PRACTICE 450-1021-103
Issued: 90 03 30
Vintage: NSR28-30 01
Standard

NETWORK OPERATIONS SYSTEM

STATION DETAIL SERVER

DNC*-50: FEATURES DESCRIPTION

CONTENTS **PAGE**

Contents continued ----------

* DNC is a trademark of Northern Telecom

CONTENTS (Continued) **PAGE**

Figures

CONTENTS (Continued) **PAGE**

Figures

1. INTRODUCTION

1.01 The station detail server (SDS), a member of the business network management (BNM) family of applications, uses a Dynamic Network Control DNC-50 system to collect station message detail records (SMDR) data concerning Meridian Digital Centrex (MDC) networks.

1.02 This Northern Telecom practice (NTP) describes the features available for the SDS product, using the DNC-50 system operating with the software release NSR28.

Structure Of This Publication

1.03 This publication is divided into the following chapters:

Introduction: Is an introduction to this publication, it introduces SDS, and identifies the applicability of this publication.

DNC-50 Configuration: Describes the DNC-50.

SDS Features: Is an overview of the features provided by the SDS.

SDS User Interface: Describes how the user can access the SDS features.

Operational Measurements: Is an introduction to operational measurements recorded in the DNC-50.

Abbreviations: Is a list of abbreviations used in this publication.

CHANGE HISTORY

1.04 This section lists the important changes that affect this publication. They are arranged by network software releases (NSR) in a descending order starting with the current release.

NSR28

1.05 The changes introduced by NSR28 are:

- a new feature on Disk Utilization and Monitor
- an enhancement to SMDR spooling to allow multiple PRU data spooling, providing 32 simultaneous spooling ports
- an enhancement to allow more than one DNC to collect data from a single DMS node

NSR27

1.06 The changes introduced by NSR27 are:

- an increase in the SMDR collection capacity to a maximum of 4M records per day
- storage of over 11 million records

NSR26

1.07 The change introduced by NSR26 is:

- enhancement to the existing data spooling feature to support multi-node spooling, this allows a customer 's SMDR records to be spooled serially from multiple DMS nodes using a single spooling schedule.

2. DNC-50 CONFIGURATION

2.01 Meridian Digital Centrex (MDC) customers require information about the calls made on their physical or virtual private networks. Such information is essential for internal accounting.The DMS* nodes (a member of the DMS-100 family of digital switches) accumulate relevant data and stores it on disk as various MDC features are used. The DMS nodes collect this data then process it to produce station message detail records (SMDR) for its MDC customers.

2.02 With a station detail server, the telco uses a DNC-50 to gather SMDR for each customer from the DMS nodes used by the customer (either 1 or all). The system then partitions the SMDR records, and spools them to customers at scheduled intervals or on demand. The logical arrangement of the system is shown in Fig. 2-1.

2.03 Station detail servers have several advantages over standard piecemeal collection and custom processing of data. These advantages include:

- central, near real-time collection and processing of data collected from the DMS node
- reduced tape and paper handling by telcos and their customers

2.04 The major component of a station detail server is a Dynamic Network Control (DNC) System located in the telco's office or computer equipment room. General information on the DNC systems is contained in 450-1011-100. A station detail server system uses a type of DNC known as the DNC-50. This is a multi-tasking communications device that collects data as it is generated by the DMS node. The DNC-50 identifies the data collected and partitions it by customer. See 450-1021-153 and 450-1021-353 for information about configuring a DNC-50 for the station detail server application.

*DMS is a trademark of Northern Telecom.

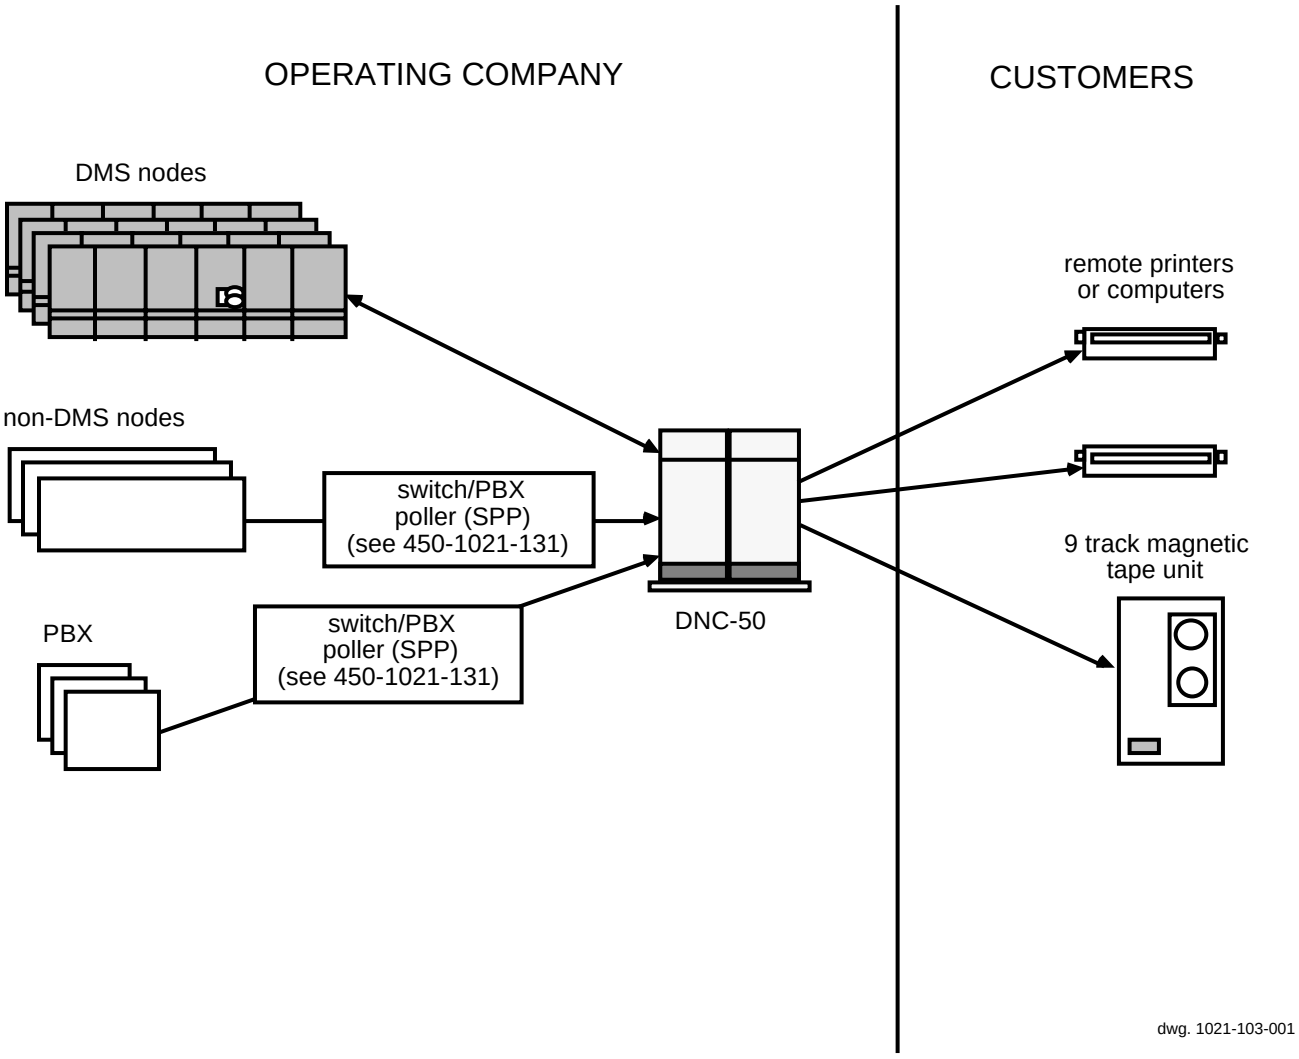

Fig. 2-1
The Logical Arrangement of a Station Detail Server

3. SDS FEATURES

SMDR DATA COLLECTION

3.01 The main function of the station detail server application is to collect station message detail record (SMDR) data from one or more switch nodes or PBXs, then format the data and transmit it to the appropriate customers through dedicated or dial-up spooling ports. If a tape unit is included with the DNC-50 system, the station detail server can also generate tapes of SMDR data that can be sent to customers for processing.

