



interskill  
learning

INTERSKILL MAINFRAME TRAINING NEWSLETTER

June 2018

## Inside This Issue:

New and Coming Releases	2
Pervasive Encryption – Why is IBM making a big thing about it?	3
Vendor Briefs	7
Tech-head Knowledge Test	11
Learning Spotlight – Blockchain Fundamentals	11
<i>Management: DASD Mirroring Part 3: Putting a Solution Together</i>	12
<i>Management: A Different Set of Goalposts for Application MQ Performance</i>	17

# Welcome to the Interskill Mainframe Training Newsletter

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

As a subscriber to the Interskill Mainframe Training Newsletter you will receive each edition. These editions will include z/OS related articles for technical experts and management, and information on upcoming releases of Interskill's Mainframe curriculum.

## New and Coming Releases:

### ★ COMING SOON

- JCL for z/OS 2.3
- z/OS 2.3 – all 40 core courses updated
- IMS Transaction Manager (TM)
- WebSphere Application Server
- Performance for Applications
- Advanced VSAM
- Application Development with MQ
- IBM Developer for z Systems
- JAVA
- z/VSE 6.3
- HPE NonStop (Tandem)

### ★ RECENT RELEASES

- IMS V15
- Tips & Tricks – advanced courses for the experienced mainframer on z/OS, JCL, JES2, TSO/ISPF, SDSF, and Utilities
- BMC – Control-M
- Agile Fundamentals
- Cryptography (Pervasive Encryption)
- Machine Learning and Spark
- Introduction to Performance Sampling Tools
- CAPM® Practice Test
- Java on z/OS for Java Programmers
- z/OS UNIX Systems Services
- Introduction to the IBM Z Mainframes (IBM z14)
- z/OS 2.2 Migration Briefings
- Blockchain
- CICS TS for z/OS 5.3
- New Assessments:
  - COBOL Basics Assessment
  - Advanced COBOL Assessment
  - PL/I Basics Assessment
  - Advanced PL/I Assessment
- **Project Management and PMBOK®**  
“PMBOK®” is a registered mark of the Project Management Institute, Inc.



## Pervasive Encryption – Why is IBM making a big thing about it?

*By Greg Hamlyn*

In late 2017, IBM released their flagship z14 mainframe, and with these types of releases the headlines normally focus on greater capacity and processing power. While this is still the case, this time around IBM is playing the data security card, highlighting the fact that the z14 has integrated pervasive encryption into its every fibre.

What is pervasive encryption, and why suddenly do we need it? What is wrong with IBM's previous encryption methods, and do you need to start looking at whether your organization should be implementing it?

This article looks at IBM's pervasive

encryption from several angles and gets you thinking about your data security needs.

### **How safe is your data?**

We all know that the amount of data stored by organizations today is enormous, and most of us are happy to believe that it is relatively safe from hackers. Surely, any organization we trust with our information is going to make sure it is safe from prying hands! In the case of banks and insurers it is fascinating that in a 2017 survey that only 3% of customers believe that their institution has experienced a data breach<sup>(1)</sup>. The reality is that in 2017 more than 2.9 billion records were leaked<sup>(2)</sup>, amounting to almost 10 billion records since 2013<sup>(3)</sup>. What is even more astounding is that it is estimated that of all those records stolen, only 4% were encrypted<sup>(3)</sup>. That's a lot of our raw data out there.

So, what are organizations doing about it? Well, they continue to strengthen firewalls, but hackers are obviously resourceful, finding gaps in security that will lead them to their gold...data. Encrypting data so that it is next to useless if they do access it seems like a worthwhile exercise, so why do the statistics above show that it is at such small numbers?

Well, until now, solutions for data encryption (mainframe and x86 environments) had many perceived drawbacks – degraded system performance (affecting customer satisfaction), too complex and costly to manage. So, what typically occurs is that organizations only encrypt the minimum amount of data to meet strict data security/privacy laws. This is called selective encryption. It is all a careful risk management process. Hope that your data is safe enough and that the bulk of it is unintelligible to hackers if it is accessed, all without spending too much.

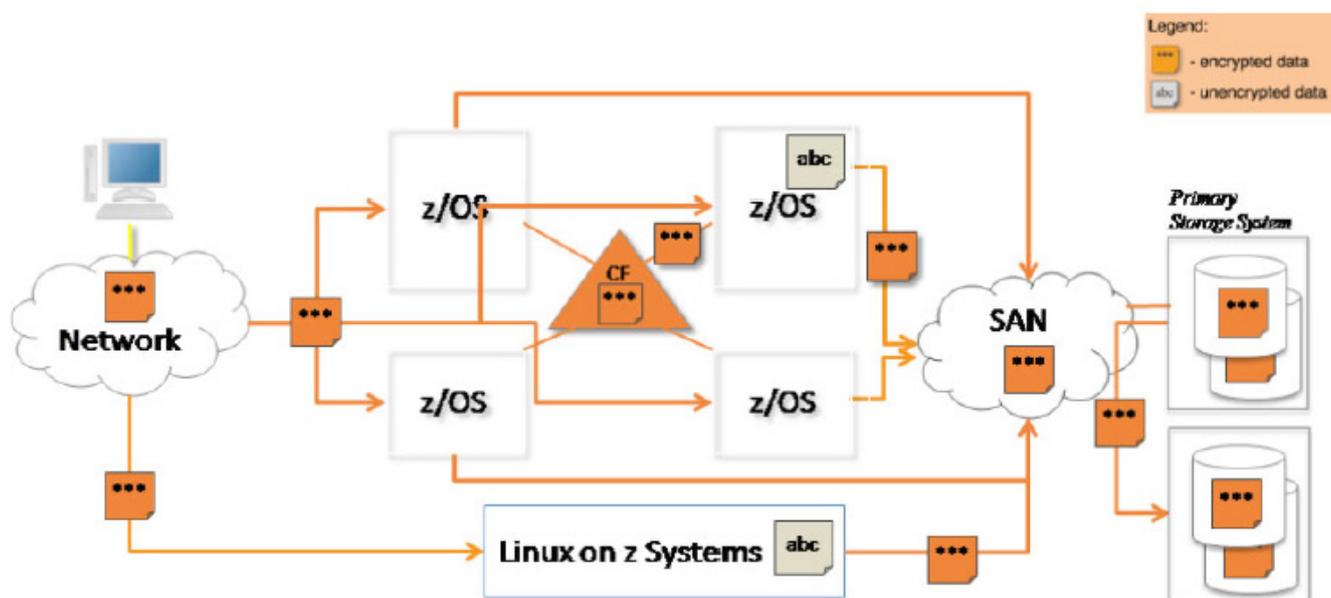
## What is pervasive encryption?

