Sterling Control Center

Implementation Guide

Version 5.4.1

Sterling Control Center

Implementation Guide

Version 5.4.1

Contents

Chapter 1. Sterling Control Center implementation overview

If you are tasked with implementing IBM® Sterling Control Center, you need to gain a high-level understanding of Sterling Control Center to assist you in planning your implementation.

To help you understand Sterling Control Center, the following information is provided:

- 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 Sterling Control Center will perform to meet your objectives. The interrelationships between the building blocks are also described.
- A high-level explanation of the types of information Sterling Control Center generates about the servers in your environment
- Sample implementation scenarios that illustrate how the building blocks are used to meet business objectives
- An ordered, best practices task list that provides planning considerations and references to planning tools and documentation to help you complete each high-level task

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

Chapter 2. Features and benefits

Sterling Control Center is a centralized monitoring and management system.

It gives operations personnel the capability to continuously monitor business activities across the enterprise for the following server types:
- IBM Sterling Connect:Direct®
- IBM Sterling Connect:Direct File Agent
- IBM Sterling Connect:Enterprise®
- IBM Sterling B2B Integrator
- IBM Sterling File Gateway
- IBM Sterling Connect:Express
- IBM QuickFile
- IBM MQ Managed File Transfer
- Many FTP servers

In addition, you can manage the configurations and licenses of Sterling 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 has three primary features:
- Service level management
- Asset management
- Configuration management

Service level management

Sterling Control Center can help you answer questions about activities in your managed file transfer environment.

It helps you answer questions such as:
- Did my business process run on time?
- Did my file transfer take place when it should have?
- Are my servers operating the way they should?

Sterling Control Center gives you tools to effectively monitor and manage your environment by giving you a common, centralized view of that environment. This insight into the environment helps you offer higher levels of service to your internal and external customers. Sterling Control Center accomplishes this service level management 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 the 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 emails, SNMP traps, and alerts.
- Consolidating information for throughput analysis, capacity planning, post-processing operational or security audits, and workload analysis. This consolidation helps 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
  - The necessity of separately configuring 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?" You can manage assets by using the asset tracking and license management features of Sterling Control Center.

Asset tracking

Track network assets by capitalizing on the server monitoring capabilities of Sterling Control Center. Using Guided Node Discovery (also known as Node Discovery), you can find all Sterling 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 Sterling Connect:Direct servers in your environment by using license management.

You can use Sterling Control Center to import updated server licenses to a central license management repository. On an ad hoc basis, you can push the server licenses to managed Sterling Connect:Direct servers. You can also configure Sterling Control Center to monitor a POP3 or IMAP mailbox for emails that contain license key file attachments. When Sterling Control Center finds a license key file attachment, it validates the license and automatically imports it into the license repository. 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 configuration questions about your Sterling Connect:Direct servers such as, "Are my servers configured correctly" and "Do they comply with our security policy?"

Sterling Control Center provides you with a centralized, simplified means of managing your Sterling Connect:Direct for UNIX, Sterling Connect:Direct for Microsoft Windows, and Sterling Connect:Direct for 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) of configuration data for Sterling Connect:Direct servers.
- Offering a way 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 IBM Sterling Connect:Direct Secure Plus on remote Sterling Connect:Direct servers.
- Alerting you about Sterling Connect:Direct Secure Plus certificate expiry.
- Offering a way for you to quickly identify Sterling Connect:Direct servers that do and do not have Sterling Connect:Direct Secure Plus installed through information displayed in the Server List View.

Benefits by server type

Sterling Control Center provides benefits for each server type supported.

Sterling Connect:Direct

- Provides centralized visibility and control of large-scale, distributed Sterling Connect:Direct server environments by enabling you to consolidate and collect data for various 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 Sterling Connect:Direct and Sterling Connect:Direct Secure Plus.
- Provides for centralized asset management (software version and license key).

Sterling Connect:Direct File Agent

- Allows you to know when a Sterling Connect:Direct File Agent submits processes to a Sterling Connect:Direct server and to what server it submits those processes.
- Allows you to know when processes are not submitted to a Sterling Connect:Direct server.
- Enables you to know when a user has updated the configuration data for a Sterling Connect:Direct File Agent.

Sterling Connect:Enterprise

- Provides visibility to files transferred into and out of mailboxes.
- Allows monitoring of daemons.

Sterling B2B 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 B2B Integrator queues.
- Enables you to view Sterling B2B Integrator license details, Java environment details, location of Sterling B2B 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).

Sterling Connect:Express

- Provides visibility of files transferred.
- Allows you to centrally monitor file transfers.
- Lets you configure notification about files transfers that did or did not occur or are late.

QuickFile

- Provides visibility of files transferred.
- Allows you to centrally monitor file transfers.
- Lets you configure notification about file transfers that did or did not occur or are late.

MQ MFT

- Provides visibility of files transferred.
- Allows you to centrally monitor file transfers
- Lets you configure notification about file transfers that did or did not occur or are early or late.

FTP servers

- 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 the IBM Sterling Managed File Transfer suite of products and solutions.

Chapter 3. Concepts and components

A basic understanding of Sterling Control Center components and how they work together will help you plan and implement Sterling 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 tiers, or layers, are 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 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 separation 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 degree of monitoring can cause a considerable load on the system. Separating core application functionality into different tiers allows Sterling 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 presentation tier displays information through graphical user interfaces (GUIs). Users log on to the Sterling Control Center engine through this tier to configure Sterling Control Center, as well as 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, an engine, and databases.

Each supported server type has a corresponding node service on the Sterling Control Center engine. Events from the servers are passed through a series of services on the Sterling Control Center engine for processing: visibility service, metadata service, rule service, and SLC service. The events trigger rules, and associated actions, such as displaying a GUI alert or sending an SNMP trap or email. Events and statistics are stored in the Sterling Control Center databases to provide a historical record of activity and reporting capabilities. User interfaces provide configuration, management, and monitoring capabilities for supported servers.
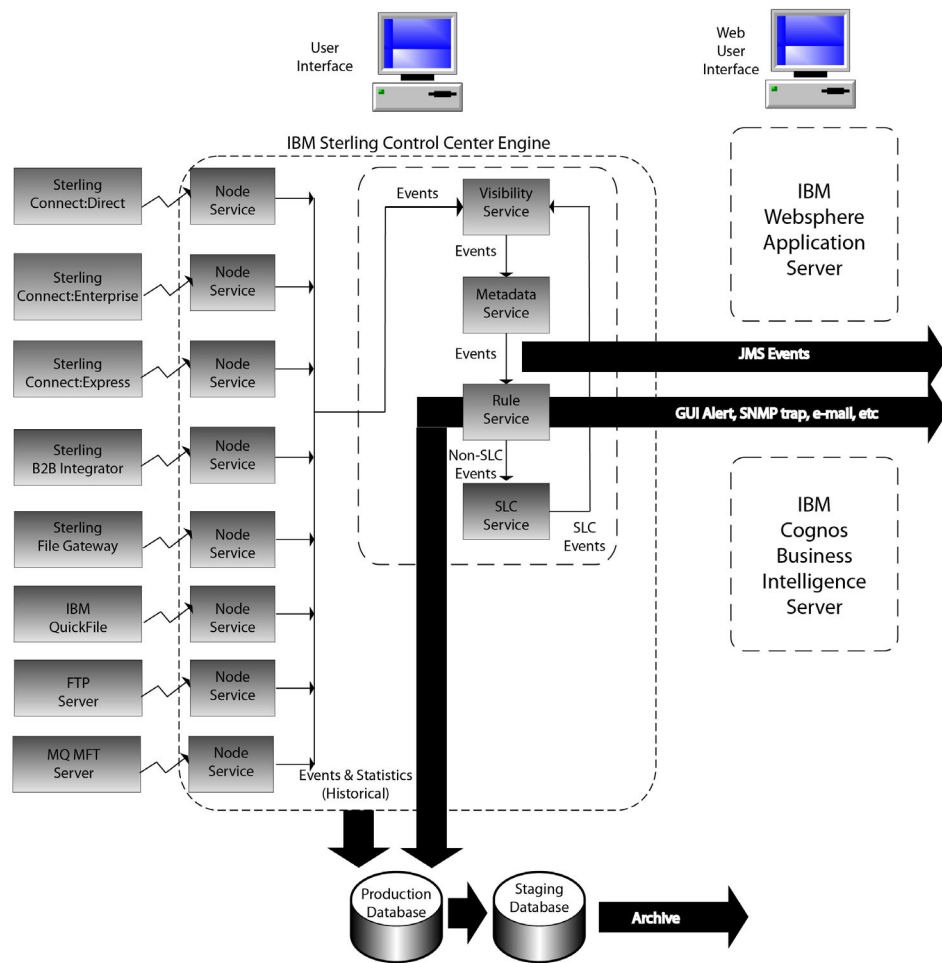
User interfaces

Sterling Control Center has two user interfaces: the console and web console. In addition, Sterling Control Center Mobile is an iPhone application that interfaces with Sterling Control Center.

The Sterling Control Center console enables you to configure Sterling Control Center and Sterling Connect:Direct nodes. Both the Sterling Control Center console and web console enable you to display information gathered from the engine via an HTTP or HTTPS connection. In addition, Sterling Control Center Mobile enables you to access to a subset of Sterling 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 Sterling 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 Sterling 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 iPhone application that allows you to receive, view, add comments to, and handle Sterling Control Center alerts. In addition, you can view the status of servers and the status of Sterling B2B Integrator adapters. For more information, see *Mobile Application*.

The following table describes the functions available in the consoles, which is limited by the interface (console versus web console) and the user's role-based privileges.

Function	Console	Web Console	Mobile
Configuring			
Configure Sterling Connect:Direct servers with configuration management.	X		
Create and maintain users, data visibility groups, rules, actions, and service level criteria (SLCs), email lists, calendars, and schedules.	X		
Change system settings.	X		
Set up reports.	X		
Managing			
Stop Sterling Connect:Direct servers.	X		
Manage Sterling Connect:Direct processes by acting on queued processes (deleting, suspending, or releasing a process).	X		
Launch the Sterling Connect:Direct Browser User Interface to access Sterling Connect:Direct, the Sterling B2B Integrator Dashboard to access Sterling B2B Integrator, and the Sterling File Gateway Console to access Sterling File Gateway.	X		
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 allowed).	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 from the Queued Activity Monitor and the Completed Activity Monitor.	X	X	
Check status from the Server Status Monitor, Daemon Status Monitor, and Adapter Status Monitor.	X	X	X

Function	Console	Web Console	Mobile
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	
Create reports.	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 in many dialog boxes displays valid parameter values. 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.