3.02 SMDR is the Meridian Digital Centrex feature that generates and collects information on telephone calls made by end users. SMDR records help Meridian Digital Centrex customers to control their communications costs because the individuals or departments that incur the costs are clearly identified. SMDR records also provide customers with a basis for allocating costs fairly.

3.03 A DMS node can generate SMDR records for calls within the node, including the following types of network calls (internode and local node) and off-network calls:

- local (dial ‘9’)
- 411
- 555-1212
- dial ‘0’
- DDD
- IDDD
- Outwats
- Inwats
- 911

3.04 In general, calls made from the following facilities in a Meridian Digital Centex system generate SMDR records:

- stations co-located with nodes
- attendant consoles
- PBX/MDC access trunks
- internodal tie trunks
- virtual facility groups

3.05 The DNC-50 normally collects SMDR data from switch nodes in continuous mode, as the data is generated. When required, a demand transfer can also be used to retrieve files.

MASKING

3.06 If the mask feature is included with the station detail server, specified stations can have the numbers they call masked, that is, hidden, by a preselected mask character before spooling. Call detail information remains intact, but the system substitutes mask characters for the last four digits in the called number. The number called by the originating station cannot be identified.

3.07 To operate the mask feature, the telephone operating company enters the telephone numbers of the stations for which calls are to be masked into a mask table. Any calls made from telephones with numbers listed in this table will have the last four digits of all called numbers masked in the resulting SMDR data.

3.08 The mask character is preselected to be the letter X, but can be changed by Northern Telecom personnel to suit other preferences. Mask table access for updating is confined to the main terminal of the DNC-50 system.

DATA SPOOLING

3.09 A DNC-50 system running the station detail server application can spool current SMDR records to customers' computers or printers. Certain records in the spooled data are masked when the mask feature is used.

3.10 The system automatically converts the data to be spooled to ASCII format. The carriage return and line feed characters are inserted after each record in the data. This allows it to be printed without further modification. A header, separated by rows of asterisks, identifies the customer, location, and type of data. New headers in the same transmission indicate data from different DMS nodes (when spooling for all nodes). A trailer, also separated by asterisks, indicates the end of the transmission and the total amount of data spooled. See 450-1021-181 for examples of spooling output.

3.11 The system spools SMDR data from the hard disk through a LAN interface unit (LIU) RS-232C port (as specified in the customer table: see 450-1021-313). Connections to the LIU RS-232C port can be direct (using a modem eliminator) if the link is less than 60m in length. For distances of more than 60m, appropriate modems must be used. The transmission speed of data may be 1200, 2400, 4800, 9600, or 19.2k bits per second.

3.12 Two methods of data spooling are available:

Dedicated The customer is continuously connected to the DNC-50, which runs a SPOOL job to transfer the data at regular intervals (as specified in the scheduler timetable, which is described later in this publication).

Interactive The customer periodically establishes a connection with the DNC-50 and signs on to request data spooling.

3.13 System parameters may be configured to permit all spooling agents to remain active for up to 24 hours at a time.

Multi-Node Spooling

3.14 Multi-node spooling is data collection from all of the customer's switch nodes and spooled in succession, one node after another through a single port, without inactive time realized on a per node schedule basis.

Dedicated Spooling Mode

3.15 The operating company may choose to assign one or more LIU ports to data links dedicated to spooling SMDR data to a single customer. In order to do this, the modem link between the LIU RS-232C and the customer's printer or computer must be established at all times. A SPOOL job then is registered on the system, specifying which customer the SMDR data is to be spooled to and the periods at which the data is to be spooled (see 450-1021-313 for details). At the scheduled times, the system formats and sends all SMDR data not previously spooled to that customer.

Interactive (Dial Up) Spooling Mode

3.16 If the operating company chooses to provide interactive spooling, a customer can dial in and connect with the DNC-50 to request the latest SMDR records at any time. No SPOOL job needs to be scheduled because the system monitors the host modem port at all times.

3.17 The customer contacts the station detail server and receives spooled data through one modem link that connects to a LIU port on the DNC-50. The customer must supply the correct customer identification, password, and options in order to be recognized by the system. If the proper identification is supplied, the system assembles, formats, and transmits all SMDR records that it has not previously spooled to that customer. During the transmission, the system recognizes standard XOFF and XON flow control characters used by the customer, and responds to an ESCAPE character by ending the session. The system also breaks the connection automatically if no data is sent or received for five minutes. See 450-1021-313 for details of the procedure for receiving spooled SMDR data from a DNC-50 in interactive mode.

TAPE GENERATION

3.18 If the tape generation feature is included with a station detail server, SMDR records on a DNC-50 disk can be written to a 9-track reel-to-reel tape at scheduled intervals or on demand. The sequence of records and fields on the tape is identical to the output produced by the DMS nodes, but all fields on the tape are in ASCII format.

3.19 For scheduled tape generation, the schedule is defined in the scheduler timetable. A TAPE job is scheduled for each customer according to how often a tape is required. (See 450-1021-313 for details of this procedure.) At the specified time, all SMDR data received for the customer since the last tape generation is dumped to tape.

3.20 The operating company can also create tapes on demand. There is a facility for listing the records on the DNC-50 hard disk and selecting them to be output to tape. See 450-1021-313 for details of this procedure.

SERVICE DATA MANAGER TABLES

3.21 An enhancement to the SMDR customer data spooling feature is the provision of a per customer LIU port configuration. This enhancement enables the async passthru initialization data to be obtained from the service data manager (SDM) tables instead of from INIT files. For details of the SDM tables for BNM, and what data they contain, refer to 450-1021-153.

DISK UTILIZATION REPORT AND MONITOR

3.22 This feature permits the operating company to:

a. Use the SDM table editor to define disk utilization levels, of SMDR data collection at pre-determined time intervals.

b. To define utilization levels for user-defined features (UDFs). See 450-1021-153.

Operational measurements, recording the number of data files and their total volume for each application, are generated (see the chapter on BNM operational measurements in this publication). The disk monitoring process takes place once an hour, or when the SDM is notified of table changes to DMOP or DMUDF.

3.23 Two BNM tables are used in this feature. The operating company user can make changes to these tables to suit the installation. The tables are called the DMOP table (disk monitoring operation table), and the DMUDF table (disk monitoring user defined features table) respectively.

DMOP Table

3.24 The DMOP table contains the parameters that can be changed by the operating company using the SDM table editor. The DMOP table follows the format shown in Fig. 3-1.

DMUDF Table

3.25 This table gives the user the facility to monitor the disk utilization of applications software other than the ones defined by default. To use this facility the user assigns a feature name to the application, and defines the feature name in terms of a directory path name.

3.26 There are three fields in the DMUDF table, Tuple_ID, Feature and Pathname. After the application is defined, the feature name assigned can be used in the DMOP table. The content of the Tuple_ID field must be the same as the one in table DMOP. See Fig. 3-1 for the format of the Tuple_ID field.

DNC PERFORMANCE MONITORING

3.27 The DNC performance monitoring feature provides the DNC user with information on the efficiency of disk and processor usage. The user is able to display the information using softkeys or typing commands on the keyboard of a terminal.

3.28 The feature consists of two programs: disk monitoring, and processor monitoring. Both programs are accessed from the DNC main menu under the option DNC services.

