BW 2.x/3.0 Performance

Alex Peter

Product Management, SAP AG

Agenda

- Data Load Performance Concepts
- Reporting Performance Analysis Tools
- Reporting Performance Analysis Techniques
- Aggregates
- Reporting Performance Tuning Concepts
- DB Connect

Agenda

- **Data Load Performance Concepts**
- Reporting Performance Analysis Tools
- Reporting Performance Analysis Techniques
- Aggregates
- Reporting Performance Tuning Concepts

Architecture: Data Load

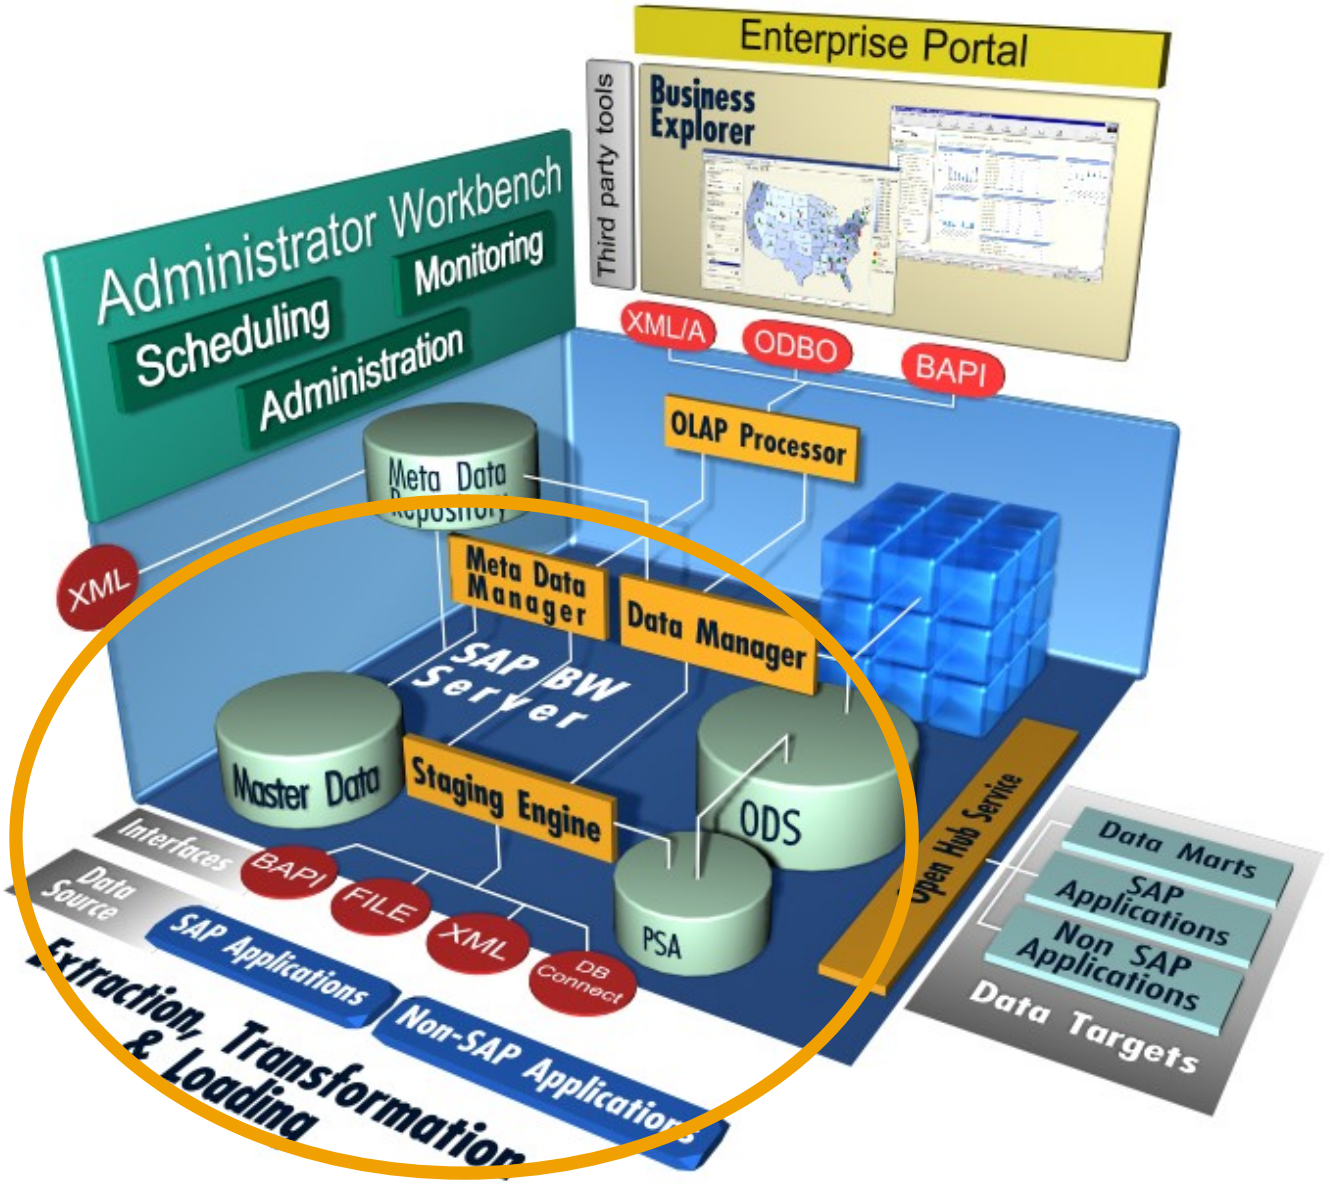

Data Load Performance: Key Concepts

The extraction and load process

Analyzing loading performance issues

- Collection and Extraction
  - SAP Content and Generic extractors
- Staging
  - PSA
- Loading
  - ODS
  - InfoCubes
  - Master Data
- Transformation
  - Transfer and update rules
- Other tuning aspects including..
  - Partitioning
  - Archiving

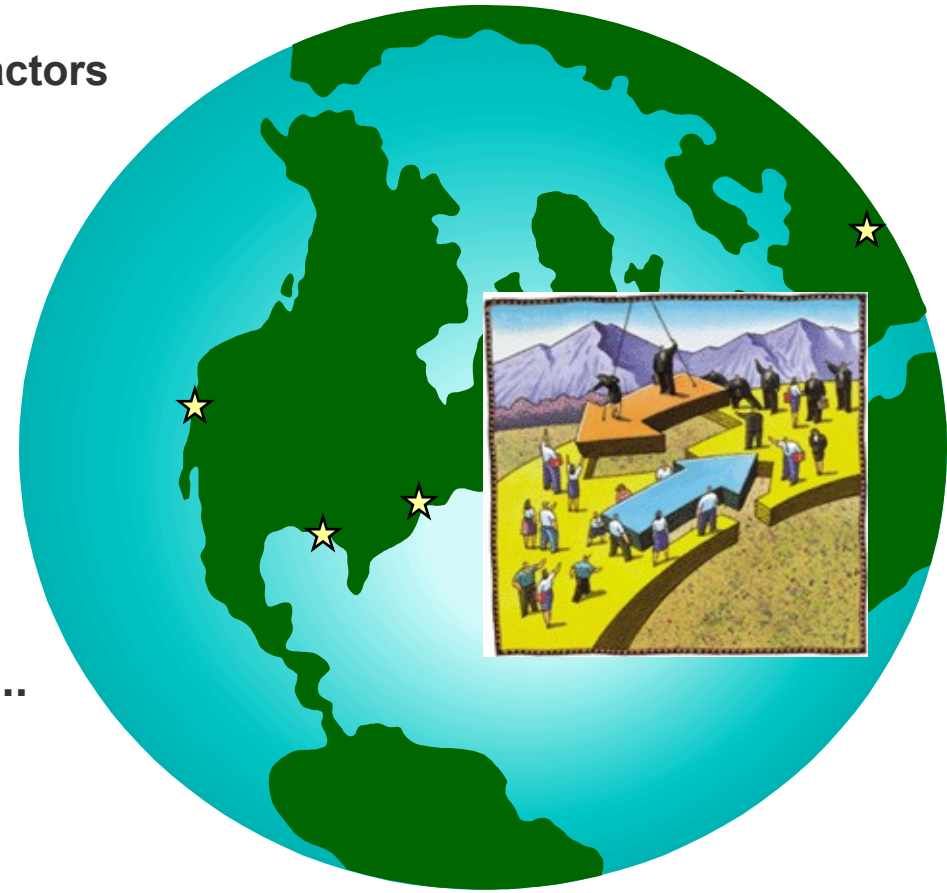

Overview: Data Load Process

Goals of performance optimization:

First tune the individual single execution and then the whole load processes.

- **Eliminating unnecessary processes**
- **Reducing data volume to be processed**
- **Deploying parallelism on all available levels**

Only Parallel processes are fully scalable!

The typical BW data load process :

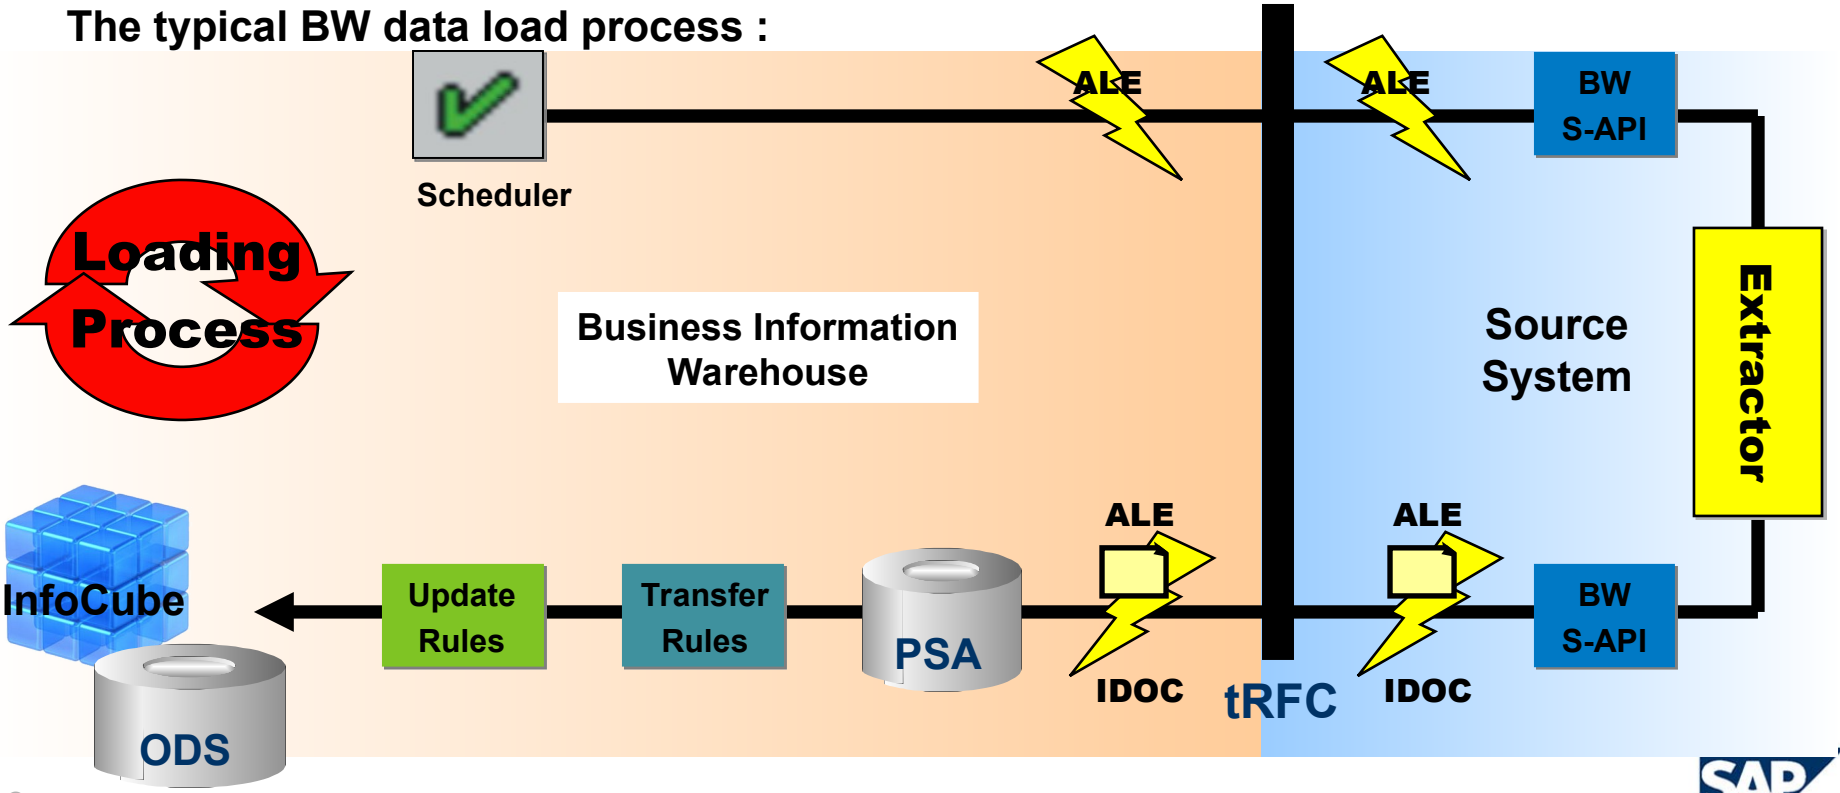

How to identify high Extraction Time ?

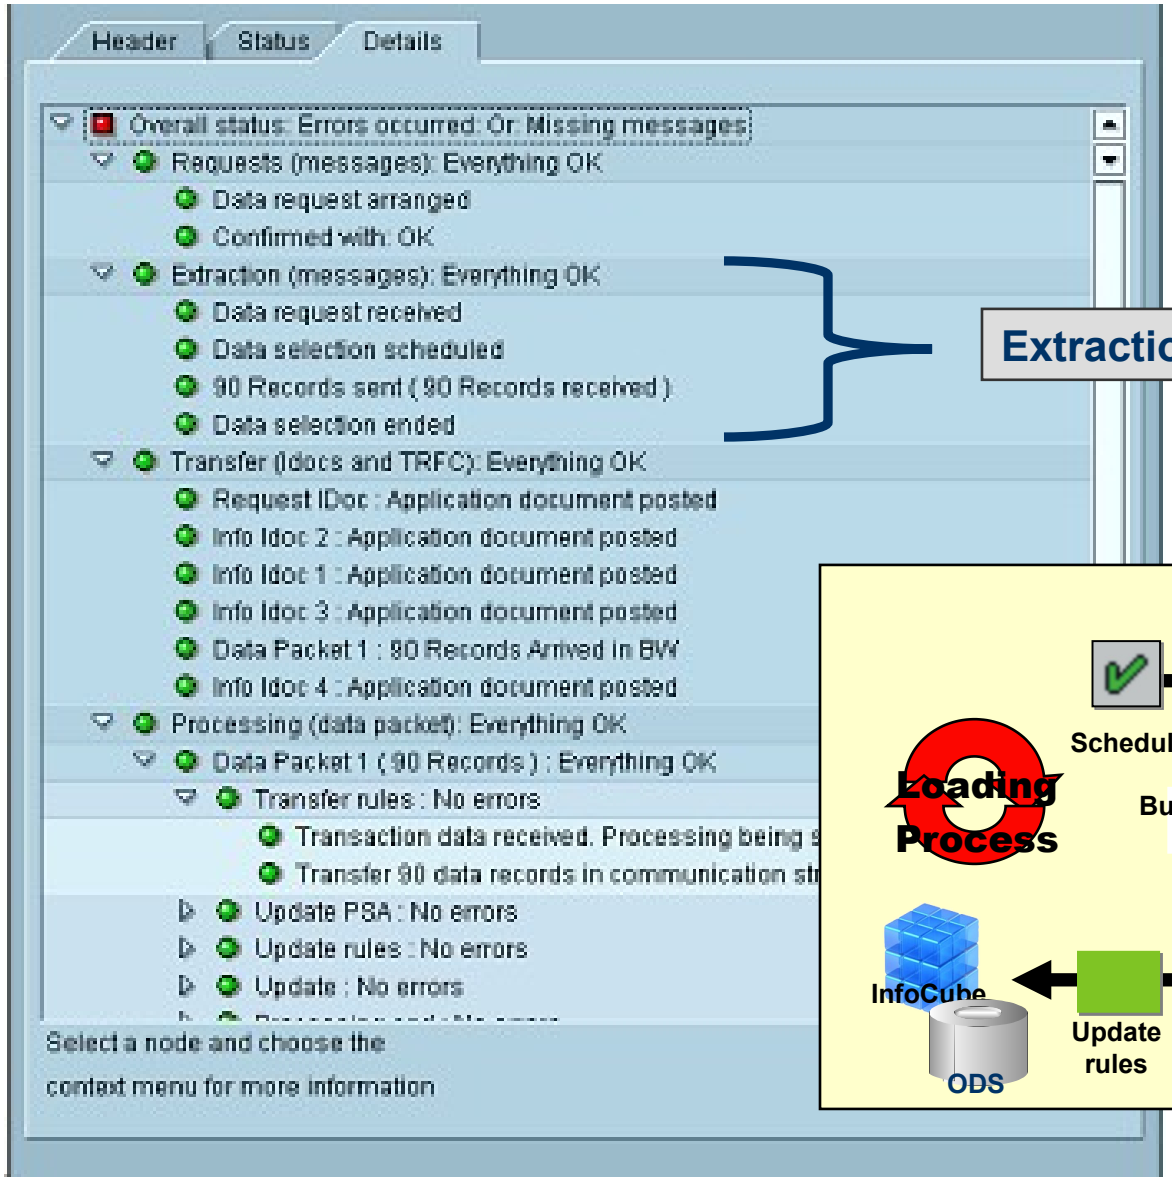

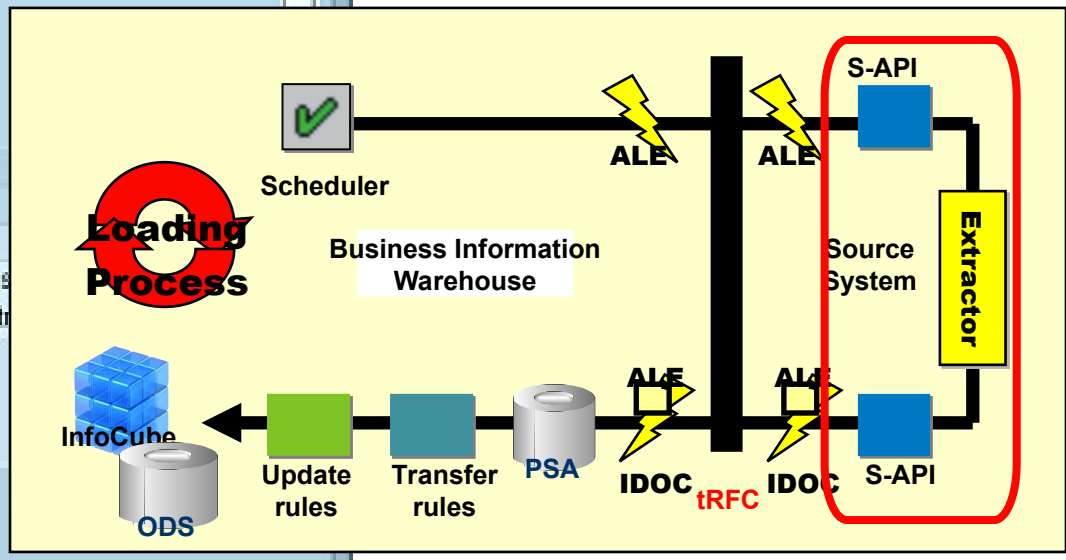

Extraction Time is too high ?

Further Analysis in case of PERFORMANCE problems extracting data...

Transaction: RSA3

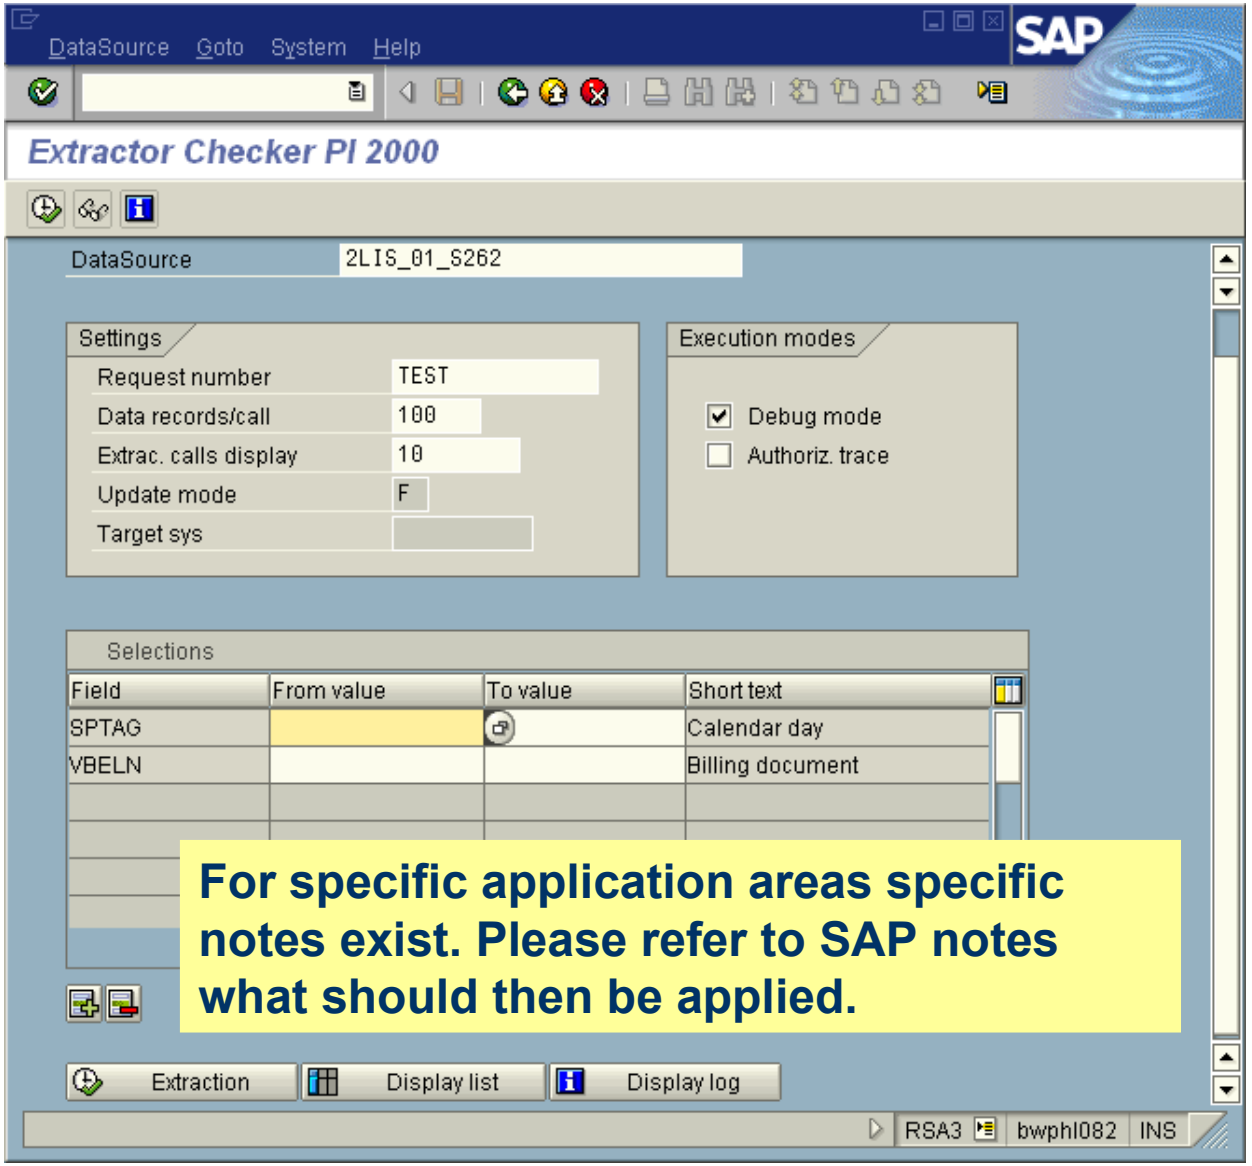

Extraction Time is too high ?

Analyze high ABAP Runtime:

Particularly Useful for User Exits

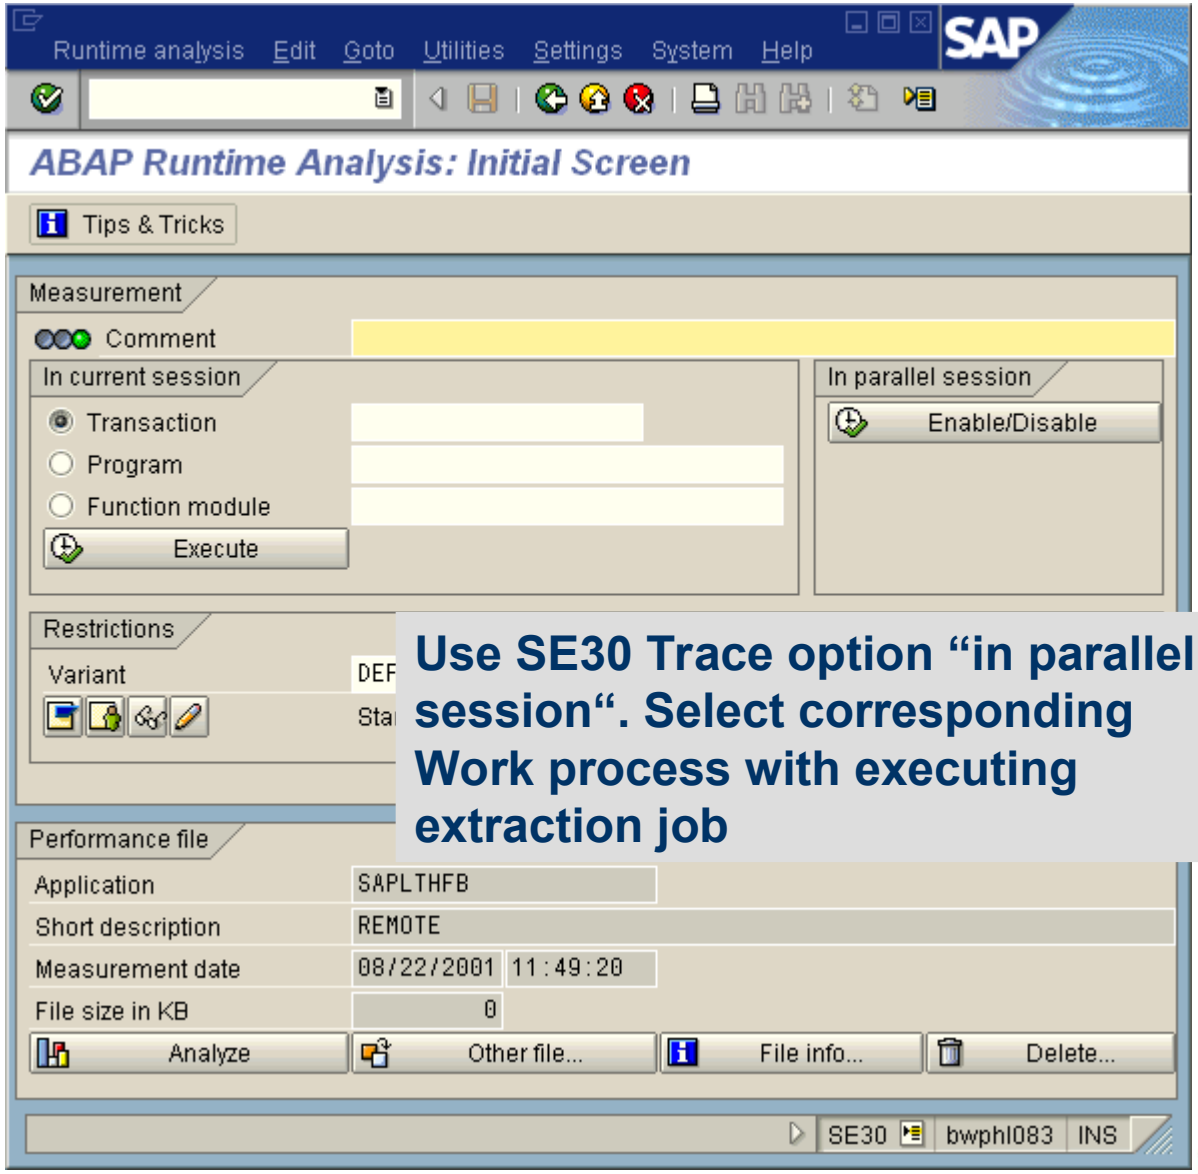

Use SE30 Trace option “in parallel session“. Select corresponding Work process with executing extraction job

Extraction Time is too high ?

Identify expensive SQL Statements

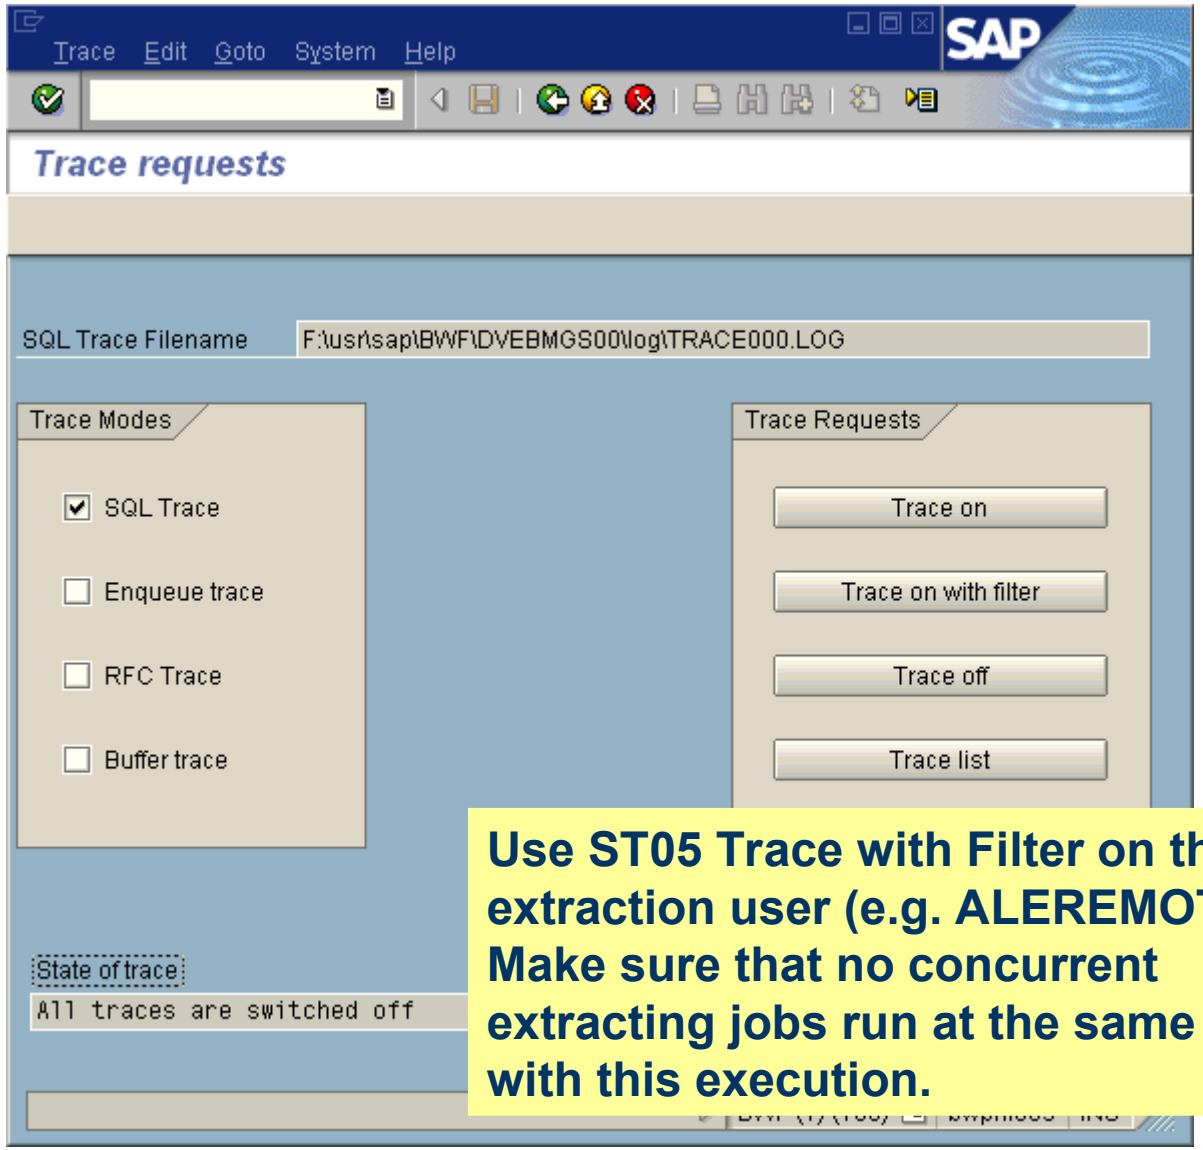

Use ST05 Trace with Filter on the extraction user (e.g. ALEREMOTE). Make sure that no concurrent extracting jobs run at the same time with this execution.

Extraction Tuning: Load Balancing

Parallel processes:

- distribute to different servers
  - avoid bottlenecks on one server

Defining a specific server:

- Configure in table ROIDOCPRMS

RFC destinations (Trx SM59):

- Example: RFC connection from BW to R/3 and R/3 to BW
- InfoPackages, event chains and Process Chains: all can be processed on specified server groups.
- XML Data loads: HTTP/HTTPS processes can be allocated to specific server groups

Expected Results:

- Avoid CPU/Memory bottlenecks on one server
- Greater Throughput: Faster time to completion per request

Extraction Tuning: Configuring Data Package Size

Size of DataPackages: Influencing Factors

- Specific to application datasource, the contents and structure of records in the extracted datasets.
- Package size: impacts frequency of COMMITs in DB .
- SAP OSS note 417307: Extractor Packet Size Collective Note for SAP Applications
- Consider both the source system and the BW system (table RSADMINC).
- Package size specified in table ROIDOCPRMS and/or InfoPackages

Scenario:

- Set up the parameters according to the recommendations; if upload performance is not improved, try to find other values that fit exactly your requirements.

Expected results:

- In a resource constrained systems, reduce the Data Package size
- In larger systems, increasing the package size to speed collection;
  - but take care not to impact communication process and unnecessarily hold work processes in SAP source system.
- Greater throughput = Faster time to completion per request

Tuning Extraction: Configuration

Further Analysis in case of Resource problems when extracting data...

Use Selection Criteria

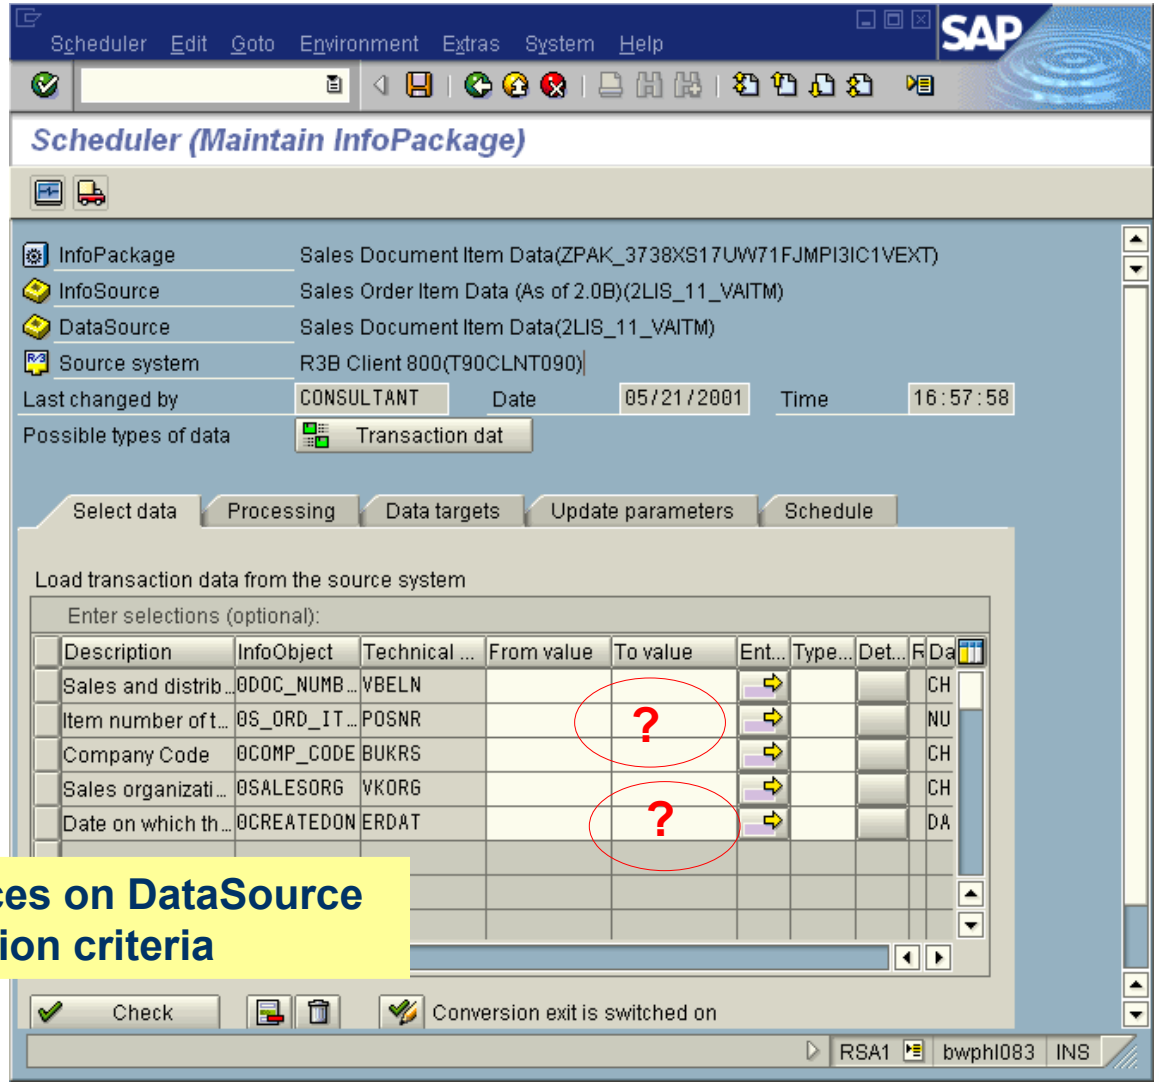

Consider building indices on DataSource Tables based on selection criteria

Extractor Tuning: SAP and Generic DataSources

SAP Content extraction

- Convert old LIS extractors to new V3 collection method
- V3 Collection jobs for different DataSources can be scheduled and executed in parallel
- Tune customer exit coding

Generic extractors:

- Collector jobs can be executed in parallel
- InfoPackages executed in parallel to extract data
  - Not possible for delta extracts from one generic DataSource
- Investigate Secondary indexes on fields used for selection
  - Too many Indexes my slow collector job
- Optimize custom ‘collector’ ABAP coding

Extraction / Load Tuning: Flat Files

- ✓ Use a **predefined record length** (ASCII file)
- ✓ File should **reside on the application server**
  i.e. not on the client PC
- ✓ Avoid large loads **across a slow networks**.
- ✓ Avoid reading load files **from tape** (copy to disk first)
- ✓ Avoid placing input load files **on high I/O disks**
  Example: same disk drives or controllers as the DB tables being loaded.

Parallel Processing Power: Extraction & Load

Automatic parallelism by the SAP system

- Data Packets - extracting from an SAP system
- Loading into PSA and data targets in parallel

You can start several processes manually (in parallel)

- InfoPackages
- InfoPackage groups (individual groups)
- Event chains
- Process chains

User-controlled parallelism

- Via InfoPackages
  - Loading from the same or different DataSource(s) with different selection criteria simultaneously
  - File Source Systems:
    - split files up for multiple InfoPackages
      Note: Individual InfoPackage groups don‘t extract in parallel for files
    - Enables parallel PSA → DataTarget load process (PSA partitions care used)

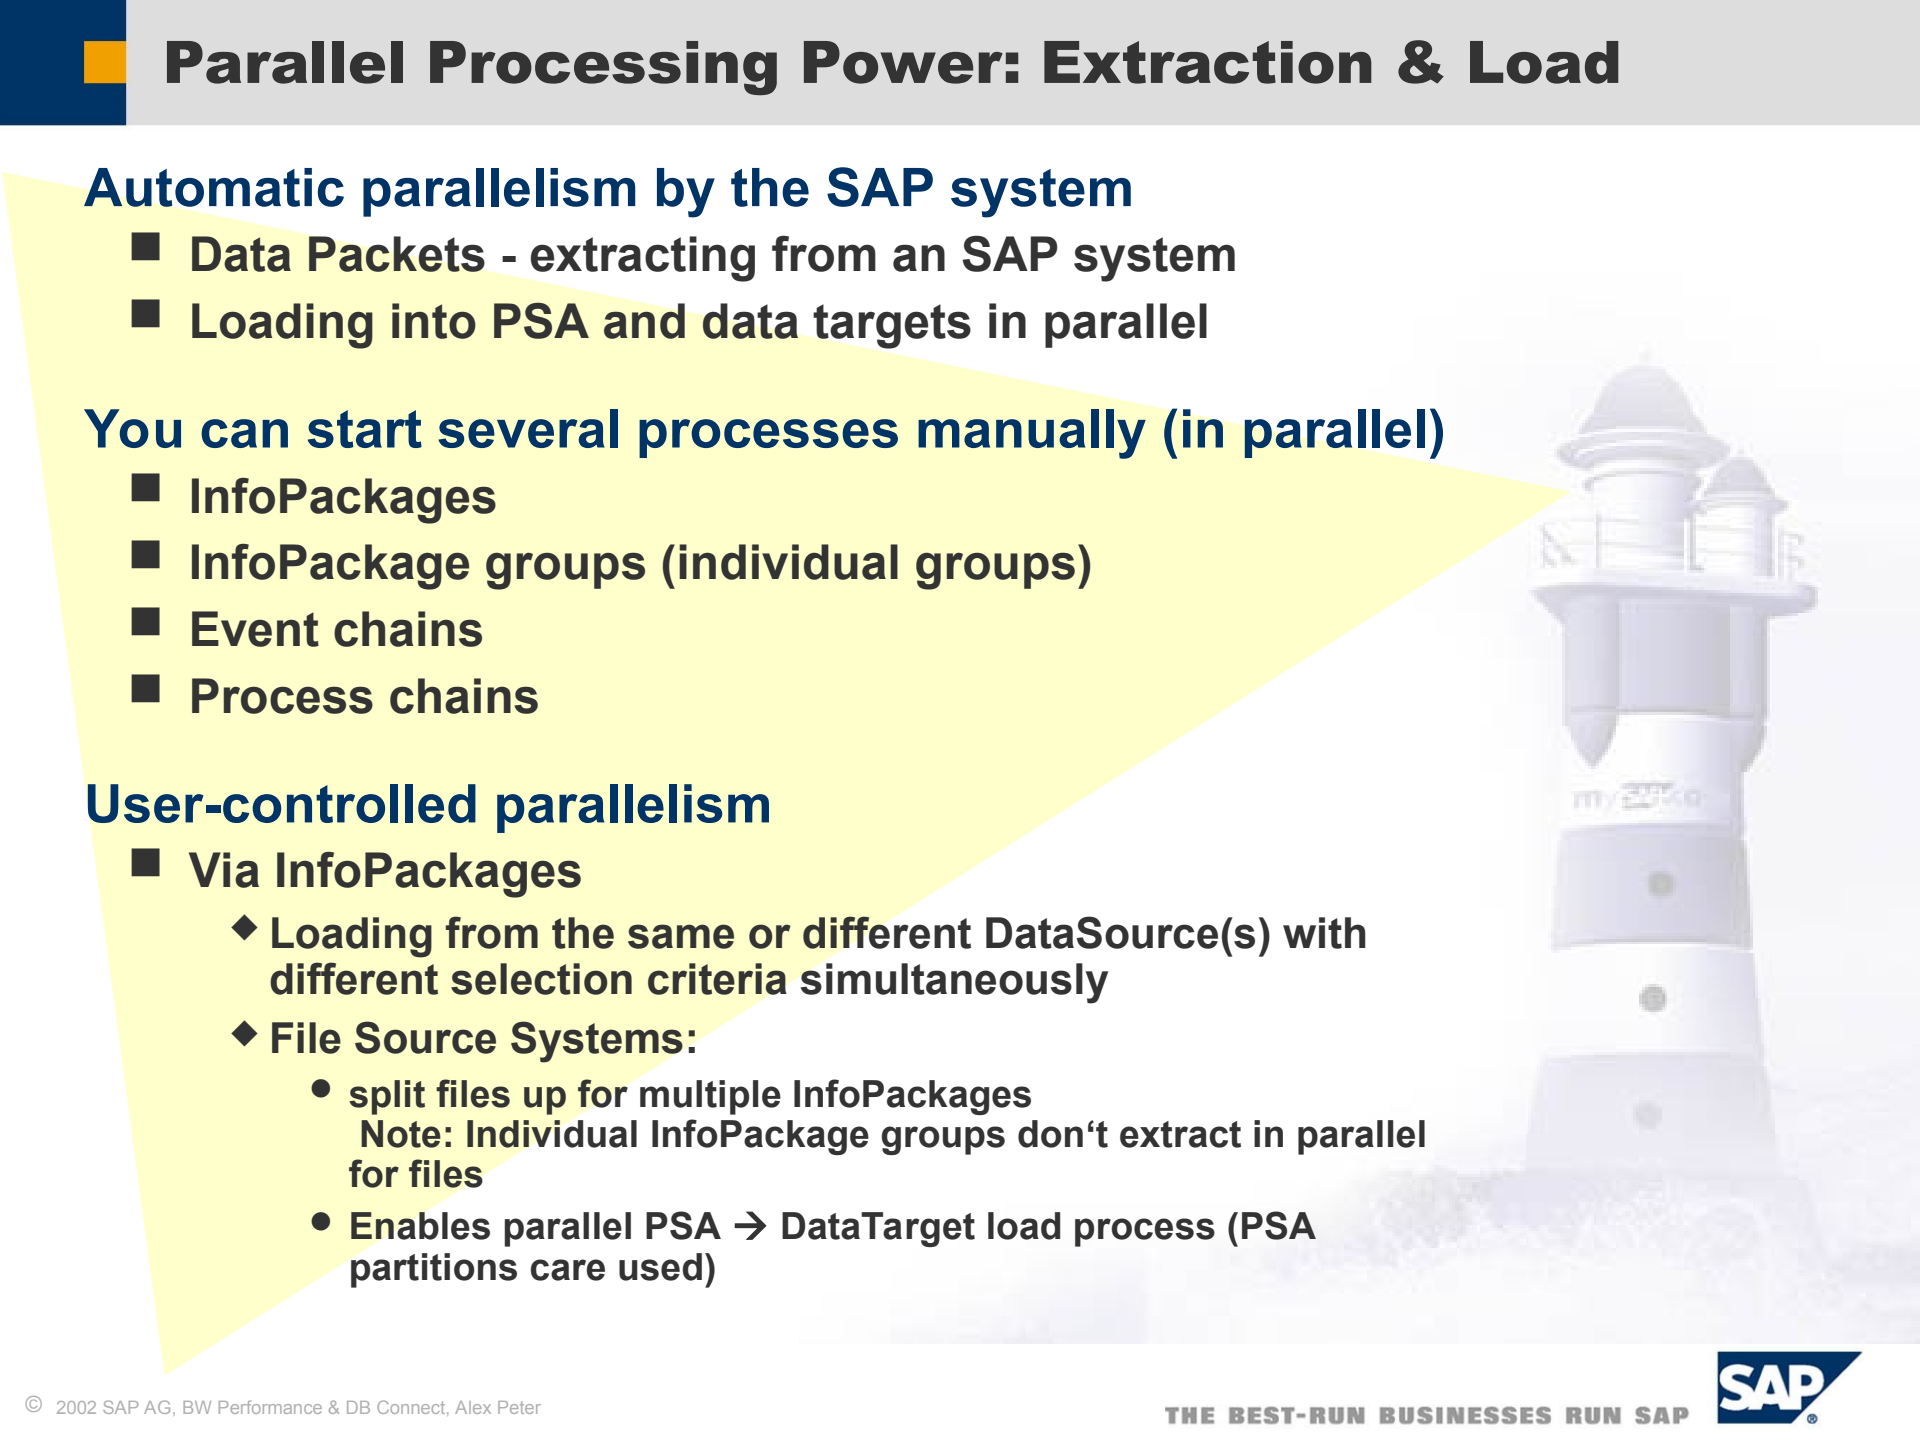

Data Transfer Times Too High ?

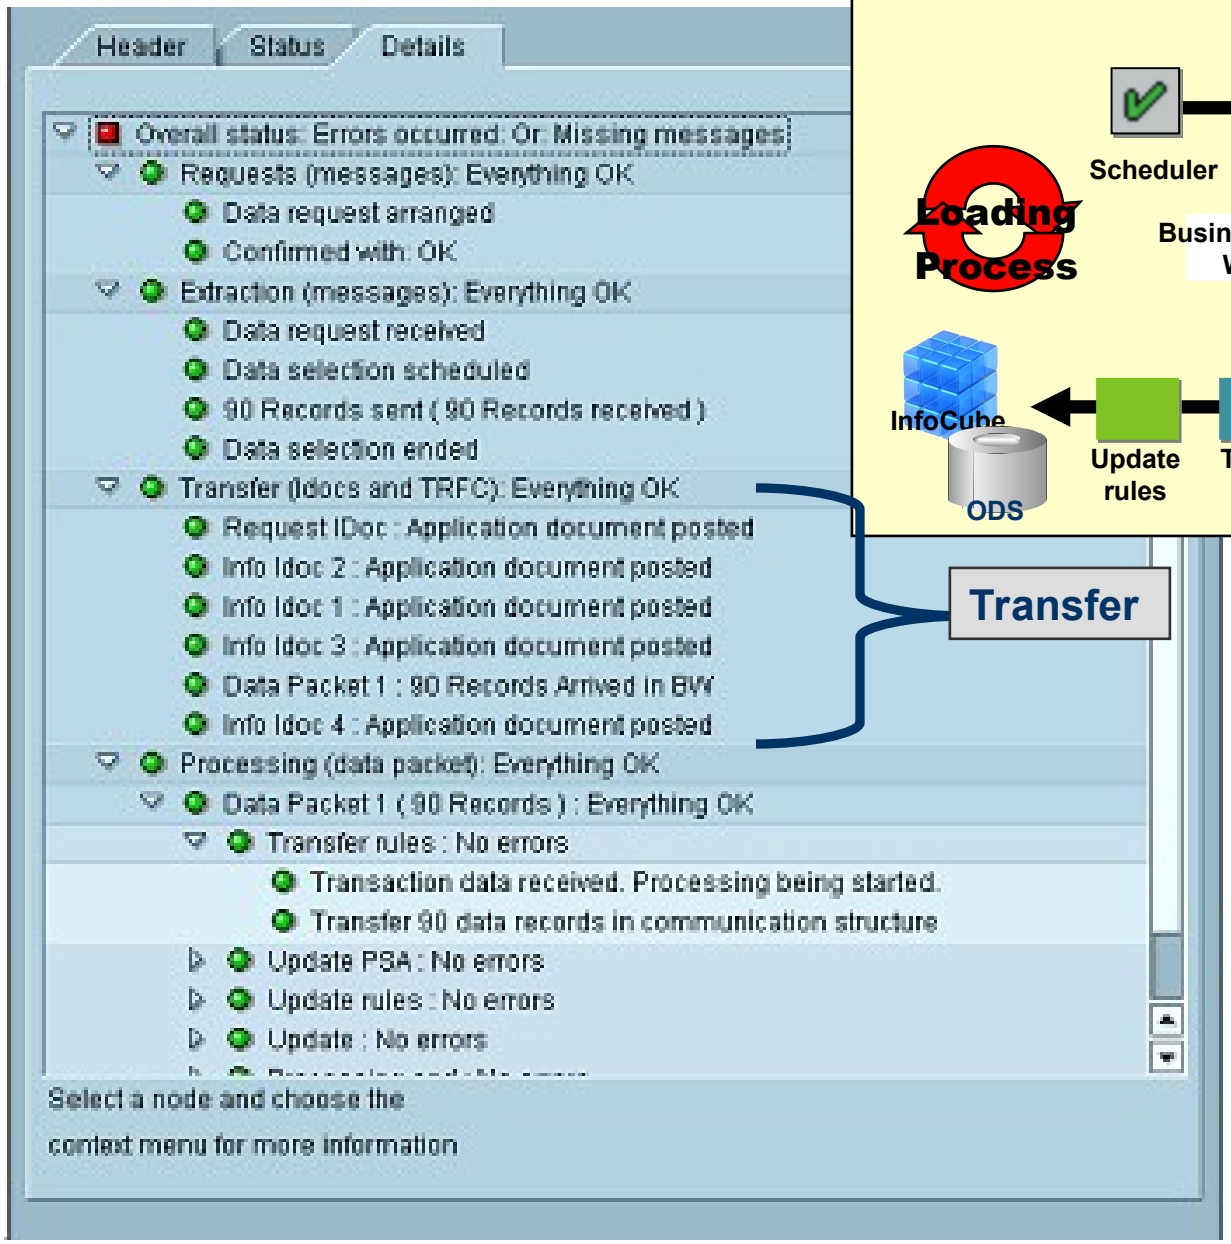

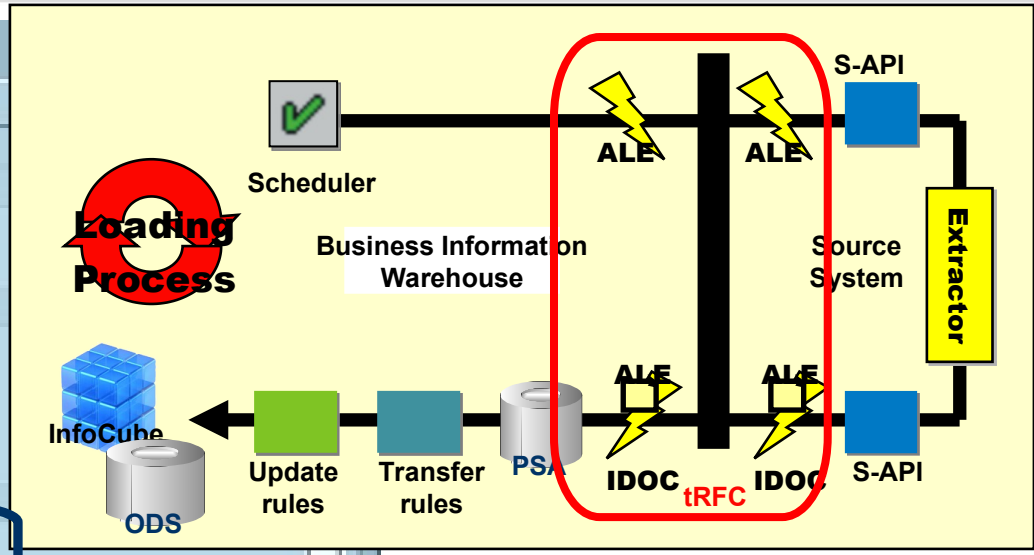

Data load Monitor : Transaction RSMO

Loading R/3 data uses too many work processes

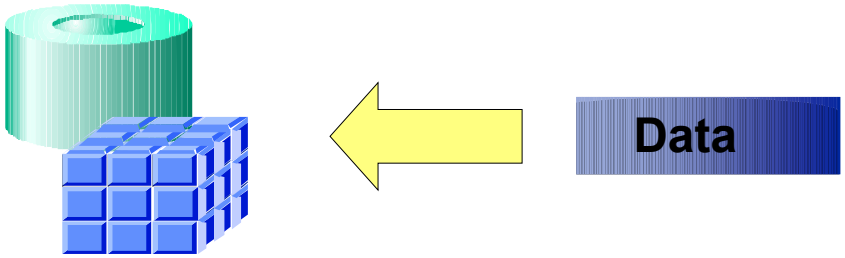

Method “Update Data Targets in Parallel“

When BW resources are constrained:

- CPU/Memory
- Network connection speed
- I/O bottleneck

Avoid large data loads with Parallel Update Method

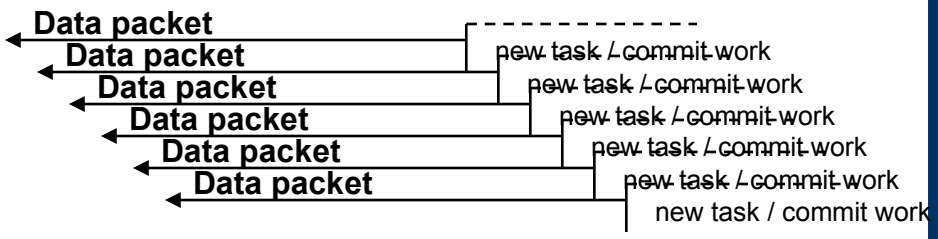

6 Dialog WPs

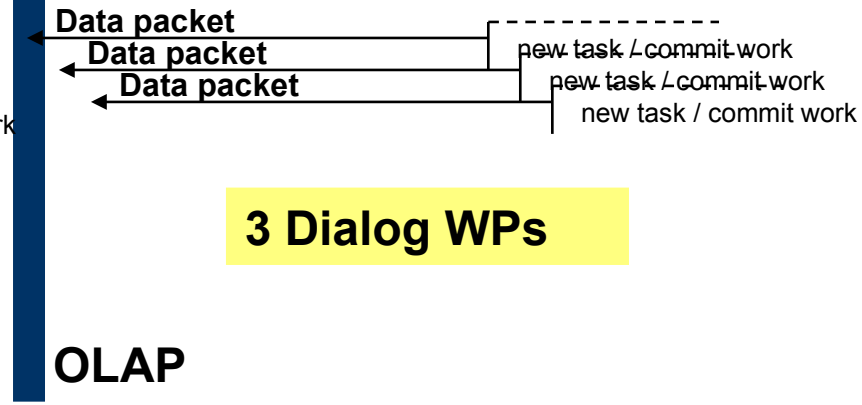

3 Dialog WPs

Analyze high PSA Upload Times

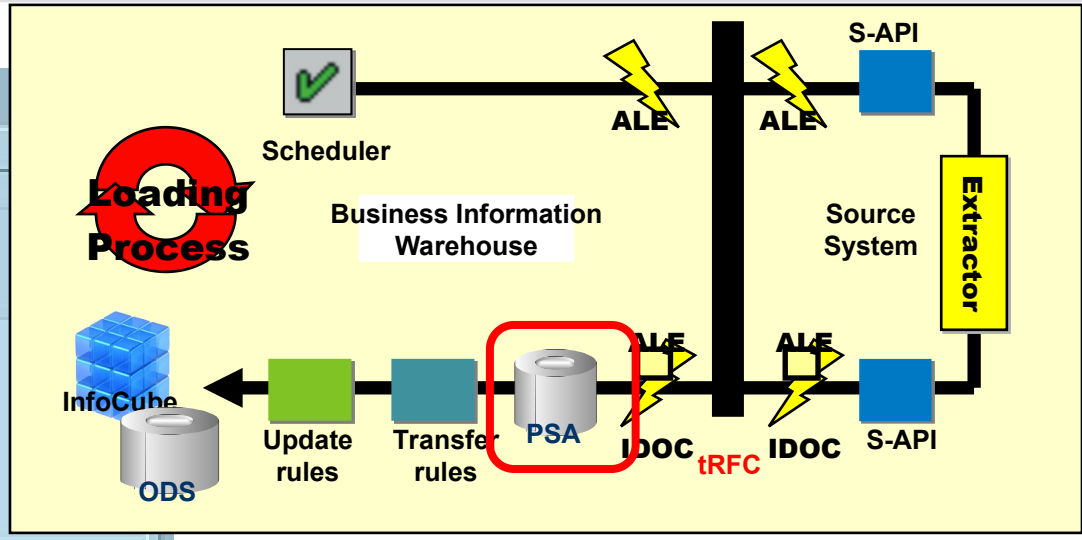

Header | Status | Details

- Overall status: Errors occurred: Or: Missing messages
  - Requests (messages): Everything OK
  - Extraction (messages): Everything OK
  - Transfer (Idocs and TRFC): Everything OK
    - Request IDoc : Application document posted
    - Info Idoc 2 : Application document posted
    - Info Idoc 1 : Application document posted
    - Info Idoc 3 : Application document posted
    - Data Packet 1 : 20 Records Arrived in BW
    - Info Idoc 4 : Application document posted
  - Processing (data packet): Everything OK
    - Data Packet 1 (20 Records) : Everything OK
      - Transfer rules : No errors
        - Transaction data received. Processing being started.
        - Transfer 20 data records in communication structure
      - Update PSA : No errors
        - Data save in PSA started
        - 20 records written in PSA 0BBP_TD_INV_1_XF
        - Reading data from PSA started
        - 20 records read from PSA 0BBP_TD_INV_1_XF
      - Update rules : No errors
      - Update : No errors
      - Processing end : No errors
    - ODS activation (change-log) : Not yet activated

Select a node and choose the context menu for more information

Load Into PSA

Data load Monitor : Transaction RSMO

PSA Partitioning

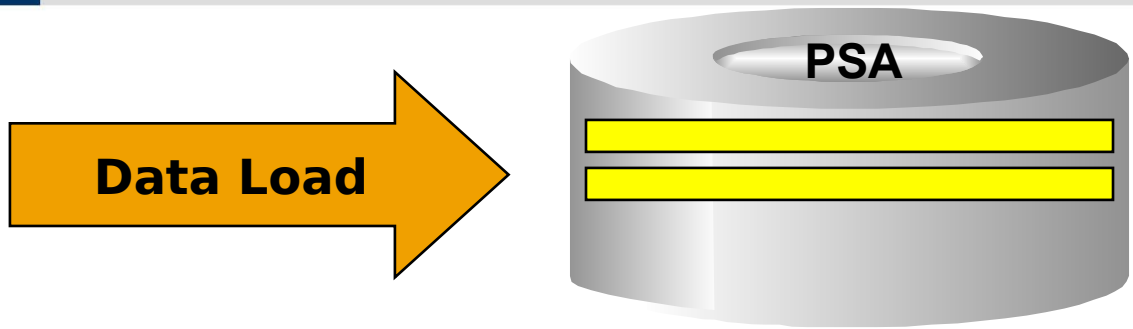

PSA Table Partitions

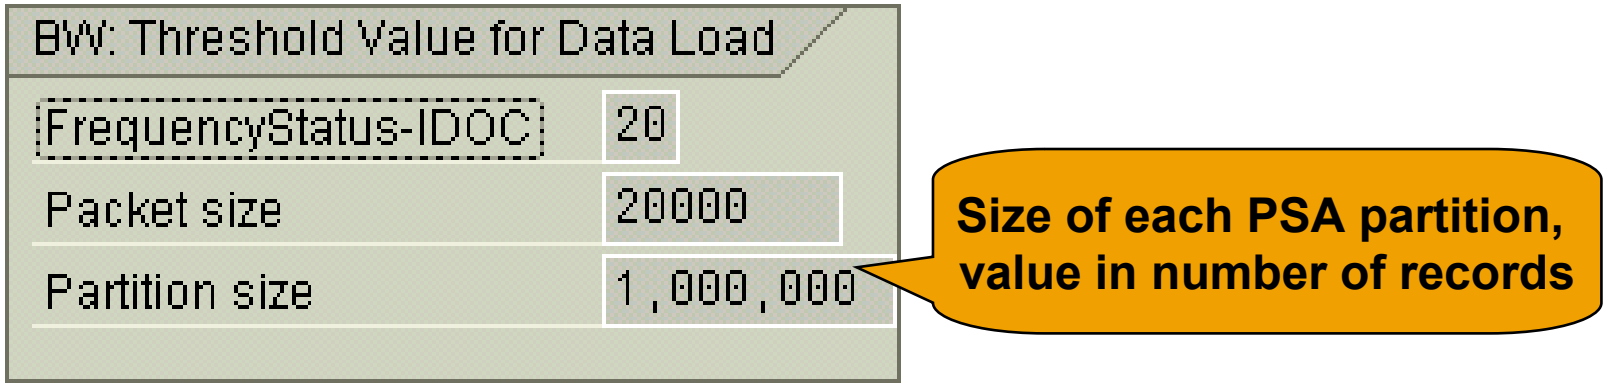

Transaction SPRO or RSCUSTV6 :

- From SPRO Business Information Warehouse > Links to Other Systems > Maintain Control Parameters for the Data Transfer
  - Note: If you start more than one load process at a time expecting to have each request in a separate partition, it probably will not work as expected; the PSA threshold is not yet reached when the second process starts writing into PSA

Potential PSA Bottleneck

Possible bottleneck (large loads or multiple parallel loads)

- I/O contention
  - High no. of writes during large data loads
  - Disk layout and striping configuration
    - What is located on the same disk or tablespace/DB Space/etc?
  - Database I/O
- Partitioning configuration
  - Partition defined too large, no parallel database sub processes used?
  - Partition too small, too many parallel database sub processes used?
- Sequential loading from PSA to DataTargets selected

