

Is Your Data Center Career Cloudy?

By Greg Hamlyn

Unless you have been living under a rock you would not have escaped the wave of “cloud” that is engulfing the IT industry. What started as hype in 2010 and 2011, and what many data center professionals saw as being just another IT fad, has grown into a full grown industry that still seems to be evolving to what will become a documented paradigm shift.

This article is the first of a **four part series** that discusses what the cloud actually is, so that when you discuss cloud computing with other IT professionals in your organization, you are all on the same page. Future installments of this article will discuss specifically how the cloud will affect you and what you can do to keep up with it:

Part 2 - What Affect Will The Cloud Have on Data Centers and You?

Part 3 - How Can You Build Your Cloud Skills Portfolio?

Part 4 - The Cloud’s Future... And Yours.

So What Is The Cloud?

Put simply, the cloud is a pooled set of virtualized resources (processors, storage and networks) that can be accessed as a service by users. The amount of resources provisioned are elastic to allow for peaks and troughs in workloads, and the user only pays for the resources that they consume.

An analogy often used is that it is like obtaining your electricity from a utility company. You access, consume and pay for, only as much as you require.

If you have been working in a Data Center for some time, you may be thinking that this type of service has been around in various forms for many years. Your virtualized environment has allowed users to create environments as required and your mainframe has for many years provided a system that for all intentional purposes provides users with the illusion that they have unlimited resources to work with. So, how is cloud computing different?

What Makes The Cloud Different?

Chances are that if you are discussing cloud computing with a colleague that you may be discussing entirely different things. The reason for that is that the cloud is not just one thing. It is a computing model that can be configured in several different ways depending on your needs and

Inside This Issue

Is Your Data Center Career Cloudy?	1
Interskill Releases – December 2011	3
Introducing the Interskill Executive Series	3
Vendor Briefs	4
What the Future Holds for the Mainframe	4
Learning Spotlight – AIX	6
Tech-Head Knowledge Test – Using DFSMS to Manage z/OS Storage	6
TECHNICAL: Lost in Translation Part 2 - ASCII and EBCDIC on z/OS	7

budget. These are the types of characteristics that are generally associated with cloud computing:

- Physical resources such as servers, processors, memory, and disk space, are used to create a pool of virtual resources that can be provisioned as required
- A cloud’s scalability allows additional resources to be provisioned as workloads increase or spike. These resources can then be released when they are no longer required.
- The value of the cloud also comes in its ability to provide multi-tenancy, where multiple customers share certain cloud provider instances. While this may not suit some organizations, it provides considerable cost savings by sharing resources.
- The flexibility of some cloud configurations allows “like” organizations to share their expertise as well as their existing data processing and storage resources. An example of this might be local public hospitals that can take advantage of specialty medical expertise and equipment spread across their group.
- With self-service provisioning, administrators or developers can request and receive new computing resources without it needing to be processed through regular IT channels. This saves considerable time and money although controls need to be implemented so that not everyone has this capability.
- The ability to access cloud resources from anywhere, using a number of different devices (desktops, laptops, mobile phones) using standard access protocols and interfaces is an important characteristic that many business demand in today’s mobile world.

Continued on page 2


- Another important characteristic of the cloud is its resiliency, which refers to its ability to make data and systems available in the event of a disaster. Without this defining attribute, no users would consider even moving to the cloud.
- The concept of billing the customer for the cloud services consumed is what makes cloud computing so attractive. You only pay for what you use.

So you can see why cloud computing has captured the attention of management in particular. Depending on the cloud model used, the touted savings can be considerable. It also provides the organization's workforce with flexible on-demand computing resources, allowing it to create new customer services and applications faster as a result of more efficient self provisioning of resources.

We'll now look at some of the more common cloud models and who is using them.


Cloud Models

Infrastructure as a Service (IaaS) is the easiest model to conceptualize as the cloud provider owns the hardware (servers, network infrastructure, storage) and the space that it resides in. The customer then creates their required virtualized environment using the self provisioning process discussed earlier. The user is then billed for the resources they consume.

