Best Practices for IBM i Performance Management

Dawn May
dmmay@us.ibm.com

Performance Disclaimer

"it depends …"

Performance information and recommendations in this presentation are based on measurements, analysis, and projections in a controlled environment for specific performance workloads.

Your results may vary significantly and are dependent on the application and configuration.

This information is provided along with general recommendations for you to better understand system performance.

Information is provided *AS IS* without warranty of any kind.

Performance Management Life-cycle

Set Performance Objectives and Create a Baseline

Collect Performance Data

Real-time Monitoring

Analyze Performance Data

Tuning

Historical Trending

Capacity Planning

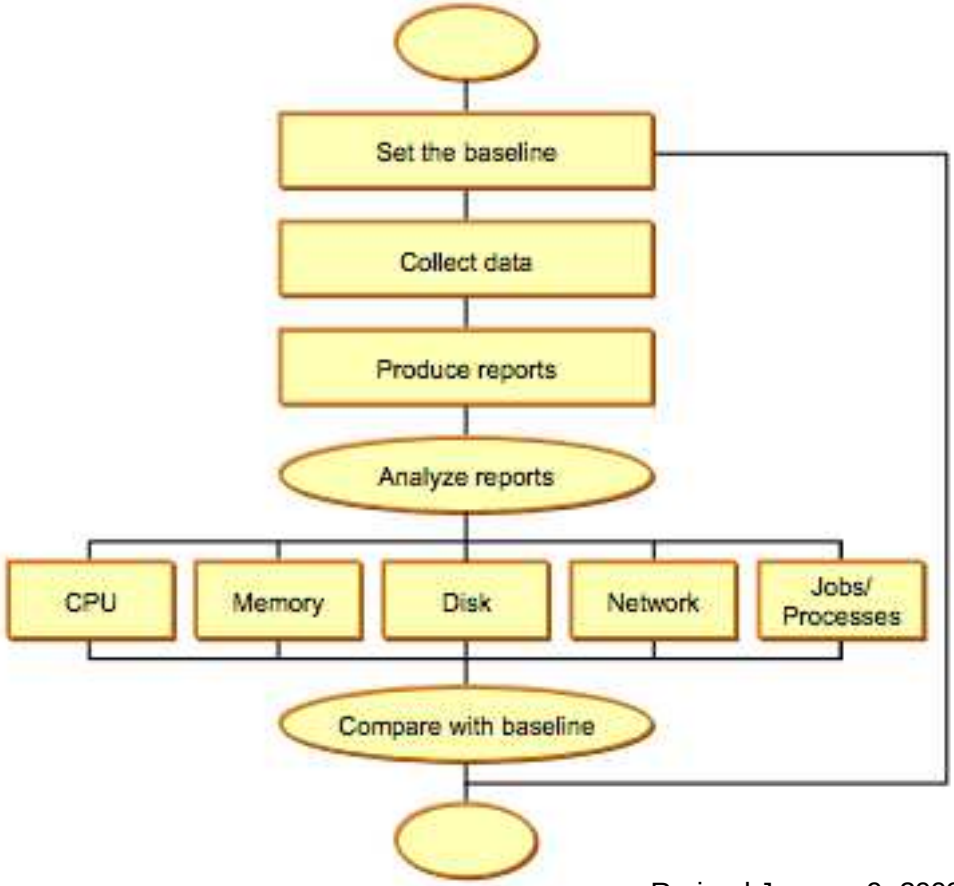

Revised January 9, 2003

Factors that Affect Performance

User expectations

Hardware capabilities and configuration

Software

Workload

Network

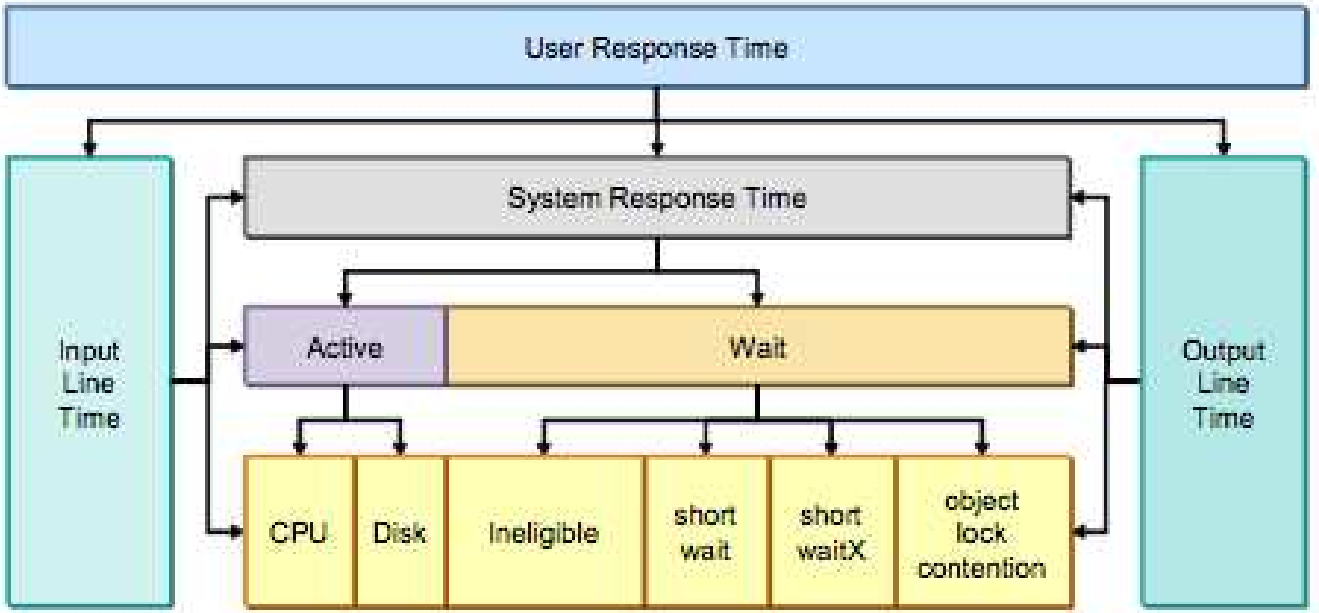

Power your planet.

Definitions

Interactive work – Generally 5250 online transaction processing (OLTP)

Batch work – non-interactive workloads

Commercial Processing Workload (CPW) – Workloads which have a relatively large amount of I/O compared to computation

Disk arms – generally one disk arm per disk drive today. More disk arms generally results in improved I/O performance

Disk capacity – the amount of disk storage space

Memory paging and faulting – the movement of data in and out of memory

Wait Accounting – the ability to determine what a job is doing when it is not running

Definitions...

Measurement: The collection of performance metrics

Transaction: A basic unit of work

Workload: An application that can drive load on a system

Benchmark: A specific workload with specific environment settings

Response Time: The average observed time to complete a transaction

Utilization: The percent of time that a resource is busy

Throughput: The rate at which transactions are completed

Capacity: The maximum throughput of a system

What is Wait Accounting?

Wait accounting is the ability to determine what a job is doing when it is not running

- When a job is not running, it is waiting
- Waits may be normal, but some waits are not normal
- Wait accounting can be used to determine if a wait condition is a problem

IBM i has instrumented most of the wait conditions

- Wait information is collected by Collection Services and Job Watcher

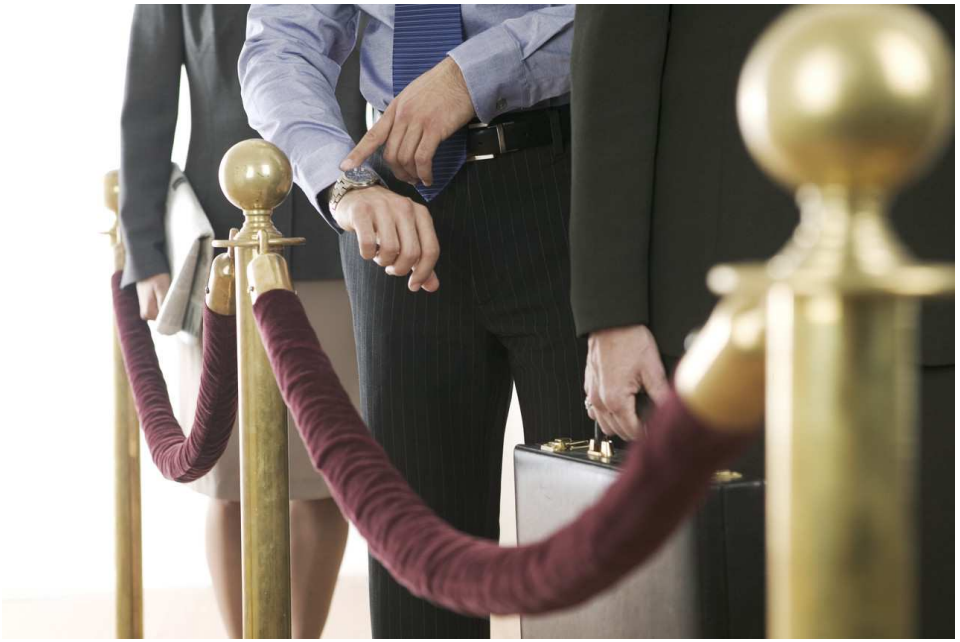

i Exclusive!

Basics of Waiting

Two basic types of waits

Idle: waiting for a work request

Typically not indicative of a problem

Waiting for the “Enter” key to be pressed on a 5250 display session

If a problem, usually external to the machine

E.g., slow arrival of work requests due to communications problem

Possible, but not typical in batch jobs

E.g., waiting for an entry to be placed on a data queue

Blocked: waits that occur while performing a work request

Blocked waits are the ones we want to take a closer look

“Outside of CPU usage and CPU queuing time, blocked waits are t
reason jobs/threads take as long as they do to complete their worl

Create a Baseline

A baseline is the expected performance characteristics over a defined period of time

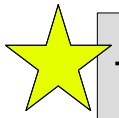

The baseline provides the data to compare with and identify changes that occur

You may need multiple baselines due to varying business periods

- Day-to-day operations
- Month-end
- Year-end

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 to create the baseline

The baseline can be simple or complex, depending upon the needs

Identify the metrics that are key to optimal performance

- Average response time
- Batch window
- Average CPU utilization
- Disk used percentage
- Number of interactive transactions

Patterns in Performance Data

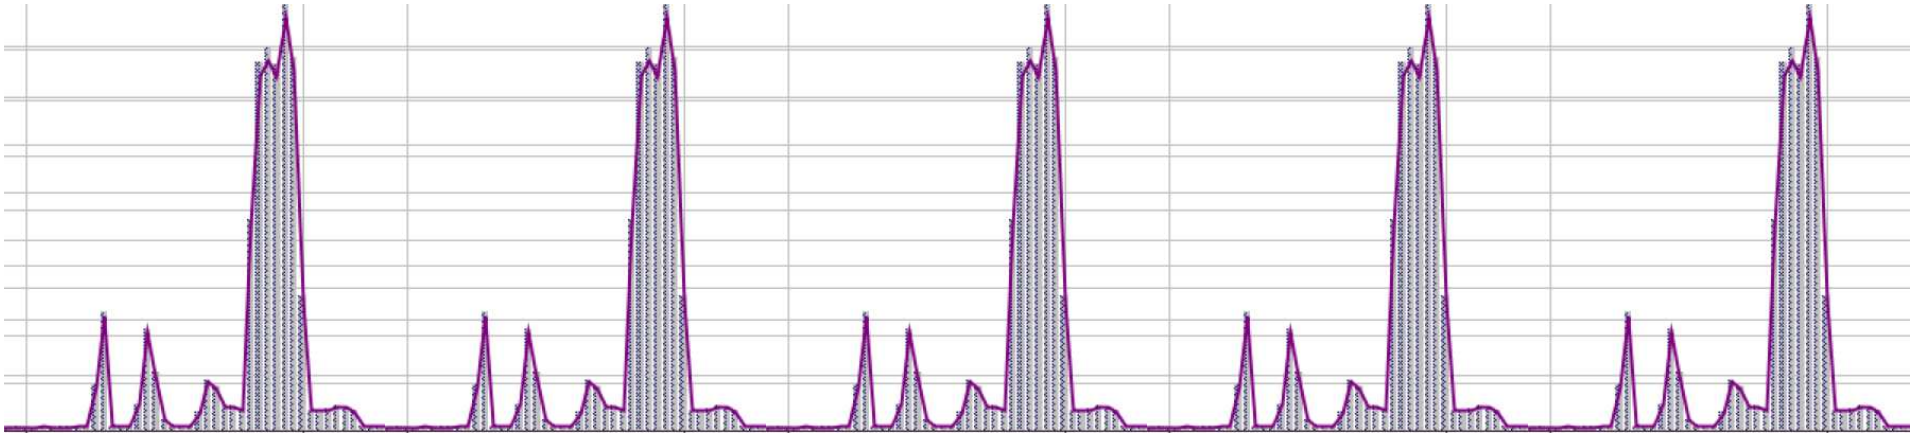

Performance data typically has patterns

- Understand your typical patterns
- Recognize change

Collect Performance Data

Collection Services

Job Watcher

Disk Watcher

Performance Explorer

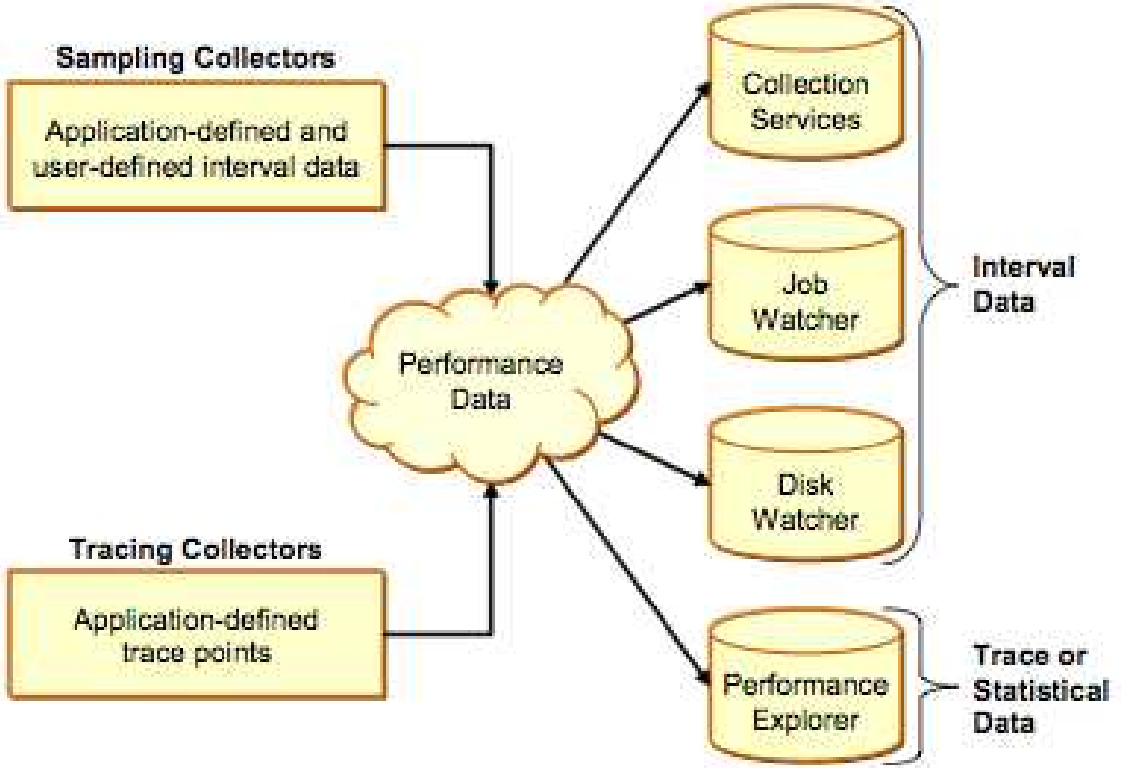

Power your planet.

