



ScienceLogic Architecture

SL1 version 12.2.1

Table of Contents

Architecture Overview	1
SL1 Configurations	2
SL1 Appliances and Nodes	2
SL1 Appliance Functions	3
User Interface	3
Database Server	3
Data Collection	4
Message Collection	4
Compute	5
Load Balancer	5
Storage	5
Management	6
The SL1 Agent	6
Backup, Recovery & High Availability Options	7
Configuration Backups	8
Scheduled Backups to External File Systems	8
Backup of Disaster Recovery Database	8
Database Replication for Disaster Recovery	9
High Availability for Database Servers	9
Differences Between Disaster Recovery and High Availability for Database Servers	10
High Availability for Data Collection	10
Restrictions	11
All-In-One Architecture	13
Unsupported Configurations	15
The SL1 Agent	15
Distributed Architecture	16
Architecture Outline	17
Database Capacity	17
Data Collection	18
Message Collection	18
Interface Layer Requirements	18

The SL1 Agent	19
Extended Architecture	21
What is the SL1 Extended Architecture?	22
Architecture Outline	22
Database Capacity	23
Message Collection	24
Interface Layer Requirements	24
Compute Nodes	24
Load Balancer	25
Storage Nodes	25
Management Node	25
The SL1 Agent	26
The SL1 Agent	27
The SL1 Agent	28
SL1 Distributed Architecture with an SL1 Agent	29
SL1 Extended with an SL1 Agent	31
Database Layer Configurations	33
Local Storage	34
Single Database Server Using Local Disk Space	34
Database with Disaster Recovery	35
Clustered Databases with High Availability using Local Disk	36
Clustered Databases with High Availability Using Local Disk and Disaster Recovery	37
Interface Layer & API Configurations	39
Single Administration Portal	40
Multiple Administration Portals	41
Load-Balanced Administration Portals	42
API Layer Configurations	43
REST API	43
GraphQL API	43
Collector Group Configurations	44
Collector Groups	45
Traditional and PhoneHome Collectors	45

Using a Data Collector for Message Collection	47
Using Multiple Data Collectors in a Collector Group	47
How Collector Groups Handle Component Devices	48
High Availability for Data Collectors	49
Using Message Collectors in a Collector Group	50
Using Virtual Machines and Cloud Instances	54
Caveats	55

Chapter



1

Architecture Overview

Introduction

This manual describes the architecture of SL1 systems, covering most common configurations of SL1 appliances. This manual can help System Administrators and staff who are responsible for planning the architecture and configuration of SL1 systems.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

<i>SL1 Configurations</i>	2
<i>SL1 Appliances and Nodes</i>	2
<i>SL1 Appliance Functions</i>	3
User Interface	3
Database Server	3
Data Collection	4
Message Collection	4
Compute	5
Load Balancer	5
Storage	5
Management	6
<i>The SL1 Agent</i>	6

This chapter covers the following topics:

<i>SL1 Configurations</i>	2
---------------------------------	---

SL1 Appliances and Nodes	2
SL1 Appliance Functions	3
The SL1 Agent	6

SL1 Configurations

SL1 includes one or more **nodes** (parts of a cluster) or **appliances** that function together to provide the SL1 platform and application.

There are three primary SL1 configurations:

- **All-In-One.** In this configuration, a single node or appliance provides all the functions of SL1. The capacity of an All-In-One instance cannot be increased by adding additional nodes or appliances. This configuration is best for smaller deployments.
- **Distributed.** In this configuration, the functions of SL1 are divided between multiple nodes or appliances. A **Distributed** instance of SL1 can be as small as two nodes or appliances or include multiple instances of each node or appliance. This configuration is best for production environments that monitor a large number of devices or that monitor a large volume of data for each device.
- **Extended.** An extension of a Distributed instance. The **Extended** configuration adds both a **Compute Cluster** and a **Storage Cluster**. The Compute Cluster includes multiple Compute Nodes. The Storage Cluster includes multiple Storage Nodes. The **Extended** configuration also adds a Management Node to install and update the Compute Cluster and Storage Cluster, and one or more Load Balancers to manage the workload to the Compute Cluster. This configuration provides scale and can take advantage of the SL1 Agent to collect detailed data about devices and applications.

Resiliency and redundancy can also be accomplished by adding additional nodes or appliances to these configurations.

SL1 Appliances and Nodes

An instance of SL1 includes one or more of the following appliances or nodes:

- **All-In-One Appliance.** A single appliance that provides all the functions of SL1. The capacity of an All-In-One instance cannot be increased by adding additional appliances.
- **Database Server.** Contains a relational database used for all policy and configuration data. This database also stores performance data and log data.
- **Data Collector.** Data Collectors run multiple services, primarily for "agent-less" collection. Each Data Collector is responsible for collecting a specific set of information from a specific set of devices. Data Collectors can also receive instructions for asynchronous tasks including discovery, user-driven device tools, and automation actions for incident enrichment and remediation.
- **Message Collector.** Collects syslog and trap messages from devices. The Message Collector also communicates with the earliest version of the SL1 Agent and sends data collected by that agent to the Database Server.
- **Administration Portal.** Provides dedicated access to the SL1 user interface and API.

- **Computer Cluster.** A cluster of Compute Nodes. Each Compute Node includes Docker, Kubernetes and a range of SL1 services for features like data pipelines, Publisher, and expanded features in the SL1 Agent.
- **Load Balancer.** Provides access to the services running on the Compute Cluster.
- **Storage Cluster.** A cluster of Storage Nodes. Each Storage Node contains a NoSQL database that stores configuration data and storage data from some SL1 services.
- **Management Node.** Allows administrators to install and update packages on the Compute Cluster, Storage Cluster, and the Load Balancer and also update services on the Computer Cluster.
- **Platform Node.** Allows administrators to install and configure the Compute Cluster, Storage Cluster, Management Node, and Load Balancer. Administrators install generic Platform Nodes and transform the platform nodes into Compute Node, Storage Nodes, a Management Node, and one or more Load Balancers.
- **PowerFlow.** Enables bi-directional communication between the ScienceLogic data platform and external data platforms to promote a unified management ecosystem. The PowerFlow allows users to translate and share data between SL1 and other platforms without the need for programming knowledge.

SL1 Appliance Functions

In a **Distributed** system, there are four general functions that an SL1 appliance can perform: user interface, Database Server, Data Collector, and Message Collectors. In large SL1 systems, dedicated nodes or appliances perform each function. In smaller systems, some nodes or appliances perform multiple functions. In the **All-In-One Appliance** system, a single SL1 node or appliance performs all four functions.

User Interface

Administrators and users access the user interface through a web browser. In the user interface, you can view collected data and reports, define organizations and user accounts, define policies, view events, and create and view tickets, among other tasks. The node or appliance that provides the user interface also generates all scheduled reports and provides access to the ScienceLogic API. The following nodes or appliances provide the user interface:

- **All-In-One Appliance.** An **All-In-One Appliance** performs all functions, including providing the user interface.
- **Database Server.** A **Database Server** can provide the user interface in addition to its database function.
- **Administration Portal.** A dedicated **Administration Portal** node or appliance can provide the user interface.

NOTE: The Administration Portal communicates only with the Database Server and no other SL1 appliance. All connections between the Administration Portal and the Database Server are encrypted in both directions.

Database Server

The node or appliance that provides the database function is responsible for:

- Storing all configuration data and policy data.
- Storing performance data collected from managed devices.
- In a distributed system, pushing data to and retrieving data from the nodes or appliances responsible for collecting data and collecting messages.
- Processing and normalizing collected data.
- Allocating tasks to the other nodes or appliances in the SL1 System.
- Executing some automation actions in response to events.
- Sending all email generated by the system.
- Receiving all inbound email for events, ticketing, and round-trip email monitoring.

The following appliances can perform these database functions:

- **All-In-One Appliance.** An **All-In-One Appliance** performs all functions.
- **Database Server.** A dedicated **Database Server** provides all database functions.

Data Collection

Data Collectors are the SL1 nodes or appliances that retrieve data from monitored devices. In a distributed system, nodes or appliances that perform the data collection function also perform some pre-processing of collected data and execute automation actions.

The following appliances can perform the collection function:

- **All-In-One Appliance.** An **All-In-One Appliance** performs all functions.
- **Data Collector.** One or more Data Collectors are configured in **collector groups** for resilience. A collector group can be configured such that if an individual collector fails, other members of the group will pick up and share the load (N+1). A Data Collector can also perform the message collection function.

NOTE: The SL1 Agent can also be used to collect data from devices on which it can be installed. See the [System Requirements](#) page of the Support Site for a complete list of operating systems and versions supported by the agent. You can collect data from devices using only Data Collectors, using only the SL1 Agent, or using a combination of both.

Message Collection

The SL1 appliances that receive and process inbound, asynchronous syslog and trap messages from monitored devices.

The following nodes or appliances can perform the message collection function:

- **All-In-One Appliance.** An **All-In-One Appliance** performs all functions.
- **Message Collector.** A dedicated **Message Collector** receives and processes inbound, asynchronous syslog and trap messages from monitored devices.

- In distributed systems that use the SL1 agent, the Message Collector passes agent data to the Database server. On these distributed systems, the **Message Collector** must be a stand-alone node or appliance, not a combination **Data Collector/Message Collector**.

- **Data Collector**. A Data Collector can also perform the message collection function in addition to the data collection function.

SL1 Extended Architecture includes additional types of SL1 nodes or appliances. The following SL1 features require the SL1 Extended Architecture:

- **Expanded Agent Capabilities**. You can configure the SL1 Agent to communicate with SL1 via a dedicated Message Collector. However, this configuration limits the capabilities of the SL1 Agent. If you configure the SL1 Agent to communicate with SL1 via a Compute Cluster, you expand the capabilities of the SL1 Agent to include features like extensible collection and application monitoring.
- **Data Pipelines**. Data pipelines transport and transform data. Data transformations include enrichment with metadata, data rollup, and pattern-matching for alerting and automation. The Data Pipelines provide an alternative to the existing methods of data transport (data pull, config push, streamer, and communication via encrypted SQL) in SL1. Data pipelines introduce message queues and communicate using encrypted web services.
- **Publisher**. Publisher enables the egress of data from SL1. Publisher can provide data for long-term storage or provide input to other applications the perform analysis or reporting.
- **Scale-out storage of performance data**. Extended Architecture includes a non-SQL database (Scylla) for scalable storage of performance data.
- **Anomaly Detection and future AI/ML developments**. Anomaly detection is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. SL1 does this by collecting data for a particular metric over a period of time, learning the patterns of that particular device metric, and then choosing the best possible algorithm to analyze that data. Anomalies are detected when the actual collected data value falls outside the boundaries of the expected value range.

SL1 Extended Architecture includes the following additional SL1 nodes or appliances:

Compute

Compute nodes are the SL1 appliances that transport, process, and consume the data from Data Collectors and the SL1 Agent. SL1 uses Docker and Kubernetes to deploy and manage these services. T

Load Balancer

A load balance is the SL1 node or appliance that brokers communication with services running on the Compute Cluster. Services running on the Compute Cluster are managed by Kubernetes. Therefore, a single service could be running on one Compute node in the Compute Cluster; to provide scale, multiple instances of a single service could be running on one, many, or all nodes in the Compute Cluster. To provide scale and resiliency, you can include multiple Load Balancers in your configuration.

Storage

SL1 Extended includes a **Storage Cluster** that includes multiple Storage Nodes and a Storage Manager. These SL1 nodes or appliances provide a NoSQL alternative to the SL1 relational database. The Storage Cluster can

store performance and log data collected by the Data Collectors and the SL1 Agent.

Management

The **Management Node** allows administrators to install, configure, and update packages on the Compute Nodes cluster, Storage Nodes, and the Load Balancer. The Management Node also allows administrators to deploy and update services running on the Computer Cluster.

The SL1 Agent

The **SL1 agent** is a program that you can install on a device monitored by SL1. There is a Windows agent, an AIX agent, a Solaris agent, and a Linux agent. The agent collects data from the device and pushes that data back to SL1.

Similar to a Data Collector or Message Collector, the agent collects data about infrastructure and applications.

You can configure an agent to communicate with either the Message Collector or the Compute Cluster.

NOTE: The following minimum agent versions are required for SL1 12.1.1: **Windows** version 131; **Linux** version 174; **AIX** version 180; and **Solaris** version 180.

For more information about configurations with the SL1 Agent, see the chapter on the [SL1 Agent](#).

For more information about monitoring devices with the SL1 Agent, see the **Monitoring with the SL1 Agent** manual.

Chapter 2

Backup, Recovery & High Availability Options

Overview

SL1 has multiple options for backup, recovery, and high availability. Different appliance configurations support different options; your backup, recovery, and high availability requirements will guide the configuration of your SL1 System. This chapter describes each option and the requirements and exclusions for each option.

This table summarizes which of the options listed in this chapter are supported by:

- All-In-One Appliances
- Distributed systems that do not use a SAN (Storage Area Network) for storage
- Distributed systems that use a SAN (Storage Area Network) for storage

	Configuration Backup	Full Backup	HA Database	Disaster Recovery	HA Collection
All-In-One	✓	✓	✗	✗	✗
Distributed (local disk)	✓	✓	✓	✓	✓

This chapter covers the following topics:

Configuration Backups	8
Scheduled Backups to External File Systems	8
Backup of Disaster Recovery Database	8
Database Replication for Disaster Recovery	9
High Availability for Database Servers	9

<i>Differences Between Disaster Recovery and High Availability for Database Servers</i>	10
<i>High Availability for Data Collection</i>	10

Configuration Backups

SL1 allows you to configure a scheduled backup of all configuration data stored in the primary database. Configuration data includes scope and policy information, but does not include collected data, events, or logs.

SL1 can perform daily configuration backups while the database is running and does not suspend any ScienceLogic services.

