



IBM i Work Management - Best Kept Secrets

Dawn May - dmmay@us.ibm.com
@DawnMayiCan



Session ID: 170151
Agenda Key: 47CM



Agenda


- Workload Groups
- Jobs held rather than ended if storage or CPU limits are exceeded
- Inactivity Timeout
- Disconnected Jobs Handling
- DSCJOB for QPADEV... devices
- Slow down of SBMJOB
- Multi-threaded job table cleanup
- Configuration of “Job Table Full” message threshold
- Print Enhancements
- PTF Enhancements



Workload Groups

Controlling work on IBM i

- IBM i has exceptional work management capabilities
- Many ways to control work within a partition:
 - Subsystems
 - Memory pools
 - Job queues
- How do you control a job from using too much of the capacity within a partition?
 - Run priority
 - Maximum CPU
 - Separate memory pools
 - Maximum number of threads in a memory pool
 - Maybe you resort to separate partitions and limit the number of processors allocated to that partition....
 - But then, you're not taking advantage of the work management capabilities!




IBM i Consolidation and Virtualization


Multiple workloads within a single operating system image has always been a hallmark of IBM i

IBM i clients often run very large, single partitions

Consolidation strategies have resulted in larger single partitions, with increasing number of cores



Power Systems hardware supports significant capacity




IBM i Workload Groups

Enhanced Management & Licensing


- Workload Groups provide the capability to manage and license workloads on IBM i
 - Limit the number of cores that are used by specific applications within single system/partition/subsystem
 - Limits placed at the whole processor/core level
 - Can cap a single job or all jobs/threads in a subsystem
 - Applications licensed for the number of capped cores

IBM i Yesterday



IBM i System / Partition / Subsystem

IBM i with Workload Groups



IBM i System / Partition / Subsystem

Workload Groups

Define the Workload Group

- Add Workload Group (ADDWLCGRP) to define the workload group
 - ADDWLCGRP WLCGRP(MYGROUP) PRCLMT(2)
 - WLCGRP - Name of the workload group you are defining
 - PRCLMT - The number of processors available to that workload group

```
                                Add Workload Group (ADDWLCGRP)

Type choices, press Enter.

Workload group . . . . . _____ Name
Processor limit . . . . . 1      1-256
```

- Can create up to 256 separate workload groups
- A workload can be associated with only 1 workload group
- A workload group can be associated with multiple workloads

Workload Groups

Associate with a Subsystem

- You associate the workload group with a subsystem
 - All jobs/threads in the subsystem cannot consume more processors than defined by the workload group
 - **Prior to IBM i 7.3**, you create this association with a data area
 - Pairs of 10-character subsystem name and workload group name
 - Up to 100 pairs can be specified

```
CRTDTAARA DTAARA(QSYS/QWTWLCGRP) TYPE(*CHAR) LEN(2000) VALUE('MYSBSNAME  
MYGROUP ' 'YOURSBS YOURGROUP ')  
TEXT('Subsystems to use workload groups') AUT(*USE)
```

- The subsystem must be restarted for any changes in the data area to take effect
 - Message CPI146C is sent to the subsystems job log during subsystem startup if the subsystem uses a workload group
 - Subsystem &1 is using workload group &2

Workload Groups

7.3

Associate with a Subsystem

- You associate the workload group with a subsystem
 - With **IBM i 7.3**, this is now part of CRTSBS and CHGSBS commands

Create Subsystem Description (CRTSBSD)

Type choices, press Enter.

Subsystem library	*NONE	Name, *NONE
Authority	*LIBCRTAUT	Name, *LIBCRTAUT, *CHANGE...
ASP group	*NONE	Name, *NONE
Workload group	*NONE	Name, *NONE

- Active subsystems can be changed
 - Any changes take effect for new jobs that are started in that subsystem
 - Message CPI146C is sent to the subsystem's job log during subsystem startup if the subsystem uses a workload group
 - Message CPI146C is sent to the subsystem's job log if the workload group parameter is changed when the subsystem is active
 - Subsystem &1 is using workload group &2

Workload Groups

for a Single Job

- You can associate a workload group with a single job
- Change Job command has the WLCGRP parameter
 - Change the workload group of a specific job with the CHGJOB command

```
Change Job (CHGJOB)

Type choices, press Enter.

Workload group . . . . . *SAME      Name, *SAME, *NONE
Maximum CPU time . . . . . *NOMAX    Milliseconds, *SAME, *NOMAX
Maximum temporary storage . . . *NOMAX Megabytes, *SAME, *NOMAX
Processor resources priority . . *SYSCTL *SAME, *SYSCTL, *NORMAL...
Duplicate job option . . . . . *SELECT *SELECT, *MSG
```

Workload Groups

Licensing

- Workload Group product entries define the license term and feature of the product that is limited by the number of processor cores defined for the workload group
- Add Workload Product Entry ([ADDWLCPRDE](#)) command
 - Workload Group - the name you gave on the ADDWLCGRP command
 - Product Identifier - the product which is being associated with this workload group
 - License Term - the release which is being associated with this workload group
 - *ALLINS for all installed releases
 - Feature - the installed features which are being associated with this workload group
 - *ALLINS for all installed features

Summary

Workload Group Commands

- Add Workload Group - ADDWLCGRP
- Remove Workload Group - RMVWLCGRP
- Change Workload Group - CHGWLCGRP
- Display Workload Group - DSPWLCGRP

- Add Workload Group Product Entry - ADDWLCPRDE
- Remove Workload Group Product Entry - RMVWLCPRDE

- Auditing
 - Workload groups operations result in a ZC (object change) journal entry being written to the QAUDJRN journal
 - Security audit records for the JS (Job Change) journal entry includes the workload group name (when jobs are started or changed)

Workload Groups

Performance Monitoring

- Once you have an application running in a workload group, how do you know what is happening?
 - Are the jobs being restricted too much?
 - Do you need to add additional CPUs to the workload group?
 - Do you have too many CPUs allocated to the workload group?

Collection Services and Workload Groups

- Collection Services collects metrics for workload groups
- Workload group performance metrics are in the [QAPMSYSWLC](#) file
- [QAPMJOBMI](#) was extended with job-level metrics for workload groups

Collection Services and Workload Groups

Fields in QAPMJOBMI:

JBFLDR2	Workload group delay time (in microseconds). The amount of time this thread could not be dispatched due to workload capping.
JBFLDR3	Workload group. The identifier for the workload capping group this thread belonged to at the time this data was sampled. A value of zero is reported when no group was assigned.

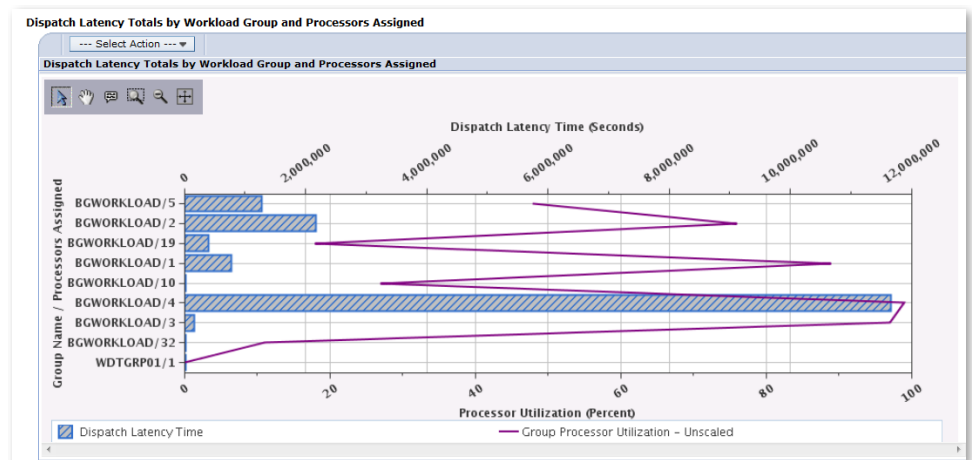
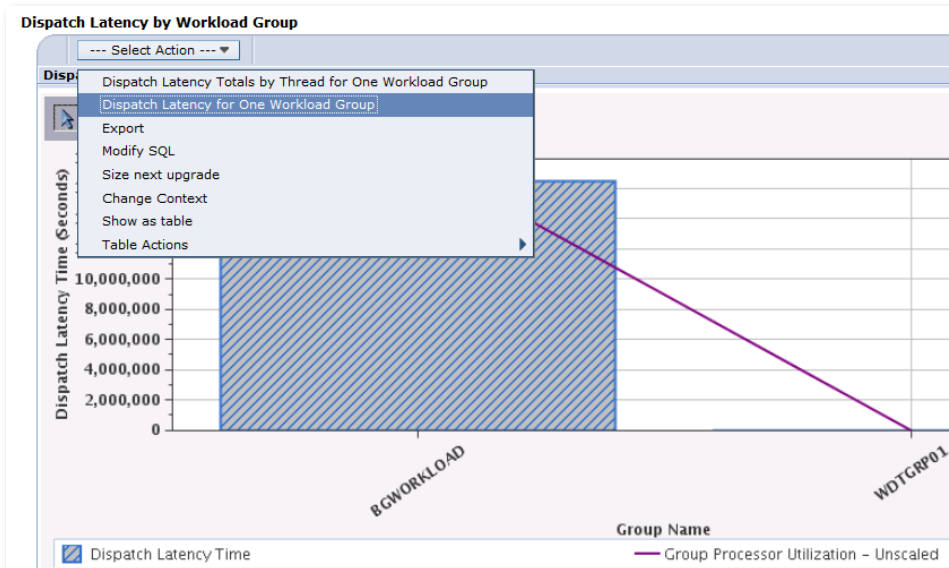