With IBM's z14 mainframe, pervasive encryption ensures that data is encrypted:

- > during processing (retrieved from files, and processed by applications)
- > in flight (transmitted internally and externally throughout the network)
- > at rest (when stored in databases, files, or data sets)
- > in store (on magnetic storage media).

Note that data is only decrypted when being processed by the operating system. This means that any vantage point used by a hacker, is met by data that needs to be decrypted before any sense can be made of it.

IBM's picture below provides a good representation of pervasive encryption in a data center environment using a z14 mainframe<sup>(4)</sup>.



In a world where technology advancements and continuous improvement practices are the norm, how do we get to that next level where everything is encrypted without impacting any other services? Before answering this, we need to look at what pervasive encryption is, and how it works.

Until now, this amount of encryption has been almost impossible because of the amount of additional processing that was required, but as a result of several advances in technology, IBM claims that this type of encryption can now be performed transparently, with no noticeable system degradation.

## Do you need it?

The first thing you are asking yourself is – this is all great, but what is the cost? As you have probably worked out, you will need to purchase or upgrade to the z14 mainframe. For existing IBM mainframe customers, this will be well worth considering, especially if it relieves headaches meeting increasing security and compliance mandates. For those currently without a mainframe, the cost of moving to this platform might just be too much. Let's take a look at some of the major factors when considering your move to a z14 and its pervasive encryption.

**Fear** – This is always a great motivating factor. A couple of the surveys discussed in this newsletter's Vendor Briefs section, discuss the top concerns for data centers and there is no surprise to see data privacy/security in the top 4 in all of them. A data breach is an organization's worst nightmare and can take many years to recover from such an event (the global average cost of a data breach is \$3.62 million<sup>(5)</sup>).

**Meeting data security regulations** – More and more stringent data security regulations are becoming commonplace, and the penalties surrounding non-compliance are becoming astronomical. For example, not meeting compliance for the EU General Data Protection Regulation (GDPR), which was introduced May 2018, could mean a fine of 4% of the organization's annual global turnover, or €20 million...whichever is greater!

No matter where you reside, there will be data security/privacy regulations or laws that affect you. For instance, in the US there is the Federal Information Processing Standards, and Federal Financial Institutions Examination Council (FFIEC), which includes the five banking regulators. Australia has its Privacy Act, and Hong Kong has its Personal

Data (Privacy) Ordinance. You may think that the GDPR mentioned earlier is just for EU organizations, but if you live outside that area and hold any data relating to customers from that region, you will still need to comply with GDPR requirements.

**Cost** – The costs of implementing pervasive encryption (through z14 purchase) obviously needs to be weighed up against the cost of not implementing it. Let's look at the pros and cons.

**Cons** – New mainframes don't come cheap! IBM do not release sales details, but depending on the configuration, it can cost between \$500,000 to \$3 million.

An alternative for smaller organizations might be the recently released IBM z14 Model ZR1, which although not as powerful as the full-blown z14, still supports pervasive encryption capabilities.

There is always going to be system overheads associated with all this type of encryption (in the z14s case, IBM indicate they are low, offset by the technical hardware and software advancements).

**Pros** – The z14 manages all your encryption needs once controls have been configured, meaning you will save on manually fine tuning what does and does not need to be encrypted.

You won't need to constantly evaluate your encryption requirements, and fine tune systems to minimize system degradation.

Come audit time, lots of time will be saved having to convince auditors that you have met your legal obligations.

Of course, you will obviously save money by not having to pay fines for any non-

compliance, and costs associated with lost business as a result of data breaches will also need to be factored in.

Note that the z14 pervasive encryption does not need you to make any application changes to cater for its implementation.

## Is pervasive encryption needed for new trends?

Trends such as the cloud, and blockchain have been around for a while, and survived without pervasive encryption, so you would have to say that it is not imperative that they go hand in hand. Having said that, if you were looking to deal with two companies offering this technology and there was little price difference you would certainly opt for the one with more security. IBM certainly agree as they have recently opened 6 new Cloud Blockchain data centers, scattered around the world.

## Is it good for everyone?

If cost was no barrier, you would think that implementing pervasive encryption is a no-brainer, but is this really the case?

Many law enforcement and intelligence agencies (think, US Homeland Security) are not fans of pervasive encryption. They believe that for them to do their job properly they should have access to all data, to enable them to confirm and track criminal activity, including terrorists. Companies like Apple (and now IBM) are firmly of the view that providing any access of this type undermines what encryption is all about.

It is not hard to side with either point.

## Conclusion

You can now see why IBM are using pervasive encryption as its primary selling point for its z14 mainframe. It is a case

where technology has finally caught up with encryption needs, to provide a secure home for our data.

I believe the z14 will have a faster uptake than previous mainframe models because of pressures from security mandates (especially GDPR), and business leaders recognizing the consequences from not implementing this type of data security. For many, it will be a case of “when”, rather than “if” pervasive encryption is implemented.

- (1) *Capgemini's Digital Transformation Institute Cybersecurity and Privacy Survey*
- (2) *IBM X-Force Threat Intelligence Index - 2018*
- (3) *Breach Level Index*
- (4) *IBM Redbook – IBM z14 Technical Introduction*
- (5) *2017 Ponemon Cost of Data Breach Study*



**Greg Hamlyn**, a senior technical course designer and developer with Interskill, has over 28 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](mailto:g.hamlyn@interskilllearning.com)

# Vendor Briefs

In this issue of Vendor Briefs we take a peek at what IBM has been up to in recent months, along with a couple of smaller mainframe vendors, and see what products they have been releasing lately. We also take a gaze at the crystal ball using the results of several surveys released throughout the year, to see where data centers, and the mainframe, are headed.

## IBM

Plenty of activity at IBM over the last 9 months, with the release of new hardware (z14 late 2017, and the recently released z14 ZR1 model), and z/OS 2.3 being made available at the end of September 2017. As I mentioned in my previous newsletter, IBM are working on some interesting projects, and are not afraid today to spread their wings and work with many others to create solutions for our fast moving, data centric lives.

