

MEMU2 for DB2 10 and 11 for z/OS *(How many threads can you run?)*

Judy Ruby-Brown

IBM Corp.

zGrowth DB2 Washington Systems Center

Tuesday, March 3, 2015

Session Number 16549



SHARE is an independent volunteer-run information technology association
that provides **education, professional networking and industry influence.**



Acknowledgements and Disclaimers

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

© **Copyright IBM Corporation 2015. All rights reserved.**

- **U.S. Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.**

IBM, the IBM logo, ibm.com, **DB2**, and **FlashCopy** are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or TM), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at

- “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml
- Other company, product, or service names may be trademarks or service marks of others.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Why this presentation?

- Users want to know: “I am in DB2 10 - NOW how many threads can I run?”
- Update or Replacement for “DBM1 Virtual Storage - How to use MEMU2” is needed
 - <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS343>
- Its MEMU2 analysis worked only for DB2 V8 and partially V9
 - Did not have QW0225BB – buffer pool stats for V9
- DB2 10/11 MEMU2 now out – very different
 - Trace record IFCID 225 rewritten and large growth



Get “MEMU2 for DB2 10 and 11” at:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS5279>

- This presentation package also includes a sample spreadsheet to make it easier to determine your DBM1 storage usage (described later)

What We Will Learn

1. Contents of MEMU2 package for DB2 10/11.
2. Using the  sample spreadsheet
3. Importing MEMU2 output into spreadsheet
4. Creating/Using  Charts
5. Summary

Goal: to teach YOU how to run MEMU2 and make sense of the output!

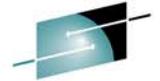
Note 1: Many thanks to Adrian Burke, IBM DB2 development “SWAT” team, for his significant contributions to these materials 

Note 2: Spreadsheets and charts produced Using Microsoft® EXCEL® 2013.

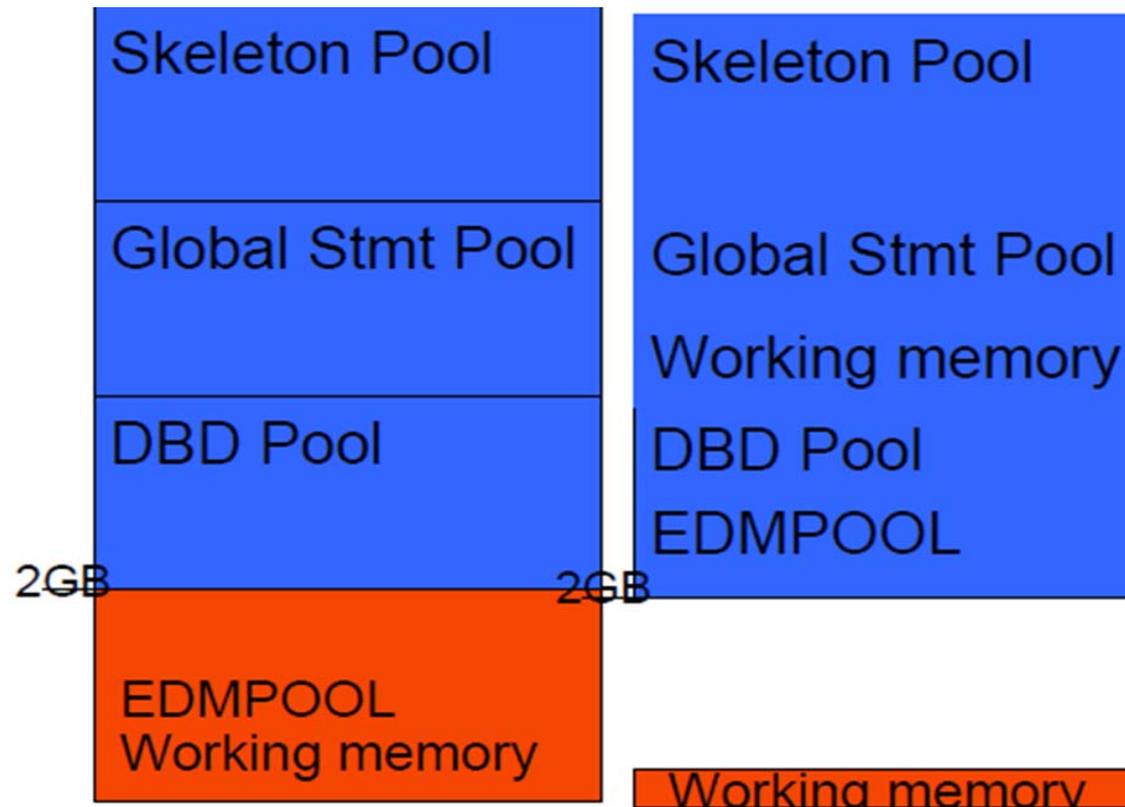
- Other versions or packages may vary
- No Macros are used

DBM1 Virtual Storage

- **Before DB2 10, large users had to restrict number of threads in a DB2 member to avoid DB2 failure due to lack of virtual storage**
 - Practical limit usually < 900 – some have 2 members per LPAR for that reason
 - DB2 10 allows up to 20K threads
- **Problem was inability to determine how many threads could be safely run in a specific DB2 – very workload dependent**
- **MEMU2 is a free package, downloadable from IBM developerworks.**
 - REXX exec runs on each DB2 member, waking up every 1 minute (default) and writing in CSV form the DB2 trace records pertaining to virtual storage
 - Output data set can be downloaded to PC and spreadsheet created that identifies:
 - Current number of local and distributed threads
 - Maximum number of threads that can be run
 - You can then set threads in DSNZPARM properly/safely.



DBM1 Virtual Storage Constraint Relief in DB2 10 +



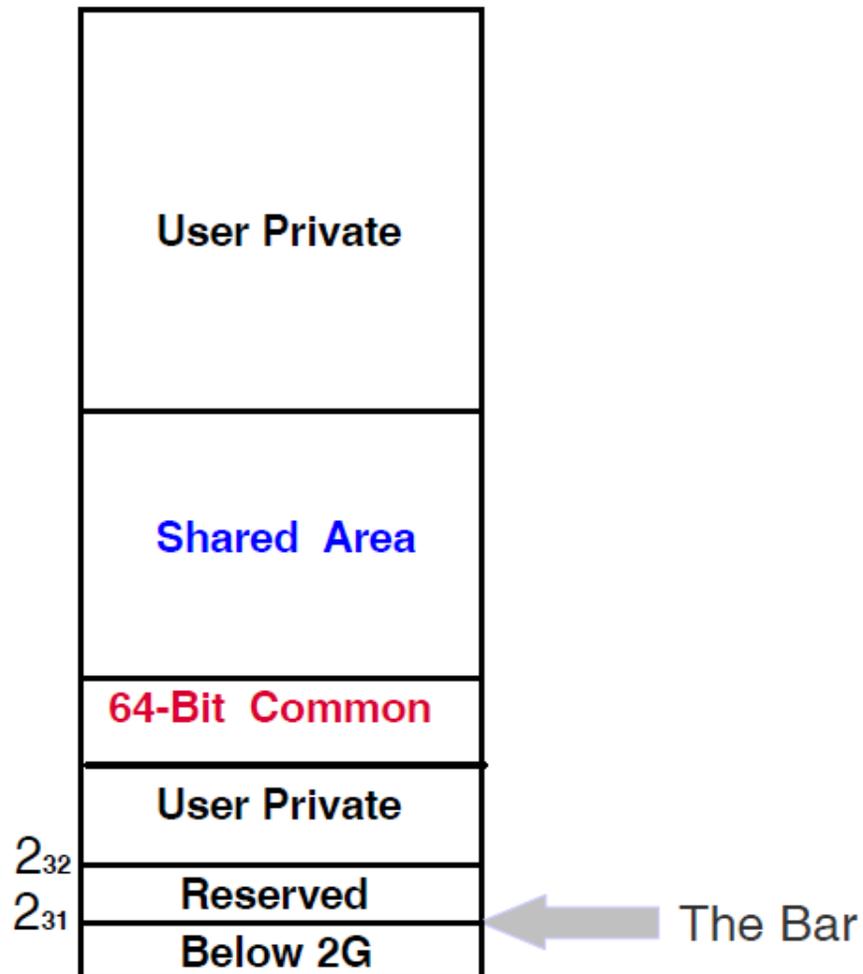
- 90% of storage below the 2GB bar was moved above it
- Realize benefits when you rebind packages
- DB2 11 moved xPROC above the bar

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

z/OS View of Virtual Storage Map

Specified in SYS1.PARMLIB
Member IEASYSxx

HVSHARE
HVCCOMMON



Virtual Storage above the Bar

- z/OS mapping of virtual storage above the 2 GB bar
 - DB2 first product to use it, beginning in DB2 Version 8
- **64 bit Private**
- No sharing across DB2 system address spaces
- Primarily local buffer pools
- Fixed RID pool
- EDM Pool

64 bit Shared Virtual Storage above the Bar

- DB2 uses in preference to 64 bit Private - now dominates DB2
 - Token required to connect to a memory object (thus a little “shared” but only among DB2 address spaces)
 - Not subject to overlays commonly found for other users of CSA/ECSA
 - DB2 requires 128G (DB2 10) and 1T (DB2 11) to start
 - Hard to track if multiple DB2s on LPAR, since HVSHARE area is by LPAR
- Uses -
 - Thread storage pools shared, no need for \$\$ cross memory moves
 - Stack storage (save areas, working storage) used by thread to execute SQL
 - Variable storage - dynamic statement caches: CT, PT, SKCT, SKPT
 - DDF and some utilities
- z/OS treats like common storage – spin lock used more for serialization