Collection Services

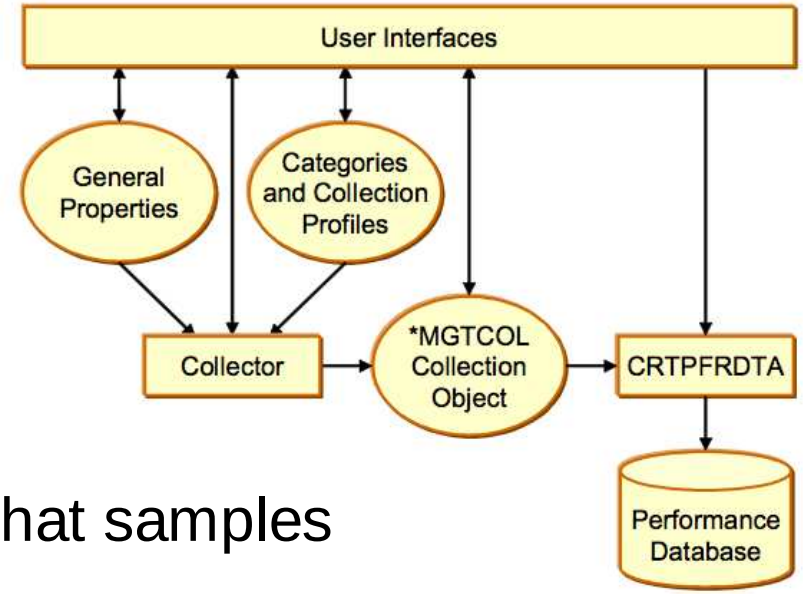

Collection Services is an IBM i function that samples system and job level performance data

IBM recommends you always run Collection Services

- Collects performance data 24x7 with minimal overhead
- Collects data from many system resources
- Collects data at regular intervals from 15 seconds to 1 hour
  - Data initially stored in a a management collection object
    - Can hold large quantities of performance data with minimal overhead
- Performance data is copied into database files

Job Watcher

Job Watcher returns real-time information about a selected set of jobs, threads, or LIC tasks

Job Watcher collects different types of data than Collection Services, and more frequent intervals

Job Watcher has more overhead than Collection Services

Data collected by Job Watcher includes

Wait times

CPU

I/O activity

Call Stacks

SQL statements

Communications statistics

Run Job Watcher when you need detailed performance data for diagnostic purposes.

There are clients that run Job Watcher 24x7
to always have diagnostic data available.

Need to manage the data carefully.

Disk Watcher

Disk Watcher is a performance data collector for detailed disk performance

Near real-time diagnosis of disk performance issues

V5R3 - SI24918

V5R4 - SI24919

Statistical and Trace data

Data is written to DB2 files

Run Disk Watcher when you need detailed performance data for disk I/O analysis.

Performance Explorer

Performance Explorer help identify the causes of performance problems that cannot be resolved using one of the other performance data collectors

- Collects more detailed information about a specific application, program, or resource

Performance Explorer is typically used for two main reasons:

- Detailed performance trace data is needed to identify the performance problem
- Analyzing the performance of applications

Performance Data Collectors - Positioning

Collection Services	Disk Watcher	Job Watcher	Performance Explorer (PEX)
Sample data Collected every 15 minutes (default) No information concerning specific I/O operations	•Statistics as well as Trace data •Focus on disk data	Sample data Support for very small intervals Focus on job data Call Stacks SQL Statements Wait buckets	Trace data Information collected for every I/O event Collection and analysis complex

Performance Data Management

Collection Services will automatically expire data – but you should keep important collections for comparison

Past seven to ten days

Key time periods & events, such as end-of-month or end-of-year processing, baseline prior to installing a new release

Store data centrally if you have multiple physical or logical partitions

Easier to analyze and backup

Resource intensive analysis won't impact production partitions

Keep some data online

Respond quickly to reports of poor or degrading performance

Can be on partition it was created on or centralized partition

Back up key performance data as you would business data

IBM i Performance Tools

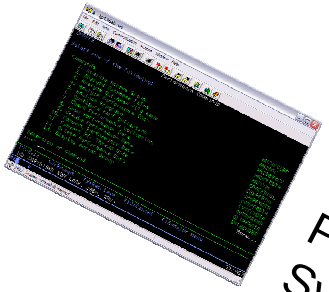

Green Screen
Performance Tools Reports
System commands
Batch Model
Performance Explorer

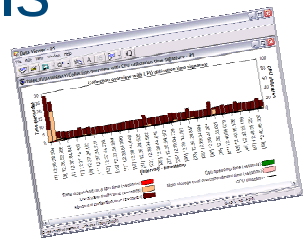

iDoctor Suite
Job Watcher
PEX Analyzer
Heap Analyzer
PDV

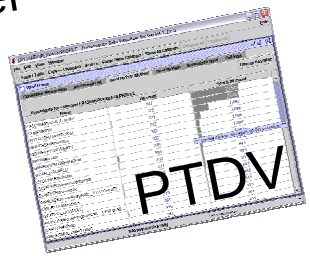

Collection Services

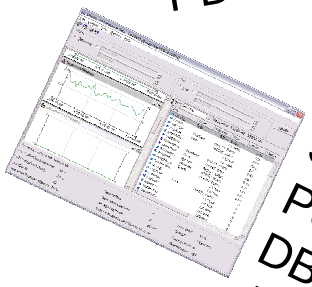

System i Navigator
Perf Tools GUI
DB/SQL Monitor
Active Jobs
System Status

Job Watcher

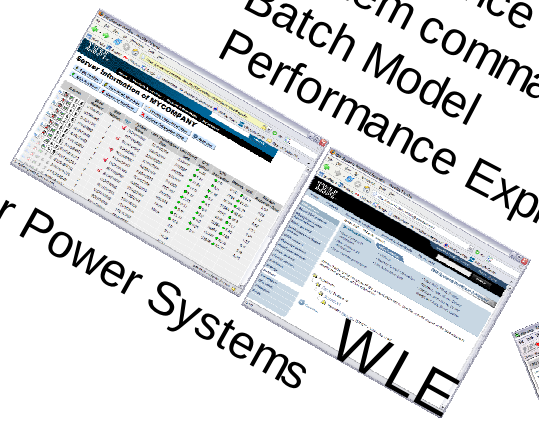

PM for Power Systems

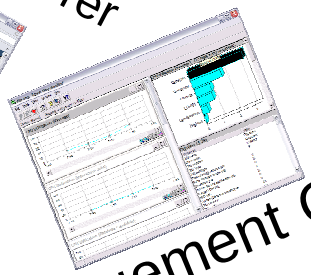

Management Central

Third Party Tools

Disk Watcher

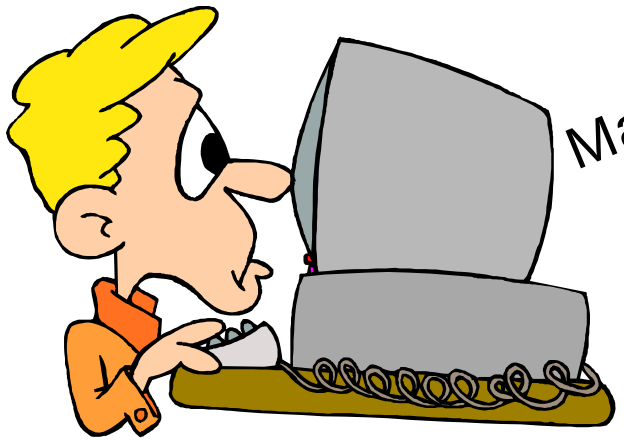

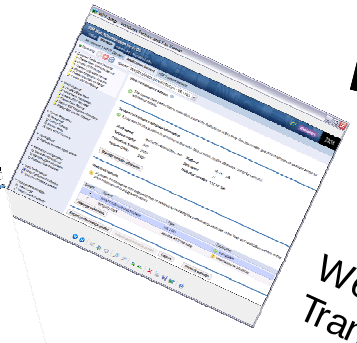

Web Performance Advisor
Transaction Performance

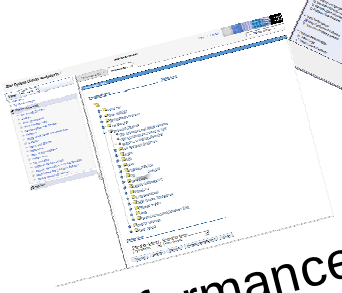

Performance Data Investigator

Performance Explorer

Power your planet.

Quick Review of Performance Tools Interfaces

5250 session
- CL Commands
- Performance menus and reports

System i Navigator
- Windows client application - aka, iSeries Navigator, Operations Navigator, Navigator for i
  - Collection Services
  - Management Central Monitors
  - Monitors – Real-time Monitoring
  - Graph History – Observer performance metrics over several months
  - Performance Tools plug-in – Basic trend and performance analysis

IBM Systems Director Navigator for i
- “Web Console” for much of what was originally in System i Navigator
- Performance Tasks: Performance Data Investigator, Collections, other performance tasks
- Performance Data Investigator:
  - Health Indicators – High level overview of key resources utilization metrics
  - Collection Services – Basic trend and performance analysis
  - Job Watcher – Job-related performance analysis
  - Performance Explorer – Trace the flow of work in an application

Tools and Usage

IBM Systems Director

- Web-based interface
- Support for multiple systems and groups
- System health and high-level monitoring
- Active Energy Manager extension

PM for Power Systems

- Web-based interface
- Current and long-term trend analysis

Workload Estimator

- Web-based interface
- Used to size a new partition, an upgrade, or a consolidation of several partitions

iDoctor suite

- Windows client application
- Detailed performance analysis and diagnostics

Relationship Between Collectors And Consumers

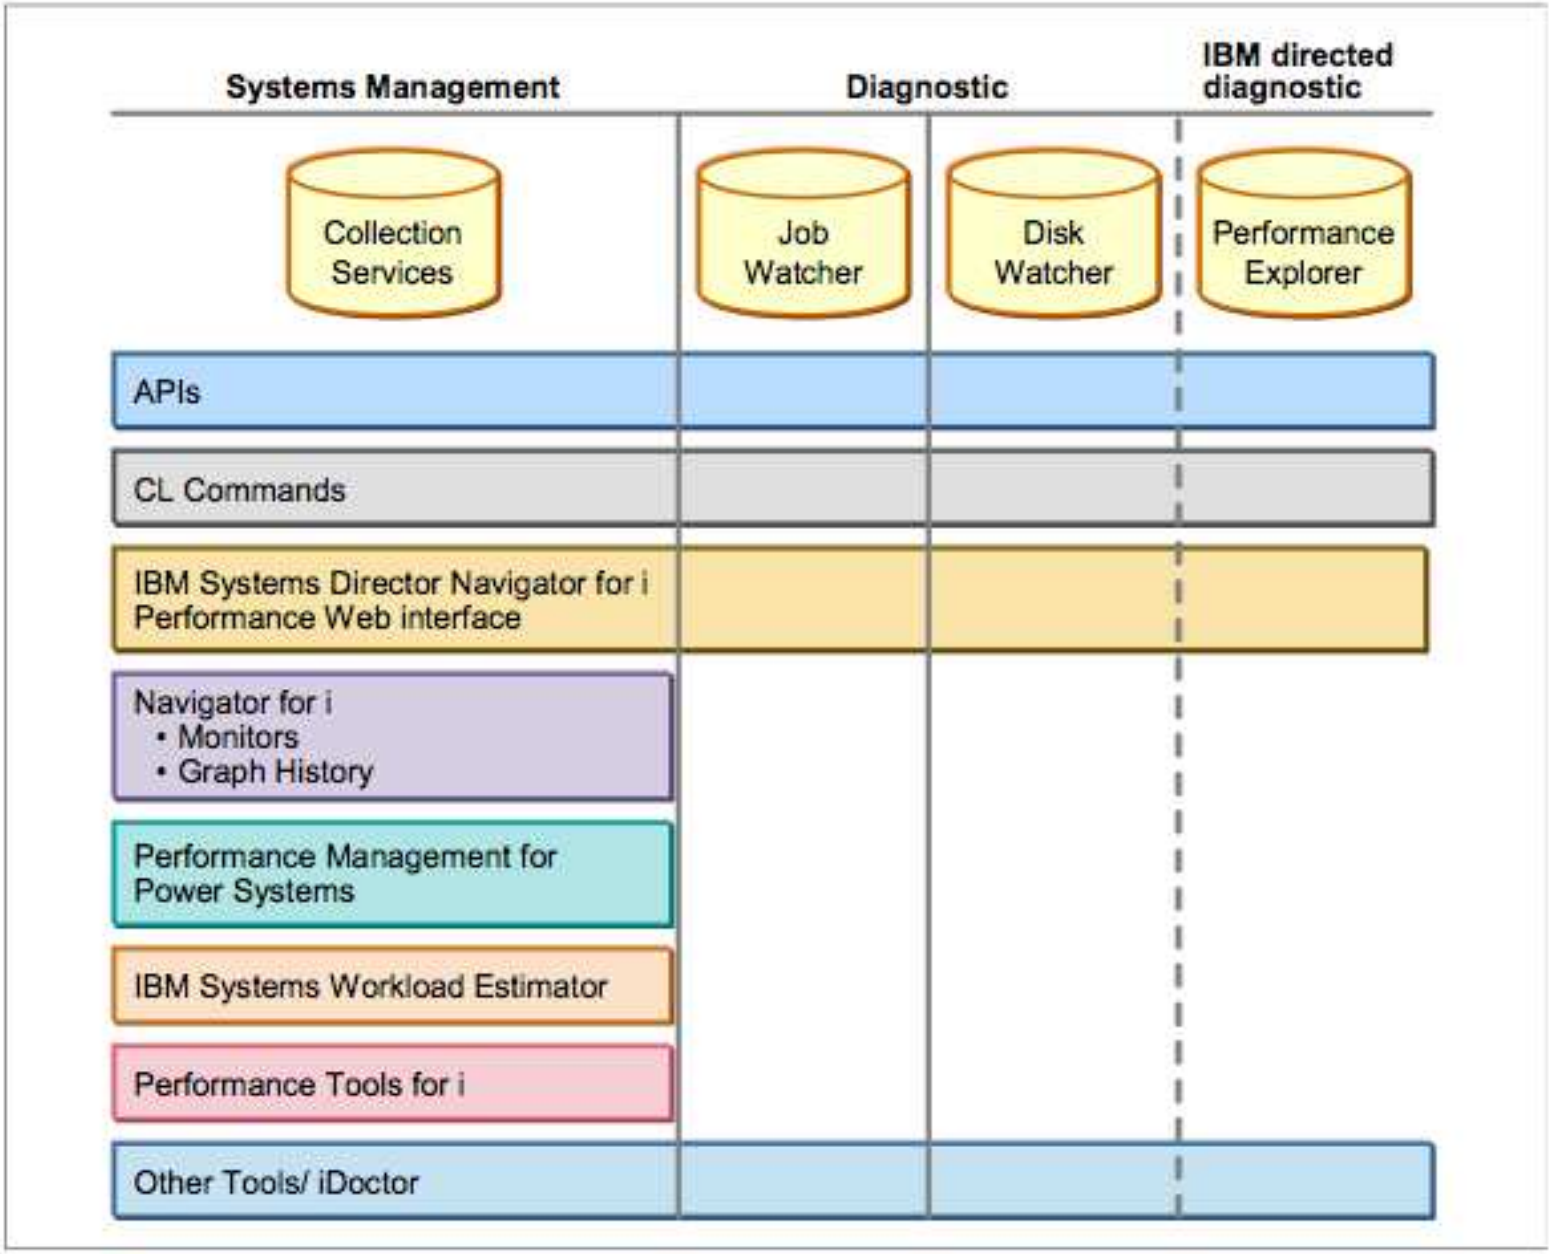

Real-time Monitoring

Real-time monitoring is essential for pro-active performance management

- Identify and correct potential issues before they become real problems
- Automated monitoring and notification

Several tools available for real-time monitoring:

- Watches
- CL Commands
- IBM Systems Director Navigator for i
- Management Central Monitors
- IBM Systems Director
- Tivoli Monitoring

Watches

Watches provide a programmatic interface to be notified when the following occur:

- Message
- Licensed Internal Code Log (LIC Log)
- Problem Activity Log Entry (PAL entry)

Start Watch (STRWCH) command or API (QSCSWCH)
End Watch (ENDWCH) command or API (QSCEWCH)

When the condition being watched occurs, your program gets control and you can take any action you want

CL Commands

Commands can be used for real-time performance monitoring

- Work with Active Jobs (WRKACTJOB)
- Work with Disk Status (WRKDSKSTS)
- Work with System Status (WRKSYSSTS)
- Work with System Activity (WRKSYSACT)
- Work with Shared Storage Pools (WRKSHRPOOL)

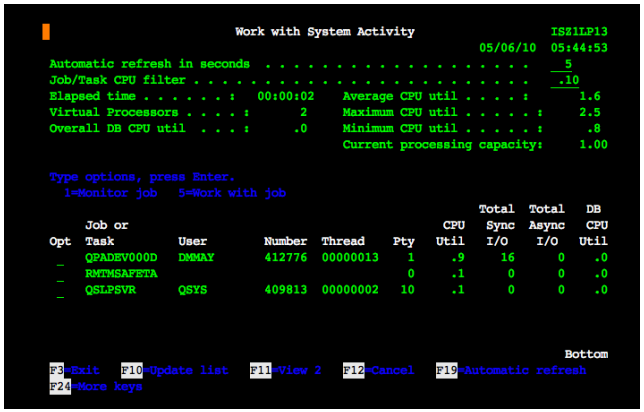

These commands are sampling based

- Enter the command
- Allow the system to collect data for a minimum of 5 minutes
- Refresh (F5) the display to see the performance data
- Restart (F10) will restart the elapsed time counter

IBM Systems Director Navigator for i

Recommended starting points for monitoring

Graphical User Interface views of

- System Status
- Disk Status
- Active Jobs

Health Indicators

- Graphical view of general health based upon Collection Services data
  - System Resources
  - CPU
  - Disk
  - Memory
  - Response Time

