

Released February 17th, these 30 IBM z Systems mainframe credentials for Mainframe Application Programming, Mainframe Operations, Mainframe Specialists, Mainframe Security and Mainframe Storage are the new mainframe computing career benchmarks.

- Read more about the IBM Open Badge Program [here](#).
- IBM Offers hundreds of badges overall, with 35 being mainframe specific. Search for the 30 new mainframe badges (and 5 old mainframe badges) [here](#) using the term “mainframe”.

Mainframe Organizations should utilize these official IBM mainframe industry credentials to exponentially increase the effectiveness and validity of their mainframe workforce training programs, as well as insuring their mainframe workforce is at or above IBM endorsed, objective industry standards.

IBM introduced IBM Open Badges to their own internal mainframe workforce training program and within a few months they saw the following benefits:

- 125% increase in new participants
- 226% increase in course completion rates
- 694% increase in course exam pass rates

Mainframe Professionals should begin earning these official IBM mainframe credentials now! They are the recognized, respected and valued benchmark for z Systems mainframe knowledge and skills in the global mainframe computing industry! They offer recognition by your organization, by your peers, and by the industry as a whole, for your mainframe skills and knowledge!

They insure you stand out from your peers in an industry that is currently searching for its best and its brightest!





Applying for a Data Center Job....5 years from now!

By Greg Hamlyn

Years ago in IT you could target a job role, do a bit of training, maybe some certification and be guaranteed of a job in that field for a large chunk of your life. But IT technology breakthroughs and trends, and how business is looking to leverage them is moving much more quickly today. So, how can you stay ahead of the game and ensure that you will have a job in 5 years time? In this article we will look at future data center jobs from a number of different perspectives, providing you with ideas on making yourself indispensable and employable.

Technology

Consider how you were using your mobile phone 5 years ago...chances are you are now streaming video, making travel bookings, or purchasing goods that you wouldn't have been doing in 2012! Data center technology and automation has also advanced during this period, allowing systems to become more self-aware, tapping into other resources that can circumvent potential problems. Other trends such as analytics, Internet of Things (IoT), mobile apps, and cloud structures have affected how many data centers now operate. So, what will a data center job look like in 5 years time?

All this automation means that fewer staff are required for monitoring and operational tasks. This is likely to extend even further in the next 5 years. Data center consolidation has been high on the agenda for many organizations and with their shrinking physical footprints there is likely to be less demand for facility engineers and related jobs. But it is not all doom and gloom.

With the trends mentioned above, there are likely to be new job roles emerging:

> Data and Analytics

Several new positions will be required to manage, and process the increasing amount of data used by the organization. **Data Architects** will be required to design how existing, and new types, of data will be stored and managed, while **Data Integration Engineers** will need to identify how various data types within the organization can be integrated. Organizations will require the services of a **Data Scientist** to make sense of all this data they have at their disposal and to identify the benefit it can bring to the organization and its customers. A number of other data-type jobs relating to its creation, storage, processing,

analysing, and destruction may also start to appear. Leadership-type roles such as **Chief Analytics Officer**, or **Head of Business Analytics** will be required to report to other business leaders within the organization, to explain how data use is being used to meet the organization's goals.

> Programming

Application programmers comfortable with API development and experienced with DevOps philosophies will most likely be in high demand. Those familiar with programming languages popular with cloud technologies (for example, Java, Python, Perl, and Ruby on Rails) will also find it easier to stay on the payroll.

> System Specialists

IT planning specialists and other system experts will be needed to continually analyse the organization's processing and network requirements, especially when dealing with cloud providers.

> Security

As organizational data is opened up to more people (using new mediums... think technology being worn on the wrist these days), the ongoing need for various security-related personnel will be required to reduce the risk of cyber-attacks (for example, cloud security).

Take away #1 - It will no longer be an asset to the business to be an expert in just one technology. Like any other creature, you will need to evolve over time if you are to survive. You will need to at least understand general concepts across different products and technologies and can link them to identify how they can best serve your needs. Keeping an eye on emerging technology trends will be a must.

Business

Chances are that in your role today, decisions on technologies used by your organization are made at a higher level and then are filtered down to you to implement at the IT level. In the future you will be called on to present possible technology solutions to various groups within your organization and show how they align to the direction that the business is taking. For those that are pure IT geeks (think Sheldon Cooper from the Big Bang Theory) this means increasing, uncomfortable as it may be, your level of interpersonal skills.

Companies will want people with foresight who can link business visions and formulate how IT can be structured to meet them. So, where many IT specialists in the past have been cocooned in their own areas, they will need branch out to determine how what they do can benefit the business.

Where work may be offloaded to cloud providers, the data center specialist of the future may be called upon to liaise with these groups on behalf of the organization to ensure they are receiving value for services, that they are meeting service level agreements. These people will also need to be proactive in identifying ways to better perform tasks, or save the organization money.

Take away #2 – For many IT people, they will need to obtain and practice soft skills such as: negotiation, assertiveness, and will need to be financially savvy. They will need to be able to present findings to various groups and clearly explain in the audience's terms how what is being presented will benefit them.

Where you need to be

It used to be that you could be the best

techie around, have no social skills but still be indispensable. Those days are fast coming to an end. You need to have a broad spectrum of skills and knowledge so that you can match current and future business requirements with either current in-house IT resources, or through vendors.

For areas outside of those discussed in this article (for example, operators, job schedulers) it might be a good time to show interest in some of the emerging disciplines. This might include obtaining some certification, or at the very least a training course.

Take away #3 – You should start to venture out from your current abode to see how other business areas use, or could use your expertise. For many IT people, this could mean stepping out of their comfort zone, which is much better than lining up on the unemployment queue.

You – Job Interview – 5 years from now

So, you have seen the types of skills you will need going forward if you want to stay a data center employee, but how would you handle an interview for one of these jobs?

What experience do you have?

One of the obvious questions, but as well as providing your IT/technology skills and knowledge, you should be prepared to show how in your previous/current job you worked with other business units and the benefits gained (from your own perspective, the business unit, and the organization). Also describe larger organizational projects you may have been involved with such as disaster recovery policies, security strategy and standards. Describing some of your trouble-shooting techniques and experiences goes a long way to showing that

you can handle anything that comes your way...no matter what the job.

