

Retaining Data Center Skills and Knowledge

By Greg Hamlyn

In the first of this **four part series** we looked at the data center skills shortage hype and some statistics on the aging and retirement of staff within this environment. In this installment we will discuss methods that can be used by organizations and data centers to prevent valuable skills and knowledge from being lost.

Part Two – How Can I Prevent Skills Loss In My Data Center?

You have several options when faced with losing experienced personnel.

Continued Employment

Firstly, look at whether you might be able to lure the person into further fulltime or part-time employment. This is really just delaying the inevitable but may be sufficient for some data centers that are looking to implement automation of tasks performed by that person. This may also buy you some time so that they can transfer their skills and knowledge to their replacement.

Recruitment

You can hire someone to take over the job role of the IT person leaving. You can bet that while the hired person may have some or all of the technical skills, including certification, they won't have the experience of knowing how things have been implemented and your organization's processes surrounding it. The cost of replacing the loss of skills and knowledge using this method can be considerable and ranges from the time and cost to find a suitable replacement, which on average is 3 to 6 months, to lost productivity while that person is up-skilling themselves.

Train-the-Trainer

You can ask the person who is leaving, to run training sessions for those people who may be taking over the tasks they are currently performing. It is rare to find a technical professional that is also good at training, so while good in theory this option could be difficult to implement. Think about it... if you were about to leave the company and retire would you be enthusiastic to learn how to present training courses and then stand up in front of people. The answer is that the majority would not.

Inside This Issue

Retaining Data Center Skills and Knowledge	1
Interskill Releases – April 2011	2
Vendor Briefs	3
z/OS 1.12... What's in it for me?	4
Tech-Head Knowledge Test –	
z/OS Problem Diagnosis and Resolution	4
Learning Spotlight – Linux on System z	5
MANAGEMENT: Is This The End of Tape?	6

Cross-Training

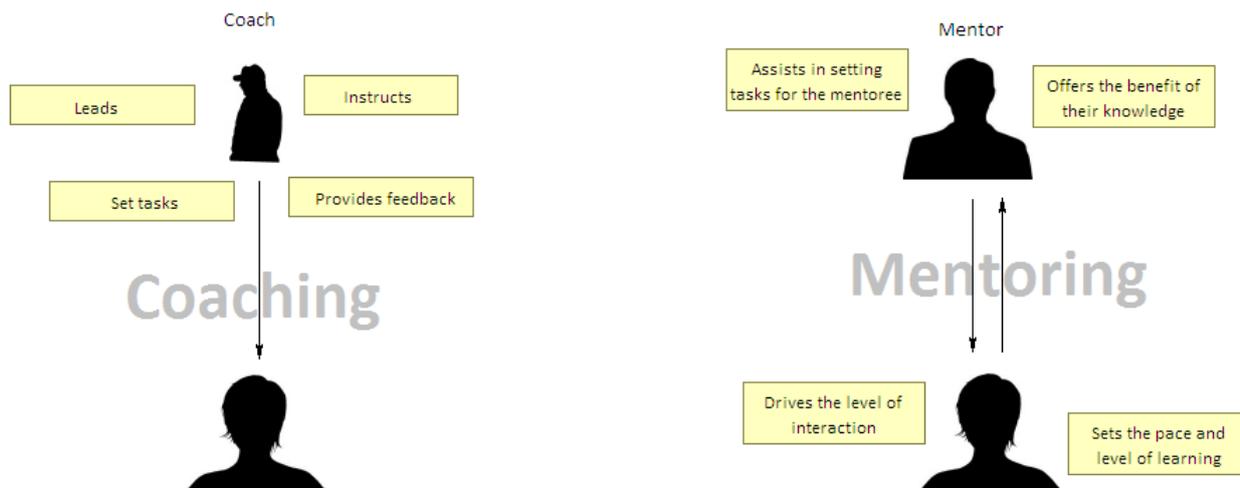
Cross training is often seen by IT departments as a way of ensuring that there is always someone around to perform the required tasks. This process is often a formal one with a list of requirements of skill and knowledge that needs to be transferred. A problem with this is that cross training may be undertaken as a "one-off" exercise and by the time that person is required to perform those tasks, processes or methods may have changed so the original information learnt may no longer be relevant. Even worse could be that the trained person may not be working in the area any more or cannot be released to perform those tasks!

