

# BRUDEN Information N' General Observations (BINGO)

**Guy Peleg**

Senior Member of the Technical Staff

Director of EMEA Operations

*[guy.peleg@bruden.com](mailto:guy.peleg@bruden.com)*

Source: OpenVMS Information Desk – October 2004

# The Golden Rules

**The best performing code is  
the code not being executed**


**The fastest I/Os are those avoided**

**Idle CPUs are the fastest CPUs**

**Look at your code....be ready to be surprised**

# Performance Results Or No Expense

- Application did not scale beyond one processor
- Redesigned performance sensitive areas
- 460% Performance improvement



Transactions Per Seconds (more is better)

# RMS

- Use larger buffers & more of 'em
- FAB/RAB parameters:
  - ASY, RAH, WBH, DFW, SQO
  - ALQ & DEQ
  - MBC & MBF
  - NOSHR, NQL, NLK
- SET RMS ...
  - /SYSTEM
  - /BUFFER\_COUNT=n
  - /BLOCK\_COUNT=n

# RMS

- `SYSGEN> SET RMS_SEQFILE_WBH 1`
- `SET FILE /STATISTICS`
  - MONITOR RMS
- After Image Journaling for data protection
  - RMSJNLSNAP freeware tool

# FTP Performance & Simple RMS Tuning

FTP into IT13 and transfer the file

```
Brutel> ftp it13
```

```
220 IT13.bruclass.com FTP Server (Version 5.6) Ready.
```

```
Connected to ALPH13.BRUCLASS.COM.
```

```
Name (ALPH13.BRUCLASS.COM:bru_guy): peleg
```

```
331 Username peleg requires a Password
```

```
Password:
```

```
230 User logged in.
```

```
FTP> cd $1$dga703:[000000]
```

```
250-CWD command successful.
```

```
250 New default directory is $1$DGA703:[000000]
```

```
FTP> put HP-I64VMS-JAVA150-V0105-1-1.PCSI_SFX_I64EXE
```

```
200 TYPE set to IMAGE.
```

```
200 PORT command successful.
```

```
150 Opening data connection for $1$DGA703:[000000]HP-I64VMS-JAVA150-V0105-1-1.PC  
SI_SFX_I64EXE; (192.168.1.7,49428)
```

```
226 Transfer complete.
```

```
local: SYS$SYSDEVICE:[BRU_GUY]HP-I64VMS-JAVA150-V0105-1-1.PCSI_SFX_I64EXE;1 rem  
ote: HP-I64VMS-JAVA150-V0105-1-1.PCSI_SFX_I64EXE
```

```
286026004 bytes sent in 00:00:49.92 seconds (5594.83 Kbytes/s)
```

```
200 TYPE set to ASCII.
```

# FTP Performance & Simple RMS Tuning

```
$ disemb $1$dga703
$ moun/over=id $1$dga703
%MOUNT-I-MOUNTED, FAX mounted on _$1$DGA703: (IT13)
$ set vol $1$dga703/nohigh
$ disemb $1$dga703
$ moun/sys $1$dga703 fax
%MOUNT-I-MOUNTED, FAX mounted on _$1$DGA703: (IT13)
$ set rms/sys/exte=60000/seq/block=127/buf=8
$ mc sysgen
SYSGEN> SET RMS_SEQ 1
SYSGEN> W A
SYSGEN> Exit
```

## Throughput increased by more than 50%

```
FTP> put HP-I64VMS-JAVA150-V0105-1-1.PCSI_SFX_I64EXE
200 TYPE set to IMAGE.
200 PORT command successful.
150 Opening data connection for $1$DGA703:[000000]HP-I64VMS-JAVA150-V0105-1-1.PC
SI_SFX_I64EXE; (192.168.1.7,49432)
226 Transfer complete.
local: SYS$SYSDEVICE:[BRU_GUY]HP-I64VMS-JAVA150-V0105-1-1.PCSI_SFX_I64EXE;1 rem
ote: HP-I64VMS-JAVA150-V0105-1-1.PCSI_SFX_I64EXE
286026004 bytes sent in 00:00:31.83 seconds (8773.78 Kbytes/s)
200 TYPE set to ASCII.
```

## Larger Multi-Block Count is Better ?!?!?

- Remember when Guy said that a larger MBC was better?



Larger MBC ?!?!?

*He lied!!!!!!!!!!!!!!!!!!!!!!*

# Smaller MBC for Random Access

- Times to read 1,000,000 records randomly (same sequence of records (where mbc passed as first parameter:

```
$ frand 1
Elapsed time == 31233ms
$ frand 2
Elapsed time == 31680ms
$ frand 4
Elapsed time == 32607ms
$ frand 8
Elapsed time == 33698ms
$ frand 16
Elapsed time == 36101ms
```

```
$ frand 32
Elapsed time == 42823ms
$ frand 64
Elapsed time == 54761ms
$ frand 96
Elapsed time == 66343ms
$ frand 124
Elapsed time == 80122ms
$
$ frand 1
Elapsed time == 31205ms
$
```

- When mostly sequential with one or two files randomly read/written specify MBC on open.

# Alignment Bytes

- Attempt to minimize alignment bytes between fields
- Compilers allocate record/structure fields in the same order they appear in the source

CHAR

INTEGER VS.

