



---

# ScienceLogic Architecture

SL1 version 8.14.0

---

# Table of Contents

<b>Architecture Overview</b> .....	<b>1</b>
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 .....	5
The SL1 Agent .....	6
User Interface .....	6
Database Server .....	6
Data Collection .....	7
Message Collection .....	7
System Updates .....	8
<b>Backup, Recovery &amp; High Availability Options</b> .....	<b>9</b>
Configuration Backups .....	10
Scheduled Backups to External File Systems .....	10
Backup of Disaster Recovery Database .....	11
Database Replication for Disaster Recovery .....	11
High Availability for Database Servers .....	11
Differences Between Disaster Recovery and High Availability for Database Servers .....	12
High Availability for Data Collection .....	13
Restrictions .....	13
<b>All-In-One Architecture</b> .....	<b>14</b>
Unsupported Configurations .....	15
The SL1 Agent .....	15
<b>Distributed Architecture</b> .....	<b>16</b>
Architecture Outline .....	17
Database Capacity .....	17
Data Collection .....	18
Message Collection .....	18
Interface Layer Requirements .....	18
The SL1 agent .....	19
<b>Extended Architecture</b> .....	<b>21</b>
What is the SL1 Extended Architecture? .....	22
Architecture Outline .....	22
Database Capacity .....	24
Message Collection .....	24
Interface Layer Requirements .....	24
Compute Nodes .....	24
Load Balancer .....	25
Storage Nodes .....	25
Management Node .....	25
The SL1 Agent .....	26
<b>SL1 Agent</b> .....	<b>27</b>
The SL1 Agent .....	28

SL1 Distributed with the SL1 Agent .....	29
SL1 Extended with the SL1 Agent .....	30
<b>Database Layer Configurations .....</b>	<b>32</b>
Local Storage .....	33
Single Database Server Using Local Disk Space .....	33
Database with Disaster Recovery .....	34
Clustered Databases with High Availability using Local Disk .....	35
Clustered Databases with High Availability Using Local Disk and Disaster Recovery .....	36
<b>Interface Layer &amp; API Configurations .....</b>	<b>38</b>
Single Administration Portal .....	39
Multiple Administration Portals .....	40
Load-Balanced Administration Portals .....	41
API Layer Configurations .....	42
REST API .....	42
GraphQL API .....	42
<b>Collector Group Configurations .....</b>	<b>43</b>
Collector Groups .....	44
"Traditional" and "PhoneHome" Collectors .....	44
Using a Data Collector for Message Collection .....	46
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 .....	52
<b>Using Virtual Machines and Cloud Instances .....</b>	<b>55</b>
Caveats .....	56

---

# Chapter

# 1



## Architecture Overview

---

### Introduction

This manual describes the architecture of SL1 systems, covering all possible 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 (.

This chapter includes the following topics:

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

User Interface .....	6
Database Server .....	6
Data Collection .....	7
Message Collection .....	7
System Updates .....	8

---

## SL1 Configurations

SL1 includes one or more appliances or nodes (part of a cluster) that function together to provide the SL1 platform and application. There are three primary configurations:

- **All-In-One.** In this configuration, a single appliance provides all the functions of SL1. The capacity of an All-In-One instance cannot be increased by adding additional appliances. This configuration is best for deployments that monitor less than 1000 devices.
- **Distributed.** In this configuration, the functions of SL1 are divided between multiple appliances. A **Distributed** instance of SL1 can be as small as two appliances or include multiple instances of each appliance. This configuration is for best for deployments that monitor more than 1000 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 appliances to these configurations.

---

## SL1 Appliances and Nodes

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

- **All-In-One Appliance.** A single appliance that provides handles all the functions of SL1. The capacity of an All-In-One instance cannot be increased by adding additional appliances.
- **Database Server.** The Database Server contains a relational database used for all policy and configuration data. This database also stores "agent-less" performance data.
- **Data Collector.** Data Collectors run multiple services for "agent-less" collection. Each Data Collector is responsible for collecting a specific set of information from a specific set of devices according to its locally stored configuration policy. Collections may also receive instructions for asynchronous tasks including discovery, user driven device querying, run book automation for incident enrichment or remediation.
- **Message Collector.** The Message Collector collects syslog and trap messages from devices.
- **Administration Portal.** The Administration Portal provides dedicated UI access and API access to SL1.

- **Computer Cluster.** A cluster of at least three Compute Nodes. Each Compute Node includes Kubernetes node and a range of services for ingestion and transformation of collected data.
- **Load Balancer.** A single node that manages the workload of the Compute Cluster.
- **Storage Cluster.** A cluster of at least three Storage Nodes. The Storage Nodes contain NoSQL databases. These Storage Nodes store high-volume performance data collected by the SL1 Agent and the Data Collector appliance.
- **Management Node.** The **Management Node** allows users to install and update packages on the Compute Node cluster, Storage Node cluster, and the Load Balancer.

---

## SL1 Appliance Functions

In a **Distributed SL1 System**, there are five general functions that an SL1 appliance can perform. In large SL1 systems, a dedicated appliance performs each function. In smaller systems, some appliances perform several functions. In an **All-In-One Appliance** system, a single SL1 appliance performs all five functions.

In an **Extended SL1 System**, there are four additional types of SL1 Appliances that support the SL1 Agent and provide scale.

### User Interface

Administrators and users access the user interface through a web browser session. 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 appliance that provides the user interface function also generates all scheduled reports. The following appliances can provide the user interface:

- **All-In-One Appliance.** In small SL1 Systems, the **All-In-One Appliance** performs all functions.
- **Database Server.** In small to mid-size SL1 Systems, a **Database Server** can provide the user interface in addition to its database function.
- **Administration Portal.** In mid-size and large SL1 Systems, a dedicated **Administration Portal** appliance provides the user interface.

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

### Database Server

The appliance that provides the database function is responsible for:

- Storing all configuration data and policy data.
- Storing performance data collected from managed devices that do not use the SL1 Agent.

- In a distributed system, pushing data to and retrieving data from the appliances responsible for collecting data.
- Processing and normalizing collected data.
- Allocating tasks to the other appliances in the SL1 System.
- Executing 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.** In small SL1 Systems, the **All-In-One Appliance** performs all functions.
- **Database Server.** In all distributed SL1 Systems, a **Database Server** provides all database functions. In some small to mid-size systems, a **Database Server** might also provide the user interface in addition to its main database function.

## Data Collection

The SL1 appliances that perform the data collection function retrieve data from monitored devices where the SL1 Agent is not installed. In a distributed system, appliances that perform the data collection function also perform some pre-processing of collected data and execute automation actions in response to events.

The following appliances can perform the collection function:

- **All-In-One Appliance.** In small SL1 systems, the **All-In-One Appliance** performs all functions.
- **Data Collector.** In a **Distributed** SL1 system, one or more Data Collectors collect data from devices where the SL1 Agent is not installed. Data Collectors can be configured 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). In smaller systems, a Data Collector can also perform the message collection function.
- **SL1 Agent.** The SL1 Agent is installed on each monitored device or application and pushes that data back to SL1. The SL1 Agent can collect more granular data than can be collected by polling the device periodically. You can install and use multiple agents, as needed.

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 perform the message collection function receive and process inbound, asynchronous syslog and trap messages from monitored devices.

The following appliances can perform the message collection function:

- **All-In-One Appliance.** In small SL1 systems, the **All-In-One Appliance** performs all functions.
- **Message Collector.** In distributed systems, the **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 appliance, not a combination **Data Collector/Message Collector**.
- **Data Collector.** In smaller distributed systems, a Data Collector can also perform the message collection function.

## Compute

SL1 Extended includes a **Compute Cluster** that includes a minimum of three Compute Nodes.

Each Compute Node run Kubernetes and includes services for ingestion and transformation of data collected by the SL1 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 relationship data for applications, infrastructure, and networks.

## Load Balancer

Provides load-balancing for the Compute Node cluster. You can include two Load Balancers, for redundancy.

## Storage

SL1 Extended includes a **Storage Cluster** that includes a minimum of three Storage Nodes.

These Storage Nodes store high-volume 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 (high-volume performance data) or the Database Server (other performance data and configuration data).

## Management

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 a **Distributed SL1 System**, there are five general functions that an SL1 appliance can perform. In large SL1 systems, a dedicated appliance performs each function. In smaller systems, some appliances perform several functions. In an **All-In-One Appliance** system, a single SL1 appliance performs all five functions.

In an **Extended SL1 System**, there are four additional types of SL1 Appliances that support the SL1 Agent and provide scale.

## User Interface

Administrators and users access the user interface through a web browser session. 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 appliance that provides the user interface function also generates all scheduled reports. The following appliances can provide the user interface:

- **All-In-One Appliance**. In small SL1 Systems, the **All-In-One Appliance** performs all functions.
- **Database Server**. In small to mid-size SL1 Systems, a **Database Server** can provide the user interface in addition to its database function.
- **Administration Portal**. In mid-size and large SL1 Systems, a dedicated **Administration Portal** appliance provides the user interface.