## Blockchain

IBM were heavily into identifying markets for blockchain technology last time I wrote, and if anything, it is ramping up further. Some of their latest involvement with this technology:

**26 June 2017** - Seven Major European Banks Select IBM to Bring Blockchain-Based Trade Finance to Small and Medium Enterprises.

**19 July 2017** - The London Stock Exchange Group is collaborating with IBM to develop a securities data blockchain solution for European SMEs.

**8 August 2017** - Sony Corporation and Sony Global Education have developed a new blockchain-based student education records platform that can be used by educational

institutions to consolidate and manage students' educational data from several schools. This new platform was developed using IBM Blockchain.

**22 August 2017** - IBM will collaborate with Dole, Driscoll's, Golden State Foods, Kroger, McCormick and Company, McLane Company, Nestlé, Tyson Foods, Unilever and Walmart to address food safety worldwide by identifying new areas in the global supply chain can benefit from blockchain.

**4 October 2017** - Bank of Montreal, CaixaBank, Commerzbank and Erste Group have joined an initiative launched by UBS and IBM to build a new global trade platform based on blockchain technology. This new platform, called Batavia, will be built to be openly accessed by organisations of all sizes anywhere in the world, and built to support trade finance for transactions across all modes of trade, whether goods are being transported by air, land or sea.

**14 December 2017** - Walmart, IBM, and Tsinghua University National Engineering Laboratory for E-Commerce Technologies have announced that they will work together using blockchain to enhance food tracking, traceability and safety in China.

**16 Jan 2018** - Transport and Logistics company - A.P. Moller –Maersk are entering into a joint venture with IBM using blockchain to build a global shipping system that provides more transparency and simplicity in the movement of goods across borders and trading zones.

**26 April 2018** - A Consortium of Jewelry Industry Leaders have announced their use of TrustChain, which uses blockchain to trace the provenance of finished pieces of jewelry across the supply chain.

At the time of writing, IBM were also offering a [“Blockchain for dummies”](#) ebook, to help demystify this technology.

### **z14 and z14 ZR1**

IBM’s z14 was made available in September 2017, and as you would expect there were the usual improvements over its z13 predecessor:

- > Increased memory (32 terabytes, up from 10 terabytes on the z13)
- > Increased number of configurable cores (35% more total capacity compared to the z13)
- > Improved compression and encryption (IBM tests show 6 times faster encryption/decryption compared to the z13)

... and the heavy promotion of pervasive encryption (comprehensive data encryption). Have a read of this newsletter’s primary article for our take on this feature.

With a major release like this, we normally wouldn’t see another new mainframe until an entry-level version was created, which is normally at least a year apart. So, what is the story with the ZR1 that has just been recently released? Well, IBM needed to cater for customers that wanted to integrate the latest mainframe into their standard built data center landscape. The physical size of the z14 made this difficult, so, the solution was to create a narrower mainframe – the ZR1, which is housed in a single frame. This new mainframe is aimed at the small to medium businesses, so it is basically the next generation on from the z13s and as you would expect it provides a number of improvements over its predecessor: faster processors, more capacity, larger cache, and a smaller footprint. For businesses that don’t need the processing capabilities of the full-blown z14, the ZR1 provides a cost-effective solution to access an entry-level mainframe containing all the technology improvements

of the z14, including pervasive encryption.

### **Magnetic Tape**

What is old is new again! Just when you think that cartridge/tape is about to become obsolete, along comes a new standard in tape technology. IBM have been working with Sony for many years in the tape industry, and have recently presented a new record in tape storage capabilities. This new capability provides a recording density of more than 20 times that used in current commercial tape drives such as the IBM TS1155 enterprise tape drive (maximum native capacity uncompressed of 15 terabytes), enabling the potential to record up to approximately 330 terabytes (TB) of uncompressed data on a single tape cartridge.

Truly impressive, but unfortunately we will need to wait a while before this technology is converted to mainstream availability!

### **Db2**

We don’t often look at Db2 in these briefs as it tends to chug along nicely, with each new release providing greater capabilities. This is what we have come to expect from this product. Recently though, several areas have deemed it being mentioned here:

- > IBM have released a Db2 Developer Community Edition that developers can quickly download and install at no cost. This full-featured version allows developers to be able to more quickly build prototypes of applications destined for the Enterprise.
- > Db2 on Cloud is a fully managed service sitting on the IBM cloud. It provides quick and simple operations to be able to scale processing power, and memory. It also comes with a new easy to use web console.
- > Support for JSON is now available with Db2’s relational database system.

There are several other enhancements, so refer to IBM's release documentation if you want to learn more.

## Hitachi

One-time IBM competitor in the mainframe hardware business, Hitachi has now stopped making mainframes, and will now provide their customers with IBM z Systems hardware technology to complement their VOS3 operating system software. Who would have thought!

## Compuware

Compuware have been around for over 40 years, providing products that many mainframers have worked with over the years (File-Aid, Abend-AID, Strobe, and Topaz to name a few). In recent years they have expanded, integrating other businesses under their umbrella, and forming partnerships with others.

Compuware and **Wipro** have launched a **Mainframe Center of Excellence**, which in short is Compuware supplying Wipro with products and their experience so that Wipro can offer customers improvement in areas such as: agile development on the mainframe, and integration of mainframe development, testing, and release processes.

And who says that you don't get anything for free! Compuware has recently announced a new product named zAdvisor, which is free to existing customers. zAdvisor uses machine learning to continuously measure and improve an organization's mainframe DevOps processes and development outcomes. It leverages a set of analytic models that uncover correlations between mainframe developer behaviors and mainframe DevOps KPIs.

Compuware has also recently released a new

web interface for their ThruPut Manager product. As many will know, this product is responsible for automating and optimizing batch processing, so the introduction of this new interface provides staff at all levels with a visual insight into batch job initialization and execution and the affect it has on mainframe software licensing costs.

## Fujitsu

We just mentioned that Hitachi was ceasing its development of mainframes, but that doesn't mean that IBM has a monopoly just yet. In April, Fujitsu announced that it is updating its GS21 series of mainframe products, launching the super-scale Fujitsu Server GS21 3600 model group and the medium-to-large-scale Fujitsu Server GS21 3400 model group.

