

Retaining Data Center Skills and Knowledge

By Greg Hamlyn

This is the final chapter of this **four part series** that briefly explains the data center skills crisis and the pros and cons of implementing a coaching or mentoring program. In this installment we will look at some of the steps to implementing a program such as this into your data center. If you missed these earlier installments, click the links below.

[Part 1 – The Data Center Skills Crisis](#)

[Part 2 – How Can I Prevent Skills Loss in My Data Center?](#)

[Part 3 – Barriers to Implementing a Coaching or Mentoring Program](#)

Part Four – Implementing a Successful Coaching or Mentoring Program

The success of any project comes down to its planning. If you already believe that your data center can benefit from skills and knowledge transfer and that coaching and mentoring will assist with this, then outlining a solid roadmap and plan is imperative. The items below describe the major steps that you can take to implement such a project within your data center.

Discuss Coaching and Mentoring with your Data Center Management Group

You will need to identify whether there is initial support for a coaching or mentoring program. You should add an agenda item to one of your data center meetings and provide some benefits of implementing such a program ([refer back to one of our earlier articles in this series](#)). You should summarize by proposing that you create a plan that describes these benefits in more detail, provide some case studies and how a pilot project could be run.

Your Formal Plan

This document should contain the following information:

- What the program is looking to achieve
- Some existing case studies
- Pros and Cons of running this program
- Costs and resources (both for the pilot and project overall if it is approved)
- Identify the key people involved in the pilot
- Identify the coaching or mentoring model to be used (a quick internet search will uncover a number of different models). One popular one you

Inside This Issue

Retaining Data Center Skills and Knowledge	1
InterSkill Releases – December 2011	2
Vendor Briefs	3
Taking Care of Storage	4
Learning Spotlight – Managing Projects	5
Tech-Head Knowledge Test – Utilizing ISPF	5
OPINION: The Case for a Fresh Technical	
Opinion	6
TECHNICAL: Lost in Translation Part 1 -	
EBCDIC Code Pages	7
MAINFRAME – Weird and Unusual!	10

should consider is the GROW model

- Determine whether an external consultant should be used (include pros and cons)
- Create a basic timeline of the project
- Identify how you will measure the effectiveness of the project
- Provide some basic steps describing the coaching and mentoring activities
- Next phase if the pilot program is deemed successful (i.e. specific data center areas that would be targeted)

Pre-Coaching and Mentoring Requirements

Before looking at the actual process used by the coach/mentor and their protégée, there are some other things for you to think about:

Pairings – In the ‘Barriers’ installment of this article I discussed that not everyone can or will want to be a coach or mentor, and as in life not everyone will get along with each other. With that in mind, the pairing of a coach/mentor and their protégée will need to be selected carefully. In a technical environment like the data center, the number of choices may be limited.

Training – Everyone involved in the project will require training so that they are all aware of the process, outcomes and expectations. Training should be provided to everyone involved and would include: sessions for management on coaching/mentoring aspects and benefits, training for people interested in becoming a coach or mentor (this would contain information on successful traits, methods, guides, checklists and problem resolution examples), training for protégées that describes what is in it for them, the data center and the organization, the process that will be used and the measurement of its effectiveness.

Continued on page 2

Guides/Checklists – There may be a number of documents used to identify, monitor and record the progress of the coaching sessions. These documents in turn will be used to provide feedback on the effectiveness of the program as well as for individual performance reviews run by your organization.

The Coaching and Mentoring Process

As mentioned earlier, the process used by the coach/mentor and their protégée depends on the model you select. The steps below are adapted from the GROW model and are based on a technical skills and knowledge requirement:

1. The protégée and coach are assigned and both parties are made aware of the relationship.
2. The coach (technical expert) identifies the topics that should be covered.
3. The protégée and coach agree on appropriate time and period in which the coaching will occur.
4. The coach will work through identified topics and discuss the items in relation to:
 - a. How they are used within the data center, citing various scenarios (this provides perspective)
 - b. Why they need to know it (this identifies importance)
 - c. How the product/action is undertaken (the actual operation)
5. The protégées progress is recorded.
6. The protégée and coach discuss the protégées next step in regards to training so that priorities can be determined.
7. The date/time of the next coaching session is confirmed.