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

## Database Server

The appliance that provides the database function is responsible for:

- Storing all configuration data and policy data.
- Storing performance data collected from managed devices that do not use the SL1 Agent.
- In a distributed system, pushing data to and retrieving data from the appliances responsible for collecting data.
- Processing and normalizing collected data.
- Allocating tasks to the other appliances in the SL1 System.
- Executing 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.** In small SL1 Systems, the **All-In-One Appliance** performs all functions.
- **Database Server.** In all distributed SL1 Systems, a **Database Server** provides all database functions. In some small to mid-size systems, a **Database Server** might also provide the user interface in addition to its main database function.

## Data Collection

The SL1 appliances that perform the data collection function retrieve data from monitored devices where the SL1 Agent is not installed. In a distributed system, appliances that perform the data collection function also perform some pre-processing of collected data and execute automation actions in response to events.

The following appliances can perform the collection function:

- **All-In-One Appliance.** In small SL1 systems, the **All-In-One Appliance** performs all functions.
- **Data Collector.** In a **Distributed** SL1 system, one or more Data Collectors collect data from devices where the SL1 Agent is not installed. Data Collectors can be configured 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). In smaller systems, a Data Collector can also perform the message collection function.
- **SL1 Agent.** The SL1 Agent is installed on each monitored device or application and pushes that data back to SL1. The SL1 Agent can collect more granular data than can be collected by polling the device periodically. You can install and use multiple agents, as needed.

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 perform the message collection function receive and process inbound, asynchronous syslog and trap messages from monitored devices.

The following appliances can perform the message collection function:

- **All-In-One Appliance.** In small SL1 systems, the **All-In-One Appliance** performs all functions.
- **Message Collector.** In distributed systems, the **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 appliance, not a combination **Data Collector/Message Collector**.
- **Data Collector.** In smaller distributed systems, a Data Collector can also perform the message collection function.

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. You can install and use multiple SL1 Agents, as needed.

Because an agent is always running on a device, the SL1 Agent can collect more granular data than can be collected by polling the device periodically with a Data Collector. You can collect data from devices using only the SL1 Agent or using a combination of the SL1 Agent and Data Collectors.

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.

---

## System Updates

SL1 includes an automatic update tool, called System Updates, that automatically installs software updates and patches. The updated software or patch is loaded onto a single ScienceLogic appliance, and the software is automatically sent out to each ScienceLogic appliance that needs it. When you apply an update to your SL1 system, SL1 automatically uploads any newer versions of the default PowerPacks.

The **System Updates** page in the user interface allows you to update the software on a SL1 system. You must first download the updates to the local computer (computer where you are running the browser). You can then load the software update to the SL1 system. When the software is loaded onto the SL1 system, you can install the software immediately or schedule the software to be installed at a later time.

When you apply an update to your SL1 system, SL1 automatically uploads any newer versions of the default PowerPacks. If a PowerPack included in an update is not currently installed on your system, the platform will automatically install the PowerPack. If a PowerPack included in an update is currently installed on your system, the platform will automatically import the PowerPack, but will not install the PowerPack.

For full details on the System Update tool, see the **System Administration** manual.

## 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 includes the following topics:

<i>Configuration Backups</i> .....	10
<i>Scheduled Backups to External File Systems</i> .....	10
<i>Backup of Disaster Recovery Database</i> .....	11
<i>Database Replication for Disaster Recovery</i> .....	11

<i>High Availability for Database Servers</i> .....	11
<i>Differences Between Disaster Recovery and High Availability for Database Servers</i> .....	12
<i>High Availability for Data Collection</i> .....	13
<i>Restrictions</i> .....	13

---

## 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

Running a full backup creates a complete backup of the ScienceLogic database. This type of backup is recommended for SL1 systems in small-to-medium enterprises. Note the following about full backups:

- A full backup comprises all configuration data, performance data, and log data
- Full backup is disabled by default. You can configure SL1 to automatically launch this backup at a frequency and time you specify.
- During full backup, the ScienceLogic database remains online.
- The full backup is stored on a remote NFS mount or a remote SMB mount.

**NOTE:** *ScienceLogic does not recommend Full Backups for large SL1 systems.* For large SL1 systems, ScienceLogic recommends you use the **DR Backup** option or a SAN with snapshot technology to backup and restore data.

**NOTE:** You can configure the staging and remote directories used for backups in the `master.system_settings_backup` database, to ensure that the backup can be placed on a directory that has enough disk space. For remote directories, the current unix time will be appended to the directory name to ensure the directory name is unique each time the backup runs.

---

## 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.

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:

- 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.
- 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.
- A High Availability configuration is not supported for All-In-One Appliances. A Disaster Recovery configuration is not supported for All-In-One Appliances.
- 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.

---

## 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.



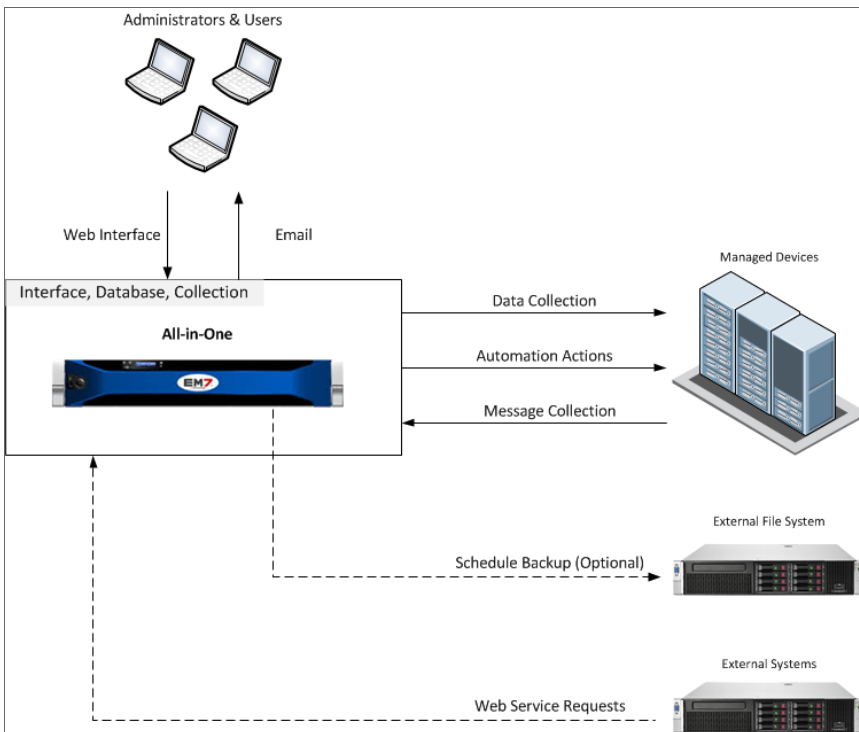
# Chapter

# 3

## All-In-One Architecture

### Overview

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



---

## Unsupported Configurations

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

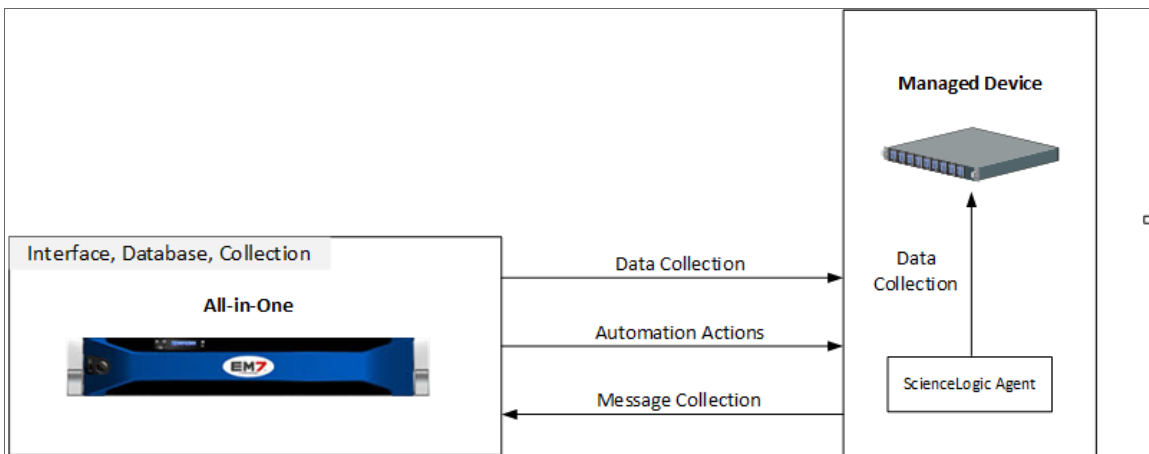
- Using a SAN for storage
- Disaster Recovery
- High Availability for Database Servers
- High Availability for Data Collectors
- Additional Data Collectors, Message Collectors, or Administration Portals

---

## 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.

The diagram below shows the interface, database, and collection layers of an All-In-One architecture in which the agent is installed on a managed device.



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 (.