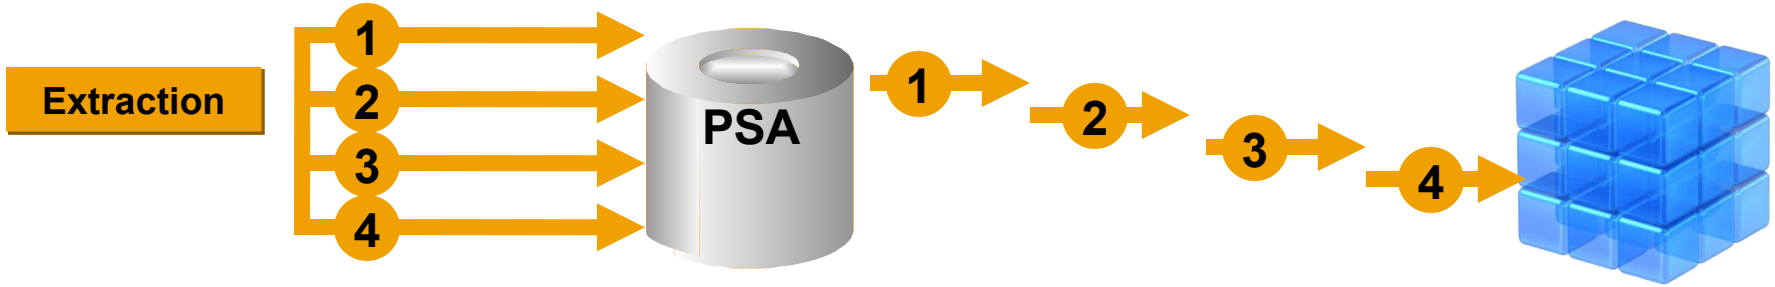

- DB Statistics
  - Usually not an issue as PSA tables are normally sequentially read

Data Processing-Transfer Rules

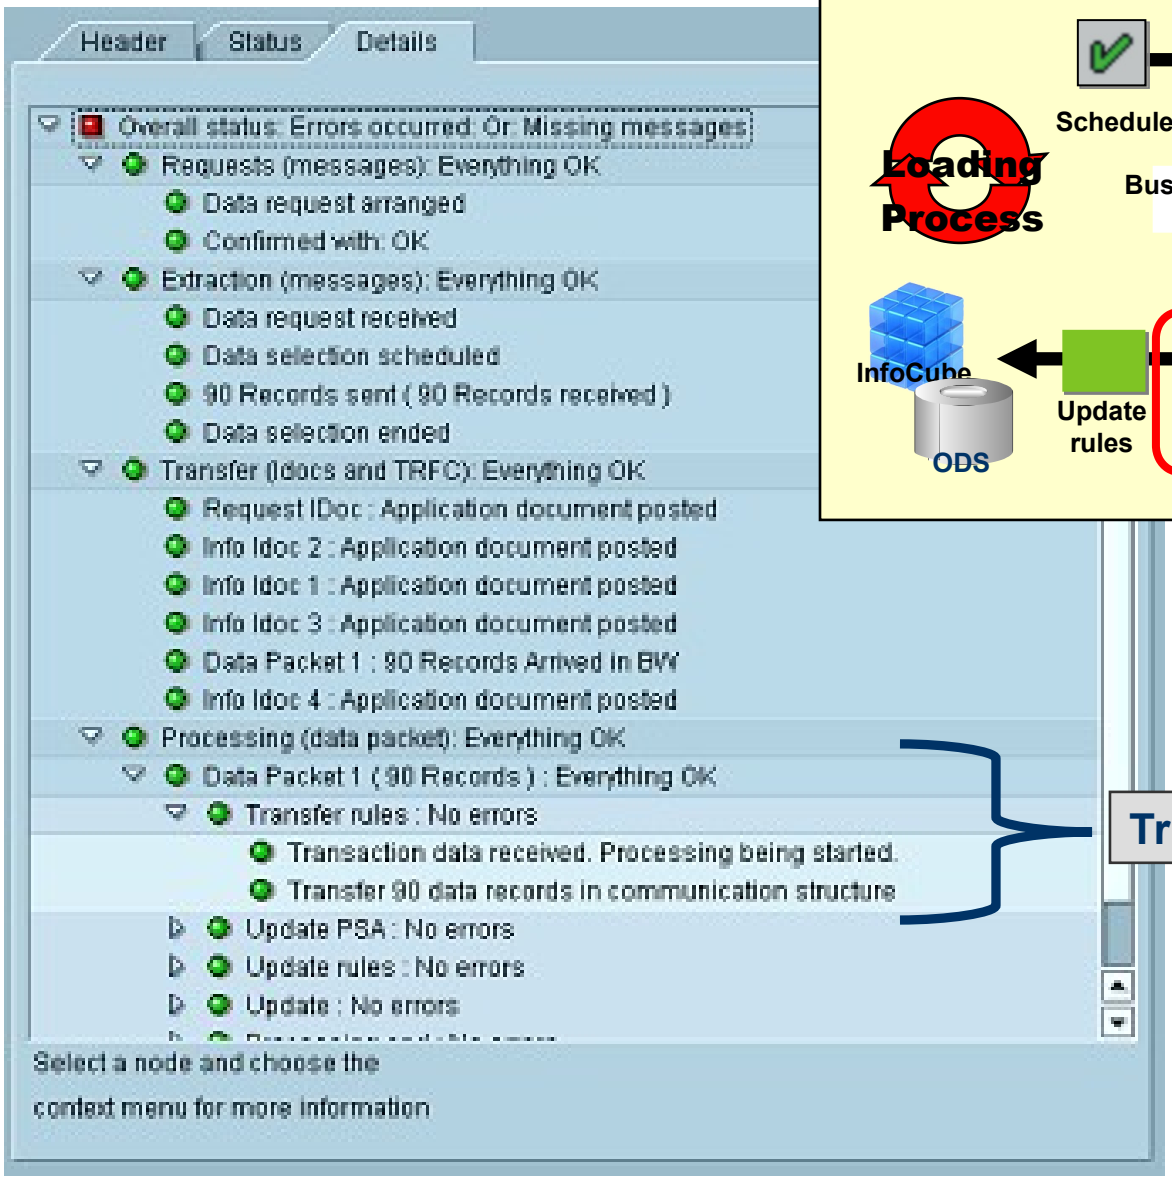

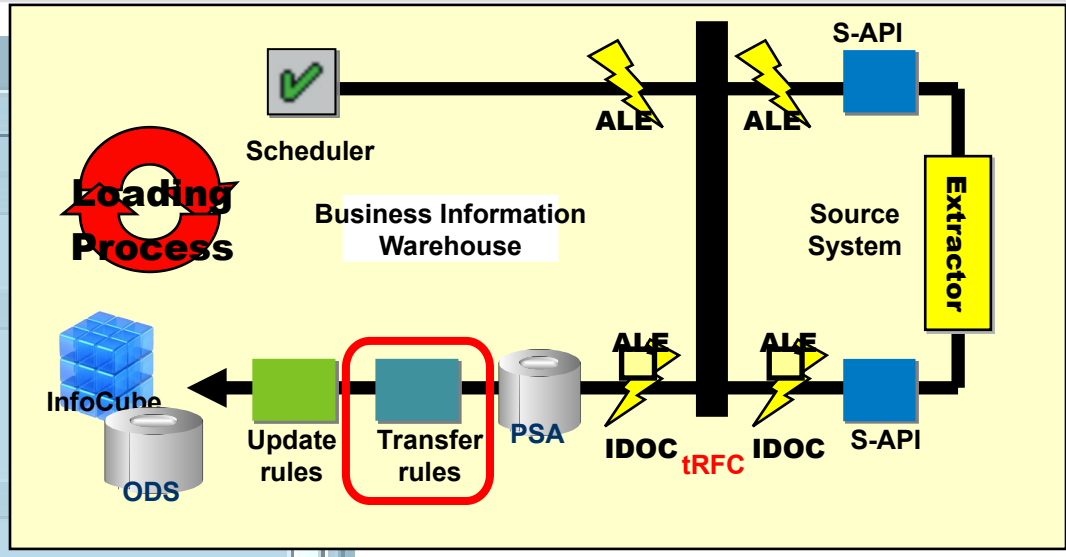

Data load Monitor : Transaction RSMO

Transfer Rules

Data Processing-Update Rules

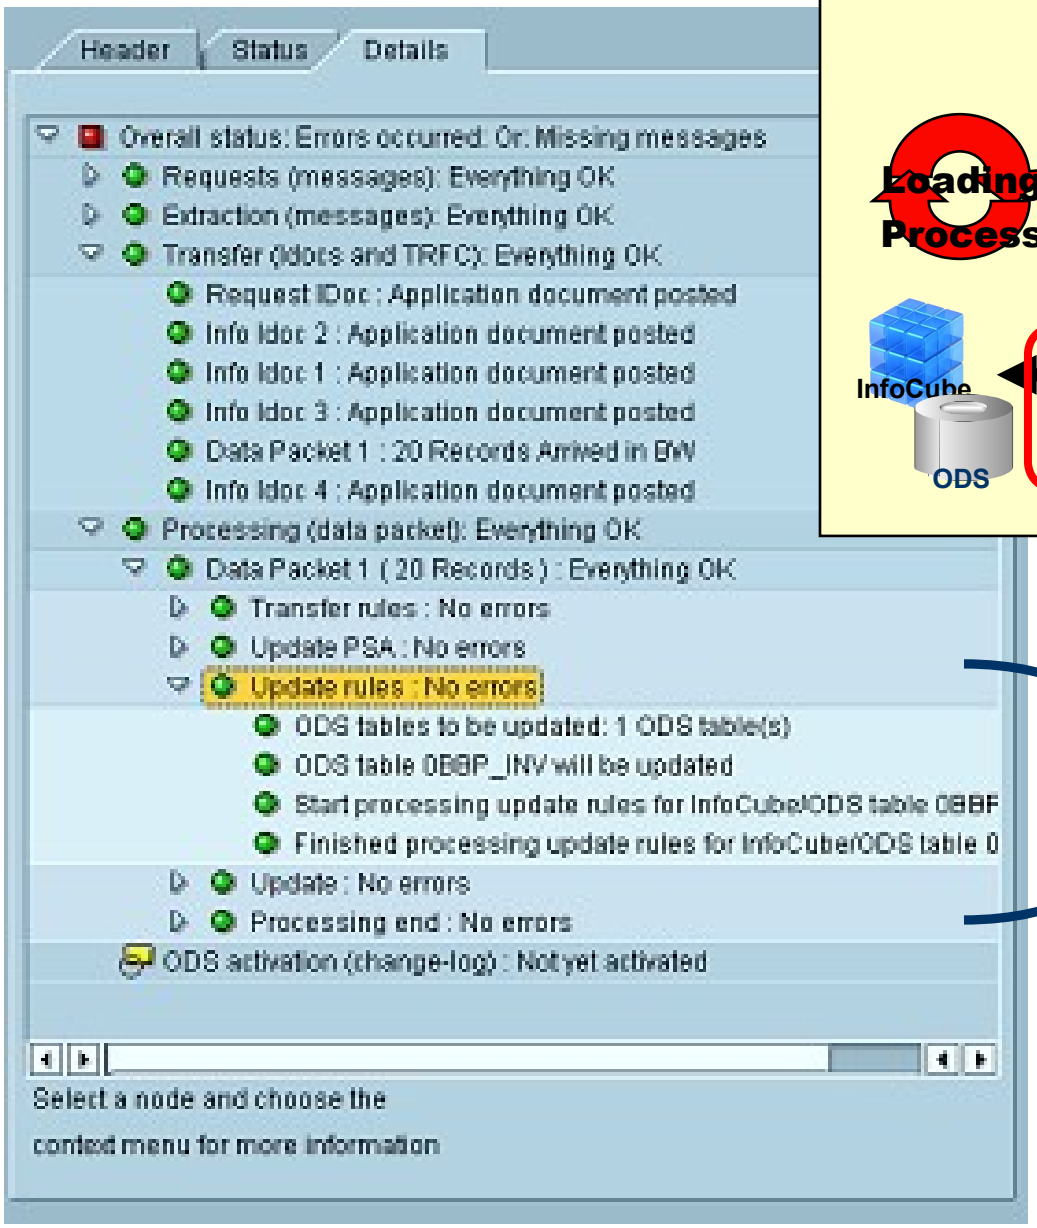

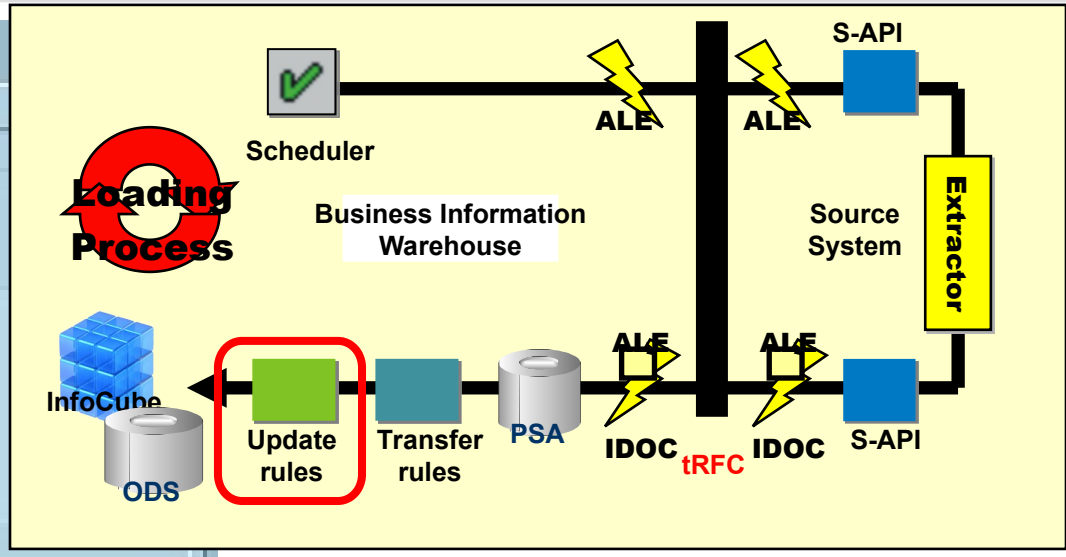

Update Rules

Data load Monitor : Transaction RSMO

Tuning Transfer and Update Rules (BW 3.0B)

Debugging and tuning Update and Transfer rules:

- Simple tool for debugging of transfer or update rules
- Improves error search and analysis – together with the enhanced error messages

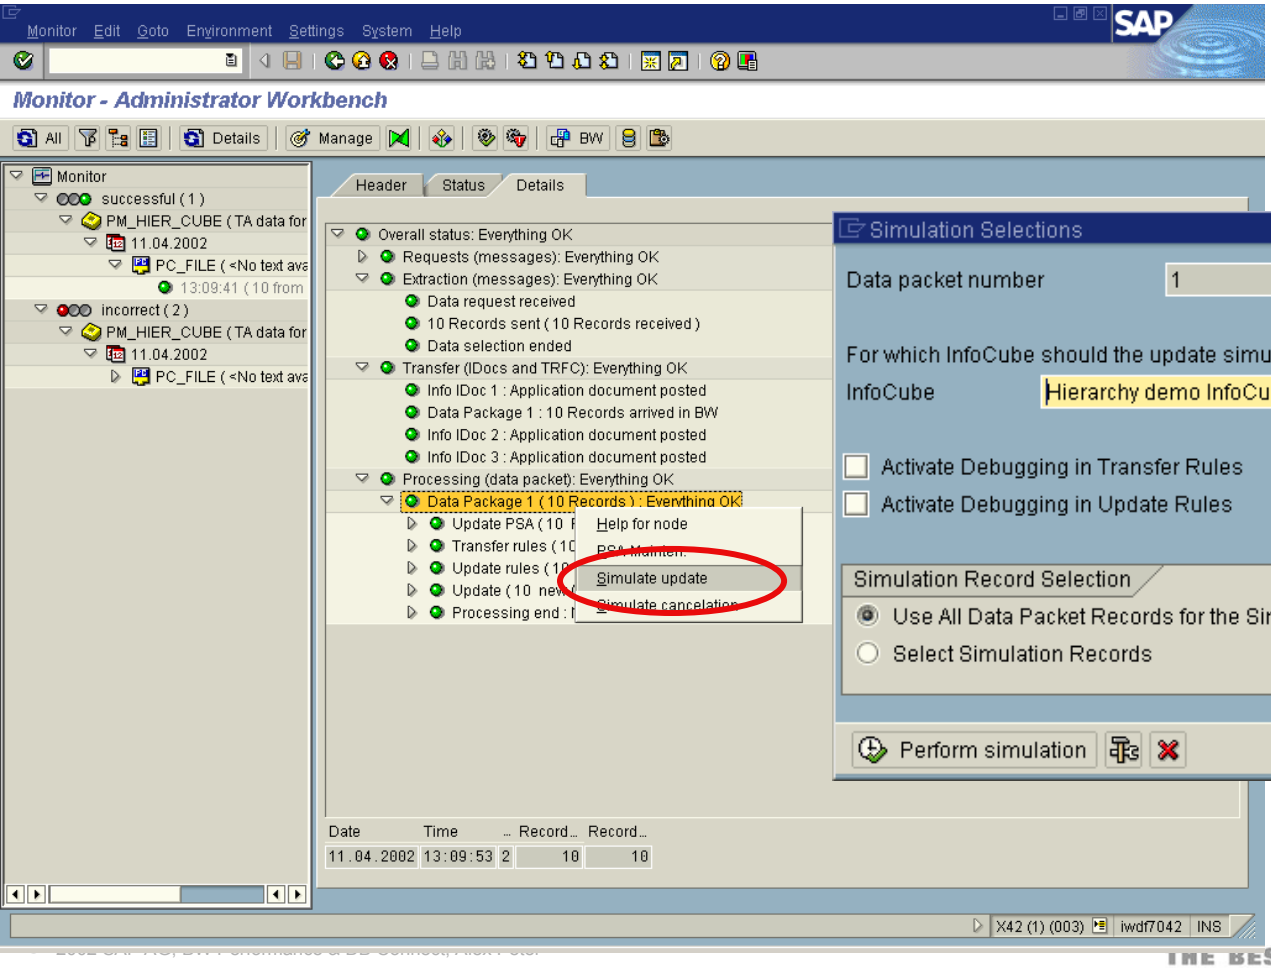

Routines: Potential Performance Bottlenecks

Identify the expensive update/transfer rules rules:

Debug from one update rule to the next update rule for each InfoObject.

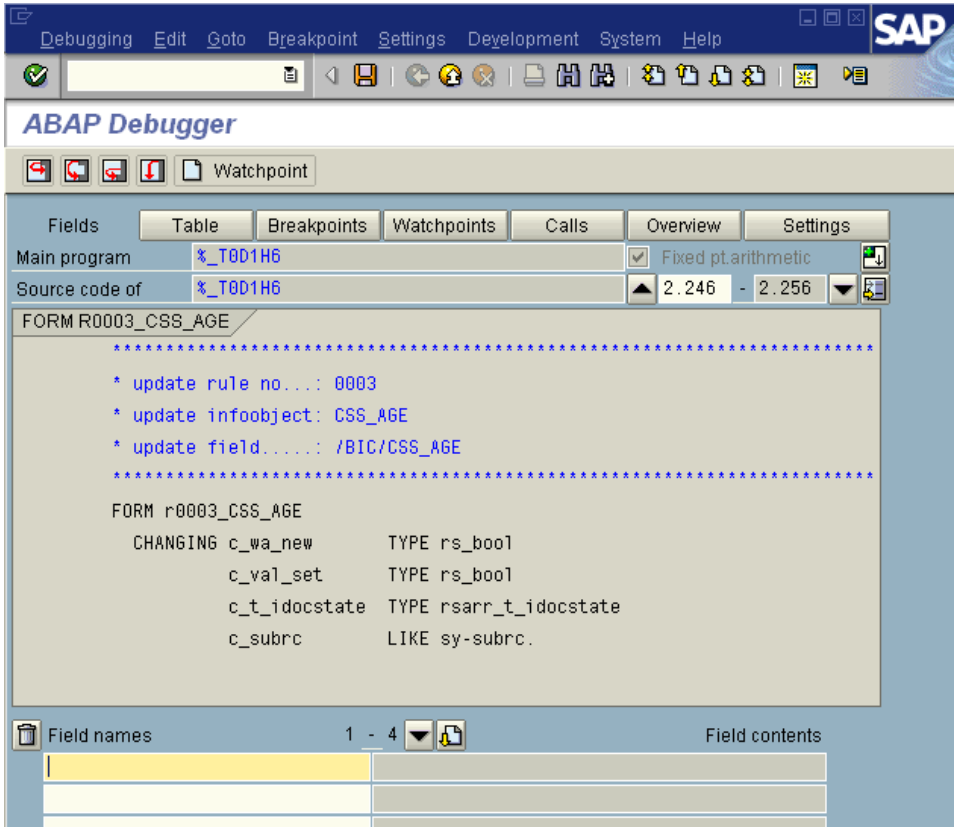

Also use ST05 or SM30

Recommendations:

- SINGLE SELECTs are one of the performance “killers” within these codings; use buffers (such as internal tables) and array operations instead.
- Avoid too many library transformations, as they are interpreted at runtime (currently not compiled like routines)
  - The transformation engine or library is new in BW 3.0

Data Load: Data Targets

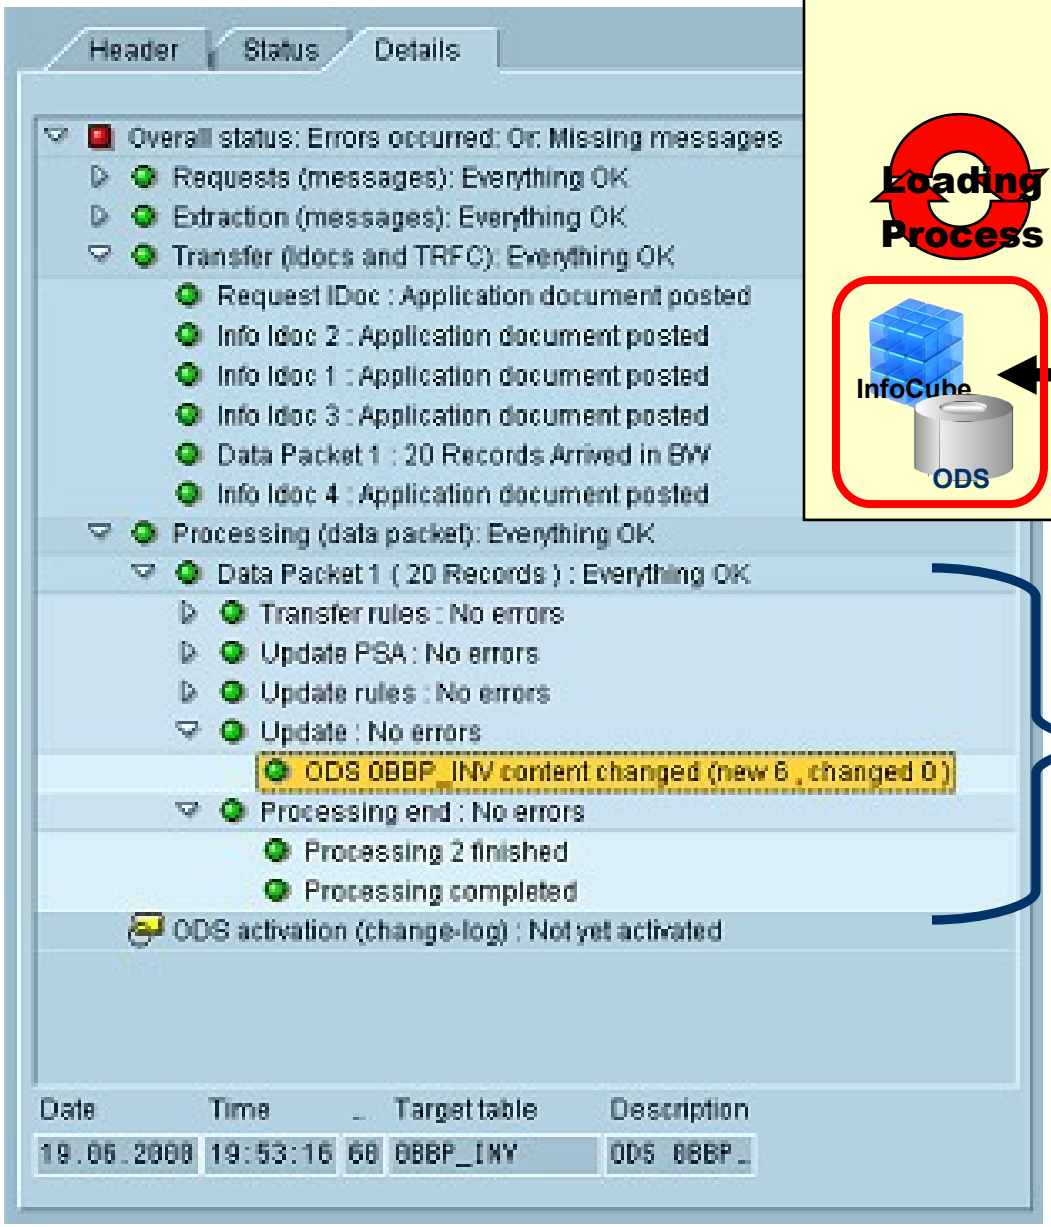

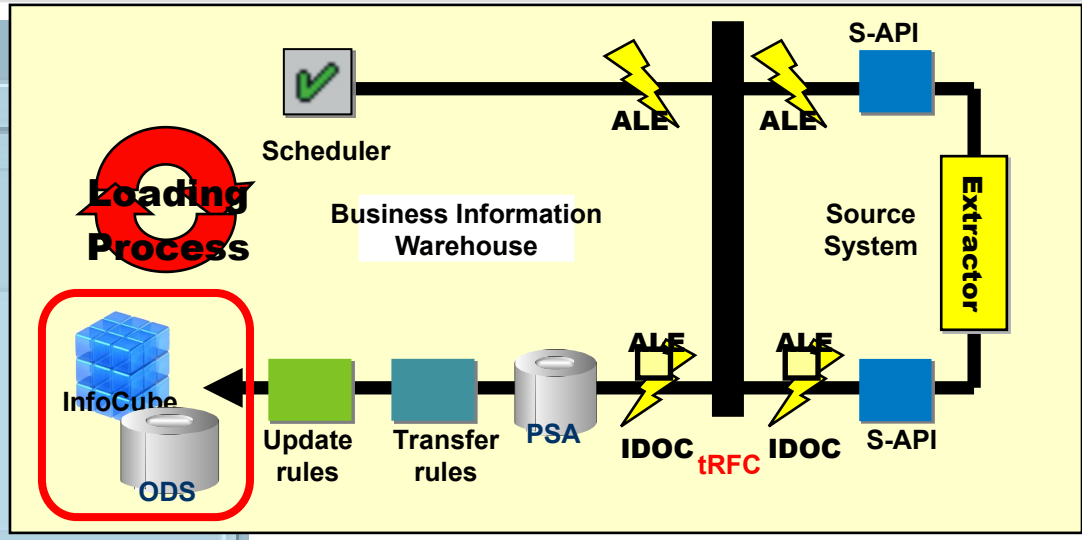

Load Into Data Targets

Data Load Monitor - RSMO

Initial and Large Data load volume tuning

Buffering Number Range (InfoCube):

Activate The number range buffer for the dimension ID's

Reduces application server access to Database.

- e.g. set the number range buffer for one dimension to 500, the system will keep 500 sequential numbers in memory
- SAP OSS note 130253: Notes on upload of transaction data into BW

Scenario:

- High volumes of transaction data: significant DB access (NRIV table) to fulfill number range requests.

Expected Results:

- Accelerates data load performance per load request.

Note:

- After the load, reset the number ranges buffer to its original state: minimize unnecessary memory allocation. Also, If the system crashes, the numbers held in the memory are gone.

Transactional Data Load Performance Tuning

Load Master data before transaction data

- Creates all SIDs and populates the master data tables (attributes and/or texts).
- SAP OSS note 130253: Notes on upload of transaction data into the BW

Scenario:

- Always load master data before transaction data.
  - ODS and InfoCube.
- When completely replacing existing data, delete before the load!

Expected Results:

- Accelerates transaction data load performance: all master data SID's are created prior to transaction load.

Transactional Data Load Performance

“Snapshot” Reporting: Data Deletion

- Some reporting scenarios require no historical data

Scenario:

- When completely replacing existing data, delete before load!

Expected Results:

- Data deleted from PSA can reduce PSA read times
- Data deleted from InfoCube reduces deletion and compression time.
  - “Drop partition...“ DDL statement instead of “delete from table...“ DML statement only takes seconds
- Deleting Data also speeds data availability (aggregates, etc)

ODS Activation in BW 2.x

Data Packets / Requests can not be loaded into an ODS object in parallel

- overwriting functionality
- Locking on Activation table

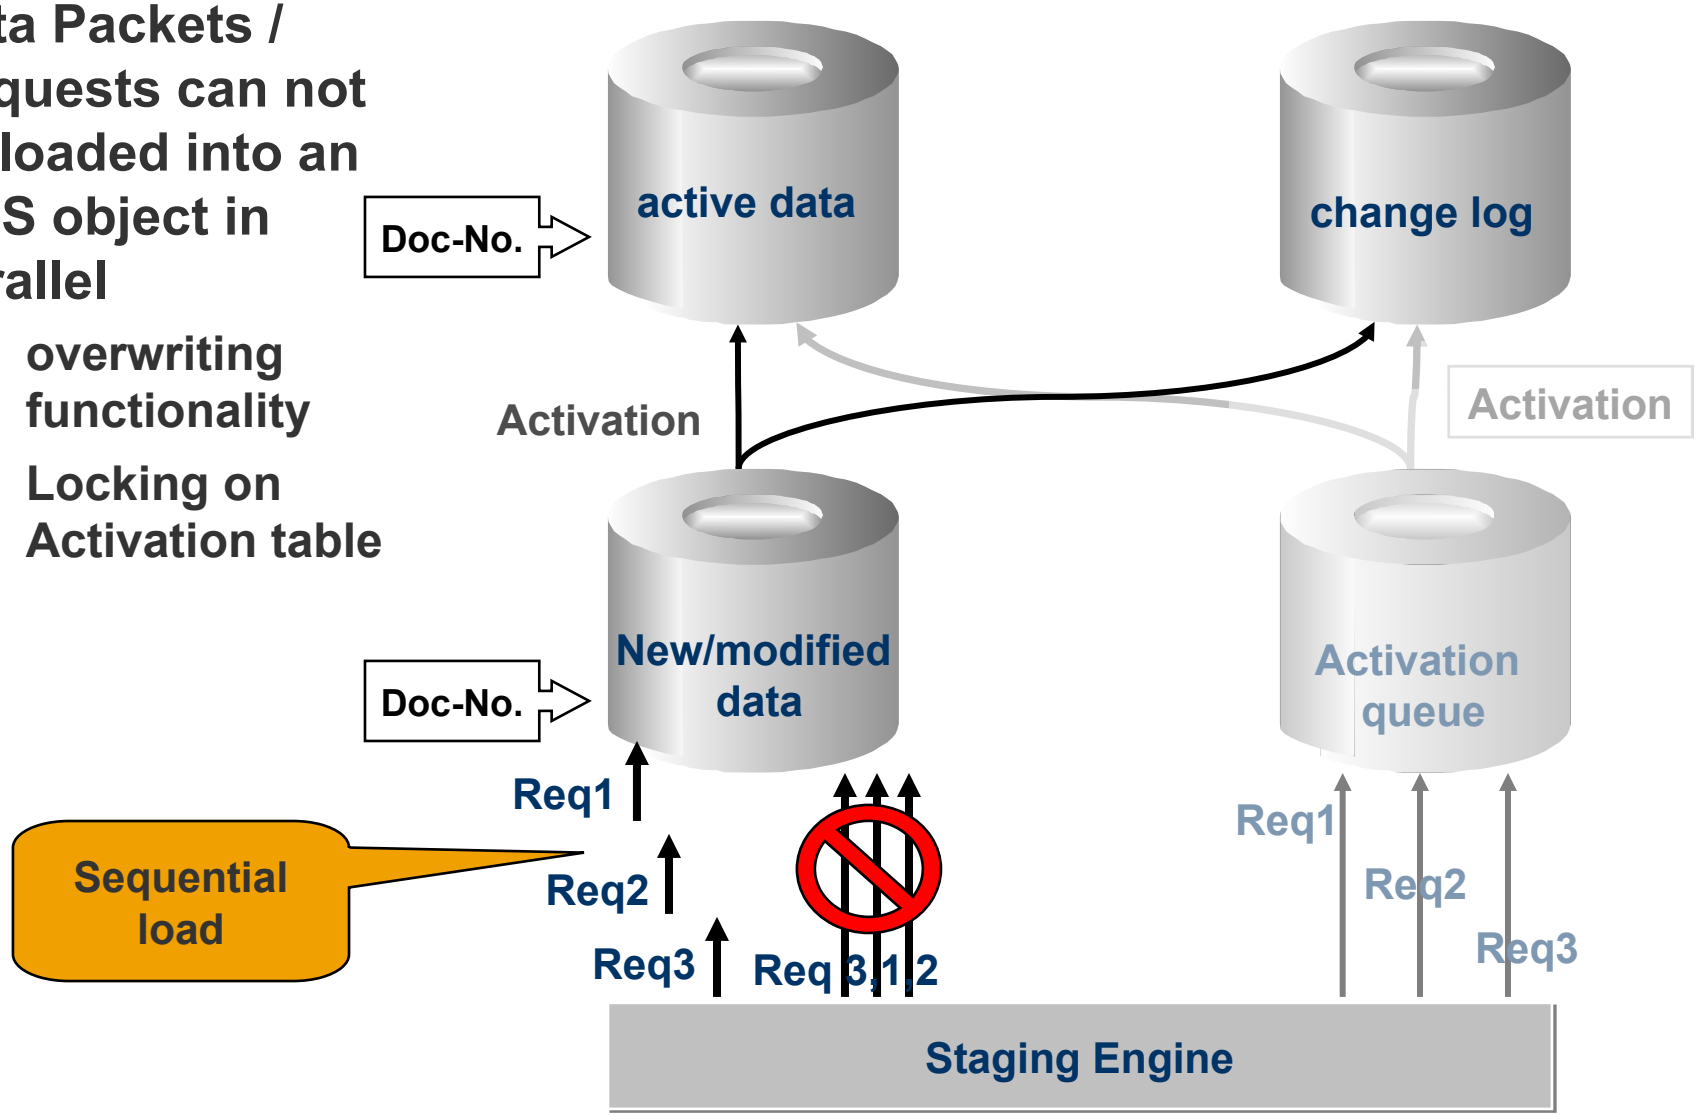

ODS Activation in BW 3.x

New queuing mechanism replacing previous Maintenance (M)table

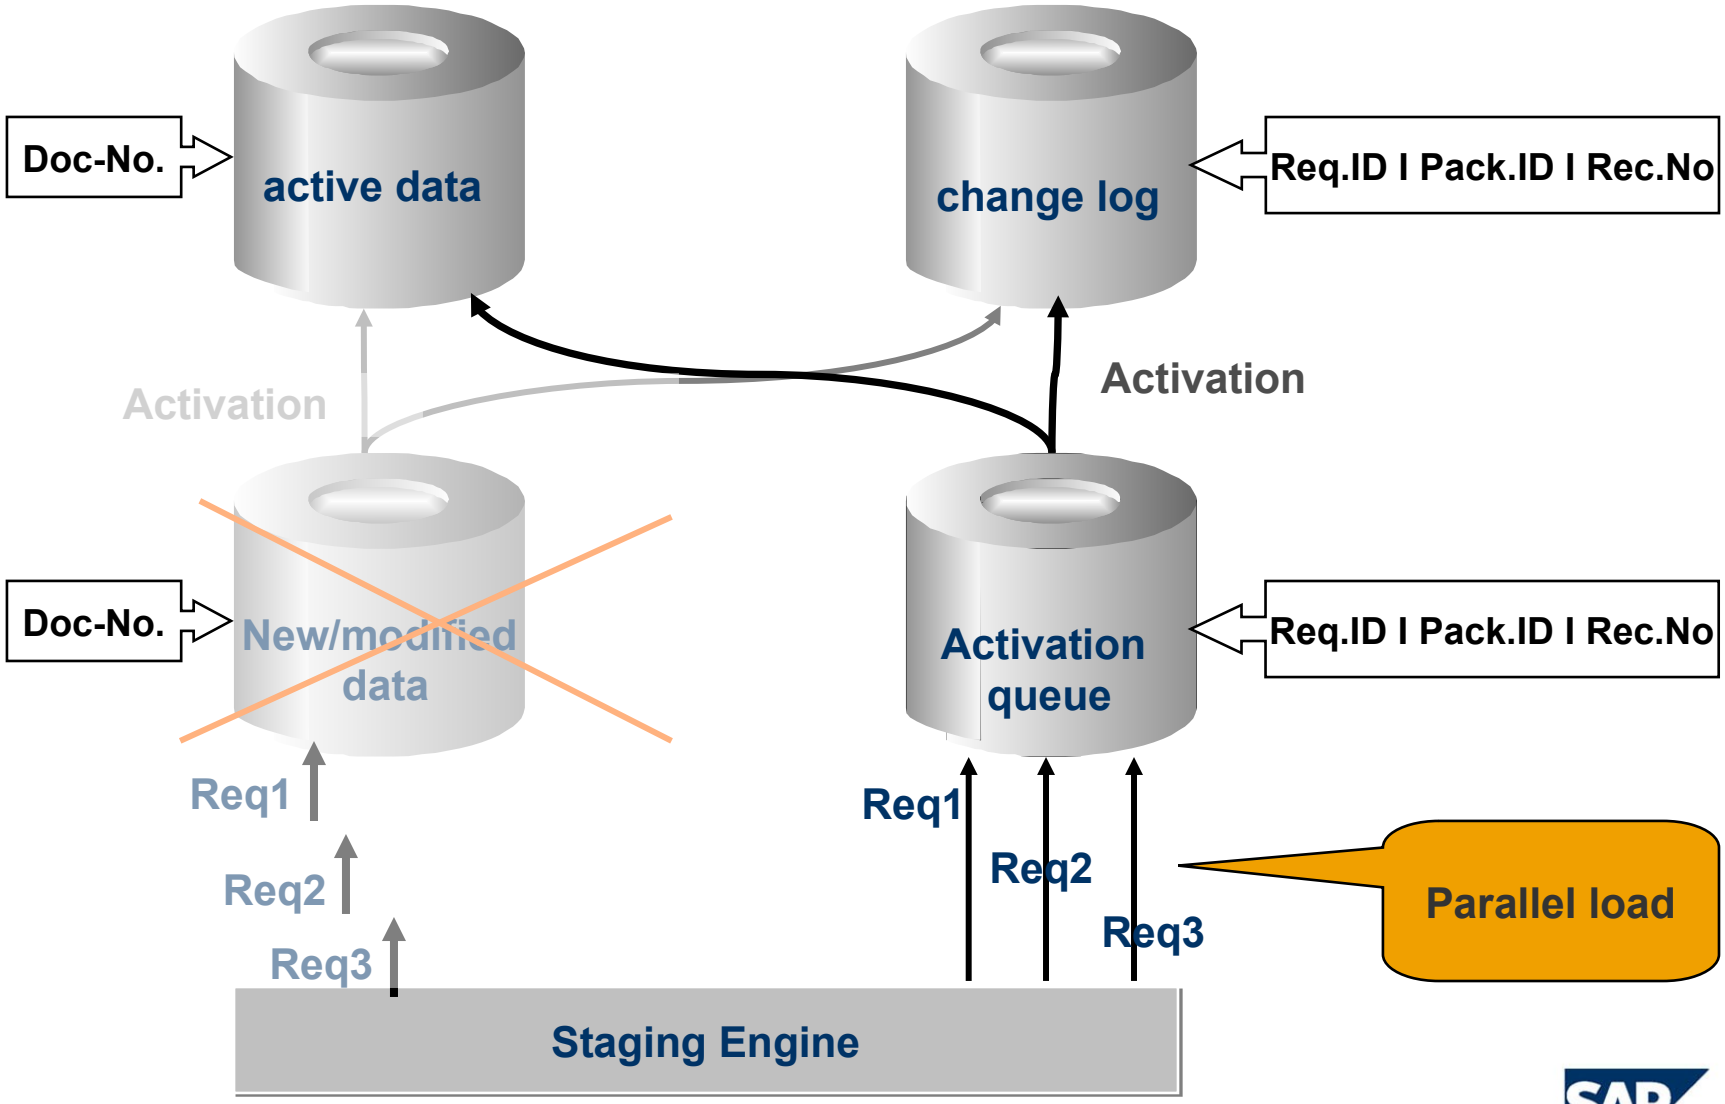

ODS Activation example (BW 3.0)

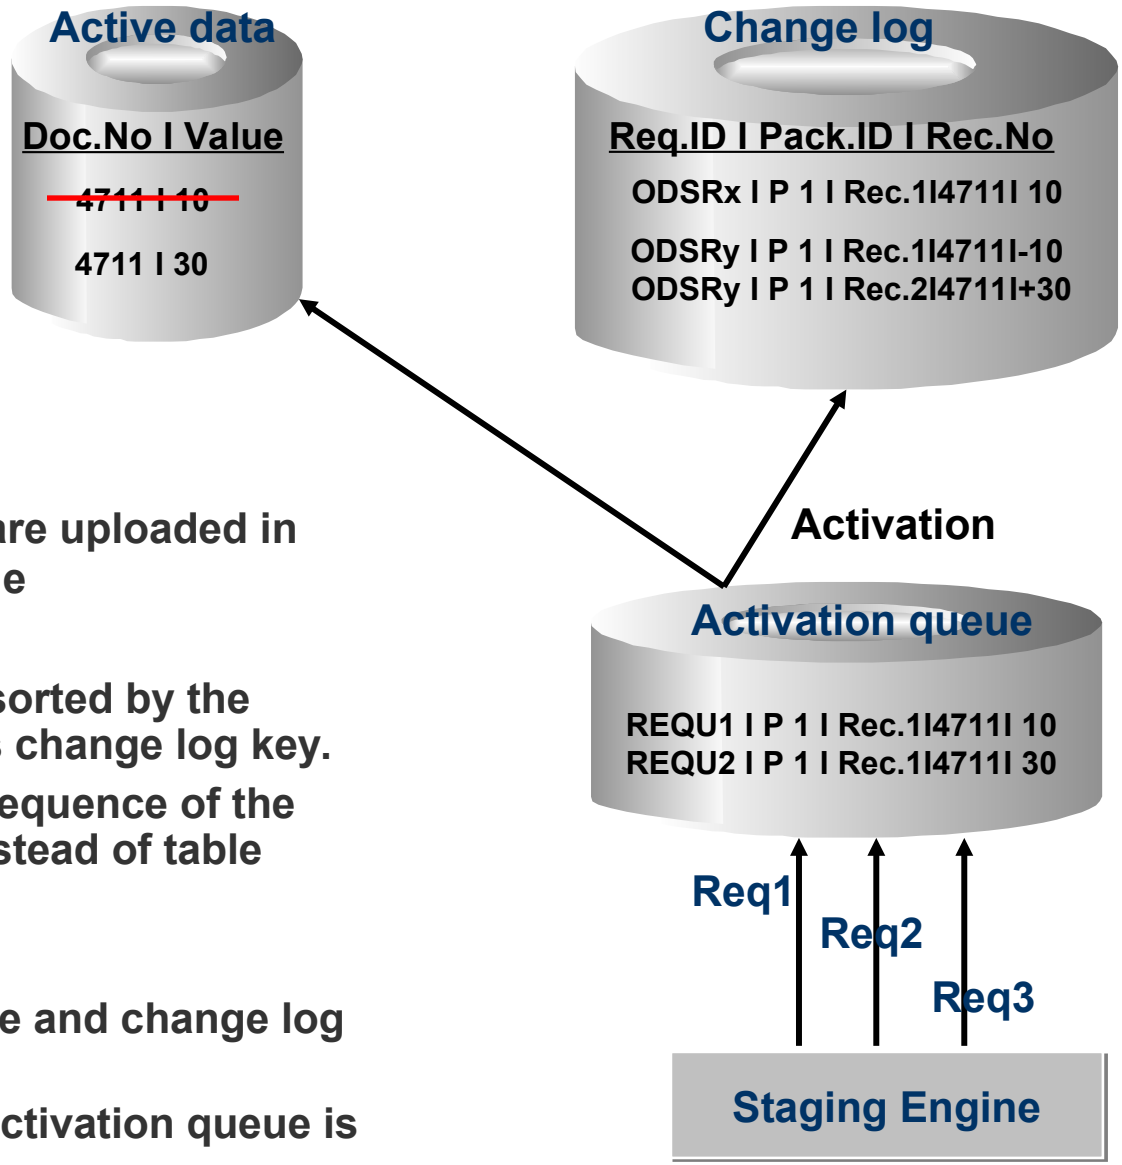

Upload to Activation queue

- Data from different requests are uploaded in parallel to the activation queue

Activation

- During activation the data is sorted by the logical key of active data plus change log key.
- This guarantees the correct sequence of the records and allows inserts instead of table locks.

Before- and After Image

- Request ID in activation queue and change log differ from each other.
- After update, the data in the activation queue is deleted.

Control of ODS Data load Packaging and Activation

Transaction RSCUSTA2

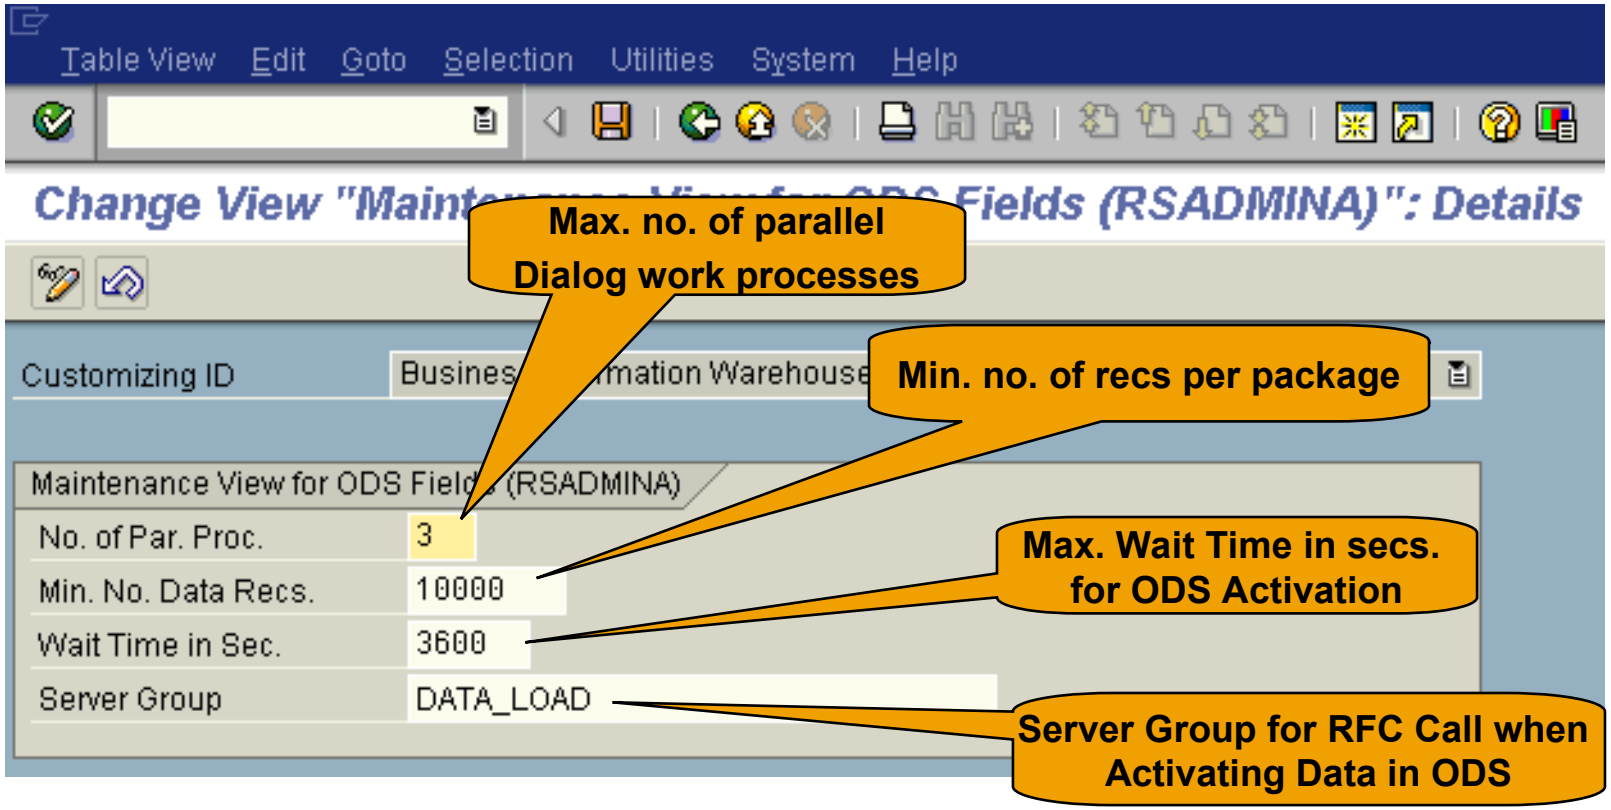

Controls data packet size utilized during parallel update/activation and number and allocation of work processes.

ODS Load/Activation Tuning Tips

Non-Reporting ODS Objects:

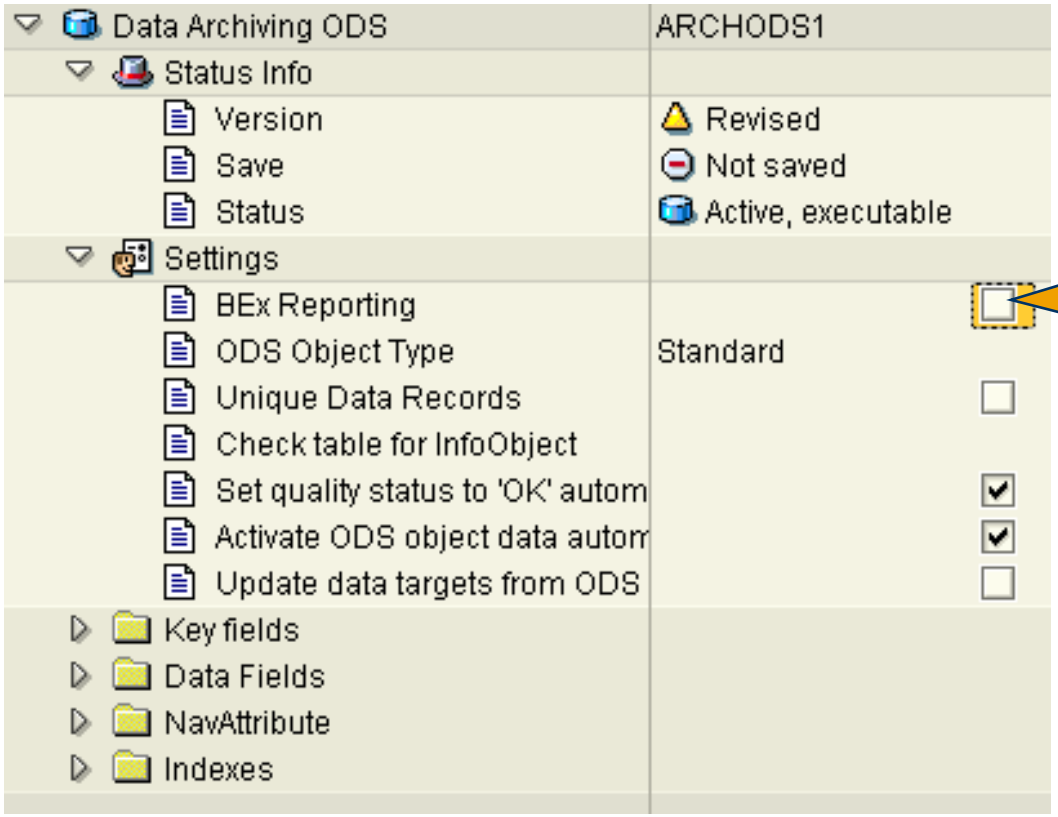

BEx Flag: Computation of SIDs for the ODS can be switched off

BEx-flag must be switched on if BEx-reporting on the ODS is executed

Loads are faster as Master Data SID tables do not have to be read and linked to the ODS data

Further ODS Loading/Activation enhancements (3.0)

Update of ODS object with unique records

- Significantly simplifies activation process
- No lookup of existing key values
- No updates in active table, only inserts
- **Note: User is responsible for uniqueness!!**

Index maintenance

- Index speed querying
- Slow down activation

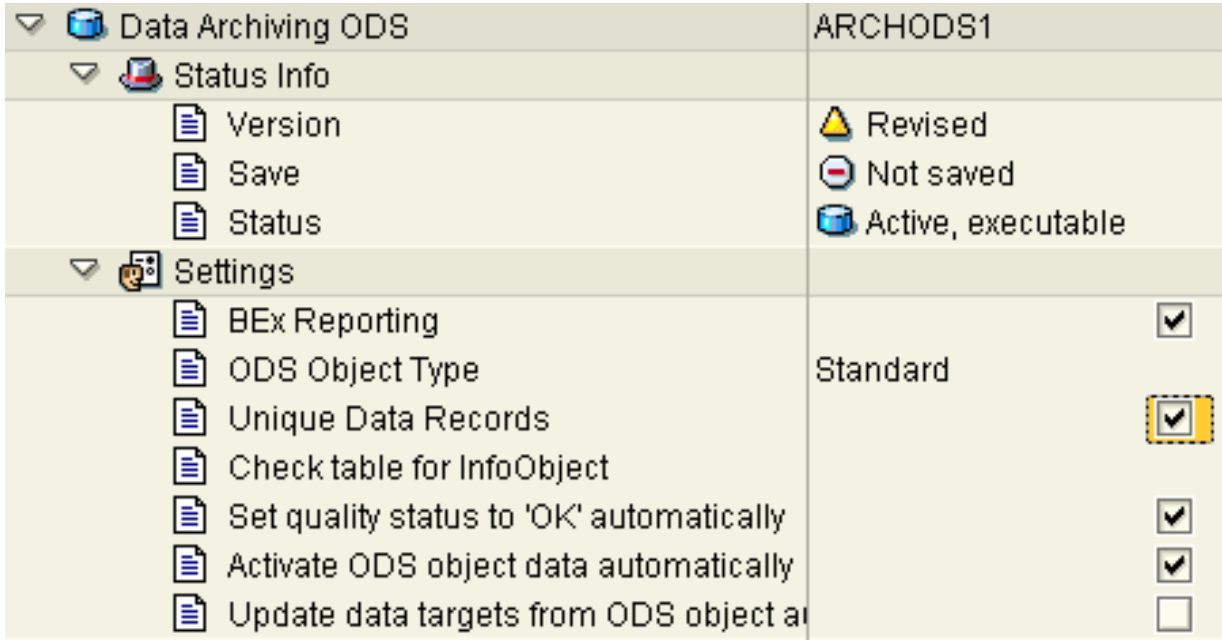

Parallel SID creation

- SIDs are created per package
- Multiple packages are handled in parallel by separate dialog processes

Using InfoCube Data Load Performance Tools

Admin WB > Modeling > InfoCube Manage > Performance Tab

Recommendation: Drop secondary Indexes for large InfoCube data loads

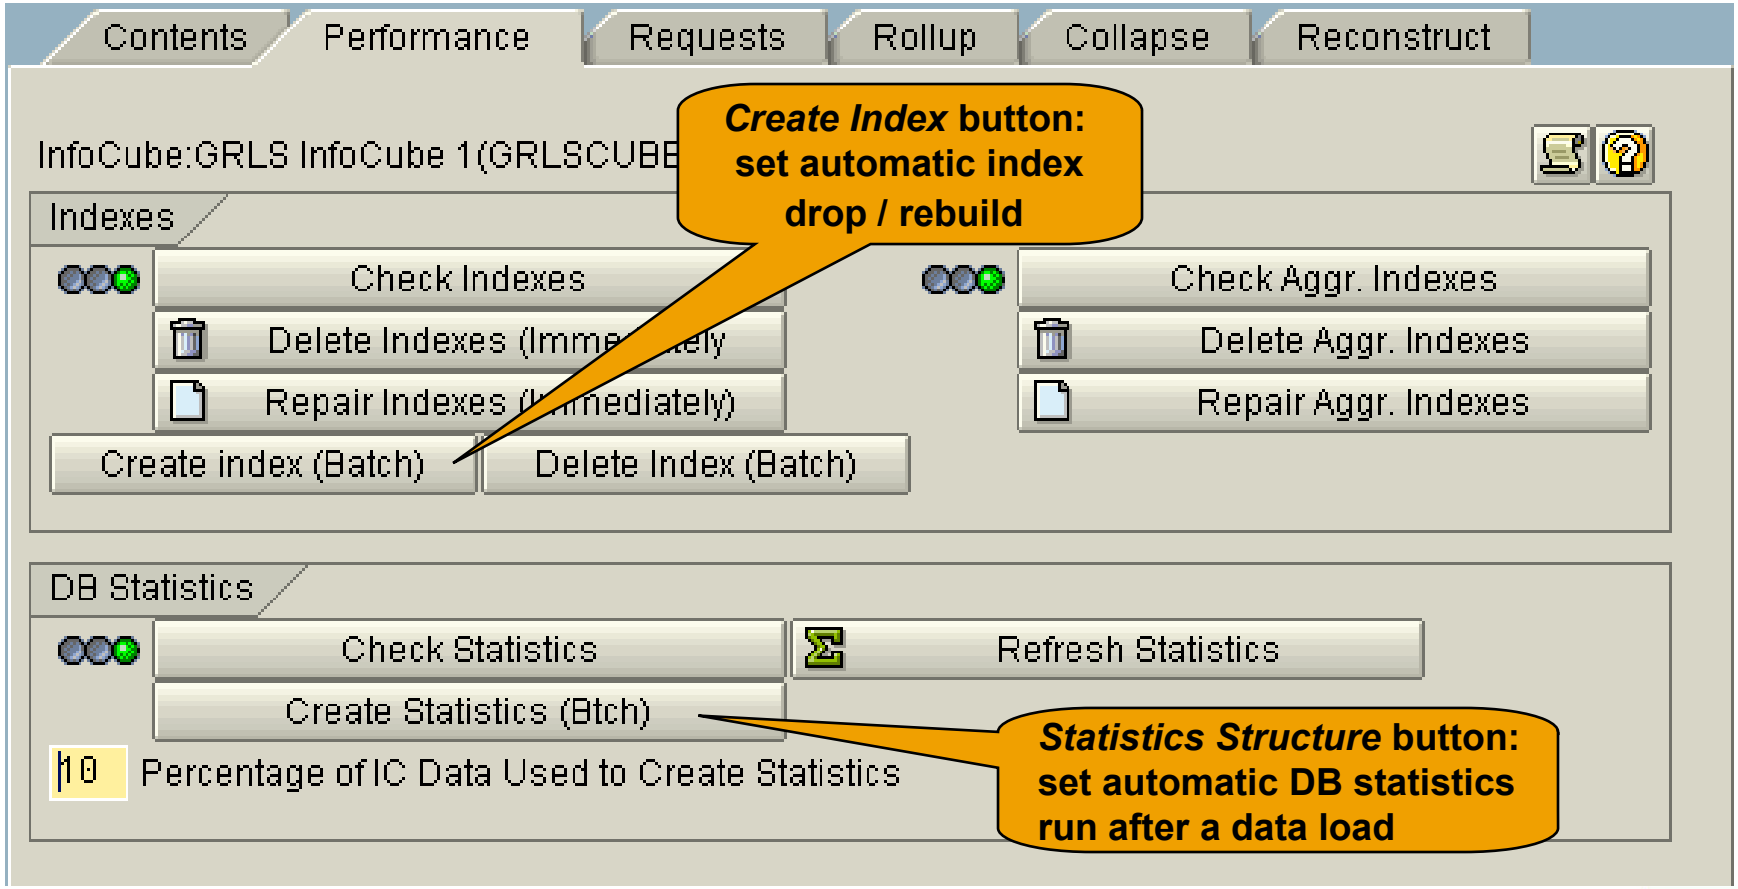

Using InfoCube Data Load Performance Tools

New in BW 3.X

Process Chains

- Replacement for Event chains

Transaction RSPC

Process type :

- Delete Index
- Generate Index
- Auto suggestion depending on InfoPackage selected.

