

Lloyd R. Perera  
lloydperera@yahoo.com

# Partition Resource Utilization & IBM i Performance

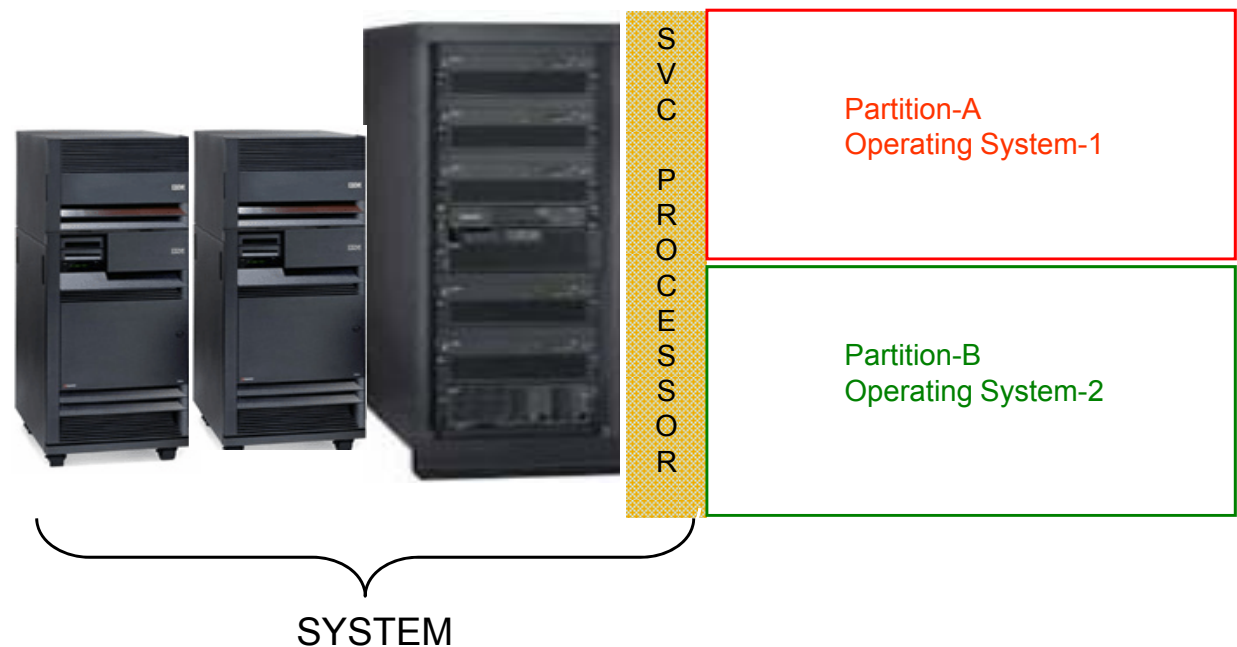


# Abstract

- **System Management**
  - Performance Analysis and Capacity Planning are an integral part of System Management that enables the IT department to meet its service level commitments to users.
- **Capacity & Performance**
  - Performance reviews and estimates of capacity requirements should ensure that adequate hardware capacity is available to meet the infrastructure demands, but it is important to ensure that the users' need for adequate response time for interactive requests and batch runtimes are satisfied.
- **Resource Utilization**
  - An appreciation of the impact of hardware resources like processor capacity and speed, memory and disk on performance expectations of the users will assist in resolving performance issues.
- **Upgrade**
  - Understanding the impact of your hardware resources will help you make better decisions in the acquisition and deployment of required hardware.


# What's a Partition?

- All the hardware resources controlled and managed by a single instance of an Operating System (like IBM i)
  - Often referred to as a Logical Partition or LPAR
- Different from a "system"
  - Single controlling Service Processor
  - Single CPU serial number
  - One or more partitions (logical)



# Computer Work Types

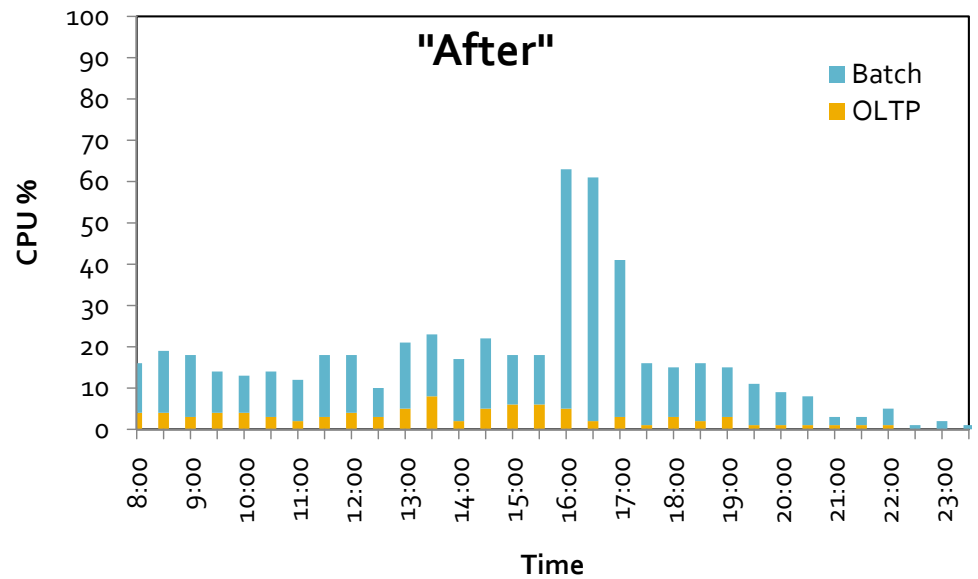

- Working with a computer
  - Interactive or Online Transaction Processing (OLTP)
    - On-going interaction between users and applications
      - “Conversational”
  - Batch/Non-interactive
    - work submitted for processing
      - Single request
      - User does not continue interacting
- Traditional IBM i reference
  - “Interactive”
    - 5250-OLTP using 5250 data-stream
  - “Batch”
    - All non-5250
      - Non-interactive, batch jobs
      - Servers Jobs - respond client requests
        - OLTP
        - Batch