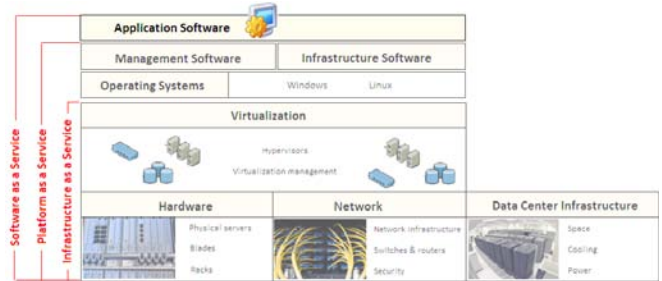


This model is popular with organizations that deal with considerable processing spikes and troughs, new computer start-ups, where investment could be spent elsewhere instead of on infrastructure, or for organizations under financial pressure to cut capital expenditure.

Platform as a Service (PaaS) offerings provide you with a hosted development platform that runs on the virtualized infrastructure of the provider. This allows you to create customized cloud applications using the online tools and APIs provided.



At the top of our building blocks is **Software as a Service (SaaS)**. This hosted model provides you with a complete operating environment (software and infrastructure) so all you need is an interface to enter and manage your data. This interface is commonly a web browser.



Now that you have these acronyms covered, it is important to identify the environments in which they can be deployed. Initially, these offerings were seen as being provided externally, which gave it the look of another form of outsourcing. But cloud configurations are evolving, mainly as a result of data security fears and regulatory compliance requirements, so that today's cloud may be:

- Public clouds (external)
- Private clouds (internal or external)
- Hybrid clouds (combination of public and private)
- Community clouds (organizational sharing)

Public Clouds

With a public cloud, the infrastructure, services and applications are all accessed from a provider who is hosting them. This environment exists outside of your organization's firewall and is configured for multi-tenanted use. Because a public cloud's vast resources are shared by many users, it provides the most benefit in relation to cost savings.

Private Clouds

There are two versions of a private cloud that can be considered, an on-site private cloud and an outsourced private cloud. Where organizations have already made a sizable investment in IT equipment, they may investigate implementing an on-site private cloud. This is basically a conversion of their traditional data center structure to one where the resources take on cloud-based characteristics. In an outsourced private cloud environment, a provider hosts a dedicated or leased portion of their facilities, and access to their services is through a private network structure that prohibits unauthorized access.

Hybrid Clouds

Where aspects of public and private clouds are undoubtedly enticing for organizations, the increasingly popular hybrid cloud deployment option is being considered. This allows the organization with a private cloud to extend their services by incorporating public cloud resources, providing them with the best of both worlds.

Continued on page 3

Community Clouds

A community cloud, is a more targeted version of the public and private cloud models, where organizations that have similar backgrounds (such as health services and government agencies) and related security, compliance and other requirements, band together to share infrastructure and services.

Who is Guiding The Cloud?

So who is involved in defining where the cloud is heading and insuring the standardization required? There are in fact numerous bodies undertaking this task including:

- National Institute of Standards and Technology (NIST)
- Distributed Management Task Force (DMTF)
- Open Cloud Consortium (OCC)
- Open Grid Forum (OGF)
- Cloud Security Alliance (CSA)
- Cloud Standards Customer Council (CSCC)
- The European Telecommunications Standards Institute (ETSI)

Summary

Cloud computing has become a reality through a combination of technological advances and economic events (global credit crunch) resulting in the need to develop a smarter and less expensive way of processing, storing and distributing data. Many standards are currently before the groups mentioned above, including a standard for ensuring interoperability between cloud vendors offerings. The cloud is an interesting place to be at the moment, but how will it affect you?

In part 2 of this series, you will see how cloud computing will transform the data center and the roles of people working within it.



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
- IBM Power Series – IBM i Curriculum released at IBM i 7.1
- Data Center Storage Management Curriculum released at z/OS V1.12
- IMS Curriculum released at IMS V12:
- The Mainframe Presenter Series have released four new Learning Pathways;
 - IMS Programming Fundamentals – Learning Pathway
 - CICS Fundamentals – Learning Pathway
 - DB2 Fundamentals – Learning Pathway
 - COBOL Fundamentals – Learning Pathway

INTRODUCING THE INTERSKILL EXECUTIVE SERIES

The Interskill Executive Series is a new curriculum of Soft Skills courses designed to provide business professionals and IT staff with the skills required to successfully manage and operate within a Data Center environment.

The **Project Management** curriculum is the first set of Interskill Executive Series courses to be released, containing 12 hours of training covering Project Management Overview, Work Content and Scope Management, Networking, and Critical Path Analysis.

A new Interskill Executive Series curriculum, **Coaching and Mentoring for Data Centers**, covering training methods, and the differences between, and benefits of coaching and mentoring programs in organisations, will also be released in 2012.