SL1 can save local copies of the last seven days of configuration backups and stores the first configuration backup of the month for the current month and the first configuration backup of the month for the three previous months. Optionally, you can configure SL1 to either copy the daily configuration backup to an external file system using FTP, SFTP, NFS, or SMB or write the daily configuration backup directly to an external file system.

All configurations of SL1 support configuration backups.

The configuration backup process automatically ensures that the backup media is large enough. The configuration backup process calculates the sum of the size of all the tables to be backed up and then doubles that size; the resulting number is the required disk space for configuration backups. In almost all cases, the required space is less than 1 GB.

Scheduled Backups to External File Systems

A **full backup** creates a complete backup of the ScienceLogic database. Full backups use a built-in tool call MariaBackup.

Note the following information about full backups:

- SL1 can launch full backups automatically at the interval you specify.
- During a full backup, the ScienceLogic database remains online.
- If you have a large system and very large backup files, you can use an alternative method to perform backups that reduces performance issues during backup. For more information, see the section [Performing Config Backups and Full Backups on a Disaster Recovery Database Server](#).

Backup of Disaster Recovery Database

For SL1 systems configured for disaster recovery, you can backup the secondary Disaster Recovery database instead of backing up the primary Database Server. This backup option temporarily stops replication between the databases, performs a full backup of the secondary database, and then re-enables replication and performs a partial re-sync from the primary.

ScienceLogic recommends that you backup to an external file system when performing a DR backup.

- DR backup includes all configuration data, performance data, and log data.
- During DR backup, the primary Database Server remains online.

- DR backup is disabled by default. You can configure SL1 to automatically launch this backup at a frequency and time you specify.
- The backup is stored on an NFS mount or SMB mount.

NOTE: The *DR Backup* fields appear only for systems configured for Disaster Recovery. DR Backup is not available for the two-node High Availability cluster.

Database Replication for Disaster Recovery

You can configure SL1 to replicate data stored on a Database Server to a Disaster Recovery appliance with the same specifications. You can install the Disaster Recovery appliance at the same site as the primary Database Server (although this is not recommended) or at a different location.

If the primary Database Server fails for any reason, you must manually perform failover. Failover to the Disaster recovery appliance is not automated by SL1.

For details on configuring Disaster Recovery for Database Servers, see the manual *Disaster Recovery*.

High Availability for Database Servers

You can cluster Database Servers in the same location to allow for automatic failover.

A cluster includes an *active* Database Server and a *passive* Database Server. The passive Database Server provides redundancy and is dormant unless a failure occurs on the active Database Server. SL1 uses block-level replication to ensure that the data on each Database Server's primary file system is identical and that each Database Server is ready for failover if necessary. If the active Database Server fails, the passive Database Server automatically becomes active and performs all required database tasks. The previously passive Database Server remains active until another failure occurs.

Each database cluster uses a virtual IP address that is always associated with the primary Database Server. No reconfiguration of Administration Portals is required in the event of failover.

IMPORTANT: High Availability for Azure deployments is supported for installations of 12.1.x and later that are running on Oracle Linux 8 (OL8). ScienceLogic recommends that customers running SL1 versions prior to 12.1.x upgrade to 12.1.x or later and then complete the High Availability setup and configuration. For more information about upgrading, see the section on "Updating SL1" in the *System Administration* manual.

The following requirements must be met to cluster two Database Servers:

- The Database Servers must have the same hardware configuration.
- Two network paths must be configured between the two Database Servers. One of the network paths must be a direct connection between the Database Servers using a crossover cable.

NOTE: All-In-One Appliances cannot be configured in a cluster for high availability.

Differences Between Disaster Recovery and High Availability for Database Servers

SL1 provides two solutions that allow for failover to another Database Server if the primary Database Server fails: Disaster Recovery and High Availability. There are several differences between these two distinct features:

- **Location.** The primary and secondary databases in a High Availability configuration must be located together to configure the heartbeat network. In a Disaster Recovery configuration, the primary and secondary databases can be in different locations.
- **Failover.** In a High Availability configuration, SL1 performs failover automatically, although a manual failover option is available. In a Disaster Recovery configuration, failover must be performed manually.
- **System Operations.** A High Availability configuration maintains SL1 system operations if failure occurs on the hardware or software on the primary Database Server. A Disaster Recovery configuration maintains SL1 system operations if the data center where the primary Database Server is located has a major outage, provides a spare Database Server that can be quickly installed if the primary Database Server has a permanent hardware failure, and/or to allow for rotation of SL1 system operations between two data centers.

NOTE: A Distributed SL1 system can be configured for both High Availability and Disaster Recovery.

NOTE: High Availability and Disaster Recovery are *not supported* for All-In-One Appliances.

High Availability for Data Collection

In a Distributed SL1 system, the Data Collectors and Message Collectors are grouped into **Collector Groups**. A Distributed SL1 system can include one or more Collector Groups. The Data Collectors included in a Collector Group must have the same hardware configuration.

In SL1, each monitored device is aligned with a Collector Group and SL1 automatically determines which Data Collector in that collector group is responsible for collecting data from the monitored device. SL1 evenly distributes the devices monitored by a collector group across the Data Collectors in that collector group. Each monitored device can send syslog and trap messages to any of the Message Collectors in the collector group aligned with the monitored device.

To use a Data Collector for message collection, the Data Collector must be in a collector group that contains no other Data Collectors or Message Collectors.

If you require always-available data collection, you can configure a Collector Group to include redundancy. When a Collector Group is configured for high availability (that is, to include redundancy), if one of the Data

Collectors in the collector group fails, SL1 will automatically redistribute the devices from the failed Data Collector among the other Data Collectors in the Collector Group. Optionally, SL1 can automatically redistribute the devices again when the failed Data Collector is restored.

Each collector group that is configured for high availability includes a setting for Maximum Allowed Collector Outage. This setting specifies the number of Data Collectors that can fail and data collection will still continue as normal. If more Data Collectors than the specified maximum fail simultaneously, some or all monitored devices will not be monitored until the failed Data Collectors are restored.

High availability is configured per-Collector Group, so a SL1 system can have a mix of high availability and non-high availability collector groups, including non-high availability collector groups that contain a Data Collector that is also being used for message collection.

Restrictions

High availability for data collection cannot be used:

- In All-In-One Appliance systems.
- For Collector Groups that include a Data Collector that is being used for message collection.

For more information on the possible configurations for a Collector Group, see the [Collector Group Configurations](#) section.


Chapter

3

All-In-One Architecture

Overview

In an SL1 System that uses an All-In-One Appliance, a single node or appliance provides the user interface, database functions, performs data and message collection (including communication with the SL1 agent, if applicable), and provides API access.



This chapter covers the following topics:

<i>Unsupported Configurations</i>	15
<i>The SL1 Agent</i>	15

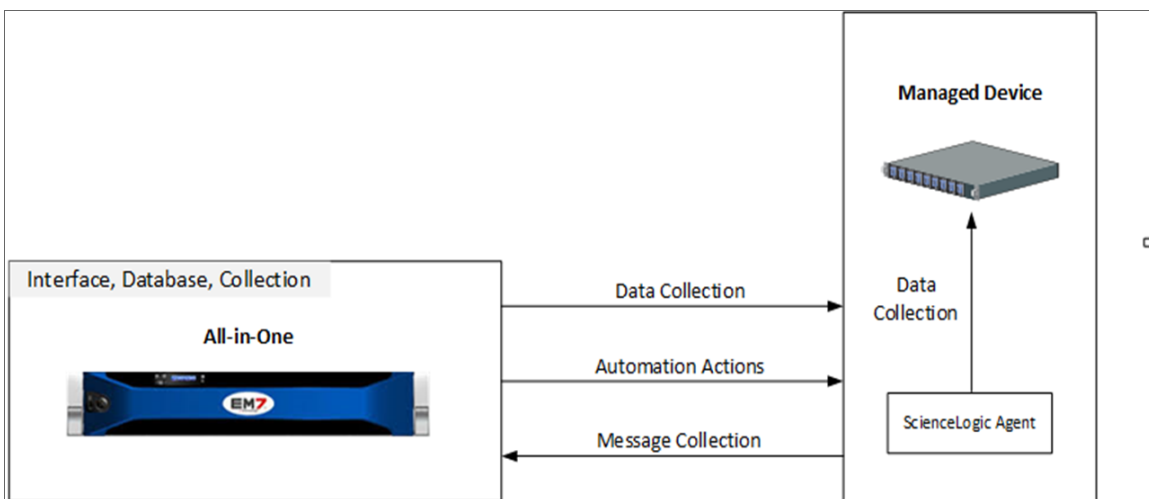
Unsupported Configurations

The following features are not supported by All-In-One Appliances:

- Using a SAN for storage
- Disaster Recovery
- High Availability for Database Servers
- High Availability for Data Collectors
- Additional Data Collectors, Message Collectors, or Administration Portals

The SL1 Agent

In an All-In-One architecture, the agent collects data from the device on which it is installed, and then sends messages to an All-In-One Appliance.



You can choose which All-In-One Appliance receives these messages from the agent.

For more information about configurations with the SL1 Agent, see the chapter on the [SL1 Agent](#).

For more information about monitoring devices with the agent, see the [Monitoring with the SL1 Agent](#) manual.

Chapter

4

Distributed Architecture



Overview

In a Distributed SL1 System, the functions of SL1 are divided between multiple appliances. The smallest Distributed SL1 System has two appliances:

- A **Database Server** that provides the user interface and performs all database functions.
- A **Data Collector** that performs data collection and message collection, including data collection from the SL1 agent, if installed.

Large SL1 systems can include multiple instances of each type of appliance. For a description of the appliance functions, see the [Overview](#) chapter.

Use the following menu options to navigate the SL1 user interface:

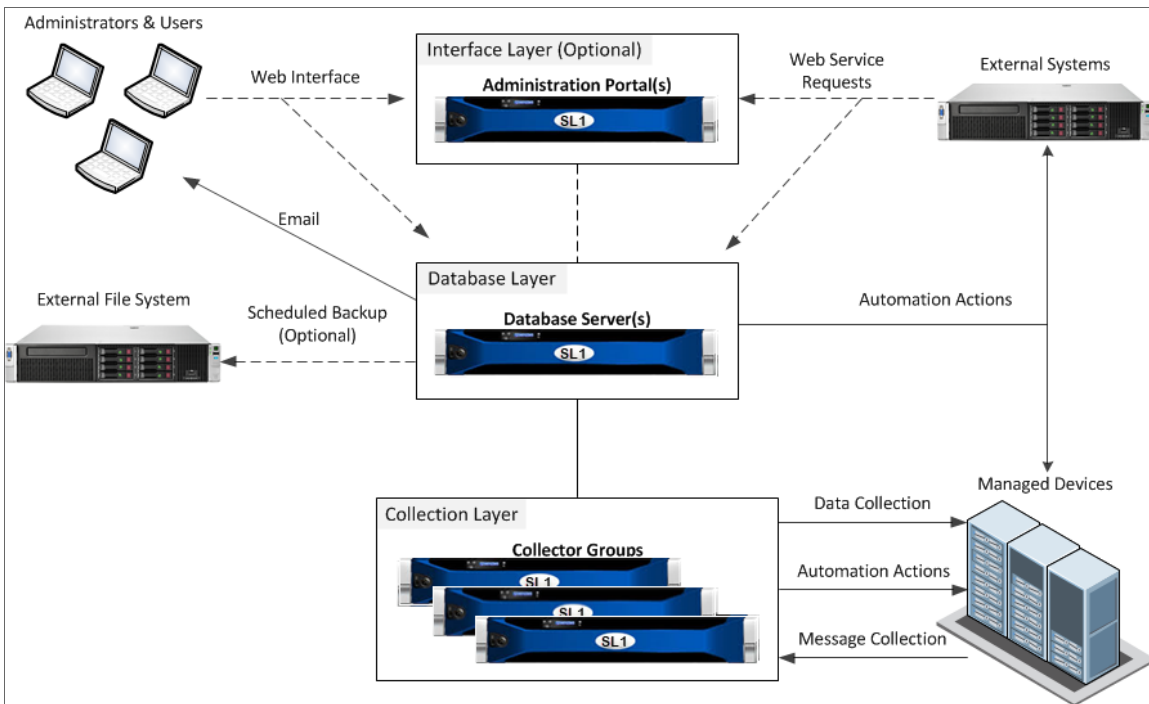
- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

Architecture Outline	17
Database Capacity	17
Data Collection	18
Message Collection	18
Interface Layer Requirements	18
The SL1 Agent	19

Architecture Outline

The general architecture of a distributed system includes two required layers, the database layer and the collection layer, and an optional layer, the interface layer.



The [Database Layer Configurations](#), [Interface Layer & API Configurations](#), and [Collector Group Configurations](#) chapters describe all the possible configurations for each layer.

Database Capacity

In a SL1 system, Database Servers and All-In-One Appliances are labeled with a Capacity. This capacity represents the total monitoring capacity of the SL1 system.

Capacity is a limit on the number of device records that can exist on an SL1 system and is defined by the license key issued to you by ScienceLogic. Each device record in a regular collector group consumes a capacity of one. The available capacity associated with your license is intended to provide adequate capacity for the baseline usage of your subscription along with room for some overage, while preventing a significant cost overrun or impacting the performance of your system.

NOTE: Virtual collector groups are *not* considered when calculating capacity.

For details on sizing and scaling your SL1 system to fit your workload, see [the ScienceLogic Support Site: https://support.sciencelogic.com/s/system-requirements](https://support.sciencelogic.com/s/system-requirements)

Data Collection

Data Collectors are arranged in collector groups for the purpose of scale and resilience. A collector group can be configured such that if an individual collector fails, other members of the group will pick up and share the load (N+1).

If a Data Collector loses connectivity with the Database Server, the collector will continue to collect data using its latest policy and will cache the results locally until collection is re-established.

