



IBM Inside Sales

International Technical Support Organization Global Content Services

Exploiting Recent IBM z/OS Features and Functions

www.ibm.com/redbooks

WRZ002

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*BladeCenter®, DB2®, e business(logo)®, DataPower®, ESCON, eServer, FICON, IBM®, IBM (logo)®, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, S/390®, System p®, System p5, System x®, System z®, System z9®, System z10®, WebSphere®, X-Architecture®, zEnterprise, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zSeries®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.





Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM has two registered trademarks for the branding of ITSO publications. These registered marks are for the text word "IBM Redbooks" and the Redbooks logo. In a nutshell, the term Redbooks must always be used in the plural form (for both text and logo) since IBM only owns the registered mark for the plural form. Usage must follow the guidelines below:

Using the term Redbooks in written text

Redbooks are only to be referred to in the plural form, NEVER in the singular.

For the initial reference (first occurrence), you must use "IBM Redbooks®" and include "IBM" as well as the ®. For instances thereafter you may use "Redbooks" without "IBM" preceding the word or ® following it.

Correct usage for written text :

In this IBM Redbooks® publication we will explore.....(® symbol required for 1st usage)

This Redbooks publication will show you.....(2nd usage or later - no ® or "IBM" needed)

Using the logo:

Redbooks (logo)



OTHER ITSO PUBLICATIONS - Marks not yet registered

Trademark registration is a lengthy process and until we are officially registered, we cannot use the ® symbol. For those terms/logos in process, we will be using the ™ symbol. In contrast to the ® symbol (placed in the lower right hand corner), the ™ symbol is placed in the upper right hand corner. Please see examples below:

Redpaper ™
 Redpapers ™
 Redwiki ™
 Redwikis ™



The following terms are trademarks of other companies:

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.




Welcome

- Our seminar topic for the day is Exploiting Recent IBM z/OS Features and Functions
- I am Marna Walle (mwalle@us.ibm.com)
 - z/OS Development, Systems and Technology Group, z/OS System Installation
 - I'm replacing for this seminar: Karan Singh: karansin@us.ibm.com
 - ITSO Redbooks Project Leader
 - Systems Programmer
- Work location is IBM's Poughkeepsie site, which is located in New York's approximately 90 miles north of New York City.



ITSO System z Workshop Tour



Yearly seminars focusing on new IBM hardware and software

2012 Topics

- IBM zEnterprise EC12 Technical Overview
- Exploiting Recent IBM z/OS Features and Functions
- Parallel Sysplex and High Availability
- DB2 v10 for z/OS and Query Accelerator
- Using z/VM v6.2 and Linux on System z Advanced Topics



Exploiting Recent IBM z/OS Features and Functions

- Typically the z/OS day focuses on a summary of new features and functions delivered with the new release of z/OS
- With the recent announcement of the 2 year release cycle, our z/OS workshop day:
 - Will focus on the features and functions of the new release during the year it is released
 - Will focus on exploiting new features and functions between z/OS releases
 - Deeper look at usage, installation and customization of selected topics

Agenda

Exploiting Recent IBM z/OS Features and Functions

<i>1</i>	Introduction
<i>2</i>	z/OS V2 announcement
<i>3</i>	zEC12 software support
<i>4</i>	Toleration support for GRS management of FICON CTCs
<i>5</i>	Flash Express and RSM enhancements
<i>6</i>	Soft Failures
<i>7</i>	IBM zAware
<i>8</i>	Runtime Diagnostics
<i>9</i>	IBM Health Checker for z/OS
<i>10</i>	Predictive Failure Analysis
<i>11</i>	Message Flood Automation
<i>12</i>	SMF Record Flood support
<i>13</i>	Auto-reply to WTOR
<i>14</i>	Auto IPL
<i>15</i>	Parallel Subsystem Initialization
<i>16</i>	Conclusion



z/OS Version 2 announcement

z/OS
A smarter operating system
for smarter computing



z/OS Version 2 Statement of Direction April 11, 2012

Summary of Statement of Direction

- Version 2 of z/OS targeted to Availability 2H2013
- A **Version 2** release is the **next** z/OS release; there is no release in 2012
- Release delivery cycle is planned to be every two years, in the second half of that calendar year.
- z/OSMF will be on the same cycle as z/OS
- N+1 or N+2 release migrations support two or four year migration cycle
- New five Year Support with optional fee based service extension
- For interim updates, IBM intends to offer updates, such as for server support, through periodic small product enhancement PTFs (SPEs) and web deliverables
- Minimum supported HW levels (z9 server or later, and 3990-3 controller or later)
- **Changes to Current Releases:**
- R12 Support extended to four years,
- R13 Support extended to five years to bridge customers to Version 2 migrations
- R13 orderable until Version 2 becomes available

Why Were the Drivers Behind the New Release Schedule?

- Customers have told IBM that they find it difficult to keep pace with annual release cycle
 - Currently, customers spend significant time preparing for their next release cycles
 - Operational costs related to planning are difficult to sustain
 - In today's economy, customers prefer to focus their skills, time and resources on applications to drive business growth
- In addition, IBM has delivered related functions over several successive releases, making it challenging for customers to derive value
- IBM intends to package its function in fewer increments designed to result in a more simplified approach for customers
- Such packaging can also have the effect of reducing field reliability exposure for APARs
 - Potential for improved quality
- The SOD intends to preview the release schedule for planning and is not intended as a content view

**Get more value out of z/OS and the mainframe
Spend less time on planning and deployment**

Annual cycles split functions over multiple releases

Functions	z/OS V1.7	z/OS V1.8	z/OS V1.9	z/OS V1.10	z/OS V1.11	z/OS V1.12	z/OS V1.13
Consoles restructuring	X	X		X			
Parallel VARY processing	X	X					
Password phrases		X	X	X			
SMF data to System Logger			X		X	X	X
NFS V4				X	X		
Message Flooding Automation			X	X			X
System REXX™ (SYSREXX)			X		X		
Large (1 MB) pages				X	X	X	
Extended Address Volumes (EAVs)				X	X	X	X
zHPF					X		X
Implementing functions over several releases makes it more difficult to see value in each step							



Support and Service Changes



- Transition Timing (to 5+3)
 - Release 11 has 3 Years Support +2 LCE
 - Release 12 has 4 Years Support + optional fee-based extended service
 - Release 13 has 5 Years Support + optional fee-based extended service .

- Migration and Support
 - Longer support lifecycles align with delivery cycles
 - Maintain N-2 release migration (accommodate a 2 or 4 year migration cycle)
 - Bridge migration from currently supported releases

	.08	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
R 1.10		3 Year Support			Life Cycle Ext.		Extended Service								
R 1.11		3 Year Support			Life Cycle Ext.		Extended Service								
R1.12			4 Year Support				Extended Service								
R1.13				5 Year Support					Extended Service						
V 2.1					5 Year Support					Extended Service					
V 2.2							5 Year Support					Extended Service			
V 2.3										5 Year Support					Ext



z/OS V2 Minimum Supported Hardware

- **z/OS V2 support of IBM System z9[®] servers and above**
 - System z9, IBM System z10[®], IBM zEnterprise™ 196 (z196) and IBM zEnterprise 114 (z114)
 - **Not supported: IBM eServer™ zSeries[®] 990, 900, 890, 800 (z900, z800, z990, z890)**
- **z/OS V2 support of 3990-3 Storage Controller and above**

Beginning with z/OS Version 2, IBM plans to require these storage control units and later control units:

- 3990-3, 3990-6
- DS8000 family (2107,2121,2122,2123,2124)
- ESS Family (2105)
- DS6000 family (1750)
- Rack mounted DASD (9340,9341, 9342, 9343)

- Provides a z/OS V2 foundation for future optimization, performance, and reliability
- Helps enable incremental exploitation of servers in future releases.
- Recompiles using more current compilers to yield potential performance benefits
- Elimination of dual path offers simplification of development efforts
- Reduces potential for errors by reducing complexity associated with supporting and testing multiple code paths

Other benefits:

- Server / Control Unit benefits for z/OS, z/OS stack, and ISVs

Customers with z900, z800, z990, and z890 servers and older storage control units should plan upgrades now in preparation for z/OS V2.



z/OS Support Summary



z/OS®	z800/ z900	z890/ z990	z9® EC z9 BC	z10 EC™ z10 BC	z196	zBX	z114	DS8000® DS6000®	TS1130	End of Service	Coexists with z/OS...	Planned Ship Date ²
R7	X	X	X	X ⁴	X ⁴			X ⁴	X	9/08	R9	
R8	X	X	X	X	X ⁴		X ⁴	X	X	9/09	R10	
R9	X	X	X	X	X ⁴		X ⁴	X	X	9/10 ¹	R11	
R10	X	X	X	X	X	X	X	X	X	9/11 ¹	R12	
R11	X	X	X	X	X	X	X	X	X	9/12 ²	R13	
R12	X	X	X	X	X	X	X	X	X	9/14 ²	V2R1 ²	
R13	X	X	X	X	X	X	X	X	X	9/16 ²	V2R2 ²	
V2R1²			X	X	X	X	X	X	X	9/18 ²	V2R3 ²	2H13

- 1. Fee-based service extension available
- 2. All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- 3. IBM Lifecycle Extension for z/OS V1.7 (5637-A01) was required
- 4. Fee-based service extension required for support, or for some features



IBM Lifecycle Extension for z/OS V1.11

IBM Lifecycle Extension for z/OS V1.11 offers an optional fee-based corrective service for users who have not completed migration from z/OS V1.11 to z/OS V1.12 or z/OS V1.13

■ What's new:

- The IBM Lifecycle Extension for z/OS service offer is now available for z/OS V1.11.
- It provides fee-based corrective service (a fix, bypass, or restriction to a problem) for up to two years starting with October 1, 2012, the earliest offering start date, up through September 30, 2014.

■ Features / Business Value:

- Offered through IBM System z, the Lifecycle Extension provides only *corrective service* for z/OS.
- Flexible terms and conditions:
 - Purchase services for any number of machines, aggregate pricing available
 - Add and delete machines as you migrate to newer releases
 - Option to add additional months
 - Keeps billing active for uninterrupted support; cancel when you are done with migration

■ Client Benefits:

- For z/OS V1.11 customers needing additional time to complete their migration to the next release, the Lifecycle Extension for z/OS V1.11 provides defect support after the end of program services.



Learn More: <http://www.ibm.com/systems/z/os/zos/>

More Information

See IBM z/OS Version 2 SOD Announcement

- <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS212-086>

See z/OS SOD FAQ

- <http://public.dhe.ibm.com/common/ssi/ecm/en/zsq03055usen/ZSQ03055USEN.PDF>

See z/OS V1.11 LCE Announcement

- <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS212-025>

See z/OS home page

- <http://www-03.ibm.com/systems/z/os/zos/>



z/OS SOD Announcement

Geo	Letter
US	212-086 -BP
AP	AP12-0019
CAN	A12-0204, B12-0084
LA	LP12-0118
EMEA	ZP12-0025
JPN	JP12-0012

zEC12 software support



IBM zEnterprise EC12

- On August 28th, 2012 IBM announced the zEnterprise EC12 (zEC12)
- General Availability on September 19, 2012

- The IBM zEnterprise EC12 (zEC12) system is comprised of:
 - The IBM zEC12 Central Processor Complex (CPC),
 - The IBM Unified Resource Manager,
 - The IBM zEnterprise BladeCenter[®] Extension (zBX) Model 003

- z/OS platform software requirements for:
 - Hardware upgrade to a zEC12 server
 - With or without zBX
 - Install a new zEC12 server
 - With or without zBX

IBM zEC12 System Functions and Features

Five hardware models
5.5 GHz processor chips
Up to 100-way on z/OS.11 or higher
Doubled HSA – 32 GB
Improved processor cache design
Improvements to pre-fetch instructions
Up to 3TB real memory (1TB per LPAR)
Improved availability with Redundant Array of Independent Memory (RAIM)
Industry standard Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) I/O drawer
Flash Express (Storage Class Memory - SCM)
1 MB Pageable Large Pages
Dynamic reconfiguration support for Flash Express ¹
2 GB Large Page Support ¹
Optional PLPA, COMMON page data sets ¹
Crypto Express4S cryptographic coprocessors and accelerators
New support for PKCS #11 Hardware Security Module (HSM)
DUKPT for MAC and Data Encryption



zEC12

zBX

(z/OS support in blue)

New and enhanced instructions
IBM zAware
OSA-Express4S (GbE LX and SX, 1000BASE-T , 10 GbE LR and SR)
FICON Express8S
Parallel Sysplex InfiniBand (PSIFB) Coupling Links
High Performance FICON for System z
CPU Measurement Facility
CFCC Level 18 enhancements
Java exploitation of transactional memory
Exploitation of new hardware instructions – XL C/C++ ARCH(10) and TUNE(10)
Common Criteria Evaluation Assurance Level 5+ (EAL5+) certification
zEnterprise BladeCenter Extension (zBX) Model 003 support of general purpose blade servers and optimizers
zBX-003 with IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z)
zBX-003 with select IBM BladeCenter PS701 Express blades, or IBM BladeCenter HX5 (7873) blades

¹ - Planned target 1Q2013. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.





z/OS support summary

Release	z900/z800 WdfM	z990/z890 WdfM	z9 EC z9 BC WdfM	z10 EC z10 BC WdfM	z196 CPC	z196 w/zBX	z114 CPC	z114 w/zBX	zEC12 CPC	zEC12 w/zBX	End of Service	Extended Defect Support ¹
z/OS V1.7 ²	X	X	X	X	X						9/08 ¹	9/10 ¹
z/OS V1.8 ²	X	X	X	X	X		X				9/09 ¹	9/11 ¹
z/OS V1.9 ²	X	X	X	X	X		X				9/10 ¹	9/12 ^{1*}
z/OS V1.10	X	X	X	X	X	X	X	X	X	X	9/11 ¹	9/13 ^{1*}
z/OS V1.11	X	X	X	X	X	X	X	X	X	X	9/12 ¹	9/14 ^{1*}
z/OS V1.12	X	X	X	X	X	X	X	X	X	X	9/14 [*]	9/16 ^{3*}
z/OS V1.13	X	X	X	X	X	X	X	X	X	X	9/16 [*]	9/19 ^{3*}
z/OS V2.1 ^{4*}			X	X	X	X	X	X	X	X	9/18 [*]	9/21 ^{3*}

Notes:

- 1 The IBM Lifecycle Extension for z/OS provides the ability for customers to purchase extended defect support for that release of z/OS for up to 24 months after the z/OS release's end of service date
- 2 See IBM GTS services for additional fee-based extended service
- 3 Optional extended service is planned to be offered
- 4 z/OS V2.1 announced as an IBM Statement of Direction
- Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

WdfM – Server has been withdrawn from Marketing

Legend

Out of Lifecycle Extension for z/OS support²

Defect support provided with Lifecycle Extension for z/OS

Generally supported



Supported Releases

- zEC12 capabilities differ depending on z/OS release
 - Toleration support provided on z/OS V1.10 and z/OS V1.11
 - The Lifecycle Extension for z/OS V1.10 or z/OS V1.11 is required to acquire toleration PTFs and for support
 - Exploitation support provided on z/OS V1.12 and higher
 - z/OS V1.12
 - Exploitation of selected functions
 - z/OS V1.13
 - Exploitation of most functions
 - z/OS V2.1*
 - Full exploitation in base



zEC12 PSP Bucket and Fix Categories

- Software requirements differ depending on z/OS release and functions exploited
 - Support provided via a combination of web deliverables and PTFs
 - zEC12 Required PTFs
 - Documented in zEC12 PSP Bucket: Upgrade = 2827DEVICE, Subset = 2827/ZOS
 - Broken out by z/OS release, identifying required and exploitation
 - Can be identified by SMP/E commands (REPORT MISSINGFIX, LIST, or APPLY)
 - Fix Categories:
 - IBM.Device.Server.zEC12-2827
 - IBM.Device.Server.zEC12-2827.Exploitation
 - IBM.Device.Server.zEC12-2827.ParallelSysplexInfiniBandCoupling
 - IBM.Device.Server.zEC12-2827.ServerTimeProtocol
 - IBM.Device.Server.zEC12-2827.zHighPerformanceFICON
 - IBM.Device.Server.zEC12-2827.UnifiedResourceManager
 - zBX Required PTFs
 - Documented in zBX PSP Bucket: Upgrade = 2458DEVICE, Subset = 2458/ZOS
 - Broken out by z/OS release, identifying required and exploitation
 - Can be identified by SMP/E commands (REPORT MISSINGFIX, LIST, or APPLY)
 - Fix Categories
 - IBM.Device.Server.zBX-2458

z/OS Toleration Support for zEC12

- z/OS V2.1*
- z/OS V1.13
- z/OS V1.12
- z/OS V1.11 (Will no longer be generally supported after September 30, 2012. The IBM Lifecycle Extension for z/OS V1.11 (5657-A01) will be required for extended defect support for z/OS V1.11 for up to 24 months after the z/OS V1.11 end of service date)
- ➔ ■ z/OS V1.10 (No longer generally supported. The IBM Lifecycle Extension for z/OS V1.10 (5656-A01) provides the ability for customers to purchase extended defect support for z/OS V1.10 for up to 24 months after the z/OS V1.10 end of service date)

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

z/OS Toleration Support for zEC12

- Provides same functionality as that on the IBM zEnterprise 196 (z196)
 - HiperDispatch, OSA-Express4S, FICON Express8S, Parallel Sysplex InfiniBand (PSIFB) Coupling Links, CF Level 17, High Performance FICON for System z (zHPF), CPU Measurement Facility, ...
- Plus for z/OS V1.10 and z/OS V1.11
 - OSA-Express4S (GbE LX and SX, 1000BASE-T, 10 GbE LR and SR)
 - Crypto Express4S toleration
 - GRS FICON CTC toleration
 - New z/Architecture® Instructions (and new OP CODE support)






z/OS Toleration Support for zEC12

Release	zEC12 PSP Bucket – 2827DEVICE 2827/z/OS																		
	Base Support									Exploitation Support									
	Base zEC12 Support	OSA-Express4S (Gbe LX and SR, 1000BASE-T, 10 Gbe LR and SR)	FICON Express8S (CHPID FC)	IFB Coupling Links	Crypto Express4S Toleration ²	High Performance FICON (ZHPPF)	CPU Measurement Facility (HIS)	GRS FICON CTC Toleration	New z/Architecture Instructions	CF Level 18	Crypto Express4S Exploitation ³	XL C/C++ ARCH(10)/TUNE(10)	IBM zAware (z/OS Monitoring)	Transactional Execution Facility	Flash Express (Storage Class Memory - SCM)	Pageable Large Pages	Dynamic reconfiguration support for Flash Express ⁴	2 GB Large Page ⁴	Optional PLPA/COMMON page data set support ⁴
z/OS V1.10 ¹	P	P	P	P	W P	P	P	P	P	N	N	N	N	N	N	N	N	N	N
z/OS V1.11 ¹	P	P	P	P	W P	P	P	P	P	N	N	N	N	N	N	N	N	N	N

- 1 – The Lifecycle Extension for z/OS V1.10 (5656-A01) is required to acquire toleration PTFs and for support. The Lifecycle Extension for z/OS V1.11 (5657-A01) is required for support after September 30, 2012
- 2 – A Crypto Web Deliverable (HCR7770 or higher) AND a PTF is required for toleration. Support differs depending on the Crypto Web Deliverable installed
- 3 – Crypto Exploitation differs based on the Crypto Web Deliverable installed
- 4 - Planned target 1Q2013. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.
- 5 - z/OS V2.1 announced as an IBM Statement of Direction

B – FMID shipped in Base product
 P – PTF is required
 W – FMID shipped in a Web Deliverable

z/OS Exploitation Support for zEC12



- z/OS V2.1*
- z/OS V1.13
- z/OS V1.12

- z/OS V1.11 (Will no longer be generally supported after September 30, 2012. The IBM Lifecycle Extension for z/OS V1.11 (5657-A01) will be required for extended defect support for z/OS V1.11 for up to 24 months after the z/OS V1.11 end of service date)

- z/OS V1.10 (No longer generally supported. The IBM Lifecycle Extension for z/OS V1.10 (5656-A01) provides the ability for customers to purchase extended defect support for z/OS V1.10 for up to 24 months after the z/OS V1.10 end of service date)

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.



z/OS Exploitation Support for zEC12

- Provides same functionality as that on the IBM zEnterprise 196
 - IBM zEnterprise Unified Resource Manager, Network and Performance Management, intranode management network (INMN) and intraensemble data network (IEDN), Static Power Save Mode, Three subchannel sets, IPL from alternate subchannel set, zDAC support, OSA-Express3 and OSA-Express4S Inbound Workload queuing (IWQ)
 - z/OS V1.13 (and higher) only
 - OSA-Express4S checksum offload for IPv6, OSA-Express4S checksum offload for LPAR to LPAR traffic (both IPv4 and IPv6), **Large Send for IPv6**, Inbound Workload queuing (IWQ) for Enterprise Extender traffic, **HiperSockets optimization for intraensemble data networks (IEDN)**
- z/OS z/OS V1.10 and z/OS V1.11 functions plus:
 - XES/XCF Support of New Hardware (CF Level 18)
 - Crypto Express4S Exploitation
 - Enterprise Security PKCS11- Hardware Security Module (HSM), DUKPT for MAC and Data Encryption, Cipher Text Translate CCA Verb, PKDS/TKDS Constraint Relief, FIPS Evaluation, Common Criteria, Random Number Cache, FIPS on Demand, Wrapping Keys with Strong Keys
 - z/OS V1.13 (and higher) only
 - Java exploitation of the Transactional Execution Facility
 - Exploitation of New Hardware Features - C/C++ Arch(10) / Tune(10)
 - IBM zAware (z/OS Monitoring)
 - RSM Enhancements
 - Flash Express Support
 - Pageable 1MB Large Page Support
 - Dynamic reconfiguration support for Flash Express - target 1Q2013*
 - 2 GB Large Page Support – target 1Q2013*
 - Optional PLPA and COMMON page data set support – target 1Q2013*

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

© 2012 IBM Corporation





z/OS Exploitation Support for zEC12


Release	zEC12 PSP Bucket – 2827DEVICE 2827/z/OS																		
	Base Support									Exploitation Support									
	Base zEC12 Support	OSA-Express4S (Gbe LX and SR, 1000BASE-T, 10 Gbe LR and SR)	FICON Express8S (CHPID FC)	IFB Coupling Links	Crypto Express4S Toleration ²	High Performance FICON (ZHPF)	CPU Measurement Facility (HIS)	GRS FICON CTC Toleration	New z/Architecture Instructions	CF Level 18	Crypto Express4S Exploitation ³	XL C/C++ ARCH(10)/TUNE(10)	IBM zAware (z/OS Monitoring)	Java exploitation of the Transactional Execution Facility	Flash Express (Storage Class Memory - SCM)	Pageable Large Pages	Dynamic reconfiguration support for Flash Express ⁴	2 GB Large Page ⁴	Optional PLPA/COMMON page data set support ⁴
z/OS V1.12	P	B	B	B	WP	P	P	P	P	P	W	N	N	N	N	N	N	N	N
z/OS V1.13	P	B	B	B	WP	P	P	P	P	P	W	P	P	P	WP	WP	WP	WP	WP
z/OS V2.1 ^{4,5}	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B

- ¹ – The Lifecycle Extension for z/OS V1.10 (5656-A01) is required to acquire toleration PTFs and for support. The Lifecycle Extension for z/OS V1.11 (5657-A01) is required for support after September 30, 2012
- ² – A Crypto Web Deliverable (HCR7770 or higher) AND a PTF is required for toleration. Support differs depending on the Crypto Web Deliverable installed
- ³ – Crypto Exploitation differs based on the Crypto Web Deliverable installed
- ⁴ – Planned target 1Q2013. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.
- ⁵ – z/OS V2.1 announced as an IBM Statement of Direction

B – FMID shipped in Base product

P – PTF is required

W – FMID shipped in a Web Deliverable



Supported z/OS Releases and ICSF Levels

- z/OS V1.10 Crypto customers can run with:
 - HCR7750 – Base z/OS V1.10
 - HCR7751 – Cryptographic Support for z/OS V1.8 through z/OS V1.10 and z/OS.e V1.8
 - HCR7770 – Cryptographic Support for z/OS V1R9-V1R11
 - HCR7780 – Cryptographic Support for z/OS V1R10-V1R12
- z/OS V1.11 Crypto customers can run with:
 - HCR7751 – Base z/OS V1.11
 - HCR7770 – Cryptographic Support for z/OS V1R9-V1R11
 - HCR7780 – Cryptographic Support for z/OS V1R10-V1R12
 - HCR7790 – Cryptographic Support for z/OS V1R11-V1R13
- z/OS V1.12 Crypto customers can run with:
 - HCR7770 – Base z/OS V1.12
 - HCR7780 – Cryptographic Support for z/OS V1R10-V1R12
 - HCR7790 – Cryptographic Support for z/OS V1R11-V1R13
 - HCR77A0 – Cryptographic Support for z/OS V1R12-V1R13
- z/OS V1.13 Crypto customers can run with:
 - HCR7780 – Base z/OS V1.13
 - HCR7790 – Cryptographic Support for z/OS V1R11-V1R13
 - HCR77A0 – Cryptographic Support for z/OS V1R12-V1R13
- z/OS V2.1* Crypto customers can run with:
 - HCR77A0 – Base z/OS V2.1

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

© 2012 IBM Corporation



z/OS Support for zEC12

- Base support for zEC12 is provided by PTFs
- Exploitation of many functions is provided by PTFs
- Exploitation of Crypto Express4S requires a web deliverable on z/OS V1.12 and higher
- Exploitation of RSM Enhancements, including Flash Express exploitation, requires the z/OS V1R13 RSM Enablement Offering web deliverable, installed on z/OS V1.13
 - Planned availability December 14, 2012

Toleration support for GRS management of FICON CTCs



New zEC12 Multisystem Considerations

- zEC12 does NOT support ESCON
- Previously, GRS could not directly manage FICON CTCs
- Two unrecommended GRS Ring configurations are affected
 - GRS Ring complex that is larger than one sysplex
 - GRS Ring that doesn't utilize sysplex signaling
- Migration Options:
 - Convert to a Parallel Sysplex exploiting GRS Star
 - GRS star is recommended over GRS Ring
 - Convert to a Basic Sysplex exploiting XCF signaling for GRS Ring
 - Sysplex communications recommended over GRS-managed CTCs
 - Install zEC12 maintenance to provide toleration for FICON CTCs
 - This support does not enhance the robustness of GRS-managed CTCs
 - The toleration must be installed across GRS complex

When to use

- If there is a need for GRS-managed CTC devices because of the need to include non-sysplex members in the GRS complex
- Apply PTFs for APAR OA38230
- Add the FICON CTC devices to GRSCNFxx
- Implement via a rolling IPL or complex wide IPL – steps are detailed in:
 - See HOT TOPICS z/OS Newsletter August 2012 issue 26
 - “Retro Ring GRS’ new support of FICON CTC communication
 - Comments for APAR OA38230

Flash Express and RSM enhancements



Introducing System z Flash Express (target 12/2012*)

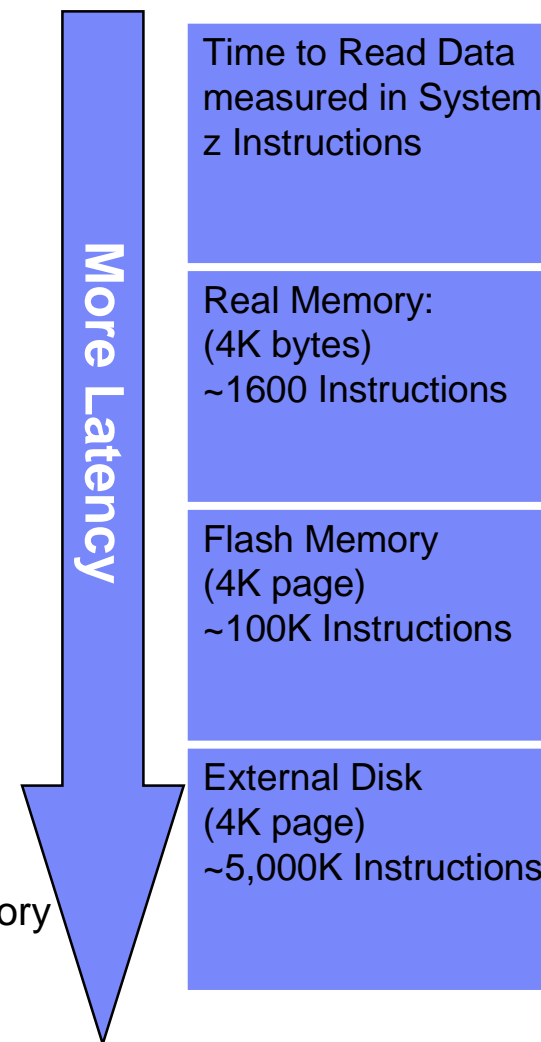
- Flash Express is intended to improve System z availability
 - Slash latency delays from paging
 - Make your start of day processing fast
 - Eliminate frustrating delays from SVC dump processing

- zEC12 will offer optional System z Flash Express memory cards
 - Supported in PCIe I/O drawer with other PCIe I/O cards
 - Pairs of cards for availability
 - No HCD/IOCP definition required

- Assign flash memory to partitions like main memory
 - Assignment is by minimum/maximum memory amount, not by feature
 - Each partition’s flash memory is isolated like main memory
 - Dynamically increase the partition maximum amount of flash
 - Dynamically configure flash memory into and out of the partition

- Options to solve many different problems
 - Flash Memory is much faster than spinning disk
 - Flash Memory is much slower than main memory
 - Flash Memory takes less power than either


- The system z Software Stack has a staged plan to exploit flash memory
 - z/OS 1.13 plus PTFs, JAVA, etc.
 - Future: IBM middleware (DB2)
 - Future: Linux on System z (IBM is working with distributors)



* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.



Relative Access Times for different technologies




Flash Express for z/OS Paging Value

- Flash Express is a faster paging device compared to a hard disk
 - The value is NOT in replacing memory with flash but replacing disk with Flash
 - Flash is suitable for workloads that can tolerate paging and will not benefit workloads that cannot afford to page
 - The z/OS design for flash memory does not completely remove the virtual storage constraints created by a paging spike in the system. (Some scalability relief is expected due to faster paging I/O with flash memory.)

Flash Express PCIe Adapter Card

Must have 4 x SSD cards. Each 400 GBs



Flash Express PCIe Adapter Card with 2 ports to form a RAID 10 Mirrored Pair



2 ports to form
A RAID 10 Mirrored
Pair with the 2nd
Flash Card


Flash Express - Twin-Cable to Form a RAID 10 Mirrored Pair

- Flash Express cards are always installed in pairs
 - Maximum 4 pairs in a System

- Installed in a PCIe I/O Drawer in 2 different I/O Domains
 - Maximum of 2 pairs installed in a drawer
 - One Flash Card per Domain only
 - Greater than 2 pairs will require a second PCIe I/O Drawer
 - eConfig will reserve a slot in each Domain in case Flash Express is ordered in the future
 - Cards first installed in the front of the installed drawers (slots 1 and 14) before using the rear slots (25 and 33)

- Flash Express Cards are cabled together to form a RAID 10 Mirror for redundancy

- Data on the Flash Card is protected with a unique key stored on the Support Element (SE) harddisk



Twin-Cable interconnect between the two Flash Express cards in slots 1 and 14

z/OS Flash Express Use Cases

■ Paging

- z/OS paging subsystem will work with mix of internal Flash and External Disk
 - Self Tuning based on measured performance
 - Improved Paging Performance, Simplified Configuration
- Begin Paging 1 MB Large Pages only to Flash
 - Exploit Flash's random I/O read rate to gain CPU performance by enabling additional use of Large Pages. Currently large pages are not pagable.
- Begin Speculative Page-In of 4K Pages
 - Exploit Flash's random I/O read rate to get Improved Resilience over Disruptions.
 - Market Open, Workload Failover

z/OS Flash Express Use Cases ...

■ Dumping

- Minimize SVC Dump duration, System impact
 - Flash performance during SDUMP
 - Flash performance after SDUMP
- Reduce Stand Alone Dump duration
 - Read time for paged out data

Flash Express Requirements

Flash Express:

Prereqs:

- zEC12, with Flash Express feature
- Flash card pair memory size is 1.6 TB
 - Min: 1 Card Pair = > Max: 4 Card Pairs
 - Mirroring across pairs for resiliency
 - Typical configuration: 1 card pair

Configuration/Setup:

- Prior to Enablement; cards will be on line, unformatted
- Firmware update delivery will be handled through zEC12 Bundle process
- Firmware update via concurrent patch Config off/on for enablement
- Array build via H/W setup (CE/SE)



z/OS (Flash, Pageable Large Pages)

Prereqs:


- V1.13 + Web deliverable 12/14 support + 1Q 2013 planned deliverables :
- 4GB Real Storage needed for pageable large pages (not specific to Flash)

Configuration/Setup:

- At IPL, z/OS detects if flash is assigned to the partition. z/OS automatically uses Flash for paging unless specified otherwise via PAGESCM=NONE
- All paging data can easily reside on Flash
- No need for fine grained capacity planning
- No definition required for placement of data on flash
- Flash Express uses sub-channels. These are allocated from the .25K reserved in sub-channel set 0

Typical Customer Configurations for Flash Express

- Flash card pair memory size is 1.6TB
 - Min: 1 Card Pair = 1.6 TB
 - Max: 4 Card Pairs = 6.4 TB
- Typical customer configuration is 6 to 8 LPARs per CEC and 40 GB - 80GB for paging configuration dataset size
- Even with 10 LPARs per CEC, each LPAR has 160 GB of flash memory available for its paging datasets, more than double the current typical customer configuration.
 - All paging data can easily reside on Flash
 - Data will preferably go to flash and only go to disk (if any) when flash is full
 - No intelligent placement of data on internal flash needed



P00MNXX4: Manage Flash Memory Allocation

Manage Flash Memory Allocation - P00MNXX4

Summary

Allocated:	32 GB	Storage increment:	16 GB
Available:	0 GB	Rebuild Complete	0 %
Uninitialized:	0 GB		
Unavailable:	0 GB		
Total:	32 GB		

Partitions

--- Select Action ---

Select	Partition Name	Status	IOCDS	Allocated (GB)	Maximum (GB)
<input checked="" type="radio"/>	LP1	Inactive	A0,A1,A2,A3	32	240
<input type="radio"/>	LP2			0	0
<input type="radio"/>	LP3			0	0
<input type="radio"/>	LP4			0	0
<input type="radio"/>	LP5			0	0
<input type="radio"/>	LP6			0	0
<input type="radio"/>	LP7			0	0
<input type="radio"/>	LP8			0	0

OK Apply Refresh Cancel Help

Allocating Flash to a partition

- The initial and maximum amount of Flash Memory available to a particular logical partition is specified at the SE or HMC via a new Flash Memory Allocation panel
- Can dynamically change maximum amount of Flash Memory available to a logical partition
- Additional Flash Memory (up to the maximum allowed) can be configured online to a logical partition dynamically at the SE or HMC
 - For z/OS this can also be done via an operator command
- Can dynamically configure Flash Memory offline to a logical partition at the SE or HMC
 - For z/OS this can also be done via an operator command
- Predefined subchannels, no IOCDS

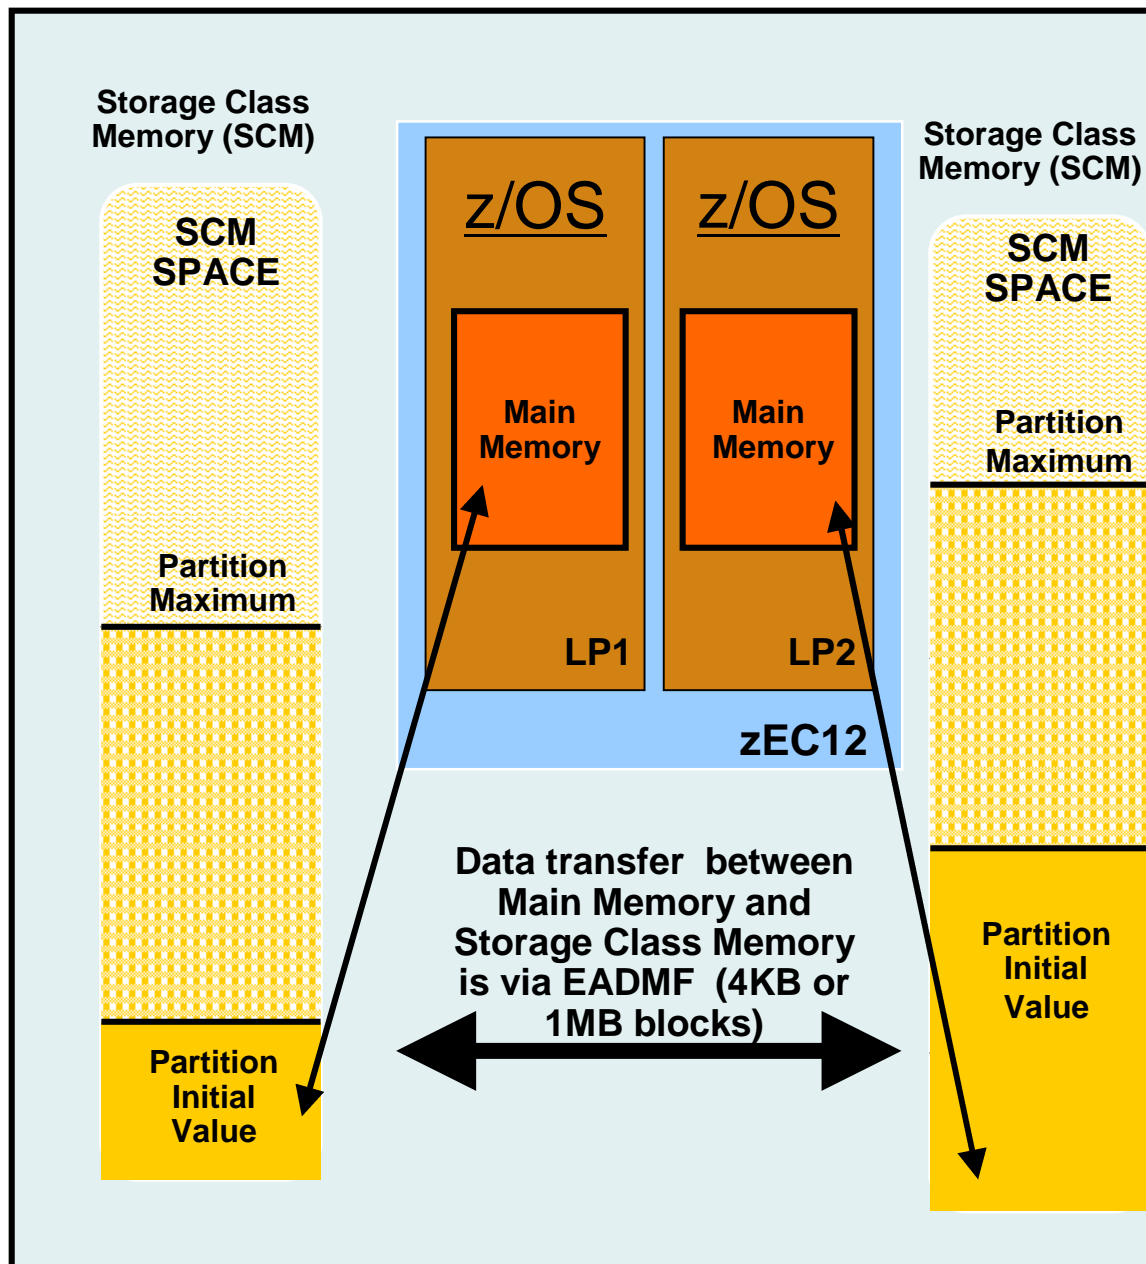


▪ **Full virtualization of physical Flash PCIe cards across partitions, software sees an Abstracted Flash Storage Space (Storage Class Memory)...**

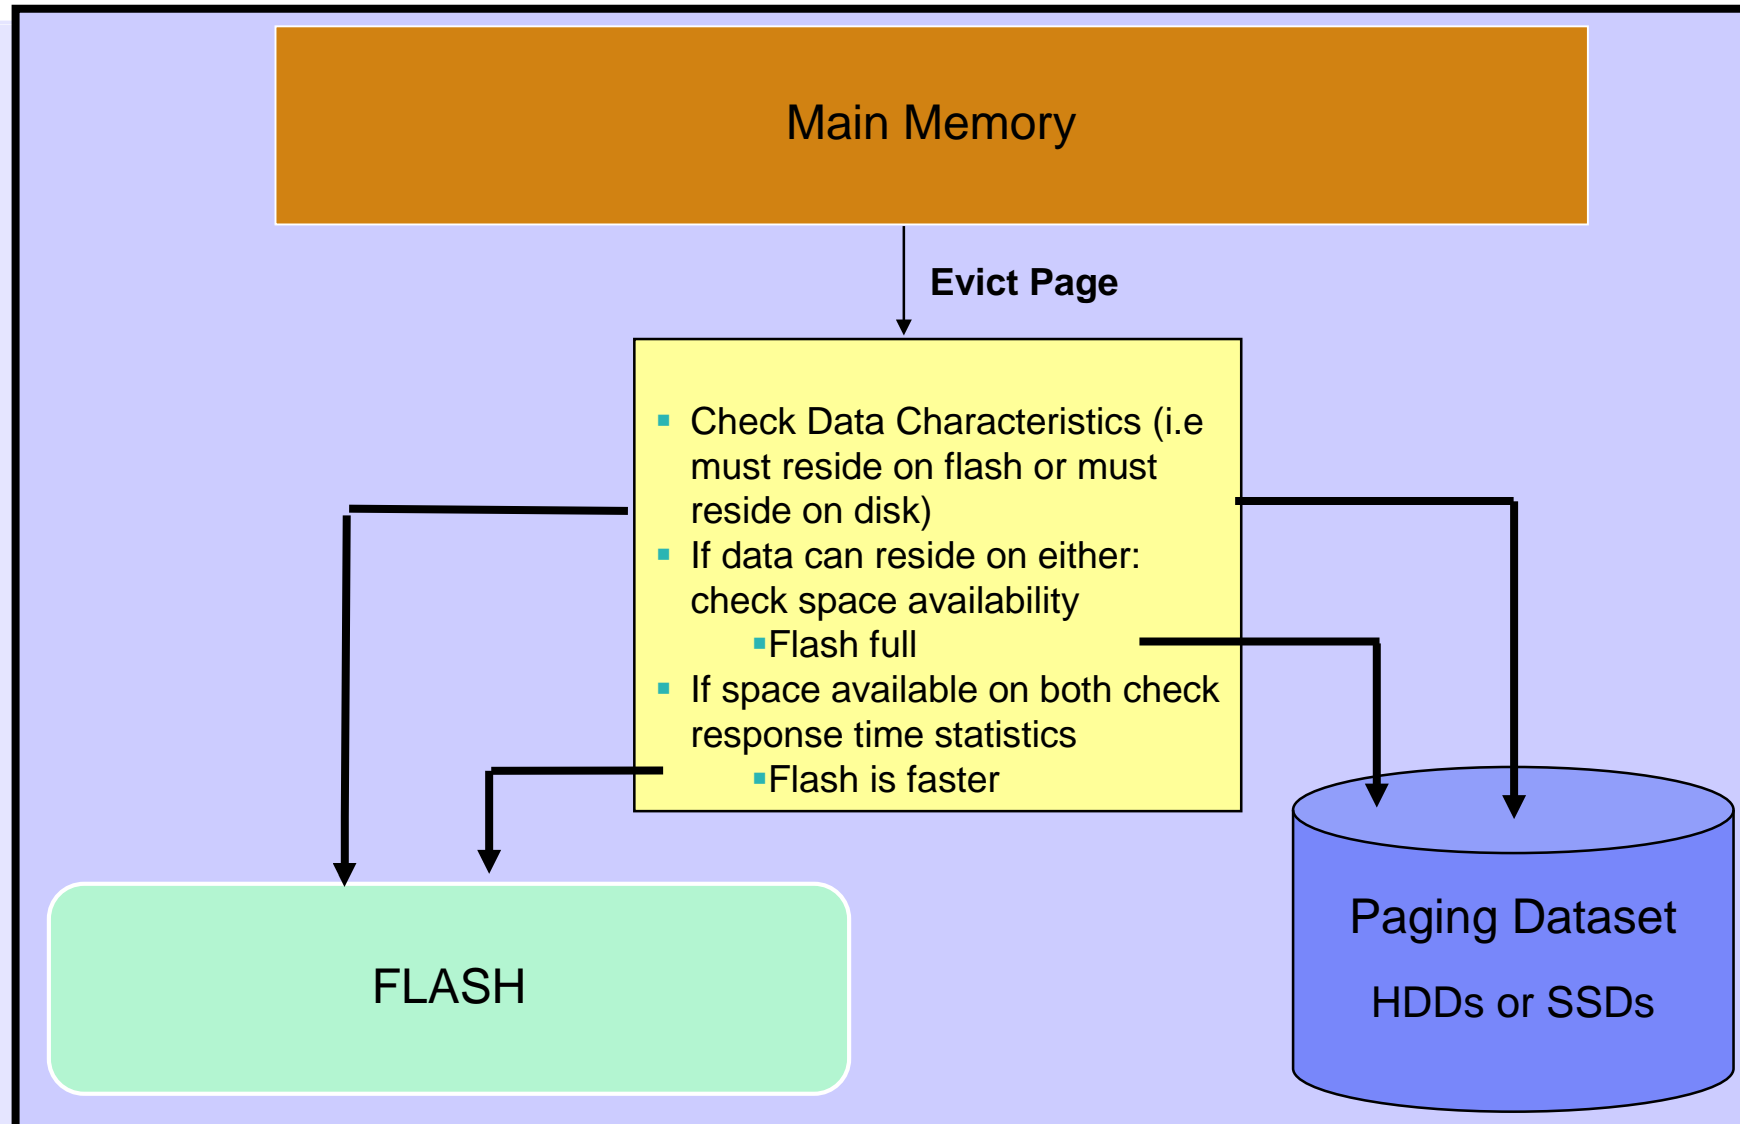
- Allows each logical partition to be configured with its own SCM address space
- Allocate Flash to partitions by amount, not card size
- Ability to change underlying technology while preserving API

▶ **No Hardware Specifics in Software.**

- Error Isolation, Transparent mirroring, Centralized diagnostics, etc.
- Hardware Logging, FRU Call, Recovery: Independent of software



Flash vs Disk Placement Criteria





Flash vs Disk Placement Criteria

Data Type	Data Page Placement
PLPA	At IPL/NIP time PLPA pages will be placed both on flash and disk.
VIO	VIO data will always be placed on disk (First to VIO accepting datasets with any spillover flowing to nonvio datasets)
Hyperswap Critical Address Space data	<p>If flash space is available, all virtual pages belonging to a Hyperswap Critical Address Space will be placed on flash memory.</p> <p>If flash space is not available, these pages will be kept in memory and only paged to disk when the system is real storage constrained and no other alternatives exist.</p>
Pageable Large Pages	<p>If contiguous flash space is available, pageable large page will be preferentially written to flash.</p> <p>If contiguous flash is not available the page will be demoted and placed on disk or flash based on which medium has the best service request response time.</p>
All other data	If available space exists on both flash and disk then make a selection based on response time.



