Chapter
1

INTRODUCTION

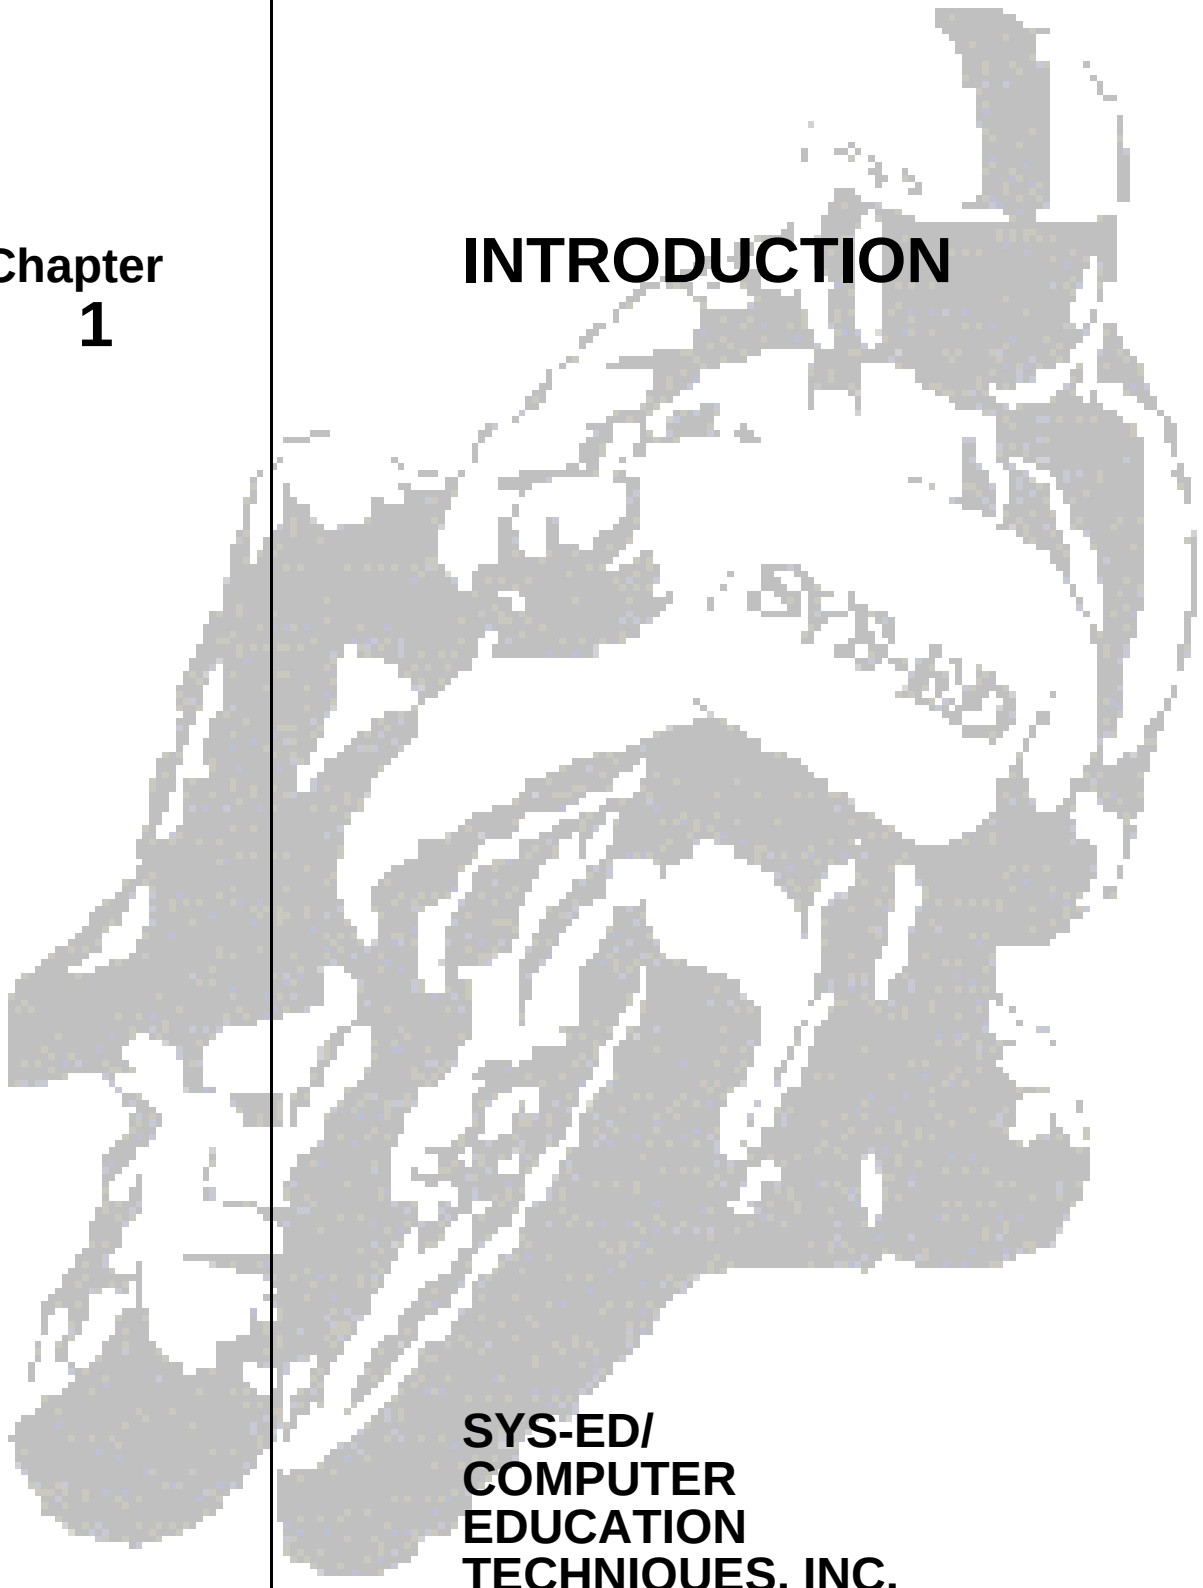

SYS-ED/
COMPUTER
EDUCATION
TECHNIQUES, INC.

Objectives:

You will learn:

- A definition of VSAM performance.
- Integrated catalogs and space management consolidation.
- Physical components: CI - Control Interval and CA - Control Area.
- VSAM structures.
- KSDS: Key Sequenced Data Set.
- ESDS: Entry Sequenced Data Set.
- RRDS: Relative Record Data Set.
- AIX - Alternate Index.
- VSAM type comparison.
- User catalog.
- ICF catalogs.
- Index levels and physical I/O's.
- DASD response time.

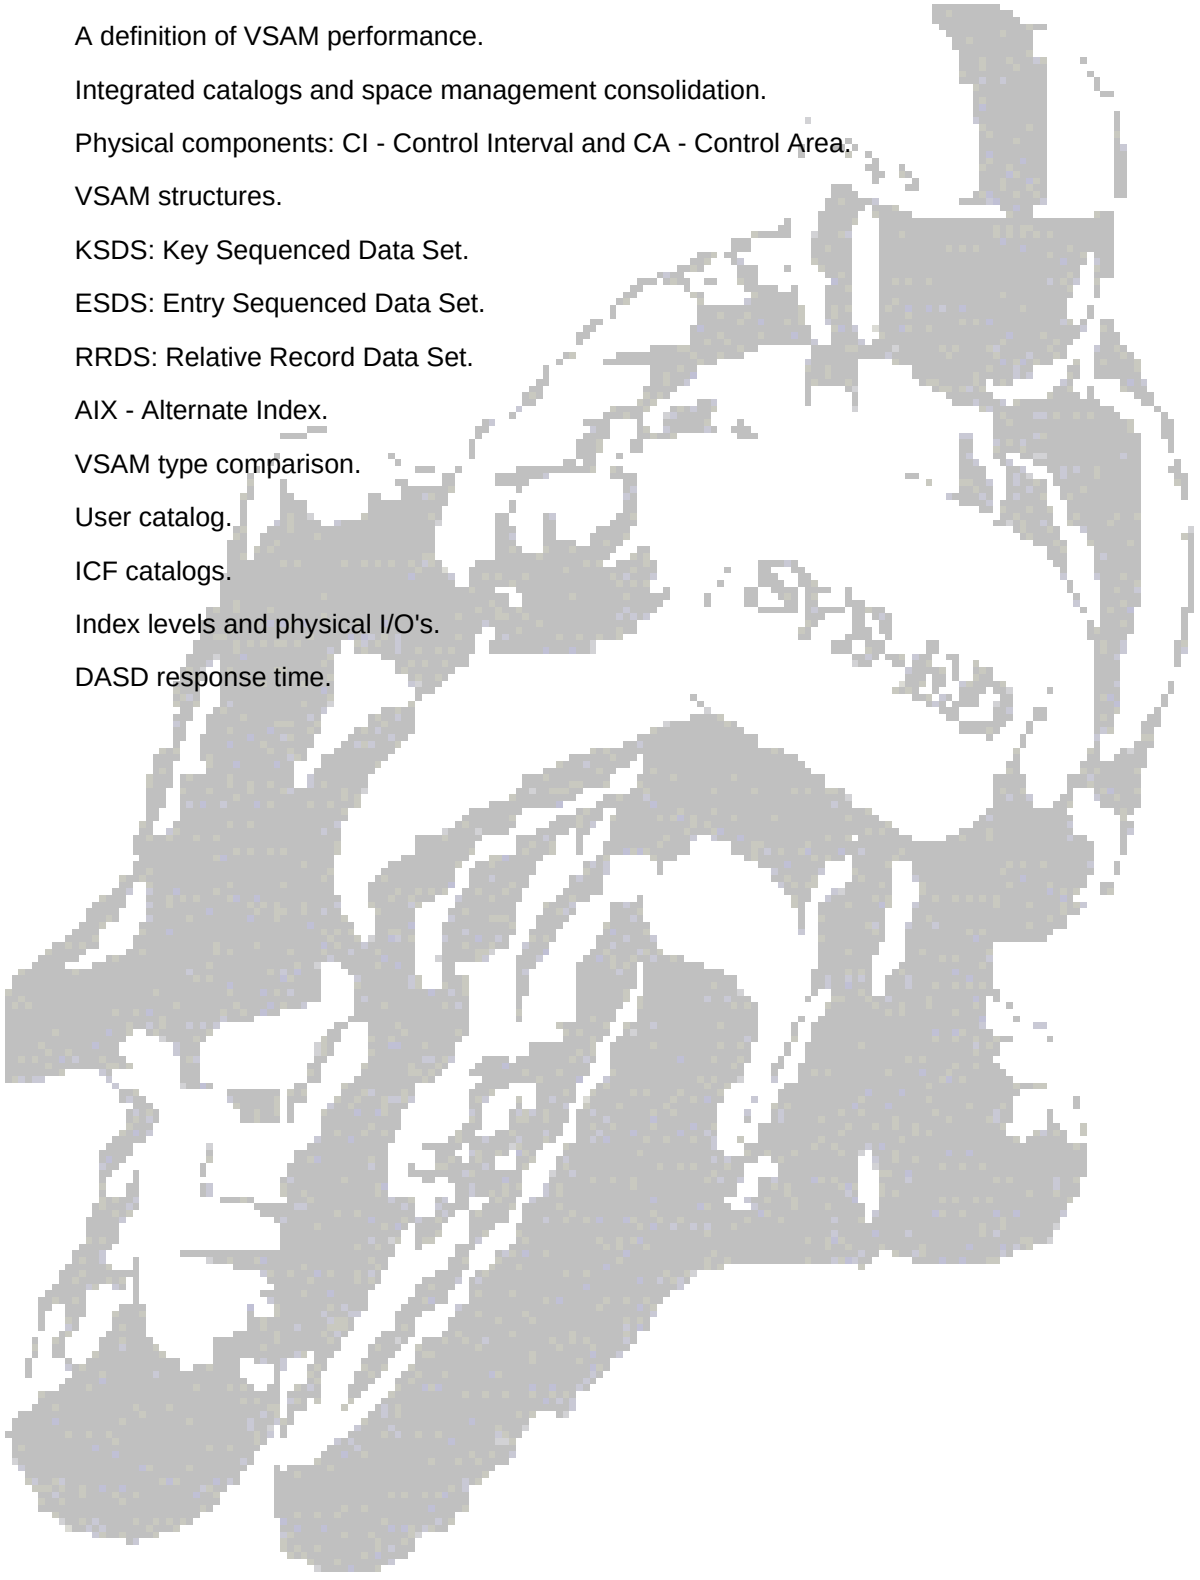

1 VSAM Performance: A Definition

Performance can be defined in several ways, as a function of:

- DASD space requirements.
- I/O 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 achieve desired results.

