z/VM Storage Analysis and Tuning

Velocity Software Inc.
196-D Castro Street
Mountain View CA 94041
650-964-8867

Velocity Software GmbH
Max-Joseph-Str. 5
D-68167 Mannheim
Germany
+49 (0)621 373844

Objectives:

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

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

Linux Traditions are Challenging

Non virtualized Linux:

- Has storage requirements
- Unlimited storage available
- Very little tuning and management required
- Simple rules, reduce I/O:
  - Swapping to disk is bad
  - More storage caches more data

Virtualized Linux

- Share storage – to what level?
- I/O much faster on z technology, less penalty for I/O
- Trade-off between storage and I/O different
- Swapping can be incredibly fast
- Teaching how to share is important for kids and admins

Linux is Challenging – 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"

Inactive storage? Linux Storage is not idle

- Extra storage used to cache data and programs

Inactive servers? 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 legitimately paged out of real storage?

The page most likely to be needed by Linux:

- Is most likely to have been paged out

Managing Linux Storage in z/VM

Linux is NOT virtual friendly, changes requirements

Determining pages for page out:

- **Active server?** Can not know if server is working or **polling**
- Take least recently used, non modified, non referenced
- If server is working, will re-use storage most likely paged out
- Fast page-in very important

Strategy / best practices in past if overcommit high

- Expanded Storage was used for "30 second test case"
- Minimum 20% of storage reconfigured to Expanded Storage
- Page-in from expanded storage was synchronous, FAST
- Pages migratable to disk after 30 seconds unreferenced

- New strategy is IBR (z/VM 6.3). VERY LIMITED. Go the max!

