

iSeries. mySeries.

V5R3 Performance Tools Update

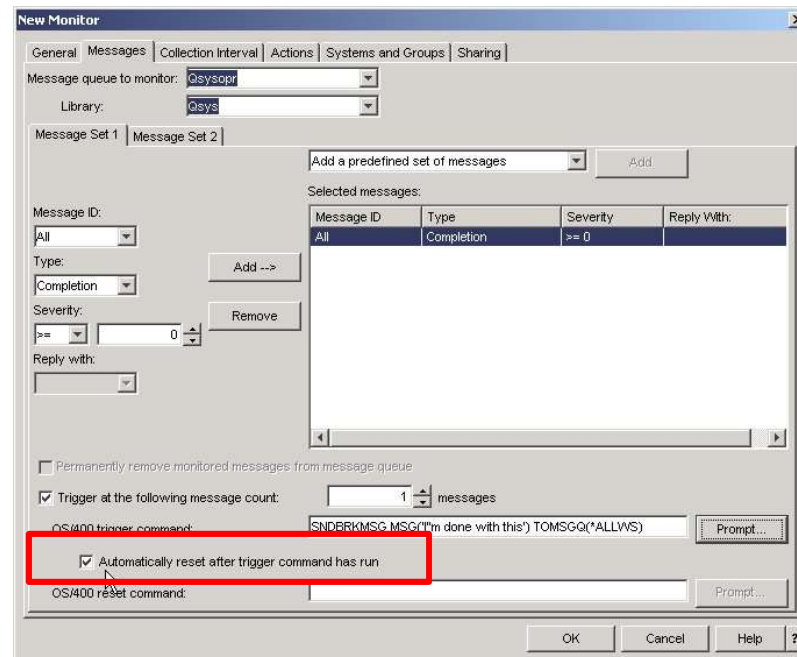

Common Belgium

June 2, 2004

iSeries Navigator monitors

iSeries Navigator monitors update

- Management Central monitors assume local system time is set correctly
 - New system value QTIMZON
 - Set system time for a specific time zone
- Automatic reset of thresholds after running trigger command
 - File and message monitors



Notes: iSeries Navigator Monitors

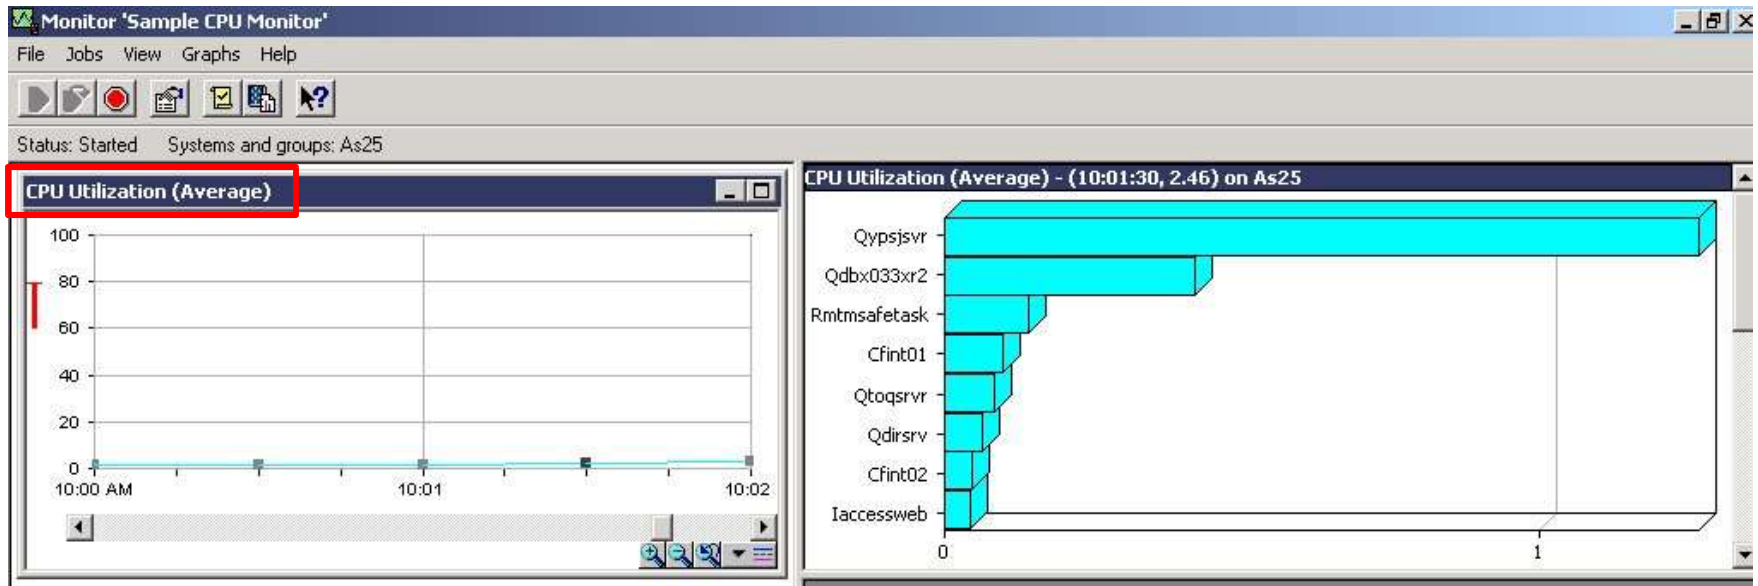
Use new V5R3 QTIMZON system value to make sure Management Central monitors are processing correct local time.

The Management Central monitors assume that the local system time on the endpoint system(s) is configured correctly. Otherwise several metrics will never trigger. In order to ensure that the Management Central monitors are processing the correct local time, you can use the new time zone system value (QTIMZON) to ensure that the current time zone on your endpoint system(s) matches your local time zone.

Starting with V5R3, you can define thresholds for file monitors and message monitors so they automatically reset when your trigger command is run. You define a threshold and specify a command to be run when the threshold is triggered, so that they Automatically reset after the trigger command has run.

CPU Utilization (Average and Interactive Feature)

- CPU Utilization (Average) and CPU Utilization (Interactive Feature) metrics
 - Now use CPU utilization values instead of CPU count
- CPU Utilization (Average) metric can go above 100%
 - for partitions that support uncapped processors
 - associated range of the graph will also accommodate percentages >100



Notes: CPU Utilization - Average and Interactive Feature

The CPU Utilization (Average) metric and CPU Utilization (Interactive Feature) metric changed their calculation methods to use the CPU utilization values instead of using the current calculation based on CPU count.

See Hardware POWER5 Technology section for additional information about the CPU Performance Metrics.

Additionally, the CPU Utilization (Average) metric can now go above 100% for partitions that support uncapped processors, so the associated range of the graph will also accommodate percentages greater than 100.

New fields now added to system monitor metrics

- Disk Arm Utilization (Average)
- Disk Arm Utilization (Maximum)
- Disk Storage (Average)
- Disk Storage (Maximum)
- Provides info about multi path disk unit
 - Multiple redundant paths from the system to the disk unit
- Provides info about remotely mirrored IASP

Notes: Disk units – New Fields

In V5R3, new fields have been added to the system monitor metrics: Disk Arm Utilization (Average), Disk Arm Utilization (Maximum), Disk Storage (Average), and Disk Storage (Maximum).

