



interskill
learning

INTERSKILL MAINFRAME TRAINING NEWSLETTER

Aug 2013

Inside This Issue:

Interskill Learning Planned New Curriculum and Updates for 2013	2
Highlights of Interskill Learning Releases from 2012	3
Are YOU Prepared For Disaster Recovery?	4
What makes a great Data Center employee?	6
Vendor Briefs	8
Tech-head Knowledge Test	11
Learning Spotlight – Data Center Storage	11
Takuzu Puzzle	12
<i>technical:</i> Learning Assembler: Where Do You Start?	13
<i>technical:</i> Working With Large Quantities of SMF Records	17
<i>opinion:</i> Do We Really Need to Compress Datasets?	20
<i>management:</i> How Many Software Products Are There?	22

Welcome to the Interskill Mainframe Quarterly

Interskill Mainframe Training Newsletter is an e-Zine published by Interskill Learning, which provides world-class elearning Mainframe training programs for the Information Communications and Technology industry.

As a subscriber to Interskill Mainframe Training Newsletter you will receive an edition several times per year which will include z/OS related articles for technical experts and management, and information on upcoming releases of Interskill's Mainframe curriculum.

Interskill Learning Planned New Curriculum and Updates for 2013

Storage Fundamentals for Data Centers

Managing the Storage Network

Parallel Sysplex

Unix System Services

CA 1 Tape Management

CA Workload Automation CA 7 Edition

CA Workload Automation Restart Option for z/OS Schedulers (formally CA11)

CA ACF2

CA SYSVIEW Performance Management

CA OPS/MVS Event Management and Automation

CA Easytrieve

CA Endevor Software Change Manager

Disaster Recovery

Performance Monitoring

Coaching and Mentoring

Linux on System z

DASD Volume Initialization

z/OS Version 2.1

IBM Power Series – iSeries

- Security Implementation
- Journal Management
- Storage Management
- Logical Partitioning and Virtualization
- Backup and Recovery
- Problem Determination and Prevention in Production Support

Highlights of Interskill Learning Releases from 2012

Interskill Learning releases for 2012 included the following:

- Cloud Computing for Data Centers
- Introduction to Mainframe Security
- The z/OS Management Facility
- The complete z/OS Series updated to V1.13 including:
 - z/OS
 - z/Enterprise
 - JCL
 - TSO/ISPF
 - JES2
 - Console Simulations
 - SDSF
 - z/OS Systems Programming Fundamentals
 - Assessments
- A new z/OS Utilities series with new courses on VSAM, General Data Set Utilities, and Data Utilities
- New courses in the Power series for IBM i Operators and IBM i System Administrators including:
 - Monitoring and Managing IBM i Workloads
 - Introduction to IBM i for System Operators
 - IBM i System Administration Fundamentals
- A new series of courses on Project Management



a number of scenarios to ensure that everyone is aware of what to do.

- Review and provide feedback on testing or actual disaster recovery events. This could include technical details outlining a better approach to data recovery, or as general as not having enough people trained in what was regarded as a simple recovery process.

Your organization will also be aware that in many disaster situations you may not be available, either because you can not physically access the work environment (weather) or remotely (power or telecommunication outage). They also understand that personal priorities such as attending to family and home matters will often take precedence over work.

Are you prepared?

Most people believe that they will never be involved in a disaster recovery situation even though they may undertake planning and testing of disaster recovery plans in the workplace. If you think you are well prepared, see how you go answering the questions below (It may prove interesting to present your answers to your boss to see if you agree!)

- Are you adequately skilled across at least two other Data Center job roles?
- What action should you take if one of the disasters mentioned previously occurs and:
 - you are at home?
 - you are at work, and the most senior onsite person?
- What would you do if you were the senior person on-site during a disaster and the majority of staff wanted to leave to attend to family/home issues?
- Are you aware of all the internal and

external data compliance requirements in relation to data security and access?

- What would you do if during a disaster the person with authority directed you to perform tasks that were in direct contrast to what was in the business continuity plan?
- How would you handle the situation if you were responsible for communicating the status and activity of your organization during a disaster, but the general telecommunication lines were down?

The questions above only represent a subset of the types of decisions that you may be called upon to make during a disaster and will give you a good idea on whether you are personally prepared for such a situation.

In our next newsletter we will continue our focus with the disaster recovery theme and rather than address the obvious disaster readiness requirements, we will look at some real life issues that you need to think about. Until then, stay safe.



Greg Hamlyn, a senior technical course designer and developer with Interskill, has over 25 years experience working in the mainframe and z/OS data center environments, having

performed in numerous operational roles ranging from computer operator to national mainframe technical training manager for a large Government agency. He has presented at AFCOM and has written a number of articles relating to training methods and trends and everything that is z/OS.

You can contact him at g.hamlyn@interskilllearning.com



technical knowledge who have shown that they can pick up new technical information quickly.

