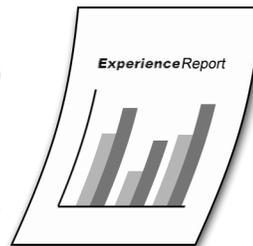


iSeries

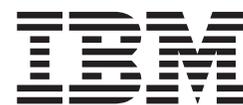


Reducing iSeries IPL time

Experience Report



iSeries



Reducing iSeries IPL time

Contents

Reducing iSeries IPL Time	1
Reducing IPL time - Jobs	1
Reducing IPL time - File Systems	3
Reducing IPL time - Configuration	4
Summary of Reducing IPL by Step.	7
Disclaimer	9

Reducing iSeries IPL Time

This report describes ways you can control how long it takes to start your iSeries^(TM) server. The Initial Program Load (IPL) for an iSeries server is made up of three major stages. First the hardware is powered up, second the Licensed Internal Code is initialized, and third i5/OS^(R) is started. This report focuses on performance aspects of i5/OS IPL that change based on how you configure and use your system.

A normal IPL occurs after the system is powered down with the Power Down System (PWRDWNSYS) command and no jobs have ended abnormally. All other IPLs are abnormal for i5/OS. An abnormal IPL takes longer because of additional recovery and verification needed. Message CPI091D (*Previous ending abnormal, reason code &1*) is sent to the history log during an abnormal IPL. This message gives the reason why the system end was abnormal. If you cannot use the Power Down System command to end the system, use option 7=Delayed power off from the panel or service tools. Avoid using 8=Immediate power off or 3=IPL restart as they can cause a significantly longer abnormal IPL.

The following areas have an impact on IPL performance:

“Reducing IPL time - Jobs”

Cleaning up jobs and spooled files. Starting jobs and preparing an execution environment.

“Reducing IPL time - File Systems” on page 3

Recovering database files, journals, and directories. Mounting file systems. ASPs and independent ASPs.

“Reducing IPL time - Configuration” on page 4

Recovering user profiles and libraries. Varying on devices. Setting IPL attributes.

Following are suggestions that could reduce your IPL time. Some of these suggestions could have an effect on runtime performance or ease of operating your system.

Reducing IPL time - Jobs

During an abnormal IPL, jobs that were active at the previous system end are cleaned up.

To reduce job and spooled file related IPL time you can:

1. Reduce the amount of joblogs produced.

The log parameter value of LOG(4 00 *NOLIST) in job descriptions and on the Submit Job (SBMJOB) command will prevent a spooled file containing the joblog from being produced when a job ends normally.

When ending jobs you can indicate that you do not want a joblog:

Command	Parameter for no joblog
End Job (ENDJOB)	LOGLMT(0)
End Prestart Jobs (ENDPJ)	LOGLMT(0)
End Group Job (ENDGRPJOB)	LOG(*NOLIST)
End Subsystem (ENDSBS)	ENDSBSOPT(*NOJOBLOG)
End System (ENDSYS)	ENDSBSOPT(*NOJOBLOG)
Power Down System (PWRDWNSYS)	ENDSBSOPT(*NOJOBLOG)
Clear Job Queue (CLRJOBQ)	LOG(*NONE)

Setting the QJOBMSGQMX system value to 8 will limit the amount of joblog messages for a job. This can reduce the recovery time for IPL after some system failures. You can set QJOBMSGQMX by using the Change System Value (CHGSYSVAL) command:

```
CHGSYSVAL SYSVAL(QJOBMSGQMX) VALUE(8)
```

You may also want to review the job message queue full action (QJOBMSGQFL) system value and job attribute. *NOWRAP will end jobs when the job message queue is full. This can be used to prevent

Table	Size	-----Entries-----			
		Total	Available	In-use	Other
1	1511168	1450	1161	289	0

You can set the QMAXJOB system value to the maximum number of jobs needed on your system during peak times. This limit will prevent additional jobs from being submitted or run on your system and using an unlimited amount of storage.