As you have probably already determined, many of the requirements discussed in this article have been watered down (there are entire books that cover these items and much more in relation to coaching and mentoring) but it should give you an idea of the amount of planning and commitment required to implement and maintain a coaching/mentoring program.

If you have run or are currently running coaching or mentoring within your data center, let us know of your experiences by logging them onto our [blog](#) site.



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. 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@au.interskill.com

Interskill Learning Releases – December 2011

The latest Interskill 11.3 Release contained the following courses;

IBM Power Series – AIX Curriculum, released at AIX 7.1:

- AIX Fundamentals
- AIX Fundamentals for UNIX Administrators
- AIX Virtualization, VIO Server and Management

IBM Power Series – IBM i Curriculum, released at IBM i 7.1:

- IBM i Fundamentals

Data Center Storage Management Curriculum, released at z/OS V1.1:

- Using DFSMS to Manage the Data Center Environment

IMS Curriculum, released at IMS V12:

- Introducing IMS
- IMS Commands
- IMS Databases

The **Mainframe Presenter Series** integrates stand-alone PowerPoint presentation templates and Webinar frameworks with Datatrain course modules to enable you to more easily conduct in-house training and share your site-specific knowledge. Four new Learning Pathways are now available;

- IMS Programming Fundamentals – Learning Pathway
- CICS Fundamentals – Learning Pathway
- DB2 Fundamentals – Learning Pathway
- COBOL Fundamentals – Learning Pathway

Vendor Briefs

A number of vendors throughout the last quarter have been busy releasing new or updated cloud offerings, while IBM has also released a new version of z/OS, and is now supporting Microsoft Windows on selected zEnterprise blades. Here is a brief run-down on what the major players have been up to.

IBM

zEnterprise - True to its promise, IBM's zEnterprise environment has now been expanded to include support for Microsoft Windows on Xeon blades installed within the zBX component. This system will use a specifically configured KVM hypervisor called RHEV-Blue for virtualization purposes. zManager, which manages the entire zEnterprise system, has also been expanded to include new APIs that allow the integration of management tools capable of performing functions such as automated discovery, monitoring and provisioning.

Cloud - IBM is keeping up with the cloud hype and has recently released its zEnterprise starter edition for cloud, which allows owners of this system to configure it with cloud-like attributes such as self service provisioning. While on the cloud, IBM's newly released Starter Kit for Cloud on Power performs a similar function for IBM's Power Systems infrastructure.

z/OS 1.13– There looks to be some more than useful additions and updates with this release ranging from DFSMS tape handling enhancements and catalog management functions designed to assist staff in identifying and resolving catalog problems, to new health checks and typical improvements in the performance of processing z/OS workloads. A new z/OS base component called Batch Runtime provides application programmers with a framework for COBOL and Java interoperability with DB2.

z/OSMF 1.13 – IBM continues to add functionality to this systems programmer tool, and in this release have implemented features such as; a Software Deployment function (for cloning z/OS images and installing system software), a DASD Management Task (adding and managing DASD to storage pools), and Capacity Provisioning (used for adding or removing System z hardware as required). This release also provides the user with a web-enabled ISPF interface.

CA

Cloud - CA Technologies have also been keeping up with cloud advances, releasing a number of new and updated products and solutions including; CA Business Service Insight 8.0, CA Automation Suite for Data Centers 12.5, CA AppLogic 3.0, CA Automation Suite for Clouds 1.0, and CA Virtual Placement Manager 1.0.

Virtualization - With many organizations encountering compliance and security issues surrounding their virtualized environments, CA Technologies have recently released CA Access Control for Virtual Environments, which tightens and automates access in this area.

Application Performance – With performance playing a visible role in today's enterprise environment, organizations are always looking for tools that can improve this aspect. CA Technologies' new Mainframe Application Tuner is one such product that can be used to identify the causes of z/OS based system application performance inefficiencies, resulting in improved response times and lower CPU consumption.

BMC

BMC have recognized the complexity of managing today's enterprise environment and have recently released a new set of Middleware Management solutions that allows customers to identify and analyze applications and their flow paths, and manage middleware involved in cloud deployments.

