SAS® IT Service Level Management 2.1

User's Guide

The Power to Know®

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2006. *SAS® IT Service Level Management 2.1: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® IT Service Level Management 2.1: User's Guide

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

2nd printing, April 2006

SAS Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at **support.sas.com/pubs** or call 1-800-727-3228.

Contents

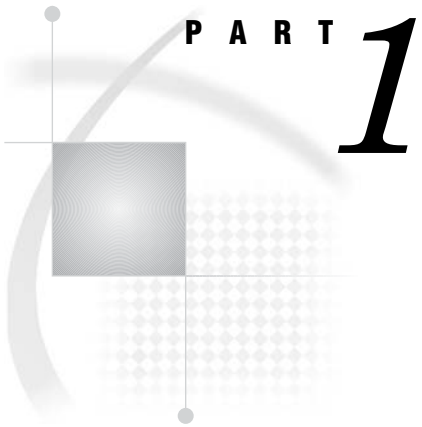

PART 1

Introduction to SAS IT Service Level Management

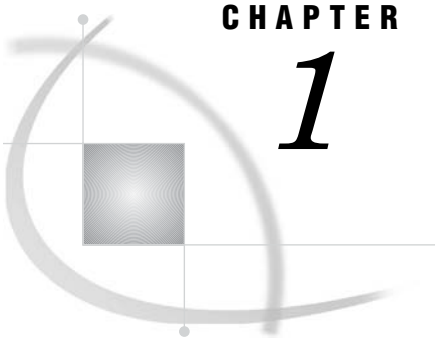

CHAPTER

1

Introduction to SAS IT Service Level Management

Introduction to SAS IT Service Level Management

The *SAS IT Service Level Management: User's Guide* consists of the following four parts:

1. This introduction, which includes these chapters:
   - "Getting Started" on page 3
   - "Understanding Contracts" on page 4
2. "Administration" on page 4
3. "Reporting" on page 5
4. "Appendixes" on page 7

This document assumes that you have reviewed *Getting Started with SAS IT Service Level Management*.

Getting Started

Information about how to get started with this solution is contained in the document *Getting Started with SAS IT Service Level Management*. You can view and download a copy of that document from the SAS Support Web page **`http://support.sas.com/documentation/onlinedoc/itslm`**.

The *Getting Started* document discusses and illustrates Service Level Management terminology. In addition, that document covers the following topics:

- Chapter 1
  - Introducting SAS IT Service Level Management
  - What Is Service Level Management?
  - SAS IT Service Level Management Software Overview

- Chapter 2
  - Overview of IT Service Level Calculations
  - Understanding the Structure of Your Contract
  - Service Level Objectives
  - Data Adjustments for Outages
  - Extract, Transform, and Load (ETL)
- Chapter 3
  - Overview: Getting Ready to Run SAS IT Service Level Management
  - Starting the Graphical User Interface (GUI)
  - Exploring the GUI
  - Using the GUI to Construct a Contract
  - Running ETL
  - Closing SAS IT Service Level Management
- Chapter 4
  - Overview of Reporting
  - Baseline Reports
  - Service Level Reports
  - SAS Enterprise Guide Reports

Understanding Contracts

Information about understanding contracts is contained in Chapters 2 and 3 of this user's guide:

- "Introduction to Understanding Contract Structure" on page 9 discusses the organization of a contract and how that organization is represented by the four types of nodes within a hierarchy. The four types of nodes are: contract nodes, service level agreement nodes, service nodes, and component nodes.
- "Introduction to Understanding Contract Objectives" on page 13 discusses the five types of objectives (availability, response time, throughput, custom, and composite objectives), and how these objectives can be applied to a component node, a service node, a service level agreement node, and a contract node.

Administration

The administrative functions provided by SAS IT Service Level Management are contained in Part 2 of this user's guide.

To use the administrative graphical user interface (GUI), you must first log on to the server and access the environment database that you want to work with. Then, you can access the five specialized workspaces provided by SAS IT Service Level Management. (There is one workspace for each group of information.)

You can use the GUI to perform administrative tasks such as viewing, adding, changing, and deleting information.

Note: When logging on and using the GUI, you might notice that some commands take a few moments to display a result because of the amount of data retrieval and processing that must be performed. For example, a progress indicator might be displayed when you attempt to access a data source PDB that is not available or when you access the data about a component's resources from the data source PDB. Tasks

such as the Validate, Synchronize, Run ETL, Run Reports, and Reconnect are marked with progress indicators such as an hourglass or a dialog box to inform you that SAS IT Service Level Management is working. △

The following chapters constitute the Administration part of the *User's Guide*:

- "Introduction to Logging On to the Software" on page 37 discusses how to invoke SAS IT Service Level Management, the environments that you can work in, and other topics, such as user names and passwords and how to work in the main window of the GUI.
- "Introduction to Accessing Data" on page 43 has information about the location of performance data for the items in your catalog. You can use the Data workspace to access and view the data. (You can also use batch jobs to access and view the data.)
- "Introduction to Recording Participants" on page 53 has information about customer and provider organizations, and the contact people at those organizations and elsewhere. You can use the Participants workspace to enter information about customers, providers, and contacts. You can then use this information in the contracts.
- "Introduction to Building the Catalog" on page 61 has information about service level agreements, services, components, range sets, and range descriptions. You can use the Catalog workspace to enter information about the entities in your catalog. You can then use the entities as building blocks in the contracts.
- "Introduction to Defining a Contract" on page 81 has information about contracts. You can use this workspace to enter contracts. You can also use the Contract workspace to create batch jobs that process and report on the data for those contracts. The reports enable you to determine whether contract objectives are being met.

Reporting

You can produce SAS IT Service Level Management reports by using any of the following methods:

- generate reports interactively from the Reports workspace
- generate and then run the report job code either in a scheduler program (in a batch) or in a SAS session
- create and then run a SAS Enterprise Guide project

Note: SAS IT Service Level Management provides views of the contract database that make it easier for customers to write their own custom reports using SAS Enterprise Guide. These views are located in the ADMIN level of the contract database. For more information about this feature, see "Accessing Views of the Contract Database" on page 237.

SAS IT Service Level Management also provides tables of metadata about the contract database. You can use the metadata to navigate the contract hierarchy and its associated objects. This makes it easier for customers to write their own custom reports using SAS Enterprise Guide, Structured Query Language (SQL), or other report writing tools. The tables are located in the ADMIN level of the contract database. For more information about this feature, see "Accessing the Metadata of the Contract Database" on page 239. △

The Reporting part of this user's guide contains these chapters:

- "Overview of Reports for Service Level Management" on page 93 provides an overview of baseline reports, service level reports, and the reports that you can create by using SAS Enterprise Guide software.

- "About Baseline Reports" on page 97 discusses how to generate and run the job code that creates a Web-based gallery of baseline reports for a specified contract, and how to view the baseline reports.
- "About Service Level Reports" on page 110 discusses the reports that are displayed by your browser as a Web page. These reports include from one to three tabs (as determined by your report specification during setup): Contracts, Compliance, and Servicing Contacts. It also describes how to specify, run, and then view these reports.

  Note: A Resource Analysis report and an optional 5–Minute Interval Table are also available and can provide details about your contracts. △
- "Introduction to Service Level Reports for Contract Providers and Customers" on page 144 provides information about service level management reports that can be generated for contract providers and customers.
- "Introduction to Working with SAS Enterprise Guide to Create SAS IT Service Level Management Reports" on page 187 provides information about using SAS Enterprise Guide to generate reports for service level management contracts.

  For an overview that discusses how to create *ad hoc* reports about the service level contracts, see "SAS Enterprise Guide Reports" on page 6.

SAS Enterprise Guide Reports

The sample SAS Enterprise Guide project that is included in the Samples/EGProject directory provides several predefined graphic reports on indicators, scores, and targets. The project was created using sample data, and only a few sample reports were produced; however, the project can be customized to produce a wide variety of reports on the data in your own contract database.

Copy the SAS Enterprise Guide project and the finance sample contract database to your computer from the samples.zip file of the SAS IT Service Level Management client installation CD. You can then open the sample project and view the supplied SAS Enterprise Guide reports. (For steps to transfer the necessary files to your computer, refer to the installation instructions that are provided in the SAS IT Service Level Management client installation CD.)

The SAS Enterprise Guide sample reports include the following:

- detail reports
- daily reports
- weekly reports
- monthly reports

A brief overview of the SAS Enterprise Guide project is provided in "SAS Enterprise Guide Reports" in *Getting Started with SAS IT Service Level Management*.

More detailed information about the SAS Enterprise Guide project and sample reports is provided in "Introduction to Working with SAS Enterprise Guide to Create SAS IT Service Level Management Reports" on page 187.

Note: SAS Enterprise Guide reports can be published on the Web. For more information, see the SAS Enterprise Guide online help. △

Appendixes

This *User's Guide* provides the following appendixes:

- "Calculation Methods" on page 215 discusses the calculation methods that are supplied with the SAS IT Service Level Management software. These include the maximum, mean, minimum, sum, percentile, redundant, serial, and weighted mean methods. This appendix also includes a topic about missing values.
- "Custom Code for the Calculation Method" on page 221 discusses the application program interface (API), its inputs and output, the topics of filtering and weighting, and various constraints that affect what you code.
- "What is an Outage?" on page 231 describes the types of outages that can affect the data and discusses how overlapping outages are prioritized. This appendix also provides information about outage start and end times.
- "The Contract Database" on page 237 discusses the availability, response time, throughput, and custom objective tables of detailed data and how they are used to summarize the daily, weekly, monthly, and yearly data for a component, service, and service level agreement. This appendix also describes the views and the tables of metadata that are created in the ADMIN level of the contract database.
- "SAS IT Service Level Management Macros" on page 266 provides complete reference documentation for these macros:
  - %CPSLMENV
  - %CPSLMCON
  - %CPSLMETL
  - %CPSLMRPT
  - %CPSLMFND
- "The ETL Process" on page 287 discusses the Extract, Transform, and Load (ETL) program.
- "Sample E-mail Notification" on page 293 displays an example of the notification that is sent when any of the servicing contact's objectives have missed their targets.
- "Batch Jobs" on page 297 discusses how to make the ETL and report jobs and then run the ETL and report jobs. In addition, this appendix discusses how to back up your databases. It also provides information that applies specifically to a z/OS operating environment.
- "Backloading Data" on page 309 discusses how to backload data from an archived PDB into the contract database.
- "Best Practices" on page 313 provides advice about topics such as defining and managing contracts and your environment.

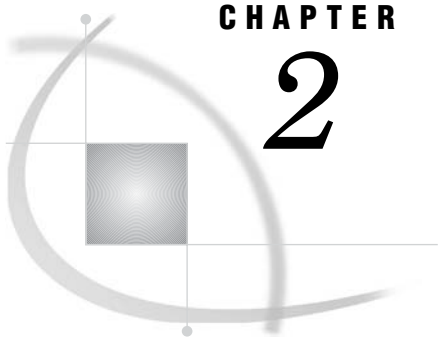

CHAPTER

2

Understanding Contract Structure

Introduction to Understanding Contract Structure

The contract structure in SAS IT Service Level Management is a representation of the elements of the actual agreements between the provider and the customer. Some of the elements, referred to as components, record data; other elements, referred to as services and service level agreements (SLAs), are used to model business processes, based on the requisite underlying IT components.

Before you start to use the software, it is a good idea to understand the organization of a contract. The organization is represented by the nodes and the hierarchy.

This chapter discusses what nodes are, and how they are related to the contract's hierarchy.

The next chapter discusses objectives, which are the most important features of nodes.

Understanding Nodes

There are four types of nodes: contract nodes, service level agreement nodes, service nodes, and component nodes.

- Each contract node is the highest-level node in the contract hierarchy. It represents a contract as a whole. Thus, each contract has one contract node.

  The contract contains basic, high-level information about the customer and the provider.

- Each service level agreement node represents an agreement between an IT provider and a customer that defines the services that are provided, the metrics that are associated with these services, and the acceptable and unacceptable service levels. A service level agreement (SLA) covers one or more business services that can be packaged together.

  Each contract can have zero, one, or more service level agreement nodes.

- Each service node represents one or more IT systems that typically enable a business to function. Services can be simple and low-level, such as offering a service for disk space. They can also be high-level and complex, such as offering a Web application hosting service that is composed of numerous subsidiary services,

such as disk space service, Web server service, network service, and help desk service.

Each contract can have zero, one, or more service nodes.

- Each component node represents one or more tangible IT hardware and software assets, such as a Web server or disk pool. Components can be low-level, such as an individual router, or high-level, such as an entire LAN/WAN.

  For each component node that you define, you must create a list of resources. The resource list is used to subset the data from the data source PDBs.

  Each contract must have at least one component node.

Note: For more information about nodes, see Chapter 2: "Transforming IT Data into SLM Intelligence" in *Getting Started with SAS IT Service Level Management* on the SAS Support Web page **`http://support.sas.com/documentation/onlinedoc/itslm`**. △

Understanding Hierarchies

Note: For information about the contract hierarchy and the child nodes that can exist in the hierarchy, see the topic "Understanding the Structure of Your Contract" in Chapter 2 of *Getting Started with SAS IT Service Level Management*. You can access this document from the SAS Support Web page **`http://support.sas.com/documentation/onlinedoc/itslm`**. △

The organization of a contract's nodes is called the contract's hierarchy. The minimum hierarchy is as follows:

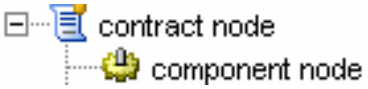

A more typical hierarchy has some intermediate nodes. For example, it might look like this:

Or it might look like this:

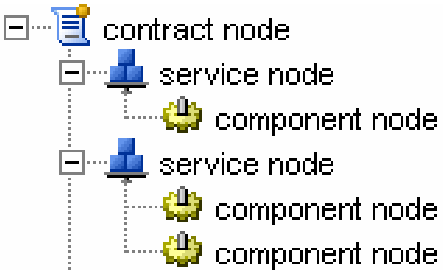

Or it might look like this:

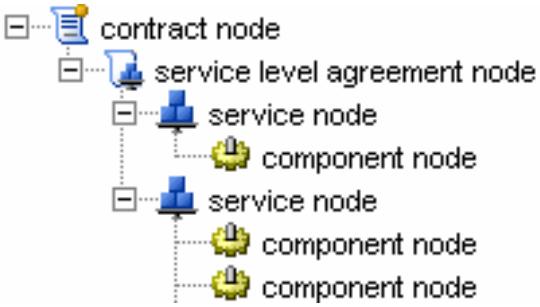

Notice that all three hierarchies are essentially the same. One site might structure its contracts by using service level agreements, another site by services, and a third site might use both service level agreements and services.

The above examples are abstract in order to describe the "shape" of a hierarchy. If you think of a hierarchy as a tree, then the following is true:

- The node that represents the top of the tree must be a contract node.
- The nodes that represent the leaves of the tree must be component nodes.
- The nodes, if any, that represent the branches of the tree can be either service level agreement nodes or service nodes or both.
  - If you identify a node as a service level agreement node, it must branch directly from the contract node.
  - If you identify a node as a service node, it can branch from the contract node, from a service level agreement node, or from another service node.

Note: When you process data for a contract, the processing job takes account of the contract's hierarchy. For more information about how the processing job (also known as the ETL job) takes account of the hierarchy, see the overview in the ETL section of Chapter 2 in *Getting Started with SAS IT Service Level Management* or see the details in "The ETL Process" on page 287. △

In order to show what a hierarchy looks like in a more practical sense, here are two concrete examples:

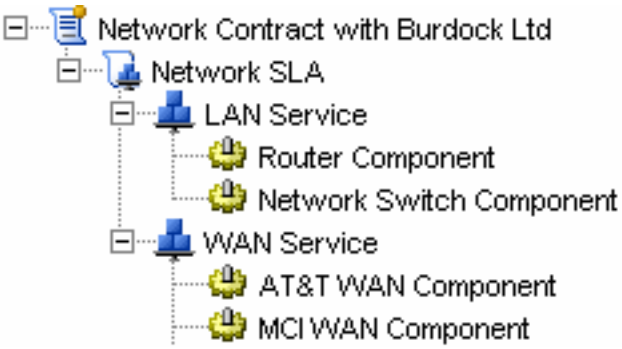

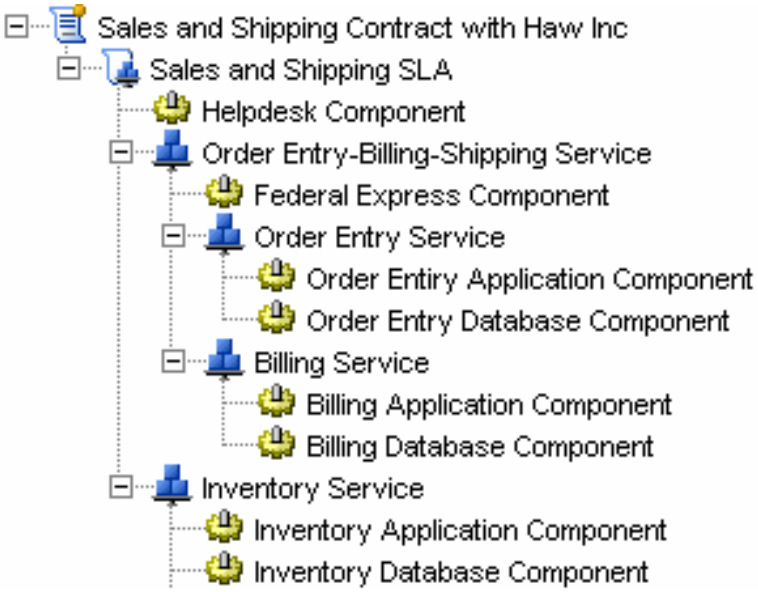
Sales and Shipping Contract with Haw Inc
Sales and Shipping SLA
Helpdesk Component
Order Entry-Billing-Shipping Service
Federal Express Component
Order Entry Service
Order Entiry Application Component
Order Entry Database Component
Billing Service
Billing Application Component
Billing Database Component
Inventory Service
Inventory Application Component
Inventory Database Component

CHAPTER 3

Understanding Contract Objectives

Introduction to Understanding Contract Objectives

Nodes have objectives. Objectives enable the contract provider and contract customer to determine whether the goals of the contract are being met. Before you start to use the software, it is a good idea to understand how objectives work.

This chapter contains these sections:

- "Understanding What Is Common to All Objectives" on page 14 discusses what objectives are and what all objectives have in common.
- "Understanding More about Objectives in Component Nodes" on page 24 discusses what is unique about the objectives in component nodes.
- "Understanding More about Objectives in Higher-Level Nodes" on page 31 discusses what is unique about the objectives in service, service level agreement, and contract nodes.

Understanding What Is Common to All Objectives

An objective describes the behavior of one aspect of the node, and it describes the goal for that behavior. Each node has three specialized objectives, one custom objective, and one overall (or composite) objective.

A node's three specialized objectives are called the availability objective, the response time objective, and the throughput objective.

- The availability objective has information about the node's availability and your goal for the node's availability.

  Availability is related to the percentage of time that the node is available for use by customers.

  For example, when a Web server is able to service requests, it is available.

  Typically, if an IT organization uses only one objective (or starts by using only one objective, with the intention of using others later), it uses an availability objective. Availability is the most important objective because nodes must be available before they have response times and throughput.

- The response time objective has information about the node's response time and your goal for the node's response time.

  Response time is the time required for the node to service a request, that is, the amount of time between the submission of a request for service and the user's receipt of the system's reply to that request.

  For example, a Web server request with a one-minute response time would be considered slow, and a help desk request with a one-minute resolution time would be considered fast.

- The throughput objective has information about the node's throughput and your goal for the node's throughput.

  Throughput is related to the number of units of work (such as packets, transactions, I/O requests, or jobs) that can be processed by that node in a given amount of time.

  For example, seven transactions per minute might be considered a heavy load for a help desk system and a light load for a Web order-entry system.

The custom objective has user-specified information about the node's performance and your goal for the node's performance.

- The custom objective is a general purpose objective that can be used to manage objectives that cannot be categorized into the standard availability, response time, and throughput objectives. Or it can be used to record a second value for availability, response time, or throughput on the same node.

  For example, the custom objective can be used to measure customer satisfaction or mean time to failure.

A node's overall objective is called the composite objective.

- The composite objective is based on a combination of the other objectives (usually, two or more of the other objectives). Thus, the composite objective has information about the node's overall behavior and your goal for the node's overall behavior.

  For example, a print server's performance might be considered excellent overall if it is "up" almost all the time *and* handles a heavy load.

The following subtopics, described below, apply to the availability, response time, throughput, and custom objectives:

- input metrics
- indicators
- scores
- targets
- range sets, range descriptions, and scores

The final subtopic describes the composite objective.

The values that are described in these subtopics can be specified on the **`Objectives`** tab of the node.

Input Metrics

The input metric defines the expected form of the input that will be used by the system for calculating the input data either from resources (that is, for components) or from child nodes in the hierarchy (that is, for contract, SLA, and service nodes). From the **`Input Metrics`** field of the GUI, the following choices are available:

- The input metric for a component is Measures.
- The input metric for a service, SLA, or contract node can be either Indicators or Scores.

Note: The input metric for any node can be defined by means of Custom Code. △

Indicators

An indicator is a numeric value that represents the performance of a node that is being measured. Each node in a contract can have five objectives: availability, response time, throughput, custom, and composite. Each objective has an indicator that represents the node's performance.

Each component, service, or SLA node must have at least one objective that is activated. The contract node is not required to have any objectives activated; this allows the contract to act as a container for a collection of SLAs, services, or components.

Note: The ETL job calculates each indicator for each five-minute interval and stores the information in the contract database. △

The node's availability indicator, response time indicator, throughput indicator, and custom objective indicator are based on data that the node obtains from the lower levels in its hierarchy. (The composite objective does not use an indicator to represent its performance. Instead, its performance is based on the scores of the indicators of the one or more objectives that make up the composite objective.)

- For a *component* node, the input metric to an objective's indicator is one value (or measure) for each of the selected resources. The values (or measures) for the selected resources are extracted and derived from the objective's data source PDB. For example, the resources for a router component might be individual routers; the resources for a network switch component might be individual switches; the resources for a help desk component might be individual people who are providing the help; and so on.

  Note: A data source PDB might have data about more resources than the ones that are relevant to a given component node. You specify which resources are relevant to the given component node, and the data from all other resources is ignored when the ETL job processes data for this node. △

Thus, each indicator (if activated) gets one value (per five-minute interval) from each of the selected resources in the objective's data source PDB for that objective.

Note: A single data source PDB can provide input to all of these component indicators. Or you can have a different data source PDB for each indicator. Or you can use any combination in between. △

□ For a *service node, service level agreement node,* or *contract node*, the input to the corresponding objective's indicator is one value from the same type objective on each of the node's child nodes. The *child nodes* are the nodes that are directly connected to the node in the contract's hierarchy. For example, look at the **`Expense Accounting`** SLA of this contract hierarchy :

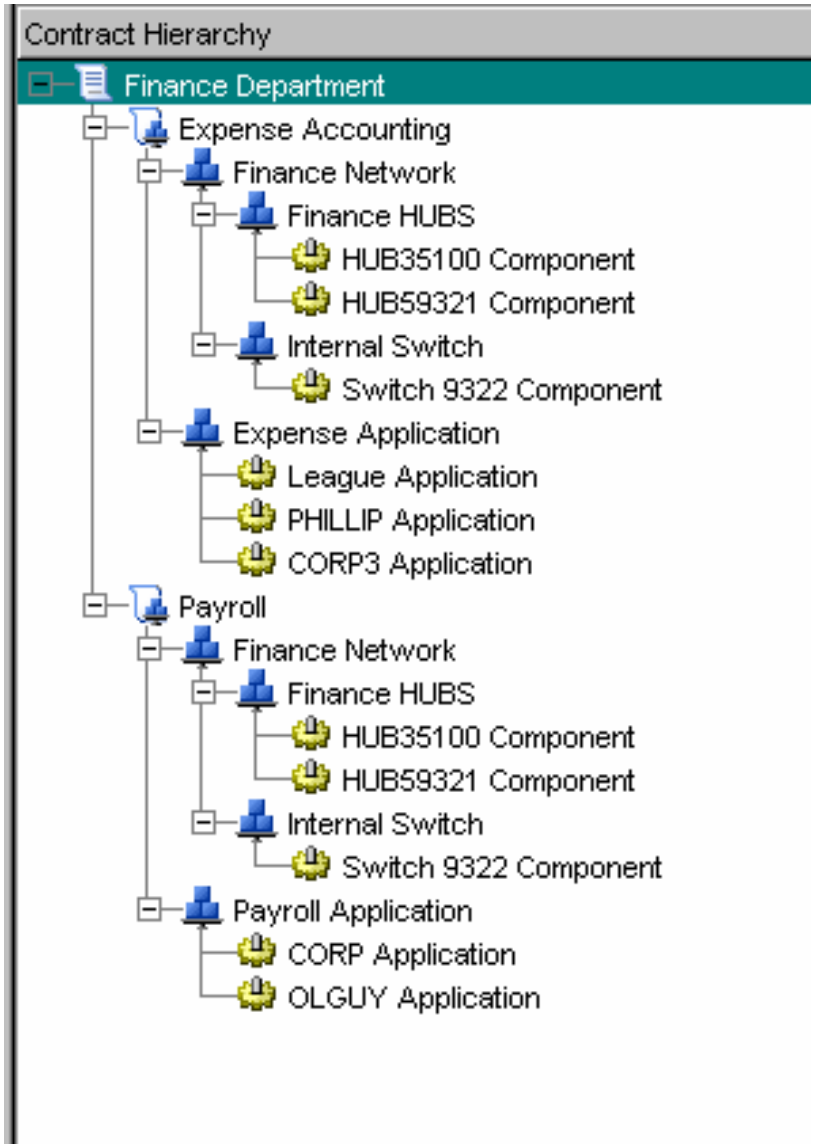

The Finance Department contract node has two SLA child nodes (Expense Accounting and Payroll)

The Expense Accounting SLA node has two service child nodes (the Finance Network and the Expense Application).

The Finance Network service node has two other service nodes (Finance HUBS and Internal Switch).

The Finance HUBS service node has two component child nodes (HUB35100 Component and HUB59321Component).

The Internal Switch Service node has one component child node (Switch 9322 Component).

The Expense Application service has three component child nodes (League Application, PHILLIP Application, and CORP3 Application).

Thus, the availability indicator gets one value from the availability objective on each of the node's child nodes, and so on for each of the indicators for the response time, throughput, and custom objectives. The indicator's one value is based on five-minute interval data from its child nodes.

Calculation Methods

From the indicator's input metrics, the ETL process calculates the indicator. You specify the calculation method that the ETL process is to use. You can choose one of the supplied calculation methods or you can define your own custom calculation method.

- Maximum, mean, minimum, percentile, and sum are supplied calculation methods for all objectives. Redundant, serial, and weighted mean are supplied calculation methods for some objectives. The calculation method that you select is applied to the input values (one from each resource, or one from each child node). For example, if you select mean, the indicator is the arithmetic mean of the input values.

For more information about calculation methods, see "Calculation Methods" on page 215.

Scores

A score is a numeric value that is derived by comparing an indicator to a previously defined range set. Scores are useful when the input data is measured in non-comparable units, for example, seconds and percentages. The availability, response time, throughput, custom, and composite objective types can be converted into a unitless value called a score by using range sets. These scores can then be combined to give an overall view of how a node is performing. The composite objective is used to combine scores for a given node. (Any objective that contributes to a composite objective is required to provide a score for the selected objective.) Similarly, the scores can be passed up the hierarchy to higher level nodes to be used in grading a level of service for a specific objective type.

For information on how to use range sets to apply scores to indicators, see "Range Sets and Range Descriptions" on page 20.

Targets

A target can be defined for each objective. A target can be based either on an indicator or on a score.

An objective's indicator or score represents the behavior for that objective. For each indicator or score, you can optionally define the objective's target, which represents at least one of your goals for the performance of the node. Thus, an availability objective has a target based on the statistical summary of either the availability indicator or the availability score. Similarly, the response time, throughput, custom, and composite objectives have targets based on their respective indicators or scores. For guidelines that can help you decide whether to use indicator-based or score-based targets, see "Selecting the Type of Target: Indicator-Based or Score-Based" on page 19.

You can use the GUI to define a target by specifying this information on the **Objectives** tab of the node:

- Compliance period, which specifies the time period that is the focus of the contractual agreement between the customer and the contract provider and that is displayed in the reports.
- Target statistic (for indicator-based targets only), which specifies what statistic (mean, minimum, maximum, or sum) is applied to the summary of indicator data to date.
- Comparison value, which specifies whether the target value for the objective is to be compared with either the indicator or the score.
- Operator, which specifies the operator to be used in the comparison of the Indicator/Score to the Target Value.
- Target value, which specifies the target for an objective.
- Label (optional), which enables you to personalize your report labels.

You can enter the following values for these fields:

- The compliance period can be *daily*, *weekly*, *monthly*, or *yearly*.
- The target statistic can be *mean* or *weighted mean*, *maximum*, *minimum*, or *sum*. The target statistic is not available for the composite objective. It is also not available for objectives whose comparison value is Scores.

  For all nodes, you can specify *weighted mean* for the target statistic only if you have specified *weighted mean* for the calculation method. (The weighted mean calculation method is not valid for availability objectives.) If you have specified *mean* for the calculation method, then you cannot specify *weighted mean* for the target statistic.

  - *For component nodes*:

    The weighted mean calculation requires that you specify a Weight By variable on the **Data** tab of the component node. *If no Weight By variable is specified for a component, or if the Weight By variable cannot be located in the datasource PDB, then the validation step of the ETL process will fail.*

  - *For the higher-level nodes (service, SLA, or contract)*:

    If you specify the weighted mean, then the weighted mean calculation will use the weights from the child nodes for this specific objective type. *If the child nodes do not provide weights, then the indicator will have a value of missing.*

  Note: If you specify the Weight By variable, you are *not* required to select the weighted mean calculation method.

  The weighted mean calculation method (and, therefore, the weighted mean target statistic) is available only for the response time, throughput, and custom objectives. It is not an option for availability objectives. △

- The operator can be >= (greater than or equal to), <= (less than or equal to), = (equal to), > (greater than), < (less than), or != (not equal to).
- The label has no effect on the target value. The label might be % for an indicator-based availability target, milliseconds for an indicator-based response time target, or transactions per second for an indicator-based throughput target. For example, if the target value is in seconds and the label is in minutes, the value is not converted. The label does not necessarily reflect the value of the target.

Examples: Specifying Indicator-Based or Score-Based Targets

Example #1: Specifying an indicator:

If your goal for the objective's indicator is that its monthly mean (that is, the mean of the five-minute-interval indicators for the month) be greater than or equal to 20, then use the GUI to make these choices in the **Objectives** tab of the node:

1. For the compliance period, select Monthly.
2. For target statistic, select Mean.
3. For the comparison value, select Indicator (for this objective).
4. For the operator, select >=.
5. For the target value, enter 20.
6. For the label, enter descriptive text of your choice. (Your text here is treated as a comment only.)

Then, if the monthly mean of the indicator is 24, the objective has met its target. And if the monthly mean of the indicator is 18, the indicator has missed its target.

Example #2: Specifying a score:

You can specify the goal that the score of the monthly indicator be greater than or equal to 3. Use the GUI to make these choices in the **Objectives** tab of the node:

- For the compliance period, select Monthly.
- For the comparison value, select Score (for this objective).
- For the operator, select >=.
- For the target value, enter 3.
- For the label, enter descriptive text of your choice. (Your text here is treated as a comment only.)
- The target statistic is not available for score-based targets.

Then, if the monthly score is 4, the objective has met its target. And if the monthly score is 2, the objective has missed its target.

Selecting the Type of Target: Indicator-Based or Score-Based

On each objective, if you want to specify a target, you can specify it as based upon either an indicator or a score. (You cannot specify the target as both an indicator and a score.)

For each objective you need to decide which type of target is more important to you and to the people who use the contract reports. Here are some factors that you might want to consider:

- The people who are interested in the reports about the component nodes might be more interested in the indicator (which is a precise value) than in the score (which represents unitless values).

  Thus, for a component node you might want to specify indicator-based targets rather than score-based targets.
- The people who are interested in the reports about the contract node might be more interested in the name of the score (which assigns a clear meaning) and the overall score than in the indicator itself.

  Thus, for a contract node you might want to specify score-based targets rather than indicator-based targets.
- For service nodes and service level agreement nodes, either target type might make sense. You might want the nodes closer to the contract node to use score-based targets and nodes closer to the components to use indicator-based targets.
- Although a node's objectives are not required to have only indicator-based targets or have only score-based targets, it might be easier to interpret reports if all of the

node's availability, response time, throughput, and custom objectives use the same kind of target.

Note: Scores can be defined for all five types of objectives (availability, response time, throughput, custom, and composite) by using a range set with scores on each of the ranges. However, scores are not required. △

Range Sets and Range Descriptions

A range set can be specified for each objective. A range description must be specified for each range. These can both be created in the Catalog workspace.

If you specify a range set for an objective, then reports about that objective will include not only the indicators and the scores (if they are defined in the range set) but also the elements of the range descriptions that are associated with the ranges in the range set, namely the color for text reports and the label to be used in the reports. For all target statistics except SUM, the colors and labels in the reports are applied and displayed for all summary periods. If the target statistic is SUM, then the colors and labels in the reports are applied and displayed for the selected summary period only and all other summary periods are displayed using the transparent color.

Range Sets

A range set consists of a set of ranges. Each set of ranges consists of two types of ranges: ranges that you specify and ranges that are predefined.

- *Ranges that you specify*:

  Each range has an upper and lower limit that you specify. You can (but are not required to) specify a score for the range. You are required to specify range descriptions for each range in the range set. Range descriptions specify the color and the label that will appear in the text reports for this range.

  Note: Gaps or overlaps in ranges are not permitted. If a gap or an overlap is detected, the GUI will display an error message. The message will advise you to verify the upper and lower limits of your ranges along with the **`Inclusive`** option for those limits. It will also display the range(s) where there is a gap. △

- *Ranges that are predefined:*

  Each range set that is created also has a set of three predefined ranges. These default ranges are High, Low, and Missing:

  - If the value of an objective's indicator or score is higher than the highest value of the ranges that you specified, then the indicator falls in the predefined High range.
  - If the value of the objective's indicator or score is lower than the lowest value of the ranges that you specified, then the indicator falls in the predefined Low range.
  - If the indicator has a missing value, the indicator is assigned to the predefined Missing range. For more information about missing values, see "Calculation Methods" on page 215 and "What is an Outage?" on page 231.

Range Descriptions

Each range description consists of a name, a color, and a label.

- The name identifies the range description. It is used when assigning the color and the label to a range. You can choose the **`Range Description Name`** that you want to use to represent whether the range is Excellent, Good, or Poor.
- The color is the color that will be displayed in the reports. Use the drop-down list to specify the **`Color for text reports`**.

- The label is the label that will be displayed in the reports. If it is not entered, the name of the range description is displayed in the reports. To specify the **`Label in reports`**, you must first check the **`Customize the report label`**.

Both the color and the label can be changed.

Range descriptions are used to provide consistency in reports. For example, if you assign the same range description to all "Excellent" performance ranges, the same colors and labels will be displayed for all excellent performance ranges in the reports. Similarly, all "Good" and "Poor" ranges can each represented by the same colors and labels.

Here is an example of a range set with three user-defined ranges and three predefined ranges.

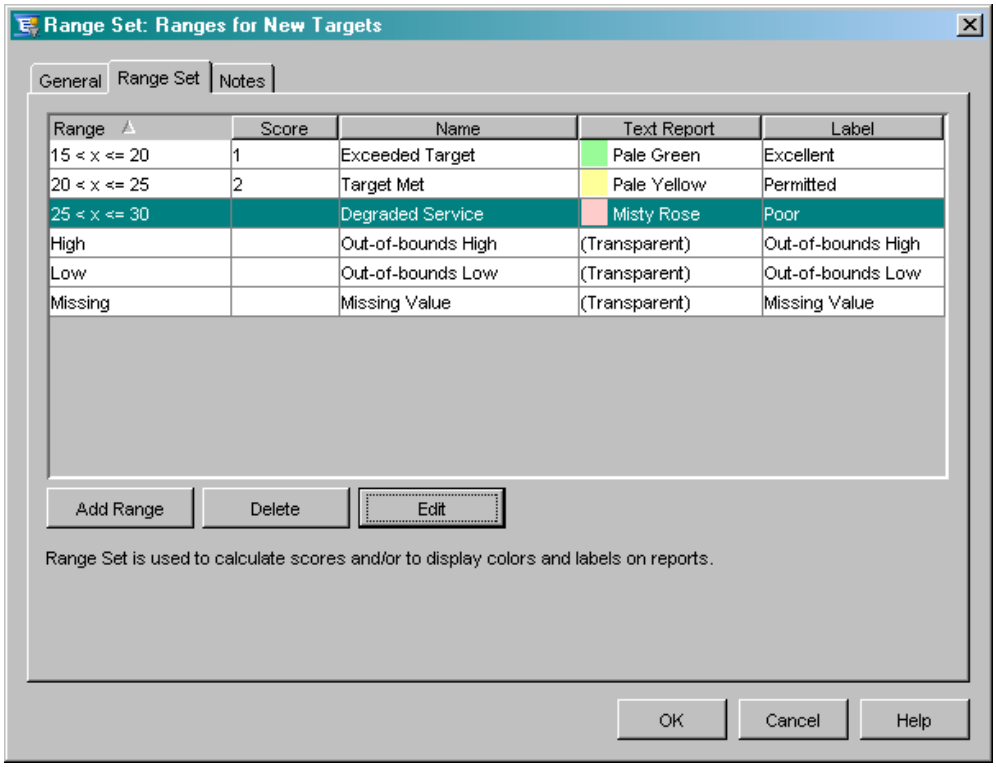

Using a Range Set

A range set is treated as a lookup table. Here are two examples of how the label, color, and score are assigned, based on the objective's indicator:

- Suppose that the indicator is 24. The indicator falls in the 20–25 range, and so the corresponding name is Target Met and the corresponding color on text reports is Pale Yellow. The corresponding score is 2 and the label to be used in reports is Permitted.
- Suppose that the indicator is 27. The indicator falls in the 25–30 range, and so the corresponding name is Degraded Service and the corresponding color for text reports is Misty Rose. The label to be used in reports is Poor.

Customizing Colors and Labels

Typically, you specify the properties of the ranges in a range set from the Catalog workspace. However, you can also customize the color and the label for these ranges for a specific set of reports when you are setting up the specification for those reports. To do so, perform these steps:

1. On the **Range Settings** tab of the Report Specification window, click **Customize range descriptions**.
2. Select the range that you want to update. Click **Edit**. (Alternatively, if you want to revert to the range description that is in the catalog, click **Reset**.)
3. In the Range Description dialog box, enter the label that you want to use for these reports. You can also select a different color to be used for the text reports that will be generated from this report specification.
4. Click **OK**.

Recommendations:

You can create a different range set for each objective or use a range set for more than one objective. In any case, it is a good idea to make all range sets consistent (that is, use the same patterns in every range set). That way, as people move from one contract report to another or drill down from one contract report to another, they can easily compare the information across reports. Here are some recommendations for range set patterns:

- If your range set has two ranges, you might want to define one range for Pass and another range for Fail.
- If your range set has three ranges, you might want to specify range descriptions that use the convention of traffic-light colors (red, yellow, green).
- If your range set has more than three ranges, you might want to specify range descriptions that use a rainbow system (that is, a system in which adjacent ranges have adjacent colors in the spectrum). A rainbow system makes reports easier to interpret visually, even if there is a large number of colors.
- In order to improve the legibility of the numbers on the reports, select muted, pastel colors for your range description colors. Stronger colors will make the numbers on the reports more difficult to read.
- Use the same convention for scores and range description names for each range set.

  Note: For availability and throughput, higher values for indicators are better. But for response time, lower values for indicators are better. It is the *meaning* that should be consistent, not whether the range corresponds to high values or low values. For example, suppose that these ranges are in a range set for availability:

INDICATOR RANGE	RANGE DESCRIPTION	SCORE
0 < X ≤ 80	Missed	1
80 < X ≤ 95	Met	2
95 < X ≤ 100	Exceeded	3

  Then suppose that these ranges are in a range set for response time:

INDICATOR RANGE	RANGE DESCRIPTION	SCORE
150 < X ≤ 200	Missed	1
100 < X ≤ 150	Met	2
50 < X ≤ 100	Exceeded	3

The range with the highest limits is best for availability and the range with the lowest limits is best for response time, but both are consistent in assigning the range description name *Exceeded* to the best range. △

□ If you have decided what the indicator-based targets are, use the indicator's target value as the border between the same two ranges in every range set.

Tip: When you are specifying a range in a range set, use the **`Inclusive`** option to govern whether or not a value is part of the range. In the example that follows, the range $1 < X \leq 40$ has been specified so that the **`Inclusive`** option is checked for the upper limit of the range. That range will therefore *include* the value 40. In that same range, the **`Inclusive`** option is not checked for the lower limit, so that the range will *not include* the value 1.

For example, suppose that a throughput objective's target is that the monthly mean indicator be greater than or equal to 43. You might use a range set with ranges like these:

INDICATOR RANGE	RANGE DESCRIPTION	SCORE
$1 < X \leq 40$	Missed	1
$40 < X < 43$	Almost	2
$43 \leq X \leq 46$	Warning	3
$46 < X \leq 80$	Exceeded	4

Suppose that a response time objective's target is that the monthly mean indicator be less than or equal to 7.075. You might use a range set with ranges like these:

INDICATOR RANGE	RANGE DESCRIPTION	SCORE
$7.15 < X \leq 8$	Missed	1
$7.075 < X \leq 7.15$	Almost	2
$7 < X \leq 7.075$	Warning	3
$6 < X \leq 7$	Exceeded	4

In both cases, the target value is the border between the Almost range and the Warning range. Thus, the Missed and Almost ranges have values that correspond to a target status of missed. And the Warning and Exceeded ranges have values that correspond to a target status of met.

Tip: It is a good idea to include or exclude the limits of the range in the same way that the indicator-based target does. So if the target specification is "indicator $\leq$ Y", then you might want to have the value Y be the high inclusive value of the range. Similarly, the 43–46 range should include 43 because the target does (and probably all of the lower limits should be included in their ranges, to be consistent); the 7–7.075 range should include 7.075 because the target includes 7.075 (and probably all of the upper limits should be included in their ranges, to be consistent).

Note: It can be helpful to have narrow ranges around the target. If behavior has been in the Poor range, then getting to Almost is an encouraging sign that the behavior is almost able to meet the target. If behavior has been in the Good range, then getting to Warning is an indication that the behavior is getting close to missing the target. △

- Consider using the Low and High ranges as traps for outlier values that may be incorrect values.

  For example, instead of having the lowest range go all the way to the lowest possible value, have it go only as low as realistic low values. Thus, the default Low range will include the lowest values (the ones that you think are more likely to be incorrect than actually low).

Composite Objective

The node's composite objective differs from the node's availability, response time, throughput, and custom objectives as follows:

- The node's composite indicator uses only the scores from some or all of its other objectives in order to determine an overall indicator for the node.

  When the data is initially input, the node's composite indicator is calculated from the scores (per five-minute interval) of the objectives that make up the composite objective. The scores for each of the objectives that make up the composite objective are then used to calculate the daily, weekly, monthly, and yearly values of the node's composite indicators.

Note: If you define a composite objective, the contributing objectives must have scores defined in their range sets. △

If you need to report on a node in an overall sense, then use the composite objective.

If you define more than one objective for each node, whether or not you define a composite objective probably depends on what people want to see in the contract reports.

- On nodes that are "low" in the hierarchy (that is, component nodes or nodes that are near them), the people who are responsible for the nodes might be more interested in the specifics than the overall behavior; if so, the composite objectives might not be useful. (However, the people who know little about these nodes might be interested in the overall behavior of these nodes.)
- On nodes that are "high" in the hierarchy (that is, contract nodes or nodes that are near them), there is likely to be more emphasis (by the provider and the customer) on the overall behavior, so composite objectives are likely to be useful.

  A composite objective at the contract node is likely to be very useful.
- If you do not know whether composite objectives might be useful, you can wait to define some or all of them until later. (Each composite objective is independent, so you could start by defining the composite objective for the contract node and then come back later to define composite objectives for other nodes in the contract hierarchy, if they are needed.)

The composite objective works the same way on every node type, so it is not included in the two following topics, which provide information about the differences between component nodes and service, service level agreement, and contract nodes.

Understanding More about Objectives in Component Nodes

The main information that you provide when you define an objective for a component node is as follows:

- what type of objective to define: availability, response time, throughput, or custom objective
- what data source PDB to use for the objective

- what table contains the relevant measure (or variable) for the objective
- what variable in that table contains the relevant measure for the objective
- what timestamp to use if the table you selected is event-based
- for response time, throughput, and custom objectives, what Weight By variable to use if the calculation method is Weighted Mean
- what resources in that table are relevant for the node (the same resources apply to all four objectives)
- what calculation method to use to calculate the objective's indicator from the one value (for each five-minute interval) per resource
- how to use the resulting indicator

Note: The **`Define Data`** option on the component's **`Data`** tab activates a wizard that guides you through the selection of the first seven of these items. △

Data Source PDB

Data source PDBs are created and populated by SAS IT Resource Management. If you are not familiar with your site's data source PDBs, ask your site's SAS IT Resource Management administrator for advice on finding the appropriate data source PDB for each objective.

You specify a data source PDB for each of the four objectives: the availability, response time, throughput, and custom objective. Each objective's indicator gets its input from the data source PDB for that objective. (You specify the input by defining the measure, the table, the data source PDB, and optionally the weight by variable by means of the Data Definition wizard.) For example, the availability indicator gets its input from the measure in the table within the data source PDB for availability, and so on.

Note: There is no requirement that each objective's data source PDB be a different PDB. The same data source PDB can have the data for all four objectives. △

Variable and Table

In the data source PDB, you can choose to use data from either type of table:

- An interval-based table is a table in which each row represents an interval of time. (A timestamp typically represents the beginning of the interval and the duration represents the length of the interval.)
- An event-based table is a table in which each row in the table represents an event. (The timestamp on the data carries the date and time of the event.)

  Note: Event-based tables can be selected for response time, throughput, or custom objectives only. They cannot be used for availability objectives. △

The data from either type of table is then summarized and used to calculate indicators and to determine whether targets have been met or missed.

The value that you choose as the measure must have a data type of Numeric. If you are not familiar with the tables and measures in your site's data source PDBs, ask your site's SAS IT Resource Management administrator for advice on finding the appropriate table and measure for the objective.

Tip: If a table has the data that you need but the data is spread across several columns in the same table, then ask the SAS IT Resource Management administrator to create a formula variable that has the information that you need. You then need to specify this formula variable as the measure.

Resources

A resource is an item of hardware or software or any other entity for which data is collected. A resource, whether by itself or combined with other resources, makes up a component.

A data source PDB table might have information about only the resources in this component or might have information about additional resources in other components.

For example, suppose the data source PDB table has information about servers. It might have information about e-mail servers, print servers, Web servers, and so on. If the data source PDB table has information about all the site's servers and this component is for a particular Web-server pool, you need to identify the servers that are included in the pool.

- You might need the value of only one column to identify the resources. For example, for a hardware resource you might use the column MACHINE; you then can select the values that name the machines in the Web-server pool.
- You might need to use two or more columns to identify the resources. For example, for the same hardware resource you might use the columns DOMAIN and MACHINE on Windows; you then can specify for each machine both its domain and name (in case duplicate machine names are allowed across domains).

Note: The data source PDB table's key variables are good candidates for columns to identify resources. (In terms of SAS IT Resource Management, these key variables are known as BY columns and CLASS columns.) △

One Value per Resource per Objective

For a given five-minute interval, the ETL job automatically goes through several steps to obtain the "one value for each resource" from the detail-level table in the objective's data source PDB. The result is written to the node's underlying table for that objective: the availability table for the availability objective, and so on. (For more information about tables in the contract database, see "The Contract Database" on page 237.)

Here are the steps, with an example:

Note: The example does not show the datetime stamps, but each row has a date and time. △

1 The ETL job filters the resources to exclude those rows that are not necessary for the processing. It then sorts the remaining rows by resource and datetime.

For example, suppose a table has data from five servers (US Operations—HAL6000, US Operations—Server/01, US Operations—ALLSERV, EU Operations—HAL6000, and SA Operations—HAL6000). Suppose also that the table has data at one-minute intervals. (SAS IT Service Level Management uses a five-minute interval, but the interval length in the data source PDB's table can be different.) In the table, the rows that correspond to a given five-minute interval in the contract might look like this when sorted by resource:

RESOURCE DOMAIN	RESOURCE MACHINE	DURATION (IN SEC.)	MEASURE VARIABLE	WEIGHT VARIABLE
US Operations	HAL6000	60	25	1
US Operations	HAL6000	60	23	2
US Operations	HAL6000	60	24	1
US Operations	HAL6000	60	26	2
US Operations	HAL6000	59	27	100

US Operations	Server/01	60	21	1
US Operations	Server/01	60	22	1
US Operations	Server/01	58	23	2
US Operations	Server/01	60	21	2
US Operations	Server/01	60	23	100
US Operations	ALLSERV	60	33	2
US Operations	ALLSERV	60	36	2
US Operations	ALLSERV	50	33	2
US Operations	ALLSERV	60	36	2
US Operations	ALLSERV	60	31	100
US Operations	WebServ	60	23	1
US Operations	WebServ	60	23	2
US Operations	WebServ	49	23	2
US Operations	WebServ	60	23	2
US Operations	WebServ	60	23	2
EU Operations	HAL6000	60	12	7
EU Operations	HAL6000	60	10	9
EU Operations	HAL6000	60	12	14
EU Operations	HAL6000	60	12	3
EU Operations	HAL6000	50	12	53
SA Operations	HAL6000	60	53	5
SA Operations	HAL6000	60	51	9
SA Operations	HAL6000	60	53	27
SA Operations	HAL6000	60	53	2
SA Operations	HAL6000	60	53	1

Note: In the preceding table, the MEASURE VARIABLE is the variable that is used as the measure for the response time, throughput, or custom objective and the WEIGHT VARIABLE is the variable that can be specified as the Weight By variable on the **Data** tab of the response time, throughput, or custom objective. △

2 The ETL job subsets the rows by resource, so that only rows from the selected resources are used.

Suppose that only three of the servers are in the component's server pool (US Operations — HAL6000, US Operations — Server/01, and US Operations — ALLSERV). Then, only those servers' rows are used in the remainder of the calculation. So at this point, the rows look like this:

RESOURCE DOMAIN	MACHINE	DURATION (IN SEC.)	MEASURE VARIABLE	WEIGHT VARIABLE
US Operations	HAL6000	60	25	1
US Operations	HAL6000	60	23	2
US Operations	HAL6000	60	24	1
US Operations	HAL6000	60	26	2
US Operations	HAL6000	59	27	100
US Operations	Server/01	60	21	1
US Operations	Server/01	60	22	1
US Operations	Server/01	58	23	2
US Operations	Server/01	60	21	2
US Operations	Server/01	60	23	100
US Operations	ALLSERV	60	33	2
US Operations	ALLSERV	60	36	2
US Operations	ALLSERV	50	33	2
US Operations	ALLSERV	60	36	2
US Operations	ALLSERV	60	31	100

3 *For response time, throughput, or the custom objective*, the ETL job writes one row for each resource and the row contains the maximum, minimum, sum, mean, and weighted mean of the above rows for the five-minute interval. Additionally, WEIGHT is summed. So the rows in the underlying table for the response time, throughput, or custom objective look like this:

		VALUES AVAILABLE FOR USE AS MEASURE FOR THE RESPONSE TIME, THROUGHPUT, OR CUSTOM OBJECTIVE					
RESOURCE DOMAIN	MACHINE	MAXIMUM	MINIMUM	SUM	UNWEIGHTED MEAN	WEIGHTED MEAN	WEIGHT
US Operations	HAL6000	27	23	125	25.0	26.9	106
US Operations	Server/01	23	21	110	22.0	22.9	106
US Operations	ALLSERV	36	31	169	33.8	31.3	108

The table contains columns that are used to accumulate the measures for each resource: maximum, minimum, sum, unweighted mean, and weighted mean. The calculation method that you specify next obtains its input from only one of these columns. So that column provides the one value per resource. (You can specify the mean or the weighted mean calculation method, but not both. They are mutually exclusive.)

Note: The underlying tables at the resource level store both the weighted mean and the unweighted mean for the response time, throughput, and custom objectives. (At the higher nodes, the only mean that is kept is the one that you specified as the calculation method.) △

For availability, the ETL job writes one row for each resource, and the row contains the uptime for the resource (in seconds) divided by the expected duration, expressed as a percent. So the rows in the availability table look like this:

RESOURCE DOMAIN	MACHINE		UPTIME
US Operations	HAL6000	(299/300*100=)	99.67
US Operations	Server/01	(298/300*100=)	99.33
US Operations	ALLSERV	(290/300*100=)	96.67

The calculation method (that you specify next) obtains its input from the UPTIME column, and that column provides the one value per resource.

Note: The example shows how original intervals shorter than five minutes are consolidated into five-minute intervals. Similarly, original intervals longer than five minutes are prorated into five-minute intervals.

Also, regardless of the interval length and even if the original intervals are five minutes), the ETL job redistributes the data so that the new intervals are on the following five-minute boundaries: :00, :05, :10, and so on. △

Calculation Method

For an objective, the calculation method obtains the appropriate values from the objective's underlying tables (Availability, Response Time, Throughput, and Custom Objective), performs the appropriate calculation, and assigns the result to the objective's indicator in the five-minute row in the component's main table in the contract database. For more information about the tables, see "The Contract Database" on page 237.

You can choose one of these calculation methods: maximum, minimum, mean, weighted mean, sum, percentile, redundant, or serial.

□ *Maximum or Minimum:*

For the availability objective, the maximum or minimum calculation method obtains the values in the UPTIME column (99.67, 99.33, and 96.67), takes the maximum (99.67) or minimum (96.67) of them, and assigns that value to the availability indicator.

For the response time, throughput, or custom objective, the maximum or minimum calculation method obtains the values in the MAX column (27, 23, and 36) or MIN column (23, 21, and 31), takes the maximum (36) or minimum (21) of them, and assigns that value to the response time indicator, throughput indicator, or custom objective indicator.

□ *Mean*:

For the availability objective, the mean calculation method obtains the values in the UPTIME column (99.67, 99.33, and 96.67), takes the mean of them (98.56), and assigns that value to the availability indicator (one value for each five-minute interval).

The response time, throughput, or custom objective calculates the mean of the child node or resource indicators. It then assigns the resulting value to the response time indicator, throughput indicator, or custom objective indicator (one value for each five-minute interval).

□ *Weighted Mean*:

The response time, throughput, or custom objective calculates the weighted mean of whatever the child node or resource indicators are, using the weights for the respective child node or resources, according to the following formula:

$$\overline{x} = \frac{\sum_{i=1}^{n} w_i \cdot x_i}{\sum_{i=1}^{n} w_i}$$

where $\overline{x}$ is the weighted mean, the numerator is the weighted sum, and the denominator is the sum of weights.

For example, the calculation of the weighted mean (TI) for the throughput objective's indicators is as follows:

```
TI = [(26.9*106) + (22.9*106) + (31.3*108)] / 106 + 106+ 108
TI = 2851.4 + 2427.4 + 3380.4 / 320
TI = 8659.2 / 320
TI = 27.06
```

It then assigns the resulting value to the response time indicator, throughput indicator, or custom objective indicator (one value for each five-minute interval).

Note: When an objective requests the calculation of the weighted mean, the child nodes are not required to calculate their indicators as weighted means. The parent node calculates both the weighted sum and the sum of the weights in order to determine the weighted mean for the objective. △

□ *Sum:*

For the availability objective, the sum calculation method obtains the values in the UPTIME column (99.67, 99.33, and 96.67), takes the sum of them (295.67), and assigns that value to the availability indicator (one value for each five-minute interval).

For the response time, throughput, or custom objective, the sum calculation method obtains the values in the SUM column (125, 110, 169), takes the sum of them (404), and assigns that value to the response time indicator, throughput indicator, or custom objective indicator (one value for each five-minute interval).

- *Percentile:*

  For the availability objective, the percentile calculation method obtains the values in the UPTIME column (99.67, 99.33, and 96.67), sorts them in ascending order (96.67, 99.33, 99.67), identifies or calculates a value that corresponds to the percentile that you specified (in the GUI), and assigns that value (99.33) to the availability indicator.

  For the response time, throughput, or custom objective, the percentile calculation method obtains the values in the MEAN column (25.0, 22.0, 33.8), sorts them in ascending order (22.0, 25.0, 33.8), identifies or calculates a value that corresponds to the percentile that you specified in the GUI (say, 50th), and assigns that value (25.0) to the response time indicator, throughput indicator, or custom objective indicator (one value for each five-minute interval.

  This method is most useful if you have a large number of inputs, ideally more than 100. In this example, because 3 is such a small number of input values, you would probably not select the percentile calculation method.

  Note: The percentile calculation uses the same default algorithm that is used by PROC UNIVARIATE (specifying option PCTLDEF5). △

  For more information about the percentile calculation method, see "Percentile" on page 217.

- *Redundant*:

  If the objective type is availability and if the relation of the component's resources is redundant, it is a good idea to select the redundant calculation method. (A redundant relation is one in which all of the resources are in a pool and serve as a backup to one another; thus, the component is available even if only one of its resources is "up.")

  For the availability objective, the redundant calculation method obtains the value in the UPTIME column (99.67, 99.33, and 96.67), applies the redundant "formula," and assigns the resulting value (100.00) to the availability indicator.

  For the response time objective, throughput objective, or custom objective, the redundant calculation method is not offered as a choice.

- *Serial*:

  If the objective type is availability and if the relation of the component's resources is serial, it is a good idea to select the serial calculation method. (A serial relation is one in which each of the resources is both unique and required; thus, the component is available only if every one of its resources is "up.")

  For the availability objective, the serial calculation method obtains the value in the UPTIME column (99.67, 99.33, and 96.67), applies the serial "formula," and assigns the resulting value (95.71) to the availability indicator.

  For the response time objective, throughput objective, and the custom objective, the serial calculation method is not offered as a choice.

- *Custom Code*:

  You define custom code that allows user-defined SAS code to be written in order to assign a value to the objective's indicator.

Note: For more information about the Maximum, Mean, Minimum, Sum, and Weighted Mean methods, and the formulas used to calculate them, see "Maximum, Mean, Minimum, and Sum" on page 217.

For more information about the redundant calculation method, see "Redundant" on page 217.

For more information about the serial calculation method, see "Serial" on page 218.

For more information about the Custom Code method, see "Custom Code for the Calculation Method" on page 221. △

Note: If the five–minute interval is in an outage period, then the indicator's value is not calculated. Instead it is assigned, based on the type of the outage (either scheduled, unscheduled, or excused). For more information about outages, see "What is an Outage?" on page 231. △

For more information about tables in the contract database, see "The Contract Database" on page 237.

How to Use the Resulting Indicator

For all nodes, after the indicator has been determined by the calculation method that you selected, you can specify that the indicator be compared to the range set in order to assign a score, a color, or a label that is displayed on the reports.

The indicator or the score can then be compared to the target value in order to determine whether it met or missed its target.

Target Values

In order to obtain realistic target values for your objectives, run the software in a test environment for a period of time so that reasonable benchmarks are established. Baseline reports are also available from the Contract window. You might also inquire whether there are industry-standard values for the component.

Understanding More about Objectives in Higher-Level Nodes

The main information that you provide when you define an objective for a higher-level node such as a service node, service level agreement node, or contract node is as follows:

- whether to calculate the objective's indicator from the indicator or score of the child node(s)
- what calculation method to use to calculate the objective's indicator from the one value (per five-minute interval) per child node
- how to use the resulting indicator or score

The following three figures show the process of converting indicators into scores.

Figure 3.1 Using the Indicators of Child Nodes to Calculate an Indicator for the Parent Node

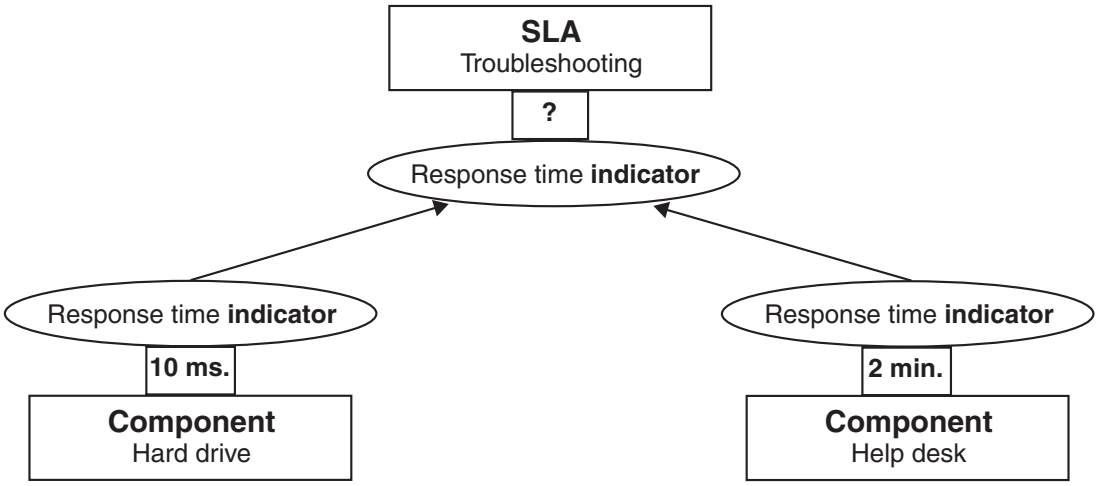

Figure 3.2 Using the Range Sets to Assign Scores Based on the Indicators

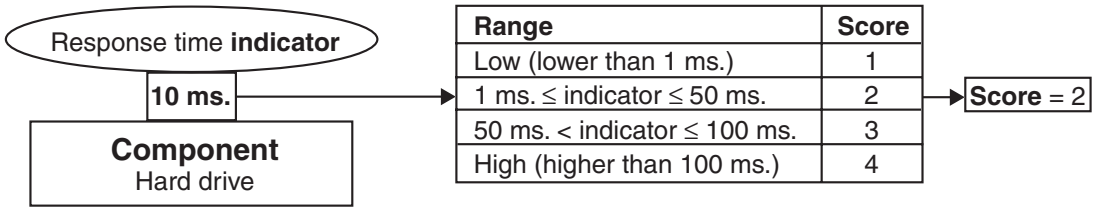

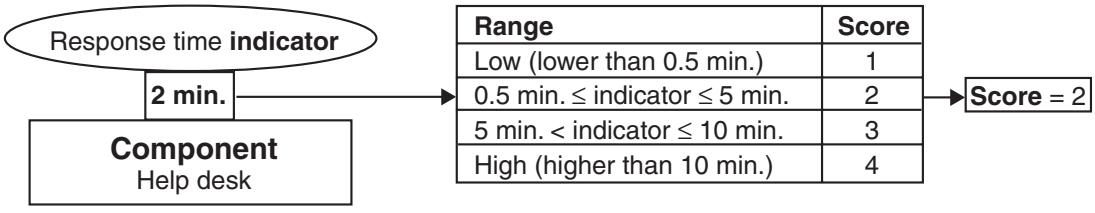

Figure 3.3 Using the Scores from the Child Nodes to Produce a Meaningful Result

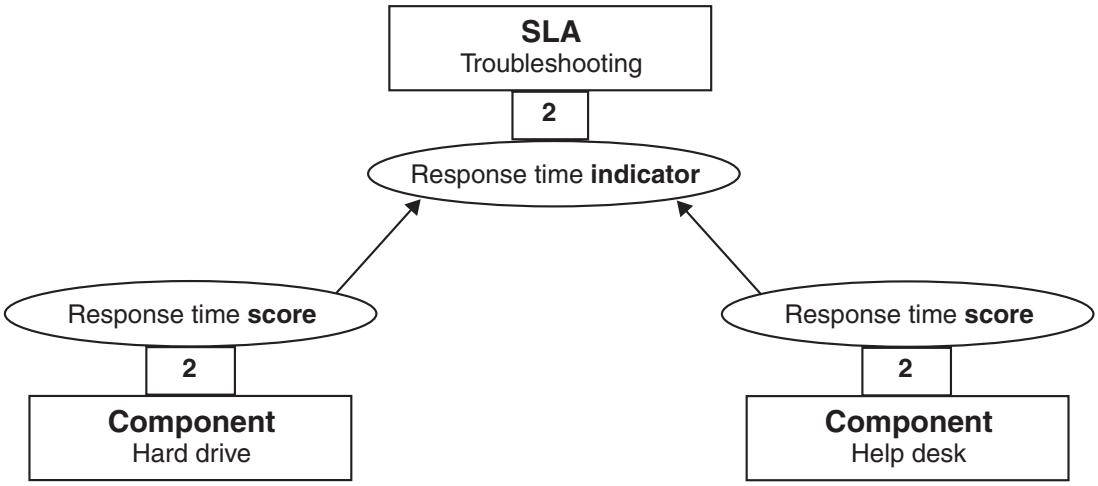

The resulting score will compared to the scores in the range set in order to determine the appropriate range description. The range description will provide the color and the label to be used in the reports for this score.

Note: The reports for objectives at the higher-level nodes will display colors and labels based on the target statistic of the indicator. If the target statistic is *sum*, then the transparent color will be used in the reports for all levels except the compliance target period. The color that is assigned for the target compliance period will be determined by the value of the sum statistic. △

Note: A contract node and its objectives are defined in the Contract workspace, not in the Catalog workspace, but the objectives work in the same way on a contract node as on a service node and service level agreement node. △

One Value per Child Node per Objective

We recommend that you make sure that the measures of the child nodes' objectives use the same scale. As shown in Figure 3.1 on page 32, combining indicators for child nodes that use dissimilar response time measures produces a meaningless indicator for the parent node. You can base the indicator on its child node indicators *or* on its child node scores, for each of the availability, response time, throughput, and custom objectives.

- If the child nodes are not similar or if the measures of your child nodes' objectives are not of a comparable scale, assign scores (by means of the range sets) to the unlike indicators in the child nodes, as in Figure 3.2 on page 32. (Even if the indicators are in different units and on different scales, their scores can be comparable if you use a consistent pattern in your range sets.)
- If the child nodes are similar, you might tend to choose:
  - the child node indicators, for the nodes that are "low" in the hierarchy
  - the child node scores, for the nodes that are "high" in the hierarchy

  As mentioned earlier, the people interested in the "low" nodes might be more interested in precise values, and the people interested in the "high" nodes might be more interested in the categories in which the values fall.

Note: Statistics are calculated across all five-minute intervals. The results are summarized to the daily tables. The same results are also summarized to the weekly

tables, to the monthly tables, and to the yearly tables. (In other words, the statistics are not "rolled up" from the daily table into the weekly table, and from the weekly table into the monthly table, and so on. Instead, the statistics that are calculated across all five-minute intervals are summarized to each of the other tables separately.)

All five-minute intervals in a day are represented by statistics in the daily table. △

Calculation Method

The calculation methods work the same way for one value per child per objective node as they do for one value per resource per objective. For example, if a node has two child nodes, the ETL job might obtain five-minute rows (one row from each child node table) like these:

CHILD NODE	AVAILABILITY INDICATOR	AVAILABILITY SCORE	RESPONSE TIME INDICATOR	RESPONSE TIME SCORE	THROUGHPUT INDICATOR	THROUGHPUT SCORE	CUSTOM OBJECTIVE INDICATOR	CUSTOM OBJECTIVE SCORE
node#1	100	4	2.7	3	357	4	152	4
node#2	95	2	2.2	4	397	4	230	6

The ETL job might use the values in the calculation method to assign values in the five-minute row for the node's own table. Here is an example:

- If you use the availability indicators (100, 95) as input to the availability objective, and if you use the mean calculation method for the availability objective, then the availability indicator is 97.5.
- If you use the response time scores (3, 4) as input to the response time objective, and if you use the mean calculation method for the response time objective, then the response time indicator is 3.5.

Note: Do not use serial or redundant calculation methods for an objective whose indicator is based on the scores of the child nodes, or where the indicators in any other node "below" the one in the hierarchy calculated the indicator based on a score. (This is because the formulas that are used to calculate the serial and redundant indicators expect to work with availability values in the range 0–100, that is, values that represent the percentage of uptime.) △

For more information about calculation methods, see "Calculation Method" in "Understanding More about Objectives in Component Nodes" on page 24.

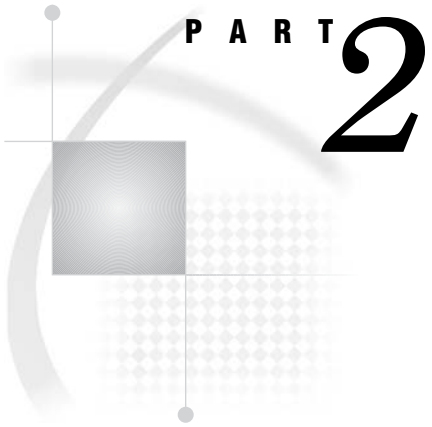

PART 2

Administration

CHAPTER

4

Logging On to the Software

Introduction to Logging On to the Software

You can work with the SAS IT Service Level Management graphical user interface (GUI) in any of the following ways:

- You can invoke SAS IT Service Level Management from a SAS IT Service Level Management client.
- You can log on from the client to a SAS IT Service Level Management server and access an environment database on that server.
- You can log off from that SAS IT Service Level Management server.

Note: If you are connecting to a SAS IT Service Level Management server via a spawner, the spawner has to be running on the server. If your server is on Windows, you must connect via a spawner. If your server is on other platforms, you can connect via a spawner or via TELNET access, but spawners are highly recommended.

At most sites the systems administrator is responsible for ensuring that the spawner is automatically started. If you are the systems administrator for the server, see the installation instructions for further details. △

Invoking SAS IT Service Level Management

To invoke SAS IT Service Level Management, select the following items from the Windows task bar: **Start ► Programs ► SAS ► SAS IT Service Level Management 2.1**.

The Logon dialog box opens.

Logging On to a SAS IT Service Level Management Server

The Logon dialog box has three items: **Environment**, **User Name**, and **Password**.

The environment connection that you select when you log on is called the *active environment*. Similarly, the SAS IT Service Level Management server that the active environment points to is called the *active SAS IT Service Level Management server* (or *active server*), and the environment database that the active environment points to is called the *active environment database*.

Environments

An environment is located on the server and represents the combination of the environment database and the SAS IT Service Level Management software on that server.

Note: The term *server* used below always refers to the server where SAS IT Service Level Management server software has been installed. △

An environment is accessed by means of an **Environment Connection**.

An environment connection is stored on the client where it is defined. A client must be on a host that is running Windows. A server can be on a host that is running Windows, Linux, UNIX, or z/OS. A client and server can be (but are not required to be) on the same host, if the host is running Windows. Otherwise, a client and server are on different hosts.

Note: If your server is on z/OS, the steps to create and initialize the environment and contract databases on the server must take place ahead of time. For information about this, see "Working with Servers on z/OS" on page 40. △

The environment connection has the following features:

- properties that enable the client to create a connection to a given server
- a property that identifies the path to the environment database that you want to use on that server.

  An environment database stores almost all the objects, such as components, range sets, range descriptions, services, and service level objectives, that are managed by the GUI, with two exceptions:

  1. The input data for contracts is separate, in a maximum of four data source PDBs per component. (The data source PDB is also known as the SAS IT Resource Management Performance Database, or PDB.)
  2. The contract database, which contains the output data for contracts, is separate. There is a separate contract database for each contract.

  However, information *about* access to and locations of those data source PDBs and contract databases is contained in the environment database, and the GUI enables you to *view* the data in these data source PDBs and contract databases.

Your client machine can have one or more environment connections, however, only one can be active at any given time. A SAS IT Service Level Management site has one or more SAS IT Service Level Management clients and at least one SAS IT Service Level Management server.

Concurrent Access with SAS/SHARE

SAS/SHARE enables multiple tasks to have concurrent access to an environment database. Concurrent access supports any combination of the following tasks:

- accessing the same environment database from multiple client GUIs
- running ETL tasks on multiple contracts that are located in the same environment database
- executing multiple reporting tasks that access contracts from the same environment database

For information about how to specify the SAS/SHARE parameters necessary for the concurrent use of the environment database, see the topic "About the Environment Properties Dialog Box" in the online help.

Information about how to start a SAS/SHARE server is available in the folder named SASMISC, which is located where SAS IT Service Level Management is installed.

Supplied Environment Connection

By default, you have one environment connection, named localhost.

- If there is an active server on the same host as your client, you can use the localhost environment connection to connect to that server and access an environment database on that server.

  If there is not an existing environment database at the specified location, one will be created at that location after you connect to that server.
- If there is not a server on the same host as your client, you can ignore the localhost environment.
- If there is a server on the same host but you want to connect to it in a different way, you can edit the localhost environment connection.

  If there is a server on the same host but you want to use a different environment database location, you can add another environment connection.

Note: You cannot share data between environment connections. Also, you cannot copy environment connections. △

Other Environments Connections

If the server that you want to log on to is not on the same host as your client, then you must add an environment for that server. On Windows and UNIX platforms, if the specified environment database does not already exist, it will be created and initialized at the specified location on the specified server when you log on. When the specified environment database is on z/OS, you must allocate space for an environment database directly on your z/OS platform. (For information about this, see "Working with Servers on z/OS" on page 40.)

You might also want to add environment connections that are variations of other environment connections. The following describes an example scenario in which you might have multiple environment connections that connect to the same server and use different environment databases:

- One environment connection might point to an environment database and to contracts that are used for production work.
- Another might point to the environment database and to contracts that are backup copies of the environment database and the contracts for production work.
- A third environment connection might point to an environment database and contracts that are used for testing.

- A fourth might point to the sample environment database (with sample contracts) that might have been copied to the server at installation time.

Although you can give arbitrary names to environment connections, it is a good idea to give each environment connection a name that is unique and meaningful. For example, if two environment connections contain properties for connection to the server named wolf, and one environment connection points to the environment database for production and the other points to its backup, you might name the environment connections wolf-production and wolf-backup.

Working with Servers on z/OS

If your server is on z/OS, the steps to create and initialize the environment and contract databases on the server must take before you start using the software. On the server, the SAS IT Service Level Management server software has a partitioned data set (PDS) named CPMISC. The default location for this data set is:
***where_you_installed_sas_at_your_site*.itrm.cpmisc**
The PDS has a member named CMSLMALC that contains a sample batch job that you can customize and submit to allocate, create, and initialize an environment database. It also contains a sample batch job that you can customize and submit to allocate, create, and initialize a contract database. For more information, see the instructions in CMSLMALC.

User Names and Passwords

For information about whether a user name and password are required and, if so, how to arrange for them, see the SAS IT Service Level Management administrator at your site.

Note: If you connect to the SLM server and if you have specified the use of SAS/SHARE, then the same authentication that you use to connect to the SLM server is used to connect to the SAS/SHARE server. If you need to specify the user ID and password to enable SAS/SHARE, then you can do so in two ways:

- by using the Environment Properties dialog box, if you are using the GUI
- by coding the parameters of the %CPSLMETL, %CPSLMENV, and %CPSLMRPT macros, if you are in batch mode

If a user name and password are not required by the SLM server, you must specify that the environment connection does not require authentication. This will disable the user name and password fields. △

Logging On

To log on to a SAS IT Service Level Management server from your client, in the Logon dialog box, follow these steps:

1. Select an environment from the list of existing environments, or add a new environment and select the new environment.
2. Provide a user name and password (if required).
3. Click **OK**.

For more information about the Logon dialog box, click **Help**.
For more information about viewing the fields in the Environment Properties dialog box, click **Edit** and then **Help**.

For more information about adding, modifying, and removing environment connections, click **Help** and then select the tasks that are listed in the **Environment** section of the table of contents.

After the logon is complete, the main window opens and its title includes the name of the environment that you selected.