6. Do not compress the job tables every IPL.

Compressing the job tables during IPL increases the time needed to IPL and also forces the system to create job structures as new work enters the system after IPL. You should only compress the job tables if you have had an abnormally large number of jobs on the system and have cleaned up the unneeded jobs. This could be caused by runaway job submission or failure in the process of cleaning up old jobs and spooled files. The Change IPL Attributes (CHGIPLA) command can be used to turn off compression:

```
CHGIPLA CPRJOBTL(*NONE)
```

In situations where job table compression is needed, it is recommended that you set this value to perform compression on the following IPL only. The value is then changed to turn off compression (*NONE) once job table compression is started. You can specify to perform job table compression on the following IPL only by:

```
CHGIPLA CPRJOBTL(*NEXT)
```

7. Do not check the job tables every IPL.

During an abnormal IPL, all of the jobs in the job tables are cleaned up and checked for damage. This can take a significant amount of time during IPL. During a normal IPL this checking should not be needed. You can turn off job checking during normal IPL with the following command:

```
CHGIPLA CHKJOBTL(*ABNORMAL)
```

8. Delete unneeded job and output queues.

During an abnormal IPL, spool recovery verifies the jobs on job and output queues. You can remove any unnecessary job queues and output queues with the Delete Job Queue (DLTJOBQ) and Delete Output Queue (DLTOUTQ) commands.

9. Do not start more system jobs than you need.

Starting system jobs takes time during both normal and abnormal IPLs. You can reduce this time by not starting more communication arbiter system jobs than you need by using the following command:

```
CHGSYSVAL SYSVAL(QCMNARB) VALUE(*CALC)
```

Reducing IPL time - File Systems

During an abnormal IPL, files and related data are recovered. This can cause IPL time to be significantly longer than for a normal IPL.

To reduce file related IPL time you can:

10. Manage storage used in basic user ASPs so that they do not overflow.

The IPL following an ASP overflow requires additional IPL time. This can be reduced by monitoring the ASP threshold message, CPI0953, and following the recovery information in the message to avoid an overflow condition. For more information see the Change ASP Attribute (CHGASPA) command.

11. Use Independent ASPs for large databases.

Independent Auxiliary Storage Pools allow you to move data from one system to another to improve the availability of the data. They can also provide improved system availability. Independent ASPs are recovered after IPL when they are varied on. This can allow the system to IPL in less time and the Independent ASP vary on can complete at the same time as other system start up functions.

12. Avoid using the *IPL recovery attribute for database files.

Except for critical files, you probably want to delay rebuilding access paths until after your system has completed the IPL. For more information see Controlling access path recovery.

13. Use *TYPE2 integrated file system directories.

For more information and procedures to convert to *TYPE2 directories, see *TYPE2 directories in the Integrated file system topic of the iSeries^(TM) Information Center.

14. Use System-Managed Access-Path Protection (SMAPP) to reduce the time spent rebuilding access paths.

When you restart your system after it ends abnormally, the system must rebuild any access paths that were open when the system stopped. This rebuild time can be lengthy. You can reduce this time by journaling access paths, either explicitly by using the Start Journal Access Path (STRJRNAP) command or implicitly by using the System-Managed Access-Path Protection support.

15. Consider application changes to reduce journal recovery time.

- Separate your applications and tables/files and spread them across more journals. This allows the recovery to process the journals in parallel.
- For applications using the integrated file system, periodically issue FSYNC to minimize the time spent on recovery.

16. Consider managing journal receivers or using RCVSIZOPT(*MAXOPT3).

Management of journal receivers can also affect IPL performance. You can set up a journal so that the system manages when to detach the current receiver and attach a new one. This simplifies your job of monitoring the size of your journal receivers. However, for each journal that specifies MNGRCV(*SYSTEM), the system may perform processing during the IPL to ensure the maximum sequence number is not reached. This includes attaching a new journal receiver and resetting the journal sequence number. If you spend considerable time during IPL at SRC C900 2976, then you could consider either using MNGRCV(*SYSTEM) and RCVSIZOPT(*MAXOPT3), or switching to MNGRCV(*USER) for most of your journals. When RCVSIZOPT(*MAXOPT3) is being used with MNGRCV(*SYSTEM), the system will not perform this extra work at IPL to reset the journal sequence number unless the sequence number is approaching the limit. This should happen rarely.

