

# Université IBM i 2018

16 et 17 mai

IBM Client Center Paris



## S29 – Best Practices for IBM i Performance Management

**Stacy L. Benfield**

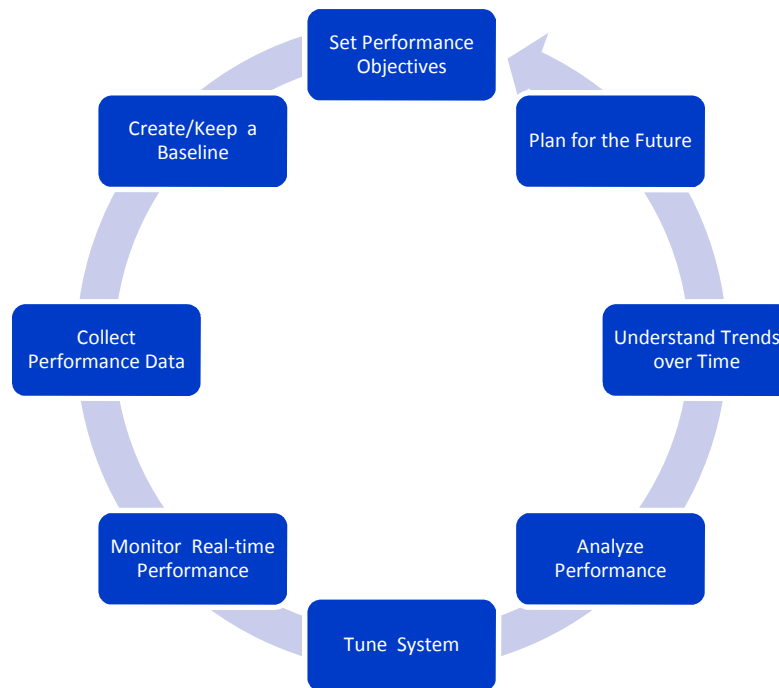
IBM i Performance Consultant - Lab Services Power Systems Delivery Practice

IBM i Large User Group (LUG) Program Manager

IBM, Rochester MN

*stacylb@us.ibm.com*

# Performance Management Life-cycle



The background of the slide is a complex network diagram. It consists of numerous small, light gray circular nodes connected by thin, light gray lines. The nodes are scattered across the page, with a higher density in the center and right side, creating a web-like structure that frames the central text.

# Objectives



## Setting Objectives

- User objectives
  
- Management objectives
  - ✓ *Overnight batch workload must complete in 6 hours....*
  - ✓ *Interactive response time must be sub-second.....*
  - ✓ *1000 widgets must be processed every hour...*
  - ✓ *Average CPU must be below 80% during core hours....*
  - ✓ *Disk used percentage must remain below 70%....*

**Keep in mind.....**




# Many Factors can Influence Performance

- Hardware capabilities
  - CPU, Memory, Disk....
  
- Configuration
  - System Values, subsystem configuration, pool sizes, etc.
  
- Software
  - Application software, IBM Software, LPPs, PTFs, etc.
  
- Workload
  - Variability → peak days, end-of-month
  
- Network
  
- *Or something external to IBM i.....*

The background of the slide is a complex network diagram. It consists of numerous small, light gray circular nodes connected by thin, light gray lines. The nodes are arranged in a somewhat irregular pattern, with a higher density in the center and fewer nodes towards the edges. The overall effect is that of a web or a social network.

# Baseline


## Patterns in Performance Data



- Performance data typically has patterns
  - Daily, weekly, monthly, yearly

- ★ Understand your typical patterns
- Recognize change

# Create a Baseline



- A baseline is the expected performance characteristics over a defined period of time
  - The baseline provides the data to compare with to identify changes
  - You may need to **keep** multiple baselines due to varying business periods
    - Day-to-day operations
    - Month-end
    - Year-end
    - Prior to any significant software/hardware changes!
- The baseline is the **reference point** for
  - Capacity planning and trending
  - Identifying impacts of changes in
    - workloads
    - applications
    - operating system
    - hardware
    - .....




## Create a Baseline

 Use Collection Services data for the baseline

- The baseline can be simple or complex, depending upon the need
  - Identify the metrics that are key to optimal performance
    - Average response time
    - Batch window time
    - Average CPU utilization
    - Disk used percentage
    - Number of interactive transactions
    - .....

## Keep Collection Services Baseline Data

- At a minimum, keep the \*MGTCOL object
  - Move to different library or change Expiration Date:

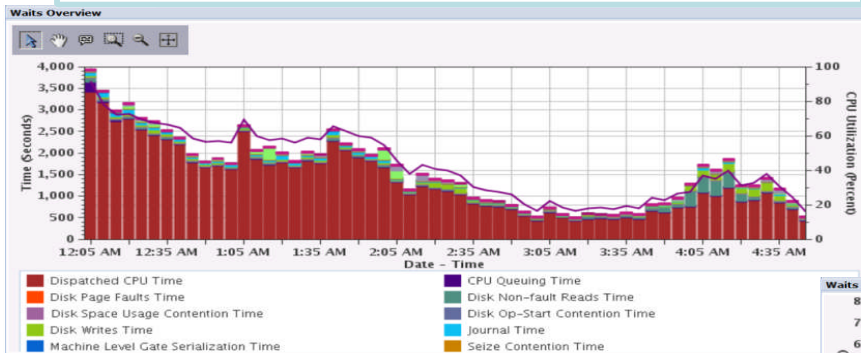


(CFGPFRCOL command)

- Database files can always be recreated if you have the \*MGTCOL
- Easy to move to another partition

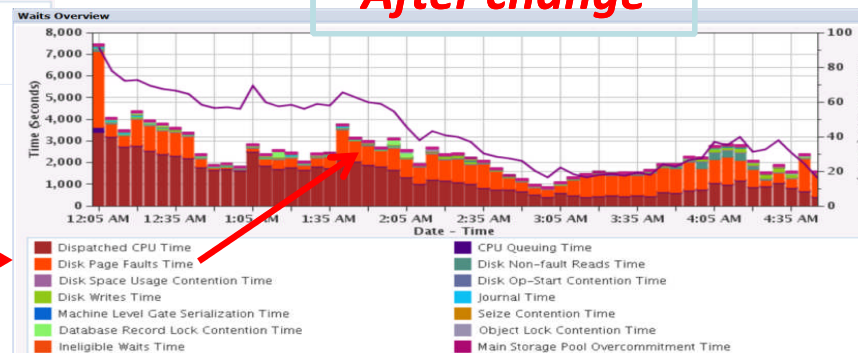
# Data Retention – Baseline benefits

**Baseline data - Before change**



Wait Accounting technology surfaced in Collection Services data can be extremely beneficial in determining what has changed from Baseline data.


**After change**



www.faccinefb.com

# Collect Performance Data

# Collect System-wide Performance Data



## ★ Collection Services Data 24/7

- 15 minutes is the default interval, 5 minute intervals provide better granularity for perform analysis
- Keep 7-10+ days for history + Baseline data

### Why??

- If something goes wrong, you have data that will help **analyze** the problem, **fix** it, and **prevent** it from happening in the future
- If you can't solve the problem, you have information that makes it easier for IBM Support to **solve the problem faster**
- To provide a **reliable baseline** so you can **understand the impact** that a software, network, or environmental **change** had on the performance of your system
- To provide historical information that enables you to **plan for future growth** based on real trends, not guesses.

The background of the slide is a complex network diagram. It features numerous small, light gray circular nodes connected by thin, light gray lines. The nodes are distributed across the page, with a higher density in the center and right side, creating a web-like structure. The overall appearance is that of a data network or a social graph.

# Monitoring

## Real-time Monitoring

- ★ Real-time monitoring is essential for pro-active performance management
- ★
  - Automated monitoring to alert you of potential problems

Identify and correct potential issues **before** they become real problems



## Real-time Monitoring

- Several tools/interfaces available for real-time monitoring:
  - CL Commands
    - WRKACTJOB, WRKDSKSTS, WRKSYSACT, WRKSYSSTS, etc...
  - IBM Navigator for i
    - Active Jobs, Disk Status, System Status, etc...
    - Dashboard
    - Monitors (7.2) – System and Message
  - Management Central Monitors
    - System, Job, Message, File



# Navigator benefits – Active Jobs

- Performance and status information for the **active jobs** in the system
- **Good high level starting point to see high resource consumers**
  - May provide “scope” information useful for other tools
- Additional performance related columns (metrics) + Integration to Collection Services/PDI

Active Jobs

Refresh Elapsed time: 00:04:17

Performance

- Investigate Data
- Manage Collections

All Tasks

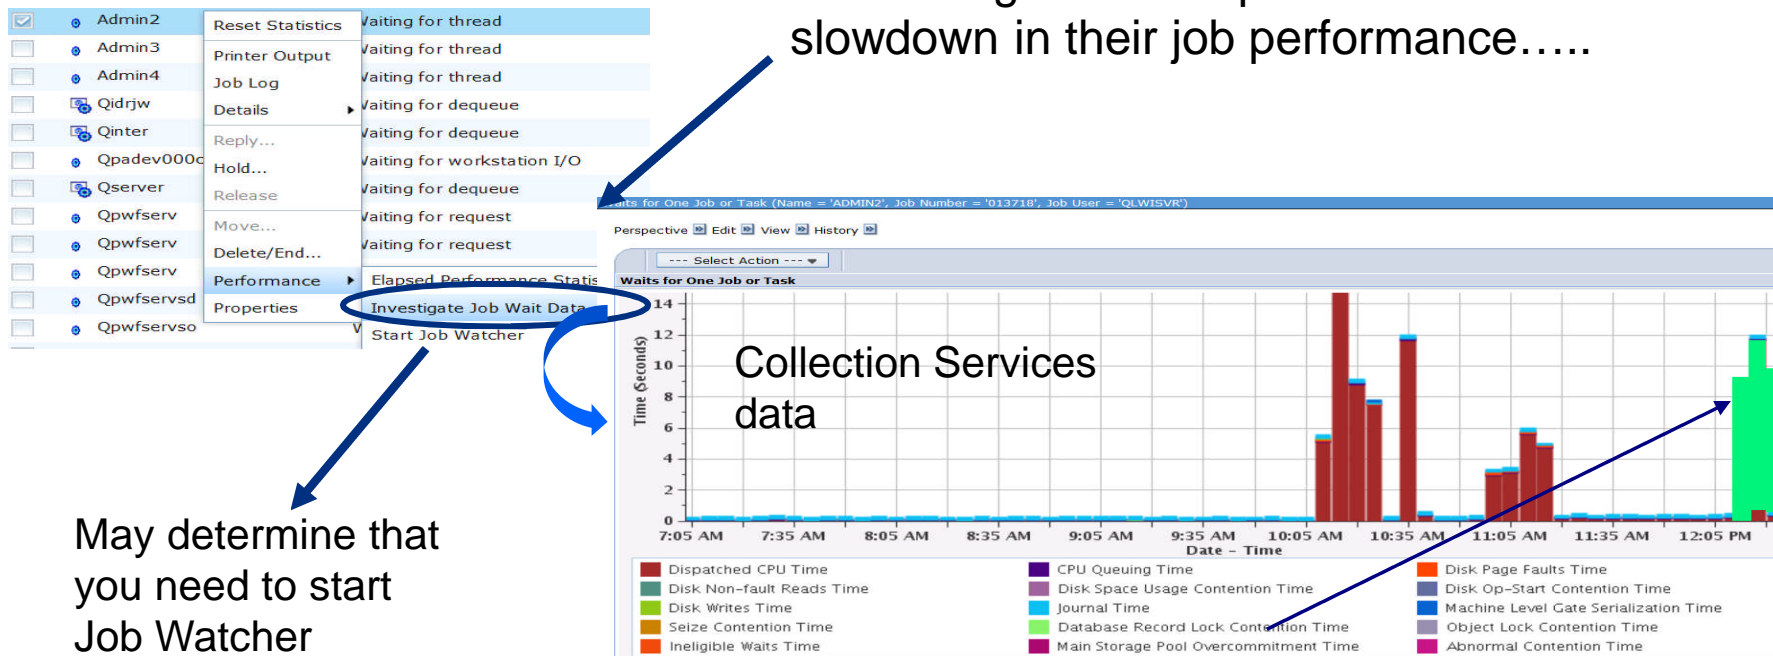
- Active Jobs
- Disk Status
- Manage Collection:
- Investigate Data
- Performance Mana
- System Status
- Collections
- Reports
- Collectors

Job Name	Detailed Status	Current User	Type	CPU %	Synchronous Disk I/O Rate	Memory Pool	Page Fault Rate	Run Priority	Thread Count
Admin2	Waiting for thread	Qlwisvr	Batch immediate - Server	0.4	4.7	Base	3.3	25	90
Qbatch	Waiting for dequeue	Qsys	Subsystem	0	0	Base	0	0	2
Qbatch2	Waiting for dequeue	Qsys	Subsystem	0	0	Base	0	0	2
Qcmn	Waiting for dequeue	Qsys	Subsystem	0	0	Base	0	0	2
Qacsotp	Waiting for request	Quser	Prestart communications	0	0	Base	0	20	1
Qlzpssrv	Waiting for request	Quser	Prestart communications	0	0	Base	0	20	1
Qnmapingd	Waiting for request	Quser	Prestart communications	0	0	Base	0	25	1
Qnmarexecd	Waiting for request	Quser	Prestart communications	0	0	Base	0	25	1
Qnpservr	Waiting for request	Quser	Prestart communications	0	0	Base	0	20	1
Qzrcsvr	Waiting for request	Quser	Prestart communications	0	0	Base	0	20	1
Qzscsvr	Waiting for request	Quser	Prestart communications	0	0	Base	0	20	1

# Real Time – Active Jobs Integration

- Launch points to other performance functions which can be useful in problem determination:

A user begins to complain about a slowdown in their job performance.....




