




ACTIVATE BUSINESS WITH THE POWER OF I.T.™



Reorganization Strategies in Depth

Peter Plevka

Consulting Brand Manager, BMC Software

Agenda



- › General Reorg Considerations
- › Reorg Availability Options
- › Reorg Automation and Avoidance
- › DB2 UDB for z/OS V8 and 9 and Reorg
- › Summary

Agenda



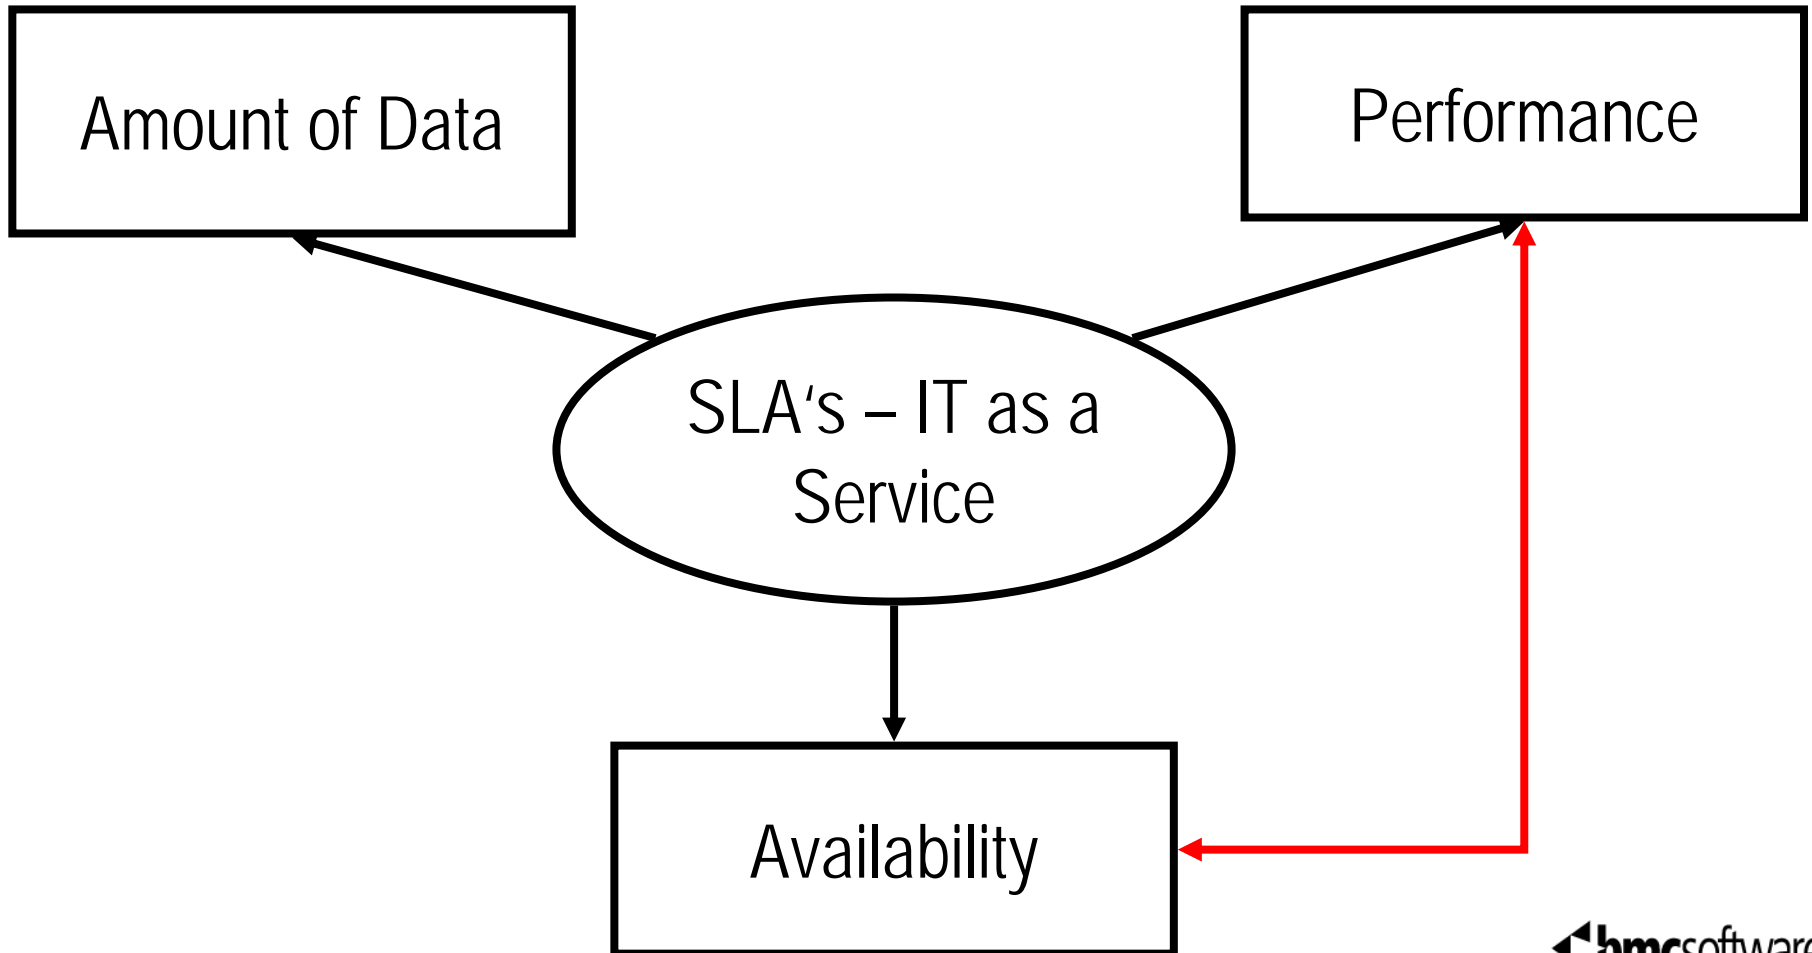
- › **General Reorg Considerations**
- › Reorg Availability Options
- › Reorg Automation and Avoidance
- › DB2 UDB for z/OS V8 and 9 and Reorg
- › Summary

