SMDS Billing System Administrator's Guide

Ascend Communications, Inc.

Product Code: 104-00026-00
Revision 00
June 1995

ASCEND COMMUNICATIONS, INC. END-USER LICENSE AGREEMENT

ASCEND COMMUNICATIONS, INC. IS WILLING TO LICENSE THE ENCLOSED SOFTWARE AND ACCOMPANYING USER DOCUMENTATION (COLLECTIVELY, THE "PROGRAM") TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. PLEASE READ THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT CAREFULLY BEFORE OPENING THE PACKAGE(S) OR USING THE ASCEND SWITCH(ES) CONTAINING THE SOFTWARE, AND BEFORE USING THE ACCOMPANYING USER DOCUMENTATION. OPENING THE PACKAGE(S) OR USING THE ASCEND SWITCH(ES) CONTAINING THE PROGRAM WILL INDICATE YOUR ACCEPTANCE OF THE TERMS OF THIS LICENSE AGREEMENT. IF YOU ARE NOT WILLING TO BE BOUND BY THE TERMS OF THIS LICENSE AGREEMENT, ASCEND IS UNWILLING TO LICENSE THE PROGRAM TO YOU, IN WHICH EVENT YOU SHOULD RETURN THE PROGRAM WITHIN TEN (10) DAYS FROM SHIPMENT TO THE PLACE FROM WHICH IT WAS ACQUIRED, AND YOUR LICENSE FEE WILL BE REFUNDED. THIS LICENSE AGREEMENT REPRESENTS THE ENTIRE AGREEMENT CONCERNING THE PROGRAM BETWEEN YOU AND ASCEND, AND IT SUPERSEDES ANY PRIOR PROPOSAL, REPRESENTATION OR UNDERSTANDING BETWEEN THE PARTIES.

1. License Grant. Ascend hereby grants to you, and you accept, a non-exclusive, non-transferable license to use the computer software, including all patches, error corrections, updates and revisions thereto in machine-readable, object code form only (the "Software"), and the accompanying User Documentation, only as authorized in this License Agreement. The Software may be used only on a single computer owned, leased, or otherwise controlled by you; or in the event of inoperability of that computer, on a backup computer selected by you. You agree that you will not pledge, lease, rent, or share your rights under this License Agreement, and that you will not, without Ascend's prior written consent, assign or transfer your rights hereunder. You agree that you may not modify, reverse assemble, reverse compile, or otherwise translate the Software or permit a third party to do so. You may make one copy of the Software and User Documentation for backup purposes. Any such copies of the Software or the User Documentation shall include Ascend's copyright and other proprietary notices. Except as authorized under this paragraph, no copies of the Program or any portions thereof may be made by you or any person under your authority or control.

2. Ascend's Rights. You agree that the Software and the User Documentation are proprietary, confidential products of Ascend or Ascend's licensor protected under US copyright law and you will use your best efforts to maintain their confidentiality. You further acknowledge and agree that all right, title and interest in and to the Program, including associated intellectual property rights, are and shall remain with Ascend or Ascend's licensor. This License Agreement does not convey to you an interest in or to the Program, but only a limited right of use revocable in accordance with the terms of this License Agreement.

3. License Fees. The license fees paid by you are paid in consideration of the license granted under this License Agreement.

4. Term. This License Agreement is effective upon your opening of the package(s) or use of the switch(es) containing Software and shall continue until terminated. You may terminate this License Agreement at any time by returning the Program and all copies or portions thereof to Ascend. Ascend may terminate this License Agreement upon the breach by you of any term hereof. Upon such termination by Ascend, you agree to return to Ascend the Program and all copies or portions thereof. Termination of this License Agreement shall not prejudice Ascend's rights to damages or any other available remedy.

5. Limited Warranty. Ascend warrants, for your benefit alone, for a period of 90 days from the date of shipment of the Program by Ascend (the "Warranty Period") that the program diskettes in which the Software is contained are free from defects in material and workmanship. Ascend further warrants, for your benefit alone, that during the Warranty Period the Program shall operate substantially in accordance with the User Documentation. If during the Warranty Period, a defect in the Program appears, you may return the Program to the party from which the Program was acquired for either replacement or, if so elected by such party, refund of amounts paid by you under this License Agreement. You agree that the foregoing constitutes your sole and exclusive remedy for breach by Ascend of any warranties made under this Agreement. EXCEPT FOR THE WARRANTIES SET FORTH ABOVE, THE PROGRAM IS LICENSED "AS IS", AND ASCEND DISCLAIMS ANY AND ALL OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTIES OF NONINFRINGEMENT.

6. Limitation of Liability. Ascend's cumulative liability to you or any other party for any loss or damages resulting from any claims, demands, or actions arising out of or relating to this License Agreement shall not exceed the greater of: (i) ten thousand US dollars ($10,000) or (ii) the total license fee paid to Ascend for the use of the Program. In no event shall Ascend be liable for any indirect, incidental, consequential, special, punitive or exemplary damages or lost profits, even if Ascend has been advised of the possibility of such damages.

7. Proprietary Rights Indemnification. Ascend shall at its expense defend you against and, subject to the limitations set forth elsewhere herein, pay all costs and damages made in settlement or awarded against you resulting from a claim that the Program as supplied by Ascend infringes a United States copyright or a United States patent, or misappropriates a United States trade secret, provided that you: (a) provide prompt written notice of any such claim, (b) allow Ascend to direct the defense and settlement of the claim, and (c) provide Ascend with the authority, information, and assistance that Ascend deems reasonably necessary for the defense and settlement of the claim. You shall not consent to any judgment or decree or do any other act in compromise of any such claim without first obtaining Ascend's written consent. In any action based on such a claim, Ascend may, at its sole option, either: (1) obtain for you the right to continue using the Program, (2) replace or modify the Program to avoid the claim, or (3) if neither (1) nor (2) can reasonably be effected by Ascend, terminate the license granted hereunder and give you a prorata refund of the license fee paid for such Program, calculated on the basis of straight-line depreciation over a five-year useful life. Notwithstanding the preceding sentence,

Ascend will have no liability for any infringement or misappropriation claim of any kind if such claim is based on: (i) the use of other than the current unaltered release of the Program and Ascend has provided or offers to provide such release to you for its then current license fee, or (ii) use or combination of the Program with programs or data not supplied or approved by Ascend to the extent such use or combination caused the claim.

8. Export Control. You agree not to export or disclose to anyone except a United States national any portion of the Program supplied by Ascend without first obtaining the required permits or licenses to do so from the US Office of Export Administration, and any other appropriate government agency.

9. Governing Law. This License Agreement shall be construed and governed in accordance with the laws and under the jurisdiction of the Commonwealth of Massachusetts, USA. Any dispute arising out of this Agreement shall be referred to an arbitration proceeding in Boston, Massachusetts, USA by the American Arbitration Association.

10. Miscellaneous. If any action is brought by either party to this License Agreement against the other party regarding the subject matter hereof, the prevailing party shall be entitled to recover, in addition to any other relief granted, reasonable attorneys' fees and expenses of arbitration. Should any term of this License Agreement be declared void or unenforceable by any court of competent jurisdiction, such declaration shall have no effect on the remaining terms hereof. The failure of either party to enforce any rights granted hereunder or to take action against the other party in the event of any breach hereunder shall not be deemed a waiver by that party as to subsequent enforcement of rights or subsequent actions in the event of future breaches.

Contents

List of Figures

About This Guide

The *SMDS Billing System Administrator's Guide*, which is intended for switching system administrators, describes the architecture of Ascend's billing system, and contains information on how to set up, install, configure, and manage the billing system hardware and software. This guide also provides basic troubleshooting solutions to assist the switching system administrator with resolving switching system problems that impact the collection, storage, transfer, and conversion of billing data.

This manual is intended to be used in conjunction with the *B-STDX 8000/9000 Hardware Installation Guide*, the *Software Release Notice for the Ascend Billing System Adjunct Processor*, and the *Ascend Network Administrator's Guide for CascadeView/UX*. You may also need the following manuals for reference:

- The *Solaris System and Network Administration Guide*
- The *Solaris 2.3 System Configuration and Installation Guide*
- The *SYBASE SQL Server Installation Guide for Sun Solaris*

Organization

The following list summarizes the information contained in this guide.

Chapter 1 — "Overview" provides general information on billing system components, and describes the general architecture of billing system software, as defined in Bellcore TR775[1].

Chapter 2 — "Ascend Billing System" provides more detailed information on Ascend's billing system architecture and its components, including the operational characteristics and functional responsibilities of each component.

Chapter 3 — "Preparing for the Installation" describes the hardware and software requirements for the Ascend billing system.

Chapter 4 — "Installing the Billing System" describes how to install and configure the hardware for the billing system. It also describes how to install the operating system and billing system software on the Adjunct Processors in your network, as well as how to install and configure SYBASE on the workstation you will be using as the billing system SYBASE server.

Chapter 5 — "Billing System Configuration" describes how to configure the Adjunct Processors in the billing system. It also describes how to enable billing on the B-STDX switches and SMDS logical ports in your switching system.

Chapter 6 — "Managing the Billing System" describes the Ascend billing system utilities, and how to use them to manage the billing system. It also describes how to display billing system information on the Adjunct Processor, as well as how to display billing system information on your NMS workstations via CascadeView/UX.

Chapter 7 — "Event Logging" describes the billing system event log file and the events that are logged to that file. It also provides possible causes and resolutions for error events that occur within the billing system.

Chapter 8 — "SNMP Traps" describes the SNMP trap events that are generated by the switches in the billing system, and sent to each NMS in the switching system. When applicable, possible resolutions are provided for those traps which indicate a problem or potential problem within the billing system.

Conventions Used

This guide uses many acronyms to describe various networking terms. Use the following list to quickly identify the meaning of the acronyms contained in this guide:

Acronym	Meaning
AP	Adjunct Processor
ATM	Asynchronous Transfer Mode
CP	Control Processor
DLCI	Data Link Connection Identifier
DXI	Data Exchange Interface
FTP	File Transfer Protocol
IOP	I/O Processor
IP	Internet Protocol
NMS	Network Management Station
NVRAM	Non-volatile Random Access Memory
PDU	Priority Data Unit
PRAM	Parameter Random Access Memory
PVC	Permanent Virtual Circuit
SMDS	Switched Multimegabit Data Services
SNI	Subscriber Network Interface
SVC	Switched Virtual Circuit
UCT	Universal Coordinated Time

This guide generically refers to the B-STDX hardware as “the switch.” Additional conventions used in this guide include:

Used to provide helpful suggestions or reference to materials not contained in this manual.

Used to warn the reader to proceed carefully in order to avoid equipment damage or personal harm.

Manual titles are shown in italics. For example:

Ascend Network Administrator’s Guide for CascadeView/UX

Program and directory names that appear within a paragraph are shown in bold type. For example:

/opt/CascadeBS
cascade.log

Menu selections that you have to make are shown as follows:

CascadeView ⇒ Logon

In this example, you would select CascadeView from the menu bar, then select Logon from the CascadeView pull-down menu.

Commands that you have to enter are indented and shown in bold Courier font. Variable parameters are shown in italics, and set off in square brackets. Also, <Return> indicates when to press the Enter key. For example:

udfdump [*usage data filename*] <Return>

In this example, **udfdump** is the command, and [*usage data filename*] is a variable parameter for which you would substitute an actual filename for a usage data file. For example:

udfdump bs_A_SMDS.152.148.50.1.30Mar95190000

Event log and SNMP trap event messages are shown in bold type and are not indented. For example:

Could not communicate with AP daemon

Where to Start

If your switching system has not been installed and configured yet, begin with the hardware installation instructions in the *B-STDX 8000/9000 Hardware Installation Guide*. Then refer to the *Ascend Network Administrator's Guide for CascadeView/UX* for instructions on how to configure the switches on your network.

Ascend recommends that you install and configure your switch network before you begin the billing system installation.

Once the switch hardware is installed and configured:

- If you not familiar with billing system components and terminology, begin with Chapters 1 and 2 of this manual. Then refer to Chapter 3 for installation prerequisites and requirements, and continue through to Chapter 5 to complete the billing system installation and configuration.
- If you are already familiar with billing system components and terminology, you can proceed directly to Chapter 3 for installation prerequisites and requirements, and continue through Chapter 5 to complete the billing system installation and configuration.

1

Overview

SMDS Billing System

Ascend's SMDS billing system software for B-STDX switches provides data usage measurements within a switching system, enabling you to easily compile intra-network SMDS usage statistics on every logical SMDS E.164 address that traverses a customer interface. The resulting usage statistics can then be transferred at your leisure, via FTP, to your own billing processing software for use in customer billing, invoice processing, network usage evaluation, and audit reconciliation.

The Ascend billing system provides high reliability, in that switching system failures which can affect the completeness of usage measurement statistics are minimized. The billing system also provides high accuracy; over-billing is prevented at all times, and under-billing of usage-based services is prevented under as many circumstances as possible. In addition, the billing system has been designed in such a way that it has a minimal impact on the performance of other components of the switching system (such as the packet processing rate).

The Ascend billing system software consists of the following components:

- Billing-enhanced B-STDX switching system software, which consists of a Usage Data Generator and Data Aggregator.
- Adjunct Processor billing system software, which consists of a Data Record Formatter and a Data Record Transmission function.
- Billing-enhanced CascadeView/UX Network Management Station (NMS) software, which provides billing system management and control functionality via the NMS console.

The Ascend billing system hardware consists of the following:

- One or more networked B-STDX switches
- One or more Sun SPARCstations to serve as the Adjunct Processor(s)
- At least one Sun SPARCstation to serve as the Network Management Station

Billing System Architecture

This section provides the general architecture of billing system software as defined in Bellcore TR775[1]. As Figure 1-1 shows, the billing system is composed of a switching system, management and control system, and RAO processing. The switching system consists of switch hardware and one or more Adjunct Processors (APs). These components are responsible for generating, compiling, formatting, and transmitting usage data. The management and control system refers to your Network Management Station (NMS). This component enables you to configure and manage the billing system. RAO (Revenue Accounting Office) Processing is internal to your own Billing Center. This component uses the compiled and formatted usage data for customer billing, invoice processing, usage measurement studies, and audit reconciliation.

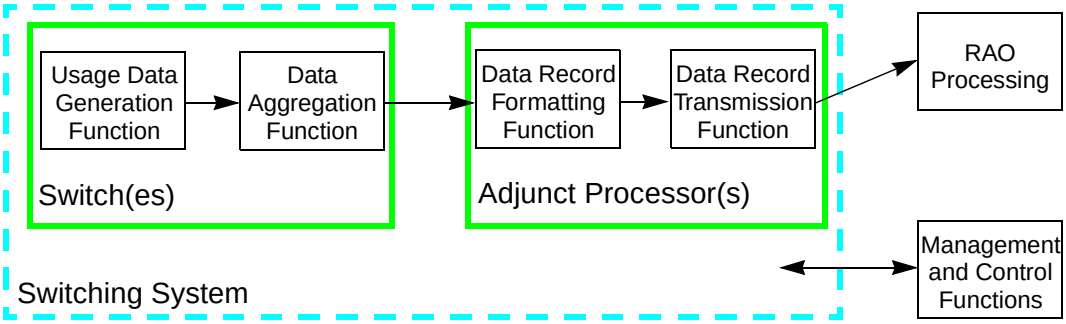

Figure 1-1. General Billing System Architecture

Figure 1-2 shows an example of a billing system hardware configuration. Please note that each switch in the system can be associated with only one AP, but one AP may service more than one switch. In the following illustration, Switch A and Switch B are connected to Adjunct Processor 1, while Switch C is connected to Adjunct Processor 2. All APs in the switching system are then connected to your RAO Processing system at your Billing Center.

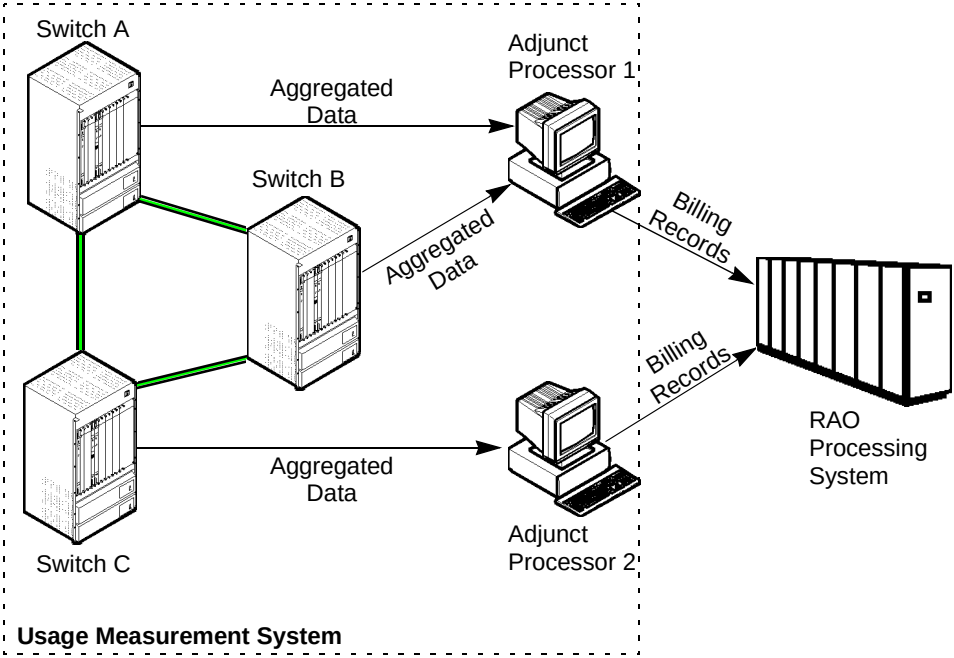

Figure 1-2. Sample Billing System Hardware Configuration

The functions performed by each component in the billing system are described in the following sections.

Switch Functions

The Usage Data Generator and Data Aggregator are both performed by the switch firmware. Within each switch in the switching system, the Usage Data Generator measures billable SMDS data traffic being switched through the network. Billable SMDS traffic is measured at egress DXI/SNI logical ports, per E.164 address subscribed to the DXI/SNI.

Since one Usage Data Generator is needed on each network I/O card on which SMDS usage statistics are to be tracked, more than one Usage Data Generator may exist on a given switch. The data produced by the Usage Data Generator is then collected by the Data Aggregator, which stores the collected data in stable storage until the data can be reliably transferred to the Adjunct Processor.

The length of time that the Data Aggregator collects data before sending it to the AP is known as the aggregation period. The aggregation period is user-defined via CascadeView/UX. For example, you may configure the billing system to send aggregated data to the AP every 30 minutes, which means you have an aggregation period of 30 minutes.

While data is being transferred to the AP, the Data Aggregator begins collecting data for the next aggregation period from the Usage Data Generators.

Adjunct Processor Functions

The Data Record Formatter and the Data Record Transmission function are both performed by one or more APs. The Data Record Formatter periodically collects the aggregated usage data from the switch, and formats the data into billing records that your RAO Processing system understands. These billing records are then stored on hard disk.

When the RAO Processing system requests data from the AP, the Data Record Transmission function transfers the billing records to the RAO Processing system. The Data Transmission function is managed by the standard UNIX File Transfer Protocol (FTP) server on the AP. (The FTP software is the standard FTP server that is bundled with the Sun OS).

RAO Processing System Functions

The RAO Processing system is responsible for initiating the billing record transfer from the AP to the RAO Processing system. Once the records have been transferred, RAO Processing performs its billing functions (e.g., produces customer invoices or performs usage measurement studies) as you have defined them.

2

Ascend Billing System

This chapter describes the billing system architecture developed by Ascend to support usage measurement on Ascend's B-STDX switches. It provides detailed information on the four billing system functions: Usage Data Generation, Data Aggregation, Data Record Formatting, and Data Transmission.

Billing System Functions

This section describes the functions that are performed by the switches and AP(s) in the billing system, and how these functions interact with each other to provide reliable and complete usage information.

Figure 2-1 illustrates how the Usage Data Generators and Data Aggregator interact within a B-STDX switch.

