

# **Disk Subsystem Capacity Management, Based on Business Drivers, I/O Performance Metrics and MASF**

Igor Trubin, Ph.D.

and Linwood Merritt

Capital One Services, Inc.

[igor.trubin@capitalone.com](mailto:igor.trubin@capitalone.com)

# Introduction: Environment

- Capital One
  - 6th largest card issuer in the United States
  - Capital One to S&P 500 in 1998
  - Fortune 500 company starting in 2000
  - Managed loans at \$71.8 billion
  - Accounts at 46.7 million
  - CIO 100 Award “Master of the Customer Connection”
  - Information Week “Innovation 100” Award Winner
  - ComputerWorld “Top 100 places to work in IT”

# The Capacity Management service


- ~1000 servers of different platforms such as
  - UNIX/Linux
  - NT/W2K
  - Tandem
  - Unisys
  - MVS
- Capacity of “Capacity Management” environment and SLA
  - a relatively small 4-way Unix server (ServerP) and several large SAS based applications should
  - provide daily web based reports of capacity and performance issues by **8 am**

Capacity Issue:  
the Capacity Management System  
needed to resolve its own capacity problem!

- SLA was broken, and the Capacity Planning web site was ready after **9 am**.
- Main reason:  
the growth in the number of servers.
- Main question:  
what subsystem needs to be upgraded?


# CPUs?

Before a “recent” upgrade the metric had reached only 80% and based on simple trend analysis, no capacity problem would occur for several months.




# DISK Subsystem ?

- SAS job is an I/O intensive workload and as shown on this chart, the Disk I/O metric had been growing as well
- The metric does not have a threshold, so, it's very hard to say this is a Disk subsystem capacity issue



# Which subsystem was upgraded?




Both charts show that an upgrade has happened and as a result, both metrics have dropped.



# “Busiest Disk utilization” (“Disk Busy” )

HP MeasureWare: *“the percentage of time during the interval that the busiest disk device had I/O in progress from the point of view of the Operating System.”*

**Busiest Disk Utilization before and after upgrade**



## Which subsystem was upgraded?

Indeed, older **disk** devices were **replaced** with faster RAID ones!



# The Presentation Objective:

**This presentation is an overview of Disk Subsystem metrics used for Capacity Management of the Capital One's large multi-platform server farm as well as discussions of how to use them to produce meaningful forecasts, simple modeling and statistical analysis.**

## Plan of the presentation:

- Introduction/Case Study - *done*
- Disk Subsystem Metrics Overview
- Disk Metric Trend Analysis and Forecast
- Overall Disk I/O Capacity Estimation
- Statistical Analysis of Disk Performance Data
- SUMMARY/ References

# Disk Subsystem Metrics Overview

- **File System Utilization**

- *Problem:* Capacity Management environment may not have the capacity to monitor and report capacity problems about all File Systems (hundred thousands).
- *Bad solution:* GLB\_FS\_SPACE\_UTIL\_PEAK (similar to “Disk Busy”) UNIX performance metric, which is  
*“the percentage of occupied disk space to total disk space for the fullest file system found during the interval.”*  
BUT (!) The file system that has OS or other UNIX system files is always “almost full” !

# Disk Subsystem Metrics Overview


- *Better solution: Concord eHealth* performance monitor system has interesting metric “**System Health Index**” which is the sum of five components (variables):
  - **SYSTEM**, which reports a CPU imbalance problem;
  - **MEMORY**, which is exceeding some memory utilization threshold or reflects some paging and/or swapping problems;
  - **CPU**, which is exceeding some utilization threshold;
  - **COMM.**, which reports network errors or exceeding some network volume thresholds;
  - And **STORAGE**, which might be a combination of
    - a. Exceeding user partition utilization threshold;
    - b. Exceeding system partition utilization threshold;
    - c. File cache miss rate, Allocation failures and
    - d. Disk I/O faults problem that can add additional points to this Health Index component.

# Disk Subsystem Metrics Overview

- Example of “System Health Index” from *Concord eHealthh*:

- **STORAGE**  
component has the biggest contribution and demonstrates some bad trending.

