35 YEARS OF PERFORMANCE MANAGEMENT

Understanding z/VM Storage Analysis and Tuning

- Barton@VelocitySoftware.com
- HTTP://VelocitySoftware.com

"If you can't Measure it,

I am Just Not Interested ™"

VELOCITY SOFTWARE

PROVEN PERFORMANCE

Storage Subsystem Objectives

Objectives:

- Understand z/VM Storage Requirements
- Understand Linux Storage Requirements
- Know how/where to measure
- Understand Demand Paging
- Determine Requirements
- Understand Measurements
- Suggest tuning measures

Mapping Storage (three types)

CP Fixed (at IPL) Storage (no tuning options)

- Nucleus
- Frame table (32 byte frame table entry per real page about 0.78%)

CP Non-Pageable: Storage made non–pageable by CP

- Free storage (control blocks, EREP)
- Prefix pages (2 per processor / IFL)
- SNTBKs (one per DCSS/NSS)
- SAVBKs (CP save areas, 32 per page)
- VMDBKs (2 pages per guest virtual processor)
- QDIO Buffers
- DIAG98 Buffers

The remaining pages are Dynamic Paging Area (DPA)

Dynamic Paging Area

Dynamic Paging Area (DPA)

- Virtual Machines “resident”
- Shares storage (DCSS/NSS)
- System address spaces
- Virtual Disks
- MDC
- Available (< 2g, < 2g)

Storage Capture Ratio

Capture ratio is accumulation of knowns / total storage

Ensure Data Accuracy of where storage is used

Drill downs for all columns:

- System execution space: ESASXS
- User Resident (by user): ESAUSPG, ESAUSR2, ESAUSP2
- NSS/DCSS Resident (shared storage by segment): ESADCSS
- System/User address Space (by address space): ESAASPC
- VDISK Resident (by user): ESAUSPG
- MDC: ESAMDC

Storage Map - ESASTR1

ESASTR1 shows Storage Map to show storage (14GB example)

- User resident should be major use

Capture ratio shows accuracy

Overcommit means paging will happen

Changes to storage are very visible

```
Report: ESASTR1      Main Storage Analysis                          Velocity Software Corporate   ZMAP 5.1.2 04/16/21   Pg   2
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78    First record analyzed: 04/15/21 00:00:00
-------------------------------------------------------------------------------------------------------------------------------
         Users <------------------------------Pages-------------------------------------------------------------> Over
         Loggd System  Fixed Non-  Free Frame <Available> Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag Commit     Capt-
Time        On Storage Store Pgble Stor Table <2gb  >2gb  ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98  Ratio      Ratio
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- ----- ---- -----      -----
17:15:00   111 3670016  2878 20882 1152 28672  3099  2421 52296  3384K    35078  75713     0  4441 18566  16K 3.653      0.988
17:30:00   111 3670016  2878 20883 1166 28672  3164  2669 52296  3383K    35077  75714     0  4307 19741  16K 3.653      0.988
17:45:00   111 3670016  2878 20872 1147 28672  3195  2389 52298  3381K    35074  75716     0  4270 21989  16K 3.653      0.988
18:00:00   111 3670016  2878 20889 1146 28672  3128  2851 52306  3383K    35079  75722     0  4103 19648  16K 3.653      0.988
18:15:00   113 3670016  2878 20876 1141 28672  3077  2508 52316  3384K    35099  75776     0  4028 19283  16K 4.609      0.988
18:30:00   116 3670016  2878 20880 1075 28672  3137  2544 52360  3349K    32071   122K     0  2118 12337  16K 7.354      0.988
18:45:00   116 3670016  2878 20808 1038 28672  3051  2234 52407  3293K    29914   196K     0     0    47  16K 8.227      0.988
19:00:00   116 3670016  2878 20765 1028 28672  3056  2245 52414  3293K    29082   196K     0     0   127  16K 8.227      0.988
19:15:00   115 3670016  2878 20797 1040 28672  3063  2232 52409  3297K    29522   192K     0    22    73  16K 8.754      0.988
```

Storage Map Analysis

Storage Map – What changed at 18:30?

- 3.67M pages – 1.5TB
- System Address space use increased – page tables
- VDISK pages stolen
- MDC Resident – went down, freed
- Over Commit Ratio – went up, more virtual machines

Result of large virtual machines logging on

- Available list is small
- Page tables in system address space grew
- User resident went down, no room

(Note option to show pages or MB)

```
Report: ESASTR1      Main Storage Analysis                    Velocity Software Corporate   ZMAP 5.1.2 04/16/21   Pg   2
-------------------------------------------------------------------------------------------------------------------------

         Users <----------------------------Pages-----------------------------------------------------------> Over
         Loggd System  Fixed Non-  Free Frame <Available> Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag Commit    Capt-
Time        On Storage Store Pgble Stor Table <2gb  >2gb  ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98  Ratio     Ratio
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- ----- ---- -----     -----

17:00:00   111 3670016  2878 20879 1153 28672  3170  2501 52291  3387K    35061  75702     0  4729 15418  16K 3.653     0.988
17:15:00   111 3670016  2878 20882 1152 28672  3099  2421 52296  3384K    35078  75713     0  4441 18566  16K 3.653     0.988
17:30:00   111 3670016  2878 20883 1166 28672  3164  2669 52296  3383K    35077  75714     0  4307 19741  16K 3.653     0.988
17:45:00   111 3670016  2878 20872 1147 28672  3195  2389 52298  3381K    35074  75716     0  4270 21989  16K 3.653     0.988
18:00:00   111 3670016  2878 20889 1146 28672  3128  2851 52306  3383K    35079  75722     0  4103 19648  16K 3.653     0.988
18:15:00   113 3670016  2878 20876 1141 28672  3077  2508 52316  3384K    35099  75776     0  4028 19283  16K 4.609     0.988
18:30:00   116 3670016  2878 20880 1075 28672  3137  2544 52360  3349K    32071   122K     0  2118 12337  16K 7.354     0.988
18:45:00   116 3670016  2878 20808 1038 28672  3051  2234 52407  3293K    29914   196K     0     0    47  16K 8.227     0.988
19:00:00   116 3670016  2878 20765 1028 28672  3056  2245 52414  3293K    29082   196K     0     0   127  16K 8.227     0.988
```

Storage Overcommit

What is the problem to solve? (cost control vs performance)

- Project storage requirements to manage paging delay
- Storage overcommit best “storage utilization” metric

Define storage overcommit

- Size of virtual machines logged on / Size of real storage
- (SUM VMDSIZE / MTRMEM.RSAGSTOR)

Ranges of overcommit?

