

# Using APA to Build a Screaming Fast Network Server Connections

High Bandwidth Management Features  
and Benefits for IT Managers

**SNSL**



**E0300**

Printed in: U.S.A.

© Copyright 2000 Hewlett-Packard Company

---

## Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this white paper, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Copyright Notices. ©copyright 2000 Hewlett-Packard Company, all rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

Trademark Notices. UNIX is a registered trademark of the Open Group. MC/Service Guard is a registered trademark of the Hewlett-Packard Company.

---

# Contents

<b>Introduction .....</b>	<b>1</b>
The IEEE 802.3AD Standard .....	1
Sample APA Applications .....	2
<i>Fat Pipe Technologies</i> .....	2
<i>Web Access</i> .....	2
<i>Data Backup and Migration</i> .....	2
APA Core Components .....	2
<b>Load Balancing with Link Aggregates .....</b>	<b>3</b>
Outbound Load Balancing .....	3
Load Balancing Algorithms .....	4
<i>Server to Switch</i> .....	4
<i>Server to Router</i> .....	4
<i>Server to Server</i> .....	4
IP-based and MAC-based Load Balancing .....	4
CPU-based Load Balancing .....	6
<b>Sample Data Center Configurations .....</b>	<b>7</b>
Hardware Considerations .....	8
Inside the Data Center .....	9
Outside the Data Center .....	10
CPU-to-CPU Link Aggregates .....	11
Automated Fail-Over and HP MC/Service Guard .....	12
<i>APA Hot Standby Single-Server Fail-Over</i> .....	13
<i>Multi-Server Fail-Over with HP MC/Service Guard</i> .....	13
<i>HP High Availability Consulting and Support</i> .....	13
<b>Auto-configuration and Administration Tools .....</b>	<b>14</b>
APA Hands-off NIC Activation Features .....	14
Administration with the HP SAM GUI .....	16

---

# Contents

---

## Introduction

The purpose of this white paper is to present network managers with a high level overview of HP's IEEE 802.3ad Auto-Port Aggregation Software (APA) for the HP 9000 from the perspective of the network manager interested in deploying Gigabit (Gb) link aggregates. The paper addresses the following four main aspects of APA:


- The IEEE 802.3ad Standard
- Load Balancing within Link Aggregates
- Sample Data Center Configurations
- Auto-configuration and Administration Tools

APA is a server-side product designed to manage both Fast Ethernet and Gigabit Ethernet IEEE 802.3 compliant networks, and works with all HP 9000 systems running HP-UX 10.20 or HP-UX 11.00.

## The IEEE 802.3AD Standard

A remarkable aspect of the IEEE 802.3ad Link Aggregation Control Protocol (LACP) standard is that, when used with standard Fast Ethernet or Gigabit Ethernet link aggregates with existing copper and fiber cabling, and offers more available bandwidth than many small to medium-sized servers can handle.

APA offers network managers a comprehensive solution with automated management of link aggregates (up to 32Gb), three methods of load balancing, automatic fail-over to hot standby links, and full integration with HP's MC/Service Guard product.



## Sample APA Applications

APA's broad range of features offer immediate solutions for applications such as:

### Fat Pipe Technologies

For enterprise and campus networks with multi-gigabit backbones, APA link aggregates offer an affordable way to maximize overall network performance and reliability at the server level.

### Web Access

The load balancing and automated fail-over features of APA offers dial-up ISP and web hosting customers the highest levels of connectivity and reliability. For a dial-up user with 56K modem, this means more responsive and efficient downloads as well as significantly fewer "server not responding" network error messages.

### Data Backup and Migration

Enterprise Resource Planning (ERP) and other workflow applications are increasing productivity while reducing the costs of doing business. For example, one server is dedicated to inventory while another at a distant location handles sales order entry and invoicing.

APA offers direct connection server-to-server link aggregates for fast and uninterrupted online processing and backups.

Server-to-server link aggregates are especially helpful when substantial data migrations are required as the result of a merger, acquisition, or reorganization.

## APA Core Components

The core components of the APA product are it's three load-balancing algorithms. Only HP offers network managers three link aggregate algorithms for use with single logical address link aggregates (up to 32Gb each).

