SystemGuard Processors for Remote Operations

By Kanti C. Shah and David F. Rittinger

This article describes the operation of the SystemGuard processor shipped with every IBM SMP system. Using the SystemGuard processor, system administrators can manage the server operations locally or remotely—even when the network, operating system, or individual hardware components are not functioning.

IBM's Symmetric Multiprocessor (SMP) products are designed to be servers in commercial environments. Commercial servers, shared by many users, typically run critical applications in which data integrity and availability are very important. Since many servers operate in unattended or remote locations, providing remote diagnostics and service is important while, at the same time, protecting access to sensitive data.

IBM's SMP family of servers includes a SystemGuard processor as a standard feature, making it possible to manage the server operations locally or remotely. In addition, it automatically alerts IBM service personnel to system problems and allows troubleshooting to be performed remotely.

Some reliability, availability, and serviceability features of IBM SMP products are as follows:

- **Data integrity**: Extensive parity checking is used on system cache memory and all internal and external buses. All IBM SMP systems use Error Checking and Correcting (ECC) memory subsystems.
- **Low downtime**: Hot pluggable/swappable disk drives are available on the Model J, which increases the uptime for the system.
- **Remote maintenance**: All facilities that can be accessed locally are accessible remotely. Menu-driven maintenance sessions aid in quick diagnosis and fault isolation.

  Critical problems are automatically reported to the technical support/service center for fast turnaround and service. The SMP system can operate with less than a full complement of resources (such as CPUs and memory), thereby increasing availability. Performance is reduced, but the system can remain operational until repairs are made.

  All remote accesses are password-protected at several levels. There is a quick disconnect feature in case of an attempted intrusion. A remote session can be mirrored at a local site or administration center.
- **Quick fault isolation**: Firmware-resident diagnostics are run when certain failures prevent the system from booting. Stand-alone diagnostics provide for quick and efficient disk and I/O fault isolation. Since online diagnostics can execute concurrently with applications and predict potential failures, preventive maintenance can be scheduled.
- **Configurability**: Permissible accesses are configurable via user settings in offline mode and from AIX service aids. An interface between internal firmware and the operating system makes these tasks transparent to users and fully executable from System Management Interface Tool (SMIT), service aids, or the command line.

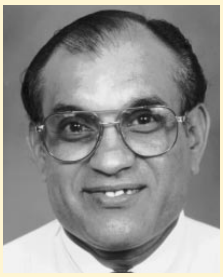

Kanti C. Shah

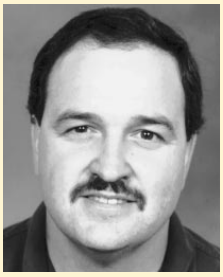

David F. Rittinger

SystemGuard Processor

The SystemGuard processor controls the server functions listed in Figure 1.

Firmware in the SMP servers (independent of the operating system) interfaces with the SystemGuard processor to perform the following types of functions:

- Power system management
- Interface with operator panel
- Vital Product Data (VPD) management
- Management of test system during POST
- IPL process control
- Offline maintenance sessions
- Remote service access

SystemGuard Processor Access

The SystemGuard processor operates on separate power from the server, so administrators can work with the SystemGuard processor even when the server is off.

The SMP servers include three asynchronous serial ports. Two ports can attach consoles, either locally or remotely, to support operations and maintenance functions; the third port can attach an Uninterruptible Power Supply (UPS). These serial ports can also be used for other customer applications when they are not being used for the functions described above.

The SystemGuard processor can be accessed with an ASCII terminal (or PC with terminal emulation) that is locally or remotely connected to one of the serial ports. With a modem on two ports, IBM service personnel can remotely connect to one serial port while a system administrator monitors activity from the other serial port. Therefore, all remote console activity is mirrored so it can be monitored locally. Customers can control all access authorization using authorization flags, passwords, and mirroring of console sessions.

A system administrator can display and change the system operating mode (Normal, Secure, Service), issue a power-on command, and monitor the hardware boot sequence from an ASCII terminal, either locally or remotely. These functions remove the need for on-site personnel to perform operations such as turning keys, pushing buttons, and reading display codes at the operator panel.

The SystemGuard processor interfaces to the main system processor via interrupts, with polling aided by a mail box in NVRAM and a JTAG bus for initialization, and logging out machine internal states if a catastrophic hardware error halts the system operation.

Server initialization
Power system management
Power-On Self Tests (POSTs)
System configuration
Boot/Initial Program Load (IPL) process management
Diagnostics testing
Firmware updating
Electronic-mode switching (Normal, Secure, Service)
Error handling
Automatic server recovery
Remote operation
Remote maintenance
Automatic problem reporting
Security parameters

Figure 1. Server functions controlled by Support Processor

The interface to the operator panel, power system, and VPD EEPROMs is through the industry-standard I2C bus. A local, proprietary system-specific bus interfaces with programmed-code ROMs, flash EEPROM, NVRAM, and other direct and system I/Os, such as Time of Day (TOD), serial ports, and diskette drive.