What this session is all about




- › General concepts
- › Reorg automation
- › Avoidance of unnecessary reorgs
- › Vendor neutral

I want to provide you with all necessary input to decide on a particular Reorg Automation Strategy, that fits your needs





Overall Tuning Potential



The Main Message



- › Reorg is NOT an art
 - Automate once, and forget about it
- › Well reorganized objects =




Reorg Reasons - Indexspace




- › Leaf Distribution – LEAFDIST, LEAFNEAR, LEAFFAR from SYSIBM.SYSINDEXPART
- › Decrease Index Levels
- › Pseudo Deleted Rows
- › Freespace low (PCTFREE/FREEPAGE)
- › Secondary Extents

Availability




IO
CPU



Reorg Reasons – Tablespace



- › **Poor Clusterratio**
 - Check if clustered index is really used by SQL
- › **Freespace low**
- › **Row relocation**
 - FARINDREF/NEARINDREF from SYSIBM.SYSTABLEPART !
- › **Dead Space – Deleted records**

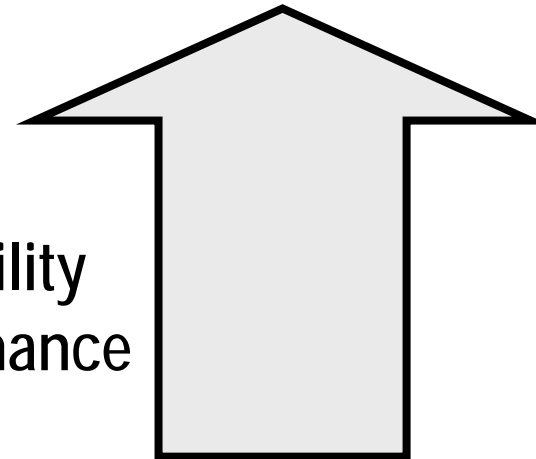


Reorg Reasons – Tablespace



- › Secondary Extents
 - ALTER PRIQTY/SEQTY before Reorg
 - Minor path length increase = CPU time increase
- › Resolve REORGP – after ALTER of part index key
- › Resolve AREO (Advisory) for optimal performance
 - After REORG REBALANCE
 - Add COL to table and index
- › Purge/Archive through REORG