May determine that you need to start Job Watcher

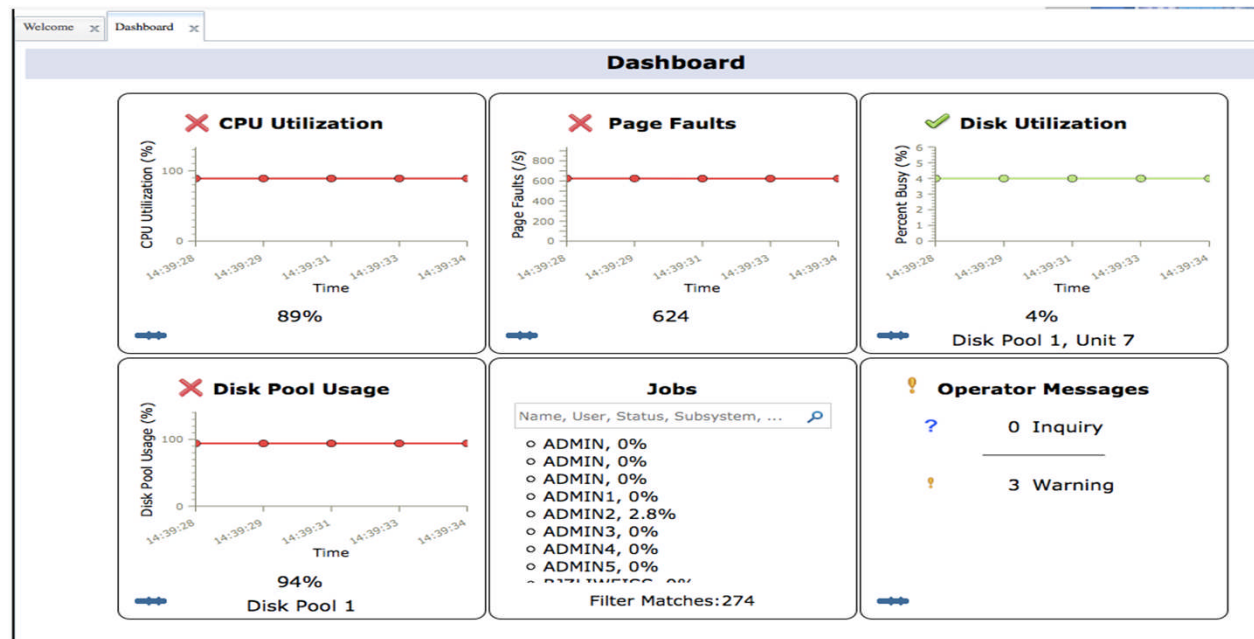
Collection Services data

# Dashboard

- Task in Navigator
- Also available with iAccess Mobile



- Key metrics updated in real time
- Set thresholds to visualize potential issues



# Configurable intervals and thresholds






Welcome x Dashboard x

**Dashboard** ▼

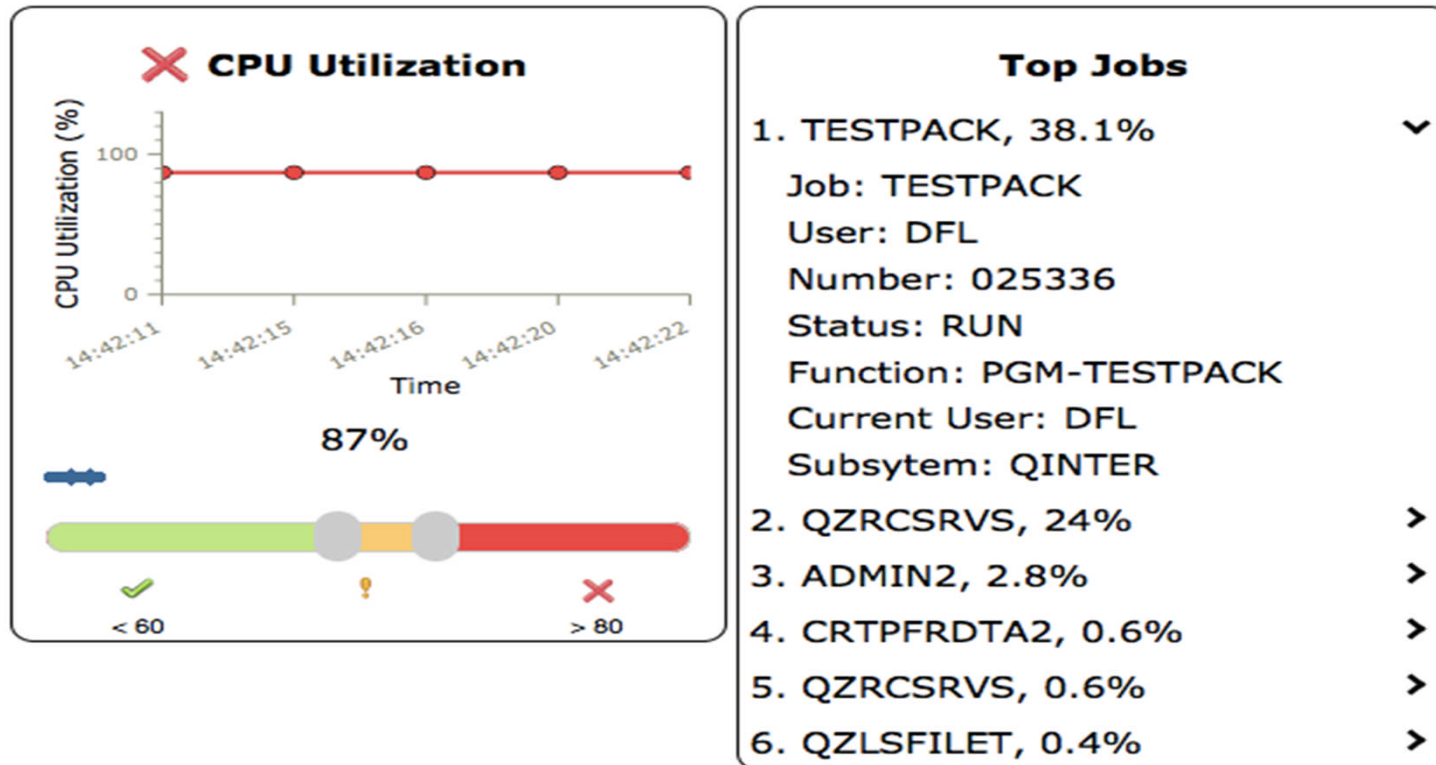
Auto

Interval: 1s Duration: 1s

Stop Reset




# Basic drill-down - CPU Utilization



# Jobs - search and drill-down

**Jobs**

dawn 

- QZDASOINIT, 0%
- QZDASOINIT, 0%
- QZDASOINIT, 0%
- QZRCSRVS, 0.6%

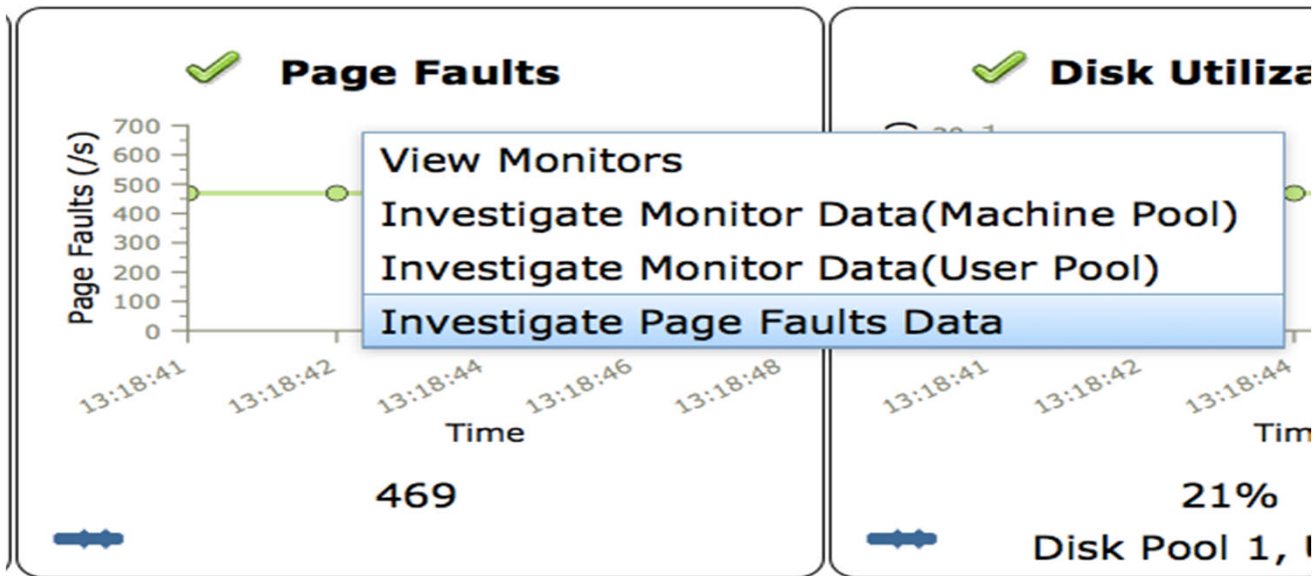
Filter Matches:4

**Jobs**  
Filter: dawn

- QZDASOINIT, 0% >
- QZDASOINIT, 0% >
- QZDASOINIT, 0% >
- QZRCSRVS, 0.6% ▾**

Job: QZRCSRVS  
User: QUSER  
Number: 041924  
Status: RUN  
Function:  
Current User: DAWN  
Subsystem: QUSRWRK

# Drill-down to System Monitors and Investigate Data



# iAccess Mobile Dashboard

View from your favorite mobile device

- Phone
- Tablet

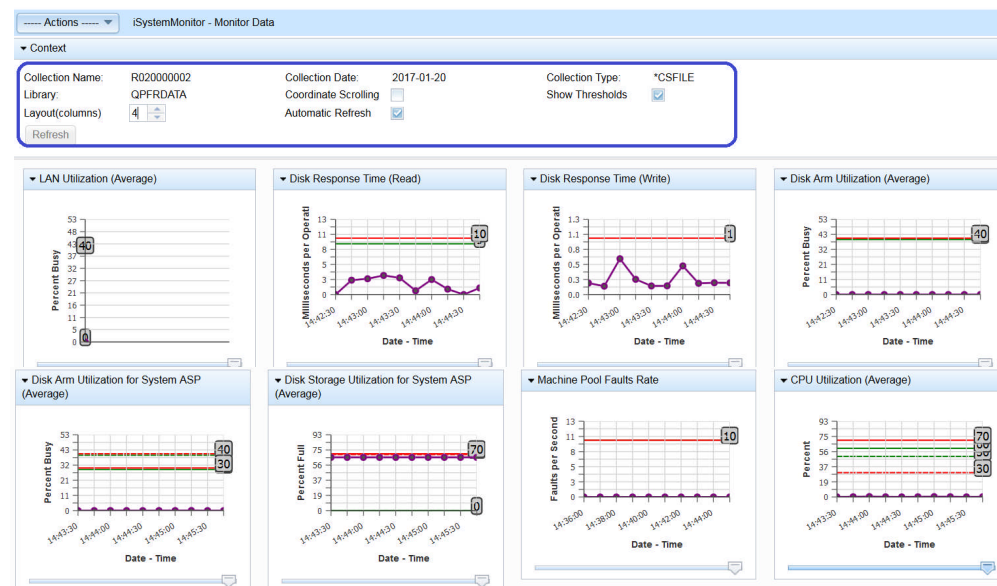
<http://system.name:2001/iamobile>





## 7.2 Navigator Monitors

- System monitors
  - Now available in IBM Navigator for i
  - Similar in concept to Management Central
  - Leverages Performance Data Investigator (PDI) function
  - Additional **new** metrics for monitoring
    - CPU utilization – uncapped
    - CPU utilization – SQL
    - Disk utilization for IASP
    - Temporary storage used
    - Spool file creation rate
    - And more....





