ZNYX

RAINlink™ for Solaris

User's Guide

Table of contents

RAINlink – Network Link Services Software

The Solaris driver for the ZNYX NetBlaster™ PCI and CompactPCI includes embedded RAINlink technology. RAINlink software provides critical link services such as link aggregation (trunking), link failover and dynamic load balancing over multiple network ports. These ports can be a group of multiple single-channel or multi-channel adapter ports, which are limited by the number of PCI slots of a host system. The link aggregation service aggregates multiple network links into one or more trunks to provide scalable bandwidth, whereas the link failover service offers a transparent failover to a redundant network link (that is part of a designated failover group) in case of a port, cable or adapter failure.

RAINlink for Solaris has been designed for the NetBlaster family of ZNYX PCI and CompactPCI network interface cards and operates with industry standard internet protocols.

RAINlink can be deployed on a variety of platforms including web, intranet, back-up, network, database and workgroup servers; engineering, scientific modeling and graphical publication workstations; and desktops.

Table 1 highlights the feature of the RAINlink for Solaris software.

Features and Services	**RAINlink for Solaris**	**Comments**
Platform Support	Sparc or Intel	
Operating System	Solaris 2.6, Solaris 7	Solaris 7 on sparc supports either 32-bit or 64-bit
Maximum ports	16	Physical network links
Maximum Failover Groups	8	System-to-system, System-to-FEC switch, System-to-non-FEC switch, and System-to-hub
Maximum Link Aggregation Groups (Trunks)	8	Supports Cisco Fast EtherChannel technology; System-to-system trunking
Dynamic Load Balancing (Layer 2, 3)	Yes	Across aggregated links
Distribution	Included with the ZNYX Solaris NetBlaster driver .Available for FREE download from www.znyx.com	
Technical Support	Yes	

Table 1. RAINlink for Solaris features

Features and Benefits

RAINlink for Solaris makes your network faster and more dependable by providing scalability, reliability and flexibility to network links. When your current network reaches its capacity, RAINlink can incrementally add more performance and reliability—all without "forklift upgrades" to a new topology. The following section describes all the features and benefits of RAINlink. Refer to the "Network Configurations" section to determine which features and benefits of RAINlink you can best exploit to enhance your network.

Link Aggregation

The link aggregation service aggregates multiple *active* LAN links (ports) and adapters into one or more link aggregation groups (often called trunks) to provide scalable bandwidth, while offering

increased fault tolerance. The system-to-switch link aggregation service of RAINlink is compatible with Cisco's Fast EtherChannel (FEC) technology.

As shown in Figure 1, unlike a link failover service, RAINlink's link aggregation service maintains all links of a trunk *active* at a time, while dynamically balancing the traffic across these links. Link aggregation provides resiliency, fast failover and convergence to your network links. If one of the aggregated links is unplugged, damaged, or fails, RAINlink automatically removes the failed link from the link aggregation group and transfers (failover) the traffic across remaining links while still maintaining load balancing. When a lost link is re-established, RAINlink detects and resumes traffic flow on that link. The seamless failover and convergence do not require user intervention and are transparent to the end user -- no host protocol timers expire, therefore no sessions are dropped.

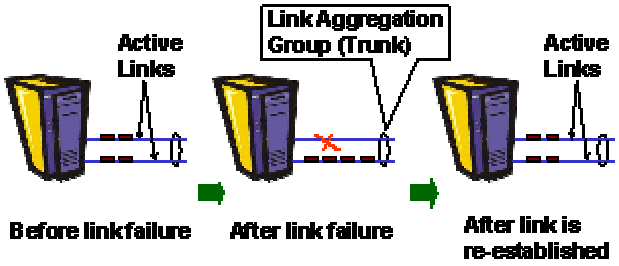

Figure 1. *Active* redundancy using Link Aggregation

Every network needs to handle increased traffic. Surveys show local backbone traffic growing at an annual rate of 40%. Instead of replacing your entire network infrastructure, RAINlink offers a more intelligent alternative. Just add more network interface cards and RAINlink will dramatically increase available network bandwidth with link aggregation service. As Figure 2 illustrates, by incrementally adding a full-duplex Fast Ethernet link in a trunk, the available bandwidth can scale from 200 Mbps per link to 400 Mbps per 2 links (i.e. 400 Mbps per trunk).

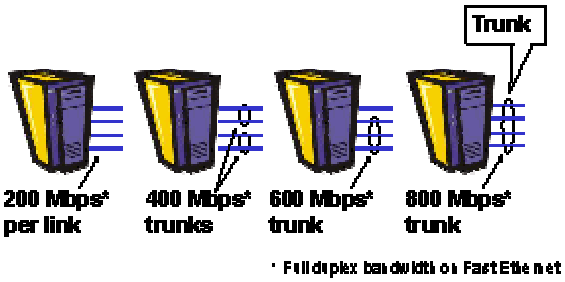

Figure 2. Scalable bandwidth using Link Aggregation

The dynamic load balancing service of RAINlink distributes two-way traffic evenly across aggregated links. As shown in Figure 3, instead of saturating a single link, RAINlink can intelligently distribute unicast, broadcast and multicast traffic across all links in a trunk. This maximizes the host systems' throughput and improves network efficiency.

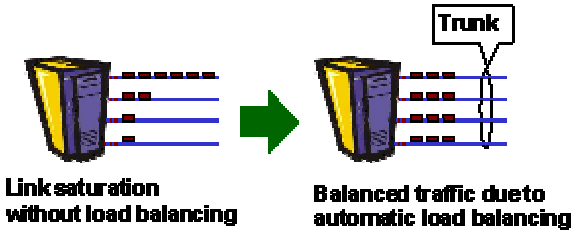

Figure 3. Dynamic Load Balancing across aggregated links

Cisco Systems Fast EtherChannel

Cisco's Fast EtherChannel (FEC) technology provides an alternative solution for implementing higher speeds in selected portions of a network. This technology aggregates parallel, 100-Mbps Fast Ethernet links to deliver high bandwidth between two end points. These point-to-point links effectively create a scalable "fat pipe (trunk)," with full-duplex bandwidth combined in increments of 200 Mbps. Cisco's Fast EtherChannel technology uses the existing Fast Ethernet physical and Media Access Control (MAC) layers and is compatible with existing equipment and cabling. In addition to creating a fat pipe, Fast EtherChannel adds functionality for load sharing and fault tolerance so that in case of a failure on one link, the remaining load will automatically be redirected across the remaining aggregated links.