- .9 for no paging, online/real time production
- 2-3 for development
- Arbitrarily create extra large servers, overcommit goes up
- Tune servers, overcommit goes down
- Its only purpose is to gauge paging requirements

Shared Storage (DCSS) Requirements

"Pages resident" show how many pages in storage

18k pages in shared storage for monitoring (zmon):

- Linux servers, z/VM
- 4 z/OS, 5 VSE, 10 CICS regions/partitions
- z/VM

DCSS storage can be protected

```
Report: ESADCSS       NSS/DCSS Analysis                                  Vel
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78          Fir
-------------------------------------------------------------------------------

                 <-Users-> <---------Number of Pages for DCSS---------->
          Spool       Non-        Non- <--resident--> <Locked> <PagedOut>
Name         ID Shrd Shrd Saved Data <2GB >2GB HOST <2G >2GB DASD XSTOR
-------- ----- ---- ---- ----- ---- ---- ---- ---- --- ---- ---- -----
18:30:00
CMS          40   74    0  1298    0   77  389    0   0    0 1298     0
CMSFILES      9    8    0   768    0   11  126    0   0    0  639     0
CMSPIPES     36   79    0   512    0   18   84    0   0    0  512     0
CMSVMLIB      6   79    0   256    0    6   46    0   0    0  256     0
EAGRTSEG     13    4    0    80    0    0   16    0   0    0   80     0
GCS           1    2    0   120  917    2   18    0   0    0 1037     0
INSTSEG       5   79    0   768    0   29   81    0   0    0  768     0
MONDCSS      11    3    0     0  16K    9  100  109   0  102    0     0
SCEE         23   11    0   256    0    0   15    0   0    0  256     0
SCEEX        39   11    0  2304    0   12  169    0   0    0 2304     0
TCPIP        24    4    0     0  256    2    3    0   0    0  255     0
ZMON         12   25    0     0  16K  15K 1797 7409   0    0    0     0
ZOSMON       16    2    0     0 4096   69 2037 2044   0    0 4096     0
ZVWS         29   25    0     0  768    1   74    0   0    0  164     0
             ---- ---- ----- ---- ---- ---- ---- --- ---- ---- -----
 Totals:      427    0  6362  52K  16K  16K 9562   0  102  12K     0
```

Dynamic Paging Area (DPA) Overview

Storage Requirements

- System functions require storage
- Work requires programs and data
- More data in storage improves response time

Overcommit (sharing)

- “Expensive” storage is shared in virtual environment
- Storage often used once (initialization), not needed after
- Unreferenced pages of virtual machine can be paged out
- Idle applications and data can be paged out
- Overcommit is the key metric for capacity planning

Linux is challenging

Maximizing Pageable Storage

Limit virtual machine sizes

- PGMBKs – cost 8mb (PTRM address space) per virtual GB
- PTRM address space is pageable, but starts resident

Limit the amount of main storage used by MDC:

- SET MDCACHE STORAGE **minM** maxM

CP Page Tables Requirements

Report: ESAASPC Shared Address Space

Owner	Space Name	<--Size---> <-(pages)-> Avg	Max	<-----Address Space Pages-----> <Resident> <2GB	>2GB	<Locked-> <2GB	>2GB	<PagedOut> DASD	XSTOR
18:30:00									
SYSTEM	ISFCDATASPACE	245K	524K	1	12	0	0	60	0
SYSTEM	PTRM0000	1049K	1049K	224	37861	0	0	9448	0
SYSTEM	PTRM0001	1049K	1049K	1434	50706	0	0	7365	0
SYSTEM	PTRM0002	1049K	1049K	46	14021	0	0	3529	0
SYSTEM	PTRM0003	1049K	1049K	123	17804	0	0	6004	0
SYSTEM	SYSTEM	315K	524K	15	67	0	0	3	0
DSA0001	VDISK$DSA0001$$0202$04F7	2730	8192	1	22	0	0	0	0
DSA0001	VDISK$DSA0001$$0203$04F8	10748	32256	87	0	0	0	0	0
MONGO01	VDISK$MONGO01$$0202$0075	15050	32256	75	337	0	0	13K	0
ORACLE	VDISK$ORACLE$$$0203$0227	15050	32256	3	42	0	0	13K	0
RHOSBOOT	VDISK$RHOSBOOT$0202$04F1	3822	8192	2	29	0	0	0	0
RHOSBOOT	VDISK$RHOSBOOT$0203$04F2	15050	32256	18	104	0	0	0	0
RHOSCP1	VDISK$RHOSCP1$$0202$04F9	1092	8192	0	9	0	0	0	0
RHOSCP1	VDISK$RHOSCP1$$0203$04FA	4302	32256	33	2	0	0	0	0
RHOSCP2	VDISK$RHOSCP2$$0202$04FB	1092	8192	0	9	0	0	0	0
RHOSCP2	VDISK$RHOSCP2$$0203$04FC	4302	32256	6	29	0	0	0	0
RHOSCP3	VDISK$RHOSCP3$$0202$04FD	546	8192	0	5	0	0	0	0
RHOSCP3	VDISK$RHOSCP3$$0203$04FE	2151	32256	0	17	0	0	0	0
RHOSWK1	VDISK$RHOSWK1$$0202$04F3	1639	8192	0	13	0	0	0	0
RHOSWK1	VDISK$RHOSWK1$$0203$04F4	6452	32256	37	15	0	0	0	0
SLES12	VDISK$SLES12$$$0203$000E	59841	128K	71	1059	0	0	60K	0
S11S2ORA	VDISK$S11S2ORA$0202$0008	29980	64256	21	68	0	0	30K	0
System Totals:		140M	140M	2200	122K	0	0	408K	0
Virtual Disk_Subset		1363K	1369K	356	1762	0	0	381K	0

Virtual disks are system-owned address spaces

- Rarely consume much storage, paged out when inactive
- PTRM are not used unless needed