Fields in QAPMSYSWLC

SWGROU	Group ID. The identifier for the workload group.
SWGNAME	Group Name. The name assigned to the workload group when allocated by License Management
SWPRCASN	Processors assigned. The maximum number of processors which may be used concurrently by all threads of all processes which are associated with the workload group. This is the value associated with the group at the time data was sampled.
SWPRCAVL	Processor time available (in microseconds). The amount of processor time that this group had available to it based on the number of processors assigned to the group over time.
SWPRCUSE	Processor unscaled time used (in microseconds). The amount of unscaled processor time used within threads assigned to this group. Note: This does not include time charged to a thread by server tasks.
SWSPRCUSE	Processor scaled time used (in microseconds). The amount of scaled processor time used within threads assigned to this group. Note: This does not include time charged to a thread by server tasks.
SWDELAY	Dispatch latency time . The amount of time ready to run threads could not be dispatched due to the group's maximum concurrent processor limit.
SWPRCADD	Processes added. The number of process instances that became associated with this group during the interval.
SWPRCRMV	Processes removed. The number of process instances that were disassociated from this group during the interval.

Performance Data Investigator

Workload Group Perspectives

Graphs to display workload group dispatch latency




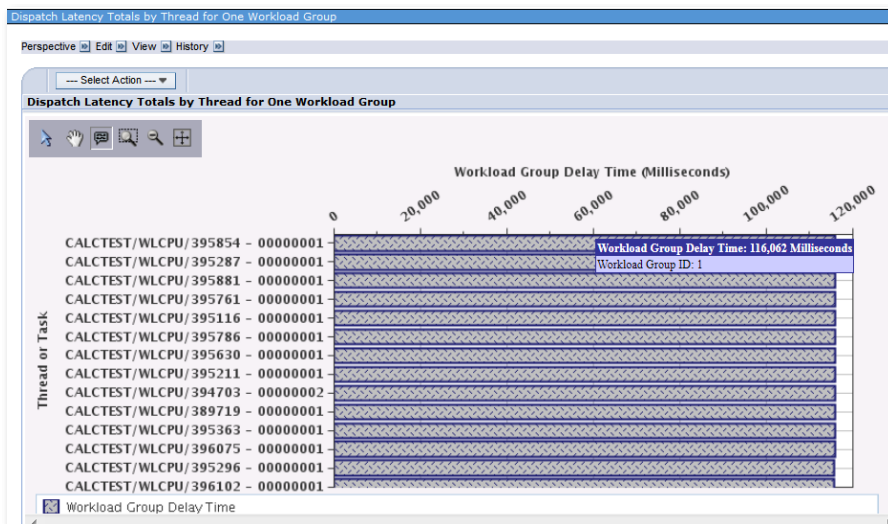
Description

This chart shows an overview of workload group dispatch latency. It shows the total delay time for each workload group. This is the amount of time threads that were ready to run could not be dispatched due to the group's maximum concurrent processor limit.

Performance Data Investigator

Workload Groups

- Drill down to see all the jobs running in the workload group
- View delay latency time per job
 - Drill down on a single job
 - View when the latency for a job occurred



Performance Data Investigator

Workload Group Data for Jobs

Workload group delay time

- The amount of time this thread could not be dispatched due to workload group.

Workload group

- The identifier for the workload group this thread belonged to at the time this data was sampled.

CPU -> CPU Utilization by Job or Task

CPU Utilization by Job or Task

Select	Job Number	CPU Utilization (Percent)	CPU Time (Milliseconds)	Scaled CPU Utilization (Percent)	Scaled CPU Time (Milliseconds)	Workload Group Delay Time (Milliseconds)	Job Type	Job Subtype
<input type="checkbox"/>	392352	0	1826	0	1826	2509313	B	
<input type="checkbox"/>	394767	0	1826	0	1826	2514772	B	
<input type="checkbox"/>	393796	0	1826	0	1826	2511459	B	
<input type="checkbox"/>	393955	0	1826	0	1826	2515482	B	

CPU -> CPU Utilization by Thread or Task

CPU Utilization by Thread or Task

Perspective Edit View History

Collection Name(s): Q341000005 Start: Dec 7, 2010 12:00:05 AM End: Dec 8, 2010 12:01:26 AM Library: WLCTEST Type: Collection Services File Based Collection File level: 36

System Name: RCHPOST3 Release: V6R1M0

Select	Scaled CPU Utilization (Percent)	Scaled CPU Time (Milliseconds)	Workload Group Delay Time (Milliseconds)	Workload Group ID	Job Type	Job Subtype	Minimum Job Pool	Maximum Job Pool
<input type="checkbox"/>	0	1822	2517868	1 B			02	02
<input type="checkbox"/>	0	1822	2519859	1 B			02	02
<input type="checkbox"/>	0	1822	2515353	1 B			02	02

developerWorks – IBM i Workload Groups

IBM i workload groups



| Updated April 19, 2012 by [ScottForstie](#) | Tags: None

Page Actions ▾

IBM i 7.1 now provides workload groups. Workload groups provide the ability to restrict a workload to a specified maximum number of processor cores within the partition it is running in.

A workload is defined as a job, subsystem, or product running on the IBM i system. The user or system administrator can define a workload group, assigning a specified number of processing cores to that group. The workload group is then assigned to a job or subsystem. Once the assignment has been done, the workload is limited to the defined number of processing cores. The system enforces this processing core assignment, ensuring that a job or all the jobs running (and threads) under the subsystem are not allowed to run on more processing cores than have been designated. The general concept is if a workload is designated to use a single core, the workload will behave as if it is truly running on a single processor core system.

Example of how workload works

A user has a batch job that is very CPU intensive. The user needs to run this job during the day but can't afford to impact the performance of their production system. By assigning this batch job to a workload group, this job can be put into a "processing container" to help ensure this job is kept to a limited amount of system capacity. If the workload group has a processor core limit of one, then the batch job and any threads running under that job will only be allowed to run on a single processor core. If this job is running on a multiple threaded core, multiple threads can be running for that designated batch job, but only a single core will be used at a time. This same concept also applies to jobs running under a subsystem that has been assigned to a workload group. All jobs and their associated threads will be limited to the number of processor cores specified in the workload group.

developer Works

Workload Groups MQ setup

sf99357	13
windows	13
esx	12
vmware	12
bladecenter	9
system_x	9
5.4	8
blade	8
sf99359	6
save	5

[View as list](#) | [cloud](#) | [all](#)


To learn more about the workload groups support, visit the IBM i Information Center at <http://publib.boulder.ibm.com/pubs/html/as400/infocenter.html>

Configure WebSphere MQ to use workload groups

Learn how to configure your systems to leverage workload groups for the WebSphere MQ product. Limit the amount of processing capacity available to this product to better fit your licensing needs based on the capacity being used for this product. To see the details, follow this link: [Limiting WebSphere MQ Licensing and Processing Capacity with Workload Groups.](#)

Updated Apr 19 by ScottForstie [Show Changes](#)

[Comments \(0\)](#) | [History of Changes](#) | [Attachments](#) | [About](#)

 [Subscribe to this page](#)

Instructions for setting up MQ to take advantage of Workload Groups

IBM i Workload Groups

References

- [Managing Workload Groups](#)
- [Setting up software licensing with workload groups](#)
- [developerWorks - IBM i Workload Groups](#)
- [developerWorks - Collection Services information](#)
- [Manage Workloads Better with Workload Groups](#) (blog)



i Can *Technical Tips for i*

Have you been reading the “[i Can](#)” blog?