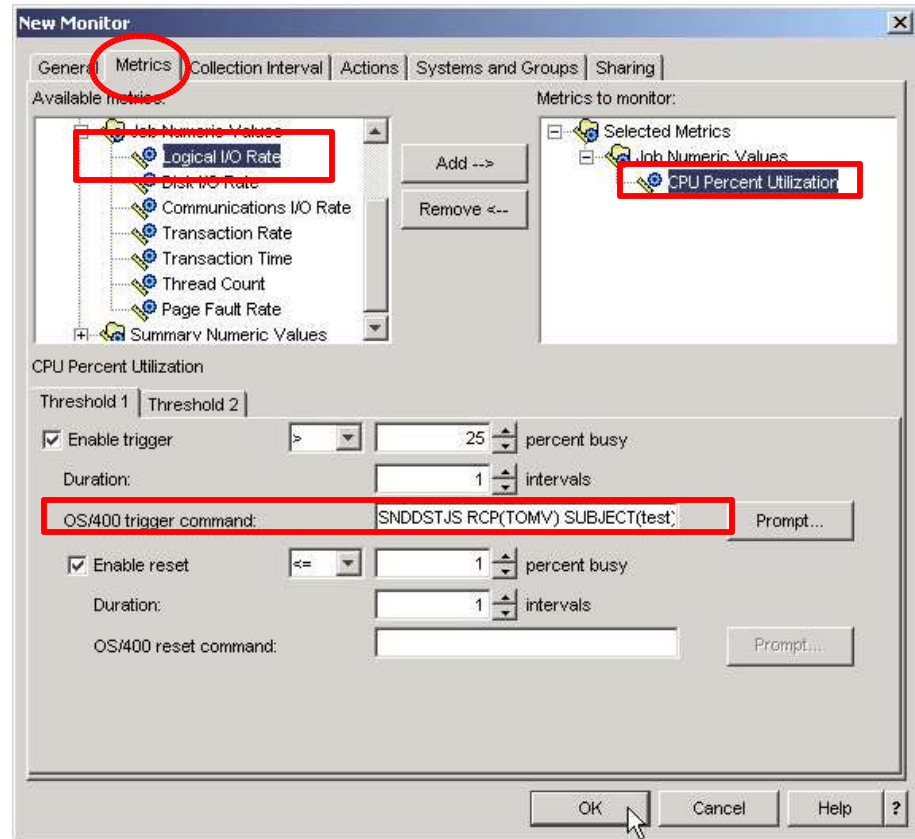
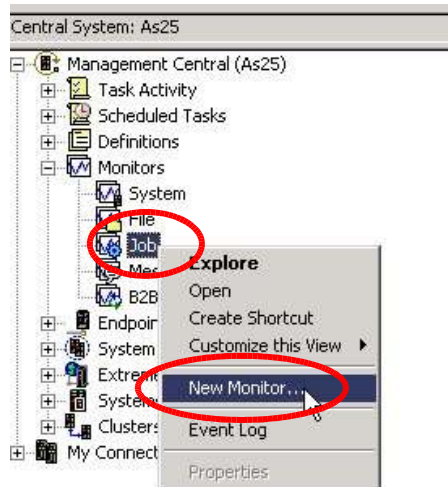
These new fields provide information about a multipath disk unit, which is a unit that has multiple redundant paths from the system to the disk unit.

Each path has a unique resource name.

- Multipath unit: indicates that the resource, represents a multipath disk unit.
- Initial path of multipath unit: indicates that the resource represents the initial path of a multipath disk unit.
- Production copy of a remotely mirrored independent ASP: indicates that the disk unit is in a production copy of a remotely mirrored independent ASP.
- Mirror copy of a remotely mirrored independent ASP: indicates that the disk unit is in a mirror copy of a remotely mirrored independent ASP.

Job Notification – Job monitor CPU % Utilization

- Advanced Job Scheduler for iSeries (5722-JS1)
 - Uses notification function in combination with job monitor
 - CPU Percent Utilization metric
 - SNDDSTJS command



Notes: Job Monitor CPU % Utilization

In the example on the previous foil we describe the possibility to define a job monitor within iSeries Navigator. You can use the notify function of the Advanced Job Scheduler (Send Distribution using JS - SNDDSTJS command) if the % CPU utilization of a job reaches a specified threshold.

In our example, we will show sending a notifying message to a user as soon as one specific job is using more than 25% of CPU.

Collection Services

CL Commands

- STRPFRCOL
 - Start a performance collection
 - Parameters
 - Default collection profile (COLPRF)
 - Cycle collection (CYCCOL)
 - Starts the Collection Services server job (QYPSPFRCOL)
- ENDPFRCOL
 - End a performance collection
 - Force end (FRCCOLEND) parameter
- CFGPFRCOL
 - Change the Collection Services properties
- CHKPFRCOL
 - Determine the current status of the Collection Services server job

Notes: CL commands

The Start Performance Collection (STRPFCOL) command starts the system-level collection of performance data by Collection Services. The properties of the system-level collection are controlled by the Collection Services configuration. The data included in the system-level collection is determined by the value specified on the Collection profile parameter. The performance data collection is conducted by the Collection Services server job (QYPSPFCOL). There are two available parameters:

- Default collection profile (COLPRF): this specifies the collection profile and determines which categories will be included in the system-level collection.
- Cycle collection (CYCCOL): this specifies whether the collection should be cycled. Cycling the collection will cause data to be collected in a new management collection object (*MGTCOL).

The End Performance Collection (ENDPFCOL) command stops the system-level collection. If there are no other client applications using Collection Services, the Collection Services server job (QYPSPFCOL) will also end. If client applications are using the Collection Services, the server job will continue to run unless you also specify the FRCCOLEND (Force end) parameter.

The Configure Performance Collection (CFGPFRCOL) command changes the Collection Services properties. These properties define certain collection attributes and determine how Collection Services will manage data collection. Some attributes changed with this command will take effect immediately, while others will take effect the next time a collection object is created. See the parameter help to determine how each attribute will be handled.

The Check Performance Collection (CHKPFRCOL) command provides a method for determining the current status of the Collection Services server job (QYPSPFCOL). If the server job is not active, the command sends escape message CPF0AA5. If the server job is active, the command sends information message CPI0A16. This information message provides the name and library of the current management collection object and the current collection profile.

New data and methodology for CPU utilization

- Within the collection interval V5R3 Collection Services reports
 - Total CPU that is consumed
 - Total CPU that is available to the partition
- In V5R3 dynamic LPAR environments with shared processors:
 - HVLPTASK concept no longer valid
 - CPU scaling to whole virtual processors does not exist
 - Collection Services no longer cycles the collection when configuration changes
- CPU utilization formula

$$\text{CPU utilization} = \frac{\text{CPU Consumed}}{\text{Available capacity}}$$

Notes: New data and methodology for CPU utilization

In V5R3, Collection Services reports the total CPU that is consumed and the total CPU that is available to the partition within the interval. The concept of HVLPTASK and CPU scaling to whole virtual processors in shared processor environments does not exist in V5R3 and Collection Services will no longer cycle the collection when the configuration changes.

Collection Services now reports the total processor time that is consumed by the partition, along with the amount of processor time that was available to be consumed within the partition, regardless of the number of virtual processors that are configured, the partition units that are configured, or how they changed during the interval.

To calculate utilization, users of this data will need to divide the reported CPU consumed, by the available capacity.

In V5R3, the Convert Performance Data (CVTPFRDTA) command performs normally. However, the data in the converted files is changed to be consistent with the unscaled system CPU data (QAPMSYSCPU database file). The results should be the same as if the data were collected on a V5R3 system, but the data is different than the values that existed in the files in a prior release.

New data and methodology for CPU utilization

- Convert Performance Data (CVTPFRDTA) command performs normally
 - Data in converted files changed to be consistent with unscaled system CPU data
 - QAPMSYSCPU DB file (reports utilization for all processing units)
 - The results should be the same as if the data were collected on a V5R3 system
 - Data is different than the values that existed in the files in a prior release
- Existing tools that calculate CPU utilization do not show the correct results
 - For shared processor partitions
 - For partitions that have had configuration changes during data collection
 - This includes those tools using
 - The performance database
 - The QPMLPFRD API
- If you copy a V5R3 *MGTCOL object to a prior release and then generate the database files on that system
 - The reported CPU data remains unscaled (shared processor environments)
 - Total system CPU reported is incorrect

Notes: New data and methodology for CPU utilization

In V5R3, the Convert Performance Data (CVTPFRDTA) command performs normally. However, the data in the converted files is changed to be consistent with the unscaled system CPU data (QAPMSYSCPU database file). The results should be the same as if the data were collected on a V5R3 system, but the data is different than the values that existed in the files in a prior release.