64 bit Common Virtual Storage above the Bar

- Like CSA/ECSA – any address space can use it
 - Currently DB2 and Communications Server
- Allows DB2 to move some common areas
 - Blocks where connect token for shared 64 bit is unknown in advance
 - DB2 11 log output buffers (OUTBUFF + 15%)
- DB2 gets a 6G object for the common objects
- Specified as HVCOMMON to z/OS by LPAR

- **Warning**
- Each DB2 on the LPAR needs 128G (10) and 1T (11) in HVSHARE
 - (default 512T is a LOT of virtual storage)
- Each DB2 on the LPAR needs 6G in HVCOMMON
 - (default 64G)
- If one or more of those values are insufficient, a DB2 fails to start.
 - When a user started his 11th DB2 subsystem, it failed due to HVCOMMON

IFCID 225 changes since DB2 9

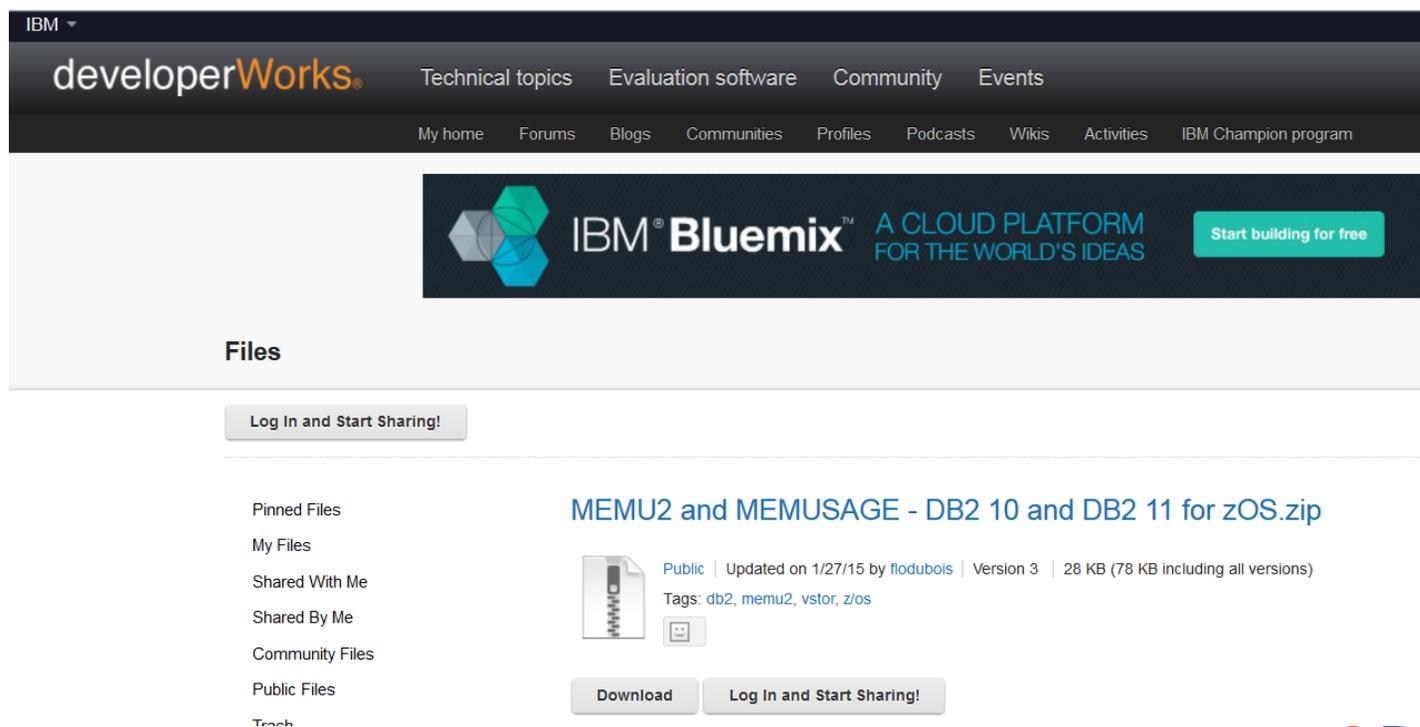
- MEMU2 output CSV (primarily IFCID 225) grew from 56 to 170 columns in DB2 11
 - Completely reformatted with repeating sections
 - This snippet taken from DSN1110.SDSNIVPD(DSNWMSGSGS)

```
_____IFCID 0225_____
(RMID 20)
THIS IFCID SUMMARIZES SYSTEM STORAGE USAGE.
IFCID 0225 RECORDS ARE DIVIDED INTO THE FOLLOWING
SECTIONS:
DATA SECTION 1: ADDRESS SPACE SUMMARY (QW0225)
                  THIS DATA SECTION CAN BE A REPEATING GROUP.
                  IT REPORTS DATA FOR THE DBM1 AND DIST
                  ADDRESS SPACES. QW0225AN IDENTIFIES
                  THE ADDRESS SPACE THAT IS BEING DESCRIBED.
DATA SECTION 2: THREAD INFORMATION (QW02252)
DATA SECTION 3: SHARED AND COMMON STORAGE SUMMARY (QW02253)
DATA SECTION 4: STATEMENT CACHE AND XPROC DETAILS (QW02254)
DATA SECTION 5: POOL DETAILS (QW02255)
DATA SECTION 6: POOL DETAILS (QW02256)

THIS RECORD IS WRITTEN AT ONE-MINUTE INTERVALS.
```

Getting MEMU2 for DB2 10/11

- **Google MEMU2 DB2 11**
 - **First in list is for Developerworks below**
- Download it (free)
- Instructions for running MEMU2 are contained in the zip file



The screenshot shows the IBM DeveloperWorks interface. At the top, there's a navigation bar with 'IBM' and 'developerWorks' logo, and links for 'Technical topics', 'Evaluation software', 'Community', and 'Events'. Below this is a secondary navigation bar with 'My home', 'Forums', 'Blogs', 'Communities', 'Profiles', 'Podcasts', 'Wikis', 'Activities', and 'IBM Champion program'. A large banner for 'IBM Bluemix' is visible, with the text 'A CLOUD PLATFORM FOR THE WORLD'S IDEAS' and a 'Start building for free' button. Below the banner, the 'Files' section is active, showing a 'Log In and Start Sharing!' button. On the left, there's a sidebar with 'Pinned Files', 'My Files', 'Shared With Me', 'Shared By Me', 'Community Files', 'Public Files', and 'Trash'. The main content area displays a file titled 'MEMU2 and MEMUSAGE - DB2 10 and DB2 11 for zOS.zip'. The file is public, updated on 1/27/15 by flodubois, version 3, and 28 KB (78 KB including all versions). It has tags 'db2, memu2, vstor, z/os'. There are 'Download' and 'Log In and Start Sharing!' buttons for this file.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

MEMU2 Package Contents

- “MEMU2 and MEMUSAGE – DB2 10 and DB2 11 for zOS.zip”
 - MEMU2V11 - REXX exec to output IFCID225 and others as CSV. Invoked as batch job (works also for V10)
 - MEMUZV11 – outputs IFCID225 if invoked from TSO Option 6
 - Returns IFCID225 immediately for current moment – one time shot
 - Memu2jcl – JCL to invoke MEMU2 REXX as a batch job
 - ReadMe.doc – documentation to install, modify, and use this package



CSVSTO11 – OMPE Spreadsheet Input Generator field list to produce a MEMU2-like .csv file

MEMU2 REXX batch job

- MEMU2 REXX –
 - Invoked via JCL on single member basis
 - Outputs IFCID225 info to comma delimited data set
 - 1 minute interval (default) independent of STATIME (you can specify more time)
 - 60 iterations default (one hour)
 - JCL specifies SSID and overrides to above
- Collection of IFCID225 is triggered by Statistics Class 1

JCL to invoke MEMU2 for 1 hour

```
//JRBMEMU JOB (????), 'JUDY R-B',MSGCLASS=0,REGION=0M,
// CLASS=A,NOTIFY=????????
/*JOBPARM SYSAFF=SYSA
/** COMMENT THE /**OUTP THE FIRST TIME, SINCE IT IS NOT CREATED
//DEL001 EXEC PGM=IEFBR14
//OUTP DD DISP=(MOD,DELETE,DELETE),
// DSN=JUDYRB.MEMU2.OUTPUT.CSV
//MEMU2 EXEC PGM=IKJEFT01,DYNAMNBR=25,ACCT=SHORT,
// REGION=4096K
//STEPLIB DD DSN=DSN1110.SDSNLOAD,DISP=SHR
//SYSEXEC DD DSN=JUDYRB.WSC.JOBS,DISP=SHR WHERE 80 CHAR REXX LIVES
//OUTP DD DISP=(,CATLG,CATLG),
// DSN=JUDYRB.MEMU2.OUTPUT.CSV,
// DCB=(RECFM=VB,LRECL=4096,BLKSIZE=0),
// SPACE=(TRK,(5,1),RLSE),UNIT=SYSDA
//SYSUDUMP DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSTSIN DD *
MEMU2V11 DBP1 60
```

Examples of input:

- MEMU2V11 DBP1 every minute for one hour (60 records output)
- MEMU2V11 DBP1 1440 every minute for 24 hours (1440 records)
- MEMU2V11 DBP1 288 300 every 5 minutes for 24 hours (288 records)

CSVSTO11 – OMPE Spreadsheet generator input File

- *For OMPE users*
- **If you use OMPE CSV Generator**
 - The default storage file is TKANSAMF(FPEPSTOR)
 - The fields do not align with MEMU2 as they are ordered differently.
 - If you decide you would like to use the MEMU2 spreadsheet we provide, you can replace the FPEPSTOR file with CSVSTO11.