You have a wide range of skills looking at your resume. How can these help you in this job that you are applying for?

Another opportunity to describe how your knowledge of various disciplines provides you with an appreciation of the IT function overall, and how aspects from each play their part providing the optimum solution for the business. Throw in the fact that your skills follow some of the important trends that have now been adopted as standard in data centers.

Describe a difficult situation and how you resolved it?

This question can be used to describe a mixture of your technical skills but just as important; the soft-skills mentioned earlier in this article. This could include negotiation, identifying value for money/competitiveness, analysing and presentation of facts and how the end result was of benefit to the business.

Where do you want to be in 5 years?

Companies will often ask a question like this to determine how motivated you are. As indicated previously, you will need to be continually learning and looking at new technologies to determine how it can benefit the organization, as well as building on your thirst for knowledge.

...and of course make sure that you don't burn any bridges. It can be quite ironic how someone that you criticized in a previous workplace ends up on an interview panel for a job you are applying for!

The Bottom Line

Many data center jobs are likely to become leaner, or disappear altogether in the next 5

years, so it is in your best interest to:

- Investigate trends and their benefit to the organization
- Network with people involved in other IT areas within the business
- Work on your soft skills

Good luck!



Greg Hamlyn, a senior technical course designer and developer with Interskill, has over 27 years experience working in the mainframe and z/OS data center environments, having performed in

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

You can contact him at:

g.hamlyn@interskilllearning.com

Vendor Briefs

In this issue of Vendor Briefs we will not only look the products that have been recently released by several mainframe players since our last update, but also gaze at the technical direction of some of the major vendors.

IBM

Not much to report in the way of major, or updated hardware or software from IBM in the last few months. What is more interesting are the projects and organizations they are working with, which says a lot about what IBM thinks about the future. Below are a few details on these initiatives/technologies.

Blockchain

For those that have not heard of blockchain, it is a distributed database/ledger that is used for authorizing/tracking items as they travel through a lifecycle (think of produce that travels from a farm, transported, stocked in a supermarket, and eventually purchased). At each stage the item can be authenticated as being handled correctly and securely, before it is moved on. IBM joined the open-source collaborative Hyperledger project to assist in improving this technology. The types of industries that would greatly benefit from Blockchain include finance, banking, IoT, supply chain, manufacturing and technology...the same as those currently using IBM's enterprise systems. IBM has even created a secure blockchain service on the IBM Cloud (residing on a LinuxONE system), that allows organizations in the fields mentioned above, to test and run blockchain applications in private. Here are some of IBM's announcements in relation to blockchain over the last few months:

07 Dec 2016 - IBM Debuts Blockchain Ecosystem To Help Accelerate Growth of Networks on Hyperledger Fabric

30 Nov 2016 - Mahindra and IBM to Develop Blockchain Solution for Supply Chain Finance

15 Nov 2016 - IBM Helps FinTechs Drive New Innovation in Payments, Lending, Blockchain and Know Your Customer

9 Nov 2016 - Kasikornbank Thailand and IBM Aim to Reduce Complexity in Corporate Credit with Blockchain

24 Oct 2016 - IBM Empowers Singapore Startup with Blockchain and Bluemix for Island-wide Federated Lockers Network

21 Oct 2016 - SBI SECURITIES Works with IBM to Test Blockchain Technology for Bond Trading Platform

19 Oct 2016 - Walmart, IBM and Tsinghua University Explore the Use of Blockchain to Help Bring Safer Food to Dinner Tables Across China

23 Sep 2016 - IBM and China UnionPay E-payment Research Institute Share Bank Card Bonus Points Among Multiple Banks Using Blockchain

16 Sep 2016 - IBM and Bank of Tokyo-Mitsubishi UFJ to Use Blockchain For Contract Management between the Two Companies

Note: IBM also have some training on this topic. [Check out the link here](#) (you will need to register, but it is free).

Watson

IBM's Watson technology is the foundation behind some pretty amazing projects in recent times. Working with this technology,

Rice University has created a robot that is designed to help assist the elderly through noninvasive heart and breathing measurements, and using face and speech recognition features. Some other things that Watson is involved in:

1 Dec 2016 - IBM and Pfizer to Accelerate Immuno-oncology Research with Watson for Drug Discovery

29 Nov 2016 - IBM Unveils Watson-Powered Imaging Solutions for Healthcare Providers

9 Nov 2016 - IBM Launches Experimental Platform for Embedding Watson into Any Device

9 Nov 2016 - IBM and Topcoder Bring Watson to More than One Million Developers

25 Oct 2016 - The Weather Channel Launches Bot For Facebook Messenger, Powered By IBM Watson

25 Oct 2016 - IBM Powers MobileFirst for iOS Apps with Watson to Boost Decision Making for Professionals

Master the Mainframe Competition

IBM's Master the Mainframe World championship has been run and won for another year. After competing at a regional level, 65 students from 26 countries were selected to participate in this world event, with the top 10 finalists gathering in San Francisco to complete and present their projects, which involved working with z/OS, Apache Spark, a large amount of financial data, and a request to link their analytical results to a business value.

The top 3 were then flown to Las Vegas for the presentation of the World Champion – Antonio Alvarez Rodriguz from Spain.

CA Technologies

CA Technologies took the opportunity at CA World (November 2016) to announce several mainframe portfolio additions and enhancements.

CA's new **Mainframe Operations Intelligence** solution uses machine learning to identify and predict system problems before they impact the organization. A new release of **CA Workload Automation** is designed to simplify running and monitoring of workloads from emerging big data, cloud, and mobile applications using an intuitive web interface. Security personnel can take advantage of new releases of CA's **Data Content Discovery**, and **Compliance Event Manager**, to identify data exposure risks and manage data in relation to compliance requirements. For DevOps teams, the new **CA Release Automation Connector for z/OS** product provides optimized delivery and deployment of application releases to mainframe environments.

BMC

BMC also took the opportunity to blow their own horn during their BMC Engage tradeshow in Las Vegas (September 2016).

