

**Chapter
1**

**VSAM COMPONENTS
AND EVOLUTION**

*Get on the
Fast Track!*



TM

**SYS-ED/
Computer
Education
Techniques, Inc.**

Objectives

You will learn:

- How to assess and define VSAM performance.
- The physical components of VSAM.
- The role of a CI: Control Interval.
- How data resides in the CI.
- Control Interval Control Information.
- The role of the CA: Control Area.
- VSAM file structures - KSDS: Key Sequenced Dataset, ESDS: Entry Sequenced Dataset, and RRDS: Relative Record Dataset, and LDS: Linear Dataset.
- AIX: Alternate Index.
- VSAM Type Comparison.
- Index Levels and Physical I/O's.
- 3380 and 3390 characteristics.
- Measuring DASD response time.

1 VSAM Performance: A Definition

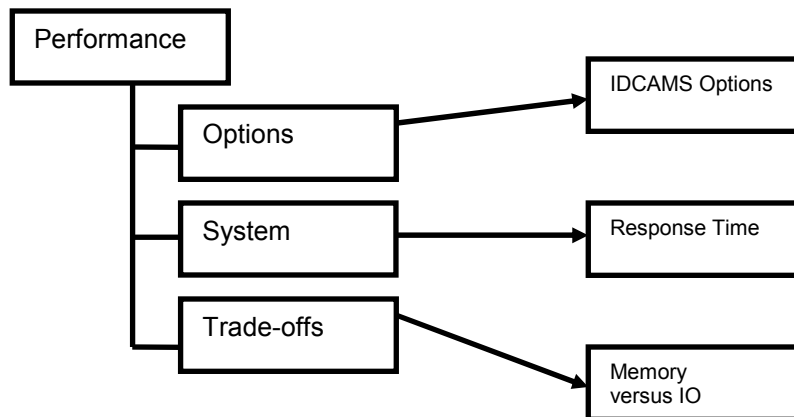
Performance can be defined as a function of:

DASD space requirements	Input/Output speed	Program design
-------------------------	--------------------	----------------

Tuning is the process of balancing options in each of these areas to improve overall performance.

Successful tuning will depend on:

- Understanding performance options.
- Understanding systems requirements.
- Selection of trade-offs to well defined results.



2 History of VSAM

2.1 Single Set of Access Method Programs

VSAM: Virtual Sequential Access Method was introduced in the early 1980's by IBM as a collection of sequential, indexed, and direct-access dataset organizations for its MVS and DOS/VSE mainframe operating systems.

IBM designed VSAM to integrate and be extended with access methods, utilities, and special processing and handling:

- For the addition of new hardware devices with older access method software.
- With multiple access methods required multiple sets of utility programs.

Access Method Services - IDCAMS is the comprehensive utility for VSAM and non-VSAM files.

2.2 Integrated Catalogs and Space Management Consolidation

Consolidation of space management facilities is required in order to reduce maintenance effort and increase compatibility with its component features.

VSAM replaced VTOCs, CVOLs, and some JCL in order to provide a central control of DASD space.

The VSAM catalog serves as a namespace repository for the entire DASD environment.

In addition to providing integrated centralization, the catalog entry structure allows for:

- More information than VTOC records.
- The relating of physical datasets to logical dataset structures such as KSDS clusters and paths.
- Multi-volume dataset handling.

3 Physical Components

VSAM builds all files with the same physical elements:

CI: Control Interval	CA: Control Area
----------------------	------------------

Control Area

CI: Control Interval
CI: Control Interval
CI: Control Interval
CI: Control Interval
CI: Control Interval

3.1 CI: Control Interval

The CI: Control Interval is the basic unit of data transfer between main storage and DASD devices.

Size:

- All CIs in a dataset component are the same size.
- Size can be selected by the user or by VSAM.
- Size is device-independent; a CI may be larger than the DASD track.
- Size determines the number and size of physical records used to hold a CI on DASD.

Content:

- One or more logical records
 - For spanned records, a part of a single spanned record.
- Unused space which may be used in the future.
- Control information for the CI and the logical records it contains.