Working in the Main Window

The main window has five specialized workspaces: Contract, Catalog, Participants, Data, and Reports. Each of these workspaces displays a different part of the active environment database. When you log on to SAS IT Service Level Management, the **Contract** workspace opens. The left panel displays a list of the workspaces that are available: Contract, Catalog, Participants, Data, and Reports.

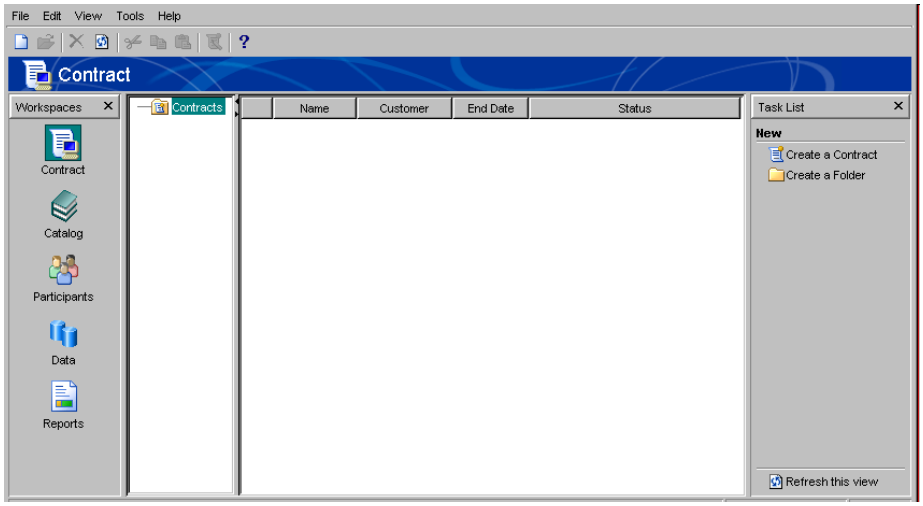

Note: All the folders and subfolders in any of the workspaces are displayed in ascending alphanumeric order. △

If you are new to SAS IT Service Level Management, you will probably need to use the specialized workspaces in this order: Data, Participants, Catalog, Contract, and Reports.

- If you select **Data** from the **Workspaces** list, the Data workspace opens and you see the information about server connections, data source PDBs, and contract databases. You use this workspace to view, add, change, and delete information that is related to data access.

  For more information about this workspace, see "Introduction to Accessing Data" on page 43.

- If you select **Participants** from the **Workspaces** list, the Participants workspace opens and you see the information about customer and provider organizations as well as the contact people at those organizations and elsewhere. You use this workspace to view, add, change, and delete information about participants.

  For more information about this workspace, see "Introduction to Recording Participants" on page 53.

- If you select **Catalog** from the **Workspaces** list, the Catalog workspace opens and you see the information about components, services, and service level agreements that the environment database has in its catalog. The workspace also has

information about range sets and range descriptions that the components, services, service level agreements, and contracts can use. You use this workspace to view, add, change, and delete information that is in the catalog.

For more information about this workspace, see "Introduction to Building the Catalog" on page 61.

- If you select **`Contract`** from the **`Workspaces`** list, the Contract workspace opens and you see the information about contracts. You use this workspace to view, add, change, and delete information about contracts. In each contract, you also create the batch ETL job that processes the contract's data and objectives, and you enter any outage information that affects the contract.

  For more information about this workspace, see "Introduction to Defining a Contract" on page 81.

- If you select **`Reports`** from the **`Workspaces`** list, the Reports workspace opens and you see the information about report specifications and the global settings that you can use to create the reports. You use this workspace to view, add, change, and delete information about the report specifications that govern what reports will be generated. (Reports are specified for two types of viewers: contract providers and customers. The report specifications for each type of viewer are placed in fixed folders in this workspace that are labelled Contract Provider and Customer.) You also use this workspace to view, add, change, and delete information about the report global settings. These global settings specify the default locations of the customer reports and the contract provider reports and the style sheet that governs the appearance of the reports.

If you are familiar with SAS IT Service Level Management, you can use the specialized workspaces in any order, depending on what you want to view, add to, change, or delete.

Logging Off

To log off the active server, from any workspace you select **File ▶ Exit**.

If your client/server connection fails, you can reconnect from any workspace by selecting **Tools ▶ Reconnect**.

Note: If the active server is running z/OS, the connection might fail if you have not used the connection for a period of time that exceeds the specified timeout period at your site. △

If you finish work with one environment and want to work with another, you can directly log off the first environment and log on to the other. Alternatively, from any workspace, you can switch environments by selecting **Tools ▶ Switch Environments**.

Note: This procedure logs you off the current environment and then logs you on to the selected environment. △

CHAPTER

5

Accessing Data

Introduction to Accessing Data

The Data workspace contains information that enables the SAS IT Service Level Management GUI, the batch jobs, and you to access input and output data for the contracts in the Contract workspace.

The Data workspace contains three supplied folders: the **`Contract Databases`** folder, the **`Data Source PDBs`** folder, and the **`Server Connections`** folder.

- In the **`Contract Databases`** folder, you can view the current list of contract databases, view their properties, and view the data in their contract databases.

  Note: In the Contract workspace, the contract database is automatically defined when a contract is activated from the Contract workspace. You do *not* create contract databases in this folder. △

- In the **`Data Source PDBs`** folder, you can view the current list of data source PDBs, view their properties, and view the data in their data source PDBs.

  If you need to access other data source PDBs, you can define other data source PDBs in this folder. You do not create data source PDBs from this workspace; instead you define the pointers to existing data source PDBs.

- In the **Server Connections** folder, you can view the current list of server connections (both supplied and user-defined) and view the properties of the user-defined server connections.

  Note: The server connection sets up the connection between the active server and the SAS IT Resource Management (ITRM) server. The SAS IT Service Level Management server is the machine where the product software is installed along with the active environment database. The ITRM server is the machine where the data source PDBs are located.
  This connection retrieves data from the ITRM server for use by SAS IT Service Level Management. △

  If you need other server connections, you can define the properties in this folder that are used to establish a connection. You can also use this folder to connect to and disconnect from servers.

Note: If you want to use a data source PDB that will be accessed by using a certain server connection, you must create that server connection *before* you access the data source PDB.
If you run ETL (or Synchronize) interactively and you are not connected to the data source PDB(s) that you need, you will be prompted for the user ID and password for each of the necessary connections. △

Contract Databases

The **Contract Databases** folder contains information about contract databases. To see the list of contract databases, double-click that folder.
The list contains the names of the contract databases that are associated with the contracts in the Contract workspace. There is one contract database for each active contract (and for each inactive contract that was formerly active). The name of each contract database is the same as the name of its contract.

Creating a Contract Database

On Windows and UNIX, a contract database is created automatically when you run ETL (or Synchronize) on a contract that is marked *Active*. For z/OS, you must allocate the space for the contract database before you run ETL (or Synchronize) on a contract for the first time. (For more information and a sample job step to allocate the space, see “Working with Servers on z/OS” on page 40.

Opening a Contract Database

To view the information about a contract database and/or to view the data in its contract database, perform these steps:

1. From the list of contract databases, select the one that you want to use.
2. From the **Task** list, click **Open**. The Contract Database dialog box opens and displays the properties of the contract database. You cannot edit the information here. If you want to edit the information, use the Contract workspace.
3. To view the data in the contract database, on the **General** tab, click **View Data**.

   Note: You do not need to open a server connection, because the contract databases are on the same host as the active server. △

For more information about the Contract Database dialog box, click **Help**.

Contract databases are created and populated by SAS IT Service Level Management. For more information about contract databases and about the Contract workspace, see "Introduction to Defining a Contract" on page 81, and "The Contract Database" on page 237.

Data Source PDBs

The **Data Source PDBs** folder contains information about data source PDBs. To see the current list of data source PDBs, double-click that folder.

The list contains the names of the data source PDBs that are used (or will be used) by the components that are nodes in a contract.

Data source PDBs are not supplied in the list because their definition, creation, naming, and location is controlled by the SAS IT Resource Management software at your site. Contact the SAS IT Resource Management administrator at your site to find out what data source PDBs are available for use with SAS IT Service Level Management.

Defining a Data Source PDB

Before you define a data source PDB, you need to define a server connection to the host where the data source PDB is located. You do not need to define a server connection if you can access the data source PDB from the active server or if you can access the data source PDB by using a server connection that is already defined. However, even though you do not need to be connected to that server, the server connection does need to be defined. (For more information about defining a server connection, see "Server Connections" on page 46.)

To define a data source PDB, perform these steps:

1. From the **Task** list, select **Define a Data Source PDB**. The Data Source PDB dialog box opens.
2. On the **General** tab, provide the following properties:
   - **Data Source PDB Name**

     Note: Although you can give arbitrary names to data source PDBs, it is a good idea to give each data source PDB a name that is unique and meaningful. △
   - **Path to Data Source PDB**
   - **Server Connection Name**

     [SLM Server] is the default.
3. Click **OK** to save the properties of this data source PDB.
4. For more information about this dialog box, click **Help**.

Opening a Data Source PDB

To view or edit the properties of a data source PDB and/or to view the data in a data source PDB, perform these steps:

1. From the list of data source PDBs, select the one that you want to use.
2. From the **Task** list, select **Open**. The Data Source PDB dialog box opens and displays the properties of the data source PDB. You can view and/or edit the properties of the data source PDB.

3 To view the data in the data source PDB, select the **General** tab and look at the **Connection Status**.

   - If the **Connection Status** is *Not Connected* [], click **Connect to Server**, enter the requested logon information for that server, and then click **View Data**.
   - If the **Connection Status** is *Connected* [], click **View Data**.

Renaming a Data Source PDB

To change only the name of the data source PDB, perform these steps:

1 From the list of data source PDBs, select the one whose name you want to change.

2 From the menu bar, select **File ▶ Rename**.

3 Enter the new name and then click **OK**.

Note: If you rename a data source PDB, all references it (both in the catalog and in the contracts) will be automatically renamed. △

Removing a Data Source PDB

To remove a data source PDB from the list of data source PDBs that are available to SAS IT Service Level Management, perform these steps:

1 From the list of data source PDBs, select the one that you want to remove.

2 From the **Task** list, select **Remove ▶ Yes**.

CAUTION:
Removing this data source PDB from the list of data source PDBs while it has connections to a contract will invalidate the contract; the data will not be collected. If this data source PDB is currently assigned to components within a contract, make sure that you update the contract to access a new data source PDB prior to deleting this one. △

Removing a data source PDB does not delete it. Instead, it removes that data source PDB from the list of data source PDBs that are available to SAS IT Service Level Management.

Data source PDBs are created and populated by SAS IT Resource Management. For more information about data source PDBs, see "Part 3: Administration" in *SAS IT Resource Management: User's Guide*. You can view (and/or print) the guide from the Web page at http://support.sas.com/documentation/onlinedoc/itsv.

Server Connections

The **Server Connections** folder contains information about one or more server connections. To see the current list of server connections, double-click that folder.

Note: The server connection is used by the ETL job to access data source PDBs. △

The Supplied Server Connection

The supplied server connection is named **[SLM Server]**. (The name is in square brackets to distinguish the supplied connection from the user-defined server connections.)

The active environment provides the properties of the supplied server connection. The supplied server connection is *from* the active SAS IT Service Level Management client host and *to* the active SAS IT Service Level Management server.

For each contract for which ETL (or Synchronize) has been run, there is one contract database on this server. The supplied server connection enables the GUI (and you) to access these contract databases.

If any data source PDBs are on the active server (or are accessible from the active server), the supplied server connection also enables the GUI (and you) to access these data source PDBs.

Note: During the current SAS IT Service Level Management session, the client is always connected to the active server until you log off or switch to a different environment or until something, such as the server rebooting, causes the client to be disconnected. If this occurs, you can reconnect to your environment by selecting **Tools ► Reconnect**. △

User-Defined Server Connections

If you want to access data source PDBs that are on hosts other than the active server, then you might need to define a server connection for each of those hosts, depending on these conditions:

- If the data source PDB was created on an operating system that is different from the active server, then you must define a server connection.
- If the data source PDB is on a file system that is not directly accessible by the SAS IT Service Level Management server, then you must define a server connection to a SAS IT Resource Management server that does have access.

Note: If the data source PDB was created on the same operating system as the active server, and is on a file system that is directly accessible by the SAS IT Service Level Management server (such as via NFS or UNC paths), then you do not have to define a separate server connection for accessing that PDB. △

A user-defined server connection is *from* the active SAS IT *Service Level Management* server and *to* a SAS IT *Resource Management* server that has one or more of the data source PDBs that you want to access.

From the Data workspace, you can perform the following tasks :

- define a server connection
- open a user-defined server connection
- remove a user-defined server connection
- copy a user-defined server connection
- paste a user-defined server connection
- rename a user-defined server connection
- connect to a user-defined server connection
- disconnect from a user-defined server connection

For more information about these tasks, click **`Help`**.

Defining a Server Connection

To define a server connection, perform the following steps:

1 From the **`Task`** list, select **`Define a Server Connection`**. A Server Connection dialog box opens.

2 On the **General** tab, enter the **Server Connection Name**.

3 Provide the properties for this server connection. On the **Definition** tab, enter the following properties:

- **Hostname** (required)
- **Port**
- **Default User Name**
- **Default Password**
- **Server OS**
- **Spawner**
- **Script**
- **SAS Command**
- **User Name Prompt**
- **User Name Prompt Timeout (sec)**
- **Password Prompt**
- **Password Prompt Timeout (sec)**
- **Command Prompt**
- **Command Prompt Timeout (sec)**

For more information about the properties in this dialog box, click **Help**.

Note: Although you can give arbitrary names to server connections, it is a good idea to give each server connection a name that is unique and meaningful. For example, if the server connection is to a host named oak and the server connection uses a spawner running on port 5414, you might name the server connection *oak-via-5414-spawner*. △

4 On the **Definition** tab, click the **Test** button to check the properties. (The test connection does not remain open.)

5 Click **OK** to save the definition of this server connection.

Opening a User-Defined Server Connection

To view or edit the properties of a user-defined server connection, you must first open it. To do so, perform these steps:

1 From the list of server connections, select the one that you want to use.

2 From the **Task** list, select **Open**. The Server Connection dialog box opens.

Note: You cannot open the [SLM Server]. △

You can view and edit the definition of the server connection.

3 When you finish, click **OK** if you want to save your changes; otherwise, click **Cancel**.

Removing a User-Defined Server Connection

To remove a user-defined server connection, perform these steps:

1 From the list of server connections, select the one that you want to remove.

Note: The [SLM Server] cannot be removed. △

2 From the **Task** list, select **Remove ▶ Yes**.

CAUTION:

If you remove a server that a component or a data source PDB depends on, you might disrupt the data collection process. △

Copying a User-Defined Server Connection

To copy a user-defined server connection, perform these steps:

1 From the list of server connections, right–click the one that you want to copy.

Note: The [SLM Server] cannot be copied. △

2 On the pop-up menu, click **Copy**.

Note: The copied server connection can be pasted only to the **Server Connections** fixed folder. △

Pasting a User-Defined Server Connection

To paste a user-defined server connection, perform these steps:

1 Right-click the **Server Connections** fixed folder.

Note: You cannot paste anything to the [SLM Server]. △

2 On the pop-up menu, click **Paste**.

Note: The copied server connection is added to the **Server Connections** fixed folder. Rename this copied server connection to give it a unique name. △

Renaming a User-Defined Server Connection

To change only the name of a user-defined server connection, perform these steps:

1 From the list of server connections, select the one that you want to rename.

Note: The [SLM Server] cannot be renamed. △

2 From the menu bar, select **File ▶ Rename**.

3 In the Rename dialog box, enter the new name in the **Rename to** field.

4 Click **OK**.

Note: If you rename a server connection, all references to the existing data sources (both in the catalog and in the contracts) will be automatically renamed. △

Connecting to a User-Defined Server Connection

If you want to view the structure and data of one or more data source PDBs that are not on or are not accessible from the active server, you can connect to the server that they are on by using a user-defined server connection.

Note: You can connect to a user-defined server without disconnecting from the other user-defined servers. △

You can connect to a server by using a server connection in one of the following ways:

□ *If you already know which server you want to connect to but you are not currently connected to it, then do the following*:

1 From the list of server connections, select the one whose server you want to connect to.

2 From the **Task** list, select **Connect to Server**, enter the requested logon information for that server, and click **OK**.

- *If you already know which server you want to connect to but do not know whether you are currently connected to it, then do the following*:
  1. From the list of server connections, select the one whose server you want to connect to.
  2. From the **Task** list, select **Open**. The Server Connection dialog box opens.
  3. On the **General** tab, do the following:
     - If the **Connection Status** is not connected [], click **Connect to Server** from the **Task** list, enter the requested logon information for that server, and click **OK**.
     - If the **Connection Status** is connected [], you do not need to do anything more to connect.
- *If you know which data source PDB you want to access but do not remember which server it is on or which server it is accessible from, then do the following*:
  1. From the list of data source PDBs, select the one that you want to use.
  2. From the **Task** list, select **Open**.
  3. On the **General** tab, do the following:
     - If the **Connection Status** is not connected [], click **Connect to Server** from the **Task** list, enter the requested logon information for that server, and click **OK**.

       This will open the server connection for *all* data source PDBs on that same host.
     - If the **Connection Status** is connected [], you do not need to do anything more to connect.

Disconnecting from a Server

You can disconnect from a server by using a user-defined server connection in one of the following ways:

- *If you already know which server you want to disconnect from and you are currently connected to it, then do the following*:
  1. From the list of server connections, select the one whose server you want to disconnect from.
  2. From the **Task** list, select **Disconnect from Server**.
- *If you already know which server you want to disconnect from but do not know whether you are currently connected to it, then do the following:*
  1. From the list of server connections, select the one whose server you want to disconnect from.
  2. From the **Task** list, select **Open**. The Server Connection dialog box opens.
  3. On the **General** tab, do the following:
     - If the **Connection Status** is connected [], click **Disconnect from Server** from the **Task** list.
     - If the **Connection Status** is not connected [], you do not need to do anything more to disconnect.

- *If you know which data source PDB you no longer want to access but do not remember which server it is on or which server it is accessible from, then do the following*:

  1. From the list of data source PDBs, select the one that you no longer want to access.
  2. From the **Task** list, select **Open**. The Data Source PDB dialog box opens.
  3. On the **General** tab, do the following:

     - If the **Connection Status** is connected [], select **Disconnect from Server** on the **Task** list.

       Note: This will close the server connection for *all* data source PDBs on that same host. If you want to continue to access the other data source PDBs on that host, wait to disconnect till you finish working with all of the data source PDBs. △

     - If the **Connection Status** is not connected [], you do not need to do anything more to disconnect.

When you exit from the GUI, you automatically disconnect from all servers. (Also, when you switch environments, you automatically disconnect from *all* server connections. This is because when you switch environments, the software logs you off and then logs you on again.)

CHAPTER

6

Recording Participants

Introduction to Recording Participants

The Participants workspace contains information about the organizations to which services are provided, information about the organizations that provide the services, and information about the contact people at those organizations. Depending on the nature of your contracts, customers (and providers) might be other companies, subdivisions of other companies, or other divisions of your own company.

The Participants workspace contains three fixed folders: the **Contacts** folder, the **Customers** folder, and the **Providers** folder.

- The *Contact* is a person or an entity who is associated with a service contract.

  In the **Contacts** folder, you can view, add, change, and delete existing definitions for contacts. If you need other contacts, you can create them here. A contact can be a customer or a provider.

- The *Customer* is the consumer of one or more services in a service contract and is part-owner (along with the Provider) of the service contract. Customers are typically business entities. They may be internal customers, for example, a department within the same business enterprise, or an external, separate business entity.

  In the **Customers** folder, you can view, add, change, and delete existing definitions for customers. If you need other customers, you can create them here.

- The *Provider* is the provider of services within a Service Contract and is part-owner (along with the Customer) of the service contract. Providers are typically business entities. For example, they may be an internal IT department or an external, separate business entity such as an outsourcer or supplier.

  In the **Providers** folder, you can view, add, change, and delete existing definitions for providers. If you need other providers, you can create them here.

Note: The information in the Participants workspace is primarily for documentation purposes. Some of the information is for internal use only and some is included in the contract reports. In addition, the information about e-mail notifications is used by the ETL process in order to send e-mails to a servicing contact when an objective's target has been missed. To view an example of this notification, see "Sample E-mail Notification" on page 293. △

If you want to organize the information in the fixed folders, you can create subfolders. Under the subfolders, you can create more folders, and so on.

- *To create a user-defined folder*, select the folder that you want to be the parent folder. Then, from the menu bar, select **File ▶ New ▶ Folder**. Enter the name and click **OK**.
- *To rename a user-defined folder*, select the folder that you want to rename. Then, from the **Task** list, select **Rename**, enter the name, and click **OK**.
- *To delete a user-defined folder*, select the folder that you want to delete. Then, from the **Task** list, select **Delete**. On the Confirm Delete message dialog box, click **Yes**.

For more information about folders, from the menu bar select **Help ▶ IT Service Level Management Help**.

Contacts

The **Contacts** folder contains information about one or more contacts. A contact is a person who is associated with a contract, SLA, service, or component. A contact can be a customer or a member of the customer's organization. Similarly, a contact can be a contract provider or a member of the contract provider's organization.

A servicing contact is a person (or entity) who is assigned the responsibility for a given node in the contract. Additionally, the servicing contact might receive e-mail notifications about missed targets.

To see the current list of contacts, double-click the folder.

Usually, for a given contract, the following is true:

- For the customer organization, you can designate:
  - the primary approver: the person who makes the final decisions about the contract
  - the primary contact: the main contact person at the organization

  The primary approver and primary contact can be the same person.

  You might also have information about others in that organization who have an interest in the service levels in the contract.
- For the provider organization, you can designate:
  - the primary approver: the person who makes the final decisions about the contract
  - the primary contact: the main contact person at the organization

  The primary approver and primary contact can be the same person.

In addition, for each node that is below the contract node in the contract hierarchy, you will probably want information about the servicing contact for that node.

You might also have information about others in the organization who have an interest in the service levels in the contract.

The term *contacts* refers to all of these people.

Creating a Contact

To create a contact, perform these steps:

1 From the **Task** list, select **Create a Contact**. The Contact dialog box opens.

2 You can (but are not required to) select one or more of the following options.

- □ Click **Customer** to associate the person with a customer organization. (For the purposes of the contract, the person is considered to be a customer.)
- □ Click **Provider** to associate the person with a contract provider organization. (For the purposes of the contract, the person is considered to be a contract provider.)
- □ Click **SLA/Service/Component** to associate the person with one or more of the following: a contract, service level agreement node, service node, or component node.
- □ If the servicing contact is to be notified by e-mail that any of his contracts (or nodes) have missed their targets, then click **Receive e-mail notification for any missed targets (default)**. (A default can be set for this servicing contact to always receive e-mail notifications when their nodes in a contract miss targets. This default can be overridden for a particular node by removing the check mark from the **E-mail this contact for any missed target definitions** on that contract, SLA, service, or component.) This default is not retroactive, and it does not affect assignments already defined.

  Note: With the e-mail notification feature, the servicing contact of any contract or node will receive an e-mail that lists the nodes that have failed to meet their targets. △
- □ Enter the requested information about the contact. This includes contact name, job title, department, telephone, pager, FAX, e-mail address, business address, city, state, and country. For information about the remaining fields, click **Help**.

  Note: If you want to use the e-mail notification feature, you must enter the e-mail address of a contact who is operating as the servicing contact for a component, service, SLA, or contract.

  CAUTION:
  This address is not validated at the time of entry by the software. If the e-mail address is not correct, then the e-mail notification cannot be delivered. △

  △

3 If you want to attach a note to this contact, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

To add an attachment, perform the following steps:

a From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
b You can, optionally, enter the name of the attachment.

c Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated contact on the reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.

f Click **OK** to save the attachment.

4 Click **OK** to save the information about this contact person.

Opening a Contact

To view or edit the properties of a contact, perform these steps:

1 Select the contact.

2 From the **Task** list, select **Open**.

3 For information about the fields, click **Help**.

4 When you finish, click **OK** if you want to save any changes; otherwise, click **Cancel**.

Renaming a Contact

To change the contact name only, perform these steps:

1 Select the contact.

2 From the **Task** list, select **Rename**.

3 Enter the new name and click **OK**.

Deleting a Contact

To delete a contact, perform these steps:

1 Select the contact.

2 From the **Task** list, select **Delete**. On the Confirm Delete message dialog box, click **Yes**.

Note: If you delete a contact, all the contact links to customers and providers in the Participants workspace are also removed. Before this occurs, a confirmation message is displayed that prompts you to decide whether to continue with the deletion. △

For more information about contacts and tasks that are related to contacts, click **Help** in the Contact dialog box.

Customers

The **Customer** folder contains information about one or more customers. To see the current list of customers, double-click the folder.

Usually, for a given contract, you need to supply a customer unless the customer already exists in the **Customer** folder.

Note: You can also create a customer from the **Parties** tab of the Contract window. △

Creating a Customer

To create a customer, perform these steps:

1. From the **Task** list, select **Create a Customer**. The Customer dialog box opens.
2. On the **General** tab, enter a customer name. You can enter (but are not required to enter) a company and business unit.

   Note: The customer name can be a person's name, a company's name, a business unit's name, a combination of one or more of these, or anything else that you use to identify the organization. In the contract's reports, the value in the **Customer Name** field is the main identifier of the customer organization. △
3. On the **Contacts** tab, select the arrow next to the **Add Existing Contact** drop-down list. This drop-down list displays customer contacts. It is derived from the contacts in the Participants workspace where the **Customer** option is selected. Select one of the contacts and then click **Add**. Repeat until you have selected and added all the contacts that you want to associate with this customer organization.

   Note: It is possible to add the contacts in the opposite order. That is, on the **Contacts** tab, you can click **New Contact**, provide the information about the contact, and then click **OK**. The information will be associated with this customer and also saved to the **Contacts** folder. △
4. When you finish adding contacts, identify the contact that you want to designate as the primary contact and the contact that you want to designate as the primary approver, if you have that information.
5. If you want to attach a note to this customer, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

   To add an attachment, perform the following steps:

   a. From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
   b. You can, optionally, enter the name of the attachment.
   c. Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

      Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △
   d. To display the attachment with its associated customer on the reports, select **Include this attachment**.

      Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.

f Click **OK** to save the attachment.

6 Click **OK**.

Opening a Customer

To view or edit the properties of a customer, perform these steps:

1 From the **Customer** folder, select the customer.

2 From the **Task** list, select **Open**.

3 When you finish, click **OK** if you want to save any changes; otherwise, click **Cancel**.

Renaming a Customer

To rename a customer, perform these steps:

1 From the **Customer** folder, select the customer.

2 From the **Task** list, select **Rename**.

3 Enter the new name and click **OK**.

Deleting a Customer

To delete a customer, perform these steps:

1 From the **Customer** folder, select the customer.

2 From the **Task** list, select **Delete**. On the Confirm Delete message dialog box, click **Yes**.

For more information about customers and tasks related to customers, click **Help** in the Customer dialog box.

Providers

The **Providers** folder contains information about one or more contract providers. To see the current list of providers, double-click the folder.

Usually, for a given contract, you need to supply a provider unless a provider already exists in the **Provider** folder.

Creating a Provider

To create a provider, perform these steps:

1 From the **Task** list, select **Create a Provider**. The Provider dialog box opens.

2 On the **General** tab, enter a provider name. You can enter (but are not required to enter) a company and business unit.

Note: The provider name can be a person's name, a company's name, a business unit's name, a combination of one or more of these, or anything else that you use to identify the provider. In the contract's reports, the value in the **Provider Name** field is the main identifier of the provider organization. △

3 On the **Contacts** tab, click the arrow next to the **Add Existing Contact** drop-down list. A list appears that displays provider contacts. This list is derived from the contacts in the Participants workspace where the **Provider** option is selected. Select one of the contacts and then click **Add**. Repeat until you have selected and added all the contacts at this organization that you want to associate with the contract as a whole.

Note: It is possible to add contacts in the opposite order. That is, on the **Contacts** tab, you can click **New Contact**, provide the information about the contact, and then click **OK**. The information will be associated with this provider and also saved to the **Contacts** folder. △

4 When you finish adding contacts, identify the contact that you want to designate as the primary contact and the contact that you want to designate as the primary approver, if you have that information.

5 If you want to attach a note to this provider, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

To add an attachment, perform the following steps:

a From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
b You can, optionally, enter the name of the attachment.
c Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated provider on the reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.
f Click **OK** to save the attachment.

6 Click **OK**.

Opening a Provider

To view or edit the properties of a provider, perform these steps:

1. From the **Provider** folder, select the provider.
2. From the **Task** list, select **Open**.
3. If you make changes that you want to save, click **OK**; otherwise, click **Cancel**.

Renaming a Provider

To change the name of the provider, perform these steps:

1. From the **Provider** folder, select the provider.
2. From the **Task** list, select **Rename**.
3. Enter the name and click **OK**.

Deleting a Provider

To delete a provider, perform these steps:

1. From the **Provider** folder, select the provider.
2. From the **Task** list, select **Delete**. On the Confirm Delete message dialog box, click **Yes**.

For more information about providers and tasks related to providers, click **Help** in the Provider dialog box.

CHAPTER

7

Building the Catalog

Introduction to Building the Catalog

The Catalog workspace contains information about service level agreements (SLAs), services, components, range sets, and range descriptions that can all be used as building blocks to create a contract.

The Catalog workspace contains five fixed folders: the **SLAs** folder, the **Services** folder, the **Components** folder, the **Range Sets** folder, and the **Range Descriptions** folder.

- In the **SLAs** folder, you can view, add, change, delete, and rename existing definitions for service level agreements. If you need other service level agreements, you can create them here.

- In the **Services** folder, you can view, add, change, delete, and rename existing definitions for services. If you need other services, you can create them here.
- In the **Components** folder, you can view, add, change, delete, and rename existing definitions for components. If you need other components, you can create them here.
- In the **Range Sets** folder, you can view, add, change, delete, and rename existing definitions for range sets. If you need other range sets, you can create them here.
- In the **Range Descriptions** folder, you can view, add, change, and delete existing definitions for range descriptions. You cannot delete any of the fixed range descriptions (Out-of-bounds High, Out-of-bounds Low, and Missing Value).

  If you need to delete a range description that is currently in use, you must replace it with another range description.

Note: If you change or delete items in the catalog, all the items that are linked to that item in the Catalog workspace are also changed or removed. Before this occurs, a confirmation message box is displayed that prompts you to decide whether to continue with the change or the deletion. △

If you want to organize the information in the fixed folders, you can create subfolders under all the fixed folders *except* the **Range Descriptions** folder. Under the subfolders, you can create more folders, and so on. All the subfolders in the Catalog workspace are displayed in alphanumeric ascending order.

- *To create a user-defined folder*, select the folder that you want to be the parent folder and then, from the menu bar, select **File ▶ New ▶ Folder**. Enter the name of the folder and click **OK**.
- *To rename a user-defined folder*, select the folder and then, from the **Task** list, select **Rename**, enter the new name, and click **OK**.
- *To delete a user-defined folder*, select the folder and then, from the **Task** list, select **Delete ▶ Yes**.

For more information about folders, from the menu bar select **Help ▶ IT Service Level Management Help**.

Service Level Agreements

The SLAs folder contains information about one or more service level agreements. To see the current list of service level agreements, double-click the folder.

For a given contract, you can use any SLA in the catalog to assist in building your contract, to reduce the effort required to build other contracts, and to minimize inconsistencies among contracts.

Creating a Service Level Agreement

To create a service level agreement, perform these steps:

1. From the **Task** list, select **Create a Service Level Agreement**. The Service Level Agreement dialog box opens.
2. On the **General** tab, enter the service level agreement name and optionally, the version number of the SLA. In the contract's reports, the value in the **SLA Name** field is the main identifier of the node.

   Note: The name must not be blank. It is a good idea to use a name that is unique across all nodes in the catalog. △

You can also select the servicing contact for this SLA (from the list of contact people for whom you selected the Servicing Contact option). The telephone number and the e-mail address of the selected contact will be displayed.

If you want this contact name to be visible on the reports that are generated for the customer, make sure that **The contact should be visible on Customer Reports** option is checked.

If an objective of this SLA misses its target, you can notify the servicing contact who is associated with this SLA. To do so, make sure that the **E-mail this contact for any missed target definitions** option is checked. You must also enter the contact's e-mail address on the appropriate Contact dialog box. Then, if the ETL process detects that a target has been missed, it will notify the servicing contact by e-mail. (If you have checked this option as the default on the Contact dialog box, the option on this tab will already have a check mark.)

3 From the **Hierarchy** tab, select **Catalog** to add items.

In the **SLA Hierarchy** column, select the name of this service level agreement, click the right mouse button, and click **Catalog**. A list of already-created services and components is displayed in the Mini-Catalog dialog box. Select a service or a component and then click **Add**. Repeat until you have identified all desired services and components to build this SLA.

The data that is displayed in the other columns is a summary of the objective settings for all items in the hierarchy. It depends on the following criteria:

- the selection in the **Filter by Objective Type** drop-down list
- the objective values of the objective settings for all items in the hierarchy

4 If you want to use an availability objective for this SLA, you can define it on the **Availability** tab by performing these steps:

a From the **Objectives** tab of the SLA dialog box, click the **Activate for Processing** option for Availability.

Note: Until the **Activate for Processing** option has been checked, the Extract, Transform, and Load job (ETL) will not process this objective and the service level data for this objective will not be gathered. △

b From the drop-down list next to the Availability option, select an objective label. (You can also enter a label.)

c Select the **Availability** tab in the lower section of the window.

d *To specify the indicator*, perform the following actions:

- If you want to use a supplied calculation method, in the **Input metric from child nodes** field, select **Indicator** or **Score** and in the **Calculation Method** field, select the method that you want to use.
- If you want to use a custom calculation method, in the **Input metric from child nodes** field click **Custom Code** and then click **Code**. Provide the SAS code that you want to use for this indicator and then click **OK**.

e *To specify the range set*, perform this action:

- In the **Range Set Definition** section of the tab, click **Catalog**. The Mini-Catalog opens and displays a list of already-created range sets. Select the appropriate range set for this objective and then click **Add**. A copy of the range set is created for this objective. The ranges that constitute this range set are displayed along with their scores, if scores have been assigned to the ranges. The name of the range set is displayed above the section that lists the ranges.

Note: For information about how to change the attributes of a range set or a range description, see “Range Sets” on page 76 and “Range Descriptions” on page 78. △

f *To specify the target*, perform the following actions:

- □ If you want to use an *indicator-based target*, select **Indicator** (for this objective) from the **Comparison Value** drop-down list, select the comparison operator that you want to use from the **Operator** drop-down list, and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** and **Target Statistic** drop-down lists, select the period and statistic that you want to use.
- □ If you want to use a *score-based target*, select **Score** (for this objective) from the **Comparison Value** drop-down list, select the comparison operator that you want to use from the **Operator** drop-down list, and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** drop-down list, select the compliance period that you want to use. (**Target Statistic** is not available for score-based targets.)

  Note: **Score** is available as a comparison value only if a range set has been defined. △

Note: If you request a report for a contract that includes a copy of this node, the report’s data about this objective is based on the target, period, and statistic that you select here, unless you change one or more of these values in the contract’s copy of this node. △

5 If you want to use the response time objective, throughput objective, or custom objective for this SLA, from the **Objectives** tab click the **Activate for Processing** option for the objective type you want to work with. Then, select that tab in the lower section of the window and make the appropriate selections on the **Response Time** tab, **Throughput** tab, or **Custom Objective** tab, respectively.

Specifying the calculation method

If you specify the Weighted Mean calculation method, then the weights that are used for the input metrics for the respective child nodes are ultimately accumulated from the **Weight By** variable for the component nodes that make up the lower levels of the hierarchy that defines the SLA. You can specify this **Weight By** variable on the **Data** tab of the respective component(s).

Specifying the target statistic

You can choose the Weighted Mean as the target statistic only if you chose the Weighted Mean as the calculation method. Otherwise, the Weighted Mean target statistic will not be available. If the Weighted Mean calculation method has been selected, the Mean target statistic will not be available.

6 If you want to use the composite objective, from the **Objectives** tab, click the **Activate for Processing** option for the composite objective type. Then, select that tab in the lower section of the window and make the appropriate selections on the **Composite** tab.

7 If you want to base the calculations for this node on a subset of the data, select the **Date Filter** tab. On the tab, select a date filter pattern.

- □ If you click **All Inclusive**, all the data is used. This is the default pattern.
- □ If you click **Continuous**, all the data from the effective date and time to the expiration date and time is used. Specify these dates and times.
- □ If you click **Daily**, all the data from the effective time to the expiration time is used on each date from the effective date to the expiration date. Specify these times and dates.

- If you click **Weekly**, all the data from the effective time to the expiration time on the selected day or days of the week is used for each week from the effective date to the expiration date. Specify these times, days, and dates.
- If you click **Monthly**, all the data from the effective time to the expiration time on the selected day of the month is used for each month from the effective date to the expiration date. Specify these times, this day, and these dates.

Note: To set up a recurring pattern for collecting data, use the daily, weekly, or monthly date filters. For more information about the **Date Filter** tab, click **Help**. △

The date filter for a node is independent of the date filters for other nodes.

For some examples that show how a date filter might be used, see the section about the **Date Filter** tab in "Creating a Component" on page 70.

8 If you want to add notes about this service level agreement, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

To add an attachment, perform the following steps:

a From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
b You can, optionally, enter the name of the attachment.
c Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated SLA on the reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.
f Click **OK** to save the attachment.

9 Then click **OK** to save the properties of this service level agreement.

Opening a Service Level Agreement

To view or edit the properties of a service level agreement, perform these steps:

1 From the **SLAs** folder, select the service level agreement.

2 From the **Task** list, click **Open**.

3 When you finish viewing or making changes, click **OK** if you want to save any changes; otherwise, click **Cancel**.

Renaming a Service Level Agreement

To rename a service level agreement, perform these steps:

1. From the **SLAs** folder, select the service level agreement.
2. From the **Task** list, select **Rename**.
3. Enter the new name and click **OK**.

Deleting a Service Level Agreement

To delete a service level agreement, perform these steps:

1. From the **SLAs** folder, select the service level agreement.
2. From the **Task** list, select **Delete**.
3. Click **Yes** in the Delete confirmation message box.

Note: If you want to discontinue data collection for this SLA but wish to retain the selections for possible future use, do not delete the SLA. Instead, on the **Objectives** tab of the SLA dialog box, clear the **Activate for Processing** option for each of the objective types. Then click **OK**. △

For information about the Service Level Agreement dialog box, click **Help**.

For information about nodes, see "Introduction to Understanding Contract Structure" on page 9.

For information about objectives, see "Introduction to Understanding Contract Objectives" on page 13.

Services

The **Services** folder contains information about services. To see the current list of services, double-click the folder.

For a given contract, you can use any service in the catalog to assist in building your contract, to reduce the effort required to build other contracts, and to minimize inconsistencies among contracts.

Creating a Service

To create a service, perform these steps:

1. From the **Task** list, select **Create a Service**. The Service dialog box opens.
2. On the **General** tab, enter the service name and, optionally, the version number of the service. In the contract's reports, the value in the **Service Name** field is the main identifier of the node.

   Note: The name must not be blank. It is a good idea to use a name that is unique across all nodes in the catalog. △

   You can also select the servicing contact for this service (from the list of contact people for whom you selected the Servicing Contact option). The telephone number and the e-mail address of the selected contact will be displayed.

   If you want this contact name to be visible on the reports that are generated for the customer, make sure that the **The contact should be visible on Customer Reports** option is checked.

If an objective of this service misses its target, you can notify the contact who is associated with this service. To do so, make sure that the **E-mail this contact for any missed target definitions** option is checked. You must also enter the contact's e-mail address on the appropriate Contact dialog box. Then, if the ETL process detects that a target has been missed, it will notify the servicing contact by e-mail. (If you have checked this option as the default on the Contact dialog box, the option on this tab will already have a check mark.)

3 From the **Hierarchy** tab, select catalog items to add.

In the **Service Hierarchy** column, select the name of this service, click the right mouse button, and select **Catalog**. A list of already-created components is displayed in the Mini-Catalog dialog box. Select a service or a component and then click **Add**. Repeat until you have identified all the desired services or components to build this service.

The data that is displayed in the other columns is a summary of the objective settings for all items in the hierarchy. It depends on the following criteria:

- the selection in the **Filter by Objective Type** drop-down list
- the objective values of the objective settings for all items in the hierarchy

4 If you want to use an availability objective for this service, you can define it on the **Availability** tab by performing these steps:

a From the **Objectives** tab in the Service dialog box, click the **Activate for Processing** option for Availability.

Note: Until the **Activate for Processing** option has been checked, the Extract, Transform, and Load job (ETL) will not process this objective and service level data for this objective will not be gathered. △

b From the drop-down list next to the Availability option, select an objective label. (You can also enter a label.)

c Select the **Availability** tab in the lower section of the window.

d *To specify the indicator*, perform the following actions:

- If you want to use a supplied calculation method, in the **Input metric from child nodes** field click **Indicator** or **Score**, and in the **Calculation Method** field select the method that you want to use.
- If you want to use a custom calculation method, in the **Input metric from child nodes** field select **Custom Code** and then click **Code**. Provide the SAS code that you want to use for this indicator and then click **OK**.

e *To specify the range set*, perform this action:

- In the **Range Set Definition** section of the tab, click **Catalog**. The Mini-Catalog opens and displays a list of already-created range sets. Select the appropriate range set for this objective and then click **Add**. A copy of the range set is created for this objective. The ranges that constitute this range set are displayed along with their scores, if scores have been assigned to the ranges. The name of the range set is displayed above the section that lists the ranges.

Note: For information about how to change the attributes of a range set or a range description, see "Range Sets" on page 76 and "Range Descriptions" on page 78. △

f *To specify the target*, perform the following actions:

- If you want to use an *indicator-based target*, select **Indicator** (for this objective) from the **Comparison Value** drop-down list, select the

comparison operator that you want to use from the **Operator** drop-down list, and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** and **Target Statistic** drop-down lists, select the period and statistic that you want to use.

- If you want to use a *score-based target*, select **Score** (for this objective) from the **Comparison Value** drop-down list, select the comparison operator that you want to use from the **Operator** drop-down list, and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** drop-down list, select the compliance period that you want to use. (**Target Statistic** is not available for score-based targets.)

  Note: **Score** is available as a comparison value only if a range set has been defined. △

Note: If you request a report for a contract that includes a copy of this node, the report's data about this objective is based on the target, period, and statistic that you select here, unless you change one or more of these values in the contract's copy of this node. △

5 If you want to use the response time objective, throughput objective, or custom objective for this service, from the **Objectives** tab click the **Activate for Processing** option for the objective type you want to work with. Then, select that tab in the lower section of the window and make the appropriate selections on the **Response Time** tab, **Throughput** tab, or **Custom Objective** tab, respectively.

Specifying the calculation method

If you specify the Weighted Mean calculation method, then the weights that are used for the input metrics for the respective child nodes are ultimately accumulated from the **Weight By** variable for the component nodes that make up the lower levels of the hierarchy that define the service. You can specify this **Weight By** variable on the **Data** tab of the respective component(s).

Specifying the target statistic

You can choose the Weighted Mean as the target statistic only if you chose the Weighted Mean as the calculation method. Otherwise, the Weighted Mean target statistic will not be available. If the Weighted Mean calculation method has been selected, the Mean statistic will not be available.

6 If you want to use the composite objective, from the **Objectives** tab, click the **Activate for Processing** option for the composite objective type. Then, select that tab in the lower section of the window and make the appropriate selections on the **Composite** tab.

7 If you want to base the calculations for this node on a subset of the data, select the **Date Filter** tab. On the tab, select a date filter pattern.

- If you click **All Inclusive**, all the data is used. This is the default pattern.
- If you click **Continuous**, all the data from the effective date and time to the expiration date and time is used. Specify these dates and times.
- If you click **Daily**, all the data from the effective time to the expiration time is used on each date from the effective date to the expiration date. Specify these times and dates.
- If you click **Weekly**, all the data from the effective time to the expiration time on the selected day or days of the week is used for each week from the effective date to the expiration date. Specify these times, days, and dates.
- If you click **Monthly**, all the data from the effective time to the expiration time on the selected day of the month is used for each month from the effective date to the expiration date. Specify these times, this day, and these dates.

Note: To set up a recurring pattern for collecting data, use the daily, weekly, or monthly date filters. For more information about the **Date Filter** tab, click **Help** in the Contract window. △

The date filter for a node is independent of the date filters for other nodes.

For some examples that show how a date filter might be used, see the section about the **Date Filter** tab in "Creating a Component" on page 70.

8 If you want to add notes about this service, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

To add an attachment, perform the following steps:

a From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
b You can, optionally, enter the name of the attachment.
c Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated service on the reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.
f Click **OK** to save the attachment.

9 Then, click **OK** to save the properties of this service.

Opening a Service

To view or edit the properties of a service, perform these steps:

1 From the **Services** folder, select the service.

2 From the **Task** list, select **Open**.

3 When you finish viewing or making changes, click **OK** if you want to save any changes; otherwise, click **Cancel**.

Renaming a Service

To rename a service, perform these steps:

1 From the **Services** folder, select the service.

2 From the **Task** list, select **Rename**.

3 Enter the new name and click **OK**.

Deleting a Service

To delete a service, perform these steps:

1 From the **Services** folder, select the service.

2 From the **Task** list, select **Delete**.

3 Click **Yes** in the Delete confirmation message box.

Note: If you want to discontinue data collection for this service but wish to retain the selections for possible future use, do not delete the service. Instead, on the **Objectives** tab of the Service dialog box, clear the **Activate for Processing** option for each of the objective types. Then, click **OK**. △

For information about the services and the tasks related to services, click **Help** in the Service dialog box.

For information about nodes, see "Introduction to Understanding Contract Structure" on page 9. For information about objectives, see "Introduction to Understanding Contract Objectives" on page 13.

Components

The Components folder contains information about one or more components. To see the current list of components, double-click the folder.

For a given contract, you can use any component in the catalog to assist in building your contract, to reduce the effort required to build other contracts, and to minimize inconsistencies among contracts.

Creating a Component

To create a component, perform these steps:

1 From the **Task** list, select **Create a Component**. The Component dialog box opens.

2 On the **General** tab, enter the component name and optionally, the version number of the component. In the contract's reports, the value in the **Component Name** field is the main identifier of the node.

Note: The name must not be blank. It is a good idea to use a name that is unique across all nodes in the catalog. △

You can also select the servicing contact for this component (from the list of contact people for whom you selected the Servicing Contact option). The telephone number and the e-mail address of the selected contact will be displayed.

If you want this contact name to be visible on the reports that are generated for the customer, make sure that the **The contact should be visible on Customer Reports** option is checked.

If an objective of this component misses its target, you can notify the contact who is associated with this component. To do so, make sure that the **E-mail this contact for any missed target definitions** option is checked. You must also enter the contact's e-mail address on the appropriate Contact dialog box. Then, if the ETL process detects that a target has been missed, it will notify the servicing contact by e-mail. (If you have checked this option as the default on the Contact dialog box, the option on this tab will already have a check mark.)

3 Use the **Data** tab to:

- define the data you want to use
- define the resources that you want to use

Defining the data

Data is specified by identifying the data source PDB, table, and measure for each availability, response time, throughput, and custom objective that you intend to use. For the response time, throughput, and custom objectives, you can optionally specify a Weight By variable as well as an Event Timestamp.

To define the data, perform these steps:

a On the **Objectives** tab, click to activate the objectives that you want to work with.

b Then, select the **Data** tab. (You can enter information on this **Data** tab for any objectives whether or not they are activated. However, only the activated objectives will be processed by the ETL step.)

The **Data** tab displays three columns of information for each of the objectives that have been activated. (Only one entry is allowed for each objective type. Thus, a component can have no more than one availability objective, and no more than one response time objective, and so on.)

To enter or change the information for any objective on this component, click **Define Data** to invoke the Data Definition Wizard. The Data Definition Wizard guides you through a series of steps that will specify a valid measure for each objective. If a weighting factor is needed, the Wizard can also ensure that a valid Weight By variable is specified. (The Weight By variable is available only on the response time, throughput, and custom objectives.)

To accomplish these tasks, the Wizard presents a series of pages that prompt you for the following information:

- *Objective*: select an objective from this list: Availability, Response Time, Throughput, or Custom Objective.
- *Data Source PDB*: select from a list of PDBs.

  The PDB where the data is stored must be connected during this process so that the Wizard can ensure that the data that is specified for each objective is complete. If you are not connected to the PDB that you select, the Server Log On page opens and requests your user name and password.
- *Measure*: select a measure from the list of all tables that are available in the selected PDB.

  To select the measure, perform these steps:

  i Select a table. The measures that are available in that table will be displayed. Tables can be either interval- or event-based. However, event-based tables are only available for response time, throughput, and custom objectives.

  ii Then select the measure that you want to use in the calculation method.

  If the measure you choose is from an event-based table, the Wizard prompts you to select a timestamp. The event timestamp can be used to determine if an event occurred within the time period specified by a target. The default timestamp is DATETIME.
- *Weight By Variable*: select from the list of variables that are available in the table you selected. (For event-based data only, a count of all events, Count, is also available for this purpose.)

If you intend to specify the Weighted Mean calculation (on the response time or throughput objective or custom objective tab only), then you must also specify the Weight By variable to be used as the weighting factor.

Note: You can also choose to accumulate the weights in order to enable the parent nodes to have weights, even if you do not intend to calculate the weighted mean at the component node. For example, for the component node objective you could choose a non-weighted calculation method like Percentile, but still specify the Weight By variable. In that case, the parent node would be able to specify a calculation method of Weighted Mean. △

c Click **Finish** in order to open the summary page. This page displays the information that you have entered to define the data for this objective.

d When you are satisfied with your selections, you can use the Wizard to enter the required information for the other objectives that are activated.

Define the resources.

A resource is an item of hardware or software or any other entity for which data is collected and which, alone or with other resources, makes up a component. Resources are specified by identifying the column and row of the table of the data source PDB in which they are located.

Note: Although you can select any variables to define your resources, the selected variables should reflect categorization information such as machine names or user IDs. It is inappropriate to select continuous analytic variables; for example, don't specify a variable that represents the disk I/O rate unless the values for that variable are categories instead of the actual raw I/O rates. Most resources can be specified with far fewer than 100 resource rows. △

To define the resources, perform these steps:

a Click **Columns** and select the columns that contain the identifiers for the resources. For example, if the resources are identified by a value of the DOMAIN column and a value of the MACHINE column, select DOMAIN and MACHINE.

Note: Columns that are available for selection are variables that are the intersection of all columns found across the defined data source PDBs for this component. (In other words, the columns that are available are those that exist *in common* across all of the data source PDBs that you have specified for this component.) If you make changes to a table name or a data source PDB name after resources are defined, you must ensure that this list is still valid. △

b Then, click **Rows** and select the values that identify the resources for this component.

Note: Rows that are available for selection are the union of all rows found across the defined columns for this component. (Rows are the observations that exist in the data source PDBs that you have specified for this component.) If you make changes to data source PDBs, table names, and measures, those changes might alter the rows that are available. You must ensure that the resource list is still valid. △

For example, if you want to use data from three resources (US Operations HAL6000, US Operations Server/01, and US Operations ALLSERV), select values for these rows:

```
      RESOURCES
DOMAIN        MACHINE

US Operations HAL6000
US Operations Server/01
US Operations ALLSERV
```

4 If you want to use an availability objective for this component, you can define it on the **Availability** tab by performing these steps:

   a From the **Objectives** tab of the Component dialog box, click the **Activate for Processing** option for Availability.

   Note: Until the **Activate for Processing** option has been checked, the Extract, Transform, and Load job (ETL) will not process this objective and performance data for this objective will not be gathered. △

   b From the drop-down list next to the Availability option, select an objective label. (You can also enter a label.)
   c Select the **Availability** tab in the lower section of the window.
   d *To specify the indicator*, perform the following actions:
      □ If you want to use a supplied calculation method, in the **Input metrics from resources** field select **Measure** and in the **Calculation Method** field select the method that you want to use.
      □ If you want to use a custom calculation method, in the **Input metrics from resources** field select **Custom Code** and then click **Code**. Provide the SAS code that you want to use for this indicator and then click **OK**.
   e *To specify the range set*, perform this action:
      □ In the **Range Set Definition** section of the tab, click **Catalog**. The Mini-Catalog opens and displays a list of already-created range sets. Select the appropriate range set for this objective and then click **Add**. A copy of the range set is created for this objective. The ranges that constitute this range set are displayed along with their scores, if scores have been assigned to the ranges. The name of the range set is displayed above the section that lists the ranges.

      Note: For information about how to change the attributes of a range set or a range description, see "Range Sets" on page 76 and "Range Descriptions" on page 78. △

   f *To specify the target*, perform the following actions:
      □ If you want to use an *indicator-based target*, select **Indicator** (for this objective) from the **Comparison Value** drop-down list, select the comparison operator that you want to use from the **Operator** drop-down list, and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** and **Target Statistic** drop-down lists, select the period and statistic that you want to use.
      □ If you want to use a *score-based target*, select **Score** (for this objective) from the **Comparison Value** drop-down list, select the comparison operator that you want to use from the **Operator** drop-down list, and enter the value that you want to use in the **Target Value** field. You can

optionally specify a label for the target value. Also, from the **Compliance Period** drop-down list, select the compliance period that you want to use. (**Target Statistic** is not available for score-based targets.)

Note: **Score** is available as a comparison value only if a range set has been defined. △

Note: If you request a report for a contract that includes a copy of this node, the report's data about this objective is based on the target, period, and statistic that you select here, unless you change one or more of these values in the contract's copy of this node. △

5 If you want to use the response time objective, throughput objective, or custom objective, from the **Objectives** tab, click the **Activate for Processing** option for the objective type you want to work with. Then, select that tab in the lower section of the window and make the appropriate selections on the **Response Time** tab, **Throughput** tab, or **Custom Objective** tab, respectively.

Specifying the calculation method

If you specify the Weighted Mean calculation method, then you must specify the **Weight By** variable on the **Data** tab of the respective component.

Specifying the target statistic

You can choose the Weighted Mean target statistic only if you chose the Weighted Mean as the calculation method. Otherwise, the Weighted Mean target statistic will not be available. If the Weighted Mean calculation method has been selected, the Mean target statistic will not be available.

6 If you want to use the composite objective, from the **Objectives** tab, click the **Activate for Processing** option for the composite objective type. Then, select that tab in the lower section of the window and make the appropriate selections on the **Composite** tab.

7 If you want to base the calculations for this node on a subset of the data, select the **Date Filter** tab. On the tab, select a date filter pattern.

- If you click **All Inclusive**, all the data is used. This is the default pattern.
- If you click **Continuous**, all the data from the effective date and time to the expiration date and time is used. Specify these dates and times.
- If you click **Daily**, all the data from the effective time to the expiration time is used on each date from the effective date to the expiration date. Specify these times and dates.
- If you click **Weekly**, all the data from the effective time to the expiration time on the selected day or days of the week is used for each week from the effective date to the expiration date. Specify these times, days, and dates.
- If you click **Monthly**, all the data from the effective time to the expiration time on the selected day of the month is used for each month from the effective date to the expiration date. Specify these times, this day, and these dates.

Note: To set up a recurring pattern for collecting data, use the daily, weekly, or monthly date filters. For more information about the **Date Filter** tab, click **Help** in the Contract window.

For examples of how you can use date filters more effectively, see "Examples of Date Filter Usage" on page 313. △

8 If you want to attach a note to this component, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

To add an attachment, perform the following steps:

a From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
b You can, optionally, enter the name of the attachment.
c Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated component on the reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.
f Click **OK** to save the attachment.

9 Click **OK** to save the properties of this component.

Opening a Component

To view or edit the properties of a component, perform these steps:

1 From the **Components** folder select the component.
2 From the **Task** list, select **Open**.
3 When you finish viewing or making changes, click **OK** if you want to save any changes; otherwise, click **Cancel**.

Renaming a Component

To change the name of a component, perform these steps:

1 From the **Components** folder, select the component.
2 From the **Task** list, select **Rename**.
3 Enter the new name and click **OK**.

Deleting a Component

To delete a component, perform these steps:

1 From the **Components** folder, select the component.
2 From the **Task** list, select **Delete**.
3 Click **Yes** in the Delete confirmation message box.

Note: If you want to discontinue data collection for this component but wish to retain the selections for possible future use, do not delete the component. Instead, on the **Objectives** tab of the Component dialog box, clear the **Activate for Processing** option for each objective type. Then click **OK**. △

For information about components and tasks related to components, click **Help** in the Component dialog box.

For information about nodes, see "Introduction to Understanding Contract Structure" on page 9. For information about objectives, see "Introduction to Understanding Contract Objectives" on page 13.

Range Sets

The **Range Set** folder contains information about one or more range sets. To see the current list of range sets, double-click the folder.

For a given contract, you can use any range set in the catalog to assist in building your contract, to reduce the effort required to build other contracts, and to minimize inconsistencies among contracts.

Creating a Range Set

To create a range set, perform these steps:

1. From the **Task** list, select **Create a Range Set**. The Range Set dialog box opens.
2. On the **General** tab, enter a range set name.

   Note: The name must not be blank. It is a good idea to use a name that is unique across all range sets in the catalog. △
3. On the **Ranges** tab, click **Add**. The **Add Range** dialog box opens. Provide the information about a range by entering the limits, whether or not the limits are inclusive, and the score. You can also select a range description for the range. The range description will assign to the range a color and a label that will be used in the reports. Then click **OK**. The range is saved and is added to this range set. Repeat until you finish adding all the ranges that you want to include with this range set.
4. If you want to change the score of the predefined ranges, select a predefined range and then click **Edit**. The **Edit Range** dialog box opens. Change the score and then click **OK**. Repeat until you finish changing scores of the predefined ranges.

   If you want to change the color and label of a predefined range, select the range description of the predefined range from the **Range Descriptions** folder in the Catalog workspace and click **Open**. You can change the color and the label of this range description, but not the name. Then, click **OK** to save your changes.
5. If you want to delete a range, select the range and then click **Delete**. When the Delete confirmation message box displays, click **Yes**.

   Note: You cannot delete or rename the predefined ranges. △
6. If you want to attach a note to this range set, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

   To add an attachment, perform the following steps:

   a. From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
   b. You can, optionally, enter the name of the attachment.
   c. Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated range set on the reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on the reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.

f Click **OK** to save the attachment.

7 Click **OK** to save the changes to the range set.

Deleting a Range Set

To delete a range set, perform these steps:

1 From the **Range Set** folder, select the range description.

2 From the **Task** list, select **Delete**.

3 Click **Yes** in the Delete confirmation message box.

Note: Deleting a range set will remove all linked appearances of this range set from the Catalog workspace. △

Opening a Range Set

To view or edit the properties of a range set, perform these steps:

1 From the **Range Set** folder, select the range set.

2 From the **Task** list, select **Open**.

3 When you finish viewing or making changes, click **OK** if you want to save any changes; otherwise, click **Cancel**.

Renaming a Range Set

To rename a range set, perform these steps:

1 From the **Range Set** folder, select the range set.

2 From the **Task** list, select **Rename**.

3 Enter the new name and click **OK**.

Viewing and Editing a Range Set

To view or edit the properties of a range set, perform these steps:

1 From the **Range Set** folder, select the range set.

2 From the **Task** list, select **Open**.

3 In the Range Set dialog box, you can add, delete, or change these fields: the **Upper Limit**, **Lower Limit**, the **Inclusive** options for the limits, and the **Score**. To do so, perform these steps:

   a Select the **Range Set** tab
   b To add a range, click **Add**. Enter the upper and lower limits of the range. Click **Inclusive** for each limit if you want those limits to include the values you entered. Optionally, enter the score. Select a range description to specify the color and label. Then click **OK**.
   c To delete a range, click **Delete**. Click **Yes** in the Delete confirmation message box.
   d To edit a range, click **Edit**. You can change the upper and lower limits, whether a limit is inclusive, the score, and the range description. Then click **OK**.

4 When you finish viewing or making changes, click **OK** to save any changes; otherwise, click **Cancel**.

For more information about range sets and tasks related to range sets, click **Help** in the Range Set dialog box.

Range Descriptions

The **Range Description** folder contains range descriptions that you have added as well as predefined range descriptions that are supplied with the software. The predefined range descriptions are: Out-of-Bounds-High, Out-of-Bounds-Low, and Missing. These fixed or predefined range descriptions cannot be deleted, renamed, or modified.

To see the current list of range descriptions, double-click the folder.

You can use any range description in the catalog to assist in building your range sets, to reduce the effort required to build other range sets, and to minimize inconsistencies among range sets. Colors can thus have the same meaning across multiple contracts as defined in the reports specifications.

Note: From the Report Specification window, you can change the color and label of a range description that you want to use for reports. △

Creating a Range Description

To create a range description, perform these steps:

1 From the **Task** list, select **Create a Range Description**. The Range Description dialog box opens.

2 Enter a range description name.

   Note: The name must not be blank. It is a good idea to use a name that is unique across all range descriptions in the catalog. △

3 Use the drop-down list to select the colors for the text reports.

   Note: The list displays the names of the colors in English. △

4 If you want to customize the label that will be used for this range in the reports, click **Customize the report label**. Then enter the report label in the **Label in reports** field.

5 Click **OK** to save this range description.

Deleting a Range Description

You cannot delete any of the three predefined range descriptions. To delete a range description, perform these steps:

1 From the **Range Description** folder, select the customer.

2 From the **Task** list, select **Delete**.

3 Click **Yes** in the Delete confirmation message box.

Note: A range description cannot be deleted if it has existing links in the catalog or the contract. However, it can be replaced. If the range description has existing links, a **Replace Range Description** message will display stating this fact. If you choose to continue with the deletion, the Replace Range Description dialog box will display a list of range descriptions that you can use to replace the one you want to delete. Select a range description from the list and click **OK**. △

Opening a Range Description

To open the properties of a range description, perform these steps:

1 From the **Range Description** folder, select the range description.

2 From the **Task** list, select **Open**.

Renaming a Range Description

You cannot rename any of the three predefined range descriptions.
To rename a range description, perform these steps:

1 From the **Range Description** folder, select the range description.

2 From the **Task** list, select **Rename**.

3 Enter the new name and click **OK**.

Note: Changing the name of a range description in the catalog will change the name of customized ranges in the report program. However, it will not change the color or the label customizations. △

Viewing and Editing a Range Description

To view or edit the properties of a range description, perform these steps:

1 From the **Range Description** folder, select the range description.

2 From the **Task** list, select **Open**.

3 On the Range Description dialog box, you can change these fields: **Range Description Name**, the **Color for text reports**, the **Customize the report label**, and the **Label in reports**.

4 When you finish viewing or making changes, click **OK** if you want to save any changes; otherwise, click **Cancel**.

For more information about range descriptions and tasks related to range descriptions, click **Help** in the Range Description dialog box.

CHAPTER

8

Entering a Contract

Introduction to Defining a Contract

The Contract workspace contains information about contracts.

The Contract workspace contains one fixed folder: the **Contracts** folder. In the **Contracts** folder, you can view, add, change and delete existing definitions for contracts. If you need other contracts, you can create them here.

If you want to organize the information in the fixed folder, you can create subfolders. Under the subfolders, you can create more folders, and so on. All the subfolders in the Contract workspace are displayed in ascending alphanumeric order.

- *To create a user-defined folder*, select the folder that you want to be the parent folder and then, from the menu bar, select **File ▶ New ▶ Folder**. Enter the name and click **OK**.
- *To rename a user-defined folder*, select the folder and then, from the **Task** list, select **Rename**, enter the new name, and click **OK**.
- *To delete a user-defined folder*, select the folder and then from the **Task** list, select **Delete**. In the Delete confirmation message box, click **Yes** to confirm your request to delete.

Note: For more information about folders, from the menu bar select **Help ▶ IT Service Level Management Help**. △

Contracts

The **Contracts** folder contains information about one or more contracts. To see the current list of contracts, double-click the folder. The contracts in that folder are displayed in the main section of the window. (To see more of that section, drag the right border of that frame to the right).

Tip: When you are defining a contract, the contract receives a *copy* of the items that are in the catalog. Thus, if you later edit the items in the contract, your changes in the contract are not transferred to the catalog. And if you later edit the items in the catalog, your changes in the catalog are not transferred to the contract.

The following topics explain how you can set up your contracts:

The following topics explain how to work with your contracts after they have been set up:

Creating a Contract

To create a contract, perform these steps:

1. From the **Task** list, select **Create a Contract**. The Contract window opens.
2. On the **General** tab, enter the contract name and, optionally, the version number of the contract. (In the contract's reports, the value in the **Contract Name** field is the main identifier of the node.)

   Note: The name must not be blank. It is a good idea to use a name that is unique across all nodes in the Contract workspace. △

   In the **Path to Contract Database** field, enter the path where you want to store the contract database (for example in Windows, you might enter *c:\myslm\cdb1*). The site must be accessible from the active server—that is, the server where you are running SAS IT Service Level Management. Also, the location must be unique (that is, each contract must have its own contract database). The system will prevent you from creating a contract database that points to an existing database.

   Note: When you are specifying the pathnames or the directory names for the contract database, do not use ampersands (&). (If you include ampersands in the pathname or directory name, the resulting placement of the index.html file cannot be predicted.) △

   If you know the contract's start and end dates, you can enter them now. (These dates do not control the ETL process. They are used for documentation only.)

   The default status is **Inactive**. (When the status is **Active**, you cannot save the contract unless you have entered the path to the contract database. If you do not wish to specify this path, you can leave the status set to **Inactive**.)

If the status of an active contract (that has already been collecting data) is changed to **Inactive**, all data collection will cease and the contract database will not be updated. However you can still make changes to the contract hierarchy and run Synchronize.

You can also select the servicing contact for this contract (from the list of contacts for whom you selected the Servicing Contact option). The telephone number and the e-mail address of the selected contact will be displayed.

If you want this servicing contact name to be visible on the reports that are generated for the customer, make sure that **The contact should be visible on Customer Reports** option is checked.

If an objective of this contract misses its target, you can notify the servicing contact who is associated with this contract. To do so, make sure that the **E-mail this contact for any missed target definitions** option is checked. You must also enter the contact's e-mail address in the appropriate Contact dialog box. Then, if the ETL process detects that a target has been missed, it will notify the servicing contact by e-mail.

(If you have checked the **E-mail this contact for any missed target definitions** option as the default on the Contact dialog box, the option on this tab will already have a check mark.)

Note: The e-mail notification feature is available for the objectives of every node of the hierarchy. △

3 If you want to base the calculations for this node on a subset of the data, select the **Date Filter** tab. On the tab, select a date filter pattern.

- If you select **All Inclusive**, all the data from the child nodes is used. This is the default pattern.
- If you select **Continuous**, all the data from the effective date and time to the expiration date and time is used. Specify these dates and times.
- If you select **Daily**, all the data from the effective time to the expiration time is used on each date from the effective date to the expiration date. Specify these times and dates.
- If you select **Weekly**, all the data from the effective time to the expiration time on the selected day or days of the week is used for each week from the effective date to the expiration date. Specify these times, days, and dates.
- If you select **Monthly**, all the data from the effective time to the expiration time on the selected day of the month is used for each month from the effective date to the expiration date. Specify these times, this day, and these dates.

Note: You can also specify details for recurrence of the daily, weekly, and monthly date filtering patterns. For more information about the **Date Filter** tab, click **Help** in the Contract window. △

The date filter for a contract is independent of the date filters for other nodes in the contract hierarchy. For examples of how you can use date filters more effectively, see "Examples of Date Filter Usage" on page 313.

4 On the **Parties** tab, you can optionally add the customer and provider that are appropriate for this contract.

To select the customer that you wish to copy, in the **Customer** section click **Catalog**, select the appropriate customer, and click **Add**. The contract now has a copy of the customer and displays its properties in the **Customer** section.

Similarly, to select the provider you wish to copy, in the **Contract Provider** section click **Catalog**, select the appropriate provider, and click **Add**. The contract

now has a copy of the provider and displays its properties in the **Contract Provider** section.

5 On the **Hierarchy** tab, you can add catalog items to create the child nodes of this contract.

In the **Contract Hierarchy** column, select the name of this contract, and right-click to select **Catalog**. A list of already-created service level agreements, services, and components is displayed. Select an SLA, a service, or a component that you want to be part of this contract and then click **Add**. Repeat until you have created a hierarchy that represents your "real world" contract with your customers.

Note: The **Catalog** provides a flexible way to define your contract hierarchy. After you have added a node from the catalog to your contract, you can modify the contract hierarchy that you are defining. For example, you can add an SLA to your contract. Then you can open that SLA and add another component to it. The catalog entry for the SLA will not be changed, but your contract will contain the SLA with the newly added component. △

The data that is displayed in the other columns depends on the selection in the **Filter by Objective Type** drop-down list and the objective values of the objective settings for all items in the hierarchy.

6 On the **Objectives** tab, click the **Activate for Processing** option for the objectives that you want to work with: **Availability**, **Response Time**, **Throughput**, **Custom Objective**, or **Composite**. The corresponding tab will be available in the lower portion of the window.

- If you want to use the availability objective, you can define it on the **Availability** tab.

  To specify the indicator, perform the following actions:

  - If you want to use a supplied calculation method, in the **Input metrics from child nodes** field select **Indicator** or **Score** and in the **Calculation Method** field select the method that you want to use.
  - If you want to use a custom calculation method, in the **Input metrics from child nodes** field select **Custom Code** and then click **Code**. Provide the SAS code that you want to use for this indicator and then click **OK**.

  To specify the range set, perform the following actions:

  - In the **Range Set Definition** section, click **Catalog** to open the Mini-Catalog dialog box. The range sets that exist in the catalog are listed. Select a range set for this objective and then click **Add**. A copy of the range set is created with for this objective, and the name and a select group of the properties of this range set is displayed.
  - You can change the limits and inclusiveness of a range, its score, and its description. To do so, click **Edit** in the Range Definition dialog box, select the range that you want to change, and click **Edit**. Make your changes in the Edit Range dialog box and click **OK**. The changes are validated to ensure there are no gaps or overlaps among the ranges in the range set. Changes are saved when you click **OK** in the calling dialog.

    Note: For information about how to change the attributes of a range set or a range description, see "Range Sets" on page 76 and "Range Descriptions" on page 78. △

  To specify the target, perform the following actions:

  - If you want to use an indicator-based target, select **Indicator** (for this objective) in the **Comparison Value** field; select the comparison

operator that you want to use in the **Operator** field; and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** and **Target Statistic** drop-down lists, select the period and the statistic that you want to use.

- If you want to use a score-based target, select **Score** (for this objective) in the **Comparison Value** field, select the comparison operator that you want to use in the **Operator** field, and enter the value that you want to use in the **Target Value** field. You can optionally specify a label for the target value. Also, from the **Compliance Period** drop-down list, select the period that you want to use.

Note: If you request a report for this contract, then the report's data about this objective is based on the target, period, and statistic that you select here. △

- To use the response time objective, the throughput objective, or the custom objective (or any combination of these objectives), make the appropriate selections on the corresponding **Response Time** tab, **Throughput** tab, or **Custom Objective** tab.

  If you want to apply a weighting factor to the numeric calculations of the mean statistic, follow these steps:

  a. Specify the **Weight By** variable(s) for each of the components that make up this objective for the contract:

     From the **Hierarchy** tab of the contract, double-click to open a component that specifies this objective of the contract. On that component's **Data** tab, enter the variable to use as the weighting factor for that component.

     Repeat this process until you have specified a **Weight By** variable for all the components that make up this objective of the contract.

  b. On the lower part of the **Objectives** tab of the contract, click to select the corresponding objective tab: **Response Time** tab, **Throughput** tab, or **Custom Objective** tab. Then select Weighted Mean from the drop-down list for **Calculation Method**.

  c. If you want to use the weighted mean as a target statistic, then select Weighted Mean from the drop-down list for **Target Statistic**.

  Note: Mean and Weighted Mean are mutually exclusive. If you choose a Weighted Mean calculation method, then you cannot choose Mean as a target statistic. Instead, Weighted Mean will be in the drop-down list. Likewise, if you choose a Mean calculation method, then you cannot choose Weighted Mean as a target statistic, because Weighted Mean will not be in the drop-down list.

  The Weighted Mean calculation is available only for the response time, throughput, and custom objectives. △

- Similarly, if you want to use the composite objective, make the appropriate selections on the **Composite** tab.

  Note: Composite objectives are useful when you need an overall measure of a component, service, SLA, or contract and when that measure takes into account various objectives. For example, a print server's performance might depend on both excellent availability and good throughput. △

7. Each contract can have outages. You might already have information about outages that apply to this contract, especially if the outages are scheduled. If so, select the **Outages** tab and then click **New**. The Outage Properties dialog box opens.

Enter the outage name. You can optionally enter the reason in the **Reason for the Outage** field. In the **Outage Type** field, click the arrow and select the type of outage. You can select one of the following types of outages:

- Scheduled (the default)
- Excused
- Unscheduled

The type of outage that you select will be displayed in the reports.

Note: Scheduled outages and excused outages do not affect the response time, throughput, custom objective, or composite indicators; these types of outages are calculated as missing. However, scheduled outages and excused outages do affect the availability indicator because the time covered by such outages reduces the total amount of time that the service provider is responsible for. Therefore, if there are unavailable periods of time, the outages become a larger percentage of the total time that the service provider is responsible for.
Unscheduled outages do affect the indicator. △

In the contract hierarchy, select the highest node to which this outage applies (that is, the node that is closest to the top of the hierarchy). For now, leave the **Is this outage in effect?** field set to **No**.

On the **Schedules** tab, select the outage pattern and enter any additional dates and times required by that pattern.

- If you select **One Time Only**, all the data from the start date and time to the end date and time is included in the outage. Specify these dates and times.
- If you select **Daily**, all the data from the outage start time to the outage end time is used on each date from the recurring start date to the recurring end date. Specify these times and dates.
- If you select **Weekly**, all the data from the outage start time to the outage end time on the selected days of the week is used for each week from the recurring start date to the recurring end date. Specify these times, days, and dates.
- If you select **Monthly**, all the data from the outage start time to the outage end time on the selected day of the month is used for each month from the recurring start date to the recurring end date. Specify these times, this day, and these dates.

If you already know of any exceptions to that pattern, select the **Scheduled Exceptions** tab, click **New**, enter the date and reason for the exception, and click **OK**.

If you want to add notes about this outage, use the **Notes** tab.

When you finish entering information about this outage, on the outage's **General** tab, select **Yes** in the **Is this outage in effect?** drop-down list. Then, click **OK**.

CAUTION:
If you modify an outage, then any data for time periods covered by the outage will be corrected the next time that the synchronize task is run, in order to accurately reflect this updated definition of the outage. △

For more information about these outage tabs, click **Help** in the **Outage Properties** dialog box. For more information about outages, see "What is an Outage?" on page 231.

8 On the **Contractual Data** tab, enter any information that you want to record about these topics. Except for the **Start of Week** field in the Contract Database Info section of the tab, the information that you enter is for your documentation purposes only; it does not affect the processing of the contract.

To specify the starting day of the week for your contract, use the drop-down menu to select the **Start of Week**.

9 If you want to add notes about the contract node, use the **Notes** tab. You can also include attachments that can be displayed on the reports. (The attachment section can contain an icon, a name, and a Web address or an address of a local file.)

To add an attachment, perform the following steps:

a From the **Notes** tab, click **New** to open the Attachment Properties dialog box.
b You can, optionally, enter the name of the attachment.
c Enter the Web address of your attachment in the form of a URL (that is, *http://*) or in the form of a file address (that is, *file://*). Spaces are not allowed in a URL. (You can click **Browse** to search your directories for the file whose address you want to record.)

Note: If the address that you specify for the attachment does not include the http:// or file:// prefix, you will not be able to view the attachment in your reports because your browser will be unable to determine its location. △

d To display the attachment with its associated contract on contract reports, select **Include this attachment**.

