HP network adapter teaming: load balancing in ProLiant servers running Microsoft® Windows® operating systems

white paper

Introduction

Balancing the network traffic load on a server can enhance the functionality of the server and the network. Load balancing within network interconnect controller (NIC) teams enables distributing traffic amongst the members of a NIC team so that traffic is routed among all available paths. HP Network Configuration Utility (NCU) software provides a user-friendly method to configure load balancing and to view load-balancing statistics. Load balancing is available to all NCU customers who team their NICs in ProLiant servers running Microsoft Windows operating systems.

This paper describes the load balancing options available with HP network adapter teaming, which is available on NICs in slots or embedded on ProLiant servers.

Overview of teaming and load balancing

By means of HP teaming software, system administrators can configure the network adapter ports within a ProLiant server as teams to provide fault tolerance. A team consists of two to eight ports that function as if they were a single port. The virtual port appears to the system as a single physical address and a single IP address. One port in the team, designated as the primary port, transmits and receives all traffic. The other ports are designated as backup ports. Should the primary port experience a disabling problem such as hardware failure, link loss, or a blocked path, HP teaming software will automatically make one of the backup ports the primary port.

Besides providing fault tolerance, NIC teams can also share the load of network traffic while presenting the unified single address to other network resources. This sharing of the load is called *load balancing*. The ports in a team can share both traffic transmitted by the server and traffic received by the server.

Transmit Load Balancing

The purpose of transmit load balancing (TLB) is to distribute transmitted traffic equally across all ports in the team. TLB allows a ProLiant server, statistically and heuristically, to balance network traffic more fully across all team members. Several configuration options allow administrators to determine which port will handle the transmit traffic. Table 1 identifies configuration options available for teams providing TLB with fault tolerance.

Table 1. Configuration options for transmit load balancing with fault tolerance in ProLiant servers

| Method	Definition	Usage notes
Automatic (the default and recommended method)	Allows the software to determine (statistically and heuristically) the best load balancing possible, dynamically and flexibly, based on per-packet packet type. A combination of Destination IP and TCP connection is used.	The default; most commonly used option.
Destination MAC address	Sends traffic through a port based on the physical address (MAC) of the destination node. Uses the least significant three bits to map to the transmission port.	Do not use a router in this traffic path.
Destination IP address	Sends traffic through a port based on the IP address of the destination node. Uses the least significant three bits of the right-most octet to map to the transmission port.	This is typically the easiest and most useful configuration.
TCP connection	Sends traffic through a port based on the network connection path (source-IP address, source_port, destination-IP address, destination Port). Uses elements of all four address fields to map to the transmission port.	This is used fairly often and can help on multi-application servers.
Round robin (packet order not guaranteed)	Sends traffic through a port based on a sequential selection model, which does not guarantee "in order" packet delivery on one path.	Web server environments are the most appropriate candidates for this option. UDP-based applications are also good candidates.

Configuring TLB

Once you have decided on the type of transmit load balancing for your system, use the Network Configuration Utility (NCU) to configure your choices. The NCU interface provides a simple configuration panel for specifying how traffic will be routed through the ports in a team (Figure 1).

To configure TLB, complete these steps:

1. Open the HP Network Configuration Utility.
2. Display the properties for the network adapter team.
3. In the Team Properties window, select the Teaming Controls tab.
4. Enter the appropriate team name.
5. From the Team Type drop-down list, select the desired type of TLB team.
6. From the Transmit Load Balancing Options drop-down list, select the TLB configuration option of your choice.
7. Click OK.
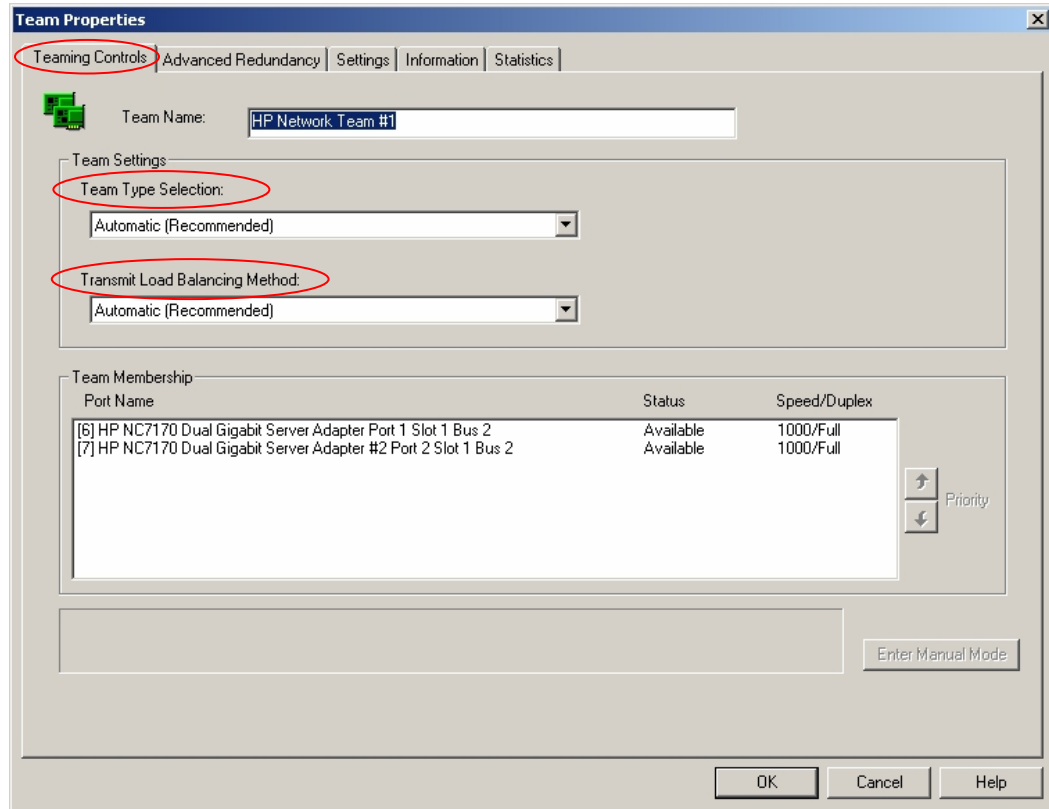
Receive Load Balancing

Switch-assisted load balancing (SLB) balances traffic transmitted and received by a ProLiant server. SLB is accomplished by means of several switch ports grouped into a single virtual port. The virtual switch selects the adapter to receive traffic based on the particular method supported by the switch.

Receive Load Balancing configuration decisions

Multiple configuration options allow administrators to determine how the ports will handle traffic. Table 2 identifies configuration options available for teams providing SLB.

Table 2. Configuration options for receive load balancing in ProLiant servers

| Method	Definition	Teams needed	When to use
802.3ad Dynamic with Fault Tolerance or LACP (no license required)	Using two virtual ports, the switch selects the ports for receive traffic as needed.	At least one two-port team	Ports are limited Traffic is normally not high. The switch supports this method.
Switch-assisted Dual Channel Load Balancing (requires PEINP license)	The switch uses two or more static ports, each designated to handle either transmit or receive traffic.	At least one four-port team with two ports per port channel	Ports are numerous. Traffic is constantly high. The switch supports this method. Switch redundancy is required.

Configuring Receive Load Balancing

Once you have decided on the type of switch-assisted load balancing for your system, use the HP NCU to configure your choices. The NCU interface provides a simple configuration panel for specifying how traffic will be routed through the ports in a team.

To configure SLB, complete these steps:

1. Open the HP Network Configuration Utility.
2. Display the properties for the network adapter team (Figure 2).
3. In the Team Properties window, select the Teaming Controls tab.
4. Enter the appropriate team name.
5. From the Team Type drop-down list, select the desired type of SLB team: *Dynamic 802.3ad* if you do not have a PEINP license, *Dual Channel* if you have a license or *both* if you have a license and want to use Dynamic Dual Channel teaming.
6. From the Transmit Load Balancing Options drop-down list, select the TLB configuration option of your choice.
7. Click OK.