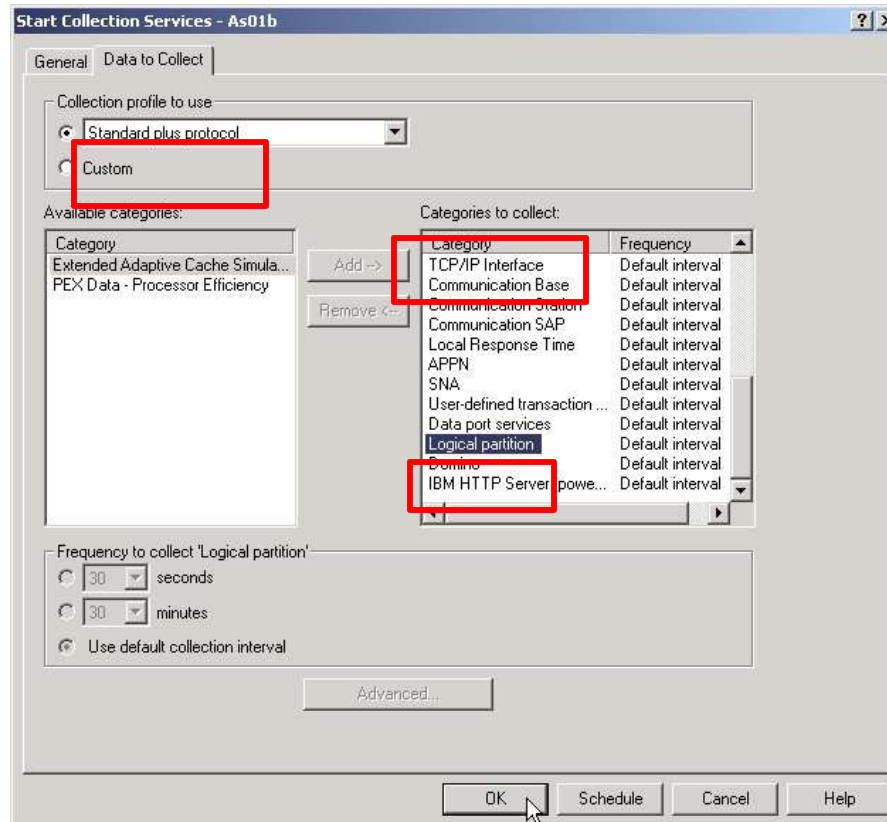
The existing and unchanged tools that calculate CPU utilization do not show the correct results for shared processor partitions or partitions that have had configuration changes during data collection. This includes those tools that use the performance database as well as those that use the QPMLPFRD API.

If you copy a V5R3 management collection object (*MGTCOL) to a prior release and generate the database files on that system, you should be aware of the following:

- The reported CPU data remains unscaled (shared processor environments). This means that the total system CPU that is reported by the tools using virtual processors (including Performance Tools) is not correct.
- A management collection object (*MGTCOL) that spans configuration changes will result in an inaccurate calculation of the percentage of CPU during those intervals after the change occurred.

Cross-Partition performance data

- Collection Services can collect performance data across partitions
 - New Logical Partition category
- IBM Director for Multiplatform product needs to be running



Notes: Cross-Partition performance data

For V5R3, Collection Services can collect data from all partitions on an iSeries system if IBM Director for Multiplatform is installed and running on the partition that is running Collection Services.

Collection Services collects the data from each partition, and PM iSeries will summarize the data.

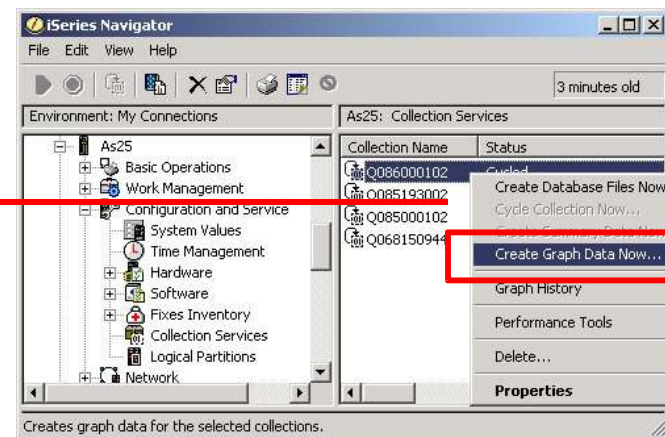
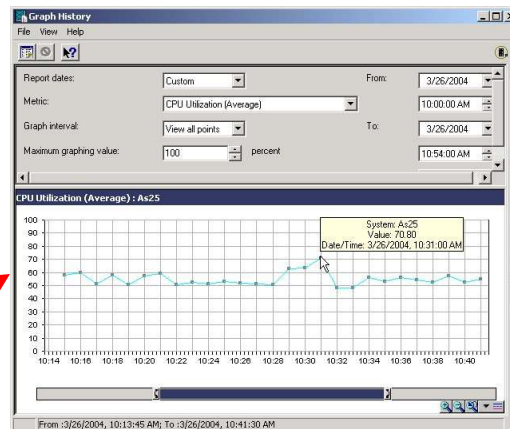
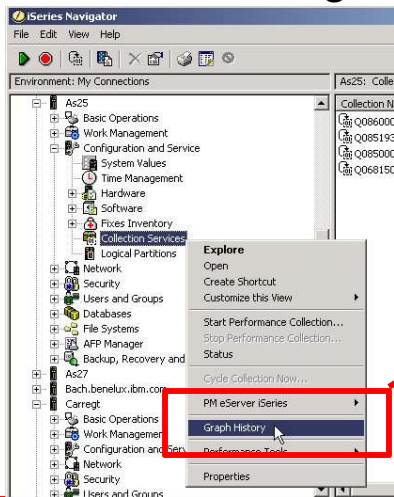
You can specify to collect logical partition data from iSeries Navigator.

A new category, Logical partition, is added to the Available categories/Categories to collect field on the Collection Services dialog.

This category is included in the *STANDARD, *STANDARDP, and *ENHCPCPLN profiles. You can also include the category in a custom profile.

Graph History

- Obtaining Graph Data prior to V5R3
 - Graph data was returned if graph data was available for a specific time range
 - If no graph data was available but raw data was available
 - Raw data was dynamically converted to graph data and returned
If retention period was great enough
 - If no graph data was still available
 - Summarized data was returned
- V5R3 eliminates dynamic conversion of raw data to graph data
 - Use Create graph data now is the option



Notes: Graph History

Prior to V5R3 you could obtain graph data in three ways.

- If graph data was available for a specific time range, it was returned.
- If no graph data was available, but the raw data was available and the retention period was great enough for the graph data to be obtained from the raw data, then the raw data was dynamically converted to graph data and returned.
- If no graph data was still available, the summarized data was returned.

Beginning in V5R3, the dynamic conversion of raw data to graph data was eliminated.

You can now use the Create graph data now option to create your graph data from the Collection name from the Collection Services in the Configuration and Service container.

Performance Tools Licensed Program

Batch and interactive CPU utilization

- Process for accounting CPU utilization
 - Non-interactive (batch) and interactive
 - based upon new bucket in QAPMJSUM
- iSeries Access jobs can run either as batch or as interactive
 - In previous releases included in CA4 bucket
 - All charged to the interactive CPU utilization
 - New category added to system report
 - iSeries Access – Batch (under non-interactive workload subsection)
- DDM server jobs moved from interactive workload subsection
 - To non-interactive workload subsection

Notes: Batch and interactive CPU utilization

The process for accounting non-interactive (batch) and interactive CPU utilization is based upon a new bucket in QAPMJSUM that is provided by Collection Services.

iSeries Access jobs can run either as batch or as interactive, but in previous releases all of the jobs were included in the CA4 bucket, which is charged to the Interactive CPU utilization column. A new category was added to the System Report, iSeries Access - Bch under the Non-interactive workload subsection.

Additionally, the DDM server jobs were moved from the Interactive workload subsection to the Non-interactive workload subsection because they do not represent interactive workload.