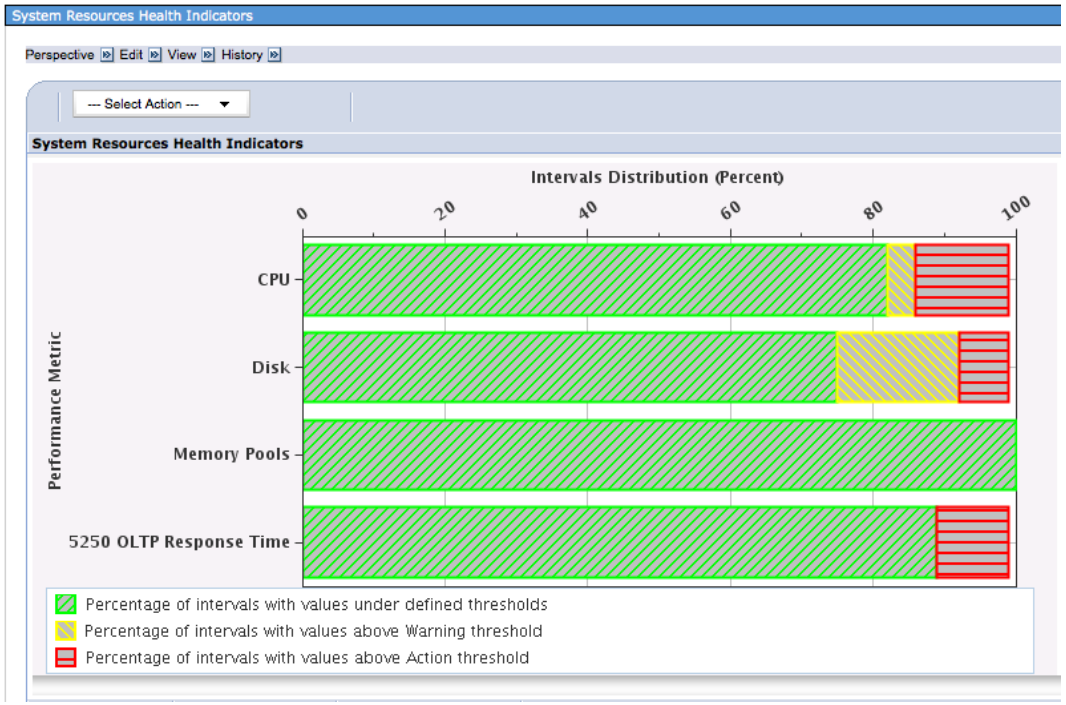

IBM Systems Director Navigator for i - Thresholds

IBM-supplied defaults for threshold levels on health indicators

- Can be customized for your requirements

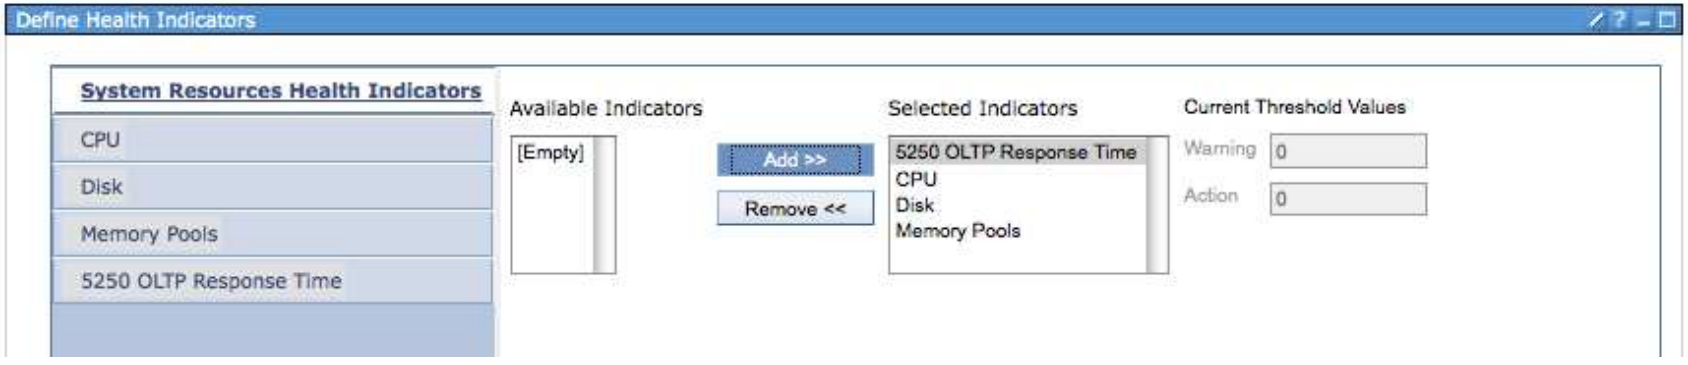

With Design Mode, Edit View, thresholds can be added to any chart

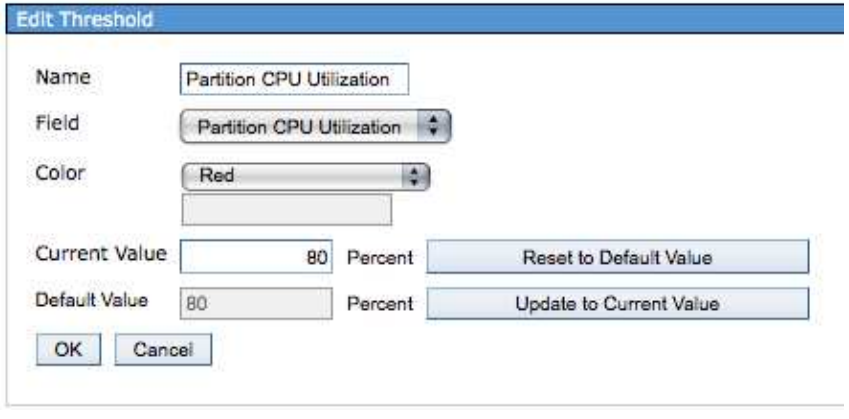

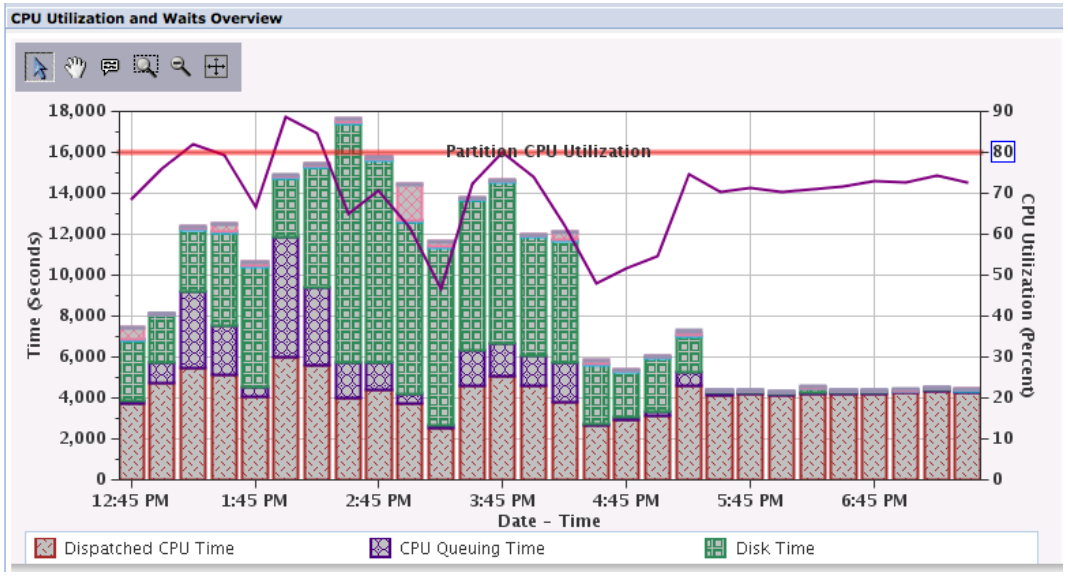

Management Central Monitors

Management Central Monitors are part of the System i Navigator

- They do not exist in the Director Navigator Web console

For performance management, use the System Monitors

- Provide real-time performance of multiple partitions
- Performance metrics
  - CPU Utilization
  - Disk Utilization
  - Faulting
  - Etc....
- Thresholds can be defined for automated notification
- Graph History allows you to see the data over an extended period of time

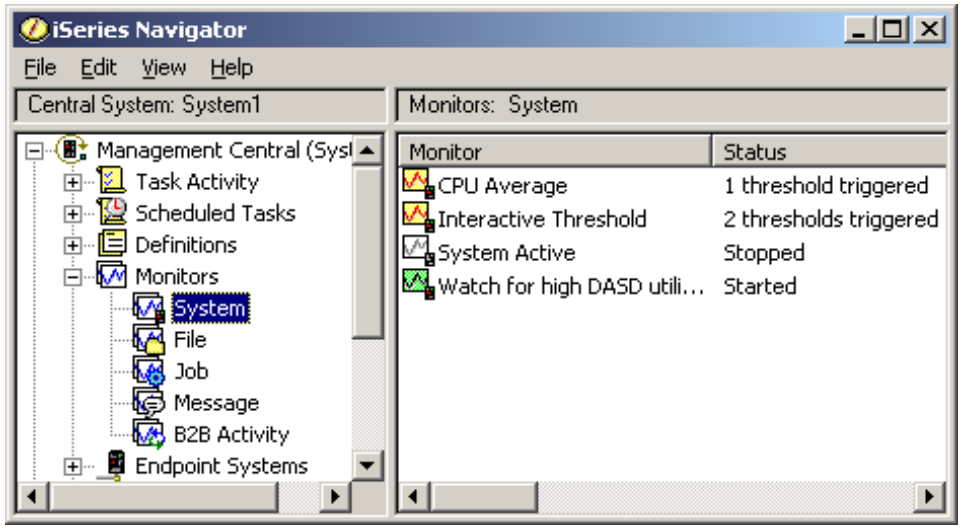

System monitors use Collection Services for the performance data

IBM Systems Director

IBM Systems Director is geared toward managing and monitoring a more complex environment

With Systems Director you can manage IBM i at 5.4 and later releases

Systems Director is integrated with Systems Director Navigator to 'drill down' into IBM i specific tasks

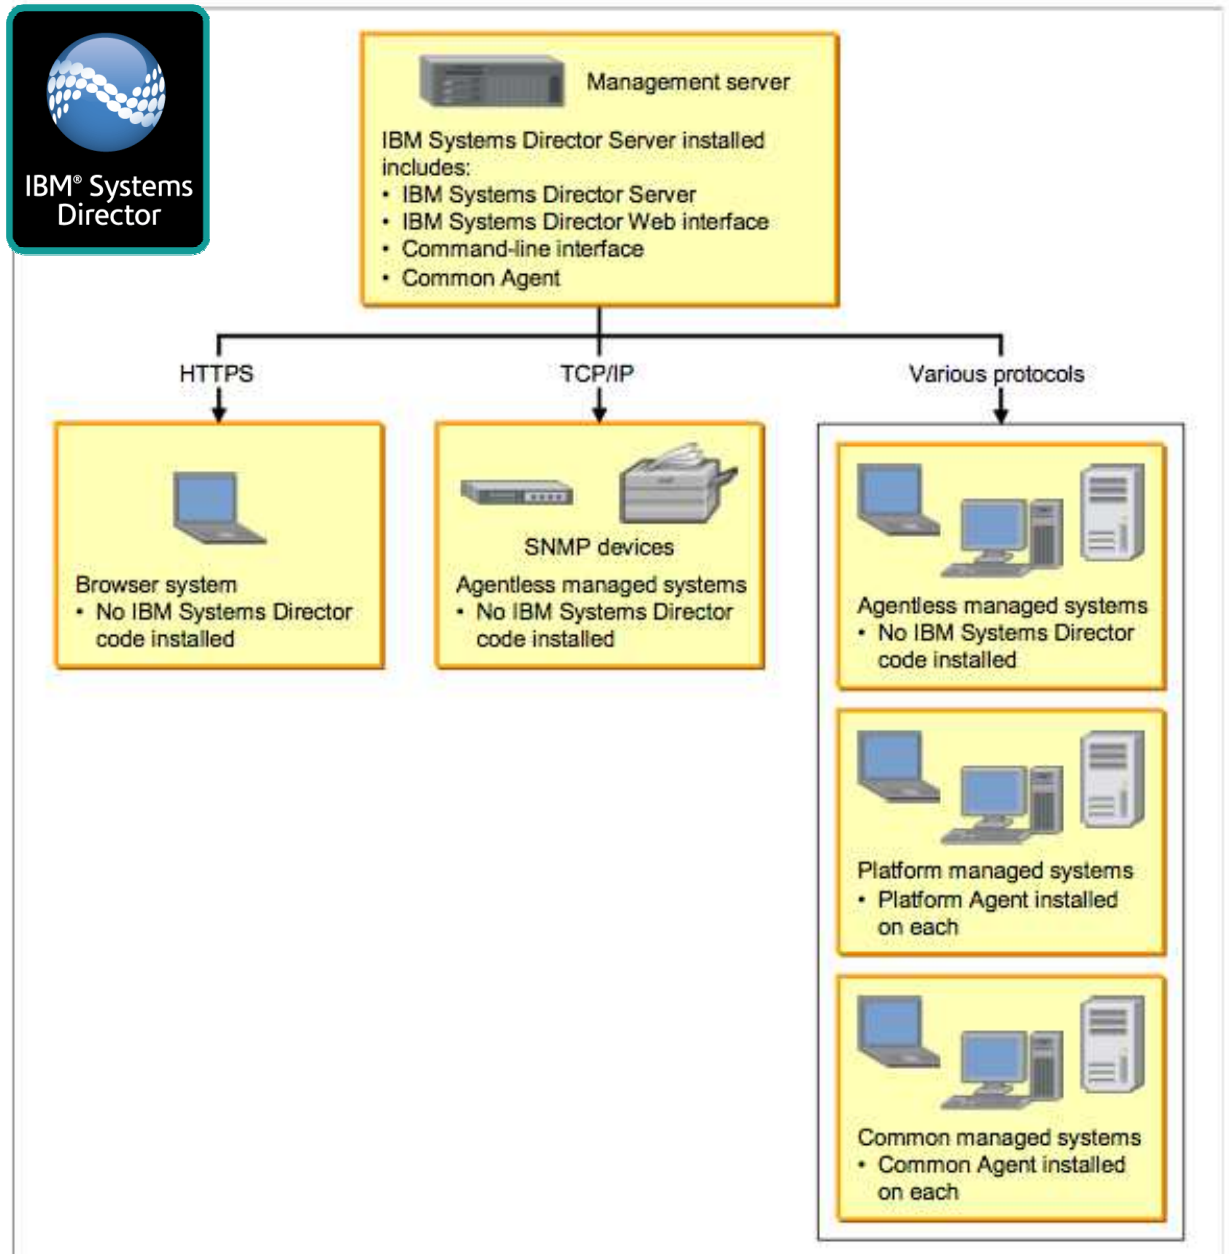

Systems Director:
Focus on Health, Status, Automation

- Health summary
- Favorite systems
- Critical monitors
- Group thumbnails
- Monitoring
- Monitor resources
- Thresholds
- Events
- Update Compliance
- Automation Plans
- Notify
- Run commands
- Trigger tasks

