The z/OS GRS Serialization "Jack of All Trades" Tools for (Performance/Contention/Monitoring...)

Session # **11717**
Nat Stevenson III
stevensn@us.ibm.com

GRS Serialization

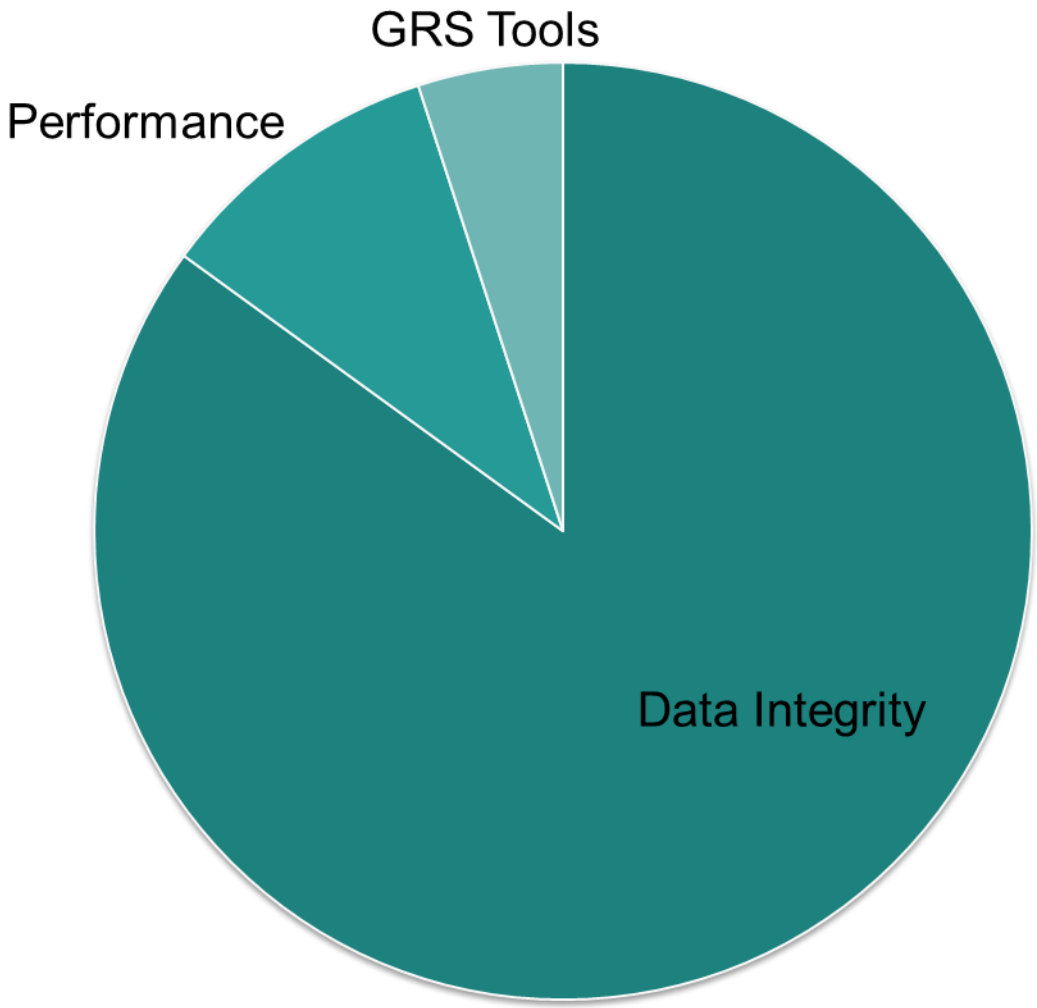

Some other improvements - Produccion

	MIM	GRS
Delete 8000 files	10 hours	30 min
PREOPEN DEDB in IMS start	30 min	10 min
DB2 Change Accumulation	5-10 hours	1-3 hours
DB2: Maintenance (delete files)	4-5 hours	10 min
IMS Change Accumulation: contention time in DBRC	20 min	Almost 0
Housekeeping RMM	45 min	1 min
Housekeeping CA-1	3 hours	30 min
SPAZIO (File Transfer for Internet): average time to catalog a file	7 sec	< 1 sec
DB2: DELAYs due to ENQUEUEs	8%	0,1 – 0,2%
Catalog Repro Mergecat (OEM)	1.400 /min	15.000/min