Flash Express Exploitation Considerations

- **New IEASYSxx parameter - PAGESCM**
 - Enables the use of Storage Class Memory (SCM) for paging, specifies the minimum amount of storage class memory to be reserved for paging
 - Value may be specified in units of M, G, or T
 - If Flash Express is installed but is not to be used for paging, then PAGESCM=NONE should be specified
 - Defaults to **ALL**

- The D ASM and D M commands are enhanced to display flash-related information/status
 - D ASM lists SCM status along with paging data set status
 - D ASM,SCM displays summary of SCM usage
 - D M=SCM display SCM online/offline and increment information
 - D M=SCM(DETAIL) displays detailed increment-level information
 -

- The CONFIG ONLINE command is enhanced to allow bringing additional SCM online
 - CF SCM(amount),ONLINE

Flash Memory Usage and Invocation

- New messages issued during IPL indicate the status of SCM

- IAR031I USE OF STORAGE-CLASS MEMORY FOR PAGING IS ENABLED -PAGESCM=ALL, ONLINE=00001536M

OR

- IAR032I USE OF STORAGE-CLASS MEMORY FOR PAGING IS NOT ENABLED - PAGESCM=NONE

Flash Express Exploitation Considerations

d asm

```

IEE200I 17.17.46 DISPLAY ASM 944
TYPE      FULL STAT   DEV  DATASET NAME
PLPA      100% FULL   02E6  SYS1.PLPA.PAGCOM
COMMON    61%   OK    02E6  SYS1.COMMON.PAGCOM
LOCAL     0%   OK    098E  SYS1.LOCAL.PAGEP2
LOCAL     0%   OK    0987  SYS1.LOCAL.PAGEP3
LOCAL     0%   OK    098F  SYS1.LOCAL.PAGEP4
SCM       11%   OK    N/A   N/A
    
```

d asm,scm

```

IEE207I 17.35.02 DISPLAY ASM 947
STATUS      FULL      SIZE              USED              IN-ERROR
IN-USE      11%    393,216          46,954              0
    
```



RSM Enhancements

- RSM Enhancements delivered in the z/OS V1R13 RSM Enablement Offering Web Deliverable (FMID JBB778H) for z/OS V1.13
 - Flash Express Support - Exploits Storage Class Memory (SCM) technology for z/OS paging and SVC dump
 - Pageable 1MB Large Page Support
 - Is expected to yield substantial improvements in SVC dump data capture time, and removes the requirement for PLPA and Common page data sets when used for cold start (CLPA) IPLs.
 - It can also be used to remove the requirement for non-VIO local page data sets when the configuration includes enough SCM to meet peak demands.
 - However, local page data sets remain required for VIO, and when needed to support peak paging demands that require more capacity than provided by the amount of configured SCM.
 - Dynamic reconfiguration support for Storage Class Memory (SCM) - target 1Q2013*
 - 2 GB Large Page Support – target 1Q2013*
 - Optional PLPA and COMMON page data set support – target 1Q2013*
 - PLPA/COMMON (*NONE*)

RSM Enhancement Considerations

- Installation of the z/OS V1R13 RSM Enablement Offering Web Deliverable (JBB778H) will:
 - Increase the size of the Nucleus by approximately 380K above the 16MB line
 - You may need to analyze your private storage usage
 - Increase of 24K (6 pages) in ESQA per CPU per LPAR
 - This increase in ESQA per CPU includes general purpose CPs, zIIPs, and zAAPs.

2 GB Large Pages (target 1Q2013*)

- Increase TLB coverage without proportionally enlarging the TLB size by using 2 GB large pages:
 - A 2 GB page is a memory page that is
 - (2048 times) larger than a Large page and
 - (524,288 times) larger than an ordinary base page
 - 2 GB Large Pages allow for a single TLB entry to fulfill many more address translations than either a large page or ordinary base page
 - 2 GB Large Pages will provide exploiters with much better TLB coverage, and therefore provide
 - Better performance by decreasing the number of TLB misses that an application incurs
 - Less time spent converting virtual addresses into physical addresses
 - Less real storage used to maintain DAT structures
 - FLASH EXPRESS is not required

2 GB Large Pages Exploitation Considerations


- Enhanced IEASYSxx parameter - LFAREA
 - Enhanced to support the use of 2 GB large pages
 - The LFAREA parameter can be specified as:
 - LFAREA = ([1M=req] [,2 GB=req] [,prompt | noprompt])
 - Note: The old form of the LFAREA keyword is still supported:
 - LFAREA = (xx% | mmmmmmM | ggggggG | tttttT)

- Usage Enhancements
 - IARV64 GETSTOR enhanced to support the request for 2 GB large pages

Soft Failures



Complex Systems



IBM Mainframes are a complex environment

WebSphere

Information Management


Tivoli

Lotus

Cognos

Java

Rational



- Hardware
- Software
- Compilers
- Database systems
- Transaction processing systems
- Middleware
- User applications

In any system there is the potential for failure

The job of the System Programmer is to deal with failure



Types of failure on System z and z/OS

Not all failure are alike
Broadly categorize failures into 3 types

Masked Failure

- Software/Hardware detects failure
- Software/Hardware corrects failure
- No impact to business
- Example: Hardware power supply failure: switch to alternate, IBM alerted, concurrent replacement

Hard Failure

- Software/Hardware detects failure
- Automations and operations restart the failing component
- Minimal impact to business
- Example: Application terminates but is restarted by ARM

Soft Failure

- User detects failure, impact to business.
- Difficult to determine recovery actions
- Example: component is failing, holds resources (locks, enqueues) required by other components, causes sysplex wide stall, leads to sysplex wide IPL.



What is a soft failure

Systems don't break, but seem to just stop working:
"sick but not dead" or soft failure

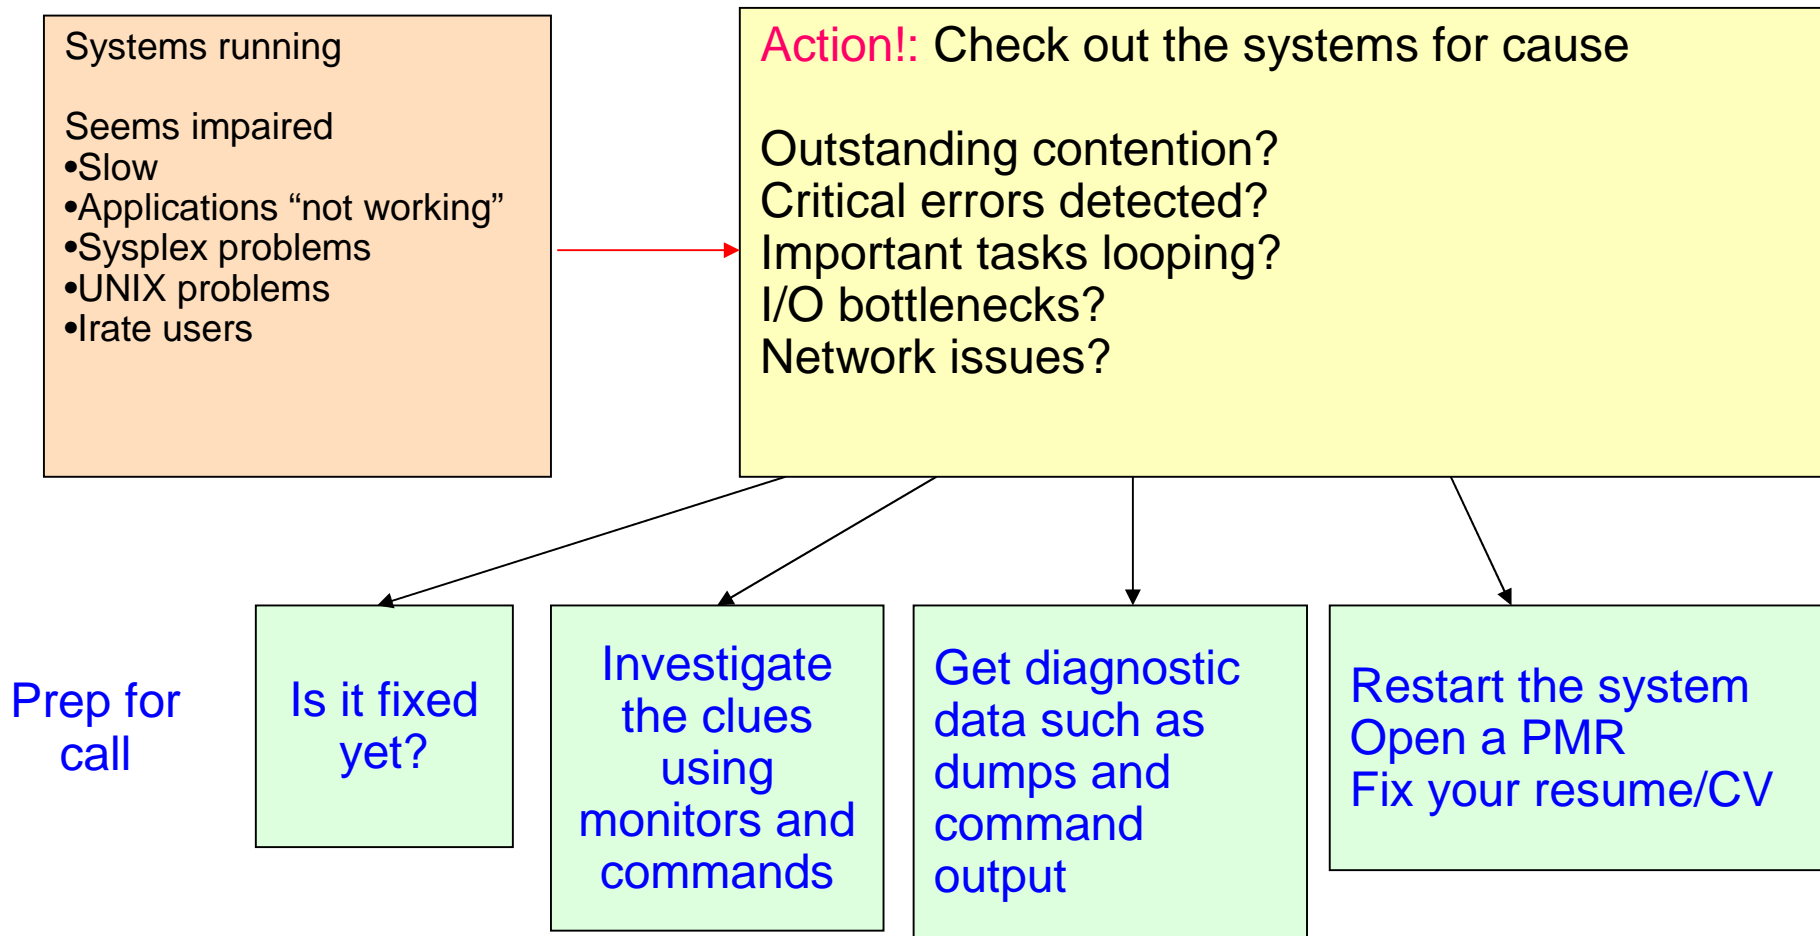
Symptoms of a Soft Failure

- 80% of business impact, but only about 20% of the problems
- Long duration
- Infrequent
- Unique
- Can be software or hardware
- Cause creeping failure and "sympathy sickness"
- Hard to determine how to isolate
- Hard to determine how to recover
- Hard for software to detect internally
- Probabilistic, not deterministic

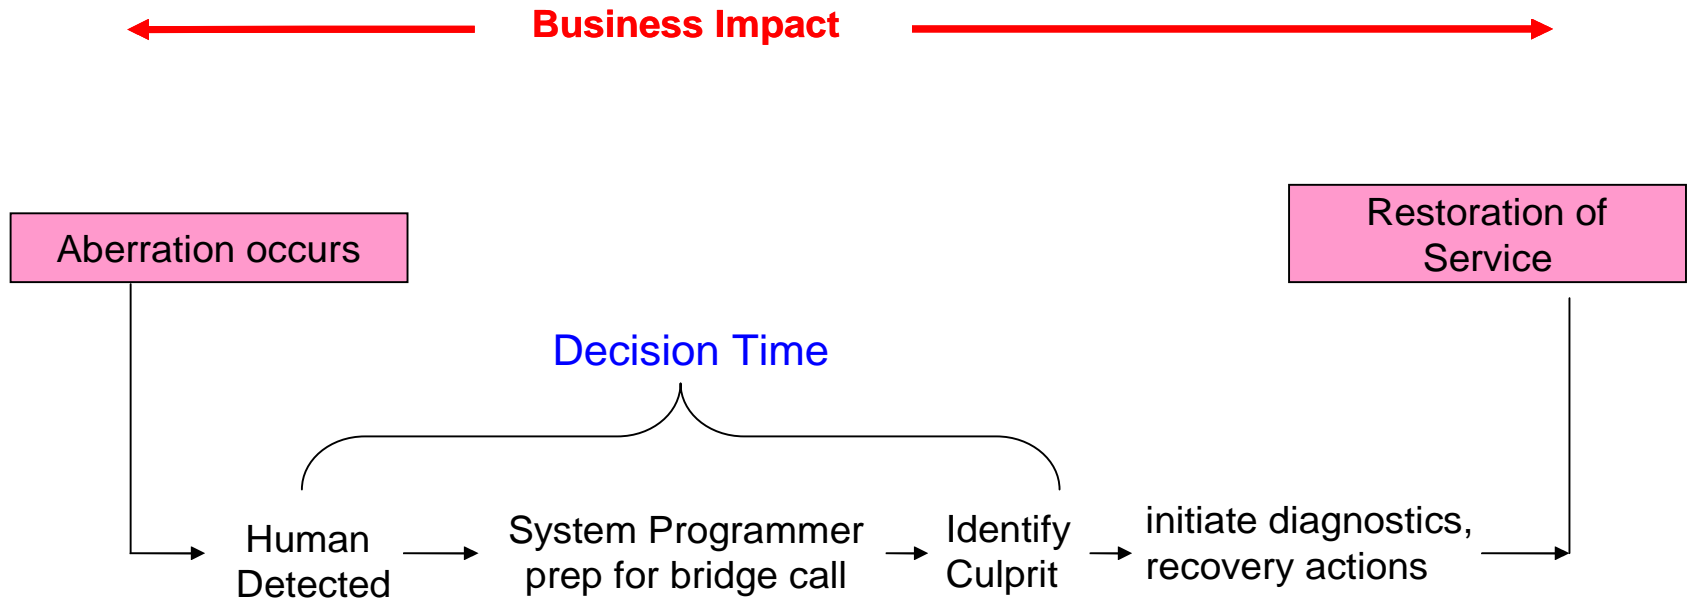
Manifested as

- Stalled / hung processes
 - Single system, Sysplex members
 - Sympathy sickness
- Resource contention
- Storage growth
- CF, CDS growth
- I/O issues (paths, response time)
- Repetitive errors
- Queue growth
- Configuration
 - Single point of failure, thresholds, cache structure size, not enabling new features

Dealing with soft failure problems



Anatomy of an Outage

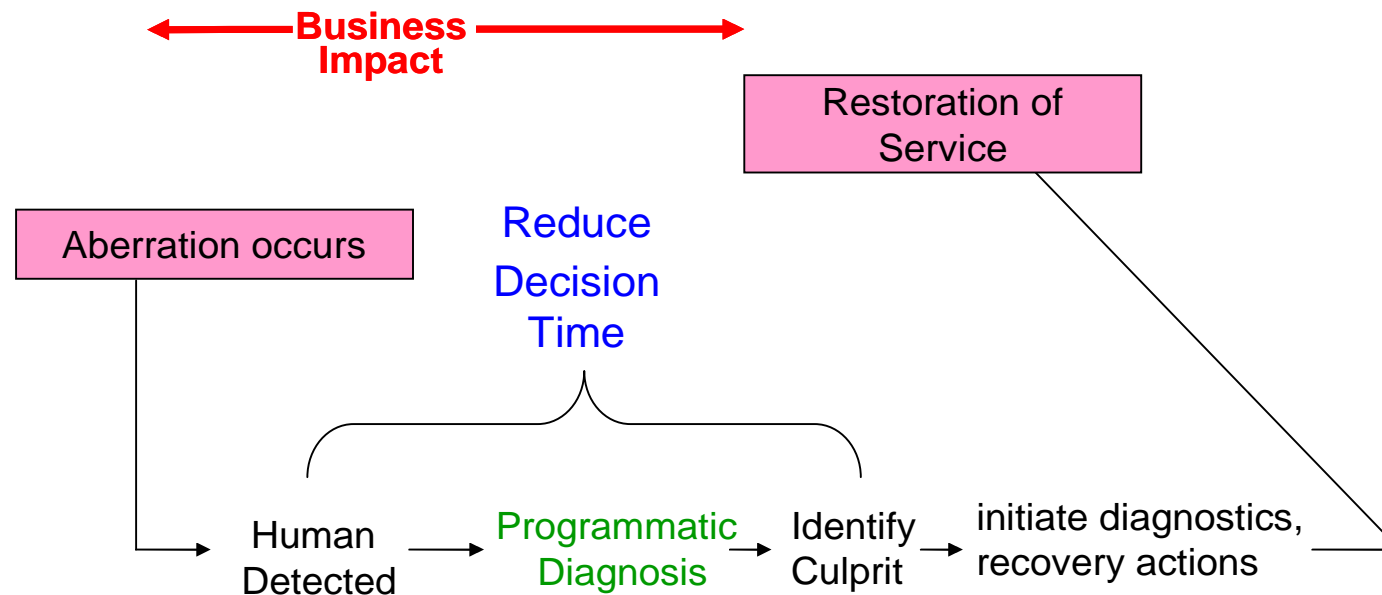


Customer Pain Points:

- *Fault occurs long before anyone notices*
- *Difficult to identify where the problem is coming from*
→ *Leads to long decision time before recovery actions*



Anatomy of an Outage – One step further



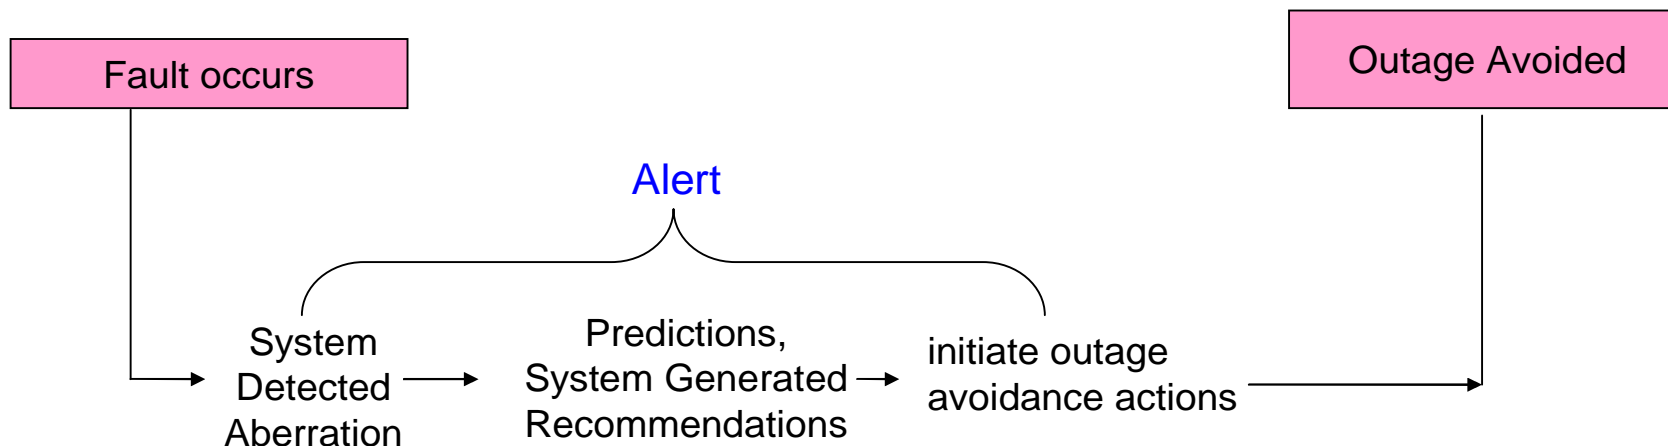
RAS Innovation --- Run Time Diagnostics:

- *Machine-speed understanding*
- *Better tooling to identify the culprit*
- *Enables faster / correct recovery actions*



Anatomy of an Outage – Objective

← **No Business Impact** →



RAS Innovations: Predictive Technologies


- *Machine Learning - Convert diagnostic data to knowledge in real time*
- *Convert soft failures to correctable incidents*



Soft failure issues and solutions

Issues

Risk to business <ul style="list-style-type: none">•The impact of the symptoms•Risk of recurrence•Impact in getting system stabilized
Complexity of performing the task
Troubleshooting a live system and recovering from an apparent failure
Data collection is very time consuming
Significant skill level required to diagnose and analyze problem

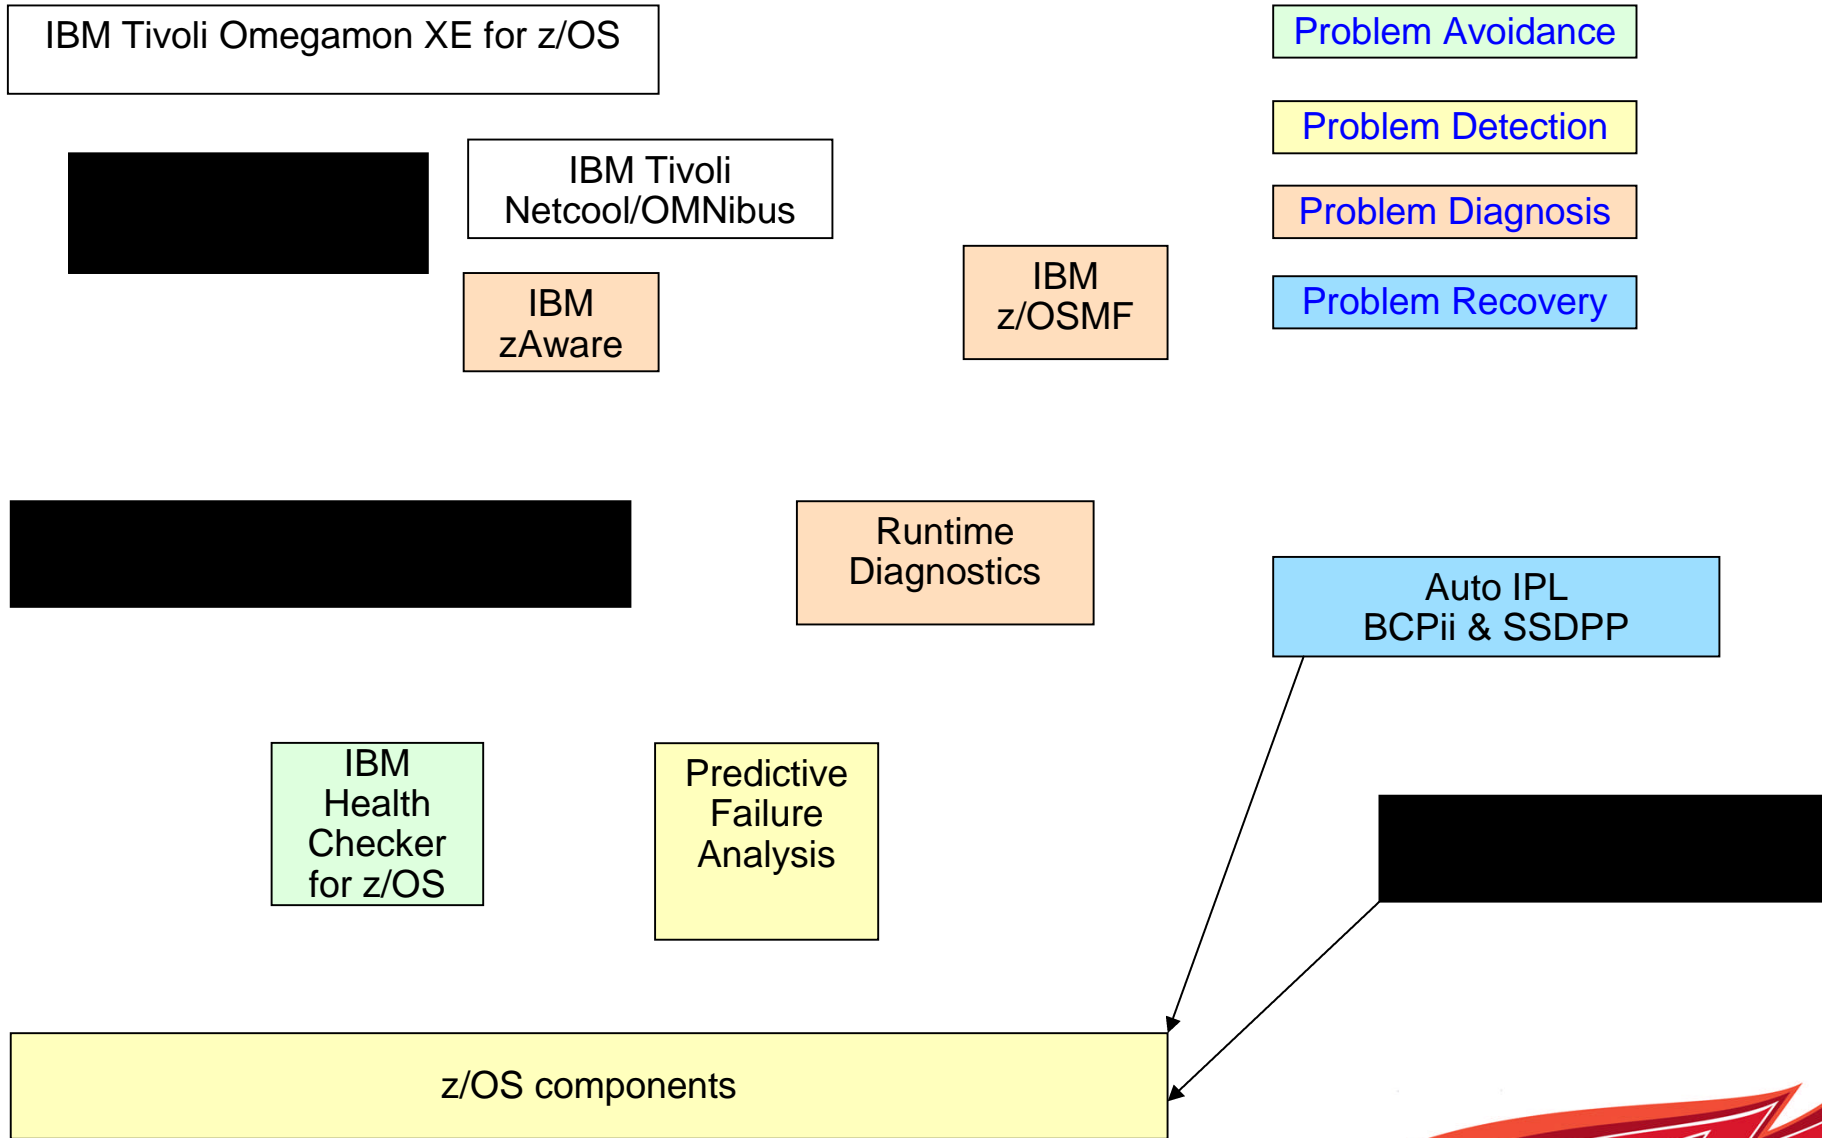


Solutions

Prevention <ul style="list-style-type: none">•provide policy based tools to prevent predictable failures
Detect/Alert <ul style="list-style-type: none">•identify “sick, but not dead” or possible conditions that could lead to larger issues
Diagnosis <ul style="list-style-type: none">•better real time diagnosis and diagnostic tools
Recovery <ul style="list-style-type: none">•improve mean time to recovery
Diagnostic data capture <ul style="list-style-type: none">•make data capture easier and less time consuming




Problem Avoidance, Detection, Diagnosis and Recovery



IBM z/OS Solutions Address Problem Determination

Solutions Available:		Rules based	Analytics / Statistical model	Examines message traffic	Self Learning	Method
z/OS Health Checker	<ul style="list-style-type: none"> ▪Checks configurations ▪Programmatic, applies to IBM and ISV tools ▪Can escalate notifications 	✓				Rules based to screen for conditions
z/OS PFA	<ul style="list-style-type: none"> ▪Trending analysis of z/OS system resources, and performance ▪Can invoke z/OS RTD 		✓		✓	Early detection
z/OS RTD	<ul style="list-style-type: none"> ▪Real time diagnostics of specific z/OS system issues 	✓		✓		After an incident
IBM zAware	<ul style="list-style-type: none"> ▪Pattern based message analysis ▪Self learning ▪Provides aid in diagnosing complex z/OS problems, including cross sysplex, problems that may or may not bring the system down 		✓	✓	✓	Diagnosis Useful before or after an incident

- IBM zAware Uniquely analyzes messages in context to determine unusual behaviors
- IBM zAware Uniquely understands and tunes its baseline to compare against your current activity
- IBM zAware does not depend on other solutions, manual coding of rules, and is always enabled to watch your system





Resiliency offering on System z


	Make sure system is likely to work	Find cause of event after event was reported	Report "first" occurrence of event (before externally visible)
Rules Based Performance	Capacity planning – RMF ¹	OMEGAMON [®] XE	OMEGAMON XE
Rules Based Non Performance	Health Checker for z/OS	RTD	NetView [®] / TSA
Analytical / Statistically Based Performance	ITM 6.2.1	Netcool [®] Tivoli Performance Analyzer	ITM 6.2.1
Analytical / Statistically Based Non Performance	IBM zAware ²	IBM zAware	PFA – control charts IBM zAware – pattern analysis

¹ RMF collects the data for customer analysis / customer rules

² Changes



IBM zAware



IBM zAware – IBM System z Advanced Workload Analysis Reporter

IBM System z: Mainframe focus

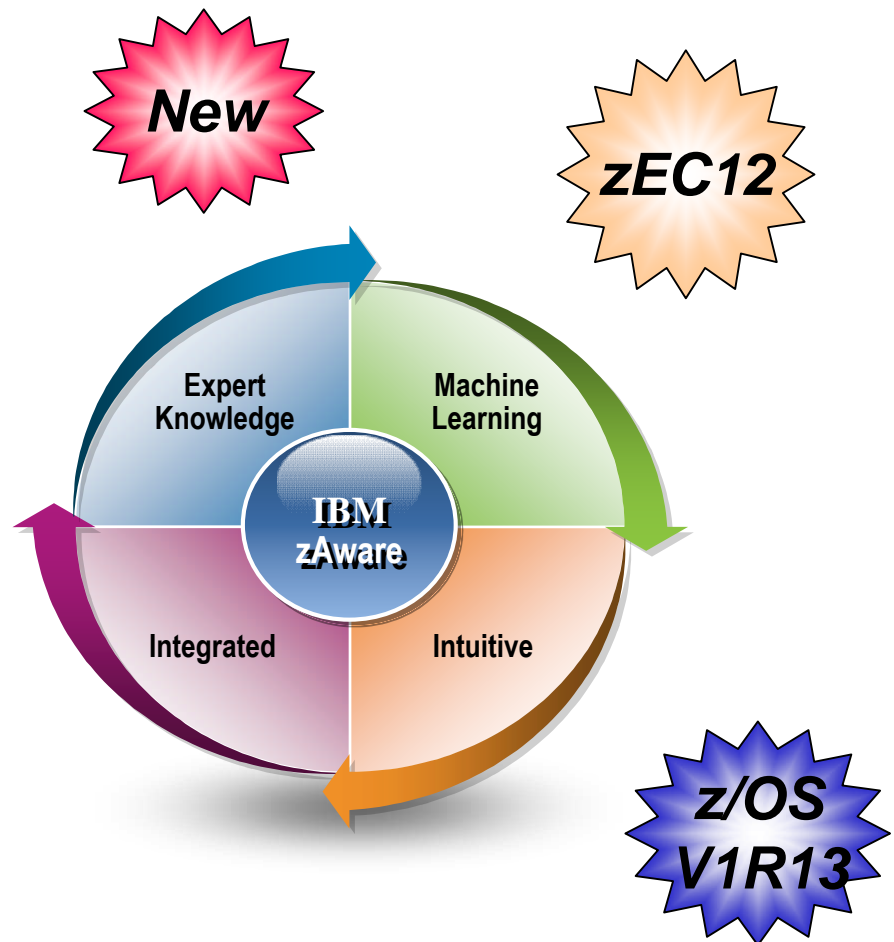
- Server runs on zEnterprise EC12
- Clients are z/OS V1R13

Advanced: Designed by smart people who understand concepts such as Poisson Distribution

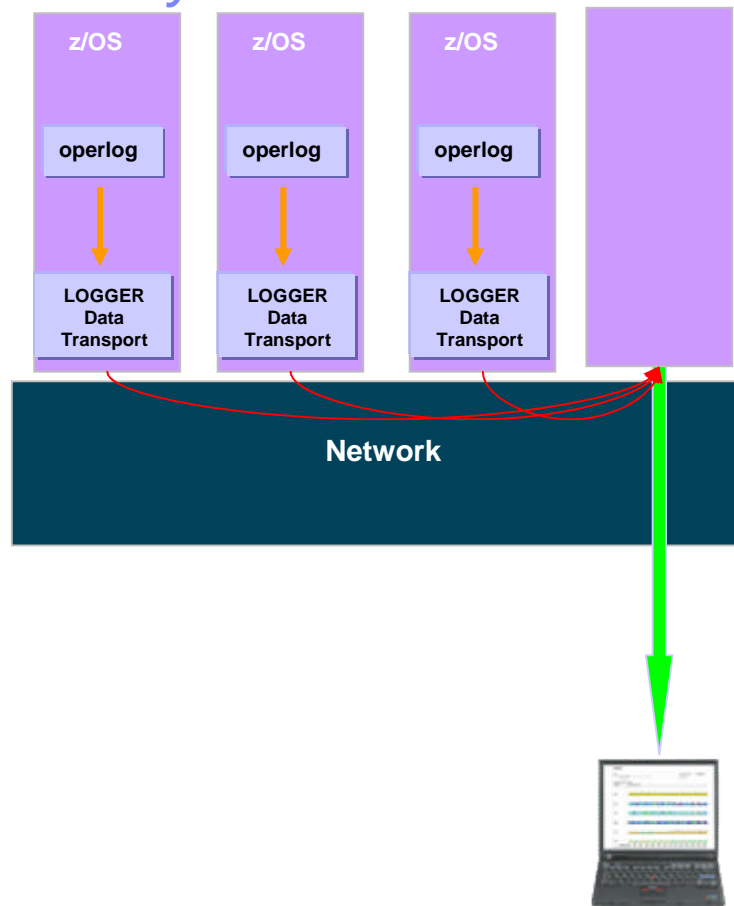
Workload: Data processed is OPERLOG

Analysis: looks at the behavior of messages and message patterns occurring now compared to a model of past behavior

Reporter: message analysis is presented in a graphical view



IBM zAware – IBM System z Advanced Workload Analysis Reporter

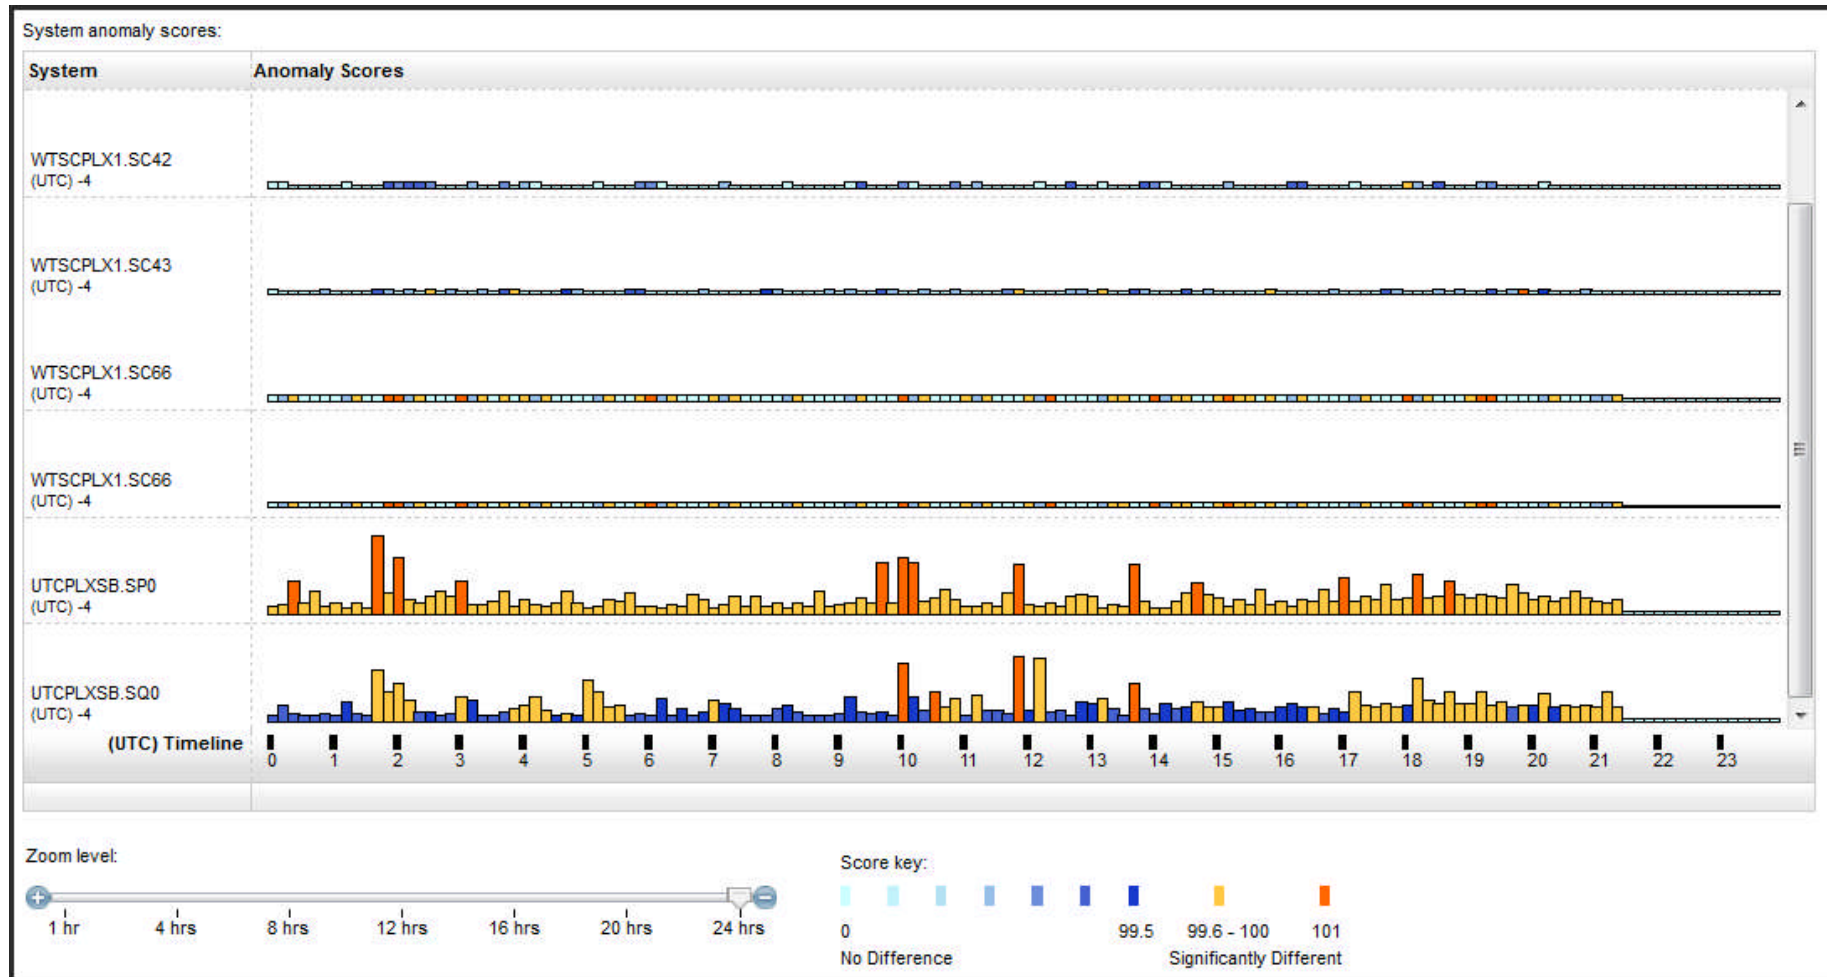


- Monitors z/OS OPERLOG messages including all z/OS console message, ISV and application generated messages
- Can monitor across a sysplex
- Uses a model of past system behavior to compare current message patterns
- Color-coded, browser-based (IE 8, Firefox)
- Visual indicators
 - Number of unique message ids for an interval
 - Anomaly score for an interval
- Detects anomalies monitoring systems miss:
 - Messages may be suppressed or rare
 - Messages may indicate a trend
- XML Output consumable through published API, can drive ISV products


Analysis

- After a model for the system has been successfully created via training, analysis on real time message traffic can occur.
- This step takes the current messages in real-time and uploads them in a manner similar to the initial upload. Every two minutes, this step performs comparisons with the model and updates the xml and the IBM zAware GUI.
- The comparisons use scoring to detect the following: unique messages (those that were not seen in the data that was used in training), rare messages, anomalous message patterns (such as finding a message out of context of a known pattern), and messages that are occurring more frequently than expected.
- The score is also calculated based on the criticality of messages which may give the message more weight if a message is deemed critical or may reduce the score if a message is deemed inconsequential.
- Messages are retrieved from the system logger prior to being excluded due to message flooding automation.
- A multi-line WTO message is counted as one message.

IBM zAware Analysis Panel



IBM zAware Analysis system timeline




Sysplex & LPAR


24 hour timeline

Visual Indicator are bars


Each bar represents a 10 minute interval
Height represents number of unique messages seen during the interval



Color represents anomaly score based on previous history - keyed



IBM zAware Interval View



Drill down to Interval View by clicking on bar

Interval View for System SQ0

The Messages table provides detailed analysis information for each message that occurred during the indicated time interval. To view message details for other intervals use the date and time interval selectors. Click the **Return to Analysis** button to go back to the Analysis view.

Date: Analysis Source: All Monitored Systems

Time interval (UTC): Interval anomaly score: 101.0

Messages

Anomaly Score	Interval Contribution Score	Message Context	Rules Status	Appearance Count	Time Line	Message ID	Message Example	Rarity Score	Component	Cluster ID
1	26.017	new	None	34		AOF313I	06:00:32 : START FOR SUBSYSTEM CICSAA01 (JOB CICSAAQ1) WAS NOT	101	AOF	-1
1	13.076	new	None	1		DFHIR379I	C13WUIQ Unable to start interregion communication because ISC=NO has been	101	DFHIR	-1
1	13.076	new	None	1		DFHPA1910	C13WUIQ SIT OVERRIDE AUTORESTTIME= IS NOT RECOGNIZED. OVERRIDE IS	101	DFHPA	-1
1	13.076	new	None	1		DFHPA1916	C13WUIQ SIT OVERRIDE DATA 47185920 IS OUT OF RANGE FOR KEYWORD EDSALIM= .	101	DFHPA	-1
1	13.076	new	None	1		EYUVS0001I	C13WUIQ CICS PLEX SM WEB USER INTERFACE INITIALIZATION STARTED.	101	EYUVS	-1



IBM zAware Interval View Message line

Messages									
Actions ▾									
▼1 Anomaly Score	Interval ▼2 Contribution Score	Message Context	Rules Status	Appearance Count	Time Line	Message ID	Message Example	Rarity Score	Component
1	26.017	new	None	34		AOF313I	06:00:32 : START FOR SUBSYSTEM CIC SAA01 (JOB CIC SAAQ1) WAS NOT	101	AOF
1	13.076	new	None	1		DFHIR379I	CI3WUIQ Unable to start interregion communication because ISC=NO has been	101	DFHIR

Analytics Information

Number of time msg Appeared in interval

Visual indicator of when Message appeared in interval

Msg ID & Link to msg manual

Msg sample

Rarity score



Analysis

- **Anomaly Score** indicates the rarity of this specific message ID within the selected interval. Higher scores indicate greater anomaly so messages with high anomaly scores are more likely to indicate a problem.
- **Interval Contribution** score Indicates the relative contribution of this message to the anomaly score for the 10-minute interval.
- **Rarity Score** indicates how rarely this message was issued within the selected 10-minute interval.
- **Appearance Count** specifies the number of times that this message was issued within the selected 10-minute interval.
- **Cluster ID** provides the identifier of the cluster to which this message belongs. When the message is not part of a recognized cluster, the cluster ID is -1.
 - Clusters are message patterns (messages appearing together as part of some system event) which define the normal context for messages. Identification of Clusters is performed when a model is created.

Analysis

Message Context indicates whether or not this message is part of an expected pattern of messages (cluster)

- **New**
 - IBM zAware has not previously detected this message in the client model.
- **Unclustered**
 - This message is not part of a defined cluster.
- **In context**
 - This message was issued as expected, within a cluster to which this message belongs.
- **Out of context**
 - This message is expected to be issued as part of a specific cluster but was issued in a different context.

Analysis

Rules Status indicates whether or not the IBM zAware server uses an IBM rule to determine the message anomaly score

- **Critical**
 - critical for diagnosing a potential system problem. For example, message IXC101I, which indicates that a system is being removed from a sysplex.
- **Important**
 - likely to indicate a problem. For example, message IEA911E, which indicates that an SVC dump has been taken.
- **Interesting**
 - indicative of a diagnostically useful event, such as a health check exception.
- **None**
 - No IBM rule is applied for this message.
- **Not interesting**
 - routine, with little or no diagnostic value. These messages are usually issued frequently and at random.

Analysis

Interval anomaly score (remember – based on analysis against model)

- 0 through 99.4
 - messages and message clusters that match or exhibit relatively insignificant differences in expected behavior
- 99.5
 - some rarely seen, unexpected, or out-of-context messages
- 99.6 – 100
 - rarely seen messages (these messages appear in the model only once or twice), or many messages that are unexpected or issued out of context
- 101
 - contain unique message IDs that the IBM zAware server has not detected previously in the client model, or messages that IBM rules define as critical messages

How can IBM zAware Improve Problem Determination?

- **Identify messages indicating a possible z/OS incident is happening**

- Which image is behaving abnormally?
 - Examines unique messages
 - High score generated by
 - unusual messages or message patterns
- When did this unusual behavior start?
 - For a selected 10 minute interval either the current 10 minute interval or past intervals
 - **Which message ids** are unusual?
 - **How often** did the message occur?
 - **When** did the message start to occur?
- Were similar messages issued in the past?
 - Similar characteristics, Same pattern?

- **After a change has been made**


- Are unusual messages being issued following changes ?
 - New software levels (operating system, middleware , applications)
 - Updated system settings / system configurations

- **When diagnosing the cause of an intermittent problem**

- Are new unusual messages being issued in advance of the problem?
- Are more messages issued then expected?
- Are messages issued out of normal pattern or context?

Vertical bar shows the number of unique messages in a 10 minute interval

Scoring of messages color coded from no difference (blue) to significantly different (orange)



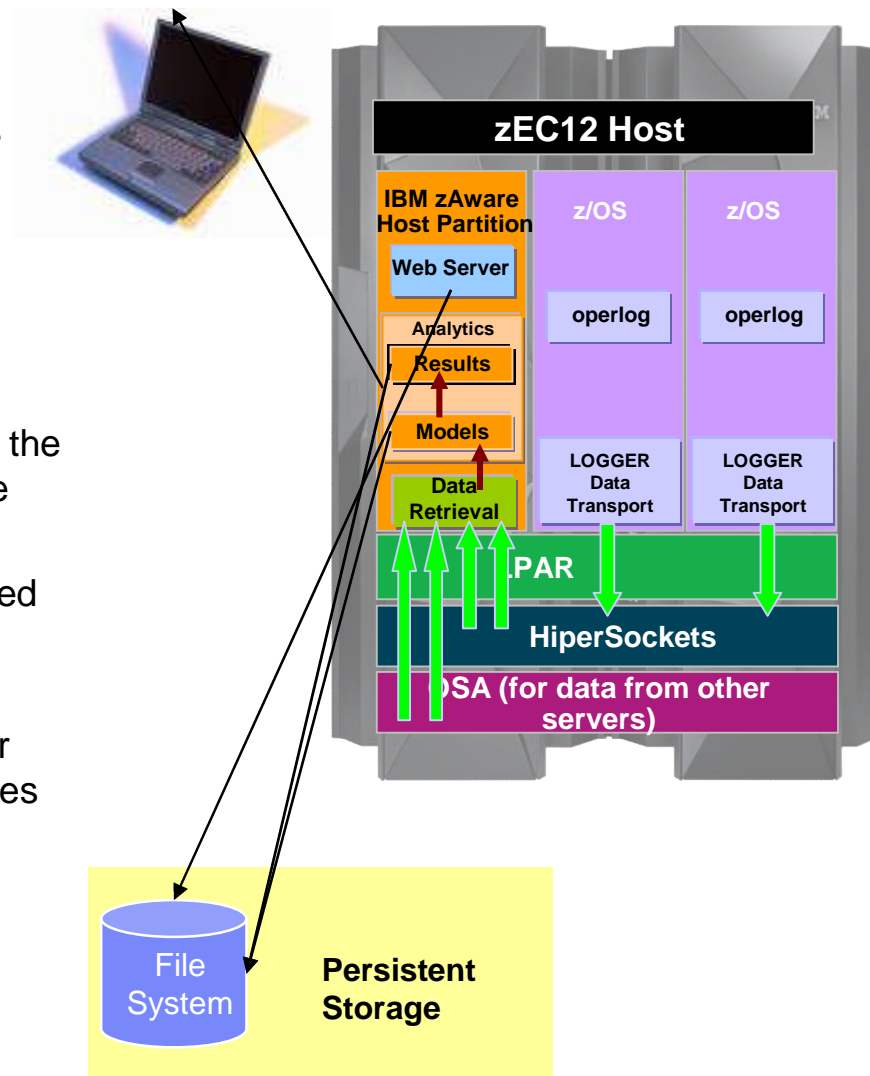
IBM zAware Terminology

- Definitions:
 - The **IBM zAware host system** is the zEC12 that hosts the IBM zAware partition. In most cases, the host server will also have partitions on it that are being monitored. [There may be multiple IBM zAware host partitions on one zEC12, but there will only be one IBM zAware FC 0011 feature \(no additional charge for multiple host partitions\).](#) FC 0101 is also required for monitored clients.
 - An **IBM zAware monitored client** is a z/OS partition that sends OPERLOG files for processing to an IBM zAware partition. The server that has partitions that are IBM zAware monitored clients, must be running z/OS 1.13 with required exploitation software. There may be multiple z/OS partitions (monitored clients) on the server.
 - The **IBM zAware environment** is the collection of the IBM zAware host system and the IBM zAware monitored clients that are sending information to the IBM zAware host system.
 - A **Disaster Recovery (DR) IBM zAware server** is a zEC12 with no-charge firmware to run IBM zAware in a disaster situation.