CP Page Tables Requirements

The CP address spaces reported on Shared Address Space Analysis

- Resident page tables relative to active pages.
- 128 PTRM (4gb) address spaces defined in z/VM 6.3
- Only PTRM address spaces in use are displayed, there are more

```
Report: ESAASPC      Shared Address Space Analysis velocity Software Corporate
Monitor initialized: 04/15/21 at 00:00:00    first record analyzed: 04/15/21
-------------------------------------------------------------------------------
                                 <--Size---> <-----Address Space Pages----->
                                 <-(pages)-> <Resident> <Locked-> <PagedOut>
Owner    Space Name                Avg   Max <2GB >2GB  <2GB >2GB DASD XSTOR
-------- ------------------------ ----- ----- ---- ----- ---- ---- ---- -----
18:30:00
SYSTEM   ISFCDATASPACE             245K  524K    1    12    0    0   60     0
SYSTEM   PTRM0000                1049K 1049K  224 37861    0    0 9448     0
SYSTEM   PTRM0001                1049K 1049K 1434 50706    0    0 7365     0
SYSTEM   PTRM0002                1049K 1049K   46 14021    0    0 3529     0
SYSTEM   PTRM0003                1049K 1049K  123 17804    0    0 6004     0
SYSTEM   SYSTEM                   315K  524K   15    67    0    0    3     0
```

Sizing Large z/VM: Pageable PGMBKs

PGMBK is page table for virtual storage

PGMBK storage per referenced 1MB segment:

- Two 4k page PGMBK per 1mb segment (8mb/gb)
- 2048 pages/gb (100mb virtual requires 800mb real)
- (1gb Linux server: 8mb PGMBKs)

Locates all user pages in

- ~~Expanded Storage (pre z/VM 6.3)~~
- DASD Paging (and IBR list)
- Main Storage

A pageable PGMBK is eligible for page-out when it maps no virtual pages into real storage

Minidisk Cache

Minidisk cache defaults to “all of it”

- MUST BE CONTROLLED!!! Very common configuration error

Example is “very constrained”, (Available list small)?

- MDC consumes almost 1m pages (4gb) out of total 2m pages….
- CP algorithms decide the size, they are broken. Default is ALL storage
- CMS makes good use of MDC (32m to 128m is good)
- zVPS will slow down if no MDC
- Note: zVPS uses MDC to enhance zVIEW, zMON performance
- Set to 128m - SET MDC STORAGE min max

```
Report: ESASTR1      Main Storage Analysis                          zVM1 Prod.
-------------------------------------------------------------------------------------------------
         Users <-----------------------------Pages-------------------------------------------------->
         Loggd System  Fixed Non-  Free Frame <Available> Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag
Time        On Storage Store Pgble Stor Table <2gb  >2gb  ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- ----- ----
01:05:12    34 2064368  2970 28013 3802 16128  3186   114  7909  1075K      783   2717 19156    24  905K 1913
01:20:12    34 2064368  2970 27954 3802 16128  2633   105  7926  1045K      569   2472 19054     5  938K 1913
01:35:12    34 2064368  2970 27721 3802 16128  3860   139  7916  1120K     1357   6683 19255   647  857K 1913
```

Storage Overcommit

What is the problem to solve? (cost control vs performance)

- Project storage requirements to manage paging delay
- Storage overcommit best “storage utilization” metric

Define storage overcommit

- Size of virtual machines logged on / Size of real storage
- (SUM VMDSIZE / MTRMEM.RSAGSTOR)

Ranges of overcommit?

- .9 for no paging, online/real time production
- 2-3 for development
- Arbitrarily create extra large servers, overcommit goes up
- Tune servers, overcommit goes down
- Its only purpose is to gauge paging requirements

Sharing Storage

To “share” requires paging out:

- Inactive storage
- Inactive applications
- Initialization pages
- Inactive servers

Linux Storage Challenging

The problem: What pages to page out?

Inactive storage? Linux Storage is not idle

- Extra storage used to cache data and programs
- Linux servers are not idle
- Linux applications poll at 200 times per second
- Which servers are actually doing work if all are “active”
- What pages can be legitimately paged out of real storage?

Determining pages for page out:

- **Active server?** Can not know if server is working or **polling**
- Take least recently used, non modified, non referenced
- Fast page-in (page recovery) very important

Storage Subsystem Overcommit

Storage Requirements

- System functions require storage
- Work requires programs and data
- More data in storage improves response time

Overcommit (sharing)

- “Expensive” storage is shared in virtual environment
- Storage often used once (initialization), not needed after
- Unreferenced pages of virtual machine can be paged out
- Idle applications and data can be paged out
- Overcommit (sharing) is a metric for capacity planning,
- (Overcommit may may mean lack of tuning, extra large virtual machines???)

Sharing Storage

To “share” requires paging out:

- Inactive storage
- Inactive applications
- Initialization pages
- Inactive servers

Linux Storage management is worst case to virtualize

- “Round robin” keeps all storage active
- Oldest unreferenced page
  - Most likely to be paged out
  - Most likely to be next used by Linux
- All storage is used to buffer data, programs
- Small “available list”

z/VM Storage and Paging Architecture - Moving Targets

z/VM shared storage / Overcommit

- Objective: Page unused pages out to allow re-use
- Need optimal test before paging to slow disk
- Optimize page-in when needed (block paging)
- ~~Page migration from fast to slow as age out (gone)~~

Architectures to choose from:

- Excessive Storage – enough so no paging (expensive)
- ~~Expanded Storage – Very fast page-in, gone…~~
- Solid State paging device – sort of fast
- Disk paging devices – not fast
- Block paging to group pages for faster load

