



IBM z Systems



# TPFUG – Various Enhancements

Michael Shershin, TPF Development Lab

Date of Presentation





# PJ42717 – GP Utilization

# PJ42717 – GP Utilization

- GP utilization is the utilization for TPF work load that is charged by workload license charging (WLC).
- Total utilization is comprised of
  - GP utilization
  - TPF work load that is not charged by WLC
- Transformation Engines (TE) execute TPF work load that is not charged by WLC.
- Key to TE engines is the amount of TPF work that is charged by WLC or GP utilization
  - Financial penalties may be incurred if the TPF work that is charged by WLC is greater than the number of GP engines on the machine.

# PJ42717 – GP Utilization

- GP utilization reported by
  - ZSTAT GPU command
  - Data collection / reduction
  - Continuous data collection

# PJ42717 – GP Utilization

- GP cap
  - Guideline for the maximum GP utilization that you want to run on a specific LPAR
    - GP utilization consistently above the GP cap may result in financial penalties
  - Not enforced
  - Used by LODIC UTILIZATION support
- GP cap value obtained from utilization configuration file
  - On TPF utilization configuration file is at /etc/ibm\_utl\_cfg.csv
  - Comma separated variable (csv) file
- Default GP cap is 100%
- To set or change the GP cap
  - Create or edit file off of TPF
  - FTP utilization configuration file to TPF
  - Validate the file with command: ZSYSL VALIDATE
  - Set or change current GP cap with command: ZSYSL REFRESH UTILCONFIG

# PJ42717 – GP Utilization

- GP cap
- Intent is to have a GP cap that is unique for each machine
  - May have different number of GP and TE engines depending on the machine
- Utilization configuration file contains:
  - Entry defines GP cap for an LPAR on a machine
    - Machine serial number
    - LPAR name
  - Wildcard \* can be used for an LPAR name or machine serial number
- Utilization configuration file example:

```
UTILIZATION_CONFIGURATION_VERSION=1  
114401,TPFP1,75  
4401,*,70  
*,*,90
```

# PJ42717 – GP Utilization

- LODIC UTILIZATION support enhanced
  - Provides ability to control processes based on GP utilization
  - No change in the LODIC API
    - LODIC CHECK,UTILCLASS=
    - LODIC SUSPEND,UTILCLASS=
  - UTILCLASS definitions set by command ZSYSL
    - ZSYSL CREATE UTILCLASS
    - ZSYSL ALTER UTILCLASS
  - UTILCLASS can check:
    - Total utilization with the TARGET option
    - GP utilization with the GPTARGET option
      - GPTARGET is a percentage of GP cap
  - ZSYSL example:

```
ZSYSL CREATE UTILCLASS-USERHI TARGET-80 GPTARGET-80
CSMP0097I 22.11.10 CPU-B SS-BSS SSU-HPN IS-01
SYSL0010I 22.11.10 UTILIZATION PRIORITY CLASS USERHI ADDED+
```

# PJ42717 – GP Utilization

- ZSTAT GPU is similar to ZSTAT U with the following additions:
  - LPAR name or VM user ID
  - Actual utilization of entire LPAR (does not remove defer / delay activity)
  - GP utilization or TPF work load that is charged by WLC (removes defer / delay activity)
  - GP cap is a guideline that represents the maximum GP utilization that should be used
    - Obtained from the utilization configuration file
  - GP engines represent the number of engines used to process TPF work load that is charged by WLC

```

CSMP0099I 15.26.32 010012-B ZSTAT GPU+
CSMP0097I 15.26.32 CPU-B SS-BSS SSU-HPN IS-01
STAT0019I 15.26.32 SYSTEM UTILIZATION DISPLAY
STATIC POWER SAVE MODE - SYSTEM AT 100 PERCENT

```

NAME	ACTUAL UTIL	GP UTIL	GP CAP	GP ENGINES
TPFP1	49.9	43.1	75.0	2

```

NUM  ADR  UTIL/ ADJ  CROSS  READY  INPUT  VCT  SUSPD  DEFER  ACT-ECB
IS-  1  00  36.2/ 36.0    0    1    0    0    1    0    40
IS-  2  01  61.0/ 60.8    1    2    1    0    0    0    37
IS-  3  02  60.8/ 60.6    0    0    0    0    0    0    33
IS-  4  03  42.0/ 42.0    0    0    0    0    0    0    36
END OF DISPLAY+

