interskill learning

INTERSKILL MAINFRAME TRAINING NEWSLETTER

Sept 2014

Inside This Issue:

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 an edition several times per year. These editions will include z/OS related articles for technical experts and management, and information on upcoming releases of Interskill's Mainframe curriculum.

Interskill Learning Planned New Curriculum and Updates for 2014

- **z/OS 2.1 Series**
  - z/OS Concepts and Components 2.1
  - Initializing and Terminating the z/OS System 2.1
  - Monitoring the z/OS System 2.1
  - z/OS Architecture 2.1
  - Identifying z/OS System Problems 2.1
  - Resolving z/OS System Problems 2.1

- **JCL 2.1 Series**
  - JCL (z/OS) – Introduction to z/OS JCL 2.1
  - JCL (z/OS) – Using Special Data Sets in Batch Jobs 2.1
  - JCL (z/OS) – Identifying and Resolving Batch Problems in JCL 2.1
  - JCL (z/OS) – Coding Procedures and JES2 Control Statements 2.1

- **IBM Environment Introduction Series 2.1**
  - Introduction to the IBM Enterprise Environment 2.1
  - z/OS System Programming Fundamentals 2.1
  - IBM Development Environment Overview 2.1

- **The z/OS Management Facility 2.1**
- **Introduction to zEnterprise - EC12, BC12 & zBX**
- **IBM Explorer for z/OS**

- **TSO/ISPF 2.1 Series**
  - Using Online System Facilities - TSO/ISPF 2.1
  - Managing Data Files and Definitions with ISPF/PDF 2.1
  - Maintaining Data in Files with the ISPF Editor 2.1

- **System Display and Search Facility (SDSF) 2.1 Series**
  - SDSF Concepts and Operation 2.1
  - Using SDSF to Control Job Processing 2.1
  - Using SDSF to Display, Manipulate and Print Job Output 2.1
  - Using SDSF to Manage System Resources and Devices 2.1

- **z/OS Console Simulators 2.1 Series**
  - Console Simulations 2.1

- **z/OS Batch Utilities 2.1 Series**
  - General Data Set Utilities 2.1
  - Data Utilities 2.1
  - Introduction to VSAM 2.1

- **JES2 2.1 Series**
  - JES2 System Initialization & Shutdown 2.1
  - Monitoring Batch Jobs with JES2 2.1
  - Using JES2 in Scheduling Batch Jobs 2.1
  - Identify and Resolve JES2 Batch Problems 2.1
  - Identify and Resolve JES2 System Problems 2.1

- **Operations Assessment Series 2.1**
  - CICS Concepts and Operation Assessment 2.1
  - JES2 Concepts and Operation Assessment 2.1
  - JES2 Problem Resolution Assessment 2.1
  - SNA/VTAM Concepts and Operation Assessment 2.1
  - z/OS Concepts Assessment 2.1
  - z/OS Problem Diagnosis and Resolution Assessment 2.1
  - z/OS System Operation Assessment 2.1
  - JCL Concepts Assessment 2.1
  - JCL Coding Assessment 2.1
  - JCL Problem Resolution Assessment 2.1
  - TSO/ISPF Concepts Assessment 2.1
  - TSO/ISPF Operation Assessment 2.1

- **Cloud Computing for Data Centers Series 2.1**
  - Cloud Computing for Data Centers 2.1

- **IBM Mainframe Communications Series 2.1**
  - IBM Mainframe Communications Concepts 2.1
  - VTAM Commands 2.1
  - Mainframe TCP/IP Commands 2.1

- **Data Center Storage Management Series**
  - Storage Fundamentals for Data Centers 2.1
  - Using DFSMS to Manage the z/OS Storage Environment 2.1
  - Storage Networks, Administration, and DASD Management Using ICKDSF 2.1

- **IBM i Programming Fundamentals Series**
  - CL Programming
  - CL Programming Functions and Messaging
  - RPG/400 Introduction
  - RPG/400 Coding
  - RPG/400 Programming
  - RPG/400 Workstation Programming

- **CA OPS/MVS® Event Management and Automation Series**
  - CA OPS/MVS® Event Management and Automation - Overview, Components, and Features
  - CA OPS/MVS® Event Management and Automation - Rules and OPS/REXX
  - CA OPS/MVS® Event Management and Automation - Automating Events Using the Relational Data Framework
  - CA OPS/MVS® Event Management and Automation - Automating Events Using the SSM
  - CA OPS/MVS® Event Management and Automation - Schedule and Group Managers for Event Management

- **CA 1® Tape Management Series**
  - CA 1® Tape Management - Using Tape Media
  - CA 1® Tape Management - Identifying and Resolving Media Problems