## System Monitors – Highlights


- See key data from a **single**, central interface
- **Visually monitor** near real-time data
- **Capability to activate thresholds**
- **Ability to define action plans**

# System Monitors - Overview

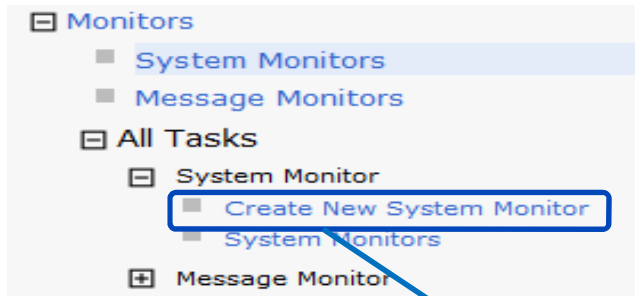
- System Monitors gather and present **near real-time performance** data
  - Monitor the health of your system
  - Identify *potential* performance problems before they become *serious* issues
  
- **Thresholds** can be defined to trigger an **action** when a system wide performance metric exceeds the defined “comfort level”
  - For example, when CPU Utilization exceeds 80%, send a message to notify the operator
  
- System Monitors provide powerful capabilities to monitor what is happening on your system

...BUT finding out what caused the problem often requires other performance analysis tools

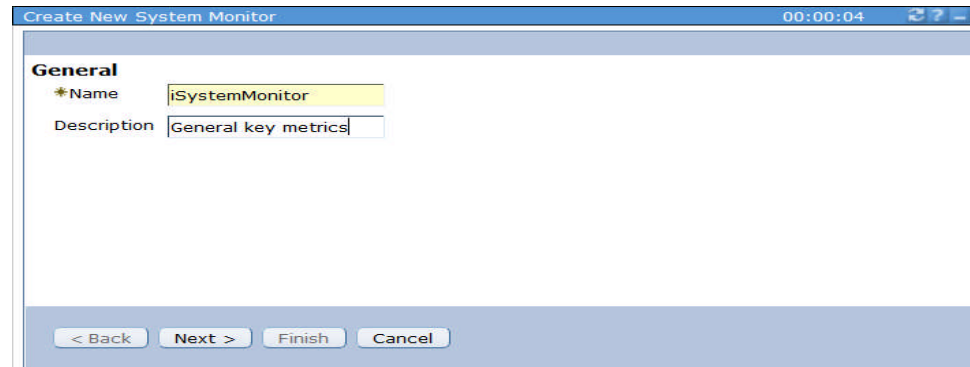
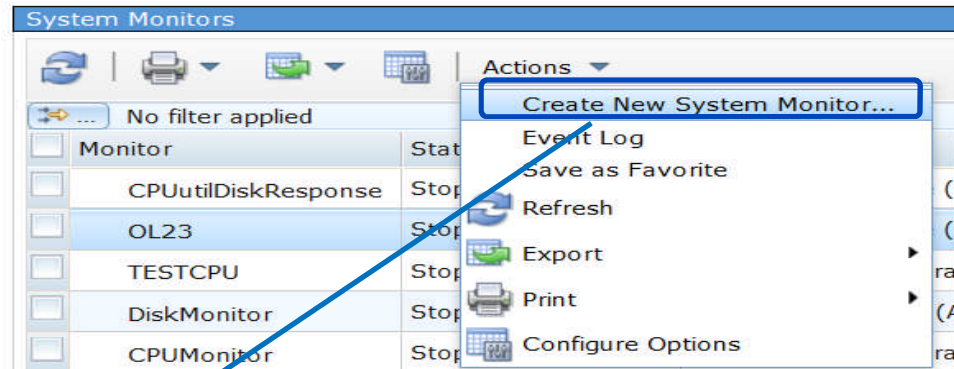
# Monitor Setup - Overview



# Create a New System Monitor (Step 1)



-OR-



- Specify Monitor name
- Description (optional)

# Select Metric(s) to Monitor (Step 2)



Links to configure thresholds and actions

Create New System Monitor

**Metrics**

Available Metrics:

- Metrics
- Communications Line Utilization (Average)
- Communications Line Utilization (Maximum)
- LAN Utilization (Average)
- LAN Utilization (Maximum)
- Disk Arm Utilization (Maximum)
- Disk Storage Utilization (Average)
- Disk Storage Utilization (Maximum)
- Disk Arm Utilization for System ASP (Average)
- Disk Arm Utilization for System ASP (Maximum)
- Disk Storage Utilization for System ASP (Maximum)
- Disk Arm Utilization for User ASP (Average)

Add >

< Remove

Metrics to monitor:

- Metrics
- [Disk Response Time \(Read\)](#)
- [Disk Response Time \(Write\)](#)
- [Disk Arm Utilization \(Average\)](#)
- [Disk Storage Utilization for System ASP \(Average\)](#)
- [Machine Pool Faults Rate](#)
- [CPU Utilization \(Average\)](#)

Note: Only partial list of Available Metrics is shown here...

< Back Next > Finish Cancel

# Full List System Monitor Metrics

- CPU Utilization (Average)
  - CPU Utilization (Interactive Jobs)
  - CPU Utilization(Uncapped)
  - CPU Utilization(SQL)
  - Interactive Response Time (Average and Maximum)
  - Transaction Rate (Interactive)
  - Batch Logical Database I/O
  - Disk Response Time (Read)
  - Disk Response Time (Write)
  - Disk Arm Utilization (Average and Maximum)
  - Disk Arm Utilization for User/System/Independent ASP (Average and Maximum)
  - Disk Storage Utilization (Average and Maximum)
  - Disk Storage Utilization for User/System/Independent ASP (Average and Maximum)
  - Communications Line Utilization (Average and Maximum)
  - LAN Utilization (Maximum and Average)
  - Machine Pool Faults
  - User Pool Faults (Maximum and Average)
  - Spool File Creation Rate
  - Shared Processor Pool Utilization (Virtual and Physical)
  - Temporary Storage Utilization
- HTTP Server Metrics
    - HTTP Requests Received Rate
    - HTTP Requests Received (Maximum)
    - HTTP Responses Sent Rate
    - HTTP Responses Sent (Maximum)
    - HTTP Non-Cached Requests Processed (Average and Maximum)
    - HTTP Error Responses Sent (Average and Maximum)
    - HTTP Non-Cached Requests Processing Time (Total and Highest Average)
    - HTTP Cached Requests Processing Time (Total and Highest Average)

**Red are new with 7.2 Navigator Monitors**

**Blue are new with 7.5**



# Metric Descriptions

Additional information on metrics available can be found here:



[https://www.ibm.com/support/knowledgecenter/en/ssw\\_ibm\\_i\\_73/rzahx/rzahxmonconmetrics.htm](https://www.ibm.com/support/knowledgecenter/en/ssw_ibm_i_73/rzahx/rzahxmonconmetrics.htm)

- Example: Descriptions of CPU Utilization metrics

Metric groups	Metric description
CPU Utilization	<p>The percentage of available processing unit time that is consumed by jobs on your system. Choose from the following types of CPU Utilization metrics for use in your monitors:</p> <ul style="list-style-type: none"><li>• <b>CPU Utilization (Average):</b> Configured CPU percent unscaled.</li><li>• <b>CPU Utilization (Interactive Jobs):</b> Configured CPU percent that is consumed by interactive jobs.</li><li>• <b>CPU Utilization (Uncapped):</b> Uncapped CPU percent unscaled. The amount of unscaled system CPU consumed relative to the maximum uncapped CPU the partition could consume based on the number of the virtual processors that are assigned to the partition and the capacity of the shared virtual pool.</li><li>• <b>CPU Utilization (SQL):</b> SQL CPU percent unscaled. The amount of unscaled system CPU consumed performing work that is done on behalf of SQL operations relative to the configured CPU time (processor units) available to the partition.</li></ul>



## Configure Metric(s) (Step 3) - Collection Interval

- Recommend **60 seconds** (or greater) 
  - Less than this will generate very large monitor collections
  
- The default collection interval on 7.3 is 60 seconds 
  
- The default collection interval on 7.2 is **15 seconds**
  - 7.2 June 2016 SP changed the default to 60 seconds
  - Manual steps to recover if you used 15 second intervals:

<http://www.ibmssystemsmag.com/Blogs/i-Can/September-2016/IBM-i-Navigator-System-Monitors--Collection-Interv/>

Configure Metric

**CPU Utilization (Average)**

Collection Interval 60  Seconds

# Configure Metric(s) – Thresholds (Step 4)

- **Threshold** – Defines acceptable variances for a specific metric being monitored
  - Optional, check to enable
  - Up to two **thresholds** may be enabled for each metric that the monitor is collecting
    - Examples: Warning and critical levels, too high and too low values

**Threshold1**

Enable Threshold

Trigger: >=  Percent

Duration:  Intervals

Operating System Command:  Prompt...

Reset: <  Percent

Duration:  Intervals

Operating System Command:  Prompt...

**Threshold2**

Enable Threshold

Trigger: >=  Percent

Duration:  Intervals

Operating System Command:  Prompt...

Reset: <  Percent

Duration:  Intervals

Operating System Command:  Prompt...

Thresholds provide capability for automatic actions to occur 

# Configure Metrics – Thresholds continued...



- An **Action** is an Operating System **Command**
  - It defines what should occur when the threshold is hit
  - Any command that can run in batch
  - CALL command exists - you can call a program to take whatever action you wish
  - Command is run in a job under the user profile that created the monitor

**Threshold1**

Enable Threshold

Trigger:   Percent

Duration:  Intervals

Operating System Command:

Reset:   Percent

Duration:  Intervals

Operating System Command:

# Configure Metric(s) – Thresholds continued...



- An **action** can be specified to be taken when:

- A **trigger** value is reached
- A **reset** value is reached

**Threshold1**

Enable Threshold

Trigger:	>=	0	Percent
Duration:		1	Intervals
Operating System Command:			Prompt...
Reset:	<	0	Percent
Duration:		1	Intervals
Operating System Command:			Prompt...

- **Duration**
  - Specify how long (in terms of collection intervals) the condition must occur before the action is taken
- An **event** is added to the **Event Log** whenever the trigger value or the reset value is reached.



## Action: Send an email example

- If SMTP (Simple Mail Transfer Protocol) is enabled, an action can be specified to send an email when a trigger value is reached:

### Example set up to send email from IBM i:

- Make sure the SMTP server is started. To start use:  
STRTCPSVR SERVER(\*SMTP)
- Add the sending user profile name to the SMTP Local Mailbox Directory  
ADDUSRSMTMP
- Use SNDSMTPEMM command to send the email  
SNDSMTPEMM RCP(ibmuser@us.ibm.com) SUBJECT('Monitor &MON triggered.') NOTE('Monitor &MON has CPU Utilization under the low value of &TVAL for &TDUR interval(s). Current value is &VAL.')

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20i%20Technology%20Updates/page/Send%20email%20as%20a%20monitor%20action>

# Configure Metric – CPU Utilization example

Configure Metric

**CPU Utilization (Average)**

Collection Interval: 60 Seconds

**Threshold1**

Enable Threshold

Trigger:  $\geq$  75 Percent

Duration: 5 Intervals

Operating System Command: HLDJOBQ JOBQ(QGPL/Q) Prompt...

Reset:  $<$  60 Percent

Duration: 10 Intervals

Operating System Command: RLSJOBQ JOBQ(QGPL/Q) Prompt...

**Threshold2**

Enable Threshold

Trigger:  $<$  30 Percent

Duration: 10 Intervals

Operating System Command: SNDMSG MSG('Low CPU) Prompt...

Reset:  $\geq$  50 Percent

Duration: 5 Intervals

Operating System Command: SNDMSG MSG('Workload) Prompt..

OK Cancel

- ✓ Specify Collection Interval
- ✓ Enable Threshold(s) and specify parameters

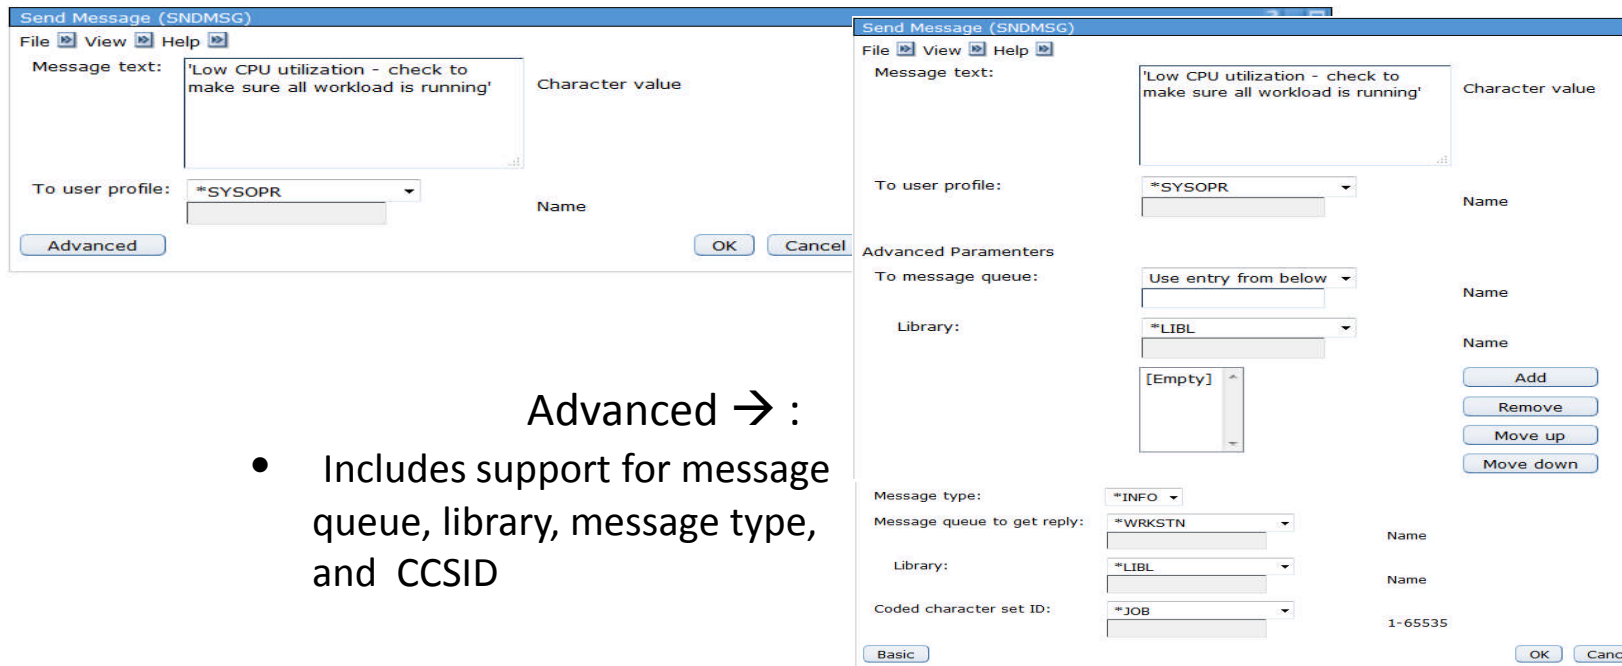
This example demonstrates:

- Enabling of Thresholds (2)
- Duration
- Actions to take for Trigger and Reset conditions

# Command Prompt Example

## Basic:

Configure Metric



**Send Message (SNDMSG) - Basic Tab**

File View Help

Message text: 'Low CPU utilization - check to make sure all workload is running' Character value

To user profile: \*SYSOPR Name

Advanced OK Cancel

**Send Message (SNDMSG) - Advanced Parameters Tab**

File View Help

Message text: 'Low CPU utilization - check to make sure all workload is running' Character value

To user profile: \*SYSOPR Name

Advanced Parameters

To message queue: Use entry from below Name

Library: \*LIBL Name

[Empty]

Add Remove Move up Move down

Message type: \*INFO

Message queue to get reply: \*WRKSTN Name

Library: \*LIBL Name

Coded character set ID: \*JOB 1-65535

Basic OK Cancel Help

Advanced → :

- Includes support for message queue, library, message type, and CCSID

# System Monitor Replacement Variables

- Replacement variables are available to customize your actions to the specific metric and values
- Specify replacement variables on the IBM i command

>> Without replacement values:

```
SNDDMSG MSG('Warning: CPU utilization has exceeded
threshold') TOUSR(*SYSOPR)
```

<input type="checkbox"/>	From User	Type	Message
<input type="checkbox"/>	Stacyb	Information	Warning: CPU utilization has exceeded threshold

>> Using replacement variables:

```
SNDDMSG MSG('Warning: &MON has exceeded threshold &TVAL for &TDUR
interval(s); the current value is &VAL.')
TOUSR(*SYSOPR)
```


<input type="checkbox"/>	From User	Type	Message
<input type="checkbox"/>	Stacyb	Information	Warning: CPUMonitor has exceeded threshold 60 for 1 interval(s); the current value is 71.71



# Documentation on Replacement Variables



- Replacement variables are documented in the Navigator **help**
- From the **Configure Metric** panel, click the question mark



## Parameters for operating system command

### Threshold trigger and Threshold reset

→ Parameters for operating system command

- Drill down:

Parameter	Passed Data
&DATE	Date
&ENDPOINT	Endpoint system name
&INTVL	Collection interval
&MON	Monitor name
&RDUR	Reset duration
&RVAL	Reset value
&TDUR	Trigger duration
&TIME	Time
&TVAL	Trigger value
&VAL	Current value (See note 4)

# System Monitor Summary



Set Monitor General Information

General  
Metrics  
**Summary**

**General**  
Name: iSystemMonitor  
Description: General Key metrics

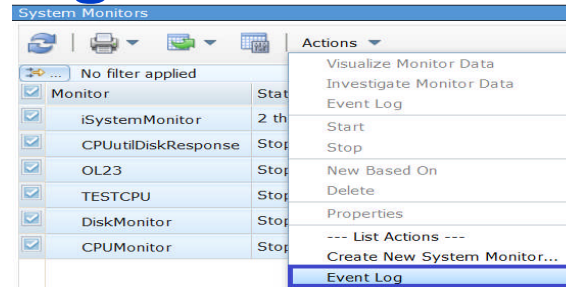

**Metrics**

Name	Trigger1	Reset1	Trigger2	Reset2
LAN Utilization (Average)	>= 40	< 0		
Disk Response Time (Read)	>= 10	< 9		
Disk Response Time (Write)	>= 1	< 1		
Disk Arm Utilization for System ASP (Average)	>= 30	< 29	>= 40	< 39
Disk Storage Utilization for System ASP (Average)	>= 70	< 0	>= 69	< 0
Machine Pool Faults Rate	>= 10	< 10		
CPU Utilization (Average)	>= 75	< 60	< 30	>= 50

OK Cancel

# Threshold Triggered - Event Logs

- Show events for one monitor or all monitors →



Event Logs

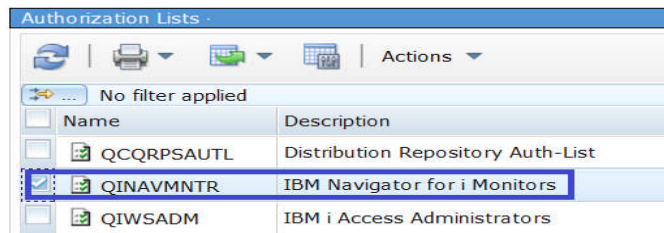
Owner: STACYB Metrics: LAN Utilization (Average),Disk Response Time (Read),Disk Response Time (Write),Disk Arm Utilization for System ASP (Average),Disk Storage Utilization for System ASP (Average),Machine Pool Faults Rate,CPU Utilization (Average) Monitor: iSystemMonitor

Event	Logged	Metrics	Monitor	Owner
✖ Trigger without command	2017-01-25 10:38:28.703	Disk Response Time (Write)	iSystemMonitor	STACYB
✔ Automatically reset	2017-01-25 10:51:58.689	Disk Response Time (Read)	iSystemMonitor	STACYB
✖ Trigger without command	2017-01-25 10:53:13.685	Disk Response Time (Read)	iSystemMonitor	STACYB
✔ Automatically reset	2017-01-25 10:55:58.689	Disk Response Time (Read)	iSystemMonitor	STACYB
✖ Trigger without command	2017-01-25 10:57:28.685	Disk Response Time (Read)	iSystemMonitor	STACYB
⚠ Trigger with command	2017-01-25 10:57:43.685	CPU Utilization (Average)	iSystemMonitor	STACYB
✔ Automatically reset	2017-01-25 10:59:58.685	Disk Response Time (Read)	iSystemMonitor	STACYB
✔ Automatically reset	2017-01-25 11:04:58.689	Disk Response Time (Read)	iSystemMonitor	STACYB
✖ Trigger without command	2017-01-25 11:06:28.688	Disk Response Time (Read)	iSystemMonitor	STACYB
✔ Automatically reset	2017-01-25 11:11:13.685	Disk Response Time (Read)	iSystemMonitor	STACYB
✖ Trigger without command	2017-01-25 11:12:28.689	Disk Response Time (Read)	iSystemMonitor	STACYB

- Trigger with no command run
- Trigger and a command was run
- Reset

# System Monitors - Authority

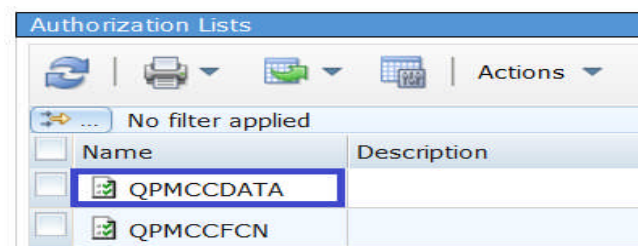
- To **create** and **start** a monitor, you need to:



Name	Description
QCQRPSAUTL	Distribution Repository Auth-List
<b>QINAVMNT</b>	<b>IBM Navigator for i Monitors</b>
QIWSADM	IBM i Access Administrators

- Have **\*ALLOBJ** special authority  
-OR-
- Be included on the **QINAVMNT** authorization list with **All** authority

- To **visualize** metrics, you need to :

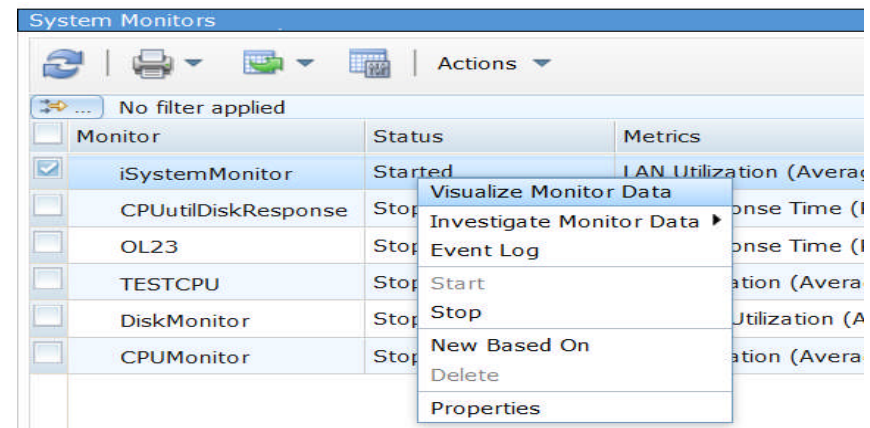


Name	Description
<b>QPMCCDATA</b>	
QPMCCFCN	

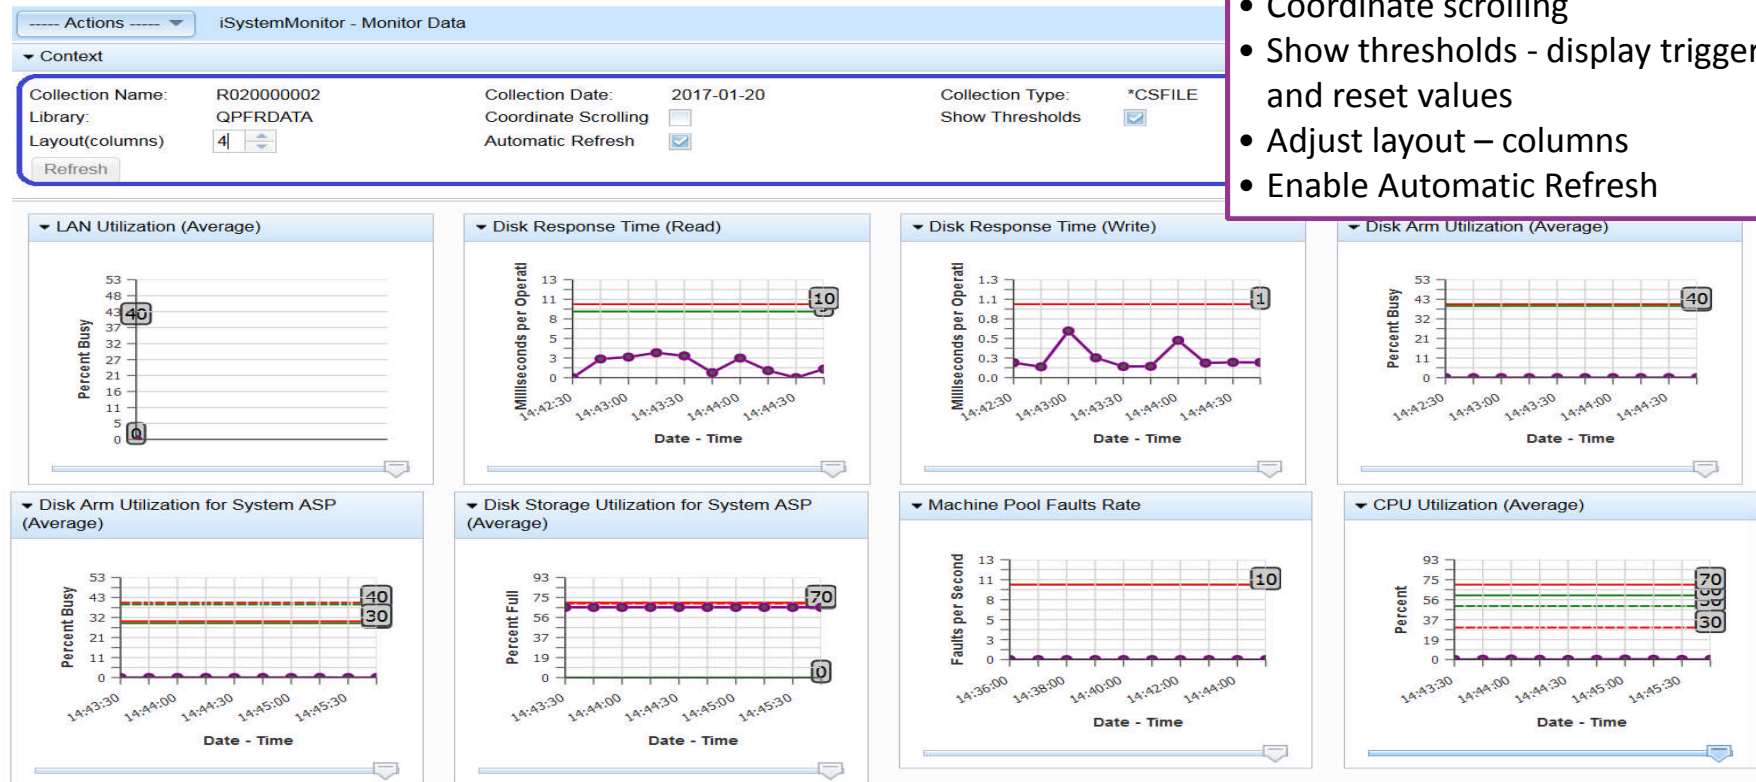
- Have **\*ALLOBJ** special authority  
-OR-
- Be included on the **QPMCCDATA** authorization list with **All** authority

# Visualize Monitor Data

- Real-time visualization of monitor data
- ★ ■ **Multiple** metrics graphed together
- Automatic refresh
- Coordinated scrolling across graphs
- Adjust graph layout



# Visualize Monitor Data - Customization



- Coordinate scrolling
- Show thresholds - display trigger and reset values
- Adjust layout – columns
- Enable Automatic Refresh

# Visualize Monitor Data – Drill down options

Actions available to get additional detailed charts for the metric selected:

The screenshot shows the 'iSystemMonitor - Monitor Data' window. A dropdown menu is open for 'CPU Utilization (Average)'. The menu items are:


- Historical Summary->CPU Utilization (Average)
- System Monitor->CPU Utilization (Average)
- System Monitor->CPU Utilization (Uncapped)
- System Monitor->CPU Utilization (SQL)
- System Monitor->CPU Utilization (Interactive Jobs)
- Collection Services->CPU Utilization Overview
- Collection Services->CPU Utilization by Thread or Task
- Collection Services->Resource Utilization Overview

Annotations with arrows point from these menu items to descriptive text on the right. A blue arrow points from the 'CPU Utilization (Average)' menu item to the 'CPU Utilization - Top 50 Contributors' chart below.

