Sterling Control Center™

Implementation Guide

Version 5.2

Sterling Commerce
An IBM Company

Sterling Control Center Implementation Guide
Version 5.2

First Edition

Contents

Chapter 1

About the *Sterling Control Center Implementation Guide*

The *Sterling Control Center Implementation Guide* is for anyone who is tasked with implementing Sterling Control Center and needs to know the "big picture" about Control Center. This guide provides the following information:

- The features and benefits of Sterling Control Center
- An overview of the concepts and components that comprise Sterling Control Center
- Questions to help you identify the high-level business objectives you have for Sterling Control Center
- An explanation of the building blocks you can use to define the work Control Center will perform to meet your objectives and the interrelationships between the building blocks
- A high-level explanation of the types of information Control Center generates about the servers in your environment
- Sample implementation scenarios that illustrate how the building blocks are used to meet business objectives
- A best practices, ordered task list that provides planning considerations and references to planning tools and documentation that will help you complete each high-level task
- Planning worksheets you can use when planning your Control Center implementation

This information is provided in a layered approach so that each piece builds on the information in the previous piece.

To help you understand how the terminology used in this guide relates to Control Center, see *Sterling Control Center Terms and Concepts* on page 89.

Chapter 2

Features and Benefits

Sterling Control Center is a centralized monitoring and management system that gives operations personnel the capability to continuously monitor business activities for Connect:Direct, Connect:Direct File Agent, Connect:Enterprise, Sterling Integrator, and Sterling File Gateway, and for many FTP servers, across the enterprise. Sterling Control Center also allows you to manage the configurations and licenses of Connect:Direct servers.

Features

Sterling Control Center enhances operational productivity and improves the quality of service for file transfer and activities in your environment from one central location. It accomplishes this through:

- Service level management
- Asset management
- Configuration management

Service Level Management

Sterling Control Center can help you answer questions about your managed file transfer environment, such as:

- Did my business process run on time?
- Did my file transfer take place when it should have?
- Are my servers running okay?

Control Center gives you tools to effectively monitor and manage your environment by giving you a common, centralized view of that environment. This view enables you to offer higher levels of service to your internal and external customers. Control Center accomplishes this by:

- Providing a real-time view of all your file transfer servers across products, platforms, and locations. To facilitate monitoring “like” servers, you can group them into server groups, by business unit or location for example, for a single view of system-wide activity.
- Monitoring activities such as business processes and file transfers.
- Monitoring the overall health of the environment in terms of server status, adapter status, and cluster health.
- Using a common set of capabilities to create an early warning system for exceptions by:
  - Ensuring critical processing windows are met through service level criteria (SLCs) you set up for your environment.
  - Reducing impact on downstream processing by verifying that expected processing occurs based on rules you define that are triggered by server events.
  - Providing proactive notification for at-risk business processes in the form of e-mails, SNMP traps sent to an SNMP host, and alerts viewable through the Control Center consoles.
- Consolidating information for throughput analysis, capacity planning, post-processing operational or security audits, and workload analysis to help ensure that your file transfer environment is functioning at a high level.
- Reducing the risk of error associated with manual system administration, including the requirement to log on to each individual server to view activity, or having to separately configure servers for error and exception notification.

Asset Management

Sterling Control Center helps you answer questions about your server assets, such as, “Where is my software installed and running?” and “Is it in compliance with license agreements?” It helps you manage assets by means of:

- **Asset tracking**—Track network assets by capitalizing on the server monitoring capabilities of Sterling Control Center. Using Guided Node Discovery (or simply, Node Discovery), you can find all Connect:Direct servers deployed in your network.
- **License management**—Ensure that your server licenses are up to date, and facilitate license distribution to the managed Connect:Direct servers in your environment.

  Sterling Control Center allows you to import updated server licenses to a central license management repository and, on an ad hoc basis, push them to managed Connect:Direct servers. You can also configure Control Center to monitor a POP3 or IMAP mailbox for e-mails containing license key file attachments. When Control Center finds a license key file attachment, it validates the license and automatically imports it into the license repository, so the next time you are ready to push licenses to your servers, you have the most current license available.

Configuration Management

Sterling Control Center helps you answer questions about your Connect:Direct servers such as, “Are my they configured correctly” and “Do they comply with our security policy?”, Control

Center provides you with a centralized, simplified means of managing your Connect:Direct for UNIX, Windows, and z/OS servers by:

- Offering a common interface for managing and auditing server configurations.
- Normalizing parameters across platforms that might have different names and value pairs.
- Allowing platform-specific syntax checking and easy-access tooltip help.
- Providing a means for updating, viewing, auditing, and tracking versions (including rollback functionality) of configuration data for Connect:Direct servers, such as netmap nodes, functional authorities, and initialization parameters.
- Allowing you to schedule when and what time configuration jobs are run.
- Generating an audit log to identify what changes were made and who made them. (This feature also applies to Sterling Control Center system configuration.)
- Managing the configuration of Secure+ Option on remote Connect:Direct servers.
- Alerting you about Secure+ Option certificate expiry.
- Allowing you to quickly identify Connect:Direct servers that do and do not have the Secure+ option installed via information displayed in a Server List View.

Benefits By Server Type

Sterling Control Center provides the following primary benefits for each server type:

Server type	Primary Benefits
Connect:Direct	◆ Provides centralized visibility and control of large-scale, distributed Connect:Direct server environments by enabling you to consolidate and collect data for a variety of purposes. ◆ Allows you to release or delete Processes from a central location. ◆ Lets you configure notification about processes, or steps in processes, that did or did not occur or are late. ◆ Lets you monitor the queue depths of the Execution, Hold, Timer, and Wait queues. ◆ Enables you to manage multiple servers on multiple platforms, including configuration of Connect:Direct and Connect:Direct Secure+. ◆ Provides for centralized asset management (software version and license key).

Server type	Primary Benefits
Connect:Direct File Agent	◆ Allows you to know when a Connect:Direct File Agent submits processes to a Connect:Direct server and to what server it submits those processes. ◆ Allows you to know when processes are not submitted to a Connect:Direct server. ◆ Enables you to know when a user has updated the configuration data for a Connect:Direct File Agent.
Connect:Enterprise	◆ Provides visibility to files transferred into and out of mailboxes. ◆ Allows monitoring of daemons.
Sterling Integrator	◆ Offers centralized visibility into the business processes and file transfer activities of your trading partners in a large, clustered, multi-node environment. ◆ Allows you to rerun business processes from a central location. ◆ Offers you the flexibility to manipulate monitoring and notification without redundant coding. ◆ Lets you configure notification about processes, or steps in processes, that did or did not occur or are late. ◆ Lets you monitor queue depths for Sterling Integrator's 10 queues. ◆ Enables you to view Sterling Integrator license details, Java environment details, location of Sterling Integrator installation, adapter properties and configuration, and perimeter service configuration.
Sterling File Gateway	◆ Provides enhanced, granular control over monitoring and alerting options compared to what is available in Sterling File Gateway. ◆ Enables monitoring of arrived file events, route events, and delivery events. ◆ Lets you configure notification about files that did or did not occur or are late. ◆ Allows monitoring of Mailbox Service and Mailbox Browser Interface (MBI).

Server type	Primary Benefits
FTP server	◆ Provides visibility of files transferred into and out of FTP servers so you can determine this activity complies with your corporate policies. ◆ Allows you to centrally monitor FTP usage. ◆ Lets you configure notification about files that did or did not occur or are late. ◆ Allows you to identify where FTP usage violates corporate policy, so you can move it to a more secure and reliable solution, such as Sterling Managed File Transfer.

Chapter 3

Understanding Sterling Control Center Concepts and Components

A basic understanding of Sterling Control Center components and how they work together will help you plan and implement Control Center in your environment.

Tiered Application Architecture

Sterling Control Center uses an application model in which different areas of functionality use separate sets of resources to operate. These areas are called tiers, or layers, and are usually arranged with central or core functionality on the bottom-most layers and GUI services that depend on core functionality at the top-most layers.

Because they can be located on physically different servers with only minor configuration changes, the layers can scale out and handle more server load. In addition, what each layer does internally is separate from the other layers. This makes it possible to change or update one layer without modifying the others.

The Sterling Control Center engine performs heavy-duty monitoring of large numbers of managed servers. This can cause a considerable load on the system. Separating core application functionality into different tiers allows Control Center to optimize each server task. For example, GUI services are designed to process in a separate tier from server monitoring.

Sterling Control Center is deployed across the following application tiers:

- **Presentation** (GUI)—This tier displays information via graphical user interfaces (GUIs). Users log on to the Sterling Control Center engine through this tier to configure Sterling Control Center monitor and manage servers. The Sterling Control Center engine and the GUIs can be installed on different operating systems.
- **Business** (Sterling Control Center Engine)—This tier generates events from data retrieved from monitored servers. The engine passes events through rules and SLC services to take appropriate action. It records data in the data tier for historical purposes and presents results to the presentation tier.
- **Data** (Databases)—Historical information is stored and retrieved in this tier. Data is kept neutral and independent from application servers or business logic. Because data has its own tier, scalability and performance are improved.

Sterling Control Center Components

Sterling Control Center consists of three primary components:

- User Interfaces
- Engine
- Databases

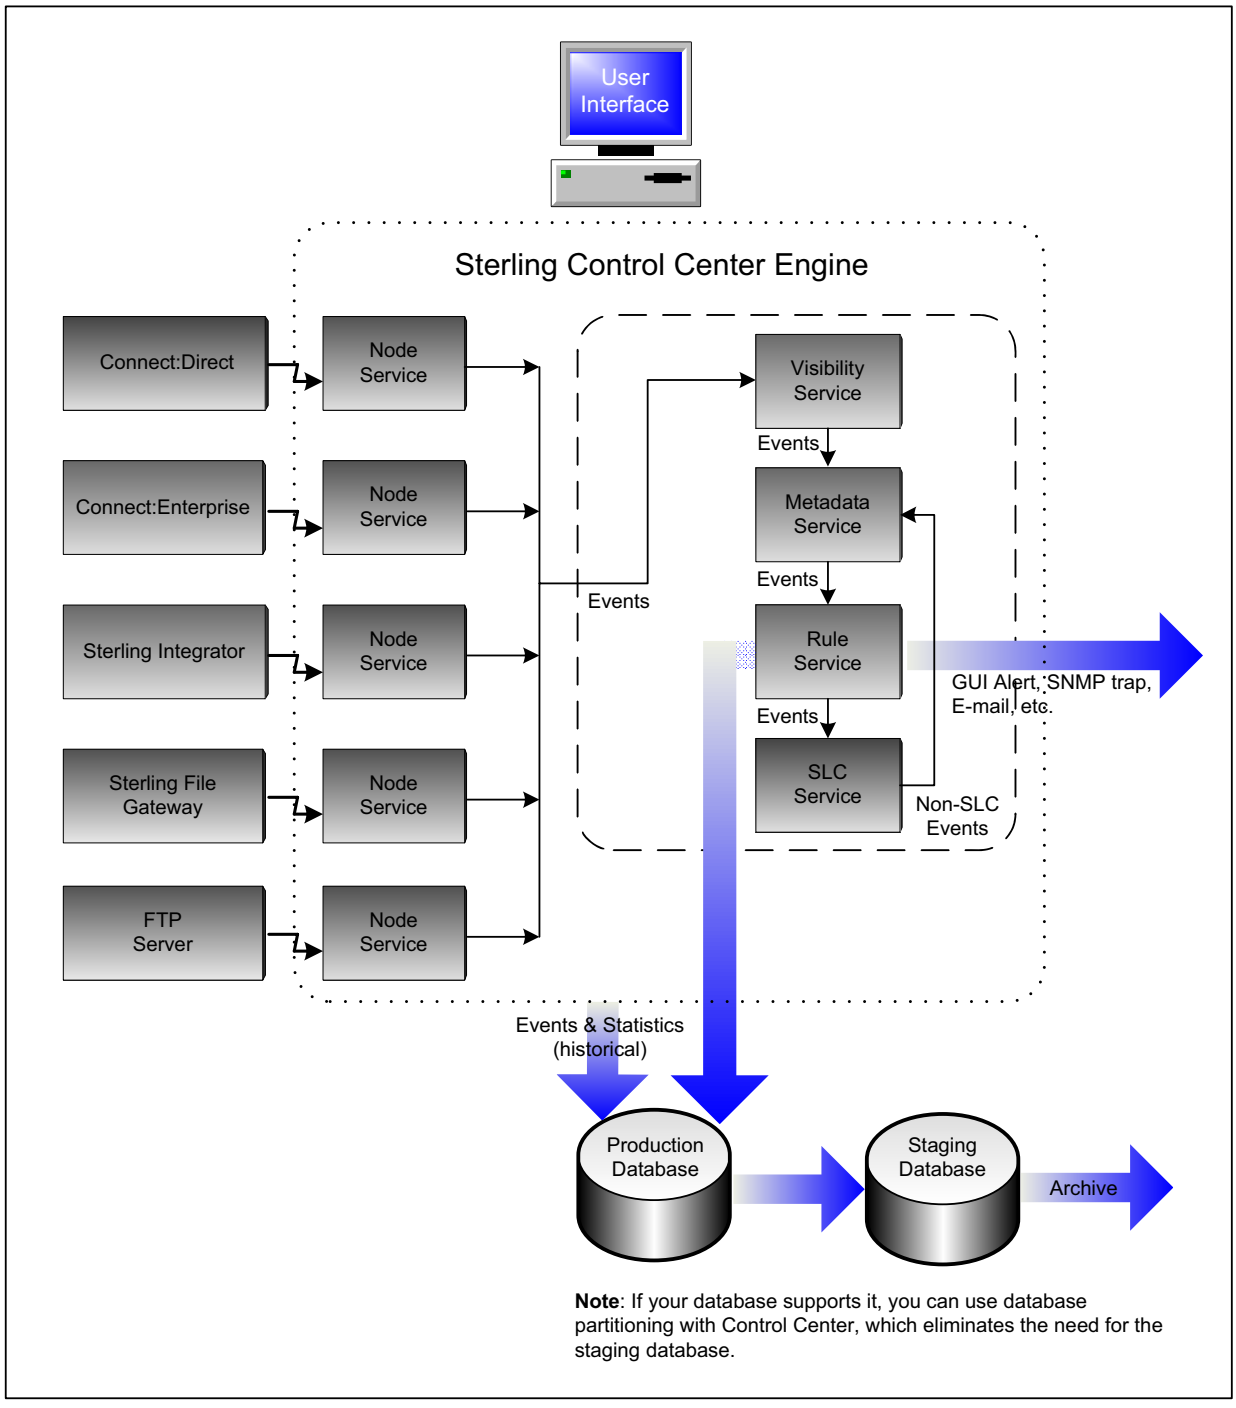

User Interfaces

The Sterling Control Center console and web console enable you to configure Sterling Control Center and Connect:Direct nodes and display information gathered from the engine via an HTTP or HTTPS connection. In addition, Sterling Control Center Mobile enables you access to a subset of Control Center functionality from your iPhone. Access to and functionality of these interfaces are limited by the role-based privileges assigned to a user.

The GUIs serve the following purposes:

- **Console**—Offers full functionality for configuring Sterling Control Center, configuring Connect:Direct servers, and monitoring/analyzing monitored servers. The console is installed locally on the computer where the engine is installed. After installation, you can access Control Center with an Internet browser by bringing up the Sterling Control Center Launch Page.
- **Web Console**—Is a lightweight version of the console that offers a subset of console functionality. It can also be launched from the Sterling Control Center Launch Page using a web browser.
- **Mobile**—Is an application that allows you to receive, view, add comments to, and handle Control Center alerts; view the status of servers; and view the status of Sterling Integrator adapters, all from your iPhone. For more information, see the *Sterling Control Center Mobile Application Guide*.

The following table describes the functionality available in the consoles. Functionality is limited by the interface (console vs. web console) and the user's role-based privileges.

Functionality	Console	Web Console	Mobile
Configuring			
Configure Connect:Direct servers using Configuration Management.	X		
Create and maintain users, data visibility groups, rules, actions, and service level criteria (SLCs), e-mail lists, calendars, and schedules.	X		
Change system settings.	X		
Set up reports.	X		
Managing			
Stop Connect:Direct servers.	X		
Manage Connect:Direct processes by taking action on queued processes (deleting, suspending, or releasing a process).	X		
Launch the Connect:Direct Browser user interface to access Connect:Direct, the Sterling Integrator Dashboard to access Sterling Integrator, and the File Gateway Console to access Sterling File Gateway.	X		

Functionality	Console	Web Console	Mobile
Monitoring and Analysis			
View server properties, such as server version and license information (expiration date, number of days before expiration), number of concurrent sessions or accounts permitted).	X	X	
View and handle alerts through the Active Alerts Monitor and the Handled Alerts Monitor. View alert properties and the rule or SLC properties associated with an alert.	X	X	X
View server activity using the Queued Activity Monitor and the Completed Activity Monitor.	X	X	
Check status using the Server Status Monitor, Daemon Status Monitor, and Adapter Status Monitor.	X	X	X
View process statistics for a particular process or for one or more servers, server groups, or a particular server type in the Statistics Viewer.	X	X	
Run and view reports.	X	X	

The consoles provide the following types of tools to assist you:

- Help—The consoles provide a full, searchable help system accessed from the Help menu. In addition, the status bar at the bottom of many dialog boxes displays valid parameter values, and tooltip help is available for parameters in the Configuration Manager. Tooltips include a short parameter definition, valid entry requirements, and default value if any.
- Wizards—When you are defining Sterling Control Center objects, such as rules, actions, and SLCs, wizards guide you through the process.