- **partitions**  
#1 and #2 were highly utilized and caused a Health Index increase.



# Disk Subsystem Metrics Overview

- *BMC Patrol Perceive* about **File Systems** metrics:
  - Percent of file system that is full;
  - Size of file system in megabytes;
  - Measure of inodes used in the file system;
  - Number of inodes in the file system;
  - Amount of free space in the file system
  - Number of free inodes in the file system
  - Amount of file system space available that is allocated for general use

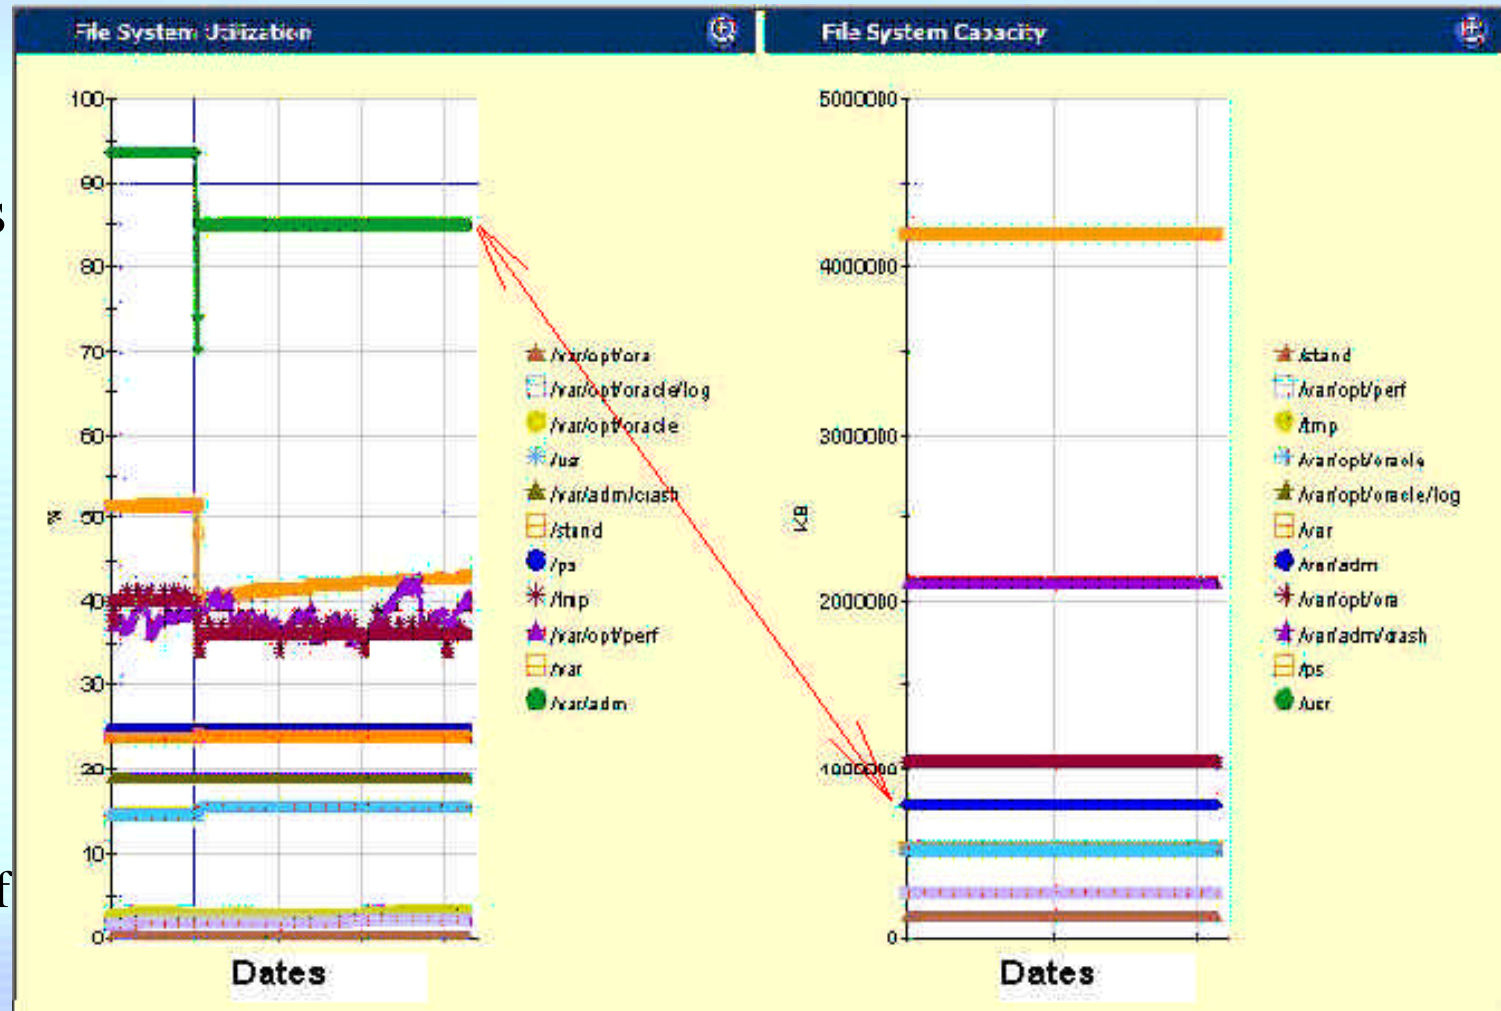


# Disk Subsystem Metrics Overview

- *BMC Patrol Perceive* report example:

Good combination is utilization and actual size of the file systems


Indeed, 1% free space of 100 GB disk is equal to 10% free space of a 10 GB disk.




# Disk Subsystem Metrics Overview

- **Disk I/O rate** is the number of physical I/Os per second during the interval.

Disk I/O rate trend for off- and work-hours and standard SAS “forecast” procedure based on one of the “time series” algorithms



24-hour profile grouped by nearest month hourly averages. Good picture to balance workload during the day






# Disk Metric Trend Analysis and Forecast

- More realistic future **Disk I/O rate** trend example:

SAS scripts should be adjustable to take in consideration upgrades, workload shifts or consolidations



Simple trend only — a detailed system model should be used for more accurate results