- Speak with your IBM Representative for ordering details

A closer look inside IBM zAware

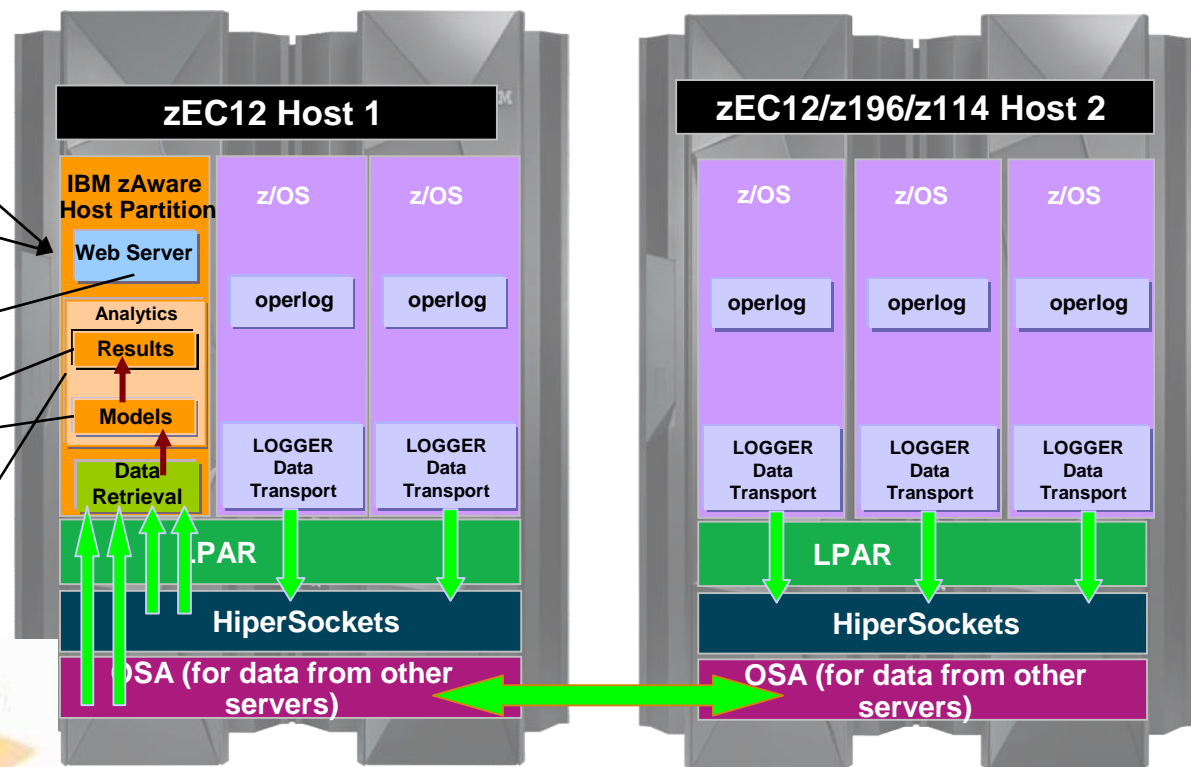
- zEC12 to host IBM zAware Server
 - IBM zAware requires it's own LPAR and runs it's own self-contained firmware stack.
 - This will reduce the number of LPARs available for customer use
 - IBM zAware processor resources can be IFL or General Purpose CP (shared or dedicated)
 - Memory and DASD resources are dependent on the number of monitored clients, amount of message traffic, length of time data retained
 - Memory - Min 4 GB + 256 MB per connected client.
 - DASD ~ 250-500 GB (ECKD)
 - IBM zAware uses Logical Volume Manager (LVM) to aggregate multiple physical devices into a single logical device
 - Network: HiperSockets or OSA ports – for both gathering of instrumentation data, and outbound alerting/communications
 - Need dedicated IP address for partition



Overview



Manage IBM zAware Firmware partition (similar to CF)



IBM zAware itself uses an LPAR. This will reduce the number of LPARs available for customer use

Note: z/OS 1.13 plus PTFs or higher for monitored client



Logical Architecture of IBM zAware

Analytics Engine: performs analysis of real time data and training which creates models of system behavior

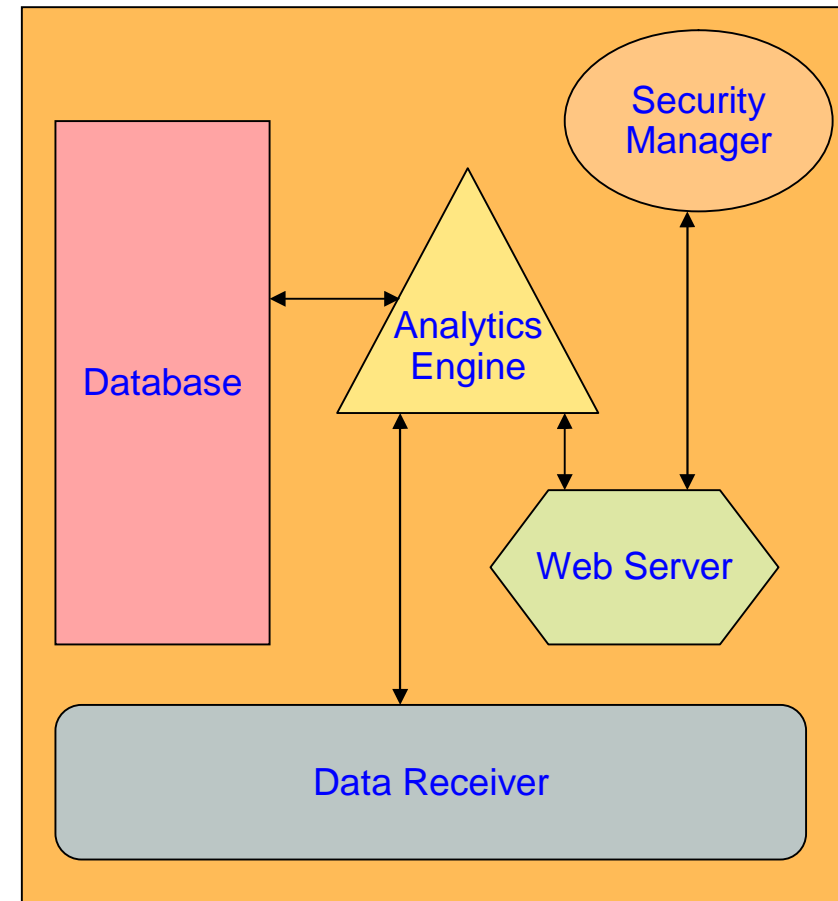
Database: stores models, data for analysis, and results of analysis for retrieval or display

Web Server: provides the interface between user (browser) and IBM zAware. Serves the web pages used to view analysis and for configuration. Provides API.

Security Manager: responsible for controlling access to IBM zAware application. Manages:

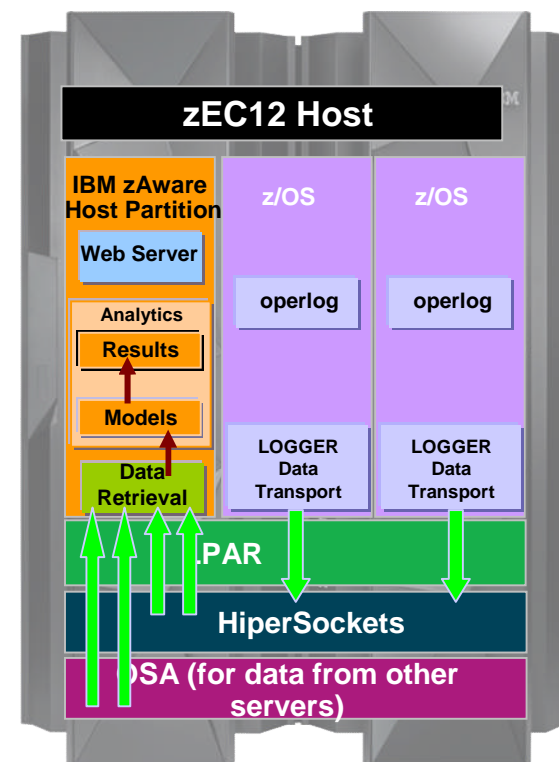
- Secure Socket Layer (SSL)
- Lightweight Directory Access Protocol(LDAP)
- Websphere Local Repository using the Integrated Solutions Console

Data Receiver: manages the TCP/IP connections between monitored clients and IBM zAware. It manages receipt of data from monitored clients.



IBM zAware Monitored Client

- IBM zAware monitored clients can be on any System z Server running z/OS 1.13 + PTFs
- IBM zEnterprise EC12 (zEC12), IBM zEnterprise 196 (z196), IBM zEnterprise 114 (z114), etc.
- TCP/IP network connection to IBM zAware server required
- Requires OPERLOG
- Changes to System Logger allow transmittal of OPERLOG logstream to IBM zAware over the TCP/IP network
- Parmlib and System command enhancements
- Suggested 90 days historical syslog or OPERLOG recommended to create initial model



IBM zAware Monitored Client Logical Architecture

Enhancements to System Logger to allow transmittal of OPERLOG data to IBM zAware.

New OPERLOG Log Stream parameters

- ZAI(YES)
- ZAIDATA('OPERLOG')

New IXGCNFxx parameters

- ZAI
 - SERVER
 - PORT
 - LOGBUFMAX
 - LOGBUFWARN
 - LOGBUFFULL

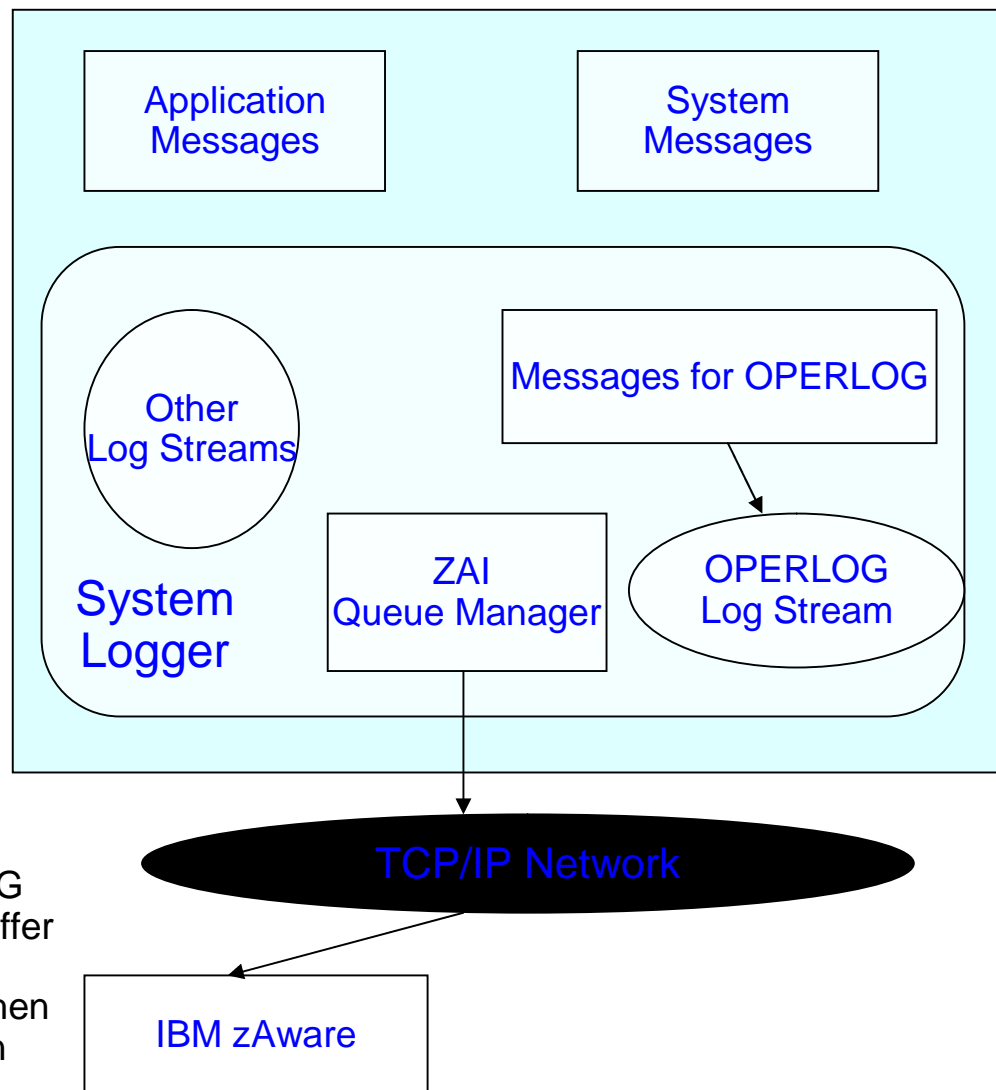
Enhancements to DISPLAY LOGGER command

- DISPLAY LOGGER,STATUS,ZAI
- DISPLAY LOGGER STATUS,ZAI,VERIFY

Enhancements to SETLOGR command

- SETLOGR FORCE,ZAICONNECT,LSN=xx
- SETLOGR FORCE,ZAIQUIESCE,LSN=xx

When a message is to be written to the OPERLOG log stream, the message is copied to a special buffer in System Logger (reserved for use by messages that are destined for IBM zAware). The buffer is then dequeued and sent to the IBM zAware application across a TCP/IP connection.



IBM zAware planning and configuration

References

- IBM Product Documentation

IBM System z Advanced Workload Analysis Reporter (IBM zAware) Guide,
SC27-2623

- IBM ITSO Redbook

Extending z/OS System Management Functions with IBM zAware,
SG24-8070

Planning considerations

- For the IBM zAware Partition
 - Requires an IBM zEnterprise EC12 with enablement and connections feature code
 - CP capacity, Memory, Network connectivity, DASD storage
 - Define Partition using HCD
 - Activate IOCDs with IBM zAware partition
 - Create/Configure IBM zAware activity (image) profile
 - Activate Partition
 - Log in to IBM zAware GUI and configure server

- For the IBM zAware Monitored Clients
 - Establish Network connectivity to IBM zAware partition
 - Define security profiles
 - Update OPERLOG log stream
 - Update IXGCNFxx and make the changes active
 - Verify connection between client and server

IBM zAware partition requirements

- IBM zAware server only runs on an IBM zEnterprise EC12
 - Feature codes 0011 (IBM zAware host system enablement feature)
 - Feature code 0101 (for monitored clients)

- A shared or dedicated Open Systems Adapter (OSA) port, with an IP address that is either dedicated or assigned through Dynamic Host Connection Protocol (DHCP).
 - OSA, HiperSockets, IEDN (intraensemble data network)
 - Consider access to zAware server GUI
 - Consider connection to zAware server from monitored clients
 - Redundancy recommended

- Shared or dedicated Integrated facilities for Linux (IFLs) or central processors (CPs).

- Storage and memory resources that are sufficient to support the IBM zAware server that runs on the partition and the z/OS® clients that the server monitors.
 - 4-6GB + 256 per monitored client
 - 500GB DASD + 4-5GB for additional monitored clients

IBM zAware client requirements

- z/OS V1R13 with ptfs:
 - The PTF for system logger APAR OA38747 and its prerequisite APAR OA38613
 - The PTF for z/OS bulk load client for IBM zAware APAR OA39256

- The system must be using the operations log (OPERLOG) as the hardcopy medium

- The system must be configured as a single-system sysplex (monoplex), a system in a multisystem sysplex, or a member of a Parallel Sysplex

Definitions for IBM zAware

Component	Definitions	Where do you define?
IOCDS	<ul style="list-style-type: none"> ▪ Define IBM zAware LPAR ▪ Add Network devices ▪ Add Storage devices 	HCD
Activation Profile	<ul style="list-style-type: none"> ▪ Configure Image Profile 	HMC
IBM zAware Server	<ul style="list-style-type: none"> ▪ Add devices ▪ Configure Security 	IBM zAware GUI
OPERLOG	<ul style="list-style-type: none"> ▪ OPERLOG log stream definition changes 	Monitored z/OS system
System Logger	<ul style="list-style-type: none"> ▪ IXGCNF changes ▪ IEASYSxx updates ▪ RACF definitions 	
TCP/IP	<ul style="list-style-type: none"> ▪ Profile updates ▪ Maybe Hosts file 	
Create Model	<ul style="list-style-type: none"> ▪ Bulk Load ▪ Topology assignment ▪ Train 	Monitored z/OS system IBM zAware GUI



Define the IBM zAware LPAR

- Add the partition
 - Name = as per your convention
 - Number = as available
 - Usage = OS

```
Specify or revise the following values.
```

```
Partition name . . . A1A  
Partition number . . A      (same as MIF image ID)  
Partition usage . . . OS      +  
  
Description . . . . zAware2
```

Define the zAware LPAR

```

----- Partition List -----
Goto Backup Query Help
-----
Command ==> _____ Scroll ==> PAGE
Row 1 of 15

Select one or more partitions, then press Enter. To add, use F11.

Processor ID . . . . : SCZP401 Helix
Configuration mode . : LPAR
Channel Subsystem ID : 1

/ Partition Name      Number Usage + Description
- A1A                 A      OS      zAware2
- A1B                 B      CF/OS   CHPID holder
- A1C                 C      OS      zAware1
- A1D                 D      CF      Trainer FACIL06
- A1E                 E      CF      COMPLEX CF38
- A1F                 F      CF      COMPLEX CF39
- A11                 1      OS      COMPLEX SC30

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F7=Backward  F8=Forward    F9=Swap      F10=Actions    F11=Add
F12=Cancel
    
```



Add channels

```

Command ==> _____ Scroll ==> PAGE

Processor ID . . . : SCZP401   Helix
Configuration mode : LPAR
Partition name . . : A1A      zAware2

ENTER to continue.

CHPID  Type  Mode  Description
00    OSD  SPAN  Exp4S 1KBaseT All LPARs 9.12.4.x
06     OSD  SPAN  Exp4S 1KBaseT
40     FC   SPAN  Exp8S 10Km LX
41     FC   SPAN  Exp8S 10Km LX
42     FC   SPAN  Exp8S 10Km LX
43     FC   SPAN  Exp8S 10Km LX
44     FC   SPAN  Exp8S 10Km LX
45     FC   SPAN  Exp8S 10Km LX
46     FC   SPAN  Exp8S 10Km LX
47     FC   SPAN  Exp8S 10Km LX
F4     IQD  SPAN  zAware
    
```





Add control units

```
Processor ID . . . : SCZP401   Helix  
Configuration mode : LPAR  
Partition name . . : A1A      zAware2
```

ENTER to continue.

Control Unit		-----Channel Path ID . Link Address-----							
Num	Type	1-----	2-----	3-----	4-----	5-----	6-----	7-----	8-----
0063	2032	46.FE							
0064	2032	47.FE							
2040	OSA	00							
7A00	IQD	F4							
9000	2107	40.1B	41.1B	42.10	43.10				
9080	2107	44.25	45.25	46.2D	47.2D				
9100	2107	40.1B	41.1B	42.10	43.10				
9180	2107	44.25	45.25	46.2D	47.2D				
9200	2107	40.1B	41.1B	42.10	43.10				





Add Devices

```
Processor ID . . . : SCZP401   Helix
Configuration mode : LPAR
Partition name . . : A1A      zAware2
```

ENTER to continue.

-----Device-----		Access	-----Attached Control Units-----								
Number	Type	Allowed	1---	2---	3---	4---	5---	6---	7---	8---	UA
0063	2032	Yes	0063								00
0064	2032	Yes	0064								00
2040, 15	OSA	Yes	2040								00
204F	OSAD	Yes	2040								FE
7A00, 8	IQD	Yes	7A00								00
9000, 16	3390B	Yes	9000								00
9000, 128	3390A	Yes	9000								80
9010, 112	3390B	Yes	9000								10
9080, 16	3390B	Yes	9080								00
9080, 128	3390A	Yes	9080								80
9090, 16	3390B	Yes	9080								10



Tips

- If you noticed, we did not use an explicit device candidate list to segregate the DASD that will be used for the IBM zAware LPAR
 - Not a good idea – you will see why in a moment
- We suggest you use an explicit device candidate list to:
 - Allow the IBM zAware LPAR access to ONLY the DASD assigned for it
 - Disallow access to the IBM zAware DASD by other LPARs
 - If you plan on backing up zAware DASD then ALSO allow the LPAR which will be doing the backups access the devices, but make sure they are offline at IPL time to the LPAR doing the backups
- IBM zAware LPARs do not use an OSCONFIG (part of IODF to define what devices an LPAR can 'see')
- IBM zAware will only 'see' devices defined to partition as part of IOCDS

Next Steps

- Once the IOCDS is activated you can configure the LPARs image profile and activate the IBM zAware partition

- Log into HMC

- Select IBM zAware Partition

- Customize Activity Profile
 - Select Default to modify

- Activate Partition

Login to HMC and select LPAR for customization

Ensemble Management > ITSO Ensemble > Members > SCZP401

Virtual Servers | Hypervisors | Topology

Filter: [] Tasks: [] Views: Images []

Select	Name	Status	Activation Profile	Last Used Profile	OS Name	OS Type	OS Level
<input type="checkbox"/>	A15	Operating	A15	LBSIPL	SC90	z/OS	V1R13
<input type="checkbox"/>	A16	Operating	A16	TEST9502	SC32	z/OS	V1R13
<input type="checkbox"/>	A17	Operating	A17	VMLINUX4	VMLINUX4	z/VM	6.1.0 - 1102
<input type="checkbox"/>	LINUXMNT	Not activated					
<input type="checkbox"/>	WASB1	Not activated			wasb1	SUSE Linux Enterprise Server	11.1
<input type="checkbox"/>	WASG1	Not activated			wasg1	SUSE Linux Enterprise Server	11.1
<input type="checkbox"/>	WASS1	Not activated			wass1	SUSE Linux Enterprise Server	11.1
<input type="checkbox"/>	A18	Operating	A18	TEST9502	SC33	z/OS	V1R13
<input type="checkbox"/>	A19	Operating	A19	TRAINER13	#@\$A	z/OS	V1R13
<input checked="" type="checkbox"/>	A1A		A1A	A1A			
<input type="checkbox"/>	A1B		A1B				
<input type="checkbox"/>	A1C		A1C	A1C			
<input type="checkbox"/>	A1D						
<input type="checkbox"/>	A1E	Operating					
<input type="checkbox"/>	A1F	Operating					
<input type="checkbox"/>	A21	Not activated	A21				
<input type="checkbox"/>	A2B	Not activated	A2B				
<input type="checkbox"/>	A31	Not activated	A31				

Max Page Size: 500 Total: 36 Filtered: 36 Selected: 1

Tasks: A1A [] [] []

Image Details
Toggle Lock
Daily
Operational Customization

Configure Channel Path On/Off
Customize/Delete Activation Profiles
Logical Processor Add

Customize/Delete Activation Profiles: Customize or delete activation profiles for selected objects - Click to launch

Image Details
Toggle Lock
Daily
Operational Customization

General

The screenshot shows the 'View Image Profiles: SCZP401:A1A : A1A : General' window. The left pane shows a tree view with 'A1A' expanded and 'General' selected. The main pane contains the following fields:

- Profile name: A1A (Assigned for activation)
- Description: zAware2 image profile
- Partition identifier: 1A
- Mode: zAware (selected from a list including ESA/390, ESA/390 TPF, Coupling facility, LINUX only, z/VM)
- Clock Type Assignment:
 - Standard time of day
 - Logical partition time offset
- Ensure that the image profile data conforms to the current maximum LICCC configuration.

Buttons for 'Cancel' and 'Help' are visible at the bottom left of the window.

- Profile Name = As per your conventions
- Description = Identify the LPAR with some text
- Partition identifier = as appropriate to match your HCD partition number
- Mode = new type is added called zAware (for IBM zAware partitions)



Processor

The screenshot shows a web browser window with the URL `https://sczhmcb.itso.ibm.com/hmc/content?taskId=91&refresh=222`. The main content area is titled "View Image Profiles: SCZP401:A1A : A1A : Processor". On the left is a tree view with "A1A" expanded, showing sub-items: General, Processor (highlighted), Security, Storage, Options, and Firmware. The main panel has a "Group Name" dropdown set to "<Not Assigned>". Under "Logical Processor Assignment", there are four radio buttons: "Dedicated central processors", "Dedicated integrated facility for Linux", "Not dedicated central processors" (selected), and "Not dedicated integrated facility for Linux". Under "Not Dedicated Processor Details", there is a text input for "Initial processing weight" with the value "10" and a checkbox for "Initial capping". Below that, there are two text inputs: "Number of processors - Initial" with the value "1" and "Reserved" with the value "0". At the bottom left of the window are "Cancel" and "Help" buttons.

Assign CPs for the LPAR



Security

View Image Profiles: SCZP401:A1A : A1A : Security

- SCZP401:A1A
 - A1A
 - General
 - Processor
 - Security
 - Storage
 - Options
 - Firmware

Partition Security Options

- Global performance data control
- Input/output (I/O) configuration control
- Cross partition authority
- Logical partition isolation

Counter Facility Security Options

- Basic counter set authorization control
- Problem state counter set authorization control
- Crypto activity counter set authorization control
- Extended counter set authorization control

Sampling Facility Security Options

- Basic sampling authorization control
- Diagnosis and basic sampling authorization control

CPACF Key Management Operations

- Permit AES key import functions
- Permit DEA key import functions

Cancel Help

Defaults are fine



Options

View Image Profiles: SCZP401:A1A : A1A : Options

- SCZP401:A1A
 - A1A
 - General
 - Processor
 - Security
 - Storage
 - Options
 - Firmware

Image Options

Minimum input/output (I/O) priority * 0

Maximum input/output (I/O) priority * 0

Defined capacity *0

CP management cluster name

Cancel Help

We took the defaults

Firmware

View Image Profiles: SCZP401:A1A : A1A : Firmware

Host name : ZAWARE2
 Master user ID : FRANK

Network Adapters

Select	CHPID	VLAN	IP address	Mask/Prefix
<input type="radio"/>	0		9.12.5.128	20
<input type="radio"/>	f4		10.1.110.3	24

Default gateway : 9.12.4.1

DNS Servers

Select	IP address
<input type="radio"/>	9.12.6.7

Cancel Help

Define your host name

Master user ID and PW used to login to IBM zAware GUI for initial configuration. ID not case sensitive but PW is case sensitive

Add network adapters with IP addresses and network masks

Assign a default gateway

Assign DNS servers

This will vary depending on you installations network topology



Add OSA

Add/Edit Network Adapters Entry - SCZP401

Select an address type and modify or fill in the details for this CHPID.

IP address type

- DHCP
- Link Local
- Static IPv4 Address
- Static IPv6 Address

Details

CHPID :

VLAN :

IP address :

Mask / Prefix :

OK Cancel Help



Add HiperSocket

https://sczhmcb.itso.ibm.com/hmc/wcl/T3ef6#tableTop_16a90

Add/Edit Network Adapters Entry - SCZP401

Select an address type and modify or fill in the details for this CHPID.

IP address type

- DHCP
- Link Local
- Static IPv4 Address
- Static IPv6 Address

Details

CHPID :

VLAN :

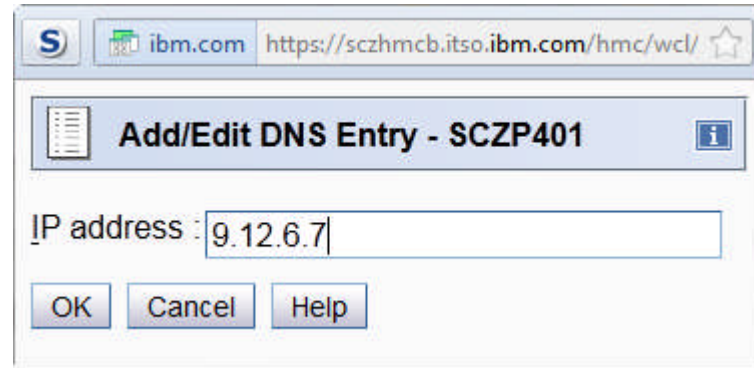
IP address :

Mask / Prefix :

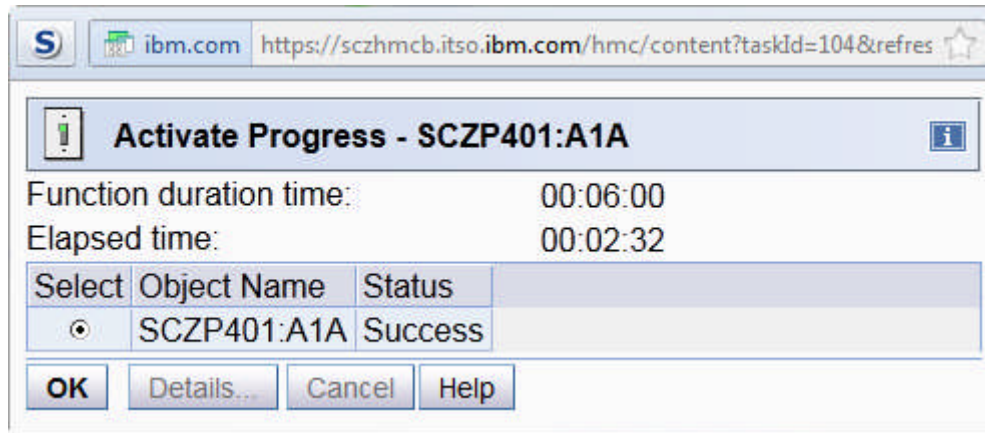
OK Cancel Help



Add DNS entry



Activate partition



After activation completes the zAware application stack is in the process of initializing – this takes a few minutes to complete after 'Activate Success' message.



Ping IBM zAware

```
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\IBM_ADMIN>ping 9.12.5.128

Pinging 9.12.5.128 with 32 bytes of data:
Reply from 9.12.5.128: bytes=32 time=18ms TTL=61
Reply from 9.12.5.128: bytes=32 time=1ms TTL=61
Reply from 9.12.5.128: bytes=32 time=1ms TTL=61
Reply from 9.12.5.128: bytes=32 time<1ms TTL=61

Ping statistics for 9.12.5.128:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 18ms, Average = 5ms

C:\Users\IBM_ADMIN>
```

Simple PING command to verify that IBM zAware network stack is up

IBM zAware initial configuration through GUI

Once the IBM zAware partition is activated several additional steps are required to complete the initial configuration:

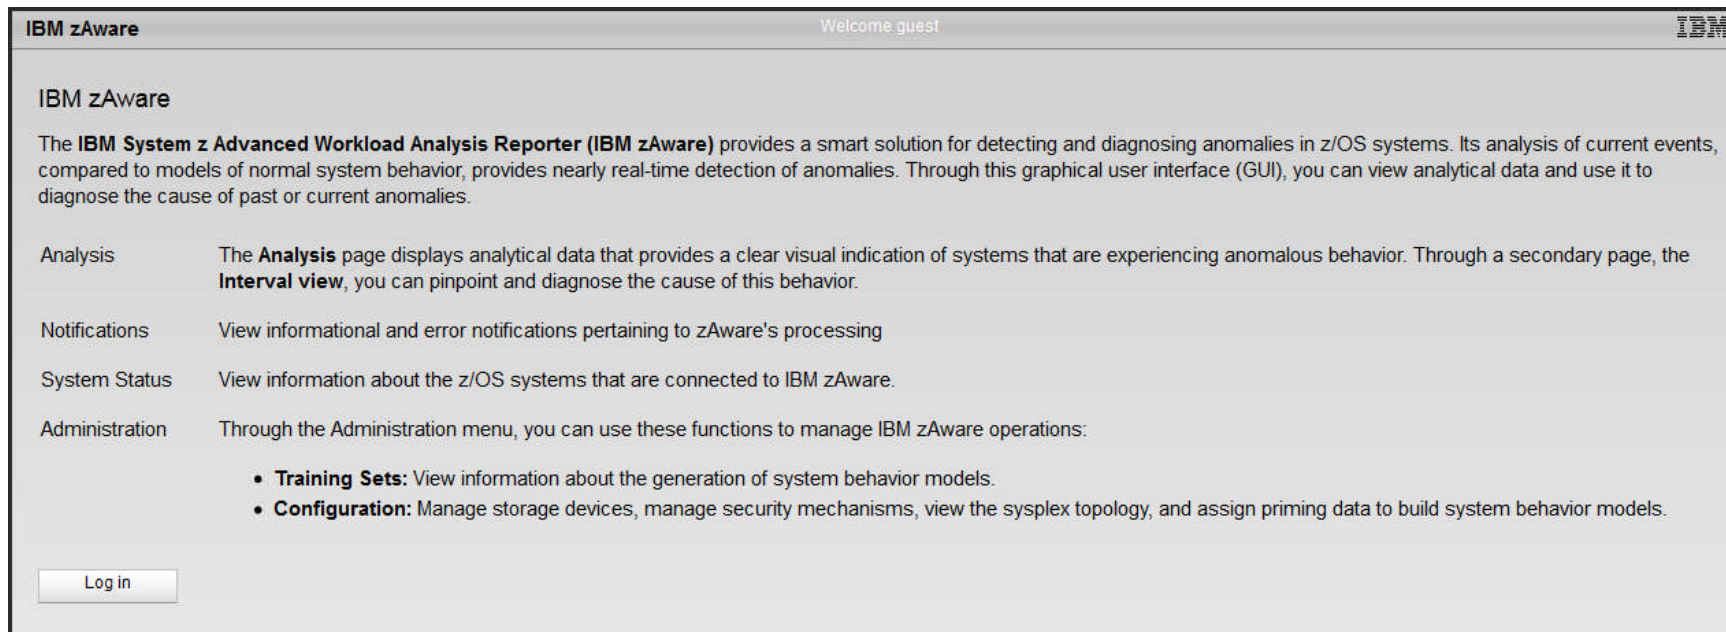

- **Configure Storage**
 - Assign DASD to IBM zAware

- **Configure Security**
 - Change default server certificate (optional)
 - Add users and assign roles

IBM zAware customization through the GUI

- Login to the IBM zAware GUI

- <http://<hostname or ip address of IBM zAware LPAR>/zAware>
 - For example: <http://9.12.5.128/zAware>
- An https connection is established
 - For example <https://9.12.5.128/zAware/Login>



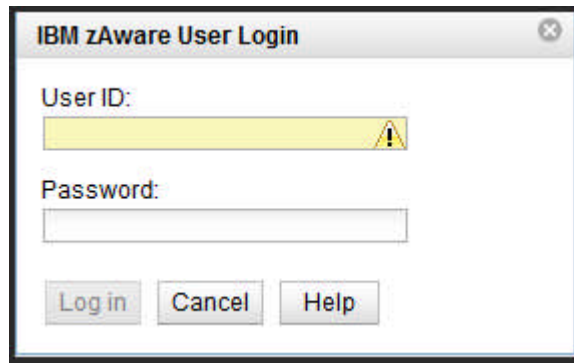
The screenshot shows the IBM zAware web interface. At the top, it says "IBM zAware" and "Welcome guest". Below this, there is a section titled "IBM zAware" with a descriptive paragraph: "The IBM System z Advanced Workload Analysis Reporter (IBM zAware) provides a smart solution for detecting and diagnosing anomalies in z/OS systems. Its analysis of current events, compared to models of normal system behavior, provides nearly real-time detection of anomalies. Through this graphical user interface (GUI), you can view analytical data and use it to diagnose the cause of past or current anomalies."

Below the description, there are four main menu items:

- Analysis:** The **Analysis** page displays analytical data that provides a clear visual indication of systems that are experiencing anomalous behavior. Through a secondary page, the **Interval view**, you can pinpoint and diagnose the cause of this behavior.
- Notifications:** View informational and error notifications pertaining to zAware's processing
- System Status:** View information about the z/OS systems that are connected to IBM zAware.
- Administration:** Through the Administration menu, you can use these functions to manage IBM zAware operations:
 - **Training Sets:** View information about the generation of system behavior models.
 - **Configuration:** Manage storage devices, manage security mechanisms, view the sysplex topology, and assign priming data to build system behavior models.

At the bottom left of the interface, there is a "Log in" button.

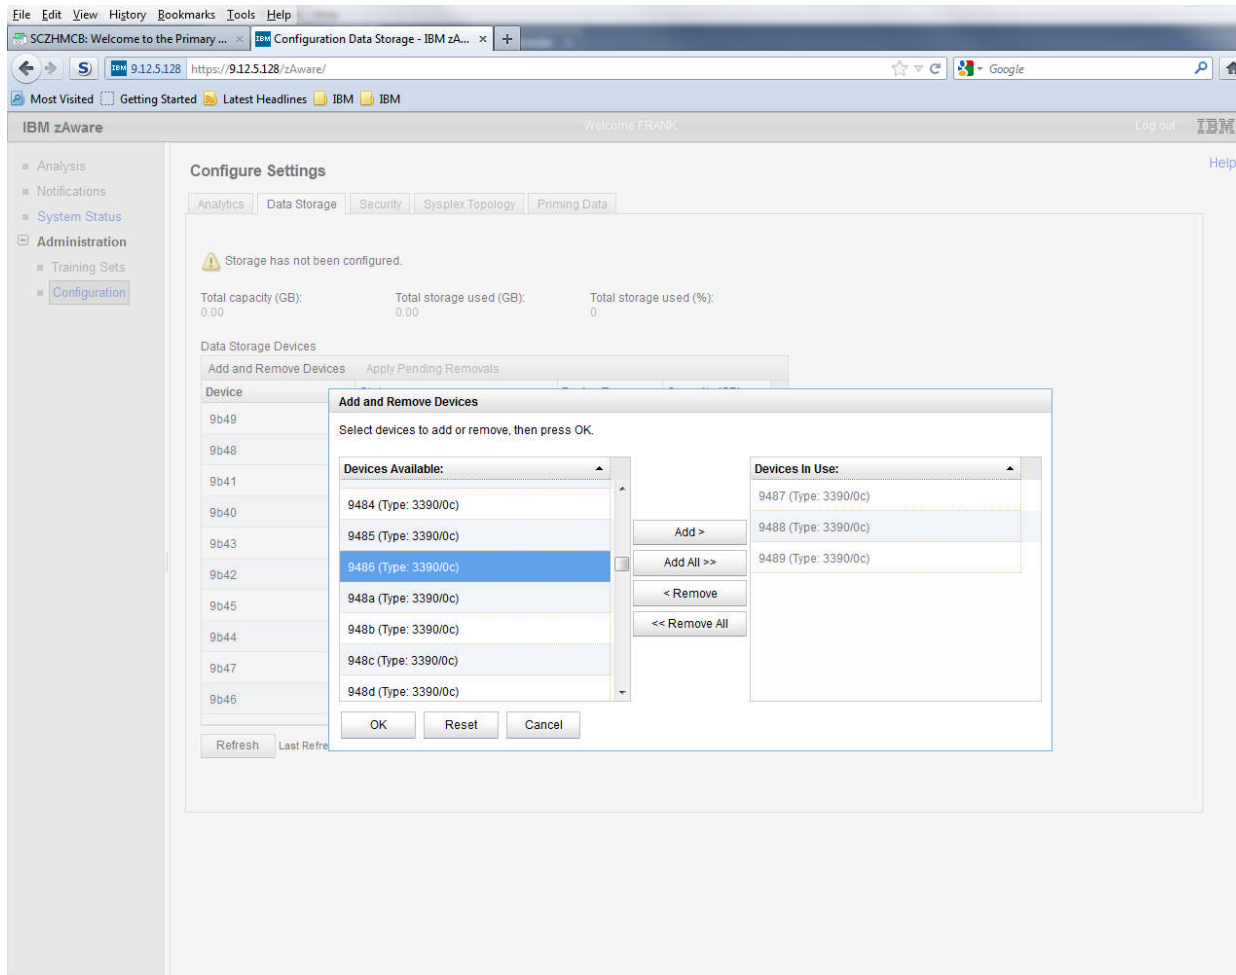
Initial Login to IBM zAware GUI



For the initial login to the IBM zAware GUI use the USER ID and PASSWORD which you defined on the FIRMWARE tab of the partitions image profile



Storage Panel



After the first login, you are automatically brought to the add or remove storage device panel.

For IBM zAware to complete its initial setup you must add storage.

The devices available window shows all devices available through the IOCDS

As you can see we defined the LPAR with shared access to most of our DASD

It would have been more convenient, and less error prone, if we had used an explicit device candidate list

Select the devices for use and add them. Click ok.

Respond to the prompts. The DASD will then be initialized for use by IBM zAware.





Storage Added

The screenshot shows the IBM zAware Configuration Data Storage interface. The 'Data Storage' tab is active, displaying summary statistics and a table of storage devices. The summary statistics are:

- Total capacity (GB): 22.15
- Total storage used (GB): 1.64
- Total storage used (%): 7.38

The 'Data Storage Devices' table is as follows:

Device	Status	Device Type	Capacity (GB)
9489	In Use	3390/0c	7.38
9488	In Use	3390/0c	7.38
9487	In Use	3390/0c	7.38
9b49	Available	3390/0a	—
9b48	Available	3390/0a	—
9b41	Available	3390/0a	—
9b40	Available	3390/0a	—
9b43	Available	3390/0a	—
9b42	Available	3390/0a	—
9b45	Available	3390/0a	—

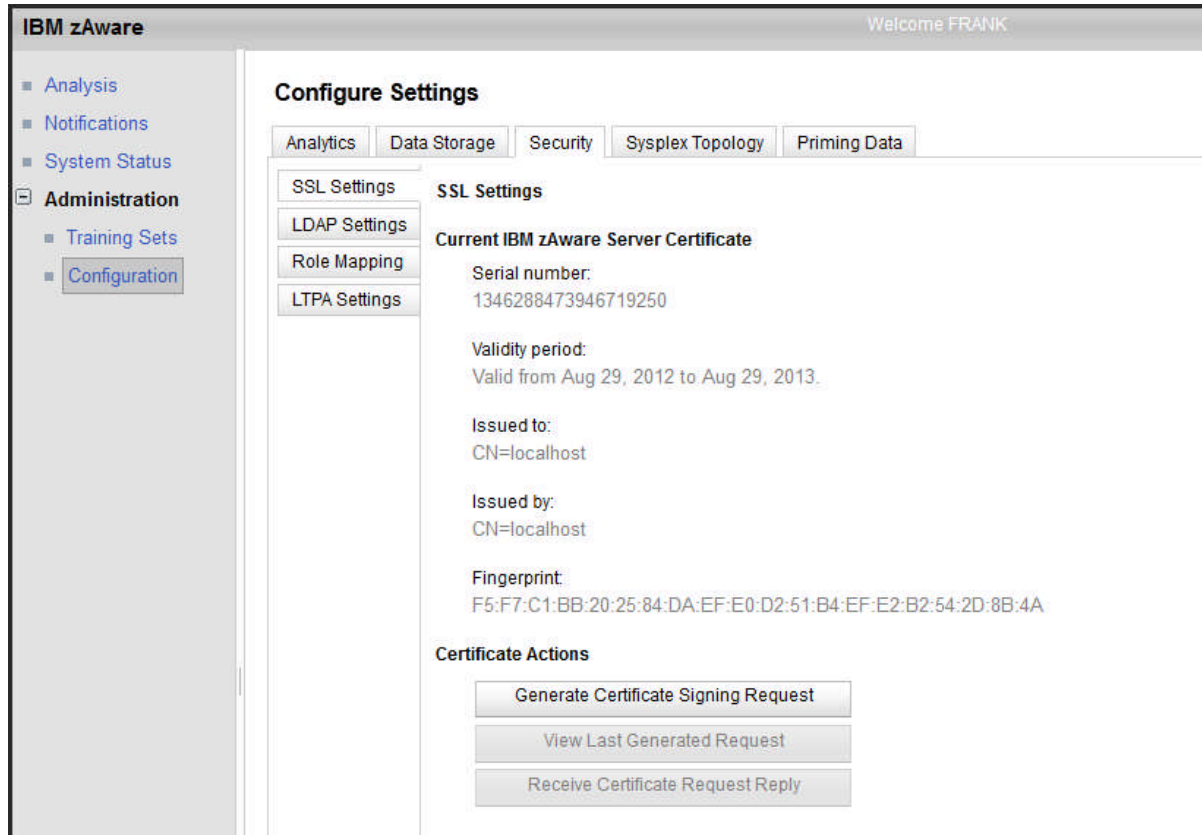
A 'Refresh' button is located below the table, with a 'Last Refresh' timestamp of Tue Sep 04 2012 10:15:50 GMT-0400 (Eastern Daylight Time).

The device status field will be updated to indicate 'in use'.

Various reporting fields for storage use will be updated.