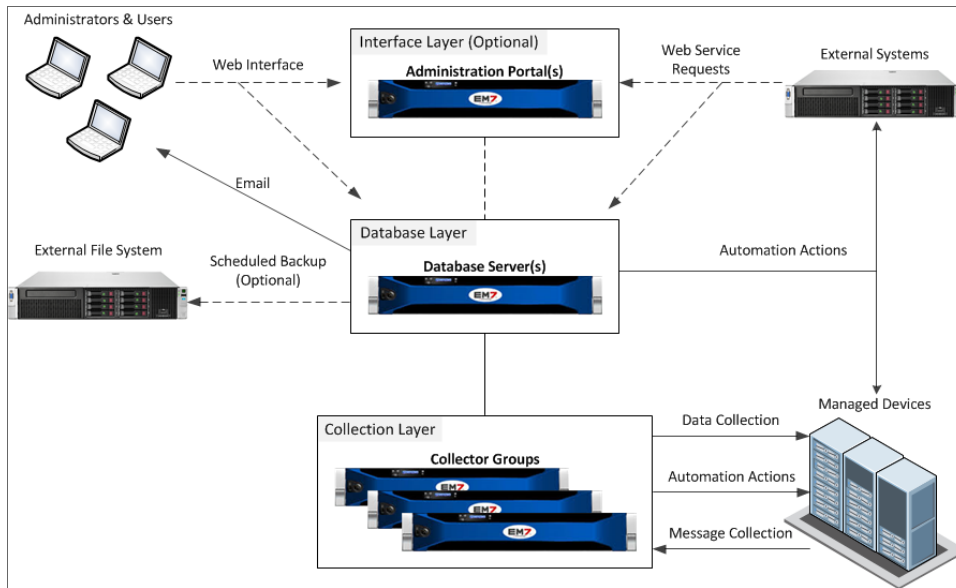
This chapter includes the following topics:

<a href="#">Architecture Outline</a> .....	17
<a href="#">Database Capacity</a> .....	17
<a href="#">Data Collection</a> .....	18
<a href="#">Message Collection</a> .....	18
<a href="#">Interface Layer Requirements</a> .....	18
<a href="#">The SL1 agent</a> .....	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.

For each discovered device, SL1 calculates a **device rating** based on the license that was used to install the SL1 system and the amount of collection SL1 is performing for the device. For example, a physical device that is polled frequently and has many aligned Dynamic Applications and monitoring policies could have a higher device rating than a component device that is polled infrequently and is aligned with only one or two Dynamic Applications.

The sum of all the rating values for all devices discovered in the system cannot exceed the Capacity for the Database Server or All-In-One Appliance. The Capacity Rating is defined by the license key issued to you by ScienceLogic. All Database Servers and All-In-One Appliances in a system must have the same capacity rating.

For details on sizing and scaling your SL1 system to fit your workload, see the Customer Portal: <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

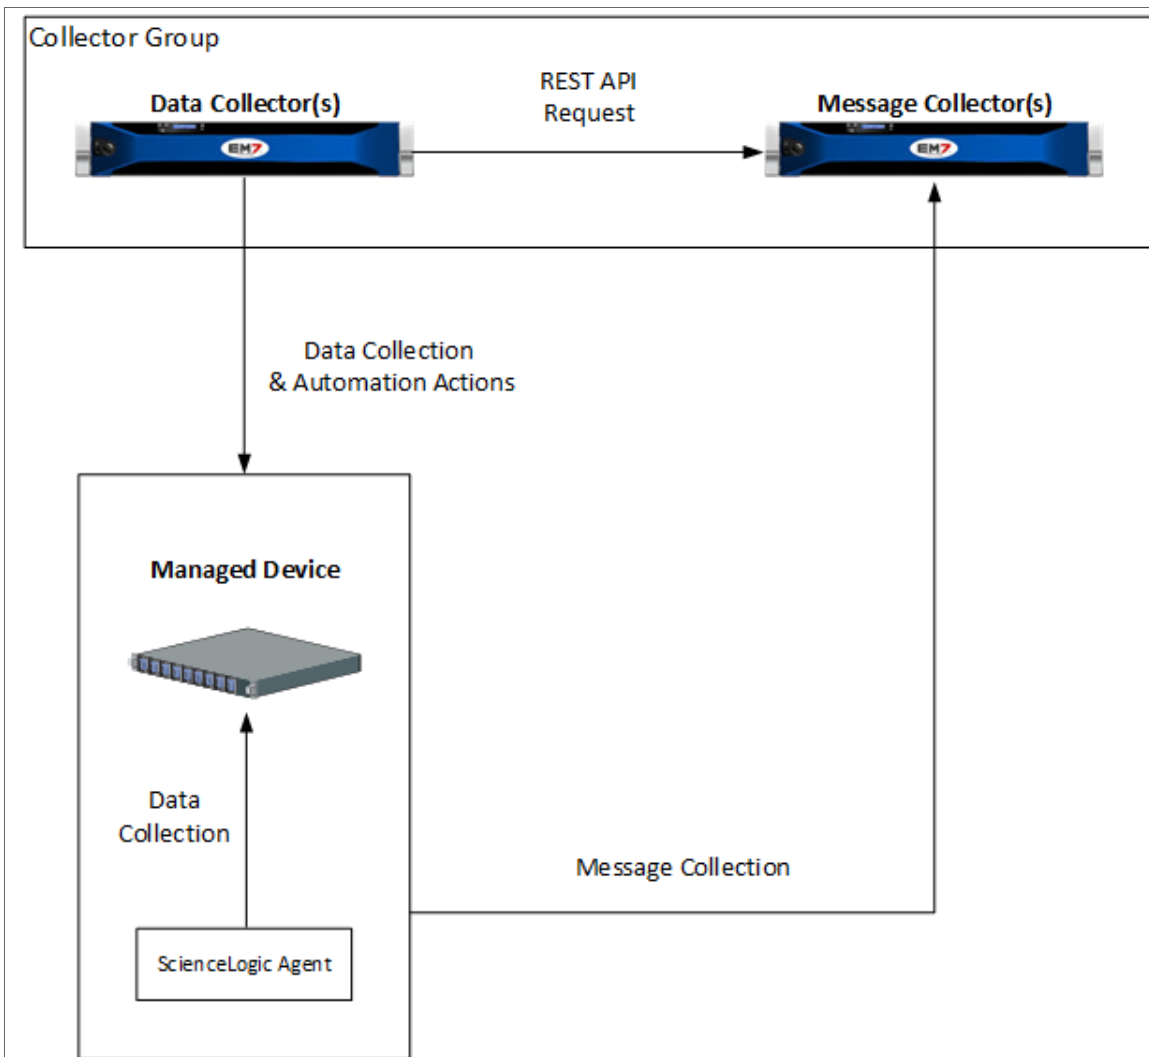
In a distributed architecture, the SL1 Agent collects data from the device on which it is installed and transfers that data to a Message Collector in an SL1 system using the HTTPS protocol. The Data Collector on which the Dynamic Applications and collection processes run then poll the Message Collector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

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

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

**NOTE:** Message Collectors that process data from the SL1 agent have different system requirements than Message Collectors that do not process data from the SL1 agent. For more information about the system requirements when running SL1 agents in a distributed architecture, see <https://support.sciencelogic.com/s/system-requirements>.

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 includes the following topics:

<i>What is the SL1 Extended Architecture?</i> .....	22
<i>Architecture Outline</i> .....	22
<i>Database Capacity</i> .....	24
<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.

