

Unleash the Capabilities of New Technologies with IMS Tools



Nick R. Griffin - IMS Tools Product Line Manager

Ron Bisceglia – IMS Tools Architect

IMS Technical Symposium 2015



Important disclaimer

© Copyright IBM Corporation 2015. All rights reserved.

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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. THE INFORMATION ON NEW PRODUCTS IS FOR INFORMATIONAL PURPOSES ONLY AND MAY NOT BE INCORPORATED INTO ANY CONTRACT. THE INFORMATION ON ANY NEW PRODUCTS IS NOT A COMMITMENT, PROMISE, OR LEGAL OBLIGATION TO DELIVER ANY MATERIAL, CODE OR FUNCTIONALITY. THE DEVELOPMENT, RELEASE, AND TIMING OF ANY FEATURES OR FUNCTIONALITY DESCRIBED FOR OUR PRODUCTS REMAINS AT THE SOLE DISCRETION OF IBM. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, **Information Management, IMS, and z/OS** are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml



Agenda

- Enabling Analytics for IMS Data
- How to manage Big Data for IMS
- Leveraging New Technologies
- Q&A



There are three important shifts fundamentally changing the way that decisions are made...

Data



Data is becoming the world's new natural resource

Cloud



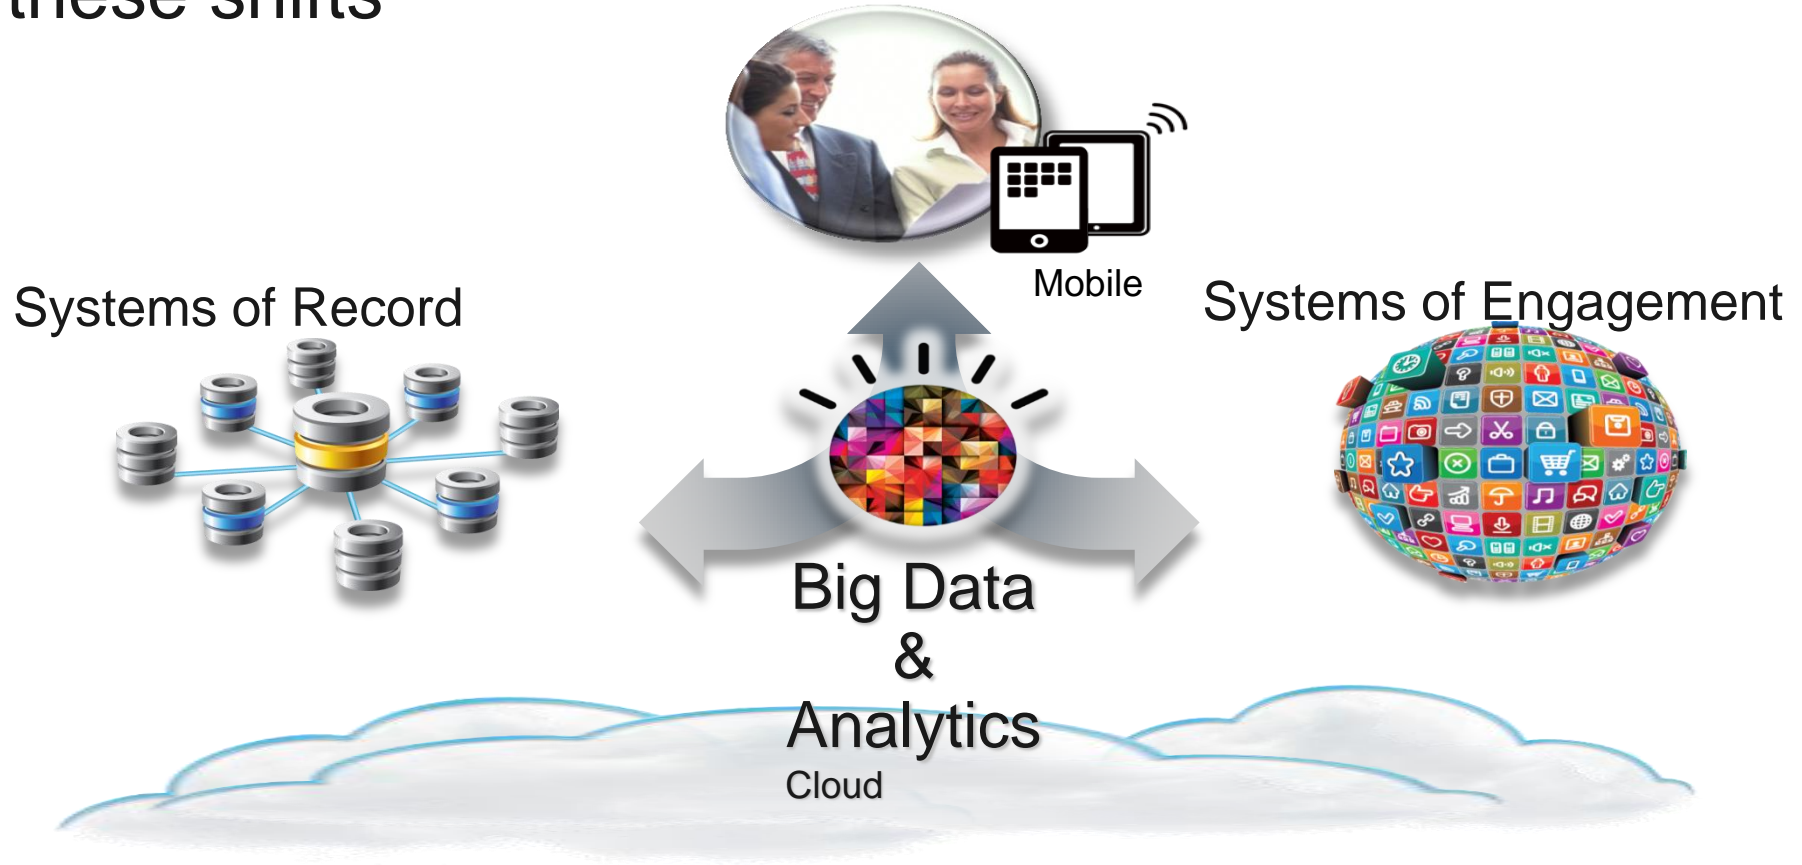
The emergence of cloud is transforming IT and business processes into digital services

Engagement



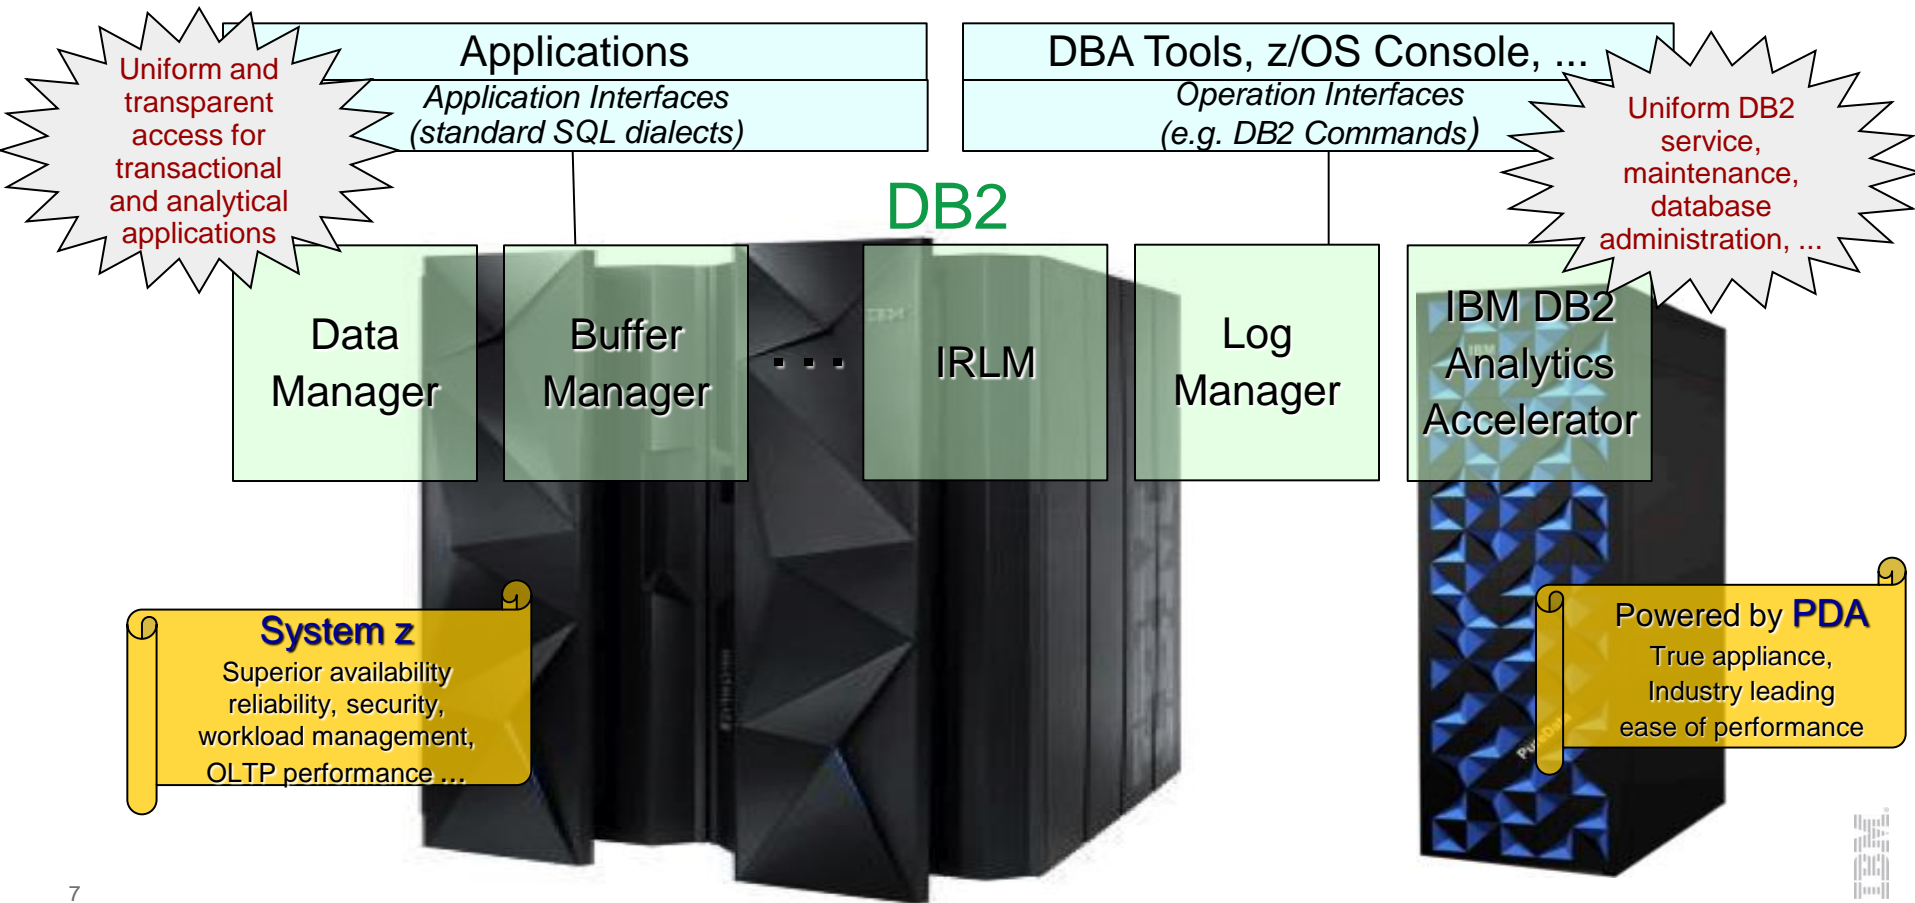
Social. Mobile. Security. Empowering people with knowledge, enriching them through networks and changing expectations.