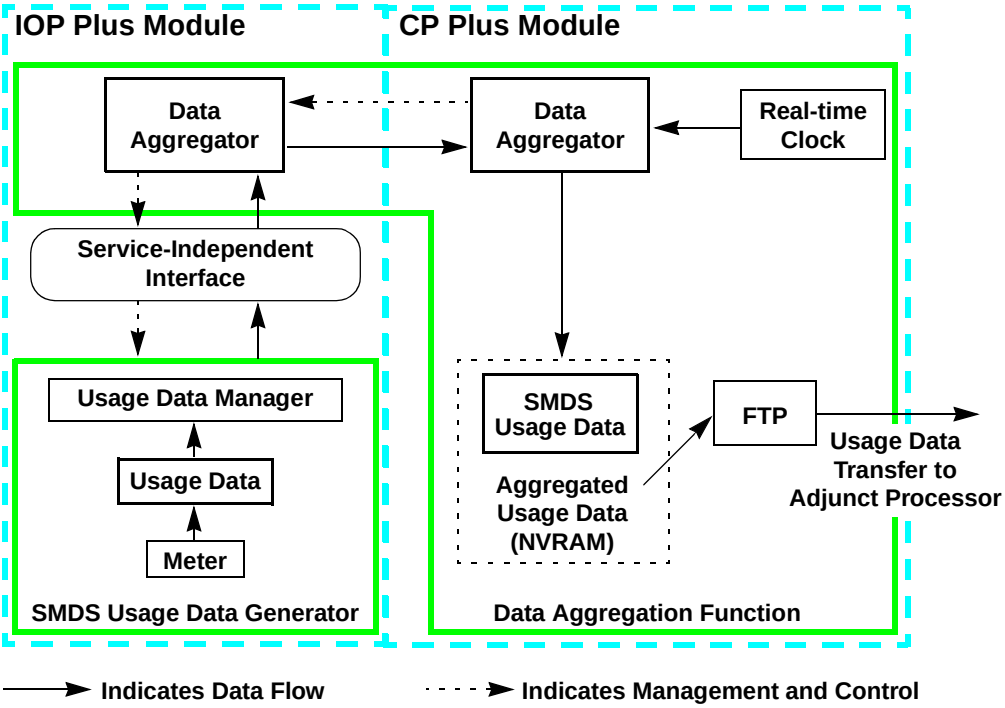

Figure 2-1. B-STDX Billing System Components

Figure 2-2 illustrates the internal processing of the Data Record Formatter and Data Record Transmission function on the AP.

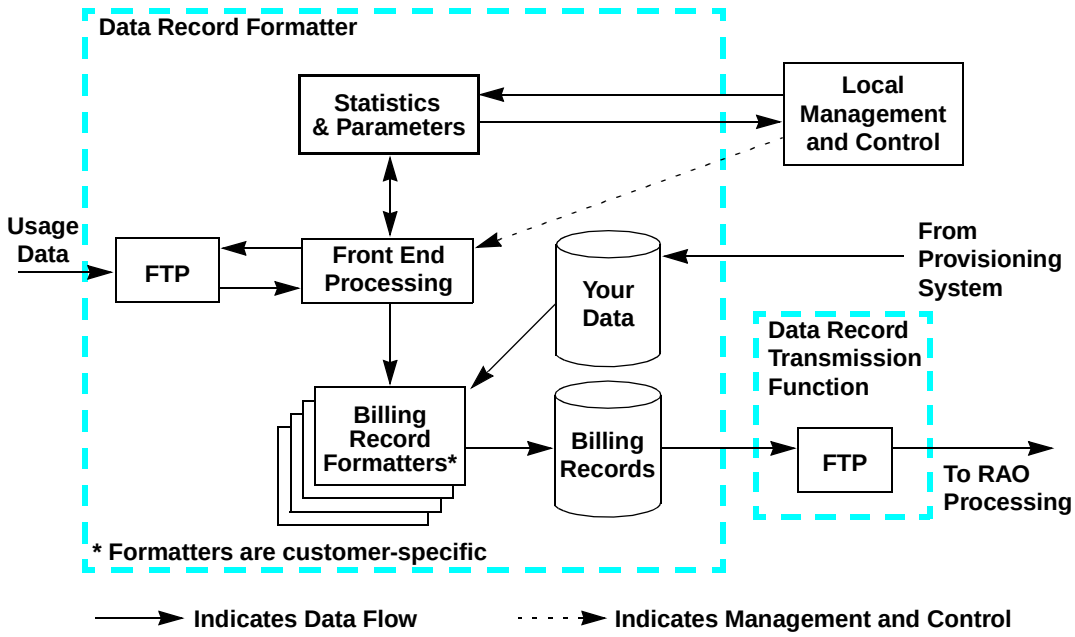

Figure 2-2. Adjunct Processor Billing System Components

For detailed descriptions of the functions shown in Figures 2-1 and 2-2, refer to the following sections.

Usage Data Generator

Ascend's Usage Data Generator is performed on each IOP Plus module in a B-STDX switch, as shown in Figure 2-1. This function measures and records all billable SMDS traffic that is processed by the switch. Traffic measurements are taken on every logical address (SMDS E.164 address) that traverses a customer access interface.

The Usage Data Generator consists of

- A Usage Data store, which stores the usage data that is being collected.
- A Meter, which is responsible for performing usage measurement.
- A Usage Data Manager, which interfaces with the Data Aggregator function.

When a Usage Data Generator begins executing (typically when the IOP Plus module is initialized, or billing is initially enabled on the switch), the Usage Data store does not contain any usage data. As SMDS L3 PDUs are received and processed by an IOP Plus module, the Meter determines whether or not the data unit needs billing treatment. If it does, the Meter makes the appropriate usage measurement and updates the Usage Data store.

The Usage Data Generator transfers data from the Usage Data store to the Data Aggregator under the following conditions:

- When a request to transfer data is made by the Data Aggregator (this occurs periodically during an aggregation period)
- When the Usage Data store is nearing its storage capacity
- When, after updating a record in the Usage Data store, a value in the record has reached a threshold and is in danger of overflowing

Two Usage Data stores are used. While data in the primary store is being transferred to the Data Aggregator, the Meter continues to update the secondary store. This prevents usage data from being lost while data is being transferred to the Data Aggregator.

Usage data maintained by the Usage Data Generator is not maintained within stable storage. Also, usage data is not mirrored on the standby IOP Plus module in a redundant configuration. As a result, when an active IOP Plus module fails, any usage data that has not yet been recorded by the Data Aggregator is lost.

Service-Independent Interface

The service-independent interface provides a consistent programming interface between the Data Aggregator and the Usage Data Generator. The Data Aggregator uses this interface to communicate with the SMDS Usage Data Generator. The interface is designed to be service-independent, which allows for future support of ATM and Frame Relay Usage Data Generators.

Data Aggregation Function

The Data Aggregation function stores usage data generated by all SMDS services that are being provided by the switch. This function is distributed between the IOP Plus modules and the CP Plus module (as shown in Figure 2-1), and consists of the following components:

- The Data Aggregator, which is responsible for collecting usage data from the Usage Data Generators, then storing the data in stable NVRAM storage until it can be transferred to the AP
- FTP, which provides reliable transfer of raw usage data from the switch to the AP
- A Real-Time Clock, which provides the Data Aggregator with the current date and time (this information is used to schedule and record aggregation period start and elapsed times)
- Aggregated Usage Data, which is maintained in NVRAM by the Data Aggregator

Real-Time Clock

Ascend's billing system software uses the CP module's real-time clock to schedule all activities at specific times of the day. When a switch is defined in CascadeView/UX, the NMS reads the value of its system clock, and sets the clock on the switch's CP Plus module to that value. Then, to ensure that all switches connected to an NMS are set to the same time, the NMS sets the time on all defined switches once each day (at 00:00:00 local time).

The time sent to each switch by the NMS is in Universal Coordinated Time (UCT), also known as Greenwich Mean Time (GMT).

Aggregated Usage Data Records

Aggregated Usage Data records can be either open or closed. Open records can be updated with new usage data. Closed records cannot be updated with new usage data.

In addition to usage data and record keys, each Usage Data record contains two timestamps, which indicate the day and time the record was opened (created) and closed. When a new record is opened, the timestamp of the open is recorded within the record. When the Data Aggregator closes a record, the elapsed time (time delta) between the open and close of the record is recorded.

Usage data records are closed under the following conditions:

- At the end of the aggregation period
- When a counter value reaches a threshold and is in danger of overflowing. In this case, the record is closed and a new record is opened. The close time of the first record and the start time of the new record are set to the current time. Also, a trap event message is generated to indicate that an overflow condition occurred (refer to "SNMP Traps" on page 8-1 for more information on trap events that are generated by the switches in the billing system).

Operational Characteristics

The following sections define the operational characteristics that the Data Aggregator maintains in order to execute the usage measurement process. These characteristics are set in CascadeView/UX, as described in "Switch and Port Configuration" on page 5-13.

Aggregation Period

The aggregation period is the length of the usage measurement periods. When an aggregation period ends, Aggregated Usage Data is transferred to the AP. The next aggregation period then starts within one second. As a result, all switches connected to a given AP must be synchronized such that they execute new aggregation periods within one second of each other. As stated earlier, this synchronization is performed by the NMS once each day.

If a switch fails during an aggregation period (causing usage data to be lost), and the switch subsequently restarts and begins generating usage data, the following records are created:

- Billing records for the period between the beginning of the aggregation period and the time of the switch failure
- Billing records for the period between the switch restart and the end of the aggregation period

If a switch stops processing data traffic, but does not lose any usage data, the current aggregation period is unaffected. In this case, only one set of billing records is created.

Collection Period

The collection period determines how long usage data is collected and stored in the Usage Data registers before being sent to the Aggregated Usage Data store. The recommended collection period is 30 seconds (which is the default setting). At the end of a collection period, all usage data in the Usage Data registers is sent to the Aggregated Usage Data store.

Billing

This is an Enabled/Disabled attribute which indicates whether or not the billing system's Usage Data Generators are enabled. This attribute can be set on both a switch-wide and customer interface (logical port) level. That is, you can enable billing on some logical ports on a switch, while having billing disabled on other logical ports on the same switch. However, no usage statistics are gathered on any logical ports until billing is enabled on the switch (even if individual logical ports are already enabled for billing). Conversely, if you enable billing on a switch, no usage statistics are gathered on a logical port until the logical port is also enabled for billing.

Daily Audit Processing

This is an Enabled/Disabled attribute which indicates whether or not daily audit statistics are to be gathered on the switch. This attribute applies only on a switch-wide basis; it does not apply to individual logical ports (that is, it cannot be set on individual logical ports). Daily audit statistics include the number of usage records created and the number of records sent to the AP.

Daily audit statistics can be used for audit reconciliation purposes.

Daily Audit Processing Time

Daily Audit Processing Time (DPT) refers to the time of day at which daily audit statistics are to be transferred to the AP. Daily audit statistics are uploaded to the AP once a day. The default setting is 00:00:00 UCT.

Adjunct Processor Network Address

The Adjunct Processor Network Address is the IP address of the AP assigned to service the switch (that is, collect aggregated usage data from the switch).

Login Credentials

The Login Credentials are the {username,password} combination which FTP uses to access the Adjunct Processor that is assigned to service the switch.

Restart Billing

The Restart Billing attribute is defined on both a switch-wide and customer interface (logical port) level.

Setting the switch-wide value to OFF terminates the current aggregation period and causes an immediate transfer of all usage records to the Adjunct Processor.

Setting a customer interface value to OFF causes the Data Aggregator to close the usage record for that customer interface.

Setting the switch-wide value from OFF to ON causes a new aggregation period to begin on the switch.

Setting a customer interface value from OFF to ON causes a new usage record to be opened for that customer interface.

Operation

To collect Usage Data, the Data Aggregator periodically polls each IOP Plus module's Usage Data Generator for the current set of billable SMDS usage statistics (the frequency with which polling occurs is determined by the collection period setting). Upon receiving the data, the Data Aggregator updates the Aggregated Usage Data store. For each usage data item received from the Usage Data Generator, the following rules are used to determine how the Usage Data store should be updated:

- If an open Usage Data record does not exist for the new data, a new record is opened and the usage data is recorded.
- If a Usage Data record exists and is open, the record is updated with the new usage data. For example, if the usage data consists of a counter, the usage data item counter is added to the counter value in the Usage Data record, and the new value is stored in the Usage Data record.
- If a Usage Data record exists, but is closed, a new record is opened and the usage data is recorded in the new record.

After a Usage Data record has been updated, if any usage data value has reached a threshold (and is in danger of overflowing), the Data Aggregator closes the record. If more usage data arrives for this record before the end of the current aggregation period, the Data Aggregator opens a new record, and stores the new information in that record.

For example, if the aggregation period is 30 minutes, an SMDS SNI that is receiving a large amount of traffic may produce two sets of Usage Data records. The first record may track usage measurement for minutes 0 through 18 of the aggregation period, while the second record would track usage measurement for minutes 19 through 30. The first record would be closed after 18 minutes, while the second record would be closed at the end of the aggregation period (after 30 minutes).

Aggregated Usage Data Store

The Aggregated Usage Data store is maintained in 2 MB of battery-backed NVRAM on the CP Plus module. This memory is separate from the CP Plus module's PRAM, and is used primarily for the billing system. The NVRAM for the Aggregated Usage Data store is separated into two areas. The primary storage area stores aggregated usage data for the current aggregation period, while the secondary storage area is used to store usage data while the FTP process is transferring data from the first area to the AP. This feature guards against the loss of usage data when updates to the usage data files are being made. Once the data in the primary storage area has been transferred, data in the secondary storage area is transferred to the primary storage area.

The Data Aggregator also maintains a hash table structure in volatile memory. This structure is used to quickly and efficiently access the Aggregated Usage Data store. When a CP Plus module reboots (due to a reset request from the NMS or a crash) in a non-redundant configuration, the Aggregated Usage Data store still exists in NVRAM, but the hash table structure is destroyed. During the system software boot, the Data Aggregator disables the billing system until after it has rebuilt the hash table structure. While the hash table is being rebuilt, the Usage Data Generator is disabled, and usage data is not being collected (this condition exists for no more than two seconds).

Recovery and redundancy of data in a redundant CP Plus module configuration is currently not supported.

Usage Data Transfers

The Data Aggregator transfers the aggregated usage data to the AP under the following conditions:

- At the end of the defined aggregation period. In this case, the Data Aggregator transfers the aggregated usage data to the AP, then begins a new aggregation period. The Data Aggregator first polls the Usage Data Generators for the latest statistics, updates the Usage Data store, then closes all records. It then signals the file transfer process to transfer the usage data to the AP.

  Usage data is transferred to the AP either via an SMDS in-band management path (which is configured from CascadeView/UX), or via a direct or indirect Ethernet connection.

Once the transfer is completed, the Data Aggregator clears the Aggregated Usage Data store and begins a new aggregation period, setting the starting time of the new aggregation period to the time at which the previous Usage Data records were closed. The scheduled end of the period is then set to the start time plus the defined length of the aggregation period. For example, if the new aggregation period starts at 10:00:00, and the defined length of the aggregation period is 30 minutes, the end of the aggregation period is scheduled for 10:30:00.

- When the Aggregated Usage Data store is near capacity. In this case, the Data Aggregator immediately closes all records and signals the FTP process to transfer the Usage Data to the AP.

  Once the transfer is completed, the Data Aggregator clears the Aggregated Usage Data store and begins a new aggregation period, setting the starting time of the new aggregation period to the time at which the previous Usage Data records were closed. However, *the scheduled end of the aggregation period does not change*. For example, if the aggregation period is 30 minutes, and is scheduled to end at 10:30:00, but data had to be transferred after 22 minutes due to a threshold condition, the new aggregation period will still be scheduled to end at 10:30:00.

If the Adjunct Processor is not Available

If communication with the AP cannot be established at the end of an aggregation period, the Data Aggregator will periodically attempt to communicate with the AP. In the meantime, the primary Usage Data store is closed, and usage data for the new aggregation period is stored in the secondary Usage Data store to keep it separate from usage data collected during the previous aggregation period.

The Data Aggregator continues adding new records to the secondary Usage Data store until the secondary store reaches capacity. Once the capacity of the secondary store has been reached, the Data Aggregator continues to update only the open Usage Data records contained within the store. This process continues until communication with the AP is established and the closed records can be transferred.

When communication is established with the AP, only the Usage Data records from the closed primary Usage Data store are transferred to the AP for billing record processing. Once these records have been transferred, the usage data in the secondary store is then moved to the primary store.

An SNMP trap event is generated by a B-STDX switch whenever the Data Aggregator cannot communicate with the AP. Similarly, an SNMP trap event is generated when the capacity of the Aggregated Usage Data store is reached. For more information on SNMP trap events for Ascend's billing system, refer to *"SNMP Traps" on page 8-1**.*

In addition to the SNMP trap events generated by switches, the AP generates event messages to inform you of normal and anomalous billing system events. These event messages are stored in the file ***/opt/CascadeBS/cascade.log****. For more information on* ***cascade.log*** *and event logging, refer to* *"Event Log Messages" on page 7-3**.*

Data Record Formatting and Transmission

The Data Record Formatter and Data Record Transmission functions are performed by the AP, as shown in Figure 2-2. In an Ascend billing system setup, the AP is a Sun SPARCstation running Solaris 2.3. For specific hardware and software requirements for the AP, refer to Chapter 3.

Any number of APs can be configured on the billing system. Any given switch must send its data to only one of the APs. However, a given AP can service multiple switches. As a rule of thumb, Ascend recommends a given AP should service no more than 10 switches.

Since the AP is a transaction processing system that uses network-, memory-, and filesystem-intensive applications, Ascend recommends that a dedicated workstation be used, and that no additional applications be resident on the system.

The Data Record Formatter is composed of the following components:

- A Front End Processor function, which is responsible for managing usage data transfers from the switch via a reliable TCP-based file transfer protocol.
- FTP, which performs the reliable transfer of the Aggregated Usage Data from the B-STDX switches in the switching system.
- A billing record formatter, which is invoked whenever usage data has been received from a switch. The billing record formatter takes this raw data, possibly combines it with customer configuration information supplied by your provisioning system, and converts it to your billing record format. The billing records are then stored in a disk file for later retrieval by your RAO Processing system at your Billing Center.

  Depending upon your billing requirements, the billing record formatter may also format statistics maintained by the AP into billing records. For example, you may need a set of audit reconciliation statistics files that summarize the day's billing operation.

The Data Record Transmission function performs reliable transfer of billing record files to your RAO Processing system. FTP is the only component of the Data Record Transmission function.

Operation

Billing records are created and stored on disk whenever usage data is received from the Data Aggregator. This section describes the billing record creation process.

First, Aggregated Usage Data is transferred to the AP via a reliable TCP-based file transfer protocol, and is buffered in a temporary file. Once the usage data has been transferred, the Front End Processor invokes the appropriate Billing Record Formatter, which then creates the billing records. The Front End Processor then waits for the next transfer of data.

The Data Aggregator on each B-STDX switch drives the scheduling of the aggregation period. The AP is a passive object in the system; it simply waits for usage data from the Data Aggregator, then creates billing records.

Billing record files are created using the naming convention S*XXX*.S*NNN*, where *XXX* is the Switch ID assigned by your Provisioning System to the switch that generated the data, and *NNN* is a value from 000-999 (this value is assigned sequentially). The Switch ID may be a number, switch name, or network address, depending on the setup of your Provisioning System. Billing record files begin with a billing file header, followed by a number of billing records, and terminated by a billing file trailer.

For more information on the format of billing records, see "Displaying Billing Record Files" on page 6-6.

Error Suspense Files

In some situations, the billing record formatter may be unable to convert some Usage Data records into the appropriate format. When this happens, the record is stored in an Error Suspense file (in the directory **/opt/CascadeBS/esf**) until it can be processed by the billing record formatter. Periodically, a UNIX cron routine re-executes the billing record formatter on these Error Suspense files to process the records that are stored in them. As a result, no usage data is lost. If you do not want to wait for the cron routine to execute, you can manually convert these files to billing record files, as described in "Manually Converting Error Suspense Files" on page 6-16.

