



CMG'07 paper

SYSTEM MANAGEMENT BY EXCEPTION: The Final Part

Igor Trubin, PhD; Ray White

09/27/2007

Introduction:

- **QUESTION 1: How to manage a large number of computer systems?**

ANSWER: Applying statistical methods for capturing different types of unusual computer system behaviors and then address them

BENEFIT: No need to look at each individual system!

- **QUESTION 2: Is that a new idea?**





ANSWER: Not really.

First let's look how some companies has already been implementing this idea

Introduction: Known implementations

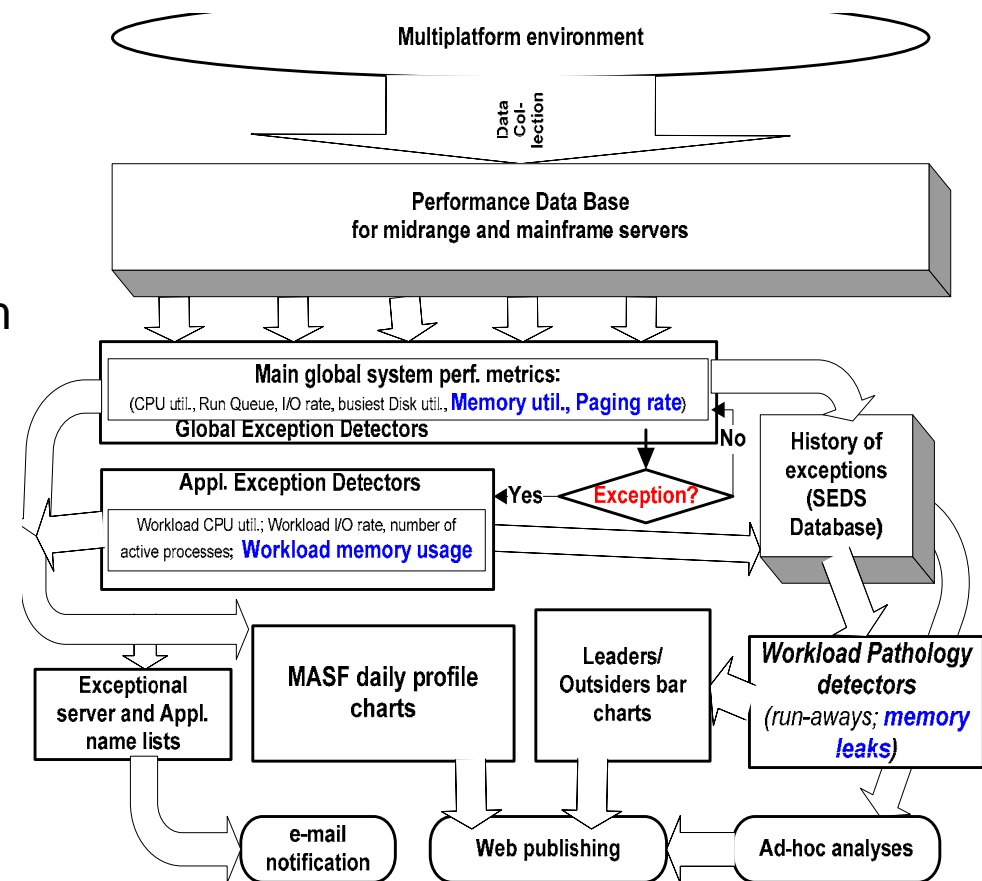
of statistical methods for capturing unusual computer system behavior:

- **BMC software** (www.bmc.com): MASF technique in Performance Analysis for Servers and Performance Assurance tools: documents.bmc.com/products/documents/51/79/65179/65179.pdf
- **CA Concord** (www.ca.com/concord/) eHealth/Fault manager tools: *DFN* (*Deviation From Normal*) www.concord.com/download/pdf/CG_fm_lh.pdf
- **Fujitsu** (www.fujitsu.com): *ACTIVE BASELINING* Technique www.fujitsu.com/downloads/AU/active_baselining_in_passive_data_environments.pdf
- **BEZ systems** (www.bez.com) for Oracle and Teradata performance www.wmoug.org/bezPresentation.pdf
- **ProactiveNet** (www.proactivenet.com): BSM and APM tools with *Abnormality Indicator*: www.freepatentsonline.com/20060217822.html
- **MonoSphere** (www.monosphere.com): *Detection of Outlier Events* www.cmg.org/membersonly/2006/papers/6105.pdf
- **McAfee** (www.mcafee.com) *Anomaly-Based Intrusion Detection* www.mcafee.com/us/local_content/white_papers/wp_ddt_anomaly.pdf
- OTHERS: **BEA** (www.bea.com), **Bank Of Montréal**, (www.bmo.com),
- **YOUR COMPANY ? - WHY NOT?**





Background: Statistical Exception Detection System (SEDS)

We developed and implemented an application called the **Statistical Exception Detection System (SEDS)** to proactively capture and notify SMEs of any statistically unusual server behavior. The system was built based on known concepts (**Statistical Process Control or Quality Control and Multivariate Adaptive Statistical Filtering**) and allows us to be focused on real performance and up-coming capacity issues and enable us to manage the capacity usage of the large computer farm using only a **few planners**.