BMC's workload automation tool, **Control-M** has had its capabilities expanded enabling new capabilities for highly-efficient managed file transfer, optimized rapid cloud deployment for AWS and Azure, and a new Automation API with expanded 'Jobs-as-Code' capabilities for DevOps teams. In December, BMC also announced that Control-M now provides support for SAP S/4HANA® and SAP Financial Closing cockpit (FCc.) workloads.

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

Click the link below to start.

[COBOL Programming](#)

Learning Spotlight – Introduction to Big Data

There is no denying that organizations are experiencing an explosion of data, emanating from a myriad of sources.

The module that we have provided for you here is [Introduction to Big Data](#), which describes today's information landscape, where data is coming from, the format of it, the size of it and most importantly the benefit that organizations are looking to obtain from it. No-one is going to escape some aspect of big data whether it is the identification of useful data, its storage, how it is filtered to the organization's requirements, or using analytical processes that result in a business edge.

This module also addresses the issues associated with properly managing big data including; security, data lifecycle, speed of processing and cost. All of these are factors when determining whether an organization can benefit from a big data strategy. The last section of this module discusses how IT professionals can forge a career in big data, by looking at certification, tradeshow, conferences, and training that can assist in building their skills and knowledge in this area.

We hope you enjoy it.



Management: The Best Mainframe DASD and Dataset Innovations

By David Stephens, Lead Systems Programmer at Longpela Expertise

I remember when I first walked into our mainframe machine room. It was 1989, and I'd just finished university. As I walked in, I saw this huge line of blue and white boxes disappearing down the back of the machine room – our 3380 DASD 'farm'. Today, the lines are a lot smaller, and that's just the beginning of the changes I've seen with DASD and datasets. So what are the best 10 I've seen since I started?

1. Catalogs

At the time when I first saw that line of 3380, our company was just finishing up implementing catalogs. Most of our application data was cataloged, but some applications and datasets weren't there yet.

But think about what this meant. Every time you wanted to allocate any uncataloged dataset, you need to know and specify the volume serial number. Every time! The idea of a catalog to keep track of datasets and volumes has been around since OS/360 where the SYSCTLG dataset kept a list of

datasets and their volume location. These CVOL catalogs weren't pretty: you used the batch utility IEHPROGM to catalog and uncatalog datasets, and IEHLIST to list catalog entries. VSAM catalogs came with VSAM datasets in the 1970s, and cataloged VSAM datasets only (still needed CVOL for non-VSAM). VSAM catalogs owned VSAM datasets, and you couldn't have VSAM datasets in different catalogs on the same volume.

In reality, many shops (including ours) cataloged their VSAM (they had to), but not much of their non-VSAM. This changed when the Integrated Catalog Facility (ICF catalogs) replaced both CVOL and VSAM catalogs in 1979. ICF catalogs were the way they should be – much easier to use and manage. So by 1989, all our catalogs were ICF.

2. Automated File Archiving and Management

We'd also just finished implementing IBM's Hierarchical Storage Manager (HSM). This brilliant product would automatically archive old data, remove unneeded temporary datasets, delete expired datasets, and even backup datasets. But that's not all. Archived datasets (to compressed disk or tape) were automatically brought back if needed. And it was so easy to restore deleted datasets using the TSO hrestore command. Just brilliant.

OK, today this doesn't seem so special. But then it was. Originally released in the 1970s, DFHSM wasn't the only product in this area. In 1989 we were actually moving away from CA-ASM2 for dataset archival, and other products like CA-Disk were also around. But most of these were batch oriented. DFHSM was automatic.

3. Cache

In 1989, we were using 3880 disk controllers with cache. And this cache was amazing. First introduced in 1981, it would store recent data read from, or written to, DASD in memory kept in the DASD controller. Subsequent reads for this data could be satisfied without a disk read. And it made a big difference. But IBM went further in 1991 with the 3990 controller.

The 3990 added Cache Fast Write (CFW), and DASD Fast Write (DFW). DFW allowed writes to DASD to be 'satisfied' at the cache level. Once received in cache, the DASD subsystem would acknowledge the write, and do the actual physical disk write a little later. This improved dataset write times as seen by z/OS (then MVS). The 3990 would write the cache data to non-volatile storage, so the data would not be lost during a power

failure or other cache failure. CFW was the same, but without the non-volatile storage part.

We were nervous about DFW long and hard when we bought our first 3990 controller. But once in, it really improved our performance, and we never had a problem.

In these early days, we didn't have a lot of cache to play with. So we tended to use it for the most important volumes, leaving the less-important to fend for themselves. Later models increased the amount of cache, to the point that today you would need a very good reason to have an uncached volume.

4. RAID

I remember a four-week period when we were getting a 3380 head disk crash every few days. We'd just moved our data centre, and had some old 3380 DASD. One such head crash would stop four devices (4 volumes per head disk assembly, or HDA). So we'd call in IBM, they'd replace the HDA, and we'd restore the four devices from backups. This was bad enough. However at a couple of points we lost critical volumes, bringing down all systems – no working z/OS system to restore from. Great disaster recovery experience, but stressful.

This would never have happened if we used RAID DASD. The IBM RAMAC DASD subsystem announced in 1994 was such a device, with emulation of 3380 or 3390 disks on a matrix of cheaper 3.5 inch PC-style disks. If one would fail, there was enough redundancy for the subsystem to continue processing with no problems.

However IBM were playing catch-up. StorageTek announced their Iceberg disk subsystem in 1992, though it took them

another two years to release it. EMC weren't far behind, announcing the Symmetrix system in 1993.

Interestingly, IBM later decided to use StorageTek Icebergs as their own, rebranding them as RAMAC Virtual Array (RVA) in 1996. And this was smart. The StorageTek Shared Virtual Array (SVA) architecture was the first to totally decouple the logical 3380/3390 architecture from the physical disk. Rather than setting aside 1GBytes for a 3380 disk device, the SVA would only use the storage if needed.

RAID greatly improved the reliability of DASD subsystems. But it also provided an additional advantage. Now disk suppliers could use a common platform for both mainframe and non-mainframe disk. Just needed an emulation layer so it could pretend to be mainframe DASD. Today all DASD subsystems can be configured on the floor to be mainframe or non-mainframe DASD.