Sterling Control Center Console—The following graphic shows the main screen of the Sterling Control Center console:

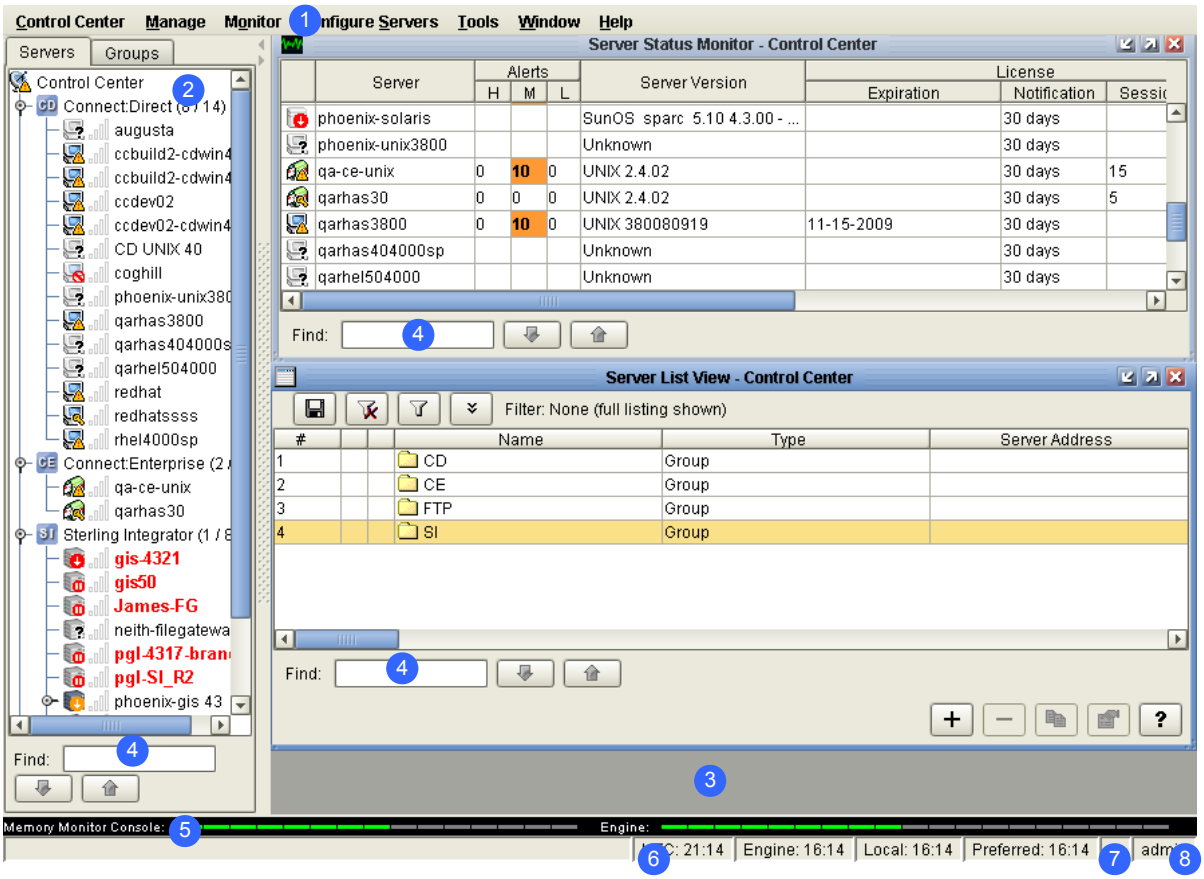

Number	Area	Description
1	Menu bar	Contains the main menu items you use in performing most Sterling Control Center tasks.
2	Node tree	Displays the listing of managed servers and server groups that have been added to Sterling Control Center. The servers are displayed on the Server tab alphabetically by server type. The server groups you have defined are displayed on the Groups tab. When you right-click a server type, individual server, or server group, you get a menu of options for that selection.
3	Work area	Displays the various available monitors and listing screens. After displaying one or more monitors or listings, you can arrange them visually by tiling them horizontally, vertically, or in a cascade.
4	Find (search feature)	Allows you to search for occurrences of text in the current window. Occurrences are highlighted in green. Click the up or down arrow to find the next (or previous) occurrence.

Number	Area	Description
5	Console and engine memory monitors	Provide a visual picture of how much available memory is being used by both the console and the Sterling Control Center engine.
6	Time indicators	Displays current time settings, including Coordinated Universal Time (UTC), as well as engine, local, and preferred times.
7	GUI connection	When a secure GUI-engine connection is used, displays a lock icon.
8	User ID	Displays the ID of the user currently logged onto the console.

Web Console—The functionality, look, and behavior of the web console differ from the Sterling Control Center console. Because the web console provides a subset of Sterling Control Center functionality (primarily monitoring), the interface has a more simplified look, with fewer items in its main screen, as shown in the following:

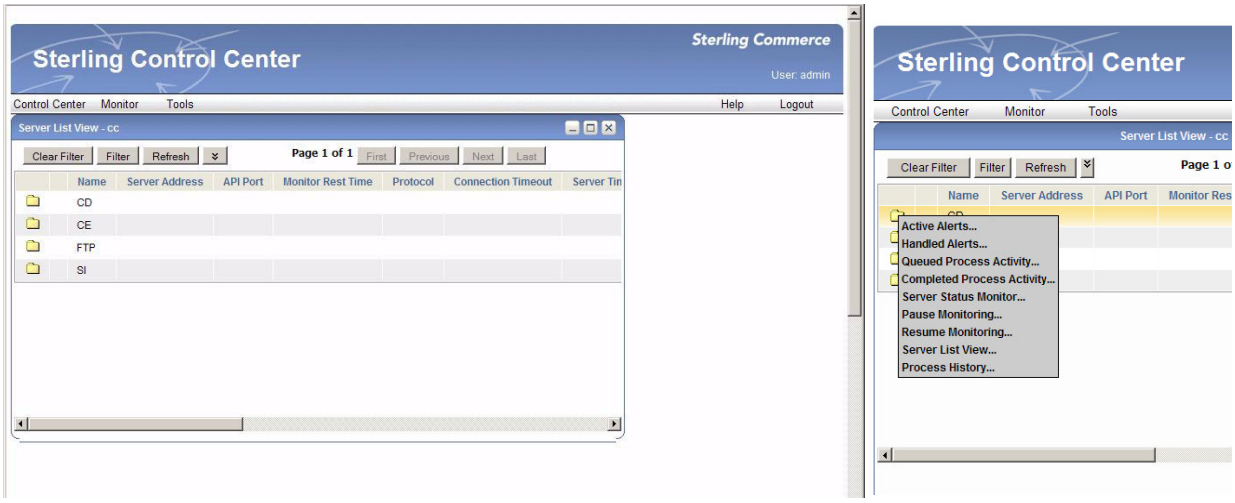

In the illustration, the screen on the left shows the web console main screen. To access the options for a server type, you left-click CD, CE, FTP, or SI. Each option you select displays a monitor window for that option. If you want to see a listing of the individual servers for a particular server type, choose the Server List View. You can then left-click a server to access the options for it. If you prefer to view servers by server group, from the menu bar, select **Monitor > Server Group View.**

Engine

The engine, which enables you to manage and monitor multiple servers on different platforms powers Sterling Control Center. It uses services, including node services, that handle the acquisition of data from the servers being monitored and managed. Other services process the visibility criteria, rules, SLCs, and metadata you define to tell Sterling Control Center what work to perform when processing the data.

Node Services—In general, there is a one-to-one mapping between node services within the Sterling Control Center engine and the servers being monitored and managed. Node services are responsible for the communications that transpire between monitored servers and Sterling Control Center.

The following table shows the server resources that Sterling Control Center accesses to retrieve data:

Server Type	Server Resource
Connect:Direct	◆ Select Statistics ◆ Select Processes
Connect:Enterprise	Information about files transferred into and out of mailboxes: ◆ Remote Connect batches ◆ Autoconnect batches ◆ Offline batches
Sterling Integrator	◆ File transfer activities ◆ Business process activities
Sterling File Gateway	◆ Arrived file events ◆ Delivery events ◆ Route events
FTP server	Information about files transferred to and from FTP servers

The servers must be configured to allow Sterling Control Center to access these resources.

Visibility Service—The visibility service applies data visibility group criteria to all events generated by the Sterling Control Center engine prior to passing them on to the metadata service.

Metadata Service—The metadata service applies enabled, active metadata rules to all events generated by the Sterling Control Center engine.

Rule Service—The rule service applies enabled, active, linked, and non-linked rules to all events generated by the Sterling Control Center engine *after* they have been processed by the metadata rule service. Rules triggered by events handled by the rule service cause associated actions to be performed: generating an e-mail, sending an SNMP trap, adding an alert indication to the event, sending a command to the server the event resulted from, or running a command or script by the Sterling Control Center engine.

SLC Service—The SLC service is responsible for generating events when things do or do not happen within a certain time frame or occur for a specified duration according to performance objectives you define.

Databases

Sterling Control Center uses the following databases to record and store information:

- **Production database**—Where Sterling Control Center records the information gathered from the monitored servers for historical purposes (for example, ad hoc select statistics and user reports). As Sterling Control Center receives information from the monitored servers, events are generated and passed through the Rule, Metadata, and SLC Services as they are being written to the database.
- **Staging database**—An optional database where Sterling Control Center stages older data for the database administrator to export to long-term storage before it is purged. If you choose not to use a staging database, you can use database partitioning. For more information, see *Database Partitioning* on page 22.

Note: If you do not create the staging database before installing Sterling Control Center or if you are going to use database partitioning, you can enter a host address of 0.0.0.0 during the Sterling Control Center product installation.

You can access the staging database for purging or archiving data without disrupting data collection into the production database. You can also establish an automated staging schedule for the production database and an automated purge schedule for the staging database.

Database Partitioning

If your database supports it, Sterling Control Center can be setup to use database partitioning, which is a best practice recommended by Sterling Commerce. This allows the data in the production database to be partitioned by date, which can improve database performance and reduce database maintenance (for example, index rebuild). When database partitioning is used, data is not moved to the staging database, eliminating the need for the staging database. For more information, see *Sterling Control Center Database Partitioning* whitepaper for more information.

Chapter 4

Defining Your Control Center Objectives

A vital part of your Sterling Control Center implementation is the process of identifying the business issues that Control Center will help you address and translating those issues into business objectives. The business objectives you have for Control Center are the beginning of your Control Center planning process. A successful Control Center implementation hinges on a comprehensive, multi-faceted plan that you formulate *before* you begin to configure the building blocks that tell Control Center the work you want it to perform.

You need to consider the following types of questions when identifying your Control Center objectives:

General

- What do I have in my environment?
  - What type of servers?
  - How many servers?
  - Performing what function?
  - Are my servers at the minimum maintenance level required to be monitored by Control Center?
  - How active are the servers in terms of number of file transfers per hour and per day?
- What historical data do I want to preserve, how much, and for how long?

Service Level Management

- What do I want to know about the health of my environment?
  - Are my servers up/down?
  - Are my server licenses about to expire?
  - How many processes are running on my servers?
  - Do I need a daily report on server activities or other information collected by Control Center?

- What is the status of my adapters?
- What is the health of my Sterling Integrator cluster?
- Have queue depths (Execution, Timer, Hold, Wait, Q0–Q9) exceeded some threshold?

- What do I need to know about my data transfers?
  - Success/failure?
  - Did the transfer happen?
    - Yes—Action to take?
    - No—Action to take?
  - Did it happen on time?
    - Yes—Action to take?
    - No—Action to take?
  - Was it bigger than X?
    - Yes—Action to take?
    - No—Action to take?
  - Was it smaller than Y?
    - Yes—Action to take?
    - No—Action to take?
  - Did it make it to the target destination?
    - Yes—Action to take?
    - No—Action to take?
  - Did it make it to the target destination on time?
  - Did it take too much or too little time whenever it ran?
  - How many transfers are failing?
  - Are my queue sizes too big?
- Overall, what actions do you want to take for the questions listed above?
  - Send e-mail
  - Send SNMP trap
  - Send e-mail to a distribution list
  - Run a program
  - Send a command to a server
- Do I want to limit the data users can view and manage?

Asset Management

- Where is my software installed and running?
- Is it in compliance with license agreements?
- When my license keys expire, do I need notifications on expiry?

Configuration Management

- Do I need to centrally manage the configurations of my Connect:Direct servers and Secure+ Option on remote Connect:Direct servers?
- What changed?
- Who changed it?
- Can I easily add a new server to multiple Connect:Direct server netmaps?
- Do I need notifications when my Connect:Direct servers' certificates used for Secure+ are going to expire?

Chapter 5

Defining the Work

After you have identified the business objectives you have for Sterling Control Center, you need to begin the process of defining how Sterling Control Center will perform the work necessary to meet those objectives. You will then configure the Sterling Control Center building blocks, or objects, that will help you accomplish that work. You define the following types of items when configuring Sterling Control Center:

- Servers in your environment and how they might logically be grouped
- Roles that specify actions users can perform and data they can view and work with on the servers and server groups defined in Sterling Control Center
- E-mail lists you can associate with rules/actions to notify groups of individuals when an event has occurred
- Rules that specify server events (such as a process starting on a monitored server) that cause actions to be taken (such as generating an alert), and when those rules should be active
- Service level criteria (SLCs) that specify processing must occur within a specific time frame and for a specific duration of time

To assist you in configuring these items, Sterling Control Center provides task lists, worksheets, scripts, utilities, wizards, and predefined items (calendars, rules, and actions). You can create items manually, or you can create multiple objects using a program, sample script, and sample templates included in Sterling Control Center. For more information, see *Creating Multiple Objects* in the *Sterling Control Center System Administration Guide*.

Laying the Foundation

To build the foundation necessary for defining the work Sterling Control Center will do, you need to plan for and define items such as servers, server groups, data visibility groups, roles/users, calendars/schedules, and e-mail lists. After you define these items, they are available for selection when you are defining rules and/or SLCs. All of the building blocks work together to tell Sterling Control Center what to do. For example, the following figure shows the function of the building blocks for Sterling Control Center rules:

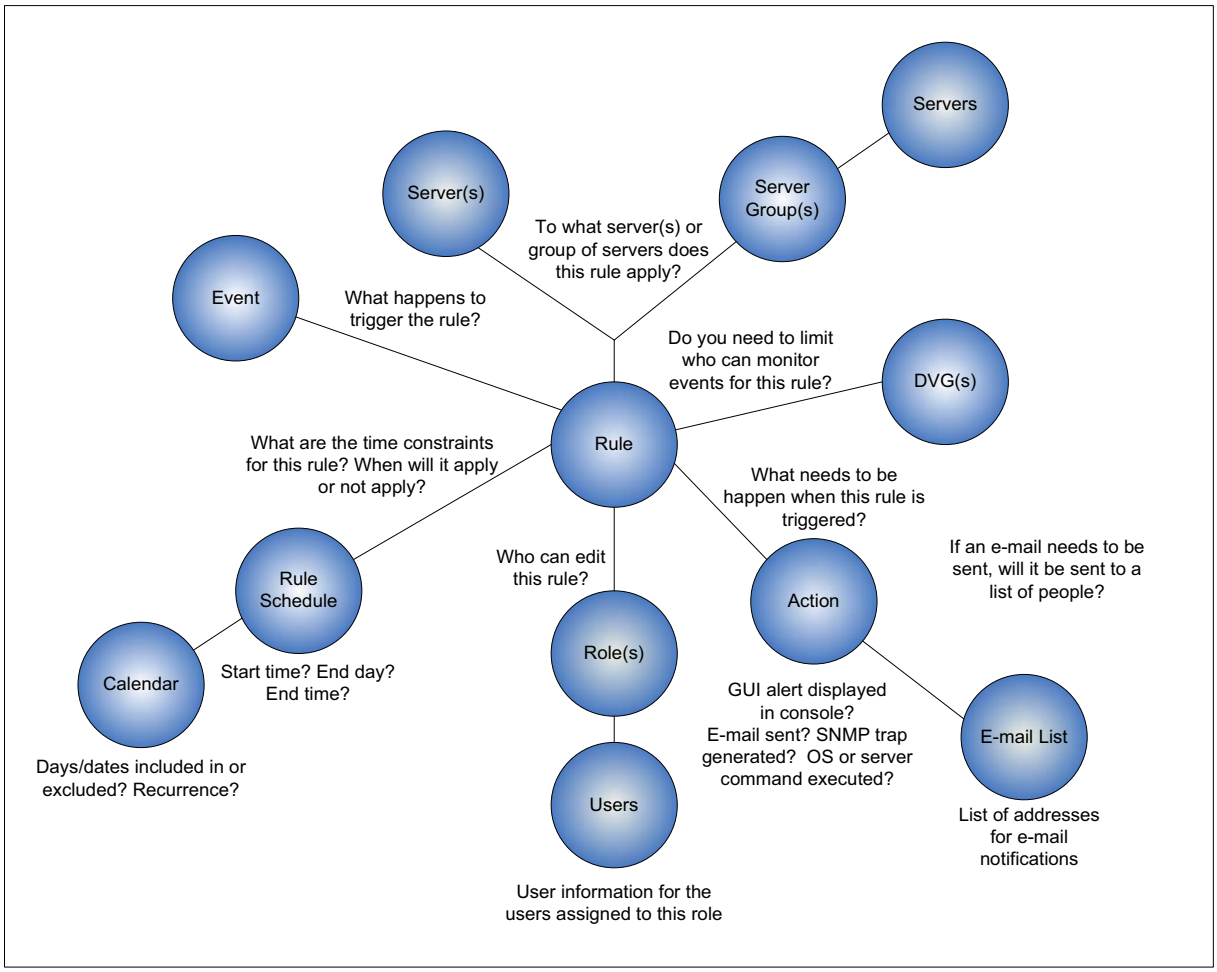

Defining Servers

When you are planning your Sterling Control Center implementation, you need to identify the number and type of servers you will be monitoring and/or managing. Whether you are planning to use Sterling Control Center to monitor servers, manage assets/license keys, or configure Connect:Direct servers, you need to supply connection information for those servers when you are setting up Sterling Control Center. Gather information regarding those servers (based on server type) that will enable Sterling Control Center to access the server resources necessary for Sterling Control Center functions. (You can use the Sterling Control Center Node Discovery feature to assist with this task.) To further define server properties, you can define settings such as monitor rest time, time zone settings (to accommodate servers in different geographic settings), and so on. Use the *Server Worksheet* on page 204 for a listing of information you need for the servers in your environment.

A wizard guides you through the process of adding servers to your Sterling Control Center configuration. After you add servers, they are displayed in the console in a tree view by server type. To access a server's properties, double-click the server. You can then review and edit the properties for that server.

In the server list, icons are displayed next to each server name to give you visual indicators of server status. For a complete listing of the server icons, see the *Sterling Control Center User's Guide*.

Defining Server Groups. As part of your planning process, decide if and how you will group servers. A server group is a user-defined grouping of servers. For example, you can group servers by processing center and/or by server type. You can even group all managed servers into one group to monitor all server activity in one monitor window. You can also put servers in multiple groups and groups within groups. Group servers in a way that makes sense for your environment based on criteria such as user access to them, the rules and/or SLCs that will be applied to those servers, and so on. Use the Server Groups Worksheet to capture information about server groups you want to set up for your environment.

When you are creating rules and SLCs that pertain to server events, you specify the server or server group(s) to which those rules and SLCs apply. You also have the option of specifying a data visibility group for rules and SLCs. If you go through the process of analyzing your monitoring needs *before* you configure the work you want Sterling Control Center to perform, you can define rules and/or SLCs that apply to a group of servers, rather than defining separate rules and/or SLCs for individual servers. This process will help cut down on the number of objects you have to define and maintain.

For example, to add a new server to Control Center and apply to it the same rules and SLCs you have set up for an existing server group, you can simply add it to the server group. If you had not implemented server groups, you would have to set up a new set of rules and SLCs for that one server.

A wizard guides you through the process of configuring server groups. The server groups you configure are displayed in a tree view on the console's Groups tab. To access a server group's properties, double-click the server group. To review and edit that group's properties, you can double-click a server in the group to access and modify its properties. Another way to display server groups is through the Server Group view. In this view, you can access a server group's properties by right-clicking the group and selecting the Properties option.

Performing Guided Node Discovery. To get a list of the Connect:Direct nodes a monitored server communicates with, you can manually track down those nodes, or do it automatically with the guided node discovery feature. To perform guided node discovery, you add servers to an explorer list, enable them for discovery, and run discovery. Control Center contacts each server to gather information about it. Sterling Control Center then scans its network map and statistics records to discover the unique servers that it communicates with. Those servers are displayed in a discovery list.

This feature is especially helpful when you are identifying all of the Connect:Direct servers you will monitor and manage in a large-scale server environment. If your Sterling Control Center license permits it, you can also add the discovered servers to the list of managed servers. (Your Sterling Control Center license limits the number of managed servers allowed.) For more information on this feature, see *Perform Guided Node Discovery* in the *Sterling Control Center System Administration Guide*.