Batch and interactive CPU utilization - continued

System Report

Workload

Non-Interactive Workload

Job Type	Number Of Jobs	Logical DB I/O Count	Printer Lines	Printer Pages	Communications I/O Count
PassThru	5.9	105	2	1.0	.0
Batch	19.4	3	1	668.7	11.6
DDM Server	.0	0	0	.0	.0
iSeries Access-Bch	8.2	0	2	2.6	.0
COLLECTION	.0	0	0	4.5	.2

Non-Interactive Resource Utilization Expansion

Job Type	Average Per Second								Logical		
	Physical Disk I/O				Data Base I/O				Read	Write	Other
	Synchronous				Asynchronous						
	DBR	DBW	NDBR	NDBW	DBR	DBW	NDBR	NDBW			
Batch	11.6	1.5	.0	2.5	2.4	650.3	.0	.0	229.2	380.2	229.2
DDM Server	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0

Notes: Batch and interactive CPU utilization - continued

The following is a simple example of a page of the system report, showing the DDM server as well as iSeries Access in the Non-Interactive Resource Utilization Expansion. It also shows the new iSeries Access -Bch under the non-interactive workload subsection.

Component report

- Metrics from QAPMCONF file added to header section
 - Virtual Processors
 - Processor Units
 - Int(eractive) Threshold - percent of total system CPU for interactive work

Component Report

Component Interval Activity

I,B,RUNSQL,DDM FEB 26

Member . . . : Q057171530 Model/Serial . . : 830/10-2257M Main storage . . : 1000.0 MB Started . .
 Library . . : QMPGDATA System name . . :RCHAS01B Version/Release : 5/ 3.0 Stopped . .
 Partition ID : 001 Feature Code . . :23DB-2403-1534 Int Threshold . . : 2.12 %

Virtual Processors: 1 Processor Units : 1.0

Itv	Tns	Rsp	DDM	-CPU Utilization-	CPU	Feat	CPU	Cpb	-----	Disk I/O	-----	High		
End	/Hour	/Tns	I/O	Total	Inter	Batch	Avail	Util	>Thld	Util	Sync	Async	Disk	Unit
17:20	140	.20	0	16.9	2.3	14.5	.0	9.1	0	11.0	5.7	459.7	7	0008

Itv End -- Interval end time (hour and minute)

Tns /Hour -- Number of interactive transactions per hour

Notes: Component report

Virtual Processors (GKEY 13), Processor Units (GKEY PU), and Int Threshold (GKEY IT) metrics fields from the QAPMCONF file were added to the header section of the component report.

- 13 - number of logical processors assigned to the partition.
- PU - processor units allocated to the partition.
- IT - Interactive threshold as a percent of the configured processor units (see GKEY PU). The Interactive Threshold is the percent of the total system CPU available for interactive work. On Enterprise Edition models this is 100 percent.

Component report - continued

- Unit column of the Disk Utilization section
 - Plus (+) sign added next to unit to identify multi path unit
 - H sign added next to unit to identify a high availability parity set
- Dash (-) sign under High disk and Unit columns
 - No performance data available for the requested interval in QAPMDISK
- New column added to Component Interval Activity section
 - Showing uncapped CPU time available for the system

Component Report													
Component Interval Activity													
Member	Q057171530 Model/Serial . . : 830/10-2257M				Main storage					1000.0	MB	Started	
Library	QMPGDATA		System name			RCHAS01B		Version/Release		5/ 3.0	Stopped		
Virtual Processors:		1	Processor Units :		1.0								
						Uncap	Int	Int	DB	-----	Disk I/O	-----	Hig
Itv	Tns	Rsp	DDM	-CPU	Utilization-	CPU	Feat	CPU	Cpb	-----	Per Second	----	-- Uti
End	/Hour	/Tns	I/O	Total	Inter	Batch	Avail	Util	>Thld	Util	Sync	Async	Disk

17:20	140	.20	0	16.9	2.3	14.5	.0	9.1	0	11.0	5.7	459.7	7

Notes: Component report - continued

A plus sign (+) was added next to the Unit column of the Disk Utilization section to identify multipath disk units. A letter H following the unit number indicates that the disk unit is in a high availability parity set.

If no letter displays then the disk unit is a regular parity set.

Dash signs (-) display under the High Disk and Unit columns when performance data is not available for the requested interval in the QAPMDISK file.

A new column was added to the Component Interval Activity section to show the uncapped CPU time available for the system. This data is reported only for performance data that is collected on uncapped partitions.

Component report - continued

- Amount of time interactive usage exceeds configured Interactive Threshold
 - Column added to Component Interval Activity section
 - New field SYIFTE in QAPMSYSTEM file
- In V5R3 Interactive Threshold can change through data collection period
 - By partition reconfiguration
 - Is in SYIFTA field in QAPMSYSTEM file

Component Report													
Component Interval Activity													
Member . . . : Q057171530			Model/Serial . . : 830/10-2257M				Main storage . . : 1000.0 MB			Started			
Library . . . : QMPGDATA			System name . . : RCHAS01B				Version/Release : 5/ 3.0			Stopped			
Virtual Processors: 1			Processor Units : 1.0										
			Uncap Int		Int		DB		----- Disk I/O -----		Hig		
Itv	Tns	Rsp	DDM	-CPU Utilization-	CPU	Feat	CPU	Cpb	-----	Per Second	----	-- Uti	
End	/Hour	/Tns	I/O	Total	Inter	Batch	Avail	Util	>Thld	Util	Sync	Async	Disk

17:20	140	.20	0	16.9	2.3	14.5	.0	9.1	0	11.0	5.7	459.7	7

Notes: Component report - continued

A column was added to the Component Interval Activity section to show the amount of time exceeding the Interactive Threshold, which is the time, in seconds, during which the interactive usage exceeds the configured Interactive Threshold.

The amount of time exceeding the Interactive Threshold is a new field, SYIFTE, in the QAPMSYSTEM file.

The Interactive Threshold is the percent of the total system CPU for interactive work.

With V5R3, the interactive threshold value can change through the data collection period by reconfiguring the partition and is now in the SYIFTA field in the QAPMSYSTEM file.

SYIFTA field definition: Interactive CPU time available. The amount of interactive CPU time that was available for use within the partition. This is the interactive capacity configured for use within the partition (also represented as interactive threshold).

System report

- DDM server job moved to Non-Interactive Workload subsection
- New category for iSeries Access – Batch jobs added
- QAPMCONF metrics added to header section
 - Virtual Processors
 - Processor Units
 - Int(eractive) Threshold - percent of total system CPU for interactive work
- DDM server job statistics are no longer shown in the first part of the Resource Utilization section
- DDM server job statistics are now shown in the Non-Interactive Resource Utilization Expansion subsection

Notes: System report

The V5R3 DDM server jobs were moved from the Interactive Workload subsection to the Non-Interactive Workload subsection.

A new category for iSeries Access - Batch jobs was added to the Non-Interactive Workload section.

Virtual Processors (GKEY 13), Processor Units (GKEY PU), and Int Threshold (GKEY IT) metrics from the QAPMCONF file were added to the header section of the system report.

The DDM server job statistics are no longer shown in the first part of the Resource Utilization section. The first part of the Resource Utilization section continues to show statistics for Interactive jobs like Interactive, System/36, MRT, iSeries Access, and Pass-through.

The DDM server job statistics are no longer shown in the Interactive Resource Utilization Expansion subsection of the Resource Utilization Expansion section. The DDM server job statistics are now shown in the Non-Interactive Resource Utilization Expansion subsection.

The foil with Batch and interactive CPU utilization earlier in this presentation, shows all these changes.

System report - continued

- Unit column of the Disk Utilization section
 - Plus (+) sign added next to unit to identify multi path unit
 - H sign added next to unit to identify a high availability parity set
- Two rows are added to the Workload section
 - To show information about amount of time exceeding the Interactive Threshold
 - Other to show the shared processors pool utilization in terms of a percentage

```

System Report
Workload

Partition ID : 001          Feature Code . : 23DB-2403-1534      Int Threshold . : 2.12 %
Virtual Processors: 1      Processor Units : 1.0

Average

Total CPU Utilization . . . . . : 35.5
Total CPU Utilization (Interactive Feature) . . : 9.2
Time exceeding Int CPU Threshold (in seconds) . : 0
Total CPU Utilization (Database Capability) . . : 24.5
Shared processor pool utilization . . . . . : .0
  
```

Notes: System report - continued

A plus sign (+) was added next to the Unit column of the Disk Utilization section to identify multi path disk units. A letter H following the unit number indicates that the disk unit is in a high availability parity set. If no letter displays then the disk unit is a regular parity set.