Business and industries are being transformed by these shifts

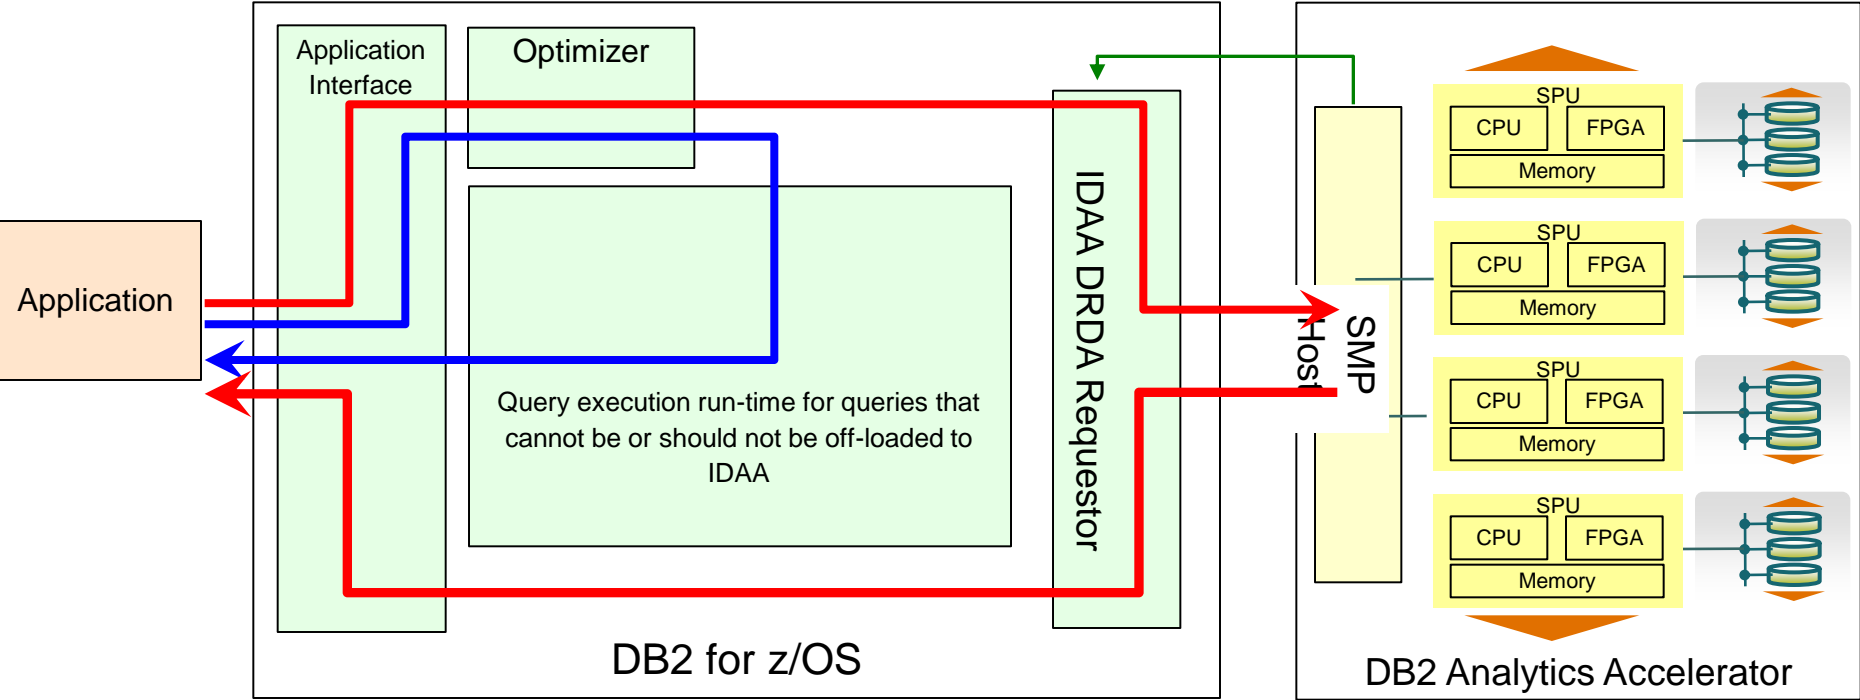






Analysis with Big Data for IMS

DB2 for z/OS Approach: Hybrid Database Management System

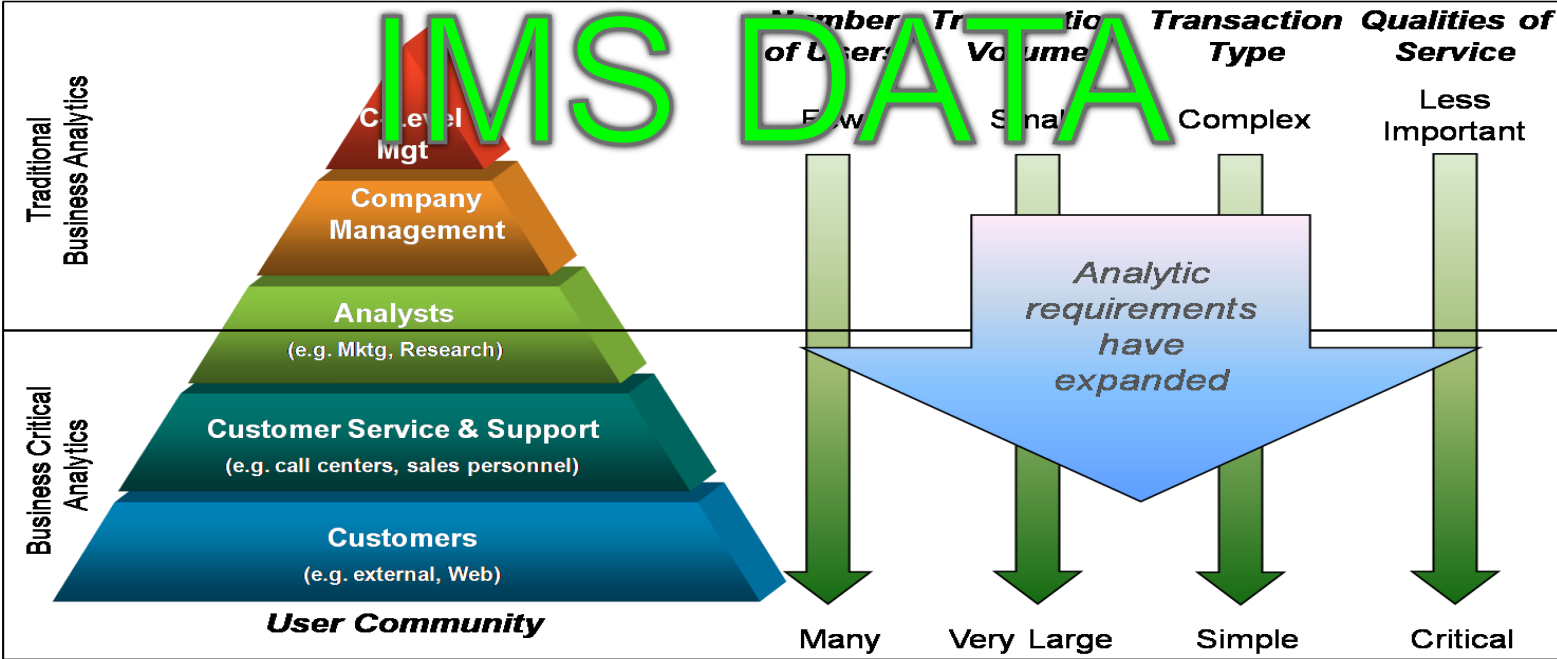


Query Execution Process Flow



-  Heartbeat (DB2 Analytics Accelerator availability and performance indicators)
-  Queries executed without DB2 Analytics Accelerator
-  Queries executed with DB2 Analytics Accelerator
-  Queries executed with value of "ALL" may receive a SQL Error Code if the query cannot run on the accelerator

More users across the organization want access to business critical analytics



History of IMS Analytics

- Desire to combine IMS data with other data
 - Social, DB2 z/OS data, SAS data, etc
- ETL IMS data into data warehouse
 - Mostly off z/OS
 - Data being sent to potentially many sources
- Security can be compromised
- Performance historically not keeping up without \$\$\$\$\$



Accelerate IMS Access - Proposed Solution

- Leverage Analytics Accelerator
 - Metadata resides in DB2
 - Copy IMS Data into Accelerator Only
- DB2 manages queries and controls access



Advantages:

- Data never leaves z/OS
- IMS workload unaffected
- Single server for z Analytics
- Join of IMS/DB2 data
- Less reason to ETL DB2/IMS data off platform


Basic Process

- Decide IMS data needed
- Decide extraction and mapping tools and process
- Currency required (Refresh Frequency)
- Map IMS data to relational model
- Create DB2 table that matches extracted record format
- Add table to Accelerator
- Extract IMS data
- Load extracted data to DB2 table
- Load data from DB2 into the accelerator
- Enable DB2 version of table for acceleration




Extraction Considerations and Methods

- Considerations
 - Availability requirements
 - Frequency of refresh?
 - Impact to OLTP workload
 - What data is needed?
 - Entire database, certain segments, multiple DBs?
 - Consistency of data?



