

What Does an IPL Do?

The easiest way to explain the IPL process is to break it into groups of related tasks. In brief, an IPL does the following tasks:

- Executes power on self-test and basic assurance tests of the input/output processors (IOPs)
- Runs diagnostics on the service processor and initializes the licensed internal code
- Initializes the system with LIC
- Displays the Attended IPL menu or Install System menu on the system console
- Executes storage management recovery, journal synchronization, and IPL cleanup
- Loads OS/400 Now, let's look at each main process in detail and list the SRC codes displayed for each. I gleaned the information in this article from my AS/400 Model 50S; the SRC codes could be different depending on which AS/400 model you use. (Any Xs in the codes indicate that multiple SRCs appear during that particular task. The preliminary procedure in an IPL merely verifies that the system unit and control panel power supplies are operational. The IPL performs these tests before any SRCs are displayed.)

Service Processor Reference Codes

(LIC)

The first main function performed after the power supplies are tested is service processor testing, represented by SRC codes C1XX BXXX. The service processor card contains a set of instructions that constitute the logic required to start the system processor and handle the error messages that may occur during initialization. Here are the SRC codes that fall into this category:

- C100 B1D2—Basic assurance Read-only Storage (ROS) testing on the control panel interface. The system first tests the control panel, and if the control panel is not functioning correctly, the IPL cannot continue and terminates. Because this testing requires little time, you may not even see this SRC.
- C10X B111—Basic assurance ROS testing on Multifunction Input/Output Processor (MFIOP) control storage. As its name suggests, the MFIOP is a multifunction card in the system. Devices that can be attached to the MFIOP vary slightly, depending on the model of AS/400 used, but generally, the MFIOP supports internal tape drives, internal disk units, and the primary workstation controller. During this step, only the control storage portion of the MFIOP is tested.
- C100 B1E9—Basic assurance ROS testing on service processor registers. Registers are storage areas in which data and addresses are held temporarily while being used by a processor.
- C1XX B18X—Basic assurance testing on MFIOP. Those functions on the MFIOP card other than control storage are now tested.

Depending on the size and type of AS/400 you have, these tests require 1 to 5 minutes. After the service processor has been tested, the LIC must be loaded onto it. As the LIC loads, SRCs C1XX XXXX are displayed.

- C1XX 1030—Loading of the service processor LIC from the load source device. A partial IPL is performed on the system bus, and the load source IOP is initialized. Basic assurance tests are performed for a second time on the I/O devices, and the LIC is loaded onto the service processor using the load source device, an internal disk drive that contains all of the LIC and operating system.

System Processor Reference Codes

With the service processor tests completed and the service processor loaded, diagnostics are now run on the system processor. These diagnostics are represented by SRCs C3XX 41XX:

- C320 4135 through C32A 4135.
- C320 4136—Array Built-in Self-test (ABIST) on the system processor. These tests may differ based on the type of processor installed. Tests performed on single processors differ from those performed on multiple processors.
- C320 4190 through C32A 4190—Main storage diagnostics (MSD). (Where information is lean, I was unable to determine exactly what type of diagnostics are performed. IBM does not share that information with the general public. In addition, for any unidentified acronym, I simply listed whatever information was in the IBM manual.) This stage of the IPL varies, depending on the processor type and type of IPL performed. On average, this stage requires 2 to 10 minutes. After the system processor diagnostics have been completed, you may see C100 2060 (a tape-read command issued to the alternate IPL tape device) and C100 2090 (acknowledgement from the alternate IPL tape device).

System Initialization

The hardware has been tested, and C100 2034 is displayed. At this point, IPL control is passed to the system processor, which continues the IPL process. The next stage of the process is the testing and initialization of the system configuration, represented by SRCs C6XX 4XXX:

- C600 4001—Start static paging.
- C600 4002—Start limited paging/call LID manager.
- C600 4003—Initialize IPL termination data area/set up node address communication area (NACA) pointer.
- C600 4004—Check and update MSD subject identifier (SID). The SID is a string that identifies a user or set of users in the distributed computing environment (DCE), a set of services that support the development, use, and maintenance of distributed applications.