EMC

Storage – Over the last few months, EMC has been releasing or updating a number of storage solutions for z/OS environments. One of them is the DLM6000, which is EMC's latest, and the industry's fastest, mainframe virtual tape library. With a logical storage capacity of 5.7PB (petabytes) and throughput reported to be twice that of its closest competitor, it is able to consolidate some organization's multiple tape and disk systems into this single environment.

Taking Care of Storage



By Greg Hamlyn

In times gone by, the majority of IT personnel didn't need to worry too much about data storage. The Storage Administrator, or someone performing the tasks associated with that role, would be responsible for ensuring that sufficient storage was purchased and that it was configured for appropriate access. Fast forward to today's information age where organizations are having trouble keeping up with data storage requirements that are growing by 20%-50% each year (depending on your organization) not to mention, compliance and security needs! On top of that, some individuals and business units are now responsible for provisioning their own data storage! In this article we will provide an insight into the data explosion, how data centers are handling it, and where enterprise data center storage is headed.

It Used To Be SOOOO Easy

In the past, it was relatively easy for organizations to determine their long-term data storage needs. They would then convert this information into a business case justifying the purchase of additional hardware and it would be included in the yearly IT budget. Some organizations didn't even require a dedicated job role to manage storage; the tasks were just incorporated into another IT person's role.

Today's Storage Issues

So what has changed?

1. Firstly, the budgets just aren't there any more that allow you to continually purchase huge chunks of storage every year.
2. Server virtualization, which is commonplace these days, places more demands on storage (performance, capacity, utilization, manageability and availability) making it a more complex environment.
3. Many vendors are offering storage outsourcing services. While this may seem on the surface, an easier and cheaper option, there are issues such as; compliance (HIPAA, SOX), data security (there have been well documented problems with cloud providers being hacked) and whether sufficient data backup and recovery services are provided.

But it is not all bad. Hardware, software and networking technologies have improved dramatically. In the 1980s, a round tape reel held around 160MB of data and the data transfer rate was 1.25MB per second. Today's cartridge media using the latest compression algorithms can store 4TB of native data with a 250MB per second data rate. Similarly, disk system technology improvements have seen the emergence of huge capacity storage systems such as IBM's SONAS (Scale Out Network Attached Storage). This system has a capacity of 21.6 PB (petabytes) and can manage up to a billion files under a single file system.

What Storage Technologies Will Help You?

Even with technology advances, storage volume growth is outstripping the dropping price of physical media. So, where is enterprise storage headed? Well it depends on whether the perceived security and compliance issues associated with cloud computing can be overcome? Many believe that over the next few years the risks associated with the provisioning of cloud storage will dissipate allowing them to reap the benefits associated with public or outsourced cloud offerings.

What about those organizations that will need to retain and manage their own storage? Where traditional hard disks and tape systems have remained king for so long, the following technologies can be seen emerging as serious data storage contenders;

- Solid State storage – Traditional disks excel at sequential read and write performance but with today's workloads demanding quick random access, a growth in deployment of solid state drives is occurring.
- Storage for virtual environments – Flexibility and rapid storage provisioning required by virtualized environments will see a combination of storage technologies (hard disks, flash, inline data deduplication) implemented by organizations.

Continued on page 5

- Data deduplication – This data compression method popular with backups and archiving for several years now, is finding support in production storage arrays and appliances.
- Thin provisioning – This method prevents over-allocating of storage by writing data when required (on-demand) as opposed to traditional up front allocation of storage blocks.
- Automated Tiered storage – Several vendors are offering tiered storage, which places often accessed data on high performance storage devices (flash) while less used data is placed on slower, usually larger capacity drives.
- Software – Advances in storage management software will continue to improve the way that it optimizes and analyzes data storage.

The future of enterprise storage seems to be an integration of several hardware and software technologies that are currently beginning to be offered by a number of vendors. The particular solution for your organization will undoubtedly be a trade-off of cost against performance, and deciding the best fit for the type of data you need to store.



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@au.interskill.com

Learning Spotlight – IBM Power Series – AIX

In this edition of the Learning Spotlight we focus on AIX on POWER. AIX is the leading proprietary UNIX system developed and marketed by IBM, and is widely used in enterprise-class corporations. In conjunction with the current range of POWER 7 based hardware, AIX can be implemented on "stand alone" servers, or more commonly as virtualized environments.