Two rows were added to the Workload section. One to show information about the amount of time exceeding the Interactive Threshold and the other to show the shared processors pool utilization in terms of a percentage. The shared processor pool utilization row displays only for performance data that is collected on partitions using a shared processor pool.

Miscellaneous reports

- Unit column of the Disk Utilization section
 - Plus (+) sign added next to unit to identify multi path unit
 - H sign added next to unit to identify a high availability parity set
- Added column to Job Interval Report
 - Interactive and Non-Interactive Job detail sections
 - Shows current user of job collected by Collection Services in QAPMJOBMI

Notes: Miscellaneous reports

A plus sign (+) was added next to the Unit column of the Disk Utilization section to identify multi path disk units. A letter H following the unit number indicates that the disk unit is in a high availability parity set. If no letter displays then the disk unit is a regular parity set.

A column was added to the Interactive Job Detail and Non-Interactive Job Detail sections of the Job Interval Report to show the current user of the job collected by Collection Services in the QAPMJOBMI file.

Performance advisor enhancements

- Changes to existing metrics to give more accurate recommendations
- Add new metrics to analyze and give recommendation to newer functions
- Provide recommendations for
 - Interactive Feature Utilization
 - Point-to-Point Protocol (PPP) activity
 - TCP/IP activity
- Update guidelines for
 - CPU utilization
 - System metrics
 - Disk service time

Notes: Performance advisor enhancements

Performance advisor enhancements include changes to existing metrics to give more accurate recommendations and to add new metrics to analyze and give recommendations related to newer functions.

- Provide recommendations for Interactive Feature Utilization, Point-to-Point Protocol (PPP) activity, and TCP/IP activity.
- Update guidelines for CPU utilization, system metrics, and disk service time.

Analyze Performance Data display – continued

- New column added to the Select Time Intervals to Analyze display to show the Interactive Feature Utilization (Int Feat Util)
- Dash (-) signs under High Dsk and Unit columns
 - No performance data available for requested interval
- Field added to the Display Recommendations display showing
 - Percentage of the processing capacity that is assigned for interactive work
 - Int(eractive) Threshold
- Two fields added to the Display Recommendations showing
 - Virtual Processors
 - Processor Units

Notes: Analyze Performance Data display

A new column was added to the Select Time Intervals to Analyze display to show the Interactive Feature Utilization (Int Feat Util). Dash signs are displayed under High Dsk and Unit columns when performance data is not available for the requested interval in the QAPMDISK file.

A field was added to the Display Recommendations display to show the percent of the processing capacity that is assigned for interactive work (Int Threshold). This information is taken from the IT field of the QAPMCONF file.

Two fields were added to the Display Recommendations display to show the Virtual Processors (GKEY 13) and Processor Units (GKEY PU) metrics from the QAPMCONF file.

Analyze Performance Data display – Example 1

Select Time Intervals to Analyze

Member : Q089000102 Library : QPFRDATA

Type options, press Enter.

1=Select

p	-Transaction-		-CPU Util--			Feat	--High--	Pool	Fault	Excp				
	Date	Time	Count	Rsp	Tot						Int	Bch	Util	Dsk
1	03/29	00:05	0	.00	2	0	2	0	--	----	0	0	00	0
	03/29	00:10	0	.00	1	0	1	0	--	----	0	0	00	0
	03/29	00:15	0	.00	1	0	1	0	--	----	0	0	00	0
	03/29	00:20	0	.00	1	0	1	0	--	----	0	0	00	0
	03/29	00:25	0	.00	1	0	1	0	--	----	0	0	00	0
	03/29	00:30	0	.00	1	0	1	0	--	----	0	0	00	0

F3=Exit F5=Refresh F11=Display histogram F12=Cancel F13=Select all

F14=Deselect all

Notes: Analyze Performance Data display – Example 1

This foil shows a new column that was added to the Select Time Intervals to Analyze display. It shows the Interactive Feature Utilization (Int Feat Util).

Analyze Performance Data display – Example 2

Display Recommendations

System: AS25

Member :	Q089000102	Library :	QPFRDATA
System :	AS25	Version/Release . . :	5/ 3.0
Start date :	03/29/04	Model :	720
Start time :	00:01:02	Serial number . . . :	10-394TM
Partition ID :	000	Feature Code :	208D-2064-1505
QPFRADJ :	2	Int Threshold . . . :	6.57 %
QDYNPTYSCD :	1	Virtual Processors :	4
QDYNPTYADJ :	1	Processor Units . . :	4.00

Type options, press Enter.

5=Display details

Option Recommendations and conclusions

Recommendations

No performance problems found in system data file

Notes: Analyze Performance Data display - Example 2

This foil shows that a field was added to the Display Recommendations display. It shows the percent of the processing capacity that is assigned for interactive work (Int Threshold). This information is taken from the IT field of the QAPMCONF file.

We can also see two more new fields showing the Virtual Processors and Processor Units metrics from the QAPMCONF file.

Display Performance Data display

- New column added on Select Time Intervals to Display screen
 - Show Interactive Feature Utilization
- Dash (-) signs under High Dsk and Unit columns
 - No performance data available for requested interval

Select Time Intervals to Display

Member : Q037121832 Library : COOK

Type options, press Enter.

1=Select

			-Transaction-		-CPU Util--			Int			Pool Fault			Excp
t	Date	Time	Count	Rsp	Tot	Int	Bch	Feat Util	--High-- Dsk	Unit	Mch	Usr	ID	Util
1	02/06	12:25	5	.40	89	0	89	0	13	0008	0	11	05	0
	02/06	12:30	20	2.00	89	10	79	13	11	0008	0	11	05	0
	02/06	12:35	15	4.73	92	21	71	0	13	0008	0	9	05	0
	02/06	12:40	51	5.49	97	65	32	111	5	0008	0	4	05	0

Notes: Display Performance Data display

Select Time Intervals to Display

A new column was added to show the Interactive Feature Utilization (Int Feat Util).

Dash signs (-) display under the High Dsk and Unit columns when performance data is not available for the requested interval in the QAPMDISK file.

Display Performance Data display – continued

- Field added showing the percentage of processing capacity assigned for interactive work
 - Int(eractive) Threshold
- Fields added showing the Virtual Processors and Processor Units metrics
- Row added showing percentage of interactive capacity used by the system
 - Interactive Feature Utilization
- Row added showing the interactive CPU time in seconds over the threshold
 - Time exceeding Int CPU Threshold (in seconds)
- Option 6 (Wait detail) added to Display Jobs display

Notes: Display Performance Data display - continued

A field was added to the Display Performance Data display to show the percent of the processing capacity that is assigned for interactive work (Int Threshold). This information is taken from the IT field of the QAPMCONF file.

Two fields were added to the Display Performance Data display to show the Virtual Processors (GKEY 13) and Processor Units (GKEY PU) metrics from the QAPMCONF file.

A row was added to the Display Performance Data display to show the percentage of interactive capacity used by the system (Interactive Feature Utilization).

A row was added to the Display Performance Data display to show the interactive CPU time in seconds over the threshold (Time exceeding Int CPU Threshold (in seconds)).

On the Display Performance Data display, you can select PF6 (Display all jobs). Option 6 (Wait detail) was added to the Display Jobs display, which shows the wait time statistics, in seconds, for the requested job or task. This information is over the life of the job within the QAPM performance database file members over the the time periods that can be selected when using Display Performance Data.

