TECHNOLOGY BRIEF

June 1998

Compaq Computer Corporation

ECG Technology Communications

CONTENTS

PCI Hot Plug Technology with SCO® Software Architecture

As an increasing number of companies migrate their business-critical applications to industry-standard servers, the need for high-availability solutions to minimize system downtime is increasing. Compaq, the world's leading server provider, has historically taken a leadership role in developing industry-standard technologies. With the development of PCI Hot Plug technology, Compaq strengthens its role in providing high-availability solutions for the enterprise.

The Santa Cruz Operation, Inc. (SCO) is the world's leading provider of UNIX server operating systems. Because of SCO's expertise in this area, their UnixWare technology is perfectly suited for integration of PCI Hot Plug capabilities, especially Compaq's leading implementation. Compaq has worked with SCO to give administrators full hot-plug capability. This technology brief provides information on the joint effort between Compaq and SCO to implement PCI Hot Plug technology.

Please direct comments regarding this communication to the ECG Technology Communications Group at this Internet address: TechCom@compaq.com

COMPAQ

NOTICE

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

PCI Hot Plug Technology with SCO® Software Architecture
Second Edition (June 1998)
ECG082/0698

INTRODUCTION

Customers today need high-availability solutions that minimize or eliminate downtime. PCI Hot Plug technology allows a PCI adapter to be added, upgraded, or replaced while the host system is running and while other adapters in the system provide uninterrupted service. PCI Hot Plug technology adds to other high-availability solutions presently available from Compaq such as Recovery Server Options, Redundant Netelligent Network Interface Controllers (NICs), hot-pluggable power supplies, hot-pluggable fans, and hot-pluggable SCSI drives.

Three PCI Hot Plug Capabilities

PCI Hot Plug technology provides three major capabilities: hot replacement, hot upgrade, and hot expansion.

Hot replacement is the process of removing a failed or failing PCI adapter and inserting an *identical* adapter into the same slot while the server is operating.

Hot upgrade is the process of replacing an existing adapter with an *upgraded* adapter or replacing the adapter's device driver with an upgraded device driver while the server is operating.

Hot expansion is the process of installing an *additional* adapter into a previously empty slot while the server is operating.

Synergistic High-Availability Techniques

PCI Hot Plug technology represents a significant advance in fault-tolerant systems. To provide even greater benefit, PCI Hot Plug technology can be combined with redundant controller configurations, such as Compaq Redundant Netelligent NIC technology.

Compaq Redundant Netelligent NIC technology allows two similar NICs to share a single instance of device driver code. One NIC becomes the active network controller, and the other NIC acts as a standby controller. If the active NIC fails, the network traffic can be switched automatically to the standby NIC. With PCI Hot Plug, the failed NIC can be replaced without shutting down the system. Thus, the end user can have continuous service and the administrator can eliminate both planned and unplanned downtime.

Progress Toward an Industry Standard

Since Compaq's June 1996 announcement to deliver PCI Hot Plug technology as an open, industry standard, Compaq has worked with key hardware and software partners and the industry-wide PCI standards committee. This effort allows customers to move forward in implementing this technology with greater flexibility. Compaq has taken several steps to ensure broad industry acceptance of the technology, including:

- Standardized PCI Hot Plug technology by initiating and chairing the PCI Special Interest Group (SIG) Hot Plug Workgroup. The workgroup included other industry leaders such as Adaptec, Inc.; Cirrus Logic, Inc.; Digital Equipment Corporation; Hewlett-Packard Company; IBM Corporation; Intel; Microsoft; Novell; Pro-Log Corporation; SCO; and Texas Instruments Incorporated. Efforts of the workgroup allowed a quick release of the PCI Hot Plug Specification—only 14 months from the group's charter. For more information on the PCI Hot Plug Specification, visit the PCI SIG website at http://www.pcisig.com.
- Partnered with operating system (OS) developers to ensure incorporation of the PCI Hot Plug standard in product releases. Compaq has taken a proactive role in ensuring broad OS support for hot-plug capabilities. Actions to date range from the development of cross-company design teams to the delivery of PCI Hot Plug platforms for testing OSs. Currently, Microsoft, Novell, and SCO support PCI Hot Plug.