As AIX is similar (but different) to other flavors of UNIX, it is important that a novice AIX system administrator, or even one who is well versed in another UNIX flavor, has a good understanding of the AIX environment. We have produced a range of modules that cater for both experience levels, and you can now access the Concepts of AIX module from our [AIX Fundamentals](#) curriculum. This module explores the background of AIX, explains how to access and interact with AIX, and details typical administrative tasks.

We hope you enjoy it!



Tech-Head Knowledge Test

With every release of this newsletter a mini online test will be provided with 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 core z/OS product ISPF, and consists of 15 questions. Click the link below to start.

[Utilizing ISPF act](#)

OPINION: The Case for a Fresh Technical Opinion

A few years back I came across a mainframe site that had a way of regularly refreshing the ideas of their systems programming group. This group consisted mainly of permanent staff that had been there for many years, and knew the systems inside out. To inject new ideas and opinions this site used contractors, but not in the usual way. Every year or two, these contractors were replaced. So someone fresh would regularly come and work with their systems.

Mainframe staff have a hard time. The mainframe has undergone many changes over the past few years that they are expected to keep up with. There are regular new software releases that not only need to be installed, but have new functions that may or may not be of benefit.

Anyone with an outsourced mainframe is in a worse position. Their software levels may be up to date, but the chances are that many new features aren't being implemented: either because it's not in the outsourcer's contract, or staff simply don't know about them.

In mainframe systems, there's also a tendency to be wary of new technology. The toleration for any mainframe outage ranges from low to none, a hard ask for mainframe support staff. So it's common to delay the implementation of newly introduced functions for a couple of years to let them settle down before throwing them into production systems. PDSEs are a classic example of why this isn't a bad idea.

This means that today there are probably lots of new features that are available for free, but not used in your systems. These may make things run more efficiently, remove the need for home grown utilities and exits, or even remove the need for other software. I'm always amazed when I see a z/OS sysplex that doesn't use Enhanced Catalog Sharing, or haven't moved their JES2 checkpoint dataset onto their coupling facility: two innovations that are now over a decade old.

But the advantages of a fresh opinion don't stop there. There are many different ways to skin a mainframe cat. With a fresh opinion, you get someone who has taken that cat skin in different ways. For example I was recently at a site that had written their own system in C. Running as a started task, it had become part of the core systems, and as such was doing a whole lot of work. Their programmer had spent a lot of time tuning this system, and squeezing the best performance he could. And it ran very well. I sat down for about six hours, and found areas that would improve performance by 15%. This doesn't mean that I was smarter. I just had a different skill-set and experience, and saw different opportunities.

In the past conferences such as Share were a great way for existing staff to get exposure to new ideas. Presentations about user experiences with new or not-so-new features, and vendors showing how to get more from their products were a gold mine of information. However it can get more difficult to justify the education expenses to send technical staff to these conferences, and they are being scheduled less and less frequently.

Conferences only provide some of the benefit. Having someone else actually looking at your system, and asking "why are you doing this", or "have you thought about doing that" is the real benefit. A fresh pair of eyes looking at your systems.

So having a fresh contractor working at your site can be a great idea. But for most sites this simply isn't viable. Any new hires have to be approved by management and HR departments. If you get this approval, finding a contractor with the skills and experience you need is no easy task. When you do find one, the chances are that you won't want to lose them. For those with outsourced systems, contractors often aren't an option at all.

A consultant is great alternative. But let's face it. The chances are that you won't be able to justify bringing in a consultant merely to do a health check of your system, no matter how good an idea this is. But what if the consultant is hired to do another project? Then you have that elusive, invaluable fresh technical opinion right there. Funding an extra day or so to let this consultant look around can provide no end of benefits.

Think about a fresh technical opinion.

David Stephens

Source: Longpela Expertise, *LongEx Mainframe Quarterly* –August 2011, Retrieved from <http://www.longpelaexpertise.com.au/ezine/FreshTechnicalOpinion.php>

TECHNICAL: Lost in Translation Part 1 – EBCDIC Code Pages

By David Stephens, Lead Systems Programmer at Longpela Expertise

In the first article of a series of three, we look at EBCDIC code pages - what they are, why they're used, and what this means.