The next 3 slides of Display Performance information, includes this new information:

- **Display Performance Data display Example 1**
- **Display Performance Data display – display jobs Example 1**
- **Display Performance Data display – job wait details**

Display Performance Data display Example 1

```
Display Performance Data

Member . . . . . Q089030704          F4 for list

Library . . . . . V5R3MPATH

Elapsed time . . . : 00:05:00      Version . . . . . : 5
System . . . . . : CARREGT        Release . . . . . : 3.0
Start date . . . . : 03/29/04     Model . . . . . : 825
Start time . . . . : 03:07:04     Serial number . . : 10-5M0FM
Partition ID . . . : 000          Feature Code . . . : 7415-2472-7415
QPFRAJ . . . . . : 0              Int Threshold . . : 10.00 %
QDYNPTYSCD . . . . : 1           Virtual Processors : 2
QDYNPTYADJ . . . . : 1           Processor Units . . : 2.00

CPU utilization (interactive) . . . . . : .00
CPU utilization (other) . . . . . : 99.76
Interactive Feature Utilization . . . . . : .00

Time exceeding Int CPU Threshold (in seconds) . . : 0
Job count . . . . . : 17
Transaction count . . . . . : 15

More...

F3=Exit      F4=Prompt  F5=Refresh  F6=Display all jobs  F10=Command entry
F12=Cancel   F24=More keys
```

Display Performance Data display – display jobs Example 1

```

                                Display Jobs

Elapsed time . . . : 00:05:00          Member . . . . . : Q089030704
                                Library . . . . . : V5R3MPATH

Type options, press Enter.

5=Display job detail  6=Wait detail

                                Job      CPU      Tns      Avg      Disk
Option Job          User      Number  Type    Util    Count   Rsp     I/O
-----
6 JOB0302    AS0302    005391  BCH     37.49    0        .0     8131
JOB0303    AS0303    005392  BCH     37.48    0        .0     7694
JOB0301    AS0301    005390  BCH     24.28    0        .0    18093
CFINT01                    LIC      .21      0        .0        0
QYPSJSVR   QYPSJSVR  005366  BCH     .07      0        .0        4
QPADEV0002 JCOOK     005395  PTH     .00     15        .1       73
QINTER     QSYS      005167  SBS     .00      0        .0        5

                                                More...

F3=Exit          F12=Cancel  F15=Sort by job  F16=Sort by job type
F19=Sort by CPU  F24=More keys
    
```

Display Performance Data display – job wait details

Wait Detail

Job : Member : Q089030704

 User : Library : V5R3MPATH

 Number :

Wait Category	Wait Time
------------------	--------------

(WAIT STATE information was not available at the time of this screen capture)

Display performance data: Job Wait classifications

Wait identifier	Description as of May 2004
0 in SysMon	Task run time (switch in to switch out) This is dispatched to run in full processor mode partition. When a partition is running with partial or shared processors the values may not reflect actual usage of a CPU. That is actual could be less in this "non-standard" partition configuration
1 in SysMon	Ready to Run (RTR) time. (This is the time from when a task switches from blocked to waiting for a processor until the time it switches in.) CPU queuing
2 in SysMon	These include conditions that the internal developers thinks should rarely occur. IBM service or development lab must get involved if this occurs.
3 in SysMon	Default (A wait of some kind that could not be classified into one of the other defined wait buckets.)
4 in SysMon	DASD1 waiting for synchronous DASD ops to complete (read or write, DB or NDB) waiting for asynchronous ops to complete before advancing (IO Pending)
5 in SysMon	DASD2 - compression, mirroring, reassign, etc waits
6 in SysMon	Seize waits
7 in SysMon	Record lock waits
8 in SysMon	Object lock waits
9 in SysMon	Gate (UNIX-like) waits
10 in SysMon	Java waits
11 in SysMon	Socket waits
12 in SysMon	Journal waits
13 in SysMon	Mutex, Semaphore waitsr
14 in SysMon	waiting for work request, key/think time, etc. (note - no block points are currently defined)

We are providing this information to help you understand the wait state information displayed in the job details introduced in V5R2 , No further information is being provided as it would be used by service personnel as part of problem determination.

Notes: Job Wait classifications

At the time this presentation screen was captured the Display Performance Data – job wait conditions was not yet operational.

However, the V5R2 System Monitor support, showing job details including the approximately 15 job wait classifications was operational and an internal lab design document that defines the wait states/classifications is summarized here.

The following slide shows a System Monitor graphic with a portion of this wait state information shown.

Display Performance Data- Wait states

Status: Started Systems and groups: As01b

CPU Utilization (Average)

Time	Utilization (%)
10:12	75
10:13	75
10:14	65
10:15	45
10:16	45

CPU Utilization (Average) - (10:12:00, 73.02) on...

Process	Utilization (%)
Qpadev0001	~23
Qpadev0002	~22
Job0303	~18
Job0301	~8
Job0302	~6
Cfint01	~4
Qcpmgtsvr	~3
Qypsjsvr	~2

CPU Utilization (Interactive Feature)

Time	Utilization (%)
10:12	200
10:13	180
10:14	125
10:15	0
10:16	0

Qpadev0001 (23.17)

Property	Value
Stream file bytes written	0
Counter Set 0 wait count	2738
Counter Set 0 wait time	36464
Counter Set 0 description	Reserved
Counter Set 1 wait count	0
Counter Set 1 wait time	0
Counter Set 1 description	CPU queuing
Counter Set 2 wait count	0
Counter Set 2 wait time	0
Counter Set 2 description	Reserved
Counter Set 3 wait count	1
Counter Set 3 wait time	5
Counter Set 3 description	Other waits
Counter Set 4 wait count	0

Perform menu

- Option 2 (Collect performance data) from the Perform menu
 - Take advantage of the new performance collection CL commands
 - Previously used the collector APIs
 - Option1 – Start performance collection
 - STRPFRCOL command
 - Option2 – Configure performance collection
 - CFGPFRCOL command
 - Option 3 – End performance collection
 - ENDPFRCOL command
- New column added to Select Time Intervals display showing
 - The Interactive Feature Utilization – Int Feat Util
- Dash (-) signs under High Dsk and Unit columns
 - No performance data available for requested interval

Notes: Perform menu

The Collection Services options from the Perform menu, Option 2 (Collect performance data), now takes advantage of the new performance collection CL commands.

Previously, these options used the collector APIs.

These options now use the following CL commands:

- Option 1 (Start Performance Collection): Start Performance Collection - STRPFRCOL
- Option 2 (Configure Performance Collection): Configure Performance Collection - CFGPFRCOL
- Option 3 (End Performance Collection): End Performance Collection - ENDPFRCOL

A new column was added to the Select Time Intervals display to show the Interactive Feature Utilization (Int Feat Util). Dash signs (-) are displayed under the High Disk and Unit columns when performance data is not available for the requested interval in the QAPMDISK file.

Perform menu

PERFORM

IBM Performance Tools for iSeries

System: AS25

Select one of the following:

1. Select type of status
2. Collect performance data
3. Print performance report

5. Performance utilities
6. Configure and manage tools
7. Display performance data
8. System activity
9. Performance graphics
10. Advisor

70. Related commands

Selection or command

===> 2

Notes: Perform menu

This foil shows the option 2 on the Perform menu – Collect performance data.

Perform menu - continued

Collect Performance Data

AS25

03/29/04 10:07:48

Collection Services status:

Status : Stopped

Select one of the following:

- 1. Start Performance Collection
- 2. Configure Performance Collection
- ~~3. End Performance Collection~~

Selection or command

===>

Notes: Perform menu - continued

This foil shows the Collect performance data screen.

There are 3 options available:

- Option 1 (Start Performance Collection) - STRPFRCOL
- Option 2 (Configure Performance Collection) - CFGPFRCOL
- Option 3 (End Performance Collection) - ENDPFRCOL

These options no longer use the collector APIs, but use the commands as described above.

Configure Collection Services command

Configure Perf Collection (CFGPFRCOL)

Type choices, press Enter.

Default interval	05.00	*SAME, .25, .50, 1.0, 5.0...
Collection library	QPFRRDATA	Name, *SAME
Default collection profile . . .	*STANDARDP	*SAME, *MINIMUM, *STANDARD...
Cycle time	000100	Time, *SAME
Cycle interval	24	*SAME, 1-24 hours
Collection retention period:		
Number of units	00024	*SAME, 1-720, *PERM
Unit of time	*HOURS	*HOURS, *DAYS
Create database files	*YES	*SAME, *YES, *NO
Change PM/400 library	*NO	*NO, *YES

Notes: Configure Collection Services

This foil shows the option 2 on the Perform menu – Collect performance data.

Note the ability to influence the performance database file library used by PM eServer iSeries.