VELOCITY SOFTWARE

z/VM STORAGE and Paging Architecture Moving Target

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

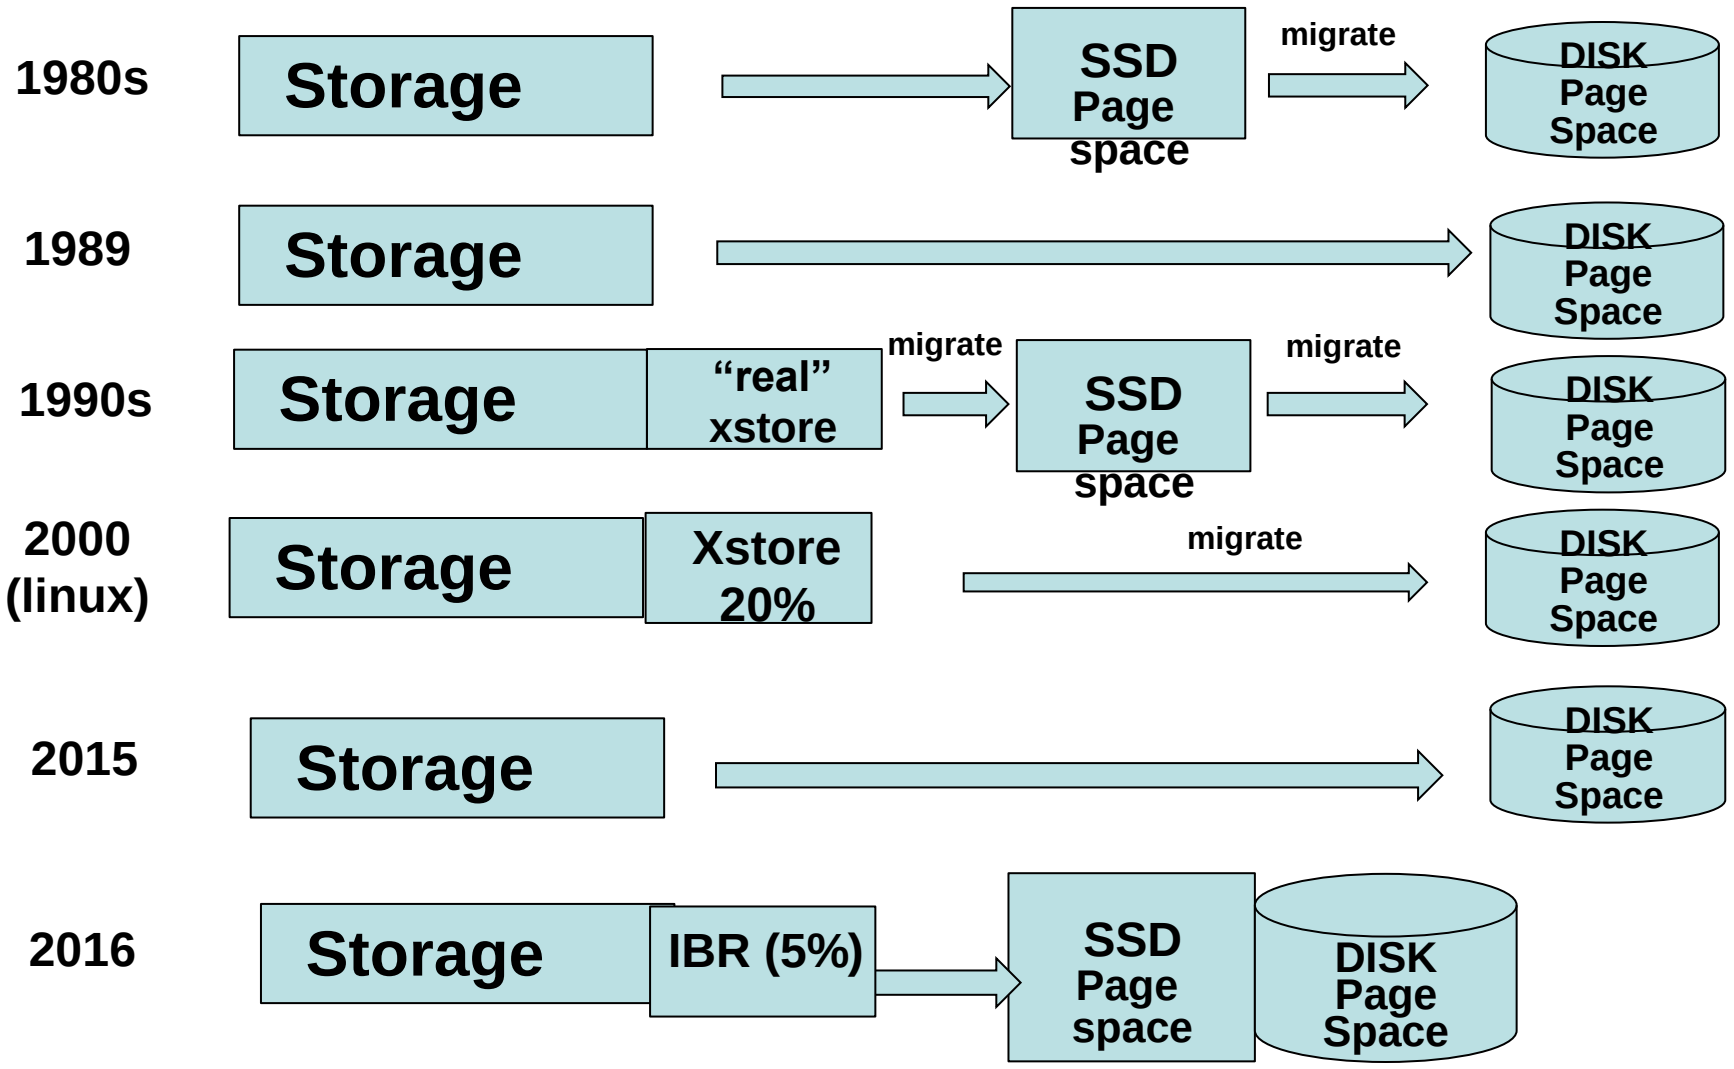

VELOCITY SOFTWARE

z/VM User Storage “test”, Replaces ExStore

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 seem indiscriminate:

- 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 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

- 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 (default)

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

VELOCITY SOFTWARE

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

CP Storage Management – Frame Table

The Frame Table (CP Fixed Storage) is the usage map of REAL (main) storage.

- One entry (32 bytes) for each page (2 page per MB Real)
- All entries chained on specific list for (current) type of page
- If 196 GB of storage, 1,536 MB free storage for page frame tables

Resident pages for a user are dynamic Frame Table entries chained together in a List: the User Frame Owned List (UFOL).

