Cheryl Watson's Tuning Letter

2013 No. 2

I hope you're enjoying our new quarterly issues. I'm thrilled because I can cover several topics in full in a single issue. The most exciting things to me in this issue are the WAS Liberty Profile, the new IBM System z Batch Network Analyzer (zBNA) tool (free!), and the large number of SHARE requirements that are satisfied in z/OS 2.1. In August I'll be heading to Boston for the SHARE conference that starts on August 12, and am looking forward to all of the sessions on z/OS 2.1. This has got to be the largest release ever, and reports from early testers are very positive. Our next *Tuning Letter* will have more in depth on 2.1, along with more *User Experiences* that I didn't have room for in this issue. Hope to see some of you in Boston!

Management Summary

Note: If you're a technical manager, this section is for you! If you're not, please read and pass this on to management. Here we try to point out things the manager should know or that the staff should explore. Please send questions to technical@watsonwalker.com.

z/OS 2.1 Preview

It's been two years since the last release of z/OS, and the upcoming (September) release of z/OS 2.1 looks to be the largest release of z/OS we've seen since z/OS 1.4. We'll have many more details in our next *Tuning Letter*, but the article in this issue will give you a hint of the scope of the release. In our last *Tuning Letter*, Cheryl described the many enhancements of 2.1 that can be implemented in your current releases by just applying some APARs. In this issue, she shows you dozens of features in 2.1 that are a direct result of customer requests, specifically through the mecha-nism of SHARE requirements. We've always been supporters of the SHARE conferences, and the response from IBM to requests from SHARE requirements shows that support to be well-founded.

Tom Walker, Publisher

Exploiting z/OS

We continue our series about exploiting enhancements that are already in the z/OS releases that you're running today by going into the benefits and the effort required for two z/OS features - Base Control Program internal interface (BCPii) and Extended Address Volumes (EAVs), two very different facilities in terms of complexity and use. EAVs are complex to implement, and are of primary interest to very large installations. BCPii, however, is simple to install, simple to use, and can provide many benefits to shops of all sizes. Cheryl's guidelines and extensive references will point you in the right direction.

Other Important Topics

WebSphere Liberty Profile is a significant enhancement to WebSphere Application Server (WAS) providing both performance improvements and a much smaller foot-print. But the most important benefit, in Cheryl's opinion, is its use in z/OSMF. z/OSMF is the z/OS Monitoring Facility that Cheryl is so enthusiastic about. It's a tool for sys-tem programmers that is designed to reduce the time to handle many typical tasks. It gets better with each release, but the main drawback has been its dependence on a WAS OEM (a special WAS for use just by z/OSMF). The problem was that WAS OEM was so large and slow and took so much DASD space that nobody could find a test LPAR to even test it on. In z/OS 2.1, z/OSMF will be using the Liberty Profile and it will be usable in every installation.

Another important item in this issue is the description of a new and free tool from the IBM Washington Systems Center called the IBM System z Batch Network Analyzer (zBNA). Because more installations are moving to more but smaller CPs, batch win-dows may start to elongate. zBNA will help you model the effect on your batch window

due to changing processors. It's one of those tools that you really need to use before a processor upgrade.

If management is considering moving from a mainframe to a distributed environment, we point out some important reading and opportunities that should be considered before making that move. Here are links:

> IBM Redbooks (www.redbook.ibm.com) - REDP-5032-00 (21Jun2013) - *The Reality of Rehosting: Understanding the Value of Your Mainframe.* Of special importance is the description of the IBM Eagle team who can perform a free TCO assessment of your installation.
> SHARE President's Corner Blog (www.share.org/p/bl/et/blogid=2) has a 3-part series called *Don't Believe the Myth-information about the Mainframe.*

Elsewhere in This Issue

You'll also find many other useful items throughout this newsletter: User experiences with increased CPU time in DB2 and DB2 zIIP eligible time • A description of the SMF Type 42, subtype 6 record for tuning data sets • Some very important IBM announcements, including a major update to COBOL • A continuing section on Social Media • Some important parmlib changes due to APARs • and our usual collection of New Function, SMF, Information, and HIPER APARs to help you identify useful maintenance. ∎

Cheryl Watson's Tuning Letter

2013 Number 2 (Vol. 23, No. 2)

Note: Implementation of any suggestions contained in this newsletter should be preceded by a controlled test and is the responsibility of the reader.

WATSON WALKER

Exploiting z/OS - Part 3

This continues our series on exploiting z/OS. As I've said before, IBM includes many enhancements in new releases and in New Function APARs, but the developers are restricted in automatically activating the enhancements because they are required to maintain the previous behavior of the system. This is to protect customers from a change in processing that might impact their production workloads. So almost every enhancement is turned off until YOU activate it. And most installations never find time to activate these new features. Ever since the beginning of our *Tuning Letter* in 1991, we've tried to explain the reason for and how (or whether) to implement each en-hancement.

Instead of simply a list, this series will explore each of the enhancements in more de-tail and will cover the following elements:

> ➤ Overview - a brief summary of the facility
> ➤ What environment is needed (hardware and software)?
> ➤ What benefit can you expect to achieve?
> ➤ User experiences
> ➤ Implementation considerations
> ➤ How do you measure the impact?
> ➤ Which are the best references?

One of the last questions on the recent SHARE MVSE survey asked the responder to list and rank the five enhancements that provided the most benefit. We started in Part 1 with the IBM Health Checker, HyperPAV, and HiperDispatch. Then, in Part 2, we cov-ered the next three most beneficial enhancements: zIIPs/zAAPs, zFS, and zHPF. In this issue we'll cover the next enhancements: BCPii (Basic Control Program internal inter-face) and EAVs (Extended Address Volumes).

Small side-note: We were asked about the change in format of the References section in this series, which changed from showing the author first to showing the source first starting in Tuning Letter 2012 No. 6. *The reason for the change is that if you go to a location, such as Redbooks, Techdocs, or SHARE, it's easier to locate all the references to that one source. Thus we list them together in the References section.*

BCPii

The Base Control Program internal interface (BCPii) runs in an address space called HWIBCPII and lets authorized programs use APIs (Application Program Interfaces) to query, modify, and perform Hardware Management Control (HMC)-like functions for any processors on the HMC network. It does this internally, and so does not need to use the IP network on the system, thus providing more security. BCPii was introduced in 2008 with z/OS 1.10. 5% of the survey replies indicated that BCPii provided some of the best benefits of all of the z/OS enhancements.

Overview

BCPii provides access to HMC functions via authorized programs. Figure 1, from SHARE session 8665 (see References), shows a typical configuration.

Figure 1 - What is BCPii? (© IBM)
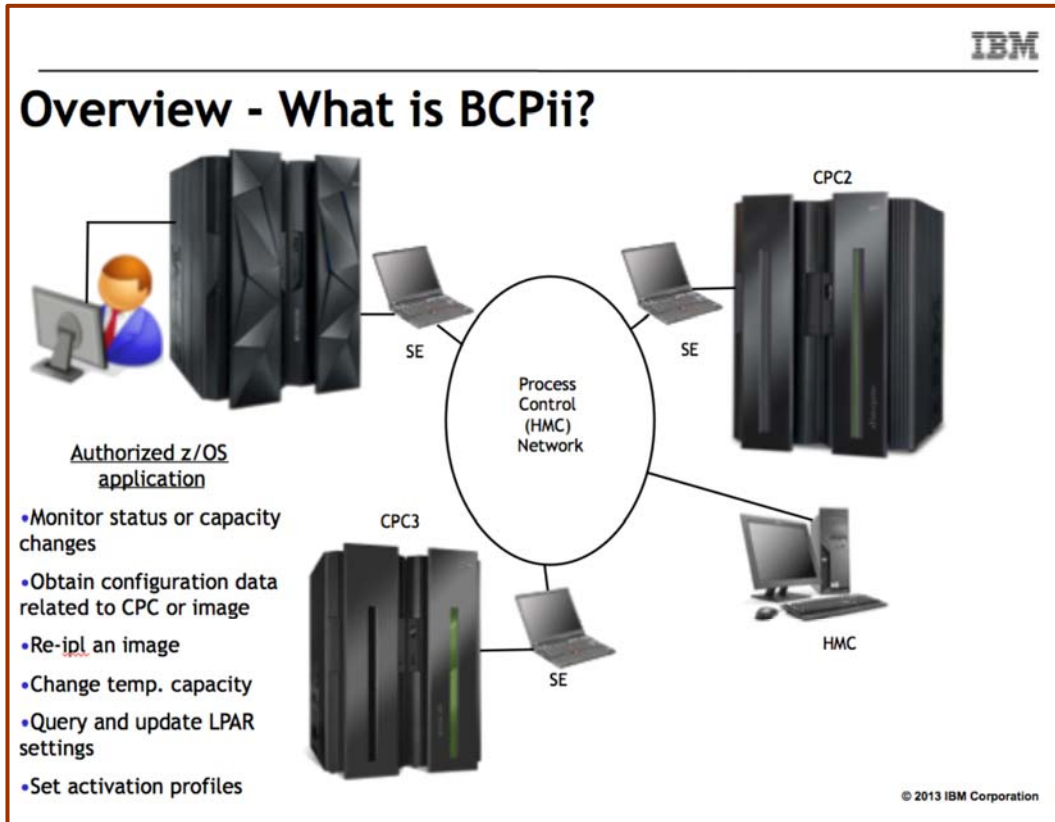
The APIs (Application Program Interfaces) that are available to these programs are:

- HWILIST - BCPii List - List CPCs interconnected with local CPC, images (LPARs), capacity records, events registered, local CPC, local image, reset activation profiles, user-defined image groups.
- HWICONN - BCPii Connect - Establish connection between CPCs, images, capacity records, activation profiles, and user-defined image groups.
- HWIDISC - BCPii Disconnect - Remove previously established connection.
- HWIQUERY - BCPii Query - Retrieve information about CPCs, images, capacity records, activation profiles, and user-defined image groups.
- HWISET - BCPii Set
- HWICMD - BCPii Command
- HWIEVENT - BCPii Event (for non-z/OS Unix callers)
- HwiBeginEventDelivery, HwiEndEventDelivery, HwiManageEvents, HwiGetEvent (for z/OS Unix callers)

The following are the types of things you can query using BCPii. (This list came from IBM ATS presentation TC000050 in *References* below and is updated from other documents.)

- CPC Information
  - General Information (name, serial, machine type, id, networking info)
  - Status information (operating status and other status values)
  - Capacity Information (capacity backup (CBU) info, capacity on demand info, processor configuration including zAAPs, IFLs, ICFs, and zIIPs)
  - Power Savings Information (z/OS 1.13 or z/OS 1.10-1.12 with APAR OA30558 / OA34001)
- User-Defined Image Group (z/OS 1.13 on System z10 and above)
  - List the image names in the group
  - Connect and disconnect from the group
  - Query values associated with the group
- Image Information
  - General Information (name, OS info)
  - Status Information (operating status and other status values)
  - Capacity Information (defined capacity, processor LPAR weights)
- Capacity Record Information
  - General Information (name, activation and expiration dates, activation days)
  - Status Information (record status)
  - Capacity Information (entire capacity record)
- Application Profile Information (most activation profile values)

> BCPii leads
> to enhanced
> availability

The following are the types of things you can modify using BCPii. (This list also came from IBM ATS presentation TC000050 in *References* below with updates.)

- CPC Information
  - Acceptable status values
  - Next Reset activation profile name
  - Processor running time
- User-Defined Image Group (z/OS 1.13 on System z10 and above)
  - Connect and disconnect from the group
  - Issue commands against all members in the group
  - All of the commands listed below under Image Commands
- Image Information
  - Various processor weights
- Activation Profile Information (any activation profile value)
- CPC Commands
  - Activate/Deactivate an entire CPC
  - CBU request (activate or undo)
  - On/Off Capacity on Demand request (activate or undo)
- Image Commands
  - SysReset, SysReset with IPL Token
  - Load
  - Start, Stop all CPs
  - Add or remove temporary capacity
  - Issue operating system command

Applications can listen for events such as command completions, status changes, capacity changes, disabled waits, BCPii status changes, BCPii communications errors, and power mode changes. BCPii uses ENF 68.

Here is a summary of BCPii enhancements in each release:

- ➢ z/OS 1.10 - Base functions (no HWISET)
- ➢ z/OS 1.11 - HWISET, IPL Token / Query PSWs, activation profiles support, minor internal serviceability enhancements
- ➢ z/OS 1.12 - CTRACE enhancements, improved storage utilization and serviceability of BCPii transport code, additional CPC/Image attributes and commands
- ➢ z/OS 1.13 - Support for user-defined image groups, additional CPC/Image attributes, new STP commands
- ➢ z/OS 2.1 - Support for REXX callers, improved performance for HWIQUERY and HWILIST

> You should DEFINITELY implement BCPii!

Environment

BCPii can run on any System z processor. It first became available on z/OS 1.11, but was rolled back to z/OS 1.10 with APAR OA25426. Programs using the APIs must be program-authorized and SAF-authorized. Initial support was available for C and Assembler languages. Unix System Services callers can get event notifications using the Common Event Adapter (CEA). z/OS 2.1 will provide support for REXX callers.

Benefit

BCPii can be used by IBM products, ISV products, and customers. Rather than a performance improvement, BCPii usually provides for stable systems. Here are some of the IBM facilities that currently use BCPii:

- ➢ Capacity Provisioning Manager (CPM) can add or delete temporary capacity based on the WLM policy. First available in z/OS 1.10 with APAR OA24945.
- ➢ Prior to BCPii, Sysplex Failure Management (SFM) (a sub-component of XCF) could remove a system if the system became unresponsive, but it couldn't determine WHY the system was unresponsive. It could issue a coupling facility ISOLATE to remove the unresponsive system from the channel subsystem, and then tell the operator to do a manual reset. Using BCPii, XCF System Status Detection (SSD) Partitioning Protocol (SYSSTATDETECT) can determine whether a system is truly dead (defined as entering a non-restartable waitstate, being reset, being reloaded, being deactivated, or being check stopped) or has been reset and can expedite partitioning during a system failure. First available in z/OS 1.11. We'll cover this more in our next Tuning Letter. An updated message is available in APAR OA40546 (z/OS 1.13, 31Oct2012, *XCF to Issue Additional IXC114I Message to Display More Information When Taking Action Via SSD*).
- ➢ Hardware Configuration Definition (HCD).

Other ISVs are using BCPii in several situations. Here are just a few that I noticed:

- ➤ CA's OPS/MVS Hardware Services (HWS) uses BCPii to show event changes (available now)
- ➤ CA's OPS/MVS Hardware Services (HWS) will use BCPii for event management and automation (in development)
- ➤ zCOST Management uses BCPii in their AutoSoftCapping product.
- ➤ BMC's MainView AutoOPERATOR uses BCPii.

User Experiences

BCPii isn't necessarily one of those enhancements that everybody raves about. They just set it up, it runs, other products exploit it, and the system is better off for it. You get a lot of enhancements under the covers of several products. I've heard of no problems with it, except that most people don't know what it's used for. We hope this article answers that question. I've heard of a few customers who are exploiting it themselves, such as in this next item.

The following is an extract from our *Tuning Letter 2010 No. 3* (page 8):

- ■ **Dave Quinton** posted the following item on IBM-Main (4Jun2010, *BCPii Services)* regarding BCPii:

  We used the following justifications for implementing BCPii (and the homegrown code to go with it):

  1. *We do not have HMC access via VPN. BCPii allows us to do selected HMC functions from home in an emergency.*
  2. *On a number of occasions, the wrong LPAR has been IPLed via the HMC. Our BCPii code will not allow an IPL of an active LPAR by default.*
  3. *Having to log on to the HMC to IPL a test LPAR is no longer required.*
  4. *We don't have to deal with the nuisance that comes with having to re-enter your HMC password (and remember to click YES) when doing anything disruptive on the HMC.*
  5. *Our operators love being able to IPL an LPAR from another LPAR with a simple operator command.*
  6. *We don't have to remember IPL addresses and parms - our home-grown code manages all of that for us.*
  7. *It was a fun learning experience for me…*

  I can't wait to get z/OS 1.11 up and running to see what the new BCPii code will let us do - I'm hoping to allow our capacity guy to automate changing soft-capping :-)

  I know there have been a number of APARs against BCPii, but we've experienced no problems using it.

Steve Warren, one of the BCPii developers added this useful information:

> *We were the most surprised of all when we had customer after customer who was not just interested in installing BCPii to exploit a z/OS operating system function or use ISV vendor applications like the ones you listed, but rather to actually write their own BCPii applications. To date, we know of hundreds of customers (including many high-profile ones) that have written their own BCPii applications to do a variety of different tasks. They have automated some aspects of their disaster recovery, automated con-figuration of their systems, automated the changing of processor weights to push through critical work on their system, used BCPii as a way to gen-erate automated reports of their configuration across their HMC network, and of course many other uses.*

The following APARs have been found to be useful:

OA38252 (z/OS 1.11-1.13, 21Mar2012) - *z/OS BCPii HWIEVENT Version HWICIC.*

OA35338 (z/OS 1.10-1.13, 28Mar2012, *BCPII ABEND042 RSN2E HWIAMCAT*) / OA40349 (z/OS 1.10-1.13, **PE**, 17Apr2013, *z/OS BCPii HWISMSTK WAIT ABENDA03 IXC336I*). The first APAR corrected a problem when more than 3500 bytes were returned to BCPii, but could cause a lost BCPii connection. Therefore, you need the second APAR applied.

OA39515 (z/OS 1.13, 4Jul2012) - *z/OS BCPii User-Defined Image Group Support Restriction.* This documents that the user-defined image group is not support-ed on a z9 processor.

OA40017 (z/OS 1.12-1.13, 29Aug2012) - *z/OS BCPii Sample Program ENF Exits.*

OA40067 (z/OS 1.12, **HIPER**, 29Aug2012) - *z/OS BCPii ABEND0C4 HWITMMNX.*

OA40392 (z/OS 1.10-1.13, **New Function**, 17Oct2012) - *z/OS BCPii Toleration Support zEC12.*

OA40884 (z/OS 1.11-1.13, **HIPER**, 12Dec2012) - *BCPii Returns a Zero Return Code When a Non-Zero Should Have Been Returned.*

OA41107 (z/OS 1.13, 13Feb2013) - *z/OS BCPii HWIGROUPPROFILECAPACITY HWICIC HWI001.* z/OS 1.13 added a new facility, the HWIGROUPPROFILECAPACITY, but it won't work because there's a spelling er-ror in SYS1.SIEAHDRV.H(WHICIC).

OA41427 (**DOC**, 17Jun2013) - *z/OS BCPii HWI014I HWI022I.* This enhances a message that might be received during setup of BCPii.

OA42165 (**OPEN**, 6May2013) - *Protect BCPii Address Space Deemed Critical to Hyperswap Operation From Page Faults.* Until this APAR has a fix, anyone us-ing BCPii and Hyperswap, should mark BCPii as Critical in the PPT and enable CRITICALPAGING on the COUPLExx parmlib member. IBM's recommendation is that you enable CRITICALPAGING only when you need it because it might cause excessive paging in other address spaces. See our *Tuning Letters 2011 No. 5* (page 14) and *2011 No. 2* (page 16).

Implementation

There are four basic steps to setting up BCPii. These are all described in the IBM manual and IBM Redbooks under *References.*

1. In the HMC you will need to enable Cross Partition Authority for each LPAR that you want BCPii to be able to see or each LPAR that is running BCPii. You also need to define BCPii community names on the SE and enable SNMP communications.
2. You then configure the BCPii address space by making sure a few data sets are in the link list concatenation and optionally changing CTRACE options.
3. Authorize applications to use BCPii.
4. (Optionally) Create the event notifications for z/OS Unix callers.

References

Cheryl Watson's Tuning Letter 2010 No. 3, page 8 - *User Experiences.*

IBM ATS Conference Presentation TC000050 (Spring2010) - *Parallel Sysplex Partitioning Using BCPii.*

IBM Manual SA22-7613-10 - *z/OS MVS Programming: Callable Services for High Level Languages.*

IBM Redbook SG24-7817-00 - *System z Parallel Sysplex Best Practices.*

IBM Redbook SG24-7946-00 (27Mar2012) - *z/OS Version 1 Release 13 Implementation.* This contains an entire chapter on BCPii and how to install it and use it. Similar Redbooks are available for z/OS 1.11 and 1.12.

IBM z/OS Hot Topics Newsletter - August 2009 - **Stephen Warren -** *The application doesn't fall far from the tree BCPii: Control your HMC and support element directly from z/OS apps.*

IBM z/OS Hot Topics Newsletter - August 2012 - **Stephen Warren -** *Seeing BCPii with new eyes.*

SHARE 2009 in Denver - Session 2209 - **Brian Peterson**, **Ed Jaffe**, **Skip Robinson,** and **Sam Knutson** - *Bit Bucket x'26'.*

SHARE 2010 in Boston - Session 7780 - **Horst Sinram** - *Introduction to z/OS Capacity Provisioning.*

SHARE 2010 in Seattle - Session 2251 - **Mark Brooks** and **Riaz Ahmad** - *Sysplex Partitioning Using BCPii.*

SHARE 2010 in Seattle - Session 2227 - **Steve Warren** - *BCPii: Secure z/OS Interface to Your HMC and SE.*

SHARE 2011 in Anaheim - Session 8665 -**Steve Warren** - *BCPii for Dummies: Start to finish installation, setup and usage.*

SHARE 2011 in Orlando - Session 9704 - **Mark Brooks** and **Nicole Fagen** - *Parallel Sysplex Resiliency.*

SHARE 2011 in Orlando - Session 9865 - **Steve Warren** - *Simple BCPii Programming for the z/OS System Programmer.*

SHARE 2012 in Anaheim - Session 12088 - **Brian Valentine** - *IBM System z HMC (Hardware management Console) Security Basics & Best Practices.*

SHARE 2012 in Anaheim - Session 11806 - **Frank Kyne** - *Recent z/OS Enhancements You Can Use to Reduce Down Time.*

SHARE 2013 in San Francisco - Session 12504 - **Mike Shorkend** - *Back to the Future: Creating Consistent Copies at Isracard.* (This is an interesting user expe-

rience where BCPii was used with some other products to create a DR solution.)