Disk Monitoring

3.29 The disk monitoring program provides the following information for each of the selected servers:

(a) the name of the server being monitored

(b) the total volume size in blocks of 1K bytes

(c) the number and percentage of blocks currently in use

(d) the number and percentage of blocks in use after the audit

(e) the date and time of the last audit

(f) the percentage and time of the last peak disk usage

(g) the time interval between disk updates

(h) the total number of files on the volume

3.30 Information on the menu items and screens available for the disk monitoring program can be found in 450-1021-311.

Processor Monitoring

3.31 The processor monitoring program provides the following information for each of the selected processors:

(a) a sequencing number that is used to sequence the processors being monitored

(b) the location of the processor being monitored

(c) the time interval between processor information updates

(d) the percentage of time the processor was busy during the last interval

(e) the amount of memory currently available

(f) the peak percentage of time the processor was busy

(g) the date and time of the peak processor usage

3.32 Information on the menu items and screens available for the processor monitoring program can be found in 450-1021-311.

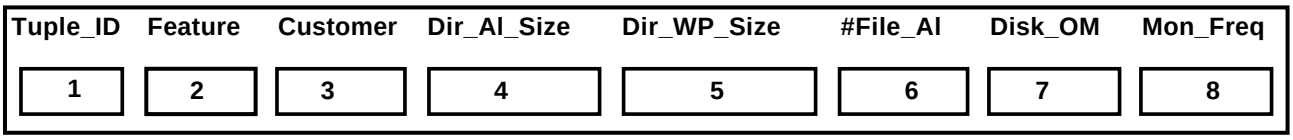

Field Description

1. This field contains an arbitrary, but unique name identifying a row of entries in the table. For user-defined features, it must have a matching pair: one in table DMOP, and one in table DMUDF.

2. This field is where the software application is identified. The field can contain up to 9 characters. The software application, SMDR Data Collection (SM), is defined by default. The user can also design a unique feature in terms of a path name of the features data directory.

3. This field contains the customer name. It is only used for the default software applications of SA and SM.

4. This field identifies the upper limit of the directory size (max. 32,767). If the limit is exceeded an alarm is raised. This field should be set to a lower number than the write protect field (number 5 in this figure). If the field is set to -1 the alarm is inhibited.

5. This field identifies another upper limit of the directory size (max. 32,767). In this case, if the limit is exceeded the directory is write protected, effectively preventing additional data from being added to the directory. If the field is set to -1 the write protect is inhibited.

6. This field contains the upper limit for the number of files for the current feature. If this limit is exceeded an alarm is raised. If the field is set to -1 the alarm is inhibited.

7. This is a YES or NO field. It determines whether OMs are generated for the feature.

8. This field contains the relative frequency of monitoring the data directories. If, for example the field is set to 4, then the data directories will be monitored every fourth time the Disk Monitoring and Utilization feature is used. If the field is set to -1, monitoring is disabled.

Dwg. 1021-101-015d

Fig. 3-1
DMOP Table Format

4. SDS USER INTERFACE

TERMINALS

4.01 Administrators normally communicate with DNC-50 systems by way of Northern Telecom M4000-series terminals. Digital Equipment Corporation VT100 terminals and other compatible ASCII terminals can also be used, provided they meet the ANSI X3.64 standard. Terminals are connected to the LANLink SRU of the DNC-50 through a LAN interface unit (LIU), and suitable modems must be used if a remote connection is used over a data line.

4.02 The station detail server accepts instructions from various "hardkeys" on terminals whose function does not vary. The SDS also sets up eight "softkeys", whose function depends on which display is on the screen. On the M4000-series terminals, the names of all the standard and special hardkeys are labeled on the keys, and the softkeys are in a row on the bottom of the screen. VT100 and other ASCII terminals must use different key combinations for the softkeys and some of the hardkeys used by the station detail server.

4.03 Refer to 450-1021-313 for descriptions of the terminals, and for the procedures for using the terminals and screens.

DATA TRANSFER FROM THE SWITCH OR PBX NODE

4.04 As the DMS node operates its various features, it generates and stores station message detail record (SMDR) data on its own holding disk. In an SDS, this data is sorted and partitioned on a per customer basis, and is transferred to a holding disk on the DNC-50 when the DNC-50 issues a request. The DNC-50 then sorts the data into separate files by customer.

4.05 Data from a non-DMS node or PBX node is converted into the DMS node format before it is input to the SDS. The SDS handles this converted data in the same way as it handles DMS node data. The conversion of data is made by a switch/PBX poller (see 450-1021-131 for details of the conversion method).

4.06 There are two ways in which a DNC-50 can receive data from a switch or PBX node:

- continuous collection
- demand transfer

4.07 A demand transfer request cannot be made during continuous collection and continuous collection cannot be started until all demand transfers are finished.

Continuous Collection

4.08 Continuous collection (also known as "near real-time transfer") is the method commonly used to transfer SMDR data. Continuous collection starts when a DNC-50 administrator makes a connection to the switch node (known as "logging on" to the node) and requests the node to start sending records to the DNC at the same time as they are being stored in the node's holding disk. The DNC-50 accumulates these records on its own hard disk. This accumulation continues until:

- a command is issued from the DNC-50 to stop sending the records, or
- an abnormality occurs

4.09 As records are collected by the DNC-50, they are sorted into files for each customer. Files are internally sectioned by time into partitions so that blocks of data can be identified and manipulated.

Demand Transfer

4.10 Demand transfer is used when recovery of individual data files is required. This type of transfer is used for unusual circumstances, such as when data already received is accidently erased.

4.11 To request a demand transfer, a DNC-50 administrator issues commands at a DNC-50 terminal to log onto the DMS node and to transfer a particular file to the DNC-50 storage disk. There are two ways to identify the file to be transferred:

- Enter its name. DMS node's file names consist of 15 to 17 characters that end in a file type; for example, U123456789012SMDR.
- List the available files and select one.

4.12 If the data file is no longer resident on the switch node disk, personnel at the switch node can manually rewrite the file onto the disk from backup tapes.

Availability of Data

4.13 The DNC-50 displays the availability of SMDR data as one of the following:

Unavailable	The DNC-50 is not logged on to any of the switch or PBX nodes, and is not collecting SMDR data.
Available	The DNC-50 is logged on to at least one of the nodes and SMDR data is available for collection.
Collect	The DNC-50 is currently in a continuous collection state.
Demand Transfer	The DNC-50 is currently engaged in a demand transfer.

4.14 A file on the DMS node holding disk can be in any of the following states:

Active An active (A) file is on the DMS node holding disk and is currently being written to by the DMS node as records are generated. File names of active files start with A.

Unprocessed An unprocessed file is an active file that has been closed, renamed so that it starts with U but otherwise retains the same filename, and has been added to the table DIRPHOLD in the DMS node. An unprocessed file becomes a processed file when it has been removed from table DIRPHOLD after being transferred to the DNC-50. It remains on the node's holding disk until the space is required for more files. The file is flagged as 'sent' when viewed by a DNC-50 user. A processed file can still be recovered by a DNC-50 user with manual assistance at the DMS node. Processed files are not visible when files are listed at the DNC-50.

Exception An exception file has a filename beginning with R. It is a processed file that has been reinserted into DIRPHOLD from the DMS node disk as a result of manual intervention at the node.

Impact on the Node

4.15 Manual intervention at a DMS node is required for the following aspects of DNC-50 operation (See 297-1001-001 for a listing of the Practices describing DMS nodes).

- **Configuration.** Various tables on the DMS node must be completed to define DIRP information for DNC-50 data collection. See 450-1021-353 for more information.
- **Resending Files.** Files that have already been sent to a DNC-50 can be resent by manual intervention at a DMS node. If the file is still on the DMS node holding disk, it can be renamed and resent by DMS node commands. If it is no longer on disk, DMS node personnel must edit table DIRPHOLD, and the tape containing the file must be mounted manually.