# Hardware Goals

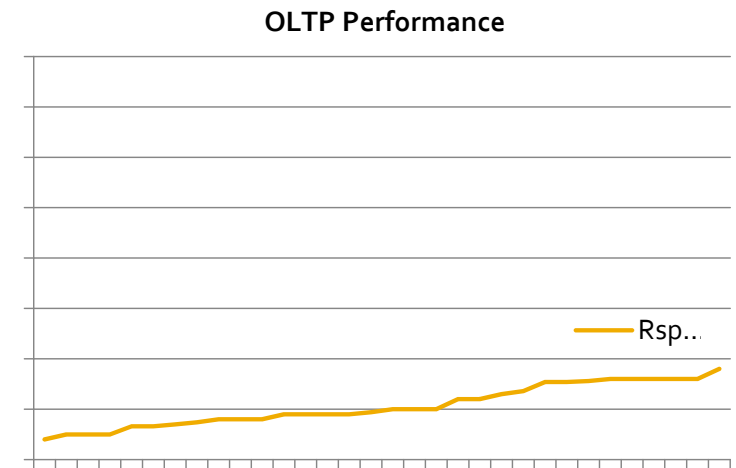
- Percentage Utilization of Capacity
  - CPU
  - Memory
    - Faulting rates
  - Disk
    - Drives busy
    - Capacity occupied
- Performance
  - Interactive-OLTP (5250 & Client/Server) requests
    - Response time
      - Time between request & result
  - Non-OLTP/batch
    - Elapsed time
      - Time between request submission & completion
        - Waiting to start (in job queue)
        - Actual run time

# Capacity Utilization?



# Good Performance?

- Currently experiencing “good” performance
  - Acceptable response time/batch run time
- Current Resource Utilization?
  - Increases in user activity
    - Unexpected peaks
  - Change in environment
    - New release of OS or Application
    - Increased functionality



**Consider Performance & Utilization!**

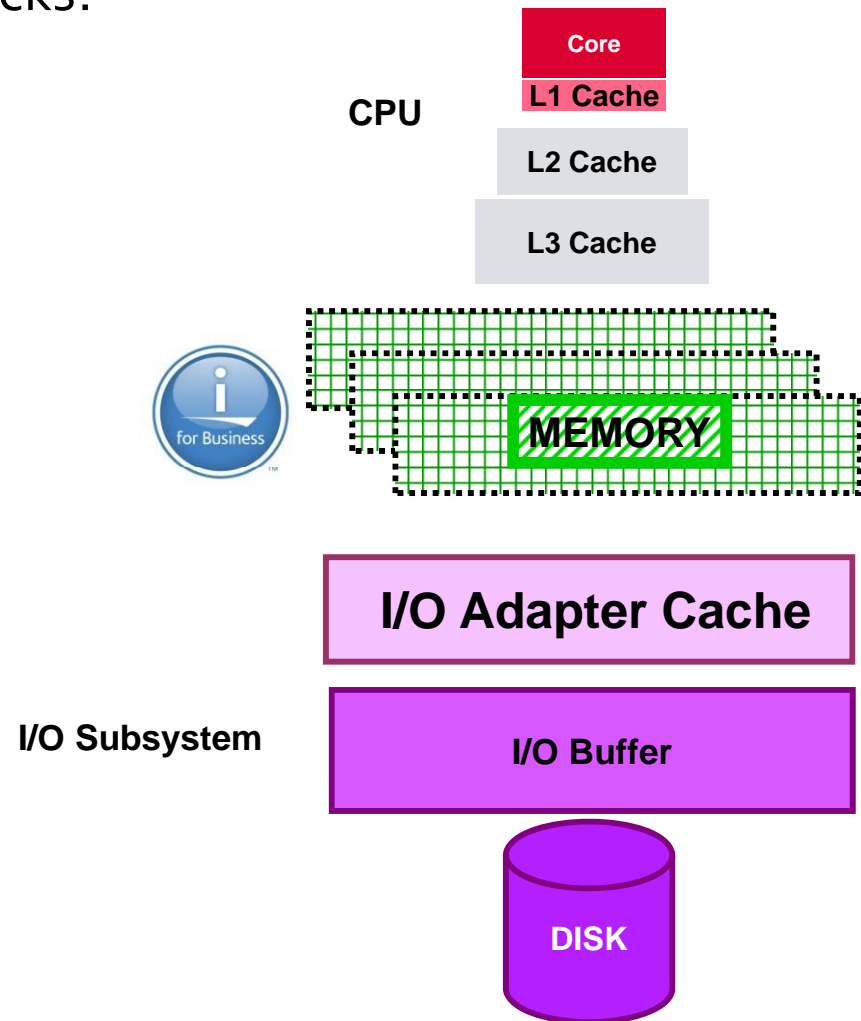
# Capacity (=throughput)