- **CA ACF2™ Series**
  - CA ACF2™ - Introduction
  - CA ACF2™ - Protecting Data Integrity
  - CA ACF2™ - Protecting System Access
  - CA ACF2™ - Defining Environment Controls
  - CA ACF2™ - Protecting General Resources
  - CA ACF2™ - Maintaining ACF2™
  - CA ACF2™ - For Auditors

Highlights of Interskill Learning Releases from 2013

Interskill Learning releases for 2013 included the following:

- **Business Continuity and Disaster Recovery Series**
  - Ensuring Data Centre Business Continuity

- **Installing and Managing z/OS Software Series**
  - Introduction to SMP/E

- **Data Center Storage Management Curriculum**
  - Storage Fundamentals for Data Centers
  - Managing the Storage Network
  - DFSMSrmm Operations

- **CA 7® Workload Automation Series**
  - Introduction to CA Workload Automation - CA 7® Edition
  - Scheduling Batch Processing
  - Monitoring and Managing the Batch Processing Environment
  - System Programmer Interaction with CA 7
  - Backup, Recovery, and Problem Resolution

- **CA 11 Workload Automation Series**
  - CA Workload Automation Restart Option for z/OS Schedulers Overview
  - Managing CA Workload Automation Restart Option for z/OS Schedulers

- **New courses in the Power series for IBM i System Administrators**
  - IBM i System Administration Fundamentals
  - Journal Management
  - Storage Management
  - Logical Partitioning and Virtualization
  - Security Implementation

Is The Cloud About To Overtake Your Mainframe?

By Greg Hamlyn

In this article I take a bit of history, a dose of clarity, and a gaze into my crystal ball to see whether the cloud and mainframe can drive happily side-by-side, or whether they are on a collision course where only one will survive. We will also look at the general direction that both the cloud and mainframe are taking and how this will affect organisations, their IT departments and you, the knowledge worker.

What the Customer Wants

Before predicting the future of the cloud and the mainframe, it is important that we take a backward step and look at today's organisational computing wants. I emphasize that these are wants as opposed to needs. Computing capabilities are changing so quickly that today's customers expect that any product or service that they can create to benefit them should be available to them now, or in the immediate future.

So, what do organisations want from their computing services?

- **Cost effectiveness** - Organisations are always looking at how to save money. However, I didn't want to mention "cheap" here, because like most things in life "you get what you pay for".
- **Agility** - Developing new products and bringing them to market more quickly in response to changes in business environments.
- **Speed** - Organisation's can never have enough of both network speed and actual computer processing speed. Customer's today expect information "on demand".
- **Security** - This can vary considerably depending on the industry, compliance requirements and the organisation's data rules and regulations.
- **Availability** - 24 x 7 is demanded by most customers and the five 9's system availability is something that is expected rather than a wish.
- **Scalability** - If business is booming or waning, organisations want to be able to quickly provision or de-provision resources, while ensuring no disruption to customers.
- **Sufficient storage** - What constitutes sufficient is difficult for any organisation, as they are finding that more and more information is required to be stored, either for compliance reasons or for internal analytical purposes.
- **Automation** - Speeding up previously manual processes through automation not only saves time, but can either free up staff for other tasks or allow the organisation to cut staff numbers and save money.
- **Less complexity** - One of the major goals for most CIO's is to reduce the complexity of their IT systems.

- **Recoverability** - Organisations want to recover IT services and data as quickly as possible following a minor or major problem.

Have a look at your organisation's business goals and you will see how the majority of the items above, either directly or indirectly, map to them. Now that the customer's needs are identified, let's look how the cloud and the mainframe are fulfilling them.

What the Cloud is Offering

Before discussing the cloud, we should identify exactly what type of hardware and software comprises a typical cloud provider's data center, so as to compare later to what the mainframe offers. Trying to find out exactly what is housed in these cloud data centers is certainly a difficult task. AWS (by far the largest cloud player) for example, tends not to buy brand name hardware, opting instead to design its own, resulting in streamlined products concentrating on their core needs. This also goes for their storage and networking equipment. AWS's leadership statement sums it up well:
"We try not to spend money on things that don't matter to customers".
The overall size of cloud data centers can be huge, with Microsoft highlighting the fact that one of its Azure data centers takes up as much space as ten football fields.

So, now that we have established that the cloud is big, let's focus on where the cloud has come from. The cloud's initial foray into the enterprise dealt mainly with cost savings associated with not having to own your own hardware or pay for software upgrades. This was the public cloud. To be truthful, there would have been many IT professionals that rolled their eyes thinking that this was just going to be another fad, or that it was just another form of out-sourcing. But then the cloud evolved.

After the initial excitement came the realisation that, after hidden costs, there were not too many real savings to be made in this area (unless you were a startup). Attention then turned to other cloud services providing greater long-term benefits such as:

- Innovation (with more capabilities, comes more ideas, and more opportunities)
- Reduced complexity (by having someone else running your computing processing)
- Collaboration and accessibility (people working together from different locations globally from platform-agnostic devices)
- Data and system recovery (for organisations lagging behind with their disaster recovery planning, the cloud has provided them with an easier path to implementation)

From the items above, you can see how the cloud is beginning to respond to the customer's wants, discussed earlier.