Link Failover

The link failover service detects an active network link failure and transfers traffic to a redundant stand-by link, which is part of a failover group. When a failed link re-establishes, RAINlink automatically detects it and designates it as a stand-by link. The new stand-by link now becomes the backup and is ready to take over traffic if another failure occurs on an active link. As shown in Figure 4, unlike a trunk, a failover group has only one link active at a time. The other redundant links of a failover group remain in a standby mode. A redundant link includes a redundant adapter (or adapter port), redundant cable, and redundant switch (or hub or router) port. A redundant link can also be made up of trunks. Note that, unlike trunking, the link failover service does not offer dynamic load balancing scheme across links, which are part of a failover group.

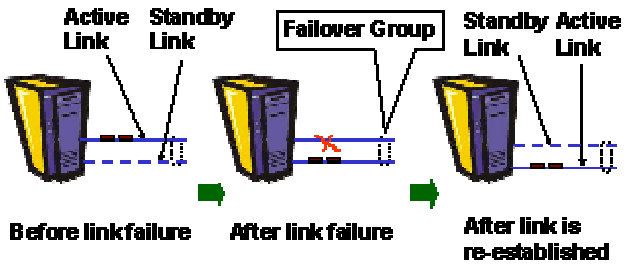

Figure 4. Reliable network connections using Link Failover

Fast Failover

By default, RAINlink uses the fast failover mode. In this mode, RAINlink moves the traffic over to a redundant stand-by link in case of a link failure in as little as 500 milliseconds. RAINlink's fast failover eliminates the need to reconfigure the Solaris network protocol stack and re-learning of layer 2 MAC addresses by clients to reestablish a network link after a failure. This sub-second failover provides network link recovery from cable, port, or adapter faults in a manner completely transparent to applications and end-users.

Timeout Mode

Typically, a failover group would only switch from one link to another if the physical link were lost. By enabling timeout mode, if no traffic is received in the specified time interval, the active port is automatically switched to another available link. The timeout mode time interval is settable by the user.

Dynamic Load Balancing

The dynamic load balancing service of RAINlink evenly distributes the two-way traffic across links, which are part of a trunk. By automatically distributing unicast, broadcast and multicast traffic evenly across aggregated links, RAINlink prevents link saturation, maximizes throughput and fully utilizes the invested resources. Choose the type of load balancing scheme from the following types, which is suitable to your needs, whether it be for system-to-switch, or system-to-system.

Layer 2 Protocol (System-to-Switch)

This method of load balancing is based on the source and/or destination MAC addresses. For example, as shown in Figure 5, Server S is connected to Switch X by Fast EtherChannel trunk and Clients A and B are connected to X.

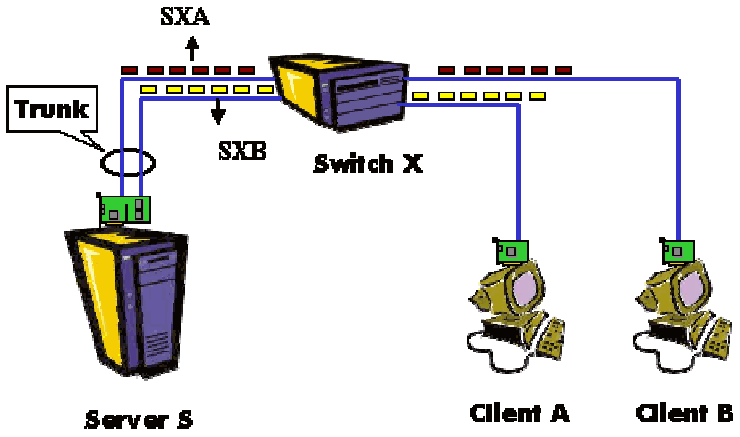

Figure 5. Layer 2 Load Balancing

This configuration has two data streams SXA and SXB. In this case, RAINlink distributes multiple data streams (SXA and SXB) over the individual links that make up the trunk to get an overall load distribution. Two component links of a trunk will carry two different data streams. Note that in a system-to-switch configuration, RAINlink does not distribute packets of a single data stream over multiple links that make up the trunk. Layer 2 mode is only valid for system-to-switch trunking

Layer 3 Protocol (System-to-Switch)

This method of load balancing is based on the source and/or destination IP addresses. . By default, RAINlink uses Layer 3 protocols for dynamic load balancing. For traffic on the same network, this balances the same as Layer 2. Layer 3 becomes necessary in configurations containing a router between endpoints. Assume the server traverses a router to get to Client A or Client B. A data stream SXA or SXB would contain the IP address (Layer 3) of A or B, but the MAC address (Layer 2) of the router. All traffic between the system and the switch would appear the same at Layer 2, and therefore would transmit across a single link. By looking at the Layer 3 information, the traffic appears as two separate data streams, and can be distributed across multiple links between the system and the switch. For non-IP protocols, Layer 2 mode is implemented by default. Layer 3 mode is only valid for system-to-switch trunking.

IP Trunking (System-to-System)

This method of load balancing is appropriate for two systems directly connected to each other without a switch. Figure 6 is an example of a trunk used to create a high-speed link between two servers.

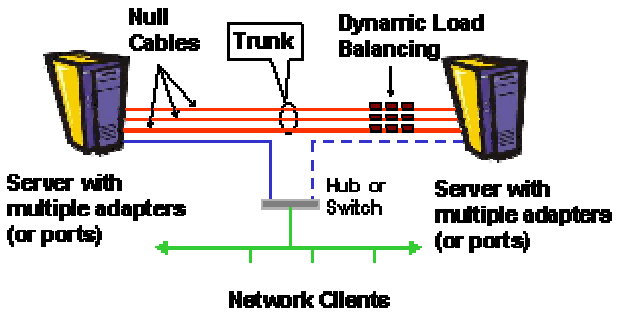

Figure 6. High-speed pipe from System-to-System using Link Aggregation

The above network configuration is commonly used to implement higher bandwidth connections between enterprise-class intranet web servers, video servers, file servers or workstations. Packets are sent to the driver as large datagrams and fragmented across the active links. Although the packets are conformant IP packets, the receiving system must also be configured for system-to-system trunking with IP trunking mode enabled to guarantee proper reception and reassembly.