Job CPU and Storage Limits

7.1

[SI42845](#)

The Class Object

Job CPU or Storage Limit

- The class object defines the processing attributes for a job, which includes:
 - Maximum processing unit time (CPUTIME)
 - Maximum temporary storage allowed (MAXTMPSTG)
 - Both have default values of *NOMAX
- If values were entered for these parameters the job would be ended if one of the limits was hit:
 - Ended due to maximum **processing** unit time
 - CPC1218 - Job ended abnormally
The job exceeded the maximum CPU time limit.
 - Ended due to maximum **temporary storage** allowed
 - CPC1217 - Job ended abnormally
The job exceeded the maximum temporary storage limit.



Class Object

MAXTMPSTG & CPUTIME values

On 7.2, MAXTMPSTG
is now Megabytes

Create Class (CRTCLS)

Type choices, press Enter.

Class	MYCLASS	Name
Library	QGPL	Name, *CURLIB
Run priority	50	1-99
Time slice	2000	Milliseconds
Eligible for purge	*YES	*YES, *NO
Default wait time	30	Seconds, *NOMAX
Maximum CPU time	*NOMAX	Milliseconds, *NOMAX
Maximum temporary storage . . .	*NOMAX	Kilobytes, *NOMAX
Maximum threads	*NOMAX	1-32767, *NOMAX
Text 'description'	*BLANK	

Was the job doing necessary work?

- The system can't know if the job...
 - was near completion of its processing
 - was doing important work that was essential to complete
 - It's possible that given a little more CPU time or temporary storage, the job would be able to run to completion
- Often these values were left at their **default** (*NOMAX) setting because
 - It is difficult to predict the upper CPU or temporary storage limits required by a job
 - The job would be ended when these limits were hit
 - Necessary work would have to be restarted

Improved control for jobs that exceed limits

- When a job exceeds its
 - Maximum CPU (CPUTIME) limit
 - Maximum temporary storage (MAXTMPSTG) or
- the system will **HOLD** the job instead of **END** it
- The job will stop consuming resources until the operator has a chance to determine if the job is running in error (ie. a runaway job) or if it should be allowed to continue.
- To let the job continue,
 - Change the CPUTIME or MAXTMPSTG job attributes via the **CHGJOB** command
 - Release the job with the **RLSJOB** command

Changed Behavior

- When a job is **held** by the system due to these conditions, a message is sent to the QSYSOPR message queue:
 - CPI112D - Job held by the system, CPUTIME limit exceeded
 - CPI112E - Job held by the system, MAXTMPSTG limit exceeded
- You can determine if the job should continue or be ended:
 - If you want the jobs to **continue to run**, you must:
 - Change the limit that was hit (CHGJOB)
 - Release the Job (RLSJOB) to let it run
 - (you cannot release a job that's above the limit)
 - End the job (ENDJOB) if you **do not want it to continue**

Change Job enhancements

- The Change Job command (and the Change Job API) were enhanced with additional parameters:
 - **Maximum CPU time (CPUTIME)**
 - The maximum CPU time parameter specifies the maximum processing unit time (in milliseconds) that the job can use. If the maximum time is exceeded, the job is held.
 - **Maximum temporary storage (MAXTMPSTG)**
 - The maximum temporary storage parameter specifies the maximum amount of temporary auxiliary storage (in megabytes) that the job can use. This temporary storage is used for storage required by the program itself and by implicitly created internal system objects used to support the job. (It doesn't include storage for objects in the QTEMP library.) If the maximum temporary storage is exceeded, the job is held.

CHGJOB command

CPUTIME & MAXTMPSTG parameters

Change Job (CHGJOB)

Type choices, press Enter.

Workload group	*SAME	Name, *SAME, *NONE
Maximum CPU time	*NOMAX	Milliseconds, *SAME, *NOMAX
Maximum temporary storage . . .	*NOMAX	Megabytes, *SAME, *NOMAX
Processor resources priority . .	*SYSCTL	*SAME, *SYSCTL, *NORMAL...
Duplicate job option	*SELECT	*SELECT, *MSG

Use the Class Object Parameters!

- Move away from the default *NOMAX values
- Set appropriate limits
 - Particularly the temporary storage limit
 - **Prevent** a system outage by setting an **upper limit** on the class object for the maximum temporary storage that a job can use
 - Be sure to keep that limit lower than the amount of storage available on the system!
 - With the new behavior *you* have the capability to *assess and determine* the best action for the job and the system

CPI112D message

Job hit CPUTIME limit

```
Message ID . . . . . : CPI112D
Message file . . . . . : QCPFMMSG
Library . . . . . : QSYS
```

```
Message . . . . . : Job &3/&2/&1 held by the system, CPUTIME limit exceeded.
```

```
Cause . . . . . : The current routing step for job &3/&2/&1 specifies a
maximum CPU time of &4 milliseconds.
```

```
Recovery . . . . . :
```

If you want the job to continue, use the Change Job (CHGJOB) command to specify a larger CPUTIME value for the job and then use the Release Job (RLSJOB) command.

If the application is in error, use the End Job (ENDJOB) command to end the job.

To avoid this condition in the future, use the Change Class (CHGCLS) or Change Job Description (CHGJOB) command to specify a larger CPUTIME value or use the Change Routing Entry (CHGRTGE) command to specify a different class.

CPI112E message

Job hit MAXTMPSTG limit

```
Message ID . . . . . : CPI112E
Message file . . . . . : QCPFMSG
Library . . . . . : QSYS
```

```
Message . . . . . : Job &3/&2/&1 held by the system, MAXTMPSTG limit exceeded.
```

```
Cause . . . . . : The current routing step for job &3/&2/&1 specifies a
maximum temporary storage limit of &4 megabytes.
```

```
Recovery . . . . . :
```

If you want the job to continue, use the Change Job (CHGJOB) command to specify a larger MAXTMPSTG value for the job and then use the Release Job (RLSJOB) command.

If the application is in error, use the End Job (ENDJOB) command to end the job.

To avoid this condition in the future, use the Change Class (CHGCLS) command to specify a larger MAXTMPSTG value or use the Change Routing Entry (CHGRTGE) command to specify a different class.

But *What if you want the old behavior?*

- Some clients want the old behavior where the jobs are automatically ended when the CPU or temporary storage limits are hit
- Use a message monitoring solution
 - CPI112D - Job held by the system, CPUTIME limit exceeded
 - CPI112E - Job held by the system, MAXTMPSTG limit exceeded
 - Use automation to end the job
- A message watch program can easily do this....

<http://www.ibmssystemsmag.com/Blogs/i-Can/Archive/How-to-End-Jobs-That-Are-Now-Held-for-Maximum-CPU-/>