You should consider the following in relation to your own technical skills and knowledge requirements:

- Don't rely on the organization to supply all your technical learning requirements. Some Data Center staff are quite happy to sit back and just go about their day-to-day business, and wait for management to decide the training (if any) they need to attend. These people are likely to be left stranded if there is a shift in IT focus away from their core tasks. You need to look at any opportunities that result in learning (not necessarily training courses) and present them as opportunities to your boss.
- You need to employ a method that you can use to keep abreast of latest products/trends. While your work may require you to be fighting fires all day, you still need to put time aside to identify trends and opportunities.
- Presenting findings to your boss/management indicating how the latest technological advances will improve the business is a great way of showing that you are adding value to the organization.
- Join specialty technical groups with the same interest. This will allow you to build your own network, which can be very beneficial if like a lot of IT areas you do not have any internal staff that you can bounce ideas off, or help resolve a problem.

What makes a great Data Center employee?

By Greg Hamlyn

Today's IT specialist needs more than just technical smarts. I have met and worked with a wide range of Data Center people over the last 25 years (the stories alone I am saving for my memoirs), the majority of which had good technical knowledge. But is that enough today to be considered a valuable employee?

In this article we look at the skills and personal traits that can make you a great Data Center employee.

1. Technical

It is logical that having technical skills and knowledge is a must for anyone wanting to work in a Data Center. Having said that, many organizations today are focusing less on employing people with a deep specific skills set, preferring those with a broader

2. Communication

You can have the greatest technical skills in the world but if you are unable to get your message across, or understand what work a

customer requires then no-one will want to deal with you.

A person that has technical knowledge and communication skills is worth their wait in gold. These people ask the right questions, clarify requirements and listen carefully to what customers (internal and external) are saying. They are able to present a need to the organization (whether this is another data center group, or management), and the solution.

Communication skills are also important if your organization has implemented a coaching or mentoring program. If you are chosen as a coach or mentor then you will need to impart your skills and knowledge to other people, something that many technical people find difficult.

3. Enthusiasm

You probably know or work with someone that just wants to keep the status quo. This type of person greets new ideas with skepticism and they are normally downplayed to show their negative side.

This is not a great scenario for people working with this type of person, especially if they are a manager. It dampens enthusiasm, often to a point where some possible beneficial IT solutions may not even be investigated. Organizations want to keep people that are innovative and enthusiastic and often place this trait higher than technical skills.

2. Openness

Some Data Center staff I have worked with throughout the years like to keep all of their knowledge to themselves. This could be so that they are seen as invaluable and

no-one can replace them (or the fact that they aren't doing too much and don't want anyone to find out). Organizations can certainly benefit by sharing their knowledge so you should jump at the opportunity to share your knowledge and in the process look to gather some yourself.

Ways that this can benefit you and the organization:

- Cross-training provides a multi-skilled workforce for the organization, ideal in the situations where people leave the organization, or move into another position internally. It could also provide an opportunity for you to move into an area you are more interested in, or that has a longer career span.
- During disaster recovery, you will never know what staff will be available, so knowing more is going to be beneficial to you and the organization.
- It is well noted that many experienced Data Center personnel are approaching retirement age. This huge bank of skills and knowledge doesn't need to be lost if proper succession planning is put in place (of course this requires that the person who is retiring is willing to impart this information).

Summary

What constitutes a great Data Center employee will probably vary between organizations, but what is clear is that versatility is becoming more preferable to pure technical expertise. When an organization has a limited number of staff and budget to go with it, they will want to be able to spread them as the need arises and still have the type of employee that can contribute new ideas. If you have all of these attributes, you can feel safe that you will be in demand.



Greg Hamlyn, a senior technical course designer and developer with Interskill, has over 25 years experience working in the mainframe and z/OS data center environments, having performed in numerous operational roles ranging from computer operator to national mainframe technical training manager for a large Government agency. He has presented at AFCOM and has written a number of articles relating to training methods and trends and everything that is z/OS.

You can contact him at g.hamlyn@interskilllearning.com

Vendor Briefs

As a result of our newsletter going biannual there has been quite a lot of vendor activity since the beginning of the year. Unless you have been living under a rock, you would not have escaped the hype surrounding cloud, analytics, mobile and the Devops movement. So, it is not a surprise that many vendor software products are focusing on these areas, making things easier, and faster (in some cases much faster). On top of this, a new zEnterprise has been born (the announcement coincidentally fell on the same day as the future King of England's birth) and the release of z/OS 2.1 is just around the corner...technology never stands still long enough for us to learn everything about it!

IBM

zEnterprise

As I am writing this, IBM has just released their latest business class mainframe server offering, the BC12. This is the successor to the z114 and as you would expect there are the usual performance improvements....but not too many surprises for those that have been working alongside the EC12.

	z114	BC12
Maximum processing capability	3100 MIPS	4958 MIPS
Processors	3.8 Ghz	4.2GHz
Maximum cores	14	18
Maximum number of configurable cores	10	13
Maximum number of capacity settings	130	156
Maximum system memory	248GB (+8 GB for the Hardware System Area)	496GB (+16 GB for the Hardware System Area)
Entry level	24, 192	48, 384

IBM's **Flash Express** and the System z Advanced Workload Analysis Reporter (zAware) that were integrated into the EC12 have also found their way into the BC12. The Flash Express feature is an I/O card that as you might have guessed provides flash memory, and is used

primarily for paging activity and SVC dumps. zAware is a system monitoring tool with a difference. It is designed to identify what it considers to be normal system behavior, which it does by analyzing system messages over a specified period. So, it is a self learning tool that once up to speed can quickly recognize any system bottlenecks, failures or just unusual messages or their frequency, and report on them immediately.