Historical Summary charts (Graph History 7.3)

Investigate Data charts showing System Monitor data

Investigate Data charts showing detailed Collection





# Best Practices for System Monitors

- ✓ 1 minute (or longer) for intervals
- ✓ Choose metrics important to your environment
- ✓ Clean up data regularly (R\* collections)
  - System monitor data retention setting
- ✓ Set thresholds (and be alerted) **prior** to negative effect
- ✓ Keep current on HTTP Group PTF
  - 7.2 SF99713
  - 7.3 SF99722

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20%20Technology%20Updates/page/PTF%20Groups>


- **NOTE:** For PTFs related to monitors, you should end the QINAVMNSRV job before applying:  
CALL PGM(QSYSDIR/QINAVMNSRV) PARM(\*STOP)



# Tuning

# CPU Considerations

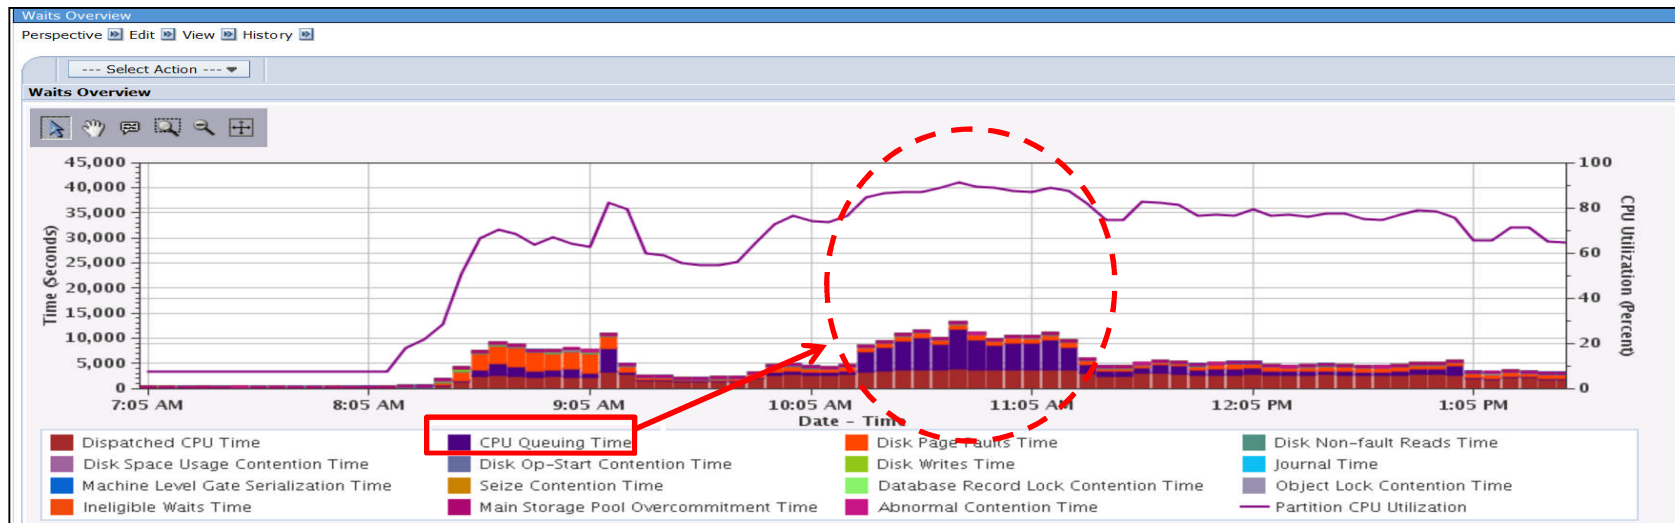
- CPU Utilization Guidelines depend upon the number of processors defined to the partition
- IBM Lab Services uses the following guidelines:
  - 70% 1-core
  - 76% 2-cores
  - 81% 4-cores
  - 85% 8-cores
  - 87% 12-cores
  - 91% 24-cores
  - 93% 32-cores



*for high-priority work, not considering lower-priority batch jobs*

# CPU Considerations – CPU Queuing

- Use Collection Services data and PDI to leverage wait accounting clues on how jobs are being affected because there is too much work for the processors and they have to wait = **CPU Queuing Wait time**




# Managing CPU

- Know which jobs are consuming the most CPU
  - Do they have business justification?

**CPU**


- CPU Utilization Overview
- Interactive Capacity CPU Utilization
- CPU Utilization by Job or Task
- CPU Utilization by Thread or Task
- CPU Utilization by Generic Job or Task



**Timeline**

- Timeline Overview for Jobs or Tasks

off peak

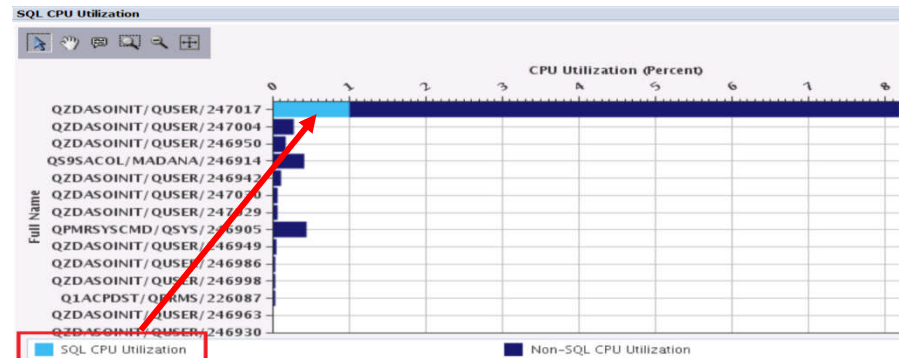


# Managing CPU

- Verify job priorities are in-line with business goals (high priority/low priority)
- Tune SQL

Database

- I/O Reads and Writes
- SQL CPU Utilization by Job or Task
- Database Locks Overview



- Consider using **Workload Groups** to limit CPU used by jobs or jobs within a subsystem →



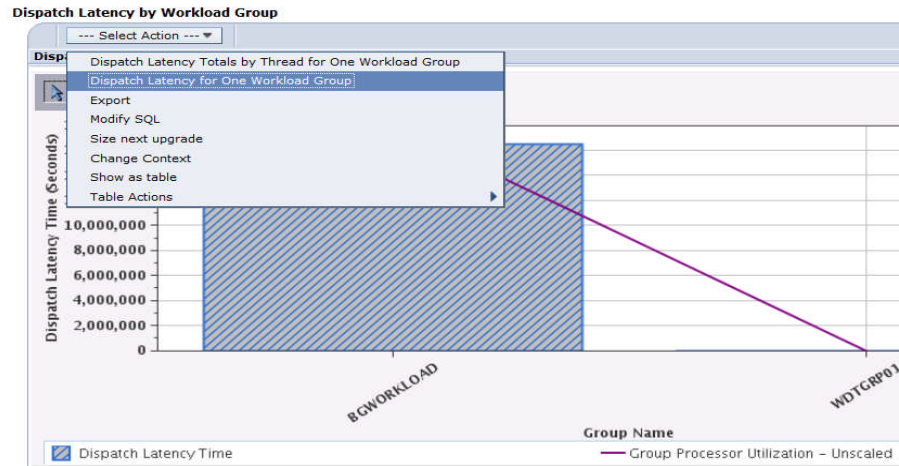

# Manage CPU with Workload Groups

- Allows you to **control the number of processor cores** that can be concurrently used by a workload
  
- Define workload groups
  - Assign the number of processor cores to the workload group
  - Associate a workload group with a job or a subsystem
  
- Example:
  - Control the resources used by the QZDASOINIT database server jobs that are in their own user-defined subsystem for ad-hoc queries

[https://www.ibm.com/support/knowledgecenter/en/ssw\\_ibm\\_i\\_72/cl/addwlcgrp.htm](https://www.ibm.com/support/knowledgecenter/en/ssw_ibm_i_72/cl/addwlcgrp.htm)  
<http://ibmsystemsmag.com/blogs/i-can/july-2016/workload-group-configuration-with-ibm-i-7-3>

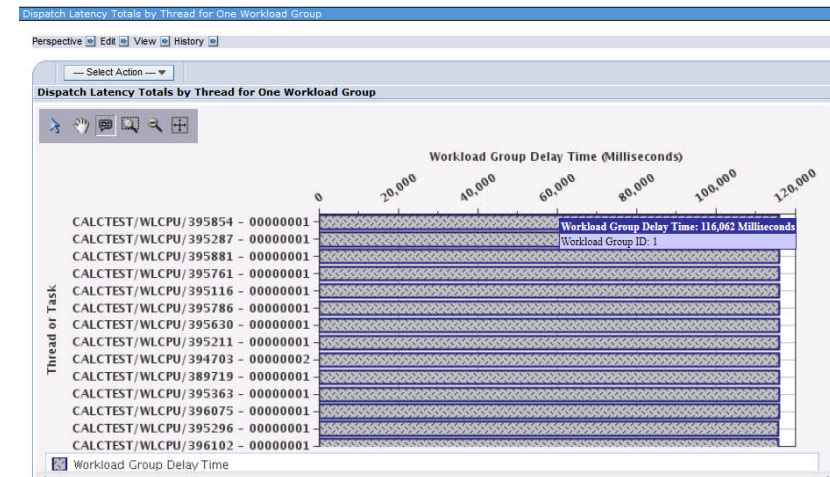
# Workload Group Monitoring

Collection Services 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.



# Memory Tuning

- Memory Faulting
  - Faulting is normal and expected
  - How much is too much?
    - ★ **The machine pool faulting rate should be less than 10/faults per second**
      - For the rest of the pools, it depends!!
  - A “fast” I/O subsystem will reduce the amount of time the page faults take
  - Put disparate workloads into separate pools
    - Java workloads, ad-hoc SQL, developers, known high faulting type workloads (replication software, RTVDSKINF)....
- Consider using the Performance Adjustor to automatically tune your pool sizes and activity levels (QPFRADJ =3 recommended)
  - Rapidly changing workloads are not well suited for the Performance Adjustor



# Memory Tuning – Work with Shared Pools

- “Tune the Tuner“ (WRKSHRPOOL / F11)
  - You can tune the values to influence the behavior of the performance adjuster
    - Focus on pool priority and setting the minimum and maximum values
    - Set according to business needs

```

Main storage size (M) . :      32476.00
Type changes (if allowed), press Enter.

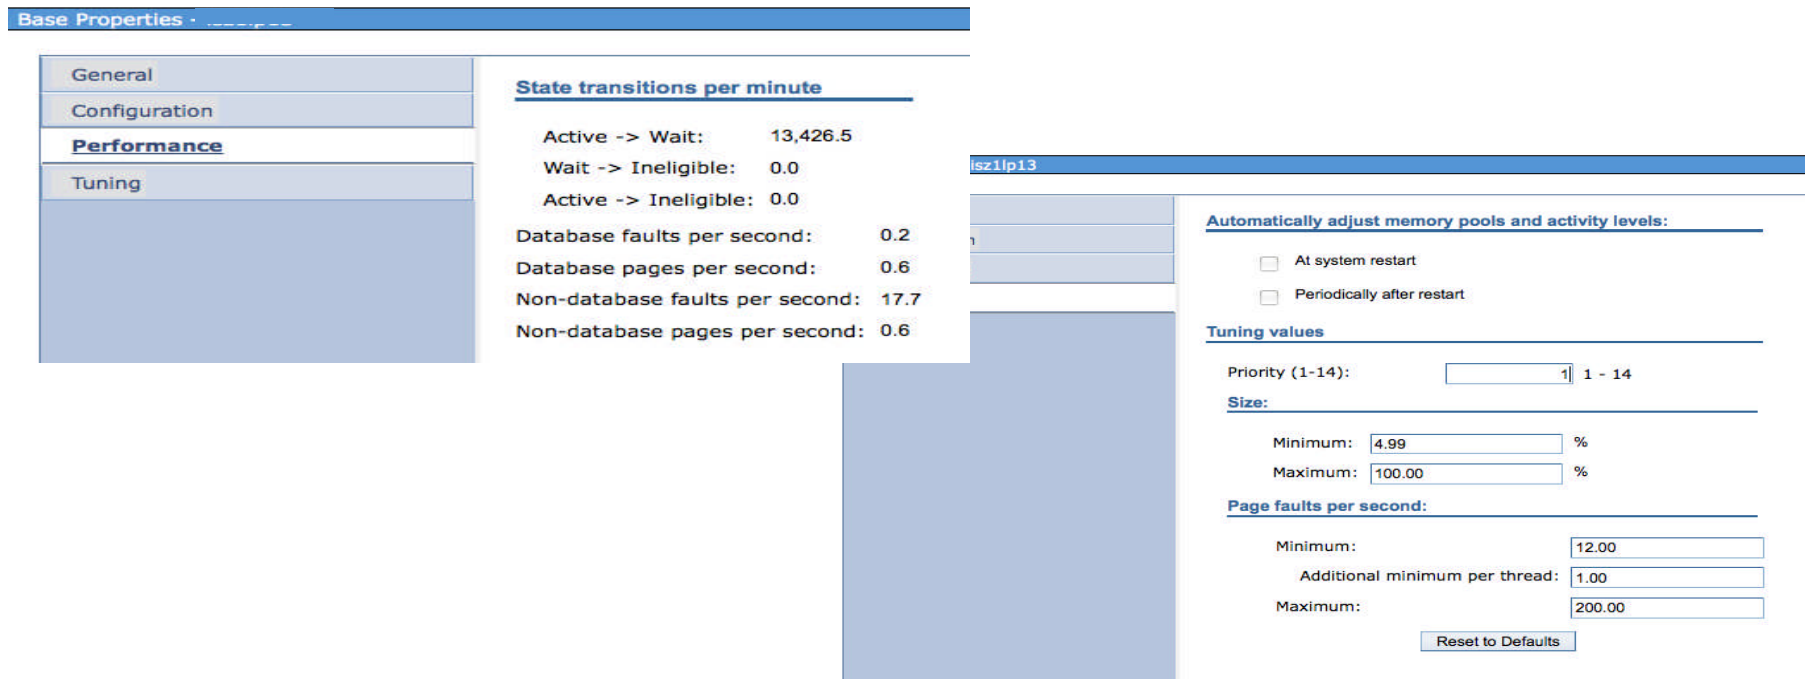
```