SHARE 2013 in San Francisco - [Session 13035](#) - **Steve Warren** - *BCPii Programming Beyond the Basics for the z/OS System Programmer.*

I'm not sure how to list this reference, but I want to make sure you see it. An IBM Information developer, **Rita Beisel**, blogs at IBM Mainframe Insights. Here is a link to her latest entry from June 19 called *[News Flash: z/OS BCPii summary tables wanted!](#)*. This also contains pointers to other BCPii blogs, with an impressive list of BCPii blogs to come from other BCPii team members such as **Karen Ransom**, **Rebecca Horner**, and **Steve Warren**.

Recommendation

==There shouldn't be any System z installation that doesn't have BCPii installed. It's a very simple process, and several products can exploit the function to provide more stability and reliability to your z/OS system. Some products will automatically use BCPii if it is installed, while others products need some type of activation before using BCPii. We highly recommend that you install this in your test system.==

EAVs

Extended Address Volumes (EAVs) provide a method to allow the amount of addressable DASD storage per volume to grow beyond traditional sizes by changing how tracks on ECKD volumes are addressed allowing capacity of up to 1 TB per volume. This allows z/OS operating systems which are limited to 4-digit device numbers (65,280 total number of devices) to grow its total addressable DASD storage per system to satisfy future DASD storage requirements or by simplifying storage management because fewer volumes for the same capacity can be used.

> EAVs benefit mostly large sites

Environment

EAVs first became available in 2008 in z/OS 1.10 with a microcode update (Release 4.0) to an IBM DS8000 storage controller. z/OS 1.13 (and z/OS 1.12 with APAR [OA28553](#)) provide 1 TB EAVs, as well as support for the DS8700, and additional non-VSAM data sets (see next paragraph). As you will see, EAVs will be most beneficial in larger installations.

Overview

Prior to EAVs, the largest volume you could have contained 65,520 cylinders (54 GB). An extended address volume (EAV) is, by definition, any volume greater than 65,520 cylinders. The space at cylinder address 65,520 and above is called cylinder-managed space. The extended addressing space (EAS) are those cylinders in cylinder-managed space that are represented by 28-bit cylinder numbers (explained later) compared to the others that are represented by 16-bit cylinder numbers. The first EAVs on z/OS

1.10 could be up to 223 GB (262,668 cylinders), but that maximum was increased to 1 TB in z/OS 1.13 (and z/OS 1.12 with APAR OA28553). Large format sequential data sets (DSNTYPE=LARGE) are limited to a maximum of 16,777,215 tracks, even though the volumes are larger.

- z/OS 1.10 - The first exploiters of EAS were VSAM (KSDS, RRDS, ESDS, and Linear) data sets used by DB2 V8+, CICS, zFS, IMS V9+, NFS, and SMP/E CSI, as well as DFSMSdss DEFRAG. CA sizes for VSAM can be 1, 3, 5, 6, 9, or 15 tracks in size. Control of which VSAM data sets can use EAS is done by use of SMS storage groups. Non-SMS data sets can be controlled by using a specific VOLSER or esoteric name. DFSMShsm recommends that you NOT use EAVs for ML1, ML2, and backup volumes.
- z/OS 1.11 provided support for sequential extended format data sets and XRC journal data sets. z/OS 1.11 also added a new EATTR data set attribute that can be specified in the AMS DEFINE CLUSTER, AMS ALLOCATE, JCL, dynamic allocation, data class, and ISPF. This attribute can override the system-determined EAS eligibility. DFSMShsm recommends that you NOT use EAVs for ML1, ML2, and backup volumes.
- z/OS 1.12 provided support for non-VSAM data sets of sequential (basic and large format), PDS, PDSE, BDAM, data sets with undefined DSORGs, XRC state data sets, and catalog VVDS and BCS data sets. DFSMShsm has a new option (SETSYS USECYLINDERMANAGEDSPACE) to allow EAS for migration copies and/or backup versions (ML1/ML2). DFSMShsm can use EAVs for its Journal, Control data sets, Logging, and PDS data sets. DFSORT can use EAS for SORTIN, SORTOUT, OUTFIL, and SORTWK (only if large format sequential).

EAS cannot contain VTOCs, VTOC indices, page data sets, VSAM data sets with imbed or keyrange attributes, HFS file system, XRC (Control, Master, or Cluster) non-VSAM data sets, or VSAM data sets with incompatible control area sizes.

The space within the first 65,520 cylinders is managed by tracks and is called the 'Track-managed space', while the space above that is managed by increments of

Sizes of current 3390 DASD volumes
3390-3 - 3 GB (3,339 cylinders)
3390-9 - 9 GB (10,017 cylinders)
3390-9 - 27 GB (32,760 cylinders)
3390-9 - 54 GB (65,520 or 64K cylinders)
3390-A - (z/OS 1.10) 223 GB (262,668 or 256K cylinders)
3390-A - (z/OS 1.13 or z/OS 1.12 with SPE) 1 TB (1,182,006 cylinders)
3390-A - (theoretical max) hundreds of TBs (1 to 268,434,453 cylinders)

cylinders (current default is 21 cylinders) and is called 'Cylinder-managed space'. This space is also referred to as EAS or 'Extended Addressing Space'. A new mapping scheme is used (CCCCcccHRR as compared to CCCHHHHRR for below 65,520), and new DSCBs are used (Format 8 and 9). A Break Point Value (BPV) is used to determine the size of data sets that should reside in EAS. The default setting for BPV is 10 cylinders. So any data set that is 10 cylinders or more will be allocated in EAS, and anything smaller will be placed in the track-managed space.

According to an IBM survey, two-thirds of their customers are on 3 GB or 9 GB volumes and most are transitioning to 54 GB devices. The track format, track length, and number of tracks per cylinder are the same for all 3390 volumes.

Benefits

The primary benefits of EAV are the relief from the 4-digit device address constraint and the reduction of DASD storage management due to the reduction of the number of devices that need to be managed.

Additional benefits include:

> ➢ Dynamic Volume Expansion allows you to expand a volume without needing to copy it.
> ➢ HyperPAV feature provides dynamic scaling for I/O rates across an LSS and reduces the number of PAV aliases needed.
> ➢ DFSMSdss DEFRAG is redesigned to provide reasonable times on larger volumes. While not directly a benefit, this enhancement was needed to keep EAVs viable.
> ➢ As users of zHPF are aware, the primary benefit of zHPF is to save channels with low utilization and applications that can take advantage of zHPF. In addition to reducing the number of channels, zHPF can result in fewer fiber, switch ports, and control unit ports. It also provides additional benefit to extended address volumes (EAVs) by increasing I/O rates as the volume sizes expand. Because of the higher I/O rate, some applications will see reduced response times. This is especially true with DB2 activity.
> ➢ Starting in z/OS 1.12, JES3 can place SPOOL, checkpoint, and Job Control Table (JCT) data sets on EAVs, thus allowing large SPOOL volumes.
> ➢ Starting in z/OS 1.12, JES2 can place SPOOL and checkpoint data sets on EAVs, thus allowing large SPOOL volumes.

User Experiences

In the SHARE MVSE program survey, only 30% of the responders indicate they had tried EAVs. The majority of non-users indicated that they did not see any benefit in the feature. Many others did not have the hardware to support it. A few people thought that only IBM supported EAVs, but today all of the big storage vendors support it.

In our last *Tuning Letter* (2013 No. 1, page 42-43), we said that **Ed Jaffe** of Phoenix Software had been having a problem with EAVs since 2009. It was tracked back to a situation when using DEFRAG on extended address volumes, and HIPER APAR OA40210 was created. We don't know whether others did not notice the problem because they weren't using DEFRAG or because they weren't using EAVs.

We have asked our readers about their experiences with EAVs, but received only a couple of responses. One of them said that it took a while to change procedures, but it was very beneficial to them.

Another replied:

> We are using EAVs in production and are undergoing a large data conversion project for specific DB2 systems. We've been working on this project for 2+ years and it seems to be going well at this point. We have a very controlled target environment for EAVs, which includes a small tool set and standard utilities. We've worked with IBM EAV Development as well as a number of ISVs to clear up known issues. We did wait for Hyper OA40210 PTFs to be installed before venturing beyond development.

Dave Crow from SAS Institute provided the following comments after going through the process of updating SAS to support EAVs:

> The multi-release approach that IBM used in providing EAV support not only allowed customers to benefit from EAV sooner, it enabled us to get full support of EAV into SAS within a few months of when the support was available in z/OS.
>
> I've also been surprised at how slowly customers have moved to it, and although we haven't seen the level of problems Ed had, just last week we hit an EAV bug in IBM's LISTDSI. IBM responded very quickly, and this is the first problem we've seen in years.
>
> Some thoughts about what it takes to use them:
> - When we first started working on EAV support we were a little surprised by the interaction between EAV and dynamic volume expansion. If the first thing a customer does with a new disk controller is to allocate all the space to pre-defined volumes, then to create an EAV you first have to delete some volumes to free up space.
> - Creating an EAV gets into needing changes to operational procedures that have probably been in place for decades.
> - For us, we're planning that any new volumes will be EAV, but we haven't done anything to combine existing volumes into a single EAV, so even though we've made a lot of use of EAV, the vast majority of our volumes are not EAV.

Implementation Considerations

Implementation of EAVs can be a challenge for some installations. So you want to make sure that you need the facility before you attempt using them. Not only do you need the correct driver level on the DS8000/DS8700 (driver 7.6.2.x), but you need to determine your procedures and policies relating to EAV usage. You need to review the required APARs and fixes for both IBM and ISVs.

The steps of implementation below are expanded in SHARE presentation 3204 (in *References* below). This should be required reading before even thinking of implementing EAVs.

- Review:
  - IBM publications - start with the z/OS 1.13 (SC26-7473) - *DFSMS Using the New Functions*.
  - Dependencies
  - Coexistence consideration
  - Migration considerations introduced with EAV support
  - Coexistence considerations introduced with EAV support
  - Migration to EAV considerations
  - EAV migration tracking facility output (this identifies applications that might fail when data sets are on EAVs)
- Configure EAV on your DS8000
- Add EAV to your storage group/pools
- Enable the use of EAV in your system (IGDSMSxx PARMLIB member, change the default USEEAV(NO) to YES)
- Migrate data

> **Important step - use EAV migration tracker**

This is NOT a simple implementation, and you should allow plenty of time to think it through before starting on the implementation. I would also check with all of your software vendors, including IBM. As recent as June 2013, EAS support was **just added** to CICS TS Performance Analyzer, and EAS support was **just added** for IMS V12 GSAM data sets. And APAR OA42723 was just opened because LISTDSI returns the wrong allocation amount if a data set has more than two extents on an EAV and is using EAS. I'm not listing the number of APARs dealing with EAVs due to the large number, but you should certainly investigate them yourself. I simply used a search on 'EAV' and found 60 APARs over the last year. IBM provides a FIXCAT of IBM.Function.EAV for APARs. Most of the APARs weren't problems with EAV itself, but with components that needed to add support for EAVs. To avoid problems even before you migrate, be sure to run the EAV Migration Assistance Tracker to find programs that use DFSMS interfaces without supporting EAS extents. You can find information about the tracker in the EAV Redbook and in the *DFSMS Advances Services* manual (both in *References* on page 19).

Many software vendors are providing EAV support in their products. IBM keeps a website updated showing the status of software developers supporting z/OS 1.13. There is also a column for EAV support. The website is - www.ibm.com/systems/z/os/zos/software/isv113.html. Here are comments from a few of the vendors I contacted:

> INNOVATION Data Processing
  > "*INNOVATION, as an EAV beta site working closely with IBM EAV development made sure at initial GA delivery of EAV, all INNOVATION products were ready to allow customers to achieve the benefits of using EAV while limiting their exposure to the difficulties they might have otherwise had managing data residing on EAV volumes.*"

> ➢ Products include: FDR / ABR data protection (backup/restore/erase); FDRInstant business continuation; FDRPAS / FDRMOVE / FDREPORT / FDRREORG storage media management (see **Figure 2** for a sample report); and IAM VSAM application accelerator.

Figure 2 - FDRREPORT - Two samples of EAV reports (© INNOVATION Data Processing)
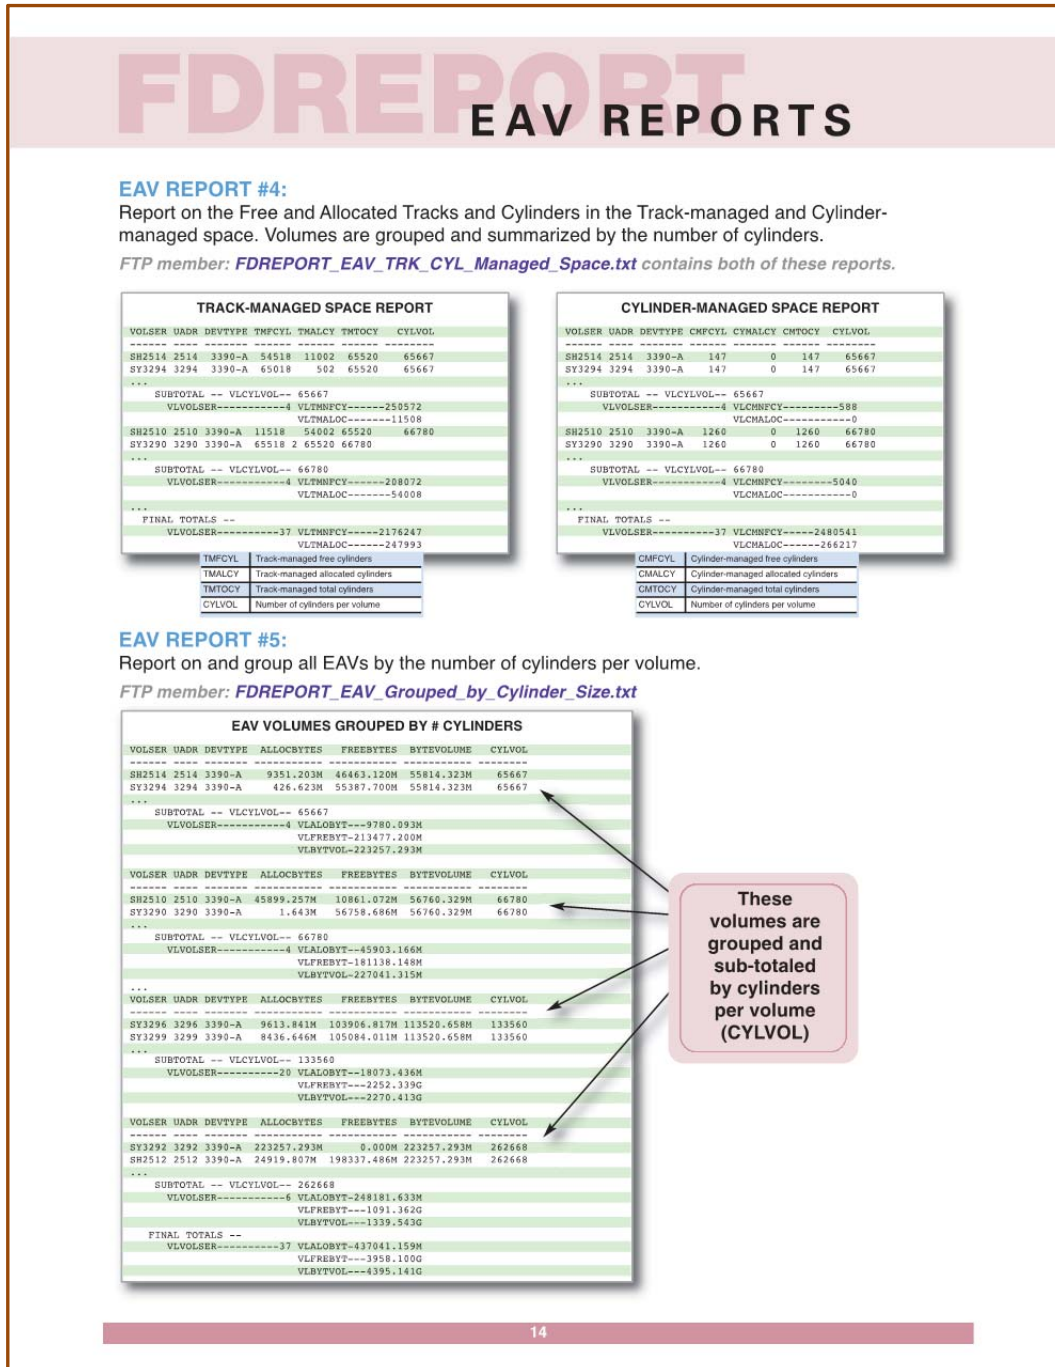
> ➢ [Allen Systems Group](#)
>   > ➢ TMON for z/OS can log to EAV volumes and provides an IO monitor for EAVs. It gathers data similar to RMF.

- ➢ Rocket Software
  - ➢ "*Too many of our products to mention support EAV, both for read and write. We were a beta site for EAV so we've supported EAV for a long time. What I don't know is how many of our customers actually use the EAV support.*"
- ➢ SAS Institute
  - ➢ A separate web page is provided to show EAV support by z/OS and SAS release, along with any limitations that should be considered - support.sas.com/kb/35/858.html.
  - ➢ For full support, a user should be running SAS 9.3 and z/OS 1.12. Note: At this time EAV is not supported for the load module libraries that are used in a CLIST or REXX exec for interactive SAS sessions. A fix is under consideration for a future release.
- ➢ CA
  - ➢ CA has a FIXCAT category CA.Function.EAV as found on their FIXCAT page https://support.ca.com/irj/portal/anonymous/phpsupcontent?contentID=%7b1426C7AF-3D50-49C1-ACE7-FCC325DCC1B0%7d
- ➢ Phoenix Software
  - ➢ "*All of our products have, for many years, fully exploited EAV by allocating/creating new data sets in the EAS at customer request, writing/reading EAS-resident data sets, reporting VTOC contents from F8/F9 DSCBs, etc. It just works.*"
- ➢ Software AG
  - ➢ A separate web page shows product compatibility for all IBM platforms (www.softwareag.com/Corporate/products/bis/platforms/ibm_availability.asp) and a direct link to a PDF shows the EAV support for their products including Adabas, Com-plete, and Natural.

Impact

The biggest impact to your installation will be the changes in policies and procedures relating to storage management. According to **Michael Graham's** (IBM) SHARE presentation 2571 (in *References* below), the performance impact should be very small. His presentation showed several benchmarks runs using different size volumes, and the performance is similar between current DASD volumes and EAVs. Here are the results from the presentation (note - the times are approximate because I estimated the values from a graph):

- ➢ DFSMSdss, while dumping the same amount of data (1.5 TB total), took longer for the 9 GB volumes (about 320 minutes) than for the 223 GB EAVs (about 280 minutes). See Figure 3.
- ➢ DFSMSdss, while copying the same amount of data (1.5 TB total), took less time for the 9 GB volumes (about 180 minutes) than for the 223 GB EAVs (about 200 minutes). See Figure 3.
- ➢ DFSMSdss had about equivalent throughput for Dump/Restore/Copy. There was an important note that you should "avoid consolidation to the point that it reduces exploitation of parallelism."
- ➢ DFSMSdss Defrag initially showed very poor performance with on EAVs prior to a major update. V2 of DEFRAG is required to have acceptable elapsed time

with EAVs. With this new version, DEFRAG can consolidate a 256K volume in a small percentage more than consolidating a 64K volume (about 430 seconds versus 390 seconds). It can consolidate with multiple extent data sets on a 256K EAV in about 410 seconds versus 300 seconds on a 64K volume.

➢ DFSMShsm backups and migrations were equivalent on 64K and 256K volumes.

➢ z/OS Global Mirror (XRC) had slightly worse response times and I/O rates for EAV volumes.

➢ Concurrent Copy (CC), FlashCopy (FC), and Virtual Concurrent Copy initialization time can be reduced by a few seconds when consolidating many smaller volumes onto fewer large volumes.

➢ TSO workload transaction rate and CPU time are equivalent when 640 users are moved from 16 9 GB volumes to 1 EAV volume.

> **IBM HIGHLY recommends HyperPAV when using EAVs**

➢ CICS workload transaction rate and CPU time are equivalent when the data for 1000 users per LPAR in three active LPARs was moved from 32 MOD3 volumes to 8 64K volumes and then to one 256K volume.

An important note Michael made in his summary was that it is highly recommended to use EAV volumes in conjunction with the HyperPAV feature. We discussed HyperPAV as part of this *Exploiting z/*OS series in *Tuning Letter 2012 No. 5* on pages 9-14.