Defining User Access

As part of the planning process you need to make decisions regarding user access to Sterling Control Center and the data it collects:

- The types of users (roles) who will access Sterling Control Center
- The permissions those roles will have when configuring and managing the servers Sterling Control Center will monitor (manage, view only, or none)
- Whether you will limit what events (data) specific users can monitor
- Whether you will implement a password policy that governs password creation

As you are planning your Sterling Control Center implementation, define the user roles you need in your environment and then configure those roles. When you add users, you can assign them the roles you have defined. When you create objects such as calendars, rules, and SLCs, you can specify the roles permitted to modify those objects.

Complete the *User Access Worksheet* when planning your Sterling Control Center implementation.

Defining Roles. When you create users in Sterling Control Center, you give them credentials, or permissions, to access the system using roles. You create roles and subordinate roles to give structure and hierarchy to permissions to meet the needs of your organization. Roles are sets of permissions that specify the data users may see and the Sterling Control Center actions users can perform and the servers and server groups they can perform these actions on. When roles are set up, data visibility groups can be used to segregate the data collected from servers or server groups that users are allowed to view. Data visibility groups facilitate the segmentation of data beyond server-level restrictions accomplished by assigning servers and server groups to roles.

Sterling Control Center is distributed with two roles: superuser and user. By default, there are no data restrictions, and all Sterling Control Center “manage” permissions are granted to the superuser role, which means users assigned this role can view all data collected and perform all Sterling Control Center functions on all managed servers. The superuser can create additional user roles or modify existing ones to serve business requirements. When additional roles are based on the default superuser role, data restrictions may be added via data visibility groups and “manage” permissions changed to “view” or “none” as needed to limit the permissions granted to them.

By default, the user role has no data restrictions and “view only” permission for all of Sterling Control Center’s functions and cannot perform management functions such as adding servers or creating SLCs or rules. If additional roles are based on the user role, data restrictions may be added via data visibility groups and permissions can be added to define those roles and expand the permissions granted those roles.

Some roles might require a mixture of permissions: they may need view-only permission regarding rules for a server or group of servers, but they may need manage permission for processes for those same servers. This would mean they could not add or edit rules for those servers, but they could delete, resume, or suspend processes. This is why it is important to consider your business requirements and plan a role hierarchy based on those requirements. Your hierarchy might be based on geographic region, server type, service line, business unit, etc.

For example, the superuser role would typically be assigned to the person ultimately responsible for Sterling Control Center. This superuser would have full permissions for setting up/configuring monitoring and full rights to all managed servers. There might be another administrator who is

responsible for configuring Connect:Direct, so a role could be defined that grants that user the permissions necessary to configure Connect:Direct servers from Sterling Control Center. For security reasons, another administrator might be responsible for configuring Connect:Direct Secure+ Option. Perhaps the business unit wants to do self-service monitoring. In this case, user roles could be defined that have view-only permissions for a restricted number of servers.

The following illustration shows a sample role hierarchy:

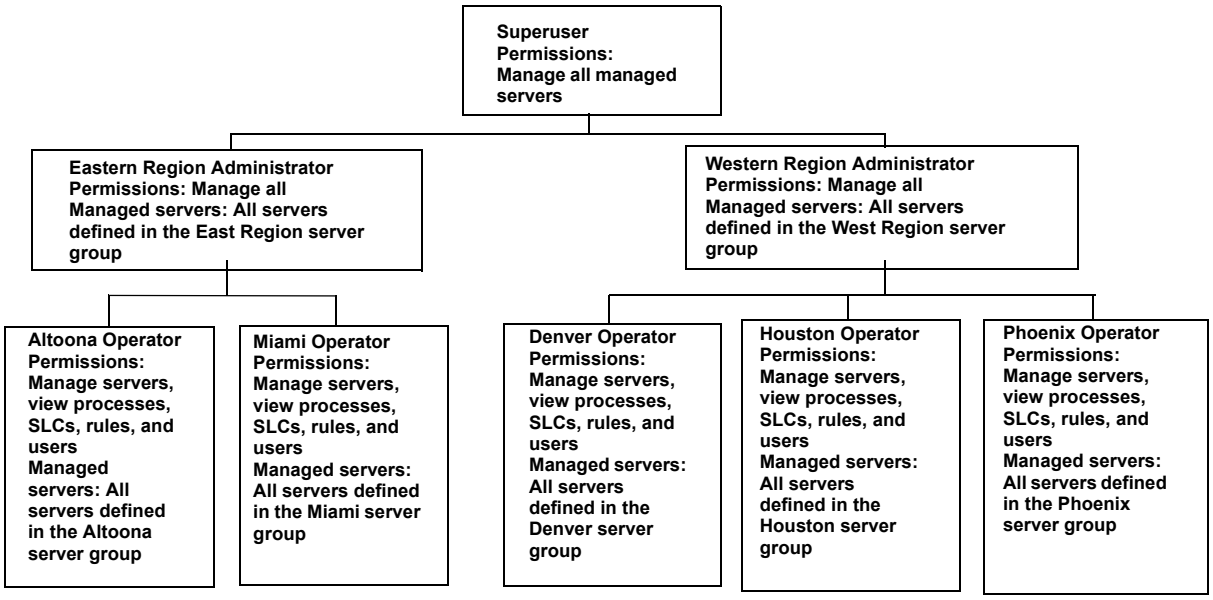

Subordinate roles cannot be given permissions higher than those of a superior role. Also, subordinate roles can only be given access to the servers or server groups a superior role can access.

For example, if an Eastern region administrator role has manage permissions on server groups A, B, and C, any of its subordinate roles can only manage or view server groups A, B, and C (or a subset). Likewise, if a Western region administrator role has only view permissions for a server group, it cannot assign manage permissions for that group to any subordinate roles.

Understanding Permissions. Users are created and given credentials, or permissions, so they can access the Sterling Control Center system. Roles and subordinate roles are defined to build a hierarchy of permissions. Roles can also be associated with server groups to further segment permissions. Data visibility groups can be assigned to a role to restrict the data (events) a user can access. For more information on data visibility groups, see *Understanding Data Visibility Groups* on page 33.

Permissions define the actions that Sterling Control Center users can perform. There are three permission levels: Manage, View Only, and None. If a role does not have permission to access a certain function (permission level for the function is set to None), that function appears dimmed on the affected user's console and cannot be selected. After you define a role with restricted permissions, you can use that role to control who can create and manage the building blocks and see, and not see, data both collected and generated by Sterling Control Center. For more information about the permissions you can grant to user roles, see *Managing Roles and Users* in the *Sterling Control Center System Administration Guide.*

Implementing Password Policies. If you require a password to authenticate users, you can configure Sterling Control Center to accept only passwords that conform to your company's

password policy. Password policies are set in Sterling Control Center by modifying the passwordPolicy.xml file located in the <*installation directory*>\ControlCenter\conf\security folder. If you implement a password policy, the policy criteria will be enforced only for existing users who change their password or new users added *after* the passwordPolicy.xml file has been modified and put into effect. If you do not edit the passwordPolicy.xml file, no password policy will be in effect for Sterling Control Center.

Password policy settings include the following:

- Minimum and maximum password length
- Requiring lowercase, uppercase, and special (non-alphanumeric) characters in the password
- Excluding lowercase, uppercase, and special characters in the password
- Using regular expressions (regex) to define specific password patterns
- Using regular expressions (regex) to define specific patterns to exclude

Choosing the Best Building Blocks for the Job—Rules vs. SLCs

A vital part of implementing Sterling Control Center in your environment is defining the building blocks that will provide the instructions Sterling Control Center follows when monitoring servers. Key to this process is an understanding of the building blocks, their differences, and how they work together.

Rules and SLCs

Two building blocks that provide structure for Sterling Control Center monitoring are rules and service level criteria (SLCs). It is important to understand the difference between the two because they serve different purposes. One of the basic questions becomes: “Do I need a rule or an SLC to accomplish my objective?”

- Rules—Can stand on their own for the purpose of processing events and taking actions when those events occur. For example, when a Sterling Integrator HTTP adapter is down, a Sterling Control Center event is generated, which triggers a rule that invokes an action that generates an alert and sends an e-mail notification to an individual or list of individuals.

  The basic process that occurs for rules is the following:

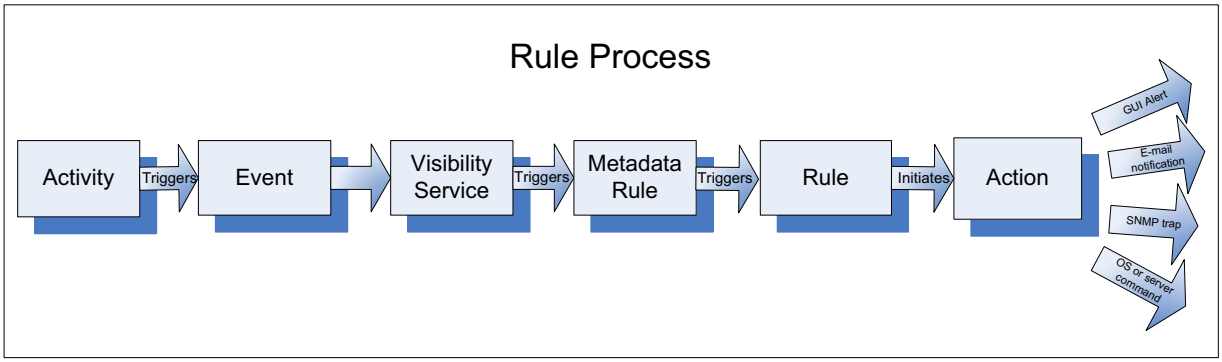

- SLCs—Are used to monitor for activities that either do not and do occur, whereas rules can only take action on events that do occur. They monitor processes based on performance objectives: the process, or process step/activity, executes or failed to execute within a certain time frame or for a specified duration. Processes can include Connect:Direct Processes, Sterling Integrator business process activities, and the movement of data into and out of Connect:Enterprise mailboxes. When the conditions in an SLC occur or fail to occur, an SLC-type event is generated. You can set up rules based on SLC event message IDs that invoke actions, such as generating an alert. SLCs can not generate alerts without rules.

  For example, Sterling Control Center monitors FTP Puts from a bank to ensure delivery to a partner within a certain time frame (SLC performance objective). When all outbound files have been sent, an outbound rule is triggered when transfers are completed successfully within the specified time frame, and an e-mail notification is sent.

  When you are dealing with SLCs, the process differs as shown in the following:

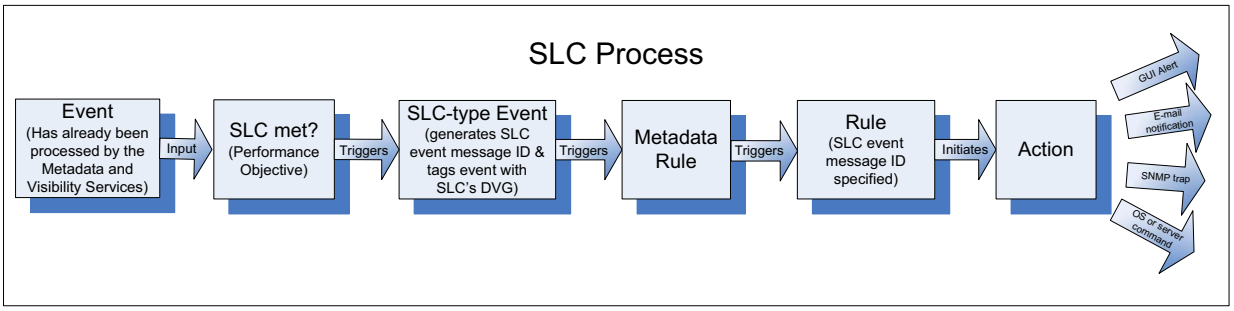

The following table shows typical situations where you would use either an SLC or a rule:

Situation	SLC	Rule
Server up/down		X
Adapter up/down		X
Step failed		X
Process didn't run	X	
Process ran on time	X	
File transfer took too long	X	
Process was not followed by another	X	

Understanding Data Visibility Groups

Whereas server groups can be used to limit what servers a user has access to, data visibility groups (DVGs) limit what data (events) a user has access to. For example, a DVG can be used to restrict a user to Accounting data on Server A and Server B. DVGs are an optional building block that can be used to limit user access to the following:

- Data shown on the following monitors:
  - Completed Process Activity Monitor
  - Queued Process Activity Monitor
  - Active Alerts Monitor
  - Handled Alerts Monitor
- Alert counts shown on the Server Status Monitor
- Rule configuration
- SLC configuration

To set up data visibility groups, you can specify criteria for segmenting data as needed for your organization. For example, you could segment data into different lines of business (LOBs) or different functional areas, such as accounting or payroll. When events match on any criteria for a DVG, that data visibility name is put into the DVG attribute of the event. Therefore, the event is "tagged" with that DVG.

Restrictions and Permissions

After you define data visibility groups (DVGs), you assign them to roles, making the roles DVG restricted. The roles are then assigned to users, making the users DVG-restricted. A role can have a server group restriction or DVG restriction or both. If a role is server or DVG restricted, it cannot be given "manage" permission for DVGs. As a result, restricted role can have only "none" or "view" permission for data visibility groups.

Note: Only Control Center administrators can manage data visibility groups. To qualify as an administrator, a role must not be server group or data visibility group restricted and must have "manage" authority to required elements. If a role qualifies as an administrator, the "manage" permission is allowed; otherwise, only "view" or "none" permissions are allowed.

A DVG-restricted user sees only a subset of all data in the Sterling Control Center monitors and a subset of all rules and SLCs. Users with roles restricted to one or more DVGs will only be able to view activity entities such as processes, statistics, alerts, and reports, that match their DVG restriction.

Rule Sets

Where previously there was only one rule set, beginning with Sterling Control Center version 5.2, there are multiple rule sets: a global rule set and one rule set for each data visibility group (DVG) defined. An event can trigger at most one rule per rule set, and each rule set has its own priorities. The following table describes the differences between the types of rule sets:

Type	Description
Global Rule Set	◆ Rules that do not specify a data visibility group (DVG) in parameters ◆ Global rules belong to the global rule set ◆ The built-in rules that are shipped with Sterling Control Center are in the global rule set
Rule Set	◆ Rules that specify a DVG ◆ Called data visibility group restricted rules ◆ If a data visibility group-restricted rule specifies ”Data Visibility Group Matches Payroll” in its parameters, the rule is in the Payroll rule set.

The Rules window has a Rule Set list that displays the rule sets the user has access to, as shown in the following:

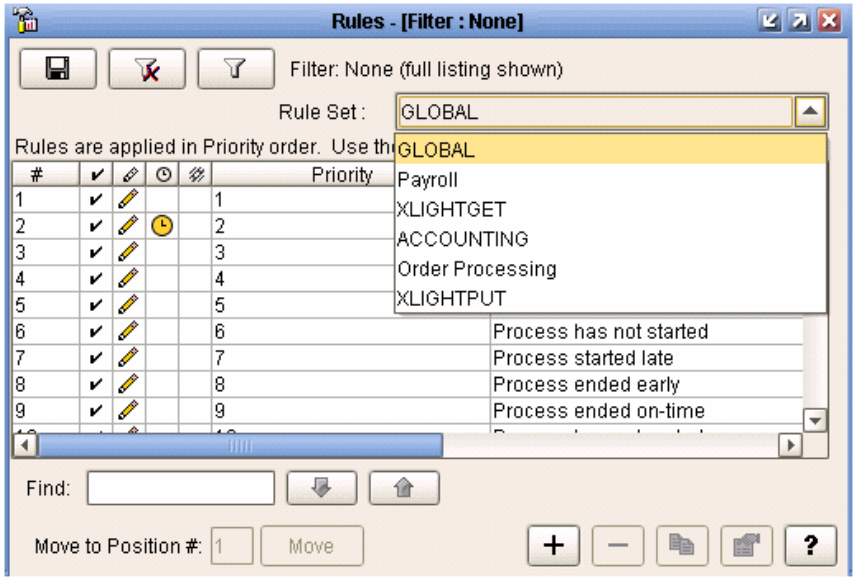

When an administrator views the Rules window, the Rule Set list displays GLOBAL as well as all the data visibility groups defined. To display a rule set, the administrator selects one from the list. The rules are filtered, and only the rules in that rule set are displayed.

When a DVG-restricted user views the Rules Listing panel, only the user’s assigned DVGs are displayed in the list. A DVG-restricted user cannot view the GLOBAL rule set.

Rules

When a DVG-restricted user creates a rule, the rule must specify a data visibility group in the rule’s parameters. The Rule wizard will pre-populate the user’s DVG in the rule’s parameters. If a user is assigned to only one DVG, the value cannot be changed or removed. If a user has authority to multiple DVGs, the value can be changed.

A non-DVG restricted user can create the following:

- A global rule
- A data visibility rule using any defined DVG

Note: Metadata rules are not restricted by data visibility groups.

SLCs

Data visibility groups (DVGs) can be added to standard, wildcard, simple, and workflow SLCs. An SLC without a DVG is a global SLC. An SLC with a DVG set is a data visibility group SLC.

When creating an SLC, DVG-restricted users can create only data visibility group SLCs. The data visibility group list will be pre-set with a data visibility group that users have authority to. If users have authority to multiple DVGs, they can select a different DVG (from the pre-set one), but a DVG must be selected.

On the SLC listing windows (standard, wildcard, simple, and workflow, the SLC Set list displays the SLCs the user has authority to: global and/or data visibility group SLCs. When administrators view the SLC listing window, the SLC Set list displays the GLOBAL SLC set and all the data visibility groups SLCs defined. When DVG-restricted users view the SLC list window, they see only the data visibility group SLCs they have authority to. DVG-restricted users cannot view the GLOBAL SLC set.

Events

Every event generated in the system (except for SLC events generated by SLCs with DVGs assigned) is compared to all criteria in all data visibility groups. If a match is found, the name of the data visibility group (DVG) is added to the event as shown in the following:

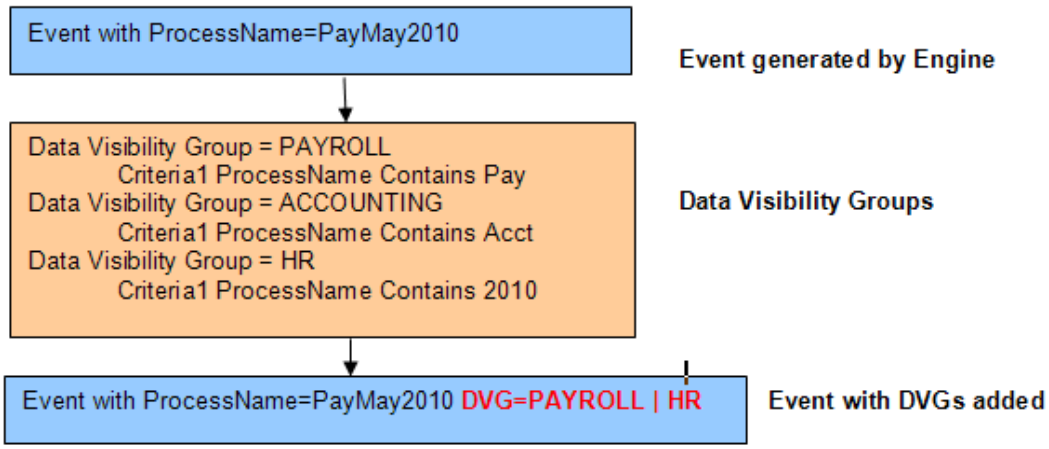

Note: Each event can be tagged with multiple data visibility groups (DVGs). Multiple DVG values are separated by vertical bars.