Extraction Considerations and Methods

- Extraction Tools and Methods
 - Custom IMS Application
 - Additional online workload
 - Data can still be changing
 - Database Clone (IMS Cloning Tool)
 - Group of databases at a point in time
 - Image Copies/Unload Files
 - Additional knowledge of structure needed
 - Mapping and ETL Tools
 - IMS Explorer
 - Data Stage, Informatica
 - Data Virtualization
 - IMS Catalog via JDBC
 - Other tools



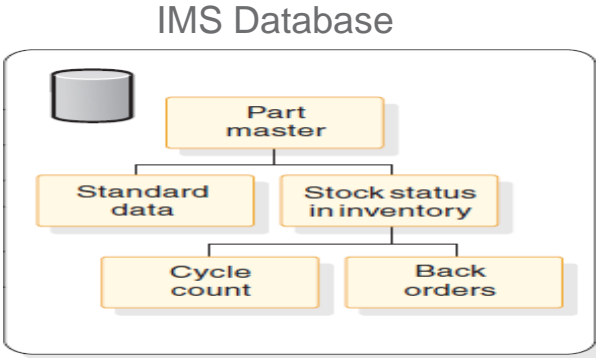
Mapping and Transforming Data

- Segment -> Table
 - Field -> Column

- Data type not required by IMS
 - Many times FIELD only defined for sequence fields
 - Data content not enforced by IMS
 - Data cleansing required?

- Where are field descriptions defined?
 - IMS Catalog
 - Copy books
 - JAVA Classes

- Non-unique or non-keyed segments



DB2 Tables

&schema.Part Master

Part_No	Part_Description
---------	------------------

&schema.Standard Info

Process Code	Invoice Code	Cost Center	PM_Part_no
--------------	--------------	-------------	------------

&schema.Stock Status

Area	Dept	Project	Division	PM_Part_no
------	------	---------	----------	------------

⋮



Mapping IMS Data to Tables

DBD:DI21PART

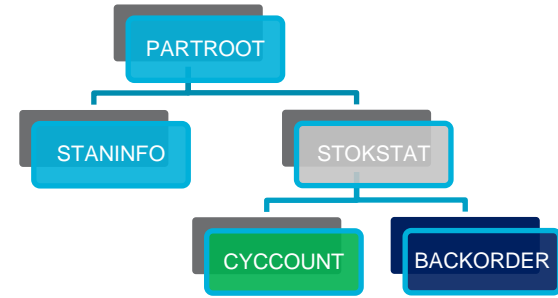


Table-name
PARTROOT
Column-names
PARTKEY
...

Table-name
STANINFO
Column-names
STANKEY
...

Table-name
STOKSTAT
Column-names
STOCKKEY
...

Table-name
CYCCOUNT
Column-names
CYCLEKEY
...

Table-name
BAKCORDER
Column-names
BACKKEY
...

DataBase Definition (DBD)

```

DBD NAME=DI21PART,ACCESS=(HISAM,VSAM)
DATASET DD1=DI21PART,DEVICE=3380,OVFLW=DI21PARO,
SIZE=(2048,2048),RECORD=(678,678)
SEGMENT NAME=PARTROOT, PARENT=0,BYTES=50, FREQ=250
FIELD NAME=(PARTKEY,SEQ),TYPE=C,BYTES=17,START=1
SEGMENT NAME=STANINFO,PARENT=PARTROOT,BYTES=85, FREQ=1
FIELD NAME=(STANKEY,SEQ),TYPE=C,BYTES=2,START=1
SEGMENT NAME=STOKSTAT, PARENT=PARTROOT, BYTES=160, FREQ=2
FIELD NAME=(STOCKKEY,SEQ),TYPE=C,BYTES=16,START=1
SEGMENT NAME=CYCCOUNT, PARENT=STOKSTAT, BYTES=25, FREQ=1
FIELD NAME=(CYCLKEY,SEQ),TYPE=C,BYTES=2,START=1
SEGMENT NAME=BACKORDR, PARENT=STOKSTAT, BYTES=75, FREQ=0
FIELD NAME=(BACKKEY,SEQ),TYPE=C,BYTES=10,START=1
DBDGEN
FINISH
END
  
```

IMS/DB definitions

Program Specification Block (PSB)

```

DBPCB01 PCB TYPE=DB,DBDNAME=DI21PART,PROCOPT=GOT,
KEYLEN=43
SENSESEG NAME=PARTROOT
SENSESEG NAME=STANINFO,PARENT=PARTROOT
SENSESEG NAME=STOKSTAT,PARENT=PARTROOT
SENSESEG NAME=CYCCOUNT,PARENT=STOKSTAT
SENSESEG NAME=BACKORDR,PARENT=STOKSTAT
PSBGEN LANG=COBOL,PSBNAME=DFSSAM07
END
  
```




Flattening IMS Database Records

- Concatenated Keys
 - Concatenated key fields not stored with segment data
 - Key fields needed for each row to maintain referential integrity

IMS Segment – Part Master



IMS Segment – Standard Data



- OCCURS clauses
 - Multiple instances of a field in a single instance of a segment
 - Multiple 'rows' should be generated

```

STANINFO (Standard data)
01 STAN-INFO.
   02 SI-PROC-CODE    PIC XX.  <-Key
   02 SI-INV-CODE     PIC X.
   02 SI-COST-CTRS   PIC X(30).
   02 SI-COST-CTRS-D REDEFINES SI-COST-CTRS.
       05 SI-COST-CTR-NO OCCURS 3 TIMES
          PIC X(10).
                    
```



Loading Transformed Data

- End result of transformation: Data in DB2 Load file format
- DB2 Load Utility can perform more transformations
- Load syntax needed to describe IMS data in file




IMS Data in DB2 Analytics Accelerator

Extracted IMS Data

DB2 Tables


DB2 Analytics Accelerator



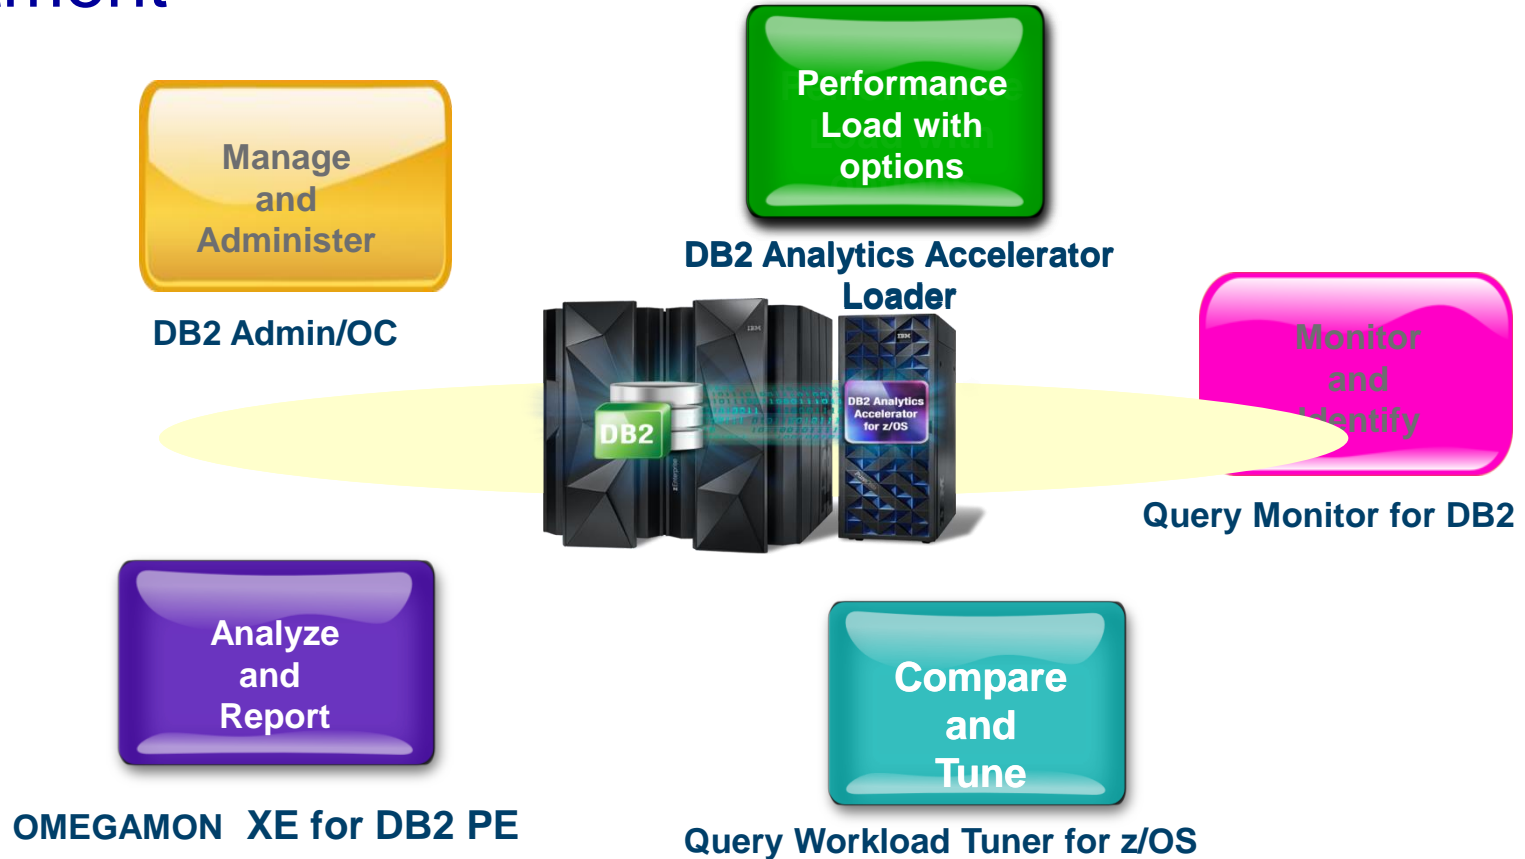
Two Step Load Process – Can be CPU Resource Intensive

How IBM Tools Can Maximize Accelerator Value

- Customers want to learn more about their investment in the Accelerator and maximize its use in their environment
 - Customer's are looking at creative ways to exploit the Accelerator....
 - IMS, VSAM, SMF Data, Non-z/OS data
 - Data Mining, IT Analytics, Reporting
- Three different areas where tools can provide value
 - Assessment
 - Do I have a workload that would benefit from the Accelerator?
 - Optimization
 - Can I optimize the workload to take advantage of the Accelerator?
 - Administration
 - Can I manage the Accelerator more effectively?