Disabling Dynamic Balance Mode
In Layer 2 or Layer 3 mode, balance mode is on by default. Balance mode causes the ports of the trunk to be dynamically load balanced. For example, in Layer 2 mode, if a certain set of MAC addresses were causing heavy traffic to be sent over a single port, enabling Balance Mode would attempt to split up the addresses across multiple ports. This mode is only valid for system-to-switch trunking. It cannot be used with system-to-system IP trunking Mode.

Setting up RAINlink for Solaris

Systems Requirements

The following requirements must be met to install and use RAINlink for Solaris in your network.

Platforms: Sun Sparc system or Intel-based PC with standard PCI or CompactPCI 2.1 compliant slots.
Operating Systems: Solaris 2.6 (SunOS 5.6), Solaris x86 2.6 (SunOS 5.6), Solaris 7 (SunOS 5.7), Solaris x86 7 (SunOS 5.7). Solaris 7 for Sparc supports either 32-bit or 64-bit. No SunOS Driver Updates are required
Network Adapter: Any combination of ZNYX NetBlaster PCI or CompactPCI adapters (ZX345, ZX345Q, ZX346, ZX346Q, ZX348. ZX348Q, ZX412, ZX414)
Driver: NetBlaster version 2.0 or later

Pre-Installation Procedure

Before installing RAINlink for Solaris, you must have at least one multi-port or two single-port ZNYX NetBlaster network interface cards (NICs) and its driver installed in your computer. Typically, the RAINlink pre-installation procedure involves the following:

1. Make sure you have Solaris installed on your host computer system.
2. Properly seat the NetBlaster adapter in the computer system.
3. Connect the network cables to appropriate adapter ports.

IMPORTANT: If you are deploying RAINlink in a network configuration where both end stations are NICs (i.e. a system-to-system configuration), you must use "null" (crossover) cables to connect two end stations. In all other network configurations, you must use "straight" cables. For more information on cabling, see the Appendix section, "Network Cabling Reference."

4. Install the driver for the NetBlaster adapters. The Solaris NetBlaster driver must be at least Version 2.0. The NetBlaster drivers and Installation instructions are available from the ZNYX website at http://www.znyx.com.

Installing RAINlink for Solaris

All the necessary installation and configuration programs for RAINlink for Solaris are placed on the system during the installation of the NetBlaster driver. The NetBlaster driver must be version V2.0 or later.

1. Verify that you have the correct version of the NetBlaster driver installed. Enter:

   ```
   pkginfo –l ZNYXnb
   ```

   The version should be at least V2.0 and the status should be "completely installed". If not, remove the package using *pkgrm*, and reinstall the package with *pkgadd*. The V2.0 NetBlaster driver and complete instructions for installation and removal are available from the ZNYX website: http://www.znyx.com.

2. RAINlink for Solaris includes utilities and example configuration files for standard configurations in order to get you up and running quickly. The necessary files are placed onto your system during the NetBlaster V2.0 driver installation in */etc/rain*. Change directories to */etc/rain* to begin the installation:

   ```
   cd /etc/rain
   ```

3. An installation script is included with RAINlink for Solaris, called *rainlink*. The *rainlink* script can be used to activate, deactivate, or display the current RAINlink for Solaris configuration. The first time the *rainlink* script is run, the script will give you a list of default configurations to choose from suitable for most basic configurations of two or four port trunks or failover groups. Detailed configuration instructions are included in the next section. Start the *rainlink* script by typing *rainlink -a* at the command prompt:

   ```
   rainlink -a
   ```

4. The script dialogue will allow you to choose from a set of default configurations.

```
RAINlink

This script initializes or reinitializes the RAINlink layer within the ZNYX
NetBlaster driver.  An input script located at /etc/rain/rainlink.conf is used
as input to ZNYX's rlconfig application.  Either choose one of the following
prepared input scripts which will be copied to /etc/rain/rainlink.conf,
or exit and create your own custom script.

    1 - System-to-System trunking between 4 ports.
    2 - Fast EtherChannel trunking between 4 ports.
    3 - Fast Failover between 4 ports.
    4 - System-to-System trunking between 2 ports.
    5 - Fast EtherChannel trunking between 2 ports.
    6 - Fast Failover between 2 ports.
    7 - Exit and create custom script.

        Choose one of the above 7 choices(1-7):rainlink -a
```

5. Make your selection and continue. A note will be displayed informing you that the *rainlink* script will create a configuration file named, */etc/rain/rainlink.conf*, which will be the default configuration file used from this point forward. Details of the contents of the configuration file are provided later. Hit Enter to continue:

```
NOTE:  /etc/rain/rainlink.conf has been created.  If a rainlink.conf file
       exists rainlink will not prompt for one of the previous 7 choices
       To modify the RAINlink configuration in the future, edit
       /etc/rain/rainlink.conf, and either run rainlink again, or reboot the
       system.

Hit any key to continue:
Bringing down all Net Blaster interfaces now....
Running rlconfig ...
#####TRUNK0 Successfully Created
TRUNK0 mode set
```

6. The next series of questions relate to bringing up the interfaces. Next, bring up the RAINlink interfaces, and then the remaining non-RAIN interfaces if there are any. You will need to supply IP names or addresses for each configured interface you bring up.

```
Do you want to configure and bring up the RAINlink interfaces [y,n,?] y
Enter hostname for interface zrl0 [barbrady-zrl0]: server1

The following host names must be in the hosts database.
Check with your Network Administrator for more information.

server1

Do you want to configure and bring up any non-RAINlink interfaces [y,n,?] y

Bring znb4 up [y,n,?] y
Enter hostname for interface znb4 [barbrady-znb4]: access1

The following host names must be in the hosts database.
Check with your Network Administrator for more information.

access1

Bringing up available boards now....
```

The script then exits. Your RAINlink and non-RAINlink interfaces should now be configured.