CHAR

INTEGER

INTEGER

INTEGER

CHAR


CHAR

# Shadowed RAM disk

- Shadowed RAM disk for applications that frequently read data from disk.
- The Shadow server will read from memory and (hopefully) will write to both devices.
- Forces data to remain resident in memory
- Significantly boosts performance when files are opened cluster wide by multiple users
  - XFC will not help
  - Beneficial if file update rate is low compared to the read rate

# Physical Disk Vs. RAM disk

- C application that processes records read from sequential file
- Each I/O 124 Blocks
- RX2600, OpenVMS V8.3, HSG80



Elapsed time to read 250MB file (less is better)

# Compression support in BACKUP

- ZLIB routines are shipping in a new shareable image
  - SYS\$LIBRARY:COMPRESS\$SHR.EXE
  - SYS\$LIBRARY:COMPRESS\$SHR\_EV6.EXE (alpha only)
- New qualifier - /DATA\_FORMAT=COMPRESSED
  - Instructs BACKUP to create a compressed save-set
  - Does not need to be specified on the restore command
- Initially viewed as a feature to save space....but turned out to be a significant performance improvement

```
IPL31> dir cxx*.exe/siz=all
```

```
Directory SYS$COMMON:[SYSEXE]
```

```
CXX$COMPILER.EXE;1    39.68MB/39.68MB  
CXX$DEMANGLE.EXE;1    41KB/42KB  
CXX$LINK.EXE;1       166KB/166KB
```

```
Total of 3 files, 39.89MB/39.89MB
```

```
IPL31> backup cxx*.exe cxx.bck/sav/log
```

```
%BACKUP-W-NOFILES, no files selected from DSA5:[SYS0.][SYSEXE]CXX*.EXE;*  
%BACKUP-S-COPIED, copied DSA5:[SYS0.SYSCOMMON.][SYSEXE]CXX$COMPILER.EXE;1  
%BACKUP-S-COPIED, copied DSA5:[SYS0.SYSCOMMON.][SYSEXE]CXX$DEMANGLE.EXE;1  
%BACKUP-S-COPIED, copied DSA5:[SYS0.SYSCOMMON.][SYSEXE]CXX$LINK.EXE;1
```

```
IPL31> dir cxx.bck/siz
```

```
Directory SYS$SYSROOT:[SYSEXE]
```

```
CXX.BCK;1            44.60MB
```

```
Total of 1 file, 44.60MB
```

```
IPL31> dir cxx*.exe/siz
```

```
Directory SYS$COMMON:[SYSEXE]
```

```
CXX$COMPILER.EXE;1    39.68MB  
CXX$DEMANGLE.EXE;1     41KB  
CXX$LINK.EXE;1        166KB
```

```
Total of 3 files, 39.89MB
```

```
IPL31> backup cxx*.exe cxx.cbck/sav/data=comp/log
```

```
%BACKUP-W-NOFILES, no files selected from DSA5:[SYS0.][SYSEXE]CXX*.EXE;*  
%BACKUP-S-COPIED, copied DSA5:[SYS0.SYSCOMMON.][SYSEXE]CXX$COMPILER.EXE;1  
%BACKUP-S-COPIED, copied DSA5:[SYS0.SYSCOMMON.][SYSEXE]CXX$DEMANGLE.EXE;1  
%BACKUP-S-COPIED, copied DSA5:[SYS0.SYSCOMMON.][SYSEXE]CXX$LINK.EXE;1  
%BACKUP-I-COMPRESS, data compressed by 66%
```

```
IPL31> dir cxx.cbck/siz
```

```
Directory SYS$SYSROOT:[SYSEXE]
```


```
CXX.CBCK;1            14.82MB
```

**Compression support in**



# BACKUP/FAST

- Fast file scan
- Large number of input files
- Small number of files meets selection criteria



Elapsed time to backup all files created today (less is better)

## XX

- Not XXX
- SDA extension for exercising LAN
- SDA> XX HELP

# PEDRIVER Data Compression

- OpenVMS V8.3
- Reduces traffic between nodes
- May be beneficial for Shadow copy and MSCP traffic
- Can be enabled system wide or per VC

# Turn on compression for one VC

```
SCACP> set vc it14/comp
```

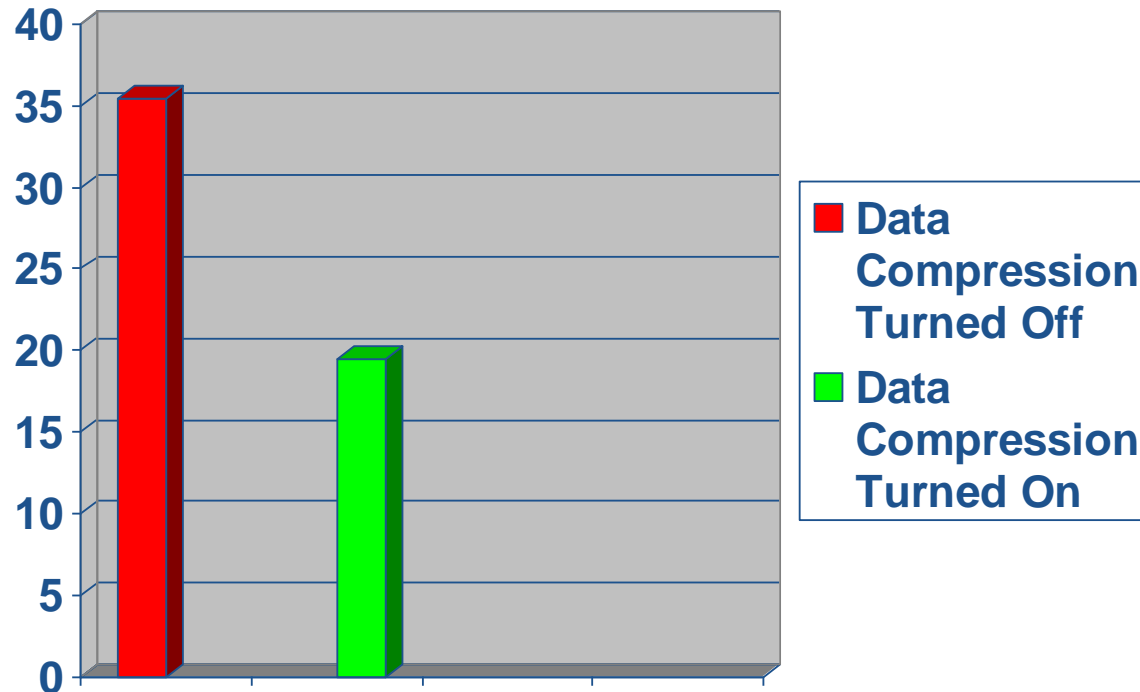
```
SCACP> sh vc
```

```
IT13 PEA0 VC Summary 30-JAN-2007 07:43:28.02:
```