Pool	Priority	-----Size %-----		-----Faults/Second-----		
		Minimum	Maximum	Minimum	Thread	Maximum
*MACHINE	1	3.07	100	10.00	.00	10.00
*BASE	1	4.99	100	12.00	1.00	200
*INTERACT	1	10.00	100	12.00	1.00	200
*SPOOL	2	1.00	100	5.00	1.00	100
*SHRPOOL1	2	1.00	100	10.00	2.00	100
*SHRPOOL2	2	1.00	100	10.00	2.00	100
*SHRPOOL3	2	1.00	100	10.00	2.00	100
*SHRPOOL4	2	1.00	100	10.00	2.00	100
*SHRPOOL5	2	1.00	100	10.00	2.00	100
*SHRPOOL6	2	1.00	100	10.00	2.00	100

★ Machine Pool is the only pool that should have Priority of 1. Evaluate business priority of other pools.

# Memory Tuning – Pool Properties via Navigator

You can also do memory pool tuning through the GUI



**Base Properties**

- General
- Configuration
- Performance**
- Tuning

**State transitions per minute**

- Active -> Wait: 13,426.5
- Wait -> Ineligible: 0.0
- Active -> Ineligible: 0.0
- Database faults per second: 0.2
- Database pages per second: 0.6
- Non-database faults per second: 17.7
- Non-database pages per second: 0.6

**Automatically adjust memory pools and activity levels:**

- At system restart
- Periodically after restart

**Tuning values**

Priority (1-14):  1 - 14

**Size:**

- Minimum:  %
- Maximum:  %

**Page faults per second:**

- Minimum:
- Additional minimum per thread:
- Maximum:

## Faulting Considerations

- You can have really high faulting rates but very fast disk drives
- ...or fairly low faulting rates and very slow response time on your drives





### Use **Disk Page Fault Wait Time bucket** information to assess

- Use Collection Services with the Performance Data Investigator
- Waits Overview at a System level to assess overall impact of disk page fault wait time
  - Can drill down to see wait time at an individual job level
  - Assess impact
    - **General guideline:** limit the amount of time spent waiting on disk faults to 25% of the runtime of an average job on the system

# Disk Page Fault Wait time



- [-] Performance
  - [-] Investigate Data
    - [-] Investigate Data Search
    - [-] Health Indicators
    - [-] Collection Services
      - [-] CPU Utilization and Waits Overview
      - [-] CPU Utilization by Thread or Task
      - [-] Resource Utilization Overview
    - [-] Job Statistics Overviews
    - [-] Waits
      - [-] Waits Overview



Disk Page Faults Time is very large component system-wide. ← It is also a large component of where some jobs are spending their time – waiting on page faults...

# Expert Cache

- Expert cache dynamically adjusts the size and type of I/Os to maximize the use of main storage and minimize the number of disk I/Os
- Paging Option controls expert cache
  - **\*CALC** enables expert cache for that pool

 Turn Expert Cache on for Shared Pools

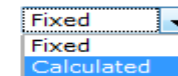
-Paging Option--

Defined Current


\*FIXED \*FIXED

\*FIXED \*FIXED

Paging option:



# Memory




- Memory perspectives are now available
- Similar information from what you get on WRKSYSSTS....

System Pool	Pool Size (M)	Reserved Size (M)	Max Active	-----DB-----	-----Non-DB-----		
				Fault	Pages	Fault	Pages
1	490.59	247.83	+++++	.0	.0	.0	.0
2	5344.71	6.07	149	.0	.0	.0	.0
3	2283.44	.00	203	.0	.0	12.3	29.0
4	.25	.00	5	.0	.0	.0	.0

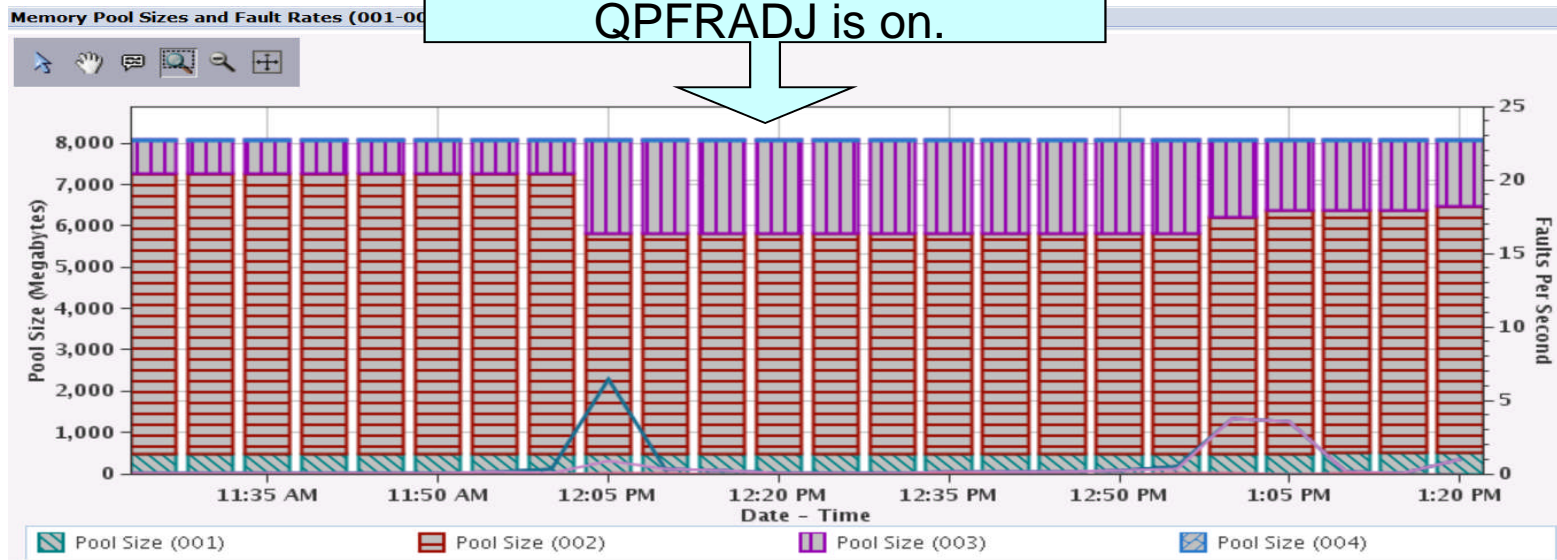


# Memory Charts



- *In a graphical view!*
- *Leverage PDI with CS data to expand your tuning tools*

Note the change in pool sizes.  
QPFRADJ is on.




# Memory - Drilldown

Memory Pool Sizes and Fault Rates (001-004)



--- Select Action ---

- Memory Metrics for One Pool
- Memory Pool Activity Levels
- DB and Non-DB Page Faults
- Page Faults by Job or Task**
- Waits by Pool
- Disk Waits Overview
- Memory Pools Health Indicators





# Disk

- Disk Response Time can be a significant component that determines how long a job will take to run




– General Guidelines for Hard-Disk-Drives (HDD) response time:

- Good: Average response time less than 5 milliseconds
- Normal: Average response time between 5 and 10 milliseconds
- Requires analysis: anything over 10 milliseconds



- ☰ Disk
  - ☰ Disk Response Time
    - ⊕ Detailed
    - Disk I/O Rates Overview
    - Disk I/O Rates Overview With Cache Statistics
    - Disk I/O Average Response Time Overview




<http://www.ibmssystemsmag.com/Blogs/i-Can/Archive/i-can-measure-disk-response-times/>


# Disk

- Understand root cause of I/O times:
  - Type:

- Waits
  - Waits Overview
  - Seizes and Locks Waits Overview
  - Contention Waits Overview
  - Disk Waits Overview



- Physical Disk I/O
  - Basic Views
    - Physical Disk I/O Overview - Basic
    - Physical Disk I/O by Job or Task - Basic



# Disk




- Disk Percent Busy
  - “Soft” Guideline = 40%
    - A disk may be >40% busy and still be able to perform additional disk operations
    - Acceptable response times can be achieved at up to 70%
    - IBM Workload Estimator uses a conservative guideline of 25%
    - OLTP environments are much more sensitive to variations in disk response time
    - Batch-oriented workloads may be able to tolerate higher disk utilizations (than 25%)
  - High disk utilization may trigger a need to inspect the disk response time
- Disk Storage
  - ASP percentage on WRKSYSSTS (% system ASP used) should be less than 80% .... Maybe.....
    - If you have a lot of disk storage, you may be able to run higher than 80%

# Analyzing Performance

# Why would you bother...?

## Go Perform

```
PERFORM IBM Performance for i
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
```



# Basic Performance Analysis


- Become familiar with your Collection Services performance data to understand your performance characteristics
  - PDI/Collection Services charts on every IBM i 6.1 or later, nothing to install!!!



Be proactive

- When a performance problem occurs you often need to use performance analysis tools to identify the cause of the problem to correct it

- Performance tasks in Navigator
  - Manage performance collections
  - Performance Data Investigator




## Analysis Reports

- Performance reports are useful for comparing your measurements over time to compare to your baseline
  - Historically, the Performance Tools LPP gave you the ability to generate text-based reports
  - You can create graphical reports in PDF format with PDI



- Create reports to track your key performance metrics over time



# PDI Report (PDF example)



Feb 28, 2013 10:03:43 AM

## Performance data report definition:

Demo Report

## Report title:

Example Report based upon COMMON performance collection

## Perspectives included in report:




CPU Utilization and Waits Overview

Page Faults Overview

Synchronous Disk I/O Overview

## Library/Collection used for report:

Common/Cs228229nd







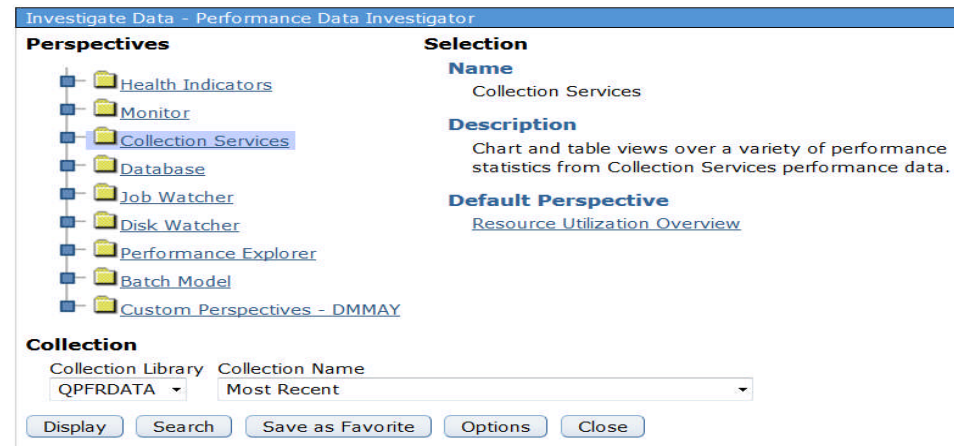
## Advanced Performance Analysis

- Advanced and detailed analysis is necessary to fully diagnose some performance problems
- IBM i has sophisticated data for this purpose
  - Job Watcher
  - Disk Watcher
  - Performance Explorer
- And two user interfaces:
  - Performance Data Investigator
  - iDoctor

## Analyze Performance Data


### ★ Start with Collection Services data

- Whether investigating a reported problem, monitoring how your system is running, or looking for improvement candidates, it can help
- Use it to understand resource usage and scope the problem:
  - What is or is not being used
  - How did it change?
  - When did it happen?
  - What is affected?