Each event is then passed through the metadata service rules and then the appropriate data visibility group rule sets and the GLOBAL rule set. If an event tagged with a data visibility group matches a corresponding rule, that rule is triggered as shown in the following:

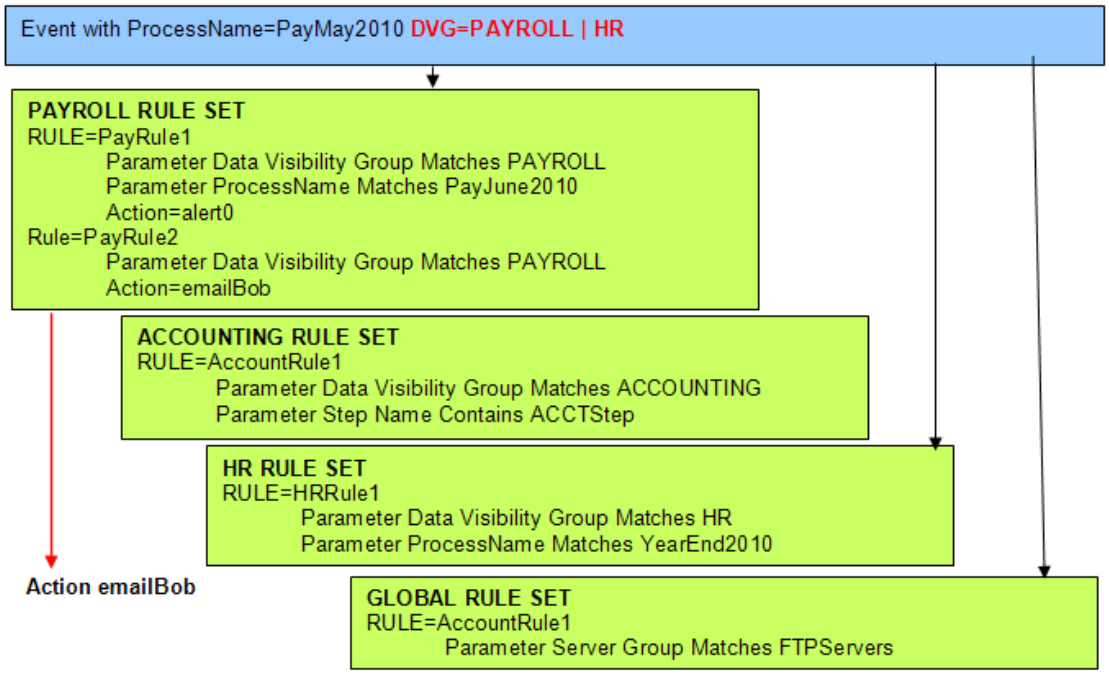

DVGs and Sterling Control Center Information

Data visibility groups (DVGs) affect the information displayed for a user if that user is DVG restricted.

Monitors

The events displayed on the Completed Process Activity, Queued Process Activity, Active Alerts, and Handled Alerts Monitors may be filtered by data visibility groups (DVGs) and their DVG values can be displayed using the optional Data Visibility Group column in each of the monitors. A non-DVG restricted user can view all events (with sufficient “view” authority). DVG-restricted users can view only events tagged with a DVG for which they have access (and sufficient “view” authority).

The following example illustrates what an administrator with no data visibility group restrictions would see in the Active Alerts Monitor:

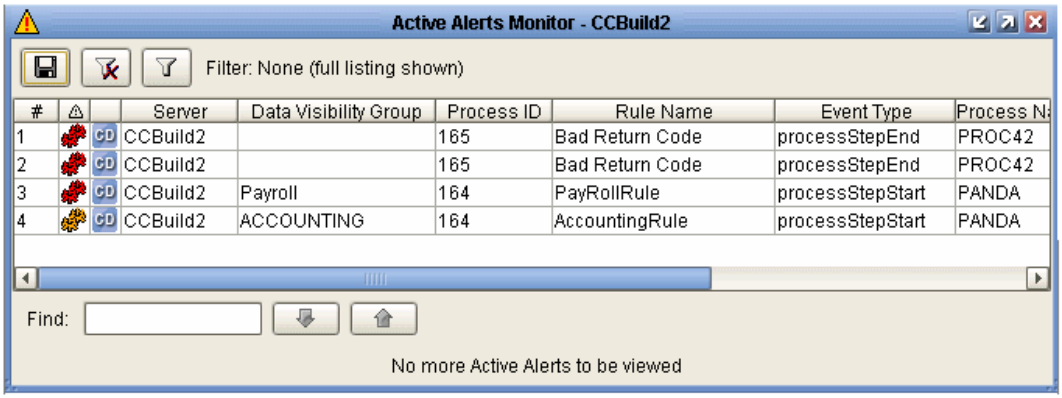

Notice Process ID 164 is one Process Step Start event that triggered both a Payroll rule and an ACCOUNTING rule and generated two alerts: an alert1 and an alert2.

In the following example, the same Active Alerts Monitor is viewed by a user restricted to the Payroll DVG. Note that the user sees only the event tagged with Payroll DVG.

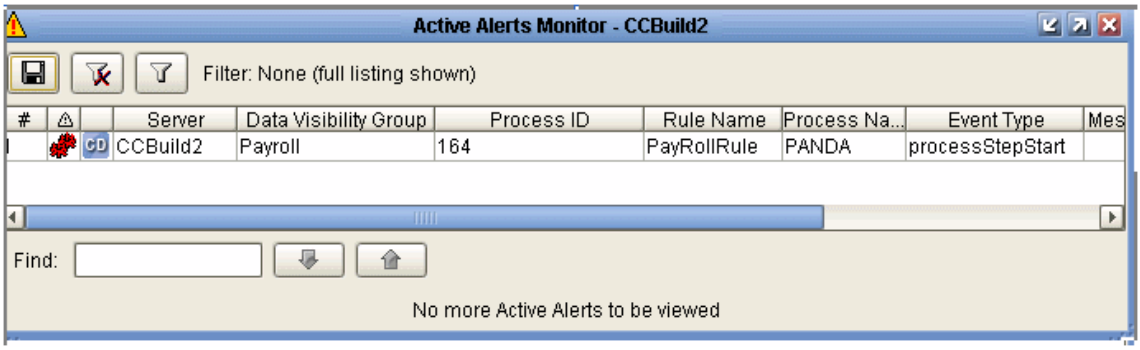

The following table describes the information displayed in the monitors based on where a user is DVG restricted:

Monitor	Non-DVG Restricted User	DVG Restricted User
Completed Process Activity	◆ Sees all processes regardless of any DVG tags on the events that comprise the process.	◆ Sees a process if any event that makes up the process is tagged with a DVG that the user has authority to. ◆ Will only be able to select statistics for the event types tagged with the DVGs the user has authority to.
Queued Process Activity	◆ Sees all processes regardless of any DVG tags on the events that comprise the process.	◆ Sees a process only if a DVG that the user is restricted to is on one of the queued process steps. **Note:** The process will be listed on the monitor when the step that the user has visibility to starts. The process remains on the monitor until the process ends.
Server Status Monitor	◆ Sees counts of all alerts	◆ Sees only the counts for the alerts that match the user's DVG restriction.

Reports

When a DVG-restricted user runs any report, the report data/results are automatically filtered by the DVG the user has authority to. For example, if a DVG-restricted user runs an Alerts Report, only the alerts tagged with the user's DVGs will show up on the report. Data Visibility Group has been added to the report filter so that DVG-restricted users who have been assigned multiple DVGs can limit the report to a certain DVG or so that non-DVG restricted users can restrict the report to a specific DVG.

Understanding Rules and Actions

Rules specify criteria that must match an event generated in instances such as the following:

- Data collected from a monitored server
- A license key file and/or server certificate is within a specified license warning configured in a parameter setting
- Polling of a node has not occurred within a reasonable amount of time
- A rule is created, updated, or deleted

When rules are triggered by events, the action they refer to is performed, such as the following:

- Generate an alert and an e-mail notification to a system administrator if a process or file transfer completes with errors
- Monitor a process or file transfer for specific message IDs, and issue an operating system command if the message is detected
- Monitor server status and generate an alert if a server error occurs
- Generate an SNMP trap when a process's return code is 8 or higher (for certain server types)

Rules have the following properties:

Property	Description
Criteria	Conditions that must be met for a rule to be applied, such as: ◆ Parameters (for example, Event type, Message ID, Server ID, and SLC Name) ◆ Server/server groups to which the rule applies ◆ Data visibility group to which the rule belongs
Actions	Action performed when all criteria are met. Actions include: ◆ Generating an alert (with different severity levels) ◆ Sending an e-mail notification ◆ Generating an SNMP trap ◆ Executing an operating system command on the system running the Control Center engine or a server command on the specified monitored server
Schedule	One or more schedules (calendar) can be associated with a rule. If a schedule is used, the rule is applied when all rule criteria have been met *and* a schedule associated with the rule matches. For more information on calendars and schedules, see *Calendars and Schedules* on page 46.

Property	Description
Linked rule	A rule with a second set of criteria that must occur within a specified time. Linked rules also include both a resolution and non-resolution action, one of which is taken depending on whether or not the second set of criteria is met within the time specified. For example, a linked rule could be used to generate an alert (non-resolution action) for a server down condition (first set of criteria) *only* if a server up event does not occur within five minutes (second set of criteria), thus giving an administrator a five minute window to restart the server before any alert is generated.

After you create a rule, it is displayed in the Rules listing in the console for the rule set it was assigned. Rules without a data visibility group (DVG) criteria are assigned to the global rule set. Rules with a DVG criteria assigned belong to the specified DVG rule set. All enabled global rules for the entire Sterling Control Center system are applied in the order in which they are listed in the Rules listing. Events with a DVG attribute are subsequently processed by each applicable DVG rule set. The basic process is as follows: An event occurs > For the global rule set, and each applicable DVG rule set, is it a match to the first rule in the listing? No. Go to the next rule. Is it a match? No. Go to the next rule, and so on, until a match occurs. Therefore, rules with specific criteria should precede rules with more general criteria (specific server vs. server group). Only one rule per rule set is triggered per event, so if the first rule is too general, a match always occurs and subsequent rules are ignored.

Understanding SLCs

You can set up service level criteria (SLCs) that help you monitor a process or file transfer based on performance objectives (processing must occur within a certain time window). For example, a Connect:Direct process may need to begin by 20:00 and end by 20:30. An SLC might monitor for the timeliness of both events (Did it begin on time? Did it end on time?). If either does not occur within its respective window, the SLC can be used to notify you of that fact.

When performance objectives are met or are not met (certain conditions occur or fail to occur because processes either execute or fail to execute as expected), Sterling Control Center generates SLC event messages. You can use SLC event message IDs in rules to trigger an action, such as one that generates an alert to display in the Alerts Monitor. As with Rules, there is a global SLC set, and there can be multiple DVG SLC sets, one per DVG defined.

SLCs can be defined to monitor for one or more of the following:

- Connect:Direct Process starts, ends, and durations
- Connect:Direct Process step starts, step ends, and durations
- Sterling Integrator business process starts, ends, and durations
- Sterling Integrator business process activity starts, ends, and durations
- Sterling File Gateway arrived file, route, and delivery starts, ends, and durations
- Connect:Enterprise batch arrivals and transmissions and durations

- File Transfer Protocol (FTP) get and put ends

Parts of an SLC

SLCs contain elements such as the following:

Element	Description
Type	Tells what kind of SLC it is. There are four types of SLCs: ◆ Standard—Standard SLCs monitor specific process names, file names, etc. Use standard SLCs when you know the specific item to monitor. ◆ Wildcard—Wildcard SLCs facilitate the specification of match criteria that can match multiple values, as opposed to one value. When dealing with batch IDs, you could either specify multiple standard SLCs (one per batch ID that may be created), or you could specify one wildcard SLC that uses a wildcard character to match the date and time portion of the batch ID, which varies from batch to batch. To specify monitoring criteria in wildcard groups, you can use the wildcard characters asterisk and question mark, or regular expressions (regex). Sterling Control Center provides an Expression Tester to test wildcard and regular expressions. ◆ Workflow—Workflow SLCs monitor the flow of related processes or process steps by tracking them as milestones in a workflow. For example, a workflow SLC can monitor a transaction consisting of three processes, all of which must finish within three hours of the first process's initiation. You can use a correlator to associate milestones with an SLC based on a value obtained at runtime. For example, if you have two different users (user ID 1 and user ID 2), who initiate two processes (P1 and P2) and those processes should not exceed a duration of 15 minutes, you could correlate them by specifying a correlator value of submitterid. This enables the engine to ensure that both instances of the workflow were submitted by the same user and that each workflow met, or did not meet, the duration specified. This prevents Sterling Control Center from associating user ID 1's second process with user ID 2's first process: 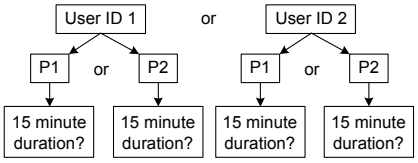

Element	Description
	◆ Simple—Simple SLCs enable you to create an SLC by answering a few basic questions, specifying values for basic parameters, and giving the SLC a name and description. When you create a simple SLC, all necessary objects to support the SLC, such as rules, actions, and schedules, are also created.
Schedule	One or more schedules (calendar or duration) can be associated with an SLC that provide the time constraints associated with performance objectives. An SLC with more than one calendar is an SLC group. For more information on calendars and schedules, see *Calendars and Schedules* on page 46.
Start/End Window Tolerance	The start and end window tolerances determine the size of the monitoring window for each SLC. They can be used to set up a wider monitoring schedule wider to detect a start that is earlier or an end that is later than expected. ◆ Start Window Tolerance—The number of hours before the expected start of processing to being monitoring. ◆ End Window Tolerance—The number of hours after the expected end of processing to stop monitoring.
Criteria	Information about processes/file transfers that Sterling Control Center is looking for when monitoring performance objectives of monitored servers: ◆ Process names/batch IDs ◆ Destination file names ◆ Submitter /sender mailbox IDs ◆ Remote servers/recipient mailbox IDs ◆ Wildcard expressions (wildcard and workflow SLCs only) ◆ Correlator (workflow SLCs only) ◆ Fire once (workflow SLCs only) ◆ Jeopardy message lists (workflow SLCs only) ◆ Servers or server groups ◆ Data visibility groups (DVGs)

SLCs by themselves cannot invoke actions. For the SLCs you create, check to see if the built-in rules meet your needs. If they do not, create rules that use SLC event message IDs to trigger the desired action. For more information on rules and SLCs, see *Choosing the Best Building Blocks for the Job—Rules vs. SLCs* on page 32.

Predefined Actions and Rules for SLCs

Sterling Control Center provides predefined actions that generate alerts and predefined rules that monitor for SLC messages. You can use these actions and rules when creating SLCs to monitor processing requirements. You can also modify these actions and rules as necessary to meet your processing requirements.Or, to take specific actions, you can create new rules with more specific match criteria.

Note: Built-in rules for SLC events only exist for the Global SLCs. You must create DVG rules to watch for events generated as a result of non-global/DVG SLCs for actions to be taken.

Understanding Metadata Rules

Metadata rules allow you to append additional elements and values to Sterling Control Center events before they are processed by both the Rule and the SLC Services. Metadata rules are applied to all Sterling Control Center events (if they occur during the schedule associated with the metadata rule), unless you explicitly set them not to be applied for statistics collected from specific managed servers.

The additional metadata type elements and values are logged in the Sterling Control Center Events database. When a metadata rule matches an event, Sterling Control Center appends metadata to the event as lists of key value pairs. You can use these metadata fields as matching criteria when defining conventional rules and SLCs. Metadata can also be used as filter criteria for reports and alert monitor or activity monitor data. So not only can metadata rules be used to simplify the specification of your rule and SLC criteria, they can also be used to simplify specification of your report criteria.

Note: Metadata can be used to analyze only new activity going forward. You cannot do retroactive analysis of existing data to which metadata tags have not already been applied.

There are four metadata fields (USER_DATA_1 – USER_DATA_4). Control Center also provides 10 server metadata fields (SERVER_DATA_1 – SERVER_DATA_10) whose values are set when a monitored server is defined. Every event generated by the server will contain the specified values for the 10 fields. The values can be evaluated by metadata rules and regular rules. You can also name the metadata fields to something more meaningful.

Metadata Example

To illustrate the use of metadata to simplify report generation, let's assume the following environment and set of requirements about a company named Ultimate Sporting Goods (USG):

Objectives

- ✦ USG wants to notify via e-mail a person in Accounting whenever a process owned by accounting fails.
- ✦ USG also wants to generate a Connect:Direct process Statistics Summary report about accounting processes.

Requirements

- There are five servers being monitored (SERVER1, SERVER2, SERVER 3, SERVER 4, and SERVER5).
- SERVER1 and SERVER5 are Connect:Direct servers that only do accounting work. So, all processes run on these two servers are accounting processes.
- SERVER2 does accounting and other departments' work. Processes submitted by Jane or Mary on SERVER2 are accounting processes.
- SERVER3 does accounting and other departments' work. Processes whose name begins with ACCT are accounting processes.
- SERVER4 does accounting and other departments' work. Processes whose remote node is SERVER100 (a non-monitored node) are accounting processes.

There are two ways to handle these objectives: 1) using regular rules or 2) using metadata rules. In most cases, using metadata rules is more efficient.

Regular Rules. To meet the requirements without using metadata rules, USG would write four regular rules that would take an action on accounting processes that fail. The four regular rules would be:

- If Server Id matches "SERVER1|SERVER5" then take Accounting E-mail Action
- If Server Id matches "SERVER2" and Submitter matches "Jane|Mary" then take Accounting E-mail Action
- If Server Id matches "SERVER3" and Process Name wildcard "ACCT*" then take Accounting E-mail Action
- If Server Id matches "SERVER4" and Remote Node matches "SERVER100" then take Accounting E-mail Action

Because Sterling Control Center does not have complex AND/OR logic for reports, a single report cannot be generated. So, USG has to write four separate reports to get the accounting reports, when they only wanted one. And, USG has to put logic similar to that for the rules into the report selection criteria.

The four reports would be:

- SERVER1 and SERVER5 Accounting Report criterion:

  Server Id matches "SERVER1|SERVER5"
- SERVER2 Accounting Report criteria:

  Server Id matches "SERVER2" and Submitter matches "Jane|Mary"
- SERVER3 Accounting Report criteria:

  Server Id matches "SERVER3" and Process Name wildcard "ACCT*"
- SERVER4 Accounting Report criteria:

  Server Id matches "SERVER4" and Remote Node matches "SERVER100"

If some other criteria are introduced (for example, SERVER6 is now being monitored and does some accounting work), then a new rule and a new report has to be defined.