This option will change the PM/400 library to match the library used by Collection Services. PM/400 has its own configuration. When PM/400 is running it will change the Collection Services configuration (on an hourly basis) so that the data it requires is collected and stored in the PM/400 configured library (default QPMGDATA). If you specify *YES, the PM/400 collection library will be changed to the Collection Services collection library. If the Collection Services server job is running when this option is specified, the collection will be cycled so that the change to both libraries will take effect simultaneously.

If the PM/400 library is changed, you should be aware of the following:

PM/400 will not be able to process any data which was collected in the previous library. In addition, PM/400 will automatically remove data in the new library based on the PM/400 setting Performance data purge days. To change the number of days performance data is retained type GO PM/400 and select the Work with PM/400 Customization option.

Work with System Activity - WRKSYSACT

- Command changed to calculate and show values
 - Consistent with other performance tools
- CPU utilization does NOT scale to whole virtual processors
- HVLPTASK is excluded from list of tasks that consume CPU
- Command shows CPU values higher than 100% for uncapped processors

Notes: Work with System Activity - WRKSYSACT

The WRKSYSACT command was changed to calculate and show values consistent with the other performance tools.

The CPU utilization in WRKSYSACT does not scale to whole virtual processors anymore.

The HVLPTASK job is excluded from the list of tasks that consume CPU.

In addition, WRKSYSACT was enhanced to show CPU values higher than 100% for uncapped processors, instead of capping the CPU utilization to 99.9%.

iSeries Navigator Performance Tools Graphics plug-in

- High Disk Utilization graph excludes intervals with no disk information
- User Pool Faults/Second graph is enhanced to graph more than 1 pool
- Library and member name added to Display Performance Data window title
- New metrics graphed:
 - New Interactive Feature Utilization Graph is added
 - New Uncapped CPU% Available metric added to Total CPU Utilization graph

Notes: iSeries Navigator Performance Tools Graphics plug-in

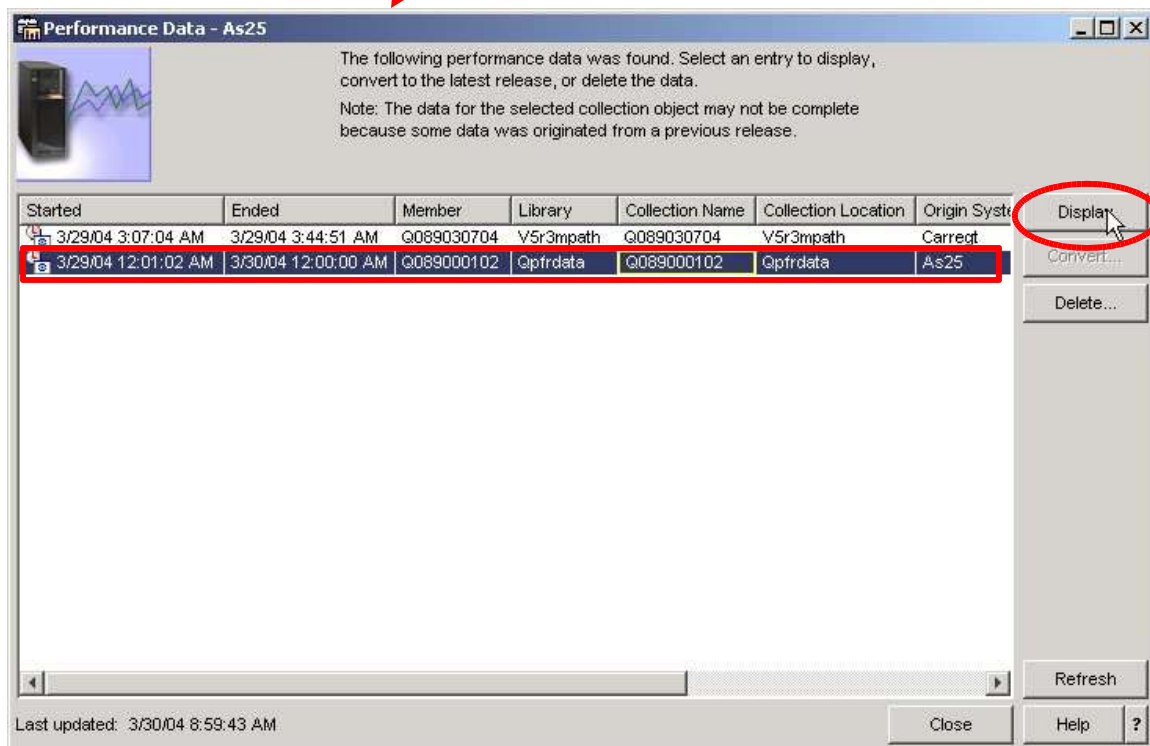
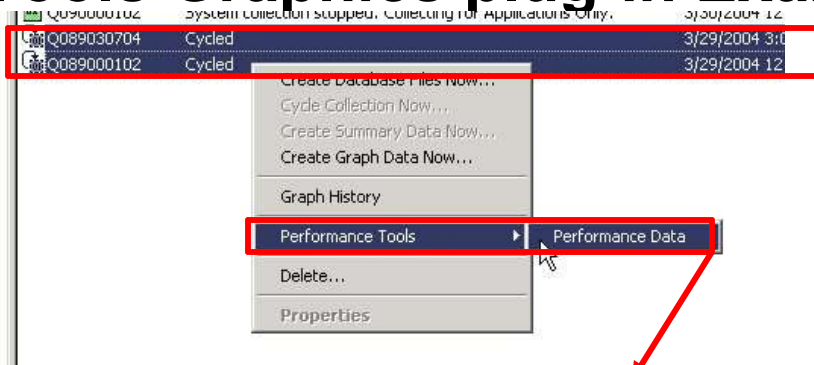
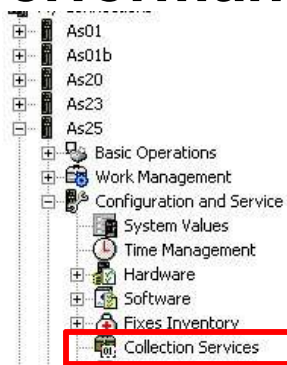
The Performance Tools -> Performance Data graphical display of various performance metrics has both some metrics graphically displayed and user interface changes:

- High Disk Utilization graph now excludes intervals with no disk information.
- The User Pool Faults/Second graph was enhanced to graph more than one pool, with one line per pool.
- Library and member name was added to Display Performance Data window title, to accommodate the situation where more than one window is opened.

New performance metrics were added to the graphical displays:

- A new graph was also added to show the Interactive Feature Utilization metric when it is provided by Collection Services.
- A new metric, Uncapped CPU% Available, was also added to the Total CPU Utilization graph.

Performance Tools Graphics plug-in Example scenario



Notes: iSeries Navigator

This foil shows the iSeries Navigator interface for Performance Tools. You can select **Collection Services** from the Configuration and Service container. You can then select the **Collection Name(s)**, right-click on it and select **Performance Tools->Performance Data** which will bring you to the Performance Data window. You may then select the Collection Name and click **Display** to show the corresponding performance data.

Performance Tools Graphics: Int Feat Utilization example

Display Performance Data (V5r3mpath/Q089030704)

File Graphs View Reports Help

From: Beginning Mar 29, 2004 3:10:00 AM
To: End Mar 29, 2004 3:40:00 AM
Update

Interactive Feature Utilization

User Pool Faults/Second - Custom List

View: All Jobs Time period: 3:10:00 AM - 3:40:00 AM

Job	User	Number	Type	CPU%	Trans
Job0303	As0303	005392	Batch	34.46	
Job0302	As0302	005391	Batch	32.99	
Job0301	As0301	005390	Batch	24.70	
Qpadev...	Jcook	005395	Pass through	0.41	
Crint01			LIC task	0.25	
Gypsjsvr	Gypsjsvr	005366	Batch	0.09	
Smxcag...			LIC task	0.07	
Crint02			LIC task	0.04	
Crtpfrdta	Gsys	005387	Batch	0.02	
loeladdent			LIC task	0.02	
ldelandev-			LIC task	0.01	
Gsysco...	Gsys	001368	System	0.01	
Gzrcsrvs	Quser	005397	Batch	0.01	
Smiorest...			LIC task	0.01	
Chainbch	Gsys	005389	Subsystem	0.00	
Cr-mgr			LIC task	0.00	
Dd-freea...			LIC task	0.00	
Dd-freea...			LIC task	0.00	
Dd-freea...			LIC task	0.00	

Total CPU Utilization

Summary - 3:10:00 AM - 3:40:00 AM

CPU Utilization (interactive):	.41 %	Transactions:	128	Disk busy:	.48 %
CPU utilization (other):	92.73 %	Transactions per hour:	234	Disk I/O:	56.9 per sec
Job Count:	56	Average response:	2.63 sec	Logical DB I/O:	0
Interactive Feature Utilization:	.14 %	Time exceeding Int CPU Threshold:	0 sec	Reads/Writes:	1.802

© 2004 IBM Corporation
PAGE 71

iSeries. mySeries.