This core components give network administrator the ability to select the most optimal load-balancing algorithm based on whether your configuration is:

- Inside the data center
- Outside the data center
- Server to Server (CPU to CPU)

---

## Load Balancing with Link Aggregates

For network managers in need of an affordable solution to bandwidth management, the APA product fills the bill because it offers Highly Available Gigabit bandwidth that is fully compatible with existing Fast Ethernet or Gigabit networks.

What this gives network managers is the flexibility to stage Gigabit or Multi-Gigabit Ethernet deployments using any combination of the following types of link aggregates (which will be discussed in greater detail later on in this white paper):

- Link aggregates inside the data center
- Link aggregates outside the data center
- CPU-to-CPU link aggregates

As a server software product, APA gives HP 9000 servers the ability to support all three types simultaneously. The net result is that more desktops can now be served with fewer unique server IP addresses, pointing to significantly fewer servers.

## Outbound Load Balancing

It is important to note that APA only load balances outbound packets. (*Inbound load balancing is handled entirely by the sending switch or server.*)


Because APA load balancing only applies to outbound packets, there are no server-to-switch load balancing interoperability issues to decide. Consequently, network managers are free to selectively choose where and how to deploy APA. Outbound Load Balancing

## Load Balancing Algorithms

HP APA, is the only IEEE 802.3ad compliant product to offer switch-simple load balancing algorithms for three types of logical addresses: MAC, IP and CPU.

### Server to Switch


MAC-based Load Balancing Algorithm: This algorithm balances the load by hashing the 8 least significant bits of the destination MAC address stored within an Ethernet packet. It then directs the packet to a specific link within an aggregate.



MAC-based Algorithm

### Sever to Router


IP-based Load Balancing Algorithm: This algorithm balances the load by hashing the 8 least significant bits of the destination IP address stored within an Ethernet packet. It then directs the packet to a specific link within an aggregate.



IP-based Algorithm

### Server to Server

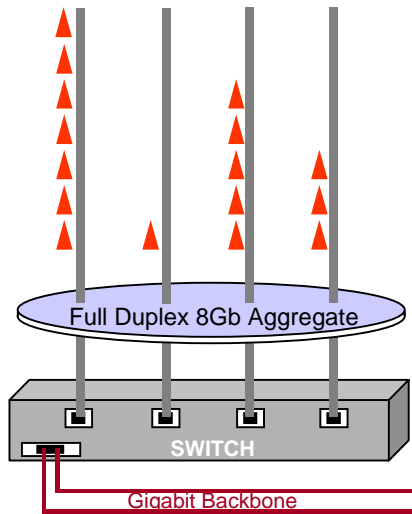
CPU-based Load Balancing Algorithm: Ethernet packets do not contain a CPU ID. For this reason, APA makes a separate system call to obtain the originating CPU ID for each thread. It then uses this CPU ID to direct the thread to an Ethernet port assigned to that CPU by the APA software.



CPU-based Algorithm

## IP-based and MAC-based Load Balancing

The IP-based and MAC-based load balancing algorithms both hash the 8 least significant bits of addressing information stored within an Ethernet packet. From the standpoint of APA, the only difference is whether the address bits come from the IP address or the DA address segment of the packet.




### FROM SWITCH

A typical switch hashes the 2 least significant bits of each address for load balancing with limited results.

Note: APA never negotiates load balancing of inbound packets with a switch or even another server running APA.

### TO SWITCH

HP APA software hashes the 8 least significant bits of each address, for switcher-style load balancing, using a data table with 256 options.