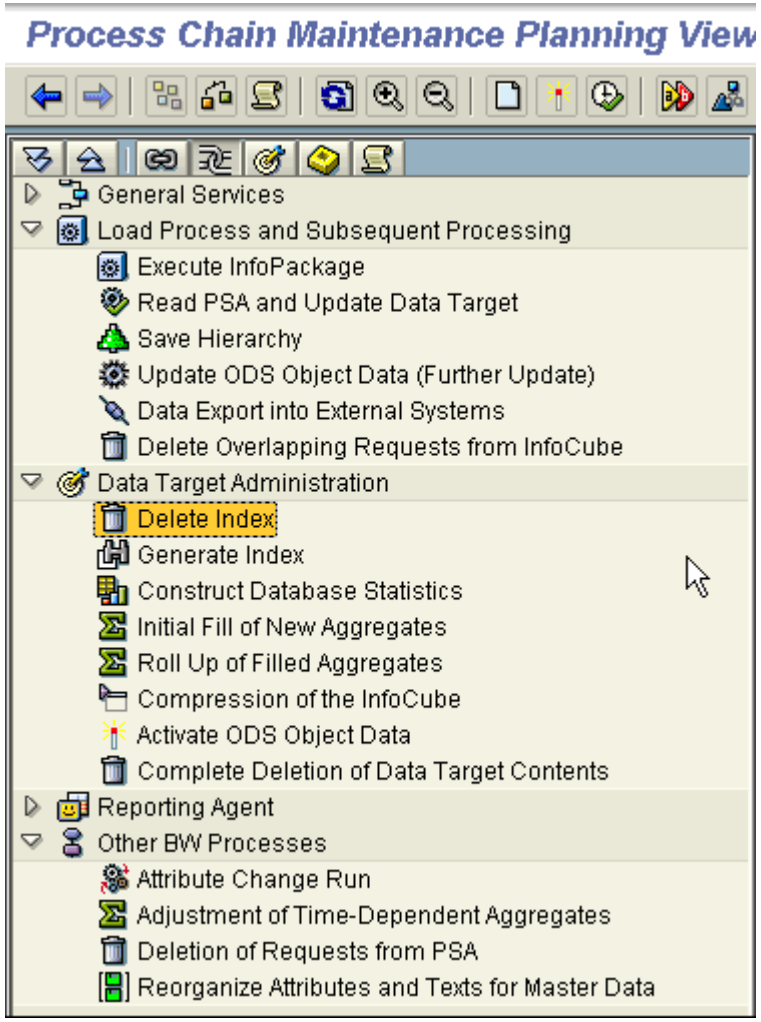

Archiving in BW

Archiving to reduction data volume

- IDOC archiving available in BW and R/3 (BW 1.2 onwards)
  - Improves extraction/load performance (including use of load monitor)
- Data Archiving (InfoCubes and ODS Objects) available in BW 3.x

Archiving TechEd session: ARCH201 Data Archiving in SAP R/3 Enterprise

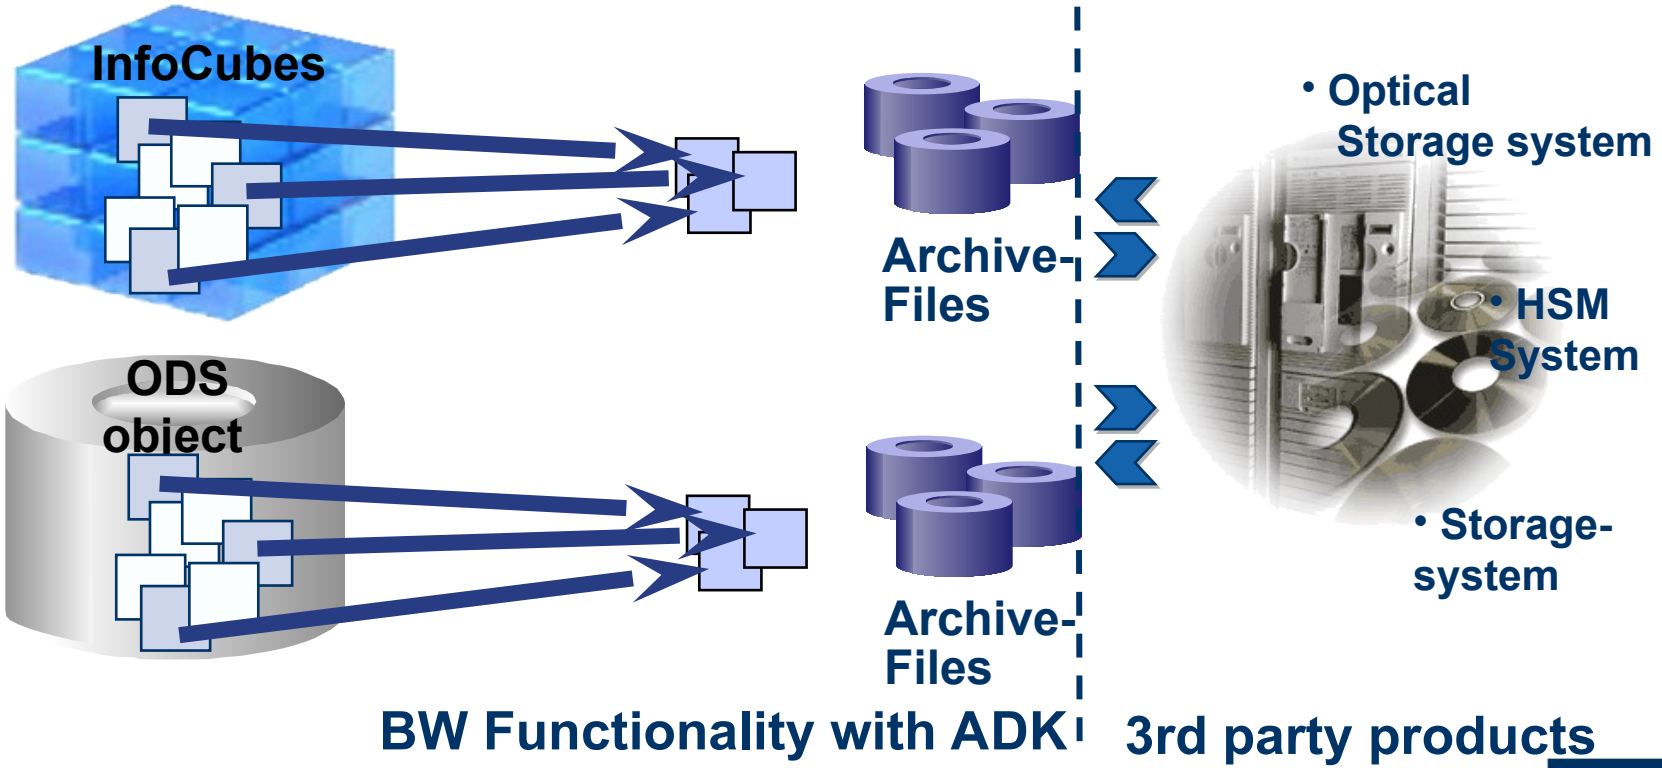

Data Load Performance Concepts

Reporting Performance Analysis Tools

Reporting Performance Analysis Techniques

Aggregates

Reporting Performance Tuning Concepts

Architecture: Reporting

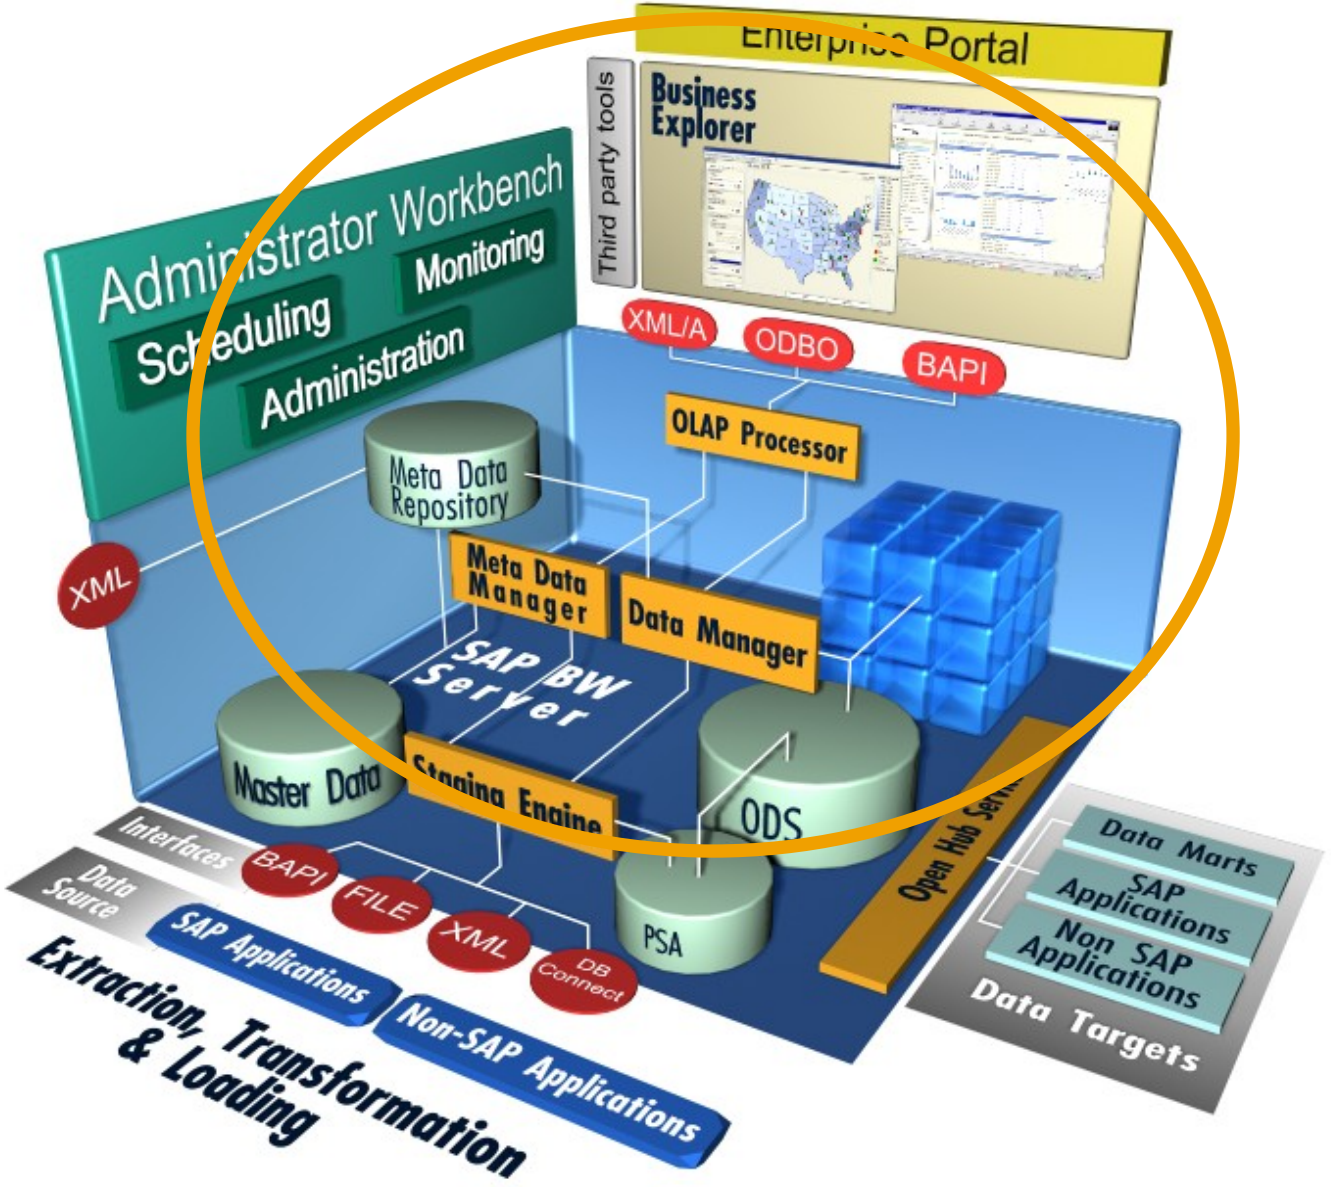

Many Aspects of Performance

Are there resource constraints at the hardware or DB level impacting system-wide performance?

Do the queries of one InfoCubes generally run slower than the queries of other InfoCubes?

Which queries are the biggest performance concerns?

- Frequently executed by many users
- Lots of navigations by analysts
- High average runtimes

Performance monitoring and tuning is part of a larger design and system maintenance strategy.

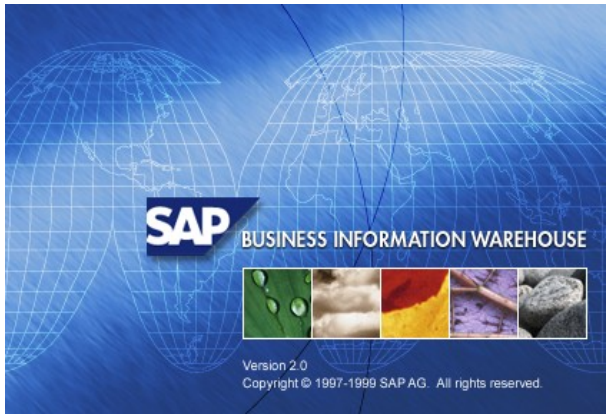

Best Practices:

- Performance should be considered in design
- Monitoring and tuning efforts are ongoing and proactive!

Analysis Tools: Operating System Monitor (ST06)

Scenario

- General system performance is bad
- You assume a hardware / OS bottleneck

Useful Key Figures

- Current Values (snapshots) and history (previous hours)
  - CPU
  - Memory
  - Swap

Check following parameter values:

- CPU utilization
- Free memory

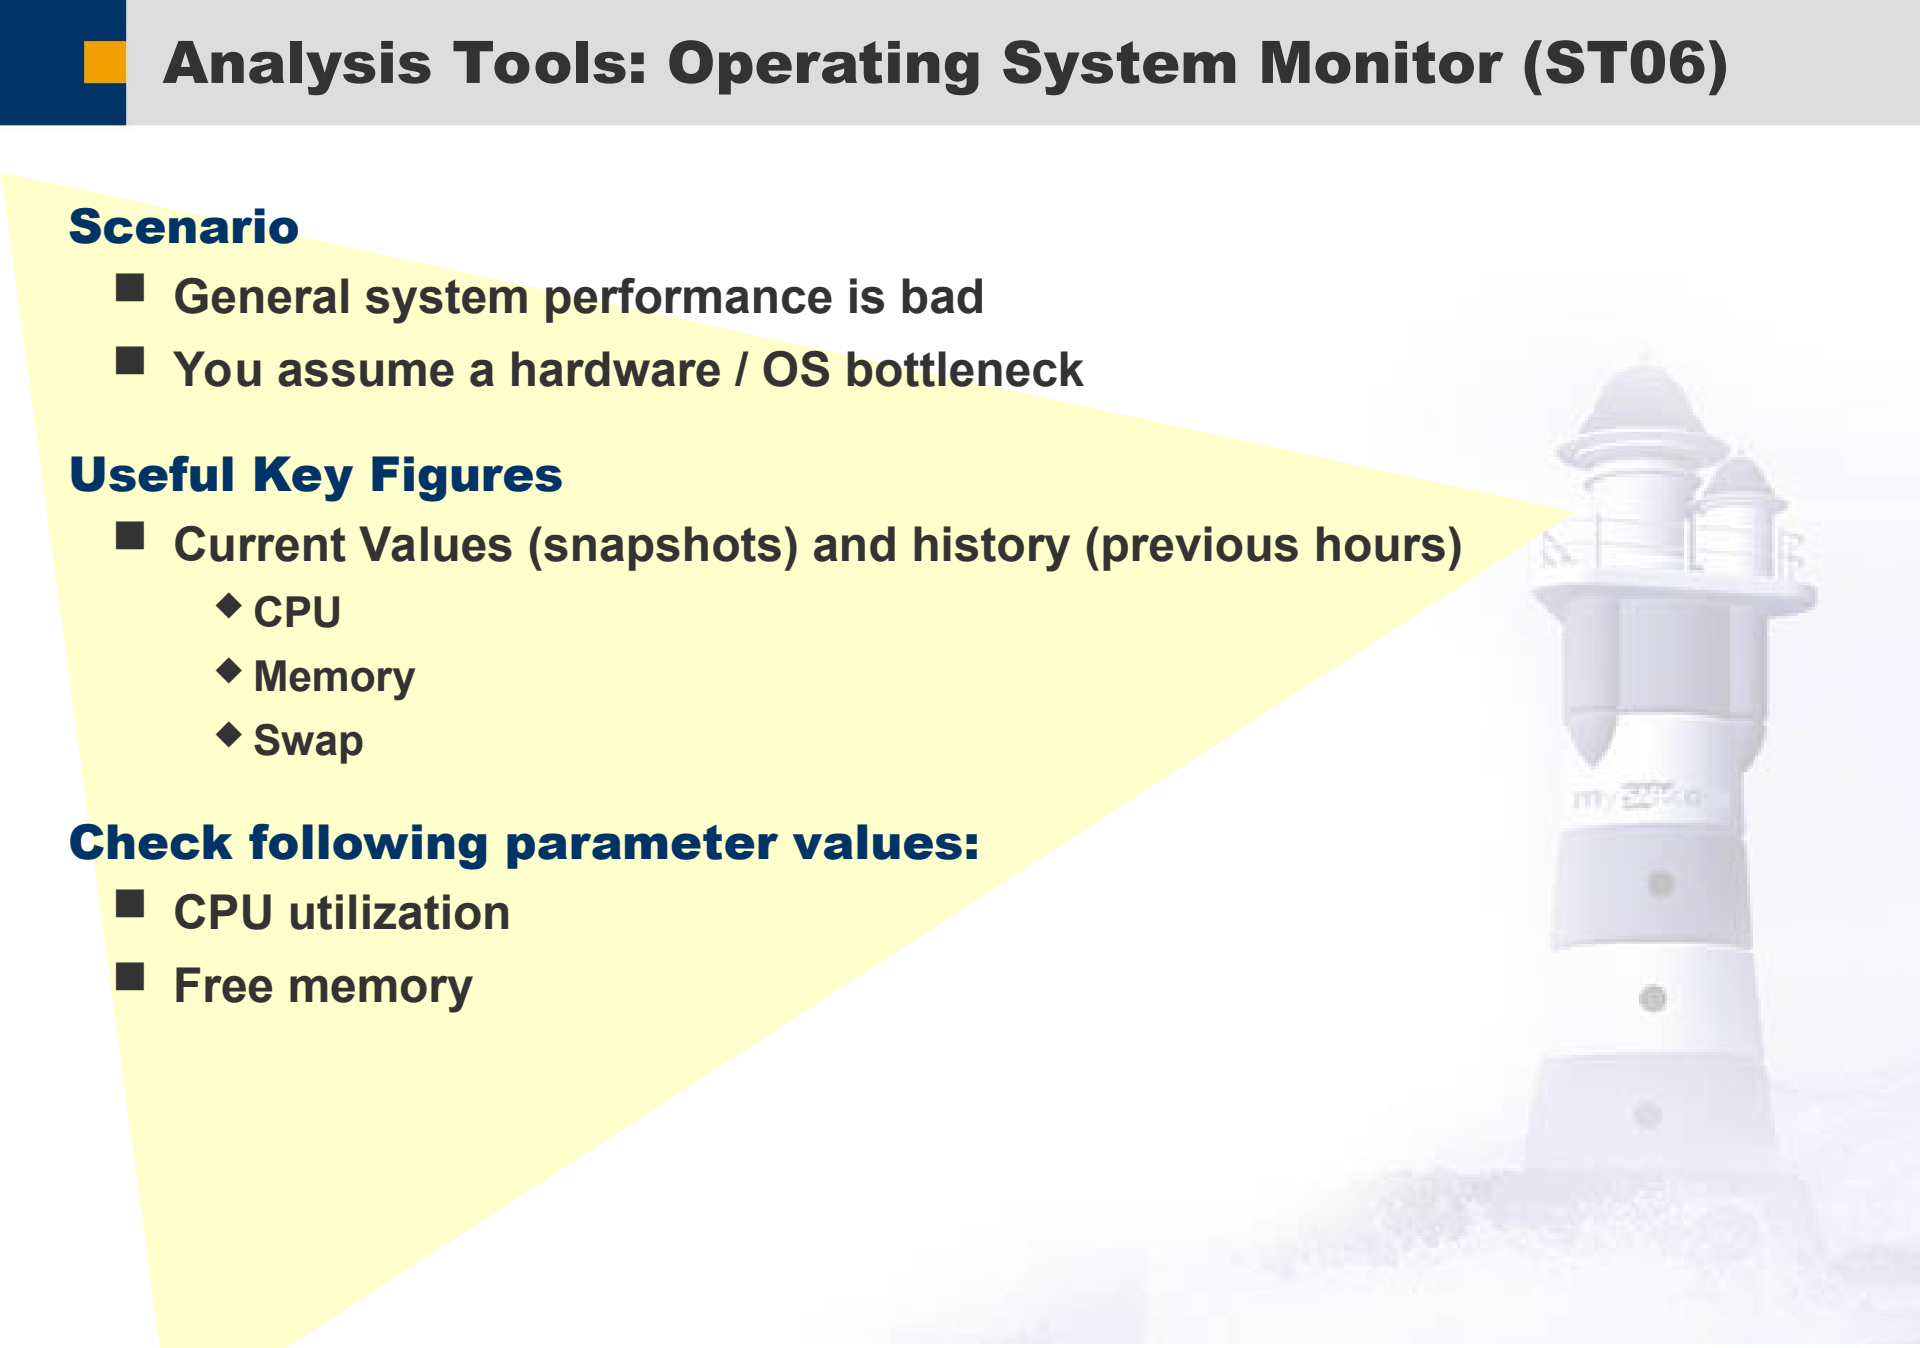

Analysis Tools: Database Monitor (ST04)

Scenario

General DB performance is bad; you assume a DB bottleneck

Useful Key Figures

- Deadlocks
- DB Buffer Hit Rates
- DB Process Overview
- Explain Query

Check following parameter values:

- Check for deadlocks
- Hit ratios on SGA buffers > 95% (or higher)
- Which access path in the explain plan?
  - Are the DB statistics up to date?
  - Are all indexes available?

Analysis Tools: Database Performance Monitor (DB02)

Scenario

- General DB performance problems
- You assume missing indexes or
  - (e.g. in ORACLE) degenerated indexes

Useful Features

- Missing Indexes
- Table, Index analysis

Check following parameter values:

- Compare index size with table size

Analysis Tools: Buffer Monitor (ST02)

Scenario

- The general system performance is bad
- You are concerned about efficient memory utilization

Important Buffers for BW

- Generic Key Table Buffer
  - Most BW-specific control tables (RS*-tables)
- Single Record Key Table Buffer
  - Master Data Tables
- Export/Import Shared Memory
  - OLAP Query Cache

Check following parameter values:

- Hit Ratio
- Swaps
- Free Space / Free Directory

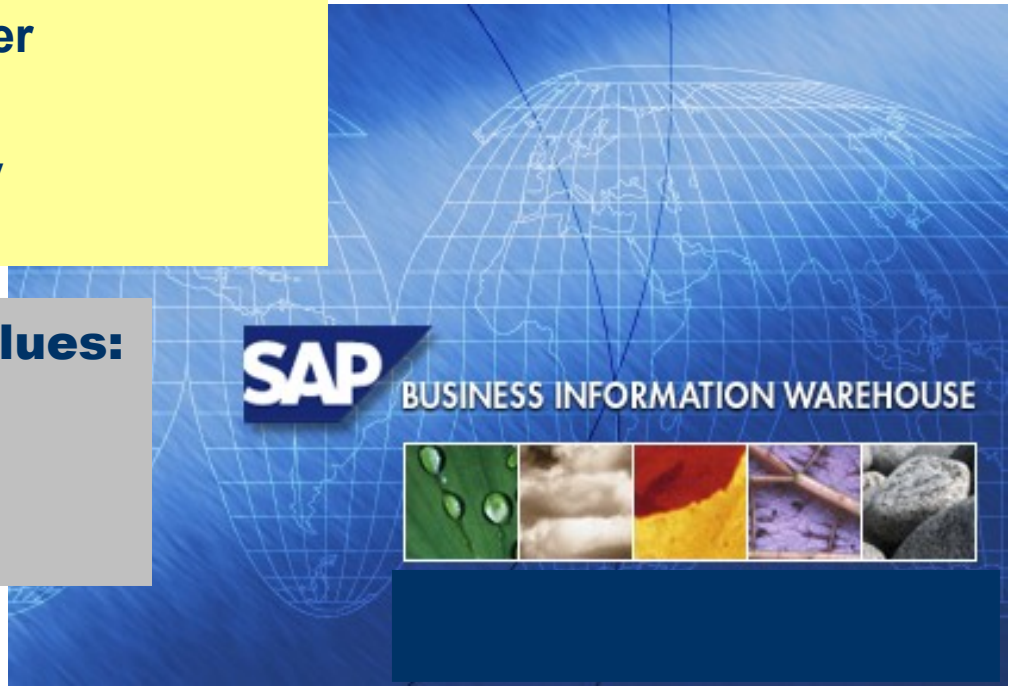

Analysis and Repair of BW Objects (RSRV)

Scenario

Performance of queries of one specific InfoCube is bad

Useful Features

- Database
  - DB statistics for an InfoCube and its aggregates
  - DB indices for an InfoCube and its aggregates
  - DB parameter settings check
  - DB Information about InfoProvider tables
    - InfoCube: Relative size of dimension tables compared to fact table
- "Repair" feature

Check following parameter values:

- Ensure DB statistics are up-to-date for an InfoCube
- Ensure indices exist for InfoCube
- Check for unbalanced InfoCubes
  - (dimension table size 10% or more compared fact table size)

RSRV – Tools for Analysis: Database tools

Admin WB > Modeling > InfoCube Manage > Goto menu > Data Target Analysis or Transaction RSRV

Analysis of BW Objects

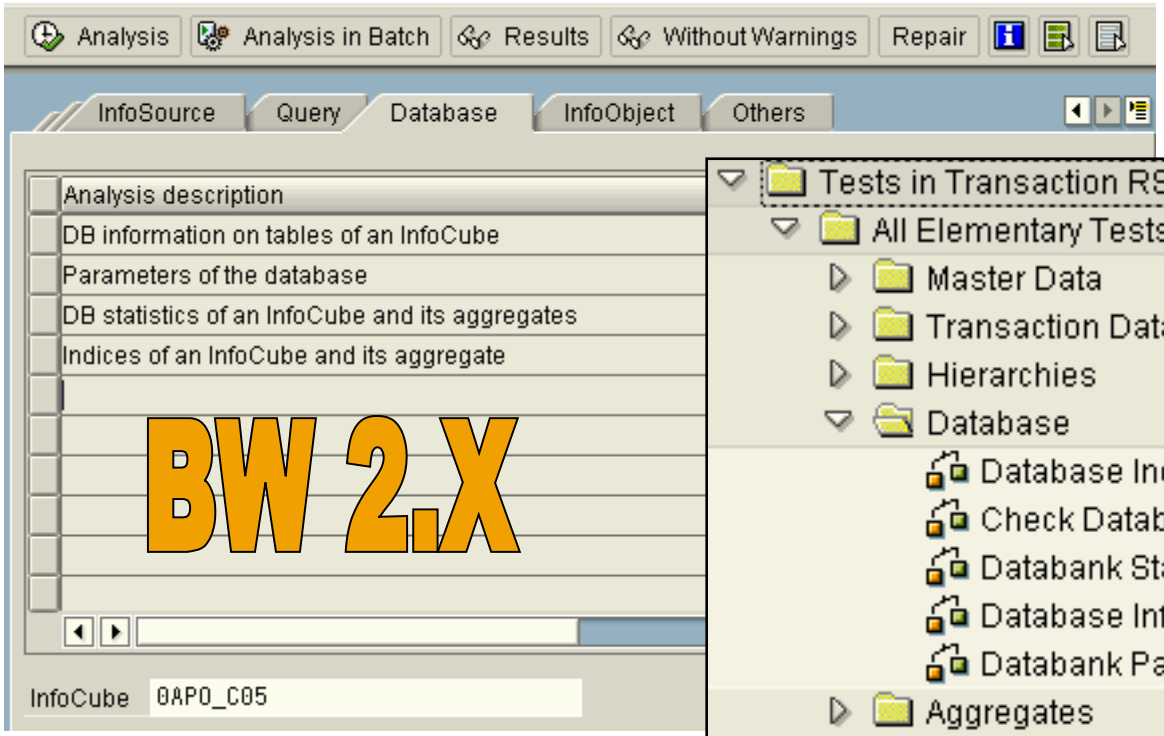

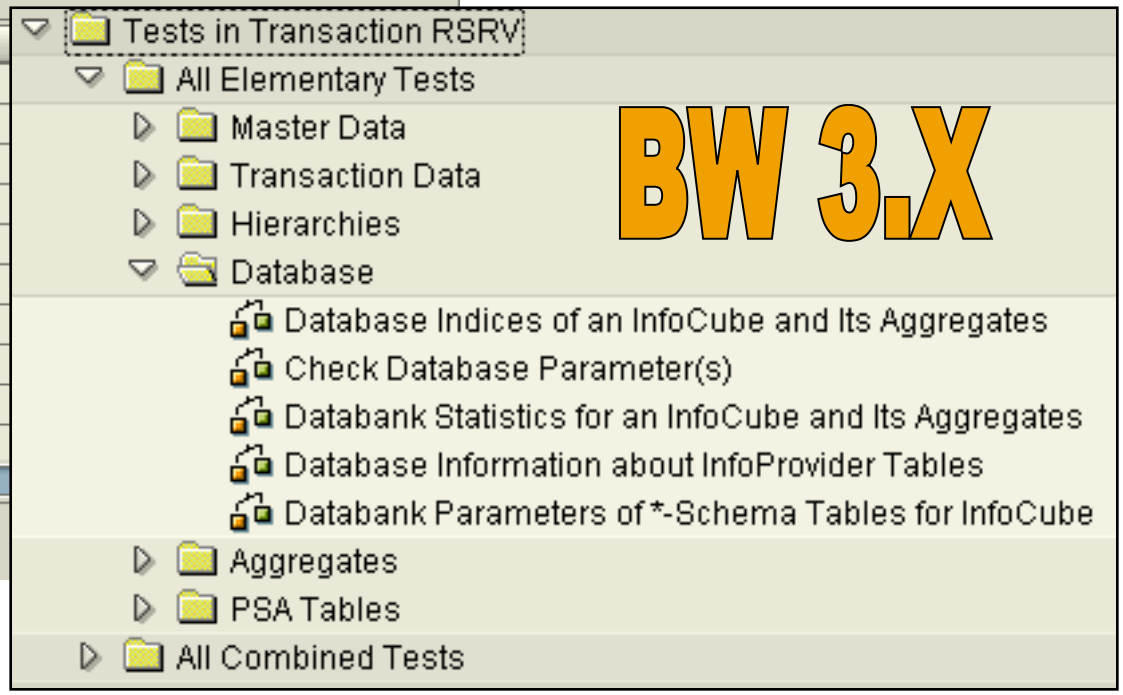

Note the checks for parameter settings, DB statistics, and Indices. The index analysis can detect most degenerated indices, and the repair tool can repair these indices.

Analysis Tools: Workload Monitor (ST03N)

Scenario

- General Query performance is bad. Find out the queries with the worst performance and try to optimize them.

Useful Features

- Expert Mode
- BW System Load → Analysis of table RSDDSTAT

Check following parameter values:

- Check queries with highest runtimes and check where most time has been consumed
  - OLAP init
  - DB
  - OLAP
  - Frontend
  - Master data time

Using Analysis Tools: ST03N - Drill Down to Query

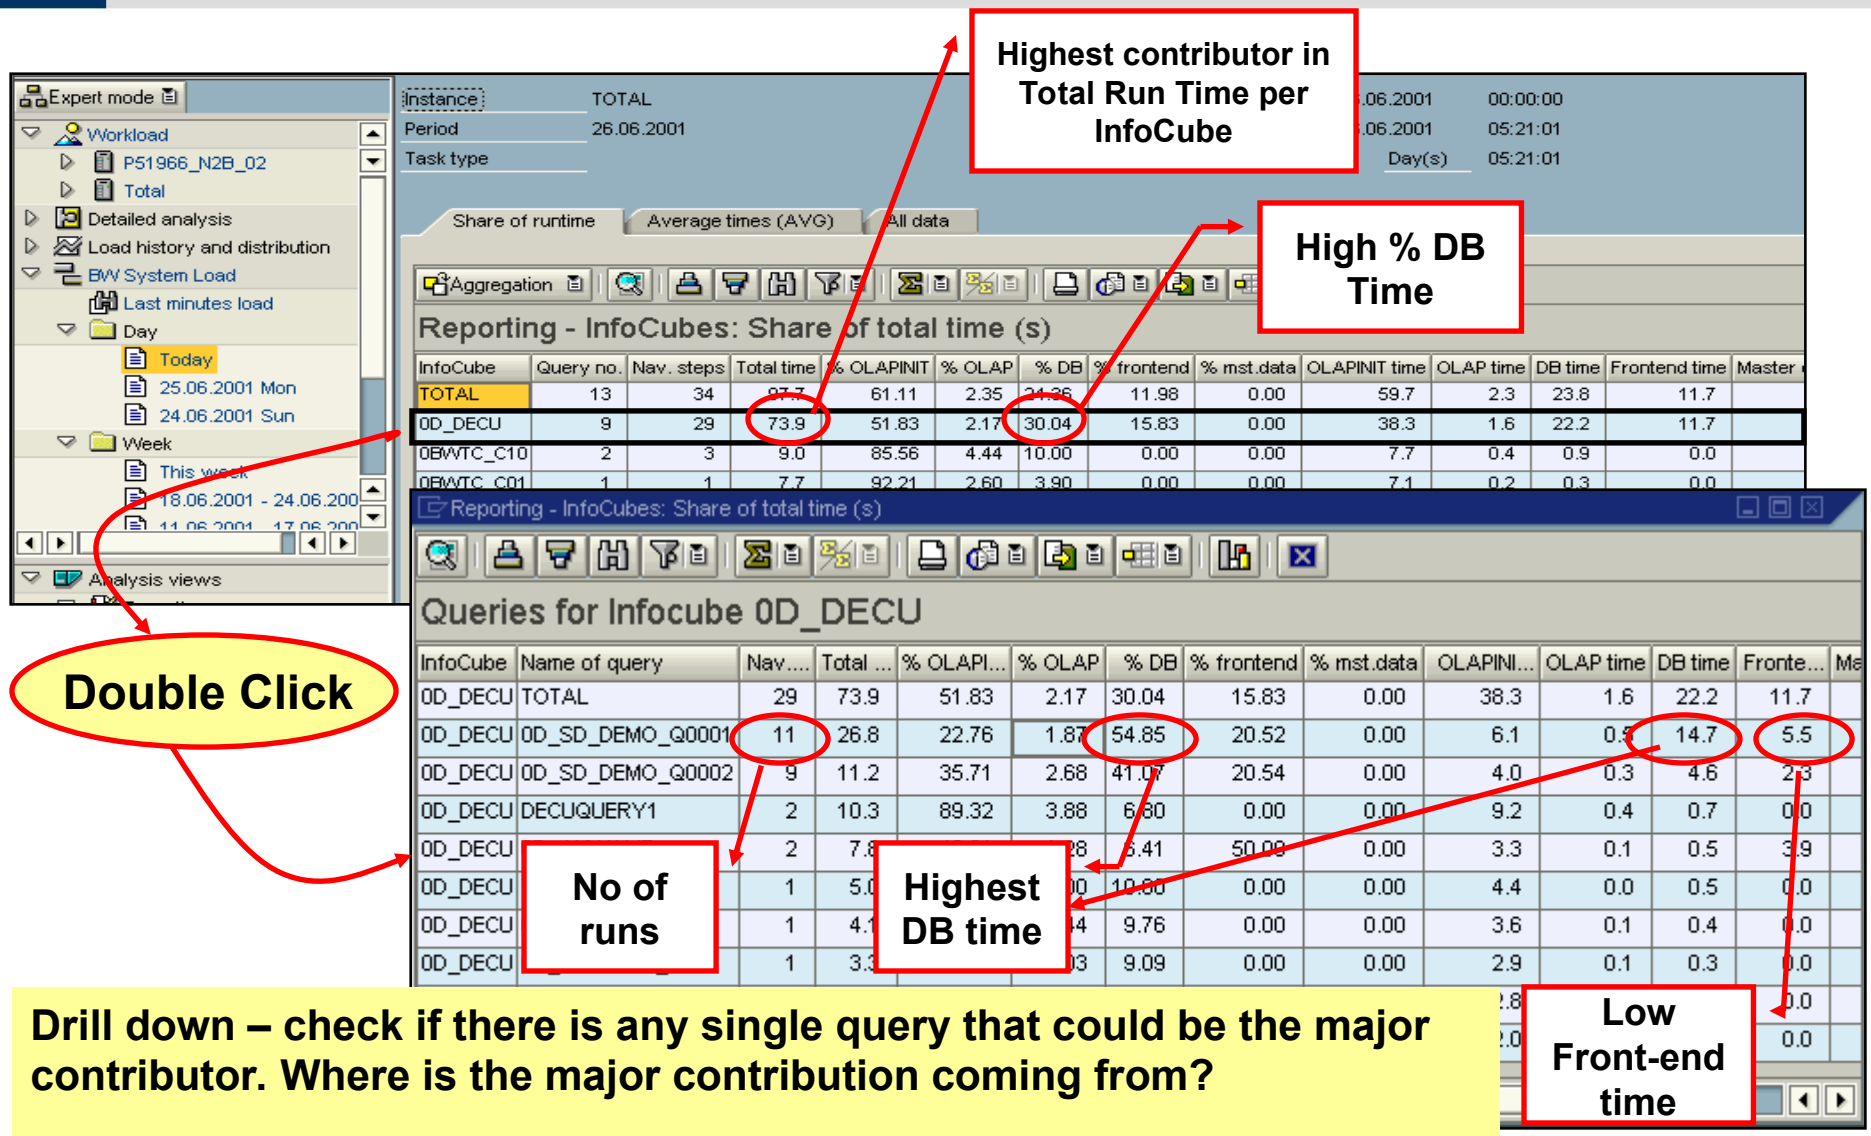

Drill down – check if there is any single query that could be the major contributor. Where is the major contribution coming from?

Example shows an active query, high DB time, with low OLAP & front-end time.

Overview: Reporting Performance Analysis Tools

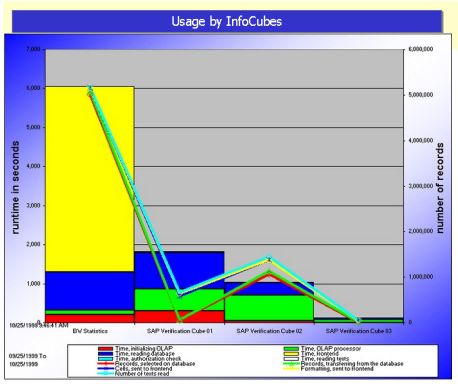

BW Statistics

Table RSDDSTAT

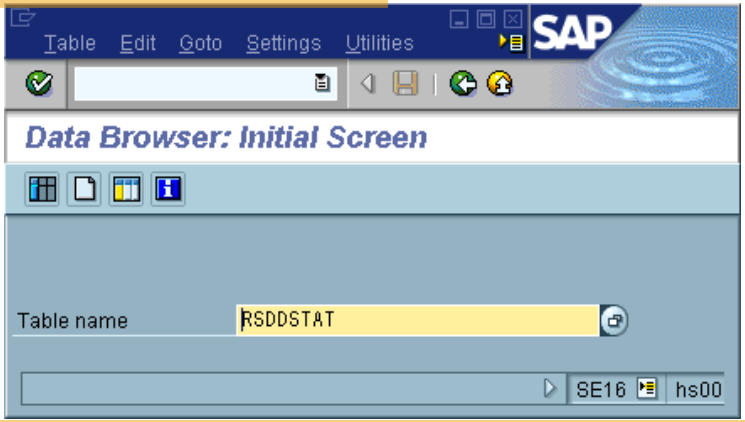

Queries of BW STATISTICS

Using table RSDDSTAT as InfoSource

BW Workload Analysis – ST03N

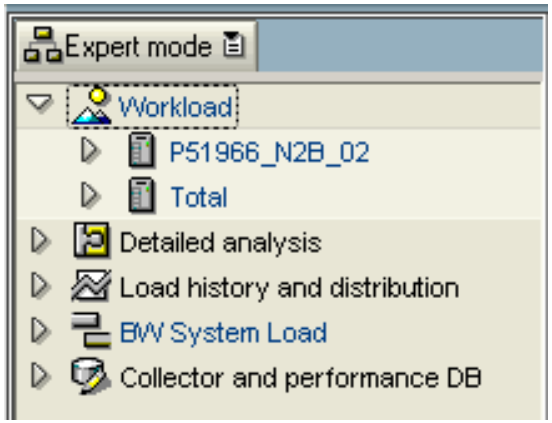

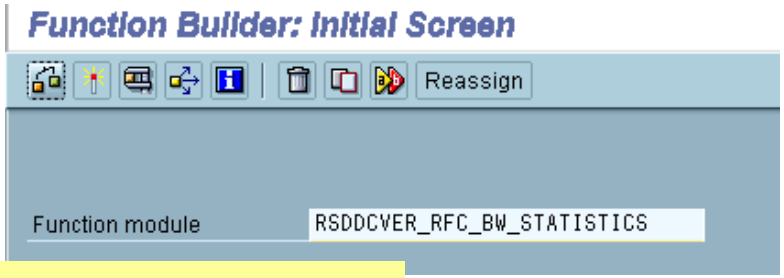

Collecting information from table RSDDSTAT

Function module RSDDCVER_RFC_BW_STATISTICS

Enable Capture of Statistical Performance Data

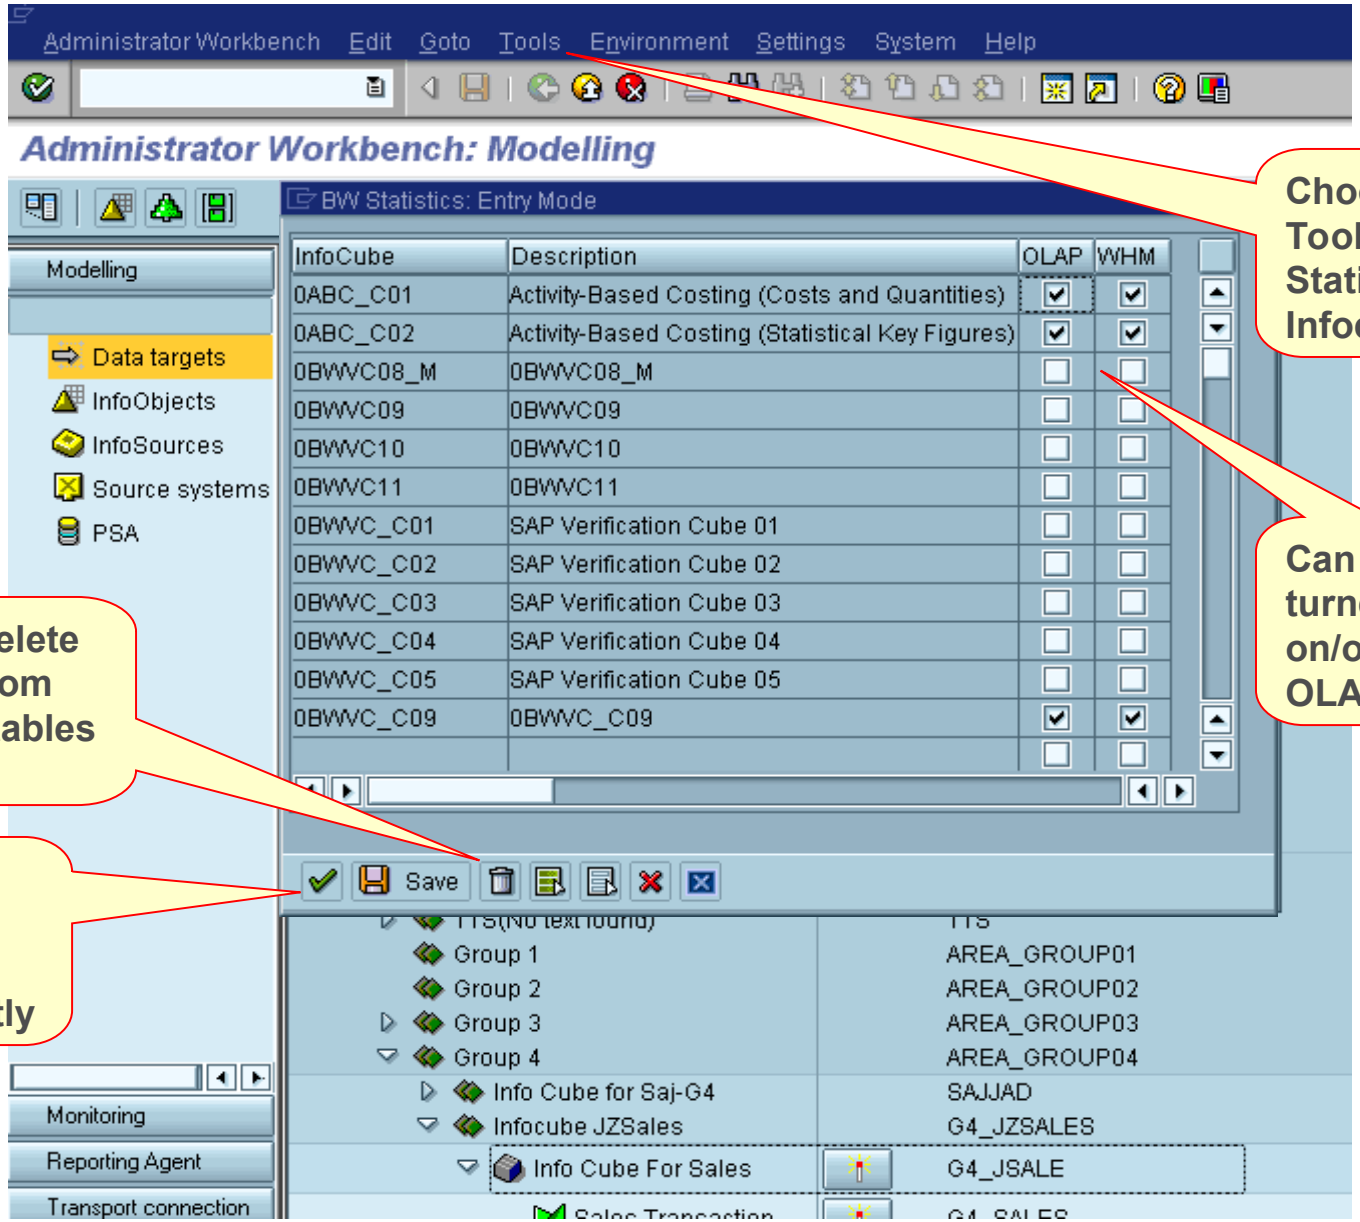

BW Statistics: Facts

SAP-delivered tools for analyzing and tuning system performance

Delivered Business Content beginning with BW 1.2A

- Redesigned for BW 2.0A

Usage of tools requires Business Content installation

Content Delivery Consists of:

- MultiCube
- InfoCubes
- Queries
- Workbook
- InfoSources
- Transfer Structures, Transfer Rules
- Update Rules
- InfoPackages, InfoPackage Groups
- Role

BW Statistics: Facts (continued)

Maintenance requires scheduling data loads and monitoring

- Delta load functionality delivered
- Transaction Data
- Master Data

Can be utilized when building aggregates

- System can “propose” aggregates from BW stats data

Object metadata of BW Statistics is transportable

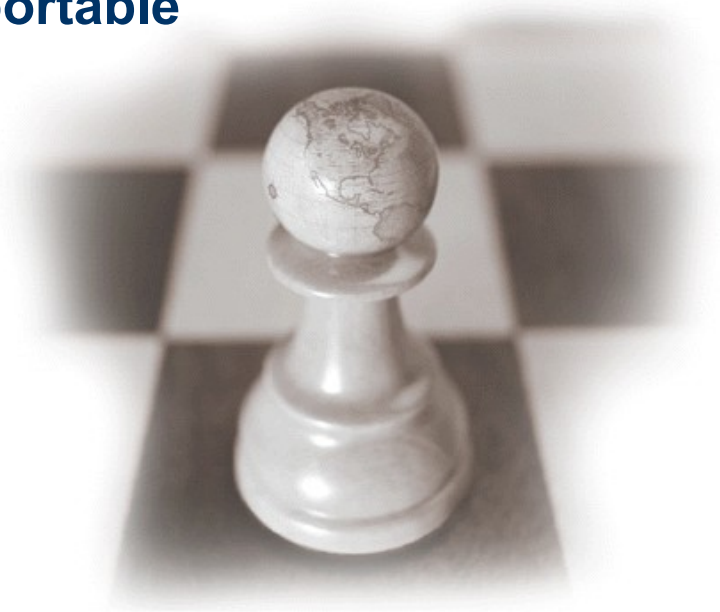

Data Load Performance Concepts

Reporting Performance Analysis Tools

Reporting Performance Analysis Techniques

Aggregates

Reporting Performance Tuning Concepts

Analyzing Reporting Performance Issues

Building Aggregates:

- Primary technique for tuning reporting performance
- Analyze to determine if building aggregates will help

Different Starting Points:

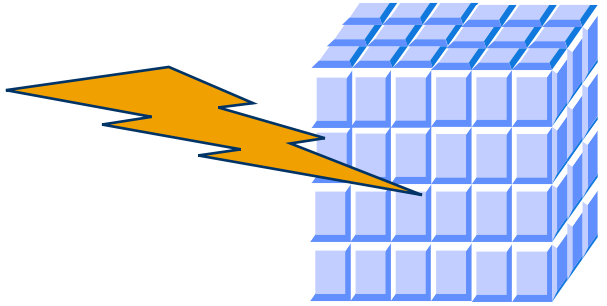

Analyze specific InfoProviders

OR

Find problem queries, then analyze them

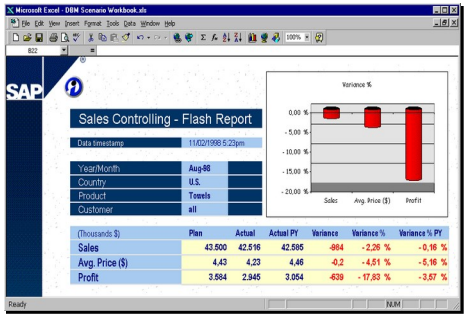

Analyzing Reporting Performance Issues

Approach: Find InfoCubes of concern & consider aggregates

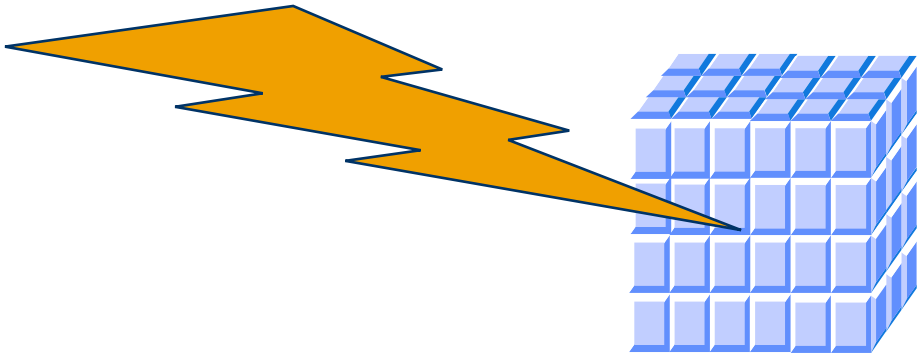

BW Statistics or ST03N: Find InfoCubes for Analysis

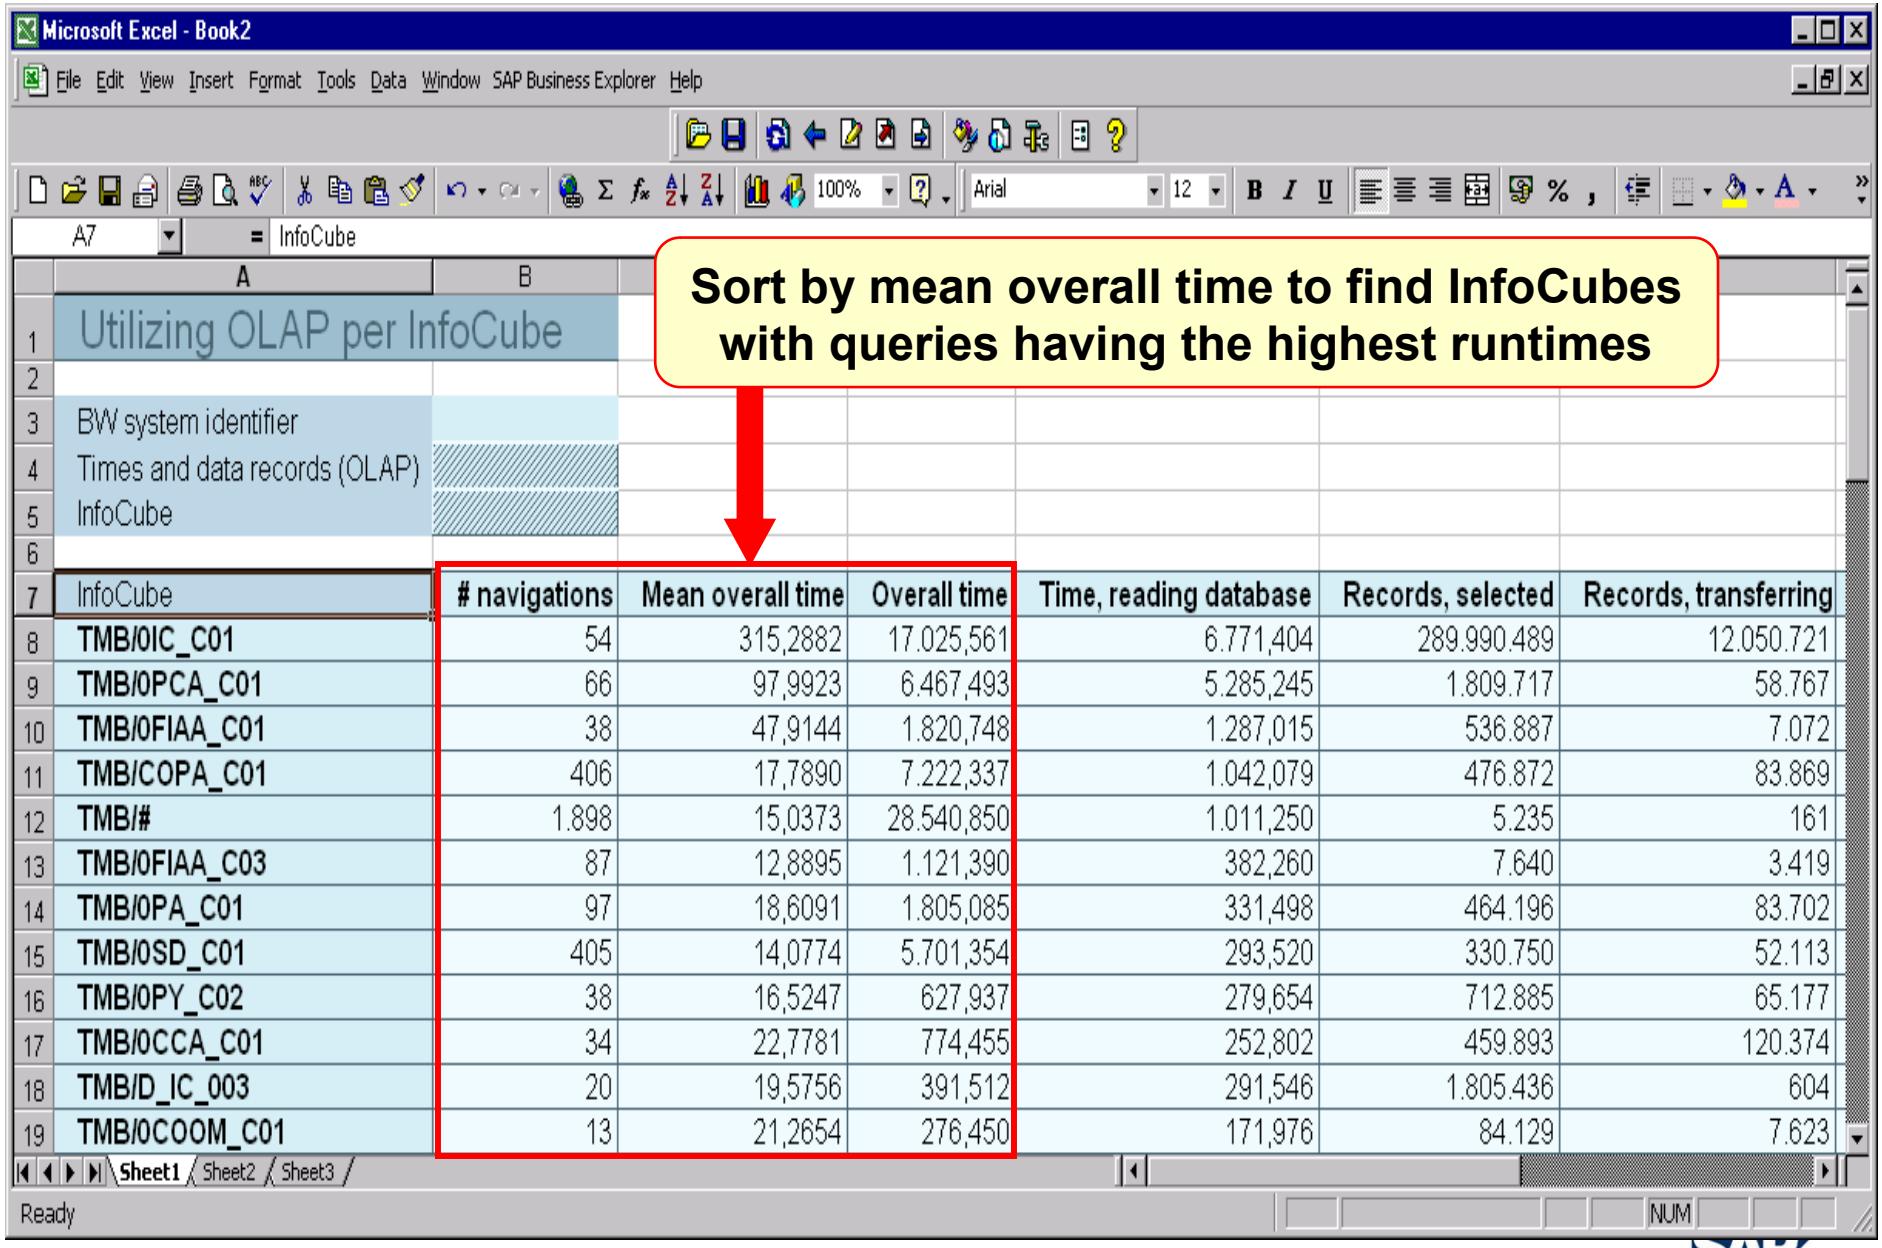

Utilizing OLAP per InfoCube

BW system identifier
Times and data records (OLAP)
InfoCube

InfoCube	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
TMB/0IC_C01	54	315,2882	17.025,561	6.771,404	289.990.489	12.050.721
TMB/0PCA_C01	66	97,9923	6.467,493	5.285,245	1.809.717	58.767
TMB/0FIAA_C01	38	47,9144	1.820,748	1.287,015	536.887	7.072
TMB/COPA_C01	406	17,7890	7.222,337	1.042,079	476.872	83.869
TMB/#	1.898	15,0373	28.540,850	1.011,250	5.235	161
TMB/0FIAA_C03	87	12,8895	1.121,390	382,260	7.640	3.419
TMB/0PA_C01	97	18,6091	1.805,085	331,498	464.196	83.702
TMB/0SD_C01	405	14,0774	5.701,354	293,520	330.750	52.113
TMB/0PY_C02	38	16,5247	627,937	279,654	712.885	65.177
TMB/0CCA_C01	34	22,7781	774,455	252,802	459.893	120.374
TMB/D_IC_003	20	19,5756	391,512	291,546	1.805.436	604
TMB/0COOM_C01	13	21,2654	276,450	171,976	84.129	7.623

Indicators That Aggregates Will Help

Look for many more records selected than necessary!