Coaching and Mentoring

Often incorporated into the method above is coaching or mentoring. This is used by many organizations today to ensure important business skills are retained and that the perspective in which they are run or implemented is passed onto the replacement. It is also often seen as an ongoing learning initiative with the additional benefit of building relationships within the organization. So what is coaching and mentoring and what is the difference between them?

Coaching and mentoring seem to be terms used interchangeably today but there are differences in the way that they work, which can affect the outcome of the learning requirement. A coach is usually someone that controls the action, deciding what needs to be done and when. A coach will provide the learner with the benefit of their experience and normally, set tasks and provides feedback to the learner on their performance. Sometimes the coaching model is seen as too restricting because the person being groomed doesn't have enough say in the learning process. Coaching in an IT environment is often associated with a need to transfer a specific technical skill quickly and spans a short period of time.

Continued on page 2



Mentoring on the other hand is seen as a long term relationship and is a two-way street where goals can be determined as a result of discussion between the two parties. The mentoree is often the one who drives the learning interaction, which is often good as it is in their best interest to know everything they should before they need to perform the tasks of a job.

In our scenario where an IT veteran is going to be leaving the organization, this person may be ideal as a coach as they may be more open to sharing their experiences and knowledge (I am sure you have people in your data center that are afraid of imparting their skills and knowledge as they believe that it may not make them as valuable to the organization, whereas a person who is retiring, may not worry about it as they will not be there).

So, if it looks as though coaching and mentoring may assist you with retaining your data center's skill and knowledge base, how do you plan and implement it?

In our next installment we look at some of the barriers that organizations have encountered when deciding to implement coaching or mentoring, and methods used to counteract them.

Join our blog that discusses your [experiences with coaching and mentoring](#) and how successful it has been.



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@elearningpeople.com

Interskill Learning Releases – April 2011

The latest Interskill 11.1 Release contained the following courses:

Database Curriculum

- The Database Series updated to V10.
- The DB2 – Database Management System Series updated to V10.

IBM Mainframe Communications Curriculum

- A course on Mainframe TCP/IP Commands to complete the IBM Mainframe Communications Series at V1.11.

zEnterprise Curriculum

- zEnterprise is a new course on the zEnterprise system.

Vendor Briefs

Product release has been relatively quiet during the last few months with many System z vendors consolidating their product links to IBM's zEnterprise offering. While there is a lull in new software and hardware, it gives us an opportunity to include some details of the major vendors' involvement with cloud computing.

IBM

zEnterprise - IBM has kept its promise of continued new offerings for its zEnterprise system by releasing the IBM WebSphere Datapower Integration Appliance XI50 for zEnterprise. This appliance sits within the zEnterprise zBX and is used to process XML and web service data, making it a valuable resource for organizations that have their data stored across different servers on multiple platforms.

Storage - Recently, IBM released its mid-range Storwize Rapid Application Storage solution, a disk system touted to improve disk utilization and application throughput, whilst providing the storage administrator with powerful tools for easier implementation and management.

Cloud – To give you an indication of how much cloud computing has infiltrated IBM, I typed “cloud” into their website search engine, and over 1.3 million results appeared. So what is IBM's current involvement with cloud computing? IBM has recently announced its intention to release “SmartCloud”, its next generation enterprise-class, secure cloud offerings for private, public and hybrid clouds based on IBM hardware, software, services and best practices. The announcement also describes other offerings:

- IBM Workload Deployer, which will provide organizations with a single platform from which they can provision their workload requirements.
- IBM's involvement in the Cloud Standards Customer Council, which will tackle issues such as security, compliance, cloud management, and reference architectures.

CA

Workload Automation - CA has released its latest version of CA Workload Automation that includes greater support for J2EE, relational databases, and web services. Through its dynamic workload placement capabilities, workloads can be processed across different physical and virtualized environments as the need arises.

Cloud – CA technologies is assisting organizations in building their own cloud infrastructure through their CA 3Tera AppLogic solution. CA are promoting their CA 3Tera AppLogic solution by stating that creating a cloud infrastructure can take as little as four hours, then by using a drag and drop feature, applications can be composed, deployed, and managed very quickly.

BMC

DB2 - BMC has upgraded a number of its products relating to the performance and tuning of data on DB2. These products include BMC SQL Performance for DB2, BMC System Performance for DB2, and BMC Recovery Management for DB2. This release of products along with IBM's improved performance in DB2 V10 for z/OS provides customers with significantly quicker access to important DB2 data.