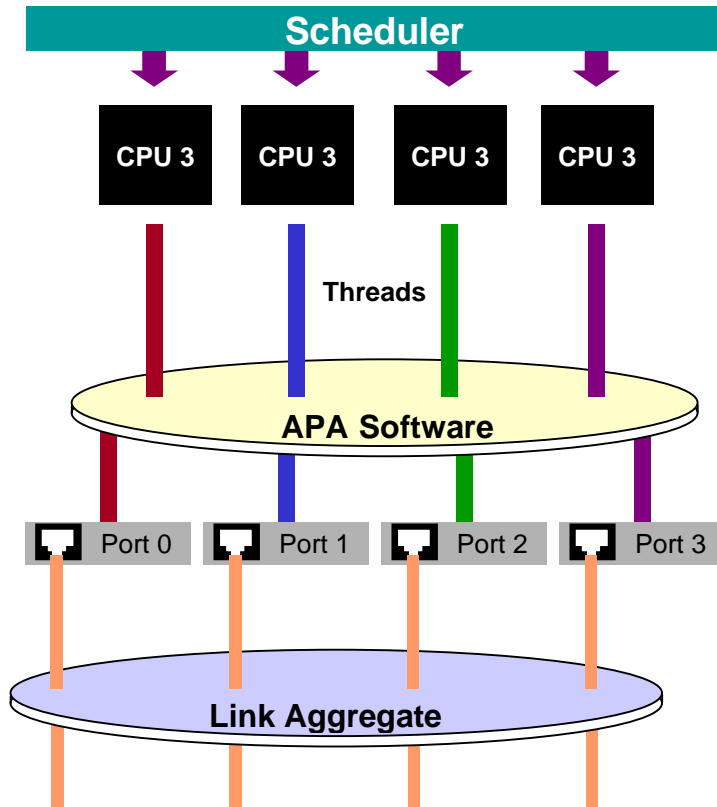


Another similarity is that the IP-based and MAC-based load balancing algorithms greatly resemble the simple, time-proven algorithms used by 802.3 Ethernet switches. The only difference here is that APA uses four times more addressing data than a switch, when assigning packets to specific links within a Fast Ethernet or Gigabit Ethernet link aggregate as shown in the MAC-based illustration above.

A consistent result in load balancing comparisons between switches and HP 9000 servers with APA is that using APA results in a more evenly balanced outbound stream. Of course, APA could be configured with even more elaborate load balancing algorithms that use more than the 8 least significant address bits, but only by sacrificing additional performance impact.

## CPU-based Load Balancing

The CPU-based algorithm offers the most direct way to connect two HP 9000 APA servers. This algorithm gives the highest level of performance for direct Server-to-Server configurations, which use CPU-intensive applications such as ERP, workflow, data warehouses, data migration, and data backup.



This algorithm relies on a CPU ID generated by the system to balance the threads coming from two or more CPUs to two or more Ethernet ports.

The purpose of a scheduler is to balance server-processing demands across one or more CPUs. The purpose of the CPU-based APA algorithm is to balance the output threads from those same CPUs across two or more Ethernet ports.

---

## Sample Data Center Configurations

The demand for network services is an ever-growing phenomenon and for the network manager, it is an absolute fact of life. Yet, the future is a murky cloud of budget meetings, acquisitions, mergers, and reorganizations.

*Can you relate to this? The annual shareholder meeting is going along exceptionally well, and the brilliant young CEO of your company has a room full of exuberant stockholders hanging on every word. Without warning, your equally enthusiastic CEO deviates from the prepared speech and begins to talk about the future of high bandwidth services. As you strain to hear the details, a colleague pats you on the shoulder and quietly says, "Is this awesome or what? I know we can depend on you to make this happen without spending too much."*

Mergers, acquisitions, and reorganizations: Each places a unique demand on enterprise data centers and the best solution is one that can address all of these needs. This is exactly what APA is designed to do!

Link Aggregates	Sample Performance Increase Areas
<b>Server to Switch</b> MAC-based Algorithm Gigabit Link Aggregates	<ul style="list-style-type: none"><li>• Client/Server LAN</li><li>• Workflow and ERP Applications</li><li>• Uninterrupted Data Backups</li><li>• High Availability Fail-over Systems</li></ul>
<b>Server to Router</b> IP-based Algorithm Gigabit Link Aggregates	<ul style="list-style-type: none"><li>• Internet Services</li><li>• VPN and Intranet Services</li><li>• Streaming Media</li><li>• Interactive Voice and Audio</li><li>• High Availability Fail-over Systems</li></ul>
<b>Server to Server</b> CPU-based Algorithm Gigabit Link Aggregates	<ul style="list-style-type: none"><li>• Uninterrupted Data Backups</li><li>• Data Warehouse Management</li><li>• Workflow and ERP Applications</li><li>• High Availability Fail-over Systems</li></ul>

With APA more desktops can be served with fewer unique server IP addresses. This in turn points to significantly fewer servers because APA can be used to create aggregate links for a wide range of needs.

What this gives network managers is the flexibility to stage link aggregate deployments in an incremental fashion, wherein they can specifically target their biggest bottlenecks first, second biggest next, and so forth.