One Key Indicator:

Ratio records selected / records transferred > 10

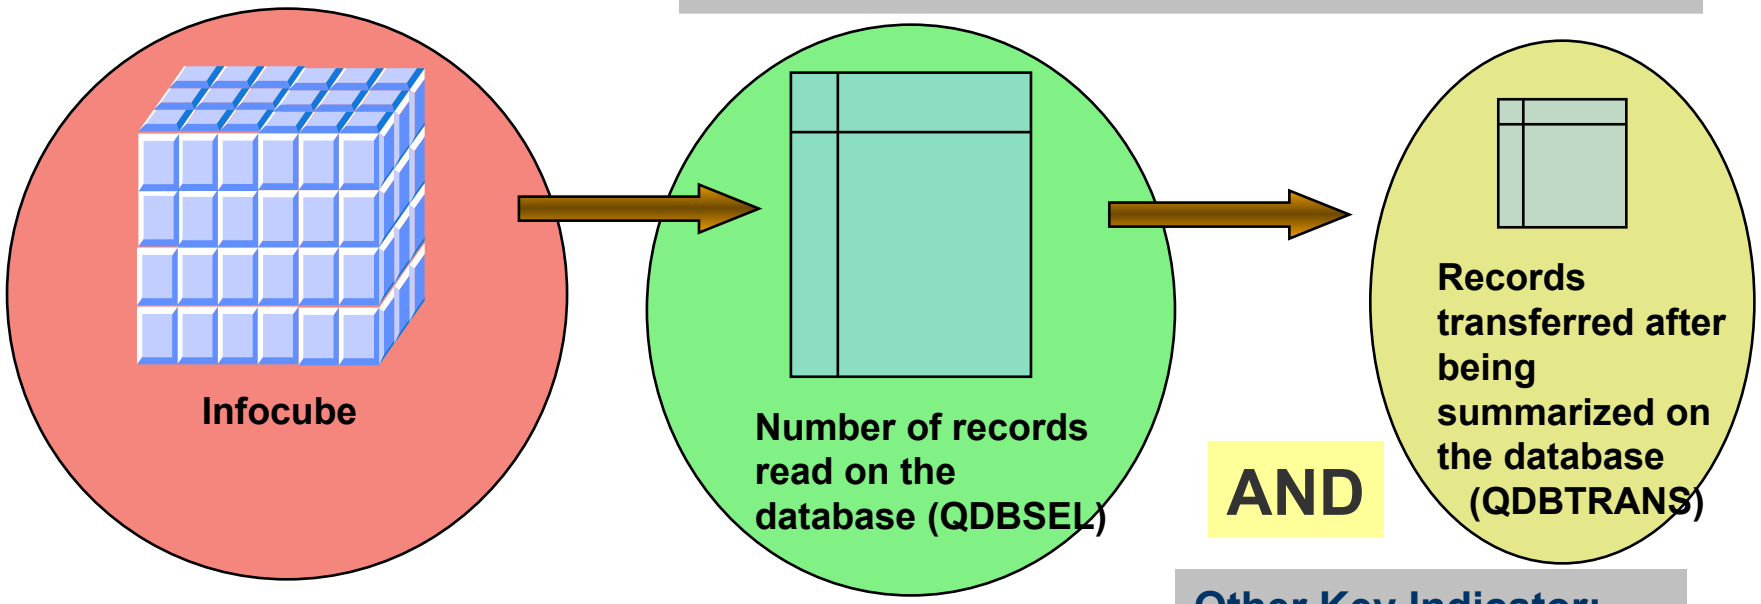

Other Key Indicator:

database time for queries > 30% of total query runtime

Total query runtime (QRUNTIMECATEGORY)

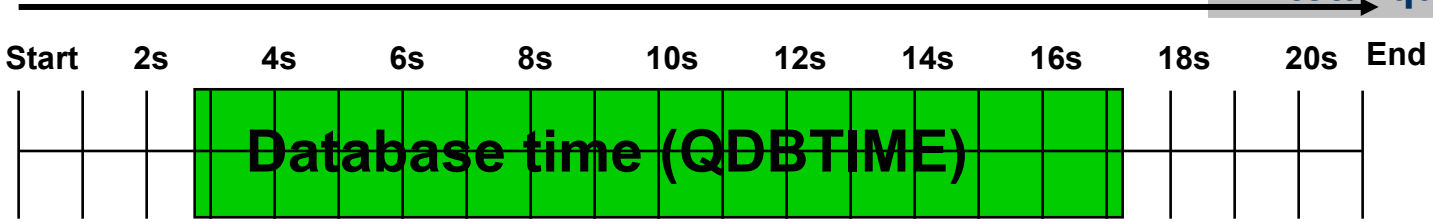

Analysis of BW Statistics at InfoCube Level (1)

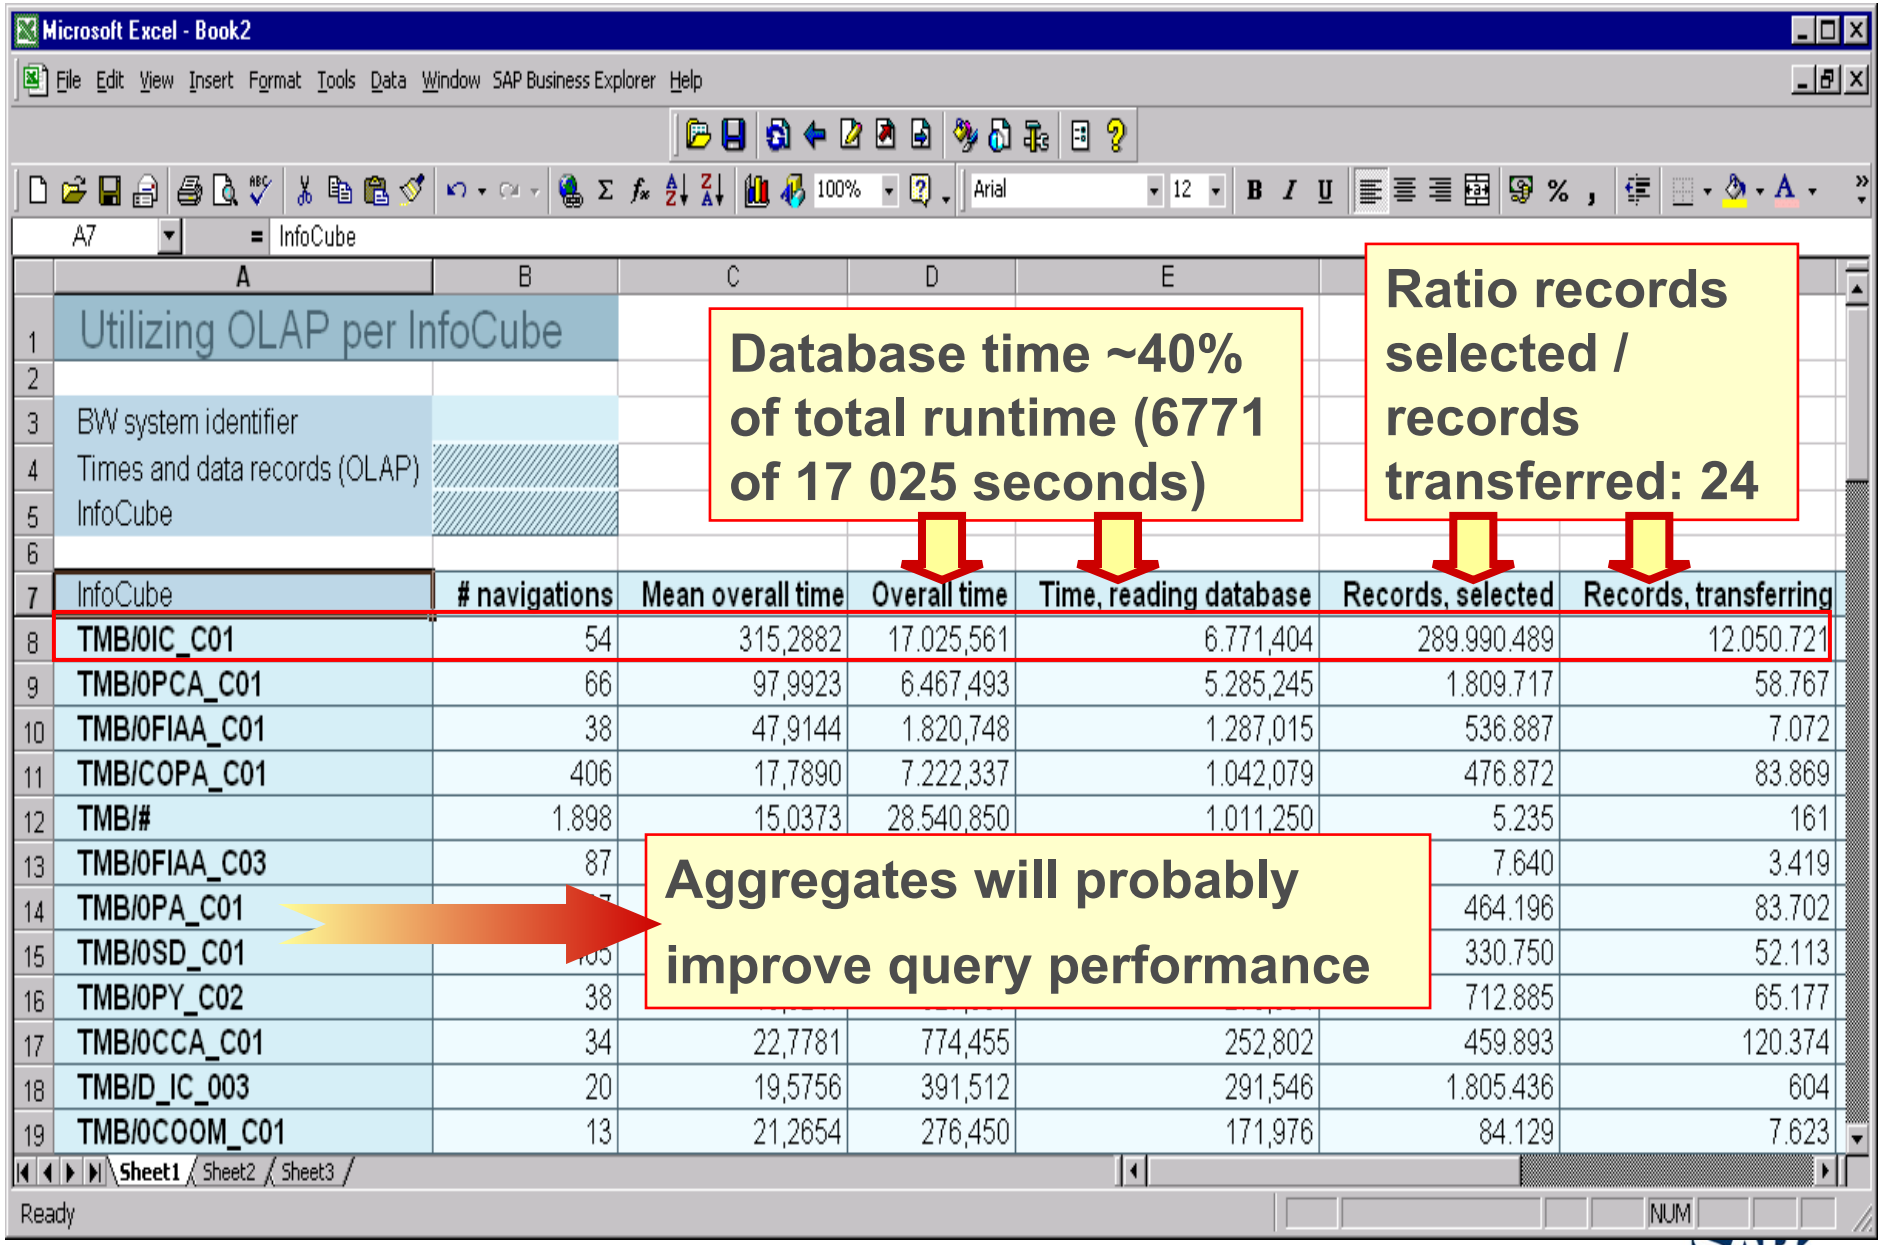

InfoCube	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
TMB/0IC_C01	54	315,2882	17.025,561	6.771,404	289.990.489	12.050.721
TMB/0PCA_C01	66	97,9923	6.467,493	5.285,245	1.809.717	58.767
TMB/0FIAA_C01	38	47,9144	1.820,748	1.287,015	536.887	7.072
TMB/COPA_C01	406	17,7890	7.222,337	1.042,079	476.872	83.869
TMB/#	1.898	15,0373	28.540,850	1.011,250	5.235	161
TMB/0FIAA_C03	87				7.640	3.419
TMB/0PA_C01					464.196	83.702
TMB/0SD_C01					330.750	52.113
TMB/0PY_C02	38				712.885	65.177
TMB/0CCA_C01	34	22,7781	774,455	252,802	459.893	120.374
TMB/D_IC_003	20	19,5756	391,512	291,546	1.805.436	604
TMB/0COOM_C01	13	21,2654	276,450	171,976	84.129	7.623

Analysis of BW Statistics at InfoCube Level (2)

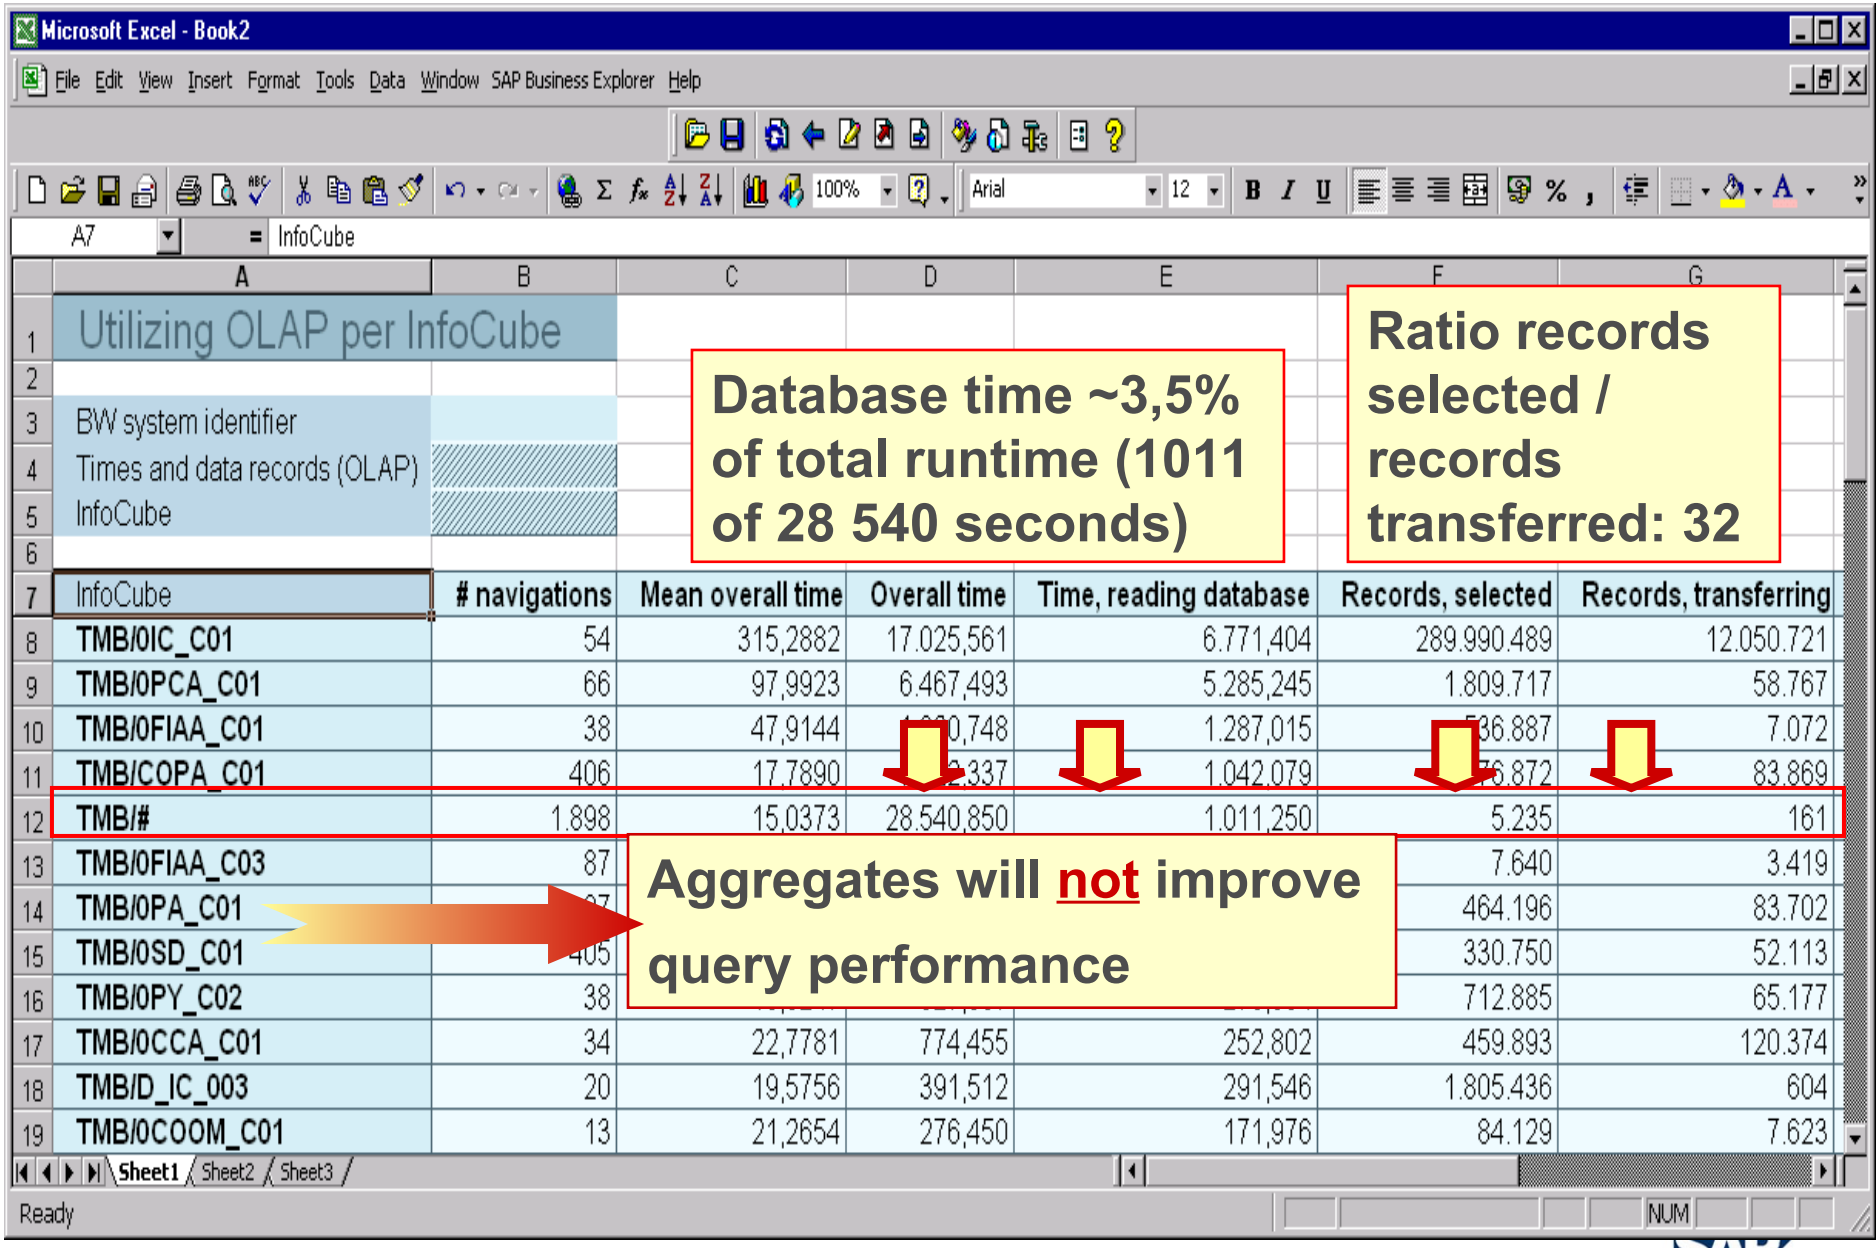

InfoCube	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
TMB/0IC_C01	54	315,2882	17.025,561	6.771,404	289.990,489	12.050.721
TMB/0PCA_C01	66	97,9923	6.467,493	5.285,245	1.809.717	58.767
TMB/0FIAA_C01	38	47,9144	[illegible]	1.287,015	[illegible]	7.072
TMB/COPA_C01	406	17,7890	[illegible]	1.042,079	[illegible]	83.869
TMB/#	1.898	15,0373	28.540,850	1.011,250	5.235	161
TMB/0FIAA_C03	87	[illegible]	[illegible]	[illegible]	7.640	3.419
TMB/0PA_C01	[illegible]	[illegible]	[illegible]	[illegible]	464.196	83.702
TMB/0SD_C01	405	[illegible]	[illegible]	[illegible]	330.750	52.113
TMB/0PY_C02	38	[illegible]	[illegible]	[illegible]	712.885	65.177
TMB/0CCA_C01	34	22,7781	774,455	252,802	459.893	120.374
TMB/D_IC_003	20	19,5756	391,512	291,546	1.805.436	604
TMB/0COOM_C01	13	21,2654	276,450	171,976	84.129	7.623

Analysis of BW Statistics at InfoCube Level (3)

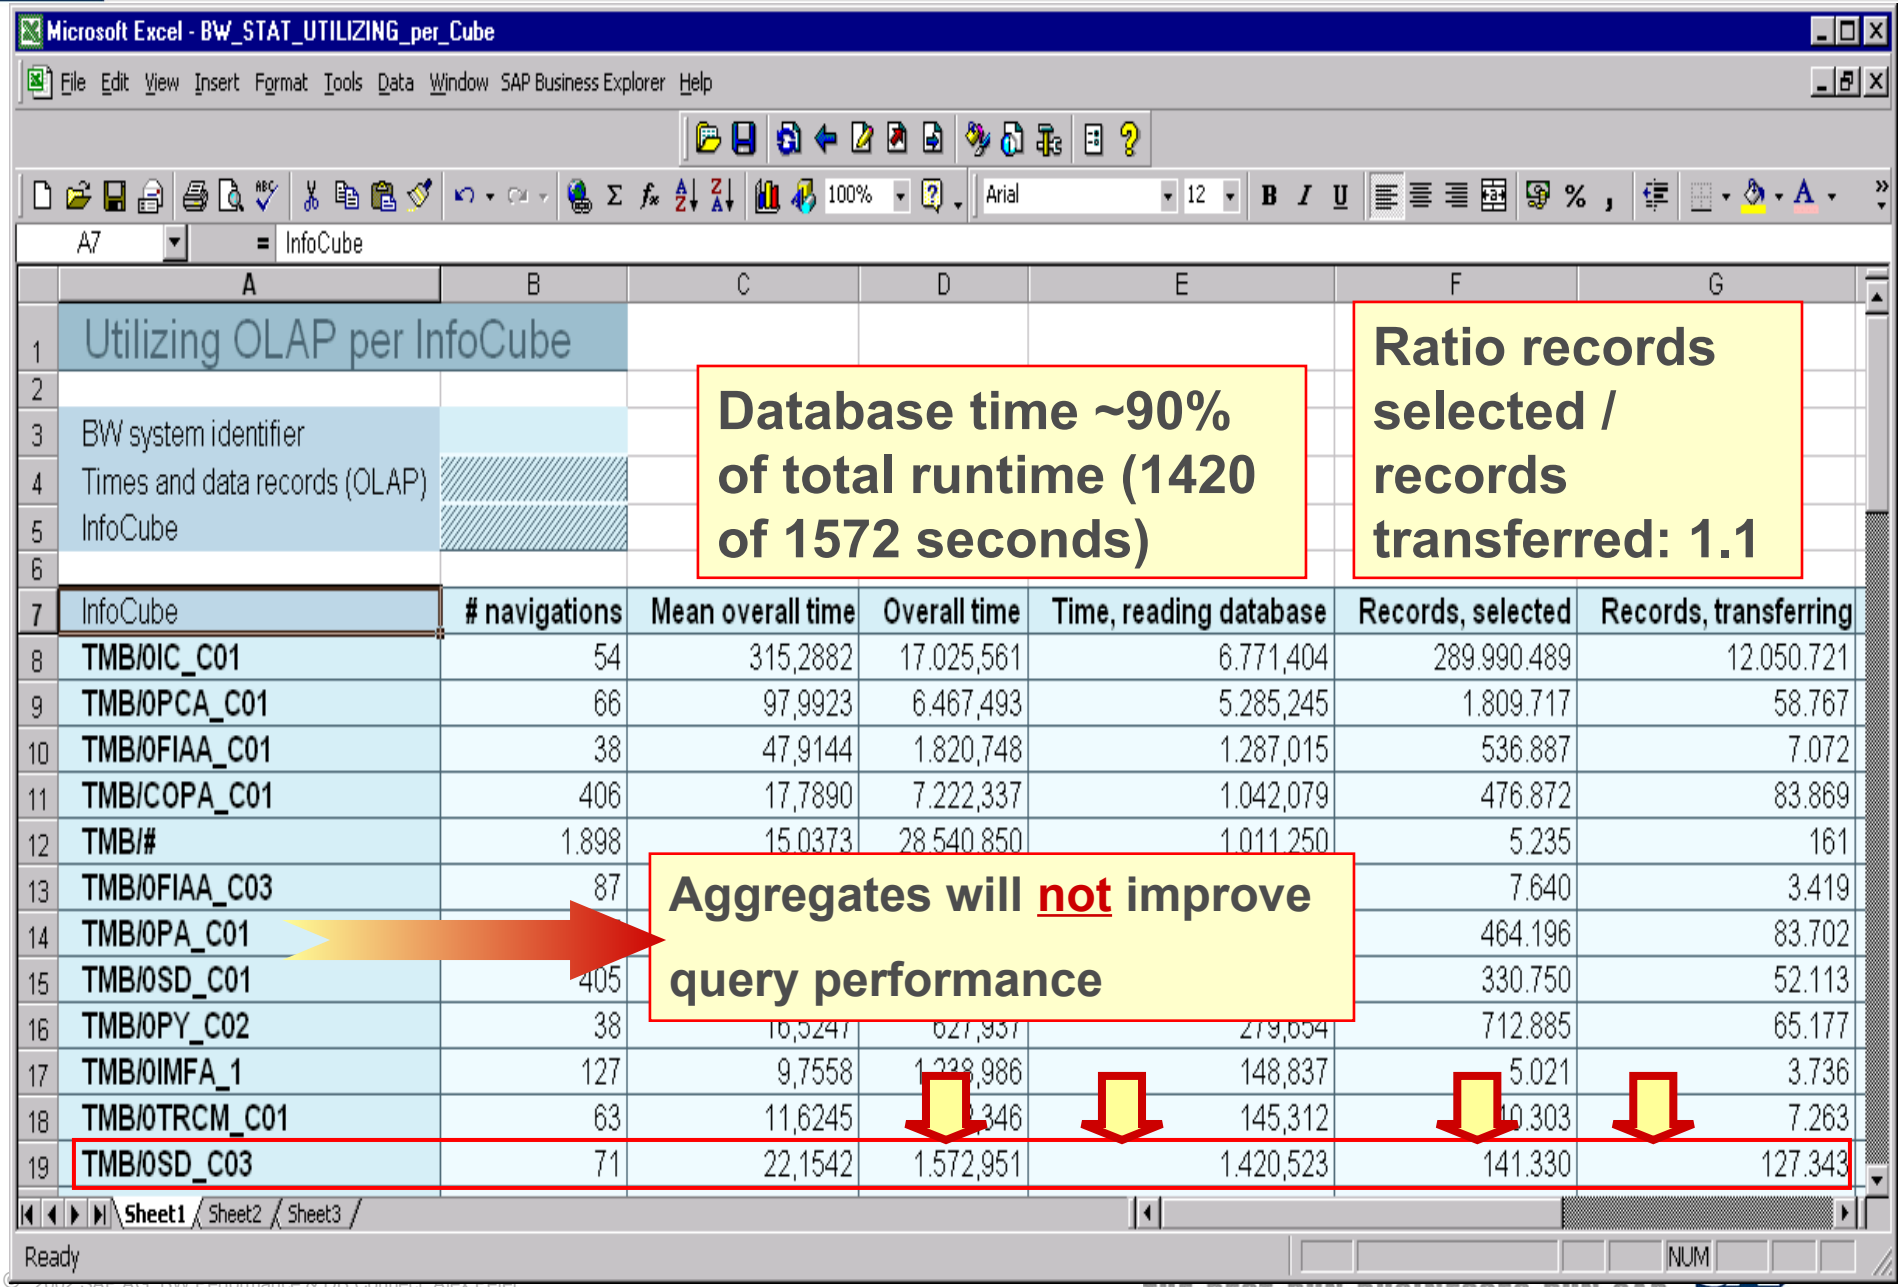

Microsoft Excel - BW_STAT_UTILIZING_per_Cube

File Edit View Insert Format Tools Data Window SAP Business Explorer Help

A7 = InfoCube

Utilizing OLAP per InfoCube

BW system identifier
Times and data records (OLAP)
InfoCube

Database time ~90% of total runtime (1420 of 1572 seconds)

Ratio records selected / records transferred: 1.1

InfoCube	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
TMB/0IC_C01	54	315,2882	17.025,561	6.771,404	289.990,489	12.050,721
TMB/0PCA_C01	66	97,9923	6.467,493	5.285,245	1.809,717	58,767
TMB/0FIAA_C01	38	47,9144	1.820,748	1.287,015	536,887	7,072
TMB/COPA_C01	406	17,7890	7.222,337	1.042,079	476,872	83,869
TMB/#	1.898	15,0373	28.540,850	1.011,250	5,235	161
TMB/0FIAA_C03	87				7,640	3,419
TMB/0PA_C01					464,196	83,702
TMB/0SD_C01	405				330,750	52,113
TMB/0PY_C02	38	16,5247	627,937	279,654	712,885	65,177
TMB/0IMFA_1	127	9,7558	1.238,986	148,837	5,021	3,736
TMB/0TRCM_C01	63	11,6245	732,346	145,312	10,303	7,263
TMB/0SD_C03	71	22,1542	1.572,951	1.420,523	141,330	127,343

Aggregates will not improve query performance

Sheet1 Sheet2 Sheet3

Ready NUM

ST03N - Workload By InfoCube

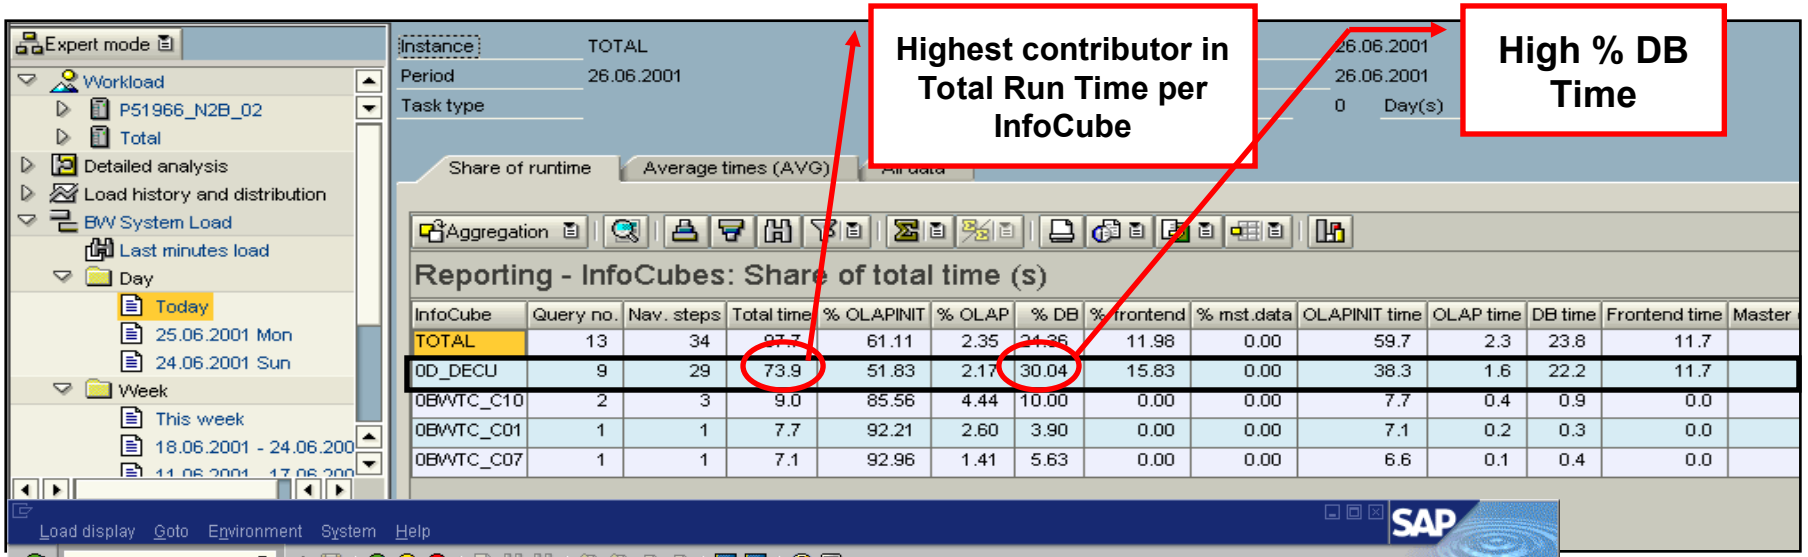

Highest contributor in Total Run Time per InfoCube

High % DB Time

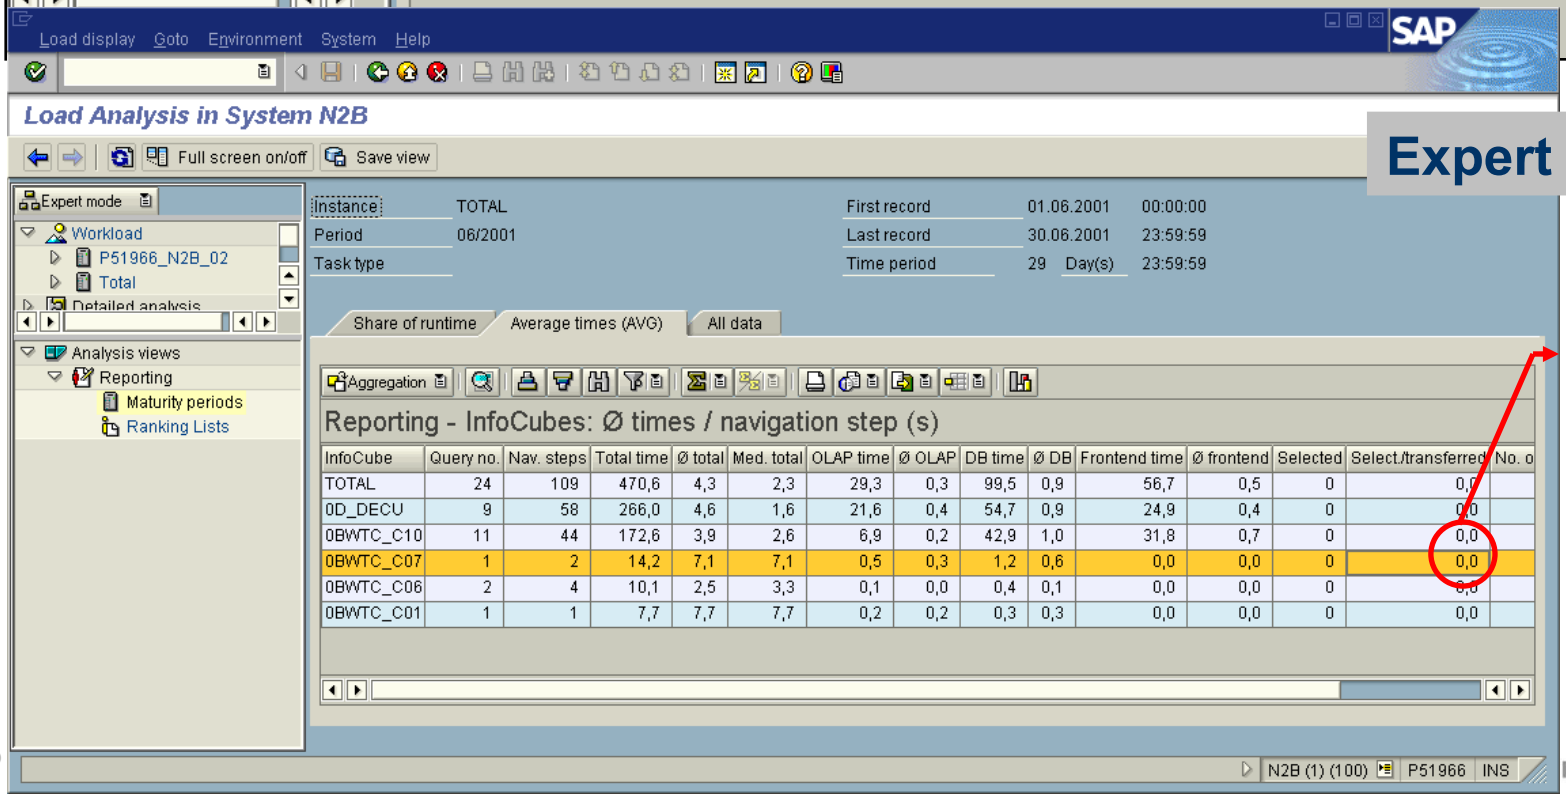

Expert Mode!!

High ratio: Selected / transferred records

Analyzing Reporting Performance Issues

Approach: Find queries of concern, then analyze them

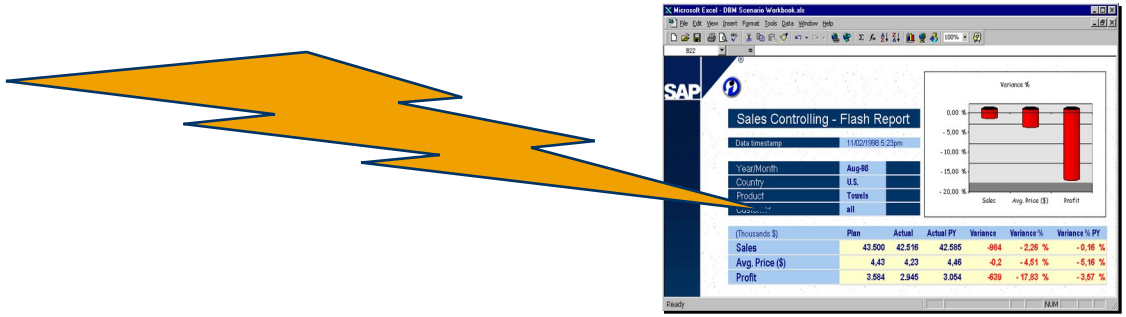

BW Stats or ST03N: Find Queries for Analysis

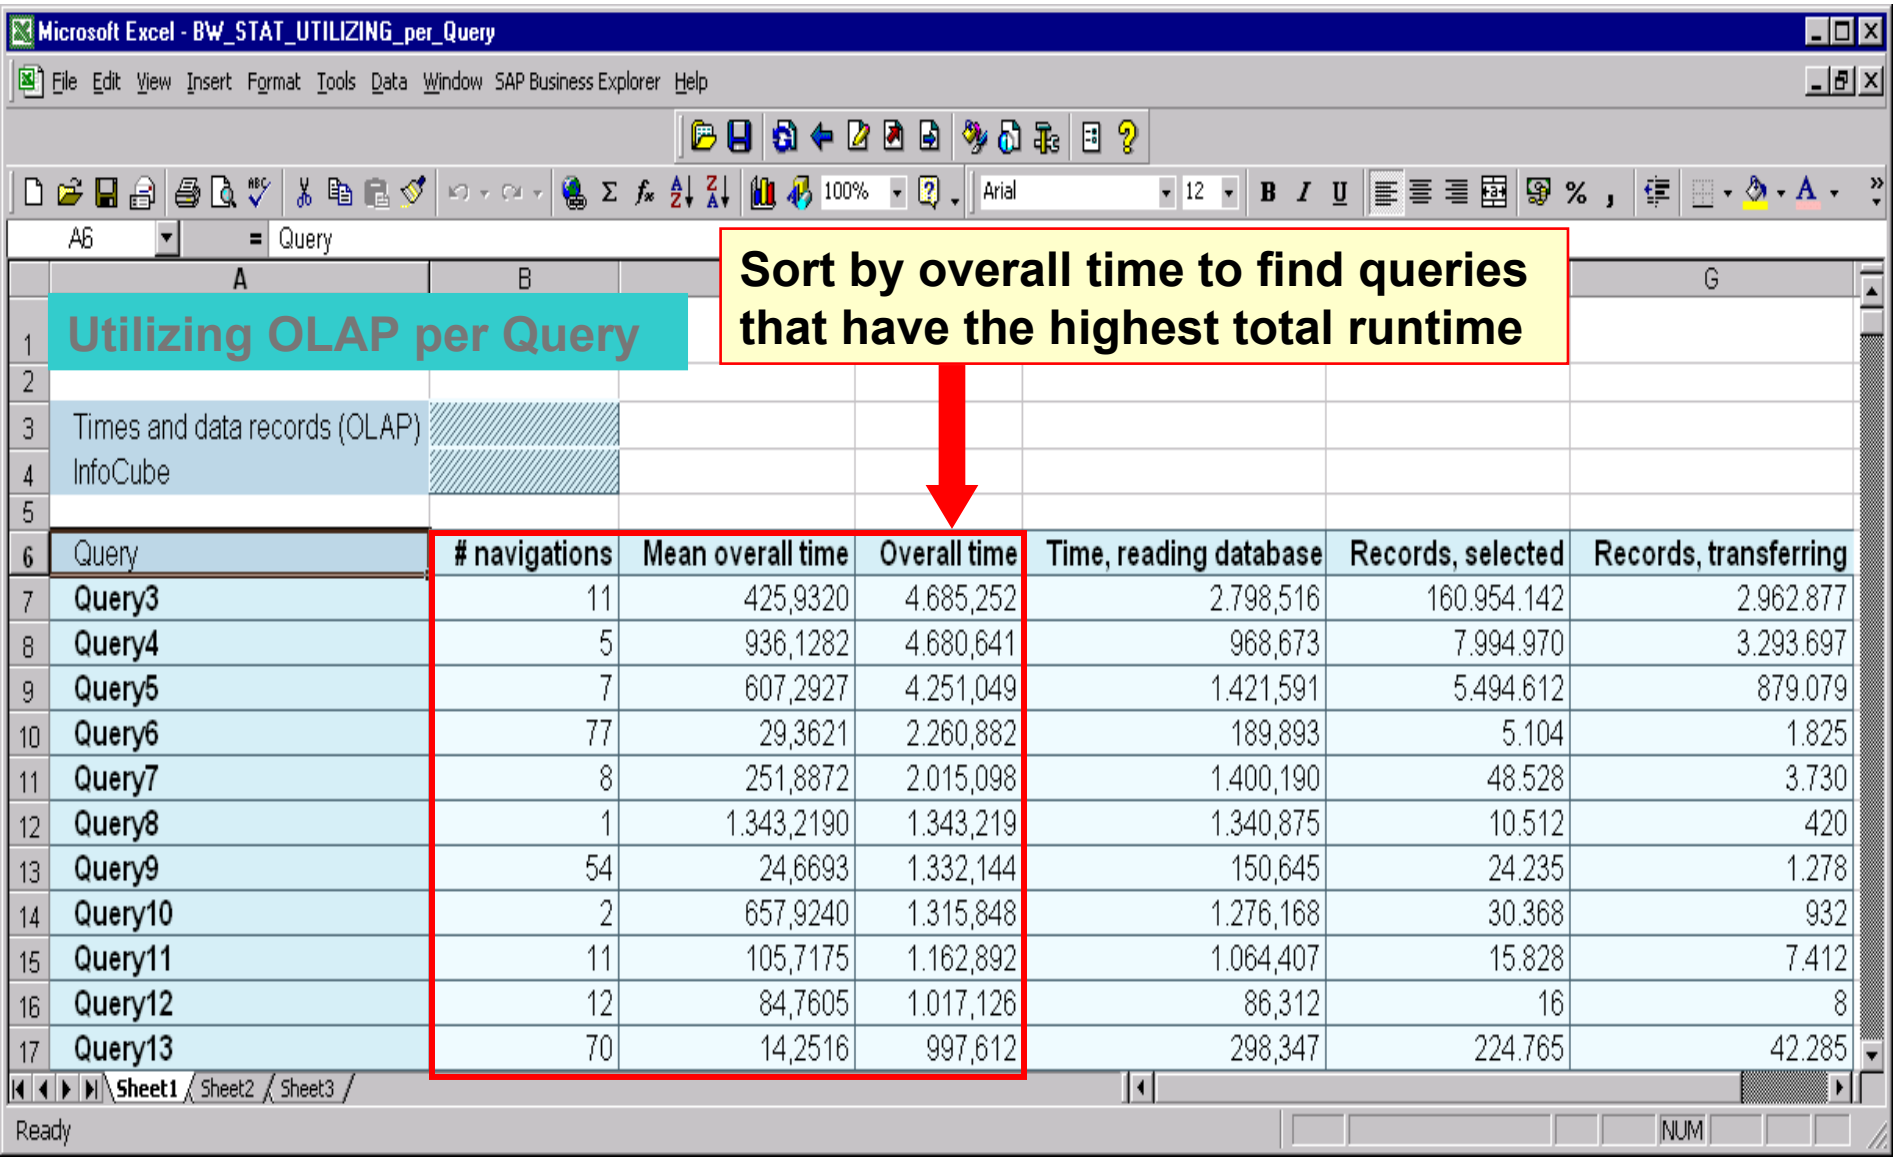

Query	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
Query3	11	425,9320	4.685,252	2.798,516	160.954.142	2.962.877
Query4	5	936,1282	4.680,641	968,673	7.994.970	3.293.697
Query5	7	607,2927	4.251,049	1.421,591	5.494.612	879.079
Query6	77	29,3621	2.260,882	189,893	5.104	1.825
Query7	8	251,8872	2.015,098	1.400,190	48.528	3.730
Query8	1	1.343,2190	1.343,219	1.340,875	10.512	420
Query9	54	24,6693	1.332,144	150,645	24.235	1.278
Query10	2	657,9240	1.315,848	1.276,168	30.368	932
Query11	11	105,7175	1.162,892	1.064,407	15.828	7.412
Query12	12	84,7605	1.017,126	86,312	16	8
Query13	70	14,2516	997,612	298,347	224.765	42.285

Indicators That Aggregates Will Help

Look for many more records selected than necessary!

One Key Indicator:

Ratio records selected / records transferred > 10

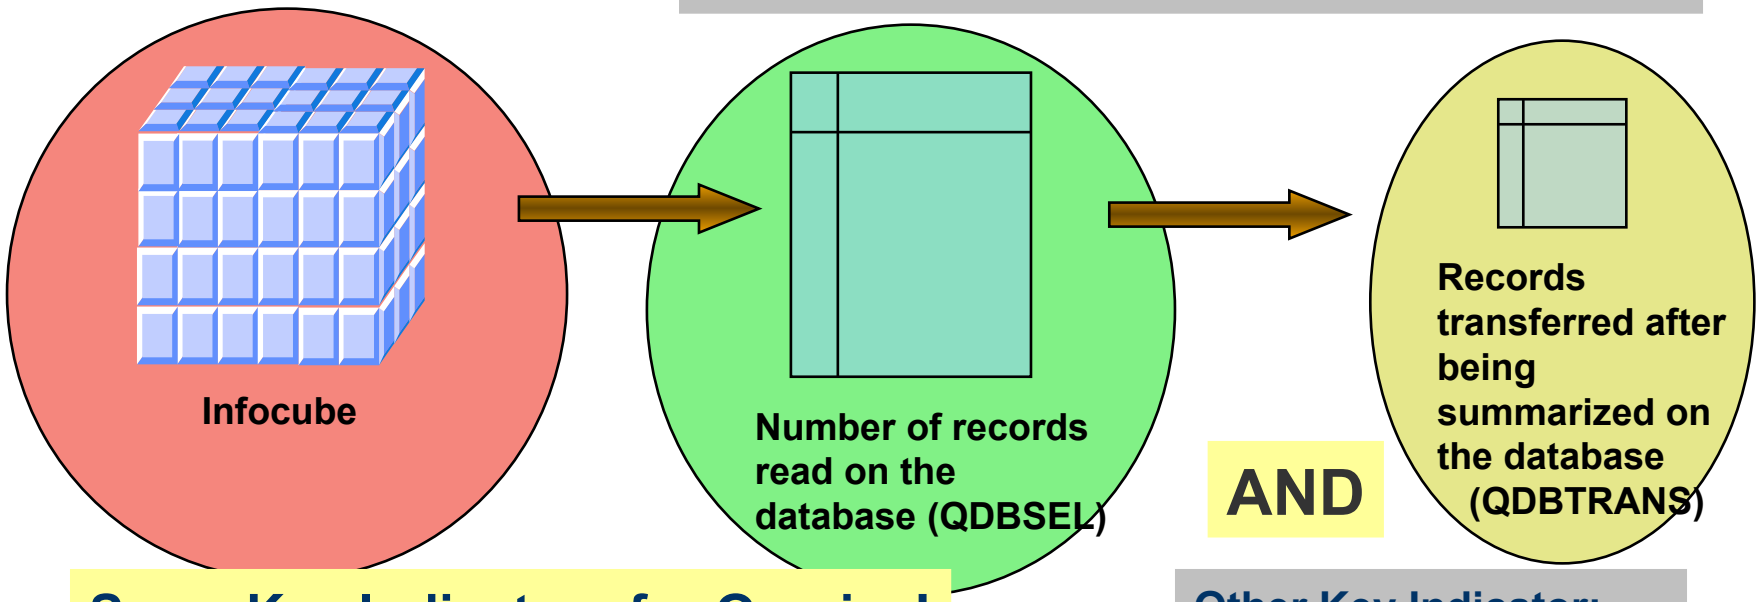

AND

Same Key Indicators for Queries!

Other Key Indicator:

database time for queries > 30% of total query runtime

Total query runtime (QRUNTIMECATEGORY)

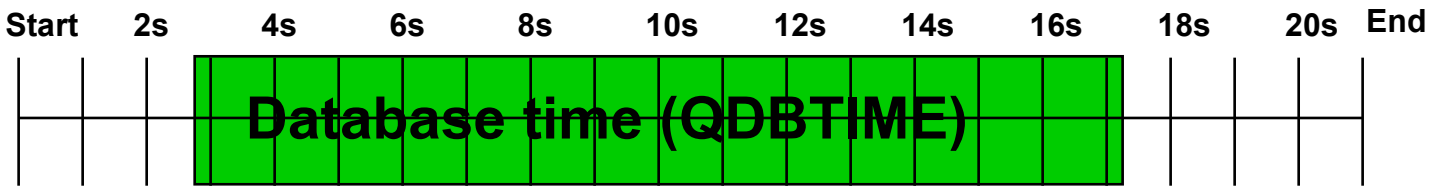

Analysis Techniques: Key Components of Runtime

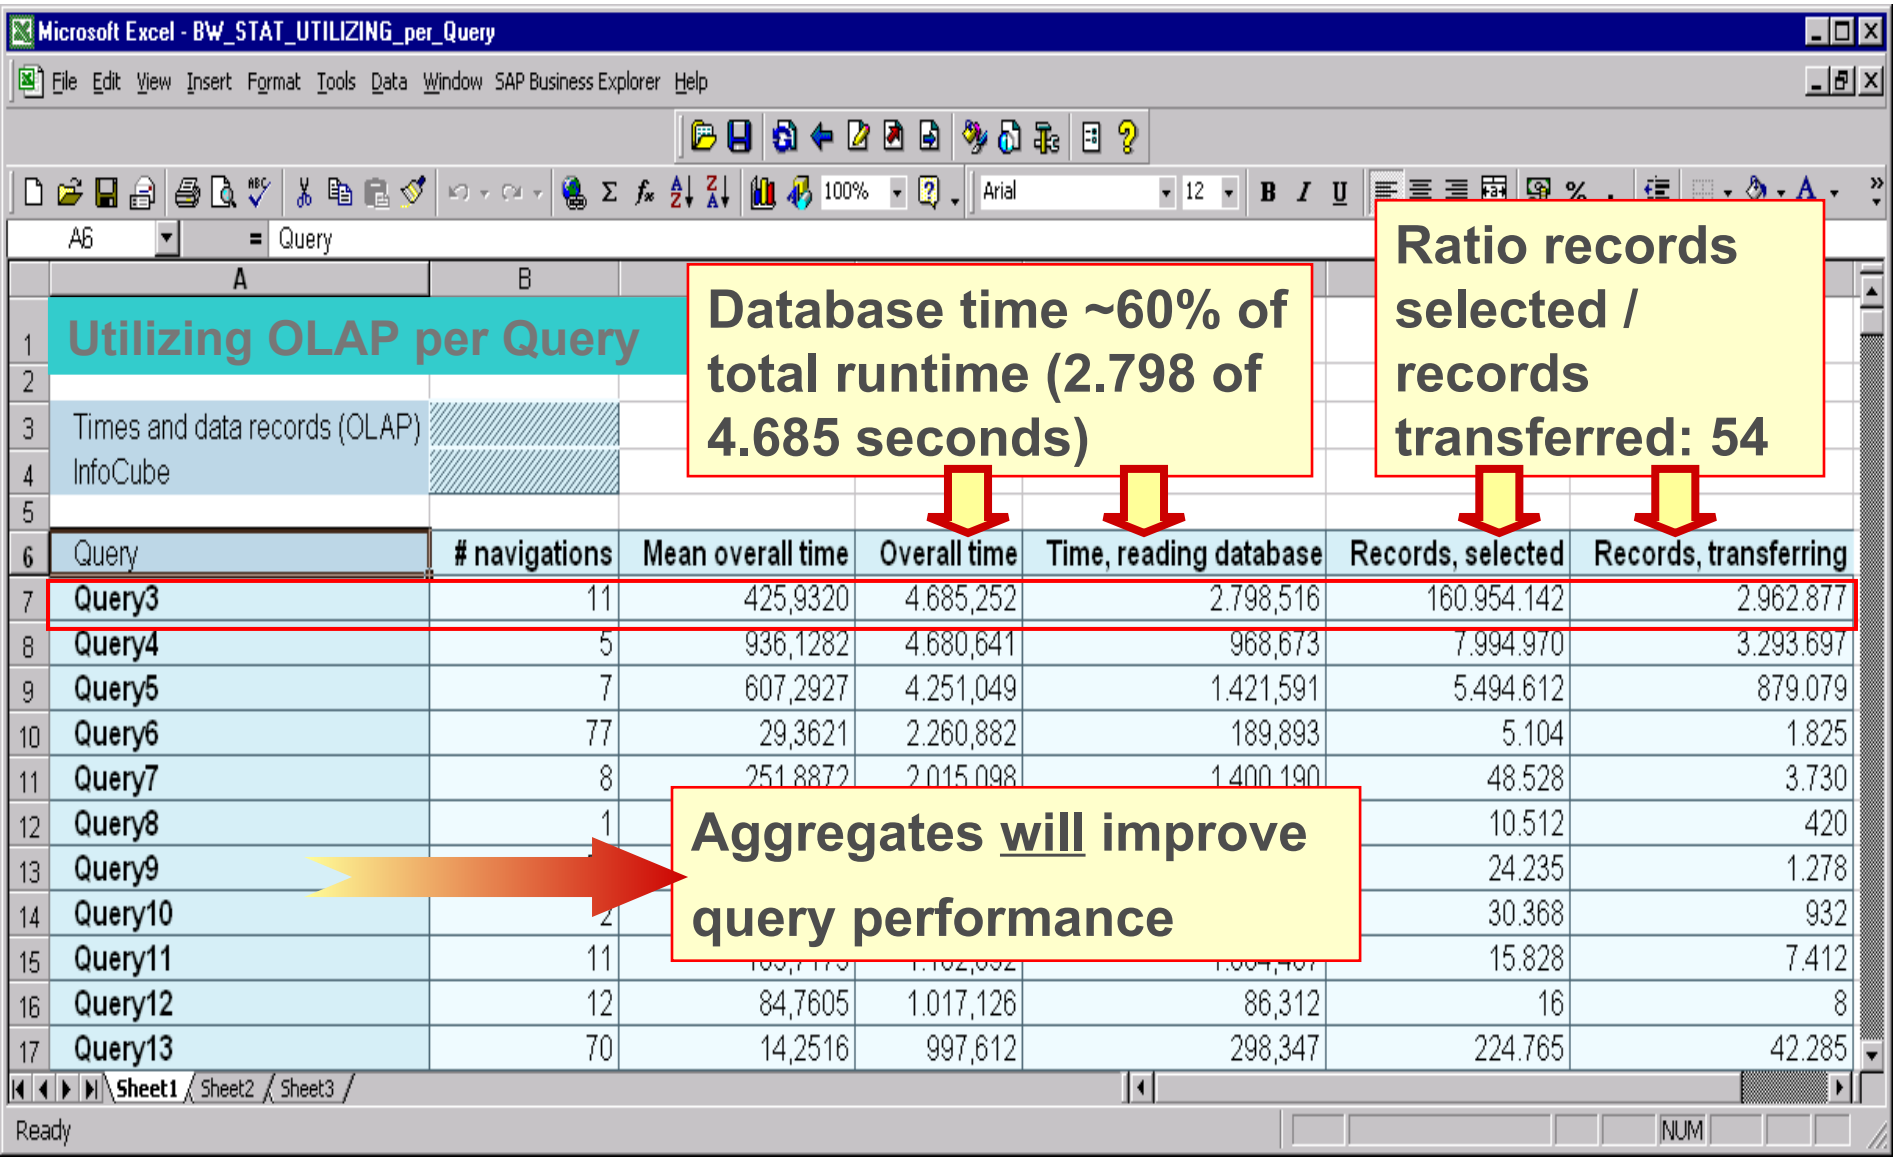

Query	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
Query3	11	425,9320	4.685,252	2.798,516	160.954.142	2.962.877
Query4	5	936,1282	4.680,641	968,673	7.994.970	3.293.697
Query5	7	607,2927	4.251,049	1.421,591	5.494.612	879.079
Query6	77	29,3621	2.260,882	189,893	5.104	1.825
Query7	8	251,8872	2.015,098	1.400,190	48.528	3.730
Query8	1				10.512	420
Query9					24.235	1.278
Query10	2				30.368	932
Query11	11				15.828	7.412
Query12	12	84,7605	1.017,126	86,312	16	8
Query13	70	14,2516	997,612	298,347	224.765	42.285

Analysis Techniques: Key Components of Runtime

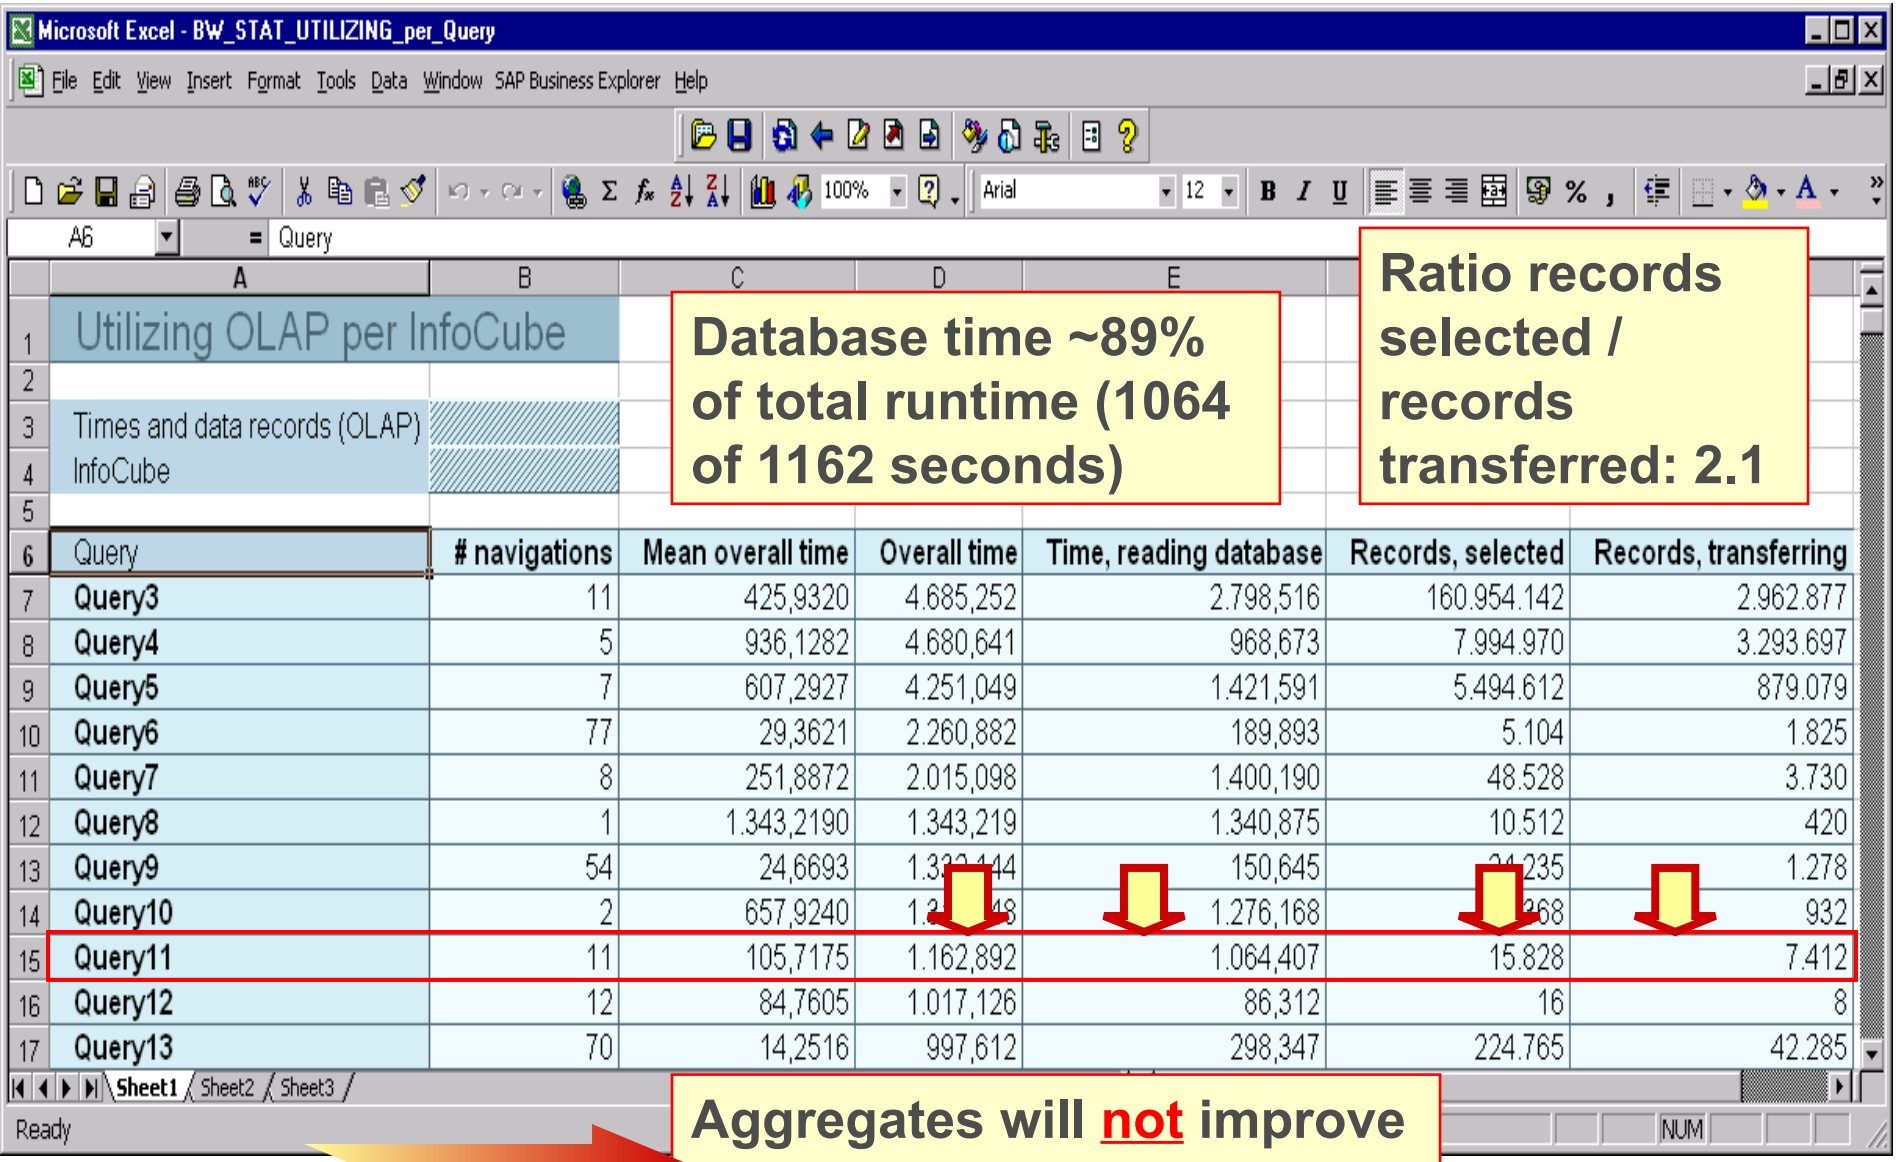

Query	# navigations	Mean overall time	Overall time	Time, reading database	Records, selected	Records, transferring
Query3	11	425,9320	4.685,252	2.798,516	160.954.142	2.962.877
Query4	5	936,1282	4.680,641	968,673	7.994.970	3.293.697
Query5	7	607,2927	4.251,049	1.421,591	5.494.612	879.079
Query6	77	29,3621	2.260,882	189,893	5.104	1.825
Query7	8	251,8872	2.015,098	1.400,190	48.528	3.730
Query8	1	1.343,2190	1.343,219	1.340,875	10.512	420
Query9	54	24,6693	1.3[illegible]	150,645	[illegible]235	1.278
Query10	2	657,9240	1.3[illegible]	1.276,168	[illegible]268	932
Query11	11	105,7175	1.162,892	1.064,407	15.828	7.412
Query12	12	84,7605	1.017,126	86,312	16	8
Query13	70	14,2516	997,612	298,347	224.765	42.285

Aggregates will **not** improve query performance

Query Monitor – Transaction RSRT

Scenario

- The performance of one specific query is bad. Debug the query to analyze possible bottlenecks

Useful Features

- Execute query in debug mode
  - SAP statistics
  - Explain Plan
  - Display found aggregates
- OLAP query cache monitor

Check following things:

- Are aggregates used?
- Is the OLAP query cache used? (DB time = 0)
- Are the statistics up-to-date?