SSL setting (security)



The screenshot shows the IBM zAware web interface. The top navigation bar includes 'Analysis', 'Notifications', 'System Status', and 'Administration'. Under 'Administration', 'Configuration' is selected. The main content area is titled 'Configure Settings' and has tabs for 'Analytics', 'Data Storage', 'Security', 'Sysplex Topology', and 'Priming Data'. The 'Security' tab is active, showing 'SSL Settings'. Under 'SSL Settings', there are sub-tabs for 'LDAP Settings', 'Role Mapping', and 'LTPA Settings'. The 'Current IBM zAware Server Certificate' section displays the following information:

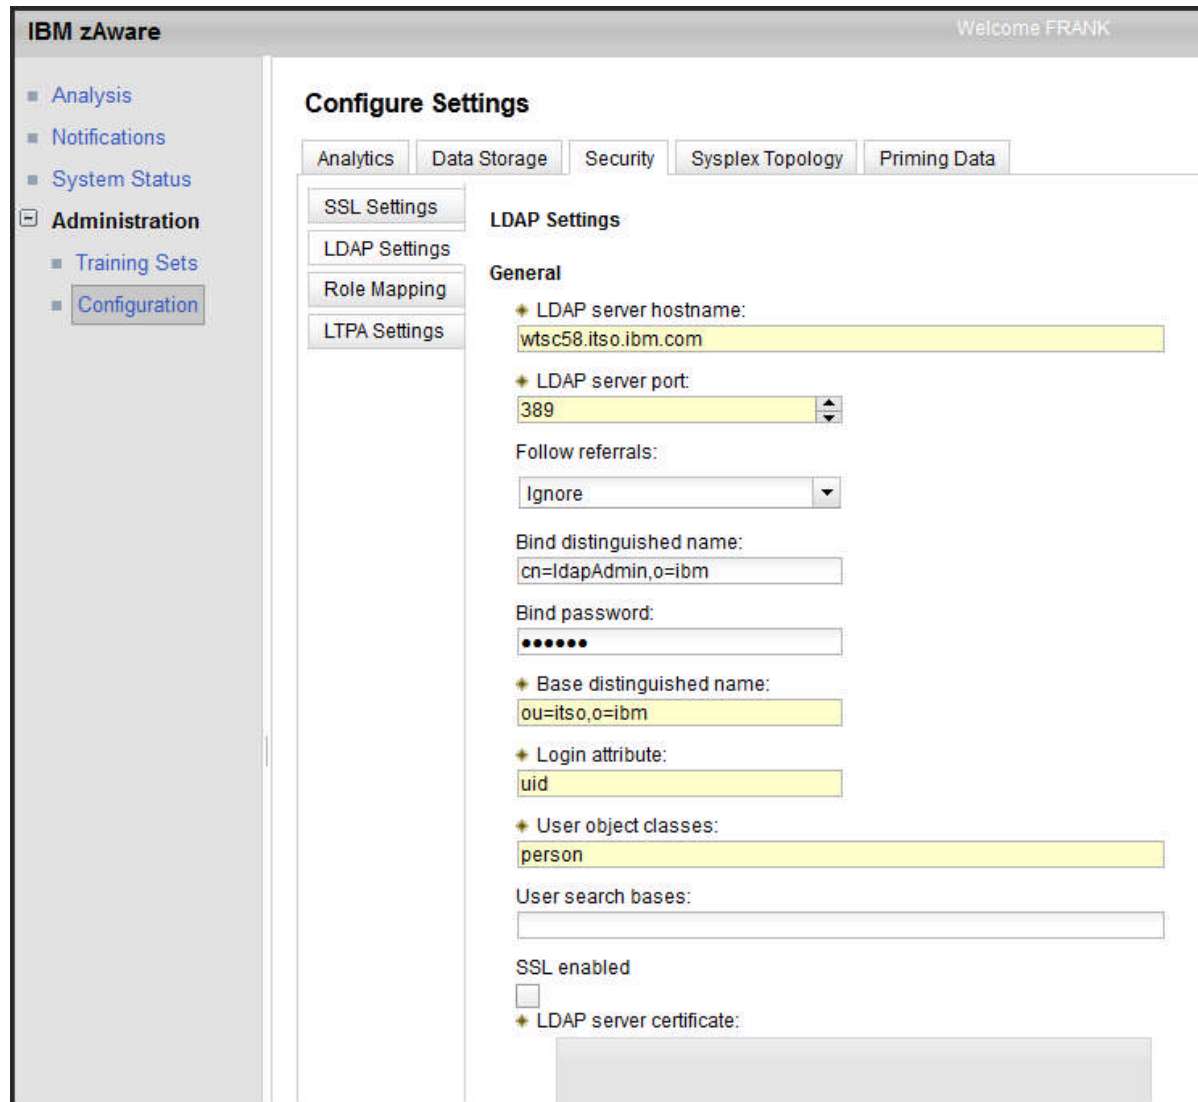
- Serial number: 1346288473946719250
- Validity period: Valid from Aug 29, 2012 to Aug 29, 2013.
- Issued to: CN=localhost
- Issued by: CN=localhost
- Fingerprint: F5:F7:C1:BB:20:25:84:DA:EF:E0:D2:51:B4:EF:E2:B2:54:2D:8B:4A

Below this information is a 'Certificate Actions' section with three buttons: 'Generate Certificate Signing Request', 'View Last Generated Request', and 'Receive Certificate Request Reply'.

- View SSL certificate data
- Generate Certificate Signing Request
- View Last Generated Certificate Signing Request
- Receive Certificate Request Reply



LDAP setting (security)



IBM zAware can be configured to use an LDAP repository

The login attribute determine what part of the LDAP users distinguished name must be used as the login ID on the zAware GUI login prompt



IBM zAware users

Two ways to add user to IBM zAware:

1. External LDAP repository

- Configure IBM zAware LDAP setting to fetch LDAP user entries
- Map fetched users to IBM zAware Roles
- Remote LDAP server must be active for IBM zAware user logins

2. IBM zAware internal repository

- Through Integrated Solutions Console
- Add a user entry to IBM zAware repository
- Map user to IBM zAware Roles

LDAP settings (security) - groups

The screenshot shows the IBM zAware configuration interface. On the left is a navigation menu with categories: Analysis, Notifications, System Status, Administration (expanded), Training Sets, and Configuration. The main area is titled 'Configure Settings' and has tabs for Analytics, Data Storage, Security (selected), Sysplex Topology, and Priming Data. Under the Security tab, there are sub-tabs for SSL Settings, LDAP Settings (selected), Role Mapping, and LTPA Settings. The LDAP Settings section includes a collapsed 'LDAP server certificate' section. Below that is the 'Group' configuration section with the following fields:

- User group membership attribute:
- User group membership scope:
- Group object classes: - Group search bases:
- Group member attributes: - Group member object classes:
- Group member scope:

At the bottom of the configuration section are 'Apply' and 'Reset' buttons.

Group entries are also supported



Configure LDAP

Configure Settings

Analytics | Data Storage | Security | Sysplex Topology | Priming Data

SSL Settings | **LDAP Settings** | Role Mapping | LTPA Settings

LDAP Settings

General

* LDAP server hostname:
wtsc58.itso.ibm.com

* LDAP server port:
389

Follow referrals:
Ignore

Bind distinguished name:
cn=ldapAdmin,o=ibm

Bind password:

* Base distinguished name:
ou=itso,o=ibm

* Login attribute:
uid

* User object classes:
person

User search bases:
[Empty field]

SSL enabled

* LDAP server certificate:

Hostname and Port required

Will bind anonymously if no bind DN/PW

Base DN sets position in LDAP hierarchy

Login Attribute is LDAP attribute that users will login to IBM zAware

If SSL is enabled then SSL certificate is required





Add user to local repository

- Login to Integrated solutions console with IBM zAware Admin ID
- <http://<IBM zAware IP address or host name>/ibm/console>
 - <https://9.12.5.128/ibm/console>

Integrated Solutions Console

Log in to the console.

User ID:

Password:

Log in



ISC manage users

Integrated Solutions Console Welcome FRANK

View: All tasks

- Welcome
- Security
- Users and Groups
 - Manage Users
 - Manage Groups
- Troubleshooting
- Console Identity

Manage Users

Search for Users

Search by: User ID, * Search for: *, * Maximum results: 100

Search

Create... Delete

Page 1 of 1 Total: 0

Select Manage users under Users and Groups menu

Select Create

From the manage users panel select create to define a new user

ISC create user

The screenshot shows a web interface titled "Manage Users". Below the title bar, there is a sub-header "Manage Users". The main content area is titled "Create a User". It contains the following fields and controls:

- * User ID**: A text input field with a "Group Membership" dropdown menu to its right.
- * First name**: A text input field.
- * Last name**: A text input field.
- E-mail**: A text input field.
- * Password**: A text input field.
- * Confirm password**: A text input field.
- Create** and **Cancel**: Two buttons at the bottom left of the form.

Enter the required fields for the new user:

User ID = login id for IBM zAware GUI

Password is case sensitive

Select Create to complete user entry

Once completed the user can be assigned a role through the IBM zAware GUI

Role Mapping

Users defined in Remote LDAP repository or through the Integrated Systems Console must be mapped to a role (USER or ADMIN) through the Role Mapping Panel.

Authority by Role

- Admin
 - Any task in the GUI

- User
 - On the Analysis page and Interval view, all controls and actions are permitted.
 - On the System Status page, all actions are disabled.
 - On the Notifications page, all actions are disabled.

Assigning users to roles

Configure Settings

Analytics | Data Storage | Security | Sysplex Topology | Priming Data

SSL Settings | LDAP Settings | **Role Mapping** | LTPA Settings

Role Mapping

Role: User

Filter: *

Search limit: 20

Search

Select Role

Select from list of sers retrieved from LDAP or take action against defined users

Set Search filter and click Search

Available users

- FRANK
- angela
- troy
- sawada
- singh
- noonan
- sawadauser
- barton

Current mapped users

- FRANK
- sawada
- noonanuser

Add >

Add All >>

< Remove

<< Remove All

Select Action and Apply



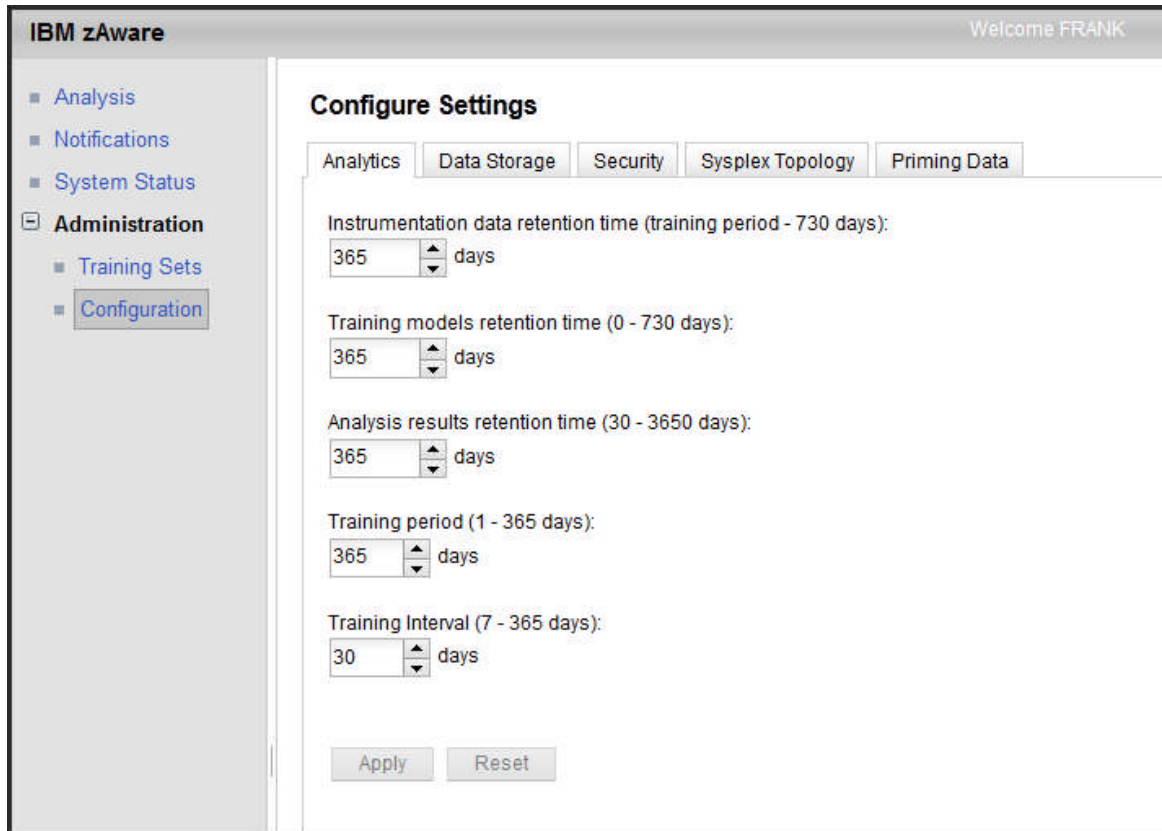
Quick Tour of IBM zAware Panels

- The following slides provide a quick look at the panels on the IBM zAware GUI

Context Help

The screenshot shows the IBM zAware web interface. On the left is a navigation menu with categories like Analysis, Notifications, System Status, Administration, Training Sets, and Configuration. The main area is titled 'Configure Settings' and has tabs for Analytics, Data Storage, Security, Sysplex Topology, and Priming Data. The 'Analytics' tab is active, showing several retention and interval settings with spinners and 'Apply'/'Reset' buttons. A context help window is overlaid on top, titled 'Specifying settings for the analytics engine', with a URL 'https://9.12.5.127/zAware/help/gui_AnalyticsTab.html'. A red arrow points from the text 'Click for HELP' to the 'Help' link in the top right corner of the main page.

Configuration Setting options



Once storage has been added and made available, you can now configure additional settings.

Analytics: Modify settings related to data retention and training settings.

Data Storage: view storage status, add or remove storage

Security: view SSL setting or change server certificate, Configure LDAP, assign roles to users, configure session timeout values

Sysplex Topology: View or reassign LPARS to SYSPLEX

Priming Data: Assign bulk loaded data to a SYSPLEX



Analytics

Configure setting for data retention and training

How long to keep data from monitored systems.

How long to keep models.

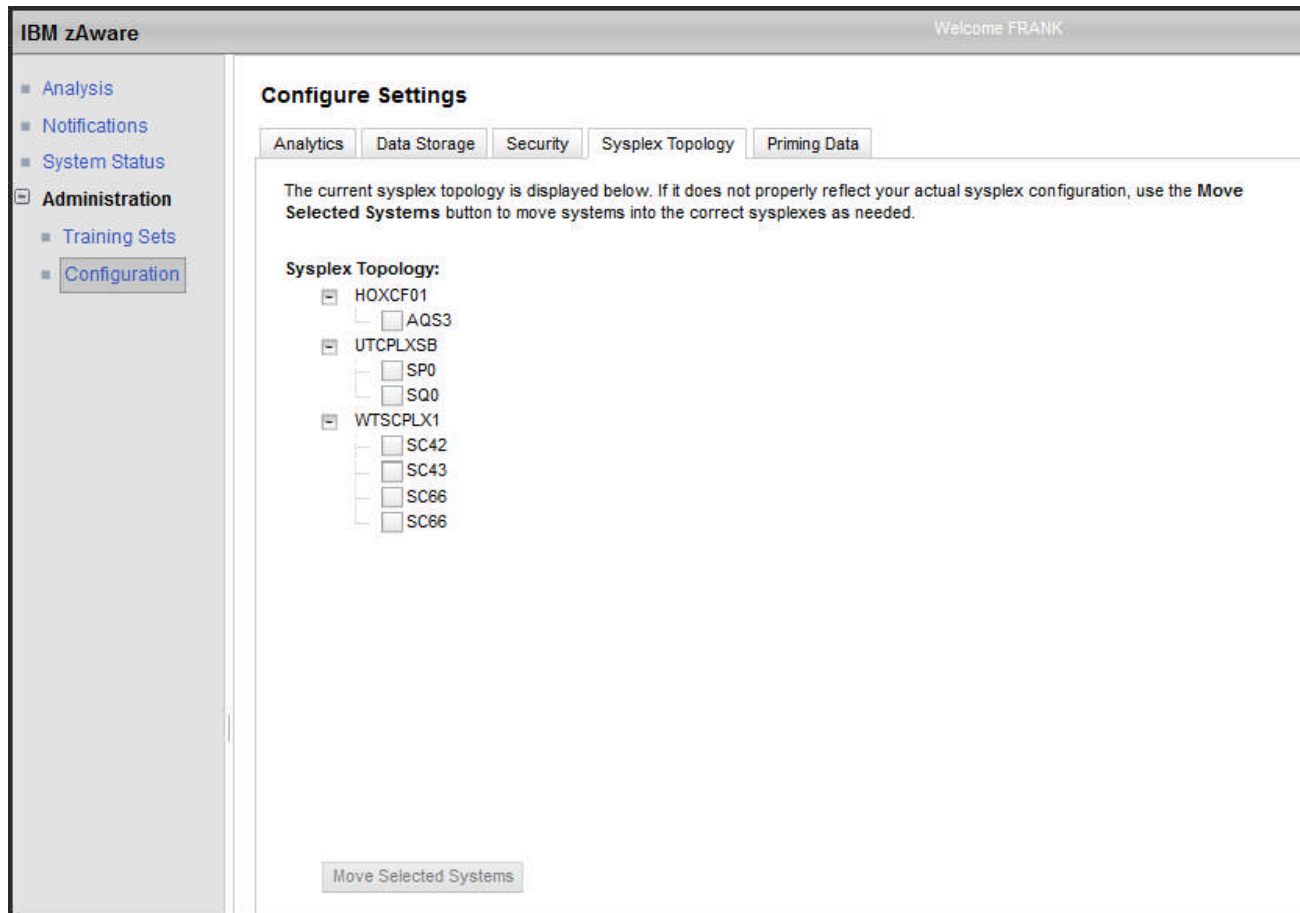
How long to keep analysis results.

Number of consecutive days to use as a training period

How many days before a new model is built (training)



Sysplex Topology



Displays the known Sysplex(es) and members

Allows moving of members from one sysplex to another

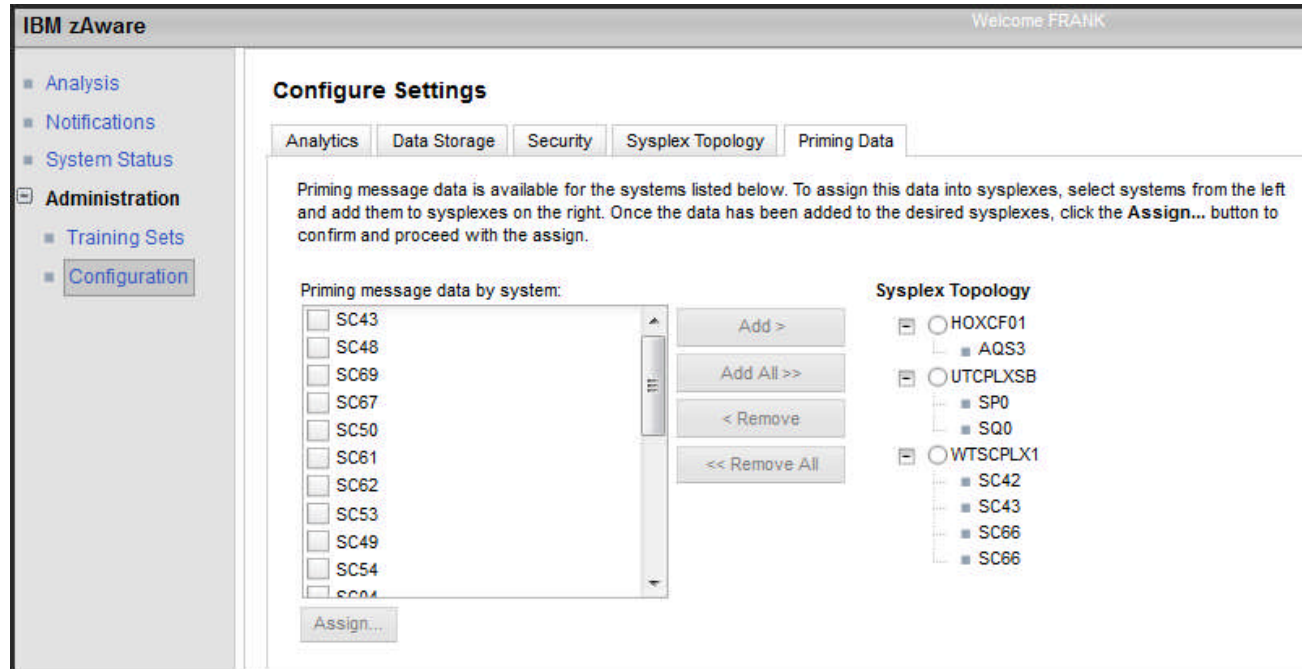
Use in the event that you incorrectly assigned a member to the wrong sysplex

Or

Actually move a member from one sysplex to another



Priming Data



Assign priming data to an appropriate Sysplex

At least one member of the Sysplex must have connected to IBM zAware



Training Sets

IBM zAware Welcome FRANK Log out

- Analysis
- Notifications
- System Status
- ▣ Administration
 - Training Sets
 - Configuration

[Help](#)

Training Sets

The Monitored Systems table provides training statuses and results for IBM zAware monitored systems. The Actions menu provides functions for managing model dates, requesting or canceling training, and managing ignored messages. Training details for a given system can be accessed by clicking on links in the Training Progress and Last Training Result columns.

Monitored Systems

Actions	System	Training Progress	Last Training Result	Last Training Result Time	Current Model Built
<ul style="list-style-type: none"> Manage Model Dates <li style="background-color: #e0e0e0;">Request Training Cancel Training 	OXCFO1	—	✔ Complete	July 25, 2012 11:29:57 AM Eastern Daylight Time	July 25, 2012 11:29:57 AM Eastern Daylight Time
<input type="radio"/> SP0	UTCPLXSB	—	✔ Complete	August 23, 2012 8:00:12 PM Eastern Daylight Time	August 23, 2012 8:00:12 PM Eastern Daylight Time
<input type="radio"/> SQ0	UTCPLXSB	—	✔ Complete	August 23, 2012 8:01:11 PM Eastern Daylight Time	August 23, 2012 8:01:11 PM Eastern Daylight Time
<input type="radio"/> SC42	WTSCPLX1	—	✔ Complete	August 29, 2012 9:32:22 PM Eastern Daylight Time	August 29, 2012 9:32:22 PM Eastern Daylight Time
<input type="radio"/> SC43	WTSCPLX1	—	✔ Complete	August 22, 2012 8:00:06 PM Eastern Daylight Time	August 22, 2012 8:00:06 PM Eastern Daylight Time
<input type="radio"/> SC66	WTSCPLX1	—	✘ Failed	September 12, 2012 8:00:08 PM Eastern Daylight Time	August 2, 2012 8:00:04 PM Eastern Daylight Time
<input type="radio"/> SC66	WTSCPLX1	—	✘ Failed	September 12, 2012 8:00:08 PM Eastern Daylight Time	August 2, 2012 8:00:04 PM Eastern Daylight Time

▶ Current Training Status Details (Click on training statuses above to view details)

Last Refresh: Thu Sep 13 2012 12:18:37 GMT-0400 (Eastern Daylight Time)

Displays results of training.

Action menu allows you to Manage Models or Request Training



System Status

IBM zAware
Welcome FRANK
Log out

- Analysis
- Notifications
- System Status
- ▣ Administration
 - Training Sets
 - Configuration

System Status

System Status displays the IBM zAware analytics engine status, as well as monitored systems information for z/OS systems connected to IBM zAware. Click the start button () to start the analytics engine, and the stop button () to stop it.

Analytics engine status: Running

IBM zAware Monitored System Data Suppliers:

System	Sysplex	Status	Instrumentation Data Type	Connect Start Time
AQS3	HOXCF01	Inactive	OPERLOG	September 7, 2012 10:57:46 AM Eastern Daylight Time
SC42	WTSCPLX1	Active	OPERLOG	September 10, 2012 10:40:47 AM Eastern Daylight Time
SC43	WTSCPLX1	Active	OPERLOG	September 10, 2012 10:38:47 AM Eastern Daylight Time
SC66	WTSCPLX1	Active	OPERLOG	September 10, 2012 10:43:26 AM Eastern Daylight Time
SQ0	UTCPLXSB	Inactive	OPERLOG	September 10, 2012 3:42:30 PM Eastern Daylight Time
SP0	UTCPLXSB	Inactive	OPERLOG	September 10, 2012 3:42:30 PM Eastern Daylight Time

Last refresh: Thu Sep 13 2012 12:20:27 GMT-0400 (Eastern Daylight Time)

Displays Status of Analytics Engine with option to start/stop.

Display Status of monitored clients



Notifications

IBM zAware Welcome FRANK

- Analysis
- **Notifications**
- System Status
- ☐ Administration
 - Training Sets
 - Configuration

Notifications

Notification messages

Actions ▼

<input type="checkbox"/> Message ID	Message Text	Message Date/Time
<input type="checkbox"/> AIFT0104I	Priming request started at Mon, 10 Sep 2012 18:36:31 +0000 processed 353 lines 82 lines had messages	Mon Sep 10 2012 14:37:53 GMT-0400 (Eastern Daylight Time)
<input type="checkbox"/> AIFT0104I	Priming request started at Mon, 10 Sep 2012 18:39:13 +0000 processed 129 lines 68 lines had messages	Mon Sep 10 2012 15:50:32 GMT-0400 (Eastern Daylight Time)

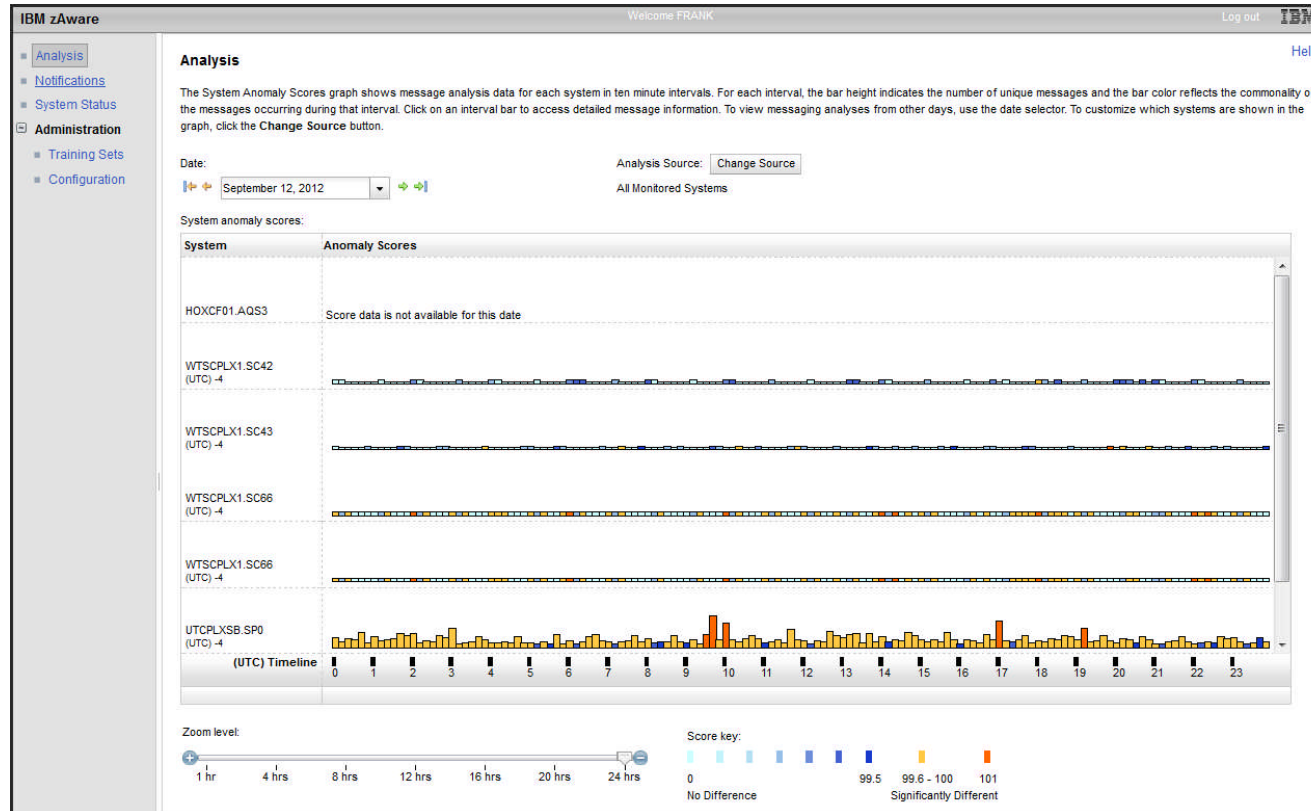
Notifications panel lists messages issued by IBM zAware

Last Refresh: Thu Sep 13 2012 12:23:14 GMT-0400 (Eastern Daylight Time)





Analysis



Displays message analysis of monitored clients

Drill down for Interval details





End Quick tour

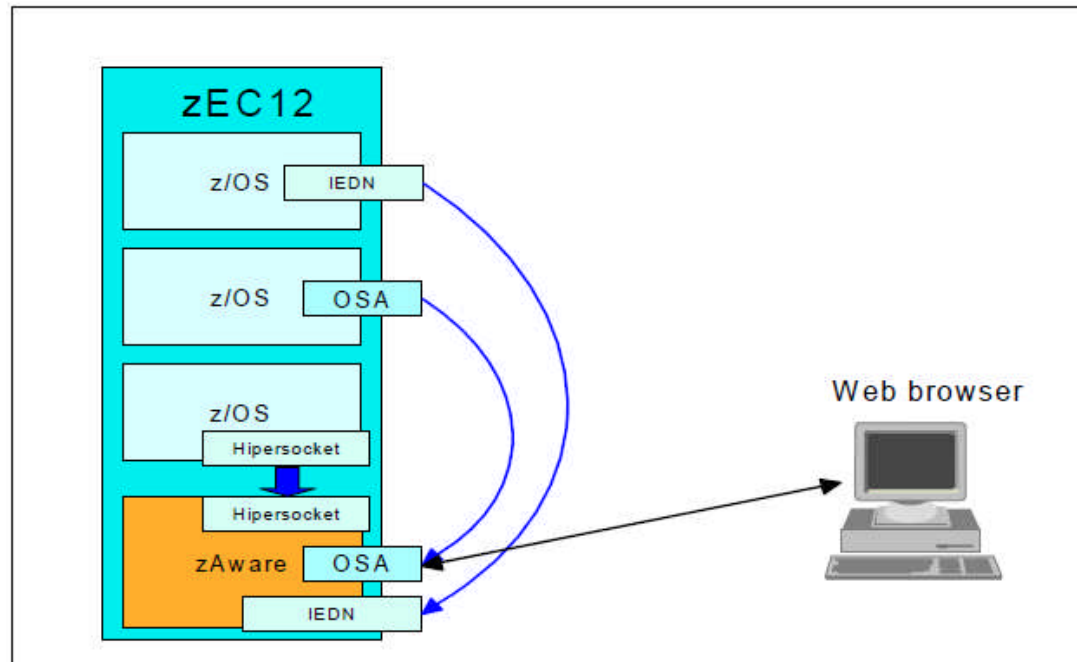
- End of Quick tour



Configuring a monitored client

- **Configure the network connection**
 - Layer 3 connection through OSD, OSX or IQD chipid
 - Make appropriate changes to TCP/IP
 - INTERFACE statements
 - DEVICE and LINK statements
 - Hosts table or DNS entry
 - Firewall changes
- **Set up OPERLOG**
 - Add zai parameters
- **Define SAF FACILITY class profile IXGAWARE_CLIENT**
 - Define/change log stream
- **Parmlib changes**
 - Add zai parameters IXGCNFxx and activate

Networking



- Setup will vary by installation – work with your network administrators
- IPv4 and IPv6 supported
- Layer 3 connection required
- If monitored clients are on the same System z server as the zAware LPAR then you can use a HiperSocket connection
- OSA Express connectivity through OSD and OSX (IEDN)
- Update TCP/IP profiles as required
- Firewall changes may be needed
- IBM zAware does not initiate connections to monitored clients; data is pushed by clients to IBM zAware

Networking

- The network bandwidth required will depend on the peak message rate of connected clients. The message rate can be determined by analyzing a representative interval of syslog or OPERLOG data using the Message Analysis Program.
- See Appendix A of Extending z/OS System Management Functions with IBM zAware, SG24-8070-00
- System Logger will buffer 99GB of data for times when the message rate exceeds the network bandwidth.

Security

- Facility class profile to protect IBM zAware related parameter updates to OPERLOG log stream
- Also when sending Bulk Data a model and temporary log stream are defined
- RDEFINE FACILITY IXGZAWARE_CLIENT UACC(NONE)
 - UPDATE access to those who
 - Will modify OPERLOG to add ZAI parameters
 - Will submit bulk load job

OPERLOG

- OPERLOG required to have messages sent to IBM zAware
- Each monitored clients sends its OPERLOG data through System Logger
- Not sure if OPERLOG Is active?
 - DISPLAY CONSOLES,HARDCOPY command

```
DISPLAY CONSOLES,HARDCOPY
CNZ4100I 16.14.47 CONSOLE DISPLAY 845
CONSOLES MATCHING COMMAND: D CONSOLES,HARDCOPY
MSG:CURR=0 LIM=3000 RPLY:CURR=1 LIM=999 SYS=#@$A PFK=00
HARDCOPY LOG=(SYSLOG,OPERLOG) CMDLEVEL=CMDS
ROUT=(ALL)
LOG BUFFERS IN USE: 0 LOG BUFFER LIMIT: 6000
```

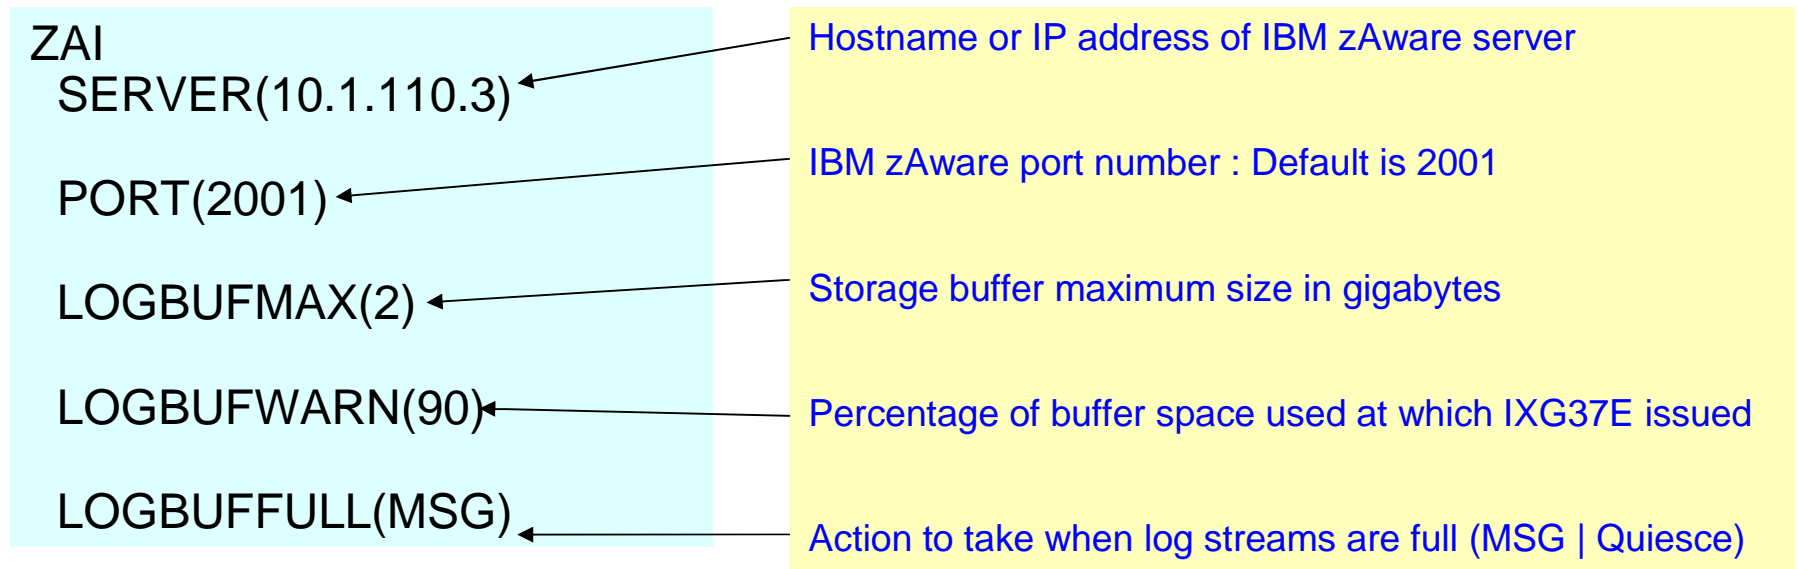
- Update existing OPERLOG LOG STREAM

```
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DATA TYPE(LOGR) REPORT(YES)
UPDATE LOGSTREAM NAME(SYSPLEX.OPERLOG)
ZAI(YES)
ZAIDATA('OPERLOG')
```



System Logger

- System Logger is responsible for sending data to the IBM zAware server
- Requires additional statements in IXGCNFxx
- Set IXGCNFxx to apply updates



Verify System Logger connection

- Updates to DISPLAY LOGGER command
- SETLOGR LOGGER,STATUS,ZAI,VERIFY
 - Opens socket connection to IBM zAware Server

```

D LOGGER,status,zai,verify
IXG601I 15.31.39  LOGGER DISPLAY 779
SYSTEM  LOGGER STATUS
SYSTEM  SYSTEM LOGGER STATUS
-----
#@$A    ACTIVE

ZAI LOGSTREAM CLIENTS: AVAILABLE
BUFFERS IN USE:    00 GB 0000 MB
ZAI VERIFY INITIATED,  CHECK FOR MESSAGES IXG37X, IXG38X

LOGGER PARAMETER OPTIONS
KEYWORD          SOURCE  VALUE
-----
ZAI
  SERVER         IPV4   SET (00) 10.1.110.3
  PORT           SET (00) 2001
  LOGBUFMAX      SET (00) 02
  LOGBUFWARN     SET (00) 90
  LOGBUFFULL     SET (00) MSG
    
```

Logger status

ZAI status

IXGCNFxx options



Verify System Logger connection

- Updates to DISPLAY LOGGER command
- SETLOGR LOGGER,STATUS,ZAI,VERIFY
 - Opens socket connection to IBM zAware Server

```
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS
FOR DISPLAY ZAI,VERIFY
STATUS: ATTEMPTING SOCKET CREATE
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS
FOR DISPLAY ZAI,VERIFY
STATUS: SOCKET CREATE SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS
FOR DISPLAY ZAI,VERIFY
STATUS: ATTEMPTING SOCKET CONNECT
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS
FOR DISPLAY ZAI,VERIFY
STATUS: SOCKET CONNECT SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS
FOR DISPLAY ZAI,VERIFY
STATUS: INITIATING SOCKET VALIDATION
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS
FOR DISPLAY ZAI,VERIFY
STATUS: SOCKET VALIDATION SUCCESSFUL
IXG380I ZAI LOGSTREAM CLIENT ESTABLISHED
FOR DISPLAY ZAI,VERIFY
```

D LOGGER,status,zai

```

D LOGGER,status,zai
IXG601I  16.02.24  LOGGER DISPLAY 808
SYSTEM  LOGGER STATUS
SYSTEM  SYSTEM LOGGER STATUS
-----  -----
#@$A    ACTIVE

ZAI LOGSTREAM CLIENTS: AVAILABLE
BUFFERS IN USE:    00 GB 0000 MB

LOGGER PARAMETER OPTIONS
KEYWORD          SOURCE      VALUE
-----  -----  -----
ZAI
SERVER          IPV4    SET (00) 10.1.110.3
PORT            SET (00) 2001
LOGBUFMAX       SET (00) 02
LOGBUFWARN     SET (00) 90
LOGBUFFULL     SET (00) MSG
    
```



Additional information in LOG STREAM display

```

-D LOGGER,c,lsn=sysplex.operlog,d
IXG601I  16.07.57  LOGGER DISPLAY 825
CONNECTION INFORMATION BY LOGSTREAM FOR SYSTEM #@$A
LOGSTREAM          STRUCTURE          #CONN  STATUS
-----          -
SYSPLEX.OPERLOG    SYSTEM_OPERLOG    000003  IN USE
  DUPLEXING: STAGING DATA SET
    STGDSN: IXGLOGR.SYSPLEX.OPERLOG.#@$A
      VOLUME=DISTS6  SIZE=007740 (IN 4K)  % IN-USE=010
  GROUP: PRODUCTION  ZAI CLIENT: YES - QUIESCED
    
```

Connecting and Disconnecting a client

Updates to SETLOGR command

- TO CONNECT:
- SETLOGR FORCE,ZAICONNECT,LSN=<OPERLOG LOG STREAM>

- TO DISCONNECT (QUIESCE)
- SETLOGR FORCE,ZAIQUIESCE,LSN=<OPERLOG LOG STREAM>

- Stopping Analytics Engine from Analysis screen of IBM zAware GUI terminates all connections from monitored systems

SETLOGR FORCE,ZAICONNECT,LSN=XX

```
SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
IXG651I SETLOGR FORCE ZAICONNECT COMMAND ACCEPTED 763
FOR LOGSTREAM=SYSPLEX.OPERLOG
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS 764
FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS: ATTEMPTING SOCKET CREATE
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS 765
FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS: SOCKET CREATE SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS 766
FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS: ATTEMPTING SOCKET CONNECT
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS 767
FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS: SOCKET CONNECT SUCCESSFUL
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS 768
FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS: INITIATING SOCKET VALIDATION
IXG386I ZAI LOGSTREAM CLIENT CONNECT ATTEMPT IN PROGRESS 769
FOR LOGSTREAM SYSPLEX.OPERLOG
STATUS: SOCKET VALIDATION SUCCESSFUL
IXG380I ZAI LOGSTREAM CLIENT ESTABLISHED 770
FOR LOGSTREAM SYSPLEX.OPERLOG
```



Display commands

```
D LOGGER,status,zai
IXG601I 10.33.55  LOGGER DISPLAY 772
SYSTEM  LOGGER STATUS
SYSTEM  SYSTEM LOGGER STATUS
-----
#@$A    ACTIVE


ZAI LOGSTREAM CLIENTS: ACTIVE
BUFFERS IN USE: 00 GB 0000 MB

LOGGER PARAMETER OPTIONS
KEYWORD          SOURCE  VALUE
-----
ZAI
SERVER          IPV4   SET (00) 10.1.110.3
PORT            SET (00) 2001
LOGBUFMAX       SET (00) 02
LOGBUFWARN      SET (00) 90
LOGBUFFULL      SET (00) MSG
```



Display commands

```
D LOGGER,c,lsn=sysplex.operlog,d
IXG601I 10.34.16  LOGGER DISPLAY 774
CONNECTION INFORMATION BY LOGSTREAM FOR SYSTEM #@$A
LOGSTREAM          STRUCTURE          #CONN  STATUS
-----          -
SYSPLEX.OPERLOG    SYSTEM_OPERLOG    000003  IN USE
  DUPLEXING: STAGING DATA SET
    STGDSN: IXGLOGR.SYSPLEX.OPERLOG.#@$A
      VOLUME=DISTS6  SIZE=007740 (IN 4K)  % IN-USE=022
GROUP: PRODUCTION  ZAI CLIENT: YES - CONNECTED
ZAI DATA: OPERLOG
LOG BLOCKS SENT TO SERVER OK: 0000000010, FAILED: 0000000000
```







Cumulative counter of log blocks sent to server
Resets if connection closed




System Status – Active client connection

System Status

System Status displays the IBM zAware analytics engine status, as well as monitored systems information for z/OS systems connected to IBM zAware. Click the start button () to start the analytics engine, and the stop button () to stop it.

Analytics engine status: Running
 

IBM zAware Monitored System Data Suppliers:

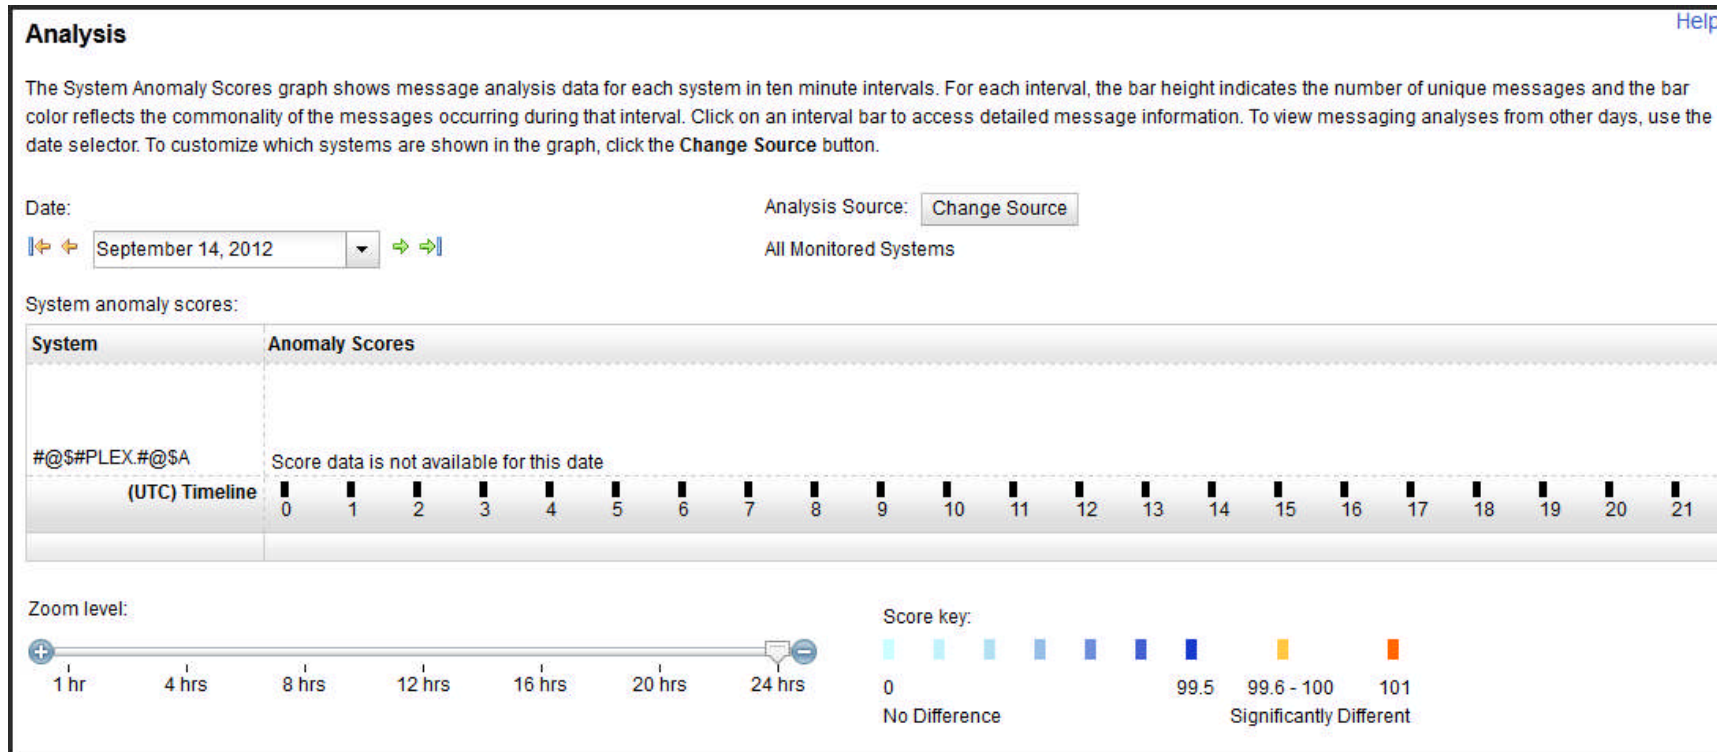
System	Sysplex	Status	Instrumentation Data Type	Connect Start Time
#@\$A	#@\$#PLEX	 Active	OPERLOG	September 14, 2012 10:30:38 AM Eastern Daylight Time

Last refresh: Fri Sep 14 2012 10:43:19 GMT-0400 (Eastern Daylight Time)

System Status Panel of IBM zAware GUI will show the client as now connected. Instrumentation type indicates OPERLOG data is being received.



Analysis



Although the newly connected client appears on the Analysis Screen, no data is reported. The client must be trained first (a model built).



Disconnecting (zaiquiesce) a client

```
SETLOGR FORCE,ZAIQUIESCE,LSN=SYSPLEX.OPERLOG
```



```
IXG651I SETLOGR FORCE ZAIQUIESCE COMMAND ACCEPTED 777  
FOR LOGSTREAM=SYSPLEX.OPERLOG
```



```
IXG382I ZAI LOGSTREAM CLIENT QUIESCED 778  
FOR LOGSTREAM SYSPLEX.OPERLOG  
REASON: SETLOGR COMMAND REQUEST.
```

```
IXG387I ZAI LOGSTREAM CLIENT CONNECTION ENDED SUMMARY 779  
FOR LOGSTREAM SYSPLEX.OPERLOG  
CONNECTION WAS ESTABLISHED AT: 09/14/2012 10:30:38 LOCAL  
LOG BLOCKS SENT TO SERVER OK: 13, FAILED: 0
```

System Status – Inactive client connection

System Status

System Status displays the IBM zAware analytics engine status, as well as monitored systems information for z/OS systems connected to IBM zAware. Click the start button () to start the analytics engine, and the stop button () to stop it.