# Disk Metric Trend Analysis and Forecast



- **Health Index** trend analysis:

*Big advantage:*

- There is a real threshold

*Disadvantages:*

- The Disk subsystem is indirectly presented here;
- The future trend tries to predict future problems of different subsystems and sounds very suspicious as an “apples to oranges” comparison




# Disk Metric Trend Analysis and Forecast

- A performance data vs. business driver **correlation analysis**:

**Take** monthly business driver data (historical and projected) from business units within the company, **configure** each server to one or more business drivers, and **perform SAS** multivariate regressions against CPU utilization or disk I/O !

May 2004




Linear  $R^2 = .990$  Regression = Linear

Business driver list: Bus1 Bus2 Bus3 Bus4 Bus5 Bus6

# Overall Disk I/O Capacity Estimation

- Could we have a threshold for Disk I/O trend chart?

Based on HP MeasureWare DISK level data, there is the possibility to estimate overall disk subsystem I/O capacity.



# Overall Disk I/O Capacity Estimation

For the sample interval (5 min) HP MeasureWare log file had DISK utilization equaled to BYDSK\_UTIL,

Time	DEVICE_NAME	BYDSK_UTIL %	BYDSK_PHYS_IO_RATE IO/sec	DISK Capacity IO/sec
11:35	7/0/0.8.0.3.1.4.6	63.0	30.9	49.05
11:35	6/0/0.8.0.3.1.0.6	60.9	41.7	68.50
11:35	...	...	...	...

the rate of I/O was equaled BYDSK\_PHYS\_IO\_RATE.


