 : Number of failures : 0
node 140 : Longest failure : 0.00 (ms)
SCICONNTEST_REPORT_END

SCI_CB_DISCONNECT:Segment removed on the other node disconnecting.....
```

The test passes if all nodes report 0 failures for all remote nodes. If the test identifies any failures, you can determine the closest pair(s) of nodes for which these failures are reported and check the cabled connection between them. The numerical node identifies shown in this output are the node ID numbers of the adapters (which identify an adapter in the N4866A PCIe interconnect).

Although this test can be run while a system is in production, you have to take into account that performance of the productive applications will be reduced significantly while this test is running. If links actually show problems, they might be temporarily disabled, stopping all communication until rerouting takes place.

5.2 Replacing Interconnect Cables

The Agilent N4866A PCIe interconnect is fully hot-pluggable. If one or more cables need to be replaced (or just need to be disconnected and reconnected again for some reason), you can do this while all nodes are up and running. However, if the cluster is in production, you should proceed cable by cable, and not disconnect all affected cables at once to ensure continued operation without significant performance degradation.

If any of the verifications did report errors, make sure that the cable is plugged in cleanly and that the screws are secured. If the error should persist, swap the position of the cable with another one that is known to be working, and observe if the problem is wandering with the cable. If it does, the cable is likely to be bad. Otherwise, one of the adapters might have a problem.

5.3 Replacing an Adapter

In case an adapter needs to be replaced, proceed as follows:

- 1 Power down the node. The network manager will automatically reroute the interconnect traffic. When you run `sciadmin` on the frontend, you will see the icon of the node turn red within the GUI representation of the cluster nodes.
- 2 Unplug all cables from the adapter. Remember (or mark the cables) into which port on the adapter each cables belongs.
- 3 Unmount the old adapter and insert the new one into the node.
- 4 Power up the node. Then connect the cables in the same way they were connected before. Make sure all LEDs on all adapters in the affected ringlets light green again.
- 5 Verify the installation of the new adapter.

Communication between all other nodes will continue uninterrupted during this procedure.

Section 6. SuperSockets

This chapter explains how to make any application use SuperSockets, how SuperSockets can be configured, and how to verify SuperSockets usage and gather usage statistics.

6.1 Make Generic Applications use SuperSockets

All applications that use generic BSD sockets for communication will be accelerated by SuperSockets. No configuration change is required for the application as the same host names/IP addresses can be used. All relevant socket types are supported by SuperSockets (TCP stream sockets as well as UDP and RDS datagram sockets).

SuperSockets will use the Agilent N4866A PCIe interconnect for low-latency, high-bandwidth communication inside the cluster.

If the Dolphin Sockets DLL resides in the same directory and the configuration file dishosts.conf is valid then all calls are redirected to the low-latency SCI transport. Otherwise, the standard NDIS (TCP/IP protocol stack) transport is used.

6.2 SuperSockets Functionality and Performance

This section describes how to verify that SuperSockets are working correctly on a cluster.

6.2.a SuperSockets Status

You can check that SuperSockets have been configured correctly by running the `dis_ping` utility:

```
>c:\Program Files\Dolphin Express DX\ SuperSockets\bin\dis_ping 10.21.0.4
PCIe link is active
```

6.2.b SuperSockets Functionality

latency_bench

A benchmark that can be used to validate the functionality and performance of SuperSockets is `latency_bench`. The basic usage requires two machines (n1 and n2). Start the server process on node n1:

```
c:\Program Files\Dolphin Express DX\ SuperSockets\demo\>latency_bench -server
```

On node n2, run the client side of the benchmark:

```
c:\Program Files\Dolphin Express DX\SuperSockets\demo\>latency_bench -client n2
```

For an Agilent setup, the private IP-addresses 10.21.0.4 is used for HOST and 10.21.0.8 is used for SCOPE when using SuperSockets.

The output for a working setup should look like this:

```
C:\ Command Shell
C:\Program Files\Dolphin Express DX\SuperSockets\demo>latency_ben
10.21.0.4
Using address family 2
Testing PINGPONG transfers (back-and-forth)
# Size      Total time (usec)    RTT (usec)           Latency (usec)
  1         7354                7.35                 3.68 Min: 3.12
  2         7242                7.24                 3.62 Min: 3.12
  3         7259                7.26                 3.63 Min: 3.12
  4         7272                7.27                 3.64 Min: 3.15
  5         7566                7.57                 3.78 Min: 3.32
  6         7553                7.55                 3.78 Min: 3.36
  7         7582                7.58                 3.79 Min: 3.34
  8         7380                7.38                 3.69 Min: 3.23
  9         7547                7.55                 3.77 Min: 3.25
 10         7553                7.55                 3.78 Min: 3.32
 11         7562                7.56                 3.78 Min: 3.31
 12         7592                7.59                 3.80 Min: 3.31
 13         7617                7.62                 3.81 Min: 3.30
 14         7709                7.71                 3.85 Min: 3.36
 15         7693                7.69                 3.85 Min: 3.36
 16         7532                7.53                 3.77 Min: 3.29
 24         7618                7.62                 3.81 Min: 3.30
 32         7709                7.71                 3.85 Min: 3.39
 48         7888                7.89                 3.94 Min: 3.47
 64         8421                8.42                 4.21 Min: 3.75
 96         8722                8.72                 4.36 Min: 3.93
 128        9565                9.56                 4.78 Min: 4.26
 256       11446               11.45                5.72 Min: 5.23
 512       15338               15.34                7.67 Min: 7.16
1024       13886               13.89                6.94 Min: 6.55
2048       17652               17.65                8.83 Min: 8.37
4096       25220               25.22               12.61 Min: 12.0
8192       37256               37.26               18.63 Min: 18.0
16384      75879               75.88               37.94 Min: 34.3
32768     130836              130.84              65.42 Min: 62.5
65536     296281              296.28              148.14 Min: 137.

C:\Program Files\Dolphin Express DX\SuperSockets\demo>_
```

The latency in this example starts around 3.6 μ s. Recent machines with Dolphin Express DX deliver latencies at 2 μ s, and on older machines the latency may be higher. Latencies above 10 μ s indicate a problem. Typical Ethernet latencies start at 20 μ s or more.

Appendix A. Configuration Files

A.1 Cluster Configuration

A cluster with the N4866A PCIe interconnect requires one combined configuration file `dishosts.conf` for the interconnect topology and the SuperSockets acceleration of existing Ethernet networks, and another file `networkmanager.conf` that contains the basic options for the mandatory Network Manager. Both of these files should be created using the GUI tool `dis_netconfig`. Using this tool vastly reduces the risk of creating an incorrect configuration file.

A.1.a `dishosts.conf`

The file `dishosts.conf` is used as a specification of the N4866A PCIe interconnect (in a way just like `/etc/hosts` specifies nodes on a plain IP based network). It is a system wide configuration file and should be located with its full path on all nodes at `%SYSTEMROOT%\system32\etc\dis\dishosts.conf` (`%SYSTEMROOT%` is where your windows installation resides, e.g. `C:\WINDOWS.`).

A syntactical and semantic validation of `dishosts.conf` can be done with the tool **testdishosts**.

The Dolphin network manager and diagnostic tools will always assume that the current file `dishosts.conf` is valid. If dynamic information read from the network contradicts the information read in the `dishosts.conf` file, Dolphin network manager and diagnostic tools will assume that components are mis configured, faulty, or removed for repair.

`dishosts.conf` is by default automatically distributed to all nodes in the cluster when the Dolphin network management software is started. Therefore, edit and maintain this file on the frontend only. You should create and maintain `dishosts.conf` by using the Dolphin Network Configurator, `dis_netconfig` GUI: `c:\Program Files\Dolphin Express Dx\util\dis_netconfig` or the command line version of the Dolphin Network Configurator, `dis_mkconf`: `c:\Program Files\Dolphin Express Dx\util\dis_mkconf`

Normally, there is no reason to edit this file manually. To make changes in `dishosts.conf` effective, the network manager needs to be restarted from the start menu.

The following sections describe the keywords used.

Basic settings

#SCI-HARDWARE

The first line of this file should contain an entry specifying what hardware you are going to use. If the `#SCI-HARDWARE` is specified, then Dolphin Express SCI is expected and the parsing of this file will fail if the hardware that you intended to use is Dolphin Express DX.

#DX-HARDWARE

The first line of this file should contain an entry specifying what hardware you are going to use. If the `#DXHARDWARE` is specified, then Dolphin Express DX is expected and the parsing of this file will fail if the hardware that you intended to use is Dolphin Express SCI.:

DISHOSTVERSION [0 | 1 | 2 | 3]

The version number of the `dishosts.conf` is specified after the keyword `DISHOSTVERSION`. The `DISHOSTVERSION` should be put on the first line of the `dishosts` file that is not a comment. `DISHOSTVERSION 0` designates a very simple configuration. `DISHOSTVERSION > 0` maps hosts/IPs to adapters, which in turn are mapped to nodeids and physical adapter numbers by means of the `ADAPTER` entries. `DISHOSTVERSION 1` or higher is required for running with multiple adapter cards and transparent fail over etc. `DISHOSTVERSION 2` provides support for dynamic IP-to-nodeId mappings (sometimes also referred to as virtual IPs). `DISHOSTVERSION 3` is for DX adapters and requires that you also specify the link width for an adapter.

HOSTNAME: <hostname/IP>

Each cluster node is assigned a unique dishostname, which has to be equal to its hostname. The hostname is typically the network name (as specified in hosts) or the node's IP-address. Examples:

```
HOSTNAME: host1.dolphinics.no
HOSTNAME: 193.69.165.21
```

ADAPTER: <physical adaptername> <nodeId> <adapterNo> <link width>

A network node may hold several physical adapters. Information about a node's physical adapters is listed right below the hostname. All nodes specified by a HOSTNAME need at least one physical adapter. This physical adapter has to be specified on the next line after the HOSTNAME. The physical adapters are associated with the keyword ADAPTER. The link width is PCI Express link width x4.

```
#Keyword  name          nodeid  adapter  link width
ADAPTER: host1_a0      4       0         4
ADAPTER: host1_a1      4       1         4
```

STRIPE: <virtual adaptername> <physical adaptername 1> <physical adaptername 2>

Defines a virtual striping adapter comprising two physical adapters, which will be used for automatic data striping (also referred to as channel bonding). Striping adapters will also be used as redundant adapters in the case of network failure.

```
STRIPE:      host1_s  host1_a0  host1_a1
```

REDUNDANT: <virtual adaptername> <physical adaptername 1> <physical adaptername 2>

Defines a virtual redundant adapter comprising two physical adapters, which will be used for automatic fail over in case one of the fabrics fails.

```
REDUNDANT:   host1_r  host1_a0  host1_a1
```

SuperSockets settings

The SuperSockets configuration is responsible for mapping certain IP addresses to N4866A adapters. It defines which network interfaces are enabled for Dolphin SuperSockets.

SOCKET: <host/IP> <physical or virtual adaptername>

Enables the given <host/IP> for SuperSockets using the specified adapter.

In the following example we assume host1 and host2 have two network interfaces each designated host1pub, host1prv, host2pub, host2prv, but only the 'private' interfaces hostXprv are enabled for SuperSockets using a striping adapter:

```
SOCKET: host1prv host1_s
SOCKET: host2prv host2_s
```

Starting with DISHOSTVERSION 2 SuperSockets can handle dynamic IP-to-nodeId mappings. In other words, a certain IP address does not need to be bound to a fixed machine but can roam in a pool of machines. The address resolution is done at runtime. For such a configuration a new type of adapter must be specified:

SOCKETADAPTER: <socket adaptername> [SINGLE | STRIPE | REDUNDANT] <adapterNo> [<adapterNo> ...]

This keyword basically only defines an adapter number, which is not associated to any nodeId.

Example:

```
SOCKETADAPTER: sockad_s STRIPE 0 1
```

Defines socket adapter "sockad_s" in striping mode using physical adapters 0 and 1. The resulting internal adapter number is 0x2003.

Such socket adapters can now be used in order to define dynamic mappings, and, in extension to DISHOSTVERSION 1 whole networks can be specified for dynamic mappings:

SOCKET: [<ip> | <hostname> | <network/mask_bits>] <socket adapter>

Enables the given address/network for SuperSockets and associates it with a socket adapter.

It is possible to mix dynamic and static mappings, but there must be no conflicting entries.

Example:

```
SOCKET: host1 sockad_s
```

```
SOCKET: host2 sockad_s
```

```
SOCKET: host3 host3_s
```

```
SOCKET: 192.168.10.0/24 sockad_s
```

Miscellaneous Notes

- Using *multiple nodeids* per node is supported. This can be used for some advanced high-availability switch-configurations.

- A short version of dishosts.conf is supported for compatibility reasons, which corresponds to DISHOSTVERSION 0. Please note that neither virtual adapters nor dynamic socket mappings are supported by this format.

Example:

```
#host/IP nodeid  
host1prv 4  
host2prv 8
```

A.1.b networkmanager.conf

The networkmanager.conf specifies the startup parameters for the Dolphin Network Manager. It is created by the dis_netconfig.

A.1.c cluster.conf

This file must not be edited by the user. It is a configuration file of the Network Manager that consists of the user-specified settings from networkmanager.conf and derived settings of the cluster (nodes). It is created by the Network Manager.

Appendix B. DX Firmware upgrade

The Agilent N4866A hardware supports software based firmware upgrade. A firmware upgrade is normally not needed unless a firmware upgrade is suggested by the dxdiag or driver installation tools. Please contact Dolphin for more information if you are requesting a firmware upgrade.