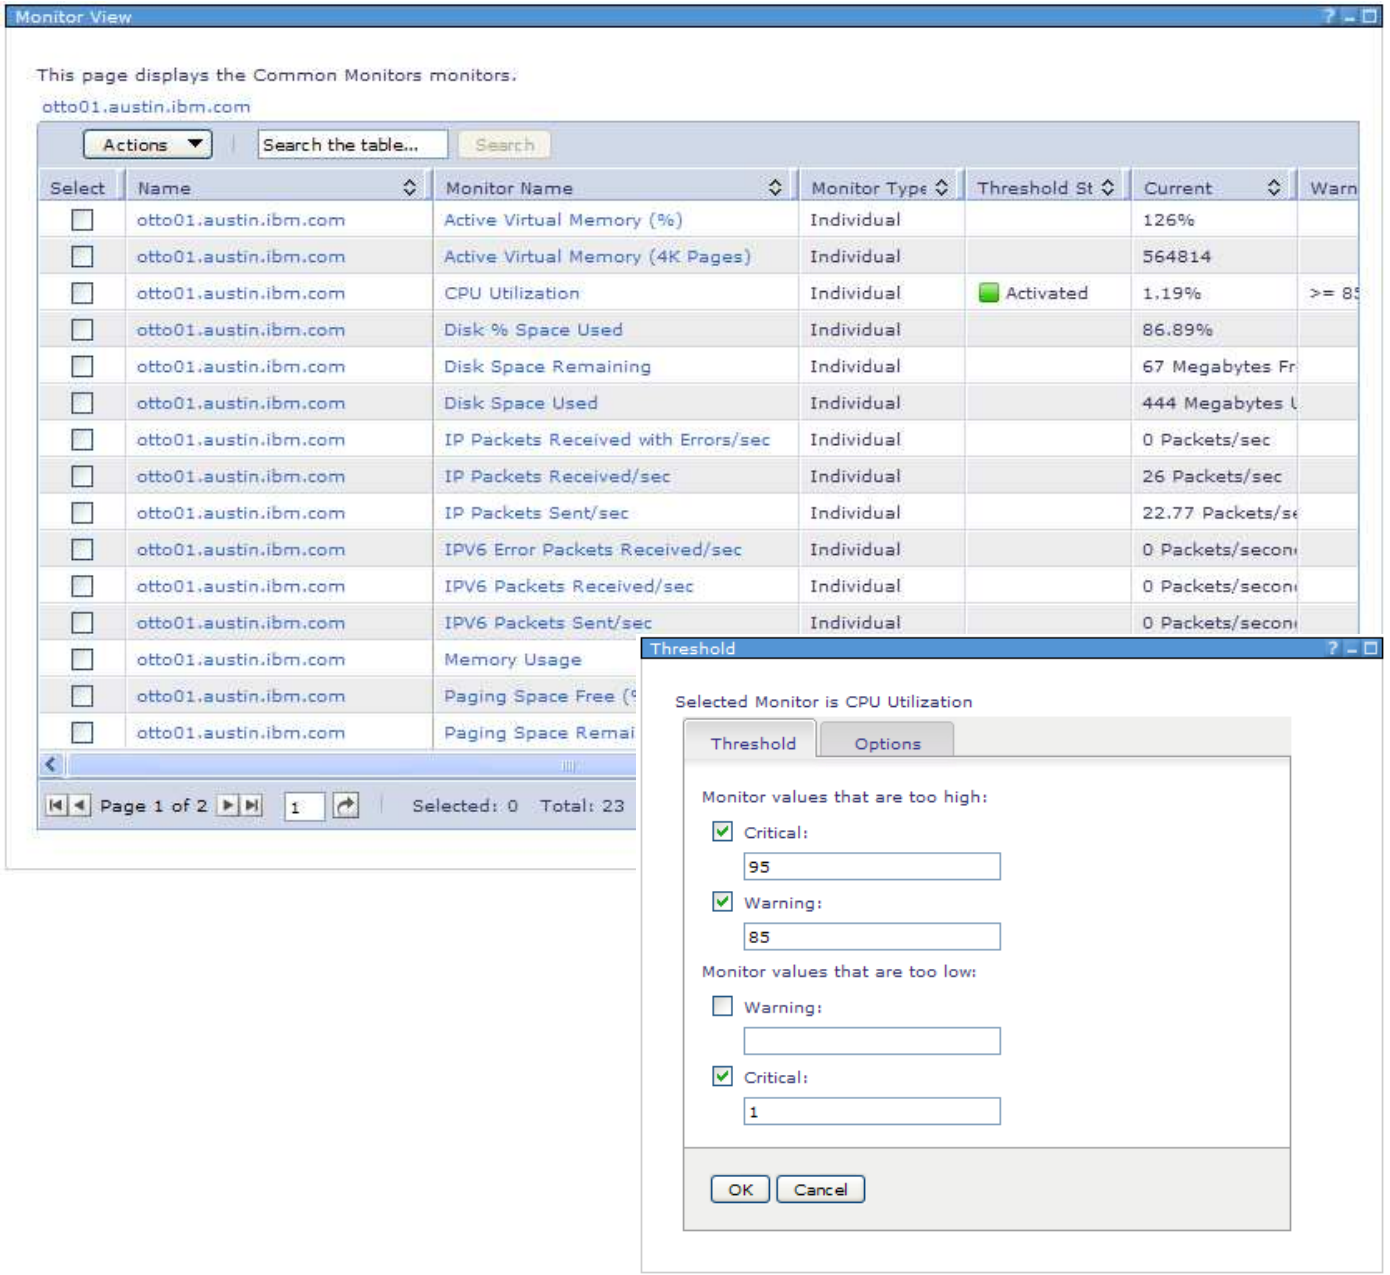

IBM Tivoli Monitoring

Enterprise-level monitoring and automation is accomplished with IBM Tivoli Monitoring

- Predefined or customized situations
- Policies for automation

Monitoring agent for IBM i provides a wealth of information that can be monitored

- Operating system
- CPU
- Disk
- Network
- Virtual and physical memory
- Paging information
- ....

and much, much more

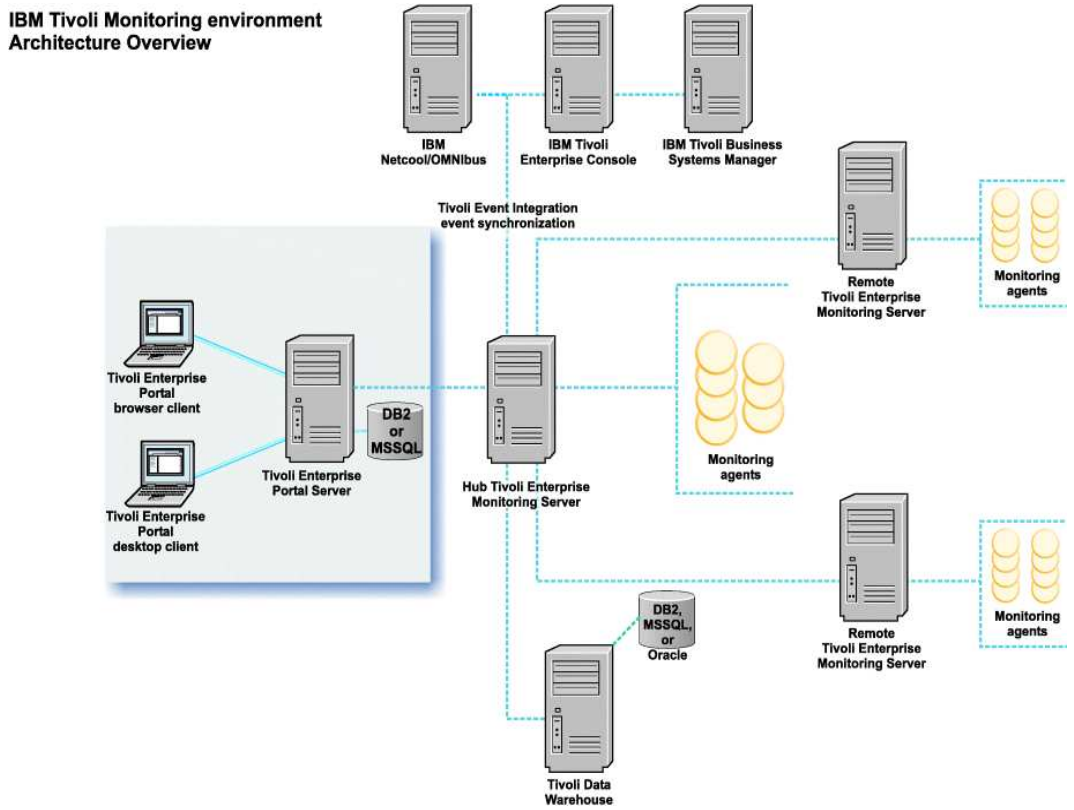

Power your planet.

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, what is or is not being used, how did it change, when did it happen, what is affected etc. Basically scope the problem

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

Tuning

Make the most efficient use of the system's resources

Iterative: measure, analyze, adjust

Basic strategy

- Set up initial tuning values
- **Measure**: observe system performance with these tools
  - WRKSYSSTS
  - WRKDSKSTS
  - WRKSYSACT
  - WRKACTJOB
- **Analyze**: review values, compare with performance goals
  - Interactive throughput and response time
  - Batch throughput for active batch jobs
  - Auxilliary input/output (AuxIO)
  - CPU percentage (CPU%)
- **Adjust**: if performance does not meet expectations, tune system based on new data
  - **Measure** and compare all key performance measurements
  - **Analyze** and evaluate adjustments one at a time
  - Review performance

Tuning, continued

General tips:

- Let the system tune itself automatically (IBM recommended)
  - Set system value QPFRADJ to 2
  - Performance adjustment and IPL and automatically

Work with Shared Pools – WRKSHRPOOL

Allows additional control over automatic adjustment of pools based on faulting and size parameters

```
                          Work with Shared Pools
                                                              System:   MYSYSTEM
Main storage size (M)  . :         4051.50

Type changes (if allowed), press Enter.

             Defined    Max    Allocated   Pool  -Paging Option--
Pool        Size (M)  Active  Size (M)     ID   Defined  Current
*MACHINE      800.00   +++++     800.00     1   *FIXED   *FIXED
*BASE        2536.01     601    2536.01     2   *FIXED   *FIXED
*INTERACT     715.22     170     715.23     3   *FIXED   *FIXED
*SPOOL           .25       1        .25     4   *FIXED   *FIXED
*SHRPOOL1        .00       0                    *FIXED
```


Power your planet.

Tuning – Work with Shared Pools

You can tune the values to influence the behavior of the performance adjuster

```
                          Work with Shared Pools
                                                          System:   ISZ1LP13
Main storage size (M)  . :         4051.50

Type changes (if allowed), press Enter.

                    -----Size %-----  -----Faults/Second------
Pool          Priority  Minimum  Maximum  Minimum  Thread  Maximum
*MACHINE             1     9.74      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
                                                                More...
```

Tuning – Pool Properties via Director Navigator

You can also do memory pool tuning through the GUI

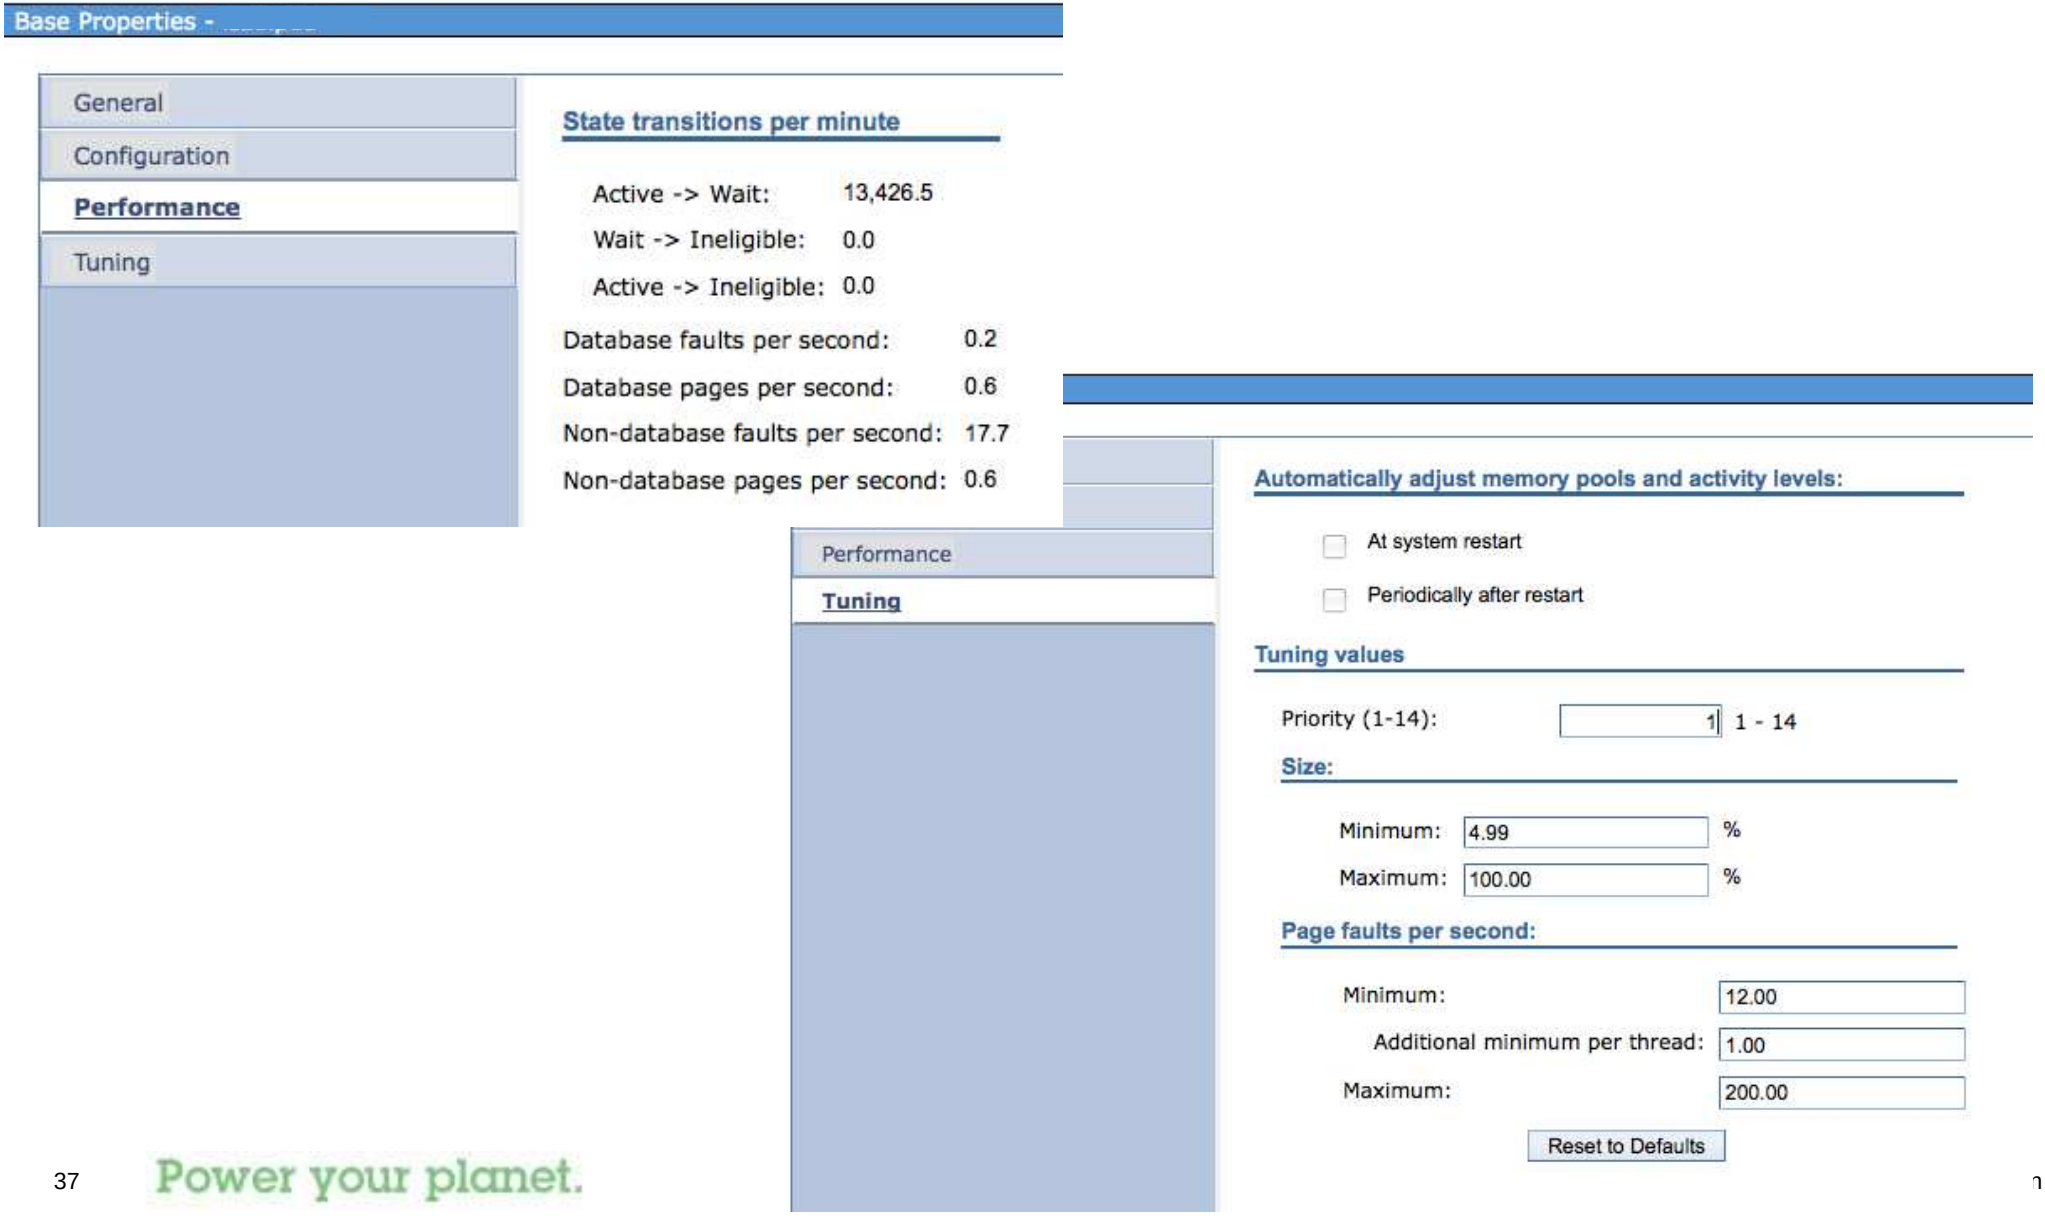

Tuning, continued

General tips:

If QPFRADJ = 0 (off) use WRKSYSSTS to manually adjust pools and activity levels

Minimize wait-to-ineligible transitions in interactive pools

< 10% of active-to-waits

Increase MAXACT by 5 or 10 to reduce wait-to-ineligible transitions

Repeat until they are less than 10% of the active-to-wait

Tune faulting in user pools

Sum of faults for all user pools should be less than the number of processors times the processors percent busy

Example: four processors running at 50% busy

Get faults to less than 200 faults/seconds (4 * 50)

Tune the machine pool to under 10 faults/second

Tuning, continued...