IBM Tools: Maximizing your Analytics Accelerator Investment



IBM DB2 Analytics Accelerator Loader: What is External (Dual) Load

- Accelerator Loader can load data from a file in one of two methods:
 1. Dual External Load
 - Loads data into both DB2 and the Accelerator in parallel
 2. Accelerator Only
 - Accelerator Loader loads directly into Accelerator (no load in DB2)
- User is responsible for building the load file
 - Extracted data can come from various sources
 - IMS, VSAM, Oracle.....etc
 - File must be compatible for input into the DB2 LOAD utility
 - Field specification must describe input data format. This must be compatible with the DB2 LOAD utility.

Minor changes
to
existing JCL

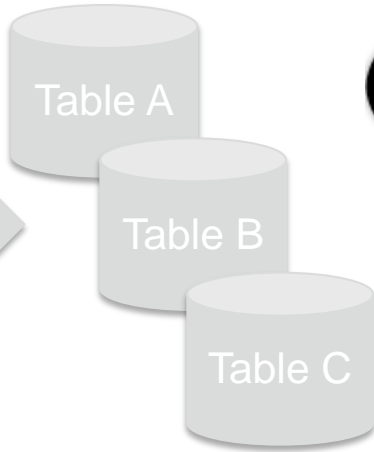



IMS Data in DB2 Analytics Accelerator

Extracted IMS Data

DB2 Tables

DB2 Analytics Accelerator



Two Step Load Process – Can be CPU Resource Intensive




DB2 Analytics Accelerator Loader: External Load (Dual Load Option)

DB2 Tables

Extracted IMS Data

DB2 Analytics Accelerator



Parallel Load into DB2 and Accelerator – Faster Load Cycles – Reduce Costs

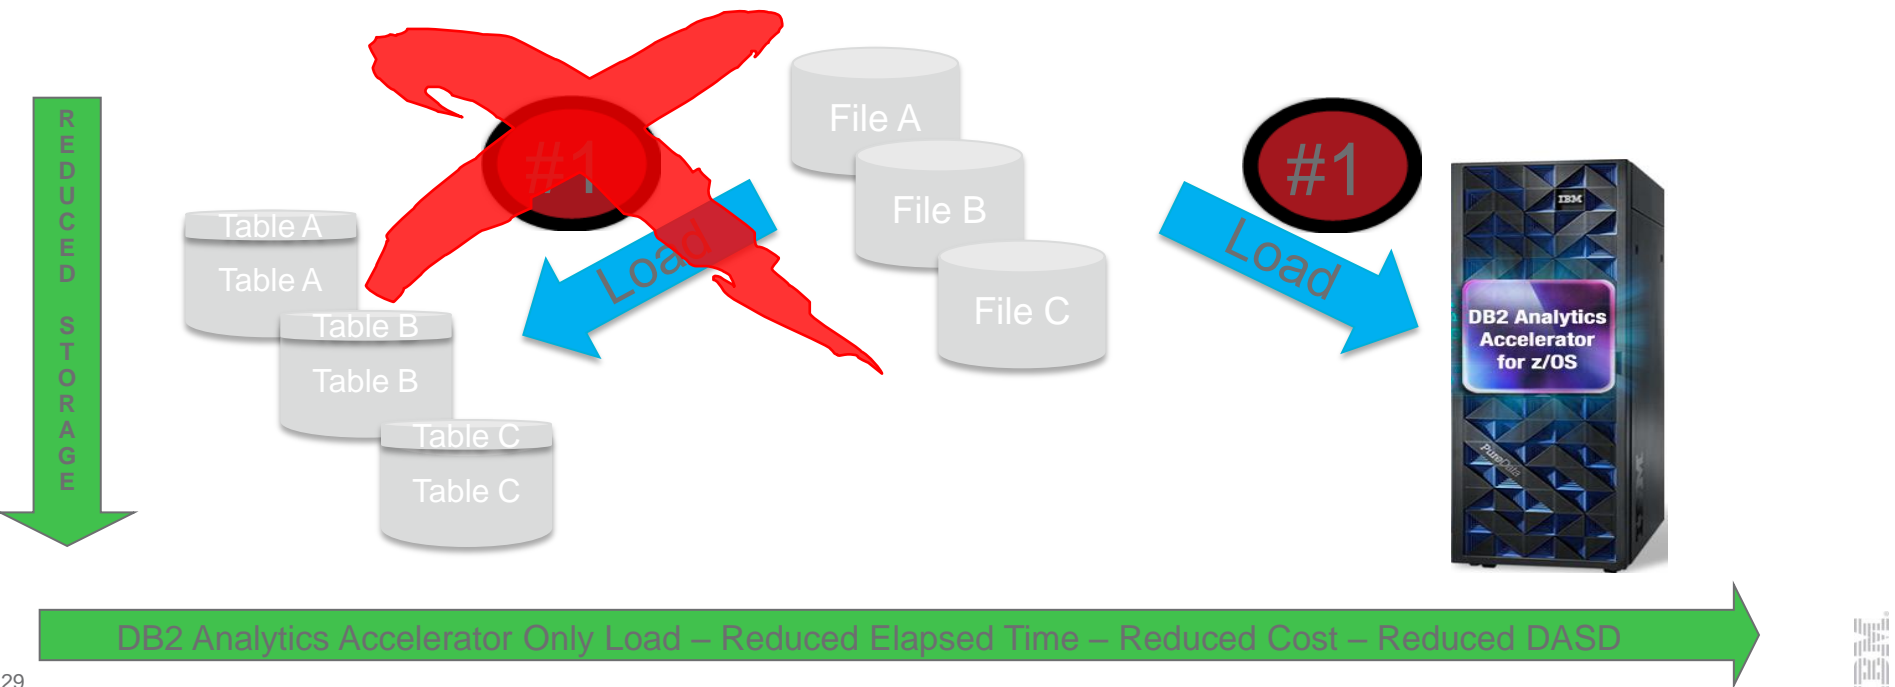


DB2 Analytics Accelerator Loader: External Load (DB2 Analytics Accelerator Only Option)