Metadata Rules. To meet the requirements using metadata rules, USG would write four metadata rules, all using the Accounting Metadata Action. The Accounting Metadata Action sets USER_DATA_1 to "ACCT." Therefore, every time a metadata rule matches, the event gets

appended with the accounting tag “ACCT”. That accounting tag is then used to match the regular rule. The four metadata rules would be:

- If Server Id matches “SERVER1|SERVER5” then take Accounting Metadata Action
- If Server Id matches “SERVER2” and Submitter matches “Jane|Mary” then take Accounting Metadata Action
- If Server Id matches “SERVER3” and Process Name wildcard “ACCT*” then take Accounting Metadata Action
- If Server Id matches “SERVER4” and Remote Node matches “SERVER100” then take Accounting Metadata Action

USG would write one regular rule whose match criterion is USER_DATA_1 matches “ACCT” and whose action is to e-mail a person in Accounting.

For the report, USG would define a single report whose match criterion is USER_DATA_1 matches “ACCT” and would get the desired consolidated report.

If some other criteria are introduced (for example, SERVER6 is now being monitored and does some accounting work), then only a new Metadata Rule would have to be defined.

Sterling Control Center allows you to name the metadata fields to something that is more meaningful using Metadata Type Mapping. In the example above, USG could map USER_DATA_1 to “Department.”

Permissible Objects

A role can have a server group restriction or data visibility group restriction or both. When restricted roles are assigned to Sterling Control Center building blocks, users with IDs assigned to those roles can manage the object. You can also elect to make the object visible either to all users, which makes it public, or only to restricted users in the selected roles, which makes it private. If users can view a permissible object, they can also use that object.

For example, if a DVG-restricted user’s role is associated with an action, the user can use that action when building a rule, regardless of whether that action (object) is private or public. If a DVG-restricted user’s role is *not* associated with an action, the user can only use that action when building a rule if the action (object) is public. The following building blocks are permissible objects:

- Actions: rule and metadata
- Schedules: rule, SLC, metadata, and report
- Lists: email and message
- Calendars

On an object’s Permissions window, if you select restricted roles for the object, and you select:

- This <object type> is visible to all users—The object is public. A public, referenced (used by another object) permissible object can have roles in “Selected Restricted Roles” removed because removing roles does not reduce who can view/use the object, only who can manage it.

- This <object type> is visible to restricted users in these Selected Restricted Roles—The permissible object is private and can only be viewed/edited/used only by users in the restricted role(s) selected and by unrestricted users. When a permissible object is private and has been referenced (used by another object), none of its roles can be removed because it cannot be made more restrictive. However, a private, referenced object can be made less restrictive.

Calendars and Schedules

When you are planning your Sterling Control Center implementation, you need to determine the day/date and time constraints that will be placed on the work that Sterling Control Center will perform. Days/dates are defined in calendars and times are defined in schedules. Calendars and schedules are defined independently for efficiency because you could have the same calendar (for example Mon.–Fri.) paired with different schedules (08:00–09:00 and 16:00–17:00). After you have set up calendars and schedules, you can select these building blocks when you are developing rules and SLCs.

Calendars

Calendars specify days/dates, how long the calendar remains in effect, and how often processing is repeated (recurrence). Sterling Control Center comes with eight predefined calendars: one for each day of the week (Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday), as well as one that includes every day of the week (Daily). You can use these predefined calendars for scheduled reports and for rule schedules and SLC schedules. You can also create additional calendars to meet your processing needs.

When you create calendars, you specify a recurrence of daily, weekly, monthly, or yearly. The planning process for your Sterling Control Center implementation should include an analysis of the calendars you will need for monitoring the servers in your environment. If you need common calendars that will be used repeatedly, you can set them up prior to defining the schedules that use them. You can also create calendars as you are defining scheduled reports, rules, and SLCs.

For example, a calendar for Monday through Friday processing might have the following values:

Panel	Field	Value
General	Name	Monday-Friday
	Description	Monday through Friday
Recurrence	Start	Today’s date
	End	No end date
	Recurrence Pattern	Every 1 week
Modification	Modifications	Remove holidays per your company schedule

Schedules

Schedules specify times and are associated with reports, rules, and SLCs. They specify the calendar that will be used, time zone for the calendar, and how long the calendar will be in effect (start time, end time, and end date). When associated with a report, via Scheduled Reports, schedules tell the Report Service when to run a report. When a schedule and an event are associated with a rule, the schedule dictates when events will be matched against the rule's criteria. As such, schedules are used to specify the criteria by which an activity is judged to be executing acceptably. For SLCs, schedules are used to specify a duration or a range for the start time and end time, or both, in which processing must occur.

Rule Schedules. A rule schedule is based on a calendar that specifies the days/dates the rule will be in effect and times it is in effect (start and end).

For example, when you are defining a rule, you could set up a calendar schedule based on a daily calendar and specify a schedule from 22:00-23:00 (the normal maintenance window for a set of servers). Then, you could set up a rule that looks for a certain message ID generated during that time on a certain group of servers. In the rule, you would also set up an action (for example, an e-mail notification) that is taken if the event occurs (message ID generated during the days/dates and times specified by any of the servers specified).

SLC Schedules. There are two types of SLC schedules: calendar and duration. An **SLC calendar schedule** is based on a calendar that specifies the days/dates that the schedule is in effect and the normal start range (NSR), normal end range (NER), or both, in which processing must occur. NSR defines a start and end time that describes when an activity is expected to begin (NSRs and NSRe). NER defines a start and end time that describes the window during which activity is expected to end (NERs and NERe).

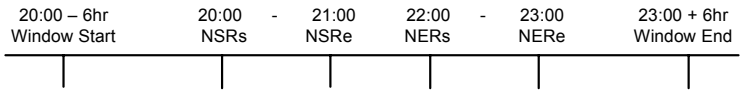

For example, you might set up an SLC schedule based on a Wednesday-only calendar and then specify that processing must start between 20:00 (NSRs) and 21:00 (NSRe) and end between 22:00 (NERs) and 23:00 (NERe).

For SLCs with calendar schedules, Sterling Control Center can monitor activity for a specified number of hours before and after the schedule requirements you define (window start and window end in the example above). The monitoring schedule can be set up wider than the schedule to detect a start that is earlier or an end that is later than expected. You set this up using the Start/End Window Tolerance.

Whereas an SLC calendar schedule specifies a range of time (with a specific start and/or end), a **duration schedule** specifies that processing can begin at any time, but once it has begun, it must be completed within a specified amount of time (hours, minutes, seconds). For example, processing must complete within 15 minutes of when a file transfer starts.

For duration schedules, Sterling Control Center can monitor for a specified number of hours after the maximum duration is reached (end window tolerance). The monitoring schedule can be set up wider than the duration schedule to detect a duration that is longer than is expected.

Note: The Concurrence count value associated with wildcard and workflow SLCs is handled differently by the SLC service depending on the type of schedule associated with the SLC. For SLCs with a calendar schedule, the concurrence count value dictates how many instances the service expects to see, and events are generated accordingly. For SLCs with a duration schedule, the concurrence count value dictates how many instances the service watches for simultaneously. Unlike SLCs with calendar schedules, “did not occur” events are not generated for SLCs with only duration schedules.

Chapter 6

Understanding Sterling Control Center Information

The result of the building blocks you define for Sterling Control Center is a wealth of information about the servers in your environment. An understanding of the types of information available in Sterling Control Center will help you to access the information you need about those servers.

Types of Information

Sterling Control Center provides you with several different types of information about the servers in your environment, such as the following:

Type of Information	Description
Status	Visual status indicators display in the consoles. You can tie rules/actions to status when you are defining work for Sterling Control Center to perform.
Server, adapter, and daemon status	The following status monitors are available: ◆ Server Status Monitor ◆ Adapter Status Monitor ◆ Daemon Status Monitor
Activity status	The **Process Activity Monitors** display completed processes (**Completed Process Activity Monitor**) and queued processes (**Queued Process Activity Monitor**) on one server or server group, multiple servers or server groups, all managed servers of one type, or all managed servers.

Type of Information	Description
Actions	When an event occurs, it triggers a rule. The rule takes an action, which may specify that an alert be displayed in the consoles, an e-mail be sent, an SNMP trap be generated, or OS or server commands initiated.
Alerts	Visual indicators (icons) displayed in the consoles with varying severity levels that you specify in actions. Alerts are displayed as active or handled. The alerts monitors (**Active Alerts Monitor** and **Handled Alerts Monitor**) provide near real-time display of alert data as it occurs. When working with alerts, you can: ◆ View the properties of an alert, and you can view the statistics associated with a process related to an alert. You can also view the SLC or rule that generated the alert. ◆ Add a comment to the alert. ◆ Move an alert from active to handled status. When you do this, you are required to add a comment about the update.
E-mails	When an event occurs, an e-mail can be sent to an individual or list of individuals defined as the action of a rule.
SNMP trap	A message generated and sent to one or more Simple Network Management Protocol (SNMP) hosts.
Logs	Information saved in log files for historical purposes that can help you troubleshoot issues.
Audit logs	Standard report of changes made to both Sterling Control Center building blocks / objects and Connect:Direct server configuration objects. Can be run as an on-demand report or displayed on screen.
Trace logs	Helpful for troubleshooting installation problems and other support-related issues.
Reports	Used to gather information about the servers in your environment.
Standard Control Center Reports	Produced from the Sterling Control Center consoles on demand or automatically using schedules. Can specify filtering criteria and can grab any field in the database but cannot manipulate the format, do calculations, perform complex queries, etc.
Sample Crystal Reports	Sample Crystal Reports are included with Sterling Control Center and are designed for use with a MySQL database, but they can be modified for use with other databases.

Monitoring Status

Sterling Control Center allows you to monitor the status of managed servers, Sterling Integrator adapters, and Connect:Enterprise master daemons in your enterprise through the status monitors:

Server Status Monitor, Adapter Status Monitor, and Daemon Status Monitor. You can open multiple monitor windows at the same time.

So, how do you find out that the status of your servers, adapters, or daemons has changed without constantly checking the status monitors? You can define rules based on status that cause a certain action to be taken when that status exists. For example, you could define a rule based on a server down condition with an action of sending an e-mail notification status.

Monitoring Server Status

The Server Status Monitor window provides a dynamic summary of managed server activity. You can view the status of the following:

- An individual server or server group
- Multiple servers or server groups
- All managed servers of one type
- All managed servers

The following types of information are displayed for servers; however, the exact information that displays depends on server type:

- Current server status, represented by a status icon
- Server's Sterling Control Center name/alias
- Number of high, medium, and low severity alerts on the server
- Version of Connect:Direct, Connect:Enterprise, FTP, or Sterling Integrator software running on a server
- License information (expiration details and concurrent number of sessions permitted or accounts defined)
- For Connect:Direct and Sterling Integrator servers, information about maximum number of concurrent sessions on the server and number of executing and non-executing processes.

Monitoring Adapter Status

The **Adapter Status Monitor** displays summary information about Sterling Integrator adapters running on Sterling Integrator servers. You can also view adapter and perimeter server properties through the Adapter Status Monitor.

The following types of information are displayed for adapters:

- Whether the adapter is turned on or off
- Whether the adapter is currently running or stopped
- User-friendly display name for the adapter
- Nodes on which the adapter is deployed
- Type of adapter
- Sterling Integrator perimeter server through which the adapter accesses the network
- State of the perimeter server

Sterling Control Center does not monitor Sterling Integrator protocol adapters when those adapters are not actively monitoring business processes or protocols. In the case of clustered Sterling Integrator servers, status is monitored for all servers in a cluster; however, to avoid duplication, only unique adapter entries are displayed.

Monitoring Daemon Status

The **Daemon Status Monitor** displays information about the master daemon status of managed Connect:Enterprise for UNIX servers, such as:

- Daemon name
- Type of daemon
- Host machine the daemon is running on
- Daemon process identifier, originator, resource, and session identifier
- Whether the daemon is up or down

Monitoring Activity

The **Process Activity Monitors** allow you to view a configurable number of completed processes (Completed Process Activity Monitor) and queued processes (Queued Process Activity Monitor) on one server or server group, multiple servers or server groups, all managed servers of one type, or all managed servers.

Note: Because Sterling Control Center does not "know" about FTP or Connect:Enterprise activity until that activity has occurred, the Queued Process Activity Monitor is disabled for those server types.

Note: In the Queued Process Activity Monitor, users who have roles with data visibility group restrictions will only have access to the process steps associated with their data visibility group. As a result, data visibility group restricted users cannot delete, suspend, or release queued processes. In the Completed Process Activity Monitor, users who have data visibility group restrictions will see only the completed processes tagged for the data visibility groups associated with their role. However, unrestricted users will see the maximum number (defaults to 200) configured by the administrator.

From the Process Activity Monitors, you can also:

- Take a snapshot of process activity for queued and completed Sterling Integrator and Connect:Direct processes that will enable you to sort the entries in the order you want to see them. You can then print or save the entries.
- View statistics related to any process listed in the Process Activity Monitors
- View process statistics related to one or more servers, server groups. or server type
- Restart Sterling Integrator business processes
- Take actions on a queued Connect:Direct process, including deleting, suspending, or releasing it

Working With Alerts

When a rule is triggered and its action is set to an alert level, the alert is displayed in the Active Alerts Monitor. To remove this alert from the Active Alerts monitor, it must be “handled” (moved with appropriate comment to the Handled Alerts monitor). For example, if a rule is in place that watches for a Server Down event and generates a Sev 1 alert for that event, when that server goes down, an alert is generated and displayed in the Active Alerts Monitor. An operator watching the Active Alerts monitor notices the alert, investigates the server, and brings the server back up. Once the serer is back up, the alert should be “handled” by specifying an appropriate comment, such as “brought the server back up.” The user ID and the date/time when the alert was “handled” are recorded in the database along with the comment, and the alert is moved to the Handled Alerts Monitor.

Options for Generating Reports

Sterling Control Center offers a variety of standard reports that can be used to gather information about the servers in your environment. In addition, you can use SQL queries or a third-party tool, such as Crystal Reports, to extract data from the Sterling Control Center databases to create reports. Audit log and trace log printouts are also available to track things such as changes made to Connect:Direct configuration objects. They provide valuable information for troubleshooting installation problems and other support-related issues.

Standard Sterling Control Center Reports

Standard Sterling Control Center reports are produced from the Sterling Control Center consoles. Reports can be produced on demand or automated using schedules. Automated reports can be sent to designated recipients via e-mail lists. When you are creating reports, you can specify criteria such as date/time range to narrow the scope of the report. The available criteria depend on the report type selected. With the built-in reports, you can display fields in the database, but you cannot manipulate the format, do calculations, perform complex queries, and so on. To accomplish that level of reporting, you must use a third-party application. The following standard report types are available in Sterling Control Center:

Report Type	Report Name	
Configuration Management	◆ Functional Authorities Report	◆ Secure+ Cipher Suites Report
	◆ Initialization Parameters Report	◆ Secure+ Key Certificates Report
	◆ Netmap Communication Paths Report	◆ Secure+ Nodes Report
	◆ Netmap Modes Report	◆ Secure+ Trusted Certificates Report
	◆ Netmap Nodes Report	◆ User Proxies Report
Monitoring	◆ Connect:Direct Process Statistics Details	◆ FTP File Transfer Report
	◆ Conncet:Direct Process Statistics Summary	◆ SI Business Process Details
	◆ Connect:Direct Statistics Log Report	◆ SI Business Process Summary
	◆ Connect:Enterprise Batch Statistics Details	◆ SI File Transfer Report
	◆ Connect:Enterprise Batch Statistics Summary	◆ High Watermark Report
	◆ Connect:Enterprise Statistics Log Report	
Node Discovery	◆ Graphical Network Topology Report	◆ Potentially Inactive Netmap Entries Report
	◆ Netmap Connections Summary Report	◆ Potentially Missing Netmap Entries Report
	◆ Node Discovery Topology Report	
System	◆ Alerts Report	◆ Server Inventory Report
	◆ Audit Log Report	◆ Server Status Report
	◆ Sterling Control Center License Report	◆ Service Level Criteria Summary Report
	◆ Database Events Report	◆ Users-Roles Summary Report
	◆ Monthly File Transfer Activity Report	
Other	SLC Debug Report **Note:** This report differs from the other reports listed. It cannot be scheduled, and it is accessed differently from the Tools menu (Tools>Run SLC Debug Report).	

Customized Reports

Sterling Control Center provides the following sample reports in Crystal Reports format:

- Connect:Direct Events
- Connect:Direct Exception Trends
- Connect:Direct Exception Trends Chart
- Connect:Direct Usage Report
- Connect:Direct Usage Report Chart
- Connect:Direct Usage by Server Pair Report
- Connect:Direct Usage by Server Pair Report Chart
- Connect:Direct Usage by Server Pair Detail/Summary Report

You can use these reports without modifications with Crystal Reports, or as templates for further customization. These reports are designed to be used with a MySQL database, although they can be modified for use with other databases. Schemas for the database tables used by Control Center are explained in the *Sterling Control Center Reports Guide.*

Chapter 7

Implementation Scenario

This implementation scenario provides a glimpse into the planning process for a fictitious company named First Country Bank (FCB). It will help you see how the Sterling Control Center building blocks are used to provide the foundation for the work Sterling Control Center will perform in meeting FCB's business objectives.

FCB plans to implement Control Center to gain greater visibility into the condition of their Connect:Direct and Sterling Integrator servers, data flow status, and information about their file transfers.

As part of their planning process, FCB identified objectives for Sterling Control Center, such as:

1. Limit user access and permissions to Sterling Control Center functions and data for FCB's file transfer personnel
2. When servers or adapters are down, generate alerts and send e-mail notifications to personnel
3. When processes are not successfully completed, generate alerts
4. Generate alerts if processes do not run at specified times to meet service level agreements with customers

To accomplish their Sterling Control Center Objectives, FCB analyzed their monitoring goals to design their Sterling Control Center implementation. This process helped them make decisions about their implementation *before* they started configuring the building blocks that defined the work Sterling Control Center would perform. FCB determined that they needed the following building blocks to support their Control Center objectives.