Note: If this option is checked, then the **Annotation Start Date** is required. △

e You can specify the start and end dates that you want the attachment to display on contract reports. Click the calendar icon to open a calendar. Then select the appropriate day to enter the date in **Annotation Start Date** and then again in **Annotation End Date**. Make sure that the end date is greater than the start date.
f Click **OK** to save the attachment.

10 To save the properties of the contract, select **File ▶ Save**.

Note: Whenever you save the contract, you are prompted to create the contract database if it does not already exist. (On Windows and UNIX, the database will be created in the location that you specified as the path to your contract database.)

However, on z/OS, you must pre-allocate the contract database before you save it. For information about how to do this, see "Working with Servers on z/OS" on page 40. △

11 To return to the Contract workspace, select **File ▶ Close**.

If you are not finished entering information about the contract, you can return later to work on the contract. For more information, see "Opening a Contract" on page 88.

When you finish entering information about the contract, you can activate the contract and begin to process data and produce reports. For more information, see "Activating a Contract" on page 87.

Activating a Contract

To activate a contract, perform these steps:

Note: Prior to activating the contract, you must have already defined the path to the contract database. △

1 In the Contract workspace, select the contract.
2 From the **Task** list, select **Open**.
3 On the **General** tab, in the **Status** field, click the arrow and select **Active**.

4 From the menu bar, select **File ▶ Save**.

This selection saves the properties of the contract. (The only new property is the active status.)

Also, because the **Status** is **Active** and there is no contract database yet, this action creates the contract database.

Note: If your server is on z/OS, the creation and initialization of the contract database must take place before you save the contract. For information about this, see "Working with Servers on z/OS" on page 40. △

Deleting a Contract

To delete a contract, perform these steps:

1 In the Contract workspace, select the contract.

2 From the **Task** list, select **Delete**.

3 Click **Yes** in the Delete confirmation message box.

Opening a Contract

To view or edit the properties of a contract, perform these steps:

1 From the Contract workspace, select the contract.

2 From the **Task** list, select **Open**.

3 When you finish viewing or making changes, if you want to save any changes, select **File ▶ Save**. Then, select **File ▶ Close**.

Renaming a Contract

To rename a contract, perform these steps:

1 From the Contract workspace, select the contract.

2 From the **Task** list, select **Rename**.

3 Enter the new name and click **OK**.

Synchronizing Contract Data

The synchronize task modifies the data in the contract database in accordance with any changes that you have made to the contract in the GUI.

From the menu bar, select **Tools ▶ Synchronize**. Then, click **Yes** to confirm that you want to run this task.

If the contract is active and there are no tables in the contract database yet, this action creates the tables in the contract database.

When the synchronize task finishes, the Contract Results dialog box opens with the **Synchronize** tab selected, and displays the log from the synchronize task. At the bottom of the log, if the synchronize task was successful, this message is displayed:

```
PASSED: Contract synchronization completed with return code 0
```

Note: If you do not successfully complete the synchronize task, review the synchronize log in the Contract Results dialog box and fix any errors noted in the log. △

Then, click **Close**.

Validating Contract Data

The validate task checks that the contract data is valid—that is, that the ETL run can complete successfully. The validate task checks that the properties of the contract are appropriately set. It also checks that the ETL job can access and extract the necessary data from the tables within the data source PDB (local or remote) that is being processed.

To run the validate task, from the menu bar, select **Tools ▶ Validate**. Then, click **Yes** to confirm that you want to run this task.

When the validate task finishes, the Contract Results dialog box opens with the **Validate** tab selected, and displays the log from the validate task. At the bottom of the log, if the validate task was successful, this message is displayed:

```
PASSED: Contract validation completed with return code 0
```

Note: If you do not successfully complete the validate task, review the validation log in the Contract Results dialog box and fix any errors noted in the log. △

Then, click **Close**.

Creating the ETL Job Code

ETL is the process of loading new data into the contract database in accordance with the current specification of the contract. Specifically, ETL takes information about the contract from the environment database and then performs the following steps:

1. extracts the data that is appropriate for this contract from the data source PDBs
2. transforms the data based on outages, targets, weights, and date filters. (Transformation of data is done in intervals of five minutes.)
3. loads the data into the contract database in preparation for viewing that data in the reports

ETL Job Code

After you create a contract and activate it, you can create the ETL Job Code from the GUI. The code that you produce runs the ETL process for a contract.

To create the ETL job code, perform these steps:

1. From the menu bar of the Contract window, select **Tools ▶ Create ETL Job Code**.
2. The Create ETL Job Code dialog box opens. Type the location of the file where you want to store the SAS code for processing data.

   Note: On Windows and UNIX, it is a good idea to use **.sas** as the file type. △
3. Click **OK**.

Customizing the ETL Job Code

You might need to customize this code by specifying (in the Server Connections Properties dialog box) a user ID and a password for each of the remote data source system connections that you use, if you have not already specified them in the GUI. It can be customized directly using SAS software or a text editor and scheduled by using your in-house scheduling software.

Tip: It is good practice to manually run a synchronize task and a validate task after initially activating a contract. The synchronize task will update the existing data in the contract database in accordance with the current specifications of the contract. It will create the necessary tables in the contract database. The validate task makes sure that all fields that are required have been properly entered.

By running these tasks, you can determine if the contract definition is ready for loading data into the contract database. For example, you can discover whether sources of data have been defined for all activated component objectives. You can correct those problems prior to running ETL.

Note: Because the Synchronize and ETL tasks can be lengthy (depending on the complexity and size of your hierarchy, your date filters, your objective definitions, and your outages), running these tasks in a batch environment is recommended. However, if you run from SAS IT Service Level Management GUI, a progress indicator will display. △

For information about running the ETL job code from the GUI and in batch mode, see "Running Job Code in Batch Mode" on page 298.

Reporting on a Contract

After you have run ETL successfully, you can generate reports. The reports will reflect the state of the contract at the time the reports are run. Information about the types of reports that you can generate can be found in the Reporting section of this user's guide. To review this information, begin with "Overview of Reports for Service Level Management" on page 93.

In order to generate reports, you must first define report specifications. Report specifications can be generated for contract providers and for customers by using a **`Report Specification`** Wizard. The Wizard will guide you through the process of defining the content of the reports you want to generate. It will also establish the settings that govern the generation of the reports you have defined, such as range settings, output locations, and the like. The Wizard saves the specifications in the appropriate Contract Provider or Customer folder in the Reports workspace.

Note: To keep the contents of each set of reports secure, specify different locations for the report output from each report group. SAS IT Service Level Management will use the security facilities of your operating system and Web server to provide appropriate access. △

After the reports specifications have been defined, they can be run from the GUI or else scheduled to run in batch mode.

To review information about how to specify the reports that you want to generate, see "Generating Service Level Reports" on page 132.

Note: For information about contracts and the tasks that are related to contracts, click **`Help`** in the Contract window.

For more information about nodes, see "Introduction to Understanding Contract Structure" on page 9. For more information about objectives, see "Introduction to Understanding Contract Objectives" on page 13. △

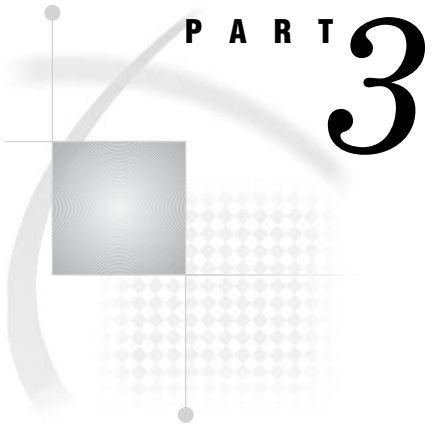

PART 3

Reporting

CHAPTER

9

Overview of SAS IT Service Level Management Reports

Overview of Reports for Service Level Management

SAS IT Service Level Management provides a flexible reporting capability. Reports can be produced by using the client GUI or the SAS Enterprise Guide application. SAS IT Service Level Management provides the following types of reports:

- *Baseline reports* are a collection of small and large graphs that are output to a Web-based gallery. A contract provider can generate these reports, which display data from components with activated objectives. These reports provide a basis for establishing realistic targets and range sets for a contract.

  For more information about these reports, see "About Baseline Reports" on page 97.
- *Service Level reports* are reports that are output to a Web browser. A contract provider can generate these reports, which display actual and targeted values for daily, weekly, monthly, and yearly time periods. These reports provide a way for contract providers and customers to assess whether contracts and their nodes are meeting or missing the targets set for their objectives. The reports that are generated for customers can be placed in separate, secure output locations.

  For more information about these reports, see "About Service Level Reports" on page 110.
- SAS Enterprise Guide reports are reports that are generated by using the Windows-based, point-and-click SAS Enterprise Guide application and that are output to a Web browser or to an Enterprise Guide window. With SAS Enterprise Guide, you can set up projects that run multiple tasks on the same group of data files and create interactive or ad hoc reports.

  For more information about these reports, see "Introduction to Working with SAS Enterprise Guide to Create SAS IT Service Level Management Reports" on page 187.

In order to facilitate writing custom reports with SAS Enterprise Guide, SAS IT Service Level Management provides views of the contract database. The views are located in the ADMIN level of the contract database. They are recreated after the synchronization task of the ETL process is successfully run. Using the views simplifies access to the information that you want to report on. For more information about this feature, see "Accessing Views of the Contract Database" on page 237.

SAS IT Service Level Management also provides tables that can be used to navigate through the contract hierarchy and its associated objects. The metadata in these tables

will facilitate writing reports using SAS Enterprise Guide, SQL, and other report writing tools. The tables are recreated in the ADMIN library of the contract database after each successful execution of the Synchronize task of the ETL process. For information about these tables, see "Accessing the Metadata of the Contract Database" on page 239.

Working with Reports

The following steps describe the typical procedure for generating and viewing reports for the contract (or contracts) that you are working with.

1 Set up and run the ETL process on the contract. (Reports will not be generated unless the ETL process has run successfully.)

   For information about this topic, see "The ETL Process" on page 287.

2 In order to determine realistic targets and ranges for the specified objectives of the contract, run the Baseline reports for that contract. (Baseline reports display the actual data for the components that are specified in the contract hierarchy.)

   For information about this topic, see "About Baseline Reports" on page 97.

3 Use the information from the Baseline reports to determine the targets and range sets that you want to use for the contract and its nodes.

   - *For the contract*: To specify these values for the targets and range sets of the contract, enter them in the Contract window of the Contracts workspace.
   - *For the nodes of a contract*: To specify these values for the targets and range sets of a node of the contract, enter them in the appropriate Service Level Agreement (SLA), Service, and Component dialog boxes of the Catalog workspace.

     If you modify values in the Catalog, the changed values are not reflected in *existing* contracts or reports unless you also make the modifications in the Contracts workspace.

4 Define or adjust the global settings for your reports. These settings enable you to specify that the following options will be used, by default, for all report specifications:

   - the default output location for your reports
   - the location of the cascading style sheet that you want to use in place of the standard SAS style sheet
   - the number of decimal places that will be used in the reports

   Note: Specifying these global report settings is not required. You can specify any, all, or none of these options. In addition, you can override the report output location and the location of the style sheet for any individual report specification. △

5 Create report specifications for each customer or provider. (A report specification can be specified for one or more contracts. Conversely, a contract can be included in one or more report specifications.)

   For more information about how to create a report specification, see "Creating or Modifying the Report Specification" on page 123.

   When you have created your report specification, you can generate your reports interactively or you can create the report job code that can be run in batch mode.

   - If you want to run the reports interactively, you can generate them immediately.

- To do so, from the **Tools** menu of the Report Specification window, click **Generate Report**.

When this report generation process is finished, the View Report Results dialog box opens.

- If you want to run the reports later, you can create the report job code.
  - To do so, from the **Tools** menu of the Report Specification window, click **Create Report Job Code**.
  - In the Create Report Job Code dialog box, enter the location of the report job code. Click **OK** to generate the report job code.

You can run the report job code by submitting it in a SAS session. You can also run this code by wrapping it in a .bat file (for Windows), a .sh file (for UNIX), or a job (for z/OS), depending on your operating system. For information about running batch jobs, see "Batch Jobs" on page 297.

When the reports have been generated, you can view them by navigating to the report output location that you designated in the report specification.

For more information about these topics, see "Working with Service Level Reports" on page 116.

6 For ad hoc reporting, Enterprise Guide reports can be run any time after the ETL process has run successfully.

CHAPTER
10
Using Baseline Reports

About Baseline Reports

This chapter consists of the following topics:

- "What Are Baseline Reports?" on page 97
- "Working with the Baseline Gallery" on page 98
  - "Creating the Job Code for Baseline Reports" on page 98
  - "Generating Baseline Reports" on page 100
  - "Viewing Baseline Reports" on page 102

What Are Baseline Reports?

Baseline reports are a collection of small and large graphs that are output as a Web-based "gallery." These reports, which are generated for the contract provider, display data from components with activated objectives over a period of time. The data can then be analyzed and used to establish target thresholds that are both attainable by the contract provider and acceptable to the customer. The data can also be used to define range sets for contracts and for the nodes in a contract.

The baseline graphs that are produced display data from the daily, weekly, and monthly summaries of the contract database. They are generated for the minimum, maximum, sum, and the mean of the availability, response time, throughput and custom objectives. (Data about composite objectives is not included in the Baseline reports.)

Baseline reports are typically generated for a specified contract when the SAS IT Service Level Management software is first implemented at a site. However, they can also be generated again to reflect any of these circumstances:

- when you add a new contract
- when you add a new component
- when you add or change resources
- when you activate an objective for an existing component
- when you process additional data

Working with the Baseline Gallery

In order to generate the reports that display in the Baseline gallery, you must create and then run the Baseline report job code. You can then access the report output in the Baseline gallery that is produced.

The following display shows the steps that are needed to create and view the baseline gallery:

Display 10.1 Working with Baseline Reports

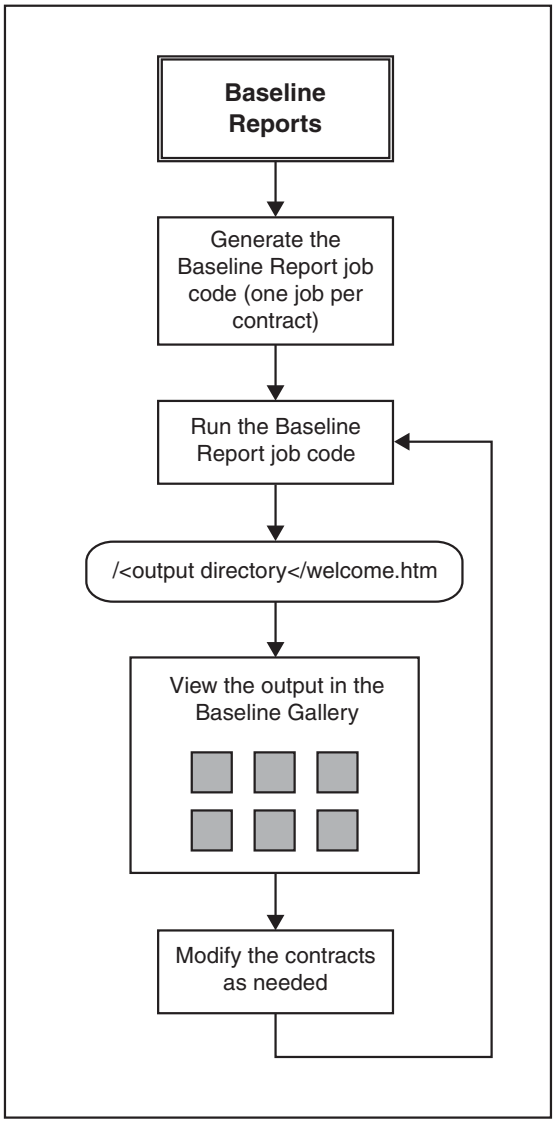

Creating the Job Code for Baseline Reports

This topic discusses how to create the job code that will generate a Web-based gallery of baseline reports for a specified contract.

To create the job code file, perform these steps:

1 From the Contract workspace, select the contract name for which you want to run the baseline reports.

2 In the toolbar of the Contract window, select **Tools ▶ Create Baseline Gallery Job Code**.

Display 10.2 Create Baseline Gallery Job Code menu

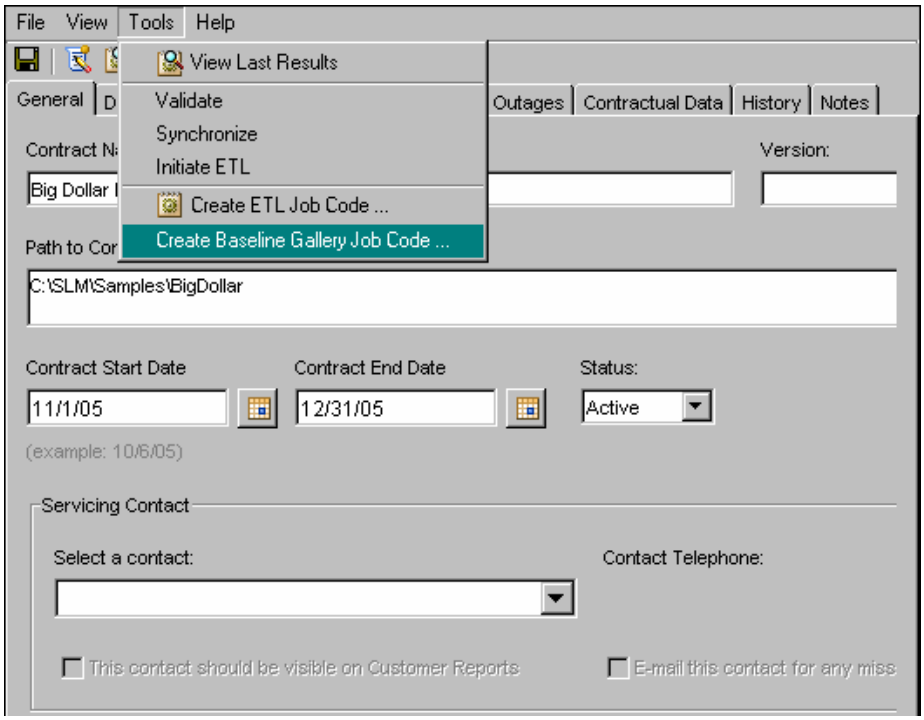

3 In the **`Create Baseline Gallery Job Code`** dialog box that opens, supply the output location for the baseline gallery and the path and filename of the job code that will be created.

Display 10.3 Create the Baseline Gallery Job Code dialog box

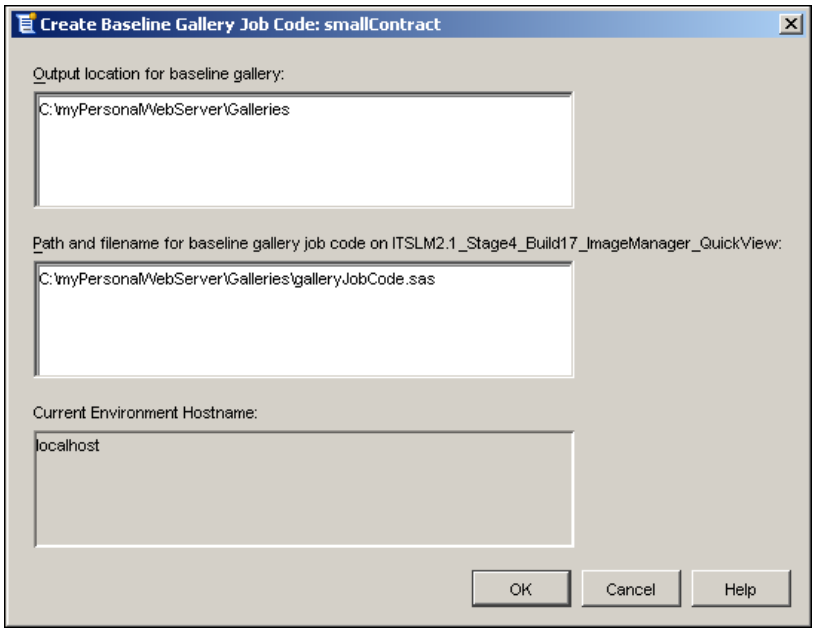

To supply the required information, perform the following steps:

a In the **`Output location for baseline gallery`**, enter the full path to the location where you want the baseline reports HTML file (**`welcome.htm`**) to be located. Do not specify a relative path. This field can contain a maximum of 200 characters.
b In the **`Path and filename for baseline gallery job code on <environment>`**, enter the full path and filename for the job code file. This field can contain a maximum of 200 characters.
c Click **`OK`** to create the job code.

Note: **`Current Environment Hostname`** displays the host name of the SAS IT Service Level Management software environment. This field cannot be modified. △

Generating Baseline Reports

This topic discusses how to run the job code that generates a Web-based gallery of baseline reports.

Before you run the job code to generate the baseline gallery for a contract, you must set up a contract that includes the components that you want to work with, activate at least one objective, and run the ETL process on the contract data.

After the baseline job code for a contract has been successfully created, you can run the job code in the following ways:

- *In batch mode*: You can generate baseline reports by scheduling a batch job or by executing the job code from a command line.

  To schedule a batch job to run, you must "wrap" the job code in the batch process that is appropriate for your operating environment. For more information about setting up and running job code in batch mode, see "Batch Jobs" on page 297.

  To execute job code from a command line, enter the name of the job code file that you created. For example, from the command line in the Windows operating environment, enter this command:

```
slmjob.bat baseline
```

where

- slmjob.bat is the name of the batch process that executes on Windows and
- baseline.sas is the name of the job code that you created

Note: To locate a copy of slmjob.bat, see "Batch Jobs" on page 297. △

- *In a SAS session*: You can generate baseline reports by running the job code in a SAS session.

  To run your job code in a SAS session, copy the job code into the Editor window of a SAS session and click **Submit**. For example, using the information that you entered in the Display 10.3 on page 100, you can run **galleryJobCode.sas**. This code generates the output in **C:\myPersonalWebServer\Galleries**.

Note: Check the SAS log file to make sure that your job code ran successfully. If an error message indicates that a synchronization action is pending, then run the Synchronize task before you rerun the baseline gallery job code. For more about the Synchronize task, see "Synchronizing Contract Data" on page 88. △

Viewing Baseline Reports

This topic discusses how to view the baseline reports that are generated.

After the baseline job code for a contract has been created and run successfully, you can view the baseline reports that were generated.

Note: If any synchronization action is pending during the execution of the baseline gallery job code, then the baseline gallery job deletes all existing reports in the gallery and writes a message about the pending synchronization action to the log file. You must run Synchronize before you rerun the baseline gallery job code. (From the Tools menu of the Contract window, click **Synchronize**.) △

To access the baseline reports, perform the following steps:

1. Navigate to the **welcome.htm** file in the **Output location for baseline gallery** that you specified when you created the baseline job code.
2. In the left panel, click **Baseline Reports** to open the Baseline Reports Welcome page that you want to view.
3. If you want to customize the appearance of the Baseline gallery page, click the **User Preferences** option that is located on the banner of the Baseline gallery page. In the User Preferences dialog box, you can enter the following values:
   - The number of small graphs that display across the screen.
   - The first day of the week. The default is Sunday.

     Note: The start-of- week setting is effective only if the gallery uses a calendar in folders that contain graphs and if those graphs use Date as a BY variable. △
4. Click the **Daily**, **Weekly**, or **Monthly** folder to open the reports for that time period.

Daily graphs	contain plots that show the mean, minimum, maximum, and sum indicator values as well as the 10th, 50th and 90th percentiles of the daily values for all activated objectives of a component.
Weekly graphs	contain plots that show the mean, minimum, maximum, and sum indicator values as well as the 10th, 50th and 90th percentiles of the weekly values for all activated objectives of a component.
Monthly graphs	contain plots that show the mean, minimum, maximum, and sum indicator values as well as the 10th, 50th and 90th percentiles of the monthly values for all activated objectives of a component.

Display 10.4 Welcome Page of the Baseline Gallery

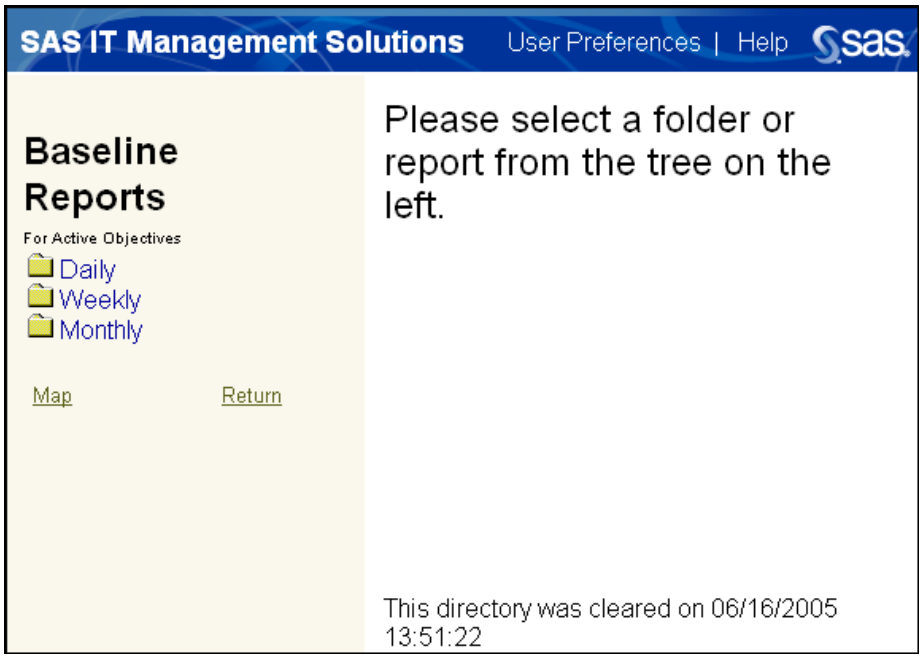

5 After you select a time period, select the statistic that you want to view. The Baseline gallery displays the thumbnail graphs for the time period and statistic that you selected.

Display 10.5 The Baseline Gallery Welcome page

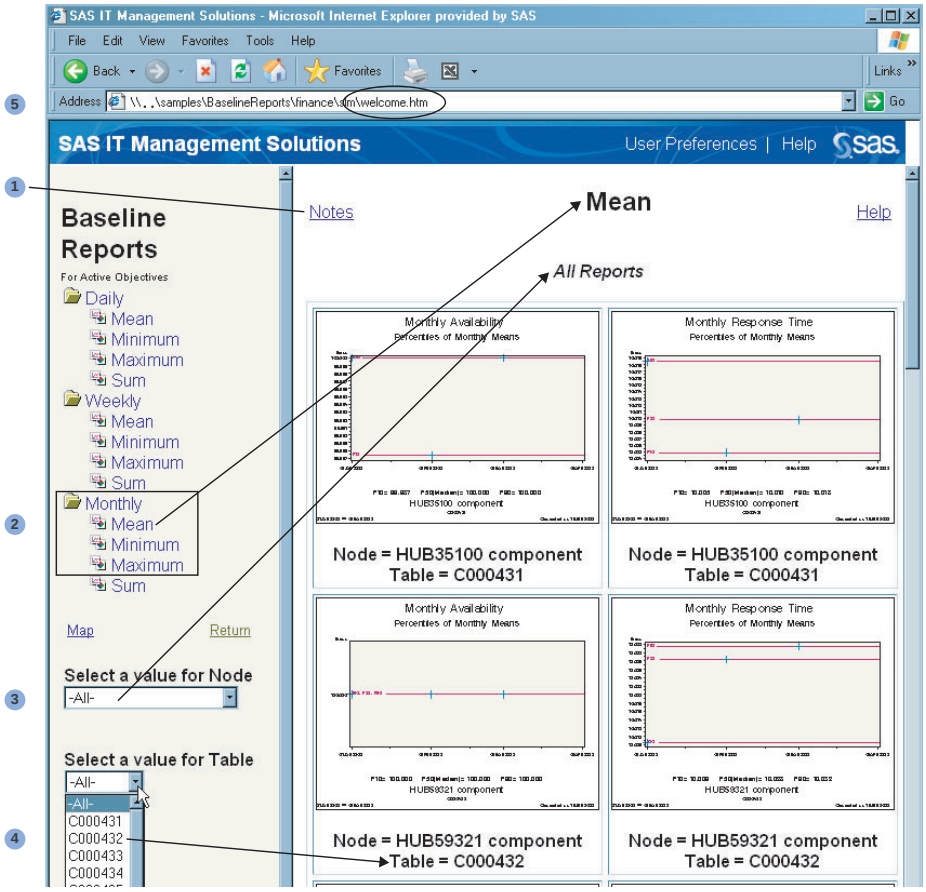

Table 10.1 Legend for Baseline Gallery Welcome page

Number	Description
1	The **Notes** link contains instructions for adding report notes to your baseline gallery.
2	The Daily, Weekly, and Monthly folders contain subfolders for eachof the report groups that are available for that time period. Each report group contains subfolders for the Mean, Minimum, Maximum, and Sum graphs.
3	**Select a value for Node** provides a drop-down list of unique node names. By default, the graphs for all nodes are displayed.
4	**Select a value for Table** provides a drop-down list of unique table names for each component of the contract. These table names are generated by the system when each component is created. The table name (such as C000432) is displayed below each graph. To distinguish components that have the same name, refer to them by their table names.
5	**Welcome.htm** is the default filename of the starting page of each baseline gallery that you generate. Therefore, if you generate a baseline gallery for more than one contract, be sure to organize the output location of the welcome.htm file so that you can easily associate the welcome.htm file with the corresponding contract.

6 You can use the drop-down menu that is labeled **`Select a value for Node`** in the navigation pane of the baseline gallery to display the graphs for a specific component.

Display 10.6 Baseline Graphs for HUB59321 Component

7 You can also use the drop-down menu that is labeled **`Select a value for Table`** in the navigation pane of the baseline gallery to display the graphs for a specific table. A component is uniquely identified by the table name. You can use the table name to distinguish between components with duplicate names.

Display 10.7 Baseline Graphs for Node HUB59321 Component

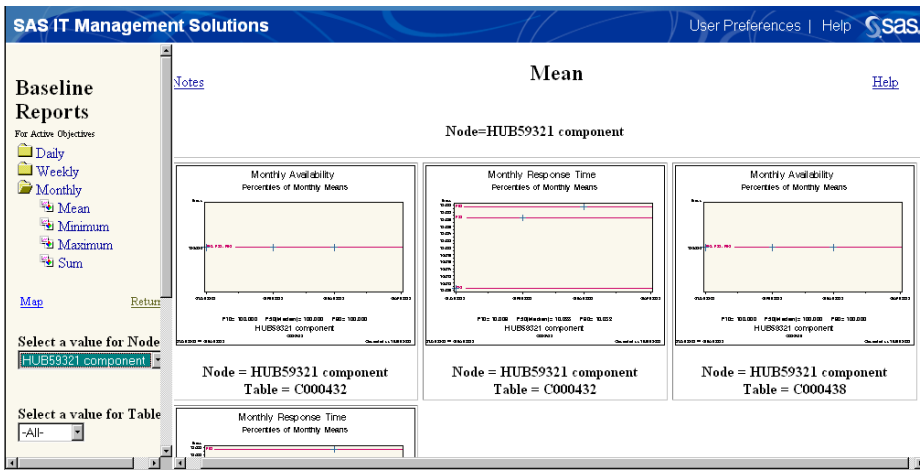

Note: When selecting values in the drop-down lists, it is possible to select a pair of values that result in no graphs in the right panel. When selecting a **`Select a value for Table`**, set the node value to **`-ALL-`**. When selecting values in the drop-down lists, it is possible to select a pair of values that result in no graphs in the right panel. When selecting a **`Select a value for Table`**, set the node value to **`-ALL-`**. △

8 To see the large version of a thumbnail graph in the baseline gallery, click on a thumbnail graph.

Display 10.8 Baseline Gallery —Thumbnail Graphs

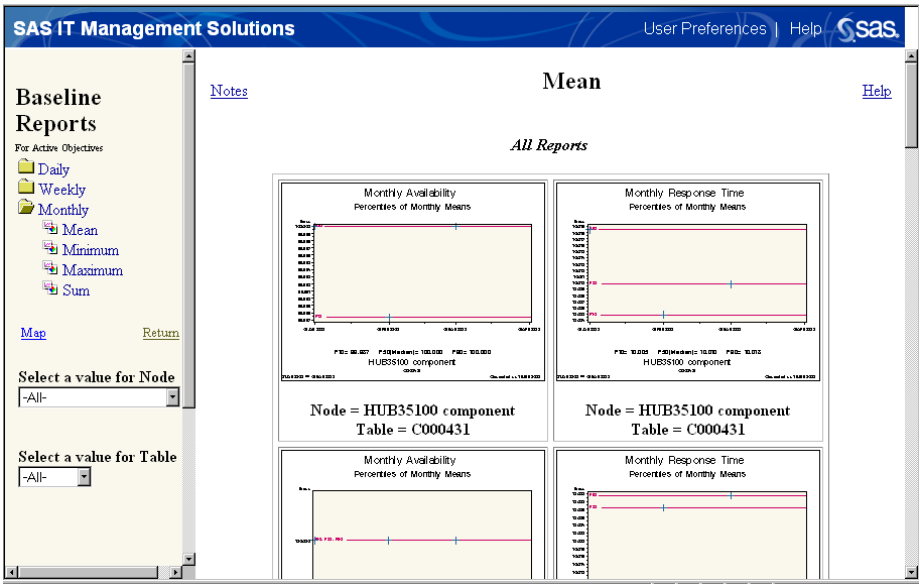

A large version of the selected graph opens. The unique table name is located below the name of the node.

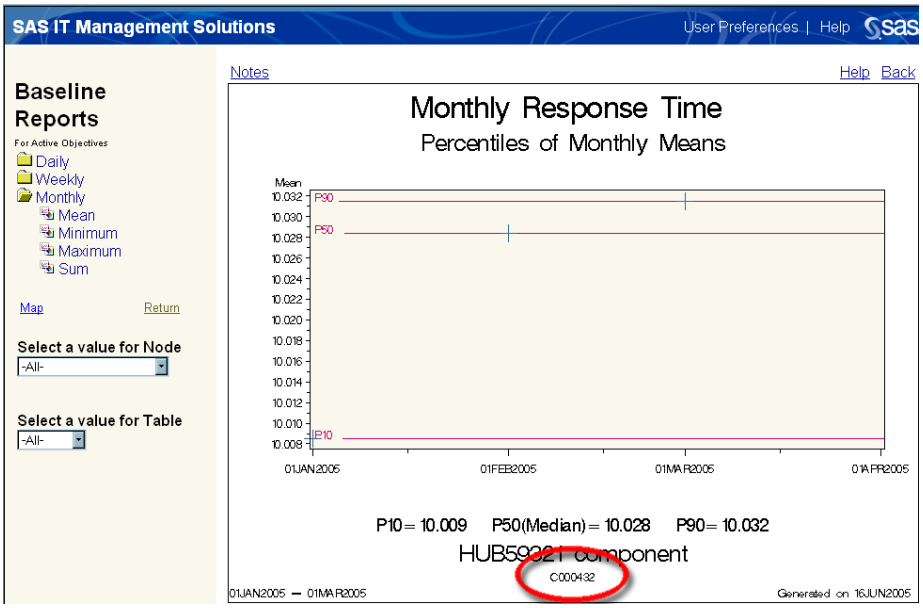

Percentile reference lines are drawn for the 10th, 50th, and the 90th percentile. Each graph contains a footnote which indicates the actual percentile values.

- In the daily graph, the percentiles are of the daily values for all the data in the report.
- In the weekly graph, the percentiles are of the weekly values for all the data in the report.
- In the monthly graph, the percentiles are of the monthly values for all the data in the report.

Note: Percentile reference lines are drawn only if two or more data points exist on a graph. For example, if a contract database includes only the data for January 1, 2005 through January 31, 2005, then on the monthly graphs for January, only one data point exists. Therefore, no percentile reference lines are drawn in the monthly graphs for January. △

CHAPTER

11

About Service Level Reports

About Service Level Reports

What Are Service Level Reports?

Service Level reports are HTML-based reports that reflect data in the contract database and that are prepared according to the report specification as defined in the SAS IT Service Level Management GUI. The tables that make up the reports are presented as trees that can be expanded by clicking the  icon; they can be collapsed by clicking the  icon to the left of the table name. Two types of reports can be generated: Contract Provider reports and Customer reports. The contract provider generates and has access to all reports. Customers have access only to those reports for which they are authorized.

Three groups of reports can be specified in the report specification, and they can each be accessed by a tab on the Web page.

Display 11.1 Service Level Reports Page (showing three report tabs)

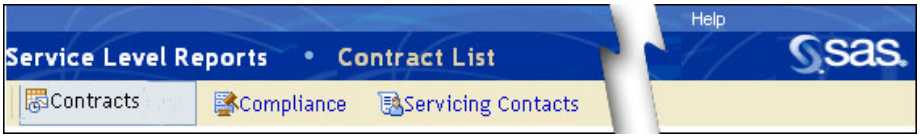

The three groups of reports are:

- The *Contracts* reports:

  These reports provide information about all the data in your contract database. They show how contracts (both active and inactive) are performing over a period of time.

  The *Contracts* reports are always generated. They are available on the **`Contracts`** tab.

- The *Compliance* reports:

  These reports provide a view of how each contract is performing for the current and previous compliance periods that are defined in the contract.

  The *Compliance* reports are generated only if they have been specifically requested in the report specification. They are available on the **`Compliance`** tab.

- The *Servicing Contacts* reports:

  These reports also provide a view of how each contract is performing for the current and previous compliance periods as defined in the contract. They are organized by the servicing contact who has been designated in the Contract window. Each Compliance by Contract report includes all contracts in the report specification that are assigned to that servicing contact.

  The *Servicing Contacts* reports are generated only if they have been specifically requested in the report specification. They are available on the **`Servicing Contacts`** tab.

Service Level Reports and Their Descriptions

The following list shows Service Level reports that can be generated, along with a brief description of each report:

Compliance by Contact Report	displays information that shows whether the contracts and the contract's nodes that are assigned to a servicing contact are meeting the targets of their objectives during a specified compliance period.
Compliance Details Report	displays information that shows whether the contract and its nodes are meeting the targets of their objectives during the current or previous compliance period. The nodes can be filtered to include only those nodes that have missed targets.
Compliance List	displays information that shows whether the contracts are meeting the targets of their objectives during the current or previous compliance period. This report provides a list of all the contracts that are included in the report specification that was run. The list of contracts can be filtered to include only those contracts that have missed targets. This report can also be filtered (or viewed) by customer.
Contract Details Report	displays settings that are associated with the contract, including information such as the calculation methods, range sets, and date filters that were used for the activated objectives of a contract. When this report is generated for a contract provider, it contains more information about contract details than when it is generated for a customer.
Contract List Report	displays a sorted list of all contracts that are defined in the report specification. The report shows how active and inactive contracts are performing across various time periods. A viewer can drill down to see the details of a specific contract along with any annotations that are associated with a contract. The **`Customer`** column is displayed only on contract provider reports. It can be used to sort the contracts by customer.
Customer Report	displays contact and customer information for a contract.
Daily Report	displays the information about one or more objective types for all contract nodes for a selected day. The report displays target and actual values.
Monthly Report	displays the month-to-date value as well as the daily values for a selected month.
Notes Report	displays the notes and annotations that are associated with the contract and with any element of a contract where notes are defined. The user can select any links in annotations that are associated with each note. If the address of the annotation is in the form of a URL (*http://*), then it can be opened as a Web page by using a browser.
Outage Details Report	displays information about an outage. This information includes a display of the date and, if any exceptions exist, the reason for the exceptions to the outage.
Outages Report	displays the outages that have been specified for a contract.

Provider Report	displays contact and provider information for a contract.
Range Set Details	displays information about the range set that is associated with a specified objective. The ranges that make up the range set are displayed along with their range descriptions.
Report Specification Details Report	displays information that was entered for the report specification that generated this set of reports.
Resource Analysis Report	displays the resource data that is associated with a component.
Servicing Contact Report	displays a list of servicing contacts for whom compliance reports have been generated.
Weekly Report	displays the week-to-date value as well as the daily values for a selected week.
Year-by-Month Report	displays the year-to-date value as well as the monthly values for the selected year. The year-to-date and the month-to-date values reflect the values in the yearly and monthly summarization tables, respectively.
Year-by-Week Report	displays the year-to-date value as well as the weekly values for a selected year. The year-to-date and the week-to-date values reflect the values in the yearly and weekly summarization tables, respectively.
5–Minute Interval Table Report	displays a daily table of the data that was collected every five minutes for a resource.

Graphs are available on some service level report pages by clicking the graph icon. Graphs enable you to view the detail and summarized service level data in a form that is a graphic representation of the Daily, Weekly, Monthly, Yearly-by-Week, and Yearly-by-Month report pages. The Daily report contains an additional graph that shows the target and the outage values when a target is set for the objective.

□ To display the graph for a node, click the  icon that is in the row of that node.

Note: If there is no data or only missing data for a given time period, then a graph will not be generated for that time period. For example, if a node does not have any data (or has only missing data) for a particular day, then no graphs will be produced for that node for that day. Similarly, if a node does not have any data (or has only missing data) for a particular week, then no graphs will be produced for that node for that week. The same rule applies to monthly and yearly data. △

How Dates Affect the Report Data

Many service level reports include a column labeled **`Latest Data`**. The value in this column indicates the date and time of the most recent data that is present for each node in the contract database. It enables you to confirm that the ETL process is processing all the data that you expect.

Display 11.2 Compliance List displaying Latest Data column

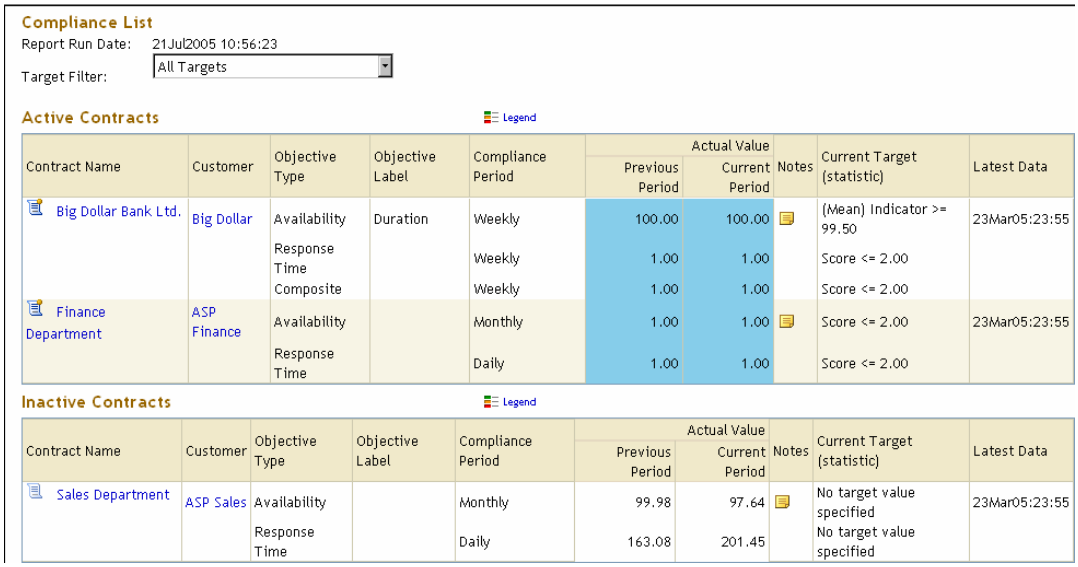

Compliance List

Report Run Date: 21Jul2005 10:56:23

Target Filter: All Targets

Active Contracts Legend

Contract Name	Customer	Objective Type	Objective Label	Compliance Period	Actual Value: Previous Period	Actual Value: Current Period	Notes	Current Target (statistic)	Latest Data
Big Dollar Bank Ltd.	Big Dollar	Availability	Duration	Weekly	100.00	100.00		(Mean) Indicator >= 99.50	23Mar05:23:55
		Response Time		Weekly	1.00	1.00		Score <= 2.00	
		Composite		Weekly	1.00	1.00		Score <= 2.00	
Finance Department	ASP Finance	Availability		Monthly	1.00	1.00		Score <= 2.00	23Mar05:23:55
		Response Time		Daily	1.00	1.00		Score <= 2.00	

Inactive Contracts Legend

Contract Name	Customer	Objective Type	Objective Label	Compliance Period	Actual Value: Previous Period	Actual Value: Current Period	Notes	Current Target (statistic)	Latest Data
Sales Department	ASP Sales	Availability		Monthly	99.98	97.64		No target value specified	23Mar05:23:55
		Response Time		Daily	163.08	201.45		No target value specified	

The value of the **`Latest Data`** is usually the same for all nodes in a contract. However, if you set a date filter for one or more but not all nodes in a contract, then the values can differ across nodes.

About Colors in Service Level Reports

Many Service Level Reports make use of colors to represent actual values. If you create range sets and apply them to nodes in your contract, the defined colors will be used to color the cell of actual values in the report. Reports that use colors will include a **`Legend`** at the bottom of the report, which will show what each color represents. The use of colors enables you to see, at a glance, the actual values that are not performing well.

The colors and labels that are displayed in the reports are affected by the target statistics that you defined when you set up your contract's objectives. For all target statistics except SUM, the colors and labels for a range are displayed for all summary periods. For the SUM target statistic, the colors and labels are displayed for the selected summary period only, and all other periods are displayed using the transparent color.

The Contracts Report Tab

On the **`Contracts`** tab, you can view reports that display actual and targeted values for daily, weekly, monthly, and yearly time periods. If range sets and range descriptions are defined, then the cells in the table are color-coded according to the colors that are specified in the range description.

The following tree shows the reports that are available on the **`Contracts`** tab. It also shows the navigational path that enables you to locate a report:

- Contract List Report
  - Report Specification Details Report
  - Daily Report
    - Resource Analysis Report
      - 5–Minute Interval Table Report

- Weekly Report
- Monthly Report
- Year-by-Week Report
- Year-by-Month Report
- Notes Report
- Contract Hierarchy Report
  - Daily Report
    - Resource Analysis Report
      - 5–Minute Interval Table Report
  - Weekly Report
  - Monthly Report
  - Year-by-Week Report
  - Year-by-Month Report
  - Contract Details Report
    - Customer Report
    - Provider Report
    - Range Set Details Report
  - Outages Report
  - Outage Details Report

The Compliance Report Tab

On the **`Compliance`** tab, you can view reports that enable you to quickly determine whether a contract and its nodes are meeting the targets for each of the objectives of that contract over the current or previous compliance period. You can drill down into the details of individual contracts in order to determine the values of the non-compliant nodes. You can view multiple contracts and see their compliance across all objective types.

The following tree shows the reports that are available on the **`Compliance`** tab. It also shows the navigational path that enables you to locate a report:

- Compliance List Report
  - Compliance Details Report
    - Resource Analysis Report
      - 5–Minute Interval Table Report
  - Customer Report

The Servicing Contacts Report Tab

On the **`Servicing Contacts`** tab, you can view reports that display compliance information for the contracts and the nodes within the contracts that are serviced by a selected contact. Select the servicing contact whose reports you want to view from the *Contact List* that is displayed on the left side of the page.

The following tree shows the reports that are available on the **Servicing Contacts** tab. It also shows the navigational path that enables you to locate a report:

- Servicing Contact Report
  - Compliance by Contact Report
    - Resource Analysis Report
      - 5–Minute Interval Table Report

Customer and Provider Reports

Service Level reports are generated from two types of report specifications: Customer report specifications and Contract Provider report specifications. Some reports that are generated for the contract provider are not available for customers. If you do not want the Contract Provider reports to be available to your customers, you can specify a secure output location that is different from the output location of the Customer reports.

For more detailed information about the reports, see "Introduction to Service Level Reports for Contract Providers and Customers" on page 144.

The following table shows the reports that can be produced by using the contract provider report specification and those that can be produced by using the customer report specification.

Table 11.1 Contract Provider Reports and Customer Reports

Report Name	Contract Provider	Customer
Reports That Are Available on the Contracts Tab		
Contract List	Yes	Yes
Report Specification Details	Yes	Yes
Notes	Yes	Yes
Contract Hierarchy	Yes	Yes
Daily Report	Yes	Yes
Resource Analysis and the optional 5-Minute Interval Table	Yes	No
Weekly	Yes	Yes
Monthly	Yes	Yes
Year by Week	Yes	Yes
Year by Month	Yes	Yes
Contract Details	Yes	Partially - see Note
Customer	Yes	Yes
Provider	Yes	Yes
Range Set Details	Yes	Yes
Outages	Yes	Yes
Outage Details	Yes	Yes
Graphs	Yes	Yes
Reports That Are Available on the Compliance Tab		

Report Name	Contract Provider	Customer
Compliance List	Yes	Yes
Compliance Details	Yes	Yes
Resource Analysis and the optional 5-Minute Interval Table	Yes	No
Reports That Are Available on the Servicing Contacts Tab		
Servicing Contact	Yes	No
Resource Analysis and the optional 5-Minute Interval Table	Yes	No

Note: The Contract Details Report for a customer excludes the following information: version, path, updated by, updated, contract PDB table, cost information, contract database, penalty data, and credit information. △

Working with Service Level Reports

After you create your contracts and successfully run ETL on them, you can set up and run reports that reflect the data that is in the contract database. The following diagram shows the typical processes that you can use to generate the reports.

Display 11.3 Working with Service Level Reports

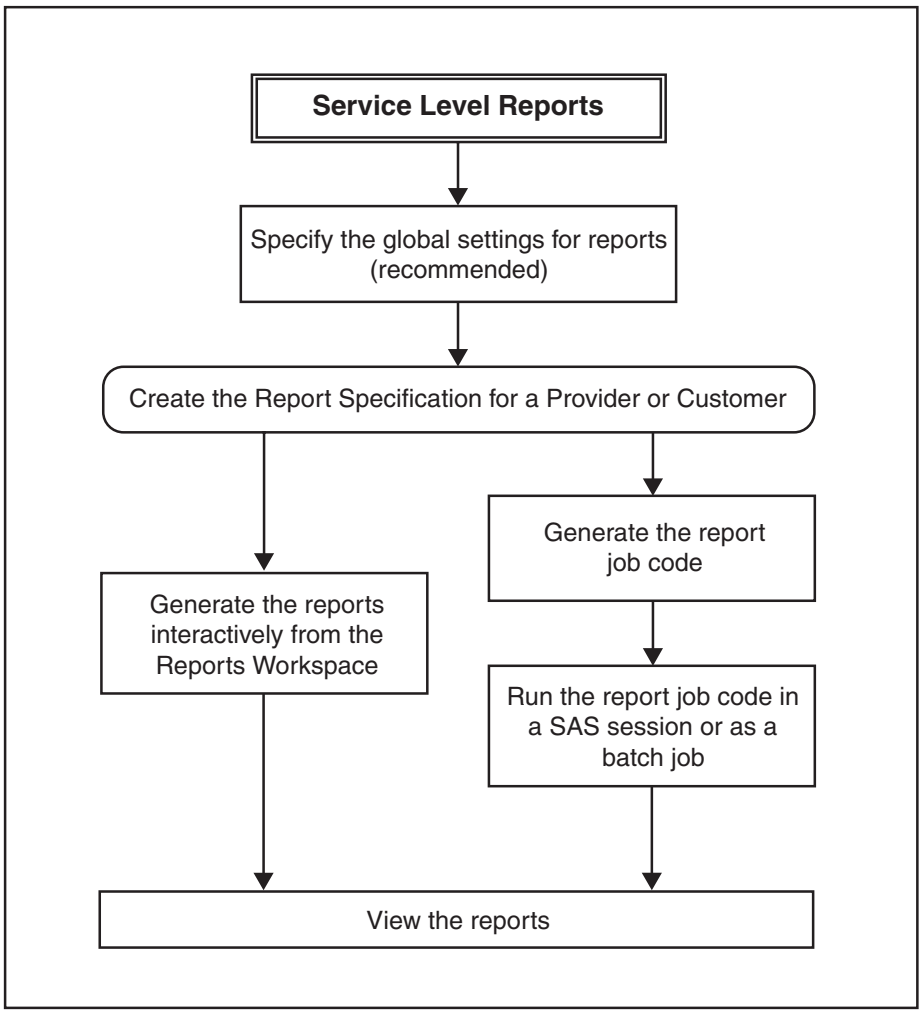

The following list provides a brief description of these processes:

1 *Specify the global settings for the reports.*

The Report Global Settings dialog box on the Reports workspace enables you to specify the default output locations (both the physical path and the URL) for the contract provider and customer reports. In addition, you can override the style sheet that governs the appearance of the reports and determine the placement of the decimal points for the actual and targeted values that are displayed for the objective types.

2 *Create the report specifications for a provider or customer.*

The Report Specification Wizard enables you to define the options that will govern the generation of the contract provider and customer reports. Once created, the report specification can be accessed and modified by using the Report Specification window.

3 Run the report specifications. They can be run interactively or in batch mode.

□ *Generate the reports interactively from the Reports workspace*:

You can run a report specification interactively from the GUI. You can then view the reports in the Web browser that you have specified.

- *Generate the reports to run in batch mode*:

  Generate the report job code: You can create job code that can be run from a SAS session. You can also "wrap" the job code in platform-specific instructions that will run the reports on your operating environment. You can view the reports by navigating to the output location that you specified.

More detailed information is presented in the following sections:

About the Reports Workspace

The Reports workspace enables you to perform the tasks relating to the global settings for the reports, the creation of the report specifications for contract providers and customers, and the generation of service level reports for contract providers and customers.

Display 11.4 The Reports Workspace

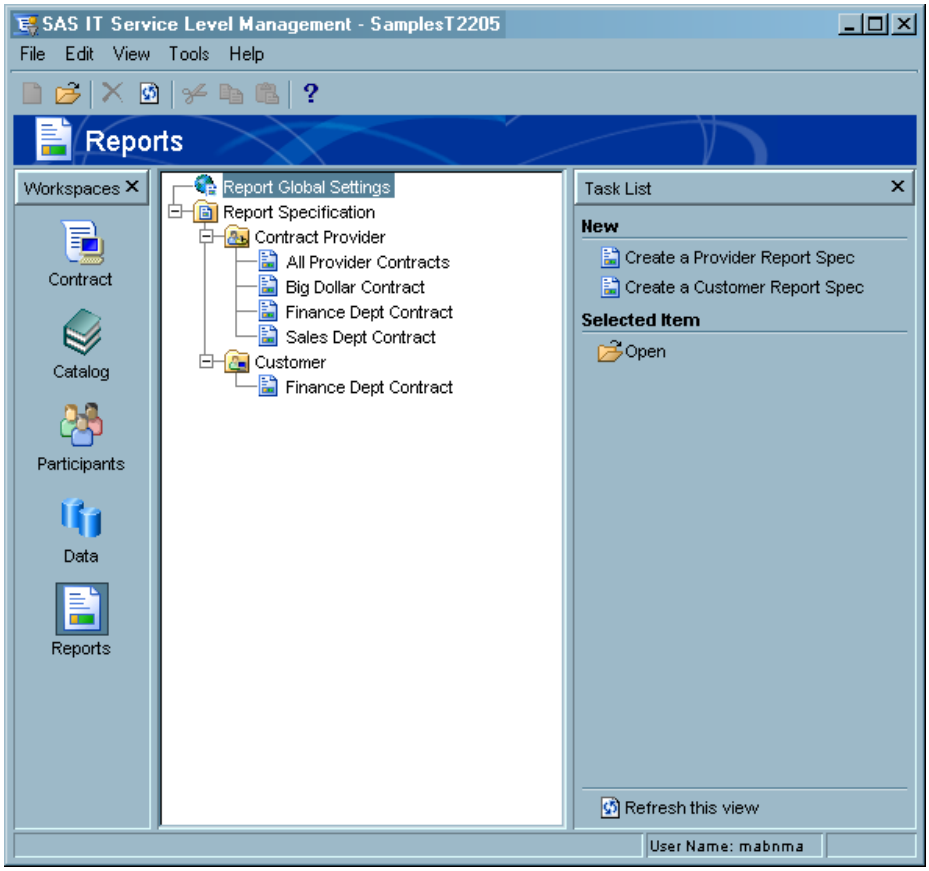

The Reports workspace contains two main folders: **Report Global Settings** and **Report Specification**. The **Report Specification** folder contains two subfolders: **Contract Provider** and **Customer**. You can use these folders and subfolders to organize the customer and contract provider report specifications that you create.

Organizing the Reports Workspace

You can use the following features to create and manage the report specifications for customers and contract providers:

- From the **Task** list of the Reports workspace, you can create a customer report specification or a contract provider report specification.
- In the Report Global Settings dialog box, you can override the global output location for customer reports by specifying a different report output location for each customer.
- In the Reports workspace, you can create a subfolder that can contain the report specifications for each customer or for the contract provider:
  1. Right-click the appropriate **Customer** or **Contract Provider** folder in the navigation tree. Select **New ▶ Folder**.
  2. In the **Folder Properties** dialog box, enter the folder name that corresponds to the customer or contract provider. Then click **OK** to create the new customer or contract provider subfolder.

Display 11.5 Creating Folders in the reports Workspace

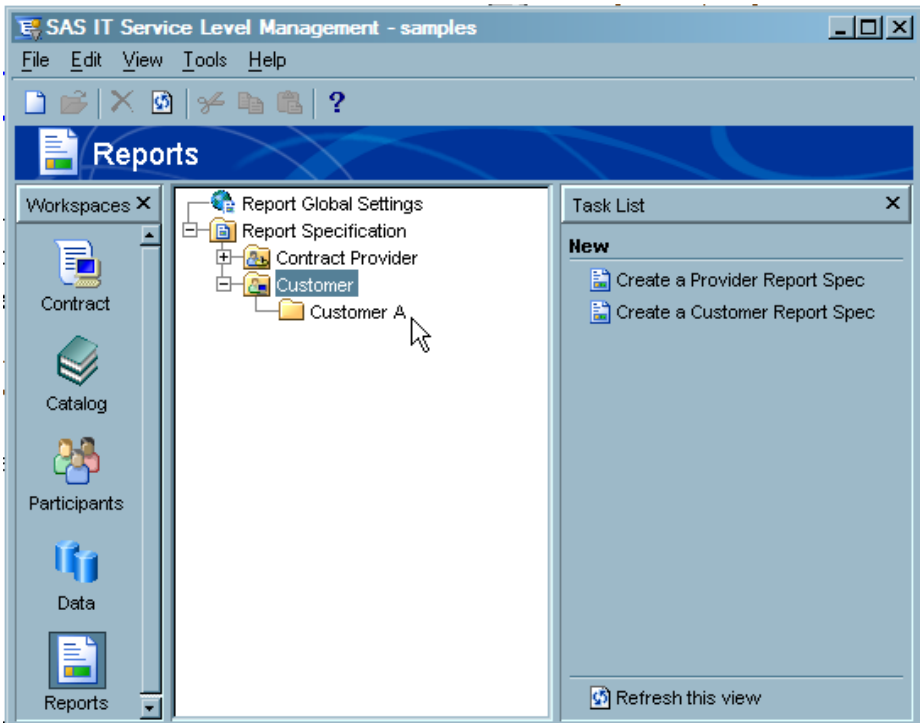

Specifying Global Settings for Reports

The Reports workspace provides the capability to set the default parameters that are applied globally during the process of creating a report specification and generating reports. The default parameters that specify locations can be overridden in an individual report specification.

The **Report Global Settings** dialog box enables you to specify the following defaults:

- Output locations for customer reports
- Output locations for contract provider reports
- Location of the customized style sheet
- Number of decimal places to be displayed for the values in the reports (These values cannot be overridden on individual report specifications.)

To specify the report global settings, perform these steps:

1 Double-click **Report Global Settings** in the Reports workspace to open the **Report Global Settings** dialog box.

Display 11.6 Report Global Settings dialog box

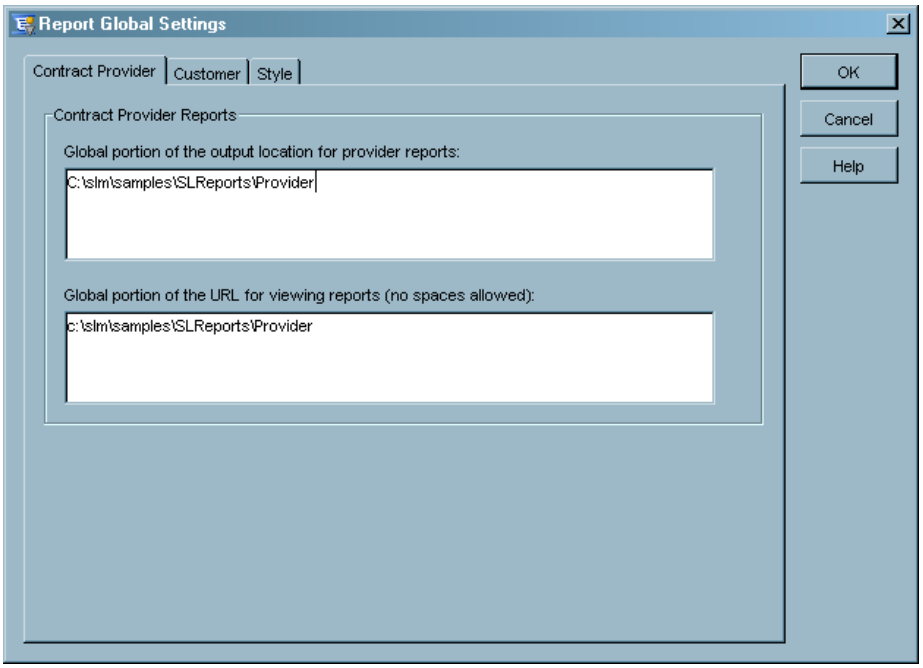

2 You can use the **Contract Provider** or **Customer** tabs of the **Report Global Settings** dialog box to create or modify the default output location for contract provider or customer reports. This output location is automatically used in any new report specifications that you create unless you override it by specifying a different location in an individual report specification. If a directory that you specify in the output location does not exist, it is created when the reports are generated.

To create or modify the default output location for contract provider or customer reports, select the appropriate tab: **Contract Provider** or **Customer**.

Then perform these steps:

a In the upper text box, enter the global portion of the physical location where you want the reports to be located.

Note: When entering pathnames, directory names, and filenames for your output location, follow the syntax rules for your server's operating environment. Avoid any special characters that are prohibited or that have special meaning for your operating environment. For z/OS, a hierarchical file system that is supported by UNIX System Services (USS) is required to store

generated report output. This can be a Hierarchical File System (HFS) or a zSeries File System (zFS). Therefore, if your SAS server is a z/OS server, be sure to enter a UNIX-style path. △

b In the lower text box, enter the global portion of the URL where the reports can be viewed interactively. Spaces are not allowed. Do not use an ending slash. The application will automatically append **`/text/index.html`** to the path that you specify to designate the starting page for the Web-based reports.

c If you have no more changes to enter in the Report Global Setting dialog box, click **`OK`** to validate and then save your changes.

3 You can use the **`Style`** tab of the Report Global Settings dialog box to point to a custom cascading style sheet (CSS) that you created for your reports. To do so, perform these steps:

a Select the **`Style`** tab.

Display 11.7 Style Tab in the Report Global Settings Dialog Box

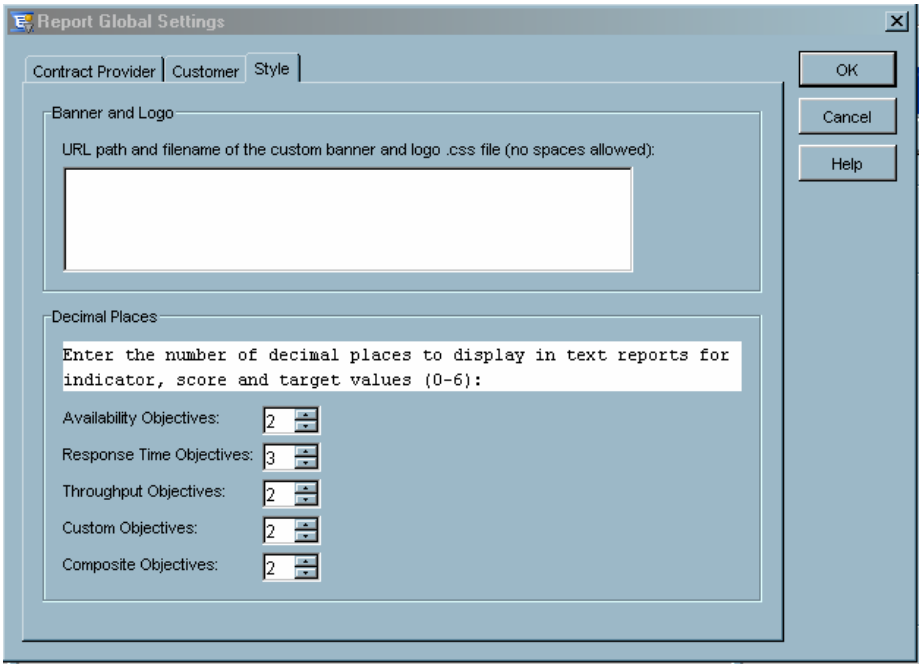

b Enter the URL path and the filename of the custom CSS file that you want to use for your reports. The path and the filename must be in the format of a URL, for example: **`http://mypath/mystyle.css`**. Make sure that this URL can be accessed by the Web-based reports. Spaces are not allowed in this path.

c If you have no more changes to enter in the Report Global Setting dialog box, click **`OK`** to validate and then save your changes.

Note: Use of the custom style sheet is optional. For information about how you can use the custom cascading style sheet, see "Customizing Reports with a Cascading Style Sheet" on page 122. △

4 You can also use the **Style** tab of the Report Global Settings dialog box to specify the number of decimal places for each selected objective type in the reports. To do so, perform these steps:

   a Select the **Style** tab.
   b In the **Decimal Places** part of the tab, use the arrows to specify the location of the decimal place in the indicator, score, and target values of each objective. You can specify from 0 through 6 decimal places.

      The following table shows the default number of decimal places for each objective type:

Table 11.2 Default for Decimal Places

Objective Type	Default Number of Decimal Places
Availability	2
Response Time	3
Throughput	0
Custom Objective	0
Composite	0

      Decimal places can be specified only in the Report Global Settings dialog box. The settings for decimal places cannot be overridden in an individual report specification.
   c If you have no more changes to enter in the Report Global Setting dialog box, click **OK** to validate and then save your changes.

CAUTION:
In service level reports, the actual data values are rounded to the decimal setting. For example, if the actual data value is 0.5 or greater and the decimal setting is set to 0, then the value is rounded up and is displayed on the report as 1. Similarly, if the actual data value is less than 0.5 and the decimal setting is set to 0, then the value is rounded down and is displayed on the report as 0. △

Customizing Reports with a Cascading Style Sheet

You can customize the banner, the logo, or the Help text of your reports by creating your own custom cascading style sheet (CSS). To customize both the banner and logo of your service level reports, you can create the following files:

- a customized CSS file in ASCII format
- a GIF image for the background (optional)
- a GIF image for the logo (optional)

  (The recommended size is 24 pixels wide by 24 pixels high.)

In the customized CSS file that you create, you can specify a value for a style definition in order to change a specific property. A style definition is prefaced by **div#**. The values for any style definition that you specify in your customized CSS file will replace the default values in the default CSS file. To view the report output from the Web, the customized CSS file must be placed on a Web server.

The following table describes the style definitions for a customized CSS file that are supported by SAS IT Service Level Management:

Table 11.3 Style Definitions That You Can Specify in a Customized CSS File

Style Definition	Description
div#slmbanner	Specifies either the color or the location of a background image (URL of a GIF file) for the banner, or both. If you are customizing the color of your banner, then you can use the name of the color (or its hexadecimal value) in combination with the location of your banner in URL format. If you change only the background color of the banner, then you do not need to specify the location of a GIF.
div#slmlogo	Specifies the location of the logo (URL of a GIF).
div#slmbannertd#helptext	Specifies the color (the color name or its hexadecimal value) of the Help text.

Example

This example shows the code in a customized CSS file. The .CSS file specifies green as the background color of the banner, points to the custom logo (logo.gif) by using a URL address, and specifies white as the color of the Help text.

```
/*slmbanner -- override this to change the banner background*/
div#slmbanner{
     background: green;
}
/*slmlogo -- override this to change the logo*/
div#slmlogo td#logo{
     background: url(http://MyWebServer/logo.gif);
}
/*slmbanner td#helptext -- override this to change the color of the help text*/
div#slmbanner td#helptext{
     color:#FFFFFF;
}
```

Creating or Modifying the Report Specification

In order to generate service level reports, you must first create the report specification that provides the information that is needed to define them. Report specifications are report definitions that contain parameters that you can set for generating a contract provider report or a customer report. Report specifications enable you to create and modify specific report configurations that can be saved and reused to generate routine reports.

- To create a report specification, invoke the Report Specification Wizard. This Wizard guides you through the steps to create a *report specification*. From each page of the Wizard you can do the following:
  - Make selections with your left mouse button.
  - When you are finished making selections, click **Next** at the bottom of the page to continue to the next step.
  - To return to a previous step, click **Back.**
  - To close the Wizard without saving your entries, click **Cancel**.
  - On the last page of the Wizard, click **Finish** to complete the report specification process.

- To get help for a page, click **Help**.

- To modify a report specification, open the appropriate tabs of the Report Specification window and change the settings.

Creating a Report Specification

To create a report specification for a contract provider or a customer, perform the following steps:

1 From the **Task** list, select **Create a Provider Report Spec** or **Create a Customer Report Spec**.

Display 11.8 Accessing the Report Specification Wizard from the Task List

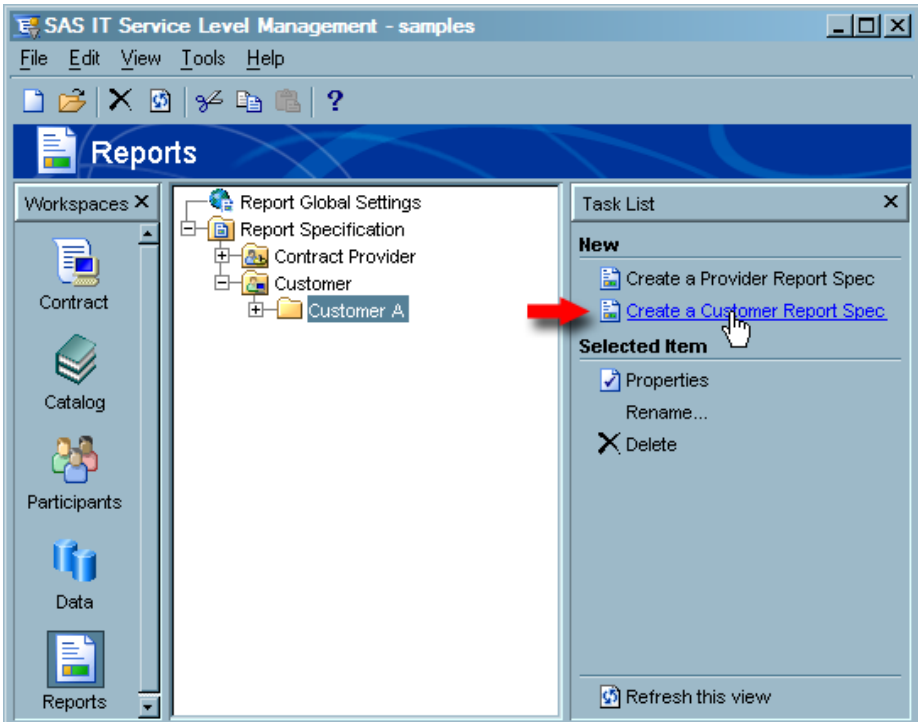

Note: You can also access the Report Specification Wizard from the menu or the toolbar. To do so, highlight the folder that corresponds to the type of report specification that you want to create. From the **File** menu or the toolbar, select **New**. Then select **Report Specification**. △

2 The **`Identify the Report Specification`** page of the Wizard opens.

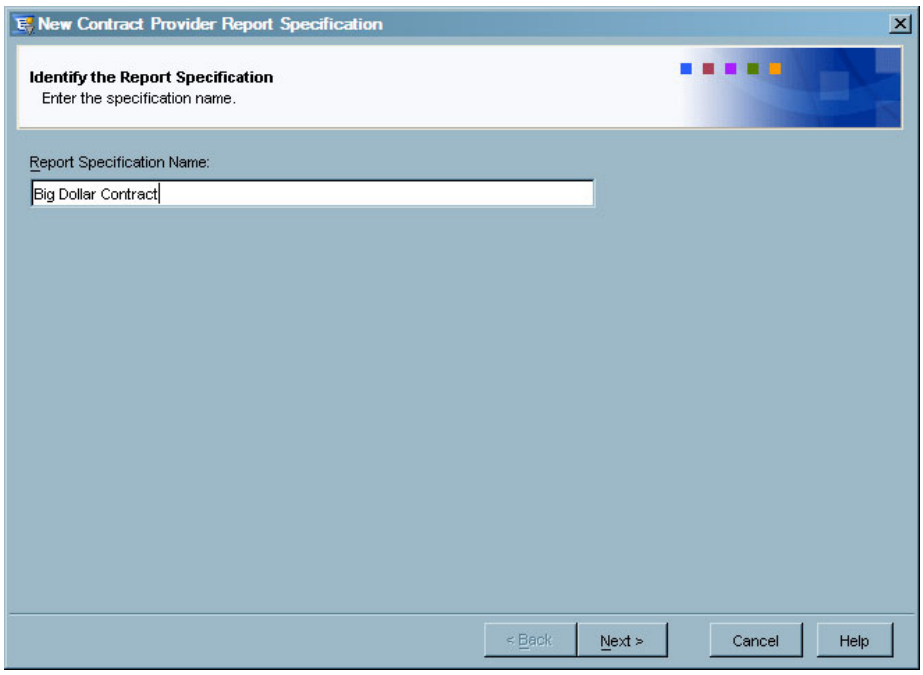

Specify a name for the new report specification that you want to create. This field is required. This field can contain a maximum of 40 characters and can include spaces.

If you opened the Report Specification Wizard from the **`File`** menu or the toolbar, the **`Identify the Report Specification`** page includes an additional option that enables you to indicate whether you are creating this report specification for a contract provider or for a customer.