The capacity of an individual collector varies in the range 300 – 6,000 end devices, depending on a range of factors such as depth of monitoring and latency of the target device. Complex video endpoints rare at the low end of the range, with basic availability monitoring at the high end of the scalability range. A collector with a typical mixed load will be able to monitor around 1,000 devices.

For details on sizing Data Collectors and Message Collectors, see <https://support.sciencelogic.com/s/system-requirements>.

Message Collection

When a Data Collector is used for message collection, the Data Collector can process approximately 20 syslog or trap messages per second.

When a Message Collector is used for message collection, the Message Collector can process approximately 100 to 300 syslog or trap messages per second. The number of syslog and trap messages that can be processed is dependent on the presence and configuration of syslog and trap event policies.

For details on sizing Data Collectors and Message Collectors, see <https://support.sciencelogic.com/s/system-requirements>.

Interface Layer Requirements

The interface layer can include one or more Administration Portals. The Administration Portal provides access to the user interface and also generates all scheduled reports. However, in some Distributed SL1 Systems, the interface layer is optional, and the Database Server can provide all functions of the Administration Portal.

If your Distributed SL1 System meets all of the following requirements, the interface layer is optional, and your Database Server can provide all functions of the Administration Portal. If your system does not meet all of the following requirements, the interface layer is required, and you must include at least one Administration Portal in your system:

- The SL1 system will have a low number of concurrent connections to the web interface.
- The SL1 system will have a low number of simultaneously running reports.

Precise requirements for concurrent connections and simultaneously running reports vary with usage patterns and report size. Typically, a dedicated Administration Portal is recommended for a SL1 System with more than fifty concurrent connections to the web interface or more than 10 scheduled reports per hour.

For details on sizing the Administration Portal, see <https://support.sciencelogic.com/s/system-requirements>.

The SL1 Agent

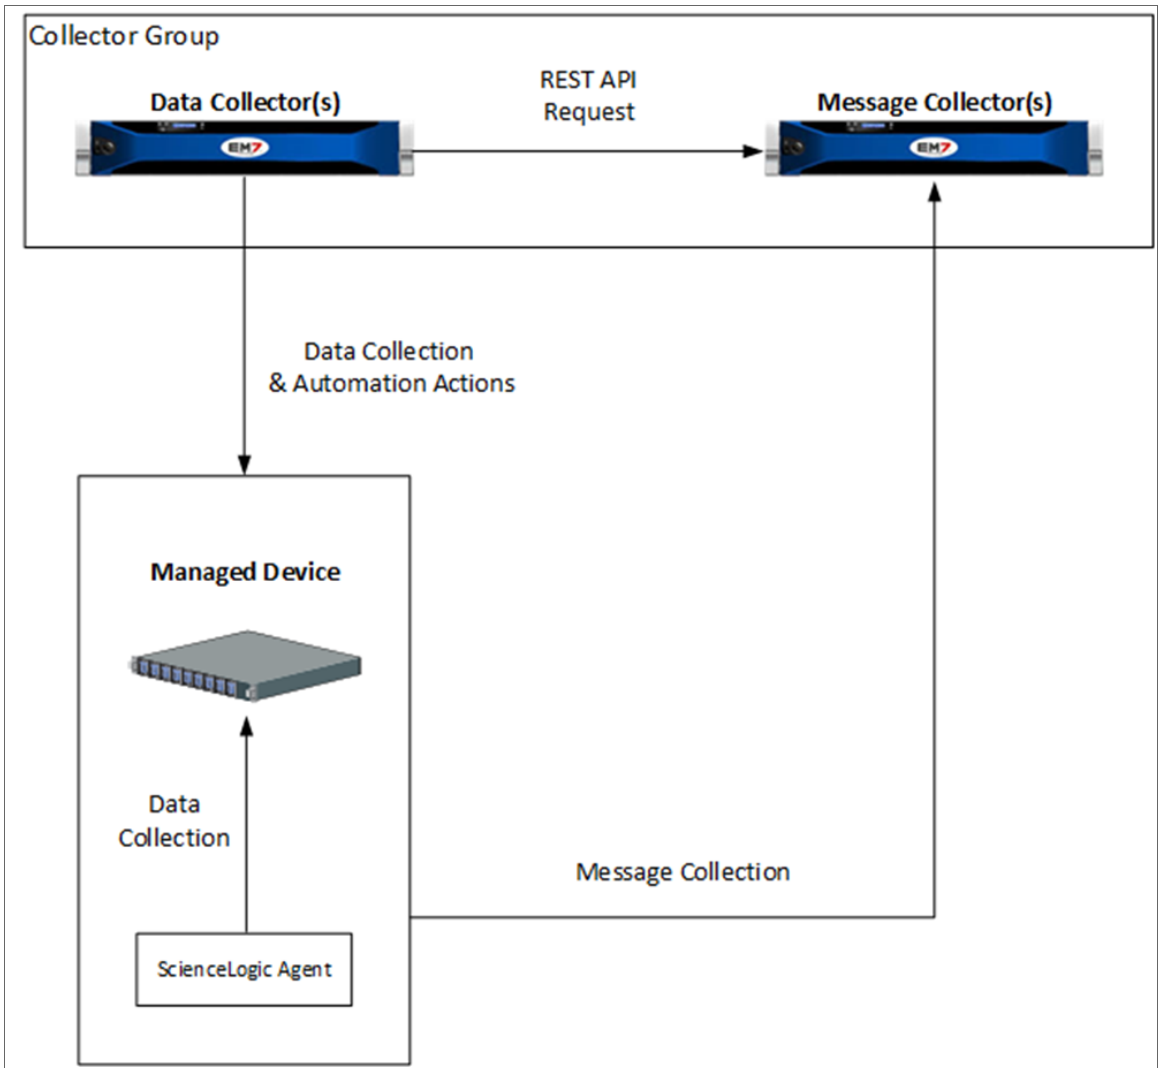
In an SL1 Distributed Architecture, the SL1 Agent collects data from the device on which it is installed and transfers that data to a Message Collector in an SL1 system using the HTTPS protocol. The Data Collector on which the Dynamic Applications and collection processes run then poll the Message Collector using the HTTPS protocol to transfer data to SL1.

TCP port 443 must be open between the Message Collector and the device on which an agent is installed.

In a Distributed Architecture, the SL1 agent requires a standalone, dedicated Message Collector. The Message Collector does not need to be dedicated to agent usage, but the Message Collector cannot be a Data Collector that also performs message collection.

NOTE: Message Collectors that process data from the agent have different system requirements than Message Collectors that do not process data from the agent. For more information about the system requirements when running agents in a Distributed Architecture, see the [System Requirements](#) page at the ScienceLogic Support Site.

The diagram below shows the collection layer of a Distributed System containing both Data Collectors and Message Collectors in which the SL1 Agent is installed on a managed device.



You can choose the Message Collector that receives these messages from the agent.

For more information about changing the Message Collector to which the Agent sends messages, see the **Monitoring with the SL1 Agent** manual.

Chapter



5

Extended Architecture

Overview

This chapter describes the SL1 Extended Architecture, which includes all of the nodes in SL1 Distributed plus additional nodes that provide scale and new functionality.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ()
- To view a page containing all of the menu options, click the Advanced menu icon ()

This chapter covers the following topics:

<i>What is the SL1 Extended Architecture?</i>	22
<i>Architecture Outline</i>	22
<i>Database Capacity</i>	23
<i>Message Collection</i>	24
<i>Interface Layer Requirements</i>	24
<i>Compute Nodes</i>	24
<i>Load Balancer</i>	25
<i>Storage Nodes</i>	25
<i>Management Node</i>	25
<i>The SL1 Agent</i>	26

What is the SL1 Extended Architecture?

SL1 Extended includes all the nodes in SL1 Distributed plus additional nodes that provide scale and new functionality.

SL1 Extended provides these new features:

- Infrastructure monitoring with the SL1 Agent
- Collectors connecting over https
- Scalable storage of performance data

In an Extended SL1 System, the functions of SL1 are divided between multiple appliances. The smallest Extended SL1 System has:

- An **Administration Portal** that provides the user interface and processes requests from REST API and GraphQL.
- A **Database Server** that stores SL1 system information and configuration data collected from managed devices.
- A **Data Collector** that performs data collection and message collection.
- An (optional) **Message Collector** that performs message collection.
- **Compute Cluster** (minimum of three appliances). Cluster of nodes that queues, transforms, and normalizes data collected by the SL1 Agent and the Data Collector appliance. More nodes can be added as needed.
- **Load Balancer**. Provides load-balancing and queues requests for the Compute Node cluster.
- **Storage Cluster** (minimum of three appliances). Scylla databases that store performance data collected by the SL1 Agent and the Data Collector appliance. More nodes can be added as needed.
- **Management Node**. Node that allows users to install and update packages on the Compute Node cluster, Storage Node cluster, and the Load Balancer.

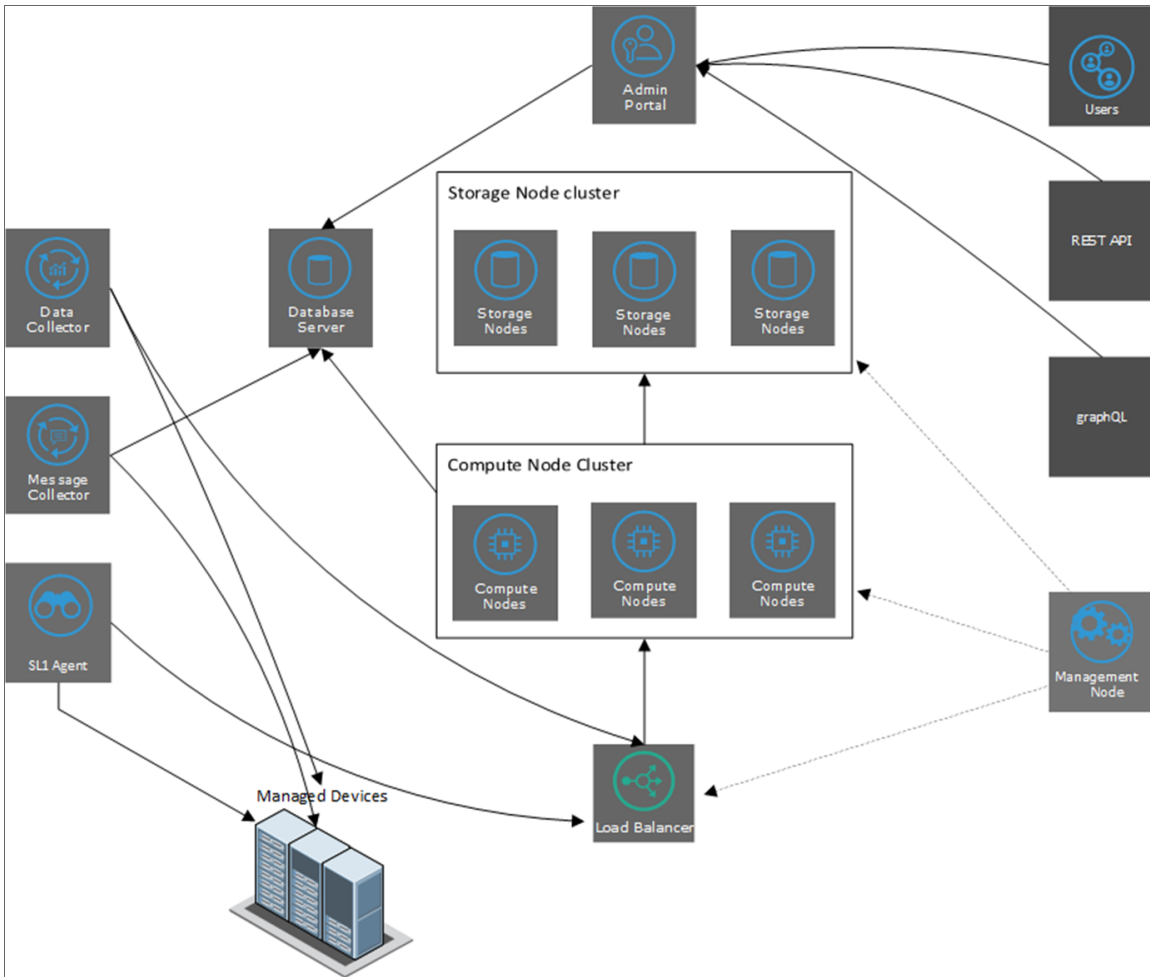
Large SL1 systems can include multiple instances of each type of appliance. For a description of the appliance functions, see the [Overview](#) chapter.

Architecture Outline

The general architecture of a distributed system includes:

- **database** layer that includes a Database Server that stores SL1 configuration data and collected configuration data.
- **collection** layer that includes a Data Collector, a Message Collector, and the SL1 Agent.
- (optional) **interface** layer that includes an Administration Portal.
- **Compute Node cluster** (minimum of three appliances) that processes incoming data from the SL1 Agent and the Data Collector. The compute node sends configuration data to the Database Server and performance data to the Storage Node cluster.

- **Load Balancer** that manages requests to the Compute Node Cluster.
- **Storage Node cluster** (minimum of three appliances) that includes Scylla databases and stores performance data.
- **Management Node** that provides updates and patches for the Compute Node cluster, the Storage Node cluster and the Load Balancer.



The [Database Layer Configurations](#), [Interface Layer & API Configurations](#), and [Collector Group Configurations](#) chapters describe all the possible configurations for each of the layers.

For details on sizing and scaling your SL1 system to fit your workload, contact ScienceLogic Support or your Account Manager.

Database Capacity

For details on sizing and scaling the Database Server in an SL1 Extended system, see <https://support.sciencelogic.com/s/system-requirements>.

Message Collection

When a Data Collector is used for message collection, the Data Collector can process approximately 20 syslog or trap messages per second.