The billing record formatter may be unable to process a record for any of the following reasons:

- The customer's network address within the record was not found in the Customer Database. In this case, additional information that is provided by the Customer Database is missing, and the billing record cannot be created. Once the Customer Database is updated to include the customer's network address, the record can then be processed by the billing record formatter.
- The Customer Database cannot be accessed (due to the database server or network being down). In this case, when the server or network is up again, the billing records in the Error Suspense file(s) will be processed by the billing record formatter.

*All conditions that result in Usage Data records being stored in an Error Suspense file are logged in the Event Log file, **cascade.log**, thereby enabling you to easily determine why an Error Suspense file has been created.*

The contents of a given Error Suspense file can be displayed using the **udfdump -u** command (as described in "Displaying the Error Suspense Files" on page 6-6).

3

Preparing for the Installation

This chapter describes the hardware and software requirements for the Ascend billing system.

Overview

When installing the billing system on your CascadeView switching system, there are two approaches you can take. The first, and recommended, approach is to have a dedicated SYBASE server for the NMS and a dedicated SYBASE server for the AP(s) in the switching system, as shown in Figure 3-1. The second approach is to have an NMS and all APs share one SYBASE server, as shown in Figure 3-2.

*In either case, the SYBASE filesystem (e.g, **/opt/sybase**) must be nfs mounted on each AP.*

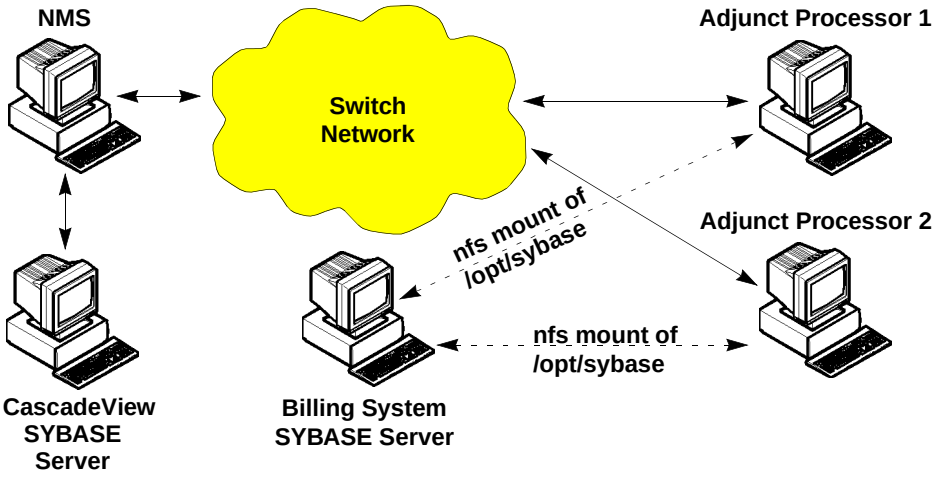

Figure 3-1. Using Dedicated SYBASE Servers

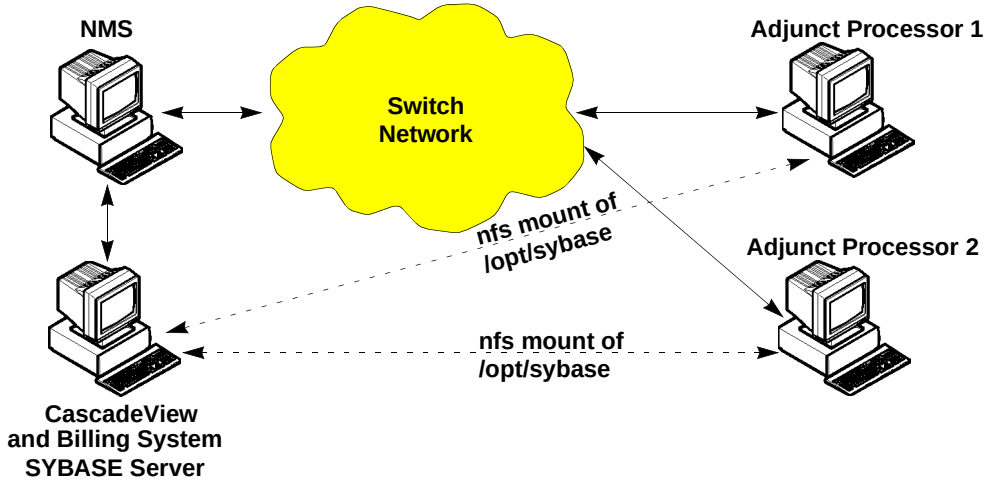

Figure 3-2. Using Combined SYBASE Server

Installation Considerations

Since the AP software is disk-intensive, special considerations have to be made when allocating disk space for usage data file and billing record file storage. Billing record files include both raw usage record files uploaded from each switch connected to the AP, as well as converted billing records. This section describes how to determine the disk space requirements for the AP, based on your network design and usage.

All AP software is installed to the **/opt** filesystem in the **/CascadeBS** directory. Also, all Usage Data files are stored in the **/opt** filesystem. Billing record files are stored in a filesystem/directory of your choice.

Before installing the Solaris operating system on the AP, you need to determine how much space has to be allocated to the **/opt** filesystem, and how much space has to be allocated to the filesystem in which converted billing record files will be stored.

Determining Disk Space Requirements

The amount of disk space you need to reserve for your partitions depends on your network design, the number of SMDS customers on the network, and the amount of traffic on the network. This section presents some background information on disk space usage, followed by a simple example to show how to calculate the required space.

When computing space requirements, you need to determine how much space is needed for raw usage data files, and how much space is needed for converted billing record files. These files will be stored in directories that you set up prior to installing SYBASE (see "Configuring the Billing System Directories" on page 4-11).

In addition to the space needed for the usage file and billing file directories, if you are installing the SYBASE server on an AP, you will need 100 MB of disk space for the SYBASE database.

Audit reconciliation (daily statistics) files take up a negligible amount of space, so they do not have to be considered.

Each switch that sends data to the AP will generate one usage data file at the end of each aggregation period. If the AP is servicing 10 switches, 10 usage data files will be generated for each aggregation period. As shown on page 5-2, these files are stored in the **/opt/CascadeBS/udfiles** directory and the **/opt/CascadBS/esf** directory.

Once the usage data files have been transferred to the AP, they are converted to billing record files. One billing record file is generated for each usage data file that was transferred to the AP during a given aggregation period. Billing record files are stored in a directory of your choice (you select this directory during the software installation).

After the conversion process, the usage data files are not deleted automatically. This must be taken into consideration when computing disk space requirements.

Space Requirements per Aggregation Period

The minimum size of a usage data file is 65538 bytes (this is the file size if the file contains no SMDS usage records). Each SMDS usage record is 24 to 36 bytes in size, depending on the record type; as a norm, an SMDS usage record is 24 bytes in size. Therefore, if a usage data file for an aggregation period contains 1500 records with an average size of 24 bytes, the file will be approximately 100 KB in size (65 KB + 36 KB).

After usage data files are transferred to the AP, the billing record formatter converts them to the appropriate billing record format, and stores the converted records in a billing record file.

Billing record files contain one header record and one trailer record. In addition, they contain one billing record for each usage record that was converted. Each record, including the header and trailer, occupies 66 bytes. Therefore, if a Usage Data file for an aggregation period contains 1500 records, the related billing record file will use about 100 KB of disk space (1502 records x 66 bytes).

Example

Suppose your network is configured as follows:

- Ten B-STDX switches connected to the AP
- 500 SMDS customer interfaces per switch
- An average of three billable events per SMDS customer interface during a given aggregation period

To determine your specific disk storage requirements for the Usage Data files generated during an aggregation period, use the following algorithm:

1. Determine the total number of billable SMDS addresses on all switches being serviced by the AP. For example, if you have 10 switches, each with 500 SMDS customer interfaces, the total number of billable SMDS addresses is 5,000.

2. Estimate the average number of other SMDS interfaces that will communicate with each billable SMDS address during an aggregation period. For example, if one billable SMDS address receives communications from an average of 4 other SMDS interfaces, and a second billable SMDS address receives from an average of 2 other SMDS interfaces during an aggregation period, this value would be 3.

3. Multiply the first value (for example, 5,000) by the second value (for example, 3). This gives you the total number of SMDS usage records that will be stored in the Usage Data file for the aggregation period. In this example, it would be 15,000 records.

4. Multiply the number of SMDS usage records by 24 bytes. For example, 15,000 records multiplied by 24 bytes is approximately 360 KB. Based on the sample network design defined above, this is the amount of storage space required for Usage Data files generated during the aggregation period.

To determine your specific storage requirements for the billing record files created during an aggregation period, use the following algorithm:

1. Determine the number of Usage Data records that will be converted to billing records, and add 2 to this number. For example, if you have 15,000 usage records, 15,002 billing records will be created (this includes the header and trailer records).

2. Multiply the resulting number by 66 bytes (the size of each billing record). For example, 15,002 x 66 bytes = 1 MB (approximately).

Therefore, in this example, you need 360 KB of storage space for Usage Data files and 1 MB of storage space for billing record files during each aggregation period.

Space Requirements per Day

To determine how much storage space is required each day, multiply the number of aggregation periods per day by the amount of space required for each aggregation period.

Following through with the example in the previous section, if the aggregation period is 15 minutes, then you have 96 aggregation periods per day. Each aggregation period requires 360 KB of storage space for Usage Data files, and 1 MB of storage space for converted billing record files. Therefore, you would need approximately 35 MB of space for Usage Data files, and approximately 96 MB of space for billing record files.

Space Requirements per Billing Cycle

To determine how much storage space you need for each billing cycle, multiply the amount of space required each day by the number of days in the billing cycle. Also, it is a good idea to increase this value by 25% to provide a safety zone.

Following through with the previous example, if your billing cycle is 30 days, your storage requirements would be about 1.1 GB for usage data files, and about 2.9 GB for billing record files. To provide for a safety zone, increase these values by 25%, to approximately 1.5 GB and 3.6 GB respectively.

Since usage data files are stored in the /opt filesystem, you need to reserve the appropriate amount of space for **/opt** when setting up your disk partitions. Since billing record files are stored in a directory of your choice, you have to decide where to store them before beginning the software installation, then reserve the needed space in the appropriate partition. For example, if you want to store billing record files in a directory called **/billing**, make sure you reserve enough space for **/billing** when setting up your disk partitions.

Hardware Requirements

This section describes the minimum configuration for the B-STDX switches in the billing system. It also describes the minimum hardware configuration for the AP. For NMS hardware requirements, see the *Ascend Network Administrator's Guide for CascadeView/UX*.

B-STDX Switch Requirements

The Ascend Billing System is supported on B-STDX 8000 and 9000 model switches that are configured with CP Plus and IOP Plus modules.

At least one CP Plus module must be installed in the B-STDX switch (for a redundant configuration, two CP Plus modules must be installed). Also, IOP Plus modules must be used in the switch for each customer interface on which billing functionality is needed. Billing system functionality is not supported on older CP Basic and IOP modules, since the older CP Basic modules do not contain the NVRAM that is needed for the Usage Data store, and older IOP modules do not support SMDS functions.

Adjunct Processor Requirements

The hardware requirements for the AP depend on your current billing system needs, and your plans for future expansion of the billing system.

At a minimum, the AP must be a Sun SPARCstation 10 with 96 MB of memory. The amount of disk space (and consequently, number of hard drives) required depends on your network design, as described in "Determining Disk Space Requirements" on page 3-3.

To support future expandability on a growing network, a low-end Sun SPARCCenter 1000 is recommended. This Sun model can be configured with up to eight CPUs (on two CPU cards) and 512 MB of memory on each CPU card. Also, this model supports up to 60 GB of disk space on each external disk adapter.

In addition, a 17-in. or larger color monitor must be installed on the AP. Also, in order to install the AP software, a 1/4-in. cartridge (QIC) tape drive must be accessible from the AP.

Software Requirements

This section lists the software requirements for the billing system.

B-STDX Switch Requirements

No additional software is needed for the switches in the billing system, since the billing system functions are integrated into the switch firmware on all CP Plus and IOP Plus modules. For information on the firmware release you should be using, refer to the *Software Release Notice for the Ascend Billing System Adjunct Processor.*

NMS Requirements

The NMS workstation being used for billing system management must have CascadeView/UX Release 1.5 or greater installed in order to support billing system functions. For additional NMS software requirements, refer to the *Ascend Network Administrator's Guide for CascadeView/UX*.

Adjunct Processor Requirements

The following *minimum* software has to be installed on each AP:

- Solaris Operating System, release 2.3
- Solaris Maintenance Release 1 for Solaris 2.3
- Disk spanning/disk management software of your choice (only if you need more than 2 GB of storage space for usage data and billing record files)
- Ascend Billing System software, release 1.0

Since each AP may require several GBs of drive space for storing usage data files and converted billing record files, you also have to install disk spanning/disk management software of your choice.

For complete installation instructions, refer to Chapter 4.

4

Installing the Billing System

This chapter describes how to install and configure the hardware for the billing system. It also describes how to install the operating system and billing system software on the APs in your network, as well as how to install the SYBASE software on a dedicated billing system SYBASE server.

Hardware Installation

This section describes how to install and configure the hardware components that will be used for the billing system, including the NMS, B-STDX switches, and Sun SPARCstation(s) you will be using as your Adjunct Processor(s). All billing system hardware must be installed and configured on the network before you begin the software installation on the APs.

NMS Installation and Configuration

The NMS workstation being used for billing system management must be installed as described in the *B-STDX 8000/9000 Hardware Installation Guide*. No additional hardware installation steps are needed to support the billing system.

To install the NMS software and configure the NMS on the network, refer to the *Ascend Network Administrator's Guide for CascadeView/UX*. The NMS software must be installed and configured before you begin the AP software installation.

B-STDX Switch Installation and Configuration

The B-STDX switches in the billing system must be installed and connected to the network as described in the *B-STDX 8000/9000 Hardware Installation Guide*. No additional hardware installation steps are needed to support the billing system.

If you have not already configured the B-STDX switches on your network, refer to the *Ascend Network Administrator's Guide for CascadeView/UX* for instructions on how to do so. For instructions on configuring the switch for billing system functionality via CascadeView/UX, see "Billing System Configuration" on page 5-1.

Do not configure your switches and logical ports for billing until all software has been installed on the APs and SYBASE server.

Adjunct Processor Hardware Installation

For instructions on installing the Sun workstation(s) you are using as the AP(s), refer to the hardware installation guide you received with the workstation.

Each AP must be connected to the network on which the switches it services reside. Currently, you can connect the APs to the switch network using a direct or indirect Ethernet connection, or via an SMDS in-band management port on any switch (preferably, the most local switch) in that network.

In addition, in order to install the billing system software, a 1/4-in. cartridge (QIC) tape drive must be accessible by the AP.

Adjunct Processor Software Installation

There are two methods that you can use to install the required CD-ROM software on your APs:

- Remotely, using a Solaris bootp server
- Locally, using an attached CD-ROM drive

Remote Installation Prerequisite

If you choose to remotely install the AP software across the network, you must use the Solaris Administration tool to properly configure the bootp protocol. To define the bootp configuration, follow the instructions in the Solaris 2.3 System and Network Administration Guide*.*

Depending upon the installation method you choose for installing the software, the commands you have to enter may vary slightly. Whenever a command is specific to a remote or local installation, it is identified in the following manner:

Remote Example

Remote installations, type

```
/cdrom/install -d /cdrom/openview <Return>
```

Local Example

Local installations, type

```
/cdrom/cdrom0/install -d /cdrom/cdrom0/openview <Return>
```

Note that UNIX commands are case sensitive. All UNIX commands must be entered exactly as they appear.

Installation Options

As stated in Chapter 3, you have several options for installing the SYBASE server:

- You can use a dedicated SPARCstation for the billing system SYBASE server.
- You can use the CascadeView SYBASE server (and the *cascview.dat* device) to support the billing system.

Regardless of the option you choose, the following software must be installed on all APs, in the order listed:

- Solaris 2.3
- Solaris Maintenance Release 1 for Solaris 2.3
- Disk spanning/disk management software of your choice (if you need more than 2 GB of storage space for usage data files and billing record files)
- Ascend billing system software

If you choose to have a dedicated billing system SYBASE server, the following software must be installed on that workstation, in the order listed:

- Solaris 2.3
- Solaris Maintenance Release 1 for Solaris 2.3
- SYBASE 4.92

If you want to use a CascadeView SYBASE server as your billing system SYBASE server, no additional software is needed on that workstation.

*In either of these situations, you have to nfs mount the **/opt/sybase** directory on each AP in your billing system.*

Installing the Solaris 2.3 Operating System

Sun Microsystems, Inc. Solaris, Version 2.3 is the operating system software to be installed on all Sun SPARCstations to be used as APs. Although you can install this software by following the installation instructions provided in the *Solaris System and Network Administration Guide*, the following set of instructions provide Ascend-recommended settings that will enable you to successfully install and run the billing system.

Use the following set of instructions in conjunction with the instructions provided in the *Solaris 2.3 System Configuration and Installation Guide*.

1. Obtain an IP address and Subnet Mask from your network administrator. (This IP address must be registered as a valid address on your network.)

2. Power on the Sun SPARCstation and insert the CD into the drive.

3. While holding down the Stop key, press the a key. This will give you the `ok` prompt.

   Remote installations:

   – At the `ok` prompt, type **boot net** and press Return. (This command sends out an ARP request to the bootp server to retrieve the kernel, configure the device drivers and load OpenWindows for operating system installation.)

   Local installations:

   – Verify that the jumper switch located on the back of the CD-ROM drive is set to SCSI ID 6.

   – At the `ok` prompt, type **boot cdrom** and press Return. (This command initiates a local boot of the operating system off of the CD-ROM.)

     After several minutes, the following messages appear:
     "Starting OpenWindows." "The system is coming up. Please wait."

4. *Remote* installations:

   – Skip to Step 11. (Steps 5 through 10 should be completed during bootp server configuration.)

   Local installations:

   – At the "What is the hostname for your workstation?" prompt, type **[*your host name*]** and press Return. (Example, ap01.)

5. At the "Will this system be connected to a network?" prompt, type **y** and press Return.

6. At the "What is your Internet Protocol (IP) address?" prompt, type **[*the IP address you received from your network administrator*]** and press Return.

7. At the "Is the following information correct?" prompt, type **y** and press Return to accept the information as displayed.

8. At the "Do you want to configure this client as a name service? If so, which name service do you want to use?" prompt, use the arrow keys to select **None** for NIS services, then press Return.

9. At the "Does this workstation's network have subnetworks" prompt, type **y** and press Return.

10. Type **y** and press Return to continue.

11. At the "Netmask:" prompt, type **[*the subnet mask you received from your network administrator*]** and press Return.

12. At the "Is the following information correct?" prompt, type **y** and press Return to continue.

13. You are now prompted to enter your geographic region and time zone, and to confirm the date and time. The defaults given should be correct as shown. You can, however, edit the fields if you need to.

14. At the "Is the following information correct?" prompt, type **y** and press Return to continue.

15. Use the Tab and Down Arrow keys to select **Custom Install**, then press Return.

16. Select **Software Selection**, then press Return.

17. Use the Tab key to highlight **Developer System Support**, then press the Space bar to select it.

18. Use the Tab key to select **Edit**, then press Return.

19. Add the following *required* new features to the Developer System Support:

    – Select **Basic Networking**, then press Return twice.

    – Select **System Accounting** (for information about who is logging into the system), then press Return twice.

    Add the following *optional* new features to the Developer System Support:

    – Select **Point-to-Point Protocol** (for optional dial-up modem use), then press Return twice.

    – Select **Automated Security Enhancement Tools** (which provides options for securing the system), then press Return twice.

20. Tab to **Done**, then press Return.

21. Select **Disk/File Systems**, then press Return to configure the disks.

22. Press Return to select the **Configure Disk** option. Press Return again. The Disk Editing Properties dialog box appears.

23. Ascend recommends that you accept the Sun defaults as shown. (Leave the size editing units in MB and *do not* allow overlapping slices).