GRS Serialization

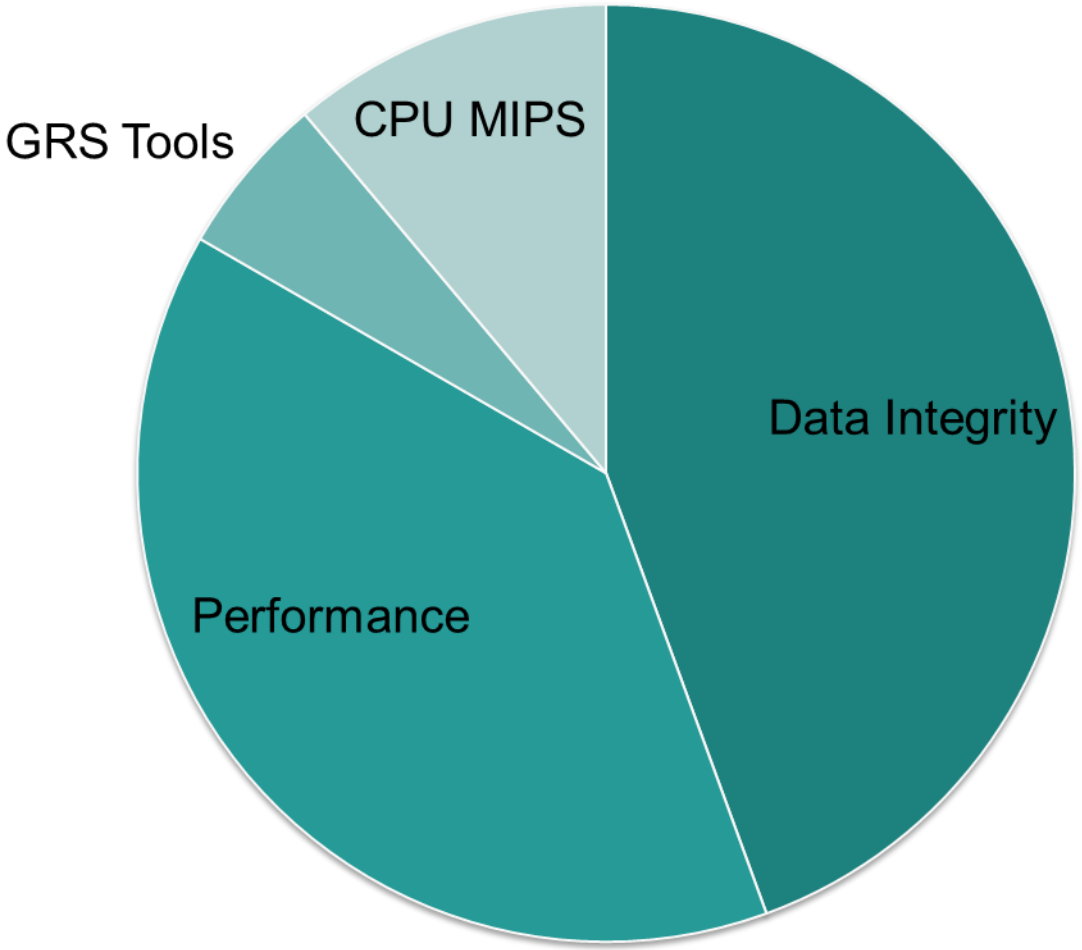

Serialization Explosion

- **Serialization (GRS) has now formed its own niche within z/OS core (Ring & STAR), supporting the following functions/areas:**
  1. z/OS Performance (for all “workloads”)
  2. GDPS (Disaster Recovery)
  3. DB2 (Middleware “expanding role”)
  4. CICS (Middleware “expanding role”)
  5. New Server Exploitation (Faster Processor Speed)
  6. Basic Sysplex -> Parallel Sysplex (Top PSLC Qualifier, Workload Performance, Higher Availability,…)
  7. ATS STAR (Shared Tape Serialization)
  8. Online Processing (TSO/E, CICS,…)
  9. Batch Processing (Allocation, DFSMS,…)
  10. Parallel Sysplex Mergers (Duplicate Data Sets, Serialization merging, Function merging,…)
  11. Debugging (Improved Contention msgs, GRS Monitor,…)

Performance

Reserve Processing

- **Reserves are old technology (over 30 years old), that is still in use today in over 80% customer sites.**
- **Ties up the volume to a single system (locks out other systems till last job is finished with the resource.**
- **Despite its age, it is still faster than current enqueue serialization (except for GRS STAR).**
- **“Pure performance” - Serialization functionality will slow down resource performance.**

Base Sysplex / Non-Sysplex

ISGNQRSP (Samplib)

- ISGNQRSP (GRS STAR and GRS Ring)

  - The samplib ISGNQRSP program issues a global enqueue, and then measures the enqueue response time from start to finish. The enqueue will be issued for a set # of times (default is “500”) and should be issued during a normal, heavy usage time of the day. I would recommend changing the “default count” to a more manageable number (e.g. “12”).

  - The results will then be placed in a data set. The time will be in microseconds.

```
03/14/2007 14:28:19.848148  005980
03/14/2007 14:28:20.073718  001843
03/14/2007 14:28:20.281341  003490
03/14/2007 14:28:20.486775  001948
03/14/2007 14:28:20.691864  001591
```

Base Sysplex -> Parallel Sysplex

- **Base Sysplex involves a slower enqueue performance environment that lacks future improvement.**
- **Parallel Sysplex qualifier (GRS STAR) – More “bang for the buck” if moving from “Base” to “Parallel”, due to improved workload completion times, availability,… among other things.**
- **Vital for any increase in workload (Online, Batch, Middleware,…).**

Parallel Sysplex

IBM Servers

Z9 BC

Z10 BC

z196

ISGLOCK “Dedicated vs Shared” Engines (CF)

- **Dedicated – Best for “enqueue performance”. Little, if any “ASYNC” requests.**
- **Shared – Will slow down “enqueue performance” resulting in increase in “ASYNC” requests.**