- IBM i CPW (Commercial Processing Workload)
  - Capacity of CPU only - in an unconstrained environment
    - Core ("processor")
      - Speed - GHz rating
      - Number of parallel active cores
      - Cache size/speed
        - L1/L2/L3
    - Parallel processing of threads
      - Hardware Multithreading
      - Simultaneous Multithreading
        - POWER5 and later

**CPW may or may not provide a close approximation for customer environment.**

# Capacity of a partition

- Hardware – balanced, without bottlenecks?
  - CPU
  - Memory
  - Buses
  - I/O Adapters
  - Disk etc
- IBM i Operating System Version
- Customer environment
  - Applications
    - I/O or Processor bound
  - Resource contention
    - Seizes/Locks
  - Workload peaks
  - User behaviour



# Components of Performance (Response)

- IBM i response time (or runtime) – excluding communication times
  - CPU time
  - Disk (Synchronous) time
  - Memory faulting time
  - Conflict/wait time
    - Seizes
    - Locks

## Example #1

Response Time	= <b>1.10</b> secs
• CPU time	= <b>0.52</b> secs
• Disk (50 I/O @ 0.010 secs)	= <b>0.50</b> secs
• Waits	= 0.08 secs

## Example #2


Response Time	= <b>1.10</b> secs
• CPU time	= 0.02 secs
• Disk (50 I/O @ 0.004 secs)	= 0.20 secs
• Waits	= <b>0.88</b> secs

# Resource Utilization Thresholds

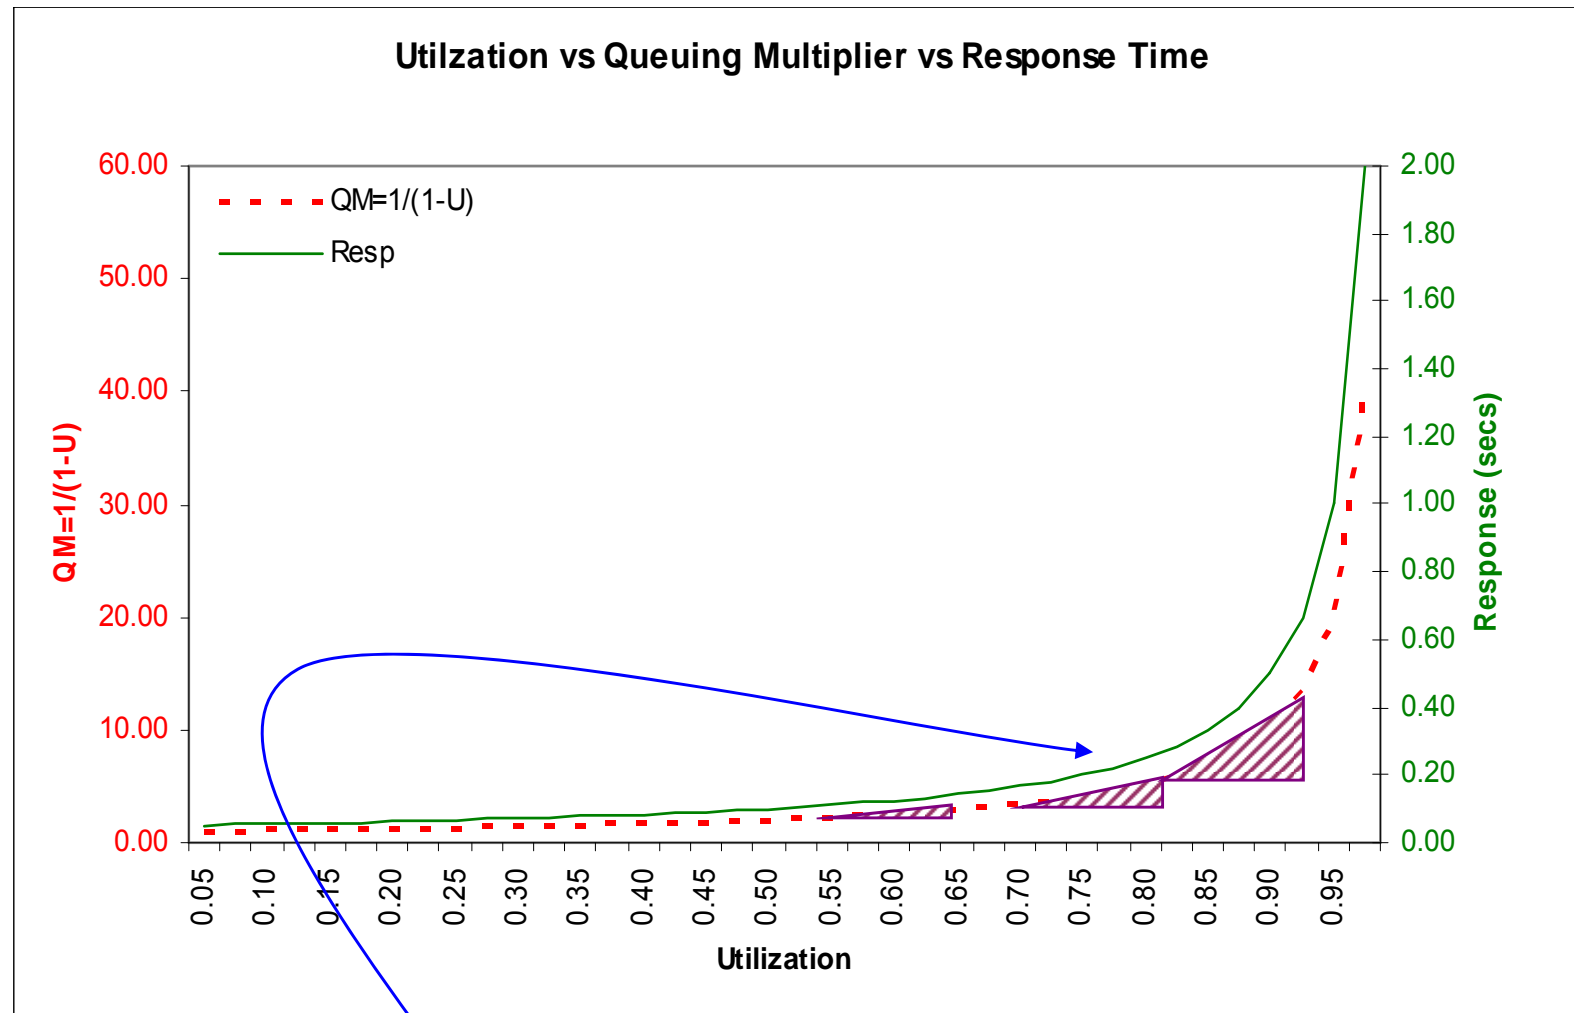
- Resource Utilization
  - Queuing : high utilization-longer queue
    - CPU
    - Disk
  - Reaction time
    - Memory faulting
    - Disk space used /disk space available
      - Extent of fragmentation
  - Disk Capacity Growth
- Thresholds of utilization
  - Impact to user perception
    - Significant or not
  - Potential for future crises
    - Upgrade decision delay