How SEDS works: Control and Top 10 Charts

- 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.
- To get detailed information of the server's behavior for the previous day, the system publishes a **control chart**
- To produce advanced capacity planning analyses, the author expanded on MASF using new derived metrics. A history of exceptions is kept in a separate exception database. The new metrics include those such as "amount of exceptions per day" or "**ExtraVolume**" (here Extra Memory Pages)
- For the most significant exceptions, a bar chart generator reports the **top 10 servers**



How SEDS works: Web Report

Server Name	Sub Business	Config: OS/Vendor/Model/Qty/Year	Perform. Charts	CPU Exceptions	Memory Exceptions	Disk Exceptions
server14	Call Center	Solaris2.6/Sun/E450/2/4	BMC MW	CPU_Usr: 0-0, Run_Queue: 0-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 0-0	
server0	Call Center	Solaris2.6/Sun/E450/2/4	BMC MW	CPU_Usr: 0-0, Run_Queue: 0-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 0-0	
server40	Call Center	Solaris2.6/Sun/E450/2/4	BMC MW	CPU_Usr: 0-0, Run_Queue: 0-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 0-0	
server1	Call Center	HP-UX11.0/HP/N4000/8/8	BMC MW	CPU_Usr: 0-0, Run_Queue: 2-0, Mem_Usr: 2-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 0-0	
server11	Call Center	Solaris2.6/Sun/E420/4/4	BMC MW	CPU_Usr: 0-0, Run_Queue: 2-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 2-0	
server12	Call Center	HP-UX11.0/HP/7400/4/4	BMC MW	CPU_Usr: 0-0, Run_Queue: 2-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 2-0	
server20	Call Center	HP-UX11.0/HP/N4000-55/4/4	BMC MW	CPU_Usr: 0-0, Run_Queue: 0-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 0-0	
server13	Credit	HP-UX11.0/HP/N5500/4/5	BMC MW	CPU_Usr: 0-0, Run_Queue: 0-0, Mem_Usr: 0-0, Page_Rate: 0-0	Disk_I/O_Usr: 0-0, Disk_I/O_Rst: 0-0	
server41	Production Services	MVS/IBM/2064-1C9/				
server44	Production Services	MVS/IBM/2064-1C9/				
server17	Production Services	HP-UX11.0/HP/7400/4/2				
server49	Production Services	MVS/IBM/2064-1C9/				
server4	Production Services	MVS/IBM/2064-1C9/				

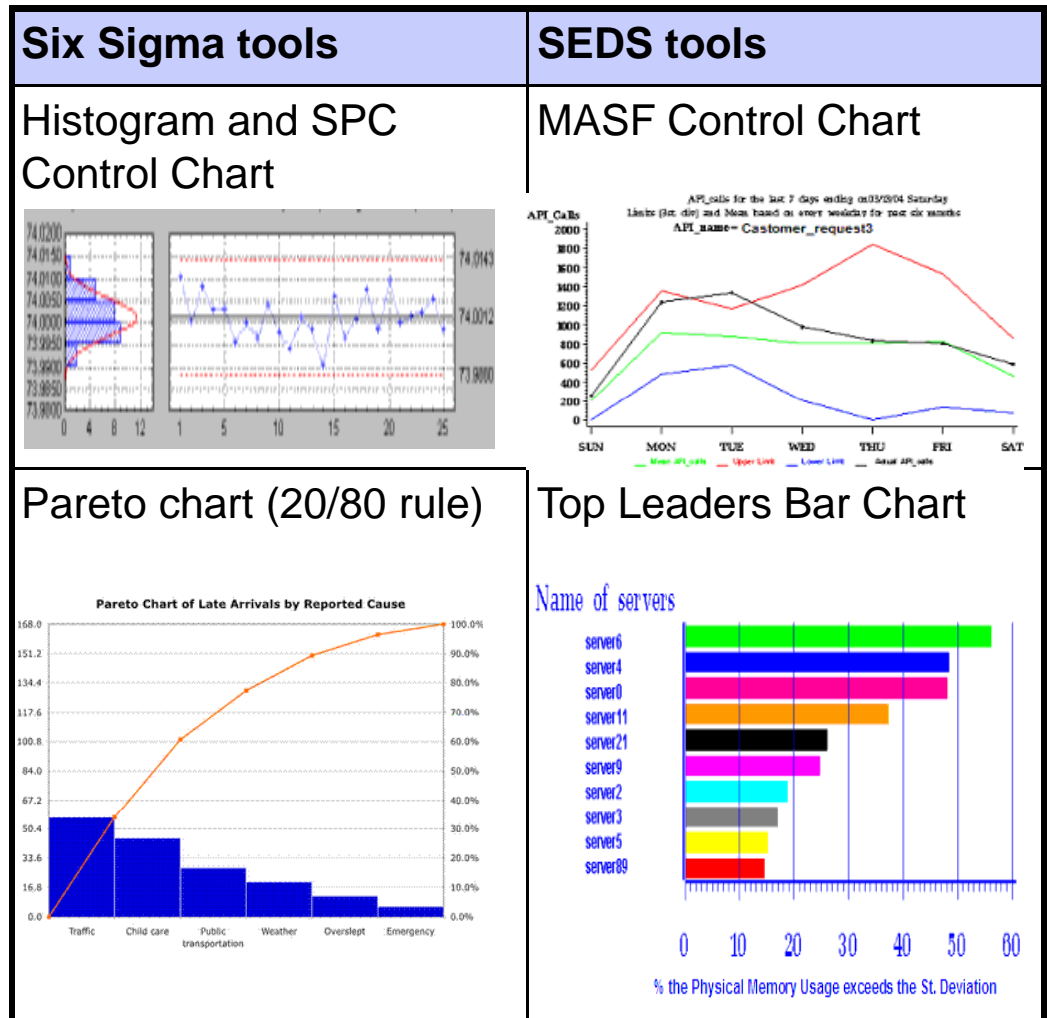
SEDS and Lean Six Sigma

LEAN – Eliminating waste

- A part of LEAN could be reducing the **list of metrics** to a small, manageable size – one that includes only those that are absolutely essential for Capacity Management. (see slide 10 with a solution)
- SEDS can be an effective tool as part of LEAN as that allows the analyst to **avoid wasting time manually scanning the entire server farm**. It allows the analyst to focus on the few important exceptional servers in terms of performance or upcoming capacity issues


Six Sigma - Eliminating defects

- SEDS has some similarity with Six Sigma methodology as it uses **similar statistical mechanisms**.
- Exception could be considered as defect!