- Worked with leading adapter vendors such as 3Com Corporation; Adaptec; Dialogic Corporation; Digi International; Mylex Corporation; QLogic Corporation; American Megatrends, Inc. (AMI); Madge Networks; SMC Networks Inc.; and SysKonnect to ensure broad acceptance and rapid implementation of the standard.
- Licensed PCI Hot Plug technology to be available in Intel architecture-based server platforms. Intel will incorporate Compaq's implementation of PCI Hot Plug, including the hot-plug controller and related system technologies, in future products.
- Developed hot-plug aware drivers for Compaq NetFlex controllers, Netelligent controllers, Smart array controllers, and Fast-Wide SCSI-2 controllers.
- Developed the Compaq ProLiant 6500 and ProLiant 7000, announced in August 1997, as the first available server platforms conforming to the PCI Hot Plug Specification and added next-generation PCI Hot Plug features to Pentium II Xeon Models of the ProLiant 7000 announced in June 1998.

Compaq is in a unique position to lead this standardization effort. Not only is Compaq a pioneer in the development of systems hardware, but Compaq also has extensive experience in systems configuration software and device driver development. Because of its unique position, Compaq has been able to gather the support of major software developers and independent hardware vendors (IHVs) to deliver PCI Hot Plug technology as an open, industry standard.

The Compaq technology brief *PCI Hot Plug Technology*, document number ECG080/0698, outlines additional details of Compaq's overall strategy regarding this technology.

Compaq and SCO Team Up

SCO is the world's leading supplier of UNIX server OSs. Compaq has expertise in the development of systems hardware, systems configuration software, device drivers, and the design of high-availability solutions. Compaq and SCO have been working together for more than a decade and recently refocused the relationship to bring customers even higher quality enterprise solutions. The partnership now places more emphasis on joint engineering and development efforts, testing, and service and support. Both companies are committed to delivering all the benefits of a flexible, easily deployed, enterprise-level UNIX system on industry-standard servers.

SYSTEM OVERVIEW

PCI Hot Plug functionality requires both hardware and software components. A fully capable hot-plug system includes the hot-plug system hardware, a hot-plug aware OS and user interface (as provided by the SCO UnixWare 7 architecture), hot-plug system drivers, and hot-plug capable device drivers.

Most industry-standard PCI adapters can be used in Compaq's hot-plug environment without modification; however, their device drivers require new features to use hot-plug capabilities. A PCI adapter can be installed without using its hot-plug functionality until all the previously mentioned components are in place. PCI Hot Plug technology is backward compatible and can be incrementally incorporated by customers as they add or replace servers or components. This allows a mixture of existing components and new hot-plug technology to be used at the same time. It also protects the substantial investment customers have made in network hardware and allows them to gradually move toward improved system availability in a cost-effective manner.

System Hardware

It is important to note that while the PCI Hot Plug Specification identifies the technical requirements of a hot-plug capable system, it does not specify the implementation of this

technology. In the specific implementation that Compaq developed, the PCI Hot Plug hardware isolates a single PCI slot from all other devices on the PCI bus. Compaq's PCI Hot Plug system hardware performs two main functions:

- Powers down a single adapter slot, allowing insertion and removal of adapters
- Protects the system and other adapters from the electrical effects of hot-plug operations

Hot-Plug Controller and Slot-Specific Power Control

The hot-plug electronics designed by Compaq consist of two separate elements: the hot-plug controller and the slot-specific power control. Compaq designed its hot-plug controller to manage the following components:

PCI Bus. The controller communicates with isolation devices on the PCI bus to electrically isolate a single PCI slot from the rest of the system. Slot isolation permits insertion or removal of an adapter without interruption to the server or other active adapters.

Power. The controller receives a command from the OS to power up or power down a single PCI slot. To perform this function, the controller uses the slot-specific power control. The slot-specific power-control electronics allow the proper power sequencing on the PCI bus and guarantee safe control of the power to the individual PCI adapters.

Slot LED Indicators. The hot-plug controller also governs the slot LEDs. In Compaq's implementation of hot-plug hardware, each slot has a green and an amber LED to indicate slot status. The green LED indicates power to the slot and flashes while performing a power state change; the amber LED indicates that the slot requires attention.

PCI Hot Plug Button. This feature is available only on Compaq's latest PCI Hot Plug servers. The button is pressed to signal the software to initiate a power state change. While the button is more convenient, the same functionality is provided through the software interface. Each slot has its own button to indicate which slot is to be addressed by the supporting software.

System Chassis Design

In addition to the electrical components, the overall system design (system chassis and other hardware) has been adapted for safety and ease of adapter installation and removal. Chassis design changes include: quick-release latches on the adapter slots to allow fast and easy removal and insertion of adapters; wider slot spacing and flexible slot separators to allow users to remove and insert adapters without electrically contacting (shorting) other components; and a top access shield that allows users access to adapters yet prevents access to other internal components. These design changes meet or exceed all regulatory safety standards.