5. Flashcopy

StorageTek's SVA added another new feature: Snapshot. This could create an instantaneous copy of a dataset or volume instantly without consuming any extra space. Pretty cool huh? IBM and StorageTek parted ways in 2000 when IBM announced their Shark (Enterprise Storage Server). Based on IBM's Seascapes architecture, the Shark was a small step backwards, without the RVA technology. IBM have brought back Snapshot functionality in their DASD line – now called FlashCopy. EMC and HDS supply something similar (TimeFinder and ShadowImage respectively).

6. Solid State Disk

Solid state disks (SSDs) are a great idea, providing amazing disk performance at a higher price. And this is brilliant when you really need the extra performance: MQ and DB2 logs, z/OS logger DASD logs, JES checkpoints and more. However SSDs – or disk without the moving disk parts, have been around for a long time. StorageTek started the ball rolling in 1979 with their RAM-based 4305 solid state disk. This worked like flash drives today with one difference: the disk was wiped when the disk unit was switched off (or lost power).

Memorex followed in 1983 with the 3864, and National Advanced Systems (NAS, later HDS) was a couple of years behind in 1985. However solid state disks were expensive, and lost their popularity until 2009. This was the year that both IBM and EMC released flash SSDs that could be used on a mainframe. Today their usage increases as storage providers make it easy to mix and match the (still) cheaper spinning disks with SSD.

7. Fibre Mainframe Channel

Moving data centres is fun. There's excitement, tension and a lot of stress. You also get to see a side of the hardware you've never seen before as boxes are broken up and lifted in and out of machine rooms. But the boxes are the easy part – it's all the connections between them that are the hard work.

Before 1991, you connected your DASD to the mainframe using Bus and Tag channels. A pair of copper wires as thick as a baseball bat, bus and tag channels had a transfer rate of around 3-4Mbytes per second up to a couple of hundred feet. That's what we had

when we moved our data centre in 1993. I remember helping our hardware engineer at 2am lay down these monsters under the false floor as we tried to put our mainframe back together. My arms hurt for a week after the move.

ESCON changed this, replacing the big copper cables with a single fibre cable. It increased transfer speeds to 17MB, and distances up to more than 43km. It also eliminated that upper body workout hardware engineers used to get when moving them around.

FICON was the next generation of fibre cable, with current transmission speeds of up to 1600Mbps bidirectional throughput on a 16Gbps link up to 100km for FICON Express 16S. Further distances can be considered with a FICON switch, of from a request price quotation. High Performance FICON (z/HPF) is the latest hardware/software feature to squeeze even more from FICON channels.

8. Disk Mirroring

In 1990, my mainframe site commenced disaster recovery tests. Assuming that our data centre was no longer available, we would get all our offsite tape backups, board a plane and fly 800km to Sydney, and proceed to restore our backups at another data centre. This worked, but took us over 48 hours. 48 hours can be too long. IBM, EMC and HDS all solved this problem with disk mirroring technologies – echoing any changes to local disks on remote disks. With this technology, that 48 hour recovery time can be shrunk to minutes. Our article “A review of Disk Mirroring: XRC, PPRC and Friends” looks at disk mirroring in more detail.

9. System Managed Storage

VOL=SER=VOL001. Must have seen that code so many times in JCL. Not to mention hard-coding volume serial numbers when allocating dataset using ISPF 3.2 or something similar. Catalogs removed this coding for existing datasets. Eligible Device Tables (EDTs) also helped, so you could replace your VOL=SER= code to UNIT=PROD1 (assuming your systems programmer has defined an EDT called PROD1 with some DASD volumes).

But coding the VOL=SER wasn't the hard part. The hard part was looking through all your disk devices for enough room for your datasets. Or keeping a list of those devices, and how full they are. Painful.

So System Managed Storage (SMS) was great. Forget about VOL=SER=, forget about keeping lists of DASD. In fact, forget about LRECL=80, RECFM=FB or DSORG=PS. Rather, let the DASD administrator look after all of that. Introduced with MVS/ESA V3.1 in 1988, SMS was a renaissance. Today we forget how hard it was to manage DASD before SMS. But now DASD administrators can not only better control DASD, they can stop users from mis-allocating datasets, using volumes not designed for them, and allocating un-cataloged datasets.

10. DFDSS

Didn't expect this one did you? For number 10, I considered Parallel Access Volumes (multiple UCBs and paths to a device for improvement), and Extended Address Volumes (EAV) – more space on logical 3380/3390 devices. I also gave serious thought to ISPF 3.4 for dataset management. But I couldn't leave out DFDSS, or DFSMSdss as it's called today. Innovation also have a

similar product: FDR/ABR – but I’ve never used it.

Before DFDSS, moving datasets was done with terrible utilities like IEHMOVE and IEBGENER. But these couldn’t move VSAM. So add IDCAMS. All are difficult to code, slow and just painful.

Enter DFDSS. Simply move or copy any dataset you want – VSAM, sequential, tape or disk. It doesn’t matter. And it will do it faster than anything else. But that’s just the start. Want to defragment a pack? Done. Want to copy an entire volume, or more than one? Done. Want to backup a volume or dataset– no problem. Want to restore a volume when you don’t have a working system up – sure, use the standalone version. Want to backup several datasets at the same time (not one after the other)? Can do. DFDSS also has kept up to date. So it uses FlashCopy automatically, and can work with CICS BWO (Backup While Open) to backup datasets while they’re open to CICS. DFDSS just has to be the most under-rated utility today.

Source:

LongEx Mainframe Quarterly - May 2016,
Retrieved from
<http://www.longpelaexpertise.com.au/ezone/BestDASDIInnovations.php>



Technical: Disk Mirroring Part 1: A Review of XRC, PPRC and Friends

By David Stephens, Lead Systems Programmer at Longpela Expertise

Ever heard of PTAM? It stands for Pickup Truck Access Method. This is where you load your backup tapes onto a truck and transport them to another site when you lose your data centre. At that remote site, you restore your production systems and continue processing.