| Remote | VC    | Total  | Channels | ECS  | MaxPkt | ReXmt | --XmtWindow-- | Xmt       | Total | ----- | Most Recent | -----   |           |                    |                |
|--------|-------|--------|----------|------|--------|-------|---------------|-----------|-------|-------|-------------|---------|-----------|--------------------|----------------|
| Node   | State | Errors | Xmt:TMO  | Open | ECS    | Pri   | Size          | TMO(uSec) | Cur   | Max   | Mgt         | Options | Pkts(S+R) | VC Opened Time     | VC Closed Time |
| ALPH50 | Open  | 4      | 115444   | 2    | 2      | 0     | 1426          | 672330.3  | 33    | 64    | 0           |         | 889107    | 21-JAN 13:34:25.78 | (No time)      |
| ALPH40 | Open  | 0      | Infinite | 2    | 2      | 0     | 1426          | 516452.3  | 16    | 32    | 0           |         | 803545    | 21-JAN 13:34:25.72 | (No time)      |
| IT14   | Open  | 1      | 790292   | 2    | 2      | 0     | 1426          | 223273.5  | 32    | 64    | 0           | CMP     | 1242954   | 21-JAN 13:34:25.93 | (No time)      |
| IT13   | Open  | 0      | Infinite | 1    | 1      | 0     | 1426          | 3000000.0 | 1     | 8     | 0           |         | 5         | 21-JAN 13:34:23.05 | (No time)      |

# PEDRIVER Data Compression

- Copy 250MB file to MSCP served SCSI disk
- Both systems are rx2600, running OpenVMS V8.3



Elapsed time to copy 250MB file (less is better)

# Free Hot File Tracking Utility

```
$ sh mem/cache=(volume=*,topqio)
```

```
System Memory Resources on 26-APR-2007 01:39:15.03
```

```
Extended File Cache Top QIO File Statistics:
```

```
_$1$DGA642: (DISK$ES40), Caching mode is VIOC Compatible
```

```
_$1$DGA642:[VMS$COMMON.SYSEXE]RIGHTSLIST.DAT;1 (open)
```

```
Caching is enabled, active caching mode is Write Through
```

|                 |    |               |      |
|-----------------|----|---------------|------|
| Allocated pages | 9  | Total QIOs    | 107  |
| Read hits       | 92 | Virtual reads | 107  |
| Virtual writes  | 0  | Hit rate      | 85 % |
| Read aheads     | 0  | Read throughs | 107  |
| Write throughs  | 0  | Read arounds  | 0    |
|                 |    | Write arounds | 0    |

```
_$1$DGA642:[VMS$COMMON.SYSEXE]VMS$OBJECTS.DAT;2 (open)
```

```
Caching is enabled, active caching mode is Write Through
```

|                 |   |               |     |
|-----------------|---|---------------|-----|
| Allocated pages | 0 | Total QIOs    | 9   |
| Read hits       | 0 | Virtual reads | 9   |
| Virtual writes  | 0 | Hit rate      | 0 % |
| Read aheads     | 0 | Read throughs | 9   |
| Write throughs  | 0 | Read arounds  | 0   |
|                 |   | Write arounds | 0   |

```
_$1$DGA642:[VMS$COMMON.SYSEXE]VMS$AUDIT_SERVER.DAT;1 (open)
```

```
Caching is enabled, active caching mode is Write Through
```

|                 |   |               |     |
|-----------------|---|---------------|-----|
| Allocated pages | 1 | Total QIOs    | 4   |
| Read hits       | 0 | Virtual reads | 4   |
| Virtual writes  | 0 | Hit rate      | 0 % |
| Read aheads     | 0 | Read throughs | 4   |
| Write throughs  | 0 | Read arounds  | 0   |

# Free Hot File Tracking Utility

\_\$1\$DGA242: (DISK\$ITANIUMVMS), Caching mode is VIOC Compatible

\_\$1\$DGA242: [VMS\$COMMON.SYSLIB]DECC\$SHR.EXE;1 (open)

Caching is enabled, active caching mode is Write Through

|                 |      |               |      |
|-----------------|------|---------------|------|
| Allocated pages | 303  | Total QIOs    | 1646 |
| Read hits       | 1561 | Virtual reads | 1646 |
| Virtual writes  | 0    | Hit rate      | 94 % |
| Read aheads     | 0    | Read throughs | 1642 |
| Write throughs  | 0    | Read arounds  | 4    |
|                 |      | Write arounds | 0    |

\_\$1\$DGA242: [VMS\$COMMON.SYSLIB]LIBRTL.EXE;1 (open)