Storage / Paging Architecture Evolution

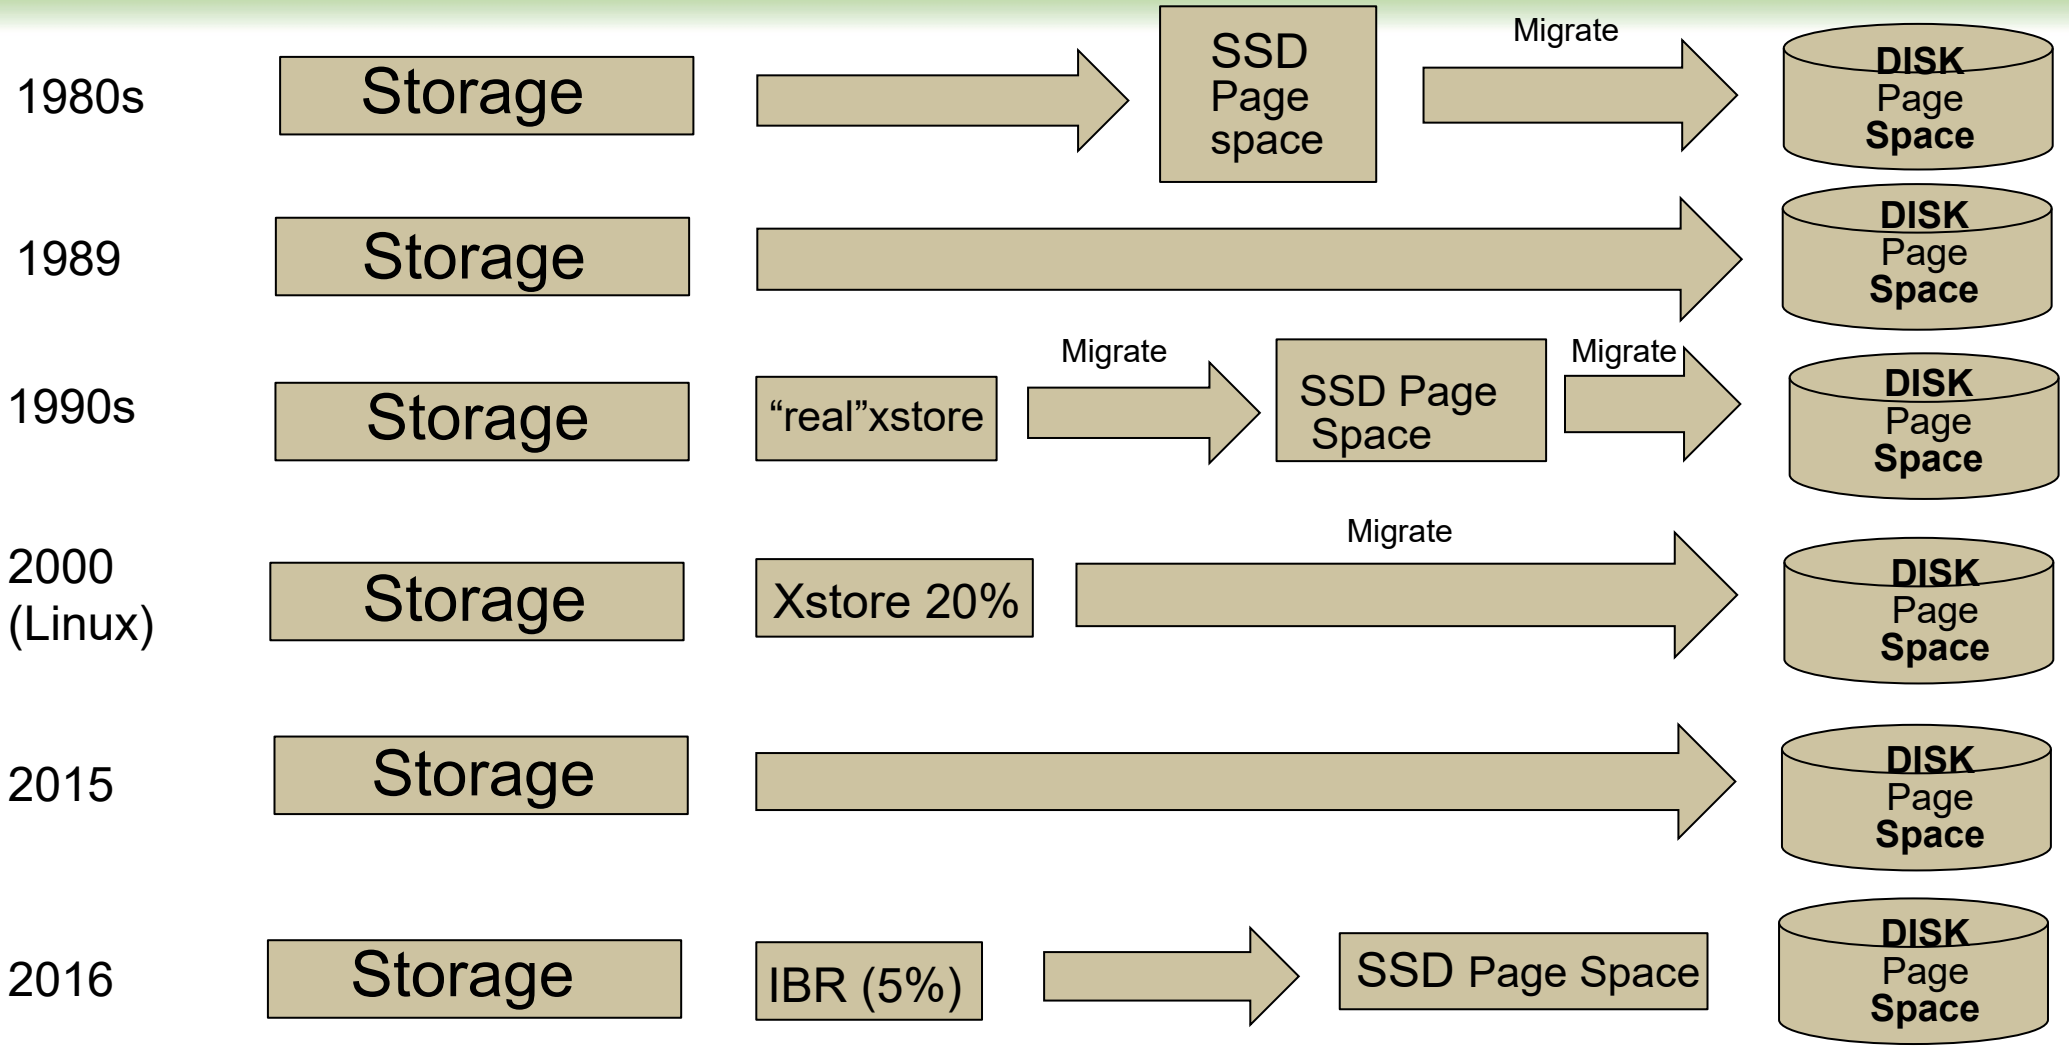

z/VM User Storage "test" Replaces XStore

Virtual machines have "resident pages"