Cloud – BMC has recognized the complexity of cloud systems by introducing their Proactive Operations Solution, which is based on the Business Service Management (BSM) platform. This solution is used to manage physical, virtual, and cloud infrastructures through planning, predictive analytics, and preventative automation facilities. BMC has also added new capabilities to their BMC Cloud Lifecycle Management product, which provides organizations with a tool to more easily control and manage cloud resources.

Editors Note on Cloud Jobs

With every new phase of IT evolution comes a ream of new jobs that are required to implement, install, manage, and solve problems associated with it. The cloud is no different, so if you are looking to keep ahead of the pack then one of these emerging cloud titles might be for you: Cloud Architect, Cloud Coordinator, Cloud Computing Manager, Cloud Integration Specialist, Cloud Security Manager, Cloud Service Manager and Cloud Solutions Manager.

Also, keep an eye out for our [Introduction to Cloud for Data Centers](#) course to be released later this year.

z/OS 1.12... What's in it for me?



By Greg Hamlyn

Now that the dust has settled from z/OS 1.12 and there has been time to absorb its content, what are the parts that are going to be most beneficial for you? In this article, rather than just spout the features and benefits of z/OS 1.12, we look at what it means to various IT professionals within the data center.

For Operators

New system functionalities have been integrated into z/OS 1.12, automating tasks which in the past may have been performed by the operator. XCF can now detect when a sysplex member becomes impaired (stalled with no signs of activity) and if it deems that it cannot recover, it can automatically vary it out of the sysplex. Outstanding WTOR messages have the ability to cause problems so z/OS 1.12 now provides an auto-reply policy, where a reply value and time delay for specific WTOR messages can be defined. Within SDSF, the SR option now displays those columns containing auto-reply definitions.

While we are on SDSF, you can now display WLM initiator details through the INIT option, and job accounting information is now displayed in several job-related SDSF screens. Support for JES3 continues to expand in this release with the availability of the printer panel within SDSF.

Several new and modified commands are included with this z/OS release. The DISPLAY AUTOR command can be used to display the auto-reply definitions mentioned above. Other display commands you may be interested in are the DISPLAY IPLINFO, *sysparm* command, which allows you to display the value of any system parameter that can be specified in IEASYSxx. Devices attached to a logical control unit are not necessarily in consecutive numerical sequence making it difficult to enter commands to this group as a whole. Enter the new VARY CU command, which can vary all devices to a control unit, online, offline or in the case of tapes, auto-switchable. The control unit number you require for this command can be obtained through the D M=DEV command, which can then be used in the D M=CU command to display devices attached to the control unit.

So for operators, z/OS 1.12 brings some automation, additional SDSF displays, and a number of new and modified commands used to display and control system components.

For System Programmers

Several new items are available in z/OS 1.12 that allow System Programmers to better manage systems. If your organization utilizes health checks, then the new CKH option within SDSF might be of interest as it displays details of previous executions of health checks. From z/OS 1.11 the IBM Health Checker has also supplied remote checks allowing the Predictive Failure Analysis (PFA) to report on possible system problems. z/OS 1.12 adds the ability for PFA to monitor the rate of SMF record generation and send you warnings if they are abnormally high.

The next step after identifying a possible problem is to analyze the system for a number of common problems. In z/OS 1.12, Runtime Diagnostics is available and can be started using the START HZR,ANALYZE command. This will invoke a number of tasks including critical message analysis, enqueue contention checking, CPU analysis, address space locking contention, and loop detection. The output produced can be sent to a file as well as the log and suggests actions you should take, which for example could be that jobs be cancelled or that specific areas require more intensive investigation.

Continued on page 5

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 z/OS Problem Diagnosis and Resolution and consists of 15 questions. Click the link below to start.

[z/OS Problem Diagnosis and Resolution](#)

For Storage Administrators

All Storage Administrators and System Programmers are aware of the requirement to run regular reorgs for VSAM KSDS so that deleted space can be reclaimed and read performance maintained. z/OS 1.12 introduces the CA Reclaim capability, which keeps track of empty Control Areas (CA) though indexing so that they are not referred to when accessing data. When new records are added to the data set, the empty CAs are reclaimed for use and the index is updated to include that CA.