Caching is enabled, active caching mode is Write Through

|                 |      |               |      |
|-----------------|------|---------------|------|
| Allocated pages | 143  | Total QIOs    | 1165 |
| Read hits       | 1123 | Virtual reads | 1165 |
| Virtual writes  | 0    | Hit rate      | 96 % |
| Read aheads     | 0    | Read throughs | 1164 |
| Write throughs  | 0    | Read arounds  | 1    |
|                 |      | Write arounds | 0    |

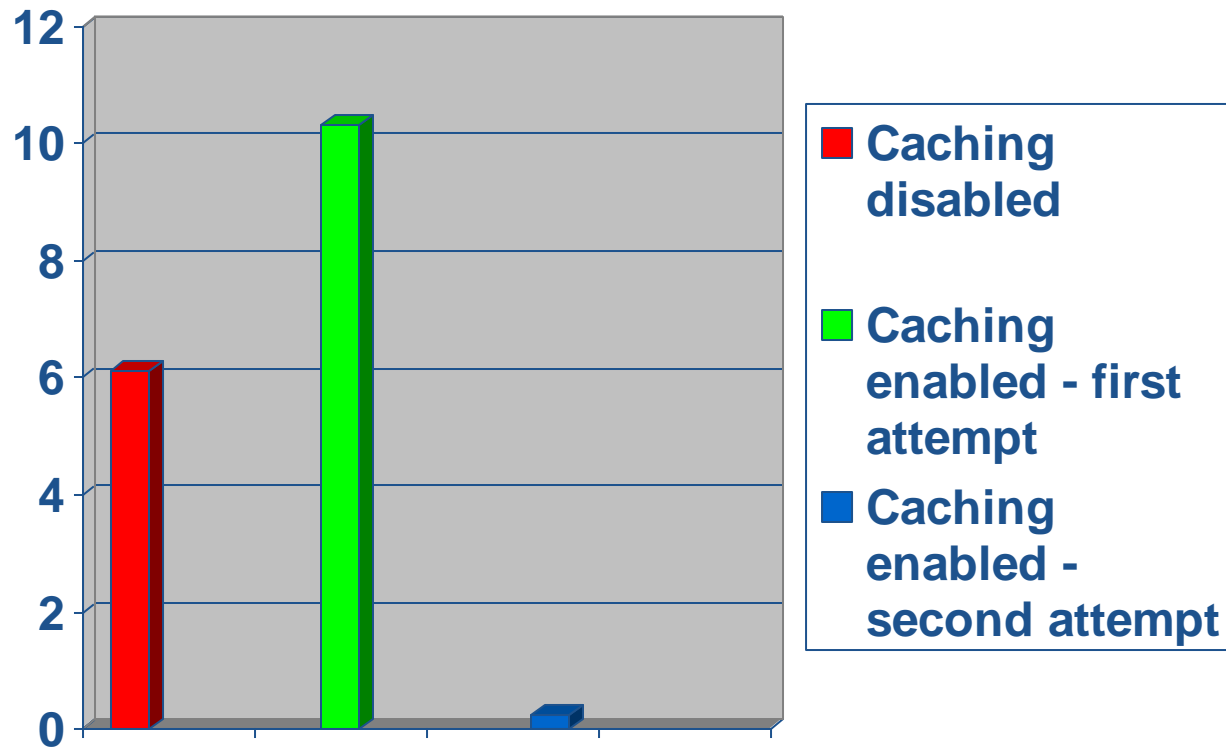
\_\$1\$DGA242: [VMS\$COMMON.SYSLIB]CMA\$TIS\_SHR.EXE;1 (open)

Caching is enabled, active caching mode is Write Through

|                 |     |               |      |
|-----------------|-----|---------------|------|
| Allocated pages | 12  | Total QIOs    | 720  |
| Read hits       | 711 | Virtual reads | 720  |
| Virtual writes  | 0   | Hit rate      | 98 % |
| Read aheads     | 0   | Read throughs | 720  |
| Write throughs  | 0   | Read arounds  | 0    |
|                 |     | Write arounds | 0    |

**Avoid caching files that  
pollute the cache**

# XFC “overhead”



**RDB users –  
consider disabling  
caching of .RDA  
files**

**Elapsed time to copy 150MB file, rx2600, HSG80, OpenVMS V8.3**



# Elapsed time for I/Os

Cache Volume Block (CVB)

-----  
Statistics Valid From: 25-APR-2007 07:43:01.07

Name: DISK\$ES40  
CVB Address: FFFFFFFF20188540  
Flink: FFFFFFFF20188300  
Blink: FFFFFFFF201887C0  
Volume (VCB): FFFFFFFF86536980  
Unit (UCB): FFFFFFFF8646B3C0  
Files Queue: FFFFFFFF20188580  
    Flink: FFFFFFFF200242E0  
    Blink: FFFFFFFF200FA9A0  
Cached Open Files: 3  
Active I/O count: 0  
Cached Closed Files: 4  
Files Ever Opened: 8  
Files Ever Deposited: 2  
Pages Allocated: 21

QIO count: 174 Total: 277504 ( 271.0 KB)

. . . . .