# Queuing

- Requests to any Provider of Service (Server)
  - Arrive -> Wait->Processed->Leave
- Response Time (T)
  - Average time to service (S) each request
  - Wait time (W) in queue
    - Determined by queue length (Q)
- Queuing Theory
  - Queue Length =  $U/(1-U^n)$
  - Queuing Multiplier - a measure of the impact of queuing
    - $QM=1/(1-U^n)$
  - Response Time=Service Time \* Queuing Multiplier
  - $T=S*QM$



# Queuing vs Response time



**"Knee" = rapid rate of change of QM / Response Time**

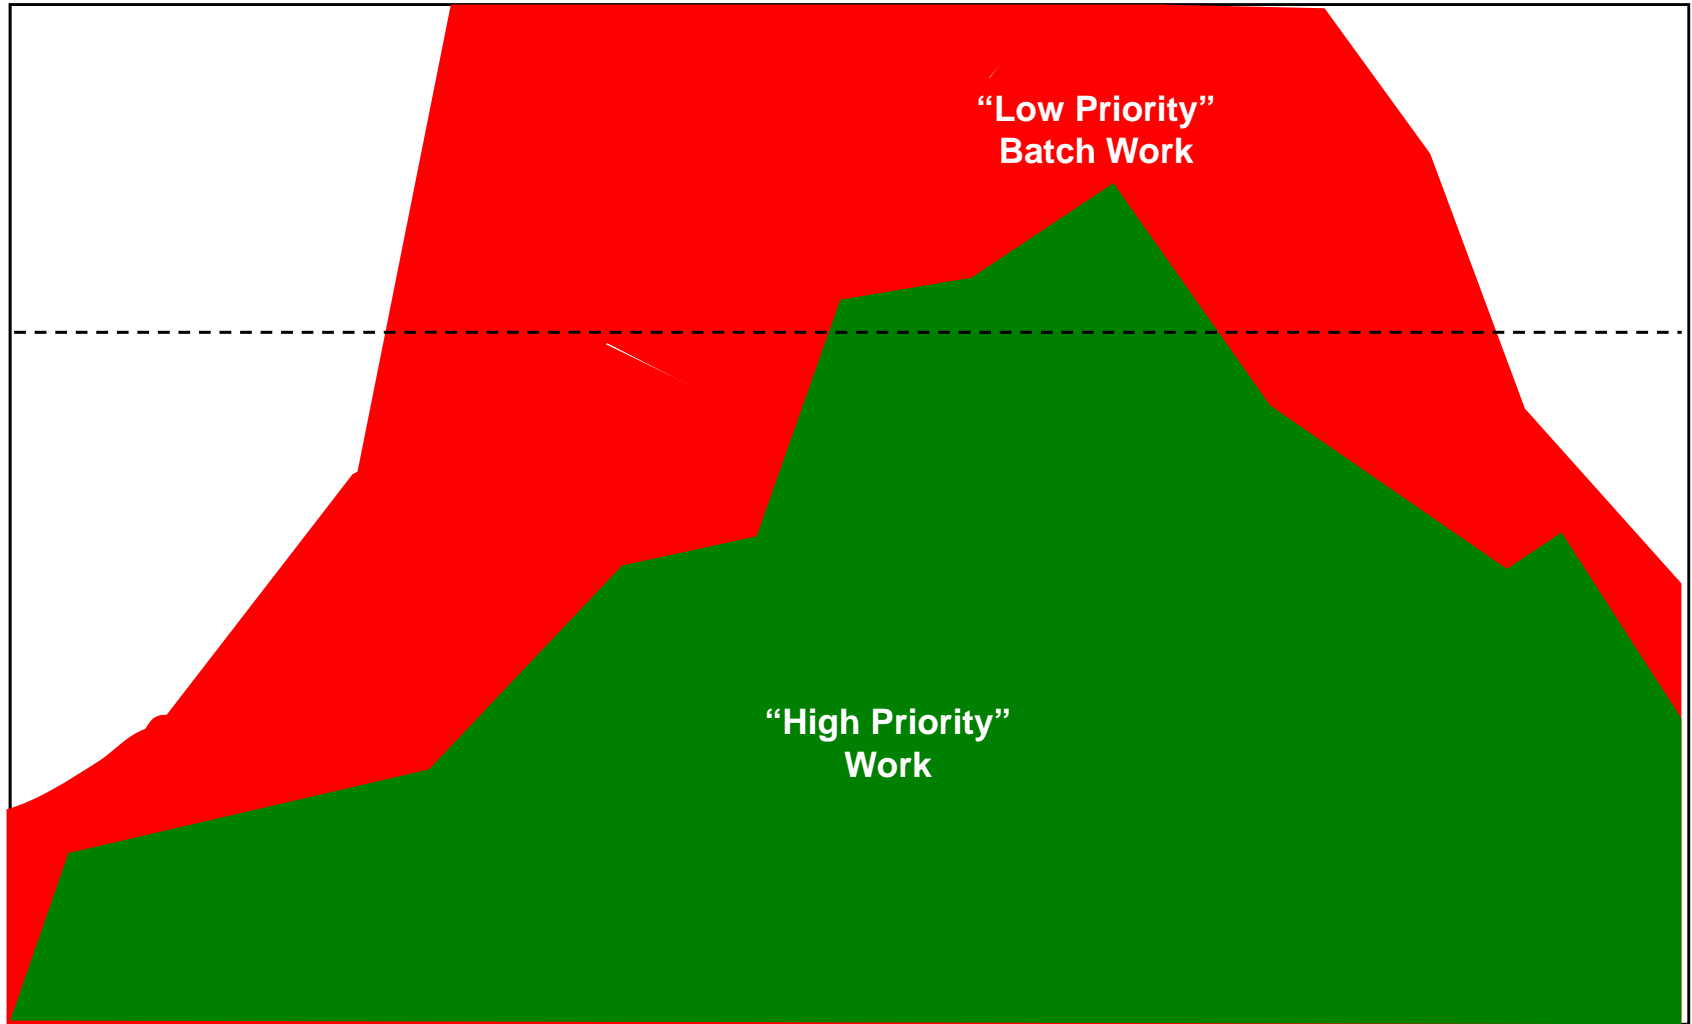
**□ change in Utilization results in a significant increase in Queuing/response**