Sample spreadsheet provided

- Some columns in spreadsheet are marked “(S)”
 - They are for serviceability and are not documented
- The meaning of most columns is not explained
 - The black columns are the output of running MEMU2
 - The blue additional columns are used to calculate the Maximum threads possible for this DB2
 - Similar to those provided for DB2 V8
 - The Orange additional columns provide real and auxiliary storage for this DB2
- “But WAIT - - There’s MORE!”
 - Sample charts are also provided

A word about Sample spreadsheet

- The “Data” worksheet is set up for up to 10K rows of data
- 3 Chart worksheets initially set up for up to 1441 rows
- It helps you more easily use the output of MEMU2
- Sample was created from run of MEMU2 on an internal IBM DB2 11 member of a data sharing group
 - One hour run at 1-minute intervals
- Charts provided have straight lines – yours will not
 - Because the DB2 subsystem is inactive
- Additional columns (blue and orange) have been added, beginning in column FP
 - (we leave column FO blank as a documentation hint)
- The following slides show the added columns/formulas used

Formulas for Max number threads (cell GA)

| | FP | FQ | FR | FS | FT |
|---|--------------------------|-----------------|-----------------|---|--|
| 1 | MVS ext region size (MB) | Storage Cushion | Storage Warning | ASID DBM1 31-bit priv total storage allocated | ASID DBM1 Getmained+Fixed+Stack+Variable |
| 2 | =D2 | =K2+L2+M2 | =D2-FQ2 | =G2+H2 | =N2+P2+Q2+R2 |

| | FU | FV | FW | FX |
|---|--|--------------------------|------------------------|-----------------------------|
| 1 | ASID DBM1 31-bit priv non-DB2 storage (MB) | Basic Cushion' | Basic Warning' | ASID DBM1 Allowable storage |
| 2 | =H2-FT2 | =IF(FQ2>0,FQ2+100,"n/a") | =IF(D2>0,D2-FV2,"n/a") | =FW2-FU2 |

| | FY | FZ | GA | GB |
|---|--------------------------------|------------------------------|-------------------------------------|-----------------------------|
| 1 | ASID DBM1 Available thread use | Average thread footprint | Max # of active threads (projected) | Current # of active threads |
| 2 | =FX2-G2-Q2-R2-DY2 | =IF(GB2>0,(P2-DY2+N2)/GB2,0) | =IF(FZ2>0,ROUND(FY2/FZ2,0),"n/a") | =BN2+BO2 |

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Real Storage Cells (GC-GL)

| | GC | GD | GE | GF |
|---|-----------------------------|-------------------------|--------------------------|-------------------------|
| 1 | Total REAL in use - DB2 | Total AUX in use - DB2 | Total REAL in use - LPAR | Total AUX in use - LPAR |
| 2 | =Y2+BD2+CU2+CW2+CY2+ET2+EU2 | =Z2+BE2+CV2+CX2+CZ2+EV2 | =Y2+BD2+CQ2 | =Z2+BE2+CR2 |

| | GG | GH | GI | GJ | GK | GL |
|---|--|---|-----------------------------------|----------------------------------|--|---|
| 1 | ASID DBM1 REAL in use for 31-bit priv (MB) | ASID DBM1 AUX in use for 31-bit priv (MB) | ASID DBM1 REAL in use for BP (MB) | ASID DBM1 AUX in use for BP (MB) | ASID DIST REAL in use for 31-bit priv (MB) | ASID DIST AUX in use for 31-bit priv (MB) |
| 2 | =Y2-AA2 | =Z2-AB2 | =AA2-AE2 | =AB2-AF2 | =BD2-BF2 | =BE2-BG2 |

- Formulas used to determine
 - Real storage in used by this DB2
 - AUX (auxiliary storage) in use by this DB2

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Create Spreadsheet from MEMU2 output

- Download MEMU2.OUTPUT.CSV to hard drive
 - The output has DB2 version in column A
 - Subsystem id SSID column B
 - Date/Time in column C
 - IFCID225 data in rest of columns, starting in column D

Two methods shown

1. Create new spreadsheet and charts using your MEMU2 output
2. Import your output from MEMU2 directly into sample spreadsheet and charts

Method 1: Creating new spreadsheet and charts

- Open MEMU2.OUTPUT.CSV file in Microsoft Excel
- Copy the blue and orange columns from the sample spreadsheet beginning in column FP (leave FO blank as a documentation hint to show where columns were added)
- Create charts as described in the “Sample Charts” section of this presentation (covered later).

Method 2: Importing your MEMU2 data (one day)

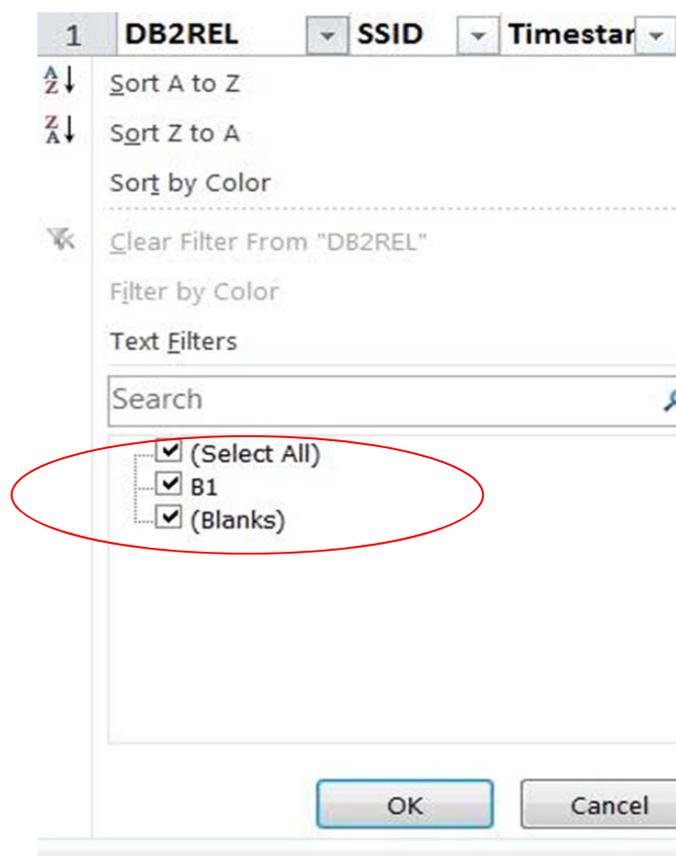
1. Bring up a copy of the Sample (safety sake!)
2. Click on “Data” tab at top of Excel
3. Click on “Refresh all” icon
 - Once you have imported a CSV, you can click “Refresh”
4. Text Import screen appears – choose your MEMU2 output
5. Click Import; your data is imported to spreadsheet, starting on row 2
 - Overlays existing data
6. The “Data” will populate the 3 charts automatically
 - The 3 charts are set up for one day or 1440 rows

Filtering Enables the Chart Population

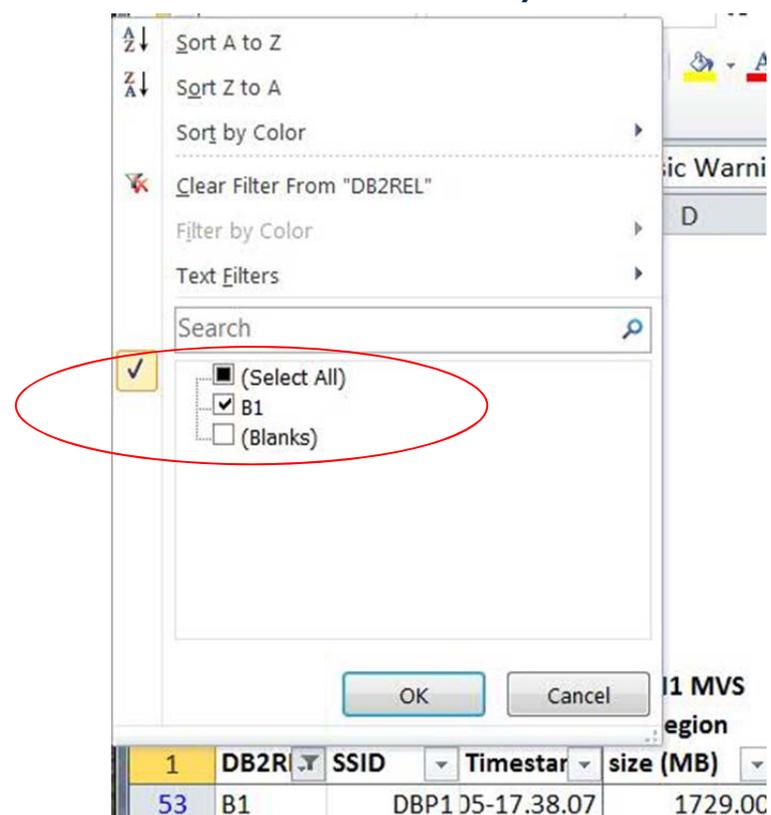
- “Data” worksheet has 10K rows initially, even without data
 - Don’t delete rows or you can eliminate the filter
- Charts can populate up to 10K rows from “Data”
 - If there are more than 10K rows, it is easier to copy the sample to a different work sheet and fill that up
 - Otherwise the charts have to be recreated
- There is a Filter on Column A, DB2REL. It is used so that only populated rows show on the charts
 - No blank rows from the “Data” worksheet appear in charts.
- Filter is originally “ON”
- Next slide shows how to turn the filter on and off
 - *Click on “down” arrow on row 1 of Column A to see the formula*

Filter Settings on DB2REL (Column A)

Filter OFF – Importing new data > row 1441 turns off the filter (“blanks” box checked)



Filter ON – uncheck “blanks” box. Charts populate correctly (Only display on chart the rows where DB2REL is NOT “blanks”).



Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Copying More Data than 1441 Rows

- The “Refresh” action to import MEMU2 places the data at the beginning of the spreadsheet
- If you want to string the MEMU2 data for a week together, you can copy it in a CSV file before importing it the first time
- Position your cursor at the next unfilled row in what will be your CSV “Master”
- On each day’s MEMU2 CSV output, copy from row 2 – 1441 on your input
- Paste it into your “master” CSV

Importing More Data

- Bring up a copy of the Sample spreadsheet
- Under “Data” click “Refresh all” and import from your Master CSV file
- Copy the blue and orange formulas and paste them in the newly added rows
- Go to the Column A filter
 - It will be turned off if you copied more data than existed previously (1441 or the first day)
 - If you look at the charts, they will be all bunched up on the left side of each chart – this is NOT what you want!
- Turn the Column A filter ON by unchecking the “blanks” box
- Now the charts will be correctly populated.

Creating Charts

- 31 bit DBM1 Private
- Real and AUX for this DB2
- Real – Available for this LPAR
- These charts automatically produced if you import your MEMU2 output directly into the sample spreadsheet (Method 2)

Charting with different spreadsheet versions

- The charts shown in this section and their formatting were created using Microsoft® EXCEL® 2013
- Microsoft Excel 2010 was used initially
- Other products/versions may have to be formatted differently
- That is why each column is listed for each chart and its characteristics: line or stacked

To create charts

1. Add a blank spreadsheet worksheet (modifying chart is easier)
2. Use CTRL key to select columns desired
3. Hold down CTRL key while pressing INSERT
4. Click Recommended charts
 - Click first LINE chart
 - Release CTRL
5. Chart will be created in “Data” worksheet
6. Right click on chart and select “move chart”
 - A popup will ask where to place it – choose radio button for the sheet from #1
7. The chart will be placed in your new worksheet
8. Make all modifications in terms of chart type and entering the correct title name in this worksheet by right clicking on the chart

For any stacked charts

- After creating the line chart, the charts need further customization
 - Good reason for using separate worksheet
- Right click on the chart and choose “Change Chart Type”
- Choose “Combo” on the left and “Stacked” on the top picture
- You then see the format on the next slide
- On the middle bottom you see the chart series columns
 - Choose each one of the columns as stated (stacked or line) as directed on slides describing each chart