Availability
Performance



Agenda



- › General Reorg Considerations
- › **Reorg Availability Options**
- › Reorg Automation and Avoidance
- › DB2 UDB for z/OS V8 and 9 and Reorg
- › Summary

SHRLEVEL Overview



› Indicates access allowed during RELOAD phase

› NONE

– UNLOAD

– READ

– RELOAD, SORT, BUILD

– **NO ACCESS**

› REFERENCE – Online Read !

– UNLOAD

– READ

– RELOAD, SORT, BUILD

– READ to original copy of TS/IX

– FASTSWITCH

– NO ACCESS

› CHANGE – Online Read/Write !

– UNLOAD

– READ/WRITE

– RELOAD, SORT, BUILD

– READ/WRITE to original copy of TS/IX

– LOGFINAL, FASTSWITCH

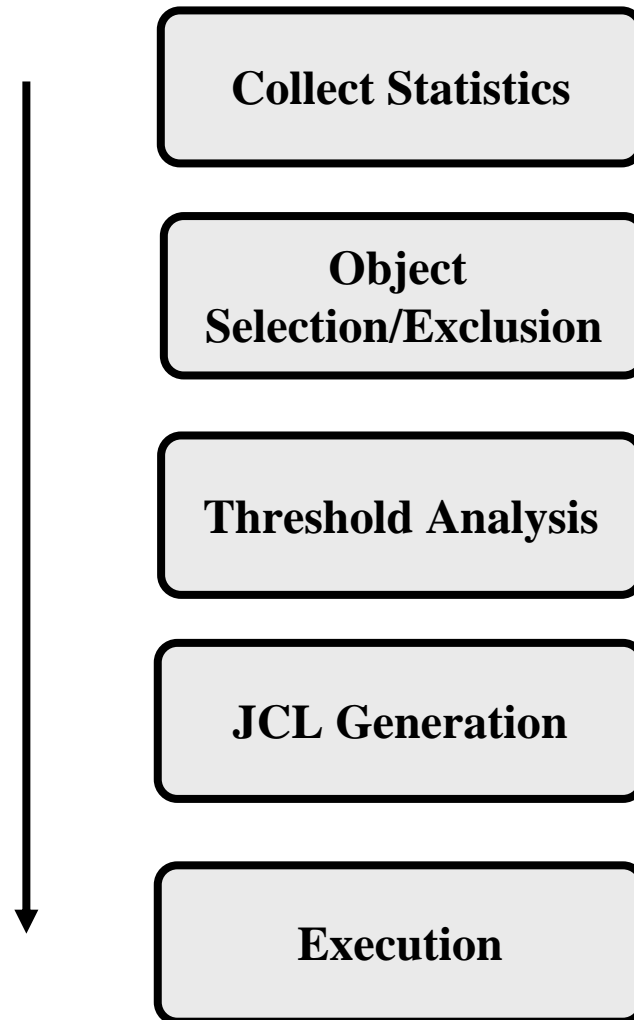
– NO ACCESS

Agenda

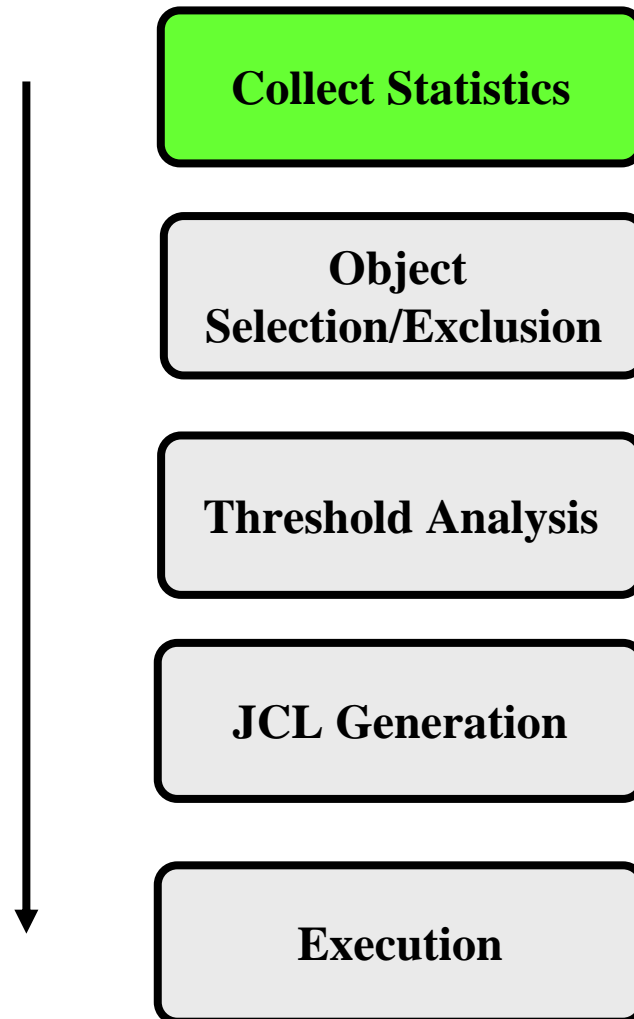


- › General Reorg Considerations
- › Reorg Availability Options
- › **Reorg Automation and Avoidance**
- › DB2 UDB for z/OS V8 and 9 and Reorg
- › Summary

To Reorg, or not to Reorg, that's the Question



To Reorg, or not to Reorg, that's the Question





› IBM RUNSTATS

- HIST Catalog Tables provide data over time
- PLEASE ! Schedule RUNSTATS for Optimizer needs
- Includes IBM STOSPACE data - Space in KB
 - RUNSTATS TABLESPACE UPDATE(ALL/SPACE)

› Real-Time-Stats tables (since V7)

- Very useful
- ‘Since’ Utility values – no trending
- Relatively complex to implement via SP