Vendor Briefs

It has been reasonably quiet across the mainframe vendor frontier since our last newsletter, with the major players focusing more on management of the ever increasing complex cross platform and cloud environments. Both IBM and CA Technologies are also concerned that you may not be getting value for money from your existing infrastructure and have introduced tools and services to help you evaluate what you have and need....how thoughtful of them.

IBM

z/VM 6.2 – With virtualization playing a major role in today's data centers, this new version of z/VM, which was released in the last quarter of 2011, contains several interesting enhancements. It has the ability to cluster together up to four z/VM instances in a single system image (SSI), which are in turn managed as a single z/VM system. This structure provides a more simplified management of z/VM systems. The new z/VM single system image (VMSSI) feature allows Linux virtual machines to be easily relocated to other z/VM systems without application disruption. IBM refers to this as Live Guest Relocation (LGR).

Smarter Computing Workload Simulator – Some time back, IBM introduced the term “Smarter Computing”, and described the three areas that data centers would need to concentrate on to gain a competitive advantage.



A holistic system

Smarter computing is realized through an IT infrastructure that is designed for data, tuned to the task, and managed with cloud technologies.

While these items are huge in their own right, IBM is beginning to provide tools that assist in these areas. The Smarter Computing Workload Simulator (<http://www-03.ibm.com/systems/data/flash/smartercomputing/simulator.html>) is a tool for CIO's that allows them to identify potential savings and efficiencies...well worth a look!

Windows on the Mainframe – As reported in our previous vendor briefs update, Microsoft Windows applications can now be integrated into the mainframe environment through the zBX component of the zEnterprise. December 16, 2011 was the official date for this support and if you are after some early feedback of uptake, then you may want to check out the following link, which describes how EUROCONTROL (European air traffic management in the Netherlands) has implemented this hybrid environment. http://www-03.ibm.com/systems/be/resources/IBM_Case_Eurocontrol_EN_HR.pdf

Mainframe Education – IBM's System z Academic Initiative continues to roll on with IBM reporting that 1067 schools from 66 countries worldwide are providing students with a number of mainframe courses and free access to enterprise computing environments.

CA Technologies

Mainframe Management – With many data centers having already moved a portion of their processing to the cloud, CA Technologies have recognized how much more difficult it is for organizations to manage this environment and have released its next-generation mainframe management strategy focusing on product improvements catering for cross platform and cloud environments.

Along similar lines to IBM's smarter computing simulator mentioned above, CA is providing the following services that allow you to evaluate whether software cost savings can be made. The Mainframe Value Program evaluates the effectiveness of CA software already implemented within your organization, while the Mainframe Software Rationalization Program (MSRP) gathers and analyzes information about your mainframe environment. This information is then used by CA to show the effectiveness and cost saving capabilities should organizations consider migrating to their competing products.

BMC

Application Dependency Mapping - BMC has recently announced the latest release of its Atrium Discovery and Dependency Mapping solution, which is used to make more visible the connection between applications and the IT infrastructure supporting them. The results allow organizations to identify application dependencies and more fully understand the impact that changes will have on their system. In the past, much of this information has been too difficult and time consuming to obtain, so this product will certainly be popular with those needing to manage and fine-tune this environment.

What the Future Holds For the Mainframe



By Greg Hamlyn

Barely a month goes by where someone hasn't written an article either pronouncing that the mainframe is dead, or that it is back in a big way. So who is right? In this article we will take an unbiased (well as much as you can be working for a company that supplies mainframe e-learning) look at what people/groups are saying about the mainframe's future.

Information Sources

Before discussing what the soothsayers have been saying about the mainframe's future, it is interesting to note where these comments have come from. There are always plenty of surveys asking organizations about their IT priorities in regards to computing, many of these already with a vested interest in one technology or another. Most of the information can be interpreted in different ways and one article I read a short time ago smacked of someone that could not interpret what the results meant at all, and were brave enough to admit this in their survey.

Other comments come from industry experts. While they hold some credibility, I still have that famous 1996 quote from Stewart Alsop II in my head, who was editor-in-chief and executive vice-president of IT technology magazine InfoWorld - "I predict that the last mainframe will be unplugged on March 15, 1996".

Other comments come from IT vendors themselves, who of course have a vested interest in the debate and are usually able to skew the results of any statistics in their favor. So with all this in mind, let's look at what is being predicted.