Figure 3 - EAV Data Movement with DFSMSdss (© IBM)
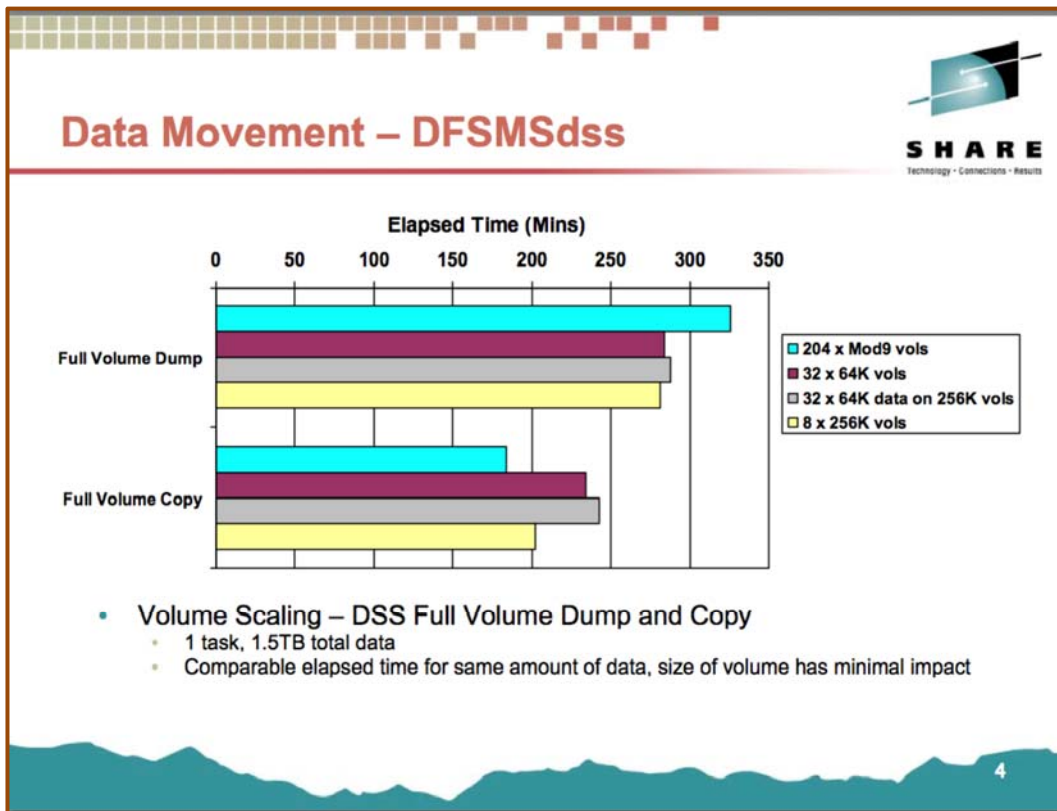
References

Cheryl Watson's Tuning Letter 2008 No. 5, Page 7 - *DS8000 - Extended Address Volumes.*

Cheryl Watson's Tuning letter 2011 No. 5, page 26 - *DS8000 Enhancements.*

Cheryl Watson's Tuning Letter 2013 No. 1, page 42 - *User Experiences.*

IBM Manual SC26-7400-14 - *z/OS 1.13 DFSMSdfp Advanced Services* (contains information about the EAV migration assistant tracker).

IBM Manual SC26-7473-11 - *z/OS 1.13 DFSMS Using the New Functions* (contains changes in each release and the implementation steps for each release).

IBM Redbook SG24-7617-00 (Updated 25Sep2009) - *DFSMS V1.10 and EAV Technical Guide.*

SHARE 2008 in San Jose - Session 2571 - **Michael Graham** - *Extended Address Volume (EAV) Performance.*

SHARE 2009 in Austin - Session 3023 - **James Cammarata** - *Extended Address Volume (EAV) - Overview, Usage and Invocation.*

SHARE 2009 in Austin - Session 3024 - **James Cammarata** - *Extended Address Volume (EAV) - Migration, Coexistence, Installation.*

SHARE 2010 in Seattle - Session 2417 - **Scott Drummond** - *What's New with Extended Address Volumes (EAV) in z/OS.*

SHARE 2010 in Boston - Session 7525 - **Tom Wasik** - *z/OS 1.12 JES2 New Functions, Features, and Migration Actions.*

SHARE 2013 in San Francisco - Session 13030 - **David Jones** - *z/OS JES3 Product Update and Review of Newer Features.*

Recommendation

I would investigate this ONLY if you are close to the limit of total devices (64K devices) or if you're finding it hard to manage the multitude of DASD devices in your installation. If this doesn't apply to you, spend your time investigating something more worthwhile. If this is true, you're probably a large installation, and one who could afford and benefit from using HyperPAV, which is almost a requirement when using EAVs. I agree with Dave Crow's comment on page 14 that it might be easiest if you implement EAVs with your next order of storage instead of trying to convert existing volumes. ∎

z/OS 2.1 Preview - Part 2

On February 5, 2013, IBM announced ([213-013](#)) a preview of the next MVS release, z/OS 2.1. The general availability (GA) for this release is expected in September 2013. Because this is a preview, there will be more details coming in July or August. In our last issue, we covered the Statements of Direction (SODs), and showed you how to exploit some of the z/OS 2.1 line items that are available as APARs for z/OS 1.12 and 1.13. In this issue, we'll cover items that will be available only in z/OS 2.1 which are the result of user requirements. And we'll give you a first look at the DASD requirements for z/OS 2.1. In our next issue, which will be after the SHARE conference in Boston (August 12-16), we'll have lots of the detailed z/OS 2.1 information that IBM will be providing at the conference.

Benefits of Requirements

As you probably know, I'm a strong advocate of the SHARE requirements process, and I'm chair of the MVSE requirements committee. The requirements process allows multiple customers to join together and vote on requirements and send them to IBM for future development features. We prioritize the things we'd like IBM to be working on, and they've been very supportive. They even highlighted at the z/OS 2.1 Preview that 25% of the SHARE MVSE project top 39 requirements had been satisfied in the new release. These are items that the SHARE users wanted, so they're probably things that you want too. For more information on requirements, see our *Tuning Letter 2011 No. 3* (pages 29-32) and my latest SHARE presentation on requirements at (www.watsonwalker.com/presentations.html). There will probably be additional satisfied requirements after the formal announcement, so we'll add those in the next *Tuning Letter.*

MVSE Top 39

MVSE is the z/OS Core Technologies project at SHARE. The Top 39 refers to a priority list of the 39 most important requirements as voted by the SHARE MVSE members. These are the line items from that list of 39 that were encouraged or initiated by users and satisfied by z/OS 2.1. Most of these are identified by a flag that says 'Top 39er' in **John Eells'** SHARE Session 12728 (see *References* below). The label at the end is the SHARE requirement number, and the link will only work if you are logged into the correct SHARE project (MVSE, MVSS, JES2, JES3) requirements section.

The following items are identified in John's handout as a 'Top 39er' in *What's New in z/OS V2.1*:

- Multiple SMP/E logical screens in ISPF. You can have one logical screen per SMP/E zone. [SSMVSE07004](#) (MVSE, 27 of 39)
- TSO/E LOGON failure messages will be issued to the terminal, such as LOGON FAILED ALLOCATION UNSUCCESSFUL preceded by the data set name that wasn't found. [SSSHARE015317](#) (MVSE, 10 of 39)

- Provide dynamic updates to system symbols. This is implemented with a new SETLOAD IEASYM keyword. SSMVSE10004 (MVSE, 1 of 39)
- z/OS console support for HMC 3270 console planned. This is intended to add an additional backup console and allow small LPARs to run without OSA-ICC. SSMVSE1107 (MVSE, 14 of 39)
- New MODIFY VLF command. Lets you specify COFVLFxx member to avoid VLF restarts. You can update VLF classes and major names, as well as change MaxVirt and AlertAge for existing classes. SSSHARE018333 (MVSE, 24 of 39)
- Add and remove MCS consoles dynamically. Command SET CON can process a CONSOLxx member to add a console, and command SETCON can remove a console. SSMVSE03026 (MVSE, 13 of 39)
- New DISPLAY PPT command. This displays the current PPT from the IBM default in CSECT IEFSDPPT, and SCHEDxx members at IPL, and any T SCH commands issued since the IPL. SSMVSE01502 (MVSE, 16 of 39)
- IEBCOPY improvement to support pattern matching. SOMVSE87101 (MVSE, 5 of 39), GO6SMG88017 (MVSS)
- TSO/E REXX enhancements to EXECIO, LISTDSI, and STORAGE. You can re-trieve information about data sets in EAS on EAVs. Support is provided for PDSE, concatenated data sets, multivolume, tape data set, I/O to undefined and spanned record format data sets. SSSHARE018757 (MVSE, 19 of 39)

> z/OS 2.1 provides dynamic updates to system symbols

More MVSE Requirements Satisfied

- Skips the 'Select Sysplex' popup on z/OSMF RMF reporting. SSMVSE11039 (MVSE)
- Reduce CPU time for z/OSMF. This is accomplished with the Liberty Profile. (See page 35.) SSMVSE11027 (MVSE)
- New operand on FORCE to terminate a TCB. SSMVSE09005 (MVSE)
- Unique 64-byte value assigned to each job in a sysplex. SSSHARE017984 (MVSE)

JES2 Requirements Satisfied

- JES2 adds support for 8-character job class names and provides for more than 36 JES2 job classes. The JCL JOB statement has been updated to allow specifi-cation of 8-character job class names for both JES2 and JES3. SSJES294003 (JES2)
- OUTPUT JCL will be honored in all cases - SSSHARE018281 (JES2)
- Converter Support for JECL cards and $HASP Messages. SSSHARE016034 (JES2)
- JES OUTPUT JESDS processing is supported even when the job fails before ex-ecution. GGMVJS94002 (JES2)
- JES2 command to display parmlib members used. SSJES299001 (JES2)

- SAF control over the use of job classes for both JES2 and JES3 environments using new profiles in the JESJOBS class is supported. SSJES297205 (JES2)

JES3 Requirements Satisfied

- JES3 dynamic spool delete. SSJES3032645 (JES3)
- Dump Job command is enhanced to let the user define the SPOOL extent to be dumped. SSSHARE01633 (JES3)
- Existing commands will be enhanced to display spool information for individual jobs or display which jobs have data on a particular spool data set. SSJES3032649 & SSJES399351 (JES3)
- The JCL JOB statement has been updated to allow specification of 8-character job class names for both JES2 and JES3. SSJES294003 (JES2)
- SAF control over the use of job classes for both JES2 and JES3 environments using new profiles in the JESJOBS class is supported. SSJES297205 (JES2)

> z/OS 2.1 - GDGs in FIFO order? I'm in heaven!

MVSS Requirements Satisfied

MVSS is the storage project within SHARE.

- GDGs in chronological order. A JCL DD statement GDGORDER keyword can specify FIFO (first in, first out). SSMVSS063068 (MVSS)
- IEAAPP00 parmlib member will allow comments. SSMVSE11045 (MVSE)
- DFSMShsm will automatically generate a new RECYCLE command to the same original tape when the original RECYCLE has to terminate due to a take-away process. SSMVSS1004 (MVSS)
- DFSMShsm provides an option for users to request a RECYCLE in response to an alternate tape copy instead of generating a tape copy request. SSMVSS1003 (MVSS)
- Provide creation date for catalog aliases. SSMVSE10018 (MVSS)
- IDCAMS commands PRINT, REPRO, IMPORT, and EXPORT will be able to open VSAM data sets using RLS with a new keyword of RLSSOURCE and/or RLSTARGET. SSMVSS01007 (MVSS)
- PDSE member sizes will be increased from 15,728,639 lines to 2,146,435,071 lines (125 times larger). SSMVSS11010 (MVSS)
- A change in Open/Close/End of volume processing of SMS-managed data sets will release space in all extents when RLSE is specified (instead of just on one volume). SSMVSS08002 (MVSS - partially addressed in z/OS 1.12)
- The user may specify whether the DS1DSCHA (data-set-changed-bit) is turned on or off during a DFSMSdss Full Volume RESTORE using a new RESET keyword on the dump. SSMVSS07002 (MVSS)

ISPF Requirements Satisfied

- Edit Compare will be able to specify a volser in order to reference uncataloged data sets. SSISPF01105 (ISPF)

- HILITE will be able to highlight the invalid use of lowercase characters in JCL. SSISPF10002 (ISPF)
- FIND and CHANGE commands will be able to use regular expressions. SSISPF98031 (ISPF)
- ISPF will support multiple logical screens created when you invoke ISPF, and end them when exiting. SSISPF99112 (ISPF)
- Increased scroll limits past 9999 are supported in Browse, Edit, View, member lists, and data set lists. SSISPF00102, SSISPF03013 (ISPF)

References

This information was obtained (or extrapolated) from the z/OS 2.1 preview presentations at the last SHARE conference:

SHARE 2013 in San Francisco - Session 12728, **John Eells** - *What's new in z/OS V2.1.*

SHARE 2013 in San Francisco - Session 13029, **Tom Wasik** - *z/OS JES2 Product Update and Latest Status.*

SHARE 2013 in San Francisco - Session 13030, **David Jones** - *z/OS JES3 Product Update and Review of Newer Features.*

SHARE 2013 in San Francisco - Session 12976, **Barbara McDonald** - *DFSMS Latest and Greatest.*

After z/OS 2.1 is announced, IBM will update their website that shows z/OS User Group Requirements (www.ibm.com/systems/z/os/zos/bkserv/user_group_reqs_all_previous.html). This shows all requirements satisfied by each z/OS release. But this list of requirements comes from other groups than SHARE, and includes customer groups, vendors, and other sources. The requirements listed above were all originated within the SHARE MVS program.

DASD Requirements

The DASD requirements for z/OS 2.1 won't be finalized until it's announced, but **Marna Walle** (IBM migration expert and guru) provided the sizes for her installation in a recent SHARE blog titled "*Get your DASD installation requirements in early!*" You can find the blog here - www.share.org/p/bl/et/blogaid=239, but we thought this was important enough that we've reprinted it here (with her permission). This was confirmed by at least one early tester. Check the blog to see if other comments have been added. There are two very important points here - an install of z/OS 2.1 will take much more DASD space (83% in her case) than 1.13, and z/OSMF has reduced it's footprint so much (87% in her case) that anyone can install it even on a small test LPAR. As Marna says - *thank you Liberty profile!* (See our article about Liberty Profile on page 35.)

> z/OSMF 2.1 just lost 87% of its weight!

Here's her blog for May 17, 2013:

> *Been a while since I've been able to blog… z/OS V2.1 is strongly under-way here, and that has been taking me away from my "treat" of blogging. For this entry, I wanted to pass along some information about planning for z/OS V2.1 that you may find helpful. (Remember, I've already mentioned to use the IBM Health Checker for z/OS right now.)*

> *I've had the question on how much DASD storage installing z/OS and z/OSMF V2.1 will take. Here's some numbers that you might find interest-ing. Of course, these are sizes taken from our systems and the usual ca-veats apply:*

> - *we use both JES's, we use all NLV's, we don't count non-z/OS program products,*
> - *the z/OS V2.1 sizes are close estimates*
> - *for the total z/OSMF sizes for R13, you need to add the "z/OSMF R13" and "WAS OEM" columns. For the total z/OSMF V2.1 sizes, you *only* need the z/OSMF V2.1 column (thank you Liberty profile!)*

Table 1 - DASD installation sizes (in 3390 cylinders)

	z/OS R13	z/OS V2.1	z/OSMF R13	z/OSMF R13 WAS OEM	z/OSMF V2.1
Target PDS and PDSE	5,868	10,000	2	2	2
DLIB	8,941	18,000	148	1,582	200
File System (zFS)	3,309	3,300	147	2,600	400
Font File System (zFS)	---	2,050	---	---	---
Grand Total	18,118		297	4,184	
		33,350			602
		(15,115 more)			(3,879 less)

> *z/OSMF V2.1 does NOT "imbed" Java SDK. Meaning that if you want to use z/OSMF V2.1 there will be an external dependency on Java SDK R7 64-bit + PTF. z/OSMF R13 did "imbed" Java SDK, so that was included in the table sizes above. How does possibly one less Java SDK to lug around sound to you??*

> *Notice that there is a new file system for z/OS V2.1 - the font file system. More about that in a later blog. For now, know that you've got another 2,000 cylinders of file system space coming at you for z/OS V2.1.*

> *We told you there was a lot coming in z/OS V2.1!*

> *-Marna* ■

SMF Type 42, Subtype 6

*My students used to joke about a phrase I often use: "**THIS** is my favorite SMF record!" I would use the phrase for the 30/4 (step termination), type 64 (VSAM stats), type 72 (workload manager), and (of course) the type 79 records. But the one that is **REALLY** my favorite is the type 42, subtype 6 record. Few people use these records, but those that do become heroes in their shops. They're incredibly delicious! (I originally discussed this record type in our Tuning Letter 2000 No. 8.)*

SMF Type 42.6

Introduction

The type 42, subtype 6 record is produced by DFSMS. It was originally intended to provide information on SMS-managed DASD volumes, but was extended to provide some excellent information for both SMS and non-SMS managed data sets.

The records are written when a data set is closed and at the end of every SMF interval (as defined in SMFPRMxx). The record format of the SMF Type 42, subtype 6 record is fairly trivial. There is a 36-byte header, then an 8-byte heading for subtype 6, a 40-byte product section, a 64-byte job header section, one or more 88-byte data set sections, a 64-byte I/O response time section, a 72-byte I/O statistics section, and a 48-byte access method statistics section. One record is written for each data set and includes the following:

Job Header Section
- ➢ Job name
- ➢ Job reader timestamp
- ➢ Service class name
- ➢ Workload name

Data Set Header Section
- ➢ Data set name
- ➢ Data set type (e.g. PS, PDS, PDSE, KSDS data component, and KSDS index component)
- ➢ Type of buffering
- ➢ Volume serial number
- ➢ Device number
- ➢ Storage class name
- ➢ Blocksize (if concatenated, then blocksize comes from the first file)
- ➢ Number of stripes

Data Set I/O Statistics Section
- ➢ Average response times (total response, connect, pend, disconnect, and CU queue times)
- ➢ Number of physical I/Os

- ➢ Cache candidates and cache hits
- ➢ Write candidates and write hits
- ➢ Number of sequential I/O operations (does not include cache read and write candidates)
- ➢ Number of record level cache I/O operations
- ➢ Number of inhibit cache load I/O operations
- ➢ Average I/O device-active-only time
- ➢ **Maximum data set I/O response time**
- ➢ **Maximum data set service time**
- ➢ Average disconnect time for reads
- ➢ Number of read operations

Access Method Statistics Section
- ➢ Number of blocks by type (sequential read, sequential write, direct read, and direct write)
- ➢ Amount of I/O delay by type (sequential read, sequential write, direct read, and direct write)
- ➢ Number of directory reads and writes
- ➢ Amount of I/O delay for directory reads and writes

A very interesting item in the I/O statistics section is the availability of the maximum response time and service time encountered during the interval. Average response times might look very good, but a large maximum response time could indicate that a user was inconvenienced at least once. It gives you a much better understanding of what the users see.

How To Use This Data
You can summarize the data in any number of ways:

- ■ First, calculate the load of each record by multiplying the number of I/Os by the average response time. The records with the highest load indicate the most time spent on behalf of I/O. Sort on descending load to find the worst data sets. You can also use the load in some of the following studies.

- ■ Find the busiest volumes by summarizing the load or I/Os for each volume, then sort on the descending load or I/O per volume.

- ■ Select records for the busiest volume and then sort on descending I/Os or load to find the highest activity data sets.

- ■ Summarize by data set name by adding the I/Os or load and then sort by descending I/Os or load to find the worst performing data sets in the center.

- ■ Analyze a specific job by extracting all records for that job. Calculate the load and sort in descending load to find which data sets contributed most to the elapsed time of the job.

- Analyze the hit ratios to find high hit ratios that should have larger buffers in order to decrease the amount of cache activity. Hit ratios are calculated as 100% * (hits / candidates).

- Analyze the hit ratios to find low hit ratios that could benefit from being behind a controller with more cache.

- Find contention on volumes by identifying high activity data sets on the same volume that are accessed concurrently.

- Summarize total I/Os and load to find the jobs responsible for the highest I/O activity in the data center.

- Sort on descending maximum data set response time or service time in order to find those data sets that experienced very poor times, even when their average looked fine.

- Compare the I/O response time to the data set service times to calculate the IOSQ time by job and data set name! Using this data, you can determine which jobs and data sets are responsible for high IOSQ time on volumes.

- If you've been analyzing the type 74 data (activity by volume) and can't understand what was happening to a volume, use the type 42.6 records to identify the jobs and data sets active on the volume during the period in question.

In many of these reports you'll want to eliminate the trivial data sets (80% of all data sets) by ignoring anything that took less than 5 minutes of total response time (i.e. load less than 5 minutes) or had less than 100 I/Os. Pick your own cutoff point.