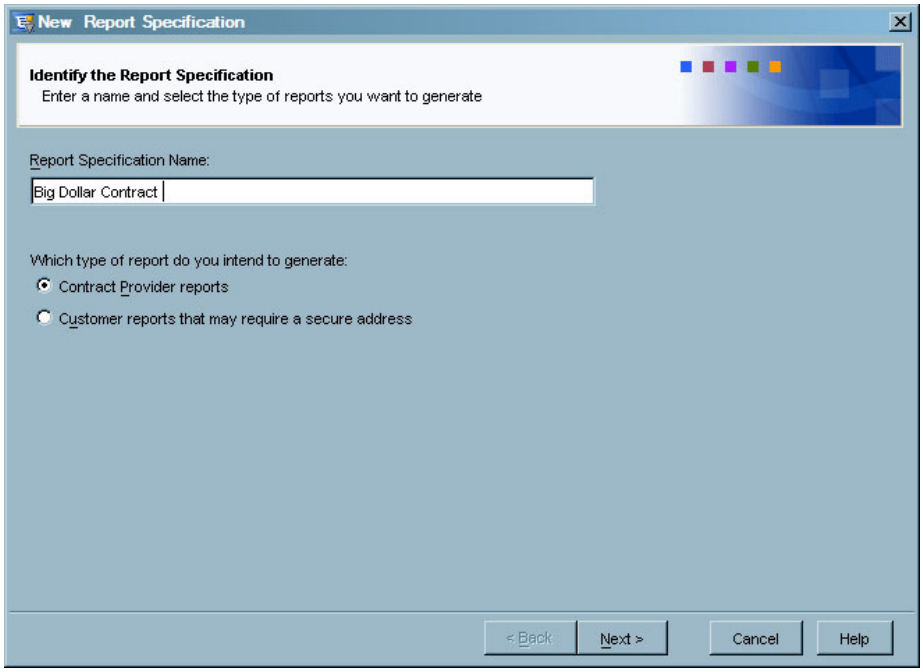

Click the type of reports that you want to generate:

- □ Contract Provider reports — selected by default.
- □ Customer reports

3 On the **`Which Contracts?`** page, select the contracts that you want to include in the report specification.

Display 11.9 Select the Contracts (Report Specification Wizard)

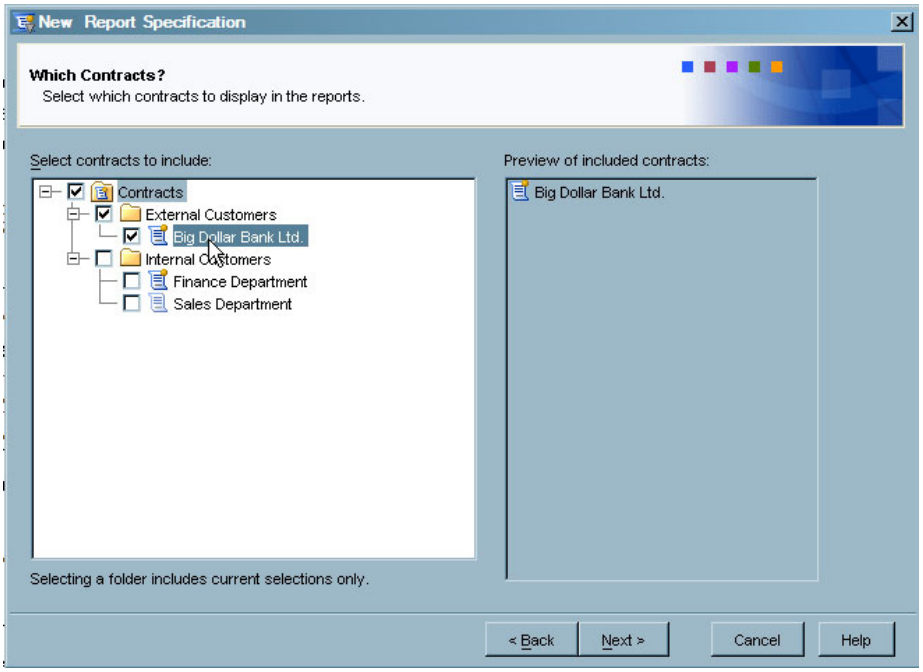

The left side of the page provides a hierarchical list of the available contracts from which you can choose. The right side of the page displays the contract(s) that have been selected. An icon displays whether the contract is active or inactive.

To select the contracts that you want to include in this report specification, click the appropriate check box. You must select at least one contract. You can also select all contracts, an entire folder, or one or more contracts within one or more folders. If a contract folder is checked, then every contract currently within that folder is checked. However, as new contracts are added to your environment, they will *not* be automatically selected even if they are contained within a checked folder.

To revise your list of contracts, change the selections that you have made on the left side of the page.

Note: The **`Preview of included contracts`** list on the right section of this page lists the contracts that will be included in this report specification as a result of the choices you have made. △

4 On the **Optional Reports?** page, select the optional reports that you want to include in your report output. If you highlight each report in the list on the left side, then its description displays on the right side of the page. For example, in the following display the second type of report, All Servicing Contacts, is highlighted, and its description appears on the right side.

Display 11.10 Select the Optional Reports (Report Specification Wizard)

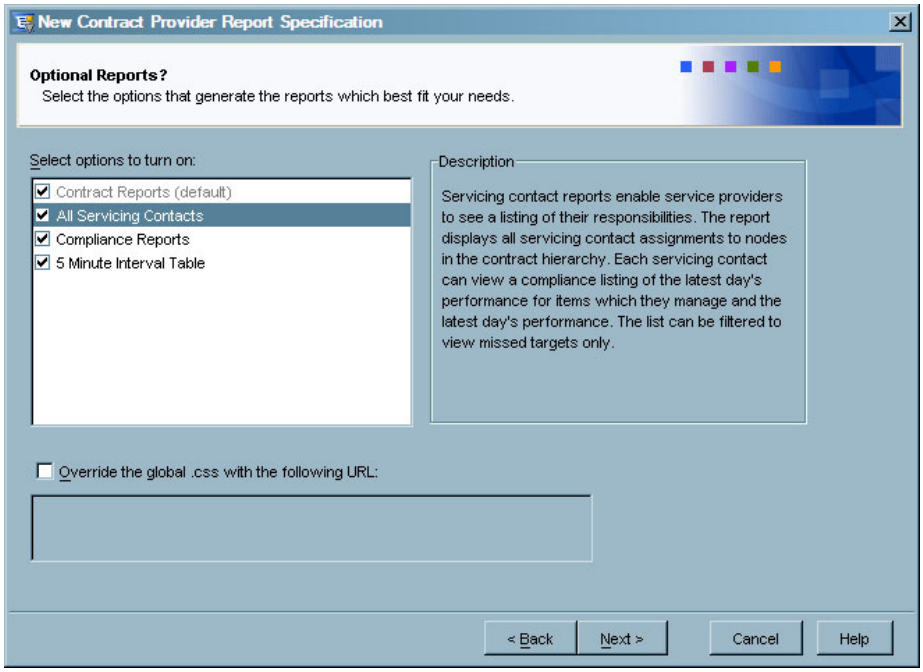

To select which optional reports you want to generate, follow these steps:

a From the list under **Select options to turn on**, click the appropriate check box to select the report(s) that you want to generate. The following table shows the types of reports that can be selected for a contract provider or for a customer report specification.

Table 11.4 Reports by Type of Report Specification

Types of Reports	Reports for Contract Providers	Reports for Customers
Contract Reports	Yes	Yes
All Servicing Contacts	Yes	No
Compliance Reports	Yes	Yes
5-Minute Interval Table	Yes	No

b To override the style sheet that is in effect, select **Override the global .css with the following URL** and enter the path to the customized style sheet that you want to use.

Note: You do not have to specify the path and filename for a custom CSS. If a custom CSS is not entered in a report specification or in the Report

Global Settings dialog box, then the standard SAS IT Service Level Management style sheet is used for the reports. △

5 On the **What Location?** page, select the location that you want to use for the report output.

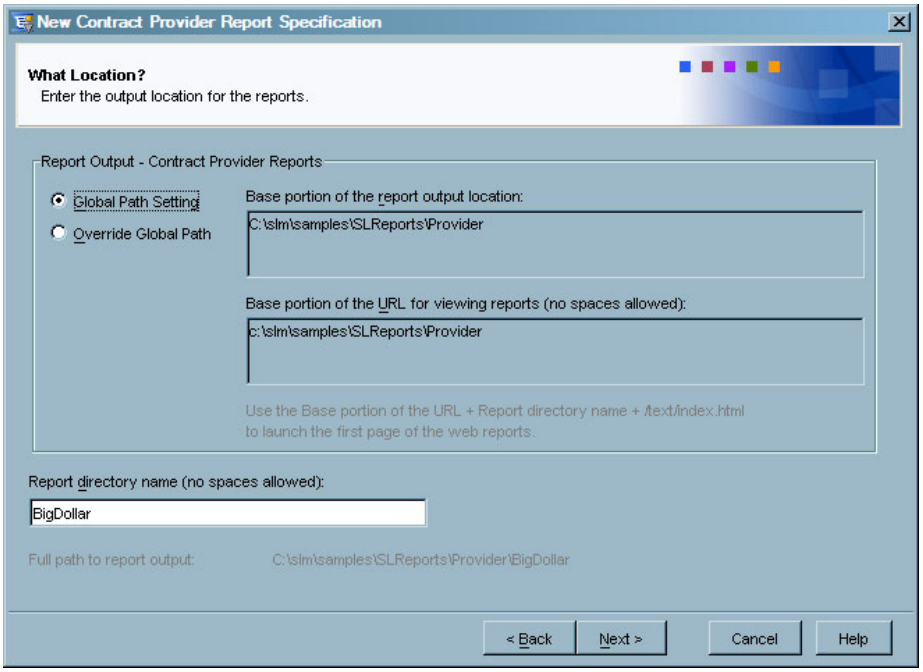

The value that you enter in the **Base portion of the report output location** field and the value that you enter in the **Report directory name** field are used to build the report output location where your reports will be located. In addition to the values that you enter, the system appends a folder named "text" to the end of the output location. A file called **index.html** in this path is the entry point to your service level reports. For example, an output location for a contract provider report specification is:

C:\slm\samples\SLReports\Provider\BigDollar\text\index.html

The output location is slightly different for reports in a customer report specification because the reports for each customer are created in a unique folder. The system uses the contract customer name to create a unique folder name. Any non-alphanumeric characters in the customer name are converted to hexadecimal values in order to create a valid file system folder name. The output location is built by inserting the unique customer name folder before the folder named "text". For example, an output location for a customer report specification is:

C:\slm\samples\SLReports\Customers\John20Smith\text\index.html

Note: When entering pathnames, directory names, and filenames for your output location, follow the syntax rules for your server's operating environment. Avoid any special characters that are prohibited or that have special meaning for your operating environment. For z/OS, a hierarchical file system that is supported by UNIX System Services (USS) is required to store generated report output. This can be a Hierarchical File System (HFS) or a zSeries File System (zFS). Therefore, if your SAS server is a z/OS server, be sure to enter a UNIX-style path. △

To select the output location that you want to use, perform these steps:

a If you want to use the base portion of the report output location that was specified in the **`Global Path Setting`** dialog box, then leave the **`Global Path Setting`** option checked (the default).

b If you want to enter a path that is different from the default global path setting, then select the **`Override Global Path`** option and enter that path in the output location field.

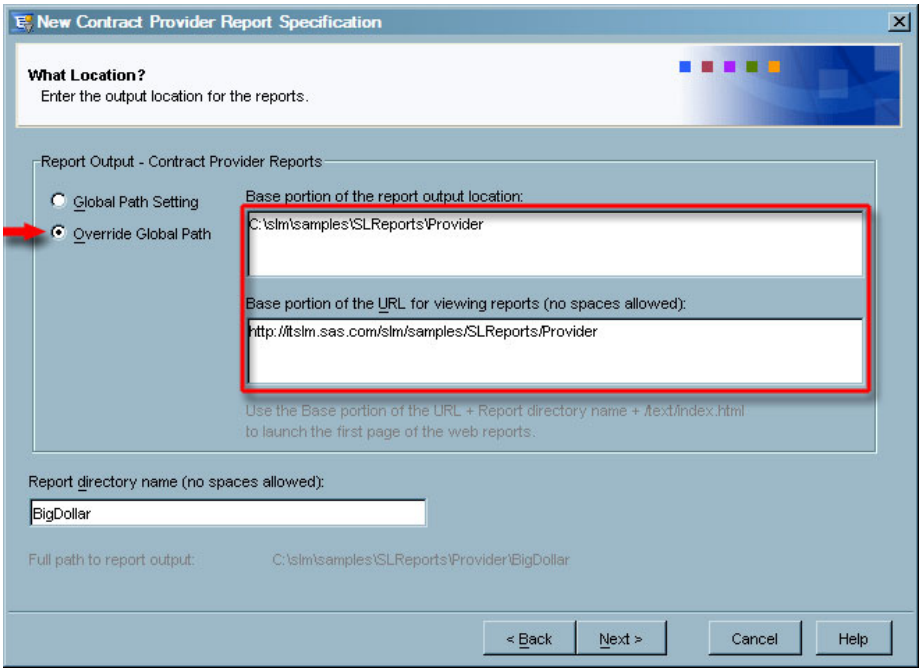

In the field that is labeled **`Base portion of the report output location`**, enter the path to the directory where you want the reports to be located.

In the field that is labeled **`Base portion of the URL for viewing reports`**, enter the path to the report output location in URL format. (This is optional).

c In **`Report directory name`**, enter the name of the directory where the reports for this report specification will be located. This field is required. This name is appended to the path that is defined above.

The **`Full path to report output`** is located below the text box for the report directory name. It displays the output location that is the result of your entries. For example, if the global path is **`C:\slm\samples\SLReports\Provider`**, then the full address is: **`C:\slm\samples\SLReports\Provider\BigDollar`**.

CAUTION:

If you attempt to create a new report output location, the underlying server process might not have permission to accomplish that task. For a description of this problem and information about how to resolve it, see your systems administrator. Refer to SAS Note SN-015642 on the Technical Support Web site for information about how to resolve the problem for the Windows 2003 Server. △

6 On the **Which Range Settings?** page, you can customize the colors in the report output so that they conform to certain user requirements. The range settings that you choose will specify the labels and the background colors of the cells in the tabular reports according to the values that they contain.

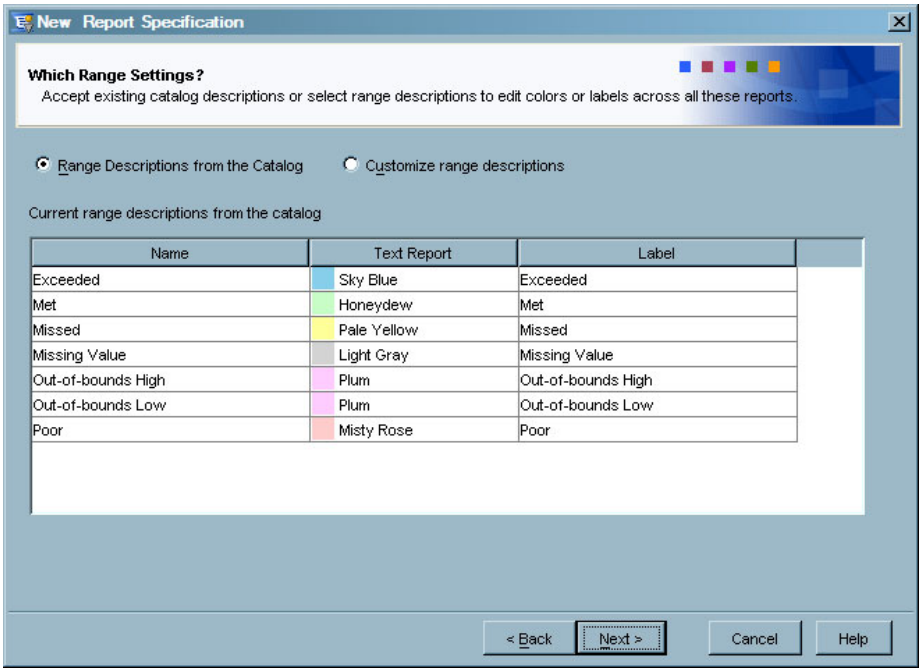

The range descriptions that display on this page are those that are currently contained in the **Range Descriptions** folder of the Catalog.

To select the range description that you want to use, follow this procedure:

a To use the range descriptions exactly as they appear in the Catalog, click **Range Descriptions from the catalog**.

b To change the selected colors or then labels of the range descriptions, click **Customize range descriptions**. When this option is selected, a new column, **Custom**, is appended to the **Customized list of range descriptions**. If a description is customized, a check mark is displayed in the **Custom** column for that description.

To change the colors or labels of a range description, perform these steps:

i From the list of range descriptions, select the range description that you want to change.

ii Click **Edit** to open the Range Description dialog box.

iii To change the color of the selected range to be used on the text reports, select the new color from the drop-down list.

iv To change the label of the selected range, click **Customize the report label**. Then enter the new label in the adjacent text field.

v Click **OK** to save the changes and return to the **Which Range Settings?** page.

To change the selected range description back to its original setting, follow these steps:

i From the list of range descriptions, select the one that you want to change.
ii Click **Reset**. The color and label of the selected description revert to the values shown in the Catalog. The check mark in the **Custom** column is removed.

7 On the **Run Reports** page of the Report Specification Wizard, you can run the reports interactively or you can create the job code that can run the reports later in a SAS session or in a batch job.

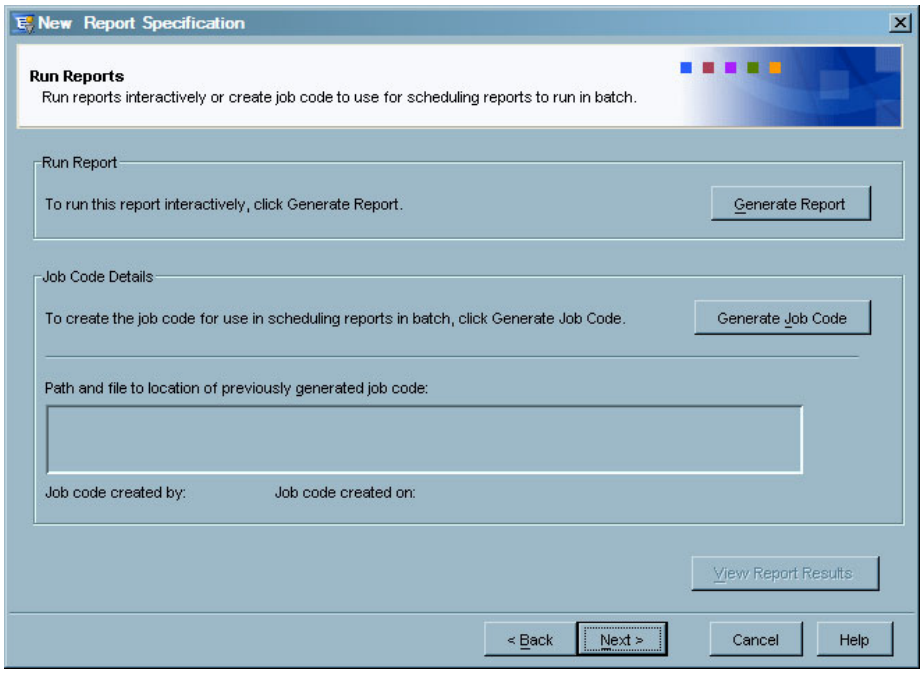

To run reports interactively from this page, follow these steps:

CAUTION:
Running the reports interactively might take a long time. It also locks the application while it runs. To save time, run only small report jobs interactively. △

a Click **Generate Report**.
b A message box appears advising you that report generation might take a long time. If you want to continue, then click **Yes**.
c When the process is finished, a **Report Results** dialog box opens. Check for errors or warning messages that might be displayed in the **Report Results** dialog box.
d If you are running contract provider reports and have specified a valid entry for the base portion of the URL for viewing reports, and if the reports are generated successfully, then the reports are displayed in your browser. If you did not specify a URL for viewing the reports, then you will need to locate the **index.html** file in the **text** directory at the report output location that you specified.
e If you are running customer reports, then you need to locate the **index.html** file in the text directory of the customer directory at the report output location that you specified.

To create the report job code in order to run reports in batch mode, see "Creating Report Job Code" on page 132.

8 After you successfully save your information and exit the wizard, the report specification name that you specified here will be placed in the appropriate contract provider or customer folder on the Reports workspace.

Modifying the Report Specification Wizard

To modify the report specification after it is created, perform these steps:

1 In the Reports workspace, navigate to the contract provider or customer folder that contains the report specification that you want to modify.
2 Double-click the report specification to open the Report Specification window.
3 Select the tab that you want to change.
4 Modify the settings.
5 Click **`Regenerate all reports with the next run`** if you want to regenerate the reports. Selecting this option is only in effect for the next time that reports are run. Because regenerating reports requires extra time and locks the application, plan accordingly.

   For information about this procedure, see "Regenerating Service Level Reports" on page 134.
6 When you have finished making your changes, click **File ▶ Save** to save them and return to the Reports workspace.

Generating Service Level Reports

After you create the report specifications for your service level reports, you can generate the reports interactively or by creating report job code and then running that job code from a SAS session or in batch mode.

Generating Reports Interactively

You can generate service level reports interactively in two ways.

□ *From the Reports workspace*:

1 Highlight the existing report specification that you want to use.
2 Select **`Generate Report`** from the **`Task`** list.
3 Click **`Yes`** in the message dialog box that displays the message about the length of time that report generation might require.

□ *From the Run Reports page of the Report Specification Wizard*:

1 Click **`Generate Report`**.
2 A message dialog box displays a message that advises you that this task might require a lot of time. To proceed with report generation, click **`Yes`**.

Note: If you generate reports interactively, you do not need to generate job code. △

Creating Report Job Code

To generate service level reports by using report job code, you must first create the job code. *Report job code* is a SAS program that generates the service level reports. It obtains the instructions for accessing, analyzing, and displaying the contract data from the report specification that you created. The job code can be used in a batch program in order to generate service level reports at regular intervals.

You can create the report job code in two ways.

- From the **Run Reports** page of the Report Specification Wizard, perform these steps:
  1. Click **Generate Job Code**.
  2. In the dialog box that appears, click **Yes** to save the report specification and to open the **Create Report Job Code** dialog box.
  3. In the text field of the **Create Report Job Code** dialog box, enter the path and filename where you want the report job code to be saved. This field is required. It can contain a maximum of 200 characters.
  4. Click **OK** to create the report job code at the location that you specified and return to the Reports workspace.

  The report job code is generated. The report ID, the user ID of the person who created the job code, and the date and time that the job code was created are displayed at the bottom of the page.

  If the report job code is not created, a message is displayed and you return to the **Run Reports** page of the Wizard.
- From the Reports workspace, perform these steps:
  1. Double-click the **Report Specification** folder to expand it.
  2. Select the contract provider or the customer report specification for which you want to create the job code.
  3. From the **Task** list, select **Create Report Job Code**.
  4. In the **Path and filename for report job code on <environment>** field the **Create Report Job Code** dialog box, enter the path and filename where you want to save the report job code. This field is required. It can contain a maximum of 200 characters.

     The **Current Environment Hostname** displays the hostname of the current environment. This field cannot be modified.

     Note: When entering pathnames, directory names, and filenames for your output location, follow the syntax rules for your server's operating environment. Avoid any special characters that are prohibited or that have special meaning for your operating environment. For z/OS, a hierarchical file system that is supported by UNIX System Services (USS) is required to store generated report output. This can be a Hierarchical File System (HFS) or a zSeries File System (zFS). Therefore, if your SAS server is a z/OS server, be sure to enter a UNIX-style path. △
  5. Click **OK** to create the report job code at the location that you specified. You then return to the Reports workspace.

     If the report job code cannot be created, a message displays. Click **Cancel** to close the dialog box without saving the report job code and return to the Reports workspace.

Running Report Job Code

To generate the service level reports by using the report job code that you just created, perform these steps:

1. Navigate to the location of the job code that you specified in the **Create Report Job Code** dialog box.
2. Copy the job code into the program editor of a SAS session and click the **Submit** icon.

   Or you can run the report job in batch mode. For more information about running your report job code in batch mode, see "Batch Jobs" on page 297.

The Report Results

Use the **Report Results** dialog box to view the report job log after generating the reports. The report job log can include report processing notes from multiple contracts. It is available whether the reports were generated interactively or from a batch job.

To access the **Reports Results** dialog box for a specific report specification, perform these steps:

1. In the Reports workspace, highlight a report specification from the contract provider folder or the customer folder.
2. From the **Task** list, click **View Report Results**. The **Report Results** dialog box opens. It displays information about the report job that was most recently run.

Regenerating Service Level Reports

When your report job runs, it will normally generate reports for new data and new time periods that have been added to the contract database(s). It will also automatically regenerate the service level reports whose contracts have been changed since the last time that reports were generated. However, there are some situations where you might want to force the regeneration of reports on all of the contracts in a report specification.

The following table describes actions that might require you to regenerate reports:

Table 11.5 When to Regenerate Reports

Action	Explanation
Change the output location	If you make changes to the report output path in the Report Global Settings dialog box or in the Report Specification window, any new reports will be placed in the new output location. The existing reports will remain in the old location. This condition results in broken links in the calendar on the dates that link to any report files that are located in the old output location. *Note:* To avoid regenerating reports in this case, you can move all the existing reports to the new location. △
Change a range description	If you change a range description, the new reports will reflect those changes. The existing reports will display the previous range descriptions.
Change a decimal place	If you change the value for the placement of the decimal in the Report Global Settings dialog box, the new reports will reflect that change. The existing reports will display values using the previous number of decimal places.
Add a new 5-Minute Interval Table report	If you add the 5-Minute Interval Table report to an existing report specification, the new reports will contain the 5-Minute Interval Table report. The existing reports will not.
Add a contract with existing data	If you add a contract to a report specification and the ETL process has been previously run for the contract before the last report run date, then the system does not automatically create reports that include the data from any previous ETL processing.

To regenerate the reports for the contracts that are included in a report specification, follow these steps:

1. Select the **`General`** tab of the Report Specification window.
2. Click the option to **`Regenerate all reports with the next run`** so that a check mark appears.

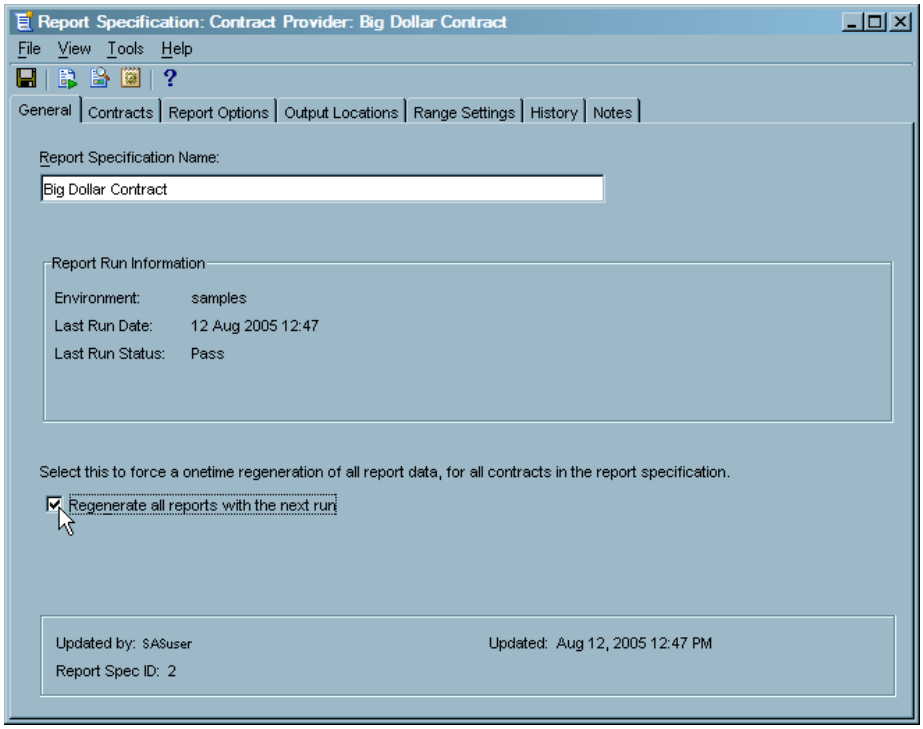

The reports are regenerated the next time that the report job code is run. The next time this tab is accessed, the check box is automatically cleared. This prevents the reports from being regenerated a second time.

How to Optimize Report Performance

Producing reports might require a significant amount of time, depending on the amount of data that exists in the contract database. You can use the following methods to minimize the time that is required to generate (or regenerate) reports:

- Limit the number of contracts that are included in a report specification.
- Use the date selection features of the **`Date Filter`** tab in the Contract window to limit the amount of data that is processed.
- For customer reports, define only one report specification per customer.
- To avoid some regenerations when something changes within the contract, perform these steps:
  1. Copy the old contract to create a new contract.
  2. On the new contract, filter the dates so that the old data will not be present. Any new data will be included in the reports for the new contract.

     Since this contract is new, it would not be necessary to regenerate all reports when adding it to the report specification.
  3. Deactivate the old contract. The historical data for the old contract will persist in the reports.

Viewing the Service Level Reports

This topic describes how to view the service level reports that you generated. It explains the general layout of the reports and includes descriptions of the standard elements of service level reports:

How to Access Service Level Reports

After you have generated the service level reports, they are available at the output location that was defined in the report specification.

For viewing Contract Provider reports:

- If you generated the reports interactively,
  - and if you specified a valid URL for viewing reports in the report specification:
    - the reports will be displayed automatically in your Web browser.
  - and if you did *not* specify a valid URL for viewing reports in the report specification:
    - navigate to the **`index.html`** file in the **`text`** directory at the report output location that you specified and view the reports in your Web browser.
- If you generated the reports by running report job code in a SAS session or in a batch job,
  - navigate to the **`index.html`** file in the **`text`** directory at the report output location that you specified and view the reports in your Web browser.

For viewing Customer reports:

- If you generated the reports interactively or by running report job code in a SAS session or in a batch job,
  - navigate to the **`index.html`** file in the **`text`** directory at the report output location that you specified and view the reports in your Web browser.

The Banner

Display 11.11 Banner of the Service Level Reports Page

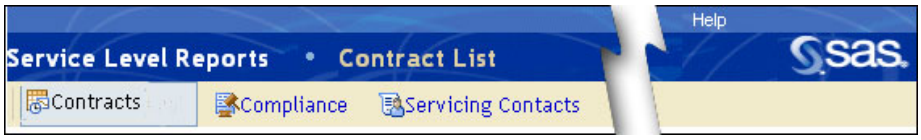

The banner includes the title bar (the two upper horizontal bars) and the tool bar that contains tabs for navigation. The title bar includes these entities:

- The **Help** link

  displays a Terminology page that includes a list of terms and a list of the icons that might be used in the service level reports.
- The **Service Level Reports** text string

  does not change among the various reports.
- The subtitle

  identifies the title of the report that is currently displayed.
- The SAS software logo

The tool bar displays a tab for each type of report that was generated. It can display up to three tabs: the **Contracts**, the **Compliance**, and the **Servicing Contacts** tabs. Select a tab to navigate to the corresponding reports.

The Breadcrumb Trail

The breadcrumb trail enables you to orient yourself within the various service level reports. You can click a link within the breadcrumb trail to jump to another page within the breadcrumb trail.

Display 11.12 The Breadcrumb Trail

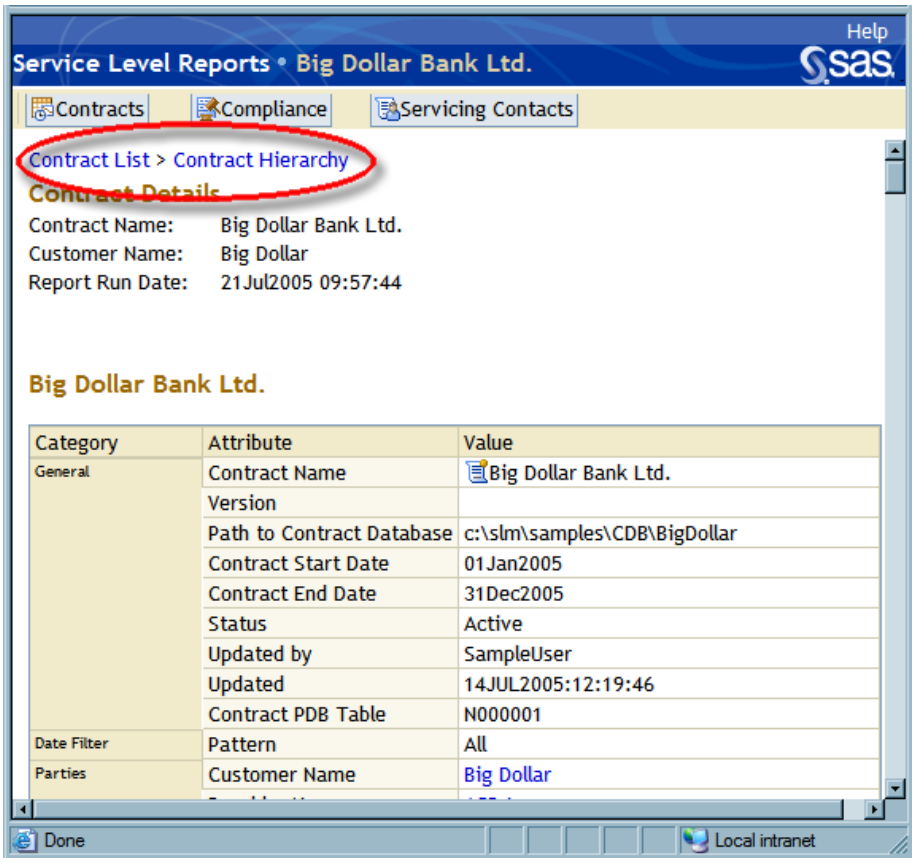

The Calendar

The service level reports that have date-based data include a calendar on the left side of the page. The **`Calendar`** provides a method by which you can select different dates to view.

Display 11.13 Parts of the Calendar

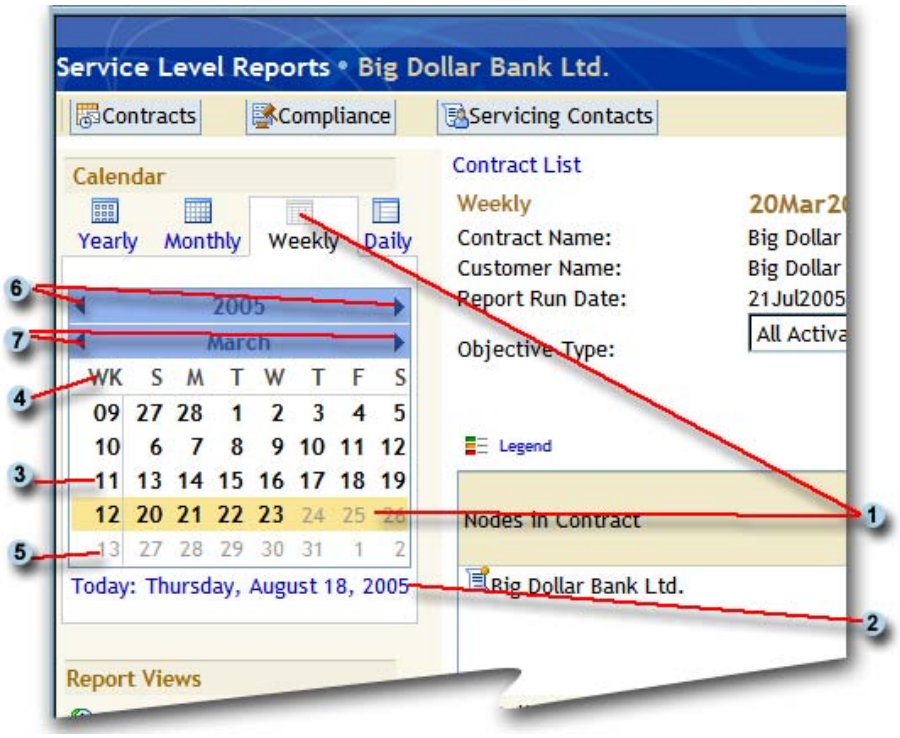

Table 11.6 Legend for the Parts of the Calendar

Number	Description
1	The selected time period for the report has a yellow background and matches the selected tab. Click a tab in the calendar to select a different time period.
2	Today's date is displayed. If today's date is visible on the calendar, then a red box is drawn around the date value and indicates today's date.
3	A date that displays in bold font means that a report is available for that date. For example, the week time period in bold means that the week is selectable and that reports for that week are available. Selecting a bold value on the calendar refreshes the report area with the appropriate view.
4	The first column displays the week number. If a week displays in a bold font, then the reports for that week are available. Select the week number to refresh the report area with the information for that week.
5	The dates that are gray are dates for which no report data exists.

Number	Description
6	The right arrow advances the calendar forward to the next year; the left arrow moves the calendar back to the previous year.
7	The right arrow advances the calendar forward to the next month; the left arrow moves the calendar back to the previous month.

The Report Views Panel

The service level reports that have date-based data display the **`Report Views`** panel on the left side of the page under the **`Calendar`**.

Display 11.14 The Report Views Panel

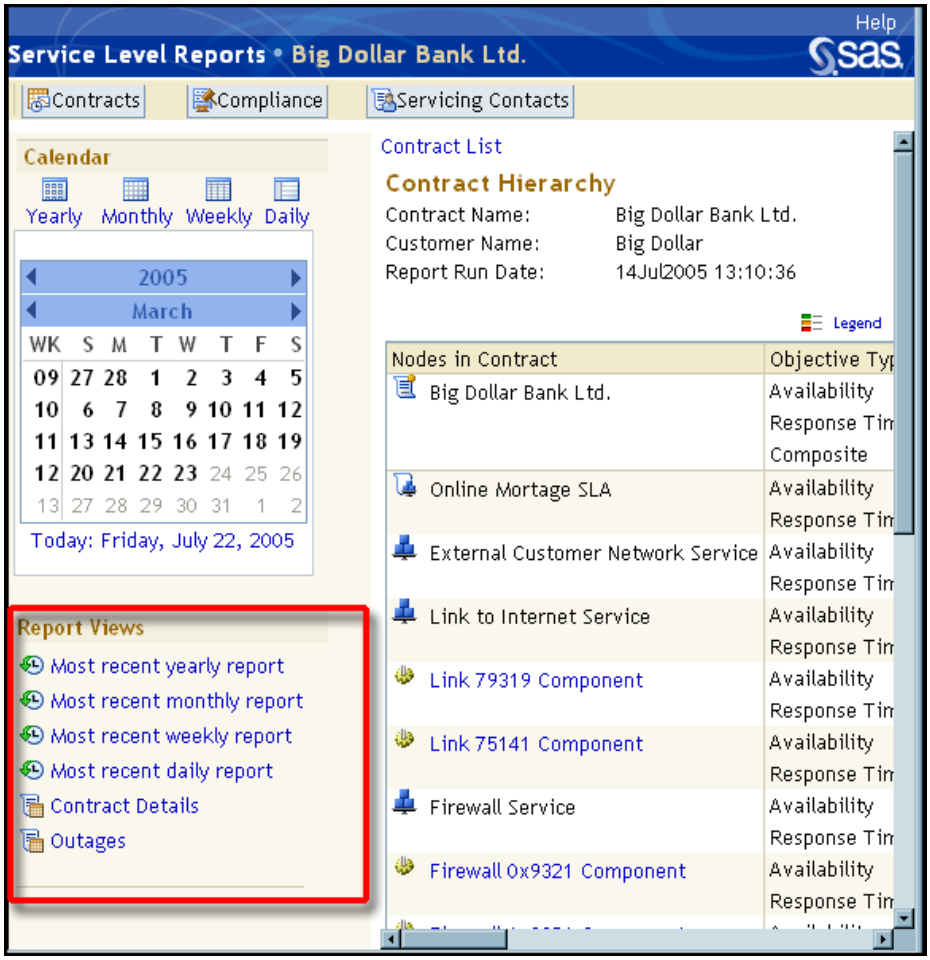

You can use the **`Report Views`** panel in the following ways:

- Select one of the **`Most recent`** links to view the report results of the most recent data that has been processed for the selected time period.
- Select **`Contract Details`** to view the Contract Details Report.
- Select **`Outages`** to view the Outages Report and to access the Outage Details Report for the selected contract.

Help for Terminology and Icons

The **`Help`** link is located in the banner of each service level report. Click **`Help`** to view a list of terms, their definitions, and the icons that are used in the reports.

Note: The list of icons is located below the Terminology list. Scroll down below the Terminology list to see a description of the icons that are used in the service level reports. △

Display 11.15 The Terminology List

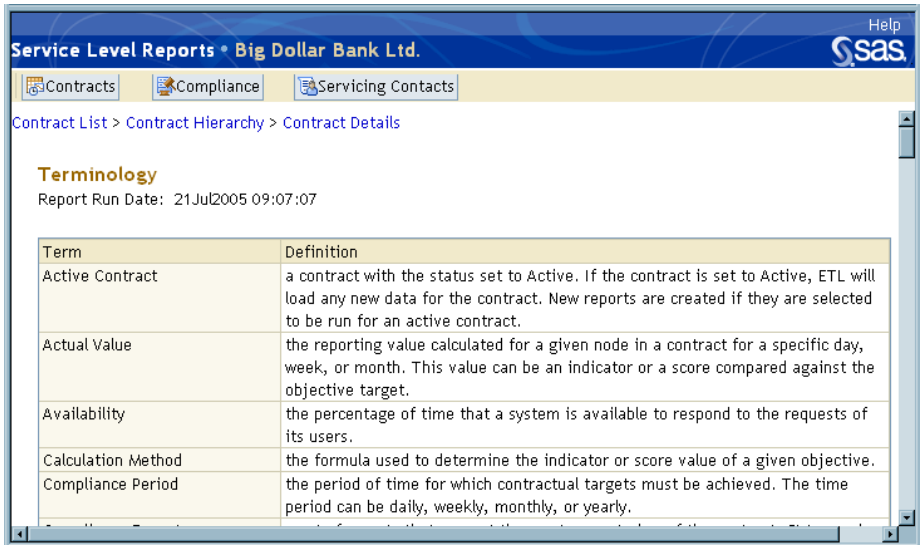

Display 11.16 Icons in the Terminology List

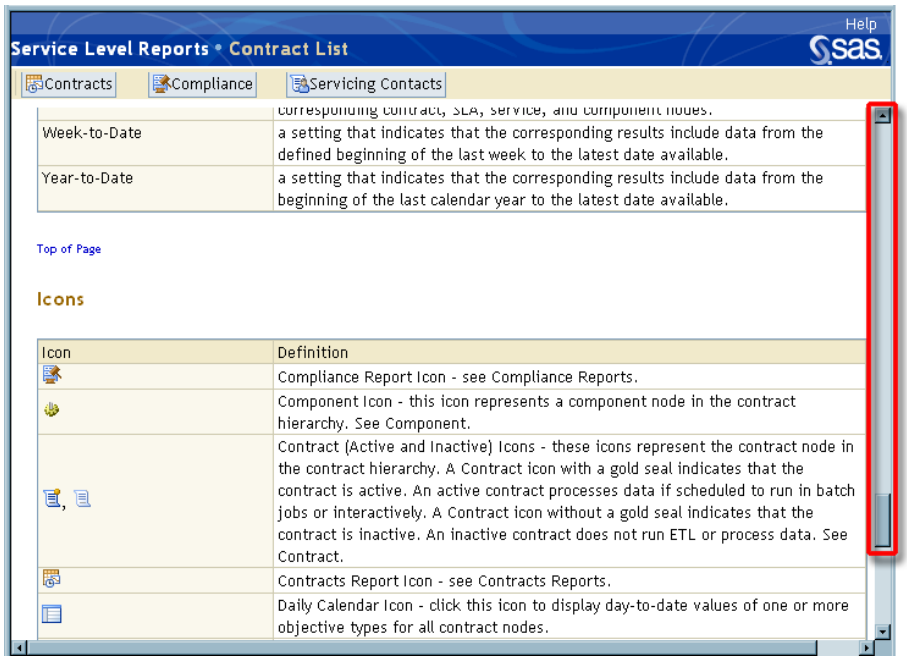

Troubleshooting

This section consists of the following topics:

- How to Allow Active Content Using Microsoft Internet Explorer
- How to Resize Fonts
- How to Print a Service Level Management Report

How to Allow Active Content Using Microsoft Internet Explorer

A security setting exists in Microsoft Internet Explorer that, if not enabled from a user's PC, might prevent a user from viewing the service level reports. If a user sees unexpected content or if a warning message appears when a user tries to access the service level reports in Internet Explorer, then the user should follow these steps:

1. In the Internet Explorer tool bar, select **Tools ▶ Internet Options**.
2. Select the **Advanced** tab.
3. In the Security section of the Settings list, make sure that the box beside **Allow active content to run in files on My Computer** is checked.
4. Click **OK**.

How to Resize Fonts

If you want to change the size of the fonts that are used for your reports, you can follow one or both of these procedures:

Internet Explorer

1. In the Internet Explorer menu bar, select **View ▶ Text Size**.
2. Select a text size that is appropriate. The text in the report immediately adjusts to the new size.

Microsoft Windows Desktop

1. Right-click on the background of your desktop.
2. In the list that appears, select **Properties**. The **Display Properties** dialog box appears.
3. Select the **Settings** tab. Move the slider indicator to the screen resolution that you want.
4. Click **OK**.

How to Print a Service Level Management Report

Although printing is not directly supported by SAS IT Service Level Management, you can use the print function of your Web browser to print a report. In Microsoft Internet Explorer, enable the **Print background colors and images** option before you print.

To enable the **Print background colors and images function**, perform these steps:

1. In the main menu bar of Internet Explorer, select **Tools ▶ Internet Options ▶ Advanced**.
2. In the **Advanced** dialog box, scroll to the **Printing** section.
3. Verify that the check box next to **Print background colors and images** is checked.
4. Click **OK**.

CHAPTER 12

Service Level Reports for Contract Providers and Customers

Introduction to Service Level Reports for Contract Providers and Customers

Service Level reports are helpful for both contract providers and customers. These two groups both need similar service level information; however, their requirements for report scope and variety differ based on their roles. For example, contract providers need reports that show service level data for all contracts and all customers for whom they provide services. Providers can use these reports to view the overall health of all contracts for all customers and can drill down into more detailed information about specific resources to troubleshoot the cause of missed targets and poor performance. Customers need reports that show service level data for their contracts only. Also, customers do not need specific information about the resources that support the contract because they typically are not responsible for troubleshooting problems with these resources.

SAS IT Service Level Management provides Service Level reports that can meet the needs of both contract providers and customers. Although many of these reports are similar for both groups, some report options and types differ slightly for contract providers and customers based on which group will use the generated reports. For example, contract providers have access to all report types with a full breadth of information about all contracts and customers. Customers have access to similar reports, yet the reports that are available to customers are limited in breadth. When generating Service Level reports, you will use report specifications to indicate if the set of Service Level reports that you are defining is for contract providers or customers.

Note: See "About Service Level Reports" on page 110 for more information about working with Service Level reports. △

About Service Level Reports for Contract Providers

Contract provider report specifications define and organize Service Level reports for contract providers. A set of reports that is generated from a contract provider report specification can include all available Service Level report types for one or more contracts. These comprehensive reports enable contract providers to view and analyze service level data for all customers, contracts, corresponding nodes, and resources. When creating and defining a contract provider report specification, you can also enable or disable various report types based on the provider's reporting needs.

Note: See "Service Level Reports for Contract Providers" on page 145 for a brief description, sample image, and navigation instructions for each Service Level report that is available to contract providers. △

About Service Level Reports for Customers

Customer report specifications define and organize Service Level reports for customers. A set of reports that is generated from a customer report specification can only include a limited set of Service Level report types for one or more contracts. When creating and defining a customer report specification, you can enable or disable some report types that are available to customers. Some Service Level report types and information are not available to customers because these report elements might contain sensitive information that could be inappropriate for customers. For example, Resource Analysis reports are not available to customers because these reports show performance data that is too detailed for a customer's needs.

Report specifications also enable you to designate a unique location for the report output. This capability enables you to assign customer-specific locations for report output and provide each customer with a link to their specific reports without allowing access to other reports for other customers or contracts.

Note: See "Service Level Reports for Customers" on page 168 for a brief description, sample image, and navigation instructions for each Service Level report that is available to customers. You can also copy this section about customer reports and provide it to your customers as a standalone document for reference. △

Note: See "Creating or Modifying the Report Specification" on page 123 for more information about creating contract provider report specifications and customer report specifications. △

Availability of Reports for Contract Providers and Customers

The report types that are available in a set of Service Level reports are determined by the type of report specification (contract provider or customer) that defines the reports and the optional reports that are enabled within the report specification. The following table shows the optional reports that you can enable in report specifications for contract providers and customers.

Table 12.1 Optional Reports Available to Enable in Report Specifications

Option	Can Enable in Contract Provider Report Specification?	Can Enable in Customer Report Specification?
Contracts Tab Reports	Always Enabled	Always Enabled
Compliance Tab Reports	Yes	Yes
Servicing Contacts Tab Reports	Yes	No
5-Minute Interval Table Reports	Yes	No

For example, Service Level reports for contract providers can have up to three tabs (Contracts, Compliance, and Servicing Contacts) that provide access to various reports. Service Level reports for customers can only have up to two tabs (Contracts and Compliance) that provide access to a limited number of reports.

For a comprehensive and comparative list of the Service Level reports that are available to contract providers and customers, see "Customer and Provider Reports" on page 115.

Service Level Reports for Contract Providers

This section provides detailed information, sample images, and navigation instructions for the following Service Level reports for contract providers:

- "Contract List Report" on page 146
- "Report Specification Details Report" on page 148
- "Notes Report" on page 148
- "Contract Hierarchy Report" on page 149

Contract List Report

The Contract List report is available for contract providers and customers from the **`Contracts`** tab. This report provides an overview of service levels for all contracts selected in the report specification. From this report, you can see an overview of how each selected contract performs at various summary levels such as daily, weekly, monthly, and yearly. You can also drill down into each summary level to explore more detailed reports about each contract. For example, if a data cell in the Contract List report is shaded in a color that indicates poor performance, you can drill down for more details about that data for each time period. The Contract List report also enables contract providers to sort data by customer name to help organize and sort the report information.

This report enables you to identify the overall health of the contract, discover problem areas by drilling down to more detailed reports, view the notes about a contract, and view the performance of a contract on a daily, weekly, monthly, year-by-week, or year-by-month basis.

The contract provider and customer versions of the Contract List report are identical except that the customer version of this report does not include a **`Sort by`** option to sort by customer or contract. The following is an example of a Contract List report for a contract provider:

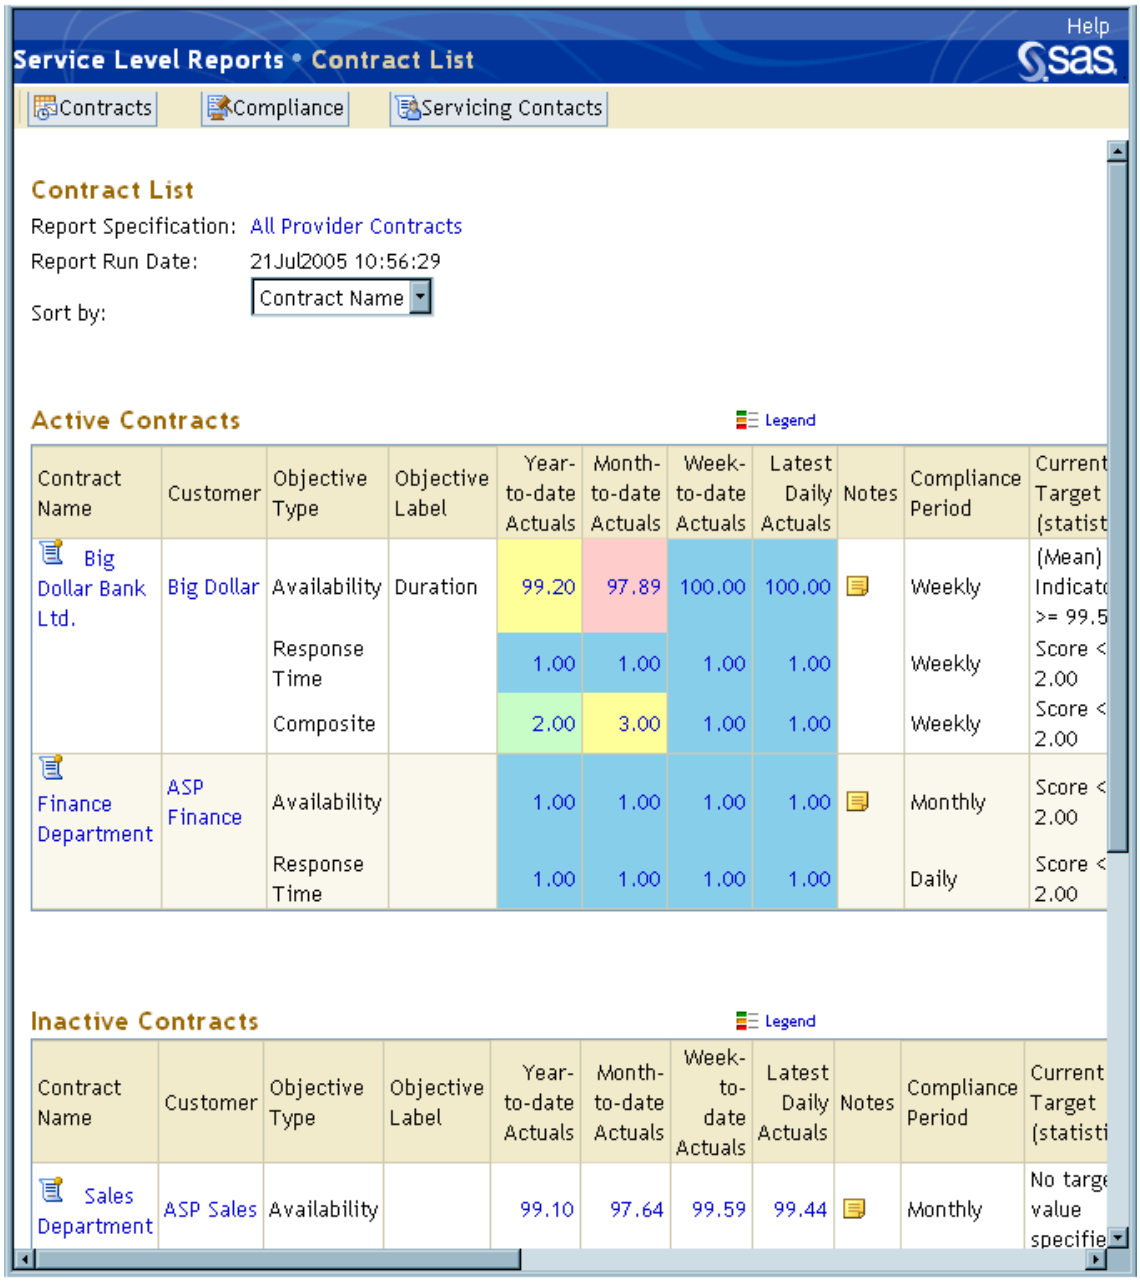

To view the Contract List report, click the **`Contracts`** tab.

From the Contract List report, you can also click various items or links to drill down to other reports such as the Report Specification Details reports and the Contract Hierarchy reports.

Report Specification Details Report

The Report Specification Details report is available for contract providers and customers. This report enables you to view the specifications that the SAS IT Service Level Management software uses to generate the reports. You can consult the values in this report to diagnose why certain elements appear (or are missing) in the generated reports.

The contract provider and customer versions of the Report Specification Details report are identical. The following is an example of a Report Specification Details report for a contract provider:

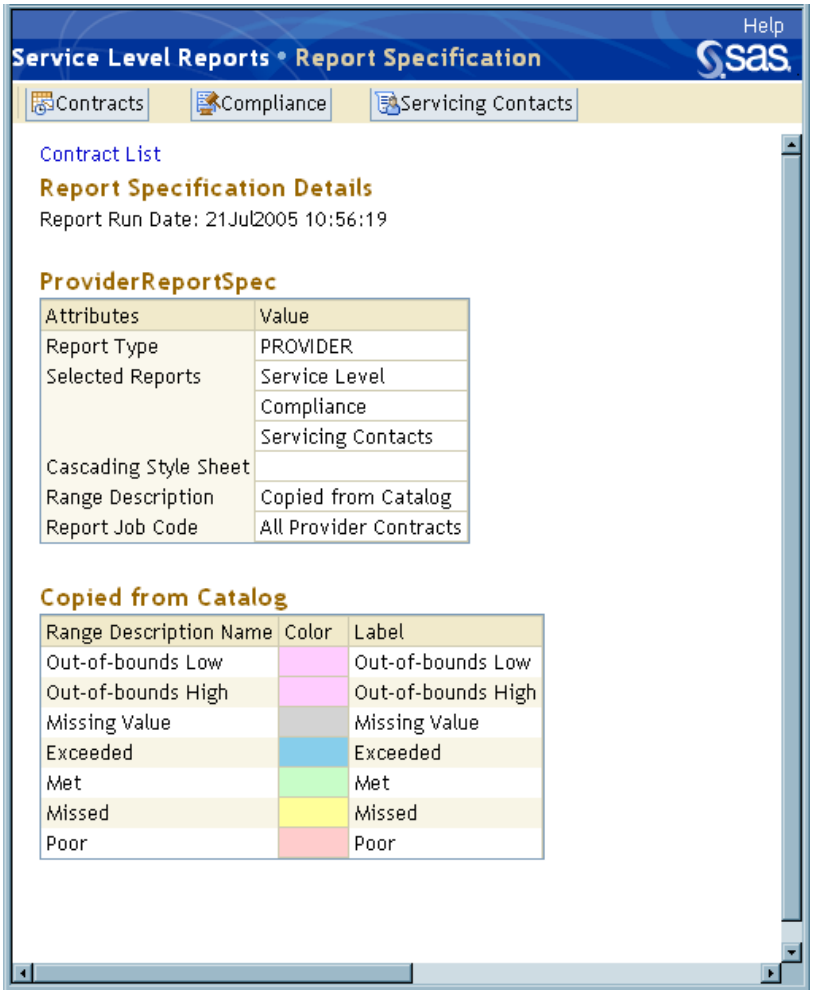

To view the Report Specification Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Click the report specification name next to **Report Specification** at the top of the Contract List report.

Notes Report

Notes reports enable contract providers and customers to view the details of a corresponding note or to navigate to a URL specified by an annotation. Contract

providers and customers can click a notes icon in a Service Level report to access the corresponding Notes report.

A notes icon appears in a Service Level report when an annotation, a note, or both exist for a corresponding contract, node, range set, outage, or contact for a specified time period. All notes will appear in the reports, but annotations can be filtered. Annotations can be excluded from reports or included for certain dates only.

The contract provider and customer versions of Notes reports are identical. The following is an example of a Notes report for a contract provider:

To view a Notes report, click a notes icon () in a Service Level report.

Contract Hierarchy Report

The Contract Hierarchy report is available for contract providers and customers. This report provides a comprehensive overview of how a specific contract is performing. The Contract Hierarchy report displays actual values and targets for each component, service, and SLA of a selected contract, including the data for the contract itself. This report enables you to view information across all summarization levels of the contract hierarchy and across all activated objectives for each node.

The contract provider and customer versions of the Contract Hierarchy report are identical except that customers cannot click on components to access the Resource Analysis reports. The following is an example of a Contract Hierarchy report for a contract provider:

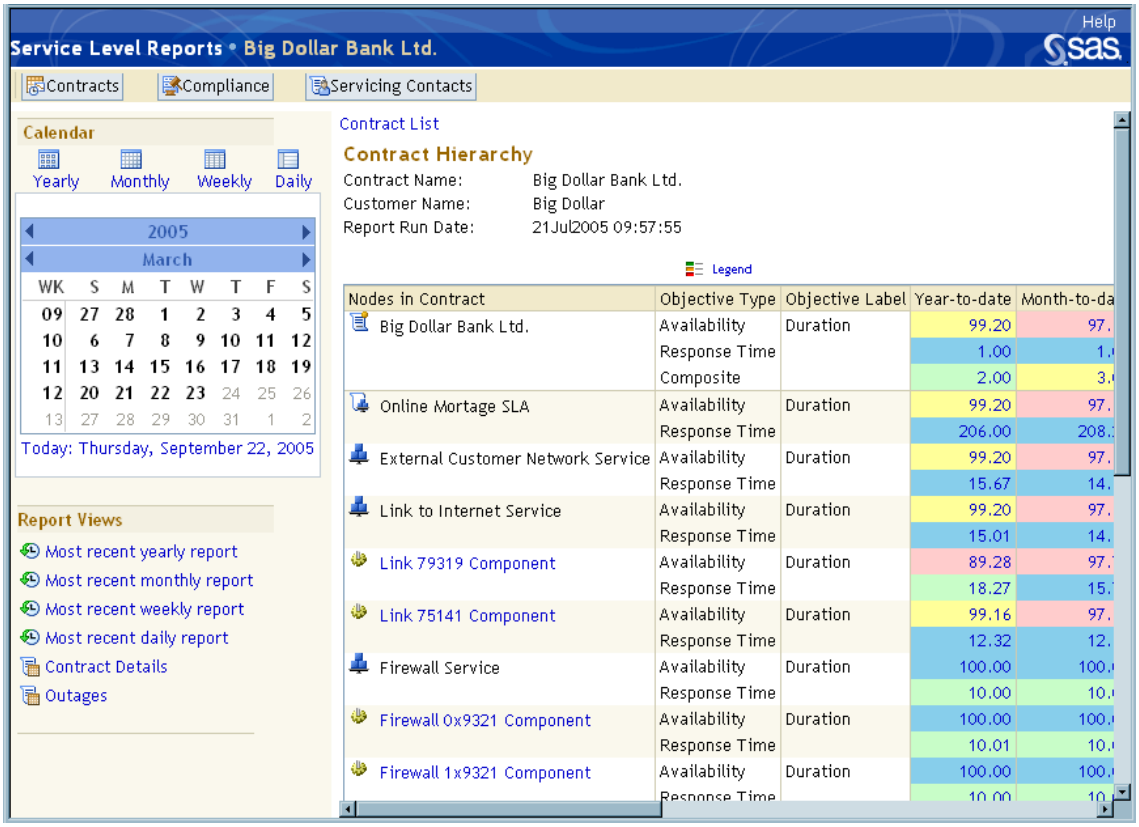

To view the Contract Hierarchy report, complete the following steps:

1 Click the **Contracts** tab.

2 Select a contract name in the table of the Contract List report.

From the Contract Hierarchy report, you can click various items and links to view other reports. These other reports can show data by various time periods (Daily, Weekly, Monthly, Year-by-Week, and Year-by-Month reports) or provide more detailed information about the contract (such as the Contract Details report and Outages report). Contract providers can also click a component in this report to access a Resource Analysis report that shows detailed information about the performance of the component's resources.

Daily, Weekly, Monthly, and Yearly Reports

Time-based reports such as the Daily, Weekly, Monthly, and Yearly reports are available for contract providers and customers. These reports display target and actual values, enabling you to view how all of the nodes of a contract are performing for a specified date. Each of these reports present similar information, differing only in the increments of time over which data is reported, as follows:

Daily reports display information about one or more objective types for all contract nodes for a selected day.

Weekly reports	display week-to-date values as well as the daily values for a selected week.
Monthly reports	display month-to-date values as well as the daily values for a selected month.
Year-by-Week reports	display the year-to-date values as well as the weekly values for a selected year.
Year-by-Month reports	display the year-to-date values as well as the monthly values for the selected year.

These time-based reports enable you to identify the objectives that are out of compliance and drill down to discover the specific dates and times during which targets were missed. These reports also provide several ways to navigate to different levels of summary. You can filter these reports by objective type, use the calendar links to select specific summary levels, and use the previous and next arrow icons to navigate to a different time period at the same summary level. See “The Calendar” on page 139 for more information about using the calendar tools to select a time-based report.

You can also click a graph icon to view a graphic report of the data for the corresponding node. See “Graphs in Time-Based Reports” on page 153 for more information about the graphs in the Daily, Weekly, Monthly, and Yearly reports for contract providers.

The contract provider and customer versions of the Daily, Weekly, Monthly, and Yearly reports are identical except that customers cannot click on components to access the Resource Analysis reports. The following is an example of a Daily report for contract providers:

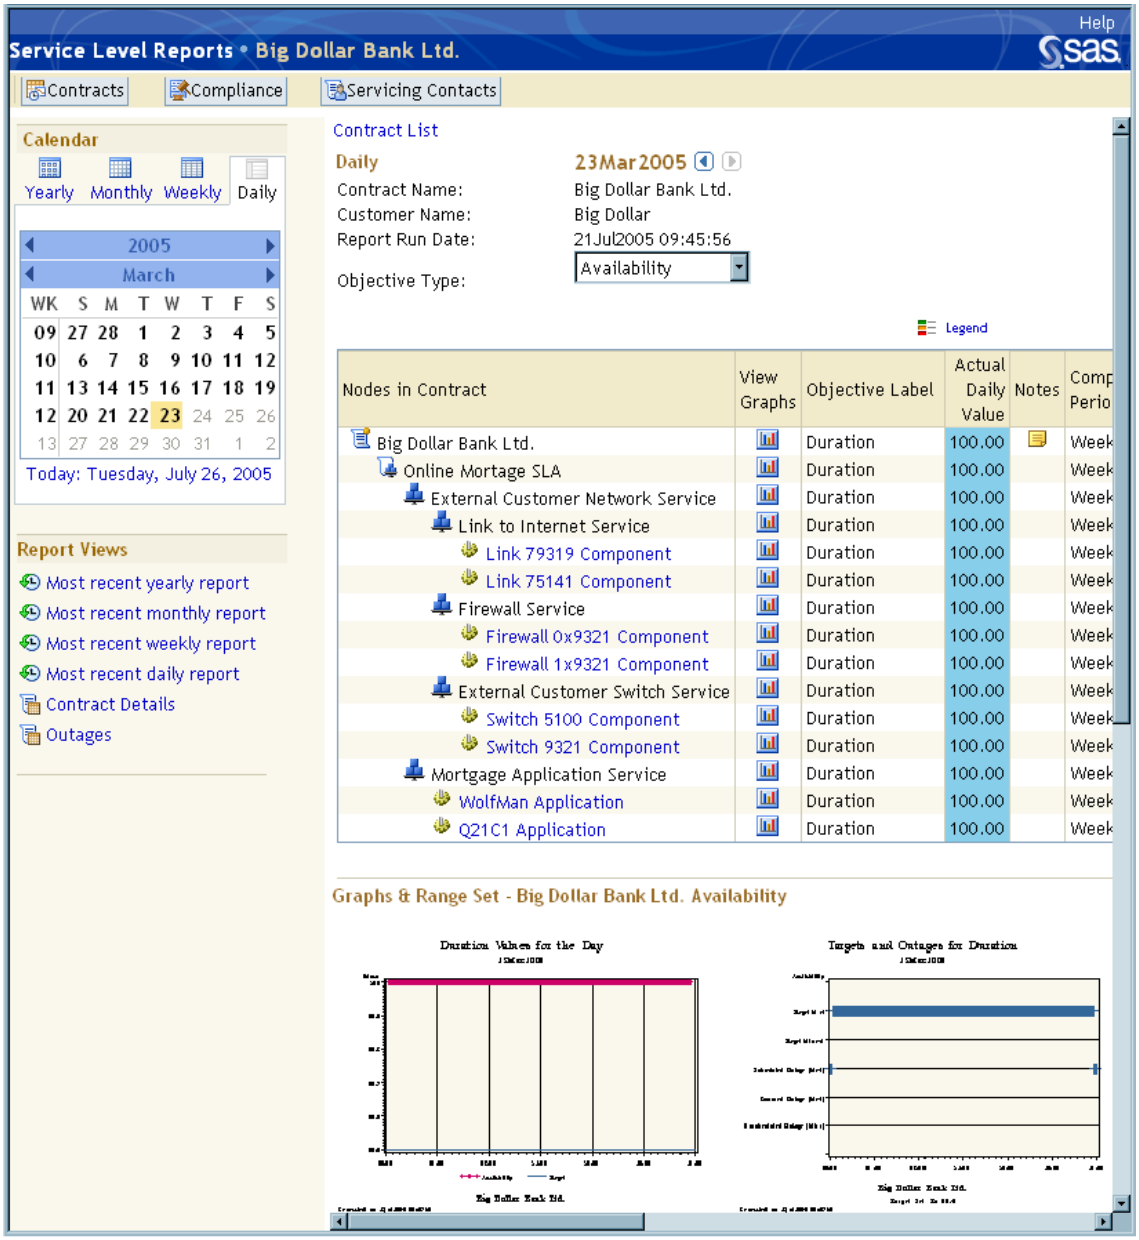

Time-based reports are accessible from both the Contract List report and the Contract Hierarchy report. The Contract List report and the Contract Hierarchy report provide different ways to access the same time-based information; the resulting time-based reports are simply displayed with different filters.

To view time-based reports from the Contract List report, complete the following steps:

1 Click the **`Contracts`** tab.

2 Select a value in the **Latest Daily Actual**, **Week-to-Date Actual**, **Month-to-date Actual**, or **Year-to-date Actual** columns.

To view time-based reports from the Contract Hierarchy report, complete the following steps:

1 Click the **Contracts** tab.

2 Select a contract name in the table of the Contract List report.

3 Choose one of the following methods to access the desired report:

- Use the calendar tools to select a time period.
- Below **Report Views** in the left panel, select **Most recent yearly report**, **Most recent monthly report**, **Most recent weekly report**, or **Most recent daily report**.
- Select a value in the **Latest Day**, **Week-to-date**, **Month-to-date**, or **Year-to-date** columns.

From the time-based reports, you can click various items or links to drill down to other reports such as the Contract Details report, the Outages report, and other time-based reports. Contract providers can also click a component to access a Resource Analysis report that shows detailed information about the performance of the component's resources.

Graphs in Time-Based Reports

The Daily, Weekly, Monthly, and Yearly reports for contract providers and customers can include up to two graphical reports for every active objective in a contract. The default graph (or graphs) that appears when you first access a time-based report represents the contract node. You can then click a graph icon in the **View Graphs** column to refresh the graph section of the report with graphs for the corresponding node.

Filtering a time-based report by objective type will also refresh the graph section of the report with the graphs for the selected objective type. However, the time-based reports do not show any graphs if the **Objective Type** filter is set to **All Activated**.

The graphs in each of the time-based reports display the data that best represents the appropriate time period. (If there is no data or only missing data for a time period, then a graph will not be generated for that time period.)

Daily Graphs	show values for every 5–minute observation during a 24–hour period of a selected date. This data is reported from the Detail level data in the contract database.
Weekly Graph	shows values for the 7 days in a selected week. The first day of the week is determined by the Start-of-Week value in the contract. This data is reported from the Day level data in the contract database.
Monthly Graph	shows values for all the days in a selected month. This data is reported from the Day level data in the contract database.
Year-by-Week Graph	shows values for all the weeks in a selected year. This data is reported from the Week level data in the contract database.
Year-by-Month Graph	shows values for all the months in the selected year. This data is reported from the Month level data in the contract database.

The Daily, Weekly, Monthly, and Yearly reports all contain a plot graph that shows actual values for the corresponding time period. If a target value was specified for a node, the plot graph will also contain a horizontal reference line that shows the target value for the node.

The following example is a plot graph of response time values from a Daily report. This sample graph shows that the Payroll Application consistently performed better (lower) than its response time target of 11000 required on March 23, 2005.

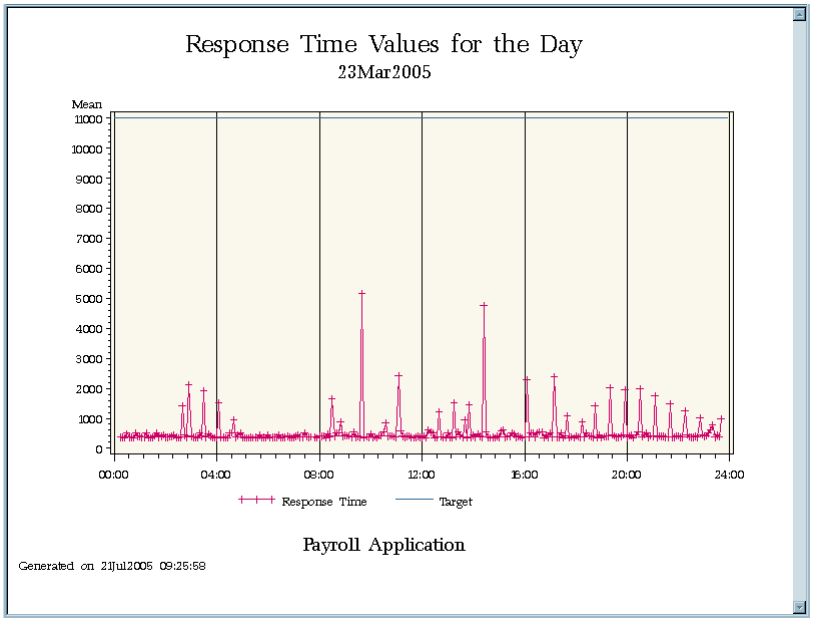

If a target value was specified for a node, then the Daily report also contains an additional plot graph that shows target and outage information for each node for the day. If a target value was specified for a node, the plot graph will contain tick marks that show if a target was met or missed. If an outage was entered for the node, then the plot will also contain tick marks that show what type of outage was entered for the node. A footnote will show the target value. This type of plot graph is only available on Daily reports.

The following example is a Targets and Outages graph for response time from a Daily report. This sample graph shows that the Payroll Application met its target for the entire 24–hour period on March 23, 2005, and experienced a scheduled outage at the beginning and end of the day.

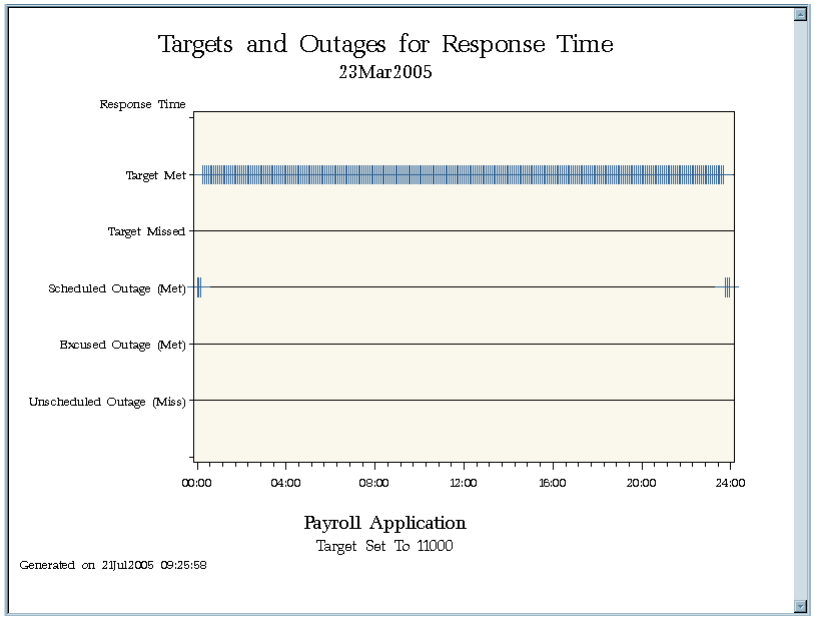

Three-Month Calendar View in Daily Reports

A three-month calendar view of the contract's daily values is available in Daily reports for contract providers and customers. This three-month calendar view is only available on Daily reports.

The three-month calendar view shows bold numbers to indicate that there is an actual value for that day. You can pause the mouse over a bold date to view the actual value for that day. In addition, the calendar cells will be color coded if a range set is specified for the period. You can use the right and left arrows to scroll to additional calendars.

The following example is a three-month calendar section of a Daily report. This sample shows an actual value of 1.00 for the response time on March 5, 2005.

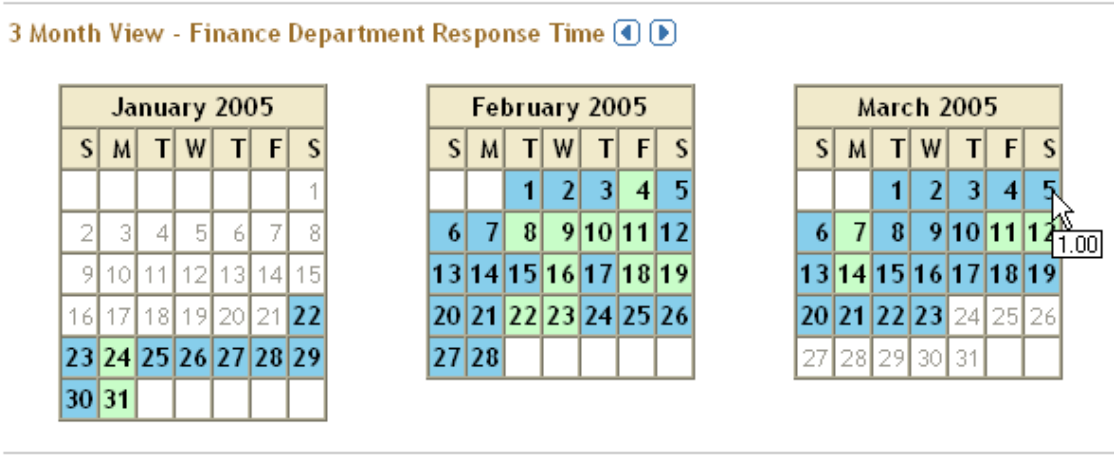

Resource Analysis Report

Resource Analysis reports are available to contract providers only. These reports enable contract providers to view the performance of all of the resources of a component on a daily basis. A flag icon in the **Potential Problem** column appears when the daily indicator value does not meet the target value that was set for the component.