Watch program

End a job that reaches CPU limit

```
/* **** */
/* Sample watch program to end maximum cpu job */
/* **** */
/* To create the watch program */
/* CRTCLPGM PGM(QGPL/WATCH) */
/* SRCFILE(QGPL/QCLSRC) */
/* USRPRF(*OWNER) */
/* **** */
/* To start the watch: */
/* STRWCH SSNID(*GEN) WCHPGM(QGPL/WATCH) */
/* WCHMSG((CPI112D)) */
/* **** */
PGM PARM(&WATCH_OPT &SESSION_ID +
&ERR_DETECT &EVENT_DATA)
DCL VAR(&WATCH_OPT) TYPE(*CHAR) LEN(10)
DCL VAR(&SESSION_ID) TYPE(*CHAR) LEN(10)
DCL VAR(&ERR_DETECT) TYPE(*CHAR) LEN(10)
DCL VAR(&EVENT_DATA) TYPE(*CHAR) LEN(600)
/* **** */
/* Message id and offset to message replacement */
/* **** */
DCL VAR(&EVENT_MSG) TYPE(*CHAR) STG(*DEFINED) +
LEN(7) DEFVAR(&EVENT_DATA 5)
DCL VAR(&OFFSET_MSG) TYPE(*UINT) STG(*DEFINED) +
LEN(4) DEFVAR(&EVENT_DATA 441)
DCL VAR(&WRKOFFSET) TYPE(*UINT) LEN(4)
/* **** */
/* Job name in message replacement data (CPI112D) */
/* **** */
DCL VAR(&MSG_DATA1) TYPE(*PTR)
DCL VAR(&JOBNAME) TYPE(*CHAR) STG(*BASED) +
LEN(10) BASPTR(&MSG_DATA1)
DCL VAR(&MSG_DATA2) TYPE(*PTR)
DCL VAR(&JOBUSER) TYPE(*CHAR) STG(*BASED) +
LEN(10) BASPTR(&MSG_DATA2)
```

```
DCL VAR(&MSG_DATA3) TYPE(*PTR)
DCL VAR(&JOBNMBR) TYPE(*CHAR) STG(*BASED) +
LEN(6) BASPTR(&MSG_DATA3)
/* **** */
/* Set no error detected, check for CPI112D message */
/* **** */
CHGVAR VAR(&ERR_DETECT) VALUE(' ')
IF COND(&WATCH_OPT *EQ 'MSGID ') THEN(GOTO +
CMDLBL(CONTINUE)) /* Message event? */
CONTINUE: IF COND(&EVENT_MSG *EQ 'CPI112D') THEN(GOTO +
CMDLBL(ENDJOB)) /* Held for CPU message? */
GOTO CMDLBL(EXIT)
/* **** */
/* Set pointers to job name in message data */
/* **** */
ENDJOB: CHGVAR VAR(&MSG_DATA1) VALUE(%ADDRESS(&EVENT_DATA))
CHGVAR VAR(&WRKOFFSET) VALUE(%OFFSET(&MSG_DATA1))
CHGVAR VAR(&WRKOFFSET) VALUE(&WRKOFFSET + %OFFSET MSG)
CHGVAR VAR(%OFFSET(&MSG_DATA1)) VALUE(&WRKOFFSET)
CHGVAR VAR(&MSG_DATA2) VALUE(%ADDRESS(&EVENT_DATA))
CHGVAR VAR(&WRKOFFSET) VALUE(&WRKOFFSET + 10)
CHGVAR VAR(%OFFSET(&MSG_DATA2)) VALUE(&WRKOFFSET)
CHGVAR VAR(&MSG_DATA3) VALUE(%ADDRESS(&EVENT_DATA))
CHGVAR VAR(&WRKOFFSET) VALUE(&WRKOFFSET + 10)
CHGVAR VAR(%OFFSET(&MSG_DATA3)) VALUE(&WRKOFFSET)
/* **** */
/* End the job that exceeded CPU */
/* **** */
ENDJOB JOB(&JOBNMBR/&JOBUSER/&JOBNAME) OPTION(*IMMED)
MONMSG MSGID(CPF0000)
EXIT: ENDPGM
```

Inactivity Timeout Changes



IBM Improves the Handling of the Inactivity Timeout

SI46398

More Changes for the Handling of QINACTIV

SI50502

7.1 PTFs, base 7.2 and later

Inactive job time-out

System Values

- Inactive job time-out (QINACTITV) system value
 - Minimum amount of time that an **interactive** job can be **inactive** before an **action** is taken
- Inactive job message queue (QINACTMSGQ) system value defines the **action** that is taken
- A related system value, time interval before disconnected job ends (QDSCJOBITV), defines **how long** a job can be disconnected before the job is ended

Inactive Job Time-out Improvements

[SI46398](#)

7.1

- Inactive job time-out (QINACTITV) system value
 - Improved accuracy
 - Subsystem checks for inactivity more often
 - Jobs ended much closer to specified time-out interval
 - Improved performance
 - Checking more frequently generally means fewer jobs will be ended at the same time (reducing spikes caused by ending many jobs).
 - System spreads the work so that subsystems don't all check for inactivity at the same time.

Inactive job time-out

(In)Accuracy

- Prior to this change, a job could be inactive for up to twice the specified interval before an action is taken
 - This level of uncertainty was too high!

- The subsystem would wait the

QINACTIV number of minutes

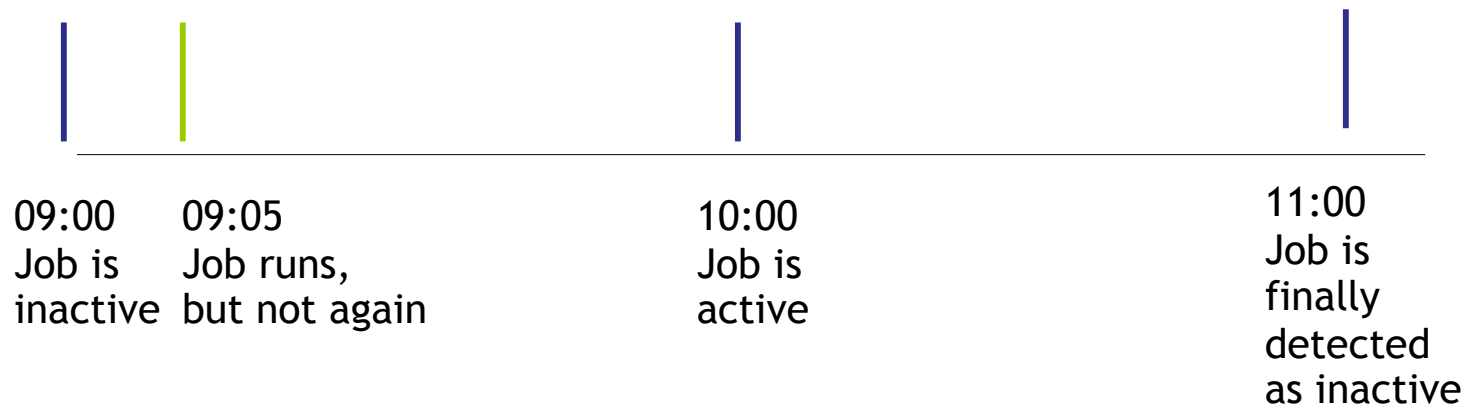
between the times it checked for inactive interactive jobs.

- If a job runs **just after** the system checks for inactivity, at the next check it will not be inactive for the full interval, so the timeout will not occur until the following interval.

Inactive job time-out

(In)Accuracy

Example with QINACITY set to 60
it can take nearly 120 minutes to detect an inactive job



Inactive job time-out

Improved Accuracy

- With this change, the subsystem will now check for inactivity **every 10 minutes**
- A job will be inactive for at least the QINACTITV value before action is taken
- Sometimes you might be interested in "at most" rather than "at least"
 - If jobs **must not be** left inactive for longer than some specified value
 - Set the QINACTITV value to something less than the limit that you want to enforce

Inactive job time-out

Improved Accuracy

Example with QINACITY set to 60

it will now take about 70 minutes to detect that inactive job

09:00
Job is
inactive

09:05
Job runs,
but not again

09:10
Job is
active

Checks every 10
minutes - job has
not been inactive
for 60 minutes

10:10
Job is
detected
as inactive

Inactive job time-out

Performance

- Subsystems that are started at the **same time** will no longer check for inactivity at the **same time**
 - When the subsystems start or when the QINACTITV system value is changed, the subsystems now use a **small but variable** delay when setting up the checking for inactivity
 - This helps reduce the number of jobs that are disconnected or ended at a single time.
 - With the old code, you could introduce delays between starting subsystems and these delays would translate more or less directly into delays between the times the different subsystems would check for inactivity.
 - With the new code, every subsystem has to check every 10 minutes, but the system takes care of spreading out the work.
 - PTF SI46398 spreads out the work for ending of inactive interactive jobs and for ending of unused prestart jobs.

What does it Mean for a Job to be Active?

[SI50502](#)

7.1

- Is there I/O occurring?
 - If yes, the job is active
- Is the job waiting?
 - DSPW means waiting for input from a user - the job is not active
- Is the job using CPU?
 - CPU time can be affected by several things that run in a job without any user interaction:
 - Java JIT threads
 - Polling applications that have something that runs in each job periodically
 - Retrieving the library list or call stack of another job, runs in that job
 - From the end user's standpoint, the job is inactive

More jobs will be ended or disconnected because of inactivity

Environment variable `QIBM_WT_QINACTIV_CPU` allows you to specify, in microseconds, the amount of CPU that can be used and still have the job considered inactive



Much of the information in this presentation was blogged about by the IBM i work management development team ...