24. Tab to **Apply**, then press Return. The Configuring File Systems on Disk box appears.

25. Edit the default values as shown in Table 4-1. These values assume that your primary drive is 650 MB or larger, and that you will need to reserve a minimum of 1 GB of space for **/opt**, and 1.5 GB of space for the directory in which converted billing record files will be stored, which means you cannot create these filesystems on the primary drive. Because of this, these instructions assume that you will be using multiple hard drives, and therefore will be installing disk spanning/disk management software after the operating system installation.

The values in Table 4-1 can be used for filesystems or raw partitions, whichever you prefer to use.

Table 4-1. Recommended Partition Settings on Primary Drive

Partition Settings on Primary Drive (Assumes 650 MB or greater, with 96 MB RAM)	
Slice 0	/ = 30
Slice 1	/swap = (3 * RAM) (*Recommend* 288 MB)
Slice 2	DO NOT CHANGE
Slice 3	DO NOT CHANGE
Slice 4	/var = 60
Slice 5	/usr = 100
Slice 6	/usr/openwin = 150
Slice 7	/export = (remaining unallocated space on drive, after all other partitions have been allocated)

26. Tab to the **Space** command button, then press Return. The filesystem space requirements appear.

27. Review the space requirements. The configuration should be greater than the suggested minimums. Tab to **Dismiss** and press Return. The Configuring File Systems on Disk box re-appears.

28. Tab to **Apply**, then press Return.

29. Tab to **Done**, then press Return. The Custom Install Configuration screen appears.

30. Tab to **Begin Install**, then press Return. (The system creates partitions and installs the operating system files.)

31. Tab to **Continue with Install**, then press Return. (The installation download takes between 45 minutes to one hour.)

32. When the installation completes, the system automatically reboots. Upon reboot, the system configures its devices and prompts you to set your root password.

33. At the root password prompt, type

    [*your desired root password*] <Return>

    The system completes the boot up procedure and returns the console login prompt.

34. Log in as root and enter the root password. The system returns a # prompt (the default shell prompt for the root user).

35. At the root prompt, type

    eject cdrom <Return>

36. The next step is to install the Solaris Maintenance Release as described in the next section.

Installing the Solaris Maintenance Release

Install the Solaris 2.3 Maintenance Release software as follows:

1. Verify that you are logged in as the root user. You should see a # prompt.

2. *Remote* installations, type

   mkdir /cdrom <Return>

   Local installations, go to Step 4.

3. *Remote* installations only, type

   mount [*remote host name*]: /cdrom/cdrom0/cdrom <Return>

4. Type the following command to enter maintenance mode:

   init s <Return>

   This command returns a # prompt. Please wait; this procedure takes a few minutes to complete.

5. At the password prompt, type **[*your root password*]** and press Return.

6. *Remote* installations, type

   cd /cdrom/Solaris_2.3_MS1 <Return>

   ./install_ms; init 6 <Return>

   Local installations, type

   cd /cdrom/cdrom0/Solaris_2.3_MS1 <Return>

   ./install_ms; init 6 <Return>

   This procedure starts the installation of the maintenance release and automatically initializes the system. The entire procedure takes approximately 90 minutes.

You will receive messages that not all of the maintenance patches apply; this is a normal condition.

7. At the root prompt, type **eject cdrom** and press Return.

8. The next step is to install the disk spanning/disk management software you will be using to manage your drives. Once you have installed the disk management software, you can configure the directories in which usage data files and billing record files will be stored. The next section provides you with guidelines for setting up these directories after your disk management software has been installed.

Configuring the Billing System Directories

Usage data files that are generated by each switch connected to an AP are stored under the **/opt** directory. Converted billing records are stored in a directory of your choice (you select this directory when installing the Ascend billing system software). Before installing the billing system software, you have to create the **/opt** directory and the directory in which you want to store billing record files.

To configure the billing system directories on your secondary drives, you first have to determine the disk space requirements for your usage data files and converted billing record files, as described in "Determining Disk Space Requirements" on page 3-3.

Once you have determined the disk space requirements, set up the **/opt** directory on your secondary drives, reserving the appropriate amount of space, plus an additional 25% as a safety zone.

Also, set up the directory in which you will be storing your converted billing record files (for example, **/billing**). As with the **/opt** directory, you should reserve an additional 25% of space for this directory as a safety zone.

Changing Directory Permissions for the Billing Directories

Each user who needs access to the billing file and audit file directories must be assigned to the *billing* group on the AP. Also, to allow these users to access these directories, the directory permissions have to be changed, using the following procedure:

1. Log in to the AP as root, and change to the **/opt** directory.
2. From the **/opt** directory, type the following commands to change the directory permissions for your billing file and audit file directories:

```
chmod -R 775 [billing file directory path]

chgrp -R billing [billing file directory path]

chmod -R 775 [audit file directory path]

chgrp -R billing [audit file directory path]
```

For example, if your billing files are stored in /opt/home/billing, substitute **/home/billing** for [*billing file directory path*].

Installing SYBASE

This section describes how to install the SYBASE database server software on the billing system SYBASE server. These instructions apply only to the workstation that you will be using as your SYBASE server, and only if you are using one SYBASE server for your billing system and a second SYBASE server for CascadeView.

*If you are using the CascadeView SYBASE server as your billing system SYBASE server, these instructions do not apply; instead, refer to "NFS Mounting the SYBASE Database" on page 4-28 for instructions on how to nfs mount the **/opt/sybase** directory on each of your APs.*

Before installing the SYBASE SQL Server programs, you have to complete the following prerequisite tasks:

- Create a SYBASE user account
- Set the shared memory allocation
- Create a tcp socket number
- Set up the master database device
- Set the environment variables for the SYBASE user account
- Load the SYBASE tape for installation

The following sections provide detailed instructions for each of these tasks. These instructions assume that you are using the Korn shell.

It is recommended that you have a copy of the *SYBASE SQL Server Installation Guide for Sun Solaris* on hand for specific references. This guide may be ordered directly from SYBASE.

Creating a SYBASE User Account

Set up a SYBASE user account as follows:

1. Log in as root and enter the root password.
2. At the # prompt, type **vi /.profile** and press Return to set the following environment variables:

   ```
   OPENWINHOME=/usr/openwin <Return>
   LD_LIBRARY_PATH=/usr/openwin/lib:/usr/lib <Return>
   DISPLAY=:0.0 <Return>
   EDITOR=/usr/ucb/vi <Return>
   export OPENWINHOME LD_LIBRARY_PATH DISPLAY EDITOR <Return>
   ```

3. Press the Escape key.
4. While holding down the Shift key, type **zz** to save and exit the file.
5. Type **exit** and press Return to log out. Then, log back in as the root user for the new environment variables to take effect.
6. Type the following command to start OpenWindows:

   ```
   /usr/openwin/bin/openwin <Return>
   ```

7. Open a new command tool window by typing the following command:

   ```
   /usr/openwin/bin/cmdtool & <Return>
   ```

8. In the command tool window, at the system prompt, type **admintool &** and press Return to invoke the Solaris administration tool.
9. Choose **Database Manager** and press Return.
10. On the Load Database screen, do the following:
    - Click on **Group**.
    - Click on **None** for the Naming Service.
    - Choose the **Load** command button to load the group database.
11. Choose Edit ⇒ Add Entry.

12. Complete the fields as follows:
    - In the Group Name field, type **dba**.
    - In the Group ID field, type any number greater than 15.
13. Choose **Add**. The dba userid is added to the list.
14. Quit the Database Manager tool.
15. Choose **User Account Manager**.
16. In the Naming Service field, click on **None**.
17. Choose **Apply**.
18. Choose Edit ⇒ Add User. The Add User dialog box appears. Complete the USER IDENTITY fields on the User Account Manager dialog box as follows:
    - In the Create User Name field, type **sybase**.
    - In the Userid field, type any number greater than 15. (Use a different number than the one you entered in Step 12 for the Group ID field.)
    - In the Primary Group field, type **dba**.
    - In the Secondary Group field, type **[*optional entry*]**.
    - In the Comment field, type **sybase dba**.
    - In the Login Shell field, select the desired login shell.

All of the instructions in this chapter assume that you are using the Korn shell as the default login shell.

19. Optionally, complete the ACCOUNT SECURITY fields as follows:
    - In the Password field, select the level desired (Ascend recommends using **Cleared until first logon**).
    - In the Min change field, keep the default value of **0**.
    - In the Max change field, type the desired maximum number of days before forcing the user password change.
    - In the Max inactive days field, *optionally* type the maximum number of days that can pass without user activity before the user id is automatically deleted.
    - In the Expiration date field, *optionally* type an **expiration date** on which the user id should be automatically deleted.
20. Complete the HOME DIRECTORY fields as follows:
    - Set **Create home dir** to **checked**
    - Set **Path** to **/opt/[*user name*]** (for example, **/opt/sybase**)
    - Set **Server** to **[*hostname*]** (for example, **saturn5**)
    - Set **Skeleton Path** to **/etc/skel**
    - Set **AutoHome Setup** to **unchecked**
    - Set **Permissions** as desired
21. Complete the MISCELLANEOUS field: Mail Server: **[*hostname*]**
22. Choose **Add**.
23. Quit the Solaris Administration tool.
24. At the # prompt, type

    su - [*user name*] <Return>

25. Rename the three files that you just created, as follows:

```
$ mv local.cshrc .cshrc <Return>
$ mv local.login .login <Return>
$ mv local.profile .profile <Return>
```

26. At the $ prompt, type **exit** <Return> to switch back to the root user.

27. Exit OpenWindows and log out.

28. The next step is to set the shared memory allocation as described in the next section.

Setting the Shared Memory Allocation

If SYBASE SQL Server does not have enough shared memory available, it will not execute. Ascend recommends that you do the following before attempting to install the SYBASE SQL Server software.

1. Log in as root and enter the root password.

2. At the # prompt, type

```
cd /etc <Return>
```

3. Type

```
vi system <Return>
```

4. While holding down the Shift key, type **$G** to scroll to the bottom of the file.

5. Type Shift-A to append a line onto the file.

6. Type the following *required* command to increase the size of the kernel:

```
set shmsys:shminfo_shmmax=131072000 <Return>
set maxusers=72 <Return>
```

If you have more than four disk drives, set maxusers to 3/4 the amount of memory. For example, if you have 96 MB of memory, set maxusers to 72.

7. Press the Escape key.

8. While holding down the Shift key, type **zz** to save the file.
9. While still logged in as the root user, complete the instructions in the next section.

Setting the TCP Socket Number

1. While logged in as the root user, type the following command:

   ```
   vi /etc/services <Return>
   ```

2. While holding down the Shift key, press **$G** to go to the end of the file.
3. Locate a number that is greater than 1024 and is not already being used in the file. The number that you select will be used as the tcp socket number.

Make a note of the number that you selected; you will need to enter this number later on during the SYBASE installation. Enter this number when prompted for the 4- or 5-digit query port number during the SYBASE installation.

4. Add the following line next to the number that you selected:

   ```
   sybase [number that you selected]/tcp
   ```

5. Press Escape.
6. Type

   ```
   :wq! <Return>
   ```

7. Type **init 6** and press Return to reboot the system. This reboot will cause all of the changes you made in the previous two sections to take effect.

Setting Up the Master Device

A partition has two device files associated with it: the character device and the block device. All I/O to the partition is done through these devices. The SYBASE SQL Server uses the character device for asynchronous I/O, so you should place your master device on the character device.

You need to know the names of both devices to set permissions for the master device. The character and block devices are in the **/dev** directory. The names of the devices differ only in that the character device names starts with "r".

Using Raw Partitions for the Master Device

When the SQL Server is installed on a raw partition, any existing files on that partition -- including SYBASE and operating system software -- are overwritten. Failure to follow the guidelines in this section may result in irrecoverable data loss.

1. Use the **prtvtoc** command to get information about the sector size of a raw partition that you will use to set up the master device. For example, if the raw device file name is /dev/rdsk/c0t1d0s4, type the command as follows:

   /usr/sbin/prtvtoc /dev/rdsk/c0t1d0s4 <Return>

2. Make note of the size of the partition (in sectors) that you selected.

Set Permissions for the Master Device

While logged in as the root user, type the following commands to set permissions for the *master*, *cascbilling*, and *log* devices. (You will be setting up these devices later on during the SYBASE installation.)

chown -R sybase /dev/rdsk/*[device name*]* <Return>
(device name example = *c0t1d0s*)*

chgrp -R dba /dev/rdsk/*[device name*]* <Return>
(device name example = *c0t1d0s*)*

The * in the device name serves as a wildcard to change all three partitions at once (*master*, *cascbilling*, and *log*).

Using an Operating System File for the Master Device

Do not place the master device on an operating system file on a "production" SQL Server. I/O to operating system files is buffered, so your data may not be recoverable in the case of a system crash or other failure.

Avoid remote-mounted filesystems. Do not create or use devices on remote nfs- or rfs-mounted directories. Certain combinations of nfs or rfs file servers and clients may cause unpredictable operating system behavior and have adverse effects on other processes on the system.

Create the Master Device Directory

If you use an operating system file for the master device to store the SQL Server system databases, complete the following steps:

1. Create the database devices directory (for example, **/opt/database**). This directory is not created automatically for you during installation.
2. Make the directory with write permissions by "sybase" using the chown and chmod commands:

   ```
   chown sybase [directory_name] <Return>

   chmod 700 [directory_name] <Return>
   ```

3. Check permissions on the database devices directory.
4. Make note of the full directory path from the root directory (for example, **/opt/database**).
5. Choose a name for your master device. The default file name is *master.dat*. If a master device file exists, it will be overwritten during the installation; if the file does not exist, it will be created.
6. The next step is to set the environment variables for the sybase user account as described in the next section.

Setting the Environment Variables

To set the environment variables for the sybase user account:

1. Log in as “sybase” and enter the appropriate password.
2. Open a new command tool window by typing the following command:

   /usr/openwin/bin/cmdtool & <Return>

3. At the system prompt, type

   vi .profile <Return>

*If you are using C-shell instead of Korn or Bourne shell, you must edit the **.login** file and the **.cshrc** file instead of **.profile**.*

4. Add the following lines to **.profile**, immediately after the “export PATH” line:

Line breaks are indicated by <Return>. Although some lines extend onto a second line, do not press <Return> until indicated.

```
set monitor <Return>
PATH=/usr/bin:/usr/sbin:/etc:/usr/ucb:/bin:/opt/sybase/bin:
/opt/sybase/install:.<Return>
SYBASE=/opt/sybase <Return>
DSQUERY=CASCADE <Return>
OPENWINHOME=/usr/openwin <Return>
LD_LIBRARY_PATH=/opt/lib:/usr/openwin/lib:/opt/sybase/lib:
/usr/lib <Return>
EDITOR=vi <Return>
```

5. Press Escape.
6. While holding down the Shift key, type **zz** to save and exit the file.
7. In order for the environment variables that you just set to take effect, you have to exit and restart OpenWindows. Log out of OpenWindows, then log back in as “sybase” to update the environment variables.

The next step is to load SYBASE for installation as described in the next section.

Loading SYBASE for Installation Using a Ascend-Supplied Tape

The following instructions assume that you are about to load the **sybload** utility for installing the SYBASE SQL Server software from a Ascend-supplied tape. If you do not have a Ascend-supplied tape for installation, follow the instructions in Chapter 3 of the *SYBASE SQL Server Installation Guide for Sun Solaris.*

If you have a Ascend-supplied tape containing a tar file for the SYBASE SQL Server software installation, complete the following steps:

1. To ensure that you are logged in as “sybase”, type

   su - sybase <Return>

2. Insert the tape into the tape drive.

3. Type the following command at the system prompt:

   cd /opt <Return>

   tar xvf /dev/*[tape device]* <Return>

   (For example: **tar xvf /dev/rmt/0m** or **tar xvf /dev/rst4**)

Once the installation files have been copied, you can install SYBASE as described in the next section.

Installing SYBASE

SYBASE SQL Server, Version 4.9.2 is a relational database program in which all database information is stored. In addition, it provides backup and recovery of database files. To install the SYBASE SQL Server software to run with the Ascend billing system software, use the recommended settings that are listed in the next section.

It is recommended that you have a copy of the SYBASE SQL Server Installation Guide for Sun Solaris *on hand for specific references. This guide may be ordered directly from SYBASE.*

Starting the SYBASE SQL Server Installation

To install the SYBASE SQL Server software:

1. Verify that you are logged in as the sybase user.

It may be helpful to reference the "Examples" section shown on page A-5 of the SYBASE SQL Server Installation Guide for Sun Solaris *when performing Steps 2 through 7.*

2. Enter the following command to create the **sybtlicf** file:

   ```
   sybtli <Return>
   ```

   The interface tool main menu appears.

3. Type **2** and press Return at the "Main Menu: Enter selection" prompt to create an interfaces file.

4. Press Return to accept the default.

5. Type **1** and press Return at the "Main Menu: Enter selection" prompt to configure network device names.

6. Type **1** and press Return to indicate tcp.

During the test series, you may notice some fail messages. Wait for the entire test series to complete and verify that the final message indicates the test has succeeded.

7. Press Return to accept the default endpoint. The software performs a series of tests and creates the **sybtlicf** file.

8. Type **3** and press Return to return to the "Main Menu: Enter selection" prompt.

9. Type **9** and press Return to exit the interfaces tool.

10. Verify that you are still logged on as the sybase user.

11. Type

    ```
    cd /opt/sybase/install <Return>
    ```

12. Type **./sybconfig** and press Return at the system prompt to Install SQL Server.
13. At the "Select Sybase product to be installed" prompt, type **1** and press Return for SQL Server.
14. Type **1** and press Return to select the Initialize software option.
15. At the "Enter the serial number supplied with your tape" prompt, type **646923** and press Return.

Use this serial number only if you are installing from a Ascend- supplied tape. Otherwise, use the serial number provided with the tape you are using.

The message "setting file permissions" appears.

16. When the menu re-appears, type **2** and press Return to Install a new SQL Server.
17. At the "Select the type of master device" prompt,
    - For *raw disk partitions*, type **1** and press Return.
    - For *file systems*, type **2** and pressReturn.
18. At the "Enter Path for Master Device" prompt,
    - For *raw disk partitions*, type **[*the full path name of the raw disk partition*]** (for example, /dev/rdsk/c0t3d0s?)
    - For *filesystems*, type **[*the full path name for the file system directory where the master database device is to be created*]** (Ascend recommends **/opt/database**). Do not include the file name at this prompt — the next prompt asks you for the master database device file name.

Make note of the full path of the device name for recovery purposes.

19. At the “Enter the File Name” prompt, press Return to accept the default *master.dat*. You are now prompted to enter the size of the master device, whether you choose a raw disk partition or operating system file installation.

    – For *raw disk partitions*, type [***the size of the raw partition in 512-byte sectors***]. The minimum is 34,816 sectors. (Refer to the number that you noted in Step 2 of “Using Raw Partitions for the Master Device” on page 4-18.)

    – For *file systems*, type **30** and press Return.

20. At the “Enter the SQL Server name” prompt, type

    ```
    CASCADE <Return>
    ```

21. At the “Enter the network device for network1” prompt, type

    ```
    /dev/tcp <Return>
    ```

    (Please wait. This process takes a few minutes to execute.)

22. Type ***[the 4- or 5-digit query port number]*** and press Return.

> *The query port number is the same number that you set for the tcp socket number. Refer to “Setting the TCP Socket Number” on page 4-17.*

23. The message “default language us_english assumed” appears.

24. At the “Enter the default character set for this SQL Server” menu, press Return to accept the default value of 3 - iso_1 - ISO 8859-1 (Latin-1) - Western European 8-bit character set.

25. At the “Select the sort order for this SQL Server” prompt, press Return to accept the default value of 1 for binary.

26. At the “Select additional character sets” prompt, type c and press Return to continue.

27. In response to the prompt, “Do you want to proceed?”, type **y** and press Return to continue. **sybconfig** completes the installation and displays a summary of the information you have supplied, including master device information, port assignments, and language configuration.

28. When the installation completes, the installation menu reappears. Type **x** and press Return to exit from the install program. You will then be returned to your operating system prompt. The Database Server should now be running.