ISGLOCK – SYNC vs ASYNC

- STRUCTURE NAME = ISGLOCK TYPE = LOCK STATUS = ACTIVE
- # REQ ---------- REQUESTS ---------- ---------- DELAYED REQUESTS ----------
- SYSTEM TOTAL # % OF -SERV TIME(MIC)- REASON # % OF ---- AVG TIME(MIC) ----- EXTERNAL REQUEST
- NAME AVG/SEC REQ ALL AVG STD_DEV REQ REQ /DEL STD_DEV /ALL CONTENTIONS
-
- SYS1 150K SYNC 23K 7.5 324.5 432.1 NO SCH 474 0.2 412.3 577.7 1.2 REQ TOTAL 198K
- 164.9 ASYNC 145K 40.7 585.2 544.3 PR WT 0 0.0 0.0 0.0 0.0 REQ DEFERRED 1123
- CHNGD 353 0.1 INCLUDED IN ASYNC PR CMP 0 0.0 0.0 0.0 0.0 -CONT 1123
- -FALSE CONT 364
-
- SYS2 200K SYNC 8154 2.7 371.9 534.6 NO SCH 599 0.4 894.3 587.1 3.5 REQ TOTAL 177K
- 272.1 ASYNC 150K 45.5 958.8 834.5 PR WT 0 0.0 0.0 0.0 0.0 REQ DEFERRED 628
- CHNGD 554 0.2 INCLUDED IN ASYNC PR CMP 0 0.0 0.0 0.0 0.0 -CONT 628
- -FALSE CONT 123

ISGLOCK - Enqueue Performance Times

- STRUCTURE NAME = ISGLOCK TYPE = LOCK STATUS = ACTIVE
- # REQ ---------- REQUESTS ---------- ---------- DELAYED REQUESTS ----------
- SYSTEM TOTAL # % OF -SERV TIME(MIC)- REASON # % OF ---- AVG TIME(MIC) ----- EXTERNAL REQUEST
- NAME AVG/SEC REQ ALL AVG STD_DEV REQ REQ /DEL STD_DEV /ALL CONTENTIONS
-
-
- SYS1 4516 SYNC 4516 0.8 25.0 8.5 NO SCH 0 0.0 0.0 0.0 0.0 REQ TOTAL 4525
- 5.02 ASYNC 0 0.0 0.0 0.0 PR WT 0 0.0 0.0 0.0 0.0 REQ DEFERRED 127
- CHNGD 0 0.0 INCLUDED IN ASYNC PR CMP 0 0.0 0.0 0.0 0.0 -CONT 127
- -FALSE CONT 11
-
- SYS2 124K SYNC 124K 21.1 3.6 5.4 NO SCH 0 0.0 0.0 0.0 0.0 REQ TOTAL 129K
- 137.6 ASYNC 0 0.0 0.0 0.0 PR WT 0 0.0 0.0 0.0 0.0 REQ DEFERRED 309
- CHNGD 0 0.0 INCLUDED IN ASYNC PR CMP 0 0.0 0.0 0.0 0.0 -CONT 309
- -FALSE CONT 63

GDPS

- **Need to eliminate all “reserves” involved with HyperSwap processing. Performance degradation possible with slower enqueues being used to replace reserves (one exception is GRS STAR).**

DB2 / GRS (Workload Example)

- **GRS (ENQMAX) –**
  1. Old Value - 250,000
  2. New Value - **99,999,999**
- **DB2 (PK29281) – 100,000 concurrently open data sets per DB2 system.**
- **Ex. 3 DB2 Subsystems operating at 100,000 enqueues per subsystem (Tot – 300,000 concurrent enqueues).**
  1. Base Sysplex - 3000 usec (3 ms) vs
  2. Parallel Sysplex – **20 usec**