For more information, see *Sterling Control Center console* and *Web console* in *User*.

Engine

The engine powers Sterling Control Center and enables you to manage and monitor multiple servers on different platforms.

It uses services, including node services, that handle the acquisition of data from the servers that are 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 monitored and managed. Node services are responsible for the communications that transpire between monitored servers and Sterling Control Center.

Sterling Control Center retrieves data from the following server resources:

Sterling Connect:Direct
- Select Statistics
- Select Processes

Sterling Connect:Enterprise

Information about files transferred into and out of mailboxes:
- Remote Connect batches
- Autoconnect batches
- Offline batches

Sterling B2B Integrator
- File transfer activities
- Business process activities

Sterling File Gateway
- Arrived file events

- Delivery events
- Route events

Sterling Connect:Express
- Queued Process Data
- Select Statistics

QuickFile

Messages with event information queued on the QuickFile server

MQ MFT

Information about files transferred to or from MQ MFT Agents

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 before 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 are processed by the metadata rule service. Events handled by the rule service trigger rules and their associated actions to be taken. Actions that can be taken include generating an email, 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 generates 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 three databases to record, store, and report information: the production, staging, and reports databases.

Production database

Sterling Control Center records the information gathered from the monitored servers for historical purposes in the production database (for example, select statistics and user reports). As information is received from monitored servers, events are generated and passed through the rule, metadata, and SLC services as they are being written to the database.

Staging database

The staging database is an optional database where Sterling Control Center stages older data.

Tip: If you use database partitioning for production database tables, a staging database is not necessary. However, if your database does not support partitioning, it is a best practice to use both production and staging databases.
The database administrator can export this data to long-term storage before it is purged. If you use database partitioning, Sterling Control Center does not move data to the staging database.

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.

Reports database

To generate Sterling Control Center reports, Cognos® Business Intelligence server is used with the production and reports databases to access and report information about monitored servers.

Database partitioning

If your database supports it, Sterling Control Center can be set up to use database partitioning. This process allows the data in the production database to be partitioned by date. Used as a best practice, database partitioning 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 IBM Sterling Control Center Database Partitioning.

Chapter 4. Define your objectives

A vital part of your Sterling Control Center implementation is the process of identifying your business issues and translating them into business objectives.

The business objectives you have for Sterling Control Center are the beginning of your Sterling Control Center planning process. A successful implementation depends on a comprehensive, multifaceted plan that you formulate before you configure the building blocks that tell Sterling Control Center the work you want it to do. You need to consider the following types of questions when identifying your Sterling Control Center objectives:

General

- What do I have in my environment?
  - What type of servers?
  - How many servers?
  - What function do the servers have?
  - Are my servers at the minimum maintenance level required to be monitored by Sterling 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 Sterling Control Center?
  - What is the status of my adapters?
  - What is the health of my Sterling B2B 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 large?
- Overall, what actions do you want to take?
  – Send email
  – Send SNMP trap
  – Send email 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 Sterling Connect:Direct servers and Sterling Connect:Direct Secure Plus on remote Sterling Connect:Direct servers?
- What changed?
- Who changed it?
- Can I easily add a server to multiple Sterling Connect:Direct server netmaps?
- Do I need notifications when my Sterling Connect:Direct server certificates used for Sterling Connect:Direct Secure Plus are going to expire?

Chapter 5. Define the work

After you identify your Sterling Control Center business objectives, begin the process of defining how Sterling Control Center will perform the work necessary to meet those objectives.

You then configure the Sterling Control Center building blocks, or objects, to 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
- Email lists you can associate with rules/actions to notify groups of individuals when an event occurs
- Rules that specify server events that cause actions to be taken, and when those rules are 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 *Create multiple objects* in *System Administration.*

Build the foundation

To build the foundation necessary for defining the work Sterling Control Center does, you need to plan for and define items such as servers, server groups, data visibility groups, roles/users, calendars/schedules, and email 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:
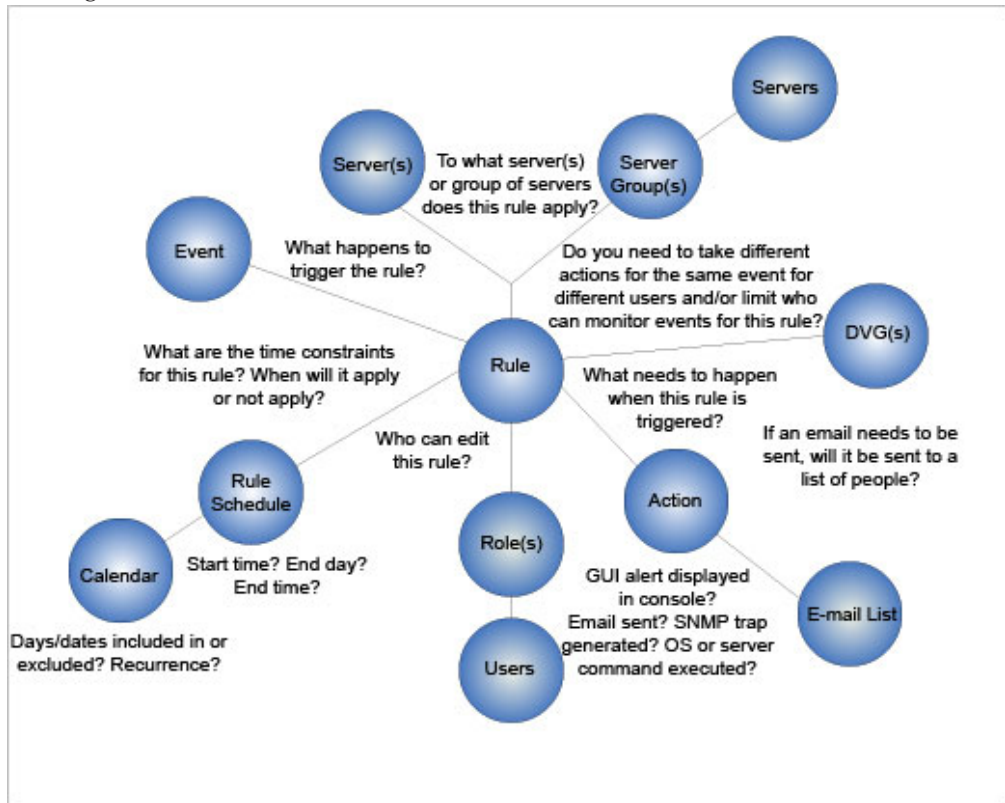
In this example, the building blocks of the rule include:

- Servers and/or Server Groups - contains which servers or groups of servers the rule applies to.
- Event: describes what happens to trigger the rule.
- Calendar - contains the included or excluded days or dates of the rule and whether they recur.
- Rule schedule - describes the start time, end time, and end day of the rule.

  The calendar and rule schedule are paired together to determine the time constraints for the rule. In other words, the calendar and rule schedule describe when the rule applies or does not apply.
- Users - contains the user information for the role or roles.
- Role or role - describes who can edit the rule.
- Email list - contains the list of addresses for email notifications.
- Action - describes what happens when the rule is triggered. Sample actions include a GUI alert displayed in the console, a sent email, a generated SNMP trap, or an executed OS or server command.

  The action is paired with the email list to describe what happens when the rule is triggered.
- Data visibility group (DVG) - describes different actions taken for the same event by different users and or limits who can monitor events for the rule.

Define servers

When you are planning your Sterling Control Center implementation, you need to identify the number and type of servers that need to be monitored, managed, or both.

Whether you are planning to use Sterling Control Center to monitor servers, manage assets or license keys, or configure Sterling Connect:Direct servers, you need to supply connection information for those servers. Gather information about the servers (based on server type) so that Sterling Control Center can access the server resources necessary for Sterling Control Center functions.

Tip: 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 and time zone settings (to accommodate servers in different geographic settings).

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 navigation tree by server type. To access server 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 provide visual indicators of server status. For a complete listing of the server icons, see *Sterling Control Center Console* in *Getting Started*.

Define server groups

A server group is a user-defined grouping of servers. As part of your planning process, decide whether and how you will group servers.

For example, you can group servers by processing center or by server type or both. 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 that is meaningful in your environment. For example, you might group servers according to user access to them, or to the rules, SLCs, or both that are applied to those servers. Use the Server Groups Worksheet to capture information about server groups you want to set up for your environment.

Important: Sterling Control Center has default system server groups that contain all of the servers of a particular type, for example, Sterling Connect:Direct for Microsoft Windows. The system server groups are strictly controlled by Sterling Control Center. Anytime you add or delete a server of a particular type, it is automatically added to or deleted from the applicable system server group. You can change the names of system server groups; however, system server groups cannot have the same names as user-defined server groups. System server group names must be changed before you define objects that reference them. For more information, see *Changing system server group names* in *Getting Started*.

When you are creating rules and SLCs that pertain to server events, specify the server or server group 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 configuring Sterling Control Center, you can define rules, SLCs that apply to a group of servers. This process is a more efficient means of defining rules and SLCs than defining separate rules, SLCs, or both for individual servers. This process helps reduce the number of objects you have to define and maintain.

For example, when you add a server to an existing server group, all the rules and SLCs defined for that server group are applied to the added server. If you do not implement server groups, you will have to set up 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 Groups tab. To access the properties for a server group, double-click the server group. To review and edit the properties of individual servers in the group, 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 server group properties by right-clicking the group and selecting the Properties option.

Perform guided node discovery

To get a list of the Sterling Connect:Direct nodes a monitored server communicates with you can find those nodes 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. Sterling Control Center contacts each server to gather information. Then, Sterling Control Center scans the server 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 Sterling Connect:Direct servers to monitor and manage in a large-scale server environment. Your Sterling Control Center license limits the number of managed servers allowed. If your Sterling Control Center license permits it, you can also add the discovered servers to the list of managed servers. For more information about this feature, see *Perform guided node discovery* in *System Administration*.

Define user access

As part of the planning process, you make decisions about user access to Sterling Control Center and the data it collects.

Consider the following when determining user access:
- The types of users (roles) who can access Sterling Control Center
- The permissions those roles have when configuring and managing the servers Sterling Control Center monitors (manage, view only, or none)
- Whether you plan to limit what events (data) specific users can monitor
- Whether you 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 define. When you create objects such as calendars, rules, and SLCs, you can specify the roles allowed to modify those objects.

Define roles

When you create users in Sterling Control Center, you give them credentials, or permissions, to access the system by 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 might 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 can 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. This means that users assigned this role can view all data collected and that they can 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 might be added with data visibility groups and "manage" permissions changed as needed to limit the permissions granted to those roles.