What IBM and the Mainframe are Doing About It

Long considered the cornerstone of enterprise computing, IBM's initial reaction to the cloud was to see if it was just going to be a fad or whether customers would take it seriously. While IBM's mainframe boasts impressive performance and security, the last few years in particular has produced some interesting directions.

In particular, mainframe costs for cloud-like functionality have reduced substantially. For example, in 2003 the IBM low-end z990 was priced at over $1 million. Fast forward to today and IBM's latest Enterprise Cloud System, can be implemented on the zEC12 or zBC12. The latter is available for purchase at around $75000. The Enterprise Cloud System integrates software, storage, and server technologies into one simple, flexible, and secure factory integrated solution.

You may have also noticed that many of IBM’s recent enterprise software releases have been aimed at the cloud. They have also acquired SoftLayer in the last year or so, to bolster their cloud offerings.

What the Statistics Tell Us

Let’s have a look at some recent cloud and mainframe statistics to see if we can glean something from them:

Cloud

- By 2015, end-user spending on cloud services could be more than $180 billion
- More than 60% of businesses utilise cloud for performing IT-related operations
- From 2012 to 2017, data center workloads will grow 2.3-fold; cloud workloads will grow 3.7-fold
- 14% of companies downsized their IT after cloud adoption

Mainframe

- Currently more than 70 percent of enterprise data and 71 percent of all Fortune 500 companies’ core businesses resides on a mainframe
- IBM sells mainframes to cloud service providers, who focus on delivering large scale Linux environments for their customers
- IBM is adding up to 60 new customers per year to its mainframe client list. These are not refresh sales or new sales to existing customers but companies that currently have no mainframe capability at all.
- The mainframe, and IBM, is still around after 50 years. The same cannot be said of other technologies or IT companies

Shake it up and what do you get

The cloud has seduced the computing world, giving it a taste of what is possible and it will obviously continue to grow based on the statistics cited above. The cloud has agility, scalability, is heavily automated, cost effective in some situations, and provides almost unlimited storage. Areas that will see improvements in the future include speed (network), and security.

The IBM mainframe on the other hand has many loyal users that are not likely, in the short term, to abandon their significant investment and totally move to the cloud. The mainframe has security, speed, automation, availability, scalability (to an extent) and is addressing its cost effectiveness. However the traditional mainframe will find it difficult to compete with cloud storage and its seemingly unlimited space.

The Wrap Up

It is clear that the cloud and mainframe are two entities coming from different directions and converging on the same goal; addressing the customer’s wants. Each has its own strengths in relation to these wants, and it is clear that both are tackling the issues to improve their weaknesses.

The mainframe isn’t going away soon, with many larger customers continuing to use it for their critical processing and compliance requirements, although other less critical processing and storage needs are beginning to be shunted to the cloud. From this scenario, what has emerged is the hybrid cloud environment, where organisations currently feel that the correct risk is being placed.

What Does this Mean for You?

Organisations will continue to analyse the viability of both the cloud and mainframe platforms as it relates to their organisation and industry, and implement what best suits their needs.

Some pundits say that **IT departments** will grow, even if computing is moved to the cloud. The reason for this is that if a hybrid cloud is created, staff from the organisation will need to continuously work with not only cloud suppliers, but their own environments as well. The statistics mentioned earlier regarding only a 14% IT staff reduction following cloud implementation, seems to support that most IT areas will not shrink.

For **IT staff**, building up your cloud knowledge is certainly going to keep you in a job long-term.

Greg Hamlyn, a senior technical course designer and developer with Interskill, has over 25 years experience working in the mainframe and z/OS data center environments, having performed in numerous operational roles ranging from computer operator to national mainframe technical training manager for a large government agency. He has presented at AFCOM and has written a number of articles relating to training methods, 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 look at not only the products that have been recently released by the major mainframe players since our last update, but also take a look at what's has been happening around mobile mainframe-related apps.

IBM

CICS Transaction Server for z/OS 5.2
This latest release of CICS TS promises functional enhancements split between areas relating to flexibility/agility, resource usage, security, performance enhancements and cloud enablement. Some of the most interesting takeaways from these enhancements are:

- CICS TS 5.2 has been enhanced to support V8.5.5 of the WebSphere Application Server Liberty Profile. This upgrade provides improved Java database connectivity, transactionality, and security features.
- The use of critical system resources by CICS is better managed through Service Level Agreement (SLA) policies and a broader range of new configurable policy thresholds.
- CICS TS 5.2 enforces the use of the TLS 1.2 cryptographic protocol and checks for conformance to the NIST-SP800-131a security standard.
- The functionality and usability of CICS platform bundles have been enhanced allowing additional resources to be more easily defined, removed or reconfigured.