CICS Online Transactions

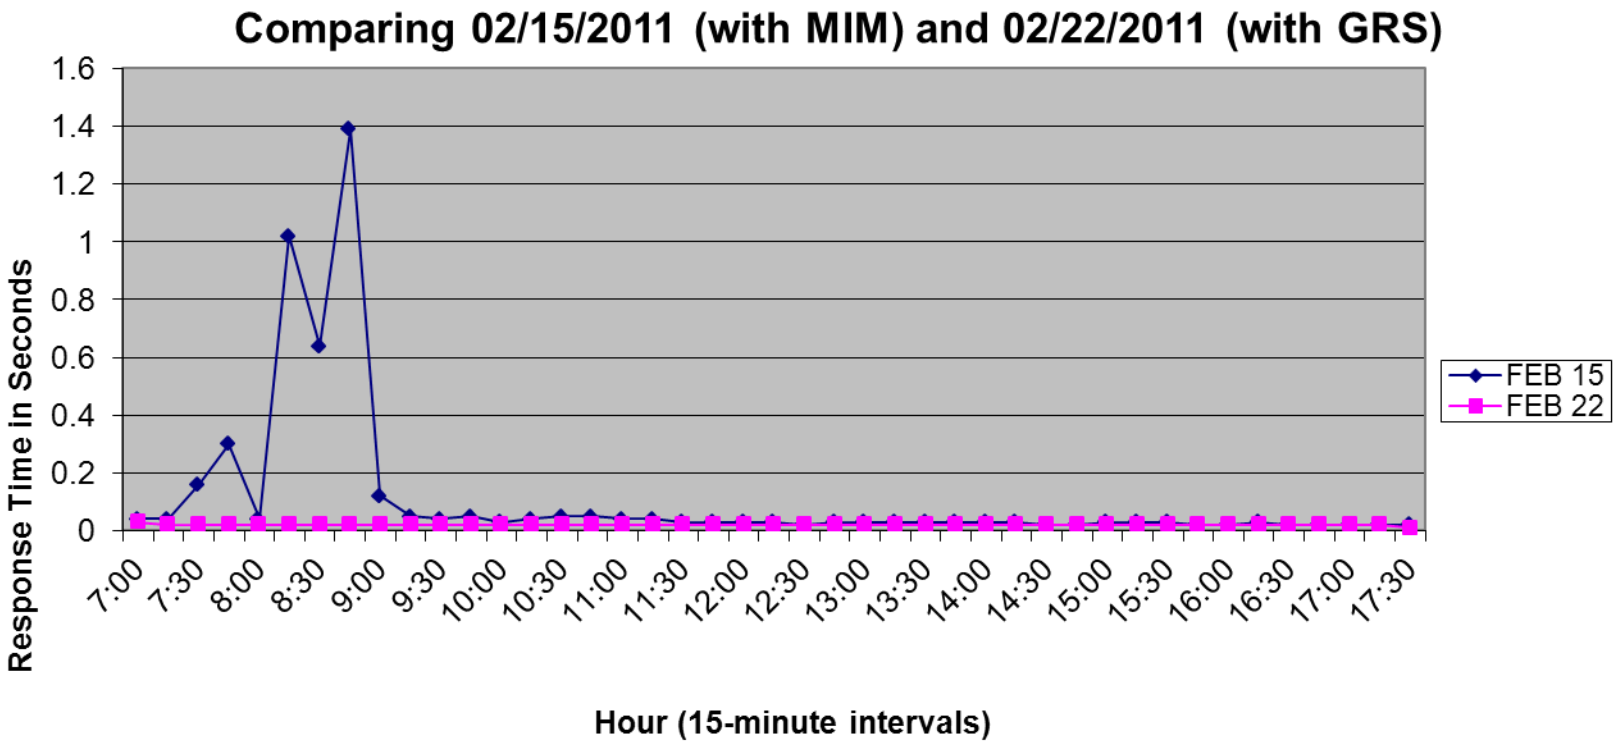

CPU Usage

System/390 Parallel Sysplex Performance (Redbook 1998)

- GRS System ITR Ratio
- **GRS Environment 2-way / 3-way / 4-way / 5-way**
- **Release 1 Ring** 1.000 / 0.979 / 0.932 / 0.873
- **Release 2 Ring** 0.977 / 0.963 / 0.907 / N/A
- **STAR** 1.078 / 1.055 / 1.017 / 0.995

MIM CPU Consumption (24h)