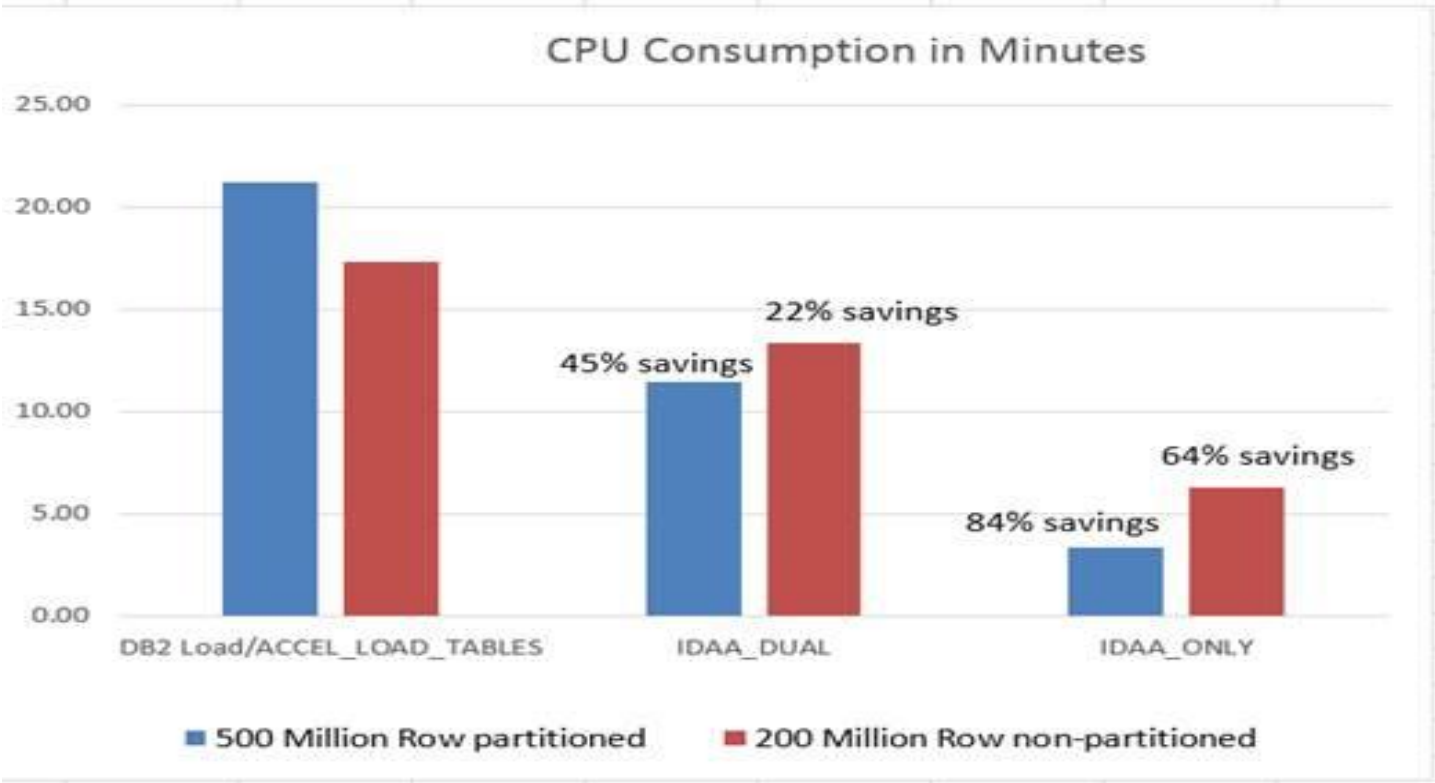
DB2 Tables

IMS Extracted External Data

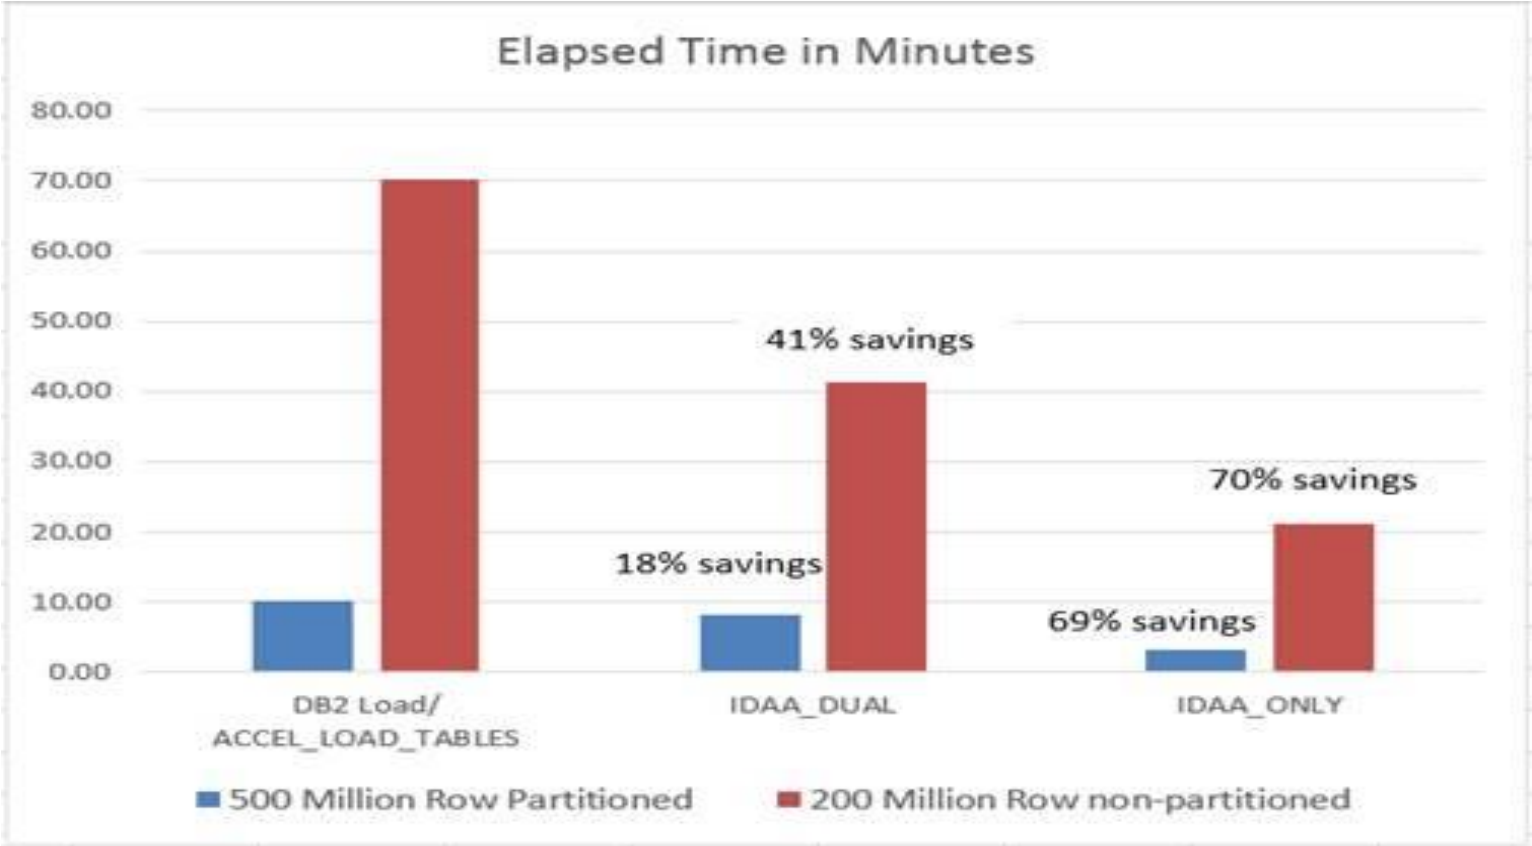
DB2 Analytics Accelerator



Performance



Performance



How to manage Big Data for IMS

IMS Automated Data Base Solutions

- Checks data base status on a regular basis
 - User specified thresholds for key indicators
- Performs reorganization only when necessary
 - Performs auxiliary functions: IC, PC, IB, etc.
- Keeps data bases performing optimally
- Saves human resources
- Saves computer resources
- Cost efficient solution for both predictable and unpredictable data base growth and activity



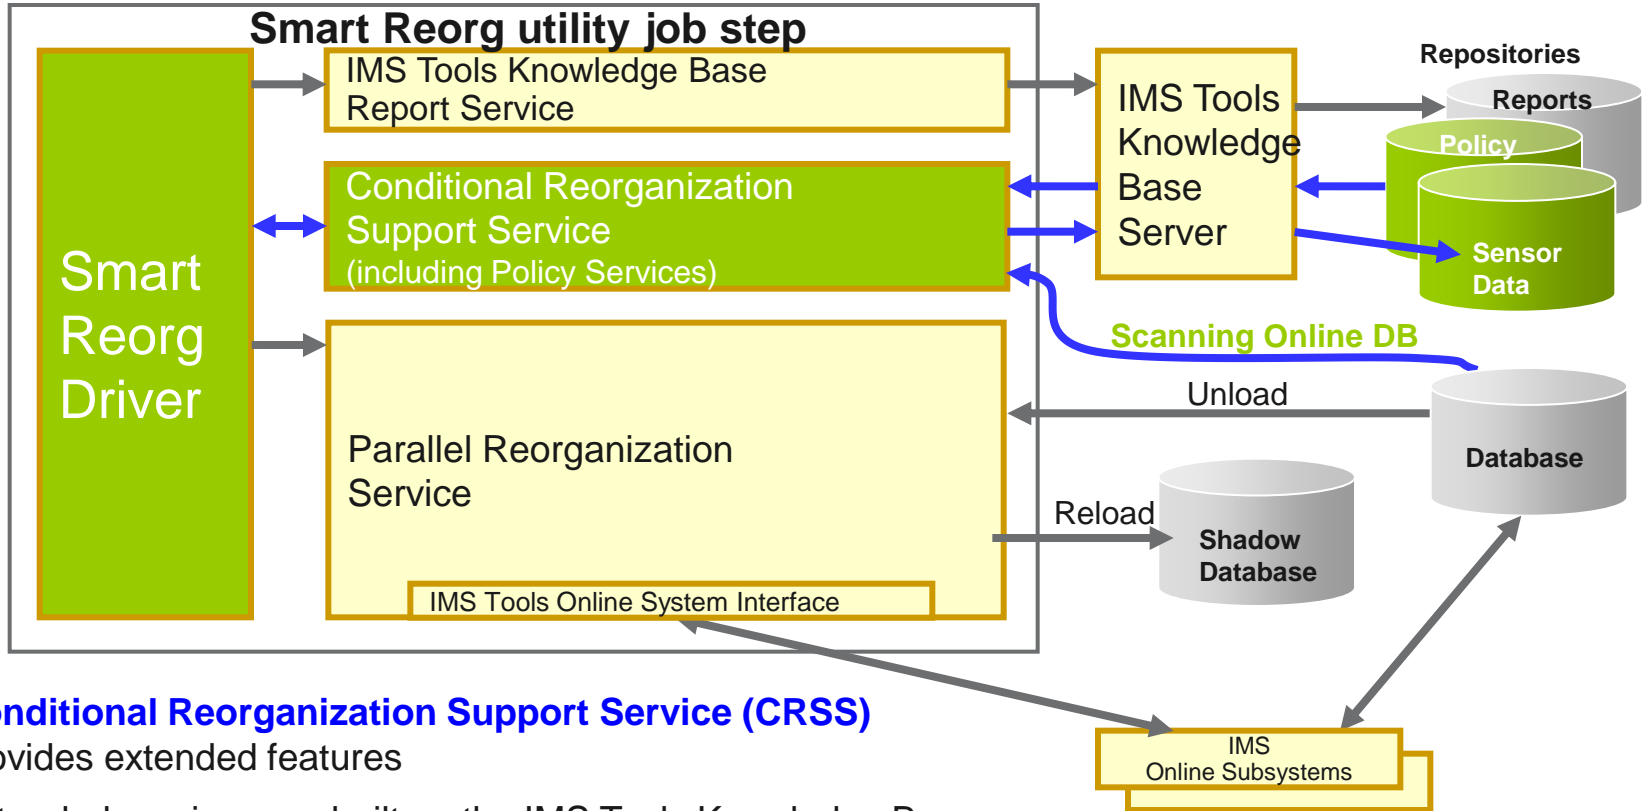
Two Approaches

- Conditional Reorganization
 - User/scheduler initiated job submission
 - Immediate Sensor Data Collection from Data Base
 - Evaluation of Sensor Data versus Policy
 - Decision on whether reorganization is needed
 - IMS Data Base Solution Pack Reorganization Expert
- Autonomic Reorganization
 - System initiated job submission
 - Periodic Sensor Data Collection from Data Base
 - Periodic Evaluation of Sensor Data versus Policy
 - Passive = Recommendations only
 - Active = Initiate and manage Autonomic Reorganization
 - IBM Base Pack IMS Autonomics Director



Getting the Most from Conditional Reorganization

Smart Reorg utility in Reorg Expert

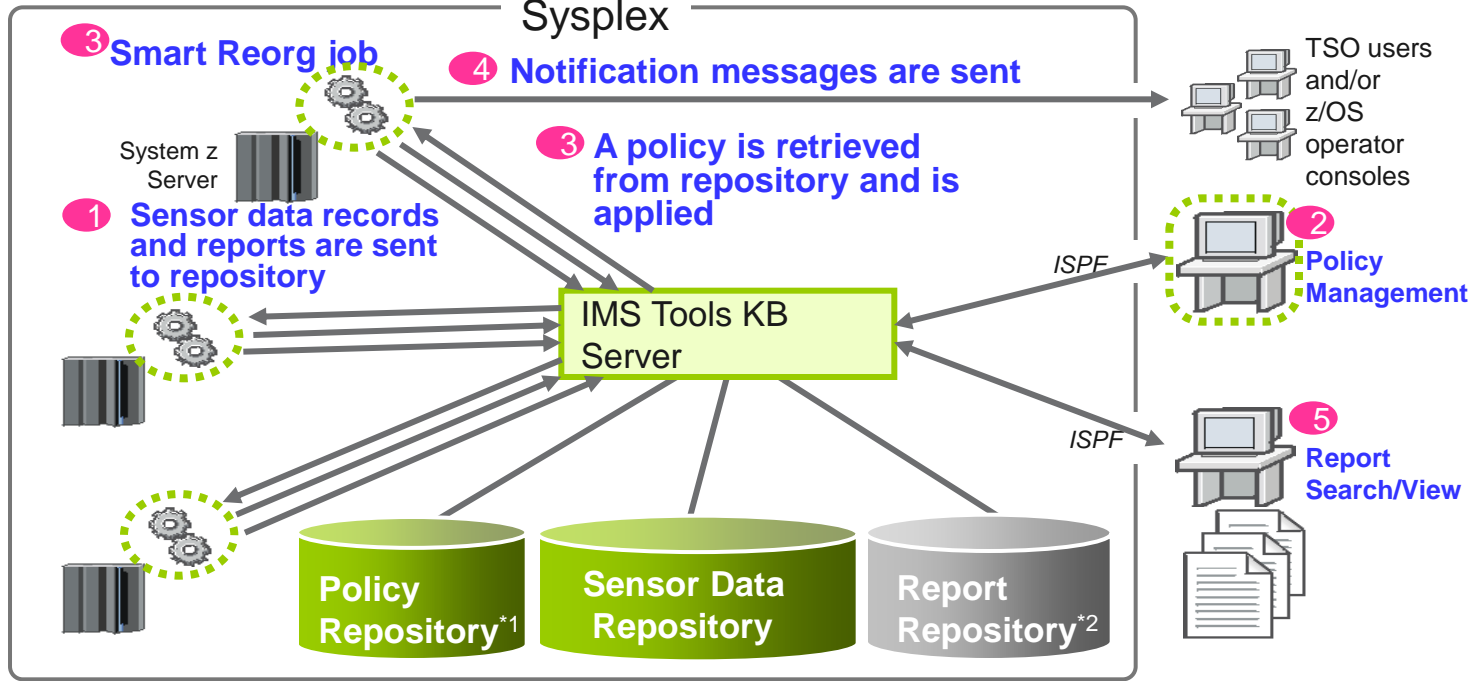


- **Conditional Reorganization Support Service (CRSS)** provides extended features
- Extended services are built on the IMS Tools Knowledge Base (IMS Tools KB) and Policy Services infrastructures



Smart Reorg utility features at a glance

1. Sensor Data Collection
2. Reorg policy Definition
3. Conditional Reorganization
4. Exception Notification and Reporting
5. Tracking exceptions and reorgs



- All information are stored in and managed by IMS Tools KB repositories
- Sysplex-wide access to these repositories is supported by IMS Tools KB Server

*1: ITKB Input Repository is used as the Policy Repository.