**Read Percentage: 96 %**

**Hit Rate: 71 %**

**I/O Response Times (This Volume)**

**Overall Average: 0.9041 milliseconds**

**Cache Hit: 0.0069 milliseconds**

**Disk: 3.1930 milliseconds**

**SDA>XFC SHOW VOLUME**

# IBM MQ series

- MQ is a heavy user of pthreads
- Set MULTITHREAD to 1
  - Thread manager upcalls are enabled; the creation of multiple kernel threads is disabled

## CA's PSDC (aka DECPS)

- On IA64 generates 2,500 – 100,000 alignment faults per second
- Disable the product during performance benchmarks (or any other important work)
- Alignment faults in:
  - IOPERFORM (VMS fix available)
  - PSDC\$LOADABLE\_IMAGE
    - Fix available from CA

# Sizing Working Sets

- It's not 1980 any more...
- Determine the size of XFC cache + MPW\_HILIMIT
- Subtract the sum from the number of fluid pages on the system (MMG\$GQ\_FLUID\_PGCNT)
- Divide by the maximum number of processes that have ever been running on the system (PMS\$GL\_PROCCNTMAX)
- Multiply the result by 16 to translate from pages to pagelets
- If you are conservative, take 70% of the result and set working set limit and quota to this value
- Working set extent should be 3 times the result
- Make sure PGFLQUOTA is properly sized

# TCP/IP & Gigabit Ethernet

- Using Gigabit Ethernet?
  - Turn on Jumbo frames
  - Frames larger than 1518 bytes, more data per frame -> less frames -> less interrupts -> better performance
  - Must be supported by the switch
  - Must be configured before TCP/IP is started
    - `mc lanep set dev ewa/jumbo`
    - Bit 6 in SYSGEN parameter LAN\_FLAGS

# *The Toolbox*

# Toolbox Overview

- Collection of highly valuable, undocumented & unsupported tools, subject to change without a notice
- Implemented as SDA extensions
- Use hooks in the VMS executive
- May be load and unloaded on the fly
  - No reboot required
- Trace data is stored in ring buffer in S2 space
  - May be viewed from a crash dump

# Toolbox Overview

|                                      | First shipped in |
|--------------------------------------|------------------|
| • CNX connection manager tracing     | V7.2-2           |
| • EXC exception tracing              | V8.2             |
| • FC Fibrechannel debug and tracing  | V7.2-2           |
| • FLT alignment fault tracing        | V8.1             |
| • IO buffered and direct I/O tracing | V7.3-2           |
| • LCK lock manager tracing           | V7.2-2           |
| • LNM logical name tracing           | V7.3-1           |
| • MTX mutex tracing                  | V7.3             |
| • PCS PC sampling                    | V7.3-2           |
| • PRF performance utility            | V8.2             |
| • PSH pshared debug utility          | V8.2-1           |



# Toolbox Overview

|                                       | First shipped in |
|---------------------------------------|------------------|
| • RDB Rdb lock decoding and tracing   | V7.3-2           |
| • RMS indexed file tracing            | V8.2-1           |
| • SPL spinlock tracing                | V7.2-1H1         |
| • TQE timer entry tracing             | V7.3-1           |
| • TR debug and trace prints           | V7.3             |
| • XFC eXtended File Cache diagnostics | V7.3             |

# Toolbox Overview

- Common commands
  - SDA> xxx ! Displays brief command help
  - SDA> xxx LOAD
  - SDA> xxx START TRACE /BUFFER=3000
  - SDA> xxx SHOW TRACE
  - SDA> xxx STOP TRACE
  - SDA> xxx UNLOAD
  - SDA> READ /EXEC /NOLOG

# PRF

- PRF is highly powerful SDA extension for monitoring various performance counters at the processor level.
- May be used for PC sampling.
- Highlights areas in the application that require performance enhancements.

# PRF

```
SDA> prf load
```

```
PRF$DEBUG load status = 00000001
```

```
SDA> prf start pc/ind=21E004DA
```

```
PC Sampling started...
```

```
SDA> prf start collect
```

```
SDA>
```

Now run the application:

```
$ r prime
```

```
ELAPSED:      0 00:00:24.16  CPU: 0:00:24.06  BUFIO: 0  DIRIO: 0  FAULTS: 0
```

```
$
```