4.16 Manual intervention on a non-DMS node is governed by the type of node and how the data is collected.

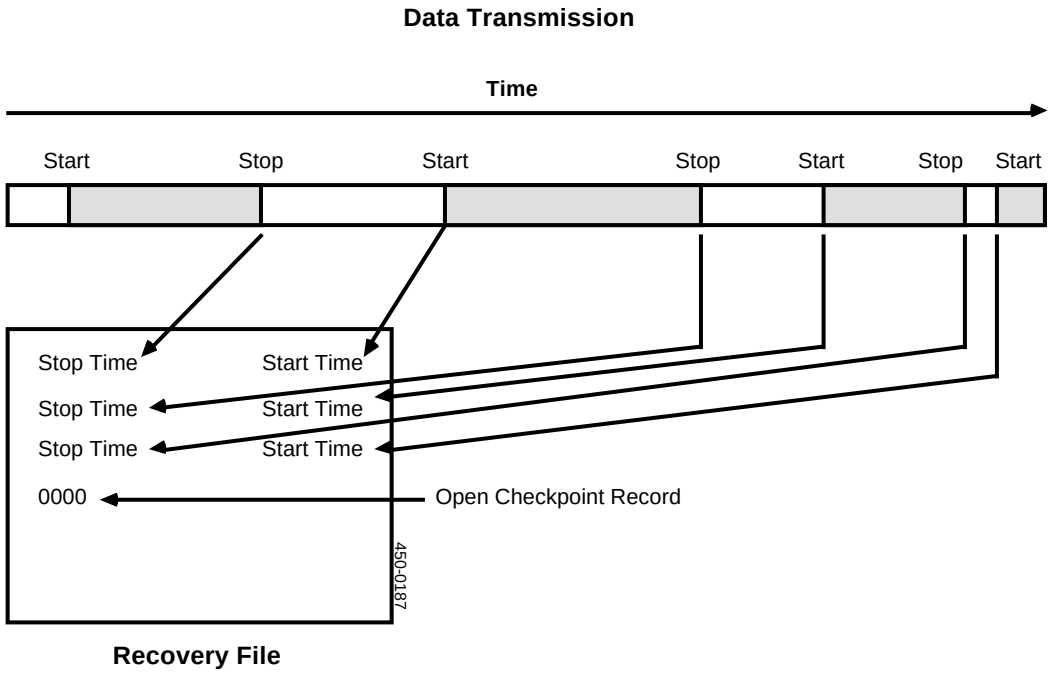

Fig. 4-1
Recovery File Creation

Recovery from Interruption of Data Transmission

4.17 If data transfer from a node to a DNC-50 is interrupted due to a faulty link, the DNC-50 recognizes the condition and tries to re-establish the connection. If unsuccessful, the system tries again after specified intervals (successively 30, 300, and another 300 seconds), which are adjustable by Northern Telecom to suit customer requirements. The system continues these automatic logon attempts (even if the user manually logs back on) until it brings the data link back up and re-establishes the connection, or until the message 'A Communication Session has failed' is displayed. In the latter case, the user must manually re-establish the logon.

4.18 Whenever data transfer from a node to a DNC-50 is interrupted due to a faulty link or error in data transmission, or even due to a 'Stop Collect' command issued by the DNC-50 administrator, the DNC-50 makes a note of the time of interruption. When transmission is restarted, the time interval during which data was missed is shown in a recovery file that the administrator can view on the screen.

4.19 This recovery file is created only after a certain amount of data has been lost. If the link recovers before the threshold is exceeded, the data is recovered automatically by the DNC-50 without intervention by the administrator. If the threshold is exceeded, the DNC-50 creates a recovery file showing the interruptions.

4.20 Fig. 4-1 shows the sequence of events when a link is interrupted. As the data is transmitted from the node to the DNC-50, a checkpoint is periodically sent and recorded by the DNC-50. The DNC-50 updates the first checkpoint block in the recovery file, and continues updating it as new data is received.

4.21 When collection stops for any reason, the checkpoint is frozen at its value at that time. When data transmission is restarted, the next checkpoint is entered into the file and paired with the first checkpoint. The two checkpoint values show the beginning and ending of the discontinuity in data transmission.

4.22 At this point, a new checkpoint is started and updated as new checkpoints are received, and the cycle is repeated if collection stops again. The resulting recovery file is a list of all of the intervals for which data collection did not take place.

4.23 From this list of intervals in the recovery file, the administrator can select any that are to be recovered. The DNC-50 will then recover all missing feature data for the intervals concerned. Once the data is successfully retrieved, the interval is automatically deleted from the recovery list.

JOB SCHEDULING

4.24 The scheduler timetable is a feature of DNC-50 used to schedule system operations, such as data spooling, that take place on a routine basis. Scheduling can be done in a flexible and timely manner with minimal operator assistance.

4.25 The DNC-50 administrator can set up schedules for

- generating spooled SMDR data for customers
- administration (internal to the DNC-50)

4.26 The administrator creates a job, then associates "timespecs" with the job. The job defines what is to be done and what data is used for it. The timespecs define:

- when the job is to start for the first time
- the time interval until the next occurrence, and subsequent occurrences
- when the job is to stop being carried out (until scheduled to start again)

4.27 The timespecs are repeated with a frequency that can be one of

- daily
- weekly
- monthly

Jobs

4.28 First, the administrator defines a job. In order to do so, the administrator enters the following information using a form displayed on the terminal:

Customer Name	The name of the customer for whom the data is to be compiled. A job can be run for only one customer at a time.
Function	The type of operation, such as SPOOL. There can be only one function for each job.
Node	The name of the DMS node or nodes, as defined in the DNC Node Table, from which data is collected.
Collection Interval	The interval of measurement to be used in compiling data. The node accumulates measurements on a daily, weekly, and monthly basis for each feature.
Frequency	The frequency with which this function is to be

carried out by the DNC-50 system: daily, weekly, or monthly.

Daily A daily job is carried out one or more times a day, and adheres to the same schedule every day, 7 days a week.

Weekly A weekly job is carried out one or more times a week, and adheres to the same schedule every week, 52 weeks a year.

Monthly A monthly job is carried out one or more times a month, and adheres to the same schedule every month, 12 months a year.

Timespecs

4.29 Associated with each job in the DNC-50 Scheduler Timetable is a table of timespecs. The timespecs show when the job starts, how often it runs, and when it stops each day.

4.30 Up to eight timespecs are allowed for each job. These timespecs may not overlap. They allow for variances in the schedule of the same job so that accommodation can be made for such things as busy periods, nonbusy periods, weekends, and holidays.

4.31 The format of the timespecs depends on the frequency that was selected for the job:

(a) **Daily:** For daily jobs, each entry in the timespec table defines a start time, end time, and time period. The start time determines the first occurrence of the job. The time period determines the time until the next occurrence of the job, and subsequent occurrences, until the end time is reached.

Example: If data is to be spooled

- (1) From 0800 through 1600 every 180 minutes
- (2) From 1600 through 2400 every 360 minutes

the DNC-50 spools at 0800 using all formatted data currently stored on the hard disk. The next spooling occurs at 1100 (180 minutes later), and contains records for all formatted data received since 0800. Another spooling takes place at 1400 (180 minutes later), and contains records for all formatted data received since 1100. The next spooling occurs at 1600 (the start of the second timespec), and contains records for all formatted data received since 1400. Another spooling takes place at 2200 (360 minutes

later), and contains records for all formatted data received since 1600. The end of the time period, 2400, occurs before the next spooling. The next day, the daily schedule starts over again at 0800.

(b) Weekly: For weekly jobs, each entry in the timespec table defines a start day (Sun-Sat), end day, start time, end time, and time period. The start day and start time determine the first occurrence of the job. The time period determines the time until the next occurrence of the job, and subsequent occurrences, until the end day and end time are reached.

- **Example:** If SMDR data is to be spooled to a customer:
- (1) Mon-Fri From 0800 through 1600 every 180 minutes
- (2) Sat-Sun From 0800 through 2400 every 360 minutes

the DNC-50 spools on Monday at 0800 all relevant SMDR data currently stored on the hard disk. The next spooling occurs at 1100 (180 minutes later), and contains records for all SMDR data received since 0800. The next set of SMDR data is spooled at 1400 (180 minutes later). The 1600 end time occurs before the next scheduled job, so the next data is spooled at 0800 the next morning, and so on. This continues until Friday at 1600. Data is spooled again on Saturday at 0800 (the start of the second timespec), and contains all relevant SMDR records received since 1400 on Friday. The next spoolings occur at 1400 and 2000 (360 minutes later in each case), and contain all SMDR records received since the last job. The end of the time period, 2400, occurs before the next spooling, so the next data spooling occurs at 0800 on Sunday, and so on until Monday morning.