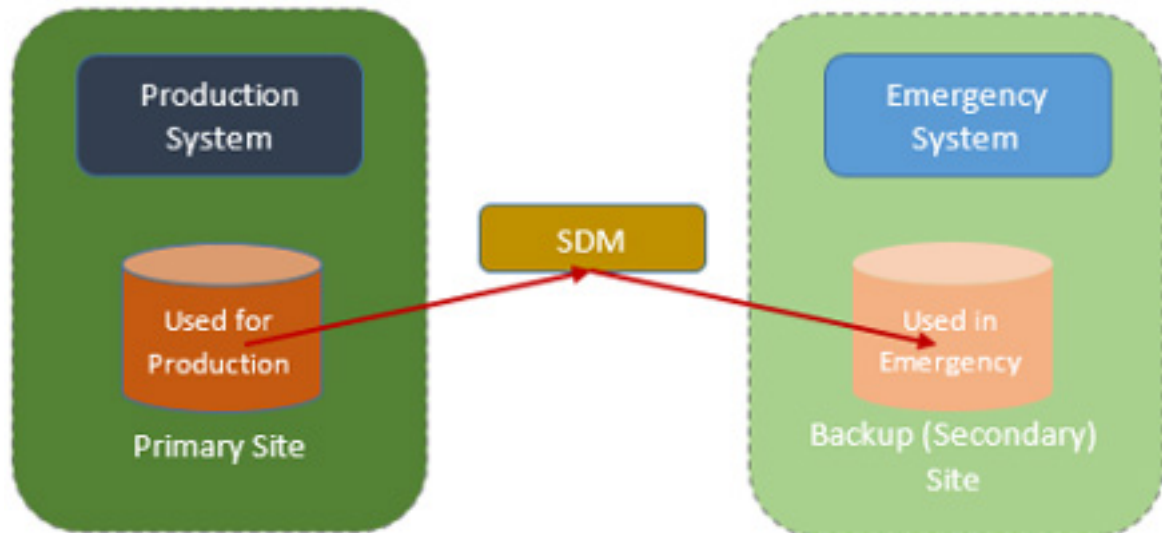
There are two problems with PTAM. Firstly, it’s slow. I used to work for a company that would regularly do PTAM DR tests. It took 48 hours. The second problem is that you only restore data to the last backup. So if you take daily backups that go offsite, you’ve lost all the changes from that last backup. If PTAM doesn’t work for you, you’re going to need to mirror your DASD – continuously create an up-date copy. In this first of three articles, we review the options to achieve this.

z/OS Global Mirror

In 1994, IBM announced the 3990-6, the latest DASD controller for 3380 and 3390 DASD. Although it took them a year to deliver the new controller, it provided a new features for mirroring a DASD device to a remote location: Extended Remote Copy (or XRC).

Today XRC is called z/OS Global Mirror, or zGM for short. But most technical people still call it XRC (including us). So here's how it works:

There are two systems: a primary system that runs production, and an emergency or



recovery system. The recovery system has one disk volume for every production volume to be mirrored. A system data mover (SDM) detects all updates to each production volume, and re-does them on the corresponding recovery volume. So you have a disk at the DR site that is pretty much the same as the production volume.

SDM is really another name for DFSMSdfp running on z/OS – this is what does the heavy lifting in terms of moving updated data between primary and backup volumes. So XRC is a combination of DASD and z/OS software. In the diagram SDM is in the middle of the two systems. In reality the SDM could run on either

- The production system (if it is z/OS)
- Another z/OS system at the primary site
- The emergency system (if it is z/OS)
- Another z/OS system at the backup site

Normally you'll want SDM at the backup site. Although the SDM runs on z/OS, it isn't limited to z/OS DASD. It can also mirror z/VM and z/Linux (both running native, and as a z/VM guest) disk volumes.

With XRC, there is no theoretical limit as to the distance between primary and backup site. Anything less than 100km can use FICON links between them. Anything over 300km or so will need a channel extender: allowing FICON to work over TCP/IP networks. There are a

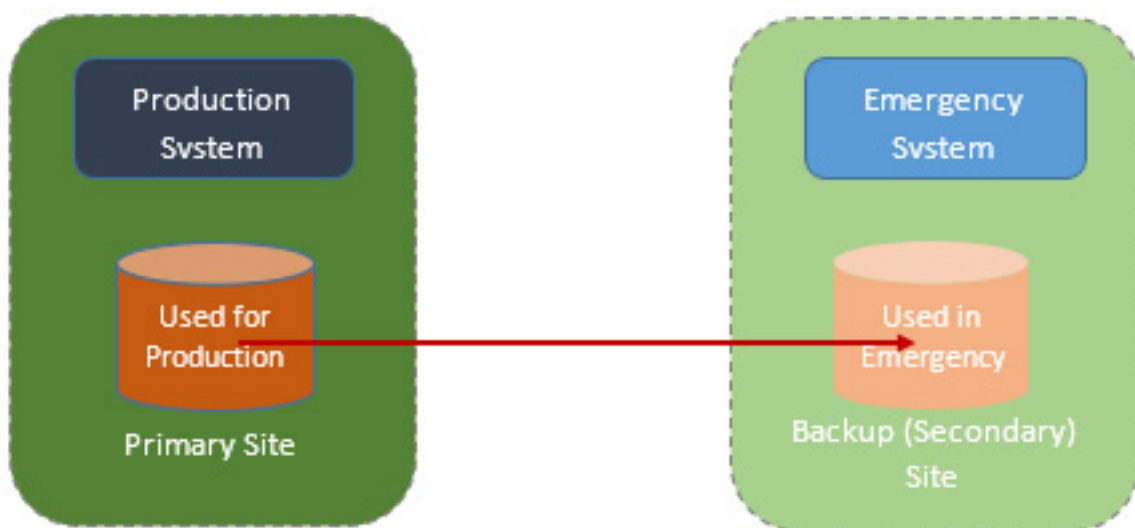
couple of vendors selling these extenders including Brocade and Cisco.

Although XRC was introduced with the IBM 3990 disk controller, it isn't an IBM-only option. Both EMC and HDS also support XRC. In fact you could mirror an IBM disk volume to an EMC volume.

XRC is an asynchronous solution. This means that when an application does an I/O, it doesn't wait for that I/O to complete on the remote system. So if the production system were to fail, the remote system maybe a little behind the times. How far behind depends on several things, including XRC performance, network bandwidth (if using channel extenders), and the distance between the data centres. It's possible for the remote disks to be a couple of minutes behind. This is data lost if a disaster happens. For some organisations, a couple of minutes isn't good enough.