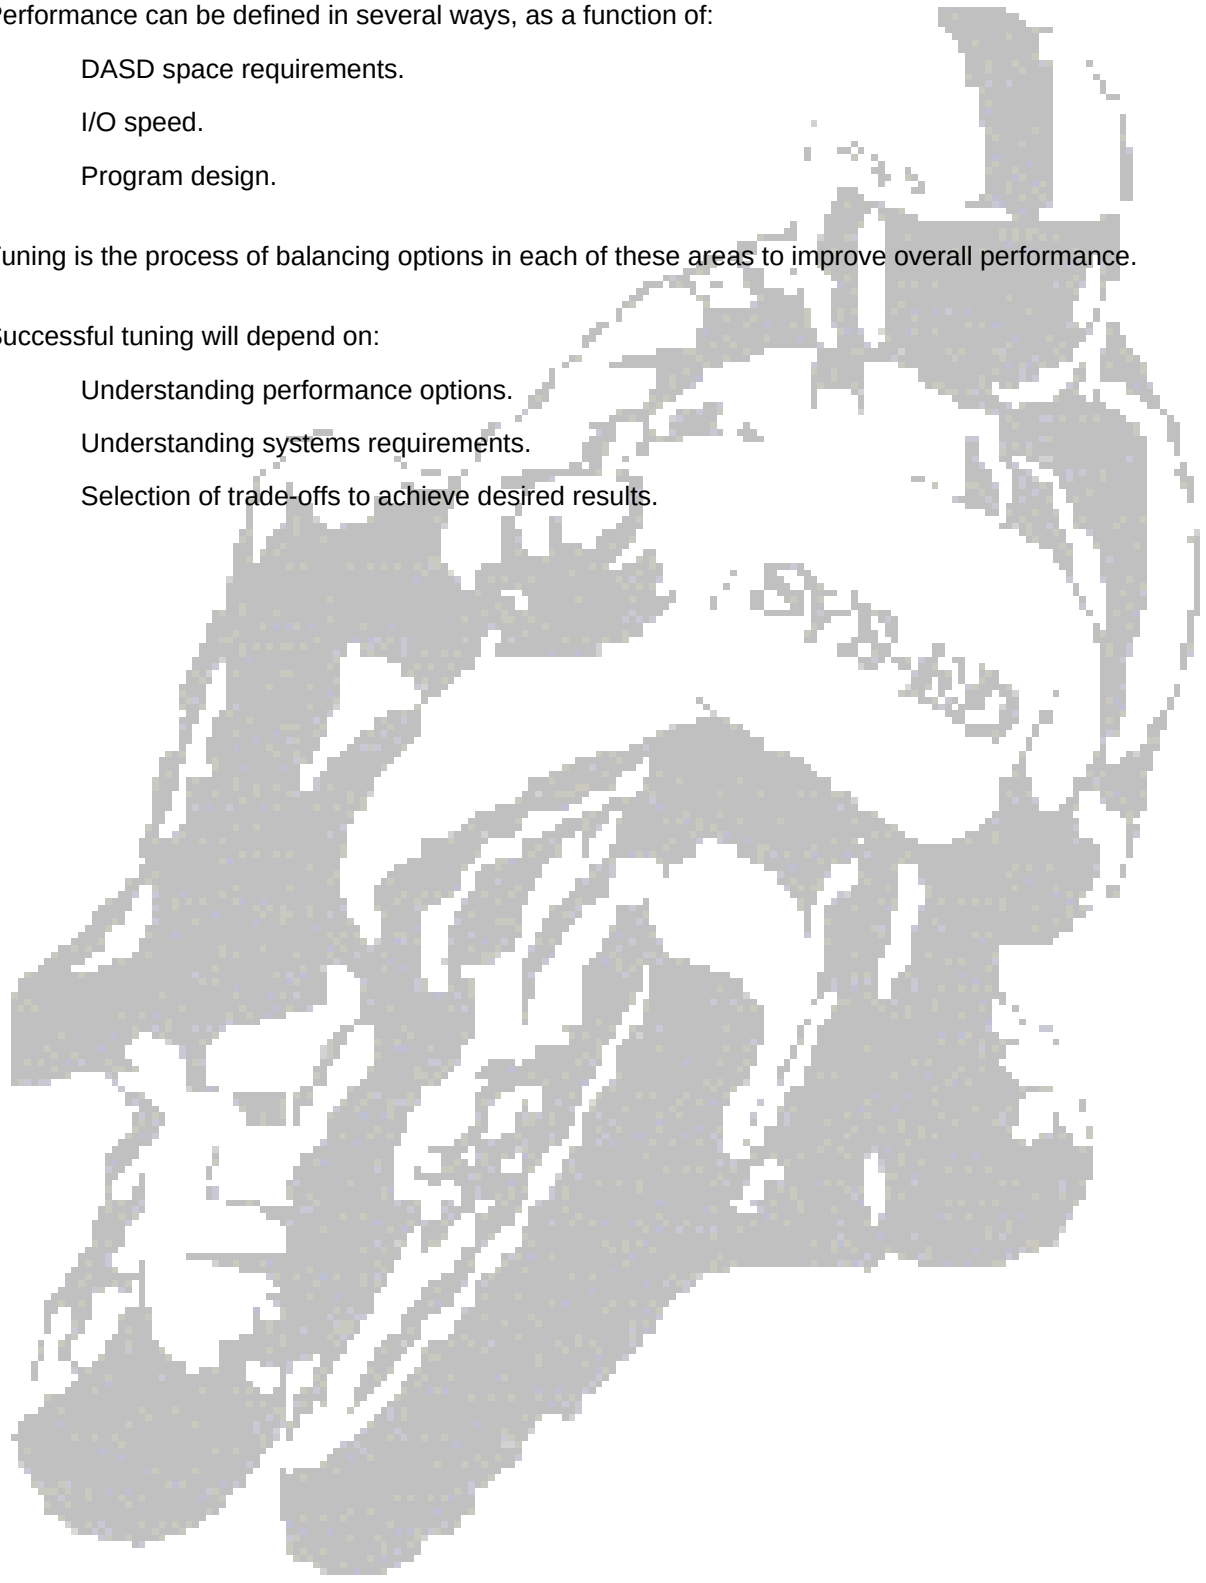

2 VSAM: History

2.1 Single Set of Access Method Programs

- Addition of new hardware devices required modifications to all old access method software.
- Plans for Fixed Block Architecture devices would require a new access method.
- Multiple access methods required multiple sets of utility programs.

Solution:

Access Method Services (IDCAMS) as overall utility for VSAM and nonVSAM files.

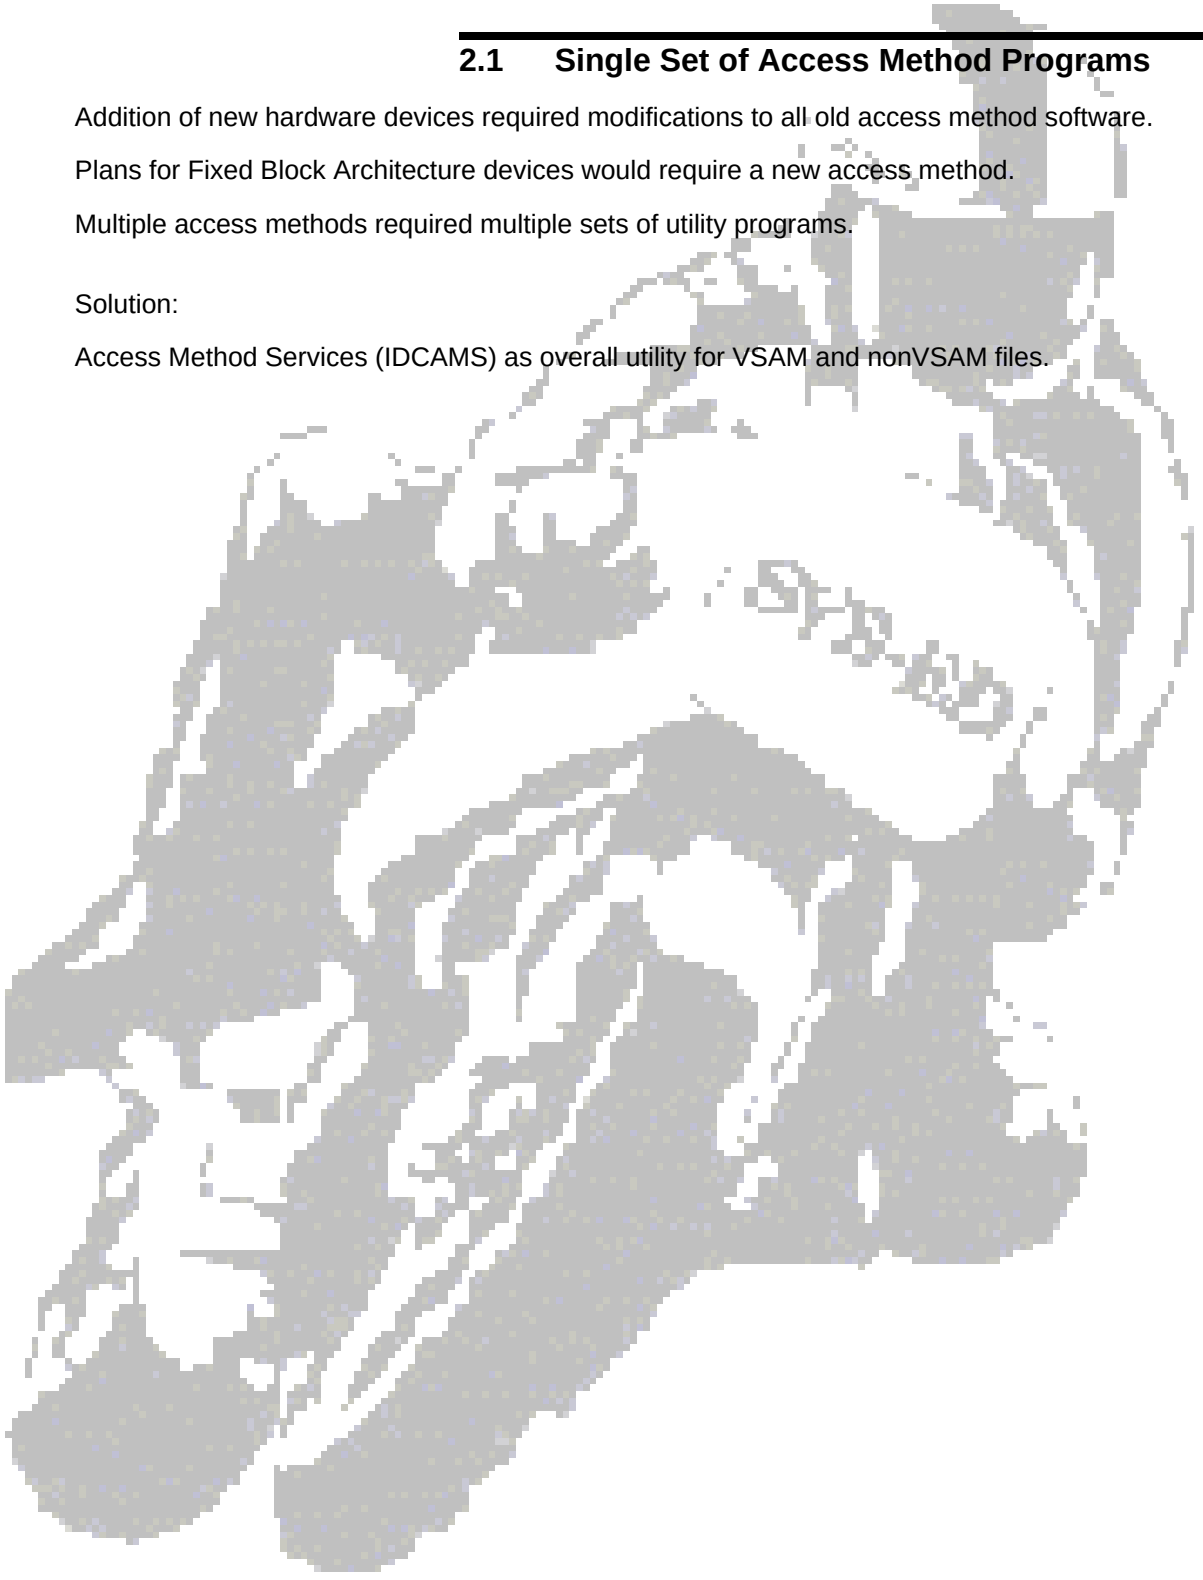

3 Integrated Catalogs and Space Management Consolidation

Consolidation of space management facilities was required in order to reduce maintenance effort and increase compatibility of the various component features.