The maximum of the I/O rate which would be executed if the disk was 100% busy is:

$$\text{DISK Capacity (IO/sec)} = \text{BYDSK\_PHYS\_IO\_RATE (IO/sec)} * 100 / \text{BYDSK\_UTIL (\%)}$$

Disclaimer: It is a very simple linear model and does not take in consideration the DISK queue and controller cache usage

# Overall Disk I/O Capacity Estimation

Yes, we have a I/O rate threshold for each Disk,  
but how to make the estimation across all Disks?



I/O Capacity: Available(calculated)

I/O Capacity: Used(The actual measured I/O rate)

		<i>Sum of ALL Disks IO Capacity IO/sec</i>			<i>Sum of ServerE: BYDSK_PHYS_IO_RATE IO/sec</i>
<i>Time</i>		<i>Total</i>	<i>Time</i>		<i>Total</i>
11:35		33,402.75	11:35		2,460.20
...		...	...		...
12:30		37,172.06	12:30		1605.79
12:35		25,957.33	12:35		2101.39
<i>Max capacity IO/sec</i>		37,172.06	<i>Max Actual IO/sec</i>		2,460.20
<i>Max capacity IO/hour</i>		133,819,406.86	<i>MAX Actual IO/hour</i>		8,856,719.97

Finally ServerE DISK IO

CAPACITY utilization is  $(\text{Max Actual IO/hour}) * 100 / (\text{Max capacity IO/hour}) = 6.62\%$

# Statistical Analysis of Disk Performance Data

- Another way to build a dynamic threshold of Disk I/O rate is SEDS - Statistical Exception Detection System based on Multivariate Adaptive Statistical Filtering (MASF) technique.
- SEDS is used for automatically scanning through large volumes of performance data and identifying measurements that differ significantly from their expected values.
- MASF is extension of Statistical Process Control or (Quality Control) which was developed by Walter Shewhart of Bell Telephone Laboratories in the 1920s.
- MASF procedure was designed and presented in CMG by BGS Systems, Inc. in 1995.
- SEDS is developed by this author and presented as the best paper in CMG 2002.






# Statistical Analysis of Disk Performance Data

- Review of the existing tools:

- **SAS/QC** (Quality Control);
- **JMP** from SAS;
- **BEZsystems**  
for Oracle and Teradata;
- **Concord eHealth** – DFN  
(Deviation From Normal)
- The Patrol Perform and Predict tool  
from **BMC software**;

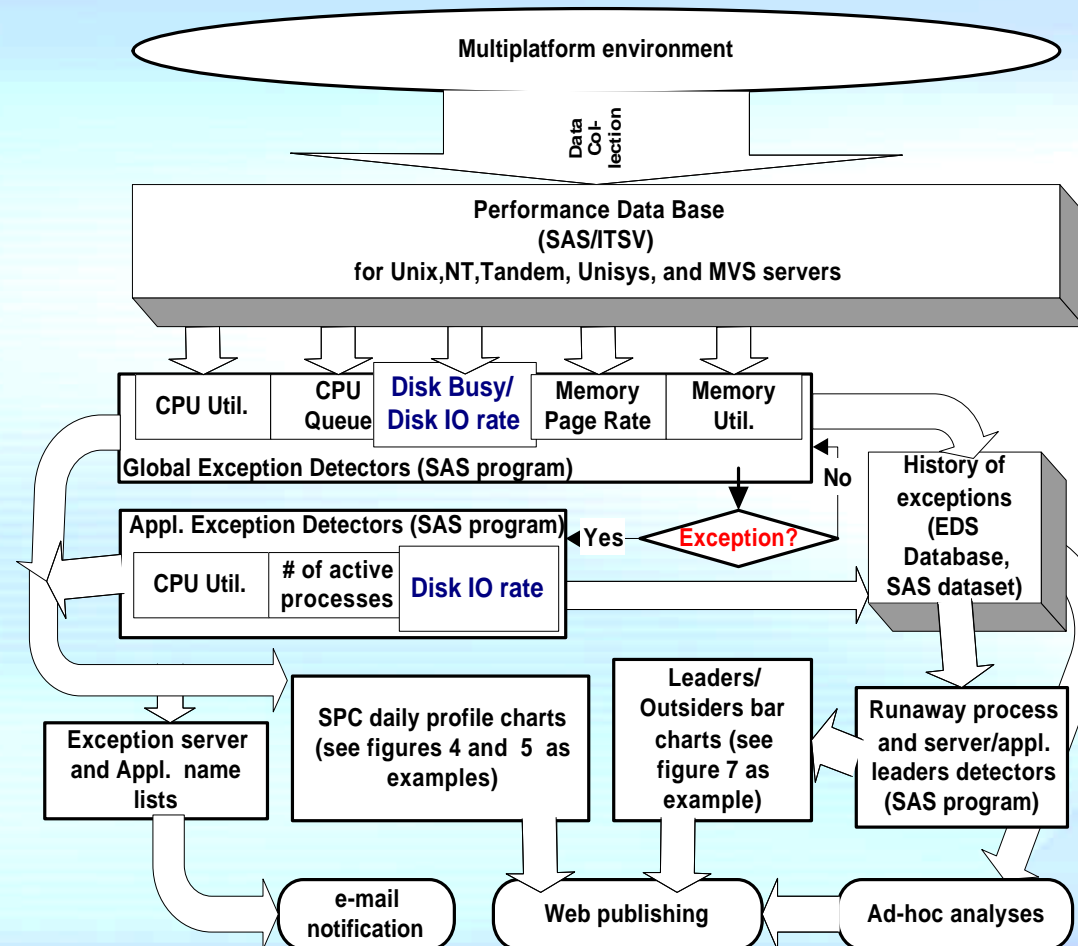
The common output is  
***Control charts*** for  
monitoring variations in  
process under statistical  
control