CP address spaces and Shared Segment (NSS/DCSS) pages are similarly chained. (SUFO) (not stolen so much)

z/VM 6.3 adds VUFO (virtual disk frame owned list)

When storage is needed by ANY process, it is acquired from the Available List, another chain of Frame Table entries.

If available list empty, big problem….

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 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
```

User Storage – Raw Metrics

User Storage analysis – ESAUS**R**2

- 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
```

VELOCITY SOFTWARE

User Storage – Rates / Percents

User Storage analysis – ESAUS**P**2

- 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

- 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"
- "NoScan", demand scan restricted by set reserved

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 /Class	Size	Alloc	Resi- dent	UFO Activ	TOT	<-----IBR----> <2gb	>2gb	<AgeList> <2gb	>2gb	<Unreferd> <2gb	>2gb	slots used	Made IBR	Refd	AgeLst Refd	<PreWrite> IBR	Agelst	Diag Rlse	SET Rsrvd	Steal 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 have locked storage?

- Servers/Users that have time dependencies
- Servers with multi-user impact

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
```

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 / lock
- Set agelist 5%

VELOCITY SOFTWARE

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

Storage Map – CP Areas

System Storage – total storage assigned to LPAR

CP Fixed Storage (no tuning options)

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

Storage Map, ESASTR1

Storage Map to show storage (14GB) utilizations

- User resident should be major use

Capture ratio shows accuracy

Overcommit means paging will happen

```
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
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- ----- ---- -----     -----

04/15/21
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
19:15:00   115 3670016  2878 20797 1040 28672  3063  2232 52409  3297K    29522   192K     0    22    73  16K 8.754     0.988
19:30:00   116 3670016  2878 20809 1031 28672  3069  2235 52450  3293K    29065   196K     0     0     6  16K 9.363     0.988
```

Capture Ratio Storage Drill Downs

Capture ratio is accumulation of knowns / total storage
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 – What changed at 18:30?

- System Address space use increased – page tables
- VDISK Resident – went down, paged out
- MDC Resident – went down, freed
- Over Commit Ratio – went up, more virtual machines
- User resident went down, no room

Result of many large virtual machines logging on

```
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
19:15:00   115 3670016  2878 20797 1040 28672  3063  2232 52409  3297K    29522   192K     0    22    73  16K 8.754     0.988
19:30:00   116 3670016  2878 20809 1031 28672  3069  2235 52450  3293K    29065   196K     0     0     6  16K 9.363     0.988
```

Shared Storage (DCSS) Requirements

```
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
```

"pages resident" show how many pages in storage

32k pages in shared storage for monitoring:

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

VELOCITY SOFTWARE

Pageable PGMBK's

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 (100GB 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.

Maximizing Pageable Storage

Limit virtual machine sizes

- PGMBKs – cost 8mb (ptrm address space) per virtual GB

Limit the amount of main storage used by MDC:

- SET MDCACHE STORAGE **minM** maxM

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
```

CP Page Tables Requirements

Virtual disks are system-owned address spaces.

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

```
Report: ESAASPC    Shared Address Space
Monitor initialized: 04/15/21 at 00:00:00
--------------------------------------------------------------------------------

                                 <--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
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
```

Minidisk cache defaults to "all of it"

- MUST BE CONTROLLED!!! Very common configuration error

Example is "very constrained", why?

- MDC consumes almost 1M pages (4GB) out of total 2M pages….
- Set to 128MB SET MDC STORAGE 128M 128M

```
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
01:50:12    34 2064368  2970 27734 3802 16128  3145   111  7897  1178K     1714  20460 19571  4882  787K 1913
```

Minidisk cache defaults to “all of it” or none of it…

Minidisk cache should be enabled for “velocity”!!!