Stand-by Mode

When the server is turned on, the SystemGuard processor executes built-in core tests to ensure all base components are functioning. By reading the VPD, it sets up a configuration table that is used to schedule Power-On Self Tests (POSTs).

In Standby mode, using menu-driven maintenance sessions, IBM service personnel—either locally or remotely, if remote access is authorized—can perform problem determination functions such as the following:

Electronic mode setting: If the machine's key is in normal position, the system can be placed in Service mode electronically from either a local or remote console. Maintenance menus are accessible only in Service mode. The Standby menu is displayed when a preset keyword is entered.

Read/display system configuration: Server configuration data can also be electronically read and displayed on the operator panel LCD by activating the reset button, or displayed on the attached console.

Set configuration: Processor and memory can be disabled or deconfigured under user control. The expansion units can be configured logically and physically.

Setting operational and test mode flags: Options can be set to bypass maintenance menus in Service mode and proceed with IPL, bypass POSTs, perform extended tests, and so on.

Other functions: Service personnel can also enable/disable remote service authorization, read/write NVRAM, read TOD, verify (read/write) operator panel display functions, check power system status, adjust voltage margins, and verify that power can be applied to the system unit and attached devices.

The following sections describe three phases completed by the SMP system before booting the operating system.

Power-On Phase

The power-on sequence can be started using one of three methods:

- Power-on push button on the operator panel
- Power-on command string from a local or remote console
- Preprogrammed TOD parameter

This last power-on feature is useful in saving energy during service-off hours, while allowing the system to be ready well before the service-on time begins.

The SystemGuard processor schedules and launches each POST on a selected system processor and gathers the results.

System Initialization Phase

When the power-on sequence begins, the SystemGuard processor starts the INIT phase. This initializes all system hardware and establishes the system configuration tables for various hardware components such as the CPU or Central Electronic Complex (CEC), memory boards, and Micro Channel adapters.

Firmware Update Phase

Normally, the system uses the code resident in the flash EEPROM. But if the SystemGuard processor determines that its checksum is not valid and cannot update the firmware (the key is not in service position or the correct diskette is not installed), the system will detour to the backup EPROM. Although the backup firmware has limited functionality, it allows the system to boot the operating system in uniprocessor mode. Then, the AIX filesystem and tools can restore or rebuild the firmware in the flash EEPROM. These operations are also possible from a remote center. In service mode, if a diskette is inserted in the drive, the SystemGuard processor will automatically update the flash EEPROM.

Power-On Self Tests

The integrity of each principal hardware component is tested by running POST diagnostics. If an error is detected by any of these quick-confidence POSTs, extended POSTs are run to further isolate the problem.

The SystemGuard processor or the main system processor runs these tests, but the test system is entirely managed by the SystemGuard processor. The SystemGuard processor schedules and launches each POST on a selected system processor and gathers the results. Any errors that occur are logged and the test sequence continues. The operator panel display and the consoles indicate test progress. The SystemGuard processor performs the self tests on directly attached I/O ports (not Micro Channel adapters or devices attached to I/O ports), the power system, and fan connections. The main processor complex executes the POST on I/O ports shared with the SystemGuard processor, cache, memory subsystem, interrupt system, SCSI adapters, and LAN adapters. Memory sharing, cache coherency, system I/O sharing, and arbitration are also verified in multiprocessor mode.

The failure of a non-critical server component, such as one of the processors or a memory segment, does not prevent the server from booting AIX. The SystemGuard processor deconfigures failing resources, so the IPL can continue using any available processor. This is a significant enhancement to system availability. Replacing a failed component can be done later.

During POST, all memory is tested, optimum memory address mapping and interleaving are determined and configured, and an IPL control block including memory map is built for the operating system.

Critical errors, such as failing access tests to minimum resources, no functioning processors, insufficient working memory, or a parity error on any bus may prevent the server from booting. Not only are such severe errors reported on the consoles attached, but a trouble call is also automatically placed to the remote service center (if the system is configured for call home and appropriate authorization flags are set).

Service Mode

If the power-on sequence was initiated with the system operating mode set to Service, the hardware boot sequence can display the maintenance menu or continue to IPL in Service mode. If the system boots in Service mode, the operator can

choose single-user mode, diagnostic mode, or the RAM filesystem. Once exited from a single-user mode, the INIT default allows the operator to enter the normal multi-user mode.

In Service mode, an offline maintenance menu provides two levels of access—general maintenance and customer-privileged access—to set passwords and remote authorization parameters.

From the offline maintenance menu, a system administrator or authorized IBM service representative (local or remote) can set hardware configuration data, display hardware error-log information, enable or disable the service console, perform a system reset, issue a power-off command, select the device from which the system should continue booting, run internal diagnostic tests, display or change system parameters, or select the language for maintenance menus (English, German, Spanish, French, Swedish, Norwegian, Belgium/Dutch, or Italian).