IBM MQ for z/OS V8.0
Released in June this year, MQ for z/OS V8.0 continues to build on MQ's offerings for platforms and environments that are quickly evolving. With this release, the

number of buffer pools has been extended from 16 to 100 providing high volume queue managers with relief from virtual storage constraints. Other benefits include: reduced performance evaluation time, and a reduction of pageset I/O for some processing types.

This release also:

- Provides the capability to increase the size of the MQ log to 16EB
- Sees additional SMF data captured in the form of channel initiator statistics, and channel initiator accounting data (this provides administrators with better tuning and capacity planning data)

The IBM Redbook "IBM MQ V8 Features and Enhancements" contains some very good scenarios on how to exploit MQ for z/OS V8.0 with zEDC and SCM capabilities added in the zEC12 and BC12 hardware releases.

IBM WebSphere Liberty z/OS Connect

Providing web, cloud and mobile users with access to mainframe transactions and data has often been difficult due to the different protocols and authorization issues. IBM has resolved some of these issues with its recent release of z/OS connect, which provides the non-mainframe platforms mentioned above with secure access to CICS and IMS using REST URIs and JSON data. It basically shields these backend systems from having to understand those protocols and formats.

Best of all, it is available at no additional cost as it is packaged with WebSphere Application Server, CICS and IMS software products.

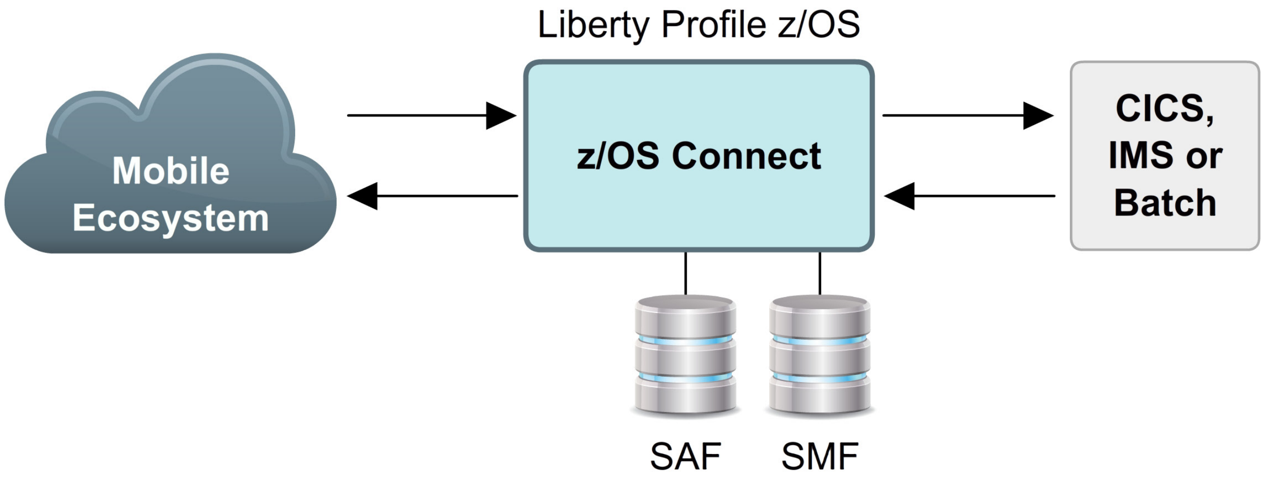

Storage

At EDGE 2014 earlier this year, IBM announced new and enhanced capabilities of it high-end storage products: Storwize V7000 Unified, XIV Cloud Storage for Service Providers, TS4500 tape library and DS8870 Flash enclosure. These high performance solutions are primarily aimed at optimizing storage for large scale cloud deployments and can provide improved:

- Data access times
- Real-time compression
- Virtualization
- Easy tiering
- Mirroring

BMC

BMC Intelligent Capping

Which data center isn't looking to save money? In our last newsletter we highlighted BMC's Cost Analyzer for zEnterprise, which recommended cost reduction strategies in relation to the organization's mainframe licensing charge (MLC). Recently, BMC announced their "intelligent capping" software, which analyzes workloads and automates workload capping, resulting in lower monthly license charges. In fact, BMC mentions that on average, cost savings of five percent or more can be achieved.

Cloud Lifecycle Management 4.0

Aimed at the hybrid cloud market, the release of BMC's Cloud Lifecycle Management 4.0 (CLM4) enables organizations using both private and public cloud to apply consistent data and network compliance, and change management policies across these complex environments.

CLM 4 also includes new support for Microsoft Azure and an expanded integration with Amazon Web Services public clouds. It provides an auto scaling capability that prevents over provisioning of cloud resources, through pre-defined service performance requirements.