When a Message Collector is used for message collection, the Message Collector can process approximately 100 to 300 syslog or trap messages per second. The number of syslog and trap messages that can be processed is dependent on the presence and configuration of syslog and trap event policies.

For details on sizing Data Collectors and Message Collectors, see <https://support.sciencelogic.com/s/system-requirements>.

Interface Layer Requirements

The interface layer can include one or more Administration Portals. The Administration Portal provides access to the user interface and also generates all scheduled reports. However, in some Distributed SL1 Systems, the interface layer is optional, and the Database Server can provide all functions of the Administration Portal.

If your Distributed SL1 System meets all of the following requirements, the interface layer is optional, and your Database Server can provide all functions of the Administration Portal. If your system does not meet all of the following requirements, the interface layer is required, and you must include at least one Administration Portal in your system:

- The SL1 system will have a low number of concurrent connections to the web interface.
- The SL1 system will have a low number of simultaneously running reports.

Precise requirements for concurrent connections and simultaneously running reports vary with usage patterns and report size. Typically, a dedicated Administration Portal is recommended for a SL1 System with more than fifty concurrent connections to the web interface or more than 10 scheduled reports per hour.

For details on sizing the Administration Portal, see <https://support.sciencelogic.com/s/system-requirements>.

Compute Nodes

SL1 Extended includes a Compute Cluster that includes a minimum of three Compute Nodes. Each Compute Node runs Kubernetes and includes services for ingestion and transformation of data collected by the ScienceLogic Agent and the Data Collector appliance.

The SL1 Agent and Data Collectors send data to an ingestion end point on the Compute Cluster. The Compute Cluster transforms the data using specific pipelines for the different types of collected data. Transformations include enrichment with metadata, data rollup, and pattern-matching for alerting and automation.

Compute nodes also run a graph database that stores relationships for applications, infrastructure, and networks.

In larger SL1 systems, you can add additional Compute Nodes to the Compute Cluster to extended overall system capacity. For details on sizing and scaling the Compute Node cluster in an SL1 Extended system, see <https://support.sciencelogic.com/s/system-requirements>.

When deployed in AWS:

- SL1 Extended uses EKS (Amazon Elastic Kubernetes Service) to provide the Kubernetes control plane.
- The cluster can be configured with nodes in three distinct availability zones within the same region, allowing continuous operation in the event of an availability zone loss.

Load Balancer

The Load Balancer is a single appliance or virtual machine that provides load-balancing for the Compute Node cluster. You can include two Load Balancers, for redundancy.

For details on sizing the Load Balancer, see <https://support.sciencelogic.com/s/system-requirements>.

Storage Nodes

SL1 Extended includes a Storage Cluster that includes a minimum of three Storage Nodes. These Storage Nodes store performance data collected by the SL1 Agent and the Data Collector appliance. The Storage Nodes contain NoSQL databases. More Storage Nodes can be added as needed.

After the Compute Node cluster has transformed the collected data, the data is written to either the Storage Node cluster (performance data) or the Database Server (configuration data).

The NoSQL database has concepts of “rack” and “datacenter” and can use these to provide fault tolerance and latency expectations for inter-data center replication.

When deployed in AWS:

- the cluster can be configured with nodes in three distinct availability zones within the same region, allowing continuous operation in the event of an availability zone loss.

NOTE: Uploads that occur in 20-second intervals, sometimes called "snapshot uploads", are no longer supported for users using the non-Scylla pipeline. These 20-second uploads were replaced with the one-minute upload default in SL1 version 11.2. ScienceLogic highly recommends that you ensure your agents are uploading in one-minute summarized uploads prior to upgrading. You can verify your uploads from the **[Settings]** tab in the current SL1 user interface. The agent pipeline is able to consume and summarize 1-minute and 5-minute payloads without the need for Scylla.

For details on sizing and scaling the Storage Node cluster in an SL1 Extended system, see <https://support.sciencelogic.com/s/system-requirements>.

Management Node

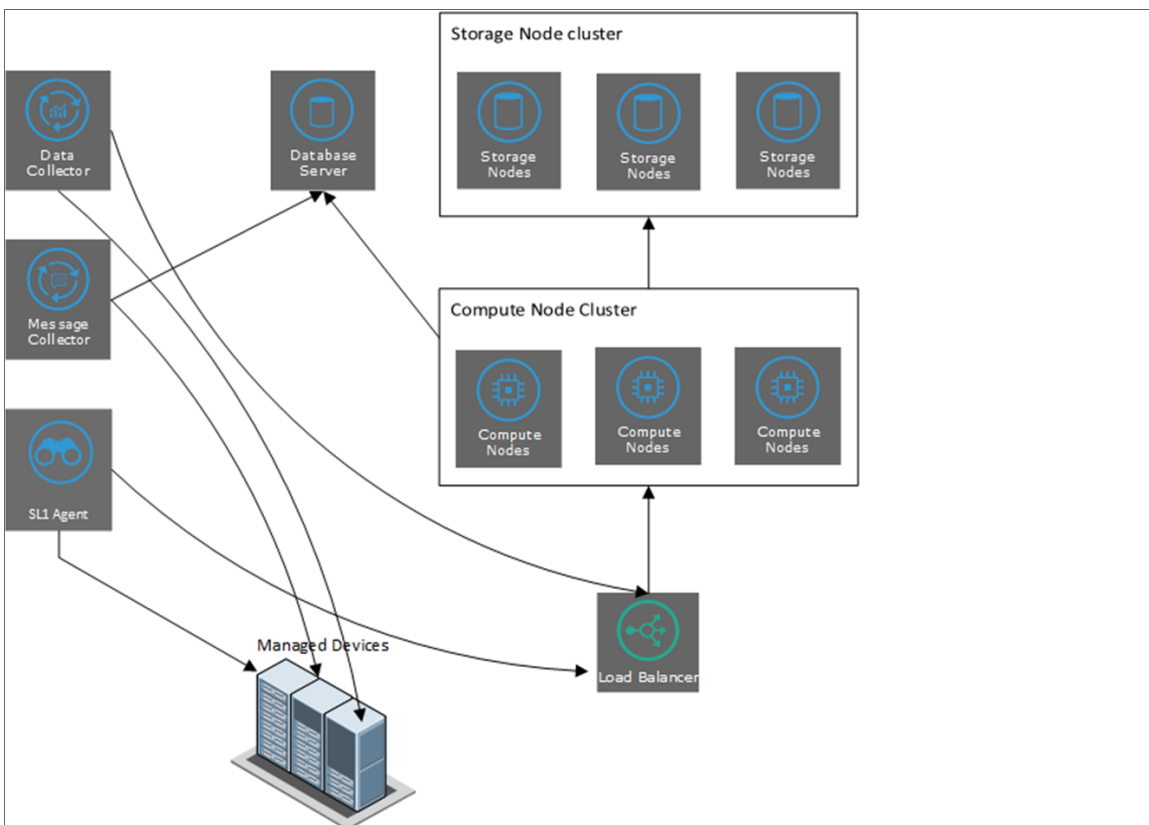
The Management Node runs on a single appliance or virtual machine. The Management Node allows users to install and update packages on the Compute Node cluster, Storage Node cluster, and the Load Balancer.

The SL1 Agent

In an SL1 Extended System, an **SL1 agent** collects data from the device on which it is installed and sends that data to a Load Balancer in front of a Compute Cluster. The Compute Cluster transforms the data and stores high-volume performance data in the Storage Cluster and other performance and configuration data in the Database Server.

In the diagram below:

- The **SL1 Agent** collects data from managed devices and sends the data to the Load Balancer and Compute Node cluster for processing and
- The optional **Message Collector** collects asynchronous traps and syslog messages and sends them to the Database Server.
- The **Data Collector** collects data from managed devices and sends the data to the Load Balancer and Compute Node cluster for processing and then storage.



Chapter

6



The SL1 Agent

Introduction

This manual describes the SL1 agent. The SL1 agent is a program that runs on a device or element monitored by SL1. The SL1 agent collects data from the device, interface, or other element and pushes that data back to SL1.

This manual is intended for System Administrators and ScienceLogic staff who are responsible for planning the architecture and configuration of SL1 systems.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

<i>The SL1 Agent</i>	28
<i>SL1 Distributed Architecture with an SL1 Agent</i>	29
<i>SL1 Extended with an SL1 Agent</i>	31

The SL1 Agent

The **SL1 agent** is a program that you can install on a device monitored by SL1. There is a Windows agent, an AIX agent, a Solaris agent, and a Linux agent. The agent collects data from the device and pushes that data back to SL1.

Similar to a Data Collector or Message Collector, the agent collects data about infrastructure and applications.

You can configure an agent to communicate with either the Message Collector or the Compute Cluster.

NOTE: The following minimum agent versions are required for SL1 12.1.1: **Windows** version 131; **Linux** version 174; **AIX** version 180; and **Solaris** version 180.

In the **SL1 Extended Architecture** (which includes Compute Nodes, Storage Nodes, and a Management Node), the Gen 3 agent collects the following data:

- **Device Availability.** SL1 can determine the availability state of a device (available or unavailable) and generate trended availability graphs based on uptime data collected by the agent.
- **Logs.** The SL1 agent can be configured to push logs to SL1 that match specific criteria from a log file or the Windows Event Log. You can view logs collected by the SL1 agent on the Logs pane of the **Device Investigator** page. The same logs also appear on the **[Logs]** tab in the **Device Properties** and **Device Summary** pages for that device. You can define event policies that specify how logs collected by an agent will trigger events.
- **Host Performance Metrics.** Using Dynamic Applications, SL1 translates data provided by an SL1 agent to trend the following metrics:
 - Overall CPU Utilization
 - CPU Utilization Breakdown
 - Disk Average Queue Length
 - Disk IO Utilization
 - Memory Utilization
 - Network Bytes Read
 - Network Bytes Written
 - Storage Available
 - Storage Total
 - Storage Utilization
 - Swap Utilization

You can view these metrics on the **Device Investigator** page and the **[Performance]** tab of the **Device Summary** panel for a specific device.

- **Host Configuration.** Using a Dynamic Application, SL1 collects the following configuration data based on data provided by the SL1 Agent:
 - The number and speed of the installed CPUs
 - The amount of installed memory
 - The overall and per-disk storage size
 - The total swap capacity (SL1 Extended Architecture only)

You can view the collected configuration data on the **[Configs]** tab of the **Device Investigator** page and the **Device Summary** panel.

- **Network Interface.** The SL1 agent collects a list of the network interfaces running on the device. You can view the list of interfaces on the **[Interfaces]** tab of the **Device Investigator** page and the **Device Summary** page. This list includes attributes such as the interface MAC address, IP address, position, and speed as well as inbound and outbound utilization, number of errors, and discard and usage percentage.
- **File System.** The SL1 agent collects data about the configuration of the file systems found within a device, such as name, size and, type as well as utilization data such as free space, size, and usage percentage. You can view the file system data on the **[Hardware?]** tab of the **Device Investigator** page and the **Device Summary** page.
- **System Processes.** The SL1 agent collects a list of all processes running on the device, such as name, process ID (PID), and state. You can view the list of processes on the **[Processes]** tab of the **Device Investigator** page and the **[Processes]** tab of the **Device Summary** page. Monitoring policies can be configured to trend and alert on process availability, process CPU usage, and process memory usage.
- **Windows Services.** The SL1 agent collects a list of all Windows services enabled on the device. This list includes attributes such as the service name and run state. You can view the list of Windows Services on the **[Services]** tab of the **Device Investigator** page and the **Device Summary** page.
- **Installed Software.** The SL1 agent collects a list of the software running on the device. This list includes attributes such as software name, version, and installation date. You can view the list of software on the **[Software]** tab of the **Device Investigator** page and the **Device Summary** page.

For more information about monitoring devices with the agent, see the **Monitoring Using the SL1 Agent** manual.