7. You can use *rainlink –s* to verify the status of your RAINlink configuration:

```
rainlink –s
```

```
trunk0<LAYER_3_MODE> = znb0<UP>, znb1<UP>, znb2<UP>,
    znb3<UP>, znb4<UP>  ==>  zrl0
```

Configuring RAINlink for Solaris

Typically you need to configure either trunks or failover groups for a given network. With RAINlink for Solaris, you can also include trunks in failover groups. It is strongly recommended that you use trunking whenever your network adapters connect directly system-to-system, or to Cisco FEC compatible switch or router or other adapters. This is because trunking provides both increased bandwidth and improved fault-tolerance, whereas the failover service only provides improved fault-tolerance.

Configuring a trunk means to create (add) a trunk between two end stations and assign the member ports (links) to a trunk from EACH end station. A trunk should have at least two member ports (links.) Each port in a trunk can be a half or full duplex.

Configuring a failover group means to create (add) a failover group between two end stations and assign the member ports (links) to a group from each station. Either ports or trunks, or both can be members of a failover group. A failover group should have at least two members. Each port in a failover group can be a half or full duplex.

Quick Start

For most configurations, you will only need to pick the correct default configuration profile and run *rainlink -a*, as described previously.

If you wish to customize configuration for trunks or choose a different scheme for the dynamic load balancing, go to the "Advanced Trunking" section. If the link aggregation service is not suitable to your network configuration, then go to the "Advanced Failover" section to configure the failover groups instead of trunks. You can also configure failover groups of trunks themselves. Always configure trunks before failover groups.

Advanced Trunking

You can create link aggregation groups (trunks), assign ports to these trunks and select a scheme of dynamic load balancing. The most straightforward method for configuring RAINlink for Solaris is to edit the */etc/rain/rainlink.conf* file and run *rainlink -a*. The script de-installs the current ZNYX interfaces, runs a configuration application *rlconfig* with the *rainlink.conf* input script, and brings up the RAIN links. The *rainlink -a* script can also bring up any regular non-RAIN interfaces, which were not configured into the RAIN links. All these steps can be done individually. Details of using the *rlconfig* application are included in a later section.

To add a trunk, edit the file */etc/rain/rainlink.conf*. The file is a plain text file that is extensively commented. Commented lines begin with a pound sign (#). Add a line for each trunk containing the list of ports to be included in that trunk. End each line with a semicolon. For example, to include ports 0 through 3 in trunk0, add an entry:

```
trunk0=znb0, znb1, znb2, znb3;
```

The order of entries is not important. You can use any available ports in any order. Port entries should not be duplicated in trunks. A trunk must consist of at least two ports, and cannot contain other trunks. The following is an acceptable way to configure two trunks:

```
trunk0=znb0, znb3;
trunk1=znb1, znb2;
```

For system-to-system trunking, enable IP Trunking mode. Both systems should be setup identically. In IP Trunking mode, packets are sent to the driver as large datagrams and fragmented across the active links. To enable IP Trunking mode, uncomment/add the following line for each configured trunk:

```
trunk0=ip_trunking_mode;
```

For system-to-switch trunking, the default method of load balancing scheme uses Layer 3 Protocols. To use only the "Layer 2 Protocols" for dynamic load balancing, comment out the "layer_3_mode" option, and uncomment the "layer_2_mode" option. The options are mutually exclusive.

```
# trunk0=layer_3_mode;
trunk0=layer_2_mode;
```

To explicitly choose Layer 3, uncomment (or add) the "Layer 3 Protocols" option, and comment out the "Layer 2 Protocols" option:

```
trunk0=layer_3_mode;
# trunk0=layer_2_mode;
```

For system-to-switch trunking, you can disable Balance mode. Balance mode is on by default, and is normally left enabled. Balance mode dynamically distributes the load across the ports equally. It can be disabled by adding a "not equals" to the option:

```
trunk0 != balance_mode;
```

An "equals" sign explicitly enables balance mode:

```
trunk0 = balance_mode;
```

Once you have assembled your *rainlink.conf* file, run *rainlink -a* to bring up RAINlink for Solaris with the new configuration.

```
rainlink -a
```

Advanced Failover

You can create link failover groups, assign ports and trunks to these groups and select a modes of failover. Failover groups are configured in the same manner as trunks: Edit the */etc/rain/rainlink.conf* file and run *rainlink -a. Rainlink -a* de-installs the current ZNYX interfaces, runs a configuration application *rlconfig* with the *rainlink.conf* input script, and brings up the RAIN links. The *rainlink -a* script can also bring up any regular non-RAIN interfaces, which were not configured into the RAIN links. All these steps can be done individually. Details of using the *rlconfig* application are included in a later section.

To add a failover group, edit the file */etc/rain/rainlink.conf*. The file is a plain text file that is extensively commented. Commented lines begin with a pound sign (#). Failover groups can contain individual ports, trunks, or both ports and trunks. Always configure trunks before failover groups. Add a line for each failover group containing the list of ports and trunks to be included in that group. End each line with a semicolon. For example, to build a failover group of ports 0 and 1:

```
failover0 = znb0, znb1;
```

To build a failover group of two trunks, first build the trunks, then the failover group:

```
trunk0 = znb0, znb1;
trunk1 = znb2, znb3;
failover0 = trunk0, trunk1;
```

The following order will not work. You must build the trunks first:

```
failover0 = trunk0, trunk1;
trunk0 = znb0, znb1;
trunk1 = znb2, znb3;
```

You can also mix trunks and ports in failover groups, as long as the trunks are built first:

```
trunk0 = znb0, znb1;
failover0 = trunk0, znb2, znb3;
```

To remove a failover group, simply remove it, or comment it out, and run *rainlink -a*.

The default mode of failover is fast failover. In this mode, RAINlink for Solaris moves the traffic over to a redundant stand-by link in case of a link failure in as little as 500 milliseconds. In addition to fast failover, Timeout mode can be enabled for a failover group. By enabling Timeout Mode, if no traffic is received in the specified time interval, the active port is automatically switched to another available link. You can choose different failover modes for different groups. You can also specify the timeout interval. To set Timeout Mode on a failover group with a timeout of 30 seconds, add a line for the failover group:

```
failover0 = timeout_mode 30;
```

Using Rlconfig

The *rlconfig* application views or changes the RAINlink for Solaris configuration. *Rlconfig* is installed into */etc/rain* during the NetBlaster driver installation. You can view the current RAINlink for Solaris configuration at any time, but the interfaces must be "down" in order to change the configuration See *ifconfig(1M)* for more explanation of how to bring an interface down.