# Analyze Performance Data

- Where to start
  - Health Indicators perspectives - Are any resources constrained?
  - Collection Services perspectives:
    - **CPU Utilization and Waits Overview**
      - Any periods of high CPU or wait time?
      - Does any type of wait stand out?
    - **CPU utilization by Thread or Task**
      - See who is using CPU time
    - **Resource Utilization Overview**
      - Time based comparison of disk busy, physical I/O, faulting, CPU, logical DB I/O and 5250 transactions
      - Does anything stand out here?
- Use drill downs and other perspectives based on what you found above and want to investigate further



## Analyze Performance Data (continued)

- Using the Collection Services data you should know what you need to investigate further
  - excessive CPU consumption, locking issues, disk utilization, ...
  - when it happened
  - what resources were involved (jobs, disk units, ....)
- Collection Services *cannot* tell you details like what programs, instructions were running, objects, files, records, SQL statements being used, who is holding locks and who is waiting...



- You need to use tools like Job watcher, Disk Watcher and/or PEX for this in depth information.
  - Based on the Collection Services information you can focus these tools to collect only the data you need.
  - Generally, Job Watcher is used when you cannot identify the root cause with Collection Services

The background of the slide is a complex network of light gray lines connecting numerous small gray circular nodes. The nodes are scattered across the page, with a higher density in the center and lower density towards the edges. The lines are thin and create a web-like structure.

# Understand Performance Trends

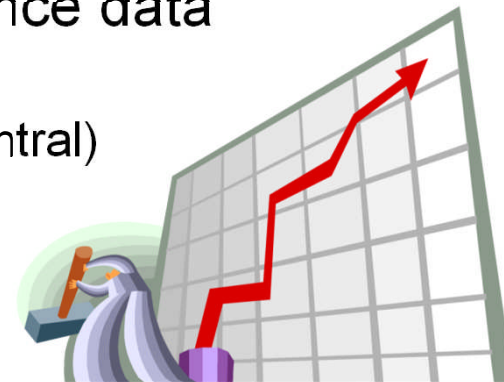
## Monitor Historical Performance Trends

Historical performance data allows you to

- Identify Trends
- Plan for the future

You have two options for historical performance data

- Graph History (7.3 Navigator for i, Management Central)
- PM for Power Systems





# Graph History

- **Graph History** allows you to view performance data over.....
  - Days, weeks, or months
  
- **Graph History with *System i Navigator* (7.1 and 7.2)**
  - If PM for Power Systems is not active, you can keep up to 7 days of detail data and 0 months of summary data.
  - If PM for Power Systems is active, then you can keep up to 30 days of detail data and 99 years of summary data.
  
- **Graph History is now available in *Navigator for i* with 7.3!**
  - If PM for Power Systems is not active, you can keep up to 7 days of detail data and 1 month of summary data.
  - If PM for Power Systems is active, then you can keep up to 60 days of detail data and 50 years of summary data.



# Historical Data

7.3



- There are two types of historical data:
  - **Summary**
    - System level or summarized metrics  
Useful in identifying trends or detecting changes in a system over a long period of time
  - **Detail**
    - Data from which the summarized metrics are derived and other relevant supplementary data.  
This data is used when looking deeper into a problem identified while looking at summary historical data.  
Only the top contributors for each metric will be stored as historical detail data.



# Enabling Historical Data

7.3



- [-] Performance
  - [+] Investigate Data
  - [+] Manage Collections
    - [-] Configure Collection Services
  - [+] Graph History

- Historical data is **OFF** by default
- Enable historical data in Collection Services Properties
  - Historical data is created at collection cycle time

**Configure Collection Services**

**General**

Library: QPFRDATA

Default collection interval:  15 seconds  5 minutes

**Cycling**

Cycle every day at: 12:00 AM Example: 12:30 PM

Cycle every: 24 hours

**System options**

- Enable system monitoring
- Create historical data when collection is cycled
- Create database files during collection
- Create performance summary data when collection is cycled

# Historical Data Retention

7.3



- Retention period for summary and detail data depends on PM Agent status
  - Radio buttons will show whether PM Agent status is \*ACTIVE or \*INACTIVE
  - Select PM Agent On or Off to Start or Stop

The screenshot shows the 'Configure Collection Services' interface. On the left is a navigation menu with 'Data Retention' selected. The main content area is divided into sections: 'Collection object' (30 days), 'Standard data' (10 days), and 'System monitor data' (2 days). The 'Historical data' section is highlighted with a purple rounded rectangle and contains two columns of settings. The left column is for 'PM Agent on' (radio button unselected) and the right column is for 'PM Agent off' (radio button selected). Under 'PM Agent on', 'Save summary data for:' is 10 year(s) and 'Save detail data for:' is 60 days (1-60). Under 'PM Agent off', 'Save summary data for:' is 1 month and 'Save detail data for:' is 7 days (1-7). A 'View disclaimer' button is at the bottom left of the highlighted area.

# Collection Services - Historical Data Properties

7.3



- Interval  
Summary and detail historical data will be saved at this interval.  
The default is 60 minutes
- Create historical detail data  
The default is to create historical detail data when summary historical

The screenshot shows the 'Configure Collection Services' interface. On the left is a navigation menu with options: General, Data to Collect, Data Retention, System Monitor Categories, and **Historical Data**. The main content area is titled 'Historical Data' and contains the following settings:

- Historical summary data**
  - Library: QPFRHIST
  - Interval: 30 Minutes (highlighted with a purple circle)
- Historical detail data**
  - Create historical detail data
  - Filter: ALL (dropdown menu) Save this many top contributors of each detailed metric

# Completing Historical Data Configuration

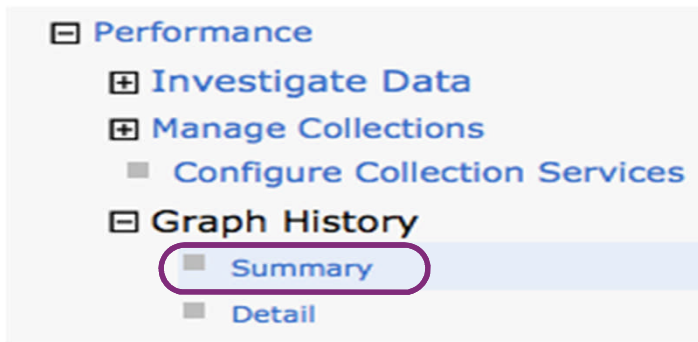
7.3



- After you enable historical data, you must cycle Collection Services to create the historical data
- Historical data will be created for all the existing management collection objects in the configured library




## Viewing Historical Data - 7.3




# CPU Utilization for the past month

7.3



# Temporary Storage for the past month



# Context

7.3



- Context is initially collapsed
- Expand it to select the metric and timeframe you wish to view

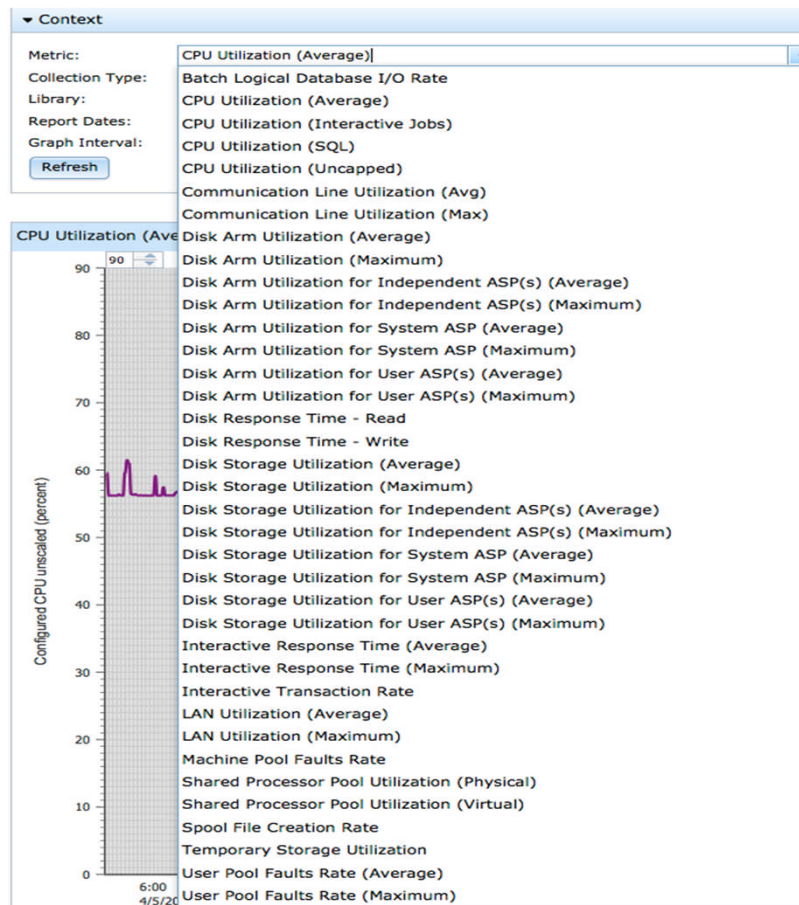
Context

Metric:	CPU Utilization (Average)	
Collection Type:	*HSTFILE	From Date: 4/4/2016
Library:	QPFRHIST	Time: 00:00
Report Dates:	1 month	To Date: 5/4/2016
Graph Interval:	1 hour	Time: 13:46

Refresh




# Metrics

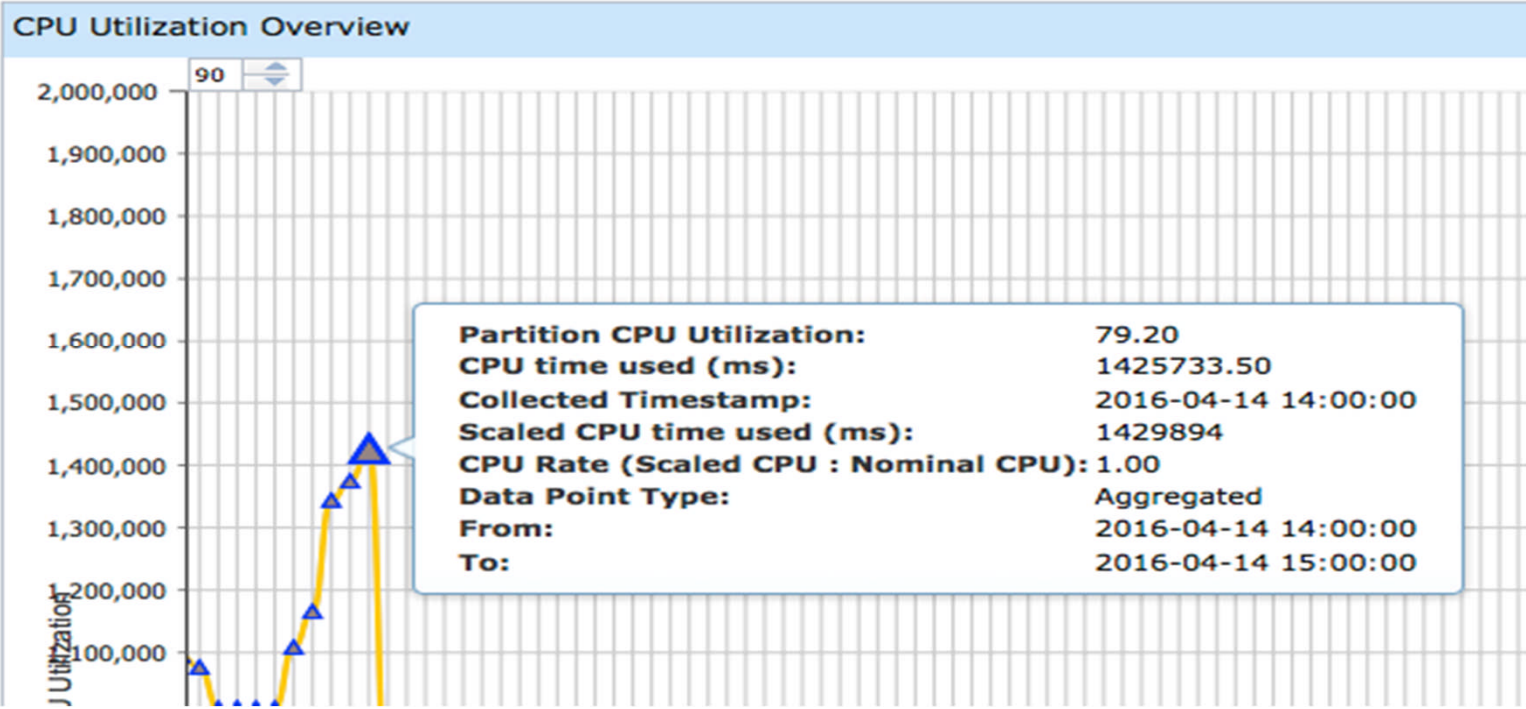


# Zoom in for Details

7.3



# Hover Over Data Points for more Information



# Historical Data - Detail



- Performance
  - Investigate Data
  - Manage Collections
  - Configure Collection Services
  - Graph History
    - Summary
    - Detail

Welcome
Dashboard
Detail

----- Actions -----
CPU Utilization Overview

Context

**CPU Utilization Overview**

Reset

Partition CPU Utilization     CPU time used (ms)