MASF, SPC control and histogram charts comparison

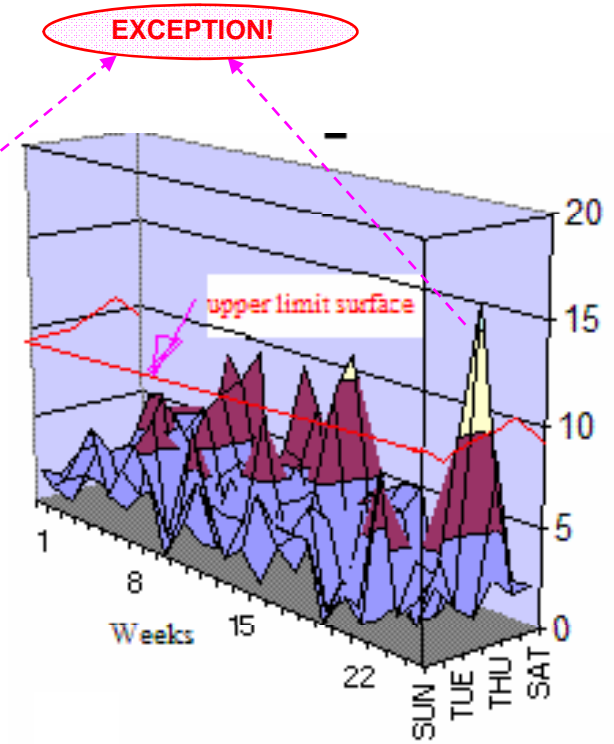

- All three charts demonstrate different views of exceptions for CPU utilization that occurred at 8 am.
- As opposed to classical control charts, **MASF charts** can be most useful for showing a 24 hour profile of CPU usage.



MASF and SPC control charts comparison

Weekly Health Index Control Chart Builder:

	A	B	C	D	E	F	G	H
1	Health INDEX by	SUN	MON	TUE	WED	THU	FRI	SAT
2	Past weeks data observations	19	14	23	15	00	11	04
3		12	02	16	20	00	12	09
4		20	20	11	30	20	20	21
5		40	40	20	38	39	53	40
6		54	40	40	66	55	40	52
7		21	10	40	70	33	28	19
8		37	20	22	21	20	35	30
9		53	65	21	38	43	92	60
10		39	04	45	69	75	81	94
11		01	16	00	34	41	33	05
12		24	41	02	00	00	14	23
13		74	20	43	16	22	21	74
14		39	28	32	40	57	99	55
15		12	30	80	40	52	80	37
16		20	00	40	38	08	20	21
17		52	22	07	47	93	111	03
18		40	31	28	20	22	43	42
19		39	50	41	41	32	42	48
20		40	00	02	52	43	40	49
21		01	20	07	02	01	00	02
22		20	02	00	16	05	20	20
23		18	80	13	00	00	08	00
24		80	30	05	20	36	40	80
25		20	44	44	14	06	23	20
26		Last week data	3.0	4.0	10.5	16.3	7.6	1.8
27	Upperlimit	9.4	8.8	8.3	9.2	10.5	13.3	11.2
28	Mean	3.2	2.6	2.4	3.1	2.9	4.0	3.4
29	Lowerlimit	0.0	0.0	0.0	0.0	0.0	0.0	0.0
30	StdDeviation	2.0	2.1	2.0	2.0	2.5	3.1	2.6



The MASF chart is a 2-D cut of a 3-D set of 24 (or 7 for weekly view) classical control charts; these correspond to the 24 hours in a day (7 days in a week)

The metrics list that absolutely essential for Capacity Management

Server Subsystem to Monitor	Level of Performance data	Metric Name	Good for SEDS ? [reference to other SEDS papers]
ALL		Server size (TPM, TPC)	No
CPU	global	CPU power/imbalance (SPEC)	No
	global	CPU and LPAR numbers	No
	global	CPU utilization	YES [1]
	global	CPUs imbalance	No
	global	CPU queue length	YES [5]
	application	CPU utilization	YES [1]
	application	Number of active processes	YES [5]
MEMORY	global	Physical Size (capacity)	No
	global	Memory utilization	YES [this paper]
	global	Paging (or Scan) rate	YES [this paper]
	global	Swap space utilization	Maybe [this paper]
	application	Resident Memory Utilization (Private + Shared)	YES [this paper]
DISK	global	Physical Size (space capacity)	No
	global	I/O rate	YES [2]
	global	Busiest Disk Utilization (I/O bandwidth)	YES [2]
	global	File System space utilization	Maybe [2]
	application	I/O rate	YES [2]
NETWORK	interface	Bandwidth Utilization (In and Out)	YES (this presentation and [5])
	Interface	Traffic volume (bytes/c In and Out)	YES