Example

Figure 4 is an example of a report sorted in descending I/Os (the actual report was several pages long). This example came from MXG's ANAL42 program. From this report you can concentrate on the most active data sets. Look at the response time for the second data set. This should definitely be reviewed! Maybe that's how that file

Figure 4 - Type 42.6 Report From MXG

```
          ANALYSIS OF TYPE 42 IO RECORDS    08:41 WEDNESDAY, SEPTEMBER 1, 1999
                            SORTED BY NUMBER OF IOS
                  SHOWING NUMBER OF IOS AND RESPONSE TIME COMPONENTS

———————————————————————————————————— SYSTEM*ID=SYSA ————————————————————————————————————

JOB NAME   PHYSICAL    VOLUME                     NO    NO    AVG    AVG    AVG    AVG    AVG    AVG
   OR      BLOCK SIZE  SERIAL   DATA SET          OF    OF    RESP   DISC   IOSQ   PEND   CU Q   CONN
TSO USER   (BYTES)     NUMBER   NAME              IOS   JOBS  TIME   TIME   TIME   TIME   TIME   TIME

AAAAYALA    27540      VVV008   ABC.DEF.GHI    950,153   14   3.4    0.0    1.0    0.3    0.0    2.0
BBBBE20A    27540      VVV007   JKL.MNO.PQR    827,647   28   279.1  110.5  1.3    0.2    0.0    167.1
CCCCDBM1    4096       VVV009   STU.VWX.YZA    791,712   21   4.3    1.2    1.8    0.3    0.0    1.0
DDDDDBM1    4096       VVV005   BCD.EFG.HIJ    762,037   21   10.7   1.2    7.1    0.1    0.0    2.3
. . .

JOB NAME   PHYSICAL    VOLUME                     PDS        SEQ       DIRECT
   OR      BLOCK SIZE  SERIAL   DATA SET          DIR        BLX        BLX         CACHE       CACHE
TSO USER   (BYTES)     NUMBER   NAME             READS      READS      READS        CANDS        HITS

AAAAYALA    27540      VVV008   ABC.DEF.GHI        0       619,668        0           0           0
BBBBE20A    27540      VVV007   JKL.MNO.PQR        0       495,412        0           6           6
CCCCDBM1    4096       VVV009   STU.VWX.YZA        0          0      1,951,481    772,007     692,595
DDDDDBM1    4096       VVV005   BCD.EFG.HIJ        0          0      5,849,689    705,066     705,052
. . .
```

works, but it should be checked out. You can also determine the reason for the majority of the response time (such as the disconnect time for our large file). With this report, you can simplify the report by including only data sets whose response time is over 30 milliseconds and whose number of I/Os was over 5000.

Figure 5 is a different example showing different information for some selected jobs and data sets. (The report was modified a little to fit onto the page.)

Figure 5 - Type 42.6 Report From MXG's ANAL42DS

```
                          Proc Print of Some ANAL42DS Variables          14:39 Tuesday, June 25, 2013 2650
                   DATA                                    INTERVAL                END TIMESTAMP           DASD
          PROGRAM  SET                                     DATA SET OR CLOSE         OF THE              IO/S PER
JOB NAME  NAME     NAME                                    TYPE RECORD             INTERVAL              SECOND

AJOBNAME  SORT     PROD.NODE21.SC28006.D130625.T090646     1:PS 0:FINAL INTERVAL 25JUN2013:10:19:38.55    5.6490
AJOBNAME  DYL280   PROD.NODE2R.PRODMULT.NOUPDT             1:PS 0:FINAL INTERVAL 25JUN2013:10:19:43.23   75.9535
AJOBNAME  SORT     PROD.NODE21.AGY28006.RFMTTAF.G0006V00   1:PS 0:FINAL INTERVAL 25JUN2013:10:19:45.62   11.4239
AJOBNAME  SORT     PROD.NODE21.AGY28006.PAYMENT.G0012V00   1:PS 0:FINAL INTERVAL 25JUN2013:10:19:48.57   20.9516
                                                                              AVERAGE
                   DATA                                    AVERAGE         CONTROL UNIT        AVG             AVERAGE
          PROGRAM  SET                                     I/O CONNECT        QUEUE       I/O DEV-ACTIVE-ONLY  I/O DISCONNECT
JOB NAME  NAME     NAME                                    MSEC PER SSCH   MSEC PER SSCH   MSEC PER SSCH    MSEC PER SSCH

AJOBNAME  SORT     PROD.NODE21.SC28006.D130625.T090646       1.920           0.000           0.0            17.792
AJOBNAME  DYL280   PROD.NODE2R.PRODMULT.NOUPDT               0.000           0.000           0.0             0.000
AJOBNAME  SORT     PROD.NODE21.AGY28006.RFMTTAF.G0006V00    46.976           0.000           0.0             0.000
AJOBNAME  SORT     PROD.NODE21.AGY28006.PAYMENT.G0012V00     3.328           0.000           0.0             0.000

                   DATA                                    AVERAGE         AVERAGE        CACHEABLE      CACHEABLE
          PROGRAM  SET                                     IOS QUEUEING    I/O PENDING       IO/S           HIT
JOB NAME  NAME     NAME                                    MSEC PER SSCH   MSEC PER SSCH   PERCENTAGE     PERCENTAGE

AJOBNAME  SORT     PROD.NODE21.SC28006.D130625.T090646       0.256           0.128            0              .
AJOBNAME  DYL280   PROD.NODE2R.PRODMULT.NOUPDT               0.384           0.128            0              .
AJOBNAME  SORT     PROD.NODE21.AGY28006.RFMTTAF.G0006V00     0.000           0.128            0              .
AJOBNAME  SORT     PROD.NODE21.AGY28006.PAYMENT.G0012V00     0.128           0.128          100            100
```

Chuck Hopf, MXG and performance consultant (chuck@chopf.com), sent an interesting note about his experiences:

> *With all of the improvements in the IO subsystem – moving things down into the channels – HyperPAV – faster channels – etc – I have seen very few real IO problems lately. Most of the time IO response time is so far down in the sub-millisecond range it is not an issue. Also, so much of the IO is being done by a DBMS that it gets hard to know who is doing what to whom from the 42 records. All it is going to tell you is DB2 did it but not who within DB2. And an IO trace in DB2 can be an ugly thing. Saw one bring a system down once.*

To illustrate this, he shared the graph in Figure 6 and these comments:

> *Look at the EXCP counts from the SMF Type 72s on a recent system I was working on. At least 75% come from DB2, ADABAS, or DFHSM. In none of those cases is the TYPE 42 going to be of much utility in tracking down a problem. It might, in the case of DB2, tell me what table was suffering but even then that is probably a known rather than a question.*

Chuck makes a good point. If the Job name is always DB2, CICS, or DFSMShsm, then there are other reports that will give better information, especially about buffering.

On the other hand, if you're looking at a critical batch window and want to find where you can reduce elapsed time, the Type 42.6 records can tell you where to spend your time in tuning I/O. This is why the zBNA team is looking at the SMF Type 42s to determine the impact of I/O of batch work. See our article on zBNA on page 48.

Figure 6 – Distribution of EXCP by Source
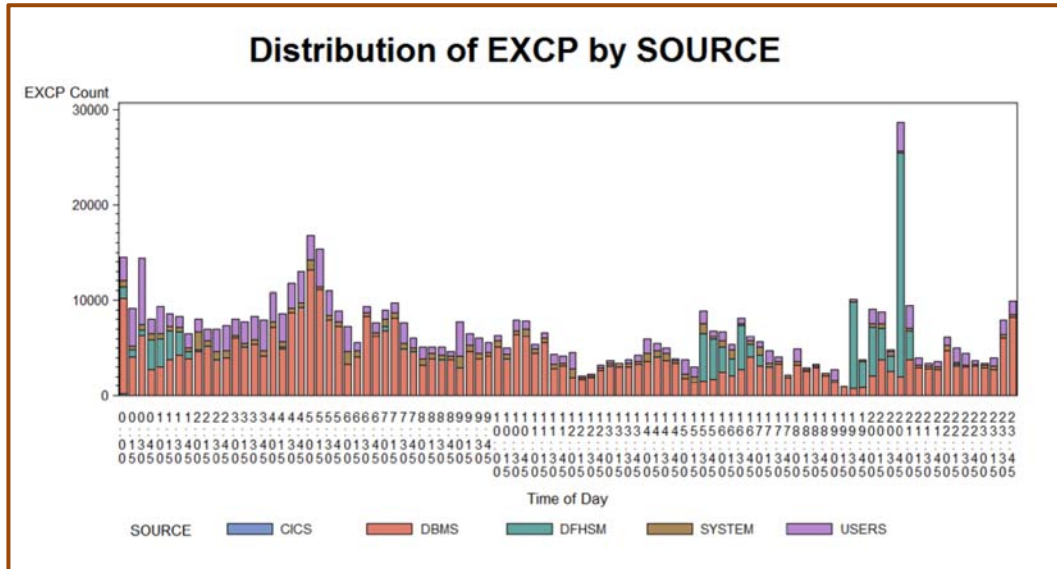
References

IBM Manual SA22-7630-25 - *z/OS V1R13.0 MVS System Management Facilities (SMF)*
IBM Techdocs Technotes – TD102069 (21Dec2004) – *Understanding the HFS and EXCP Counts in RMF and SDSF.*

Summary

I think you can see why I like these records so much. They can help tune your batch I/O system better than almost any other data available to you. Have fun!

zFS and HFS Counts

Dr. Barry Merrill writes an MXG Newsletter for his customers that is also available to the public. While many of the items are related to recent changes in the MXG software, he often has some fascinating articles on analysis that he's performed. The following article was included in his *MXG Newsletter Number Fifty-Five*, from January 20, 2010. It's been reformatted for our *Tuning Letter*, but the content is the same, and we're printing it with his permission. The purpose of his analysis was to identify how "EXCP" counts are recorded for zFS and HFS filesystems.

"EXCP" counts recorded for access to HFS & ZFS filesystems

An HFS file, 10,000 50-byte records, 496K, or 123 4096-byte blocks, and a ZFS file, 1,000 50-byte records, 49K, or 13 4096-byte blocks, was created/copied on z/OS 1.9 by different programs.

Total "EXCPs" were between 50 and 23,710 for HFS.
Total "EXCPs" were between 37 and 5,416 for ZFS.

These "EXCP" counts are displayed on JOBLOG and are included in the SMF 30 Address Space Total EXCP count, EXCPTOTL (SMF30TEP).

Job Description	HFS 496K EXCPTOTL	ZFS 49K EXCPTOTL
TEST92LD –SAS92 LOAD	23710,23290	5416
TEST91LD –SAS91 LOAD	21856,21785	3867
TEST92RD –SAS92 READ	13364,13295	4464
TEST91RD –SAS91 READ	11787,11763	2891
TESTGENR –IEBGENER READ	309, 306	n/a
TESTFAST –FASTGENR READ	298, 209	65
TESTSORT –SYNCSORT READ	209, 209	70
TZOS92LD –SAS92 LOAD z/OS	3301	3324
TZOS91LD –SAS91 LOAD z/OS	1764	1771
TSTWGENR –IEBGENER WRITE	301	n/a
TSTWFAST –FASTGENR WRITE	268	53
TSTWSORT –SYNCSORT WRITE	252	62
ZOSCGENR –IEBGENER COPY	113	53
ZOSCFAST –FASTGENR COPY	50	28
ZOSCSORT –SYNCSORT COPY	46	37

All of the SMF records written for two of these test jobs were analyzed in detail: the SAS-TEST91LD and FASTGENR-TESTFAST are analyzed in detail below.

SAS was used to create a 10,000 record text file of 50 byte records, written to an dynamically allocated HFS1 Filename. FASTGENR was then used to copy that HFS file, with a static SYSUT1 DD, to a disk data set.

- EXCP counts in DD Segments in SMF type 30 subtype 2, 3, 4, and 5):

  ➢ There was no DD segment created for the dynamically allocated HFS1 DDNAME in the SAS job.
  ➢ While there was a SYSUT1 DDNAME in the type 30 records for the FASTGENR job, it contained ONLY the DDNAME; there were no EXCPs recorded, and there was no DEVNR nor DEVCLASS/DEVTYPE information.

- HFS "EXCP" counts ARE captured in the SMF 30 record; but only in the address space total EXCP Count EXCPTOTL(SMF30TEP/TEX). RMFEXCP are the EXCPs counted in IO Service Units (SMF30IO/IOL), and the HFS EXCP count IS included in RMF IO Service Units.

- *EXCPTODD is the sum of all EXCPs in the DD segments.*
- *EXCPNODD is the EXCPs count NOT counted in the DD segments, calculated as EXCPTOTL minus EXCPTODD.*
- *EXCPDASD is the total DD EXCPs count to DASD devices.*
- *SMF30AIS is the total count of DASD SSCH's (NOT BLOCKS/EXCPs)*
- *IOTM variables are the IO Connect Time durations, as above.*

```
JOB          EXCPTOTL RMFEXCP  EXCPTODD  EXCPNODD EXCDASD SMF30AIS
SAS             21785   21778      1379     20406    1379     704
FASTGENR          285     280       101       184     101     213

JOB          IOTMTOTL RMFIOTM  IOTMTODD  IOTMNODD
SAS              0.51     n/a      0.37      0.14
FASTGENR         0.02     n/a      0.02      0.01
```

Observations:

➢ *SAS wrote 10,000 blocks of 50 bytes each, but counted 20,000 EXCP, and that count was also shown on the SAS log; why 20000 was the count will be investigated, but that count of 20000 was passed to IEASMFEX, as it does show up in the EXCPTOTL and RMFEXCP.*

➢ *FASTGENR, the SYNCSORT replacement for IEBGENER, counted 101 "EXCPs" to the 3390 output disk device in the EXCP segment for SYSUT2; the "EXCPs" reading the HFS file were counted as 184 in the EXCPNODD (i.e., included in EXCPTOTL and RMFEXCP). But FASTGENR and SYNCSORT have NEVER counted EXCPs, but, instead count SSCHs, and that is what it passed to IEASMFEX. (I was involved in legal issues between DFSORT and SYNCSORT because SYNCSORT published false I/O comparisons that used the SIOs for SYNCSORT but BLOCKS/EXCPs for DFSORT, many years ago. There was a "Special Core Zap" from SYNCSORT that would change their count to BLOCKS, but I don't know if it still exists, and I suspect no one uses it now!). In addition, the FASTGENR log shows that it read and wrote 10,000 logical records; however it shows a total size of 800,000 bytes, whereas only 500,000 bytes were written, so even FASTGENR can't correctly count HFS activity.*

➢ *While HFS EXCP counts are in the EXCPNODD field, there are other I/O counts included in EXCPNODD, for all file I/O that does not have a DD: Catalog I/O, LinkList I/O, and JES2 SPOOL I/O, and the JES2 Spool I/O count can be significant.*

◾ *HFS-only EXCP counts do exist in the OMVS Segment of type 30s.*

The old "OMVS" segment is now known as "z/OS UNIX System Services Process Section" in the SMF manual. I LOVE the fact that UNIX is in CAPITAL LETTERS! MXG's first technical note on measuring unix, by Chuck Hopf, was subtitled "or how i learned to type in lower case".

The SAS job created one "OMVS" segment, while the FASTGENR created two segments, having apparently spawned/forked/whatever unix does that created a second PID for their copy program. The first three columns are the only

block count fields that were non-zero; the last columns are the only other metrics that were non-zero.

| | DIR READ BLOCKS | DATA READ BLOCKS | DATA WRITE BLOCKS | PATHNAME LOOKCALL LOGICAL FILES | PATHNAME LOOKCALL PHYSICAL FILES | SYSCALLS REQUESTED BY PROCESS |
JOB	OMVSODR	OMVSOFR	OMVSOFW	OMVSOLL	OMVSOLP	OMVSOSC
SAS	65	0	20000	8	37	21
FASTGENR–1	16	0	0	2	8	3
FASTGENR–2	26	125	0	3	13	4
FASTGENR	42	125	0	5	21	7

Comparing the type 30 with the type 30 OMVS segment:

	Total I/O Blocks OMVS	Total NODD IO COUNT
SAS	20065	20406
FASTGENR	167	184

Observations:

➢ *The UNIX segment EXCP counts can indeed be subtracted from the address space EXCP counts, for sites that do NOT want to include HFS EXCPs in their billing, if they are using the EXCPTOTL field.*
➢ *I polled MXG users, and most said that when EXCP counts were used in chargeback, they used only the EXCPDASD and EXCPTAPE counts (MXG sums DD EXCP counts by device type); the use of EXCPTOTL that includes HFS (and SPOOL) counts are not commonly used in billing.*

■ *HFS-only EXCP counts do exist in the Type 92 records.*

	BYTES READ SMF92CBR	BYTES WRITTEN SMF92CBW	DIR I/O BLOCKS RD/WR SMF92CDI	DATA I/O BLOCKS READ SMF92CIR	DATA I/O BLOCKS WRITTEN SMF92CIW	READCALL ISSUED SMF92CSR	WRITECALL ISSUED SMF92CSW
SAS:	0	498K	12	0	20000	0	20000
SYNC:	498K	0	10	125	0	9	0

Observations:

➢ *While FASTGENR reported 800,000 bytes copied, the SMF 92 shows that FASTGENR is wrong (it used a default LRECL=80 times 10,000 logical records), and that SAS was right (it showed 10,000 logical records with the correct 50 byte LRECL).*
➢ *The EXCP counts for HFS activity, 20012 for SAS and 135 for SYNC in the SMF 92 are consistent with the counts in the OMVS segment and the EXCP counts passed into the type 30 step records, but the values are the counts*

passed by the application, blocks for SAS, and SSCH for FASTGENR, and there's no way to tell which is which.

■ *No SMF 42 subtype 6 records were created for HFS for these jobs.*
And I did NOT expect to see those records, as they are documented in the SMF manual "records DASD data set level I/O statistics", and, for these two jobs, HFS was NOT a DASD data set.

There were type 42 subtype 6 records created for the DASD DDnames for the two jobs, and they captured these counts, for comparison with the SMF 30s:

JOB TOTAL	TOTAL NUMBER OF IOS IOCOUNT (S42DSION)	CACHE CANDIDATES (S42DSCND)	Sequential blocks read (S42aMSRB)	Read Operations to Dataset (S42DSRDT)	Sequential I/O's (S42DSSEQ)
SAS	442	187	27	431	5
FASTGENR	101	1		1	100

Observations

➢ *Whereas the EXCP counts in the TYPE 30 are whatever the application access method passed to SMF, type 42 subtype 6 counts are direct from the hardware, independent of the access method, etc.*
➢ *The FASTGENR SSCH count of 101 SSCHs in the type 42 is the same as the SSCH count passed by FASTGENR into the SYSUT2 DD segment, and that was the only DD allocated to DASD, since SYSUT1 points to the hfs file. But the (relatively new) SMF30AIS field, documented as "DASD I/O Start Subchannel Count for address space and dependent enclaves" count of 213 appears to me to be in error. The SAS EXCPDASD count of 1379 is consistent with SMF30AIS of 704, as half-track blocking is normally used by SAS.*
➢ *I believe there would be type 42 subtype 6 records created for the z/OS VSAM file that "contains" the HFS file system, but those data would have the JOB name of the address space from which the actual physical I/O occurs, and those counts would be for all users of the file system, with no counts for the actual jobs that cause the I/O.*

■ *Data on the banner page may include HFS counts in the "EXCP Count"*

This site uses the IBM "banner page" to print EXCP counts on Job Log; the EXCP count that is printed is, indeed, that EXCPTOTL/SMF30TEP Address Space Total Count, and which we now know DOES include the HFS EXCPs, and those counts are only slightly larger than the two products reported on their execution logs:

	Banner Page EXCPs	Product Log's EXCPs
SAS	21785	20240
FASTGENR	285	240

Observation:

➢ *This is very likely the source of the large EXCP counts that have been re-ported, since it requires no analysis of the SMF 30 records (and I think this is also the EXCP count displayed by SDSF).*

Conclusions

■ *Whatever is counted by the application as an "EXCP" for HFS access whether blocks or SSCHs (at the whim of the I/O programmer!) is included in the EXCPTOTL field in the SMF 30 records, and is the count that is displayed by banner pages and SDSF.*

■ *The type 30 OMVS segments are now used in MXG 27.08 Change 27.213, to create the new USSEXCPS count variable that could be used to "back-out" these large counts, if the site is actually using the EXCPTOTL field in charge-back and has significant USS activity.*

■ *With the inaccuracies in counting HFS and zFS EXCPs, and because they are included in the RMF IO Service Units, alternative ways to count, including di-viding the total bytes in the type 92s by 4096 are under consideration by IBM. This research is in progress and this note will be updated if corrections are made.* ■

What's New?

In this section we try to provide our view on some of the latest announcements that we consider important for our readers. The direct link to IBM USA announcements has been changed to a generic lookup, and is now located at www.ibm.com/common/ssi/index.wss.

WebSphere Liberty Profile

The Liberty profile is an ultra lightweight profile of the WebSphere Application Server (WAS) that is optimized for developer productivity and web application deployment. It is available on multiple platforms including z/OS, and provides significant performance benefits in elapsed startup time, CPU usage, reduced memory requirements, and reduced DASD resources. It provides a subset of traditional WAS features, and can be configured with selected components using a single process instead of three for traditional WAS. It was first made available over a year ago, but more subsystems are starting to exploit it now.

> Liberty Profile will make z/OSMF available to all

My interest in the Liberty Profile has been its potential for z/OSMF. Although I believe that everyone should be exploiting z/OSMF, it's been difficult to even test because it needed a fully-loaded WebSphere OEM product that often couldn't run in a test or development configuration. As mentioned at recent conferences, z/OSMF for z/OS 2.1 will be built on a Liberty base. That will make z/OSMF available to many more installations. See the z/OS 2.1 sizing estimates on page 22, where the DASD footprint for z/OSMF was reduced by 87% due to the Liberty Profile.