Metro Mirror

When the IBM 3990-6 controller was announced in 1994, it didn't just introduce XRC. It also introduced Peer-to-Peer Remote Copy (PPRC). Today PPRC is called Metro Mirror, but we still call it by its old name. PPRC was a little different to XRC. Here's how it worked:



The diagram looks pretty boring doesn't it? Just a line from the production volume to the backup volume. PPRC doesn't use DFSMSdfp or z/OS. In fact, it is completely removed from the operating system (though there are z/OS commands to control PPRC). Rather, it is a solution implemented by the DASD controllers. So the DASD controllers (included in DASD subsystems today) detect changes to production volumes, and instruct the remote controller to re-do them.

PPRC is a synchronous solution. So when an application writes to a production volume, it must wait for that write to be made to the backup volume. In fact, the DASD controller won't even pass control back until the write is complete on both primary and backup systems.

The good news is that you won't lose a minute or two of data if the primary system goes down: the emergency system will always be up to date. The bad news is that this is expensive. In one site, we have seen I/O service times increase by a factor of 5 for a PPRC link over 30km (about 20 miles).

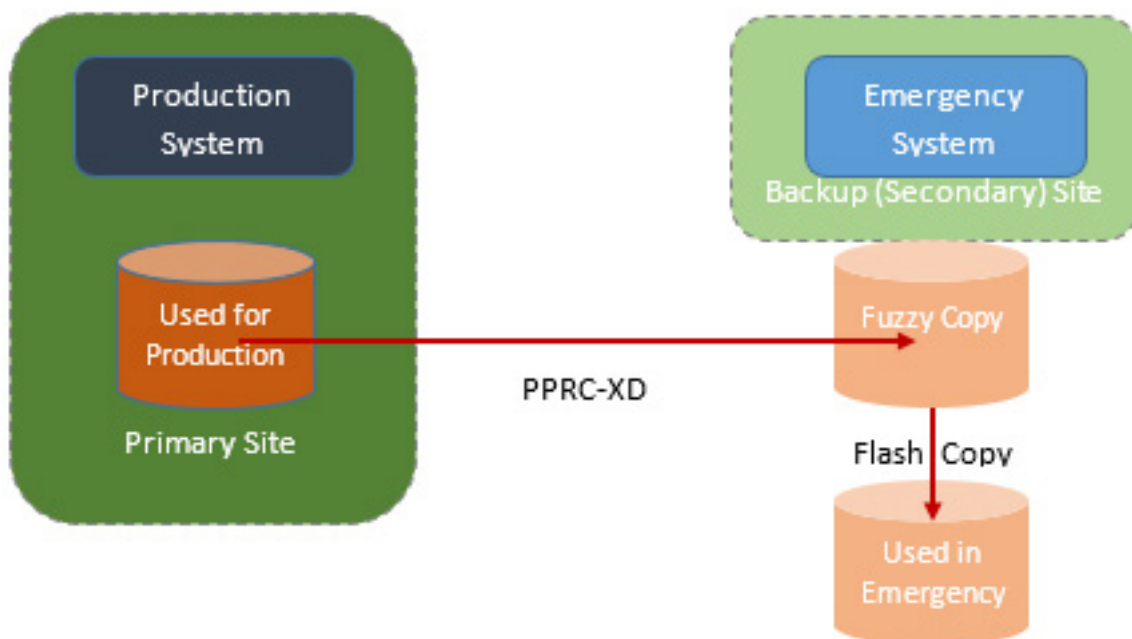
This brings a second limitation to PPRC – distance. Theoretically, you can have a PPRC link over any FICON channel (no channel extenders) – so you can have the two sites up to 300km apart. In practice, this is too far, so your backup site must be closer.

Both HDS and EDS have their own facilities to do the same processing as PPRC: TrueCopy for HDS (also known as Hitachi Remote Open Copy, or Hitachi Remote Copy), and SRDF (Symmetrix Remote Data Facility) respectively. However they don't mix. So you can't mirror an IBM volume to an HDS volume at a remote site.

Global Copy

In 2002, IBM announced PPRC-Extended Distance (PPRC-XD). This is basically PPRC, but asynchronous. So the DASD controller at the primary site will replicate DASD updates at the backup site without z/OS intervention. However the application doing the I/O doesn't need to wait. Today this is called Global Copy. However PPRC-XD has a problem – it's a fuzzy backup. The DASD controllers can't guarantee integrity. So if the primary site fails, the backup site may not be workable.

To fix this, PPRC-XD is used with Flashcopy (IBMs way of instantly making a copy of a disk). Here's how it works:



PPRC-XD makes asynchronous backup. However periodically the primary and backup controllers communicate to make a consistent copy at the backup site. When this completes, a copy of this consistent copy is made using FlashCopy. This is the volume to use in an emergency.

Today PPRC-XD is called Global Copy. When combined with FlashCopy it is called Global Mirror (previously Asynchronous PPRC). Again, HDS offers a similar TrueCopy Extended Distance (also called Hitachi Universal Replicator), and EMC SRDF/A (SRDF/S is sometimes used for the synchronous copy feature of SRDF).

Conclusion

So now we can get an up-to-date (or almost) copy of our primary disk at a remote site. However the performance of XRC and PPRC can be a big issue. In our follow-on article, we'll look at the performance of DASD mirroring, and its impact on production systems.