# Statistical Analysis of Disk Performance Data

- SEDS structure:

- Exception detectors for the most important metrics including Busiest Disk Utilization and Disk I/O Rate;
- SEDS Database with history of exceptions;
- statistical process control daily profile chart generator;
- exception server name list generator;
- Leader/Outsider servers detector and detector of runaway processes; and
- Leaders/Outsiders bar charts generator.




# Statistical Analysis of Disk Performance Data

- SEDS implementation:
  - **Performance database (PDB):** SAS/ITSV; BMC Visualiser Database
  - **Home made programs:** SAS 8.2; Unix scripting (awk/sed/perl); VisualBasic.NET/SQL;
  - **Reporting:** Intranet web server; HTML, Email
  - **Special features:**
    - a.* Two level exception estimation: Global and Application.
    - b.* statistical exception alerts (e-mail notification);
    - c.* spetial database to keep history of exceptions
  - **The rules to avoid taking into consideration:**
    - a.* noise (collector errors, runaway processes);
    - b.* insignificant exceptions (like slight increases of workloads for underutilized servers);
    - c.* other insignificant patterns, based on the analyst's interpretation.

# Statistical Analysis of Disk Performance Data

- DISK I/O Control Chart for Web Publishing:


The full  
"7 days X 24 hours"  
adaptive filtering  
policy is applied to  
calculate the  
average, upper, and  
lower statistical  
limits of a  
particular metric for  
each weekday for the  
past six months.



# Statistical Analysis of Disk Performance Data

- Application Level DISK I/O Control Charts

- SEDS captured a Disk I/O rate exception at about 4:00 PM on **ServerB**,
- and the Application detector found that the Workload **“Appl2”** had an exception as well.