(c) Monthly: For monthly jobs, each entry in the timespec table defines a start date and end date (for example, 1-31), a start time, an end time, and a time period. The start date and start time determine the first occurrence of the job. The time period determines the time until the next occurrence of the job, and subsequent occurrences after that, until the end day and end time are reached.

- **Example:** If SMDR data is to be spooled to a customer:
- (1) 1-20 From 0800 through 1600 every 180 minutes
- (2) 21-31 From 0800 through 2400 every 360 minutes

the DNC-50 spools on the first of the month at 0800 with all formatted data currently stored on the hard disk. The next spooling is at 1100 (180 minutes later), and contains records for all formatted data received since 0800. Another spooling takes

place at 1400 (180 minutes later), and contains records for all formatted data received since 1100. The 1600 end time occurs before the next spooling, so the next job is at 0800 the next morning, and so on. This continues until the 20th at 1600.

- The next spooling takes place on the 21st at 0800 (the start of the second timespec), and contains records for all formatted data received since 1400 on the 20th. Another spooling occurs at 1400 (360 minutes later), and contains records for all formatted data received since 0800. The next spooling is at 2000 (360 minutes later), containing records for all formatted data received since 1400. The end of the time period, 2400, occurs before the next spooling, so the next spooling is not until 0800 on the 22nd. Another spooling occurs at 1400 on the 22nd, and so on until the 31st (or when the month expires, if sooner); then the first timespec comes into effect again on the morning of the first of the next month, and so on.

Scheduled Job Queues

4.32 The job queue is a list of events, each of which is the next immediate occurrence of the jobs entered into the Scheduler Timetable.

4.33 Each event in the queue can be removed from the queue or set to take place at a new time. Note that such a change in setting affects only one event, not all the events in the queue.

- In the case of deletion, the event is removed and the administrator exits out of the queue menu. If the administrator reaccesses the queue, the event will have been replaced by the event that was scheduled after the deleted event
- In the case of reordering, the administrator can reschedule the event for a new date and time. Once this event has been carried out by the system, the next event takes place at its regularly scheduled time.

SMDR DATA STORAGE ON DNC-50

4.34 The DNC-50 administrator can access the files stored on the DNC-50 internal holding disk. The files are shown in the form of partitions for each customer. With each partition is shown its status, one of Unformatted (hexadecimal) or Formatted (ASCII). While inspecting the list of files, the administrator can select a partition and

- delete it
- dump it to tape

4.35 Partitions are identified by start and end time. They can be 1-minute, 2-minute, 15-minute, 30-minute, hourly, daily, weekly or monthly partitions, depending on the type of data and how the DMS node stores it.

5. OPERATIONAL MEASUREMENTS

BNM OPERATIONAL MEASUREMENTS

5.01 This chapter describes the BNM operational measurements (OM) that are collected by the DNC-50. At this time, the only BNM feature that generates OMs is the disk utilization report and monitor feature.

Basic DNC Operational Measurement Concepts

5.02 Each OM is a count of how often a particular event happens. OMs are collected in registers, which are segments in the DNC-50 storage memory. There is one register for each type of event. Each register is identified by a name with a maximum of 8-characters. This name is fixed and cannot be changed by the user. OM counts take place during time intervals, known as collection intervals. The collection interval can be set to 5, 15, 30, or 60 minutes.

5.03 Groups. OM registers are collected into OM groups. A register can belong to only one group, but groups can have one or more registers. Like the register name, a group is identified by a name with a maximum of 8-characters, which also cannot be renamed by the user. Also, the user cannot transfer a register from one group to another.

5.04 Key Fields. Certain OM groups may be subdivided into key fields. In these groups, OMs are collected for multiple instances of a number of events. An example is an OM group for customer activity; such a group may have a separate key field for each customer. The key field is identified by an integer, which is unique to each set of events in the group. In reports, the printout appears as a table, with the registers across the top defining columns and the keys down the left hand side defining the sets of events (one event per column).

5.05 Classes. An OM class is a customer-defined set of OM groups. Users can create OM classes and enter OM groups into those classes. The user can then define how the classes are to be treated.

5.06 Collection. For any OM register, the system software sets aside two areas in memory. One of these areas is used as an "active" register. The active register is the register currently in use for counting. When the accumulating interval, usually 5 minutes, is up, the active register is frozen and becomes the "holding" register. The second register takes over the counting and becomes the active register. Meanwhile, the contents of the holding register is written to memory, and on completion the register is reset to zero.

5.07 When the accumulation interval is again completed, the new active register is frozen (to become the holding register), and the first register starts the count function, completing a cycle.

5.08 Accumulation. Data, for the purpose of reporting and trend studies, must be accumulated over longer periods than the basic period of active register counting. The accumulating registers, of which there are several types for various time periods (hourly, daily, weekly,

monthly), are used for this purpose. The data accumulation process adds the contents of holding registers to the accumulation register just prior to the next active-to-holding register transfer. The accumulated data is not available until the end of the accumulating period.

OM Report Generation

5.09 OM reports are produced by a number of Program Resource Units (PRUs) which are collectively called the OM Report Generator. These PRUs are described in the following list, and their relationship to each other is shown in Fig. 5-1.

- **OM Collector Agent (OMCA).** All OM collection points are associated with an OMCA.
- **OM Collector (OMC).** The OMC periodically collects data from the OMCAs and stores it on the DNC-50 system disk. The configuration of the OMC is controlled by data in table OMC_CONFIG.
- **OM Reporter (OMR).** The OMR retrieves operational measurement data from the system disk and generates an operating company defined report.
- **DNC Report Generator.** This report generator prepares the reports using data supplied by the OMR.
- **Service Data Manager (SDM).** The SDM stores information to access the table definition (illustrated in Fig. 5-2) that was prepared by the operating company.

Operation Of OM Report Generator

5.10 The OMC collects data from the the OMCAs at the end of each collection interval, and stores the data on the system disk. The collection interval is set during installation, or when the load is built.

5.11 The OMR uses the SDM to access the appropriate tables for determining the types of reports to be generated, and the location of the OM data. After this information has been obtained, the OMR retrieves the OM data from the system disk and processes it in accordance with the schedule and format established in the OM definition.

5.12 The OMR then creates new data files formatted to be compatible with the DNC Report Generator, then invokes the DNC Report Generator to produce a specific report.