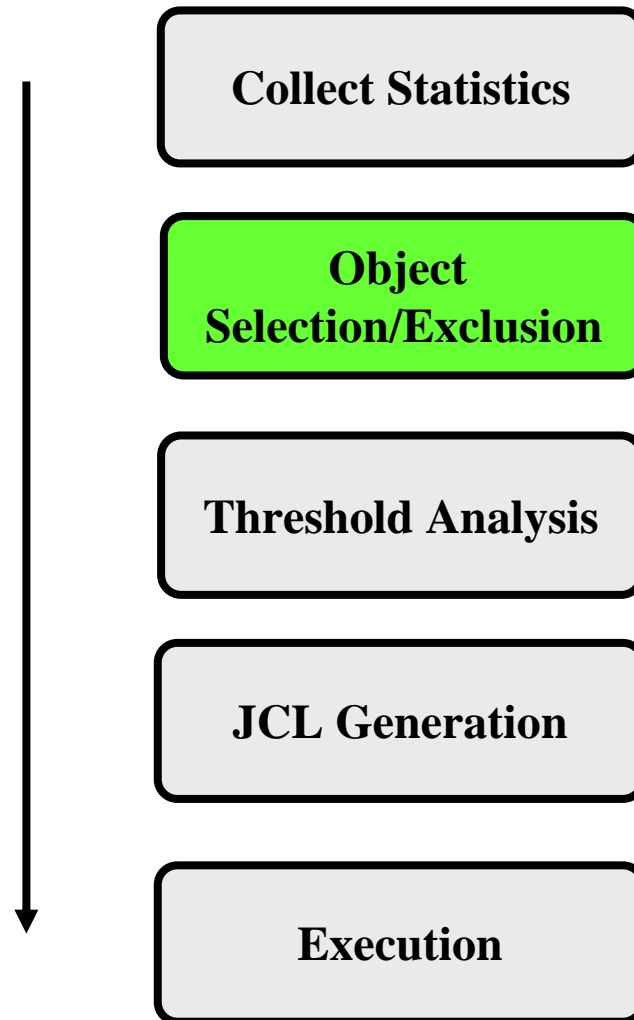
› VSAM LISTCAT

- Its the truth ! Use it for space related reorgs

› ISV Tools

- Do not rely on RUNSTATS/STOSPACE
- Optionally utilize RTS, LISTCAT and VTOC data
- Enhanced stats database and reports (Trending, Forecasting)

To Reorg, or not to Reorg, that's the Question





- › Use IBM LISTDEF or query DB2 Catalog individually
 - Can also be used for stats collection with RUNSTATS
- › Exclude particular objects if needed
 - Very large/small objects
 - Hot tables – when is a table hot ?
 - LOB Tablespaces



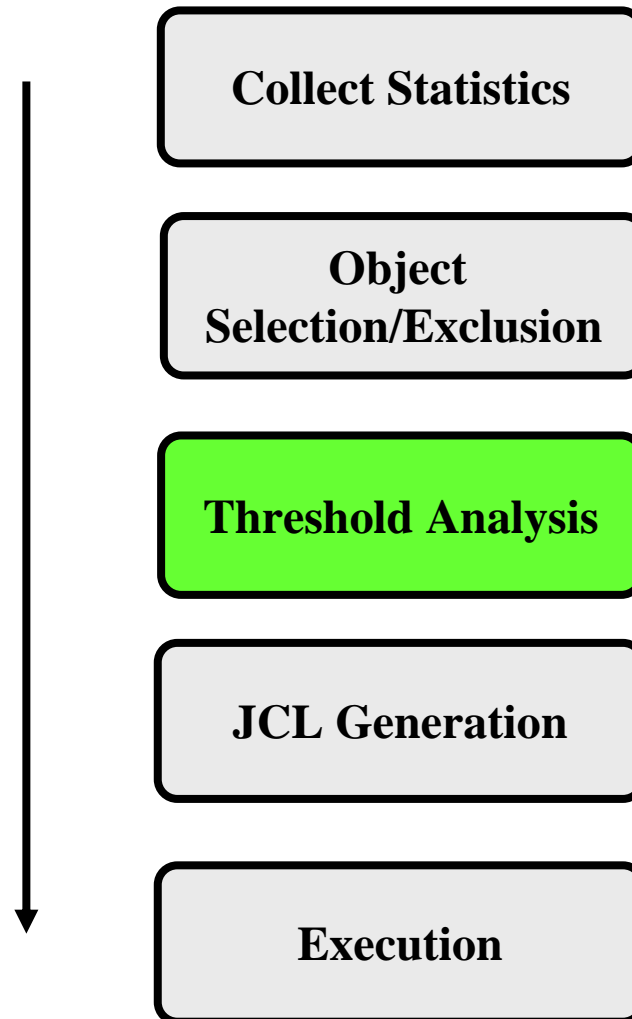
› Consider

- DEFINE NO objects
- Indexes
- RI relationships

› Grouping of objects by Application/Business Service

- Collection ID
- Package/Plan Name List
- Use SYSPACKDEP or SYSPLANDEP to find objects

