



interskill  
learning

INTERSKILL MAINFRAME TRAINING NEWSLETTER

August 2016

## Inside This Issue:

Interskill Learning New Releases for 2016	2
Interskill Learning Recent Releases	3
Data Center surveys...what are they really saying?	4
Vendor Briefs	8
Tech-head Knowledge Test	11
Learning Spotlight – DB2	11
<i>Management: A Five-Minute Guide to Monitors</i>	12
<i>Opinion: How Much Monitoring Do You Need?</i>	15
<i>Technical: Monitoring by Sampling</i>	16

# 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 New Releases for 2016:

### ★ AVAILABLE NOW

- **COBOL Programming Series – including:**
  - COBOL Programming Basics
  - COBOL Data and Datafile Definitions
  - COBOL Programming - Manipulating Data
  - COBOL Programming - Advanced
  - Accessing IMS Databases from COBOL
  - COBOL for z/OS
- **Linux on System z**
- **Enterprise Big Data and Hadoop**

### ★ COMING SOON

- **CICS TS for z/OS 5.3**
- **IMS V14**
- **z/Enterprise – IBM z13**
- **z/OS 2.2 Migration Briefings**
- **REXX**
- **PMBOK5 Project Management Curriculum**
- **Java on z/OS for Java Programmers**
- **Agile and Application Development**
- **COBOL Performance**
- **Improving Application Performance with Strobe, FreezeFrame and APA**
- **IMS TM (Transaction Manager)**
- **COBOL Basics Assessment**
- **COBOL Advanced Assessment**
- **PL/I Basics Assessment**
- **PL/I Advanced Assessment**

# Interskill Learning Recent Releases

- **IBM Development Environment**
  - Introduction to Language Environment
- **DevOps Curriculum**
  - DevOps Fundamentals
- **z/OS Connect Curriculum**
  - z/OS Connect
- **C Programming Curriculum**
  - C/C++ on z/OS for C Programmers
- **DB2 Curriculum**
  - The DB2 Operations Series updated to DB2 11.
    - DB2 Fundamentals
    - Managing DB2 Operations
  - The DB2 – Database Management System Series updated to DB2 11.
    - Introduction to RDBMSs and DB2
    - Manage Data Definitions with DB2
    - DB2 SQL Fundamentals
    - Advanced DB2 SQL
    - Create and Maintain DB2 Programs
    - DB2 Stored Procedures
    - Optimize DB2 Application Performance
- **Rational Developer for System z Series**
  - Rational Developer for System z Basics
  - Creating and Managing Applications Using RDz



# Data Center surveys...what are they really saying?

*By Greg Hamlyn*

Instinctively, most of us strive to find out what other people or organizations are doing successfully, in order to validate what we may be already doing, or to determine whether we should be heading down that track. In our work life, surveys are a good way of obtaining this type of information, but you need to be careful interpreting the survey's results!

## Who produces Data Center surveys?

Before analyzing the results from the numerous data center-related surveys produced, it is worthwhile exercise to look at the surveys themselves; who is producing them, whether the company has a vested interest in the results, what size was the survey and so on. If we don't have confidence in their validity then we are more likely to dismiss their results. So, let's take a look at some of data center surveys produced over the last year or two.

Survey	Topics	No. of respondents	Obtained
2015 BMC Annual Mainframe Research	<p>Top factors influencing mainframe growth.</p> <p>Top Mainframe priorities.</p> <p>Modernization of apps.</p> <p>Perception of the Mainframe's future.</p>	1202	Free after providing details.
AFCOM 2015 State of the Data Center Survey	<p>Growth and expansion in the data center world.</p> <p>Power, cooling, services and redundancy requirements.</p> <p>Cloud, storage, virtualization and security.</p> <p>New types of staffing and personnel dynamics.</p> <p>Evolving CapEx and OpEx models.</p>	Approx. 350	Need to be a member of AF-COM (although a summary is available through AFCOM's free Data Center Management digital magazine)
451 Research Voice of the Enterprise	<p>Details on the data center market, its size, and the rate it is growing.</p> <p>Who is building new data centers, and where are they building them.</p>	900 +	Need to become a subscriber (but there is trial access if you want to test it out)