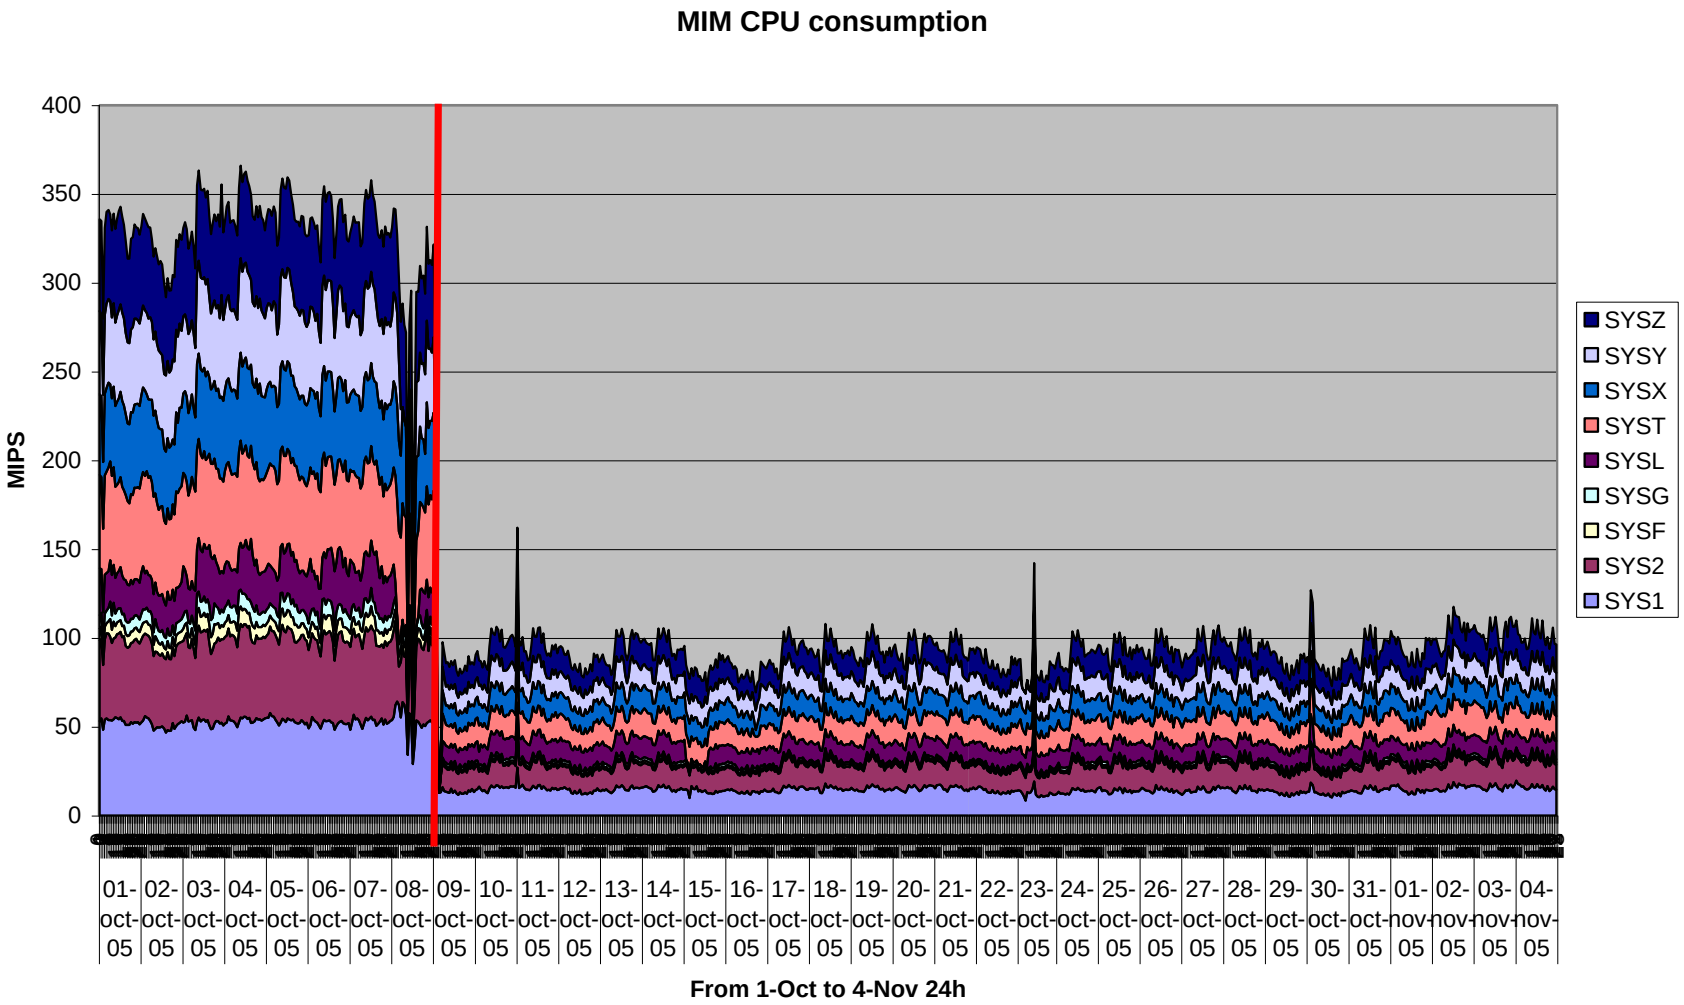

GRS CPU consumption (24h)