To Reorg, or not to Reorg, that's the Question



Threshold Analysis



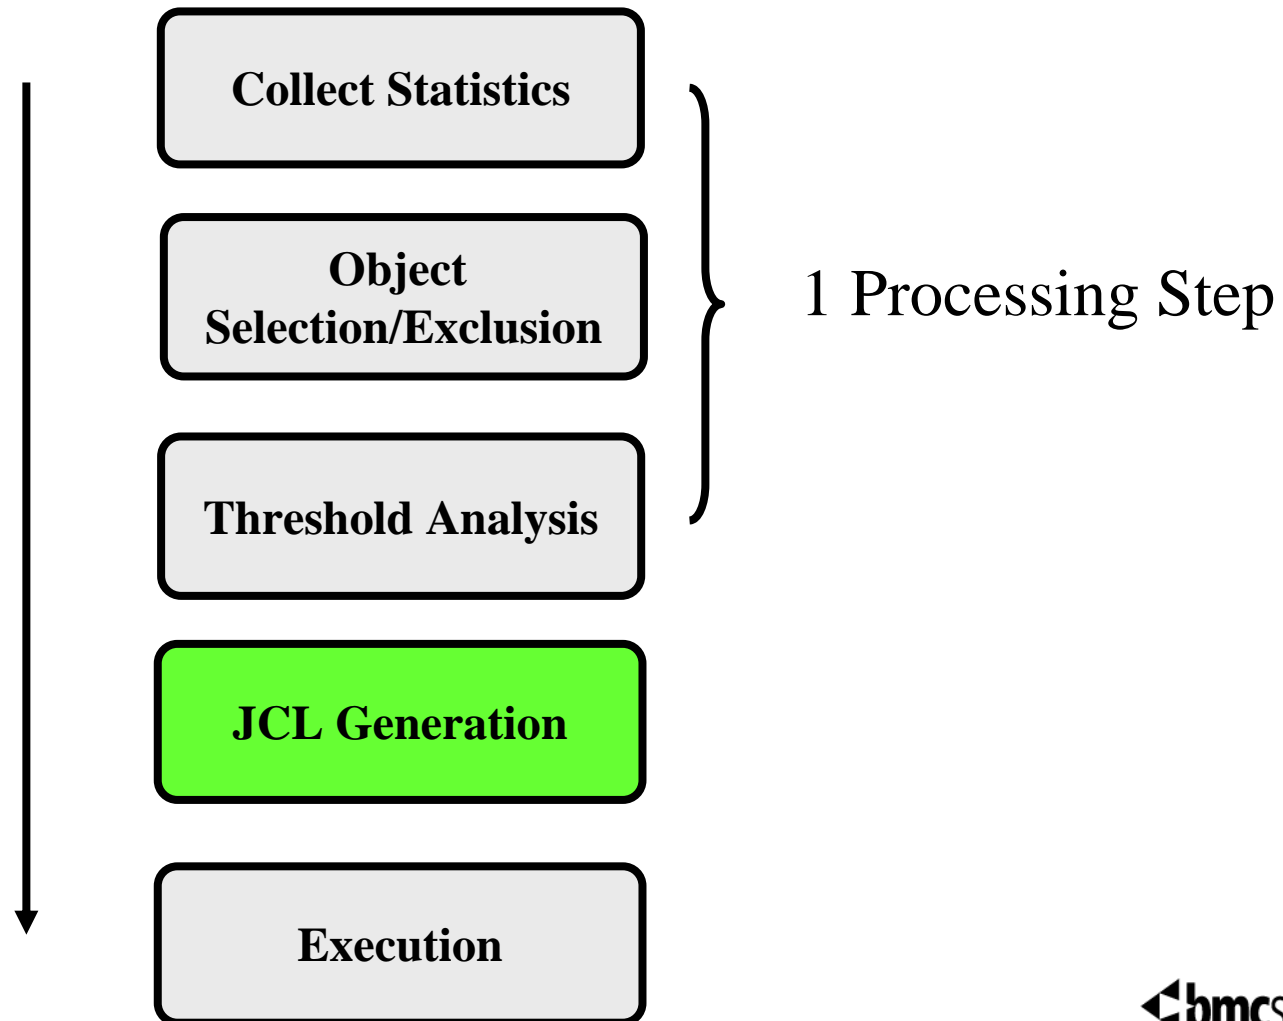
- › Find out which objects qualify for reorg
 - Via SQL: SELECT COUNT(*) WHERE EXISTS...
 - Returns 0 for FALSE and 1 for TRUE
 -
- › Consider AND'ing and/or OR'ing conditions
 - Physical Extents **and** growth rate over time
- › Trigger reorgs based on historical trends
 - I.E. more than 25% increase of rows since last stats
 - Reorg to avoid exceptions in the future
 - Stats data helps with capacity planning too