Source:

LongEx Mainframe Quarterly - May 2016, Retrieved from
<http://www.longpelaexpertise.com.au/ezone/DASDMirroring1.php>

Technical: Disk Mirroring Part 2: Performance

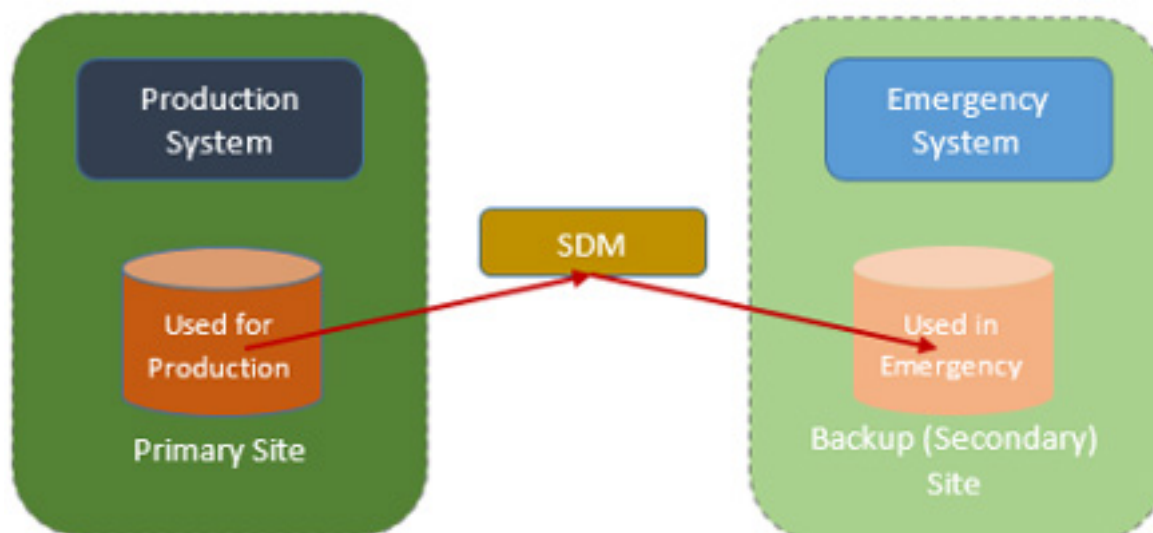
By David Stephens, Lead Systems Programmer at Longpela Expertise

In Part 1 of this series of three, we looked at the basics of DASD mirroring: XRC and PPRC (or Metro Mirror and Global Mirror as they're now known). And this technology is great, providing a continuously updated copy of a disk on another subsystem at a remote location.

What you may not have thought of is that performance of XRC and PPRC are important, but for very different reasons. So in this article, we're going to look at XRC and PPRC performance.

XRC Performance

So we have our XRC systems: the z/OS data mover (normally at the remote site) automatically replicating data updates from one set of DASD to another:



When we talk of XRC performance, we are concerned about two issues:

- The amount of data in-transit
- The effect on production performance

Data In-Transit

Assuming the updates are constantly being performed on the primary disk subsystems, there will be data that is 'in-transit': that has been updated on the primary systems, but not the remote systems. This data is the data that will be lost in the case of a failure of the primary systems. Ideally, we want to minimise the amount of this potential lost data. This can be done by transferring data as fast as possible between the primary and backup systems. There are a couple of factors that impact this:

- The distance between the primary and remote subsystems.
- The speed of the data link between the sites.
- The performance of the z/OS data mover.

There isn't much we can do about the distance issue. However the speed of the data link between the two sites can be improved: by improving the telecommunications infrastructure such as channel speeds and hardware. Both Brocade and Cisco offer optional features to further accelerate XRC transmission and improve the performance of the data link.

As XRC uses the DFSMS SDM as the data mover, the performance of z/OS and the DFSMS SDM are understandably really important. Most sites will use DFSMS on the remote site as the data mover for two reasons: production workloads cannot impact the performance of this remote SDM, and similarly, SDM processing cannot impact production performance.

Impact on Production

The second area of XRC performance is its impact on production and production costs.

If running DFSMS SDM on the production systems, this processing could possibly impact production as DFSMS SDM chews up CPU and other resources. The CPU consumption may also increase software licensing costs, though some of the XRC load can be offloaded to zIIP processors.

Even if running DFSMS SDM on a remote system, XRC performance could impact production performance. The worst case scenario for XRC is if XRC cannot keep up with production updates. For example, suppose the production system performed 10,000 updates a minute on a production system. If XRC can only transfer 8,000 updates a minute, then it simply cannot keep up with the production workload.

In some cases this will be normal for short periods of time: for example disk defragmentation or full disk copy operations. Other times it will spell trouble. XRC has a couple of ways of solving this.

The simplest way is to do nothing and leave things as they are. This way a site has no control over the data 'in-transit' – it could vary dramatically. If the in-transit data becomes too high, the 'link' between primary and backup device pairs could be suspended. In practice most sites with XRC have some target recovery point objective (RPO) – the maximum period where data could be lost.

For a long time XRC has provided a facility to block all I/O to a device for short periods if the amount of in-transit data exceeds a threshold: data blocking. Data blocking thresholds can be controlled by parmlib settings, and switched on or off for different

volumes as appropriate. A more elegant solution is write pacing, which slows down writes (rather than stopping all I/O). As the amount of waiting updates increases, the performance of production updates can be further decreased. Again, write pacing is controlled in parmlib members, and can be enabled or disabled as required. The XRC command XQUERY VOLUME_PACE produces a report on write pacing.

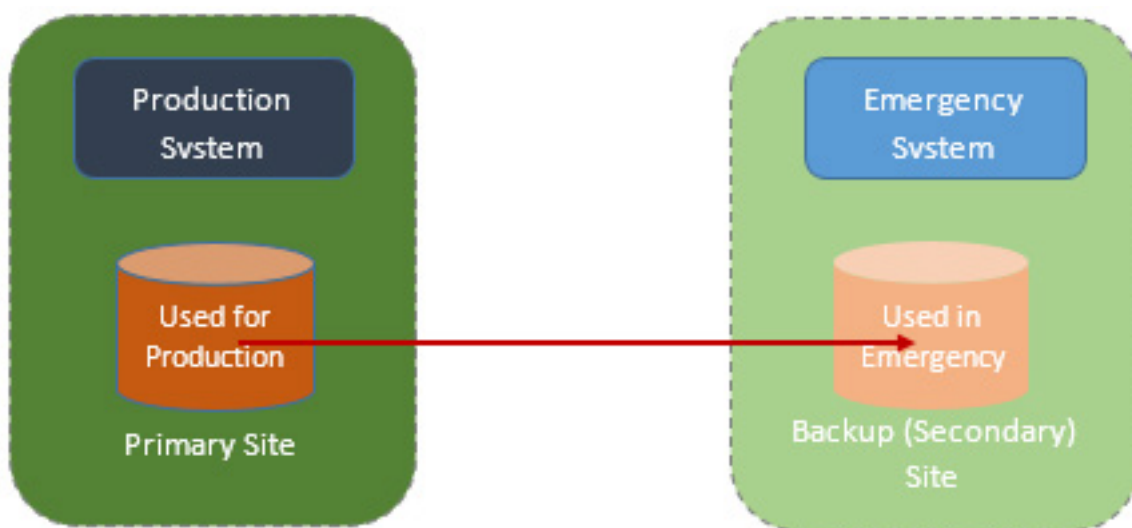
Workload-based write pacing is another step up, using WLM to tailor pacing options to subsets of address spaces or applications.