## Hardware Considerations

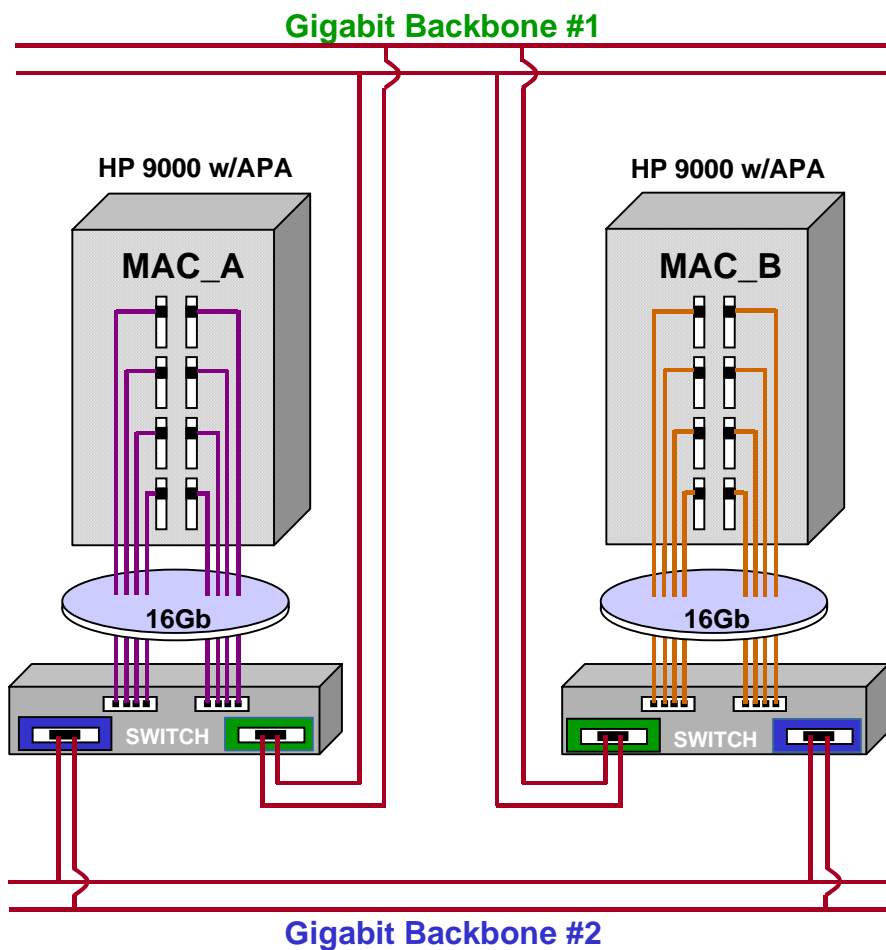
The APA software is the affordable solution for Highly Available Gigabit bandwidth and is designed to be fully compatible with existing Fast or Gigabit Ethernet networks. When planning a Gigabit link aggregate deployment, the following rules apply:

Rules	Description
Link Aggregate Type	Full Duplex only. Half duplex is not supported.
Fast Ethernet Trunk Sizes	2Kb to 32 Kb
Gigabit Ethernet Trunk Sizes	2Gb to 32 Gb
Maximum Number of IEEE 802.3ad Ethernet Link Aggregates per Server	50
Ethernet NIC	All 100BT Fast Ethernet  -- OR -- All 1000BT Gigabit Ethernet. 100BT + 1000BT NIC combinations are not supported.
FDDI and Token Ring	Fail-over links only. Load-balancing is not supported with FDDI and Token Ring.
Cable Type	UTP or fiber.
HP-UX 11.00 Operating System	Fail-over and load-balancing.
HP-UX 10.20 Operating System	Fail-over only.