Disconnected Jobs Handling for QPADEV... devices

7.1

[SI50502](#)

Disconnecting QPADEVxxxx Jobs

- System-named virtual devices - E.g., QPADEVxxxx - were not allowed to be disconnected
- System enforced the capability to reconnect to the job
 - When the session was gone it was unlikely for the same user to get the same device to reconnect the session
 - In order to reconnect to the job, the *same user* must sign back on to the *same device description*
- With this PTF -
 - If QINACTMSGQ is set to *DSCJOB and the user is using a QPADEVxxxx device
 - the job will now be disconnected rather than ended
 - the user can sign on and continue if they acquire the same device description

Reconnecting to a QPADEVxxxx device

```
—                               Display Program Messages

Job 231975/DMMAY/QPADEV0004 started on 04/28/14 at 08:26:55 in subsystem QINTER
All jobs at work station QPADEV0004 disconnected.
CPI1131 Job 231975/DMMAY/QPADEV0004 disconnected by user QSYS.
Job connected again. Sign on information ignored.
```

```
                               Sign On

System . . . . . : Z      1
Subsystem . . . . : QINTER
Display . . . . . : QPADEV0004

User . . . . . : DMMAY
Password . . . . : 
Program/procedure . . . . : 
Menu . . . . . : 
Current library . . . . : 
```

```
QINTER in QSYS: job entered system on 04/28/14 at 08:26:55
> dscjob
Job connected again. Sign on information ignored.
```

Disconnecting QPADEVxxx Jobs *(details)*

- This PTF allows more jobs to be disconnected. If your system is set up with [QINACTITV\(*NONE\)](#) or [QINACTMSGQ\(*ENDJOB\)](#), this PTF will probably have very little effect on the workload. So, for discussion, we'll assume that jobs are being disconnected for inactivity.
- There is a cost to start a new job, to end a job, to disconnect a job, or to reconnect to a job. Different applications will have different costs. While the overall amount of work is important, many systems are more strongly affected by how that work is distributed over time. Disconnected jobs continue to hold locks and use system resources. This is also true of inactive jobs.
- When a user signs on and is reconnected to a disconnected job, the system avoids the cost of ending the old job and creating a new job. When a large number of users reconnect, the work is generally well spread out over time.
- When a user does not reconnect before the time limit defined by the [time interval before disconnected job ends \(QDSCJOBITV\)](#) system value, the disconnected job is ended. By comparison to [QINACTMSGQ\(*ENDJOB\)](#), the system sees the extra work of a disconnect, but there is no work avoided. All jobs see the same QDSCJOBITV value, so the work is spread out the same way the work for QINACTITV is spread out.
- When a user does not reconnect and the subsystem is ended before the QDSCJOBITV time limit, the work of ending the disconnected jobs gets done all at once during the ending of the subsystem. This can significantly increase the stress on the system. Jobs that are inactive and not disconnected are also ended during the ending of the subsystem.
- The QDSCJOBITV system value should be set high enough to allow users enough time to reconnect to their disconnected jobs, but low enough that the disconnected jobs are likely to be ended before the subsystem ends. If users are not going to reconnect, the QINACTMSGQ should be set to *ENDJOB rather than *DSCJOB.
- One of the common problems with QINACTITV occurs when a user returns to a session at the same time the system is checking for inactivity. The system does not know that the user is there until the system sees input on that session and the system only sees input when the user presses enter or a function key.
- If QINACTMSGQ is set to *DSCJOB and the user is using a virtual device selected by the system (a QPADEVxxxx device), the job will now be disconnected rather than ended. The user can sign on and continue. The PTF should be very helpful, even though the handling of QINACTITV will never be perfect.

Slow down repetitive starting of new jobs

PTF Cover letter documents this changes
[SI40568](#) for 7.1

SBMJOB

- Have you ever had a bug in your code that does a SBJJOB (or maybe spawn) over and over and over and over.....
 - Until you used up all the job table entries ...
- Yeah, there's that message in QSYSOPR
 - Which you didn't see in time....
- And ...
 - Your partition failed because the job table got full
 - » Because no more jobs could be started



Speed bump for runaway jobs

[SI40568](#)

7.1

- Automatic slowdown of job submissions
 - SBMJOB, SPAWN()
- Delays start at a half second (at 90%)
- Delays increase as the number of jobs in the system approaches the QMAXJOB limit
 - 1 second delays at 95%
 - 2 second delays at 97.5%



Slow down starting New jobs

(additional detail)

- QMAXJOB system value - maximum number of jobs
 - 32,000 to 485,000 (or 970,000 on 7.2)

- Slow down the creation of new jobs
 - The default value for the threshold is 90 percent of the QMAXJOB value.

You can specify a threshold by creating a data area named QWTSBMDLY in library QSYS.

- The value is a percentage where the system needs to react to the fact that it is approaching the QMAXJOB limit.
- If you specify a value of 100, the system will provide no delays
 - With no protection against hitting the QMAXJOB limit.
- To set a threshold at 95 percent:

```
CRTDTAARA DTAARA(QSYS/QWTSBMDLY) TYPE(*DEC) LEN(3 0) VALUE(95) AUT(*USE) TEXT('Theshold for  
delays when creating jobs.')
```

- Changes to the data area are effective immediately.

Slow down Starting New jobs

(additional detail...)

- Delays start at a half second
- Delays increase (twice) as the number of jobs in the system approaches the QMAXJOB limit
 - The first increase occurs half way between the specified threshold and the QMAXJOB limit.
 - The second increase occurs half way between the first increase and the QMAXJOB limit.
- By default, the system uses a value of 90, giving a 0.5 second delay at 90 percent of the QMAXJOB value, a 1.0 second delay at 95 percent, and a 2.0 second delay at 97.5 percent
 - Specifying 94 causes the system to react at 94 percent, 97 percent, and 98.5 percent.
- If your system commonly has a very large number of jobs and if this is more than 90 percent of the QMAXJOB limit, you should create the QSYS/QWTSBMDLY data area and specify a higher threshold.



Miscellaneous things

Collection Services

7.2

Enhancements for Jobs

- Collection Services now tracks the number of:
 - Spooled files created by a job
 - Jobs submitted or spawned by a job

QAPMJOBOS file

» » JBSPLFC « «	» » Number of spooled files created by this job. This includes those spooled files which are spooled or restored under this job as well as all spooled files that are spooled under QPRTJOBS by this job. « «
» » JBSBMJOBS « «	» » Number of jobs submitted or spawned by this job. « «

Determining the Size of QTEMP

- There are many ways to determine the size of a job's QTEMP library
 - Display your own job and view your own QTEMP
 - DSPLIB QTEMP OUTPUT(*PRINT) - the total size of the library is at the bottom of the listing
 - You can use the [Retrieve Library Description API](#) to get the size of your job's QTEMP library
 - Work with an active job (WRKJOB) and look at the job's library list
You can then display the contents of the QTEMP library; you have to calculate the total size from the displayed list
 - Use the [Open List of Objects API](#) to get a list of objects in another job's QTEMP library
You need to call the API using optional parameter group 1 and specify the other job
The program would also need to sum the sizes of all objects returned to get the total size
 - Extending upon the option above, you can use the [Open List of Jobs API](#) to get a list of all active jobs and then for each job, use the Open List of Objects API for QTEMP
Be aware that the jobs can start or end after the Open List of Jobs API returns its results
 - Use the Call Job Interrupt Program Exit Program to run a program in a different job, which could be a program to determine the size of the QTEMP library for that job
- New with 7.2, the Retrieve Job APIs (QUSRJOBI, QGYOLJOB) have a new field with the size of QTEMP

7.2

<http://www.ibmssystemsmag.com/Blogs/i-Can/Archive/qtemp-temporary-or-permanent-storage/>
<http://www.ibmssystemsmag.com/Blogs/i-Can/March-2017/Determining-the-Size-of-QTEMP/>

Job Table Cleanup

- Job table cleanup happens during an abnormal IPL as well as the processing that is done when job tables are compressed
 - The job table cleanup IPL step is also known as Work Control Block Table cleanup (C900 2C40) IPL step.
- Prior to 7.1, the job table was processed by one thread during the IPL step; when the job table became large, it could take this one thread quite some time to process all the entries in the job table
 - In the 7.1 release, this single-threaded processing was changed to be multi-threaded. This change decreases the time the system is unavailable after an unexpected termination as well as when the job tables are compressed during an IPL
- The performance improvement is seen on systems with a large number of entries in the job table.
 - You can use the [Display Job Table \(DSPJOBTL\)](#) command to view the number of entries on your system.
 - The [Change IPL Attributes \(CHGIPLA\)](#) command allows you to influence how much cleanup is done at IPL with the Check job tables (CHKJOBTL) and Compress job tables (CPRJOBTL) parameters.
 - How the system powered down is the other influence on how much cleanup is done at IPL.
 - You can check the history on your system to see if you have had any recent abnormal IPLs with the Display Log command: DSPLOG PERIOD((*AVAIL *BEGIN)) MSGID(CPI0990 CPI091D)

IPL Progress Information

- There was also a 7.1 change that included a minor performance improvement for job table cleanup and a correction for the IPL status estimated time remaining

```
Operating System IPL in Progress                                10/24/12 15:21:10

IPL:
Type . . . . . : Unattended
Start date and time . . . . . : 10/24/12 15:14:43
Previous system end . . . . . : Abnormal
Current step / total . . . . . : 46 49
Reference code detail . . . . . : C900 2C40 20 D0 0810

IPL step                                Time Elapsed  Time Remaining
Device configuration                    00:00:01
Check user console                      00:00:07
Job table clean up                      00:03:03  00:02:31
SNA distribution recovery
MSF recovery

Item:
Current / total . . . . . : 149722 163470
Identifier . . . . . : I_CAN DAWN MAY 343449
Activity . . . . . : Checking job tables
```