Note: The flag icon warning in the **Potential Problem** column is not in effect in two cases: when there is no target for the corresponding component and when the corresponding component's actual value is score-based. In these instances, data comparisons cannot be made to determine if a flag icon is appropriate to indicate a problem. The flag icon warning will not be in effect in these cases and a flag icon will not appear in the **Potential Problem** column, although there might be a problem with the resource.

For example, if no target is specified, IT Service Level Management cannot determine if the resource value meets or does not meet expectations.

Also, if a component's actual value is score-based, IT Service Level Management cannot compare the component's score-based value with the resource value, because the resource value is always indicator-based, not score-based. In this case, the **Latest Daily Indicator** column will not display a value and the flag icon warning will not be in effect. △

In the Resource Analysis report, contract providers can filter the data by objective type and investigate the root cause of missed targets based on the poor performance of resources that support contract nodes. From this report, contract providers can also drill down to 5–Minute Interval tables that provide more information about resources.

Note: When you drill down to a Resource Analysis report from a component node, you may find that the **Resource Indicator** value in the Resource Analysis report differs from the **Actual Daily Value** in the corresponding Daily report. This difference in specific values might occur if you have entered an outage for the corresponding node.

Outages are entered on five-minute boundaries but actual outages can be slightly longer or shorter than a five-minute increment. Because *actual* outages are factored into the **Resource Indicator** values, and *entered* outages are factored into the **Actual**

Daily Value, a difference between the two values might occur. In these cases, you might choose to drill down into the 5-Minute Interval Table report to locate and identify the actual outages. See "What is an Outage?" on page 231 for more information about how outages can affect data that is reflected in Service Level reports. △

The following is an example of a Resource Analysis report for contract providers:

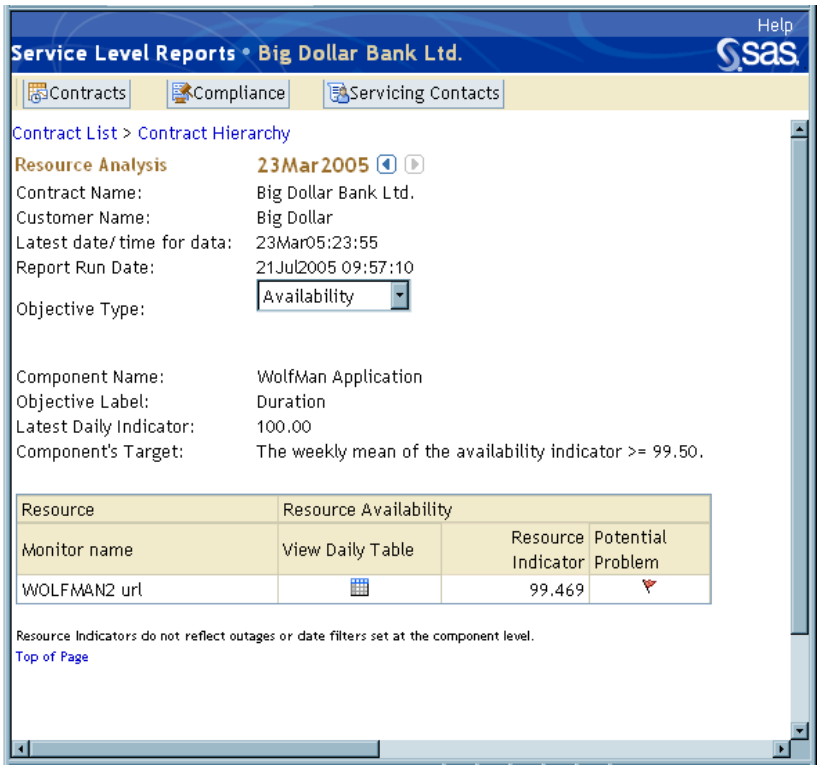

Resource Analysis reports are accessible to contract providers from reports in the **Contracts**, **Compliance**, and **Servicing Contacts** tabs. To view a Resource Analysis report, click a component name in a Service Level report.

For example, one way to access a Resource Analysis report is to complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the **Contract Name** column.
3. Select a component in the **Nodes in a Contract** column.

From the Resource Analysis report, contract providers can also click a table icon in the **View Daily Table** column to access the 5–Minute Interval Table report for the corresponding resource.

5–Minute Interval Table Report

The 5–Minute Interval Table reports are optional reports that are available to contract providers only. These reports detail numeric values from the resource data before they are interpreted with a calculation method for an objective. This information shows contract providers when (during the day) the performance of a resource complied or failed to comply with a component's target for an objective.

The following is an example of a 5–Minute Interval Table report for contract providers:

Resource Duration
5 Minute Interval Table
WolfMan Application

Date	Time	Indicator	MONITOR
23MAR2005	0:00:00	20.666666667	WOLFMAN2 url
23MAR2005	0:05:00	100	WOLFMAN2 url
23MAR2005	0:10:00	100	WOLFMAN2 url
23MAR2005	0:15:00	100	WOLFMAN2 url
23MAR2005	0:20:00	100	WOLFMAN2 url
23MAR2005	0:25:00	100	WOLFMAN2 url
23MAR2005	0:30:00	100	WOLFMAN2 url
23MAR2005	0:35:00	100	WOLFMAN2 url
23MAR2005	0:40:00	100	WOLFMAN2 url
23MAR2005	0:45:00	100	WOLFMAN2 url
23MAR2005	0:50:00	100	WOLFMAN2 url
23MAR2005	0:55:00	100	WOLFMAN2 url
23MAR2005	1:00:00	100	WOLFMAN2 url

To view a 5–Minute Interval Table report, click a table icon in the **View Daily Table** column of any Resource Analysis report. For example, one way to access a 5–Minute Interval Table report is to complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the **Contract Name** column.
3. Select a component in the **Nodes in a Contract** column.
4. Click a table icon in the **View Daily Table** column.

Contract Details Report

The Contract Details report is available to contract providers and customers. This report enables you to view all the settings that are associated with the contract, including detailed information about the customer and provider, calculation methods, range sets, and date filter information used for activated objectives of a specified contract.

The Contract Details report for customers does not include as much comprehensive and potentially sensitive information as the contract provider report. When a set of reports is generated for customers via a customer report specification, the Contract Details report will not include the following information fields: **Version**, **Path to**

Contract Database, **Updated by**, **Updated**, **Contract PDB Table**, **Cost Info**, **Term Info**, **Contract Database Info**, **Penalty Info**, and **Credit Info**.

The following is an example of a Contract Details report for contract providers:

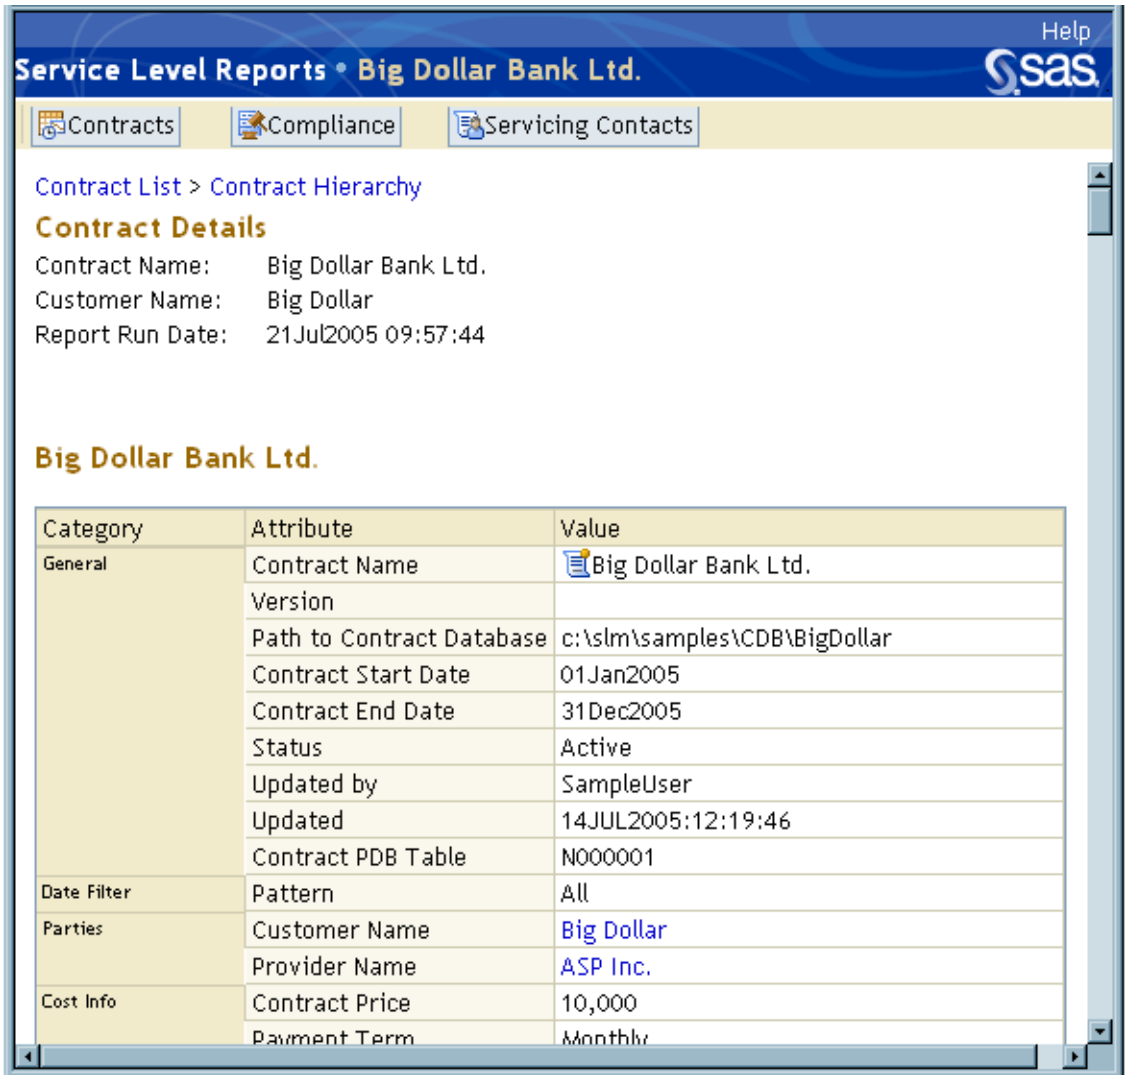

To view the Contract Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the Contract List report.
3. Below **Report Views** in the left panel, click **Contract Details**.

From the Contract Details report, you can also click various items or links to drill down for more information about the terms of the contract. Some of the other reports you can access from this report include Customer reports, Provider reports, and Range Set Details reports.

Customer Report

The Customer report is available for contract providers and customers. This report enables you to view the customer and contact information for a selected contract.

The contract provider and customer versions of the Customer report are identical. The following is an example of a Customer report for contract providers:

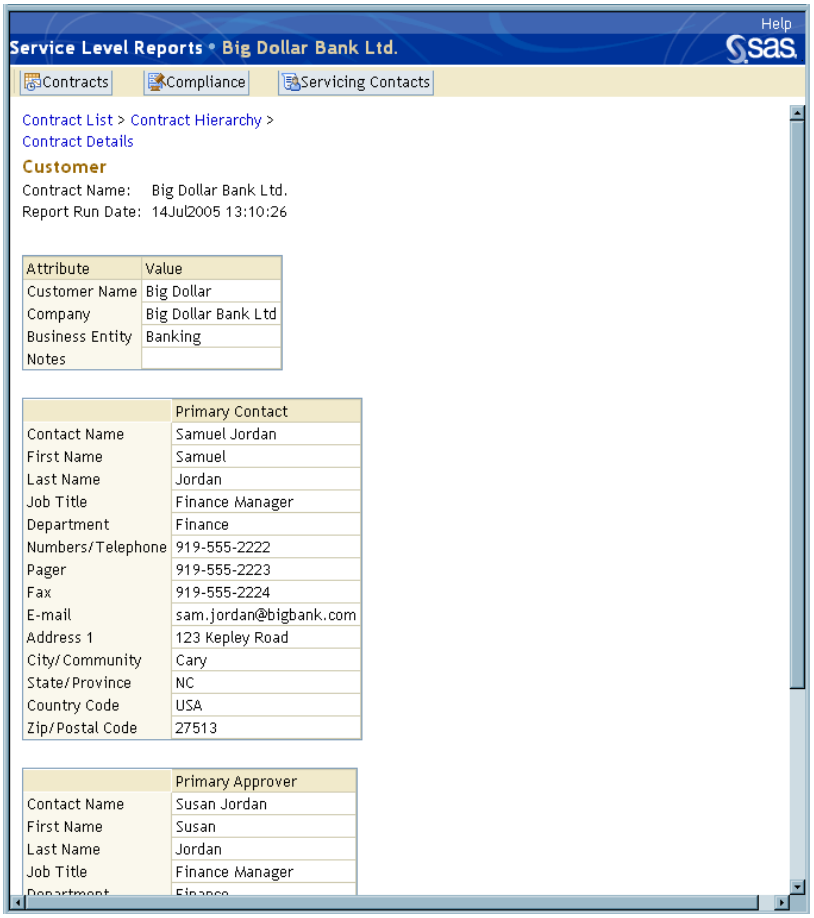

To view a Customer report, you can complete the following steps:

1. Click the **Contracts** tab.
2. Select the name of a customer in a **Customer** column of the Contract List report.

Provider Report

The Provider report is available for contract providers and customers. This report displays information about the provider for a selected contract. The layout of the Provider report is identical to that of the Customer report.

The contract provider and customer versions of the Provider report are identical. The following is an example of a Provider report for contract providers:

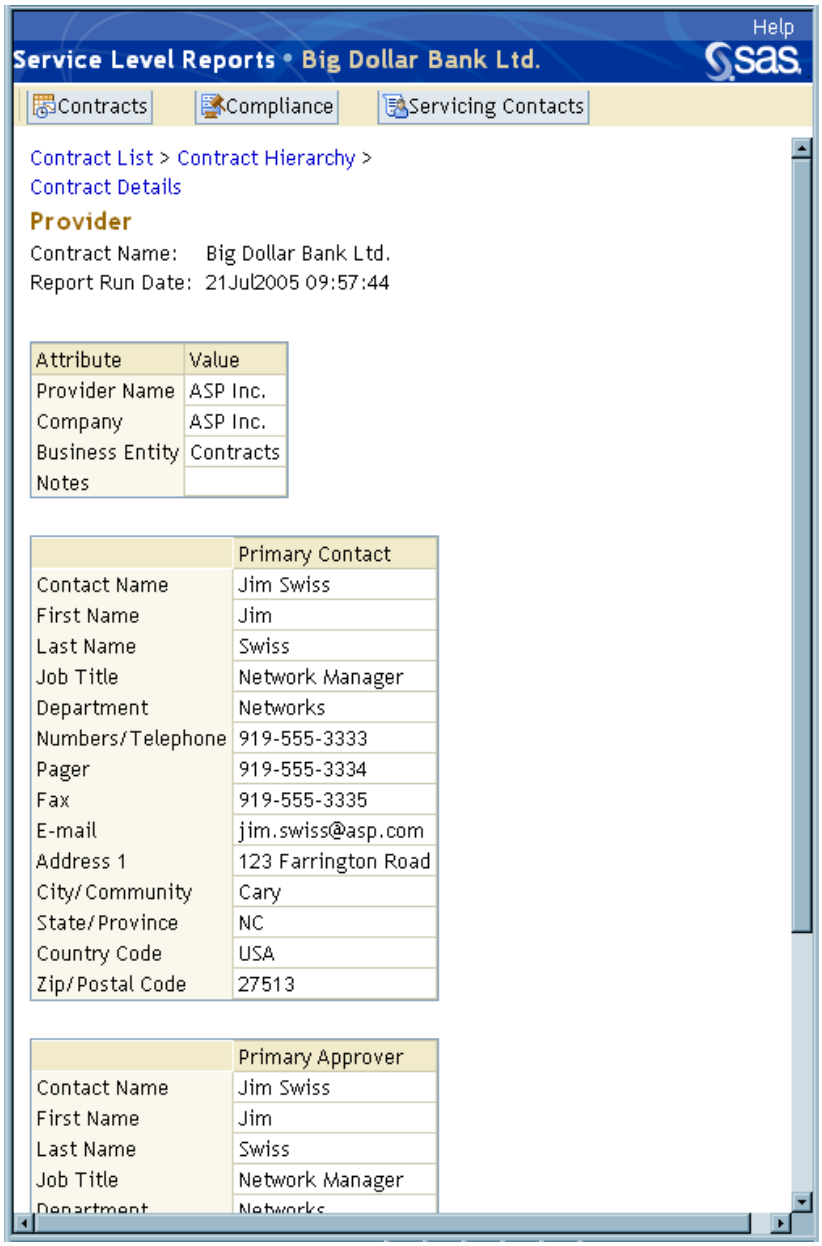

To view a Provider report, complete the following steps:

1 Click the **`Contracts`** tab.

2 Select a contract in a **`Contract Name`** column of the Contract List report.

3 Below **Report Views** in the left panel, select **Contract Details**.

4 Select a provider name in the **Provider Name** row of the Contract Details report.

Range Set Details Report

The Range Set Details report is available for contract providers and customers. This report enables you to see all of the details for the range set that is associated with a selected objective in a contract. This report displays the ranges that make up the range set as well as their range descriptions so that you can review the colors and boundary conditions that are assigned to each range.

The contract provider and customer versions of the Range Set Details report are identical. The following is an example of a Range Set Details report for contract providers:

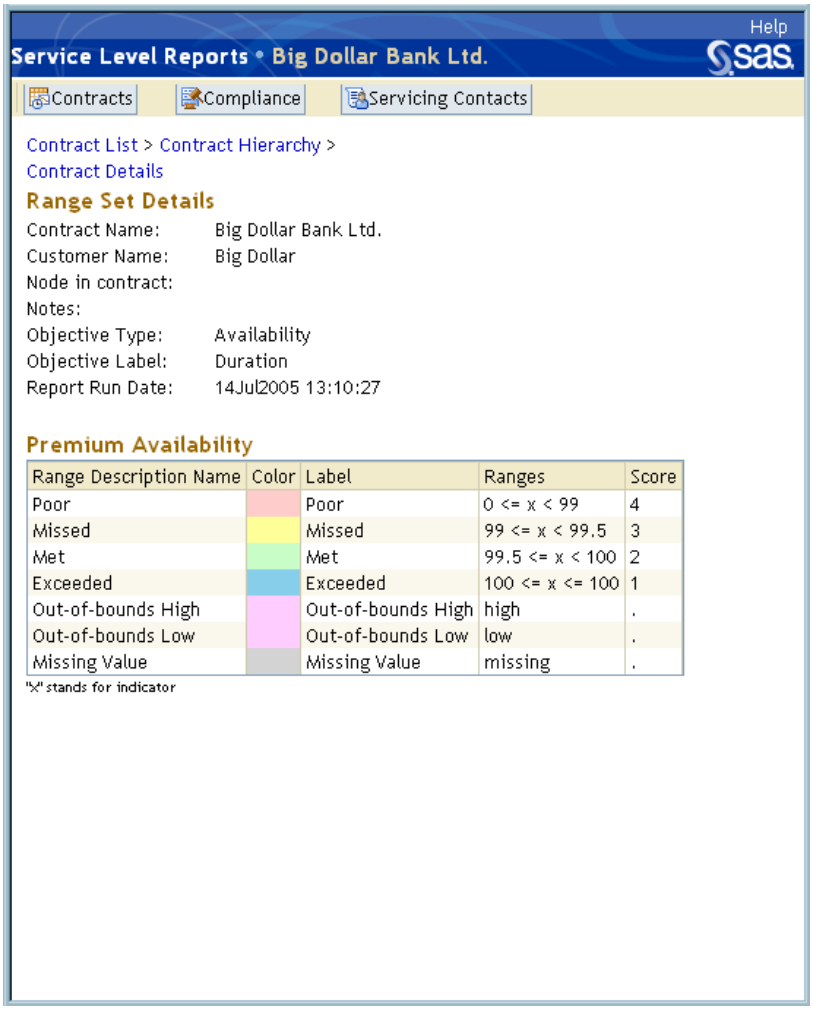

To view a Range Set Details report, complete the following steps:

1 Click the **Contracts** tab.

2 Select a contract in a **Contract Name** column of the Contract List report.

3 Below **Report Views** in the left panel, click **Contract Details**.

4 Select a range set name in a **Range Set** column of the Contract Details report.

Outages Report

The Outages report is available for contract providers and customers. This report enables you to view all of the outages for every node (and its child nodes) in a selected contract. Contract providers might use this report to view all the outages that have been entered for their nodes.

The contract provider and customer versions of the Outages report are identical. The following is an example of an Outages report for contract providers:

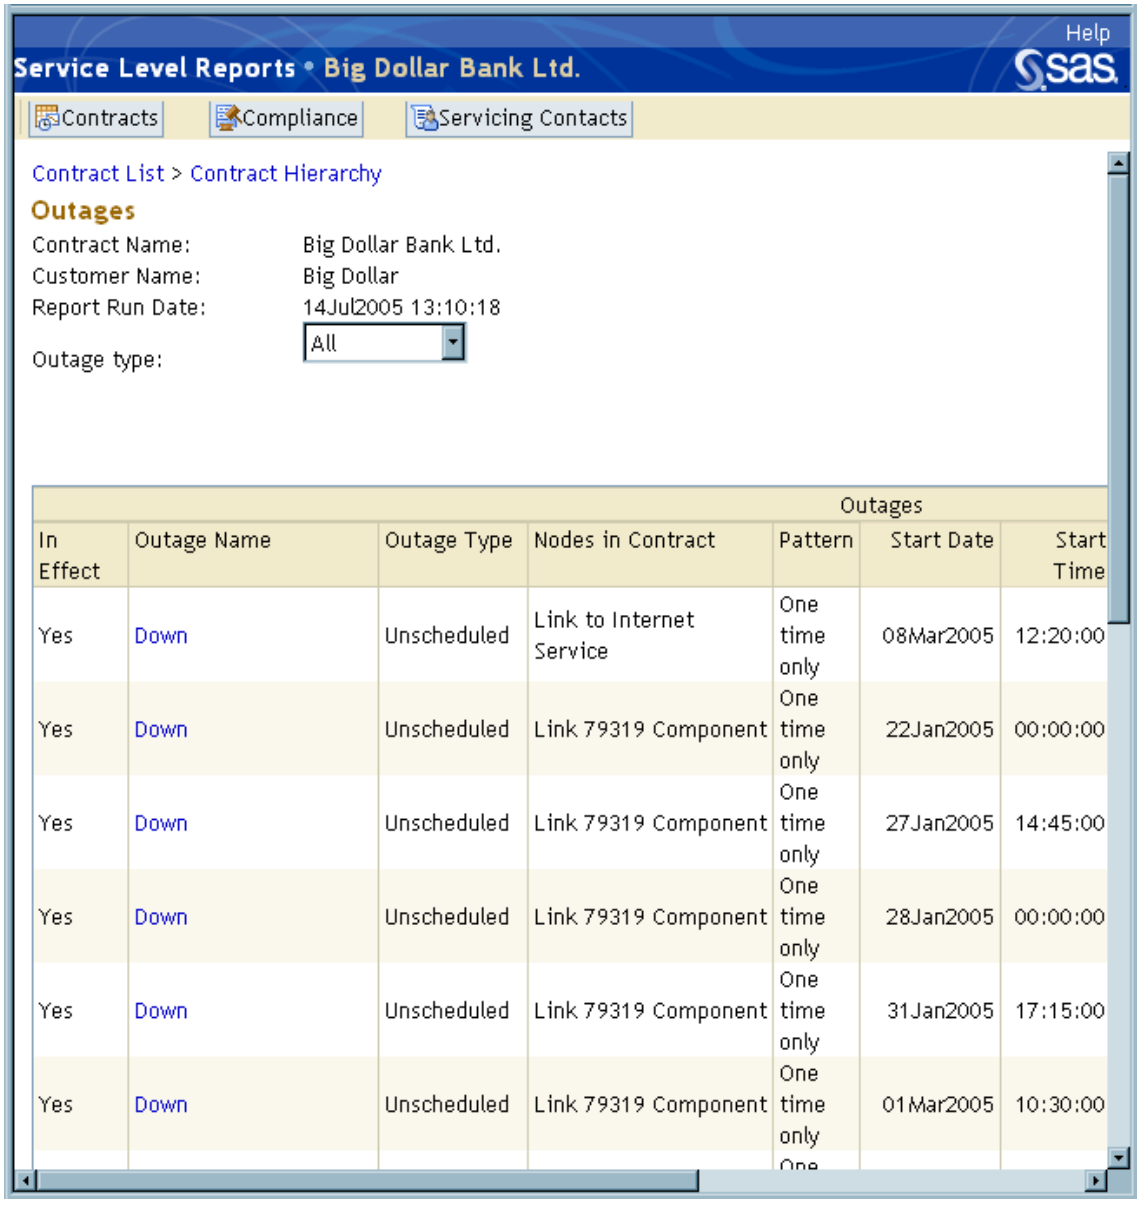

To view an Outages report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the **Contract Name** column.
3. Below **Report Views** in the left panel, click **Outages**.

From the Outages report, you can also select an outage name to access the Outage Details report for more information about the outage.

Outage Details Report

The Outage Details report is available for contract providers and customers. This report enables you to view all of the details, including the exceptions, for a selected outage. This information includes a display of the date and reason for all the exceptions to the outage.

The contract provider and customer versions of the Outage Details report are identical. The following is an example of an Outage Details report for contract providers:

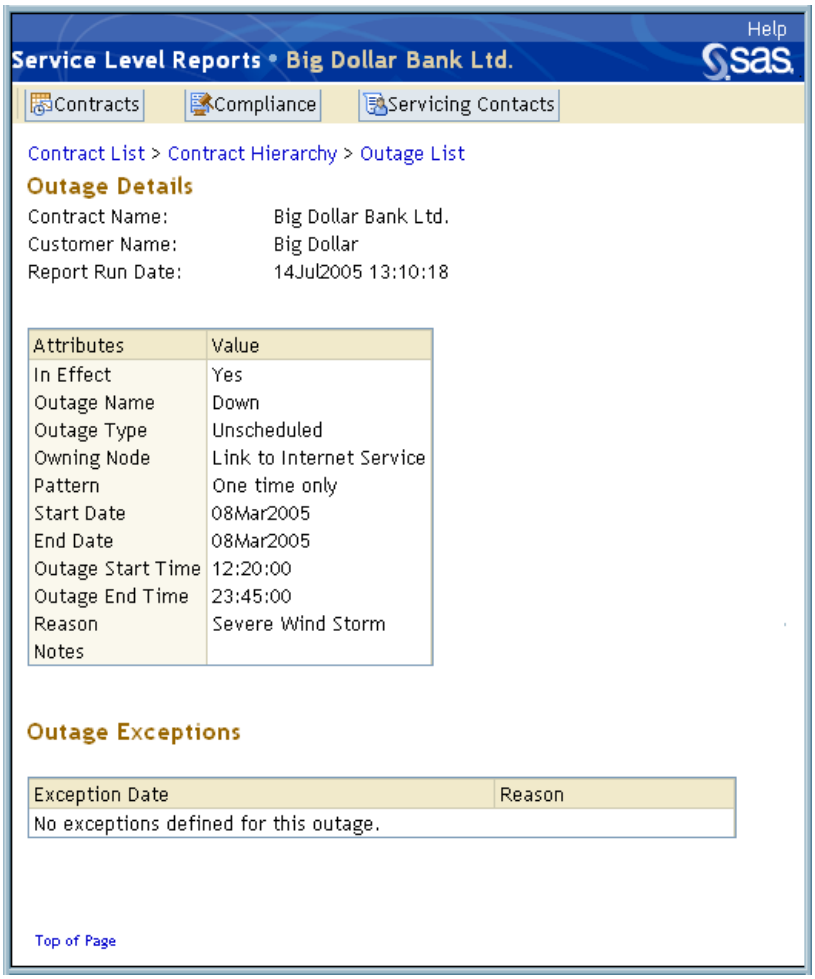

To view an Outage Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the **Contract Name** column.
3. Below **Report Views** in the left panel, click **Outages**.
4. Select an outage in the **Outage Name** column of the Outages report.

Compliance List Report

The Compliance List report is an optional report that is available for contract providers and customers from the **Compliance** tab. Because the Compliance List report is optional, based on the report specification settings, a set of Service Level reports for contract providers or customers might or might not include the **Compliance** tab and its corresponding reports.

The Compliance List report enables you to determine if contracts are meeting their targets during the compliance period for each active objective. You can also filter this report by customer, all targets, missed targets for the current period, and missed targets for the previous period.

The contract provider and customer versions of the Compliance List report are identical. The following is an example of a Compliance List report for contract providers:

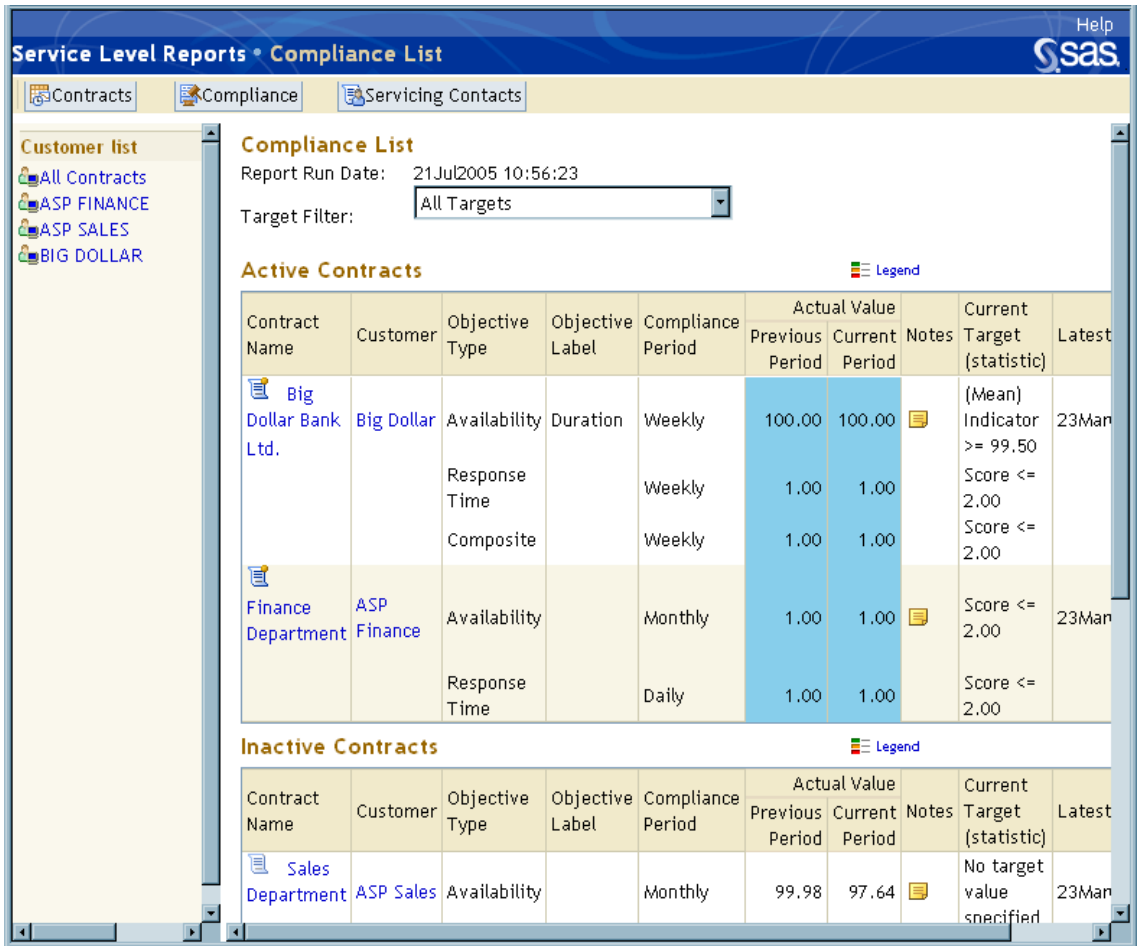

To view the Compliance List report, click the **Compliance** tab.

From the Compliance List report, you can also click various items or links to drill down to other reports such as the Compliance Details reports and the Customer reports.

Compliance Details Report

The Compliance Details report is available for contract providers and customers from the Compliance List report on the **Compliance** tab. Because the Compliance List report is optional, based on the report specification settings, a set of Service Level reports for contract providers or customers might or might not include the **Compliance** tab and its corresponding reports.

The Compliance Details report for contract providers enables you to determine if nodes within a specific contract are meeting their targets during the compliance period for each active objective. This report also enables contract providers to view the details for the individual components in order to uncover any problems with resources. You can filter this report by contract, all targets, missed targets for the current period, and missed targets for the previous period.

The contract provider and customer versions of the Compliance Details report are identical except that customers cannot click on components to drill down to Resource Analysis reports. The following is an example of a Compliance Details report for contract providers:

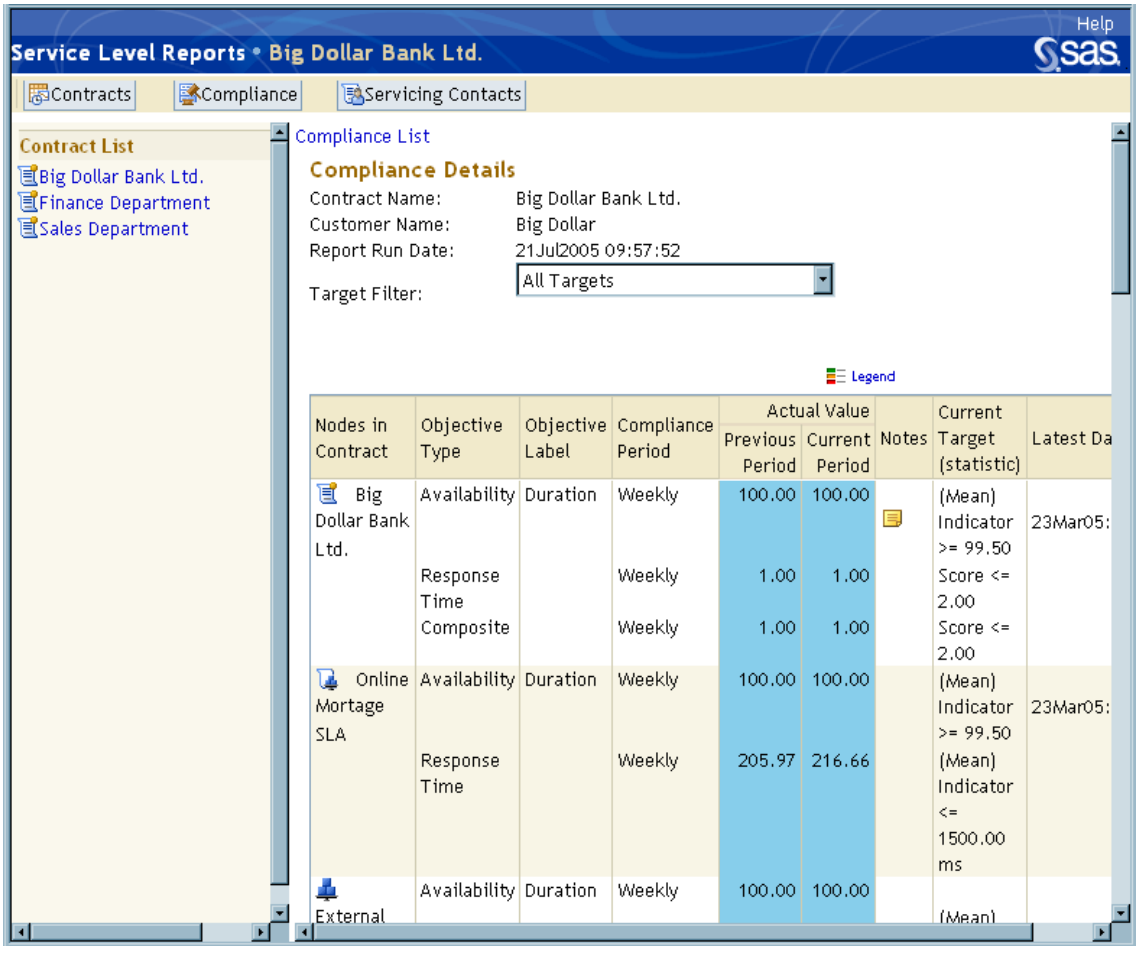

To view the Compliance Details report, complete the following steps:

1 Click the **Compliance** tab.

2 Select a contract in a **Contract Name** column of the Compliance List report.

From the Compliance Details report, contract providers can also drill down to view Resource Analysis reports to see resource details.

Servicing Contacts Report

The Servicing Contacts report is an optional report that is available to contract providers only. Contract providers can access the Servicing Contacts report from the **`Servicing Contacts`** tab.

This report lists servicing contacts and the compliance data of the contracts for which the servicing contacts have responsibility. The data in this report provides a snapshot of the current and previous compliance period, enabling you to view the compliance information for contracts that are serviced by a selected contact person or business entity. Contract providers can also filter this report by contact, all targets, missed targets for the current period, and missed targets for the previous period.

The following is an example of a Servicing Contacts report for contract providers:

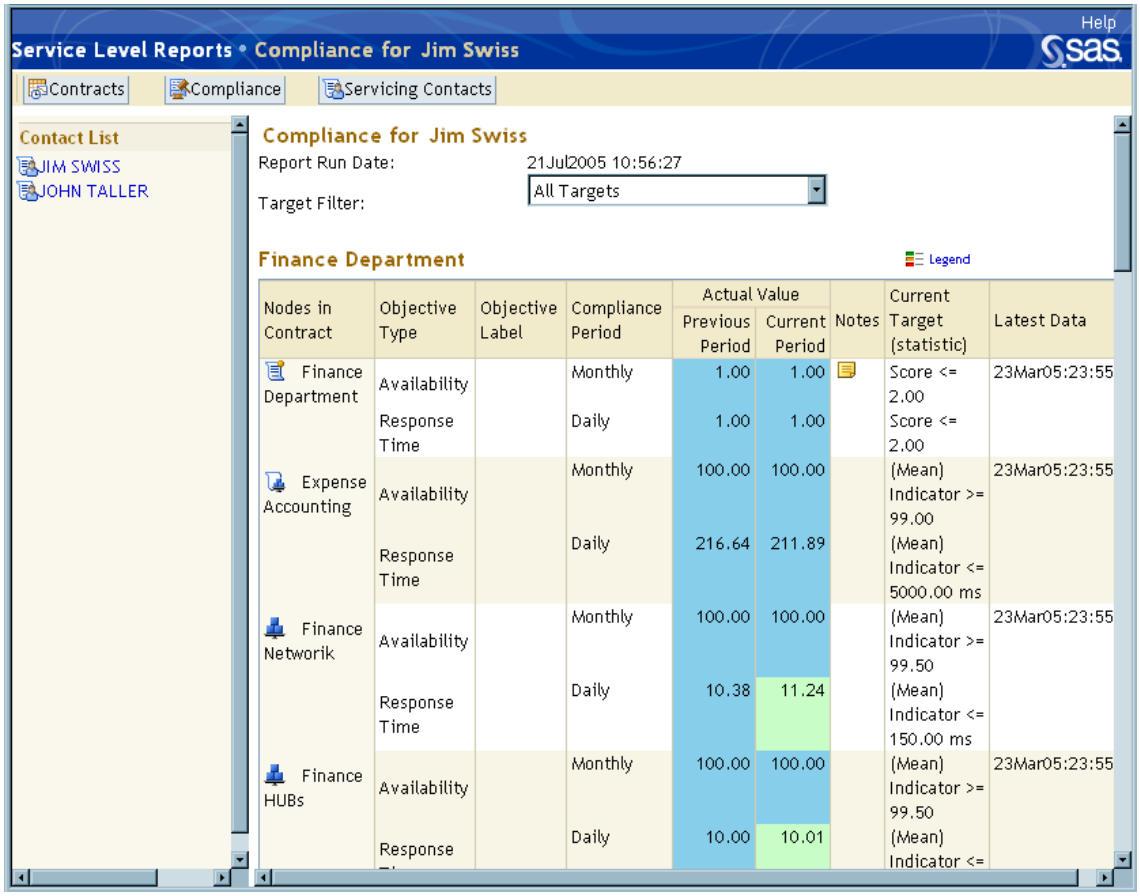

To view the Servicing Contacts report, click the **`Servicing Contacts`** tab.

From the Servicing Contacts report, contract providers can also select components to view Resource Analysis reports for resource details.

Service Level Reports for Customers

This section provides detailed information, sample images, and navigation instructions for the following Service Level reports for customers:

Contract List Report

The Contract List report is available for customers from the **Contracts** tab. This report provides an overview of service levels for selected customer contracts. From this report, you can see an overview of how each selected contract performs at various summary levels such as daily, weekly, monthly, and yearly. You can also drill down into each summary level to explore more detailed reports about each contract. For example, if a data cell in the Contract List report is shaded in a color that indicates poor performance, you can drill down for more details about that data for each time period.

This report enables you to identify the overall health of the contract, discover problem areas by drilling down to more detailed reports, view the notes about a contract, and view the performance of a contract on a daily, weekly, monthly, year-by-week, or year-by-month basis.

The following is an example of a Contract List report for a customer:

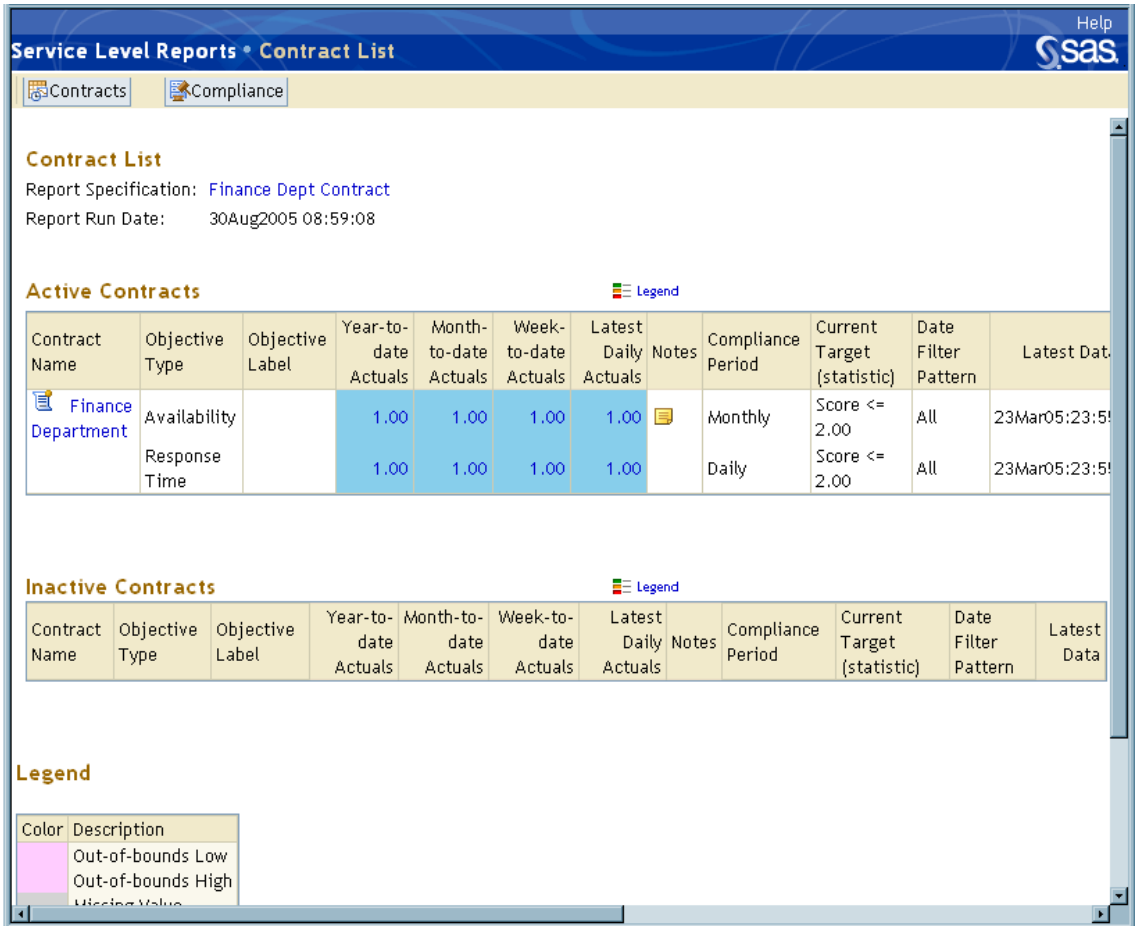

To view the Contract List report, click the **`Contracts`** tab.

From the Contract List report, you can click various items or links to drill down to other reports such as the Report Specification Details reports and the Contract Hierarchy reports.

Report Specification Details Report

The Report Specification Details report for customers enables you to view the specifications that the SAS IT Service Level Management software uses to generate the reports. You can consult the values in this report to diagnose why certain elements appear (or are missing) in the generated reports.

The following is an example of a Report Specification Details report for customers:

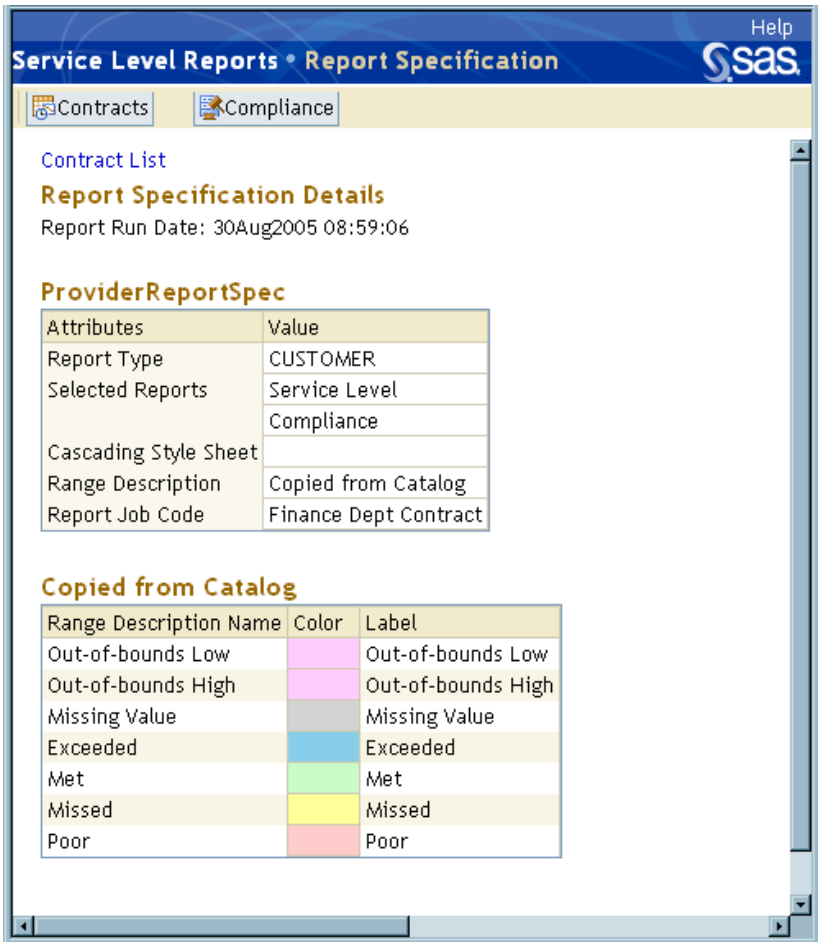

To view the Report Specification Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Click the report specification name next to **Report Specification** at the top of the Contract List report.

Notes Report

Notes reports enable customers to view the details of a corresponding note or to navigate to a URL specified by an annotation. You can click a notes icon in a Service Level report to access the corresponding Notes report.

A notes icon appears in a Service Level report when an annotation, a note, or both exist for a corresponding contract, node, range set, outage, or contact for a specified time period.

The following is an example of a Notes report for a customer:

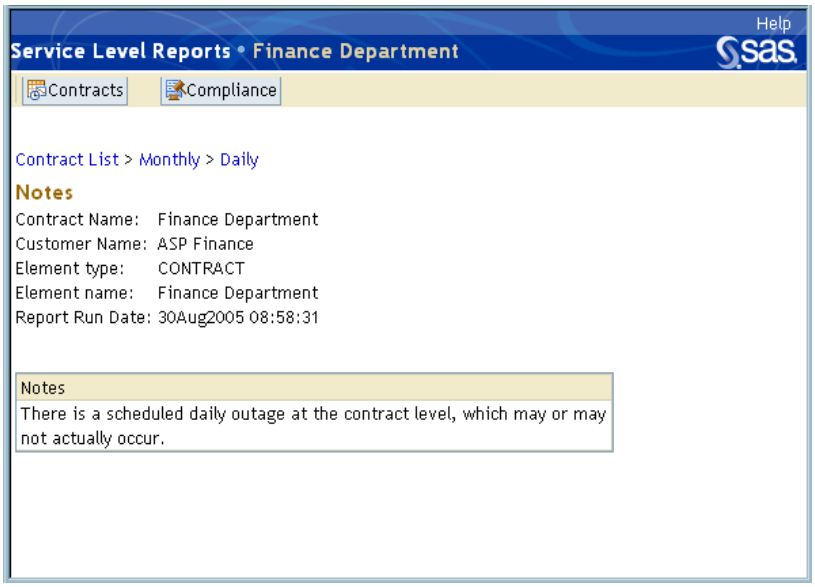

To view a Notes report, click a notes icon () in a Service Level report.

Contract Hierarchy Report

The Contract Hierarchy report for customers provides a comprehensive overview of how a specific contract is performing. This report displays actual values and targets for each component, service, and SLA of a selected contract, including the data for the contract itself. This report enables you to view information across all summarization levels of the contract hierarchy and across all activated objectives for each node of the contract.

The following is an example of a Contract Hierarchy report for a customer:

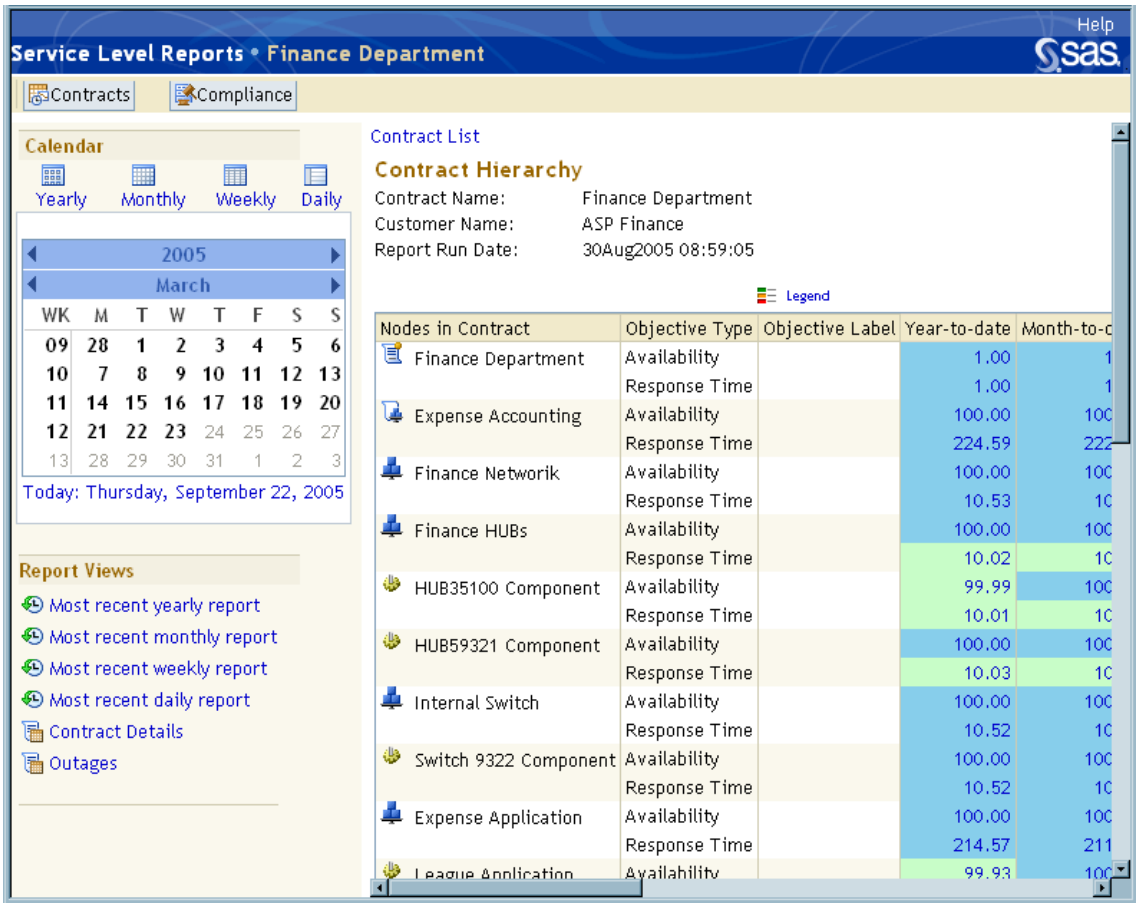

To view the Contract Hierarchy report, complete the following steps:

1 Click the **Contracts** tab.

2 Select a contract name in the table of the Contract List report.

From the Contract Hierarchy report, you can click various items and links to view other reports. These other reports can show data by various time periods (Daily, Weekly, Monthly, Year-by-Week, and Year-by-Month reports) or provide more detailed information about the contract (such as the Contract Details report and Outages report).

Daily, Weekly, Monthly, and Yearly Reports

Time-based reports such as the Daily, Weekly, Monthly, and Yearly reports for customers display target and actual values to enable you to view how all of the nodes of a contract are performing for a specified date. Each of these reports present similar information, differing only in the increments of time over which data is reported, as follows:

Daily reports	display information about one or more objective types for all contract nodes for a selected day.
Weekly reports	display week-to-date values as well as the daily values for a selected week.
Monthly reports	display month-to-date values as well as the daily values for a selected month.
Year-by-Week reports	display the year-to-date values as well as the weekly values for a selected year.
Year-by-Month reports	display the year-to-date values as well as the monthly values for the selected year.

These time-based reports enable you to identify the objectives that are out of compliance and drill down to discover the specific dates and times during which targets were missed. These reports also provide several ways to navigate to different levels of summary. You can filter these reports by objective type, use the calendar links to select specific summary levels, and use the previous and next arrow icons to navigate to a different time period at the same summary level.

You can also click a graph icon to view a graphic report of the data for the corresponding node. See “Graphs in Time-Based Reports” on page 175 for more information about the graphs in Daily, Weekly, Monthly, and Yearly reports for customers.

The following is an example of a Daily report for customers:

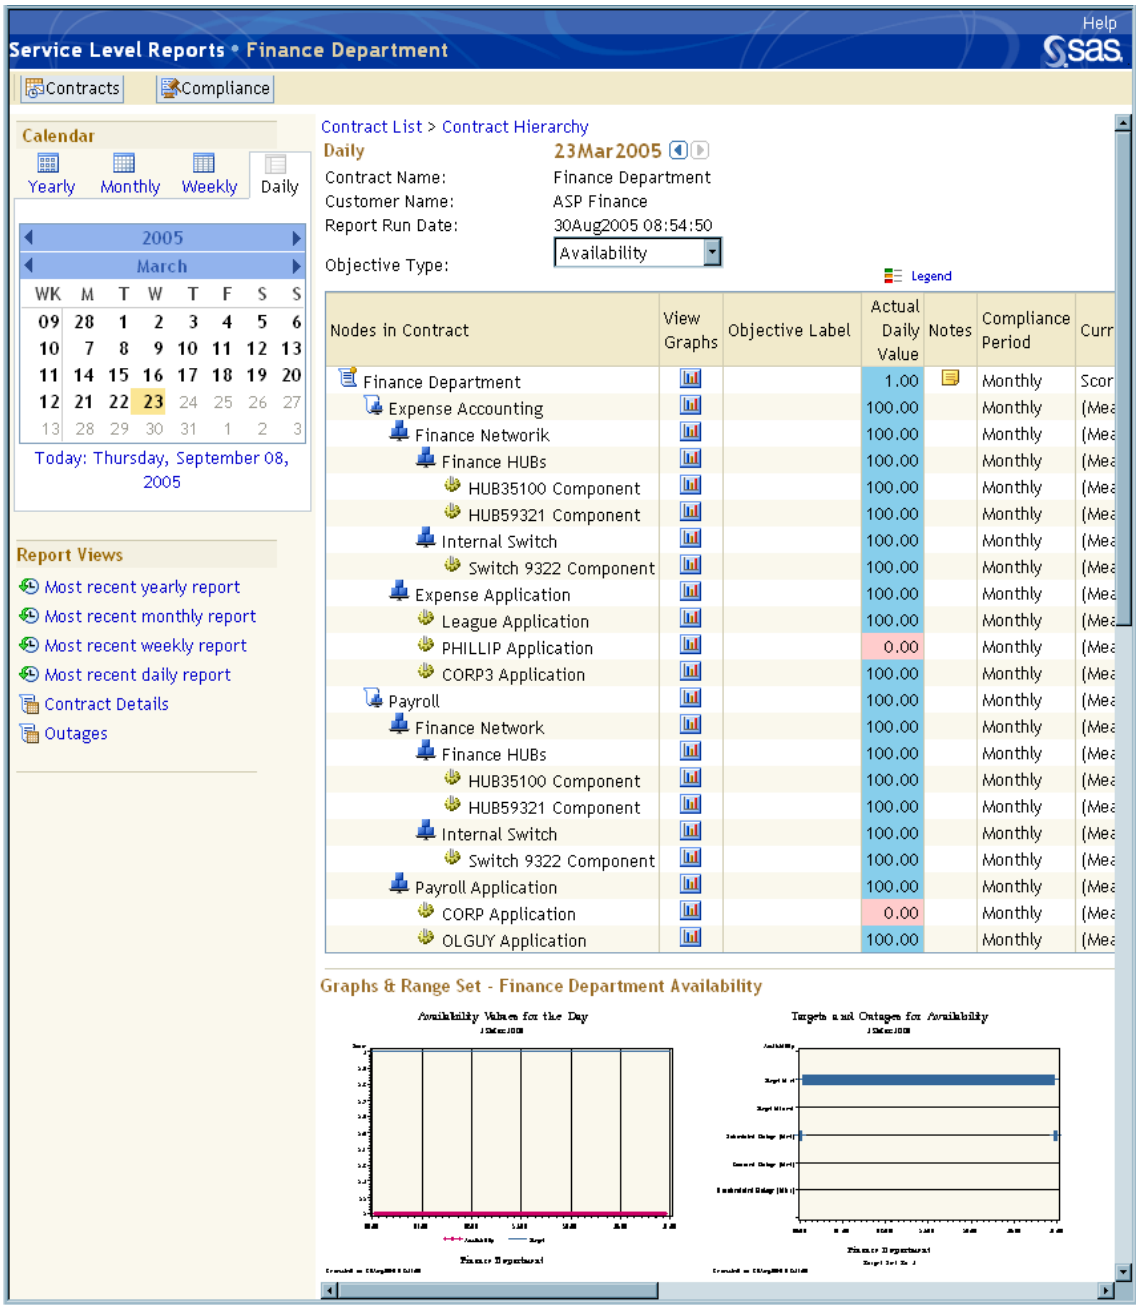

Time-based reports are accessible from both the Contract List report and the Contract Hierarchy report. The Contract List report and the Contract Hierarchy report provide different ways to access the same time-based information; the resulting time-based reports are simply displayed with different filters.

To view time-based reports from the Contract List report, complete the following steps:

1. Click the **`Contracts`** tab.
2. Select a value in the **`Latest Daily Actual`**, **`Week-to-Date Actual`**, **`Month-to-date Actual`**, or **`Year-to-date Actual`** columns.

To view time-based reports from the Contract Hierarchy report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract name in the table of the Contract List report.
3. Choose one of the following methods to access the desired report:
   - Use the calendar tools to select a time period.
   - Below **Report Views** in the left panel, select **Most recent yearly report**, **Most recent monthly report**, **Most recent weekly report**, or **Most recent daily report**.
   - Select a value in the **Latest Day**, **Week-to-date**, **Month-to-date**, or **Year-to-date** columns.

From the time-based reports, you can also click various items or links to drill down to other reports such as the Contract Details report, the Outages report, and other time-based reports.

Graphs in Time-Based Reports

The Daily, Weekly, Monthly, and Yearly reports for customers include up to two graphical reports for every active objective in a contract. The default graph (or graphs) that appears when you first access a time-based report represents the contract node. You can then click a graph icon in the **View Graphs** column to refresh the graph section of the report with graphs for the corresponding node.

Filtering a time-based report by objective type will also refresh the graph section of the report with the graphs for the selected objective type. However, the time-based reports do not show any graphs if the **Objective Type** filter is set to **All Activated**.

The graphs in each of the time-based reports display the data that best represents the appropriate time period. (If there is no data or only missing data for a time period, then a graph will not be generated for that time period.)

Daily Graphs	show values for every 5–minute observation during a 24–hour period of a selected date. This data is reported from the Detail level data in the contract database.
Weekly Graph	shows values for the 7 days in a selected week. The first day of the week is determined by the Start-of-Week value in the contract. This data is reported from the Day level data in the contract database.
Monthly Graph	shows values for all the days in a selected month. This data is reported from the Day level data in the contract database.
Year-by-Week Graph	shows values for all the weeks in a selected year. This data is reported from the Week level data in the contract database.
Year-by-Month Graph	shows values for all the months in the selected year. This data is reported from the Month level data in the contract database.

The Daily, Weekly, Monthly, and Yearly reports all contain a plot graph that shows actual values for the corresponding time period. If a target value was specified for a node, the plot graph will also contain a horizontal reference line that shows the target value for the node.

The following example is a plot graph of response time values from a Daily report. This sample graph shows that the Payroll Application consistently performed better (lower) than its response time target of 11000 required on March 23, 2005.

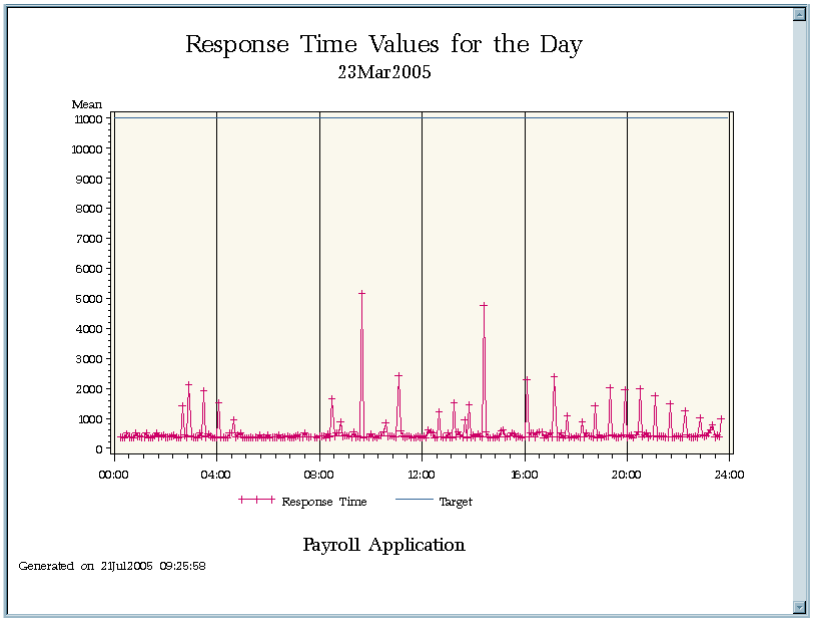

If a target value was specified for a node, then the Daily report also contains an additional plot graph that shows target and outage information for each node for the day. If a target value was specified for a node, the plot graph will contain tick marks that show if a target was met or missed. If an outage was entered for the node, then the plot will also contain tick marks that show what type of outage was entered for the node. A footnote will show the target value. This type of plot graph is only available on Daily reports. (If there is no data or only missing data for a time period, then a graph will not be generated for that time period.)

The following example is a Targets and Outages graph for response time from a Daily report. This sample graph shows that the Payroll Application met its target for the entire 24–hour period on March 23, 2005, and it experienced a scheduled outage at the beginning and end of the day.

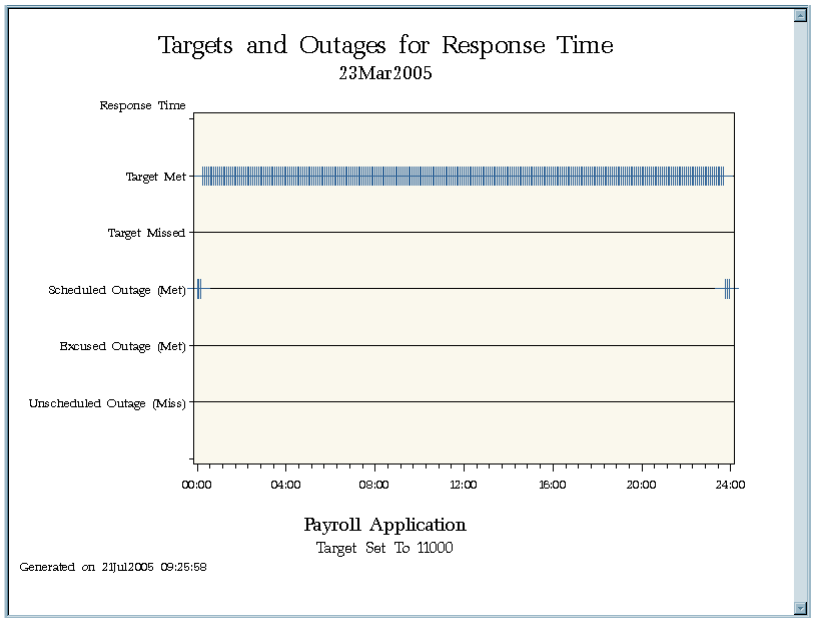

Three-Month Calendar in Daily Reports

A three-month calendar view of the contract's daily values is available in Daily reports for customers. This three-month calendar view is only available on Daily reports.

The three-month calendar view shows bold numbers to indicate that there is an actual value for that day. You can pause the mouse over a bold date to view the actual value for that day. In addition, the calendar cells will be color coded if a range set is specified for the period. You can use the right and left arrows to scroll to additional calendars.

The following example is a three-month calendar section of a Daily report. This sample shows an actual value of 1.00 for the response time on March 5, 2005.

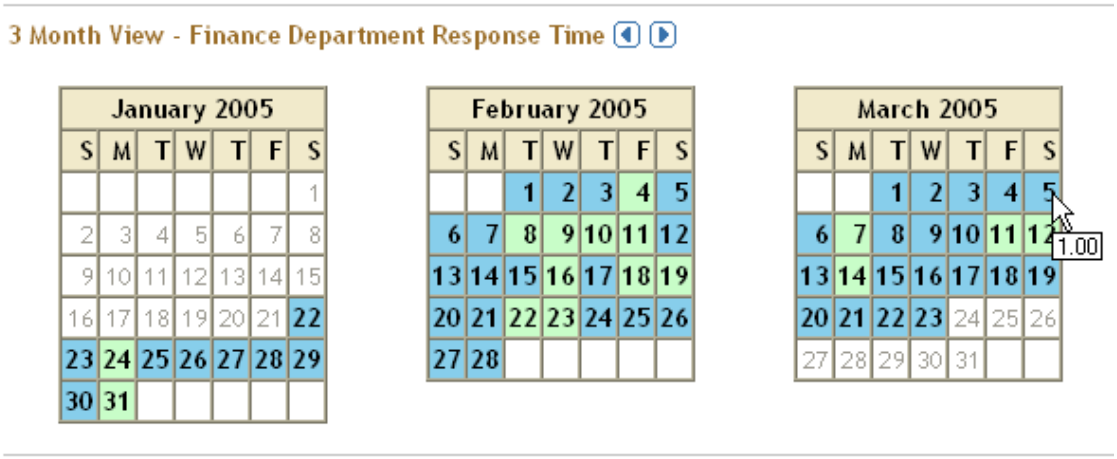

Contract Details Report

The Contract Details report for customers enables you to view all the settings that are associated with the contract, including information about the customer and provider, calculation methods, range sets, and date filter information used for activated objectives of a specified contract.

The following is an example of a Contract Details report for customers:

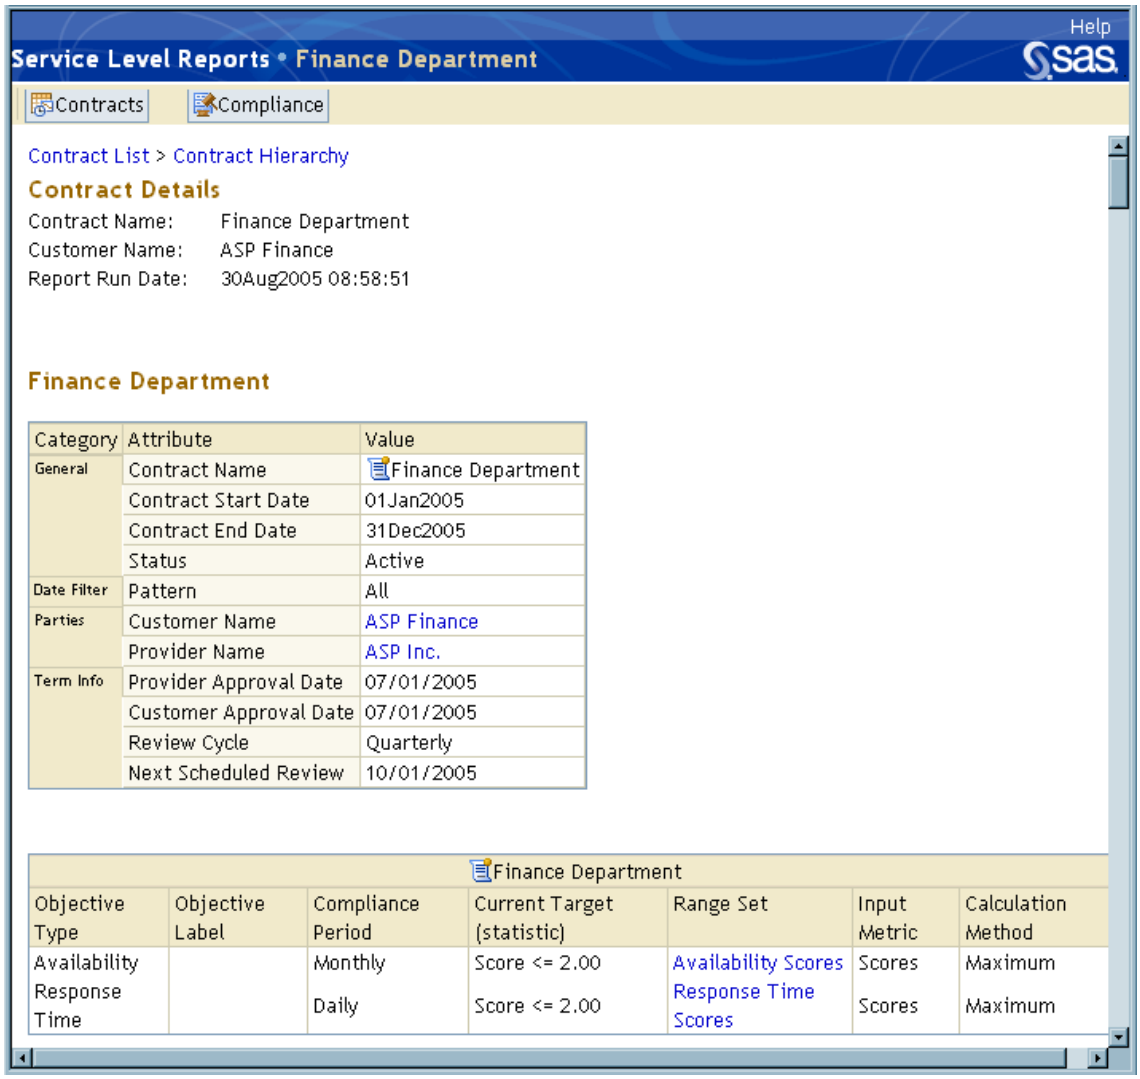

To view the Contract Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the Contract List report.
3. Below **Report Views** in the left panel, click **Contract Details**.

From the Contract Details report, you can also click various items or links to drill down for more information about the terms of the contract. Some of the other reports you can access from this report include Customer reports, Provider reports, and Range Set Details reports.

Customer Report

The Customer report for customers enables you to view the customer and contact information for a selected contract.

The following is an example of a Customer report for customers:

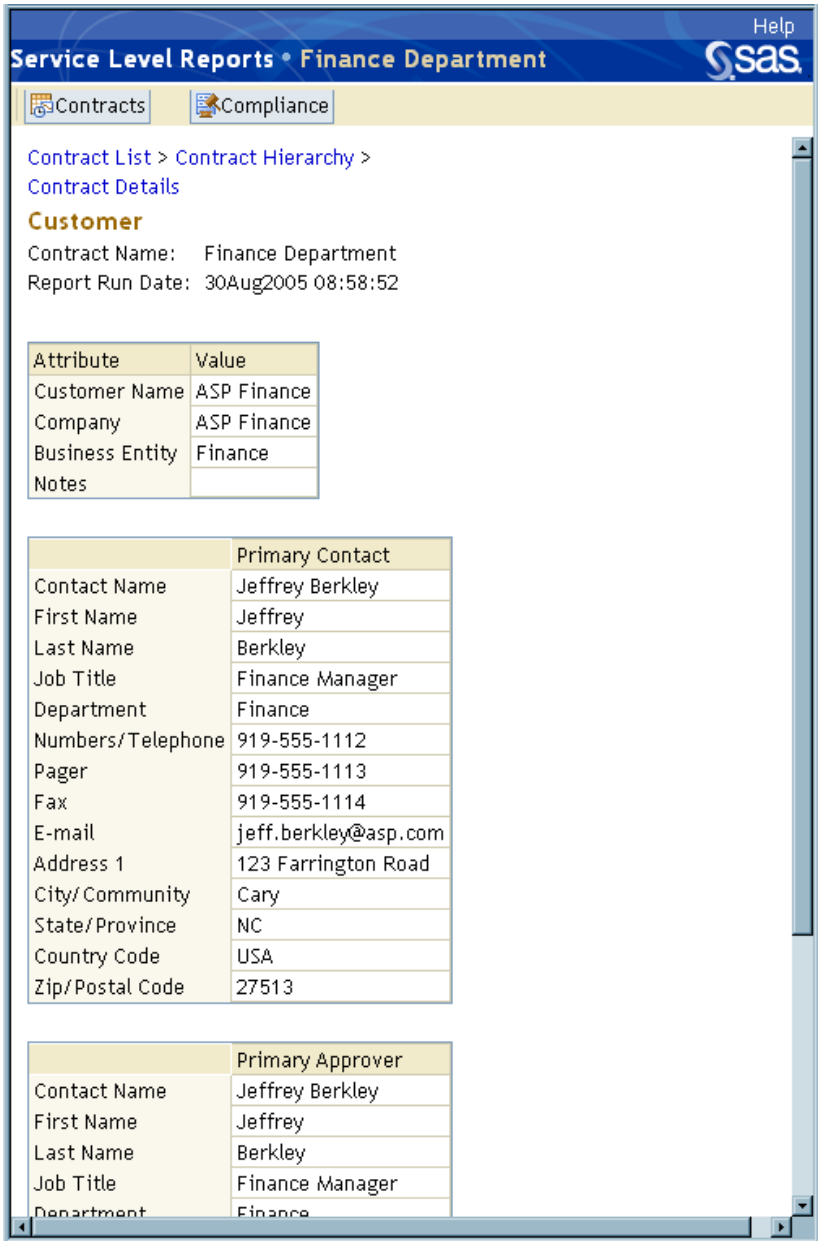

To view a Customer report, you can complete the following steps:

1 Click the **Contracts** tab.

2 Select the name of a customer in a **Customer** column of the Contract List report.

Provider Report

The Provider report for customers displays information about the provider for a selected contract. The layout of the Provider report is identical to that of the Customer report.