## Inside the Data Center

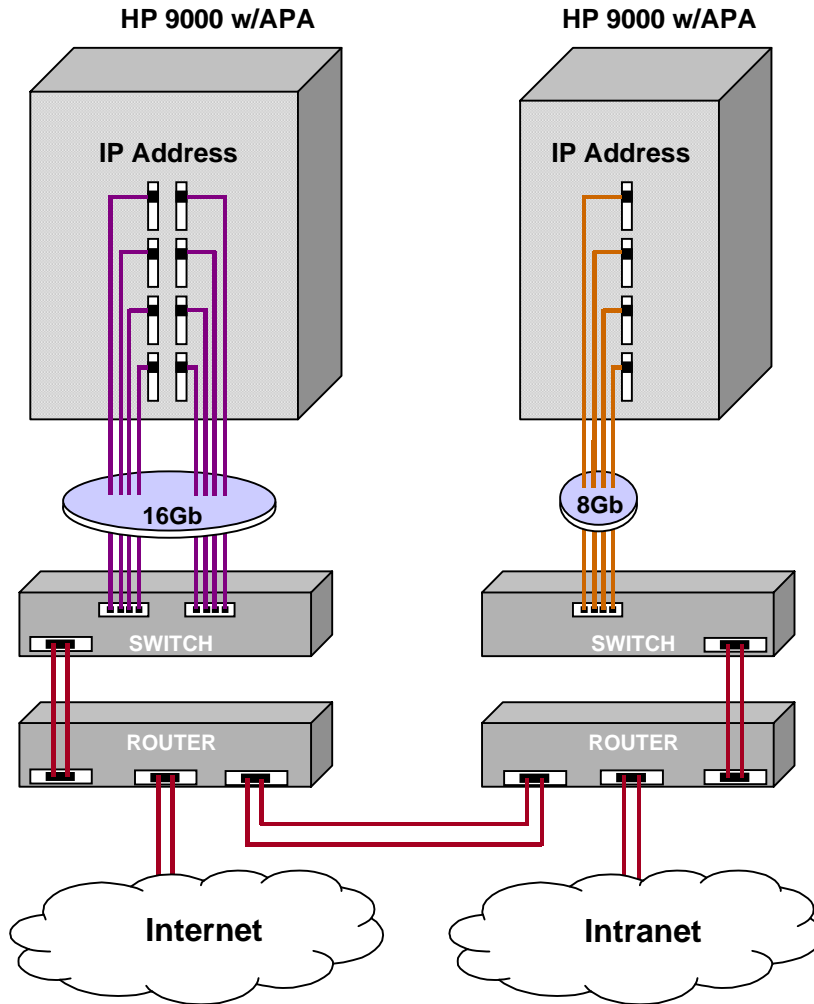
This sample configuration demonstrates how APA can be used to create link aggregates for a client/server configuration where side-by-side switch and server pairs serve two separate gigabit backbones.

Each server in this illustration has a single MAC address, supported by a 16Gb link aggregate. (APA can support link aggregates of up to 32Gb per MAC address.)