- You should not have to run job table cleanup or compression very often
- It is one of those things that when it is necessary, time is most important

Some additional work-management topics

These are not new ... but maybe they are *new to you*

- CPF1240 and CPF1241 Messages Contain Summary Performance Data
- Handy utilities that are part of the Application Runtime Expert runtime (these utilities are included for free with the operating system)
 - Damaged Object Detection
 - Network Health Checker



Print Enhancements

Copy an Open Spooled File

- Previously you could not copy an open spooled file
- A new parameter, Copy open spooled file ([OPNSPLF](#)), has been added to the [CPYSPLF](#) command
 - Open spooled files may have data that is in internal buffers
 - When you copy an open spooled file, the data that is in internal buffers will not be copied

7.1

SI56003

7.2

SI57338

7.3

Included

[OSP-PRT Allow CPYSPLF to Copy Open Spooled Files.](#)

<http://www.ibmssystemsmag.com/Blogs/i-Can/March-2017/Copy-an-Open-Spooled-File/>

New Information on Display Open Files

7.3

- An additional panel available with F11

Display Open Files			
Job . . . : QPADEV0003		User . . . : DAWN MAY	Number . . . : 049002
Number of open data paths		2	
File	Library	File Type	Text
QDUI80	QSYS	DSP	Display file for 24x80 UIM Panels
QDDSP0F	QSYS	DSP	Display file for WRKJOB command

Initial Panel for Display Open Files

- WRKJOB, option 14, display open files, initially shows activation group information
- Create the data area specified below to display the file I/O data first, and use the F11 key to toggle to the activation groups only if needed.

Display Open Files								
Job . . . : DAWNMAYB		User . . . : DMMAY		Number : 026370				
Number of open data paths : 2								
File	Library	Member/ Device	Record Format	File Type	I/O Count	----Open--- Opt Shr-Nbr	Relative Record	
QDUI132	QSYS	DAWNMAYB	USRRCD	DSP	85	IO NO		
QDDSPOF	QSYS	DAWNMAYB	DETAIL2	DSP	23	IO NO		

- CRTDTAARA DTAARA (QUSRSYS/QDMDSPOFIO) TYPE(*LGL) LEN(1) AUT(*USE)

<http://www.ibmssystemsmag.com/Blogs/i-Can/November-2016/Displaying-Open-Files-with-Work-Job/>

QMAXSPLF limit - Different Limits by Job

- QMAXSPLF system value can be used to limit the number of spooled files each job can create
 - QMAXSPLF is a **system-wide** setting with a range of 9999 -999999

- Now you can set a **different** limit for a job
- To set a **higher** or **lower** limit in a specific job:

```
ADDENVVAR ENVVAR(QIBM_SP_MAXSPLF) LEVEL(*JOB) VALUE(99999)
```

7.1	SI54499
7.2	SI55214
7.3	Included

- The **Value** parameter can be from 1000 to 999999
- Each job needs to set the environment variable after it is started, in its initial program or during program setup
- **QIBM_SP_MAXSPLF** is ignored for spooled files created under a QPRTJOB jobs

WRKSPLF - *CURRENT Date

- WRKSPLF command was enhanced with a special value of *CURRENT for the date on the PERIOD parameter

Work with Spooled Files (WRKSPLF)

Type choices, press Enter.

Time period:

Start time and date:

Beginning time *AVAIL

Beginning date *BEGIN

End time and date:

Ending time *AVAIL

Ending date *END

Time, *AVAIL

Date, *BEGIN, *CURRENT

Time, *AVAIL

Date, *END

7.1

[SI55117](#)

7.2

Included

Copy Spooled Files to PDFs

SI43471

7.1

www.ibmssystemsmag.com/Blogs/i-Can/Archive/copy-spoiled-files-to-pdfs/

- PTF [SI43471](#) enhanced Copy Spooled File (CPYSPLF)
 - Parameters to specify the location the file is to be stored in the IFS
 - and a workstation customizing object (WSCST) value of *PDF
 - Transform Services (5770-TS1) licensed program product must be installed.

```
CPYSPLF FILE(MYSPLF) TOFILE(*TOSTMF)
      TOSTMF('/tmp/test')
      WSCST(*PDF)
```

- Additional WSCST values of QWPGIF, QWPTIFFPB, and QWPTIFFG4 (in the QSYS library)
 - Convert your spooled file to a GIF or TIFF image, using two different compression methods for TIFF. These WSCST values don't require the Transform Services product.

Sending Spooled Files via TCP

- You can work with spooled files through a variety of command interfaces
 - Work with Spooled Files (WRKSPLF)
 - Work with Output Queue (WRKOUTQ)
 - Work with Job (WKRJOB) option 4 (aka WRKJOB OPTION(*SPLF))
- These interfaces have option **1=Send** to send spooled files to another system
 - This send option invokes [Send Network Spooled File](#) (SNDNETSPLF)
- 7.1 PTF [SI50975](#), was delivered to provide the capability to modify the **1=send** option when working with spooled files to use [Send TCP Spooled File](#) (SNDTCPSPLF)

Sending Spooled Files via TCP

- To change the behavior to use SNDTCPSPLF:
 - Add an environment variable named QIBM_SNDSPFLF_OPT with a value of 'TCP'
 - The environment variable name and its value must be uppercase, and may be set at the *SYS level or *JOB, or both

Example:

```
ADDENVVAR ENVVAR(QIBM_SNDSPFLF_OPT) VALUE(TCP) LEVEL(*SYS)
```

- After you have added the environment variable, when you select option 1 on the WRKSPLF command, you will now see the prompt for SNDTCPSPLF

(No PTFs needed for the GUI)

QIBM_QSP_SPLF_LSTACT Exit Point

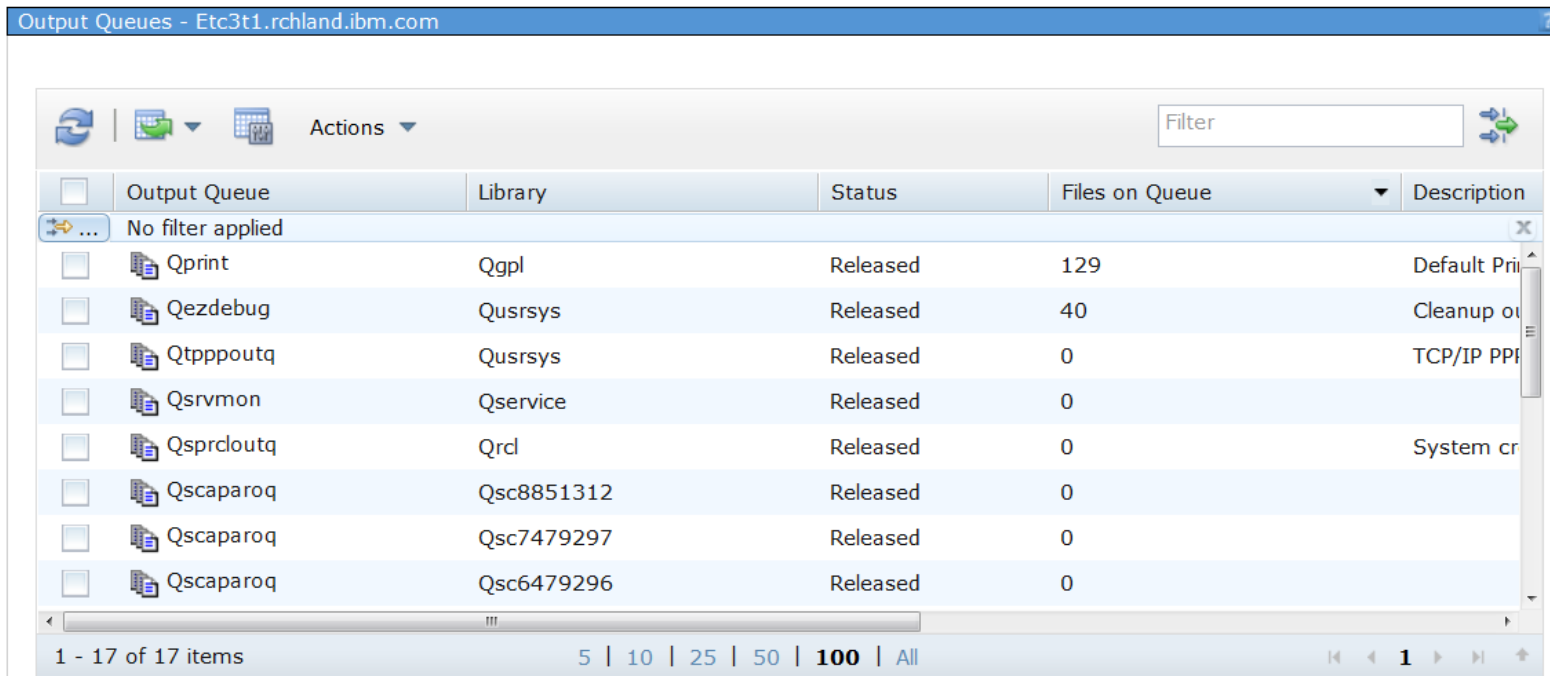
- The UIM Spooled File User-defined List Action exit point can be used to **define additional list actions** on the following green screen interfaces
 - Work with Printer Output (WRKSPLF ASTLVL(*BASIC))
 - Work with Spooled Files (WRKSPLF ASTLVL(*INTERMED))
 - Work with Job Spooled Files (WRKJOB OPTION(*SPLF))
 - Work with Output Queue (WRKOUTQ output queue name)
 - Work with Spooled File Status (D P or WRKSPLF DSPFMT(*S36FMT))
- For example, you could add an option to convert the spooled file to a PDF