Figure 2. Using the HP Network Configuration Utility to configure SLB Transmit and Receive Load Balancing in a ProLiant server
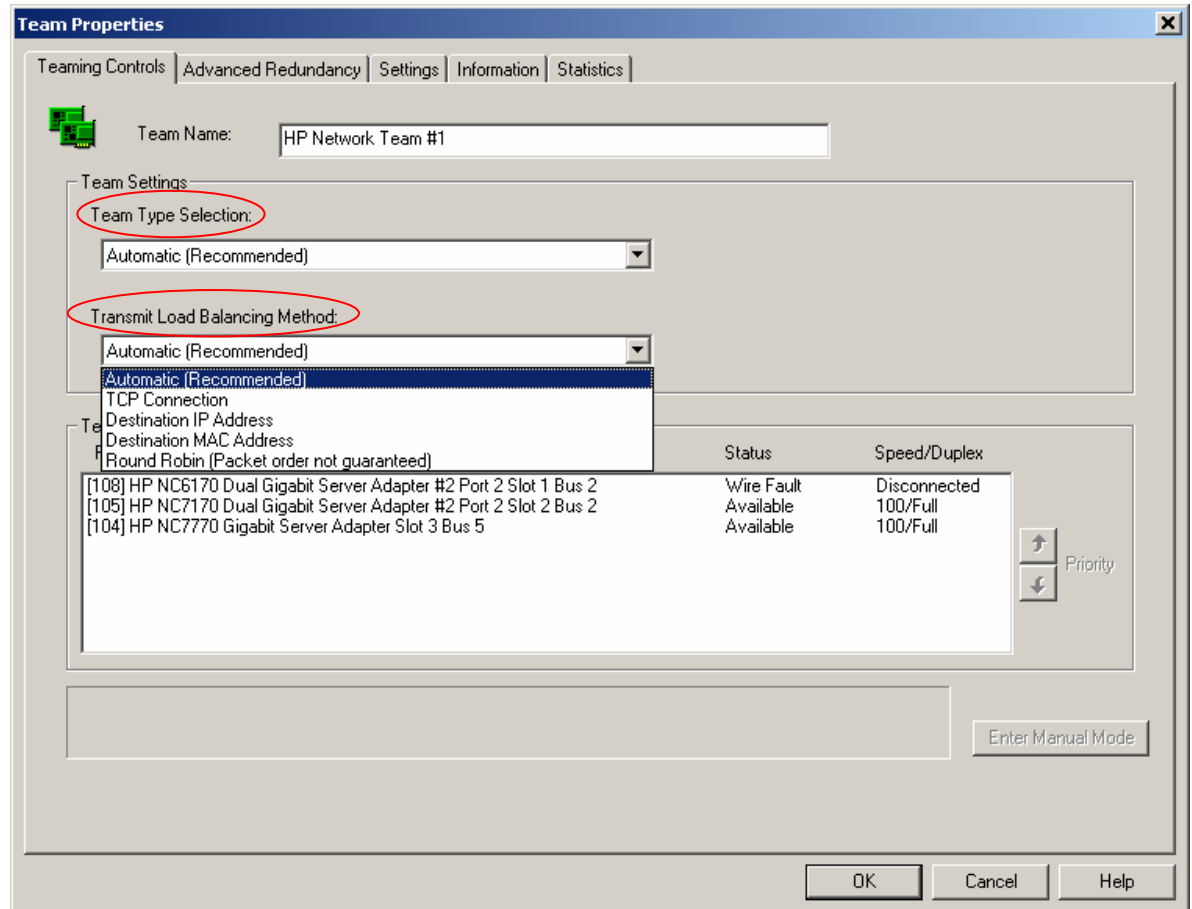
Conclusion

Load balancing is an important way to maximize network efficiency. HP provides practical and effective load balancing options and smooth configuration for harnessing the power of networks that contain ProLiant servers.

For more information

For more information about ProLiant network adapters and adapter teaming, see the following URLs:

http://h18004.www1.hp.com/products/servers/networking/whitepapers.html

http://h18004.www1.hp.com/products/servers/networking/teaming.html#1