Here is the subset of traditional WAS features (as of WAS 8.5.0) that are supported by Liberty:

- Bean validation
- Blueprint
- Java API for RESTful Web Services
- Java Database Connectivity (JDBC)
- Java Naming and Directory Interface (JNDI)
- Java Persistence API (JPA)
- Java Server Faces (JSF)
- Java Server Pages (JSP)
- JMX
- Monitoring
- OSGi JPA
- Remote connector
- Secure Sockets Layer (SSL)

- Security, supported by basic user registry or a Lightweight Directory Access Protocol (LDAP) user registry
- Session Persistence
- Transaction
- Web application bundle (WAB)
- Web security
- z/OS Security (SAF)
- z/OS Transactions (RRS)
- z/OS Workload Management

WAS 8.5.5 (see announcement 213-174 below) adds the following support:

- Certification to the Java EE 6 Web Profile, which includes new programming models such as EJB-Lite, Managed Beans, and Context and Dependency Injection (CDI)
- Web services support through the JAX-WS Java EE API
- New messaging capabilities, including MQ client support, support for Java Message Service (JMS) and message-driven beans and a new single server message provider
- Ability to add Liberty features through a new system programming interface
- Support for the NoSQL MongoDB client
- Enhanced Liberty profile security capabilities
- High Performance Extensible Logging (HPEL) for Liberty servers
- Encryption of passwords in server configuration
- New Liberty administration features
- Dynamic Web Cache support
- Clustering of server instances
- Asynchronous Work Management

> Liberty Profile uses about 55 MB versus 2.5 GB (!!) of T-WAS

Here are some initial performance results (see the SHARE session 12183 below by **David Follis**):

- Startup CPU time reduced from 21.2 seconds to 4.1 seconds
- Startup elapsed time reduced from 21.6 seconds to 3.2 seconds
- Memory footprint reduced from 420 MB to 80 MB
- Throughput of TradeLite (a benchmark program) improved by 35% using Liberty over traditional WAS
- Although not in his handout, David says that the disk footprint in 8.5.0 was just 55 MB. It grows in 8.5.5, but they've add a 'minify' feature that let's you build deployable packages that contain only the .jar files you need to support the features you are using. So a developer might need the whole pile, but deployed production servers can be smaller by eliminating runtime pieces it doesn't need. When you compare this to a traditional WAS that needs about 2.5 GB (!!!), you can see why everyone is so excited about this enhancement.

One of the main features of the Liberty Profile is that it can be deployed on a variety of platforms, not just z/OS, so that it can be used for development with no change of features. An early exploiter of the Liberty Profile is CICS (see session 12445 below by **Catherine Moxey**). Each installation / subsystem / product will provide different configurations for the Liberty Profile. I like Catherine's slide (see Figure 7) showing the difference between traditional WebSphere Application Server (tWAS) and Liberty.

Figure 7 — CICS and Liberty Configurations (© IBM)
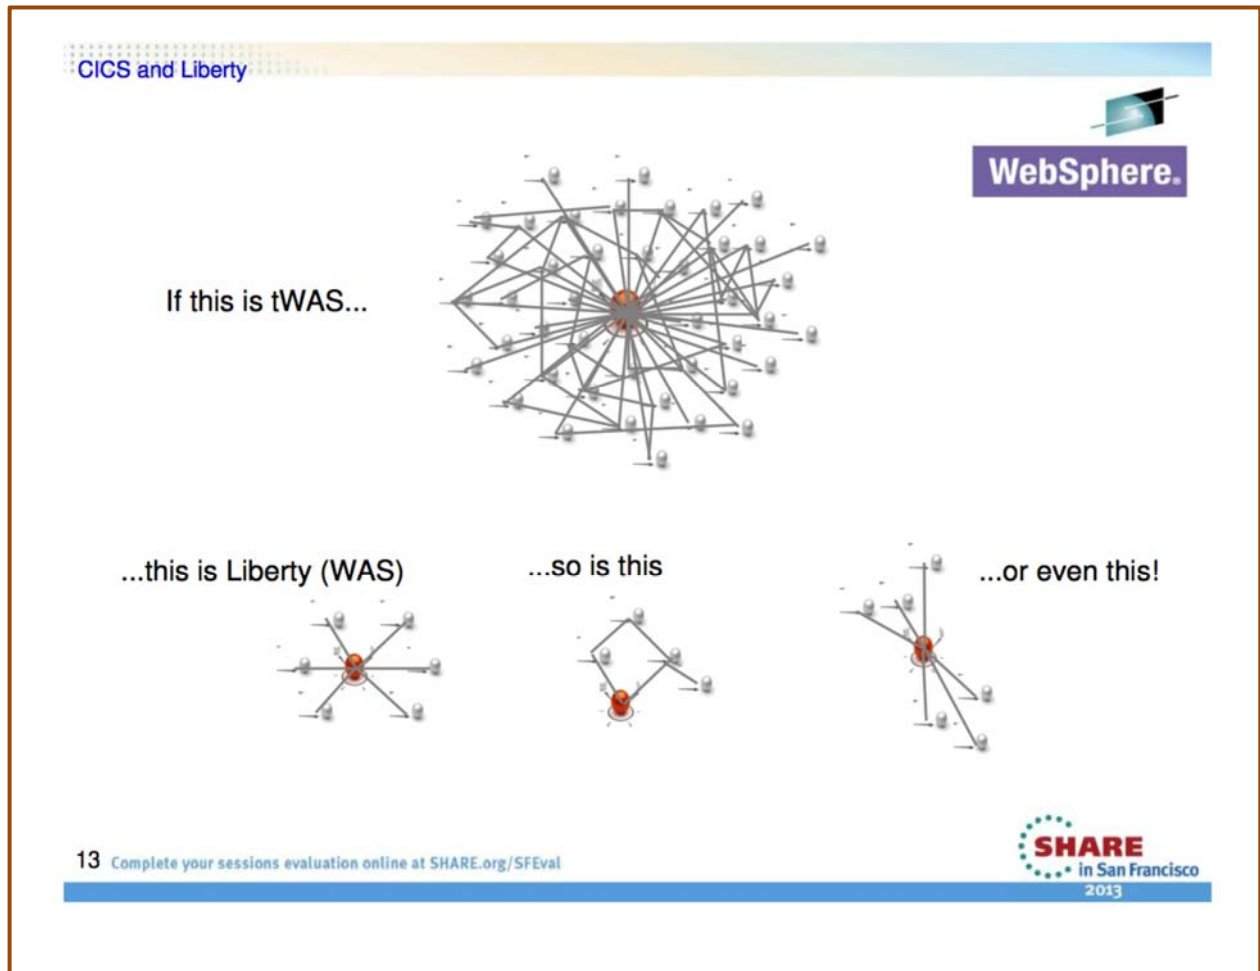
We don't plan to go into more detail at this time about the Liberty Profile, but we want you to know about it. Please see the references below:

IBM DeveloperWorks - _Introducing the Liberty Profile_, **Ian Robinson**
IBM DeveloperWorks - WASdev Community Site - www.ibmdw.net/wasdev/
IBM Redpaper REDP-4855-01 (draft 28May2013) - _WebSphere Application Server V8.5.5 Technical Overview._
IBM Redbook SG24-8076-00 (9Nov2012) - _WebSphere Application Server Liberty Profile Guide for Developers._ This also has an introductory video.
IBM Redbook TIPS1024 (10Jun2013) - _From Development to Production with the IBM WebSphere Application Server Liberty Profile._
IBM White Paper WP102110 (7Dec2012) - _WebSphere Liberty Profile for z/OS._

SHARE 2013 San Francisco - Session 12182 - **Michael Stephen** and **John Hutchinson** - *WebSphere Application Server - What's New?*

SHARE 2013 San Francisco - Session 12183 - **David Follis** - *WebSphere Application Server Liberty Profile - Rumors Dispelled.*

SHARE 2013 San Francisco - Session 12445 - **Catherine Moxey** - *Web Apps Using Liberty Profile Technology in CICS.*

Figure 8 is from the SHARE presentation by **David Follis** and shows the performance benefit of the Liberty Profile. The picture tells the whole story!

Figure 8 - Performance Benefit of Liberty Profile (© IBM)
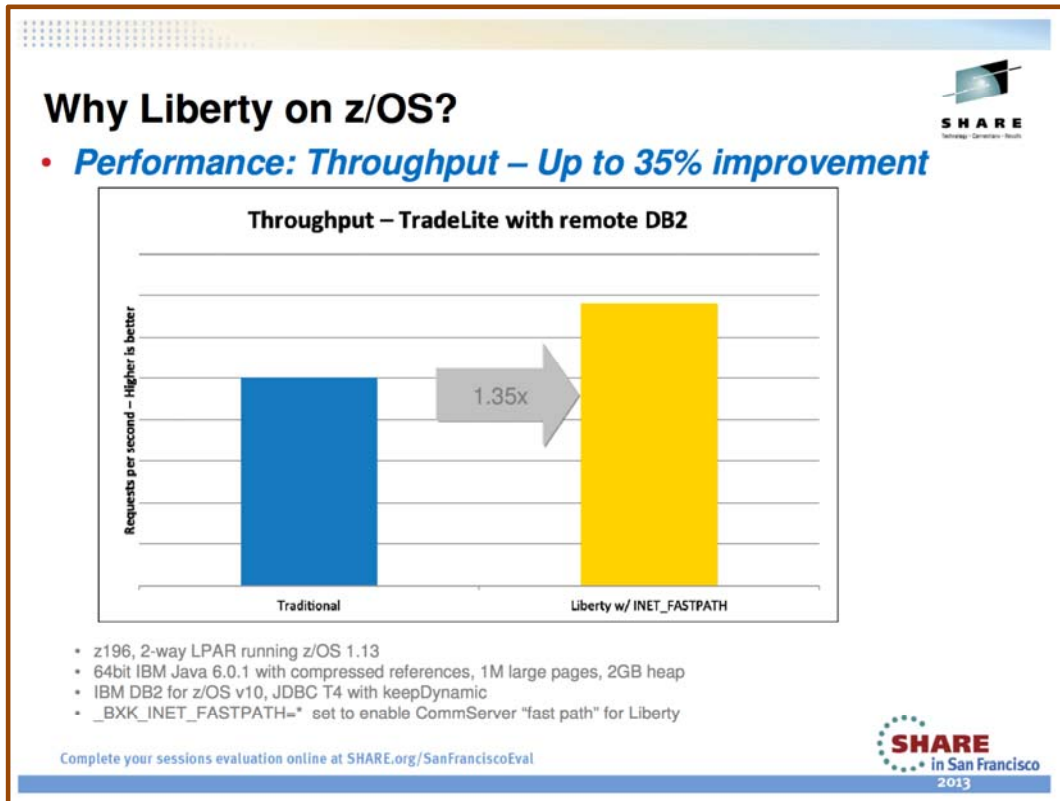
WAS Announcements

213-174 (23Apr2013) - *IBM WebSphere Application Server for z/OS, V8.5 is enhanced with additional Liberty profile programming models, security, management, and extensibility.* Also see IBM Redbook SG24-8056-00 (updated 18Mar2013) - *WebSphere Application Server V8.5 Administration and Configuration Guide.*

IBM Announcements

Here are the announcements that interest us the most.

CICS Announcements

213-141 (23Apr2013) - *IBM Explorer for z/OS, V2.1 delivers extensible work-station access to Key IBM z/OS functions.* z/OS Explorer is a z/OS Eclipse integration platform to provide workstation connectivity to z/OS datasets, zFS files, JES jobs, and output. This includes being able to do many ISPF-type functions without logging on to TSO, such as submitting jobs, editing members of a PDS or PDSE, or creating and deleting datasets. The z/OS Explorer features, that were previously available in CICS Explorer, are now available for use by other products, such as CICS administration, DB2, WebSphere MQ, and IMS. The product is also embedded within other IBM products, including CICS Explorer 5.1.1 and the CICS Tools family, IBM Rational Developer for System z, IMS Enterprise Explorer, and IBM Problem Determination Tools. The z/OS Explorer Rich Client Platform (RCP) runs on any machine that supports Microsoft Windows, Redhat Linux, or SUSE Linux. There are plenty of APIs for ISVs to exploit. There is no charge for z/OS Explorer, but you must have a license for a supported version of z/OS. It's downloadable at www.ibm.com/software/htp/cics/ibmexplforzos/ starting June 14, 2013.

213-177 (23Apr2013) - *IBM CICS Transaction Server (4.2 and 5.1) for z/OS feature packs further enhance service agility and operational efficiency by adding support for mobile devices and modern batch.* The mobile extensions add native support for the JSON (JavaScript Object Notation) data interchange format, along with runtime conversion to and from COBOL, C/C++, and PL/I language structures. The modern batch pack is a batch container in CICS to receive and process job requests along with online transactions, in order to reduce the batch window needed for CICS. Additionally, CICS TS 5.1 adds support for the JAX-RS and JSON WebSphere Liberty features. CICS TS 4.2 requires z/OS 1.11, while CICS TS 5.1 requires z/OS 1.13. The availability date was June 14, 2013. The z/OS Explorer 2.1 (described above) plug-ins for the CICS Tools Family are also included in this update.

213-168 (23Apr2013) - *IBM CICS Transaction Server for z/OS Value Unit Edition V5.1 offers a one-time-charge price metric for net new Java workloads.* This provides a substantial discount for CICS TS 5.1 when deployed on zNALC LPARs. zNALC is System z New Application License Charge, and we discussed this in *Tuning Letter 2007 No. 1* (pages 39-41).

COBOL Announcement

213-144 (23Apr2013) - *IBM Enterprise COBOL for z/OS, V5.1 allows you to generate your applications for higher levels of the z/Architecture and higher levels of*

performance optimization. This is a major update to COBOL that will appeal to most users. Here's a summary of the enhancements:

1. New compiler option ARCH to identify the z/Architecture level of z/OS that you expect to execute on. ARCH(6) for z890/z990 through ARCH(10) for zEC12 are supported.
2. Enhanced compiler option OPTIMIZE to select level of optimization.
3. Support for full AMODE(31) addressability for large data. In previous versions, the limit for each data division section and data items was 128 MB. The data division section limits were eliminated, and the limit for a data item is 999,999,999 bytes.
4. XML enhancements PARSE and GENERATE statements give you control over the form of XML documents.
5. Support for latest middleware, including CICS, DB2, and IMS.
6. Improved programming with new UTF-8 Unicode intrinsic functions.
7. Improved Pseudo-assembly listings when using the LIST compiler option.
8. Interoperability with Java 7.
9. New DWARF/Common Debug Architecture (CDA) interface. DWARF is a standard for a debugging file format and is described at www.dwarfstd.org.
10. Support for unbounded variable-length tables and groups. A new keyword of UNBOUNDED is added to the OCCURS...DEPENDING ON clause. This is especially useful for XML documents.
11. New floating comment indicator ('*>') to create a comment anywhere in the program-text area.
12. New level of SMF tracking support to reduce administrative overhead. COBOL 5.1 now creates SMF type 89 records for use in sub-capacity pricing. Previously, you had to manually indicate which LPARs supported COBOL. Now, the SMF 89 records and SCRT 21.2.0 or later (the Sub-Capacity Reporting Tool) can identify where COBOL runs. You can also disable COBOL from running on certain LPARs via the IFAPRDxx parmlib member.
13. Other new compiler options are: AFP(VOLATILE/NOVOLATILE), DISPSIGN(SEP/COMPAT), HGPR(PRESERVE/NOPRESERVE), MAXPCF, and STGOPT/NOSTGOPT.
14. Other enhanced compiler options are: NOTEST(..DWARF/NODWARF) and EXIT.
15. Addressed 24 marketing requirements.

IBM expects most existing COBOL programs can be re-compiled with no changes using the new compiler, but you should check the COBOL Migration Guide to confirm what changes, if any, are needed. The new compiler has an increase in price of about 13% for PSLC and VWLC licensing.

The z/OS COBOL website is at www.ibm.com/software/products/us/en/entecoboforzos. **Tom Ross** discussed a few of the new features in SHARE in San Francisco session [12334](12334) - *Full Speed Ahead with COBOL Into the Future.* ∎

Parmlib Changes

In our *Tuning Letter 2010 No. 6*, we started a two-part series that described parmlib changes by release, from z/OS 1.8 to z/OS 1.12. z/OS 1.13 was covered in *Tuning Letter 2011 No. 6*, pages 3-13. For Tuning Letter articles about parmlib members as a group, please see the following issues (or search our DVD archive for a particular member or parameter):

- Tuning Letter 2009 No. 6, pages 3 - 21 - *Focus: IEASYS - Part 1*. This includes general recommendations for managing parmlib members, as well as a description of how parmlib processing is handled during IPL.
- Tuning Letter 1998 No. 1 (pages 22-35) and Tuning Letter 1998 No. 2 (pages 24-38) - *Parmlib Analysis*. These two issues described changes to all parmlib members that might be needed when you change any part of your configuration including: replacing processors, changing storage, changing software releases, changing an LPAR configuration, moving applications from one system to another, changing the number of systems in a sysplex, or changing the number of users in a system.

We'll have another major article after the final announcement of z/OS 2.1, probably announced in July, to include all parmlib changes in that release. Until that time, here's an update from some APARs.

ALLOCxx

ALLOCxx members specify allocation values for the system. These members contain defaults for unit names, TIOT size, space attributes, how to handle allocation requests, and catalog error policies. Some sites do not use ALLOCxx members because they instead use installation exit routines that control allocation. The exit routines would be specified in the EXITxx parmlib member or the PROGxx parmlib member.

OPTCDB_SPLIT
- A New Function APAR was recently closed: [OA40542](#) (z/OS 1.12-1.13, 6Feb2013) - *New Function - OPTCD=B Enhancement.* This provides a new ALLOCxx parameter and a change in multi-volume tape processing.
- A new keyword OPTCDB_SPLIT on the SYSTEM parameter of ALLOCxx indicates how the DCB=OPTCD=B DD parameter should be handled when the data set is a multi-volume tape whose volumes are in tape libraries. The values are OPTCDB_SPLIT(<u>EXPLICIT</u>/CATALOG).
- Prior to this APAR, if volumes for a multi-volume tape data set resided on different tape libraries, you had to explicitly code the volsers in the JCL, with OPTCD=B. Allocation would then treat the single DD as concatenated DDs in order to process the data set correctly. After the application of this APAR, the behavior is the same if you use the default or specify EXPLICIT for this parameter in ALLOCxx. If CATALOG is specified in ALLOCxx and OPTCD=B is speci-

fied on the DD statement, then you do NOT need to specify the volsers. There are several examples of uses of this new option in the APAR, which I recommend you read.

- *==RECOMMENDATION: This will apply only to sites who have multiple tape libraries, where volumes of a data set can be cataloged in more than one library. Leave the default of EXPLICIT unless you need this facility. Then change to CATALOG and test to confirm that it is handling the data sets properly.==*

IEAOPTxx

MANAGENONENCLAVEWORK

This is from our *Tuning letter 2012 No. 4* (page 21):

> *z/OS 1.12 introduced a new parameter in parmlib member IEAOPTxx, ManageNonEnclaveWork, to indicate whether work in an address space that is not associated to an enclave be managed towards the goals of the external service class of the servant address space. Currently, this applies primarily to WebSphere Application Server (WAS) applications. The default is NO, which follows the previous behavior where the service class for the servers was only used during startup, shutdown, and recycles.*
>
> *We then recommended that, if you run WAS, you change the default value of NO to YES in three different Tuning Letters: 2012 No. 3 (page 17, WAS Recommendations), 2011 No. 5 (page 8, WLM Enclave Enhancements), and 2011 No. 1 (page 22, z/OS 1.12 Changes).* **Kathy Walsh** *of the Washington Systems Center suggested that we rescind that recommendation because it could cause some performance problems, and ==we agree with Kathy that you should leave the default of NO unless you have a specific problem to solve.==*
>
> *In Kathy's SHARE 2012 in Anaheim presentation 11463, she listed both the advantages and disadvantages of using 'YES'. Please see her session for additional information.*

A recent New Function APAR supports our changed recommendation to leave the default of NO:

OA42178 (z/OS 1.12-1.13, **FIN**, 10Jun2013) - *With MANAGENONENCLAVEWORK=YES Servers in SYSSTC May Unilaterally or Exchange Swap.* A WebSphere servant region classified into service class SYSSTC may swap. A local fix is to classify the server in a service class other than SYSSTC. FIN means that a solution is expected in a release (if any) to be available within thee years. ∎

User Experiences

This section incorporates user experiences, as well as items of the type that we previously carried under the topics of Reader Feedback, From the Forums, and Q & A. Please feel free to share anything you've found that might be of interest to our readers. Just send an email to technical@watsonwalker.com.

DB2 zIIP Eligible Time

Perseverance pays off...eventually. We're including this article for two reasons. The first is to notify DB2 10 users (or potential users) about the omission and restoration of the zIIP eligible field in DB2 10 SMF 101 records. The second is to show that if you pursue a valid requirement with IBM, it may take awhile, but it just might pay off.

In March 2011, a reader wrote to us concerned about hearing that DB2 V10 no longer populated the zIIP eligible fields in the SMF 101 records. The company was using the fields for chargeback to give users a discount for specialty engines. When I queried IBM about it, I was given the following (unsatisfactory) response:

> DB2 10 omitted the zIIP eligible time in SMF 101s

> *Rumors of the death of the information have been greatly exaggerated. The information which we can't get and can't report is the projection. For that information, customers need to use RMF. DB2 accounting includes separate times for CP and zIIP used.*

> *SG24-7892-00 (Updated 6Dec2011)* - DB2 10 for z/OS Technical Overview *on page 38 states:*

> > *DB2 for z/OS no longer provides information about zIIP eligible time that was processed on a central processor (CP) due to unavailability of sufficient zIIP resources. As a consequence the information provided by DB2 for z/OS can no longer be used to determine whether further zIIP capacity might be required for TCO optimization. To monitor zIIP eligible time that was dispatched to be processed on a CP you therefore need to use the RMF report options. In this part of this book we outline how to set up WLM and to use RMF to monitor zIIP usage for DRDA and z/OS batch workloads using RMF report class definitions.*

> > *z/OS 1.12 RMF includes information in the CPU Activity Report about how many units of work are running or waiting for a processor (CP, zIIP, or zAAP). Additionally, RMF provides this information in SMF Type 70 records. This new information is expected to be helpful for determining how much latent demand there is for processor time.*

As our reader put it - "RMF is not a charge-back tool." I strongly agree. My response to the IBM developers was this:

> It's a chargeback issue for some installations. If you charge a different price for use of a zIIP, for example, you want to keep the charges fairly consistent. If work ran on a zIIP one day, but didn't run on it another day because other work was keeping the zIIP busy, you want to charge the same for both days. Because it was eligible to run on a zIIP, it should still get the discount. That's what this site has been doing in the past, and I know other installations that do the same. Therefore, DB2 10 has re-moved capability that was previously available in DB2 9. And RMF does NOT provide the same information.
>
> In addition to chargeback, capacity planners also want to know which workloads and applications are growing their zIIP usage. Having to guess and create special report classes is not going to be a reasonable solution. They need the information at the plan and/or package level in the DB2 SMF records. The RMF type 70 record or the various service classes do not provide the detail that is required.