Fujitsu are also implementing a formal training academy, which is aimed at a new generation of mainframe engineers. It is to provide various types of training, from beginner to advanced level, focusing on Enterprise Platform Services (EPS). So, Fujitsu obviously believe that the mainframe has a solid future and are backing it up by providing the expertise to keep them running in the future.

## Surveys

A few interesting surveys have surfaced since our last newsletter.

### **BMC's 12th annual worldwide survey of mainframe users**

This survey was completed by over a 1000 technical professionals and executives, spread globally. Some of the key findings:

Top mainframe priorities (in order):

- > Cost reduction/optimization
- > Data privacy/compliance
- > Application availability

## > Application modernization

The survey also revealed the breakdown of age groups working in data centers and dispelled the myth that it was primarily manned by those eyeing retirement. In fact, 53% of respondents were under 50! What also showed up was that a key challenge (44%) for many organizations was managing their staffing/skills shortage. Related to this was the statistic that only 39% of respondents provide training for their employees. I guess the other 61% have to organize their own training?

Another interesting statistic was that just over 50% of companies are seeing increased data and transaction volumes (personally, I would have thought this would have been higher with current trends associated with big data, and analytics).

For interest sake, I had a look back at BMC's 8th annual mainframe survey, and many of the priorities, and future usage of the mainframe statistics had changed very little. It is also worth mentioning that the survey didn't specifically touch on many of today's data center trends relating to items such as; cloud, data analytics, DevOps, Blockchain, IoT, and cognitive computing.

A copy of the survey can be downloaded [here](#).

## **Uptime Institute's 2017 Data Center Survey**

This annual survey also involves around 1000 respondents, so with that number the results are likely to be more reliable than some other smaller surveys we reviewed. The areas canvassed for feedback concentrated on movement to the cloud, IT recoverability and resilience.

It was interesting to see where enterprise IT assets are currently deployed:

- > 65% reside in an Enterprise-owned data center
- > 22% reside in a colocation or multi-tenant data center
- > 13% reside in the cloud

Over the last few years these percentages have not changed much meaning that the initial thought of moving the majority of enterprise work to the cloud has not emerged as many thought. So, with the hype of the cloud it seems like there is still some way to go before fully utilizing this resource.

This survey also identified that around 70% of respondents had a multi-site failover strategy for live IT applications, which is what you want to hear from organizations you are doing business with. What I found interesting was that in a follow-up survey question that only 68% were confident that their IT functionality would work as it should in such an emergency. This uncertainty could be the result of not being able to test for all problem situations, which is not possible in real life.

Another interesting statistic related to the above is that 25% of survey respondents experienced a data center outage in the last 12 months. That seems quite high, but what constitutes an outage is not clear, so they may have been minor incidents with little impact on customers.

You can sign up for the webinar discussing this survey [here](#).

## **2018 Arcati Mainframe Yearbook**

The Arcati mainframe yearbook is eagerly awaited by mainframe professionals each year and contains an annual user survey. This year's survey was completed by 100 individuals, so somewhat smaller than other surveys mentioned above.

Looking at the results, this small sample base certainly seems to have skewed the results. For example, the survey identified that only 8% of respondents still viewed the mainframe as strategic. This is in stark contrast to BMC's survey mentioned above where 90% of respondents have indicated that the mainframe continues to be a strategic, long-term platform for growth. I guess it depends on who you are asking?

In relation to recent trends, the survey found that 5% of sites are using Big Data (with 19% planning to do so); a third of sites have implemented DevOps (with approximately another third planning to do so); and 13% of sites were already publishing and re-using APIs (with 56% planning to do so).

Looking at some of IBM's new offerings, the survey identified that 60% of respondents would not be looking at purchasing a LinuxOne (all Linux mainframe) system in the foreseeable future. Having said that, 4% do already have one of these systems, with 20% expecting to purchase one in the future. Another interesting statistic was that no-one in the survey indicated that they were using the mainframe for cloud computing.

Because of the small number of respondents, I wouldn't be using this survey as the basis for selling a new idea to management, but it doesn't hurt to pool all the surveys available to piece together the mainframe puzzle.

A full version of the survey can be downloaded using the link [here](#).



## 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 some of the new SDSF functions and features that were introduced with z/OS 2.2, and consists of 15 questions.

Click the link below to start.

[SDSF \(z/OS 2.2\)](#)

## Learning Spotlight – Blockchain Fundamentals

Most people today have heard of Bitcoin and are excited by the potential of its underlying technology: Blockchain. While the concept behind it is pretty simple, the actual workings get more complicated the more you delve into it.

In this newsletter we have provided access to the first of our three-module Blockchain Fundamentals course - [Introduction to Blockchain](#) - which discusses the evolution of Blockchain and its use with Bitcoin, its potential business value, and how Blockchain is structured, and implemented.

We hope you enjoy it.

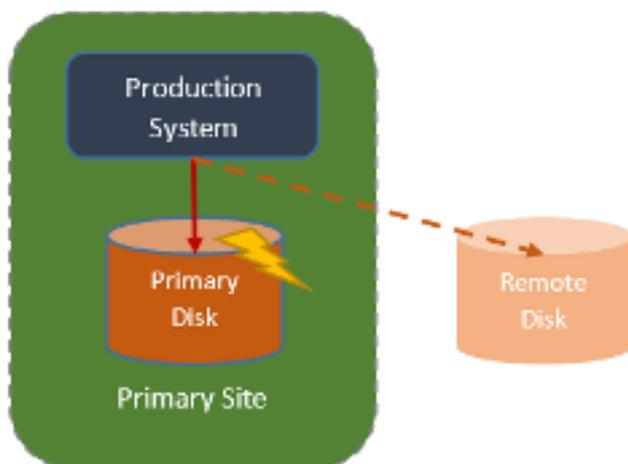
# Management: DASD Mirroring Part 3: Putting a Solution Together

By David Stephens, Lead Systems Programmer at Longpela Expertise

In the [previous article of this three-part series](#), we continued our look at DASD mirroring by looking at performance issues with PPRC and XRC mirroring. PPRC and XRC provide a second ‘emergency copy’ of your critical disk volumes. However simply having this copy isn’t enough. You need the procedures and abilities to switch to these copies when you need. So how can this be done?

## Basic Switching

Let’s take an example: your disk volumes crash, and you want to switch to your alternate DASD. This sounds simple: you issue an XRC or PPRC command to switch from the primary to the alternate volume, and you’re away.