17. Consider manually deleting journal receivers.

Similarly, you can set up the system to automatically delete journal receivers that it no longer needs for recovery by specifying DLTRCV(*YES) for the journal. This avoids excessive disk usage for journal receivers but it also increases your IPL time. If you spend considerable time during IPL at reference code C900 2AA2, you could consider switching to DLTRCV(*NO). However, you will need to monitor your system and remove the journal receivers manually to avoid disk usage problems.

Reducing IPL time - Configuration

There are a number of configuration choices that can affect IPL time.

18. Delete unnecessary user profiles and objects.

During some abnormal IPLs, authority recovery must be performed to verify the objects owned by a user profile and the amount of storage used by a user profile. You can reduce the storage used on your system and abnormal IPL processing time by deleting objects and user profiles that are no longer needed. The Work with User Profiles (WRKUSRPRF) command can be used to help manage the user profiles on your system.

19. Delete unnecessary device descriptions.

Whenever you perform an IPL, the system processes every device description on the system several times (to make performance calculations and to prepare for varying on the device). You can reduce this processing time by removing obsolete device descriptions from your system. The Work with Device Description (WRKDEVDD) command displays and allows you to work with device descriptions.

20. Use ONLINE(*NO) for network server, line, controller, and device descriptions if possible.

You can change configuration descriptions so that they do not vary on during IPL by using the ONLINE(*NO) parameter. After IPL, configuration descriptions can be varied on with the Vary


```
Previous system end . . . . . : Normal
Current step / total . . . . . : 14 48
Reference code detail . . . . . : C900 2965 20 50 0040

IPL step                               Time Elapsed   Time Remaining
Console configuration                   00:00:01
Sign on processing                      00:00:14
> SMS initialization                    00:00:00
Internal journal recovery
SAG initialization
```

This can be helpful during long IPLs. If your normal IPLs are consistently short, you may want to have IPL status only displayed during abnormal IPLs. This can save the time used to vary on the console and display the status during normal IPLs. You can use the following command to make this change:

```
CHGIPLA DSPSTS(*ABNORMAL)
```

Summary of Reducing IPL by Step

"How do I tell where the IPL is taking time on my system?"

You can use System Service Tools (STRSST command) to Display partition information and then Display secondary partition reference code history. This option shows a list of reference codes and the time they were displayed, so you can see how long each step took.

Below is a table of reference codes that are affected by the topics in this report.

Reference code	IPL Type	Report topic
C600 4272	All	File Systems # 10
C600 4025	Abnormal	Configuration # 18
C600 4058	All	File Systems # 15
C900 2830	All	Configuration # 19
C900 2920	All	Configuration # 21
C900 2940	Unattended	Configuration # 28
C900 2965	Unattended	Configuration # 26
C900 2973	Abnormal	File Systems # 11
C900 2976	Abnormal	File Systems # 16
C900 29B0	Abnormal	Jobs #1,2,4
C900 29C0	Abnormal	Jobs #8
C900 2A95	Abnormal	Jobs # 1,2,3,5
C900 2A90	All	Jobs # 9
C900 2AA5	Abnormal	File Systems # 13
C900 2AA2	Abnormal	File Systems # 17
C900 2AA3	Abnormal	File Systems # 11
C900 2AB0	Abnormal	File Systems # 12,14
C900 2B30	All	Configuration # 19
C900 2B40	All	Configuration # 19
C900 2C10	Unattended	Configuration # 28
C900 2C40	Abnormal	Jobs # 1,2,3,5,6,7
C900 2C25	Abnormal	Configuration # 25

Disclaimer

Information is provided "AS IS" without warranty of any kind. Mention or reference to non-IBM products is for informational purposes only and does not constitute an endorsement of such products by IBM.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve the performance improvements stated here.



Printed in USA