- Favor output over input so existing jobs are not adversely impacted
- Do not mix different types of jobs and priorities in the same pool
  - Remove batch jobs from *BASE by creating another batch pool
  - Route batch jobs to *SHRPOOL1
  - One batch job to a pool is ideal. This can be done by modifying the job description for the routing data, and the subsystem description for the routing entry
- On a shared production/development partition
  - Create a second interactive pool *SHRPOOL2 for programmers
  - Change programmers' job description routing data and subsystem description routing entries so their jobs run in SHRPOOL2

Tuning .... Work With System Status

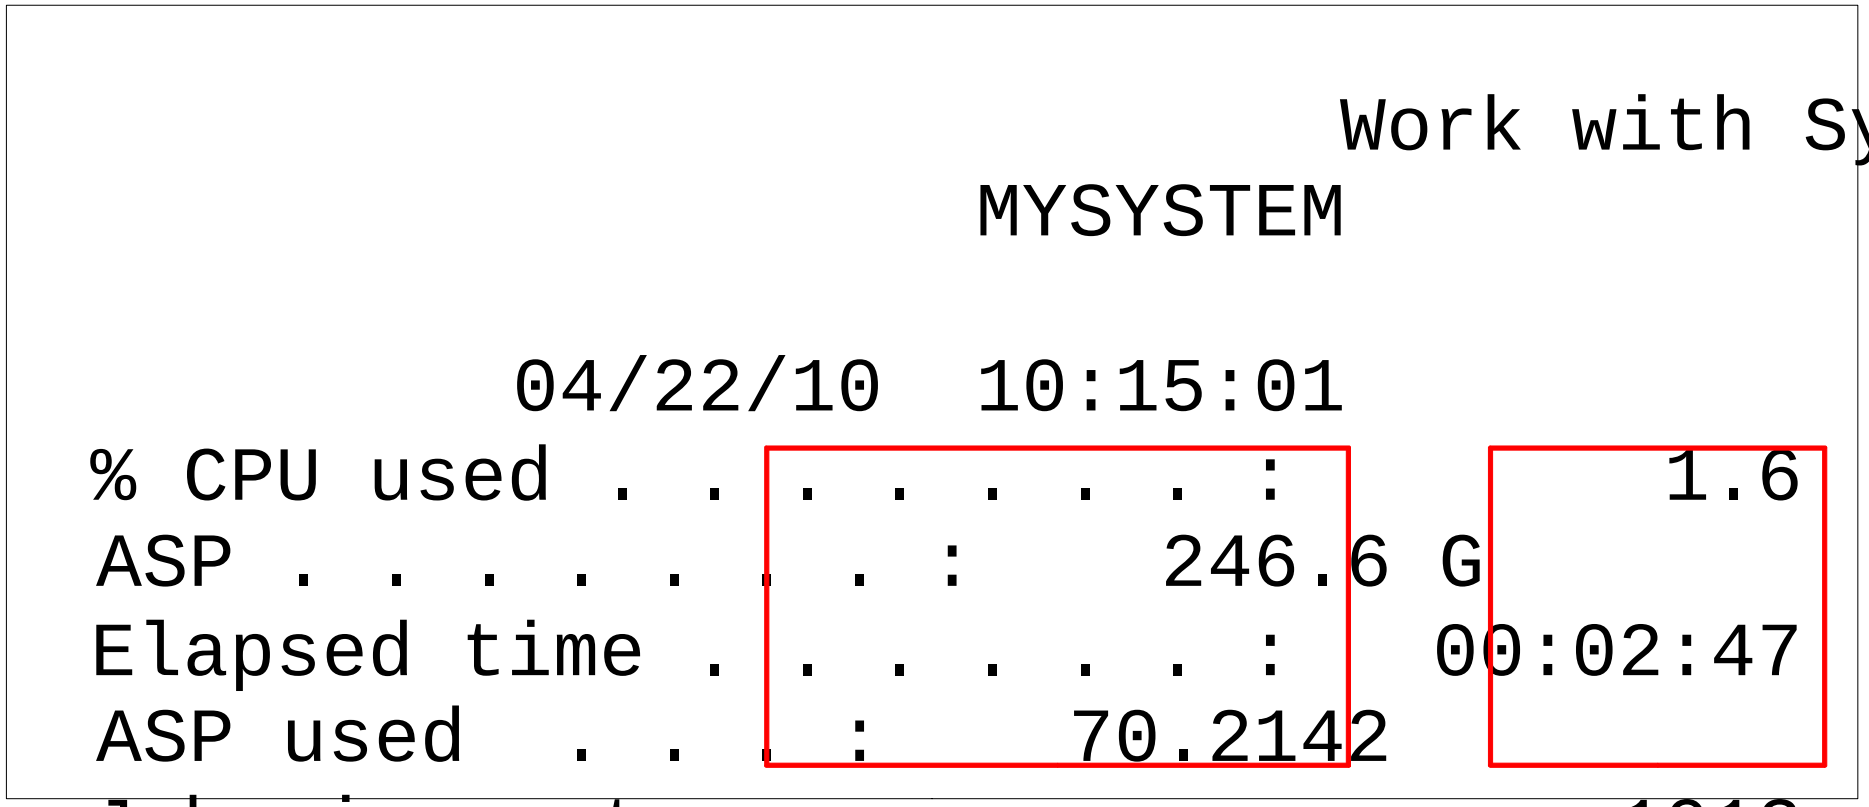

Historical Trending

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

- 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

Graph History

- Graph History allows you to see that performance data over time.
  - 7 days if you are not collecting performance data with PM
  - One year if you are collecting performance data with PM

PM for Power Systems

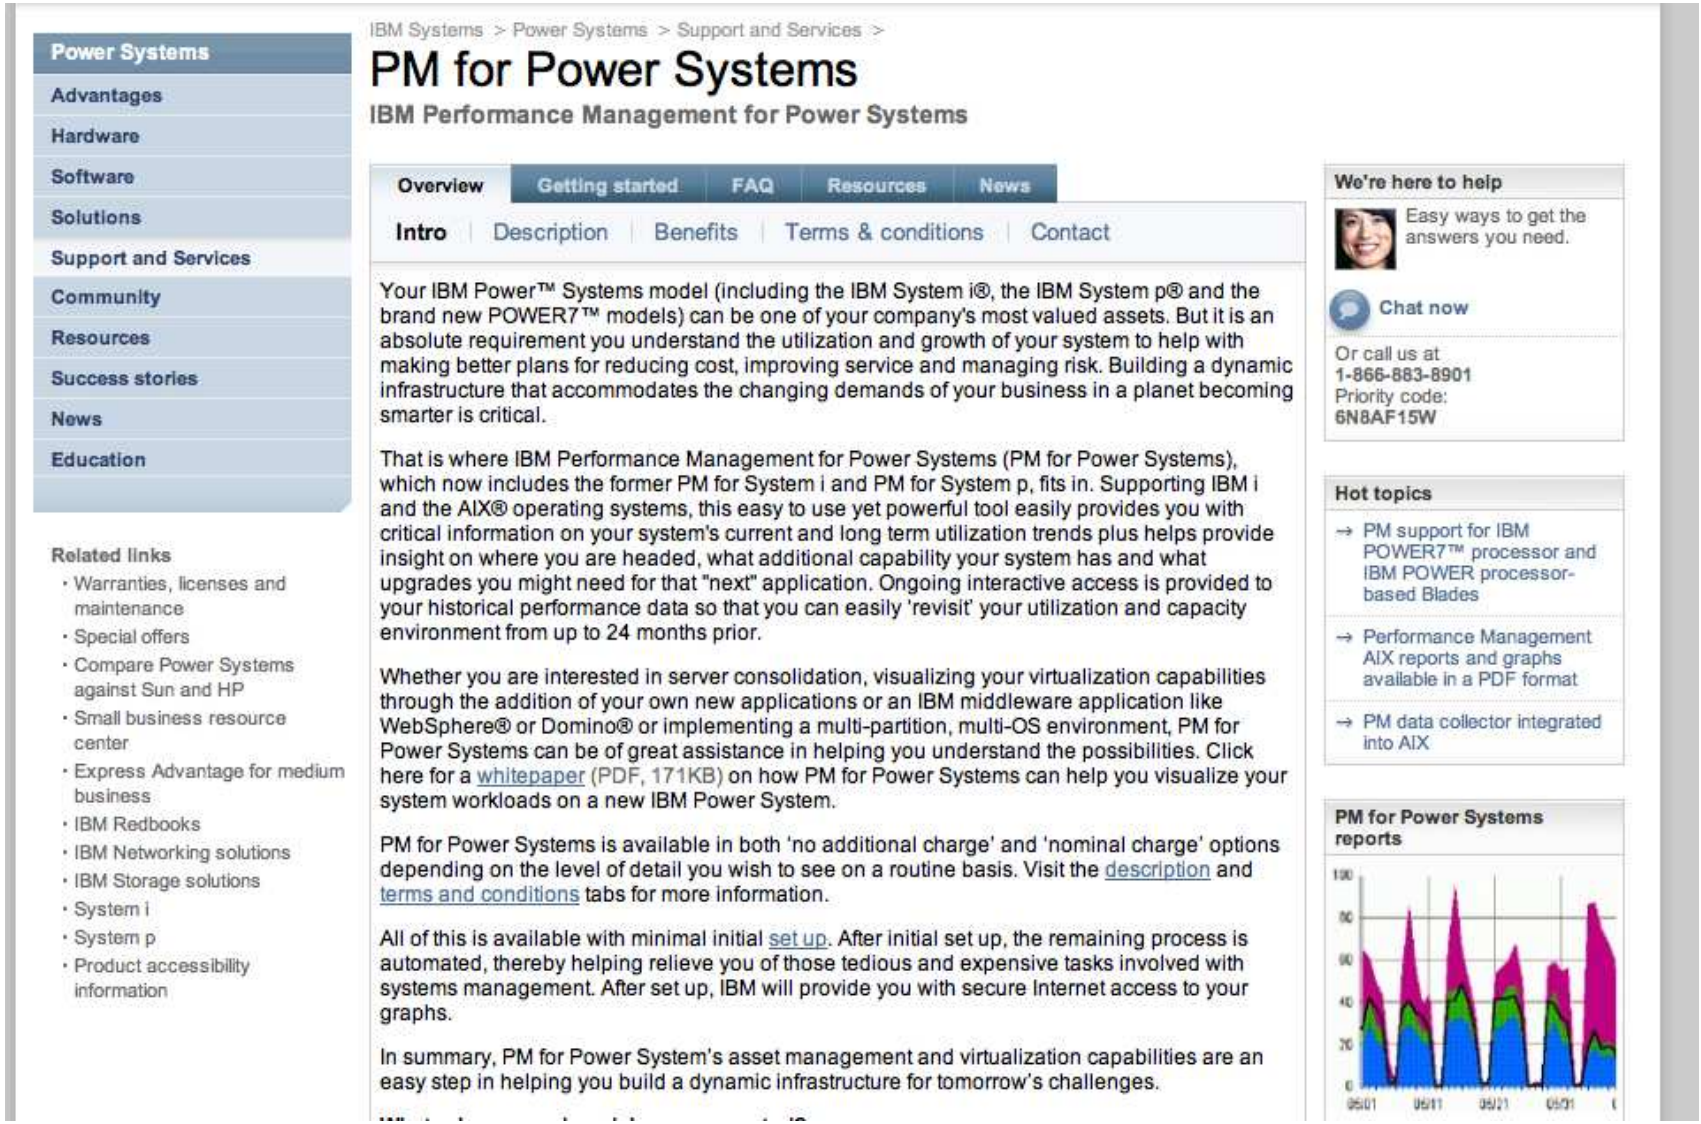

Power Systems
Advantages
Hardware
Software
Solutions
Support and Services
Community
Resources
Success stories
News
Education

Related links
- Warranties, licenses and maintenance
- Special offers
- Compare Power Systems against Sun and HP
- Small business resource center
- Express Advantage for medium business
- IBM Redbooks
- IBM Networking solutions
- IBM Storage solutions
- System i
- System p
- Product accessibility information

IBM Systems > Power Systems > Support and Services >

PM for Power Systems

IBM Performance Management for Power Systems

Overview | Getting started | FAQ | Resources | News

Intro | Description | Benefits | Terms & conditions | Contact

Your IBM Power™ Systems model (including the IBM System i®, the IBM System p® and the brand new POWER7™ models) can be one of your company's most valued assets. But it is an absolute requirement you understand the utilization and growth of your system to help with making better plans for reducing cost, improving service and managing risk. Building a dynamic infrastructure that accommodates the changing demands of your business in a planet becoming smarter is critical.

That is where IBM Performance Management for Power Systems (PM for Power Systems), which now includes the former PM for System i and PM for System p, fits in. Supporting IBM i and the AIX® operating systems, this easy to use yet powerful tool easily provides you with critical information on your system's current and long term utilization trends plus helps provide insight on where you are headed, what additional capability your system has and what upgrades you might need for that "next" application. Ongoing interactive access is provided to your historical performance data so that you can easily 'revisit' your utilization and capacity environment from up to 24 months prior.

Whether you are interested in server consolidation, visualizing your virtualization capabilities through the addition of your own new applications or an IBM middleware application like WebSphere® or Domino® or implementing a multi-partition, multi-OS environment, PM for Power Systems can be of great assistance in helping you understand the possibilities. Click here for a whitepaper (PDF, 171KB) on how PM for Power Systems can help you visualize your system workloads on a new IBM Power System.

PM for Power Systems is available in both 'no additional charge' and 'nominal charge' options depending on the level of detail you wish to see on a routine basis. Visit the description and terms and conditions tabs for more information.

All of this is available with minimal initial set up. After initial set up, the remaining process is automated, thereby helping relieve you of those tedious and expensive tasks involved with systems management. After set up, IBM will provide you with secure Internet access to your graphs.

In summary, PM for Power System's asset management and virtualization capabilities are an easy step in helping you build a dynamic infrastructure for tomorrow's challenges.

We're here to help

Easy ways to get the answers you need.

Chat now

Or call us at
1-866-883-8901
Priority code:
6N8AF15W

Hot topics

→ PM support for IBM POWER7™ processor and IBM POWER processor-based Blades

→ Performance Management AIX reports and graphs available in a PDF format

→ PM data collector integrated into AIX

PM for Power Systems reports

http://www-03.ibm.com/systems/power/support/perfmgmt/

Power your planet.