InformationWeek - 2014 State of the Data Center	Data center budgets and what they are being spent on. Virtualization drivers and virtualization. Uptake of cloud services. Trend that will have the most impact in the next 12 months.	217	Free after providing details.
Syncsort - Mainframe's Role in the Big Data Ecosystem	Identifying the importance of mainframe and products to the organization. Top concerns over the next 12 months. Big data strategies in place. How mobile applications are being used/accessed.	187	Free after providing details.
SANS - The State of Dynamic Data Center and Cloud Security in the Modern Enterprise - 2015	How respondents are handling security aspects in today's dynamic computing environment. Types of data center attacks that have occurred and ones that organizations are most concerned with. Time between containment and recovery of security issues. Deployment types, compliance initiatives and security controls.	430	Free
Uptime Institute 2015 Data Center Industry Survey	New data center construction and co-location. Concerns today and tomorrow. Managing IT and data center costs. Energized electrical equipment.	1000	Free after providing details.
IDC Futurescape: Worldwide Datacenter 2016 Predictions	Data center infrastructure costs. Co-location and cloud data center. Network transformations. Green energy use in the data center.	undisclosed	Web presentation provides overview. Actual report at a cost.
ARCATI Mainframe Yearbook 2015	Mainframe growth. Configuration (operating hardware and software, middleware, processor types installed). IT Budgets. Data archiving strategies.	100	Free
Compuware - "The new world of mainframes" CIO Survey	CIO perspectives on the future on the mainframe, and strategic value in relation to new workloads. Mainframe developer shortages. Mainframe outage costs.	350	Free after providing details.

Unisphere Research - Key Shifts: Transforming the Data Center - 2016	Data growth within Data Centers and how Data Centers are handling it.  Big Data Projects and their reason for implementation.  Barriers and issues in relation to initiating big data projects.  How Apache Hadoop is being used.	319	Free after providing details.
--	---	-----	-------------------------------

**Digging beneath the survey surface**

There are certainly quite a few surveys, so how do you identify ones that are worthwhile? I regularly use a set of criteria to judge what I am looking at:

**Survey size** - A couple of these surveys do not seem to have many respondents, which could make it easier to skew the results.

**Survey statistics** - Many of the surveys disclose a breakdown of respondents (occupation, geographical location, industry), which is good but unless you are able to break down the survey results into the group you belong to, then they may not be as significant to you. Be suspicious of any surveys or summary results that don't provide demographic information.

**Survey questions divulged** - I like to see the actual questions that were asked as part of the survey, rather than a summary of the survey. Having been involved with questioning/assessments I know how easy it can be to lead respondents with questions and possible answers (i.e. through selective multiple choice answers). For me, the quality of questions provided in the survey is a main factor to its validity.

**Vested interest** - I like to work out whether the survey owner/sponsor has a vested interest in the outcome of the survey. For example, if a survey indicated that automation was considered an advantage over com-

petitors, and the survey was sponsored by a vendor selling automation solutions, then I would be wary of the results.

**Different points of view** - The majority of surveys are stand-alone in that they provide a point of view based on the questions they asked. I am usually keen to cross-reference this data with other reports to confirm these outcomes. A good example where this is already done for you is with AFCOM's survey, where aspects have been made available as part of magazine articles. These articles not only use their own survey results but substantiate their claims citing other surveys or results (e.g. 451 Research, Uptime Institute, and IDC).

You can see that I am a suspicious person by nature, but I truly believe that you should not trust everything you read. So, by now you are probably waiting for my analysis on what all these surveys are saying.

**Survey analysis**

Most of you will already have a good idea what is happening in data centers now, and where they are heading. See how your thoughts stack up against the compiled survey data from the sources mentioned above.

**Data center budget** - Certainly some conflicting data on this front. 451 Research reports that budgets are on the way up (87% report that they are maintaining or increasing). The Uptime Institute seems to support this, in-

dicating that 74% of third-party/co-location data center respondents received budget increases and 47% of other enterprise companies, although this seems to contradict itself as it mentions in a couple of other sections of this survey about “shrinking budgets” and “flattening enterprise IT budgets”.

The Informationweek survey mentions that while demand for resources is up (this is confirmed in several other surveys also) that this does not necessarily mirror additional funding. In fact they report that total budget for data center facilities, hardware and operations is down from previous years.

ARCATI’s survey takes a slightly different slant on IT budgets emphasizing the budget split between mainframe and other architectures, and highlights the value that organizations are getting from the mainframe (without more granular details on those budget spending costs it is difficult for me to agree or disagree). Compuware’s report also leads to the conclusion that non-mainframe platform costs are higher.

**Data center consolidation** - Several of the surveys mention that smaller data centers are under pressure to maintain cost competitiveness and are seriously looking at co-location or cloud as an alternative to buying or building a new data center. The Uptime Institute in particular highlights the continued growth of data center co-location, with a number of enterprises looking to build new data centers in the next 12 months. The AF-COM survey also reports that 45% of survey respondents expect to be building new data centers in the next 3 years.