- Active pages
- IBR (Marked Invalid But Resident)

IBR Pages (the heart of current paging architecture)

- Invalidated, so not addressable by address space
- Pages "stolen" from end of sorted IBR page list
- IBR pages optionally "pre-written"

Re-Referenced pages "validated"

- Page fault causes page to move to "Active Pages"
- Reclaim of "test invalidated page" very fast

Storage Management

Storage management functions

- IBR – Invalid But Resident
- Global Aging list – sorts pages LRU
- Saved segments (shared storage) can be protected

IBR List – pages belonging to user in "test" mode

- Pages marked IBR (Invalid But Resident) as a "test" - Pages are reclaimed if referenced
- Referenced pages go to top of frame list
- **Size or IBR list is configurable (to 5%)**

Global Aging List – list ready to page out

- Pages moved from IBR to sorted Global Aging list
- Pages can be reclaimed if referenced
- Pages will be oldest pages to be paged out

z/VM Storage Management Options

System Age List Important if Overcommit

- Maximum 5%
- Recommend 5% always
- SET AGELIST SIZE 5% EARLYWRITES YES KEEPSLOT YES

```
-Set--AGELIST---.-SIZE--.-n.n--PERCent-.-.
                |       |-n.n%---------| |
                |       '-storsize-----' |
                |-EARLYWrites--.-Yes-.---|
                |              '-No--'   |
                '-KEEPSlot--.-Yes-.------'
                            '-No--'
```

- CP QUERY AGELIST (defaults)

```
Target size    =          280576K (274M)      2.0% of pageable storage
In use         =          271712K
Pending writes =          120296K
Early writes   = Yes
Sizing         = Variable
Keep slot      = Yes
```

Storage Management - IBR

Each page of storage has a key (4 bits)

- 3 status bits - usually expressed as two hex digits with the last bit always 0
- For storage alteration, the key on the page must match the key in the PSW (bits 8-11)

Flags

- F – fetches are protected as well
- R – some location has been referenced
- C – some location has been changed

IBR: Invalid, but resident, fetch protected

To make valid, switch the “F” bit

z/VM Storage "Case Study"

Project to install OpenStack:

- VSIVM4 is demonstration LPAR
- Alerts set for high paging rate
- Alerts set for page space thresholds

Started with an emailed alert

- `Alert on 2021/04/15 at 20:02:11`
- `System paging rate is 113855 (above 10000 for 30 minutes)`

Process:

- ESAXACT to see impact
- Storage analysis to understand why

User Storage – Case Study

User Wait analysis – ESAXACT

- System page wait spikes from one interval to next
- Openstack server consuming storage and CPU

```
Report: ESAXACT       Transaction Delay Analysis                    Velocity Sof
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78     First record
-------------------------------------------------------------------------------
                      <-----Percent non-dormant (Wait states)------------------
UserID    <-Samples->                          E-   D-   T-      Tst <Asynch>
/Class    Total  In Q Run Sim CPU SIO Pag SVM SVM SVM  CF Idl I/O Pag Ldg Oth
-------- ----- ----- --- --- --- --- --- --- --- --- --- --- --- --- --- ---
18:15:00  1697   610 2.3 0.5 3.8   0 1.1   0   0 5.1   0  79 3.0   .   .   1
Hi-Freq:  110K 34022 2.9 0.1 1.0   0 0.0   0  10 3.7   0  87 0.0   0 0.2 0.0
-------------------------------------------------------------------------------
18:30:00  1738   658 2.0 0.8 2.0   0  14   0   0 3.8   0  69 1.8   .   .   2
Hi-Freq:  119K 40387 3.5 0.8 3.4 0.0 9.4   0 9.3 2.2 0.0  58 0.2 0.9  16 0.0
 ***Key User Analysis ***
TCPIP      900   376   0   0 0.8   0 2.1   0   0   0   0  96   0   0 1.3   0
 ***User Class Analysis***
Velocity  7200  1812 3.1 0.8 0.8   0 4.0   0 5.4  32   0  44   0   0  15   0
REDHAT    8100  6222 0.5 0.0 1.7   0  10   0   0   0   0  78   0 1.5 8.6   0
SUSE      6300  5190 3.8 0.4 1.9   0  11   0   0   0   0  48   0 0.6 8.5   0
ORACLE    2700  2694 1.3 0.1 4.9   0  12   0   0   0   0  61   0 2.3  18   0
GPFS      2700  2700 5.9 0.0 2.6   0  21   0   0   0   0  48   0 0.1  22   0
TheUsrs  37605 17585 5.2 0.7 4.8 0.0 8.4   0  10 0.2 0.0  56 0.4 1.0  20 0.0
 ***Top User Analysis***
RHOSBOOT  3600  1873  22 0.3  13   0 2.5   0   0   0   0  52 1.0 4.5 5.0 0.1
```

User Storage – Case Study

Overall Subsystem Activity: ESASSUM / ESAMAIN

- Changes: Paging rates at 18:30 / CPU 18:45
- User "Resident" drops?

```
Report: ESASSUM      Subsystem Activity                                Velocity Sof
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78    First record
-------------------------------------------------------------------------------
         <---Users----> Transactions <Processor> Storage (MB) <-Paging--> <----
         <-avg number->    Per  Avg. Utilization Fixed Active <pages/sec> <-DAS
Time       On Actv In Q Minute  Resp Total Virt.  User Resid. XStore DASD Rate
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ---- ----

17:00:00  111   73 41.1  439.4 0.468   141   133  81.5  13219      0   90  139
17:15:00  111   74 38.1  449.5 0.410   132   125  81.5  13209      0  114  147
17:30:00  111   72 40.7  427.7 0.391   142   134  81.5  13202      0   75  137
17:45:00  111   72 39.5  446.8 0.563   146   139  81.5  13194      0   83  139
18:00:00  111   74 41.7  449.3 0.420   143   136  81.6  13202      0  120  157
18:15:00  113   75 40.7  445.1 0.494   101    95  81.5  13206      0  119  146
18:30:00  116   76 43.9  385.4 0.631   179   134  81.5  13078      0  43K  12K
18:45:00  116   80 57.2  303.8 0.833   341   270  81.2  12861      0  69K  22K
19:00:00  116   79 57.1  275.2 0.930   445   352  81.1  12863      0  81K  29K
19:15:00  115   80 54.7  307.7 0.748   313   249  81.2  12878      0  50K  17K
19:30:00  116   80 61.7  204.7 1.691   465   349  81.2  12863      0 101K  38K
```