5.13 The DNC Report Generator then uses the formatted data to generate and send a report to the user's printer.

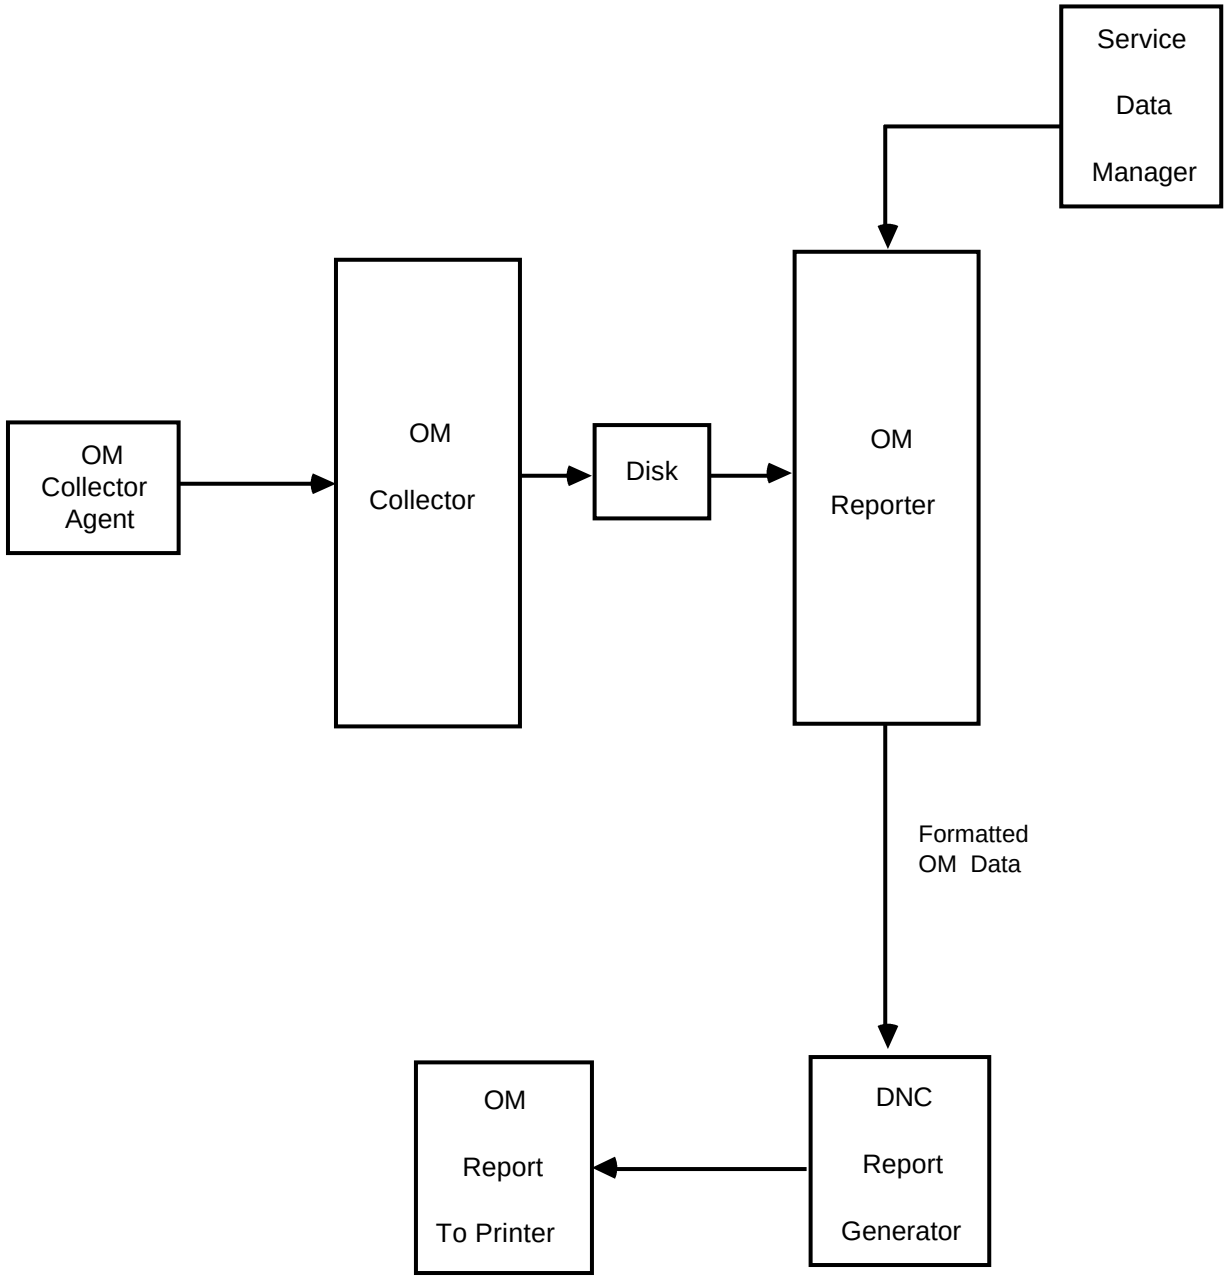

Fig. 5-1
OM Report Generator

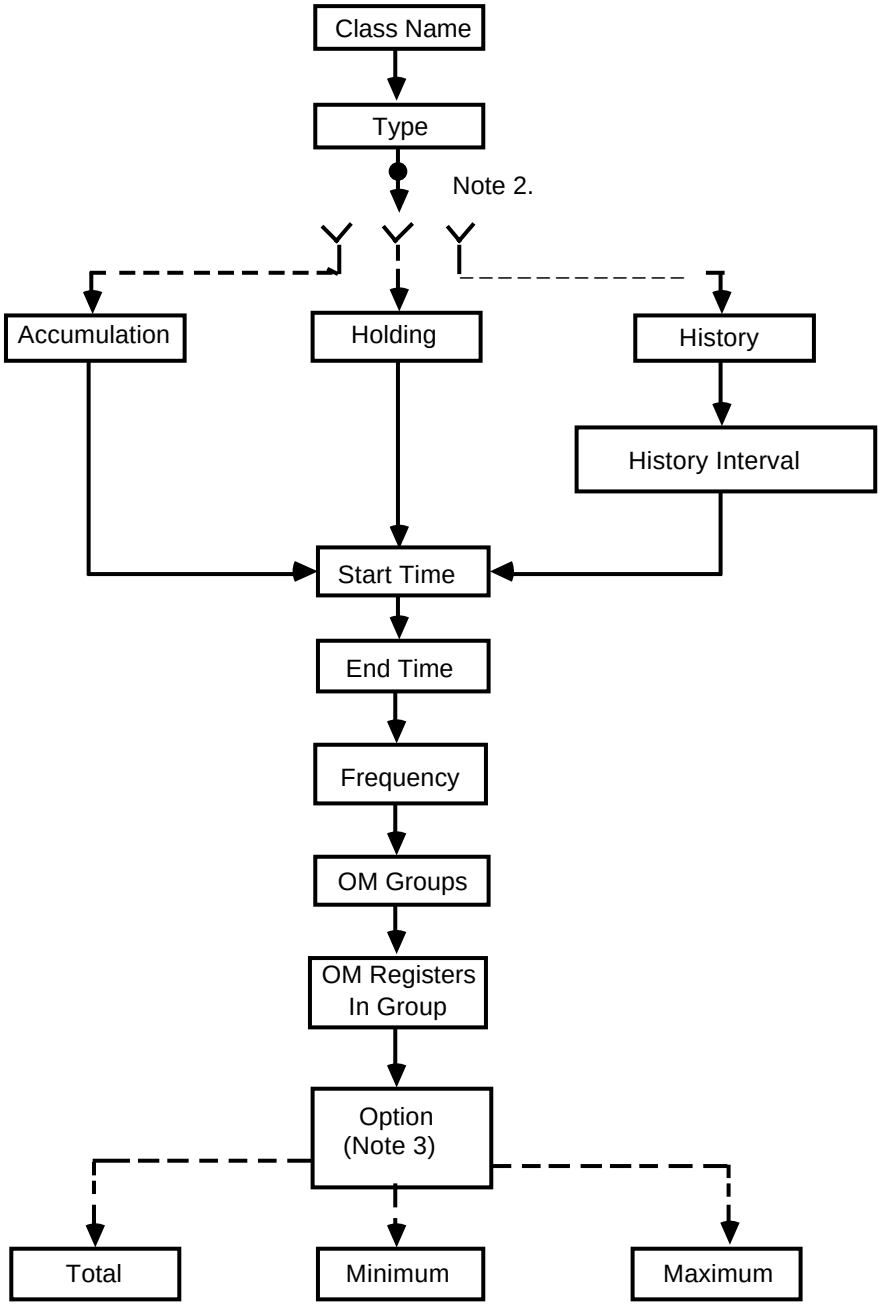

Notes:

1. ▬ ▬ ▬ indicates that a choice is required
2. One type of report is required.
3. One or more options can be selected, or none can be selected.

Fig. 5-2
OM Table Definition

Report Definition

5.14 The OMR requires that each report be defined so that the data is assembled and displayed to meet the needs of the operating company. Fig. 5-2 shows the factors that must be defined for every report.