**Cost optimization** - A number of surveys focused in on cost optimization as it was identified as a high priority for organizations to tackle. The BMC survey highlighted the reduction of resource usage during peak times,

and migrating workloads off the mainframe as the two main priorities, while the Syncsort report identified the CIO’s top priority was to reduce CPU usage and related costs. Related to this, the Uptime institute questioned whether organizations allocate IT and data center costs back to internal customers (for example, charge back).

**Staffing** - There were some interesting takes on the mainframe staffing front. The Syncsort report mentions that roughly a third of organizations won’t expect any significant loss of mainframe staff in the next 5 years, a third do but will be able to cover with existing staff, and the other third indicate a concern about covering the loss of staff. In the Compuware report, a high number of CIOs (61%) indicate that they are not concerned with mainframe developer staff shortages, while in the AFCOM survey they report that 71% of respondents will need to increase investment in data center IT and facilities staff within the next 3 years. Certainly some fluctuating results here.

**Security** - In the Syncsort report, Security is the 3rd highest priority for the CIO (behind Reducing CPU usage, and meeting Service Level Agreements). The SANS survey focuses on the types of attacks and threats that organizations are encountering, and how quickly they are able to contain them and recover from them. It continues by describing the various architectures found in the respondent’s organization (i.e. cloud, co-location, private data center) and the compliance and security controls that they have implemented. This information is usually hard to come by and makes for interesting reading. The Uptime Institute provides information on physical security (maintenance activities on energized electrical equipment) which is interesting, because it is not a topic that is often discussed in surveys. The AFCOM survey mentions that security is the biggest

concern for organizations when implementing a cloud architecture (32%).

**Longevity** - Several surveys asked respondents questions relating to their thoughts on the lifespan of the mainframe within their organization. Syncsort reports around 60% have no foreseeable plans to get rid of their mainframe, while in the BMC survey, 90% predict the mainframe's long term viability. In the Compuware report, 88% of CIOs indicated that they believe the mainframe will remain a key business asset for at least the next ten years.

### Summary

The good news for those within the mainframe industry is that the mainframe is not going away anytime soon based on these surveys. While a number new data center builds are forecast, many others seem to be seriously looking at co-location or movement to the cloud for some or all of its services. Details of data center budgets and staffing requirements are quite varied depending on the survey, which may reflect where the main survey respondents hailed from (i.e. fast growing cloud provider, or government based in-house environment). Security will always be near the top of priorities for any data center, but even cost savings in this area is becoming the norm.

Whichever surveys you subscribe to, ensure that you take a critical look at the criteria mentioned in this article first before determining whether it may be a suitable direction for you and if possible, obtain several surveys or reports to confirm trends or facts.



Happy investigations!

**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](mailto:g.hamlyn@interskilllearning.com)

## Vendor Briefs

Our Vendor Briefs section introduces products released by vendors over the last few months, with this issue primarily covering the many IBM releases that have occurred during this period. We also look at CA Technologies and BMC new offerings and round off the article with a look at some interesting trends and thoughts on the future of the data center.

### IBM

IBM has been extremely busy in a number of different areas. Here's an overview of what they have released.

#### z/OS 2.2

With two years now between z/OS releases, there is plenty packed into z/OS 2.2 which was made generally available the end of September 2015. Not surprisingly, it is designed to get the most value when coupled with the z13 mainframe, which was released earlier in that year. Some of the product enhancements are:

**z/OSMF** - This is now a base element so no need to order this product separately. Enhancements have been made to allow calling programs to retrieve, update, and

delete data from all systems in a sysplex or CPC. Some minor changes have also been made to the Workflows function in relation to substitution variables, and the ability for users to add their own steps to a workflow.

**JES2** - JES2 is one of those products that never seems to change but in this z/OS release there are some interesting enhancements. One of these is the new JES2 execution control statements, which allows you to configure JCL containing several jobs, so that they are run in a particular order. In other words, quick ad-hoc scheduling capabilities.

**SDSF** - New SDSF screens are able to display the virtual memory and devices used by a job, and identify any delays that are being experienced by jobs. There is a new panel to display job group information mentioned in the JES2 enhancement above. Where SDSF is running under ISPF you can now create a customized action character to invoke a REXX exec on a row.

**Other** - Other significant changes have also been implemented for ISPF, Generation Data Groups, SMP/E, Infoprint Server, and Communications Server.