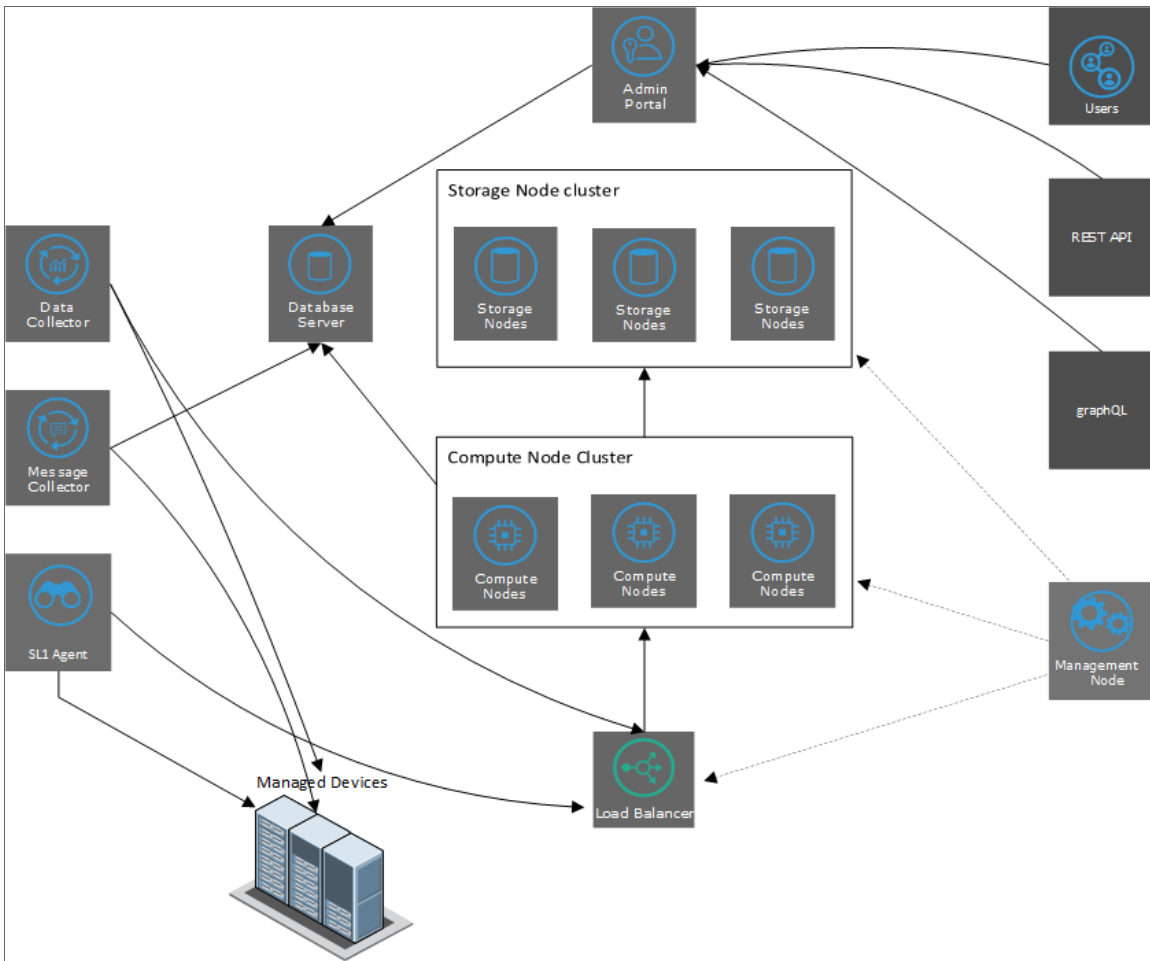
---

## 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 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

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.

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

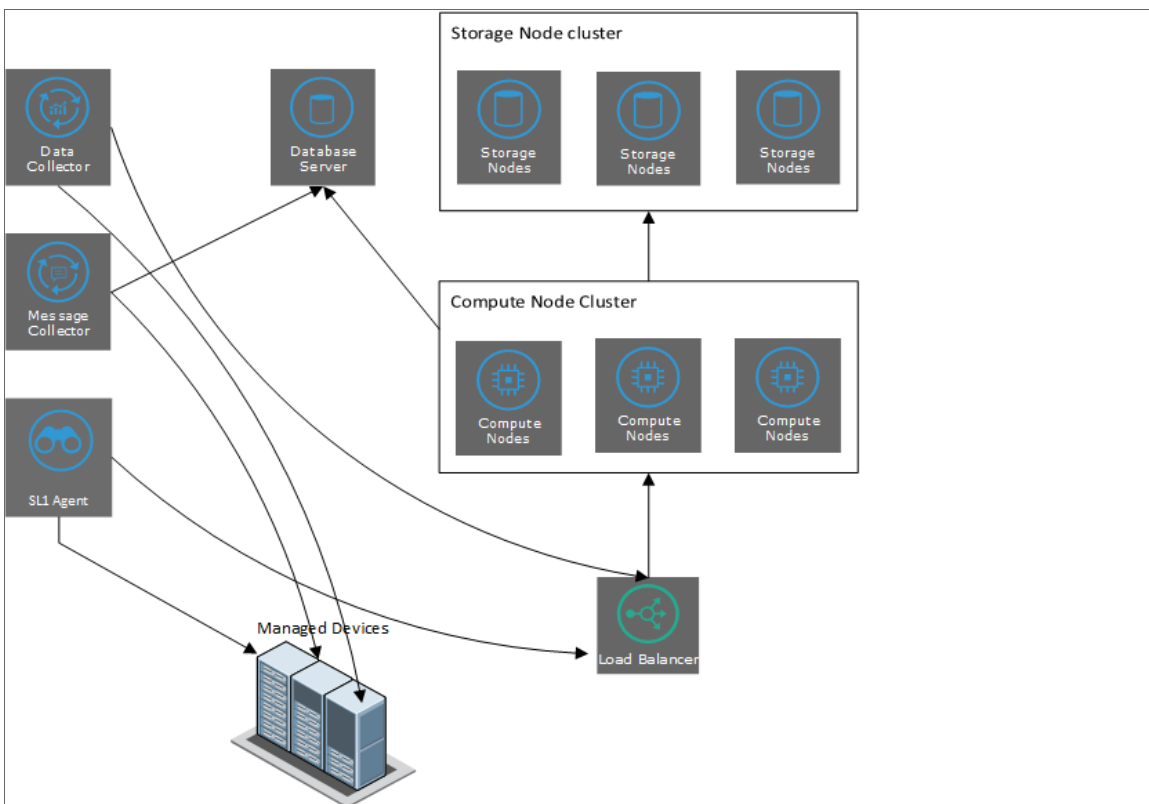
The Management Node runs on a single appliance or virtual machine. The Management Node allows users to install and update packages on the Compute Node cluster, Storage Node cluster, and the Load Balancer.

## 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

## 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 includes the following topics:

<i>The SL1 Agent</i> .....	28
<i>SL1 Distributed with the SL1 Agent</i> .....	29
<i>SL1 Extended with the SL1 Agent</i> .....	30

---

## 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. You can install and use multiple SL1 Agents, as needed.

Because an agent is always running on a device, the SL1 Agent can collect more granular data than can be collected by polling the device periodically with a Data Collector. You can collect data from devices using only the SL1 Agent or using a combination of the SL1 Agent and Data Collectors.

In SL1 **Extended**, the SL1 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.
- **Host Performance Metrics.** Using a Dynamic Application, SL1 translates data provided by an SL1 agent to trend the following metrics:
  - Overall CPU Utilization
  - Per-Processor CPU Utilization
  - Disk Average Queue Length
  - Disk 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 Interfaces.** 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** panel. This list includes attributes such as the interface MAC address, IP address, position, and speed.
- **System Processes.** The SL1 Agent collects a list of all processes running on the device. You can view the list of processes on the **[Processes]** tab of the **Device Reports** panel and the **[Processes]** tab of the **Device Investigator** 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** panel.
- **Open Ports.** The SL1 Agent collects a list of open TCP and UDP ports on the device. You can view the list of open ports on the **[TCP/UDP Ports]** tab of the **Device Reports** panel and the **[Ports]** tab of the **Device Investigator** page. Monitoring policies can be configured to trend and alert on port availability.
- **Logs.** The SL1 Agent can be configured to push logs that match specific criteria from a log file or the Windows Event Log to SL1. You can view logs collected by the SL1 Agent on the **Device Investigator** page and the **Device Logs** page (Registry > Monitors > Logs) for a device. Logs can be configured to trigger events.

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

---

## SL1 Distributed with the SL1 Agent

In a 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.

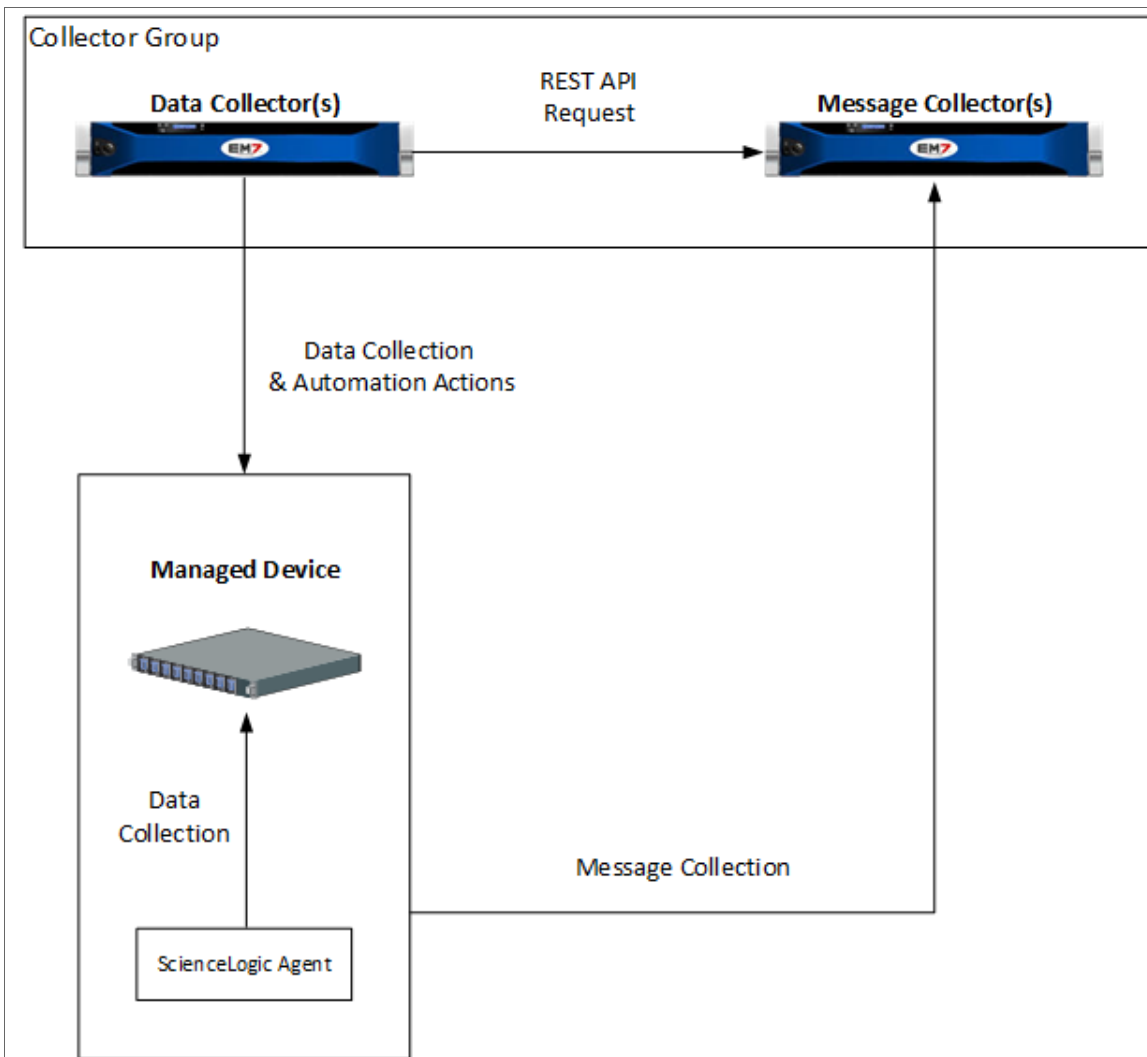
TCP port 443 must be open between the the Message Collector and the device on which an agent is installed.