The SEDS morning e-mail report example for Network devices

E-mail notification report with links to SEDS charts:

Untitled - Message (Plain Text)

File Edit View Insert Format Tools Actions Help

Reply Reply to All Forward

Statistical Exception Detection System -SEDS- Network Performance Exception Report for 09/13

List of interfaces with daily maximum of BW utilization >49%

interface	daily mean	daily max
devswt82-RH-6/6-m	33.32%	49.22%
dmzcor03-2/6	45.81%	54.12%



Bandwidth_Utilization exceptions of 5 network interfaces and links to MASF control chart:

rtr102-POS2/1	http://it/par/content/eds/thelast/BWutil_SEDSchart.0.0913.gif
mls01-1/1	http://it/par/content/eds/thelast/BWutil_SEDSchart.1.0913.gif
mls01-3/1	http://it/par/content/eds/thelast/BWutil_SEDSchart.2.0913.gif
rtr01-Serial0/0:0	http://it/par/content/eds/thelast/BWutil_SEDSchart.3.0913.gif
rtr102-POS2/2	http://it/par/content/eds/thelast/BWutil_SEDSchart.4.0913.gif

NOTE: The 1st exception in the list above is the most severe, the last one is the least severe.
 ++++++

NB: Dear user, note please:
 NB: The SEDS Top 10 elements with unusually HIGH BW usage:
 NB: for the 1 last day: http://it/par/content/eds/leaders/SEDS_network_1_day_leaders.gif
 NB: for the 3 last days: http://it/par/content/eds/leaders/SEDS_network_3_day_leaders.gif
 NB: for the 7 last days: http://it/par/content/eds/leaders/SEDS_network_7_day_leaders.gif
 NB: for the 30 last days: http://it/par/content/eds/leaders/SEDS_network_30_day_leaders.gif

NB: Exceptions are captured if actual yesterday data exceeded historical mean +/-3 st.dev,
 NB: plus some non statistical rules are applied like if actual data < 20%
 NB: or only 1 hour had a spike that is not considered as an exception.
 NB: To create Top 10 chart the mean +/- 1 st.dev rule is used only,
 NB: plus ExtraVolume metric is calculated and used to take in account magnitude of exceptions.
 NB: More detailed explanation can be found in the SEDS paper

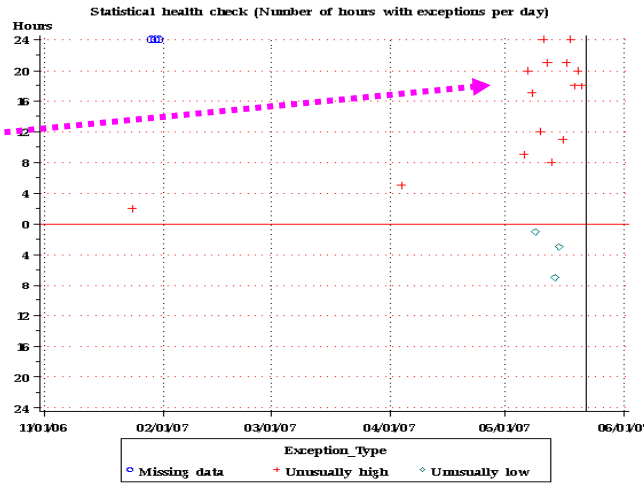


Note SEDS benefit: there are only 4 devices out of ~2000(!) need some attention!


The SEDS and Memory Metrics (Case study)

On one morning the TOP leaders chart showed **server9** with 25% unusual memory used:

- The exceptions “cloud” on the SEDS statistical health check chart shows a real change in the environment



- The classical trend also shows a memory usage spike not measured before

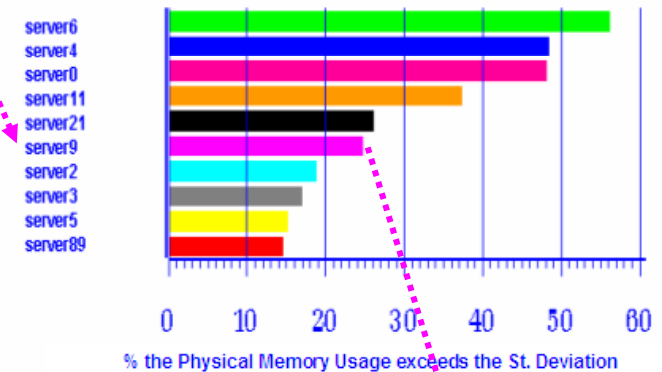


?? So what particular application caused this jump ??

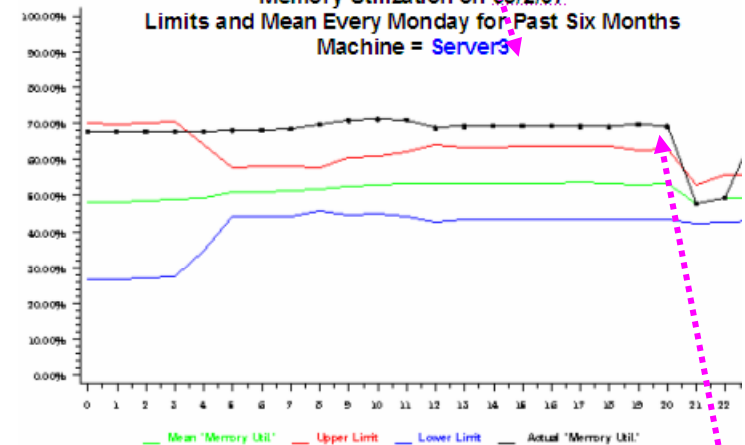
Simple trend only — a detailed system model should be used for more accurate results

Name of servers

TOP Servers
% of Physical Memory Size beyond 1 St. Deviation
For Last 1 days since 05/21/07



Memory Utilization on 05/2/07
Limits and Mean Every Monday for Past Six Months
Machine = Server3





- Memory utilization exceeded the statistical upper-limit for 18 hours.