But it’s not that simple. For a start, you’ll need to issue a command for every volume. If you’ve got a couple of hundred volumes that have failed, this will take some time.

What’s more, you’ll want to confirm the status of each volume before switching over, and perform any corrective action on a failure. Or in other words, it’ll take a few hours, and you’ll probably have to IPL your production systems.

Automation is the key here. Some automation products such as CA Ops/MVS High Availability option provide an automation interface to PPRC and XRC. So you can create scripts to automate some of the hard work.

## Basic Hyperswap

Perhaps a better solution is IBM’s Hyperswap - providing you’re using synchronous PPRC. A combined z/OS / hardware feature Hyperswap can automatically switch from a primary to an alternate volume based on criteria such as I/O errors; boxing a device offline, a control unit failure, or a user defined I/O timing threshold. And the good news is that this can all happen in a way that is transparent to applications: only a couple of seconds delay.

More importantly, Hyperswap will also automatically ‘freeze’ data replication to ensure that a valid copy of the data exists at the remote site before any cutover: the only thing worse than losing data is corrupting it.

Interestingly, Hyperswap uses either Tivoli Storage Productivity Center for Replication (TPC-R) or IBM’s GDPS as a user interface, so you’ll need one of these to work. TPC-R Basic Edition can be obtained for free, and provides single-site processing. The full product (not free) is needed when the DASD is in multiple sites.

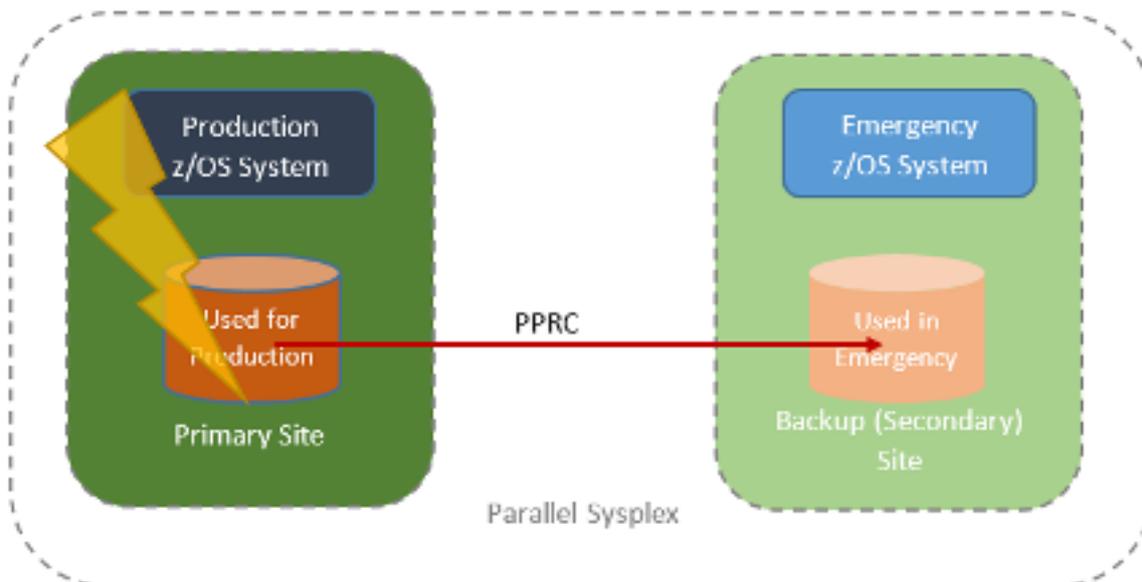
The GDPS offering for Basic Hyperswap is called GDPS/PPRC Hyperswap Manager

(GDPS/PPRC HM), and is the entry-level GDPS product.

Basic Hyperswap works great when there is a disk failure. However it can't help if there's a total system failure - it can't swap from one mainframe to another. For that, we need some extra smarts.

## GDPS/PPRC

Suppose things are more serious, and your entire primary site (systems and disk) are destroyed, and you need to switch to your backup site. Let's take PPRC for a start.



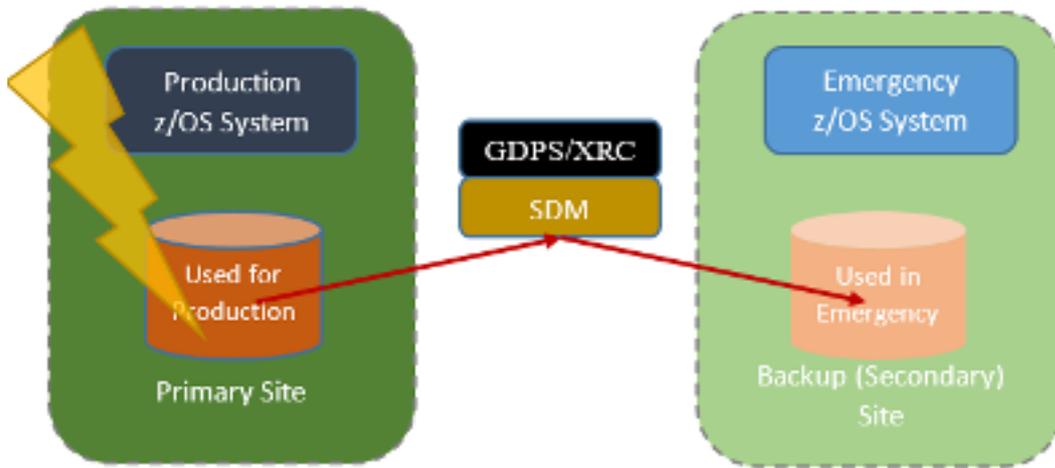
Hyperswap can't help you - it only switches disks, not systems. So we need to move up the GDPS totem pole to GDPS/PPRC.

For it to work, we need production and disaster recovery z/OS systems in a single parallel sysplex. GDPS/PPRC can do everything that GDPS/PPRC HM can do. However it can also move all processing from the production site to the remote site. It will automatically perform any IPLs or other processing necessary.

Like all GDPS products, GDPS/PPRC uses IBM System Automation for the automation part, though this can work alongside other automation products. And of course System Automation requires IBM Netview.

## GDPS/XRC

In simplest terms, GDPS/XRC does what GDPS/PPRC does, but for XRC. With XRC you going to get some data loss as XRC is an asynchronous solution - a Recovery Point Objective of a few seconds. However it will do all the automation to switch from a primary site to a remote site.