# CPU Utilization by Priority

**Impact of Queuing  
On CPU component  
Of response time**

- 93% (32-way)
- 91% (24-way)
- 90% (18-way)
- 87% (12-way)
- 85% ( 8-way)
- 81% ( 4-way)
- 76% ( 2-way)
- 71% ( 1-way)

CPU%



**Watch CPU Utilization by HIGH PRIORITY work .....Not total ONLY (peaks & duration)**

**Performance objectives satisfied ?**

- Batch runtime
- OLTP response time

# Example – Impact of CPU Queuing

- CPU Response Time = QM \* CPU Service Time
  - Overall impact of CPU queuing on USER response time
    - Review CPU per OLTP request
      - Small CPU use per request = low impact

CPU		
Service Time (secs) S	Queuing Multiplier QM	Resp Time (secs) T
0.05	3	0.15
0.50	3	1.50

Depends on Service Time per request too!

# Benefit of faster CPU

- Faster CPU
  - Improves CPU component only
- Higher CPU capacity (CPW)
  - Increased throughput
- Comparative CPU "Speed"
  - $\sim$ Total CPW / # of cores
- "IBM i" response time
  - CPU Time + Synchronous Disk Time + Other
    - (Synch Disk Time = Sync Disk I/O \* Disk Response Time)

	CPU	CPU time	Disk time	Total
Example-1	current	0.70 secs	0.50 secs	1.20 secs
	2*faster	0.35 secs	0.50 secs	0.85 secs
Example-2	current	0.04 secs	0.50 secs	0.54 secs
	2*faster	0.02 secs	0.50 secs	0.52 secs



Faster CPU will not always improve user perceived response time!

# Impact of Memory Faulting

- Faulting “guidelines”
  - Machine Pool - Less than 10 per second
  - User Pools - ?
- Impact of a Memory Fault
  - Time to service a fault  $\approx$  time for one disk I/O
- Effect on user response time? Depends .....
- faulting rate
- number of requests
- disk response time

## Consider 1000 faults per sec generated by

- 10 requests: delay per request = time to service 100 faults
- 1000 requests: delay per request = time to service 1 fault

## Consider 50 faults per request

- Disk resp = 40 msec: delay per request = 2000 msec
- Disk resp = 4 msec: delay per request = 200 msec

“Guideline”:  $(100 * \text{CPU util})$  per core

# Example – Memory Faulting

- Utilization
  - CPU Util = 36%
  - Processing Units = 28
- “Rule of Thumb”
  - # of Processors \* Utilization \* 100
  - $28 * 0.36 * 100 = 1006$  faults/sec
- But measured ....
  - Faults/sec = **2875** 😞
  - Response Time = 0.19 seconds



