What's *REALLY* New with VSE/VSAM

Dan Janda
The Swami of VSAM

The Swami of VSE/VSAM

Dan Janda
VSE, VSAM and CICS
Performance Consultant

RR 2 Box 49E
Hardin Road
Montrose, PA 18801-9624

(570) 934-2862
theswami@epix.net

World Alliance of VM and VSE
Winston-Salem, North Carolina
April, 2003

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Abstract and References

- **Abstract:**
  New Features in recent VSE/ESA releases add significant function. Critical issues -- batch window processing, handling of large data sets, and VSE constraints on number of devices have been addressed by IBM developers.

  In this presentation, we'll highlight a few of these with the most impact on your system's performance.

  VSE/VSAM in VSE/ESA 2.5, 2.6 and 2.7 are examined in some detail.

- Bibliography:
  - VSE/VSAM Commands
    Version 6 Release 4
    SC33-6731
  - VSE/VSAM User's Guide and Application Programming
    Version 6 Release 4
    SC33-6732-00
  - VSE/ESA Version 2 Release 6
    Release Guide
    SC33-6718-02
    and earlier versions of these manuals
- Trademarks:
  - IBM, VSE, VSE/ESA, ESA, CICS, ESS and RAMAC are trademarks or registered trademarks of the IBM Corporation
  - The Swami of VSAM is a trademark of Dan Janda mailto:theswami@epix.net

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New VSAM Features

- New features
  - Extra Large Data Set Support
  - LSR Buffer Hashing (VSE/ESA 2.5)
  - 3390-9 Support (VSE/ESA 2.6)
  - IXFP/Snap Shot for RVA (VSE/ESA 2.5)
  - Flash Copy for ESS (VSE/ESA 2.5)
  - Fast Copy (FCOPY) via Flash Copy or Snap Shot (VSE/ESA 2.6)
- Features impacted
  - IMBED/NOIMBED for clusters
  - REPLICATE/NOREPLICATE for clusters
  - IMBED/NOIMBED for catalogs
  - The interpretation of these options changes with VSE/ESA 2.6

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

New Features Extra Large Data Set

- VSE/VSAM Extra Large Data Set support provides
  - Relief from 4 GB limit for KSDS files accessed by key
  - Still supports up to 123 extents
  - Each extent could be a whole volume
  - Now could be up to 123 volumes, in theory
  - With 3390m3 volumes ~ 492 GB. 3390m9 ~ 1.4 TB.
- No measurable additional CPU processing cost
- Can be used with Compressed KSDS if desired
- No support for ESDS, RRDS, SAM-ESDS, or KSDS accessed by RBA
- Internal changes -- use CI number instead of RBA number as internal pointer in index
- No visible external changes other than specification
  DEFINE CLUSTER ... EXTRALARGEDATASET (or XXL)
  and display of CI numbers instead of RBAs in LIST/PRINT

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

New Features
LSR Buffer Hashing

- VSE/VSAM LSR Buffer Hashing provides
  - Relief from increased CPU cost as LSR pool grows larger
  - Ability to drive VSAM physical I/O costs very low for direct processing
  - Significant changes in the trade-offs between extended use of 31-bit storage and CPU cost of managing the buffer searches
- No external changes -- support is automatically included in VSE/ESA 2.5 (with or without using CICS Transaction Server)
- Merely specify additional LSR buffers (with CICS TS, use the RDO specifications that permit separate subpools for data and index buffers, too)

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

New Features
LSR Buffer Hashing

- LSR Hashed Search vs. LSR Sequential Search

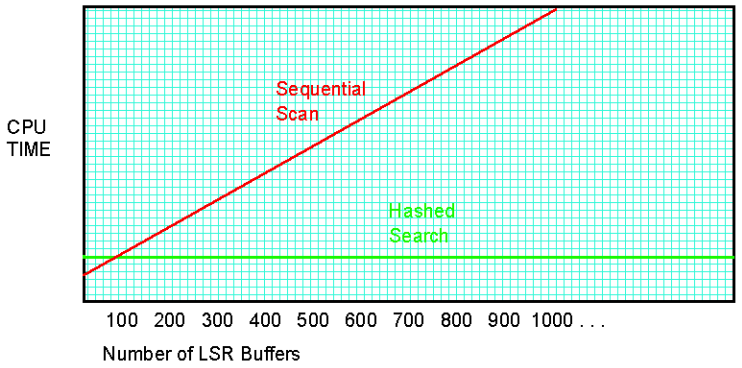

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

New Features
LSR Buffer Hashing

- LSR benefits for multi-string access
  - e.g. CICS file accesses
- LSR benefits for direct processing
  - No read-ahead, no write-behind (sequential processing)
  - Use NSR for sequential processing
- NSR does not use hashing search
  - Limit BUFND and BUFNI to less than 150 unless a few more BUFNI or BUFND will contain entire index or data
  - Ideal BUFND is the number of CIs per CA (CIs/CA in LISTCAT)
  - Ideal BUFNI (per string -- for multi-string files, multiply!) is
    - for sequential - 1
    - for direct - minimum acceptable - Number of index levels
    - for direct - maximum useful - Number of index records
    - Dependent on access pattern