3.2 How Data Resides in the CI

Logical records are placed in the data CI in key sequence from left to right. RDF fields are placed in the data CI in the same sequence as their related records, but from right to left, beginning with the first byte to the left of the CIDF.

Adjacent records of the same length are represented by a pair of RDFs indicating length and record count. The CIDF indicates the amount of FREESPACE remaining in the CI. In empty CAs, the CIDF of the first CI is set to zero.

Spanned records occupy more than one CI and always start at the beginning of a CI.

- Are represented by 2 RDFs containing segment length and update number segment.
- An update number is used to secure the integrity of all segments of the record.
- Unused space at the end of the last CI for a spanned record can be used only to lengthen the existing record.

Key-Sequenced Datasets

Data Component

Control Interval Structure

logical record	logical record	logical record	unused space	R D F	R D F	R D F	CIDF
----------------	----------------	----------------	--------------	-------------	-------------	-------------	------

No two consecutive records with same length.


logical record	logical record	logical record	logical record		R D F	R D F	CIDF
----------------	----------------	----------------	----------------	--	-------------	-------------	------

All records same length.




3.3 Control Interval Control Information

One CIDF: Control Interval Definition Field per CI contains the total length of the data in the CI, and the total length of the unused space. There is no requirement for there to be any data.




One or more RDF: Record Definition Fields per CI, which specify a record, record slot, or spanned record segment. They are placed in the CI from right to left, in order that the right-most RDF specifies the left-most record.



When two or more records of the same length are adjacent in the CI, one RDF will specify the record, and the next RDF will indicate how many of these records are in sequence in the CI.

An RDF consists of a control field and a binary number. The number indicates the control field content.



3.4 Record Definition Field

Bit	Meaning
0	Reserved.
1	0 - There is not a paired RDF to the left of this RDF. 1 - There is a paired RDF to the left of this RDF.
2,3	00 - The record is not spanned. 01 - It is the first segment of a spanned record. 10 - Last segment of a spanned record. 11 - Intermediate segment of a spanned record.
4	0 - The binary number in this RDF provides the length of the record 1 - The binary number in this RDF provides the number of unspanned records of the same length, or the update number of a spanned record segment.
5	0 - For an RRDS, the slot described by this RDF is empty. 1 - The slot is occupied.

3.5 Control Field

Control Field	Number
bit 4 = 0	Provides the length of record, segment, or slot described by this RDF.
bit 4 = 1 bits 2,3 = 0	Provides the number of records with the same length.
bit 4 = 1 bits 2,3 not 0	Provides the update number of the spanned record segment in this control interval.

3.6 CA: Control Area

The Control Area is the:

- Unit of file organization containing a fixed number of CIs for a VSAM file.
- Determining factor in index structure for KSDS.
- Always a contiguous area of DASD storage.

All CIs reside in CAs; a CA may contain CIs with data as well as empty CIs.

Size is determined implicitly by space allocation:

- minimum - 1 track
- maximum - 1 cylinder (recommended)

4 VSAM Structure

4.1 File Structures

VSAM manages four basic file structures, or clusters:

- KSDS: Keyed Sequenced Dataset
- ESDS: Entry Sequenced Dataset
- RRDS: Relative Record Dataset
- LDS: Linear Dataset

The KSDS replaces ISAM; the ESDS replaces QSAM; the RRDS replaces BDAM.

There also is an additional file organization:

- AIX: Alternate Index

4.2 KSDS: Key Sequenced Dataset

A KSDS: Key Sequenced Dataset is composed of two logical components: data and index.

- Index component contains records which relate logical keys to physical locations.
- Records are identified by unique logical key, or can be addressed by a RBA: Relative Byte Address which may change.
- Data records may be variable or fixed length.
- If data records are variable, update may change record length.
- Retrieval may be sequential or direct.
- Records are added in place by key sequence using distributed FREESPACE.
- Records deleted are physically erased, and occupied space is reclaimed as reusable FREESPACE.
- Spanned records are supported.