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

The VSAM catalog was designed to serve as the 'dictionary' of the entire DASD environment. In addition to this centralization, the catalog entry structure allowed:

- More information than VTOC records.
- Relation of physical datasets to logical dataset structures like KSDS clusters and paths.
- Multi-volume dataset handling.

Solution:	VSAM catalog
Second Solution:	Enhanced VSAM
Third Solution:	Data Facility/Extended Function

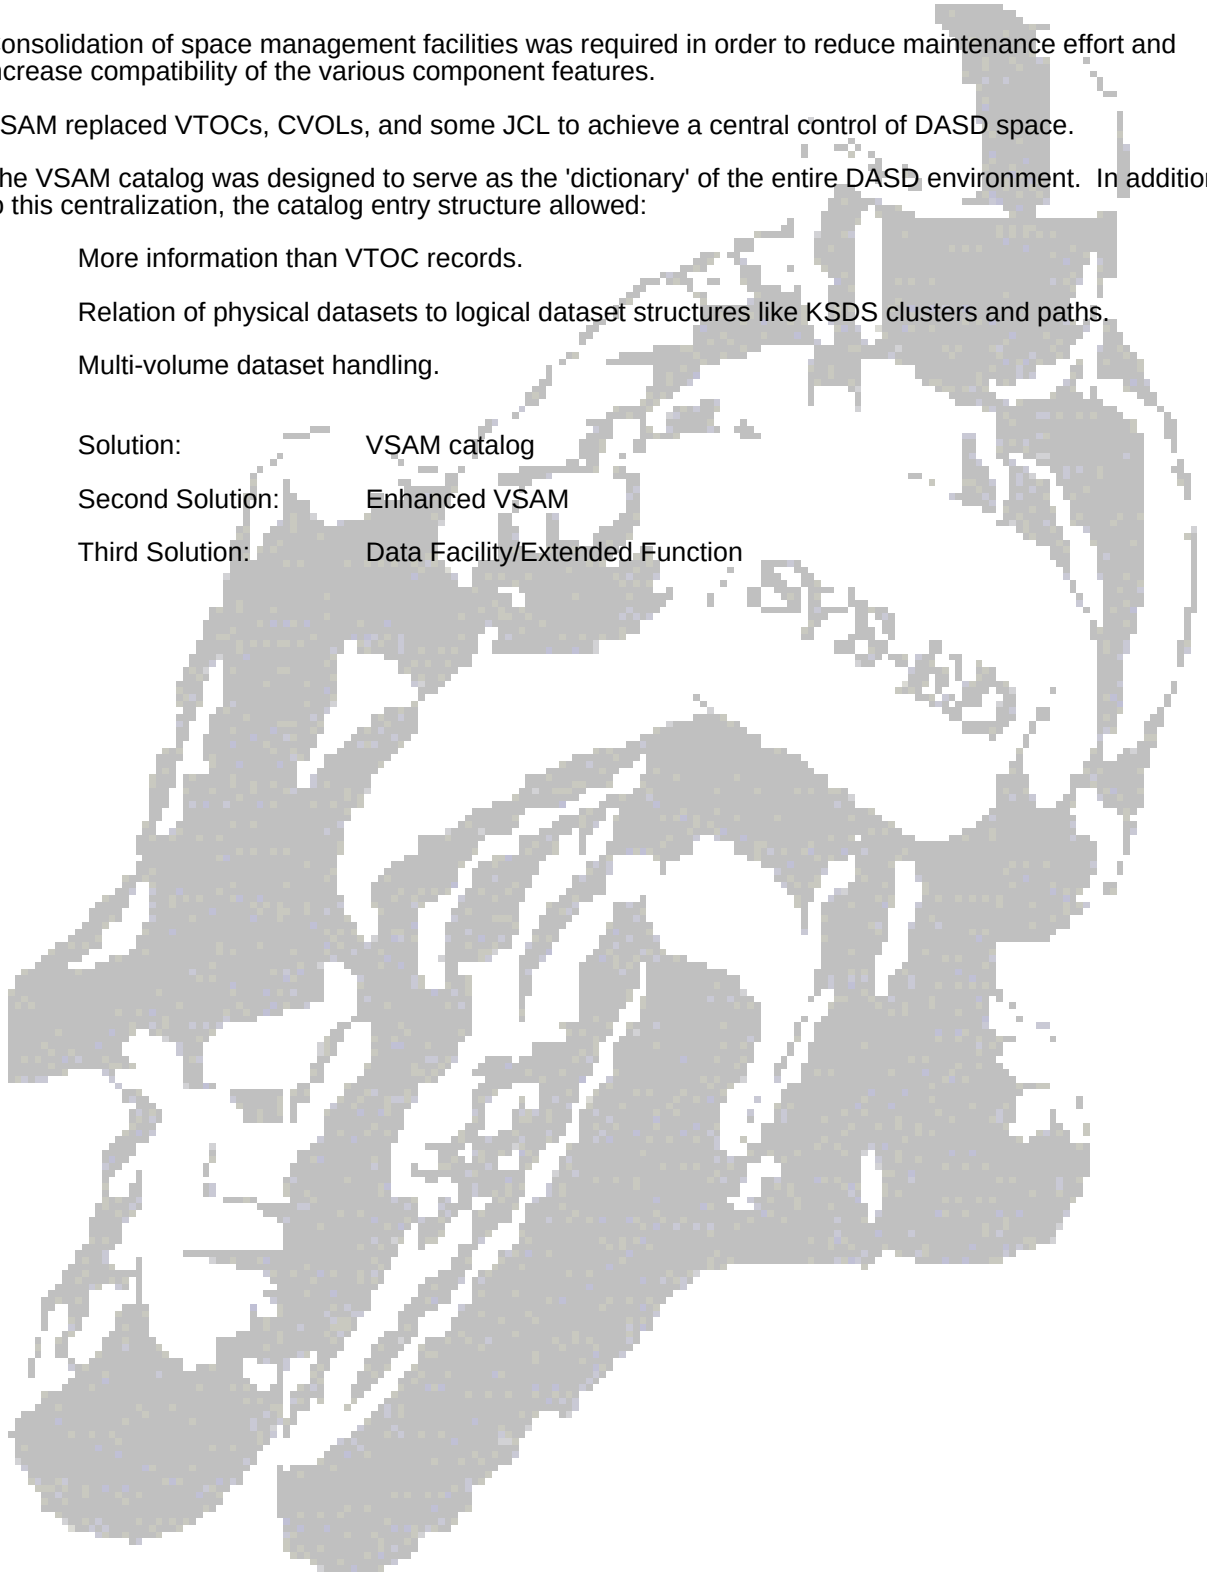

4 Physical Components

VSAM builds all files with the same physical elements:

- Control Interval: CI
- Control Area: CA

4.1 CI: Control Interval

The 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.

4.2 How Data Sits 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.
- Always start at the beginning of a CI.
- Are represented by 2 RDFs containing segment length and update number segment.
- The 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.

