



Security Features

SL1 version 10.1.4

Table of Contents

Introduction	4
Who Should Read This Manual?	4
Built-In Security for Appliances and Data	5
Hardened Operating System	6
Limited Open Ports	6
Firewalls and White Lists	6
Hardened Configuration on Each Appliance	7
Root Access	7
API	7
All-In-One	7
Administration Portal	8
Database Server	8
Data Collectors and Message Collectors	8
Multiple Tenancy and Segregation of Duties	9
Account Types	9
Access Keys	9
Segregation by Organization	10
Credential Management	10
User Policies	11
Protection Against Injections and Cross-Site Scripting	12
Operating System Scan	12
Data Integrity	12
Backups	13
Disaster Recovery and High Availability	13
Audit Logs	13
Manage the Security of Your Network	15
Monitoring IDS, Firewalls, and Security Hardware	16
Security Events	16
Monitoring Changes to Device Configuration	16
Monitoring for Illicit Behavior	17
Blueprinting Windows Services	17
Blueprinting System Processes	17
Blueprinting DNS	18
Monitoring Open Ports	18
Monitoring Bandwidth Usage	18
Monitoring Hardware Performance	19
Managing Patches and Hot Fixes	21
Using Standard Deviation To Calculate "Normal" Conditions and Abnormal Conditions	21
Using Run Book Automation to Automate Responses to Security Events	21
Reports	22
Proxied Web Services	23
Security Settings	24
Access Control	25
Authentication	30
Multiple Tenancy and Segregation of Duties	31
Protection of Shared Content	33
Data Integrity	33
Security Events	34
Monitoring Changes to Device Configuration	35
Monitoring for Illicit Behavior	35

Blueprinting DNS, System Processes, and Windows Services	36
Monitoring Open Ports	37
Monitoring Bandwidth Usage	37
Monitoring Hardware Performance	39
Monitoring Patches and Hot Fixes	40
Using Run Book Automation to Automate Responses to Security Events	41
Reports	41
Proxied Web Services	41
Audit Logs	42

Chapter

1

Introduction

Overview

SL1 addresses two major aspects of system and network security:

- SL1 appliances are lean, hardened, and configured for maximum security.
- SL1 integrates with and complements your existing network and system security (policies, software, and hardware) and adds powerful new features to help you monitor and maintain all the devices in your network.

Who Should Read This Manual?



This manual is intended for administrators who manage ScienceLogic systems and any other users who are involved in system and network security.

Built-In Security for Appliances and Data

Overview

SL1 is specifically designed to provide a secure environment for monitoring your network. This chapter describes the security features included in SL1.

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 the menu options, click the Advanced menu icon ().

This chapter includes the following topics:

<i>Hardened Operating System</i>	6
<i>Limited Open Ports</i>	6
<i>Firewalls and White Lists</i>	6
<i>Hardened Configuration on Each Appliance</i>	7
<i>Multiple Tenancy and Segregation of Duties</i>	9
<i>Protection Against Injections and Cross-Site Scripting</i>	12
<i>Operating System Scan</i>	12
<i>Data Integrity</i>	12
<i>Audit Logs</i>	13

Hardened Operating System

Each SL1 appliance includes a hardened version of the Linux operating system that is designed and engineered for high-security deployments. For example:

- The operating system installs only those software packages that are required.
- The operating system operates under a strictly limited or minimum number of ports open.
- The operating system updates are owned by ScienceLogic
- All data in transit between SL1 appliances is encrypted.
- Insecure protocols such as, FTP, Telnet, and Rlogin, are disabled by default

Limited Open Ports

Each SL1 appliance uses a strictly limited number of ports. For the list of current required ports, see the manual *Installation*.

Firewalls and White Lists

SL1 includes multiple features that tightly control access to SL1 appliances.

Each SL1 appliance includes a built-in firewall. For Data Collectors and Message Collectors, the firewall allows communication to the central database only from other SL1 appliances. Even communication between other SL1 appliances is limited with the ScienceLogic "store and retrieve" architecture.

The built-in firewall for each SL1 appliance allows users to create "white lists". For example, because each Data Collector can be placed in an unsecured part of the network, the built-in firewall defines a very limited number of communications that will be accepted by the Data Collector. These limited communications make up the "white list" for each Data Collector. Devices that are not part of the white list cannot communicate with a Data Collector.

SL1 also allows administrators to create white lists for web access. System administrators can specify that users can connect to the user interface only from specified IP addresses. System administrators can also create "black lists" and prohibit web access (to a SL1 System) based on user name, the user's IP address, or both.

For example, suppose the system has performed the initial discovery session to discover the devices in your network. Suppose a device that was not discovered tries to connect and send information to a Data Collector, Message Collector, or All-In-One Appliance. The system will block the connection and incoming data. ScienceLogic firewalls and white lists prevent rogue devices from connecting to a SL1 appliance and sending information. SL1 Appliances accept asynchronous traffic only from devices that have been discovered in a discovery session.

Hardened Configuration on Each Appliance

Each SL1 appliance is configured to maximize security and minimize security risks.

Root Access

The ScienceLogic platform has been designed to operate without root privileges and command line access. By customer request, root access is enabled on all appliances. ScienceLogic strongly recommends that the root account be secured and enabled only when necessary for troubleshooting or custom configuration tasks. To disable root access, see the *System Administration* manual.