5.15 An example of a report is shown in Fig. 5-3. This example is an Accumulation type of report. When the quantity of registers is more than can be accommodated on a single page, additional pages are printed with the key designation on each page.

5.16 When the report type is History, an additional column is displayed to the left of the KEY column, displaying the time of the report for each set of data.

Report Types

5.17 The following reports are generated for OMs

- **Accumulation Reports.** The raw OM data that is collected by the OMR at the end of each collection interval is stored at the same location on the system disk, for the defined duration of the accumulation register report.
- **History Reports.** The data for each report measurement interval is reported separately for the duration of the report (the report measurement interval is a multiple of the collection interval).
- **Holding Report.** The data at the end of each collection interval is reported as a holding report.

Report Options

5.18 Options are available for a History Report or an Accumulation Report only. These options are as follows

- Total Values
- Minimum Values
- Maximum Values

5.19 When an option is set, the report is supplemented with the information contained in Table 5-A.

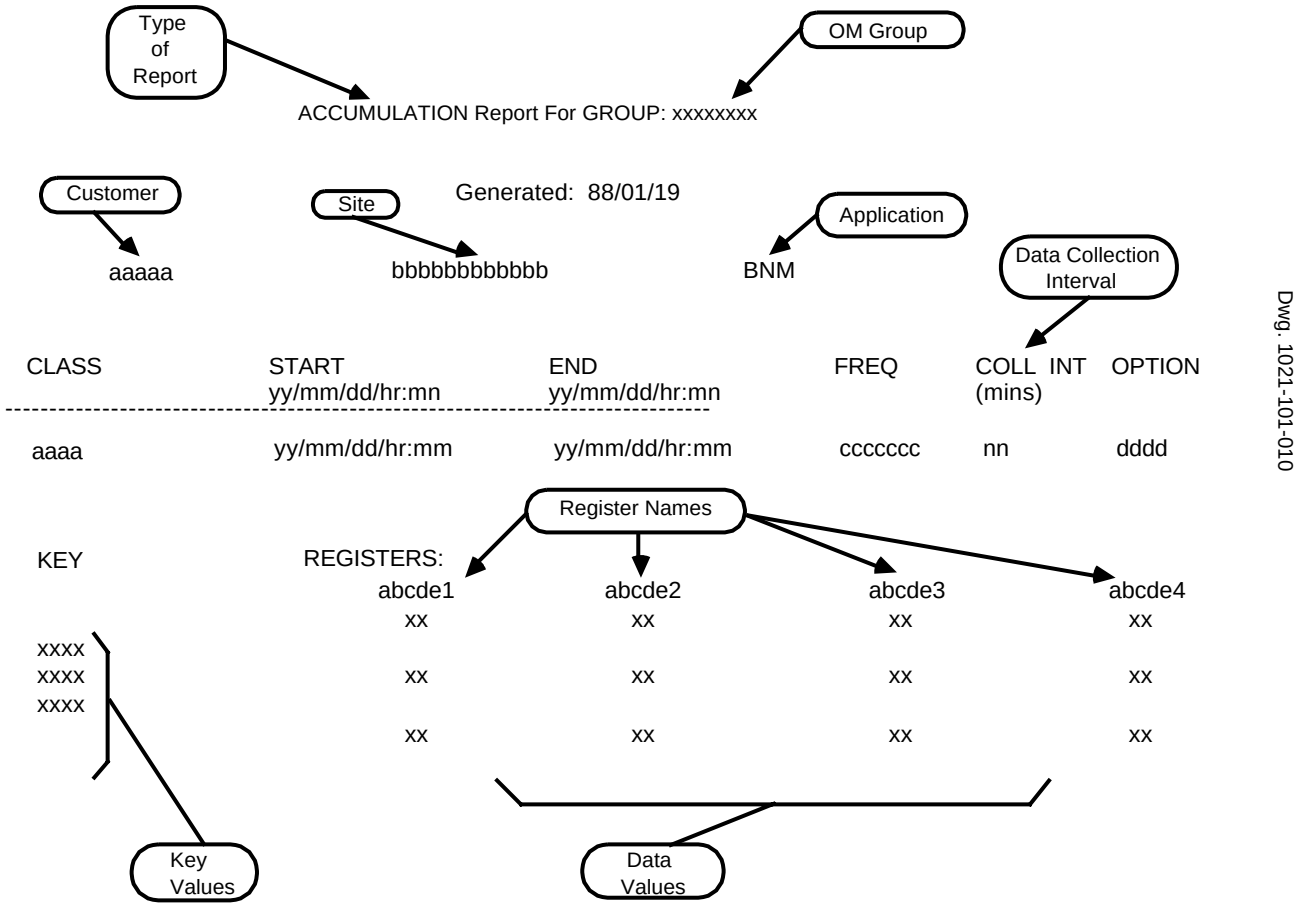

Fig. 5-3
Format Of A Typical OM Report

Table 5-A
REPORT OPTIONS

REPORT	TOTAL	MINIMUM	MAXIMUM
Accumulation	Valid for groups with two or more keys that are associated with registers.	The minimum value for each register, and the collection times, are reported on a separate page.	The maximum value for each register, and the collection time, is reported on a separate page.
	Register totals are displayed on an extra line at the end of the report.		
History	A separate Accumulation report is prepared for each OM group that is selected.	The minimum value for each register, and the collection times, at the end of each report interval are reported on a separate page.	The maximum value for each register, and the collection times, at the end of each report interval are reported on a separate page.
	If the group has keys associated with the registers, a subtotal is provided at the end of each report interval.		

Start And End Times

5.20 The start time and end time for all reports are expressed as follows

yy/mm/dd hh:mm

Where:

yy is the last two digits of the year

mm is the month expressed as digits

dd is the day expressed as digits in the following way:
Monthly Reports -- is two digits identifying the day of the month
Weekly Reports -- is a single digit within the range 1 to 7 representing Monday through Sunday
Daily Reports -- left blank

hh is two digits representing the hours of a 24 hour clock

mm is two digits representing minutes

5.21 When the start time is before the time that the report was defined, then the first report will not contain data collected between these two times. The actual start time of the report is recorded in the report header the report header.

5.22 When the start time is not accompanied by an end time, only one report is prepared.

5.23 Reports can be prepared to be run at the following frequencies

- **Single Report.** The start and end times are the same, and must be stated.
- **Hourly Reports.** Reports are generated according to the start time. When the minutes are not stated, reports are generated on the hour.
- **Daily Reports.** Reports are generated according to the start and end times. Periods within each day can be selected. If hours and minutes are not defined, 24 hour reports will be generated. The minimum daily report period is one hour.
- **Weekly Reports.** Reports are generated for the parts of the week defined in the start and end times. If no day, hour, or minute is defined, reports will be generated for Monday through Sunday starting at midnight on Monday. The minimum weekly report period is one day.
- **Monthly Reports.** Reports are generated for the parts of the month defined in the start and end times. If no values are put in the start and end times, the monthly report will contain data for the complete month. The minimum monthly report period is one day.
- **Hold.** Using this frequency the report is generated once, at the end of the report period provided a start and end time is defined.

Setting Up OM Reports

5.24 OM reports are set up using software tables. The tables may be edited using the Fully-Generic MMI Program (FGMP). The tables and their interaction are illustrated in Fig. 5-4.

5.25 The following tables are used to set up OM reports

(a) OM Group Name. This table defines, for each OM register, whether it is a peg, usage, or snapshot register. Users cannot access this register.

(b) Report Class Definition. These tables are named according to the customer, application, and site. Users access these tables to input characteristics of the OM class.

(c) Table OMREPORT. This table is accessed by the user to associate the Class Definitions tables with the customer parameters. They contain information used to send the reports to the printer management feature, and to print site and customer information on the report headers.

Disk Utilization Report And Monitor OMs

5.26 There are three OM groups for the disk utilization report and monitor feature but only the SMDROM is used by the Station Detail Server. Refer to 450-1021-101 for details of the other OM groups.