Notes: iSeries Navigator Plug-in continued

This slide shows the new Interactive Feature Utilization graph.

The Interactive Feature Utilization graph is shown along with Time Exceeding Interactive CPU Threshold metrics displayed in the Summary window with Interactive Threshold to Data Properties page if the information is available in the collection.

iSeries Navigator Plug-in: Storage Pools example

The screenshot shows the 'Display Performance Data - As25 (Qpfrdata/Q089000102)' window. The 'Preferences...' menu is open, and the 'Graph Preferences' dialog is shown. In the 'Graph Preferences' dialog, the 'User Pools' tab is selected, and the 'Select pools to graph' radio button is chosen. The 'Pools in the system' table is as follows:

Pool ID	Faults/Sec...
003	0
004	2.36

The 'Pools to graph' table is as follows:

Pool ID	Faults/Sec...
002	24.84

Below the graph, a summary of performance data is displayed for the time period 12:05:00 AM - 12:00:00 AM:

CPU Utilization (interactive):	.29 %	Transactions:	116	Disk I/O:	.0 per sec
CPU utilization (other):	1.69 %	Transactions per hour:	116	Logical DB I/O:	72
Job Count:	761	Average response:	.73 sec	Reads/Writes:	.000
Interactive Feature Utilization:	.06 %	Time exceeding Int CPU Threshold:	0 sec		

Notes: iSeries Navigator Plug-in: Storage Pools example

The User Pool Faults/Second graph was enhanced to graph more than one pool, with one line per pool.

This foil shows the two options for the User Pool Faults/second graph, which you can access by selecting Preferences from the Graphs menu.

The possible options are:

- Graph the top 10 pools with the highest fault rate, depending on the time period you select.
- Graph specific pools that you select.

This enhancement helps performance analysts view the impact that changes in the size of pools has on fault rates, and also gives them a better understanding of pool activity in general.

Performance Explorer

Performance Explorer – ADDPEXDFN

- Add threads/tasks option (ADDTHDOPT)
 - Specifies what types of threads and tasks should be included
 - Based on creation time of threads and tasks relative to start of PEX session
- Added Randomize element to INTERVAL parameter
 - *FIXED
 - *VARY
- Added Event format element to
 - Base events (BASEVT) parameter
 - Communications events (CMNEVT) parameter
 - Describes what data is collected for this event
 - *FORMAT1 provides the data used for most data analysis
- Added Save/Restore events (SAVRSTEVT) parameter
 - Specifies which save/restore events are included in the definition

Notes: Performance Explorer – ADDPEXDFN

Add threads/tasks option (ADDTHDOPT)

Specifies what types of threads and tasks should be included in the Performance Explorer session based on the creation time of the threads and tasks relative to the start time of the Performance Explorer session.

Added Randomize element to INTERVAL parameter

You can specify *FIXED or *VARY. If *FIXED is specified, the sampling interval will always be the value specified for element 1 of this parameter. If *VARY is specified, the sampling interval will be approximately the value specified for element 1 of this parameter. The specified interval will be changed each time by adding or subtracting a small random percentage of the interval. This may be necessary to eliminate harmonics when the set of threads/tasks in a collection have settled into a steady state sequence

Added Event format element

This is added to Base events (BASEVT) parameter and Communications events (CMNEVT) parameter.

The event format describes what data is collected for this event.

*FORMAT1 provides the data used for most data analysis. The other formats allow for collection of other data related to these events. Values other than *FORMAT1 are valid for only the *PMCO and *SWOQ events.

For all other events, *FORMAT1 will be used regardless of what format is specified.

Added Save/Restore events (SAVRSTEVT) parameter

Specifies which save/restore events are included in the definition.

Performance Explorer – ADDPEXFTR

- Added Java trigger parameter (JVATRG)
 - If a Java method entry event occurs that matches this trigger specification
 - All events specified in PEX definition are collected
 - only for the thread where the trigger occurs
- Added Java class filter parameter (JVACLSFTR)
 - Specifies the Java package and class
 - To be used as compare values for the Java class filter

Notes: Performance Explorer – ADDPEXFTR

Added Java trigger parameter (JVATRG)

If a Java method entry event (*JVAENTRY) occurs that matches this trigger specification, then performance explorer collects all events specified in the performance explorer definition used for the active performance explorer session.

The events are collected only for the thread where the trigger occurs.

Added Java class filter parameter (JVACLSFTR)

Specifies the Java package and class to be used as compare values for the Java class filter.

iDoctor for iSeries update

the three modules *for V5R3*



PEX Analyzer

System Job Watcher

Heap Analyzer

Notes: iDoctor for iSeries update

You call on iDoctor for iSeries when the more commonly used iSeries performance tools are not able to identify the cause of a performance bottleneck.

In a generalized scenario you may Management Central monitors to identify high resource utilizations and the combination of Collection Services and the Performance Tools for iSeries, licensed product 5722-PT1 capabilities to identify high resource utilizations and point to jobs that appear to be related to the problem. At the job level you can use Performance Explorer (PEX) and the associated 5722-PT1 reports to zero in on programs within one or more jobs. In many cases all of these tools provide sufficient problem identification information.

When these tools cannot identify the problem the set of iSeries tools available under iDoctor need to be considered.

To access iDoctor for iSeries, perform the following from your browser:

- Enter <http://www.ibm.com/eserver/series/support>
- Select iSeries Tools -> Performance Analysis (iDoctor for iSeries)

iDoctor is both a service and a set of "detail level" performance tools (used by the service) for analyzing OS/400-provided Performance Explorer trace data, which includes a "higher than PEX trace data analyzer level" tool called Job Watcher.

Performance Explorer trace data (provided by OS/400) is "condensed" into summary reports and graphs to help isolate performance issues that other tools were not able to do for you.

The Basic Performance Assessment is a fee-based service. By clicking a web link, an IBM representative will contact you with a no obligation price quote to assess your business need. You can get a copy of the iDoctor support for a trial period. Consult the web site for details.

The remainder of this section describes only the "Job Watcher" function on iDoctor – at a V5R3 planned set of enhancements level.

Job Watcher can be used when ...

- fits between regular performance tools and PEX analysis tools
- doesn't replace nor depend on either tool set
- *and* does not replace WRK... commands

makes everyone more productive.

- experts get focused more quickly on the specifics of a complaint
- regular users can capture the specifics easier, quicker
- *and* many complaints can be validated and worked without the delays and expense of getting experts involved
- *or* the right expert can be engaged more promptly

Job Watcher progress...

Job specific (V5R1)

- limited number of Job_Threads/Tasks
- frequent, detailed snapshots of selected jobs

System wide (V5R2)

- all Job_Threads/Tasks
- less frequent, less detailed snapshots of all jobs


System Job Watcher (V5R3)

- all Job_Threads/Tasks
- new, more efficient, versatile snapshot engine

A Job Watcher is ...

- a real time, non-intrusive, snapshot oriented, run/wait analysis of a job_thread/task

The Run/Wait Signature (% of time by component by interval)



Summary of all intervals (% of time by component)



Job Watcher DB record lock example

- infrequent, occasional, happen most often when system is busy
- complainer is the victim and just says response is long, sometimes
- WRK... commands not handy, timely as problem does not persist
- performance tools 15 minute intervals don't provide enough detail
- Don't know when or how long to run traces (and don't want to)
- DB record locks are generally application oriented
- they tend to occur randomly, during high volume requests from multiple users and/or multiple batch jobs scenarios
- system-wide even better than job-specific Job Watcher