- CMS makes good use of MDC
- zVPS will slow down if no MDC
- CP SET MDC STORAGE 32M 32M

```
Report: ESAMDC       Minidisk Cache Analysis                          Ve
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78       Fi
-------------------------------------------------------------------------
         <----Load---->     <IO per><Insertions> <-----Main Storage M
         <-Users-> Tran Hit <second> Usr Per Not <-Sizes (MB)--> </Se
Time     Actv In Q /sec Pct rds hits Max Min Ald Avg MIN MAX Obj Stls
-------- ---- ---- ---- --- --- ---- --- --- --- --- --- --- --- ----

18:00:00   74 41.7  7.5  96  32 30.8  2K 1.0   0  77   0 256 222  0.1
18:15:00   75 40.7  7.4  98  27 26.2  2K 1.7   0  75   0 256 244  0.1
18:30:00   76 43.9  6.4  91  27 24.6 913 2.6 1.3  48   0 256 112 25.6
18:45:00   80 57.2  5.1  72  21 15.5 150 4.2  10 0.2   0 256 1.5 20.7
19:00:00   79 57.1  4.6  75  21 16.0 150 4.0 5.3 0.5   0 256 5.8 27.5
19:15:00   80 54.7  5.1  76  23 17.5 150 4.5 7.1 0.3   0 256 1.4 15.7
19:30:00   80 61.7  3.4  68  20 13.3 150 4.2 5.3 0.0   0 256 1.3 25.4
```

VELOCITY SOFTWARE

CP Storage Management - Free Storage

Free storage (now in SXS) used for recording for:

Erep, Accounting, Symptom

Issue CP QUERY RECORDING

- Very common configuration error

To stop recording, free up storage:

- CP RECORDING ACCOUNT OFF PURGE
- Or disable it in SYSTEM CONFIG

```
q recording
RECORDING     COUNT      LMT USERID   COMMUNICATION
EREP      ON  00000088   002 EREP     INACTIVE
ACCOUNT   ON  00044232   020 DISKACNT INACTIVE
SYMPTOM   ON  00000000   002 OPERSYMP ACTIVE
```

Free Storage Recording

To stop recording, free up storage:

- CP RECORDING ACCOUNT OFF PURGE
- Or disable it in SYSTEM CONFIG.

```
cp recording account off purge
HCPCRC8058I User BARTON has purged 00044234
    records from the *ACCOUNT queue.
Command complete
```

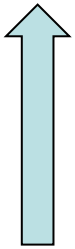

```
Screen: ESASTR1  Velocity Software
2 of 2  Main Storage Analysis
<----------------------Pages-------
          Systm  User  NSS/DCSS  <-AddSpace>
Time      ExSpc Resdnt Resident  Systm User
-------- ----- ------ --------  ----- -----
16:05:00  1059  28184     9210   3386     0
16:04:00  1060  24539     9215   3386     0
16:03:00  2279  27702     9207   3386     0
16:02:00  2274  25189     9209   3374     0
16:01:00  2259  23617     9209   3374     0
```

VELOCITY SOFTWARE

Free Storage analysis – zmon esastr1 (split screen)

```
Screen: ESASTR1  Velocity Software                    ESAMON 4.240
2 of 2  Main Storage Analysis
         <----------------------Pages----------------------->
         Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag
Time     ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98
-------- ----- ------ --------  ----- -----  ----- ----- ----
16:05:00  1059  28184     9210   3386     0      0 64966 1037
16:04:00  1060  24539     9215   3386     0      0 64919 1037
16:03:00  2279  27702     9207   3386     0      0 64905 1037
16:02:00  2274  25189     9209   3374     0      0 65407 1037
16:01:00  2259  23617     9209   3374     0      0 65407 1037
====>
Screen: ESASTR1  Velocity Software                    ESAMON 4.240
1 of 2  Main Storage Analysis
         <------------------Pages------------------> Capt-
         System  Fixed Non-  Free Frame <Available >   ure
Time     Storage Store Pgble Stor Table <2gb  >2gb    Ratio
-------- ------- ----- ----- ---- ----- ----- ------ -----
16:05:00 1310720  2244  3143    4 10240  473K 694793 0.995
16:04:00 1310720  2244  3143    4 10240  473K 694848 0.995
16:03:00 1310720  2244  3146    4 10240  473K 693692 0.995
16:02:00 1310720  2244  3167    4 10240  473K 693716 0.995
16:01:00 1310720  2244  3166    4 10240  473K 693714 0.995
```

CP Storage Management - Trace Table

Trace table for Master processor and each other real CPU

Trace Table size: Master = Min(100 pages)

Trace table size = master + (NCPU-1) * .75 * Master

"CP SET TRACEFRAMES MASTER 100 ALTERNATE 75 PERCENT"