Query Monitor – Transaction RSRT

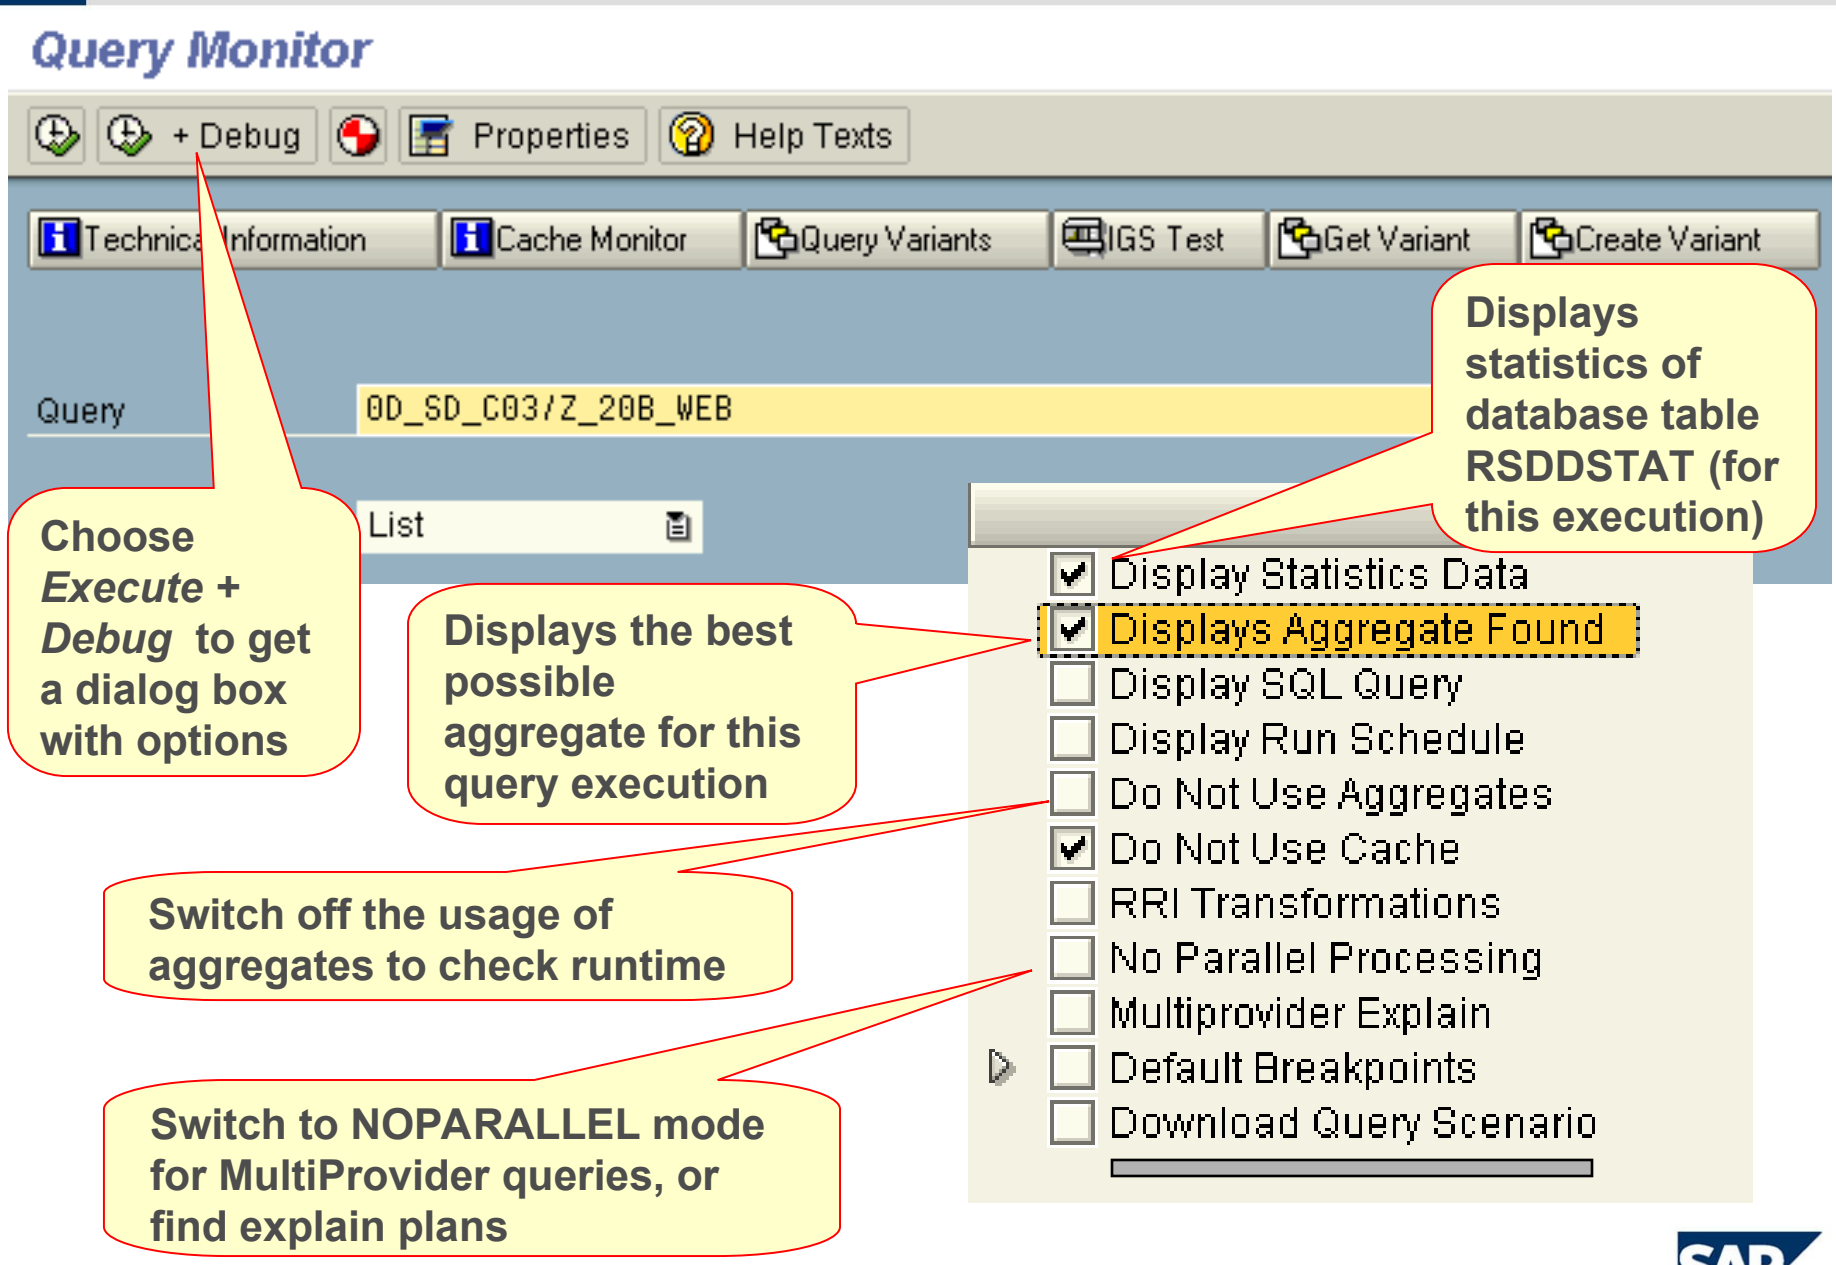

Best Possible Aggregate (1)

Number of database accesses (if all best possible aggregates are available)

Suggested aggregation type (*, F, H, blank)

Fixed value used (if aggregation F is used)

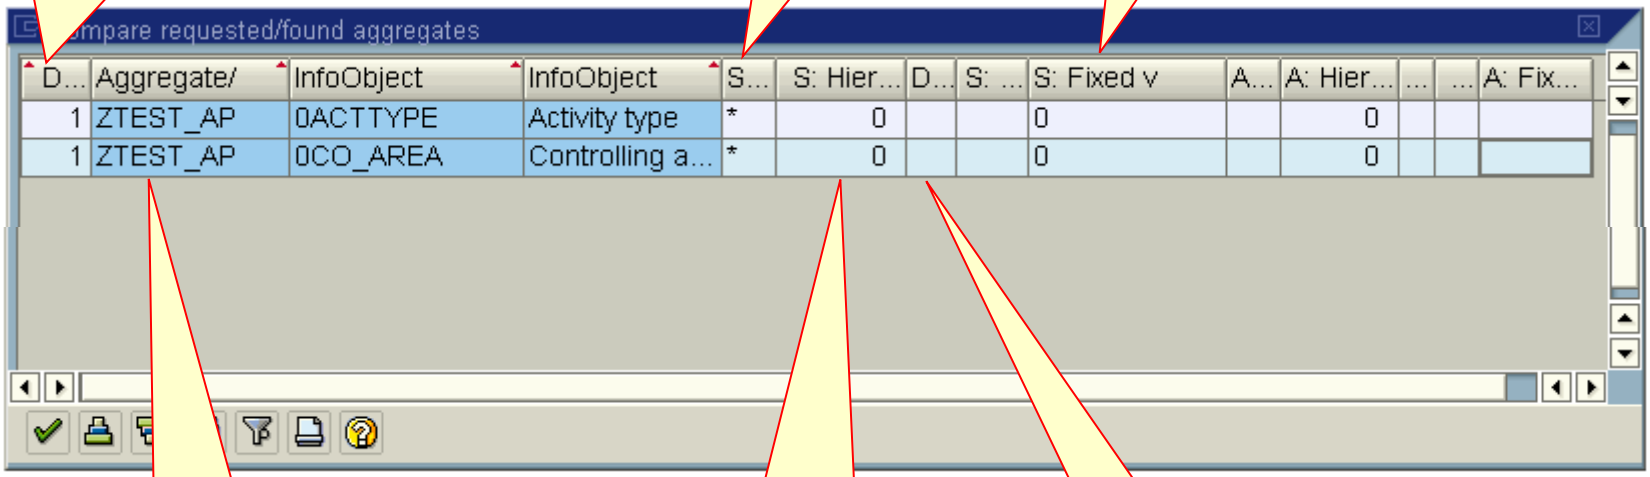

Structure to be read (InfoCube or aggregate)

Hierarchy used (if aggregation H is used)

Hierarchy level (if aggregation H is used)

Best Possible Aggregate

Structure to be read (InfoCube or aggregate)

Suggested aggregation type (*, F, H, or blank)

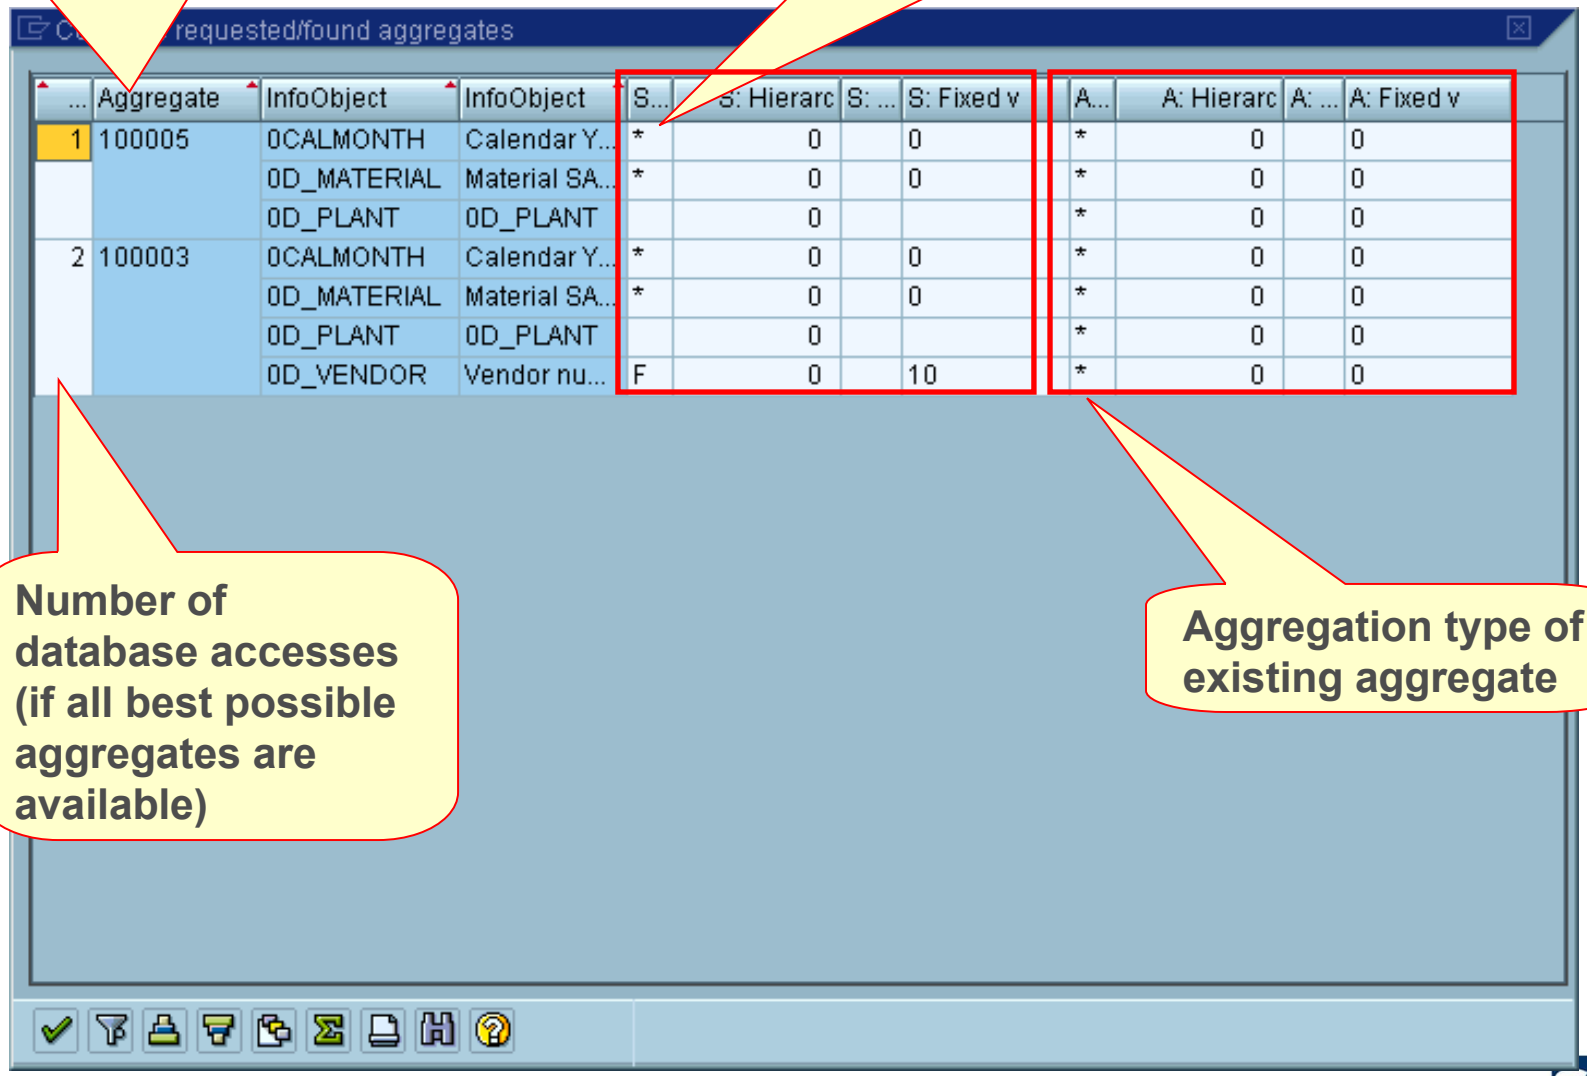

...	Aggregate	InfoObject	InfoObject	S...	S: Hierarc	S: ...	S: Fixed v	A...	A: Hierarc	A: ...	A: Fixed v
1	100005	0CALMONTH	Calendar Y...	*	0		0	*	0		0
		0D_MATERIAL	Material SA...	*	0		0	*	0		0
		0D_PLANT	0D_PLANT		0			*	0		0
2	100003	0CALMONTH	Calendar Y...	*	0		0	*	0		0
		0D_MATERIAL	Material SA...	*	0		0	*	0		0
		0D_PLANT	0D_PLANT		0			*	0		0
		0D_VENDOR	Vendor nu...	F	0		10	*	0		0

Number of database accesses (if all best possible aggregates are available)

Aggregation type of existing aggregate

RSRT Query Monitor: BW Statistics Data

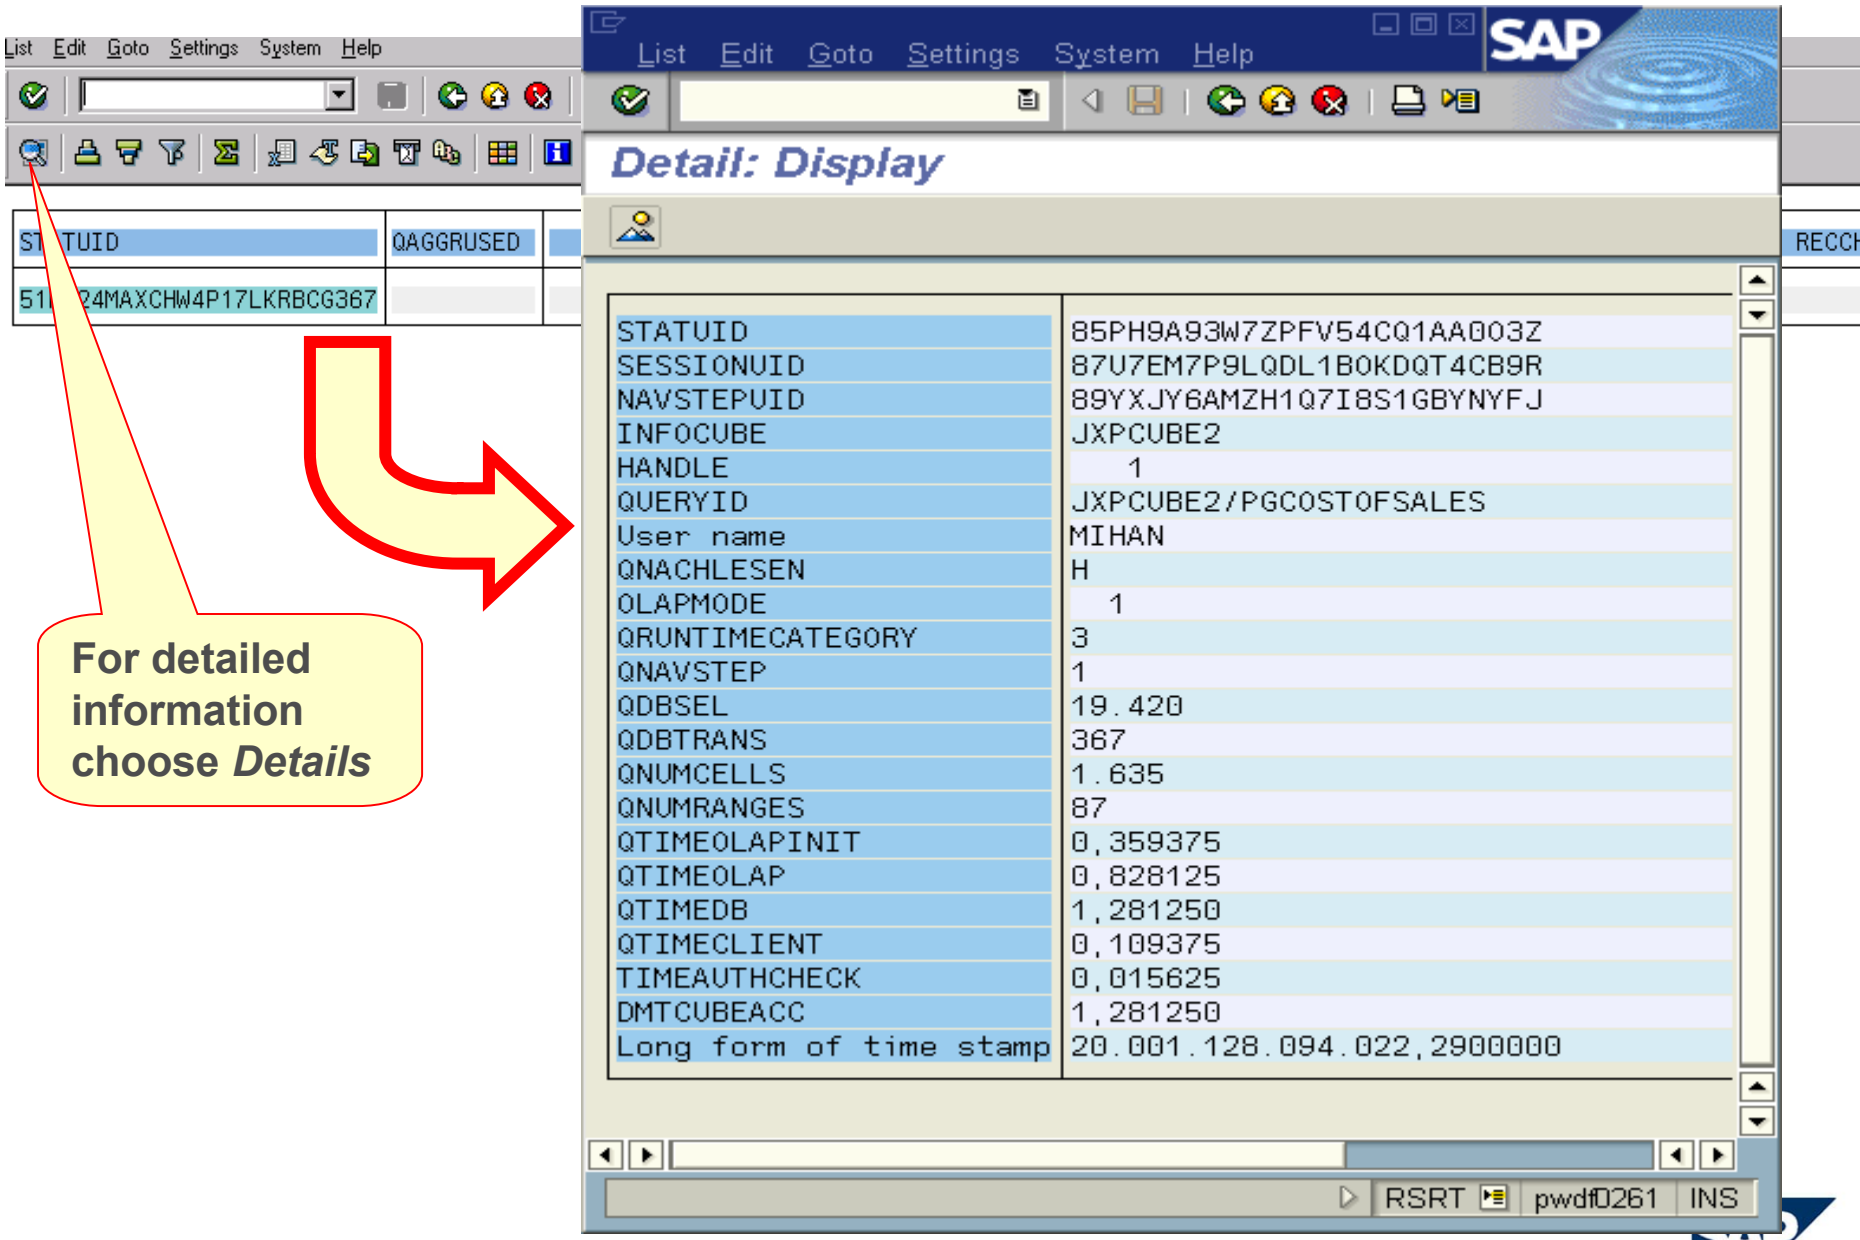

For detailed information choose *Details*

Trace Tool: RSRTRACE

Activates logging for specific user id

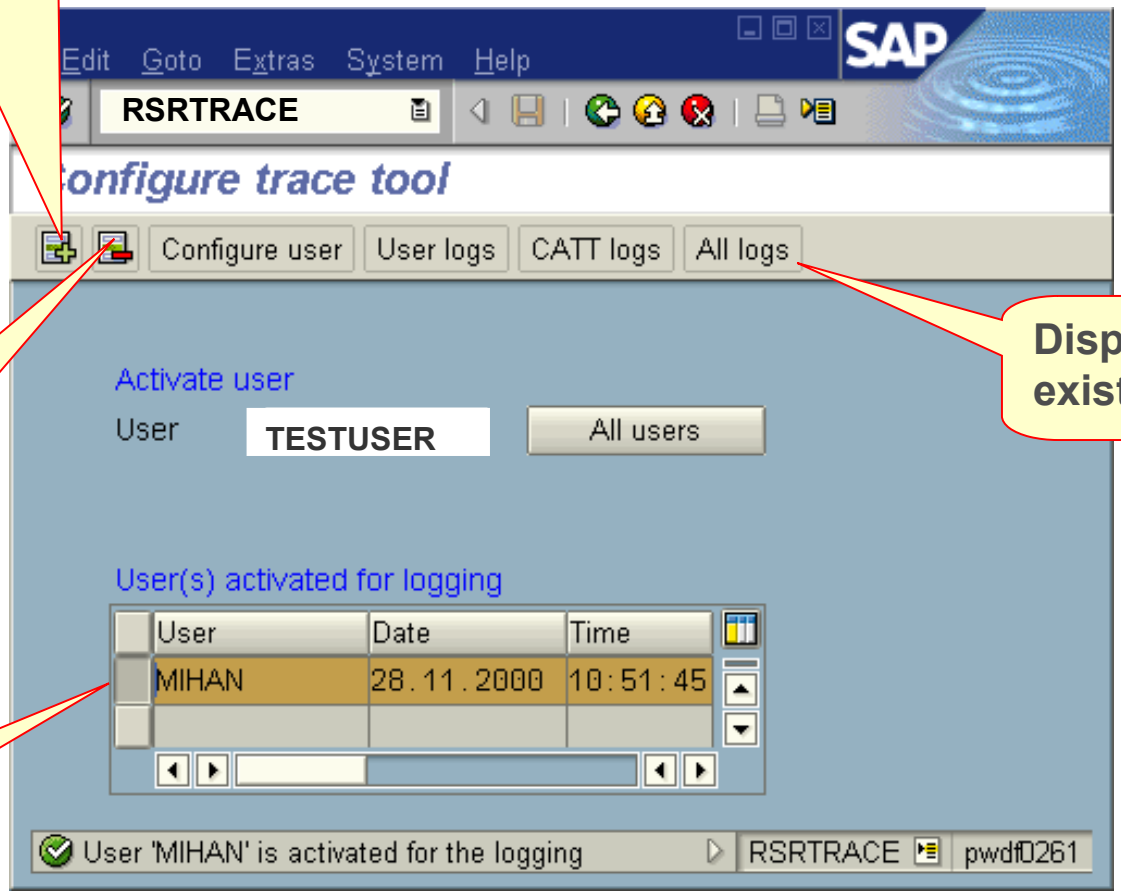

Displays all existing logs

Deactivates logging for user id

User already activated for logging

Trace Tools: RSRTRACE > RSRCATTTRACE

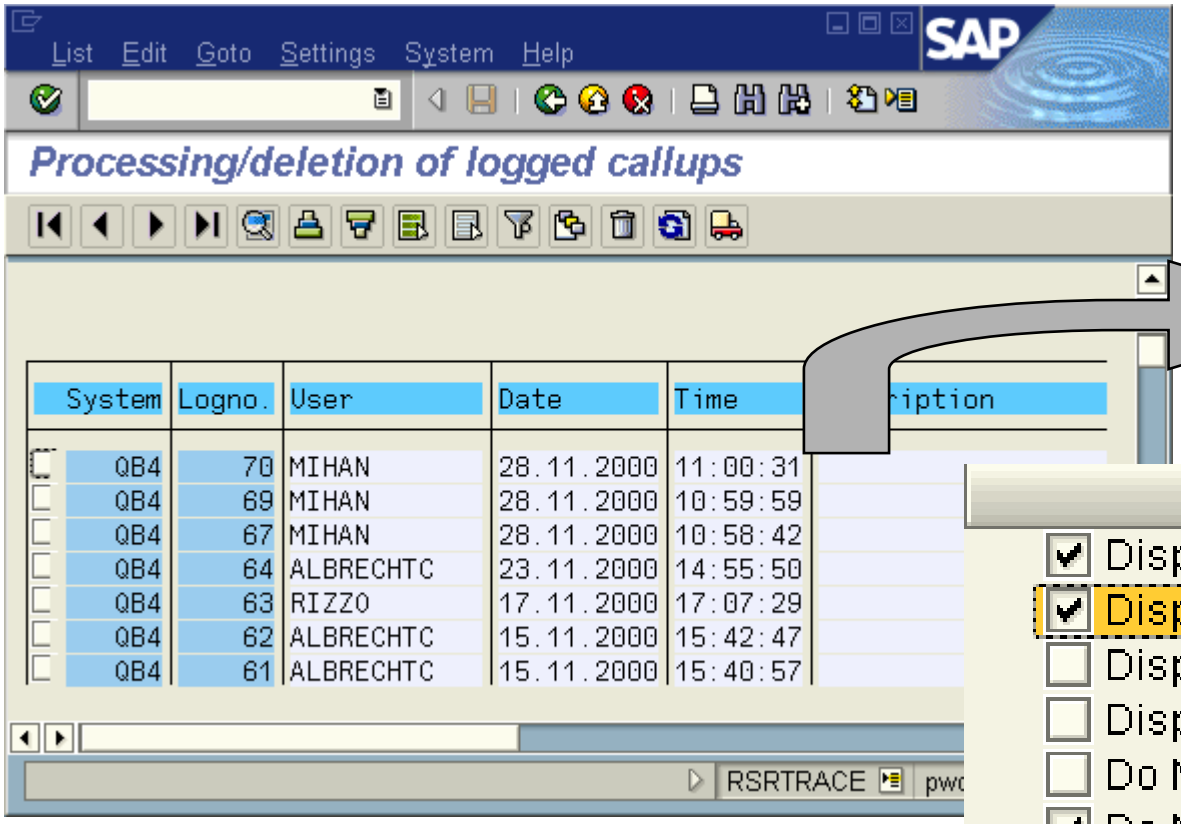

Copy Log Number from RSRTRACE into Transaction RSRCATTTRACE

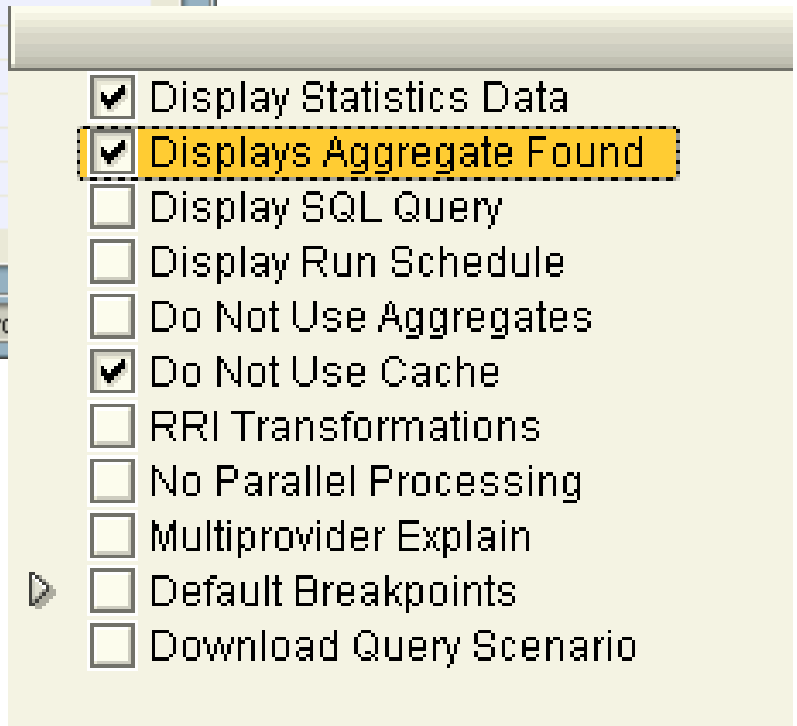

RSRCATTTRACE: specific log from RSRTRACE is input, tool gives aggregates suggestions for first execution AND all further navigations performed

Data Load Performance Concepts

Reporting Performance Analysis Tools

Reporting Performance Analysis Techniques

Aggregates

Reporting Performance Tuning Concepts

Aggregates: Definition

Aggregates are materialized subsets of fact table data. They are independent structures where summary data is stored within separate, transparent InfoCubes.

Aggregates: Concept

Data transfer from InfoCube to query

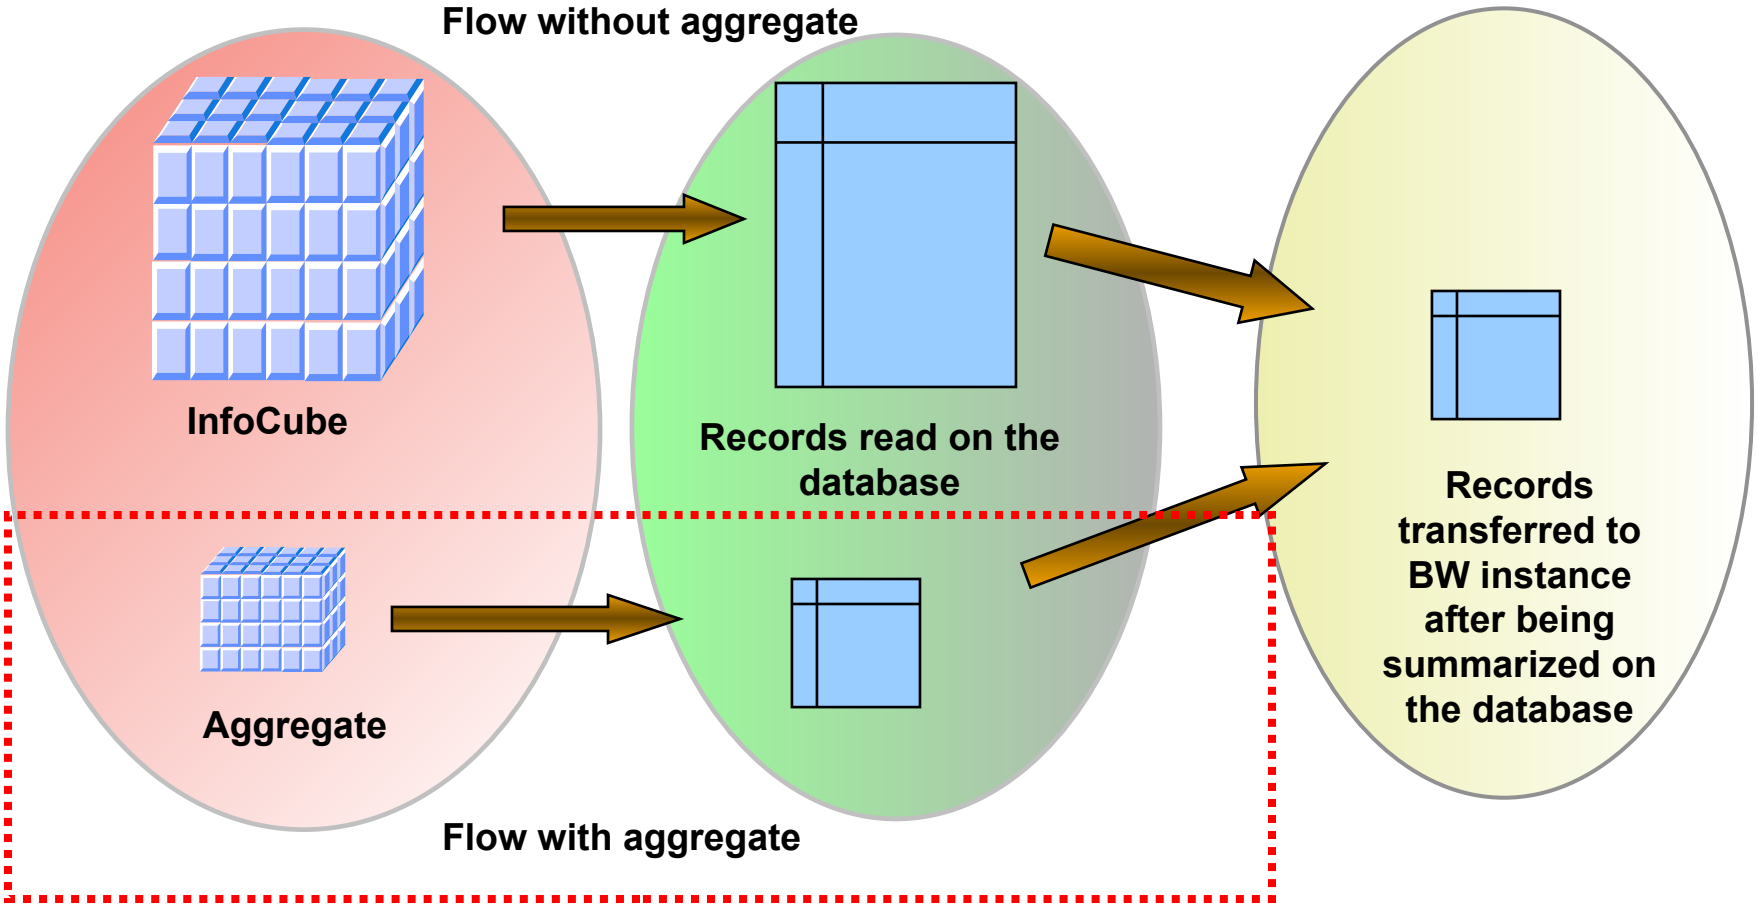

Aggregates: Concept

Data transfer from InfoCube to query (revenue grouped by month)

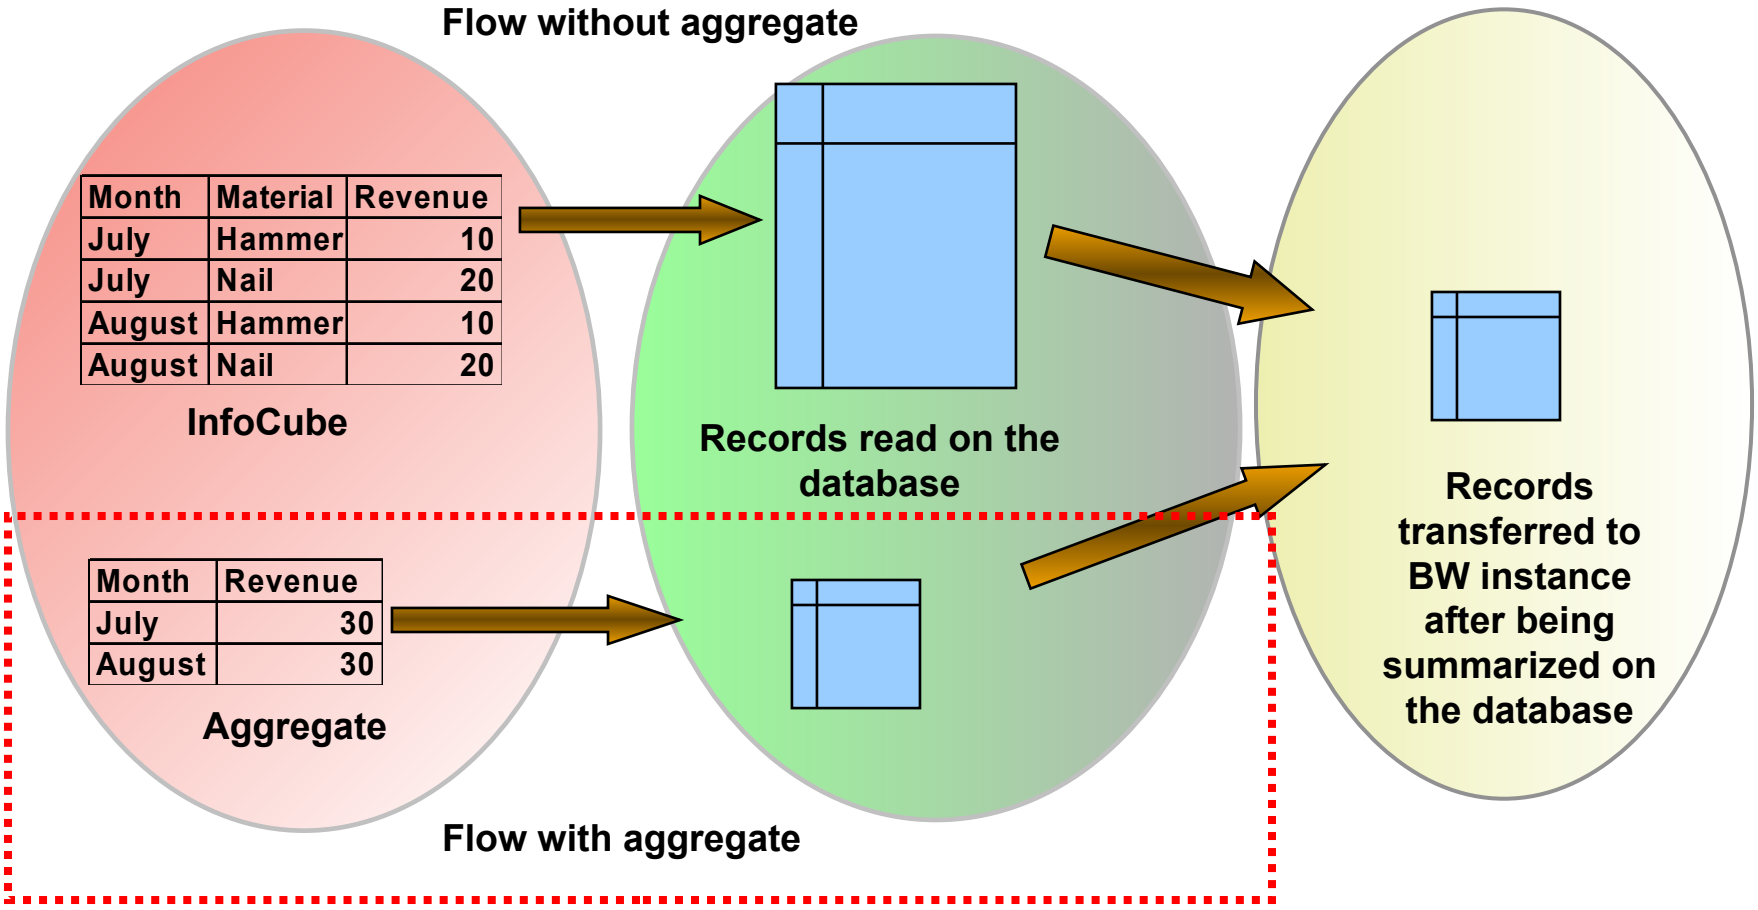

Aggregates: Concept

Data transfer from InfoCube to query (revenue grouped by month)

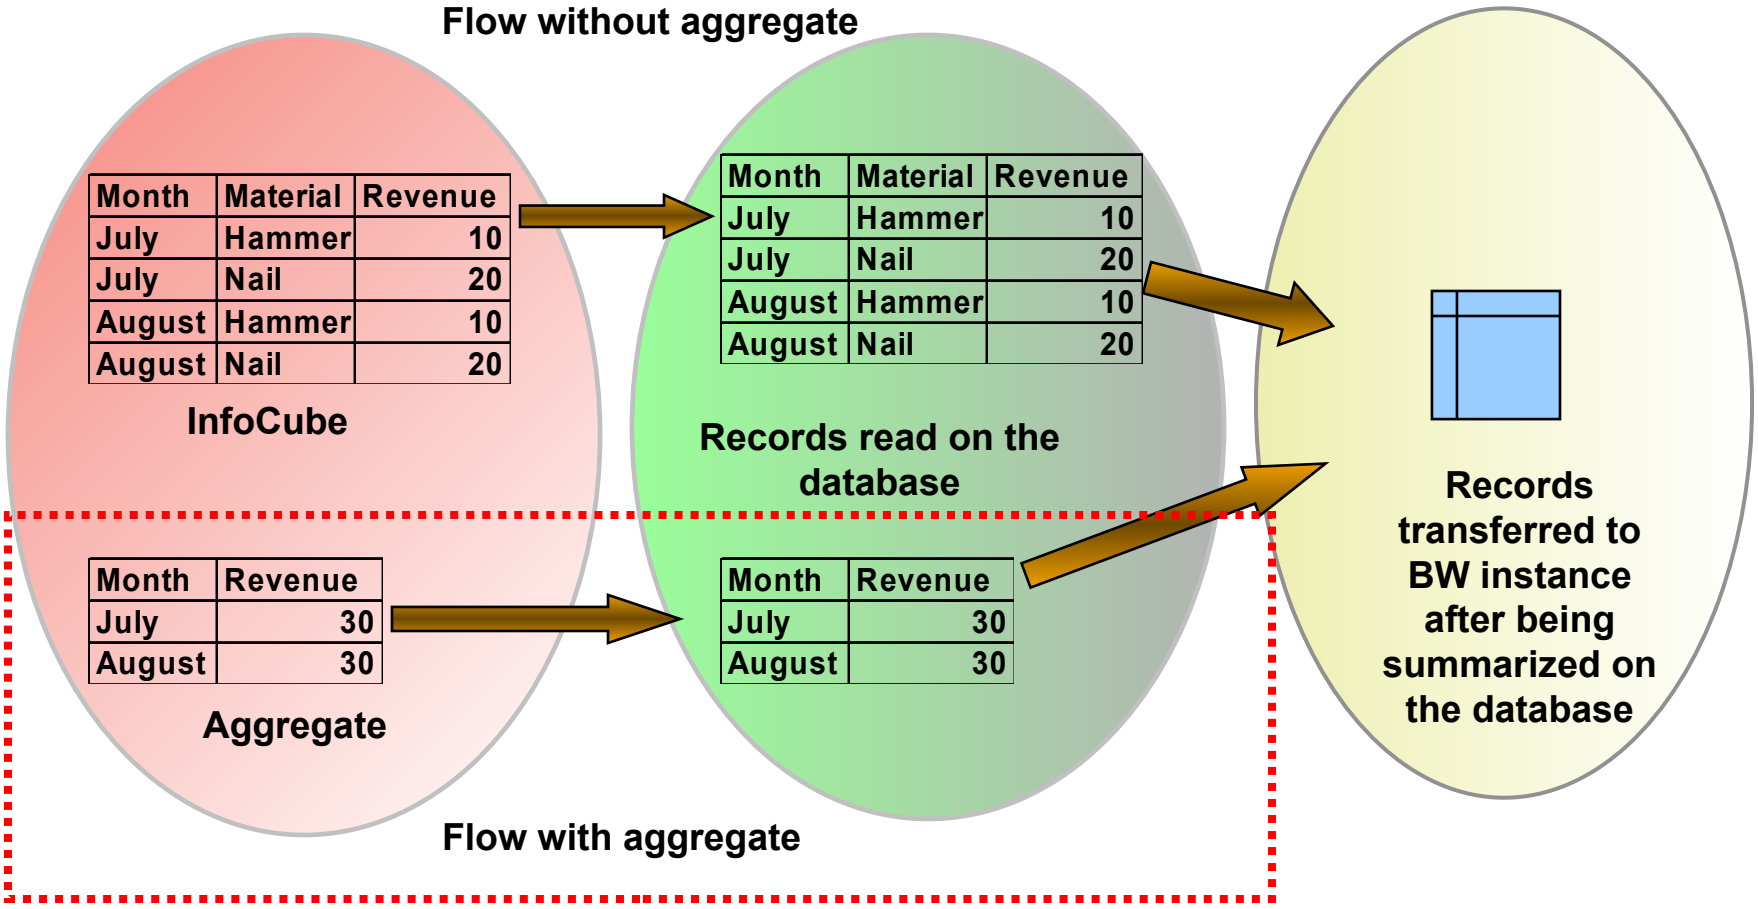

Aggregates: Concept

Data transfer from InfoCube to query (revenue grouped by month)

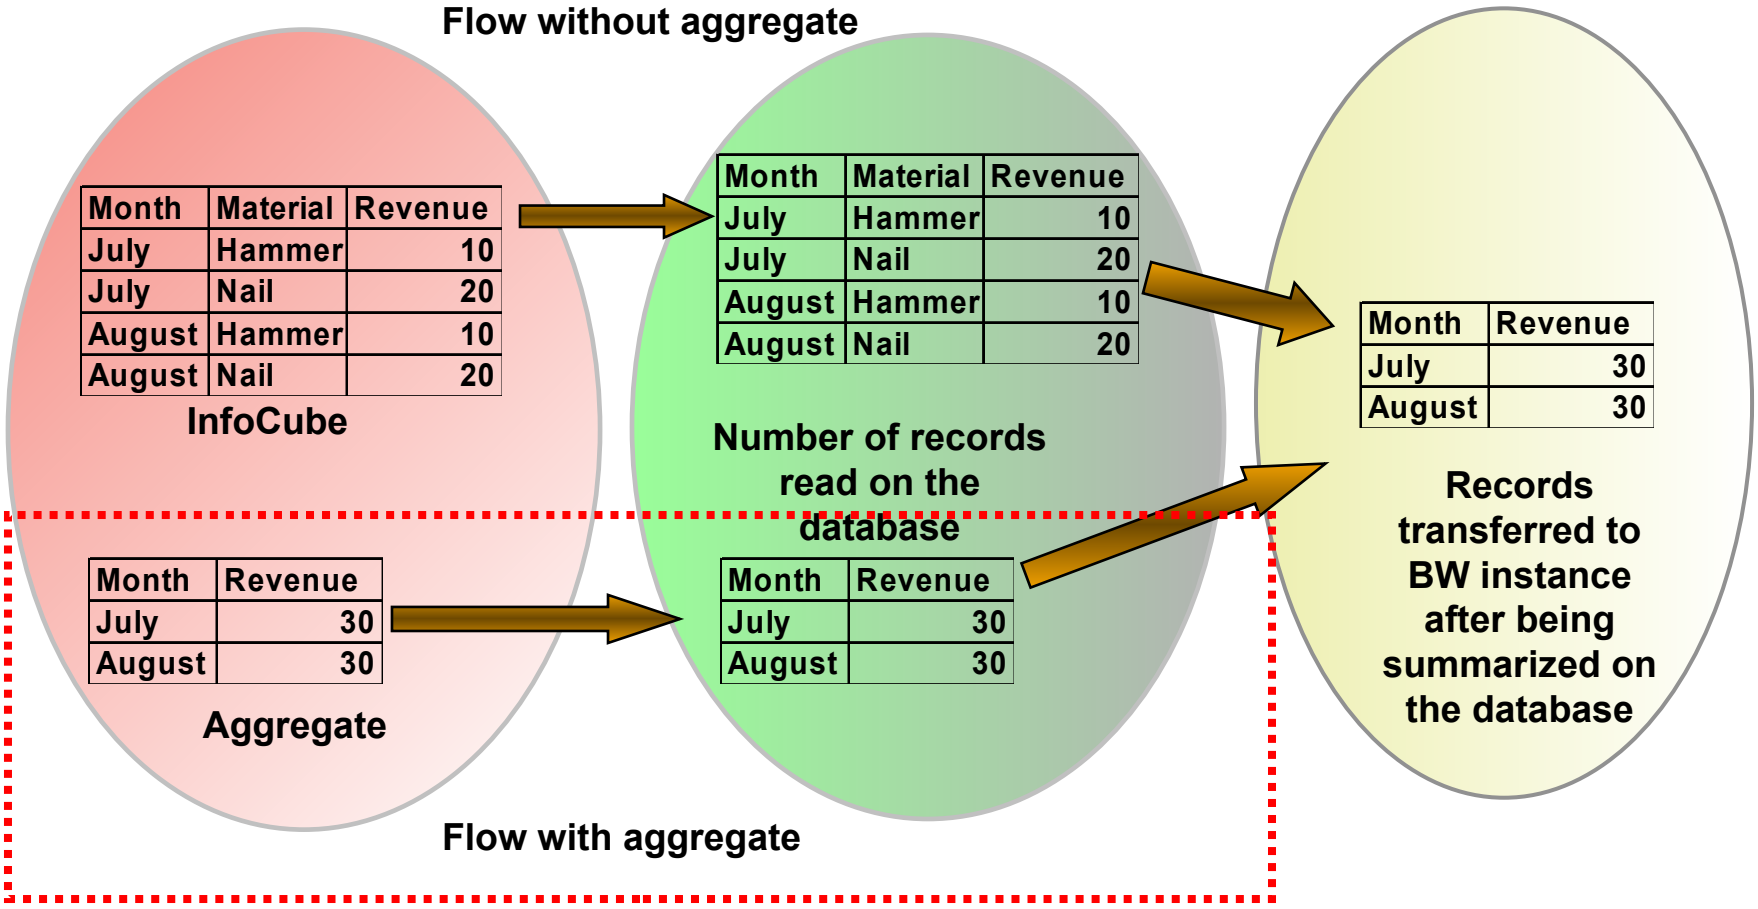

Aggregates: Purpose

To accelerate the response time of queries, by reducing the amount of data that must be read in the database for a given query navigation step.

Query Design: General Recommendation

Using Aggregates

No Aggregates available

Create queries to start at an summarized level

Aggregate 1

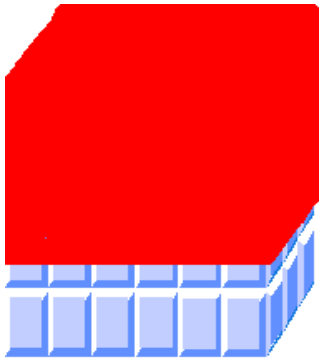

InfoCube

First drilldown may have to read a larger structure

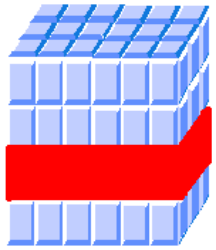

Aggregate 2

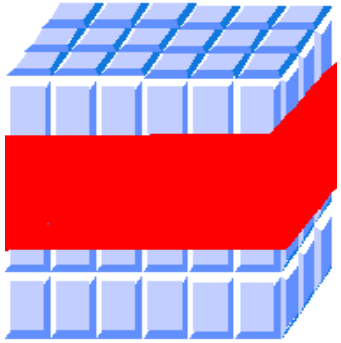

InfoCube

Second drilldown may have to read a large structure
but with increasing filter restrictions

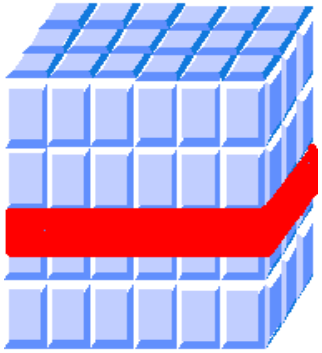

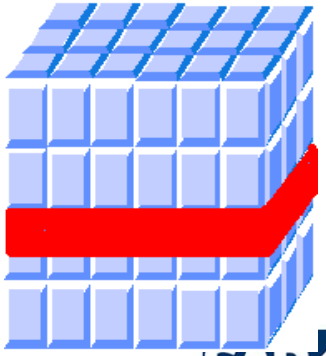

InfoCube

Amount of data that has to be read

Aggregates – Properties

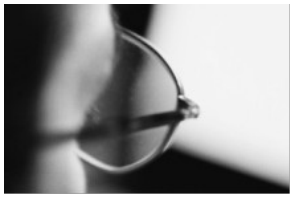

Aggregates can be created:

- For Basic InfoCubes
- For dimension characteristics
- Using navigational attributes
- On hierarchy levels
- Using time-dependent navigational attributes
- On hierarchy levels where the structure is time-dependent

New for 3.0!

New for 3.0!

Aggregate Concepts: Summarize by Characteristic

Fact Table: Sales Data

Country	Customer	Sales
USA	Winsoft Inc.	10
Germany	Internetworks	15
USA	Funny Duds Inc.	5
Austria	Internetworks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Winsoft Inc.	25

Aggregate Tables: Sales Data

Country	*
Customer	Space

Country	Sales
USA	40
Germany	35
Austria	20

Aggregate Concepts: Fixed Value (Subset of Data)

Fact Table: Sales Data

Country	Customer	Sales
USA	Winsoft Inc.	10
Germany	Internetworks	15
USA	Funny Duds Inc.	5
Austria	Internetworks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Winsoft Inc.	25

Aggregate Tables: Sales Data

Country	F, Germany
Customer	*

Country	Customer	Sales
Germany	Internetworks	15
Germany	Funny Duds Inc.	20

Aggregate Concepts: Sum. By Navigational Attributes

Navigational Attribute for Characteristic Customer

Customer	Industry
Winsoft Inc.	Technology
Funny Duds Inc.	Consumer Products
Internetworks	Technology
Thor Industries	Chemical

Fact Table: Sales Data

Country	Customer	Sales
USA	Winsoft Inc.	10
Germany	Internetworks	15
USA	Funny Duds Inc.	5
Austria	Internetworks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Winsoft Inc.	25

Aggregate Tables: Sales Data

Country Customer__Industry	Space *

Industry	Sales
Technology	60
Consumer Products	25
Chemical	10

Aggregate Concepts: Summarize on Hierarchy Levels

Hierarchy for Country

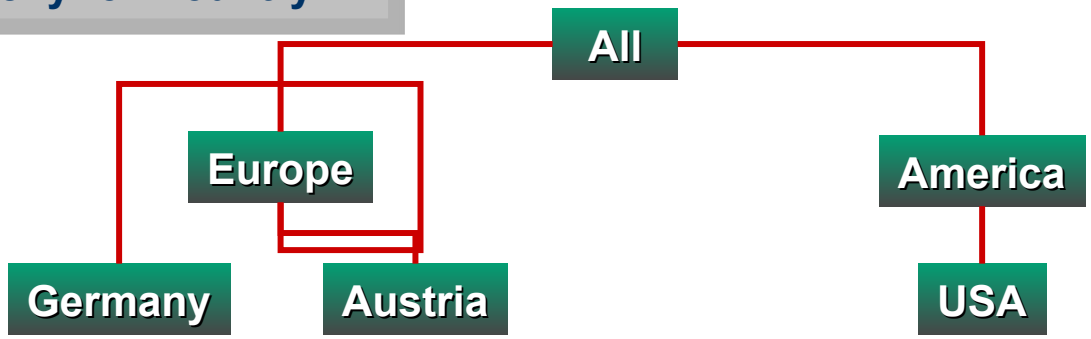

Time-independent hierarchies are stored outside the dimension, in a table called /BI0/ICOUNTRY

Fact Table: Sales Data

Country	Customer	Sales
USA	Winsoft Inc.	10
Germany	Internetworks	15
USA	Funny Duds Inc.	5
Austria	Internetworks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Winsoft Inc.	25

Aggregate Tables: Sales Data

Country H, Level 2
Customer Space

Country	Sales
America	40
Europe	55

Query Behavior: Multiple aggregates per navigation

Queries may be automatically split up over several aggregates, as many queries are comprised of multiple select statements

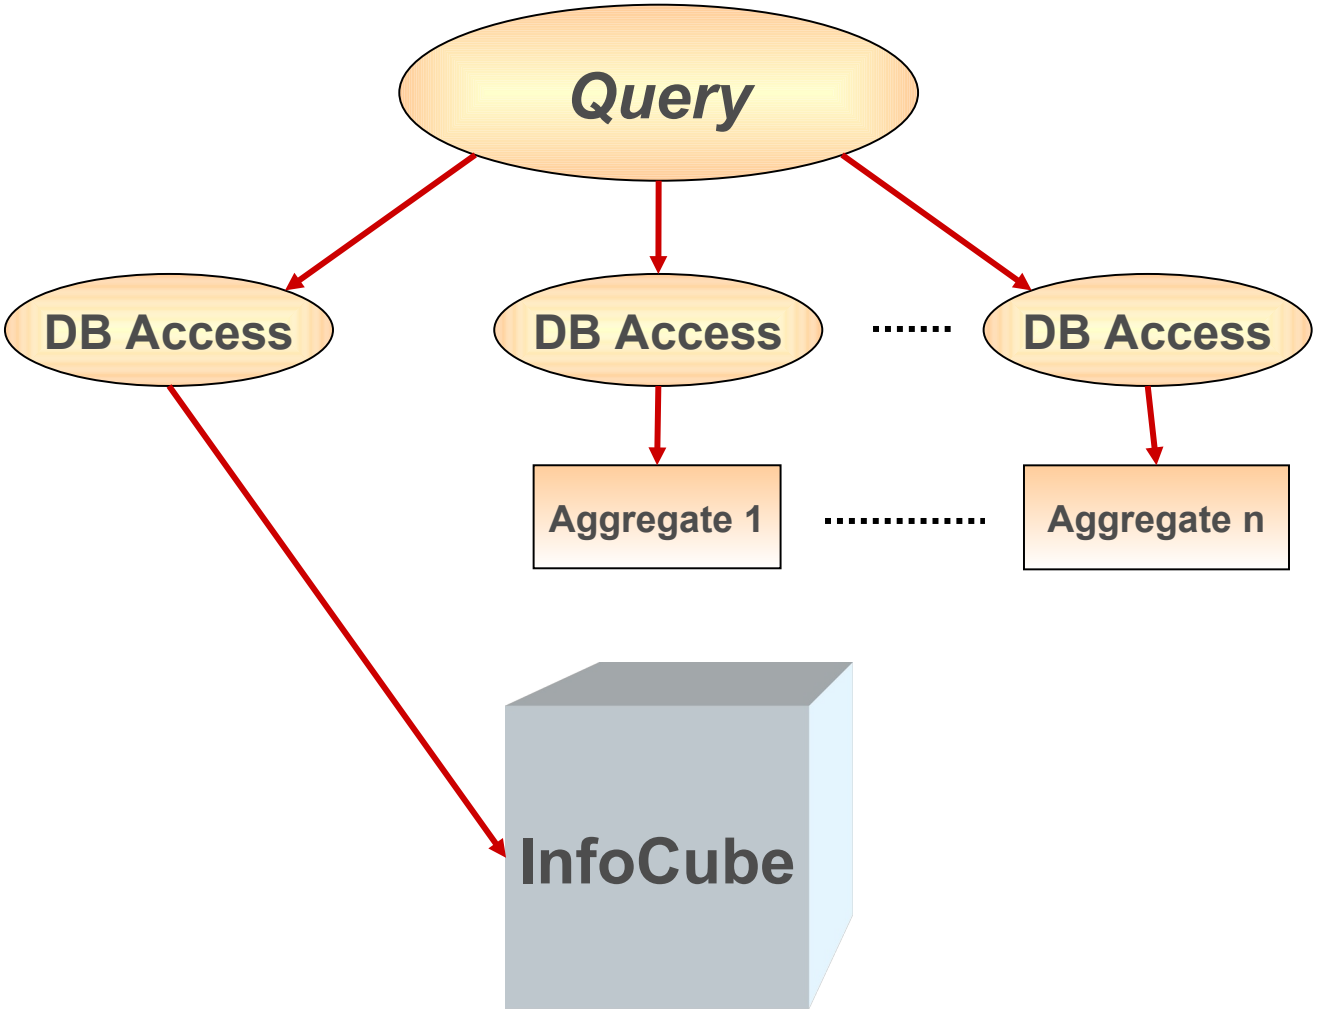

Rollup and Change run (Master Data Activation)

Definition: Rollup

To apply the newly loaded transaction data to the aggregate

Definition: Change Run

To activate the changes of master data and hierarchies

- **During the change run, all aggregates containing navigational attributes and/or hierarchies are realigned**

Aggregates Concepts: Rollup (1)

Rollup applies the newly uploaded transaction data to the aggregate

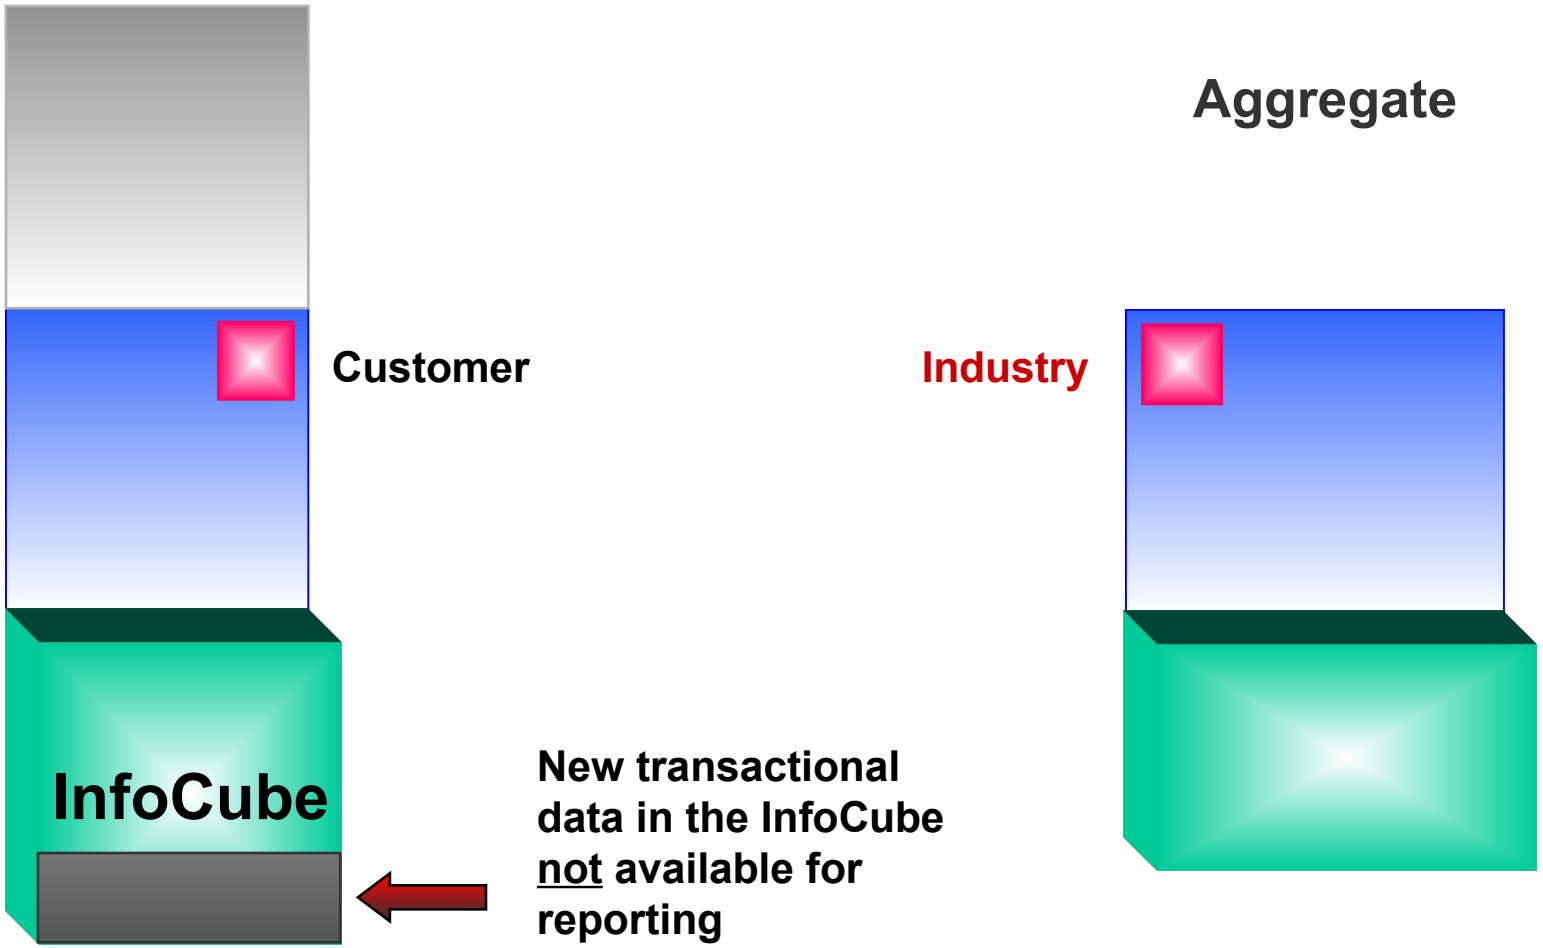

Aggregates Concepts: Rollup (2)

Rollup applies the newly uploaded transaction data to all aggregates of an InfoCube

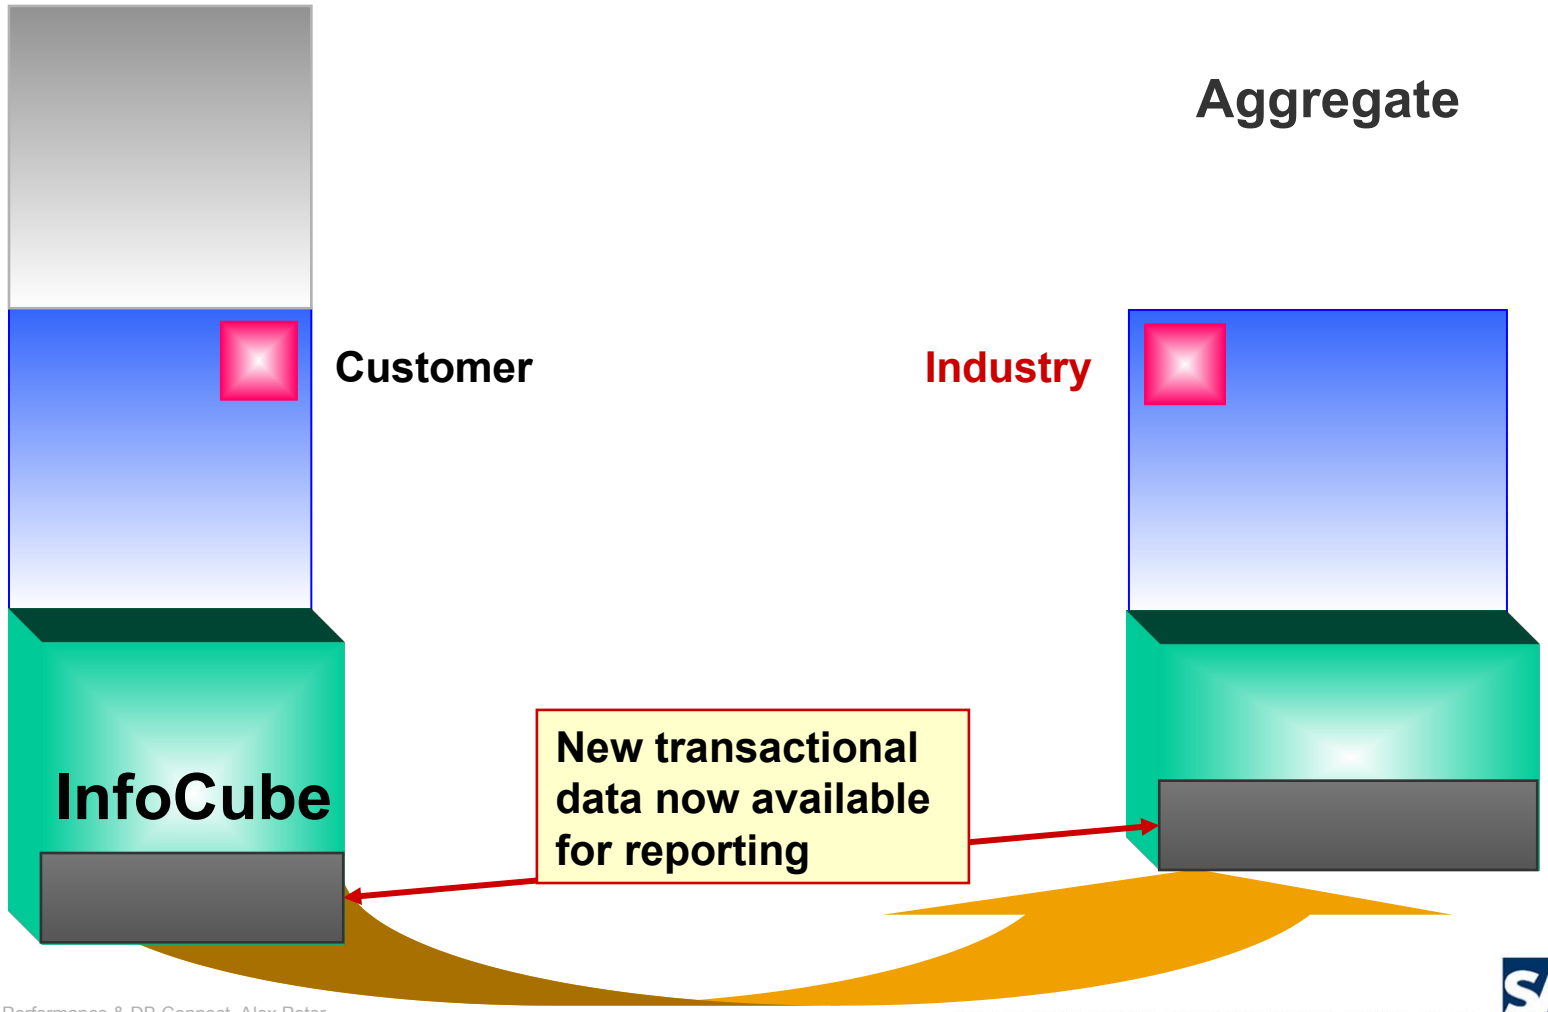

Aggregate Rollup: Create Batch Job

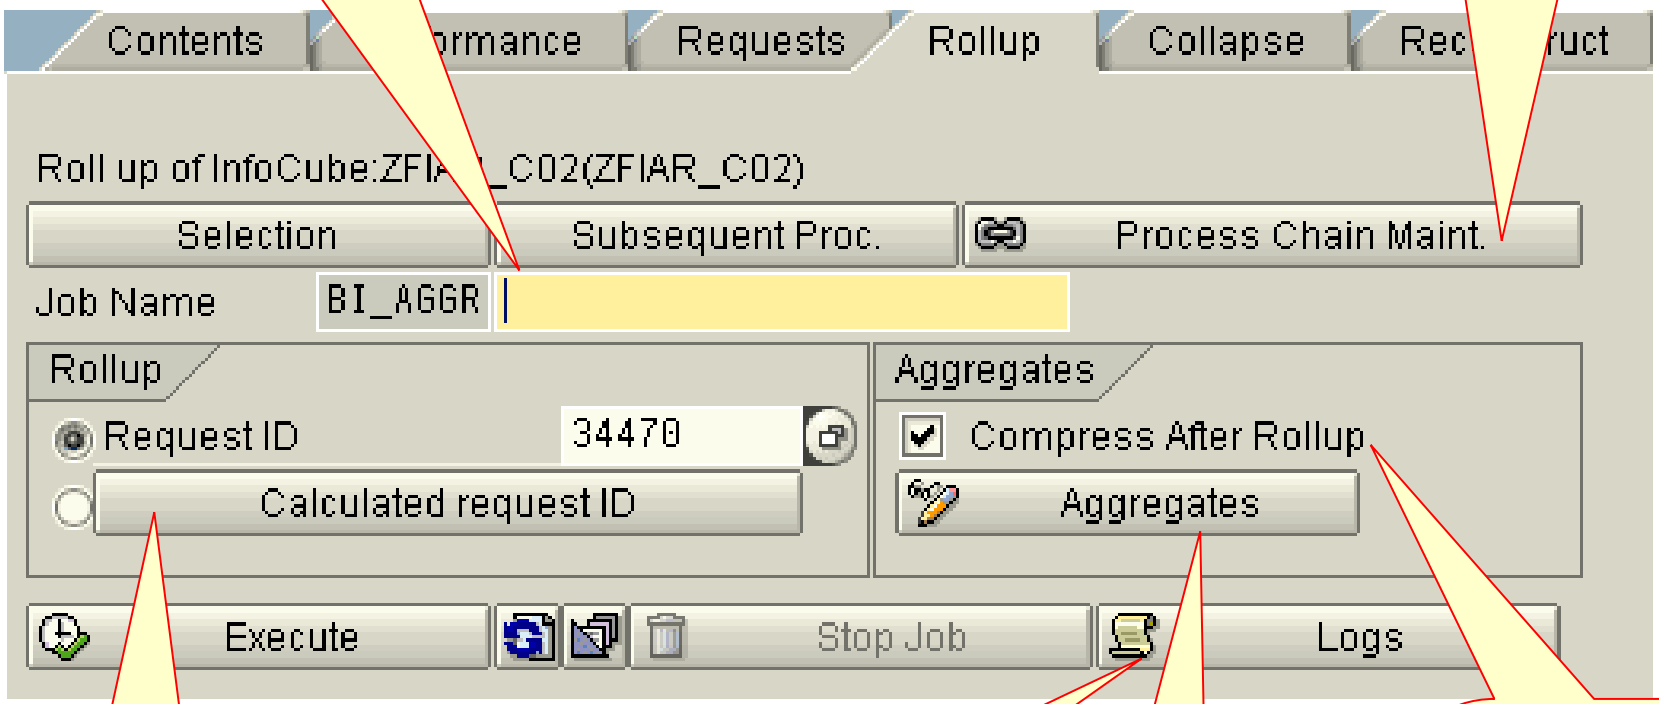

Once selection is made, system-generated batch job name goes here

Aggregate Rollup can be performed via Process Chain

Parameters for system to calculate requests to rollup

View batch job logs for aggregate rollup

View InfoCube aggregates

Compress aggregates with rollup or with IC compression?