29. The next step is to prepare the SYBASE database for the billing system, as described in the next section.

Preparing SYBASE for the Billing System

This section describes how to prepare the dedicated billing system SYBASE server for the billing system.

1. Verify that you are logged in as “sybase”.

2. Initiate an interactive SQL session to create two devices for the billing system by typing

```
isql -U sa -P <Return>
```

3. At the 1> prompt, type

```
sp_password null, superbase <Return>
go <Return>
```

   This command changes the system administrator password for SYBASE.

4. Type the following sequence of commands to create a 100 MB *cascbilling.dat* device.

The [device name] format for a raw partition is /dev/rdsk/[dev_name]; The format for a filesystem is /opt/database/cascbilling.dat.

```
1> Disk init NAME = "cascbilling_device", <Return>
2> physname = "[device name]", <Return>
3> vdevno = 3, <Return>
4> size = 51200 <Return>
5> go <Return>
```

5. Type the following commands to create the *log.dat* device:.

*The device name format for a raw partition is **/dev/rdsk**; the device name format for a UNIX file system is **/opt/database/log.dat**.*

```
1> Disk init NAME = "log_device", <Return>
2> physname = "[device name]", <Return> (e.g., /opt/database/log.dat)
3> vdevno = 4, <Return>
4> size = 51200 <Return>
5> go <Return>
```

6. Expand the master device by typing the following command:

```
1> alter database master on master = 15 <Return>
2> go <Return>
```

This process takes a few minutes to complete.

7. Expand the **tempdb** file by typing the following commands:

```
1> alter database tempdb on log_device = 28 <Return>
2> go <Return>
```

8. Quit the Interactive SQL session by typing

```
1> quit <Return>
```

Fine Tuning the SYBASE Configuration

This sections contains instructions on how to increase the memory allocation for SYBASE to increase performance of the database.

Increasing the Memory

To increase the performance of SYBASE, Ascend recommends that you increase the memory allocation as follows:

1. Verify that you are logged in as the sybase user.
2. Initiate an interactive SQL session by typing the following command:

   ```
   isql -U sa -P superbase <Return>
   ```

3. Type the following:

   ```
   1> sp_configure memory, 20480 <Return>
   2> go <Return>
   ```

   The message “Configuration option changed. Run the RECONFIGURE command to install” appears.

4. Type the following:

   ```
   1> reconfigure <Return>
   2> go <Return>
   ```

5. Type the following to shut down the SYBASE Server:

   ```
   1> shutdown <Return>
   2> go <Return>
   ```

6. Restart the SYBASE Server as follows:

   ```
   cd install <Return>
   startserver -f RUN_CASCADE <Return>
   ```

7. If the startserver procedure fails due to a lack of available memory, type the following command to reset the SYBASE configuration back to its original state:

   cd $SYBASE <Return>

   buildmaster -d [*master_device_name*] -r <Return>

8. Restart the server by following the instructions in Step 6. Then proceed to the next section, "NFS Mounting the SYBASE Database".

NFS Mounting the SYBASE Database

This section describes how to nfs mount the **/opt/sybase** directory on any AP that needs to remotely access the billing system database. The **/opt/sybase** directory must be nfs mounted on an AP before you can begin installing the billing system software on that AP.

Set Up the SYBASE Server

The following steps need to be performed on the workstation you are using as the SYBASE server.

1. Edit the following file to set up nfs mounts and export the file system:

   vi /etc/dfs/dfstab <Return>

2. While holding down the Shift key, press **$G** to scroll to the end of the file.

3. Press Shift-A to append the following line to the file:

   share -F nfs -o rw -d "sybase 492" /opt/sybase

4. Press the Escape key.

5. While holding down the Shift key, type **zz** to save and exit the file.

6. At the command prompt, type

   shareall <Return>

Set Up the Adjunct Processor(s)

The following steps have to be performed on each AP before you begin the billing system software installation on the AP.

1. Edit the following file to mount the filesystem:

   vi /etc/vfstab <Return>

2. While holding down the Shift key, press **$G** to scroll to the end of the file.

3. Press Shift-A to append the following line to the file:

   [*SYBASE machinename*]:/opt/sybase - /opt/sybase nfs - yes -

4. Press the Escape key.

5. While holding down the Shift key, type **zz** to save and exit the file.

6. At the command prompt, type

   mkdir /opt/sybase <Return>

   mount /opt/sybase <Return>

7. Type **init 6** <Return> to reboot the system for the changes to take effect.

   Once the system is done rebooting, proceed to the next section for instructions on how to create the file **/etc/rc2.d/S97sybase**.

Creating the SYBASE Startup File

Before starting the billing system installation, you have to manually create the file **/etc/rc2.d/S97sybase** on your billing system SYBASE server. This file contains a script that automatically starts the SYBASE server when the workstation boots.

To create the file:

1. Log in to the system as root.

2. Type the following command to create the file:

   vi /etc/rc2.d/S97sybase <Return>

3. Type Shift-A, and add the following lines to this file:

```
#!/bin/sh <Return>
# <Return>
# script to start sybase server automatically <Return>

if [ -x /opt/sybase/install/startserver ] ; then <Return>
        /opt/sybase/install/startserver -f
        /opt/sybase/install/RUN_CASCADE <Return>
fi <Return>
```

4. Press Esc, then type **:wq!** and press Return to quit and save the file.

5. Type the following commands to change the owner and group for the file:

```
chown sybase /etc/rc2.d/S97sybase
chgrp dba /etc/rc2.d/S97sybase
```

6. Type the following commands to change the access privileges for this file:

```
chmod u+s /etc/rc2.d/S97sybase
chmod g+s /etc/rc2.d/S97sybase
```

When done, proceed to the next section for instructions on how to install the billing system software on the APs in your network.

Installing the Billing System Software

Before installing the billing system software on the AP:

- Verify that Solaris 2.3 and the Solaris Maintenance Release are installed and operating properly on the AP.
- Verify that SYBASE has been installed and is operating properly on the SYBASE server. If you are using the CascadeView SYBASE server as your billing system SYBASE server, verify that SYBASE is operating properly on that workstation.
- Verify that **/opt/sybase** on the SYBASE server is nfs mounted on the AP.
- Verify that a 1/4-in. cartridge (QIC) tape drive is accessible by the AP.
- Verify that the directory **/opt** exists and is mounted on the AP. In addition, verify that you have reserved the appropriate amount of disk space for **/opt**, based on the instructions in "Determining Disk Space Requirements" on page 3-3.
- Verify that the directories in which the billing record files and audit reconciliation files are to be stored have been created. Also, ensure that you have reserved the appropriate amount of disk space for the directory in which billing record files will be stored, based on the instructions in "Determining Disk Space Requirements" on page 3-3.

Also, you have to create a *billing* group on the AP, and configure at least one user in the *billing* group (to be used by your RAO Processing System to access the billing system). These procedures are described in the following two sections.

Creating a Billing Group

To create a *billing* group on the AP:

1. Log in as root and enter the root password.
2. Type the following command to run OpenWindows:

   ```
   /usr/openwin/bin/openwin <Return>
   ```

3. In a command tool window, run the Administrator Tool by typing

   ```
   admintool & <Return>
   ```

4. Select **Database Manager**.
5. In the Databases box, select **Group**, then select **None** for Naming Service.
6. Select **Load** to display the Group Database screen.
7. Select **Add Entry** from the Edit menu.
8. Enter *billing* in the Group Name field, enter an unused number greater than 15 in the Group ID field, then select **Add**.
9. Select **Close**, then select **Quit** to return to the main Administration Tool menu. Continue to the next section to add users to the *billing* group.

Configuring Users for Billing System Access

You have to configure at least one user that the RAO Processing System will use for accessing billing system files. Also, if you have other users who need access to the billing system, you need to configure them on the AP as described in this section.

To configure users for billing system access:

1. While still in the Administration Tool, select **User Account Manager**.
2. In the Naming Service field, click on **None**.
3. Choose **Apply**.
4. Choose Edit ⇒ Add User. The Add User dialog box appears. Complete the USER IDENTITY fields on the User Account Manager dialog box as follows:
   - In the Create User Name field, type the user id for the user.
   - In the Userid field, type any unused number greater than 15.
   - In the Primary Group field, type **billing**.
   - In the Secondary Group field, type **[*optional entry*]**.
   - In the Comment field, type **billing user**.

– In the Login Shell field, select the desired login shell.

All of the instructions in this chapter assume that you are using the Korn shell as the default login shell.

5. *Optionally,* complete the ACCOUNT SECURITY fields as follows:
   – In the Password field, select the level desired (Ascend recommends using **Cleared until first logon**).
   – In the Min change field, keep the default value of **0**.
   – In the Max change field, type the desired maximum number of days before forcing the user password change.
   – In the Max inactive days field, *optionally* type the maximum number of days that can pass without user activity before the user id is automatically deleted.
   – In the Expiration date field, *optionally* type an **expiration date** on which the user id should be automatically deleted.
6. Complete the HOME DIRECTORY fields as follows:
   – Set **Create home dir** to **checked**
   – Set **Path** to **/opt/[*user name*]**
   – Set **Server** to **[*hostname*]**
   – Set **Skeleton Path** to **/etc/skel**
   – Set **AutoHome Setup** to **unchecked**
   – Set **Permissions** as desired
7. Complete the MISCELLANEOUS field: Mail Server: **[*hostname*]**
8. Choose **Add**.
9. Quit the Solaris Administration tool.

10. At the # prompt, type

    su - [*user name*] <Return>

    where [*user name*] is the user id you just configured.

11. Rename the three files that you just created, as follows:

    ```
    $ mv local.cshrc .cshrc <Return>
    $ mv local.login .login <Return>
    $ mv local.profile .profile <Return>
    ```

12. At the $ prompt, type **exit** <Return> to switch back to the root user.

13. Exit OpenWindows and log out.

14. Log back in as the user whose account you just created.

15. At the system prompt, type

    vi .profile <Return>

*If you are using C-shell instead of Korn or Bourne shell, you must edit the **.login** file and the **.cshrc** file instead of **.profile**.*

16. Add the following lines to **.profile**, immediately after the "export PATH" line:

Line breaks are indicated by <Return>. Although some lines extend onto a second line, do not press <Return> until indicated.

```
set monitor <Return>
SYBASE=/opt/sybase <Return>
DSQUERY=CASCADE <Return>
PATH=$PATH:/opt/CascadeBS/bin:/opt/CascadeBS/etc:
$SYBASE/bin:$SYBASE/include/bin <Return>
OPENWINHOME=/usr/openwin <Return>
LD_LIBRARY_PATH=/opt/lib:/usr/openwin/lib:/opt/sybase/lib:
/usr/lib <Return>
EDITOR=vi <Return>
```

17. Press Escape.

18. While holding down the Shift key, type **zz** to save and exit the file.

19. In order for the environment variables that you just set to take effect, you have to exit, then log back in as the user id you just configured (in order to update the environment variables).

20. Repeat the steps in this section to add any additional users who need access to the billing directories.

Installing the Software

Once you have created a *billing* group, and have added at least one billing system user to the *billing* group, you can install the billing system software, using the following procedure:

1. Log in as root and switch to the **/opt** directory.

2. Type the following command to extract the billing system files from the installation tape:

   /bin/tar xvpf /dev/[*tape device*] <Return>

   where [*tape device*] is the device designation for your tape drive (e.g., rst4).

3. Type the following command to initiate the configuration script:

   /opt/CascadeBS/etc/ap-install.sh <Return>

4. When prompted for the name of the SYBASE database server, press Return to accept the default name (CASCADE). If you specified a database server name other than CASCADE during the SYBASE installation, enter that name instead.

5. When prompted for the SYBASE system administrator name, type **sa** and press Return. If you specified a different sybase administrator name during the SYBASE installation, enter that name instead.

6. When prompted for the SYBASE system administrator password, type **superbase** and press Return. If you specified a different password during the SYBASE installation, enter that password instead.

At this point, you are prompted to create the CascBilling database.

You only need to create the CascBilling database when installing the billing system software on the first AP that is connected to a given SYBASE server. You do not have to create it when installing the software on subsequent APs that are connected to that same SYBASE server. If you have multiple SYBASE servers, you have to create the database once for each server (that is, create the database when installing the software on the first AP that is connected to a given SYBASE server).

The following messages may be displayed during the database creation process; these messages can be ignored:

```
Cannot drop the database “CascBilling” because it doesn’t
exist on the system.

Cannot drop the table “SmdsCustomer” because it doesn’t
exist on the system.

Cannot drop the table “SwitchInfo” because it doesn’t exist
on the system.
```

7. At this point, the following operations are performed automatically:

   - Billing system services are added to **/etc/services** and **/etc/inetd.conf**.
   - The **inetd** daemon is signaled to reread its configuration.
   - Billing system procedures are added to **/etc/inet.d** and **/etc/rc3.d**.

   Once these operations are completed, you are prompted for the following information. At each prompt, enter the appropriate information and press Return:

   "Enter the Adjunct Processor ID number"

   

   The Adjunct Processor ID must be a unique value in the range 1-255. For example, if you have a primary AP and a backup AP, assign ID 1 to the primary AP, and ID 2 to the backup AP. No two APs in the network can have the same ID.

   "Enter the Audit Reconciliation File directory path"

   "Enter the Billing Usage Record File directory path"

8. After entering this information, the following operations are performed automatically:

   - The AP daemon, **apd**, is started.
   - The Error Suspense File crontab entries are created.
   - The Daily Audit Reconciliation processing **at** job is created.

9. Once these operations complete, the AP software installation is done. You now have to update the CascBilling database with your network's B-STDX switch and SMDS customer data, as described in the *Software Release Notice for the Ascend Billing System Adjunct Processor.*

When done updating the database, go to Chapter 5 for billing system configuration instructions.

5

Billing System Configuration

This chapter describes the ASCII configuration file, **brf.cfg**, which contains basic AP configuration settings. It also describes how to set up each AP for communicating with the switch network it will be servicing. In addition, this chapter describes how to configure each switch and SMDS logical port for billing.

Adjunct Processor Configuration

The AP is configured via the file **/opt/CascadeBS/data/brf.cfg**, which is created during the billing system software installation. This file contains the AP identifier (APID), audit record file and billing record file directory names, and the environment variable DSQUERY, which identifies the SYBASE SQL server name for the billing system. In addition, other variables may be included in this file, based on your billing system requirements. A sample **brf.cfg** file is shown here:

```
APID = adjunct1
AUDIT_DIR = /billing/audit
BILLING_DIR = /billing/usage
DSQUERY = CASCADE
```

There must be one NMSTrapAddr entry for each NMS to which you want SNMP Trap Event messages sent. These entries consist of the network address of the NMS, as well as the Community Name **(**and Community ID) of the NMS workstation. These entries must match the NMS network address/Community strings already defined in CascadeView (for more information, refer to the *Ascend Network Configuration Guide for CascadeView/UX*).

The **/opt/CascadeBS/data/brf.cfg** file is created when the billing system software is installed on the AP. The audit file directory and billing directory shown here are examples; you may choose to store audit and billing files anywhere. Please note that the APID may be a network name, network address, or some other unique identifier, based on your Provisioning System setup.

Directory Structure and Contents

The AP billing system files are installed to the directory **/opt/CascadeBS**. The directory structure is shown in Figure 5-1. A sample directory structure for storing billing and audit record files is also shown.

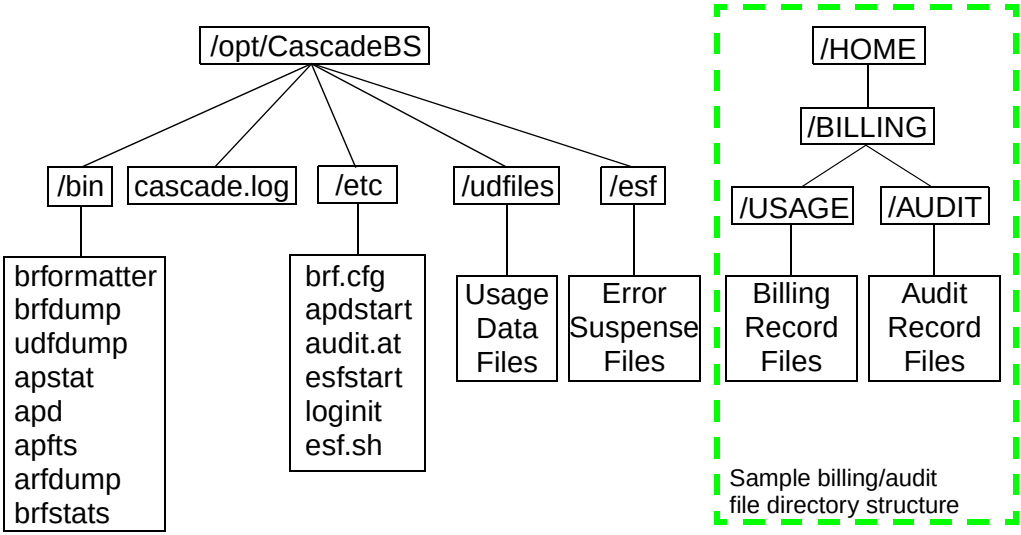

Figure 5-1. Directory Structure of Billing System Files

The following components and sub-components are contained in the **/opt/CascadeBS** directory:

- The **bin** directory, which contains the following billing system executables:
  - **apd**, which is the AP daemon
  - **apstat**, which is the AP statistics utility
  - **apfts**, which manages/performs usage file transfers between each switch and the AP
  - **brformatter**, which is the Billing Record Formatter program
  - **udfdump**, which produces a dump of all usage data files currently stored on the AP
  - **brfdump**, which produces a dump of any billing record file currently stored on the AP
  - **arfdump**, which produces a dump of any audit reconciliation (daily processing statistic) file currently stored on the AP

  - **brfstats**, which enables you to summarize billing statistics from a large set of billing record files

- The **udfiles** directory, which contains the usage data files transferred from each switch in the billing system that this AP services
- The **data** directory, which contains the AP configuration file, **brf.cfg**. This file is populated with the proper configuration information by the installation procedure.
- The **cascade.log** file, which is the ASCII log file dedicated to billing system processing. All billing system events are logged to this file.

Configuring Adjunct Processors on the Network

In order for the APs in your switching system to communicate with their respective switch networks, you have to do the following:

- Define the SMDS in-band management port the AP will be using to communicate with the switch network (this is required only if you are managing via SMDS in-band management).
- Add the network address of each AP that will be servicing a switch network. The NMS Path mechanism is used to define the route from the switch network to the AP.
- Disable the SNMP trap function of CascadeView to prevent NMS SNMP traps from being sent to the APs in your network.
- Create a network route from the AP to the gateway switch or router through which the AP is connected to the switch network.
- Save routing entries in a file on the AP to avoid having to recreate the routes each time you restart the AP.

Defining an SMDS In-Band Management Port

This section describes how to configure the SMDS logical port that the AP will use to manage the switches it services. This enables an AP that is remotely connected to the network via SMDS services to transport packets. You need to follow the instructions in this section only if the AP will be managing the switches via SMDS in-band management. If the AP is managing its switches via a direct or indirect Ethernet connection, skip to the next section.

The instructions in this section assume that your SMDS logical ports have already been configured in CascadeView. For complete logical port configuration instructions, see the *Ascend Network Administrator's Guide for CascadeView/UX*.

To define an SMDS in-band management port:

1. On an NMS workstation, bring up the CascadeView network map for the switching system on which the AP is installed. (If you have to start CascadeView, see the *Ascend Network Administrator's Guide for CascadeView/UX*).
2. Select Misc ⇒ CascadeView Logon and log in to the network map.
3. On the network map, select the gateway switch for this switch network.
4. From the Administer menu, choose Cascade Parameters ⇒ Set All Management Addresses. The Set All SMDS Management Address dialog box appears.