Alive and Well

The past is often considered a good guide to the future and with mainframes there is solid evidence to suggest that IBM can reinvent themselves and provide an infrastructure and services to cater for the fast changing IT requirements of today's mobile savvy world. While their zEnterprise offering provides a cloud look and feel, the money they are spending on mainframe education and upgrading tools to give them a modern interface is enough to make you think they wouldn't be spending that much unless they saw a future.

BMC Software's Sixth Annual Worldwide Mainframe Survey received feedback from 1300 IT professionals. Statistics indicated that mainframe capacity requirements for the majority would only continue to grow and that the mainframe still played a pivotal role in their organization.

During 2011, IBM reported the best four quarters of growth in mainframe sales in over five years with many new customers coming from emerging economies such as Senegal, Namibia and Brazil. These signals indicate that new and existing organizations with critical mission workloads still see the mainframe as their future.

Fading Away

The Standish Group surveyed a number of CIO's with a mainframe environment and produced some very interesting results. One statistic showed that 70% of CIO's believed that the mainframe played a strategic role in their IT environment today but in 5-10 years, NONE (that is zero percent!), saw the mainframe as being a strategic platform. Standish puts this down to the emergence of the cloud that obviously CIO's see as having no link to the mainframe...very strange seeing that IBM's zEnterprise plays the role of a hybrid cloud.

There are plenty of individuals that will tell you that the modernization requirements of data center's today require equally savvy hardware and software, and that the mainframe is just around to cater for legacy applications. Another argument in support of non-mainframe technology is that it is less expensive...or is it?

Why the Difference in Opinion?

The problem is comparing apples with apples. While UNIX boxes or x86 machines may be cheaper up front than purchasing a new mainframe, all the costs need to be factored in. This is fascinating when reading through blogs on this discussion. Many comments have a superficial view on what constitutes the total cost of ownership (TCO) or return on investment (ROI), often just taking into account hardware and software costs. While this forms a large component, what about the number of personnel required to support the infrastructure, and other incidentals such

Continued on page 6

as; power/cooling costs, real estate costs to house the equipment, networking costs, recovery costs, and the impact of planned outages?

The Bottom Line

Many organizations have spent years building a reliable and secure IT environment and in my opinion are not likely to dump all this overnight for a new infrastructure. So, in the foreseeable future, mainframes will continue to provide many large organizations in particular with their IT needs, most likely in a hybrid environment that incorporates other technologies such as x86 servers, which are better at processing some workloads. Smaller organizations on the other hand will be more likely to obtain mainframe-like processing requirements through the cloud, or in some cases may pool their technologies and expertise to run more efficiently (my article on cloud computing that begins in this newsletter sheds some light on this technology and what it means to the data center).

Let us know what you think about the future of mainframes by logging 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, 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 – AIX

The AIX operating system has evolved from humble beginnings on the IBM 6150 RISC workstation back in 1986 to form part of many data center's strategic plans, running alongside IBMi and Linux, on IBM Power Systems hardware.

The module that we have provided for you here is the [Concepts of AIX](#), which as the name suggests provides you with an overview of the AIX environment, the role of the AIX administrator, along with a description of tasks commonly performed by this person.

We hope you enjoy it.



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 DFSMS, and consists of 15 questions.
Click the link below to start.

[Using DFSMS to Manage z/OS Storage](#)