Servers/Server Groups

FCB identified the Connect:Direct and Sterling Integrator servers they need to manage/monitor in their environment and collected information about those servers using the appropriate Sterling Control Center server worksheets for their server types. They used this information when they defined the following servers in Control Center:

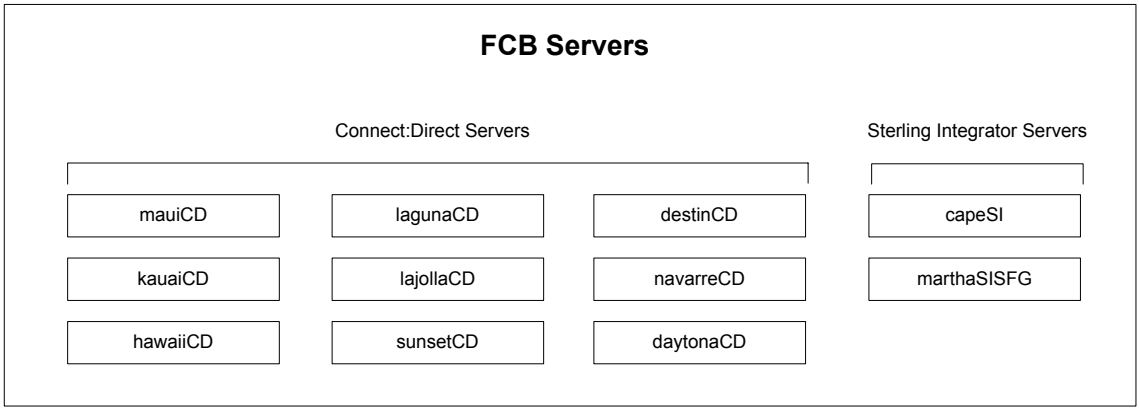

They also decided to group the servers so that user roles can be associated with particular server groups and to make e-mail notifications easier when notifying users about events regarding their assigned servers. They defined the following server groups:

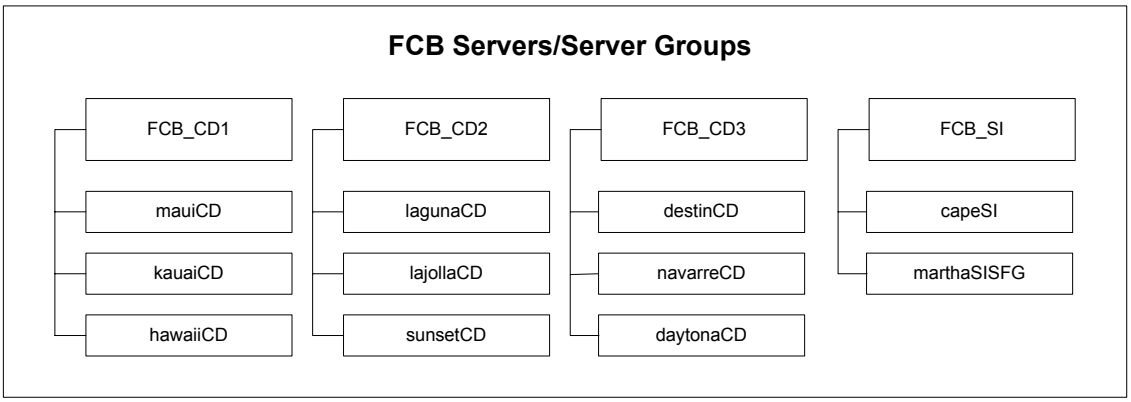

Objective 1—Limit User Access

FCB identified the following roles to control user access to Sterling Control Center:

Role	Responsibilities
Sterling Control Center Admin	A super user who is responsible for installing, configuring, and maintaining Sterling Control Center, including startup/shutdown of the engine and defining subordinate admin roles.
File Transfer Operations Admin	Is a subordinate role that has the “manage” access to configure Connect:Direct nodes and Sterling Integrator adapters, as well as create SLCs, rules, and reports.

Role	Responsibilities
File Transfer Monitoring Staff	Is a subordinate role that has “view only” access for monitoring file transfers and cannot edit artifacts or objects, only view them. They are given view only access to only those objects necessary to their monitoring responsibilities. For example, they do not need to view SLC or rule configuration.

To show the relationships between these roles, FCB developed the following role hierarchy and then defined the roles in Control Center:

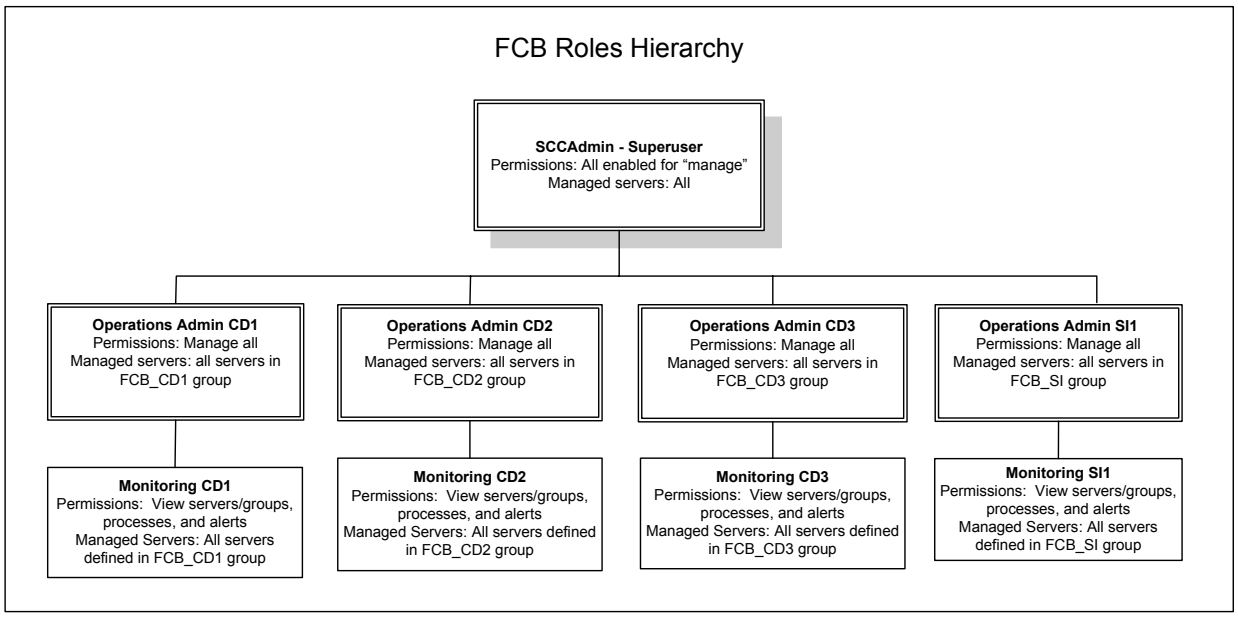

When FCB’s Control Center admin configured users who can access the Sterling Control Center console, an appropriate role was assigned to each user. As a result, when a user logs on to the console, he will have access to the servers and functions associated with his role.

All of the Sterling Control Center console users will be running on a Windows platform. FCB will not require Sterling Control Center to maintain any passwords in its user file and will use the signed on user and the Windows domain as the credentials to allow signon to Sterling Control Center.

Objective 2—Server Down

One of the monitoring objectives FCB identified for Sterling Control Center is the notification of personnel when a server down condition occurs. To accomplish this objective, FCB defined a rule that will be triggered when a server is down. When this rule is triggered, an alert will be generated in the Active Alerts Monitor and an e-mail sent to notify personnel that a server is down.

E-mail List

Because FCB wants to send an e-mail to multiple individuals when an alert occurs, they created an e-mail list in Sterling Control Center. The list, Monitoring staff, contains the e-mail addresses of all personnel who need to be notified when events occur on the servers Sterling Control Center is monitoring. This list was selected when the "Server down" action was defined.

Note: To support e-mail notifications, FCB also had to configure the Sterling Control Center engine's System Setting value for e-mail to specify the location of the SMTP (e-mail) server that Sterling Control Center used to send e-mail.

Field	Value
Name	Monitoring staff
Description	List of all monitoring staff e-mail addresses
To	dan_brown@fcb.com,mike_jones@fcb.com,martha_white@fcb.com,james_kim@fcb.com,gunther_tag@fcb.com,marcus_hite@fcb.com
Permissions	SCCAdmin

Server Down Action

When a server is down, FCB wants Control Center to generate an alert and send an e-mail to the "Monitoring staff" e-mail list. To accomplish this, they created the following action. (This action was selected when the "Server down" rule was defined.)

Field	Value
Name	Server down
Description	Action taken when a server is down.
Email	
To	Monitoring staff (e-mail list)
From	SCCAdmin@fcb.com
Subject	&nodeName; (node name variable) server is down
Message	&nodeName; server is currently down. View this alert in the Active Alerts Monitor and handle as needed.
Alert Severity	1 - High
Permissions (Roles that can edit this action)	SCCAdmin

Server Down Rule

FCB created a “Server down” rule that provides the following instructions for Control Center:

When a server down message id is detected for a monitored server, a high alert will be generated in the Active Alerts Monitor and an e-mail notification will be sent to the monitoring staff.

(When FCB defined the rule, they selected the “Server down” action they had defined earlier.)

Field	Value
Name	Server down
Description	This rule is triggered when a monitored server goes down.
Parameters	
Key	Message Id
Operator	Matches
Value	CCTR034E (Server is down - for monitored servers)
Action	Server down

Tip: A schedule was not required because the rule needs to be in effect 24/7.

The following graphic shows the building blocks that comprise the “Server down” rule:

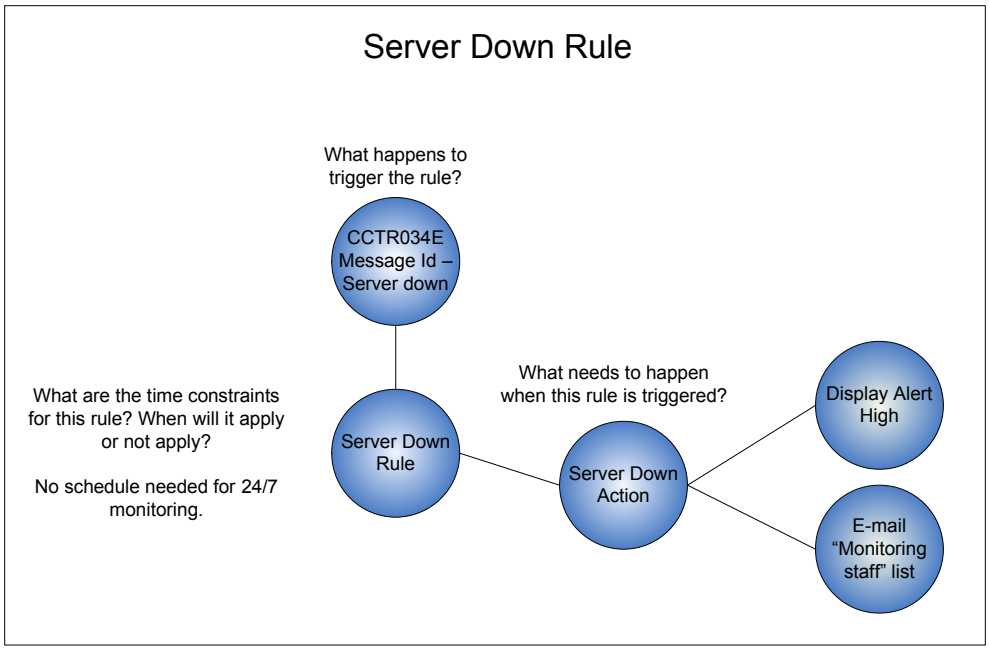

Modifying the Server Down Rule

Because monitoring personnel found that servers were often back up before they handled the alerts, FCB decided to modify the rule to add a second condition (using a linked rule) so that if the server comes back up within 5 minutes, no action will be taken. If the server does not come back up, the "Server down" action will be taken. FCB made the following changes (noted in bold) to the "Server down" rule to accomplish this objective:

Field	Value
Name	Server down
Description	This rule is triggered when a monitored server goes down.
Parameters	
Key	Message Id
Operator	Matches
Value	CCTR034E (Server is down - for monitored servers)
Action	**No operation**
Linked Rules	
Enabled	**Yes**
Parameters	
Key	**Message Id**
Operator	**Matches**
Value	**CCTR033E (Server is up - for monitored servers)**
Resolution Action	**No Operation**
Non-Resolution Action	**Server down**
Timeout	**5 minutes**

Objective 3—Process Completes in Error

Another monitoring objective FCB identified for Sterling Control Center is the notification of personnel when a process completes in error on a monitored server. To accomplish this objective, FCB defined the rule that will be triggered when an error occurs. When this rule is triggered, an alert will be generated in the Active Alerts Monitor.

Process Error Rule

FCB created a "Process error" rule that provides the following instructions for Control Center:

When a process competes with any return code other than 0 on any managed server, a medium alert will be generated in the Active Alerts Monitor.

FCB defined the following information for the rule:

Field	Value
Name	Process error
Description	This rule is triggered when a process completes in error on any monitored server.
Parameters	
Key	Event Type
Operator	Matches
Value	Process Ended
Key	Return Code
Operator	Not Equal To
Value	0
Action	alert2 (predefined action shipped with Control Center)

Tip: A schedule was not required because the rule needs to be monitored 24/7.

The following graphic shows the building blocks that comprise the “Process error” rule:

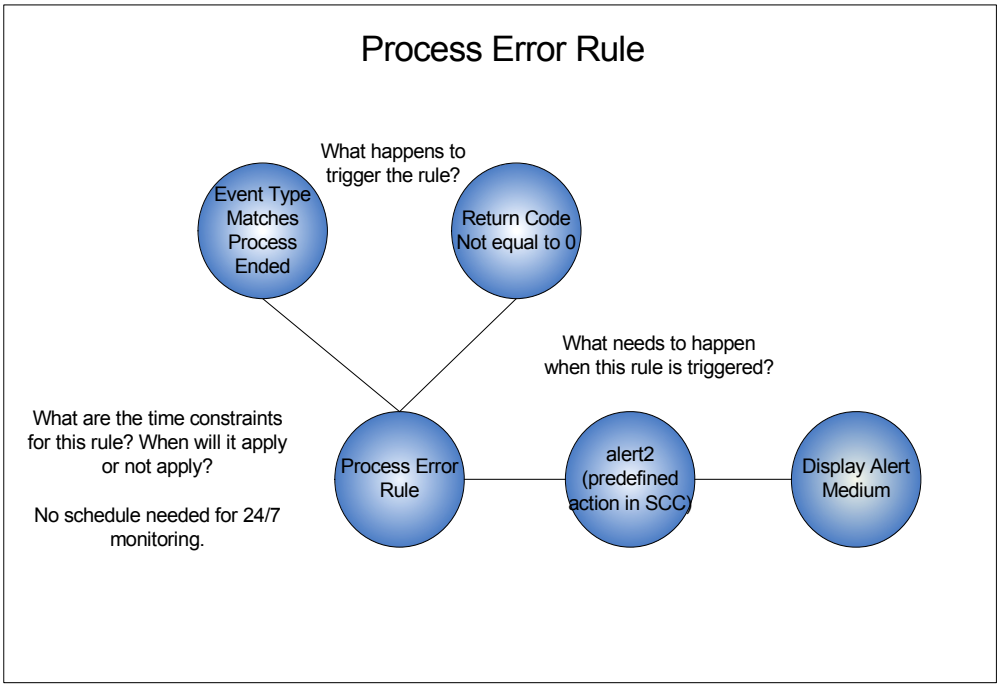

Followup Objective—Specific Process Completes in Error on a Specific Server

Mike Jones, the Control Center administrator for the FCB_CD1 server group, needs to know when a specific process (DailyGrind) on a specific server (hawaiiCD) completes in error. He wants to generate a high alert and be notified by e-mail when an error occurs. To accomplish this objective, he created the following building blocks:

Daily Grind Error Action. When the Daily Grind process encounters an error on the hawaiiCD server, Mike wants Control Center to generate an alert and send an e-mail to him. To accomplish this, he created the following action. (This action was later selected when the "Daily Grind" rule was defined.)

Field	Value
Name	Daily Grind Error
Description	Action taken when the DailyGrind process completes in error on hawaiiCD server.
Email	
To	mike_jones@fcb.com
From	SCCAdmin@fcb.com
Subject	&processName; error
Message	&processName; completed in error on &nodeName;
Alert Severity	1 - High
Permissions (Roles that can edit this action)	Operations Admin CD1

Daily Grind Error Rule. Mike defined the following information for the rule and then placed the Daily Grind error rule higher in the priority sequence than the Process error rule:

Field	Value
Name	Daily Grind Error
Description	This rule is triggered when the DailyGrind process completes in error on any monitored server
Parameters	
Key	Event Type
Operator	Matches
Value	Process Ended
Key	Process Name
Operator	Matches

Field	Value
Value	DailyGrind
Key	Servers and Server Groups
Operator	Matches
Value	hawaiiCD
Key	Return Code
Operator	Not Equal To
Value	0
Action	Daily Grind Error

The following graphic shows the building blocks that comprise the “Daily Grind error” rule:

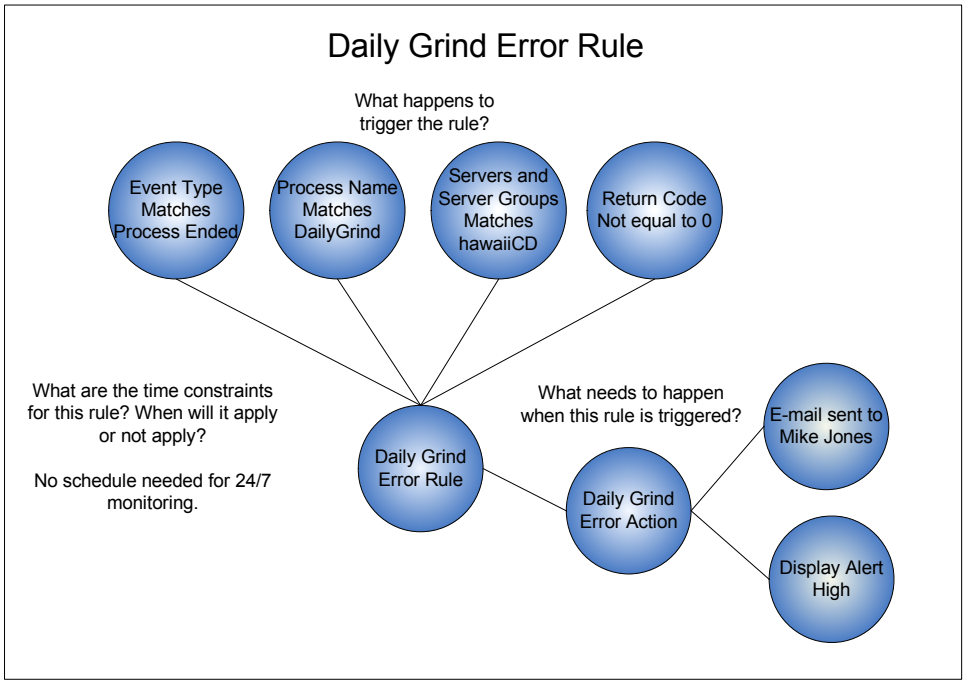

Objective 4—Process Did Not Start at Specified Time

FCB wants to ensure that the EndOfDay process runs on the FCB_CD3 server group at 6:00 p.m. (Central Time), or 18:00, each day. If it does not, they want to be notified. First, they disabled all of the built-in SLC rules to eliminate actions other than those they would setup to meet this objective. To specify a window of time in which this process must begin, they created a wildcard SLC. To tie that SLC to an SLC event, they created two rules and also disabled all of the built-in rules:

- When the EndOfDay process does not start by 6:00 p.m. (18:00), a rule with an action to generate a high alert and send an e-mail to the administrator of the FCB_CD3 server group
- When the EndOfDay process starts late, a rule with an action to clear all alerts associated with the SLC and send an e-mail to the FCB_CD3 server group administrator notifying him that the process started late

To accomplish this objective, FCB created the following building blocks.

End of Day SLC Calendar Schedule

Because FCB wants to monitor the start of the EndOfDay Process, they created an SLC calendar schedule with the following information. (This schedule was later selected when the “End of Day” SLC was defined.)

Field	Value
Name	End of Day
Description	Monitor for start failure of EndOfDay process
Schedule Type	Calendar Schedule
Parameters	
Calendar Name	Daily
Time Zone	(UTC 5:00) Central Time (US & Canada)
Normal Start Range (NSR	
Start Time	17:55
End Time	18:00
Permissions (Roles that can edit this action)	Operations Admin CD3

End of Day SLC

Because FCB wants to monitor for the start of a specific process on a specific server, they created a wildcard SLC with the following information:

Field	Value
Name	End of Day
Description	Monitor start failure for EndOfDay process
Start Window Tolerance	1 hour
End Window Tolerance	1 hour
Generate notification if event has not occurred	Enabled
Enabled	Yes
Server groups	FCB_CD3
Schedules	End of Day
Process Names/Batch IDs	EndOfDay

End of Day Rules

To monitor whether the EndOfDay process starts by 6:00 p.m. (18:00), FCB disabled all of the built-in SLC rules and created two End of Day rules.