SCO Software Architecture

SCO's UnixWare 7 supports all three hot-plug capabilities (hot replace, hot upgrade, and hot expansion), allowing users to eliminate downtime when they replace faulty adapters, upgrade their systems, or dynamically add new adapters. UnixWare 7 does not currently support full PCI Hot Plug Button functionality, but support for the button is forthcoming. Figure 1 shows primary software component enhancements that are included in UnixWare 7: the hot-plug user interface, the Hot Plug Controller Interface, the Autoconfig Subsystem, the device driver-specific interface, and the hot-plug controller driver.

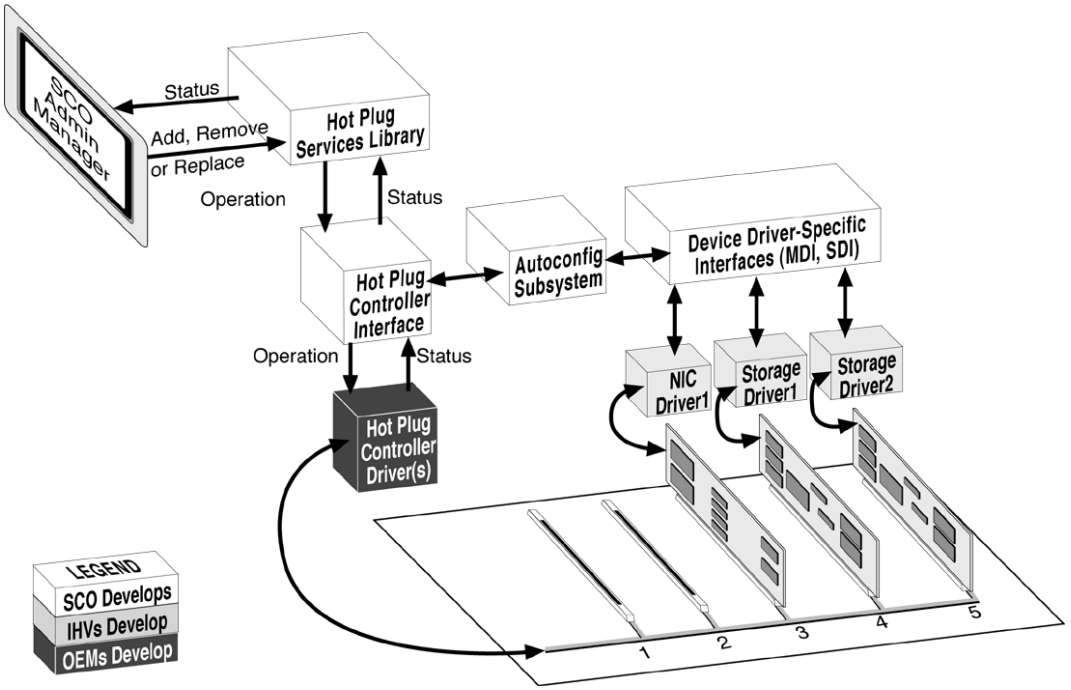

Figure 1: SCO UnixWare PCI Hot Plug Software Architecture

User Interface

As shown in Figure 1, two components comprise the user interface layer: the Admin Manager and the Hot Plug Services Library (HPSL).

Admin Manager – The Admin Manager is the graphical user interface for PCI Hot Plug technology. All user commands to add, remove, or replace an adapter funnel through the Admin Manager, which reports the status of these actions back to the user. The Hot Plug Admin Manager is incorporated into SCO's system management framework called SCO Admin.

Hot Plug Services Library – The HPSL provides isolation between the Admin Manager and the UnixWare kernel. In essence, it hides all the specific hot-plug operations performed by the Config Driver and allows the user interface to remain constant even if future changes occur in the OS. The HPSL processes all of the user requests and accesses the state of the adapter ports through the Hot Plug Controller Interface. The HPSL has access to the resource manager database, which allows it to get device-specific information for the user.

Hot Plug Controller Interface

The Hot Plug Controller Interface (HPCI) does the real work: it is the primary component for the underlying management of hot-plug capability. It tracks system component information as well as manages the actions required for hot-plug functions. The HPCI processes HPSL requests by interacting with the Autoconfig Subsystem and various hot-plug drivers in the system.

Autoconfig Subsystem

Building on the existing SCO UnixWare 2.1 auto-configuration framework, hot-plug events are handled through the resource manager database. The resource manager database has been updated for hot-plug capability, specifically in the area of dynamically adding and removing individual device driver instances.

Device Driver Interface Specification

The device driver interface (DDI) specification defines how device drivers communicate with the UnixWare kernel and eventually with higher-level system software and hot-plug software. SCO's Device Driver Interface version 8 (DDI-8) specification and other device driver interfaces are enhanced to provide the functions for adding or removing driver instances and suspending or resuming operation in conjunction with hot-plug operations.