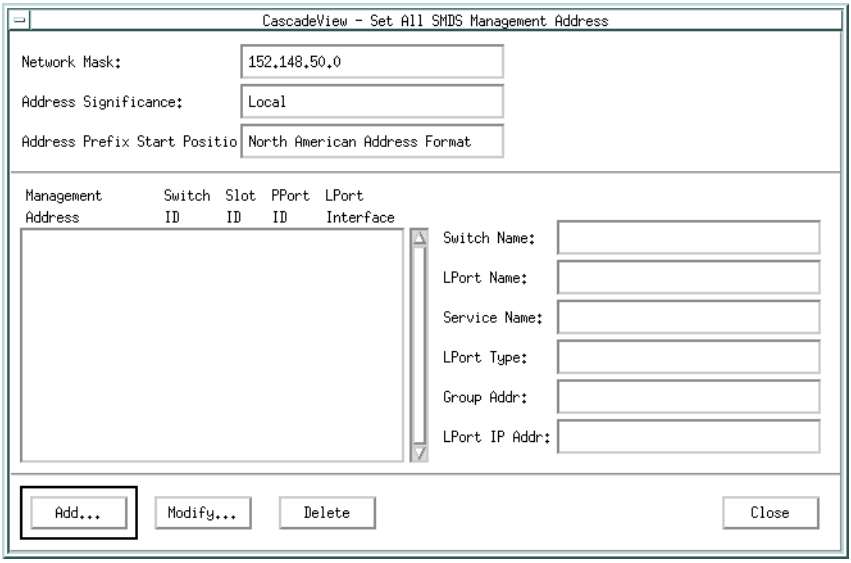

If you have already configured management addresses, the dialog box displays this information. From the Set All SMDS Management Address dialog box, you can use the Modify. . . *or* Delete *commands to modify or delete management address configurations.*

5. Choose Add. The Select End Logical Port dialog box appears.

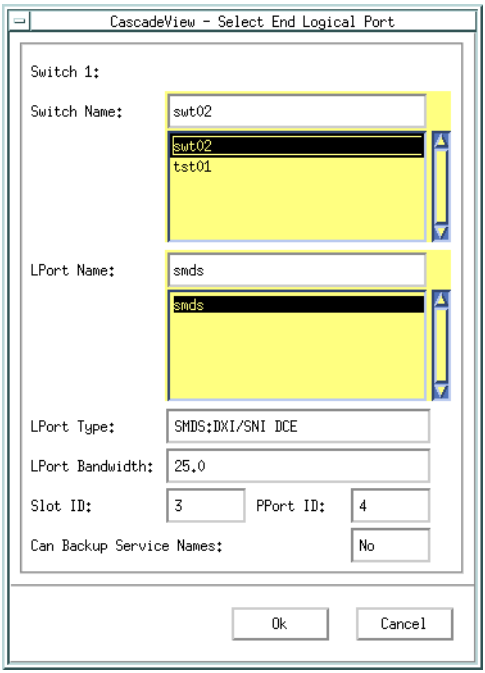

6. Complete the Select End Logical Port dialog box fields as follows:

   Switch Name — Select the name of the switch that contains the Address pool.

   LPort Name — Select the name of the logical port for which you are defining the In-Band management address.

   LPort Type — Displays the logical port type.

   LPort Bandwidth — Displays the logical port bandwidth.

   Slot ID — Displays the I/O slot number in which the card resides.

   PPort ID — Displays the port number for the port you are configuring.

7. Choose OK. The Add SMDS Management Address dialog box appears.

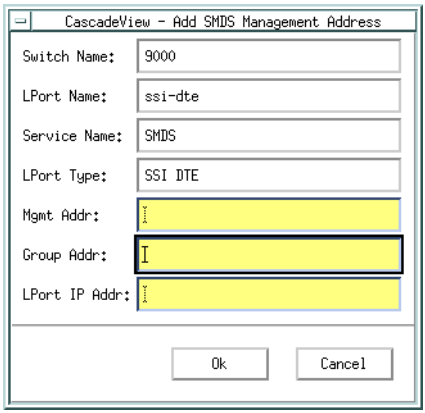

8. Complete the Add SMDS Management Address dialog box as follows:

   Switch Name — Displays the name of the selected switch.

   Service Name — Displays the type of service (SMDS).

   LPort Name — Displays the name of the SSI-DTE logical port configuration that contains the Address pool.

   LPort Type — Displays the type of logical port configuration (SMDS DTE).

   Mgmt Addr — If the LPort Type is SSI, the management address is already entered automatically in this field. If the LPort Type is DXI, you have to manually enter the individual address to which the logical port subscribes.

   Group Address — If applicable, enter the Group address configured on the router to which the management address is provisioned.

   LPort IP Address — Enter the configured IP address of the selected logical port.

9. Choose OK. The Set In-Band Management Addresses dialog box reappears and displays the In-band Management Address you just defined.

You can now add the AP addresses to the Set NMS Path dialog box, as described in the next section.

Creating Routes from the Switch Network to APs

In order for your APs to communicate properly with the switch network, you have to create routes form the switch network to each AP that is servicing that network. This is currently done via the NMS Path function of CascadeView.

1. On an NMS workstation, bring up the CascadeView network map for the switching system on which the AP is installed. (If you have to start CascadeView, see the *Ascend Network Administrator's Guide for CascadeView/UX*).

2. Select Misc ⇒ CascadeView Logon and log in to the network map.

3. On the network map, select the switch you are using as the gateway switch (this is the switch on which you have defined the NMS path).

4. From the Administer menu, choose Cascade Parameters ⇒ Set NMS Path. The Set NMS Path dialog box appears.

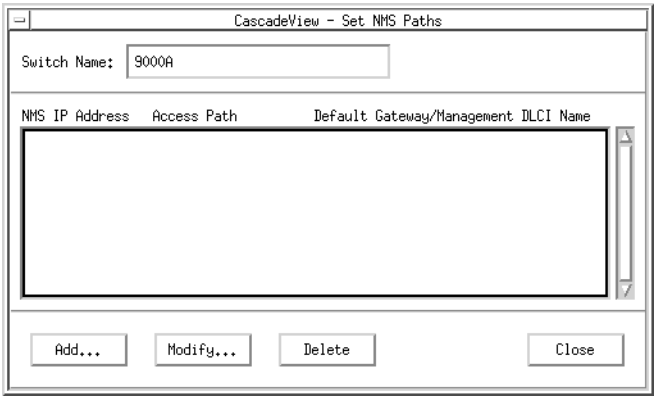

5. Choose Add. The Add NMS Path dialog box appears (this example shows an Indirect Ethernet NMS Path dialog box).

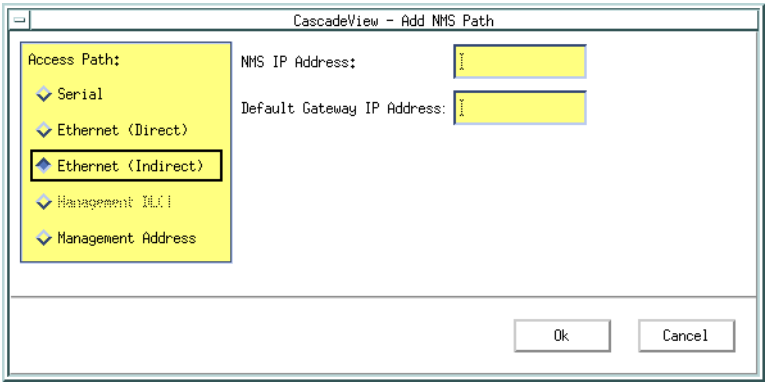

To complete this dialog box:

a. In the Access Path field, select the connection method you used to connect the AP to the switch network (direct Ethernet, indirect Ethernet, or Management Address).

b. In the NMS IP Address field, enter the IP address of the AP.

c. If you connected the AP via indirect Ethernet, enter the IP address of the router in the Default Gateway IP Address field.

   If you connected the AP via SMDS in-band management, select the management address from the displayed list.

d. Choose OK to add the AP to the NMS Path list.

6. Repeat Step 5 for each AP that is connected to this switch network. When done, choose Close from the Set NMS Path screen to return to the network map.

7. If you have additional APs connected to another switch network, repeat Steps 1 through 6 to configure them.

Disabling the SNMP Trap Mechanism

Because you have added your APs to the NMS Path screen, CascadeView will attempt to send SNMP trap messages to each AP, as if they were NMS workstations. You have to disable this mechanism to prevent the SNMP trap messages from being sent to your APs. To do so, you have to modify the NMS entry that was automatically created when you added the AP to the NMS Path screen, by doing the following:

1. From the Administer menu in CascadeView, choose Cascade Parameters ⇒ Set Parameters to access the Switch Back Panel screen.
2. From the Switch Back Panel screen, choose the Set Sw Attr command button.
3. From the Set Switch Attributes dialog box, choose the NMS Entries command button.
4. For each AP address listed on this screen, do the following:
   a. Choose the entry for the AP.
   b. Choose the Modify command button.
   c. Choose No in the Receiving Traps field.
   d. Choose OK, then choose Close until you are back to the network map.

Creating a Route from the AP to the Switch Network

Before defining the switch and logical port configuration in CascadeView, you have to create a network route from each AP to the switch network that the AP is servicing.

On each AP in your billing system, type the following command to create the route:

```
/usr/sbin/route add net [switch network address]
[gateway address] 1 <Return>
```

where **[*switch network address*]** is the IP address of the switching network on which the switches being serviced by the AP reside, and **[*gateway address*]** is the IP address of the router or gateway switch through which the AP is connected to the switch network.

Enter a router address if you used indirect Ethernet or SMDS in-band management to connect the AP to the network; enter a gateway switch address if you used direct Ethernet to connect the AP to the network.

In order to avoid recreating these routes each time you restart the AP, you now have to add the AP's routes to a file on the AP, as described in the next section.

Saving Route Entries

To save each routing entry so that it is automatically reloaded into the Sun operating system routing table when you restart the AP, add the routing entries you created in the previous section to the file **/etc/rc3.d/S99ap.routes**.

The commands in this file will be executed each time you restart the AP. If the routing entries are not recorded in this file, the configured route will be lost when the AP is restarted, and communication from the AP to the switch network will fail unless you manually enter the routes after the restart (as described in the previous section).

Switch and Port Configuration

This section describes how to configure each switch and logical port in the billing system. Switches and logical ports are configured by running CascadeView/UX on any NMS workstation you are using to manage the network. The procedure in this section assumes that each switch and logical port has already been configured on the network (that is, you have already created your CascadeView network map). If you have not created the network map yet, or if you need more detailed instructions on these steps, refer to the *Ascend Network Administrator's Guide for CascadeView/UX*.

Switch Configuration

This section describes how to define the billing system parameters for each switch in the billing system, as well as how to enable/disable billing on a switch-wide basis.

To configure the switches in the billing system

1. Log in to the NMS workstation.
2. Run OpenView.
3. From the Misc menu, choose CascadeView ⇒ Logon, and log in to the network map.
4. From the file menu, choose Open/List Maps. Select the appropriate map name, then choose OK to display the map.
5. Once the map is displayed, select the switch on which you want to enable billing services.
6. From the Administer menu, select Cascade Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
7. Select the CP card (on the far right). Then choose the Set Attr ... command button. The Set Card Attributes dialog box appears.
8. Verify that "Capability" is set to CP Plus. Then choose OK to return to the Switch Back Panel dialog box.
9. Choose the Set Sw Attr ... command button. The Set Switch Attributes dialog box appears.

10. From the Set Switch Attributes dialog box, choose the Billing ... command button. The Set Billing Attributes dialog box appears.

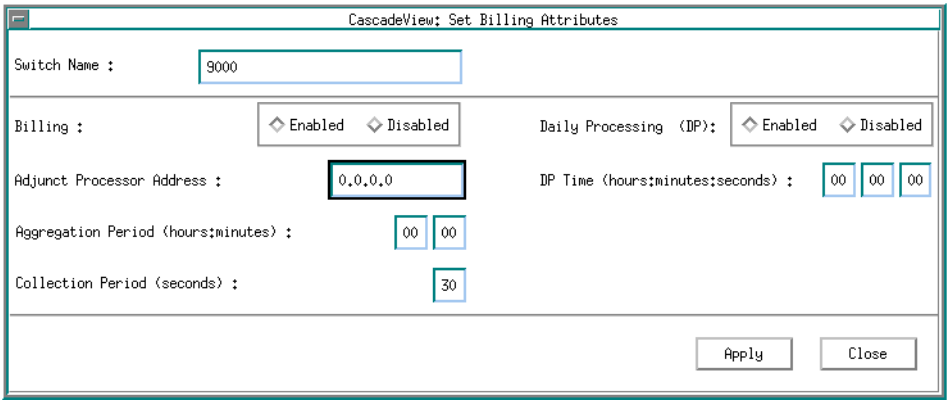

The settings in this dialog box are described as follows:

Billing — This field determines whether or not billing is enabled on the switch. Usage data is not collected on any logical ports on the switch until this field is set to Enabled. When set to Disabled, billing is disabled on all logical ports on the switch, even on those ports that are currently enabled for billing.

Adjunct Processor Address — The network address of the AP that will process billing records for this switch.

Aggregation Period — Defines the length of time (in hours:minutes) that usage data is collected and stored in the Aggregated Usage Data store before being transferred to the AP. The default aggregation period is 15 minutes; the valid range is 1-1440 minutes. If you notice that aggregated usage data is frequently being transferred to the AP before the aggregation period expires, you may want to define a shorter aggregation period.

Collection Period — The number of seconds that usage data is collected on an IOP Plus module before being transferred to the Aggregated Usage Data store. The default value is 30 seconds.

Daily Processing — Specifies whether or not daily audit statistics are to be gathered on the switch. When set to Enabled, daily audit statistics are collected over a 24-hour period, and sent to the AP at the time specified in the Daily

Processing Time field. When set to Disabled, daily audit statistics are not collected on the switch.

DP Time — Specifies the time of day (in hours:minutes:seconds) at which daily audit statistics are transferred to the AP. The valid range is 00:00:00 through 23:59:59. The value is entered in Universal Coordinated Time (UCT). The default value is 00:00:00 UCT.

DP Time should be configured to coincide with the end of an aggregation period. For example, if the aggregation period is 30 minutes, schedule the DP Time at *nn*:00 or *nn*:30 (as opposed to *nn*:15 or *nn*:25). The reason for this is that once the DP Time is reached, daily usage statistics are sent to the AP when the current aggregation period ends; therefore, it is preferable that the aggregation period end at the same time that daily statistics are scheduled to be transferred to the AP.

11. Once you have made the appropriate entries in this dialog box, choose the Apply button to return to the Set Switch Attributes dialog box.

12. From Set Switch Attributes, choose the Apply command button, then choose Close to return to the Switch Back Panel dialog box.

For instructions on how to enable/disable the logical ports on your billing system switches, refer to the next section.

Usage data is not collected on a given logical port until both the logical port and the switch are enabled for billing.

Enabling/Disabling Billing on Logical Ports

This section describes how to enable/disable billing on each logical port on a switch.

To enable/disable a logical port for billing

1. If it is not already displayed, access the Switch Back Panel dialog box.
2. From the Switch Back Panel dialog box, select a physical port that is configured with logical ports on which you want to enable billing. Then select the Set Attr ... command button to access the Set Physical Port Attributes dialog box.
3. From this dialog box, select the Logical Port ... command button to access the Set All Logical Ports in PPort dialog box. This dialog box contains a list of the logical ports that are defined for the selected physical port.
4. For each SMDS logical port on which you want to enable billing

   a. If you are defining a new logical port, select the Add button. For instructions on how to add a logical port, refer to the *Ascend Network Administrator's Guide for CascadeView/UX*.

      If you are modifying a logical port that has already been configured, select the logical port in the Logical Port Names section. Then select the Modify button to display the Modify Logical Port dialog box.

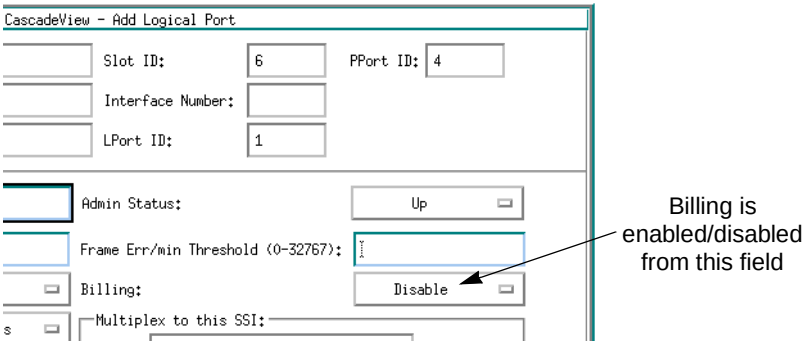

   b. If you want to enable billing on the logical port, set the Billing field to Enabled. If you want to disable billing on the logical port, make sure you set the Billing field to Disabled. Currently, billing can be enabled only on logical ports that are configured for SMDS service.

c. When done, select OK to return to the Set All Logical Ports dialog box.

d. Repeat Steps 4a through 4c for each logical port on the current physical port. When done, choose Close from the Set All Logical Ports dialog box to return to the Set Physical Port Attributes dialog box. Then choose Close to return to the Switch Back Panel dialog box.

5. Repeat Steps 2 through 4 for each physical port that is configured with logical ports on which you want to enable billing.

6. When done configuring billing on all logical ports, choose Close from the Switch Back Panel dialog box.

6

Managing the Billing System

This chapter describes how the AP functions and B-STDX switch functions for the billing system are managed and controlled. AP functions are currently managed locally on the AP. The B-STDX switch functions are managed from within CascadeView/UX on the NMS.

This chapter also describes the daily statistics that are gathered from the billing system, and how to access them.

Adjunct Processor Management

This section describes the Ascend utilities that are available for displaying various information that is collected and stored on the AP. The utilities that Ascend supplies include

apdstart — This utility enables you to start and stop the **apd** daemon.

apstat — This utility enables you to obtain statistics that were collected on each switch being serviced by the AP.

arfdump — This utility enables you to display audit reconciliation files generated by the AP and the switches in the billing system.

audit.at — This utility allows you to enable or disable the AP's daily audit processing.

brfdump — This utility enables you to display a dump of the billing record files currently stored on the AP for a specified switch.

brfstats — This utility enables you to summarize and display billing statistics from a specified set of billing record files.

esfstart — This utility is used to enable or disable the periodic processing of files in the Error Suspense file directory, **/opt/CascadeBS/esf**.

loginit — This utility enables you to reset the **/opt/CascadeBS/cascade.log** file, which is used to log billing system events.

udfdump — This utility enables you to display a dump of a specified Usage Data file for a specified switch.

Displaying Adjunct Processor Statistics

Ascend's **apstat** utility, which resides on the AP, allows the billing system administrator to view billing statistics for each switch in the system. Using **apstat**, you can display statistics for the current aggregation period, or you can display all statistics collected since the beginning of the current daily period.

The following statistics are displayed by **apstat**:

- The switch ID of each switch in the billing system that is connected to the AP on which **apstat** is being run
- The type of service (currently, only SMDS is supported)
- The number of usage records received by the AP
- The number of billing records created by the AP
- The number of billing record creations that failed

To display the statistics for a given AP

1. Log in to the workstation, either locally or remotely. You do not have to log in as root in order to display these statistics.

2. Type one of the following commands to execute the **apstat** utility:

   apstat -p — Displays statistics for the current aggregation period.

   apstat -t — Displays statistics collected since the beginning of the current day.

A display similar to the following appears. The items that are displayed depend on your billing system setup. For example, Switch ID may be a switch name, a unique character string, or a network address.

```
Mon Apr 10 1995  10:02:01 UCT (Mon Apr 10 1995  5:02:01 Local Time)

  Switch ID   Service        Usage Rec     Billing Rec     Creation
               Type          Received       Created       Failures
   Boston1     SMDS            847125         748375             0
   Boston2     SMDS           1124417        1032531             0
   Atlanta1    SMDS           2141171        2070024             0
   New York1   SMDS            653392         599841             0
```

Note that the current date and time are shown in the display header. Date and time is shown in UCT, as well as local time.

The columns on this screen are described as follows:

Switch ID — This is a unique identifier assigned by the Provisioning System to each switch connected to this AP. The switch ID identifies each switch to the AP. Depending on your Provisioning System setup, Switch ID may be a switch name, switch address, or a unique character string.