There are too many other enhancements to mention here, so check out IBM's website and the zEnterprise BC12 redbook for more information on these and the following:

- zEDC (new data compression card)
- 10GB Ethernet RoCE (memory to memory communication between CPCs)
- IBM Mobile Systems Remote (mobile/tablet access to the HMC using an app, which allows remote monitoring and management of BladeCenters and Rack Mounted systems)

IBM is pricing its hardware and software offerings associated with the BC12 competitively especially in relation to running Linux (there is a Linux only based model of the BC12 available) and running cloud and analytics applications. A great opportunity for medium sized businesses looking to upgrade or consolidate its x86 sprawl, or new businesses requiring reliable, powerful and secure services to offer new customers. At a \$75,000 starting price I might start looking at one for the home office!

IBM SmartCloud - DB2 Analytics using BLU Acceleration

IBM's Smartcloud offerings are growing almost daily, providing products and services for organizations looking to take advantage of private, public and hybrid cloud environments. One of IBM's latest releases in this area is around the hot topic of analytics

- IBM Smartcloud DB2 Analytics using BLU Acceleration. BLU's combination of memory optimization, data storage techniques and hardware exploitation provides speeds up to 25 times faster when dealing with the analyzing and reporting of analytical data..

IBM MobileFirst – Rational Test Workbench

Earlier this year IBM launched MobileFirst, aimed at providing organizations with the capability to move into the mobile world by providing solutions related to: mobile application development, management of mobile devices linked to the business, mobile security, and mobile usage analytics. Let's face it if your organization doesn't use mobile applications to allow some customer communication or interaction, then you are going to be left behind.

IBM's latest release in this area is the Rational Test Workbench, which provides application developers with automated processes for testing highly complex and integrated mobile applications. This is an area that can be very time consuming – creating test scenarios using hardware, software and cloud based variables, emulating workloads accurately whilst needing to support current standards and protocols.

z/VM 6.3

Tweaks to this system allow it to exploit the latest enhancements to the zEnterprise EC12 and BC12 systems. Real memory support up to 1TB for those running very large virtual machines is now provided while there are a number of performance/efficiency improvements such as z/VM HiperDispatch, which manages z/VM LPAR placement and processor usage. Combined, these items now provide customers with the ability to consolidate up to 520 virtual servers in a single footprint.

IBM Enterprise COBOL for z/OS 5.1

Long live COBOL! With an estimated 200 billion lines of COBOL code still in existence, IBM has recently freshened its capabilities with the release of IBM Enterprise COBOL for z/OS 5.1. A new compiler option (ARCH) allows you to identify the hardware that the code will run on, and the compiler will optimize the code for that environment. For example, the ARCH(10) option produces code that utilizes instructions available on the zEnterprise EC12. XML processing enhancements and support for Java 7 is also provided in this release and shows IBM's continued push towards modernizing COBOL capabilities.

Upcoming releases

The much anticipated z/OS 2.1 will be made available at the end of September, 2013.

CA Technologies

CA Chorus 3.0

CA Technologies' integrated mainframe management product, CA Chorus, has recently been expanded with the release of new network and system capabilities. These changes provide the ability to display a wide range of mainframe-related performance information in a dashboard format, which can be customized and shared between colleagues.

Compuware

Abend-Aid

For many mainframe application developers, Compuware's Abend-Aid has often been there to assist with identifying and analyzing faults and abends, whether it is in PL/I, COBOL or Assembler. In July, Compuware announced enhancements to Abend-Aid making it easier to be used by staff unfamiliar with legacy systems, shortening

the time it takes to analyze and resolve problems.

APM for Mainframe

Compuware's APM for Mainframe (released late 2012) is an end-to-end transaction performance monitoring solution capable of tracking transaction activity across browser, distributed and mainframe platforms. Compuware have recently enhanced this product to extend its support for CICS to include CICS Transaction Gateway (CTG) and CICS Web Services (SOAP).

Global Executive Cost of IT Failure Survey

A recently presented Compuware sponsored survey highlighted the long-term and short-term costs to business when IT fails. It also discusses the difficulty in measuring this information as it will be different depending on the criteria used (which is slightly different for all of these surveys I have read).

The magic figure for short-term costs..... around US \$10.8 million. Check out the link provided here for details about this survey.

http://www.compuware.com/en_us/about/press-releases/2013/7/short-term-price-tag-per-it-failure-averages--10-8m--according-t.html

Rocket Software

One of the hot trends at the moment is identifying how to turn all your valuable data into actionable intelligence. One issue is likely to be that your mainframe data may be stored in a number of different formats (flat files, hierarchical databases, relational databases).

Well, Rocket Software is here to help with their recent release of Rocket z/SQL. This product provides organizations with access

to non-relational mainframe data through standard SQL calls by performing all the in-place mainframe data transformation requirements in the background.

Brocade

Network and storage specialists, Brocade has recently announced the release of their latest Storage Area Networking (SAN) solution – Fabric Vision. This hardware and software technology allows data transfer rates of 16 Gbps. The solution provides administrators with the tools to more quickly monitor and identify potential data flow issues, displaying this information through a customizable health and performance dashboard.

Brocade has also released a new high density SAN switch – the 6250. This device supports 96 Fibre channel ports providing today’s highly virtualized, cloud enabled data centers with the throughput required to keep pace with data access needs.