```





# PJ41930 – Enhanced dump controls

# PJ41930 – Enhanced dump controls

- Improved ability to limit the number of dumps that are written to tape
- PJ41930 addresses situation when large number of unique dumps are happening.
- Before PJ41390
  - ZASER NODUPL - Duplicate dumps are NODUMPed.
    - OPR dumps
  - ZASER MAXCTL - maximum number of CTL dumps allowed within one minute.
    - If the number of CTL dumps exceeds MAXCTL, additional CTL dumps will be NODUMPed
- PJ41930 adds
  - ZASER MAXOPR – maximum number of OPR dumps allowed within one minute.
    - If the number of OPR dumps exceeds MAXOPR, additional OPR dumps will be NODUMPed
  - ZASER DBACTL – maximum percent of the DBA that can be in use when a CTL dump is taken
    - If the percentage of in use area in the DBA is greater than DBACTL when a CTL dump happens, the CTL dump will be NODUMPed.
  - ZASER DBAOPR – maximum percent of the DBA that can be in use when an OPR dump is taken
    - If the percentage of in use area in the DBA is greater than DBACTL when an OPR dump happens, the OPR dump will be NODUMPed.

# PJ41930 – Enhanced dump controls

- Additional change
  - Messages CPSE0156I and COSK0399I are sent when a dump is written to tape
  - Now includes dump type (CTL vs OPR) and dump number


COSK0399I 13.56.48 \*CP\* BSS DUMP SEQUENCE NUMBER 3561 STARTED ON VSN T1G047 [CTL-I0000E1](#)




PJ32404 – Enable ZOLDR load of CFVS

# PJ32404 – Enable ZOLDR load of CFVS

- CFVS is the file system support
- Previously any change to file system logic required an outage to change.
  - TLDR or ALDR load was required.
- PJ32404 allows changes to file system logic to be loaded without an outage.
  - Most changes can be loaded using OLDR.
  - If a future change to CFVS requires an IPL, the APEDIT will indicate this requirement.



# PJ41914 – Capture / Restore enhancement



# PJ41914 – Capture / Restore enhancement

- Before PJ41914
  - Capture and restore will not complete if one of the prime modules of a fully duplicated environment is offline.
- With PJ41914
  - Capture and restore allow processing of duplicate modules in place of prime modules that are offline in a fully duplicated environment.
  - To enable this support SIP macro DDCCAP CAPDUP=YES must be coded.
    - Default is DDCCAP CAPDUP=NO

# Disclaimer

- Any reference to future plans are for planning purposes only. IBM reserves the right to change those plans at its discretion. Any reliance on such a disclosure is solely at your own risk. IBM makes no commitment to provide additional information in the future.





Proposed: New data reduction tool

# Proposed: New data reduction tool

- New tool
  - Will run on linux
  - Will not require PL/I
  - Will contain new reports
  - Gradually old reports will be migrated to new tool
- Current tool
  - Will continue to run on z/OS
  - Will continue to require PL/I
  - Will not contain new reports
  - Only updated to correct bugs
- Output for new tool would be in two formats
  - Traditional reports
  - Data in a .csv file
- No change to TPF data collection



# Proposed: ECB resource monitor enhancements

# Original ECB resource monitor

- Each ECB counts resource usage in an ECB unique resource count table (RCT).
- System default resource limit table (RLT).
  - First level action
  - Second level action
- Each ECB has its own RLT.
  - When ECB is created, ECB RLT is copied from system default RLT.
- ECBs can change resource limits using the ECBMC macro.
- System default RLT can be changed using ZECBM command
- Actions:
  - When resource usage reaches the first level limits, the action taken is the same for all resources.
  - When resource usage reaches the second level limits, the action taken is the same for all resources.

# PJ40979 adds named limit set support

- Each ECB counts resource usage in an ECB unique resource count table (RCT).
- System default resource limit table (RLT).
  - First level action
  - Second level action
- → Named limit set is an alternate resource limit table
- Each ECB has its own RLT.
  - When ECB is created, ECB RLT is copied from system default RLT.
- → ECBs can use the alternate resource limits by using a named limit set
  - Use ECBMC or `tpf_ecbmc_set_lsetname()` to use a named limit set
  - Updates ECB RLT with values from named limit set.
- ECBs can change resource limits using the ECBMC macro.
- System default RLT can be changed using ZECBM command
- → Named limit set RLT can be changed using ZECBM command

# PJ40979 adds named limit set support

- Actions:
  - → First level actions to be taken and second level action to be taken can be set by resource.
  - System default actions to be taken if the first level and second levels limits are reached continue to exist.
  - When resource usage reaches the first level limits, the action taken is the action set for the specific resource or if no action is set, the system default action for all resources.
  - When resource usage reaches the second level limits, the action taken is the action set for the specific resource or if no action is set, the system default action for all resources.