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

Some roles might require a mixture of permissions: they might need view-only permission for rules for a server or group of servers, but manage permissions for processes for those same servers. In other words, they cannot add or edit rules for those servers, but they can delete, resume, or suspend processes. This example illustrates 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, or business unit.

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

The following illustration shows a sample role hierarchy:

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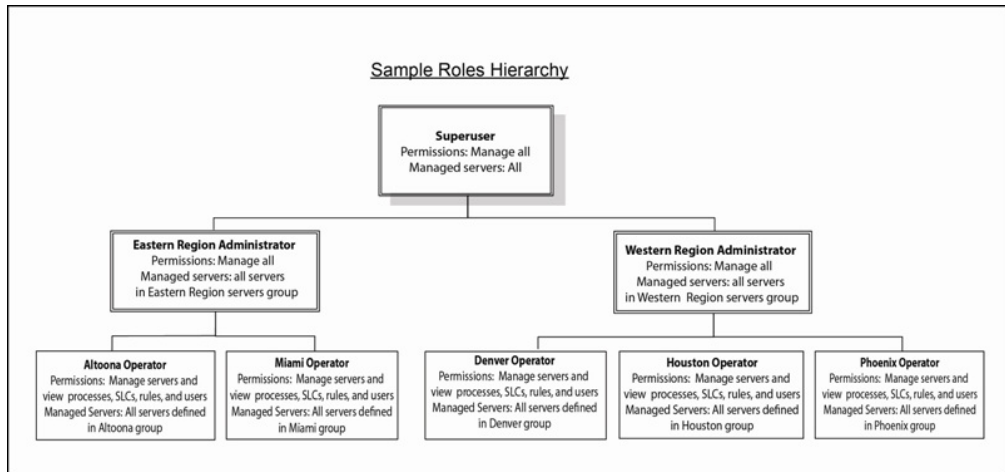
Subordinate roles cannot be given permissions higher than the permissions of a superior role. Also, subordinate roles can be given access only to the servers or server groups a superior role can access.

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

Understand 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 about data visibility groups, see *Understand Data Visibility Groups*.

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 is dimmed in the 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 *Manage Roles and Users* in *System Administration*.

Implement 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 password policy.

Password policies are set in Sterling Control Center by modifying the passwordPolicy.xml file in the *InstallationDirectory*\IBM\ControlCenter\conf\ security folder. If you implement a password policy, the policy criteria is enforced only for the following users:

- Existing users who change their password

- New users who are added after the passwordPolicy.xml file is modified and put into effect

If you do not edit the passwordPolicy.xml file, no password policy is in effect for Sterling Control Center. Password policy settings include the following criteria:
- 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

# Choose the best building blocks for the job: rules or 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.

Rules

Rules can stand on their own for processing events and acting when those events occur.

For example, when a Sterling B2B Integrator HTTP adapter is down, a Sterling Control Center event is generated. This event triggers a rule that starts an action that generates an alert and sends an email notification to an individual or list of individuals.

The basic process that occurs for rules is shown in the following figure:
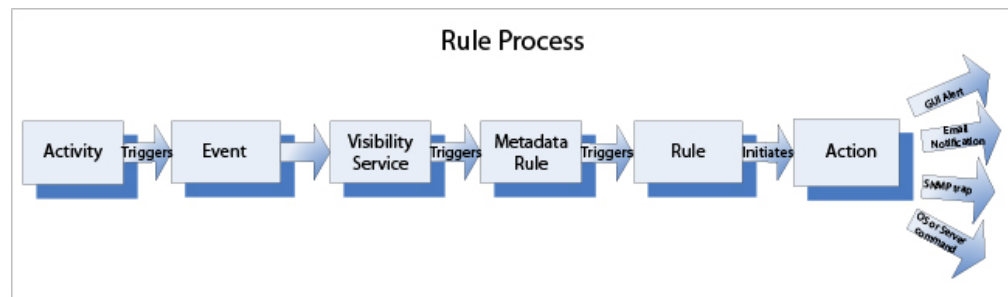
SLCs

Service level criteria (SLCs) are used to monitor for activities that either do or do not occur. Whereas, rules can act on only the events that do occur.

They monitor processes based on performance objectives: the process or process step/activity executes or fails to run within a certain time frame or for a specified duration. Processes can include Sterling Connect:Direct Processes, Sterling B2B Integrator business process activities, and the movement of data into and out of Sterling Connect:Enterprise mailboxes. When the conditions in an SLC occur or fail to occur, an SLC event is generated. You can set up rules based on SLC event message IDs that start actions, such as generating an alert. SLCs cannot 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 are sent, an outbound rule is triggered when transfers are completed successfully within the specified time frame. An email notification is sent.

When you are dealing with SLCs, the process differs as shown in the following figure:
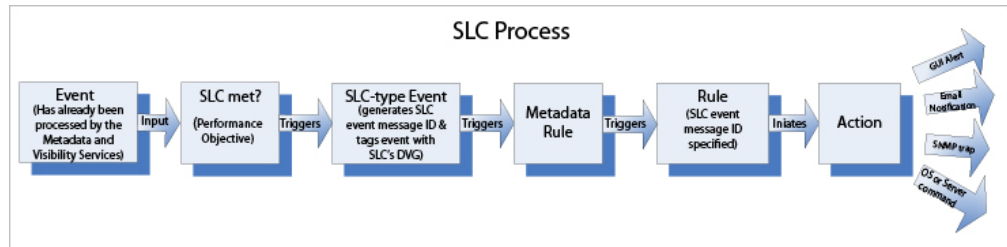
The SLC process begins with an event. This event is processed by the metadata and visibility services. The event is input for the SLC, or performance objective. The SLC triggers an SLC-type event which generates an SLC event message ID and tags the event with its DVG. The SLC-type event triggers a metadata rule which triggers a rule. This rule is specified by the SLC event message ID. The rule initiates an action, such as a GUI alert, email notification, SNMP trap, or OS or server command.

Situations for SLCs or rules

The situation dictates whether you use an SLC or a rule.

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
Process or Process Step failed		X
Process did not run	X	
Process ran on time	X	
File transfer took too long	X	
Process was not followed by another	X	

Data visibility groups overview

Whereas server groups can 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:
- 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.

In addition, each DVG may have its own set of rules and SLCs associated with it, allowing a single event to be processed by multiple rules and SLCs, resulting in different actions.

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.

Important: Only Sterling 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 can view only activity entities such as processes, statistics, alerts, and reports, that match their DVG restriction.

Rule sets

Sterling Control Center has 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 describes the differences between the types of rule sets:

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.

When an administrator views the **Rules** window, the **Rule Set** list displays GLOBAL and 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 and DVGs

When a DVG-restricted user creates a rule, the rule must specify a data visibility group in the rule parameters.

The Rule wizard populates the user's DVG in the rule 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 user who is not DVG restricted can create the following types of rules:
- A global rule
- A data visibility rule with any defined DVG

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

SLCs and DVGs

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 is 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 and DVGs

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.

The following figure shows how an event is tagged with DVGs:
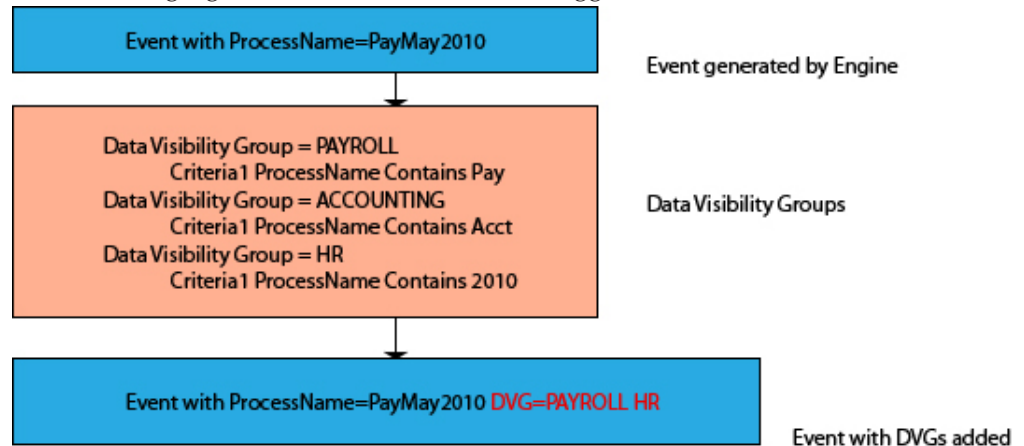
An event with a ProcessName of PayMay2010 is generated by the engine. The PAYROLL DVG Criteria1 ProcessName contains Pay which matches the event ProcessName. The ACCOUNTING DVG Criteria1 ProcessName contains Acct which does not match the event ProcessName. The HR DVG Criteria1 ProcessName contains 2010 which matches the event ProcessName. The PAYROLL and HR DVGs are added to the event.

Tip: 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 figure:
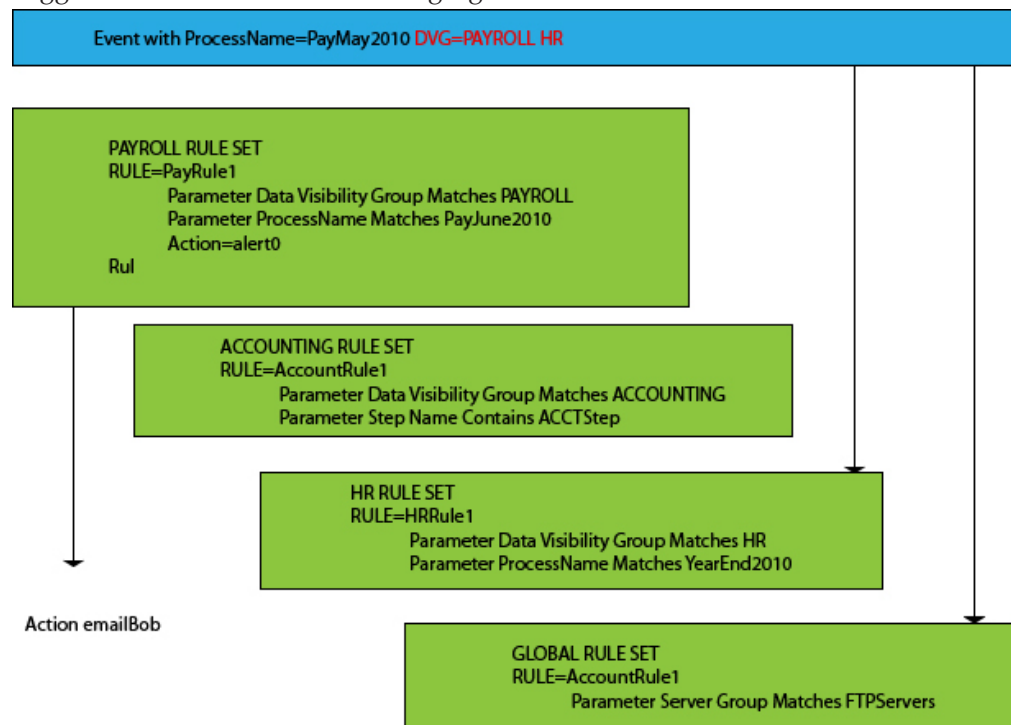
An event with a ProcessName of PayMay2101 and PAYROLL and HR DVGs is passed through the metadata service rules. The event is then passed to the

PAYROLL rule set. In this rule set, the first rule, PayRule1, contains the parameter DVG which matches PAYROLL. The parameter ProcessName matches PayJune2010. The action is alert0. In the second rule, PayRule2, the parameter DVG matches PAYROLL. The action is emailBob.

