PERKIN-ELMER

Series 3200

RESILIENT SYSTEM

High Availability Transaction Processing

Full Resilient System

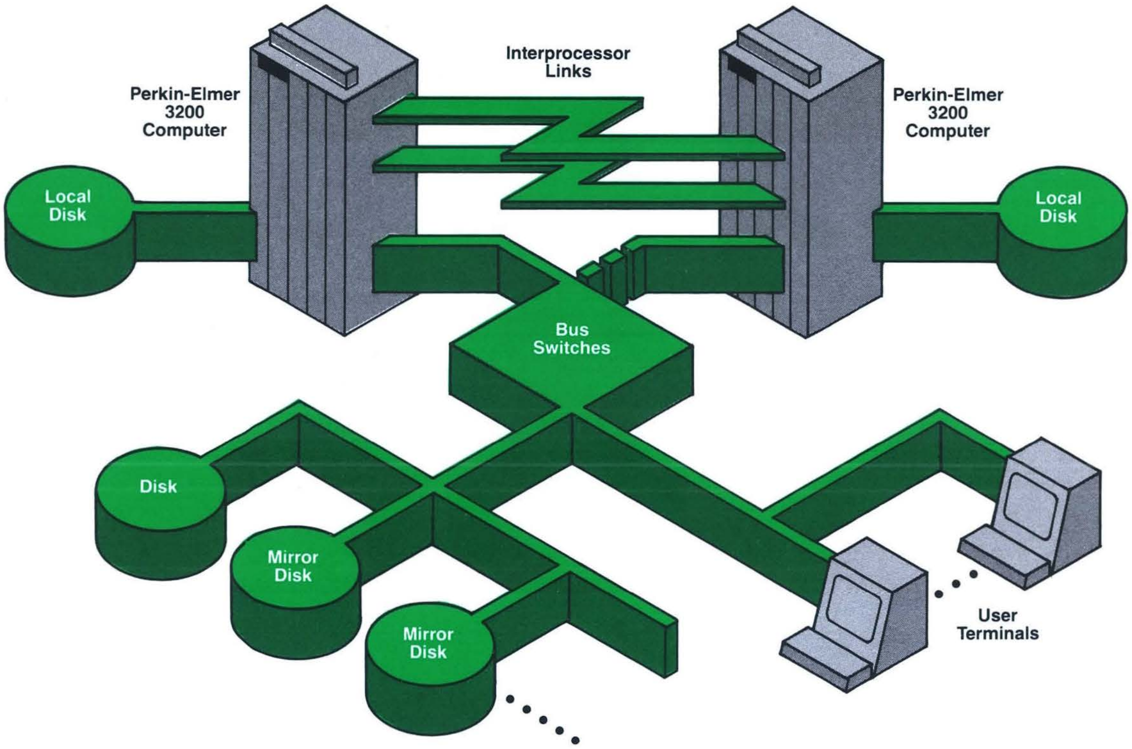

Product Overview

The Resilient System is Perkin-Elmer's answer to the increasing need for higher availability transaction processing. Resilience is achieved by combining redundant hardware components and employing two software features: the Reconfiguration Monitor and Mirror Disks.

The **Reconfiguration Monitor** is a software package which enables two processors running OS/32 to be connected in a dual computer configuration. Each system is capable of independently running a number of application systems and simultaneously monitoring the other to take control of all the applications, if necessary.

The Reconfiguration Monitor runs in both computers, and reassurance messages are exchanged between the two monitors over an interprocessor link. On detection of a failure in the other system, the Reconfiguration Monitor activates bus switches to take over peripherals from the failed system and performs specific actions to load and restart the failed application system or systems. It does this by submitting Command Substitution System (CSS) files to the Operating System, OS/32. Additionally, the Reconfiguration Monitor is able to accept operator commands to reconfigure the system in the event that part of the system is required for normal engineering maintenance. This means there will be a minimum effect on the production environments.