Tech-head Knowledge Test

With every release of this newsletter a mini online test will be provided of a product that you are likely to be using in the workplace. You may want to challenge your work colleagues or just want to confirm your own knowledge!

The test for this newsletter focuses on the use of SDSF and consists of 15 questions. Click the link below to start.

[Monitoring and Managing System Activity Using SDSF](#)

Learning Spotlight – Data Center Storage

Incredibly, 90% of all data in the world has been created in the last two years, with expectations that the total amount of digital data will reach 40 zettabytes by 2020 (this is forty billion terabytes). The majority of this data will be stored in a data center environment and will need to adhere to compliance and security requirements. It is little surprise then that the importance of data center storage has seen resurgence, and today is considered a “hot topic” for conversation around corporate boardrooms.

With advances in storage hardware and software technology moving just as quickly, it is important for data center technical staff to have an understanding of these components and how they can be leveraged to provide organizations with their needs.

The module that we have provided for you here is [Overview of Data Center Storage](#) which describes the storage needs facing data centers and provides an overview of the types of hardware and software that can help provide the solution.

We hope you enjoy it.

Takuzu Puzzle

Looking for something different now that you have mastered Sudoku? Well, how about a version of this popular puzzle for mainframers where instead of using 1 to 9, the puzzle consists of binary 0s and 1s....introducing Takuzu.

Here are the rules you need to follow:

1. Each box should contain a 0 or a 1.
2. You are not allowed to have more than two 0s or 1s next to each other
3. Each row and each column should contain an equal number of 0s and 1s.
4. Each row and column should be unique.

Each Takuzu puzzle has only one solution.

Good luck.

							1		0
1							1		
	0			0					0
		1	1		0				
					0				
			0				0	0	
0									
			1					1	
1						1		1	
							0		

technical: Learning Assembler: Where Do You Start?

By David Stephens, Lead Systems Programmer at Longpela Expertise

So you want to learn to program in assembler. Congratulations, and welcome to the select few. Most programmers are terrified of assembler language, and stick to more comfortable high-level languages like Java, C and COBOL. But in the mainframe world, there are times when a problem needs assembler. And by learning to program in assembler, you're going to gain other benefits: the basics of dump analysis, a better understanding of how programs work, and even hints on improving program efficiency.

The problem is that for the beginner, learning assembler is hard. Hard, but not impossible. Twenty-odd years ago I did it with less resources than are available today. And I'm not alone. So if I were starting over again, how would I learn assembler today?

Step 1: Get the Tools

First of all, you need somewhere where you can create, assemble and run your programs. The best place is a test z/OS system. However if this isn't an option, all is not lost. You can use the excellent and free z390 Assembler Emulator from www.z390.org.

The rest of your tools are books:

- IBM Principles of Operation – I call it POPs. This is a scary looking book describing the internal workings of the System z processors. But don't be afraid. It lists all the assembler instructions,

and what they do. Once you get used to the format, this is the best assembler instruction reference. You can find it on the front page of the z/OS Infocenter.

- Bill Quall's [Mainframe Assembler Programming](#). Originally written by Bill in 1998, and now available as a free download. This is a great introduction to assembler programming. Bill refers to the PC/370 emulator in this book – this is now the z390 emulator mentioned above.
- HLASM Programmers Guide – you will use this together with POPs above. POPs details instructions to the mainframe processor. This book explains about assembler language, and how to use the High Level Assembler. Also available from the HLASM section in the z/OS Infocenter.
- MVS Programming: Assembler Services Guide and the two MVS Programming: Assembler Service References . These are three books that describe the z/OS system services you can call from assembler. Again, found in the MVS section in the z/OS Infocenter.
- DFSMS Macro Instructions for Data Sets – this describes the assembler macros essential for access VSAM and non-VSAM datasets. Found in the DFSMS section in the z/OS Infocenter.

Step 2a: If You Can, Do a Course

Here's where the problems start. Ideally, a great first step in learning assembler would be to do a course. In the past a few vendors offered a 5 day course to introduce assembler programming, and get you started. Today, most of these have dried up, though [The Trainers Friend](#) and [Verhoef](#) still advertise classroom-based assembler courses.

If you're lucky enough to get to a [Share](#)

conference, then attend their Assembler Bootcamp: a series of five sessions introducing assembler run every conference.

All the above options assume you have a travel budget. For most of us, this isn't possible. [Interskill](#) offer a great range of online assembler courses you can do anywhere, from introductory level up to advanced concepts such as cross-memory and 64 bit programming. If you have access, these are a great place to start. [zNextGen members](#) are eligible for free courses, and Interskill will offer an assembler course in September. [Marist College's IDCP](#) also hold assembler training classes that can be studied remotely.

Another alternative is to find someone who knows assembler, and is willing to be your mentor. Longpela Expertise offers a similar service through our [Systems Programming Mentoring training](#).

Step 2b: Read

If Step 2a doesn't work for you, it's not all over. Here's how you start.