End of Day - Not Started Rule. FCB also created an “End of Day - Not Started” rule that provides the following instructions for Control Center:

When a CSLC034E message (Process did not start by NSRe) is generated when the EndOfDay process does not start by 18:00 on any of the servers in the FCB_CD3 server group, a high alert will be generated in the Active Alerts Monitor and an e-mail sent to James Kim, the administrator responsible for the FCB_CD3 server group.

FCB defined the following information for the rule. (They created the action for the rule as they were defining the rule.)

Field	Value
Name	End of Day - Not Started
Description	Alerts when the EndOfDay process has not started by 18:00.
Parameters	
Key	SLC Name
Operator	Matches
Value	End of Day
Key	Message Id

Field	Value
Operator	Matches
Value	CSLC034E
Action	End of Day - Not Started
Description	Action taken when the EndOfDay process has not started by 18:00 on any server in FCB_CD3 server group.
Email	
To	james_kim@fcb.com
From	SCCAdmin@fcb.com
Subject	&processName; not started
Message	&processName; has not started on &nodeName;
Alert Severity	1 - High
Permissions (Roles that can edit this action)	Operations Admin CD3

End of Day - Late Start Rule. FCB also created an "End of Day - Late Start" rule that provides the following instructions for Control Center:

When a CSLC035E message (Process started after NSRe) is generated because the EndOfDay process starts late on any of the servers in the FCB_CD3 server group, an e-mail will be sent to James Kim. No alert will be generated, and all alerts generated for that SLC will be deleted.

FCB defined the following information for the rule. (They created the action as they were defining the rule.)

Field	Value
Name	End of Day - Late Start
Description	Notifies James Kim when EndOfDay process starts late.
Parameters	
Key	SLC Name
Operator	Matches
Value	End of Day
Key	Message Id
Operator	Matches
Value	CSLC035E
Action	End of Day - Late Start

Field	Value
Description	Generates e-mail when the EndOfDay process starts late in FCB_CD3 server group.
Email	
To	james_kim@fcb.com
From	SCCAdmin@fcb.com
Subject	&processName; late start
Message	&processName; has started late on &nodeName;
Alert Severity	0 - In compliance (does not generate an alert and deletes all previously generated alerts for that SLC)
Permissions (Roles that can edit this action)	Operations Admin CD3

Note: For a listing of the message IDs associated with the timeline of this type of SLC, see *SLC Notifications for Calendar Schedule-Based SLCs* in the *Sterling Control Center How-To Guide*.

The following graphics show the building blocks that comprise each of these rules:

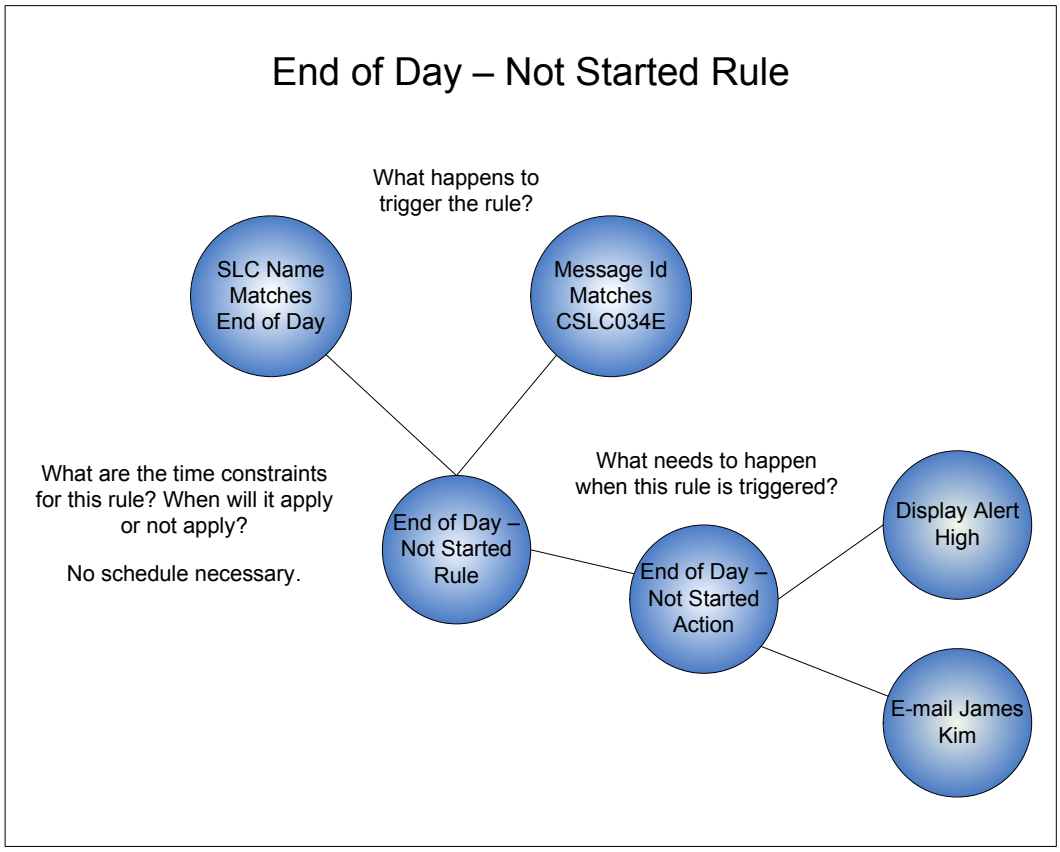

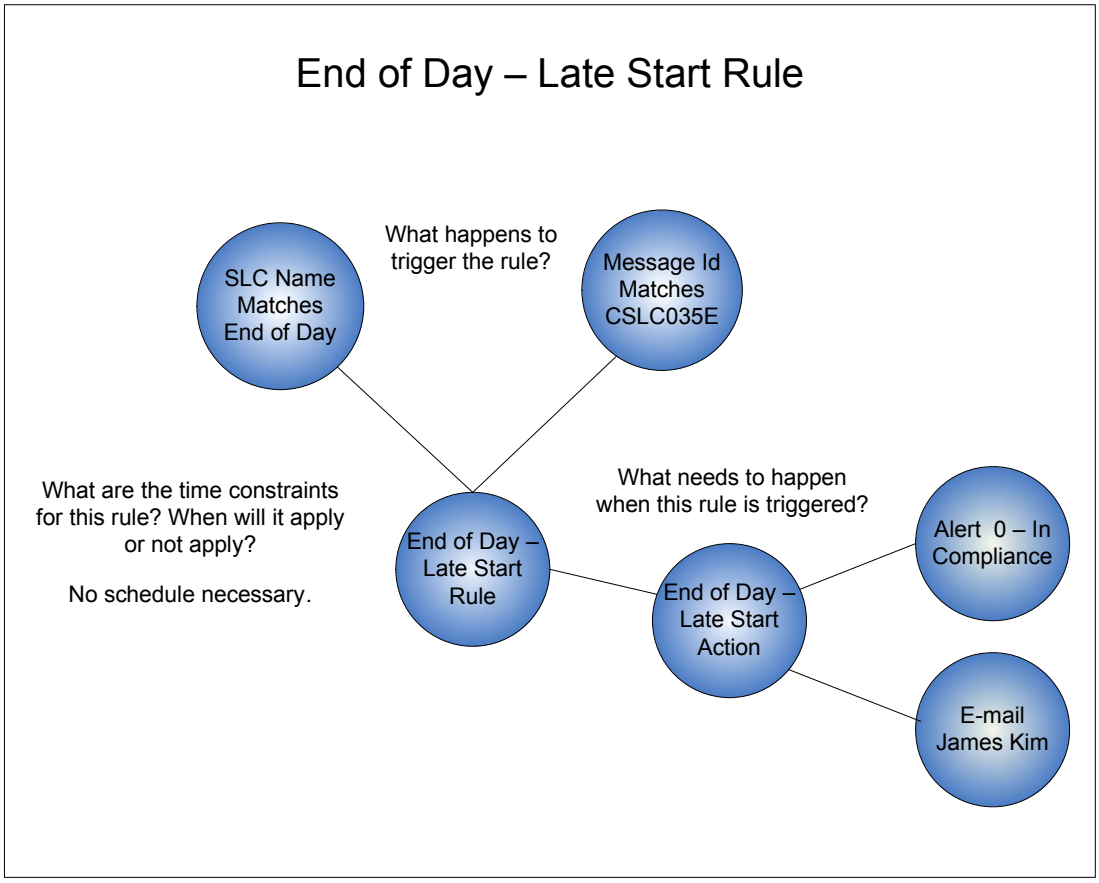
End of Day – Late Start Rule
SLC Name Matches End of Day
What happens to trigger the rule?
Message Id Matches CSLC035E
What are the time constraints for this rule? When will it apply or not apply?
No schedule necessary.
End of Day – Late Start Rule
What needs to happen when this rule is triggered?
End of Day – Late Start Action
Alert 0 – In Compliance
E-mail James Kim

Chapter 8

Best Practices Task List

To more effectively implement Sterling Control Center, you should complete the tasks associated with that implementation in a certain order. This section contains the following information to assist you with implementation tasks for Control Center:

- Description of the product documentation that is referenced in the task list
- Best practices, ordered task list that contains a listing of planning and documentation resources that support each high-level task

Product Documentation

The following product documentation provides you with the information you need to plan for, install, configure, operate, and maintain Sterling Control Center:

Title	Audience and Purpose
Sterling Control Center Release Notes	Provides programmers, network operations staff, and system administrators with the latest release-specific information including last-minute changes and product requirements, as well as other information on installing and implementing Sterling Control Center. Use this guide to get an overview of the current release of the product and any last minute-information you need to know prior to installing Control Center. Read the document in its entirety before installation.

Title	Audience and Purpose
Sterling Control Center Implementation Guide	Provides decision makers, administrators, and users of Sterling Control Center with a conceptual overview of the product. It also provides a best practices task list to aid personnel responsible for implementing Control Center. Use this guide to get the big picture of the concepts, components, and building blocks that comprise Control Center. An understanding of this information will help you identify business objectives for Sterling Control Center and plan how Control Center will meet those objectives. The *Best Practices Task List* guides you through the suggested order that tasks should be performed to effectively implement Control Center.
Sterling Control Center Worksheets	Provide personnel responsible for implementing Sterling Control Center with worksheets they can use to plan their implementation and gather information used when configuring Control Center building blocks.
Sterling MFT License Key Guide	Provides information on using the license keys for all MFT products for personnel responsible for installing software. Use this guide when you are installing Control Center.
Sterling Control Center Getting Started Guide	Provides installation and configuration information for personnel responsible for installing software and maintaining databases. Use this guide when you are planning your Control Center implementation, as well as installing and configuring Control Center and the database software that supports it.
Sterling Control Center System Administration Guide	Provides programmers, network operations staff, and system administrators with the information they need to configure and maintain Sterling Control Center. Use this guide when you are configuring the building blocks that define the work Control Center will perform, maintaining Control Center, and troubleshooting issues.
Sterling Control Center User Guide	Provides operations staff with the information they need to monitor server activity and oversee routine functioning of Sterling Control Center. Use this guide to help you access the wealth of information available through Control Center regarding the servers in your environment.
Sterling Control Center Reports Guide	Provides programmers, network operations staff, and system administrators with the information they need to create and run Sterling Control Center reports. Use this guide to understand the types of reports available in Sterling Control Center and how to create and run those reports in Control Center, as well as working with third-party reporting tools to generate customized reports. This guide also defines Control Center's database table schemas.

Title	Audience and Purpose
Sterling Control Center How-To Guide	Provides programmers, network operations staff, and system administrators with the answers to questions they may have about Control Center functions (arranged in a Q&A format). Use this guide when you are planning your Control Center implementation to gain more insight into how Control Center functionality can be used in very specific situations. As you work with Control Center, if you have a question about how to get it to perform a particular function or how to troubleshoot an issue, scan the table of contents of the How-To Guide. There's a good chance your question is addressed in this helpful guide.
Sterling Control Center Asset Tracking Guide	Provides system administrators with information on using Sterling Control Center's Guided Node Discovery feature with Connect:Direct servers. Use this guide to understand how to find other Connect:Direct servers that a managed Connect:Direct server communicates with.
Sterling Control Center Configuration Management Guide	Provides system administrators with the information they need to manage Connect:Direct server configurations from Sterling Control Center. Use this guide to maintain the following Connect:Direct configuration objects: functional authorities; initialization parameters; Netmap nodes, modes, and communication paths; Secure+ nodes; and user proxies, and to track configuration object changes and versions.
Sterling Control Center Mobile Application Guide	Provides programmers, network operations staff, and system administrators with the information to download and set up Sterling Control Center Mobile on the iPhone.
Sterling Control Center Database Partitioning whitepaper	Provides information on using database partitioning for Sterling Control Center.

Best Practices Task List

The Best Practices Task List is an aid that outlines the high-level steps necessary to implement Sterling Control Center. It describes planning that needs to be done for each task and the documentation references you can access for more information.

Some of these steps may have already been completed by someone in your organization. For example, someone may have already analyzed the environment and installed the hardware and software. You may be tasked with configuring Control Center, so you would start with planning steps that identify your organization's Control Center objectives. The task list will help you identify the high-level tasks that need to be performed to implement Control Center.

Task		Planning Information and Best Practice Notes	Documentation Resources
1	**Planning Your Implementation**		
1a	Identify your business objectives	Identify your Control Center business objectives. Part of this process involves analyzing your environment to determine what you have, what you need, and what you want to do with those resources. This process is vital to an effective Control Center implementation. **Tip:** Locate network diagrams, documents that state Control Center objectives (which may have been identified when your organization purchased Control Center), etc., that might help you in identifying your Control Center business objectives. Use the following planning aid: ◆ High-Level Business Objectives Worksheet	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining Your Control Center Objectives ◆ Understanding Control Center Information ◆ Implementation Scenario **Planning Worksheets**
1b	Analyze your environment	Determine your hardware and software needs and system requirements for engine(s), consoles (GUIs), databases, reports. Use the following planning aids: ◆ Events per Second Worksheet (GSG) ◆ Events per Day Worksheet (GSG) ◆ Event Counter Utility ◆ Platform Configurations Table (GSG) ◆ Database Sizing Worksheet (GSG) ◆ Server worksheet for server type(s)	**Release Notes** ◆ Hardware and Software Requirements **Getting Started Guide** ◆ Before You Install Sterling Control Center ◆ Determining Engine Requirements ◆ Database FAQ **Planning Worksheets**
1c	Gather information for the building blocks that will define Control Center system-level objects: servers, users, e-mail lists, calendars, etc.	You will use this information later when you are creating these building blocks.	**Planning Worksheets**

Task	Planning Information and Best Practice Notes	Documentation Resources
Gather server information	Gather information on the servers in your environment using the server worksheet for your server type(s). **Tip:** Because you must have a valid user ID and password to access each server, you may want to setup a generic user ID and password to access servers more easily. **Tip:** For Connect:Direct servers, you can add one server and then use the Control Center Guided Node Discovery feature to get a list of the nodes that server communicates with based on netmap entries and statistics logs. See *Performing Guided Node Discovery* in the *System Administration Guide*.	**Implementation Guide** ◆ Defining Servers **Planning Worksheets** **System Administration Guide** ◆ Manage Servers
Decide how you will group servers	Decide how you will group servers in a way that makes sense for your environment, for example, geographic location, service line, etc. Use the following planning aids: ◆ Server worksheet for server type(s) ◆ Server Groups Worksheet	**Implementation Guide** ◆ Defining Servers **Planning Worksheets** **System Administration Guide** ◆ Manage Servers
Decide if you will use data visibility groups (DVGs)	Decide if you will use data visibility groups to limit the data users can monitor. I so, determine a way to segment access to data that makes sense for your environment, for example, by functional department such as accounting or payroll. Use the following planning aids: ◆ Data Visibility Groups Worksheet ◆ User Access Worksheet	**Implementation Guide** ◆ Understanding Data Visibility Groups ◆ Defining User Access **Planning Worksheets** **How-To Guide** **System Administration Guide** ◆ Manage Data Visibility Groups ◆ Manage Roles and Users

Task	Planning Information and Best Practice Notes	Documentation Resources
Decide how you will implement user roles and permissions	Develop a role hierarchy to capture the user roles/permissions needed in your environment using the superuser and user roles as a basis. Use the following planning aid: ◆ User Access Worksheet	**Implementation Guide** ◆ Defining User Access **Planning Worksheets** **System Administration Guide** ◆ Manage Roles and Users
Determine if you need to implement your organization's password policy	If your organization enforces a password policy, get a copy of it. To implement your password policy, you will edit the passwordPolicy.xml file located in the <installation directory>\ControlCenter\conf\security folder when you configure user access to Control Center in a later task.	**Implementation Guide** ◆ Defining User Access **Planning Worksheets** **System Administration Guide** ◆ Manage Roles and Users
Decide which users need to access Control Center and decide which role they will be assigned and whether any of their permissions will differ from their assigned role	Gather a list of users who will need access to Control Center from the consoles and decide the permissions they need to be granted. Refer to the user permissions table in the *Manage Roles and Users* section of the *System Administration Guide* for a listing of the permissions users can be assigned. Use the following planning aid: ◆ User Access Worksheet	**Implementation Guide** ◆ Defining User Access **Planning Worksheets** **System Administration Guide** ◆ Manage Roles and Users
Decide what type(s) of notifications you will use to handle alerts	Gather information you will use to define SMTP e-mail settings and SNMP settings for host computers where SNMP traps are sent. The specific elements a trap contains are dictated by the contents of the Control Center configuration file named SnmpAdaptorWrapper.xml. The Management Information Block (MIB) that defines the Object Identifiers (OIDs) contained in the trap is shipped with Control Center. It is named SterlingMIB.mib and is located in the esm installation directory folder.	**Getting Started Guide** ◆ Configure SMTP Settings for E-mail Messages ◆ Configure SNMP Settings

Task		Planning Information and Best Practice Notes	Documentation Resources
1d	Identify the building blocks needed to meet your monitoring objectives	For the objectives you identified, plan how you will use the building blocks to define the work Control Center will perform. **Tip:** As you are going through this process, be sure to plan for reuse of building blocks where possible such as an action used in more than one rule or a calendar used in multiple SLCs. This will simplify the ongoing maintenance of these objects. Use the following planning aid you completed earlier in the task list: ◆ Business Objectives - Service Level Management Worksheet	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining the Work ◆ Implementation Scenario **Planning Worksheets** **How-To Guide**
	Decide which rules you need to meet your business objectives for monitoring servers	Your rules need to be based on the business objectives you identified earlier in the planning process. The rules you define will provide Control Center with instructions that specify the event to listen for. Use the following planning aids to help you determine the rules you need: ◆ Business Objectives - Service Level Management Worksheet	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining the Work ◆ Implementation Scenario **Planning Worksheets** **How-To Guide** ◆ Setting Up Actions **System Administration Guide** ◆ Manage Rules and Actions