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

**NOTE:** Message Collectors that process data from the SL1 agent have different system requirements than Message Collectors that do not process data from the SL1 agent. For more information about the system requirements when running SL1 agents in a distributed architecture, see <https://support.sciencelogic.com/s/system-requirements>.

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 the SL1 Agent

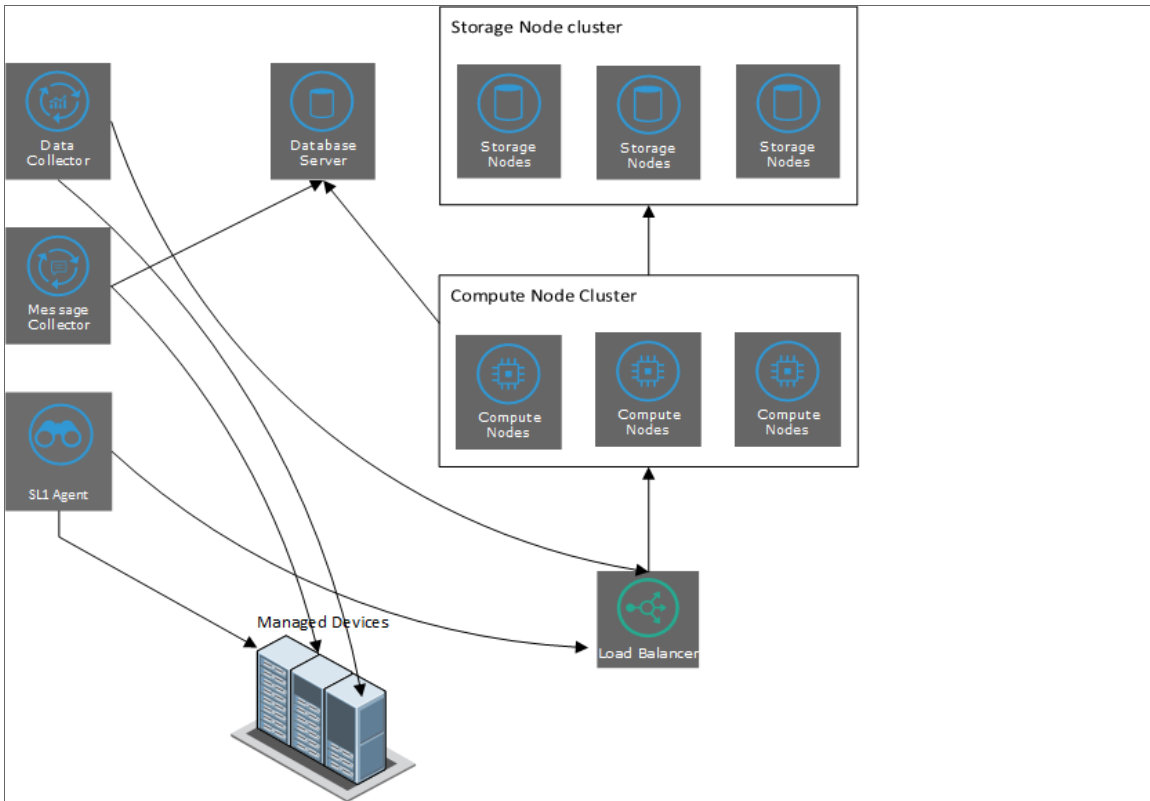
In an SL1 Extended System, an **SL1 agent** collects data from the device on which it is installed and sends that data to a Load Balancer in front of a Compute Cluster. The Compute Cluster transforms the data and stores high-volume performance data in the Storage Cluster and other performance and configuration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

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

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.



## 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 includes the following topics:

<i>Local Storage</i> .....	33
<i>Single Database Server Using Local Disk Space</i> .....	33
<i>Database with Disaster Recovery</i> .....	34
<i>Clustered Databases with High Availability using Local Disk</i> .....	35
<i>Clustered Databases with High Availability Using Local Disk and Disaster Recovery</i> .....	36

---

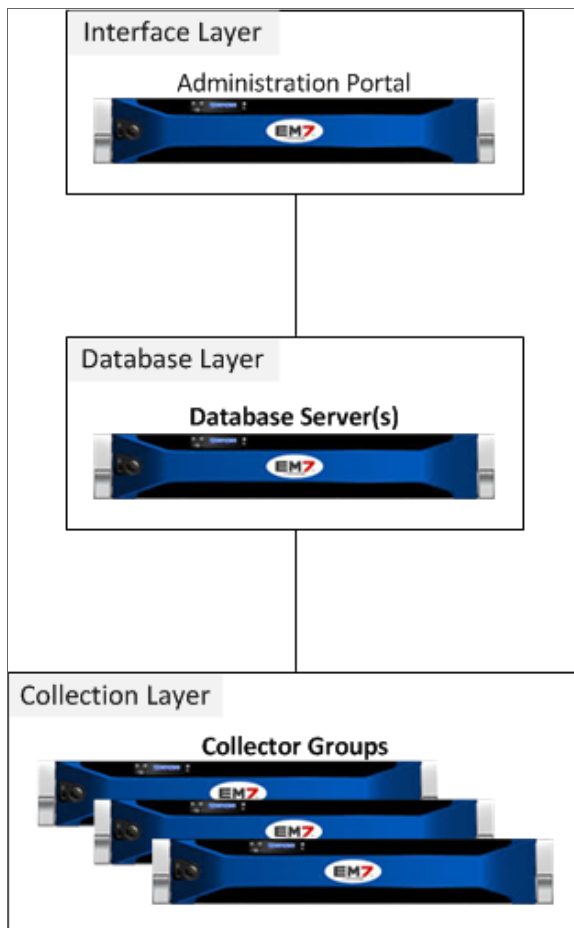
## Local Storage

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 synced 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 synced 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 ().

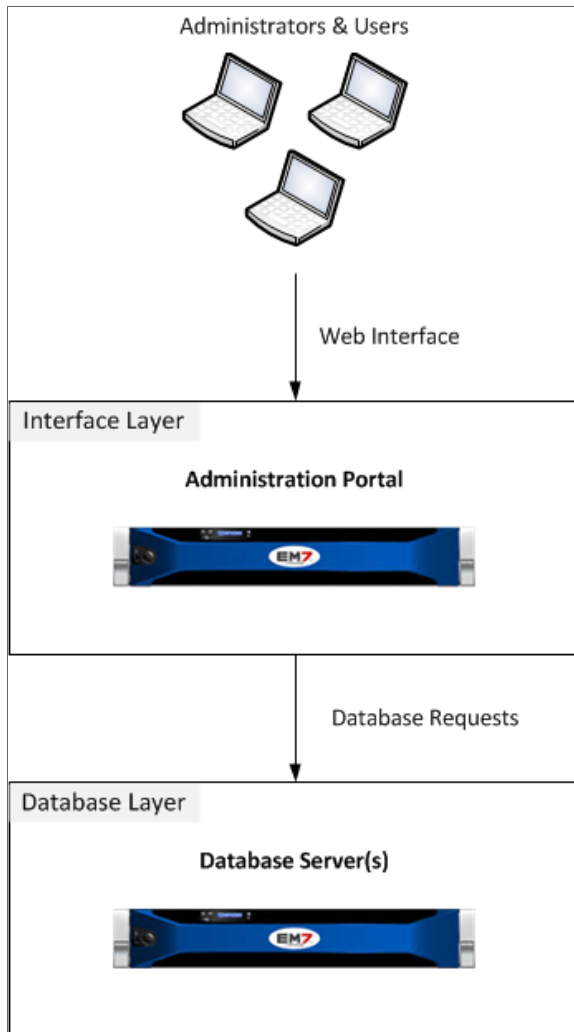
This chapter includes the following topics:

<a href="#">Single Administration Portal</a> .....	39
<a href="#">Multiple Administration Portals</a> .....	40
<a href="#">Load-Balanced Administration Portals</a> .....	41
<a href="#">API Layer Configurations</a> .....	42
<a href="#">REST API</a> .....	42
<a href="#">GraphQL API</a> .....	42