The SEDS and Memory Metrics (*Application level data*)


- Fortunately a special memory metric for application level data is available for this server: “Resident Memory Utilization (Private + Shared)”

APP_MEM_RES “The size (in KB) of resident memory for processes in this group that was alive at the end of the interval. This consists of text, data, stack, as well as the process' portion of shared memory regions (such as, shared libraries, text segments, and shared data).” – *HP OV Performance Agent Technical Guide*

- The rest of the applications like **AppsOK** on the server did not have a significant deviation from normal memory usage.



- only **AppsN** application memory usage exceeded the upper-limit for 17 hours .

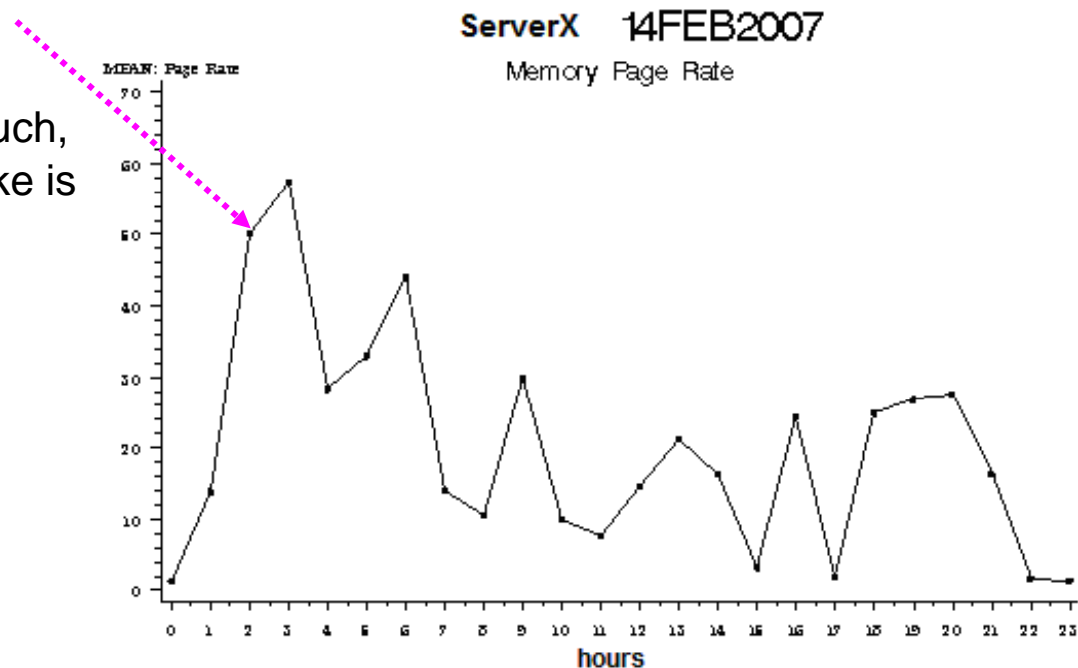


- CONCLUSION: SEDS** can detect what workload is responsible for global exception.

The SEDS and Memory Metrics (*Paging exceptions*)

- Some modern Operating Systems (AIX and Linux) and some large applications (MS SQL) grab almost all memory to allow use by other subsystems or parts of applications. (E.g. file cache).
- This is a smart and wise way to use memory capacity but the performance collector agent has to report the global memory utilization at about 100% all the time. Thus this metric becomes useless.
- If application level data is not available there is only one way to capture unusual memory usage – to analyze some other global metric such as the Memory Paging Rate

This metric has the following problem: there is no simple calculated threshold and, as such, it is hard to say if the 2 am spike is big enough to worry about