Paging Happens When Storage Overcommitted

Paging subsystem, high rates, high queue, SSD (Need PAV, prewrite)

```
Report: ESAPSDV      Page And Spool Device Activity                  Velo
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78      Firs
-------------------------------------------------------------------------
            <----Paging/Spooling--------->       </Sec><Device->
 Dev        <------Slots-------> <per sec>        SSCH Serv Resp %Alloc
 No. Serial Avail Used %Use  Max Read Writ Queue +RSCH Time Time Select
---- ------ ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
18:30:00
Page Devices
2270 VM4P1A 1803K 806K   45   1M 2329 1559   9.1 242.8  0.0  1.6  100.0
2181 VM4P11 1803K 993K   55   1M 2788 1557  28.5 238.5  0.1 13.6  100.0
2381 VM4P12 1803K 993K   55   1M 2816 1555   4.3 239.4  0.0  0.2  100.0
2182 VM4P13 1803K 997K   55   1M 2829 1556  11.7 247.5  0.0  1.4  100.0
2382 VM4P14 1803K 993K   55   1M 2787 1556  13.1 250.5  0.0  1.9  100.0
2183 VM4P15 1803K 994K   55   1M 2830 1555   2.1 240.8  0.0  0.0  100.0
2383 VM4P16 1803K 995K   55   1M 2833 1555  17.3 241.9  0.1  6.6  100.0
2184 VM4P17 1803K 990K   55   1M 2749 1556  12.3 235.6  0.0  0.4  100.0
2384 VM4P18 1803K 992K   55   1M 2815 1556  33.9 236.0  0.0  2.7  100.0
207F VM4P19 1803K 809K   45   1M 2361 1558   2.4 242.6  0.0  0.0  100.0
            ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
Total Page    18M  10M   53  13M  27K  16K

Spool Devices
2180 VM4S11 1803K 414K   23 439K 42.5 31.2     0  96.3  0.1  0.1  100.0
2380 VM4S12 1803K 286K   16 327K  103 82.2     0 131.2  0.1  0.1  100.0
            ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
Total Spl   3606K 701K   19 766K  145  113
```

User Storage – Raw Metrics

User Storage analysis – ESAUSR2 ("R" is for "raw")

- Totals of all users to understand system impact
- Changes: resident, paging allocated, paging I/O
- User resident dropped?
- "Reserved" is for zWRITE to ensure performance data

Report: **ESAUSR2** User Resource Utilization Velocity S

	<---CPU time-->			<**Main Storage (pages)**>				<**------Paging (pages)-------**>				
UserID	<(seconds)>		T:V	<**Resident**>		Lock		<---**Allocated**--->		<---I/O--->		
/Class	Total	Virt	Rat	Totl	Activ	-ed	Resrvd	Total	ExStg	**Disk**	**Read**	**Write**
17:30:00	1230	1209	1.0	3.4M	3380K	6956	5000	7853K	0	7853K	17840	7098
17:45:00	1268	1247	1.0	3.4M	3378K	6906	5000	7852K	0	7852K	24890	9267
18:00:00	1243	1221	1.0	3.4M	3380K	6896	5000	7854K	0	7854K	44857	19867
18:15:00	877.0	857.8	1.0	3.4M	3381K	6906	5000	7857K	0	7857K	39079	26947
18:30:00	1307	1205	1.1	3.3M	3348K	7048	5000	9147K	0	9147K	24M	14M
18:45:00	**2589**	**2431**	**1.1**	**3.3M**	**3292K**	**7182**	**5000**	**13M**	**0**	**13M**	**43M**	**19M**
19:00:00	3384	3171	1.1	3.3M	3293K	7103	5000	16M	0	16M	54M	19M
19:15:00	2406	2244	1.1	3.3M	3297K	7198	5000	17M	0	17M	32M	12M
19:30:00	3391	3137	1.1	3.3M	3293K	7278	5000	17M	0	17M	66M	24M

User Storage – Rates / Percents

User Storage analysis – ESAUSP2

- Same data - rates / percents vs total counts
- Totals of all users
- Changes: resident, paging allocated, paging I/O

```
Report: ESAUSP2      User Resource Rate Report                        Velocity
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78       First re
-------------------------------------------------------------------------------
         <---CPU time--> <----Main Storage (pages)-----> <-Paging (pages)->
UserID   <(Percent)> T:V <Resident> Lock <-----WSS-----> Paged <Pgs/Second>
/Class   Total  Virt Rat Totl Activ  -ed Totl Activ  Avg 2Disk  Read Write
-------- ----- ----- --- ---- ----- ---- ---- ----- ---- -----   ---- -----
17:30:00 136.6 134.3 1.0 3.4M 3380K 6956 4.0M 3957K  36K 7853K  19.8   7.9
17:45:00 140.9 138.6 1.0 3.4M 3378K 6906 4.0M 3962K  36K 7852K  27.7  10.3
18:00:00 138.1 135.6 1.0 3.4M 3380K 6896 4.0M 3948K  36K 7854K  49.8  22.1
18:15:00 97.43 95.29 1.0 3.4M 3381K 6906 4.0M 3973K  35K 7857K  43.4  29.9
18:30:00 145.3 133.9 1.1 3.3M 3348K 7048 3.9M 3909K  34K 9147K 27057 15496
18:45:00 287.6 270.1 1.1 3.3M 3292K 7182 5.4M 5363K  46K   13M 48060 21144
19:00:00 376.0 352.3 1.1 3.3M 3293K 7103 4.3M 4263K  37K   16M 59528 20806
19:15:00 267.3 249.3 1.1 3.3M 3297K 7198 3.7M 3682K  32K   17M 35972 13731
19:30:00 376.8 348.6 1.1 3.3M 3293K 7278 3.2M 3161K  27K   17M 73843 27041
```