Analytics engine status: Running
 

IBM zAware Monitored System Data Suppliers:

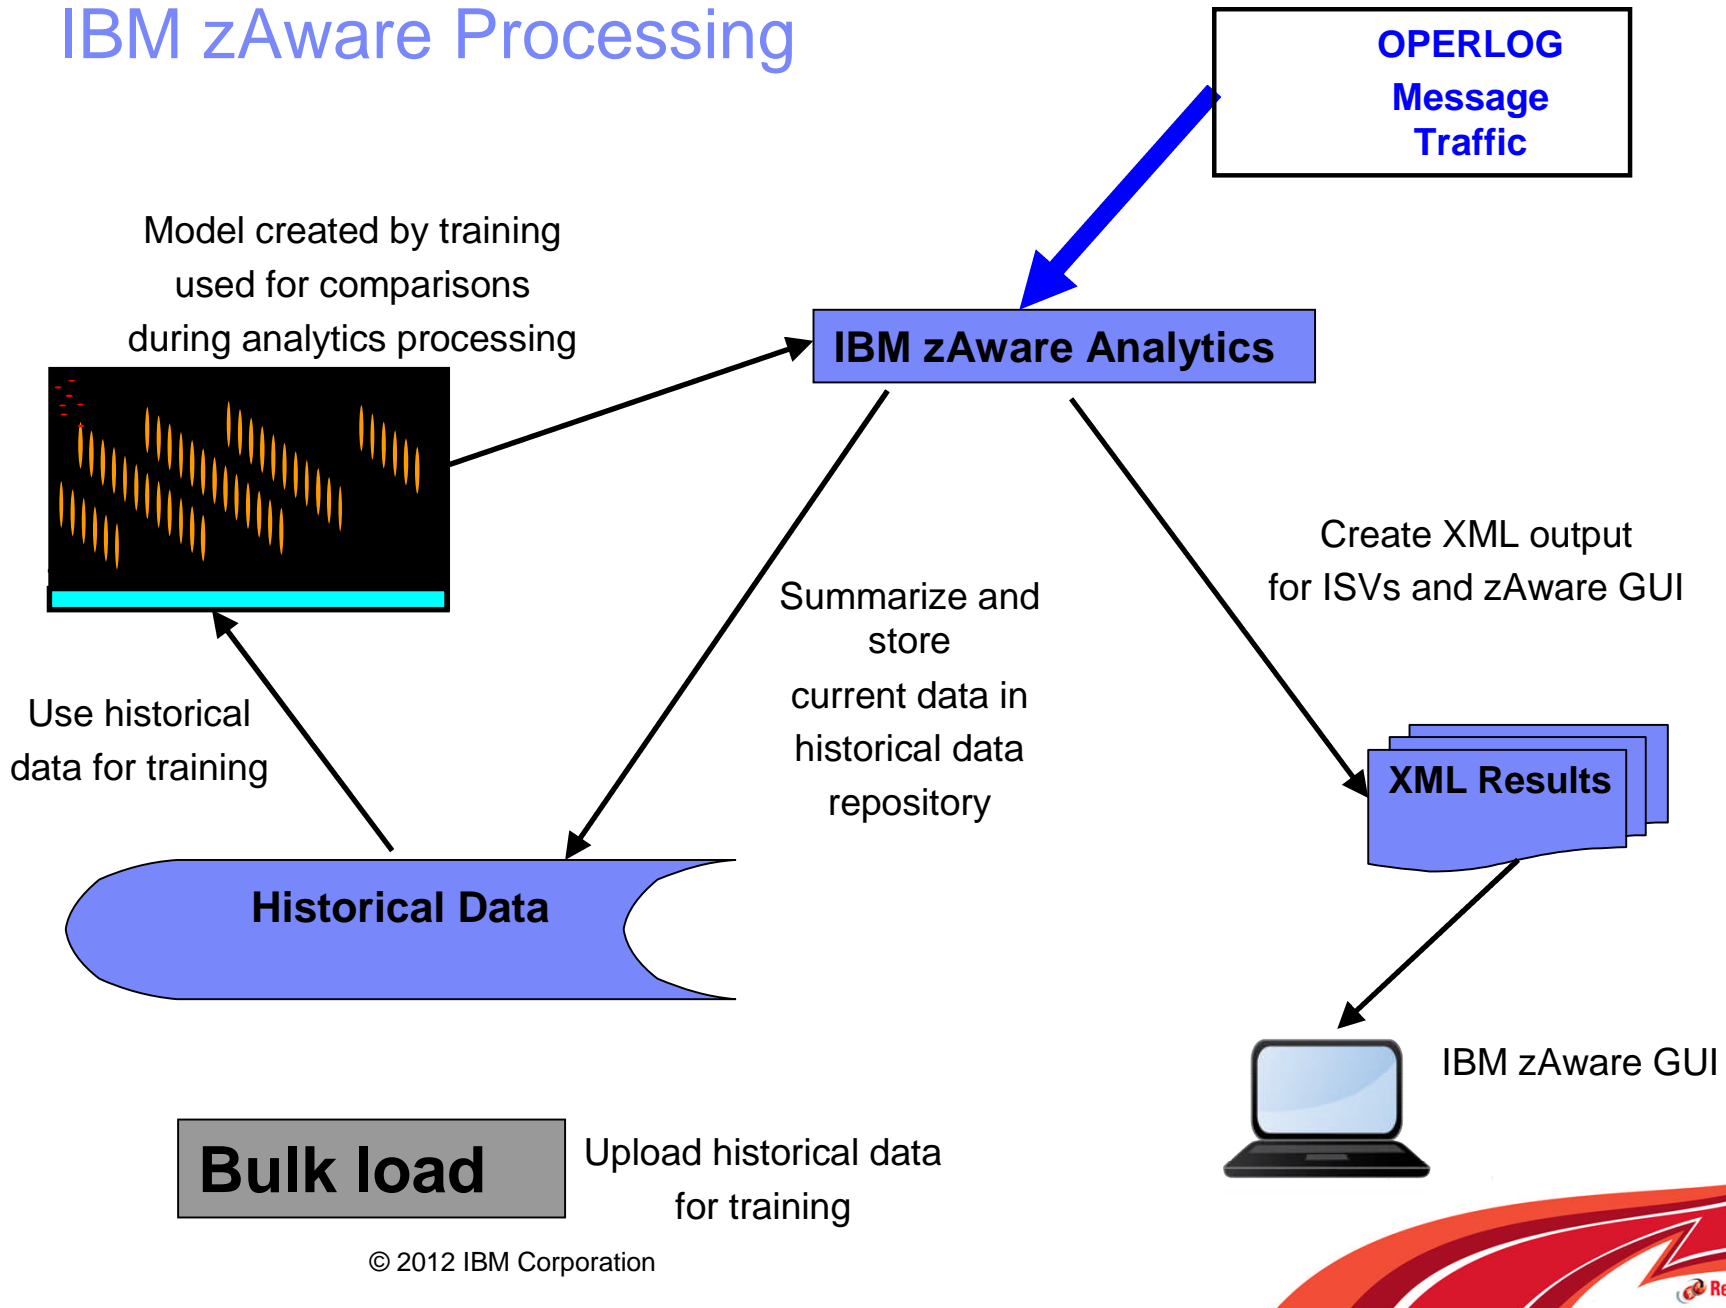
System	Sysplex	Status	Instrumentation Data Type	Connect Start Time
#@\$A	#@\$#PLEX	<input checked="" type="checkbox"/> Inactive	OPERLOG	September 14, 2012 10:30:38 AM Eastern Daylight Time

Refresh the Screen

Last refresh: Fri Sep 14 2012 10:48:57 GMT-0400 (Eastern Daylight Time)



IBM zAware Processing



IBM zAware Processing

- Prime database: The database can be primed by performing a bulk load or by allowing zAware to collect data in real time until enough data exists for training. When the message data is received by zAware, it is parsed, summarized, and stored in the instrumentation data database.
- Training: The training creates the model of historical data. Enough data must exist in order for the training to occur. The training can be run after the data is bulk loaded or you can wait for the analytics phase to upload enough data for training and run it manually at that time or wait for the next scheduled training to occur.
- Analytics: The analytics has two tasks:
 - The analytics task continuously processes incoming message data. The current data is parsed, summarized, and stored in the instrumentation data database to be used for training in the future.
 - When a model exists, every two minutes the analytics step uses advanced analytics to compare the current data to the model. The results of the comparisons are output to the xml files and to the zAware GUI.
- Each bar on the graph represents a ten minute interval. Once a ten minute interval is complete, the results for that interval are hardened and saved in the zAware file system. That is, the color and height of the bar no longer changes.
- If the interval does not yet contain ten minutes of analytics, the color and height of the bar can change every two minutes based on the results of the analytics.

Priming (Bulk Load) and Training

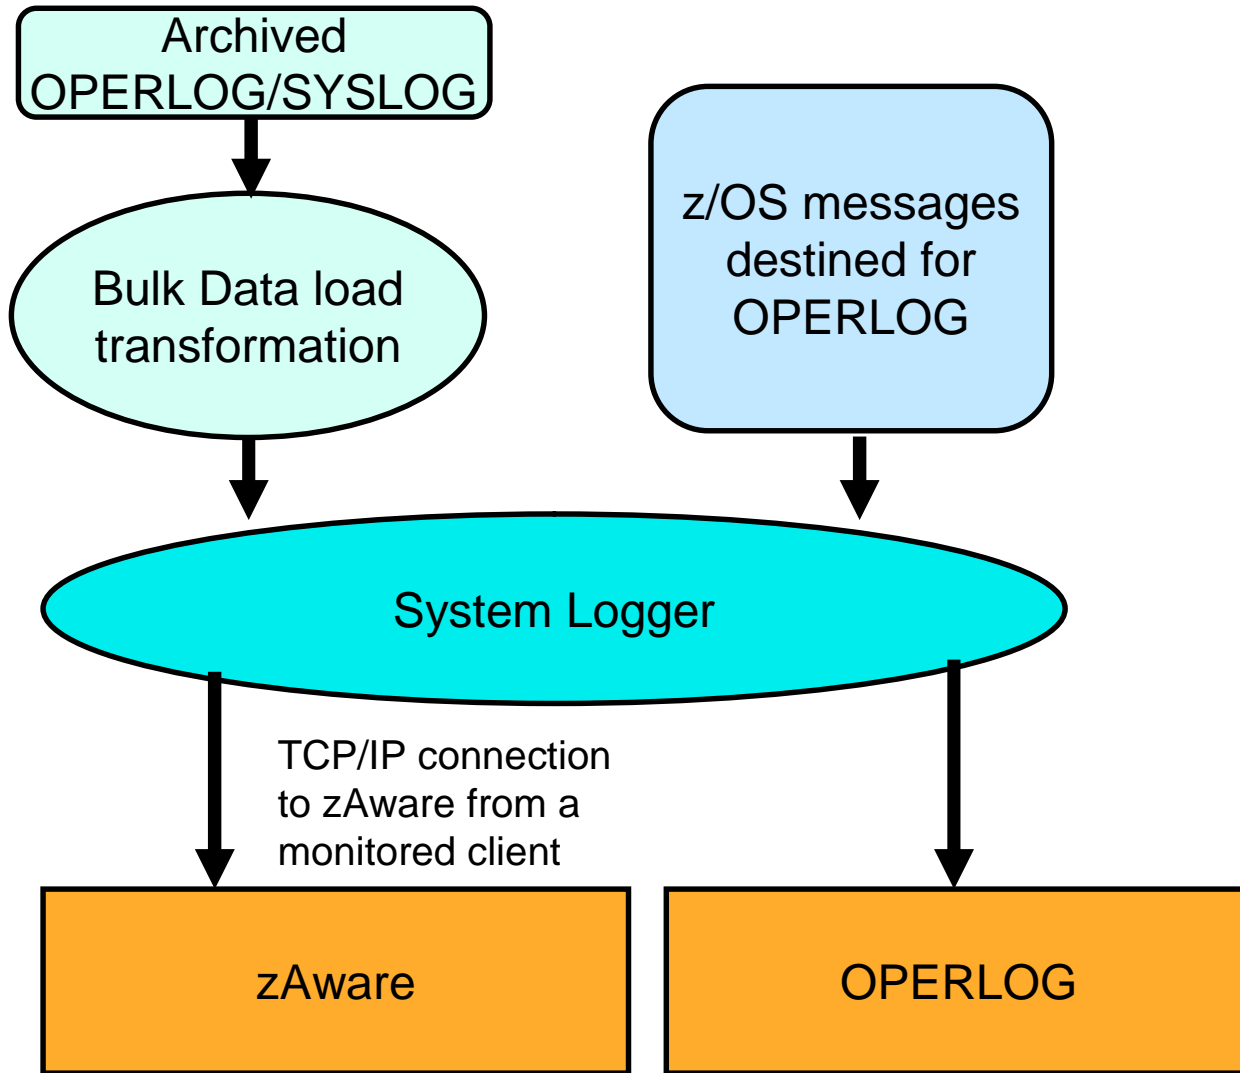
- IBM zAware focuses on identifying message anomalies
 - Input data is OPERLOG messages
 - Analyzed against a model of previous system message behavior
 - Output are messages that are identified as ‘different’ from previous behavior
 - Analytics requires some history of previous behavior to create a model
 - Training is the process of creating a model using historical data
 - Choices to PRIME
 - Connect a system and ‘force’ train with no previous history
 - Let it run for 90 days (accumulate real time data) then ‘train’ it again
- OR
- ‘Prime’ it with historical data – bulk loading
 - Upload 90 days of OPERLOG/SYSLOG records
 - Train the system

Note: The Analysis results available through the web browser will not be populated with any information until both of the following events have taken place:

- The monitored client has been connected to zAware.
- A model of the monitored client has been built, using a mechanism known as “training”.



Bulk Loading



The Bulk Data Load utility consists of a batch job that runs a REXX exec to perform the transformation and load of the message data into a temporary log stream for transmission to the zAware application.

The monitored client that transmits the bulk data must have a connection to the zAware application.

The Bulk Data Load utility can process sequential data sets that contain data from syslog or OPERLOG in HCL (two-digit year) or HCR (four-digit year) format.



Steps to Bulk Load and Train

- Prepare the Hard Copy Log
- Connect the client
- Submit the bulk load job
- Disconnect all systems
- Assign systems to the sysplex
- Reconnect a system(s)
- Train the system(s)

Prepare the Hard Copy Log

- Data must be in sequential data sets that contain SYSLOG data stored in hardcopy log 2-digit year (HCL) or 4-digit year (HCR) format.
 - Can be concatenated data sets containing SYSLOG data in proper format
 - Can be OPERLOG data converted to SYSLOG format using IEAMDBLG

- SYS1.SAMPLIB(IEAMDBLG)
 - reads records from an OPERLOG log stream and convert them to SYSLOG format

Connect the Client

- You can perform the bulk load from any monitored z/OS client
- Before doing the bulk load, check the connection with the following commands.
 - DISPLAY LOGGER,STATUS,ZAI,VERIFY
 - DISPLAY LOGGER,C,LSN=SYSPLEX.OPERLOG,D
- If not connected, then connect the client
 - SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG
- Check the IBM zAware System Status Panel to verify the client is connected

Perform the Bulk Load

- Sample bulk load job is in SYS1.SAMPLIB(AIZBLK)
- 3 step job
- Creates a model logstream
- Runs a REXX exec to upload the data
- Deletes the model logstream

SYS1.SAMPLIB(AIZBLK) Step 1

```
//DEFINE EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSIN DD *
DATA TYPE (LOGR) REPORT(NO)
  DEFINE LOGSTREAM NAME(ZAWARE.ZAI.LSMODEL)
  MODEL(YES)
  ZAI(YES)
  ZAIDATA('SYSLOG')
  DASDONLY(YES)
  MAXBUFSIZE(65532)
  LS_SIZE(2560)
  STG_SIZE(2560)
  STG_DATACLAS(STGLOG)
LIST LOGSTREAM NAME(ZAWARE.ZAI.LSMODEL)
```

Defines the model log stream

SYS1.SAMPLIB(AIZBLK) Step 2

```
//BULKLOAD EXEC PGM=IKJEFT01,COND=(0,NE,DEFINE)
//SYSEXEC DD DISP=SHR,DSN=KYNEF.ZAI.REXX
//SYSTSPRT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSTSIN DD *
AIZBLKE KYNEF.ZAI.CONTROL CREATECNTLS
WHEN SYSRC(EQ 0) +
  AIZBLKE KYNEF.ZAI.CONTROL ADDSYSLOGDSN +
    KYNEF.WTSCPLX1.OPERLOG1
  AIZBLKE KYNEF.ZAI.CONTROL ADDSYSLOGDSN +
    KYNEF.WTSCPLX1.OPERLOG2
WHEN SYSRC(EQ 0) +
  AIZBLKE KYNEF.ZAI.CONTROL DISPLAYCNTLS
WHEN SYSRC(EQ 0) +
  AIZBLKE KYNEF.ZAI.CONTROL IMPORT DUMMY +
    ZAWARE.ZAI.LSTREAM ZAWARE.ZAI.LSMODEL
```

Data being bulk loaded



Comments in JCL provide instructions on how to customize the JCL

SYS1.SAMPLIB(AIZBLK) Step 3

```
//DELETE EXEC PGM=IXCMIAPU,COND=(0,NE,DEFINE)
//SYSPRINT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSIN DD *
  DATA TYPE (LOGR) REPORT(NO)
    DELETE LOGSTREAM NAME(ZAWARE.ZAI.LSMODEL)
```

Delete model log stream

Submit the Bulk Load Job

- Once the JCL has been edited to meet your installation requirements, submit the job.
- Display commands can be issued to monitor the status; this upload can be lengthy depending on the amount of data being transferred.



Disconnect all systems

zAware Welcome sawada

- Analysis
- Notifications
- ▣ Administration
 - System Status
 - Training Sets
 - Configuration

System Status

Analytics Engine status: Running



←
Click here to stop the engine

zAware managed system data suppliers:

System	Sysplex	Status	Instrumentation Data Type	Connect Start Time
AQS3	HOXCF01	Active	OPERLOG	August 2, 2012 4:54:22 PM Eastern Daylight Time
SC42	WTSCPLX1	Active	OPERLOG	August 2, 2012 12:01:28 AM Eastern Daylight Time
SC43	WTSCPLX1	Active	OPERLOG	August 2, 2012 4:05:45 PM Eastern Daylight Time
SC66	WTSCPLX1	Active	OPERLOG	August 2, 2012 4:05:55 PM Eastern Daylight Time
SQ0	UTCPLXSB	Active	OPERLOG	August 2, 2012 11:56:08 AM Eastern Daylight Time
SP0	UTCPLXSB	Active	OPERLOG	August 2, 2012 11:56:04 AM Eastern Daylight Time

Last refresh: August 3, 2012 10:25:32 AM Eastern Daylight Time

- Disconnect all monitored z/OS clients
- Stopping the Analytics Engine will close all monitored z/OS client connections
- When you assign priming data the analytics engine will be recycled anyway so this is optional

Assigning bulk loaded data for systems to Sysplex

- The bulk load data are syslog records from systems (either monoplex or sysplex)
- IBM zAware can determine which systems the bulk loaded messages belong to but...
- IBM zAware can not determine which Sysplex the bulk loaded data belongs to
- You must manually assign the bulk loaded data to a Sysplex
- At least one member of each SYSPLEX for which bulk loaded data is uploaded must have successfully connected to IBM zAware in order to assign data to that SYSPLEX

Assign bulk loaded data to Sysplex

Configure Settings

Analytics | Data Storage | Security | **Sysplex Topology** | Priming Data

Priming message data is available for the systems listed below. To assign this data into sysplexes, select systems from the left and add them to sysplexes on the right. Once the data has been added to the desired sysplexes, click the **Assign...** button to confirm and proceed with the assign.

Priming message data by system:

- #@\$2
- #@\$3

Add >
Add All >>
< Remove
<< Remove All

Sysplex Topology

- #@\$#PLEX
- #@\$A

Assign...

Bulk loaded data for each system displayed here.

Known Sysplexes displayed here. At least one client from a sysplex must have connected for this topology to be established

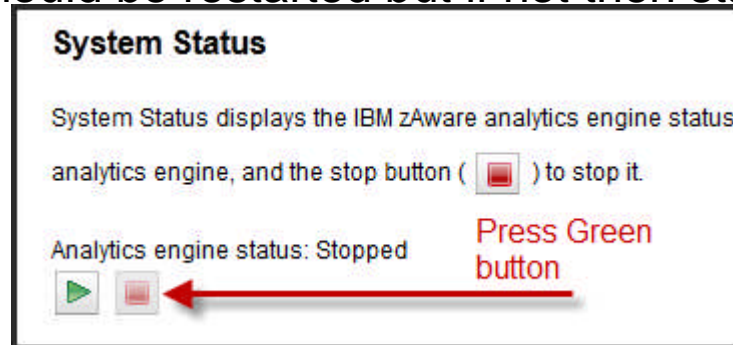
Select one or more systems then select the sysplex where the systems belong then add them

Priming Data panel of IBM zAware GUI



Restart Analytics Engine and Connect Systems

- From the System Status panel of the IBM zAware GUI, restart the analytics engine. Should be restarted but if not then start it.



- Reconnect the z/OS monitored clients which are to be trained
 - SETLOGR FORCE,ZAICONNECT,LSN=SYSPLEX.OPERLOG

Training Period and Training Interval

Configure Settings

Analytics | Data Storage | Security | Sysplex Topology | Priming Data

Instrumentation data retention time (training period - 730 days):
365 days

Training models retention time (0 - 730 days):
365 days

Analysis results retention time (30 - 3650 days):
365 days

Training period (1 - 365 days):
90 days

Training Interval (7 - 365 days):
30 days

Apply Reset

The *training period* is the number of consecutive calendar days that the IBM zAware server uses to identify the instrumentation data to include in training models.

The *training interval* is the number of consecutive calendar days between automatic builds of system behavior models.



Training

- Once bulk data (priming data) for a system is assigned to a sysplex and that monitored client is connected:
- The first training is kicked off based on the first date that is available (loaded in the DB) vs the current date (actually yesterday). If this difference satisfies the training period (default 90 days), then a training attempt will be performed.
- Otherwise the training has to be initiated manually.
- The training interval applies only to re-training after the first model is built successfully.

Training Sets

You can use the Training Sets page provided in IBM zAware to:

- request training for a monitored system
- to display the current training status for each monitored system
- to view the current and future training dates for a monitored system

Training Sets

The Monitored Systems table provides training statuses and results for IBM zAware monitored systems. The Actions menu provides functions for managing model dates, requesting or canceling training, and managing ignored messages. Training details for a given system can be accessed by clicking on links in the Training Progress and Last Training Result columns.

Monitored Systems

System	Sysplex	Training Progress	Last Training Result	Last Training Result Time	Current Model Built
<input type="radio"/> #@\$A	#@\$#PLEX	—	<input type="checkbox"/> Not Trained	—	—

▶ Current Training Status Details (Click on training statuses above to view details)

Last Refresh: Tue Sep 18 2012 13:12:00 GMT-0400 (Eastern Daylight Time)



Training Example - Assign

Configure Settings

Analytics | Data Storage | Security | Sysplex Topology | Priming Data

Priming message data is available for the systems listed below. To assign this data into sysplexes, select systems from the left and add them to sysplexes on the right. Once the data has been added to the desired sysplexes, click the **Assign...** button to confirm and proceed with the assign.

Priming message data by system:

- #@\$2
- #@\$3

Buttons: Add >, Add All >>, < Remove, << Remove All

Assign...

Sysplex Topology

- #@\$#PLEX
- #@\$A

Priming Data for #@\$A has been assigned



Training Example – configure settings

Configure Settings

Analytics | Data Storage | Security | Sysplex Topology | Priming Data

Instrumentation data retention time (training period - 730 days):
365 days

Training models retention time (0 - 730 days):
365 days

Analysis results retention time (30 - 3650 days):
365 days

Training period (1 - 365 days):
365 days

Training Interval (7 - 365 days):
30 days

Apply Reset

The Training Period has been changed from the default (90) to 365.

365 days of data must be available before the first training will be automatically started



Training Example – Manage Model Dates

Training Sets

The Monitored Systems table provides training statuses and results for IBM zAware monitored systems. The Actions menu provides functions for managing model dates, requesting or canceling training, and managing ignored messages. Training details for a given system can be accessed by clicking on links in the Training Progress and Last Training Result columns.

Monitored Systems

Actions	System	Training Progress	Last Training Result	Last Training Result Time	Current Model Built
<ul style="list-style-type: none"> Manage Model Dates Request Training Cancel Training 	ysplex	—	<input type="checkbox"/> Not Trained	—	—

From Training Sets select Manage Model Dates from Actions to get details on

- Next Period Model Dates
- Current Model Dates



Manage Model Dates – Summary View

Training Sets ▶ Manage Model Dates Help

Manage Model Dates

Model dates for the selected system are summarized below. To switch between *Current Model Dates* and *Next Training Period Model Dates*, use the 'Model dates' selector. When displaying *Next Training Period Model Dates*, days can be included or excluded from future training periods days can be excluded or included from future training periods by adding or removing them from the list of excluded days.

Training System: #@\$#PLEX.#@\$A

Model dates: Next Training Period Model Dates Current Model Dates

Today's date (UTC): September 18, 2012

Manual training period begin date (UTC): September 18, 2011

Next training period begin date (UTC): October 30, 2011

Next scheduled training date (UTC): October 30, 2012

Excluded dates: September 18, 2012 Add > < Remove

Excluded Days (UTC)

No days have been excluded.

Unavailable Days (UTC)

- September 18, 2011
- September 19, 2011
- September 20, 2011
- September 21, 2011

Return to Training Sets

This is the summary view. You can toggle a calendar view.

Switch to Calendar View

Drop Down allows you to switch to viewing information on Current Model Dates

Manual training period begin date is based on training period defined by training period in configure setting panel. Earliest date to be used when manually training

Next period training date will be based on next scheduled training date minus the training period

Next scheduled training date is when the next automatic build of a model will occur

Days can be excluded from the model.

Days for which no data is available for modeling is listed

Manage Model Dates – Calendar View

Training Sets > Manage Model Dates

Manage Model Dates

Model dates for the selected system are displayed in the calendar below. To switch between *Current Model Dates* and *Next Training Period Model Dates*, use the 'Model dates' selector. When displaying *Next Training Period Dates*, days can be included or excluded from future training periods by clicking on them.

Training System: #@\$#PLEX.#@\$A

Model dates: Next Training Period Model Dates

Switch to Summary View

Training calendar:

August						
S	M	T	W	T	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

2011 2012 2013

September						
S	M	T	W	T	F	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30						

2011 2012 2013

October						
S	M	T	W	T	F	S
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

2011 2012 2013

Scrollable Calendar View

Color Key

- Excluded days
- Unavailable days
- Today's date (system date)
- Next training period begin date
- Next scheduled training date

Return to Training Sets



Manage Model Dates – Current Model Dates

Manage Model Dates

Model dates for the selected system are summarized below. To switch between *Current Model Dates* and *Next Training Period Model Dates*, use the 'Model dates' selector. When displaying *Next Training Period Model Dates*, days can be included or excluded from future training periods days can be excluded or included from future training periods by adding or removing them from the list of excluded days.

Training System: #@\$#PLEX.#@\$A Model dates:

Today's date (UTC): September 18, 2012

Current model trained date (UTC): — Current model begin date (UTC): —

Excluded Days (UTC)

No days have been excluded.

Unavailable Days (UTC)

No days are unavailable.

This system has never been trained so no information on a the current model is displayed.



Training Example – Request Training



Training Sets

The Monitored Systems table provides training statuses and results for IBM zAware monitored systems. The Actions menu provides functions for managing model dates, requesting or canceling messages. Training details for a given system can be accessed by clicking on links in the Training Progress and Last Training Result columns.

Monitored Systems

Actions	System	Training Progress	Last Training Result	Last Training Result Time	Current Model Built
<ul style="list-style-type: none"> Manage Model Dates <li style="background-color: #0056b3; color: white;">Request Training Cancel Training 	sysplex OXCF01	—	<input checked="" type="checkbox"/> Complete	July 25, 2012 11:29:57 AM Eastern Daylight Time	July 25, 2012 11:29:57 AM Eastern Daylight Time
<input type="radio"/>	SP0 UTCPLXSB	—	<input checked="" type="checkbox"/> Complete	August 23, 2012 8:00:12 PM Eastern Daylight Time	August 23, 2012 8:00:12 PM Eastern Daylight Time
<input type="radio"/>	SQ0 UTCPLXSB	—	<input checked="" type="checkbox"/> Complete	August 23, 2012 8:01:11 PM Eastern Daylight Time	August 23, 2012 8:01:11 PM Eastern Daylight Time
<input type="radio"/>	SC42 WTSCPLX1	—	<input checked="" type="checkbox"/> Complete	August 29, 2012 9:32:22 PM Eastern Daylight Time	August 29, 2012 9:32:22 PM Eastern Daylight Time
<input checked="" type="radio"/>	SC43 WTSCPLX1	—	<input checked="" type="checkbox"/> Complete	August 22, 2012 8:00:06 PM Eastern Daylight Time	August 22, 2012 8:00:06 PM Eastern Daylight Time
<input type="radio"/>	SC66 WTSCPLX1	—	<input checked="" type="checkbox"/> Complete	August 2, 2012 8:00:04 PM Eastern Daylight Time	August 2, 2012 8:00:04 PM Eastern Daylight Time

Select a connect monitored client. Select Request Training. Reply to confirmation popup.



Training Example – Training In Progress

Monitored Systems						
Actions ▾						
System	Sysplex	Training Progress	Last Training Result	Last Training Result Time	Current Model Built	
<input type="radio"/> AQS3	HOXCF01	—	<input checked="" type="checkbox"/> Complete	July 25, 2012 11:29:57 AM Eastern Daylight Time	July 25, 2012 11:29:57 AM Eastern Daylight Time	
<input type="radio"/> SP0	UTCPLXSB	—	<input checked="" type="checkbox"/> Complete	August 23, 2012 8:00:12 PM Eastern Daylight Time	August 23, 2012 8:00:12 PM Eastern Daylight Time	
<input type="radio"/> SQ0	UTCPLXSB	—	<input checked="" type="checkbox"/> Complete	August 23, 2012 8:01:11 PM Eastern Daylight Time	August 23, 2012 8:01:11 PM Eastern Daylight Time	
<input type="radio"/> SC42	WTSCPLX1	—	<input checked="" type="checkbox"/> Complete	August 29, 2012 9:32:22 PM Eastern Daylight Time	August 29, 2012 9:32:22 PM Eastern Daylight Time	
<input type="radio"/> SC43	WTSCPLX1	<input checked="" type="checkbox"/> In Progress	—	—	August 22, 2012 8:00:06 PM Eastern Daylight Time	
<input type="radio"/> SC66	WTSCPLX1	—	<input checked="" type="checkbox"/> Complete	August 2, 2012 8:00:04 PM Eastern Daylight Time	August 2, 2012 8:00:04 PM Eastern Daylight Time	
<input type="radio"/> SC66	WTSCPLX1	—	<input checked="" type="checkbox"/> Complete	August 2, 2012 8:00:04 PM Eastern Daylight Time	August 2, 2012 8:00:04 PM Eastern Daylight Time	

► Current Training Status Details (Click on training statuses above to view details)

Last Refresh: Tue Sep 18 2012 14:29:20 GMT-0400 (Eastern Daylight Time)

Training in Progress the will switch to Complete Status



BEGIN IBM zAware customization DEMO

- **Begin DEMO**
 - Assume that LPAR definition has been completed and IOCDS activated
 - Configure IBM zAware LPAR definitions
 - Configure IBM zAware GUI
 - Storage & Security
 - Configure and connect a monitored client
 - Run Bulk Load
 - Assign Systems
 - Train a system
 - Look at analysis screen

Runtime Diagnostics



Runtime Diagnostics

- Analyze system to provide diagnostic data in a timely manner
- Performs analysis similar to an experienced Systems Programmer
 - Faster – goal is 60 seconds or less
 - More comprehensive
 - Looks for specific evidence of “soft failures”
 - Provide suggested next steps
- It is not:
 - Not automations or a monitor
 - Takes no corrective action, but recommends next steps
 - Minimal dependencies on system services

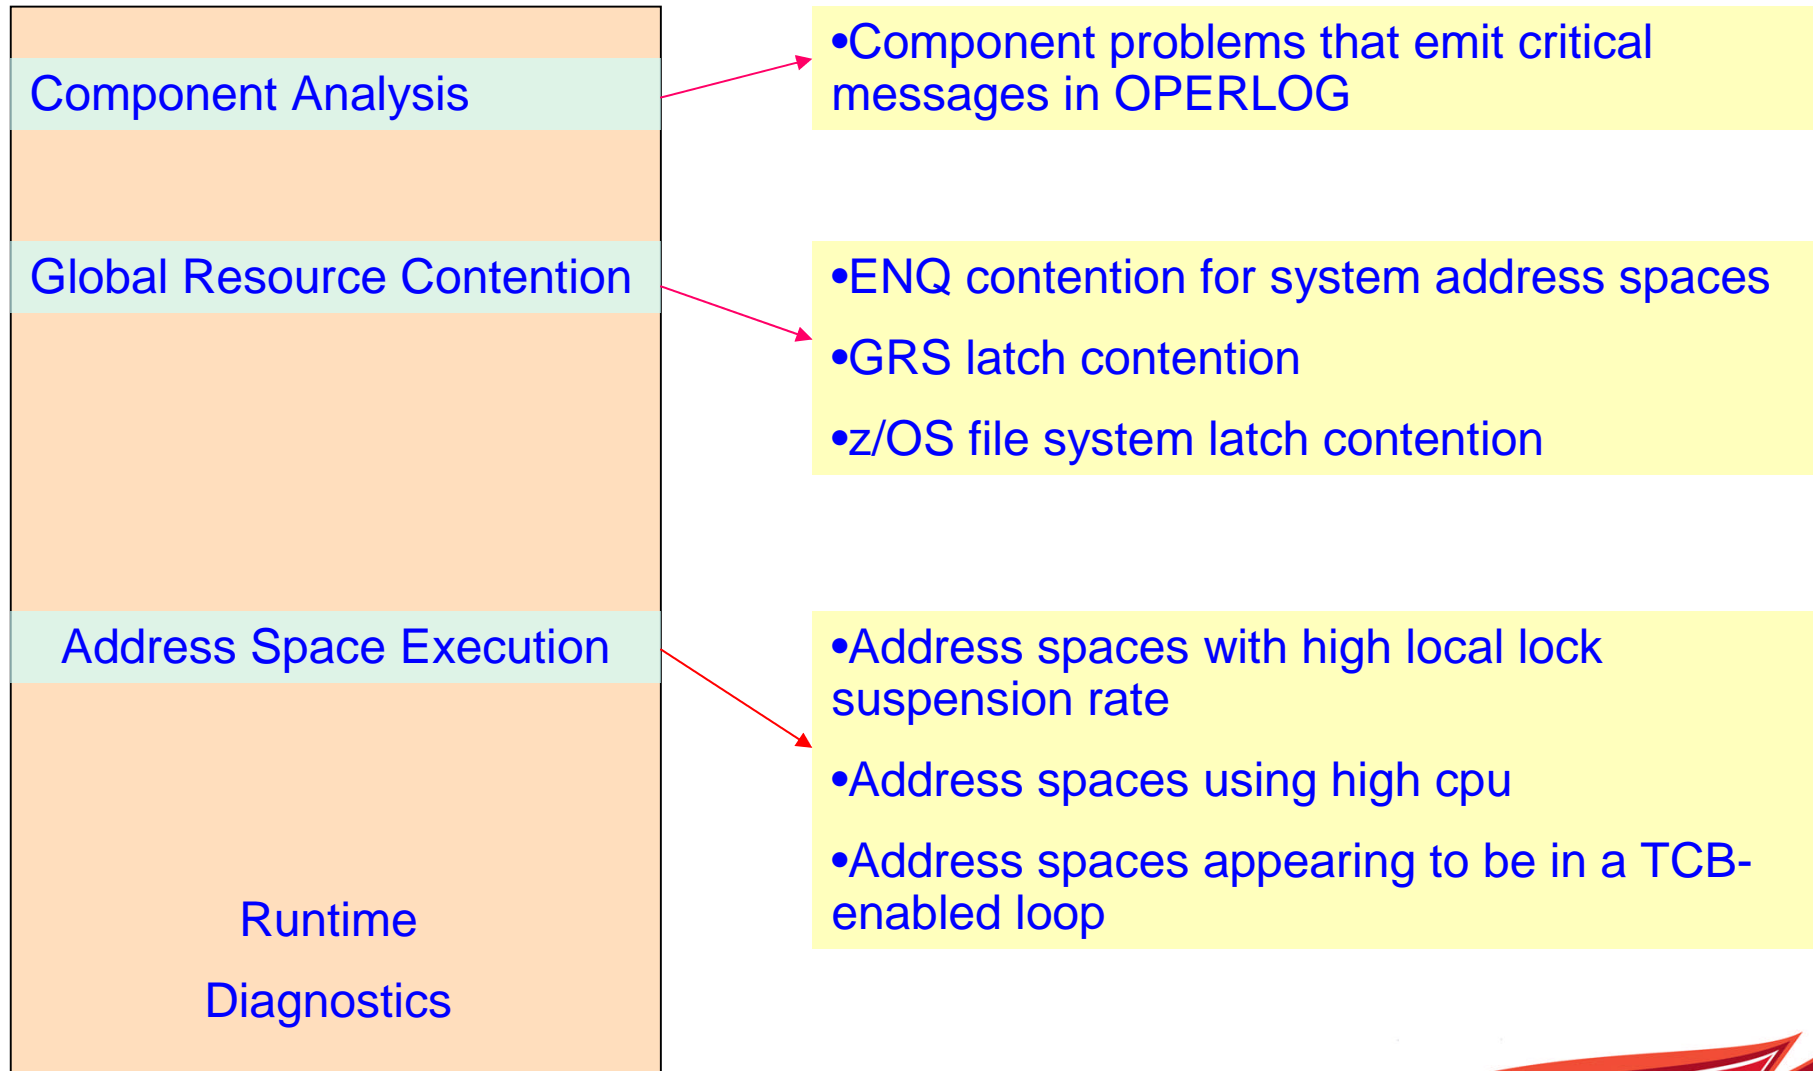
Benefits

- Reduces the skill level needed by a system programmer for investigating soft failures
 - Provide timely , comprehensive analysis when needed
 - Productivity aid for experienced system programmers

- Quickly discover next actions to take
 - Which jobs to cancel
 - What to investigate further using other tools

- Use when
 - Problem reported on system
 - For problem resolution calls
 - When Predictive Failure Analysis reports a problem

Categories for diagnostics



z/OS R12

z/OS

```
Runtime Diagnostics
START HZR,SUB=MSTR

START
HZR,SUB=MSTR,OPTIONS=(SYSNAME=SYSB)

HZR PROC
PGM=HZRMAIN

HZROUT DD DUMMY
or
HZROUT DD <sequential dataset>
```

Start and run Analysis

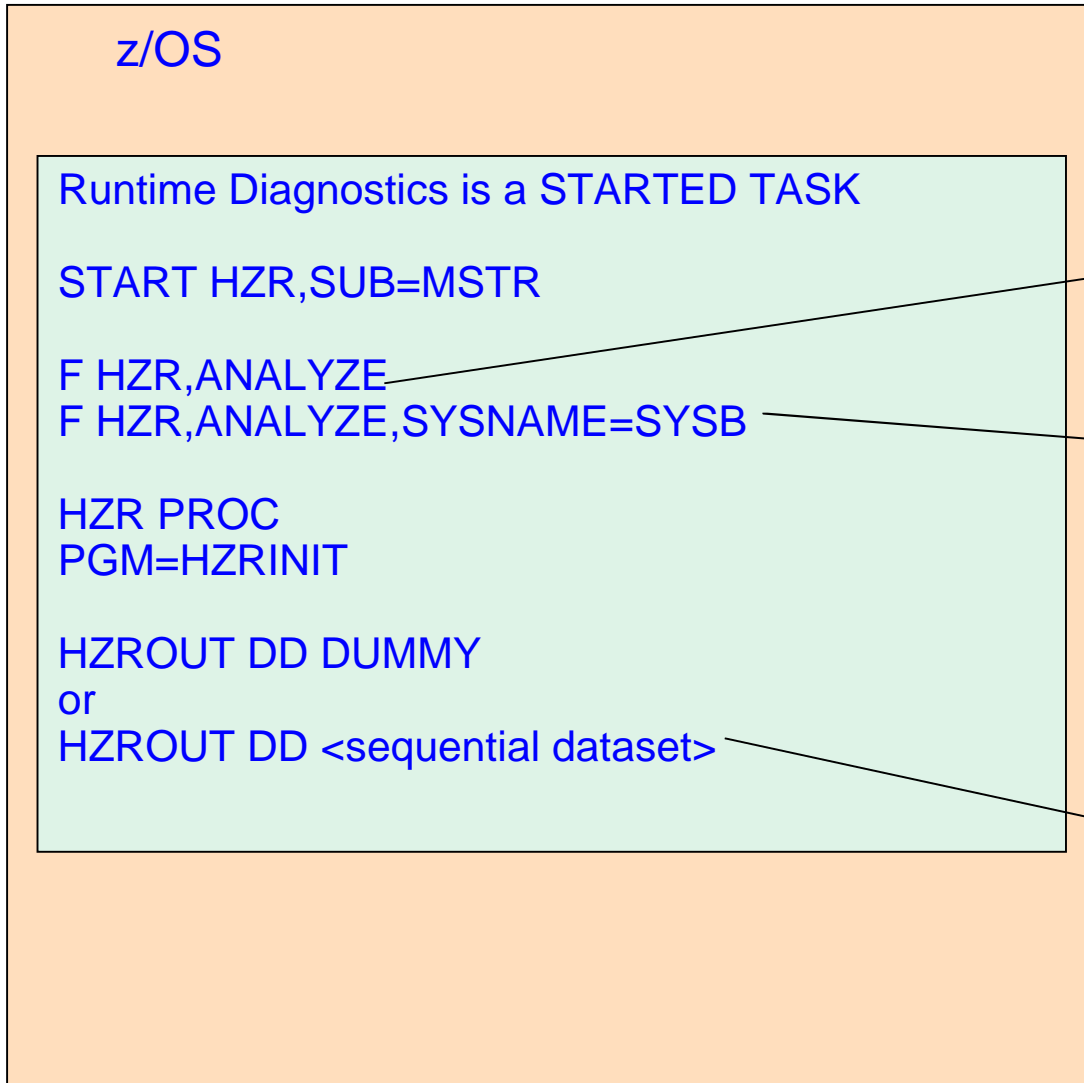
Can specify another system but only ENQ and OPERLOG scans done

z/OS R12 pgm is HZRMAIN

Output can be written to physical sequential dataset; Output goes to console and hardcopy log



z/OS R13



Start Address Space

Modify command to start analysis

Can specify another system but only ENQ and OPERLOG scans done

z/OS R13 pgm is HZRINIT

Output can be written to physical sequential dataset
Output goes to console and hardcopy log



Runtime Diagnostics output

- The output of Runtime Diagnostics is a WTO to the operator console (actually, a multiline WTO) that is issued to the console that originally issued the **F HZR,ANALYZE** command.
- If the MCS console that issued the **F HZR** command has an out-of-line display area setup (through a K A,xx) then the output will be displayed in the display area.
- The output of Runtime Diagnostics can also be directed to a sequential data set. DISP=SHR to view output without stopping HZR started task.

Target System Analysis – No Events Found

Analyze the system where Runtime Diagnostics has started – “All is well”

```
F HZR,ANALYZE
HZR0201I RUNTIME DIAGNOSTICS SUCCESS. TIME (2012/09/25 - 12:55:19).
NO EVENTS WERE FOUND FOR SYSTEM: #@$A
```

Analyze another system in sysplex – only ENQ and OPERLOG scan

```
F HZR,ANALYZE,SYSNAME=#@$2
HZR0200I RUNTIME DIAGNOSTICS RESULT 830
SUMMARY: SUCCESS - NO EVENTS FOUND
REQ: 045 TARGET SYSTEM: #@$2      HOME: #@$A      2012/09/25 -
13:07:37
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 00 - PRIORITIES: HIGH:00  MED:00  LOW:00
PROCESSING BYPASSED:
OMVS.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
LATCHES....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
LOOP.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
HIGHCPU....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
LOCK.....SPECIFIED TARGET SYSTEM IS NOT THE HOME SYSTEM.
```


Processing failure messages

- Runtime Diagnostics also reports when part of its processing fails (that is, it is unable to complete processing for one or more events) as QUALIFIED SUCCESS in the SUMMARY: portion of the report

f hzr,analyze

HZR0200I RUNTIME DIAGNOSTICS RESULT 751

SUMMARY: QUALIFIED SUCCESS - SOME PROCESSING FAILED

REQ: 001 TARGET SYSTEM: SY1 HOME: SY1 2010/12/21 - 11:25:55

INTERVAL: 60 MINUTES

EVENTS:

FOUND: 01 - PRIORITIES: HIGH:01 MED:00 LOW:00

TYPES: HIGHCPU:01

PROCESSING FAILURES:

OPERLOG...IXGCONN REQ=CONNECT ERROR.....RC=00000008 RS=0000080B

EVENT 01: HIGH - HIGHCPU - SYSTEM: SY1 2010/12/21 - 11:25:56

ASID CPU RATE:99% ASID:002E JOBNAME:IBMUSERX

STEPNAME:STEP1 PROCSTEP: JOBID:JOB00045 USERID:IBMUSER

JOBSTART:2010/12/21 - 11:22:51

ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.

ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

In this case, although Runtime Diagnostics was unable to connect to OPERLOG to examine messages, it continues to find other soft failures (HIGHCPU)

High CPU Analysis

```
HZR0200I RUNTIME DIAGNOSTICS RESULT 568
SUMMARY: SUCCESS
REQ: 003 TARGET SYSTEM: SY1 HOME: SY1 2010/12/21 - 13:45:49
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 01 - PRIORITIES: HIGH:01 MED:00 LOW:00
TYPES: HIGHCPU:01
-----
-
EVENT 01: HIGH - HIGHCPU - SYSTEM: SY1 2010/12/21 - 13:45:50
ASID CPU RATE:99% ASID:002E JOBNAME:IBMUSERX
STEPNAME:STEP1 PROCSTEP: JOBID:JOB00045 USERID:IBMUSER
JOBSTART:2010/12/21 - 11:22:51
ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
```

- Point-in-time check of any address space that is using more than 95% of the capacity of a single CPU. Takes 2 samples over a 1 second interval.
- The analysis is a one-second sample interval based on the capacity of a single CPU within the LPAR. Be aware that it is possible for the usage to be reported greater than 100% if the address space has multiple TCBs and several of the TCBs are individually using a high percentage of the capacity of a CPU.

Loop Detection

```
HZR0200I RUNTIME DIAGNOSTICS RESULT 581
SUMMARY: SUCCESS
REQ: 004 TARGET SYSTEM: SY1 HOME: SY1 2010/12/21 - 13:51:32
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 02 - PRIORITIES: HIGH:02 MED:00 LOW:00
TYPES: HIGHCPU:01
TYPES: LOOP:01
-----
-
EVENT 01: HIGH - HIGHCPU - SYSTEM: SY1 2010/12/21 - 13:51:33
ASID CPU RATE:99% ASID:002E JOBNAME:IBMUSERX
STEPNAME:STEP1 PROCSTEP: JOBID:JOB00045 USERID:IBMUSER
JOBSTART:2010/12/21 - 11:22:51
ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
-
EVENT 02: HIGH - LOOP - SYSTEM: SY1 2010/12/21 - 13:51:14
ASID:002E JOBNAME:IBMUSERX TCB:004FF1C0
STEPNAME:STEP1 PROCSTEP: JOBID:JOB00045 USERID:IBMUSER
JOBSTART:2010/12/21 - 11:22:51
ERROR: ADDRESS SPACE MIGHT BE IN A LOOP.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----
-
```

- Investigates all tasks in all address spaces looking for TCB loops

- Takes a snapshot of the system trace

- Looks for consistent, repetitive activity that typically indicates a loop

Runtime Diagnostics looks through all tasks in all address spaces to determine whether a task might be looping. Runtime Diagnostics does this by examining various system information for indicators of consistent repetitive activity that are typical when a task is in a loop.

When both a HIGHCPU event and a LOOP event list the job name, there is a high probability that a task in the job is in a loop. The normal corrective action is to cancel the job name listed.

