Case Study 4

Slow Response Times

VELOCITY
SOFTWARE

PROVEN PERFORMANCE

Velocity Software solves performance problems.

- As a valued customer, we want to pass this knowledge on to you.
- The following is a case study of a solved real-life performance issue.
- This case study will show:
    - The problem as reported by users
    - The problem observations
    - What was found in the Velocity Software data
    - What was suggested to the customer
    - If provided, follow up from the customer

VELOCITY
S O F T W A R E

P R O V E N   P E R F O R M A N C E

The Problem:

An LPAR running very large Linux guests using MongoDB were experiencing performance issues/slow response times

Problem Observations:

- Customers were reporting slow response times for their MongoDB applications

What the Data Showed (Configuration data)

ESAHDR – System Configuration showed:

- Up to date on z/VM release
- SMT is enabled
- Which is the master processor
- Running on IFLs

```
Report: ESAHDR        z/VM Monitor Analysis
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08
Monitor period:        7200 seconds (   2:00:00)
-----------------------------------------------------------------
ZMAP Release                                  5.1.3.1
Monitor file created:          04/25/22  00:00:00

z/VM Version:  7              Release 2.0 SLU 2102
TOD clock at termination                    20:00:00
.
System Identifier                           Vxxxxxx8
Machine Model/Type                      Z15:8561/401
Multithreading Status(SMT): Enabled
   Core Thread count:        2
   Enabled Count:            2

System Sequence Code            0000000000047A08
Processor   0 model/serial          8561-401   /087A08 Master
Processor   1 model/serial          8561-401   /087A08

Processor 39 model/serial          8561-401   /087A08

 CPU(GP) Capability Factor:     3085
 CPU(IFL) Capability Factor:     416
 CPU Cycles/ns:               5200
 CPU Cycles/ns (GP):           781
 Operating on IFL Processor(s)
 Channel Path Measurement Facility(CPMF) Extended is installed
```

ESAHDR – System Configuration (cont.) showed:

- Of the 108 IFL processors, only 15 were 'in use' – SMT may not be needed
- 616.6% out of 4000% 'in use' (only 6 threads out of 40)
- Lxxxx3F8 was the top CPU user at 70% (or 4.4 IFLs)

```
Totals by Processor type:
<--------CPU-------> <-Shared Processor busy>
Type Count Ded shared total assigned Ovhd Mgmt
---- ----- --- ------ ----- -------- ---- ----
CP      1    0      1   0.4      0.3  0.0  0.1
IFL   108    0    108  1501   1478.6 21.1 22.2

Number of logical partitions defined:          21

Main Storage installed (MB):              2867199
Main Storage Generated (MB):              2867199

Horizontal/Vertical Scheduling Configuration IFL CPUs
 UNPARKING set to Medium
 EXCESSuse moderate

CPU utilization:         616.6  of  4000%
CPU charged to users: 598.3%
System overhead:         18.4%
Capture ratio:          100.00%
.

Top users and user classes by CPU consumption:
                <-Relative->    <---Absolute Percent CPU--->
      UserID    <-Pct CPU-->   <Out of 100%>  <Out of 4000%>
      /Class    Used  Cum       Util   Cum      Util    Cum
      --------  ----  ----      ----   ----     -----   -----
   1. Lxxxx3F8  71.6   72       11.0    11      441.4    441
   2. Lxxxx3FD  13.3   85        2.1    13       82.2    524
   3. System     3.0   88        0.5    14       18.4    542
   4. Lxxxx3FA   2.2   90        0.3    14       13.7    556
   5. Lxxxx411   1.7   92        0.3    14       10.2    566
```

VELOCITY
SOFTWARE

PROVEN PERFORMANCE

ESAUSRC – User Configuration showed:

- There are multiple servers with 32 vCPUs with REL 3200
- The storage for each server adds up to more than real storage – which is more than necessary for each server

```
Report: ESAUSRC     User Configuration                Velocity Software Corporate   ZMAP 5.1.3
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08   First record analyzed: 04/25/22 15:00:00
Monitor period:         7200 seconds (   2:00:00)                 Last record:           04/25/22 17:00:00
---------------------------------------------------------------------------------------------------------
                                           <-------SHARE-------> <---CPU---><Status> <-MDC-> <Storage>
                  Account  ACI Grp <CP POOL> CPU <Normal> <--MAX-> Lim <Count>     Qck NO  NO <VM Size>
UserID   ClassID  Code     Name     PoolName Type Rel Abs Typ Shre -it Def On Mode SVM Dsp FS INS Dflt  Max
-------- -------- -------  --------  -------- ---  --- --- --- ---- --- --- -- ---- --- --- -- --- ---- ----
Lxxxx3B2 TheUsrs  Lxxxx3B2    .         .     IFL 1600  .   .    .   .  16 16 ESA   N   N  N   N  244G 244G
Lxxxx3FA TheUsrs  Lxxxx3FA    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
Lxxxx3FD TheUsrs  Lxxxx3FD    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
Lxxxx3F5 TheUsrs  Lxxxx3F5    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
Lxxxx3F8 TheUsrs  Lxxxx3F8    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
Lxxxx40A TheUsrs  Lxxxx40A    .         .     IFL  400  .   .    .   .   4  4 ESA   N   N  N   N   61G  61G
Lxxxx40D TheUsrs  Lxxxx40D    .         .     IFL  400  .   .    .   .   4  1 ESA   N   N  N   N   61G  61G
Lxxxx400 TheUsrs  Lxxxx400    .         .     IFL  400  .   .    .   .   4  4 ESA   N   N  N   N   61G  61G
Lxxxx404 TheUsrs  Lxxxx404    .         .     IFL  200  .   .    .   .   2  2 ESA   N   N  N   N   31G  31G
Lxxxx408 TheUsrs  Lxxxx408    .         .     IFL  400  .   .    .   .   4  4 ESA   N   N  N   N   61G  61G
Lxxxx41B TheUsrs  Lxxxx41B    .         .     IFL  800  .   .    .   .   8  8 ESA   N   N  N   N  122G 122G
Lxxxx411 TheUsrs  Lxxxx411    .         .     IFL  400  .   .    .   .   4  4 ESA   N   N  N   N   61G  61G
Lxxxx51E TheUsrs  Lxxxx51E    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
Lxxxx529 TheUsrs  Lxxxx529    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
Lxxxx538 TheUsrs  Lxxxx538    .         .     IFL 3200  .   .    .   .  32 32 ESA   N   N  N   N  488G 488G
```