Aggregates Concepts: Change Run

Master data activation = Activating the changes of master data and hierarchies. During the change run, all aggregates containing navigational attributes and/or hierarchies are realigned.

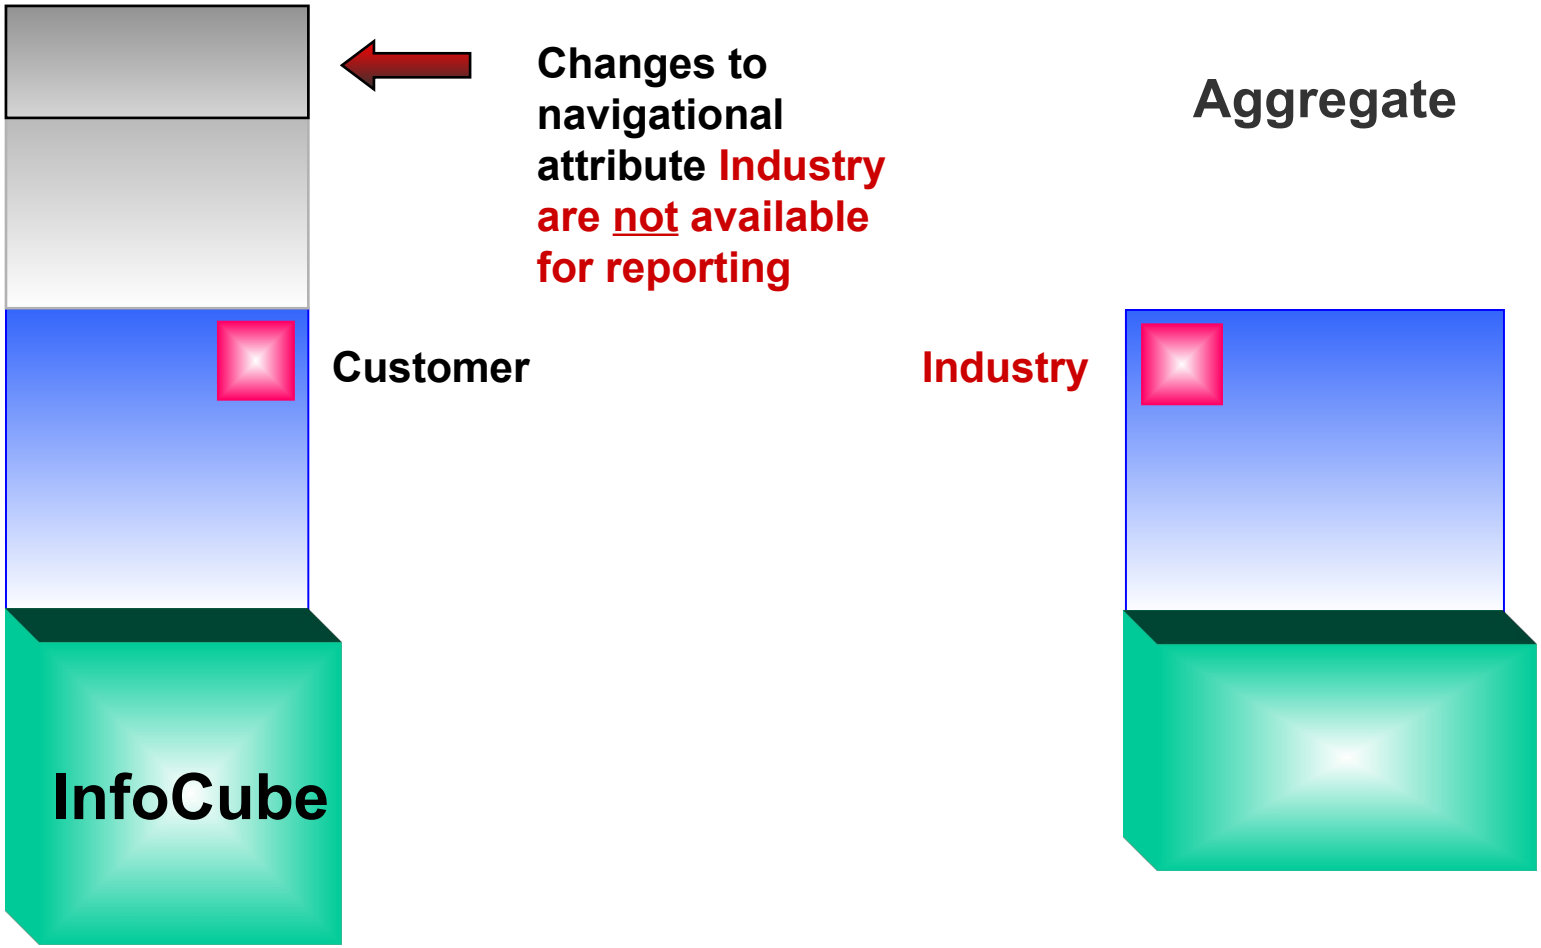

Aggregates Concepts: Change Run

Master data activation = Activating the changes of master data and hierarchies. During the change run, all aggregates containing navigational attributes and/or hierarchies are realigned.

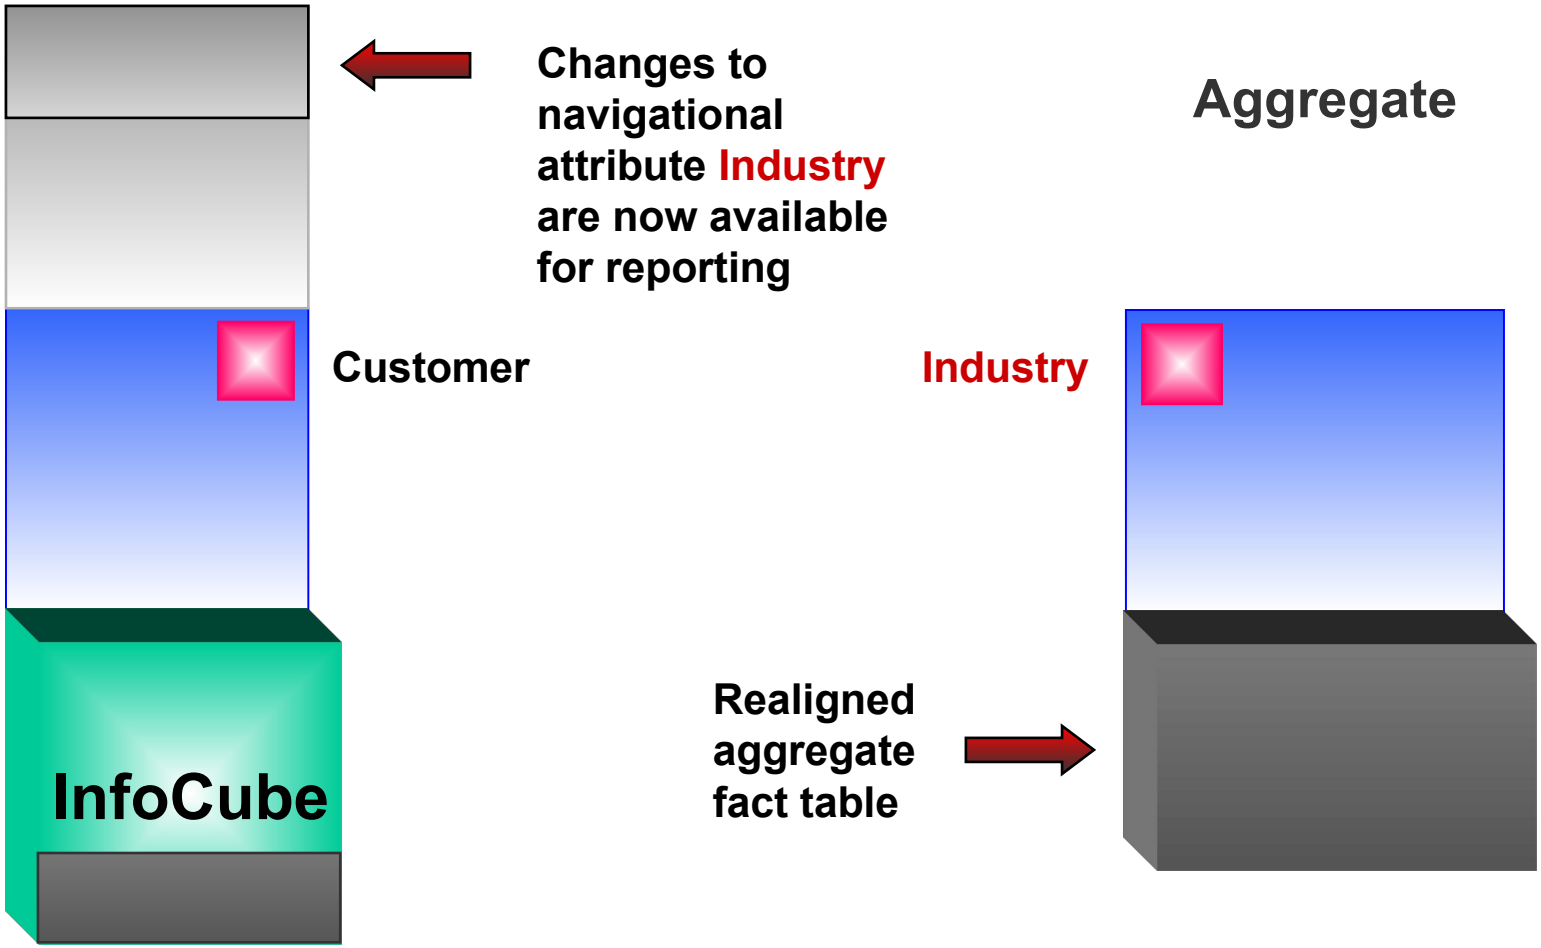

Aggregates Concepts: Change Run

Navigational Attribute for Characteristic Customer

Customer	Industry
Winsoft Inc.	Technology
Funny Duds Inc.	Consumer Products
Old: Internetworks	Technology
Thor Industries	Chemical
New: Internetworks	Consumer Products

Changed master data not available for reporting

Fact Table: Sales Data

Country	Customer	Sales
USA	Winsoft Inc.	10
Germany	Internetworks	15
USA	Funny Duds Inc.	5
Austria	Internetworks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Winsoft Inc.	25

Aggregate Tables: Sales Data

Country	Space
Industry	*

Industry	Sales
Technology	60
Consumer Products	25
Chemical	10

Aggregates Concepts: Change Run

Navigational Attribute for Characteristic Customer

Customer	Industry
Winsoft Inc.	Technology
Funny Duds Inc.	Consumer Products
~~Internetworks~~	~~Technology~~
Thor Industries	Chemical
New: Internetworks	**Consumer Products**

Old: ~~Internetworks~~ / ~~Technology~~

Changed master data now available for reporting

Fact Table: Sales Data

Country	Customer	Sales
USA	Winsoft Inc.	10
Germany	Internetworks	15
USA	Funny Duds Inc.	5
Austria	Internetworks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Winsoft Inc.	25

Aggregate Tables: Sales Data

Country / Industry	Space
	*

Industry	Sales
Technology	**35**
Consumer Products	**50**
Chemical	10

Aggregates Concepts: Change Run and Consistency

Situation before master data activation:

Navigational Attribute for Characteristic Customer

	Customer	Industry	Object version
	Winsoft Inc.	Technology	Active
	Funny Duds Inc.	Consumer Products	Active
Old:	Internetworks	Technology	Active
	Thor Industries	Chemical	Active
New:	Internetworks	Consumer Products	Modified

Changed master data, not available for reporting

Master data activation, including change run, of affected aggregates

Situation after master data activation:

Navigational Attribute for Characteristic Customer

	Customer	Industry	Object version
	Winsoft Inc.	Technology	Active
	Funny Duds Inc.	Consumer Products	Active
	Thor Industries	Chemical	Active
New:	Internetworks	Consumer Products	Active

Changed master data, now available for reporting

Why Can Too Many Aggregates Be Harmful?

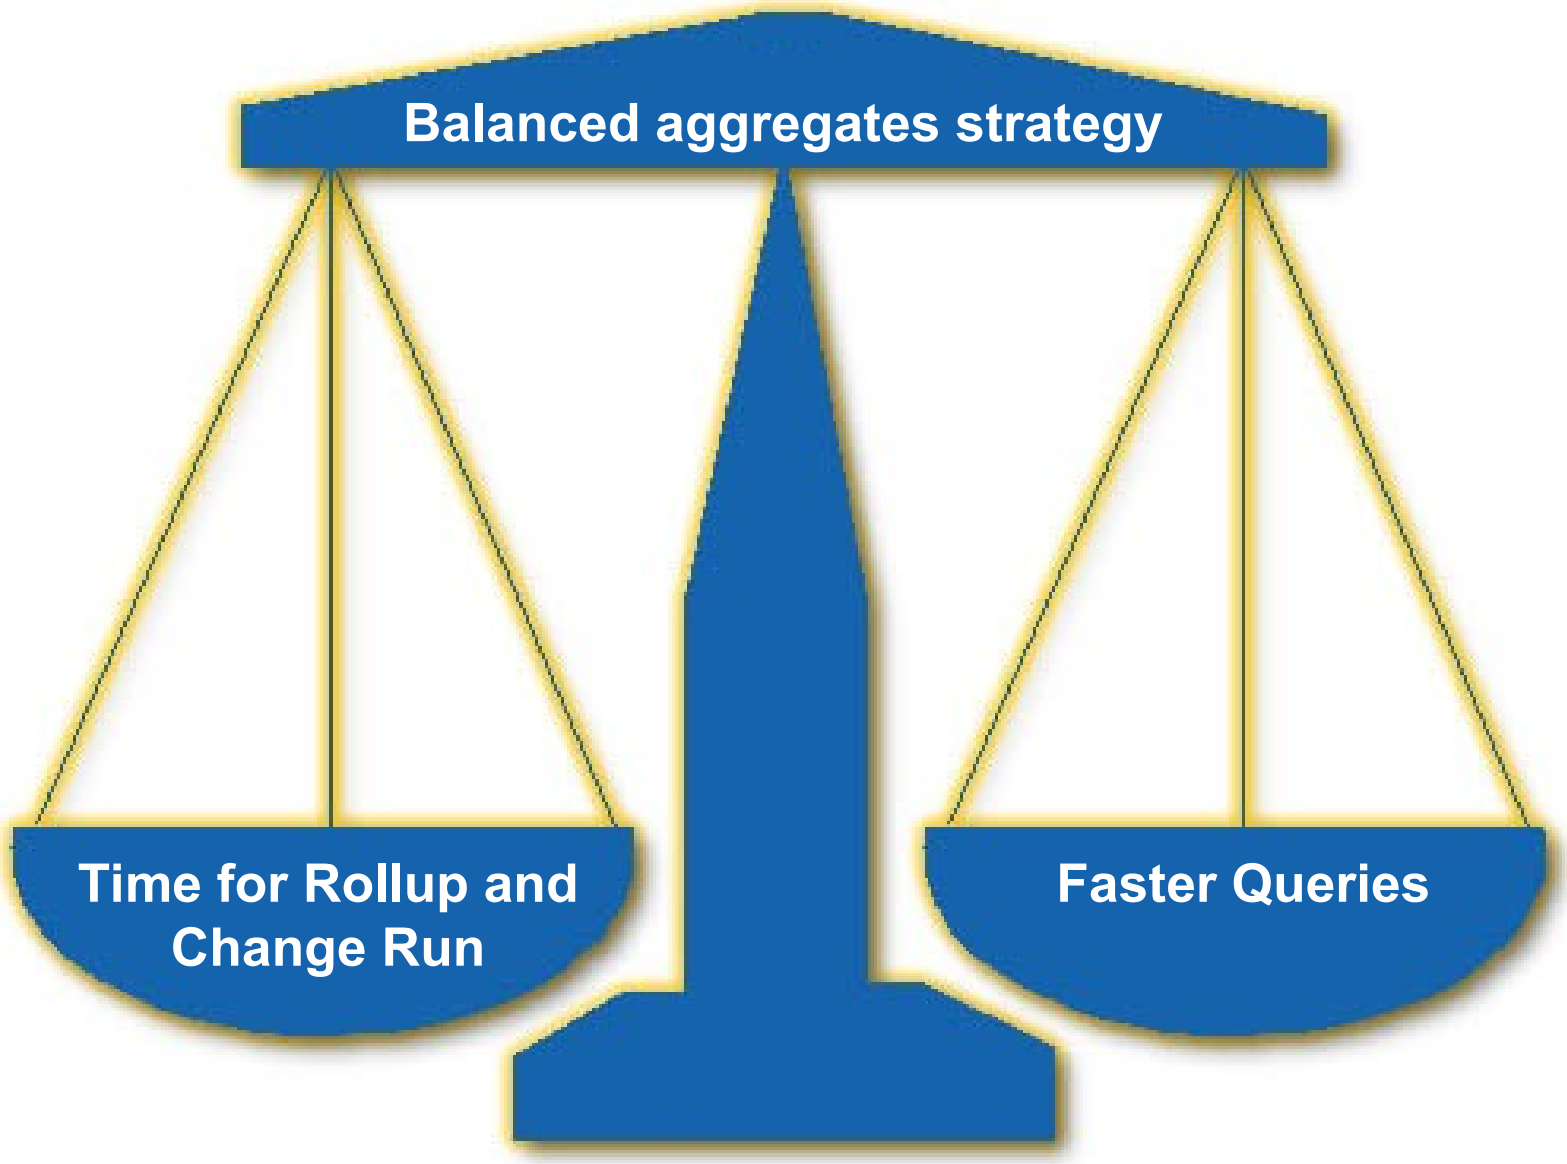

Time Frame

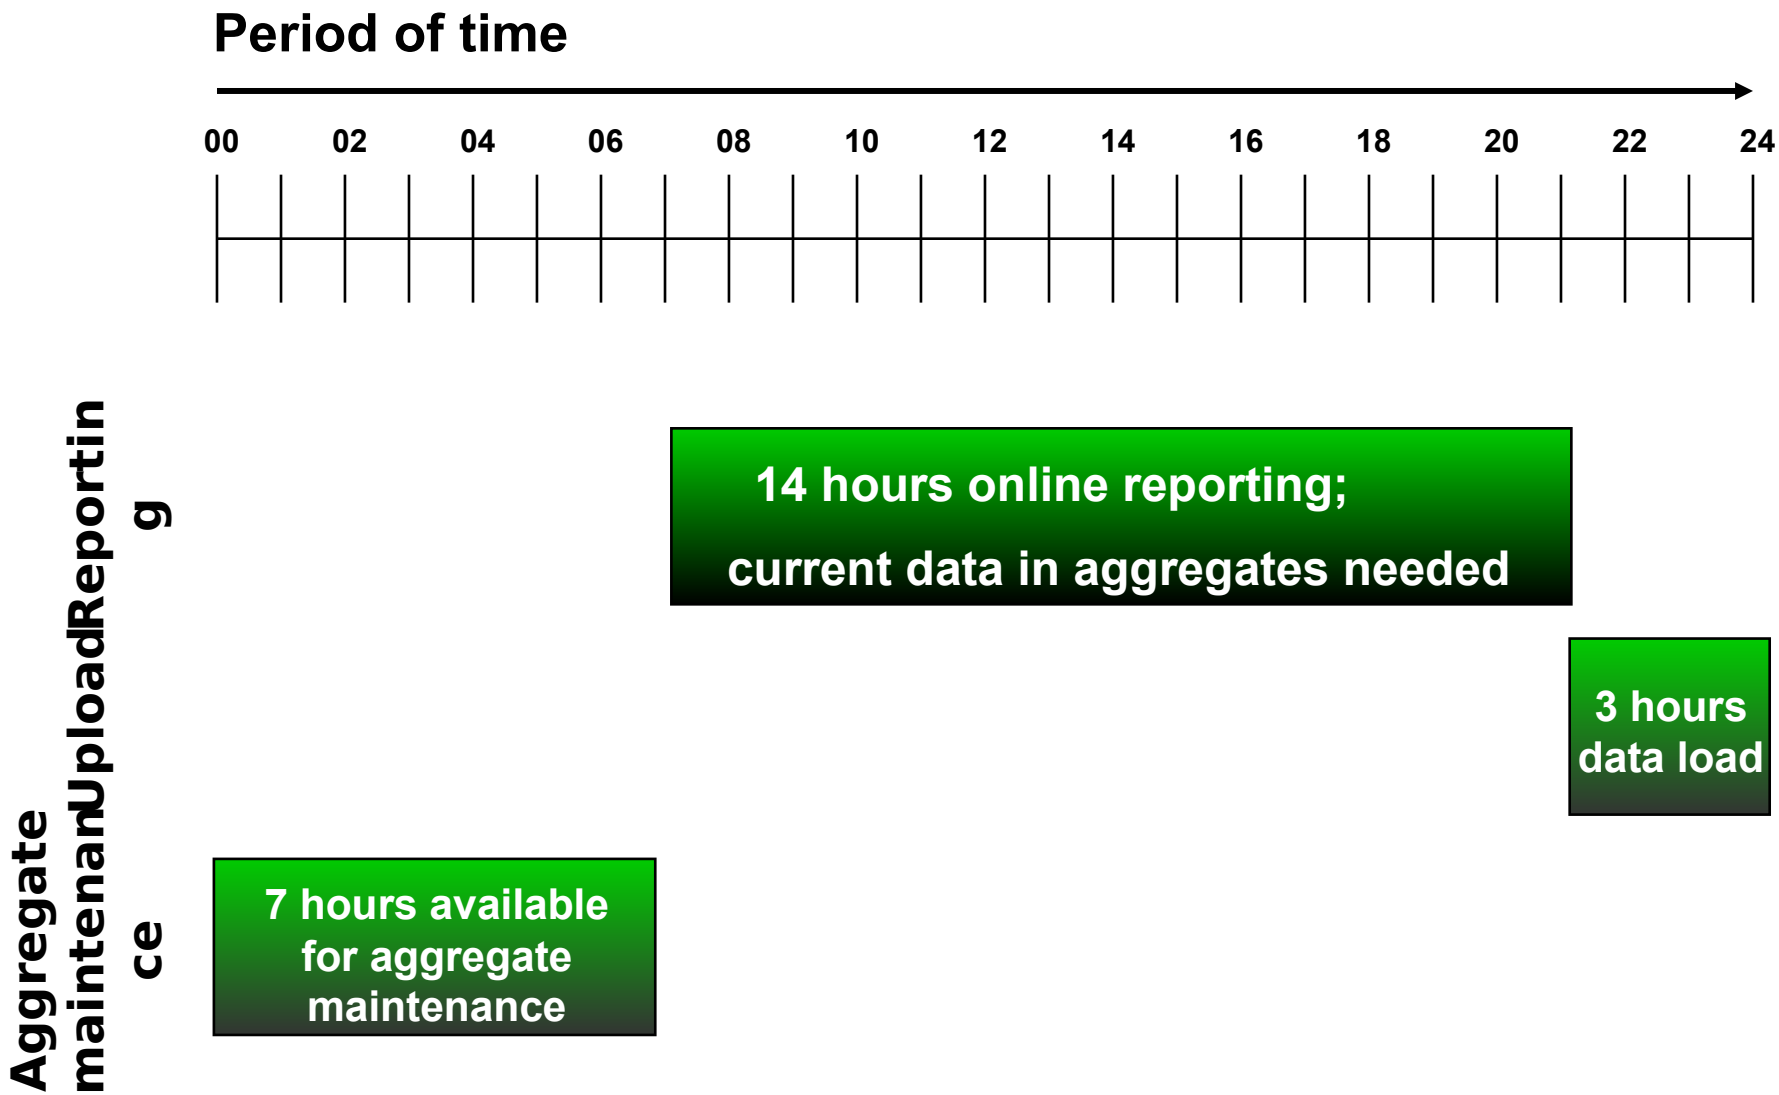

Time Frame: Problem

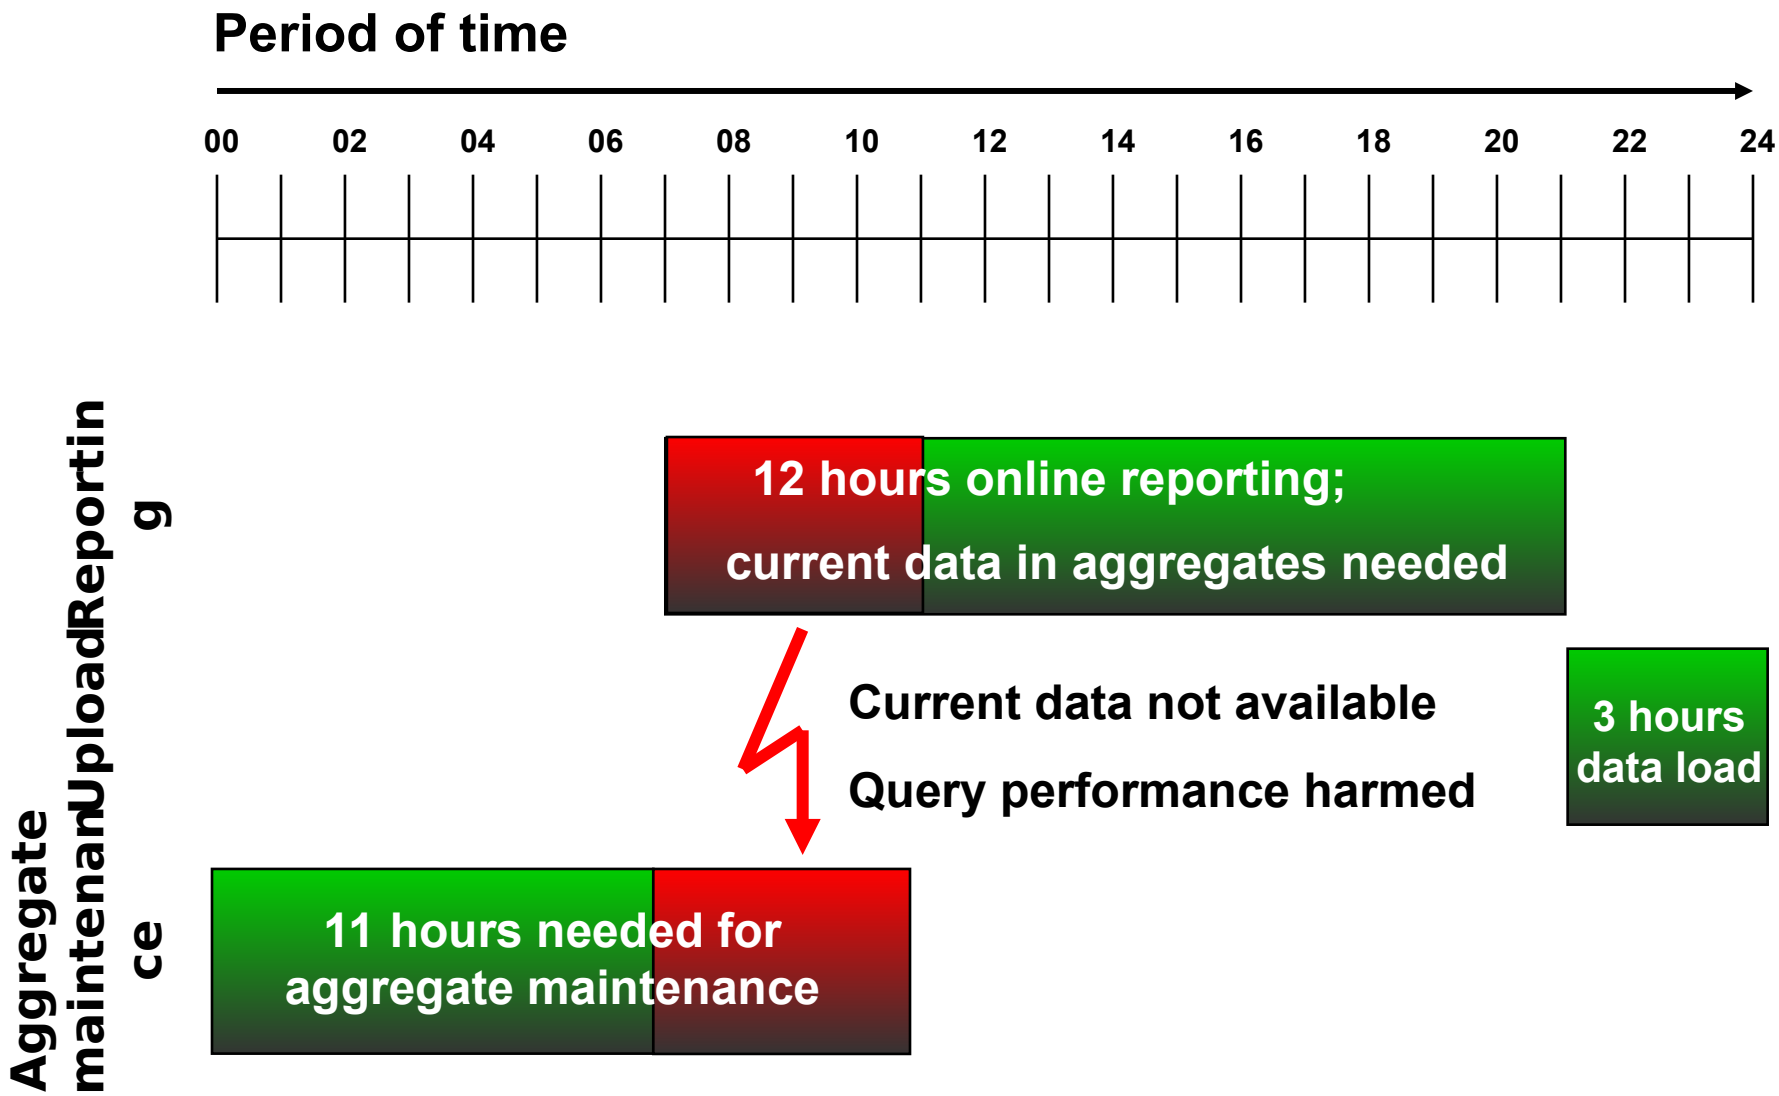

Aggregates: Tuning Tips

Tips for Maintaining Good Aggregates:

- Relatively small compared to parent InfoCube
- Try for summarization ratios of 10 or higher
- Find good subsets of data (frequently accessed)
- Build on some hierarchy levels, not all
- Not too specific, not too general – should serve many different query navigations
- Consider “component” aggregates
- Should be frequently used and used recently (except basis aggregates)

Characteristics of Bad Aggregates:

- Too many very similar aggregates
- Aggregates not small enough (compared to parent cube)
- Too many “for a certain query” aggregates, not enough general ones
- Old aggregates, not used recently
- Infrequently or unused aggregates

Aggregates in BW 3.0

- Aggregates: New Features
  - Time-dependent aggregates
  - Compression & Data Request
  - Filling Aggregates: Block Size
  - Flat Aggregates
  - MOLAP Aggregates

Aggregates with Time-Dependent Components

Aggregates with a time-dependent component (navigational attribute or hierarchy) are calculated for a specific Keydate

A Keydate can be specified in the query definition.

The Keydate can be determined by:

- **a BEx variable which is filled via a SAP- or User- Exit.**
- **a fixed date.**

When the aggregate is filled, data is stored in a manner representative of the state of the data as of the specified Keydate

Using Time-dependent Aggregates in Queries

If time-dependent components are used in a query, it can only use aggregates with the same Keydate.

Important for the use of aggregates is not the variable, but the processed Keydate: only data for this Keydate is available in the aggregate.

Adjustment of Time-dependent Aggregates

Since the Keydates the users are interested in change, the time-dependent aggregates need to be updated regularly.

Process *"Adjustment of Time-Dependent Aggregates"* adjusts data of all aggregates with variables for the Keydate to the changes of the Keydate.

This process is only available in the process chains.

Only aggregates are recalculated if their Keydate has changed.

This adjustment can take a long time and use many system resources.

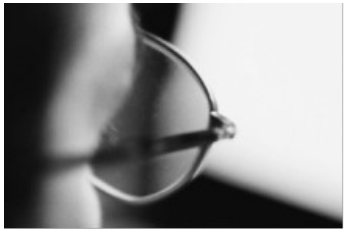

Time-Dependency: Example Using Master Data

Master Data Table: Country

Country	Valid from	Valid to	Sales Person
Austria	1/1/2000	12/31/2000	Huber
Austria	1/1/2001	12/31/2001	Meyer
Germany	1/1/2000	03/31/2001	Meyer
Germany	4/1/2000	12/31/2001	Huber
USA	1/1/2000	12/31/2001	Smith

Fact Table: Sales Data

Country	Customer	Revenue
USA	Buggy Soft Inc.	10
Germany	Ocean Networks	15
USA	Funny Duds Inc.	5
Austria	Ocean Networks	10
Austria	Thor Industries	10
Germany	Funny Duds Inc.	20
USA	Buggy Soft Inc.	25

Aggregate Tables: Sales Data

Sales Person	*
Key date	9/1/2001

Sales Person	Revenue
Huber	35
Meyer	20
Smith	40

Aggregates – Maintenance with Key date

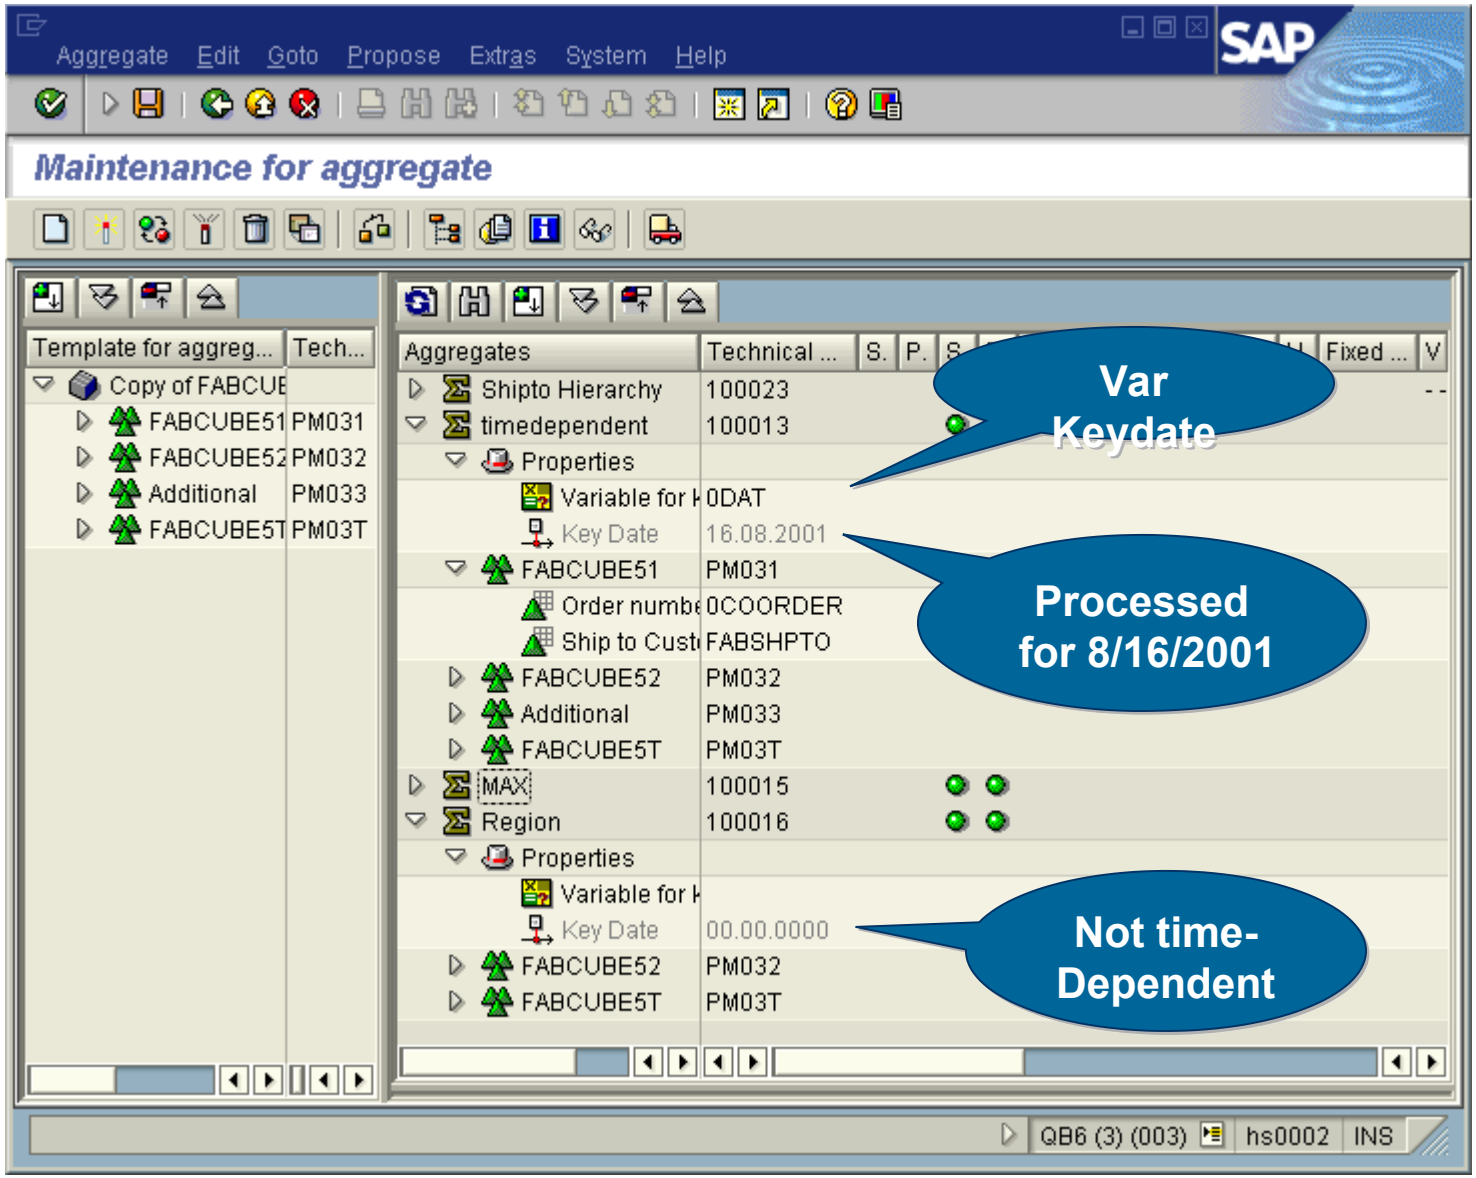

Aggregates in BW 3.0

- **Aggregates: BW 3.0 New Features**
  - Time-dependent aggregates
- **Compression & Data Request**
  - Filling Aggregates: Block Size
  - Flat Aggregates
  - MOLAP Aggregates

Aggregates – Automatic Compression vs. Delayed

InfoCubes can be marked such that the request is kept in its aggregates. (Administrator Workbench; Tabstrip 'Rollup')

New

Requests which are not compressed can then be deleted out of the InfoCube and its aggregates, without completely rebuilding the aggregates.

Aggregates of marked InfoCubes are compressed together with the InfoCube.

This feature should only be used for InfoCubes for which deletion of rolled up requests is necessary -> Performance penalty when aggregates are not compressed.

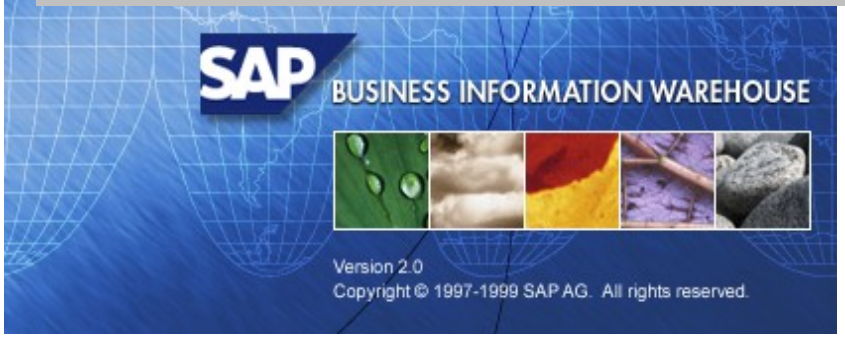

"Two Fact Table" Concept

E fact table

- ... contains consolidated data
- ... is optimized for reading
- ... might be huge
- ... is partitioned by the user
- … cannot be partitioned once InfoCube contains data!

F fact table

- ... contains data on request level
- ... is optimized for writing / deleting
- ... should be small
- ... is partitioned by the system

Aggregates in BW 3.0

- Aggregates: New Features
  - Time-dependent aggregates
  - Compression & Data Request
- Filling Aggregates: Block Size
  - Flat Aggregates
  - MOLAP Aggregates

Aggregates – Block Size

Data of large InfoCubes is read in several blocks to prevent resource problems when filling an aggregate. (Transaction SPRO)

New

Potential problems are, for example, temporary tablespace, memory.

Blocks are distinguished by characteristic values.

The block size can be customized (system-wide).

The best block size depends on the sizing of your system

Aggregates in BW 3.0

- **Aggregates: New Features**
  - Time-dependent aggregates
  - Compression & Data Request
  - Filling Aggregates: Block Size
- **Flat Aggregates**
  - MOLAP Aggregates

Flat Aggregates

If an aggregate has less then 15 components, each component is put into a separate dimension (“Flat Aggregates”).

The dimensions (except the package and unit) are marked as “Line Item”.

“Flat Aggregates” are filled / rolled up without loading the data into the application server (performance gain).

Line item / high cardinality dimensions already specified in the InfoCube are also used for aggregates for corresponding dimensions (also for non-flat aggregates).

Flat Aggregates

Transp. table	/BIC/F100061	Active
Short text	KWFAB5	

Attributes | Delivery and Maintenance | Fields | Entry help/check | Currency/Quantity Field

Srch help | Built-in type

Field	Key	Initi...	Data element	DTyp	Length	Dec.p...	Short text
KEY_100061P	☑	☑	RSDIMID	INT4	10	0	Dimension Table Key
KEY_100061T	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_KWFAB5U	☑	☑	RSDIMID	INT4	10	0	Dimension Table Key
KEY_1000611	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000612	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000613	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000614	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000615	☑	☑	RSSID	INT4	10	0	Master data ID
CRMEM_QTY	☐	☐	/BI0/OICRMEM_QTY	QUAN	17	3	Credit memos quantity

Data Load Performance Concepts

Reporting Performance Analysis Tools

Reporting Performance Analysis Techniques

Aggregates

Reporting Performance Tuning Concepts

Reporting **Performance Tuning Concepts**

- MultiProvider
- Partitioning
- Data Modeling and Line Item Dimension
- Web Reporting
- OLAP Query Cache

Performance Tuning Concepts

- MultiProvider
- Partitioning
- Data Modeling and Line Item Dimension
- Web Reporting
- OLAP Query Cache

Performance Tuning Concepts: MultiProvider

Combination of all types of InfoProviders

No additional data storage

Queries are split automatically and distributed to InfoProviders

Transparent usage for reporting

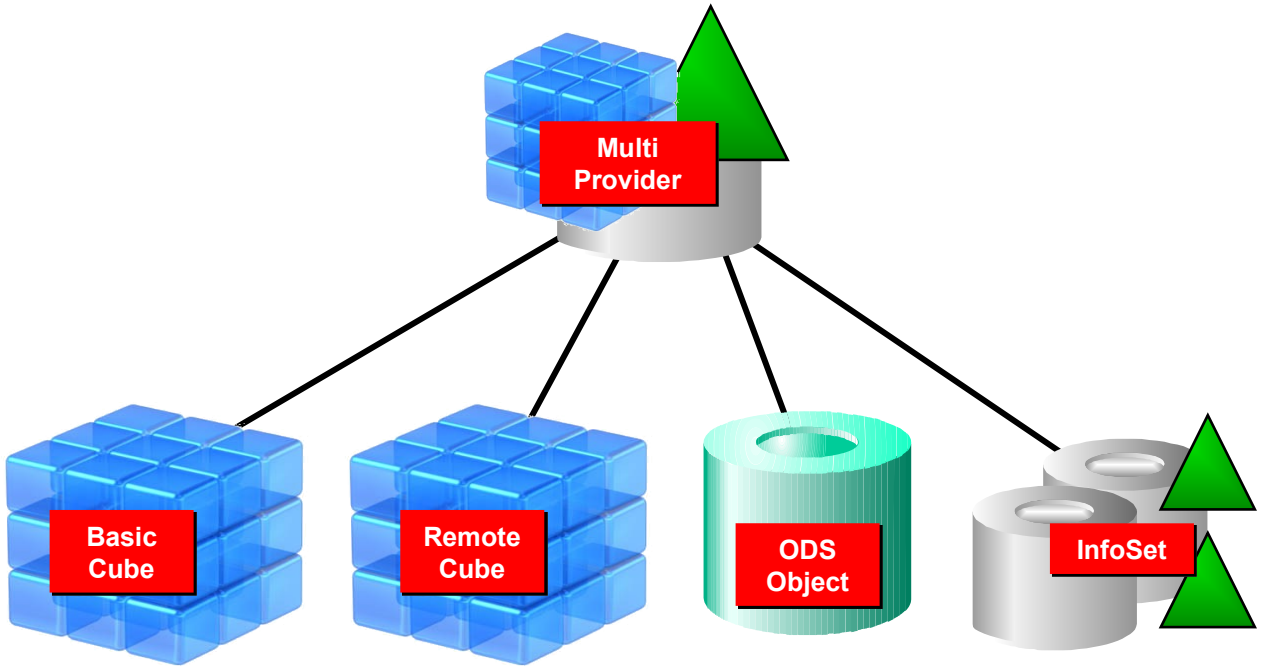

MultiProvider: Overview

- An InfoCube represents a reporting scenario.
- Reporting scenarios might comprise sub-scenarios.
- Example (1): plan and actual data.
- Example (2): order, delivery, billing data.

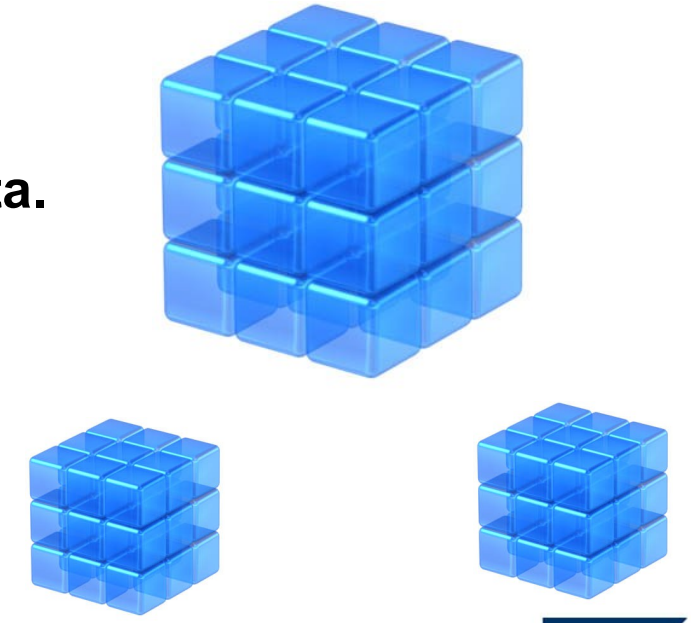

Logical Partitioning via Data Marts

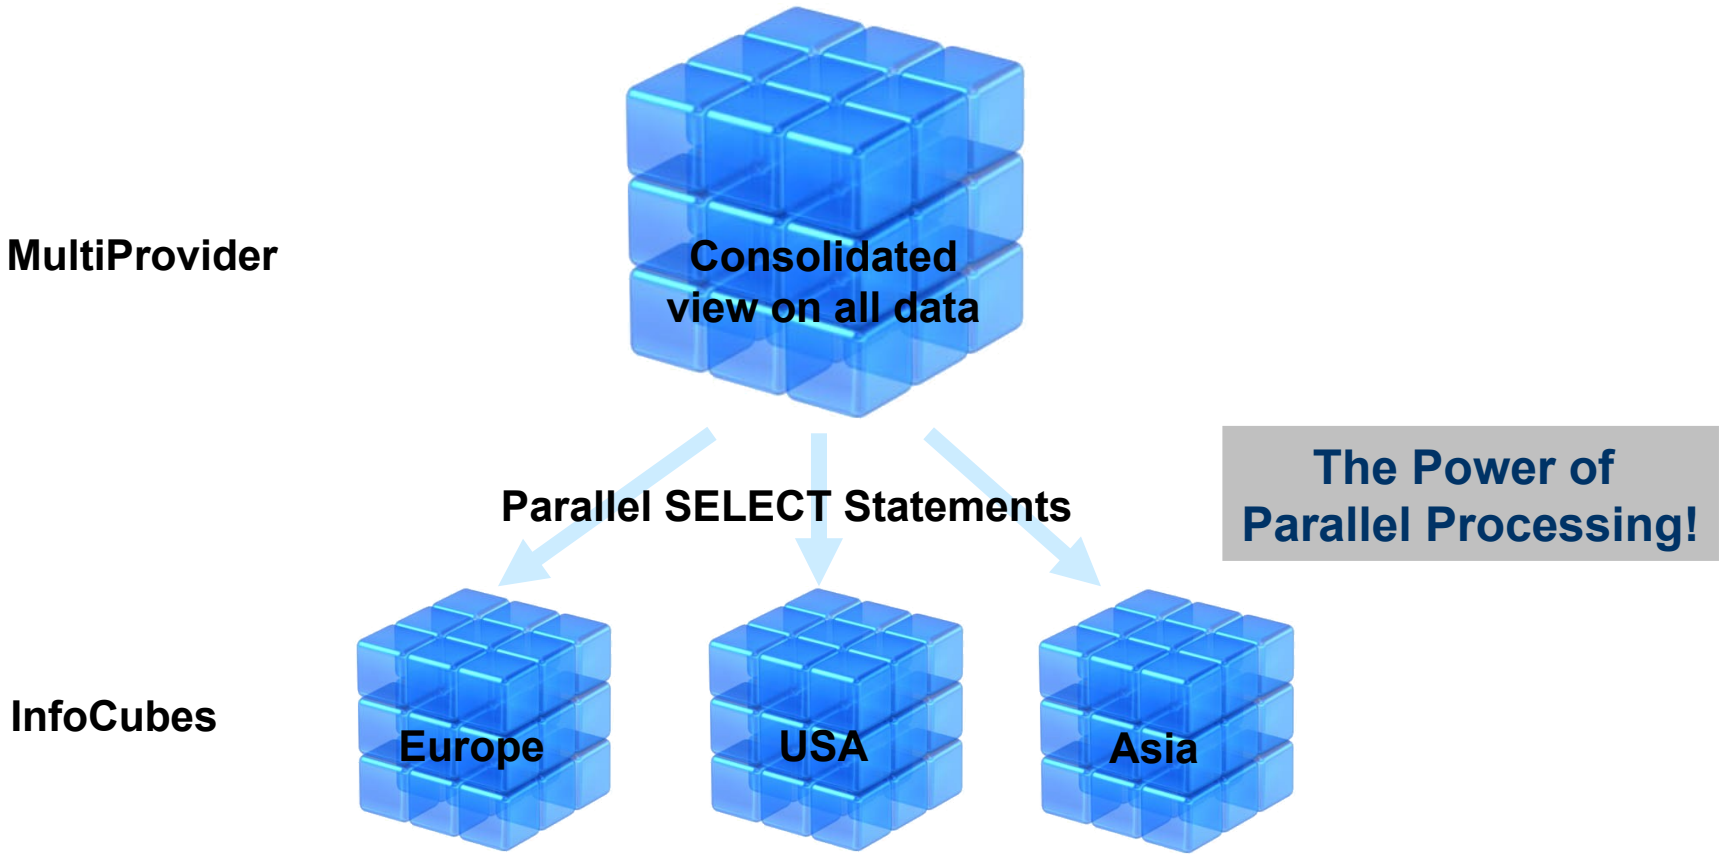

Multi-Cube Example

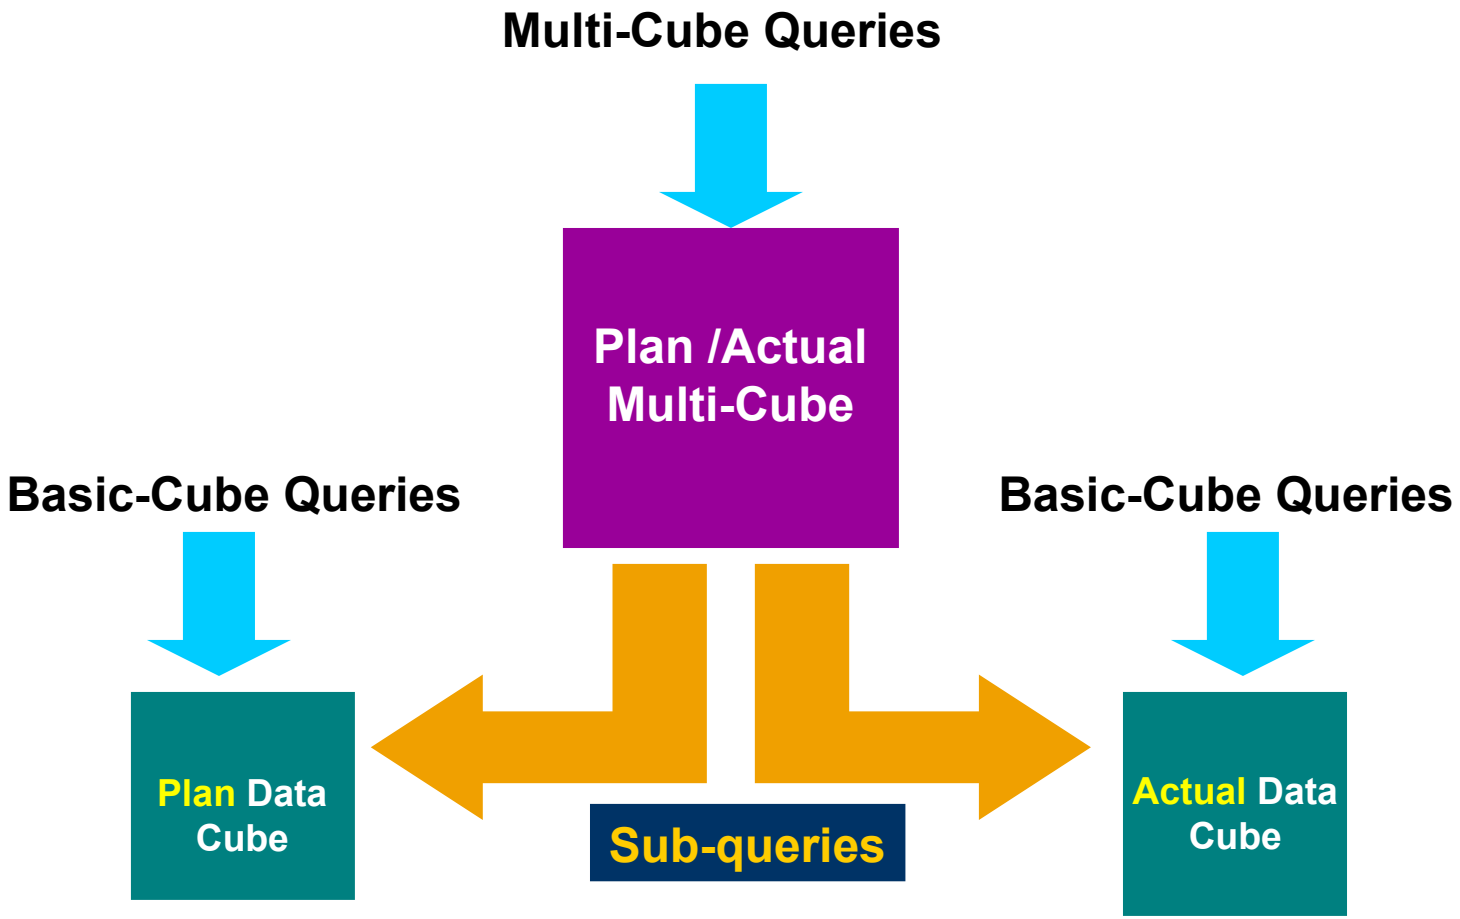

Multi-Cube Example

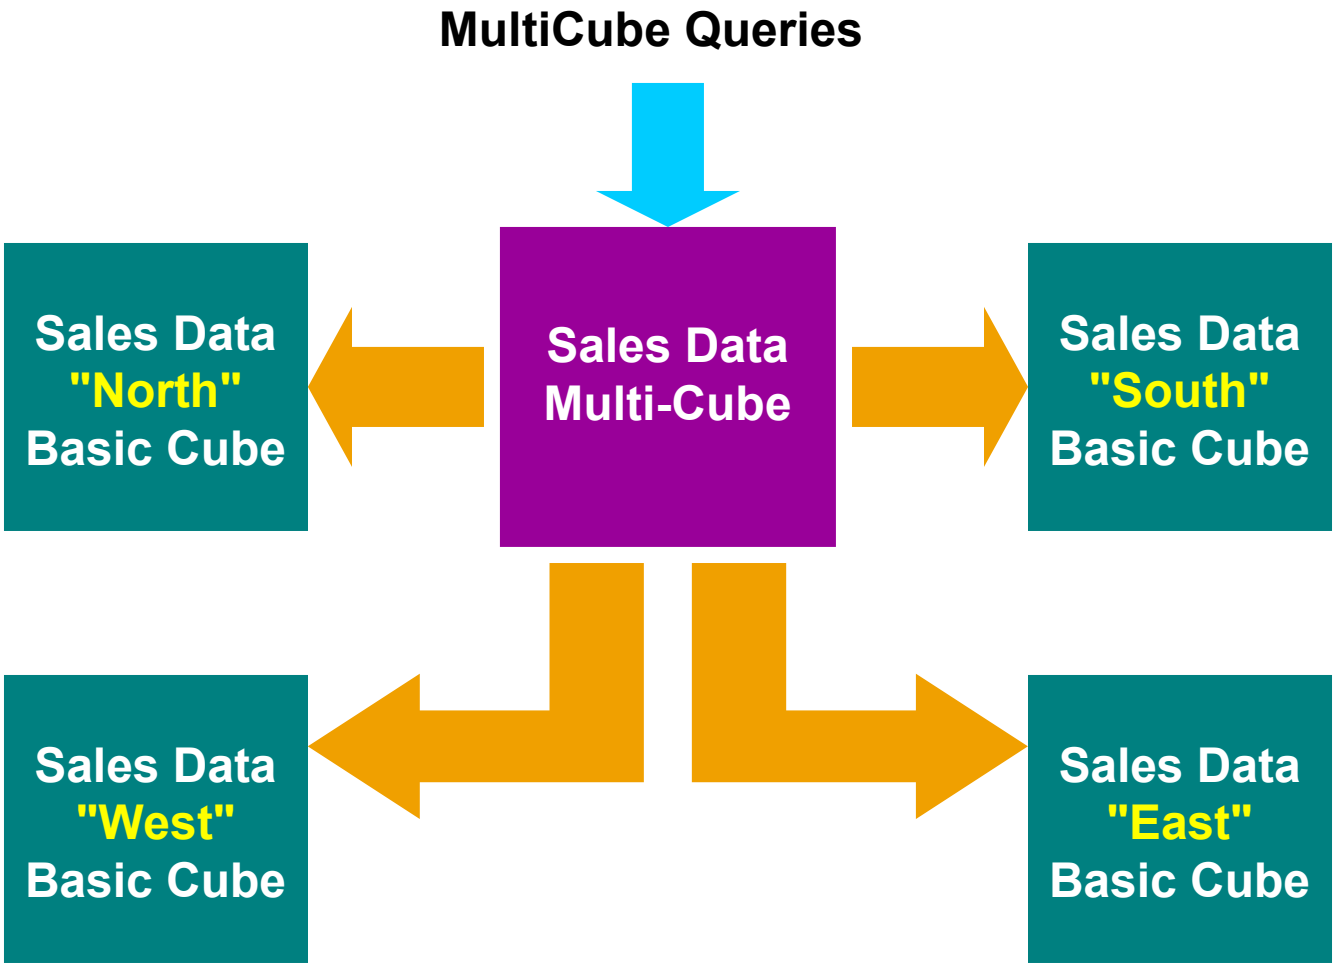

Query on MultiCubes - PARALLEL

Queries on MultiCubes are split up over the different Basis Cubes with parallel access to the InfoCube or maximal one aggregate per Cube