### **LinuxONE**

IBM has recently expanded their traditional mainframe offerings to the point where data centers can now implement a Linux only solution, which is based on z Systems technology. This portfolio of hardware, software, and services is named LinuxONE and is available in two configurations - Emperor, and Rockhopper, of which the former is capable of scaling up to 8000 virtual machines. IBM has also opened up LinuxONE to open source software such as Apache Spark, Node.js, MongoDB, MariaDB, PostgreSQL, Chef and Docker, which strategically fits with its involvement

with the “Open Mainframe Project”. The Linux Foundation, IBM and a collaboration of nearly a dozen organizations across academia, government and corporate sectors have pledged their allegiance to this “Open Mainframe Project” to advance development and adoption of Linux on the mainframe.

### **z13s**

Earlier this year, IBM released its entry-level mainframe, the z13s. It is specifically aimed at the mid-sized market, and in particular, those players looking to implement a secure hybrid cloud environment. In a world where hacking of data is becoming an almost common occurrence, the z13s contains cryptographic and tamper-resistant hardware-accelerated coprocessor cards providing encryption at twice the speed of previous generations. IBM and its partners are also offering various other customized security solutions for z13s clients.

### **Storage**

IBM has been extremely busy on the storage hardware and software front. They have released IBM XIV Storage System Model 314 and associated IBM XIV Storage System V11.6.1 software, providing high-end, gridstyle storage with doubled RAM and CPU resources over its predecessor. A new DS8880 hybrid flash storage system has also hit the market, providing a mixture of flash and traditional disks to cater for each organization’s varied storage requirements. On the tape/cartridge front, many IBM tape systems now support the new LTO Ultrium 7 data cartridge that has a native 6 TB capacity (and 15TB with data compression).

### **z/VSE**

z/VSE is actually 50 years of age and still kicking strongly. Version 6.1 of this operating system was released back in November 2015 and where previous releases have

focused mainly on scalability and resource improvements, this one addresses improvements with online transaction processing, security, and networking. Associated with this release is the new CICS Transaction Server for z/VSE V2.1.

## IMS

IMS V14 was released in December 2015 with its main aim being to support the explosion of transaction and data that has occurred as a result of the ever increasing cloud, web-based and mobile applications. The Database manager contains improvements relating to the recovery of work and data, and increased scalability, while the Transaction Manager promises improved throughput, and increased application availability. A number of IMS Connect command enhancements have also been introduced.

## Other announcements

So many releases...so little newsletter space! Some other products that IBM has released lately include CICS TS V5.3 for z/OS, IBM Explorer for z/OS V3.0 Aqua (aka z/OS Explorer), and IBM Security zSecure V2.2 suite.

## CA Technologies

CA Technologies has recently upgraded their DevOps portfolio to include a number of new and improved products:

- CA API management products aimed at fast-tracking API creation (CA Live API Creator) and mobile app development (CA Mobile App Services)
- CA Service Virtualization (on Microsoft Azure) streamlines application simulation and testing
- CA's Virtual Network Assurance operations tool provides greater visibility of SDN/NFV network configurations, allowing advanced performance analysis

to identify bottlenecks and other network issues

- CA Unified Infrastructure Management for z Systems provides users with a unified view across mainframe and distributed systems

## BMC

BMC's Control-M product has come a long way since its initial simple mainframe batch scheduling days. Version 9, which has recently been released, is now part of BMC's Digital Enterprise Management strategy and encompasses a GUI interface that is able to track workloads across the entire IT environment. This new version in particular contains predictive analytical capabilities (used to improve runtime estimations), automated application workflow promotion between test and production environments, automated agent and client deployment for faster upgrades, and improved data security.

## Searchdatacenter.techtarget.com

To round off all this vendor activity, I thought that I would throw in a couple of links to interesting articles on the searchdatacenter.techtarget.com site. They focus on trends and the perceived future of data centers.

### [Top data center industry trends of 2016](#)

discusses the move from the traditional data center to one that is hybrid in nature...not just in relation to the cloud, but also managing other assets that are geographically spread and contain a mixture of staff and outsourced workers.

### [Data center uptime pressure mounts as IoT takes hold](#)

looks at the pressure mounting on data centers to handle the amount of digital data in the future and the technology coming our way. It has some interesting statistics on how the technology world was only nine short years ago.

## Tech-head Knowledge Test

With every release of this newsletter a mini online test will be provided of a product that you are likely to be using in the workplace. You may want to challenge your work colleagues or just want to confirm your own knowledge!