1. Work through the five Enterprise Server Intro to Programming – Assembler PowerPoint presentations from IBM. This provides an easy-to-digest introduction to assembler.
2. Read the first five chapters of the High Level Assembler Language Reference. This builds on what you've already seen in the PowerPoint presentations in step 1.
I know, there's a lot of reading here. But before you can program in assembler, you need to know some of the basics about memory, registers, hexadecimal and other concepts. Once you have these, you're ready to get into it:
3. Assemble your first program. You may think that your first program will be a 'Hello World' program called from TSO, but this isn't as simple as you'd think. And few assembler programs run in TSO. So a better platform to start with is batch. Modify the following JCL to suit your site, and run it:

```
//JOB1 JOB (ACCT),'ASSEMBLE',CLASS=A,MSGCLASS=X
/* Assemble
//C EXEC PGM=ASMA90
//SYSLIB DD DSN=SYS1.MACLIB,DISP=SHR
//SYSUT1 DD DSN=&&SYSUT1,SPACE=(4096,(12,12),,,ROUND),
// UNIT=SYSDA,DCB=BUFNO=1
//SYSPRINT DD SYSOUT=*
//SYSLIN DD DSN=&&OBJ,SPACE=(3040,(40,4),,,ROUND),
// UNIT=SYSDA,DISP=(MOD,PASS),
// DCB=(BLKSIZE=3040,LRECL=80,RECFM=FB,BUFNO=1)
//SYSIN DD DATA,DLM=QQ
*=====
* ASSEMBLER PROGRAM
*=====
TST1 AMODE 31
```

```

TST1 RMODE ANY
TST1 CSECT
    BAKR R14,0          SAVE CALLERS ENVIRONMENT
    LR  R12,R15         LOAD OUR ADDRESS IN R12
    USING TST1,R12      USE R12 AS BASE REG
    LA  R15,53          LOAD 53 AS RETCODE
    PR                      RETURN TO CALLER
    END                  END OF PROGRAM

QQ
//*
//* Link-Edit
//L  EXEC PGM=HEWL,PARM='MAP,LET,LIST,NCAL',COND=(8,LT,C)
//SYSLIN DD DSN=&&OBJ,DISP=(OLD,DELETE)
//    DD DDNAME=SYSIN
//SYSUT1 DD DSN=&&SYSUT1,SPACE=(1024,(120,120),,,ROUND),
//    UNIT=SYSALLDA,DCB=BUFNO=1
//SYSPRINT DD SYSOUT=*
//SYSLMOD DD DSN=MY.LOAD.DATASET(TST1),DISP=SHR
//* Execute
//R  EXEC PGM=TST1,COND=(8,LT,L)
//STEPLIB DD DISP=SHR,DSN=MY.LOAD.DATASET

```

Don't forget to change the load library (MY.LOAD.DATASET) to your own load library. This job assembles, binds and executes a simple assembler program that returns the number 53. It's a great first step.

4. By now you've assembled and run your first program. You're away. Go to Bill Quall's book and start reading. Tweak the program in step 3 as you work through.

Step 3: Code

The only way to really learn assembler is to write assembler. So write assembler programs. You could write programs as you work through Bill Quall's book. Or you could write the following programs in order, building on the simple program in Step 2b.

1. Return the number $(10+15-45)*2$ to introduce mathematical options.
2. Return a number entered as a parameter (EXEC PGM=TST1,PARM='22' in the jobs in Step 2a) to introduce parameter handling.
3. Return the number of letters of a string input to introduce strings and loops.
4. Output a parameter entered to OPERLOG/SYSLOG to introduce z/OS service routines, and in particular, WTO.
5. Write a string into a sequential dataset to introduce datasets and allocation.
6. Write some information into a VSAM dataset to introduce VSAM.
7. Abend the program using the ABEND macro to introduce errors and error handling.
8. Insert the following code into your program:

XR R1,R1
BR R1

Try to explain what this does, and why. This will introduce you to addresses, and memory management.

9. Output the current day and time to introduce data handling, and some more system routines.
10. Output the name of the job executing – see our Control Blocks for Beginners article for some sample code. This introduces control blocks, mapping DSECTs, and internal z/OS structure.

If you successfully managed to create these 10 programs, you're well on your way.

Step 4: Research

Once you get confidence, start reading and researching how better to program in assembler. Here are some good places to start:

- [IBMs Enterprise Server Intro to Programming – Assembler](#) has more advanced assembler programming PowerPoint presentations.
- [Considerate Programming: Reducing the Maintenance Costs of Commercial Quality Code](#) –David Cole, Colesoft, 2003 – a great discussion on how to write good assembler programming.
- [Eight Ways to Improve Your Assembler Programs](#) – David Stephens, Longpela Expertise, 2013.
- [Understanding Re-entrant Programming](#) – David Stephens, Longpela Expertise, 2010
- [Using USINGS](#) – David Cole, Colesoft, 2005
- [Coding Assembler for Performance - David Bond, Tachyon Software, 2007](#)
- [Why Assembler is a 21st Century Language](#) – Kristine Harper, NEON

Enterprise Software, 2007

- [A tutorial by Jim Morrison](#) on the CBT website. In particular, the section on POPs will help you wrestle with the Principles of Operation manual.
- [SimoTime](#) have a suite of sample JCL and programs that can help you write, assembler and run your first assembler programs

Also look for other assembler programs, and see how they work. IBM provides many in sys1.samplib. One of my favourite sources for assembler is the brilliant CBT website.