The optional **Mirror Disk** capability allows pairs of disks to be maintained in a completely consistent state totally transparent to users and applications. All information is recorded twice, once on each disk.

Product Overview (Continued)

In the event of a single disk failing for any reason, the system will detect the disk failure, alert the system operator, and automatically switch to use only the remaining disk. As soon as the disk drive failure is corrected, the system provides the operator with the utilities to guide the rate at which it will catch up to be in full synchronization again. The recovery process is able to proceed in parallel to normal running and only uses that portion of the system resources indicated by the operator.

Features

Automatic System Reconfiguration and Restart Following a System Failure

- Continuous operation by the user.
- Reconfiguration Monitor checks status of another machine by exchanging reassurance messages.
- Failure detection in another machine.
- Automatic Switching of peripherals from failed to working system.
- Automatic restart of system and applications on new machine after switching.
- Operator requested switching for orderly reconfigurations.

Data Redundancy Via Optional Mirror Disks

- Automatic continuation following a single disk failure with uninterrupted availability of data.
- Automatic mirroring of pairs of disks.
- On-line resynchronization without break of service and without interrupting the availability of data.
- Transparent to users and applications.
- Operator requested switching of mirrored disks for orderly reconfigurations.

Flexible Configurations With Modular Growth

- Will accommodate any paired combination of 32-bit processors.
- Start with one system today and upgrade to a second system and Mirror Disks tomorrow for resiliency when you need it.
- Both systems operate totally independent of each other—for example, one does production work and the second does development work.
- Takes full advantage of the Reliance transaction processing systems' inherent recovery and reliability features.

Benefits

The Resilient System combines a number of features and functions that make it a unique offering in the high availability marketplace. The major advantages of the Resilient System are immediately beneficial to customers using Perkin-Elmer's on-line transaction processing system, Reliance PLUS. Reliance PLUS employs many recovery and database integrity features which are exploited by the high availability characteristics of the Resilient System. The Reconfiguration Monitor, the Mirror Disk option, and Reliance PLUS run independently of each other. Together, they form an extremely high-level of reliability and system availability for the user.

Additional benefits of the Resilient System include:

High Performance/High Availability

Perkin-Elmer offers one of the highest performance Resilient Systems in the industry. By exploiting the strengths of the 32-bit architecture, customers can extend the availability and reliability of these systems thus protecting their investment and improving productivity.

Complete, Resilient Transaction Processing Environment

The combination of the Resilient System and Perkin-Elmer's transaction processing and database management system, Reliance PLUS, provides an extremely high-level of reliability and system availability for the transaction processing user.

Simple Implementation

Installation and implementation of the Resilient System does not require major hardware or software reconfigurations. This is a critical consideration for customers migrating into a Resilient System who cannot afford to disturb their production environment with costly, time-consuming upgrades. Additionally, application software written for a Reliance transaction processing environment does not have to be changed since the Reconfiguration Monitor and Mirror Disks run independently and are transparent to application programs.

Applications written for Reliance PLUS automatically take advantage of the database integrity and reliability features built into the Reliance PLUS software. Applications written for non-Reliance PLUS environments have the capabilities to add their own data checkpointing and unique application recovery procedures when and where needed.

Independent Operation of Both Systems

No system in a resilient configuration is dedicated, either fully or partially, to monitoring or duplicating its counterpart system. Both systems operate totally independent of each other permitting one system to be used exclusively for high priority production work while the second system can be used for development, for instance. In the event of a failure in the production system, the development system will automatically reconfigure itself, switch over all necessary peripherals from the failed system and take over the production environment. This approach eliminates the need to have a hot stand-by sitting idle.

Benefits (Continued)

Modular and Incremental Growth

A Resilient System can actually begin as a single system. Once the development efforts are completed and the system is ready for production, additional hardware and system software can be added to migrate to the specific level of resilience that is needed. No advanced programming is required in preparation for these capabilities since Mirror Disks and the Reconfiguration Monitor are transparent to user applications.