The test for this newsletter focuses on the IBM development environment as it relates to application programming on the mainframe, and consists of 13 questions.

Click the link below to start.

[IBM Development Environment](#)

## Learning Spotlight – DB2

Our DB2 courses have recently been revamped and updated to the current version of DB2 for z/OS, so we thought that you might enjoy looking at a module that contains one of the newer GUI products associated with it....IBM Data Studio.

The module that we have provided for you here is [Using SPUFI and IBM Data Studio to Access DB2](#), which describes the use of traditional ISPF SPUFI greenscreens to process SQL statements against DB2 data, and then moves to the IBM Data Studio GUI interface to achieve the same.

We hope you enjoy it.



# Management: A Five-Minute Guide to Monitors

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

We all have them. A monitor like TMON for CICS to monitor our CICS. Or Omegamon XE for IMS. And they're not cheap.

Aside from the cost of the software itself, they use CPU, and consume man-hours to commission and maintain. But what are all these monitors, and why do we have them? In this article, we give the essential five-minute guide to mainframe monitors.

## The Basic Categories

There are a lot of monitors on the market, with different purposes and features. However they can be divided into four basic categories:

- Operational
- Analysis
- Centralized
- Niche

Let's look at these categories.

## Operational Monitors

Operational monitors are exactly that – software designed to 'keep the trains

running'. They look at the system or subsystem they're monitoring, and report on the current health as it stands now. Let's look at operational CICS monitors. These provide information on the CICS region they are monitoring like:

- Current workloads – current transactions running. Information like how long they've been running, or how many file, DB2, temporary storage and transient data accesses they've made is shown. They will also show statistics such as CPU consumed, DSA (memory) used, and more.
- Resource status – resource usage for each transaction. A summary for the entire region is also shown: total memory usage, total CPU usage, and other metrics such as file I/O rates, DB2 activity etc.
- Problems – A screen that highlight problems. For example, a message may displayed if the CICS region hits its maximum task limit, or if transactions regularly run for longer than a set limit. Other possible problems include files or

connections that are closed, program abends/dumps, excessive storage use, or excessive file I/Os.

- Performance – performance statistics for each transaction. For example the service time of a transaction broken down into journal, DB2, Websphere MQ, file access and more. Similar information broken down by resource is also shown. For example, I/O performance by file, or LSR performance by LSR pool.
- General info –allow administrators to view resources values, and sometimes even change them. So a CICS Systems Programmer can see a program's definition, or system configuration values in use.

Most operational monitors throw in a few handy tools and gadgets. For example tracing features, screens to issue system commands, or tools to do things not possible with system commands. Some also include historical data recording and processing.

But perhaps the biggest strength of operational monitors is how they present this information. They all have screens that generally present information in a clear and easy-to-digest format. Screens are easy to navigate, normally with drill-down menus to allow problems to be identified and analysed fast.

Many monitors also include facilities to notify operators and administrators of any problems.

Operational monitors include IBM Omegamon XE, CA SYSVIEW, ASG TMON, and BMC MAINVIEW. They come in flavours for z/OS, CICS, IMS, TCP/IP, Websphere MQ, DB2 and more.

## Analysis Monitors

These monitors are used for in-depth

analysis. There's a lot of different types, so let's look at a couple of examples.

### Sampling Monitors

Sampling monitors take a lot of snapshots of an address space as it runs – usually around 1000 per minute. From these snapshots, or samples, they display what's happening over the sampling period. For example, they will show programs that consumed CPU (usually from highest to lowest), with options to further break this down to the location within each program. From this we can identify the programs causing high CPU usage, and target them for tuning. Similar analysis is provided for wait times, so service/elapsed times can be analysed.

Because sampling monitors take many samples, they have a record of this information over time. So they can show how waits, I/O or CPU usage change as a program runs from minute to minute. These monitors also include additional options to show information on things like DB2 calls, CICS transactions, MQ performance, file performance and more.

Our article [Monitoring by Sampling](#) looks at these in more detail.

### SQL Monitors

There are a few products that analyse DB2 SQL statement performance. These trace SQL calls as they are made, and measure their performance. So they provide information on poorly performing SQL statements, and the corresponding SQL source. These are great when working on DB2 application performance. BMC AppTune, CA Detector, and IBM Query Monitor are examples.

### Application Performance Monitors

Some monitors look more closely at application performance with tracing

or monitoring facilities. For example ASG TriTune provides much of the SQL monitoring above, and adds VSAM, IMS and Websphere MQ monitoring. Compuware PurePath solutions provide monitoring for CICS and Java applications.