The event is passed to the ACCOUNTING rule set. In this rule set, the rule, AccountRule1, contains the parameter DVG, which matches ACCOUNTING. The parameter step name contains ACCTStep.

The event is passed to the HR rule set. In this rule set, the rule, HRRule1, contains the parameter DVG which matches HR. The parameter ProcessName matches YearEnd2010.

The event is passed to the GLOBAL rule set. In this rule set, the rule, AccountRule1, contains the parameter server group, which matches FTPServers.

DVGs and Sterling Control Center information

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

Monitors and DVGs

The events displayed on the **Completed Process Activity**, **Queued Process Activity**, **Active Alerts**, and **Handled Alerts Monitors** might be filtered by data visibility groups (DVGs).

The DVG values for an event can be displayed by 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**:

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. The user sees only the event tagged with Payroll DVG.

The following table describes the information displayed in the monitors based on whether 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. **Tip:** 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 and DVGs

When a DVG-restricted user runs any report, the report data 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 shows up on the report. Data Visibility Group is an option in the report filter so that DVG-restricted users assigned to multiple DVGs can limit the report to a certain DVG. Non DVG-restricted users can restrict the report to a specific DVG.

Rules and actions overview

Rules are triggered by events on servers monitored by Sterling Control Center. When a rule is triggered by a server event, the action specified in the rule is performed.

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

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

When rules are triggered by events, the action specified in the rule is performed, such as:

- Generate an alert and an email 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 return code is 8 or higher (for certain server types)

Rules have the following properties:

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 email notification
- Generating an SNMP trap
- Executing an operating system command on the system where the Sterling Control Center engine is running or executing 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 are met and a schedule associated with the rule matches. For more information, see *Calendars and schedules*.

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 these actions is taken depending on whether the second set of criteria is met within the time specified.

For example, a linked rule can 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 5 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 later 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 to the next rule, until a match occurs.

Therefore, rules with specific criteria should precede rules with more general criteria (specific server versus 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

Introduction to SLCs

You can set up service level criteria (SLCs) that help you monitor a process or file transfer based on performance objectives.

For example, a Sterling Connect:Direct process needs to begin by 20:00 and end by 20:30. An SLC might monitor for the timeliness of both events to answer questions such as, 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 events:
- Sterling Connect:Direct Process starts, ends, and durations
- Sterling Connect:Direct Process step starts, step ends, and durations
- Sterling B2B Integrator business process starts, ends, and durations
- Sterling B2B Integrator business process activity starts, ends, and durations
- Sterling File Gateway arrived file, route, and delivery starts, ends, and durations
- Sterling Connect:Enterprise batch arrivals and transmissions and durations
- Sterling Connect:Express Process starts, ends, and durations
- Sterling Connect:Express Process step starts, step ends, and durations
- QuickFile upload and download ends
- MQ MFT Process starts, ends, and durations
- MQ MFT Process step starts, step ends, and durations
- File Transfer Protocol (FTP) get and put ends

Parts of an SLC

There are multiple components that comprise service level criteria (SLCs).

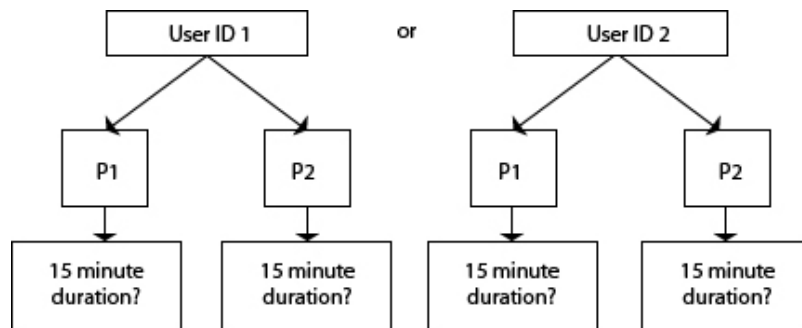SLCs contain elements as shown in following list:

Type

Tells what type of SLC it is. There are four types of SLCs:

- Standard - Standard SLCs monitor specific process names, file names, and other criteria. 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 can either specify multiple standard SLCs (one per batch ID that might be created), or you can 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 that consists of three processes. All three process must finish within three hours of the start of the first process.

  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 must not exceed a duration of 15 minutes, you can correlate them by specifying a correlator value of "submitterid." This correlation 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:

User ID 1	or	User ID 2
P1 / P2		P1 / P2
15 minute duration? / 15 minute duration?		15 minute duration? / 15 minute duration?

- Simple - With simple SLCs, you can 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.

  Note: Do not use simple SLCs for QuickFile.

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 about calendars and schedules, see *Calendars and Schedules*.

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 to detect an earlier start or a later end 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 start actions. For the SLCs you create, check to see whether the built-in rules meet your needs. If they do not, create rules that use SLC event message IDs to trigger the wanted action. For more information about rules and SLCs, see *Choose the Best Building Blocks for the Job—Rules vs. SLCs*.

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 exist only 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.

Metadata rules

You can use metadata rules 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 simplify the specification of rule and SLC criteria, they can also be used to simplify specification of report criteria.

Note: Metadata can be used to analyze only new activity. You cannot do retroactive analysis of existing data.

There are four metadata fields (USER_DATA_1 – USER_DATA_4). Sterling 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 contains 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

This scenario helps illustrate the use of metadata to simplify report generation for a large sporting goods company.

Objectives
* The company wants to notify a person in Accounting by sending email whenever a process owned by accounting fails.
* The company also wants to generate a Sterling Connect:Direct process Statistics Summary report about accounting processes.

Requirements
* There are five servers that are being monitored (SERVER1, SERVER2, SERVER 3, SERVER 4, and SERVER5).
* SERVER1 and SERVER5 are Sterling Connect:Direct servers that do only accounting work. So, all processes run on these two servers are accounting processes.
* SERVER2 does accounting and work for other departments. Processes submitted by Jane or Mary on SERVER2 are accounting processes.
* SERVER3 does accounting and work for other departments. Processes whose name begins with ACCT are accounting processes.
* SERVER4 does accounting and work for other departments. 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, metadata rules are more efficient than regular rules.

Regular rules

To meet the requirements without using metadata rules, the company writes four regular rules to take action on accounting processes that fail. The four regular rules are:
* If Server Id matches "SERVER1|SERVER5" then take Accounting Email Action
* If Server Id matches "SERVER2" and Submitter matches "Jane|Mary" then take Accounting Email Action

- If Server Id matches "SERVER3" and Process Name wildcard "ACCT*" then take Accounting Email Action
- If Server Id matches "SERVER4" and Remote Node matches "SERVER100" then take Accounting Email Action

Because Sterling Control Center does not have complex "AND/OR" logic for reports, a single report cannot be generated. So, the company writes four separate reports to get the accounting reports, when they only wanted one. And, the company puts logic similar to the rules logic 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, the company writes 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

The company writes one regular rule whose match criterion is USER_DATA_1 matches "ACCT" and whose action is to email a person in Accounting.

For the report, the company defines a single report whose match criterion is USER_DATA_1 matches "ACCT" and to 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 this example, USG could map USER_DATA_1 to "Department."

Permissible objects

Permissible objects are configuration objects to which restricted roles can be assigned. Permissible objects have both use and visibility (private or public) aspects controlled by user roles.

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 (public), or only to restricted users in the selected roles (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 use that action only 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 a **Permissions** window for an object, 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 or use the object. It affects only who can manage the object.
* This *<object type>* is visible to restricted users in these Selected Restricted Roles. The permissible object is private and can be viewed, edited, or used only by users in the restricted role or roles selected and by unrestricted users. When a permissible object is private and is 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 to place on the work that Sterling Control Center performs.

Days and dates are defined in calendars and times are defined in schedules. Calendars and schedules are defined independently for efficiency. They are defined independently because the same calendar (for example, Monday-Friday) can be paired with different schedules (08:00-09:00 and 16:00-17:00). After you 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 has eight predefined calendars. There is one calendar for each day of the week (Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday). Then, there is one calendar 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. In the planning process for your Sterling Control Center implementation, be sure to include an analysis of the calendars you needed for monitoring the servers in your environment. If you need common calendars for reuse, you can set them up before 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 to use, time zone for the calendar, and how long the calendar is in effect (start time, end time, and end date). When associated with a report, through 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 are matched against the rule 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 and dates the rule is in effect and times it is in effect.

For example, when you are defining a rule, you can set up a calendar schedule based on a daily calendar. Then, you can specify a schedule from 22:00-23:00 (the normal maintenance window for a set of servers). Next, you can 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 email 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 and 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).

20:00 - 6hr Window Start	20:00 NSRs	-	21:00 NSRe	22:00 NERs	-	23:00 NERe	23:00 + 6hr Window End
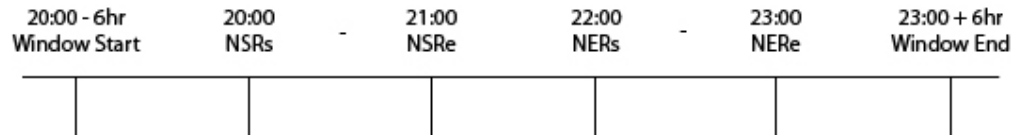
For SLCs with calendar schedules, Sterling Control Center can monitor activity for a specified number of hours before and after the schedule requirements you define (window start and window end in the previous example). The monitoring schedule can be set up wider than the schedule to detect a start that 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, end, or both), a **duration schedule** specifies that processing can begin at any time. However, when it begins, 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. How the count is handled depends 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 so. 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. Sterling Control Center information overview

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.