Options for *rlconfig* include:

Option:	**Use:**
-s	Displays the current RAIN configuration. Commands are not read from standard input with this option.
-t	Tear down all previously configured RAIN ports. Commands are not read from standard input with this option.
-l [#] file_name	Creates a file containing a list of the configured ZRL devices and the available ZNB devices. The "#" is used for version identification. This option is not intended for use by the user; it is used by the rainlink script to configure RAIN

To view the current configuration at any time, enter:

```
rlconfig  -s
```

The resulting display shows you the status of the RAIN configuration.

```
trunk0<IP_TRUNKING> = znb0<UP>, znb1<UP>, znb2<UP>,
        znb3<UP>, znb4<UP>  ==>  zrl0
```

Lines can be input directly to *rlconfig* from standard input, or from a script like */etc/rain/rainlink.conf*. To start *rlconfig*, enter:

```
rlconfig
```

Enter commands one line at a time and conclude with Ctrl-d.

To use a previously prepared file of commands, redirect the file into rlconfig. For example:

```
rlconfig < /etc/rain/rainlink.conf
```

Typical usage would be to redirect a script of commands into *rlconfig* which instruct *rlconfig* to build and set modes on trunks and failover groups. Semicolons delimit commands. Spaces and new lines are ignored. Commands take the following form:

```
trunk<number> = znb<number> [ ,  znb<number> … ] ;
failover<number> = znb | trunk<number> [ , znb|trunk<number> …]

trunk<number> = mode ;
trunk<number> != mode ;

failover<number> = mode ;
failover<number> != mode ;
```

where <number> is a value between 0 and the maximum number of ports in the system. Trunks can consist of ports denoted by a comma-delimited list of znb<number>. A failover group can consist of ports or trunks. You must build a trunk prior to its use in a failover group.

Acceptable modes for trunks are:

Mode:	**Appropriate for:**
ip_trunking_mode	System-to-system trunking
layer_2_mode	System-to-switch trunking
layer_3_mode	System-to-switch trunking
balance_mode	System-to-switch load balancing enabled/disabled
timeout_mode <time>	Failover mode; time parameter in seconds

To tear down the existing RAIN configuration, the interface must first be down. See ifconfig(1M):

```
ifconfig zrl0 down
rlconfig –t
```

Network Configurations

RAINlink for Solaris can be deployed in selected portions of a network for implementing either higher bandwidth or increased fault-tolerance or both. The link aggregation together with dynamic load balancing services offer both higher speeds (bandwidth) and increased fault tolerance, whereas the link failover service only provides increased fault tolerance. Therefore, it is recommended to use link aggregation service wherever possible in a network to make it faster and more reliable.

Refer to Table 2 to find out supported RAINlink for Solaris services and network topologies for various network configurations. Each network configuration is described in detail in the following section.

Network Configurations	**Link Aggregation** (with Dynamic Load Balancing)	**Link Failover**	**Comments**
System-to-System	√	√	Offers scalable bandwidth and increased reliability to network links between two computing stations (i.e. servers, workstations or desktops)
System-to-FEC* Switch or Router	√	√	Cisco supports FEC on Catalyst 5000, 2900XL, 2800, 1900 series Switches and 7500 series Routers. Many other switch vendors have also announced support for FEC.
System-to-Other Switch or Router		√	Protects against network port/NIC, cable and switch or router port failure on all topologies
System-to-Hub		√	Protects against network port/NIC, cable and hub port failure on all topologies

* FEC is an acronym for Cisco Fast EtherChannel technology. The "FEC Switch or Router" means Cisco Fast EtherChannel compatible Switch or Router.

Table 2. RAINlink for Solaris services and network topologies for various network configurations

System-to-System

You can configure trunks or failover groups between two computing stations (i.e. servers, workstations or desktops). Each computer may consist of multiple single-channel or multi-channel adapter ports.

Figure 7 illustrates a system-to-system configuration between two servers. Each server has four network ports. The three network links between servers (as represented by null i.e. crossover cables) are aggregated to implement a high-speed pipe between servers. The fourth link of each server is connected to a hub, which in turn is connected to multiple clients. One of these links will be forced to act as a standby link (as shown by the dotted line) if your hub or switch is running the IEEE 802.1d Spanning Tree algorithm.

> **IMPORTANT:** The current release of RAINlink for Solaris does not support IEEE 802.1d Spanning Tree algorithm. Therefore, you should avoid creating a network loop, if your switch does not support IEEE 802.1d Spanning Tree. Otherwise, your network may go down due to multiple redundant paths created by a loop.

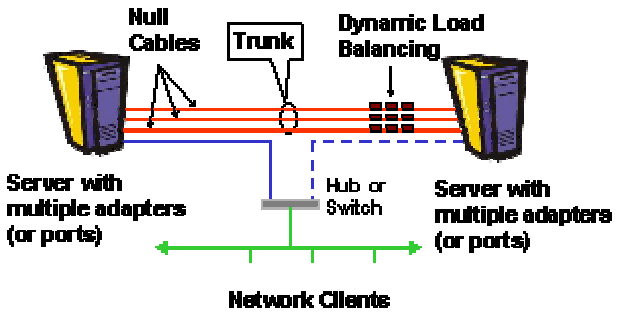

Figure 7. High speed pipe from server to server, using system-to-system trunking

The above network configuration is commonly used to implement higher bandwidth connections between enterprise-class intranet web servers, file servers or workstations.

System-to-FEC Switch or Router

This can be used to configure trunks or failover groups between a computing station and network switch (i.e. backbone, workgroup) or router. To configure trunks, the network switch (or router) must be a Cisco switch (or router) with FEC module or any other switch (or router) which supports Cisco's FEC compatible trunking scheme. To configure trunks on a switch (or router), refer to the technical manual from your switch (or router) vendor. Note that many FEC compatible switches or routers limit the maximum number of ports per trunk to four. The number of ports per failover group is unlimited.

Figure 8 illustrates a system-to-FEC Switch configuration. A server in data center is connected to Cisco Fast EtherChannel compatible backbone switch. As shown, by aggregating four full-duplex Fast Ethernet links (as represented by straight cables), you can implement 800 Mbps of pipe between a server and switch. You can also implement a similar type of trunk between a server and FEC compatible router. The Fast EtherChannel with router provides load balancing across aggregated links within a trunk based on IP addresses.