Service Type — Identifies the type of service for which billing is enabled on the switch identified by Switch ID. Currently, SMDS is the only service for which billing is supported.

Usage Rec Received — Specifies the number of usage records that have been sent to the AP by the Data Aggregator on the switch identified by Switch ID.

Billing Rec Created — Specifies the number of billing records that have been created and stored on the AP for the switch identified by Switch ID.

Creation Failures — Specifies the number of billing record creations that failed for the switch identified by Switch ID. Typically, this is the number of records sent to the Error Suspense file.

Displaying Usage Data Files

Ascend's **udfdump** utility, which resides on the AP, enables you to display the Usage Data files for any switch being serviced by the AP. The files are displayed in their original format (that is, before the billing record formatter converts them to the appropriate billing record format).

Usage data files are stored in two directories:

/opt/CascadeBS/udfiles — Usage data files received from all switches connected to the AP are stored in this directory. The billing record formatter is run periodically on each file in this directory to convert the files into the appropriate billing record format for your RAO Processing system. This directory contains usage data files that are awaiting processing by the Billing Record Formatter.

/opt/CascadeBS/esf — This directory contains the Error Suspense files, which store usage data records that could not be converted into the proper billing record format by the Billing Record Formatter (see "Error Suspense Files" on page 2-14 for more information on Error Suspense files).

You can use **udfdump** to display the usage data files in either of these directories, as described in the following sections.

Displaying Current Usage Data Files

To display the usage data files currently stored in the **/udfiles** directory, type the following command:

udfdump [*usage data filename*] <Return>

where [*usage data filename*] is in the format shown here.

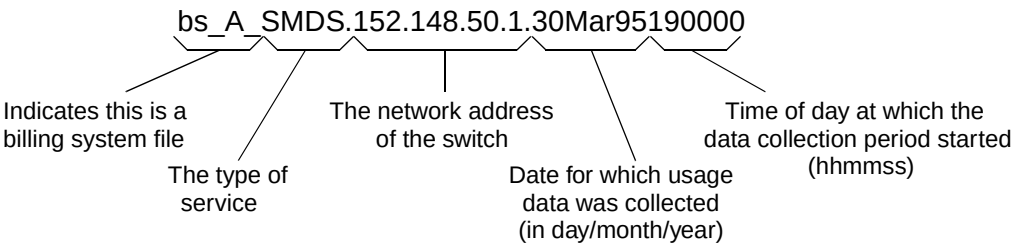

The following illustration shows a sample output of the **udfdump** command. Again, some field entries, such as Switch ID, may differ from what is shown here.

```
Usage Data File: bs_A_SMDS.152.148.40.1.30Mar95190000

File header record:
     Usage data type SMDS Aggregated Usage Data
     Switch ID       152.148.40.1
     Start time      Thu Mar 30 19:00:00 1995
     End Time        Thu Mar 30 19:15:00 1995
     File Size       65672 bytes
     CRC             0xbc60c119
     Version         1

SMDS Header:
     Header version                    1
     Number of usage records in file   4
     Number of entries in Lookup Table 4
     Record Table Offset               65576
     E.164 Lookup Table Offset         40

Record 1:
     DA c1 40 39 70 0 0 ff ff (LUT index 0)
     SA c1 40 39 71 0 0 ff ff (LUT index 0)
     Start time      Thu Mar 30 19:00:01 1995
     End time        Thu Mar 30 19:15:00 1995
     L3 PDU count    550116
     Byte count      169435728 (0x0a196250)
     Flags           3
     Processed       1

     .
     .
     .

Record 4
     DA c1 40 39 70 0 0 ff ff (LUT index 0)
     SA c1 40 39 71 0 0 ff ff (LUT index 0)
     Start time      Thu Mar 30 19:00:01 1995
     End time        Thu Mar 30 19:15:00 1995
     L3 PDU count    550116
     Byte count      169435728 (0x0a196250)
     Flags           3
     Processed       1
```

Displaying the Error Suspense Files

udfdump also enables you to display billing records that could not be converted to the appropriate billing record format. These records are stored in Error Suspense files, and are periodically re-processed by the billing record formatter to see if they can be converted. An Error Suspense file is created for each aggregation period during which one or more billing records could not be processed by the billing record formatter.

To display a usage data Error Suspense file, type the following command:

```
udfdump -u [error suspense filename] <Return>
```

where **[*error suspense filename*]** is in the same format as described in the previous section. The display is also the same as shown in the previous section, except that the text "** Record has not been converted" is displayed after each record that has not been converted to the appropriate billing record format.

Displaying Billing Record Files

Ascend's brfdump utility enables you to display billing records that have been converted into the appropriate format for your RAO Processing system. These files are stored in the directory you specified during the installation process.

Displaying Formatted Records

To display formatted SMDS billing records on the AP terminal, type the following command:

```
brfdump [billing record filename] <Return>
```

The format of **[*billing record filename*]** depends on your billing system requirements. In general, the filename prefix identifies the type of service (e.g., S for SMDS) and a switch identifier for the switch on which the data in the file was collected. The filename extension is in the format S*NNN*, where *NNN* is a sequential number that is incremented by one each time a new billing record file is created for a given switch.

For example, to display the second SMDS billing file created on the switch that is assigned Switch ID Boston1, type the following command:

```
brfdump SBoston1.S001 <Return>
```

The following illustration shows a sample **brfdump** display. All values are in decimal, with the exception of Source, Destination, and Group addresses, as well as Switch ID.

```
Billing Record File: SBoston1.S001

Header Record
     Switch id             Boston1
     Total Length          66
     Sequence Number       1
     Total Records         4
     Start Time            Thu Mar 30 19:00:00 1995
     Timezone Offset       0
     Time Delta            0 days 0 hours 15 min 0 sec

SMDS Record 1
     Switch id             Boston1
     Total Length          66
     Data Service Type     2 (SMDS DXI/SNI)
     Start Time            Thu Mar 30 19:00:00 1995
     Time Delta            0 days 0 hours 14 min 45 sec
     Source Address        c1 40 39 71 00 00 ff ff
     Destination Address   c1 40 39 70 00 00 ff ff
     Timezone Offset       1
     Group Address         00 00 00 00 00 00 00 00
     Data Units (L3 PDUs) 550116
     Data Bytes            169435728 (0xa196250)

SMDS Record 2
     .
     .
     .

Trailer Record
     Switch id             Boston1
     Total Length          66
     Sequence Number       1
     Total Records         4
     Start Time            Thu Mar 30 19:00:00 1995
     Timezone Offset       0
     Time Delta            0 days 0 hours 15 min 0 sec
     Adjunct Processor     adjunct1
     Local/Regional ID     0
```

The items listed on this screen are described as follows:

Record Identifier — Indicates whether the record is a header, trailer, or SMDS usage data record.

Switch ID — Uniquely identifies the switch for which the billing record file was created. The Switch ID is defined on your Provisioning System. The format of Switch ID depends on your Provisioning System setup.

Total Segments Used — Indicates the number of segments that come after the billing record header.

Total Length — Indicates the total length of the billing record header/trailer (in bytes). This value should always be 42h (66 decimal) for billing record files.

Segment ID Number — Indicates the type of billing record. Currently, only SMDS is supported. SMDS billing records are identified by a Segment ID of 0060h (96 decimal).

Segment Version — Enables a new version of a given segment (i.e., a segment to which changes have been made) to be identified. Whenever a segment is changed, this number is incremented by one.

Total Segment Length — Indicates the length of the segment. This value should always be 38h (56 decimal).

Sequence Number — Indicates where in the sequence the billing record was created. Billing record files are created sequentially, starting with zero (000).

Total Records — Indicates the total number of billing records contained in the file. This value does not include the header and trailer records.

Start Time — Indicates the day and time at which the aggregation period for this billing file began. In the example, the aggregation period started on March 30, 1995 at 7:00 PM.

Timezone Offset — Indicates the timezone offset for the switch on which the billing data was collected. If the customer's access connection is physically located in a different time zone than the one in which the switch is located, the difference is noted in this field.

A value of 0 through 12 (decimal) indicates the number of time zones ahead. A value of 17 through 27 (decimal) indicates the number of time zones behind (where 17 is one time zone behind, and 27 is eleven time zones behind). For example, if the access connection is located in the Eastern Time Zone, and the switch is located in the Central Time Zone, this value would be 1; if the access connection is located in the Central Time Zone, and the switch is located in the Eastern Time Zone, this value would be 17.

This field is always 0 in the billing record header and trailer.

Time Delta — Indicates the amount of time that has elapsed between the start and end of data collection. Normally, this is the length of the aggregation

period. If, however, there are two data collection periods in a single aggregation period, this value would be less than the length of the aggregation period. For example, if you have a 30 minute aggregation period, and there are two data collection periods (0-18 minutes, and 19-30 minutes) in the aggregation period, the first set of records would have a Time Delta of 18 minutes, and the second set of records would have a Time Delta of 12 minutes.

This value is stored in raw billing records as the number of seconds that have elapsed between the start and end of data collection.

Adjunct Processor — Uniquely identifies the AP that created these billing record. The AP identifier is defined in the file **/opt/CascadeBS/brf.cfg**. The format of this identifier depends on your Provisioning System setup.

Local/Regional ID — Identifies the billing file's source regional/local billing processor. This field is used only when the billing system contains regional processors that reside near or at switch sites. Otherwise, it is set to 0.

Data Service Type — Identifies the type of service (SMDS SNI/DXI, ATM PVC, Frame Relay PVC, etc.) for which these billing records were created. Currently, only SMDS SNI/DXI is supported; therefore, the value of this field should always be 02h (2 decimal).

Address Type — Indicates the address type of the logical port (customer interface) for which the billing record was created. Address type can be E.164 (Type 4), 3-Byte DLCI (Type 7), 4-Byte DLCI (Type 9), or VPI/VCI (Type 10). Currently, only Type 4 is supported.

Source Address — Identifies the source E.164 address of the L3 PDU to which the billing record applies.

Destination Address — Identifies the destination E.164 address to which the L3 PDU was sent. If the L3 PDU was sent to a group address, this field is all zeros.

Billing Company — Identifies the country (carrier partner) responsible for generating the customer bill. These values are specific to each carrier, and are defined on the Provisioning System.

Carrier ID — Identifies the Carrier Prime (owner of the customer for whom the bill is to be generated). These values are specific to each carrier, and are defined on the Provisioning System.

Group Address — Identifies the destination E.164 group address to which the L3 PDU was sent. If the L3 PDU was sent to a single destination address instead of a group address, this field is all zeros.

Data Units (L3 PDUs) — Indicates the total number of terminating L3 PDUs sent to the destination (or group) address specified in the record.

Data Bytes — Indicates the total number of L3 PDU data bytes sent to the destination (or group) address specified in the record. **brfdump** displays this value in both decimal and hexadecimal format.

Displaying a Raw Dump of Billing Record Files

You can also display a raw dump of a billing record file on the AP. The billing records are displayed in the format defined by the billing record formatter you are using.

To display a raw dump of a billing record file, type the following command:

```
brfdump -r [filename] <Return>
```

where **[*filename*]** refers to the billing record file you want to display (for example, SBoston1.S001).

The following illustration shows an example of the **brfdump -r** command. The actual format depends on the structure of your raw billing records.

Displaying Audit Reconciliation Files

Ascend's **arfdump** utility enables you to display both AP and Switch Billing audit reconciliation files, which are stored in the directory you specified during the billing system software installation. AP audit reconciliation files contain combined data on all switches connected to the AP, whereas Switch Billing audit reconciliation files only contain information for a particular switch.

```
Billing Record File: SBoston1.S001

Header record:
      0:  0  0  0  0 33 42 47  1  0 42  0
     11: 60  0 38  0  0  0 82  0  4  e 18
     22: e6 c8 c0  0  0  1 5f 90  a  0  0
     33:  0  0  0  0  0  0  0  0  0  0  0
     44:  0  0  0  0  0  0  0  0  0  0  0
     55:  0  0  0  0  0  0  0  0  0  0  0

SMDS record 1:
      0: 33 42 47  1  0 42  0 60  0 38  0
     11:  2  e 18 e6 ce 9c  4  0  1 59 b4
     22: c1 40 39 71  0  0 ff ff c1 40 39
     33: 70  0  0 ff ff  0  a  0  1  0  0
     44:  0  0  0  0  0  0  0  0  0  8 64
     55: e4  0  0  a 19 62 50  0  0  0  0

SMDS record 2:
     .
     .
     .

Trailer record:
      0: ff ff ff  0 33 42 47  1  0 42  0
     11: 60  0 38  0  0  0 82  0  4  e 18
     22: e6 c8 c0  0  0  1 5f 90  a  0  0
     33:  0  0  0  0  0  0  0  0  0  0  0
     44:  0  0  0  0  0  0  0  0  0  0  0
     55:  0  0  0  0  0  0  0  0  0  0  0
```

Displaying Adjunct Processor Reconciliation Files

To display an AP reconciliation file, type the following command:

arfdump [*AP reconciliation filename*] <Return>

where **[*AP reconciliation filename*]** is in the format shown here:

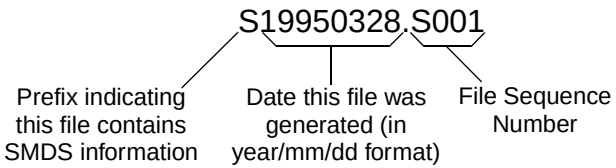

The information that is displayed is similar to the information displayed by the **brfdump** utility (see page 6-7), with the exception that Source and Destination addresses are not listed and the Switch ID field is all zeros (since these files contain combined switch information). Also, four additional fields are displayed. These fields are described as follows:

Records Created — The total number of billing records created on the AP for all switches connected to the AP during the aggregation period to which the file applies.

Records Transported — The number of billing records transported to the RAO Processing system for all switches connected to the AP during the aggregation period to which the file applies.

Records Received — The total number of billing records received by the AP from all switches connected to the AP during the aggregation period to which this file applies.

Records Dropped — The number of billing records received by the AP that were dropped because the record could not be processed by the billing record formatter.

For a description of the other fields on this display, see page 6-7.

Displaying Audit Reconciliation Files

To display an audit reconciliation file for a switch, type the following command:

```
arfdump [audit reconciliation filename] <Return>
```

The format of [*audit reconciliation filename*] depends on the requirements of your billing system. In general, the filename prefix identifies the type of service (e.g., S for SMDS) and the switch on which the data in the file was collected. The filen extension is in the format S*NNN*, where *NNN* is a sequential number that is incremented by one each time a new reconciliation file is created for a given switch.

For example, to display the second reconciliation file created on the switch that is assigned Switch ID Boston1, type the following command:

```
arfdump SBoston1.S001 <Return>
```

The information is displayed in the same format as AP reconciliation files. The difference is that the statistics apply only to the specified switch (as opposed to being combined statistics from all switches).

Manually Starting the Scheduler

When the system is booted, a scheduler routine is automatically started to schedule the processing of audit reconciliation files. All files are scheduled to be converted at 00:00:00 UCT. This value cannot be changed.

If you notice that daily audit processing is not taking place at the scheduled time, it most likely means that the scheduler is not running. In this event, you can manually start the audit processing scheduler. To do so, log in as the superuser, then type the following command:

```
/opt/CascadeBS/etc/audit.at <Return>
```

Under normal circumstances, you should never have to manually start the scheduler.

Displaying a Summary of Usage Statistics

Ascend's **brfstats** utility enables you to display a summary of billing information for a specified group of billing record files. This utility displays the total PDU and byte counts for each source/destination address combination and group/source/destination address combination catalogued in the specified group of billing record files.

To display a summary of billing usage information, type the following command and press Return:

```
brfstats -i [Switch ID] -s [sequence #] -e [sequence #] -v
```

The switches/parameters for this command are described as follows:

-i [*Switch ID*] — Specifies the Switch ID of the switch for which you want to display information. For example, -i Boston1 indicates that you want information from the switch with Switch ID Boston1. This item is required.

-s [*sequence #*] — Specifies the sequence number of the first file in the group of files for which you want to display information. For example, to display information for the files in the range SBoston1.S001 through SBoston1.S034, you would enter -s 001 in the command. This item is optional; if you do not specify a starting sequence number, 000 is assumed.

-e [*sequence #*] — Specifies the sequence number of the last file in the group of files for which you want to display information. For example, to display information for the files in the range SBoston1.S001 through SBoston1.S034, you would enter -i 034 in the command. This item is optional; if you do not specify an ending sequence number, 999 is assumed.

If the starting sequence number is greater than the ending sequence number, the sequence will "wrap around". That is, if you specify -s 997 and -e 002, the files with the following sequence numbers would be processed: 997, 998, 999, 000, 001, and 002.

-v — This optional switch causes the name of each file to be displayed on the terminal as it is being processed.

A sample **brfstats** display is shown here:

```
brfstats -i Boston1 -s 335 -e 530

Generating summary statistics for SBoston1.S335 to SBoston1.S530 ...

Summary statistics

DA                  SA                  L3 PDU CT      Byte CT
-----------------   -----------------   ------------   ------------
c14039690000ffff    c14046680000ffff         2319324      352537180
c14039700000ffff    c14039710000ffff        70684538    0x511a49ec8
c14039810000ffff    c14039800000ffff         3076058      369008160

GA                  DA                  SA                   L3 PDU CT      Byte CT
----------------    ----------------    ----------------     ------------   ------------
e18005550000ffff    c14039710000ffff    c14039700000ffff         70684555    0x511a4b33c
e18009690000ffff    c14039690000ffff    c14046680000ffff               23            184
```

The items on this screen are destination address (DA), source address (SA), group address (GA), L3 PDU count, and total byte count. All values in the L3 PDU CT and Byte CT columns are shown in decimal. If, however, the L3 PDU count or total byte count values exceeds four bytes (decimal value 4,294,967,296), the value is converted to hexadecimal and preceded by ‘0x’ to indicate it is a hex value. For example, in the sample display, there are two values that have been converted to hex in the Byte CT column.

Starting the apd Daemon

If it appears that the AP daemon, **apd**, is not running, you can use the **apdstart** utility to start the daemon. Normally, you should not have to do this, since apd is started automatically by the system initialization sequence when you boot the system.

To start the **apd** daemon, log in as the superuser, then type the following command:

/opt/CascadeBS/etc/apdstart <Return>

Managing Error Suspense Files

When a usage data record cannot be converted to a billing record, the data is stored in an Error Suspense file until the billing record formatter has the information it needs to convert the record (for more information, see “Error Suspense Files” on page 2-14).

Normally, a cron routine is scheduled to periodically attempt converting these files to billing record files. If you’d like, you can disable the cron routine or manually convert these files at your convenience, as described in the next two sections.

Enabling/Disabling Error Suspense File Processing

If you want to disable the cron routine that schedules Error Suspense file processing, type the following command:

```
/opt/CascadeBS/etc/esfstart stop <Return>
```

If you want to re-enable the cron routine, type the following command:

```
/opt/CascadeBS/etc/esfstart start <Return>
```

Under normal circumstances, you should not need to run either of these utilities. They are provided to give you manual control over Error Suspense file processing (for example, if you want to perform system tests).

Manually Converting Error Suspense Files

If you don’t want to wait for the cron routine to automatically re-execute the billing record formatter on these files, or if you have disabled Error Suspense file processing, you can manually convert them to billing records. When you do so, the billing record formatter re-attempts converting the records in the Error Suspense files to billing records.

There are two ways to manually convert the Error Suspense files. You can convert a specific file, or you can convert all files in the directory.

To convert a specific Error Suspense file, log in as the superuser, then type the following command and press Return:

```
brformatter /opt/CascadeBS/esf/[error suspense filename]
```

where **[*error suspense filename*]** is the name of the file you want to convert.

To convert all files in the Error Suspense file directory, type the following command:

```
sh /opt/CascadeBS/etc/esf.sh <Return>
```

If all records in a given Error Suspense file are successfully converted, the Error Suspense file is removed from the **/esf** directory, and is moved to the **/udfiles** directory. If any record in a given Error Suspense file cannot be converted, the file is not moved to the usage data file directory. In both cases, a new billing record file is created; however, in the case where the billing records could not be converted, the billing record file that is created contains only a header and a trailer record.