- To look at the collected data:

```
SDA> prf show collect
```

# PRF SHOW COLLECT

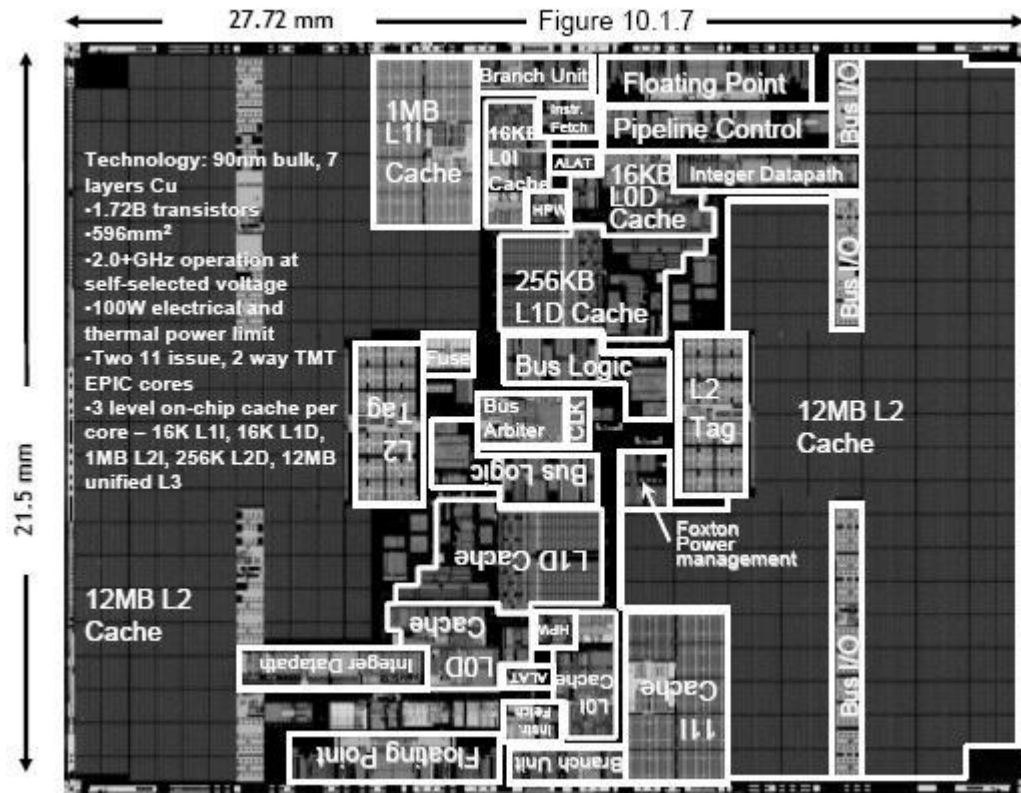
| Start VA           | End VA             | Image                      | Count         | Percent       |
|--------------------|--------------------|----------------------------|---------------|---------------|
| -----              | -----              | -----                      | -----         | -----         |
| FFFFFF802.11F00000 | FFFFFF802.11F01FFF | PRIME                      | <b>305113</b> | <b>99.85%</b> |
| FFFFFF802.A1000000 | FFFFFF802.A1015FFF | Kernel Promote VA          | 1             | 0.00%         |
| FFFFFFFF.80000000  | FFFFFFFF.800000FF  | SYS\$PUBLIC_VECTORS        | 2             | 0.00%         |
| FFFFFFFF.80000100  | FFFFFFFF.800111FF  | SYS\$BASE_IMAGE            | 2             | 0.00%         |
| FFFFFFFF.80011200  | FFFFFFFF.800651FF  | SYS\$PLATFORM_SUPPORT      | 258           | 0.08%         |
| FFFFFFFF.800A0000  | FFFFFFFF.801DD6FF  | SYSTEM_PRIMITIVES          | 88            | 0.03%         |
| FFFFFFFF.801DD700  | FFFFFFFF.80243BFF  | SYSTEM_SYNCHRONIZATION_MIN | 9             | 0.00%         |
| FFFFFFFF.80254600  | FFFFFFFF.8026EFFF  | SYS\$EIDRIVER.EXE          | 5             | 0.00%         |
| FFFFFFFF.8026F000  | FFFFFFFF.802895FF  | SYS\$LAN.EXE               | 2             | 0.00%         |
| FFFFFFFF.80289600  | FFFFFFFF.802BA1FF  | SYS\$LAN_CSMACD.EXE        | 2             | 0.00%         |
| FFFFFFFF.80440E00  | FFFFFFFF.8052B2FF  | IO_ROUTINES                | 1             | 0.00%         |
| FFFFFFFF.8053A600  | FFFFFFFF.80670DFF  | PROCESS_MANAGEMENT         | 7             | 0.00%         |
| FFFFFFFF.80670E00  | FFFFFFFF.807759FF  | SYS\$VM                    | 11            | 0.00%         |
| FFFFFFFF.80779500  | FFFFFFFF.807C76FF  | LOCKING                    | 1             | 0.00%         |
| FFFFFFFF.807C7700  | FFFFFFFF.807F9CFF  | MESSAGE_ROUTINES           | 1             | 0.00%         |

# PRF SHOW COLLECT

SDA> prf show coll/thresh=2

| PC                | Count | Rate   | Symbolization  | Module | Offset   |
|-------------------|-------|--------|--|--------|----------|
| FFFFF802.11F00170 | 63410 | 20.07% | PRIME+10170<br>[GENERATE_PRIME+00000170 / GENERATE_PRIME+00000170] | PRIME  | 00010170 |
| FFFFF802.11F00190 | 6138  | 2.01%  | PRIME+10190<br>[GENERATE_PRIME+00000190 / GENERATE_PRIME+00000190] | PRIME  | 00010190 |
| FFFFF802.11F001A0 | 6761  | 2.21%  | PRIME+101A0<br>[GENERATE_PRIME+000001A0 / GENERATE_PRIME+000001A0] | PRIME  | 000101A0 |
| FFFFF802.11F00200 | 6296  | 2.06%  | PRIME+10200<br>[GENERATE_PRIME+00000200 / GENERATE_PRIME+00000200] | PRIME  | 00010200 |
| FFFFF802.11F00220 | 8102  | 2.65%  | PRIME+10220<br>[GENERATE_PRIME+00000220 / GENERATE_PRIME+00000220] | PRIME  | 00010220 |
| FFFFF802.11F00290 | 6804  | 2.23%  | PRIME+10290  | PRIME  | 00010290 |

# Montecito



Source: Wikipedia

# HyperThreads

- HyperThreads have the potential of improving performance
- Application has to meet the following criteria:
  - COM Queue
  - Poor locality (L2/L3 misses)
  - No pagefulating
- PRF may be used to track L2 misses