Step 5: Keep Programming

Many people do a quick assembler course, and let it lapse. Fast forward five years and they've forgotten almost everything. So keep on using assembler. Program, debug, and explore.

Twenty-odd years ago I taught myself assembler in a similar way to what I've outlined here. And I found it hard work. But it was definitely worth it, and today I love assembler. Anyone working with me will recognise when I've seen a chance to use assembler by the large grin on my face. I hope you enjoy it as much as I do. Good luck in your assembler adventure.

Source:

LongEx Mainframe Quarterly - August 2013, Retrieved from <http://www.longpelaexpertise.com.au/ezone/LearnAssembler.php>

technical: Working With Large Quantities of SMF Records

By David Stephens, Lead Systems Programmer at Longpela Expertise

In the last edition, I wrote an article [that described how I would start a CPU reduction project](#) by producing a table of the highest CPU consuming programs during a peak 1-2 hour period. This table is created by processing SMF Type 30 interval records, which introduces a problem. One SMF Type 30 record is produced by every address space every recording interval (usually 15-30 minutes). So there are a lot of these records to work through. What's worse, these records are amongst the millions of SMF records produced constantly by z/OS. So this processing will take time and CPU seconds. Lots of it.

This is a very common problem when processing any SMF records. From producing security related reports to DASD performance or DB2 capacity planning. When you work with SMF records, you will almost always need to manage and process very large quantities of SMF records. So let's look at some tricks to make our lives easier.

Processing Raw SMF Records

z/OS continually writes SMF records to in-memory SMF buffers. These buffers are then regularly written to either VSAM datasets or log streams by z/OS (depending on how the systems programmer has set it up). Every site will have regular jobs that empty these VSAM datasets/log streams, archiving the SMF data to disk or (more likely) tape. So it sounds easy doesn't it? To get our SMF records, we just scan through these archived SMF records. And products like Merrill's MXG, IBM's TDSz and CA-MICS help us out. They provide tools and facilities to map each SMF record, and extract the information we need. So if I'm at site with MXG, I can extract the SMF Type 30 records I need using a job like:

```
//SAS          EXEC SAS,
// CONFIG='MXG.SOURCLIB(CONFIG) '
//SASLOG      DD SYSOUT=*
//SASLIST     DD SYSOUT=*
//WORK        DD LRECL=27648, BLKSIZE=27648, SPACE=(CYL, (150, 100))
//SOURCLIB    DD DSN=MXG.SOURCLIB, DISP=SHR
//LIBRARY     DD DSN=MXG.FORMATS, DISP=SHR
//SMF         DD DISP=SHR, DSN=SMF.ARCHIVE
//SYSIN       DD *
%INCLUDE SOURCLIB(TYPE30);

DATA JOBS1;
  SET TYPE30_V;
  WHERE (TYPETASK='JOB' OR TYPETASK='STC') AND CPUTM > 1 AND
  HOUR(SMFTIME) > 9 AND HOUR(SMFTIME) < 11;

RUN;
```

Those familiar with SAS will see that I'm working through an SMF archive (the dataset SMF.ARCHIVE – either on tape or disk), and putting a list of all SMF Type 30 records produced between 9am and 10am into the JOBS1 SAS file.

When working with SMF records, the secret is to work with as few as possible. And this is what I'm doing in this job. If you look more closely, you can see that I'm only listing jobs and started tasks (no USS processes, or APPC address spaces), and skipping every SMF Type 30 record with less than one second CPU time (There will be many records with 0 CPU if the task has not been dispatched. I'm not interested in those).

However I'm doing this after the %INCLUDE SOURCLIB(TYPE30). This means that MXG has already moved SMF Type 30 records into SAS files before my JOBS1 data step starts. To limit the records that MXG will process, you can use the MXG macro facilities.

Now, most sites I've seen archive their SMF records into daily tape datasets. So if you were to run this job against the SMF archive for one day, you'll be working through every SMF record for that day. Or in other words, you'll be waiting a couple of hours, with a corresponding CPU usage bill. There are a few ways to reduce the CPU overhead, but it's never going to be cheap. An excellent reason to run these jobs out of peak-periods: usually overnight.

If this is an ad-hoc job, then this is no problem. However if this job is to run regularly, then it's going to be expensive. A classic example of this is your regular IBM SCRT job that process SMF Type 70 and 89 records.

Limiting Raw SMF Records

If a site is going to regularly process specific SMF records, it makes sense to create a separate SMF archive holding only those SMF record types. This way any job only has to work through the SMF record types they need, rather than scanning through them all. This is also important if you need to archive some SMF records types for long periods for compliance (like RACF Type 80). IBM helps us out here. The IBM utility that archives SMF records: IFASMFDP (for SMF VSAM) and IFASMF DL (for log streams) can be setup to do just this. For example, look at the following JCL:

```
//STEP1      EXEC PGM=IFASMFDP
//SYSPRINT  DD  SYSOUT=*
//DUMPIN    DD  DISP=SHR,DSN=SYS1.MAN1
//SMFOUT    DD  DSN=DUMP.SMF.ALL,DISP=MOD
//SCRT      DD  DSN=DUMP.SMF.SCRT,DISP=MOD
//SEC       DD  DSN=DUMP.SMF.SECURITY,DISP=MOD
//SYSIN     DD  *
            INDD(MAN,OPTIONS(ALL))
            OUTDD(SMFOUT,TYPE(Ø:255))
            OUTDD(SCRT,TYPE(7Ø,89))
            OUTDD(SEC,TYPE(8Ø))
```

This job dumps all SMF records into the dataset DUMP.SMF.ALL. Additionally, SMF records

needed for the IBM SCRT (Type 70 and 89) go into DUMP.SMF.SCRT, and RACF records (Type 80) go into DUMP.SMF.SECURITY. The good news is the additional work writing to additional SMF dump datasets doesn't incur a large CPU overhead. There are also a few ways to further reduce CPU overhead of SMF dump jobs.