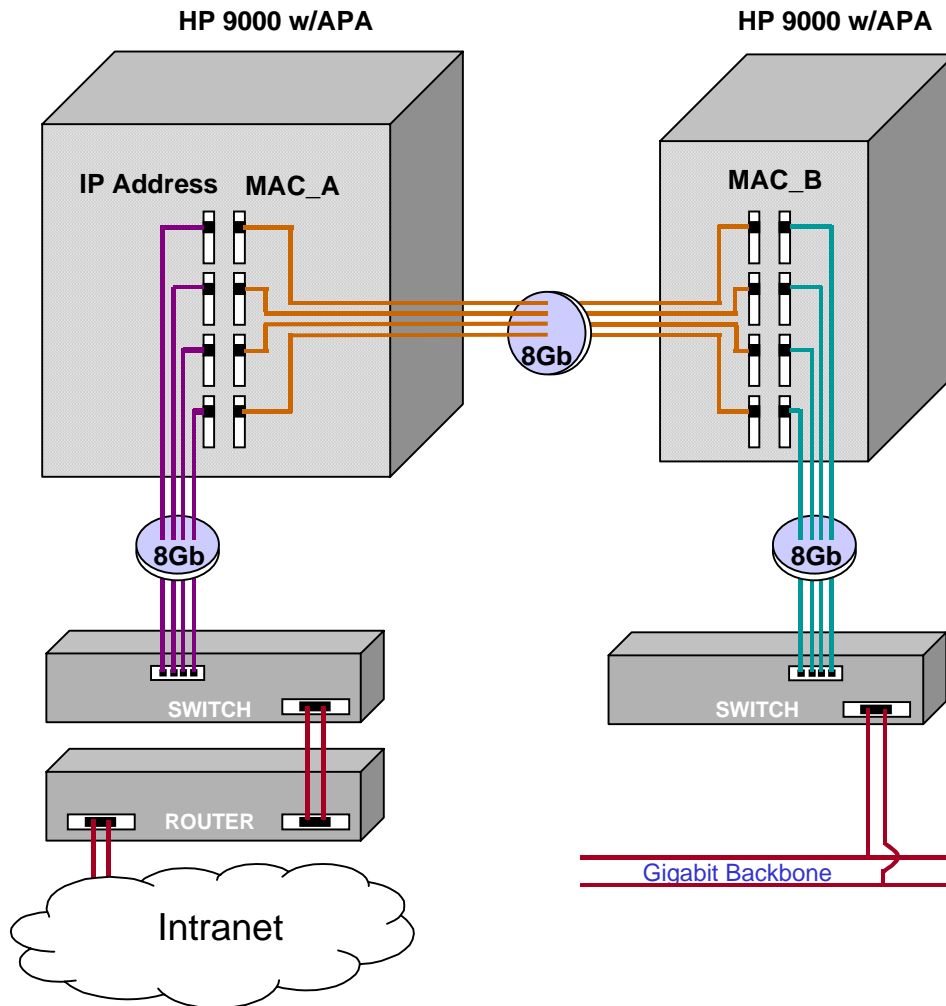
## Outside the Data Center

This sample configuration demonstrates how APA can be used to create link aggregates for common network services configurations. In this illustration, the Internet server is configured with an IP-based 16Gb link aggregate connection to the Internet. The Intranet server is configured with an IP-based 8Gb connection. (APA can support link aggregates of up to 32Gb per IP address.)



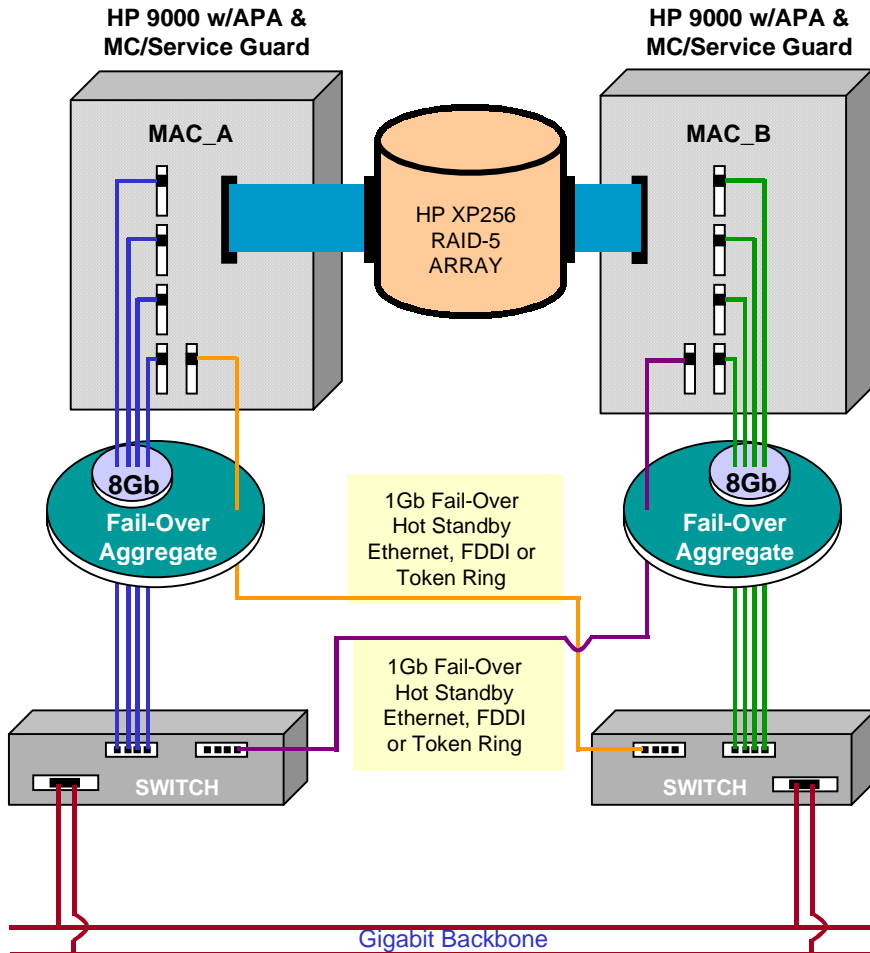
## CPU-to-CPU Link Aggregates

A key difference between CPU-based link aggregates and MAC-based link aggregates is that MAC-based aggregates are designed to work with configurations involving a large number of client accesses to the server. The CPU-based link aggregate is designed to manage the uninterrupted transmission of data directly between servers.