Local Lock suspension

```
HZR0200I RUNTIME DIAGNOSTICS RESULT 581
SUMMARY: SUCCESS
REQ: 004 TARGET SYSTEM: SY1 HOME: SY1 2010/12/21 - 13:51:32
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 01 - PRIORITIES: HIGH:01 MED:00 LOW:00
TYPES: LOCK:01
-----
EVENT 01: HIGH - LOCK - SYSTEM: SY1 2010/12/21 - 13:51:33
HIGH LOCAL LOCK SUSPENSION RATE - ASID:000A JOBNAME:WLM
STEPNAME:WLM PROCSTEP:IEFPROC JOBID:+++++++ USERID:+++++++
JOBSTART:2010/12/21 - 11:15:08
ERROR: ADDRESS SPACE HAS HIGH LOCAL LOCK SUSPENSION RATE.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
```

Suspend:

Local
Lock

Cross
Memory
Lock

CMS lock

- Runtime Diagnostics provides a point-in-time check of local lock suspension for any address space.
- For the local lock suspension, Runtime Diagnostics calculates the amount of time an address space is suspended waiting for the local lock. If an address is suspended more than 50% of the time waiting for a local lock,

ENQ contention checking

```
HZR0200I RUNTIME DIAGNOSTICS RESULT 568
-----
EVENT 01: HIGH - ENQ - SYSTEM: SY1 2010/12/21 - 11:43:33
ENQ WAITER - ASID:0038 - JOBNAME:IBMUSER2 - SYSTEM:SY1
ENQ BLOCKER - ASID:002F - JOBNAME:IBMUSER1 - SYSTEM:SY1
QNAME: TESTENQ
RNAME: TESTOFAVERYVERYVERYVERYLOOOOOOOOOOOOOOOOOOOOONGRNAME1234567...
ERROR: ADDRESS SPACES MIGHT BE IN ENQ CONTENTION.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE BLOCKING JOBS AND
ACTION: ASIDS.
-----
```

- ENQ contention equivalent to issuing the D GRS,AN,WAITER command
- Lists both waiter and blocker
- It compares the list of job names that are waiters with the list of system address spaces that are started at IPL to determine if any system address spaces are waiters. Looks for system address spaces that is in ENQ 'waiter' for > 5 seconds

z/OS UNIX latch contention – z/OS R13

```
EVENT 01: HIGH - OMVS - SYSTEM: SY1 2010/07/07 - 13:07:32
ASID:000E - JOBNAME:OMVS
MOUNT LATCH WAITERS: 2
FILE SYSTEM LATCH WAITERS: 0
XSYS AND OTHER THREADS WAITING FOR z/OS UNIX: 3
ERROR: z/OS UNIX MIGHT HAVE FILE SYSTEM LATCH CONTENTION.
ACTION: ISSUE D OMVS,W,A TO INVESTIGATE z/OS UNIX FILE SYSTEM
ACTION: LATCH CONTENTION, ACTIVITY AND WAITING THREADS.
```

- If z/OS UNIX file system latch contention or waiting threads exist for > 5 minutes in z/OS UNIX, a Runtime Diagnostics OMVS event is created.
- Normal action is to issue D OMVS,W,A to get the ASID and job names of the waiters

GRS latch contention – z/OS R13

```
-----  
EVENT 01: HIGH - LATCH - SYSTEM: SY1 2010/12/21 - 14:32:01  
LATCH SET NAME: SYSTEST.LATCH_TESTSET  
LATCH NUMBER:3 CASID:0039 CJOBNAME:TSTLATCH  
TOP WAITER - ASID:0039 - JOBNAME:TSTLATCH - TCB/WEB:004E2A70  
TOP BLOCKER- ASID:0039 - JOBNAME:TSTLATCH - TCB/WEB:004FF028  
ERROR: ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.  
ACTION: D GRS,AN,LATCH,DEP,CASID=0039,LAT=(SYSTEST.L*,3),DET  
ACTION: TO ANALYZE THE LATCH DEPENDENCIES. USE YOUR SOFTWARE  
ACTION: MONITORS TO INVESTIGATE BLOCKING JOBS AND ASIDS.  
-----  
EVENT 02: HIGH - LATCH - SYSTEM: SY1 2010/12/21 - 14:32:01  
LATCH SET NAME: SYSTEST.LATCH_TESTSET  
LATCH NUMBER:3 CASID:003B CJOBNAME:TSTLATC2  
TOP WAITER - ASID:003B - JOBNAME:TSTLATC2 - TCB/WEB:004E2A70  
TOP BLOCKER- ASID:003B - JOBNAME:TSTLATC2 - TCB/WEB:004FF028  
ERROR: ADDRESS SPACES MIGHT BE IN LATCH CONTENTION.  
ACTION: D GRS,AN,LATCH,DEP,CASID=003B,LAT=(SYSTEST.L*,3),DET  
ACTION: TO ANALYZE THE LATCH DEPENDENCIES. USE YOUR SOFTWARE  
ACTION: MONITORS TO INVESTIGATE BLOCKING JOBS AND ASIDS.
```

- Obtains Latch contention information from GRS
- Omits z/OS UNIX file system latch contention
- Returns the longest waiter for each latch set

Critical Messages

```

HZR0200I RUNTIME DIAGNOSTICS RESULT 361
SUMMARY: SUCCESS
REQ: 120 TARGET SYSTEM: #@$A      HOME: #@$A      2012/09/25 - 15:26:52
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 01 - PRIORITIES: HIGH:01  MED:00  LOW:00
TYPES: CF:01
    
```

```

-----
EVENT 01: HIGH - CF                - SYSTEM: #@$A      2012/09/25 - 15:10:59
IXL158I PATH B8 IS NOW NOT-OPERATIONAL TO CUID: FFD4
                COUPLING FACILITY 002817.IBM.02.0000000B3BD5
                PARTITION: 2E  CPCID: 00

ERROR: INDICATED CHANNEL NOT OPERATIONAL.
ACTION: RUN EREP TO DUMP DATA FROM SYS1.LOGREC AND PROVIDE IT TO IBM
ACTION: SUPPORT.
-----
    
```

- Checks previous 1 hour of OPERLOG
- For some messages additional analysis done
- Groups related messages into single event
- Weeds out shortage and relieved messages
- In some cases only shows last message if message repeated

Additional analysis done for:

IXC101I, IXC105I, IXC418I

IXL013I

IXC431I

IXC246E

IXC585E



Runtime Diagnostics Summary

- **z/OS 1.12**
- **Component-specific, critical messages in OPERLOG**
 - Looks one hour back, if available
 - Additional analysis for some msgs
 - Message summary found in output
 - Can analyze messages in other systems in sysplex
- **Enqueue Contention Checking**
 - Looks for system address space waiting > 5 seconds
 - Lists both waiter and blocker
 - Can detect contention in other system in sysplex
- **Local Lock Suspension**
 - Any address space whose local lock suspension time is > 50%
- **z/OS 1.12 (continued)**
- **CPU Analysis**
 - Takes 2 samples over 1 sec. interval
 - Any task using > 95% is considered a potential problem
- **Loop Detection**
 - Investigates all tasks in all address spaces looking for TCB loops
- **z/OS 1.13**
- **z/OS UNIX Latch Contention**
 - Looks for z/OS UNIX latch contention or waiting threads that exit for > 5 minutes.
- **GRS Latch Contention**
 - Obtains latch contention info from GRS
 - Omits z/OS UNIX file system latch contention
 - Returns longest waiter for each latch set

IBM Health Checker for z/OS



Soft Failure Avoidance with IBM Health Checker

Health checker's role is to keep subtle configuration errors from resulting in Soft Failures

- Performance
 - System effects
 - Check configuration for best practices
 - Single points of failure for log structures, data sets, CDS
 - Storage utilization, running out of resources
 - How many ASIDs do I have left? LXs? When will I run out?
 - Whether DAE is inactive
 - VSAM RLS latch contention, CF Cache size, CDS SPOF, etc.
 - System Logger structure usage
 - I/O timing, protection
 - ... many others
- **Used by Preventive Failure Analysis to emit alerts**
 - Warnings of detected soft failures
 - 187 z/OS Health Checks in z/OS R13
 - Just a reminder to enable IBM Health Checker for z/OS

z/OS Health Checker – History & Objectives

■ History

- Multi-system outage analysis
- 15-20% system outages attributed to Setup/Configuration
- A tool developed by ITSO to address component configuration and setup errors commonly made by installations
- Originally available as web download tool,
- Is available with its own FMID since z/OS 1.7
- Migration checks included since z/OS 1.9

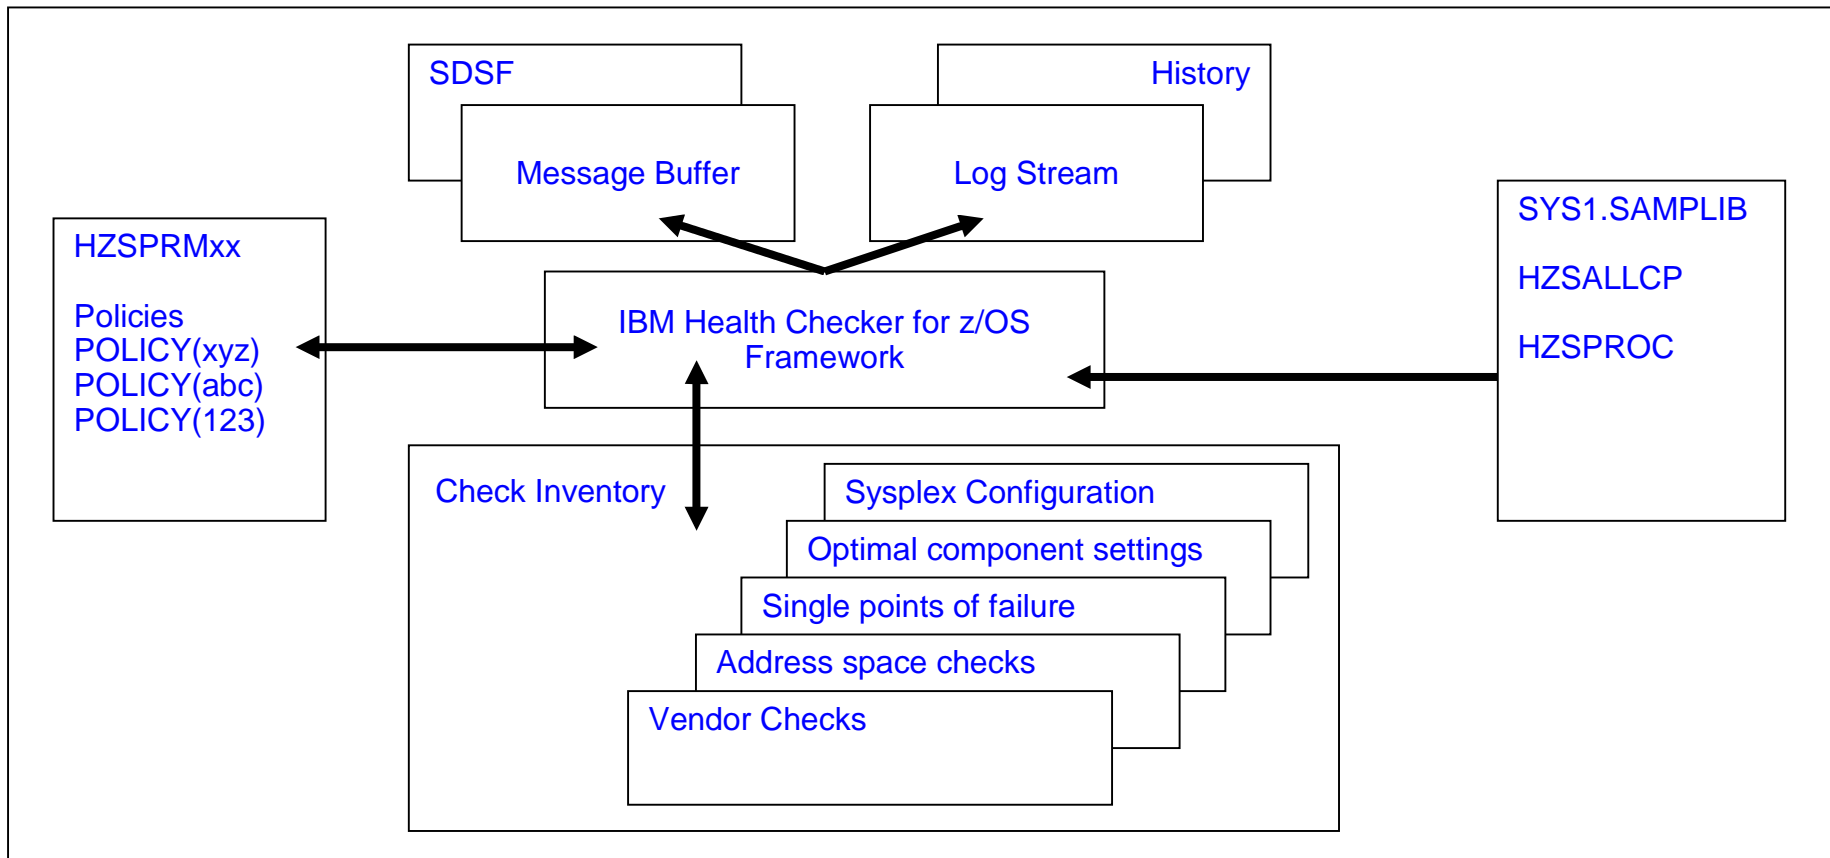
■ Objectives

- Identify potential problems before they impact availability or cause outages
- Checks the current active z/OS and sysplex settings
- Checks definitions for a system and compares the values to those suggested by IBM or defined by you
- Run continuously to find deviations from best practices
- Not a diagnostic or monitoring tool

■ Produces output in the form of detailed messages (SDSF support)

- Indicates both potential problems and suggested actions
- Does not mean that Health Checker has found problems that you need to report to IBM!
- Messages simply informs you of potential problems so that you can take action on your installation, before they become critical, thereby avoiding an IPL

IBM Health Checker for z/OS



The *framework* is an interface that manages services like check registration, messaging, scheduling, command processing, logging, and reporting.

Checks are programs or routines that evaluate component, element, or product-specific settings and definitions, looking for potential problems on a running system. Checks are independent of the framework.

- Not new:
- Base function with z/OS 1.7 onwards
- Web download + PTFs for z/OS 1.4 – 1.6

- But:
- Can help avoid soft failures
- Over 180 checks



z/OS Health Checker - Overview

- Each check has 3 parts
 - The dynamic exit routine that identifies the check to the Health checker
 - The check itself
 - A message table to define messages that are issued by the check

- Each check includes a set of pre-defined values, such as:
 - Interval, or how often the check will run
 - Severity of the check, which influences how check output is issued - (high, medium, and low)
 - Routing and descriptor codes for the check

- You can override some check values:
 - Statements in the HZSPRMxx parmlib member
 - The MODIFY command - F HZSPROC,.....
 - SDSF interactive command panel

Important Considerations

1. Don't just change the configuration ... investigate the exception and then take appropriate action
2. There are 187 Health Checks in z/OS R13
 1. Don't think that you must activate all health checks at once to get benefit
 2. Goal should be to remove all exceptions
 1. by fixing the condition
 2. by tuning the check so that it looks for what you need it to look for
 3. (as a last resort) by deactivating the check
 3. Once you can run cleanly, you will be in the ideal position to know that an exception indicates something has changed
 4. Consider categorizing health checks by
 1. Checks I expect no exceptions from
 2. Checks not turned on because exceptions not cleaned up yet
 3. Plan to move checks to group 1 as you clean up exceptions
3. GDPS recommendations for changing z/OS checks trump z/OS in a GDPS environment
 1. Some z/OS Health Check recommendations conflict with GDPS function, so follow GDPS guidelines



Resources

- IBM Health Checker's User Guide, SA22-7994
- Exploiting the IBM Health Checker for z/OS Infrastructure
 - IBM Redpaper REDP-4590-01



Predictive Failure Analysis



Predictive Failure Analysis (PFA)

- The goal of predictive analysis and early-detection analysis is to notify the system programmer when the system can see the problem is occurring internally rather than later when it is visible externally.
- PFA detects problems before they are visible externally by using resources and metrics at different layers of the software stack that can indicate that resource exhaustion, damaged address spaces, or damaged systems could be occurring.
- PFA is not intended to find immediate problems (on a machine-time scale) that will bring the system down. Rather, it can detect potential problems on a human-time scale.

How PFA Detects Soft Failures

- Causes of “sick, but not dead”
 - **Damaged systems**
 - Recurring or recursive errors caused by software defects anywhere in the software stack
 - Serialization
 - Priority inversion
 - Classic deadlocks
 - Owner gone
 - **Resource exhaustion**
 - Physical resources
 - Software resources
 - Indeterminate or unexpected states
- Predictive failure analysis uses
 - *Historical data*
 - *Machine learning and mathematical modeling*to detect abnormal behavior and the potential causes of this abnormal behavior
- Objective
 - Convert “sick, but not dead” to a correctable incident

PFA focus

- PFA focuses on the **damaged address space or system** and **resource exhaustion** categories that may lead to soft failures.
 - Damaged address space or system:** The indication of a damaged system is typically when there are recurring or recursive errors anywhere in the software stack.
 - Physical or software resource exhaustion** of a shared system resource
- PFA collects the data from the individual system and models it to determine what is normal for that system as well as to determine what kinds of resource trends are occurring.

Abnormal behavior detection

- There are three types of abnormal behavior detection that PFA's algorithms incorporate:
 - **Future prediction:** This processing does trend analysis and models the behavior **into the future** to predict if the current trend will exhaust a common resource.
 - **Expected value:** This processing does trend analysis and models the behavior to determine **what value should be expected at the current time** to determine if an abnormal condition is occurring.
 - **Expected rate:** This processing does trend analysis and models the behavior to determine **if the current rate when compared to the rates for multiple time periods** indicates that an abnormal condition is occurring. The rate is often created by normalizing a count of the metric being analyzed by the CPU being used. By normalizing the rate and by comparing against multiple time period predictions, normal workload changes do not appear as abnormal conditions.
- For both the expected value and the expected rate types of predictions, PFA clusters the historical data so that trends within the data can be identified. It then determines which trend is currently active and uses the prediction for that trend in its comparisons.


How PFA Determines Expected Values

- **Behavior of z/OS system is a function of many factors such as**
 - Workload, type of work, hardware and software configuration, system automation, etc....

- **Use historical data to calculate future or expected value to eliminate factors**

- **Same type of work runs at approximately same time or runs consistently**
 - **Expected rate = $fn(\text{workload, time})$** -- Can compare different time ranges such as 1 hour ago, 24 hours ago, 7 days ago
 - **Expected value = $fn(\text{workload})$** – based on past and current trends
 - **Future prediction = $fn(\text{workload, time projected into the future})$**

- **Cluster metric by time to calculate expected or future value**

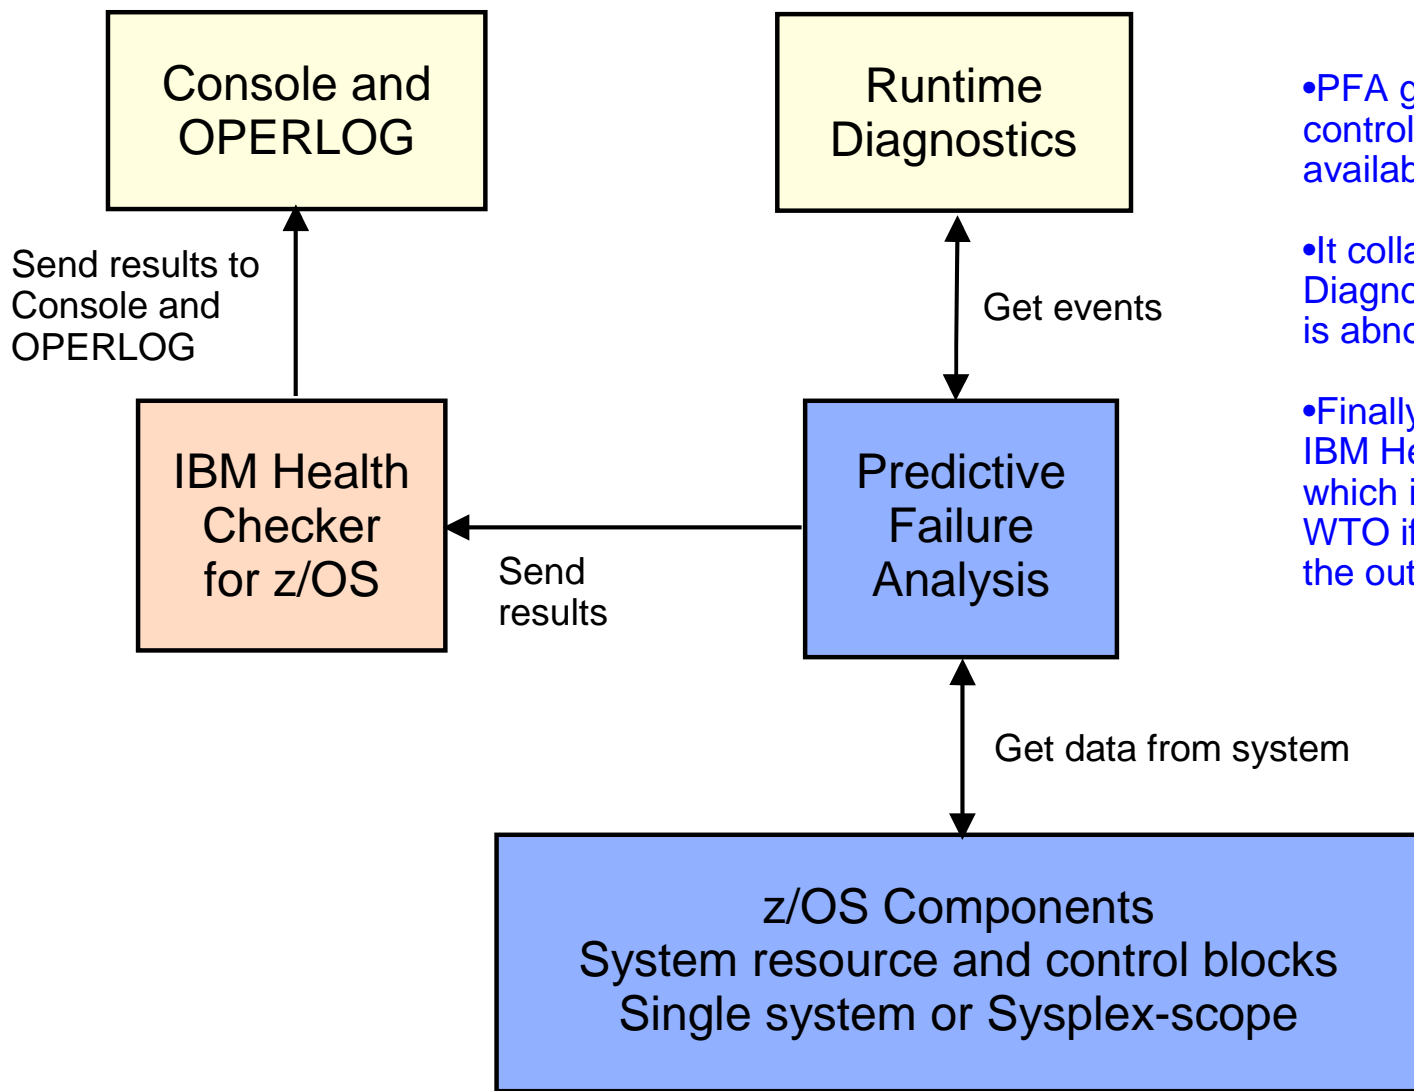


Abnormal Behavior

Hourly Logrec Arrival Rate — 30 per. Mov. Avg. (Hourly Logrec Arrival Rate)



PFA input and output



- PFA gets its data from z/OS control blocks and externally available interfaces.

- It collaborates with Runtime Diagnostics to detect if a metric is abnormally low.

- Finally, it sends the results to IBM Health Checker for z/OS which issues the message as a WTO if necessary and makes the output viewable in SDSF.





PFA Serviceability

Modify command to display status

STATUS examples:

```
f pfa,display
f,pfa,display,status
```

```
AIR017I 10.31.32 PFA STATUS
NUMBER OF CHECKS REGISTERED      : 5
NUMBER OF CHECKS ACTIVE          : 5
COUNT OF COLLECT QUEUE ELEMENTS: 0
COUNT OF MODEL QUEUE ELEMENTS   : 0
COUNT OF JVM TERMINATIONS       : 0
```

SUMMARY examples:

```
f pfa,display,checks
f pfa,display,check(pfa*),summary
```

```
AIR013I 10.09.14 PFA CHECK SUMMARY
```

CHECK NAME	ACTIVE	LAST SUCCESSFUL COLLECT TIME	LAST SUCCESSFUL MODEL TIME
PFA_COMMON_STORAGE_USAGE	YES	04/05/2008 10.01	04/05/2008 08.16
PFA_LOGREC_ARRIVAL_RATE	YES	04/05/2008 09.15	04/05/2008 06.32

(all checks are displayed)

DETAIL examples:

```
f pfa,display,check(pfa_logrec_arrival_rate),detail
f pfa,display,check(pfa_*),detail
```

```
AIR018I 02.22.54 PFA CHECK DETAIL
CHECK NAME: PFA_LOGREC_ARRIVAL_RATE
ACTIVE : YES
TOTAL COLLECTION COUNT : 5
SUCCESSFUL COLLECTION COUNT : 5
LAST COLLECTION TIME : 04/05/2008 10.18.22
LAST SUCCESSFUL COLLECTION TIME: 04/05/2008 10.18.22
NEXT COLLECTION TIME : 04/05/2008 10.33.22
TOTAL MODEL COUNT : 1
SUCCESSFUL MODEL COUNT : 1
LAST MODEL TIME : 04/05/2008 10.18.24
LAST SUCCESSFUL MODEL TIME : 04/05/2008 10.18.24
NEXT MODEL TIME : 04/05/2008 16.18.24
CHECK SPECIFIC PARAMETERS:
COLLECTINT : 15
MODELINT : 360
COLLECTINACTIVE : 1=ON
DEBUG : 0=OFF
STDDEV : 10
EXCEPTIONMIN : 25
EXCLUDED_JOBS:
(excluded jobs list here)
```

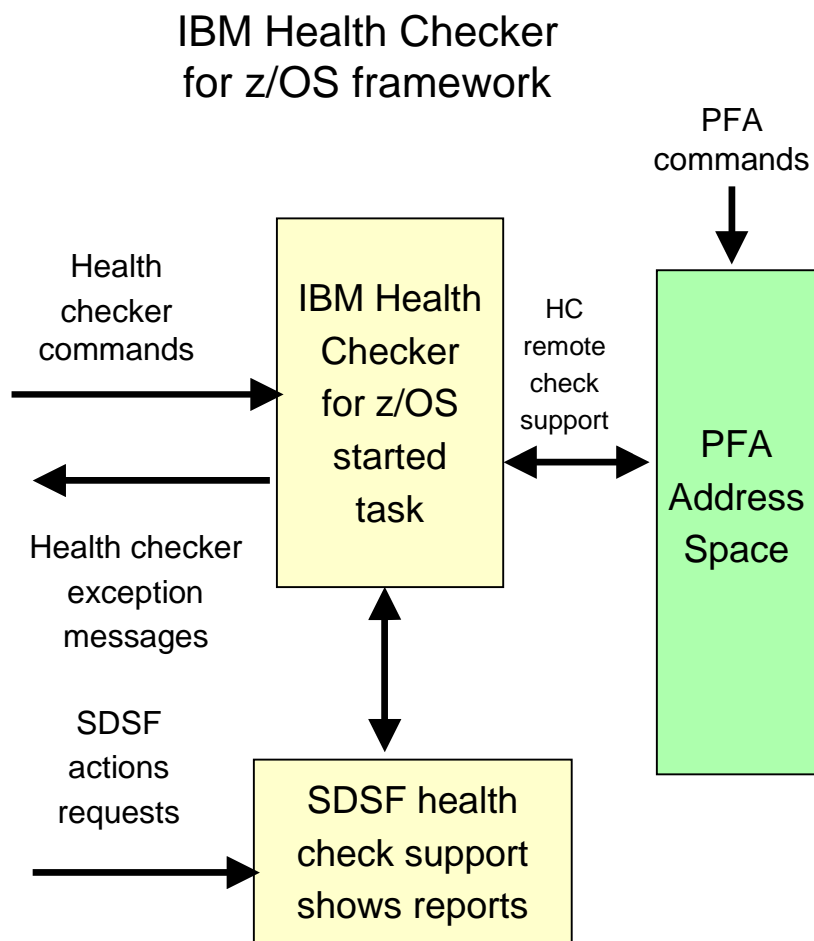

The PFA checks

Resource or Metric	Causes detected	Type of Analysis	Release
Common Storage usage	Detects resource exhaustion in common storage by looking for spikes, leaks and creeps	Future prediction	1.10 SPE
LOGREC arrival rate	Detects a damaged address space or system	Expected rate	1.10 SPE
Message arrival rate	Detects a damaged or hung address space based on abnormal rate in the WTOs and WTORs issued, normalized by the amount of CPU being used	Expected rate	1.11
SMF arrival rate	Detects a damaged or hung address space based on an abnormal rate of SMF arrivals, normalized by the amount of CPU being used	Expected rate	1.12
JES spool usage	Detects a damaged persistent address space based on an abnormal increase in the number of track groups used	Expected value	1.13
Enqueue request rate	Detects a damaged or hung address space or system based on an abnormal rate of enqueue requests, normalized by the amount of CPU being used	Expected rate	1.13

PFA_FRAMES_AND_SLOTS_USAGE

- If you are familiar with PFA from other sources, you may have noticed that the check for frames and slots usage (PFA_FRAMES_AND_SLOTS_USAGE) was omitted from the previous table.
- This check has been permanently removed from PFA with APAR OA40065 due to the fact that it caused unwarranted exceptions that could not be avoided with available mechanisms.
- It is recommended that you apply the PTF for APAR OA40065 as soon as possible. To remove the check from PFA prior to applying the PTF, use the following IBM Health Checker for z/OS command:
- **f hzsproc,delete,check(ibmpfa,pfa_f*)**

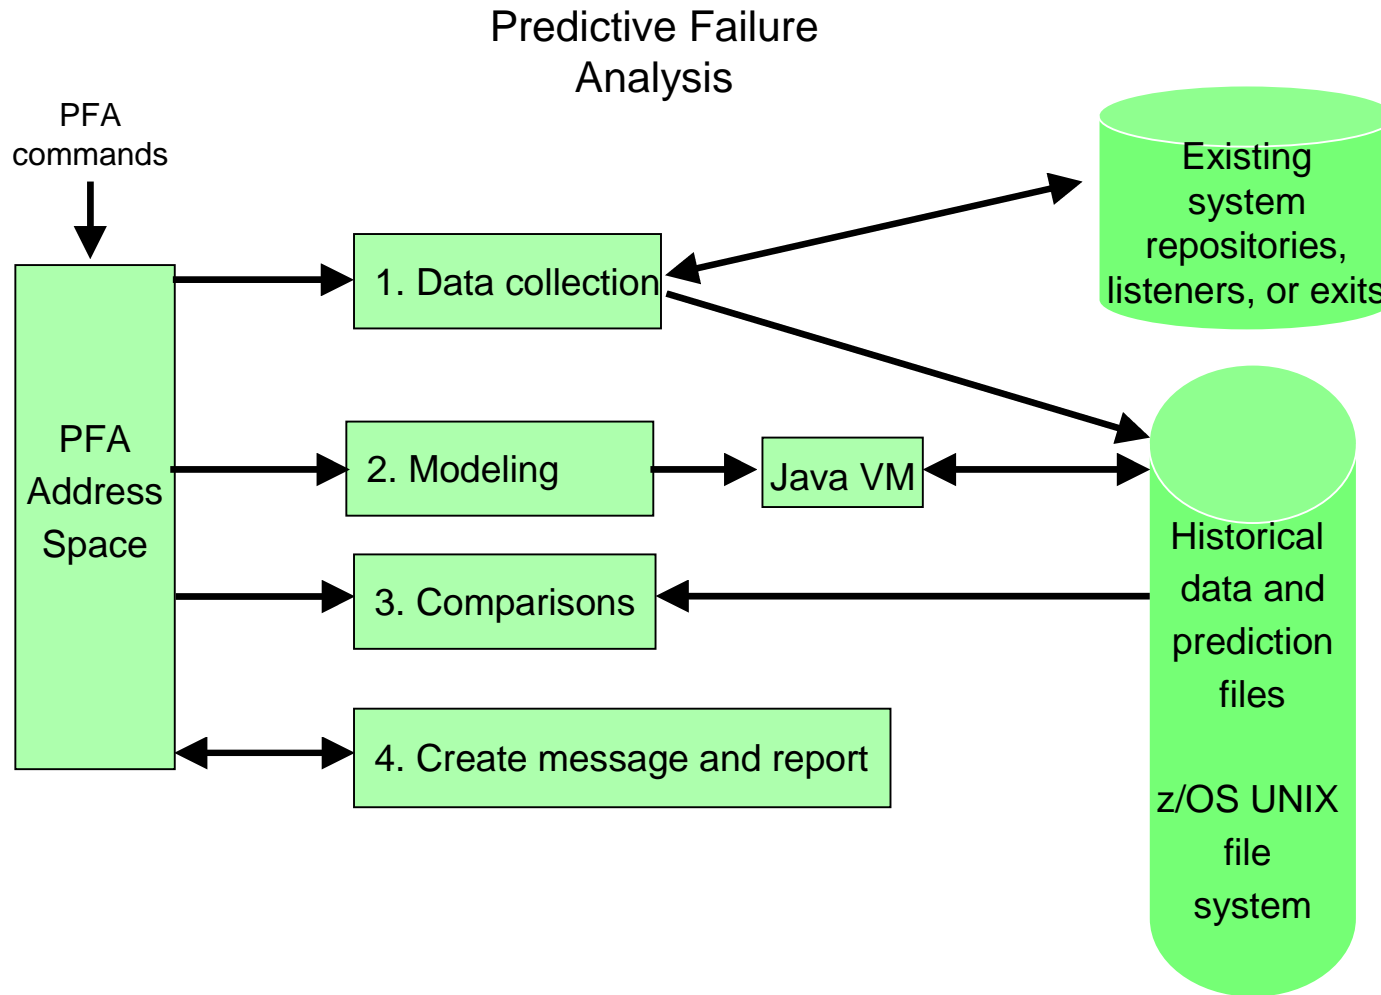
PFA and IBM Health Checker for z/OS



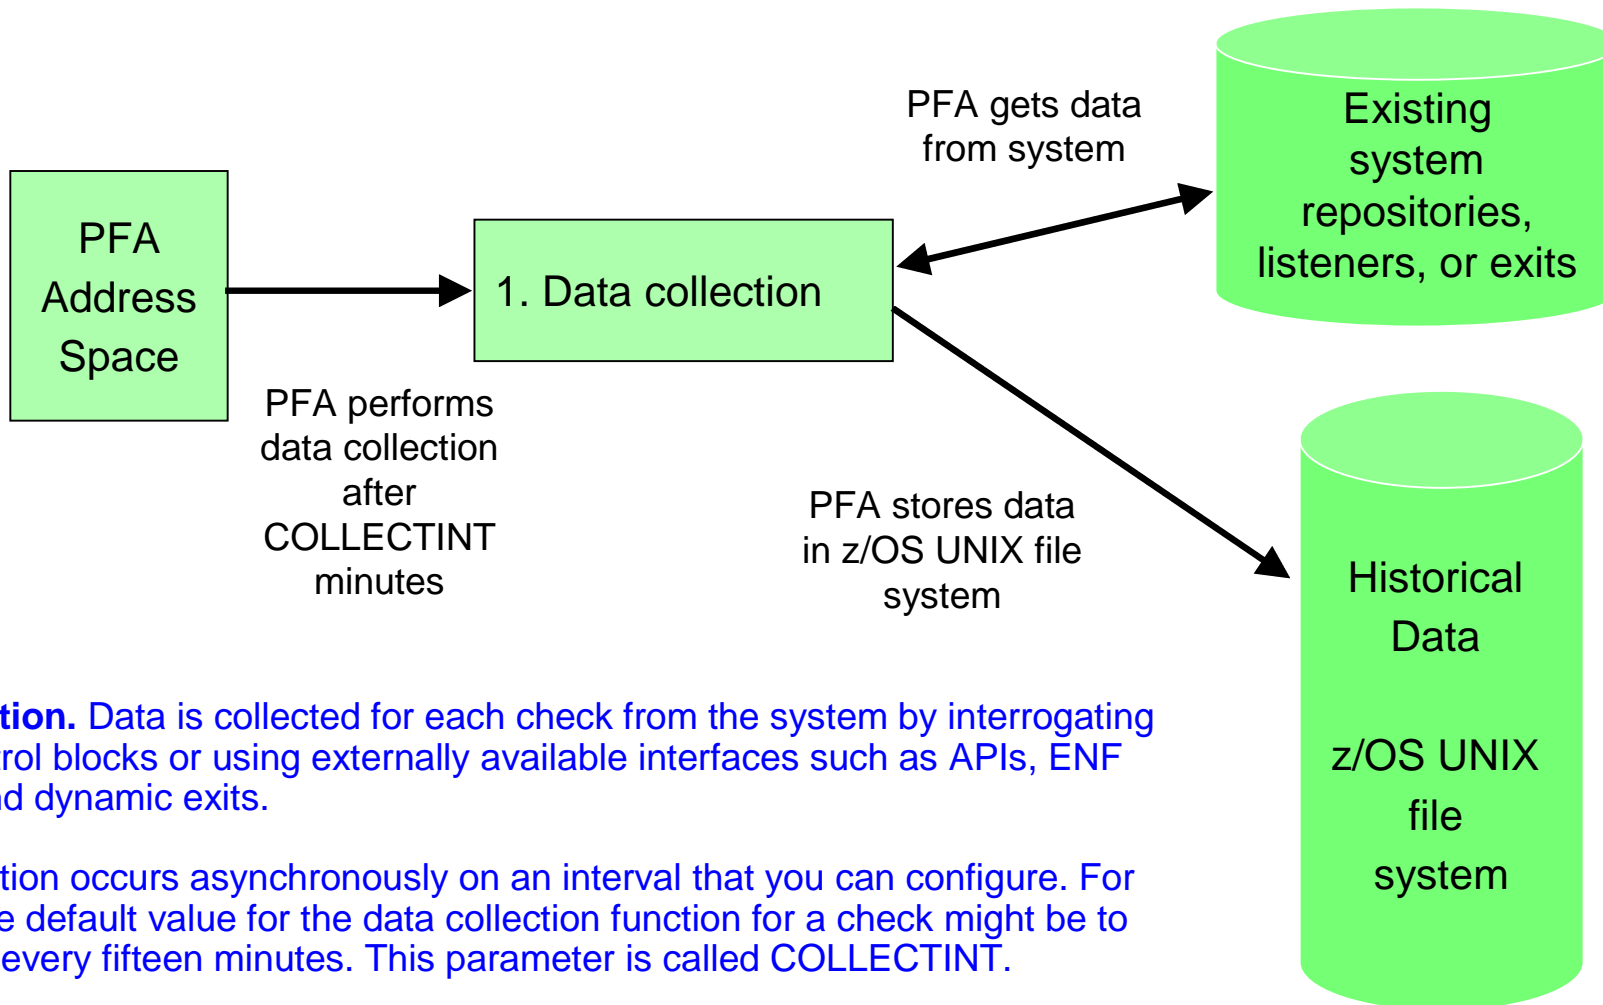
- PFA is built using the remote check feature of the IBM Health Checker for z/OS framework
- The IBM Health Checker for z/OS commands are available for all PFA checks
- The results of the PFA comparisons are available through IBM Health Checker for z/OS interfaces such as the health check support in SDSF
- The results of PFA's comparisons are sent back to IBM Health Checker for z/OS which writes the report. If an exception occurs, a WTO is issued by default
- If IBM Health Checker for z/OS is not active prior to PFA starting, PFA collects and models data, but waits for IBM Health Checker for z/OS to start before performing comparisons or issuing results because those functions are dependent upon IBM Health Checker for z/OS.



PFA processing



PFA data collection



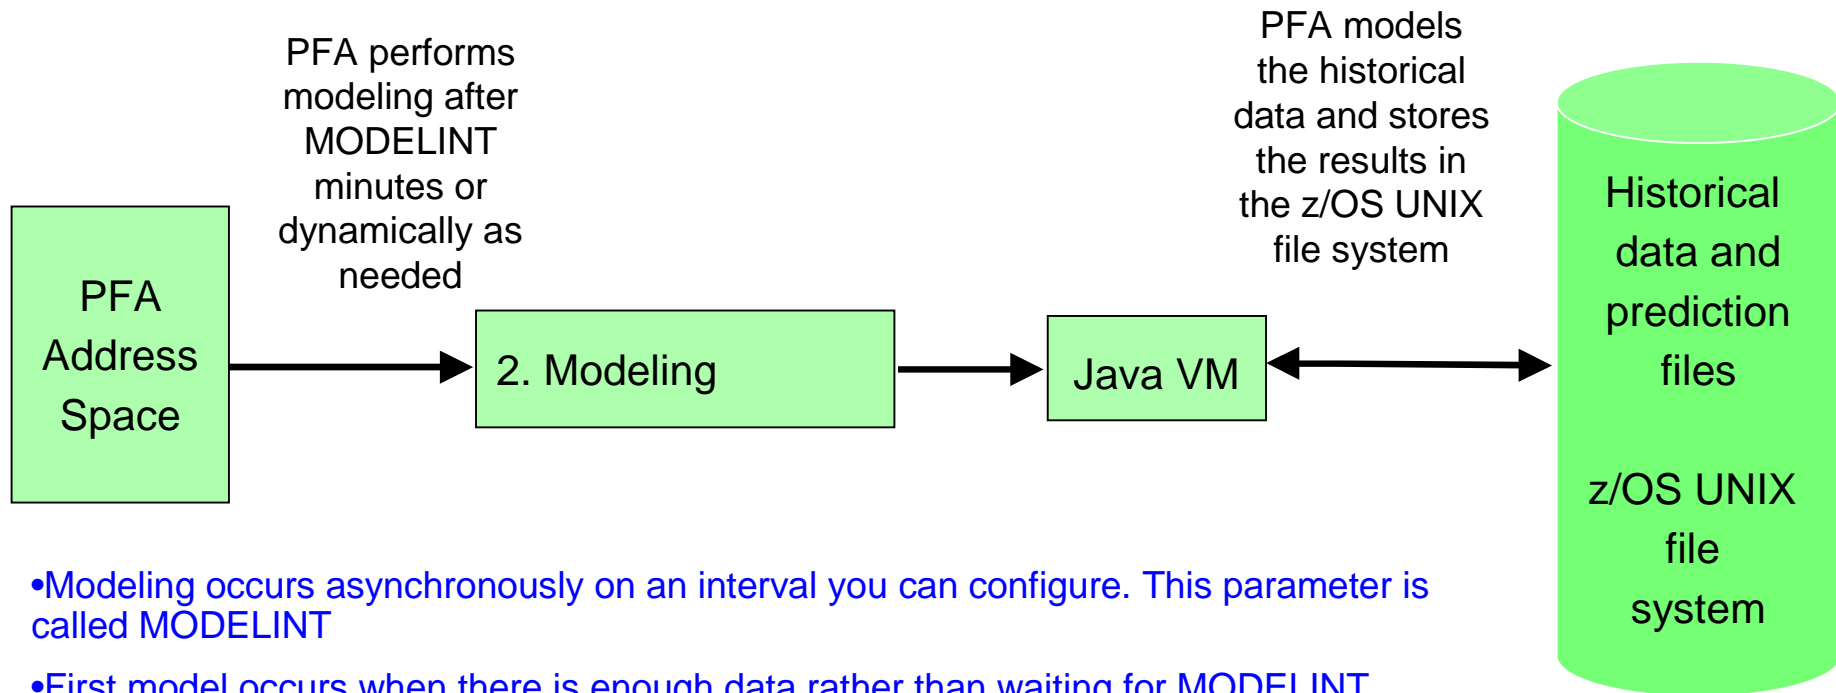
Data collection. Data is collected for each check from the system by interrogating system control blocks or using externally available interfaces such as APIs, ENF listeners, and dynamic exits.

- Data collection occurs asynchronously on an interval that you can configure. For example, the default value for the data collection function for a check might be to collect data every fifteen minutes. This parameter is called COLLECTINT.

- The data is stored in files in the z/OS UNIX file system and the files are self-maintained by PFA. Data that is deemed too old by PFA for use in making new predictions is deleted.



PFA Modeling



- Modeling occurs asynchronously on an interval you can configure. This parameter is called MODELINT
- First model occurs when there is enough data rather than waiting for MODELINT
- Some checks require 7-13 hours of data collection prior to a model being created
- Four successful collections required before modeling is attempted
- Models are then created based in MODELINT or if PFA determines current model is too old (trends changing)
- Results of modeling are stored in z/OS UNIX file systems – self-maintained by PFA






PFA comparison processing


IBM Health Checker for z/OS framework

Comparison initiated via INTERVAL parameter or by PFA after a collection

Predictive Failure Analysis



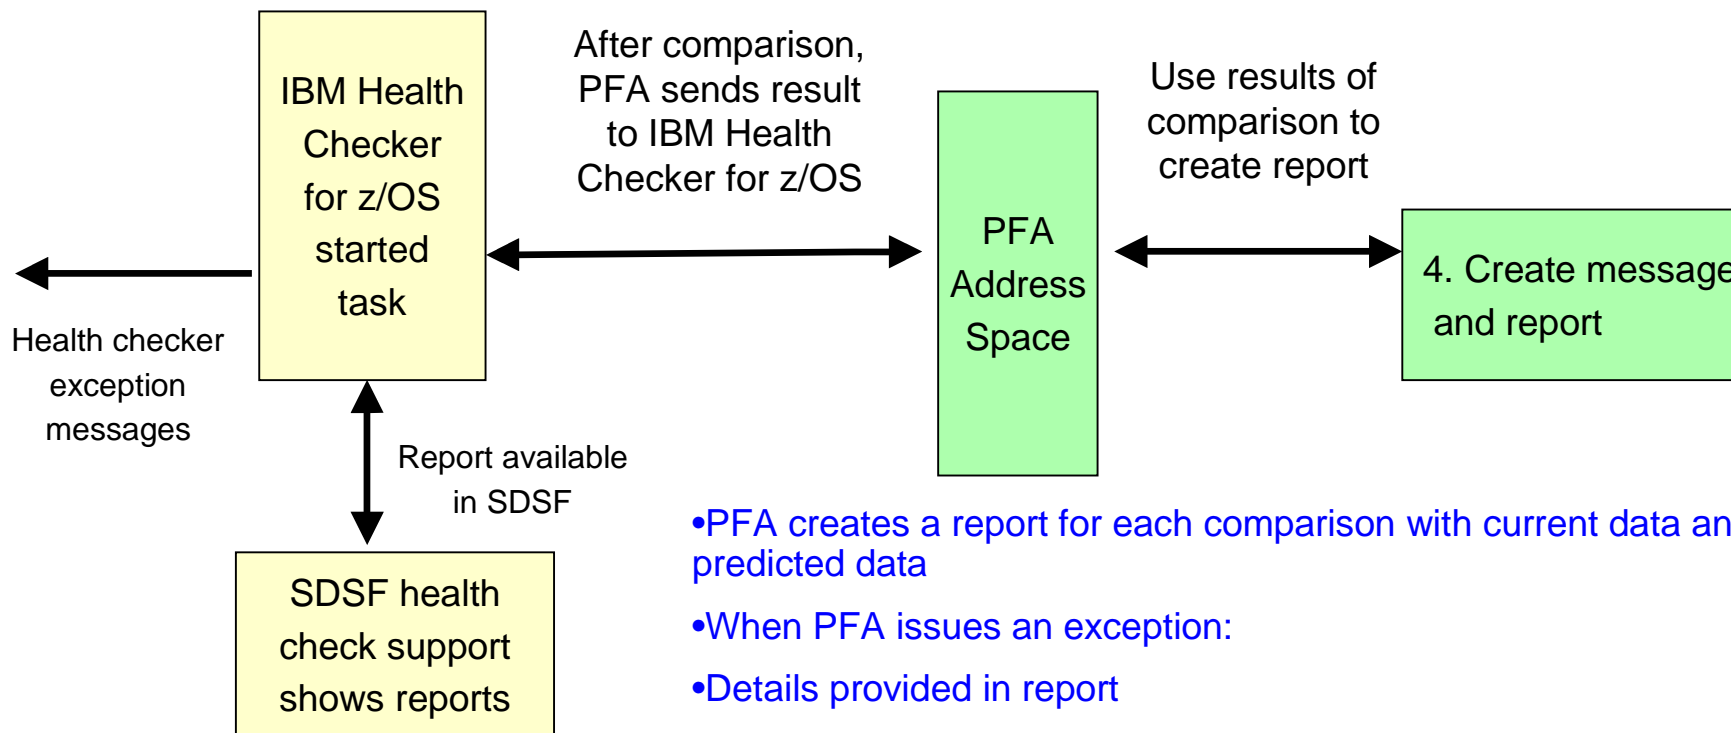
Use predictions and current usage values in statistical algorithms



- PFA performs comparisons needed to determine if an exception should be issued
- Applies mathematical algorithms and z/OS domain knowledge and user-defined parameters
- Determines if values are normal, abnormal or becoming abnormal
- If abnormal or becoming abnormal may cause a remodel to occur before next MODELINT
- Comparisons are initiated either by INTERVAL parameter of Health Check or by PFA



PFA reporting



- PFA creates a report for each comparison with current data and predicted data
- When PFA issues an exception:
- Details provided in report
- WTO issued – future WTOs for same exception suppressed until new data available
- Sub-directory created under PFA check’s UNIX filesystem directory containing diagnostic data for IBM service
- Modeling may occur more frequently until data stabilizes



PFA report format

Message Arrival Rate Prediction Report

```

Last successful model time      : 04/05/2012 07:08:01
Next model time                 : 04/05/2012 19:08:04
Model interval                  : 720
Last successful collection time : 04/05/2012 08:22:15
Next collection time            : 04/05/2012 08:37:16
Collection interval             : 15

```

Heading information: data on modeling and collection

```

Message arrival rate
  at last collection interval      :      83.52
Prediction based on 1 hour of data :      98.27
Prediction based on 24 hours of data:      85.98
Prediction based on 7 days of data :     100.22

```

System-level information: includes current values and predictions

Top persistent users:

Job Name	ASID	Message Arrival Rate	Predicted Message Arrival Rate		
			1 Hour	24 Hour	7 Day
TRACKED1	001D	58.00	23.88	22.82	15.82
TRACKED2	0028	11.00	0.34	11.11	12.11
TRACKED3	0029	11.00	12.43	2.36	8.36
...					

