



# Catalog Address Space: What, Why, Where

For GSE Large Systems Group  
March 1998

Paul Arnerich  
Tsd (UK) Ltd  
Paul\_TSD@compuserve.com

# Acknowledgements



- DFSMS/MVS: Managing Catalogs  
IBM - SC26-4914-nn
- Data Facility Extended: Catalog Diagnosis Reference  
IBM - SY26-3886-0 (March 1981)
- Data Facility Extended: Catalog Diagnosis Guide  
IBM - SY26-2887-0 (March 1981)

er...that's it

Except to say, there isn't much 'real' doco out there, more like OCO than doco



# CAS: Agenda



- WHAT
- WHY
- WHERE
- ISC
- CDSC
- 'STUFF'



# CAS: What ?

- It's cache for cats !
- A big fat chunk of processor storage !
- Typically top five 'earner' in the CPU and storage consumption game !
- aka 'CAS'
  - "CAN ABUSE STORAGE" ???
  - "CPU ABOUT STUFFED" ???

A BLOODY MYSTERY !



# CAS: What ?



- An MVS subsystem, started at IPL
- Can be stopped and restarted
- Has operator interface - "F CATALOG,....."
- A high CPU consumer
- A high storage consumer

...PERFORMANCE AT A PRICE....



# CAS: WHY ?



- ICF catalogs designed in 1979
- Not modified since (well...just a bit)
- Dataset open is a performance issue
- The ICF catalog structure has to be performance oriented as all dataset opens are via the ICF catalog
- Delivers performance for dataset open
- Simply, it reduces the number of I/O's needed to open datasets
- How ? Keeps most recently referenced catalog records in storage



# CAS: WHERE



- Uses processor storage as a massive cache
- Two types of cache;
  - ISC - In store cache
  - CDSC - Catalog data space cache
- A catalog can be in ISC or CDSC, not both
- Caches all eligible records, from all eligible catalogs in ISC or CDSC
- Default maximum number of catalogs is 50 but can be changed by modify command



# CAS: Eligibility

## CATALOGS

- ICF & TAPE Catalogs
- Not explicitly ineligible

## RECORDS

- For records accessed via direct read
- Not for update (DEFINE/DELETE/ALTER)
- Not for key GT or EQ (LISTC LVL)
- Not for sequential (LISTC)
- Not for non sequential

→ "Records will not be cached for any catalog when the request is for update, key greater-or-equal, non-sequential or sequential processing." - APAR OY14870



# CAS: Processing

- At dataset open, catalog record is placed in cache (ISC or CDSC)
- When catalog is updated, a timestamp is placed in the VVR for the catalog
- At end of update, timestamp is zeroed
- When catalog is referenced, VVR is checked... Always !
- If timestamp is present:
  - For ISC, all records are flushed from cache and read is from DASD
  - For CDSC, wait for timestamp to be zeroed



# CAS: Processing

- CAS has a maximum number of open catalogs, once reached they are discarded on least recently used basis
- Maximum is set by site, F CATALOG,CATMAX(nn) where nn is a hex value
- CATMAX(00) means no limit
- If CATMAX(00), catalogs are allocated in the user's address space, not CAS



# CAS: ISC



- In Storage Cache
- Uses processor storage to maintain a copy of catalog records that are most frequently read
- All catalogs are eligible unless in CDSC or explicitly removed from ISC
- At first catalog open, a 'chunk' of storage is acquired from CAS private
- This 'chunk' is fixed, records are discarded when the 'chunk' is full



# CAS: ISC

- At CAS startup, all aliases are placed in ISC
  - Guess 100 bytes each
- At CAS startup, all master catalog records are placed in ISC
  - Guess 140 bytes per non-vsam
  - Guess 1024 bytes per vsam (512 to 4096)
- These records are remain in CAS for the life of the MVS !!
- If you have junk in your Master Cat, you have junk in CAS !!



# CAS: ISC



- Use F CATALOG,UNALLOCATE to remove catalogs from ISC
- Good candidates are catalogs that are not frequently updated
- Any update caused all records for that catalog, to be flushed from CAS, if catalog is shared
  - Definition of shared to follow



# CAS: CDSC



- Catalog Data Space Cache
- Exploits VLF to maintain data space for storing catalog records
- To enable, update COFVLFxx member
- Catalog must be explicitly put in CDSC
- Master catalog is not a good candidate for CDSC, as “unpredictable errors may occur”



# CAS: CDSC

- Unlike the ISC, there is no limit per catalog
- Fixed 'chunk' of storage based on 'MAXVIRT' parameter of COFVLFxx
- Catalog records reside in same 'chunk', no isolation
- Once the CDSC is full, records are discarded on LRU algorithm
- Otherwise, records are only removed when the record is deleted or the catalog is unallocated from CAS
- Default size of data space is 1Mb



# CAS: SHARED CATALOGS

- Beware of shared catalogs in ISC, any update causes the entire catalog to be flushed from the ISC on all 'sharing' systems
- A catalog is 'shared' if it resides on a shared device and is defined SHR(3 4)
  - i.e. every Catalog in the known world !!
- If a shared catalog is updated and is in CDSC, all sharing CAS's are notified and the changed record is reloaded in all CDSC's
- This process uses the VVDS sharing subcell as a 'no entry' flag until update is complete.



# CAS: RECOVERY



- Remove catalogs from CAS before attempting recovery
- F CATALOG,UNALLOCATE
  - Removes Catalog from CAS until next IPL
- F CATALOG,CLOSE
  - Flushes all records from CAS, then starts caching new opens



# CAS: TUNING

## ■ COFVLFxx

- MAXVIRT
- CDSC eligible Catalogs

## ■ CAS Service Task Lower Limit

- Specified in SYSCAT operand of LOADxx
- Allows CAS to dynamically lower the number of service tasks (open Catalogs) if CAS is not busy
- Presumably allows dynamic upping ?

## ■ CATMAX

- System default is 50 (x'32')
- Modify via F,CATALOG,CATMAX(xx)
- nn is a hex number



# CAS: PERFORMANCE

- For the CDSC, VLF overhead is high, so use only where hit ratio can justify it
- IBM recommends a 20% hit ratio is “worth” it
  - 30% may be more realistic
- Monitor via “F CATALOG,REPORT,VLF
- Monitor over a period of time, it is a cache after all
- Limit Jobcat and Stepcat, these place high workload on CAS
- Monitor CAS usage over batch and online day cycles



# CAS: PERFORMANCE

- Consider two catalogs
- CATALOG.DB2T: 1000's of DB2 tablespaces
  - Test DB2
  - Used on one MVS only
  - Frequent updates (defines)
- CATALOG.BATCH: 100000's of GDG's.
  - Write once, read never
- Are either of these catalogs good ISC candidates ?
- Are either of these catalogs good CDSC candidates ?



# CAS: PERFORMANCE



- Is the ISC a good thing ?
- Is the CDSC a good thing ?

....Performance, at a price.....