# Proposed ECB resource monitor enhancements

- Problem: How to determine reasonable limit values for each resource?
- Provide ability to collect resource usage information
  - Similar to collection of resource usage by owner names (ZMOWN)
  - Counts include:
    - Counts of ECBs
    - Counts of all resources that are monitored by the ECB resource monitor
    - Include high water marks
  - Counters categorized by named limit set that is being used at ECB EXITC time.
  - Collect usage over a specified time period
  - Write output to a file in the file system.

# Proposed ECB resource monitor enhancements

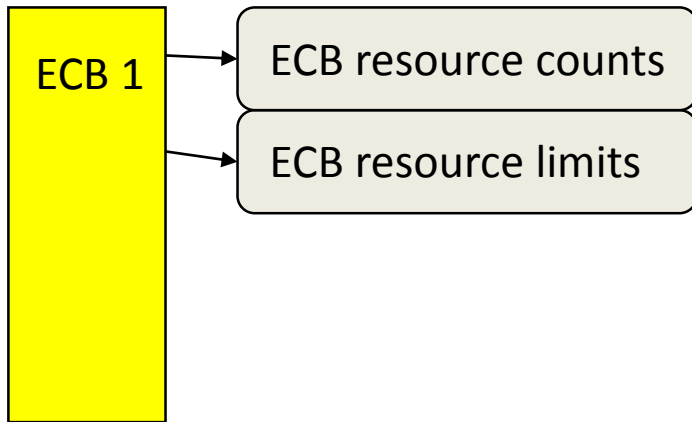
- Problem: How to monitor a utility that consists of a group of ECBs?
  - Single ECBs may behave well.
  - A group of ECBs in a utility combined may abuse resources.
  - For example:
    - Pool usage for any one ECB may be acceptable.
    - Pool usage for a group of 100 ECBs in a utility may use be too high.
    - Need to report or stop group of ECBs from using too many pools.
- Provide monitoring for a group of ECBs
- Group monitoring would happen when:
  - ECB has used a named limit set
  - Group monitoring is enabled for the named limit set
  - ECB creates a child ECB
- Group consists of:
  - The original ECB
  - Any child ECB that is created by the original ECB
  - Any child / grandchild / great-grandchild ECBs created by child ECBs.