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 single or multiple servers or server groups, all managed servers of one type, or all managed servers.

Actions

When an event occurs, it triggers a rule. The rule initiates an action, such as displaying an alert in the consoles, sending an email, generating an SNMP trap, or initiating OS or server commands.

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.

Note: The Active Alerts monitor shows only the latest alert when more than one alert is triggered by the same SLC instance or linked rule. 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 move an alert to handled, you are required to add a comment about the update.

Emails

When an event occurs, an email 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 Sterling Connect:Direct server configuration objects. Can be run as an on-demand report or displayed on screen.

Reports

Used to gather information about the servers in your environment.

Standard Sterling Control Center Reports

Produced from the Sterling Control Center consoles on demand or automatically from schedules. Standard reports can grab any field in the database, and you can specify filtering criteria on report data. However, you cannot manipulate the format of standard reports or use them to perform complex queries, and they cannot be used to do calculations.

Sample Crystal Reports

Sample Crystal Reports are included with Sterling Control Center and are designed for use with databases supported by Sterling Control Center.

Monitor status

You can monitor the status of managed servers, Sterling B2B Integrator adapters, and Sterling Connect:Enterprise master daemons in your enterprise through the Sterling Control Center status monitors.

The status monitors include: 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 changes 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 can define a rule based on a server down condition with an action of sending an email notification status.

Monitor server status

The Server Status Monitor window provides a dynamic summary of managed server activity.

You can view the following types of server status:
- 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 Sterling Control Center name or alias
- Number of high, medium, and low severity alerts on the server
- Version of Sterling Connect:Direct, Sterling Connect:Enterprise, Sterling Connect:Express, QuickFile, FTP, or Sterling B2B Integrator software on a server
- License information (expiration details and concurrent number of sessions allowed or accounts defined)
- For Sterling Connect:Direct, Sterling Connect:Express, QuickFile, and Sterling B2B Integrator servers, information about maximum number of concurrent sessions on the server. You can also see the number of executing and non-executing processes for these server types.
- For MQ MFT servers, information about maximum number of concurrent transfers on the server.

Monitor adapter status

The **Adapter Status Monitor** displays summary information about Sterling B2B Integrator adapters that are running on Sterling B2B 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
- Display name for the adapter
- Nodes on which the adapter is deployed
- Type of adapter
- Sterling B2B Integrator perimeter server through which the adapter accesses the network
- State of the perimeter server

Sterling Control Center does not monitor Sterling B2B Integrator protocol adapters when those adapters are not actively monitoring business processes or protocols. For clustered Sterling B2B Integrator servers, status is monitored for all servers in a cluster. However, to avoid duplication, only unique adapter entries are displayed.

Monitor daemon status

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

The following types of information are displayed:
- Daemon name
- Type of daemon
- Host workstation the daemon is running on
- Daemon process identifier, originator, resource, and session identifier
- Whether the daemon is up or down

Monitor activity

You can view a configurable number of completed processes and queued processes on single or multiple servers or server groups, all managed servers of one type, or all managed servers through the **Process Activity Monitors**.

Important: Sterling Control Center does not "know" about FTP or Sterling Connect:Enterprise activity until it occurs. Therefore, the Queued Process Activity Monitor is disabled for those server types. In addition, held and deferred processes do not display in the Queued Process Activity Monitor for Sterling Connect:Express for z/OS.

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

For QuickFile, although Sterling Control Center can detect when a packageStart event occurs, which is generated prior to an upload or download, it does not generate events for uploads or downloads until they have been completed. As a result, although the Queued Process Activity Monitor is not disabled for QuickFile servers, you may not see any QuickFile activity in the monitor due to the timing of the events generated by Sterling Control Center.

From the Process Activity Monitors, you can also:
* Take a snapshot of process activity for queued and completed Sterling Connect:Direct, Sterling Connect:Express, QuickFile, MQ MFT, and Sterling B2B Integrator processes. You can 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 B2B Integrator business processes
* Act on a queued Sterling Connect:Direct and Sterling Connect:Express process, including deleting, suspending, or releasing it

Work 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). Consider the following scenario:
* 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 who is watching the Active Alerts monitor notices the alert, investigates the server, and brings the server back up.

- After the server is back up, the alert needs to be "handled" by specifying an appropriate comment, such as "brought the server back up."
- The user ID and the date and time when the alert was "handled" are recorded in the database along with the comment.
- The alert is moved to the Handled Alerts Monitor.

Options for generating reports

Sterling Control Center offers various reports that provide information about the servers in your environment.

Sterling Control Center includes 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 Sterling Connect:Direct configuration objects. They provide valuable information for troubleshooting installation problems and other support-related issues.

Standard 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 using email 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 • Initialization Parameters Report • Netmap Communication Paths Report • Netmap Modes Report • Netmap Nodes Report • Sterling Connect:Direct Secure Plus Cipher Suites Report • Sterling Connect:Direct Secure Plus Key Certificates Report • Sterling Connect:Direct Secure Plus Nodes Report • Sterling Connect:Direct Secure Plus Trusted Certificates Report • User Proxies Report • Versions Report

Report Type	Report Name
Monitoring	• Sterling Connect:Direct Process Statistics Details Report • Sterling Connect:Direct Process Statistics Summary • Sterling Connect:Direct Statistics Log Report • Sterling Connect:Enterprise Batch Statistics Summary • Sterling Connect:Enterprise Batch Statistics Details Report • Sterling Connect:Enterprise Statistics Log Report • Sterling Connect:Express Process Statistics Details Report • Sterling Connect:Express Process Statistics Summary • Sterling File Gateway Route Detail by Producer • Sterling File Gateway Route Detail by Consumer • Sterling B2B Integrator Business Process Details • Sterling B2B Integrator Business Process Summary • Sterling B2B Integrator File Transfer Report • QuickFile File Transfer Report • QuickFile Users Details Report • MQ MFT File Transfer Report • FTP File Transfer Report • High Watermark Report • Sterling Connect:Direct File Agent Process Submission Report
Node Discovery	• Netmap Connections Summary Report • Node Discovery Topology Report • Potentially Inactive Netmap Entries Report • Potentially Missing Netmap Entries Report
System	• Alerts Report • Audit Log Report • Database Events Report • Monthly File Transfer Activity Report • Server Inventory Report • Server Status Report • Service Level Criteria Summary Report • Users-Roles Summary 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 by using **Tools** > **Run SLC Debug Report**.

Sample reports

Sample reports in Crystal Reports format are included in Sterling Control Center. You can use these reports without modifications with Crystal Reports, or as templates for further customization.

The following sample reports are provided:
• Sterling Connect:Direct Events
• Sterling Connect:Direct Exception Trends
• Sterling Connect:Direct Exception Trends Chart

- Sterling Connect:Direct Usage Report
- Sterling Connect:Direct Usage Report Chart
- Sterling Connect:Direct Usage by Server Pair Report
- Sterling Connect:Direct Usage by Server Pair Report Chart
- Sterling Connect:Direct Usage by Server Pair Detail/Summary Report

These reports are designed for use with databases supported by Sterling Control Center. Schemas for the database tables used by Sterling Control Center are explained in *Reports*.

Chapter 7. Sample implementation scenario

This implementation scenario provides a glimpse into the planning process for a mid-sized bank in the United States. The scenario helps illustrate how the Sterling Control Center building blocks are used to provide the foundation for Sterling Control Center to meet the business objectives of the bank.

The bank plans to use Sterling Control Center to gain greater visibility into the condition of their Sterling Connect:Direct and Sterling B2B Integrator servers. They want to know their data flow status and information about their file transfers.

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

1. Limit user access and permissions to Sterling Control Center functions and data for the bank file transfer personnel
2. When servers or adapters are down, generate alerts and send email 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, they analyzed their monitoring goals to design their Sterling Control Center implementation. This process helped them design their implementation before they started configuring the building blocks that defined the work Sterling Control Center would perform. They determined that they needed the following building blocks to support their Sterling Control Center objectives.

Servers and server groups

The bank identified the Sterling Connect:Direct and Sterling B2B Integrator servers they need to monitor in their environment.

They also collected information about those servers by using the appropriate Sterling Control Center server planning worksheets for their server types. They used this information when they defined the following servers in Sterling Control Center:

Type of Server	Server Name
Sterling Connect:Direct	mauiCD
Sterling Connect:Direct	lagunaCD
Sterling Connect:Direct	destinCD
Sterling Connect:Direct	kauaiCD
Sterling Connect:Direct	lajollaCD
Sterling Connect:Direct	navarreCD
Sterling Connect:Direct	hawaiiCD
Sterling Connect:Direct	sunsetCD
Sterling Connect:Direct	daytonaCD

Type of Server	Server Name
Sterling B2B Integrator	capeSI
Sterling B2B Integrator	marthaSISFG

They also decided to group the servers. With server groups, user roles can be associated with particular server groups. This method makes email notifications easier when notifying users about events in regard to their assigned servers. They defined the following server groups:

Server Group	Server Name
CD1	mauiCD kauaiCD hawaiiCD
CD2	lagunaCD lajollaCD sunsetCD
CD3	destinCD navarreCD daytonaCD
SI	capeSI marthaSISFG

Objective 1: Limit user access

The bank identified roles to control access to Sterling Control Center.

The roles that the bank identified and developed is shown in the following list:

Sterling Control Center Admin

A super user who is responsible for installing, configuring, and maintaining Sterling Control Center, including startup and shutdown of the engine and defining subordinate administrator roles.

File Transfer Operations Admin

A subordinate role that has "manage" access to configure Sterling Connect:Direct nodes and Sterling B2B Integrator adapters, and can create SLCs, rules, and reports.

File Transfer Monitoring Staff

A subordinate role that has "view only" access for monitoring file transfers and cannot edit artifacts or objects. The role can view artifacts or objects. Uses in this role have view only access to objects necessary to their monitoring responsibilities. For example, users in this role do not need to view SLC or rule configuration.