The following is an example of a Provider report for customers:

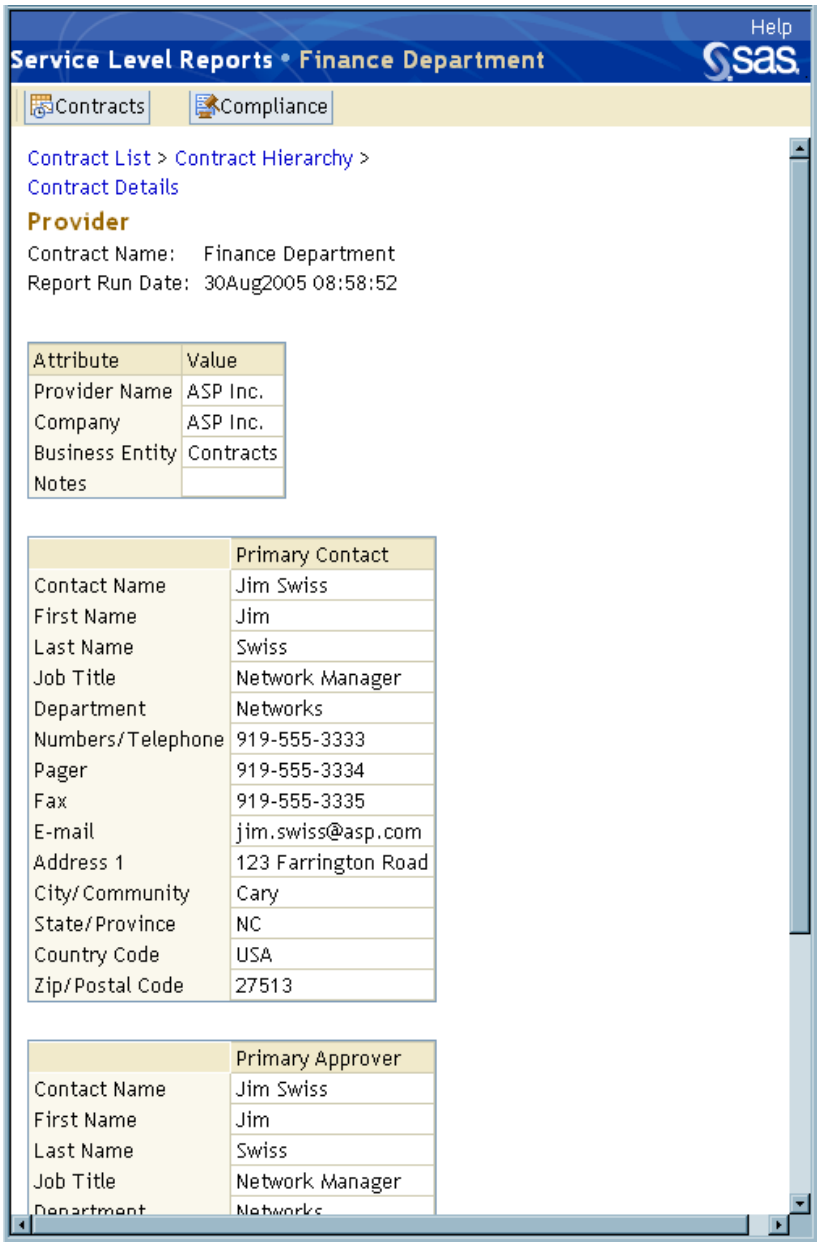

To view a Provider report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in a **Contract Name** column of the Contract List report.
3. Below **Report Views** in the left panel, click **Contract Details**.
4. Select a provider name in the **Provider Name** row of the Contract Details report.

Range Set Details Report

The Range Set Details report for customers enables you to see all of the details for the range set that is associated with a selected objective in a contract. This report displays the ranges that make up the range set as well as their range descriptions so that you can review the colors and boundary conditions that are assigned to each range.

The following is an example of a Range Set Details report for customers:

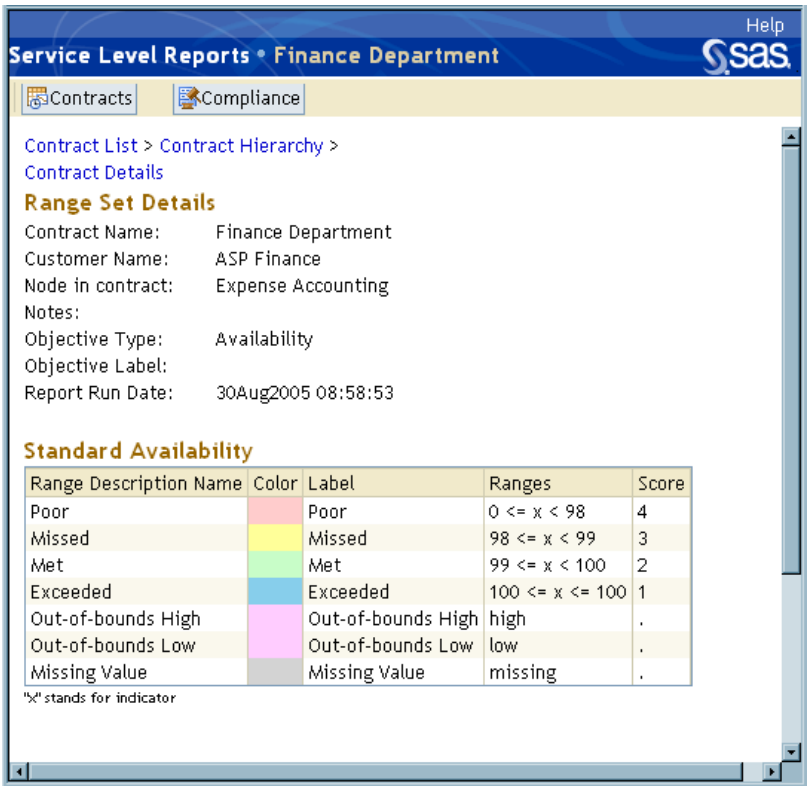

To view a Range Set Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in a **Contract Name** column of the Contract List report.
3. Below **Report Views** in the left panel, click **Contract Details**.
4. Select a range set name in a **Range Set** column of the Contract Details report.

Outages Report

The Outages report for customers enables you to view all of the outages for every node (and its child nodes) in a selected contract. Customers might use this report to view the outages that have been entered for their contract(s).

The following is an example of an Outages report for customers:

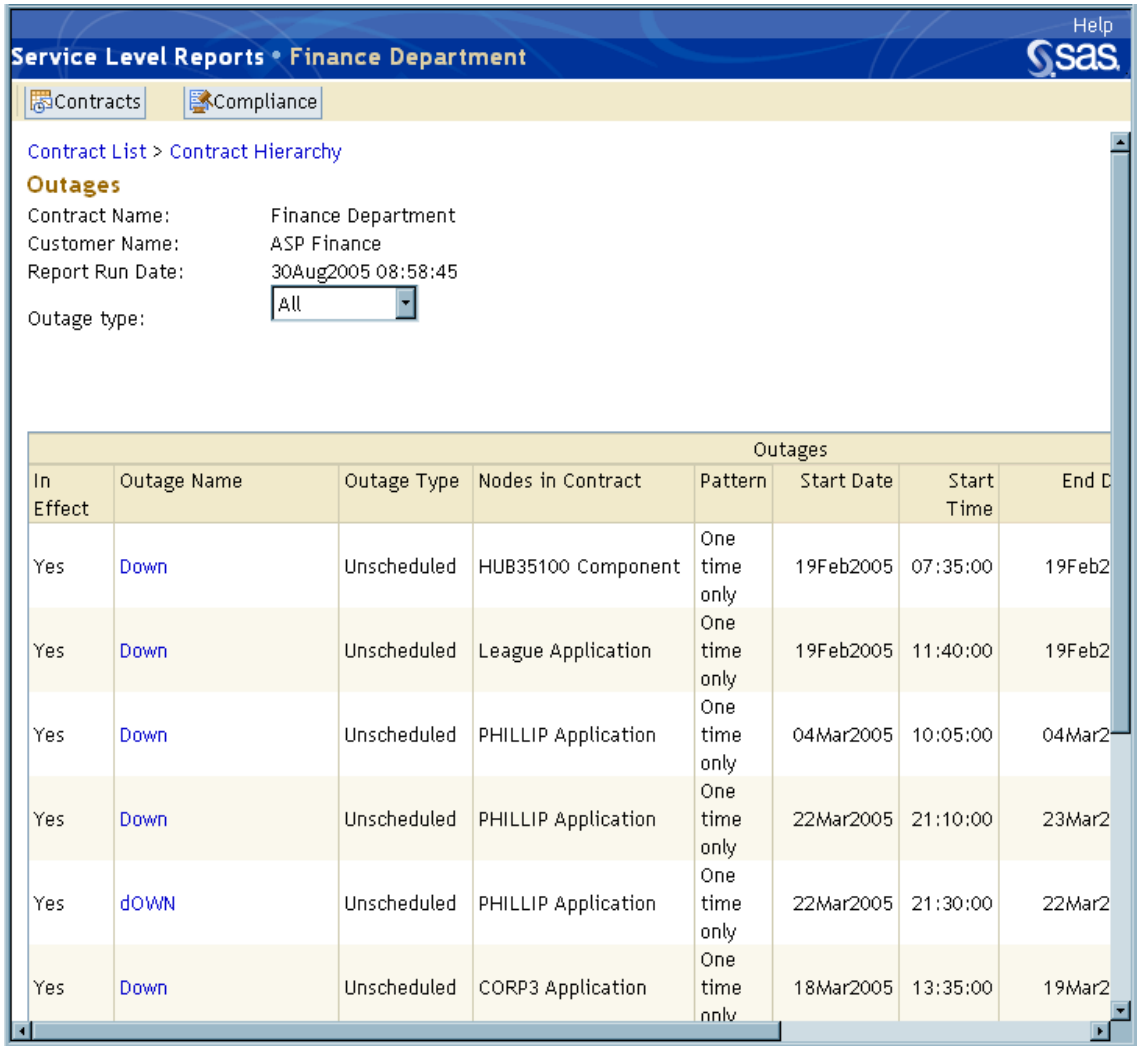

To view an Outages report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the **Contract Name** column.
3. Below **Report Views** in the left panel, click **Outages**.

From the Outages report, you can also select an outage name to access the Outage Details report for more information about the outage.

Outage Details Report

The Outage Details report for customers enables you to view all of the details, including the exceptions, for a selected outage. This information includes a display of the date and reason for all the exceptions to the outage.

The following is an example of an Outage Details report for customers:

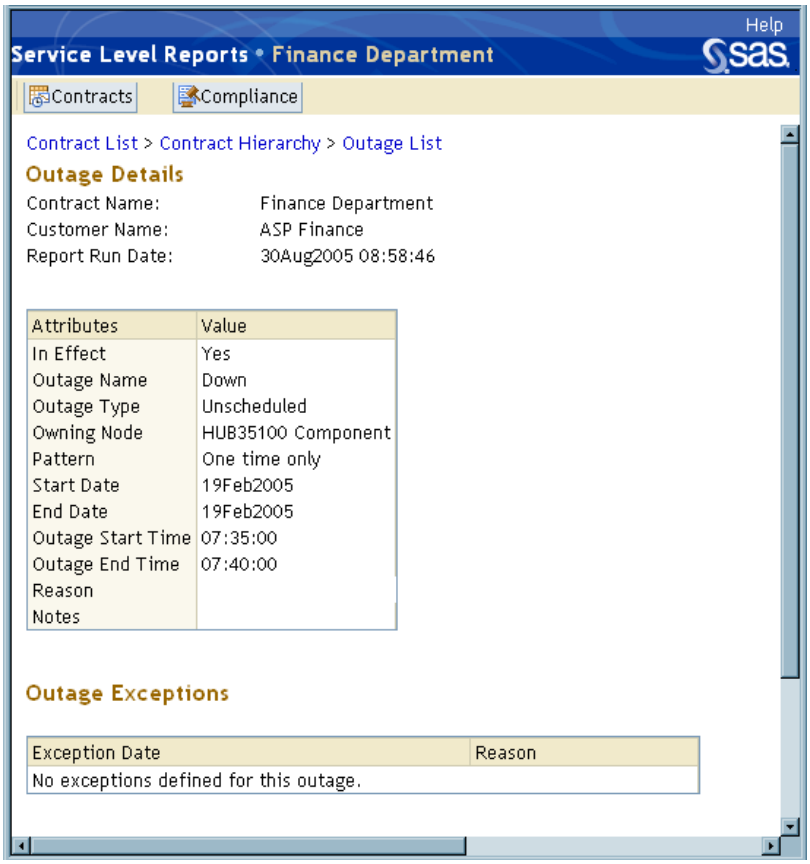

To view an Outage Details report, complete the following steps:

1. Click the **Contracts** tab.
2. Select a contract in the **Contract Name** column.
3. Below **Report Views** in the left panel, click **Outages**.
4. Select an outage in the **Outage Name** column of the Outages report.

Compliance List Report

The Compliance List report is an optional report that is available for customers from the **Compliance** tab. Because the Compliance List report is optional, a set of Service Level reports for a customer might or might not include the **Compliance** tab and its corresponding reports.

The Compliance List report for customers is available from the **Compliance** tab. This report enables you to determine if contracts are meeting their targets during the compliance period for each active objective. You can also filter this report by all targets, missed targets for the current period, and missed targets for the previous period.

The following is an example of a Compliance List report for customers:

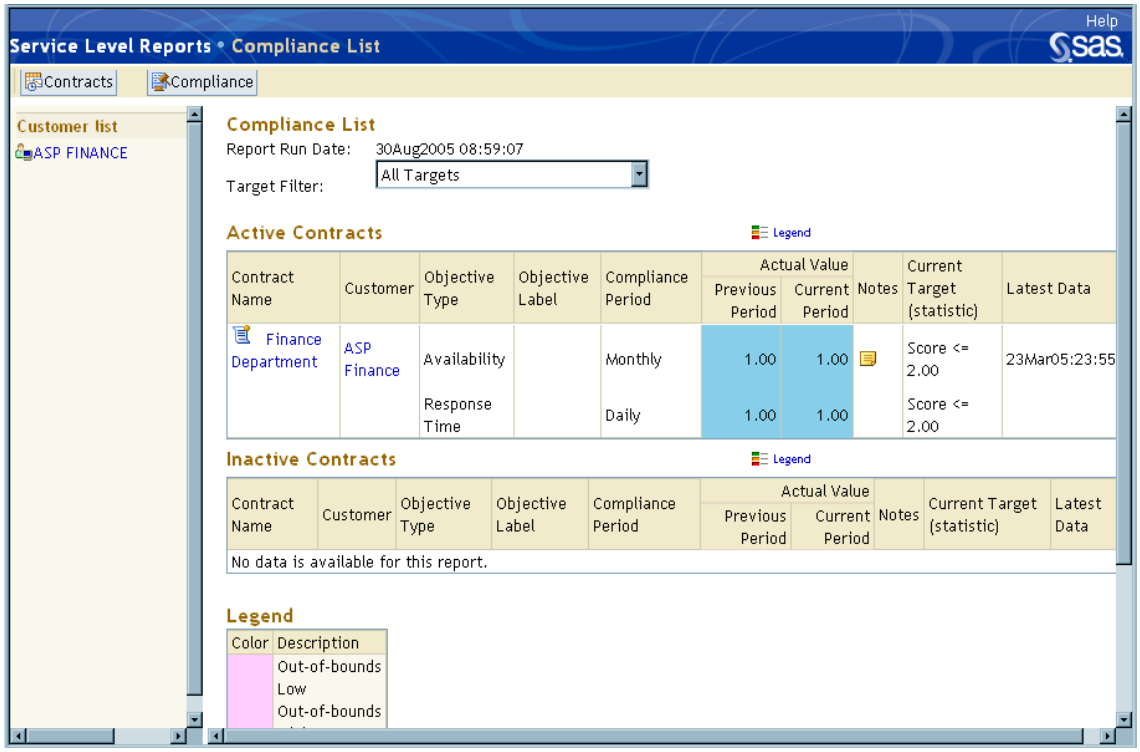

To view the Compliance List report, click the **Compliance** tab.

From the Compliance List report, you can also click various items or links to drill down to other reports such as the Compliance Details reports and the Customer reports.

Compliance Details Report

The Compliance Details report is available for customers from the Compliance List report on the **Compliance** tab. Because the Compliance List report is optional, a set of Service Level reports for customers might or might not include the **Compliance** tab and its corresponding reports.

The Compliance Details report for customers enables you to determine if nodes within a specific contract are meeting their targets during the compliance period for each active objective. This report also enables you to view the details for the individual components in order to uncover any problems with resources. You can filter this report by contract, all targets, missed targets for the current period, and missed targets for the previous period.

The following is an example of a Compliance Details report for customers:

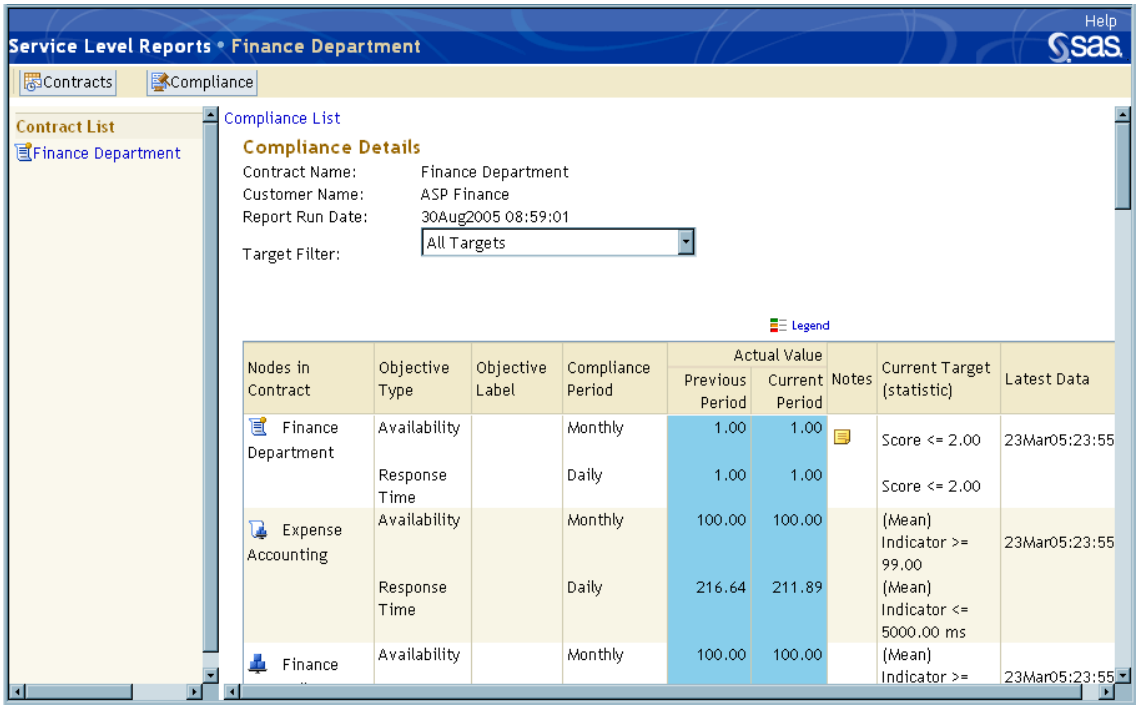

To view the Compliance Details report, complete the following steps:

1. Click the **Compliance** tab.
2. Select a contract in a **Contract Name** column of the Compliance List report.

CHAPTER

13

Working with SAS Enterprise Guide to Create Service Level Reports

Introduction to Working with SAS Enterprise Guide to Create SAS IT Service Level Management Reports

SAS Enterprise Guide uses projects to manage a collection of related data, tasks, code, and results. With projects, you can run multiple tasks on the same group of data files and create interactive or ad hoc reports. This chapter provides details and instructions that are specific for using SAS Enterprise Guide to create and view SAS IT Service Level Management reports. For additional information about how to use SAS Enterprise Guide, see *SAS Enterprise Guide Administrator: User's Guide* or the SAS Enterprise Guide online help.

The Sample Project

The SAS IT Service Level Management client installation CD contains sample data and a sample SAS Enterprise Guide project. You can use the sample project immediately to view sample SAS Enterprise Guide reports. You can also copy the sample project and customize it to work with the new contract databases that you create. Because SAS Enterprise Guide provides various ways to report on your data, you can also customize the project to create different kinds of reports.

To use and customize the sample project, you should first unzip the samples.zip file on the SAS IT Service Level Management client installation CD and copy the contents of the zip file to a chosen location on your hard drive. The contents of the zip file include several folders that contain sample data and a sample SAS Enterprise Guide project. The sample project, ITSLM21.egp, is in the **`samples/EGProject`** folder. It is designed to work with the *finance* sample contract database in the **`/slm/samples/cdb/finance`** folder.

Note: For more information about installing the sample databases from the installation CD, refer to the installation instructions to unzip and import the *finance* sample contract database to your SLM Server. △

Complete the following steps to open the project:

1. Start SAS Enterprise Guide by double-clicking the SAS Enterprise Guide icon on your desktop or by selecting **Start ▶ Programs ▶ SAS ▶ Enterprise Guide** from the **`Start`** menu. The Welcome to SAS Enterprise Guide dialog box opens.
2. Click **`More Projects`** to browse for the sample project.
3. Locate your copy of ITSLM21.egp on your hard drive and click **`Open`**. The sample project will open.

Note: This project was created using SAS Enterprise Guide 3.0.2 with hot fix 3 applied. Projects that are created or saved with SAS Enterprise Guide 3.0 or later cannot be opened with any earlier version of SAS Enterprise Guide. △

Viewing the Sample Project

SAS Enterprise Guide provides a Project window and a Process Flow window to display projects in different ways. You can choose to work in the window that displays the project in the format that you prefer.

The Project window displays the project items in a hierarchy tree format. This window shows small icons in an expandable list that demonstrates how each object is connected to or nested within other objects of the project. Select the **`Project`** tab to view the project in this format.

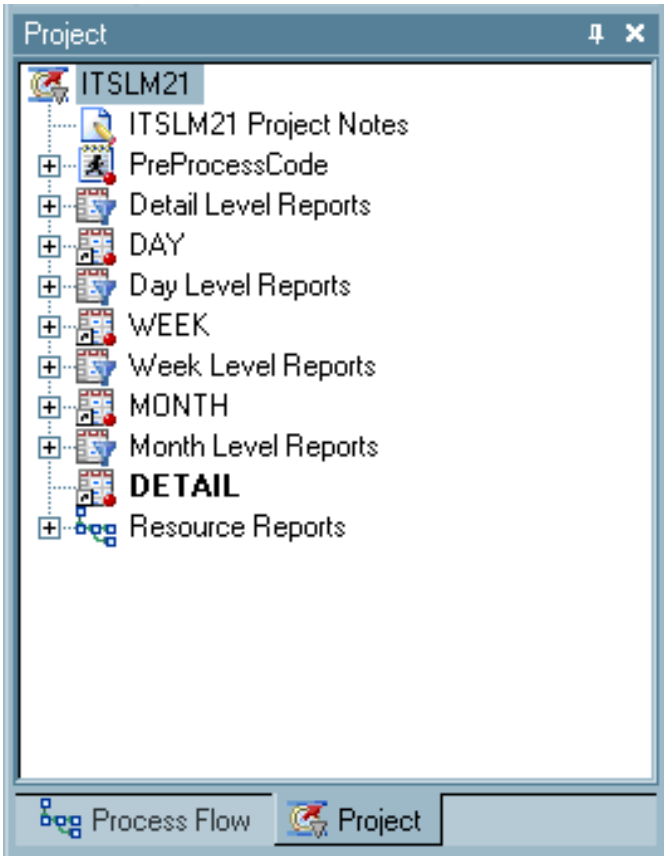

The following top-level items will appear in the Project window for the sample project:

- ITSLM21 Project Notes (notes)
- PreProcess Code
- Detail Level Reports (query)
- DAY (data table)
- Day Level Reports (query)
- WEEK (data table)
- Week Level Reports (query)
- MONTH (data table)
- Month Level Reports (query)
- DETAIL (data table)
- Resource Reports

The Process Flow window enables you to view the project in a process flow format. This window shows icons or objects connected by arrows that show the relationships and connections between objects in the project. Click the **`Process Flow`** tab to view the project in this format.

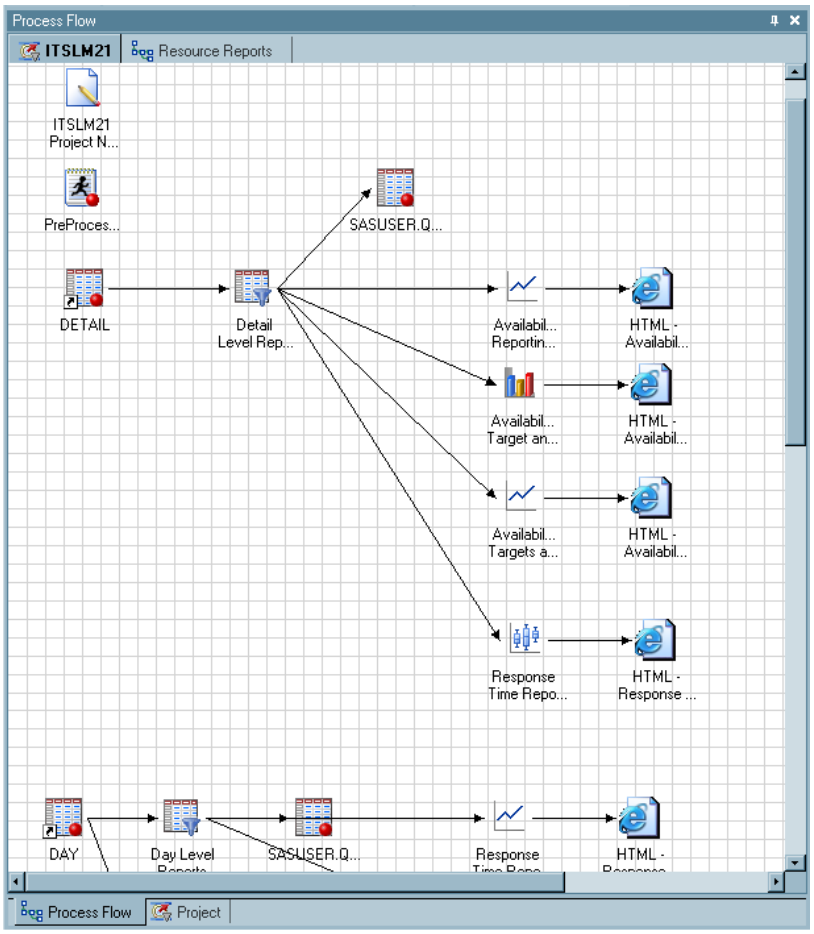

A single project can have multiple process flows. The sample project includes two process flows (ITSLM21 and Resource Reports) that are each accessible from corresponding tabs at the top of the Process Flow window.

This chapter provides instructions for working in the Project window of SAS Enterprise Guide. For more information about the Project and Process Flow windows, refer to the SAS Enterprise Guide online help.

Components of the Sample Project

SAS Enterprise Guide projects contain several types of objects such as notes, code, data, queries, report tasks, and HTML reports. The following list describes some of the object types that are in the sample project.

Notes — provide detailed instructions for customizing the sample project so that you can report on your own data.

Preprocess Code — contains the information necessary for SAS IT Service Level Management to connect to a contract database, set up macro variables, and create user-defined formats that are used by the queries and reports in the sample project.

Data Tables	contain raw data from the *finance* sample database. SAS Enterprise Guide uses this data to generate reports.
Queries	filter data from the data tables for generating reports. You can set parameters for queries to filter data in various ways based on your needs. Queries can also prompt you to create another output table with the subset data. The queries in the sample project are linked to reports that have already been generated for you. You will be able to view these pre-existing sample reports when you first open the sample project. However, you must alter and run the preprocess code before you can actually run any of the queries and generate additional reports on the sample data or on your own contract database.
Report Tasks	define the HTML reports. When the report tasks are run, the corresponding HTML report output is generated.
Report Output	shows specified data from the contract database in an HTML graphic display.

Customizing the Sample Project

You can customize the sample SAS Enterprise Guide project to report on your contract database. The following sections describe how to use the sample project notes and preprocess code to customize the project for your specific contract database.

The Sample Project Notes

After you have successfully run the ETL job in SAS IT Service Level Management to create your contract database, you must complete several tasks in order to use SAS Enterprise Guide to report on your contract.

The notes file in the SAS Enterprise Guide sample project has two purposes. It provides instructions to help you complete the tasks that enable you to produce SAS Enterprise Guide reports on data in your contract database. The notes file also provides information about the sample project, such as how the project was created and what is included in it.

Double-click **ITSLM21 Project Notes** in the Project window to display the notes file in the main SAS Enterprise Guide window.

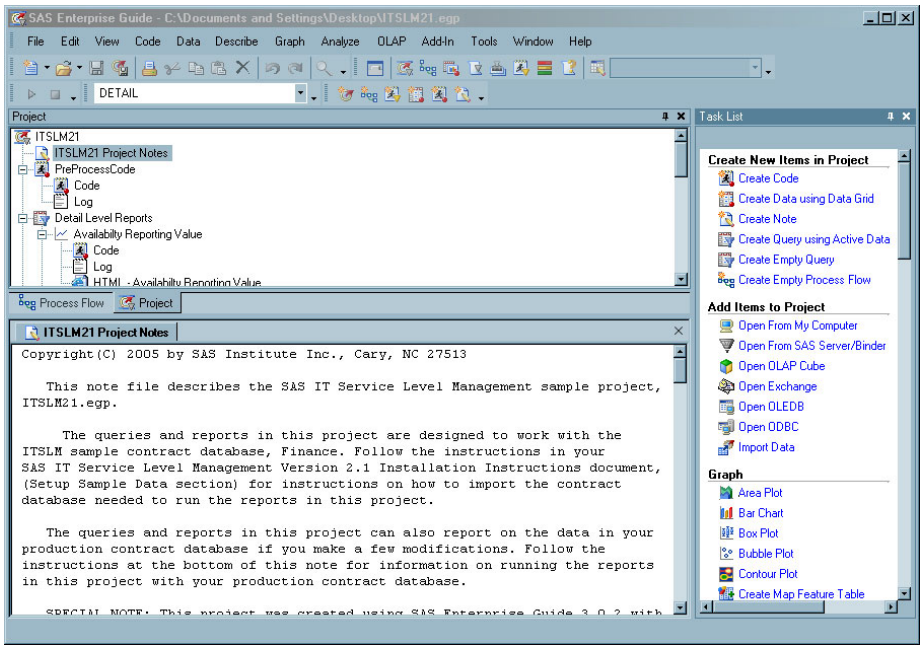

Read through the notes to familiarize yourself with the sample project. The steps that appear at the end of the notes file are also included below with additional suggestions. The following steps provide instruction for setting up the project to work with either the sample data or your contract data:

1. Ensure that you have successfully run the ETL job to create your contract database.
2. Make a copy of the sample project and follow the customization instructions included within the preprocess code. (See “The Preprocess Code” on page 195 for more details about customizing the preprocess code.)

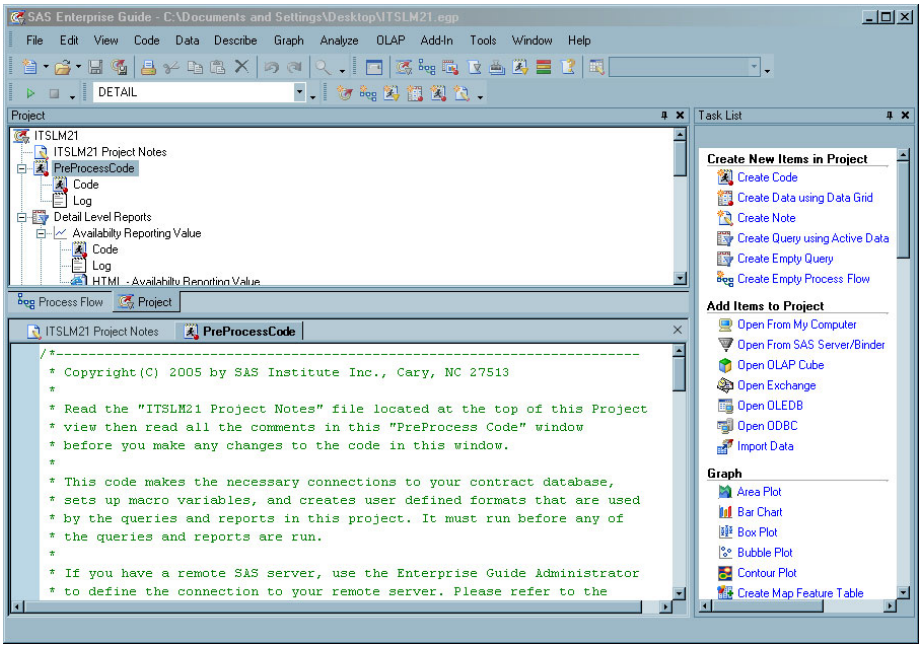

3 Submit the preprocess code by using one of the following two methods based on your system configuration:

- If your SLM Server is on the same machine that is running the SAS Enterprise Guide client, right-click **`PreProcess Code`** in the SAS Enterprise Guide project window and select **`Run on Local`**. Make sure that you are connected to your contract database by verifying that there are no errors in the generated log.
- If your SLM Server is not on the same machine that is running the SAS Enterprise Guide client, right-click **`PreProcess Code`** in the SAS Enterprise Guide project window and select **`Select Server`** to choose the server that holds your SLM data. Make sure that you are connected to your contract database by verifying that there are no errors in the generated log.

  Note: For information about setting up a remote server, see *SAS Enterprise Guide Administrator: User's Guide*. △

4 Double-click each of the queries (Detail Level Reports, Day Level Reports, Week Level Reports, and Month Level Reports) to update the list of actual values that are associated with each query parameter. For example, to update the Day Level Reports query, you would complete the following steps:

a Double-click the Day Level Reports query icon or its label to display the Query Builder window.

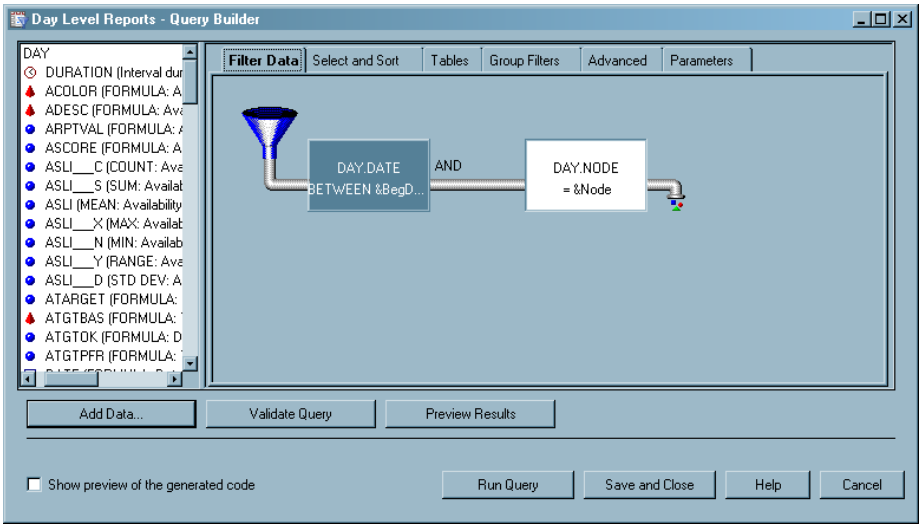

b Click the **`Parameters`** tab.

c Click the first parameter to highlight it in the large white window.

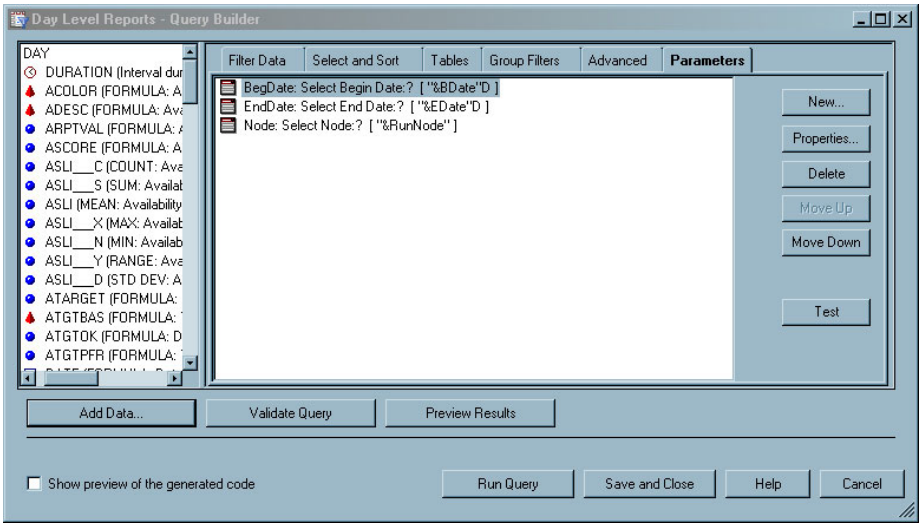

d Click **`Properties`**.

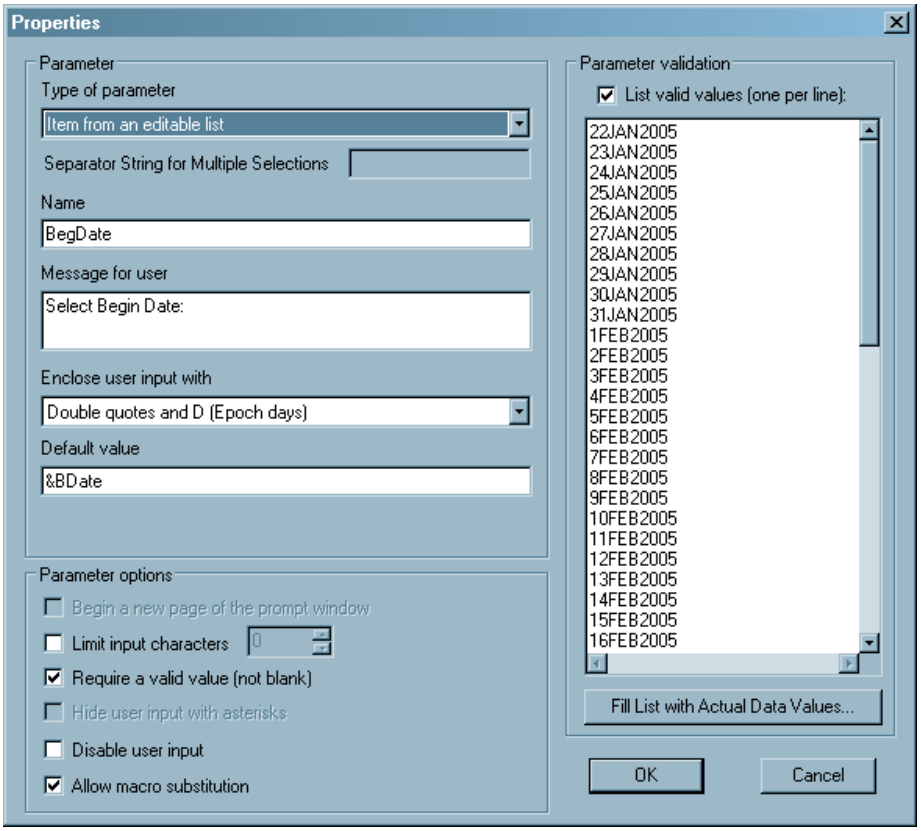

e If there are any values in the parameter validation window, then click **`Fill List with Actual Data Values`**.

f Click **`OK`**.

g Repeat for every parameter in all four queries and each time there is new data in the contract database.

Note: These illustrations and instructions provide an example of query parameters. Not all queries use the **`Item from an editable list`** parameter and **`Parameter validation`** field to specify date variables. For example, some queries provide a calendar widget for selecting date values. You can use the options in the Properties dialog box to choose a different filter that best accommodates your specific type of data. △

5 Save the project.

The Preprocess Code

The SAS Enterprise Guide sample project enables you to run various reports without requiring that you perform extensive programming. You only need to check the preprocess code to verify or edit the location of the database that you want to use and confirm the default dates for the reports. Instructions for customizing the preprocess code are detailed in the preprocess code and in the project notes.

To verify and edit the PATH statement and default dates in the preprocess code, complete the following steps:

1 Double-click **`PreProcess Code`** in the Project window to open the preprocess code file in the SAS Enterprise Guide main window.

2 Modify the appropriate PATH statement for your site so that it points to the directory that contains the ADMIN library folder for the contract that you want to work with. The LIBNAME statement looks like this:

```
%let PATH = c:\slm\samples\cdb\finance\;    /* Windows Platform */
*%let PATH = /slm/samples/cdb/finance;      /* UNIX Platform    */
*%let PATH = SLM.SAMPLES.CDB.FINANCE;       /* OS/390 Platform  */
libname ADMIN  "&PATH.ADMIN";
```

3 If you are using UNIX or z/OS, remove the comment of the appropriate PATH statement by removing the single asterisk before the percent symbol. Then, add a comment to the Windows PATH statement by adding a single asterisk before the percent symbol.

4 Modify the default values for the macro variables that are used by some of the queries in the sample project. You can also modify the values to match the data in any contract database that you create. You can override the default values when you run the queries and reports interactively. You can also change the default values of each query in the parameter properties, or you can disable macros. The default values section looks like this:

```
%let BDate=31JAN2005;                       /*  Begin Date  */
%let EDate=22FEB2005;                       /*  End Date    */
%let BWeek=23JAN2005                        /*  Begin Week  */
%let BMonth=01JAN2005;                      /*  Begin Month */
%let RunNode=OLGUY Application;             /*  Node        */
%let Node=OLGUY Application;                /*  Node        */
```

The preprocess code that is provided in the sample project also contains the following two formats that are used by target and outage reports:

- TrgtFmt (target format) is used by target reports at the DAY and MONTH levels, where the only possible values are 0 and 1. The TrgtFmt section looks like this:

```
proc format;
    value TrgtFmt
     1 = 'Target Met'
     0 = 'Target Missed'
    -1 = '.'
 run;
```

  The value -1 is used as a placeholder in the plot.

- OtgFmt (outage format) is used by target and outage reports at the DETAIL level, where various negative numbers represent types of outages. The OtgFmt section looks like this:

```
proc format;
    value OtgFmt
    1 = 'Target Met'
    0 = 'Target Missed'
   -1 = 'Scheduled Outage (Met)'
   -2 = 'Excused Outage (Met)'
   -3 = 'Unscheduled Outage (Miss)'
   -4 = '.'
run;
```

  The value -4 is used as a placeholder in the plot.

  You can also add other formats here to support new report types.

After you have completed the customization steps in the preprocess code and notes, you are ready to run queries and reports.

Sample Reports

Each of the data tables and queries in the SAS Enterprise Guide sample project are associated with sample reports that have already been generated. Click the plus sign (+) next to a table or query label to expand the tree and see what sample reports are available.

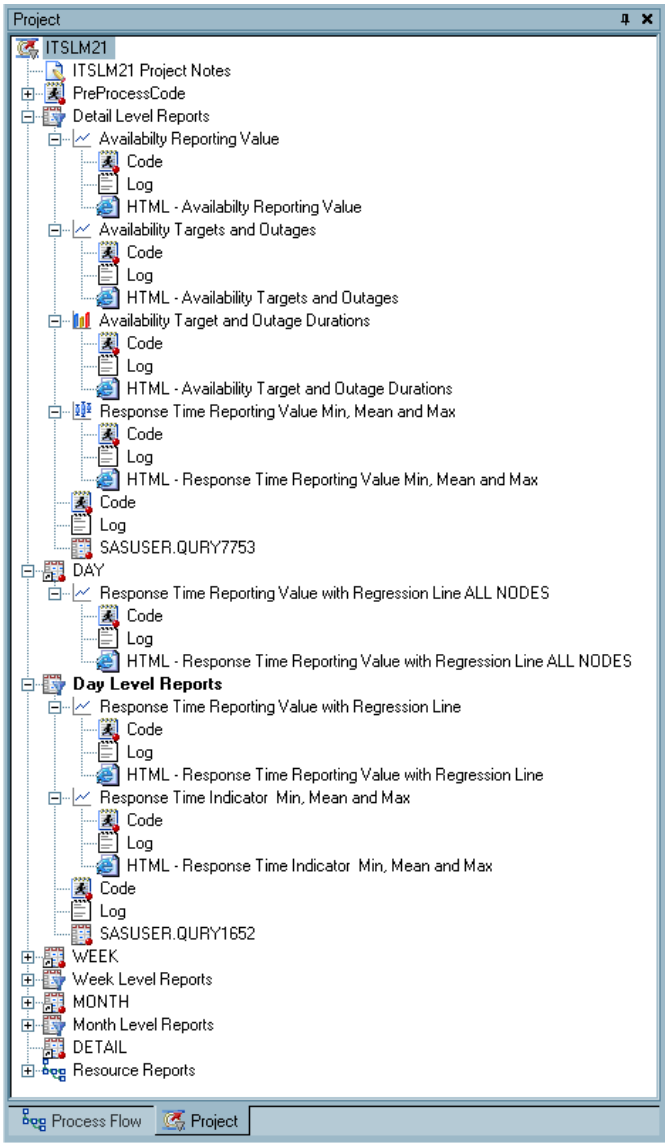

The SAS Enterprise Guide sample project includes report samples that show service level information and resource information. These reports use data from two kinds of *views* of the contract database. A view is a pointer to a grouping of the tables in the contract database. Views contain no data; they merely define data that is stored elsewhere. The sample reports primarily use node views and the resource views of the contract database.

Sample reports that present service level information get data from the node views in the ADMIN level of the contract database. These node views get compliance information by comparing data in the resource tables with the objective targets configured in SAS IT Service Level Management. Sample resource reports, located in the **`Resource Reports`** branch of the hierarchy tree, get data from the resource views in the ADMIN level of the contract database. The data from the resource views is raw data and it does not include accommodations such as filtering or outages. The sample resource reports show resource data for a given resource before it is interpreted with a calculation method for an objective.

Note: For more information about using views in the contract database for SAS Enterprise Guide reports, see "The Contract Database" on page 237. △

You can view reports by fully expanding the tree and then double-clicking on the report description that you want to view. For example, to view the Detail Availability Reporting Values report in the sample project, complete the following steps:

1. Select the plus sign (+) next to Detail Level Reports.
2. Select the plus sign (+) next to Availability Reporting Value.
3. Double-click **`HTML --- Availability Reporting Value`** to view the report in HTML format.

The following sections describe some of the sample reports that are in the sample SAS Enterprise Guide project. These sample reports represent only a few of the reports that SAS Enterprise Guide can generate. The samples are provided to help you get acquainted with SAS Enterprise Guide and to give you some ideas about how you can use this tool to generate reports on your own contract data. "Working with SAS Enterprise Guide" on page 208 shows you how to set up reporting options and get help within SAS Enterprise Guide.

Detail Availability Reporting Values Report

The Detail Availability Reporting Values report gives you a quick overview of availability of a node for a selected 24–hour period. The sample report shows that on March 18, 2005, the CORP Application (C000440) node experienced 100% availability (above its target) for most of the day. However, after approximately 6:30 p.m., the node experienced 0% availability. This report helps you research problems by enabling you to focus quickly on the specific hours during which problems occur.

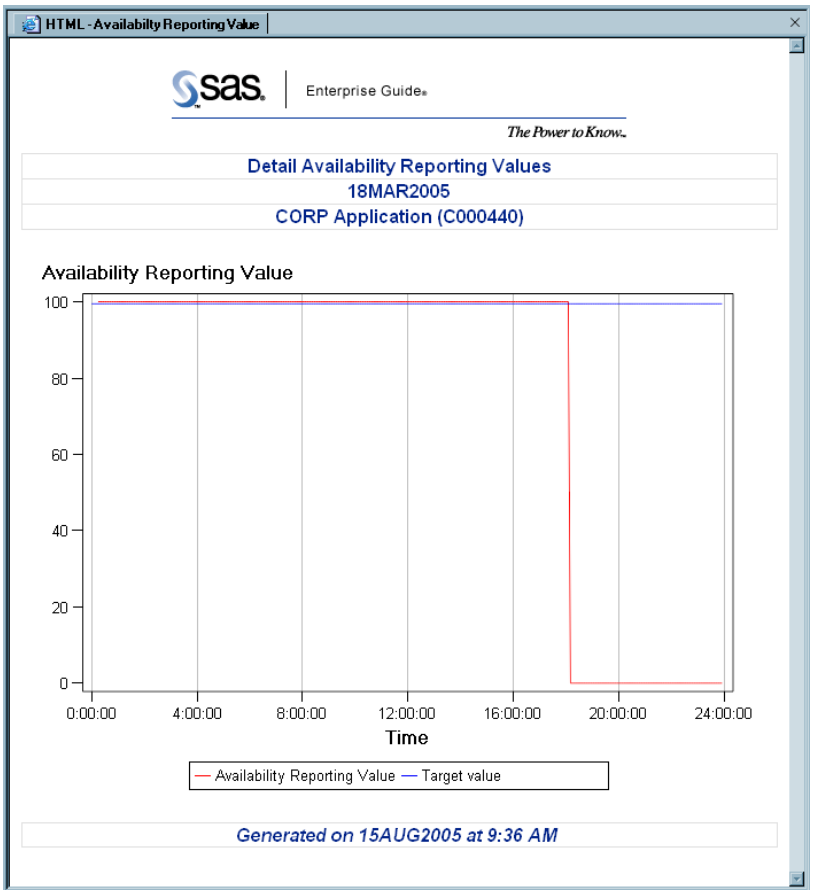

Detail Availability Targets and Outages Report

The Detail Availability Targets and Outages report displays a simple graphic overview of whether availability targets were met or missed for each hour in a selected 24-hour period for a particular node. Also depicted are hours during which scheduled, excused, and unscheduled outages occurred. This report will only show targets and outage information if a target is set for the given node.

The sample report shows that on March 18, 2005, the CORP Application (C000440) node met availability targets for most of the day. This report also shows that the node experienced a scheduled outage for a few minutes at the beginning of the day. An unscheduled outage also occurred at approximately 6:30 p.m., causing the node to miss its availability target for the remainder of the day.

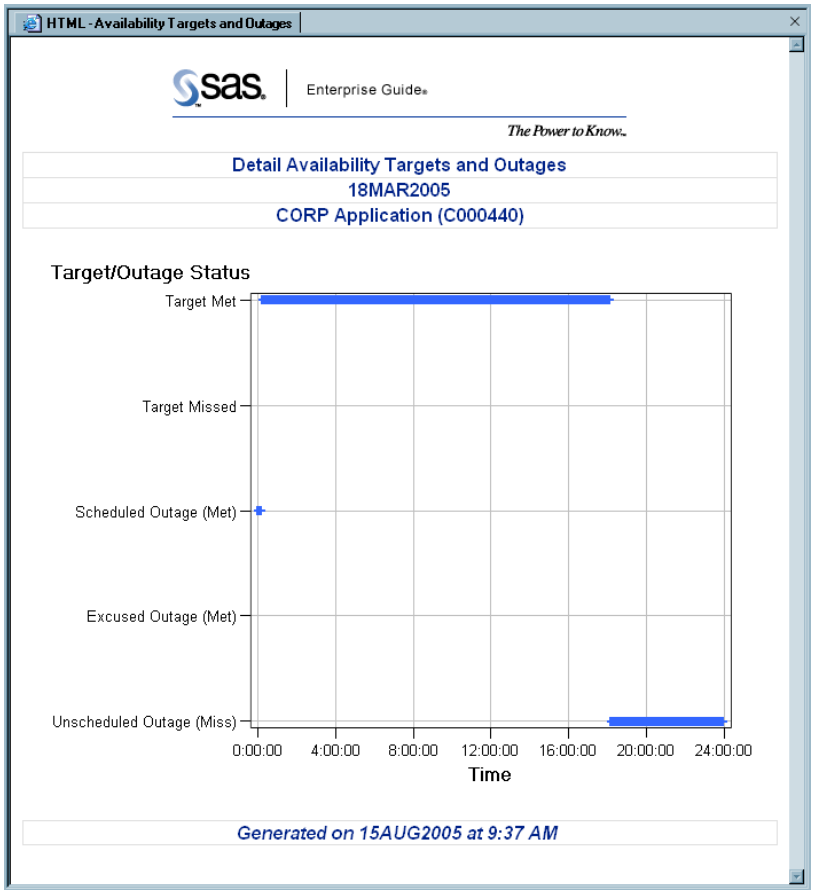

Detail Availability Target and Outage Durations Report

The Detail Availability Target and Outage Durations report uses shaded bars to represent the percentage of the 24-hour period during which targets were met and missed. This report will only show target met and missed information if a target is set for the given node.

The shading in the **`Target Met`** bar indicates what percentage of *target met* time is attributed to good performance of the node, excused outages, and scheduled outages. The shading in the **`Target Missed`** bar shows how much of the *target missed* time is attributed to poor performance of the node or the result of unscheduled outages.

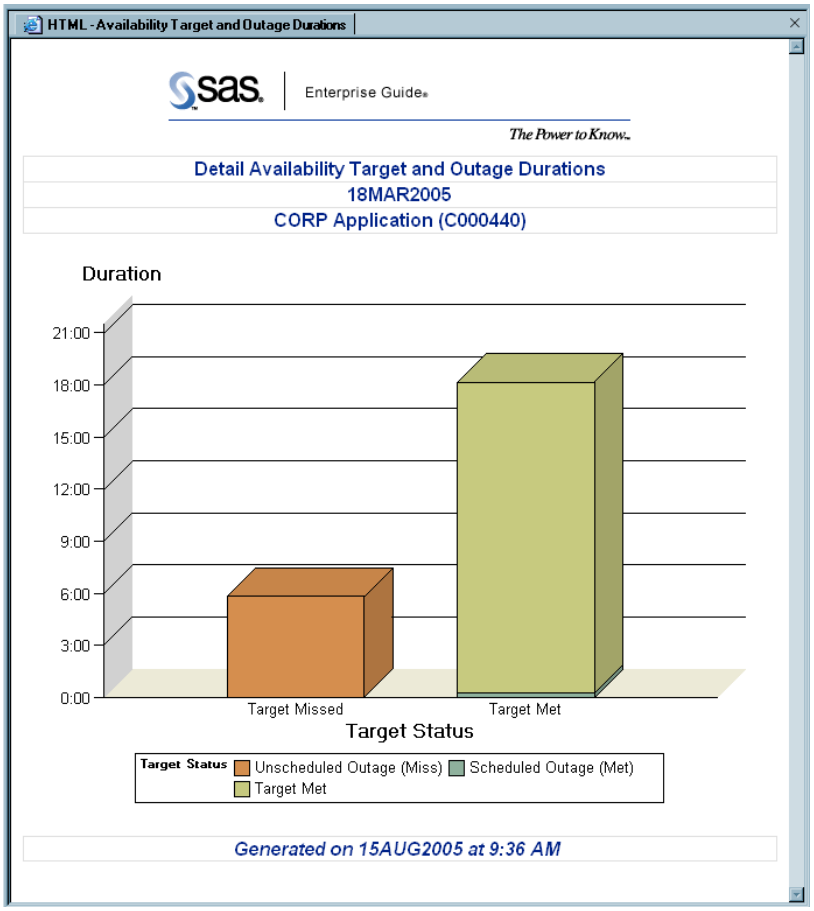

Detail Response Time Reporting Values Report

The Detail Response Time Reporting Values report is a high-low graph. The tops of the vertical lines indicate the highest value for each hour. The bottoms of the lines show the lowest values that were recorded for each hour. The middle horizontal line on each vertical line shows the mean of all response time values that were recorded during that hour. In this sample report for March 18, 2005, you can see that the lowest response time values were usually between 3900 and 4050. The highest values ranged from around 4000 to 5200, and the mean generally fluctuated between 3950 and 4350.

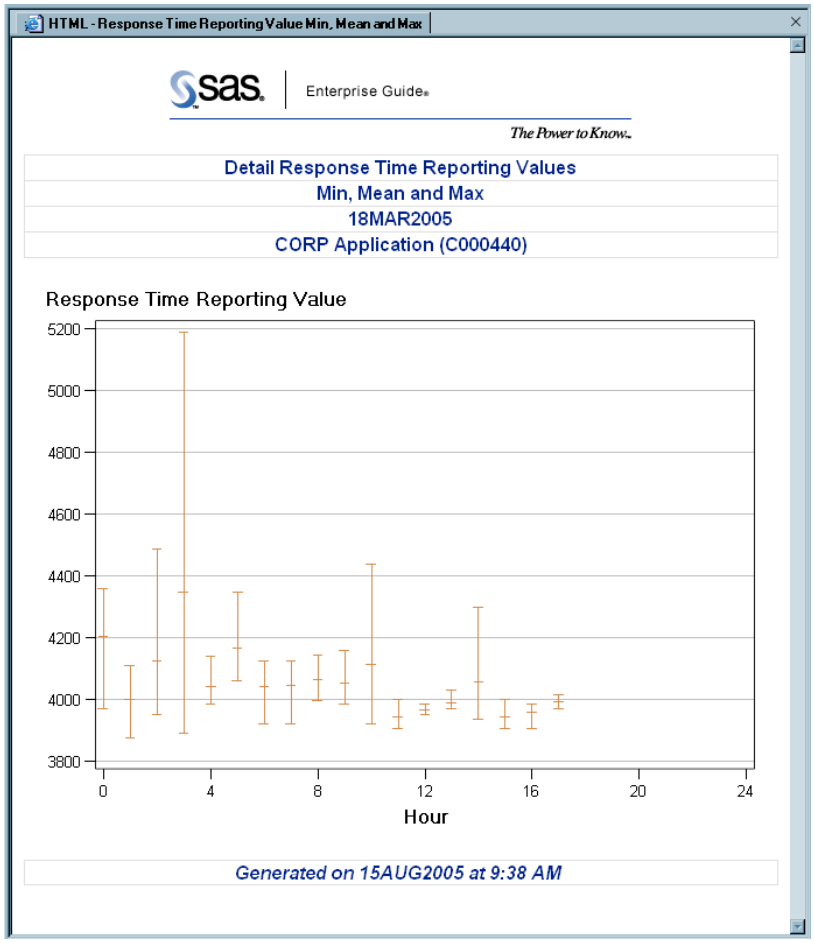

Daily Response Time Reporting Values Report

The Daily Response Time Reporting Values report is a trend report in which each plot represents a node's response time value that was recorded for each day in a given time period. The regression line in this sample report shows that the OLGUY Application (C000441) node experienced a decreasing trend in response time values for the first two weeks of February 2005.

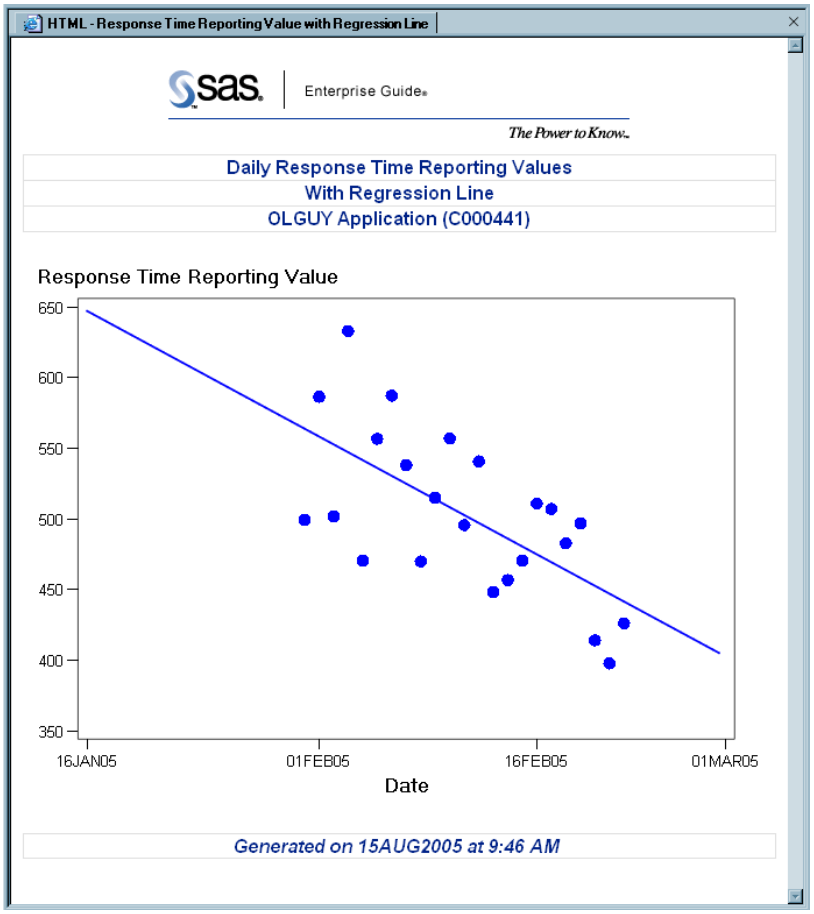

Daily Response Time Indicator Report

The Daily Response Time Indicator report is a line plot report in which three lines represent a node's maximum, mean, and minimum response time indicator that was recorded for each day in a given time period. This sample report shows that the mean response time of the OLGUY Application (C000441) node was consistently low in the period shown. A glance at the maximum response time indicator shows that the highest value fluctuated between approximately 800 and 5400.

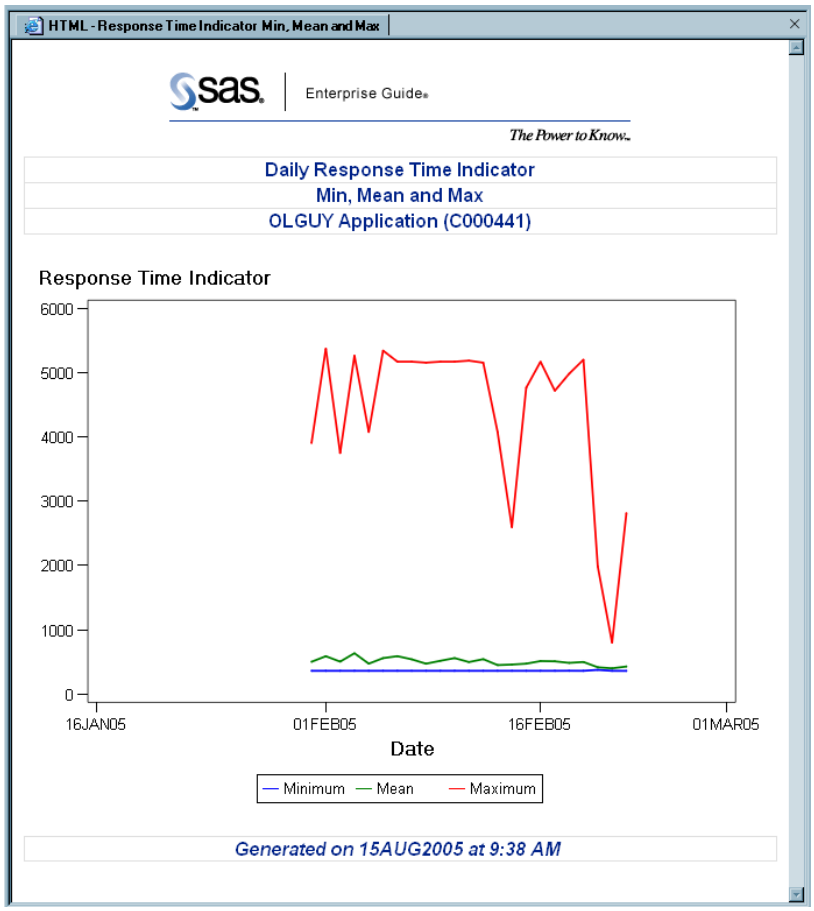

Weekly Availability Reporting Values Report

The Weekly Availability Reporting Values report shows the target for availability as a solid blue line. For each week, the availability reporting value for the node is displayed as a green dot. In this sample report for the CORP Application (C000440), you can quickly see that the weekly availability reporting value for this node only missed the target for two weeks between January 17, 2005, and March 21, 2005.

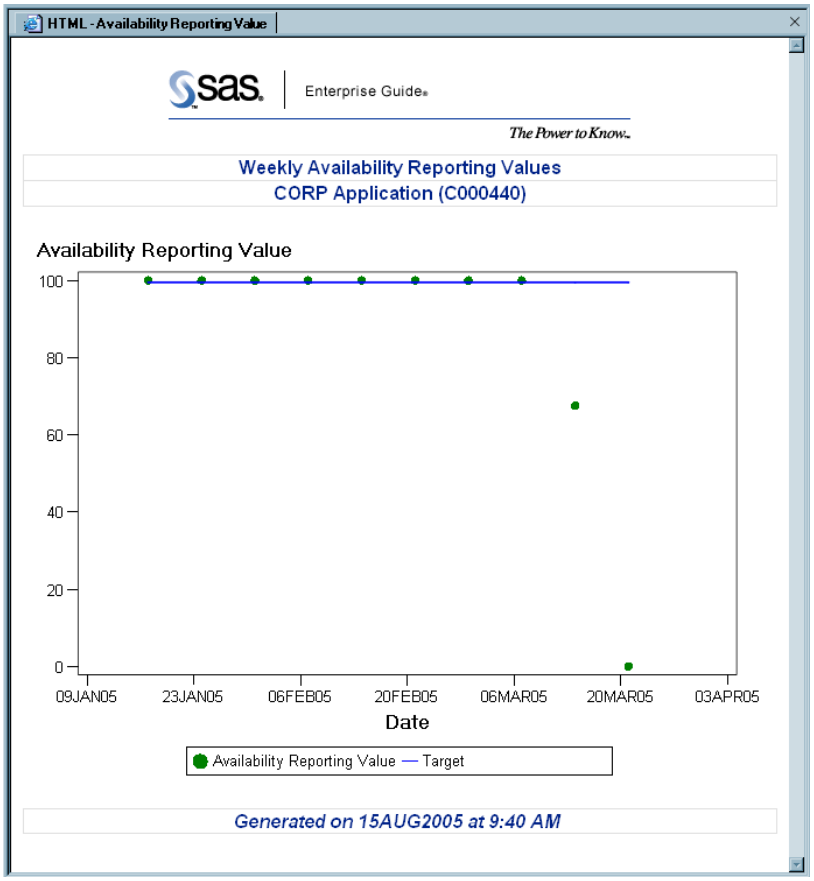

Monthly Response Time Reporting Values Report

The Monthly Response Time Reporting Values report is a bar chart that displays the mean of response time reporting values for a node for each month. In the example below, you can see that the mean response time reporting value for the OLGUY Application (C000441) node for January and February was just above 500, and it increased to approximately 540 for March.

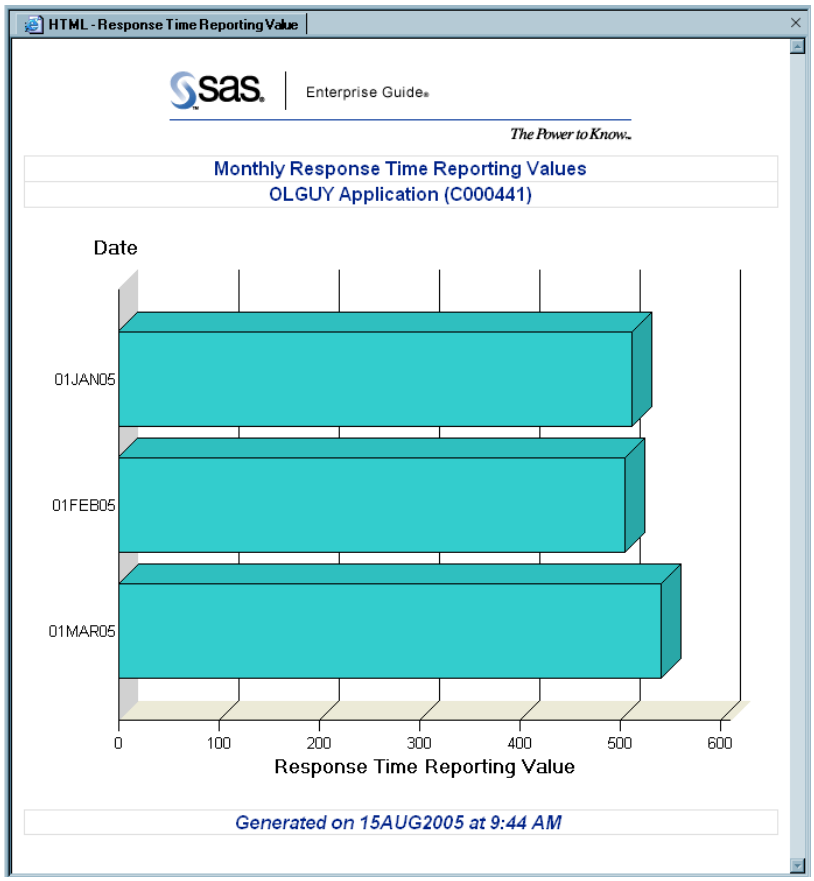

Resource Report: Detail Response Time Mean Report

The Detail Response Time Mean report is a sample resource report. This report shows the mean response time of the CORP3 url resource every 5 minutes over the course of approximately one week. In this example, you can see that mean response time did not fall below approximately three minutes and it did not go higher than 31 minutes.

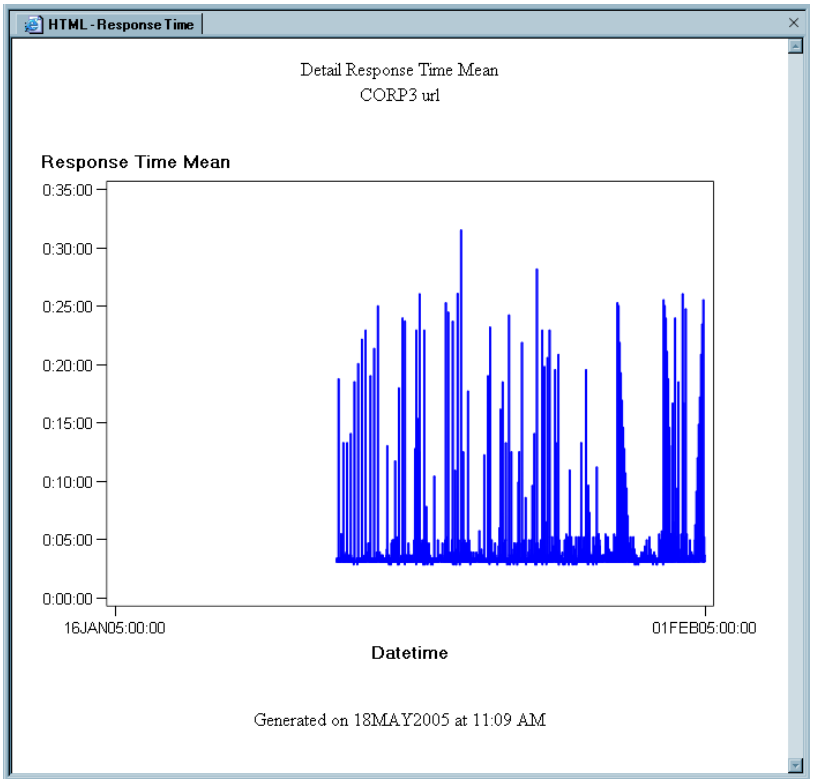

Resource Report: Detail Response Time and Availability Means Summary Table

The Detail Response Time and Availability Means summary table is a sample resource report that shows the raw data for the availability and response time means of the CORP3 url resource. This summary table details numeric values from the resource data for the resource before it is computed with a calculation method for a service level objective. This information shows you the resource data in 5-minute intervals during a given day so that you can isolate a specific time in the day when a resource performed successfully or unsuccessfully.

			Availability	Response Time
			Mean	Mean
MONITOR	Date	Time		
CORP3 url	27FEB2005	0:00:00	100.00	187.00
		0:05:00	100.00	203.00
		0:10:00	100.00	187.00
		0:15:00	100.00	187.00
		0:20:00	100.00	797.00
		0:25:00	100.00	750.00
		0:30:00	100.00	203.00
		0:35:00	100.00	188.00
		0:40:00	100.00	203.00
		0:45:00	100.00	187.00

Working with SAS Enterprise Guide

Setting Reporting Options

To change your reporting options in SAS Enterprise Guide, select **Tools ▶ Options** to open the Options dialog box. This dialog box enables you to change options that affect the entire SAS Enterprise Guide application.

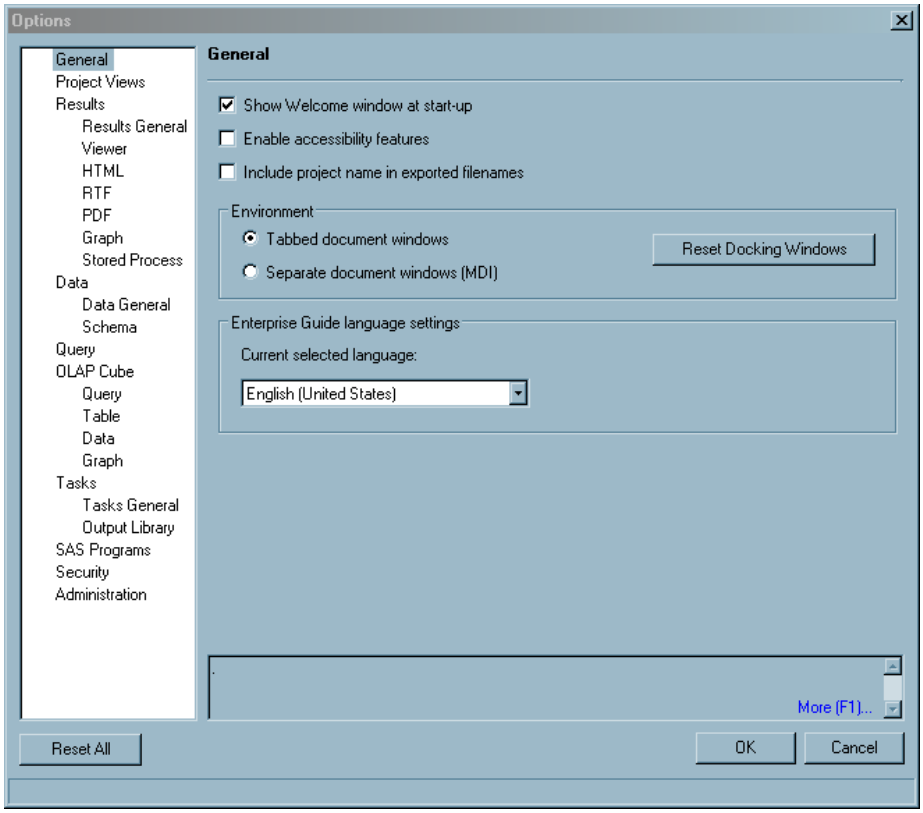

The Options dialog box includes several categories of settings that you can view and edit as needed. These settings are accessible from the menu options listed in the navigation pane on the left of the dialog box. The following information details some of the parameters you might consider when you are working with SAS Enterprise Guide to create and view SAS IT Service Level Management reports.

- Select **Results ▶ Results General** to select one of the following formats for your reports:
  - HTML
  - PDF
  - RTF
  - Text output

  You can access and set additional parameters such as report style for each of these formats by selecting the corresponding options in the navigation pane.
- Select **Results ▶ Viewer** to specify your preferred viewer for the various report formats.

- Select **Results ▶ Graph** to view the current settings for the graphic format of reports. In the **`Graph Format`** field, you can select from the following formats for report graphs:
  - ActiveX
  - Java
  - GIF
  - JPEG
  - ActiveX image (SAS V9)
  - Java image (SAS V9)
  - SAS EMF (SAS V9)

  Note: ActiveX is the default output format and can only be viewed in the SAS Enterprise Guide window or on Microsoft Windows machines using Microsoft Internet Explorer. The ActiveX control output format is interactive. You can right-click on any graphic output that is generated in this format and change the chart type and many other options. △

Getting Help

You can learn more about using SAS Enterprise Guide by exploring the SAS Enterprise Guide online Help and the Getting Started Tutorial that are included with the application. To view the Help files, from the SAS Enterprise Guide menu bar select **Help ▶ SAS Enterprise Guide Help**.