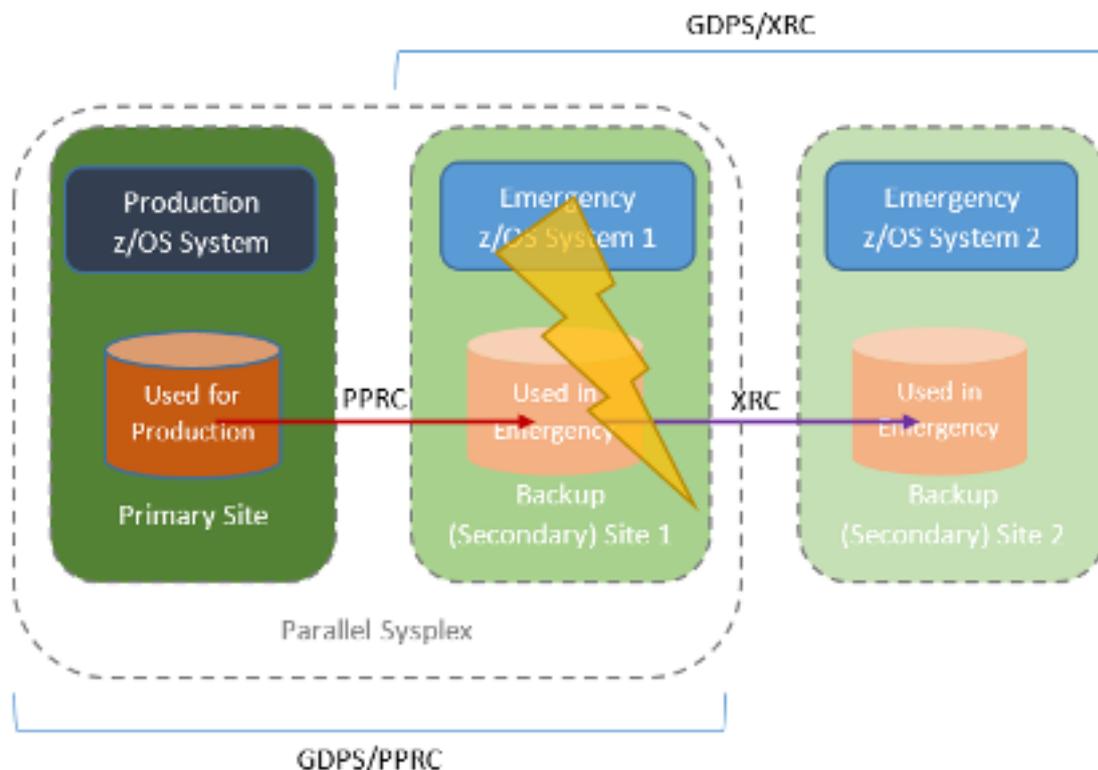


There are a couple of differences. For a start, GDPS/XRC cannot automatically switch over if there is a primary system failure: an operator must manually start the process. GDPS operates alongside the SDM - either in the production or remote system. No parallel sysplex required between the production and remote systems.

## GDPS/MzGM

GDPS/PPRC is excellent. Almost no data loss (a low recovery point objective, or RPO), and a quick cutover to the remote system if a disaster strikes (a low recovery time objective, or RTO). The catch is that GDPS/PPRC only works over short distances: a theoretical maximum of 200km, and a more realistic maximum of around 30km to minimize performance issues.

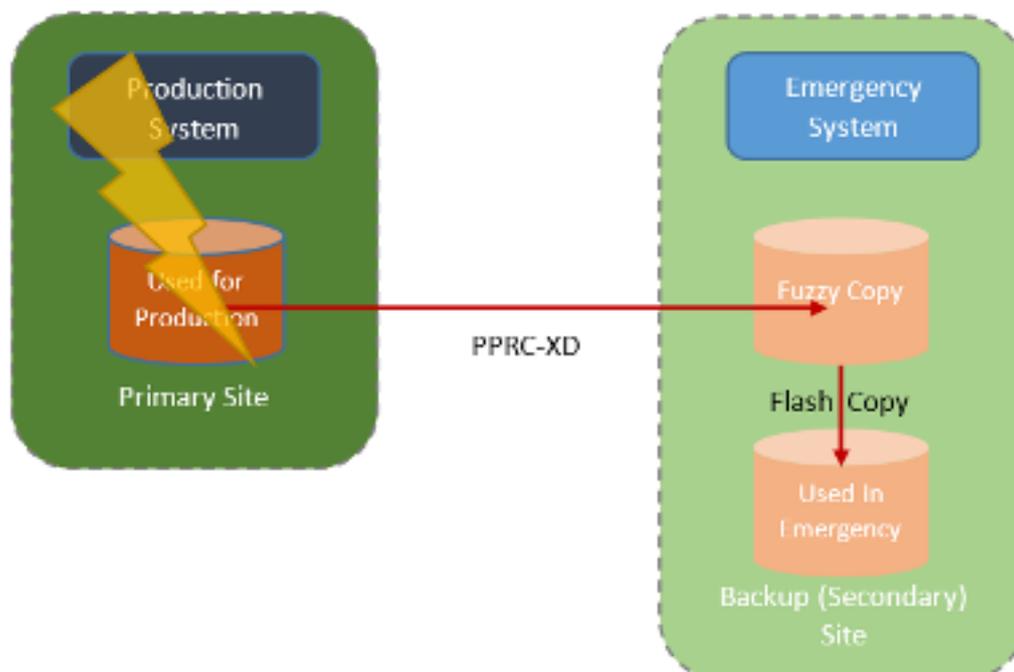
So if a disaster covers a 200km area (such as a hurricane), both sites are in danger. Those sites needing heavy-duty disaster recovery solve this problem by mixing PPRC and XRC. Here's how it works.



The production site uses PPRC to mirror DASD to a remote site less than 200km away. This remote site may be a full hot site with a second z/OS system, or just DASD. XRC is then used to mirror this intermediate DASD to a third hot site that is further away.

IBM GDPS supports the automation of this configuration by using both GDPS/PPRC and GDPS/XRC together into GDPS/MzGM (GDPS Metro z/OS Global Mirror).