You can go further and split up SMF records by application. So CICS1 data would go into one SMF archive; CICSB into another. However this takes some assembler programming skills to create exits for IFASMFDP/IFASMF DL. separate SMF archive.

Summarizing SMF

Another options is to avoid accessing the raw SMF records altogether. This can be done by creating databases holding summarised data. For example, summarising SMF Type 70 records by hour, or by day. So instead of working through every record, you can just access the summarised results. This approach is done by most sites when handling performance and capacity planning records, such as SMF Type 70/72, 110 (CICS) and 101 (DB2). These databases can be stored in DB2 (for TDSz), SAS files (for MXG and CA-MICS), or any other database system you choose. TDS, MXG and CA-MICS all include features to produce these summarised databases, which are often called Performance Databases, or PDBs.

These PDBs are great, and usually satisfy most needs. However I find that I'm regularly going back to the raw SMF records. Usually this is because I need information from a non-summarised record (such as Type 30 for CPU usage, Type 80 for security, or Type 14/15 for a dataset open), or I need more detail.

Processing SMF On a Workstation

Even after getting my Type 30 records, the chances are that I'll have tens of thousands of them to process. In the old days, I would sit down and code SAS statements to produce the tables and graphs that I want. And this is still an option. But tools like Microsoft Excel produce results much faster. The trick is to get the data to them.

Let's take my Type 30 extract job above. What I tend to do is create a CSV file with the Type 30 records. I can do this in SAS by adding this code to the bottom of the job:

```
ODS _ALL_ CLOSE;
ODS CSV FILE=CSV RS=NONE;

PROC PRINT;
  VAR JOB DATE SYSTEM
  CPUSRBTM CPUTCBTM CPUTM PROGRAM CPUIFETM CPUZIETM;

RUN;

ODS CSV CLOSE;
```

This will take the SMF records from the SAS file and output the figures I want to a CSV file. Transfer this to my PC, open with Excel, and I'm away. For smaller sites, this works great.

However I find for larger sites, there is still too much data to process with Excel. In these cases, I fire up the free MySQL database I have installed on my laptop. I'll import the CSV created above to MySQL using SQL like:

```
LOAD DATA LOCAL INFILE 'C:/SMF30V.csv' INTO TABLE SMF30 FIELDS
TERMINATED BY ',' ENCLOSED BY '@' LINES TERMINATED BY '\n'
(job, datetime, system, cpu_srb, cpu_tcb, cpu_total, program,
cpu_zaa_eligible, cpu_ziip_eligible);
```

I can then create subsets of the data using SQL statements, and import them into Excel for processing.

If this all seems too hard, then the Australian based Black Hill Software can help. Their software product Easy SMF for z/OS does all of this for you. Download the raw SMF, and Easy SMF will process it, and produces graphs or tables. You can also import this data into Microsoft Excel for your own processing. The disadvantage is that Easy SMF only supports a subset of SMF records.

Conclusion

If you start working with SMF records, the chances are that you will be handling large quantities of records. The good news is that there are ways to ease the pain of processing these records.

Source:

LongEx Mainframe Quarterly - February 2013, Retrieved from

<http://www.longpelaexpertise.com.au/ezine/LargeQuantitiesSMF.php>

opinion: Do We Really Need to Compress Datasets?

By David Stephens, Lead Systems Programmer at Longpela Expertise

In many sites that I work, I see my clients compressing datasets. And they do this for two reasons. Some compress datasets before transmitting them somewhere else - and this makes sense. Others automatically compress datasets on disk. This doesn't make as much sense.

Many years ago, DASD space was expensive, and CPU seconds were cheap. So many sites decided to save money by compressing some of their disk datasets and databases using products such as Shrink/MVS (now CA Compress), Infotel InfoPak, or BMC Data Accelerator. And these products were great. They would transparently compress datasets as they were

written to disk, and decompress them as they were read. No change needed to any application or JCL.

More recently, IBM has added this functionality to z/OS without the need for ISV products. Systems programmers can setup data classes to automatically compress VSAM/E and sequential datasets.

IBM and other vendors also released products to automatically compress inactive datasets. So DFHSM (now DFSMSHsm) would automatically move inactive datasets to disk volumes (called Migration Level 1, or ML1) as compressed files, or even records in a VSAM dataset for very small datasets. Other products such as Innovation's FDRARC can do the same.

However all this compression comes at a cost, firstly in terms of increased dataset access times. In the old days, data compression would reduce I/O times, compensating for the processing time required for compression and decompression. However with today's sophisticated disk subsystems, this is less of an issue. The good news is that processors are also faster, so the overhead is less than it used to be. So unless the compressed dataset is heavily used, the chances are that you won't notice too much performance degradation.