*PRF START PROFILE/CPU=n/CACHE=L2/INDEX=PID*

*PRF START COLLECT*



# L2 Cache Misses on TC\_CF (13.2% improvement)

| I-Cache Misses     |                    | D-Cache Misses      |  | Branch Trace Buf |       | Latency | Percent | Latency |
|--------------------|--------------------|---------------------|--|------------------|-------|---------|---------|---------|
| Start VA           | End VA             | Image               |  |                  |       |         |         |         |
| Percent            | Count Percent      |                     |  |                  |       |         |         |         |
| 00000000.00000000  | 00000000.7ADCBFFF  | Process Space       |  |                  | 17062 | 1.73%   | 6072893 |         |
| 96.52%             | 244963 8.62%       |                     |  |                  |       |         |         |         |
| 00000000.7ADCC000  | 00000000.7AEF7FFF  | DCL                 |  |                  | 101   | 0.01%   | 0       |         |
| 0.00%              | 242 0.01%          |                     |  |                  |       |         |         |         |
| FFFFFF802.0806C000 | FFFFFF802.0825DFFF | LIBRTL              |  |                  | 4104  | 0.42%   | 1217    |         |
| 0.02%              | 21753 0.77%        |                     |  |                  |       |         |         |         |
| FFFFFF802.0825E000 | FFFFFF802.08283FFF | LIBOTS              |  |                  | 2150  | 0.22%   | 123     |         |
| 0.00%              | 240662 8.47%       |                     |  |                  |       |         |         |         |
| FFFFFF802.082E8000 | FFFFFF802.0837FFFF | SMGSHR              |  |                  | 52    | 0.01%   | 10      |         |
| 0.00%              | 211 0.01%          |                     |  |                  |       |         |         |         |
| FFFFFF802.08404000 | FFFFFF802.0840DFFF | CMA\$TIS_SHR        |  |                  | 281   | 0.03%   | 0       |         |
| 0.00%              | 1504 0.05%         |                     |  |                  |       |         |         |         |
| FFFFFF802.08444000 | FFFFFF802.084F7FFF | DPML\$SHR           |  |                  | 5     | 0.00%   | 0       |         |
| 0.00%              | 1 0.00%            |                     |  |                  |       |         |         |         |
| FFFFFF802.084F8000 | FFFFFF802.085A9FFF | PTHREAD\$RTL        |  |                  | 2657  | 0.27%   | 294     |         |
| 0.00%              | 6315 0.22%         |                     |  |                  |       |         |         |         |
| FFFFFF802.085AA000 | FFFFFF802.090B3FFF | DECC\$SHR           |  |                  | 24027 | 2.43%   | 6258    |         |
| 0.10%              | 369765 13.02%      |                     |  |                  |       |         |         |         |
| FFFFFF804.0E000000 | FFFFFF804.0E015FFF | Kernel Promote VA   |  |                  | 2232  | 0.23%   | 0       |         |
| 0.00%              | 5191 0.18%         |                     |  |                  |       |         |         |         |
| FFFFFFFF.80000000  | FFFFFFFF.800000FF  | SYS\$PUBLIC_VECTORS |  |                  | 403   | 0.04%   | 41      |         |

## L2 Cache Misses on PRIMES\_1 (Slight Degradation)

| Cache Misses       | Branch Trace Buf   |               | Image |         | Latency |
|--------------------|--------------------|---------------|-------|---------|---------|
| Start VA           | End VA             |               | Image |         | Latency |
| Percent            | Latency            | Percent       | Count | Percent |         |
| -----              | -----              | -----         | ----- | -----   | -----   |
| 00000000.00000000  | 00000000.7ADCBFFF  | Process Space |       |         | 5077    |
| 2.77%              | <b>29968</b>       | <b>52.88%</b> | 26607 | 5.27%   |         |
| 00000000.7ADCC000  | 00000000.7AEF7FFF  | DCL           |       |         | 19      |
| 0.01%              | 0                  | 0.00%         | 22    | 0.00%   |         |
| FFFFFF802.0806C000 | FFFFFF802.0825DFFF | LIBRTL        |       |         | 949     |
| 0.52%              | 570                | 1.01%         | 3816  | 0.76%   |         |
| FFFFFF802.0825E000 | FFFFFF802.08283FFF | LIBOTS        |       |         | 63      |
| 0.03%              | 0                  | 0.00%         | 201   | 0.04%   |         |
| FFFFFF802.082E8000 | FFFFFF802.0837FFFF | SMGSHR        |       |         | 20      |
| 0.01%              | 0                  | 0.00%         | 46    | 0.01%   |         |
| FFFFFF802.08404000 | FFFFFF802.0840DFFF | CMA\$TIS_SHR  |       |         | 0       |
| 0.00%              | 0                  | 0.00%         | 6     | 0.00%   |         |

# LNМ

- The LNМ extension allows tracking logical name translations.
- Logical name translations are expensive from a performance point of view and should be avoided when possible.
- MONITOR IO displays the total number of logical name translations per second

# LNМ Example

```
SDA> lnm show collect
```

```
Logical Name Trace Information:
```

```
-----  
Count      Logical Name  
-----  
5000      SYS$SCRATCH      !SYS$SCRATCH is being translated 5000 times  
10        SYS$SHARE  
10        SYS$SYSROOT  
5         GBL$INS$8DDE9730  
5         SYS$COMMON  
4         GBL$INS$8DDAE310  
4         SYS$OUTPUT  
3         GBL$INS$8DDC20D0  
3         GBL$INS$8DDD1A60  
3         IPC$ACP_NETMBX  
2         CMA$TIS_SHR  
2         DPML$SHR  
2         LIBOTS  
2         LIBRTL  
2         PAS$RTL  
1         GBL$INS$8DDB0B50
```