From a very early age, most of us are taught about ASCII, and how this is used by computers to convert single byte numbers to the characters we see on our screen. So an 'a' is really 97 as far as the computer is concerned. So imagine my surprise when I found out that mainframes don't use ASCII, but EBCDIC. I remember my reaction: "You've got to be kidding! Didn't EBCDIC die out years ago?"

Nope. And it's not just z/OS that uses it. IBM i, Fujitsu BSD2000/OSD, Unisys MCP, z/VSE, z/VM and z/TPF all happily continue to use EBCDIC today. To them an 'a' is really 129, not 97.

This all worked out fine for many years. In fact EBCDIC was the most popular encoding system in the world until the Personal Computer revolution brought ASCII to the limelight. But EBCDIC falls down when we need to display languages other than English. Words like "på" (Swedish) and "brève" (French) need special characters not necessarily available in the standard EBCDIC table. What's worse, there's no way that all these special characters for all the languages in the world are going to fit into the 255 places that an eight bit number has. To get around this, IBM created code pages

EBCDIC Code Pages

Today there's no such thing as a single EBCDIC code table. You can find a few websites that claim to convert from ASCII to EBCDIC. But the chances are that they're really converting from ASCII to EBCDIC code page 37, or EBCDIC 0037. EBCDIC 0037 is the default code page used by the United States and other English speaking countries when working with MVS: the traditional side of z/OS. It has all the normal a-z, A-Z, 0-9 characters, and other symbols like +, () and *. It also includes a few of the foreign characters for when we've borrowed foreign words like "résumé".

However if you're in France, the chances are that you'll be using EBCDIC 0297. In EBCDIC 0297, the standard a-z, A-z and 0-9 characters are the same as EBCDIC 0037. But to see French words, other characters are used for other numbers. For example, 177 is a pound sign (£) in EBCDIC 0037, and a cross-hash (#) in EBCDIC 0297. Longpela's [free EBCDIC code converter](#) tool shows common EBCDIC code pages.

There are many different code pages for all the different regions. From Spain and Iceland to Thailand and Japan. This is not a lot different to ASCII, which has gone from the original 7-bit code ASCII to ISO8859 with its different sub-definitions. For example ISO8859-1 is the standard 'Extended ASCII' that we all love, ISO8859-2 is better for Eastern Europe, and ISO8859-4 for countries like Latvia, Lithuania, Estonia and Greenland.

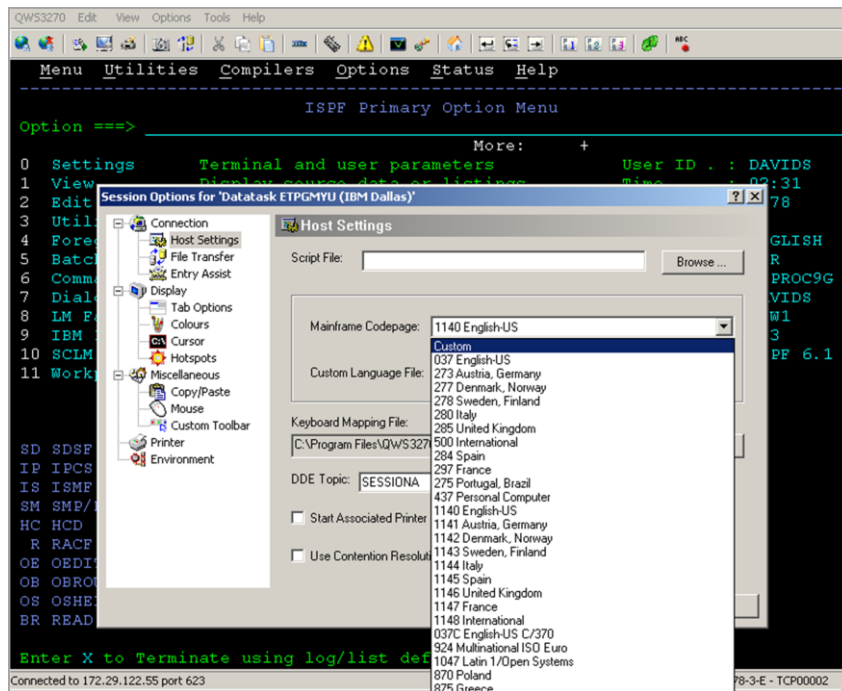
Every EBCDIC code page uses the same numbers for the standard a-z, A-z, and 0-9 characters, along with a few other standard symbols. So you can see COBOL and JCL code, standard z/OS messages, and update sys1.parmlib the same way, regardless of the code page used. Code pages mess with characters that aren't normally needed when programming or administering the mainframe (remember this statement - we'll talk about it more in a moment).