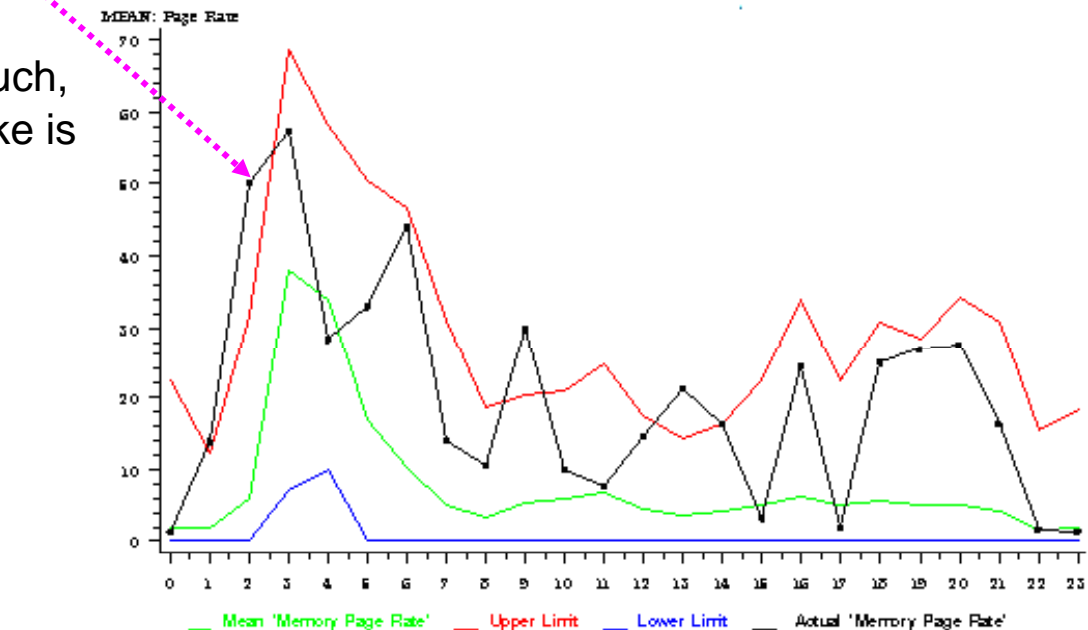


The SEDS and Memory Metrics (*Paging exceptions*)

- Some modern Operating Systems (AIX and Linux) and some large applications (MS SQL) grab almost all memory to allow use by other subsystems or parts of applications. (E.g. file cache).
- This is a smart and wise way to use memory capacity but the performance collector agent has to report the global memory utilization at about 100% all the time. Thus this metric becomes useless.
- If application level data is not available there is only one way to capture unusual memory usage – to analyze some other global metric such as the Memory Paging Rate

This metric has the following problem: there is no simple calculated threshold and, as such, it is hard to say if the 2 am spike is big enough to worry about


'Memory Page Rate' for 02/14/07
Limits and Mean based on every Wednesday for past six months
Machine=ServerX




The SEDS and Memory Metrics (*Paging exceptions*)

- Some modern Operating Systems (AIX and Linux) and some large applications (MS SQL) grab almost all memory to allow use by other subsystems or parts of applications. (E.g. file cache).
- This is a smart and wise way to use memory capacity but the performance collector agent has to report the global memory utilization at about 100% all the time. Thus this metric becomes useless.
- If application level data is not available there is only one way to capture unusual memory usage – to analyze some other global metric such as the Memory Paging Rate

The control chart shows unusual paging activity. That is confirmed by reviewing the historical paging trend:



'Memory Page Rate' for 02/14/07
Limits and Mean based on every Wednesday for past six months
Machine=ServerX