User Storage analysis – ESAUSP2 (percent/rate)

- Analyze by user
- RHOS* users paging too much to get work done
- RHOS* is OpenShift installation

```
Report: ESAUSP2       User Resource Rate Report                      Velocit
-------------------------------------------------------------------------------
         <---CPU time--> <----Main Storage (pages)-----> <-Paging (pages)-
UserID   <(Percent)> T:V <Resident> Lock <-----WSS-----> Paged <Pgs/Second
/Class   Total  Virt Rat Totl Activ  -ed Totl Activ  Avg 2Disk  Read Write
-------- ----- ----- --- ---- ----- ---- ---- ----- ---- -----  ---- ----
18:30:00 145.3 133.9 1.1 3.3M 3348K 7048 3.9M 3909K  34K 9147K 27057 15496
 ***Key User Analysis ***
TCPIP     0.15  0.05 3.0 1422  1422  601  817 817.3  817  7750  43.4   8.6
***User Class Analysis***
Velocity  5.82  5.43 1.1 3763  3598    5 4593  4271  534 14472 137.4  57.0
SUSE     20.17 19.28 1.0 112K  112K 1534 193K  193K  32K 1048K  2754 828.5
ORACLE    4.66  3.84 1.2 195K  195K  734 381K  381K 190K  473K  2895 936.7
GPFS     12.51 11.68 1.1 195K  195K  975 439K  439K 146K 1332K  4008  1383
TheUsrs  95.37 89.07 1.1 2.6M 2615K 1145 2.5M 2472K  80K 5017K 12958 11022
 ***Top User Analysis***
RHOSBOOT 39.91 38.51 1.0 727K  727K   30  99K 98642  99K  454K  1175  2346
RHOSCP2   8.92  8.20 1.1 250K  250K   19 116K  116K 174K  201K 997.0  1965
RHOSCP1   8.78  8.05 1.1 252K  252K   19 126K  126K 189K  205K 967.6  2005
RHOSCP3   7.83  7.04 1.1 161K  161K   28  48K 47842  80K  125K  1230  1157
```

User Storage – Classify Servers of Interest

ESAUSPG – Group relevant users, show MB

- OpenShift class takes all storage, flushes all others
- OpenShift is just the installation, no work “yet”
- Can now quantify the impact of the “OpenShif” servers as group

```
Report: ESAUSPG      User Storage Analysis                    Velocity Software
-------------------------------------------------------------------------------
         <-Storage Occupancy in MegaBytes-> <--Main Storage page Read/Write-->
UserID   <---Main Storage---> <--Paging---> <-Page Writes to:--> <Page Reads:>
/Class    Total    >2gb    <2GB  Xstor   DASD   Xsto   Disk   Migr  Xstor  Disk
-------- ------ ------ ----- ----- ----- ----- ----- ------ ---- ------
19:55:00  12883  10948   1935      0  55185      0 406320      0       0  1130K
 ***Key User Analysis ***
TCPIP         3      3      1      0     32      0    757      0       0   4056
 ***User Class Analysis***
Servers       5      4      2      0    185      0    526      0       0   1358
KeyUsrs       4      3      1      0     97      0    400      0       0   1165
ZVPS         10      9      2      0     77      0   2948      0       0   6435
Lnx         122    102     20      0   4186      0  47575      0       0 190345
OpenShif  12308  10474   1834      0  34811      0 206426      0       0 451662
TheUsers    428    352     76      0  15748      0 147274      0       0 473648
```

User Storage Fully Instrumented (ESAUSTR)

User Resident Storage = Active + IBR + Agelist

- System thresholds maintained, source of pages changed
- UFO Active: User frame list
- IBR: Invalid But Resident
- AgeList: Ready for “steal”
- Unreferd: (Block read, no reference)
- “NoScan”: Demand scan restricted by set reserved
- DiagRlse: Diagnose 10, 214
- PreWrite: IBR already backed before reclaim
- PreWrite: Agelst – pagefaulted (heading confusing)
  - “Wasted write”, was written, but then page faulted and was needed

```
Report: ESAUSTR      User Storage Analysis                Velocity Software Corporate   ZMAP 5.1.2 04/16/21   Page  232
-------------------------------------------------------------------------------------------------------------------------
         <--------Virtual Server Storage (Pages)--------->  <Resident> Page  <-----Page Rates / Second-----> NoScan
UserID   Size  Alloc Resi- UFO   <-----IBR----> <AgeList> <Unreferd> slots Made IBR AgeLst <PreWrite> Diag  SET  Steal
/Class               dent Activ TOT  <2gb >2gb <2gb >2gb <2gb >2gb  used  IBR Refd Refd   IBR Agelst Rlse Rsrvd Weight
-------- ----- ----- ----- ----- ---- ---- ---- ---- ---- ---- ---- ----- ---- ---- ----- ---- ------ ---- ----- ------
```

User Storage Fully Instrumented (ESAUSTR)