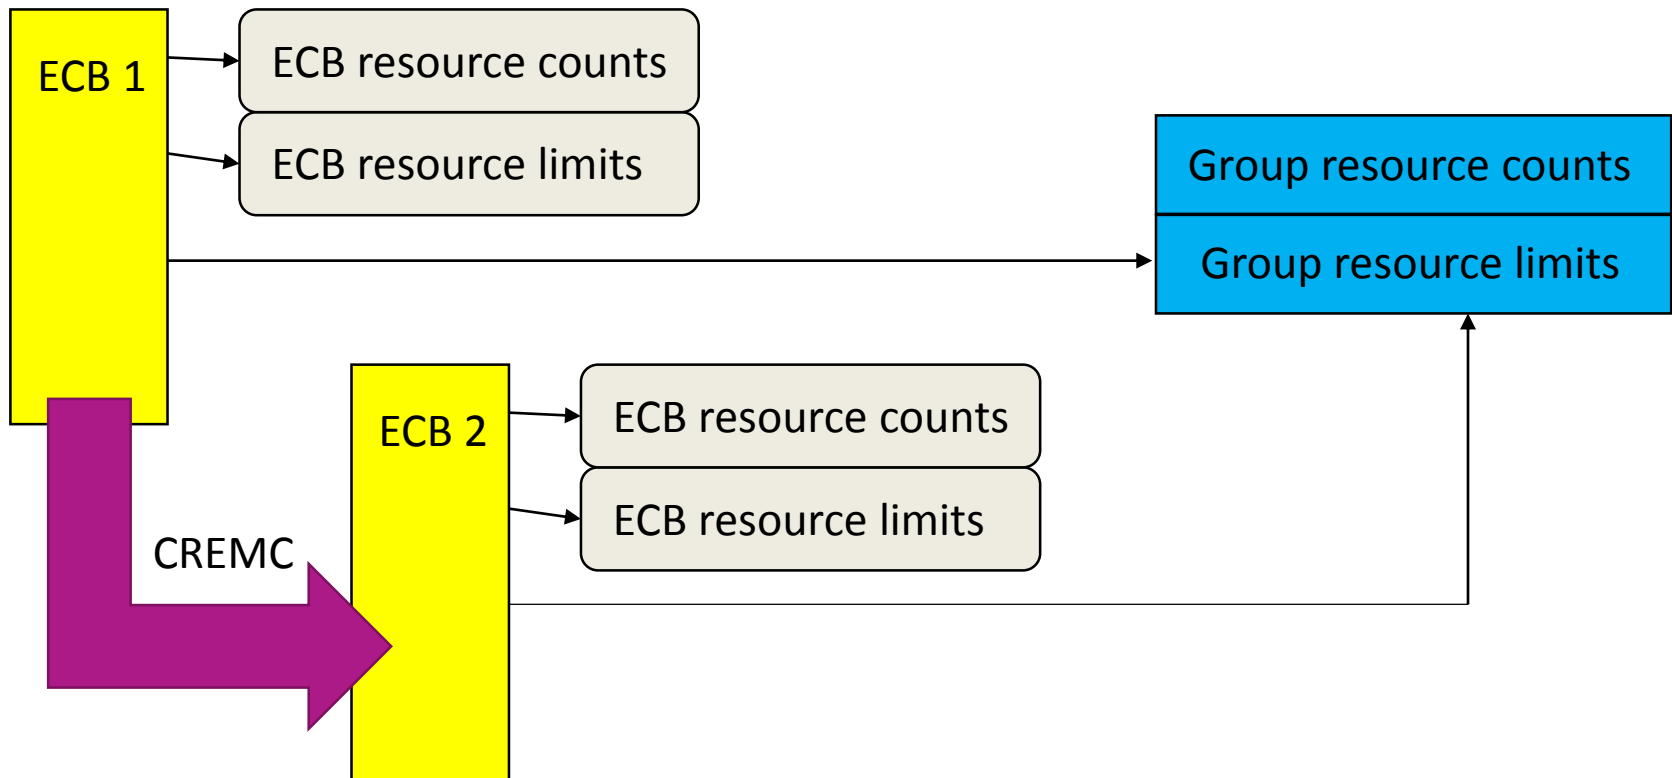
# Proposed ECB resource monitor enhancements

- Group would have a separate set of resource counts.
  - Group resource counts would be set when the group is created.
  - Initial group resource counts would be set to the counts in the original ECB when the group is created.
  - Group resource counts would be incremented in addition to incrementing counts for individual ECBs.
- Group would have separate resource limit table (RLT)
  - Based on named limit set
  - Managed by ZECBM