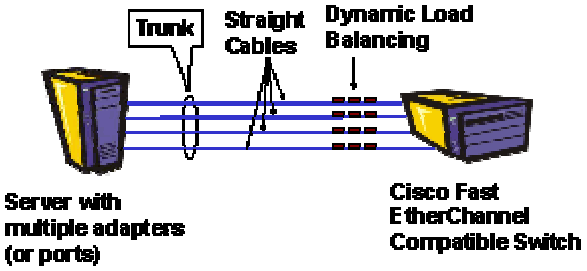

Figure 8. High speed pipe from Server-to-FEC switch using system-to-switch trunking

The above configuration provides network managers a reliable high-speed solution for the campus network backbone.

System-to-Other Switch or Router

In this configuration, you can configure failover groups (but not trunks) between a computing station, and network switch (i.e. backbone, workgroup) or router **which is not compatible with Cisco FEC**.

Figure 9 illustrates a system-to-Switch configuration, where the switch is not compatible with Cisco FEC. As shown, a failover group is implemented with two member links (as represented by straight cables) using RAINlink for Solaris installed on a server. The dotted line represents the stand-by link. If an active link goes down due to a server port/NIC, cable or switch port failure, RAINlink will switch over the traffic to a stand-by link without user intervention.

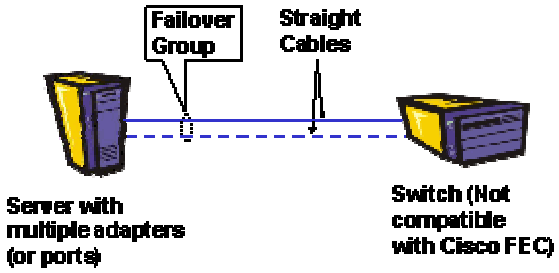

Figure 9. Reliable network connections from Server-to-Other Switch using Link Failover

System-to-Hub

This configuration allows you to configure failover groups (but not trunks) between a computing station and network hub. The computing station may consist of multiple single-channel or multi-channel adapter ports. You must use the same type of media and adapters in the host computer for optimal performance.

Figure 10 illustrated a system-to-Hub configuration. Again a failover group is implemented with two member links (as represented by straight cables) using RAINlink installed on a server. The dotted line represents the stand-by link. If an active link goes down due to a server port/NIC, cable or hub port failure, RAINlink's fast failover will automatically transfer the traffic to a stand-by link in as little as 500 ms, thus making your failover transparent.

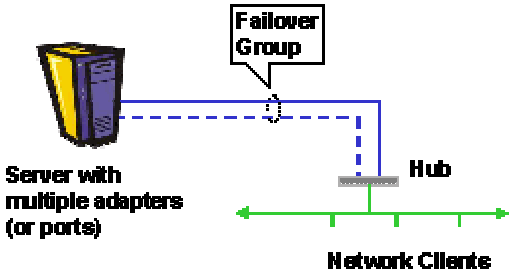

Figure 10. Reliable network connections from Server-to-Hub using Link Failover

Applications

RAINlink for Solaris can be used in a variety of applications to improve the bandwidth and/or reliability of your network. The following two examples depict the widely used applications of RAINlink to improve fault-tolerance in a server farm and to speed server interconnects and network backbones.

Server Farm

A typical server farm features several servers connected to a network switch, which, in turn, is connected to multiple clients. This configuration of server farm is vulnerable to single points of failure at the servers, the network links and the switch as shown in Figure 11. If a server, network link or switch fails for any reason, the client would loose connectivity to the server and the critical applications that run on the server.

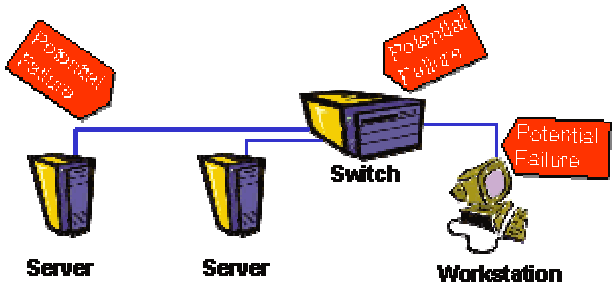

Figure 11. Typical Server Farm

RAINlink in conjunction with redundant network ports (or links) offers reliability to server farms by eliminating all single points of failure. As shown in Figure 12, one server and one workstation use two single-port adapters, while the other server uses a 2-port adapter. All systems are equipped with RAINlink for Solaris.

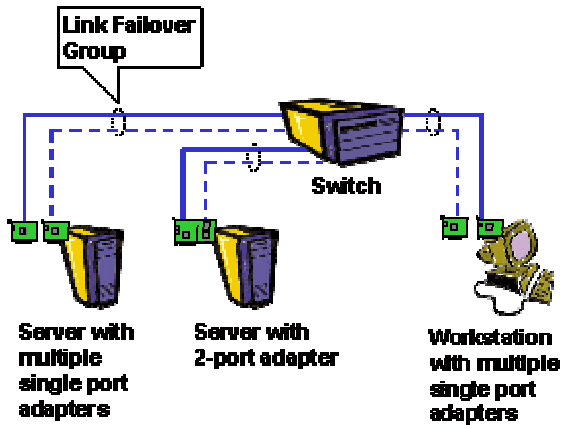

Figure 12. Reliable Server Farm

In the illustration, the link failover groups, with two network links, use RAINlink from server-to-switch and switch-to-client. The link failover feature offers continuous reliable network connectivity to client. In the event one of the links fails, the server farm can still service the client through the other links.

Server Interconnects and Network Backbones

Figure 13 illustrates how RAINlink and a 4-port NIC create an 800 Mbps "pipe" for server-to-server interconnects and server-to-switch backbone. In the illustration, four individual Fast Ethernet links are aggregated into one high-speed logical link to create the 800 Mbps "pipe" to relieve congestion in

a dedicated area of a network. The similar 800 Mbps "pipe" can also be deployed between Cisco FEC-compatible switches as shown by the switch-to-switch connection.