**“Eye Catcher”**

~~Rule of Thumb:~~ (100 \* CPU util) per core

# Disk Performance – Frequently used gauge

```

Work with Disk Status
                                XXXXXXXX
                                01/14/08 15:55:38
Elapsed time: 00:05:42

Unit  Type  Size  %  I/O  Request  Read  Write  Read  Write  %
      (M)  Used  Rqs  Size (K)  Rqs  Rqs  (K)  (K)  Busy
  1  6713  8589  92.1  1.8    4.3    1.6   .1   4.3   4.0    0
  2  6713  8589  92.2  2.4    4.3    2.2   .1   4.3   4.0    3
  3  6713  8589  92.2  2.3    5.0    2.0   .2   5.0   4.7    1
  4  6713  8589  92.2  3.4    5.1    3.2   .2   5.1   4.8    1
  6  1750  35165 12.6   .2     .8     .1   .0    .0    .0    0
  6  1750  35165 12.6   .4     .9     .3   .0    .0    .0    0
                                Bottom

Command
===>
F3=Exit  F5=Refresh  F12=Cancel  F24=More keys

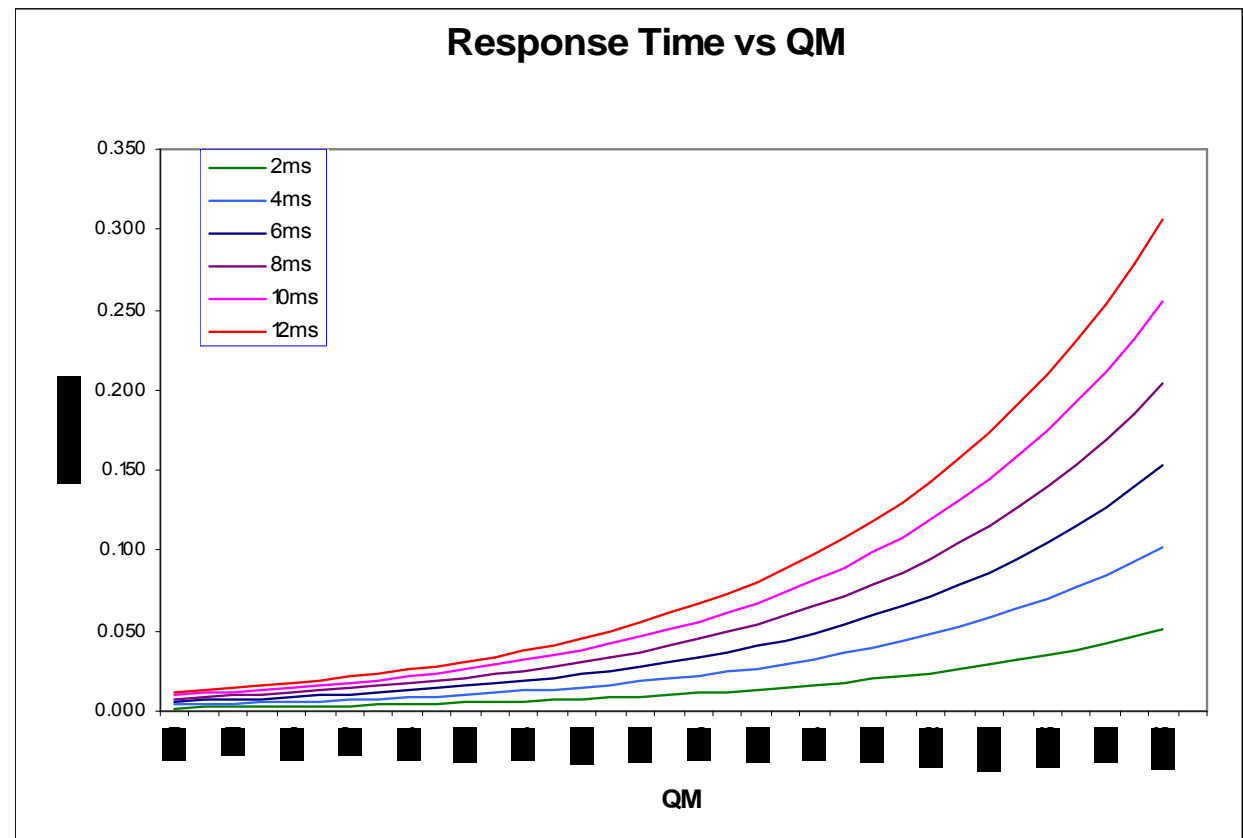
```