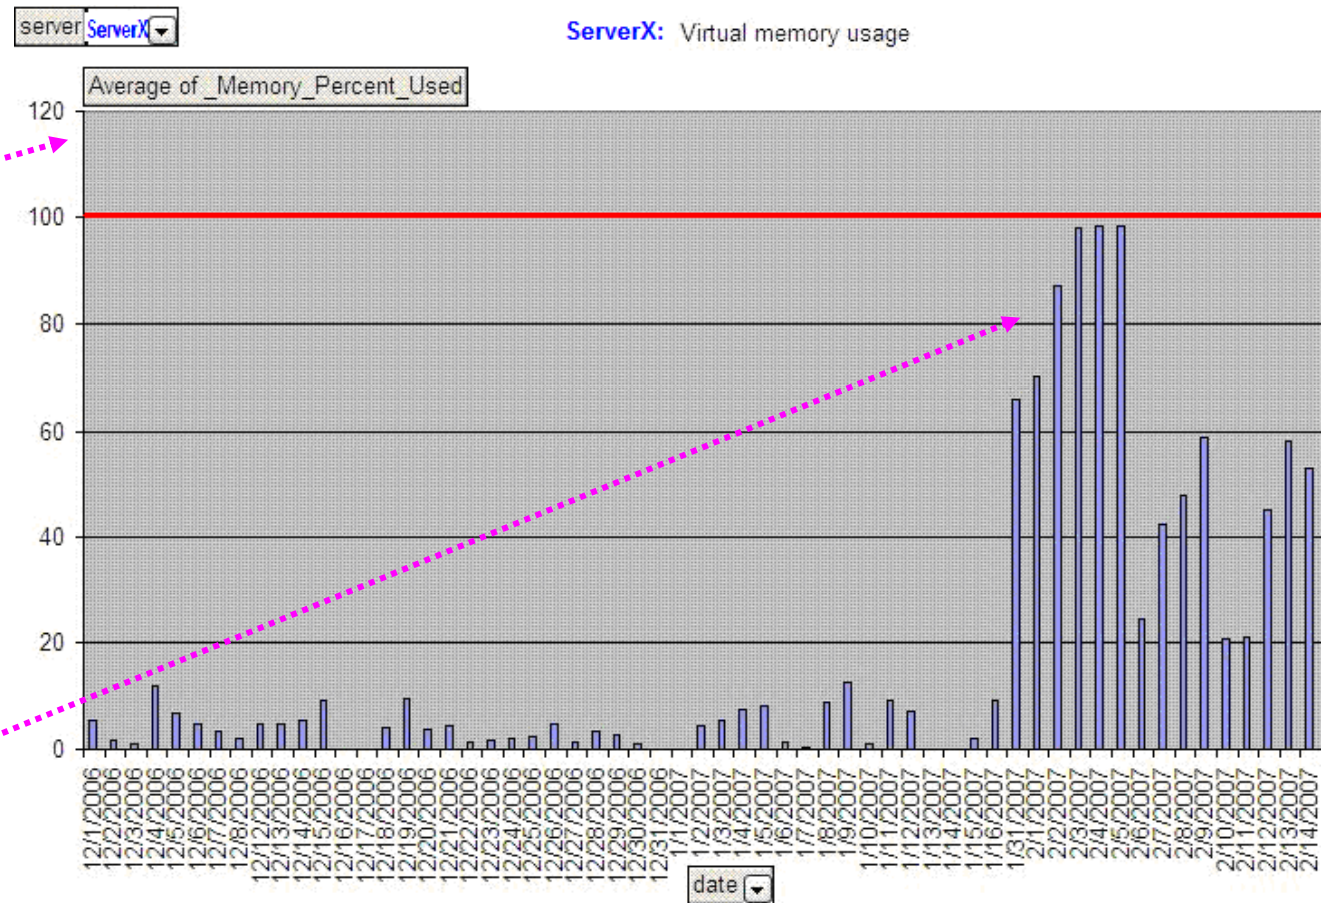


The SEDS and Memory Metrics (*Virtual Memory – end of case*)

- Finally, the Memory issue becomes obvious after downloading virtual memory utilization data to a spreadsheet for analysis

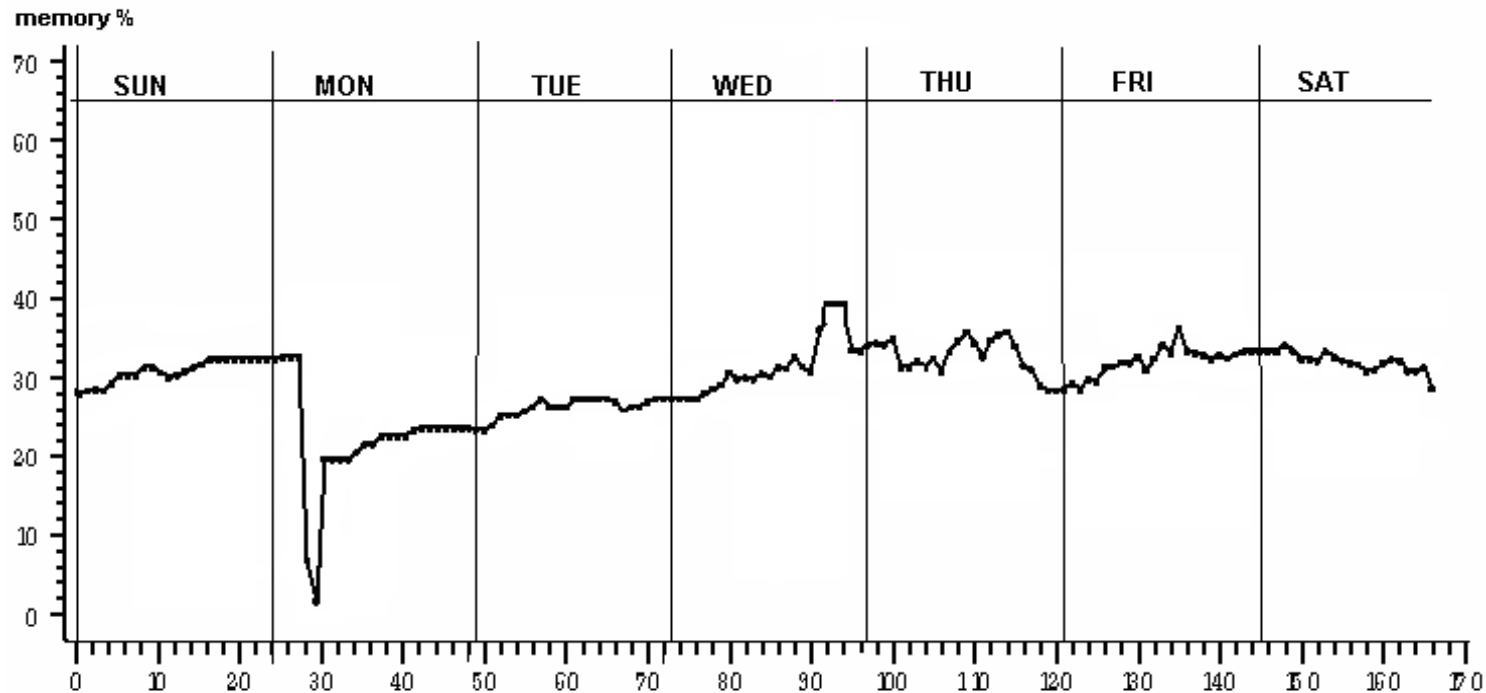
- CONCLUSION:**

Now we see the issue is serious. In February virtual memory usage jumped to almost 100% and continued to be much higher than before!




The SEDS and Memory Metrics (*Weekly control charts*)

- It is fairly common that the memory utilization metric does not fluctuate very much, unlike memory page-rate metrics. Instead of using 24-hour control charts, **7-day charts of hourly data** can reveal weekly patterns for memory utilization "at a glance":