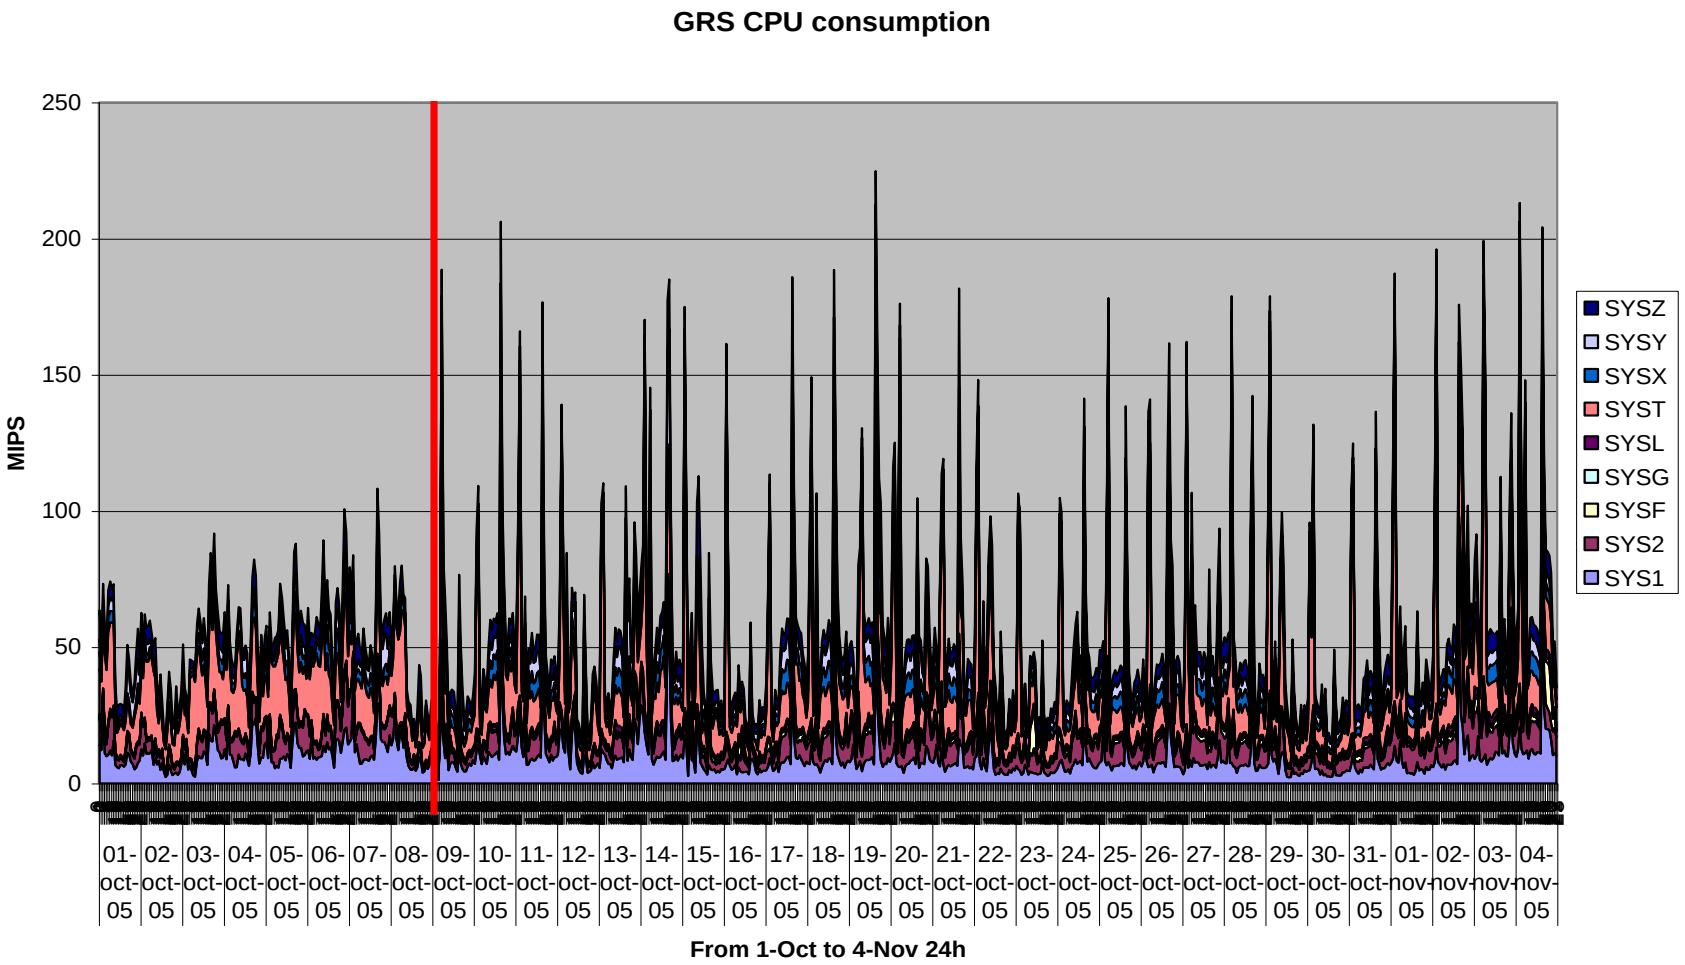

Almost the same with some spikes related to catalog activity

GRS Tools

GRS Monitor

Main Menu: Star Mode

```
                ENQ/DEQ Monitor - Main Menu

Select an option:

__   1.  MAJOR Names                    Date &   Time       : 99.049   09:36
     2.  Resource Name List             Monitor started at : 99.049   09:10
     3.  Volume List                    Elapsed seconds    :           1557
     4.  Filter List                    SMF System ID      :           SP21
-------------------------------------------------------------------------------
GRS Star -> Number of Lock Entries: 1048576                NUMSYS: 2
-------------------------------------------------------------------------------
                                   Time of Delay High. . : 99.049 09:32:36
Global Requests . . . :      364   Enqueue Delay Hi - Low: 18241262    1030
Local  Requests . . . :     1494   Enqueue Delay  mic-sec:     1030

Major Names . . . . . :       29
Minor Names . . . . . :      196
Volumes . . . . . . . :        3   Data Space Used .bytes:    28664    0 %
Number of Events. . . :     3716   Active Filter. . . . .:       08
Lost Events . . . . . :        0   Events Rate  . . . . .:       22
```

ENQ/DEQ Monitor - Major Name List Row 1 to 21 of 31

Enter S to select a Major Name for details .
L major on command line to locate a Major. Elapsed seconds: 5076

Sel. Field	Major Name	RNL	Scope	Counter	-average- msec	-Reserved- seconds
_	SPFEDIT		RESERVE	36	2135	77
_	SYSVTOC		RESERVE	8	202	1
_	SYSZABCD	EXCLUDED	*RESERVE	7	0	0
_	SYSZABCD		RESERVE	9	12	0
_	SYSZCOMM	NO	SYSTEMS	2		
_	SYSZDSCB		SYSTEMS	6		
_	SYSZENQM	NO	SYSTEMS	228		
_	SYSIGGV2		SYSTEMS	258		
_	SPFEDIT		SYSTEMS	42		
_	SYSDSN	INCLUDED	*SYSTEMS	136		
_	SYSIGGV2		SYSTEM	15		
_	SYSIKJBC		SYSTEM	4		
_	SYSIEFSD		SYSTEM	152		
_	SYSIKJPL		SYSTEM	4		
_	SYSIKJUA		SYSTEM	9		
_	SYSIEA01		SYSTEM	30		
_	SYSDSN		SYSTEM	4		
_	SYSZ#SSI	NO	SYSTEM	2		
_	SYSZWTOR		SYSTEM	13		

ISGNQRSP

ISGNQRSP (Samplib)

- ISGNQRSP (GRS STAR and GRS Ring)

The ISGNQRSP program issues a global enq, and then measures the enq response time from start to finish. The enq will be issued for a set # of times (**12**) and should be issued during a normal heavy usage time of the day. The results will then be placed in a data set. The time will be in microseconds.

```
03/14/2007 14:28:19.848148  005980
03/14/2007 14:28:20.073718  001843
03/14/2007 14:28:20.281341  003490
03/14/2007 14:28:20.486775  001948
03/14/2007 14:28:20.691864  001591
03/14/2007 14:28:20.897170  002042
03/14/2007 14:28:21.103801  001585
03/14/2007 14:28:21.317708  002135
```

ISGECMON / ISGECJES

ISGECMON / ISGECJES

- ISGECMON (Samplib) / ISGECJES (Advanced Technical Support)
  - **This program runs as a never ending task that checks dataset contention at periodic intervals, and sends messages to TSO users and/or system operators for batch jobs causing contention.**

1 - Issue notification messages either to the operator or the TSO user for the following conditions for a resource:

OWN	WAIT	
TSO/E	BATCH	
BATCH	BATCH	(ISGECJES Only)