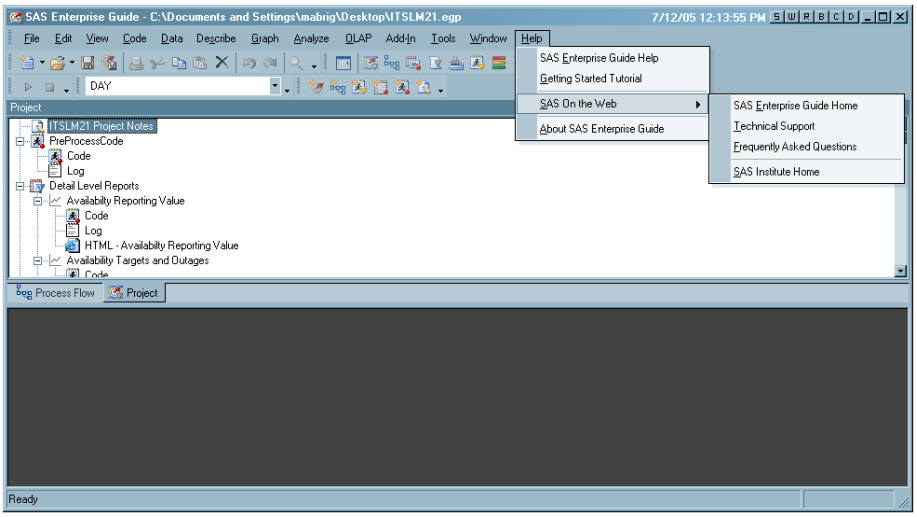

The SAS Enterprise Guide Help provides instructions for creating new projects, building new queries, and customizing SAS Enterprise Guide, along with many other useful topics.

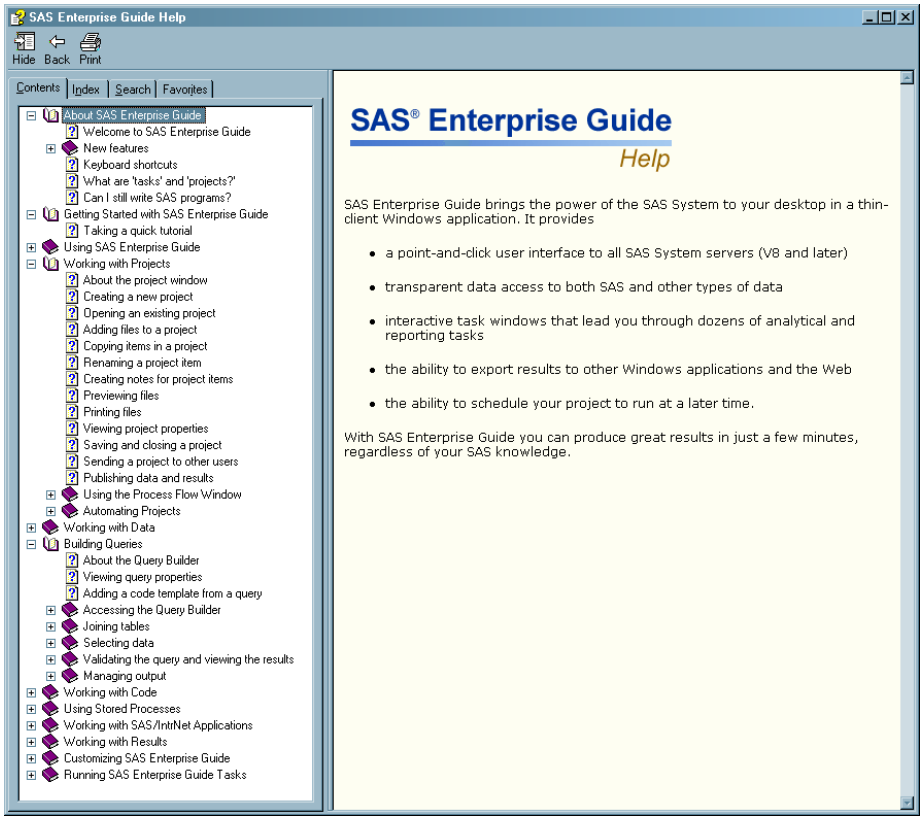

To view a SAS Enterprise Guide tutorial, from the SAS Enterprise Guide menu bar select **Help ▸ Getting Started Tutorial**. You will get step-by-step instructions for working through a hands-on tutorial that uses sample data that is provided during the SAS Enterprise Guide installation.

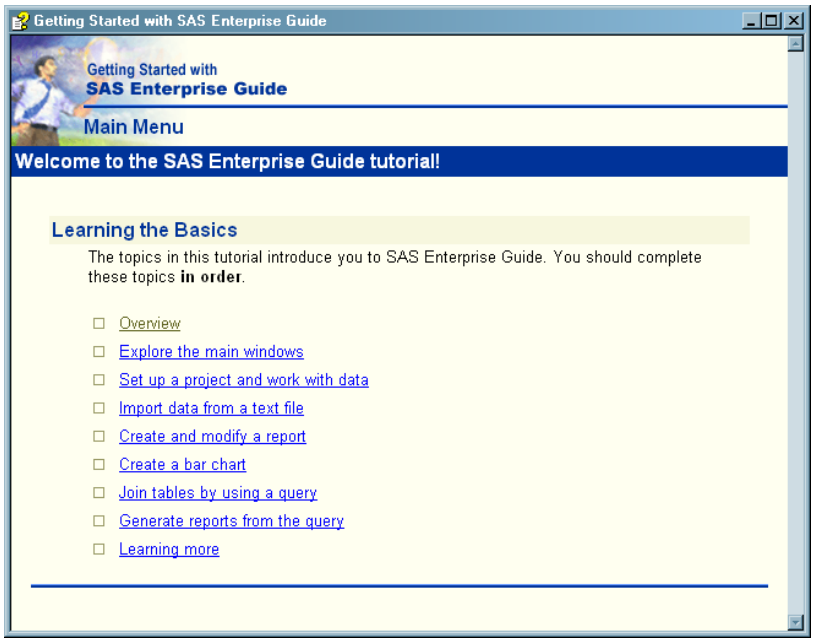

See **`http://www.sas.com/technologies/bi/query_reporting/guide/`** to read more about SAS Enterprise Guide and **`http://support.sas.com/documentation/onlinedoc/guide/`** to find additional documentation for SAS Enterprise Guide.

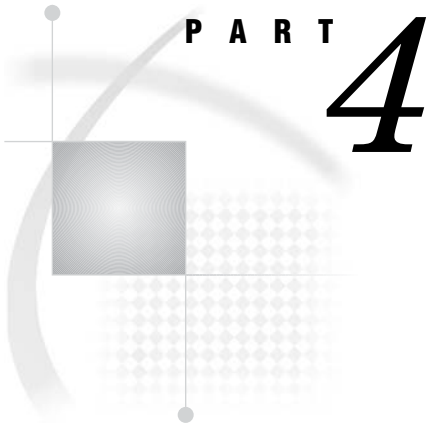

PART 4

Appendixes

APPENDIX

1

Calculation Methods

Calculation Methods

SAS IT Service Level Management supports two types of calculation methods:

- supplied calculation methods

  For the supplied calculation methods, see "Supplied Methods" on page 216. For recommendations about which calculation method to use, see "Recommendations" on page 219.

- user-defined calculation methods

  For user-defined calculation methods, see "Custom Code" on page 218.

For information about how missing values are handled, see "Missing Values" on page 218.

Outages

Service outages can impact the calculation of a node's indicators.

- If an outage *is* in effect for a given node during a given interval, the outage takes precedence over the calculation of the node's indicators. In that case, these rules apply:
  - The indicators are not based on calculations. The calculation methods are ignored.

- The indicators are assigned a value based on the outage type. For more information, see "Effect of Outages on Data" on page 233.

- If an outage *is not* in effect for a given node during a given interval, then for each objective the following statements apply.
  - If the objective is *active*, then its indicator is calculated.
  - If the objective is *inactive*, then its indicator is set to .I, which is a special missing value.

Supplied Methods

The supplied methods are maximum, mean, minimum, percentile, redundant, serial, sum, and weighted mean. The calculations are based on data from the detail tables *in the contract database*.

Note: The availability of a resource must always be expressed as a value greater than or equal to zero. If the measure that you have chosen for your availability objective has a negative value in your source PDB, then the availability indicator for the resource will be set to zero in the contract PDB. This prevents a resource from having a negative availability. △

- *A component node's calculation* is based on the values of the measure variable in the appropriate A, R, T, or U table in the contract database. (For more information about A, R, T, and U tables, see "Detail Tables" on page 243.)

  Note: For all these supplied methods (redundant, serial, and so on), the indicator is calculated separately at each five–minute interval of data in the contract database. This indicator will be summarized into a daily, weekly, and monthly indicator. △

- *A service node, service level agreement node, or contract node's calculation* is based on the values of the child node indicators or scores that you specified when you selected the input metric for the objective's indicator. For information about the component (C) node, service (S) node, or service level agreement (L) node tables of the contract database, see "Detail Tables" on page 243.

Inputs

- For a component node, the input for an availability, response time, throughput, or custom objective indicator is the same-objective-type values from the resources for that node.
- For a service, service level agreement, or contract node, the input for an availability, response time, throughput, or custom objective indicator is the same-objective-type indicators or same-objective-type scores from the node's child nodes.
- For all nodes, the input for a composite indicator is the availability, response time, throughput, or custom objective scores for that same node.

Changing the calculation method

If you change the calculation method, then the data is recalculated by the Synchronize task. In the case where you change to the Weighted Mean calculation method, you will receive missing values for your indicator if the underlying components did not already specify a Weight By variable.

Maximum, Mean, Minimum, and Sum

The Maximum is the maximum of all non-missing values.
The Mean is the arithmetic mean of all non-missing values.
The Minimum is the minimum of all non-missing values.
The Sum is the sum of all non-missing values.

Weighted Mean

You can select this method only for the following objectives: response time, throughput, or custom objective.

The Weighted Mean is the weighted mean of all non-missing values.

Percentile

The percentile function is implemented in code that operates as if it uses the PCTLDEF=5 version of the SAS UNIVARIATE procedure (PROC UNIVARIATE). You can read more about that procedure in the documentation for your version of SAS, but here are some notes that you might find helpful:

- The code ignores the missing values and sorts the remaining values in ascending order.
- You provide the number p. The minimum value of p is 0. The maximum value of p is 100. For example, p can be 12.5 or 95.
- The code returns the value that corresponds to that number. That is, the value that is p/100ths along the list. For example, if you have 200 values, and you request the value at the 90th percentile, the code returns the 180th value in the sorted order, assuming no missing values.

  If the requested percentile does not correspond to an actual value on the list, the code takes the mean of the value on each side.
- If you request the value at the 50th percentile, the code returns the median value.
- If you request the value at the 0th percentile, the code returns the minimum value; if you request the value at the 100th percentile, the code returns the maximum value.

Note: The percentile calculation method is not useful if you are working with a small number of inputs. △

Redundant

You can select this method only for availability objectives. The Redundant statistic is calculated from availability values (in the range 0 through 100) that are input to the calculation (in any order) from that interval.

Redundant means that all of the contributors are in a pool and serve as a backup to one another. Thus, if any one of the contributors is available, then the entire pool of contributors is considered to be available. An example of a redundant relation of contributors is a Web-server pool.

Serial

You can select this method only for availability objectives.

Serial means that all of the contributors are required; that is, that the parent of the contributors is dependent on all of the contributors. (It does not matter whether the business uses the contributors in a fixed sequence or in any sequence, and it does not matter whether the business uses each contributor one time or some contributors are used more than one time.)

- An example of a serial relation of contributors is an end-to-end network connection. In this case, the contributors are used in sequence.
- Another example of a serial relation of contributors is a Web server, a network, and a database server in an e-commerce application. In this case, the contributors are used as needed, and every contributor is used at least once.

Custom Code

Instead of using a supplied calculation method, you can write your own custom code to assign a value to the availability, response time, throughput, or custom objective indicator for any node.

Custom

For the custom calculation method, your code must produce the objective type's indicator for an interval. You can calculate the value of this indicator or simply define it to be a specific value.

For more information about custom code, see "Custom Code for the Calculation Method" on page 221.

Missing Values

In Supplied Methods

All calculation methods ignore missing values. For example:

```
maximum (1, 3, ., 8) = 8
mean (1, 3, ., 8) = 4
minimum (1, 3, ., 8) = 1
percentile (using the 50th percentile (., 1, 2) = 1.5
redundant (97, ., 97) = 99.91
serial (99, ., 99) = 98.01
sum (1, 3, ., 8) = 12
weighted-mean
  (values: 1, 3, ., 8; with these corresponding weights: 1, ., 2, 3) = 6.25
```

Note: Throughout this document, a period (.) is used to represent missing values. △

In the example above, there is a missing value for both measure and weight, but not for the same observation. *Weights accumulate only if the indicator value is non-missing*.

Therefore, the sum of the weighted sums is 25 and the sum of the weights is 4; the weighted mean is 6.25.

If the input to a calculation method consists entirely of missing values (standard or special), the calculation method returns the standard missing value. For example:

```
maximum (., .S, .U, .X, .I) = .
mean (., .S, .U, .X, .I) = .
minimum (., .S, .U, .X, .I) = .
percentile (., .S, .U, .X, .I) = .
redundant (., .S, .U, .X, .I) = .
serial (., .S, .U, .X, .I) = .
sum (., .S, .U, .X, .I) = .
weighted-mean (values: ., .S, .U, .I; corresponding weights: 1, 2, ., 3) = .
```

In Custom Code

You can use the same conventions as for supplied methods, or not.

Recommendations

- For an availability objective that is neither directly nor indirectly based on scores, if the relation of the resources or child nodes is redundant, it is strongly recommended that you choose redundant.
- For an availability objective that is neither directly nor indirectly based on scores, if the relation of the resources or child nodes is serial, it is strongly recommended that you choose serial.
- For a response time, throughput, or custom objective whose resource or child node inputs are rates or percentages, it is strongly recommended that you choose weighted mean instead of mean. The weighted mean calculation method is recommended as long as you have a value at the component objective level that functions as a weight.

  For example, if your measure is a value that represents the percent of successful transactions, then you would want to use a Weight By variable that represents the total number of transactions for the interval that is being reported.
- If you want to emphasize best or worst cases, consider using maximum or minimum, as appropriate.

  Or you might want to use a low or high percentile (say, 95th or 5th) if you have a large number of input values and you want to lean in the direction of best case or worst case, but do not want to take the most extreme value (which might be an outlier or error).
- If you want a method that is less sensitive to extreme values, consider using mean or 50th percentile (median), as appropriate.
- Be consistent. For example, if your choice for one calculation is to emphasize the best case, consider emphasizing the best case in other calculations.
- Use custom code when none of the supplied calculation methods meets your needs and you are calculating an availability, response time, throughput, or custom objective indicator.

APPENDIX

2

Custom Code for the Calculation Method

Custom Code for the Calculation Method

Any particular custom code is specific to the objective for which it is defined.

Custom code represents a custom calculation method that operates in a way that is similar to a supplied calculation method. For example, for an objective on a component that uses the custom calculation method, the following statements apply:

- The custom calculation method might (but is not required to) use as input the same one-value-per-resource that a supplied calculation method uses.
- The custom calculation method must assign a value that will be used as that objective's indicator.

For more information about all calculation methods, see "Calculation Methods" on page 215.

The custom code runs one time for each five-minute interval. For more information about intervals, see "Indicators" on page 15.

Input Variables

For an Objective on a Component Node

You can refer to the following variables and arrays for input in your code. (However, you are not required to do so. Your code can use some or all or none of these variables and arrays.)

Table A2.1 API Input Variables for Component Nodes

Variable Name	Contents
COUNT	The count of events in the interval (for response time, throughput, and custom objectives only).
DATETIME	The date and time of the start of this interval.
DURATION	The length of this interval (300 sec).
OUTAGE	Outage type: □ NONE - not an outage □ SCHEDULED - Scheduled Outage □ EXCUSED - Excused Outage □ UNSCHEDULED - non-excused (user-entered) outage.
AVAILABILITY_OBJECTIVE_ENABLED	Is the availability objective enabled for this node? (Y|N)
RESPONSETIME_OBJECTIVE_ENABLED	Is the response time objective enabled for this node? (Y|N)
THROUGHPUT_OBJECTIVE_ENABLED	Is the throughput objective enabled for this node? (Y|N)
CUSTOM_OBJECTIVE_ENABLED	Is the custom objective enabled for this node? (Y|N)
RESOURCE_COUNT Note: This value is also available in a macro variable named RESOURCE_COUNT.	The number of resources. Any number from 1 through the value of RESOURCE_COUNT can be used as a valid subscript into an array.
first resource identification column for this node[] (See note below this table.)	The value in this column for this resource.
...	...
last resource identification column for this node[]	The value in this column for this resource.
AVAILABILITY_INDICATOR[]	The availability indicator for this resource.
RESPONSETIME_INDICATOR[]	The response time indicator for this resource.

Variable Name	Contents
RESPONSETIME_IS_EVENT_BASED	Indicates whether the objective is extracted from event-based data. (Y|N)
RESPONSETIME_TOTAL_EVENT_COUNT[]	If this objective is event-based, contains an array of the total count of events per interval for each of the resources for this objective.
RESPONSETIME_WEIGHT[]	The response time weight for this resource.
THROUGHPUT_INDICATOR[]	The throughput indicator for this resource.
THROUGHPUT_IS_EVENT_BASED	Indicates whether the objective is extracted from event-based data. (Y|N)
THROUGHPUT_TOTAL_EVENT_COUNT[]	If this objective is event-based, contains an array of the total count of events per interval for each of the resources for this objective.
THROUGHPUT_WEIGHT[]	The throughput weight for this resource.
CUSTOM_INDICATOR[]	The custom indicator for this resource.
CUSTOM_IS_EVENT_BASED	Indicates whether the objective is extracted from event-based data. (Y|N)
CUSTOM_TOTAL_EVENT_COUNT[]	If this objective is event-based, contains an array of the total count of events per interval for each of the resources for this objective.
CUSTOM_WEIGHT[]	The custom weight for this resource.

Note: For more information about outages, see "What is an Outage?" on page 231. △

Note: The first resource identification column for this node[], ..., the last resource identification column for this node[]. Here is an example of how these columns can be used:

Suppose the resource identification columns for the node are DOMAIN and MACHINE and suppose they have the following values.

```
   RESOURCE
DOMAIN  MACHINE

domain1 machine1
domain1 machine2
domain1 machine3
```

Then, you will have an array named DOMAIN[] (with values domain1, domain1, domain1) and an array named MACHINE (with values machine1, machine2, machine3). △

Note: For more information about the values in AVAILABILITY_INDICATOR [], RESPONSETIME_INDICATOR [], RESPONSETIME_WEIGHT [], THROUGHPUT_INDICATOR [], THROUGHPUT_WEIGHT [], CUSTOM_INDICATOR [], and CUSTOM_WEIGHT [], see "One Value per Resource per Objective" on page 26. △

Note: Your code for a given objective is not restricted to using input values that are directly related to the objective. For example, your code for a response time objective might have some reason to use the availability, throughput, or custom objective values. △

For an Objective on a Higher Level Node

A higher level node can be a Service node, a Service level Agreement (SLA) node, or a Contract node. You can refer to the following variables and arrays for input in your code. However, you are not required to refer to them. Your code can use some or all or none of these variables and arrays.

Table A2.2 API Input Variables for Service Node, Service Level Agreement Node, Contract Node

Variable Name	Contents
DATETIME	The date and time of the start of this interval.
DURATION	The length of this interval (300 sec).
OUTAGE	Outage type: □ NONE - not an outage □ SCHEDULED - Scheduled Outage □ EXCUSED - Excused Outage □ UNSCHEDULED - non-excused (user-entered) outage.
AVAILABILITY_OBJECTIVE_ENABLED	Is the availability objective enabled for this node? (Y|N)
RESPONSETIME_OBJECTIVE_ENABLED	Is the response time objective enabled for this node? (Y|N)
THROUGHPUT_OBJECTIVE_ENABLED	Is the throughput objective enabled for this node? (Y|N)
CUSTOM_OBJECTIVE_ENABLED	Is the custom objective enabled for this node? (Y|N)
RESOURCE_COUNT	The number of child nodes.
Note: This value is also available in a macro variable named RESOURCE_COUNT.	Any number from 1 through the value of RESOURCE_COUNT can be used as a valid subscript into an array.
NODENAMES[]	The name of each child node.
AVAILABILITY_OBJECTIVES_ENABLED	Number of child nodes for which the availability objective is enabled.
RESPONSETIME_OBJECTIVES_ENABLED	Number of child nodes for which the response time objective is enabled.
THROUGHPUT_OBJECTIVES_ENABLED	Number of objectives for which the throughput objective is enabled.
CUSTOM_OBJECTIVES_ENABLED	Number of objectives for which the custom objective is enabled.
COMPOSITE_OBJECTIVES_ENABLED	Number of objectives for which the composite objective is enabled.
AVAILABILITY_INDICATOR[]	The availability indicator for this child node.
AVAILABILITY_SCORE[]	The availability score for this child node.

Variable Name	Contents
AVAILABILITY_TARGETMET[]	The availability target's status for this child node. (1=met, 0=missed)
RESPONSETIME_INDICATOR[]	The response time indicator for this child node.
RESPONSETIME_WEIGHT[]	The response time weight for this child node.
RESPONSETIME_SCORE[]	The response time score for this child node.
RESPONSETIME_TARGETMET[]	The response time target's status for this child node. (1=met, 0=missed)
THROUGHPUT_INDICATOR[]	The throughput indicator for this child node.
THROUGHPUT_WEIGHT[]	The throughput weight for this child node.
THROUGHPUT_SCORE[]	The throughput score for this child node.
THROUGHPUT_TARGETMET[]	The throughput target's status for this child node. (1=met, 0=missed)
CUSTOM_INDICATOR[]	The custom indicator for this child node.
CUSTOM_WEIGHT[]	The custom objective weight for this child node.
CUSTOM_SCORE[]	The custom score for this child node.
CUSTOM_TARGETMET[]	The custom target's status for this child node. (1=met, 0=missed)
COMPOSITE_INDICATOR[]	The composite indicator for this child node.
COMPOSITE_SCORE[]	The composite score for this child node.
COMPOSITE_TARGETMET[]	The composite target's status for this child node. (1=met, 0=missed)

Note: For more information about outages, see "What is an Outage?" on page 231. △

Note: For more information about the values in AVAILABILITY_INDICATOR [], RESPONSETIME_INDICATOR [], THROUGHPUT_INDICATOR [], CUSTOM_INDICATOR [], RESPONSETIME_WEIGHT [], THROUGHPUT_WEIGHT [], CUSTOM_WEIGHT [], AVAILABILITY_SCORE [], RESPONSETIME_SCORE [], THROUGHPUT_SCORE [], and CUSTOM_SCORE [], see "One Value per Child Node per Objective" on page 33. △

Note: In your code for a given objective, you are not restricted to using input values that are directly related to the objective. For example, in your code for a response time objective, you might have some reason to use the availability, throughput, or custom objective values. △

Output Variable

For an Availability Objective

If you use custom code for an availability objective, your code must assign a value to this variable:

Table A2.3 API Output Variable for an Availability Objective

Variable Name	Value will be used as
AVAILABILITY	This objective's availability indicator.

Note: For a component node, the value is normally greater than or equal to 0 and less than or equal to 100.

For a service, service level agreement, or contract node, the value is normally one of the following:

- greater than or equal to 0 and less than or equal to 100, if it is based on the child nodes' indicators (and these indicators are not based on scores, or based on indicators that are based on scores)
- greater than or equal to the lowest score and less than or equal to the highest score, if it is based on the child nodes' scores

△

For a Response Time Objective

If you use custom code for a response time objective, your code must assign a value to this variable:

Table A2.4 API Output Variable for a Response Time Objective

Variable Name	Value will be used as
RESPONSETIME	This objective's response time indicator.

For a Throughput Objective

If you use custom code for a throughput objective, your code must assign a value to this indicator:

Table A2.5 API Output Variable for a Throughput Objective

Variable Name	Value will be used as
THROUGHPUT	This objective's throughput indicator.

For a Custom Objective

If you use custom code for a custom objective, your code must assign a value to this indicator:

Table A2.6 API Output Variable for a Custom Objective

Variable Name	Value will be used as
CUSTOM	This objective's custom objective indicator.

The Custom Code

Your code must be written in SAS. It runs within a SAS DATA step prior to the implicit OUTPUT statement. Some statements should not be used in a SAS DATA step. For example, we recommend that you avoid SAS statements such as DELETE, ABORT, and OUTPUT. We also recommend that you avoid changing the data type of any variable (for example, from numeric to character).

Note: If an outage *is* in effect for a given node during a given interval, the outage takes precedence over the calculation of the node's indicators. In this case, the following statements apply:

- The indicators are not based on calculations. The calculation methods are ignored.
- The indicators are assigned a value based on the outage type. For more information, see "Effect of Outages on Data" on page 233.

If an outage *is not* in effect for a given node during a given interval, then for each objective, the following statments apply:

- If the objective is *active*, its indicator is calculated.
- If the objective is *inactive*, its indicator is set to .I, which is a special missing value.

△

Filtering Example

Suppose that for a component's availability indicator you want to use the calculation for serial contributors but ignore any data from Saturday and Sunday.

The following code shows how to do this.

```
if weekday(datepart(DATETIME)) in (1,7)
then
   AVAILABILITY = . ;
else do;
   AVAILABILITY = 100 ;
   do i = 1 to RESOURCE_COUNT ;
      AVAILABILITY = AVAILABILITY * AVAILABILITY_INDICATOR[i]/100 ;
   end;
end;
```

This code obtains the numeric day-of-week (1=Sunday, 2=Monday, ..., 7=Saturday) from the datetime of the interval.

- If the day-of-week is a Sunday or Saturday, the component's availability indicator is set to a standard missing value for a numeric, which SAS code represents as a decimal point.
- Otherwise, the code uses the calculation for serial contributors to produce the component's availability indicator.

  For more information about the calculation for serial contributors, see "Calculation Methods" on page 215.

Differential Emphasis Example

Suppose that you have a router service and it has two components: backbone routers and campus routers.

```
Router Service
   Backbone Component
   Campus Component
```

Suppose that for the service's response time indicator you want to use a calculation that is similar to the mean of the child node's response time scores, but you want to give more emphasis to the backbone routers than to the campus routers. For example, suppose that you want to give four times the emphasis to the backbone routers that you give to the campus routers.

Here is one way that you might write your code:

```
MySumOfScores = . ;

do i = 1 to RESOURCE_COUNT ;

   if NODENAME[i] =: 'Backbone'
   then
      emphasis = 0.8 ;
   else
      emphasis = 0.2 ;
   MySumOfScores =
```

```
      sum (MySumOfScores, emphasis*RESPONSETIME_SCORE[i]);

end;

if RESOURCE_COUNT gt 0 then do;
   RESPONSETIME = MySumOfScores / RESOURCE_COUNT ;
end;
else do;
   RESPONSETIME= . ;
end;
```

For each interval, the code performs the following actions:

1. It initializes the variable that will contain the running sum.
2. For each of the two components, the code performs these tasks:
   a. It assigns an emphasis of 0.8 if the name of the component contains the word "Backbone" and an emphasis of 0.2 otherwise.
   b. It updates the running sum.
3. The code calculates the mean (as sum/count) by dividing the final value of the running sum by the number of component nodes.

APPENDIX

3

Outages

What is an Outage?

This appendix contains the following information about outages. An outage is a time period during which a node in the contract hierarchy is not available. An outage can be specified for any type of node:

- a component
- a service
- a service level agreement
- a contract

An outage can be either scheduled, excused, or unscheduled.

- A *scheduled* outage is a predetermined and accepted time period, such as during a routine maintenance event, when a node (and its child nodes) in the contract hierarchy is expected to be unavailable. A scheduled outage reduces the amount of time that a service provider is accountable for providing service. Corresponding service level indicators during this time period are displayed as missing and are ignored. For example, down time for preventive maintenance might be a scheduled outage.
- An *excused* outage is a time period, such as during a power outage, when a node (and its child nodes) in the contract hierarchy was unavailable without fault from the service provider. An excused outage reduces the amount of time that a service provider is accountable for providing service. Corresponding service level indicators during this time period are displayed as missing and are ignored. An example of an excused outage might be the down time for a multiple-day power outage that is caused by a severe ice storm, if the backup generator does not have enough fuel to provide power for that many days.
- An *unscheduled* outage is a time period, such as during a server failure, when a node in the contract hierarchy is not available. This type of outage incurs a penalty value of zero for the availability objective of the affected nodes during the outage time period. For example, down time for a several-minute power outage that is caused by a windstorm might be an unscheduled outage if the backup

generator has enough fuel to provide power for at least a day but did not start. As another example, if the outage is not expected and the contract classifies all outages that are not scheduled as unscheduled outages, regardless of "acts of nature," then the outage is an unscheduled outage.

Note: For more information about how any type of outage affects the data, see "Effect of Outages on Data" on page 233.

For more information about intervals, see "Understanding What Is Common to All Objectives" on page 14. △

How Outages Affect Availability

The type of outage determines how availability is affected. SAS IT Service Level Management calculates availability as a percentage of time that is not covered by a scheduled or excused outage.

Scheduled and excused outages

In the five-minute intervals that occur during a scheduled or an excused outage, the indicators of the node for which the outage is specified (and of the nodes that are "below" this node in the hierarchy) are assigned missing values and the target status of their objectives is set to *met*.

A scheduled or an excused outage does not have a negative effect on availability with the following exception: *scheduled and excused outages reduce the possible time that the provider is responsible for availability.* Therefore, a node with any downtime that also has a scheduled or excused outage will have a downtime that is a larger percentage of the whole. For example, a node with a single 6-hour unscheduled outage for the day will result in an availability actual daily value of 75%. If instead, the node has a 12-hour scheduled outage, followed by a 6-hour unscheduled outage, then the actual daily value will be 50%. This is because the node was down 50% of the remaining 12 hours.

Unscheduled outages

Unscheduled outages affect the indicators because data cannot be collected during that time. In the five-minute intervals that occur during an unscheduled outage, the indicators of the node for which an unscheduled outage is specified (and of the nodes that are "below" this node in the hierarchy) are not assigned missing values and the target status of their objectives is set to *missed*. An unscheduled outage has a negative effect on availability.

Outage Start and End Times

The start time of an outage must be entered on a five-minute boundary: :00, :05, :10, :15, :20, :25, :30, :35, :40, :45, :50, or :55. The same is true for the end time of an outage. Thus, an outage covers one or more continuous five-minute intervals of time.

The outage starts at *(and includes)* the specified start date and start time. The outage ends at *(and includes)* the end date but *excludes* the end time. For example, an outage that starts at 11:00 on a given day and ends at 11:45 on the same day will cover any data starting at 11:00 up to but *not* including data with a datetime stamp of 11:45 or later.

Discuss with your customer what "rule of thumb" to use when entering an outage. For example, suppose an outage begins at 13:59 and ends at 14:12. You could enter the outage from 13:55 to 14:15 (in order to "round out" to the five–minute boundaries), from 14:00 to 14:10 (in order to "round in" to the boundaries), or from 14:00 to 14:15 (in order to more accurately represent the length). In the case of an unscheduled outage, you can choose not to enter the outage because the data naturally reflects this.

Regardless of the method used, note that outages are not applied to the data in the resource tables. Outages are taken into consideration whenever data is rolled up into the component, service, SLA, or contract nodes. Therefore, you might see differences between a component's actual daily values and a resource's indicator values.

Outage Priorities

For each five-minute time interval, determination of outage type is made in two steps.

1. *The ETL task searches for the nearest node that has one or more outages.* The search starts at the given node and goes upward in the hierarchy. If the given node has one or more outages, then the given node is the nearest node. If the given node has no outages but the node's parent node has one or more outages, then the node's parent node is the nearest node. And so on, moving up the hierarchy.
2. *The ETL task applies the nearest node's outage type to the given node.* If the nearest node has only one outage, then that outage type applies to the given node. If the nearest node has multiple outages, then the outage type with the highest priority takes precedence and applies to the given node.

   Here is the priority for each outage type:

   - Unscheduled (highest priority)
   - Excused
   - Scheduled (lowest priority)

Here are two examples of determining the outage type for a component node to which multiple outages apply:

- Suppose that a *scheduled* outage is specified for every third Tuesday from 02:00 to 03:00 and that there is an *excused* outage for that same component on a particular third Tuesday from 01:30 to 04:30 (say, for a one-time power outage).

  Then, in the component's main table in the contract database for this contract, an *excused* outage is represented in the data from 01:30 to 04:30.
- Suppose that a *scheduled* outage is specified for every third Tuesday from 01:30 to 04:30 and that there is an *excused* outage for that same component on a particular third Tuesday from 02:00 to 03:00 (say, for a one-time power outage).

  Then, in the component's main table in the contract database for this contract, a *scheduled* outage is represented in the data from 01:30 to 02:00, an *excused* outage is represented in the data from 02:00 to 03:00, and a *scheduled* outage is represented in the data from 03:00 to 04:30.

Effect of Outages on Data

When a node has an outage, then the contract database that is associated with the node's contract is affected in several places.

At the Detail Level

- *At the node for which the outage is specified:* In the node's table, all rows whose data occurred during the outage have their outage variable set to the name of the outage type and their variables for indicators (if active) are set to special missing values, as shown in the following table.

Table A3.1 Effect of Outage Type on Outage and Indicator Variables

Type of Outage	Value of the Outage Variable	Value of the Availability Indicator Variable	Value of the Response Time Indicator Variable	Value of the Throughput Indicator Variable	Value of the Custom Objective Indicator Variable
Scheduled	SCHEDULED	.S	.S	.S	.S
Excused	EXCUSED	.X	.X	.X	.X
Unscheduled	UNSCHEDULED	0	.U	.U	.U

Note: The value of the composite indicator is calculated and, depending on the calculation method, might be affected. For more information about calculation methods and how their results are affected by missing values, see "Calculation Methods" on page 215. △

If a range set is specified for a given objective and the objective is active, then, if they are defined, the score and the range description for that objective are those that are defined for the Missing range with the following exception. Availability for an unscheduled outage is zero, which is *not* in the Missing range.

If a target is specified for a given objective and the objective is active, then the target status and the target pass-fail reason for the objective are set as described in this table.

Table A3.2 Effect of Outage Type on Target Status and Target Pass-Fail Reason

Type of Outage	Target Met-or-Missed Status	Target Pass-Fail Reason
Scheduled	1 (met)	-1
Excused	1 (met)	-2
Unscheduled	0 (missed)	-3

Note: A component node's main table (Cnnnnnn table) *is* affected as described above, but its availability table (Annnnnn table), response time table (Rnnnnnn table), throughput table (Tnnnnnn table), and custom objective table (Unnnnnn table) *are not* affected.

Because the component's availability, response time, throughput, and custom objective tables are not affected, if the outage information changes, then the indicators can be recalculated. △

□ *At the nodes that are "below" (lower in the hierarchy than) the node at which the outage is specified:* The same changes take place as described above. That is, the child nodes inherit the outages of their parent nodes, regardless of how deep the hierarchy is.

Note: If the node at which the outage is specified is a component node, there are no child nodes, so only the component's main table (Cnnnnnn table) is affected. △

□ *At the nodes that are "above" (higher in the hierarchy than) the node at which the outage is specified:* The indicators are recalculated.

Note: The special missing values *can be* distinguished from the "." missing value.

- In the supplied calculation methods for indicator (maximum, mean, minimum, percentile, redundant, serial, sum, and weighted mean), special missing values *are not* distinguished from the "." missing value. All types of missing values are ignored in the calculation.

  Note: If the calculated value of the indicator is a missing value, then it is represented by a period (.). △

- In the custom calculation method for the indicator, *you can decide* whether to detect any special missing values and handle them differently or to handle them in the same way as the "." missing value.

△

If a range set is specified for a given objective, then the score and the range description for that objective are based on the recalculated indicator.

If a target is defined for the objective, the target status and the target pass-fail reason for the objective are based on the recalculated indicator or on the recalculated score.

For more information about the data in the detail level of the contract database, see "The Contract Database" on page 237.

At the Summary Levels (Day, Week, Month, and Year)

The summary statistics are recalculated if there are changes in the node's indicators at the detail level. For example, if an objective's indicator changes at the detail level, the mean of the objective's indicator is recalculated at the day, week, month, and year levels.

If the recalculation changes the mean for an objective's indicator and if a range set is defined for that objective, then the mean's score (if any) and the mean's range description (if any) automatically update to correspond to the new value of the mean.

If a target is defined for that objective, the target status for the objective automatically updates, based on the recalculated mean for the indicator or on the recalculated score.

For more information about the data in the summary levels of the contract database, see "The Contract Database" on page 237.

APPENDIX

4

The Contract Database

The Contract Database

Each contract that has been activated has its own contract database. In a contract database, the data is in these tables:

- “Detail Tables” on page 243
- “Daily Tables” on page 252
- “Weekly Tables” on page 262
- “Monthly Tables” on page 262
- “Yearly Tables” on page 263

Accessing Views of the Contract Database

When you are generating reports with SAS Enterprise Guide you will need to refer to specific tables within the contract database. In order to simplify the access to these tables, SAS IT Service Level Management provides *views* of the contract database. A view is a pointer to a group of the tables in the contract database. Views are recreated after every successful execution of the synchronization task of the ETL process.

The types of "views" that are available in the ADMIN library of the contract database are as follows:

- *Views for resources:*

  The data is merged together by datetime and resources into one view per component per level of the contract database. The views are as follows:

  - ADMIN.C*nnnnnn*_Det
  - ADMIN.C*nnnnnn*_Day
  - ADMIN.C*nnnnnn*_Wk
  - ADMIN.C*nnnnnn*_Mon
  - ADMIN.C*nnnnnn*_Yr

  where C*nnnnnn* is the system-generated identification number of the component. To determine the view that corresponds to a specific component, navigate to the ADMIN._ComponentLookupTable in the contract database. This table consists of two columns:

  - componentID — contains the identifier of the view, for example, C000013.
  - componentDescription — contains the name of the component as specified in the GUI, for example, Switch 4830.

  Note: Within the view for resources, the tables of resource data are still individually accessible. In other words, you can select any or all of the data for a specific resource.

  The data for each resource that is contained in the component can be identified in the following manner:

  - Availability data — variables are AVAIL, UNAVAIL, and UPTIME
  - Response time data — variables are prefixed by R
  - Throughput data — variables are prefixed by T
  - Custom objective data— variables are prefixed by U

  △

- *Views for nodes:*

  The objective type data for all of the nodes (component, service, SLA, and contract) are *concatenated* into one view per level of the contract database. The views are as follows:

  - ADMIN.DETAIL
  - ADMIN.DAY
  - ADMIN.WEEK
  - ADMIN.MONTH
  - ADMIN.YEAR

Note: Within the view for nodes, the tables of objective type data are concatenated together into one view per node within the contract. You can select any or all of the objective type data for a specific node. The data for each node can be identified by the prefixes:

- A — for availability data
- R — for response time data
- T — for throughput data
- U — for custom objective data
- Z — for composite data

△

Accessing the Metadata of the Contract Database

In order to provide a simplified method for navigating a contract's hierarchy and its associated objects, SAS IT Service Level Management enables you to refer to specific tables within a contract database by accessing the metadata of the contract database. This method is especially useful when you are generating reports with SAS Enterprise Guide, Structured Query Language (SQL), or other report writing tools. For an example of how you can use these tables of metadata, see "Example Using Contract Database Metadata" on page 242.

The tables are recreated in the ADMIN library of the contract database after every successful execution of the synchronization task of the ETL process. A contract database has the following ADMIN-level tables:

Structure	contains one row for each node in the contract hierarchy.
NodeObjective	contains one row for each activated objective for each non-deleted node in the contract.
RangeSets	contains one row for each range that is used in each rangeset for each activated objective for each non-deleted node in the contract.
NodeOutage	contains one row for each outage that is assigned to a given node.

The Structure Table

The Structure table contains the following elements:

Table A4.1 The Structure Table

Column Name	Description	Data Type
ID	The identification number of this element in the structure. The ID increases monotonically across all levels. (length: 8)	Numeric
Level_ID	The level of this node. The Contract level is 1. Each subsequent level of nodes is assigned the next number. The order of peers is indeterminiate. The combination of ID and Level ID can be used to order the nodes in a relative sense. (length: 8)	Numeric
Canonical_ID	A character string that defines the traversal of ID values that ends up at this node. For example, 001.002.003 refers to the third child node of the second child node of the contract node. Each distinct ID value is formatted with leading zeroes (Z3). (length: 200)	Character

Column Name	Description	Data Type
NodeType	Valid values are: □ CONTRACT □ AGREEMENT □ SERVICE □ COMPONENT The NodeType can be used to determine which data set in the environment database to query for additional attributes about this node. (length: 9)	Character
Node_ID	The ID of the node. It can be used to search for the specific node in the *SLM.<nodetype>*table. (length: 8)	Numeric
NodeName	The name of the node. (length: 40)	Character

Note:

- The contract node always has an ID of 1, a Level ID of 1, and a Canonical ID of 001.
- This data set can be input to the *com.sas.sasserver.tree.LevelTreeInterface* webAF component, using the ID and Level_ID columns.

△

The NodeObjective Table

The NodeObjective table contains the following elements:

Table A4.2 The NodeObjective Table

Column Name	Description	Data Type
ParentNode_ID	The identification number of the parent node. It can be used to search for the specific node in the *SLM.<ParentNodeType>* table. (length: 8)	Numeric
ParentNodeType	The type of the parent node. Valid values are: □ CONTRACT □ AGREEMENT □ SERVICE □ COMPONENT It can be used in conjunction with the ParentNode_ID column to query the NodeOutage table. (length: 9)	Character

Column Name	Description	Data Type
NodeType	The type of node. Valid values are: □ CONTRACT □ AGREEMENT □ SERVICE □ COMPONENT It can be used in conjunction with the Node_ID column to query the NodeOutage table. (length: 9)	Character
Node_ID	The identification number of the node. It can be used to search for the specific node in the *SLM.<NodeType>* table. (length: 8)	Numeric
Contact_ID	The identification number of the servicing contact. It can be used to search for the specific servicing provider in the *SLM.Contact* table. (length: 8)	Numeric
Objective_ID	The identification number of the objective. It can be used to search for the specific node in the *SLM.Objective* table. (length: 8)	Numeric
ObjectiveType	The type of the objective. Valid values are □ AVAILABILITY □ RESPONSETIME □ THROUGHPUT □ CUSTOM □ COMPOSITE (length: 12)	Character
RangeSet_ID	The identification number of the RangeSet that is associated with the objective. It can be used to search for the respective ranges and range descriptions in the *ADMIN.RANGESETS* table. (length: 8)	Numeric

The RangeSets Table

The RangeSets table contains the following elements:

Table A4.3 The RangeSets Table

Column Name	Description	Data Type
RangeSet_ID	The identification number of the RangeSet. It can be used to search for the specific RangeSet in the *SLM.RangeSet* table. (length: 8)	Numeric
Range_ID	The identification number of the Range within the RangeSet. It can be used to search for the specific RangeSet in the *SLM.Range* table. (length: 8)	Numeric
RangeDescription_ID	The identification number of the RangeDescription. It can be used to search for the specific RangeSet in the *SLM.RangeDescription* table. (length: 8)	Numeric

The NodeOutage Table

The NodeOutage table contains the following elements:

Table A4.4 The NodeOutage Table

Column Name	Description	Data Type
Node_ID	The identification number of the node.	Numeric
NodeType	The type of the node. Valid values are: □ CONTRACT □ AGREEMENT □ SERVICE □ COMPONENT (length: 9)	Character
Outage_ID	The identification number of the outage. It can be used to search for the specific outage in the *SLM.OUTAGE* table. (length: 8)	Numeric

Example Using Contract Database Metadata

This example produces an eight-column table that fully describes all target definitions for a specific node. In the example, the *findTargetsForNode* macro uses metadata from the contract database. The macro takes one numeric input (the

identifier of the node for which targets will be located) and extracts the requested information from three tables (the *ADMIN.NodeObjective* table, the *ADMIN.Structure* table, and the *SLM.Objective* table) in order to produce a table called *nodeTargets*.

```
%macro findTargetsForNode(nodeID);
proc sql;
create table work.nodeTargets as
  select objective.period            as TargetPeriod,
         objective.reportstatistic   as TargetStatistic,
         objective.targettype        as TargetType,
         objective.operator          as Operator,
         objective.value             as TargetValue,
         adminNodeObj.objectivetype  as ObjectiveType,
         adminNodeObj.nodetype       as NodeType,
         adminStructure.nodename     as NodeName
  from   slm.objective               as objective,
         admin.nodeobjective         as adminNodeObj,
          admin.structure             as adminStructure
  where  adminNodeObj.objective_ID eq objective.ID and
         adminNodeObj.node_ID      eq &nodeID and
         adminStructure.node_ID    eq &nodeID
  order by NodeName;
quit;
%mend findTargetsForNode;
```

Note: Prior to using the *findTargetsForNode* macro, allocate the environment database and the contract database. You can perform this allocation by calling the %CPSLMENV and %CPSLMCON macros. For information about these macros, see "%CPSLMENV" on page 266 and "%CPSLMCON" on page 271 . △

Detail Tables

A contract database has the following detail tables:

- one detail table for the contract node
- one detail table for each service level agreement node
- one detail table for each service node
- one main detail table for each component node
- four underlying detail tables (the availability table, the response time table, the throughput table, and the custom objective table) for each component node

Note: For more information about the structure of the detail tables, see "Contract, Service Level Agreement, Service, and (Main) Component Tables" on page 244.

For more information about the structure of the availability table, see "Availability Table for Component Node" on page 249.

For more information about the structure of the response time, throughput, and custom objective tables, see "Response Time, Throughput, and Custom Objective Tables for a Component Node" on page 250. △

Contract, Service Level Agreement, Service, and (Main) Component Tables

A table is named according to the following standard:

- N*nnnnnn*, for a contract node
- L*nnnnnn*, for a service level agreement node
- S*nnnnnn*, for a service node
- C*nnnnnn*, for a component node

where *nnnnnn* is the node number, which is unique within the node type across all contracts in the same environment database. In the GUI, you can see the table names in the *Contract* workspace by double-clicking the contract. The name of the contract node's table is in the lower-left corner of the **`General`** tab. To see the names of the other nodes' tables, select the **`Hierarchy`** tab and double-click the node whose table name you want to see. The name of the node's table is in the lower-left corner of the *General* tab.

Note: You can also use the views that are defined in the ADMIN level of the contract database to refer to the contract node (which combines all the component, service, SLA, and contract nodes). Similarly, you can access resource tables by using views. For more information, see "Accessing Views of the Contract Database" on page 237. △

For every date in the table, there can be a maximum of 288 values of DATETIME (one every five minutes), and for each of these values there is one row. The rows in this table are sorted in ascending order by the following field:

```
DATETIME
```

The row has the columns that are shown in the following table:

Table A4.5 Row in a Contract, Service Level Agreement, Service, or Component Node's (Main) Detail Table

Column Name	Description	Data Type
DATETIME	Datetime of the beginning of this interval.	Numeric
DURATION	Length of this interval, in seconds. .(300).	Numeric
HOUR	Hour part of DATETIME.	Numeric
DATE	Date part of DATETIME.	Numeric
TIME	Time part of DATETIME.	Numeric
SHIFT	SAS IT Resource Management SHIFT setting for this DATETIME. (length: 1 character)	Character
About node:		
NODTYPE	The type of this node: □ CONTRACT □ AGREEMENT □ SERVICE □ COMPONENT	Character
NTABLE	The table name for this node. (length: 7 characters)	Character

Column Name	Description	Data Type
PARENT	The name of this node's parent node. (length: 40 characters)	Character
PTABLE	The name of the table for this node's parent node: N*nnnnnn*, for a contract node L*nnnnnn*, for a service level agreement node S*nnnnnn*, for a service node C*nnnnnn*, for a component node where *nnnnnn* is the node number	Character

About outage:

SCHDUR	Amount of DURATION that is included in a scheduled outage during this five-minute interval, in seconds.	Numeric
EXCDUR	Amount of DURATION that is included in an excused outage during this five-minute interval, in seconds.	Numeric
UNSCDUR	Amount of DURATION that is included in an unscheduled outage during this five-minute interval, in seconds.	Numeric
SLMDUR	Amount of DURATION excluding scheduled and excused outages during this five-minute interval, in seconds.	Numeric

About the availability objective:

ASCORE	The score, from the range set for the availability objective, based on the above ASLI value.	Numeric
ACOLOR	The six-character RGB color, from the range set for the availability objective, based on the above ASLI value.	Character
ADESC	The description, from the range set for the availability objective, based on the above ASLI value. (length: max 80 characters)	Character
ARPTVAL	Value (for reports) of the **`Comparison Value`** field in the **`Target Definition`** section of the **`Availability`** objective tab.	Numeric
ATGTBAS	The availability objective's target is based on the following: SLI = the indicator SCORE = the score NONE = target not fully specified	Character
ATARGET	The availability objective's target value.	Numeric
ATGTOK	The target status for the availability objective, based on the above ASLI value: 1=met 0=missed	Numeric

Column Name	Description	Data Type
ATGTPFR	The target pass-fail reason for the availability objective, based on the above ASLI value: 1 = passed target criteria 0 = failed target criteria -1 = passed due to scheduled outage -2 = passed due to excused outage -3 = failed due to unscheduled outage	Numeric

About the response time objective:

RWEIGHT	The sum of the values of the response time weight. For supplied calculation methods: for a component node, based on WEIGHT values from component's detail table for response time; for the other types of nodes, based on RWEIGHT values from the main detail tables for this node's child nodes.	Numeric
RSCORE	The score, from the range set for the response time objective, based on the above value of RSLI.	Numeric
RCOLOR	The six-character RGB color, from the range set for the response time objective, based on the above value of RSLI.	Character
RDESC	The description, from the range set for the response time objective, based on the above value of RSLI.	Character
RRPTVAL	Value (for reports) of the **`Comparison Value`** field in the **`Target Definition`** section of the **`Response Time`** objective tab.	Numeric
RTGTBAS	The response time objective's target is based on: SLI = the indicator SCORE = the score NONE = target not fully specified	Character
RTARGET	The response time objective's target value.	Numeric
RTGTOK	The target status for the response time objective, based on the above value of RSLI: 1 = met 0 = missed	Numeric
RTGTPFR	The target pass-fail reason for the response time objective, based on the above value of RSLI: 1 = passed target criteria 0 = failed target criteria -1 = passed due to scheduled outage -2 = passed due to excused outage -3 = failed due to unscheduled outage	Numeric

About the throughput objective:

Column Name	Description	Data Type
TWEIGHT	The sum of the values of the throughput weight. For supplied calculation methods: for a component node, based on WEIGHT values from component's detail table for throughput; for the other types of nodes, based on TWEIGHT values from the main detail tables for this node's child nodes.	Numeric
TSCORE	The score, from the range set for the throughput objective, based on the above value of TSLI.	Numeric
TCOLOR	The six-character RGB color, from the range set for the throughput objective, based on the above table's value of TSLI.	Character
TDESC	The description, from the range set for the throughput objective, based on the above table's value of TSLI.	Character
TRPTVAL	Value (for reports) of the **`Comparison Value`** field in the **`Target Definition`** section of the **`Throughput`** objective tab.	Numeric
TTGTBAS	The throughput objective's target is based on: SLI = the indicator SCORE = the score NONE = target not fully specified	Character
TTARGET	The throughput objective's target value.	Numeric
TTGTOK	The target status for the throughput objective, based on the above value of TSLI: 1 = met 0 = missed	Numeric
TTGTPFR	The target pass-fail reason for the throughput objective, based on the above value of TSLI: 1 = passed target criteria 0 = failed target criteria -1 = passed due to scheduled outage -2 = passed due to excused outage -3 = failed due to unscheduled outage	Numeric

About the custom objective:

UWEIGHT	The sum of the values of the custom objective weight. For supplied calculation methods: for a component node, based on WEIGHT values from component's detail table for the custom objective; for the other types of nodes, based on UWEIGHT values from the main detail tables for this node's child nodes.	Numeric
USCORE	The score from the range set for the custom objective, based on the above value of USLI.	Numeric
UCOLOR	The six-character RGB color, from the range set for the custom objective, based on the above value of USLI.	Character

Column Name	Description	Data Type
UDESC	The description, from the range set for the custom objective, based on the above value of USLI.	Character
URPTVAL	Value (for reports) of the **Comparison Value** field in the **Target Definition** section of the **Custom Objective** tab.	Numeric
UTGTBAS	The custom objective's target is based on: SLI = the indicator SCORE = the score NONE = target not fully specified	Character
UTARGET	The custom objective's target value.	Numeric
UTGTOK	The target status for the custom objective, based on the above value of USLI: 1 = met 0 = missed	Numeric
UTGTPFR	The target pass-fail reason for the custom objective, based on the above value of USLI: 1 = passed target criteria 0 = failed target -1 = passed due to scheduled outage -2 = passed due to excused outage -3 = failed due to unscheduled outage	Numeric

About the composite objective:

ZSCORE	The score, from the range set for the composite objective, based on the above value of ZSLI.	Numeric
ZCOLOR	The six-character RGB color, from the range set for the composite objective, based on the above value of ZSLI.	Character
ZDESC	The description, from the range set for the composite objective, based on the above value of ZSLI.	Character
ZRPTVAL	Value (for reports) of the **Composite Value** field in the **Target Definition** section of the **Composite** objective tab.	Numeric
ZTGTBAS	The composite objective's target is based on: SLI = the indicator SCORE = the score NONE = target not fully specified	Character
ZTARGET	The composite objective's target value.	Numeric

Column Name	Description	Data Type
ZTGTOK	The target status for the composite objective, based on the above value of ZSLI: 1 = met 0 = missed	Numeric
ZTGTPFR	The target pass-fail reason for the composite objective, based on the above value of ZSLI: 1 = passed target criteria 0 = failed target criteria -1 = passed due to scheduled outage -2 = passed due to excused outage -3 = failed due to unscheduled outage	Numeric

Availability Table for Component Node

The name of the table is

□ A*nnnnnn*

where *nnnnnn* is the node number.

For every date in the table, there are 288 values of DATETIME (one every five minutes), and for each of these values there is one row per resource. The rows in this table are sorted in ascending order by the following fields:

```
DATETIME
first-resource-identification-column
...
last-resource-identification-column
```

The row has the columns that are shown in the following table.

Table A4.6 Row in a Component's Detail Table for Availability

Column Name	Description	Data Type
DATETIME	Datetime of the beginning of this interval.	Numeric
DURATION	Length of this interval, in seconds. (300)	Numeric
HOUR	Hour part of DATETIME.	Numeric
DATE	Date part of DATETIME.	Numeric
TIME	Time part of DATETIME.	Numeric
SHIFT	SAS IT Resource Management SHIFT setting for this DATETIME. (length: 1 character)	Character

About resource:

...	...	...

Column Name	Description	Data Type
Last resource-identification column	Resource's value for this column.	depends on column
UPTIME	The amount of time during this interval that this resource was available, in seconds.	Numeric
AVAIL	Percent of time during this interval that this resource was available. From UPTIME and DURATION, above: ((UPTIME / DURATION)*100)	Numeric
UNAVAIL	Percent of time during this interval that this resource was unavailable. From AVAIL above: (100-AVAIL)	Numeric

Note: UPTIME is measured in seconds. If missing, UPTIME is treated as zero seconds of uptime.

In addition, the availability of a component must always expressed as a value greater than or equal to zero. If the measure that you have chosen for your availability objective has a negative value in your source PDB, then the availability indicator for the component will be set to zero in the contract PDB. This prevents a component from having a negative availability. △

Response Time, Throughput, and Custom Objective Tables for a Component Node

The name of the table is

- R*nnnnnn* for a response time table
- T*nnnnnn* for a throughput table,
- U*nnnnnn* for a custom objective table

where *nnnnnn* is the node number, which is unique within the node type across all contracts in the same environment database.

For every date in the table, there are 288 values of DATETIME (one every five minutes), and for each of these values there is one row per resource. The row has the columns that are shown in the following table. The rows in these tables are sorted in ascending order by the following fields:

```
DATETIME
first-resource-identification-column
...
last-resource-identification-column
```

Table A4.7 Row in a Component's Detail Table for Response Time, Throughput or Custom Objective

Column Name	Description	Data Type
DATETIME	Datetime of the beginning of this interval.	Numeric
DURATION	Length of this interval, in seconds. (300)	Numeric

Column Name	Description	Data Type
HOUR	Hour part of DATETIME.	Numeric
DATE	Date part of DATETIME.	Numeric
TIME	Time part of DATETIME.	Numeric
SHIFT	SAS IT Resource Management SHIFT setting for this DATETIME. (length: 1 character)	Character
About resource:		
...	...	...
Last resource-identification column	Resource's value for this column.	depends on column
COUNT	For this resource and this interval, the count of events in the interval. If the data was not extracted from an EVENT table, then this is set to missing.	Numeric
MAX	For this resource and this interval, the maximum value for the measure in this component's data source PDB for this objective.	Numeric
MIN	For this resource and this interval, the minimum value for the measure in this component's data source PDB for this objective.	Numeric
SUM	For this resource and this interval, the sum for the measure.That is, the sum of the values for the measure in this component's data source PDB for this objective.	Numeric
WSUM	For this resource and this interval, the weighted sum for the measure. That is, the sum of the (value-of-the-weight*value-of-the-measure) products in this component's data source PDB for this objective.	Numeric
MEAN	For this resource and this interval, the arithmetic mean of the values for the measure in this component's data source PDB for this objective; based on SUM, above: SUM/number of resources that do not have missing measure values.	Numeric
WEIGHT	If the corresponding measure is not missing, then for this resource and this interval, the sum of the values of the weight in this component's data source PDB for this objective.	Numeric

Column Name	Description	Data Type
WMEAN	For this resource and this interval, the weighted arithmetic mean of the values for the measure in this component's data source PDB for this objective; based on WSUM and WEIGHT, above: (WSUM/WEIGHT).	Numeric
SLI	Same value as the MAX, MIN, SUM, MEAN, or WMEAN above based on the calculation method selected for this objective of this component.	Numeric

For more information about COUNT, MAX, MIN, SUM, WSUM, MEAN, WEIGHT, WMEAN, and SLI, see "One Value per Resource per Objective" on page 26 and "Calculation Methods" on page 215.

Daily Tables

A contract database has these tables:

- one daily table for the contract node
- one daily table for each service level agreement node
- one daily table for each service node
- one daily table for each component node
- four underlying daily tables (the availability table, the response time table, the throughput table, and the custom objective table) for each component node

Note:

- For information about the structure of these daily tables, see "Contract, Service Level Agreement, Service, and (Main) Component Tables" on page 252.
- For information about the structure of the availability table, see "Availability Table for a Component Node" on page 258.
- For information about the structure of the response time, throughput, and custom objective tables, see "Response Time, Throughput, and Custom Objective Tables for a Component Node" on page 260.

△

The values in the daily tables are based on the values in the detail tables.

Contract, Service Level Agreement, Service, and (Main) Component Tables

The values in this table are based on the values from the node's (main) *detail* table. The rows in this table are sorted in ascending order by the following field:

```
DATETIME
```

Table A4.8 A Contract, Service Level Agreement, Service, or Component's (Main) Daily Table

Column Name	Description	Data Type
DATETIME	Datetime of the beginning of this day.	Numeric
DURATION	Sum of DURATIONs for all five-minute intervals for this day in this node's main detail table.	Numeric
DATE	DAte part of DATETIME.	Numeric
About node:		
NODTYPE	The type of this node: □ CONTRACT □ AGREEMENT □ SERVICE □ COMPONENT	Character
NTABLE	The table name of this node. (length: 7 characters)	Character
PARENT	The name of this node's parent node. (length: 40 characters)	Character
PTABLE	The name of the table for this node's parent node: N*nnnnnn*, for a contract node L*nnnnnn*, for a service level agreement node S*nnnnnn*, for a service node C*nnnnnn*, for a component node where *nnnnnn* is the node number	Character
About outage:		
SCHDUR_C	Count of five-minute SCHDUR values for this day in this node's main detail table.	Numeric
SCHDUR_D	Standard deviation of five-minute SCHDUR values for this day in this node's main detail table.	Numeric
SCHDUR_N	Minimum of five-minute SCHDUR values for this day in this node's main detail table.	Numeric
SCHDUR_S	Sum of five-minute SCHDUR values for this day in this node's main detail table.	Numeric
SCHDUR_X	Maximum of five-minute SCHDUR values for this day in this node's main detail table.	Numeric
SCHDUR_Y	Range of five-minute SCHDUR values for this day in this node's main detail table.	Numeric
EXCDUR	Arithmetic mean of five-minute EXCDUR values for this day in this node's main detail table.	Numeric
EXCDUR_C	Count of five-minute EXCDUR values for this day in this node's main detail table.	Numeric

Column Name	Description	Data Type
EXCDUR_D	Standard deviation of five-minute EXCDUR values for this day in this node's main detail table.	Numeric
EXCDUR_N	Minimum of five-minute EXCDUR values for this day in this node's main detail table.	Numeric
EXCDUR_S	Sum of five-minute EXCDUR values for this day in this node's main detail table.	Numeric
EXCDUR_X	Maximum of five-minute EXCDUR values for this day in this node's main detail table.	Numeric
EXCDUR_Y	Range of five-minute EXCDUR values for this day in this node's main detail table.	Numeric
UNSCDUR	Arithmetic mean of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
UNSCDURC	Count of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
UNSCDURD	Standard deviation of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
UNSCDURN	Minimum of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
UNSCDURS	Sum of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
UNSCDURX	Maximum of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
UNSCDURY	Range of five-minute UNSCHDUR values for this day in this node's main detail table.	Numeric
SLMDUR	Arithmetic mean of five-minute SLMDUR values for this day in this node's main detail table.	Numeric
SLMDUR_C	Count of five-minute SLMDUR values for this day in this node's main detail table.	Numeric
SLMDUR_D	Standard deviation of five-minute SLMDUR values for this day in this node's main detail table.	Numeric
SLMDUR_N	Minimum of five-minute SLMDUR values for this day in this node's main detail table.	Numeric
SLMDUR_S	Sum of five-minute SLMDUR values for this day in this node's main detail table.	Numeric
SLMDUR_X	Maximum of five-minute SLMDUR values for this day in this node's main detail table.	Numeric
SLMDUR_Y	Range of five-minute SLMDUR values for this day in this node's main detail table.	Numeric

About the availability objective:

ASLI___C	Count of five-minute ASLI values for this day in this node's main detail table.	Numeric

Column Name	Description	Data Type
ASLI___D	Standard deviation of five-minute ASLI values for this day in this node's main detail table.	Numeric
ASLI___N	Minimum of five-minute ASLI values for this day in this node's main detail table.	Numeric
ASLI___S	Sum of five-minute ASLI values for this day in this node's main detail table.	Numeric
ASLI___X	Maximum of five-minute ASLI values for this day in this node's main detail table	Numeric
ASLI___Y	Range of five-minute ASLI values in this node's main detail table for this day in this node's main detail table.	Numeric
ASCORE	The score, from the range set for the availability objective, based on the above ASLI value.	Numeric
ACOLOR	The six-character RGB color, from the range set for the availability objective, based on the above ASLI value.	Numeric
ADESC	The description, from the range set for the availability objective, based on the above ASLI value.	Character
ATGTBAS	The availability objective's target is based on: SLI = the indicator SCORE = the score NONE = target not enabled	Character
ATARGET	The availability objective's target value.	Numeric
ATGTOK	The target status for the availability objective, based on the above ASLI value.	Numeric
ATGTPFR	The target pass-fail reason for the availability objective, based on the above ASLI value.	Numeric

About the response time objective:

RSLI___C	Count of five-minute RSLI values for this day in this node's main detail table.	Numeric
RSLI___D	Standard deviation of five-minute RSLI values for this day in this node's main detail table.	Numeric
RSLI___N	Minimum of five-minute RSLI values for this day in this node's main detail table.	Numeric
RSLI___S	Sum of five-minute RSLI values for this day in this node's main detail table.	Numeric
RSLI___W	If and only if the selected calculation method for response time is weighted mean, the weighted sum of five-minute RSLI values for this day in this node's main detail table; that is, the sum of the RSLI*RWEIGHT products for this day in this node's main detail table. Otherwise, this column does not exist.	Numeric

Column Name	Description	Data Type
RSLI___X	Maximum of five-minute RSLI values for this day in this node's main detail table.	Numeric
RSLI___Y	Range of five-minute RSLI values for this day in this node's main detail table.	Numeric
RWEIGHT	Arithmetic mean of five-minute RWEIGHT values for this day in this node's main detail table.	Numeric
RWEIGHTS	Sum of five-minute RWEIGHT values for this day in this node's main detail table.	Numeric
RSCORE	The score, from the range set for the response time objective, based on the above RSLI value.	Numeric
RCOLOR	The six-character RGB color, from the range set for the response time objective, based on the above RSLI value.	Numeric
RDESC	The description, from the range set for the response time objective, based on the above RSLI value.	Character
RTGTBAS	The response time objective's target is based on: SLI = the indicator SCORE = the score NONE = target not enabled	Character
RTARGET	The response time objective's target value.	Numeric
RTGTOK	The target status for the response time objective, based on the above RSLI value.	Numeric
RTGTPFR	The target pass-fail reason for the response time objective, based on the above RSLI value.	Numeric

About the throughput objective:

TSLI___C	Count of five-minute TSLI values for this day in this node's main detail table.	Numeric
TSLI___D	Standard deviation of five-minute TSLI values for this day in this node's main detail table.	Numeric
TSLI___N	Minimum of five-minute TSLI values for this day in this node's main detail table.	Numeric
TSLI___S	Sum of five-minute TSLI values for this day in this node's main detail table.	Numeric
TSLI___W	If and only if the selected calculation method for throughput is weighted mean, the weighted sum of five-minute TSLI values for this day in this node's main detail table; that is, the sum of the TSLI*TWEIGHT products for this day in this node's main detail table.	Numeric
TSLI___X	Maximum of five-minute TSLI values for this day in this node's main detail table.	Numeric
TSLI___Y	Range of five-minute TSLI values for this day in this node's main detail table.	Numeric

Column Name	Description	Data Type
TWEIGHT	Arithmetic mean of five-minute TWEIGHT values for this day in this node's main detail table.	Numeric
TWEIGHTS	Sum of five-minute TWEIGHT values for this day in this node's main detail table.	Numeric
TSCORE	The score, from the range set for the throughput objective, based on the above TSLI value.	Numeric
TCOLOR	The six-character RGB color, from the range set for the throughput objective, based on the above TSLI value.	Numeric
TDESC	The description, from the range set for the throughput objective, based on the above TSLI value.	Character
TTGTBAS	The throughput objective's target is based on: SLI = the indicator SCORE = the score NONE = target not enabled	Character
TTARGET	The throughput objective's target value.	Numeric
TTGTOK	The target status for the throughput objective, based on the above TSLI value.	Numeric
TTGTPFR	The target pass-fail reason for the throughput objective, based on the above TSLI value.	Numeric

About the custom objective:

USLI__C	Count of five-minute USLI values for this day in this node's main detail table.	Numeric
USLI__D	Standard deviation of five-minute USLI values for this day in this node's main detail table.	Numeric
USLI__N	Minimum of five-minute USLI values for this day in this node's main detail table.	Numeric
USLI__S	Sum of five-minute USLI values for this day in this node's main detail table.	Numeric
USLI__W	If and only if the selected calculation method for the custom objective is weighted mean, the weighted sum of five-minute USLI values for this day in this node's main detail table; that is, the sum of the USLI*UWEIGHT products for this day in this node's main detail table.	Numeric
USLI__X	Maximum of five-minute USLI values for this day in this node's main detail table.	Numeric
USLI__Y	Range of five-minute USLI values for this day in this node's main detail table.	Numeric
UWEIGHT	Arithmetic mean of five-minute UWEIGHT values for this day in this node's main detail table.	Numeric
UWEIGHTS	Sum of five-minute UWEIGHT values for this day in this node's main detail table.	Numeric

Column Name	Description	Data Type
USCORE	The score, from the range set for the custom objective, based on the above USLI value.	Numeric
UCOLOR	The six-character RGB color, from the range set for the custom objective, based on the above USLI value.	Numeric
UDESC	The description, from the range set for the custom objective, based on the above USLI value.	Character
UTGTBAS	The custom objective's target is based on: SLI = the indicator SCORE = the score NONE = target not fully specified	Character
UTARGET	The custom objective's target value.	Numeric
UTGTOK	The target status for the custom objective, based on the above USLI value.	Numeric
UTGTPFR	The target pass-fail reason for the custom objective, based on the above USLI value.	Numeric

About the composite objective:

ZSCORE	The score, from the range set for the composite objective, based on the above ZSLI value.	Numeric
ZCOLOR	The six-character RGB color, from the range set for the composite objective, based on the above ZSLI value.	Numeric
ZDESC	The description, from the range set for the composite objective, based on the above ZSLI value.	Character
ZTGTBAS	The composite objective's target is based on: SLI = the indicator SCORE = the score NONE = target not enabled	Character
ZTARGET	The composite objective's target value.	Numeric
ZTGTOK	The target status for the composite objective, based on the above ZSLI value.	Numeric
ZTGTPFR	The target pass-fail reason for the composite objective, based on the above ZSLI value.	Numeric

Availability Table for a Component Node

The rows in this table are sorted in ascending order by the following fields:

```
DATETIME
first-resource-identification-column
...
last-resource-identification-column
```

Table A4.9 Daily Table for a Component's Availability Data

Column Name	Description	Data Type
DATETIME	Datetime of the beginning of this day.	Numeric
DURATION	Sum of DURATIONs for all five-minute intervals for this day.	Numeric
DATE	Date part of DATETIME.	Numeric
About resource:		
...	...	
Last resource-identification column	Resource's value for this column.	depends on column
UPTIME	For this resource and this day, arithmetic mean of five-minute UPTIME values in this component's detail table for availability.	Numeric
UPTIME_D	For this resource and this day, standard deviation of five-minute UPTIME values in this component's detail table for availability.	Numeric
UPTIME_N	For this resource and this day, minimum of five-minute UPTIME values in this component's detail table for availability.	Numeric
UPTIME_S	For this resource and this day, sum of five-minute UPTIME values in this component's detail table for availability.	Numeric
UPTIME_X	For this resource and this day, maximum of five-minute UPTIME values in this component's detail table for availability.	Numeric
UPTIME_Y	For this resource and this day, range of five-minute UPTIME values in this component's detail table for availability.	Numeric
AVAIL	Percent of time during this day that this resource was available from UPTIME and DURATION, above: ((UPTIME/DURATION)*100).	Numeric
UNAVAIL	Percent of time during this day that this resource was unavailable from AVAIL, above: (100-AVAIL).	Numeric

Note: The availability of a component must always expressed as a value greater than or equal to zero. If the measure that you have chosen for your availability objective has a negative value in your source PDB, then the availability indicator for the component will be set to zero in the contract PDB. This prevents a component from having a negative availability. △

Response Time, Throughput, and Custom Objective Tables for a Component Node

The rows in these table are sorted in ascending order by the following fields:

```
DATETIME
first-resource-identification-column
...
last-resource-identification-column
```

Table A4.10 Daily Table for a Component's Response Time, Throughput, or Custom Objective Data

Column Name	Description	Data Type
DATETIME	Datetime of the beginning of this day.	Numeric
DURATION	Sum of DURATIONs for all five-minute intervals for this day in this component's detail table for this objective.	Numeric
DATE	Date part of DATETIME.	Numeric
About resource		
...	...	...
Last resource-identification column	Resource's value for this column.	depends on column
COUNT	For this resource, the count of events during this day in this component's detail table for this objective. If the data was not extracted from an EVENT table, then this is set to missing.	Numeric
MAX	For this resource and this day, arithmetic mean of the five-minute MAX values in this component's detail table for this objective.	Numeric
MAX____D	For this resource and this day, standard deviation of the five-minute MAX values in this component's detail table for this objective.	Numeric
MAX____N	For this resource and this day, minimum of the five-minute MAX values in this component's detail table for this objective.	Numeric
MAX____S	For this resource and this day, sum of the five-minute MAX values in this component's detail table for this objective.	Numeric
MAX____X	For this resource and this day, maximum of the five-minute MAX values in this component's detail table for this objective.	Numeric
MAX____Y	For this resource and this day, range of the five-minute MAX values in this component's detail table for this objective.	Numeric

Column Name	Description	Data Type
MIN	For this resource and this day, arithmetic mean of the five-minute MIN values in this component's detail table for this objective.	Numeric
MIN____D	For this resource and this day, standard deviation of the five-minute MIN values in this component's detail table for this objective.	Numeric
MIN____N	For this resource and this day, minimum of the five-minute MIN values in this component's detail table for this objective.	Numeric
MIN____S	For this resource and this day, sum of the five-minute MIN values in this component's detail table for this objective.	Numeric
MIN____X	For this resource and this day, maximum of the five-minute MIN values in this component's detail table for this objective.	Numeric
MIN____Y	For this resource and this day, range of the five-minute MIN values in this component's detail table for this objective.	Numeric
SUM	For this resource and this day, arithmetic mean of the five-minute SUM values in this component's detail table for this objective.	Numeric
SUM____D	For this resource and this day, standard deviation of the five-minute SUM values in this component's detail table for this objective.	Numeric
SUM____N	For this resource and this day, minimum of the five-minute SUM values in this component's detail table for this objective.	Numeric
SUM____S	For this resource and this day, sum of the five-minute SUM values in this component's detail table for this objective.	Numeric
SUM____X	For this resource and this day, maximum of the five-minute SUM values in this component's detail table for this objective.	Numeric
SUM____Y	For this resource and this day, range of the five-minute SUM values in this component's detail table for this objective.	Numeric
WSUM___S	For this resource and this day, sum of the five-minute WSUM values in this component's detail table for this objective.	Numeric
MEAN	For this resource and this day, arithmetic mean of the five-minute MEAN values in this component's detail table for this objective.	Numeric
MEAN___D	For this resource and this day, standard deviation of the five-minute MEAN values for response time.	Numeric

Column Name	Description	Data Type
MEAN___N	For this resource and this day, minimum of the five-minute MEAN values in this component's detail table for this objective.	Numeric
MEAN___S	For this resource and this day, sum of the five-minute MEAN values for response time.	Numeric
MEAN___X	For this resource and this day, maximum of the five-minute MEAN values in this component's detail table for this objective.	Numeric
MEAN___Y	For this resource and this day, range of the five-minute MEAN values in this component's detail table for this objective.	Numeric
WEIGHT	For this resource and this day, arithmetic mean of the five-minute WEIGHT values in this component's detail table for this objective.	Numeric
WEIGHT_S	For this resource and this day, sum of the five-minute WEIGHT values in this component's detail table for this objective.	Numeric
WMEAN	For this resource and this day, weighted arithmetic mean of the five-minute MEAN values in this component's detail table for this objective; based on WSUM___S and WEIGHT_S above: (WSUM___S/WEIGHT_S).	Numeric
SLI	Same value as MAX____X, MIN____N, SUM____S, MEAN, or WMEAN, above. (The value that is used is based on the calculation method selected for this objective of this component.)	Numeric

For more information about SLI, MAX, MIN, SUM, MEAN, WEIGHT, WSUM, and WMEAN, see "Understanding More about Objectives in Component Nodes" on page 24 and "Calculation Methods" on page 215.

Weekly Tables

The weekly tables have the same structure as the daily tables, except that the values are for a *week* instead of for a day, and the value of the DATETIME variable is the start of the *week* instead of the start of the day.

The values in the weekly tables are based on the values in the detail tables.

Monthly Tables

The monthly tables have the same structure as the daily tables, except that the values are for a *month* instead of for a day, and the value of the DATETIME variable is the start of a *month* instead of the start of a day.

The values in the monthly tables are based on the values in the detail tables.