31-bit Private / Maximum Number Threads Setup

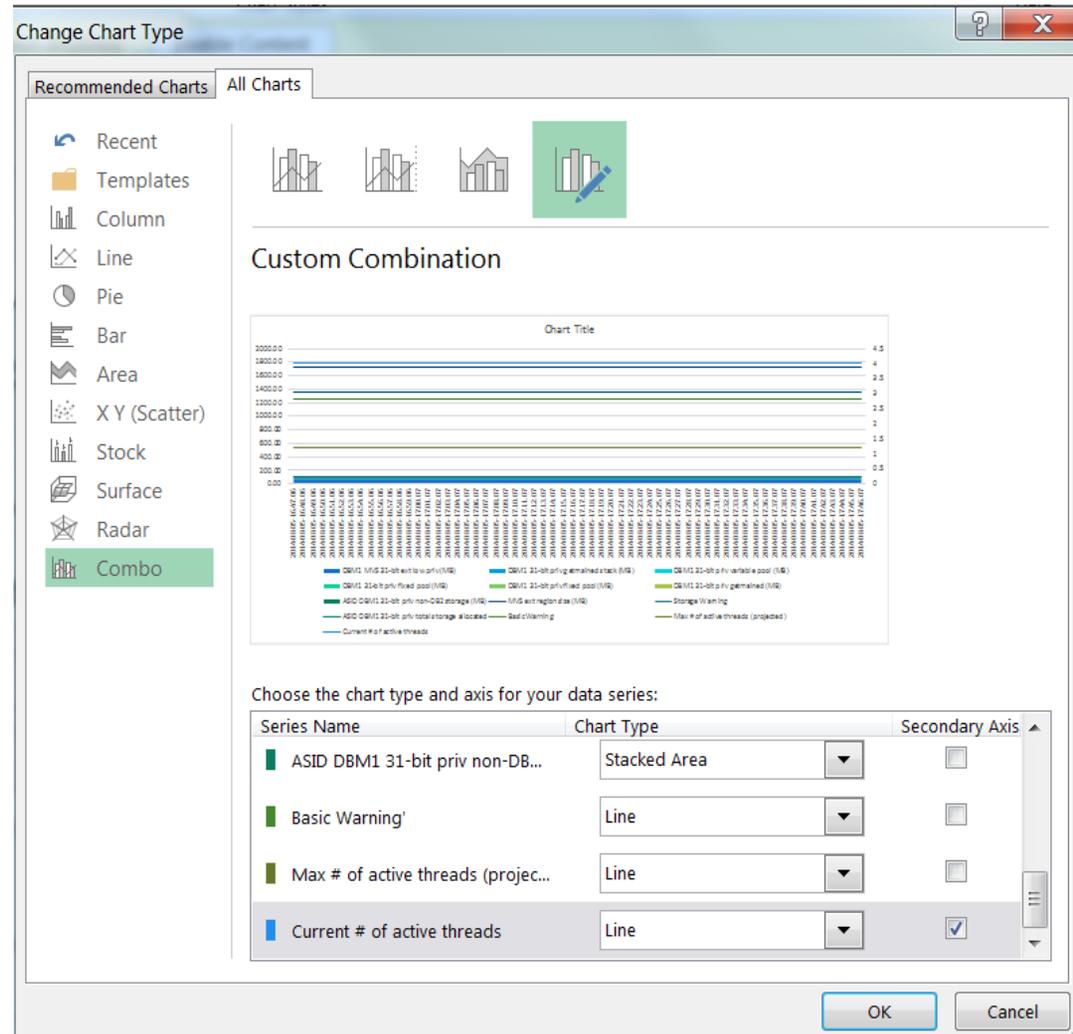
Chart type is “Combo” stacked

Column C (timestamp)
Stacked area columns

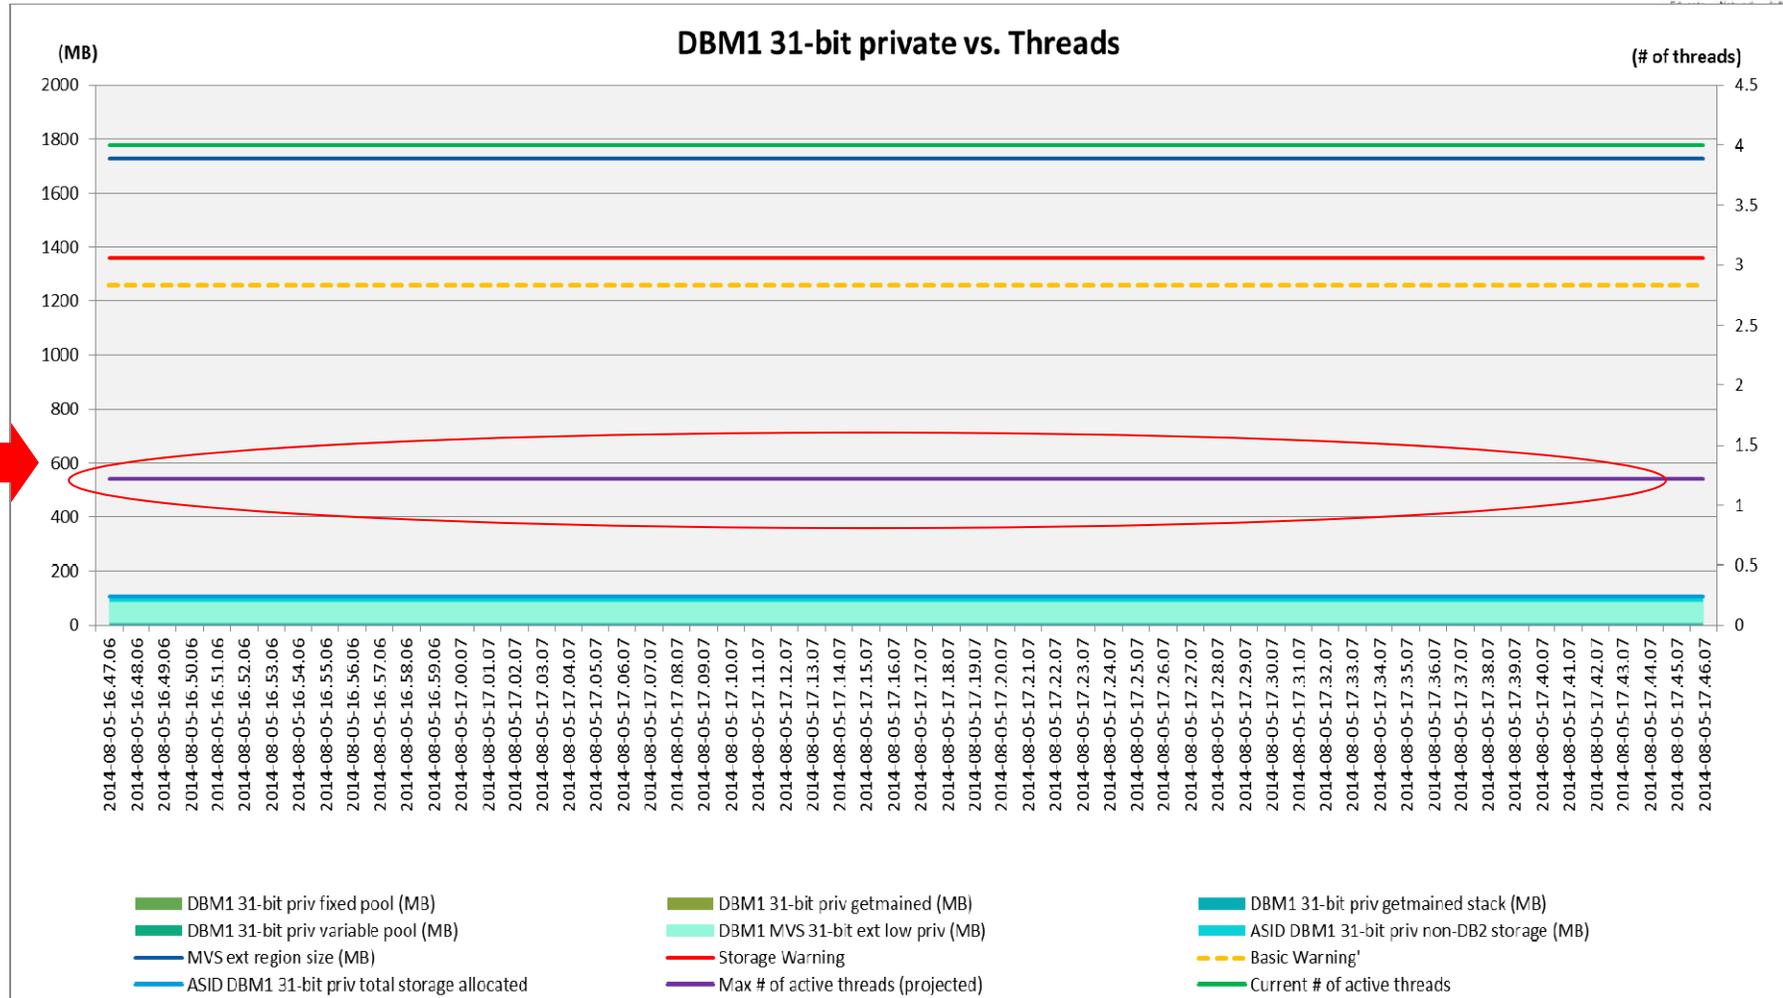
- Q, R, N, P, G, FU

Line columns

- FP – ASID 31-bit priv storage allocated
- FS – MVS Extended Region
- FR – Storage warning
- FW – Basic Warning
- GB – Current # Active Threads
- GA – Max # Threads (projected)



31-Bit Private Virtual Storage Chart



Purple line is the Maximum Projected Threads – (red oval)

Green line is the current # active threads – scale shown on secondary axis (right)

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

31-bit Private Storage Chart

- This chart provides the answer to the question asked: How many threads can I run?
 - It is similar to the one from DB2 V8/9
- • The answer is the “Max # of active threads” (GA)
- As you rebind more packages in DB2 10, you can expect to see this maximum number increase, since thread storage will migrate from the 31-bit to 64-bit storage.
- Chart also maps the main components of DB2 31-bit storage
 - Lines shown for critical health
 - MVS Extended Region Size
 - Storage warning – Threads cancelled with E2003 and E20016 when reached
 - Basic warning – Storage contraction occurs when reached
 - ASID 31-bit private storage allocated – for DBM1

Real and AUX DB2 Charts

- Real chart shows
 - Components of real storage for this DB2
- Real applies to this DB2
 - Exception: real in use for 64-bit common (MB)
 - It is shared across the entire LPAR
 - Shared among all DB2 subsystems on this LPAR
 - Communications Server uses
- AUX chart shows
 - Components of AUXiliary storage for this DB2

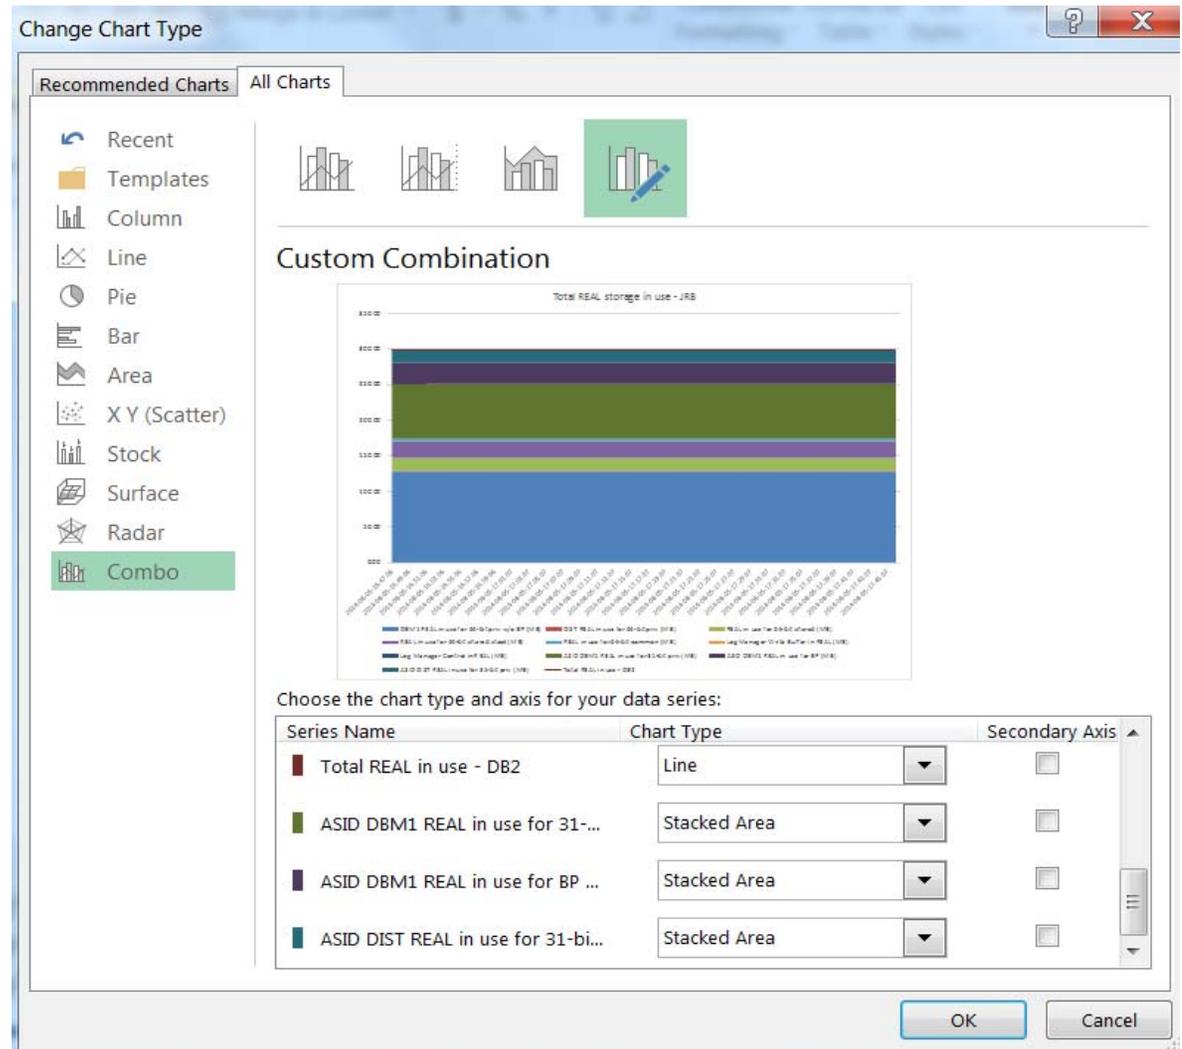
Real Storage Chart Setup

After creating the line chart, change Chart type to “Combo”

- Stacked Bar

The columns to select are

- C (time horizontal axis)
- Stacked Area
 - **AE, BF, CU, CW**
 - **CY, ET, EU, GG**
 - **GI, GK**
- GC – line



Change Chart Type

Recommended Charts | All Charts

Recent
Templates
Column
Line
Pie
Bar
Area
X Y (Scatter)
Stock
Surface
Radar
Combo