- Must be at least 3
- **Storage is "below the line" - What if 40 threads in LPAR? (12m)**
- Size displayed on ESASTRC, included in SXS

```
CP SET TRACEFRAMES MASTER 2000
```

```
Screen: ESASTR1  Velocity Software
1 of 2  Main Storage Analysis
<------------------Pages-----------
          System  Fixed Non-  Free Frame <Avai
Time      Storage Store Pgble Stor Table <2gb
-------- ------- ----- ----- ---- ----- -----
13:56:00 1310720  2244  5088    4 10240  468K
13:55:00 1310720  2244  3188    4 10240  470K
13:54:00 1310720  2244  3197    4 10240  470K
```

Locked Storage for QDIO

Locked Storage > 2GB + < 2GB (see ESAUSPG)

- Originally QDIO only below 2gb line, now > 2gb
- 8mb QDIO Buffers locked per server

```
Report: ESAUSR2       User Resource Utilization
-----------------------------------------------
         <---CPU time--> <Main Storage (pages)>
UserID   <(seconds)> T:V <Resident> Lock
/Class    Total  Virt Rat Totl Activ  -ed Resrvd
-------- ----- ----- --- ---- ----- ---- ------
15:03:00 31.62 30.80 1.0  50M   50M  12K   5000
 ***Top User Analysis***
S1C7NA2Z 17.42 17.38 1.0 4.8M 4784K 2156      0
S1C7NA1Z  5.28  5.22 1.0 8.4M 8378K 1666      0
S1SSSA1Z  1.70  1.25 1.4 1.8M 1761K 53.0      0
S1BSAA1Z  1.33  1.31 1.0 2.2M 2199K 2086      0
S1SLTA1Z  1.19  1.17 1.0 3.8M 3842K  588      0
D1SAPA1Z  1.19  1.18 1.0  15M   15M  407      0
S1BSAA2Z  0.97  0.95 1.0 2.1M 2104K 2445      0
ZWRITE    0.06  0.06 1.0 1896  1896  1.0    5000
ZTCP      0.06  0.05 1.1 1430  1430  1.0      0
```

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
- It's only purpose is to gauge paging requirements

Linux Storage Analysis

High Level, UCD

Standard Linux system storage – high level, ESAUCD2

Linux system storage

Linux system storage details - ESALNXR

Linux process storage

By process

Linux System Storage Reporting

Preview, Linux Storage

- Storage overview (ESAUCD2)
- Storage Details (ESALNXR)
- Process Storage (ESALNXP)

```
Report: ESAUCD2      LINUX UCD Memory Analysis Report                Velocity Sof
-------------------------------------------------------------------------------
Node/    <-------------------------Storage Sizes (in MegaBytes)----------------
Time/    <--Real Storage--> <-----SWAP Storage----> Total <-------Storage in Us
Date     Total  Avail Used  Total Avail Used  MIN   Avail CMM    Buffer Cache O
-------- ------ ----- ----- ----- ----- ----- ----- ----- ------ ------ ----- -
18:30:00
*** Nodes *******
lxsugar   999.4   7.6 991.8 154.9 151.3   3.6  15.6 158.9      0   85.7 648.1 2
mail     8112.8  2318  5795     0     0     0  15.6  2318      0  639.8 907.9
mongo01  3849.8 983.3  2866 371.9 309.6  62.3  15.6  1293      0  150.6  1130
opensuse  15846 160.1 15686  8192  8192   0.3  15.6  8352      0 1524.5  8392
REDHAT6X  996.8  13.8 983.0 495.8 380.4 115.5  15.6 394.2      0  114.7 724.1 1
redhat7   994.0 411.5 582.4  1124  1124     0  15.6  1535      0    1.1 472.6 1
rhel64v   996.1  66.3 929.8  2047  2034  12.5  15.6  2101      0  103.3  39.6 7
rhel7v   2002.3 101.2  1901  2064 766.0  1298  15.6 867.2      0      0 253.0
sles11v3  868.8  88.0 780.8  2046  1406 639.7  15.6  1494      0    3.3  27.7 7
sles11x3  493.2 132.8 360.4 867.9 867.9     0  15.6  1001      0  141.6 149.5
```