GT Software

Data Unification Solution

Today's environment demands that an organization's data, no matter what platform it sits on, is available to the people who need it. This complex need is made that much easier with GT Software's Data Unification Solution, which allows users to deploy, merge and better access data within or outside the mainframe and across disparate platforms, programming languages or formats. The solution incorporates three components:

- **Ivory Data Hubv5.3** (SQL support for all data types, including VSAM, DB2, Oracle and SQL Server)
- **Ivory Service Architect** (extends access to mainframe applications, data and processes through web services)
- **Portfolio Assessment Services** (provides assessment of mainframe environments).

Mobile Apps for the Mainframe

With the proliferation of mobile apps available, the mainframe involvement with this genre has been surprisingly quiet, until recently. We will now look at how IBM is looking to change this, and at some of the mainframe-related apps that are currently available.

At the time of writing, IBM is running an app building competition and looking for the best ideas for connecting your mobile device to the mainframe world. It is open until September 17, 2014 with prizes such as an iPad, entrance to IBM Enterprise 2014, or a week with the IBM zMobile experts on offer.

As a related side-note to this, IBM and Apple have recently entered into an exclusive agreement to create, among other things, more than 100 industry-specific enterprise solutions including native apps, developed exclusively from the ground up, for iPhone and iPad. This could be big.

Existing Apps

There are a number of existing iOS and android apps available that allow you to access information to keep you up to date with the mainframe world. Here is a sprinkling of them:

- IBM's Systems Magazine
- IBM Redbooks
- 3270 emulators (Mocha, TN3270, Glink 3270, QWS3270 tablet, and TinyTERM to name a few)
- Mainframe questions for self testing & interview purposes (Mainframe IBM Interview QA, Test Your Mainframe Skills!!!)
- Error descriptions (COBOL Abends, Mainframe Translator)
- IBM Technical Event Portal (if you are lucky enough to get to one of the IBM Technical University events, this app can plan your attendance for sessions and allow downloads of presentations)

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 new features and functions associated with the introduction of z/OS 2.1, and consists of 13 questions. Click the link below to start.

z/OS 2.1 Features

Learning Spotlight – ICKDSF

ICKDSF remains a critical utility to initialize and manage mainframe disk units for z/OS, z/VM and z/VSE. Storage Administrators and Systems Programmers need to know how this utility is invoked in these different environments, and its capabilities.

The module that we have provided for you here is Introducing ICKDSF which describes how ICKDSF is run, and how it is used to initialise disks and perform disk maintenence.

We hope you enjoy it.

JCL Word Search

Unlike traditional word searches where the words are supplied and you attempt to find them, this is a version with a twist; I don't supply any words!

Hidden in the maze below are the names of 4 JCL statements and 20 JCL parameters. I've provided some blank boxes at the bottom of the page to record your findings.

R	T	I	X	O	E	Z	I	S	K	L	B	S	C	S
J	E	N	O	T	I	F	Y	U	C	C	M	E	E	E
P	I	F	E	T	S	E	D	E	O	S	L	I	X	T
B	S	Q	D	M	U	E	R	N	G	A	E	P	E	S
L	O	I	T	D	M	L	D	C	S	A	B	O	I	P
D	C	B	D	L	P	O	L	F	Y	F	A	C	V	A
N	K	U	J	D	R	A	C	E	S	K	L	T	E	C
T	T	D	L	O	S	T	M	X	O	W	L	J	T	E
R	X	O	N	S	B	A	S	C	U	Q	U	Y	Y	T
T	H	A	V	R	N	N	E	I	T	J	P	O	I	K
P	D	W	P	S	E	B	K	Z	Q	R	B	J	T	X
X	I	P	D	Q	C	G	F	A	U	Y	M	M	U	D
Z	B	H	X	D	W	W	I	N	T	I	C	R	Z	U
L	U	M	R	E	C	F	M	O	B	N	Z	R	K	A
D	N	E	Y	B	D	K	D	U	N	X	G	J	E	Z

You can contact me at g.hamlyn@interskilllearning.com if you need any help. Results will be in the next newsletter. Good luck.

management: Why are you Tuning?

By David Stephens, Lead Systems Programmer at Longpela Expertise

Many sites or technical groups tune simply because they think they need to. A better performing system/application/database is better for everyone, right? However how does this help your business? Does it reduce costs, meet compliance, make customers happier or give a competitive edge? Too often the answer is 'no', or 'maybe'. Too often tuning dollars are wasted.

What Is Tuning, Really?

It sounds obvious doesn't it? Tuning is making things run faster. Except that it's not. Tuning is the entire process of making things run better. There are two things here to think about.