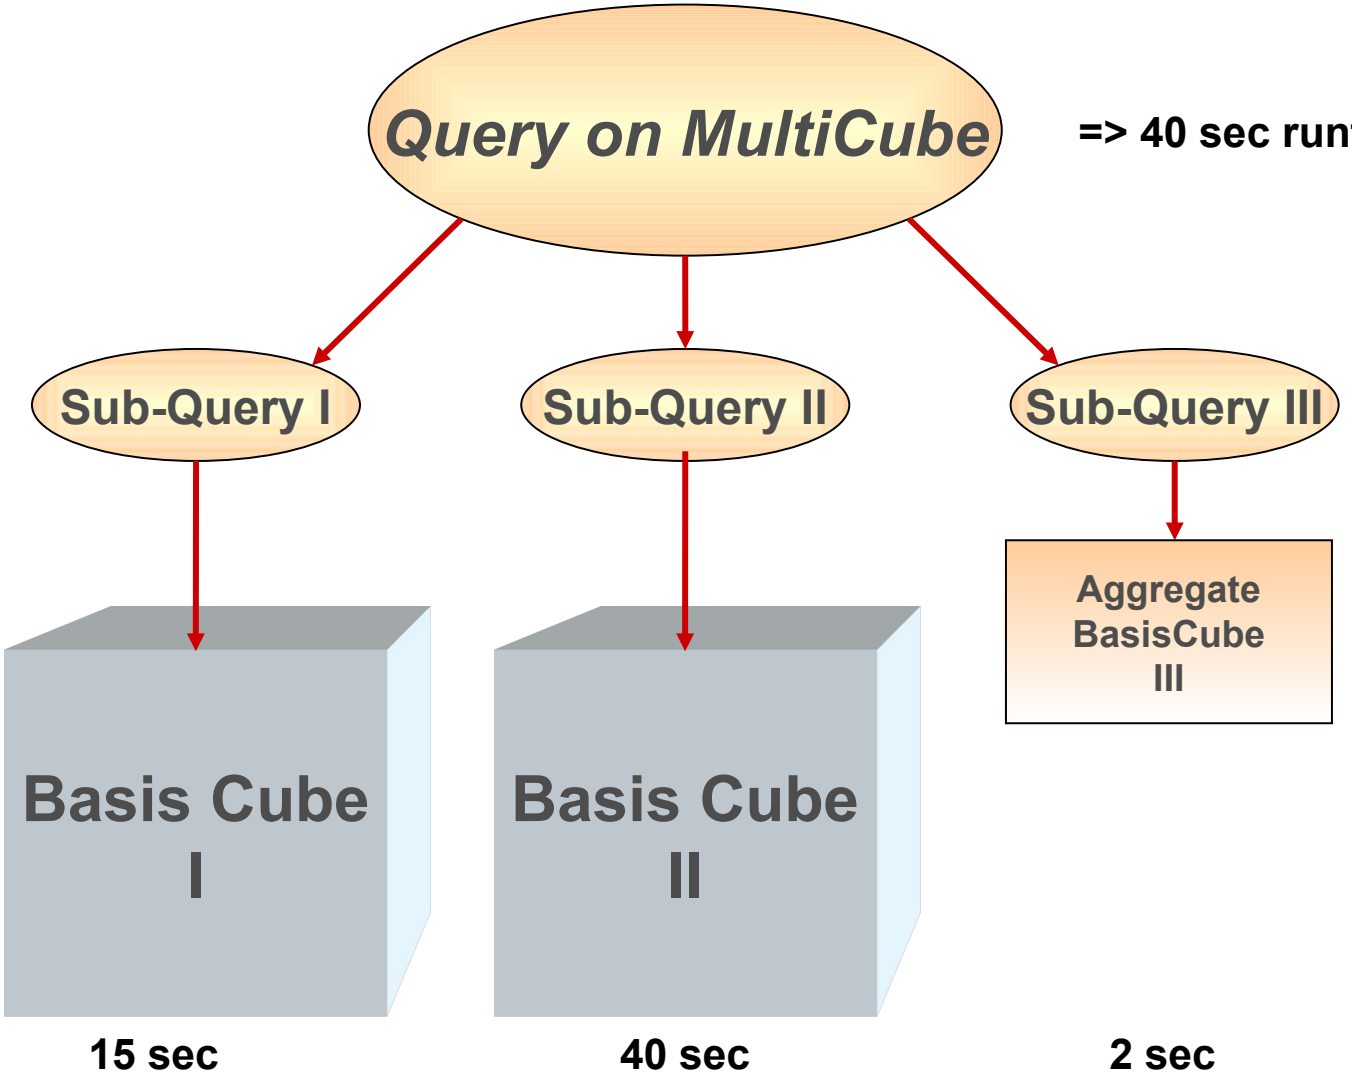

Query on MultiCubes - NOPARALLEL

NOPARALLEL execution is automatic via the data manager in BW 3.0

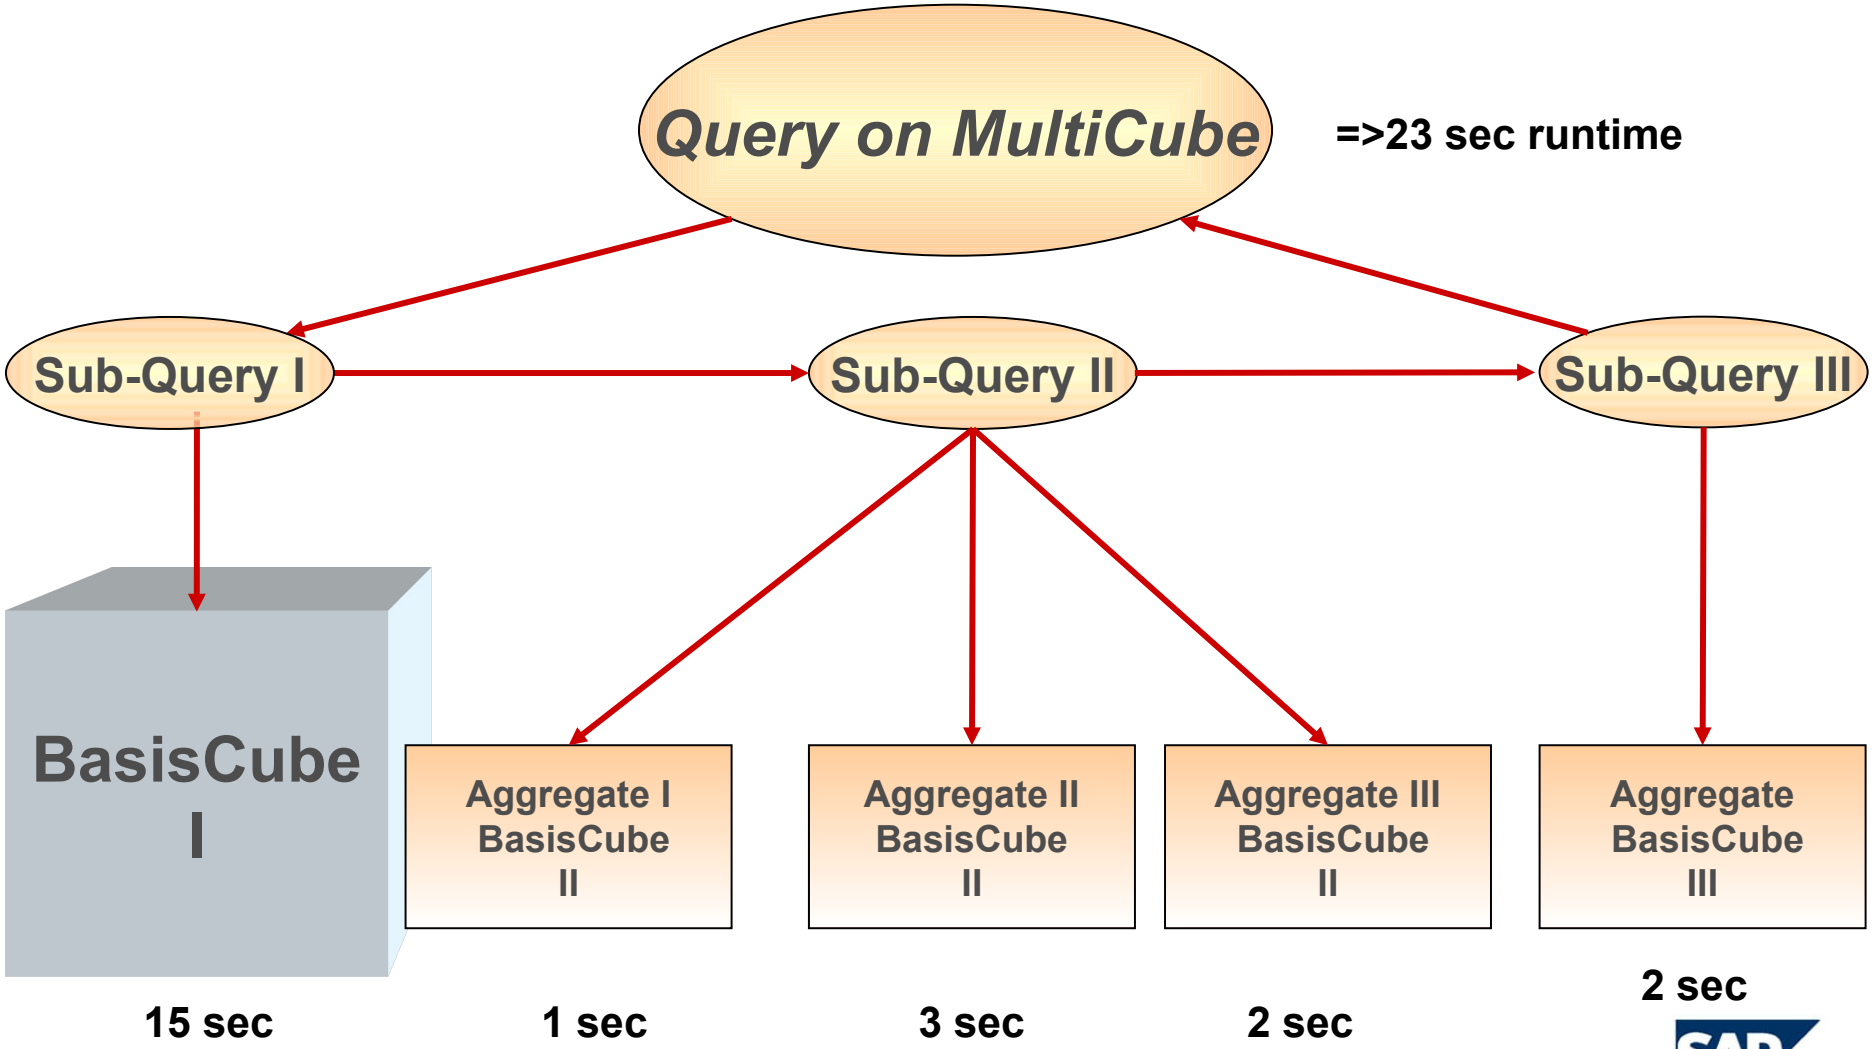

MultiProvider Performance / Scalability

Reporting

- Local queries (on each InfoProvider)
- Global queries (parallel execution)

Data load

- Independent (parallel) into individual InfoProviders
- Small total data volumes (less redundancy)

Database tables / views

- Smaller
- Less complex
- Less sparsely filled

Performance Tuning Concepts

- MultiProvider
- Partitioning
- Data Modeling and Line Item Dimension
- Web Reporting
- OLAP Query Cache

BW Partitioning: Platforms

Platforms with Table Partitioning

- IBM DB2/UDB*
- IBM DB2/390
- Informix
- Oracle

* hash partitioning

Platforms without Table Partitioning

- IBM DB2/400
- Microsoft SQL Server
- SAP DB

InfoCube Structure: Two Fact Table Concept

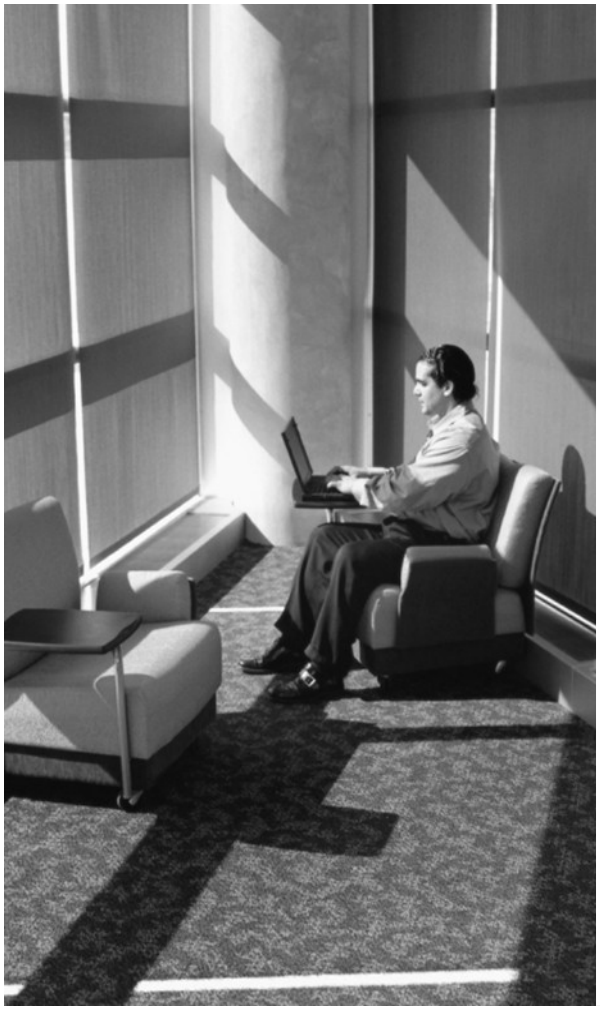

E fact table

- ... contains consolidated data
- ... is optimized for reading
- ... might be huge
- ... is partitioned by the user
- … cannot be partitioned once InfoCube contains data!

F fact table

- ... contains data on request level
- ... is optimized for writing / deleting
- ... should be small
- ... is partitioned by the system

InfoCube Range Partitioning: Example

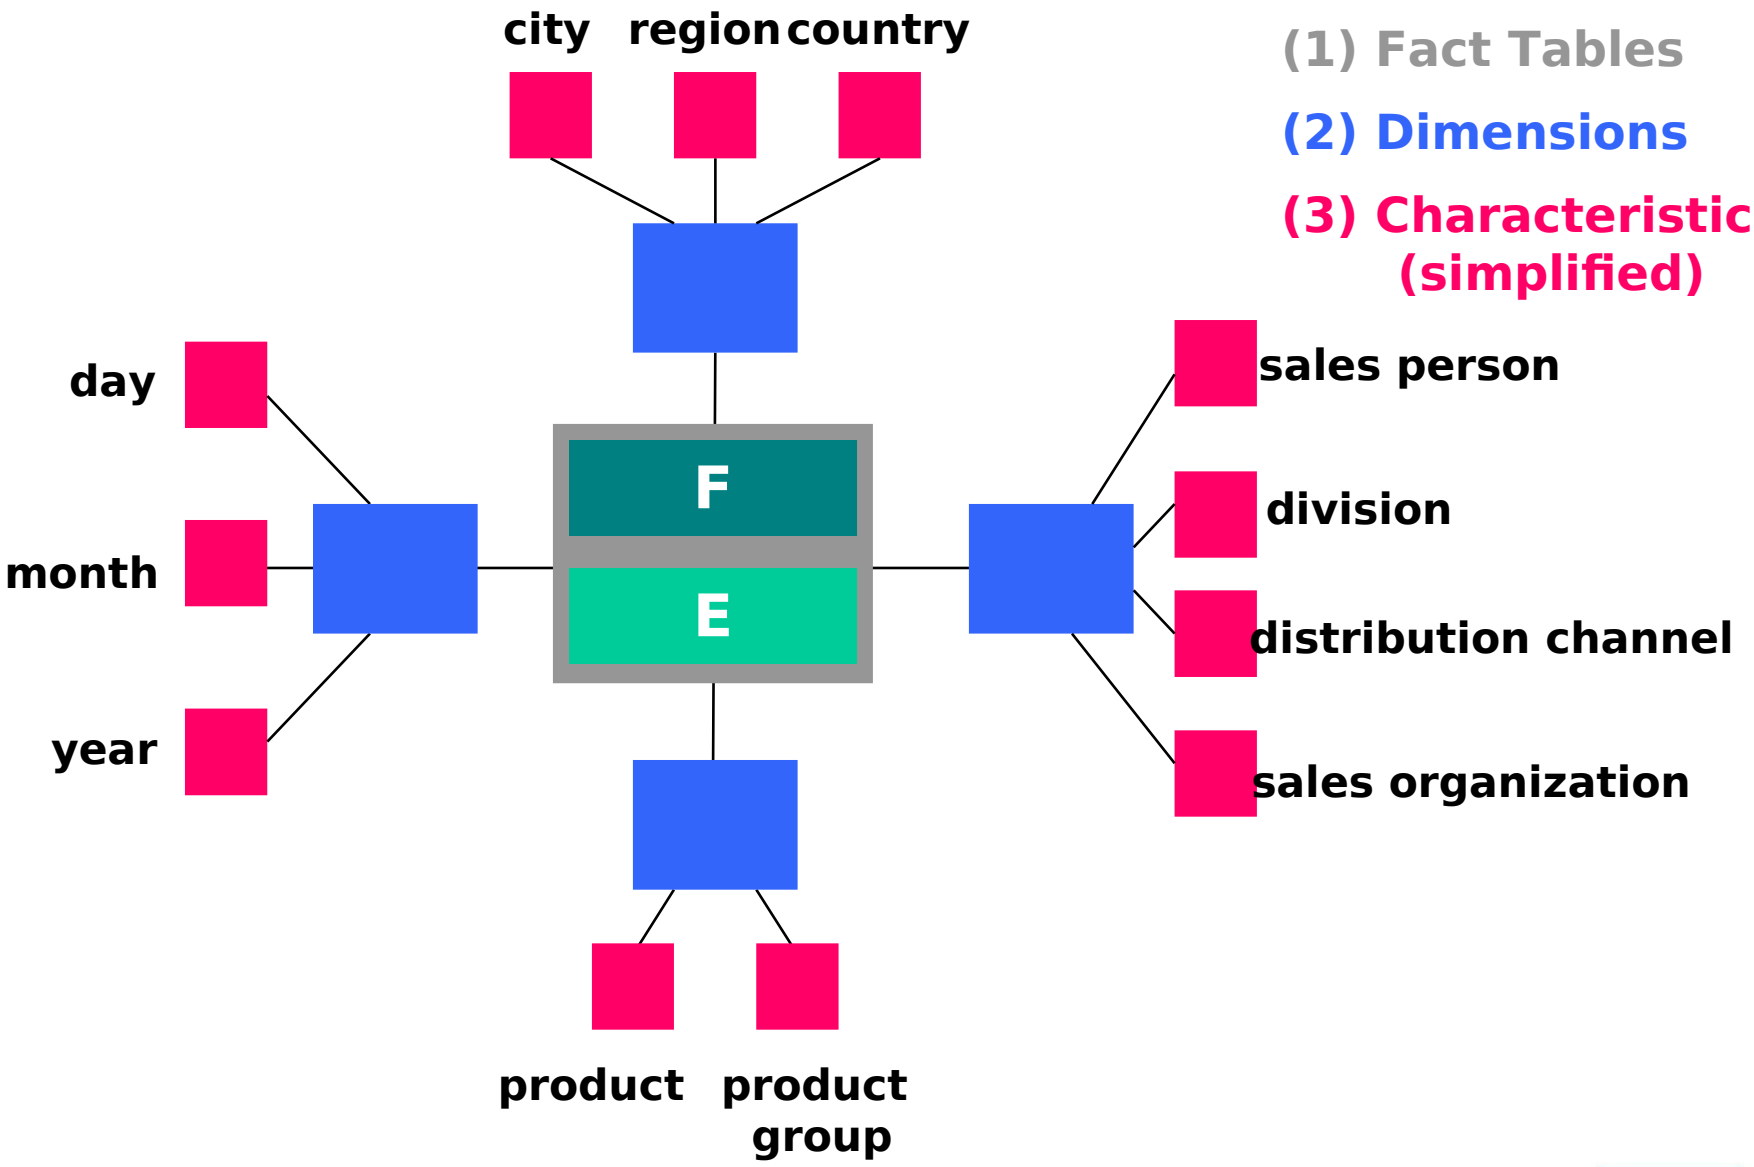

InfoCube Partitioning example: Two fact tables

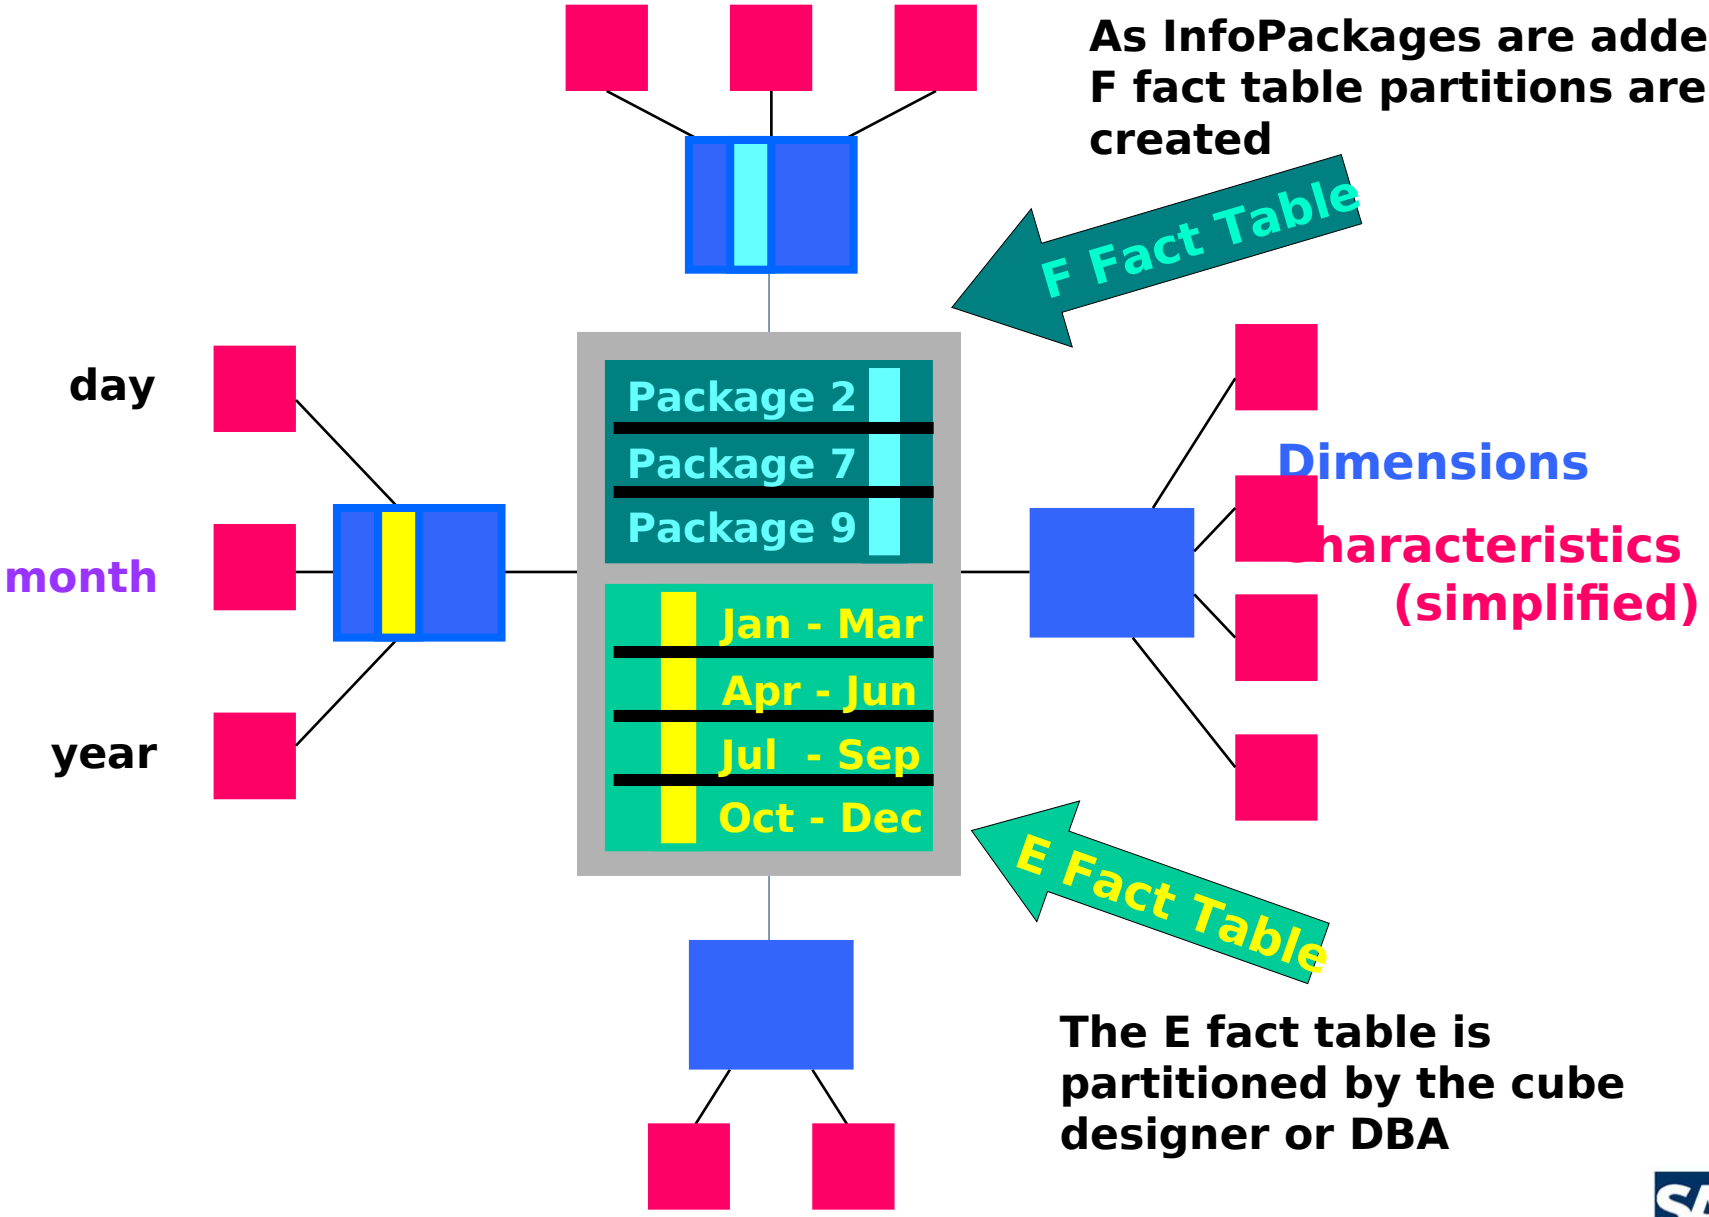

E Table Partitioning: Read a smaller dataset

Query Example

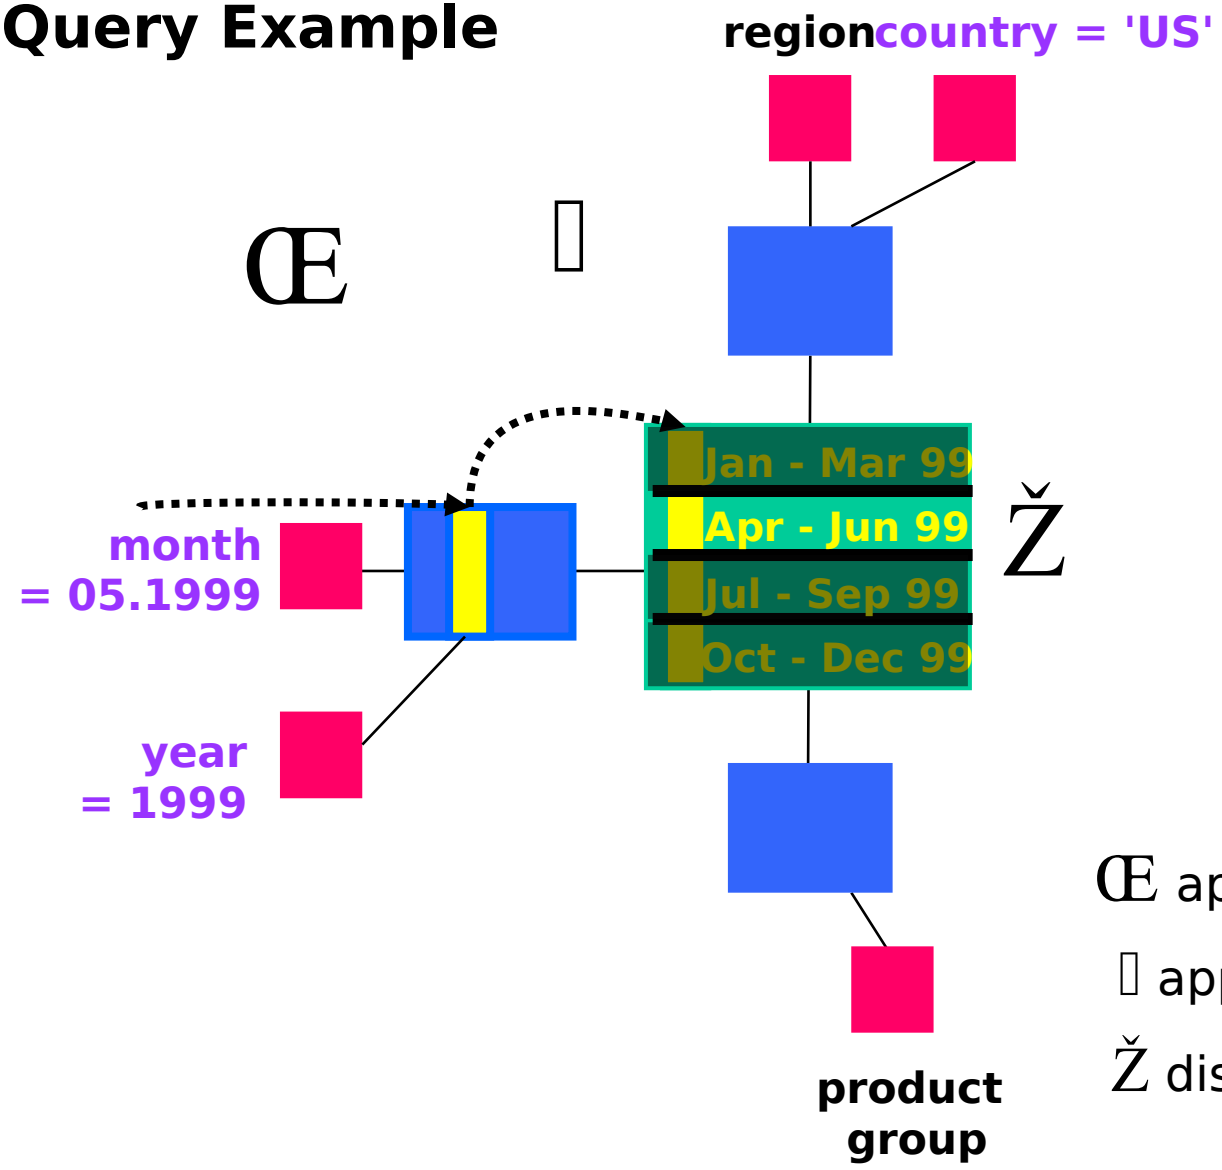

(1) E Fact Table

(2) Dimensions

(3) Characteristics (simplified)

Œ apply restriction to dimension

 apply restriction to fact table

Ž discard irrelevant partitions

User Interface: E Table Partitioning

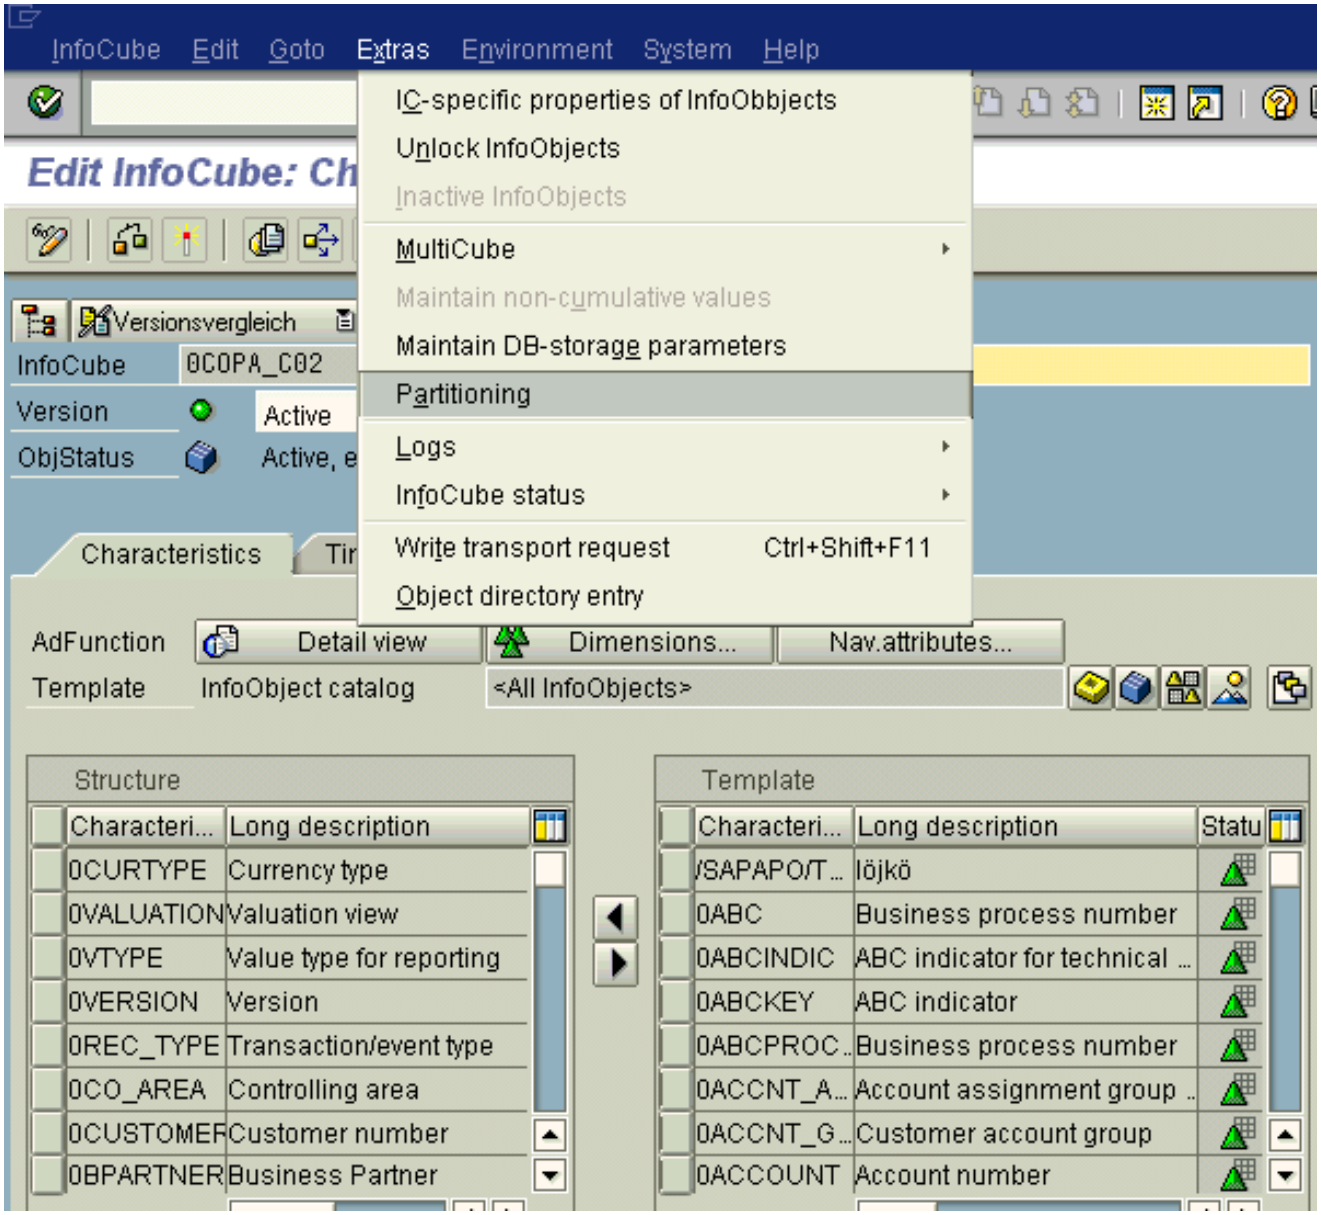

User Interface: E Table Partitioning: Slide 2

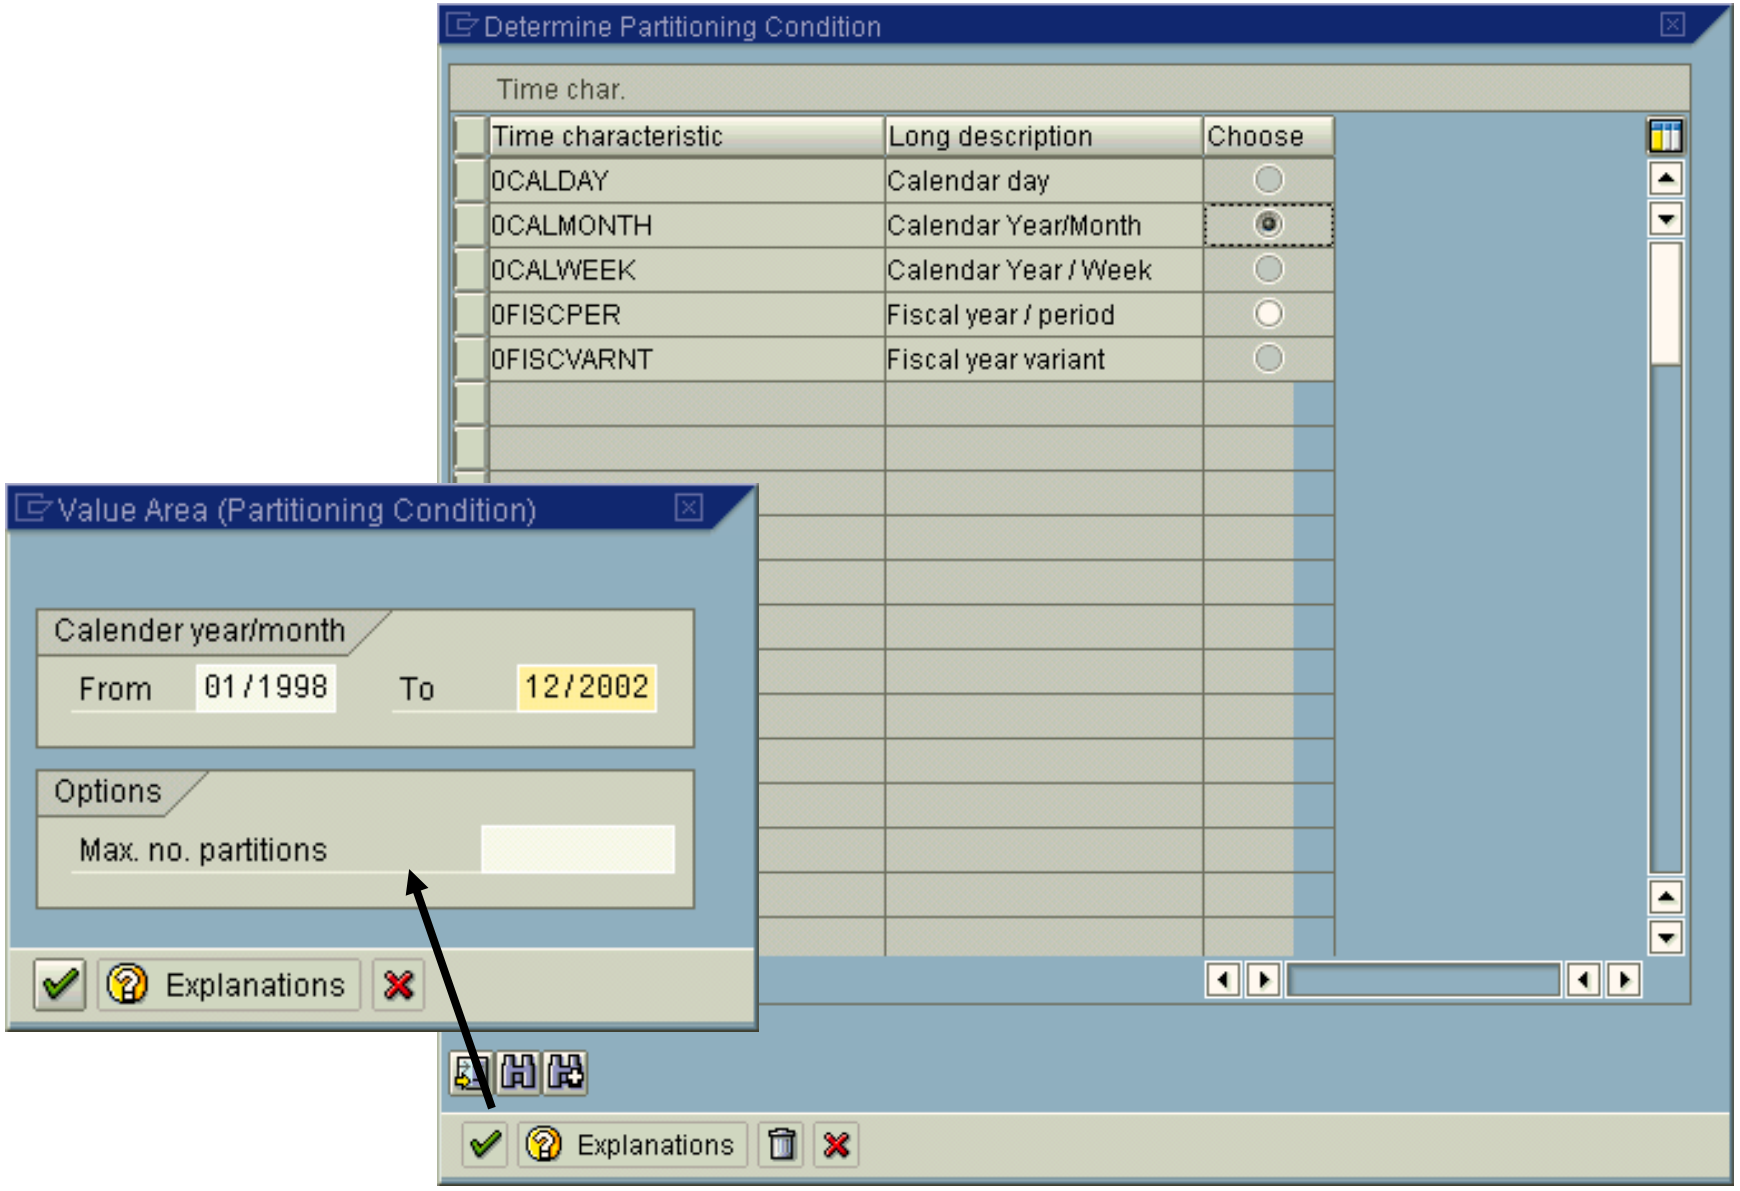

Partitioning/ MultiCube Example: A combination

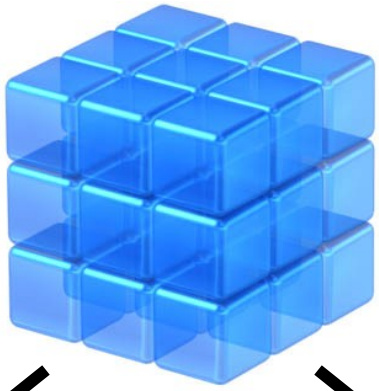

Multi-Cube

MultiCube

"North America"

„Europe"

Logical Partitioning

Parallel Processing

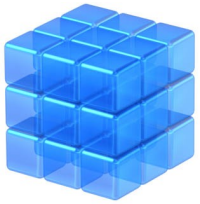

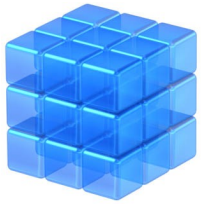

Basic-Cubes

Table Partitioning

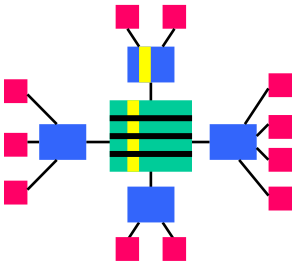

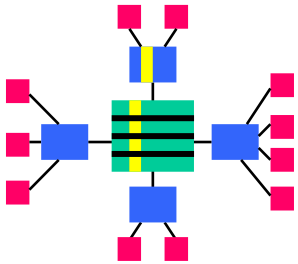

Tables

Reduces Number of Records to Read

Performance Tuning Concepts

- MultiProvider
- Partitioning
- **Data Modeling and Line Item Dimension**
- Web Reporting
- OLAP Query Cache

SAP BW InfoCube Data Model

- An InfoCube is designed, or “modeled” to meet a set of business reporting requirements.
- Modeling is the process by which reporting requirements are structured into an object with the facts and characteristics that will meet the reporting needs.
- Characteristics are structured together in related branches called “Dimensions” .
- The key figures form the “Facts”.
- The configuration of dimension tables in relation to the fact table results in what is known as the “star schema”

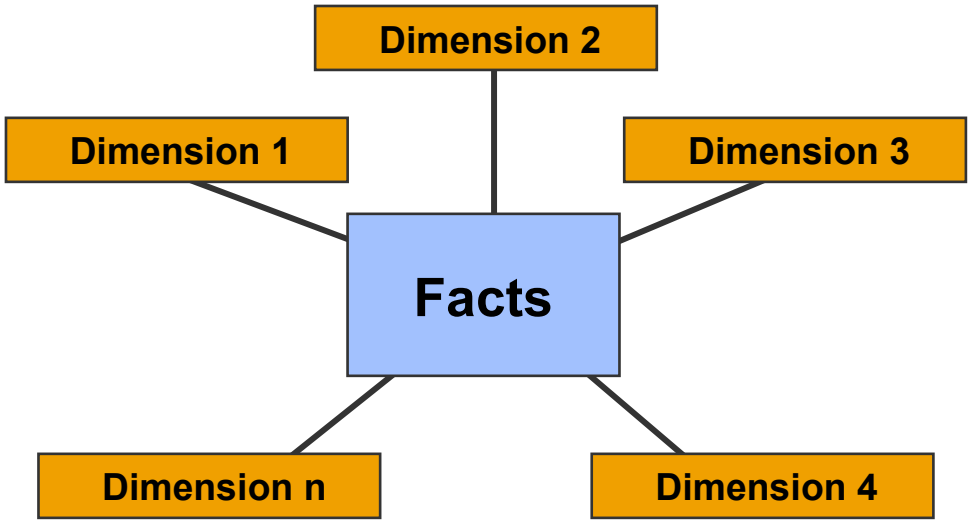

Characteristics of Dimensions

- Dimensions are groupings of related characteristic attributes.
- The keys of the dimension tables are foreign keys in the fact table. The key is a "DIMID" used to map related values.
- A dimension table contains a primary key (DIMID), characteristic values , and SID values to link to master data tables.

Customer dimension

C	Customer #	Region	…
	13970522	West	...

Product Service Line dimension

P	Product #	Product group	…
	2101004	Displays	...

Time dimension

T	Period	Fiscal year	…
	10	1997	...

Example: Multidimensional Schema in BW

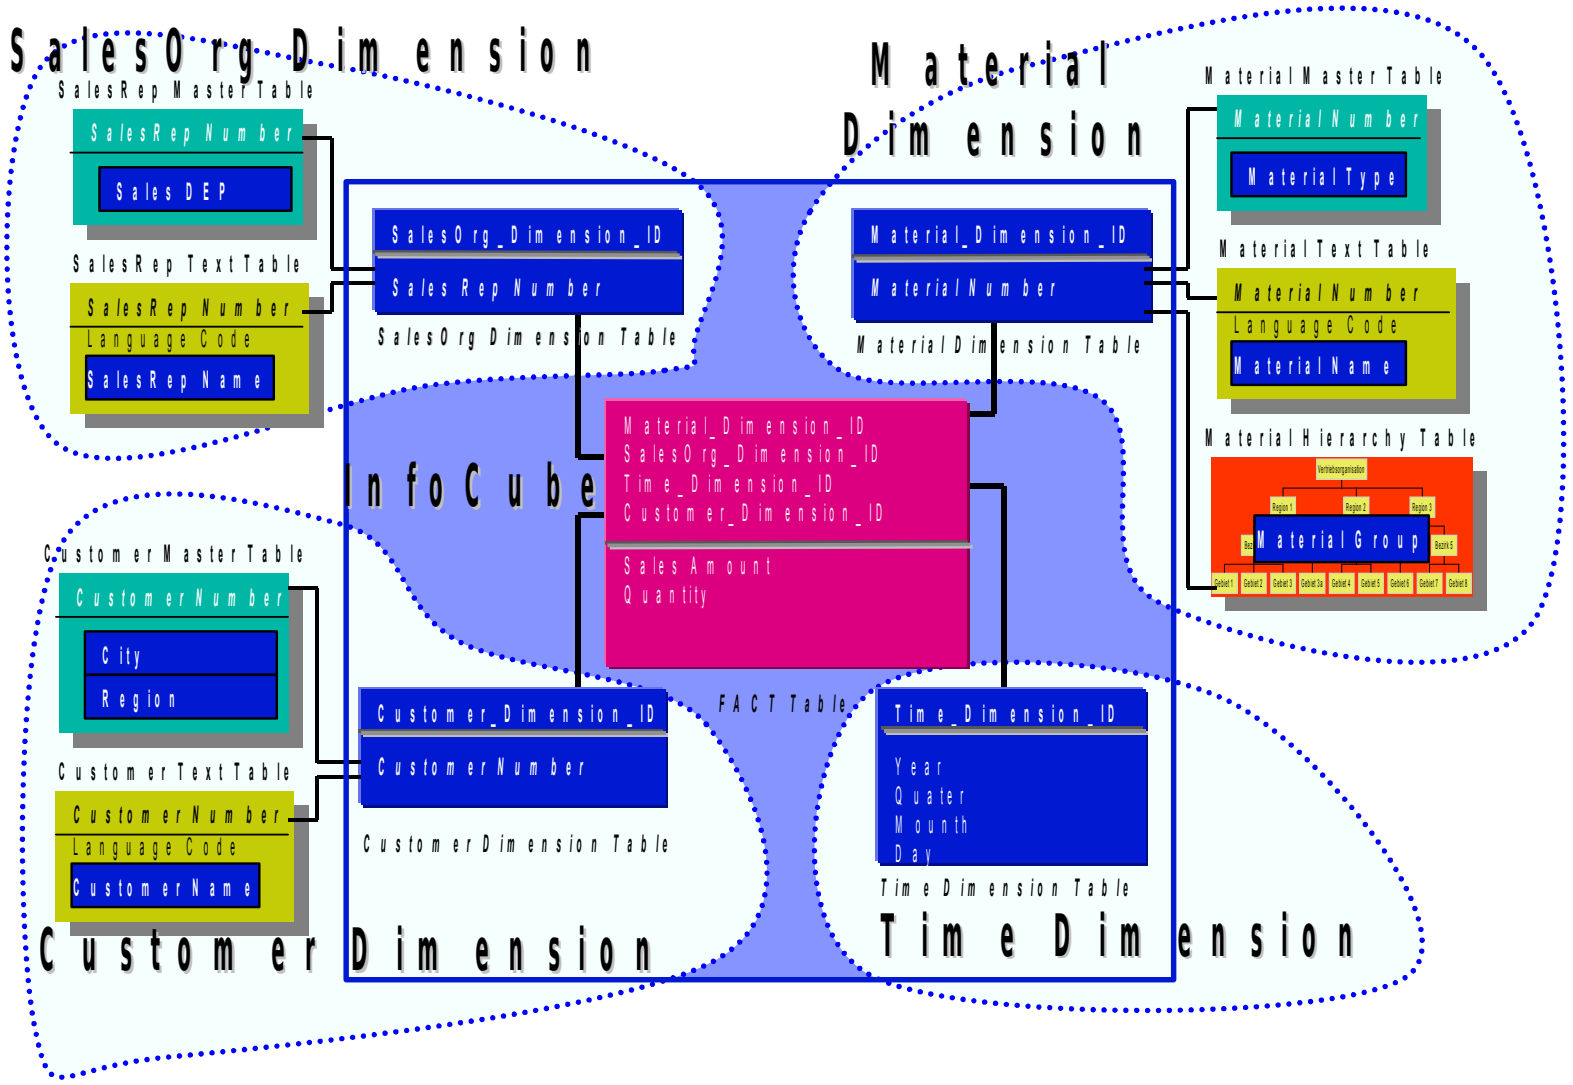

Query Execution Example: Table Joins

Table : /BIC/FIUSALES
Displayed fields: 8 of 8 Fixed columns: 6 List width 0250

Fact Table

KEY_IUSALESP	KEY_IUSALEST	KEY_IUSALESU	KEY_IUSALES1	KEY_IUSALES2	KEY_IUSALES3	/BIC/IUPROFIT	/BIC/IUQUAN
5	731	0	29	62	62	982-	820
5	732	0	30	63	63	1.051	140
5	732	0	31	66	89	3.389	390
5	732	0	41	79	82	31-	140
5	732	0	46	70	79	2.144-	540

Dimension 1

Table : /BIC/DIUSALES1
Displayed fields: 4 of 4 Fixed co

DIMID	SID_IUCITY	SID_IUCOUNTRY	SID_IUREGION
0	0	0	0
29	2	2	2
30	3	3	3
31	4	4	4
32	5	5	5

Table : /BIC/SIUPOP
Displayed fields: 5 of 5 Fixed columns:

/BIC/IUPOP	SID	CHCKFL	DATAFL	INCFL
000000000	0	X	X	X
000500000	500.000		X	X
000600000	600.000		X	X
000800000	800.000		X	X
001000000	1.000.000		X	X

S (Population)

X (City)

Table : /BIC/XIUCITY
Displayed fields: 5 of 5 Fixed columns:

SID	OBJVERS	/BIC/IUCITY	CHANGED	S__IUPOP
0	A			0
2	A	LYON		2.000.000
4	A	MONTREAL		2.000.000
3	A	NEWCASTLE		600.000
5	A	SAN DIEGO		2.150.000
1	A	WALLDORF		0

Line Item Dimensions

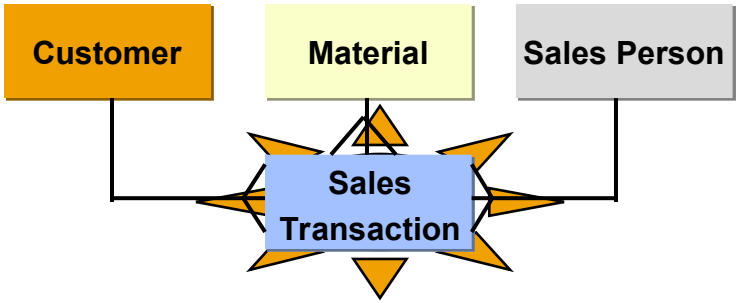

- Line-item or *degenerate* dimensions arise in nearly every case where the granularity of the fact table represents an actual working document like an order number, invoice number, sequence number, etc.
- Why are line-item dimensions a concern?
  - DIMID column values are often high cardinality (slow read times if bitmap index is used)
  - Join of large dimension table with large fact table is a big overhead for reporting response times.

Line-Item Dimension: The Star Schema

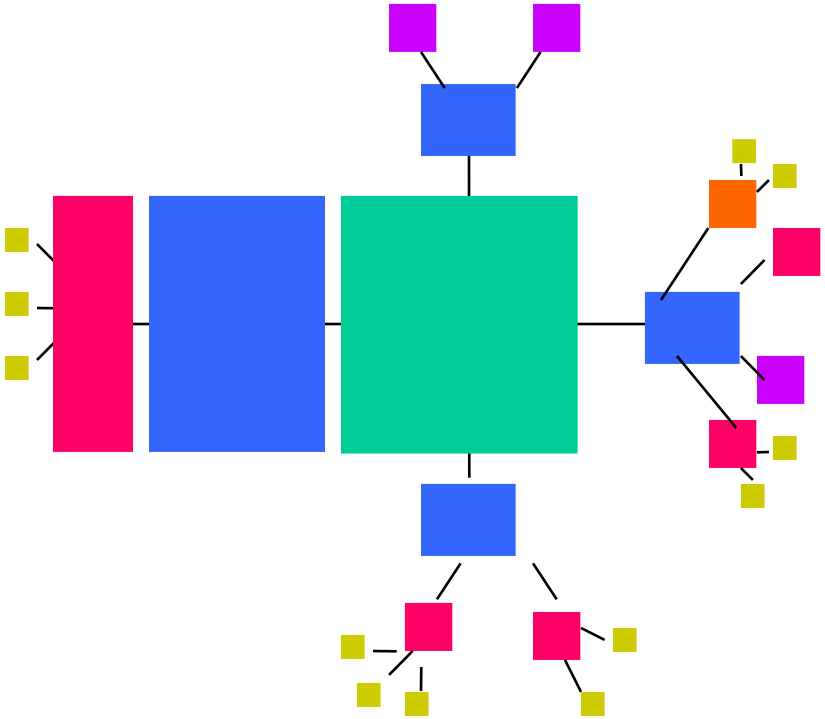

(1) Fact Table

(2) Dimension

(3) time-independent-SID
time-dependent-SID
traditional SID Char

(4) SID Attr

Without line-item dimension designated

With line-item dimension

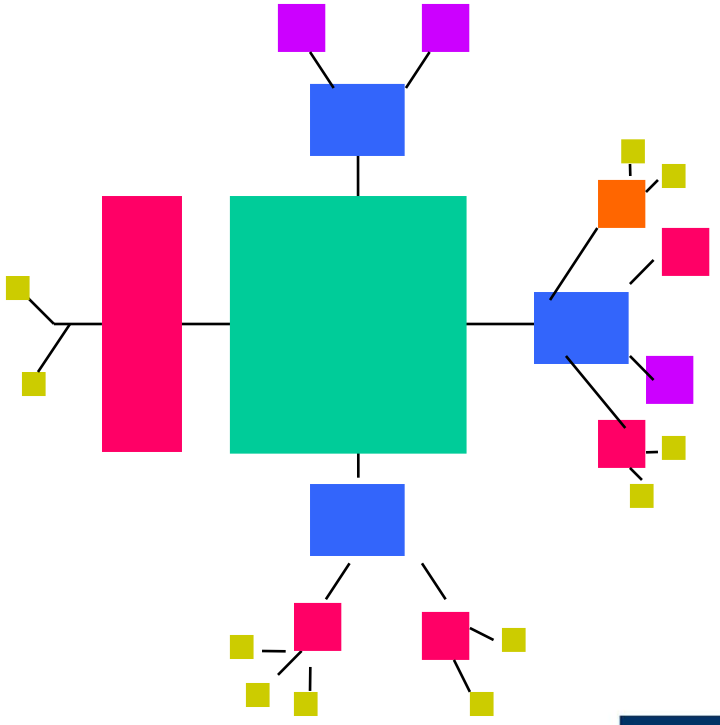

Where to set Line Item in InfoCube

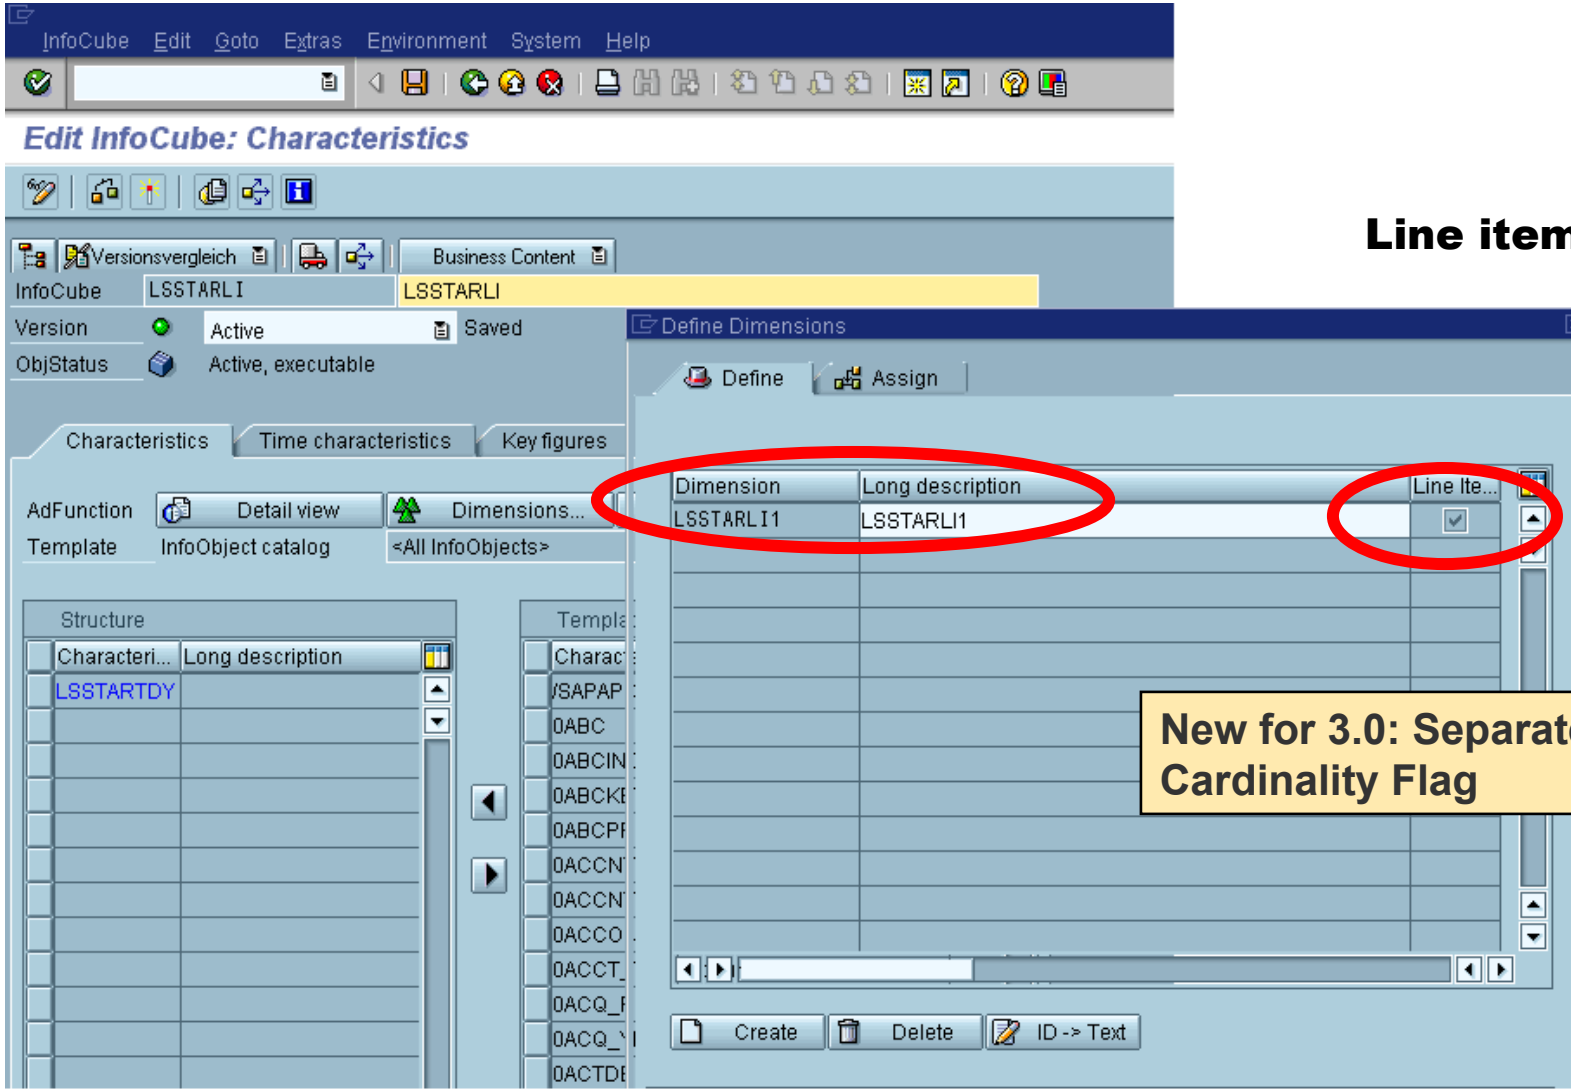

Fact table with line item dimension

Transp. table	/BIC/F100061	Active
Short text	KWFAB5	

Attributes | Delivery and Maintenance | Fields | Entry help/check | Currency/Quantity Field

Srch help | Built-in type

Field	Key	Initi...	Data element	DTyp	Length	Dec.p...	Short text
KEY_100061P	☑	☑	RSDIMID	INT4	10	0	Dimension Table Key
KEY_100061T	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_KWFAB5U	☑	☑	RSDIMID	INT4	10	0	Dimension Table Key
KEY_1000611	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000612	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000613	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000614	☑	☑	RSSID	INT4	10	0	Master data ID
KEY_1000615	☑	☑	RSSID	INT4	10	0	Master data ID
CRMEM_QTY	☐	☐	/BI0/OICRMEM_QTY	QUAN	17	3	Credit memos quantity

Performance Tuning Concepts

- MultiProvider
- Partitioning
- Data Modeling and Line Item Dimension
- **Web Reporting**
- OLAP Query Cache

Different types of users in OLAP reporting

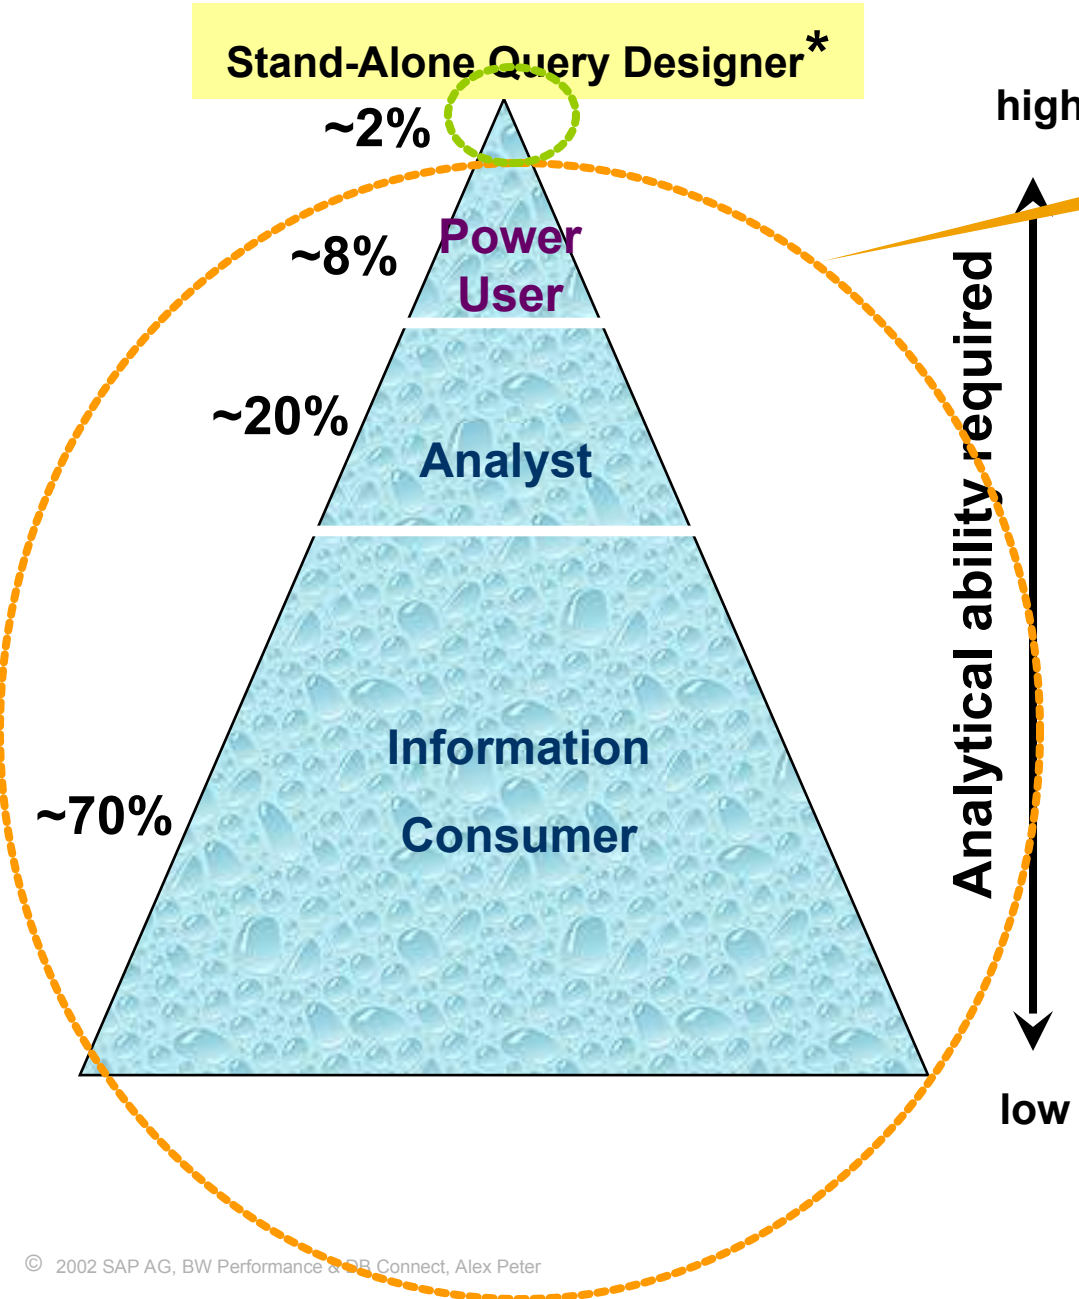

~98% Potential web reporting users

- Report creation
- Ad-hoc reporting
- Multi-dimensional analysis
- Self-explanatory, predefined navigational paths
- Predefined data collections
- Static reports

***Only needed for Global Report Creation**

Performance: Web Reporting VS. BEX Analyzer

Query activity using the Business Explorer Analyzer requires many more “round trips” for data packets than the Browser

- Communication between the GUI and BW
- Multiple round trips reduces efficiency of data transfer
- Some communication needed for screen controls, MS-interface, etc.