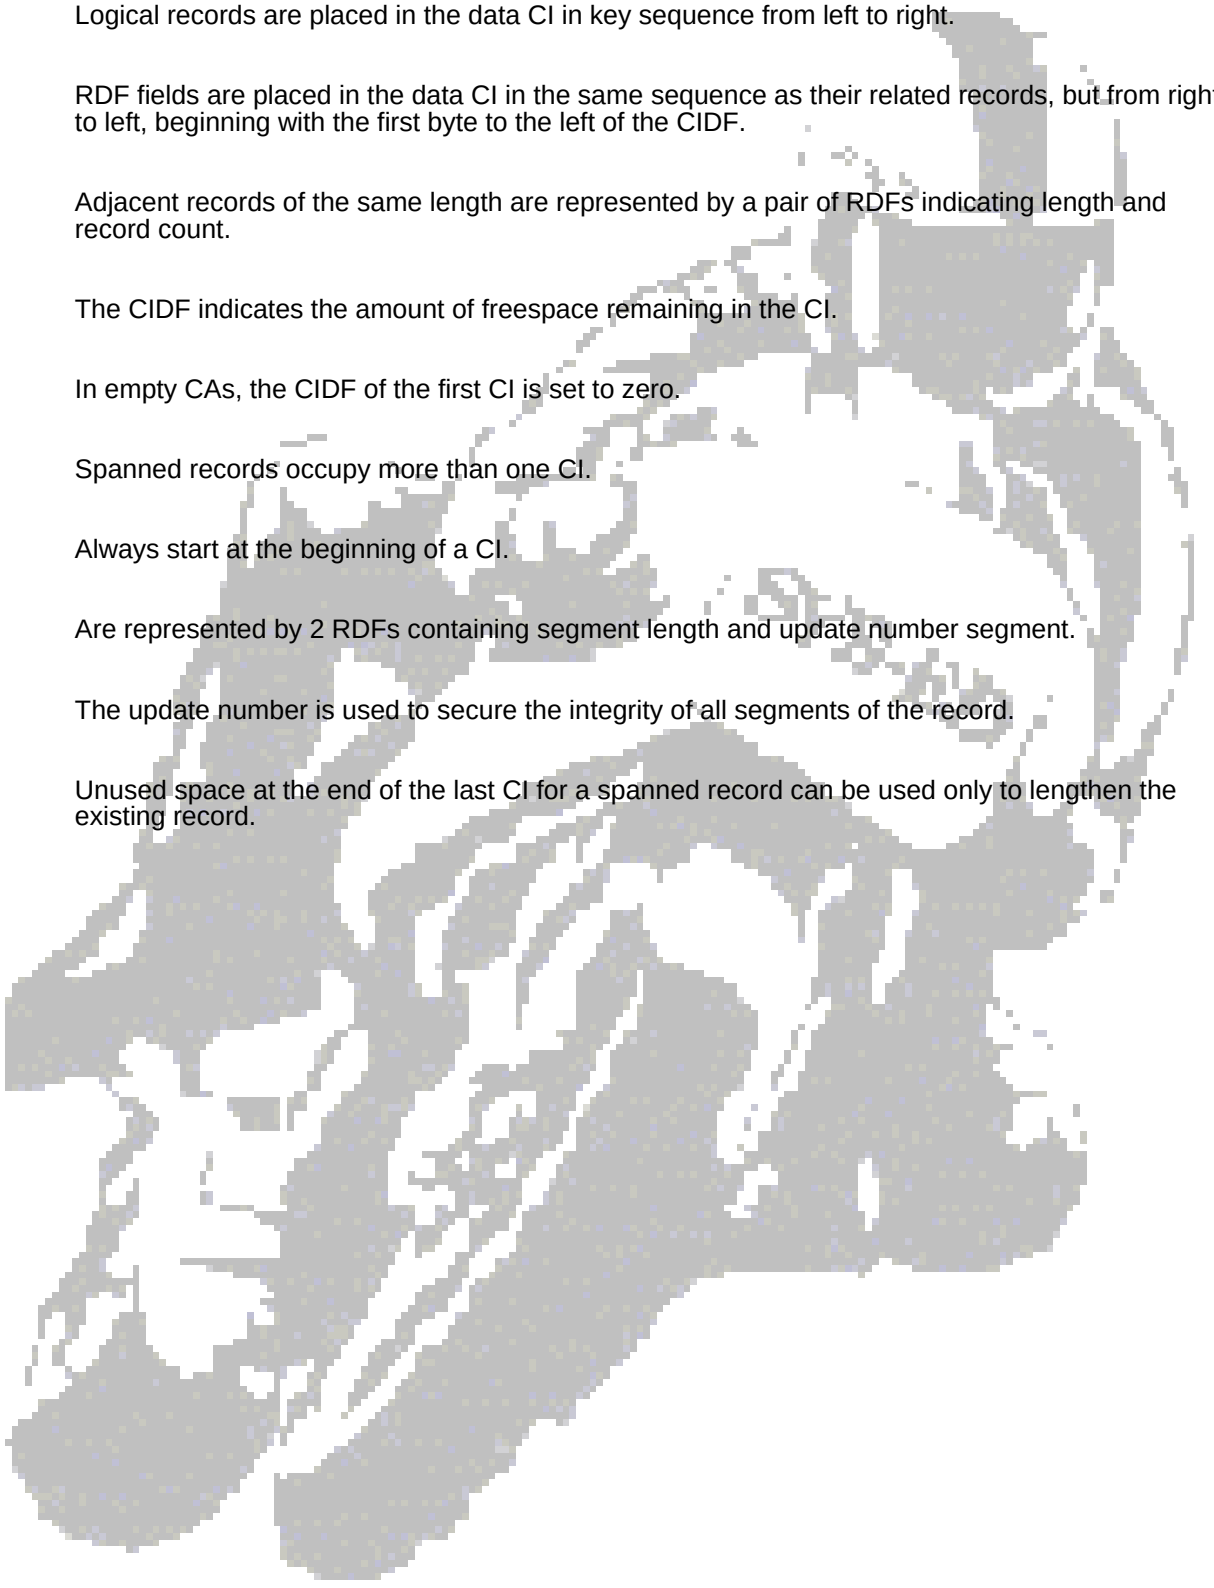

Key-Sequenced Datasets

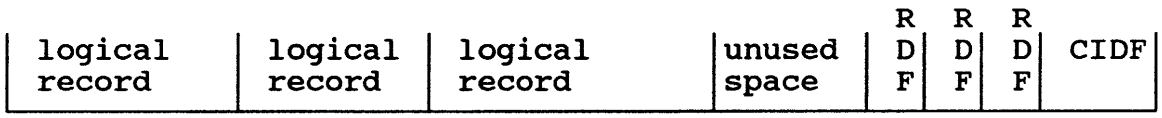

No two consecutive records with same length.

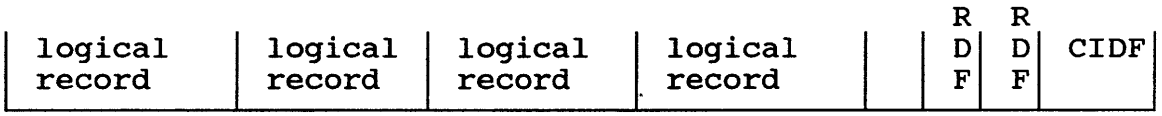

All records same length.

Unused space

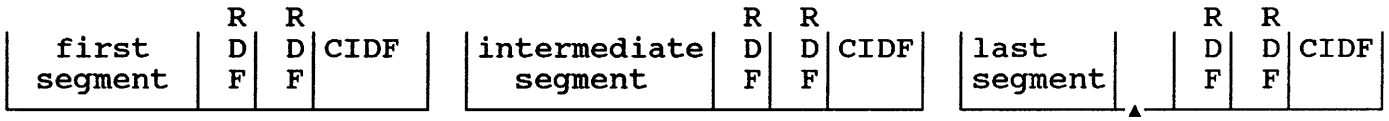

Spanned records

unused space

4.3 Control Interval Control Information

One control interval definition field (CIDF) per CI, containing the total length of the data (if any) in the CI, and the total length of the unused space (if any).

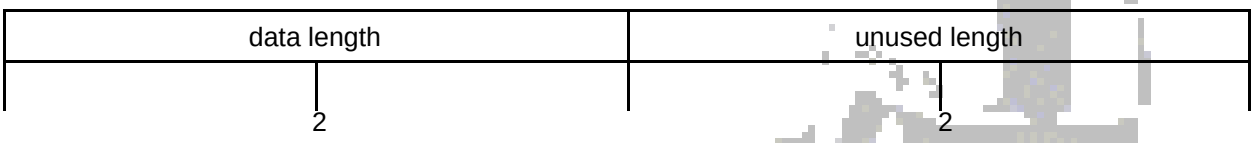

One or more record definition fields (RDF) per CI, which describes a record, record slot, or spanned record segment. They are placed in the CI from right to left, so that the right-most RDF describes the left-most record.

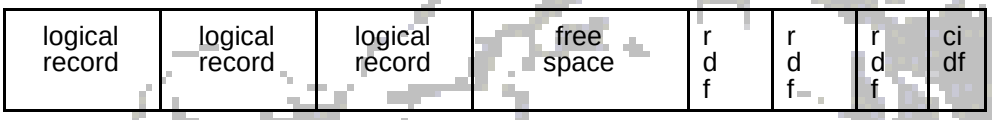

When two or more records of the same length are adjacent in the CI, one RDF will describe the record, and the next will indicate how many of these records are in sequence in the CI.

An RDF consists of a control field and a binary number. The meaning of the number is indicated by the control field contents.

4.4 Record Definition Field

Control 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 gives the length of the record.
	1 - The binary number in this RDF gives 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.

4.5 Control Field

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

4.6 CA: Control Area

- Unit of file organization containing a fixed number of CIs for a VSAM file.
- Determining factor in index structure for KSDS.
- Always contiguous area of DASD storage.
- All CIs reside in CAs; a CA may contain CIs which contain data as well as empty Cis.
- Size is determined implicitly by space allocation:

  minimum - 1 track
  maximum - 1 cylinder

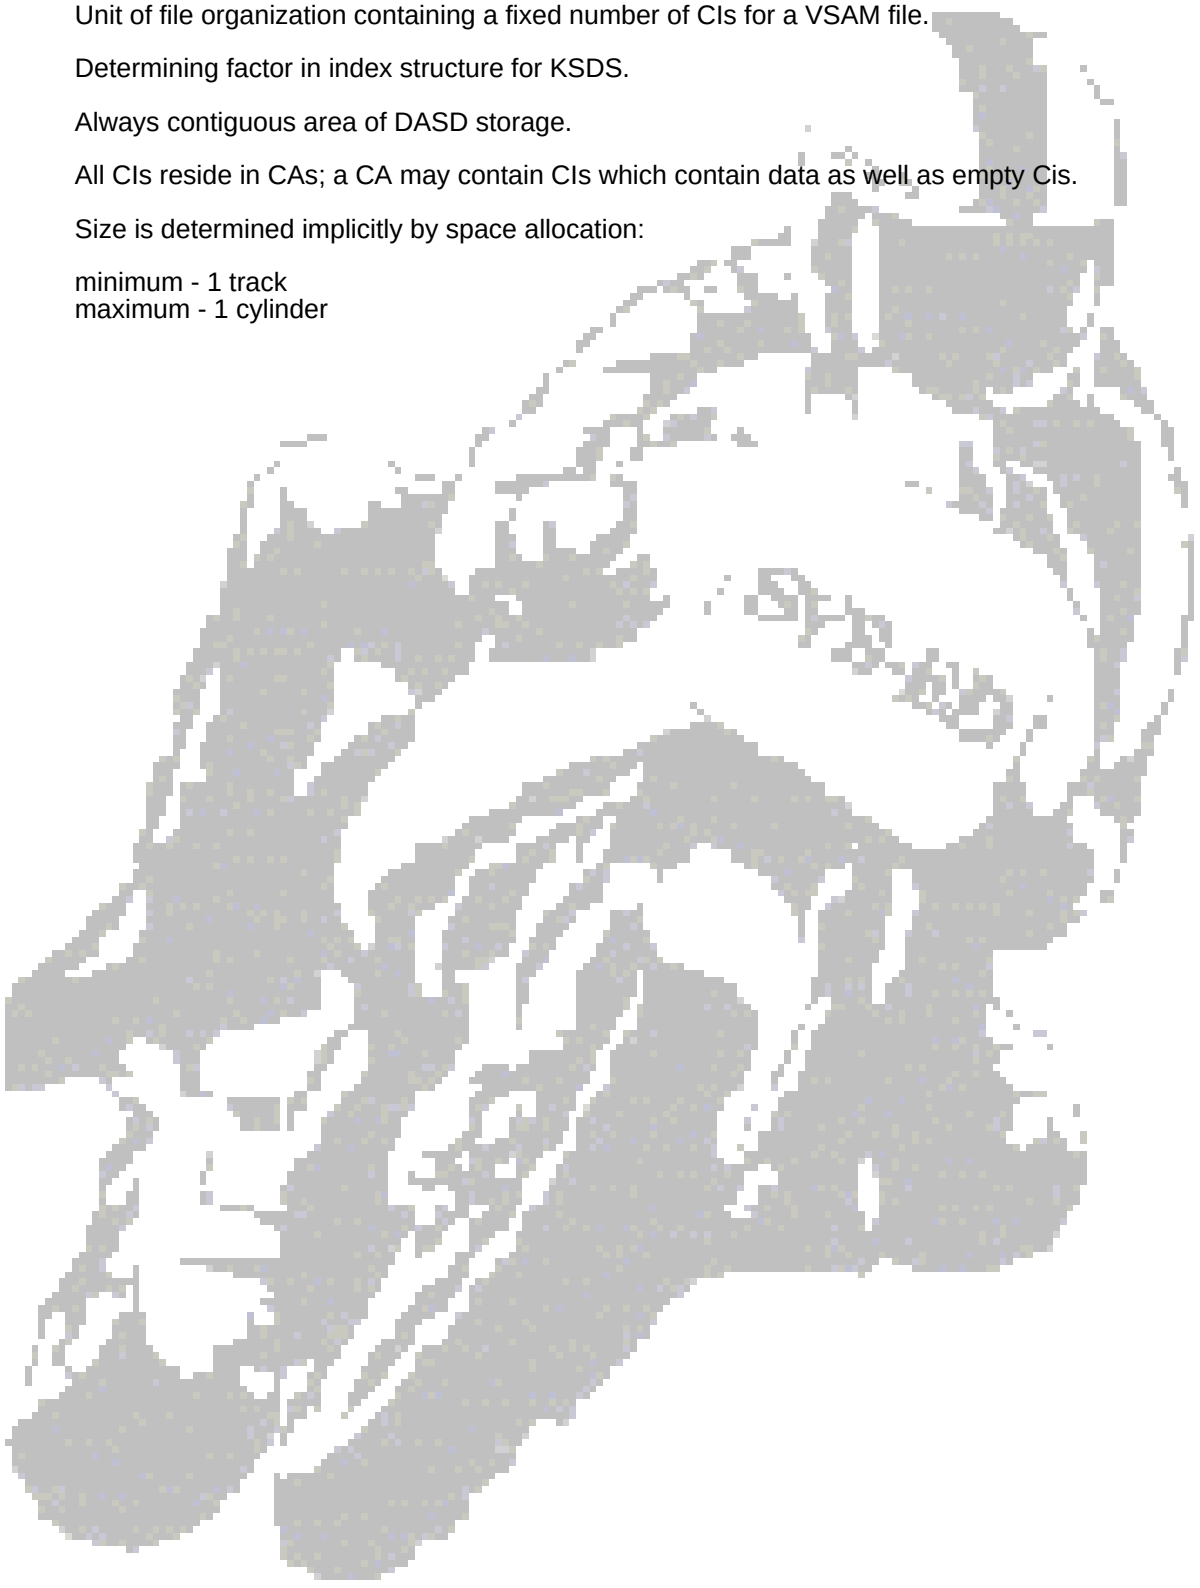

5 Review of VSAM Structures

5.1 Files

VSAM manages three basic file structures, or clusters:

KSDS: Keyed Sequenced Data Set

ESDS: Entry Sequenced Data Set

RRDS: Relative Record Data Set

And an additional file organization:

AIX: Alternate Index

5.2 KSDS: Key Sequenced Data Set

A KSDS 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 RBA (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 added in place (key sequence) using distributed freespace.
- Records deleted are physically erased, and occupied space is reclaimed as reusable freespace.
- Spanned records are supported.

5.3 ESDS: Entry Sequenced Data Set

An ESDS 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 added at the end of the file.
- Individual records may not be deleted.
- Spanned records are supported.