- C600 4005—Initialize event manager.
- C600 4006—IPL all buses. The AS/400 supports different bus structures, two of which are Peripheral Component Interconnect (PCI) and System Products Division (SPD). PCI is growing more popular because PCI cards are less costly than SPD. However, because not all devices can be attached with PCI cards, SPD cards still exist in high-end RISC models. During this step, all buses are initialized for all I/O devices.

- C600 4007—Start error log ID. An error log ID is created to log hardware and software errors that may occur.

- C600 4008—Initialize I/O service, and C600 4009—Initialize I/O machine. These two processes prepare the I/O devices to be used.

- C600 4010—Initialize interactive device exerciser (IDE).
- C600 4011—Initialize remote services.
- C600 4012—Initialize RMAC data values.
- C600 4013—Initialize context management.
- C600 4014—Initialize RM seize lock.
- C600 4015—Initialize MISR.
- C600 4016—Set time of day.
- C600 4017—Initialize RM process management.
- C600 4018—Initialize error log. The error log is prepared to receive log entries.
- C600 4019—Reinitialize the service processor. The service processor is used to start the system processor. This step resets the service processor.

- C600 4020—Initialize machine services.
- C600 4021—Initialize performance data collector. The performance data collector is prepared to gather information about the system regarding response times and throughputs. An example of such a job is job name QPFCOL running in the QCTL subsystem.

- C600 4022—Initialize event manager.
- C600 4023—Create Machine Interface (MI) boundary manager tasks. The Technology Independent Machine Interface (TIMI) is a logical rather than physical interface to the system hardware. The MI architecture provides a complete set of APIs for OS/400 and all application programs. The boundary manager provides the method of communication between the hardware and system software. Frank Soltis' Inside the AS/400 contains a detailed explanation of the MI. (See the References section at the end of this article.)

- C600 4024—Disable Continuously Powered Main Storage (CPM). This step is a little confusing. CPM is available on certain AS/400 models to supply main storage power for a short time to allow an orderly system shutdown in the event of power failure. CPM is disabled during this step and is made available at each IPL. CPM is enabled only when utility power is interrupted. It may be necessary to disable CPM to make specific repairs to the system.

- C600 4025—Initialize battery test. If the system has an internal battery, it is tested at this point. If the test fails, the system remains operational, but the system attention light may be lit and an SRC code may be displayed while the system is running.
- C600 4026—Hardware card checkout.
- C600 4028—Start dedicated service tools (DST). During an attended IPL, the DST menu is displayed at this point, allowing DST options to be used. Some of the options that might be used at this time are to start or suspend mirroring, add or remove disk units from the auxiliary storage pool (ASP), start or stop device parity protection or RAID, and other similar tasks where the system must be in a dedicated state.
- C600 4030—Free static storage.
- C600 4031—Destroy IPL task. The system performs a cleanup, removing unnecessary IPL job steps from the system.
- C600 4205—Synchronization of mirrored data. The system checks the integrity of data on mirrored pairs of disk units. If the last power-down was normal, this operation can take just a minute or so per each set of drives. However, if the last power-down was abnormal or you opted to start or stop mirroring from the DST menu, this step can take several hours, depending on the storage size of the drives.
- C600 4056—Journal recovery. If the system ends abnormally, database files in the journal are automatically recovered during this procedure. The database files are updated to reflect all activity recorded in the journal receivers. If the system ends abnormally, this may be a lengthy procedure.
- C600 4065—Start operating system. This function starts the operating system, which is loaded onto the AS/400. OS/400 is the operating system of choice, but, for the advanced 36, SSP is also part of the operating system.