TECHNICAL: Lost in Translation Part 2 – ASCII and EBCDIC on z/OS

By David Stephens, Lead Systems Programmer at Longpela Expertise

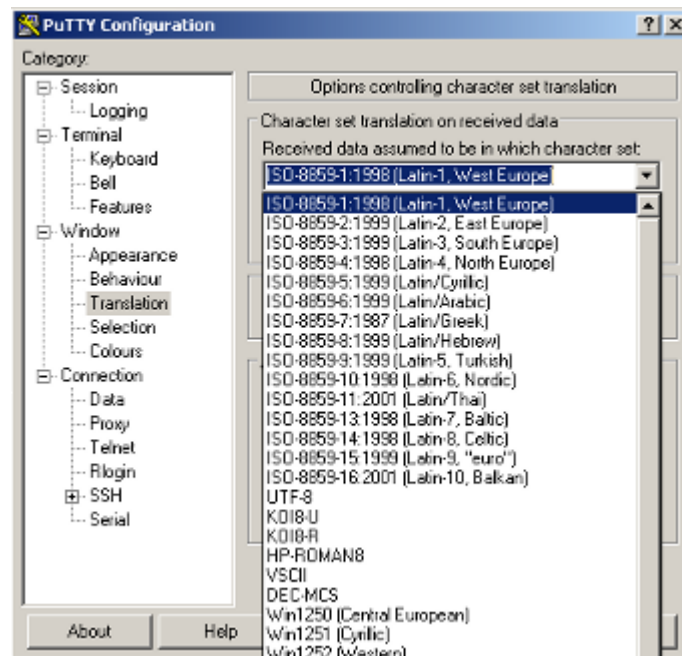
In the [first article](#) in this series, we introduced EBCDIC code pages, and how they work with z/OS. But z/OS became more complicated when UNIX Systems Services (USS) introduced ASCII into the mix. Web-enablement and XML add UTF-8 and other Unicode specifications as well. So how do EBCDIC, ASCII and Unicode work together?

ASCII Code Pages

If you are a mainframe veteran like me, you work in EBCDIC. You use a TN3270 emulator to work with TSO/E and ISPF, you use datasets, and every program is EBCDIC. Or at least that's the way it used to be.

UNIX Systems Services (USS) has changed all that. Let's take a simple example. Providing it has been setup by the Systems Programmer, you can now use any Telnet client such as PuTTY to access z/OS using telnet, SSH or rlogin. From here, you get a standard UNIX shell that will feel like home to anyone who has used UNIX on any platform. telnet, SSH and rlogin are ASCII sessions. So z/OS (or more specifically, TCP/IP) will convert everything going to or from that telnet client between ASCII and EBCDIC.

Like EBCDIC, expanded ASCII has different code pages for different languages and regions, though not nearly as many. Most English speakers will use the ISO-8859-1 ASCII. If you're from Norway, you may prefer ISO-8859-6, and Russians will probably go for ISO-8859-5. In UNIX, the ASCII character set you use is part of the locale, which also includes currency symbol and date formats preferred. The locale is set using the setenv command to update the LANG or LC_* environment variables. You then set the Language on your Telnet client to the same, and you're away. Here is how it's done on PuTTY.



From the USS shell on z/OS, this is exactly the same (it is a POSIX compliant shell after all). So to change the locale to France, we use the USS setenv command:

```
Setenv LANG fr_FR
```

The first two characters are the language code specified in the ISO 639-1 standard and the second two the country code from ISO 3166-1.

ASCII in Files

Converting to and from ASCII on z/OS consumes resources. If you're only working with ASCII data, it would be a good idea to store the data in ASCII, and avoid the overhead of always converting between EBCDIC and ASCII.

The good news is that this is no problem. ASCII is also a Single Byte Character Set (SBCS), so all the z/OS instructions and functions work just as well for ASCII as they do for EBCDIC. Database managers generally just store bytes. So if you don't need them to display the characters in a readable form on a screen, you can easily store ASCII in z/OS datasets, USS files and z/OS databases.

The problem is displaying that information. Using ISPF to edit a dataset with ASCII data will show gobbldy-gook - unreadable characters. ISPF browse has similar problems.

With traditional z/OS datasets, there's nothing you can do. However z/OS USS files have a tag that can specify the character set that used. For example, look at the ls listing of the USS file here:

```
>ls -T file1.txt
```

```
t IBM-1047 T=on file1.txt
```

You can see that the code page used is IBM-1047 - the default EBCDIC. The t to the left of the output indicates that the file holds text, and T=on indicates that the file holds uniformly encoded text data. However here is another file:

```
>ls -T file2.txt
```

```
t ISO8859-1 T=on file2.txt
```

The character set is ISO8859-1: extended ASCII for English Speakers. This tag can be set using the USS ctag command. It can also be set when mounting a USS file system, setting the default tag for all files in that file system.

This tag can be used to determine how the file will be viewed. From z/OS 1.9, this means that if using ISPF edit or browse to access a dataset, ASCII characters will automatically be displayed if this tag is set to CCSID 819 (ISO8859-1). The TSO/E OBROWSE and OEDIT commands, ISPF option 3.17, or the ISHELL interface all use ISPF edit and browse.