Browser-based web reporting, using http, represents a significant gain in query response time!

- The number of round trips is significantly less than the BEx analyzer!
- HTML compression, MIME compression, MIME caching all possible with Web Reporting only

Data Compression and MIME Compression

Data Compression is a benefit to web reporting performance.

- Data volume is reduced, consuming less bandwidth, enabling faster throughput

SAP BW development recommends enabling compression for the ICF services used by BW: BEX Service, MIME Service, others.

Implement the latest BW 3.0B support packages: SP 6 and SP 7 contain important performance improvements

MIME (images – company logo, etc) objects are also compressed

- Images can be large, and compression speeds data transfer!

WAS ICF: ICF Services

Transaction SICF

- services defines URL for HTTP handlers
- Default user settings for public sites
- Default client, language
- definition of Aliases

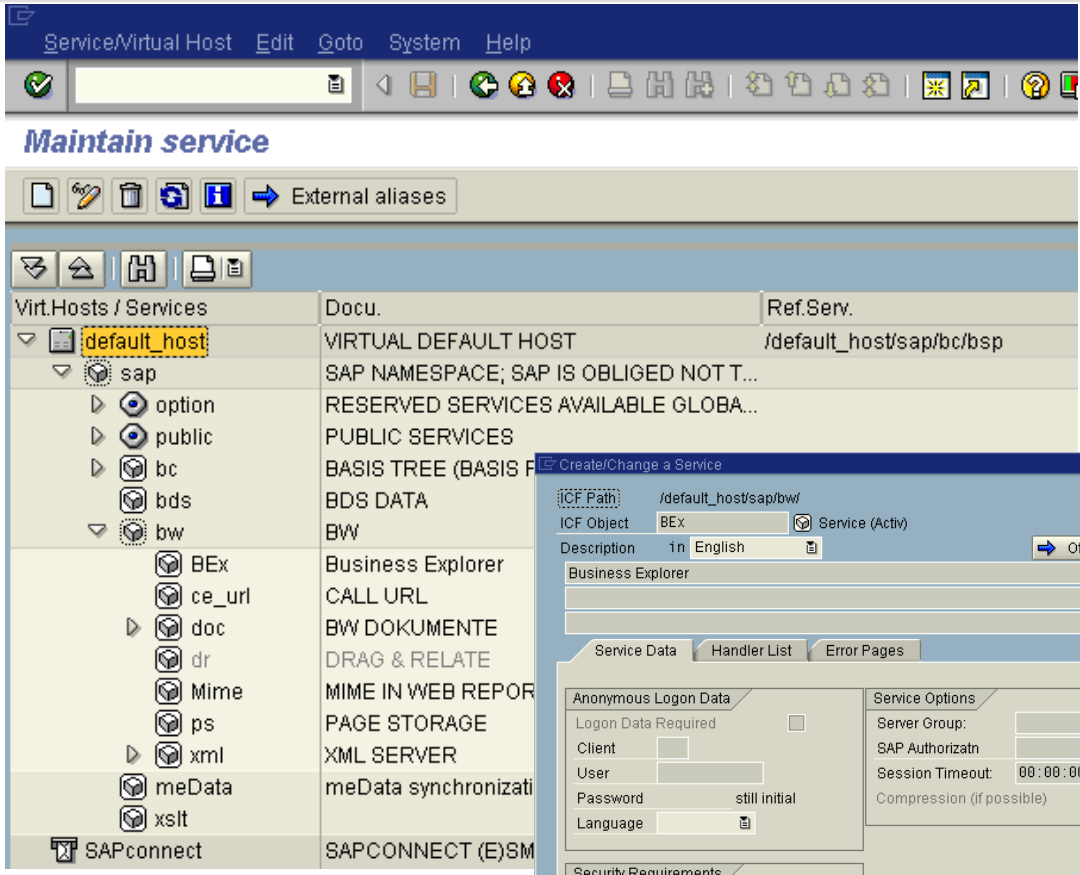

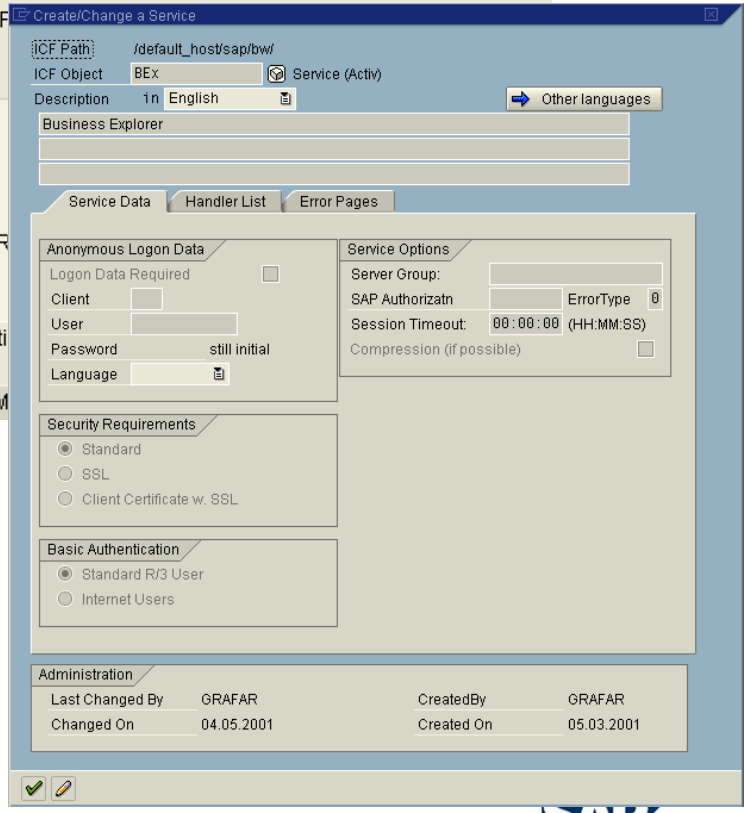

- Individual ICF Services for BSP and HTTP Applications
- BW services delivered as standard

The BEx Service is utilized for BW Web reporting

ICF BEX Service: Enable Compression (trans SICF)

The BEx Service and MIME service: Enable Compression!

See OSS notes:
550669 & 553084
Enable compression for the BW-relevant ICF services:

sap->BW->BEx
sap->BW->Mime
sap->BW->doc->browser
sap->BW->doc->hier
sap->BW->doc->mast

sap->BW->doc->meta
sap->BW->doc->metadata
sap->BW->doc->tmpl
sap->BW->doc->tran

Also, "http 1.1" should be configured in the browser settings

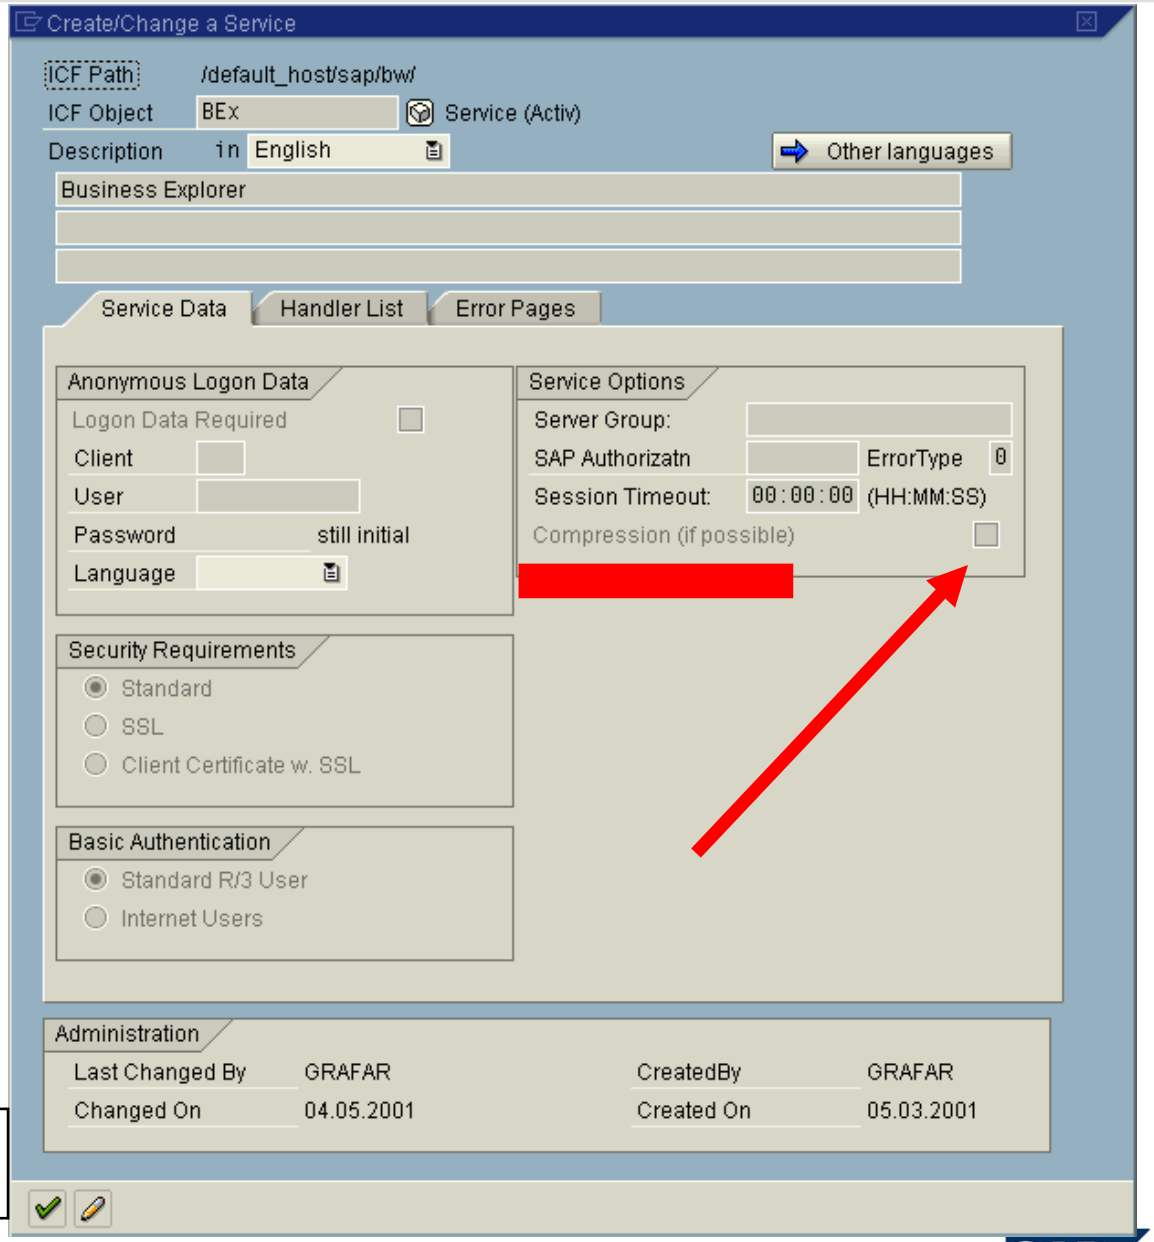

Browser Cache of MIME objects

MIME objects are now cached in the browser, see OSS note 561792 - BW 3.0B SP 7

Performance Gain: Static images (i.e. company logo, etc) are not reloaded (from the server MIME cache) every time a query result set is rendered in the browser – less data to send!

Note:

The ICM will attempt to force a new version of the image into the browser's cache periodically

- based on an instance profile parameter setting.
- A performance gain can be achieved by extending the time MIMEs will sit in the browser cache before the ICM reloads them from the MIME server cache.

The related profile parameter is 'icm/HTTP/server_cache_0/expiration'.

Browser Settings for Client Caching of MIMEs

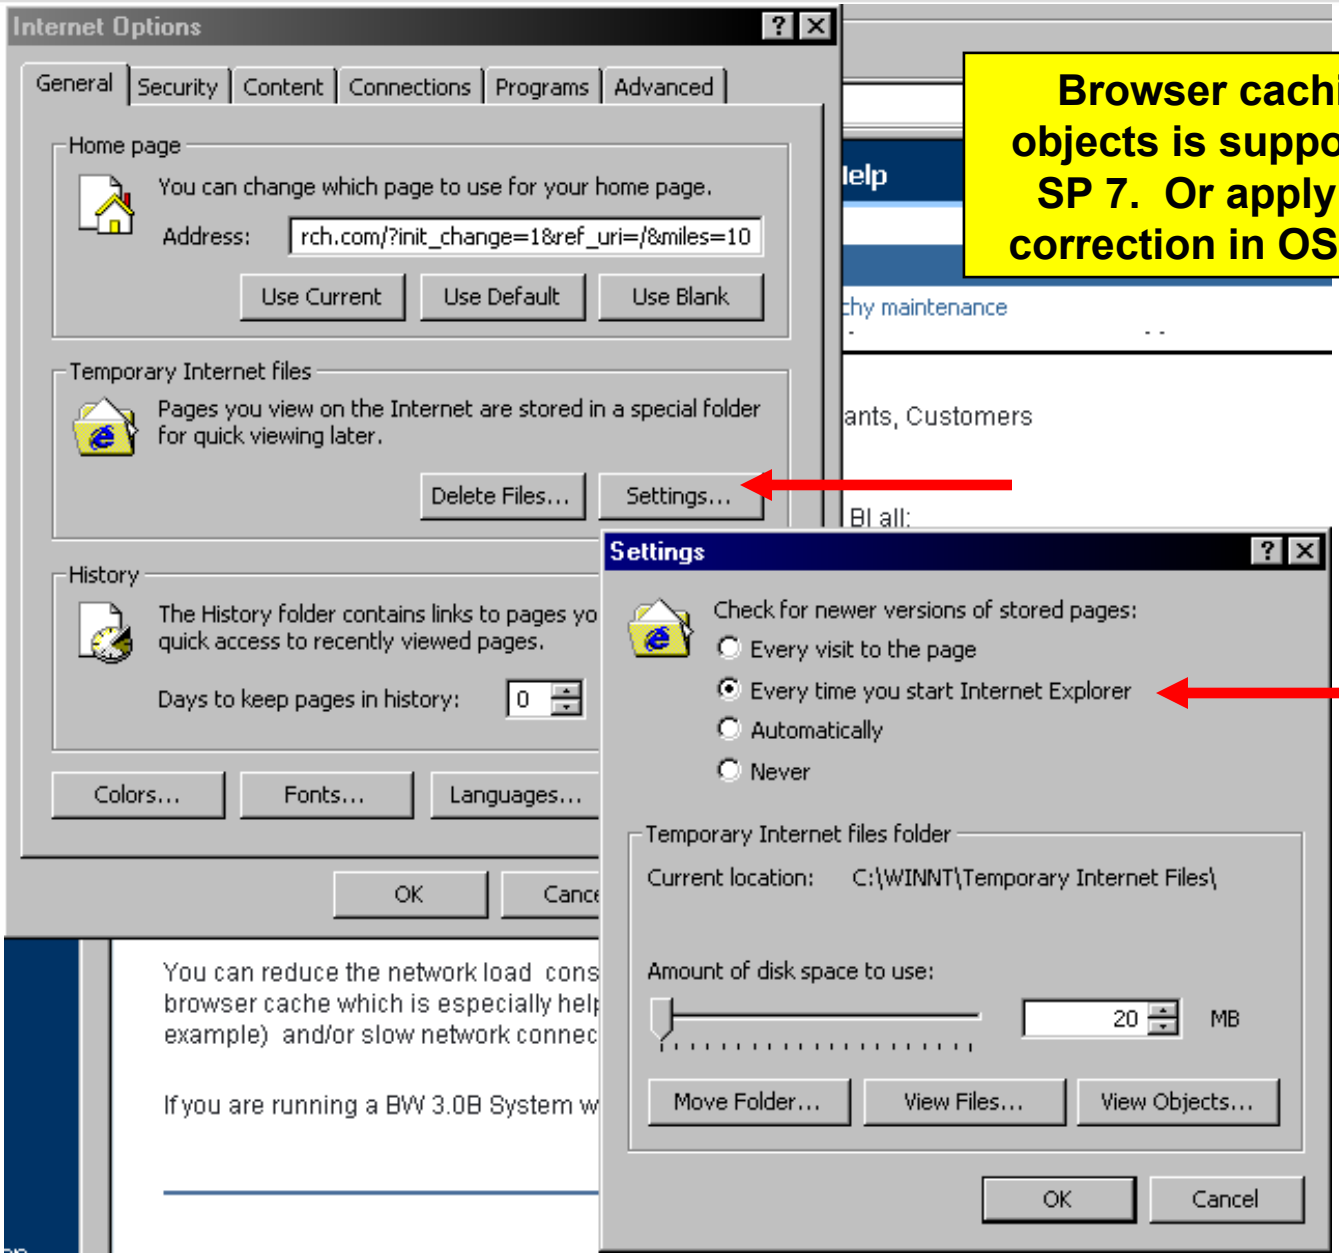

Browser caching of MIME objects is supported as of BW SP 7. Or apply the advance correction in OSS note 561792

Reporting Agent: Overview

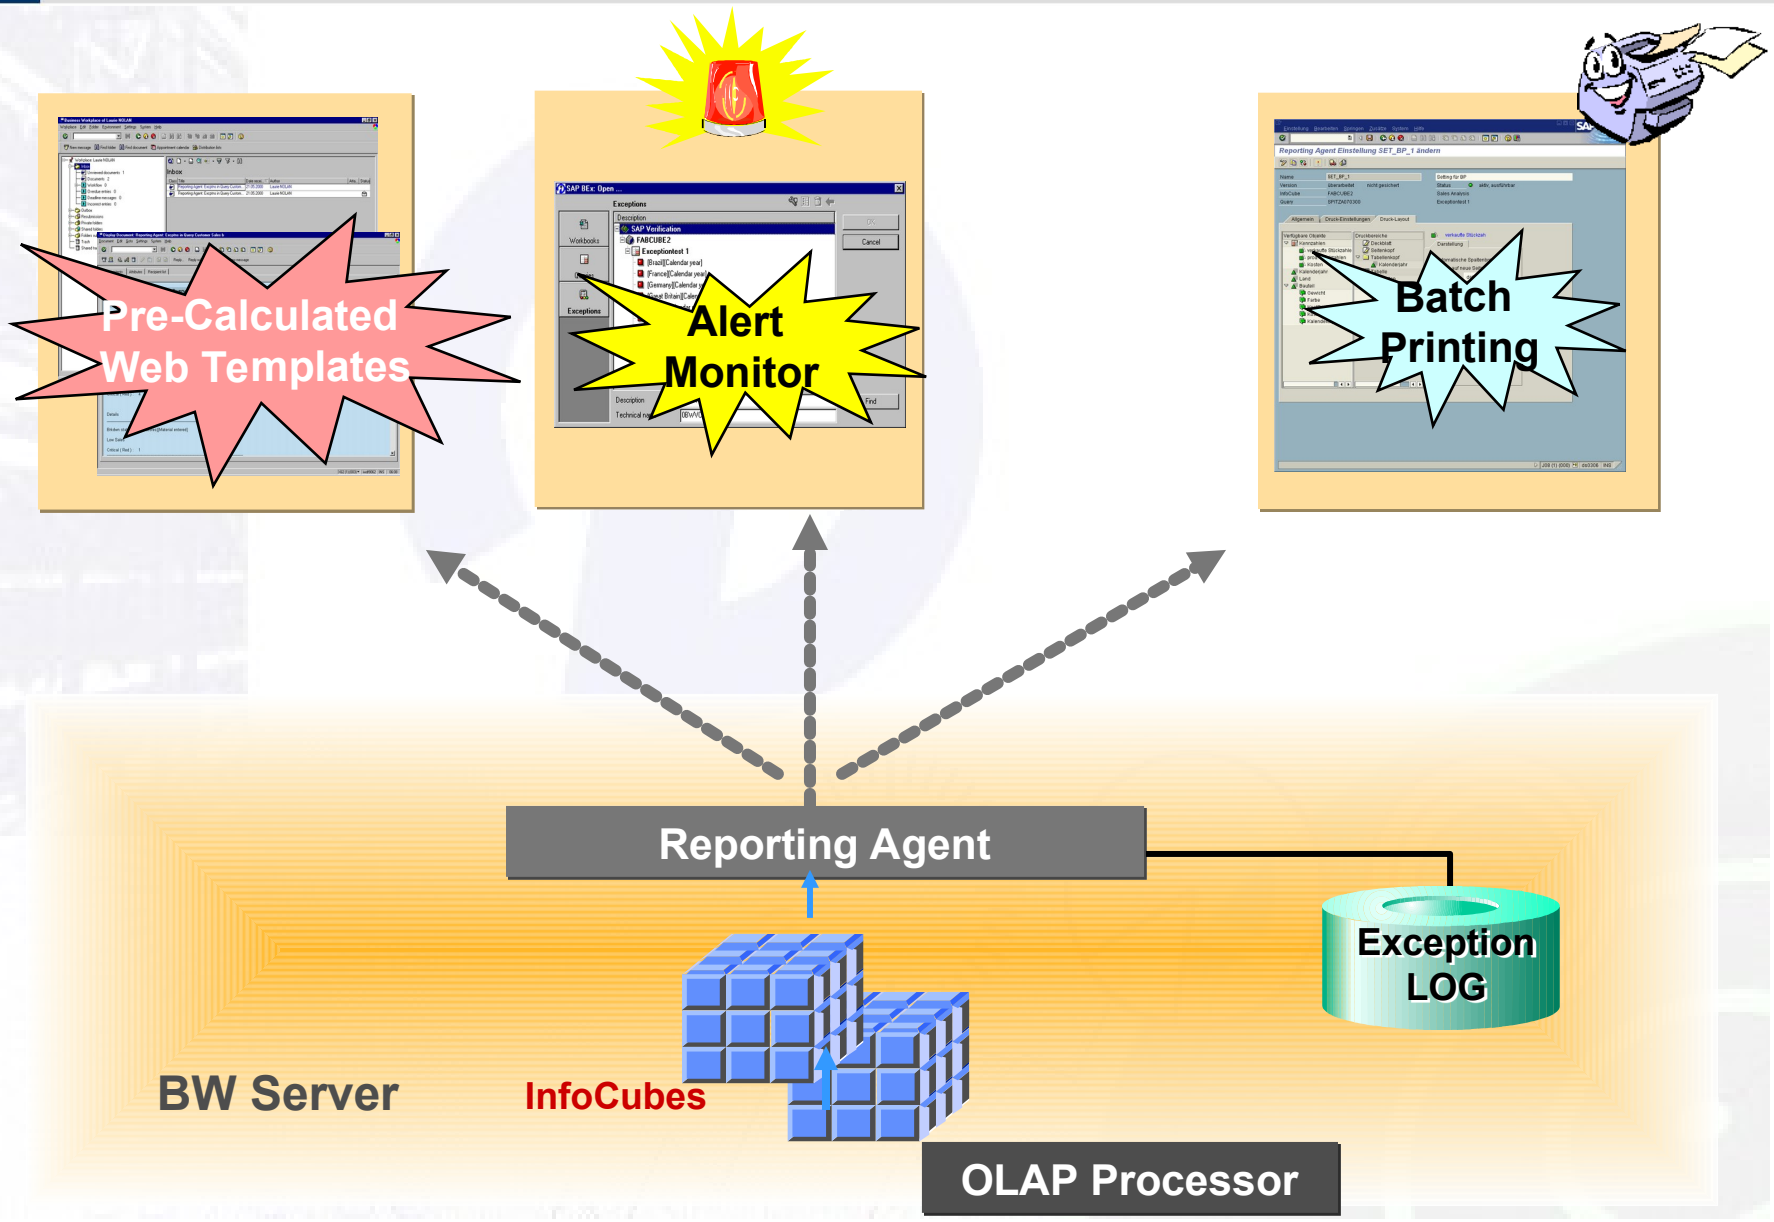

Reporting Agent: Precalculated Template

The Reporting Agent has performance enhancing functionality!

Precalculation is a set of techniques where you can distribute the workload of running the report to off-peak hours, and have the report result set ready for very fast access to the data

Response-time performance is greatly improved to the end user

…And overall system workload is reduced during peak hours

Pre-Calculating Web Templates – Business Scenario

Scope of Application

- Web reports that are accessed by many users.
- Web reports that are static, or involve limited navigation.
- Web reports that should be made available offline.

Goal

- Reduce server load significantly.
- Provide faster data access.
- Re-use data that goes into many web reports.

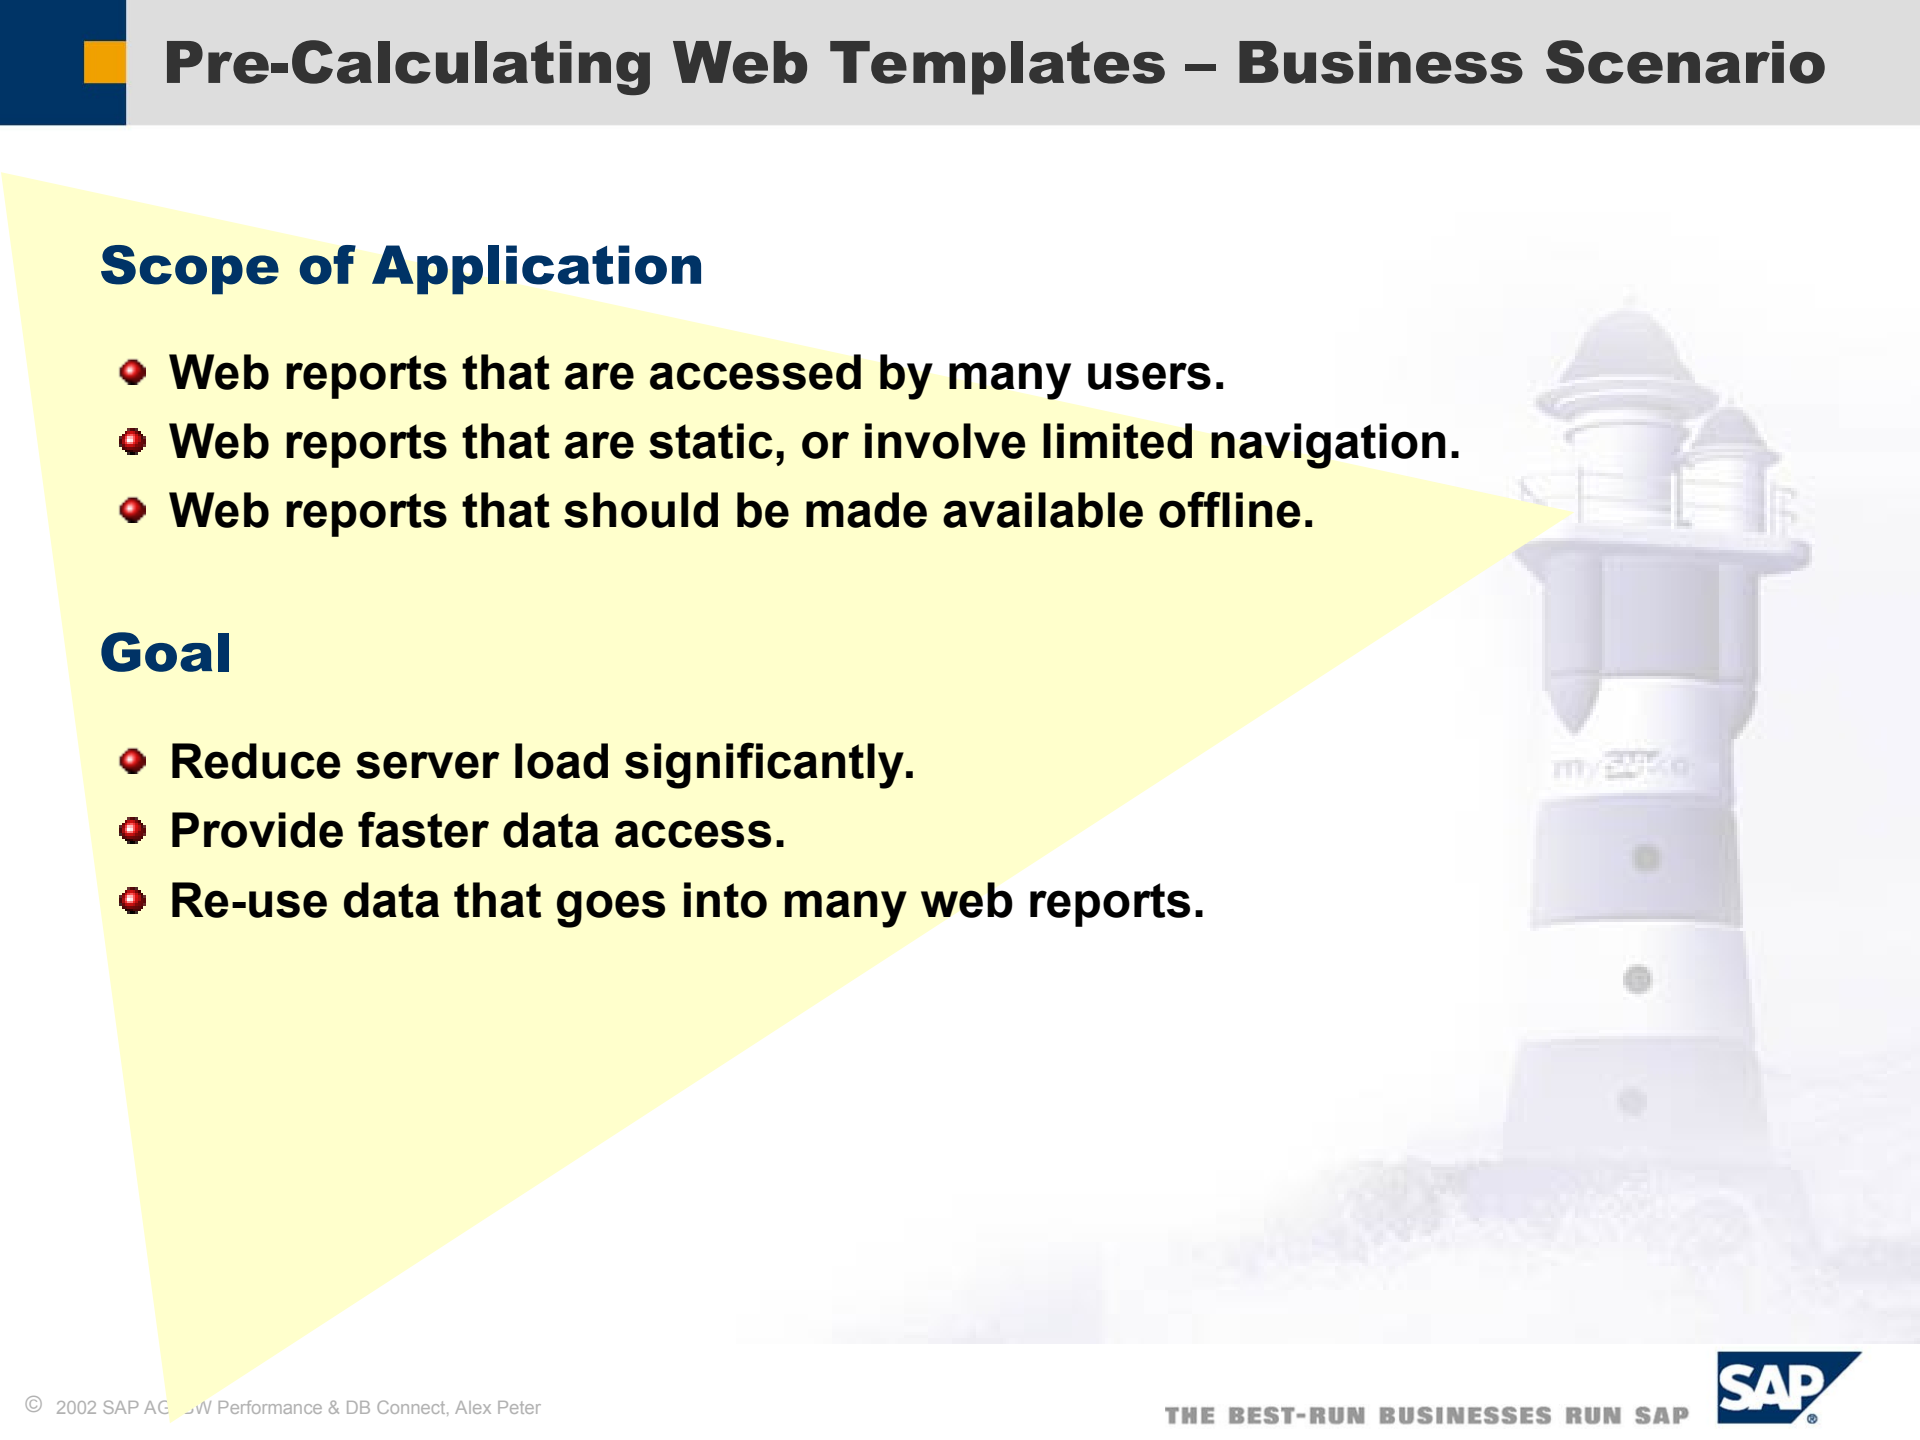

Pre-Calculating Web Templates – Setup

Realization

Settings

Output formats

- Data
- HTML for Web Browser
- HTML for Pocket IE
- Excel

Parameterization

- specify filter values
- pre-assign variables

Access

Access modes

- NEW = Default
- STORED
- HYBRID
- STATIC
- STATIC_HYBRID

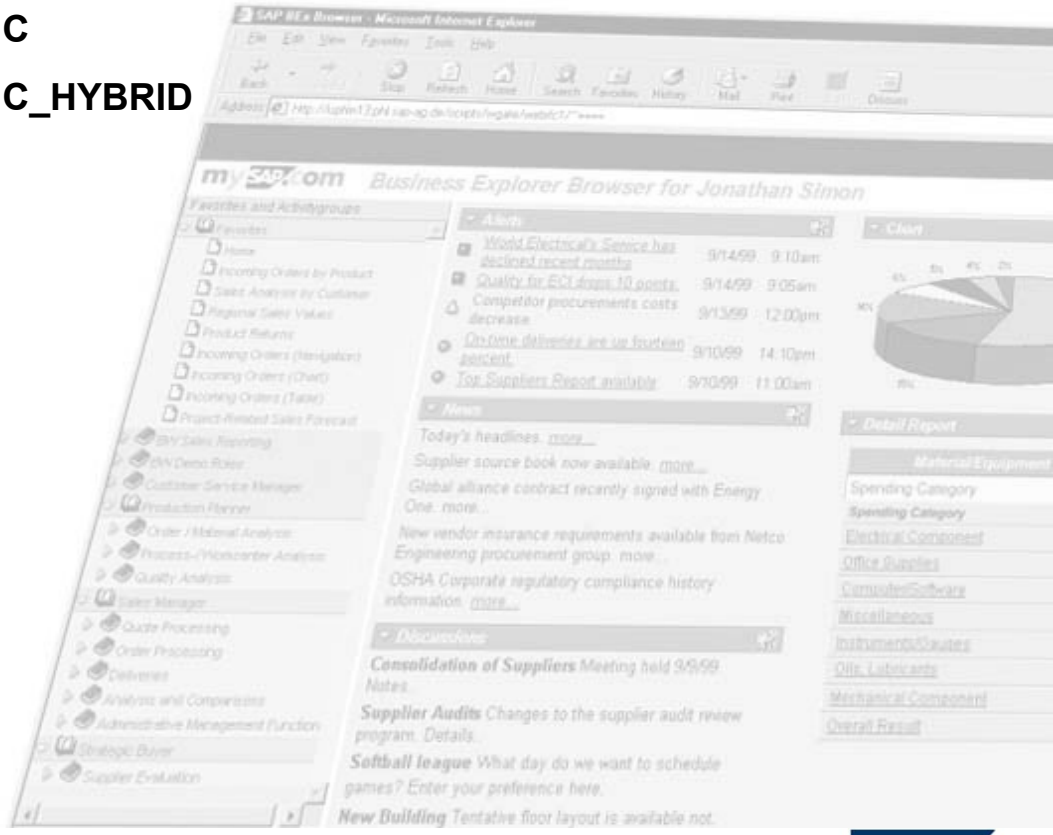

Access Modes for Web Templates

DATA MODE

NEW

STORED

HYBRID

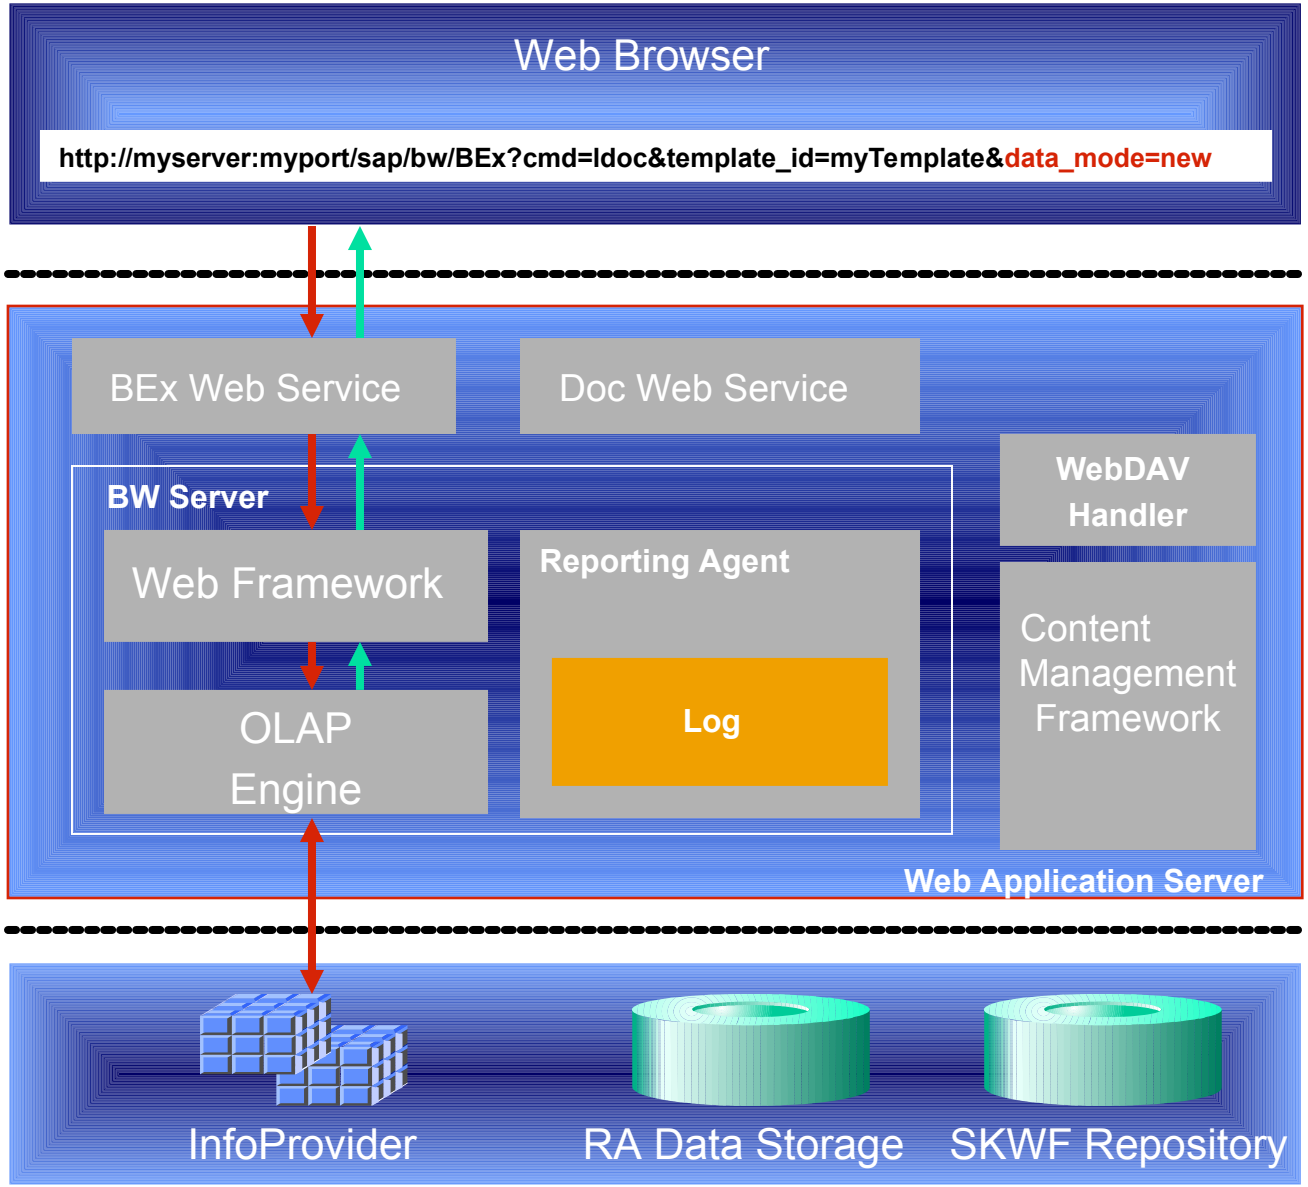

Access Modes for Web Templates

DATA MODE

NEW

STORED

HYBRID

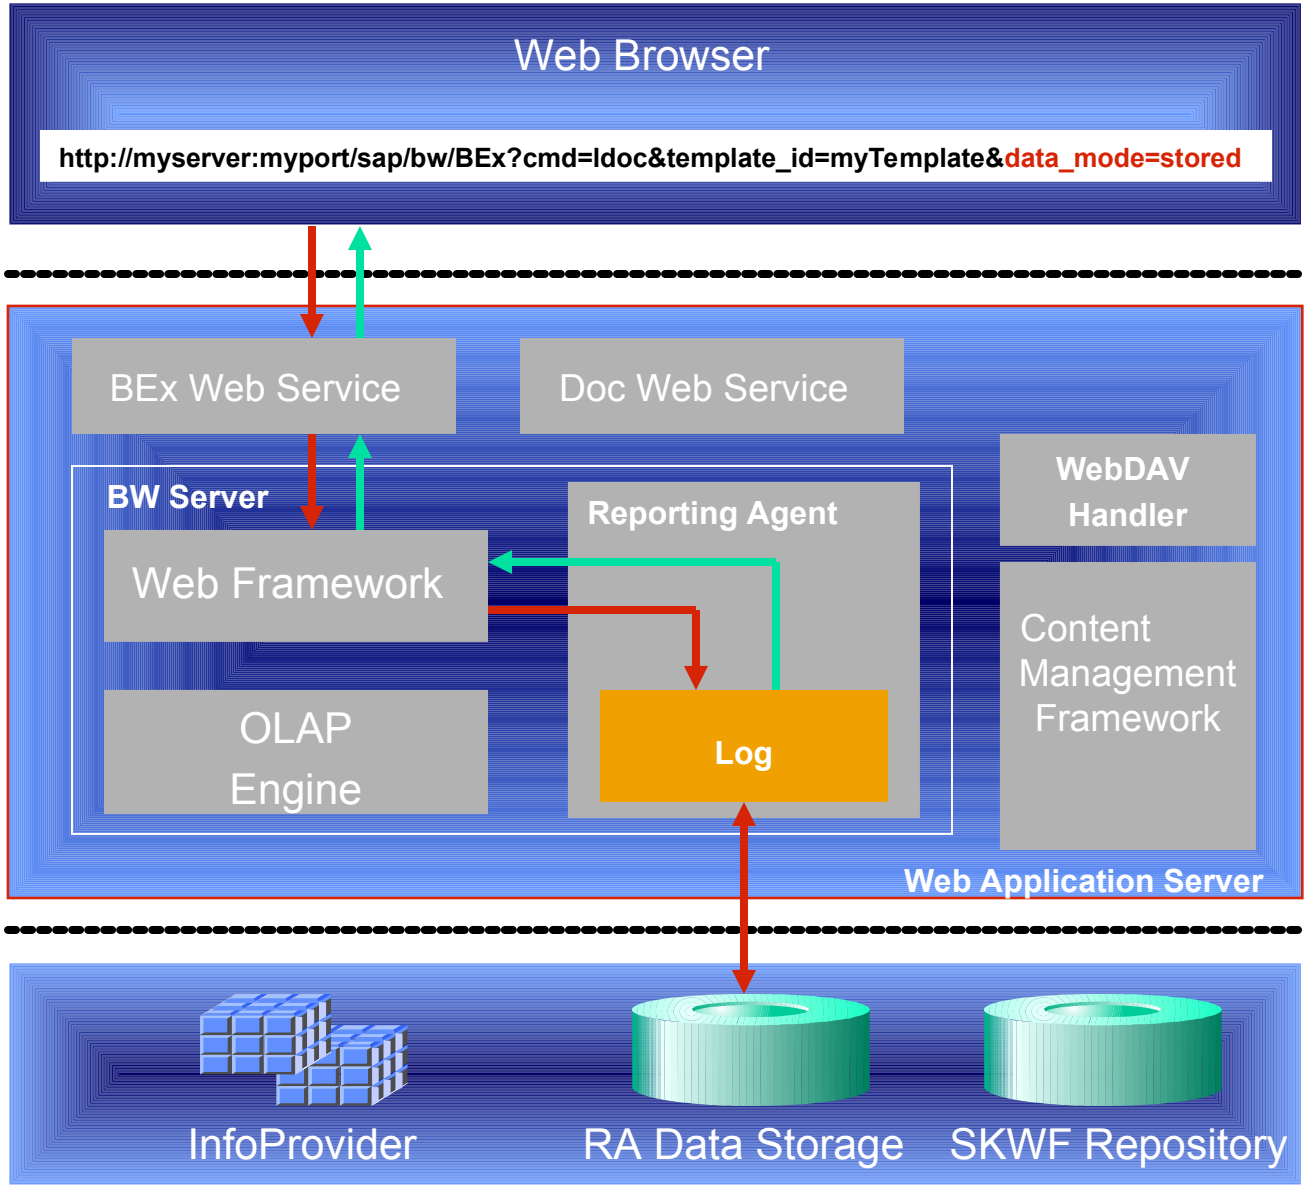

Performance Tuning Concepts

- MultiProvider
- Partitioning
- Data Modeling and Line Item Dimension
- Web Reporting
- OLAP Query Cache

OLAP Caching

Query results and navigation states of the OLAP processor are stored in a cache area of the application server as clustered data

Caching in 2.x and 3.0A:

- One local cache for each session
- No cache access across sessions
- No possibility to set size of cache

Caching in 3.0B:

- Additional global cache which is accessible from all sessions
- Global and local cache sizes can be adjusted
- Global cache can be switched off entirely

Query Execution: From Memory, Aggregate, or Cube

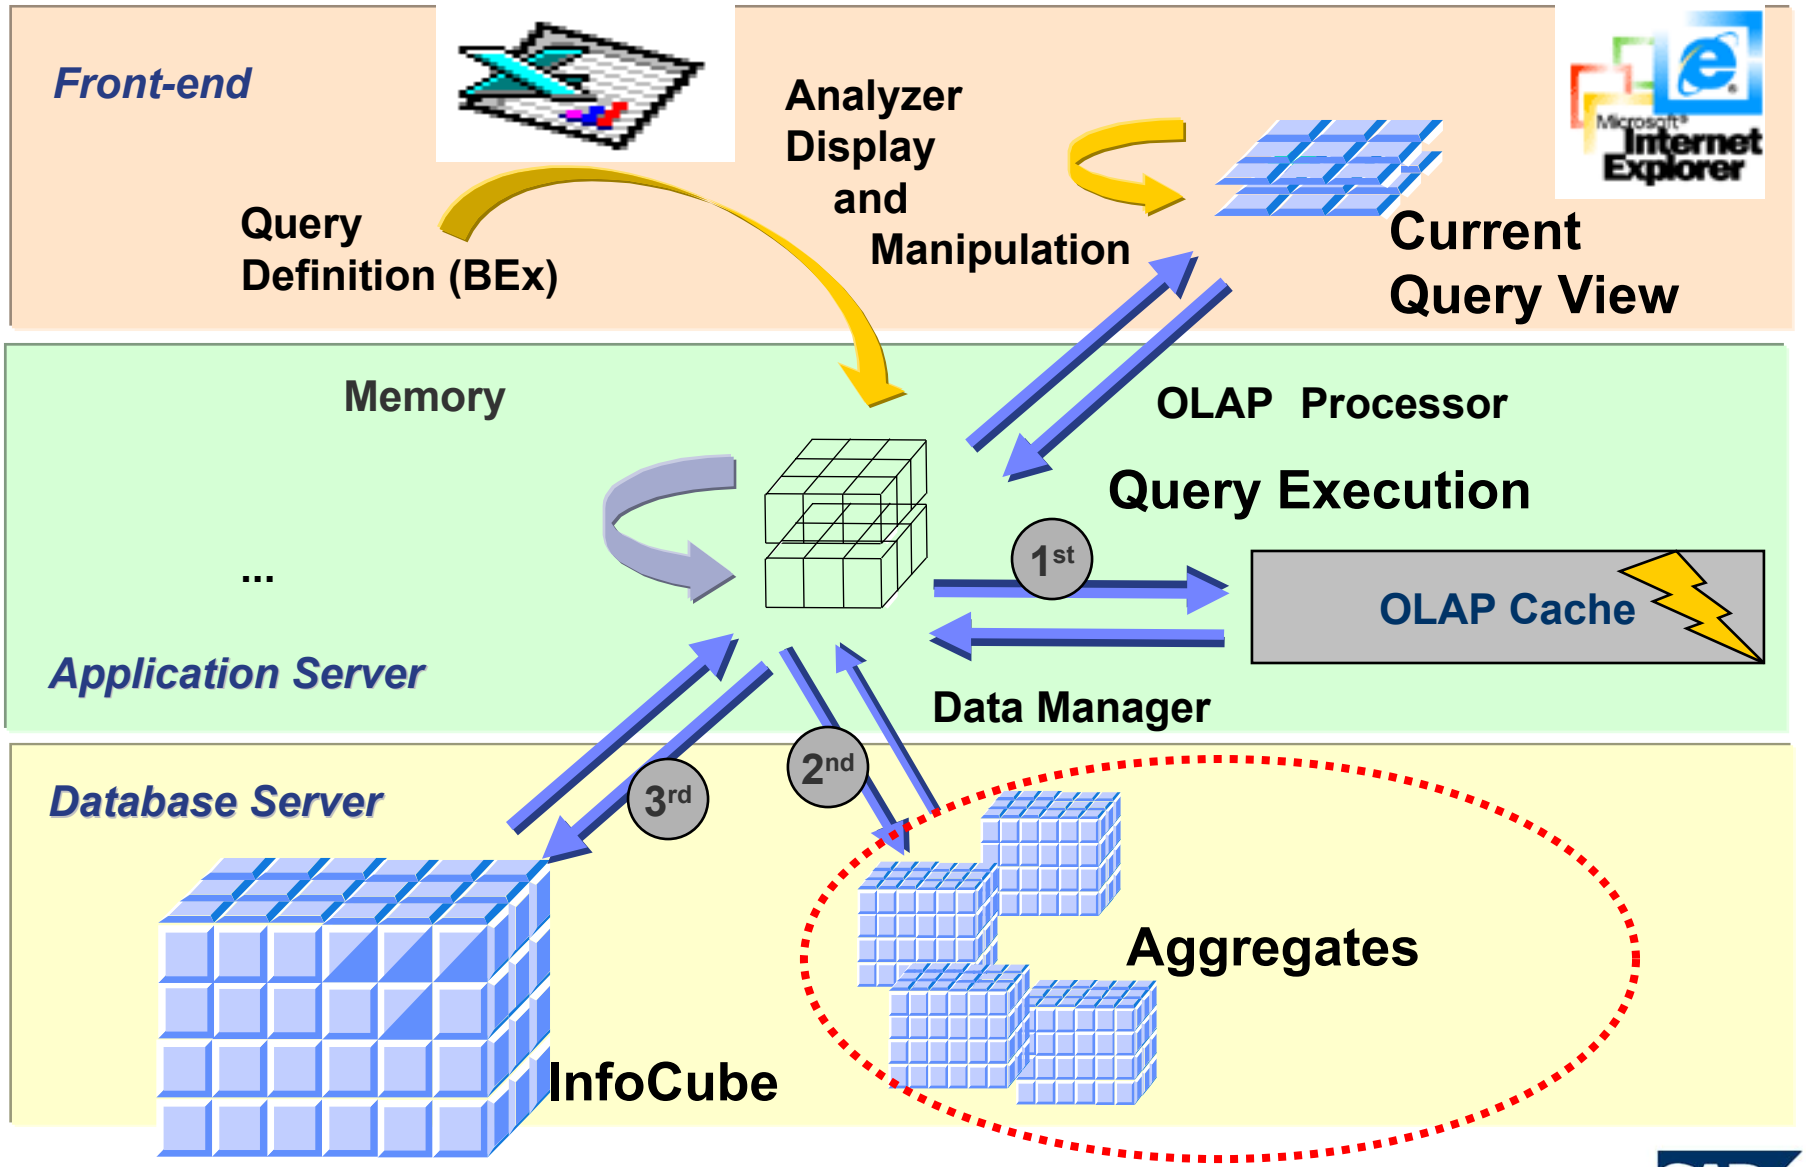

OLAP Caching

Cache is part of application buffer (Import / Export SHM)

Specific Instance Profile parameters:

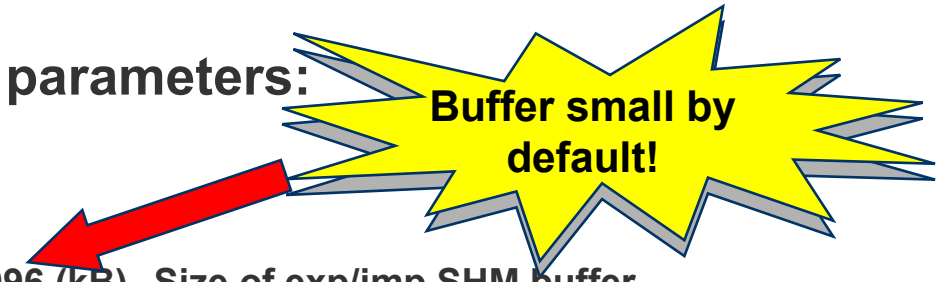

rsdb/esm/buffersize_kb	4096 (kB)	Size of exp/imp SHM buffer
rsdb/esm/max_objects	2000	Max. number of objects in the buffer
rsdb/esm/large_object_size	8192 (byte)	Estimation for the size of the largest object
rsdb/esm/mutex_n	0	Number of mutexes in Exp/Imp SHM buffer

Local cache will be used if data cannot be stored in global cache (e.g. if global cache has been switched off)

Benefits: reduced workload on database and application server

OLAP Caching Global Settings

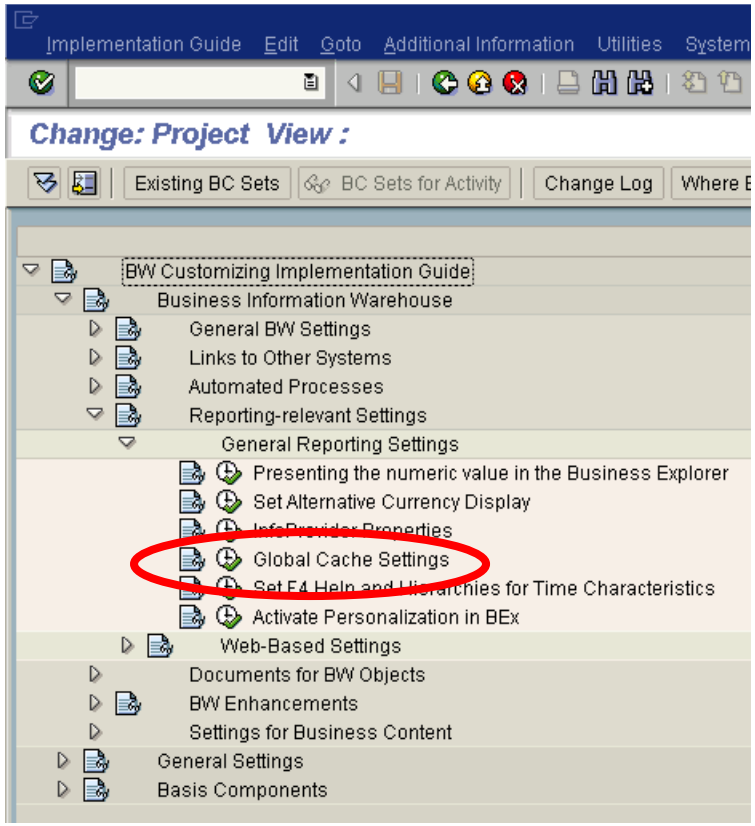

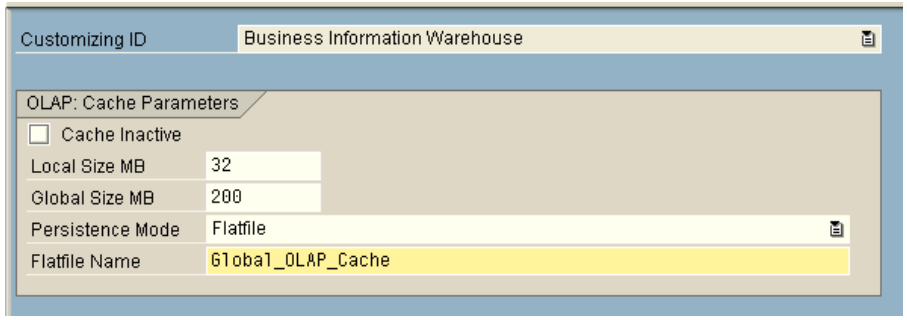

Customization menu (SPRO)

- Cache inactive:
  - Switch *global* caching off
- Local size:
  - Max. size of *local* cache (session results only)
- Global size:
  - Max. size of *global* cache (used for all sessions)
- Persistence Mode
  - Optional: Flat file or clustered table

Note: total size of (uncompressed) runtime objects, actual memory requirements are lower due to compression

OLAP Caching Settings for InfoProvider

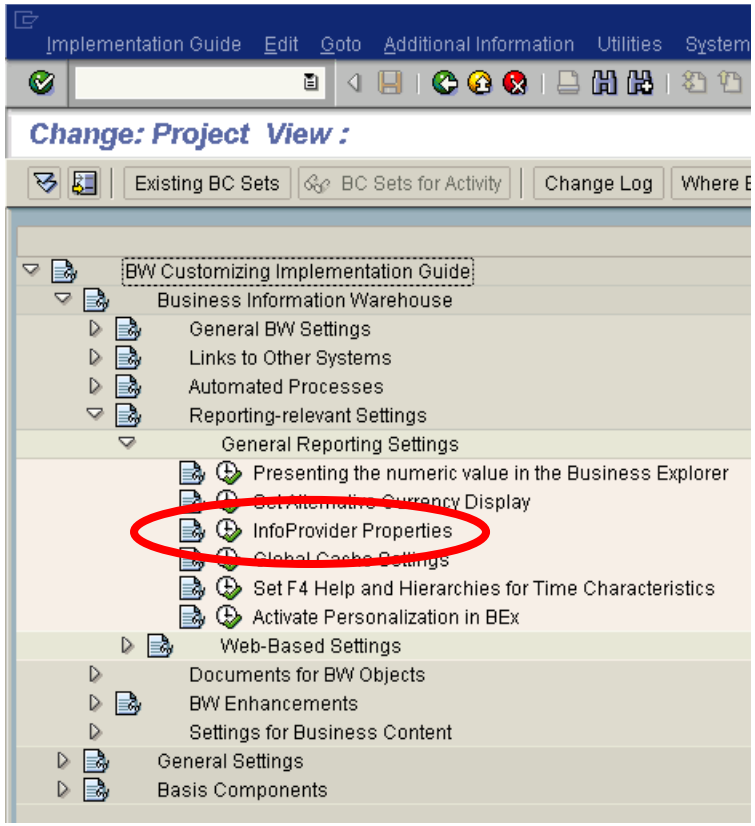

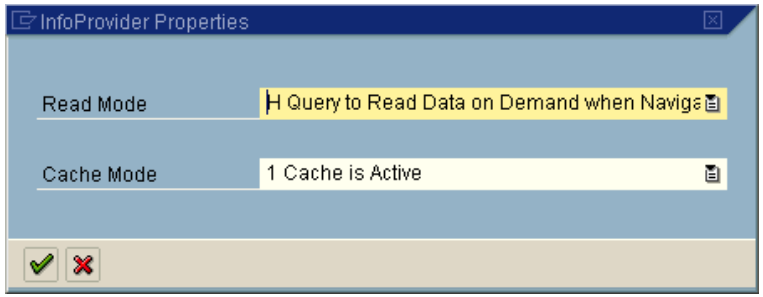

For each InfoProvider:

- Set cache mode BEFORE defining a query (inactive / active)
  Cache is active by default

Settings can also be changed for individual queries (RSRT):

run query in debug mode, select „Do not use cache“

OLAP Cache Monitor in Transaction RSRT

Query Cache

- Part of application buffer
- Query results and navigation status stored
- Similar queries (from any user) can use the cache

Query Cache Monitor