Flexible Configurations

There are no restrictions of CPU model types in a resilient configuration provided each system can support the other environment(s) it will be required to upon a recovery procedure. Any combination of 32-bit computers running OS/32 can be mixed or matched to provide the appropriate level of resilience at a cost that is more manageable and controllable by the customer.

Operator Initiated Reconfiguration

The system operator can easily cause a system takeover via a simple command thus allowing scheduled engineering maintenance to occur with minimum effect on the production environment.

The Reconfiguration Monitor

The Kernel Task

The Reconfiguration Monitor is split into a number of discrete tasks. The heart of the system is the Kernel Task which monitors the current machine configuration and communicates with the other processor. In addition to the Kernel Task there are a number of other tasks which make up the Reconfiguration Monitor. These tasks perform various functions including loop detection, I/O bus switching and system definition.

The Kernel Task monitors the status of the system as a whole. To do this, it communicates using intertask messages with all the components of the Reconfiguration Monitor. In addition, it continually monitors the communication lines with the other processor and communicates with the Kernel Task in the other machine.

Interprocessor Communications

The Kernel can operate with a single or dual RS232 interprocessor link. A dual link, on separate communication ports, is the normal configuration for the monitor since a single link provides a single point of failure and therefore cannot be considered as resilient as with a dual link. On initiating the Reconfiguration Monitor, the operator specifies the name of the link or links.

When a failure is detected on one of the links, the Kernel issues a message and continues operating on the other link. The Kernel periodically retries the link and when it determines that the link has recovered, it issues a message and resumes normal operation.

Reconfiguration Procedures

Initially the Kernal Task sends a Configuration Request message to the other processor indicating that it needs details of the current configuration. If no reply is received or the original write fails, the Kernel assumes that the other processor is not running.

In normal running the Kernel sends and receives Reassurance Messages. If a Reassurance Message is not received, then the Kernel Task takes control of the other system's applications. This is done by running CSS files to acquire the required peripherals and then running other CSS files to load and start the applications.

If a configuration change is requested by the operator indicating, for instance, an application is being started, the corresponding message is sent to the other processor. On receipt of a Configuration Request, the configuration details are sent to the other system. On receipt of a confirmation of changeover, the Kernel updates its own configuration file. These interactions cause each system to be fully aware of the other systems' configuration and operational status.

If the Kernel detects that it did not terminate normally, it automatically enters a restart mode. This situation will occur following a fault which caused the Kernel to terminate abnormally such as a system failure or power failure. When the Kernel enters restart mode, it establishes communications with the other processor and attempts to restart any applications that it was running before the failure that are not now being run by the other processor. Certain critical applications can be recovered quickly and re-run on the second processor in the case of a failure. However, the other lower priority tasks do not need to be run on the recovery system. When the failed system becomes available again, those remaining applications will automatically be run on that system.

Mirror Disks

Operation

During normal operation, any write to a mirrored disk will result in writes to both disks. This not only includes writes to the files themselves, but also file allocation, file renaming, etc. Reads will be scheduled from one disk only. Disks are optionally designated as being mirrored by the MARK ON command. When a disk is marked on, or made available to, the operating system, an

Mirror Disks (Continued)

optional parameter can be supplied which designates the mirrored disk. At this point, certain mirroring criteria must be met. The synchronization stamps must agree and a synchronization bit must be set on both disks.

If these criteria are met, then both drives will be marked on and normal mirroring operations will proceed. If the mirroring criteria are not met, the operator will be advised that synchronization is necessary and the Synchronization Utility must be run.

In the event of an unrecoverable I/O failure during operation, the mirror of the disk on which the I/O failed will be marked as unsynchronized. Whichever disk fails, the operator will be alerted that an error has occurred and that the disks are now unsynchronized. Depending on the reason for the failure, the operator may choose to synchronize the failed disk or synchronize a completely different disk which must have the same name, but may be on a different drive. The normal I/O will continue to the remaining disk without affecting the users or applications.