## Centralized

The third category of monitors provide a centralized information point. Larger sites will have several z/OS regions, DB2 subsystems, Websphere MQ queue managers, and more. It's not uncommon to see a site with 100 separate CICS regions. Managing all of these can be very difficult. Enter the Centralized monitor.

Once example is IBM's Tivoli Enterprise Portal (TEP). This is a centralized GUI for managing different IBM monitors, both on and off the mainframe. So an administrator from one PC screen can see information from different IBM Omegamon monitors on different systems – as well as other IBM monitors on non-mainframe systems.

ConicIT also provides a centralized view. However it can take information from several monitoring products from different vendors, including Omegamon XE and BMC MAINVIEW. Like IBM TEP, screens can be tailored as needed.

Some monitors are simply availability checkers. They regularly confirm that a system or subsystem is available, and can also provide response time information. These monitors usually include historical information. Most are not mainframe based, though Inside Products Availability Checker is an exception.

Monitors such as Nastel AutoPilot M6 TransactionWorks, BMC Application Transaction Tracing and IBM ITCAM for

Transaction Tracking tackle a different problem. These display routing and performance information for composite applications crossing system and subsystem boundaries. For example a J2EE application on a UNIX system that accesses DB2 tables on z/OS. We talk more about composite applications in our articles [The Challenges Monitoring Composite Applications](#) and [Mainframe Products for Monitoring Composite Applications](#).

## Niche Monitors

There are many special purpose monitors that provide performance and other information for very specific areas not covered by normal operational monitors. For example, ARCH Consulting markets BATSTAT for CA-IDMS batch performance. Similarly CA-TSOMON provides detailed performance information for TSO, and IBM Financial Transaction Manager monitors financial transactions routed through Websphere Message Broker.

## Summary

No site will have every possible monitor. All will have a basic operational monitor for CICS, DB2 and IMS. Few will run without an operational z/OS monitor, and most will also have a sampling tool like Compuware Strobe. Other monitors are bought by companies as they are needed.

It's not unusual for a site to acquire a monitor to solve an immediate need, and then maintain that monitor even after the need has gone. For example, many sites are reviewing their SNA network monitors as SNA networks are replaced with TCP/IP. However new problems are encouraging new monitors such middleware monitors for Websphere MQ and SOA environments, or Java and Websphere Application Server monitors.

Monitors can be an invaluable aid in improving performance and ensuring the continued health of systems and applications. However poorly-customized or maintained monitors can be a drain on finances and system resources, without providing the value that they are capable of.

Source:

LongEx Mainframe Quarterly - August 2015,  
Retrieved from

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

## Opinion: How Much Monitoring Do You Need?

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

Monitors are wonderful tools. They provide all sorts of functionality to trace, sample and report. The problem is that monitors don't come for free. And I'm not talking about the cost of the software. The more you monitor, the more CPU you consume. They'll also chew up memory, and add disk load.

Whenever I'm on a CPU reduction project, monitors are one of the first things I look at. And the pickings can be good. Some monitor features can be enormously expensive when switched on. And this cost may not be obvious. For example, the CPU overhead of some DB2 monitors is recorded against each address space attaching to DB2. Looking at the DB2 address space CPU usage won't help.

I also look for duplicate monitoring. Do you really need to record TMON for CICS statistics to SMF when you're recording normal CICS SMF 110 records? Probably not.

But there is of course a flip side. Systems programmers, DBAs and application programmers still need some monitoring. Some of these can be switched on for a short period of time (like IBM APA). Other need to be continuously running to gain any benefit.

So when I'm on a tuning project improving response times, I'm sometimes looking to get tracing turned on. It can be a tricky line.

Which begs the question: how much monitoring do you really need? The answer is of course a frustrating "it depends". But here are some guidelines.

The first guideline is to estimate the CPU overhead. If it doesn't cost much, then there's no harm in turning it on. Sounds pretty basic, however the real CPU overhead can be difficult to estimate. Sometimes vendors will publish figures, but often they won't. In this case looking at the CPU overhead of the monitor address space may help. Some monitors such as CA SYSVIEW also provide screens estimating their own overhead which is nice.

The second guideline is to determine what you really need. OK, this sounds trivial, but it isn't. If you will probably never need a feature, then switch it off. I've seen many sites with features enabled that they will never need just in case someone else wants it in the future. Vendors will never need monitor features – they'll use the basic operating system/CICS/IMS features to get their information. If they do need extra information, they'll ask you to enable tracing or something else for short periods of time.