Because these interfaces are well defined, existing drivers can be easily modified in support of PCI Hot Plug technology. The DDI specification evolved from both UNIX System5 Release 4 and UnixWare 2. The DDI-8 specification provides a high degree of compatibility for UnixWare users. Device drivers that meet DDI-8 requirements will also work for PCI Hot Plug. Thus, investment protection is assured by adhering to DDI-8. DDI-8 also allows device instance independence, which makes it possible to add or remove a single device instance without requiring the driver to load and unload when it is managing multiple devices.

DDI-8 includes new driver functions for add, remove, suspend, and resume operations. The add operation is required for all DDI-8 drivers. The remove, suspend, and resume operations are optional, but they are necessary if the driver is to support hot-plug operations. An add operation is the process of adding a new device instance to a driver. This entry point will be called at system initialization for each device that the driver supports and that was detected by the UnixWare auto-configuration process. Additionally, this entry will be called after an adapter has been installed into a hot-plug slot and a new device driver has been loaded into a system. A remove operation unbinds an instance of a driver from being associated with a particular device. This occurs prior to an adapter being physically removed from the system.

The suspend and remove operations are used primarily when the user intends to replace an adapter. The suspend operation tells the driver not to access the adapter while it is being replaced. This action allows the higher-level system services such as protocol stacks and file systems to remain in place, although suspended, during the hot-plug replacement of the adapter. Once the adapter has been replaced, a resume operation occurs, which signals the driver that it is safe to resume operations to the adapter.

Hot-Plug System and Device Drivers

Hot-Plug System Drivers

A hot-plug system driver is the software driver that controls and monitors the hot-plug controller hardware. Compaq provides the system level hot-plug driver to support UnixWare 7 on Compaq server platforms that provide PCI Hot Plug. As other original equipment manufacturers develop hot-plug platforms that comply with the PCI Hot Plug Specification, they will need to supply a hot-plug system driver to support the UnixWare 7 OS on their server platforms.

Device Drivers

Device drivers that support PCI Hot Plug with UnixWare include Host Bus Adapter drivers, which incorporate storage device drivers, and Media Access Control (MAC) Driver Interface (MDI) drivers, which incorporate NIC drivers. These drivers are developed by IHVs. Compaq provides hot-plug-capable drivers for its own leading PCI server adapters such as NetFlex controllers, Netelligent controllers, Smart array controllers, and Fast-Wide SCSI-2 controllers. Leading IHVs committed to modifying their device drivers to be hot-plug aware include 3Com, Adaptec, Dialogic, Digi International, Mylex, QLogic, AMI, Madge, SMC, and SysKonnect.

HARDWARE AND SOFTWARE CONTROL

In a typical scenario in which the administrator adds an adapter to an empty slot, the following steps occur while the system is running:

1. The administrator prepares the slot for installation of the adapter by opening the appropriate slot release lever and removing the expansion slot cover.
2. The administrator installs the adapter into the appropriate expansion slot.
3. The administrator closes the slot release lever.
4. The administrator uses the Admin Manager or the PCI Hot Plug Button to notify UnixWare that power can be applied to the slot.
5. UnixWare directs the hot-plug controller to turn on power to the slot. The green LED will flash while UnixWare is performing the power state change.
6. The administrator uses the Admin Manager to notify UnixWare to load the adapter's device driver.
7. UnixWare instructs the device driver interface to select and load the appropriate driver.

CONCLUSION

With the introduction of PCI Hot Plug technology, Compaq once again set the standard in high-availability and fault-tolerant solutions for the enterprise. PCI Hot Plug offers unprecedented server availability by allowing users to replace, upgrade, and add PCI adapters to the PCI local bus without powering down the server. With the release of the ProLiant 6500 and ProLiant 7000, Compaq made the chassis and electronics design changes required for safe and efficient removal and insertion of PCI adapters without powering down the server. Next-generation PCI Hot Plug features are included in Compaq's newest enterprise server, the Pentium II Xeon Model of the ProLiant 7000, released in June 1998.

SCO is a leading OS developer whose UnixWare technology is perfectly suited for integration of PCI Hot Plug capabilities. Compaq has worked closely with SCO to develop hot-plug capabilities for UnixWare. UnixWare 7 supports all three hot-plug capabilities, and future releases will include additional support for next-generation features like the PCI Hot Plug Button.

PCI Hot Plug technology brings significant advances to the other high-availability solutions Compaq currently offers. By bringing this technology to the market as an open, industry standard, Compaq continues to strengthen its position as a leader in enterprise computing.