5.27 SMDROM. This OM group contains the customers names as the key, and six registers for each key. There is a maximum of 64 customers to a group. The registers contain the following data:

- Register 1 is called fvol. It contains the value of the total volume size of SMDR data directory for a customer.
- Register 2 is called numfile. It contains the number of files that exist in the SMDR data directory for a customer.
- Register 3 is called DFrate. It contains the value of the SMDR data file creation rate for a customer
- Register 4 is called DRrate. It contains the rate of SMDR record creation for a customer.
- Register 5 is called MaxDF. It contains the maximum DFrate of the day.
- Register 6 is called the MaxDR. It contains the maximum DRrate of the day.

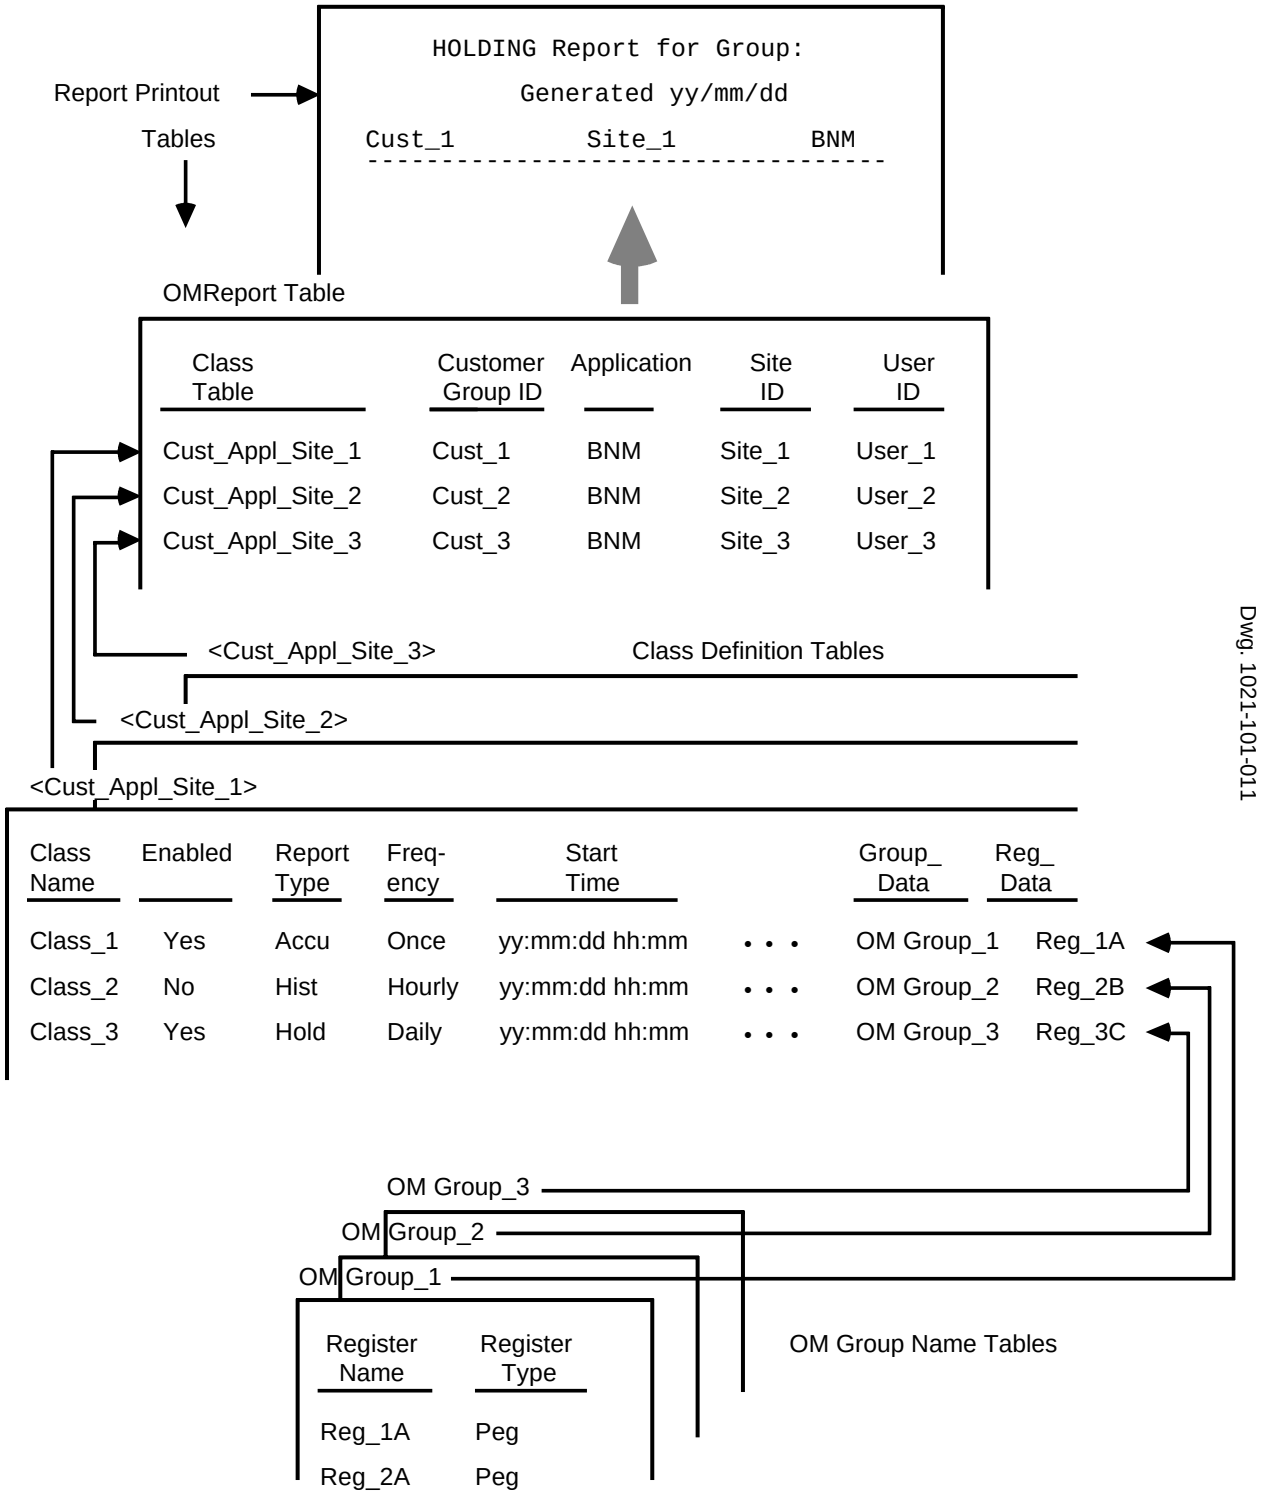

Fig. 5-4
OM Table Interactions

6. ABBREVIATIONS

6.01 The following abbreviations are used in this Practice:

Accu	Accumulative
ANSI	American national standards institute
Appl	Application
ASCII	American standard code for information interchange
BNM	Business network management
Cust	Customer
DDD	Direct distance dialing
DIRP	Device independent recording package
DMS	Digital Multiplex System
DMS node	A member of the DMS-100 family of digital switches. It includes the variants DMS-100, DMS-200, DMS-250, and DMS-300.
DNC	Dynamic Network Control System
FGMP	Fully-generic MMI program
Hist	History
IDDD	International direct distance dialing
LAN	Local area network
LIU	Local area network interface unit
MDC	Meridian Digital Centrex
MMI	Man-machine interface
NSR	Network software release
NTP	Northern Telecom practice
OM	Operational measurement
OMC	OM collector
OMCA	OM collector agent
OMR	OM reporter
PBX	Private branch exchange

PRU	Program resource unit
Reg	Register
SDM	Service data manager
SDS	Station detail server
SMDR	Station message detail record
Telco	Telephone company, also called operating company