During the last few z/OS releases, Extended Address Volume (EAV) disks have been providing support for different types of data sets. In z/OS 1.12 EAVs can now be used for basic and large format sequential data sets, PDS and PDSEs, BDAM, VVDS and BCS. JES2 and JES3 can also take advantage of these large volumes, using them for checkpoint and spool volumes.

For Organizations

z/OS 1.12 also heralds in IBM's zEnterprise system, which is able to manage workloads across System z, and select POWER7 and System x servers. IBM has already shipped the first of these systems to global reinsurance company, Swiss Re. As you would expect from a major system release such as this, the processing power and workload capabilities have also been considerably improved, making it an attractive proposition for data intensive processing organizations.

So as usual, there are improvements on a number of levels. Some will affect you greatly and others may resolve issues that you have had with previous z/OS versions, while other items you will find will just make your life that little bit easier.

Let us know what you have benefited most as a result of your migration to z/OS 1.12, through our blog [z/OS 1.12... What's your take?](#).



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@elearningpeople.com

Learning Spotlight – Linux on System z

In this edition of the Learning Spotlight we focus on Linux on System z. The acceptance and implementation of this product within the data center structure has now become mainstream with a multitude of case studies testifying to benefits ranging from cost savings associated with the consolidation of equipment, reduced complexity, and greater server utilization.

Understanding how Linux is implemented and managed in a System z environment is imperative for those IT personnel in your organization that are required to interact with it. For those starting off on the Linux on System z path, we are providing access to the [Introduction to Linux on System z](#) module from our Virtualization curriculum. This module discusses the use of the popular Red Hat and Novell SuSe Linux distributions and describes how they are configured and can be accessed.



MANAGEMENT:

Is This The End of Tape?

By David Stephens, Lead Systems Programmer at Longpela Expertise

In past years there was no choice: mainframe sites needed tape. However improvements in disk and virtual tape technologies today offer ways to run a tapeless mainframe. So is this the end of tape?

Tape has always been an essential and integral part of every mainframe environment. From old reel-to-reel 3420s to the current cartridge-based units, every mainframe site has relied on tape. For decades it has provided a cheap, safe, reliable way to store large quantities of backups and archival data, and to transfer data from site to site. In fact until the last few years, the only way to receive z/OS software was on tape.



However tape today is less important than it once was. Disk subsystems offer more storage for the same money, and in recent years have become almost bulletproof. Many sites mirror their production disk to remote disaster recovery sites, and virtual tape servers seamlessly store data intended for tape on disks. So are tapes now obsolete?

The Decreasing Importance of Tape

Traditionally, tape has been used for four main tasks:

- **Backups:** of individual databases, application data and entire disks.
- **Storing archival data** – data that must be retained, but is rarely used.
- **Information transfer** – moving data from one site to another.
- **Emergencies** – backups on tape can be used to restore a system on a remote site, and standalone utilities residing on tape can be used to recover a failed system.

However the importance of tapes in these areas isn't as great as it was.

Backups

Most sites still perform regular “full pack” backups of their disk subsystems as protection in case of a disk failure. However current disk subsystems use RAID technology and other technical features to the point that disk failures almost never occur. The “head crash” that in the past would kill a disk today is not even noticed. Some argue that such full pack backups are no longer required.

Regular backups of application data, databases and database logs are also made in case of application or operator failure. With the decreased cost of disk, these backups could be stored on disk.

For those with a remote disaster recovery (DR) site, the cheapest way to restore data at that site is to physically transport tape backups (the Pickup Truck Access Method, or PTAM) and restore from them. However sites with a need for faster recovery can use mirror facilities available in most disk subsystems such as IBM PPRC and HDS TrueCopy. These will automatically copy or mirror the disk contents onto a remote subsystem. The DR site can have up-to-date information stored on already-running disks, ready to be used in the case of a disaster. This technology can reduce the downtime of a full production site disaster to minutes.

Storing Archival Data

The cost per byte of disk has continued to plummet over the past years – faster than the decreasing cost of tape. Nowhere has this been faster than in the mainframe arena, where the introduction of RAID systems based on standard low cost SCSI drives in the 1990s dramatically reduced mainframe disk costs.

Continued on page 7

This has been matched by reduced environmental and footprint requirements for disk subsystems (remember the old DASD farms of the 1980s?). So it's now a very real option to store all archival data on disk.