Our reader and his team continued to pursue the issue by submitting a requirement to IBM in late April 2011 and an APAR was created in February 2012.

> The whole specialty engine conundrum is that you need to somehow en-courage your customers to do the work to zIIP enable their workloads so that you have enough eligible workload to justify multiple engines. If the GPs have a pool of 12 engines and the zIIP only 1 or 2, the GP pool will vacuum up all the work and the zIIPs will not get driven to high utiliza-tion. Discounts directed at the application level rather than reducing the overall cost of computing could potentially provide that incentive. Since we are about to roll out DB2 V10, we are going to have to take the dis-count away from our customers before it can have much impact on be-havior. DB2 is taking away something which we are just about to ask for under IMS and CICS, where Java enablement has the same barrier. I think they have a much higher bar to jump because the CPU time and memory requirements may eat up any savings achieved on the CPU en-gine cost.
>
> Without those statistics, we will not be able to do ensure the discounted rate, making our offerings less competitively priced.

They were rewarded for their perseverance in June 2012 with a fix for APAR PM57206.

PM57206 (DB2 10, 23Jun2012) - *QWACZIIP_ELIGIBLE Serviceability Field.* The amount of time that work was eligible for a specialty processor is added back in IFCID3 in DB2 10. It contains the eligible time for distributed DBATs, paral-lel query parent threads, and zIIP eligible utilities. "All other cases of IBM spe-cialty engine offload are NOT reflected in this serviceability field." The APAR

states that it was closed as a NEW FUNCTION, but was not marked with a SPECIAL ATTENTION flag, so few people noticed this.

PM74888 (DB2 10, 19Jan2013) - *Wrong Data in SMF Record 101 for QWACZIIP_ELIGIBLE.* This occurs when running DB2 utilities.

DB2 10 Increased CPU Time

Greg Dyck, of IBM's Next Generation Systems, sent us this note to share with our readers. Thanks, Greg!

> *Some DB2 customers, depending on environmental considerations, have seen a 1-2% increase in uncaptured CPU time after migrating from DB2 V9 to DB2 10. I have also seen high (captured) RSM spin lock contention in high n-way environments. The root cause of the uncaptured time increase is a significant increase (as high as 3x) of 1st reference page faults, as reported by RMF. This will be corrected by APAR PM88804, which is currently in the process of regression testing prior to closure. It will be a good fix to install for anyone running DB2 10.*

Here's the info on the APAR:

PM88804 (DB2 10, **OPEN,** 10May2013) - *Change DISCARDDATA to be triggered by Monitor Status Rather Than Always Being Done in Certain Cases.* DISCARDDATA of real storage frames results in SRM serialization to manage the frame status and page faults for those frames which were unbacked with KEEPREAL(NO).

MSUs per MIPS

Those of you who are familiar with our CPU Charts may have seen the column called MIPS/HW MSU. The higher the MIPS per hardware MSU, the more value you are getting from a processor. Figure 9 is an extract from our last CPU Charts.

We often get questions about H/W MSUs and the MIPS/HW MSUs, and would like to add some more explanation. First, here's a description of two fields from our CPU Chart document:

H/W MSUs

The hardware (H/W) MSUs (millions of service units) rating is used to indicate the capacity of a machine. The original calculation of MSUs was based on the SU/Sec value. It represented the number of possible service units consumed in an hour (3600 seconds) and was calculated as:

$$\frac{\text{number of CPUs } * \text{ the SU/Sec value } * \text{ 3600}}{1,000,000}$$

Figure 9 - Extract of Selected Columns From Cheryl Watson's CPU Charts

Model	# of CPs	IBM PCI	Avg RNI MIPS	Avg RNI MIPS/ CP	SU/Sec	Common Name	Proc Grp	S/W MSUs	MIPS/ S/W MSU	H/W MSUs	MIPS/ H/W MSU	GA
2827-701	1	1514	1513.7	1513.7	78048.7805	zEC12	IMLC	188	8.1	281	5.4	09/12
2827-702	2	2853	2853.4	1426.7	73394.4954	zEC12	IMLC	352	8.1	528	5.4	09/12
2827-703	3	4151	4151.5	1383.8	71428.5714	zEC12	IMLC	511	8.1	771	5.4	09/12
2827-704	4	5409	5409.4	1352.3	69868.9956	zEC12	IMLC	664	8.1	1006	5.4	09/12
2827-705	5	6628	6628.2	1325.6	68085.1064	zEC12	IMLC	813	8.2	1226	5.4	09/12
2827-706	6	7809	7809.3	1301.6	66945.6067	zEC12	IMLC	957	8.2	1446	5.4	09/12
2827-707	7	8954	8953.8	1279.1	65843.6214	zEC12	IMLC	1092	8.2	1659	5.4	09/12
2827-708	8	10063	10062.8	1257.9	64777.3279	zEC12	IMLC	1224	8.2	1866	5.4	09/12
2097-701	1	902	902.1	902.1	47619.0476	z10-EC	IMLC	115	7.8	171	5.3	09/08
2097-702	2	1705	1704.9	852.4	44692.7374	z10-EC	IMLC	215	7.9	322	5.3	09/08
2097-703	3	2468	2467.6	822.5	43010.7527	z10-EC	IMLC	312	7.9	465	5.3	09/08
2097-704	4	3192	3192.1	798.0	41666.6667	z10-EC	IMLC	401	8.0	600	5.3	09/08
2097-705	5	3895	3895.0	779.0	40404.0404	z10-EC	IMLC	488	8.0	727	5.4	09/08
2097-706	6	4577	4576.7	762.8	39603.9604	z10-EC	IMLC	571	8.0	855	5.4	09/08
2097-707	7	5238	5238.0	748.3	38834.9515	z10-EC	IMLC	651	8.0	979	5.4	09/08
2097-708	8	5879	5879.4	734.9	38004.7506	z10-EC	IMLC	729	8.1	1095	5.4	09/08
2094-701	1	560	559.8	559.8	29520.2952	z9-EC	IMLC	81	6.9	106	5.3	09/05
2094-702	2	1086	1086.0	543.0	28368.7943	z9-EC	IMLC	158	6.9	204	5.3	09/05
2094-703	3	1607	1606.6	535.5	27538.7263	z9-EC	IMLC	229	7.0	297	5.4	09/05
2094-704	4	2122	2121.6	530.4	26845.6376	z9-EC	IMLC	298	7.1	387	5.5	09/05
2094-705	5	2601	2600.8	520.2	26143.7908	z9-EC	IMLC	363	7.2	471	5.5	09/05
2094-706	6	3062	3061.8	510.3	25559.1054	z9-EC	IMLC	422	7.3	552	5.5	09/05
2094-707	7	3505	3505.2	500.7	25000.0000	z9-EC	IMLC	479	7.3	630	5.6	09/05
2094-708	8	3932	3931.8	491.5	24427.4809	z9-EC	IMLC	532	7.4	704	5.6	09/05

If the result was more than one, the value was rounded to the nearest integer; if less than one, the value was rounded to the nearest tenth of an MSU. IBM initially published hardware MSUs, but doesn't any more. These MSUs had previously been used for pricing, but lately have been replaced with the software MSUs described above.

MIPS/HW MSU

This is the Avg RNI MIPS divided by the H/W MSU value. It shows the comparison between MIPS and MSUs because many installations are beginning to use MSUs instead of MIPS for capacity reporting. Notice, however, that there is quite a variation - between 5.1 and 6.3 MIPS per MSU for the older machines. The z9, z10, z114, z196, and zEC12 seem to average in the 5.3 to 5.8 MIPS per H/W MSU range.

Here is a question we recently received from a reader:

Thank you for the recent CPU capacity information. We appreciate you providing the hardware MSU/MIPS information. That was hard to come by in the last series or two of processors. Can I ask how you determined the hardware ratings? Was is through IBM's capacity formula or??? The reason I ask is I see the ratio of MIPS to hardware MSU ranges from 5.3 to 5.5 with the new processors. When I look back at the z9 series, the ratio was approximately 5.8. Has the workload mix that determines these ratings changed again or was this just a simple calculation using 5.3 or 5.5? Does that make sense? Just curious…

How we calculate the values for the HW MSUs is shown on page 8 of the CPU Chart:
(# of CPs * SU/Sec * 3600) / 1,000,000. We then divide this value into the average
MIPS to get the 5.x values for MIPS / HW MSU.
Here are the ranges of MIPS / HW MSU for each processor series:

    zEC12 - 5.3 to 5.5
    z114 - 5.4
    z196 - 5.3 - 5.5
    z10-EC - 5.3 - 5.7 (GAed in February 2008)
    z10-BC - 5.2 - 5.4 (GAed in September 2008)
    z9-BC - 5.2 - 5.5 (GAed in 2006)
    z9-EC - 5.3 - 5.9 (GAed in 2005/2006)

The GA date is important because it's when the SU/Sec is actually set. IBM runs their
benchmarks, sees how each workload runs, calculates their own 'average' (either the
MIX MIPS or the Avg RNI), and sets the SU/Sec accordingly. What changed was this:

- In 2005, IBM published two sets of ratings - single-image and multi-image.
  Prior to this, they only published single-image. Multi-image is about 5% less in
  MIPS. At this point in time, the SU/Sec was based on the higher single-image
  view, so the z9s had high SU/Sec values.

- In October 2008, IBM changed to z/OS 1.9 and only published multi-image
  mode. This essentially reduced the average MIPS, and thus the SU/Sec by
  about 5% on the z10s. So a z9 would have SU/Sec based on single image, and
  a z10 would have SU/Sec based on multi-image.

- A major change occurred in 2010, before the z196s were GAed. IBM discov-
  ered that the average MIPS (their MIX workload) was too high for most cus-
  tomers. Most customers were seeing the Low/IO MIPS. So when IBM changed
  from a MIX workload to an average RNI for their average LSPRs, the effective
  MIPS went down.

The result of all these changes means that you can't easily compare very old machines
with the newest machines. But you can usually compare models that have been GAed
within just a few years (maybe going back two generations). ■

IBM System z Batch Network Analyzer (zBNA)

This new zBNA tool from the Capacity Planning Support team in Gaithersburg, MD is designed to help you analyze the impact on batch jobs when moving them to different processors. There is no charge for the Windows tool, and support is provided as time is available. This team is the same team that provides and supports zPCR (System z Processor Capacity Reference) and zSoftCap (Software Migration Capacity Planning), two of my favorite IBM tools. I think this new tool is even niftier!

Background

First, here's a little background on why this tool is needed. You usually change CECs (i.e. machines) because you want more capacity, less capacity, or you want to divide the capacity into multiple CECs. More capacity is usually achieved with faster CPs or more CPs. Less capacity and division of capacity is usually achieved with slower CPs or fewer CPs. There are three questions with any of these changes: is there enough capacity for the work (this can be determined with zPCR), how will this impact my online systems (which traditionally like the fastest CPs available), and how will this impact my batch window (this is the reason for zBNA).

In describing the rationale for this new tool, IBM admits that due to technology limits, the days of large speed increases are now gone. While new processors might be slightly faster, you won't see these huge speed increases again (61% for z9 to z10; 54% for z10 to z196, and 26% for z196 to zEC12). Figure 10 was provided by **Harv Emery** in a SHARE 2013 San Francisco session [13077](#) (*zEC12 Hardware: I/O Subsystem, I/O*

Figure 10 - IBM System z: Design Comparison for High End Systems
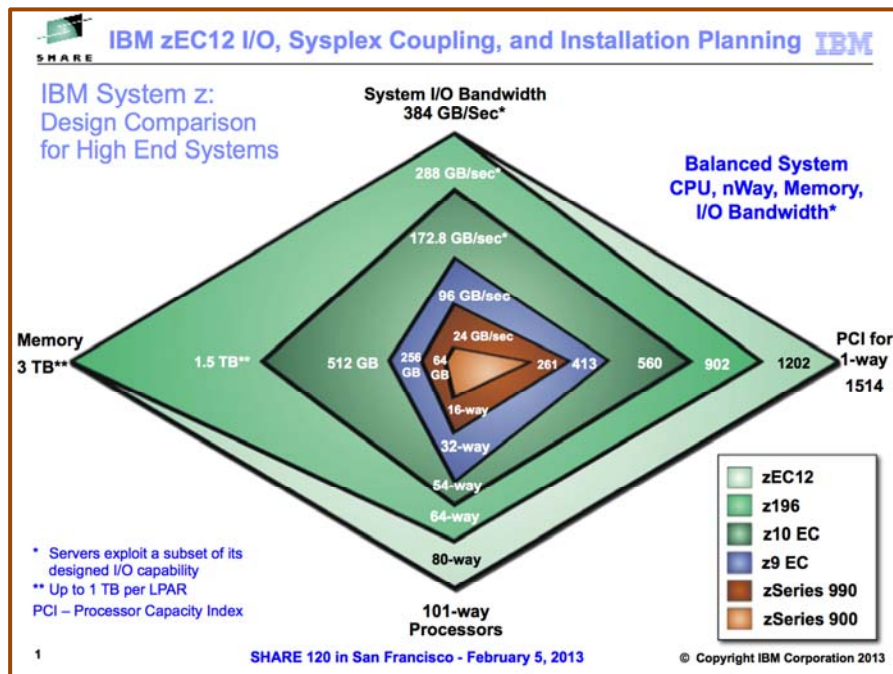
and Parallel Sysplex Coupling Features, and Installation Planning). In order to grow the capacity of new CECs, you'll probably see more of a growth in the number of CPs (bottom of graph) rather than in the speed of a CP (right side of graph).

That means you should start thinking of more parallelism in future processing. For example, you might need more CICS regions to take advantage of more CPs instead of growing your current regions. For batch work, you might want to grow your capacity with more slower CPs than fewer faster CPs. That will cause the jobs to take longer and might shrink your current batch window. But, by how much? That's where zBNA comes to the rescue, as we'll see below.

How Does zBNA Work?

As you upgrade to fewer and faster CPs, or more but slower CPs, your batch window will most likely be affected. zBNA can help you determine just HOW the window will change. In this article, I'll show you how to use zBNA, and also provide some other ways I might use this new tool.

Getting Started

You can download the zBNA tool and User's Guide from the Techdocs website. It's under the Presentations & Tools section, and is called [PRS5132](#) - *IBM System z Batch Network Analyzer (zBNA) Tool* (Updated 20Jun2013). Under the Education link at the bottom, you'll find the link to the Customer zBNA webcast delivered on June 20, 2013, a lab exercise, and some test files for the lab exercise. I recommend that you start with this exercise. It's excellent and takes you through using this new tool. And it only takes a few minutes of your time. After you're familiar with the tool, you can go on to use your own data. But *do* take time for this step!

Using zBNA

The first thing you need to do is to obtain data for your batch jobs. The z/OS job, CP3KEXTR, that has been used to extract data for zPCR, now has an additional PDS member called EXTRZBNA that is used to extract data for zBNA. It collects data from SMF types 30 and type 70 records, and writes them to two PDS members which you download to your PC. There is some date selection that you can apply to the extraction step. All work from this point is done on the PC.

The zBNA User's Guide is excellent at describing how to get started, but it still needs some updating on how to best use zBNA. I'll try to add some pointers below. The User's Guide is still being updated, so there will be more added soon. Because the User's Guide gives a summary of all of the steps and selection screens, I won't duplicate that information here, but I will show you some screens to help you understand how zBNA can be used.

After you import the data from EXTRZBNA, you can display the jobs that were extracted. You can filter them, sort them, and obtain more detail about each job. Figure 11 shows the screen that is used to filter the jobs after downloading them. (All of the ex-

amples below are using the zBNA test data provided by IBM.) This allows you to work with just certain jobs, certain classes, or certain service classes.

Figure 11 - zBNA Filter Screen
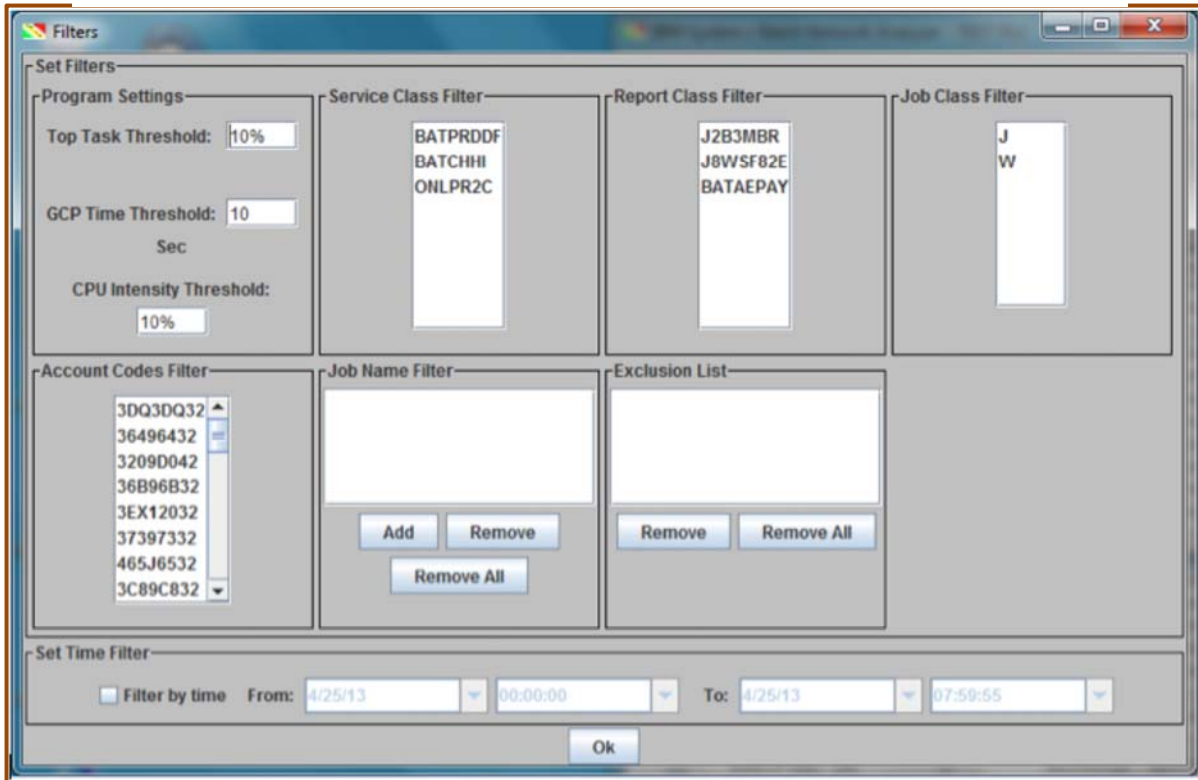
The 'Top Task Threshold' (first filter) is only available if you've applied APAR OA39629 (see *References*), which provides two new fields in the SMF type 30 record - the program name for the task using the highest percent of a CP during the interval, and the percent of the CP it used. For zBNA to be really effective, you need this APAR! For this article, I changed the filter of Top Task Threshold to 60% and saw the resulting set of jobs in Figure 12. The reason this information is important is that most batch is run under a single TCB. Thus, a task using 80% of a CP would take much longer if the speed of the CP were cut in half, and would be more affected than tasks using 5% of a CP.

I knew to use 60% in the filter because I had previously sorted all of the jobs on the Top Task % in descending order (second column from the right). (This is very easy in zBNA because all you have to do is to click on a column heading and the entire table is sorted in sequence by that column. This can be VERY useful when dealing with hundreds of jobs.) In this case, I was left with seven jobs where at least one task consumes at least 60% of a CP. These tasks can't be split across multiple CPs, and if they move to a slower CP they'll take longer. If I want to get more detail about any job, I can double-click the job to get the screen in Figure 13. In this case I selected the second job and found out when it ran, how much CPU time it took, and details about the steps.

Figure 12 - zBNA Listing of Jobs

Another thing you'll notice in Figure 12 is that I selected two of the jobs as 'Key Jobs'. Information about these key jobs will be printed with any reports, and the colors will be different in the graphs. We'll see those later.

Figure 13 - zBNA Job Detail Screen

zBNA gives you lots of options for output. You can extract the information to a CSV file, you can display a graph on the screen, you can display a Gantt chart, or you can create an html file with an explanation, the table information, and an optional graph. We've included the generated html report with the graph. (We had to divide this into

three parts for reprinting.) This html file can then be posted on an internal website, printed or sent to other staff.

The first part of the report (Figure 14) provides some guidelines for using the report and some information about the configuration. The second part (Figure 15) shows the processor the data came from plus any alternative processors that you might be considering. It also provides more details for the jobs marked as Key Batch. Finally, the third part (Figure 16) provides a graph of all of the jobs selected. I'll discuss the graph in more detail in the next section.