Error Detection

In order to support error detection within the billing system, a dedicated logfile is created at installation on the AP. This logfile, **/opt/CascadeBS/cascade.log**, contains a listing of normal and anomalous events that have occurred within the Ascend billing system. For more information, see “Event Log Messages” on page 7-3.

Switching System Management

The B-STDX switches in the billing system are managed remotely via the NMS by running CascadeView/UX. From CascadeView/UX, you can set the billing system attributes for each switch in the system, and enable or disable billing on some or all logical ports that are defined on a switch.

Maintaining and Managing Billing Records on the System

The length of time that billing records are maintained on the system depends on your requirements, the amount of data being stored, and the storage capacity of the AP. You may want to store the records on the AP for a week before deleting them, or you may want to store them for a longer period. Ascend recommends that you set up an archiving process that can be run every night to delete records that you no longer need. For example, if you want to maintain billing records for 30 days, you could set up an archiving process that deletes files older than 30 days, and run this process on a daily basis.

Daily Audit Statistics

In addition to the usage measurements that are sent to the AP at the end of each aggregation period, daily audit statistics are gathered and reported in order to support audit control and data reconciliation. These statistics are gathered over a 24-hour period and reported at the end of the period. The start/end of the period is defined by the Daily Processing Time attribute that is maintained on the switch (via CascadeView/UX).

Table 6-1 lists the daily audit statistics that are maintained on each switch. At the end of the daily period, only the Total Usage Records Created and Usage Records Sent values are transferred from the switch to the AP for billing record generation. Likewise, at the end of the daily period, the AP creates a similar set of statistics. The statistics created by the AP are listed in Table 6-2.

The daily audit statistics (audit reconciliation) files generated by the AP are stored in the audit reconciliation file directory you specified during the billing system software installation. The filenames are in the format shown on page 6-12.

Table 6-1. Daily Statistics Maintained on Each B-STDX Switch

Statistic	Description	Trap[a]
Usage Register Overflows	A count of the number of usage register overflows that have occurred during the current aggregation period.	Yes
Total Usage Register Overflows	The total number of usage register overflows that have occurred during the current day.	No
Billable Usage Events	The number of PDUs processed for billing treatment.	No
Non-billable Usage Events	The number of PDUs not considered for billing treatment.	No
Usage Records Created	The number of usage records created in the current aggregation period.	No
Total Usage Records Created	The total number of usage records created during the current day.	No
Usage Record Creation Failures	The total number of usage records that could not be created during the current aggregation period.	Yes
Total Usage Record Creation Failures	The total number of usage records that could not be created during the current day.	No
Usage Records Sent	The number of usage records sent to the AP during the current aggregation period.	No
Total Usage Records Sent	The total number of usage records sent to the AP during the current aggregation period.	No
Usage Data Store Full	A count of the number of times that the Aggregated Usage Data store became full during the current aggregation period.	Yes
Total Usage Data Store Full	A count of the number of times that the Aggregated Usage Data store became full during the current day.	Yes

Table 6-1. Daily Statistics Maintained on Each B-STDX Switch (Continued)

Statistic	Description	Trap[a]
AP Communication Failures	A count of system failures on the switch that caused Switch-to-Adjunct Processor Transmission failure.	Yes

a. If Yes, an SNMP trap is generated each time the statistic is updated. For more information on SNMP traps, see "SNMP Traps" on page 8-1.

Table 6-2. Statistics Maintained on Adjunct Processor

Statistic	Description	Trap[a]
Usage Records Received	The number of usage data records received from each switch that is being serviced by the AP. This statistic is measured on a per-switch basis.	No
Billing Records Received	The total number of billing records created. This statistic is measured on a per-switch basis	No
Billing Record Creation Failures	The total number of usage data records dropped due to an error condition that occurred during billing record creation. This statistic is measured on a per-switch basis.	No

a. The AP does not generate SNMP traps.

Adjunct Processor Audit Reconciliation File Maintenance

The following counters are maintained on a daily basis by each B-STDX switch in the billing system, and transferred to the AP at the time specified by the Daily Processing Time attribute:

- Number of usage data records created by the switch
- Number of usage data records transferred to the AP

The AP creates two sets of audit reconciliation files on a daily basis. One set summarizes the transactions for the day. The second set of files (one file per switch) summarizes the day's activity on each switch.

For each set of daily audit statistics uploaded from the switches, the AP creates a statistics file that contains the following information:

- Switch ID
- Service Type (currently, only SMDS is supported)
- Date
- Hour (in UCT)
- Switch Time Zone Offset
- Number of Usage Data records created
- Number of Usage Data records transported to the AP

The AP also creates a statistics file that summarizes the day's transactions (as viewed by the AP). The file contains the following information for each switch connected to the AP:

- Switch ID
- Service Type (currently, only SMDS is supported)
- Date
- Hour (in Greenwich Mean Time)
- Time Zone Offset
- Number of usage data records received from each switch
- Number of usage data records dropped by the AP
- Number of billing records created
- Number of billing records transported to the RAO Processing system

Daily statistics are maintained on each B-STDX switch in non-volatile RAM (NVRAM). When the Daily Processing Time is reached, the Data Aggregator closes the current set of statistics, and opens a new set for the next 24-hour period. The closed data is saved in NVRAM until it has been successfully transferred to the AP.

Using CascadeView/UX to Display Statistics

Statistics gathered on each switch (see Table 6-1) can be displayed via CascadeView/UX. From within CascadeView/UX, you can display these statistics for a selected switch. The displayed statistics are for the current aggregation period, as well as totals for the current 24-hour daily statistics period.

To display switch statistics:

1. Log in to the appropriate network map in CascadeView/UX.
2. From the Map view, select the switch for which you want to display statistics.
3. From the Monitor menu, select Cascade Objects ⇒ Show Detail to access the Show Switch Details screen.
4. From the Show Switch Details screen, select the Get Sw Attr... command button. A display similar to the following appears:

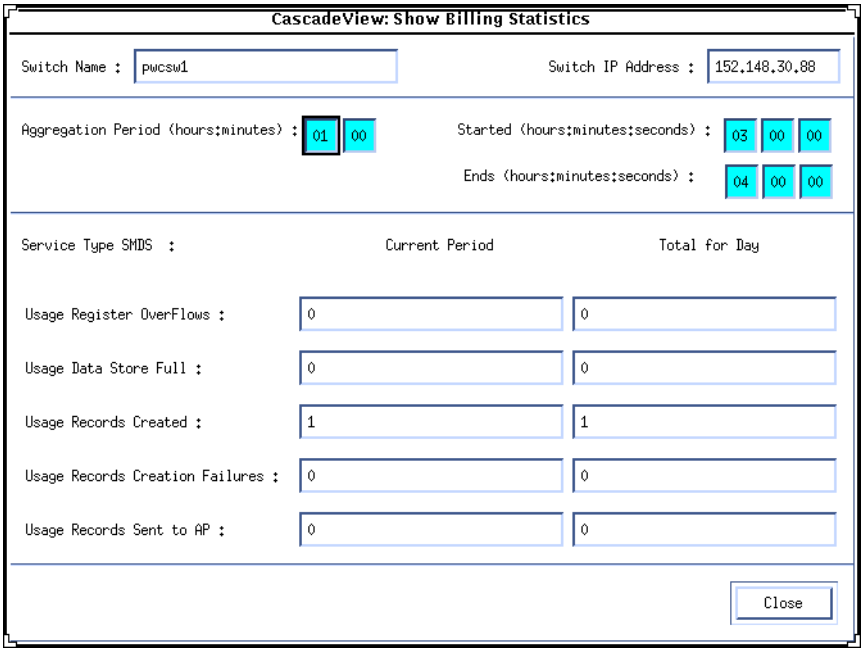

The fields on this screen are described as follows:

Switch Name — The name assigned to the switch on the CascadeView map.

Switch IP Address — The IP address of the selected switch.

Poll Interval (sec) — The frequency with which the Adjunct Processor polls the switch for Aggregated Usage Data. This is *not* the same as the Collection Period, which is defined on each switch.

Aggregation Period — The defined length of the data aggregation period for the switch (see "Aggregation Period" on page 2-6 for more information).

Started — The scheduled start time of the current aggregation period.

Ends — The scheduled end time of the current aggregation period.

Service Type — Currently, SMDS is the only service type supported. In future releases of the software, ATM and Frame Relay will be supported. Statistics will then be displayed for each service type supported on the selected switch.

The following statistics are shown for both the current aggegration period and the current day. The daily totals include all statistics gathered since the most recent Daily Processing Time.

Usage Register Overflows — The number of usage record overflow conditions that have occurred. An overflow condition exists when an attempt is made to update a usage record counter, but the update, if made, would have caused the counter to overflow. In this case, the usage record is closed, and a new one opened (providing there is sufficient space for the new record in the Aggregated Usage Data store).

Usage Data Store Full — The number of times that the Usage Data store reached capacity, therefore causing all data in the store to be transferred to the AP before the end of the current aggregation period. If this value is consistently greater than zero, you should consider using a shorter aggregation period.

Usage Records Created — The number of usage records created on a switch.

Usage Record Creation Failures — The number of times that an attempt to create a usage record failed. Usage record creation failures usually occur when

an attempt to communicate with the AP fails, or when an existing connection to the AP fails.

Usage Records Sent to AP — The number of usage records sent from the switch to the AP.

7

Event Logging

This chapter describes the set of asynchronous events that are generated by each AP in the billing system to indicate anomolous conditions or task completions that have occurred within the billing system.

Overview

Some of the events that are generated are simply informational. Other events indicate that a problem or potential problem may exist in the billing system (either a hardware problem or a configuration problem).

These events are logged in the file **/opt/CascadeBS/cascade.log**. If you seem to be experiencing problems with the billing system, or if you want to monitor billing system activity, you can display the list of billing system events that have been logged to this file. Also, check the SNMP trap events that are listed in CascadeView/UX for each switch (see “SNMP Traps” on page 8-1).

Table 7-1. Billing System Events Generated by the Adjunct Processor

Event	Description
Billing Record Creation Error	An attempt to create a billing record has failed. In this case, the billing information is stored in an Error Suspense file until it can be properly processed. For more information, see "Error Suspense Files" on page 2-14.
Billing Record File Created	This is a task completion message indicating that a billing record file has been created.
General System Failure	When the AP generates this message, it indicates that the AP is not operating normally. In this case, view the **/var/adm/messages** file on the AP to determine what the problem is on your network.

Resetting the Log File

If you want to reset the **/opt/CascadeBS/cascade.log** file, and save the current log file to a new file, log in as the superuser, then type the following command:

```
/opt/CascadeBS/etc/loginit reset <Return>
```

When you enter this command, **cascade.log** is copied to the file **cascade.log.*N***, where N increments by one with each reset. Once the file has been copied, cascade.log is reset.

You can also reset **cascade.log** without copying the existing file to a new file. To do so, log in as the superuser, then type the following command:

```
/opt/CascadeBS/loginit init <Return>
```

Event Log Messages

This section contains a listing of the event messages that are logged in **cascade.log**, as well as the condition that generated the message, possible causes, and suggested resolutions.

All event log messages appear in the following format:

[*timestamp*] [*program*] [*pid*] [*severity*] [*message text*]

These items are described as follows:

timestamp — Indicates when the message was generated, in mm/dd/yy hh:mm:ss format.

program — Indicates the billing system component to which the message applies.

pid — The process ID of **[*program*]**.

severity — The severity level of the message. If no severity level is displayed in the message, then the message is simply informational. Severity may be any of the following:

ERROR	A non-fatal error has occurred
FATAL	A fatal error has occurred
WARNING	Indicates the existence of a non-fatal condition

Messages Generated by the Billing Record Formatter

This section lists the event messages that are generated by the billing record formatter program, **brformatter**. Each message is followed by the condition that generated the message, possible cause(s), and suggested resolutions.

Could not communicate with AP daemon

Condition:
brformatter could not send conversion statistics to the AP daemon, **apd**.

Possible cause:
The **apd** daemon may not be running.

Suggested resolution:
Use the **ps** command to determine if the **apd** daemon is running. If not, log on as root, and use the command **/opt/CascadeBS/etc/apdstart &** to restart **apd**.

Terminating. Cannot recover from error; aborting [*usage data filename*]

Condition:
brformatter could not recover from a previous error; as a result, the listed usage data file (or Error Suspense file) could not be converted to a billing record.

Possible cause:
Refer to the previous message in the log file for a possible cause.

Suggested resolution:
Refer to the previous message in the log file, then perform the suggested resolution for that message.

Could not find environment file brf.cfg

brf.cfg : Syntax error on line [*#*]

Error in brf.cfg; one or more definitions is missing

Condition:
brformatter could not find the **/opt/CascadeBS/etc/brf.cfg** file, or detected a syntax error in this file on the indicated line.

Possible cause:
Either the **brf.cfg** file does not exist, has been moved, or contains incorrect syntax.

Suggested resolution:
Verify that the **brf.cfg** file exists in the **/opt/CascadeBS/etc** directory. If it does, check the file for correct syntax.

Could not open Customer database (CascBilling)

Condition:
An error occurred while **brformatter** was performing a disk access operation.

Possible cause:
Either SYBASE is installed incorrectly on the AP, the database server is down, or the DSQUERY environment variable in **brf.cfg** contains an incorrect value.

Suggested resolution:
Verify that SYBASE has been installed correctly. If it is, verify that the database server is operating. If it is, check the DSQUERY variable in **brf.cfg**, and make sure it is set to the SYBASE SQL server name (normally CASCADE), as defined in the SYBASE interface file, **/opt/sybase/interfaces**, on the SYBASE server.

Error writing record to disk

Error marking record as processed

Could not open usage data file [*usage data filename*]

Could not create new OSCIR record file for [*usage data filename*]

Condition:
An error occurred while **brformatter** was performing a disk access operation.

Possible cause:
The disk is full, or the disk is corrupted.

Suggested resolution:
Check the integrity of the disk on which billing record files are being stored.

[#] unprocessed records in [*usage data filename*]

File [*usage data filename*] has been scheduled for Error Suspense processing

Condition:
The CascBilling customer database is not consistent with the B-STDX switch.

Possible cause:
Either a switch definition is missing from the SwitchInfo table, or one or more E.164 customer addresses is missing from the SmdsCustomer table. If this condition exists, the specified usage data file is moved into the Error Suspense directory **/opt/CascadeBS/esf**, and will be reprocessed.

Suggested resolution:
Verify that the Provisioning System is scheduled to update the CascBilling database.

Messages Generated by File Transfer Process

This section lists the event messages that are generated by the AP file transfer process, **apfts**. This process manages and performs usage data file transfers between each switch and the AP.

Connection to [*switch IP address*] terminated unexpectedly

Failed to send acknowledgment to [*switch IP address*]

Failed to receive file from switch [*switch IP address*]

Condition:
The file transfer protocol between the AP and the B-STDX switch whose IP address is listed in the message is failing.

Possible cause:
Network congestion or a communication failure is disrupting the file transfer process.

Suggested resolution:
Check CascadeView/UX to determine if the indicated switch is still reachable. Then 'ping' the switch from the AP to determine if the switch is reachable from the AP.

Rejected file from [*switch IP address*]: File integrity was bad

Condition:
The file integrity of a usage data file sent from the indicated switch is bad. When this condition occurs, **apfts** deletes the file, and the switch attempts to send it again.

Possible cause:
A disk write error may have occurred.

Suggested resolution:
If this error continues to occur, check the integrity of the disk on which usage data files are being stored.

Could not create file [*usage data filename*]

Error writing file from switch to disk [*switch IP address*]

Condition:
apfts could not create the indicated usage data file, preventing the transfer of the usage data file from the indicated switch to the AP.

Possible cause:
The disk on which you are storing usage data files may be full.

Suggested resolution:
Check the integrity of the disk on which usage data files are being stored.

Failed to start formatting program brformatter

Condition:
apfts could not execute **/opt/CascadeBS/bin/brformatter** to convert the newly-received usage data file to the appropriate billing record format.

Possible cause:
Either the integrity of the disk is bad, or **brformatter** has been deleted or moved from the **/opt/CascadeBS/bin** directory.

Suggested resolution:
Verify that **brformatter** exists in the **/opt/CascadeBS/bin** directory, and that it can be executed.

8

SNMP Traps

This chapter describes the SNMP trap events that are generated by the B-STDX switches and sent to each NMS in the switching system that is enabled to receive traps.

A list of currently-logged SNMP trap events can be displayed at any time on any NMS connected to the switching system, using CascadeView/UX. Some events are simply informational, while other events indicate a problem or potential problem within the billing system or the switching system on which it is running. Events are categorized based on the standard HP OpenView Event Categories: *Critical*, *Major*, *Minor*, *Warning*, or *Normal*.

All SNMP trap messages are preceded by the timestamp (in mm/dd/yy hh:mm:ss format) and the type of trap (e.g., LPort Status, Node Status). The timestamp indicates when the trap message was generated. The trap type indicates the network component that generated the trap event; for example, Node Status traps are generated by B-STDX switches, and LPort Status traps are generated by logical ports on a B-STDX switch.

Normal Events

This section describes the *Normal* SNMP trap events that are generated by the billing system. These events are simply informational. They do not indicate a problem or potential problem within the billing system.

Billing has been (*enabled*/*disabled*) for service SMDS

This is a node-type trap signifying that the billing system state on a switch has been changed, either from enabled to disabled, or from disabled to enabled. The billing state is changed when a switch boots (billing becomes enabled), or when the state is changed manually via CascadeView/UX.

Billing has been (*enabled/disabled*) for service SMDS on [*LPort ID*]

This is an lport-type trap signifying that the billing system state on the indicated logical port has been changed, either from enabled to disabled, or from disabled to enabled. The billing state is changed on a logical port when the switch boots (billing becomes enabled), or when the state is changed manually via CascadeView/UX.

Anomalous Events

This section lists the trap event messages that signify an anomolous event has occurred within the billing system. Preventative measures or resolutions to these events are also provided. If any of these events appear on the CascadeView Events screen for a switch, you should also display the **/opt/CascadeBS/cascade.log** file on the AP, as this file contains information that will be helpful in diagnosing the cause of the anomolous event.

A usage record counter-value overflow condition has occurred for service SMDS

This is a *Warning* node-type trap signifying that a usage data counter in a usage data record is about to overflow, or has overflowed, on the indicated switch. Normally, when a usage record is near overflow, the record is closed and a new record is opened; in this case, this trap event message is not generated. If, however, the new usage data record cannot be created, this message is generated, along with a Usage Record Creation Failure message (see the next section).

The generation rate of this message is controlled such that only one of these traps is generated per Collection Period. For example, if your Collection Period is set to 30 seconds, and there are more than one Usage Record Creation Failures, only one Usage Record Overflow Warning (and Usage Record Creation Failure) message is generated. Also, if more than one usage data register is in danger of overflowing, only one Usage Record Overflow Warning is generated.

If the condition still exists when the next Collection Period starts, another Usage Record Overflow Warning and Usage Record Creation failure is generated at that time.

For information on how to resolve the situation, see the next message description.

A usage record could not be created for service SMDS

This is a *Major* node-type trap signifying that a Usage Data record could not be created due to the Aggregated Usage Data store being at capacity. This should occur only when the secondary store has reached capacity, and the primary store has not been transferred to the AP (possibly due to a communication failure with the AP). This trap is not generated when the system is in a normal state.

As with the Usage Record Overflow Warning, the generation rate of this message is controlled such that only one of these traps is generated per Collection Period.

If this message appears on the Trap Event screen, check the AP to see if it is on-line; if it is, check to see if the AP is connected to and communicating with the switching system.

Billing communications failure to adjunct processor [*IP address*]

This is a Critical node-type trap issued by the indicated switch when the switch cannot communicate with the AP, and is unable to transfer usage data records to the AP. This message is generated when a connection cannot be made to the AP, or when an existing connection fails.

If this message appears on the Trap Event screen, check the AP to see if it is on-line; if it is, check to see if the AP is connected to and communicating with the switching system.