## GDPS/Metro Mirror



Metro Mirror uses a combination of PPRC-XD (asynchronous PPRC) and Flashcopy in a similar way to our three-site scenario mixing PPRC and XRC. This has the advantage of unlimited distances between primary and remote site, no SDM needed, and can also support non-mainframe disk devices. GDPS also supports this scenario using GDPS Metro Mirror.

## Summary

In this article we looked at the basic GDPS offerings. There's further scope to create more robust environments with duplicated backups and remote data centres (3-site and 4-site options). If you look at the GDPS redbooks and articles, you'll only see them talk about IBM DASD. However the other storage vendors (HDS, EMC and HP) also support most of the GDPS and Hyperswap features.

In most cases, simply using DASD mirroring

technology isn't enough. Procedures and automation need to be in place to quickly recovery from a disaster when remote DASD is needed. IBM GDPS today is really the only solution on the market that provides out of the box automation of many possible DASD mirroring scenarios to achieve the required RPO and RTO needed.

Source:

LongEx Mainframe Quarterly -  
<http://www.longpelaexpertise.com.au/ezone/DASDMirroring3Solutions.php>



## Management: A Different Set of Goalposts for Application MQ Performance

*By David Stephens, Lead Systems Programmer at Longpela Expertise*

When you've been working with mainframes as long as I have, you'll have spent a lot of time working on application performance: tuning applications so they work as fast as necessary. However recently I've been working on an application using MQ. And I've found that tuning applications and MQ can be a little different to the tuning I've been used to. Let me explain.

### Traditional Application Tuning

The tuning I've been used to (which I'll call 'traditional application tuning') can be divided into two categories: online and batch. For batch, we're worried about the job's elapsed time, usually to meet a schedule or deadline. For online, we're worried about transaction response times,

a user hits a key (or clicks a mouse), and expects something back fast - say, less than a second.

And the path is well travelled. For batch, you look at the step elapsed time, for online the transaction response times. If they're not good enough, you can start your tuning.

### An MQ Application

For batch, MQ doesn't make much different. We're still worried about the step elapsed time. We'll be looking at MQ-related things like persistence, message size and syncpoint frequency to speed our MQ gets and puts.

In some cases, online is similar. You may have a CICS transaction that is triggered

from an MQ message, so you're looking at the CICS transaction response time. Nothing new. However, the application I've been working on is a bit different. It's a CICS-based MQ application. But rather than submitting a CICS transaction for each MQ message, one transaction processes a block of incoming MQ messages. Makes sense if you've a lot of incoming messages - no overhead of starting a new CICS transaction every time.

But now we have a situation where a single CICS transaction could process one, 100 or 1000 incoming MQ messages. We can no longer rely on the CICS transaction response times alone.

## MQ Goalposts

I know what you're thinking: the CICS SMF records (or more precisely, the CICS Performance Class data written by the CICS Monitoring Facility to SMF Type 110 records) include MQ statistics: number of MQ requests and response times. So, we're good - all we need to do is divide the CICS transaction response by the number of MQ requests, and we're in business.

True response time = CICS 110 Response Time / Number of MQ operations

There is a small problem with this approach. Our application gets messages until there are no more, and then does an MQ Getwait for up to one minute. If more messages arrive in this minute, the same transaction processes them. If not, the transaction ends, and the next incoming message triggers a new CICS transaction. This Getwait time is included in the CICS response times, making our 'true' response time misleading. We can get around this, as the 110 records include the Getwait time (I love CICS SMF records). So now we have:

True response time =  
(CICS 110 Response Time - MQ Getwait Time) / Number of MQ operations

This approach could work well if the MQ messages were coming from a person at a screen: they need a set response time, and we can figure it out. But for such 'batching' transactions it's different. Incoming messages are usually from other applications: from within the organisation, or from an external organisation. In many cases, MQ has become a common way for applications in different organisations to communicate.

So, here's the question. When looking at application tuning, the first question is always "where are my goalposts?" What performance do I need? For a CICS online transaction, this is usually a CICS transaction response time. So, a Service Level Agreement (SLA) would have something like 90% of transactions complete within 0.1 seconds. But what about incoming MQ messages from other applications? Often these are not real-time, but more like "near real-time."

In my example, each message needs to be processed within a few minutes (no specific time limit). But all messages have to be processed by a set time of day. There are many incoming MQ queues with varying processing deadlines.

## A Different Kind of Goalpost

So, we could have listed all incoming queues, added their deadlines, and worked from there.

Queue	Deadline
APC.MSGIN.Q1	2pm every weekday
APC.MSGIN.Q2	5pm every weekday
APC.MSGIN.Q3	One hour after end of day message received

This would have worked out fine but would have been time-consuming and complicated. So, we took a different approach. From experience, MQ performance problems are rarely from a single message on an otherwise empty queue not being processed fast enough. Rather, the problems occur when queues fill up: when messages enter the queue faster than they're processed. For example, if we have an application that can receive (and process) 10 messages a second, but the queue is receiving 20 messages a second, then we're not keeping up. Allow this to occur for a few minutes and we've got lots of messages on the queue.

When queues fill up, MQ performance suffers as the MQ queue manager does more work. For example, more messages may go to or from the pagesets, and not stay in buffer pools. This can be made worse if the receiving application is not simply getting the next message but searching through the queue for a specific message (though queue indexes on z/OS can help here).

More importantly, a message put on the queue will need to wait for all the messages before it to be processed before it will get its turn. For example, I've seen a case where the application team believed their application was performing well, not realising that at peak times, messages could wait up to 20 minutes in a queue before being processed. In the end, we decided that a suitable yardstick was this: messages

should be processed as fast as they're received. We don't need such performance on a second-by-second basis; averaged over an hour will do. So, if an application receives 60,000 messages in one hour, it should be able to process at least 60,000 messages in one hour to keep up with demand. This is sufficient for current workloads, but we also wanted to allow some room for growth or spikes in traffic, say 30%. So, the application now needs be able to process 80,000 (130% of 60,000) messages per hour.