Threshold Analysis



- › Trigger different actions by condition type
 - SHRLEVEL REFERENCE/NONE REORG
- › Test it – find threshold values that fit your needs
 - 80% or 90% Clusterratio ?
- › Think about triggering Reorgs when performance declines

To Reorg, or not to Reorg, that's the Question





- › Avoid index reorg if tablespace needs reorg too

- › Automatically resize (ALTER) objects before reorg
 - Calculate PRIQTY (SEQTY ?)
 - Consider sliding scale method to always reach max dataset size (DB2 V8)

- › Calculate Workdataset size
 - SORTWK, SYSUT1, SYSCOPY DDs etc...



- › **IBM Reorg uses Mappingtable and Index**
 - Create TS/TB/IX in STEP before Reorg
 - Drop TS after Reorg
 - Sizing of Mappingtable and Index depends on table size (110%)
 - Run MODIFY on Mappingtable TS to reduce DBD size

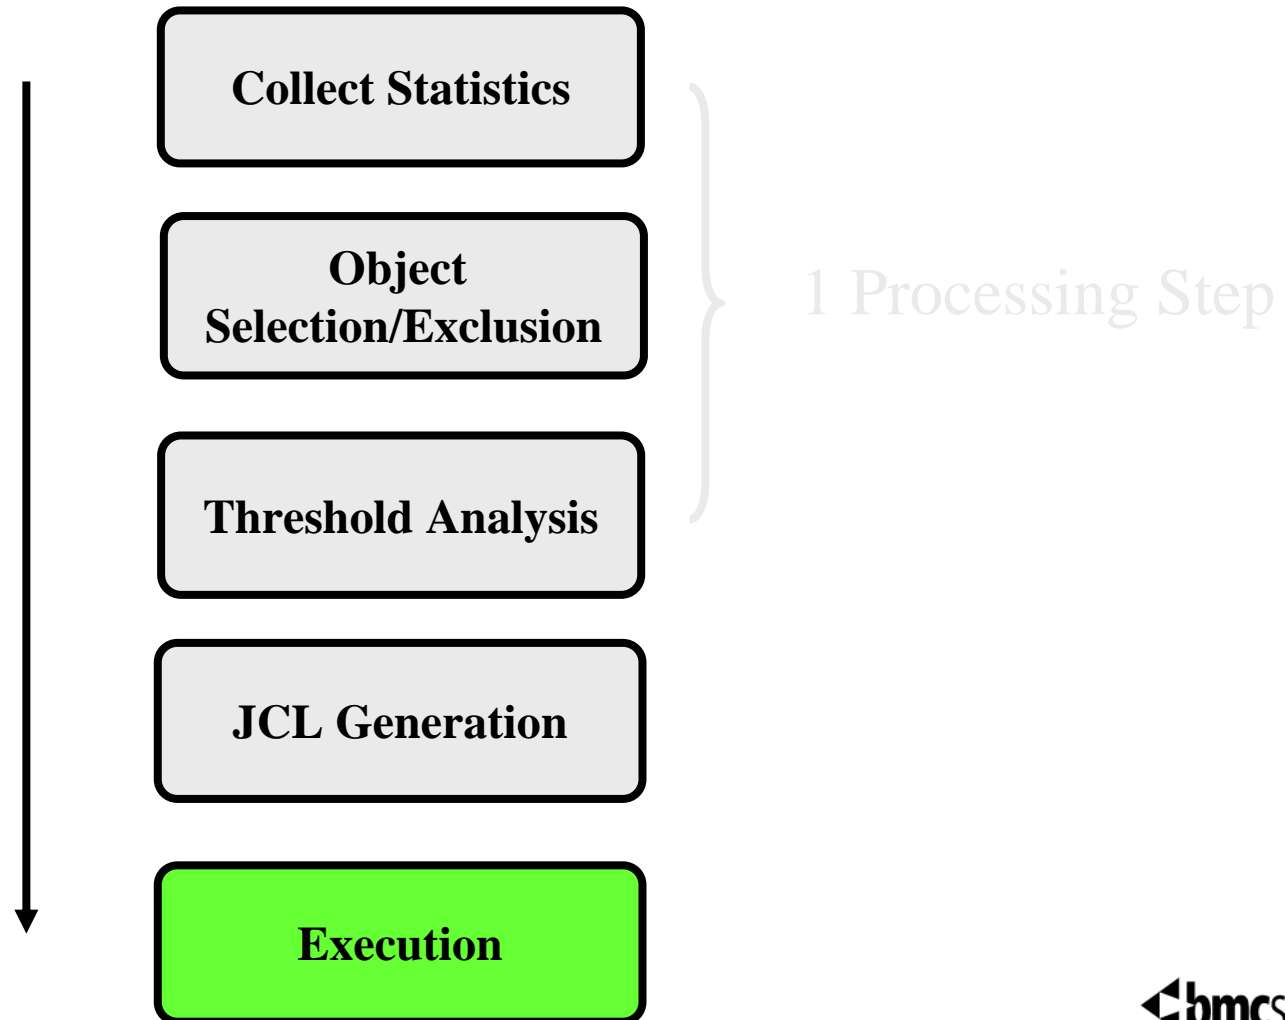
- › **Assign priorities by object and exception**
 - Do worst organized and largest object first