SL1 Distributed Architecture with an SL1 Agent

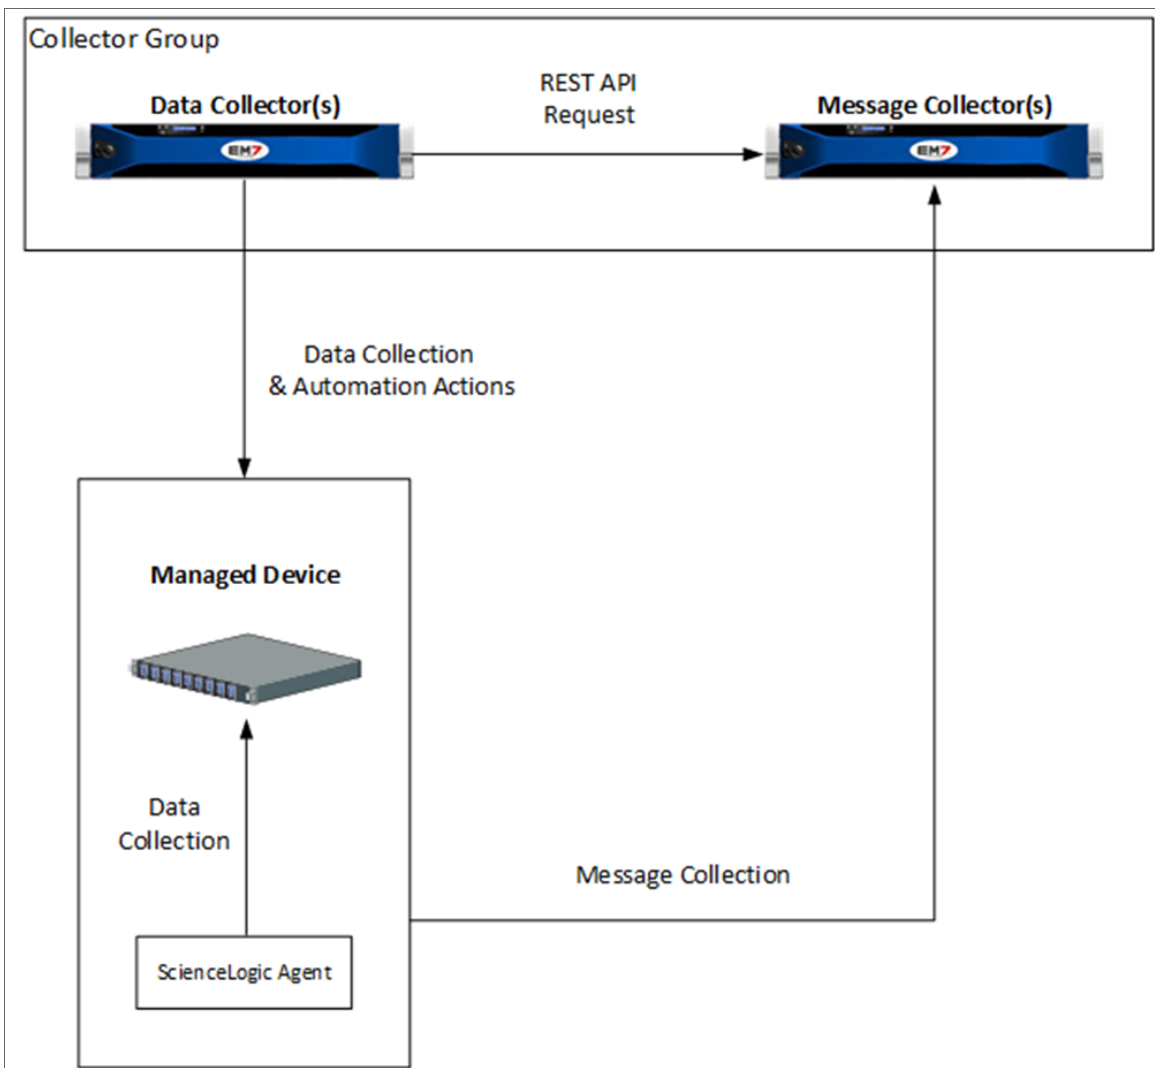
In an SL1 Distributed Architecture, the SL1 Agent collects data from the device on which it is installed and transfers that data to a Message Collector in an SL1 system using the HTTPS protocol. The Data Collector on which the Dynamic Applications and collection processes run then poll the Message Collector using the HTTPS protocol to transfer data to SL1.

TCP port 443 must be open between the Message Collector and the device on which an agent is installed.

In a Distributed Architecture, the SL1 agent requires a standalone, dedicated Message Collector. The Message Collector does not need to be dedicated to agent usage, but the Message Collector cannot be a Data Collector that also performs message collection

NOTE: Message Collectors that process data from the agent have different system requirements than Message Collectors that do not process data from the agent. For more information about the system requirements when running agents in a Distributed Architecture, see the [System Requirements](#) page at the ScienceLogic Support Site.

The diagram below shows the collection layer of a Distributed System containing both Data Collectors and Message Collectors in which the SL1 Agent is installed on a managed device.



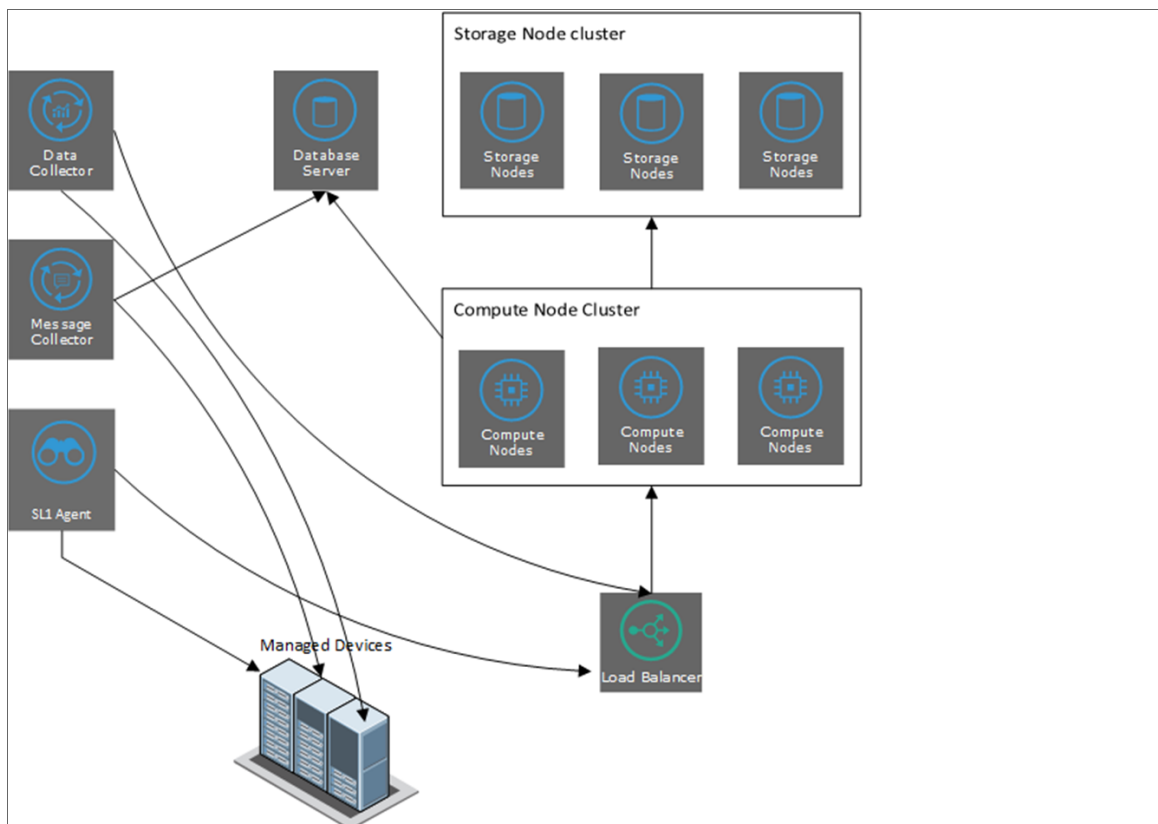
SL1 Extended with an SL1 Agent

In the SL1 Extended Architecture, an **SL1 agent** collects data from the device on which it is installed and sends that data to a Load Balancer in front of a Compute Cluster. The Compute Cluster transforms the data and stores high-volume performance data in the Storage Cluster and other performance and configuration data in the Database Server.

If required, agents can use an HTTP proxy server as an intermediate step in sending data to SL1.

In the diagram below:

- The **SL1 agent** collects data from managed devices and sends the data to the Load Balancer and Compute Node cluster for processing.
- The optional **Message Collector** collects asynchronous traps and syslog messages and sends them to the Database Server.
- The **Data Collector** collects data from managed devices and sends the data to the Load Balancer and Compute Node cluster for processing and then storage.



Using an agent in the SL1 Extended Architecture provides more configuration and performance data than using an agent in a Distributed Architecture. This additional data includes system vitals, log data, and extensible collection.

NOTE: Uploads that occur in 20-second intervals, sometimes called "snapshot uploads", are no longer supported for users using the non-Scylla pipeline. These 20-second uploads were replaced with the 1-minute upload default in SL1 version 11.2. ScienceLogic highly recommends that you ensure your agents are uploading in one-minute summarized uploads prior to upgrading. You can verify your uploads from the **[Settings]** tab in the current SL1 user interface. The agent pipeline is able to consume and summarize 1-minute and 5-minute payloads without the need for Scylla.

NOTE: For more information about the system requirements when running agents in an Extended Architecture, see the [System Requirements](#) page at the ScienceLogic Support Site.



Database Layer Configurations

Overview

This chapter contains diagrams for all possible configurations of the database layer in a distributed SL1 System. The possible configurations are:

- A single Database Server using local disk space for storage.
- One primary Database Server and one secondary Database Server for disaster recovery. Both Database Servers use local disk space for storage.
- Two Database Servers in a high availability cluster using local disk space for storage.
- Two Database Servers in a high availability cluster using local disk for storage with an additional Database Server for disaster recovery.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:


<i>Local Storage</i>	34
<i>Single Database Server Using Local Disk Space</i>	34
<i>Database with Disaster Recovery</i>	35
<i>Clustered Databases with High Availability using Local Disk</i>	36
<i>Clustered Databases with High Availability Using Local Disk and Disaster Recovery</i>	37

Local Storage

ScienceLogic Database Servers include one of two types of local storage: HDD or SSD.

- For SL1 systems that monitor 5,000 or more devices, Database Servers with SSD will provide significantly improved performance.
- For SL1 systems that monitor 10,000 or more devices, Database Servers with SSD are the required design standard.

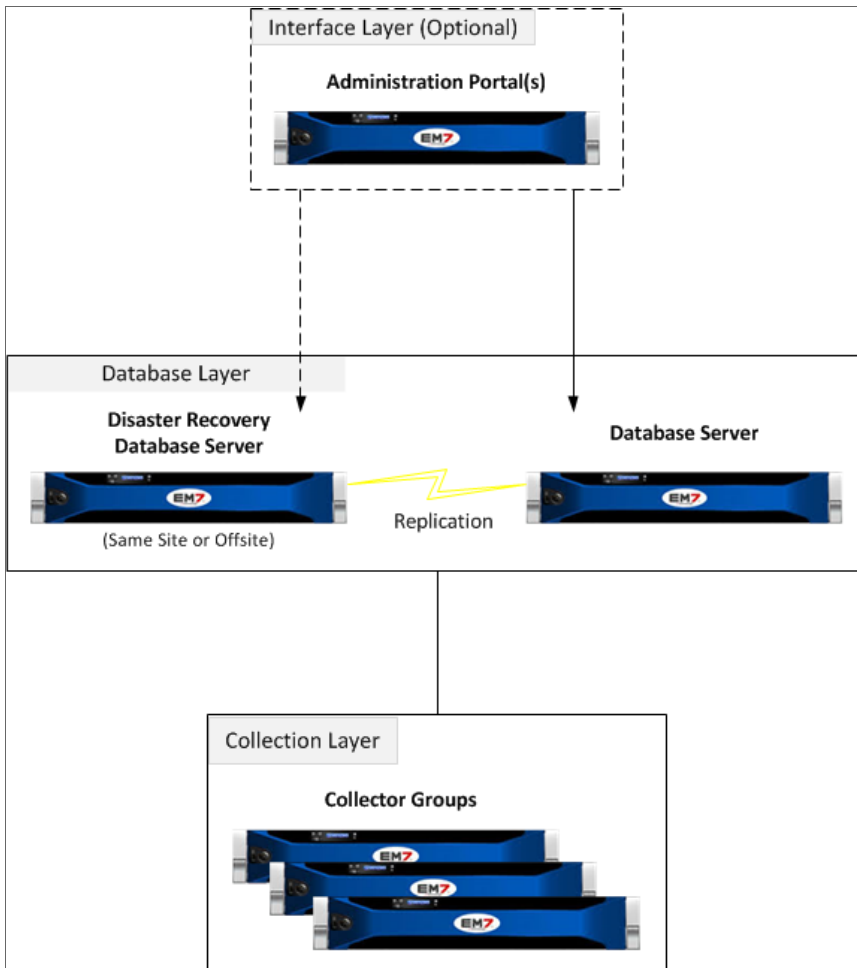
Single Database Server Using Local Disk Space



The following restrictions apply to this configuration:

- The interface layer is optional only if the system meets all of the requirements listed in the **Interface Layer Requirements** section in the **Distributed Architecture** section.
- To use the scheduled full backup feature with this configuration, the maximum amount of space used for data storage on the Database Server must be less than half the total disk space on the Database Server.

Database with Disaster Recovery



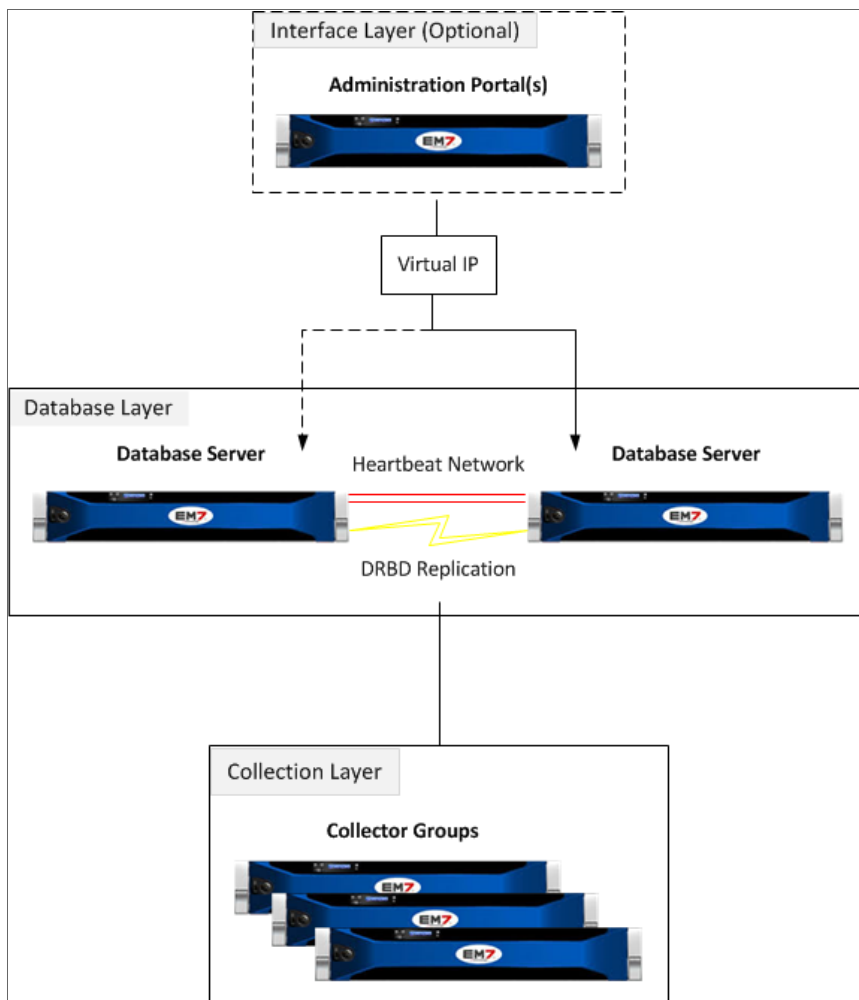
The following restrictions apply to this configuration:

- To use the scheduled full backup feature with this configuration, the maximum amount of space used for data storage on the Database Server must be less than half the total disk space on the Database Server.
- To use the scheduled online full backup feature, a Distributed SL1 system deployed on a database with disk storage must have a database less than 250 GB in size.