Yearly Tables

The yearly tables have the same structure as the daily tables, except that the values are for a *year* instead of for a day, and the value of the DATETIME variable is the start of the *year* instead of the start of the day.

The values in the yearly tables are based on the values in the detail tables.

APPENDIX

5

SAS IT Service Level Management Macros

SAS IT Service Level Management Macros

The following macros are available in SAS IT Service Level Management:

Administrative Macros:

- %CPSLMENV: starts a SAS IT Service Level Management session and activates an environment database, or stops the session and deactivates the environment database.

  For more information about this macro, see "%CPSLMENV" on page 266.
- %CPSLMCON: activates a contract database.

  For more information about this macro, see "%CPSLMCON" on page 271.
- %CPSLMETL: runs the validate, synchronize, and ETL tasks for a contract.

  For more information about this macro, see "%CPSLMETL" on page 273.

Reporting Macro

- %CPSLMRPT: Produces reports for SAS IT Service Level Management.

  For more information about this macro, see "%CPSLMRPT" on page 279.

Special Purpose Macro

- %CPSLMFND: Searches contracts to locate where an object is used.

  For more information about this macro, see "%CPSLMFND" on page 283.

Note: For more information about these macros, contact the SAS IT Service Level Management administrator at your site. △

%CPSLMENV

Starts a SAS IT Service Level Management session and activates an environment database, or stops the session and deactivates the environment database

%CPSLMENV Overview

If the value of the ACTION= parameter is START, then the %CPSLMENV macro starts a SAS IT Service Level Management session on the server.

If the value of the ACTION= parameter is STOP, the %CPSLMENV macro stops the session and the environment database is deactivated. Any contract database that was activated will also be deactivated. (The %CPSLMENV and the %CPSLMCON macros can be used in batch jobs that produce reports. To run ETL in a batch job, use the %CPSLMETL macro instead.)

Note: If you invoke the macro in batch mode and receive a message that indicates that the SAS IT Resource Management software cannot be located, you can correct the problem in this way:

- Prior to invoking the macro, assign a libref to PGMLIB in order to specify the location of the SAS IT Resource Management software.

(This occurs only if you have installed SAS IT Resource Management software in a non-standard location.) △

%CPSLMENV Syntax

%CPSLMENV(
ACTION=START | STOP
,**ENVIRONMENTDBPATH=***environment-database-path*
<,**_RC=***macro-var-name*>
<,**SHAREPASSWORD=***password*>
<,**SHARESERVICE=***share-service-or-port*>
<,**SHAREUSER=***userid*>
<,**SLMPATH=***path-to-software*>
);

%CPSLMENV Details

The following parameters apply to the %CPSLMENV macro:

ACTION=START | STOP
specifies the action that you want the %CPSLMENV macro to perform. There is no default value.

The valid values are START and STOP. In a call to the %CPSLMENV macro, use only one of these values.

- If ACTION=START, the macro invokes the SAS IT Service Level Management software that is at the location that is specified by the SLMPATH= parameter and activates the environment database that is specified by the ENVIRONMENTDBPATH= parameter.
- If ACTION=STOP, the macro deactivates the environment database.

ENVIRONMENTDBPATH=*environment-database-path*
specifies the path of the environment database that is to be activated. The environment database must already exist and must be accessible from the server on which this SAS session is running.

There is no default value for this parameter.

Note: If the environment is shared and has been pre-allocated to the SAS/SHARE server under the libname SLM, then this parameter is not required.

For more information about SAS/SHARE, see the *SAS/SHARE User's Guide* in SAS Help and Documentation. △

_RC=*macro-var-name*
specifies the name of a macro variable that is to contain the return code from this macro. The name must be a valid SAS macro variable name. There is no default value for the name of the macro variable.

Note: Do not use _RC as the name of the macro variable. △

If the macro variable does not exist, then it will be created and will be given a value. If the macro variable whose name you specify already exists, then its value will be overwritten. To prevent conflicts with the names of SAS IT Management Solutions macros, avoid specifying any name that is already in use in your SAS session.

A value of zero indicates a successful execution of the macro. A nonzero value indicates that the macro failed; in this case you can check the explanatory message in the SAS log or in the SLM logs, which are written to the SASLIST file.

Note: You can test the value of the return code by using one of these macros from SAS IT Resource Management: %CPFAILIF or %CPLOGRC. For information about these macros, see the *SAS IT Resource Management: Macro Reference*. △

SHAREPASSWORD=*password*
specifies the password that you want to use to access the SLM server if you are using the SAS/SHARE server.

SHARESERVICE=*share-service-or-port*
specifies the SAS/SHARE service name or port number that you want to use to provide concurrent access to the SLM server. This field can contain a maximum of eight characters. There is no default. If you do not specify a value in this field, SAS/SHARE is not used.

Note: SAS/SHARE provides concurrent access to the environment database. This enables one or more clients to access the environment database while one or more ETL or reporting jobs is running. △

SHAREUSER=*userid*
specifies the user ID that is required in order to authenticate your access to the SAS/SHARE server.

SLMPATH=*path-to-software*
specifies the physical location where the SAS IT Service Level Management software is installed on a server. The location includes the complete directory path or high-level qualifiers where PGMLIB is installed, including the name PGMLIB.

This parameter is required if you install SAS IT Service Level Management in a location other than the default installation location. If you install SAS IT Service Level Management in the default installation location, then you do not need to specify this parameter.

For more information about the location of the software, see the installation instructions that accompany your software or ask the SAS Installation Representative at your site.

Note: In batch tasks, if the SAS/SHARE server requires authentication, then you must specify a user ID and password separately as parameters from the macro invocation of %CPSLMETL and %CPSLMENV. △

%CPSLMENV Notes

If you run a SAS IT Service Level Management batch job, you must run it on a SAS IT Service Level Management server. The environment database that is specified by the ENVIRONMENTDBPATH= parameter must be accessible from that server.

To run a job that generates user-written reports that were not created in SAS IT Service Level Management, perform these steps:

1. Begin a SAS IT Service Level Management session by invoking SAS.
2. Call the %CPSLMENV macro (with ACTION=START).
3. Call the %CPSLMCON macro.
4. Call the SAS code to generate reports.
5. Call the %CPSLMENV macro (with ACTION=STOP).
6. End the SAS IT Service Level Management session by calling the %CPSLMENV macro (with ACTION=STOP).

Note: You can omit the ENVIRONMENTDBPATH= and SLMPATH= parameters when ACTION=STOP. △

Note: If you want to switch from one environment database to another while the batch job is running, call %CPSLMENV (with ACTION=STOP) and then call %CPSLMENV (with ACTION=START and with a new value for the ENVIRONMENTDBPATH= parameter). △

%CPSLMENV Examples

Example 1: Report on one contract

This example is a job that can be used to produce reports on the ABC contract in the edb1 environment database.

```
// ...
%CPSLMENV (
    ACTION=START,
    ,ENVIRONMENTDBPATH=c:\slm\edb1
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;

%CPSLMCON (
   CONTRACT=ABC INC. Contract
   ,_RC=cpcorc
   );
%PUT CPSLMCON return code is &cpcorc ;

.
. (calls to the SAS code that generates reports go here.
.  This can include report job code that was generated in the GUI.)
.
%CPSLMENV (
    ACTION=STOP,
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;
```

Example 2: Report on multiple contracts in the same environment database

This example is a job that can produce reports on the ABC contract (in the edb1 environment database) and on the DEF contract (also in the edb1 environment database).

```
%CPSLMENV (
    ACTION=START,
    ,ENVIRONMENTDBPATH=/u/slm/edb1
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;

%CPSLMCON (
   CONTRACT=ABC INC. Contract
   ,_RC=cpcorc
   );
%PUT CPSLMCON return code is &cpcorc ;
```

```
.
. (calls to the SAS code that generates reports for ABC contract go here.)
.
%CPSLMCON (
   CONTRACT=DEF INC. Contract
   ,_RC=cpcorc
   );
%PUT CPSLMCON return code is &cpcorc ;

.
. (calls to the SAS code that generates reports for DEF contract go here.)
.
%CPSLMENV (
    ACTION=STOP,
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;
```

Example 3: Report on multiple contracts in different environment databases

This example is a batch job that can produce reports on the ABC contract (in the edb1 environment database) and on the DEF contract (in the edb2 environment database).

```
%CPSLMENV (
    ACTION=START,
    ,ENVIRONMENTDBPATH=sas.slm.edb1
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;

%CPSLMCON (
   CONTRACT=ABC INC. Contract
   ,_RC=cpcorc
   );
%PUT CPSLMCON return code is &cpcorc ;

.
. (calls to the SAS code that generates reports for ABC contract go here.)
.
%CPSLMENV (
    ACTION=STOP,
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;
%CPSLMENV (
    ACTION=START,
    ,ENVIRONMENTDBPATH=sas.slm.edb2
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;

%CPSLMCON (
   CONTRACT=DEF INC. Contract
   ,_RC=cpcorc
```

```
   );
%PUT CPSLMCON return code is &cpcorc ;

.
. (calls to the SAS code that generates reports for DEF contract go here.)
.
%CPSLMENV (
    ACTION=STOP,
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;
```

%CPSLMCON

Activates a contract database

%CPSLMCON Overview

The %CPSLMCON macro activates a contract database in read-only mode. (The contract is specified by the CONTRACT= or CONTRACTID= parameter.)

If a contract database has already been activated, then it will be deactivated and the contract that is specified by this invocation of %CPSLMCON will be activated.

Note: If you invoke the macro in batch mode and receive a message that indicates that the SAS IT Resource Management software cannot be located, you can correct the problem in this way:

- Prior to invoking the macro, assign a libref to PGMLIB in order to specify the location of the SAS IT Resource Management software.

(This occurs only if you have installed SAS IT Resource Management software in a non-standard location.) △

%CPSLMCON Syntax

%CPSLMCON(
CONTRACT=*contract-name* | **CONTRACTID=***contract-identifier*
<,**_RC=***macro-var-name*>
);

%CPSLMCON Details

The following parameters apply to the %CPSLMCON macro:

CONTRACT=*contract-name*
specifies the name of the contract. The contract must already exist. This name is the same as the name that you assign to the contract in the SAS IT Service Level Management client GUI.

You must specify a value for the CONTRACT= parameter or the CONTRACTID= parameter, but do not specify a value for both.

There is no default value for this parameter.

CONTRACTID=*contract-identifier*
specifies the identifier of the contract. The contract must already exist.

You must specify a value for the CONTRACTID= parameter or the CONTRACT= parameter, but do not specify a value for both.

There is no default value for this parameter.

_RC=*macro-var-name*
specifies the name of a macro variable that is to contain the return code from this macro. The name must be a valid SAS macro variable name. There is no default value for the name of the macro variable.

Note: Do not use _RC as the name of the macro variable. △

If the macro variable does not exist, then it will be created and will be given a value. If the macro variable whose name you specify already exists, then its value will be overwritten. To prevent conflicts with the names of SAS IT Management Solutions macros, avoid specifying any name that is already in use in your SAS session.

A value of zero indicates a successful execution of the macro. A nonzero value indicates that the macro failed; in this case you can check the explanatory message in the SAS log or in the SLM logs, which are written to the SASLIST file.

Note: You can test the value of the return code by using one of these macros from SAS IT Resource Management: %CPFAILIF or %CPLOGRC. For information about these macros, see the *SAS IT Resource Management: Macro Reference*. △

%CPSLMCON Example

Example: Activate the environmental database

This example accomplishes the following tasks:

- activates the environmental database (by the call to %CPSLMENV)
- activates the contract called "ABC INC. Contract 1," which is located in the edb1 environment database (by the call to %CPSLMCON)
- closes the environment database (by the final call to %CPSLMENV).

```
// ...
%CPSLMENV (
    ACTION=START
    ,ENVIRONMENTDBPATH=c:\myenvironments\edb1
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;

%CPSLMCON (
   CONTRACT=ABC INC. Contract 1
   ,_RC=cpcorc
   );
%PUT CPSLMCON return code is &cpcorc ;


 [If you want to report on this contract, you can insert a call to the report
   macro here.]
```

```
%CPSLMENV (
    ACTION=STOP,
    ,_RC=cpenrc
    );
%PUT CPSLMENV return code is &cpenrc ;
```

%CPSLMETL

Runs the validation, synchronization, and ETL task for a contract

%CPSLMETL Overview

The %CPSLMETL macro runs the validation task, the synchronization task, and the ETL task for a contract.

The validation task checks whether the contract and the data source PDBs from which the contract obtains data are consistent and complete so that an ETL task can be run.

The synchronization task updates the contract database to match its contract definition. This is necessary whenever you make changes (such as excusing an outage or changing a calculation method) to a node or objective in the contract hierarchy.

The ETL task updates the contract database with new data (from the data source PDBs). (The %CPSLMETL is used in batch jobs that run one or more ETL tasks. In a batch job that produces reports, use the %CPSLMENV and %CPSLMCON macros instead.)

Note: If you invoke the macro in batch mode and receive a message that indicates that the SAS IT Resource Management software cannot be located, you can correct the problem in this way:

- Prior to invoking the macro, assign a libref to PGMLIB in order to specify the location of the SAS IT Resource Management software.

(This occurs only if you have installed SAS IT Resource Management software in a non-standard location.) △

%CPSLMETL Syntax

%CPSLMETL(
,**ACTION=**VALIDATE | SYNC | ETL
,**CONTRACT=***contract-name* | **CONTRACTID=***contract-identifier*
,**ENVIRONMENTDBPATH=***environment-database-path*
<,**_RC=***macro-var-name*>
<,**SHAREPASSWORD=***password*>
<,**SHARESERVICE=***share-service-or-port*>
<,**SHAREUSER=***userid*>=
<,**SLMPATH=***path-to-software*>);

%CPSLMETL Details

The following parameters apply to the %CPSLMCON macro:

ACTION=VALIDATE | SYNC | ETL
specifies the action that the %CPSLMETL macro is to perform. There is no default value.

Valid values for this parameter are VALIDATE or SYNC or ETL. (In a call to the %CPSLMETL macro, use only one of these values.)

- If ACTION=VALIDATE, the macro checks whether all the information is available that is needed to run *if* ACTION=ETL. For example, validation will check to confirm the following:
  - The path to the environment database has been specified and the environment database exists.

    Note: If you have specified the parameters for SAS/SHARE, you need to specify the path to the environment database only if the SAS/SHARE server has not pre-allocated the desired environment database. △
  - The data source PDBs from which the ETL will extract data exists. The macro checks that it can connect to the data source PDB and that the specified tables exist in that PDB and that the specified columns exist in those tables.
  - The contract PDB exists.

    If the contract database that was specified by the CONTRACT= or CONTRACTID= parameter does not exist, then the macro will create one.

    Note: In the z/OS operating environment, you must allocate the space for the contract database before you call the %CPSLMETL macro for the first time for a contract. For more information and a sample job step to allocate the space, see the CMSLMALC member in the CPMISC PDS.

    The CPMISC data set is a standard z/OS partitioned data set (PDS). The default location for this data set is as follows:

    ```
    where_you_installed_sas_at_your_site.itrm.cpmisc
    ```

    △

    If the contract database that was specified by the CONTRACT= or CONTRACTID= parameter does not contain all of the tables that the contract requires, this macro will create them.

  The macro writes a validation log to the SASLIST file that is associated with the batch job. (SASLIST is the standard SAS output area.)

  You can also view the validation log from the SAS IT Service Level Management client GUI.
- If ACTION=SYNC, synchronization will perform any necessary calculations of the data in order to match the current definition of the hierarchy.

  The macro does not read in new data from the data source PDBs that are specified in the contract.
- If ACTION=ETL, the macro updates the components in the contract database with new data from the data source PDBs. This data is then "rolled up" through the higher nodes in order to match the current definition of the contract hierarchy.

If you specify ACTION=ETL, then the following tasks are run in this order:

1. Synchronization
2. Validation
3. ETL

If an error is discovered in the synchronization task, then the validation and ETL tasks are not run.

If an error is discovered in the validation task, then the ETL task is not run.

Otherwise, the macro will continue processing.

For more information about this topic, see "The ETL Process" on page 287.

Note: Calculation of the parent node's values from the values of its child nodes is sometimes called *rollup*. This does *not* mean that the child nodes pass up the values that their parent node is to use. The child nodes do *not* determine which values the parent node uses. Instead, the parent node determines which values it will obtain from its child nodes. In other words, it *reaches down* to the child nodes. △

CONTRACT=*contract-name*
specifies the name of the contract. The contract must already exist. The name is the same as the name that you assign to the contract in the SAS IT Service Level Management client GUI.

You must specify a value for either the CONTRACT= parameter or the CONTRACTID= parameter, but do not specify a value for both.

There is no default value for this parameter.

CONTRACTID=*contract-identifier*
specifies the identifier of the contract. The contract must already exist. You must specify a value for either the CONTRACTID= parameter or the CONTRACT= parameter, but do not specify a value for both.

There is no default value for this parameter.

ENVIRONMENTDBPATH=*environment-database-path*
specifies the path of the environment database that is to be activated. The environment database must already exist and must be accessible from the server on which this SAS session is running.

There is no default value for this parameter.

Note: If the environment is shared and has been pre-allocated to the SAS/SHARE server under the libname SLM, then this parameter is not required.

For more information about SAS/SHARE, see the *SAS/SHARE User's Guide* in SAS Help and Documentation. △

_RC=*macro-var-name*
specifies the name of a macro variable that is to contain the return code from this macro. The name must be a valid SAS macro variable name. There is no default value for the name of the macro variable.

Note: Do not use _RC as the name of the macro variable. △

If the macro variable does not exist, then it will be created and will be given a value. If the macro variable whose name you specify already exists, then its value will be overwritten. To prevent conflicts with the names of SAS IT Management Solutions macros, avoid specifying any name that is already in use in your SAS session.

A value of zero indicates a successful execution of the macro. A nonzero value indicates that the macro failed; in this case you can check the explanatory message in the SAS log or in the SLM logs, which are written to the SASLIST file.

Note: You can test the value of the return code by using one of these macros from SAS IT Resource Management: %CPFAILIF or %CPLOGRC. For information about these macros, see the *SAS IT Resource Management: Macro Reference*. △

SHAREPASSWORD=*password*
specifies the password that you want to use to access the SLM server if you are using the SAS/SHARE server.

SHARESERVICE=*share-service-or-port*
specifies the SAS/SHARE service name or port number that you want to use to provide concurrent access to the SLM server. This field can contain a maximum of eight characters. There is no default. If you do not specify a value in this field, SAS/SHARE is not used.

Note: SAS/SHARE provides concurrent access to the environment database. This enables one or more clients to access the environment database while one or more ETL or reporting jobs are running. △

SHAREUSER=*userid*
specifies the user ID that is required in order to authenticate your access to the SAS/SHARE server.

SLMPATH=*path-to-software*
specifies the physical location where the SAS IT Service Level Management software is installed on a server. The location includes the complete directory path or high-level qualifiers where PGMLIB is installed, including the name PGMLIB.

This parameter is required if you install SAS IT Service Level Management in a location other than the default installation location. If you install SAS IT Service Level Management in the default installation location, then you do not need to specify this parameter.

For more information about the location of the software, see the installation instructions that accompany your software or ask the SAS Installation Representative at your site.

Note: In batch tasks, if the SAS/SHARE server requires authentication, you must specify a user ID and password separately as parameters from the macro invocation of %CPSLMETL and %CPSLMENV. △

%CPSLMETL Notes

Each action (VALIDATE, SYNC, and ETL) has a separate log. In the batch job, the three logs are written to the SASLIST file. If you want to view the logs, then from the Contract window in the SAS IT Service Level Management client GUI, select **Tools ▶ View Last Results**.

ACTION=ETL runs the VALIDATE, SYNC, and ETL actions, so you do not need to run the VALIDATE and SYNC actions separately.

However, you can run the the actions separately in order to accomplish the following goals:

- You can run the VALIDATE action separately in order to make changes in the GUI and to confirm (during the day) that the changes are valid. This enables you to confirm that the ETL job will be able to run (typically at night) to completion.
- You can run the SYNC action separately in order to make changes in the GUI that will require a large amount of recalculation. Examples are defining an excused outage or changing the calculation method in a component.

  Running the SYNC action separately enables the recalculation to be done before the regular ETL processing begins, thus ensuring that the nightly ETL job will not run out of time before it completes.

- You can run the SYNC action separately and then produce reports that are based on the changed data, before going ahead and processing new data for the contract.

%CPSLMETL Examples

Note: For an example of using a SAS/SHARE server to share a SAS IT Service Level Management environment database among multiple users or jobs, see the Installation Instructions. △

The following examples show code that runs after invoking SAS in a batch job. For more information about building and submitting batch jobs that invoke SAS, see "Batch Jobs" in Appendix 7 of the *SAS IT Service Level Management: User's Guide*.

Example 1: Run only the validation task

This call to the %CPSLMETL macro runs the validation task for the contract with the ABC Corporation.

```
%CPSLMETL (
    ACTION=VALIDATE
    ,ENVIRONMENTDBPATH=c:\slm\envDB
    ,CONTRACT=ABC Corporation
    ,_RC=cpetrc
    );
%PUT CPSLMETL (ABC Corporation contract validation) return code is &cpetrc ;
```

You can see the results of this run by checking the validation log.

Example 2: Run only the synchronization task

This call to the %CPSLMETL macro runs the synchronization task for the contract with the ABC Corporation.

```
%CPSLMETL (
    ACTION=SYNC
    ,ENVIRONMENTDBPATH=/u/slm/envdb
    ,CONTRACT=ABC Corporation
    ,_RC=cpetrc
    );
%PUT CPSLMETL (ABC Corporation contract synchronization) return code is &cpetrc ;
```

You can see the results of this run by checking the synchronization log.

Example 3: Run the validation, synchronization, and ETL tasks for one contract

The following code runs the validation, synchronization, and ETL tasks for the contract with the ABC Corporation.

```
%CPSLMETL (
    ACTION=ETL
    ,ENVIRONMENTDBPATH=sas.slm.envdb
    ,CONTRACT=ABC Corporation
    ,_RC=cpetrc
    );
%PUT CPSLMETL (ABC Corporation contract ETL) return code is &cpetrc ;
```

If this call runs after the code in Example 1 and Example 2, then the VALIDATE and SYNC actions are skipped and the macro runs only the ETL task. If the code in Example 1 and Example 2 does not run before this call, then this call runs all three tasks (VALIDATE, SYNC, and ETL). You can see the results of this run by checking the validation, synchronization, and ETL logs.

Example 4: Run the validation, synchronization, and ETL tasks for more than one contract

The following code runs the validation, synchronization, and ETL tasks for the contract with the DEF Corporation. Then it runs the validation, synchronization, and ETL tasks for the contract with the GHI Corporation.

```
%CPSLMETL (
    ACTION=ETL
    ,ENVIRONMENTDBPATH=\\server\slm\envdb
    ,CONTRACT=DEF Corporation
    ,_RC=cpetrc
    );
%PUT CPSLMETL (DEF Corporation contract validation) return code is &cpetrc ;

%CPSLMETL (
    ACTION=ETL
    ,ENVIRONMENTDBPATH=\\server\slm\envdb
    ,CONTRACT=GHI Corporation
    ,_RC=cpetrc
    );
%PUT CPSLMETL (GHI Corporation contract validation) return code is &cpetrc ;
```

Example 5: Run the validation, synchronization, and ETL tasks for a contract and then produce the reports for the contract

The following code runs the validation, synchronization, and ETL tasks for the contract with the ABC Corporation. It also produces service level reports from the report specification that is named MyServiceLevelReports.

```
%CPSLMETL (
    ACTION=ETL
    ,ENVIRONMENTDBPATH=c:\myenvironments\envdb
    ,CONTRACT=ABC Corporation
    ,_RC=cpetrc
    );
%PUT CPSLMETL (ABC Corporation contract ETL) return code is &cpetrc ;

%CPSLMRPT (
       ACTION=SERVICE_LEVEL
       ,ENVIRONMENTDBPATH=c:\myenvironments\envdb
       ,_RC=slrptrc
       ,REPORT=MyServiceLevelReports
       ,SLMPATH=c:\SAS 91\pgmlib
       );
%PUT CPSLMRPT return code is &slrptrc;
```

%CPSLMRPT

Produces reports for SAS IT Service Level Management

%CPSLMRPT Overview

This macro can produce two types of reports: a set of HTML-based drill-down reports for service level reports or a gallery of graphs for the baseline reports. (All ETL processes must be completed before this macro is submitted.)

Note: If you invoke the macro in batch mode and receive a message that indicates that the SAS IT Resource Management software cannot be located, you can correct the problem in this way:

- Prior to invoking the macro, assign a libref to PGMLIB in order to specify the location of the SAS IT Resource Management software.

Receipt of this message occurs only if you have installed SAS IT Resource Management software in a non-standard location. △

%CPSLMRPT Syntax

%CPSLMRPT(
<**ACTION=**SERVICE_LEVEL | BASELINE>
<,**CONTRACT=***contract-name* | **CONTRACTID***contract-identifier*>
<,**ENVIRONMENTDBPATH=***environment-database-path*>
<,**OUTPUTLOCATIONPATH=***path-to-output-location*>
<,**_RC=***macro-var-name*>
<,**REPORT=***report-name* | **REPORTID=***report-identifier*>
<,**SHAREPASSWORD=***password*>
<,**SHARESERVICE=***share-service-or-port*>
<,**SHAREUSER=***userid*>
<,**SLMPATH=***path-to-software*>
<,**TEXTENCODING=***encoding-for-HTML-text*>
<,**TEXTTERMSTR=***end-of-line-character-for-HTML-text*>
);

%CPSLMRPT Details

The following parameters apply to the %CPSLMRPT macro:

ENVIRONMENTDBPATH=*environment-database-path*
specifies the path of the environment database that is to be activated. The environment database must already exist and must be accessible from the server on which this SAS session is running.

There is no default value for this parameter.

Note: If the environment is shared and has been pre-allocated to the SAS/SHARE server under the libname SLM, then this parameter is not required.

For more information about SAS/SHARE, see the *SAS/SHARE User's Guide* in SAS Help and Documentation. △

ACTION=SERVICE_LEVEL | BASELINE
specifies the type of report that you want the %CPSLMRPT macro to generate. The default is SERVICE_LEVEL.

CONTRACT=*contract-name*
specifies the name of the contract for which you are running baseline reports. The contract must already exist. This name is the name that you assign to the contract in the SAS IT Service Level Management client GUI.

If you are running baseline reports, then you must specify a value for the CONTRACT= parameter or the CONTRACTID= parameter, but do not specify a value for both.

There is no default value for this parameter.

CONTRACTID=*contract-identifier*
specifies the identifier of the contract for which you are running baseline reports. The contract must already exist. This identifier is automatically generated by SAS IT Service Level Management when you create the contract.

If you are running baseline reports, then you must specify a value for the CONTRACTID= parameter or the CONTRACT= parameter, but do not specify a value for both.

There is no default value for this parameter.

OUTPUTLOCATIONPATH=*path -to-output-location*
specifies the path to the directory where the baseline reports will be stored. Do not specify this parameter for service level reports.

_RC=*macro-var-name*
specifies the name of a macro variable that is to contain the return code from this macro. The name must be a valid SAS macro variable name. There is no default value for the name of the macro variable.

Note: Do not use _RC as the name of the macro variable. △

If the macro variable does not exist, then it will be created and will be given a value. If the macro variable whose name you specify already exists, then its value will be overwritten. To prevent conflicts with the names of SAS IT Management Solutions macros, avoid specifying any name that is already in use in your SAS session.

A value of zero indicates a successful execution of the macro. A nonzero value indicates that the macro failed; in this case you can check the explanatory message in the SAS log or in the SLM logs, which are written to the SASLIST file.

Note: You can test the value of the return code by using one of these macros from SAS IT Resource Management: %CPFAILIF or %CPLOGRC. For information about these macros, see the *SAS IT Resource Management: Macro Reference*. △

REPORT
specifies the name of the report specification that you want to run. This name is the same name that you assign to the report specification in the SAS IT Service Level Management client GUI.

If you are running service level reports, then you must specify a value for the REPORT= parameter or the REPORTID= parameter, but do not specify a value for both.

There is no default value for this parameter.

REPORTID
specifies the identifier of the report specification that you want to run. This identifier is automatically generated by SAS IT Service Level Management when you create the report specification.

If you are running service level reports, then you must specify a value for the REPORTID= parameter or the REPORT= parameter, but do not specify a value for both.

There is no default value for this parameter.

SHAREPASSWORD=*password*
specifies the password that you want to use to access the SLM server if you are using the SAS/SHARE server.

SHARESERVICE=*share-service-or-port*
specifies the SAS/SHARE service name or port number that you want to use to provide concurrent access to the SLM server. This field can contain a maximum of eight characters. There is no default. If you do not specify a value in this field, SAS/SHARE is not used.

Note: SAS/SHARE provides concurrent access to the environment database. This enables clients to access the environment database while one or more ETL jobs are running. △

SHAREUSER=*userid*
specifies the user ID that is required in order to authenticate your access to the SAS/SHARE server.

SLMPATH=*path-to-software*
specifies the physical location where the SAS IT Service Level Management software is installed on a server. The location includes the complete directory path or high-level qualifiers where PGMLIB is installed, including the name PGMLIB.

This parameter is required if you install SAS IT Service Level Management in a location other than the default installation location. If you install SAS IT Service Level Management in the default installation location, then you do not need to specify this parameter.

For more information about the location of the software, see the installation instructions that accompany your software or ask the SAS Installation Representative at your site.

TEXTTERMSTR=NONE | NL | CR | LF | CRLF | LFCR | CRNL
specifies the type of record separator character sequences to use to terminate records in the file for reports that are generated with the parameter ACTION=SERVICE_LEVEL.

The TEXTTERMSTR= parameter accepts the following values to use as record terminators:

Table A5.1 Table of Record Terminator Values

Value of TEXTTERMSTR=	Description
NONE	Record terminators are not used.
NL	Record terminator is the new line character (x'15').
CR	Record terminator is the carriage return character (x'0c').
LF	Record terminator is the line feed character (x'25').
CRLF	Record terminator is the sequence CR followed by LF.
LFCR	Record terminator is the sequence LF followed by CR.
CRNL	Record terminator is the sequence CR followed by NL.

If the value you enter is not valid for your installed image of SAS, an error message will be written to the SAS log and the report processing will be terminated.

For more information about the values, filenames, and dataset options that can be used to specify the TEXTTERMSTR= value, see the documentation for TERMSTR= in the *SAS National Language Support (NLS): User's Guide*.

TEXTENCODING=*encoding-value*
specifies which type of encoding is to be used for textual reports that are generated with the parameter ACTION=SERVICE_LEVEL. The default is the encoding value of your SAS session.

For more information about the values, filenames, and dataset options that can be used to specify the TEXTENCODING= value, see the documentation for ENCODING= in the *SAS National Language Support (NLS): User's Guide*.

Tip: If you are generating your reports on a z/OS server and wish to deliver those reports on a Web server running under Windows or UNIX, specify

```
TEXTENCODING=UTF8
```

%CPSLMRPT Notes

If you specify ACTION=SERVICE_LEVEL, then you must also specify either REPORT= or REPORTID=. Do not specify CONTRACT=, CONTRACTID=, or OUTPUTLOCATIONPATH=.

If you specify ACTION=BASELINE, then you must also specify either CONTRACT= or CONTRACTID=. Additionally, you must specify OUTPUTLOCATIONPATH=. Do not specify REPORT= or REPORTID=.

%CPSLMRPT Examples

Example 1: Generate service level reports

This example produces service level reports from the report specification that is named MyServiceLevelReports.

```
// ...
%global slrptrc;
%let slrptrc=0;

%CPSLMRPT (
       ACTION=SERVICE_LEVEL
       ,ENVIRONMENTDBPATH=c:\myenvironments\edb1
       ,_RC=slrptrc
       ,REPORT=MyServiceLevelReports
       ,SLMPATH=c:\SAS 91\pgmlib
       );
%PUT CPSLMRPT return code is &slrptrc;
```

Example 2: Generate baseline reports

This example produces baseline reports for a contract that is named MyFirstContract.

```
// ...
%global baselnrc;
%let slrptrc=0;

%CPSLMRPT (
       ACTION=BASELINE
       ,ENVIRONMENTDBPATH=c:\myenvironments\edb1
       ,_RC=baselnrc
       ,CONTRACT=MyFirstContract
```

```
       ,OUTPUTLOCATIONPATH=C:\baseline_rpts\gallery1\
       ,SLMPATH=c:\SAS 91\pgmlib
       );
%PUT CPSLMRPT return code is &baselnrc;
```

Example 3: Generate service level reports using the TEXTENCODING= parameter

This example produces service level reports from the report specification that is named MyServiceLevelReports for a customer who is running under z/OS. If the customer specifies the TEXTENCODING=UTF8 parameter, the generated reports can be displayed on a UNIX or Windows server.

```
// ...
%global slrptrc;
%let slrptrc=0;
%CPSLMRPT (
       ACTION=SERVICE_LEVEL
       ,ENVIRONMENTDBPATH=SLM.MYENVIRONMENT.EDB1
       ,_RC=slrptrc
       ,REPORT=MyServiceLevelReports
       ,SLMPATH=c:\SAS 91\pgmlib
       ,TEXTENCODING=UTF8
);
%PUT CPSLMRPT return code is &slrptrc;
```

%CPSLMFND

Searches contract hierarchies to locate specific nodes.

%CPSLMFND Overview

The %CPSLMFND macro searches all the contract databases in the current environment in order to find specific nodes (either a contract, an agreement, a service, or a component node). The macro can search either by the name of the node or by the parent identifier of the node, that is, the identifier of the catalog node from which the contract nodes were derived.

In order to establish the IT Service Level Management environment, execute the %CPSLMENV macro prior to executing the %CPSLMFND macro. For information about this, see "%CPSLMFND Examples" on page 285.

%CPSLMFND Syntax

%CPSLMFND(
objectType=CONTRACT | AGREEMENT | SERVICE | COMPONENT | RANGESET
,**objectName=***object-name* | **masterID=***object-identifier*
<,**outputLocation=***output-location*>
<,**resultType=**LOG | TEXT | HTML | CATALOG>
<,**_RC=***macro-var-name*>
);

%CPSLMFND Details

The following parameters apply to the %CPSLMFND macro:

objectType=CONTRACT | AGREEMENT | SERVICE | COMPONENT | RANGESET
specifies the type of object that you want to search for. Valid values for this parameter are CONTRACT, AGREEMENT, SERVICE, COMPONENT, or RANGESET. (In a call to the %CPSLMFND macro, use only one of these values.)
The default is COMPONENT.

objectName=*object-name*
specifies the name of the object that you want to search for. (Searches by name are case-sensitive.) Use this parameter to locate all the places in the current environment where this object is used.
You can search for either an object by name or by identifier, but not both.

masterID=*object-identifier*
specifies the identifier of the node from which the objects were derived. The identifier can be found in the property sheet of the catalog node. Use this parameter to locate all the objects that were derived from this original object.
You can search for either an object by name or by identifier, but not both.

outputLocation=*output-location*
specifies the location where you want to display the output of your search. If not specified, the default is to display the output on the SAS log.
For resultType=TEXT or resultType=HTML, *output-location* is an external, physical filename. For resultType=CATALOG, *output-location* is a fully qualified (four-level) SAS SLIST catalog entry name.

resultType=LOG | TEXT | HTML | CATALOG
specifies the type of output. The default is LOG.

LOG	specifies that the output is to go to the SAS log.
TEXT	specifies that the text output is to be written to an external file.
HTML	specifies that the HTML output is to be written to an external file.
CATALOG	specifies that the output is to go to a SAS catalog SLIST entry.

Note: If you specify resultType= TEXT or HTML or CATALOG, you must also specify the outputLocation= parameter. △

_RC=*macro-var-name*
specifies the name of a macro variable that is to contain the return code from this macro. The name must be a valid SAS macro variable name. There is no default value for the name of the macro variable.

Note: Do not use _RC as the name of the macro variable. △
If the macro variable does not exist, then it will be created and will be given a value. If the macro variable whose name you specify already exists, then its value will be overwritten. To prevent conflicts with the names of SAS IT Management Solutions macros, avoid specifying any name that is already in use in your SAS session.
A value of zero indicates a successful execution of the macro. A nonzero value indicates that the macro failed; in this case you can check the explanatory message in the SAS log or in the SLM logs that are written to the SASLIST file.

Note: You can test the value of the return code by using one of these macros from SAS IT Resource Management: %CPFAILIF or %CPLOGRC. For information about these macros, see the *SAS IT Resource Management: Macro Reference*. △

%CPSLMFND Notes

If you specify resultType=CATALOG, then specify the name of a catalog entry for the outputLocation. At this time, SLIST is the only type of catalog entry that is supported.

%CPSLMFND Examples

Example 1: Search by name

In this example, you might have created the component called "Web Server" and added it to several contracts without changing its name. This example locates all components in all contracts in the current environment that have the name "Web Server."

```
%CPSLMENV(slmPath=locationOfSoftware,
          environmentDBPath=locationOfEnvironment,
          action=START);
%CPSLMFND(objectType=COMPONENT,
          objectName=Web Server);
```

Example 2: Search by identifier

In this example, you might have created the component called "Web Server" and added it to several contracts changing its name in each of the contracts. In order to locate all the components that are related to the component that was originally added to the master catalog, you must first obtain the identifier of the catalog's component. You can then use the identifier to search the contracts in the current environment for the nodes that were derived from the component that was originally entered in the master catalog.

This example locates all components in all contracts in the current environment that were derived from the originating catalog component whose identifier is 23.

```
%CPSLMENV(slmPath=locationOfSoftware,
          environmentDBPath=locationOfEnvironment,
          action=START);
%CPSLMFND(objectType=COMPONENT,
          masterID=23);
```

APPENDIX

6

The ETL Process

The ETL Process

The ETL process consists of these three tasks:

- "The Synchronization Task" on page 287
- "The Validation Task" on page 288
- "The ETL Task" on page 289

Note: As changes are made to a contract through the SAS IT Service Level Management GUI, any changes in the contract database that are required to reflect the changes in the contract are put in a list. The items on the list are called *pending synchronization actions*.

If any synchronization actions are pending, the report job will not run. No reports will be generated. A message will be displayed in the Report Results dialog box saying that there is a pending synchronization action on one or more of the contracts that are being used in the report. The message will be prefixed by "ERROR." This will also be written to the report log. △

The Synchronization Task

The synchronization task updates the existing data in the contract database in accordance with the current specifications of the contract.

Note: If any step of the synchronization task is not successful, the task stops and gives a non-zero return code. If all steps of the synchronization task are successful, the task gives a zero return code. △

As its first steps, the synchronization task performs the following verifications:

1. It verifies that the contract status is set to *Active*.
2. It verifies that the day that represents the start of the week is valid.
3. It verifies that the contract database's path is not blank.
4. It verifies that the contract database exists and can be accessed in *Write* mode.

Next, the synchronization task begins traversing the contract hierarchy, starting at the top (that is, at the contract node). As each node is encountered, the synchronization task performs the following actions if they are necessary:

1. It constructs the code to apply the appropriate outages to this node. This code is used whenever data is repopulated in the table for this node or in the tables for all descendant nodes (regardless of how far down the hierarchy those nodes are).

   For more information about the effect of the outage code on data, see "Effect of Outages on Data" on page 233.
2. If this node's table does not exist in the contract database, it creates the table. It also clears any pending synchronization actions for this node, because a newly created table does not have any data that requires changing.
3. It synchronizes any child nodes. Thus, conceptually, the synchronize task consists of a series of mini-synchronizations.
4. It examines the list of pending synchronization actions for this node, and performs the appropriate steps for each action.

   a. It might purge the table and repopulate it from the child nodes (if this node is a contract, service level agreement, or service node) or purge the main table and repopulate it from the resources (if this node is a component node), using the current definitions for this node.
   b. It might resummarize the data to the day, week, month, and year tables.
   c. It might regenerate formulas and views, as needed.
   d. If this node is a component node and data must be repopulated in the component's main table, it might verify that the resource identification columns and the resource rows have been selected. If they have not been selected, stop with a non-zero return code.
   e. After the changes are finished, it might clear the list of pending resource actions.

Then, after the entire hierarchy has been traversed and successfully synchronized, the synchronization task performs this action.

1. It constructs the data sets for reporting.

   The report job code that was generated by the **`Report Specification`** Wizard in the GUI specifies which data sets to construct. The data sets are stored in the ADMIN library of the contract database.

If the batch job for reporting is resubmitted after this step, the reports that are produced incorporate the changes that were made to the contract through the GUI.

The Validation Task

The validation task makes sure that an ETL task can succeed.

Note: If any step of the validation task is not successful, the task attempts to continue but will give a non-zero return code. If all steps of the validation task are successful, the task gives a zero return code. △

As its first steps, the validation task performs the following verifications:

1. It verifies that the Start of Week day is valid.
2. It verifies that the Contract Database path is not blank.
3. It verifies that the Contract Database exists and can be accessed in WRITE mode.

Next, the validation task begins traversing the contract hierarchy, starting at the top (that is, at the contract node). As each node is encountered, the validation task performs the following actions as necessary:

1. It verifies that this node's (main) table exists in the contract database. For more information about a node's (main) table, see "The Contract Database" on page 237.
2. It verifies that there are no pending synchronization actions for the node.
3. It verifies the node's objectives, and checks whether the availability, response time, throughput, and custom objective types that are associated with the node are active.

   Note: If an SLA, service, or component node has no activated objective, then the validation process will fail and an error message is displayed. If there are no other failures, the validation will be considered successful. △
4. It verifies any child nodes. Thus, conceptually, the validate task consists of a series of mini-validations.

Then, if the node currently being validated is a component node, the validation task performs these actions:

1. It verifies that the component's underlying tables (for availability, response time, throughput, and custom objective) exist in the contract database.

   For more information about a component's underlying tables, see "The Contract Database" on page 237.
2. For the availability, response time, throughput, and custom objectives, it verifies that the objective's data source PDB, table, and measure are selected.

   It also verifies that the data can be read from the data source PDB. (This might involve establishing a network connection if the data source PDB is accessed remotely.)
3. It verifies that the resource list has columns and rows.

   It also verifies that the variables selected for the columns exist in all the data source PDBs that are selected for this node.

The ETL Task

The ETL task loads new data into the contract database in accordance with the current definitions in the contract.

Note: If any step of the ETL process is not successful, ETL stops and gives a non-zero return code. If the ETL process is successful, ETL gives a zero return code. △

As its first steps, the ETL task performs the following actions:

1. It runs the synchronization task.
2. It runs the validation task, if the synchronization task is successful.

Next, if the validation task is successful, the ETL task begins to traverse the hierarchy starting at the top (that is, at the contract node).

For each component node, the ETL task performs the following steps:

1. For each objective of the component, it performs the following actions:

   a. For the selected resources, it extracts the appropriate data from the data source PDB.

      Note: If an observation has missing values for the datetime stamp or for the measure, then it is not extracted. It is omitted from the data for that resource. △

b It adjusts the data to the interval length that is specified in the contract; this might involve combining shorter-length observations, prorating longer-length observations, and shifting the beginning and end of the interval to be a multiple of the interval length.
c It calculates the requested measure.
d It writes the data to the component's table in the contract database that is associated with this contract.

2 For each component, it performs the following actions:

a It calculates the value of the objective's indicator by using the data from the component's table and the calculation method that was specified for the objective.

Note: Date filtering, if it is in effect for this node, is applied. △

b If a range set is specified in the objective, the ETL process obtains the value of the score, the color to be used for the text reports, the label to be used in the reports, and the description of the range set for the objective.
c It calculates whether the target has been met or missed for the objective.
d It calculates the value of the composite indicator by using the availability, response time, throughput, and custom objective indicator or score, as specified on the composite objective and the calculation method that is specified on the composite objective.
e If a range set is specified in the composite objective, the ETL process obtains the value of the composite score, the color to be used for the text reports, the label to be used in the reports, and the description of the range set for the composite objective.
f It calculates whether the target has been met or missed for the composite objective.
g It writes the results as an observation in the detail level of the component's table in the contract database that is associated with this contract.
h It updates the summary statistics in the day, week, month, and year levels of the component's main table in the contract database that is associated with this contract.

Note: Date filtering can be specified for any node in the contract. △

□ When each node's calculations are completed, its values are available to its parent node. When the calculations are completed for all child nodes of a parent, the calculations begin on the parent's node.

□ When the contract node's calculations are completed, the ETL task performs post-processing work, such as building data sets for use by the report job code.

Note: Outages are applied to the data at the contract node. For information about the effect of outages on the data, see "Effect of Outages on Data" on page 233. △

If the batch job for reporting is submitted after this step, the reports that are produced incorporate the changes that were made to the contract through the GUI and the changes due to acquiring new data from the data source PDBs.

E-mail Notifications Setup

A contract can be specified so that an e-mail notification is sent to the servicing contact if an objective of the contract or one of the nodes under it has not met its target. When the ETL process runs, it determines if an e-mail notification needs to be sent. If so, the SAS system is used to dispatch the e-mail notifications.

Note: If SAS is not configured to use Simple Mail Transfer Protocol (SMTP), then the e-mail notifications will not be sent and a warning message will be written to the ETL and Validate log files.

(The default for Windows is Messaging Application Programming Interface (MAPI) and it is not supported.) △

APPENDIX

7

Sample E-mail Notification

Sample E-mail Notification

This is a sample of the notice that is sent by e-mail to any contact who services a contract, SLA, service, or component that has failed to meet the target of one or more of its objectives.

```
From: Jim Taller@ReliableNets.com
Sent: Friday, July 01, 2005 5:44 PM
To: John Taller
Subject: Missed targets for Contract: Big Dollar Bank Ltd.
- SAS IT Service Level Management

The node(s)in this Contract: Big Dollar Bank Ltd. are assigned to
Servicing Contacts using this e-mail:
john.Taller@ReliableNets.com. The following objectives failed to
meet their targets:
Node  Objective/      Objective Label/         Actual         Target
   Compliance Period   Target Details           Value          Value
-------------------------------------------------------------------------
Contact: John Taller
CONTRACT: Big Dollar Bank Ltd.
  ---------------------------------------------------------------------
  Availability        Duration                  97.89           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ---------------------------------------------------------------------
  Response Time                                  1.05            1.00
  Daily   22MAR2005   (MEAN) I < 1
-------------------------------------------------------------------------
Contact: John Taller
AGREEMENT: Online Mortage SLA
  ---------------------------------------------------------------------
  Availability        Duration                  97.89          100.00
  Monthly 01MAR2005   (MEAN) I > 100
  ---------------------------------------------------------------------
  Response Time                                215.79            1.00
  Daily   22MAR2005   (MEAN) I < 1 ms
-------------------------------------------------------------------------
Contact: John Taller
```

```
SERVICE: External Customer Network service
  ------------------------------------------------------------------------
  Availability        Duration                  97.89           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ------------------------------------------------------------------------
  Response Time                                 12.49            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
SERVICE: Link to Internet service
  ------------------------------------------------------------------------
  Availability        Duration                  97.89          100.00
  Monthly 01MAR2005   (MEAN) I > 100
  ------------------------------------------------------------------------
  Response Time                                 11.72            1.00
  Daily   22MAR2005   (MEAN) I < 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: Link 79319 component
  ------------------------------------------------------------------------
  Availability        Duration                  97.76           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ------------------------------------------------------------------------
  Response Time                                 11.23            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: Link 75141 component
  ------------------------------------------------------------------------
  Availability        Duration                  97.89           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ------------------------------------------------------------------------
  Response Time                                 12.21            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
SERVICE: Firewall service
  ------------------------------------------------------------------------
  Response Time                                 10.02            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: Firewall 1x9321 component
  ------------------------------------------------------------------------
  Response Time                                 10.02            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
SERVICE: External Customer Switch service
  ------------------------------------------------------------------------
  Availability        Duration                  97.89           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ------------------------------------------------------------------------
  Response Time                                 10.82            1.00
```

```
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: Switch 5100 component
  ------------------------------------------------------------------------
  Availability        Duration                  97.89           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ------------------------------------------------------------------------
  Response Time                                 10.23            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: Switch 9321 component
  ------------------------------------------------------------------------
  Availability        Duration                  97.89           99.50
  Monthly 01MAR2005   (MEAN) I >= 99.5
  ------------------------------------------------------------------------
  Response Time                                 10.96            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
SERVICE: Mortgage Application service
  ------------------------------------------------------------------------
  Response Time                                203.30            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: WolfMan application
  ------------------------------------------------------------------------
  Response Time                                 205.99            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
Contact: John Taller
COMPONENT: Q21C1 application
  ------------------------------------------------------------------------
  Response Time                                 200.24            1.00
  Daily   22MAR2005   (MEAN) I <= 1 ms
------------------------------------------------------------------------
"I" stands for Indicator, "S" stands for Score.
"Actual Value" is subject to the effects of rounding
due to format restrictions.
```

APPENDIX

8

Batch Jobs

Batch Jobs

From the SAS IT Service Level Management GUI, you can run the ETL process or a report specification. However, you can also set up these tasks so that they run in batch mode. You can then use a scheduling tool (specific to your operating environment) to run your ETL and report-generating tasks when you require. In order to accomplish this, you must do the following:

1. Execute one successful run of the ETL and report-generating tasks in order to create the job code for these tasks.
2. Create the batch file, script, or job that will execute the ETL and report job code.
3. Schedule the ETL and report tasks to run in batch mode.

These steps are explained in the following sections:

- “Running Job Code in Batch Mode” on page 298
- “Scheduling Jobs” on page 304

In addition, this appendix provides information about the following topics:

Running Job Code in Batch Mode

SAS IT Service Level Management uses job code to run the ETL process and to run the report generation process.

□ The ETL job code processes the new data into the contract database in accordance with the current definitions in the contract. For information about *creating* the ETL job code in the GUI, see "Creating the ETL Job Code" on page 89.

□ The reporting job code produces report output that reflects the report specifications that you defined in the Reports workspace of the GUI. For information about *creating* the report job code in the GUI, see "Creating Report Job Code" on page 132.

Depending on your operating environment, you must make a file, script, or job in order to run the ETL job code and the report job code. (It is a good idea to make separate files, scripts, or jobs to run these processes. This enables you to set up the report process to be run conditionally, only if the ETL process executed successfully.)

The file, script, or job must be specific to the operating system on the SAS IT Service Level Management server where the environment database and contract database(s) are located.

□ To make a file for the Windows operating system, see "Batch Jobs for Windows" on page 298.

□ To create a script to use for the UNIX operating system, see "Batch Jobs for UNIX" on page 300.

□ To make a job for the z/OS operating system, see "Batch Jobs for Z/OS" on page 302.

Note: For assistance, see the SAS IT Service Level Management administrator at your site. At some sites, alternative commands or SAS options may be required. △

Batch Jobs for Windows

To make a batch program that runs ETL job code on the Windows operating environment, perform these steps:

1 Run the ETL job code one time in order to generate the SAS code that will be inserted into your batch file. From the command line in the Windows operating environment, use this command to run the job code:

```
sas.exe -sysin filename -icon -nosplash -noxwait
```

where *filename* is the path and name of the file that contains the ETL job code that you wish to run.

2 Check that the job code runs without error messages.

3 Then follow the instructions in "Making a Batch File for Windows" on page 299.

After you have created the file that will run the ETL job code, create the file that runs the report job code. To do so, perform these steps:

1. Run the report job code one time in order to generate the SAS code that will be inserted into your batch file. From the command line in the Windows operating environment, use this command to run the job code:

   ```
   sas.exe -sysin filename -icon -nosplash -noxwait
   ```

   where *filename* is the path and name of the file that contains the report job code that you wish to run.
2. Check that the job code runs without error messages.
3. Then follow the instructions in "Making a Batch File for Windows" on page 299.

Making a Batch File for Windows

To make a file that you can schedule (from the command that has just successfully run the ETL or the report job code), modify the following .bat file (sasjob.bat). This file is in a folder that can be found at this location: ...*the_location_where_sas_is_installed_at_your_site*\itslm\sasmisc.

Note: This .bat file starts SAS in batch mode from the DOS command line. The comments within the code describe how to modify and use the file. △

```
@echo off
Rem - -- - - - - - - - - - - - - - - - - - - - - - - - - - - -- -- -- --
Rem Copyright (c) 2005 by SAS Institute Inc., Cary, NC 27513
Rem Name:  slmjob.bat
Rem Doc: Run SAS on Windows with proper options for SAS IT
Rem      Service Level Management batch mode
Rem
Rem Usage: sasjob.bat infile
Rem    where "infile" is the simple (path is not required
Rem    if in the current working directory)
Rem    first name of the files read and written by SAS:
Rem       .sas is appended to "infile" for the input program
Rem       .log is appended to "infile" for the SAS log
Rem Example: cd \mypgms
Rem          sasjob.bat mysaspgm
Rem - -- - - - - - - - - - - - - - - - - - - - - - - - - - - -- -- -- --
set INFILE=%1%

Rem The following line is broken into three lines so that you can
Rem  see the right end of the command. Make these three lines into
Rem  a single line before you run.
"C:\program files\sas\sas 9.1\sas.exe"
    -sysin %INFILE% -awstitle "SAS ITSLM Batch"
    -icon -nosplash -noxwait -noterminal

Rem Note common exit codes
Rem The following line is broken into two lines so that you can
Rem  see the right end of the command. Make these two lines into
Rem  a single line before you run.
for %%c in (0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16)
    do if errorlevel %%c set EXITCODE=%%c
```

```
Rem Did SAS terminate gracefully?
Rem The following line is broken into two lines so that you can
Rem  see the right end of the command. Make these two lines into
Rem  a single line before you run.
find "NOTE: SAS Institute Inc., SAS Campus Drive"
      %INFILE%.log > nul
if not errorlevel 1 goto NORMTERM
echo SAS was interrupted before completing termination processing.
goto EXIT

:NORMTERM

Rem Go scan log when exitcode is 0 or 1.
if %EXITCODE%==0 goto SCANLOG
if %EXITCODE%==1 goto SCANLOG

:ERRORS
echo SAS terminated with errors (EXITCODE ge %EXITCODE%).
goto EXIT

:SCANLOG
Rem Scan the SAS log since we might still have errors with
Rem  zero exitcode.
find "Errors printed on page" %INFILE%.log > nul
if not errorlevel 1 goto ERRORS
echo SAS terminated without errors (EXITCODE eq %EXITCODE%).
goto EXIT

:EXIT
%echo on
```

Batch Jobs for UNIX

To make a UNIX shell script that runs ETL job code in batch mode on the UNIX operating environment, perform these steps:

1. Run the ETL job code one time in order to generate the SAS code that will be inserted into your batch script. From the command line in the UNIX operating environment, use this command to run the job code:

   ```
   sas filename -noterminal -xrm 'SAS.startSessionManager:false'
   ```

   where *filename* is the path and name of the file that contains the ETL job code that you wish to run.
2. Check that the job code runs without error messages.
3. Then follow the instructions in "Making a Script for UNIX" on page 301.

After you have created the script that will run the ETL job code, create the file that runs the report job code in batch mode. To do so, perform these steps:

1. Run the report job code one time in order to generate the SAS code that will be inserted into your batch script. To do so, from the command line in the Windows operating environment, use this command:

   ```
   sas filename -noterminal -xrm 'SAS.startSessionManager:false'
   ```

where *filename* is the path and name of the file that contains the report job code that you wish to run.

2 Check that the job code runs without error messages.

3 Then follow the instructions in "Making a Script for UNIX" on page 301.

Making a Script for UNIX

To make a UNIX shell script that you can schedule (from the command that has just successfully run the ETL or the report job code), modify the following UNIX shell script (slmjob.sh). This script is in a folder that can be found at this location: ...*the_location_where_sas_is_installed_at_your_site*/misc/itslm.

Note: This UNIX shell script (slmjob.sh) starts SAS in batch mode from the UNIX command line and displays a message that is based on the return code from the SAS session. The comments within the code describe how to modify and use slmjob.sh. △

```
#!/bin/ksh
# - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -+
# Copyright (c) 2004 by SAS Institute Inc., Cary, NC 27513
# Name:    slmjob.sh
# Doc:     Runs SAS with proper options for SAS IT Service Level
#          Management batch mode
#
# The input file contains the job code.
# - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -+

PROGNAME='sas'

# - - - - - - - - - - - - - - - - - - - +
# First, be helpful
# - - - - - - - - - - - - - - - - - - - +
if test "$#" = "0"
   then echo " "
        echo "Usage: $PROGNAME (sasfile)"
        echo " "
        echo "where 'sasfile' is the simple (no dots, no dirname)' "
        echo "first name of the files read and written by SAS:"
        echo " "
        echo "  .sas is appended to 'sasfile' for the input program"
        echo "  .log is appended to 'sasfile' for the output SAS log"
        echo " "
        echo "Example: cd/u/myid/mypgms"
        echo "         slmjob.sh mysasprog"
        echo " "
        echo "This reads the SAS program in /u/myid/mypgms/mysasprog.sas "
        echo "and creates a SAS log in /u/myid/mypgms/mysasprog.log."
        exit 0
        fi

# - - - - - - - - - - - - - - - - - - - - - - -- -- -- -- -- --  +
# Begin
# - - - - - - - - - - - - - - - - - - - - - - -- -- -- -- -- --  +
# Change the single quotes around date to tip leftward +
```

```
# so that date will be evaluated.                       +
echo $PROGNAME: Starting at ‘date‘

 INFILE=$1

rm $INFILE.log 2> /dev/null

# - - - - - - - - - - - - - - - - - - - - +
# Start SAS
               sas -noterminal $INFILE.sas

# - - - - - - - - - - - - - - - - - - - - +
# Advertise exit code
# - - - - - - - - - - - - - - - - - - - - +
EXITCODE=$?

# Did SAS terminate gracefully?
grep "NOTE: SAS Institute Inc., SAS Campus Drive" $INFILE.log > /dev/null
if test $? -ne 0
   then echo SAS was interrupted before completing termination processing.
fi

# Complain appropriately based on the exit code.
if test $EXITCODE -lt 2
   then # We can still have errors even with exit code of 0 or 1
        grep "Errors printed on page" $INFILE.log > /dev/null
        if test $? -eq 0
           then echo "SAS terminated with errors."
           else echo "SAS terminated without errors."
        fi
   else echo "SAS terminated with errors: EXITCODE eq $EXITCODE"
fi

# - - - - - - - - - - - - - - - - - - - - -- -- -- -- -- --  +
# That is all
# - - - - - - - - - - - - - - - - - - - - -- -- -- -- -- --  +
# Change the single quotes around date to tip leftward +
# so that date will be evaluated.                       +
echo $PROGNAME: Ending at ‘date‘

exit $EXITCODE
```

Batch Jobs for Z/OS

To make a job that runs ETL job code in batch mode on the z/OS operating environment, perform these steps:

1. From the GUI, run the ETL job code one time in order to generate the SAS code that you will insert into your batch job.
2. Check that the job code runs without error messages.
3. Then follow the instructions in “Making a Job to Run in Batch Mode on Z/OS” on page 303.

After you have created the job that will run the ETL job code, create the job that runs the report job code in batch mode. To do so, perform these steps:

1. From the GUI, run the report job code one time in order to generate the SAS code that you will insert into your batch job.
2. Check that the job code runs without error messages.
3. Then follow the instructions in "Making a Job to Run in Batch Mode on Z/OS" on page 303.

Note: This z/OS job starts SAS in batch mode and displays a message that is based on the return code from the SAS session. The comments within the job describe how to modify and use CMSLMJCL. △

Making a Job to Run in Batch Mode on Z/OS

To make a z/OS job that you can schedule (from the command that has just successfully run the ETL or the report job code), modify the following job. This job is a member of the CPMISC PDS that can be found at:
THE.LOCATION.WHERE.SAS.IS.INSTALLED.AT.YOUR.SITE

Note: Be sure to specify an appropriately large work size in your SAS invocation command. △

```
//CMSLMJCL JOB (Room101),'SLM Batch Job',NOTIFY=USERID,TIME=(20,00)
/*JOBPARM FETCH
//*
//*
//* - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
//*
//* Copyright (C) 2003 by SAS Institute Inc., Cary, NC 27513
//*
//* - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
//* CMSLMJCL: SAS(R) IT Service Level Management job template for
//* submitting z/OS batch work for ETL processing and reporting.
//*
//* This batch job is named CMSLMJCL. You can use it to submit
//* jobs generated by SAS IT Service Level Management.
//*
//* Because this batch job and the other programs require
//* customizations for your site, consider making copies of the
//* originals and modifying the copies.
//*
//* Before you submit this job, search for all occurrences of CUSTOMIZE.
//* At each occurrence, there is a description of what to customize.
//*
//* - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
//*
//* CUSTOMIZE: The sample JOB statement above runs the job. The sample
//* EXEC statement below runs the SAS procedure. The SAS procedure runs
//* the code that is in the SYSIN stream. The SYSIN stream includes one
//* or more of the three programs.
//*
//* Make any changes to the JOB and EXEC statements that are necessary
//* for the JCL to be appropriate for your site.
```

```
//*
//* - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
//SYSIN DD DATA, DLM='$$'
//*
/* CUSTOMIZE: Place your SAS code after this comment. */
/* The last line of this job is the $$ in column 1.   */
//*
$$
//
```

Scheduling Jobs

After you have created the Windows .bat file, the UNIX .sh script, or the z/OS job for the ETL and reporting job codes, you can schedule them to run on the SAS IT Service Level Management server.

Typically, sites schedule the jobs to run one time a day, at night. However, you can schedule them to run more or less frequently, and at whatever time works best for your site.

Note: Schedule the ETL job to run before the reporting job, and use conditional scheduling to make sure that the ETL job completes successfully before the reporting job starts. △

For more information about scheduling:

- On UNIX: refer to your *cron* and *crontab* man pages.
- On Windows: refer to the documentation for your task scheduler.
- On z/OS: refer to your scheduling system's documentation.

Be sure to schedule backups of the contract database(s) and environment database(s). For more information, see "Backing Up Databases" on page 304.

Backing Up Databases

The fastest way to back up a database is to use an operating system utility. You can back up a database by using the following utilities:

- An appropriate Windows utility on Windows. For more information, see "Backing Up a Database (Windows)" on page 304.
- An appropriate UNIX utility on UNIX. For more information, see "Backing Up a Database (UNIX)" on page 305.
- An appropriate z/OS utility on z/OS. For more information, see "Backing Up a Database (Z/OS)" on page 305.

Backing Up a Database (Windows)

There is no preferred Windows utility. On Windows, you can use your site's standard backup utility.

Backing Up a Database (UNIX)

There is no preferred UNIX utility. Here is an example of backing up a contract database on UNIX:

- To make the backup for a contract database named "mycdb" to a file "/u/sasabc/mycdb.tar.Z", you can use these commands:

```
cd <parent-directory-of-mycdb>
tar cf - mycdb | compress > /u/sasabc/mycdb.tar.Z
```

  where

  - <parent-directory-of-mycdb> is the full pathname of the contract database
  - mycdb is the contract database name (without the path)
  - /u/sasabc/mycdb.tar.Z is an arbitrary filename, in which the .Z indicates that the file is compressed and the .tar indicates that the tar utility was used to convert a directory structure to a flat file.

- To restore the contract database above from the tar file above, you can use these commands:

```
cd <parent-directory-of-mycdb>
uncompress -c /u/sasabc/mycdb.tar.Z | tar xf -
```

Backing Up a Database (Z/OS)

There is no preferred z/OS utility. Here is an example of backing up a contract database on z/OS:

```
//BACKUP  EXEC PGM=ADRDSSU                          #1
//TAPE     DD DSN=your.database,DISP=(OLD,KEEP),    #2
//            UNIT=cart,RETPD=30
//SYSPRINT DD SYSOUT=A
//SYSIN    DD *
  DUMP DATASET(INCLUDE(your.database.*))            #3
              OUTDD(TAPE) - SHARE COMPRESS
//
```

Notes:

1 This example uses the IBM DFDSS program. You might prefer to use the backup procedure that is commonly used at your site.

2 You might want to change characteristics on the TAPE allocation according to your needs or according to site-specific rules.

3 You can use the DFDSS DUMP command to back up volumes and data sets, and you can then use the RESTORE command to recover them. To make incremental backups of your data sets, perform a data set DUMP with RESET specified and filtering on the data set-changed flag. For more information about this facility, refer to the IBM manual *Data Facility Data Set Services: Reference V2R5 (SC26–4389)*. You might want to change the DUMP command according to your needs or according to site-specific rules.

Special Considerations for Z/OS Operating Environments

Some special considerations apply to sites that run SAS IT Service Level Management on a z/OS operating environment:

Locating Sample Jobs

The miscellaneous file that contains sample jobs is located on the z/OS server that is running SAS IT Service Level Management. The file is in a PDS named as follows:
<*prefix*>.ITRM.CPMISC
(where <*prefix*> is the location where SAS is installed at your site.)

Preallocating Space for the Contract Database

Before creating a contract database, you must preallocate the space for it. CMSLMALC contains a job step that allocates and initializes space for an environmental database or a contract database on z/OS. For information about the CMSLMALC job, see the instructions in the CMSLMALC job.

The CMSLMALC job is in the CMSLMALC member in the PDS that can be found at this location: <*prefix*>.ITRM.CPMISC where <*prefix*> is the location where SAS is installed at your site.

Converting Text Files from EBCDIC to ASCII

You can convert text files that were created with the EBCDIC code set (on z/OS) to files that use the ASCII code set. This enables the report files that are generated by SAS IT Service Level Management to be copied to a Windows or UNIX Web server for convenient viewing. To convert the file, perform the following steps:

1. Copy the CMSLME2A member from the CPMISC PDS to a location in the UNIX Systems Services (USS) environment on the z/OS.

   Note: The CPMISC PDS can be found at this location: <*prefix*>.ITRM.CPMISC (where <*prefix*> is the location where SAS is installed at your site.) △

2. Access the USS environment on z/OS either from your console or by logging into your z/OS system from a UNIX shell.

3. Use the chmod command to change the permissions on the CMSLME2A script so that you (the owner of the file) have execute permissions.

   Note: The command to let the owner of the file have read, write, and execute permission is:

   ```
   chmod 744
   ```

   With this setting, everyone else has read-only permission. △

4. Follow the instructions in the comments sections of the cmslme2a script to modify the syntax that is needed.

Sharing the Environment Database

If you wish, you can use a SAS/SHARE server to share a SAS IT Service Level Management environment database among multiple users or jobs. For information about the program that sets up a SAS/SHARE server for use with a SAS IT Service Level Management environment data base, see "Using CPSLMSHR" on page 307.

If you set up a SAS/SHARE server for use with a given SAS IT Service Level Management environment database, then the following statements apply:

- The client GUI, the ETL program, and the reporting program can access that environment database concurrently as long as each task is working with a different contract.

If you do not set up a SAS/SHARE server for use with a given SAS IT Service Level Management environment database, then the client GUI, the ETL program, and the reporting program must not access that environment database concurrently, even if the tasks are working with different contracts.

Note: For more information about SAS/SHARE, see the *SAS/SHARE User's Guide*. △

Using CPSLMSHR

For more information about the CPSLMSHR program, which sets up a SAS/SHARE server for use with a SAS IT Service Level Management environment database, see the instructions in the CPSLMSHR job.

- On a z/OS server, the SAS IT Service Level Management server software has a PDS named *THE.LOCATION.WHERE.SAS.IS.INSTALLED.AT.YOUR.SITE*.ITRM.CPMISC. The CPSLMSHR job is in the CPSLMSHR member in that PDS.
- On a UNIX server, the SAS IT Service Level Management server software has a directory named ...*the_location_where_sas_is_installed_at_your_site*/misc/itslm. The CPSLMSHR job is in the cpslmshr.sas file in that directory.
- On a Windows server, the SAS IT Service Level Management server software has a folder named ...*the_location_where_sas_is_installed_at_your_site*\itslm\sasmisc. The CPSLMSHR job is in the cpslmshr.sas file in that folder.

APPENDIX

9

Backloading Data

Backloading Data

This appendix explains how to backload archived data from a SAS IT Resource Management PDB into a SAS IT Service Level Management contract database. This procedure can benefit existing customers who have been using SAS IT Resource Management. These customers might have collected several months of resource management data that has aged out of the DETAIL level of the SAS IT Resource Management PDB but has been archived in some form. Backloading data enables customers to use this archived data in SAS IT Service Level Management by rapidly filling in the historical values for the contracts they have specified.

In order to bring data into SAS IT Service Level Management, the data must currently reside in the DETAIL level of a SAS IT Resource Management PDB. Typically, the components that are specified in SAS IT Service Level Management will pull the newest data from the production SAS IT Resource Management source PDB.

Note: We recommend that you move the archived SAS IT Resource Management DETAIL data to a separate PDB. If the archived data is too voluminous to retrieve and use in one large operation, this can be done in chunks. △

The following figures show how backloading works.

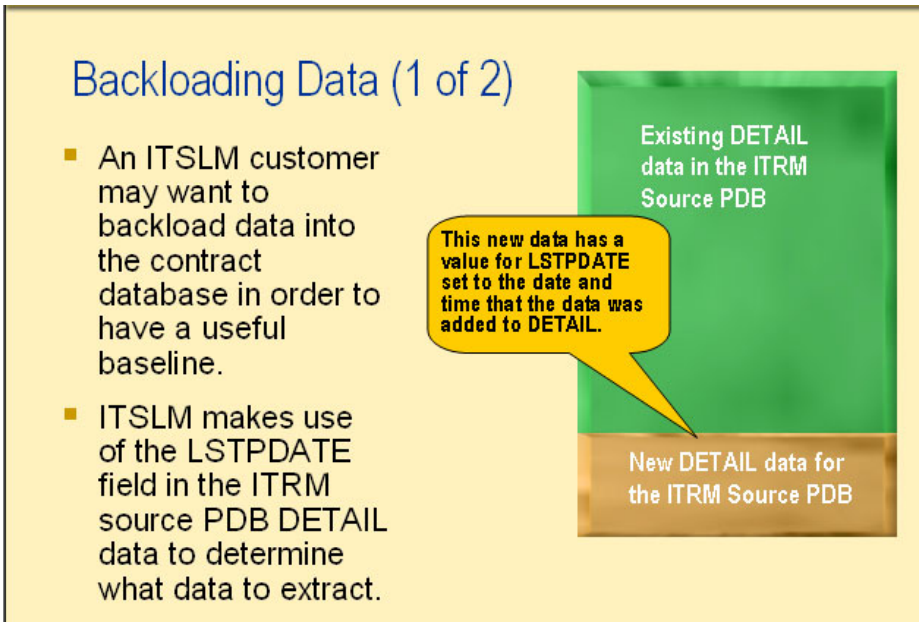

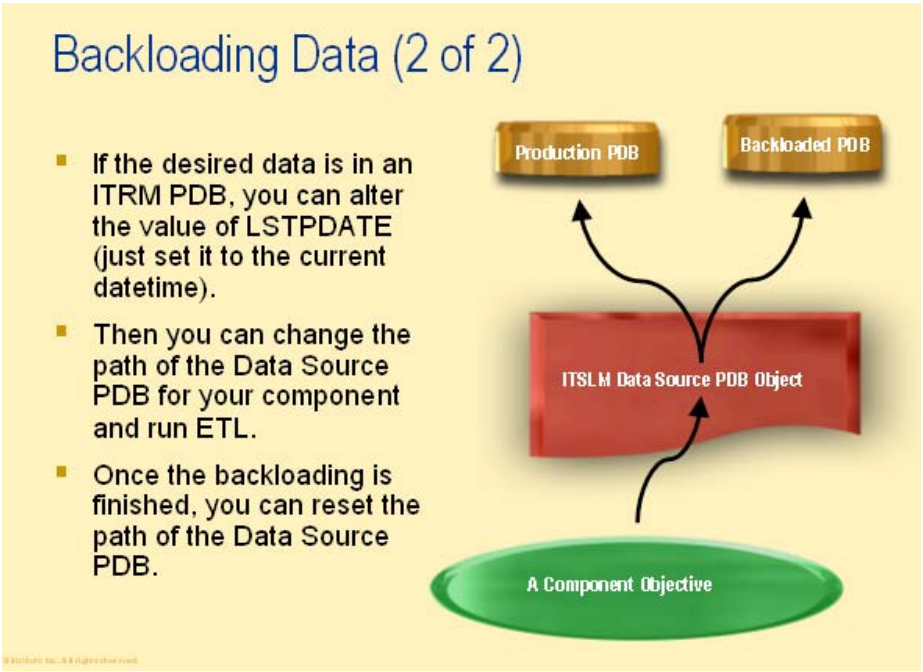

To backload data, perform these steps:

1 In SAS IT Resource Management, restore the archived data to a separate PDB. In the following example, this PDB is called ArchivePDB.

2 Reset the LSTPDATE column in the data that was restored in the previous step. The SAS IT Service Level Management administrator can edit the value of the LSTPDATE column in the restored data.

The LSTPDATE column specifies when this data was originally stored in the DETAIL level of the PDB. SAS IT Service Level Management keeps track of the value of LSTPDATE each time it extracts data. This enables it to pull only new data from the data source PDB.

In order to bring the archived data into SAS IT Service Level Management regardless of the value of LSTPDATE, edit LSTPDATE in the PDB that contains the archived data.

The following example shows how to use the %CPEDIT macro in SAS IT Resource Management to set the value of LSTPDATE to the current system datetime.

Example;

a Put the following SAS code into an external file. (In this example, the code is put into c:\lstpdate.sas.)

```
lstpdate = dhms("&sysdate9"d,
           hour("&systime"t),
           minute("&systime"t),
           second("&systime"t));
```

b Execute SAS code to change the value of LSTPDATE.

```
%cpstart(mode=batch,access=write,pdb=ArchivePDB);
%cpedit(table=testtbl,level=detail,
        srcfile=c:\lstpdate.sas,zobsok=n);
```

where ArchivePDB is the name of the PDB that contains the restored archived data.

3 Change the path of the SAS IT Service Level Management data source PDB to point to the PDB that contains the archived data. (In this example, the PDB is called ArchivePDB.)

4 Once you have done this for all the data source PDBs that are used by the SAS IT Service Level Management contract, run the ETL job for this contract.

5 Repeat this procedure as necessary for each chunk of archived data.

When you are finished, change the path of the SAS IT Service Level Management data source PDB to point back to the production PDB path.

APPENDIX
10

Best Practices

Best Practices

This appendix contains recommendations that can help you achieve the most efficient use of the SAS IT Service Level Management software. It consists of "Examples of Date Filter Usage" on page 313.

Examples of Date Filter Usage

You can use date filters to provide additional flexibility to your use of SAS IT Service Level Management. They can have an important effect on the way data is collected. Here are some examples:

- Customers who are mainly concerned with high-volume days might focus on Monday through Friday from 9:00 to 17:00.

  In this case, the date filter pattern might be **`Weekly`** with **`Monday`**, **`Tuesday`**, **`Wednesday`**, **`Thursday`**, and **`Friday`** selected. The effective and expiration times might be 9:00 and 17:00, respectively.

- The customer organization might have one kind of load on weekdays and another kind of load on weekends; so the contract might specify one set of targets for weekday data and another set of targets for weekend data.

  In this case, there might be a pair of components. One component might have a **`Weekly`** pattern with **`Monday`**, **`Tuesday`**, **`Wednesday`**, **`Thursday`**, and **`Friday`** selected. The other component might have a **`Weekly`** pattern with **`Saturday`** and **`Sunday`** selected.

  And there might be a service that has these two components as child nodes, and that calculates the value of each service indicator as the mean of the corresponding two component indicators. Additionally, these two component objectives could be summarized by having a parent node, such as an SLA, service, or contract, that calculates mean value.

Index

Numbers

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

W

Y

Z

Your Turn

If you have comments or suggestions about *SAS IT Service Level Management 2.1: User's Guide*, please send them to us on a photocopy of this page, or send us electronic mail.

For comments about this book, please return the photocopy to

SAS Publishing
SAS Campus Drive
Cary, NC 27513
E-mail: **yourturn@sas.com**

For suggestions about the software, please return the photocopy to

SAS Institute Inc.
Technical Support Division
SAS Campus Drive
Cary, NC 27513
E-mail: **suggest@sas.com**