The Entire Process

Tuning isn't just changing a program, or system parameter, or database structure. It's the whole process:

1. Measure - Measure current performance
2. Evaluate - Determine if this is sufficient
3. Change - Find and fix problem areas if it isn't
4. Re-measure – Measure current performance again

And it doesn't stop – ever. Step 4 above is the start of the whole cycle again. Tuning is a constant circle of measurement, evaluation, and sometimes (and only sometimes) making changes for improvement. And it's long term.

Tuning doesn't stop when things are working fine. The best tuning is where performance is regularly measured and monitored, and changes rarely made. In an ideal environment, changes are performed before problems occur. So if an application's response time is measured over a year, trends can be seen. If the response time is seen to be slowly increasing as transaction loads increase, this can be addressed before those times become unacceptable; tuning vs fire-fighting.

Tuning is the entire long-term, on-going process of making things run better.

Run 'Better'

Tuning isn't just about making things run faster. The most common projects I work on are about reducing the CPU used. Other reasons could be to reduce disk space, tape, or network bandwidth. You may use an API that costs money every time it is called, and want to minimise this cost. Or maybe you may want an application to affect other applications less. For example, use less memory so another application can get more. There are hundreds of possible reasons to tune.

Tuning is the entire process of making things run better, not necessarily faster.

This doesn't mean that tuning is the only work done to improve performance. Your applications staff will be developing new applications with performance in mind. Hopefully, your development process concludes with a complete tuning cycle (measure, evaluate, change and re-measure) in a test environment under load.

Similarly, your DBAs will be designing databases for performance, and your systems programmers will be thinking about performance when installing new software or configuring systems. Everyone wants to improve performance. However often these tasks are done in isolation, without set targets. Tuning is the process of making existing things run better.

Getting Value From Tuning

I remember one site with a team tasked with reducing CPU usage across the board. Targets were set, and bonuses paid when they were met. The problem was that most of the CPU savings gained didn't benefit the organisation. They didn't make customers happier, give any advantage, or satisfy any compliance requirement. In fact the team's purpose was to reduce CPU usage so an upgrade to a larger processor would not be needed; a larger processor costing a lot of money to buy and run. And this was the problem. They needed to reduce the peak CPU usage: only a few hours every month during month-end processing. However the application team were reducing CPU usage at any time they could find.

So before even starting a measurement phase, the reason for tuning must be settled: the mission. This is Project Management 101. So let's work through an example.

Imagine we want to tune to reduce our running costs. Great, but what are those costs? In my field these are usually software licensing, staff, hardware and environmentals. So which one do we go after first? Software licensing is usually the largest by far, so let's start there.

But don't start reaching for your measuring gear. There are many ways to reduce software licensing costs. A review of the software portfolio is an excellent place to start – check that all software is still needed, used, and not duplicated. In my area of mainframes, this is particularly relevant as software has often been installed for decades. Let's also check your software licenses – there may be ways to improve your licensing or configuration to reduce costs.

Now let's take another look at that software licensing and see how it is affected by CPU. It can be determined by the size of the processor, the number of systems where it is used, or possibly even how it is used. Let's also look at your systems and see if there any ways of reducing this CPU cost.

Now we can start measuring. In my field software licensing is often tied to the highest CPU usage in a four-hour period in each month. So we need to measure to find this peak, then measure again to find the big users in this peak, and start targeting. And for each target, the same measure, evaluate, change, and re-measure cycle starts.

And this goes on. Once changes are made to reduce CPU usage, CPU usage is re-measured. Sometimes changes don't give us what we want, and sometimes our environment changes. A system change can affect things, or maybe our peak four-hour period moves. So after a first round, we measure again, determine the peak period, look for candidates and work from there.

So tuning isn't simply the measure, evaluate, change and re-measure cycle. It's also evaluating what to tune, and what is needed from tuning. Tuning is the considered, planned, targeted process of making thing run better.

When Is It Bad, and When Is It Better?

Another common reason I see for tuning is response time: the time between a user hitting the Enter key, and the screen coming back. Sometimes batch elapsed time is also an issue. However the goal “reduce response time” isn’t enough. If response time or batch elapsed time is an issue, specific targets are needed: Service Level Agreements (SLAs). For example: “90% of transactions completed within 1 second”, or “batch schedule completed by 6am.”

Without such specific targets, it isn’t possible to properly tune. If response times are sufficient, you’re wasting time tuning. If they aren’t, you don’t know how much you need to improve. In fact, excessive tuning can make things worse: impacting other areas or applications, increasing costs and impacting availability by change. Smart managers only change things when they need to be changed.

Tuning is the process of determining why things must be better, and making them so.

Summary

Tuning is an essential part of computing: ensuring that computer resources do what they need to do for the business. However tuning isn’t quick, isn’t knee-jerk and isn’t fast change. Tuning is careful, considered and long-term. Tuning is evolution: not revolution and not crisis management.

Tuning is the entire long-term process of making things run better.