- For large SL1 Systems, you must use the backup options for offline backups (that is, the backup is not created locally and copied but instead is written directly to offline storage) from the Disaster Recovery database.
- The interface layer is optional only if the system meets all of the requirements listed in the **Interface Layer Requirements** section in the [Distributed Architecture](#) chapter.

For details on configuring Disaster Recovery for Database Servers, see the manual [Disaster Recovery](#).

Clustered Databases with High Availability using Local Disk




The following restrictions apply to this configuration:

- The two clustered Database Servers must be located in the same facility and attached directly to each other with a network cable.

- The two clustered Database Servers must use DRBD Replication to ensure data is synched between the two servers.
- The interface layer is optional only if the system meets all of the requirements listed in the **Interface Layer Requirements** section in the [Distributed Architecture](#) chapter.
- To use the scheduled online full backup feature, a Distributed SL1 system deployed on a database with disk storage must have a database less than 250 GB in size.
- The backup from DR feature cannot be used with this configuration.

For details on configuring High Availability for Database Servers, see the manual [Database Clustering](#).

Clustered Databases with High Availability Using Local Disk and Disaster Recovery



The following restrictions apply to this configuration:

- The two clustered Database Servers must be located in the same facility and attached directly to each other with a network cable.
- The two clustered Database Servers must use DRBD Replication to ensure data is synched between the two servers.
- The scheduled full backup feature cannot be used with this configuration.

- For large SL1 Systems, ScienceLogic recommends using the backup options for offline backups (that is, the backup is not created locally and copied but instead is written directly to offline storage) from the Disaster Recovery database.
- The interface layer is optional only if the system meets all of the requirements listed in the **Interface Layer Requirements** section in the [Distributed Architecture](#) chapter.

For details on configuring High Availability for Database Servers, see the manual **Database Clustering**.

For details on configuring Disaster Recovery for Database Servers, see the manual **Disaster Recovery**.



Interface Layer & API Configurations

Overview

In a Distributed SL1 system, the interface layer is optional if the system meets all the requirements listed in the [Distributed Architecture](#) section. If the interface layer is required, it must include at least one Administration Portal.

For all Distributed SL1 systems, browser session information is stored on the main database, not on the Administration Portal currently in use by the administrator or user.

Use the following menu options to navigate the SL1 user interface:

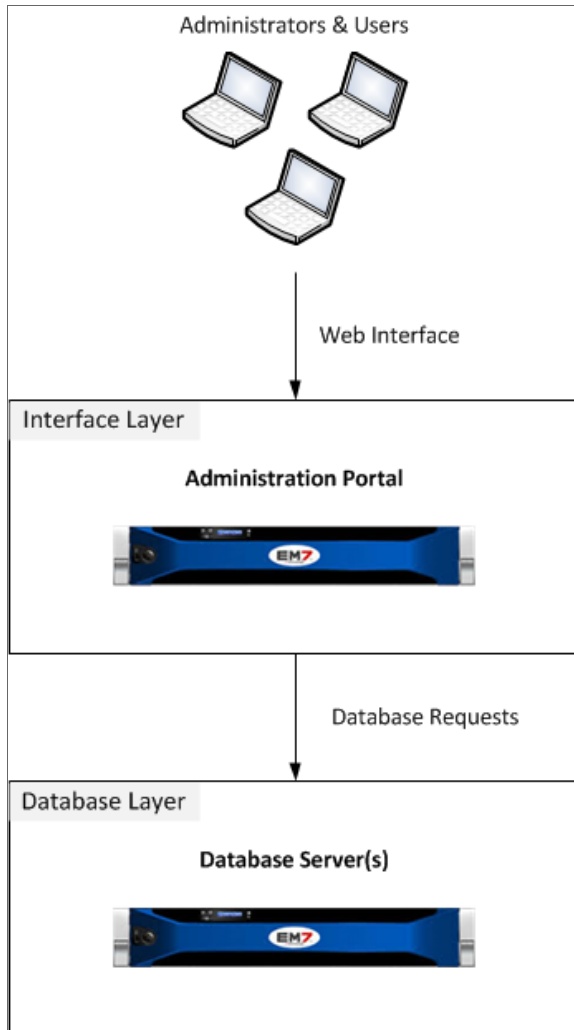
- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

Single Administration Portal	40
Multiple Administration Portals	41
Load-Balanced Administration Portals	42
API Layer Configurations	43

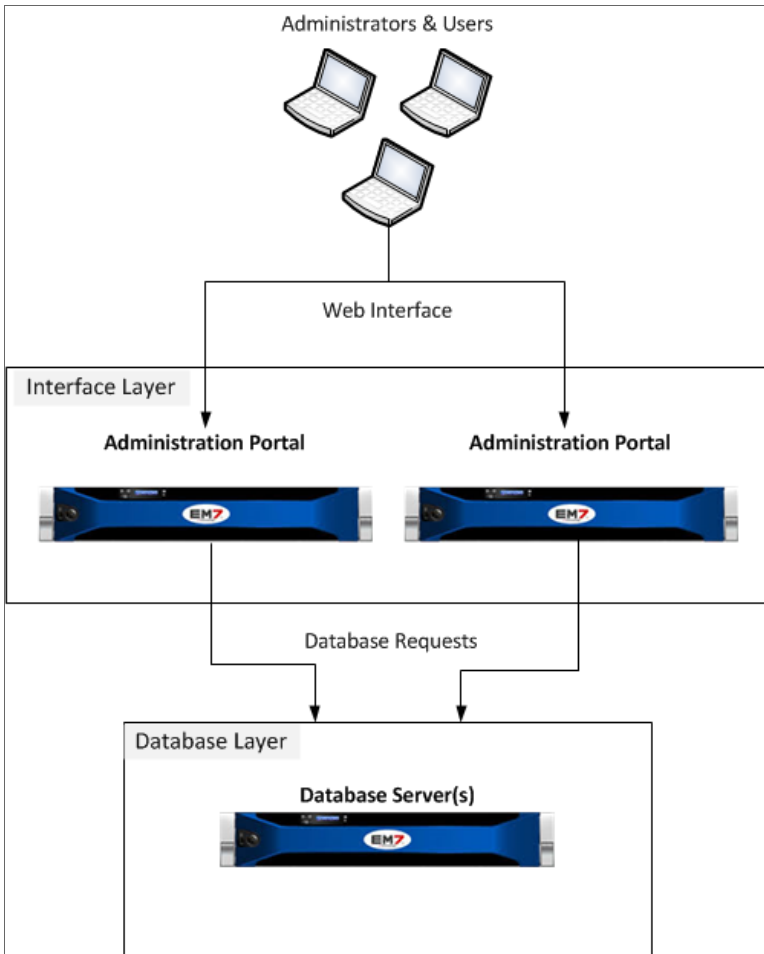
Single Administration Portal

In this configuration, the interface layer includes a single Administration Portal. An administrator or user can log in to SL1 using the Administration Portal:



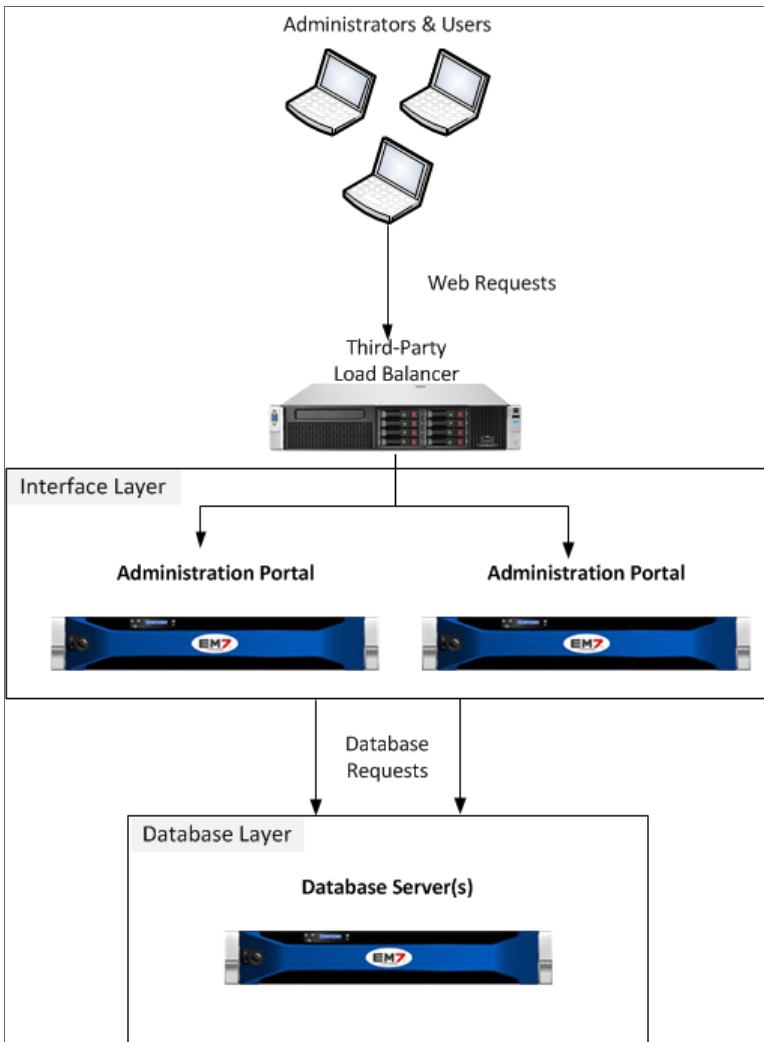
Multiple Administration Portals

The interface layer can include multiple Administration Portals. An administrator or user can log in to SL1 using any of the Administration Portals:



Load-Balanced Administration Portals

A third-party load-balancing solution can be used to distribute traffic evenly among the Administration Portals:



NOTE: ScienceLogic does not recommend a specific product for this purpose and does not provide technical support for configuring or maintaining a third-party load-balancing solution.

API Layer Configurations

REST API

SL1 provides an optional REST-based API that external systems can use to configure SL1 and access collected data.

The REST API is available through the **Administration Portal**, the **All-In-One Appliance**, and the **Database Server**.

The REST API is also used within the SL1 platform in Dynamic Applications and Run Book Automations.

Support for the REST API will continue after 8.14.0, However support will be limited to bug fixes and ensuring the API continues to provide visibility to data residing in the relational data store (MariaDB).

Performance and log data stored in a NoSQL store (the Storage Node cluster) will not be visible via the ScienceLogic REST API.

GraphQL API

SL1 provides an optional GraphQL-based API that external systems can use to access collected data.

Customers planning to upgrade to a version of SL1 that stores data in a NoSQL store or a graph database should plan to update any integrations to use the GraphQL API instead.

Performance and log data stored in a NoSQL store (the Storage Node cluster) is visible only via the GraphQL API.

The GraphQL API is available through the **Administration Portal**, the **All-In-One Appliance**, and the **Database Server**.

Collector Group Configurations

Overview

For Distributed SL1 Systems, a collector group is a group of Data Collectors. Data Collectors retrieve data from managed devices and applications. This collection occurs during initial discovery, during nightly updates, and in response to policies and Dynamic Applications defined for each managed device. The collected data is used to trigger events, display data in SL1, and generate graphs and reports.

Grouping multiple Data Collectors allows you to:

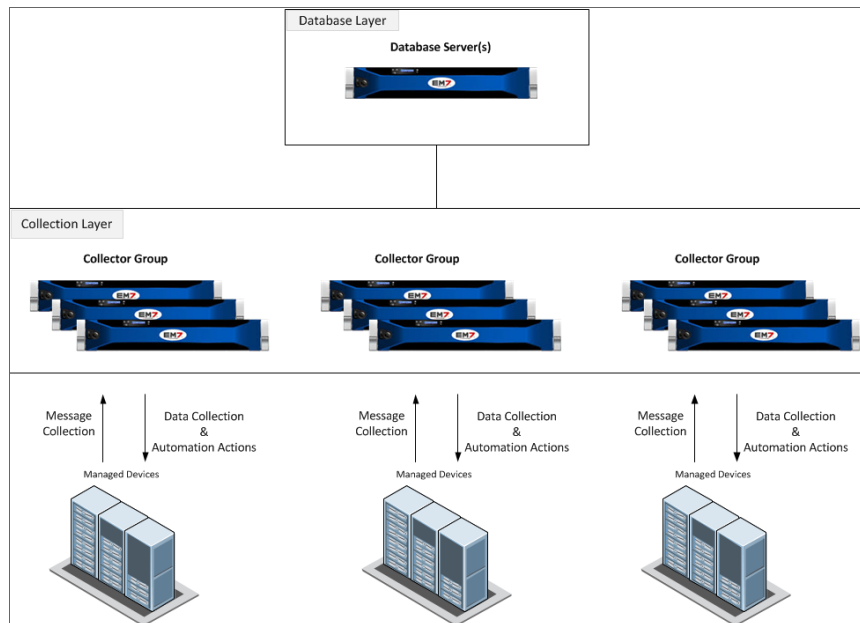
- Create a load-balanced collection system, where you can manage more devices without loss of performance. At any given time, the Data Collector with the lightest load handles the next discovered device.
- Create a redundant, high-availability system that minimizes downtime should a failure occur. If a Data Collector fails, another Data Collector is available to handle collection until the problem is solved.