The following contention situation can be handled by ISPF. Enter the PF1 help key (twice) for enqueue contention involving only TSO/E users:

OWN	WAIT
TSO/E	TSO/E

GRS – “Display Cmds” (Contention)

GRS - “Display Cmds” (Contention)

- GRS “Display Contention” commands.
  - Old Contention Msg –
    - D GRS,C
  - New Contention Msgs –
    - D GRS,AN,**WAIT**
    - D GRS,AN,**BLOCK**”
    - D GRS,AN,**DEP**”

GRS - “Display Cmds” (Contention)

D GRS,C

```
ISG343I 16.38.21 GRS STATUS 298
S=SYSTEMS SYSDSN   STEVENS.TEST.DATA1
SYSNAME      JOBNAME      ASID    TCBADDR   EXC/SHR    STATUS
SYSA     STEVENSA        0031     008FF350 EXCLUSIVE   OWN
SYSB     NAT45B          0032     008FF350 EXCLUSIVE   WAIT
S=SYSTEMS SYSDSN   STEVENS.TEST.DATA2
SYSNAME      JOBNAME      ASID    TCBADDR   EXC/SHR    STATUS
SYSA     STEVENSA        0031     008FF350 EXCLUSIVE   OWN
SYSB     NAT45B          0032     008FF350 EXCLUSIVE   WAIT
```

D GRS,AN,WAIT

```
ISG349I 17.04.38 GRS ANALYSIS 340
LONG WAITER ANALYSIS:  ENTIRE SYSPLEX
WAITTIME  SYSTEM  JOBNAME E/S SCOPE QNAME    RNAME
00:27:12 SYSB     NAT45B*E*  SYSS SYSDSN   STEVENS.TEST.DATA1
BLOCKER  SYSA     STEVENSA E
00:27:12 SYSB     NAT45B*E*  SYSS SYSDSN   STEVENS.TEST.DATA2
BLOCKER  SYSA     STEVENSA E
```

GRS - “Display Cmds” (Contention)

D GRS,AN,BLOCK

```
ISG349I 17.11.02 GRS ANALYSIS 377
LONG BLOCKER ANALYSIS:  ENTIRE SYSPLEX
BLOCKTIME SYSTEM  JOBNAME E/S SCOPE QNAME    RNAME
0:00:22 SYSA     STEVENSA*E*  SYSS SYSDSN0   STEVENS.TEST.DATA1
                 OTHER BLOCKERS: 0  WAITERS: 1
00:00:22 SYSA    STEVENSA*E*  SYSS SYSDSN    STEVENS.TEST.DATA2
                 OTHER BLOCKERS: 0  WAITERS: 1
```

D GRS,AN,DEP

```
ISG349I 16.53.27 GRS ANALYSIS 327
DEPENDENCY ANALYSIS:  ENTIRE SYSPLEX
----- LONG WAITER #1
WAITTIME  SYSTEM  JOBNAME E/S SCOPE QNAME    RNAME
00:16:00 SYSB     NAT45B*E*  SYSS SYSDSN    STEVENS.TEST.DATA1
BLOCKER  SYSA     STEVENSA E
----- LONG WAITER #2
WAITTIME  SYSTEM  JOBNAME E/S SCOPE QNAME    RNAME
00:16:00 SYSB     NAT45B*E*  SYSS SYSDSN    STEVENS.TEST.DATA2
BLOCKER  SYSA     STEVENSA E
```

GRS – “Display Cmds” (CNS)

GRS - “Display Cmds” (CNS)

- GRS “CNS” (Contention Notification System)
  - `D GRS,ALL`

```
NO ENQ RESOURCE CONTENTION EXISTS
NO REQUESTS PENDING FOR ISGLOCK STRUCTURE
NO LATCH CONTENTION EXISTS
SYSTEM    STATE           SYSTEM    STATE
SYSA      CONNECTED       SYSD      CONNECTED
SYSB      CONNECTED       SYSC      CONNECTED
THERE IS NO RNL CHANGE IN PROGRESS.
GRS STAR MODE INFORMATION
  LOCK STRUCTURE (ISGLOCK) CONTAINS 4194304 LOCKS.
  THE CONTENTION NOTIFYING SYSTEM IS SYSD
```

SETGRS CNS=XXXX (to change the CNS system)