Information Transfer

If information needed to be moved from one mainframe site to another, tape used to be the only option. This could be:

- To receive software or software updates.
- To send dumps or other diagnostic information to software and hardware suppliers for problem diagnosis.
- To send information between companies.

Today this is no longer required. Software and hardware vendors all provide a way to download software and software updates, and upload diagnostic information over the internet. Company to company data transfers can be done simply and securely over the internet using TCPIP and related tools available for the mainframe.

Emergencies

Tapes are handy in an emergency. Leaving aside their use to restore data from backups, there are a few utilities that can be IPLed from a tape when z/OS is down. For example:

- Both IBM DFSMSdss and Innovation FDRABR have a standalone program that can be IPLed to restore a disk from a backup.
- New Era Software [Stand Alone Environment \(SAE\)](#) is a utility that can be used to edit datasets and make other changes when z/OS is down.
- Jan Jaegar has also written [a utility that can edit datasets](#), but this one is free.

However, once installed these utilities can also be IPLed from disk.

Virtual Tape Servers

Virtual Tape Servers (VTS, sometimes also called Virtual Storage Managers or VSM) are designed to reduce or eliminate tape mounts. They achieve this by pretending to be a tape subsystem but

storing the data, at least initially, on disk. For a data write, a VTS would work like this:

1. The application writes data to a device. No change required.
2. To z/OS this device is a tape. z/OS sends a write operation to this device. Again, no change required.
3. The device is actually a VTS that is pretending to be a tape drive. The data is written first to the VTS cache and then to disk.

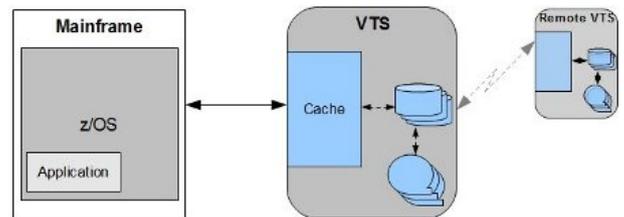


Figure 1: A VTS System

Some VTS subsystems include tape subsystem functionality: either as part of the unit or as an external tape library. These systems could subsequently migrate data on the VTS disk to tape, clearing VTS disk space.

The reverse occurs for a read from the VTS:

1. The application creates a request for data. No change required.
2. To z/OS, this device is a tape. z/OS sends a read operation to this device. Again, no change required.
3. The VTS looks at its internal catalog. If the information is in cache, it is quickly sent to z/OS. If the information is on disk, it's moved to cache and sent to z/OS. If it is on tape, then that tape is mounted, the information copied to disk and cache, and then sent to z/OS.

In short, everyone at the z/OS end think that they are writing and reading to tape. However they are really writing and reading to a subsystem pretending to be tape. The advantages of VTS are impressive:

- Because cache and disk are far faster than tape, z/OS tasks accessing tape can run much faster.

Continued on page 8

- A VTS can 'pretend' to be many tape drives: up to 256 in most cases. No more waiting for idle tape drives.
- A VTS can greatly reduce the number of tapes required by intelligently stacking information on each tape. Of course some tape management software such as DFSMSrmm and BMC Control-M/Tape can also achieve this.
- VTS have sophisticated management utilities to control the VTS. This allows users to perform tasks such as creating rules for writing data to tape and creating tapes for offsite storage.
- VTS have mirror facilities that are similar to disk subsystems. So a production VTS can have a mirror VTS at a DR site.

Using a VTS isn't the only option to reduce or minimise tape usage. A budget alternative is to use DFSMSsms to route tape requests to disk, possibly for later (and more efficient) archival to tape using IBM DFSMSshsm or Innovation FDRABR.

A Route to Tape Removal

So today it is possible to eliminate tape from a mainframe site. This could be done by:

1. Buying enough disk to hold the entire tape library, copy all tape data to disk, and change the DFSMS configuration (or all JCL) to remove tape reads and writes. An alternative would be to implement a VTS to replace all tape drives with enough internal disk space to hold the entire tape library.
2. Implement mirroring of the production disk subsystems and VTS to a DR site.
3. Make sure that all standalone utilities are on at least one disk, and their location well documented.
4. Implement electronic file transfer facilities for any transfers between companies or organisations.
5. Order all software electronically.
6. Send all dumps and diagnostic information electronically to software vendors.

But is this really a good idea?