This chapter covers the following topics:

<i>Collector Groups</i>	45
<i>Traditional and PhoneHome Collectors</i>	45
<i>Using a Data Collector for Message Collection</i>	47
<i>Using Multiple Data Collectors in a Collector Group</i>	47
<i>How Collector Groups Handle Component Devices</i>	48
<i>High Availability for Data Collectors</i>	49
<i>Using Message Collectors in a Collector Group</i>	50

Collector Groups

In a Distributed SL1 System, the Data Collectors and Message Collectors are organized as **collector groups**. Each monitored device is aligned with a collector group:



A Distributed SL1 system must have one or more collector groups configured. The Data Collectors included in a collector group must have the same hardware configuration.

A Distributed SL1 system could include collector groups configured using each of the possible configurations. For example, suppose an enterprise has a main data center that contains most of the devices monitored by the SL1 system. Suppose the enterprise also has a second data center where only a few devices are monitored by the SL1 system. The SL1 system might have two collector groups:

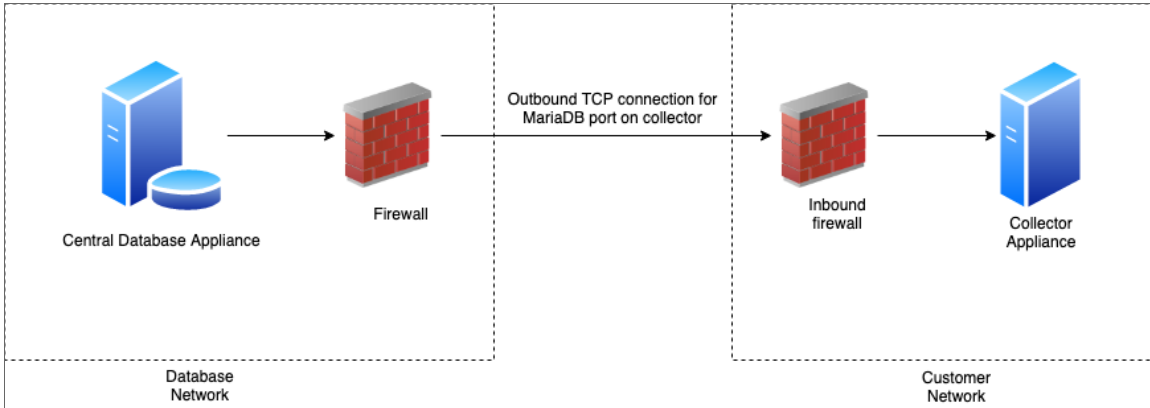
- In the main data center, a collector group configured with high availability that contains multiple Data Collectors and Message Collectors.
- In the second data center, a collector group that contains a single Data Collector that is also responsible for message collection.

Traditional and PhoneHome Collectors

SL1 supports two methods for communication between a Database Server (an SL1 Central Database or an SL1 Data Engine) and the Data Collectors and Message Collectors:

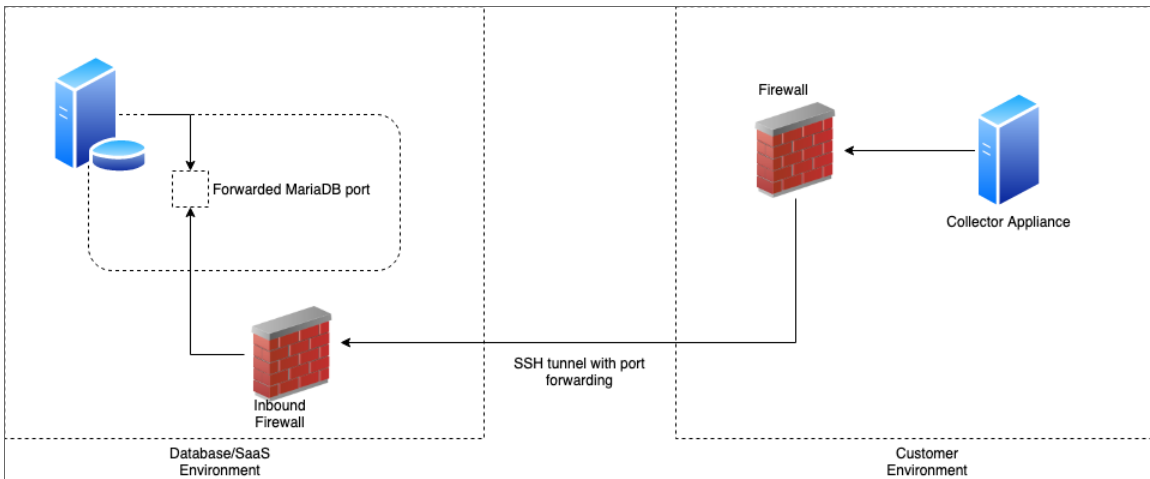
- Traditional
- PhoneHome

In the **Traditional** method, the SL1 services on the Database Server initiate a new connection to the MariaDB port on the collector to read and write data. The connection request traverses the network, including the Internet if necessary, eventually reaching the collector. For this approach to work, the collector administrator must allow ingress communication from the Database Server on port 7707 (the MariaDB port on the collector). The communication is encrypted using SSL whenever possible.



The benefit of the traditional method is that communication to the Database Server is extremely limited, so the Database Server remains as secure as possible.

In the **PhoneHome** method, the collectors initiate an outbound connection to the Database Server over SSH. After authenticating, the client forwards the local MariaDB port onto the Database Server using a loopback remote IP address. A corresponding SL1 appliance is added using the loopback IP. When the SL1 services on the database try to make a connection to the collector's MariaDB, they connect locally to the loopback IP address, in contrast to reaching out to the collector's IP or DNS name. The communication is encrypted.

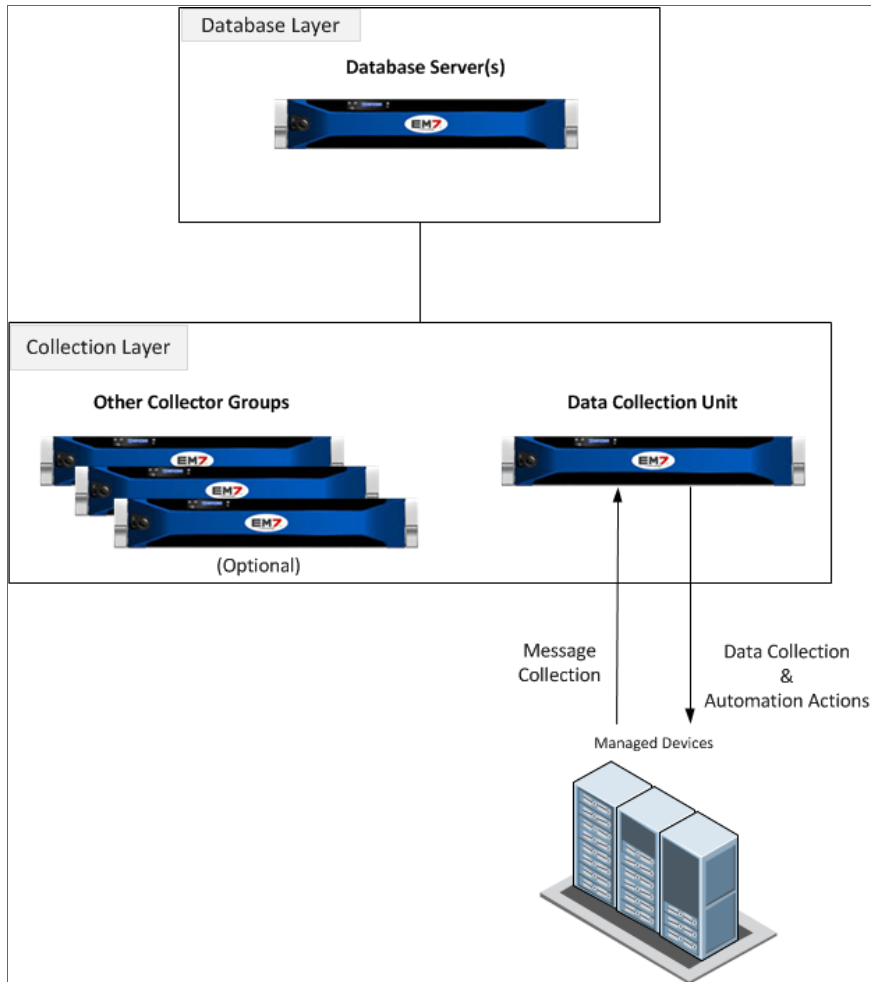


The benefits of this method are that no ingress firewall rules need to be added as the collector initiates an outbound connection, and no new TCP ports are opened on the network that contains the Data Collectors.

The PhoneHome configuration uses public key/private key authentication to maintain the security of the Database Server. Each Data Collector is aligned with an SSH account on the Database Server and uses SSH to communicate with the Database Server. Each SSH account on the Database Server is highly restricted, has no login access, and cannot access a shell or execute commands on the Database Server.

Using a Data Collector for Message Collection

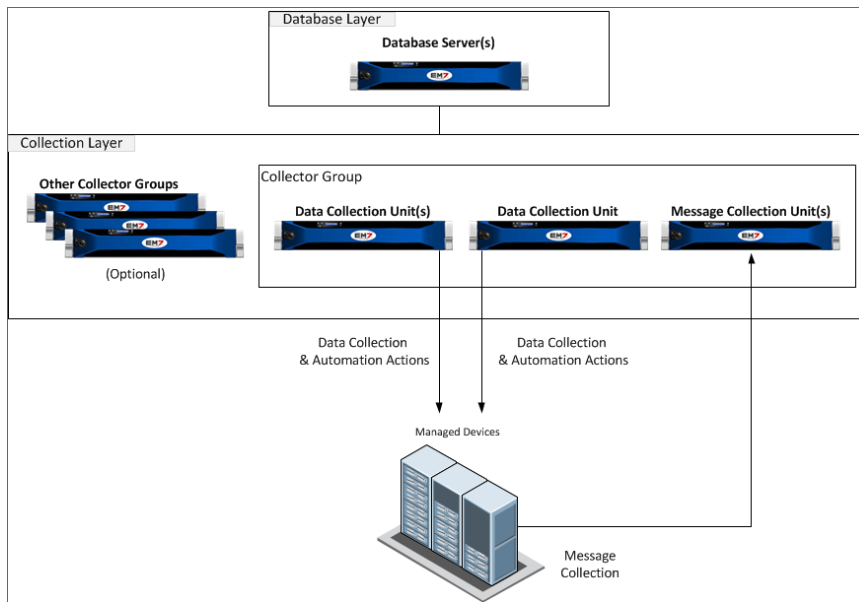
To use a Data Collector for message collection, the Data Collector must be in a collector group that contains no other Data Collectors or Message Collectors.



NOTE: When a Data Collector is used for message collection, the Data Collector can handle fewer inbound messages than a dedicated Message Collector.

Using Multiple Data Collectors in a Collector Group

A collector group can include multiple Data Collectors to maximize the number of managed devices. In this configuration, the collector group is not configured for high availability:



In this configuration:

- All Data Collectors in the collector group must have the same hardware configuration
- If you need to collect syslog and trap messages from the devices aligned with the collector group, you must include a Message Collector in the collector group. For a description of how a Message Collector can be added to a collector group, see the [Using Message Collection Units in a Collector Group](#) section.
- SL1 evenly distributes the devices monitored by a collector group among the Data Collectors in the collector group. Devices are distributed based on the amount of time it takes to collect data for the Dynamic Applications aligned to each device.
- **Component devices** are distributed differently than physical devices; component devices are always aligned to the same Data Collector as its root device.

NOTE: If you merge a component device with a physical device, the SL1 system allows data for the merged component device and data from the physical device to be collected on different Data Collectors. Data that was aligned with the component device is always collected on the Data Collector for its root device. If necessary, data aligned with the physical device can be collected on a different Data Collector.

How Collector Groups Handle Component Devices

Collector groups handle component devices differently than physical devices.

For physical devices (as opposed to component devices), after the SL1 system creates the device ID, the SL1 system distributes devices, round-robin, among the Data Collectors in the specified collector group.

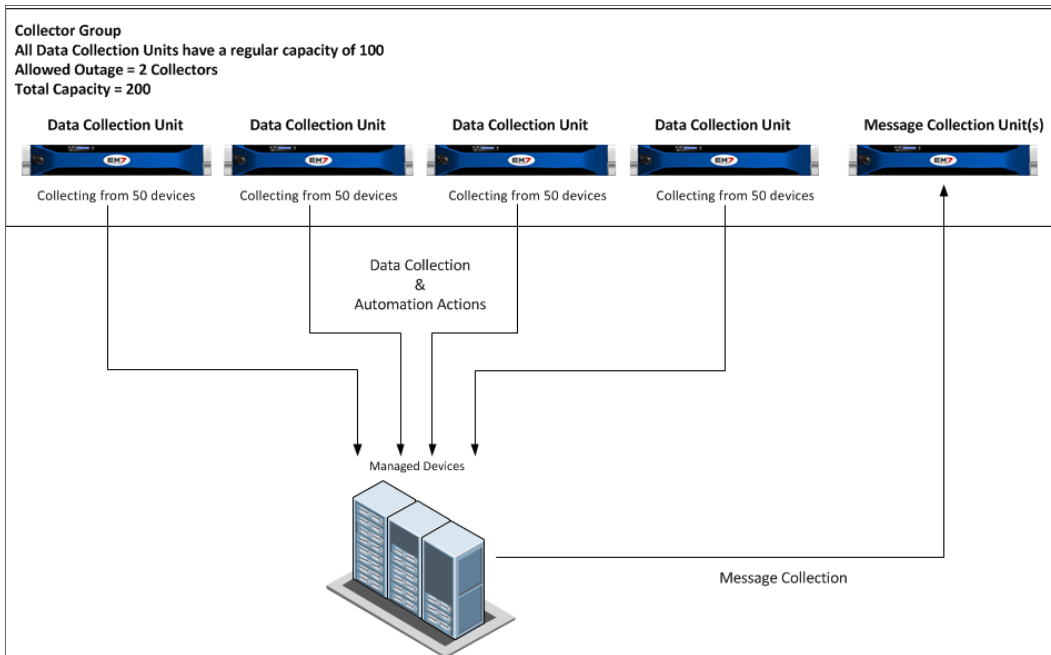
Each component device must use the same Data Collector used by its root device. For component devices, the SL1 System must keep all the component devices on the same Data Collector used by the root device (the

physical device that manages the component devices). SL1 cannot distribute the component devices among the Data Collectors in the specified collector group.