Custom Combination

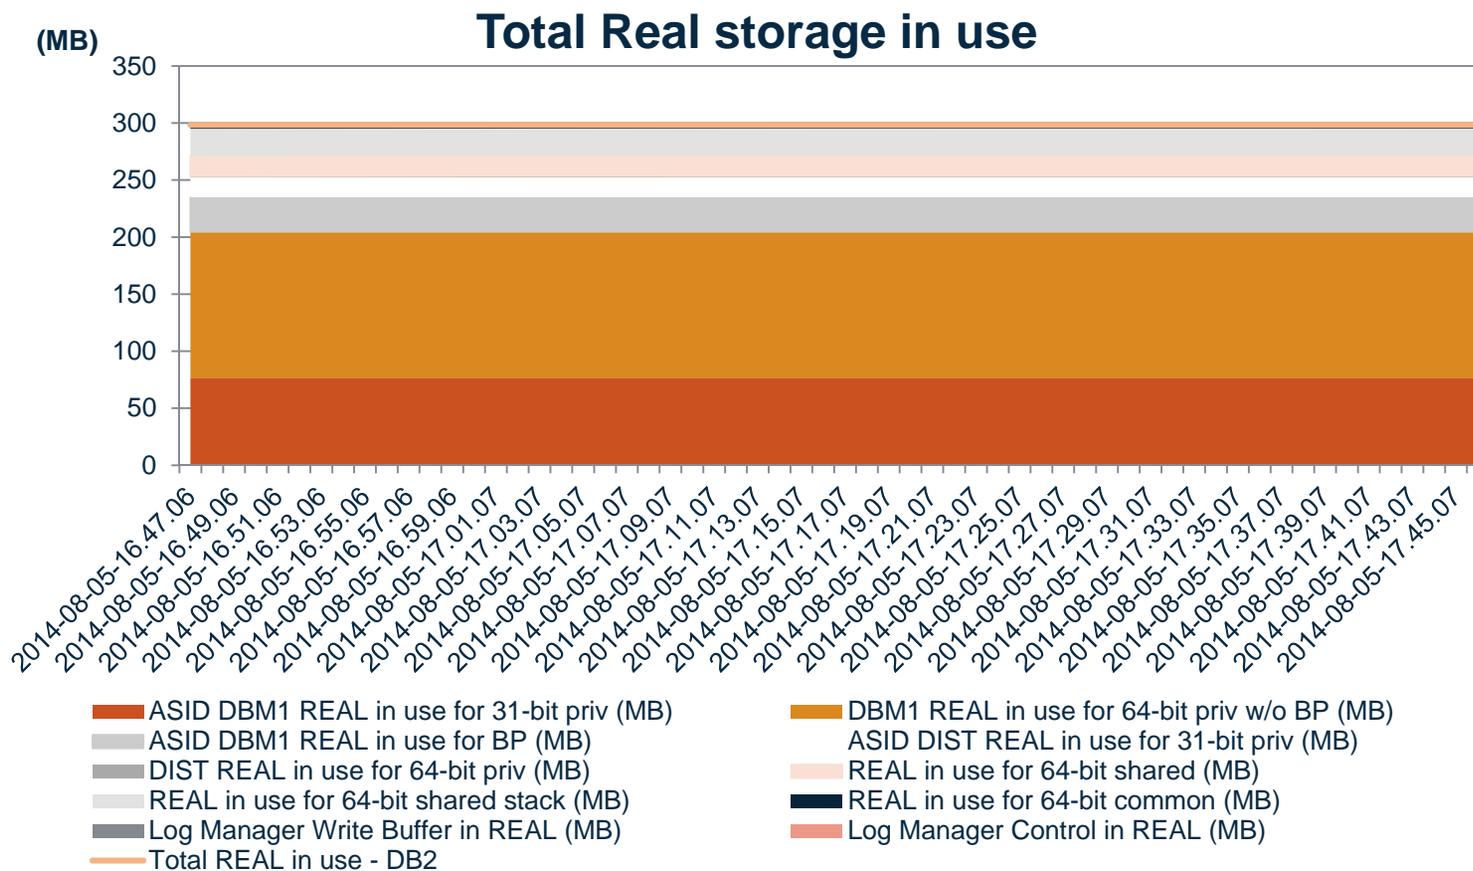
Total REAL storage in use - JRS

Choose the chart type and axis for your data series:

| Series Name | Chart Type | Secondary Axis |
|------------------------------------|--------------|--------------------------|
| Total REAL in use - DB2 | Line | <input type="checkbox"/> |
| ASID DBM1 REAL in use for 31-... | Stacked Area | <input type="checkbox"/> |
| ASID DBM1 REAL in use for BP ... | Stacked Area | <input type="checkbox"/> |
| ASID DIST REAL in use for 31-bi... | Stacked Area | <input type="checkbox"/> |

OK Cancel

Real Storage for this DB2 Chart



Columns can be specified in any order as this chart is stacked
(in terms of storage consumption, not its physical location)

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Paging (AUX) Basics

- Paging rate is a critical performance measure for any z/OS system
 - When shortage of REAL frames occurs, frames are moved to AUX (DASD)
 - Having DB2 paged out is a not good thing for performance
 - Paging should be minimized
- Page fixing buffer pools is a good idea for performance
 - It avoids page fix and page free for high activity buffer pools (heavy I/O)
 - Page fixing 1M size real storage page frames reduces TLB misses (saves CPU)
 - Bad if paging – z/OS converts 1M to 256 x 4K small frames
 - Cannot be used for long term page fix; z/OS expects to recombine later
- If insufficient REAL storage provisioned for the LPAR
 - LPAR begins to page and DB2 is a candidate for page stealing
 - Thread and EDM Pool storage is paged out
 - Performance problems as data is rapidly paged back in (no prefetch)

AUX DB2 Stacked Combo Chart

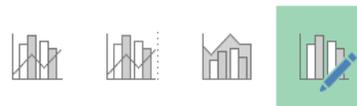
After creating the line chart, change Chart type to “Combo”

- Stacked Bar

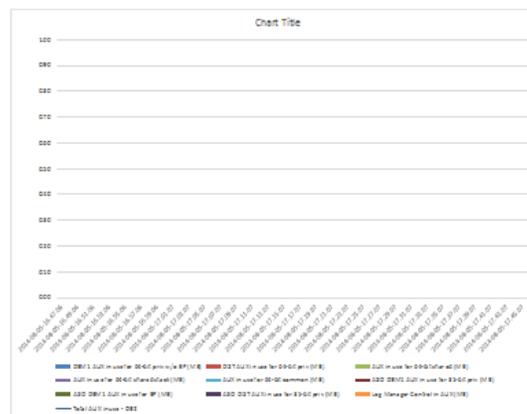
The columns to select are

- C (timestamp)
- Stacked Area
 - **GH**
 - **GJ**
 - **GL**
 - **AF**
 - **BG**
 - **CV**
 - **CX**
 - **CZ**
 - **EV**
- GD - line

- Recent
- Templates
- Column
- Line
- Pie
- Bar
- Area
- XY (Scatter)
- Stock
- Surface
- Radar
- Combo**



Custom Combination



Choose the chart type and axis for your data series:

| Series Name | Chart Type | Secondary Axis |
|------------------------------------|--------------|--------------------------|
| Total AUX in use - DB2 | Line | <input type="checkbox"/> |
| ASID DBM1 AUX in use for 31-... | Stacked Area | <input type="checkbox"/> |
| ASID DBM1 AUX in use for BP (...) | Stacked Area | <input type="checkbox"/> |
| ASID DIST AUX in use for 31-bit... | Stacked Area | <input type="checkbox"/> |

Total AUX DB2 Storage in use Chart

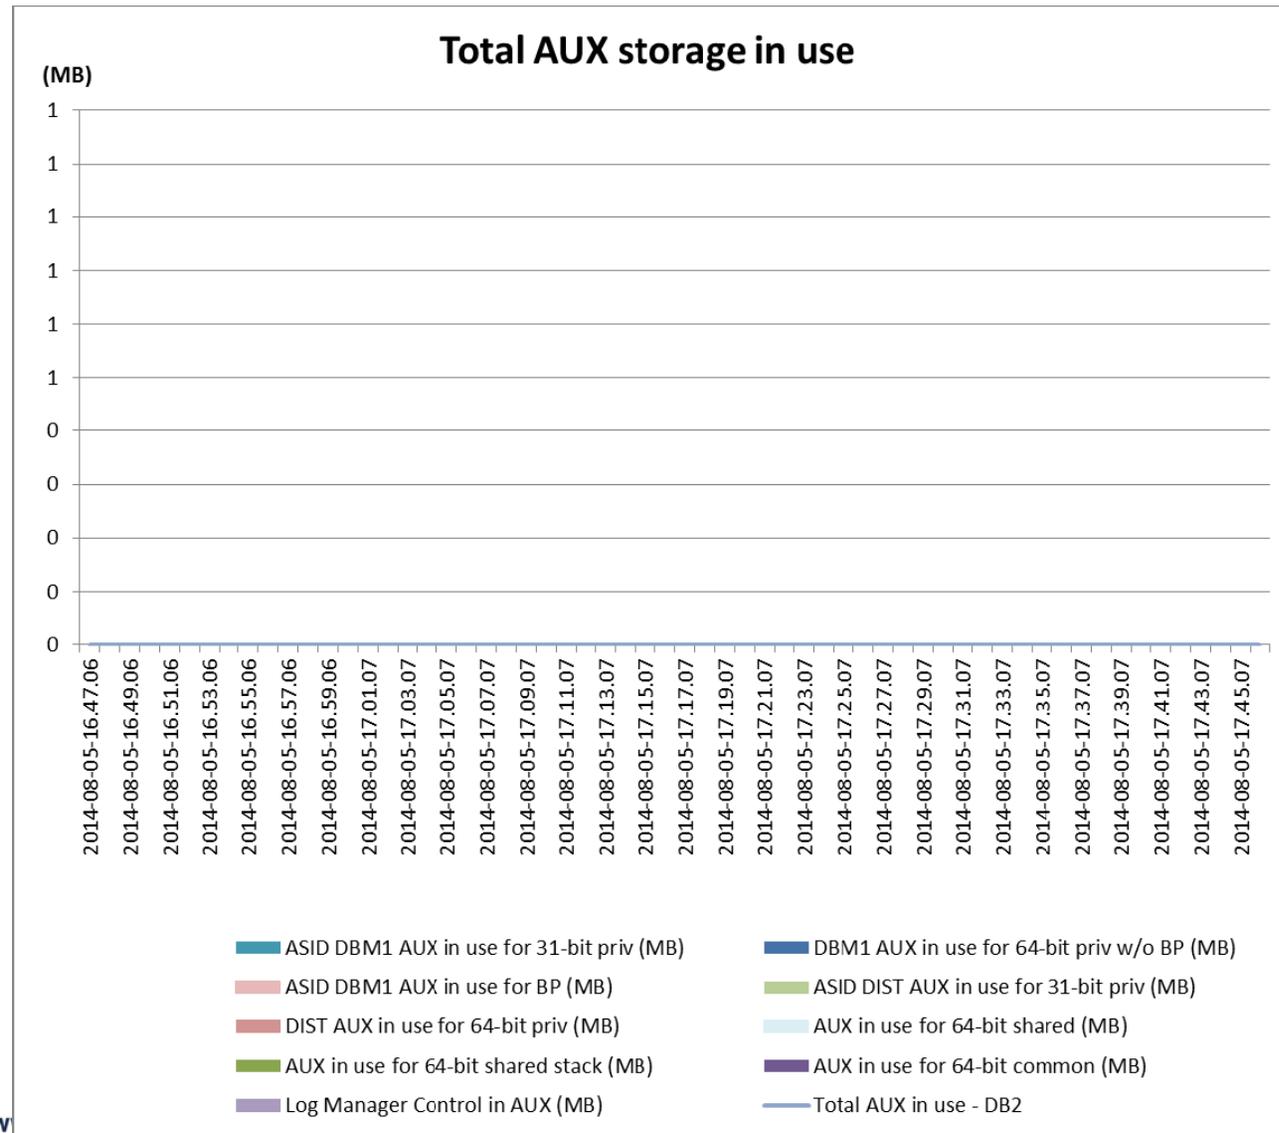
Chart uninteresting for sample because no AUX is in use