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

Roles Hierarchy

SCCAdmin - Superuser
Permissions: All enabled for "manage"
Managed servers: All

Operations Admin CD1
Permissions: Manage all
Managed servers: all servers in
CD1 group

Operations Admin CD2
Permissions: Manage all
Managed servers: all servers in
CD2 group

Operations Admin CD3
Permissions: Manage all
Managed servers: all servers in
CD3 group

Operations Admin SI1
Permissions: Manage all
Managed servers: all servers in
SI group

Monitoring CD1
Permissions: View servers/groups,
processes, and alerts
Managed Servers: All servers
in CD1 group

Monitoring CD2
Permissions: View servers/groups,
processes, and alerts
Managed Servers: All servers in
CD2 group

Monitoring CD3
Permissions: View servers/groups,
processes, and alerts
Managed Servers: All servers in
CD3 group

Monitoring SI1
Permissions: View servers/groups,
processes, and alerts
Managed Servers: All servers in
SI group

The SCCAdmin - Superuser role has all permissions enabled to "manage" for all servers. The Operations Admin CD1 role has all permissions enabled for all servers in the CD1 group. The Monitoring CD1 role has permissions to view servers and groups, processes, and alerts for all servers in the CD1 group. There are similar Operations Admin and Monitoring roles for the CD2 and CD3 groups. The Operations Admin SI1 role has all permissions enabled for all servers in the SI group. The Monitoring SI1 role has permissions to view servers and groups, processes, and alerts for all servers in the SI group.

When the Sterling Control Center admin configured users who are allowed access to the Sterling Control Center console, an appropriate role was assigned to each user. As a result, when a user logs on to the console, they have access to the servers and functions associated with their role.

All of the Sterling Control Center console users run on a Microsoft Windows platform. The bank does not require Sterling Control Center to maintain any passwords in its user file and uses the signed on user and the Microsoft Windows domain as the credentials to allow signon to Sterling Control Center.

Objective 2: Server down

One of the monitoring objectives the bank identified for Sterling Control Center is the notification of personnel when a server down condition occurs.

To accomplish this objective, FCB defined a rule is triggered when a server is down. When this rule is triggered, an alert is generated in the **Active Alerts Monitor** and an email sent to notify personnel that a server is down.

Email list

Because the bank wants to send an email to multiple individuals when an alert occurs, they created an email list in Sterling Control Center. The list contains the email 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 email notifications, the bank also had to configure the Sterling Control Center engine System Setting value for email. This value specifies the location of the SMTP (email) server that Sterling Control Center uses to send email.

Filed	Value
Name	Monitoring staff
Description	List of all monitoring staff email addresses
To	Ops1@bank.com,Ops2@bank.com, Ops3@bank.com,Ops4@bank.com, Ops5@bank.com,Ops6@bank.com
Permissions	SCCAdmin

Server down action

When a server is down, the bank wants Sterling Control Center to generate an alert and send an email to the "Monitoring staff" email list.

To accomplish this objective, 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 (email list)
From	SCCAdmin@bank.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

The bank created a "Server down" rule for a monitored server.

The rule provides the following instructions for Sterling Control Center:

When a server down message id is detected for a monitored server, a high alert is generated in the **Active Alerts Monitor**. An email notification is sent to the monitoring staff.

(When the bank defined the rule, they selected the "Server down" action they defined earlier.)

Field	Value
Name	Server down
Description	This rule is triggered when a monitored server goes down.

Field	Value
Parameters	
Key	Message Id
Operator	Matches
Value	CCTR034E (Server is down - for monitored servers)
Action	Server down

A schedule was not required because the rule needs to be in effect at all times.

The following graphic shows the building blocks that comprise the "Server down" rule:
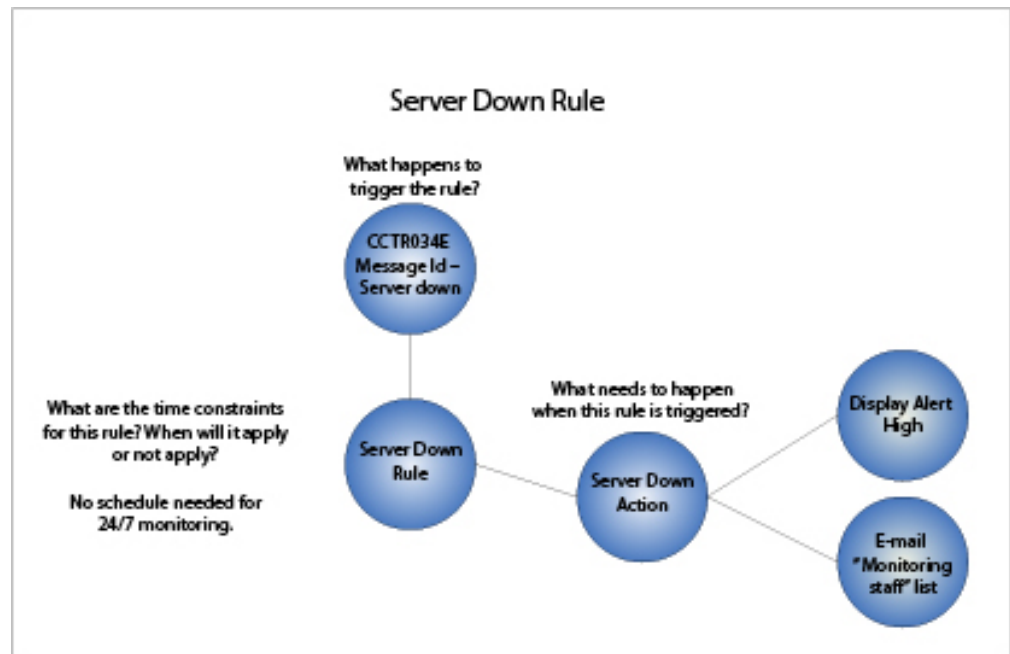
When the Message ID indicating that a server is down, CCTR034E, is detected, the "Server down" rule is triggered. Since there are no time constraints for this rule, no schedule is needed. When the rule is triggered, Sterling Control Center displays an alert and emails the "Monitoring staff" list.

Modify the server down rule

Monitoring personnel found that servers were often back up before they handled the alerts. As a result, the rule was modified to add a second condition (with a linked rule). With this change, if the server comes back up within 5 minutes, no action is taken.

If the server does not come back up, the "Server down" action is taken. The bank made the following changes (noted in bold) to the "Server down" rule to accomplish this objective:

Field	Value
Name	Server down

Field	Value
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)
Changes made by the bank	
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 the bank identified for Sterling Control Center is the notification of personnel when a process completes in error on a monitored server.

To accomplish this objective, they defined the rule that is triggered when an error occurs. When this rule is triggered, an alert is generated in the **Active Alerts Monitor**.

Process error rule

The bank created a "Process error" rule that provides the following instructions for Sterling Control Center:

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

They 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

Field	Value
Key	Return Code
Operator	Not Equal To
Value	0
Action	alert2 (predefined action shipped with Sterling Control Center)

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:
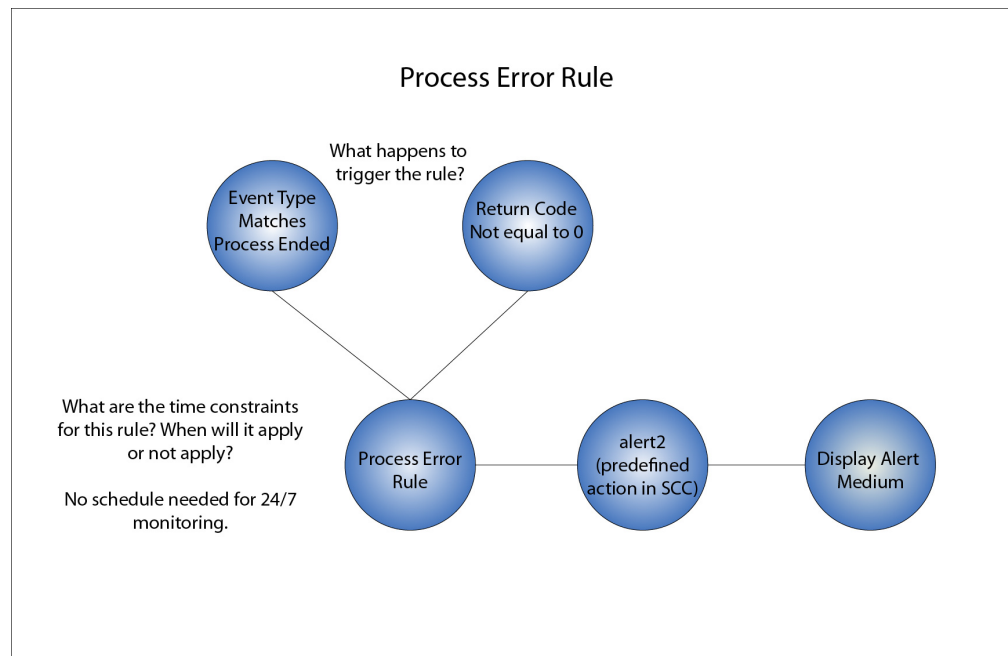
Followup objective: Specific process completes in error on a specific server

The Sterling Control Center administrator for the CD1 server group, needs to know when a specific process (DailyGrind) on a specific server (hawaiiCD) completes in error.

The CD1 administrator wants to generate a high alert and be notified by email when an error occurs. To accomplish this objective, the administrator created the following building blocks:

Daily Grind Error action

When the Daily Grind process encounters an error on the hawaiiCD server, the administrator for the CD1 server group wants Sterling Control Center to generate an alert and send an email. To accomplish this objective, the administrator 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	Ops1@bank.com
From	SCCAdmin@bank.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
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 from the values in the table:

Daily Grind Error Rule

What happens to trigger the rule?

Event Type Matches Process Ended

Process Name Matches DailyGrind

Servers and Server Groups Matches hawaiiCD

Return Code Not equal to 0

What are the time constraints for this rule? When will it apply or not apply?

No schedule needed for 24/7 monitoring.

Daily Grind Error Rule

Daily Grind Error Action

Email sent to CD1 Admin

Display Alert High

Objective 4: Process did not start at specified time

The bank wants to ensure that the EndOfDay process runs on the CD3 server group at 6:00 p.m. (Central Standard Time), or 18:00, each day. If it does not, they want to be notified.

To accomplish this objective, they first disabled all of the built-in SLC rules. This task was done to eliminate any actions other than the actions they would set up 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 email to the administrator of the 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 email to notify the CD3 server group administrator that the process started late

To accomplish this objective, they created the following building block:
- End of Day SLC calendar schedule
- End of Day SLC
- End of Day rules

End of Day SLC calendar schedule

Because the bank wants to monitor the start of the EndOfDay Process, they created an SLC calendar schedule.

The specified the following information in the calendar schedule. (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

To monitor for the start of a specific process on a specific server, the bank created a wildcard SLC.

They specified the following information for the SLC:

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	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), the bank disabled all of the built-in SLC rules and created two End of Day rules.