# Consider - Disk Response Time

Unit	Unit Name	Type	Size (M)	IOP Util	IOP Name	Dsk CPU Util	--Percent-- Full Util		Op Per Second	K Per I/O	- Average Time Service Wait		Per I/O -- Response
ASP ID/ASP Rsc Name: 1/													
0001	DD013	4328	132,309	.0	CMB01	10.3	43.2	17.1	44.99	55.0	.0038	.0092	.0130
0002	DD016	4328	132,309	.0	CMB01	10.3	43.2	16.7	44.27	58.3	.0037	.0081	.0118
0003	DD004	4328	132,309	.0	CMB01	10.4	43.2	17.1	47.10	54.1	.0036	.0074	.0110
0004	DD015	4328	132,309	.0	CMB01	10.4	43.2	17.3	48.54	53.5	.0035	.0077	.0112
0005	DD012	4328	132,309	.0	CMB01	10.4	43.2	16.7	45.06	56.1	.0037	.0081	.0118
0006	DD007	4328	132,309	.0	CMB01	10.4	43.2	16.3	43.14	58.4	.0037	.0083	.0120
0007	DD002	4328	132,309	.0	CMB01	10.3	43.2	17.1	53.19	48.9	.0032	.0070	.0102
0008	DD006	4328	132,309	.0	CMB01	10.3	43.2	16.7	43.19	59.1	.0038	.0084	.0122
0009	DD020	4328	123,488	.0	CMB01	10.4	43.2	15.9	39.38	59.5	.0040	.0076	.0116
0010	DD003	4328	123,488	.0	CMB01	10.4	43.2	15.9	41.09	58.1	.0038	.0075	.0113
0011	DD018	4328	132,309	.0	CMB01	10.3	43.2	17.0	47.14	54.5	.0036	.0079	.0115
0012	DD019	4328	123,488	.0	CMB01	10.4	43.2	15.8	40.24	59.9	.0039	.0077	.0116
0013	DD010	4328	123,488	.0	CMB01	10.3	43.2	15.7	39.80	58.8	.0039	.0078	.0117
0014	DD022	4328	132,309	.0	CMB01	10.3	43.2	16.7	44.33	57.8	.0037	.0085	.0122
0015	DD024	4328	132,309	.0	CMB01	10.4	43.2	16.6	45.09	57.1	.0036	.0081	.0117
0016	DD005	4328	132,309	.0	CMB01	10.4	43.2	16.9	44.76	57.8	.0037	.0086	.0123
0017	DD008	4328	132,309	.0	CMB01	10.4	43.2	17.0	47.55	54.5	.0035	.0070	.0105
0018	DD001	4328	132,309	.0	CMB01	10.4	43.2	16.4	42.88	59.3	.0038	.0084	.0122
0019	DD021	4328	132,309	.0	CMB01	10.3	43.2	16.1	43.46	58.0	.0037	.0086	.0123
:													
0026	DD050	4328	123,488	.0	CMB10	3.7	43.2	10.0	51.50	49.2	.0019	.0020	.0039
0027	DD051	4328	123,488	.0	CMB10	3.7	43.2	10.1	52.09	48.9	.0019	.0031	.0050
0028	DD052	4328	123,488	.0	CMB10	3.7	43.2	9.9	50.23	49.3	.0019	.0021	.0040
0029	DD053	4328	123,488	.0	CMB10	3.7	43.2	10.5	55.56	44.9	.0018	.0021	.0039
0030	DD054	4328	123,488	.0	CMB10	3.7	43.2	10.3	53.15	47.7	.0019	.0020	.0039
0031	DD055	4328	123,488	.0	CMB10	3.7	43.2	10.8	53.17	48.2	.0020	.0035	.0055
0032	DD056	4328	105,847	.0	CMB10	3.7	43.2	8.2	43.14	46.4	.0019	.0015	.0034
0033	DD057	4328	105,847	.0	CMB10	3.7	43.2	8.4	48.05	42.4	.0017	.0017	.0034
:													
Total for ASP ID: 1 4,516,140													
Average							43.2	14.2	46.23	52.9	.0030	.0058	.0088
Total			4,516,140										
Average							43.2	14.2	46.23	52.9	.0030	.0058	.0088

# Disk Performance

- Service Time \* Queuing Multiplier = Response Time
  - $S * QM = T$ 
    - $QM = 1 / (1 - U^n)$
    - U = Utilization