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features
3390m9 Support

- Before VSE/ESA 2.6, VSE/VSAM volume size use was limited by Space Mapping limit of 65,535 tracks
  - VSAM used halfword pointers, which set bitmap size limit
  - Internal changes were made to permit larger bitmap sizes
  - Full support of 3390m9 devices
    - 10,017 cylinders
    - 150,255 tracks
  - In fact, even larger devices could be supported by this change
    - Potentially,
      - 65,535 (64K-1) cylinders
      - 65,535 (64K-1) tracks per cylinder
      - 65,535 (64K-1) bytes per track
      - 4,294,836,225 (4 G) tracks per device
      - 281,462,092,005,375 (280 T) bytes per device
    - See the VSE/ESA 2.7 Preview announcement

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

2003

Really New Features 3390m9 Support

- With virtual devices, number of physical devices limit within VSE (255 devices with old supervisor, 1024 devices with VSE/ESA changes) is often a constraint
- Also provides relief for 492 GB size constraint for Extra Large Data Set support!
- Notes:
  - VSE/ESA tends to work better when spreading I/O across more devices -- due to VSE device queueing -- if you don't need more than 1024 I/O devices
  - It would be beneficial to use as many devices as possible, spreading activity and reducing queuing delays
  - Extensive use of hardware and software caching and VSAM LSR buffer pools reduces I/O activity and queuing delays

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

2003

Really New Features 3390m9 Support

- Catalog Index Space Allocation Increased
  - Primary allocation
    - Increased to at least 4% of catalog primary data allocation
  - Secondary allocation
    - To be same size as index primary allocation amount
- Catalog Volume Space Map kept and increased in size
  - A catalog CI (512 bytes) keeps a space map segment mapping 3520 tracks
  - Number of CIs used for space mapping will increase with large DASD size
  - Previous limit was 19 segments, enough to map 65,535 tracks per volume
  - 3390-9 has 150,255 tracks, requiring 43 catalog records for the space map

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

2003

Really New Features 3390m9 Support

- Old (65,535 track) algorithms used for
  - DASD volumes smaller than 65,535 tracks, OR
  - Space on the volume was defined in the current catalog before VSE V2.6
- New (65,535 cylinder) algorithms used for
  - DASD volumes larger than 65,535 tracks, AND
  - No space on the volume was defined in the current catalog before VSE 2.6
- New algorithm effects:
  - Space allocations that resolved to tracks are rounded up to cylinders
    - Cluster with allocation of TRACKS(1 1) becomes CYLINDERS(1 1)
    - Track-level boundaries of extents will be rounded to cylinder boundaries
  - Catalogs residing on large DASD will have minimum allocations of
    - 4 cylinders for the catalog's data component
    - 1 cylinder for the catalog's index component

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

2003

Really New Features 3390m9 Support

- Large DASD Performance Considerations for KSDS
  - KSDS performance depends upon Data Control Area Size -- larger is better
  - One Index Control Interval controls the sequence for each Data Control Area
  - Better keyed performance when Index I/O is minimized
    - Smaller Data Control Interval sizes require larger or more Index CIs
  - KSDS CA size on large DASD will be one cylinder whenever possible.
    - Primary and secondary allocations are rounded up to cylinder multiples
  - Minimum Data CI size for KSDS will be increased for large key lengths:
    - Key length between 7 and 35 bytes requires at least 1024 byte Data CI
    - Key length between 36 and 55 bytes requires at least 2048 byte Data CI
    - Key length greater than 55 bytes requires at least 4096 byte Data CI
  - Too small BUFFERSPACE value can restrict the size of a Data CA
    - Two Data CIs plus One Index CI must fit.
    - If too small, the Index CI is reduced to fit
    - This can result in CA size being reduced
- RECOMMENDATION: Don't specify BUFFERSPACE unless it is much larger than what VSAM calculates, for special AIX issues

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features 3390m9 Support

- RECOVERABLE Catalog option
  - Support for RECOVERABLE catalogs will be changed
  - No Catalog Recovery Area can be defined on Large DASD
  - Existing RECOVERABLE catalogs and CRAs continue to be supported
  - New RECOVERABLE catalogs can be created on small DASD
    - A RECOVERABLE catalog cannot own space on Large DASD
    - Otherwise, a CRA would be required on a Large DASD
  - RECOVERABLE catalogs cannot reside on Large DASD
  - RECOVERABLE catalogs cannot own space on Large DASD
  - RECOMMENDATION: Don't use recoverable catalogs anyway --
    - They provide recovery capabilities if the catalog is damaged
    - They do NOT provide recovery capabilities if the CRA is damaged