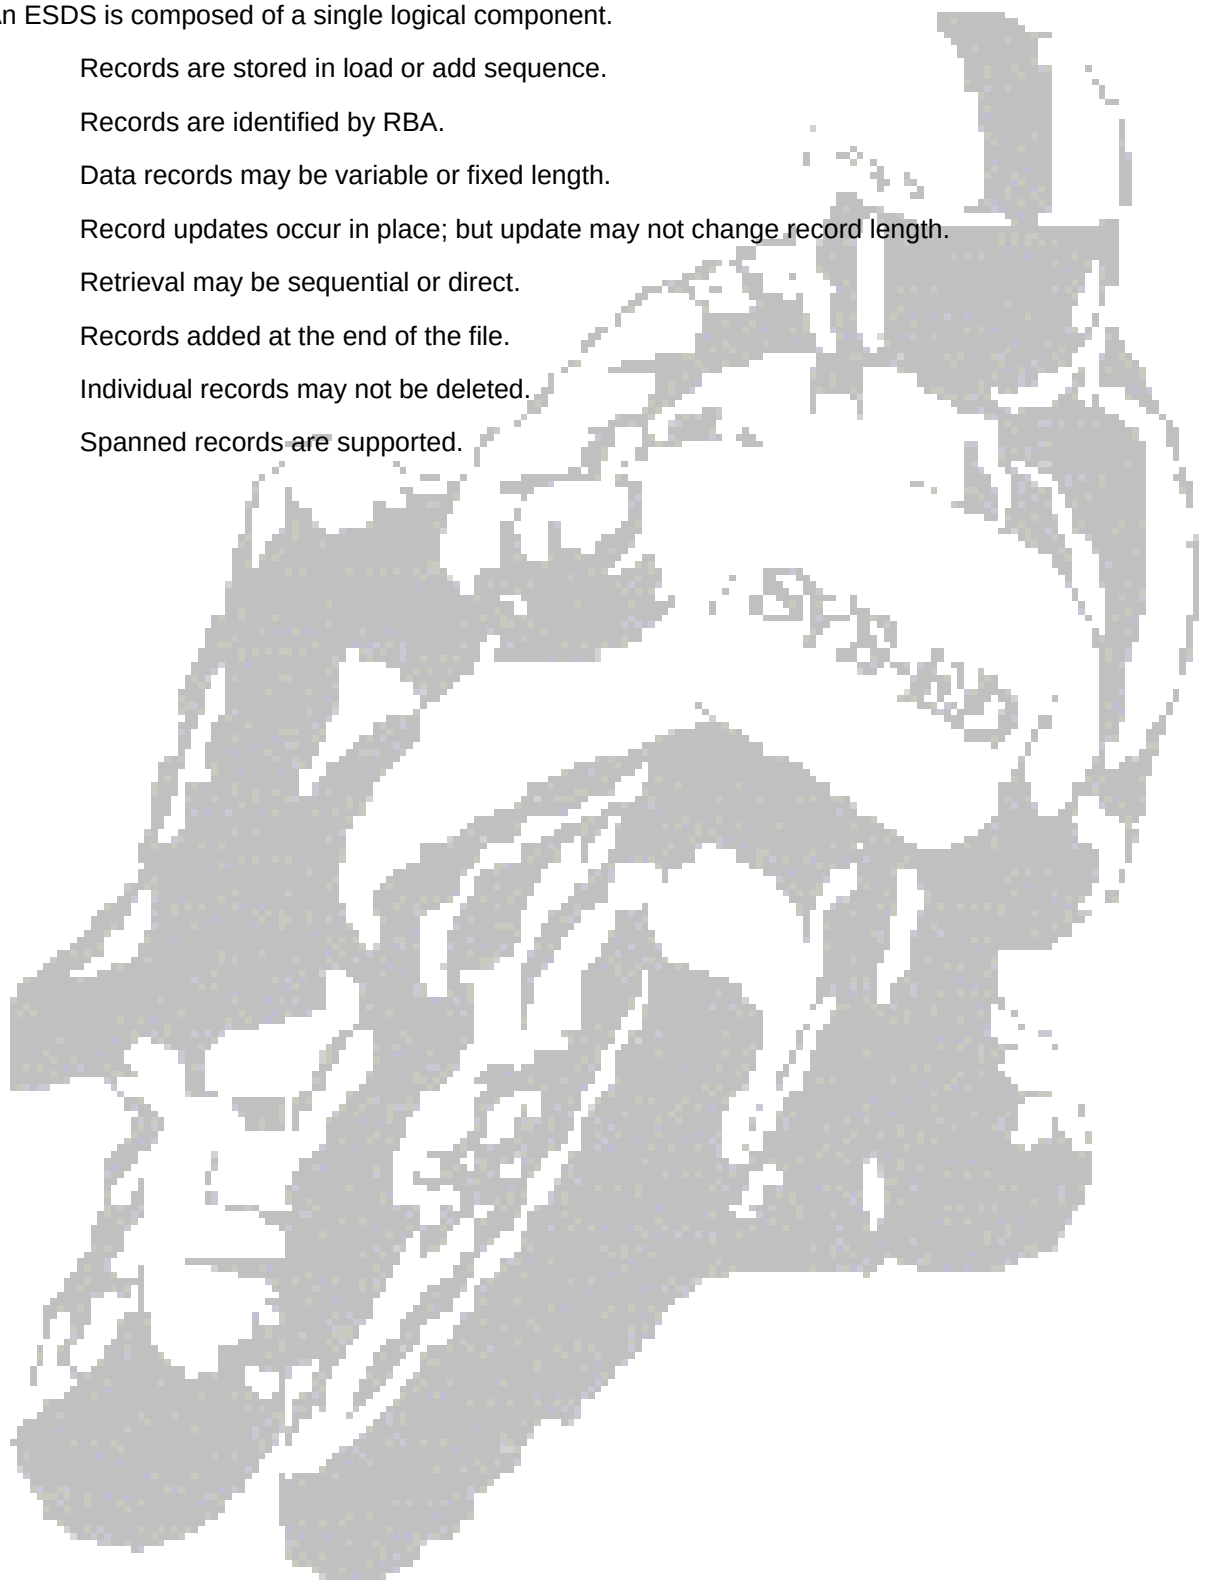

5.4 RRDS: Relative Record Data Set

An RRDS is composed of a single logical component, formatted into fixed length slots.

Records are identified by RRN (slot).

Data records may be fixed length or variable length.

Record updates occur in place; but update may not change record length, unless variable length.

Retrieval may be sequential or direct.

Records added in empty slots.

Record deletion marks a slot as empty.

Spanned records are NOT supported.

In concept:

- the KSDS replaced ISAM.
- the ESDS replaced QSAM.
- the RRDS replaces BDAM.