Task		Planning Information and Best Practice Notes	Documentation Resources
	Define the service level criteria that will enable you to determine whether service level agreements have or have not been met	Use the following planning aid: ◆ Business Objective Worksheet- Service Level Management	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining the Work ◆ Implementation Scenario **Planning Worksheets** **How-To Guide** ◆ Setting Up SLCs ◆ Troubleshooting SLCs **System Administration Guide** ◆ Manage Service Level Criteria
2	**Preparing the Environment**		
2a	Identify required patches, service packs, and releases for managed servers	Download any product updates from the Sterling Commerce Customer Center web site.	**Release Notes** ◆ Hardware and Software Requirements
2b	Verify that the system hardware is in place and ready for the installation		**Release Notes** ◆ Hardware and Software Requirements
2c	Install the database software	Record the following information from the database administrator for use later during the Sterling Control Center installation: ◆ Database name ◆ User ID and password for databases ◆ Location of the JDBC driver for the database Make sure you have the correct objects for the database you chose (for example, DB2 page size). For more information, consult the product documentation for the database you chose. If	**Other** Product documentation for the database you chose **Getting Started Guide** ◆ Before You Install Sterling Control Center

Task		Planning Information and Best Practice Notes	Documentation Resources
2d	Create the staging and production databases for your database type	If your database supports partitioning, you should take advantage of it. If you implement database partitioning, you will only need to create the production database, not the staging database.	**Getting Started Guide** ◆ Before You Install Sterling Control Center
2e	If consoles will access the Sterling Control Center engine using a secure connection, configure the HTTPS connection		**Getting Started Guide** ◆ Determine HTTPS Information for the Engine and Console Connection
3	**Installing Control Center**		
3a	Install Sterling Control Center		**Release Notes** **Getting Started Guide** ◆ Before You Install Sterling Control Center ◆ Install the Engine and Console on UNIX or Installing the Engine and Console on Windows
3b	Install the license key	Obtain the license key. Make sure that your license supports the number and type of servers you plan to manage/monitor.	**Sterling MFT License Key Guide**
3c	After you install the engine, if you plan to use the Control Center console locally on computers where Control Center is installed, install the console.	The Console can be started by accessing the engine's ”launch" page via a browser once the engine is started.	**Getting Started Guide** ◆ Working With the Console
3d	If the console will access the engine using a secure connection, configure a secure connection between the engine and the console.		**Getting Started Guide** ◆ Configuring a Secure Connection
3e	Install an FTP agent (if you will manage FTP severs using Control Center)	One agent per FTP server must be installed (except for the z/OS FTP server).	**Getting Started Guide** ◆ Set Up Control Center to Monitor FTP Agents
4	**Starting and Accessing Sterling Control Center**		
4a	Start Control Center		**Getting Started Guide** ◆ Start and Stop Control Center

Task		Planning Information and Best Practice Notes	Documentation Resources
4b	Invoke the console to gain access to Control Center	After you log into the console, be sure to change the admin password. If you are unable to log into the Console, you should check the Engine log file for problems that may have caused the engine to not start.	**Implementation Guide** ◆ User Interfaces **Getting Started Guide** ◆ Working With the Console
5	**Changing settings and tuning Control Center**	To ensure optimal operation, you can change system settings to tune Control Center. For example, if the engine time is different from the server time, adjustments need to be made in the time Preferences Settings.	**System Administration Guide** ◆ Sterling Control Center Settings ◆ Tuning Sterling Control Center
5a	Configure Control Center to send e-mails or traps	Configure the SMTP e-mail settings and SNMP settings for host computers where SNMP traps are sent.	**Getting Started Guide** ◆ Configure SMTP Settings for E-mail Messages ◆ Configure SNMP Settings
6	**Defining Servers**		
6a	Add servers	Use the Server Worksheet you completed earlier in the task list. **Tip:** For Connect:Direct servers, you can add one server and then use the Guided Node Discovery feature to get a list of the nodes the server communicates with based on netmap entries and statistics logs. See *Performing Guided Node Discovery* in the *System Administration Guide*. **Tip:** After adding a server of a specific type, you may use that definition as a template for other servers to be added if you use the duplicate capability from the Server List view.	**System Administration Guide** ◆ Manage Servers ◆ Perform Guided Node Discovery **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
6b	Create server groups	Use the Server Groups Worksheet you completed.	**System Administration Guide** ◆ Manage Servers **Implementation Guide** ◆ Defining the Work **Planning Worksheets**

Task		Planning Information and Best Practice Notes	Documentation Resources
7	**Defining User Access**		
7a	Define data visibility groups (DVGs)	Use the information you defined on the Data Visibility Group (DVG) Worksheet	**System Administration Guide** ◆ Manage Data Visibility Groups **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
7b	Define user roles	Use the information you defined on the User Access Worksheet.	**System Administration Guide** ◆ Manage Roles and Users **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
7c	Define a user password policy	Use the information you defined on the User Access Worksheet.	**System Administration Guide** ◆ Manage Roles and Users **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
7d	Add users	Use the information you defined on the User Access Worksheet.	**System Administration Guide** ◆ Manage Roles and Users **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
8	**Setting Up Calendars and Schedules**	**Tip:** When defining calendars and schedules, use a descriptive naming convention that allows users to know what they are by name. This will save users the time of looking at the object's properties to determine their use and facilitate object reuse.	

Task		Planning Information and Best Practice Notes	Documentation Resources
8a	Set up calendars	Use the Business Objective Worksheet - Service Level Management you completed earlier in the task list.	**System Administration Guide** ◆ Manage Calendars and Schedules **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
8b	Set up schedules	Use the Business Objective Worksheet - Service Level Management you completed earlier in the task list.	**System Administration Guide** ◆ Manage Calendars and Schedules **Implementation Guide** ◆ Defining the Work **Planning Worksheets**
9	**Create actions that will be invoked by a rule when an event occurs**	Will the default actions that are included in Control Center work for you? If not, create actions that will meet your needs. These actions will be used when you are creating rules. **Tip:** You can create actions while you are creating rules; however, a planned, methodical approach of creating them up front will help you create only the actions you absolutely need. You can then select these actions when creating rules. (One action might be used in many different rules.) This process will result in fewer actions to maintain. **Tip:** Be sure to give any actions you create a descriptive naming convention that allows users to know what they are by name. This will save users the time of looking at the object's properties to determine their use and facilitate object reuse. Use the following planning aids you completed earlier in the task list: ◆ Business Objectives - Service Level Management Worksheets	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining the Work ◆ Implementation Scenario **Planning Worksheets** **How-To Guide** ◆ Setting Up Actions **System Administration Guide** ◆ Manage Rules and Actions

Task		Planning Information and Best Practice Notes	Documentation Resources
10	**Create rules that specify an event that Control Center will listen for**	Your rules need to be based on the business objectives you identified earlier in the planning process. Use the following planning aids as you create rules: ◆ Business Objectives - Service Level Management Worksheets **Tip:** As you are creating rules and SLCs, you should approach this process in a slow-roll, phased approach where you define a rule or an SLC for a server. If you get the results you want, then add another rule or an SLC and test, and add another and test, and so on. This will allow you to troubleshoot rules and SLCs as you go before adding more complexity.	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining the Work ◆ Implementation Scenario **Planning Worksheets** **How-To Guide** ◆ Setting Up Rules ◆ Troubleshooting Rules **System Administration Guide** ◆ Manage Rules and Actions
11	**Create the Service Level Criteria that will enable you to determine whether service level agreements have or have not been met**	Use the following planning aids: ◆ Business Objectives - Service Level Management Worksheets **Tip:** As you are creating rules and SLCs, you should approach this process in a slow-roll, phased approach where you define a rule or an SLC for a server. If you get the results you want, then add another rule or an SLC and test, and add another and test, and so on. This will allow you to troubleshoot rules and SLCs as you go before adding more complexity.	**Implementation Guide** ◆ Understanding Sterling Control Center Concepts and Components ◆ Defining the Work ◆ Implementation Scenario **Planning Worksheets** **How-To Guide** ◆ Setting Up SLCs ◆ Troubleshooting SLCs **System Administration Guide** ◆ Manage Service Level Criteria

Task		Planning Information and Best Practice Notes	Documentation Resources
12	**Ongoing Administration of SCC**		
	Changing SCC Settings	In the ongoing administration of Control Center, you may need to change system settings, such as the following: ◆ Database settings ◆ E-mail settings ◆ SNMP host settings ◆ Engine connection settings ◆ Console settings ◆ License key settings ◆ Console preferences	**System Administration Guide** ◆ Sterling Control Center Settings
	Downloading and Setting Up Sterling Control Center Mobile	If you want to use Sterling Control Center Mobile, you must download and set up the application.	**Sterling Control Center Application Guide**
	Improving Performance	You can perform general tuning to improve the performance of the Control Center engine.	**System Administration Guide** ◆ Tuning Sterling Control Center
	Updating licenses		**System Administration Guide**
	Managing Configuration Objects	If you want to create, update, or delete multiple configuration objects such as actions, rules, schedules, e-mail addresses, and SLCs without manually entering each one or using the duplicate function in the console, you can use the Batch Creation Utility), sample script, and sample templates provided with Sterling Control Center. **Tip:** You can also copy configuration objects between Control Center configurations to prepare for disaster recovery or to copy a production instance to a test instance.	**System Administration Guide** ◆ Create Multiple Objects ◆ Copy Configuration Objects Between Control Center Installations
	Troubleshooting SCC		**System Administration Guide** ◆ Administrative Troubleshooting

Appendix A

Sterling Control Center Terms and Concepts

The following table defines terms and concepts used in Sterling Control Center. An understanding of this information will assist you in working with Control Center.

Term	Definition
Action	An activity (or activities) that Sterling Control Center performs when an event triggers a rule. One action may be referred to/utilized by multiple rules.
Alerts	Providing proactive notification for at-risk business processes in the form of alerts viewable through the Sterling Control Center Console. Alerts are a Control Center event with an alert element that has a severity value between 0 and 3. They are created when an event triggers a rule that has an alert action.
Arrived file	A message in a mailbox that Sterling File Gateway monitors, causing Sterling File Gateway to perform some activity on it.
Asset tracking	Also referred to as guided node discovery, or node discovery. A feature of Sterling Control Center that enables you to find the other servers that a Connect:Direct server communications with.
Batches	Data file residing in a mailbox of the repository on the Connect:Enterprise host computer. When a batch is added to the repository, it is assigned a unique number (from 1 to 9,999,999).
Business process	A series of activities that accomplishes a business objective.
Calendar	A calendar describes a range of dates. After they are defined, calendars can be used to tell Control Center both when and when not to monitor. This includes how long the calendar remains in effect, and how often processing is repeated (recurrence), such as patterns within the days of the week, patterns within the days of the month, or specific days of the year. Sterling Control Center comes with predefined calendars for each weekday, as well as a daily schedule. You can use these calendars for both rule schedules and SLC schedules, and you can also create additional calendars to meet your processing needs. Calendars can also be used for Scheduled Reports.

Term	Definition
Calendar schedule	Associates a calendar with a schedule, and between the two of them, defines both the processing window (start and end) and the days those processing windows occur. For example, a process must start between 19:00 and 19:30 and end between midnight and 00:30. Calendar schedules are useful for monitoring processing that occurs at fixed times, whereas duration schedules are used when an event can start at any time but must complete within a specified amount of time once started, for example, 15 minutes.
Cluster	A group of Sterling Integrator nodes that share the workload.
Configuration management	A feature of Sterling Control Center that enables you to manage the configurations of your Connect:Direct for UNIX, Windows, and z/OS servers.
Connect:Direct	Point-to-point file transfer software for high-volume, assured data delivery of files within and between enterprises.
Connect:Direct Browser	Allows you to create, submit, and monitor Connect:Direct processes from an Internet browser. You can also perform Connect:Direct system administration tasks, such as viewing and changing the network map or initialization parameters, from the Browser User Interface if you have the appropriate authority.
Connect:Direct Secure+ Option	Provides secure, reliable high-volume data exchange between business-critical applications, both within the enterprise and with external business partners. It provides data confidentiality, message integrity checking, and server and client authentication.
Connect:Direct Process	A group of statements that provide instructions for transferring files, running programs, submitting jobs on the adjacent node, and altering the sequence of process step execution.
Connect:Direct Select	Provides reliable and secure unattended data delivery between remote sites (where Connect:Direct Select is installed) and a Connect:Direct server. In its basic configuration, a Connect:Direct Select node sends files from a watch directory or e-mail inbox to a Connect:Direct server and receives files from the Connect:Direct server.
Connect:Enterprise	Provides a secure means of collecting and distributing files using mailboxing, automation, and utilizing an extensive, open protocol set.
Console	The graphical user interface for Sterling Control Center. There is a console and a web console. The console provides the full set of capabilities for Control Center administrative functions, such as configuring Control Center, defining servers to be monitored, configuring rules and notifications for monitoring, and receiving and handled alerts. The web console has a subset of the console functionality, which is useful for self-service monitoring and generating reports.

Term	Definition
Correlator	Enables you to associate multiple milestones within a workflow. For example, a business process that creates an order can be associated with the business process that creates the invoice for the order by using the order number as the correlator.
Database partitioning	Allows the data in the production database to be partitioned by date, which can improve database performance and reduce database maintenance (for example, index rebuild). When database partitioning is used, data is not moved to the staging database, eliminating the need for the staging database.
Data Visibility Group	Limit what events (data) a specific user can monitor. When a data visibility group (DVG) is assigned to a role, the role is restricted. Any users who are assigned that role are restricted users. DVGs can be specified in rules and SLCs. When events match on any criteria for a DVG, that DVG name is put into the DVG attribute of the event. Therefore, the event is “tagged” with that DVG.
Delivery	A record of the activities Sterling File Gateway takes to deliver a file to a specific consumer endpoint.
DVG	Data visibility group
Duration schedule	Uses a minimum duration and a maximum duration to define the processing window. Processing can begin at any time but should end within a specified time frame, for example 15 minutes. A duration schedule is also used to identify the percentage of processing that is complete for an SLC.
Engine	The server portion of Sterling Control Center that is installed on a network computer.
Event	Records of activity occurring on servers managed and monitored by Control Center as well as within Control Center itself, which are persisted by Control Center in a central location. Events may trigger Control Center rules.
Guided Node Discovery	The process of checking the network map and the statistic records of a selected Connect:Direct server and identifying other servers it has communicated with.
High availability	A configuration where a backup Sterling Control Center engine and/or database is used for failover when the primary engine and/or database is down.
Java	A programming language that allows Sterling Control Center to be developed to run on a variety of machines. The Java language makes it possible to develop software that is portable, modular, and secure.
Java Web Start	Allows users to start the Control Center console directly from the Sterling Control Center Launch page using their browser and ensures that the most current version of the console is deployed. Accessing the console in this manner gives you full console functionality.

Term	Definition
Linked rule	A regular rule that specifies two actions: one that is performed when the second condition is met and one that is performed when the second condition is not met, within the resolution time specified. a second condition and an action that is performed if the second condition is or is not resolved within a user-specified timeout period.
Metadata	Metadata is information about data. In Sterling Control Center, you can add metadata to events using metadata rules and metadata fields associated with monitored servers. For more information on metadata, see *Understanding Metadata Rules* on page 43.
Metadata rule	Allows you to append additional elements and values to Sterling Control Center events before they are processed by the SLC service or rule service.
Netmap	Also known as network map. The Connect:Direct netmap is a file that contains configuration information for a Connect:Direct server. This information describes the local node and the remote nodes that server can communicate with in the network. The guided node discovery feature in Sterling Control Center accesses the netmap of a managed server to generate a list of all nodes that server communicates with.
Notification	Providing proactive notice for at-risk business processes in the form of e-mails, SNMP traps, and alerts viewable through the Control Center console.
Object	Sterling Control Center components such as actions, e-mail lists, rules, schedules, Connect:Direct netmap nodes, and Connect:Direct functional authorities.
Password policy	In user authentication, settings that specify criteria that must be met when setting passwords, such as minimum and maximum password lengths. Password policies are often defined within an organization as part of their overall security policy.
Permissible Objects	Control Center building blocks to which restricted roles can be assigned. Visibility of the object can also be specified for the object: visible to only the users in the selected restricted roles (private) or to all users (public). Permissible objects include rule and metadata actions; rule, SLC, report, and metadata schedules; message and email lists; and calendars.
Polling	When Sterling Control Center repeatedly requests data from a monitored server.
Process data	In reference to Sterling Integrator business processes, process data defines a context that exists for each process instance and can be used to hold or reference information accumulated during the life of the process.

Term	Definition
Production database	Where Control Center records the information gathered from the monitored servers for historical purposes (for example, ad hoc select statistics and user reports).
Route	In Sterling File Gateway, a route is a record of all the activities performed on a routable payload, once it known who the consumer is. Each routable payload is associated with a route.
Rule	A system instruction you create and that Sterling Control Center executes automatically. An event triggers a rule, which in turn invokes the action associated with that rule, such as generating an alert.
Schedule	Dictates time-based constraints for when events will be matched against a rule's or SLC's criteria and when they won't be. Schedules are also used to tell when to run an automated report. There are multiple types of schedules: calendar and duration for SLCs, rule schedules for rules, and report schedules for automated reports.
Service Level Criteria (SLC)	Performance objectives that require processing to occur within a certain time window. You can use Sterling Control Center to set up SLCs and monitor processing against them. There are three types of SLCs: standard, wildcard, workflow, and simple.
Simple SLC	Enable you to create an SLC by answering a few basic questions, specifying values for basic parameters, and giving the SLC a name and description. When you create a simple SLC, all necessary objects to support the SLC, such as rules, actions, and schedules, are also created.
Sizing	The process of determining how many Sterling Control Center engines and databases you need to handle the workload for your server environment.
SNMP trap	Notification generated as part of processing criteria defined in a rule and sent to one or more Simple Network Management Protocol (SNMP) hosts.
Staging database	Where Sterling Control Center stages older data for the database administrator to export to long-term storage before it is purged.
Standard SLC	Enables monitoring using specific process names, file names, etc. Use standard SLCs when you know the specific item to monitor.
Sterling Control Center Launch Page	The Java Web Start page used to access the Sterling Control Center console, web console, documentation, and other supporting applications and interfaces.
Sterling File Gateway	An application for transferring files between partners using different protocols, file naming conventions, and file formats.
Sterling Integrator	A business process-centric transaction engine for modeling and managing business processes. Sterling Integrator supports high-volume electronic message exchange, complex routing, translation and flexible integration, and real-time interaction with multiple internal systems and external business partners.

Term	Definition
Visibility Criteria	A set of parameters, one or more of which are used to construct/define Data Visibility Groups (DVGs).
Web Console	A browser-based version of the Control Center console with limited functionality, which can also be started from the Sterling Control Center Launch Page. Tasks such as viewing and handling alerts and generating reports can be performed using the web console; however, tasks to create the rules that generate alerts and to set up the reports would be performed using the Control Center console.
Wildcard SLC	Enables monitoring of processes, file names, etc., that may have multiple values such as Batch IDs with the date and time in their names. To specify monitoring criteria in wildcard groups, you can use the asterisk and question mark wildcard characters, or regular expressions.
Workflow SLC	Enables monitoring of the flow of related processes or process steps by tracking them as milestones in a workflow. For example, a workflow SLC can monitor a transaction consisting of three processes, all of which must finish within three hours of the first process's initiation.