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features 3390m9 Support

- IMBED
  - Causes first track of each CA to contain a replicated copy of the low level index (Sequence Set) CI for that CA
  - Objective was to reduce seek delay and rotational delay for sequence set
  - NOT RECOMMENDED today (for the same reasons as for REPLICATE)
  - Definition of IMBEDded indexes will no longer be done for
    - Catalogs
    - Clusters (including Alternate Indices)
  - Existing catalogs and clusters with IMBED defined prior to VSE/ESA 2.6 will still be accepted for processing
  - Existing DEFINE commands with IMBED specified will be accepted for compatibility, but the IMBED specification will be ignored
  - Default specification for new catalog definitions will be NOIMBED

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features 3390m9 Support

- LISTCAT changes
  - DEVTYPE-----BIG-3390
  - Volume Group (DATA, INDEX) TRACKS shows CYLINDERS for Large DASD
  - Space Map (SPC-MAP) will become CYL-SPC-MAP for Large DASD
- REPLICATE
  - Causes each index CI to be repeated as many times as possible on a disk track
  - Objective was to significantly reduce latency (rotational) delay for index set
  - NOT RECOMMENDED today if
    - You have any kind of disk cache
    - One copy is enough if the track is in cache
    - A cache block will contain many index records if not replicated, but many copies of only one index record if replicated
    - You can afford enough index buffers (BUFNI=nn) to contain much of the index (caching directly in VSAM buffers is most efficient)

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features Snap Shot & Flash Copy

- Snap Shot & Flash Copy Advantages
  - After Snap of disk volume is complete, backup processing can be started immediately and can run concurrently with online processing
  - Copy time is very short -- minutes, even seconds, shortening batch window
  - No risk of losing data
  - Faster access to snapped (target) datasets -- less contention for device
  - No catalog changes on the target catalog
  - Complex, error prone, time consuming changes not needed on snapped catalog and datasets
  - When Snap is made, the contents of the snapped (target) volumes remain identical to the contents of the source volumes
  - Backups are made from the "frozen", snap (target) data while other processing continues
  - IDCAMS SNAP together with IDCAMS BACKUP using a synonym list significantly reduces time your system and data is unavailable

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features Snap Shot & Flash Copy

- Snap Shot Example

```
//JOB SNAP and BACKUP from snapped Volumes
//ASSGN SYS005,180
//DLBL IJSYSUC,'VSESP.SNAP.CATALOG',,VSAM
//EXEC IDCAMS,SIZE=AUTO
/*First:do the SNAPSHOT */-
SNAP
SOURCEVOLUMES(SYSWK1,DOSRES)-
TARGETVOLUMES(VOLSN1,VOLSN2)
/*Second:Synonym Name for the snapped Catalog */-
IMPORT CONNECT OBJECTS((VSESP.SNAP.CATALOG -
VOLUMES(VOLSN1)DEVT(3390)))-
CATALOG(VSAM.MASTER.CATALOG)
/*Third:Backup from snapped volumes */-
BACKUP (*)-
SYNONYMLIST(-
SOURCEVOLUMES(SYSWK1,DOSRES)-
TARGETVOLUMES(VOLSN1,VOLSN2)-
CATALOG(VSESP.USER.CATALOG)-
SYNCATALOG(VSESP.SNAP.CATALOG))
/*
/&
```

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Really New Features Snap Shot & Flash Copy

- FASTCOPY exploitation of Snap Shot & Flash Copy
  - Prerequisites:
    - IXFP/Snap Shot priced feature for RAMAC Virtual Array (RVA)
    - Included with VSE/ESA V2.6 for ESS Flash Copy
  - COPY ALL
  - COPY VOLUME
    - Snap Shot or Flash Copy will NOT be used if the following options are used:
      - EXCLUDE
      - NOVSAM
      - NOEXPIRED
  - If a full volume COPY request is made, Snap Shot or Flash Copy is installed, and restricted options are NOT used, FCOPY will automatically exploit the new and faster copy technique
  - Otherwise, normal FCOPY processing will be done

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM

Winston-Salem, NC 2003

Contacting the Presenter

- For more information...
  - You can contact the Swami by e-mail

    theswami@epix.net

  - He's building a web site about VSE/VSAM issues

    http://business.epix.net/~theswami

  - His knowledge and experience can help you too

Copyright 2001--2003 by Dan Janda
The Swami of VSE/VSAM