The menu-driven interface allows the administrator to define the command strings or passwords that must be entered to remotely power-on the system. They can set voltage margins, security access passwords, phone numbers, and service-line speeds. Special reserved commands modify VPD EEPROMs after a field upgrade. Administrators can also disable or enable remote power-on command strings.

Offline diagnostics are resident in the system—independent of the operating system—and can be performed even if the system fails to boot the operating system. Offline diagnostics allow the administrator to build a test suite consisting of a sequence of resident tests, and to select specific test parameters for stressing the system or isolating a fault. Extensive diagnostic tests can verify whether the server can boot AIX, whether all installed resources are functioning, and whether the resources can be shared by the server's processors. Tests to verify the wiring interconnects across printed wiring assemblies can be run.

These offline diagnostics tests can be run in a loop mode, or they can be halted on every error and continue upon user command. These tests can display detailed test status messages on the consoles. If the test results are accumulated, they can be analyzed later using maintenance aid procedures to identify a failing replaceable unit.

Automatic Problem Reporting

During the hardware boot sequence, the SystemGuard processor controls the running of Built-In Self Test (BIST) and POST routines to ensure that hardware components, such as CPUs, ECC memory, and data paths are operating properly. If the test routines encounter a failure that prevents the machine from booting, the SystemGuard processor automatically reports a problem to the IBM Service organization. A problem management record will automatically be opened in the IBM RETAIN system (an automated system for tracking problems and dispatching service personnel), which is monitored by the IBM Support Center. The IBM Support Center can log in remotely, perform maintenance, and run diagnostics.

Security

All remote operations, maintenance, and automatic problem-reporting functions of the SMP product family have access authorizations that the customer can control, including the following:

- Attachment of remote consoles
- Activation of automatic problem reporting
- The ability to dial-in to the server to perform remote operations or maintenance
- Customer-assigned command strings to enable power-on and power-off
- Customer and maintenance passwords to control function access
- Console mirroring that allows the activity on the remote service console to be visible on the customer console
- The ability to disable remote sessions immediately by resetting flags or changing system operating mode

The SMP system also maintains an event log of all call-out and login attempts.

The menu-driven interface allows the administrator to define the command strings or passwords that must be entered to remotely power-on the system.

Console Mirroring

During offline sessions before AIX is up, and during the online runtime phase in Service mode, the serial port device drivers allow mirroring of both consoles. Input is accepted from both of them, and output is sent to both.

SystemGuard Processor Interaction With AIX

The SystemGuard processor communicates with AIX using a specialized protocol. Interacting with AIX, the SystemGuard processor can control automatic server recovery, allow a technician to remove and replace hot-plug disk drives (on Model J), select server operating mode (Normal, Secure, or Service), configure CPUs (processor

start, stop, disable, and enable), alter the SystemGuard processor operating parameters, manage flash EEPROM update, log single-bit correctable memory errors and failing addresses, monitor status of server fans and temperature sensors, and control system shutdown, power off, reset, and restart.

If a severe hardware component failure is detected in a CPU or an unrecoverable memory error occurs, the SystemGuard processor will automatically scan the internal and boundary states of the processor and memory subsystem complex for later analysis. It will initiate server recovery by running internal tests, deconfiguring the failing resource, and then rebooting AIX. Using the `surv_m` command, the SystemGuard processor can monitor the AIX heartbeat. If the operating system hangs and cannot generate the heartbeat, the SystemGuard processor will detect this condition and automatically initiate a reboot of AIX, if the operator has enabled surveillance mode.

Using selections from the AIX Diagnostic Service Aids menus, a system administrator can display and alter numerous server parameters, such as server operating mode, CPU allocation and deallocation, security parameters (flag settings, passwords, console visibilities), and Support Processor dial-in and dial-out telephone numbers.

Summary

The SystemGuard processor ensures that system administrators and authorized service personnel have access and control of the server, even when the network, operating system, or individual hardware components are not functioning.

The capability of remote operations for power-on, power-off, reset, system operating mode, and display of operator panel messages eliminates the need for personnel on-site, eliminates travel to unattended locations, enhances system administrator productivity, and improves system availability.

The automatic recovery functions of the server reduce the time needed to identify problems and improve system availability.

With hardware failures reported automatically—directly to the IBM Service organization—and the ability of IBM service personnel to remotely access configuration data, error logs, and run diagnostics for problem determination, server downtime and serviceability timeframes are reduced and system availability is improved.

Kanti C. Shah, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Shah, an advisory engineer, has been involved in various phases of design and testing of the SMP family of products. He has a BS in Electrical Engineering from Sardar Patel University in India and an MS in Electrical Engineering from the University of California at Berkeley.

David F. Rittinger, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Rittinger is a senior service planner involved in the remote support strategy and Reliability, Availability, and Serviceability (RAS) architecture testing of SMP. He has an MBA from Fairleigh Dickinson University in New Jersey.