Address space Information: Details on tracked address space or address space causing the problem



How PFA Chooses Address Spaces to Track

- **Some metrics require data for the *entire system* to be tracked**
 - Exhaustion of common storage for entire system
 - LOGREC arrivals for entire system grouped by key

- **Some metrics call for tracking only *persistent address spaces***
 - Those that start within the first hour after IPL.
 - For example, tracks “track groups usage” by persistent address spaces to detect damaged address spaces

- **Some metrics are most accurate when using several categories**
 - *“Chatty” persistent* address spaces tracked individually
 - Start within the first hour after IPL and have the highest rates after a warm-up period
 - Data from first hour after IPL is ignored.
 - After an IPL or PFA restart, if all are running, same address spaces are tracked.
 - Duplicates with the same name are not tracked
 - Restarted address spaces that are tracked are still tracked after restart.
 - *Other persistent* address spaces as a group
 - *Non-persistent* address spaces as a group
 - *Total system rate* (“chatty” + other persistent + non-persistent)

PFA_COMMON_STORAGE_USAGE (z/OS 1.10 SPE)

- Predicts *exhaustion of common storage* by the z/OS image
- z/OS 1.10 and 1.11 -- Models **two** locations → CSA+SQA and ECSA+ESQA
- z/OS 1.12 -- Models **six** locations, handles expansion, and performance improved → CSA, SQA, ECSA, ESQA, CSA+SQA, and ECSA+ESQA
- *Not a monitor* of individual address spaces
- Does not detect
 - Fragmentation or really rapid growth
 - Usage exceeds a specific threshold (done by VSM_COMMON_STORAGE_USAGE)
 - An address space abnormally consuming common storage without impacting the z/OS image



Example Report: z/OS 1.12 Common Storage Usage Report

- Top predicted users

- ▶ Tries to pinpoint potential villains
- ▶ Those whose usage has *increased* the most in the last hour

- Other information

- ▶ Expansion information
- ▶ IBM Health Checker for z/OS message in its entirety

Common Storage Usage Prediction Report
(heading information intentionally omitted)

Storage Location	Current Usage in Kilobytes	Prediction in Kilobytes	Capacity When Predicted in Kilobytes	Percentage of Current to Capacity
*CSA	2796	3152	2956	95%
SQA	455	455	2460	18%
CSA+SQA	3251	3771	5116	64%
ECSA	114922	637703	512700	22%
ESQA	8414	9319	13184	64%
ECSA+ESQA	123336	646007	525884	23%

Storage requested from SQA expanded into CSA and is being included in CSA usage and predictions. Comparisons for SQA are not being performed.

Address spaces with the highest increased usage:

Job Name	Storage Location	Current Usage in Kilobytes	Predicted Usage in Kilobytes
JOB3	*CSA	1235	1523
JOB1	*CSA	752	935
JOB5	*CSA	354	420
JOB8	*CSA	152	267
JOB2	*CSA	75	80
JOB6	*CSA	66	78
JOB15	*CSA	53	55
JOB18	*CSA	42	63
JOB7	*CSA	36	35
JOB9	*CSA	31	34

* = Storage locations that caused the exception.



PFA_LOGREC_ARRIVAL_RATE (z/OS 1.10 SPE)

- Detects *a damaged system* by predicting and comparing LOGREC arrival rates in a collection interval
- Models expected number of LOGRECs in *time ranges by key*
 - Exceptions produced for any key grouping for any time range
- Not looking for individual LOGRECs, bursts of failures or patterns of failures
- Must be a software LOGREC with a usable SDWA

```
LOGREC Arrival Rate Prediction Report
(heading information intentionally omitted)
                                     Key 0      Key 1-7      Key 8-15
                                     _____  _____  _____
Arrivals in last
  collection interval:                1          0          2
Predicted rates based on...
  1 hour of data:                     1          0          1
  24 hours of data:                   0          0          1
  7 days of data:                     0          0          1
  30 days of data:                    0          0          1
Jobs having LOGREC arrivals in last collection interval:
Job Name      ASID      Arrivals
-----      -
LOGREC08      0029      2
LOGREC00      0027      1
```



z/OS 1.11 PFA Checks

- **Frames and slots usage check** → recently deleted with OA40065

- **Message arrival rate check** → detects *damaged address spaces or a damaged LPAR* by tracking WTO and WTORs normalized by CPU across time ranges
 - Counted prior to possible exclusion by Message Flooding Automation
 - Tracks the four categories: “chatty” persistent; non-chatty, persistent; non-persistent; and total system
 - Does not detect abnormal patterns or single critical messages
 - Performs comparisons after every collection rather than on an INTERVAL schedule in IBM Health Checker for z/OS
 - An appropriate report is printed for each type of exception (see slide 12)

PFA_SMF_ARRIVAL_RATE (z/OS 1.12)

- Detects a *damaged system* based on an SMF arrival rate (normalized by CPU) across time ranges that is too high
- Same four categories as the Message Arrival Rate check
- *Not looking* for abnormal SMF record arrival patterns or single SMF record arrivals
 - *If SMF is not running or stops,*
 - previously collected data is *automatically discarded* so that predictions aren't skewed.
 - *If you change the SMF configuration,*
 - delete the files in the PFA_SMF_ARRIVAL_RATE/data directory or your data will be skewed.
- Report very similar to Message Arrival Rate



PFA_ENQUEUE_REQUEST_RATE (z/OS 1.12)

- Detects a *damaged address space* or *damaged system* by comparing the number of enqueue requests per CPU millisecond (OA39924) to the rate expected.
- Two categories compared across three time ranges
 - “Chatty” persistent address spaces tracked individually and total system rate
 - 1 hour, 24 hour, and 7 day comparisons

```

                Enqueue Request Rate Prediction Report

(Heading information intentionally omitted.)

Enqueue request rate
  at last collection interval      :      83.52
Prediction based on 1 hour of data :      98.27
Prediction based on 24 hours of data:      85.98
Prediction based on 7 days of data  :     100.22

Top persistent users:

                Predicted Enqueue
                Request Rate

Job           Enqueue
Name          ASID      Request
-----
TRACKED1 001D      58.00      23.88      22.82      35.82
TRACKED2 0028      11.00      10.34      11.11      12.11
TRACKED3 0029      11.00      12.43      12.36      8.36
    
```



PFA_JES_SPOOL_USAGE (z/OS 1.13)

- Detects a *damaged address space or system* based on persistent jobs usage of the number of track groups
- Models 15 persistent jobs with the highest *increase* in their track group usage from one collection to the next
 - The number of actual track groups used is irrelevant due to the fact we are looking for a damaged address space or system *rather than exhaustion of track groups*.
 - Dynamic modeling occurs when “top jobs” change significantly to model new top jobs
 - JES2 only



PFA_JES_SPOOL_USAGE (z/OS 1.13)


- The exception is issued based on an *unexpected increase in the number of track groups used* from one collection to the next
 - Often confused – it is looking for a *damaged address space* based on an abnormal *increase* in track groups usage NOT exhaustion!
- The current number of track groups used is provided as additional information and is *totally irrelevant* to the comparisons.

JES Spool Usage Prediction Report

(Heading information intentionally omitted.)

Address spaces causing exception:

Job Name	ASID	Current Change in Number of Track Groups Used	Expected Change in Number of Track Groups Used	Current Number of Track Groups Used
JOB1	0019	252	10	892
JOB55	000E	129	3	400



PFA Integration with Runtime Diagnostics (z/OS 1.13)

- Detects a *damaged or hung address space or system* based on rates being too low
- When PFA detects an abnormally low condition, Runtime Diagnostics is executed
 - If the results of Runtime Diagnostics indicate a problem,
 - the PFA exception is issued
 - the PFA prediction report includes the Runtime Diagnostics output
- **Supported by three checks** → Message Arrival Rate, SMF Arrival Rate, and Enqueue Request Rate
- **Supported by three categories** (if supported by the check) → “Chatty” persistent jobs, other persistent jobs as a group, and total system
- **The Runtime Diagnostics address space (HZR) must be active**

Exception Report for PFA Integration with Runtime Diagnostics

- **“Too low” exception** message sent as WTO by default
- **Runtime Diagnostics output** included in PFA report
- Prediction report and result message **available in SDSF** (sdsf.ck)
- **PFA current rates and predictions** relevant to category causing exception

```

Message Arrival Rate Prediction Report
(Heading information intentionally omitted.)

Persistent address spaces with low rates:

Job Name      ASID      Message Arrival Rate      Predicted Message Arrival Rate
              ASID      Rate      1 Hour      24 Hour      7 Day
-----
JOBS4      001F      1.17      23.88      22.82      15.82
JOBS5      002D      2.01      8.34      11.11      12.11

Runtime Diagnostics Output:

Runtime Diagnostics detected a problem in job: JOBS4
EVENT 06: HIGH - HIGHCPU - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID CPU RATE: 96% ASID: 001F JOBNAME: JOBS4
STEPNAME: PFATEST PROCSTEP: PFATEST JOBID: STC00042 USERID:
+++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MAY BE LOOPING.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
-----

EVENT 07: HIGH - LOOP - SYSTEM: SY1 2009/06/12 - 13:28:46
ASID: 001F JOBNAME: JOBS4 TCB: 004E6850
STEPNAME: PFATEST PROCSTEP: PFATEST JOBID: STC00042 USERID:
+++++++
JOBSTART: 2009/06/12 - 13:28:35
Error:
ADDRESS SPACE APPEARS TO BE IN A LOOP.
Action:
USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.

(Additional output intentionally omitted.)
    
```



PFA Dependencies and Installation

- IBM Health Checker for z/OS – recommend to start at IPL
- z/OS UNIX file system – where we store the data
- Java (31-bit only) – used primarily during modeling
 - Java 5.0 or later for z/OS 1.11 and z/OS 1.12
 - [zAAP eligible](#) (recommended)
- z/OS 1.13 – The Runtime Diagnostics address space must be active for “too low” detection
- Follow the install steps in *z/OS Problem Management*

How to Get the Most Out of PFA

- Use a zAAP to offload PFA's Java Processing
- Start z/OS Resiliency functions at IPL
 - IBM Health Checker for z/OS
 - PFA
 - Runtime Diagnostics (z/OS 1.13)
- Automate the PFA IBM Health Checker for z/OS exceptions
 - Simplest: Add exception messages to existing message automation product
 - More complex: Use exception messages and other information to tailor alerts
 - See *z/OS Problem Management* for exceptions issued for each check
- Create a policy in an HZSPRMxx member for persistent changes
 - Not all check-specific parameters are required on an UPDATE of PFA checks!
 - UPDATE CHECK=(IBMPFA,PFA_COMMON_STORAGE_USAGE)
PARM('THRESHOLD(3)')

How to Get the Most Out of PFA (continued)

- **Get the latest PTFs!**
 - Configuration value default changes
 - Comparison algorithm tuning changes
 - Changes to design
 - Exclude interactive users from being persistent jobs for some checks
 - Skip comparisons for ESQA
 - zFS space reduction

- Help us to make PFA's results better for everyone!

How to Get the Most Out of PFA (continued)

- Use PFA's modify command to display parameters
 - Modify command in IBM Health Checker for z/OS does not display cumulative values
- Change the type of WTO of a check if default is not appropriate for your installation
 - Default is SEVERITY(MED) which issues an eventual action WTO
 - `f hzsproc,update,check(ibmpfa,pfa_j*),wtotype=info`
- Quiesce rather than delete PFA checks
 - Optional (only if you want to stop collections and modeling):
 - `f hzsproc,update,check(ibmpfa,pfa_j*,parm('collectinactive(0)')`
 - `f hzsproc,deactivate,check(ibmpfa,pfa_j*)`

How to Get the Most Out of PFA (continued)

- z/OS 1.12 -- Implementing supervised learning
 - Example: Exclude test programs that issue many LOGRECs and cause exceptions.
 - Example: Exclude address spaces that issue many WTOs, but are inconsistent or spiky in their behavior and cause message arrival rate exceptions. Supported by all checks except Common Storage Usage
 - Create EXCLUDED_JOBS file in the check's /config directory
 - Simple comma-separated value format
 - JobName,Systems,Date,Reason
 - Supports wildcards in both job name and system name
 - KKA*,*,04/05/2011,Exclude all KKA* jobs on all systems
 - Use `f pfa,update,check(check_name)` if PFA running
 - PFA creates an EXCLUDED_JOBS file for some checks during installation



How to Get the Most Out of PFA (continued)

- **Use check-specific tuning parameters** to adjust *sensitivity of comparisons* if needed
 - To minimize customer configuration
 - Default parameter values constructed from in-house and external data
 - Some defaults changed via PTFs using customers' data

Parameter	Description
STDDEV	<ul style="list-style-type: none"> > Increase value to decrease sensitivity. > Not available on the Common Storage Usage check.
EXCEPTIONMIN	<ul style="list-style-type: none"> > Increase value to decrease exceptions issued for relatively low rates. > Not available on the Common Storage Usage check
THRESHOLD	<ul style="list-style-type: none"> > Increase value to decrease sensitivity. > Common Storage Usage check only
STDDEVLOW	<ul style="list-style-type: none"> > Increase value to decrease sensitivity for “too low” checking. > Available on checks where “too low” checking is supported.
LIMITLOW	<ul style="list-style-type: none"> > Defines the maximum rate where “too low” checking is performed > Available on checks where “too low” checking is supported.



How to Get the Most Out of PFA (continued)

- Use PFA check-specific parameters to *affect other behavior*

Parameter	Description
COLLECTINT	Number of minutes between collections
MODELINT	Number of minutes between models <ul style="list-style-type: none"> PFA automatically and dynamically models more frequently when needed z/OS 1.12 default updated to 720 minutes. First model will occur within 6 hours (or 6 hours after warm-up)
COLLECTINACTIVE	Defines whether PFA should collect and model if check not active/enabled in IBM Health Checker for z/OS
DEBUG	Use only if IBM service requests it
CHECKLOW	z/OS 1.13 – Turns on/off “too low” checking with RTD for checks that support it
TRACKEDMIN	Requires a persistent job to have this minimum rate at the end of the warm-up in order to be tracked (where supported)
Health Checker parameters	For example, SEVERITY -- All PFA checks default = SEVERITY(MED): Eventual action WTO



Summary

- PFA uses *historical data* and *machine learning algorithms* to detect and report soft failures *before they can impact your business*
 - PFA is focused on *damaged systems* and *resource exhaustion*
- Tips
 - Use the PFA reports to help diagnose problems
 - Use a zAAP for PFA's java processing
 - Tune the PFA checks using the configuration parameters and the EXCLUDED_JOBS list if necessary.
 - Stay current on PTFs
 - Start IBM Health Checker for z/OS, Runtime Diagnostics, and PFA at IPL
 - Automate exception messages
 - *z/OS Problem Management* is the main source for documentation
- Use other products to do deep investigation of system or address space problems.

Additional Resources

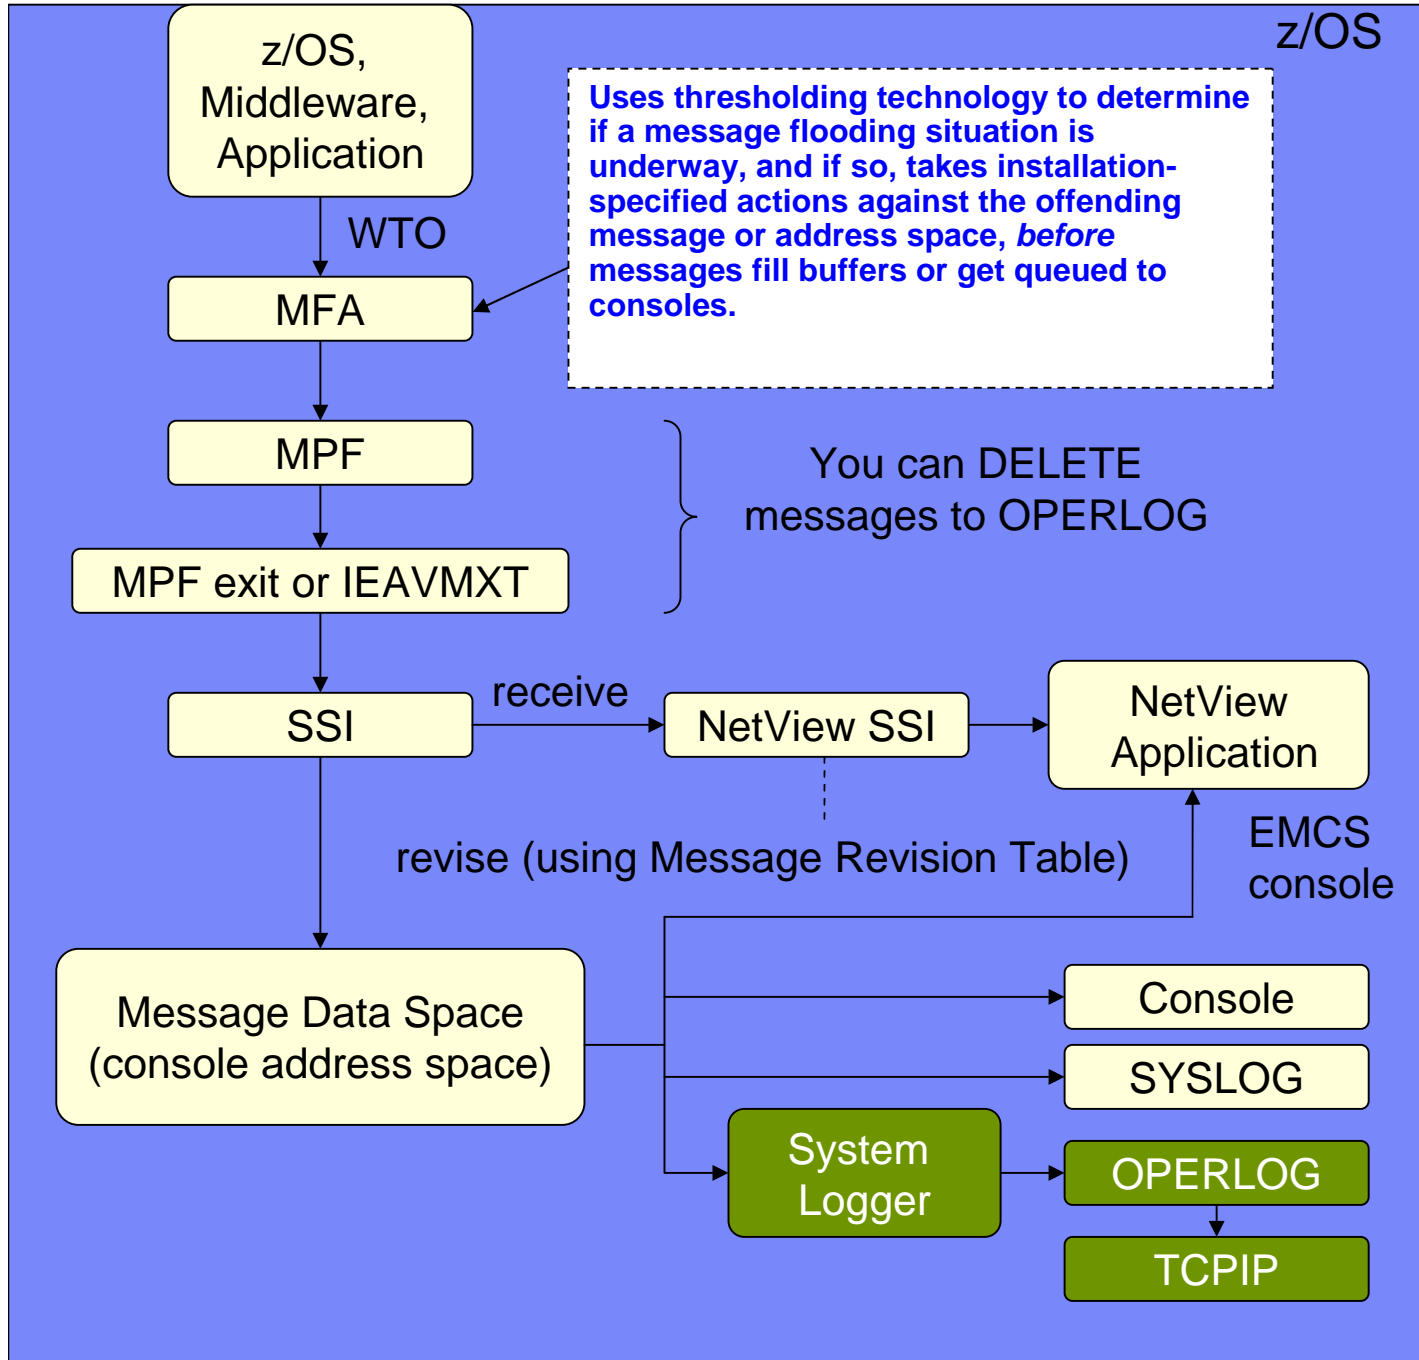
- One main source of information: *z/OS Problem Management G325-2564-XX*
- PFA IEA presentations
 - ▶ http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp?topic=/com.ibm.iea.zos/zos/1.11/Availability/V1R11_PFA/player.html
 - ▶ http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp?topic=/com.ibm.iea.zos/zos/1.12/Availability/V1R12_Availability_PFA_Enhancements/player.html
- *z/OS Hot Topics Newsletters*: http://www.ibm.com/systems/z/os/zos/bkserv/hot_topics.html
 - ▶ #20 (GA22-7501-16) -- *Fix the Future with Predictive Failure Analysis* by Jim Caffrey, Karla Arndt, and Aspen Payton
 - ▶ #23 (GA22-7501-19) – *Predict to prevent: Let PFA change your destiny* by Jim Caffrey, Karla Arndt, and Aspen Payton
 - ▶ #23 (GA22-7501-19) – *Runtime to the Rescue! Using Runtime Diagnostics to find out your problems fast* by Bob Abrams, Don Durand, and Dave Zingaretti
- *IBM Systems Magazine - Mainframe Edition*
 - ▶ PFA *A Soft Touch* by Karla Arndt, Jim Caffrey, and Aspen Payton
 - ▶ http://www.ibmssystemsmagmainframedigital.com/nxtbooks/ibmsystemsmag/mainframe_20101112/index.php#/48

Message Flood Automation



z/OS Message Flood Automation

- **Not New**
- **Delivered as an SPE (OA17514) for z/OS R6, R7 and R8.**
- **Delivered in the base of z/OS R9.**
- **Part of z/OS Consoles component.**
- **Shipped by GDPS since 1Q2003 – installed at some of their largest customers. Has successfully dealt with a number of message flooding situations.**
- **Another tool available to reduce outages**



Message Flood Automation goals:

Identify runaway conditions that can cause disruptions:

- Large number of messages that obscure important messages or delay automation processing
- Large number of messages that lead to excessive CPU and storage utilization.
- Exhaustion of console buffers
- Large console queues that require lengthy manual intervention



Message classes

Messages are divided into three classes, with independent policy and independent processing for each class:

- SPECIFIC messages – messages by message ID
- ACTION messages – messages with action descriptor codes
 - 1, 2, 3, 11
- REGULAR messages – all other messages
 - For each class, the class policy can be overridden for specific message IDs (SPECIFIC messages) or JOBS (ACTION and REGULAR messages)

Actions

Actions that can be taken:

- LOG | NONLOG – writing of the message to SYSLOG/OPERLOG
- DISPLAY | NODISPLAY – queuing of the message to consoles
- AUTO | NOAUTO – queuing of the message for automation
- RETAIN | NORETAIN – (action messages only) retention in AMRF
- CMD | NOCMD – issuing a command
- IGNORE | NOIGNORE – whether to take action or not

Specifying NOLOG, NODISPLAY and NOAUTO together will delete the message. Deleted messages are not sent to the console address space or to other systems in the sysplex.

Commands are issued from the ADDRESS SPACE causing the flood, so must have sufficient authority to issue the specified command.

MFA modes

Normal Mode (low processing overhead)

If # of messages counted \geq (MSGTHRESH) and time to count $<$ INTVLTIME then go to intensive mode.

Intensive Mode – extra processing

If any address space issues JOBTHRESH/MSGLIMIT number of messages within INTVLTIME then this address space or message is subject to specified action (act upon mode).

At end of each MSGTHRESH if time since last message is $>$ JOBIMTIME or MSGIMTIME then end act upon mode

At end of each MSGTHRESH if time since last message is $>$ SYSIMTIME then end intensive mode

```

/* Sample MSGFLDxx parmlib member */

REGULAR MSGTHRESH=50 ,JOBTHRESH=20 ,INTVLTIME=1
REGULAR SYSIMTIME=2 ,JOBIMTIME=2
DEFAULT LOG,NOAUTO,NODISPLAY,NOCMD
DEFAULTCMD '&,CANCEL & -- cancelled by MFA'
JOB AOC%NV* AUTO
JOB LLA* AUTO
JOB ZAP1 CMD

ACTION MSGTHRESH=50 ,JOBTHRESH=20 ,INTVLTIME=1
ACTION SYSIMTIME=2 ,JOBIMTIME=2
DEFAULT LOG,NOAUTO,NODISPLAY,NOCMD,NORETAIN
DEFAULTCMD '&,CANCEL & -- cancelled by MFA'
JOB AOC%NV* AUTO,RETAIN
JOB LLA* AUTO
JOB ZAP2 CMD

SPECIFIC MSGTHRESH=50 ,INTVLTIME=1
SPECIFIC SYSIMTIME=2
SPECIFIC MSGIMTIME=2
SPECIFIC MSGLIMIT=20
DEFAULT LOG,NOAUTO,NODISPLAY,NORETAIN
MSG IOS001E
MSG IOS003A
MSG IOS050I
MSG IOS051I
    
```



Normal Mode thresholds

- **MSGTHRESH**
 - A positive, non-zero integer in the range of 1 to 999,999,999
 - Represents a number of messages (counts only 1st line of multiline WTO)
- **INTVLTIME**
 - A positive, non-zero integer time in seconds in the range 1 to 999,999,999
 - Represents an 'interval' of time

- Messages are counted and stored in a counter.
- Once counter is equal to MSGTHRESH, the time taken to reach MSGTHRESH is determined.
- If the time taken to reach MSGTHRESH < INTVLTIME then MFA is put into INTENSIVE mode
- Since a flood may begin at any point during the count, it is possible that one or more MSGTHRESH intervals may be needed before MFA detects the flood.
- The bottom line is that the triggering of intensive mode may not occur precisely after MSGTHRESH flood messages have occurred.

Intensive Mode thresholds

- JOBTHRESH (regular / action messages) | MSGLIMIT (specific messages)
 - A positive, non-zero integer in the range of 1 to 999,999,999 - represents a number of messages
 - MUST be less than MSGTHRESH
- If any one address space issues JOBTHRESH messages within INTVLTIME, it is subject to defined action from then on (act upon).
 - 30-40% of MSGTHRESH is a good starting point
- For specific messages algorithm is similar but applies to individual messages and not jobs/address spaces. Threshold used is MSGLIMIT.
 - 15-20% of MSGTHRESH is a good starting point

Detecting end of message flood

- Inter-message thresholds
 - a positive, non-zero floating point time in seconds in the range 0.000001 to 16777215.0
 - Time between messages
 - Inverse of message rate: for example, a message rate of 2.0 messages/second means that messages arrive on average every .5 seconds (so the inter-message time is .5 seconds)

- SYSIMTIME: system inter-message time
- JOBIMTIME: job inter-message time
- MSGIMTIME: message inter-message time

- In intensive mode if time since last message is $>$ SYSIMTIME then intensive mode is discontinued.
- If an address space is in act-upon mode and time since last message is $>$ JOBIMTIME then the address space is removed from act-upon mode
- If a message is in act-upon mode and time since last message is $>$ MSGIMTIME then the message is removed from act upon mode

Determining thresholds

The key to setting up Message Flood Automation is determining the REGULAR MSGTHRESH value.

- Should be set high enough that Message Flood Automation is not constantly “tripping” in and out –
- *But* low enough to catch real message floods in their infancy

Message Rate Monitoring function provided to measure the message rate as Message Flood Automation sees it

- Message Rate Monitoring provides suggested values for REGULAR MSGTHRESH
- Message Rate Monitoring should be run periodically to readjust REGULAR MSGTHRESH value

Message Rate Monitoring

Turn on/off via operator command

- SETMF MONITORON
- SETMF MONITOROFF

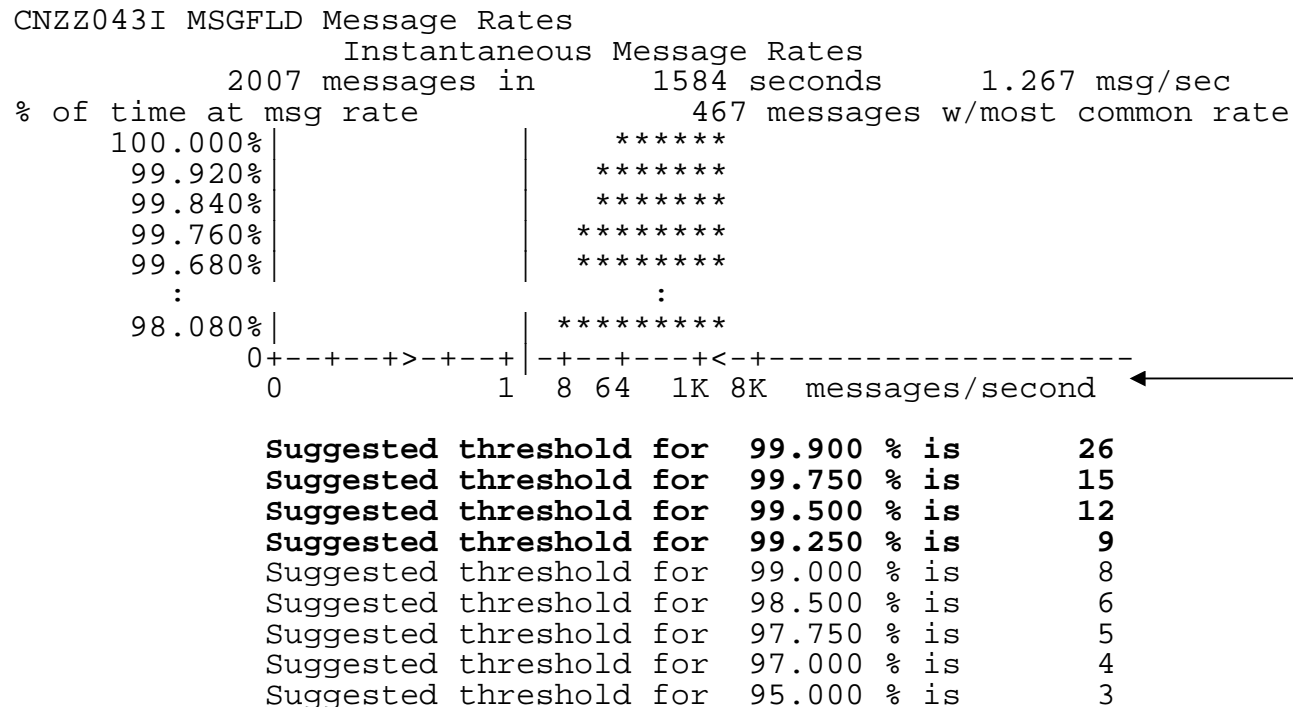
- It uses the same filtering as the rest of Message Flood Automation and therefore “sees” messages at the same rate that Message Flood Automation “sees” them.

- The message rate information captured by the Message Rate Monitoring function can be displayed in the form of a graph that is produced in response to a DISPLAY MSGFLD,MSGRATE command.

- Take a 24-hour sample every few months during a busy period and adjust your policy if needed.

Monitor Report to derive MSGTHRESH

```
D MF,MSGRATE,25,2      -- display top 2% of graph in 25 lines
```



The Y-axis is the percent of time at a particular message rate (or less).

The X-axis scale is logarithmic – each position is power-of-2 greater or lesser than its neighbor

Result of interpolation, rounded to the nearest integer value

- The suggested thresholds are interpreted as follows: “99.5% of the time, the instantaneous message rate will not exceed 12 messages per second”.
- Use to set the MSGTHRESH value in PARMLIB MSGFLDxx (recommended 99% threshold value)
- For general recommendations see “Setting thresholds based on message rates” in z/OS MVS Planning: Operations





MSGFLDxx statements

STATEMENT	REGULAR	ACTION	SPECIFIC
msgtype	REGULAR	ACTION	SPECIFIC
DEFAULT	LOG NOLOG AUTO NOAUTO DISPLAY NODISPLAY CMD NOCMD	LOG NOLOG AUTO NOAUTO DISPLAY NODISPLAY CMD NOCMD RETAIN NORETAIN	LOG NOLOG AUTO NOAUTO DISPLAY NODISPLAY RETAIN NORETAIN IGNORE NOIGNORE
DEFAULTCMD	'cmdchr[ASIDchar],cmdtext'	'cmdchr[ASIDchar],cmdtext'	N/A
JOB 0-64 possible	jobname [action][,action]	jobname [action][,action]	N/A
MSG 0-1024 possible	N/A	N/A	msgid [action][,action]



MSGFLDxx REGULAR and ACTION example

```
REGULAR MSGTHRESH=50,JOBTHRESH=20,INTVLTIME=1  
REGULAR SYSIMTIME=2,JOBIMTIME=2  
DEFAULT LOG,NOAUTO,NODISPLAY,NOCMD  
DEFAULTCMD '&,CANCEL & -- cancelled by Message Flood Automation'  
JOB AOC%NV* AUTO
```

```
ACTION MSGTHRESH=50,JOBTHRESH=20,INTVLTIME=1  
ACTION SYSIMTIME=2,JOBIMTIME=2  
DEFAULT LOG,NOAUTO,NODISPLAY,NOCMD,NORETAIN  
DEFAULTCMD '&,CANCEL & -- cancelled by Message Flood Automation'  
JOB ZAP2 CMD
```

MSGFLDxx SPECIFIC example

```
SPECIFIC MSGTHRESH=50,INTVLTIME=1  
SPECIFIC SYSIMTIME=2  
SPECIFIC MSGIMTIME=2  
SPECIFIC MSGLIMIT=20  
DEFAULT LOG,NOAUTO,NODISPLAY,NORETAIN  
MSG IOS001E  
MSG IOS003A  
MSG IOS050I  
MSG IOS051I  
MSG IOS071I  
MSG IOS251I  
MSG IOS444I  
MSG IOS450E
```

MSGFLDxx defaults

For REGULAR, ACTION and SPECIFIC messages:

- MSGTHRESH=50
- JOBTHRESH=20
- MSGLIMIT=20 for specific
- MSGIMTIME=2 for specific
- INTVLTIME=1
- SYSIMTIME=2
- JOBIMTIME=2

REGULAR actions

- LOG, AUTO, NODISPLAY, NOCMD

ACTION actions

- LOG, AUTO, NODISPLAY, NOCMD, NORETAIN

SPECIFIC actions

- LOG, AUTO, NODISPLAY, NORETAIN, NOIGNORE

Defaults activated if:

MSGFLD=(NONE,(ON)) on a
CONSOLxx INIT statement

SETMF ON issued before a
SET MSGFLD=xx

Activate MFA with defaults

SETMF ON prior to issuing a SET MSGFLD=xx starts MFA with defaults

```
SETMF ON  
CNZZ041I Message Flood Automation ENABLED.  PARMLIB member:internal
```

```
D MF,STATUS  
CNZZ042I MSGFLD Status: ENABLED. 742  
Policy INITIALIZED.      Using PARMLIB member: internal  
Message rate monitoring DISABLED.          0 msgs          0 secs  
No message flood is underway.
```

Default thresholds

```

D MF,PARAMETERS
CNZZ901I MSGFLD Parameters 744
Message type    REGULAR      ACTION      SPECIFIC
INTVLTIME =          1          1          1
JOBIMTIME =    2.000000    2.000000
JOBTHRESH =          20          20
MSGCOUNT =          0          0          0
MSGIMTIME =                    2.000000
MSGLIMIT =                    20
MSGTHRESH =          50          50          50
NUMJOBS =          128          128
SYSIMTIME =    2.000000    2.000000    2.000000
    
```



Action defaults

```

D MF,DEFAULTS
CNZZ904I MSGFLD Defaults 746
Message type REGULAR ACTION SPECIFIC
LOG = Y Y Y
AUTO = Y Y Y
DISPLAY = N N N
CMD = N N
RETAIN = N
IGNORE = N
REGULAR CMD action command text
'&%,CANCEL &,A=% Default cmd action from Message Flood Automation '
ACTION CMD action command text
'&%,CANCEL &,A=% Default cmd action from Message Flood Automation '
    
```



Example 1 – WTO flood from JOB

```

11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000018 CNZZ007E MESSAGE RATE EXCEEDED 50 MESSAGES
IN <1 Seconds
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000010 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00000010 +MSG11 EST
11:08:06.61 JOB31874 00000018 CNZZ002E MESSAGE THRESHOLD REACHED FOR JOB
KARANMF
11:08:06.61 JOB31874 00080619 IEC130I SYSPRINT DD STATEMENT MISSING
11:08:06.61 JOB31874 00080619 +MSG11 EST
    
```

JOB loops writing WTO, flooding console

MFA detects and goes into intensive mode (MSGTHRESH reached)

JOBTHRESH reached and address space is in act upon mode

Messages suppressed from console



Example 1

Status display will show is a flood is underway

```

D MF,STATUS
CNZZ042I MSGFLD Status: ENABLED. 636
Policy INITIALIZED.      Using PARMLIB member: internal
Message rate monitoring DISABLED.          0 msgs          0 secs
A message flood is underway.
JOBNAME      ASID T MSGS-ACTED-ON  --DURATION  -----STARTED-----
KARANMF      0034 R           24336      0:02.71    2012277 11:12:25.01
    
```

D MF,MODE will tell you whether intensive mode is active

```

D MF,MODE
CNZZ040I Intensive modes: REGULAR-ON  ACTION-OFF  SPECIFIC-OFF
    
```

CNZZ008E message issued when MFA leaves intensive mode

```

CNZZ008E REGULAR MESSAGE RATE ACCEPTABLE. 24336 MESSAGES ACTED UPON.
    
```



Example 2 – set a policy

```
/*-----*/
/* Sample MSGFLDxx PARMLIB member REGULAR MESSAGES ONLY */
/*-----*/
REGULAR MSGTHRESH=30 ,JOBTHRESH=20 ,INTVLTIME=1
REGULAR SYSIMTIME=.33 ,JOBIMTIME=2
DEFAULT LOG ,AUTO ,NODISPLAY ,NOCMD
DEFAULTCMD '&%,CANCEL &,A=% cmd action from Message Flood Automation'
JOB KARANMF CMD
```

```
SET MSGFLD=01
CNZZ016I Message Flood Automation policy initialized.
CNZZ401I Message Flood Automation loading: MSGFLD01
IEE252I MEMBER MSGFLD01 FOUND IN SYS1.PARMLIB
CNZZ410I Message Flood Automation loading of MSGFLD01 complete.
```

```
D MF,JOBS
CNZZ905I MSGFLD JOB Actions 907
REGULAR messages LOG AUTO DISPLAY CMD
      JOB KARANMF  Y  Y  N  Y
ACTION messages LOG AUTO DISPLAY CMD RETAIN
No entries
```

Example 2

```
CNZZ007E MESSAGE RATE EXCEEDED      30 MESSAGES IN <1 SECONDS.
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
IEC130I SYSPRINT DD STATEMENT MISSING
+MSG11 EST
CANCEL KARANMF,A=001D CMD ACTION FROM MESSAGE FLOOD AUTOMATION
IEE301I KARANMF                CANCEL COMMAND ACCEPTED
```

Additional Resources

- **z/OS MVS Planning: Operations SA22-7601**
 - Details on MFA and setting thresholds

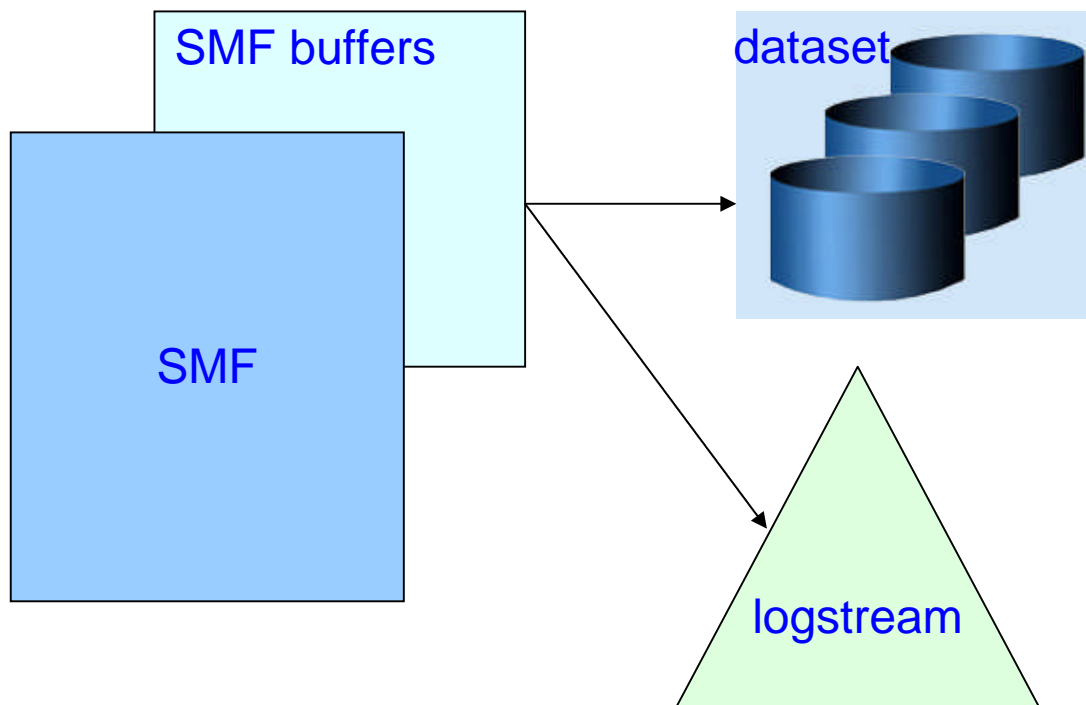
- **z/OS MVS Initialization and Tuning Reference SA22-7592**
 - MSGFLDxx parmlib member specification
 - Starting MFA from CONSOLxx member

- **z/OS MVS System Commands SA22-7627**
 - Display MF options
 - SET MF command option
 - SETMF command options

SMF record flood support



SMF record flood support



Regardless of the recording medium, either DASD or LOGSTREAM, space is finite and there is always the possibility of SMF buffers filling up, resulting in a system WAIT state.

For specified SMF records an SMF flood policy can be set up to warn about possible flood conditions and to have the SMF record suppressed.

SMF record flood support is enabled through the FLOOD statement in SMFPRMxx.

SMF record flood support policies are defined through the FLOODPOL statement in SMFPRMxx



Usage and Invocation

- Turning on and off the SMF Flood support facility
 - A new option is available for use in the SMFPRMxx member and SETSMF command which will turn the facility on and off.
 - FLOOD(ON|OFF)
 - When the facility is turned off all active counts are reset and any active flood actions are ended
 - By default the facility is turned off
 - The facility applies to both SMF data set record and SMF logstream recording

Usage and Invocation

- Setting up flood policies using the FLOODPOL keyword

```
FLOODPOL (
  {TYPE( {aa,bb} ) {aa,bb:zz} {aa,bb:zz,...} ) } ,
  RECTHRESH( xxxx ) ,
  INTVLTIME( xxxx ) ,
  MAXHIGHINTS( xxxx ) ,
  ENDINTVL( xxxx ) ,
  ACTION( {MSG|DROP} ) )
```

- Only available via SMFPRMxx, policy can not be updated using the SET SMF command.
- Policy action is either DROP or MSG.
- A single record type can have one MSG policy and one DROP policy.
- When a MSG and DROP Policy are specified for a single record type the MSG policy must be satisfied before processing of the DROP policy will being.
- Each SMF record type specified in single a FLOODPOL statement is treated separately. Doing this is equivalent to have multiple FLOODPOL statements each for a single record.

Usage and Invocation

- Several sub options exist for the FLOODPOL keyword
 - TYPE
 - The SMF record types this policy applies to.
 - RECTHRESH
 - The number of records that comprise an interval.
 - INTVLTIME
 - The minimum amount of time it can take to match RECTHRESH number of records to consider this interval a flooding interval. Specified in 10's of seconds.
 - MAXHIGHINTS
 - The number of intervals that must occur at or above the flooding rate before an action is taken.
 - ENDINTVL
 - The amount of time, in 10th of seconds, that must pass before RECTHRESH number of records are generated to assume the flood has ended
 - ACTION
 - The action this policy should take, either DROP or MSG.

Usage and Invocation

How the FLOODPOL processing works

- Once an interval (*RECTHRESH* number of records) are written within *INTVLTIME*, SMF will move the policy into an intensive monitoring mode.
- In intensive monitoring mode when an interval occurs within *INTVLTIME* SMF keeps a counter. If the counter reaches *MAXHIGHINTS* then the specified *ACTION* is taken. If an interval occurs slower than *INTVLTIME* SMF leaves intensive mode and returns to normal processing.
- When the *ACTION* is taken SMF waits for the rate to return to *RECTHRESH* records occurring in at least *ENDINTVL* amount before stopping the *ACTION*.
- If an *ACTION(MSG)* and *ACTION(DROP)* policy are active for a record type once the MSG policy becomes active the DROP policy will then be in effect. The DROP policy will then start running in normal mode and move through the same progression.

Example

```
FLOODPOL ( TYPE ( 4 , 5 ) , RECTHRESH ( 1000 ) , INTVLTIME ( 50 ) , MAXHIGHINTS ( 15 )  
          , ENDINTVL ( 120 ) , ACTION ( MSG ) )
```

For SMF record types 4 and 5:

If 1000 records are generated within 5 seconds ($50 * .10 = 5$ seconds = INTVLTIME)

Then enter intensive monitoring:

If 1000 records are generated within 5 seconds for 15 intervals (MAXHIGHINTS) then take specified action, in this case issue message (IFA780A)

When it takes more than 12 seconds ($120 * .10 = 12$ seconds = ENDINTVL) to generate 1000 records then end flood state and generate message (IFA781I)

Messages

- IFA780A SMF Record Flood MSG filter for type XX exceeded at time=hh.mm.ss
Issued when a flood is detected for an ACTION(MSG) type policy.
- IFA781I SMF Record Flood MSG filter for type XX returned to normal at time=hh.mm.ss
Issued when a flood subsides for an ACTION(MSG) type policy.
- IFA782A SMF Record Flood MSG filter for type XX returned to normal at time=hh.mm.ss
Issued when a flood subsides for an ACTION(DROP) type policy.
- IFA783I SMF Record Flood DROP filter for type XX returned to normal at time=hh.mm.ss, RECORDS DROPPED=yyyyy
Issued when a flood subsides for an ACTION(DROP) type policy.
- IFA784I SMF Record Flood Automation disabled due to errors
Issued if repeated ABENDs occur during flood processing.