User Resident Storage = Active + IBR + Agelist

- Just user totals, easy to see significant changes
- System thresholds maintained, source of pages changed
- UFO Active – User Frame list
- IBR – Invalid But Resident
- AgeList – Ready for “steal”
- Only few of “PreWrite AgeIst” were “wasteful writes”

```
Report: ESAUSTR      User Storage Analysis                Velocity Software Corporate   ZMAP 5.1.2 04/16/21   Page  232
-------------------------------------------------------------------------------------------------------------------------

         <--------Virtual Server Storage (Pages)--------->  <Resident> Page  <-----Page Rates / Second-----> NoScan
UserID   Size  Alloc Resi- UFO   <-----IBR----> <AgeList> <Unreferd> slots Made IBR AgeLst <PreWrite> Diag  SET  Steal
/Class                dent Activ TOT  <2gb >2gb <2gb >2gb <2gb >2gb  used  IBR Refd Refd   IBR Agelst Rlse Rsrvd Weight
-------- ----- ----- ----- ----- ---- ---- ---- ---- ---- ---- ---- ----- ---- ---- ----- ---- ------ ---- ----- ------
17:30:00 13.4M 8785K 3383K 3307K 6334 1145 5189  13K  57K 7193  26K 7853K  3.6  1.6   1.3  1.4    0.0  8.3   2.1    111
17:45:00 13.4M 8786K 3381K 3305K 6301 1158 5143  12K  58K 6251  23K 7852K  5.5  2.4   1.8  2.4    0.0  7.9   3.3    111
18:00:00 13.4M 8790K 3383K 3307K 6698 1312 5386  12K  57K 4387  17K 7854K  7.5  3.5   2.9  3.3    0.0 17.1   4.5    111
18:15:00 16.9M 8797K 3384K 3307K 7267 1428 5839  13K  57K 4691  19K 7857K  8.6  3.2   3.2  4.2    0.0 18.9   5.3    113
18:30:00 27.3M 10.8M 3349K 3274K 5988 1178 4809  12K  57K 3759  13K 9147K 3559 1038  1321 2401   19.1  154  5984    116
18:45:00 30.2M 14.6M 3293K 3215K 7317 1394 5923  12K  58K 4502  17K   13M 5188 1540  2131 3623   38.8  213  8500    116
19:00:00 30.2M 17.2M 3293K 3218K 5637 1198 4439  13K  57K 5445  18K   16M 6438 1891  2797 4213   53.6  277 10697    116
19:15:00 32.1M 17.5M 3297K 3221K 5878 1196 4681  13K  57K 4709  16K   17M 4182 1245  1854 2613   27.0  182  6246    115
```

Tuning User Storage

Tuning process:

- Evaluate requirements of virtual machine
- Evaluate value from business perspective

Options for tuning user storage:

- Use SET REServed to fence storage for a user
- If TCPIP is paged out, who is delayed?

If ZWRITE paged out, what happens?

- CP SET RESERVED ZWRITE 500
- CP LOCK USERID ZWRITE 0 1000 TO LOGICAL

What servers or users should be protected?

- Servers/Users that have time dependencies
- Servers with multi-user impact (RACF, TCPIP)

Reserved Storage

Protect a user:

- CP LOCK USERID ZWRITE 0 1000 logical
- CP SET RESERVED ZWRITE 4500

```
Screen: ESAUSR2  Velocity Software - VSIVM4      ESAMON 5.121 05/15 17:55-17
1 of 3  User Resource Utilization                CLASS VELOCITY USER 8562 04

                  <---CPU time--> <--------Main Storage (pages)---------->
          UserID  <(seconds)> T:V <Resident>   Lock <-----WSS----->
Time      /Class   Total  Virt Rat Total Activ   -ed Total Actv  Avg Resrvd
-------- -------- ----- ----- --- ----- ----- ----- ----- ---- ---- ------
17:56:00 ZALERT    0.95  0.94 1.0   301   301     0   301  301  301      0
         ZVWS      0.51  0.50 1.0  1229  1229     1  1228 1228 1228      0
         ZWRITE    0.12  0.11 1.0  4102  4102  3576   506  506  506   4500
         ZTCP      0.07  0.06 1.3  2154  2154     1  2132 2132 2132      0
         ESAWEB    0.00  0.00 1.4    68    68     2    66   66   66      0
         VSIMAINT  0.00  0.00 2.1   200   200     0   200  200  200      0
         ZADMIN    0.00  0.00 2.0   379   379     0   379  379  379      0
         ZSERVE    0.00  0.00 1.1   154   154     0   154  154  154      0
         ZMON         0     0   0     4     0     0     4    0    0      0
```

Saved Segments

Monitor DCSS requirements, not “82”

Monitor “storage” domain 3 is early in monitor cycle

Why is “CMS” there twice?

- If you resave CMS, only on IPL will new one be used

Report: **ESADCSS** Velocity

		<-Users->		<-------Number of Pages fo				
	Spool		Non-		Non-	<--resident-->		
Name	ID	Shrd	Shrd	Saved	Data	<2GB	>2GB	HOST
03/13/23 16:01:00								
CMS	48	9	0	1298	0	132	957	0
CMS	1	10	0	1298	0	4	1078	0
MONDCSS	38	2	0	0	16K	0	**82**	81
ZALERT	22	1	0	0	256	0	1	0
ZMON	11	3	0	0	16K	6094	0	0
Totals:		118	0	7460	33K	9930	2842	81

Saved Segment: MONDCSS

ESAHDR shows final count of each monitor sample

Protect the monitor DCSS

- SET RESERVED DCSS MONDCSS 5000

```
Report: ESAHDR       z/VM Monitor Analysis                          Velo
-------------------------------------------------------------------------
Monitor Segment Analysis
Sample Configuration DCSS used:   1676 pages, 147456K -154159K  ( 91%)
Sample DCSS storage used:         4268 pages, 154800K -171869K  ( 67%)
Event Configuration DCSS used:       1 pages, 180224K -180225K  (  1%)
Event  DCSS storage used:         1680 pages, 180248K -186965K
Times Event storage wrapped:         1
```

User Summary Storage Measurement

Virtual Machine Storage requirements provided:

- For system
- For user class / workload
- For virtual machine

Page rates provided

Process:

- ESAXACT – review impacts
- ESAUSR2 / ESAUSP2 / ESAUSPG for VM requirements

Review:

- Virtual machine sizes (many are too large)
- VDISK (very very fast)
- Set reserved / locked
- Set Agelist 5%

Mapping Storage Summary

Storage types:

- CP Fixed Storage (no tuning options)
- CP Non Pageable: Storage made non–pageable by CP (some tuning)
- The remaining pages are Dynamic Paging Area(DPA)

Dynamic Paging Area:

- Virtual Machines
- Shares storage (DCSS/NSS)
- System address spaces
- Virtual Disks
- MDC
- Available

Maximizing Pageable Storage

Limit virtual machine sizes

- PGMBKs – cost 8mb (PTRM address space) per virtual GB
- PTRM address space is pageable

Limit the amount of main storage used by MDC:

- SET MDCACHE STORAGE **min**M maxM

Storage Summary

Full storage map available (ESASTR1)

- System
- User
- Address space / VDISK
- MDC

Tune users as able

- Reduce sizes saves control blocks, real storage
- Reserve storage for critical functions