(This is what you want!)

Causes for AUX are

- Paging due to general storage constraint
 - Dump of another ASID that caused DB2 storage to be paged temporarily
 - DFSORT with no storage bounds also culprit

- You can increase real storage



Complete your session evaluations online at w

Real Storage Available for LPAR Chart

It is a line chart

- The Real Storage Available in the chart is itself a serviceability field
 - May be obtained from z/OS metrics
 - It is real storage still available, not the defined storage for this LPAR

The other two lines are

- Real Storage in use (blue) by DB2
- Auxiliary storage in use (red/orange) by DB2
 - This shows if some DB2 storage is being paged to AUX and you can tell which pieces looking at the Real-DB2 chart
 - Other ASIDs could be paging; other work can be affected
 - Take actions (increase real storage) to avoid paging
 - If > 1 DB2 on this LPAR, separate MEMU2 / charts for each needed

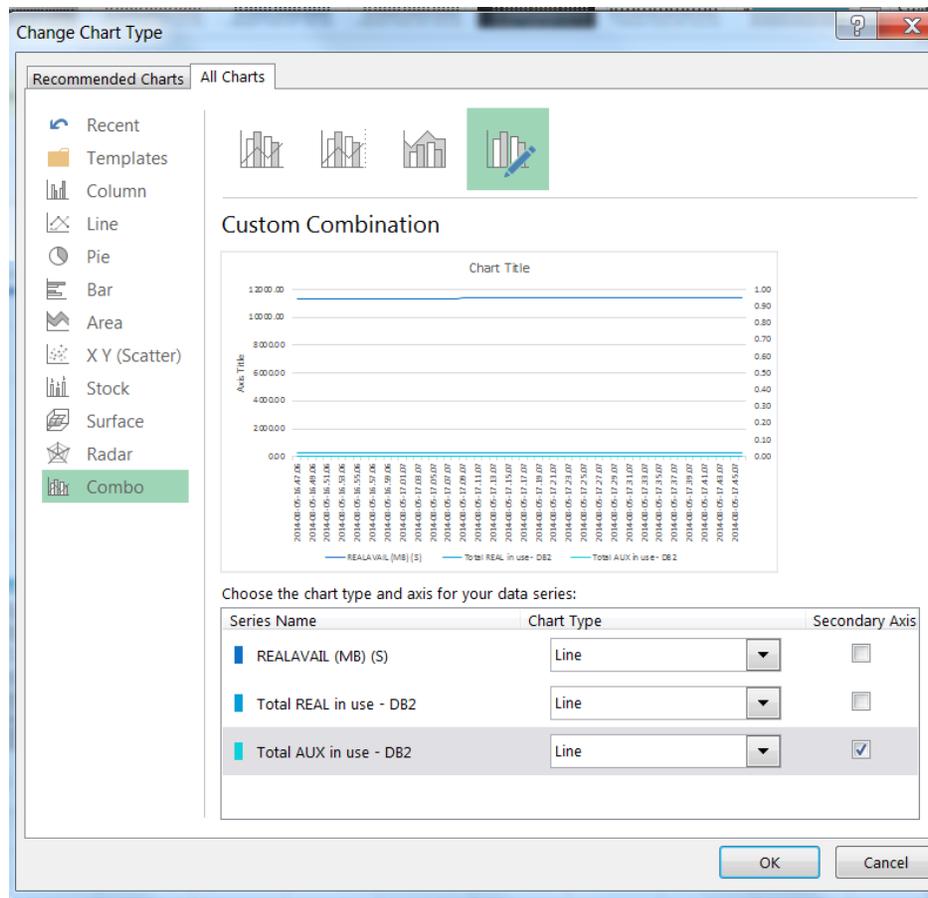
Real Available Chart Modifications

From Line chart, change chart type to “Combo” and “Stacked”