## Automated Fail-Over and HP MC/Service Guard

APA offers single-server fail-over automation. When a link fails within a link aggregate, APA automatically redistributes the load across the remaining links within that aggregate. If an entire link aggregate fails (because of a switch outage), APA will automatically fail-over to an APA configured hot standby link or link aggregate.



## **APA Hot Standby Single-Server Fail-Over**

APA offers fully automated single-server fail-over using an Ethernet, FDDI or Token Ring hot standby.

*Can you relate to this? The phone rings, and you sit up in your chair. "You know who had a catastrophic failure last month. They lost their web presence for two whole days." You remember the news as the voice continues. "Our people in acquisitions just picked them up for a song and we need you to fly down there first thing in the morning and check it out. I think the problem is in their fail-over configuration. Let us know your arrival time. We'll have someone meet you at the airport."*

It is important to note, that smaller data centers can enjoy real cost savings with APA, because it offers single-server fail-over support with MC/Service Guard levels of protection at no additional cost.

## **Multi-Server Fail-Over with HP MC/Service Guard**

For multi-server fail-over support, APA is designed to work with HP MC/Service Guard for HP 9000 servers and supports Ethernet, FDDI and Token Ring fail-over links.

## **HP High Availability Consulting and Support**

HP is a true single-source solution for world class support and professional consulting services. Our dedicated mission-critical support teams and consultants are ready to assist you in every detail of your planning, deployment and operations.

---

## Auto-configuration and Administration Tools

The shortfall in the current labor markets of seasoned network administrators is placing a heavy strain on IT staffing and personal budgets. A proven way to reduce these costs is by using technology that anticipates human error so as to automatically prevent catastrophic failures wherever possible.

*Can you relate to this? Half of your system is down and slowly coming back on line. You go to your best engineer for answers. The initial response to your hasty request is a glaring look that creeps out at you from across the top of a well-chewed pair of reading glasses. "It was the new hire, but in all honesty, anyone could have made that mistake." The rest is just unpleasant details.*

Should human intervention be required, automated solutions must offer intuitive problem-resolution tools designed for use by novice and expert alike.

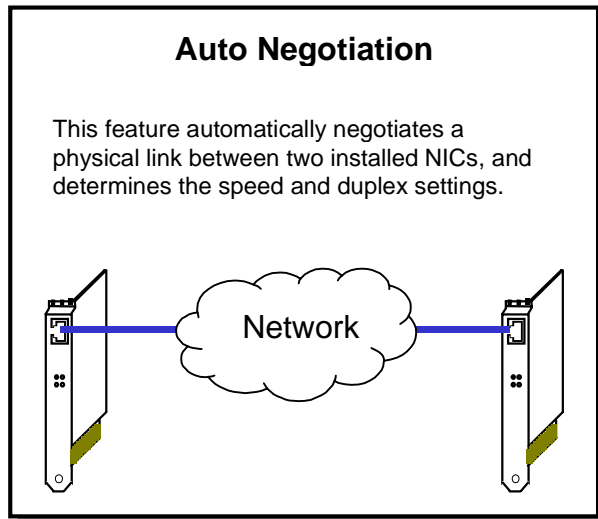
### APA Hands-off NIC Activation Features

Once a Network Interface Card (NIC) has been physically installed, the following automated features speed the process of building link aggregates.

- Auto Negotiation
- Auto Detection
- Auto Configuration


It is important to note that these features work together help prevent those types of failures that typically result from human error or component failure:

In the near future, 1000BT and 100BT HP 9000 server NIC installs and uninstalls will only require minutes of downtime or scheduled off-hours maintenance time per year, thanks to HP's soon to be released Online Addition and Replacement (OLAR) technology.