Figure 14 - zBNA Sample HTML Listing (Part 1 of 3)

IBM System z Batch Network Analyzer

Please Be Aware
This tool provides a means of estimating the elapsed time for batch jobs solely based on the differences in CPU speeds for a base processor and a target processor, the number of engines on each system, and system capacities. Data sharing is not considered. This tool is meant to be used to give the customer a starting point for addressing jobs that could be candidates for running on processors other than the processor they currently run on. The results of this tool must be used with caution when evaluating jobs to be moved to the target processor. There is no guarantee that your results will correspond to the projections. Many factors, including but not limited to the following, may result in significant variances between the actual and projected results:

- Different I/O configurations between base and target systems.
- Differences greater than 10% between actual processor utilization and the utilization parameters used when running the tool.
- Customer workload characteristics different from the workloads used to generate LSPR values.

IBM does not guarantee the results from this tool. This information is provided "as is", without warranty, expressed or implied. You are responsible for the results obtained from your use of this tool.

zCP3000 Power Numbers
zCP3000 capacity projections are based on zPCR algorithms that project capacity for any LPAR Host processor model with its specific LPAR configuration. For z/OS and z/VM partitions, the workload assignment is based on CPU MF data if available (preferred method); otherwise "Average" will be used. All capacity results are scaled based on the Reference-CPU setting. For this study, the Reference-CPU is set to a 2094-701 with an assumed capacity of 593.0, representing the productive capacity of a shared single-partition configuration.

Specialty Engines
The information contained in this document provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SEs only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html . No other workload processing is authorized for execution on an SE. IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Trademarks
IBM is a registered trademark of the International Business Machines Corporation in the United States or other countries and may be used throughout this publication. Complete trademark information can be found at http://www.ibm.com/legal/copytrade.shtml .

Study Specifics
This report was generated by the IBM System z Batch Network Analyzer Tool V1.0.0 on Jul 9, 2013 3:31:19 PM.

The data was extracted by CP3KEXTR06/12/13.

In this instance data was read from file C:\Users\Cheryl\Documents\zBNA Files\test1s.dat and file C:\Users\Cheryl\Documents\zBNA Files\test1s.edf.

These files are from a 2817-711 processor, in partition PROD running in a zOS partition with a SYSID of SYS1. For the time of this study this partition had a GCP logical utilization of 93.7% and the entire processor had a GCP utilization of 93.70%.

This data spans from Apr 25, 2013 12:00:00 AM until Apr 25, 2013 7:59:55 AM.

Figure 15 - zBNA Sample HTML Listing (Part 2 of 3)

The processors considered in this analysis are the following:

Note: There is no effort to determine if the alternate processor has the total capacity to run this workload. The analysis is simply comparing the single engine speed of base versus the alternate processor.

Name	Processor	Single GCP Mips	Ratio
Base (B)	2817-711	1,102	
Alternate 1 (A1)	2827-615	866	-21.4%
Alternate 2 (A2)	2827-709	1,433	30.0%

The analysis follows:

Key Batch Jobs
These are jobs that the user specifically selected for display.

There are 2 jobs in the following table.

Name	Line	Key	Job Name	Program Name	Start	End	Steps	Job Class	Acct Code	Serv Class	Duration	CPU Time	Top Program	Top Pct
B	1	X	M373XQ3		4/25/13 12:00 AM	4/25/13 1:30 AM	5	J	37397332	BATPRDDF	5,426	3,394	DSNECP10	87
A1	1	X	M373XQ3		4/25/13 12:00 AM	4/25/13 1:47 AM	5	J	37397332	BATPRDDF	6,434	4,402		
A2	1	X	M373XQ3		4/25/13 12:00 AM	4/25/13 1:17 AM	5	J	37397332	BATPRDDF	4,661	2,629		
B	2	X	M3YHK7SE		4/25/13 12:00 AM	4/25/13 1:30 AM	26	J	3YH3YH32	BATPRDDF	5,459	2,199	DSNECP10	64
A1	2	X	M3YHK7SE		4/25/13 12:00 AM	4/25/13 1:41 AM	26	J	3YH3YH32	BATPRDDF	6,112	2,851		
A2	2	X	M3YHK7SE		4/25/13 12:00 AM	4/25/13 1:22 AM	26	J	3YH3YH32	BATPRDDF	4,963	1,703		

Figure 16 - zBNA Sample HTML Listing (Part 3 of 3)

Modeling

Now it's time to do some rudimentary modeling. What will happen to these jobs when you move to other processors? For this example, the jobs had been run on a z196/711 (11-way, 1,102 MIPS per CP). The 1,102 MIPS per CP is from zBNA and is based on the configuration (number of LPARs, number of specialty processors, etc.).

If this site wants to upgrade to a zEC12, they could use zPCR to model some other options. But even using zPCR, you need to know what those options are. In our CPU Charts, we recommend that you use our Excel table of all processors, sort them by Average RNI MIPS (or IBM CPI), and look for processors in the range of MIPS you'd like. Eliminate any processors that you don't want, then look for the lowest MSU that will give you the needed capacity. I did that (see Figure 17) and found that two processors, the zEC12-615 (15-way) and the zEC12-709 (9-way) would be the closest match. I would want to check these options with zPCR before I got really serious.

Figure 17 - Extract From Cheryl Watson's CPU Chart

Model	# of CPs	IBM PCI	Avg RNI MIPS/CP	S/W MSUs	MIPS/ S/W MSU	Name
2827-708	8	10063	1258	1224	8.2	zEC12/700
2827-614	14	10336	739	1256	8.2	zEC12/600
2817-711	11	10609	965	1286	8.2	z196/700
2827-615	15	10938	730	1327	8.2	zEC12/600
2827-709	9	11137	1238	1350	8.2	zEC12/700
2827-616	16	11524	720	1397	8.2	zEC12/600
2827-617	17	12096	712	1464	8.3	zEC12/600
2827-710	10	12179	1218	1473	8.3	zEC12/700

So now I'm ready to do my modeling, so I went to the zBNA Alternates screen (Figure 18) and added two alternate processors. Then I ran the reports again to get the html report shown in Figure 14 through Figure 16.

Figure 18 - zBNA Alternates Screen

Results

Now I'm ready to talk about the results shown in the html report. Figure 15 shows that zBNA estimates that MIPS/CP for the zEC12-615 is 866 (21.4% slower), and for the zEC12-709 is 1,433 (30% faster). In Figure 16, the solid bars represent the current jobs, with pink representing the CPU time, blue representing the CPU queue time, and yellow showing the rest of the time, and therefore the duration. The other time is usually I/O. For Key Jobs, instead of yellow, the bar is red.

Below each job, you'll find the estimate for the two alternate processors. In this case, the queue time is eliminated because there are either more CPs or faster CPs. But the jobs on the zEC12-615 all take more elapsed time due to the slower CP. And, because I selected jobs that had a task ALREADY taking over 60% of a CP, I know that multi-tasking is NOT going to help these jobs, even if we add another ten CPs.

> **Start using zBNA now!**

Coming Enhancements

Kathy Walsh and **Alex Patterson** (IBM) gave an excellent webinar on June 20[th]. They provided a demo of the tool and described their expectations of how people will use this new tool. They are looking at using the SMF type 42 records (which is why we included an article on SMF 42.6 on page 25 of this issue) to obtain information on data sets and response times. They are also looking into working with job scheduler information so that they can project changes on a critical batch window. Because this is a new tool, please be sure to provide any suggestions to the team.

My wish list includes the ability to select non-batch jobs such as started task intervals. The reason is that the output from zBNA is so easy to use, it could be used for many other things (as I mention on the next page). I would also like the User's Guide updated to include more of the philosophy of how to use the tool.

References

Cheryl Watson's Tuning Letter 2012 No. 3, page 34 - *APAR OA39629.*
> ➢ OA39629 (z/OS, 26Jul2012) - *New Function to Report the Highest Percent of CPU Time Used by a Single Task in an Address Space.* This adds two new fields in the type 30 record. SMF30_Highest_Task_CPU_Percent contains the highest percent of CPU time used by any task in the address space and is calculated as 'TCB time * 100 / interval time'. The program name of the task using that percent is in the SMF30_Highest_Task_CPU_Program field.

IBM Techdocs - PRS5132 (Updated 20Jun2013) - *IBM System z Batch Network Analyzer (zBNA) Tool.* This is where you download zBNA, the User's Guide, and the test files.

IBM Techdocs - PRS4229 (Updated 5May2013) - *Data Extraction Program (CP3KEXTR) for zPCR.* (and zBNA!)

IBM zBNA Webinar (replay of June 20, 2013) - ibm-stg.adobeconnect.com/_a819987824/replaysyszbna06202013/.

Other Uses for zBNA

There are a lot of installations who would like to quickly look at SMF type 30 data, but don't have the necessary tools. zBNA is a free tool to look at key information for batch jobs. I especially like the fact that you can create an CSV/Excel file as soon as you download it. All of the columns in Figure 12 and Figure 13 are included in the CSV file. Once in a spreadsheet, you can do all sorts of analysis. These are some of the things you might want to do:

➢ create your own graphs
➢ find all of the batch jobs that use zIIPs (or zAAPs) and sort them in descending order
➢ find all the jobs for a specific application (by account code or name) and give the data to those application programmers to create a history file
➢ identify all the batch jobs in a given service class or report class
➢ identify all the batch jobs in a given job class
➢ find the jobs with the highest CPU times
➢ find the jobs with the highest EXCP times

The WSC team does such nice work, and this tool will be one of my favorites. It would be especially useful for consultants going into installations that have no tools to look at SMF type 30 data. Thanks, WSC! ■

News

This section provides interesting bits of news, important APARs, and descriptions of new documents. The parentheses after an APAR contain the applicable product releases and the date the PTFs were closed (if available) or the date that the APAR was closed (if PTFs aren't available yet), or OPEN (if it's still open). Note that when an APAR is still open, you will need to log in with your own userid and password. Please see our *Tuning Letter 2009 No. 5* newsletter that explained the anatomy and life of an APAR. The majority of these APARs were closed or opened in April or May.

New Function APARs

Regular software maintenance APARs are designed to correct defects in existing code, while new function APARs are designed to introduce new facilities between z/OS releases. Some of these new functions may provide significant performance improvements, so it's worth your time to pay attention to them. In a few cases, the "New Function" is provided to fix a bug.

Because most new functions are installed to match the current behavior (off), they will not provide any benefit until you activate them.

- OA40542 (z/OS 1.12-1.13, 6Feb2013) - *New Function - OPTCD=B Enhancement.* This provides a new ALLOCxx parameter and a change in multi-volume tape processing. Please see page **Error! Bookmark not defined.** for a full explanation.
- OA40548 (z/OS 1.12-2.1, 31May2013) - *New Function APAR to Support the Switch CUP Diagnostics Function.* This provides two functions. The first is the addition of ROUTE and HEALTH keywords that are added to the D M=DEV command to show the route through the fabric. The second is the addition of two IBM Health Checks to identify problems with the fabric. They are applicable to all switches on Brocade Fabric Operating Systems (FOS) 7.1.0c and Brocade Network Advisor (BNA) 12.0.2. See the APAR for more information on the health checks.

SMF APARs

These are APARs that can affect the SMF and RMF data that you are collecting. Many of these are not considered HIPER or performance related, so you might miss them. But they could have a significant impact on your reporting or chargeback systems.

If you have any unusual experiences with these or similar SMF APARs, please let me know. **Dr. Barry Merrill** identifies most of these APARs in his MXG newsletters (www.mxg.com), but not all of our readers monitor that site. His forum, which often

discusses various performance, capacity, and chargeback issues, is available to everyone, not just MXG customers. We highly recommend it.

Please note that there were several APARs relating to SMF problems on z/OS 2.1. I did not include them because they apply only to the Beta test sites who have easy access to all APARs on this new release.

OA41876 (z/OS 1.13-2.1, 15May2013) - *Message IFA023I Logstream is Empty Not Produced at z/OS 1.13.* A message should be written, but isn't.

OA41933 (z/OS 1.13, **OPEN**, 12Apr2013) - *MSGIEE388I Incorrect Data Set Shown.* During an SMF switch, the message showed the incorrect data set name. The data was fine but the message was wrong.

OA42006 (z/OS 1.12, **OPEN**, 19Apr2013) - *Many SMF Log Stream Offload Data Sets Are Created in a Short Time and Contains Many Duplicate Records.* In this situation, the SMF log stream was duplexed to a DRXRC staging data set which became full. Reissuing the write request resulted in duplicates to the original log. The temporary fix is to avoid filling the staging data sets. You can also try lowering the HIGHOFFLOAD parameter.

OA42312 (VSAM RLS, **OPEN**, 15May2013) - *ABEND0C4 in SMSVSAM Module IGWMROUT or IGWMRTE3 Due to Freed Block Being Referenced.* This occurs when processing SMF records using the CDATA interface.

OA42440 (VSAM RLS, **OPEN**, 30May2013) - *RMF Monitor III RLSLRU and SMF42-19 values for CPU Time Used.* Until a fix is available, you need to divide the reported value by 4,096,000. This was routed from OA40596, and is related to OA41815.

PM80172 (TCP/IP 1.12-1.13, 26Feb2013) - *AE PM68039 Fix Completion.* The UDPTERM records (SMF type 119.0) do not have the SMF119UDUCRIP or SMF119UDUCRPort set. PM68039 fixed this problem, but this APAR corrects the problem for Fast UDP sockets.

PM83558 (HLASM, 14Mar2013) - *Negative Value in ADATA Record Type 30 Operand Length.*

PM88663 (DB2 10, **OPEN**, 8May2013) - *SMF 101 Record Elapsed Times Unrealistically Large.* That's all the information that's available, but it sounds serious.

PM89899 (DB2 10, **OPEN**, 28May2013) - *Package Name Truncated to 18 Characters in IFCID110 SMF Records.* Several IFCIDs do not currently provide long name support. IFCID108, 109, 111, and 177 are others that don't provide long name support.

PM90169 (File Manager, **OPEN**, 31May2013) - *File Manager Support for Mapping SMF Records Using Segmented Template and Distributed SMF Assembler Macros.* Although we don't normally track APARs in File Manager, we thought this was important because it might be providing a low cost method of looking at SMF data.

HIPER APARs

HIPER (High Performance and Pervasive) APARs, can cause serious damage if the situation is encountered. These are not all the HIPER APARs, but we look for those that lose data, contaminate the catalog, or might pertain to many of our subscribers.

OA41029 (z/OS 1.13-2.1, **HIPER**, 8May2013) - *ABEND0C1 in IECVPST and IOS071I Idle With Work Queued Messages After a Hyperswap.* When ENABLE(PPRCSUM) is specified in parmlib member DEVSUPxx, an ABEND0C1 may occur after a successful Hyperswap was triggered by a device boxing condition.

OA41138 (DFSMS 1.12-2.1, **HIPER**, 13June2013) - *ADR789W on Logical Data Set Restore of a VSAM KSDS and Only Restores a Portion of the Data Set.* This situation occurs when a KSDS was defined with cylinder allocation that was not a multiple of its CA size. Accessing the data set and dumping it worked successfully, but restoring it failed. The solution is to fail the dump in this situation. Two other APARs were created to prevent the problem from occurring in the first place:

> ➢ OA41256 (DFSMS 1.12, 1.13, 26Apr2013) - *Improvements to IDCAMS Diagnostics to Detect non Tracks/CA ALLOC.* IDCAMS DIAGNOSE will now identify this condition. Until this is applied, you can visually identify this condition by dividing the TRACKS/CA (usually 15) into the number of TRACKS.

```
VOLUME
....
VOLFLAG------------PRIME     TRACKS/CA-------------15
EXTENTS:                                          **
LOW-CCHH---X'00120000' LOW-RBA------------0 TRACKS----------7498
HIGH-CCHH--X'0205000C' HIGH-RBA---378122239              ****
```

> ➢ OA41279 (z/OS 1.11-1.13, **HIPER**, 27Feb2013) - *Extent Not on a Cylinder Boundary During a Cylinder Allocation When VVDSSPACE is Altered and DS4VVDSA is Not On MSGADR789W.* This condition does not occur often due to the combination of situations that cause it, but this APAR will ensure that it does not occur in the future.

OA41898 (z/OS 1.12-2.1, **HIPER**, 15May2013) - *Exit CNZWTOMDBEXIT is not Called When Specified in PROGxx Parmlib Member at IPL Time.* This exit does not get called as expected after IPL, but it should.

Miscellaneous APARs

This section includes APARs that aren't provided elsewhere, but that we found interesting.

OA41979 (DFSMS 1.12-2.1, 30May2013) - *zHPF IO May Fail if More Than 255 Tracks Are Accessed on zHPF Phase 0 Capable Device.* This occurs when using FICON Express 8S channels. The request will be redriven and succeed, but message IOS050I will still be issued. This condition on IBM DS8000 devices occurs at micro code levels starting at 4.3 and below 6.2. The APAR ensures that multiple DCWs are supported before building a second DCW.

News From Techdocs

IBM's Technical Sales Library, also known as simply 'Techdocs', can be found at www.ibm.com/support/techdocs. For this issue of the *Tuning Letter*, all updates for April through June 2013 are included.

Flashes

Select "Flashes" at the ATS website for new or updated flashes.

> FLASH10815 (13May2013) - *Data Capture HIPER Fix for DB2 for z/OS.* This applies to data sharing for DB2 and recommends that users apply APAR PM84864.

> > ➤ PM84864 (DB2 9-10, **HIPER**, 12Apr2013) - *IFI306 Reads Returns the Wrong Log Range.* This could result in lost data.

Presentations

Select "Presentations & Tools" at the ATS website for the following new or updated presentations.

> PRS5117/PRS5127 - Do a Health Check on your network

> PRS1381 (Updated 31May2013) - *Getting Started with zPCR (IBM's Processor Capacity Reference).*

> PRS3279 (Updated 14Apr2013) - *IBM zEnterprise System Running DB2 for z/OS with Oracle PeopleSoft Applications.* This summarizes the value of running PeopleSoft on z/OS.

> PRS3961 (Updated 17May2013) - *Security Workshop: WebSphere Application Server for z/OS Version 8 Handouts for WSW08 Wildfire Workshop.*

> PRS4229 (Updated 15May2013) - *Data Extraction Program (CP3KEXTR) for zPCR.*

> PRS4686 (Updated 10Jun2013) - *WAS z/OS Version 8 Configuration Spreadsheets.*

> PRS4840 (Updated 17May2013) - *Generate a Key Check Value for KCV4ISPF.*

> PRS4944 (Updated 10Jun2013) - *WAS z/OS Version 8.5 Configuration Spreadsheets.*

> PRS5117 (1Apr2013) - *Network Part of ATS Parallel Sysplex Health Check.* From Techdocs: "Advanced Technical Skills (ATS) provides a Parallel Sysplex Health Check to some customers. While the ATS Health Check depends upon the z/OS release, implemented features, current applications, and any specific customer comments about any recent problems, this presentation may be of interest as a starting point for the z/OS Communications Server piece of the ATS Health Check. It lists the current z/OS Health Checker checks that check z/OS Communications Server components and some Best Practices for z/OS Communications Server."

> PRS5127 (14Apr2013) - *z/OS Networking Health Check: Is Your z/OS Network Up-to-date?* This is an excellent presentation given by **Gwen Dente** at the last

SHARE conference and shows how to use tools to analyze TCP/IP. This can be used with PRS5117 above.

PRS5128 (14Apr2013) - *IPv6 on z/OS: Avoiding Common IPv6 Implementation Mistakes.* This is another presentation by **Gwen Dente** from the last SHARE. I haven't met Gwen yet, but I'm impressed her work.

PRS5132 (20Jun2013) - *IBM System z Batch Network Analyzer.* Please see our description of this neat new tool on page 48.

PRS5135 (30Apr2013) - *Understanding and Analyzing TS7700 Performance - Accelerate with ATS.* This provides a replay of a webinar given on April 18, 2013 which provided a description of performance controls on the tape subsystem TS7700.

PRS5139 (21May2013) - *IBM Tape Library Monitoring and Reporting.* This presentation describes multiple tools used for tape library monitoring and reporting, such as Tape Systems Reporter (TSR), SNMP Audit Logging, and the Crossroads ReadVerify Appliance (RVA).

PRS5152 (27Jun2013) - *An Overview of Mainframe Security for Non-Mainframe Personnel.* I think the title of this excellent presentation by **Mike Smith** is slightly misleading. I know a lot about mainframes, but this presentation taught me several new things, especially about the interaction of various security products.

Technotes & Tips

Select "Technotes & Tips" at the ATS website for the following new or updated tips:

TD104300 (Updated 21May2013) - *An Example of a NetView for z/OS HTTP Client.*

TD105477 (Updated 16Apr2013) - *TS7700 Virtualization Series - VEHSTATS Decoder.*

TD105671 (Updated 16Apr2013) - *DS8000 Host Ports and Installation Sequence.*

TD106002 (Updated 25Apr2013) - *WebSphere AppServer V8.5 - SMF Browser for Job Usage Data.* The SMF type 120, sub-types 9 and 20 for Job Usage can now be displayed in the Detailed report of the SMF Browser sample program. From Techdocs: *"WebSphere Application Server Version 8.5 for z/OS provides the equivalent 'Batch' function that was previously part of the 'IBM WebSphere Extended Deployment Compute Grid for z/OS' product. One of the z/OS exclusive capabilities is to provide for Charge-back accounting records for 'Job Usage.' This document provides guidance and examples for recording the job usage statistics, and displays some sample programs."*

TD106098 (Updated 10Jun2013) - *Understanding WAS z/OS Thread Tuning Considerations.* "Performance tuning may be considered a balancing act between many competing factors. This article provides a description of the relationship between requests, threads, processors, and the behavior of the application itself."