# LNМ Example

SDA> lnm show trace

Logical Name Trace Information:

```
-----
```

| Timestamp              | CPU | EPID     | Main Image   | CallerPC          | Logical Name                            |
|------------------------|-----|----------|--------------|-------------------|---|
| 25-JAN 06:22:15.530026 | 01  | 21E0040E | IPCACP       | FFFFFFFF.80514560 | IOC\$TRANDEVNAM_C+007C0 IPC\$ACP_NETMBX |
| 25-JAN 06:22:05.530027 | 01  | 21E0040E | IPCACP       | FFFFFFFF.80514560 | IOC\$TRANDEVNAM_C+007C0 IPC\$ACP_NETMBX |
| 25-JAN 06:21:30.440094 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$OUTPUT                             |
| 25-JAN 06:21:30.440010 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | PAS\$OUTPUT                             |
| 25-JAN 06:21:30.439846 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439835 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439825 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439814 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439803 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439792 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439782 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439771 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439760 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439750 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |
| 25-JAN 06:21:30.439739 | 00  | 21E004DA | MANY_TRNLNMS | 00000000.00000000 | SYS\$SCRATCH                            |

## LNМ & Cobol

- Do you have an application written in Cobol?
- COB\$5644
- Contact your support center to obtain a fix

# RMS

- New since V8.2-1
- Uses lock manager trace hooks to trace RMS bucket and RMS record lock operations.
- Either LCK or RMS tracing at a time, but not both simultaneously
- Identifies busy RMS indexed files and which buckets are heavily used

# TR\_PRINT

- Advanced debugging tool
- “printf” to memory
- Works from any mode (user, super, exec and kernel mode)
- “tr\_print” macros available for Macro32, C, Bliss, IAS (IA64 only)
- Control string with formatting directives (no width) and variable number of arguments (Macro32 only allows 5 parameters)



# TR\_PRINT

- Formatting directives:
  - %S, %s                   – zero-terminated string
  - %A, %a                   – ascii string (pointer and length)
  - %D, %d                   – decimal value
  - %X, %x                   – hexadecimal longword
  - %L, %l                   – hexadecimal quadword
  - %W, %w                   – hexadecimal word (V7.3-2 or better)
  - %B, %b                   – hexadecimal byte (V7.3-2 or better)

# TR\_PRINT Example

```
//  
// The main action routine  
//  
int ToggleOnlineBit (struct _ucb *dev_ucb, int BitState)  
{  
    tr_print (("We are now in Kernel mode, UCB address:  
              %X",dev_ucb));  
    dev_ucb->ucb$v_online = BitState;  
    return SS$ _NORMAL;  
}  
  
$ CC module + SYS$LIBRARY:SYS$LIB_C/LIB  
$ LINK/SYSEXE
```

# TR\_PRINT Example

```
$ ana/sys
```

```
OpenVMS system analyzer
```

```
SDA> tr load
```

```
SDA> tr start trace
```

```
Run the application....
```

```
SDA> tr sh trace
```

```
Debug Trace Information:
```

```
-----
```

```
Timestamp
```

```
CPU Buffer
```

```
-----
```

```
-----
```

```
25-JAN 13:54:50.681193 00 We are now in Kernel mode, UCB address: 8646C300
```

```
SDA>
```

# FLT

- Powerful tool for detecting and logging alignment faults.
- For each alignment fault logs:
  - Time of fault
  - CPU encountered the fault
  - Unaligned Virtual Address
  - Access mode
  - Program Counter
  - Process ID
- Allows determining the exact location in the application that generated the alignment fault.

# FLT Example

SDA> flt show trace

Unaligned Data Fault Trace Information:

```
-----
```

| Timestamp              | CPU | Exception         | PC                      | Unaligned VA       | Access | EPID     | Trace Buffer      |
|------------------------|-----|-------------------|-------------------------|--------------------|--------|----------|-------------------|
| 24-JAN 04:31:28.640575 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDDB59 | User   | 21E0049A | FFFFFFFF.7E69B818 |
| 24-JAN 04:31:28.640571 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDDB4A | User   | 21E0049A | FFFFFFFF.7E69B7B0 |
| 24-JAN 04:31:28.640568 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDDB3B | User   | 21E0049A | FFFFFFFF.7E69B748 |
| 24-JAN 04:31:28.640565 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDDB2C | User   | 21E0049A | FFFFFFFF.7E69B6E0 |
| 24-JAN 04:31:28.640562 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDDB1D | User   | 21E0049A | FFFFFFFF.7E69B678 |
| 24-JAN 04:31:28.640559 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDDB0E | User   | 21E0049A | FFFFFFFF.7E69B610 |
| 24-JAN 04:31:28.640556 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDAFF  | User   | 21E0049A | FFFFFFFF.7E69B5A8 |
| 24-JAN 04:31:28.640553 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDAE1  | User   | 21E0049A | FFFFFFFF.7E69B540 |
| 24-JAN 04:31:28.640550 | 01  | 00000000.000100E1 | SYS\$K_VERSION_16+000A1 | 00000000.7ACCDAD2  | User   | 21E0049A | FFFFFFFF.7E69B4D8 |

```
{ .mmi
```

```
0119FA4F0900      00E0      adds      r36 = -8, r36 ; ;
0080C2400900      00E1      ld8       r36 = [r36]
000008000000      00E2      nop.i    0 ; ;
}
```

## Decoding PCs

- New routine to decode PC into module and routine names with offsets (IA64 only)
  - `tf$get_mod_rtn` in module `TRACE_ELF` in `SYSS$SHARE:VMS$VOLATILE_PRIVATE_INTERFACES.OLB`

```
tf$get_mod_rtn ( entry->spltre$q_pc, &mod_name, &rtn_name, &mod_rel_pc, &rtn_rel_pc  
);
```

# Questions?

BRUDEN-OSSG thanks you for attending this session.

See us at [www.BRUDENOSSG.com](http://www.BRUDENOSSG.com) for:

- *Performance analysis*
  - *(Performance Results Or No Expense)*
- *Porting assistance*
- *Special OPS (OpenVMS Programming Services)*
- *Support*