Blocking and write pacing are features to manage the I/O degradation if XRC cannot keep up with the number of incoming DASD updates. Either way, ideally XRC should comfortably be able to handle all updates, even in peak times. A first step in doing this is to look at the telecommunications infrastructure, and making sure that it is up to the job.

However there are other resources that need to be thought of. The XRC address spaces (including the DFSMS SDM) need to have sufficient CPU to operate – so their priority must be set correctly in WLM. Similarly the DFSMS SDM address spaces needs memory. The maximum number of buffers assigned can be set using parmlib settings – ensuring that enough buffers are available is important. Page fixing these buffers and other memory using the parmlib PermanentFixedPages parameter can also improve performance. There are several other parmlib settings that can impact XRC performance. A list of current XRC parameters can be obtained using the XQUERY command.

XRC writes updates to a journal before applying them on a remote system. So journal performance can impact XRC performance. There are things that can be done to maximise the performance of these journals, including defining them as striped datasets, putting them onto their own disk devices, and spreading them across disk volumes.

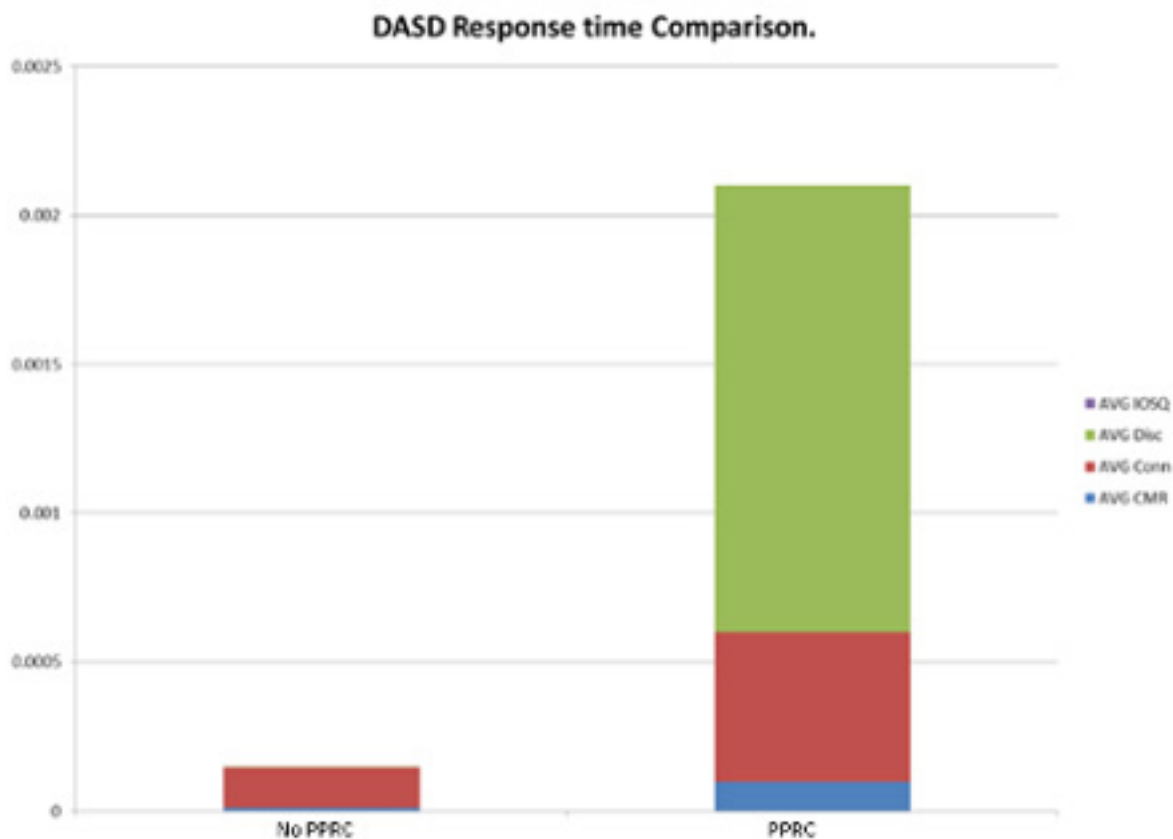
PPRC Performance



This is the same boring diagram from the last article, showing how PPRC works: a simple link where the data movement is handled by the DASD controllers in primary and remote sites.

So although the telecommunications infrastructure between these controllers is important, there is less to tune than XRC.

For synchronous PPRC, the impact on production performance is the key issue. The graph below compares the average I/O time of a write with, and without synchronous PPRC.



It's a big difference isn't it? The synchronous PPRC has increased the total average I/O wait by a factor of 5. You can see that most of the time increase is from Disconnect time (in green). The connect time (brown) has also increased, though this is probably because the no-PPRC example used z/HPF, whereas the PPRC example did not. This example is for a site where the primary and backup site were approximately 50 km (30 miles) apart. The difference (overhead) would increase markedly if the distance between the two sites was further increased.

So what can we do? The bottom line is: not much (unless you can make the distance between the two sites closer). We can make sure that the telecommunications infrastructure is the best we can make it. We can make sure that we configure the I/O as best as possible, and use features such as z/HPF to improve performance on the primary system as much as possible. And we can make sure that only disk volumes that really need the availability provided by PPRC are mirrored. But that's about it.

Conclusion

That's right – PPRC and XRC aren't necessarily straightforward (is anything with the mainframe?). One of the important issues when configuring this mirroring technology is performance – managing the potential data loss, and the impact on production systems.

Source:

LongEx Mainframe Quarterly - August 2016, Retrieved from

<http://www.longpelaexpertise.com.au/ezone/DASDMirroring2Performance.php>

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