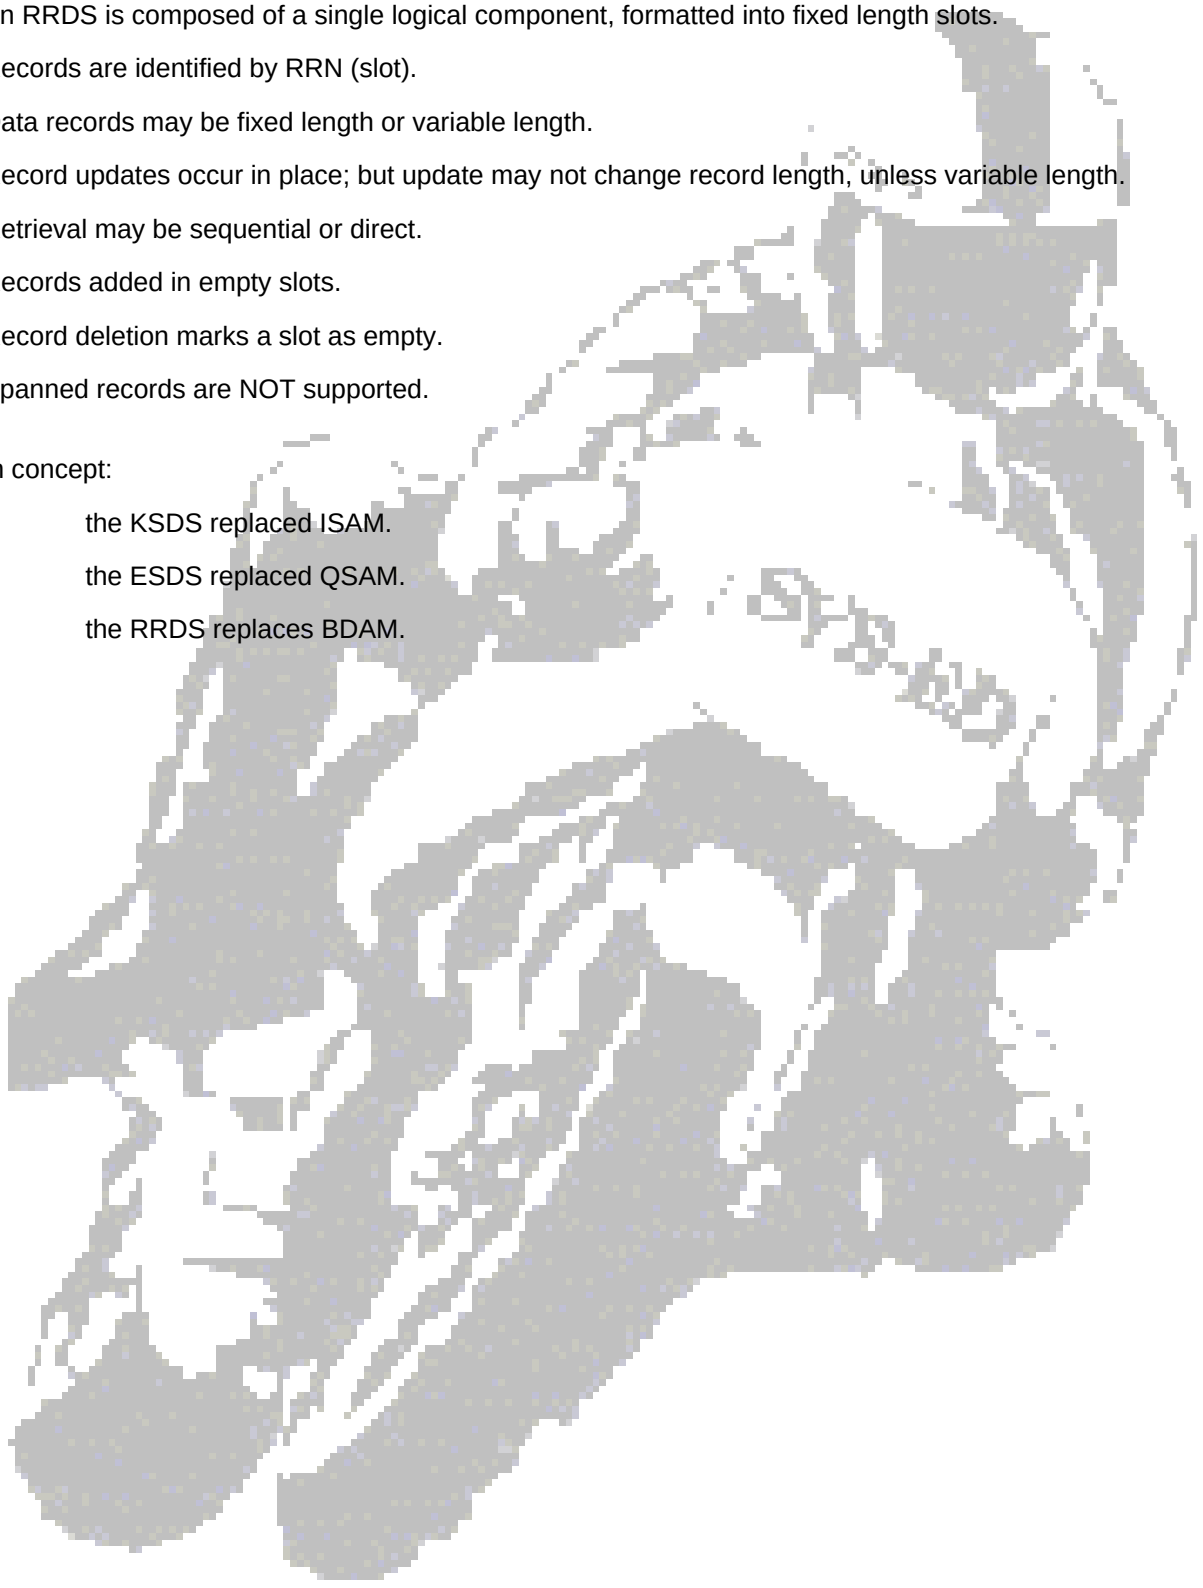

KSDS

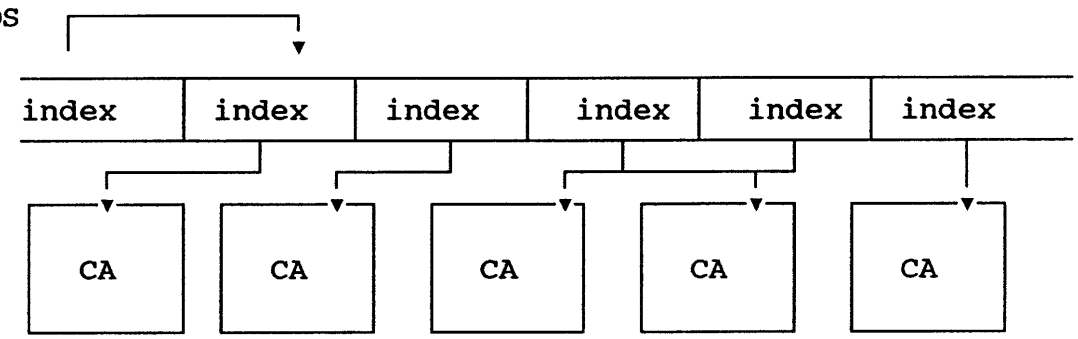

CI

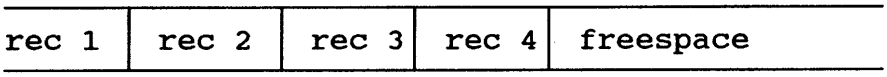

ESDS and RRDS

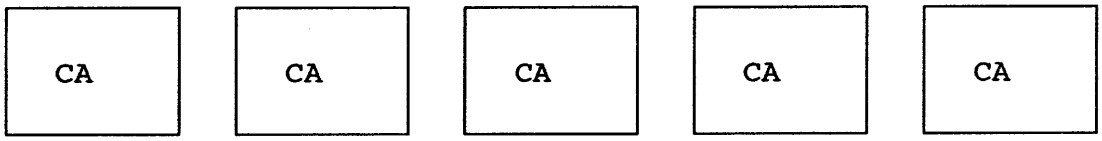

ESDS

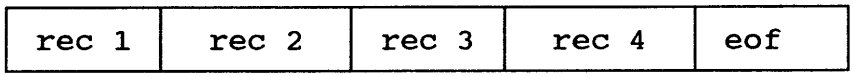

RRDS

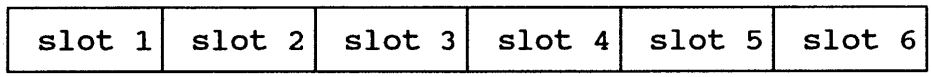

5.5 AIX: Alternate Index

An AIX is structured as a KSDS.

- Is an index to data in another base file, using an alternate key.
- The base file may be a KSDS or an ESDS.
- Is connected to its base cluster by a PATH definition in the VSAM catalog.
- Alternate keys may be unique or non-unique.
- Record updates can occur synchronously with the base file or not.
- The file may be processed independently, or with its base through a PATH.
- A base cluster may have multiple alternate indexes.

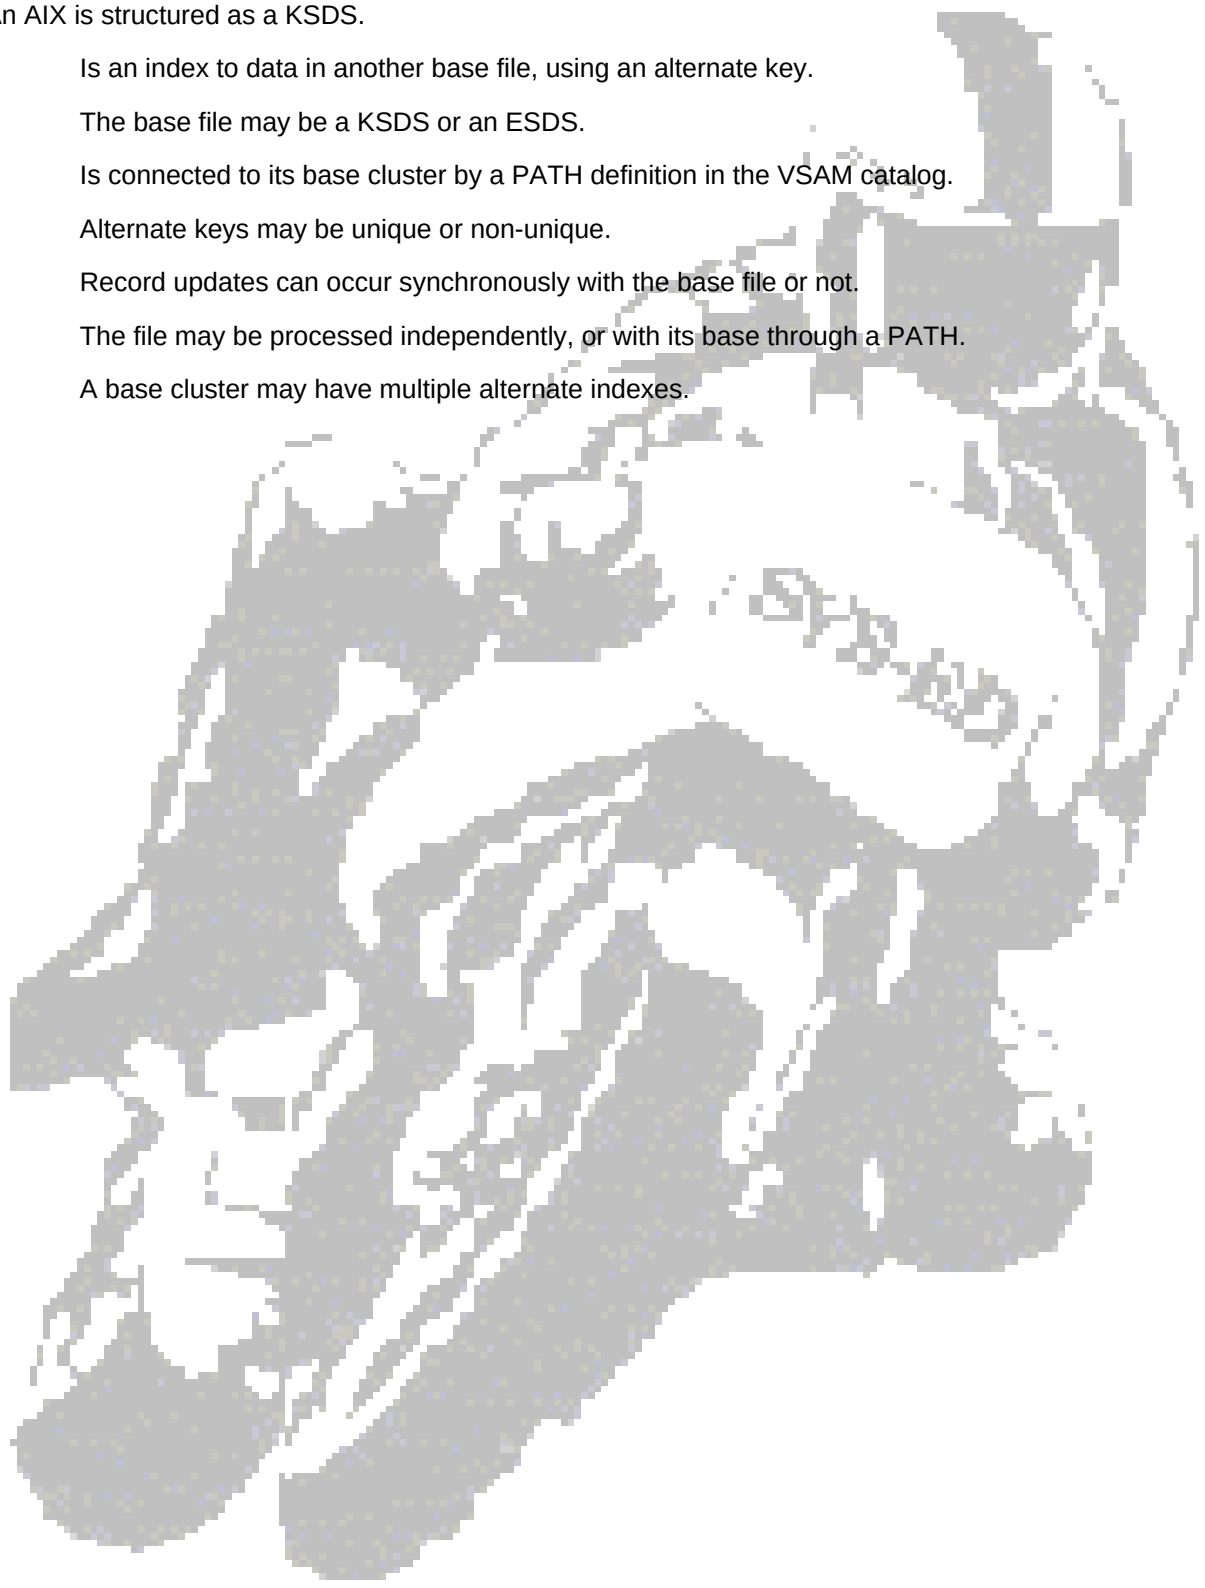

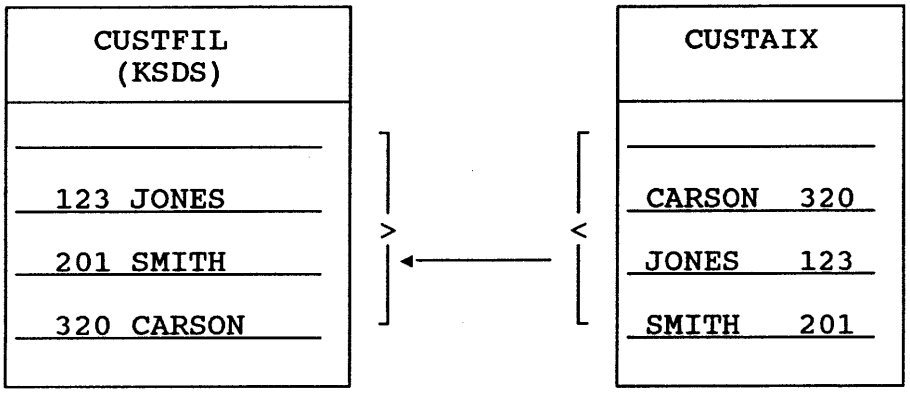

Base Cluster

Base Key = customer #

Alternate Index

Alternate key = customer name

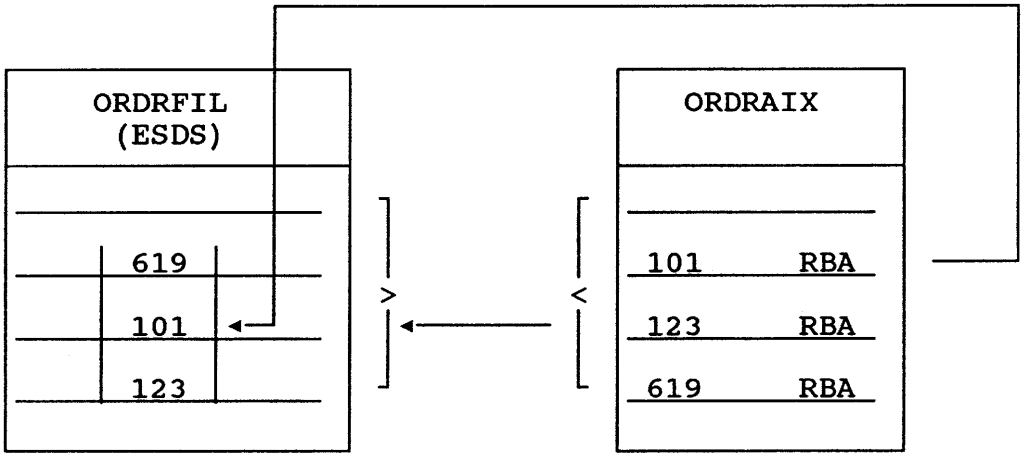

Base Cluster

No base key
Access is by RBA

Alternate Index

Alternate key = customer #

6 VSAM Type 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

7 Catalogs

- There must be one MASTER catalog per system.
- For the OS/390 or z/OS operating systems, the System Catalog is the VSAM Master Catalog.
- One or more user catalogs may exist, but need not exist.