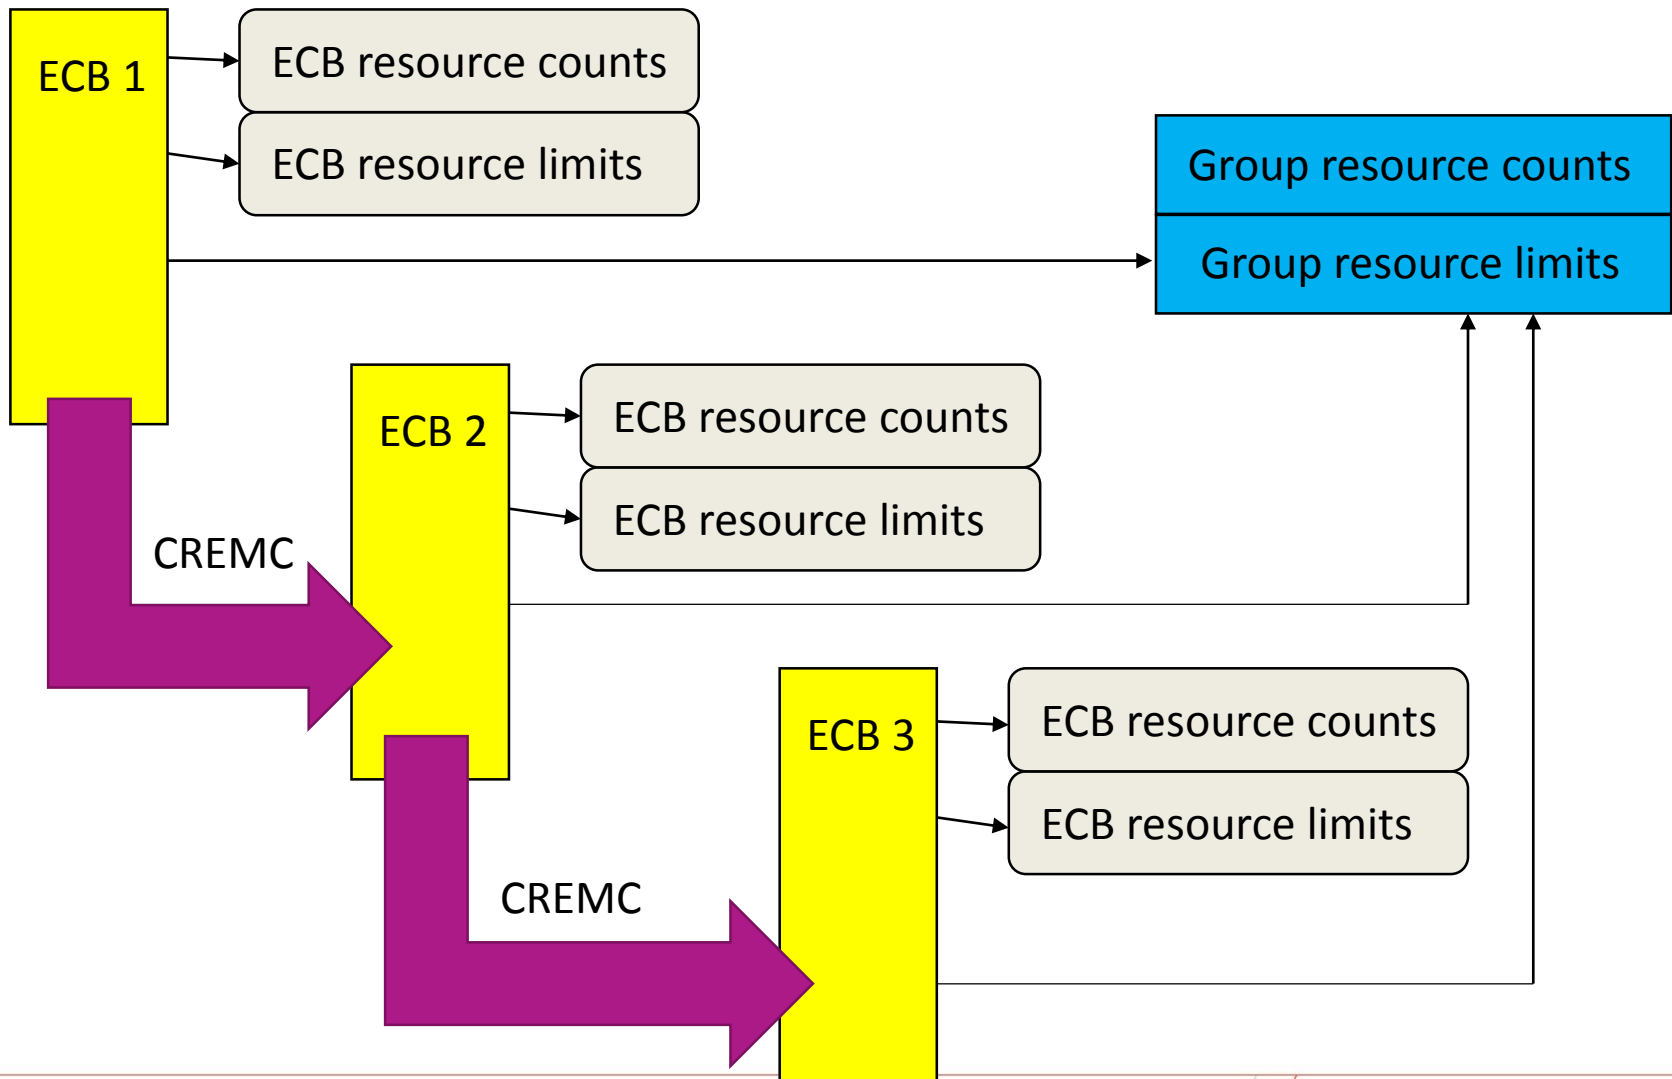
# Single ECB monitoring



# Group ECB monitoring starts



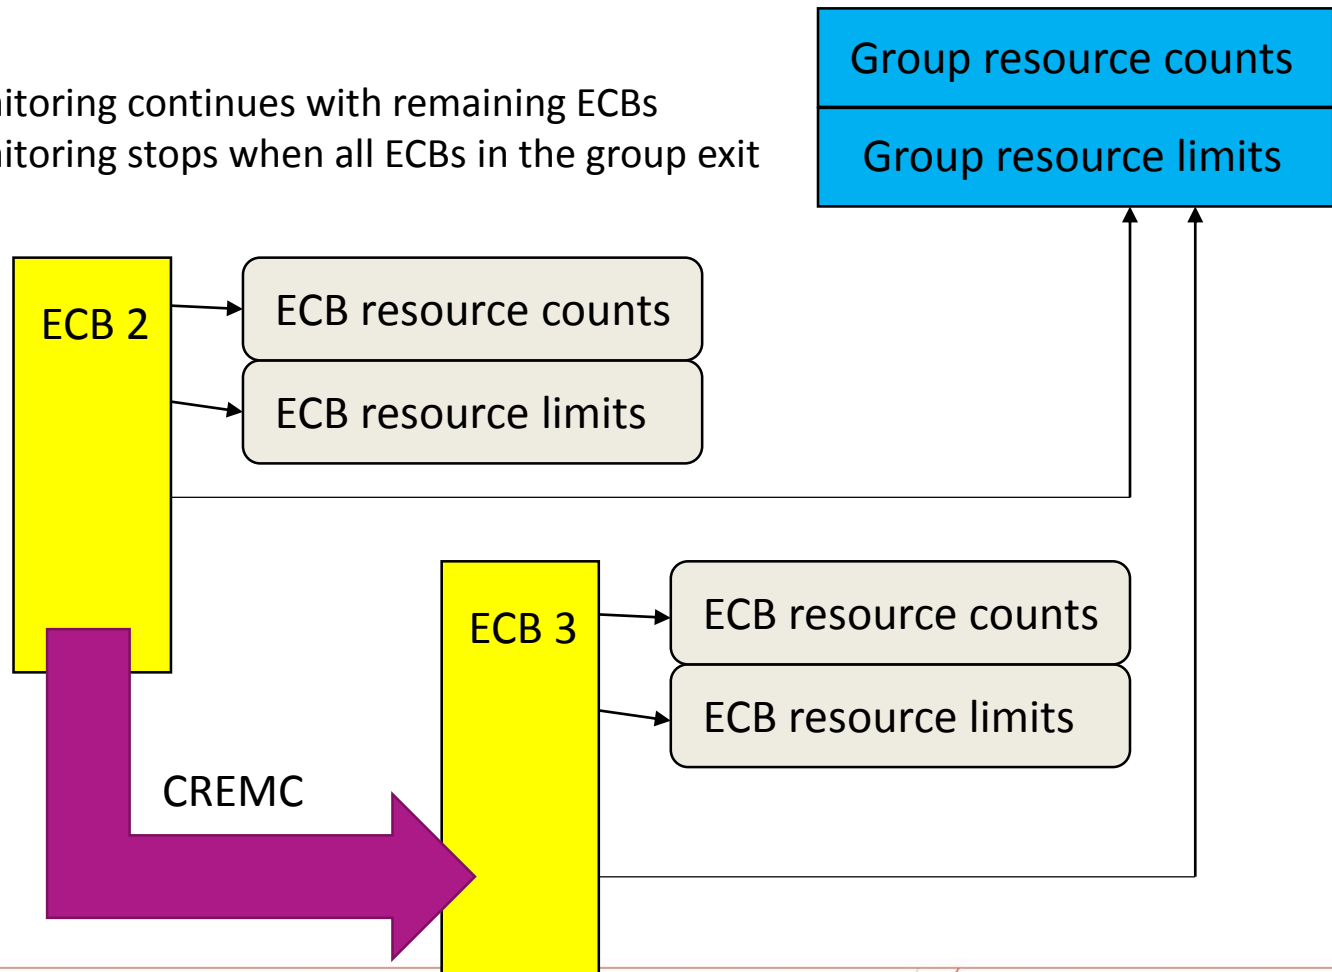
# Group ECB monitoring



# Group ECB monitoring

ECB 1 exits

- Group ECB monitoring continues with remaining ECBs
- Group ECB monitoring stops when all ECBs in the group exit



# Trademarks

- IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).
- *(Include any special attribution statements as required – see Trademark guidelines on <https://w3-03.ibm.com/chq/legal/lis.nsf/lawdoc/5A84050DEC58FE31852576850074BB32?OpenDocument#Developing%20the%20Special%20Non-IBM%20Tr>)*

## Notes

- Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
- This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
- Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
- This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.