Of course having characters that move around can be awkward. For example lawyers always want a copyright symbol © displayed on screens. But a © on your screen could be another character entirely in a different character set. This is why you'll often see a (c) rather than © on 3270 screens.

IBM controls these EBCDIC code pages, and assigns an ID to them called the Coded Character Set Identifier (CCSID). The CCSID for EBCDIC 0037 is, you guessed it, 37. IBM also has set CCSIDs for other character sets - CCSID 1208 is Unicode (UTF-8). You can see them all on [IBM's website](#).

Continued on page 8

Here's an interesting fact: the original IBM System/360 mainframe could run in either EBCDIC or ASCII. Unfortunately, the operating system couldn't, and this feature was dropped for the System/370. Today there's still no switch or definition to tell z/OS what code page it is using - it happily uses the standard characters that don't change. It leaves you to figure out how you want to view the rest. For a TN3270 user, this is done in the emulation software. Below shows the screen to change the codepage for Giant Software's QWS3270.



How C and REXX Mess Things Up

Before I said that code pages only change characters that aren't needed when programming or administering the mainframe. And that works for traditional z/OS features, and programming languages such as COBOL and PL/1. But there are a couple of exceptions, and you can safely bet that more will follow.

Take REXX for example. A common thing to do in REXX is concatenate two strings. This is done using the two vertical bars ('||'). For example, our sample REXX to get z/OS information has the line:

```
say 'z/OS version:' zos_ver || '.' || zos_re1
```

The bad news is that the vertical bar character '|' isn't one of those 'standard' EBCDIC characters. So a vertical bar in EBCDIC 0037 looks like an exclamation mark (!) in Sweden. The REXX interpreter doesn't care what this character looks like, as long as it has a code of 79. So if you're in Sweden and using EBCDIC 0278, the above line becomes:

```
say 'z/OS version:' zos_ver !! '.' !! zos_re1
```

C is another problem child. It uses funky characters like the square brackets '[]', curly brackets '{}', and broken vertical bar '|'. These move around (or disappear) depending on your code page. But with C there's another catch: it's designed to use EBCDIC 1047, not EBCDIC 0037. So if you're using arrays in C, the line:

```
char cvtstuff[140];
```

is fine if you're using IBM1047. For IBM0037, it becomes:

```
char cvtstuffŸ140´
```

If you're using EBCDIC 0050, another common EBCDIC code page, it becomes:

Continued on page 9

```
char cvtstuffY140"
```

How USS Doesn't Follow the Rules

So why is C designed for EBCDIC 1047? Because z/OS Unix Systems Services (USS) is also designed for it.

When IBM created USS for z/OS, it makes sense that it had to work in EBCDIC. The POSIX standard for UNIX doesn't require the use of ASCII, and z/OS is an EBCDIC operating system. IBM really didn't have a choice.

The problem is that UNIX, and its core programming language C, rely on characters that don't exist in some EBCDIC codepages. EBCDIC 1047 is designed to include all the characters USS needs - effectively all the characters from Extended ASCII: ISO8859-1. So EBCDIC 1047 is the default EBCDIC codepage used in USS. All parameter and help files are usually supplied in EBCDIC 1047, the C compiler expects code in EBCDIC 1047, and all UNIX file contents default to EBCDIC 1047. If you decide to use something else, they may look a little funny.

Dealing with Larger Alphabets

Up to now, our character sets have had one byte per character: Single Byte Character Sets (SBCS). However anyone who speaks Japanese, Chinese or Korean will laugh at the idea of their character set fitting into a mere 255 places. The EBCDIC solution: using two bytes for each character - double Byte Character Sets (DBCS). In reality, the chances are that these character sets will use a combination for better performance: Latin characters as single-byte, Asian as double-byte. This is called a Mixed Byte Character Sets (MBCS).