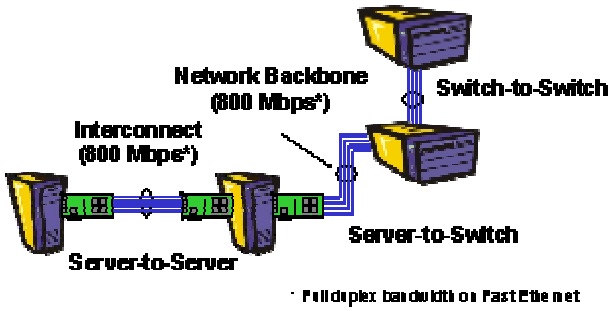

Figure 13. High-speed pipe for Server Interconnects and Network Backbones

Appendix

A: RAINlink for Solaris Man Page Reference

RAINlink includes the following files and utilities to aid in configuration and monitoring:

- rainlink
- rlconfig

This section includes Unix-style man pages to be used as reference.

rainlink(1M)

NAME

rainlink – Activate/Deactivate RAINlink for Solaris

SYNOPSIS

```
/etc/rain/rainlink [ -a | -d | -s | -u ]
```

DESCRIPTION

The *rainlink* command script initializes and reinitializes the RAINlink layer within the ZNYX NetBlaster driver. An input script located in */etc/rain/rainlink.conf* is used as input to ZNYX's rlconfig application. A list of default choices is presented if the */etc/rain/rainlink.conf* input script does not exist. The script can be used to bring up the RAINlink and non-RAINlink interfaces. It can also be used to display the current RAINlink configuration.

OPTIONS

-a	Activate RAINlink features. Creates an /etc/rain/rainlink.conf file if it doesn't exist from a list of default choices, brings down the existing NetBlaster interfaces, tears down the existing RAINlink configuration, builds the new RAINlink configuration, brings up the RAINlink interfaces, and finally any other NetBlaster interfaces.
-d	Deactivate RAINlink features. Tears down the existing RAINlink interfaces, and brings up NetBlaster interfaces.
-s	Show RAINlink configuration (same *as /etc/rain/rlconfig –s*)
-u	Display usage (same as no parameters)

FILES

/etc/rain/rainlink.conf

SEE ALSO

rlconfig(1M)

rlconfig(1M)

NAME

rlconfig - Configure ZNYX Redundant Array of Independent Netports (RAIN).

SYNOPSIS

```
/etc/rain/rlconfig [ -s ] [ -t ] [ -l # list_file ] < input_file
```

DESCRIPTION

The *rlconfig* application is used to configure multiple ZNYX NetBlaster ports into trunks or failover groups. With exception of the -s option, all interfaces affected by *rlconfig* must be in the down states. See *ifconfig(1M)* for explanation on bringing down interfaces. The application reads standard input for commands.

A trunk is a grouping of two or more ports that can do port aggregation. Port aggregation can be put in a mode where multiple clients connecting to a server can utilize the bandwidth of multiple ports seamlessly, or in a mode where a server connecting to a server can seamlessly utilize the higher bandwidth of multiple ports.

A failover group is a grouping of ports where only one port is active at any one time. The remaining ports in the failover group are in standby in case the active port goes down. Typically, a failover group would only switch from one link to another if the physical link were lost. By enabling timeout mode, if no traffic is received in the specified time interval, the active port is automatically switched to another available link. The timeout mode time interval is settable by the user.

OPTIONS

-s	Displays the current RAIN configuration. Commands are not read from standard input with this option.
-t	Tear down all previously configured RAIN ports. Commands are not read from standard input with this option.
-l [#] file_name	Creates a file containing a list of the configured ZRL devices and the available ZNB devices. The "#" is used for version identification. This option is not intended for use by the user; it is used by the *rainlink* script to configure RAIN

USAGE

Typically, a script of commands is directed into *rlconfig*, which instruct *rlconfig* to build and set modes on trunks and failover groups. Semicolons delimit commands. Spaces and new lines are ignored. Commands take the following form:

```
trunk<number> = znb<number> [ ,  znb<number> … ] ;
failover<number> = znb<number>  | trunk<number> [ , znb<number> |
                   trunk<number> …];

trunk<number> = mode ;
trunk<number> != mode ;
```

failover<number> = mode ;

```
failover<number> != mode ;
```

where <number> is a value between 0 and the maximum number of ports in the system. Trunks can consist of ports denoted by a comma-delimited list of znb<number>.A failover group can consist of ports or trunks. You must build a trunk prior to its use in a failover group.

Acceptable modes for trunk and failover configurations:

Mode:	**Appropriate for:**
ip_trunking_mode	System-to-system trunking
layer_2_mode	System-to-switch trunking
layer_3_mode	System-to-switch trunking
balance_mode	System-to-switch load balancing enabled/disabled
timeout_mode <time>	Failover mode; time parameter in seconds

DISPLAYS

The "-s" option of the *rlconfig* displays the current status of the RAINlink interfaces. The display includes any configured modes for trunks or failover groups, and the status of the interfaces in the form:

trunk<number> <mode> = znb<number>, znb<number> … ==> zrl<number>

Where number is the trunk, group, or interface identifier, and mode is one of the acceptable modes. For example:

trunk0<IP_TRUNKING> = znb0<UP>, znb1<UP> ==> zrl0

The interface status can be:

Name:	**Meaning:**
UP	Interface is UP
DOWN	Interface is down
ACTIVE	The interface is UP, and is the ACTIVE member of a failover group
STANDBY	The interface is UP, but is not the ACTIVE member of a failover group.

FILES

/etc/rain/rainlink.conf

SEE ALSO

rainlink(1M)

B: Compatible Products

RAINlink software should be compatible with any network interface card, hub and FEC compatible switch or router. ZNYX tested the compatibility and interoperability of RAINlink for Solaris software with the following products. Contact ZNYX Technical Support for the latest information on additional products tested by ZNYX labs.