# Statistical Analysis of Disk Performance Data

Server name	Sub Business	Config: OS/Server model/Qty/Item	Perform. Charts	CPU Exceptions	Memory Exceptions	Disk Exceptions
				CPU Util. Run Queue	Mem. Util. Page Rate	Disk I/O Util. Disk I/O rate
server10	Operations	HP-UX11.0/HP/2240/2/4/5	BMC MW	0-0 4-0	0-0 0-0	0-0 0-0
server1	Call Center	Solaris 2.6/Sun/E480/2/4	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0
server2	Call Center	Solaris 2.6/Sun/E480/2/4	BMC MW	4-0 0-20	0-0 0-0	0-0 4-0
server3	Call Center	HP-UX11.0/HP/N440/4/6	BMC MW	0-0 5-0	0-0 0-0	0-0 0-0
server11	Call Center	Solaris 2.6/Sun/E420/4/4	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0
server12	Call Center	HP-UX11.0/HP/7400/4/4	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0
server20	Call Center	HP-UX11.0/HP/E4000-55/4/4	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0
server13	Credit	HP-UX11.0/HP/E380/4/6	BMC MW	2-0 0-0	0-0 0-0	0-0 0-0
server15	Call Center	HP-UX11.0/HP/7400/4/2	BMC MW	2-0 0-0	0-0 0-0	0-0 0-0
server16	Credit	HP-UX11.0/HP/7400/6/9	BMC MW	0-0 0-0	5-0 0-0	0-0 0-0
server17	Production Services	HP-UX11.0/HP/7400/4/2	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0
server18	Production Services	HP-UX11.0/HP/V2400/4/6	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0
server4	Production Services	MVS/IBM/2064-1C9/1	BMC MW	0-0 0-0	0-0 0-0	0-0 0-0

- System performance daily web report based on EDS database




'Disk IO' for 06/18/03  
Limits and Mean based on every Wednesday for past six months  
Machine= ServerB



# Statistical Analysis of Disk Performance Data

- **ExtraVolume** is the numeric estimation of the exception magnitude

- It calculates **the area** between the limit curve and the actual data curve (for periods when the exceptions occurred).
- **Physical meaning** is the number of I/Os the server has taken that exceeds a standard deviation.






# Statistical Analysis of Disk Performance Data


- TOP I/Os Leaders Charts (ExtraIOs>0)

- The system automatically produces **ExtraIOs calculation** for the last day and records that in the SEDS database.
- This data is used for generating **Leaders/Outsiders charts** for the last day, last week, last month, and publishing the bar charts



# Statistical Analysis of Disk Performance Data


- Overall company wide picture of all servers that had Disk I/O exceptions
- The colored “**Treemap,**” or “**heat chart.**” has been already used to publish an overall capacity status
- SEDS produces the similar chart **for IO exceptions;** here the ServerB is presented as pretty large red box inside of “M Department”, because the unusual I/O usage was bigger than 40,000,000



# Statistical Analysis of Disk Performance Data

- History of exceptions can give very interesting data for a trend analysis

- This is **history of unusual Disk I/Os** on ServerB for the last two weeks.
- The **disk performance issue was escalating** and the server fell into the "Top 10" server list and then the issue was addressed and resolved.



# SUMMARY

- **Understand the metrics.** There can be a large amount of data, from different sources. The Capacity Planner must first know which metrics are captured, and understand reporting and analysis nuances around the metrics.
- **Forecast demand.** This presentation has discussed the use of trend analysis and business driver based forecasting to predict future demand.
- **Determine capacity thresholds** for action. This presentation discusses the calculation of maximum I/O rates as well as a method using Statistical Process Control concepts.
- **Reporting.** This presentation gives examples of utilization and trend charts, exception reporting, “Top 10” reporting, and “Treemap” heat charts..

# References

- Merritt, Linwood, "Capacity Planning for the Newer Workloads," Proceedings of the Computer Measurement Group, 2001
- Merritt, Linwood, " Seeing the Forest AND the Trees: Capacity Planning for a Large Number of Servers," Proceedings of the United Kingdom Computer Measurement Group, 2003
- Shneiderman, Ben, "Treemaps for space-constrained visualization of hierarchies," <http://www.cs.umd.edu/hcil/treemaps>, December 26, 1998 and November 8, 2000
- Trubin, Igor, Ph. D. and Mclaughlin, Kevin, "Exception Detection System, Based on the Statistical Process Control Concept," Proceedings of the Computer Measurement Group, 2001
- Trubin, Igor, Ph. D., "[Global and Application level Exception Detection System, Based on the MASF Technique,](#)" Proceedings of the Computer Measurement Group, 2002

Thanks!

Igor Trubin  
IT Capacity Planning  
Capital One Services, Inc.  
[igor.trubin@capitalone.com](mailto:igor.trubin@capitalone.com)