*2: ITKB Output Repository is used as the Report Repository.




Getting the Most from Autonomics

IMS Tools Autonomics Vision



Putting information to work


- Sensors collect resource statistics
- Policies evaluate sensor data and identify potential problems
- Automation orchestrates the collection and evaluation of sensor data
- Modernization presents an interactive modern interface for managing the system



Sensors: Collecting the Basic Information You Need

- Statistical point-in-time sensor data on your FF/FP Databases
 - Stored in IMS Tools Knowledge Base repository
 - Historically maintained per user specifications
 - Over 60 separate data elements related to space usage, optimization, and fragmentation
 - data set extents, DASD volume usage, data set free space, roots distribution, RAP usage, CI/CA splits, and IMS free space, etc
- Two methods of collection:
 - Standalone database Sensor utilities for full-function and Fast Path databases
 - Integrated with existing IMS Tools
- Integrated Tools support
 - High Performance Image Copy, High Performance Pointer Checker
 - Fast Path Analyzer, Fast Path Online Pointer Checker

List of Full Function sensor data collected



Database Record Statistics (per database or HALDB partition)

▪ Nbr. of DB records	▪ Avg. DB record length		
----------------------	-------------------------	--	--

Randomizer Statistics (per HDAM or PHDAM partition)

▪ Nbr. of total RAPs	▪ Nbr. of unused RAPs	▪ % of number of unused RAPs	▪ Nbr. of synonyms
▪ % of number of synonyms	▪ Nbr. of root not on home block	▪ % of root not on home block	▪ % of segment data in overflow
▪ Nbr. of roots in overflow	▪ % of number of roots in overflow	▪ Bytes of segments in RAA	

Volume/Extents Statistics (per data set)

▪ Allocation type (CYL, TRK, ...)	▪ Primary allocation amount	▪ Secondary allocation amount	▪ SMS-managed or not
▪ Max. nbr. of extents for the d.s.	▪ Max. nbr. of extents for the volume	▪ Nbr. of extents allocated	▪ Nbr. of volumes used
▪ Nbr. of unused volumes	▪ Nbr. of unused assigned volumes	▪ Nbr. of unused candidate volumes	
▪ Nbr. of available remaining extents determined by the max. nbr. of data set extents and the max. nbr. of extents available on volumes assigned to the data set			

Data Set Space Usage Statistics (per data set)

▪ Block/CI size	▪ Nbr. of blocks/CIs used	▪ Max. size of the data set	▪ % of data set size against the max.
▪ High-Allocated-RBA	▪ High-Used-RBA		

IMS Space Utilization Statistics (per data set)

▪ Total bytes of segment data	▪ Total bytes of free spaces	▪ Total bytes of slack bytes	▪ % of free spaces
▪ % of segment data	▪ % of unused bytes in the data set	▪ Total nbr. of segments	▪ Total nbr. of VL segments
▪ Total nbr. of VL-split segments	▪ % of nbr. of VL-split segments	▪ Total nbr. of slack bytes	▪ Avg. nbr. of slack bytes per block
▪ Total nbr. of FSEs	▪ Avg. nbr. of FSEs per block	▪ Nbr. of FSEs valid for shortest segments	▪ Nbr. of FSEs valid for longest segments
▪ Avg. nbr. of non-reusable FSEs	▪ Total nbr. of pointers	▪ Total nbr. of ptrs pointing external block	▪ % of nbr. Of ptrs pointing ext. block

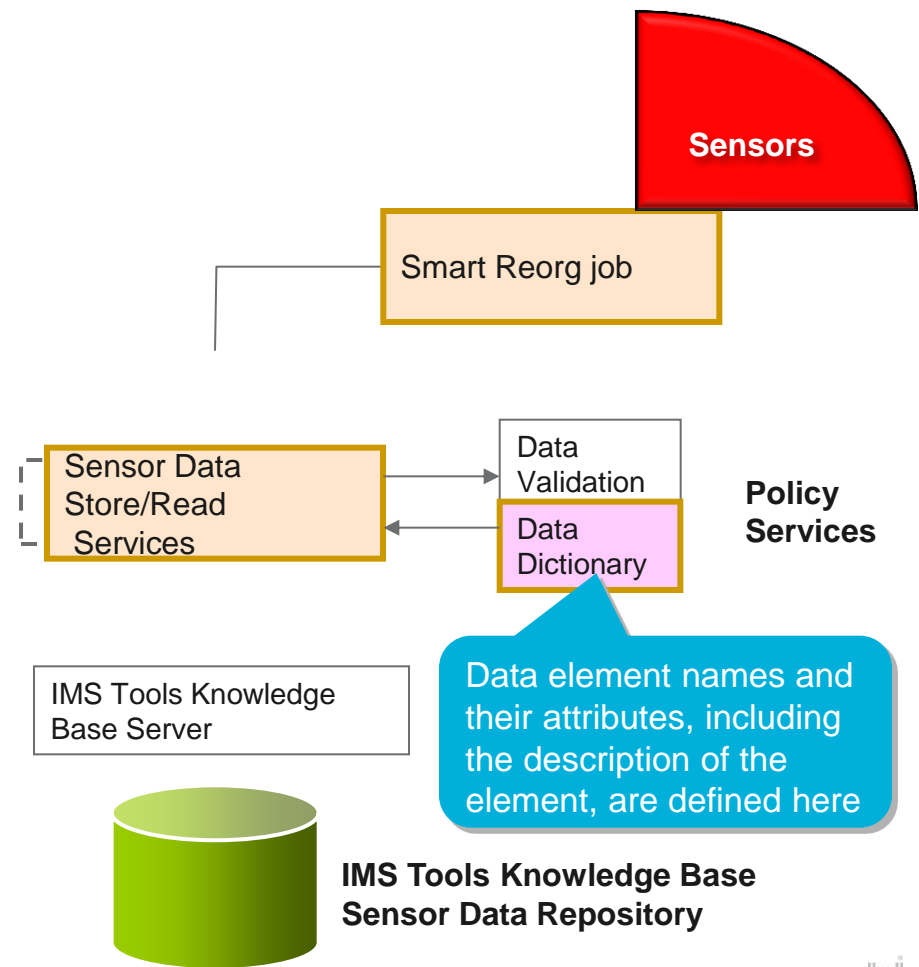
HISAM/SHISAM Statistics (for HISAM)

▪ Logical record length	▪ Total nbr. of CI splits	▪ % of nbr. of CI splits	▪ Total nbr. of CA splits
▪ % of nbr. of CA splits	▪ Total nbr. of HISAM delete bytes	▪ % of nbr. of HISAM delete bytes	



Sensor Data Repository

- The sensor data is stored in the *Sensor Data Repository* as records made up of *data elements*
- The data record is stored in a well-understood and flexible format
 - This allows its use years and multiple product releases later in time
- The data and its format is understandable between products and releases to ensure reliable functionality




Policies: Using Sensor Data to Make Decisions

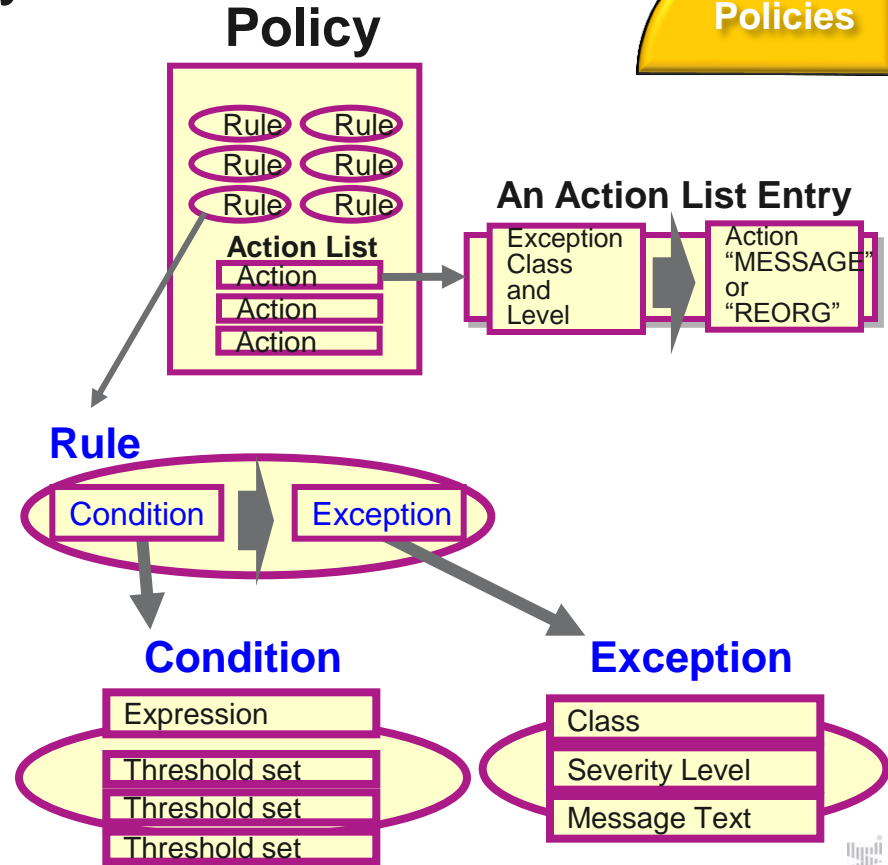
- Policy definitions are used to evaluate specific database states
 - Threshold values are compared against sensor data for a given database or group of databases
 - When thresholds are met or exceeded, exceptions occur
- Works “out of the box”
 - Ships with predefined policies and threshold values
 - Full ISPF interface provided for policy management
- Customizable to fit your shop
 - You can define your own sets of threshold values
 - Customize the messages sent when exceptions do occur
 - Specify who receives which messages and how
 - WTO, e-mail, or text



Major components of a policy



- Policy has two major components:
 - **Rules** that detect **exceptions**
 - **Exception-to-Action** mapping
- Rule Set for exception detection
 - Rule has two elements:
 - **Condition** (a threshold check formula)
 - **Exception** (a named state of a DB)
- Action List for action mapping
 - An Action List entries defines an exception-action mapping
 - The sequence of Action List entries defines whether to reorganize the subject database



Exception detection condition is defined in a rule

Sample Data Elements

DB_PCT_OF_MAX_DS_SIZE

The percentage of allocated bytes (bytes for High Allocated RBA) compared to the maximum size (4 GB or 8 GB).