7.1 User Catalog

- Reduces the size of and contention for the Master Catalog.
- Isolates potential damage problems.

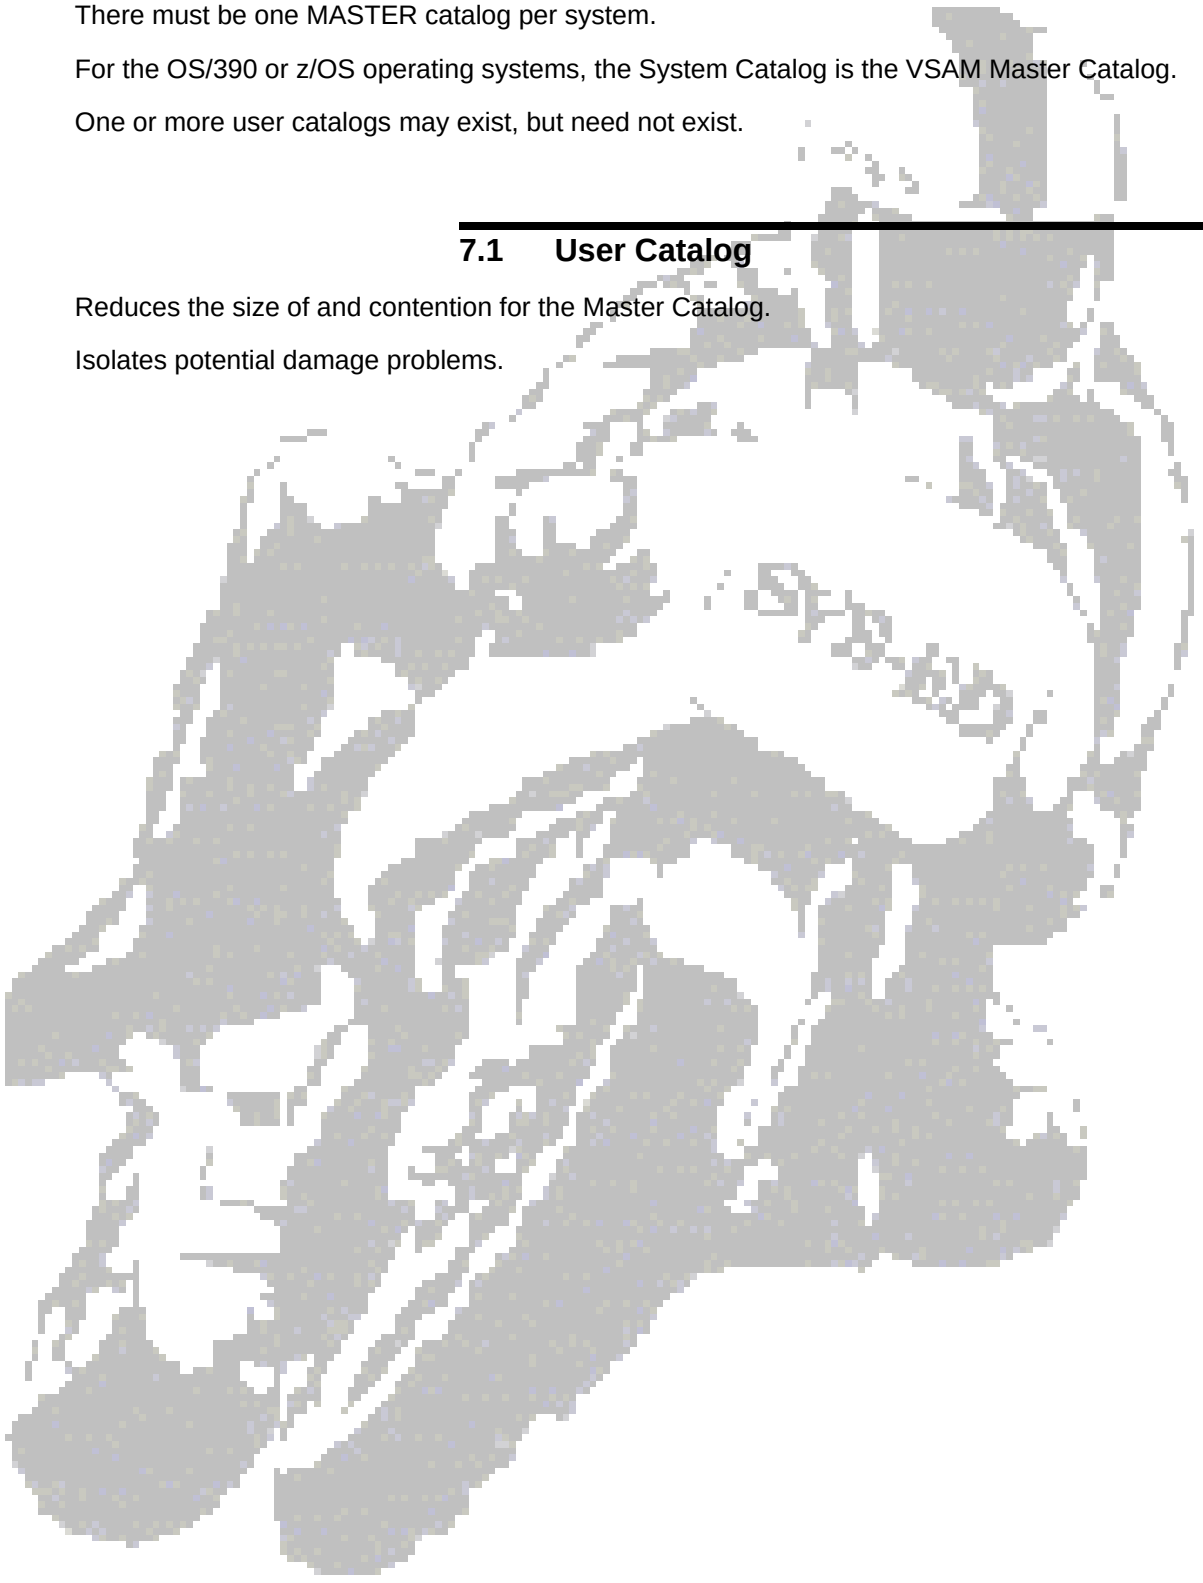

7.2 ICF Catalogs

ICF catalogs are composed of two elements:

BCS - Basic Catalog Structure

One BCS exists for each ICF catalog. It is a standard KSDS using the defined dataset name as the cluster key.

Each record in the BCS contains the history, protection and association information for the dataset, and volume information.

VVDS - VSAM Volume Data Set

One VVDS exists for each VSAM volume. It is a standard VSAM ESDS, and contains records for VSAM datasets on that DASD pack.

Each record contains attribute, statistical, allocation and volume information groups.

- A VVDS is not 'owned' by the BCS; as many as 36 separate BCS files can point to the same VVDS.
- As many as 36 BCS files can exist on the same pack.
- Since the VVDS always resides on the same pack as its datasets, recovery is simplified.

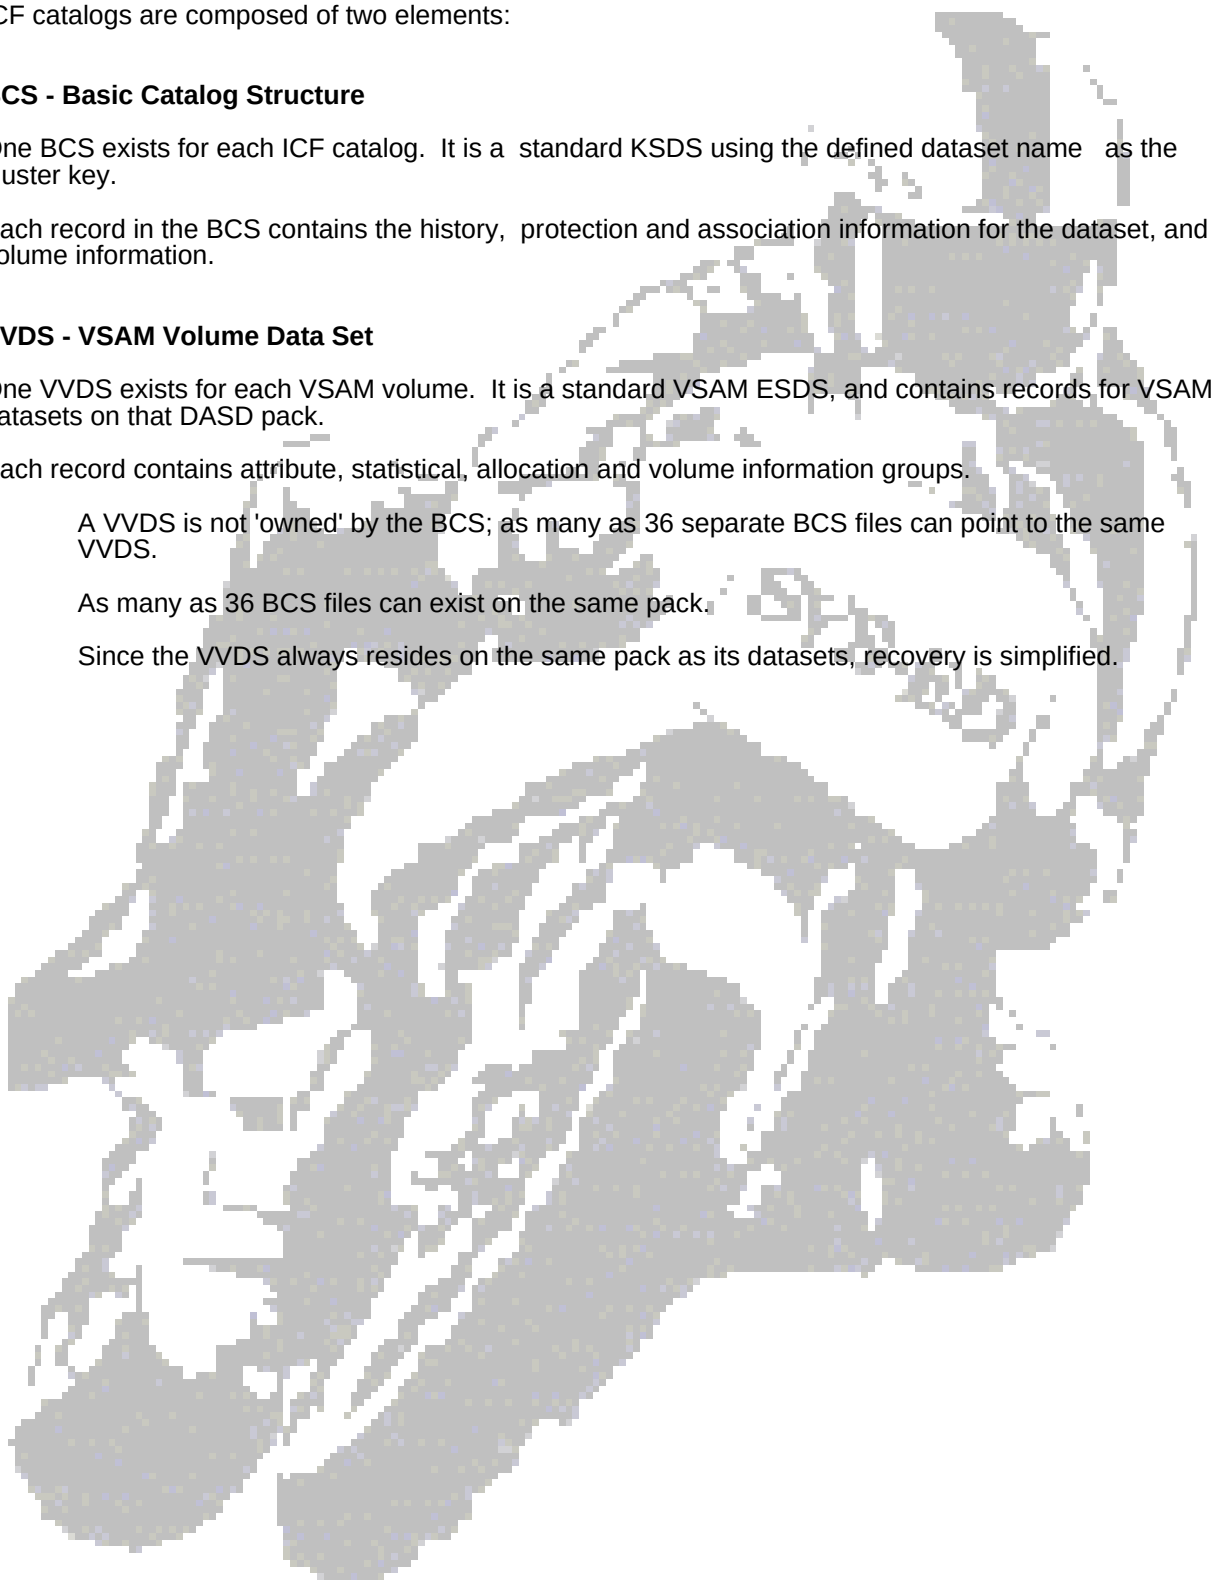

7.3 Index Levels and Physical I/O's

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
Readnext	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