White Papers

Select "White Papers" at the ATS website for the following new or updated documents.

WP100810 (Updated 3Apr2013) - *A Synopsis of System z Crypto Hardware.*

WP101092 (Updated 13May2013) - *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide.*

WP101094 (Updated 25Apr2013) - *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide.*

WP101249 (Updated 9Apr2013) - *Configuring and Clustering IBM WebSphere Portal Enable for z/OS.*

WP101374 (Updated 5May2013) - *WebSphere Application Server for z/OS - Dispatch Timeout Improvements.* WAS V7 (and, therefore, WAS 8.0 and 8.5) provided new configuration options and facilities for avoiding time-out related abends. This paper describes these new facilities.

WP101382 (Updated 12Apr2013) - *IBM Virtualization Engine TS7700 Series Best Practices - Cache Management in the TS7720.* From Techdocs: *"Similar to the TS7740, whose capacity is limited by the number of physical tapes, the TS7720's capacity is limited by the size of its cache. The customer in both cases needs to manage keeping the amount of data stored below the limits imposed by the number of physical tapes or cache size. The intent of this document is to make recommendations for managing the cache usage in the TS7720. It details the monitoring of cache usage, messages and attentions presented as the cache approaches the full state, consequences of reaching the full state and the methods used for managing the amount of data stored in the cache."*

WP101490 (Updated 13May2013) - *WebSphere z/OS Optimized Local Adapters.*

WP101532 (Updated 4Apr2013) - *Why WebSphere Application Server for z/OS.*

WP102247 (10Apr2013) - *IBM Virtualization Engine TS7700 Release 3.0 Performance White Paper.*

WP102267 (13Apr2013) - *Implementing the Output APAR (PM74923) Enhancements in WebSphere Application Server on z/OS.* From Techdocs: *"While the APAR addresses getting the files into the filesystem, this document describes an approach that enables the viewing of those files, associated with specific cells, nodes, clusters or servers to be viewed by designated groups of people in a simple manner, without providing access by those groups to the physical file system components. This is done by exploiting the functionality of the APACHE HTTP server included in WebSphere on z/OS."*

> PM74923 (WAS 7.0, 8.0, 8.5, 20Jun2013) - *Servant Output Redirected to Unix System Services File System Could Collide When Servants Start at the Same Time.* Without this APAR, WAS running multiple servants that start at the same time might get output from multiple servants written to the same file. This APAR provides a new WebSphere environment variable to associate the output to a specific path name.

WP102279 (13May2013) - *Big Data Introduction on the Why and What Does it Mean for Sales and Marketing Teams.* This paper describes how companies can use analytics to provide tools to their internal sales and marketing people.

WP102281 (13May2013) - *Couple Data Sets Best Practices and Avoiding Disasters.* This should be required reading in any data center that uses couple data sets. There are excellent recommendations here.

Conferences

Select "Auxiliary Material" at the ATS website. This listing contains additional IBM presentations that were given at the February 2013 SHARE conference in San Francisco.

TC000140 (16Apr2013) - **Mike Bonett** - *Using NetView for z/OS for Enterprise-Wide Event Management and Automation.*

TC000148 (2Apr2013) - **Kenneth Blackman** - *IMS Database Access Enhancements.*

TC000149 (3Apr2013) - **Kenneth Blackman** - *IMS Enterprise Suite 2.2 SOAP Gateway.*

Social Media

WP102281 provides great CF Best Practices

This is a new section we've added to our *Tuning Letter.* We follow several technical blogs, twitter feeds, and Facebook/LinkedIn pages. We'll let you know when we find something especially interesting (and educational).

SHARE Blogs

Most of the SHARE blogs require a login, although you don't need to be a SHARE member to obtain a login. Just go to www.share.org to obtain a userid and password.

Marna's Musings (**Marna Walle**, migration guru extraordinaire)
Get your DASD installation requirements in early! (17May2013) - See page 22.
What's your favorite dynamic function in z/OS??? (19Jun2013) - Please respond to her survey.
Thinking of Brad Carson (23Jun2013) - In memory of a wonderful SHARE volunteer.

MVS Program Blog (Hosted by **Mary Anne Matyaz**, SHARE MVSE Project Manager):
When is 20 milliseconds too long? (23Apr2013) - This is a very valuable article to help your understanding of parallel access volumes (PAV).
My Blue Monday (13May2013) - This is great for some laughs.
A Fascinating History of JES2 (17May2013) - This history by **Jack Schudel** really is fascinating.
SHARE in Boston Preview (14Jun2013) - This lists several of the really wonderful presentations that will be sponsored by the MVS program at the next SHARE conference in Boston.

MVS EWCP Project Blog (Hosted by **Meral Temel**, SHARE EWCP Deputy Project Manager)
*Knowledge Is The Treasure! And SHARE IS good at providing it * AS ALWAYS * since 1966 !!!* (21Feb2013) - Meral's reflections on SHARE in San Francisco.

> *A check list that leads us to a Mainframe platform which performs as perfect as it can?* (15Mar2013) - This presents Meral's performance checklist. I also learned about some excellent presentations on the zNextGen project files page (www.share.org/p/do/si/topic=39&source=7):

> ➢ **Norman Hollander** (12Mar2013) - *New Age Performance Management: Is It Worth It?* - Norman makes some excellent points.
> ➢ **Vit Gottwald** (26Jun2013) - *z/OS UNIX for all, Part 1.*
> ➢ **Glenn Anderson** (5Jun2012) - *Understanding Your Workload Manager (WLM) Service Policy and Associated RMF Data*.

SHARE President's Corner

> *IBM's zAware: The Tip of the Iceberg* (16Apr2013)

> *A Conversation with Doug Balog, General Manager for System z, IBM: Part 1* (17Apr2013)

> *A Conversation with Doug Balog, General Manager for System z, IBM: Part 2* (18Apr2013)

> *The Case for Big Data Analytics – And the Infrastructure to Support It* (24Apr2013)

> *The Mainframe: Bridging One Generation to Another* (30Apr2013)

> *Don't Believe the Myth-information about the Mainframe: Part 1* (7May2013) - This and the next two parts are great articles relating to the strengths of the mainframe over distributed platforms.

> *Don't Believe the Myth-information about the Mainframe: Part 2* (9May2013)

> *Don't Believe the Myth-information about the Mainframe: Part 3* (15May2013)

> *BlackBerry Vs. Apple Vs. Android: Should the Mainframe Team Know (or Care) What Employees Use?* (13Jun2013)

Other Interesting Blogs

Ray Mullins blog - www.catherdersoftware.com/blog.html
> *A little caveat if you are playing with DCB= on LINK/LOAD* (22Apr2013)
> *Pitfalls of reusing storage with ATTACH(X)* (25Jun2013)

Willie Favero (IBM Data Warehouse for System z Swat Team, DB2 SME) - DB2 blog at it.toolbox.com/blogs/db2zos/
> *A little more about real time statistics* (2Apr2013) - Part 2 in a series on real time statistics. (I listed Part 3 in the previous Tuning Letter.)
> *When it rains, it pours… Redbook time…* (8Apr2013) - Willie will be working on a new Redbook about analytics. Keep posted.
> *Stuff that could be interesting to attend* (11Apr2013) - An excellent list of webinars, conferences, and publications related to DB2.
> *Do you find a Redbook useful?* (27Apr2013) - Helpful tips on using Redbooks.
> *What I should have said today at IDUG about EXPLAIN and DB2 10 for z/OS* (1May2013) - Lots of information about EXPLAIN.
> *Let's review DB2's sort via my blog* (2May2013) - This contains links to ten of his previous blogs about DB2's sort.
> *DB2 10 and DB2 9 for z/OS Product Publication refreshes for April 2013* (3May2013) - This gives direct links to the refreshed manuals.

May 2013 (RSU1304) service package has been tested and is now available (9May2013)

Red Alert: Potential loss of data for DB2 v10 users on z/OS 1.12 and 1.13 re-leases (29May2013) - Also see page 69 where we announced this in our *Cheryl's List 165.*

It's June 7, 1983 and IBM introduces three new licensed programs - Database 2, DXT, and QMF - Happy 30th Birthday DB2 (3Jun2013) - Willie's memories include the original announcement of DB2.

APAR Friday: A couple of HIPERS and Red Alerts (7Jun2013)

APAR Friday: Part 2, a few more that might be of interest (7Jun2013)

DB2 Analytics Accelerator for z/OS V3.1 prereqs (14Jun2013)

APAR Friday: Special edition; a friends suggestion (19Jun2013)

DB2 10 for z/OS and DB2 9 for z/OS Product Publications have been refreshed for the month of June 2013 (28Jun2013)

Martin Packer -

www.ibm.com/developerworks/mydeveloperworks/blogs/MartinPacker/?lang=en

Martin Packer hits two home runs! See his blogs on 4/8 and 6/24

How Many Eggs In Which Baskets? (8Apr2013) - This is a very important discussion relating to the number of CICS regions you run, especially before you migrate to the newest release of CICS TS (5.1).

Analysing a WLM Policy - Part 1 (1May2013) - Drawing out the hierarchy.

Analysing a WLM Policy - Part 2 (1May2013) - Setting goals.

Playing Spot The Different With WLM Service Definitions (20May2013) - How to compare service definitions using XML.

Discovering Report Class / Service Class Correspondence (22May2013)

REXX That's Sensitive To Where It's Called From (26May2013)

Data Collection Requirements (1Jun2013) - This describes that type of data that Martin asks for in order to do an analysis, but it also shows the type of data that you want to make sure you keep.

Dragging REXX Into The 21st Century? (7Jun2013)

Recent Conference Presentations (24Jun2013) - Martin has always been generous with his presentations, and the three that he provides here are outstanding as usual. *Time For Dime* provides a selection of data in memory techniques; *Life And Times Of An Address Space* shows how much you can learn about an address space; and *DB2 Data Sharing Performance* is one of the best presentations I've seen lately on DB2 data sharing. Thanks for sharing, Martin!

Why Do I Program In <insert language name here>? (30Jun2013) - A history of his programming languages, which brought back (un)fond memories. My favorite item was a link to an article I love titled "On 'Geek' Versus 'Nerd'" at slackprop.wordpress.com/2013/06/03/on-geek-versus-nerd/.

Timothy Sipples of IBM Singapore writes on the *Mainframe Blog* at main-frame.typepad.com/blog.

California Keeps its Mainframes (Thank Goodness) (16Feb2013) - A sad tale of attempting to change platforms. They should know better.

Unplanned 3 Hour Global Mainframe Outage Scheduled (1Apr2013) - Tongue in cheek - notice the date!

The PC Market is Shrinking Fast (11Apr2013) - Did you know that mobile phones are literally more popular than toilets? This is an interesting take on the demise of PCs. I especially appreciated a link to a ZDNet article with the graphic in Figure 19 (www.zdnet.com/idc-global-pc-shipments-plunge-in-worst-drop-in-a-generation-7000013839/).

IBM Announced 1Q2013 Earnings: zEnterprise Shines (19Apr2013)

Big Data, Big Tape (28Jun2013)

Figure 19 — Falling PC Market (© ZDNet)
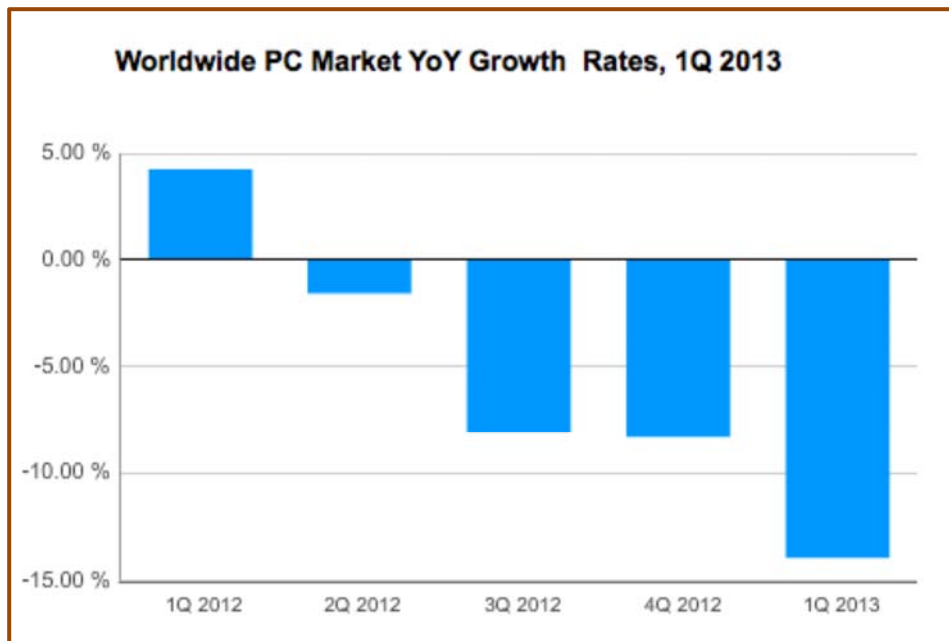
IBM Mainframe Insights provides blogs from a variety of IBM specialists. In the last couple of months, you can find blogs on BCPii (see page 4), zFlash, Redbooks, system automation, Rational development, z/VM, and Tivoli products. Here are a few that should be of interest to our readers:

System z eLearning for IBM Customers (26Mar2012) by **Scott Harris** – this describes some free online courses.

Mainframe's Role in the Cloud (3Apr2013) by **Bill Meinhardt**.

Enterprise COBOL for z/OS V5.1 at Innovate2013 (16May2013) by **Roland Koo** (see page 39). He points to another good blog from **Kurt Bittner** called *COBOL Application Development: (Still) Not Dead Yet* (3May2013). I personally never thought it was even sick!

Big data - past, present and future (18Jun2013) by **Mark Simmonds**. This is announcing a July 11 webinar on Big Data that you can replay when you have time.

Publications

Redbooks

The following IBM Redbooks and Redpapers are either newly published (-00 extension) or are updates to previous editions. There are between 30 and 80 publications added or updated each month, but we've picked the ones that are most applicable to our readers. Redbooks can be found at www.redbooks.ibm.com.

REDP-4460-00 (16May2016) - *Synchronizing IBM RACF Data by Using IBM Tivoli Directory Integrator.* This shows how to synchronize a RACF user id, password, and password phrase data between z/OS and z/VM.

REDP-4855-01 (25Jun2013) - *WebSphere Application Server V8.5.5 Technical Overview.*

REDP-4870-01 (28May2013) - *WebSphere Application Server: New Features in V8.5.5.*

REDP-4875-00 (Updated 3Apr2013) - *Performance Implications of Cloud Computing.* This provides multi-platform considerations, and provides an excellent description of service level agreements (SLAs) for cloud providers. Here are two fun points from this Redpaper:

> ➢ Up to 80% of data today is unstructured content, including email, video, and images.
> ➢ Medical images take up 30% of the available worldwide storage. The field of medical imaging is driving breakthroughs in diagnosis and treatment in medicine. As a result, there is exponential growth in the number and size of digital medical images. Medical images that were two-dimensional and 1 MB in size a few years ago are now typically four dimensional and 1 TB in size.

REDP-4893-00 (Updated 21May2013) - *IBM SmartCloud: Building a Cloud Enabled Data Center.*

REDP-4938-00 (4May2013) - *Rethink Your Mainframe Applications: Reasons and Approaches for Extension, Transformation, and Growth.*

REDP-4953-00 (20Jun2013) - *Modernize Your Application Infrastructure with IBM CICS Tools V5.1.*

REDP-4974-00 (3Jun2013) - *Artificial Intelligence: Learning Through Interactions and Big Data.*

REDP-4975-00 (19Apr2013) - *The Interconnecting of Everything.* This is a fun read that makes you appreciate the mainframe. One quote: "*The industry predicts that, by 2020, possibly 50 billion devices will be connected, which is 10 times the number of all current Internet hosts, including connected mobile phones.*"

REDP-5008-00 (29May2013) - *Maximize IBM System z ROI with Open Cloud Management and Policy-based Optimization.*

REDP-5026-00 (Updated 3Jun2013) - *Unlock Big Value in Big Data With Analytics.*

REDP-5032-00 (21Jun2013) - *The Reality of Rehosting: Understanding the Value of Your Mainframe.* If anyone in your installation is considering moving work from a mainframe to a distributed platform, rush this Redpaper to them. It shows several examples where the cost to stay on the mainframe is much

less. And there is a pointer to an IBM "Eagle Team" for a free assessment of your installation.

SG24-5250-09 (4May2013) - *IBM System Storage Solutions Handbook.*

SG24-6105-02 (Updated 26Jun2013) - *VSAM Demystified.*

SG24-6374-08 (31May2013) - *GDPS Family An Introduction to Concepts and Facilities.* This the 9th revision of this handbook on GDPS (Geographically Dispersed Parallel Sysplex).

SG24-6787-06 (Updated 5Apr2013) - *IBM System Storage DS8000 Copy Services for IBM System z.*

SG24-7280-03 (28Jun2013) - *Server Time Protocol Planning Guide.*

SG24-7281-03 (28Jun2013) - *Server Time Protocol Implementation Guide.*

SG24-7380-01 (28Jun2013) - *Server Time Protocol Recovery Guide.*

SG24-7615-01 (Updated 12Jun2013) - *WebSphere Application Server V7 Administration and Configuration Guide.*

SG24-7721-05 (Updated 24Jun2013) - *IBM System z Personal Development Tool: Volume 1 Introduction and Reference.*

SG24-7722-05 (18Jun2013) - *IBM System z Personal Development Tool: Volume 2 Installation and Basic Use.*

SG24-7723-05 (Updated 28Jun2013) - *IBM System z Personal Development Tool: Volume 3 Additional Topics.*

SG24-7859-02 (31May2013) - *IBM System z Personal Development Tool Vol. 4 Coupling and Parallel Sysplex.*

SG24-8008-00 (23Apr2013) - *Using zEnterprise for Smart Analytics: Volume 2 Implementation.*

SG24-8034-00 (28May2013) - *IBM zEnterprise EC12 Configuration Setup.* Almost 700 pages covers everything you need to know about configuring your new EC12.

SG24-8049-00 (Updated 3Apr2013) - *IBM zEnterprise EC12 Technical Guide.*

SG24-8066-00 (11Apr2013) - *Real-time Fraud Detection Analytics on IBM System z.* This new Redbook should be of interest to banking and insurance companies where payment fraud is especially high.

SG24-8069-00 (Updated 2Apr2013) - *Secure Messaging Scenarios with WebSphere MQ.*

SG24-8093-00 (24Apr2013) - *Improving Productivity with IBM ISPF Productivity Tool V7.1 for z/OS.*

SG24-8116-00 (23May2013) - *New Ways of Running IBM z/OS Batch Applications.* This is part of a four-volume series on running hybrid batch.

Redbooks Web Docs, Product Guides, and Solution Guides are small documents (usually less than ten pages) that provide product overviews, solution overviews, or technical hints and tips. The link to all these publications is www.redbooks.ibm.com/redbooks.nsf/tips. Here are the System z publications through March of this year:

TIPS0967 (5Apr2013) - *Automated Conversions to IBM DB2 for z/OS.*

TIPS1024 (10Jun2013) - *From Development to Production with the IBM WebSphere Application Server Liberty Profile.* See our article on the Liberty Profile on page 35. ■

Cheryl's List

This is a summary of the last transmissions sent to subscribers of our free electronic Cheryl's List. We've eliminated any sections printed in a previous newsletter. Past issues of Cheryl's List can be obtained in full at www.watsonwalker.com/archives.html. To sign up for the list, go to www.watsonwalker.com and fill out the form under "Cheryl's List."

Cheryl's List #165 – June 4, 2013

1. About *Cheryl Watson's Tuning Letter 2013 No. 1*
2. IBM Red Alerts

1. About Cheryl Watson's Tuning Letter 2013 No. 1

The 68-page 2013 No. 1 Tuning Letter was emailed to paid subscribers on May 21, 2013. You may visit our website at www.watsonwalker.com to obtain subscription information and the table of contents.

2. IBM Red Alerts

IBM Red Alerts should be tracked and monitored by someone at each installation. There are four new alerts:

2013.03.25 - *PTF UK91435 (APAR PM79520) is required for DB2 10 for z/OS customers in a Data Sharing environment.* This addresses a potential data loss during data sharing.

2013.05.29 - *Potential loss of data for DB2 v10 users on z/OS 1.12 and 1.13 releases.* Data loss can be seen in some instances during Media Manager IO recovery, especially users of zDMF running zHPF. See APAR OA42277 (28May2013, HIPER) and the next two alerts for more information because there is a problem with OA42277.

2013.05.31 - *Potential loss of data for DB2 v10 users and IMS FastPath DEDB users on z/OS 1.12, 1.13 and 2.1 releases.* The same APAR is referenced, but the alert indicates that the same problem can occur for IMS FastPath DEDB or PDSE data sets. See the next alert also.

2013.06.03 - *An error has been discovered in a Pre Req for Red Alert APAR OA42277.* IBM's recommendation is to NOT install the PTFs for APARs OA42277 and OA39870 (26Apr2013, New Function). This Red Alert also indicates a problem with MQSeries and other users accessing Extended Function VSAM data sets greater than 2 GB.

If you don't get these Red Alerts automatically, you can sign up to receive them at www.software.ibm.com/webapp/set2/sas/f/redAlerts/home.html. ∎