End of Day - Not Started rule

They also created an "End of Day - Not Started" rule that provides the following instructions for Sterling Control Center:

When a CSLC034E message (Process did not start by NSRe) is generated because the EndOfDay process does not start by 18:00 on any of the servers in the CD3 server group, a high alert is generated in the Active Alerts Monitor and an email sent to the administrator responsible for the CD3 server group.

They 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 did not start by 18:00.
Parameters	
Key	SLC Name
Operator	Matches
Value	End of Day
Key	Message Id
Operator	Matches
Value	CSLC034E
Action	End of Day - Not Started
Description	Action taken when the EndOfDay process did not start by 18:00 on any server in CD3 server group.
Email	
To	Ops3@bank.com
From	SCCAdmin@bank.com
Subject	&processName; not started
Message	&processName; did not start on &nodeName;
Alert Severity	1 - High
Permissions (Roles that can edit this action)	Operations Admin CD3

The following graphic shows the build blocks that comprise the End of Day - Not Started Rule:

End of Day – Not Started Rule
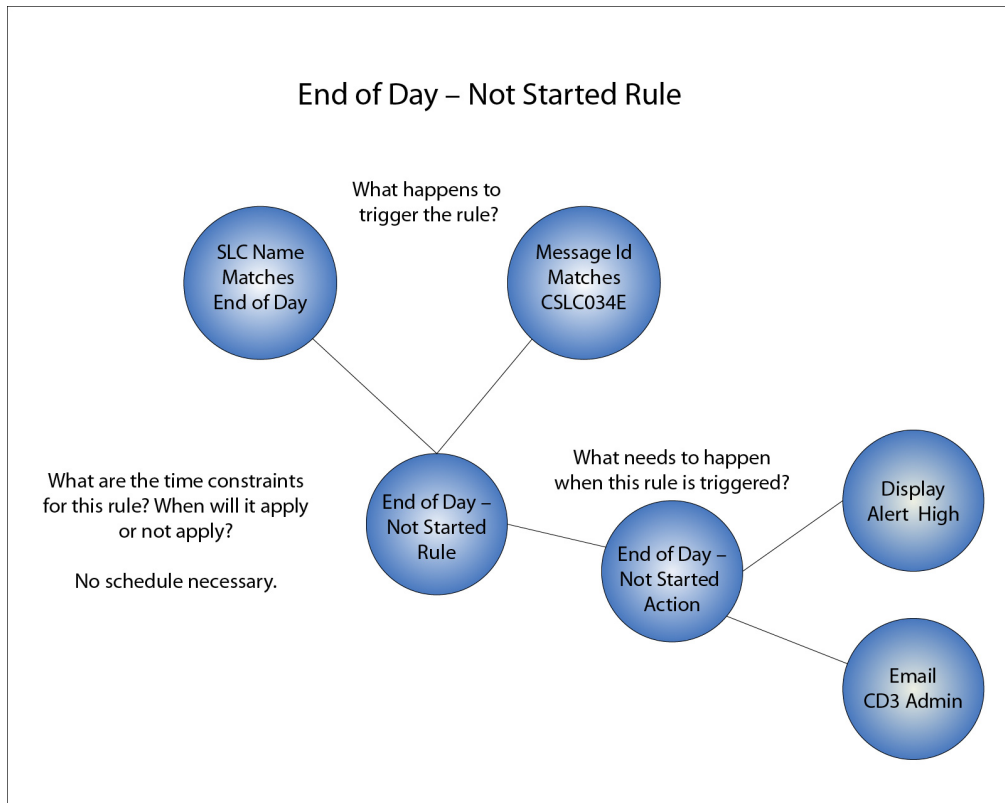
End of Day - Late Start rule

They bank also created an "End of Day - Late Start" rule that provides the following instructions for Sterling 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 CD3 server group, send an email to the administrator of the CD3 server group. No alert is generated, and all alerts generated for that SLC are 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 CD3 server group administrator 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 email when the EndOfDay process starts late in CD3 server group.
Email	
To	Ops3@bank.com
From	SCCAdmin@bank.com
Subject	&processName; late start
Message	&processName; 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 *How-To*.

The following graphics show the building blocks that comprise each of these rules:
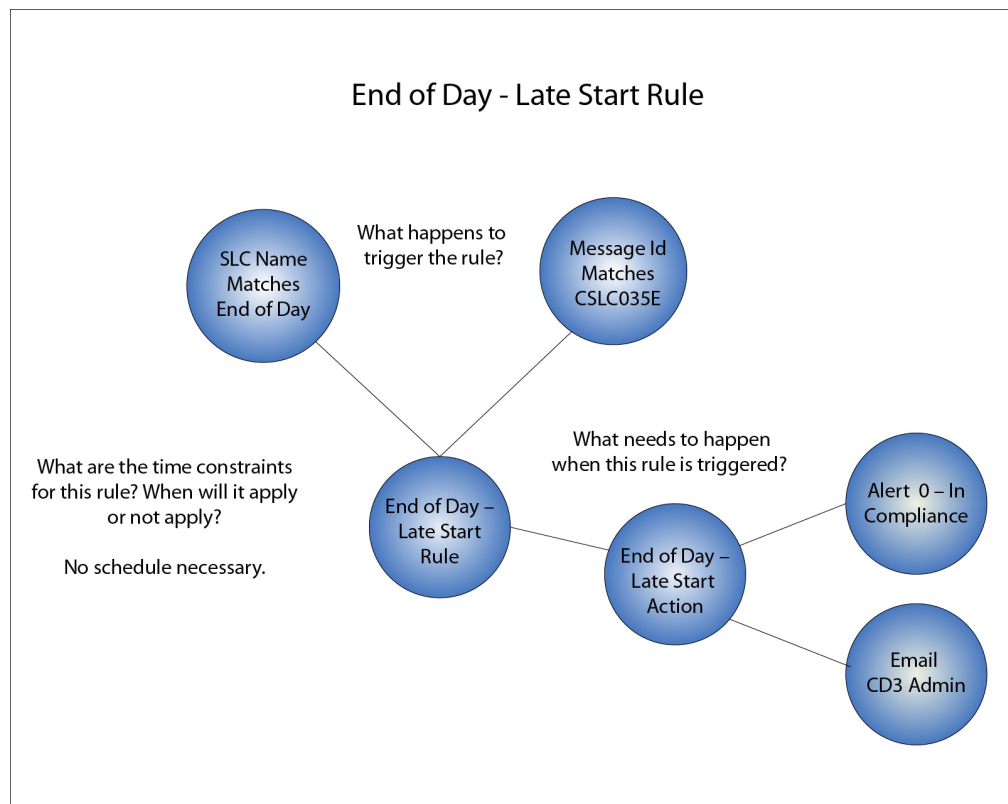
Chapter 8. Documentation and best practices

To more effectively implement Sterling Control Center, complete the tasks associated with that implementation in a certain order. Refer to the documentation resources for more information when planning your implementation.

The following information is provided to assist you with implementation tasks for Sterling 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

Product documentation provides you with the information you need to plan for, install, configure, operate, and maintain 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 Sterling Control Center. Read the document in its entirety before installation. For more information, see Release Notes.

Implementation

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 Sterling Control Center.

Use this information to get the big picture of the concepts, components, and building blocks that comprise Sterling Control Center. An understanding of this information will help you identify business objectives for Sterling Control Center and plan how Sterling 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 Sterling Control Center.

Planning Worksheets

Provide personnel responsible for implementing Sterling Control Center with worksheets they can use to plan their implementation and gather information used when configuring Sterling Control Center building blocks. For more information, see Planning Worksheets.

Getting Started

Provides installation and configuration information for personnel responsible for installing software and maintaining databases.

Use this information when you are planning your Sterling Control Center implementation, as well as installing and configuring Sterling Control Center and the database software that supports it.

System Administration

Provides programmers, network operations staff, and system administrators with the information they need to configure and maintain Sterling Control Center.

Use this information when you are configuring the building blocks that define the work Sterling Control Center will perform, maintaining Sterling Control Center, and troubleshooting issues.

User

Provides operations staff with the information they need to monitor server activity and oversee routine functioning of Sterling Control Center.

Use this information to help you access the wealth of information available through Sterling Control Center regarding the servers in your environment.

Reports

Provides programmers, network operations staff, and system administrators with the information they need to create and run Sterling Control Center reports.

Use this information to understand the types of reports available in Sterling Control Center and how to create and run those reports in Sterling Control Center, as well as working with third-party reporting tools to generate customized reports. This guide also defines Sterling Control Center's database table schemas.

How-To

Provides programmers, network operations staff, and system administrators with the answers to questions they may have about Sterling Control Center functions (arranged in a question and answer format).

Use this information when you are planning your Sterling Control Center implementation to gain more insight into how Sterling Control Center functionality can be used in very specific situations. As you work with Sterling 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 contents of *How-To*. There's a good chance your question is addressed in this helpful guide.

Configuration Management

Provides system administrators with the information they need to manage Sterling Connect:Direct server configurations from Sterling Control Center.

Use this information to maintain the following Sterling Connect:Direct configuration objects: functional authorities; initialization parameters; netmap nodes, modes, and communication paths; Sterling Connect:Direct Secure Plus nodes; and user proxies, and to track configuration object changes and versions.

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.

Database Partitioning article
> Provides information on using database partitioning for Sterling Control Center. For more information, see the IBM Sterling Control Center Database Partitioning paper.

Using External Directory Services for IBM Sterling Control Center User Authentication
> Provides information for using IBM Sterling External Authentication Server for validating Sterling Control Center users. For more information, see the Using External Directory Services for IBM Sterling Control Center User Authentication paper.

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 Sterling Control Center, so you would start with planning steps that identify your organization's Sterling Control Center objectives. The task list will help you identify the high-level tasks that need to be performed to implement Sterling Control Center.