www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_72/apis/xsp_splf_list_act.htm?lang=en

Navigator

Output Queues, Files on Queue

Easily find the output queues with the largest number of spooled files
(*this is not new*)



	Output Queue	Library	Status	Files on Queue	Description
...	No filter applied				
	Qprint	Qgpl	Released	129	Default Pri
	Qezdebug	Qusr sys	Released	40	Cleanup o
	Qtppoutq	Qusr sys	Released	0	TCP/IP PP
	Qsrvmon	Qservice	Released	0	
	Qsprcloutq	Qrd	Released	0	System cr
	Qscaparoq	Qsc8851312	Released	0	
	Qscaparoq	Qsc7479297	Released	0	
	Qscaparoq	Qsc6479296	Released	0	

1 - 17 of 17 items 5 | 10 | 25 | 50 | 100 | All 1



A few things on PTFs

Automatic Download of Group PTFs

7.3

- Your IBM i must be configured to connect to IBM to download PTFs
- Change your service attributes - [Change Service Attributes](#) (CHGSRVA) command
 - Order PTFs Automatically (ORDPTFAUTO)
 - *YES / *NO parameter - default is *NO
 - Tells the system that you do (or do not) want to order PTFs automatically
 - You must change this parameter to have the system automatically download PTFs for you
 - Order PTF Groups (ORDPTFGRP)
 - This parameter allows you to specify which PTF groups you want to automatically download
 - *CUMPKG
 - *ALLGRP
 - *HIPERGRP
 - *DB2GRP
 - *BRSGRP
 - *JVAGRP
 - *HTTPGRP
 - *PFRGRP

<http://www.ibmssystemsmag.com/Blogs/i-Can/December-2016/Automatic-Download-of-Group-PTFs-with-IBM-i-7-3/>

Automatic Download of Group PTFs

7.3

- A scheduled job will be added
 - The first of the month by default

```
Display Job Schedule Entry Details
Job:  QORDPTFGRP      Entry number:  000001      Status:  SCD      System:  DMI2
Last attempted submission:
  Status . . . . . :  Job successfully submitted.
  Date . . . . . :  05/01/17
  Time . . . . . :  00:01:00
Last successful submission:
  Job . . . . . :  QORDPTFGRP
  User . . . . . :  QSECOFR
  Number . . . . . :  026172
  Date . . . . . :  05/01/17
  Time . . . . . :  00:01:00
Schedule date . . . . . :  *MONTHSTR
Schedule time . . . . . :  00:01:00
Frequency . . . . . :  *MONTHLY
Recovery action . . . . . :  *SBMRLS
More...

Press Enter to continue.

F3=Exit  F6=Display last successful submission  F12=Cancel
```

Automatic Download of Group PTFs

7.3

- PTFs are downloaded in save files and stored on your system, ready to be installed
- PWRDWNSYS command has a new parameter:
 - Install PTF Device (INSPTFDEV)
 - Allows PTFs to be installed before the system is powered down

Navigator for i

Support for PTFs

[-] Configuration and Service

- System Values
- Time Management
- Disk Units
- Disk Pools
- Partitions
- Create Partition

[+] Software

[-] Program Temporary Fixes

- [+] PTFs
- [+] PTF Groups

[-] Program Temporary Fixes

[-] PTFs

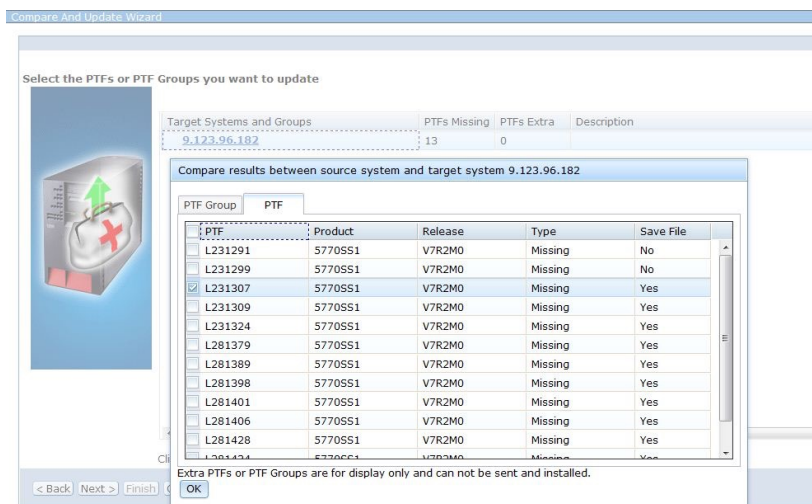
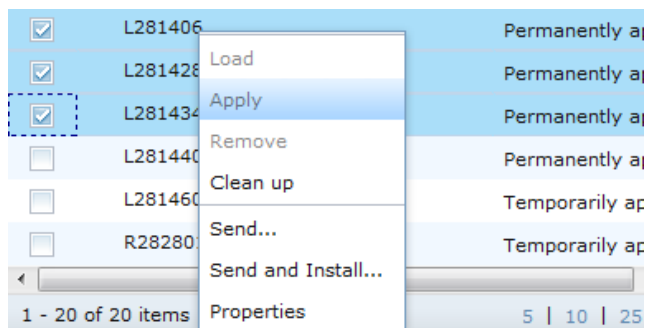
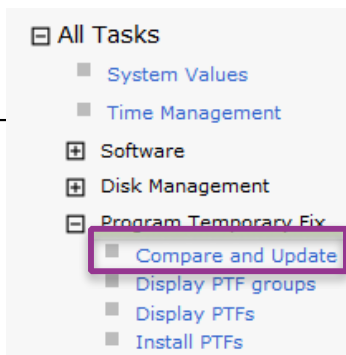
- [+] 5770999
- [+] 5770SS1
- [+] 5700NT1
- [+] 5722IP1
- [+] 5724C04
- [+] 5724J23
- [+] 5733ARE
- [+] 5733B45
- [+] 5733CY3
- [+] 5733LD9
- [+] 5733LEI
- [+] 5733OMF
- [+] 5733QU2

[-] 5770SS1

- All options
- (Option Base, V7R2M0) IBM i
- (Option 1, V7R2M0) Extended Base Support
- (Option 2, V7R2M0) Online Information
- (Option 3, V7R2M0) Extended Base Directory Support
- (Option 5, V7R2M0) System/36 Environment
- (Option 6, V7R2M0) System/38 Environment
- (Option 7, V7R2M0) Example Tools Library
- (Option 8, V7R2M0) AFP Compatibility Fonts
- (Option 9, V7R2M0) PRV CL Compiler Support
- (Option 12, V7R2M0) Host Servers

PTF Compare and Update

- Compare PTFs and PTF groups on different systems
- Send and install PTFs/PTF groups to multiple systems
- Multi-system PTF comparison



Application Runtime Expert - Managing PTFs

- The IBM i Application Runtime Expert can be used to manage PTFs

ARE can help you answer the following questions:

- Are all of my systems at the same fix level?
- How can I keep my systems up-to-date?
- What are the latest PTF fix levels released by IBM?
- Can I be notified when something on my system changes

Display PTF Group Command

- Displays or prints information for a selected PTF group

7.1

SI54641

7.2

SI53559

7.3

Included

Display PTF Group (DSPPTFGRP)

Type choices, press Enter.