Capacity Planning

Goal of capacity planning is to make your system most efficient with regard to
Utilizing resources
Fulfilling user requests

IBM Systems Workload Estimator (WLE) is the IBM tool for capacity planning
Submit input from the PM for Power Systems data
Select a specific 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

Capacity Planning Cycle

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 projects 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

IBM Systems Workload Estimator (WLE)

Size Next Upgrade action from Performance Data Investigator

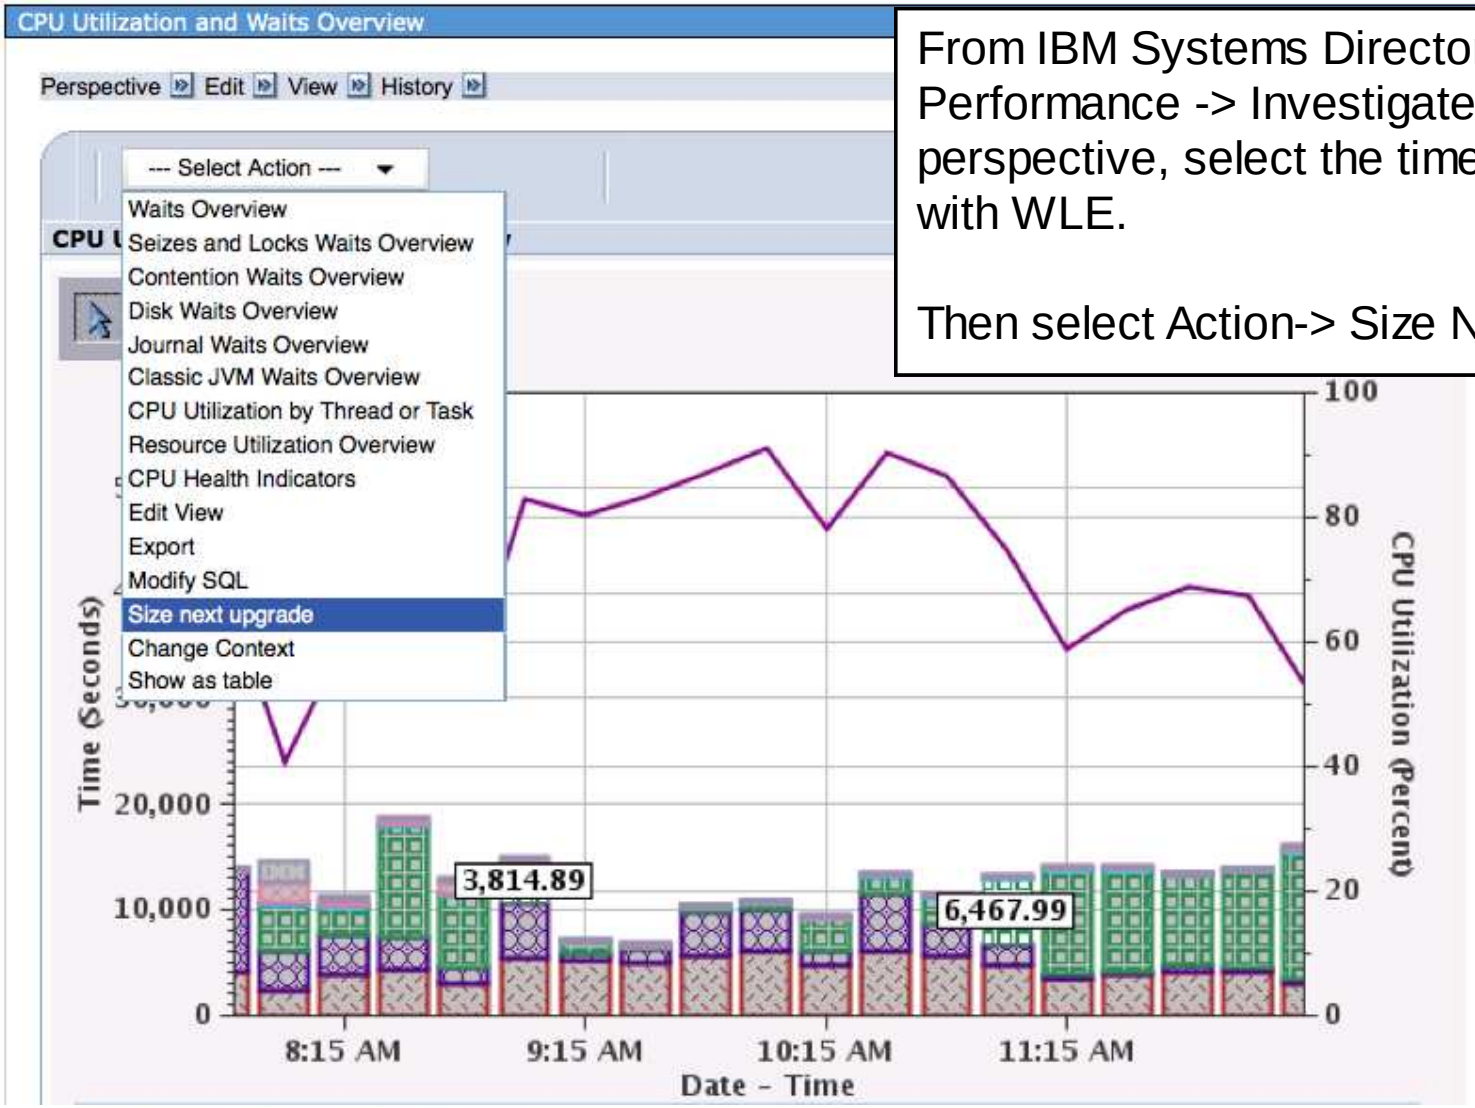

From IBM Systems Director Navigator for i -> Performance -> Investigate Data -> Display perspective, select the timeframe you want to size with WLE.

Then select Action-> Size Next Upgrade

IBM Systems Workload Estimator (WLE)

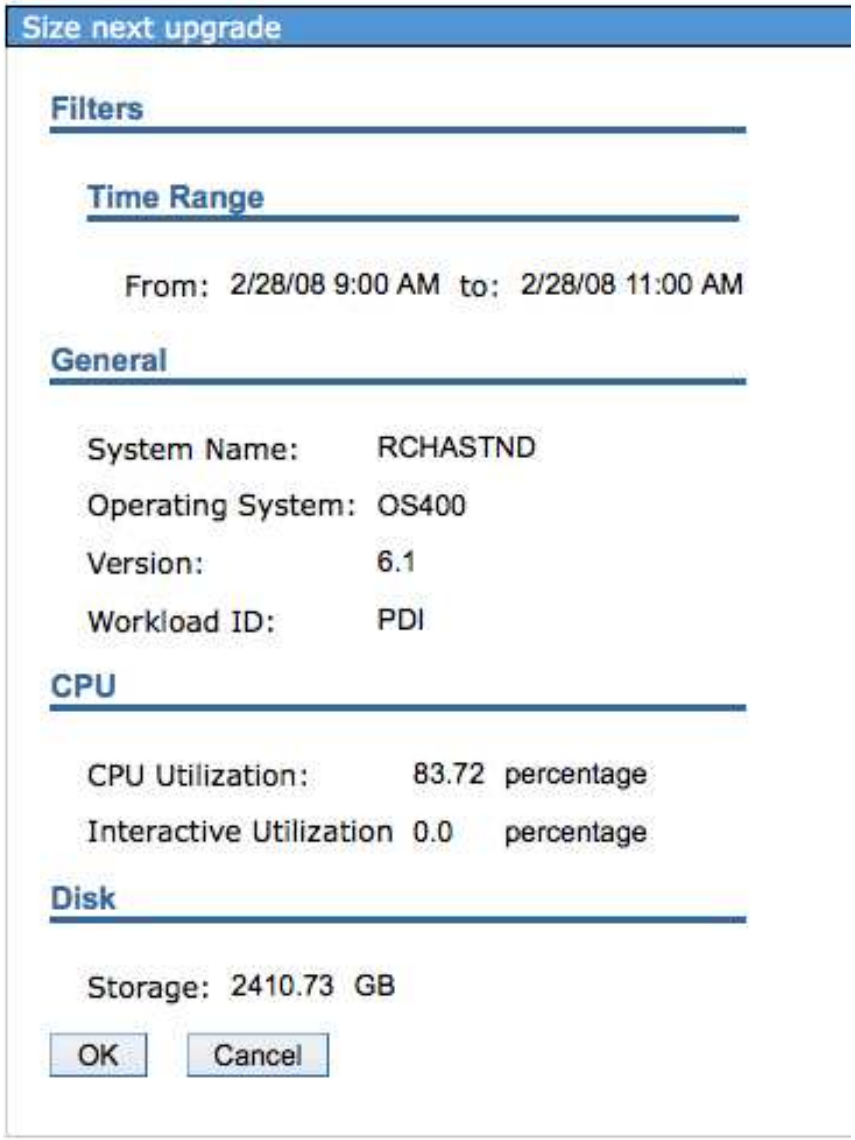

Power your planet.

WLE – Workload Definition Panel

Definition page - PDI Workload Definition

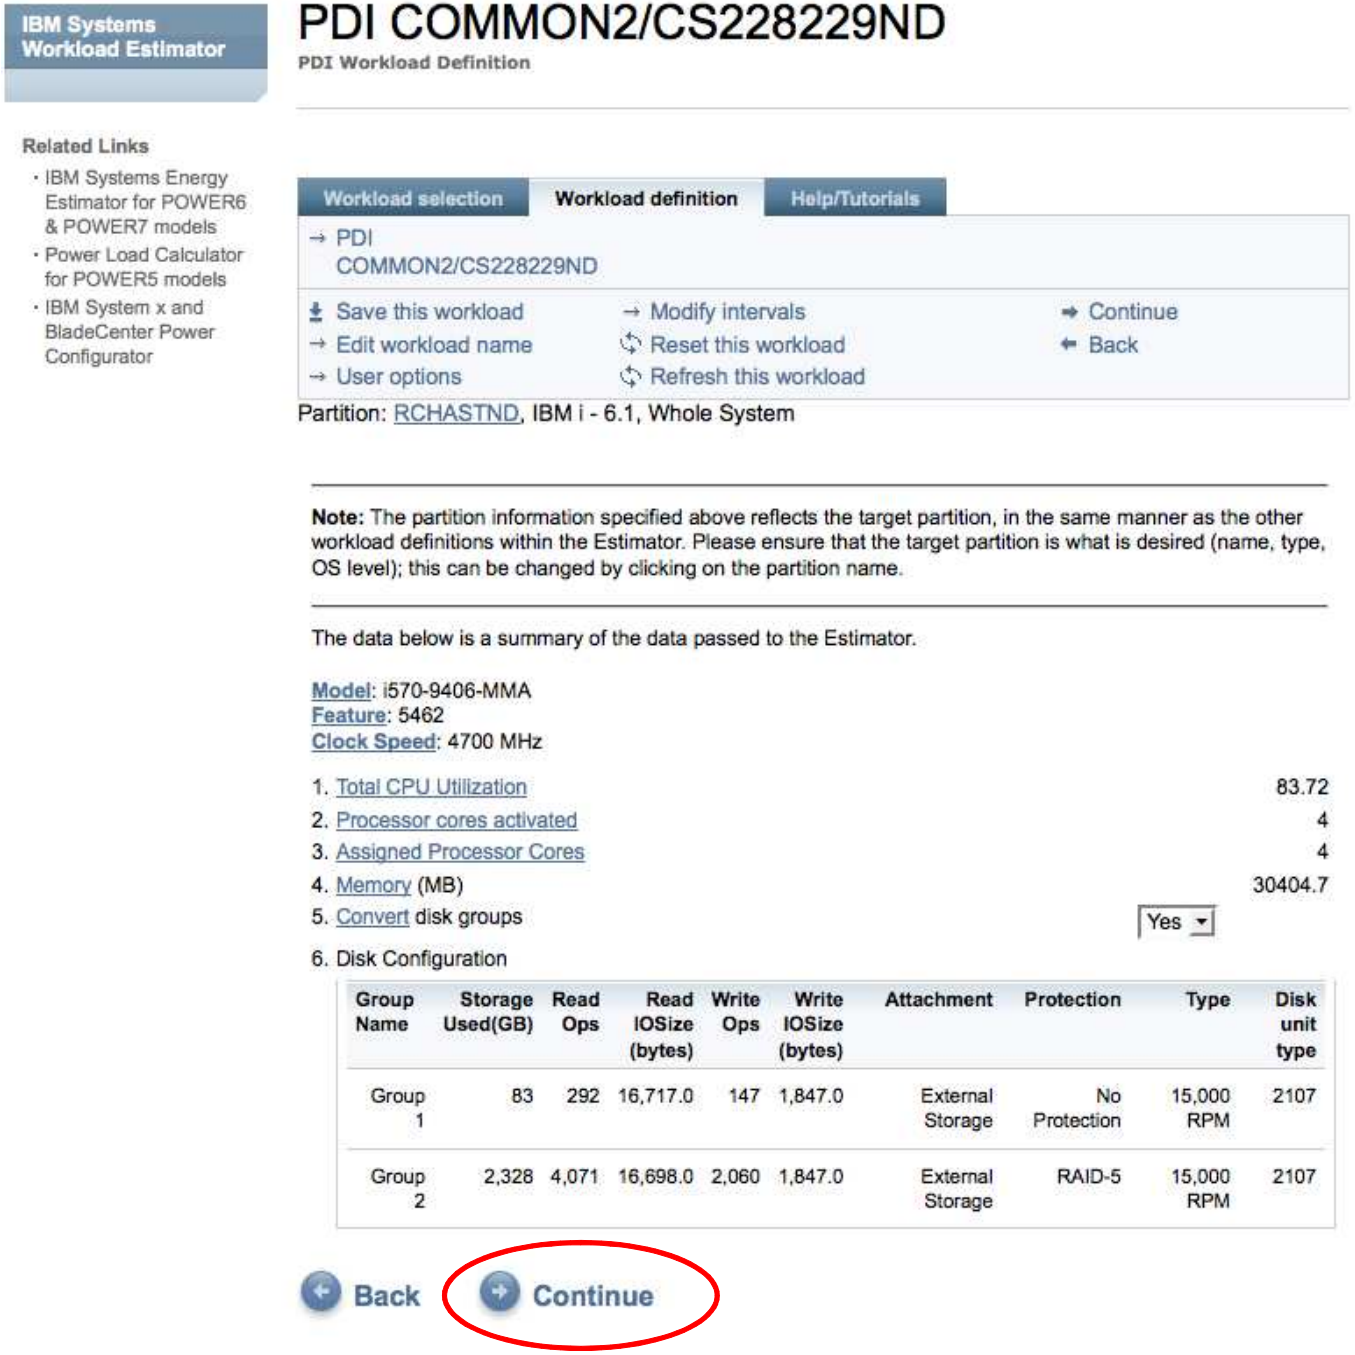

IBM Systems Workload Estimator

Related Links
- IBM Systems Energy Estimator for POWER6 & POWER7 models
- Power Load Calculator for POWER5 models
- IBM System x and BladeCenter Power Configurator

PDI COMMON2/CS228229ND

PDI Workload Definition

Workload selection | Workload definition | Help/Tutorials

→ PDI COMMON2/CS228229ND

Save this workload | → Modify intervals | → Continue
→ Edit workload name | Reset this workload | ← Back
→ User options | Refresh this workload

Partition: RCHASTND, IBM i - 6.1, Whole System

Note: The partition information specified above reflects the target partition, in the same manner as the other workload definitions within the Estimator. Please ensure that the target partition is what is desired (name, type, OS level); this can be changed by clicking on the partition name.

The data below is a summary of the data passed to the Estimator.

Model: i570-9406-MMA
Feature: 5462
Clock Speed: 4700 MHz

1. Total CPU Utilization 83.72
2. Processor cores activated 4
3. Assigned Processor Cores 4
4. Memory (MB) 30404.7
5. Convert disk groups Yes
6. Disk Configuration

Group Name	Storage Used(GB)	Read Ops	Read IOSize (bytes)	Write Ops	Write IOSize (bytes)	Attachment	Protection	Type	Disk unit type
Group 1	83	292	16,717.0	147	1,847.0	External Storage	No Protection	15,000 RPM	2107
Group 2	2,328	4,071	16,698.0	2,060	1,847.0	External Storage	RAID-5	15,000 RPM	2107

Back | Continue

WLE – Selected System

Selection System Tab:

Immediate Solution and Growth Solution are shown

Physical systems capable to support the measured workload

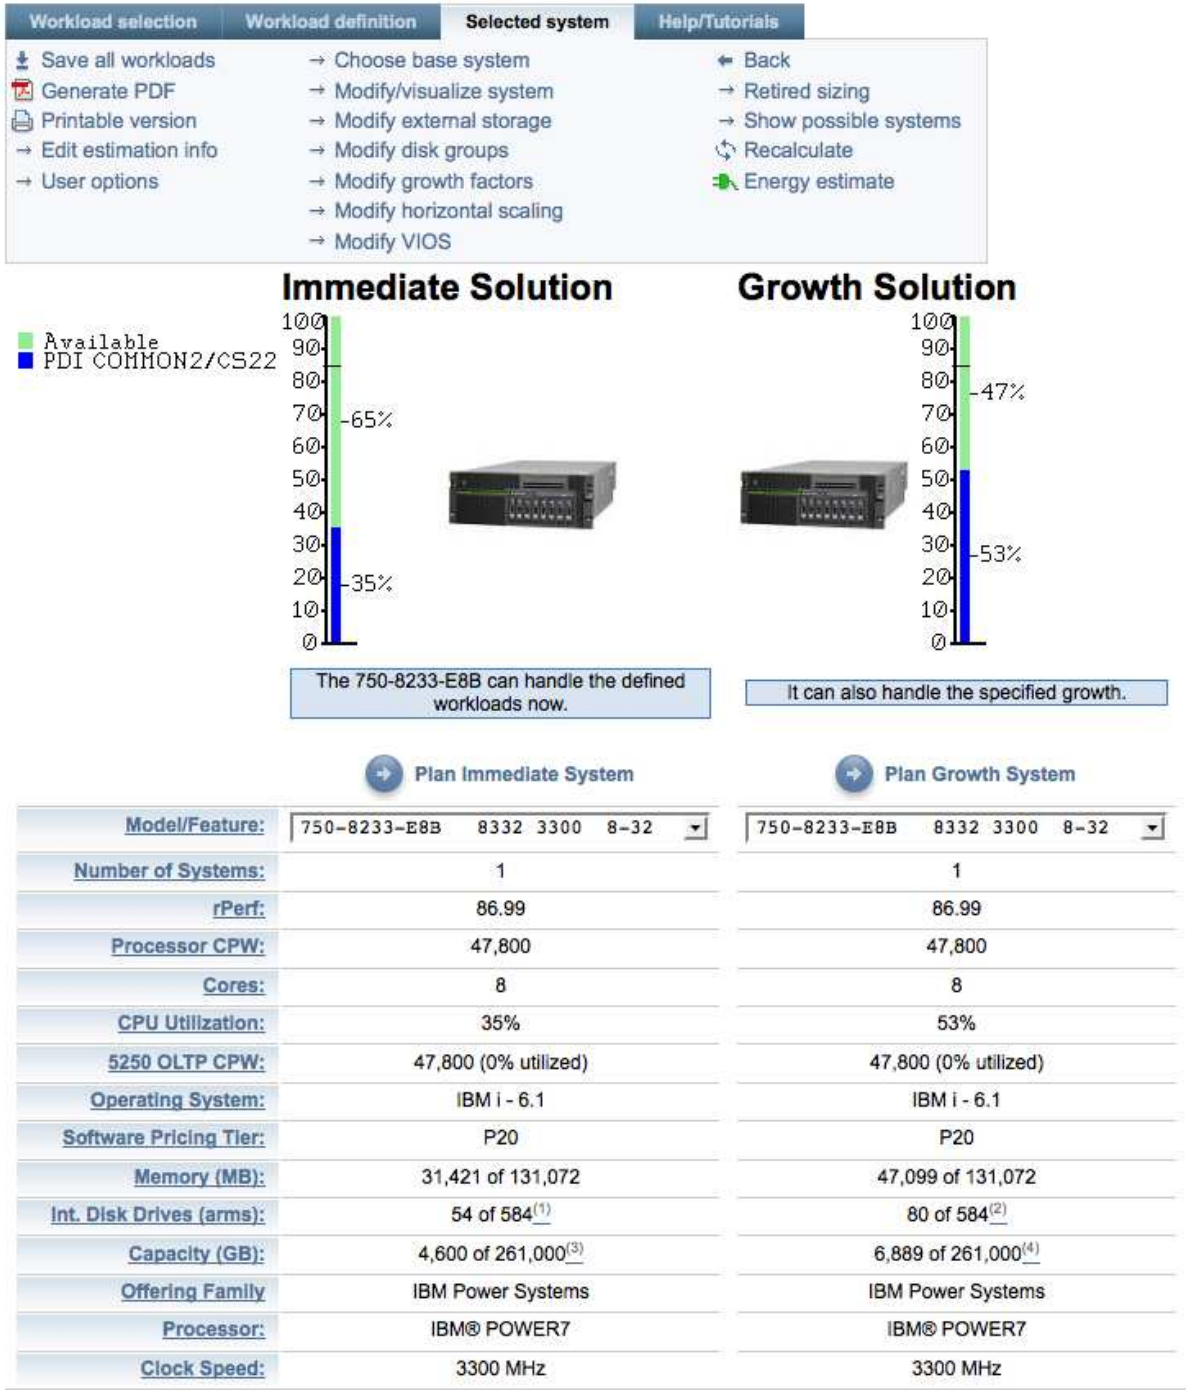

Guidelines - CPU

CPU Utilization Guidelines depend upon the number of processors defined to the partition

- 50% 1-way
- 70% 2-way
- 85% 8-way
- 90% 32-way

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

Guidelines - Memory

Rules of thumb – starting point from which you can fine tune to determine appropriate thresholds for your environment

Memory

- aulting is normal and expected

How much is too much?

Generally …

(100 * (CPU utilization by pool / 100)) * number of processors)

Use the Performance Adjustor if your workloads are slow to change

E.g., Daytime versus over night

Rapidly changing workloads are not well suited for the Performance Adjustor

QPFRADJ

Automatically manages the shared memory pools without any user interaction

Set to 2 or 3

Memory - Pool Faulting Guidelines

Machine Pool Faulting Guidelines

The machine pool faulting rate should be less than 10/faults per second

Disk utilization should be less than 40%

IOP utilization should be less than 60%

User Pool Faulting Guidelines

Paging guidelines no longer published

Track response time, throughput, and CPU utilization against the faulting rates to determine a reasonable faulting rate

Determine how much the faulting is affecting the interactive response time or batch throughput.

Data is available with PRTSYSRPT

Guidelines - Disk

Disk Response Time may be the most important component in overall system performance in today's system

Very fast processor technology, while disk performance has not yet made as much progress

Solid State Devices will have some impact here

Rules of thumb for categories:

Good: Average response time less than 5 milliseconds

Normal: Average response time between 5 and 10 milliseconds

Requires analysis: anything over 10 milliseconds

Alternatively - Establish disk response time objectives for a given application

This becomes your benchmark

Dividing the total disk response time per transaction by the average number of synchronous disk I/Os per transaction

Online Transaction Processing (OLTP) objectives will vary from batch processing because it may be more disk I/O intensive

http://ibmsystemsmag.blogs.com/i_can/2010/05/-i-can-measure-disk-response-times.html

Disk Response Time Groups
7.1 Performance Data Investigator enhancement

- Disk
  - Disk Response Time
    - Detailed
      - Disk I/O Rates Overview - Detailed
      - Disk I/O Rates Overview With Cache Statistics - Detailed
      - Disk I/O Average Response Time Overview - Detailed
      - Disk I/O Total Response Time Overview - Detailed
      - Disk I/O Total Service Time Overview - Detailed
    - Disk I/O Rates Overview
    - Disk I/O Rates Overview With Cache Statistics
    - Disk I/O Average Response Time Overview
    - Disk I/O Total Response Time Overview
    - Disk I/O Total Service Time Overview

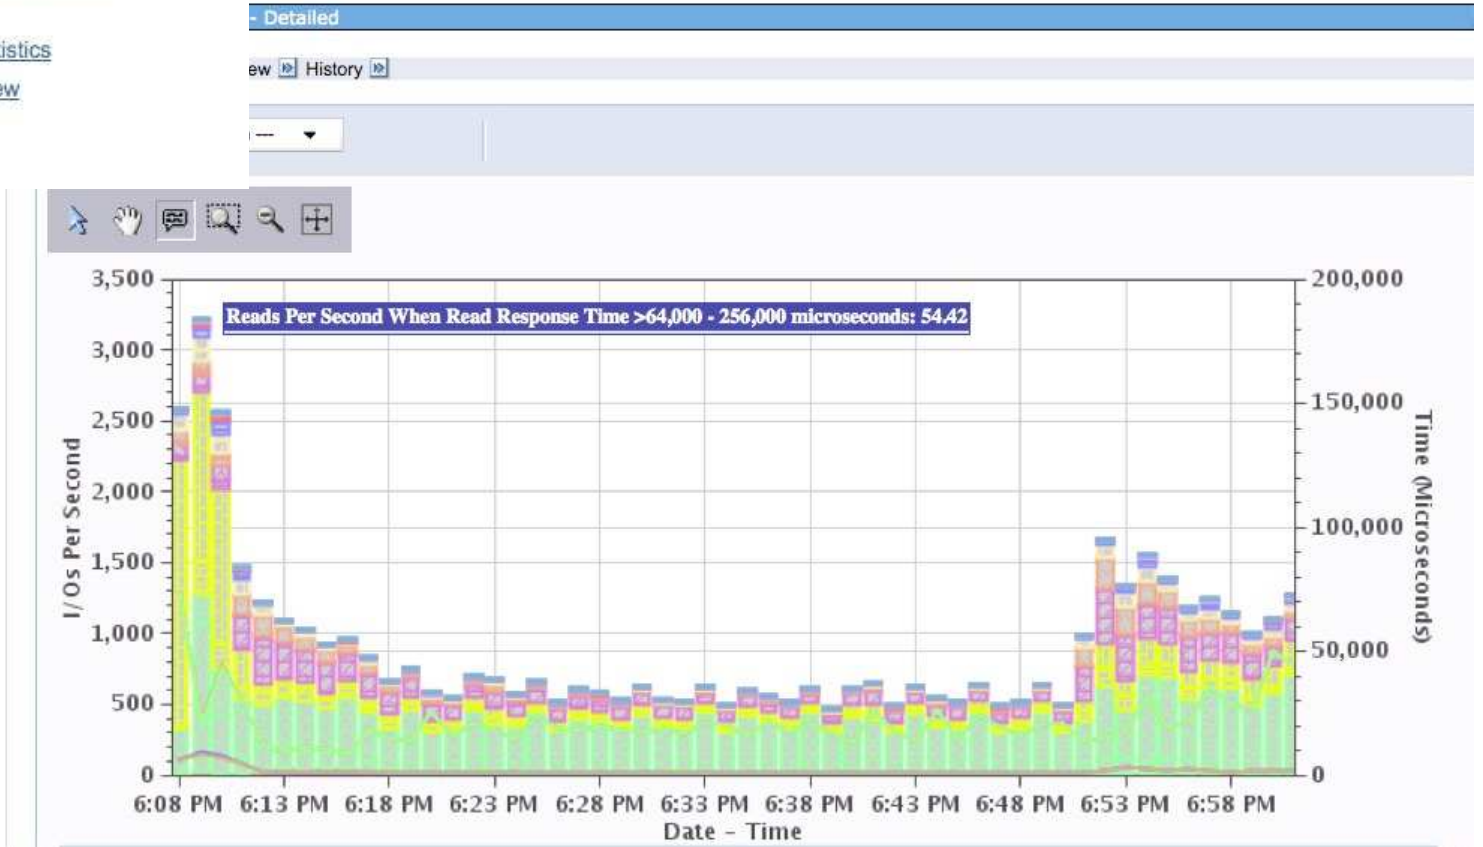

Power your planet.

Guidelines - Disk

Disk Percent Busy

Disk Percent Busy is not equal to Disk Utilization

A disk may be 100% busy and still be able to perform additional disk operations

Arbitrary:

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%)

Other batch jobs can be time sensitive

High disk utilization may indicate a need to inspect the disk response time

Disk Storage

ASP percentage on WRKSYSSTS (% system ASP used) should be less than 80%

Key Performance Indicators

Performance Indicators gauge how well a system is performing over time

Key Performance Indicators (KPIs)

- selection of metrics that make sense to your business

Easily measured

Track over time to determine if they are static, increasing or decreasing

Track against predefined thresholds

Two recommended categories:

Hardware resource utilization

Throughput

Hardware Resource Utilization & Throughput

Hardware Resource Utilization

Total CPU Utilization

CPU usage by job priority (batch & interactive)

Memory faults per second by memory pool (database & non-database)

Disk (arm utilization, response time, wait time, space utilization)

Throughput

A business transaction varies per business

An application can be built to capture these metrics and collect them such as is done by Collection Services.

IBM i defined metrics for components that typically make up a business transaction:

OLTP transactions (5250 and/or HTTP) per hour (peak & average)

Average response time per transaction

CPU usage per transaction

Disk I/O per transaction

Elapsed time of batch jobs

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

See this **Redbooks** publication for complete coverage of IBM i Performance Management Best Practices.

Based upon the 6.1 release and the October 2009 update to the Performance Data Investigator.

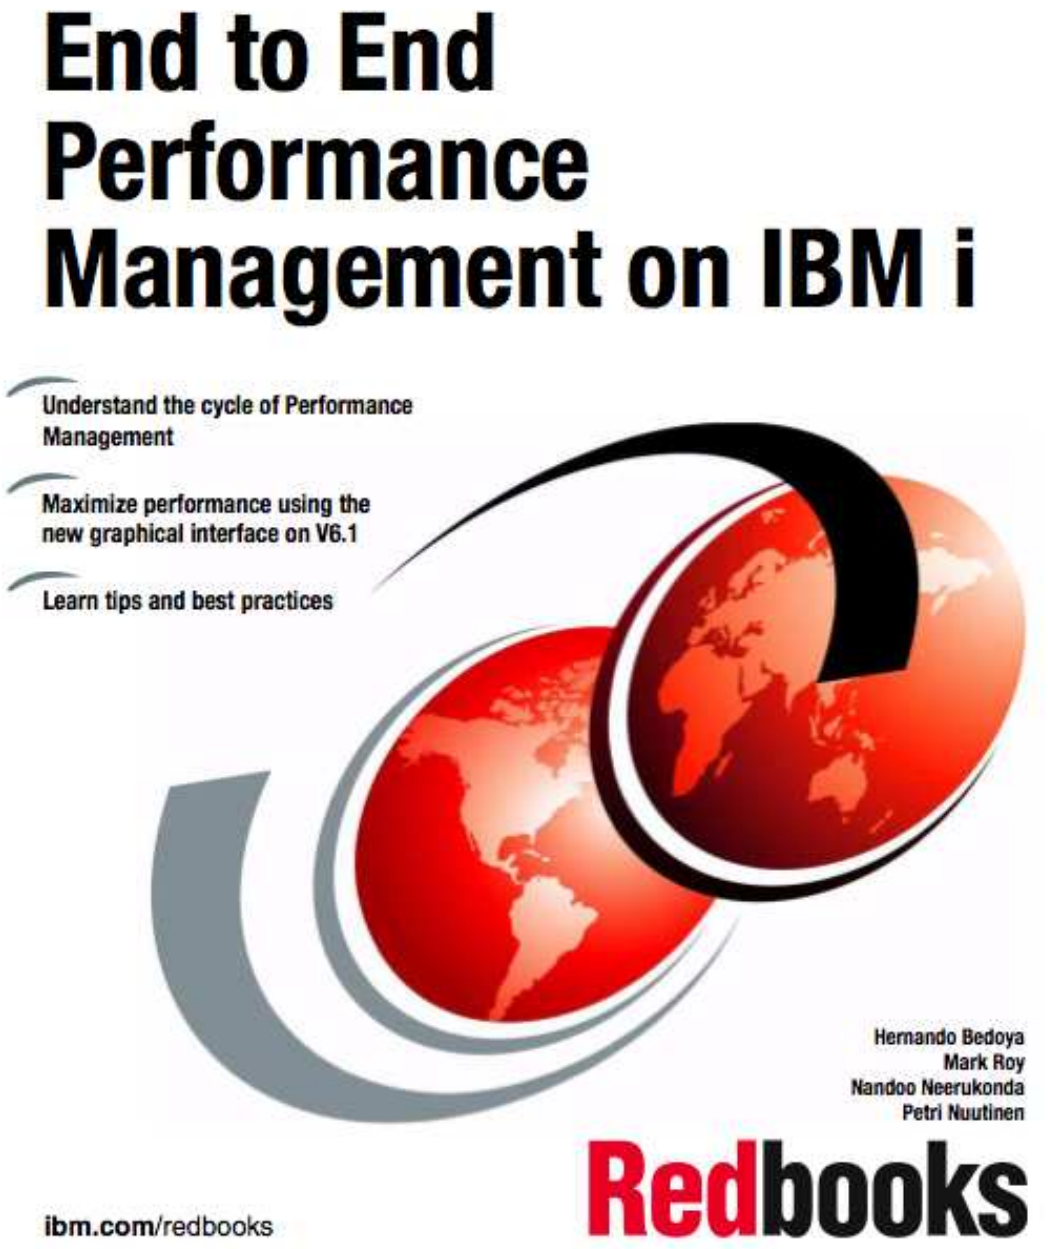

Power your planet.

Performance and Scalability Services

The IBM i Performance and Scalability Services Center can provide facilities and hardware **IN ROCHESTER** to assist you in testing hardware or software changes

"Traditional" benchmarks

Release-to-release upgrades

Assess and tune application and database performance

Stress test your system

Determine impact of application changes

Proofs of Concept (e.g. HA alternatives; SSD analysis, external storage, etc.)

