

**508**

90A060820

COMBINED EVENT PERFORMANCE TRACE FOR AIX

Disclosed is a design that permits event and performance tracing of system software execution natively in the AIX\* environment. Prior to this implementation, performance tuning on AIX was done primarily by the use of additional hardware and a connection to a separate processor. Key to the design is the efficient availability of a high resolution timer. Such a timer has been architected as a special-purpose register for future RISC processors.

The trace data collection mechanism consists of a device driver, trace recording routines, a process that reads the data and a process that monitors and modifies the trace state. The device driver allocates and pins trace buffers during open and configures the trace for the specified management of those buffers (circular fill, single sweep, last buffer, or continuous to file). Trace data is passed in registers to trace recording routines which record the data for the event atomically in the buffer along with a time stamp. The use of routines, rather than inline code, permits the routines to be different for various trace configurations (for example, the data can be directed to the bus rather than to memory, or events can be real time filtered). When trace is inactive, these routines simply return. Activating trace incurs a very small delta to system overhead such that execution flow is not appreciably altered. The reading process reads the trace buffers and writes them to file. Its reads are blocked until the buffers are ready to be cleared. The device driver provides a set of ioctl functions to start and stop collection of events. These controls may be issued around sections of code to be traced. Additionally, a monitor process is active that is designed to catch keyboard signals and issue these ioctl's on behalf of the user. This permits the user to start and stop collection of data with a simple keyboard sequence.

The trace report facility formats trace entries according to rules provided in a trace format file. This file contains a stanza for each hook id in use. This approach permits users to add additional trace events and have them formatted by the report facility by adding stanzas to this file. A set of id's are reserved for users for this purpose.

\* Trademark of IBM Corp.