


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 OS/400^(R) is started. This report focuses on performance aspects of OS/400 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 OS/400. 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

short running jobs from looping forever in error situations. *WRAP will allow new messages to overlay old messages when the job message queue is full. This should be used for long running jobs.

2. Delete unneeded spooled files.

Often, jobs create spooled files that remain on the system for days or weeks. As long as those spooled files remain, an entry for the job remains in the job table. Deleting unneeded spooled files reduces IPL pool and job recovery time. The system can help you manage some system created spooled files such as joblogs and dumps. The Change Cleanup (CHGCLNUP) command can be used to change how long these spooled files are retained.

3. Detach spooled files.

If you cannot delete spooled files, consider detaching the spooled files, to reduce the number of jobs on the system. When spooled files are detached, the job that produced the spooled files is removed from the system, and can no longer be found with job commands such as Work with Submitted Jobs (WRKSBMJOB). However, the detached spooled files are still available in the output queue on the system. They can be printed or viewed using spool commands such as Work with Output Queue (WRKOUTQ) and Work with Spooled Files (WRKSPLF).

You can change the spooled file action system value to have new spooled files detached as jobs end by using the following command:

```
CHGSYSVAL SYSVAL(QSPLFACN) VALUE(*DETACH)
```

You can also change the spooled file action only for jobs using the same job description by using the following command:

```
CHGJOB JOB(QDFTJOB) SPLFACN(*DETACH)
```

For existing spooled files, you can use the Change Job (CHGJOB) command to get rid of the job without deleting spooled files that you need to keep.

4. Reclaim spool storage more frequently.

Set system value QRCLSPLSTG to 1 or 2 days to reduce the number of spool objects to be recovered at IPL by using the following command:

```
CHGSYSVAL SYSVAL(QRCLSPLSTG) VALUE('1')
```

This can reduce the number of spool objects to be processed during spool recovery during an abnormal IPL. Setting the system value QRCLSPLSTG to a value of *NONE can cause adverse runtime performance problems and is not recommended. For more information see the Reclaim Spool Storage (RCLSPLSTG) command.

5. Limit the size of the job tables.

When you perform an IPL, the system processes all of the jobs in the job tables. A very large number of jobs can significantly increase your IPL time. You can use the Display Job Tables (DSPJOBTL) command to monitor the size of the job tables on your system.

```

+-----+-----+
                Display Job Tables                                SYSTEM
                                                08/30/03 13:50:43
Permanent job structures:      Temporary job structures:
  Initial . . . . . : 30          Initial . . . . . : 20
  Additional . . . . : 10         Additional . . . . : 10
  Available . . . . . : 1161      Available . . . . . : 18
  Total . . . . . : 1450
  Maximum . . . . . : 163520
-----
Table      Size      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 CPRJOBTBL(*NONE)
```

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 CHKJOBTBL(*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:

1. **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.

2. **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.

3. **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.

4. **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.

5. **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.

6. **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.

7. **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.

1. **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.

2. **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 (WRKDEVD) command displays and allows you to work with device descriptions.

3. **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 Configuration (VRYCFG) command. Let TCP/IP vary on all the configuration descriptions that it requires. Also, let passthrough or telnet vary on all the virtual device descriptions they use.

4. **Delete unnecessary libraries and objects.**

During some abnormal IPLs, context recovery must be performed to verify the objects in a library. You can reduce the storage used on your system and abnormal IPL processing time by deleting objects and libraries that are no longer needed. The Retrieve Disk Information (RTVDSKINF) and Print Disk Information (PRTDSKINF) commands can be used to help manage storage use on your system.

5. **Turn off IPL performance tuning.**

The performance adjustment (QPFRADJ) system value controls when the system does automatic performance tuning. You can set up your system to dynamically tune performance, but not perform IPL tuning by using the following command:

```
CHGSYSVAL SYSVAL(QPFRADJ) VALUE('3')
```

6. **Use minimum hardware diagnostics.**

You can set the hardware diagnostics IPL attribute by using the following command:

```
CHGIPLA HDWDIAG(*MIN)
```

Then the system performs only a minimum, critical set of hardware diagnostics each IPL. This reduces IPL time. You should run the full set of hardware diagnostics, HDWDIAG(*FULL), when you add to or change your hardware configuration.

7. Allow the system to restart hardware only when needed.

You can set the restart IPL attribute by using the following command:

```
CHGIPLA RESTART(*SYS)
```

Then the system restarts the System Licensed Internal Code and the operating system. However, the system does not completely stop and then restart all the hardware functions. Usually, the IPL time for RESTART(*SYS) is less than the IPL time for RESTART(*FULL).

8. Run mail server framework recovery after IPL.

You can set your system to run mail server framework recovery after IPL by using the following command:

```
CHGIPLA MSFRCY(*NONE)
```

For more information see the Start Mail Server Framework (STRMSF) command.

9. Do not rebuild the product directory.

The system maintains an internal product directory of all installed licensed programs. The system rebuilds this product directory automatically when you install the operating system. For most systems, this directory should be stable. When you upgrade to a new release or install new licensed programs the product directory is updated. You can reduce your IPL time by setting the IPL attribute to not rebuild the product directory by using the following command:

```
CHGIPLA RBDPRDDIR(*NONE)
```

10. Use a high severity filter for the operator's message queue to reduce interruptions during attended IPLs.

You can use the Change Message Queue command to change the operator's message queue so that it only breaks for high severity messages:

```
CHGMSGQ MSGQ(QSYSOPR) SEV(90)
```

11. Change IPL status to only be displayed when you use it.

During IPL, status is displayed to show how the IPL is progressing:

```
+-----+
|                                     OS/400 IPL in Progress                                     |
|                                                                 8/30/03  16:09:22 |
| IPL: |
|   Type . . . . . : Unattended |
| Start date and time . . . . . : 08/30/03  16:07:15 |
| Previous system end . . . . . : Normal |
| IPL step . . . . . : PTF processing |
+-----+
```

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 # 17
C900 2830	All	Configuration # 18
C900 2920	All	Configuration # 20
C900 2940	Unattended	Configuration # 27
C900 2965	Unattended	Configuration # 25
C900 2973	Abnormal	File Systems # 11
C900 2976	Abnormal	File Systems # 15
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 # 16
C900 2AA3	Abnormal	File Systems # 11
C900 2AB0	Abnormal	File Systems # 12,14
C900 2B30	All	Configuration # 18
C900 2B40	All	Configuration # 18
C900 2C10	Unattended	Configuration # 27
C900 2C40	Abnormal	Jobs # 1,2,3,5,6,7
C900 2C25	Abnormal	Configuration # 24

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