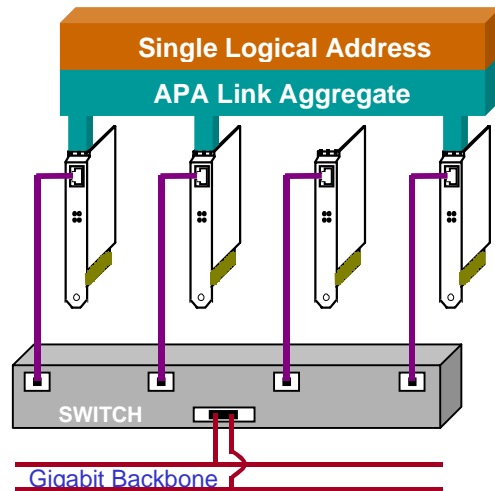
## Auto Detection

This feature identifies NICs that have been successfully configured with the Auto Negotiation feature, for use within an APA link aggregate.



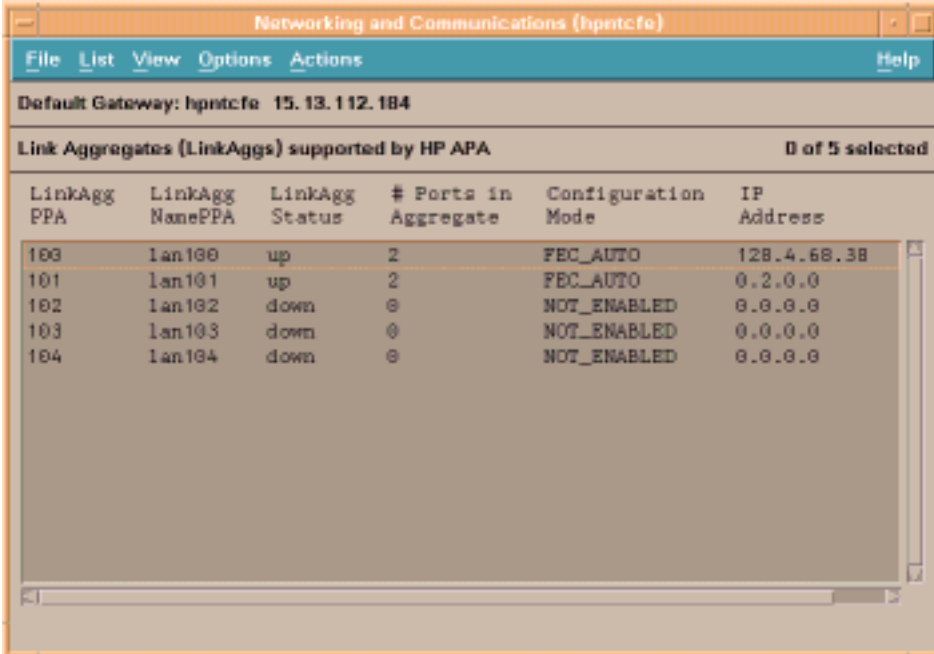
## Auto Configuration

This feature automatically builds link aggregates out of links that have been successfully identified by the Auto Detection feature



## Administration with the HP SAM GUI

The SAM GUI administration tool offers a single common interface for all HP servers. Once APA has been installed it appears as an additional system administration tool on the SAM menu option page. (*Network Administrators can also enter APA commands at the HP-UX prompt.*)



The screenshot shows a window titled "Networking and Communications (hprtcfe)". The window has a menu bar with "File", "List", "View", "Options", "Actions", and "Help". Below the menu bar, it displays "Default Gateway: hprtcfe 15.13.112.184". The main content area is titled "Link Aggregates (LinkAggs) supported by HP APA" and shows "0 of 5 selected". A table lists the following data:

LinkAgg PPA	LinkAgg NamePPA	LinkAgg Status	# Ports in Aggregate	Configuration Mode	IP Address
100	lan100	up	2	FEC_AUTO	128.4.68.38
101	lan101	up	2	FEC_AUTO	0.2.0.0
102	lan102	down	0	NOT_ENABLED	0.0.0.0
103	lan103	down	0	NOT_ENABLED	0.0.0.0
104	lan104	down	0	NOT_ENABLED	0.0.0.0

### SAM for APA Offers:

- Well Established Tool: SAM is a familiar administration tool for HP 9000 series system administrators.  
Quick Study: For new-hires, the intuitive simplicity of the SAM interface makes it quick and easy to learn.
- Simple-by-Design: SAM features a simple-by-design interface intended to reduce the chance of human error.

HP looks forward to the day when customers will be able to administer their systems using a single cross-platform 802.3ad Management Information Base (MIB) administration tool. For this reason, HP is actively working with other IEEE 802.3ad standard vendors, to promote the development of standardized MIB administration tools.