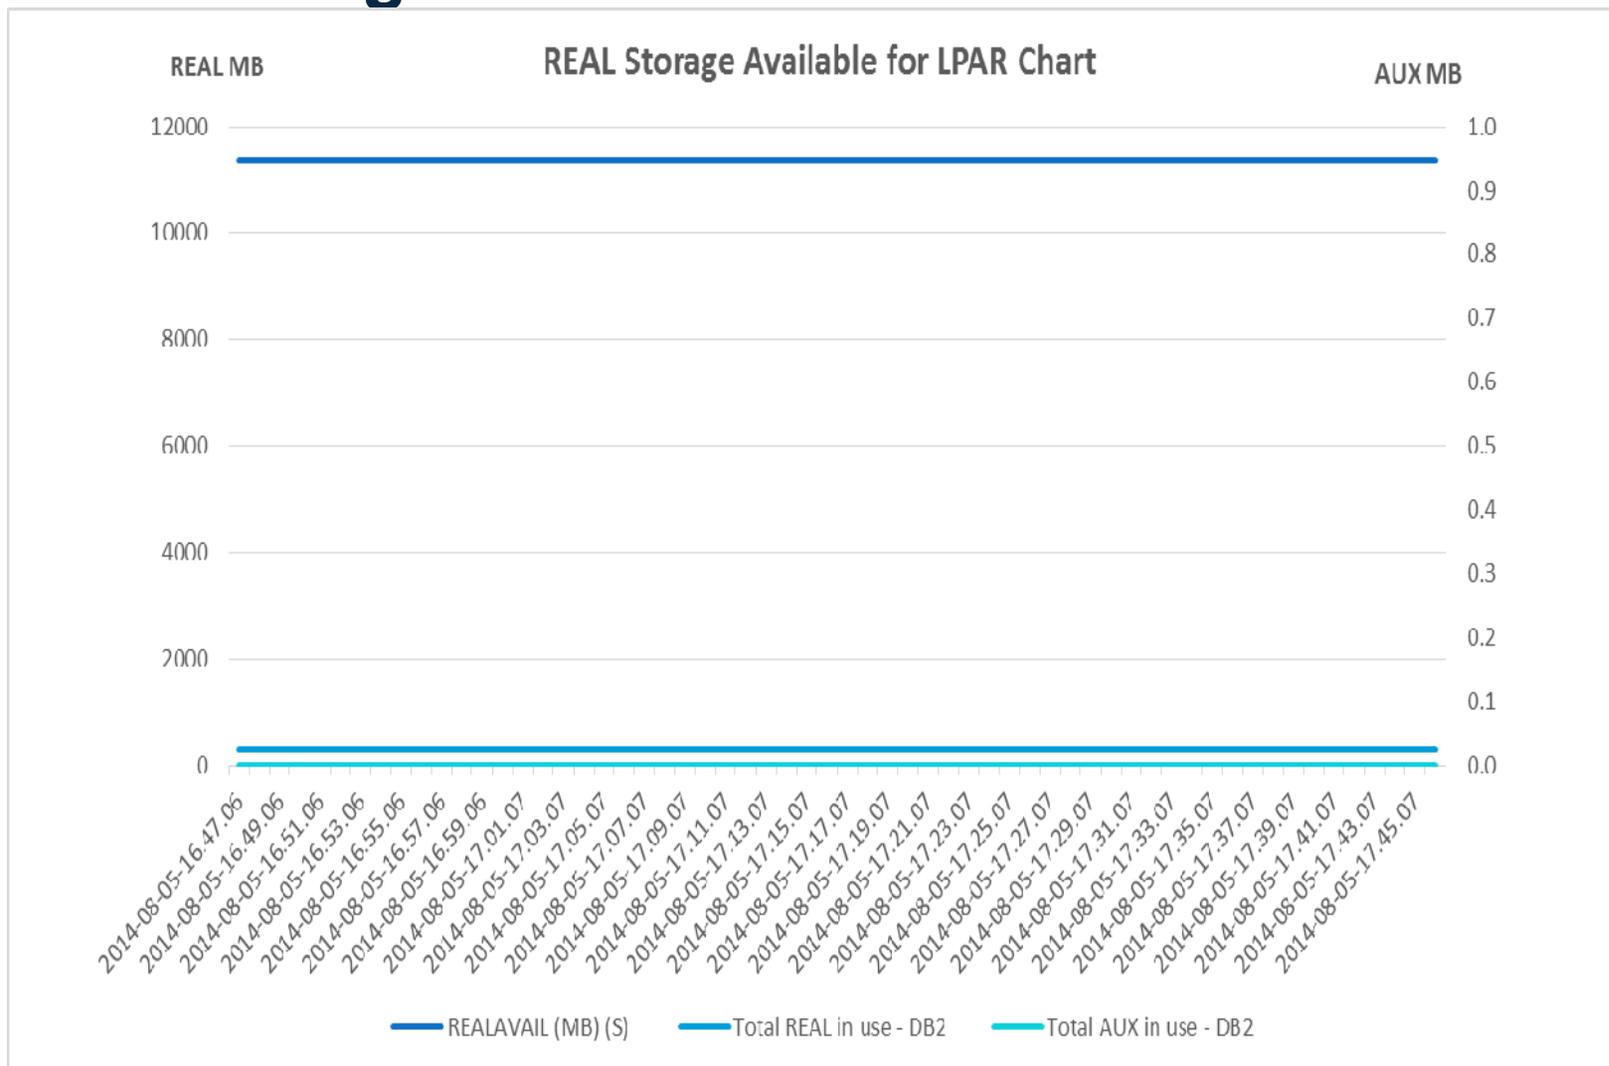
These are LINE columns:
C, DB, GC, GD

Create secondary axis
(right hand side)

- Change Total AUX in use to “line” and check “secondary axis”
- Make sure the other two columns are “line” data types



Real Storage Available for LPAR



Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Use of Real Available Chart

- The DBM1 31-bit Private chart gives you the maximum threads you can run for this workload based on virtual storage constraints
- Do you have enough real storage to support that virtual storage?
 - That spreadsheet only reports on the 31-bit footprint
 - Roughly 90% of the footprint is above the bar now

Real storage

- Additional real storage is the cheapest and easiest performance upgrade
 - Real storage shortage not only can cause performance issues but if DUMPs are needed then it can cause a small issue to become a SYSPLEX failure
 - Cheapest, as MLC and other charges are not based on amount of real storage
 - Vendors do not charge by the amount of real on the CEC/CPC processor
 - zSystem 13 real storage is very attractively priced
- Specify z/OS WLM STORAGE CRITICAL for DB2 system address spaces
 - Tells WLM to not page these address spaces
 - Keeps the thread control blocks, EDM and other needed parts of DB2 in REAL
 - Prevents the slow response that occurs if DB2 has to be rapidly paged back in
 - If BPs are page fixed, it safeguards rest of DB2
 - Other workload will page if real storage shortage

When the Cupboard is Bare...

- The example on next slide shows in < 24 hours
 - Real storage in Use on the LPAR is about 15G
 - Real in Use for DB2 is 6G
 - Real Available is < 2G
 - Illustrates that as real storage dips (blue), the AUX increases (purple)
 - Real storage on the left vertical axis
 - AUX for DB2 on the right vertical axis (secondary axis) 914-928M

Example of Insufficient Real Storage

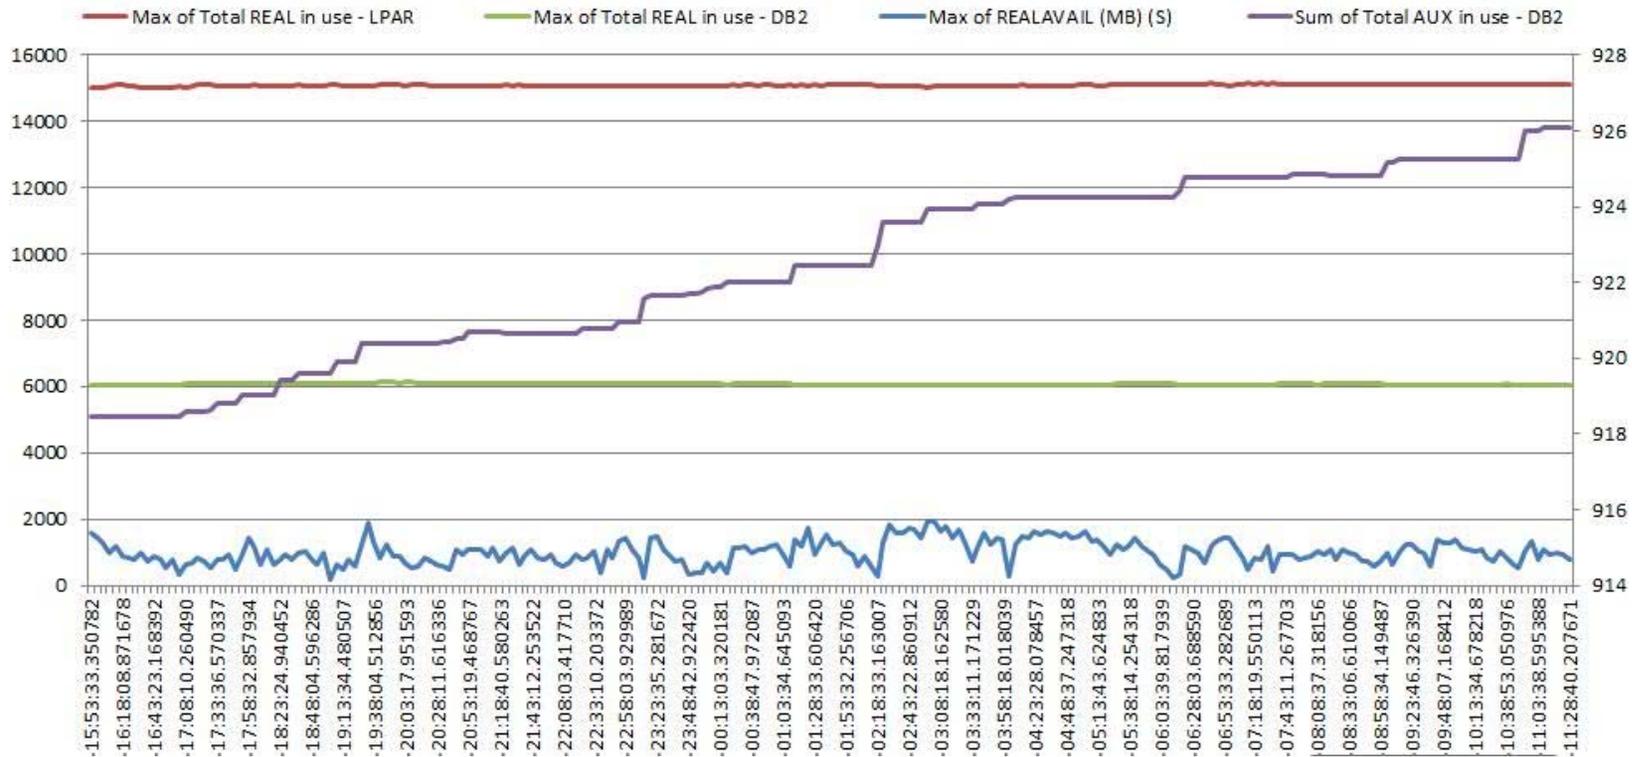


Chart (different from Real Available Chart) provided only to show egregious lack of real storage)

Formula for red line not validated as of 9/2014

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Summary

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

General Guidelines for Charts

- Run MEMU2 from DB2 start up
- Ignore storage on the first day (while DB2 settles in)
- Throw away abnormally low values for Current Active Threads (especially after an interval of high values).
 - If #users is small, footprint is abnormally large as current threads and max DBATs active are the divisor.
 - Max # Threads will be lower.
- Pay attention to month end or other known peaks, weekends where REORG and heavy I/C can increase VSTOR stress.
- Look for peaks after DB2 has been up for awhile. That gives a good feeling that the Max can still be maintained.

General Guidelines for Charts

- The most accurate storage is obtained as you approach the maximum number of active threads, because DBM1 storage is a more accurate reflection of capacity. The converse is also true!
- Apportion chosen maximum thread value between CTHREAD and MAXDBAT.
- NOTE: Assumption is charted workload does not change. Changes mean that these measurements must be redone.
- Provision enough real storage to support threads
 - Avoid paging
 - Reserve enough for fast DUMP (16G/prod DB2/LPAR)

Summary

- MEMU2 shows how many times critical storage contraction occurred
- Use several common CTHREAD values for cloned, diverse workloads?
 - 2 Batch members, 3 online members, 2 distributed members
 - A common CTHREAD value for all can leave a lot of capacity on the table
 - ✓ Use one CTHREAD for each of 3 workload types to provide failover
- ✓ To arrive at IDFORE/IDBACK and validate CTHREAD/MAXDBAT, run OMPE Statistics Trace before shutdown. It displays HWM for life of a DB2.
 - IDFORE
 - IDBACK
 - CTHREAD
 - ACTIVE DBATS
 - Ignore Thread Footprint if system is already drained

QR Code



Complete your session evaluations online at www.SHARE.org/Seattle-Eval