A REL 3200 setting gives each of the 32 vCPUs only REL 100 (the z/VM default

ESASSUM – showed:

- Spikes in Processor Utilization
- A sudden change in the I/O subsystem
- Looking for spikes and sudden changes can show where and when problems start – and where to start investigations


```
Report: ESASSUM        Subsystem Activity
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08      First record anal
-----------------------------------------------------------------------------------------
         <---Users----> Transactions <Processor> Storage (MB) <-Paging--> <-----I/O-
         <-avg number->    Per  Avg. Utilization Fixed Active <pages/sec> <-DASD-->
Time     On Actv In Q Minute  Resp Total Virt.  User Resid. XStore DASD Rate Resp
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ---- ---- ----
04/25/22
13:15:00   78   45  191  139.3 0.019   662   613  74.9 2765K      0    1  165  0.2
13:30:00   78   45  190  142.1 0.019   625   580  74.9 2774K      0    1  168  0.2
13:45:00   78   45  196  138.9 0.020  1303  1248  74.9 2778K      0    0  163  0.2
14:00:00   78   46  178  140.6 0.020   848   811  74.9 2783K      0    0  168  0.2
14:15:00   78   45  190  140.3 0.016   312   287  74.9 2791K      0    0  164  0.2
14:30:00   78   45  195  139.8 0.020   337   308  74.9 2805K      0    0  165  0.2
14:45:00   78   45  188  140.8 0.020   943   891  74.9 2811K      0 2269  187  1.2
15:00:00   78   45  189  140.0 0.023   596   556  74.9 2812K      0  17K  470  2.9
15:15:00   78   46  179  141.7 0.023   374   341  74.9 2812K      0  20K  661  2.4
15:30:00   78   45  194  139.7 0.020   413   381  74.9 2813K      0 8889  489  1.3
15:45:00   78   47  193  141.8 0.023   417   386  74.9 2813K      0  16K  664  2.2
16:00:00   78   48  191  134.6 0.032   247   232  74.9 2813K      0 5049  379  1.0
16:15:00   78   47  169  122.3 0.025   216   202  74.9 2813K      0 5164  359  1.1
16:30:00   78   47  180  121.5 0.018   210   196  74.9 2813K      0 4470  331  1.0
```

ESAXACT – Transaction Delay Analysis showed:

- Lxxxx3F8 has a large number of Running samples
- Only 24% of 32 vCPUs are running – don't need that many

```
Report: ESAXACT        Transaction Delay Analysis
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08    First record analyz
--------------------------------------------------------------------------------------
                       <-----Percent non-dormant (Wait states)------------------->
UserID    <-Samples->                     E-  D-  T-     Tst <Asynch>       Lim  Pct
/Class     Total   In Q  Run Sim CPU SIO Pag SVM SVM SVM  CF Idl I/O Pag Ldg Oth Lst Elig
--------  -----  -----  --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- ----
04/25/22
14:45:00   1170   2815 7.1 0.1 2.1   0   0   0   0   0   0  90   0   .   .   1   0   0
Hi-Freq:   302K   174K 4.6 0.1 1.2 0.0 0.0   0 1.0 0.0   0  94   0 0.0 0.0 0.0   0   0
 ***User Class Analysis***
Servers   10800    459   0 0.2 0.4   0   0   0  10 8.9   0  90   0   0   0   0   0   0
Velocity   9000    115 1.7 0.9   0   0   0   0 0.2 8.7   0  89   0   0   0   0   0   0
CATech     4500      1   0   0   0   0   0   0   0   0   0   0   0   0   0   0   0   0
TheUsrs    278K   174K 4.6 0.1 1.2 0.0 0.0   0 0.7 0.0   0  94   0 0.0 0.0 0.0   0   0
 ***Accounting Code Analysis**
Lxxxx3B2  14400  12319 0.3   0 0.4   0   0   0   0   0   0  99   0   0 0.1   0   0   0
Lxxxx3FA  28800  15945 1.1 0.1 0.4   0   0   0   0   0   0  98   0   0   0   0   0   0
Lxxxx3FD  28800  27916 4.3 0.2 2.5   0 0.0   0   0   0   0  93   0   0   0 0.0   0   0
Lxxxx3F5  28800  16568 0.2 0.0 0.4   0   0   0   0   0   0  99   0   0 0.0 0.0   0   0
Lxxxx3F8  28800  25036  24 0.4 3.2   0   0   0   0   0   0  73   0   0   0 0.0   0   0
Lxxxx411   3600   3600 2.9 0.1 3.2   0   0   0   0   0   0  94   0 0.0 0.1   0   0   0
Lxxxx51E  28800  15990 0.2   0 0.3   0   0   0   0   0   0  99   0   0   0   0   0   0
Lxxxx529  28800  15122 0.2 0.0 0.2   0   0   0   0   0   0 100   0   0   0   0   0   0
Lxxxx538  28800  15500 1.9 0.0 0.3   0   0   0   0   0   0  98   0   0   0   0   0   0
```

ESALPAR – Logical Partition Analysis showed:

- Shows there are too many vCPUs assigned
  - There are 20 vCPUs on the LPAR – mostly Vertical High and Vertical Medium in use but still with idle threads

```
Report: ESALPAR     Logical Partition Analysis              tware Corporate    ZMAP 5.1.3 05/01/22   Pg    6452
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08      analyzed: 04/25/22 15:00:00
----------------------------------------------------------------  ------------------------------------------------
        CEC  <-Logical Partition-> <----------Logical Processor--- <--------CPU (percentages)-------> <Multi-thread>
        Phys                  Pool   VCPU <%Assigned> VCPU Weight/  Total  Emul  User   Sys  Idle  Stl Idle  cp1/cp2
Time    CPUs Name    No  Name   Addr Total   Ovhd TYPE Polar        util   time  ovrhd ovrhd time  Pct Time
-------- ---- -------- --- -------- ---- ----- ---- ---- --- ---   ----- ----- ----- ----- ----- ---- ------ --- ---
04/25/22
14:45:00  109 Vxxxxxx8  08          .    0  62.5   0.7 IFL   60 VHi   86.7  81.3   2.3   3.0 110.6 2.64 35.32   0 / 1
                                        1  63.7   0.7 IFL   60 VHi   89.2  84.3   2.2   2.6 108.4 2.48 35.46   2 / 3
                                        2  63.1   0.7 IFL   60 VHi   87.7  82.6   2.4   2.7 109.8 2.57 35.58   4 / 5
                                        3  65.4   0.7 IFL   60 VHi   91.0  86.2   2.2   2.6 106.6 2.40 37.09   6 / 7
                                        4  65.5   0.7 IFL   60 VHi   91.0  86.2   2.2   2.6 106.5 2.48 37.07   8 / 9
                                        5  65.8   0.7 IFL   60 VHi   91.5  86.6   2.3   2.7 105.9 2.57 37.24  10 /11
                                        6  65.0   0.7 IFL   60 VHi   90.2  85.2   2.3   2.7 107.1 2.65 36.76  12 /13
                                        7  64.9   0.7 IFL   60 VMe   90.1  85.2   2.3   2.7 107.2 2.66 36.70  14 /15
                                        8  64.0   0.7 IFL   60 VMe   89.5  84.5   2.4   2.7 107.8 2.65 35.55  16 /17
                                        9  36.7   0.4 IFL   60 VLo   52.4  49.9   1.1   1.4  47.7  100 19.46  18 /19
                                       10  20.7   0.3 IFL   60 VLo   29.6  27.5   1.1   1.0  44.6  126 10.75  20 /21
                                       11  13.9   0.1 IFL   60 VLo   20.6  19.5   0.6   0.6  25.2  154  6.70  22 /23
                                       12   9.6   0.1 IFL   60 VLo   14.6  14.0   0.3   0.3  14.5  171  4.38  24 /25
                                       13   5.5   0.0 IFL   60 VLo    8.4   8.1   0.1   0.2  10.1  181  2.47  26 /27
                                       14   3.4   0.0 IFL   60 VLo    5.2   5.1   0.1   0.1   5.4  189  1.41  28 /29
                                       15   3.0   0.0 IFL   60 VLo    4.9   4.7   0.1   0.1   4.9  190  1.17  30 /31
                                       16   0.3   0.0 IFL   60 VLo    0.4   0.4   0.0   0.0   2.2  197  0.19  32 /33
                                       17   0.0   0.0 IFL   60 VLo    0.0   0.0   0.0   0.0   0.4  200  0.00  34 /35
                                       18   0.0   0.0 IFL   60 VLo    0.0     0     0   0.0     0  200  0.00  36 /37
                                       19   0.0   0.0 IFL   60 VLo    0.0     0     0   0.0     0  200  0.00  38 /39
                                           ----- ----                ----- ----- ----- ----- ----- ---- ------ --- ---
                                 LPAR 673.2   7.3                   943.1 891.3  23.9  27.9  1125 1932 373.3   0 / 0
```

ESALPARS – Logical Partition Analysis Summary showed:

- There are too many vCPUs causing overhead
  - Thread Idle percentage is high and
  - Shared Processor busy Ovhd/Mgmt are also high – usually indicative of a thrashing situation

```
Report: ESALPARS      Logical Partition Summary              Velocity Software Corporate    ZMA
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08     First record analyzed: 04/25/22 1
---------------------------------------------------------------------------------------------------

           <--Complex--> <--------Logical Partition-------->  <-Assigned Shares---->
           Phys Dispatch            Virt CPU  <%Assigned>  <---LPAR--> <VCPU Pct>  Wait <-Thread->
Time       CPUs     Slice Name   Nbr CPUs Type Total  Ovhd  Weight  Pct /SYS /CPU  Comp  Idle  cnt
--------   ----  -------- ----   --- ---- ---- -----  ----  ------ ---- ---- ----  ---- ------ ---
04/25/22
15:30:00                 Vxxxxxx8  08   20 IFL  347.0  6.2     60  7.5 0.37 40.2   No  243.3   2
15:31:00                 Vxxxxxx8  08   20 IFL  341.6  6.5     60  7.5 0.37 40.2   No  233.0   2
15:32:00                 Vxxxxxx8  08   20 IFL  397.9  6.4     60  7.5 0.37 40.2   No  235.4   2
15:33:00                 Vxxxxxx8  08   20 IFL  351.7  6.2     60  7.5 0.37 40.2   No  216.5   2
15:34:00                 Vxxxxxx8  08   20 IFL  340.3  6.4     60  7.5 0.37 40.2   No  241.6   2
15:35:00                 Vxxxxxx8  08   20 IFL  372.3  6.3     60  7.5 0.37 40.2   No  263.3   2
15:36:00                 Vxxxxxx8  08   20 IFL  334.9  6.3     60  7.5 0.37 40.2   No  238.8   2
15:37:00                 Vxxxxxx8  08   20 IFL  356.3  6.3     60  7.5 0.37 40.2   No  243.6   2
15:38:00                 Vxxxxxx8  08   20 IFL  362.2  6.1     60  7.5 0.37 40.2   No  248.1   2
15:40:00                 Vxxxxxx8  08   20 IFL  374.3  5.9     60  7.5 0.37 40.2   No  241.1   2

Totals by Processor type:
<---------CPU-------> <-Shared Processor busy->
Type Count Ded shared   Total    Logical  Ovhd  Mgmt
---- ----- --- ------  ------  --------- ----- ----
IFL   108   0    108   1693.2    1642.8  25.1  25.2
IFL   108   0    108   1662.9    1605.7  28.6  28.6
IFL   108   0    108   1697.7    1641.5  28.3  27.8
IFL   108   0    108   1680.6    1626.9  27.1  26.6
IFL   108   0    108   1730.3    1675.0  27.9  27.4
```

VELOCITY SOFTWARE

PROVEN PERFORMANCE

ESACPUU – CPU Utilization Analysis showed:

- The vCPUs from the z/VM perspective (40 threads)
- Shows the parking of Vertical Low vCPUs (threads 18-39)

```
Report: ESACPUU        CPU Utilization Report              Ve  05/01/22    Pg    6150
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08           Fi
-----------------------------------------------------------------------------------------
        <----Load---->               <--------CPU (percentages)--------> MThread <-Vertical-->
        <-Users-> Tran        CPU   Total  Emul  User   Sys  Idle Steal Core/   Entitle- Park
Time    Actv In Q /sec CPU    Type  util   time  ovrhd ovrhd time  time Thread  ment  seconds
------- ---- ---- ---- -      ----  -----  ----- ----- ----- ----- ---- ------  ----  -----
04/25/22
14:45:00  45  188  2.3  0     IFL   42.8   39.9   1.2   1.8  55.9  1.3  0/ 0    0.79      0
                         1     IFL   43.9   41.5   1.2   1.3  54.7  1.3  0/ 1    0.79      0
                         2     IFL   45.1   42.7   1.1   1.3  53.7  1.2  1/ 0    0.79      0
                         3     IFL   44.0   41.6   1.1   1.3  54.7  1.3  1/ 1    0.79      0
                         4     IFL   44.3   41.8   1.2   1.4  54.4  1.3  2/ 0    0.79      0
                         5     IFL   43.3   40.8   1.2   1.3  55.4  1.3  2/ 1    0.79      0
                         6     IFL   46.2   43.8   1.1   1.3  52.6  1.2  3/ 0    0.79      0
                         7     IFL   44.8   42.4   1.1   1.3  54.0  1.2  3/ 1    0.79      0
                         8     IFL   46.0   43.5   1.1   1.3  52.8  1.2  4/ 0    0.79      0
                         9     IFL   45.1   42.6   1.2   1.3  53.7  1.3  4/ 1    0.79      0
                        10     IFL   46.3   43.8   1.1   1.4  52.4  1.3  5/ 0    0.79      0
                        11     IFL   45.2   42.8   1.1   1.3  53.5  1.3  5/ 1    0.79      0
                        12     IFL   45.6   43.1   1.1   1.4  53.1  1.3  6/ 0    0.79      0
                        13     IFL   44.6   42.1   1.2   1.3  54.1  1.3  6/ 1    0.79      0
                        14     IFL   45.8   43.3   1.1   1.4  52.9  1.3  7/ 0    0.79      0
                        15     IFL   44.3   41.9   1.2   1.3  54.3  1.3  7/ 1    0.79      0
                        16     IFL   45.1   42.6   1.2   1.4  53.6  1.3  8/ 0    0.79      0
                        17     IFL   44.4   41.9   1.2   1.3  54.2  1.3  8/ 1    0.79      0
                        18     IFL   26.7   25.4   0.6   0.7  23.4  49.9  9/ 0   0.36    443.8
                        19     IFL   25.8   24.5   0.6   0.7  24.3  49.9  9/ 1   0.36    443.8
                        20     IFL   14.9   13.8   0.5   0.6  22.1  62.9  0/ 0   0.25    561.9
                        21     IFL   14.7   13.6   0.5   0.5  22.4  62.9  0/ 1   0.25    561.9
                        22     IFL   10.3    9.7   0.3   0.3  12.6  77.1  1/ 0   0.15    691.9
                        23     IFL   10.3    9.7   0.3   0.3  12.6  77.1  1/ 1   0.15    691.9
                        24     IFL    7.3    7.0   0.1   0.2   7.2  85.4  2/ 0   0.09    767.9
                        25     IFL    7.3    7.0   0.1   0.2   7.3  85.4  2/ 1   0.09    767.9
                        26     IFL    4.2    4.1   0.1   0.1   5.0  90.7  3/ 0   0.06    816.0
                        27     IFL    4.2    4.1   0.1   0.1   5.1  90.7  3/ 1   0.06    816.0
                        28     IFL    2.6    2.5   0.0   0.0   2.7  94.7  4/ 0   0.03    852.0
                        29     IFL    2.6    2.5   0.0   0.0   2.7  94.7  4/ 1   0.03    852.0
                        30     IFL    2.4    2.4   0.0   0.0   2.4  95.1  5/ 0   0.03    856.0
                        31     IFL    2.4    2.3   0.0   0.0   2.4  95.1  5/ 1   0.03    856.0
                        32     IFL    0.2    0.2   0.0   0.0   1.1  98.7  6/ 0   0.01    888.0
                        33     IFL    0.2    0.2   0.0   0.0   1.1  98.7  6/ 1   0.01    888.0
                        34     IFL    0.0    0.0   0.0   0.0   0.2  99.8  7/ 0   0.00    898.0
                        35     IFL    0.0    0.0   0.0   0.0   0.2  99.8  7/ 1   0.00    898.0
                        36     IFL    0.0    0     0     0.0    0  100.0  8/ 0   0.00    900.0
                        37     IFL    0.0    0     0     0.0    0  100.0  8/ 1   0.00    900.0
                        38     IFL    0.0    0     0     0.0    0  100.0  9/ 0   0.00    900.0
                        39     IFL    0.0    0     0     0.0    0  100.0  9/ 1   0.00    900.0
                              ----- ----- ----- ----- ----- ----- ----  ----  -----
System:                       943.1 891.3  23.9  27.9  1125  1932  ./ .  0.40  17151
```

ESADSD2 – DASD Performance Analysis showed:

- High response and service times – indicates queueing
- PAV is turned off (all zeroes) – paging devices are single-threaded, response times will suffer

```
Report: ESADSD2      DASD Performance Analysis                 Velocity Software Corporate    ZMAP 5.1.3 05/01/22    Pg
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08      First record analyzed: 04/25/22 15:00:00
-----------------------------------------------------------------------------------------------------------------------
                                                     <--------DASD Response times (ms)------>             Interupts <-PAV/Hipe
 Dev        Device <--SSCH--> <%DevBusy> <SSCH/sec->      <--Service times--> <--Queueing-> QLengths </Second> <-Rate/Sec
 No. Serial Type    Total ERP Avg  Peak  avg   peak  Resp Serv Pend Disc Conn DASD Cntl THR avg max Sol Unsol Q-d  SSCH
---- ------ -----   ----- --- ---- ----  ----- ----- ---- ---- ---- ---- ---- ---- ---- --- --- --- --- ----- ---- ----
15:01:00
***Top DASD by Device busy***
5111 V5111B 3390-9   2031   0  8.7  8.7  33.8  33.8   7.0  2.6  0.1  0.7  1.7  4.4    0    0 0.1   2  34    0    0    0
5222 V52227 3390-9   2078   0  8.3  8.3  34.6  34.6   5.3  2.4  0.1  0.8  1.5  2.9    0    0 0.1   1  35    0    0    0
5333 V5333C 3390-9   2071   0  8.2  8.2  34.5  34.5   3.4  2.4  0.1  0.7  1.5  1.0    0    0 0.0   0  34    0    0    0
5444 V54449 3390-9   2166   0  8.2  8.2  36.1  36.1   7.3  2.3  0.1  0.6  1.5  5.1    0    0 0.2   2  36    0    0    0
5555 V55558 3390-9   2004   0  8.1  8.1  33.4  33.4   4.9  2.4  0.1  0.7  1.6  2.5    0    0 0.1   1  33    0    0    0
5666 V56666 3390-9   2029   0  7.9  7.9  33.8  33.8   3.8  2.3  0.1  0.7  1.5  1.5    0    0 0.0   1  34    0    0    0
5777 V5777A 3390-9   2028   0  7.7  7.7  33.8  33.8   2.3  2.3  0.1  0.7  1.5    0    0    0   0   0  34    0    0    0
5888 V58885 3390-9   2027   0  7.5  7.5  33.8  33.8   2.2  2.2  0.1  0.7  1.4    0    0    0   0   0  34    0    0    0
5999 V59994 3390-9   2050   0  7.4  7.4  34.2  34.2   6.6  2.2  0.1  0.6  1.4  4.4    0    0 0.1   2  34    0    0    0
5aaa V5aaa3 3390-9   2016   0  7.4  7.4  33.6  33.6   4.2  2.2  0.1  0.7  1.4  2.0    0    0 0.1   1  34    0    0    0
```

Velocity
SOFTWARE

PROVEN PERFORMANCE

ESAUSP2 – User Percent Utilization showed:

- Lxxxx3F8 was using a lot of CPU at the time of the issue

ESAUSPG – User Storage Analysis showed:

- Multiple servers have a huge amount of storage
- Lxxxx538 started holding storage below the 2G line
- Paging started thrashing (Megabytes Paged Out)

```
Report: ESAUSPG      User Storage Analysis
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08
------------------------------------------------------------------------------
         <Storage Occupancy (MegaBytes)> <-Page I/O--> Pages    Page
UserID   <---Main Storage--->     Paged Page    Page   Moved    Faults
/Class   Total   >2gb    <2GB     Out  Writes   Reads  <2GB
-------  ------  ------  ------   ------- ------ ------  ------  ------
04/25/22
14:30:00  2806K   2805K    234        0      0      0       0    1817
TheUsrs   2805K   2805K    181        0      0      0       0    1769
 ***Top User Analysis***
Lxxxx3FD 497610 497610      0        0      0      0       0      11
Lxxxx3F8 497609 497609      0        0      0      0       0       4
Lxxxx538 421606 421505    101        0      0      0       0      34
Lxxxx3FA 360765 360759      6        0      0      0       0      93
Lxxxx411  45343  45343      0        0      0      0       0     135
------------------------------------------------------------------------------
14:45:00  2811K   2809K   1698     1241  1548K 102635       0  233532
TheUsrs   2811K   2809K   1668     1147  1517K 101501       0  198997
 ***Top User Analysis***
Lxxxx3F8 497540 497540      0       69  94026    955       0     270
Lxxxx3FD 497558 497558      0       52  81951  22318       0   15886
Lxxxx538 427501 426021   1480       73  97357    183       0      43
Lxxxx3FA 361117 360972    144       73  97493     33       0     142
Lxxxxx411 45291  45291      0       50  89325  15935       0    3955
------------------------------------------------------------------------------
15:00:00  2812K   2810K   1914    20765 11217K  4105K       0   1278K
TheUsrs   2812K   2810K   1907    20588 11202K  4090K       0   1168K
 ***Top User Analysis***
Lxxxx3F8 496265 496265      0     1344 658447 183132       0    9162
Lxxxx3FD 497114 497114      0      496   1247K  1139K       0  174027
Lxxxx538 446916 445180   1736     1349 546979  87328       0   29226
Lxxxx3FA 360201 360033    168     1489 547280   5704       0    1662
Lxxxx411  44433  44433      0      903  1159K 883815       0   91732
------------------------------------------------------------------------------
15:15:00  2812K   2811K   1915    49630 12540K  5000K       0   1531K
TheUsrs   2812K   2811K   1909    49453 12514K  4979K       0   1409K
 ***Top User Analysis***
Lxxxx3FD 497092 497092      0      518  1100K  1071K       0  359847
Lxxxx538 475487 473748   1739     3126 617158 169221       0   40030
Lxxxx3F8 494369 494369      0     3239 801799 335159       0   12930
Lxxxx3FA 358567 358399    168     3592 751006 218475       0   12484
Lxxxx411  43078  43078      0     2252  1021K 636711       0   74060
------------------------------------------------------------------------------
```

ESALNXS – Linux VSI System Analysis showed:

- Lxxxx3F8 has 32 vCPUs – Linux uses all of them, even though only one major process was running (from ESALNXP)
- Spin locks result from too many vCPUs

```
Report: ESALNXS      LINUX VSI System Analysis Report
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08    First record
----------------------------------------------------------------------------

Node/     <---Load Numbers--> CPU <Processor Pct Util> NICE <CPU Overhead%> IO
Time      Users Procs MaxProc NBR Total Syst User Idle Time Krnl IRQ Steal Wait
--------  ----- ----- ------- --- ----- ---- ---- ---- ---- ---- ---- ----- ----
LINUX3F8      0   441       0 Tot 616.4 13.1  598 2541    0  0.7  4.2  42.5  0.3
                                1  20.0  1.2 16.2 77.5    0  0.1  2.5   2.5  0.0
                                2  24.9  0.6 24.2 73.5    0  0.0  0.1   1.6  0.0
                                3  24.6  0.5 24.0 73.7    0  0.0  0.1   1.7  0.0
                                4  27.2  0.5 26.5 71.0    0  0.0  0.1   1.8  0.0
                                5  27.9  0.5 27.3 70.3    0  0.0  0.1   1.8  0.0
                                6  22.6  0.5 22.0 75.9    0  0.0  0.1   1.5  0.0
                                7  22.1  0.5 21.5 76.4    0  0.0  0.1   1.5  0.0
                                8  23.5  0.5 22.9 74.9    0  0.0  0.1   1.6  0.0
                                9  22.5  0.5 21.9 76.0    0  0.0  0.1   1.6  0.0
                               10  20.9  0.4 20.4 77.6    0  0.0  0.1   1.5  0.0
                               11  21.6  0.4 21.1 76.9    0  0.0  0.1   1.5  0.0
                               12  18.6  0.4 18.2 80.1    0  0.0  0.1   1.3  0.0
                               13  17.3  0.4 16.8 81.4    0  0.0  0.1   1.2  0.0
                               14  18.5  0.4 18.1 80.1    0  0.0  0.1   1.3  0.0
                               15  19.1  0.4 18.6 79.6    0  0.0  0.1   1.3  0.0
                               16  19.7  0.4 19.3 78.9    0  0.0  0.1   1.4  0.0
                               17  17.6  0.4 17.1 81.2    0  0.0  0.1   1.3  0.0
                               18  16.8  0.4 16.3 82.0    0  0.0  0.0   1.3  0.0
                               19  18.7  0.4 18.2 80.0    0  0.0  0.1   1.3  0.0
                               20  20.3  0.4 19.9 78.3    0  0.0  0.1   1.4  0.0
                               21  16.8  0.4 16.4 82.1    0  0.0  0.1   1.1  0.0
                               22  16.3  0.4 15.8 82.7    0  0.0  0.1   1.0  0.0
                               23  15.9  0.3 15.6 83.1    0  0.0  0.0   1.0  0.0
                               24  19.4  0.3 19.0 79.3    0  0.0  0.1   1.3  0.0
                               25  15.6  0.2 15.3 83.4    0  0.0  0.0   1.0  0.0
                               26  18.8  0.3 18.4 79.9    0  0.0  0.1   1.2  0.0
                               27  18.3  0.4 17.8 80.5    0  0.0  0.1   1.2  0.0
                               28  15.1  0.3 14.8 83.9    0  0.0  0.0   1.0  0.0
                               29  15.6  0.3 15.2 83.3    0  0.0  0.1   1.0  0.0
                               30  12.5  0.3 12.1 86.7    0  0.0  0.0   0.8  0.0
                               31  13.2  0.3 12.9 85.9    0  0.0  0.0   0.9  0.0
```

VELOCITY SOFTWARE

PROVEN PERFORMANCE

ESALNXP – Linux HOST Process Statistics showed:

- Lxxxx538 had a MongoDB process that was ramping up (CPU/Storage)
- This ended up needing more storage than was available

```
Report: ESALNXP       LINUX HOST Process Statistics Report        ZMAP 5.1.3
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08    First rec
--------------------------------------------------------------------------------
node/     <Process Ident> Nice PRTY <------CPU Percents----> <Storage (MB)
 Name          ID    PPID Valu Valu  Tot  sys user syst usrt  Size  RSS Peak
--------- ------- ------- ---- ---- ---- ---- ---- ---- ---- ---- ---- ----
Lxxxx538        0       0    0    0 2.82 0.57 1.56 0.08 0.61  13K  487 211K
Lxxxx538        0       0    0    0 12.4 0.67 7.85 0.40 3.44  14K  533 219K
Lxxxx538        0       0    0    0 61.8 5.48 52.7 0.16 3.45  46K  31K 698K
Lxxxx538        0       0    0    0 21.8 2.24 19.5 0.01 0.02  94K  80K 1.4M
Lxxxx538        0       0    0    0 53.4 5.06 47.6 0.08 0.64 126K 112K 1.9M
Lxxxx538        0       0    0    0 17.2 1.63 15.5 0.01 0.02 149K 134K 2.2M
Lxxxx538        0       0    0    0 54.7 5.39 49.3 0.01 0.00 173K 159K 2.6M
Lxxxx538        0       0    0    0 62.1 4.18 57.2 0.07 0.58 204K 189K 3.1M
Lxxxx538        0       0    0    0 69.5 4.50 65.0 0.00 0.00 207K 193K 3.1M
Lxxxx538        0       0    0    0 11.3 1.18 10.2 0.01 0.00 208K 193K 3.1M
Lxxxx538        0       0    0    0 41.8 2.59 38.5 0.08 0.61 208K 193K 3.1M
Lxxxx538        0       0    0    0 20.6 1.67 19.0 0.01 0.00 208K 193K 3.1M
14:45:00
Lxxxx538        0       0    0    0 23.6 1.76 21.2 0.08 0.55 208K 193K 3.1M
 mongod-g   45390       1    0   20 20.6 1.16 19.4    0    0 195K 193K 2.9M
Lxxxx538        0       0    0    0 72.2 4.80 67.4 0.01 0.00 208K 193K 3.1M
 mongod-g   45390       1    0   20 69.5 4.00 65.5    0    0 195K 193K 2.9M
```

VELOCITY
SOFTWARE

PROVEN PERFORMANCE

ESAUCD2 – Linux UCD Memory Analysis showed:

- Way too much real storage is allocated but not being used
- Very little swap space is being used

ESAUSTR – User Storage Analysis showed:

- The Made IBR hit the 2% line (agelist default is 2%)
- A few minutes later, the Made IBR increased exponentially
- (The correct users aren't releasing pages)

```
Report: ESAUSTR      User Storage Analysis
Monitor initialized: 04/25/22 at 00:00:00 on
-----------------------------------------------------------------------
         <--------Virtual Server Storage (MB)----------->  <Resident> Mbyte <--
UserID   Size  Alloc Resi- UFO  <------IBR---->  <AgeList> <Unreferd> slots Made
/Class               dent Activ TOT  <2gb >2gb <2gb >2gb <2gb >2gb used  IBR
-------- ----- ----- ----- ----- ---- ---- ---- ---- ---- ---- ---- ----- ----
04/25/22
TheUsrs  4231K 2745K 2745K 2745K  3.5  0.5  3.0 21.4  413    0    0     0 33.3
TheUsrs  4231K 2681K 2681K 2681K  0.9  0.1  0.8  5.4  103    0    0     0 31.3  <- 2% line
TheUsrs  4231K 2765K 2765K 2763K 10.9  1.5  9.4 64.3 1533    0    0     0  0.8
TheUsrs  4231K 2774K 2774K 2772K 10.9  1.5  9.4 64.3 1532    0    0     0  1.8
TheUsrs  4231K 2778K 2778K 2776K 10.6  1.5  9.1 64.3 1540    0    0     0    0
TheUsrs  4231K 2783K 2783K 2781K 10.6  1.5  9.1 64.3 1536    0    0     0  0.0
TheUsrs  4231K 2775K 2775K 2773K 10.8  1.5  9.3 64.3 1535    0    0     0  2.4
TheUsrs  4231K 2791K 2791K 2789K 10.6  1.5  9.1 64.3 1532    0    0     0    0
TheUsrs  4231K 2805K 2805K 2804K 10.6  1.5  9.1 64.3 1530    0    0     0    0
TheUsrs  4231K 2812K 2811K 2808K 12.9  1.5 11.4 31.5 2023  0.0 37.0  1147  137  <- Fall off cliff
TheUsrs  4231K 2833K 2812K 2810K 36.9  1.5 35.4  0.1 2319  0.0  737 20588  632
TheUsrs  4231K 2810K 2805K 2803K 17.8  1.5 16.3 40.0 1851  0.0  193  5434  720
TheUsrs  4231K 2862K 2812K 2810K 45.1  1.5 43.6  0.2 2319  0.2  744 49453  687
```

VELOCITY
S O F T W A R E

P R O V E N   P E R F O R M A N C E

ESAPSDV – Page and Spool Device showed:

- There are an adequate amount of paging devices, but they are on differently-sized devices – which can cause issues
- The average SSCH/RSCH queuing was very high
- Again, this shows that PAV/HPAV is off (all zeros)

```
Report: ESAPSDV        Page And Spool Device Activity          Velocity Software Corporate    ZMAP 5.1.
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08    First record analyzed: 04/25/22 15:00:0
-----------------------------------------------------------------------------------------------------------
         <----Paging/Spooling---------->        </Sec><Device->        <--Queued pgs--> HPAV <--Record-->
Dev      <------Slots-------> <per sec>          SSCH Serv Resp %Alloc        <-Block--> I/O <Locate/Sec>
No. Serial Avail Used %Use  Max Read Writ Queue +RSCH Time Time Select Singl Read Write /Sec Read  Write
---- ----- ----- ---- ----  --- ---- ---- ----- ----- ---- ---- ------ ----- ---- ----- ---- ----- -----
Page Devices
5111 V5111A  12M 744K    6 744K  505 1267     0  33.4  0.0  0.0  100.0     0    0     0    0   36    11
5222 V5222B  12M 743K    6 743K  531 1274     0  33.3  0.0  0.0  100.0     0    0     0    0   37    10
5333 V5333C  12M 744K    6 744K  560 1290     0  33.8  0.1  0.1  100.0     0    0     0    0   37    11
5444 V54440 5897K 739K  13 739K  533 1285     0  33.4  0.0  0.0  100.0     0    0     0    0   37    10
5555 V55551 5897K 743K  13 743K  498 1286     0  33.6  0.0  0.0  100.0     0    0     0    0   36    10
5666 V56662 5897K 745K  13 745K  517 1268     0  33.9  0.0  0.0  100.0     0    0     0    0   37    10
5777 V57773 5897K 744K  13 744K  512 1286     0  33.3  0.1  0.1  100.0     0    0     0    0   37    10
5888 V58884  12M 744K    6 744K  526 1282     0  33.8  0.1  0.1  100.0     0    0     0    0   37    11
5999 V59995  12M 746K    6 746K  515 1279     0  33.6  0.0  0.0  100.0     0    0     0    0   37    10
5aaa V5aaa6  12M 745K    6 745K  523 1286     0  33.4  0.0  0.0  100.0     0    0     0    0   37    11
5bbb V5bbb7  12M 742K    6 742K  528 1287     0  34.5  0.1  0.1  100.0     0    0     0    0   38    10
5ccc V5ccc8  12M 744K    6 744K  509 1288     0  32.8  0.0  0.0  100.0     0    0     0    0   36    10
5ddd V5ddd9  12M 741K    6 741K  552 1249     0  35.8  0.1  0.1  100.0     0    0     0    0   39    11
          ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
Total Page 130M  10M    7  10M 6808  17K
```

ESABLKP – Block Paging Analysis showed:

- Service times were climbing

- Block reads and size were climbing (optimal size is 10 pages)

- Block Steal and Unreferenced pages climbing

- Single User page reads climbing

- All show stress on the storage system

```
Report: ESABLKP      Block Paging Analysis
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 087A08    First record analyzed: 04/25/22
-------------------------------------------------------------------------------------------
          <----Load---->  Serv <-Block-> <-Blocks Formed By-> Block <--Block Exceptions/sec-->
          <-Users-> Tran  Time <-Reads-> <-Steal-> <Migrate>  Fault <Single Read>  <No Refers>
Time      Actv In Q /sec  (ms) /sec Size /sec Size /sec Size   /sec  User System   Migr Steal
--------  ---- ---- ----  ---- ---- ---- ---- ---- ---- ----  -----  ------ ------  ----- -----
14:00:00    46  178  2.3   .      0    0    0    0    0    0      0       0    0.0      0     0
14:15:00    45  190  2.3   .      0    0    0    0    0    0      0       0    0.0      0     0
14:30:00    45  195  2.3   .      0    0    0    0    0    0      0       0    0.0      0     0
14:45:00    45  188  2.3  0.0   4.4 18.7 64.5 26.9    0    0    4.4     0.1    0.2      0  37.9
15:00:00    45  189  2.3  1.2 123.8 18.9 391.3 27.8   0    0  123.8    73.3    7.7      0  3668
15:15:00    46  179  2.4  3.0 148.2 18.7 434.4 27.5   0    0  148.2   207.3   18.5      0  4582
15:30:00    45  194  2.3  2.6  99.3 18.3 111.7 30.0   0    0   99.3   173.4   11.2      0  3236
15:45:00    47  193  2.4  3.3 128.7 17.6 329.3 28.7   0    0  128.7   185.4   41.2      0  3692
16:00:00    48  191  2.2  4.1  59.9 19.1 53.6 31.3    0    0   59.9    78.3    7.5      0  1900
```