Source:
LongEx Mainframe Quarterly - August 2014, Retrieved from
http://www.longpelaexpertise.com.au/ezine/WhyTuning.php

technical: Limiting the Damage from Loops and Spikes

By David Stephens, Lead Systems Programmer at Longpela Expertise

It was a shock when I first found the loop. I was at a client site to reduce their CPU usage when I noticed that they were burning far more CPU than normal. And it had been going on for days –using almost one entire processor engine. This was really going to hurt their software licensing costs.

Spikes and loops can be nasty creatures. They’re symptoms of problems or errors, which in themselves are a concern. But often worse, they affect other systems and applications around them. From increased costs to reduced performance from other systems or applications, spikes and loops can seriously affect processing. So how can you limit the damage they cause?

What are Loops and Spikes?

Any long-term application programmer has seen a looping program – a program that executes a loop an excessive number of times – sometimes infinitely. This is a simple example of a loop. Put this into a batch job and you have a batch job using excessive CPU until it is cancelled. Bad, but not catastrophic. Put this into an online CICS transaction, and things get worse. Other CICS applications can be starved of CPU, affecting their response time. If the transaction accesses a database inside the loop, the database workload jumps up. If the application continues to acquire storage inside the loop, then CICS storage can become constrained. This can cause short on storage conditions, and seriously impact CICS processing.

This sounds bad, but it could be even worse. Put an abend inside the loop, and you have an application that seriously stresses the CICS region and its dump handling. Put that loop in z/OS SRB code, and you can lose a processor. Put it into a z/OS exit and you may lose your z/OS system. Loops can be bad.

But loops aren't limited to application programs and assembler routines. They can appear in REXX execs, be triggered from automated job schedules (think of two jobs automatically scheduling each other), automated operations rules or even DFSMS ACS routines.

Loops are the most obvious cause of a spike – a short term, unexpected, large increase in CPU usage. But they're not the only causes. Consider a user accidentally running a transaction that scans all transactions for the past 12 months rather than the past 30 minutes. Or a business analyst setting a Business Rule that accidentally increases processing ten-fold. Maybe even an automated batch schedule that starts a batch suite needed only monthly every day.

When to Limit the Damage

Most sites concentrate on minimising loops and spikes, and this is an excellent idea. Proper testing and change management goes a long way. However many think less about ways of minimising the damage should one occur. One of the major reasons for this is that it is hard.

For starters, it is difficult to find the line: to determine when something is operating normally and when it is a loop/spike. Some systems may very well double their CPU usage at irregular intervals.

Another problem is the detection. A looping CICS transaction or batch job may be obvious. But a looping DFSMS ACS routine may only increase the CPU usage of the z/OS SMS address space – an address space rarely reported or monitored. Similarly a scheduling error may not register as a problem because all jobs are completing within their parameters: the problem may be the number of jobs submitted.

How to Limit the Damage

The key to limiting damage from loops and spikes is early warning. And this cannot be done without prior preparation. There are two approaches that can be used:

1. Kill the offending task
2. Alert

Kill the Offending Task
It may sound harsh, but a tactic used for decades on mainframes is to kill or cancel tasks using excessive CPU or memory. In the old days this was essential to prevent mainframes from crashing. Today it is unlikely that a looping or spiking task will crash z/OS, however the same principles apply. Killing an offending task achieves a few things:

- It stop the offending task from causing further damage
- It often produces a dump that can be used for problem determination
- It activates automatic notification procedures for abends, so the relevant technical staff will know about the problem quickly

So how can this be done? Let's take a few examples:

- Batch job – JCL JOB and STEP statements can be used to limit the resources that a step of job can consume. For example REGION and MEMLIMIT can limit the memory used – a useful fuse for a looping application continually acquiring memory. Smart sites limit who can specify REGION=0M (no limit). Similarly the TIME parameter specifies the maximum CPU time a job or step can use, and the PAGES parameter limits the amount of output sent to JES spool. Again, smart sites limit who can specify TIME=1440 (no limit). Many set reasonable defaults for all the above parameters using features such as the z/OS IEFACTRT exit.
- CICS – The CICS ICVR value specifies the runaway task time for all transactions. This can be overridden for individual transactions in the RDO Transaction resource RUNAWAY value.
- IMS – The IMS LOCKMAX parameter in the PSB or JCL override can limit the number of locks an application can acquire. The SEGNO parameter of a TRANSACT definition can limit the number of segments a message-driven transaction can acquire. Unfortunately, there is no SEGNO default – many IMS systems programmers use an IMS exit to set a default here.
- DB2 Stored Procedures – The ASUTIME parameter specifies the maximum CPU time in the stored procedure definition.
- Websphere Application Server – There are a few timeout values that can be specified for threads, together with a CPU time limit.
- SRBs – Few will need to code an SRB, but SRBs can and should be defined with a set CPU limit when scheduled.