The third guideline is to avoid tracing. Tracing is the largest CPU consumer by far, and is almost never needed on an ongoing basis. If you need tracing, turn it on for

a short period, then turn it off. DB2 and Websphere MQ have some traces you will need for accounting or basic performance monitoring. However in general, monitor tracing is never required to be permanently enabled. Perhaps you could switch it on for 10 minutes outside of the highest peak times to get a 'feel' of what is happening. Or perhaps automation can automatically enable tracing when a problem or pre-defined set of conditions is met. But never keep it on all the time.

The final guideline is to minimise historical recording. As a rough rule, I never recommend historical recording for monitors – usually the basic SMF and similar statistics will give you what you need.

When all is said and done, monitors are a wonderful tool that can make problem diagnosis easier. Sensible configuration and use will maximise their benefit, while minimising their overhead.

Source:

LongEx Mainframe Quarterly - August 2015,  
Retrieved from

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

## Technical: Monitoring by Sampling

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

Often when we think about monitors we think about software like IBM Omegamon XE, BMC MAINVIEW and ASG TMON. These software products do things like regularly hop around storage and control blocks to find out what's happening right now. They'll often also insert 'hooks' or other magic to get information on current performance and workloads. They may even have tracing features to give you more details about what's happening, and some historical features to find out not only what's happening now, but what was happening in the past.

However there's another type of monitor: a sampling monitor. And the chances are that you have one.

### The First Sampling Monitor

In the 1970s, a company called Programart developed Strobe. If this sounds familiar, then it should. Compuware acquired Programart in 1999, and continue to market Strobe. Strobe is a sampling based monitor. This means that it will 'sample' an address space, meaning it will take lots of snapshots, to see what's happening. How often it samples is set when you define a sampling session. As a rough guide, I like to sample around 1,000-1,500 times per minute when analysing something for more than 5 minutes. So if we sample 1,000 times per minute, Strobe will check every 0.06 seconds to see what an address space is doing. Compuware isn't publishing how it does this, but we can make a few educated guesses. For a start, Strobe is checking to see if any task in the address space is currently

executing on a processor. If it is, it will find out from the PSW the address, and from this (in most cases), the module.

So we can get a report that looks something like this:

MODULE NAME	SECTION NAME	LINE NUM	PROCEDURE/FUNCTION NAME	STARTING LOCATION	CPU SOLO	TIME % TOTAL
.COMMONX			EXTENDED COMMON AREA	1823FB48	5.34	5.71
.XES	IXLR1LKR		LOCK W/RECORD STRUCT		5.09	5.50
.SUPERVS	IGC0013{		SUPERVISOR SERVICES		2.38	2.42
.XMEMORY	.XMEMORX		EXTENDED CROSS MEM		2.05	2.26
.CICS	DFHAIP	DFHEIP	EXEC(CMD-LEVEL)INT		2.05	2.05
.SVC	SVC 011		TIME		1.93	1.93
.LELIB	CEEPLPKA	CEEV#GTS	GET HEAP STORAGE		1.81	1.85
.COBLIB	IGZCPAC		COBPACK GENERAL		1.81	1.85
.DB2	DSNXGRDS	DSNXEPM	PGM MGR		1.68	1.81
.DB2	DSNK2DM		IDX MGR		1.77	1.81

OK, but what does this all mean? What it tell us is where the CPU is being burned – in which modules. For example, we can see that the TIME SVC (SVC11) is consuming around 2% of the total CPU, and Language Environment GET HEAP almost the same. From this, we can target the CPU users, and start to tune them. Whenever I’m in a project looking to reduce CPU, this report is an excellent first step. Strobe will even go into more detail. So you can find the DB2 DBRM, the offset in a COBOL module, or the CICS transaction consuming the CPU. Nice huh?

## Competition for Strobe

In 2004, Macro4 announced a new product as a direct competitor to Compuware Strobe: Freeze Frame. Shortly after in 2005, IBM announced an almost identical product: IBM Application Performance Analyzer. Like Compuware Strobe, these products are sampling monitors, producing reports on CPU usage for an address space. Reports break down CPU usage by module, and like Strobe, can go down further into the DBRM, Transaction and offset.

However all three products go further. They analyse wait times – why an address space was waiting. FreezeFrame can product a report similar to the following.