DB_PCT_BYTES_FREE_SPACE

The percentage of bytes of total free spaces to the total used bytes for the data set.

A Sample Condition Description

```
Help
REORG/OPERATION                               Evaluation Formula Descrip Row 1 to 10 of 10
Command ==>
Rule name . . . . . : IBM.DBDS_GROWTH.20      Locale . . : $IVP
Value set for threshold . : MED
&1=85, &2=20,
Evaluation formula description
Both of the following thresholds have been reached or
exceeded in a database data set. This condition indicates
the possibility that high percentage of unusable free
spaces has caused the growth in data set size.
- Threshold on the percentage of data set size against
its allowable maximum size:
  &1(85)
- Threshold on the percentage of total free spaces against
the used space that is allocated for the data set:
  &2(20)
***** Bottom of data *****
```

A Sample Set of Threshold Values

```
Commands Help
DOMAIN: REORG                               View Threshold Values      Row 1 to 2 of 2
Command ==>
View threshold values and press End to exit.
Locale . . . . . : $IVP              Rule name : IBM.DBDS_GROWTH.20
Value set for threshold : MED
ID#  Value          Description
&1  85              Numeric, range: 0 to 100
                          The percentage of allocated bytes (bytes for High allocated
                          RBA) in the maximum size (4 GB or 8 GB).
&2  20              Numeric, range: 0 to 100
                          The percentage of bytes of total free spaces compared to the
                          total used bytes for the data set.
***** Bottom of data *****
```

Threshold Set

A named set of threshold values for the threshold variables that are referred to in the condition description above is called a *threshold set*.

“MED” =

&1 = 85
&2 = 20

 ← You can tweak these threshold values

Attributes of an exception

- **Exception class**
 - Represents the specific database event category being monitored
- **Exception severity level**
 - Is a category representing the severity of the detected exception
 - There are fixed three levels:
 - WARNING
 - SEVERE
 - CRITICAL
- **Exception message**
 - Is the text that can be used by the resulting policy action to describe the database event that crossed a rule threshold set
 - Users can modify the message text

Policies

An Example of Exception Class

Exception Class:

FRAGMENTED_FREE_SPACES

* Name of the rule that detects the this exception:

IBM.FRAGMENTATION.10

Exception

Class

Level

Message Text

- CRITICAL
- SEVERE
- WARNING

An Example of Exception Message

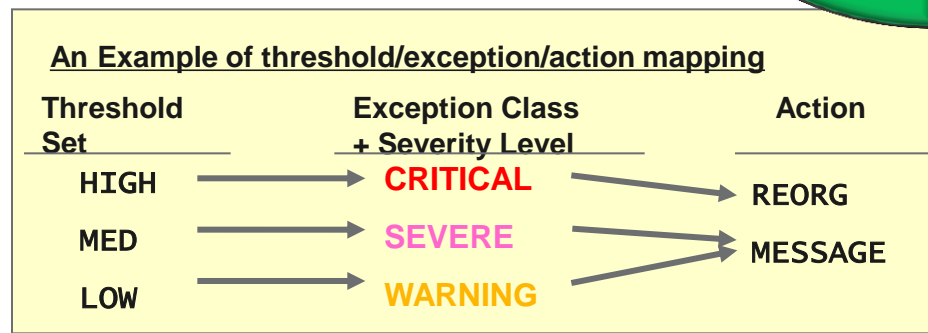
“The fragmentation of free space in %RESOURCE% has increased”

* The symbol %RESOURCE% is replaced by a DBD name or a partition name.

Exception-to-Action mapping

- An *action* is the result of a rule condition being reached or exceeded during a policy evaluation
- A rule threshold set has been mapped to a severity level for the exception class associated with the rule
- In turn, the severity level is mapped to an action

Note: In IBM-provided REORG policies, severity-level-to-action mappings are fixed for each exception class and are not customizable.



```

Commands  Help
REORG/OPERATION          Associate Actions With Rule Thres Row 1 to 3 of 3
Command ==>

Select actions. Then press Enter to be prompted to choose the associated rule
thresholds. Press End to cancel all selections.
Locale . . : BSNGLBL  Policy name . . : SYS.DBDBTYPE.HDAM
Locale . . : BSNGLBL  Rule name . . . : IBM.DBDS.GROWTH.20
Description : Simple rule on the size of data sets that have certain

A: Row Actions: S - Select Actions. (You will then be prompted to choose
                  thresholds from a list.)
                  U - Unselect.

S: Status:       S - Selected.
                  O - Pre-selected from original policy. (Update only).

A  S  Action          Level          Threshold
-  -  -
-  O  REORG           CRITICAL     HIGH
-  O  MESSAGE        SEVERE       MED
-  O  MESSAGE        WARNING      LOW

F1=Help  F3=End  F5=RFind  F7=Up    F8=Down  F10=Actions
F12=Cancel
  
```

Automation: Delivering on our Vision

- IBM Tools Autonomics Director 1.3 (Passive)
 - Automates collection and analysis of Sensor Data
 - Recommends when databases should be reorganized
 - With email or text notifications
 - Provides a scheduling feature that allows you to control how frequently sensor data is collected and how frequently policies are evaluated
 - Flexible scheduling around pre-defined PEAK times

- IBM Tools Autonomic Director 1.4 (Active)
 - Actively initiate recommended actions on user-defined database groups
 - Discovery feature for identifying related database groups
 - Ability to manage and coordinate reorganization of multiple IMS database groups as if reorganizing a single database
 - Flexible scheduling only in pre-defined Maintenance windows




IMS Management Console Goals

Modernization

- Provide a single holistic, easy-to-use interface to manage IMS systems and databases
 - Consolidate information from various tools to paint a more complete picture of IMS systems and databases
 - Leverage the latest web technologies for a richer user experience
 - Access from anywhere via the Internet using standard web browsers
 - Prepare the next generation of IMS DBAs and System Programmers through an integrated, context sensitive help system

- Begin with integration of our Autonomics Solutions
 - ...Extend integration to support all our IMS Tools



Drill down on Exceptions from an Enterprise-wide View

IBM Tools Base Administration Console for z/OS

Resources with Exceptions > Critical (3)

Summary

Resources	Type	Overall Health	Critical	Severe	Warning	Recommendations	Time Since Synchronized
HDAMVSAM (ACDEMOFF)	HDAM	■	3	1	2	1	10 minutes
HDAMVSAM (IMSPLEX)	HDAM	■	3	1	2	1	7 minutes
DBJ1AR0 (IMSPLEX)	DEDB	■	2	0	0	0	7 minutes

Oldest synchronization: 7 minutes

- Resources with Exceptions
 - Critical (3)
 - HDAMVSAM (ACDEMOFF)
 - HDAMVSAM (IMSPLEX)
 - DBJ1AR0 (IMSPLEX)
 - Severe (2)
 - HDAMVSAM (ACDEMOFF)
 - HDAMVSAM (IMSPLEX)
 - Warning (2)
 - HDAMVSAM (ACDEMOFF)
 - HDAMVSAM (IMSPLEX)

Finished retrieving child resources



Holistic View of IMS Databases

Modernization

The screenshot displays the IBM Tools Base Administration Console for the database HDAMVSAM (ACDEMOFF). The interface is divided into several sections:

- Properties:** Environment alias: STLABE2, Locale alias: ACDEMOFF, Database name: HDAMVSAM, Database type: HDAM, Segment levels: 2, Segment types: 3, External databases: 0, Logical children: 0, Access type: VSAM.
- Exceptions:** Reorganization recommended, Exceptions as of Fri Oct 19 15:55:25 PDT 2012. Critical (4): Excessive number of synonyms on RAPs, Excessive number of roots not in home blocks, Excessive number of variable-length split segments. Severe (0), Warning (3).
- Reports:** 152 reports listed, including dates like 2012-10-29 (2), 2012-10-28 (2), 2012-10-27 (2), 2012-10-26 (2), 2012-10-25 (2), 2012-10-24 (2), 2012-10-23 (2), 2012-10-22 (2), 2012-10-20 (2), 2012-10-19 (19), 2012-10-18 (2), 2012-10-16 (2).
- Space Use:** Bar chart showing Number of Segments for HDAMVSD1 and HDAMVSD2. HDAMVSD1 has approximately 12,000,000 segments, and HDAMVSD2 has approximately 10,000,000 segments.
- Optimization:** Line chart showing Number of Database Records over time. The number of records increases from approximately 500,000 on 8/11/12 to approximately 1,800,000 on 10/26/12.
- Fragmentation:** Bar chart showing Variable-Length Segment Splits for HDAMVSD1 and HDAMVSD2. HDAMVSD1 has approximately 40% splits, and HDAMVSD2 has approximately 40% splits.

...from Auto Discovery

...from Autonomics Director

...from Various HP Tools

...from Sensors



Integrated Help Throughout

Modernization

The screenshot displays the IBM Tools Base Administration Console for z/OS interface. The main content area is titled "Resources" and shows the configuration for the "HDAMVSAM (ACDEMOFF)" database. The interface is divided into several panels:

- Properties:** Environment alias: STLABE2, Locale alias: ACDEMOFF, Database name: HDAMVSAM, Database type: HDAM, Segment levels: 2, Segment types: 3, External databases: 0, Logical children: 0, Access type: VSAM.
- Exceptions:** 7 exceptions listed, including "Reorganization recommended" and "Excessive number of variable-length data sets are full and..."
- Reports:** 152 reports listed, including "2012-10-29 (2)", "2012-10-28 (2)", "2012-10-27 (2)", "2012-10-26 (2)", "2012-10-25 (2)", "2012-10-24 (2)", "2012-10-23 (2)", "2012-10-22 (2)", "2012-10-20 (2)", "2012-10-19 (19)", "2012-10-18 (2)", "2012-10-16 (2)".
- Space Use:** A bar chart titled "Highest Used and Highest Allocated RBA chart (Index)" showing "DB RBA HIGH ALLOC" and "DB RBA HIGH USED" for "HDAMVSBDBAMVS02".
- Optimization:** Number of Database Records.
- Fragmentation:** Variable Length Segment Splits.