If storing ASCII in traditional z/OS datasets, ISPF BROWSE and EDIT will not automatically convert from ASCII. Nor will it convert for any other character set other than ISO8859-1. However you can still view the ASCII data using the SOURCE ASCII command. Here is an example of how this command works.

```
File Edit Edit_Settings Menu Utilities Compilers Te
EDIT          P058834.JUNK1(JUNKY) - 01.00
Command ==> source ascii
***** Top of Data *****
000001 \a<<! !!â<@
***** Bottom of Data *****

File Edit Edit_Settings Menu Utilities Compilers Te
EDIT          SHR.P058834.CNTL(JUNKY) - 01.00
Command ==>
***** Top of Data *****
000001 HELLO WORLD.
***** Bottom of Data *****
```

ASCII in Databases

Some database systems also store the code page. For example, have a look at the following output from the SAS PROC _CONTENTS procedure. This shows the definitions of a SAS table. You can see that it is encoded in EBCDIC 1142 (Denmark/Norway):

THE CONTENTS PROCEDURE

DATA SET NAME	DS1.CICSTXN	OBSERVATIONS	81023
MEMBER TYPE	DATA	VARIABLES	114
ENGINE	BASE	INDEXES	0
CREATED	28. APRIL 2011 TORSDAG 03: 38: 16	OBSERVATION LENGTH	840
LAST MODIFIED	28. APRIL 2011 TORSDAG 03: 38: 16	DELETED OBSERVATION	
PROTECTION		COMPRESSED	NO
DATA SET TYPE		SORTED	NO
LABEL			
DATA REPRESENTATION	MVS_32		
ENCODING	EBCDIC1142 DENMARK/NORWAY (EBCDIC)		

DB2 also plays this kind of ball. Every DB2 table can have an `ENCODING_SCHEME` variable assigned in `SYSIBM.TABLES` which overrides the default `CCSID`. This value can also be overridden in the SQL Descriptor Area for SQL statements, or in the stored procedure definition for stored procedures. You can also override specify `CCSID` when binding an application, or in the DB2 `DECLARE` or `CAST` statements.

Or in other words, if you define tables and applications correctly, DB2 will do all the translation for you.

The Problem with Unicode

Unicode has one big advantage over EBCDIC and ASCII: there are no code pages. Every character is represented in the same table. And the Standards people have made sure that Unicode has enough room for a lot more characters - even the Star Trek Klingon language characters get a mention.

But of course this would be too simple. There are actually a few different Unicodes out there:

- UTF-8: Multi-Byte character set, though most characters are just one byte. The basic ASCII characters (a-z, A-Z, 0-9) are the same.
- UTF-16: Multi-byte characters set, though most characters are two bytes.
- UTF-32: Each character is four bytes.

Whichever Unicode you use, they all have the same drawback: one character does not necessarily use one byte. z/OS string operations generally assume that one byte is one character.

Most high level languages have some sort of Unicode support, including C, COBOL and PL/1. However you need to tell these programs that you're using Unicode in compiler options or string manipulation options.

z/OS also has instructions for converting between Unicode, UTF-8, UTF-16 and UTF-32. By Unicode, IBM means Unicode Basic Latin: the first 255 characters of Unicode - which fit into one byte.

A problem with anything using Unicode on z/OS is that it can be expensive in terms of CPU use. To help out, IBM has introduced some new Assembler instructions oriented towards Unicode. Many of the latest high level language compilers use these when working with Unicode instructions - making these programs much faster. If you have a program that uses Unicode and hasn't been recompiled for a few years, consider recompiling it. You may see some performance improvements.

What This Means

z/OS is still EBCDIC and always will be. However IBM has realised that z/OS needs to talk to the outside world, and so other character encoding schemes like ASCII and Unicode need to be supported - and are.

z/OS USS files have an attribute to tell you the encoding scheme, and some databases systems like SAS and DB2 also give you an attribute to set. Otherwise, you need to know yourself how the strings are encoded.

This second article has covered how EBCDIC, ASCII and Unicode work together with z/OS. In the final article in this series of three. I'll look at converting between them.

David Stephens

Source: Longpela Expertise, *LongEx Mainframe Quarterly – November 2011*, Retrieved from <http://www.longpelaexpertise.com.au/ezine/LostinTranslation2.php>

In the next edition of Interskill Learning Mainframe Quarterly:

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

Learning Spotlight – We look at TSO/ISPF, the bread and butter product for z/OS Mainframers.

TECHNICAL: Lost in Translation Part 3 – The final article of the three series that looks at converting between EBCDIC, ASCII, Unicode and z/OS.