#	Task	Planning Information and Best Practice Notes	Documentation Resources
1	**Planning Your Implementation**		
1a	Identify your business objectives	Identify your Sterling 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 Sterling Control Center implementation. Locate network diagrams, documents that state Sterling Control Center objectives (which may have been identified when your organization purchased Sterling Control Center), and other documents, that might help you identify your Sterling Control Center business objectives. Use the following planning aid: • *High-Level Business Objectives Worksheet*	Implementation • Concepts and components • Define your objectives • Information overview • Chapter 7, "Sample implementation scenario," on page 47 Planning Worksheets

#	Task	Planning Information and Best Practice Notes	Documentation Resources
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)*	System Requirements Getting Started • Before you install Sterling Control Center • Determining engine requirements • Database FAQ
1c	Gather information for the building blocks that will define Sterling Control Center system-level objects: servers, users, email lists, calendars, etc.	You will use this information later when you are creating these building blocks.	Planning Worksheets
	Gather server information	Gather information on the servers in your environment using the server worksheet for your server type(s). 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. For Sterling Connect:Direct servers, you can add one server and then use the Sterling 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 *Perform guided node discovery* in *System Administration*.	Implementation • Define servers Planning Worksheets System Administration • Manage servers
	Decide if you will rename system server groups	You must rename system server groups before you define any objects that will reference them. Use the following planning aid: • *System Server Groups Worksheet*	Planning Worksheets Release Notes Getting Started Implementation System Administration
	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 or service line. Use the following planning aids: • Server worksheet for server type(s) • *Server Groups Worksheet*	Implementation • Define servers Planning Worksheets System Administration • Manage servers

#	Task	Planning Information and Best Practice Notes	Documentation Resources
	Decide if you will use data visibility groups (DVGs)	Decide if you will use data visibility groups to limit the data users can monitor. If 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 • Understand data visibility groups • Define user access How-To Planning Worksheets System Administration • Manage data visibility groups • Manage roles and users
	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 • Define user access Planning Worksheets System Administration • 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 Sterling Control Center in a later task.	Implementation • Define user access Planning Worksheets System Administration • Manage roles and users
	Decide which users need to access Sterling 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 Sterling Control Center from the consoles and decide the permissions they need to be granted. Refer to the user permissions table in *Manage Roles and Users* in *System Administration* for a listing of the permissions users can be assigned. Use the following planning aid: • *User Access Worksheet*	Implementation • Define user access Planning Worksheets System Administration • Manage roles and users
	Decide what type(s) of notifications you will use to handle alerts	Gather information you will use to define SMTP email 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 Sterling 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 Sterling Control Center. It is named SterlingMIB.mib and is located in the esm installation directory folder.	Getting Started • Configure SMTP settings for email 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 Sterling Control Center will perform. 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 • Concepts and components • Define the work • Chapter 7, "Sample implementation scenario," on page 47 System Administration How-To
	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 Sterling 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 • Concepts and components • Define the work • Chapter 7, "Sample implementation scenario," on page 47 System Administration How-To • Setting up actions System Administration • Manage rules and actions
	Identify 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 • Concepts and components • Define the work • Chapter 7, "Sample implementation scenario," on page 47 How-To • Setting up SLCs • Troubleshooting SLCs System Administration • Manage Service Level Criteria
2	**Preparing the Environment**		
2a	Identify required patches, service packs, and releases for managed servers	Download any product updates from IBM Sterling Customer Support	System Requirements
2b	Verify that the system hardware is in place and ready for the installation		System Requirements

#	Task	Planning Information and Best Practice Notes	Documentation Resources
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.	Product documentation for the database you chose Getting Started • Before you install Sterling Control Center
2d	Create the production, staging, and reports 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 and the reports database, not the staging database.	Getting Started • Before you install Sterling Control Center
2e	Configure secure connections for use between Sterling Control Center and a Sterling Connect:Direct orSterling B2B Integrator server.	Before you configure the connections, obtain and import certificates into the Sterling Control Center truststore.	Getting Started • Configuring a Secure Connection
2f	If consoles will access the Sterling Control Center engine using a secure connection, configure the HTTPS connection		Getting Started • Determine HTTPS information for the engine and console connection
3	**Installing Sterling Control Center**		
3a	Install Sterling Control Center		Release Notes Getting Started • Before you install Sterling Control Center • Install the engine and console on UNIX or Installing the engine and console on Microsoft Windows
3c	After you install the engine, the Sterling Control Center console is available locally on the computer where Sterling Control Center is installed. For remote computers, the console is available via the engine's "launch" page.	The Console can be started by accessing the engine's "launch" page via a browser once the engine is started.	Getting Started • Working with the console

#	Task	Planning Information and Best Practice Notes	Documentation Resources
3d	If the console will access the engine using a secure connection, configure a secure connection between the engine and the console.		Getting Started • Configuring a secure connection
3e	Install the FTP agents (if you will manage FTP servers using Sterling Control Center)	One agent per FTP server must be installed (except for the z/OS FTP server).	Getting Started • Set up Sterling Control Center to monitor FTP agents
4	**Starting and Accessing Sterling Control Center**		
4a	Start Sterling Control Center		Getting Started • Start and stop Sterling Control Center
4b	Invoke the console to gain access to Sterling 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 • Working with the console
5	**Changing settings and tuning Sterling Control Center**	To ensure optimal operation, you can change system settings to tune Sterling Control Center. For example, if the engine time is different from the server time, adjustments need to be made in the time Preferences Settings. **Note:** To change SMTP email settings, run the configCC utility. If you change email settings, you must also configure Cognos Business Intelligence server using the configCC utility.	System Administration • Sterling Control Center Settings • Tuning Sterling Control Center Getting Started • Configuring SMTP settings for email messages
5a	Configure Sterling Control Center to send traps	Configure SNMP settings for host computers where SNMP traps are sent.	Getting Started • Configure SNMP Settings
5b	Configure JMS message retention time	Because Sterling Control Center is a durable subscriber for messages, it receives all messages published on a topic, including those published while it is not running. Configure the message provider to retain messages longer than the longest amount of time Sterling Control Center may be taken down.	IBM WebSphere MQ

#	Task	Planning Information and Best Practice Notes	Documentation Resources
5c	Change the default system server group names	If you want to change the default names for the system server groups, do it before you start defining servers, server groups, or any other Sterling Control Center objects. The engine must be restarted to put the changes into effect. **CAUTION:** **After you define objects, such as rules and SLCs, that reference system server groups, do not change system server group names.**	System Administration • System server groups • Changing system server group names • Resolving system server group naming conflicts
6	**Defining Servers**		
6a	Add servers	Use the Server Worksheet you completed earlier in the task list. For Sterling 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 *System Administration*. 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 • Manage servers • Perform guided node discovery Implementation Guide • Define the work Planning Worksheets
6b	Create server groups	Use the *Server Groups Worksheet* you completed.	System Administration • Manage servers Implementation • Define the work Planning Worksheets
7	**Defining User Access**		
7a	Define data visibility groups (DVGs)	Use the *Data Visibility Group (DVG) Worksheet* you completed.	System Administration • Manage data visibility groups Implementation Guide • Define the work Planning Worksheets
7b	Define user roles	Use the *User Access Worksheet* you completed.	System Administration • Manage roles and users Implementation • Define the work Planning Worksheets

#	Task	Planning Information and Best Practice Notes	Documentation Resources
7c	Define a user password policy	Use the *User Access Worksheet* you completed.	System Administration • Manage roles and users Implementation • Define the work Planning Worksheets
7d	Add users	Use the *User Access Worksheet* you completed.	System Administration • Manage roles and users Implementation Guide • Define the work Planning Worksheets
8	**Setting Up Calendars and Schedules**	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.	
8a	Set up calendars	Use *Business Objective Worksheet - Service Level Management* you completed earlier in the task list.	System Administration • Manage calendars and schedules Implementation • Define the work Planning Worksheets
8b	Set up schedules	Use *Business Objective Worksheet - Service Level Management* you completed earlier in the task list.	System Administration • Manage calendars and schedules Implementation • Define the work Planning Worksheets

#	Task	Planning Information and Best Practice Notes	Documentation Resources
9	**Create actions that will be invoked by a rule when an event occurs**	Will the default actions that are included in Sterling Control Center work for you? If not, create actions that will meet your needs. These actions will be used when you are creating rules. 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. 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 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 • Concepts and components • Define the work • Chapter 7, "Sample implementation scenario," on page 47 Planning Worksheets How-To • Setting Up Actions System Administration • Manage rules and actions
10	**Create rules that specify an event that Sterling 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* 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 • Concepts and components • Define the work • Chapter 7, "Sample implementation scenario," on page 47 Planning Worksheets How-To • Setting up rules • Troubleshooting rules System Administration • Manage rules and actions

#	Task	Planning Information and Best Practice Notes	Documentation Resources
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* 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 • Concepts and components • Define the work • Chapter 7, "Sample implementation scenario," on page 47 Planning Worksheets How-To • Setting up SLCs • Troubleshooting SLCs System Administration • Manage service level criteria
12	**Ongoing Administration of Sterling Control Center**		
	Changing Sterling Control Center Settings	In the ongoing administration of Sterling Control Center, you may need to change system settings, such as the following: • Database settings • Email settings • SNMP host settings • Engine connection settings • Console settings • License key settings • Console preferences **Note:** To change SMTP email settings, run the configCC utility. If you change email settings, you must also configure Cognos Business Intelligence server using the configCC utility. **Note:** If emails are sent in any rule actions, the email system setting must be configured first Likewise, the SNMP host settings must be configured appropriately for traps generated due to rule actions to be sent to the correct locations.	System Administration • 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.	Mobile Application Guide
	Improving Performance	You can perform general tuning to improve the performance of the Sterling Control Center engine.	System Administration • Tuning Sterling Control Center
	Updating licenses		System Administration

#	Task	Planning Information and Best Practice Notes	Documentation Resources
	Managing Configuration Objects	If you want to create, update, or delete multiple configuration objects such as actions, rules, schedules, email 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. You can also copy configuration objects between Sterling Control Center configurations to prepare for disaster recovery or to copy a production instance to a test instance.	System Administration • Create multiple objects • Copy configuration objects between Sterling Control Center installations • Create and maintain objects with the CCAPI
	Troubleshooting SCC		System Administration • Administrative troubleshooting

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing

IBM Corporation

North Castle Drive

Armonk, NY 10504-1785

U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing

Legal and Intellectual Property Law

IBM Japan Ltd.

19-21, Nihonbashi-Hakozakicho, Chuo-ku

Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be

incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

J46A/G4

555 Bailey Avenue

San Jose, CA 95141-1003

U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© IBM 2013. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2013.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium and the Ultrium Logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Connect Control Center®, Connect:Direct®, Connect:Enterprise, Gentran®, Gentran®:Basic®, Gentran:Control®, Gentran:Director®, Gentran:Plus®, Gentran:Realtime®, Gentran:Server®, Gentran:Viewpoint®, Sterling Commerce™, Sterling Information Broker®, and Sterling Integrator® are trademarks or registered trademarks of Sterling Commerce®, Inc., an IBM Company.

Other company, product, and service names may be trademarks or service marks of others.

Index

IBM ®

Product Number: 5725-D02

Printed in USA