Auto-Reply to WTOR



Overview

- Customers no longer have operators closely monitoring the system waiting to immediately reply to a WTOR.
- Operators typically do not have authority, experience, or system understanding to make their own decision on what to reply for uncommon WTORs.
- Reasonable to assume a WTOR may take at least 30-45 minutes to be answered.
 - No longer feasible to expect fast, accurate, knowledgeable answers from system operators.
 - Some WTORs are so infrequent that they are not automated and operators may never have seen them before.
- Reply delays can affect all systems in sysplex
- **Auto-reply to WTORs** provides a policy based means to automate a response to a WTOR when:
 - There is no automations
 - Operator is unaware of the outstanding request
 - Operator spends a long time determining what response should be given
 - **Auto-Reply to WTORs** provides a policy definition (via a parmlib member) which specifies a reply value and a time delay for a WTOR.
 - The system will issue the reply if the WTOR has been outstanding longer than the time delay.

Auto-reply process



If WTOR is in policy, then monitor.
Policy is specified in AUTORxx
member of PARMLIB.

Wait for specified delay for operator
to take action. Operator can stop
auto-reply from monitoring a specific
WTOR.

If delay time exceeded then respond
to WTOR using the response
specified in policy.

Overview

■ Benefit:

- While traditional automation products already provide this support, Auto-Reply can be used for most **WTORs issued during NIP** when inboard automation is not available.
- Simple replies easily handled without involving complex automation
- Default policy is pro-active in that it contains 78 WTOR replies
 - Replies & delays provided by component experts
- If policy exists, Auto-Reply activated automatically during IPL.
- Delay value allows operators/automation to reply first
- Automating replies:
 - Takes burden off operators
 - Allows systems to continue processing sooner

AUTORxx syntax

```
[NOTIFYMSG(S {HC|CONSOLE})]  
[MSGID([']msgid[']) {NOAUTORREPLY}  
{DELAY(nnn{M|S}) REPLY([']replytext['] [, [']replytext[']...)]}]
```

- The following parameters can be specified with the AUTORxx parmlib member:
- NOTIFYMSG (HC)
 - Indicates that the auto-reply notification messages (CNZ2605I, CNZ2606I and CNZ2608I) only appear in the hardcopy log - If NOTIFYMSGS is not specified, HC is the default value
- NOTIFYMSG(CONSOLE)
 - Indicates that the auto-reply notification messages (CNZ2605I, CNZ2606I and CNZ2608I) are displayed on consoles receiving routing codes 2 (operator information) or 10 (system programmer information) and also appear in the hardcopy log
- MSGID()
 - The first keyword of the message definition. This keyword must be in the range of 1 to 10 characters. The message ID must be enclosed in quotes if it contains non-alphanumeric characters, like equal sign or parenthesis

AUTORxx syntax

```
[NOTIFYMSG(S ({HC | CONSOLE}))]  
[MSGID([' ]msgid[' ]){NOAUTOREPLY}  
{DELAY(nnn{M|S}) REPLY([' ]replytext[' ] [, [' ]replytext[' ]...)]}]
```

■ MSGID()

- The first keyword of the message definition. This keyword must be in the range of 1 to 10 characters. The message ID must be enclosed in quotes if it contains non-alphanumeric characters, like equal sign or parenthesis (not converted to upper case).
- Partial support for wildcards

■ NOAUTOREPLY

- Can be specified to cause subsequent MSGID() specifications of this message ID (in this member or subsequent members) to be ignored. With the NOAUTOREPLY option, you can have your own AUTORxx member and use it to remove messages specified in the AUTOR00 member without actually removing the AUTOR00 statements. No auto-reply processing is performed for this message ID.

■ DELAY(nnn{M | S})

- Can appear in any order and on different lines from REPLY() and MSGID()
- nnn is the minimum amount of time, in minutes (M) or seconds (S) to delay after the WTOR is issued but before the system issues a reply. Only three digits are supported, which gives a limit of 999 minutes (16.65 hours) or 999 seconds (16.65 minutes). The delay value of 0 is supported, but specifying the value of 0 can prevent automation or an operator from providing a reply

AUTORxx syntax

```
[NOTIFYMSGSGS( {HC|CONSOLE} ) ]  
[MSGID([' ]msgid[' ]){NOAUTORREPLY}  
{DELAY(nnn{M|S}) REPLY([' ]replytext[' ] [, [' ]replytext[' ]...)]}]
```

•REPLY

- *replytext* is limited to 64 characters. If the reply contains blanks, non-alphanumeric characters, or is not to be folded to uppercase, enclose the reply in single quotes. The single quotes count towards the total of 64 characters.
- If the reply is too long for the WTOR requestor, the notification message CNZ2608I is issued when the WTOR is issued.
- Where message CNZ2608I appears depends on the setting of NOTIFYMSGSGS. It is displayed on consoles receiving routing codes 2 (operator information) or 10 (system programmer information), and also appears in the hardcopy log.
- Symbolic substitution supported

AUTORxx syntax example

```
/* **** */
/* $HASP070 SPECIFY RECOVERY OPTIONS ('RECOVER' OR 'TERMINATE' OR */
/*          'SNAP' AND, OPTIONALLY, ',NODUMP') */
/*          */
/* Rule: 2 */
/*          */
      Msgid(?HASP070)   Delay(30S) Reply(TERMINATE)
/* **** */
/* $HASP294 WAITING FOR RESERVE (VOL volser). REPLY 'CANCEL' TO */
/*          END WAIT */
/*          */
/* Rule: 1 */
/*          */
      Msgid(?HASP294)   Delay(30S) Reply(CANCEL)
/* **** */
```

Auto-Reply Policy Definition

Wildcards in MSGID

- Single character wildcard (question mark) is supported
- Wildcard may be for any character in the message id
- Multiple wildcards may be specified in a message id
- Useful for JES2 messages
 - First character (\$) can be installation defined
 - Poly-JES2 uses different first character for each instance
 - Use of ? in Auto-Reply policy provides support for any flavor of JES2 message

```
Msgid(?HASP070) Delay(30S) Reply(TERMINATE)
```

Auto-Reply Policy Definition

Use of symbolics

- Specify in quotes and in uppercase
- Symbolic to be resolved when the reply is issued:
 - Specify two ampersands
`REPLY('&&LWDAY' , ' ,&&LMON' , ' /&&LDAY' , ' /&&LYR4')`
 - the policy will be defined as:
``&LWDAY ,&LMON /&LDAY /&LYR4'`
 - When the reply is issued
`SAT,12/25/2010`
- Symbolic resolved at policy definition:
 - Specify one ampersand
`REPLY('&LWDAY' , ' ,&LMON' , ' /&LDAY' , ' /&LYR4')`
 - The policy will be defined and Auto-Reply will always reply:
`WED,12/01/2010`

Auto-Reply Policy Definition

- Default policy in parmlib member AUTOR00 (supplied by IBM)
 - Not recommended for customer to modify AUTOR00
- The policy will be activated during IPL if an AUTOR00 member is found in parmlib
- If there is no AUTOR00 member in parmlib, auto-reply is not activated and the following messages are produced:

```
CNZ2600I AUTO-REPLY POLICY ATTEMPTING TO USE AUTOR=00.  
IEA301I AUTOR00 NOT FOUND IN PARMLIB  
CNZ2601I AUTO-REPLY POLICY NOT ACTIVATED.  
NO ENTRIES SPECIFIED
```

- Not recommended that you remove AUTOR00 from parmlib, service or new releases might reinstall AUTOR00
- Activated before XCF is initialized

AUTOR parameter in IEASYSxx

- System parameter AUTOR= specified in IEASYSxx parmlib member or in response to message IEA101A SPECIFY SYSTEM PARAMETERS

The parameter options are:

```
AUTOR={xx }
```

```
AUTOR={{xx[,xx]...}}
```

```
AUTOR={OFF }
```

```
AUTOR={{(OFF) }
```

If AUTOR= is specified in response to IEA101A SPECIFY SYSTEM PARAMETERS, the AUTOR= value overrides any AUTOR= specification in the IEASYSxx parmlib member.

Value Range: Any two characters (A-Z, 0-9, @, #, and \$).

Default Value: AUTOR=00

Use of AUTOR Parmlib member

- Installation can provide set of parmlib members to contain Auto-Reply policy, specified in IEASYSxx.
 - AUTOR=(00,01,02)
- Installation can request that Auto-Reply processing not be activated via IEASYSxx.
 - AUTOR=OFF
- Default policy will be used unless requested not to activate or another policy is specified
 - NO AUTOR parameter coded in IEASYSxx then default policy is used if it exists (AUTOR00)
- Customer can add their own policy or override AUTOR00 policy
- Vendor products can provide their own parmlib members

Operator Commands

- Activate or change policy

```
SET AUTOR=(xx,yy,zz,...)
```

- Display policy or outstanding WTORs being monitored

```
DISPLAY AUTOR[,POLICY][,WTORS]
```

- Turn Auto-Reply off

```
SETAUTOR OFF
```

- Cause Auto-Reply to stop monitoring an outstanding WTOR

```
SETAUTOR IGNORE,RID={rpId}{0,SYS=sysname}
```

- If Auto-Reply is going to provide a reply to a WTOR but the operator wants more time to determine what should be done, the SETAUTOR IGNORE command can be used to tell Auto-Reply to leave the WTOR alone. It is then up to the operator to provide a reply.
- The system uses a reply id of zero when there are no free reply ids to use. There can be a max of one reply id zero per system in the sysplex. If a reply id of zero is to be ignored, the name of the system where the WTOR is outstanding must also be specified.



Default Auto-reply Policy

D AUTOR,P
 CNZ2603I 14.09.56 AUTOR POLICY 234
 POLICY ACTIVATED AT 09.13.26 ON 07/06/2010 NOTIFYMSG(S)(HC)
 FROM PARMLIB MEMBERS 00
 --MSG ID-- DELAY MEM ----REPLY TEXT----

ANTU2220D	60S	00	C
ANTX8925A	60S	00	N
ANTX8926A	60S	00	N
ANTX8942A	60S	00	N
ANTX8943A	60S	00	N
ANTX8944A	60S	00	N
ANTX8973A	60S	00	N
ANTX8978A	60S	00	N
ANTX8981A	60S	00	N
ARC0310A	15M	00	N
ARC0311A	15M	00	N
ARC0314A	15M	00	Y
ARC0346A	30S	00	Y
ARC0380A	60S	00	CANCEL
ARC0387A	60S	00	N
ARC0505D	60S	00	Y
ARC0803A	60S	00	Y
ARC0825D	60S	00	Y
ARC0962A	30S	00	C
ARC6254A	30S	00	N
BPXI078D	60S	00	N
BPXI083D	30S	00	N
BPXM055D	60S	00	N
BPXM061D	60S	00	N
BPXM063D	60S	00	N
BPXM120D	60S	00	N

CBR9810D	60S	00	QUIT
CNZ9009D	60S	00	N
CPO4205I	60S	00	2
CPO4206I	60S	00	2
EDG0103D	60S	00	RETRY
EDG1107D	2M	00	RESTART
EDG1200D	30S	00	CANCEL
EDG1203D	30S	00	CANCEL
EDG2103D	30S	00	L
EDG2106D	30S	00	L
EDG3213D	30S	00	RETRY
EDG4000D	2M	00	RETRY
EDG4001D	2M	00	RETRY
EDG4010D	60S	00	RETRY
EDG8008D	60S	00	RETRY
EDG8010D	60S	00	RETRY
EDG8011D	60S	00	RETRY
EDG8013D	60S	00	RETRY
EDG8102D	60S	00	RETRY
EDG8108D	60S	00	RETRY
EDG8110D	60S	00	RETRY
EDG8113D	60S	00	RETRY
EDG8121D	60S	00	RETRY
EDG8122D	60S	00	RETRY
EDG8123D	60S	00	RETRY
ERB306D	60S	00	GO
IAT2855	60S	00	RETRY
IAT3155	60S	00	CONTINUE
ICH15041A	30S	00	CANCEL
IEA029D	60S	00	YES
IEA893A	30S	00	NO

IEE799D	30S	00	CANCEL
IEE800D	60S	00	NO
IEF739D	30S	00	TERM
ISG017D	60S	00	NO
ISG027D	60S	00	NO
ISG082D	60S	00	NO
ISG101D	60S	00	NO
ISG117D	60S	00	NO
ISG186D	60S	00	KEEP
ISG220D	60S	00	C
ISG366D	60S	00	C
ISG880D	60S	00	C
IXC222D	60S	00	U
IXC289D	60S	00	U
IXC394A	60S	00	N
IXC403D	60S	00	W
?HASP070	30S	00	TERMINATE
?HASP294	30S	00	CANCEL
?HASP360	60S	00	N
?HASP457	60S	00	Y
?HASP811	30S	00	Y



SETAUTOR command examples

Assume the following VARY command is issued to generate message IEE800D:

- V 3D0,OFFLINE,FORCE

```
0009 IEE800D CONFIRM VARY FORCE FOR 3D0 - REPLY NO OR YES
```

- If message IEE800D (reply ID 0009) is to be ignored by auto-reply processing, issue:

- SETAUTOR IGNORE,RID=9

```
CNZ2607I AUTO REPLY WILL NO LONGER OCCUR FOR THE FOLLOWING WTOR: 0009  
06.46.55 2008130 IEE800D CONFIRM VARY FORCE FOR 3D0 - REPLY NO OR YES
```

- If auto-reply processing is to be deactivated, issue:
- SETAUTOR OFF

D AUTOR,W example

If WTORs IEE800D and IXC371D are issued and are waiting to be replied to by auto-reply processing, the output of the D AUTOR,W command

```
D AUTOR,W
```

```
CNZ2604I 06.47.02 AUTOR WTORS
```

```
0009 STATUS=06.47.58 SYS=SY1
```

```
MSG=IEE800D CONFIRM VARY FORCE FOR 3D0 - REPLY NO OR YES  
REPLY=NO
```

```
0008 STATUS=06.47.46 SYS=SY1
```

```
MSG=IXC371D CONFIRM REQUEST TO VARY SYSTEM SY1 OFFLINE. REPLY  
SYSNAME=SY1 TO REMOVE SY1 OR C TO CANCEL.  
REPLY=C
```

How to tell if Auto-Reply is monitoring a WTOR

- A new option in the Message Processing Facility (MPF) has been added
 - MPFLSTxx parmlib member
 - .MSGCOLOR **AUTOR**(c,h,i)
 - Allows WTORs monitored by Auto-Reply to be displayed in a particular color, highlighting or intensity
 - By default, Auto-Reply WTORs are not displayed any different than normal WTORs
 - For example AUTOR(G,R,H)
 - Green, Reverse Video, High Intensity
- Response to DISPLAY AUTOR,WTORS command lists outstanding WTORs being monitored

How to tell if Auto-Reply is monitoring a WTOR

- Message CNZ2605I (by default issued to hardcopy)

```
CNZ2605I At 09.35.17 the system will automatically reply: NO to the following WTOR: 0016 IEE800D  
CONFIRM VARY FORCE FOR 1800 - REPLY NO OR YES
```

- Response to DISPLAY REQUESTS command (message IEE112I)

- First character of message text is:

- * WTOR issued by an authorized program (existing)
- @ WTOR issued by a problem program (existing)
- & WTOR issued by an authorized program and Auto-Reply is monitoring (new)
- % WTOR issued by a problem program and Auto-Reply is monitoring (new)

How to tell if Auto-Reply is monitoring a WTOR

- SDSF adds support for the Consoles Auto-Reply function with new columns and a new action character on the SR panel.

- New columns on the SR panel
 - AutoReply
 - AutoRDelay
 - AutoReplyTime
 - AutoReplyText

- New action character AI (Ignore auto reply for the message)

Error Detection

- Auto-Reply will attempt to detect if the reply is too long for the requestor.
 - If so, the reply is not issued, a warning message is issued (by default to hardcopy) and WTOR is no longer monitored
 - If the reply contains symbolics, there is no length detection

```
CNZ2608I REPLY FOR WTOR msgid IS TOO LONG FOR REQUESTOR.  
AUTO-REPLY WILL NOT PROCESS THIS WTOR:  
rpid msgtxt
```

Error Detection

- Auto-Reply attempts to detect a “run away”
 - If Auto-Reply attempts to reply to the same WTOR more than 20 times in a second, monitoring is stopped.
 - For example:
 - Policy contains a typo so the reply is always invalid
 - Policy requests a delay of zero
 - WTOR loop (“run away”) occurs

```
CNZ2609I RECURSIVE AUTO-REPLY DETECTED FOR MESSAGE xxxxxxxxxxxx.  
AUTO-REPLY WILL NOT PROCESS THIS WTOR  
rpid msgtext
```

Messages

- When Auto-Reply replies during NIP
`CNZ2602I REPLY TO 00 IS:replytext <- Auto replied`
- Post NIP
`CNZ2605I At hmmmss the system will automatically
reply:replytxt to the following WTOR:
rpid msgtxt`

`CNZ2606I System has automatically replied:
replytxt to the following WTOR:
rpid hmmmss yyyyddd msgtxt`

Resources

- **MVS Planning: Operations, SA22-7601**
 - Overview

- **MVS Initialization and Tuning Reference, SA22-7592**
 - AUTORxx member of parmlib

- **MVS Systems Commands, SA22-7627**
 - Display AUTOR
 - SETAUTOR

Auto IPL



z/OS
**A smarter operating system
for smarter computing**

The advertisement banner features a dark blue background with a glowing globe on the right side. The globe is illuminated from behind, creating a bright blue glow and radiating light rays. The text on the left is in white, bold, sans-serif font.

AutoIPL

- AutoIPL support provides the capability to request (after a wait state is requested by a system component) that the system automatically:
 - IPL standalone dump
 - IPL of z/OS
 - IPL standalone dump and IPL z/OS
- The AutoIPL function improves the serviceability of z/OS
 - Take a SADMP with pre-defined parameters (so that the operator doesn't have to make decisions while z/OS is in a wait state), achieving faster failure data capture
 - The AutoIPL function improves the availability of z/OS by reducing the time to re-IPL (either following a SADMP or instead of a SADMP), achieve fast recovery of systems

AutoIPL

- ❑ AutoIPL can reduce latency of operator response time by automatically initiating a dump to capture data for analysis and a restart based on z/OS diagnostics
- ❑ z/OS customers have requested that dumping for disabled wait state problems be driven without operator intervention due to the infrequency of taking SADMPs
- ❑ AutoIPL provides an automated policy-driven means to allow the system to initiate either a SADMPs, re-IPL or both.
 - Improved serviceability via quick, automatic SAD initiation
 - Improved availability via quick, automatic re-IPL of z/OS
 - Minimizing downtime in the event of a disabled wait state

AutoIPL Policy

- ❑ AutoIPL policy is specified via DIAGxx parmlib member
- ❑ A hard-coded wait state action table designates disabled wait state codes and reason codes for which AutoIPL processing should or should not be performed, and which actions to be taken
- ❑ Most non-restartable disabled wait states/reasons will be eligible for AutoIPL processing
- ❑ VARY XCF commands support new options to:
 - IPL stand alone dump
 - Re-IPL z/OS or both after the target image has been partitioned out of the sysplex

DIAGxx Parmlib Member

AUTOIPL SADMP(sadmp info) MVS(mvs info)

□ sadmp info

- Specifies either (device,loadparm) or NONE
- When z/OS is about to enter a wait state, SADMP will be loaded from this volume with this load parameter - If NONE is specified an SADMP will not be taken
- Specifying loadparm will execute without operator intervention
- **AUTOIPL SADMP(D056,SNSYSC4) MVS(LAST)**

DIAGxx Parmlib Member

AUTOIPL SADMP(sadmp info) MVS(mvs info)

❑ mvs info

- Specifies either (device, loadparm) or (LAST) or (NONE)
- When z/OS is about to enter a wait state, z/OS will be IPLed from this device using this load parameter, either immediately or following the completion of SADMP processing (if SADMP with a device and load parameter was also coded)
- If MVS with LAST was specified, MVS will be IPLed from the same device and load parameter used for the current IPL, either immediately or following the completion of SADMP processing, (if SADMP with a device and load parameter was also coded).
- If MVS with NONE was specified, MVS will not be IPLed, either immediately or following the completion of SADMP processing (if SADMP with a device and load parameter was also coded).

Using AutoIPL Policy

□ Activate AutoIPL policy

- IPL with a DIAGxx parmlib that specifies AUTOIPL options

□ To deactivate the AutoIPL policy

- Specificity of SADMP(NONE) MVS(NONE) in DIAGxx and issue, then
 - SET DIAG=xx

□ To display the AutoIPL Policy in effect

- Issue z/OS command
 - D DIAG

AutoIPL Processing with AutoIPL Policy

- ❑ AutoIPL processing is performed when LOADWAIT is entered and compares the wait state code and the reason code against the Wait State Action Table (WSAT) to determine whether to initiate a SADMP and/or initiate IPL of z/OS
- ❑ In order for AutoIPL processing to occur:
 - ❑ the wait state that is being loaded on the failing system must be eligible for some form of AutoIPL processing as designated in the WSAT, and
 - ❑ AutoIPL processing must be enabled by the customer as designated in the DIAGxx AutoIPL Policy
 - ❑ If policy includes taking a SADMP, SADMP should be generated on a specific volume

AutoIPL

- **Wait State Action Table (WSAT)** – part of nucleus
- Hardcoded table of wait state codes and reason codes with specific AutoIPL actions designated for them
 - Possible actions include:
 - None. Let the system enter its wait state condition. This is the default if no AutoIPL policy is specified.
 - Initiate Standalone Dump.
 - Initiate Standalone Dump, followed by IPL of z/OS
 - When the SADMP completes, it checks whether to IPL z/OS and obtains the z/OS IPL information to IPL z/OS
 - Initiate IPL of z/OS only.

AutoIPL

- WSAT is not a programming interface, now its externalized so customers can decide whether to create a policy and activate the AutoIPL function
- Each wait state and reason code entry has a flag to indicate whether the SADMP part of the policy should be honored, and a flag to indicate whether the MVS part of the policy should be honored

WSAT Table

- ❑ The WSAT contains four bytes entries for pre-determined wait states and reason codes.
- ❑ Looking at the four byte entries as eight 4-bit “nibbles”:
- ❑ Nibble 1 - contains the AutoIPL action in the last two bits:
 - b'0000' - no action
 - b'0010' - StandAlone dump
 - b'0001' - IPL z/OS
 - b'0011' - StandAlone dump followed by IPL of z/OS
 - b'1000' – any reason code for this wait state is a match
- ❑ Nibbles 2 to 5 - contain the reason code
- ❑ Nibbles 6 to 8 - contain the wait state code

WSAT Table Entries

WSAT entry	Wait state-reason code	AutoIPL action
X'000040A2'	0A2-0004	No action
X'1017C0A2'	0A2-017C	IPL z/OS
X'201800A2'	0A2-0180	SADMP
X'301840A2'	0A2-0184	SADMP and IPL z/OS
X'200010B5'	0B5-0001	SADMP
X'200020B5'	0B5-0002	SADMP
X'A0000001'	001-0000	SADMP
X'A0000007'	007-0000	SADMP
X'A0000008'	008-0000	SADMP
X'A0000009'	009-0000	SADMP
X'A0000010'	010-0000	SADMP
X'A0000037'	037-0000	SADMP
X'A0000039'	039-0000	SADMP
X'A0000056'	056-0000	SADMP
X'A0000079'	079-0000	SADMP

Overview – Restriction removed

- The PTF [UA45446](#) for ([APARs OA26993/OA26995](#)), along with underlying LPAR firmware support, now enables AutoIPL to be used in configurations where an SFM policy is active.
 - With this support, requested AutoIPL actions will be performed in accordance with the DIAGxx parmlib member, even when an SFM policy is active in the sysplex.
 - Use of SFM and system isolation to quickly and automatically remove a failed system from the sysplex remains a highly-recommended "best practice" for sysplex availability.
 - AutoIPL might delay for several minutes before actually initiating the SADMP or the re-IPL. During this time, the failed system will appear to be hung.

AutoIPL and SFM Policy LIC Levels

- ❑ The use of AutoIPL on a system in a multisystem-capable sysplex configuration where an SFM policy is active requires the following LIC levels be installed:
 - For z9 Systems at Driver-67: MCL006 in EC Stream (LPAR) G40954 (Bundle 38) – Feature code 9904
 - For z10 Systems at Driver-73: MCL009 in EC Stream (LPAR) F85901 (Bundle 45b)
 - For z10 Systems at Driver-76: MCL003 in EC Stream (LPAR) N10965 (Bundle 8)

Important: The AutoIPL function must not be used in a sysplex with SFM active until both the LPAR firmware support and the software support have been installed on all affected systems

VARY XCF Command Support

- ❑ VARY XCF command supports new options to IPL a standalone dump, re-IPL z/OS or both
 - Two new optional keywords, REIPL and SADMP, are accepted on the VARY XCF command
- ❑ VARY XCF is used to request the removal of a system from the sysplex via sysplex partitioning
 - The new keywords indicate that when a system is varied out of the sysplex, AutoIPL processing will be done when VARY XCF loads a wait state on that system
 - This assumes that the target system is not already in some other wait state, and has not already been system reset, manually re-IPLed, or fenced at the time of the VARY XCF processing

VARY XCF Command Support

- ❑ The new syntax of VARY XCF is:
 - VARY XCF, sysname, OFFLINE [, FORCE] [RETAIN=YES|NO] [, SADMP] [, REIPL]
- ❑ Either SADMP or REIPL, or both, may be requested after the target image is partitioned out of the sysplex
- ❑ SADMP and REIPL options are not allowed with FORCE
- ❑ The REIPL option is not allowed with RETAIN=NO
- ❑ VARY XCF with Auto-IPL options is rejected if the target system image does not support the requested options and message IXC372I is issued with a new insert:
 - ❑ DOES NOT SUPPORT THE REQUESTED AUTOIPL OPTION(S)

Parallel Subsystem Initialization



Mean Time To Recovery

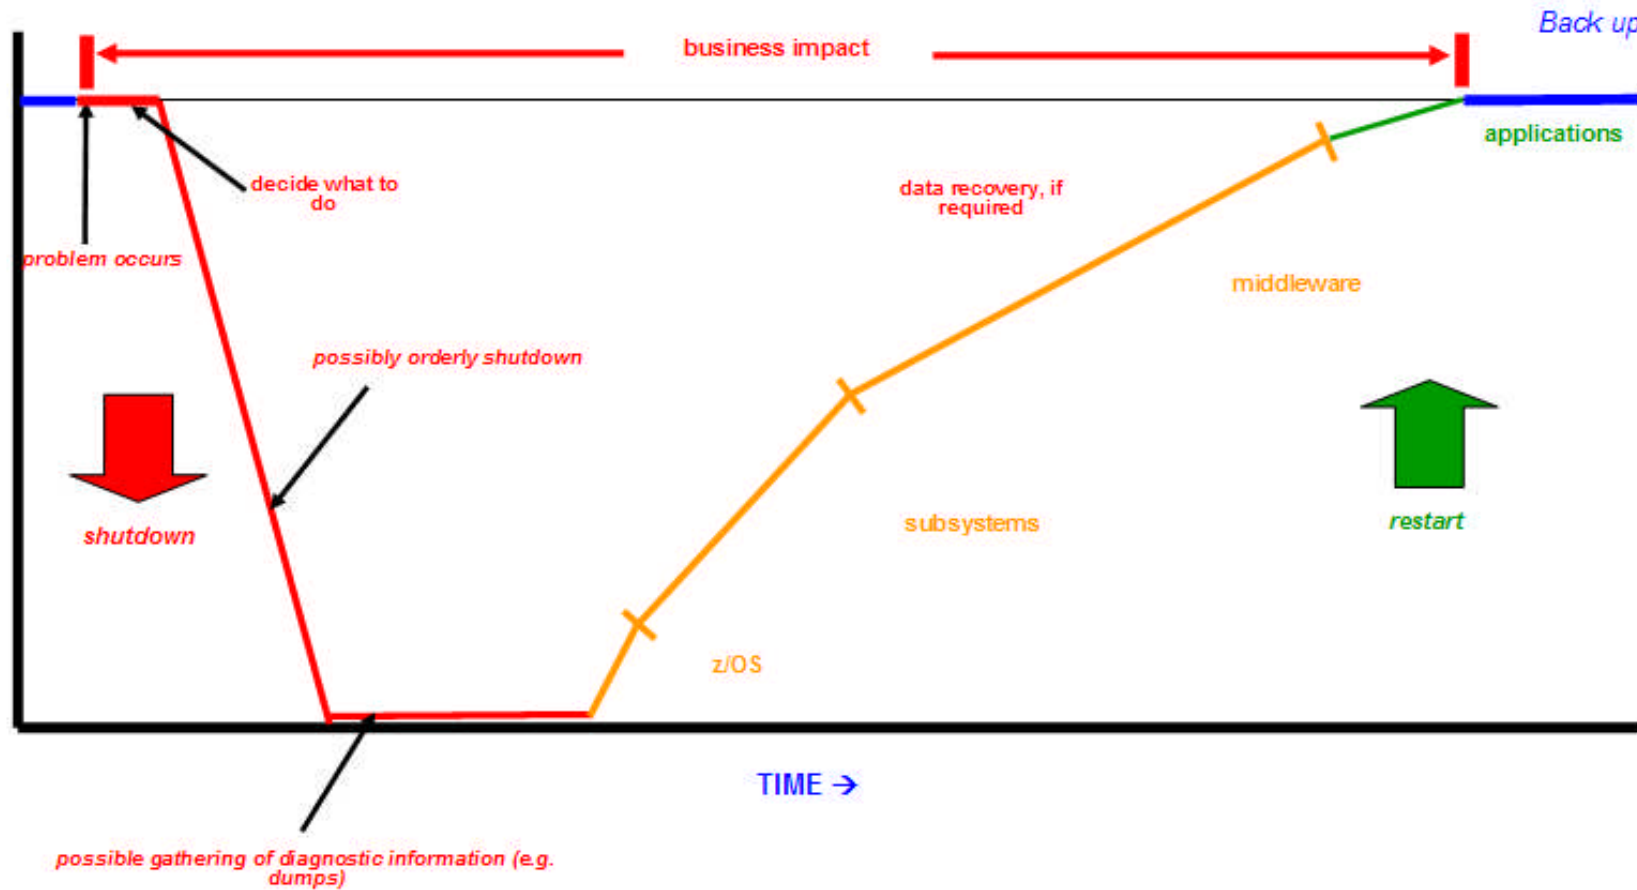
■ Problem Statement / Need Addressed:

- Mean Time To Recovery (MTTR)
 - The time from when problem is detected until system is back doing productive work.

Goal:

- Reduce the time that it takes the operating system, subsystems and middleware to go down and come back up
 - Reducing initialization path lengths where possible
 - Exploiting parallelism where possible
- Reduce the cost of operating system services used by subsystems and middleware for initialization

MTTR



Although the emphasis is on reducing restart time, reducing the time to shutdown and gather diagnostic information is also critical to reducing the overall duration of the outage.



Parallel subsystem initialization

- **Problem Statement / Need Addressed:**
 - Subsystem Initialization Routines identified as having potential MTTR issues

- **Solution:**
 - Allow SSI Initialization Routines to be run in parallel

- **Benefit:**
 - Subsystem initialization time may decrease

Subsystems

- Subsystems are defined in IEFSSNxx parmlib member
 - Optionally, an initialization routine can be specified with the subsystem definition

- Initialization routines invoked as IEFSSNxx definitions are processed
 - Routines invoked in order specified in IEFSSNxx
 - Routines are invoked serially
 - Routines run under a never ending task

BEGINPARALLEL

- To support parallelism, IEFSSNxx now supports a new keyword (BEGINPARALLEL)
- All Subsystem Initialization Routines specified after BEGINPARALLEL will be invoked in parallel.
- Initialization routines specified before BEGINPARALLEL (or if keyword is not specified), will continue to be invoked serially.
- Benefit depends on
 - Number of initialization routines
 - Complexity of routines
 - Serialization requirements of routines
 - Available CPs

BEGINPARALLEL

- Should be specified after the SMS definition
 - If not, message IEFJ009E issued and BEGINPARALLEL honored
 - **IEFJ009E BEGINPARALLEL KEYWORD SPECIFIED BEFORE SMS SUBSYSTEM DEFINITION**
- Has **no** effect on subsystems that don't specify an initialization routine
- If specified multiple times or in different concatenated IEFSSNxx members, all but the first will be rejected with message ASA011I

**ASA011I ERROR IN PARMLIB MEMBER=IEFSSNR2 ON LINE
2,POSITION 1: DUPLICATE SPECIFICATION OF 'BEGINPARALLEL'
FIRST SPECIFICATION IS USED. DETECTING MODULE IS IEFJPACT.
INPUT LINE: BEGINPARALLEL**

IEFSSNxx

```
SUBSYS SUBNAME(SMS)
INITRTN(IGDSSIIN)
INITPARM(ID=ZX)
SUBSYS SUBNAME(JES2) /* JES2 AS PRIMARY SUBSYSTEM */
PRIMARY(YES) START(YES)
BEGINPARALLEL
SUBSYS SUBNAME(RACF) /* RACF SUBSYS */
  INITRTN(IRRSSI00)
  INITPARM('#,M')
SUBSYS SUBNAME(LOGR) /* LOGR */
  INITRTN(IXGSSINT)
SUBSYS SUBNAME(RRS) /* RESOURCE RECOVERY SERVICES */
SUBSYS SUBNAME(ID9Y) /* IRLM DB2 9.1 FOR DB9Y */
SUBSYS SUBNAME(DB9Y) /* DB2 V910 DB9Y */
  INITRTN(DSN3INI)
  INITPARM('DSN3EPX,-DB9Y,S')
```

BEGINPARALLEL considerations

For the z/OS Communications Server TNF and VMCF subsystems, subsystem definitions must be specified before you specify the BEGINPARALLEL statement if the INITRTN parameter is included with the subsystem definitions.

For IBM products or a vendor-supplied subsystem, check the product's installation or configuration documentation to determine the placement of the BEGINPARALLEL keyword.

Remember: BEGINPARALLEL only affects subsystem initialization routines. It has no affect on subsystems that do not specify a routine.

BEGINGPARALLEL considerations

Behavior Differences

- Execution order of initialization routines running in parallel now unpredictable
 - Routines **must not** have any execution order dependencies
- When running in parallel, routines execute under a task that terminates
- Subsystem owners should verify initialization routines can run under a task that terminates. Look for:
 - Data spaces created but not deleted (particularly CADS)
 - Task-related storage obtained but not released
 - ENQ obtained but not DEQ'd
 - ALESERV ADD without DELETE
 - ESTAE CREATE without delete
 - Joining XCF groups without leaving
 - Connections to Coupling Facility structures obtained but not released
 - Task level Name Tokens created without deleting

CNZ_MSIEXIT -- Master Scheduler Initialization Dynamic Exit

- Some vendor products use the SSI initialization routines as an “exit point”:
 - These “exits” need to complete before true subsystems are initialized
 - Parallelization could cause these products to have serious problems

- In order to keep those “exit points” out of the IEFSSNxx parmlib member process, a new dynamic exit point is provided during master scheduler initialization (MSI), after the security product is initialized and before the IEFSSNxx parmlib member is processed.

CNZ_MSIEXIT -- Master Scheduler Initialization Dynamic Exit

- In z/OS V1R12, a MSI dynamic exit routine can be used in the following ways:
 - Defined in PROGxx parmlib member
 - EXIT ADD EXITNAME(CNZ_MSIEXIT) MODNAME(xxx) PARAM(parm)
 - PROGxx is enhanced to allow parameter specification to be passed to exit routine
 - Code must reside in LPA, LNKLST or Nucleus
 - Dynamic allocation not available at this point in MSI
 - No exit routine return codes are processed
 - Once exit routine is invoked, exit point will become “undefined”
 - Releases storage occupied by the exit routine

Conclusion



z/OS
**A smarter operating system
for smarter computing**

The advertisement banner features a dark blue background with a glowing blue globe on the right side. The globe is illuminated from behind, creating a bright glow and lens flare effects. The text on the left is in a clean, white, sans-serif font.

End of Session

- Thank You for attending and for your patience
- Please email me if you have any questions
- While you complete the evaluation form, I'll spend a few minutes to share some information about the ITSO, IBM Redbooks and Residencies
- Thanks to all who knowingly or unknowingly contributed material for this presentation

Michael Grossman

Parwez Hamid

Karla Arndt

Bob Abrams

Mark Noonan

Ryotaro Sawada

Frank Kyne

Stephen Warren

Marna Walle

Kevin Kelley

Greg Daynes

Michael Ferguson

Rita Pleus


Mario Bezzi

John Troy




What is ITSO and what are IBM Redbooks?

- IBM Redbooks are developed and published by IBM International Technical Support Organization (ITSO) Global Content Services (GCS).
- No charge in-depth technical documents:
 - Positioning, guidance, implementation experiences
 - Solution scenarios.... How-to perform guides
 - IBM's largest job rotation program
 - ~250 projects & -700 authors/residents per year
- For IBM software, systems, and solutions
- IBM's second largest Facebook community (10K)
 - <http://www.facebook.com/home.php#!/IBMRedbooks>
- Download from: <http://www.redbooks.ibm.com/>



Corporate



2010 metrics

294 Redbooks published

6.72 million downloads

98% client satisfaction

Redbooks are creating using the **residency model**



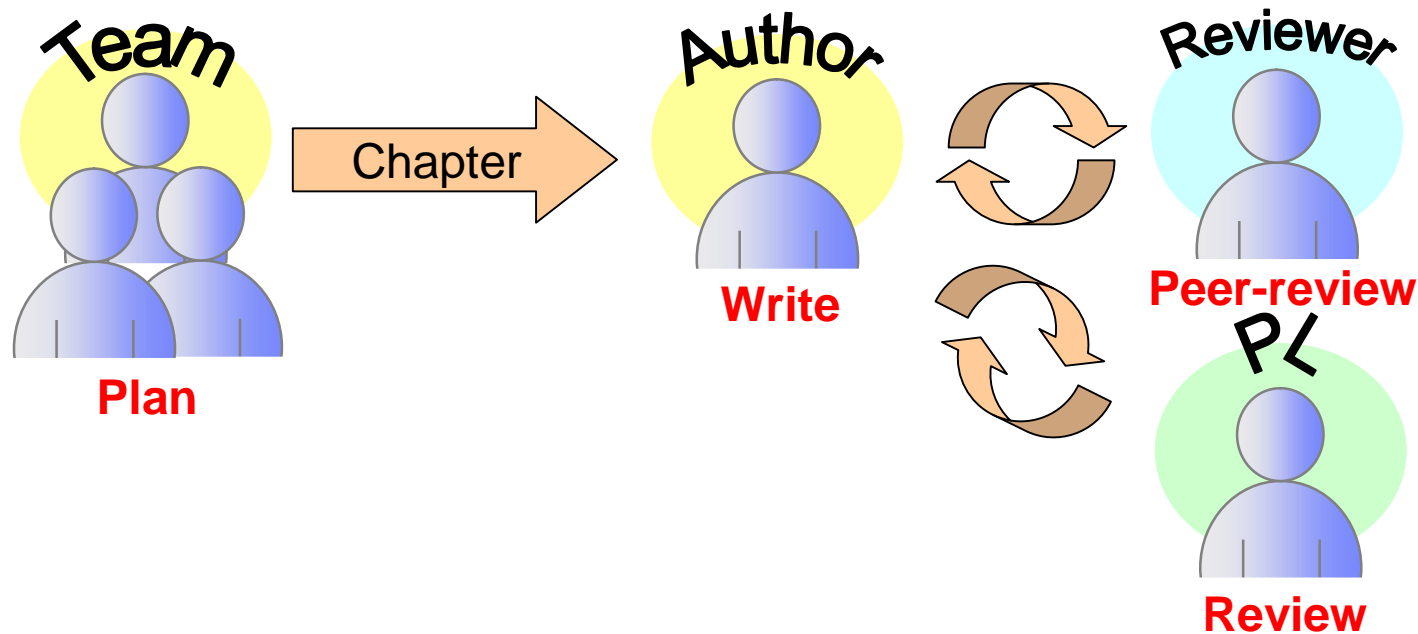
Residents
(subject matter experts)



ITSO Center
of Excellence

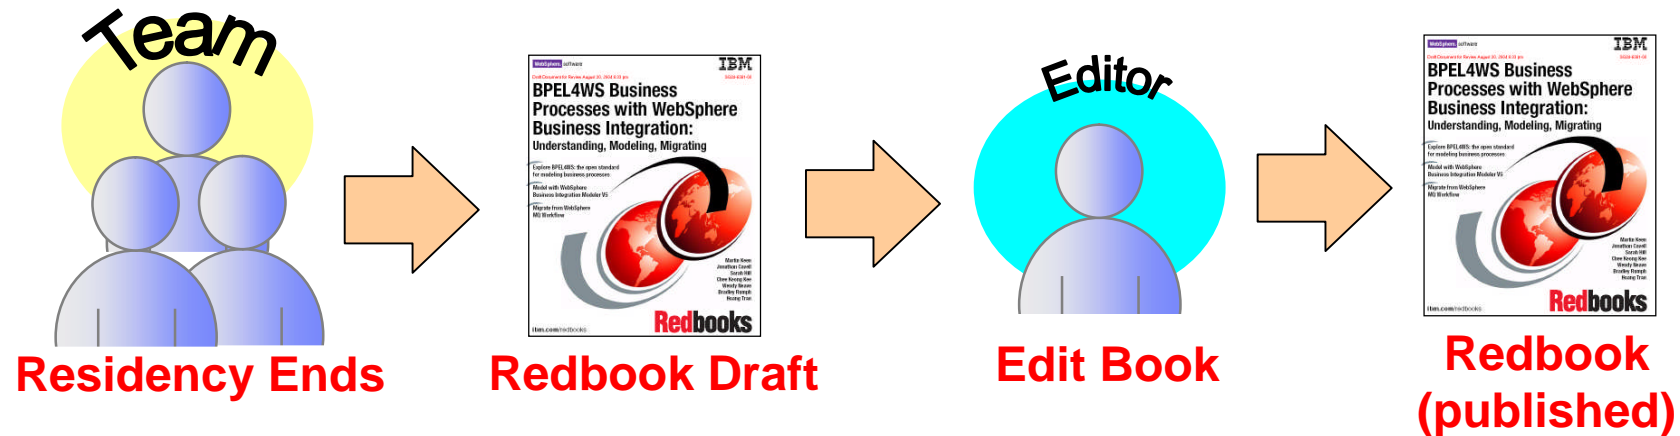
- Subject matter experts are selected for a residency and work full-time in a single location for the residency duration (3-6 weeks, most commonly 5 weeks)
- Some remote participation is also common
- A residency is typically run at an ITSO Center of Excellence (Raleigh NC or Poughkeepsie NY) but other worldwide locations can also be used by special request

Residency process



- The residency team plan the content of the Redbook based on sponsor feedback, breaking the work down into chapters
- A resident is responsible for authoring one or more chapters
- These chapters are reviewed by the Project Leader, a Peer Reviewer (another resident on the team), and Technical Reviewers (external to the team)

Redbook process



- After the Redbook ends, the Project Leader prepares the Redbook to be published as a **Redbook Draft**
 - Redbook Drafts are published externally, with the caveat that they may contain small technical and grammatical errors
- Any additional technical review comments are incorporated into the Redbook Draft
- Content is then frozen, and the book is assigned to the IBM Publication Team for editing
 - This editing process ensures the book meets IBM’s publication guidelines
- The Redbook Draft is then replaced with the edited, final Redbook publication

Participating in a Redbooks residency (1/3)

- Redbook residencies are posted on the Redbooks Web site
- <http://www.redbooks.ibm.com/redbooks.nsf/home?ReadForm&page=residencies>
- Redbooks weekly newsletter (drafts, Redbooks, and residencies)
 - <https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

Country/region [select]

Redbooks

Home Solutions Services Products Support & downloads My IBM

IBM Redbooks®

Advanced Search

Software

Storage

Systems & Servers

Solutions

IT Business Perspectives

Residencies

Workshops

Additional Materials

How to order

About Redbooks

Contact us

Newsletter

RSS feeds

Related links

- IBM Publications
- Technical training
- Developers
- developerWorks wikis
- IBM Business Partners
- IBM Press books

IBM Redbooks >

IBM Redbooks Residencies

Index of Residencies

IBM Redbooks are developed through a unique residency program that teams IBM technical professionals with Business Partners, clients, and IBM product development staff. Depending on the market needs of a product or solution, a team of residents is chosen through a competitive nomination process and works at one of the ITSO centers for two to eight weeks, devoted to developing IBM Redbooks. **Would you like to join a team writing an IBM Redbooks publication? Find an upcoming residency in the list below and submit an online nomination!**

Note: The ITSO pays authorized travel and living expenses for all residents, but does not compensate for the time or services of its residents. For IBMers, ITSO residency travel is pre-approved. See [How Residencies Work](#) for more details. .

1 to 30 of 37 documents ← Previous | Next →

1. [SAN Volume Controller: Best Practices and Performance Guidelines, ST-1S12-R01](#) **Closed**
Starts 11 Jul 2011, ends 19 Aug 2011 (6 weeks). **Accepting back-up nominations.**
2. [Creating an Integrated Web Experience \(with Connections, Quickr, Portal, Sametime, and Domino\), LO-1W01-R01](#) **Closed**
Starts 25 Jul 2011, ends 09 Sep 2011 (7 weeks). **Accepting back-up nominations.**
3. [Endpoint Protection and Management Design Guide using IBM Tivoli Endpoint Manager, TI-1S04-R01](#) **Closed**
Starts 25 Jul 2011, ends 26 Aug 2011 (5 weeks). **Accepting back-up nominations.**
4. [Products and Technology Book for New Offerings, BL-1201-R01](#) **Closed**
Starts 01 Aug 2011, ends 26 Aug 2011 (4 weeks). **Accepting back-up nominations.**
5. [Domino Development Best Practices, LO-1W03-R01](#) **Open**
Starts 08 Aug 2011, ends 09 Sep 2011 (5 weeks) and requires 6 residents.

IBM

August 8 2011

Redbooks Weekly Newsletter

Welcome to the IBM Redbooks weekly newsletter. This mailing lists the Redbooks, Drafts, Redpapers, Hints and Tips and IBM Press books published in the past week as well as the Residencies and Workshops announced during the same period. As always please visit the IBM Redbooks web site at: <http://www.redbooks.ibm.com> for more details and up to the minute technical information. For IBM Press books please visit <http://www.ibm.com/ibmpress>

Drafts

SAN Storage Performance Management Using Tivoli Storage Productivity Center
Revised: August 1, 2011
More details are available at <http://w3.itso.ibm.com/redpieces/abstracts/sq247364.html?Open>

IBM System Storage N Series Hardware Guide
Revised: August 3, 2011
More details are available at <http://www.redbooks.ibm.com/redpieces/abstracts/sq247840.html?Open>

Security Functions with DB2 10 for z/OS
Revised: August 5, 2011
More details are available at <http://www.redbooks.ibm.com/redpieces/abstracts/sq247959.html?Open>

Redbooks

Server Time Protocol Planning Guide
Revised: August 1, 2011 ISBN: 0738434310 216 pages
Explore the book online at <http://www.redbooks.ibm.com/abstracts/sq247280.html?Open>

Server Time Protocol Implementation Guide
Revised: August 1, 2011 ISBN: 0738434655 416 pages
Explore the book online at



Participating in a Redbooks residency (2/3)

- Residency announcements detail skills required

Skills Needed	
Skill Areas	Requested Skill Level
AIX Internals (Kernel, Device Drivers)	4
AIX Version 6.1 Knowledge	5
AIX Version 5.3 Knowledge	4
IBM Power System Hardware and Product Knowledge	4
Sales and Marketing Skill	3
AIX support of POWER6	4
Debugging skills	3

- Nominees fill out a nomination form, specifying their skills
- Project Leader reviews all nominations and sends out acceptance emails to the strongest candidates



Participating in a Redbooks residency (3/3)

- All selected residents stay in an apartment for the duration of the residency
- ITSO pays all travel and living costs
 - Residents receive one round trip flight, so weekend trips home are not funded
 - Some residents participate remotely and never travel
- Residents are provided with all required hardware and software:
 - A residency workstation
 - Adobe FrameMaker license (tool used to write Redbooks)
 - Lab servers where needed
- Residents extend their skills:
 - Network with lab experts, enhance their technical writing skills, gain author credibility, return home to teach others, etc.





THANK YOU!