PTF group		
PTF group level	*LATEST	1-99999, *INSTALLED, *LATEST
Type of information	*PTF	*PTF, *NOTINSPTF...
Output	*	*, *PRINT

Additional PTF Enhancements

- PTF install history is now tracked in the Security Audit Log
 - Will track both PTF operations and/or PTF object changes.
- More immediate Apply options for PTFs

IBM has added support for *conditional immediate* PTFs.

Conditional immediate PTFs allow an

immediate apply PTF

which supersedes a delayed PTF

to be immediately applied

if the superseded PTF has already been applied to the system

<http://www.ibmssystemsmag.com/Blogs/i-Can/May-2015/IBM-i-Conditional-Immediate-PTFs/>

Additional PTF Enhancements - *DLTPTF*

7.1 - [SI54641](#)

7.2 - [SI54590](#)

- Delete PTF (DLTPTF) command - [PTF selection criteria](#)
 - Additional options for the PTF parameter
 - *ALL
All PTFs for the specified product are deleted
 - *PRMAPY
All PTFs for the specified product that are permanently applied are deleted. This includes PTFs with a status of Superseded where the "Superseded-by" PTF is permanently applied.
 - *SAVONLY
All PTFs for the specified product with a status of Save file only are deleted.
 - *ONORDONLY
All PTFs for the specified product with a status of On order only are deleted.
 - *COVERONLY
All PTFs for the specified product with a status of Cover letter only are deleted.
 - character-value
Specify the identification number of the PTF to be deleted. A maximum of 50 PTF numbers can be specified.

Additional PTF Enhancements - DSPPTFAPYI

7.1 - SI52034
or SF99708 level 8

- DSP PTF Apply Information (DSPPTFAPYI)
 - Have the PTFs or PTF Groups you downloaded be applied without an IPL?
 - The possible values for "Apply immediate" are:
 - *YES
The selected PTFs can be installed immediately.
 - *NO
An IPL is required to install the selected PTFs.
 - *PRECND
The selected PTFs can be installed immediately if all of the preconditions are met.

DSPPTFAPYI is also on the Navigator GUI

Configuration and Service

- System Values
- Time Management
- Disk Units
- Disk Pools
- Partitions
- Create Partition

Software

Program Temporary Fixes

All Tasks

- Time Management

Disk Management

Program Temporary Fix

- Compare and Update
- **Display PTF apply information**
- Display PTF groups
- Display PTFs
- Install PTFs

Display PTF Apply Information

Product: *ALL

PTF number: *NOTAPY

PTF Group: *NONE *LATEST

Release: *ONLY

Display Cancel Reset

Display PTF Apply Information		
Product:	*ALL	
PTF number:	*PTFGRP	
PTF Group:	SF99722	*LATEST
Release:	*ONLY	
Display	Cancel	Reset

ID	Text	Help
CPF3610	All selected PTFs are already applied.	Cause : The Display PTF Apply Information (DSPPTFAPYI) command is not valid when all selected PTFs are temporarily or permanently applied or when no PTF activity exists for the selected product. If a PTF has a status of superseded, a superseding PTF is either temporarily or permanently applied. If a PTF group was specified, all PTFs for installed products are applied. PTF apply information was requested for Product *ALL, Release *ONLY, PTF numbers to select *PTFGRP, PTF group SF99722, and PTF group level 7 (level -1 indicates the latest level of the PTF group). Recovery . . . : Specify a PTF or PTF group that is not installed, and then try the request again.

IBM i Services

Use SQL to Obtain PTF Information

PTF Services				
QSYS2.GROUP_PTF_INFO	View	Base	Base	SF99701 Level 6
QSYS2.PTF_INFO	View	Enhanced in Base	Base	SF99701 Level 23
SYSTOOLS.GROUP_PTF_CURRENCY	View	Base Enhanced: SF99703 Level 3	SF99702 Level 3 Enhanced: SF99702 Level 14	SF99701 Level 32 Enhanced: SF99701 Level 41
SYSTOOLS.GROUP_PTF_DETAILS	View	Base Enhanced: SF99703 Level 3	SF99702 Level 9 Enhanced: SF99702 Level 14	SF99701 Level 38

SQL to obtain PTF Group information

- You can use SQL to interrogate PTF Group information
 - Use SQL instead of the WRKPTFGRP command

-- Query the PTF Groups that are about to change

```
SELECT * FROM QSYS2.GROUP_PTF_INFO
WHERE PTF_GROUP_STATUS = 'APPLY AT NEXT IPL'
ORDER BY PTF_GROUP_LEVEL DESC
```

-- What PTF Cumulative level is installed?

```
SELECT MAX(PTF_GROUP_LEVEL) FROM
QSYS2.GROUP_PTF_INFO
WHERE PTF_GROUP_NAME IN ('SF99610', 'SF99710')
AND PTF_GROUP_STATUS = 'INSTALLED'
```

DB2 PTF Group enablement was shipped in:

IBM i 7.2 - base release

IBM i 7.1 - DB2 PTF Group SF99701 Level 6

SQL to obtain PTF information

- You can use SQL to interrogate PTF information
 - Use SQL instead of the DSPPTF command

-- I'm about to IPL, will any PTFs change?

```
SELECT PTF_IDENTIFIER, PTF_IPL_ACTION, A.*  
FROM QSYS2.PTF_INFO A  
WHERE PTF_IPL_ACTION <> 'NONE'
```

-- Which PTFs are loaded, but not applied?

```
SELECT PTF_IDENTIFIER, PTF_PRODUCT_DESCRIPTION, A.*  
FROM QSYS2.PTF_INFO A  
WHERE PTF_LOADED_STATUS = 'LOADED'  
ORDER BY PTF_PRODUCT_ID;
```

DB2 PTF Group enablement was shipped in:

IBM i 7.2 - base release

IBM i 7.1 - DB2 PTF Group SF99701 Level 23

SQL to to determine Installed PTF Group Currency

- You can use SQL to determine if there are more recent PTF groups available

```
SELECT * FROM SYSTOOLS.GROUP_PTF_CURRENCY  
ORDER BY ptf_group_level_available - ptf_group_level_installed DESC
```

DB2 PTF Group currency was shipped in:

IBM i 7.2 - DB2 PTF Group SF99702 Level 3

IBM i 7.1 - DB2 PTF Group SF99701 Level 32

SQL Output

Page 1 of 1 Jump to page: 1 Total: 14

PTF_GROUP_CURRENCY	PTF_GROUP_ID	PTF_GROUP_TITLE	PTF_GROUP_LEVEL_INSTALLED	PTF_GROUP_LEVEL_AVAILABLE	PTF_GROUP_LAST_UPDATED_BY_IBM	PTF_GROUP_RELEASE	PTF_GROUP_STATUS_ON_SYSTEM
INSTALLED LEVEL IS CURRENT	SF99720	Current Cumulative PTF Media Documentation	14276	14276	11/11/2014	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99719	720 Group Hiper	29	29	04/07/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99747	720 DB2 Web Query for i V2.1.0	11	11	03/25/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99715	720 Backup Recovery Solutions	10	10	03/20/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99713	720 IBM HTTP Server for i	6	6	04/01/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99775	720 Hardware and Related PTFs	6	6	04/01/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99481	720 WebSphere App Server V8.5	4	4	03/13/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99702	720 DB2 for IBM i	4	4	02/15/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99716	720 Java	4	4	12/10/2014	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99766	720 Print PTFs	3	3	01/07/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99714	720 Performance Tools	2	2	12/08/2014	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99717	720 Technology Refresh	1	1	11/24/2014	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99767	720 720 TCP/IP PTF	1	1	03/02/2015	R720	INSTALLED
INSTALLED LEVEL IS CURRENT	SF99776	720 High Availability for IBM i	1	1	11/11/2014	R720	INSTALLED

Page 1 of 1 Jump to page: 1 Total: 14




www.ibm.com/power/i

References



References

- [developerWorks](#)
 - [Technology Updates](#)
- [IBM i 7.2 Documentation - Knowledge Center](#)
- [IBM i 7.3 Documentation - Knowledge Center](#)
- [i Can](#) blog
- [Best Practices for Job and Output Management](#) Redbook
- [IBM i 7.1 Technical Overview with Technology Refreshes](#) Redbook
- [IBM i 7.2 Technical Overview with Technology Refreshes](#) Redbook
- [Application Runtime Expert](#)



IBM Knowledge Center



Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries. Altivec is a trademark of Freescale Semiconductor, Inc.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

Revised December 2, 2010