Name	NTxns	Description	Error	Execution	Mean Time in Seconds + Suspend + Delay = Service
CSMI	25464	Mirror transactio	0.6%	0.000	0.165 0.046 0.213
> DFHMIRS		CICS Program		0.000	0.161 0.044 0.206
> CICS		System Services		0.000	0.161 0.039 0.201
> IRLINK		Wait on InterRegion Link		0.000	0.081 0.000 0.081
> FCIOWAIT		wait on VSAM File I/O		0.000	0.064 0.000 0.064
> CICSdly		CICS Dispatch Delay		0.000	0.000 0.039 0.039
> ZCIOWAIT		wait on Terminal I/O		0.000	0.015 0.000 0.015
> FCXCSUSP		FILE001 FILE001		0.000	0.000 0.000 0.000
> LG_FORCE		wait on Log Buffer		0.000	0.000 0.000 0.000
> LG_DEFER		wait on Log Buffer		0.000	0.000 0.000 0.000
> LGWRITE		wait on Writing Log		0.000	0.000 0.000 0.000
		Record			

So you can see that for a CICS region, FreezeFrame has broken down the service time by

transaction. So our CSMI transaction has most of its service time waiting on the InterRegion Link LINKMM1, with VSAM File I/O on file FILE001 coming a close second.

## Other Information

Our three sampling products go further than this. They include reports on things as varied as:

- DASD and file I/O values, counts and performance
- SQL activity
- Websphere MQ statistics

And more. A gold mine of tuning information.

## Cost vs Accuracy

All this information comes at a cost. For a start, a sampling session must be manually defined. You define the address space (job, region or started task), the monitoring options to enable, the number of samples to take, and how long to run. The number of samples and length of sampling run are important. Let's look at them:

You want the length of the sample to be long enough for a representative sample if monitoring online or recurring workloads. For example, I like a good 30 minute sampling period when monitoring a CICS region. For batch, the chances are that you want to measure the entire step you're looking at. All products have a flag that can be set to monitor until step-end. As an aside, they also all offer scheduling controls so you can setup a sampling session today that will run tomorrow, or next week, or every time a batch job runs.

The number of samples is also important. You don't want this too high, or you'll burn a lot of CPU monitoring the address space. If you have too many samples, the monitoring may also get in the way of the address space you're monitoring, skewing the results. On the other hand, you want enough samples to get a good look. If you only get 10 samples where a batch job is using CPU, the CPU reports are useless.

All three products include a report that is something like this:

Total samples	45,000	Duration	29 min 59.96 sec
Sampling rate	25.00 per sec	Report dataspace	505.81MB
CPU/WAIT samples	36,545	Sample dataspace	863.95MB
TCB samples	45,000	Meas significance	81.21%
CPU queued samples	8,455		

In this example, we took 45,000 samples over a 30 minute period (well, 29 minutes, 59.96 seconds). During this period, we obtained 36,545 samples where the address space was executing, or waiting. This is good. However if this values was something like 50, you don't have enough samples.

If I'm trying to perform a very short sample (say, for a job step running for 15 seconds), I'll increase the number of samples to get a good number of CPU/WAIT samples. If I can't get a few hundred samples, I won't trust the output.

## System z CPU Measurement Facility

In their z10 Systems z processors, IBM added the System z CPU Measurement Facility (CPU MF for short). This allows you to perform your own sampling without any of the above products. The HIS started task in z/OS manages this facility: operator commands start, stop and define sampling session parameters. This started tasks also collects these samples, and writes them to a z/OS UNIX dataset for analysis. CPU MF also provides some counter functionality and a new SMF Type 113 record. But we'll concentrate on the sampling side.

The CPU MF has the advantage over the previous products in that it is a hardware feature, so the sampling won't affect the results obtained. It will also sample everything that executes – all address spaces. It also uses less CPU than the sampling monitors.

The disadvantages start with the fact that this facility cannot be limited to one address space. You sample the entire z/OS system, or nothing. This will provide a huge amount of information, so it's rare to exceed 10 minutes in a sampling session. So sampling that 30 minute batch job step won't help. Further, this facility dumps information into a z/OS HFS. You need to do the analysis. Phoenix Software International market the zHISR product that can format and analyse this information.

CPU MF also won't provide much of the DB2, CICS Transaction, Websphere MQ, file I/O and other information provided by the above sampling monitors.

## Summary

Sampling monitors are excellent tools for analysing the performance of batch and online applications and systems, providing information on the top CPU consumers, and why things are waiting. The CPU MF facility is a free alternative. However it provides different information that is harder to analyse, and does not provide many of the add-on benefits of the other sampling monitors.

Source:

LongEx Mainframe Quarterly - August 2015,  
Retrieved from  
<http://www.longpelaexpertise.com.au/ezone/SamplingMonitors.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.