Linux System Storage Details Reporting

Preview, Linux Storage details

```
Report: ESALNXR      LINUX RAM/Storage Analysis Report                Velocity Sof
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78       First record
-------------------------------------------------------------------------------
Node/    <--------Memory in megabytes-------------> <-Kernel(MB)-> <-Buffers(MB
                  <---Cache----><---Anonymous---> Stack<-Slab-->
Time     Total Free Size Actv Swap Total Actv Inact Size Size SRec Size Dirty B
-------- ----- ---- ---- ---- ---- ----- ---- ----- ---- ---- ---- ---- ----- -
18:30:00
mongo01   3850  983 1130  939 26.9  1464 1333 201.3  3.5 57.3 46.3  151   0.7
opensuse 15846  160 8392 4346  0.1 915.4  426 512.0  6.2  554  477 1525   0.0
REDHAT6X 930.4 13.0  676  308  2.5  41.8 62.0 154.7  2.7 51.5 41.0  107   0.0
redhat7  994.0  412  473  328    0  40.8 40.9  56.0  2.6 46.8 26.9  1.1     0
rhel64v  996.1 66.2 39.6 74.6  1.2  14.0  1.1  13.9  1.8  101 42.9  103   0.0
rhel7v    2002  101  253  105 10.0  1437 1142 407.7  4.0  112 67.7    0   0.0
sles11v3 868.8 88.0 27.7 17.4 51.6 106.0 44.6  69.6  2.6 35.6  8.5  3.3   0.0
sles11x4 492.8  102  235  160    0  26.8 26.8   0.7  1.4 31.2 23.2 78.1   0.0
sles12    3374  124 2259 1557  2.7 534.0  483 459.6 30.8  153 51.8  110   0.1
sles12v  995.6  101  440  206  8.1 339.2  162 230.1  2.0 73.9 51.2  0.0   0.0
sles12x3 820.9  182  334  377    0  38.5 38.7  42.2  2.5 88.9 70.5  154   0.0
```

VELOCITY SOFTWARE

Linux System Storage Details Reporting

Preview, Linux Process Storage details

```
Report: ESALNXP      LINUX HOs Statistics Report          Velocity Software Co
Monitor initialized: 04/15/21:00 on 8562 serial 040F78    First record analyze
-------------------------------------------------------------------------------
node/     <-Process Ident-> N<------CPU Percents----> <---------------Storage
 Name      ID     PPID   GRP  V Tot  sys user syst usrt  Size RSS Peak Swap Data
--------- ----- ----- ----- ----- ---- ---- ---- ---- ---- ---- ---- ---- ----
18:30:00
mongo01       0     0     0  14.8 1.18 13.2 0.03 0.31 7248 1544 113K  727  78K
 mongod   10889     1 10887  5.75 0.60 5.15    0    0 2653 1307  40K  429  37K
 java     51013  8515  8515  4.94 0.31 4.62    0    0 1665  155  16K    0  14K
 java     51596  8515  8515  3.61 0.20 3.41    0    0 1743  186 8985    0 8053
opensuse      0     0     0  10.0 8.75 1.26 0.00 0.01  33K 5900 537K    0  38K
 gsd-colo  1909  1791  1776  1.13 0.00 1.13    0    0  706   84  11K    0 1773
 VBoxHead 24298 24280 24298  8.61 8.61    0    0    0 5824 4237  87K    0 2089
REDHAT6X      0     0     0  0.72 0.34 0.27 0.07 0.05  16K 1205 227K  641  14K
rhel7v        0     0     0  2.46 0.41 1.69 0.25 0.11  43K 1643 676K  20K 252K
 java      2028     1  1321  1.22 0.04 1.18    0    0 3848  865  58K 2054  55K
sles11v3      0     0     0  0.65 0.19 0.46    0    0 6526  117 105K 9009  27K
sles12        0     0     0  4.60 0.72 3.84 0.03 0.02  76K 5518 1.0M 2918 178K
 ora_mmon  2596     1  2596  3.61 0.32 3.29    0    0  896  403  11K 16.3 1155
sles12v       0     0     0  0.52 0.16 0.32 0.01 0.03  15K  379 239K  10K 144K
```

Paging Subsystem Details Reporting

Paging subsystem, high rates, high queue, ssd

```
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
```

VELOCITY SOFTWARE

Storage Summary

Full storage map available (ESASTR1)

- System
- User
- Address space / vdisk
- MDC

Manage storage to meet requirements

- Trace table reduce
- Accounting off

Tune users as able

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