JCL Generation



- › Partition level reorgs
- › ISFP Skeleton technology useful

To Reorg, or not to Reorg, that's the Question





- › Talk to your job scheduling person
- › Know your application data usage or use a tool
 - Read vs. Update
 - Frequent Delete Jobs
 - LOAD RESUME – Clusterratio !
 - Day vs. Night vs. 24/7



- › Do REORGs during off-peak time
 - Even SHRLEVEL CHANGE
 - Use regular maint. window for hot tables

- › Consider Workload balancing
 - Submit largest and worst reorganized object first
 - Use stats about past executions to schedule reorg better

- › What if reorg fails ?
 - Online Reorg relieves



- › Synchronization with transactions is the issue !

- › 2 Joices for LOGFINAL phase:
 - Transactions must complete, **OR** REORG must complete
 - Acceptance that there may be some application failures
 - Use `DRAIN_WAIT int` keyword (for IBM Reorg)
 - tells REORG how long to wait for not committing transactions
 - `RETRY` keyword multiplies `DRAIN_WAIT`

Agenda



- › General Reorg Considerations
- › Reorg Availability Options
- › Reorg Automation and Avoidance
- › **DB2 UDB for z/OS V8 and 9 and Reorg**
- › Summary



- › Type of enhancement depends on vendor
- › IBM related
 - V8
 - DPSI – no BUILD2 Phase for NPIs for IBM Online Reorg
 - SCOPE PENDING keyword for REORP/AREO status
 - ALTER ROTATE PART – Reorg does mass delete
 - SHRLEVEL CHANGE for DISCARD Reorg
 - 9
 - No BUILD2 Phase anymore
 - CPU reduction for reorgs of varying length data
 - LOB Reorg
 - Original LOB TS is drained from writers
 - Extract from original LOB and insert into shadow dataset
 - Drain readers
 - FASTSWITCH

Agenda



- › General Reorg Considerations
- › Reorg Availability Options
- › Reorg Automation and Avoidance
- › DB2 UDB for z/OS V8 and 9 and Reorg
- › **Summary**

Summary



- › Reorg is not an art
- › Use Online Reorgs
- › Automate as much as possible
- › Get it off your desk



Peter Plevka

BMC Software

Peter_plevka@bmc.com