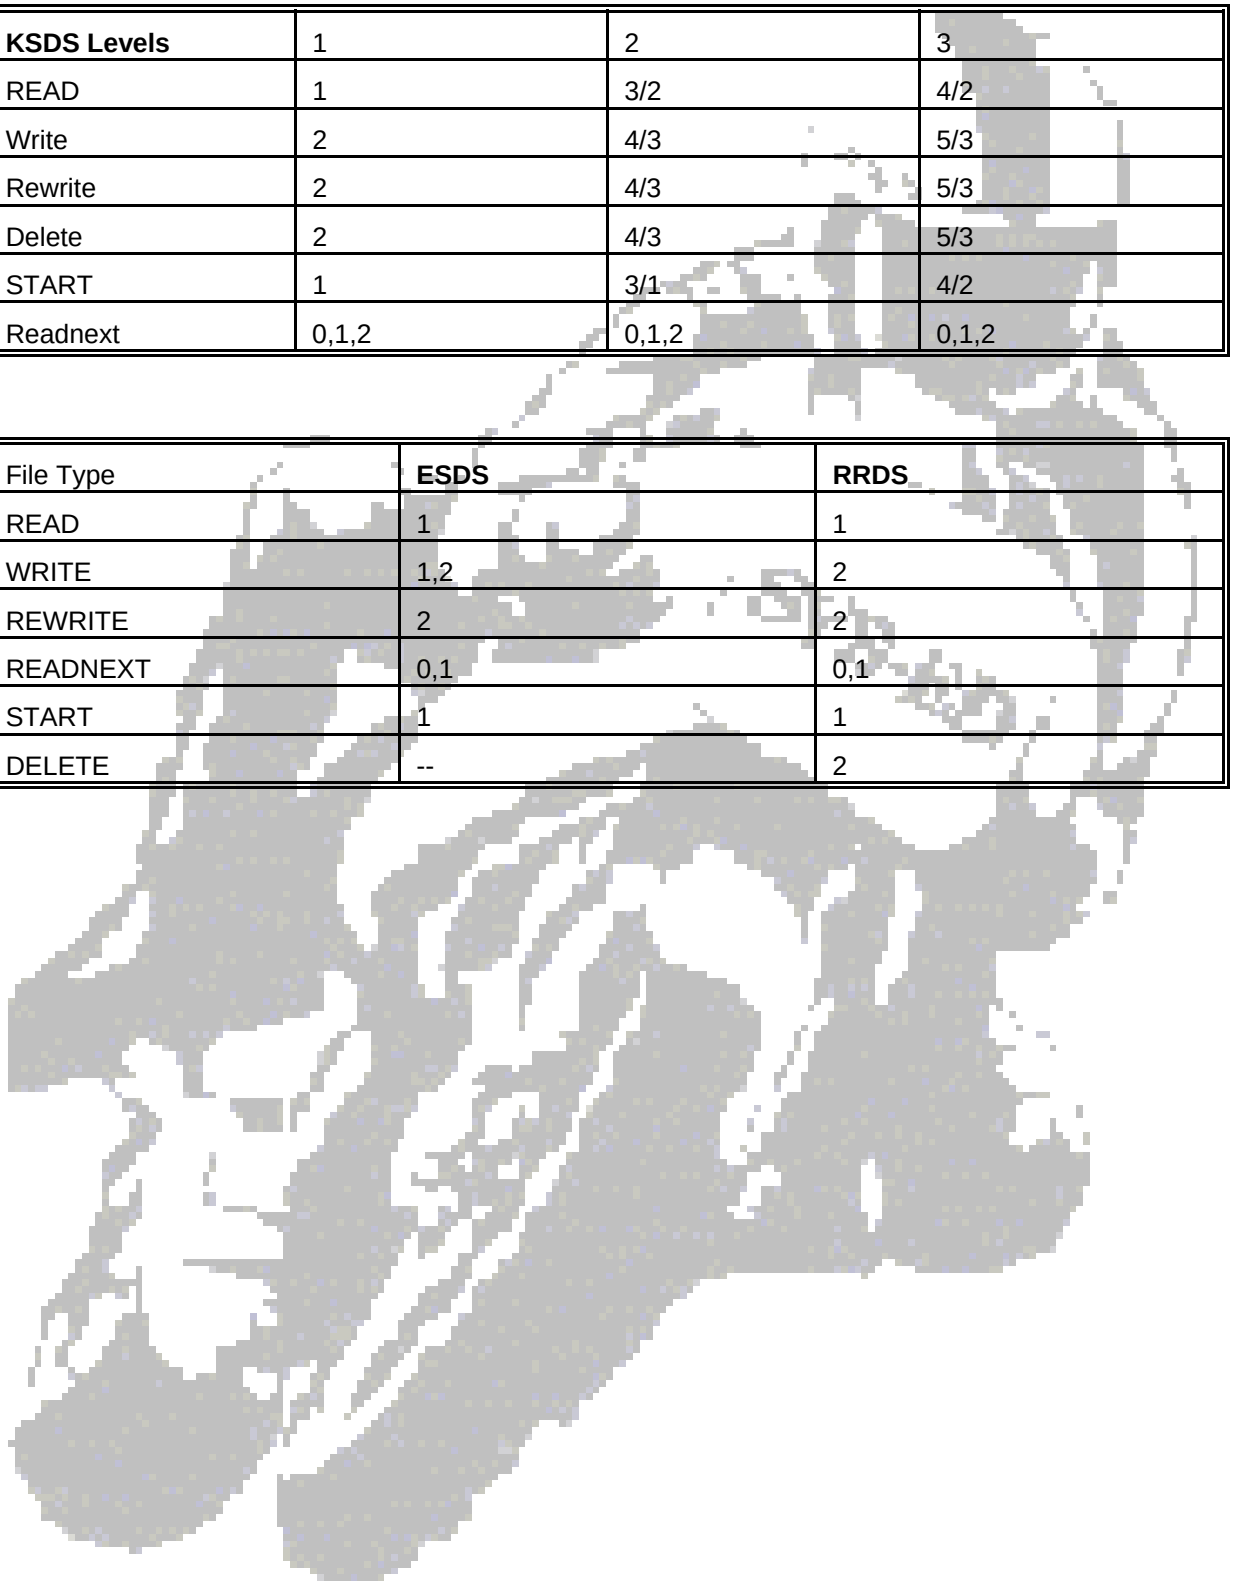

7.4 3380 and 3390 Characteristics

DESCRIPTION	3380	3390
Tracks/Cylinders	15	15
Available Bytes Per Track	47,476	56,664
Actual Bytes Per Track	47,968	
Cylinders/Volume	1770	1113/2226
Minimum Seek Time	3 msec	
Average Seek Time	16 msec	9.5/12.5 ms
Maximum Seek Time	30 msec	18/23 ms
Rotational Delay	8.35 msec	7.1 ms
Transfer Rate (MB Per Second)	3.0	4.2
Total I/O Time	27.35	20.8/23.8
Data Overhead/Physical Rcd	492	

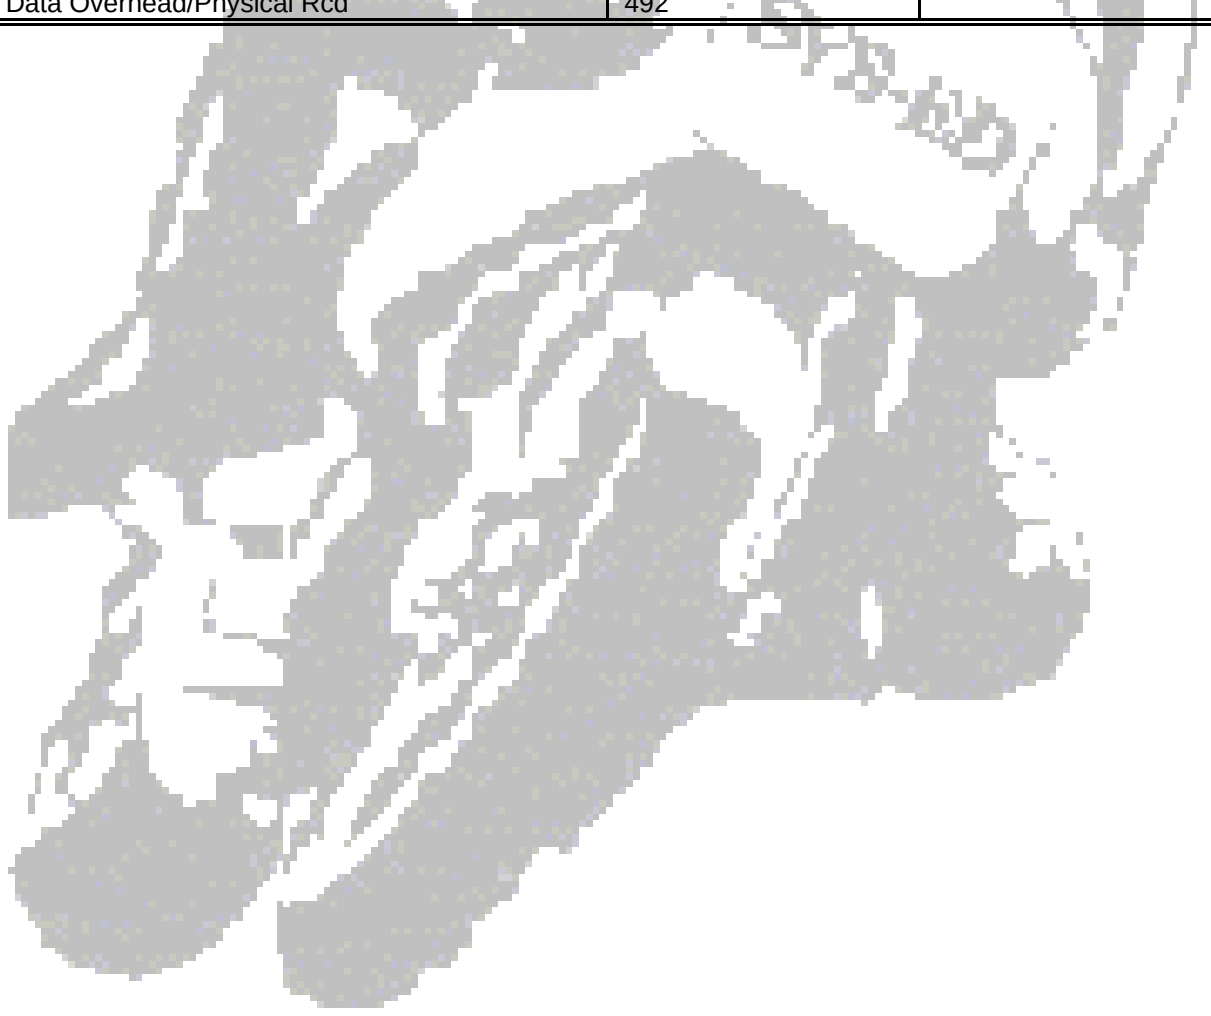

8 DASD Response Time

DASD response time is measured in four components:

IOS Queue time	Is the delay experienced within the 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 experiences while the device positions itself to the proper cylinder and record. These delays can be elongated by RPS misses (that is, extra revolution) that occur due to control unit and channel contention.
Connect time	Is the 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.