From: 4/14/2016 5:00; To: 4/17/2016 9:00

**Top Contributors (04/15/2016 13:00:00 - 04/15/2016 14:00:00)**

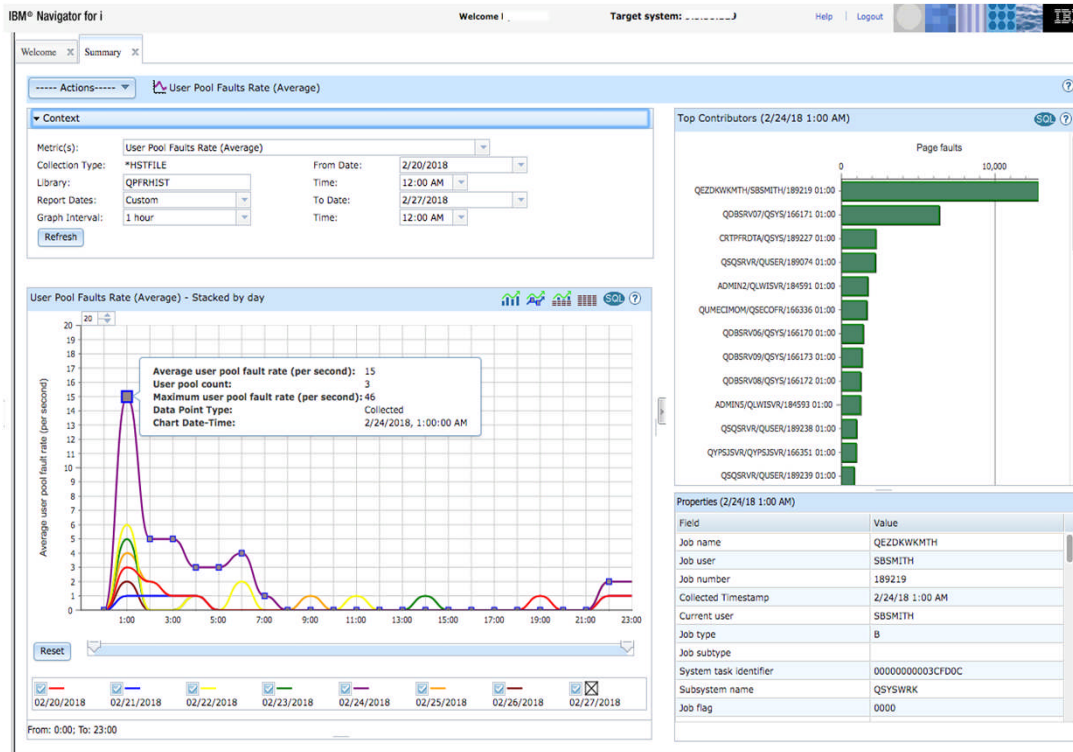
Unscaled CPU time charged (ms)

Job Name	Time	Unscaled CPU time charged (ms)
QSLPSVR/QSYS/005871	13:30	~1,850,000
QTSMTPSRVD/QTCP/005872	13:30	~1,800,000
QSLPSVR/QSYS/005871	13:00	~1,750,000
QZRCRSVS/QUSER/041019	13:30	~1,650,000
QZDASOINIT/QUSER/040956	13:30	~1,600,000
QYPSJSVR/QYPSJSVR/005822	13:30	~1,550,000
QZRCRSVS/QUSER/041019	13:00	~1,500,000
QZDASOINIT/QUSER/040956	13:00	~1,450,000
QTMSSMTPD/QTCP/005880	13:30	~1,400,000
QUMEPRVAGT/QSECOFR/042130	13:00	~1,350,000
QTMSSMTPD/QTCP/005880	13:00	~1,300,000
QTSMTPLD/QTCP/005881	13:30	~1,250,000
QTSMTPLD/QTCP/005881	13:00	~1,200,000
QUMEPRVAGT/QSECOFR/042131	13:00	~1,150,000
QUMEPRVAGT/QSECOFR/042129	13:00	~1,100,000
QUMEPRVAGT/QSECOFR/042127	13:00	~1,050,000
QUMEPRVAGT/QSECOFR/042128	13:00	~1,000,000
ADMIN/QTMHHTTP/040972	13:30	~950,000

**Properties (04/15/2016 13:30:00)**

Field	Value
Job name	QSLPSVR
Job user	QSYS
Job number	005871
Collected Timestamp	2016-04-15 13:30:00.000000
Current user	QSYS
Job type	B
Job subtype	
System task identifier	000000000000087D
Subsystem name	QSYSWRK
Job flag	0000
Server type	QIBM_SLP_SERVER
Job priority	010

# Historical Data – New Stacked Chart Support



<http://ibmsystemsmag.com/blogs/i-can/may-2017/graph-history-%E2%80%93-stacked-charts/>



# PM for Power Systems

★ Performance Management for Power Systems is an IBM offering that provides a historical view of performance

[https://www.ibm.com/support/knowledgecenter/en/ssw\\_ibm\\_i\\_72/rzahx/rzahxplangrow1pm1.htm](https://www.ibm.com/support/knowledgecenter/en/ssw_ibm_i_72/rzahx/rzahxplangrow1pm1.htm)

- Send your performance data to IBM
- IBM will store up to two years of performance data
  - You can use the supplied graphs to view the trends of key performance metrics
- Performance data sent to IBM can later be used to size your next upgrade

Report samples



# Plan for the Future



## **Strategy**

**1. Understand where your (potential) bottlenecks are today**

**2. Capacity Planning**





## Why know your what your workload is waiting on?

- Helps you know where to focus your infrastructure and investment
  - Hardware, software, etc.
  
- Understand how or if additional or improved resources will help
  - Invest only in resources that will benefit your environment
  - May differ by workload

## What is your workload waiting on today?


- Hardware (may indicate unbalanced system)
  - Processor bound
  - I/O bound
  - Memory bound
  
- Software
  - Application
  - Operating system
  
- Something else
  - Network



# Processor Bound System




Dispatched CPU  
+  
CPU Queuing




# I/O Bound System

Disk Time



# Lock Contention Time Bottleneck


Lock Contention wait time




# Wait Accounting at a Job Level




- Would this job benefit from additional memory?



- Would this job benefit from additional memory? CPU? Disk?



- Would this job benefit from an improved I/O subsystem?



# Capacity Planning Strategy

- Review current performance
- Select representative data
  - Is the data variable due to business cycle or seasonal economic conditions?
    - Identify seasonal processing peaks
  - Are there peaks within peaks?
    - Eliminate exceptional and unusual conditions
- Estimation
  - Gather realistic projections of the business factors that affect performance
  - Focus on how the business factors will cause changes in the workload
  - Account for pent up demand
  - Using a sizing tool such as the Workload Estimator ->

## Capacity Planning - WLE



- IBM Systems Workload Estimator (WLE) is the IBM tool for capacity planning
  - Submit input from the PM for Power Systems data
  - Select workload to size from the Performance Data Investigator
  - Manual input
  
- **Use WLE to:**
  - ✓ **Size a new system**
  - ✓ **Size an upgrade**
  - ✓ **Size a consolidation of several systems**
  
- WLE provides current and growth recommendations for
  - Processor
  - Memory
  - Disk

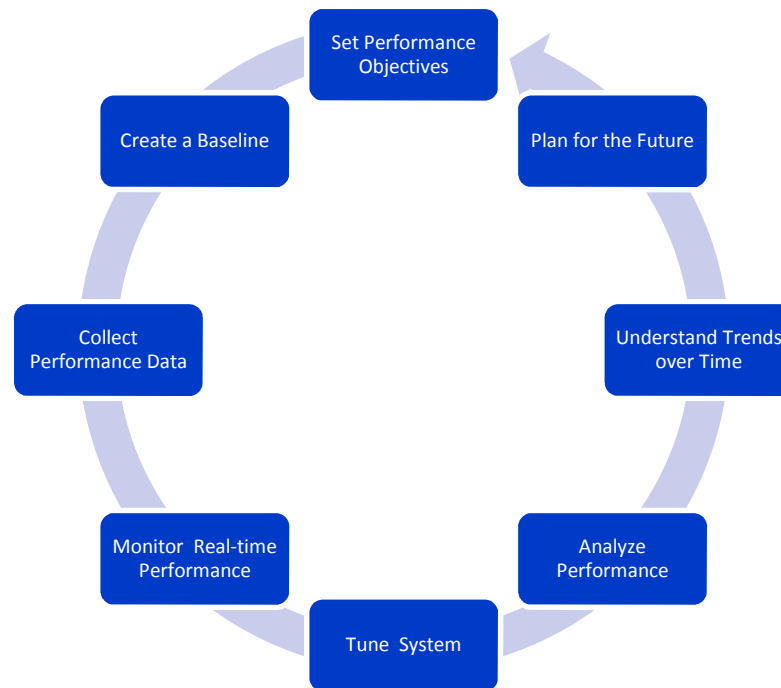


## A final thought on Best Practices....

- It's always good practice to keep current on the latest fixes from IBM
  - PTFs address **defects**
  - PTFs introduce **new capabilities**
  - ✓ IBM i Technology Refresh Updates
  - ✓ IBM i Group PTFs
    - Database
    - Performance tools
    - Java
    - HTTP Server
      - HTTP Server Group PTF for latest Navigator for i functionality
  - ✓ Individual PTFs for performance data collectors
    - Collection Services, Job Watcher, Disk Watcher, Performance Explorer



# Performance Management Life-cycle





# References



## **IBM i Performance FAQ a MUST read!**

October 2017 update (*watch for a Spring 2018 soon!*):

<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=POW03102USEN>

IBM Power Systems Performance



## **IBM i on Power - Performance FAQ**

*October 9, 2017*

## IBM i Web Sites with Performance Information

- IBM Knowledge Center:
  - [7.2 Performance](#)
  - [7.3 Performance](#)
- IBM i Performance Management:
  - [i Performance Management](#)
- developerWorks:
  - IBM i Performance Tools: [developerWorks Performance Tools](#)
  - IBM i Performance Data Investigator: [developerWorks PDI](#)
- IBM iDoctor for IBM i: [iDoctor](#)
- IBM i Wait Accounting information:
  - [Job Waits Whitepaper](#)
  - [KnowledgeCenter: The basics of Wait Accounting](#)
  - [developerWorks: IBM i Wait Accounting](#)

A **Redbooks** publication!



# End to End Performance Management on IBM i

Understand the cycle of Performance Management

Maximize performance using the new graphical interface on V6.1

Learn tips and best practices



Hernando Bedoya  
Mark Roy  
Nandoo Neerukonda  
Petri Nuutinen

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247808.pdf>

[ibm.com/redbooks](http://ibm.com/redbooks)

**Redbooks**

# IBM i 7.2 Technology Refresh Updates



Covers the 7.2 content through  
Technology Refresh 1

Section 2.8 – Performance

Section 8.6.7 – Job level SQL stats in  
Collection Services

Draft Document for Review December 10, 2014 2:51 pm



## IBM i 7.2 Technical Overview with Technology Refresh Updates

- Covers new functions and enhancements through IBM i 7.2 TR1
- Easy to use web-based system management
- Integrated Data-Centric approach



# IBM i Performance and Optimization Services



The IBM i Performance and Optimization team specializes in resolving a wide variety of performance problems. Our team of experts can help you tune your partition and applications, including:

- Reducing batch processing times
- Resolving SQL query and native IO performance problems
- Tuning RPG, COBOL, C, and Java (including WebSphere Application Server) programs
- Removing bottlenecks, resolving intermittent issues
- Resolving memory leaks, temporary storage growth problems, etc.
- Tuning memory pools, disk subsystems, system values, and LPAR settings for best performance
- Optimizing Solid State Drive (SSD) performance
- Tuning client interfaces such as ODBC, JDBC, .Net and more

Skills transfer and training for performance tools and analysis also available!

Contact Eric Barsness at [ericbar@us.ibm.com](mailto:ericbar@us.ibm.com) for more details.

[www.ibm.com/systems/services/labservices](http://www.ibm.com/systems/services/labservices)



# IBM i Performance Analysis Workshop

*Learn the science and art of performance analysis, methodology and problem solving*


Managing and analyzing the data can be quite complex. During this workshop, the IBM Systems Lab Services IBM i team will share useful techniques for analyzing performance data on key IBM i resources, and will cover strategies for solving performance problems. It will aid in building a future foundation of performance methodology you can apply in your environment.

## Overview:

- Topics covered include:
  - Key performance analysis concepts
  - Performance tools
  - Performance data collectors (Collection Services, Job Watcher, Disk Watcher, and Performance Explorer)
  - Wait accounting
- Core methodology and analysis of:
  - Locks
  - Memory
  - I/O subsystem
  - CPU
- Concept reinforcement through case studies and lab exercises
- Discussions on theory, problem solving, prevention and best practices

## Workshop details:

- Intermediate IBM i skill level
- 3-4 day workshop, public or private (on-site)
  - For general public workshop availability and enrollment:  
[IBM i Performance Analysis Workshop](#)
  - For public workshop availability and enrollment in France, please contact Philippe Bourgeois at [pbourgeois@fr.ibm.com](mailto:pbourgeois@fr.ibm.com) or Françoise Laurens at [f\\_laurens@fr.ibm.com](mailto:f_laurens@fr.ibm.com)
  - For additional information, including private workshops, please contact Eric Barsness at [ericbar@us.ibm.com](mailto:ericbar@us.ibm.com) or Stacy Benfield at [stacylb@us.ibm.com](mailto:stacylb@us.ibm.com), members of Systems Lab Services



***And finally.....***





**Thank you**

**Don't forget to fill-in the  
feedback form!**



**ithankyou**

[www.ibm.com/power/i](http://www.ibm.com/power/i)



**End of Presentation material.....**



**i**thankyou

[www.ibm.com/power/i](http://www.ibm.com/power/i)