NOTE: If you merge a component device with a physical device, the SL1 System allows data for the merged component device and data from the physical device to be collected on different Data Collectors. Data that was aligned with the component device is always collected on the Data Collector for its root device. If necessary, data aligned with the physical device can be collected on a different Data Collector.

High Availability for Data Collectors

To configure a collector group for high availability, the collector group must include multiple Data Collectors:



In this configuration:

- All Data Collectors in the collector group must have the same hardware configuration.
- If you need to collect syslog and trap messages from the devices monitored by a high availability collector group, you must include a Message Collector in the collector group. For a description of how a Message Collector can be added to a collector group, see the [Using Message Collection Units in a Collector Group](#) section.
- On the **Collector Group Management** page (System > Settings > Collector Groups), each collector group that is configured for high availability must have the **Collector Failover** field set to *On (Maximize Reliability)* and a value in the **Collectors Available for Failover** field that specifies the minimum number of Data Collectors that must be available (i.e., with a status of "Available [0]") before a Data Collector failover can occur. For example:

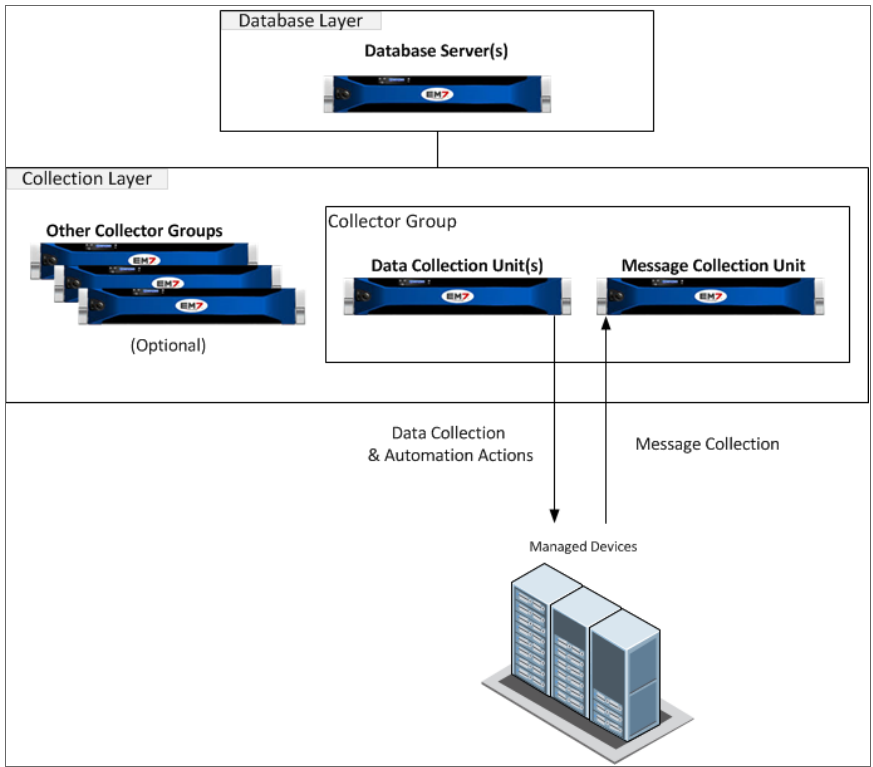
- For collector groups with only two Data Collectors, the **Collectors Available for Failover** field will contain the value "1 collector".
- For collector groups with more than two Data Collectors, the **Collectors Available for Failover** field will contain values from a minimum of one-half of the total number of Data Collectors up to a maximum of one less than the total number of Data Collectors. For example, for a collector group with eight Data Collectors, the possible values in the **Collectors Available for Failover** field would be 4, 5, 6, and 7.
- SL1 will never automatically increase the maximum number of Data Collectors that can fail in a collector group. For example, suppose you have a collector group with three Data Collectors and the **Collectors Available for Failover** field is set to "2". If you add a fourth Data Collector to the collector group, SL1 will automatically set the **Collectors Available for Failover** field to "3" to maintain the maximum number of Data Collectors that can fail as "one". However, you can override this automatic setting by manually changing the value in the **Collectors Available for Failover** field.

NOTE: If you set the **Collectors Available for Failover** field to a value equaling half of your available Data Collectors and a 50% Data Collector outage occurs and the remaining Data Collectors are down by one, no rebalance will occur. If you specify a value equaling one-third of the total number of Data Collectors, then a rebalance will be attempted until your overall capacity falls below one-third of your Data Collectors, thereby maximizing your resiliency but minimizing the opportunity for your system to enter an unproductive rebalancing loop.

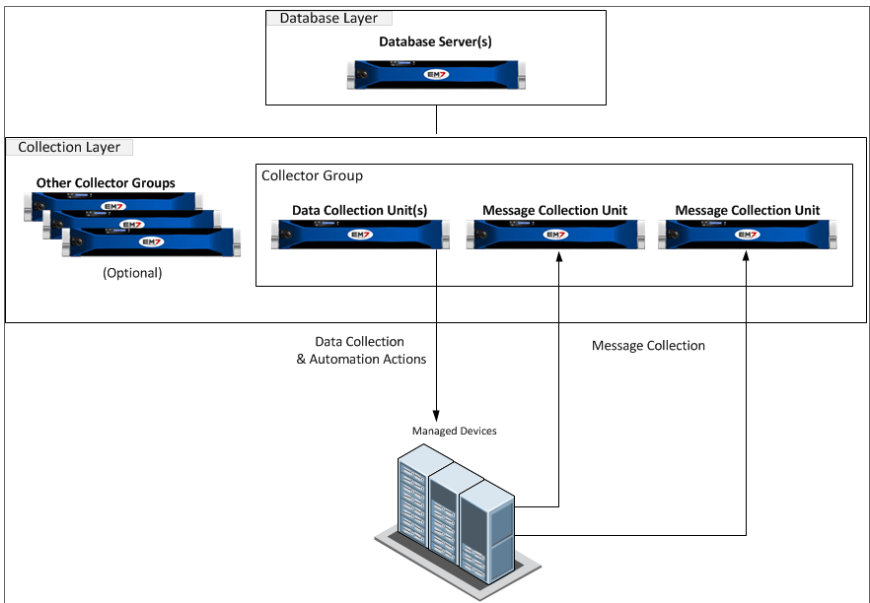
CAUTION: If the number of available Data Collectors is less than the value in the **Collectors Available for Failover** field, SL1 will not fail over within the collector group. SL1 will not collect any data from the devices aligned with the failed Data Collector(s) until the failure is fixed on enough Data Collector(s) to equal the value in the **Collectors Available for Failover** field. SL1 will generate a critical event.

Using Message Collectors in a Collector Group

If you need to collect syslog and trap messages from the devices monitored by a collector group that includes multiple Data Collectors, you must include a Message Collector in the collector group:



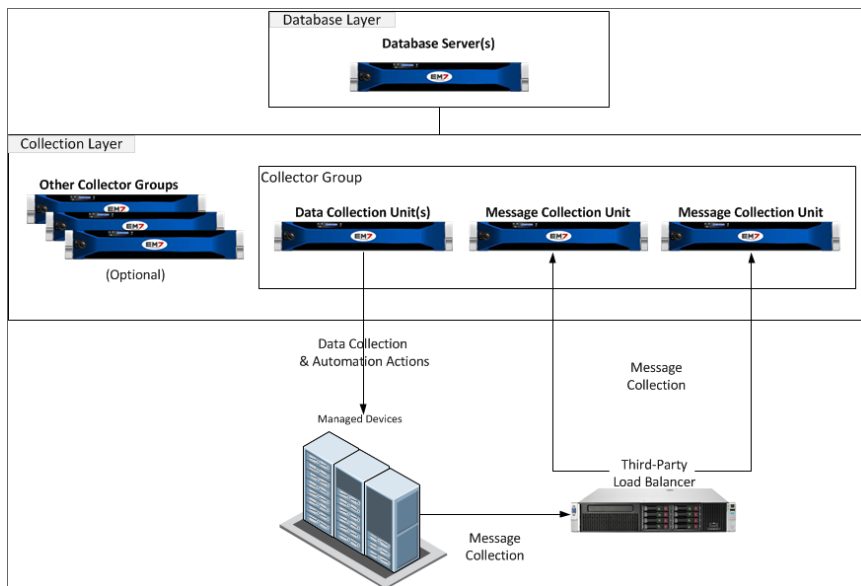
If your monitored devices generate a large amount of syslog and trap messages, a collector group can include multiple Message Collectors:



In this configuration, a monitored device can send syslog and trap messages to either Message Collector.

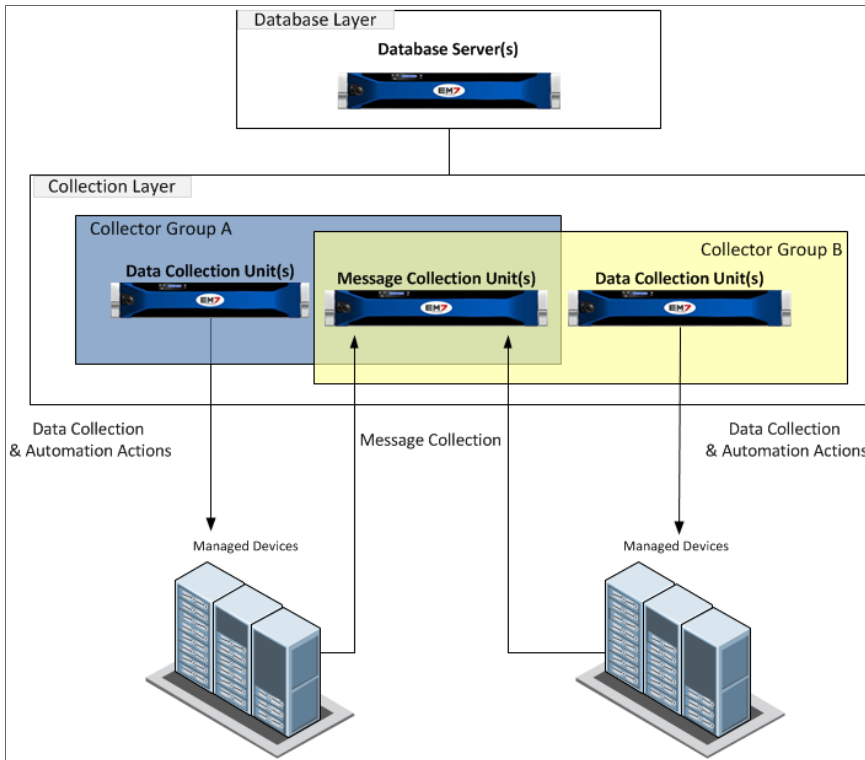
NOTE: Each syslog and trap message should be sent to only one Message Collector.

A third-party load-balancing solution can be used to distribute syslog and trap messages evenly among the Message Collectors in a collector group:



NOTE: ScienceLogic does not recommend a specific product for this purpose and does not provide technical support for configuring or maintaining a third-party load-balancing solution.

One or more Message Collectors can be included in multiple collector groups:



In this configuration, each managed device in collector group A and collector group B must use a unique IP address when sending syslog and trap messages. The IP address used to send syslog and trap messages is called the **primary IP**. For example, if a device monitored by collector group A and a device monitored by collector group B use the same primary IP address for data collection, one of the two devices must be configured to use a different IP address when sending syslog and trap messages.

A collector group can have multiple Message Collectors that are also included in other collector groups. It is possible to include every Message Collector in your SL1 System in every collector group in your SL1 System.



Using Virtual Machines and Cloud Instances

Overview

Each appliance in your ScienceLogic configuration can be run as a virtual machine.

For details about supported hypervisors and the requirements and specifications for each SL1 appliance, see the ScienceLogic Support Site: <https://support.sciencelogic.com/s/system-requirements>

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

Caveats	55
-------------------------------	----

Caveats

NOTE: Microsoft Hyper-V Linux Integration software cannot be installed on any SL1 appliances.

NOTE: ScienceLogic databases have a very high bandwidth of memory changes under normal operations, often in excess of 10Gb/sec. This rate of memory change limits the feasibility of VM live migration methods (such as vMotion) for SL1 appliances because on moderately large databases, the rate of memory change is too high to be synchronized between hosts over a 10Gb/sec ethernet link.

The following requirements apply to all virtualized appliances:

- Fixed storage is required. Dynamically-expanding storage is not supported.
- Memory over-commit is not supported. In the case of VMware, this means that 100% of memory must be reserved for all ScienceLogic appliances.
- Running SL1 on a virtualization server that is close to capacity will result in unexpected behavior.

© 2003 - 2024, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: legal@sciencelogic.com. For more information, see <https://sciencelogic.com/company/legal>.

ScienceLogic

800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010