The Case for Tape

The big advantages enjoyed by tape have always been cost and reliability. Although the cost of disk storage is decreasing, the amount of data that needs to be stored is increasing just as fast. Megabyte for megabyte, tape remains the cheapest way of storing data – cheaper than disk by a factor of 10 or so.

With the push for greener data centres, tape wins again. While spinning disks are continually consuming power and generating heat, tapes have no such overhead.

Although disk subsystems are very reliable, the possibility of failure is still present. This can be a media failure, subsystem failure (such as a microcode error) or environmental failure (such as a machine room flooding, as happened to Dallas County in the United States a couple of months ago). For essential application and database backups, tapes are still ideal. Sites without access to expensive mirroring facilities can use tapes and PTAM to restore systems at the DR site.

Sites with tape can also use their tape drives for standalone utilities – useful for worst case situations, such as when the disk holding the utility is the disk in trouble.

Tape media development has by no means ceased, and tape storage densities continue to increase (currently up to 1TByte uncompressed per cartridge). In January IBM, in partnership with Fujifilm, announced a new prototype tape that stores 39 times more data than of today's tape. But perhaps the biggest thing going for tapes is that they give piece of mind. Tape is the most reliable way of storing information. It's much harder to accidentally delete a file on tape, and there's no need to worry about power failures, disk crashes, microcode problems, or network outages. Data on tape can quietly stay in a room at the back for decades (up to 30 years for current cartridges).

Continued on page 9

The Case Against Tape

As anyone who has worked in a data centre knows, tapes require a lot of management. The more tapes, the more management required. Many data centres have tape libraries ranging into the tens of thousands – and growing quickly.

These tapes quite often move between data centres, and between a data centre and an off-site location. This is all done manually, with the potential of operator error. What's more, tapes won't help if no-one knows what is on them. So tape management software is required, and off-site tapes must have accompanying documentation showing what is where.

Tapes can break and fail, particularly heavily used tapes. This is usually managed by:

- Regularly replacing older or higher-used tapes.
- Keeping track of any data errors on tapes.
- Ensuring high-use data remains on disk.
- Keeping backups of important tapes.

It's also much easier to lose a tape. Late last year Zurich announced that it had lost a backup tape in South Africa holding details of some 51,000 customers and other parties. This shows that tapes must be secured, and preferably encrypted to prevent the data landing in the wrong hands.

When migrating to a new tape format, there may be significant work in migrating legacy tapes to the new format. Although current tape libraries all support many older formats, it's unlikely that all formats will be supported forever.

Conclusion

Although technically a tapeless mainframe site is possible, it will continue to be the exception for some time to come. With the increased cost and risk involved, most sites will stick to the tried-and-true. Tapes are low-tech, and have far fewer failure points than the alternative tapeless solutions. Having regular, up-to-date backups and archival data on tapes that are well managed and regularly checked will give any manager piece of mind.

In many mainframe sites, tape is enjoying a renaissance with data centres under pressure to reduce costs. However options to reduce the size of tape libraries and the number of tape drives will always be appealing. Increased usage of disk for frequently accessed archival data, together with VTS or other technology to intelligently pack and manage tapes will be on the wish-list for many mainframe managers.

David Stephens

Source:

Longpela Expertise, *LongEx Mainframe Quarterly – August 2010*, Retrieved from <http://www.longpelaexpertise.com.au/ezone/GoodbyeToTape.php>

In the next edition of Interskill Learning Mainframe Quarterly:

Retaining Data Centre Skills and Shortage – Part Three – We look at the barriers organizations encounter when deciding to implement coaching or mentoring, and methods used to counteract them.

Learning Spotlight – DB2 V10 – We look at DB2 V10 and provide you with access to our new modules.

Improving your Training Skills and Knowledge Retention – Introducing Interskill's Mainframe Presenters Series, which integrates stand-alone PowerPoint presentation templates and Webinar frameworks with up-to-date product e-learning.

Vendor Briefs – What have the major mainframe vendors on offer?

The 10 Benefits of Lifelong Learning – An article discussing the culture of lifelong learning, and the benefits it brings to individuals for both work and home environments.

Tech-Head Knowledge Test – DB2 V10

A Sample z/OS Disaster Recovery Plan –

Mainframes are important. Very important. Critical, in fact. So what do you do when they become unavailable – when there's a disaster?

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.