4.3 ESDS: Entry Sequenced Dataset

An ESDS: Entry Sequence Dataset is composed of a single logical component.

- Records are stored in load or add sequence.
- Records are identified by RBA.
- Data records may be variable or fixed length.
- Record updates occur in place; but update may not change record length.
- Retrieval may be sequential or direct.
- Records are added at the end of the file.
- Individual records may not be deleted.
- Spanned records are supported.

4.4 RRDS: Relative Record Dataset



A RRDS: Relative Record Dataset is composed of a single logical component, formatted into fixed length slots.

- Records are identified by RRN: Relative Record Number slot.
- Data records may be fixed length or variable length.
- Record updates occur in place; however, an update may not change record length, unless it is a variable length record.
- Retrieval may be sequential or direct.
- Records are added in empty slots.
- Record deletion marks a slot as empty.
- Spanned records are not supported.


4.5 LDS: Linear Dataset

A LDS: Linear Dataset contains data that can be accessed as byte-addressable strings in virtual storage. It is a VSAM dataset with a control interval size multiple of 4096 bytes.


- An LDS has no imbedded control information in its CI; there is no RDFs and CIDs.
- Logical records must be blocked and deblocked by the application program.
- Logical records are transparent from VSAM's point of view.
- An LDS is a non-VSAM file with some of the VSAM facilities, such as the use of IDCAMS and VSAM specific information in the catalog.
- The most common LDS utilization is by DB2.




ESDS and RRDS



ESDS



RRDS



4.6 AIX: Alternate Index

An AIX: Alternative Index is structured as a KSDS. It is an index to data in another base file, using an alternate key. The base file may be a KSDS or an ESDS and the connection to base cluster is by a PATH definition in the VSAM catalog.

- Alternate keys may be unique or non-unique.
- Record updates can occur synchronously with or without the base file.
- The file may be processed independently, or with its base through a PATH.
- A base cluster may have multiple alternate indexes.

4.8 VSAM File Structure Comparison

Category	KSDS	ESDS	RRDS
SEQUENCE OF DATA RECORDS	Primary Key	Entry	Relative Record Number
TYPE OF ACCESS	Sequential or Direct Key or RBA	Sequential or Direct RBA	Sequential or Direct Relative Number
RECORD FORMAT	Fixed Variable Spanned	Fixed Variable Spanned	Fixed
USE FREE SPACE TO:	Insert Update	Not Used	Not Used
ADD NEW RECORDS	Anywhere	EOF	At specified slots
CHANGE RECORD SIZE	Yes	No	No
DELETE RECORD	Yes Space is reused	No	Yes. Slot is set to X '00'
ALTERNATE INDEX	Yes	Yes	No

4.9 Index Levels and Physical I/O's

KSDS Levels	1	2	3
READ	1	3/2	4/2
WRITE	2	4/3	5/3
REWRITE	2	4/3	5/3
DELETE	2	4/3	5/3
START	1	3/1	4/2
READ NEXT	0,1,2	0,1,2	0,1,2

File Type	ESDS	RRDS
READ	1	1
WRITE	1,2	2
REWRITE	2	2
READNEXT	0,1	0,1
START	1	1
DELETE	--	2

4.10 3380 and 3390 Characteristics

Description	3380	3390
Tracks/Cylinders	15	15
Available Bytes Per Track	47,476	56,664
Cylinders/Volume	1770	1113/2226

5 DASD Response Time

In z/OS, DASD response time is measured in four components:

IOS queue time	Is the delay experienced within the z/OS operating system waiting for a device to become logically available?
Pending time	Is the delay experienced in the channel system waiting for a path to the device to become physically available? A device may be logically available but not physically accessible to shared DASD, control unit or channel contention.
Disconnect time	Is the physical delay experienced while the device positions itself to the proper cylinder and record? These delays can be elongated by RPS misses, which is an extra revolution that occurs due to control unit and channel contention.
Connect time	Is time when the channel, control unit and I/O device are busy for the I/O request? This time includes data transfer, hardware protocols and some search delays.