Synchronization Utility

The Synchronization Utility is run to synchronize a new disk to an existing disk as well as to resynchronize a disk which has been temporarily unavailable. System and data availability is uninterrupted during the synchronization process. The utility allows writes to both disks but reads are from the "good" disk only. A verify option is also provided to check the integrity of the procedure.

Physical synchronization means that bit for bit, the two disks of a mirrored pair are identical. The Synchronization Utility reads one disk and writes the information to the other. Both disks must be of the same device type, e.g. 300MB disks. The minimum data transferred is a track, with increments of one track up to a maximum of a cylinder. The memory segment size provided when the utility is loaded determines the amount of data transferred. The Synchronization Utility merges the bad sector information of both disk packs before synchronization begins in order to keep the units identical.

System Requirements

Minimum Software Requirements

- OS/32 Revision 7.2 or higher
- Reliance PLUS recommended for resilient transaction processing applications.

Minimum Hardware Requirements

- Any two (2) Perkin-Elmer 32-bit systems each with 50K bytes of memory over and above that required by the operating system and applications.
- A dedicated system disk per CPU.
- Appropriate disks, bus switches, peripherals, etc. to support the required configurations.

Product Numbers

S70-070 Resilient System
Includes one (1) copy each, plus one (1) additional CPU license of the Reconfiguration Monitor and the Mirror Disk option.

S70-071 Mirror Disk Option Only
(for a single CPU)

Related Documentation

48-030–OS/32 Operator Reference Manual
48-040–System Programmer's Reference Manual
48-123–Reconfiguration Monitor User Manual
48-124–Resilient System Overview

Worldwide Sales Offices

U.S.A Offices
ALABAMA: Huntsville; ARIZONA: Phoenix; CALIFORNIA: Los Angeles, Sacramento, San Diego, Santa Clara, Tustin; COLORADO: Denver; CONNECTICUT: Fairfield, Hartford; FLORIDA: Orlando; GEORGIA: Atlanta; ILLINOIS: Chicago, Springfield; KANSAS: Kansas City; MARYLAND: Rockville; MASSACHUSETTS: Boston; MICHIGAN: Detroit; MISSOURI: St. Louis; NEW JERSEY: Cherry Hill, West Long Branch; NEW MEXICO: Albuquerque; NEW YORK: Binghamton, Lake Success, New York City, Rochester; NORTH CAROLINA: Charlotte; OHIO: Cleveland, Dayton; OKLAHOMA: Oklahoma City, Tulsa; PENNSYLVANIA: Pittsburgh; TEXAS: Dallas, Houston; VIRGINIA: Richmond; WASHINGTON: Seattle.

Major Subsidiaries
AUSTRALIA: Adelaide, Albury, Brisbane, Canberra, Melbourne, Perth, Sydney; and NEW ZEALAND: Wellington; BELGIUM: Brussels; CANADA: Calgary, Montreal, Ottawa, Toronto, Vancouver; ENGLAND: Manchester, Slough; FRANCE: Arcueil, Bordeaux, Grenoble, Lille, Lyon, Perigueux, Toulouse; GREECE: Athens; ITALY: Milan; WEST GERMANY: Dusseldorf, Frankfurt, Munich, and AUSTRIA: Vienna; NETHERLANDS: Gouda; SINGAPORE; SWITZERLAND: Zurich; HONG KONG; JAPAN: Tokyo. Other countries are served by a network of distributors.

The information contained herein is intended to be a general description and is subject to change with product enhancement.

EVERYWARE...EVERYWARE...EVERYWARE...EVERYWARE...

PERKIN-ELMER
Data Systems Group
2 Crescent Place • Oceanport, N.J. 07757
(201) 870-4712 • (800) 631-2154

October 1983
PB321103
Printed in U.S.A