Backbone	
Switches, Adapters	Fast EtherChannel (FEC) Compatible: Cisco Catalyst 5000/2900XL/2800/1900 Series switches, ZNYX's RAINcluster Family adapters
Routers	Fast EtherChannel (FEC) Compatible: Cisco 7500 Series
Network Interface Card (NIC)	
Fast Ethernet Adapters	ZNYX NetBlaster Single Channel PCI (ZX345, ZX345Q) ZNYX NetBlaster Dual Channel PCI (ZX348) ZNYX NetBlaster Quad Channel PCI (ZX346, ZX346Q) ZNYX NetBlaster Dual Channel CompactPCI (ZX412U3, ZX412U6) ZNYX NetBlaster Quad Channel CompactPCI (ZX414U3, ZX414U6)

C: Network Cable References

If you are deploying RAINlink in a network configuration where both end stations are NICs (i.e. system-to-system configuration), you must use "null" (crossover) cables to connect two end stations. In all other network configurations, you must use "straight" cables. Refer to the "Network Configurations" section for information on all possible network configurations supported by RAINlink.

Straight vs. Null (Crossover) Cable

Figure 1 shows the pin-out of a RJ45 connector for Ethernet and Fast Ethernet.

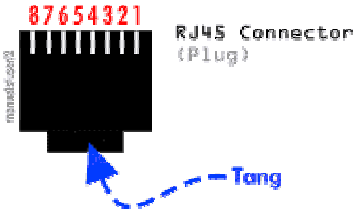

Figure 1. Pin-out of a RJ45 connector

As shown in Table 1, the wiring for a *straight cable* is simply pin 1 to pin 1, pin 2 to pin 2, etc. Notice that Pins 4, 5, 7 and 8 are not used.

RJ45 Pins	**Straight Connections**
1	1
2	2
3	3
4	-
5	-
6	6
7	-
8	-

Table 1. Ethernet and Fast Ethernet Wiring for a straight cable

The wiring for a *null cable* is shown in the Table 2 below.

RJ45 Pins	**Null (Crossover) Connections**
1	3
2	6
3	1
4	-
5	-
6	2
7	-
8	-

Table 2. Ethernet and Fast Ethernet Wiring for a null cable

D: Troubleshooting

Why don't I get a link light on my Fast Ethernet or Ethernet adapter?

- Verify that you are using a null (crossover) cable when your adapter is connected directly to another adapter. Otherwise, use a straight cable.

I don't see improved throughput to my Cisco switch even if I configured my adapter for link aggregation. Why?

- To see improved throughput, the ports on Cisco switch must be configured for Fast EtherChannel. Configure Fast Ethernet ports into Fast EtherChannel groups containing 2 or 4 segments, yielding 400- or 800-Mbps bi-directional bandwidth, respectively. Channels (Fast EtherChannels) can be configured using the command-line-interface (CLI) or Simple Network Management Protocol (SNMP).

  For instance, to set up a Fast EtherChannel, login to Cisco switch via a Telnet session. Then type:

  ```
  enable
  ```

  To aggregate ports 1 and 2 on the line card module 2 to be a Fast EtherChannel, type:

  ```
  set port channel 2/1-2 on
  ```

  To see how ports are configured, type:

  ```
  show port
  ```

  To see how Fast EtherChannels are configured, type:

  ```
  show port channel
  ```

Refer to the user's manual of Cisco switch for more information on Fast EtherChannel configuration.

- Multiple clients must be used to realize higher bandwidth. In system-to-switch trunking, a data path is established based on either the MAC or IP address, so the data between a particular client and server will always traverse the same path. More clients insure better distribution of data over all the ports in a trunk, leading to better total bandwidth utilization between the system and the FEC switch.

When I have two RAINlink systems configured system-to-system, why does it not work with more than one port plugged in between them?

- Check to see if you have a configuration mismatch. If one system is configured for trunking and the other for failover, having multiple ports between them will cause the link to fail. Ensure that system-to-system trunking is configured. The mode is displayed with the rlconfig –s command. Make sure the trunk includes IP_TRUNKING_MODE.

In NIC-to-Cisco FEC Switch or Router configurations, the process of unplugging the "last working port" and failing over to a "new working port" takes a very long time.

- The spanning tree "relearning" on the Cisco FEC switch may take long time to relearn (~ 20 seconds). There is a parameter on the Cisco FEC switch, which you can tweak on a set of ports to make the relearning happen much sooner. Use the following SNMP (Simple Network Management Protocol) command on your Cisco switch:

  ```
  set spantree fastport <mod/port> enable
  ```

  For example, "set spantree fastport 2/5-8 enable" enables fast learning on the ports 5 through 8 on line card module 2 of a Cisco FEC switch. Now you can unplug and plug a network link (port) quickly. You can also verify this using the ping command. You should notice that a network link establishes in just a few ping timeouts (<3 tries) during the process of unplugging and plugging a link.

Where to go for Technical Support

Electronic versions of this *RAINlink for Solaris User's Guide* are available off the ZNYX website:

http://www.znyx.com

✓ *RAINlink for Solaris User's Guide* (HTML) may be viewed using your web browser. From your browser's File menu, choose Open, select `index.htm` and click on OK. Follow the on-screen instructions.

✓ *RAINlink for Solaris User's Guide* (PDF) is available in `rl4sol.pdf`.

✓ Both versions of the User's Guide contain the same information. The HTML edition is designed for interactive browsing and reference. The PDF edition contains additional formatting information; it may be printed using Adobe Acrobat Reader.

Resource	Address
Telephone	(510) 249-0800
Toll-Free	(800) 724-0911 (USA Only)
FAX	(510) 656-2460
Website	www.znyx.com
E-mail	support@znyx.com

ZNYX is a leading worldwide supplier of scalable, reliable network solutions for servers, workstations and desktop systems. First in the industry to ship single and multi-channel PCI Ethernet and Fast Ethernet adapters, ZNYX offers a full line of products for PC, UNIX and Macintosh computers, including:

- Network Link Services Software
- High Availability and High Performance Network Cluster Adapters
- Open Integrated Network Switches
- Single and Multi-Channel Fast Ethernet Adapters
- Configuration and Network Management Software

For more information i.e. data sheets, white papers, FAQ, visit our website. ZNYX products are available worldwide through its OEM, distributors, systems integrators, and VAR partners. Sales inquiries can be directed to:

ZNYX Corporation
48501 Warm Springs Blvd., Suite 107
Fremont, CA 94539 USA

Tel.: (510) 249-0800
Toll-Free: (800) 724-0911
Fax: (510) 656-2460
E-mail: sales@znyx.com
Website: www.znyx.com

Document # DC0104-03
Version # July 7, 1999