---

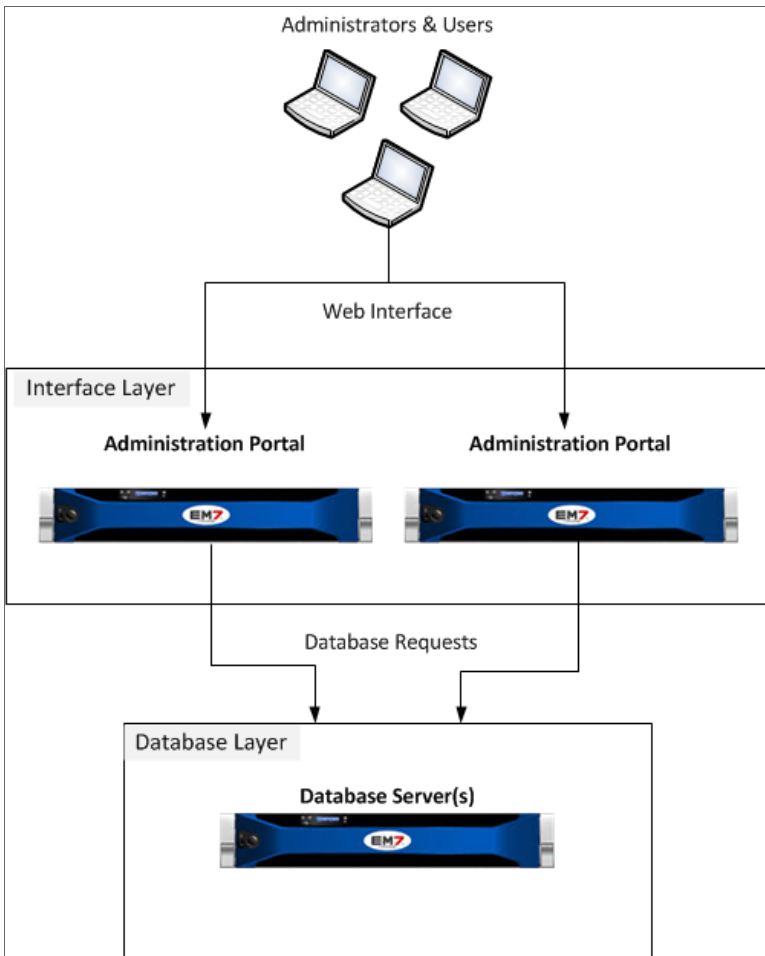
## Single Administration Portal

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.

Relationship data stored in a Graph database (the Compute 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.

Relationship data stored in a Graph database (the Compute Node cluster) is visible only via the ScienceLogic REST 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.

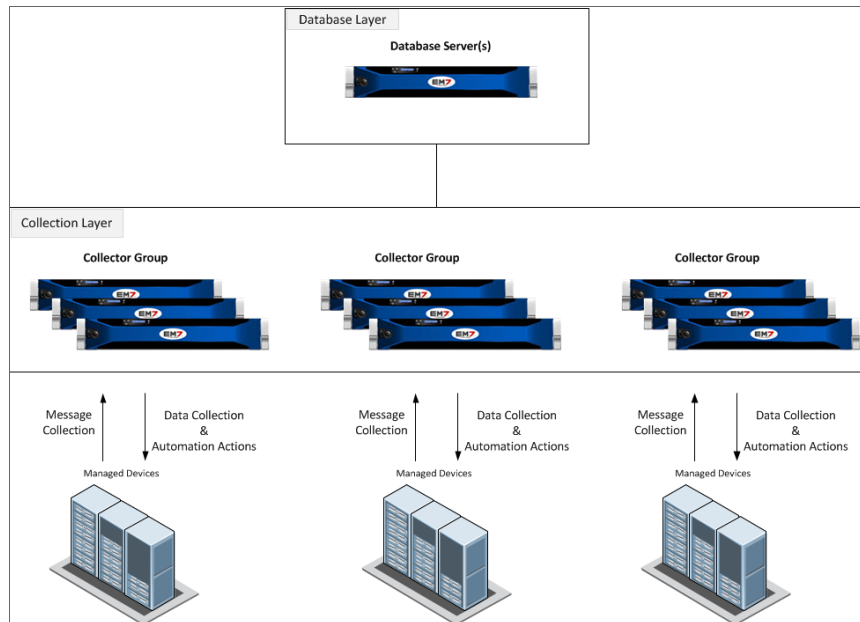
This chapter includes the following topics:

<i>Collector Groups</i> .....	44
<i>"Traditional" and "PhoneHome" Collectors</i> .....	44
<i>Using a Data Collector for Message Collection</i> .....	46
<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> .....	52

---

## 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:

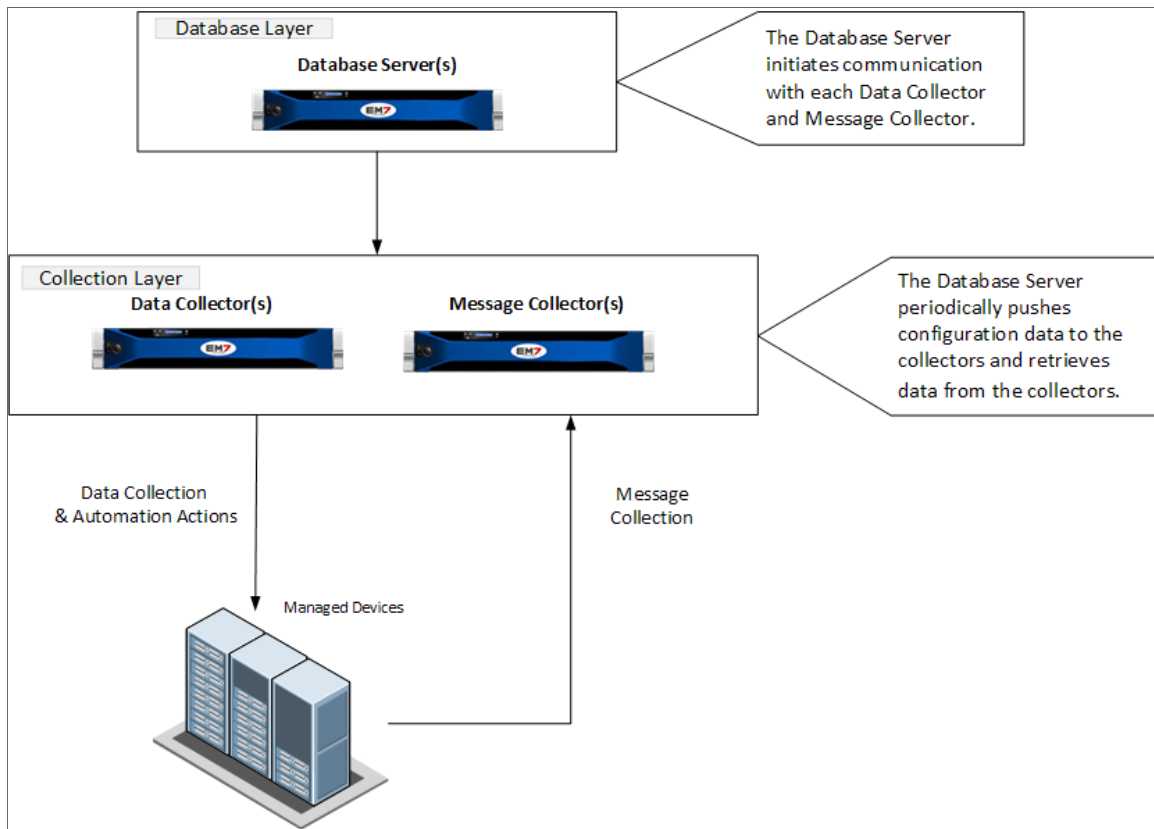
- In the main data center, a Collector Group configured with high availability that contains multiple Data Collectors and Message Collectors.
- In the second data center, a Collector Group that contains a single Data Collector that is also responsible for message collection.

---

## "Traditional" and "PhoneHome" Collectors

SL1 supports two methods for communication between Database Servers and the Data Collectors and Message Collectors in a system:


**The traditional method**, where the Database Server initiates communication with each Data Collector and Message Collector. The Database Server periodically pushes configuration data to the Data Collectors and Message Collectors and retrieves data from the Data Collectors and Message Collectors.



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



**The PhoneHome method**, where the Data Collectors and Message Collectors initiate communication with the Database Server. The Database Server then creates an SSH tunnel. The Database Server uses the SSH tunnel to periodically push configuration data to the Data Collectors and Message Collectors and retrieve data from the Data Collectors and Message Collectors.




The benefits of this method are that no firewall rules must be added on the network that contains the Data Collectors, 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.
- 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.