DBCS support is still very limited, and isn't enabled by default. DB2 supports DBCS, as does TCP/IP and many TCP/IP applications such as FTP and SMTP. Unicode Services supports DBCS, and includes some DBCS conversion tables. Websphere MQ is also up for DBCS.

What This Means

For many years I was ignorant of code pages and language issues. My first wakeup call was when I tried to program in C from ISPF Edit. Another wakeup call happened when I first needed to convert from EBCDIC to UTF-8, a conversion that needs to know the EBCDIC code page used.

The fact is that there are some basic rules that aren't written down, but are very important:

1. Only use 'standard' EBCDIC characters if your dataset may be used somewhere else, and you want it to look the same. For example, any product parameter datasets should keep it to 'a-z, 'A-Z' or '0-9'.
2. Be aware that the standard EBCDIC code pages may be different between TSO/ISPF and USS.
3. If you are in different countries, be aware of the character set your emulator uses.
4. Know your EBCDIC code page whenever converting between EBCDIC and another character set like ASCII or UTF-8.

A basic understanding of code pages is more than handy, it's important. Particularly today where computers supply information to users worldwide, and where information must be converted between the mainframe and other computers.

This first article has taken a quick look at EBCDIC code pages, and what they mean. In the next part of this series I'll look at how ASCII, Unicode and EBCDIC work together on z/OS.

David Stephens

Source: Longpela Expertise, *LongEx Mainframe Quarterly –August 2011*, Retrieved from <http://www.longpelaexpertise.com.au/ezine/LostinTranslation1.php>

David is the Lead Systems Programmer at Longpela Expertise. He has over 18 years Mainframe Systems Programming experience. He has been the sole site Systems Programmer twice, and has six years systems software development and level 3 support experience with IBM. He is the author of [What On Earth is a Mainframe?](#)

MAINFRAME – Weird and Unusual!

By Greg Hamlyn

While weird and unusual aren't adjectives you would normally expect in the same sentence as 'mainframe', we thought that we would take a tangent from this usually tech heavy topic to look at some mainframe web links with a twist.

Reboot

The following site provides some background on a CGI television show called "Reboot", which first aired in 1994. Produced by Mainframe Entertainment (I kid you not), it told the story of Mainframe, a city in cyberspace, and the various "programs" (people) that lived within its walls. Watch out for the love interest (Dot Matrix)!

<http://tvtropes.org/pmwiki/pmwiki.php/Main/Reboot>

Joe Clabby

Many of you will know of Joe Clabby through many of his articles on servers and storage, cloud computing and service management. This link tracks back to one of his trips with his son, to an island off of Brazil to visit Hoplon Infotainment, a company using IBM mainframes to support its online gaming environment. While a couple of years old now, I love the mixing of business with pleasure.

<http://searchdatacenter.techtarget.com/tip/Joe-Clabbys-Brazilian-mainframe-adventure-travel-log>

OBJECT

The following sites have some fascinating pictures linked to a number of unusual topics. Some of the soviet graphics show computers not that dissimilar to ones that I worked on in my early days!

<http://www.oobject.com/category/12-soviet-block-computers/>

<http://www.oobject.com/category/top-supercomputers/>

The Computer History Museum

The Computer History Museum is the world's leading institution exploring the history of computing and its ongoing impact on society. With that in mind, it is not surprising that there is a reasonable display/exhibition relating to the humble mainframe.

<http://www.computerhistory.org/revolution/mainframe-computers/7/intro>

Build Your Own Mainframe

I've left the best till last... Five steps describing how to build your own mainframe. Who knew it could be that simple!

http://www.ehow.com/how_2121798_build-mainframe.html

In the next edition of Interskill Learning Mainframe Quarterly:

Cloud Computing – Part One – The first article of a three part series that discusses what cloud computing is and isn't, and how it is being embraced.

Learning Spotlight – Project Management – We look at Project Management and provide you with access to our new modules.

z/OSMF – The Future of Systems Management? – This article looks at typical systems programming tasks, how much time they take, and whether z/OSMF is a savior or a hindrance.

TECHNICAL: Lost in Translation Part 2 – The second article of a three part series that discusses how ASCII, Unicode and EBCDIC work together on z/OS.

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.