GRS Serialization

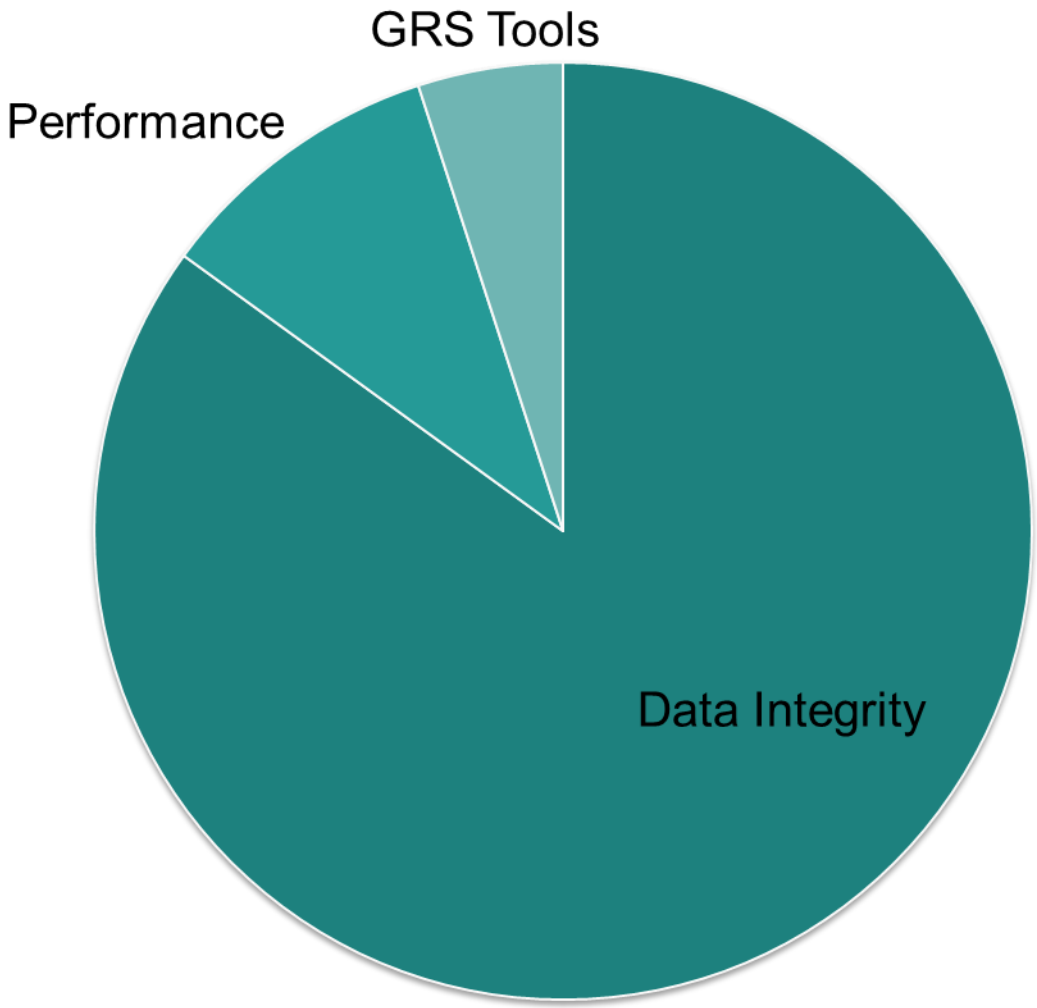

GRS Serialization

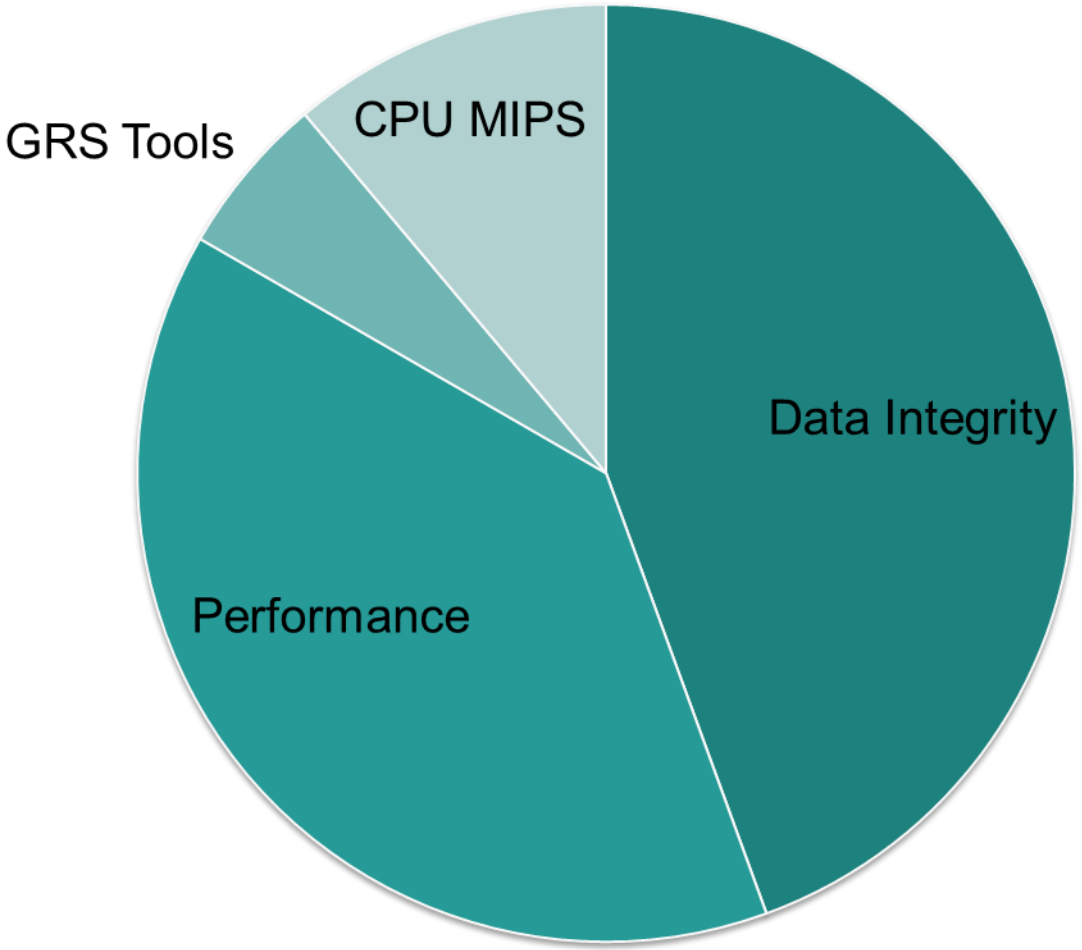

Questions