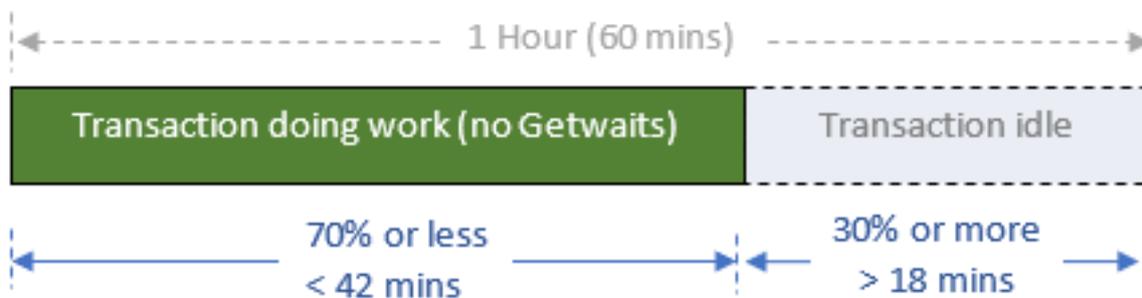
## Using These Goalposts

Once we have these goalposts, we can start using our statistics intelligently to determine if we're achieving our goals. One way is simply to look at the maximum queue depths for our MQ queue. If it starts getting high, we have a performance problem.

The problem is that this approach can only tell you if there's a problem now. Sure, you can also get historical information if you're keeping the right MQ statistics (we talk more about this in our partner article). But it won't tell you if there's a problem about to happen with a small increase in traffic. Or in other words, it doesn't help with analysis of the capacity of an application servicing an incoming MQ queue.

Another way is to look at the CICS transaction response times again. Take away times when the transaction is waiting for

MQ traffic (Getwaits), and we can figure out how much time in an hour the transaction is actually performing work. To be able to process 30% more than the peak incoming transaction rate, our processing transaction should be doing real work for a maximum of 70% of the hour (42 minutes - 70% times 60 minutes) and be idle for 30% of the time (18 minutes).



We can even plot this on a heat chart to see when the processing transaction is busy, and if it's getting close to our limit. Here's an example for one transaction.

Date:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
1-Aug	39	43	309	142	261	1257	796	1062	836	1887	2727	3355	3351	2215	2131	1492	2183	2456	589	355	531	163	409	107
2-Aug	361	32	219	199	256	316	776	764	560	1329	1101	2627	2729	1564	1581	1551	2463	1951	803	821	524	178	154	101
3-Aug	173	25	302	131	282	259	637	882	535	1252	1528	2444	2379	992	1805	1085	2217	1899	748	884	274	38	452	69
4-Aug	368	43	168	180	176	252	710	729	487	1049	1491	2430	2390	1017	1416	901	2072	1099	499	550	206	82	139	74
5-Aug	126	41	210	183	285	280	1480	1022	566	1310	1526	2570	2584	1207	1144	1100	2026	1234	664	16	262	29	40	14
8-Aug	35	87	157	119	258	199	1241	852	469	886	1343	2137	1729	1116	1544	1298	1880	1366	727	120	200	171	334	55
9-Aug	171	35	141	150	225	211	779	690	499	906	1512	2105	2609	1160	1443	805	1887	1180	712	211	139	394	288	69
10-Aug	390	50	142	184	351	245	594	798	482	1065	1342	2620	2354	951	1331	880	2291	2059	363	331	140	338	199	72
11-Aug	194	58	82	183	224	274	700	615	501	692	1155	2800	2778	1028	1580	948	1970	1765	354	263	204	319	259	79
12-Aug	210	55	140	136	214	229	639	827	532	1019	1512	2277	2323	1031	1644	715	2045	2013	341	3	0	139	18	70
15-Aug	37	34	376	51	257	1034	684	902	542	1298	1994	3132	3021	1429	1379	819	2366	2344	596	190	289	145	41	59
16-Aug	304	97	125	167	204	245	1335	723	464	1460	1362	2202	2354	1508	1471	1118	2679	1966	301	314	221	206	251	111
17-Aug	329	34	166	150	210	263	709	675	463	1215	1380	2381	2354	853	2245	884	1815	1159	315	325	324	27	236	100
18-Aug	324	43	130	128	290	331	730	871	505	1114	1129	2217	2445	1187	1178	852	2232	1363	196	475	208	165	65	68
19-Aug	309	24	159	114	198	226	927	734	380	984	1736	2275	2428	1176	925	1152	2142	1723	179	53	284	174	26	18
22-Aug	46	56	110	89	163	219	1043	699	513	917	1305	1770	2126	934	1106	921	2001	850	214	82	279	333	243	51
23-Aug	284	31	125	137	226	178	599	635	407	989	1335	2374	1554	788	1136	998	2256	1603	364	407	441	792	210	71
24-Aug	266	39	131	144	170	239	693	601	312	847	1172	2213	1888	1244	801	1444	2095	1826	208	535	239	41	605	200
25-Aug	444	114	799	372	321	364	1068	696	853	1654	1402	2082	2229	1010	1342	1083	2191	1404	485	341	133	82	564	135
26-Aug	414	82	196	122	222	232	944	671	438	942	1579	2026	1860	754	1473	589	2045	1233	166	3	333	122	32	28
29-Aug	62	55	106	87	138	205	590	448	315	680	1014	1554	1534	943	938	747	2161	795	107	300	47	167	71	378
30-Aug	76	387	114	147	162	241	543	571	393	1009	1108	2356	1561	1055	971	1301	1869	925	198	554	196	57	282	71
31-Aug	277	26	167	124	254	355	1250	645	556	1022	1423	2653	1924	1296	1390	1375	2248	1922	475	653	241	109	176	637

All the numbers are seconds of actual processing (excluding Getwaits). So, at midnight, our application is only processing for 30-400 seconds in any one hour: no performance issues. However, on 01-Aug between 11:00 and 12:59, the processing is getting very close to our one-hour (3600 seconds) limit - some tuning required.

From here, we followed the normal tuning process for a CICS transaction: looking at the reasons for the response times from the CICS SMF 110 records, and working from there.

## Conclusion

There's a lot of information about monitoring queue manager performance and MQ networks. There are also a lot of tools to look at queue depths, status and more. However, I haven't seen anything that talks about application performance, and in particular finding

out if the performance of a service task processing incoming MQ messages is sufficient. To my mind, such MQ processing programs need a new set of goalposts to determine if the current performance is good enough. If it isn't, then the normal performance tuning tasks can be used to get the performance to where it needs to be.

Source:

LongEx Mainframe Quarterly - <http://www.longpelaexpertise.com.au/ezone/MQGoalposts.php>