However the big cost is in CPU seconds. The CPU overhead for this compression and decompression can be high. I've seen it double the CPU usage in some cases. Now this is OK when disk is expensive and CPU seconds are cheap. However today, the exact opposite is true. CPU seconds are expensive, driving up your software licensing costs. On the other hand, DASD costs have plummeted over the past years, diluting any cost justification for compression.

The benefits of compression may also be less than you think. I've always used a rough rule of thumb of 50%. Meaning a compressed dataset will, on average, be half the size of a similar uncompressed dataset. Of course compression values will vary from dataset to dataset.

So let's look at HSM. Today, I cannot see any reason to continue to use DFHSM ML1. Datasets should remain on disk, or be migrated straight to tape. You'll need to keep a couple of ML1 volumes for ad-hoc backups and CDS backups. However that's all they should be used for.

Similarly, in most cases datasets don't need to be compressed. There may be some cases where it makes sense, but they will be in the minority. Of course moving away from compression isn't a simple task. In most cases, more disks will need to be acquired, and storage groups and other configurations will need to be changed. But in the long run, this will pay for itself in reduced CPU usage. You may also see some performance improvements to sweeten the pot.

Source:

LongEx Mainframe Quarterly - February 2013, Retrieved from <http://www.longpelaexpertise.com.au/ezone/CompressDatasets.php>

management: How Many Software Products Are There?

By David Stephens, Lead Systems Programmer at Longpela Expertise

1906. That's right, there are around 1906 currently marketed software products that run on, or work with, z/OS, z/VM and z/VSE. But that's not all. There's another 227 such products that are supported, but not marketed. But what is perhaps more interesting is the number of vendors currently selling mainframe software: 246. To me that looks like a very healthy mainframe software market.

At the end of 2010, we at Longpela Expertise asked this same question: "how many mainframe software products are there?" At that time, we had been around the mainframe world for a long time, and together believed we knew most of the major mainframe software products on the market.



But at every new site we would find one or two more that we hadn't heard of. And more importantly, we would have to dig around to find out what it was, what it did, and who supported it. So we decided to do a survey to see how many we could find.

What followed was two years of hard work. Fitting in with our other commitments, we searched for mainframe products. And by "mainframe product", we meant any

software product that runs on the three major IBM mainframe operating systems (z/OS, z/VM and z/VSE), or other software that works with these operating systems. We searched through vendor websites and documentation, forums and user groups, advertisements, blogs, articles, trademark registrations, patent documents, and even court documents to find out past and present mainframe software. As we searched, we saw that some of this software had a rich history: moving from vendor to vendor, changing names, merging and separating. So we decided to include some of this information as well.

As the list grew, we saw that it would be invaluable to be able to compare software. So if you had one software product, you could quickly find other products that did something similar. This would be ideal for searching for a replacement for a product that is too expensive, or is discontinued. This functionality could be increased to search for other products that relate to a software product. We could even categorise each product, so you could look for all software to backup DB2 databases or all security administration tools available.

With every vendor website reviewed, we saw that in many cases, it was difficult for someone with no knowledge of the product to quickly find out what it did. In some cases, it took hours of research before we were comfortable we understood the product's purpose and features. So we needed another feature: a brief, plain English summary of each product.

As we continued to work, we acknowledged that the information we had to date was only our opinion. To be balanced, and confirm that we got it right, we needed vendor input. So we have approached every active vendor, offering them the opportunity

to review and modify every entry before announcing the website – for free. We have also provided an opportunity for them to add their own opinion in a ‘From the Vendor’ section – again for free.

It’s been a fascinating ride. We’ve unearthed some amazing products. For example, we found several products to execute mainframe workloads on UNIX or Microsoft Windows machines, products to manage and monitor the HMC, and software to manage and maintain postal addresses in databases. We also unearthed four products to migrate from CA-Easytrieve, one alternative product to SAS, and products to allow z/VSE applications to execute on z/OS.

We’ve also been amazed by the choice of mainframe products. For example, we discovered 14 database management systems, 29 programming languages, and 33 banking and finance applications.

The final result is www.lookupmainframesoftware.com. This site has grown from an initial question, to a rich directory of current and past mainframe software products. We have included the status of each product (active, supported, or not supported), categories,

similar products, past and present vendors, and past and present names. We’ve included a search facility, and quick drop-downs to list software by starting letter, vendor or category. Where possible, we’ve included links to vendor pages, links to relevant Wikipedia pages, and replacement products where appropriate.

However the work doesn’t end there. Software can change quickly, so our site is only valuable if the information is up to date. We think this website is a valuable tool for anyone working with mainframes, and so we have made a commitment to maintain it. We continue, and will continue, to welcome vendor and other SME comments. We have also developed tools to regularly and automatically check each entry. We will also be manually reviewing pages regularly.

We hope everyone enjoys, and find useful, lookupmainframesoftware.com.

Source:
LongEx Mainframe Quarterly - May 2013,
Retrieved from
<http://www.longpelaexpertise.com.au/ezone/HowManySoftware.php>

Copyright Disclaimer: Interskill Learning product names and images referenced herein are registered trademarks, or trademarks of Interskill Learning, or one of its subsidiaries. All other trademarks, trade names, service marks, content, images, and logos referenced herein belong to their respective companies.