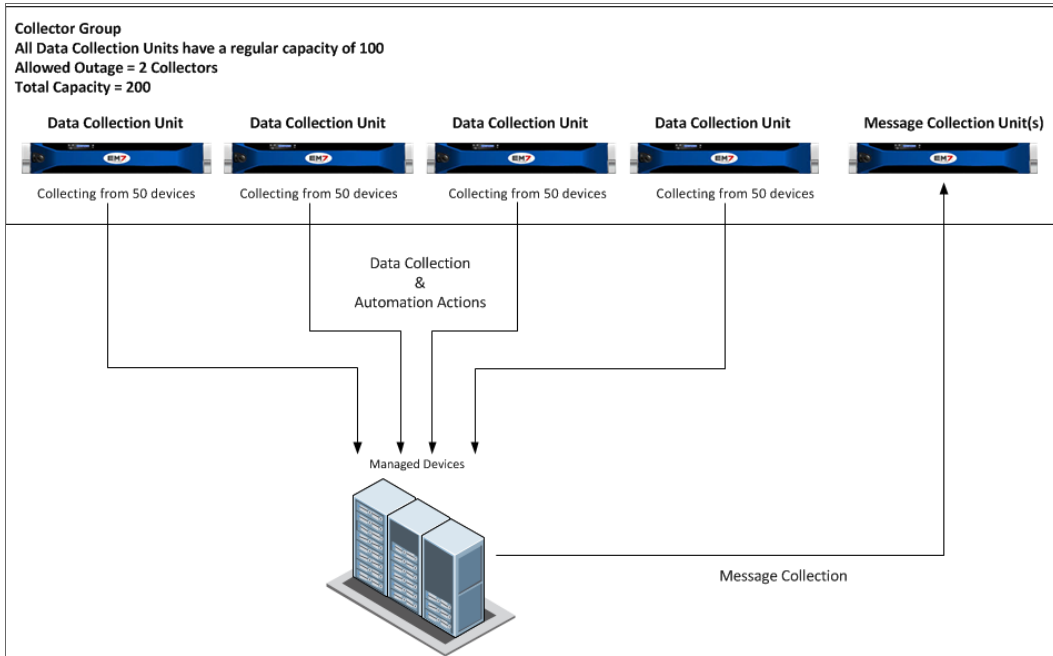
Each component device must use the same Data Collector used by its root device. For component devices, the SL1 System must keep all the component devices on the same Data Collector used by the root device (the physical device that manages the component devices). SL1 cannot distribute the component devices 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.

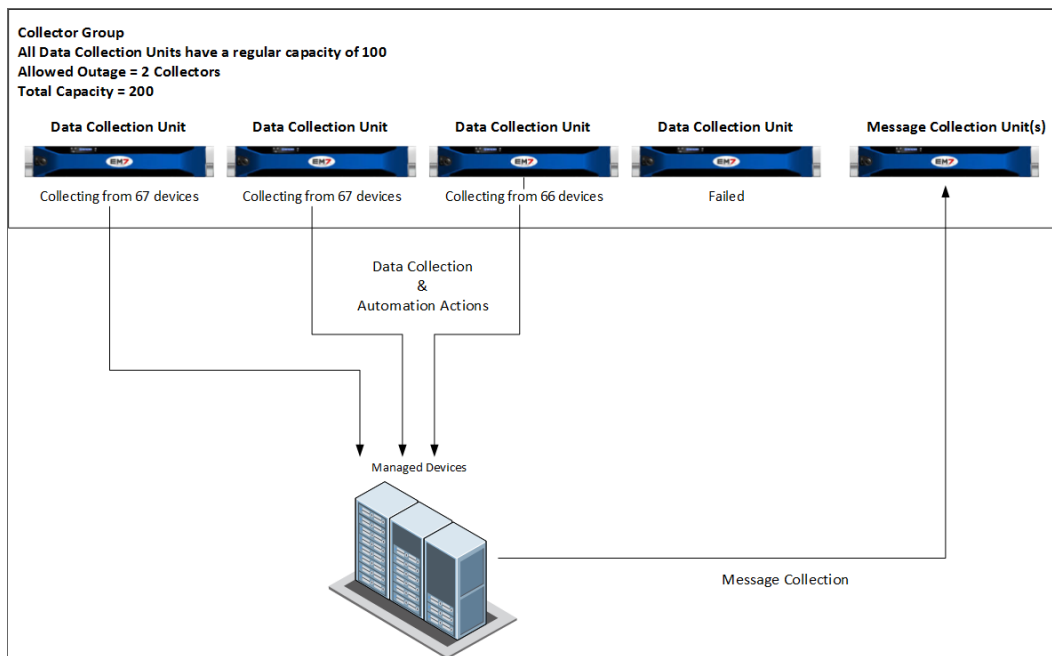
- Each collector group that is configured for high availability includes a setting for Maximum Allowed Collector Outage. This setting specifies the number of Data Collectors that can fail and data collection will continue. If more Data Collectors than the specified maximum fail simultaneously, some or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

**WARNING:** If a collector group is configured for high availability and the number of failed Data Collectors in that collector group becomes greater than the Maximum Allowed Collector Outage setting, SL1 will not failover within the Collector Group. SL1 will not collect or store any data from the devices aligned with the failed Data Collector(s) until the failure is fixed, and SL1 will generate a critical event. This is true regardless of whether the individual Data Collectors are able to collect data.

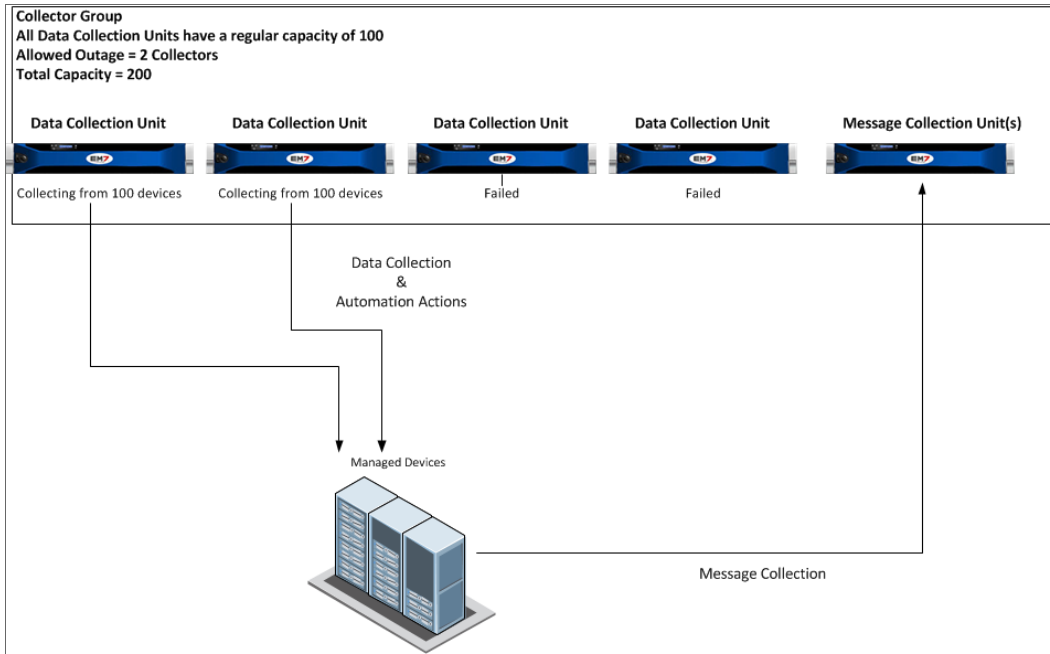
In this example, the Collector Group includes four Data Collectors. The Collector Group is configured to allow for an outage of two Data Collectors.

When all Data Collectors are available, the SL1 System evenly distributes the devices monitored by a Collector Group among the Data Collectors in that Collector Group. In this example, there are 200 devices monitored by the Collector Group, with each of the four Data Collectors responsible for collecting data from 50 devices. For simplicity, this example assumes that SL1 spends the same amount of time collecting Dynamic Application data from every device; therefore, the devices are divided evenly across the four collectors.

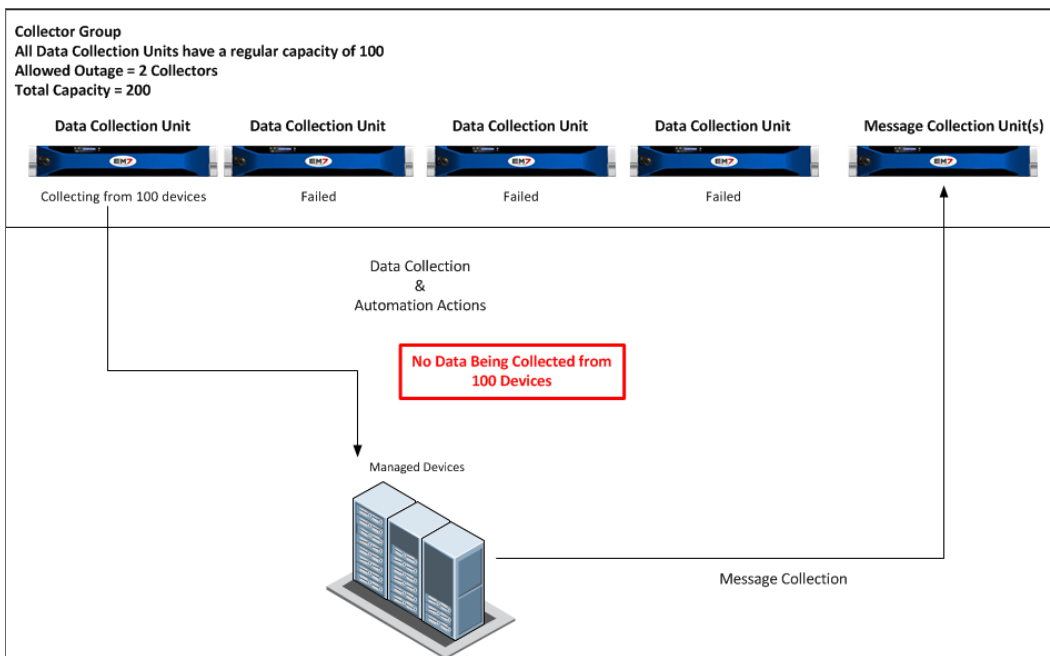
If one of the Data Collectors in the example Collector Group fails, the 50 devices that the Data Collector was monitoring are redistributed evenly between the other three Data Collectors:



If a second Data Collector in the example Collector Group fails, the 50 devices that the Data Collector was monitoring are redistributed evenly between the other two Data Collectors:

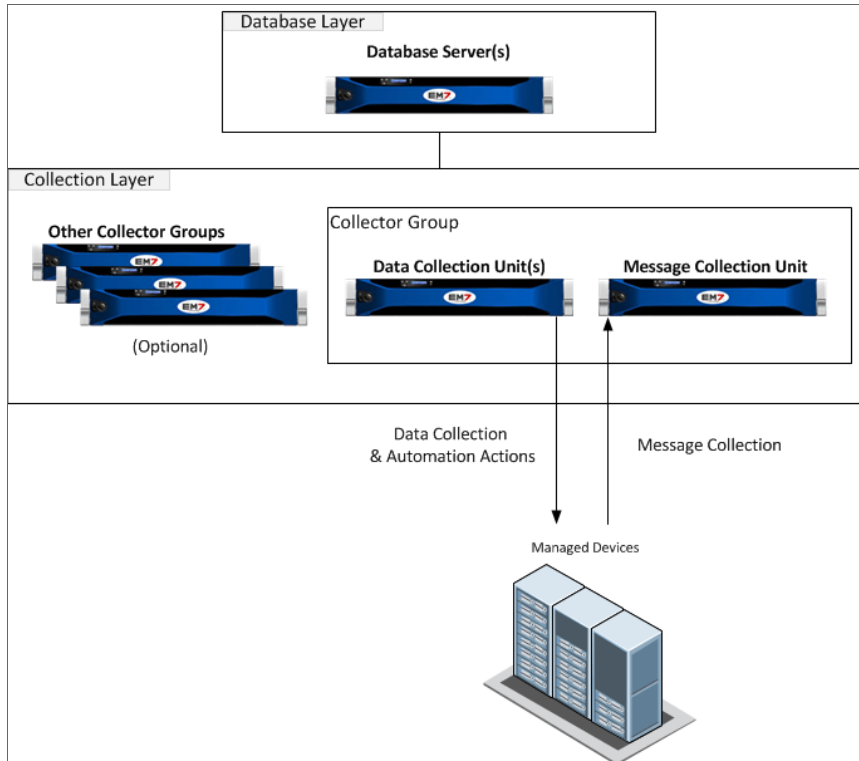


If a third Data Collector in the example Collector Group fails, the Collector Group has exceeded its maximum allowable outage. Until one of the three failed Data Collectors becomes available, 100 devices are not monitored:

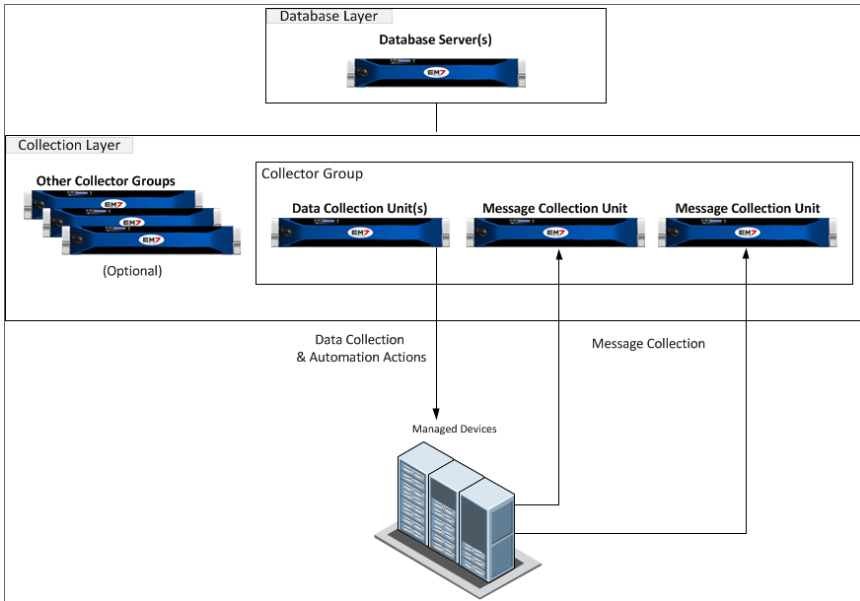


## Using Message Collectors in a Collector Group

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:



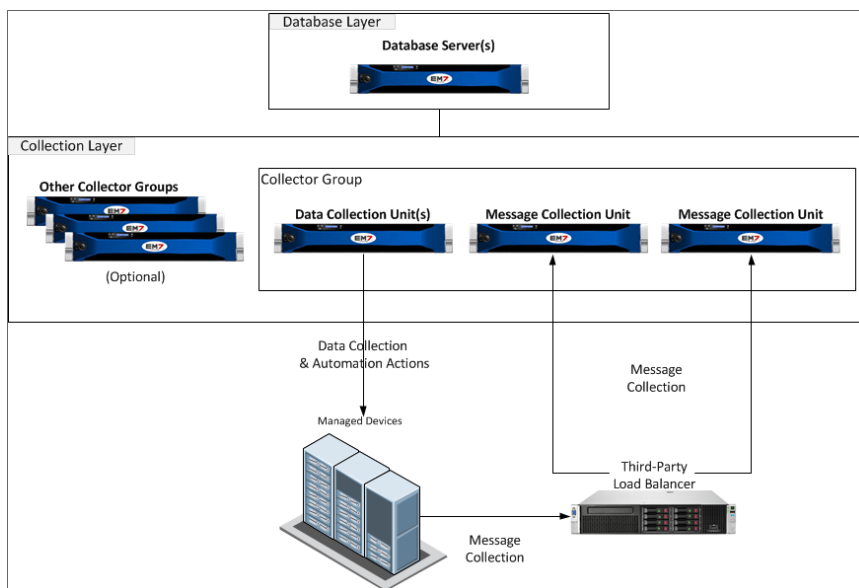
If your monitored devices generate a large amount of syslog and trap messages, a Collector Group can include multiple Message Collectors:



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.

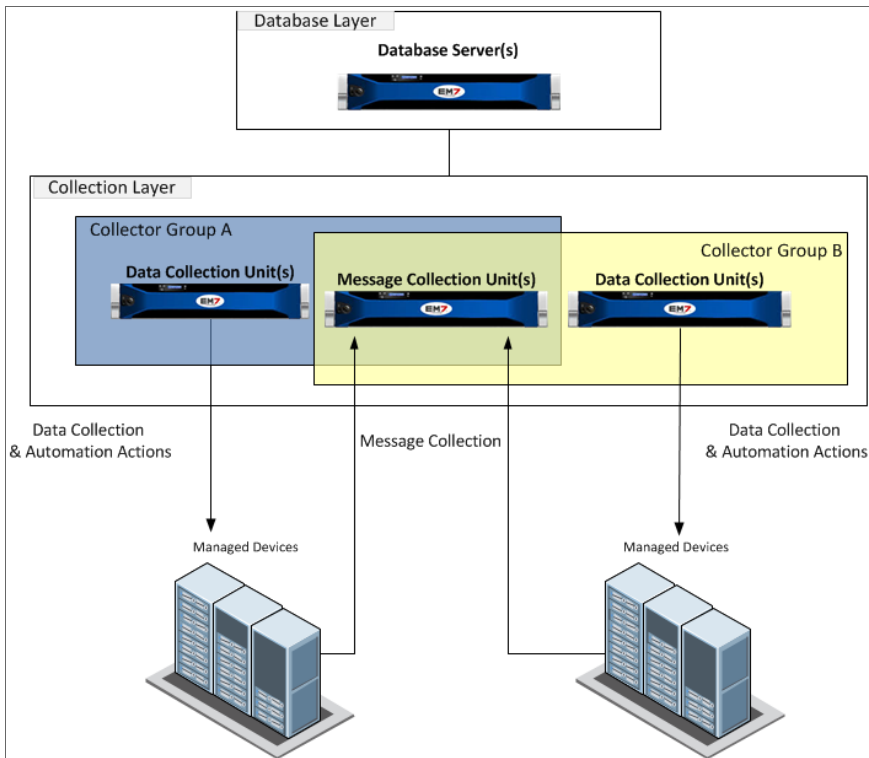
A third-party load-balancing solution can be used to distribute syslog and trap messages evenly among the Message Collectors in a Collector Group:





**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 SL 1 System in every Collector Group in your SL 1 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 Customer Portal: <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 includes the following topics:

<i>Caveats</i> .....	56
----------------------	----

---

## 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 - 2020, 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](mailto:legal@sciencelogic.com)



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

International: +1-703-354-1010