A help panel on the right side of the screen is open, displaying the "Highest Used and Highest Allocated RBA chart (Index)" help text. The text explains that the storage space between the highest-used relative byte address (RBA) and the highest-allocated RBA is unformatted space. It also includes a note: "The space between the highest-used and the highest-allocated RBA does not represent unformatted space for a VSAM KSDS. See the Unformatted Space chart for this information." and a section titled "Data elements in this chart" which states: "This chart displays the following data elements from the sensor data service. If any of these data elements cross thresholds that are defined with Tools Base Policy Services, an exception notification might be generated. For information about resolving an exception, see the topic for this information." and "DBX_RBA_HIGH_USED The highest value of relative byte".

A callout box with red text points to the chart, stating: "Integrated help educates new and experienced DBAs on database concepts and how to interpret charts".

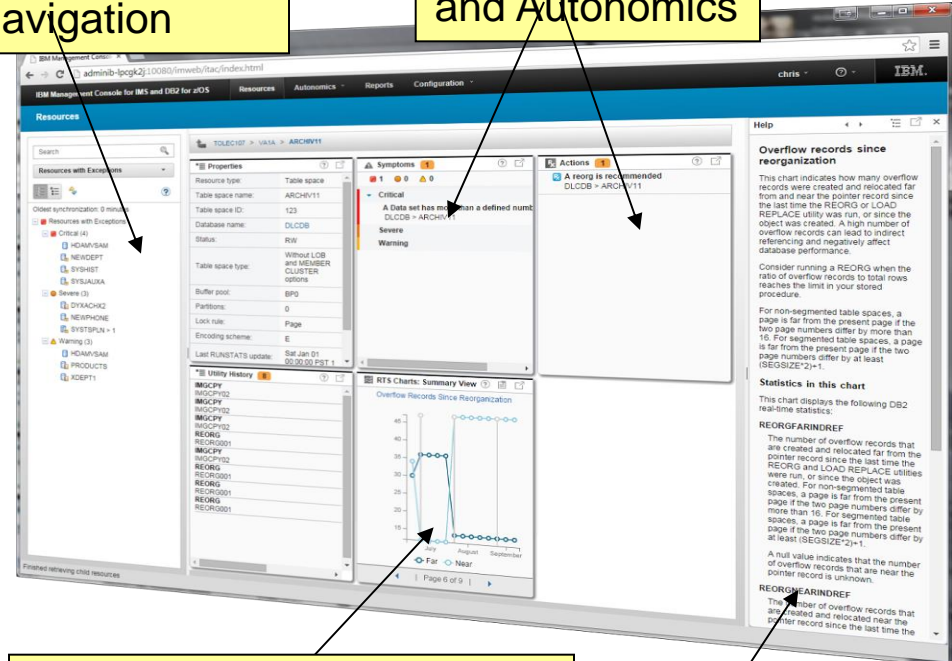
At the bottom left, it says "Finished retrieving child resources".

IBM Management Console for IMS and DB2 for z/OS

- Provides a single, holistic easy-to-use interface to manage **IMS and DB2**
 - Zero-install web-based interface
 - Consolidate information from IMS, DB2 and tools to from across the entire enterprise
 - Reduced time for problem identification and resolution through tight integration with IMS and DB2 Autonomics
 - Dramatically reduced learning curve for new users of IMS and DB2
- Now a separate product available **no-charge** (5655-TAC)
 - Extensible by growing number of products and solution packs adding additional value

Enterprise-wide Navigation

Object Health and Autonomics



Graphical Visualization of data not possible in ISPF

Integrated Help

IBM Management Console for IMS and DB2 for z/OS

Modernization

- Progressive drill down through variety of **DB2 object dashboards**
- Autonomics Director for DB2 for z/OS** (in the *no-charge* Tools Base) enables:
 - Charting of DB2 object statistics through RTS snapshots with
 - Autonomics control to define profiles and maintenance windows
 - Integrated support for the DB2 Admin Task Scheduler




IBM Management Console for IMS and DB2 for z/OS

Modernization

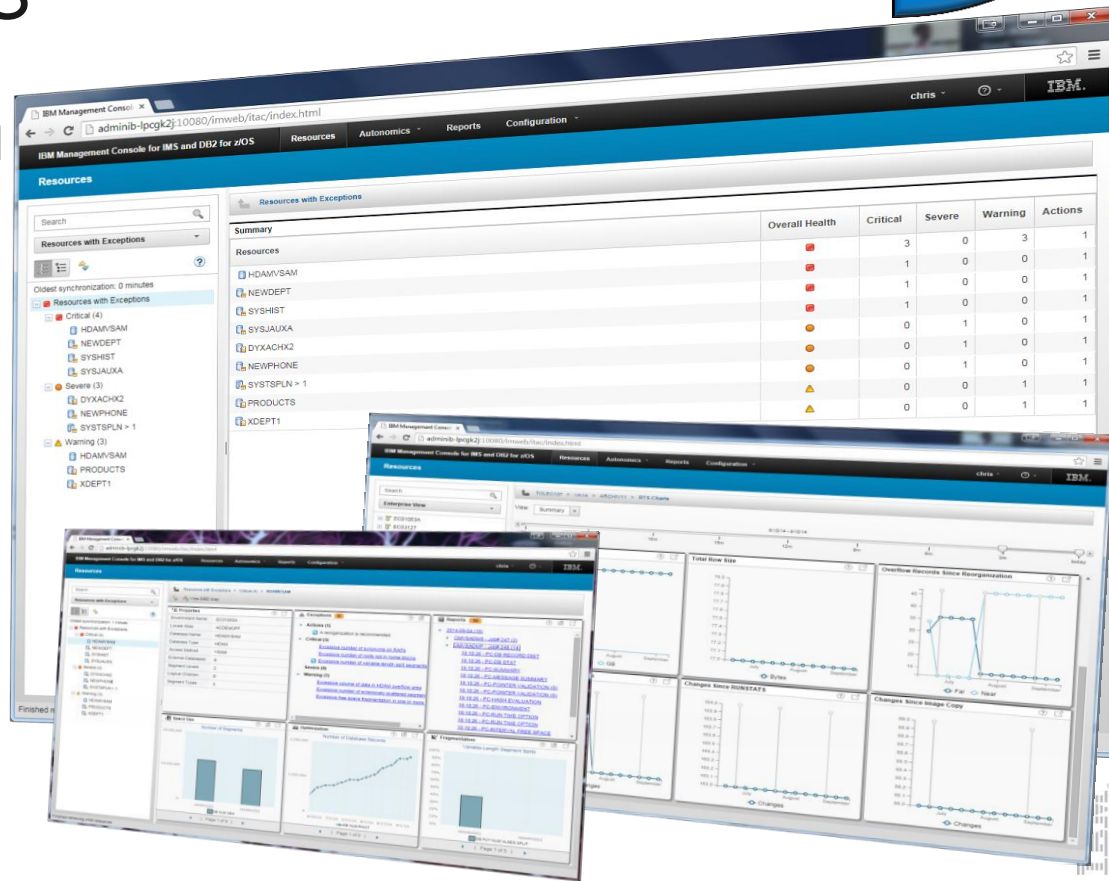
Extended with the DB2 Utility Solution Pack

- Identification and Diagnosis of symptoms and recommended actions for REORGs, ICs, Runstats
- Reporting on historical utility execution including timestamp, elapse time, system output, etc
- Graphical interfaces to define Automation Tool Object, Utility, Exception, and Job Profiles



IBM Management Console for IMS and DB2 for z/OS


- Quickly identify and drill down to databases and objects that need your attention from a single starting point
- Easily manage by exception and recommendation, taking action before problems occur
- Rapidly interpret statistical trends to verify and project
- Shorten the learning curve for new administrators
- ...all from a unified IMS and DB2 interface



Consider a combined strategy

Use for appropriate situation

- Conditional Reorganization for environmental compatibility issues
 - Mainly Job scheduler mandates
- Autonomic Director
 - In Passive Mode for health check between scheduled reorganizations to detect anomalies
 - For On Demand requests for DB status to address perceived performance issues
- Phased approach is best
 - Gain experience with a small subset of data bases
 - Consider using passive mode first



Leveraging New Technologies

Evolution of IMS Tools

- Reduce elapsed time of data management processes
 - > More application availability
- Easier, more intuitive interfaces
 - > Multi-tasked staff
- Self managing – Autonomics
 - > Increased workloads
- Exploit new technologies
 - > Growth in data, transactions
 - > Enhance business value of data



Fast-Replication Storage Processors

- Creates an instant copy of a volume or a dataset at a specific point-in-time
 - Often referred to as Point-in-Time copy, instantaneous copy or time zero (t0) copy
- Advantages
 - Can copy huge volumes of data very quickly
 - > Lab tests
 - Minimal disruption for the running applications
 - Copy process is offloaded to the storage subsystem
 - > opportunity to lower host CPU and host I/O
- Reduce backup costs and time
- Reduce recovery times



zIIP Processors

- Offloads General Mainframe Processor Work
- Originally Developed for DB2 Processing Loads
 - DB2 V8 was the first application to exploit zIIP processors
- Now widely used to offload many zOS workloads
- Must switch between TCB and SRB execution mode
- How IMS Tools can leverage
 - ‘Sorting’
 - Reorganization, Index Rebuild, Change Accumulation
 - Computing
 - Pointer Checker
- BSAM and VSAM I/O can now be offloaded



z13 Processors

- More Instructions in Hardware
- Larger Memory
 - Up to 10tb
- Simultaneous Multi-Threading (SMT)
 - Similar to IDAA processors
- Intelligent I/O system
 - 830 gb/sec bandwidth
 - IBM zHiperWrite
 - DB2 log write performance



Summary

- IBM is continuing to invest in IMS Tools
 - Analytics Accelerator – means to do real analytics without moving data off of z platform.
 - Autonomics – adding intelligence into tools to help DBAs manage growing workloads more efficiently
 - Leveraging new technologies to help DBAs perform tasks more efficiently and reducing costs
- Enablement of new technologies
 - > Enhance business value of application data
 - > Assist in workload demands
 - > Optimize cost of data management



Thank
You