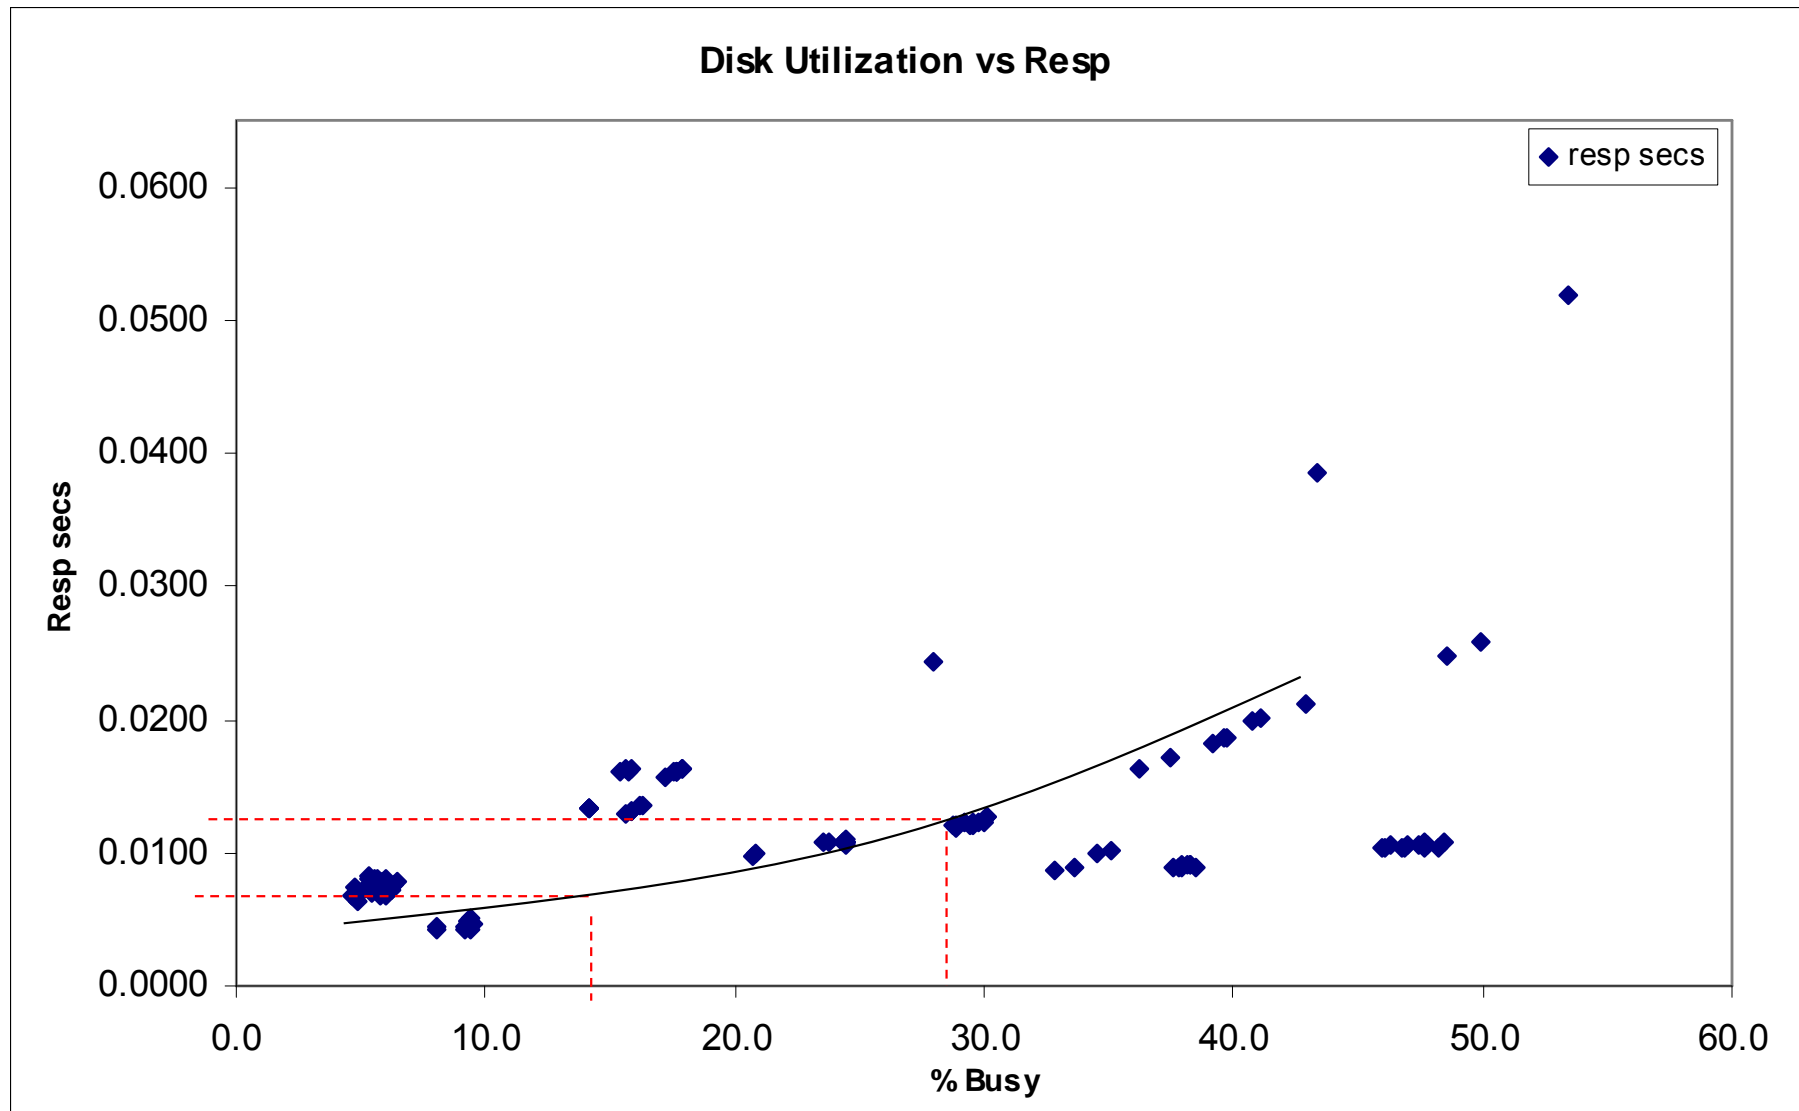


Faster disk can provide “good response” at higher levels of utilization

# Establish Disk Response Objectives

- OLTP Example
  - “System” response objective < 1.00 seconds
    - CPU secs per request = 0.25 secs
    - Disk secs per request < 0.75
      - Average Synchronous Disk I/O per request = 50
      - Disk response objective
        - $0.75/50 < 0.015$  seconds = **15 msec**
- Batch Example - High Disk I/O
  - Run time objective < 1000 seconds
    - CPU secs = 600 secs
    - Disk secs < 400 secs
      - Average Synchronous Disk I/O = 100,000
      - Disk response objective
        - $400/100,000 < 0.004$  seconds = **4 msec**

# Disk Utilization vs Disk Response Time



# Summary - Establish Performance Goals

- Time
  - User Response/Runtime
    - OLTP
    - Non-OLTP (batch)
- Components of User Response
  - CPU
  - Disk
  - Other Waits

# Summary – Review Current Data

- Resource Utilization
  - CPU
    - Utilization
      - Capacity for growth
      - Contingency for peaks
      - Impact of queuing – CPU secs per request
    - Memory
      - Faulting rate per request
      - Disk response time
    - Disk
      - Response time
      - Space
        - For growth
        - Reduce fragmentation
  - Performance
    - OLTP Response Time
    - Critical Batch Runtime
  - Consider Resource Upgrade Requirements

# Questions?

Thank you for your time & interest

If you have additional questions

[lloydperera@yahoo.com](mailto:lloydperera@yahoo.com)