Loading the Operating System

At this point, LIC initialization is complete, and the operating system has started. All of the hardware has been tested and verified. C9XX 2XXX are the tasks required to start the operating system:

- C900 2830—Resolve system objects. The first step in starting the operating system is to locate all of the system objects needed to start the operating system. In the system exists a resolve instruction that uses the name, type, and authority being requested from the unresolved pointer. The libraries on the library list are then searched until the object is found. Once located, the object is said to be resolved.
- C900 28C5—Initialize system objects. After all objects required to load the operating system have been located, they can then be used or initialized.

- C900 2910—Start system log. The system starts logging messages to the log file. If you display the QHST log after the IPL is complete, you can view messages logged from this point forward.
- C900 2920—Library and object information repository (OIR) cleanup. In SystemView System Manager/400, OIR consists of information about each object that identifies its associated product, such as release level, option, and load identifier.
- C900 2925—Verify POSIX root directories. POSIX is a collection of international standards for UNIX-style operating system interfaces. An example of where POSIX standards are used is the AS/400 Integrated File System (IFS) announced for V3R1.
- C900 2930—Database cross-reference.
- C900 2960—Sign-on processing. The system prepares for user access.
- C900 2965—Software Management Services (SMS) initialization. SMS provides the user with consistent distribution, installation, and service strategy. It allows you to save and install user-written application software as though it were licensed.
- C900 2A85—Load POSIX SAG.
- C900 2967—Applying PTFs. When PTFs are loaded onto the system, some of them are applied immediately while others affect hardware and system software and require an IPL to be applied.
- C900 2968—IPL options.
- C900 2970—Database Recovery, Part 1: Journal commit. If the last power-down was normal, this step should be fairly quick. If the last power-down was abnormal, the system recovers what it can from the journal receivers and automatically performs a rollback if a commit was not processed for files that were under commitment control. This option also rebuilds access paths if the system determines that logical files were open when the abnormal power-down occurred. This step can be time-consuming.
- C900 29B0—Spool initialization.
- C900 29C0—Write control block table. A control block is a storage area used by a program to hold control information. In this instance, the system sets up a table for system jobs to use.
- C900 2A90—Start system jobs. Some of the jobs that the system starts at this time are in the QSYSWRK and QALERT subsystems.
- C900 2AA0—Damage notification. Every system object contains header information pertaining to the object. The first header is called the segment header, and the second header is the Encapsulated Program Architecture (EPA) header. The EPA header contains an attribute byte that defines the object as permanent or temporary and determines whether or not the object is suspended or damaged. There are two types of object damage: hard or soft. An object with hard damage is not usable; it can only be removed. Soft damage indicates that some data can still be extracted from the object. One source of damage is bad sectors on a disk drive. If storage management cannot read these sectors, it uses the EPA header to flag the object as damaged.

- C900 2AA5—IFS directory recovery. The same function performed on the DB2/400 database is performed for the IFS. If an abnormal power-down occurs, this step may be extended.
- C900 2AC0—DLO recovery. The system recovers objects that may have been in use during an abnormal power-down or system crash. Folders are examples of DLOs.
- C900 2B10—Establish event monitors. An event is an activity during a machine operation that may be of interest to a user. An example of an event is an I/O operation, such as reading a record from a disk initiated by a read operation from an application program. The mechanism used to report completion of the I/O process is an event because it is caused by an action outside the application program currently executing. The actual I/O processing takes place at the MI level. System arbiter jobs are an example of event monitor. The system arbiter, identified by job name QSYSARB and QSYSARB2 through QSYSARB5, is the central and highest-priority job within the operating system. Each system arbiter responds to systemwide events that must be handled immediately and those that can be handled more efficiently by a single job rather than multiple jobs.
- C900 2B30—Start QCLUS job. The logical unit services, identified by job name QCLUS, support communication devices. The system arbiter starts QCLUS even if no communication devices are configured on the system. QCLUS is the event handler for logical unit (communication) devices and also acts as their manager.
- C900 2B40—Device configuration.
- C900 2C40—Work control block table cleanup. At this point, the system performs a cleanup on the control block table written in step C900 39C0.