Evaluate application scalability

Capacity planning

… all with the availability of Lab Services IBM i experts and development personnel

To request any of these services, submit at:

http://www-03.ibm.com/systems/services/labservices/psscontact.html

IBM i Web Sites with Performance Information

IBM i Information Center
Systems Management → Performance
http://publib.boulder.ibm.com/iseries/

IBM i Performance Management
http://www-03.ibm.com/systems/i/advantages/perfmgmt/
Performance Management for Power Systems
http://www-03.ibm.com/systems/power/support/pm/index.html
IBM Workload Estimator
http://www.ibm.com/systems/support/tools/estimator

iDoctor
http://www-912.ibm.com/i_dir/idoctor.nsf
Job Watcher Whitepaper
https://www-912.ibm.com/i_dir/idoctor.nsf/3B3C112F7FBE774C86256F4000757A8F/$FILE/Job_Waits_White_Paper_V5R4.pdf

IBM Systems Director
http://www-03.ibm.com/systems/management/director/index.html
IBM Tivoli Monitoring

Redbooks and Redpapers on IBM i Performance Tools

IBM eServer iSeries Performance Management Tools
http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4026.html?Open
A Systems Management Guide to Performance Management for System i and System p servers
http://www.redbooks.ibm.com/abstracts/sg247122.html?Open
Sizing IBM i5/OS Work on IBM System i5 Partitions
http://www.redbooks.ibm.com/abstracts/sg246656.html?Open
Application and Program Performance Analysis Using PEX Statistics
http://www.redbooks.ibm.com/abstracts/sg247457.html?Open
Managing OS/400 with Operations Navigator V5R1 Volume 5: Performance Management
http://www.redbooks.ibm.com/abstracts/sg246565.html?Open
IBM iDoctor iSeries Job Watcher: Advanced Performance Tool
http://www.redbooks.ibm.com/abstracts/sg246474.html?Open

Best Practices for Managing IBM i Jobs and Output (and a few other special tips)
http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4454.html?Open
i5/OS Diagnostic Tools for System Administrators: An A to Z Reference for Problem Determination
http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg248253.html?Open
IBM eServer iSeries Systems Management Handbook
http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4070.html?Open

Blog Posts

http://ibmsystemsmag.blogs.com/i_can/

http://ibmsystemsmag.blogs.com/i_can/performance/

http://ibmsystemsmag.blogs.com/i_can/2010/07/i-can-collect-more-performance-data-in-71.html

http://ibmsystemsmag.blogs.com/i_can/2010/05/-i-can-measure-disk-response-times.html

http://ibmsystemsmag.blogs.com/i_can/2010/03/i-can-understand-scaled-cpu-time.html

http://ibmsystemsmag.blogs.com/i_can/2010/02/i-can-use-power7-features-with-ibm-i-611.html

Articles

IBM Systems Magazine, IBM i - “Sky High Performance “, Aug 2009
http://www.ibmsystemsmag.com/ibmi/august09/coverstory/26021p1.aspx

SystemiNetwork - “Performance Data Investigator Consolidates Functions in One Place”, June 2009
http://systeminetwork.com/article/performance-data-investigator-consolidates-functions-one-place

SystemiNetwork - “IBM Systems Director Navigator for i: Performance Tasks Overview”, June 2009
http://systeminetwork.com/article/ibm-systems-director-navigator-i-performance-tasks-overview

IBM Systems Magazine, IBM i – “A Command Performance”, Nov 2008
http://www.ibmsystemsmag.com/ibmi/november08/administrator/22426p1.aspx

IBM Systems Magazine, IBM i - “Introducing IBM Systems Director Navigator for i5/OS”, Aug 2008
http://www.ibmsystemsmag.com/ibmi/august08/administrator/21503p1.aspx

IBM Systems Magazine, IBM i – “A Collective Effort”, Nov 2006
http://www.ibmsystemsmag.com/ibmi/november06/trends/7201p1.aspx

IBM Systems Magazine, IBM i - “Mission: Performance Management”, Nov 2006

Articles on Job Watcher

iDoctor versus IBM i 6.1 Performance Tools
http://www.ibmsystemsmag.com/i5/october08/trends/21990p1.aspx

“Web Power”
http://www.ibmsystemsmag.com/i5/november08/administrator/22431p1.aspx

Introduction to Job Watcher Green Screen Commands
http://www.ibmsystemsmag.com/i5/november08/tipstechniques/22521p1.aspx

Top 10 Hidden iDoctor Gems
http://www.ibmsystemsmag.com/ibmi/enewsletterexclusive/23868p1.aspx

Using iDoctor for iSeries Job Watcher to Determine Why Jobs Wait
http://www.ibmsystemsmag.com/ibmi/october05/technicalcorner/8896p1.aspx

Articles on Disk Performance

A New Way to Look at Disk Performance
http://www.ibmsystemsmag.com/i5/may07/administrator/15631p1.aspx

Analyzing Disk Watcher Data
http://www.ibmsystemsmag.com/i5/may08/tipstechniques/20662p1.aspx

Using Wait State Accounting to Determine Disk Performance
http://www.systeminetwork.com/artarchive/20700/index.html

Understanding Disk Performance, Part 2: Disk Operation on i5/OS
http://www.systeminetwork.com/artarchive/20870/Understanding_Disk_Performance__Part_2__Disk_Operation_on_i5_OS.html

Understanding Disk Performance, Part 3: Metrics of Disk Performance
http://systeminetwork.com/article/understanding-disk-performance-metrics

A Look at System i Integrated DASD Configuration and Performance under i5/OS

Back-up

Additional Wait Accounting Information

Run/Wait Signature

Typical batch job run/wait signature

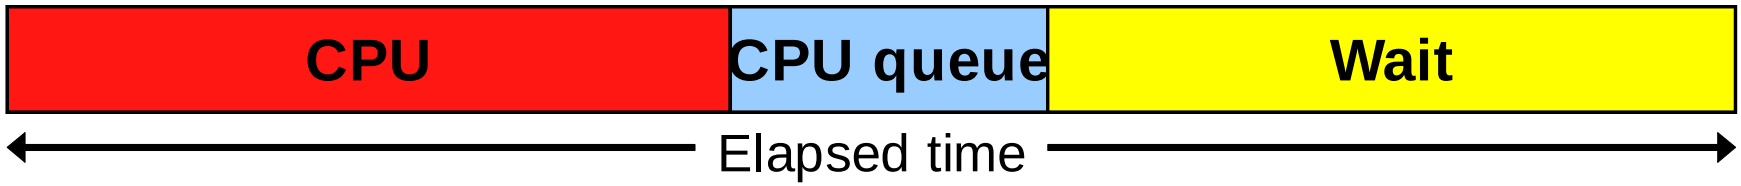

Interactive job run/wait signature

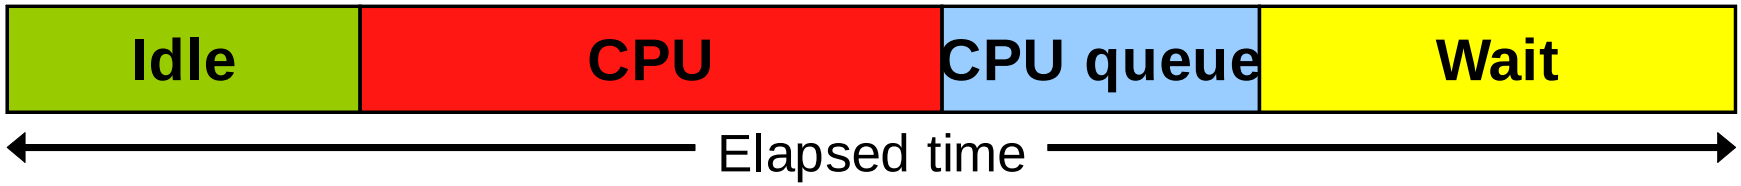

Detailing wait time

Determine the components of time spent waiting

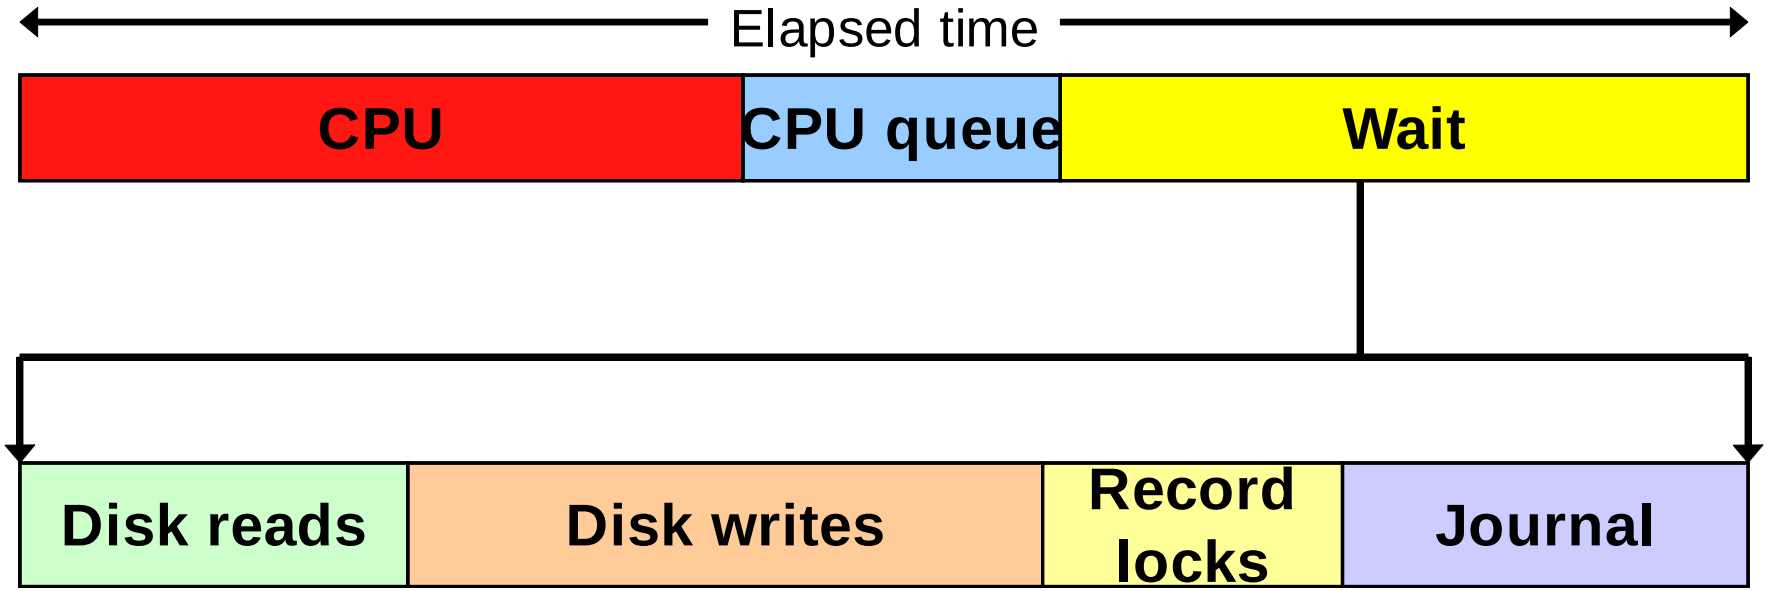

Detailing wait time - metrics

Then get metrics related to these components of wait time

	Disk reads	Disk writes	Record Locks	Journal
Total count	3,523	17,772	355	5,741
Total time	42 sec	73 sec	45 sec	44 sec
Avg time per wait	0.012 sec	0.004 sec	0.126 sec	0.007 sec

Even at this level, we can already see possible questions to ask:

How many of the reads are page faults? Could memory/pool changes help?

What programs are causing reads? Could they be reduced or made async?

What programs are causing writes? Could they be reduced or made async?

What DB2 files are involved with the record locks?

What files are being journaled? Are journals needed and optimally configured?

Waits Overview – Collection Services Data

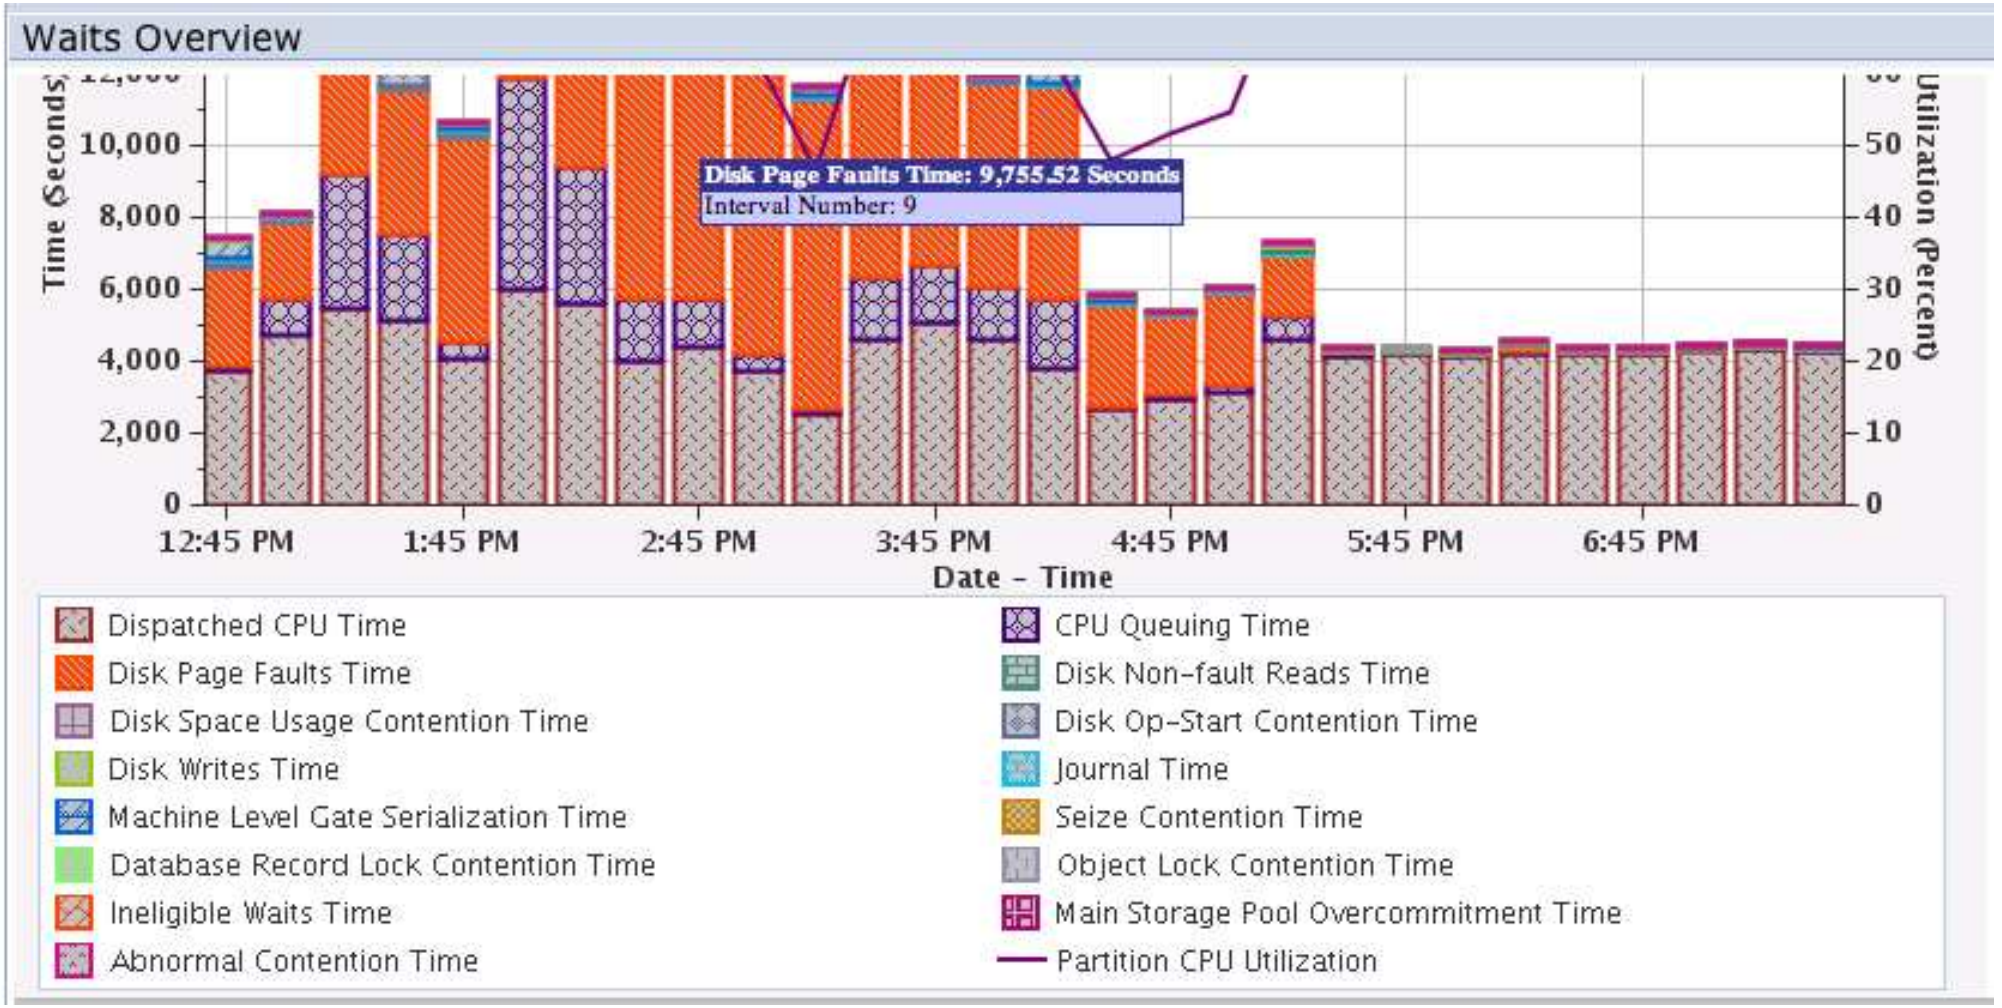

Power your planet.

Special notices

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

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

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

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

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

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

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

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

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

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

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

Revised September 26, 2006

Special notices (cont.)

Revised February 9, 2010