ESAOPER – Operator/System Console showed:

- Parking was very active at the time of the issue
- Parking thrashing causes cache issues
- Too many vCPUs and incorrect weighting causes parking thrashing

```
Report: ESAOPER     Operator/System Log
Monitor initialized: 04/24/22 at 23:45:04 on 8561 serial 087A08    First
-----------------------------------------------------------------------
14:40:01 CPU Unpark from  26 to  30 CPUUtil= "12.1", Projected= "13.0"
14:40:03 CPU Park from 30 to  28 CPUUtil= "11.6", Projected= "12.1"
14:40:05 CPU Park from 28 to  26 CPUUtil= "11.1", Projected= "12.0"
14:40:09 CPU Unpark from  26 to  28 CPUUtil= "11.7", Projected= "12.0"
14:40:11 CPU Park from 28 to  26 CPUUtil= "11.0", Projected= "12.0"
14:40:13 CPU Unpark from  26 to  28 CPUUtil= "11.7", Projected= "12.4"
14:40:17 CPU Park from 28 to  26 CPUUtil= "11.0", Projected= "12.0"
14:40:19 CPU Unpark from  26 to  28 CPUUtil= "10.0", Projected= "12.2"
14:40:33 CPU Unpark from  28 to  32 CPUUtil= "13.2", Projected= "14.4"
14:40:35 CPU Unpark from  32 to  34 CPUUtil= "14.2", Projected= "15.1"
14:40:37 CPU Park from 34 to  32 CPUUtil= "14.3", Projected= "14.5"
14:41:05 CPU Unpark from  32 to  34 CPUUtil= "7.78", Projected= "15.5"
14:41:09 CPU Unpark from  34 to  36 CPUUtil= "6.78", Projected= "16.1"
14:41:11 CPU Park from 36 to  34 CPUUtil= "5.50", Projected= "15.9"
14:41:15 CPU Park from 34 to  32 CPUUtil= "5.38", Projected= "14.2"
14:41:17 CPU Park from 32 to  30 CPUUtil= "5.20", Projected= "13.0"
14:41:19 CPU Park from 30 to  26 CPUUtil= "5.33", Projected= "11.3"
14:41:21 CPU Park from 26 to  20 CPUUtil= "4.63", Projected= "8.79"
14:41:23 CPU Park from 20 to  18 CPUUtil= "5.54", Projected= "7.50"
14:41:35 CPU Unpark from  18 to  22 CPUUtil= "7.62", Projected= "9.16"
14:41:37 CPU Park from 22 to  20 CPUUtil= "8.27", Projected= "8.93"
14:42:05 CPU Park from 20 to  18 CPUUtil= "6.42", Projected= "7.89"
14:42:59 CPU Unpark from  18 to  20 CPUUtil= "6.73", Projected= "8.46"
14:43:03 CPU Park from 20 to  18 CPUUtil= "5.64", Projected= "7.06"
14:43:07 CPU Unpark from  18 to  20 CPUUtil= "7.31", Projected= "8.41"
14:43:09 CPU Park from 20 to  18 CPUUtil= "6.82", Projected= "7.82"
14:43:13 CPU Unpark from  18 to  22 CPUUtil= "7.52", Projected= "9.07"
```

VELOCITY SOFTWARE

PROVEN PERFORMANCE

What was the actual problem?

- Lxxxx538 started a process that was ramping up its Mongo database
- Lxxxx538 was holding a lock but got paged out
- The other large systems were spinning waiting on that lock (Lxxxx3F8 was a victim, not the culprit)
- There wasn't enough storage for that system to get paged back in
- Once everything starts backing up, the problems grow exponentially
- Eventually that lock was released and things recovered – but it had the potential to happen again

Performance Enhancement Suggestions:

1 – The "T-Shirt" size approach that is often used when moving Linux servers from xSeries boxes to the z/VM platform causes them to have excessive resources

- Several large Linux servers had more vCPUs than were needed
  - Verify only the necessary number of vCPUs are allocated
- Several large Linux servers had more storage than was needed
  - Verify only the necessary storage is allocated
- Use swap space
  - Use swap space to allow servers to use very fast and efficient virtual disk when extra storage is needed

Performance Enhancement Suggestions:

2 – Lower the number of vCPUs of the Linux servers

- This will help reduce processor parking and cache issues
- Each large server had 32 vCPUs when 16 would suffice

- Verify Parking settings
  - If needed, set unparking to large – **SET SRM UNPARKING LARGE** – this leaves more cores unparked which helps with processor cache issues
  - If needed, set excessive use to high – **SET SRM EXCESSUSE HIGH** – this allows the system to use more unentitled CPU capacity

Performance Enhancement Suggestions:

3 – Verify the SRM agelist settings are correct for the environment

- The replacement for xstore is the agelist
- Instead of up to 20% xstore available, now the default is 2% pageable storage – this works better for smaller servers
  - If needed, set the SRM agelist size to 5% - **SET SRM AGELIST SIZE 5.0%** - to give a bigger buffer area
  - If needed, set the early writes to yes – **SET SRM AGELIST EARLYW YES** – to allow unused pages to be written out early
  - If needed, set the keep slot to yes – **SET SRM AGELIST KEEPS YES** – to keep storage addresses longer

What the customer reported:

- The suggestions were implemented and no further issues have been reported

VELOCITY
SOFTWARE

PROVEN PERFORMANCE