In addition to these facilities, some monitors can be used to terminate tasks when set criteria are met. For example, Omegamon XE for CICS can cancel CICS transactions that consume more than a specified quantity of memory. However these facilities can be expensive in terms of CPU usage.

Alert

Most z/OS sites should have functions in place to automatically send alerts when a situation could impact z/OS processing. For example, alerts when system memory is low, WTO buffers are low, CPU usage is high, or free DASD space is low. However many sites only have a minimum of these alerts, if any.

Few sites will not have several monitoring tools for various systems and products. Often these can be configured to send alerts on set conditions. So TMON for CICS could be set to generate alerts when a CICS transaction exceeds a set CPU usage. Similarly automated operations can be used to generate alerts for all sorts of conditions: from batch processing missing windows to started tasks failing.

z/OS and related systems also include many facilities that can help. For example, the z/OS Runtime Diagnostics feature can generate an alert when an address space is using more than 95% of the capacity of a single CPU, or performing consistent repetitive processing (such as looping). Exits can also help. For example, the IMS DFSQSPC0/DFSQSSP0 exit can be used to report on IMS queue space issues – IBM provides samples.

Conclusion

The looping problem I found was caused by an output archiving product. It had a report definition that created an infinite loop, and it had been doing so for over a week. Luckily it didn't impact any other processing, but my client's software licensing costs took a hit.

Looping or spiking tasks can hurt, and they're not just in applications. Prevention is an important step in reducing such tasks, however it's wise to supplement this with procedures and functions to limit the damage should they occur. Setting your z/OS and related systems, together with their related monitors correctly can stop these tasks before they can perform too much damage, or at least notify support staff as soon as they occur.

Source:
LongEx Mainframe Quarterly - August 2014, Retrieved from
http://www.longpelaexpertise.com.au/ezine/SpikesLoops.php

opinion: Stop Firefighting and Start Tuning

By David Stephens, Lead Systems Programmer at Longpela Expertise

Recently I was at a site with an over-worked performance team. This isn't unusual, as performance is a large, time-consuming task. This small group was responsible for many z/OS systems in different Parallel Sysplexes. They not only were responsible for z/OS systems performance, but were involved in any performance issues: batch overruns, CICS response times, Websphere MQ delivery times. However the big problem was that they spent all their time dealing with problems: they were fire-fighting, not tuning.

Let's take an example. An application team rings up: "Our Websphere MQ performance between 10:00 and 10:30 this morning was 1.5 seconds. This is more than the 1 second specified in the Service Level Agreement. Tell us why and fix it!" So let's see what our performance group needs to do now:

1. They need to confirm what the application team has said. So they must go into the monitoring tools they have and take a look at the response times. Hopefully, they're familiar with the SLA, and can quickly confirm that the response time is too high.
2. They must confirm that this is unusual, by looking at response times at similar times in the past. If the Websphere MQ response time is always higher than the SLA, then there should already be a project working on this.
3. If they confirm there's a problem, then they need to find out why. So they'll look at the z/OS and related performance during that period to find out what happened. They'll also look to see what, if anything has changed.

All this is a lot of work. That 10 second phone call has taken out a few hours of a performance staff member's day. If this happens regularly, then that is all our performance team will be able to do. Full time fire-fighting

A much better approach would be to monitor performance. So let's look at the perfect scenario:

Our performance team has setup automated monitoring systems. Performance tools have been configured with SLAs and expected performance, so screens quickly show when things are outside of normal. Automated notifications (like emails) are sent to performance staff when something doesn't perform as it should.

Daily batch jobs analyse SMF records, and produce performance reports that are archived. Our performance team can quickly look at the past performance of critical systems. Trends can be seen, and potential issues addressed before they become problems.

If all this were the case, then our scenario above would be a little different:

1. Performance team is notified by automated systems that Websphere MQ performance between 10:00 and 10:30 wasn't sufficient.
2. Automated systems also notify the performance team that a CICS transaction is looping at the same time.
3. By the time the application teams rings, the performance team has already confirmed that the looping CICS transaction consumed excessive CPU, starving Websphere MQ. The problem transaction was terminated, and the relevant application team notified.

Our perfect approach took less than 30 minutes, and our performance team were fixing the problem before it was reported.

The problem is that I rarely see our perfect scenario. Simply put: performance isn't seen as important by management until there's a problem.

Setting up the automated procedures and configuring monitoring tools takes time, and is an ongoing process as things change. Many performance groups are too busy fire-fighting to setup and maintain this monitoring.

Performance monitoring and management is long-term. An investment to create and maintain infrastructure for effective, ongoing, automated monitoring will pay off again, and again, and again.

Source:
LongEx Mainframe Quarterly - August 2014, Retrieved from
http://www.longpelaexpertise.com.au/ezine/FirefightingAndTuning.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.