Why Is My System Slow?

That was a high-level look at just about every SRC code you're likely to see during an IPL. When you see 01 B N displayed on your AS/400, you may think the IPL is finished. Well, not quite. Although the operating system initialization is complete when the sign-on screen appears on the console, internal procedures are still happening that are part of the overall IPL process. If you log on during this stage, you may discover that your response time is slower than normal. This slowdown happens because the last IPL event, running the startup program identified by system value QSTRUPPGM, occurs at this point. The startup program determines which subsystems should be started as well as any other functions you wish to run. The runtime for this program depends on the number of subsystems started and the number of devices under each subsystem that must be activated.

Use CHGIPLA to Customize Your IPL

To make IPL operation faster, you can specify the level of diagnostic testing. Starting with V4R1, a change was made to the Power Down System (PWRDWNSYS) command. There are three restart types that may now be specified:

- *IPLA—The value specified on CHGIPLA is used.

- ***SYS**—The operating system is restarted, and the hardware is restarted only if a PTF that requires a hardware restart is to be applied. In other words, the I/O processors are not IPLed unless a patch has been made to the software running on these processors.
- ***FULL**—All portions of the system are restarted, including the hardware. CHGIPLA, shown in Figure 1, has several options that you can use to reduce IPL time even further:
 - Restart Type is the same as PWRDWN SYS. You can specify *SYS or *FULL. The initial value of the command is *SYS. Hardware diagnostics specify whether certain hardware diagnostics should be performed during the IPL. The list of diagnostics is predetermined by the system and cannot be modified by the user. There are two options for these diagnostics: *MIN, whereby the system performs a minimum set of critical hardware diagnostics, and *ALL, whereby the system performs a complete set of hardware diagnostics (the shipped value for this attribute is *MIN).
 - Compress Job Tables specifies when job tables should be compressed to remove unused entries. Excessive unused entries can result in poor performance during IPL steps that process the table and during runtime functions that work with jobs.
 - Check Job Tables specifies when a damage check on job tables should be performed. The possible values are:
 - ***ABNORMAL**—job tables are checked during abnormal IPL only. This is the recommended setting.
 - ***ALL**—job tables are checked during all IPLs.
 - ***SYNC**—the job table checks are performed synchronously during all IPLs. The system maintains a product directory of all installed licensed programs. Normally, it is not necessary to rebuild this directory after initial installation of the system; it is rebuilt automatically when the operating system is installed. The possible values are:
 - ***NONE**—indicates the product directory is not fully rebuilt.
 - ***NORMAL**—rebuilds the product directory during normal IPLs only.
 - ***ABNORMAL**—rebuilds the product directory after an abnormal IPL.
 - ***ALL**—rebuilds the product directory after all IPLs.

Reducing Required IPL Time

Another method for reducing IPL time is to set the automatic performance adjustment system value to 0 (no adjustment) or 3 (automatic adjustment). A setting of 1 or 2 performs adjustments at IPL time. When you set your system to make adjustments at IPL time, performance settings are calculated based on the number of devices and network interfaces and the total amount of main storage. If your system is stable, these calculations have the same result each time and adjustments are not made.

To reduce the amount of time required to rebuild access paths in the event of an abnormal power-down, logical files may be kept in a journal.

Although this article may not make the rather dull process of an IPL seem interesting, I hope that I have provided some insight into the process and explained the new IPL options for Version 4.

References

AS/400 Basic System Operation, Administration, and Problem Handling (SC41- 5206-01, CD-ROM QB3AGO00)

AS/400 Master Glossary (SC41-5006-01, CD-ROM QB3AIG00) AS/400 Service Functions (SY44-5902-01) Inside the AS/400, 2nd Edition. Soltis, Frank G. Loveland, Colorado: 29th Street Press, 1997

Figure 1: CHGIPLA offers you several options for reducing IPL time.