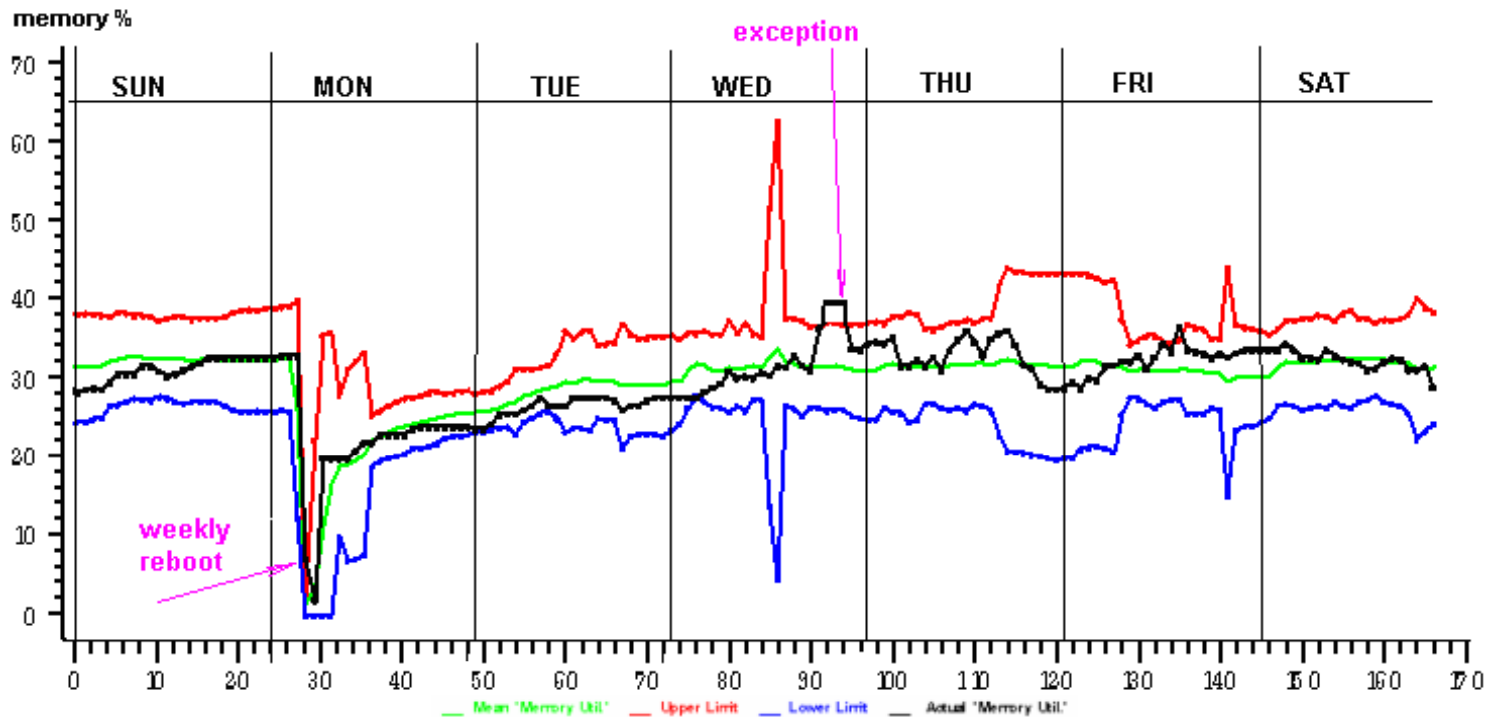
The SEDS and Memory Metrics (Weekly control charts)

- It is fairly common that the memory utilization metric does not fluctuate very much, unlike memory page-rate metrics. Instead of using 24-hour control charts, **7-day charts of hourly data** can reveal weekly patterns for memory utilization "at a glance":



The SEDS and Memory Metrics (Weekly control charts)

- It is fairly common that the memory utilization metric does not fluctuate very much, unlike memory page-rate metrics. Instead of using 24-hour control charts, **7-day charts of hourly data** can reveal weekly patterns for memory utilization "at a glance":



- This example shows the **weekly scheduled server reboot** (to avoid memory leak issues). This kind of graph is also useful since, even if there were no exceptions from yesterday, it may **show exceptions from previous days**.

The SEDS – Final Part

- All the main subjects related to SEDS are covered in previous CMG papers. **This paper is the last one** in a series that has been published in CMG since 2001.
- The successful implementation and usage of the system as a capacity management tool proves the concept that **SPC and MASF** can be applied **for Computer System Management**.
- The numerous responses from CMG audiences show **the positive influence made by SEDS** related publications on Capacity Management.
- The author invites all interested readers, MASF users and enthusiasts to visit and respond to the following **technical Blog**:

<http://itrubin.blogspot.com/>

References

1. **Kevin McLaughlin, Igor Trubin: “Exception Detection System, Based on the Statistical Process Control Concept”,** Proceedings of the Computer Measurement Group, **2002**.
(http://www.cmg.org/measureit/shared/trubin_02.pdf)
2. **Linwood Merritt, Igor Trubin: “Disk Subsystem Capacity Management Based on Business Drivers I/O Performance Metrics and MASF”,** Proceedings of the Computer Measurement Group, **2003**.
(http://regions.cmg.org/regions/ncacmg/downloads/june162004_session3.doc)
3. **Igor Trubin: “Mainframe Global and Workload Level Statistical Exception Detection System, Based on MASF”,** Proceedings of the Computer Measurement Group, **2004**.
(<http://www.cmg.org/membersonly/2004/papers/4179.pdf>)
4. **Igor Trubin: “Capturing Workload Pathology by Statistical Exception Detection System”,** Proceedings of the Computer Measurement Group, **2005**.
(<http://www.cmg.org/membersonly/2005/papers/5016.pdf>)
5. **Igor Trubin: “2006 - System Management by Exception, Part 6”,** Proceedings of the Computer Measurement Group, **2006**.
(<http://www.cmg.org/membersonly/2006/papers/6120.pdf>)

Questions?

Thank you!

Igor Trubin, Ray White

<http://itrubin.blogspot.com/>