API

- The API provides a single point of communication between an external site and the Database Server. The API insulates the Database Server; the external site cannot communicate directly with the Database Server but can communicate with the API.

All-In-One

- The All-In-One Appliance supports TLSv1.2 services for secure sign-on.
- Administrators can require that browser access to the All-In-One Appliance be authenticated and sent over HTTPS.
- Administrators can require that authentication checks the originating IP address, to control where client requests originate during user login.
- Users can be authenticated via LDAP, Active Directory (using the LDAP/Active Directory integration features), SSO, or by local SL1 authentication.
- SL1 has built-in support for multi-factor authentication (MFA) using RSA SecurID.
- Administrators can configure the All-In-One Appliance to use DoD certificates or your own certificates. You can install server-side certificates on the ScienceLogic All-In-One Appliance and then authenticate access to the All-In-One Appliance with a CAC or a client-side certificate associated with a user's web browser.
- System administrators can also define blacklisting and whitelisting to further control access to the All-In-One Appliance.
- All-In-One Appliances use unique session cookies to validate client requests.
- System administrators can finely customize user access to the All-In-One Appliance to create a multi-tenant environment that meets the needs of management, system administrators, NOC personnel, and customers and clients. Each type of user can view and access only the information that is relevant to them. System administrators can define user access per user or create access templates to apply to each type of user.

Administration Portal

- The Administration Portal supports TLSv1.2 services for secure sign-on.
- Administrators can require that browser access to the Administration Portal be authenticated and sent over HTTPS.
- Administrators can require that authentication checks the originating IP address, to control where client requests originate during user login.
- Users can be authenticated via LDAP, Active Directory, SSO, or by local SL1 authentication.
- Administrators can configure the Administration Portal to use DoD certificates or your own certificates. You can install server-side certificates on an Administration Portal appliance and then authenticate access to the Administration Portal with a CAC or a client-side certificate associated with a user's web browser.
- System administrators can also define blacklisting and whitelisting to further control access to the Administration Portal appliance.
- Administration Portal appliances use unique session cookies to validate client requests.
- System administrators can finely customize user access to the Administration Portal to create a multi-tenant environment that meets the needs of management, system administrators, NOC personnel, and customers and clients. Each type of user can view and access only the information that is relevant to them.
- 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 Database Server initiates one-way communication with Data Collectors and Message Collectors through "data pull". Data pull uses a single TCP outbound port.
- The Database Server receives inbound communication only from the Administration Portal. All connections between the Database Server and the Administration Portal are encrypted in both directions.
- The Database Server uses X509 certificates with RSA 2048 encryption to establish sessions and authenticate communication with other SL1 appliances.

Data Collectors and Message Collectors

- All communication between Data Collectors and Message Collectors and the Database Server is initiated by the Database Server and is unidirectional.
- On the Data Collectors and Message Collectors, only port 7707 is used for communication with the Database Server. The Database Server uses X509 certificates with RSA 2048 encryption to establish sessions and authenticate communication with Data Collectors and Message Collectors.
- Communication between the Database Server and each Data Collector uses TLS v1.2 negotiated through the database connection protocol. Both client and server software use the secure binaries and libraries, which are built with OpenSSL and is not vulnerable to the Heartbleed vulnerability.

- SL1 uses the DHE-RSA-AES256-SHA cipher for OpenSSL. Thus, SL1 uses ephemeral Diffie-Hellman key exchange with RSA key authentication, 256-bit AES, and SHA-2 message digests. The Data Collector uses a 2048-bit RSA key for certificate verification.
- By default, there is no way for the Data Collector(s) or Message Collector(s) to initiate communication with the Database Server(s).
- Data Collectors and Message Collectors store information about only the devices that the Collector Group is configured to monitor. If a Collector Group is not configured to monitor a given device in the platform, the database on the Data Collector or Message Collector will not contain any information about that device.
- The SL1 agent collects data from the device on which it is installed and transfers that data to the appliance performing message collection (usually, a Message Collector) using the HTTPS protocol. TCP port 443 must be open between the device on which the agent is installed and the Message Collector. To transfer data gathered by the agent to SL1, a Data Collector polls the Message Collector with the HTTPS protocol.

Multiple Tenancy and Segregation of Duties

Account Types

ScienceLogic administrators can enforce segregation of duties using account types. The platform includes the following account types:

- **Administrator.** A user who is automatically assigned all access privileges in the platform.
- **User.** A user who must be assigned access privileges in the platform. Regardless of assigned Access Keys:
 - Users cannot modify administrator accounts.
 - Users cannot make themselves or another user an administrator.
 - Users cannot grant or remove Access Keys that they have not been granted.
- **External Contacts.** People who can receive Emails sent from the platform but do not have user accounts.

For details on creating and configuring user accounts, see the manual **Organizations and Users**.

Access Keys

ScienceLogic administrators can enforce segregation of duties using access privileges. In SL1, **Access Keys** allow administrators to define highly granular, customized access privileges for regular users. Access keys define what each user can view in the platform and which actions each user can perform.

An access key has two parts: the access key definition and the access hooks that are aligned with the access key.

- An **access hook** is a granular privilege. For example, an access hook might be called "View Asset." This asset hook would allow a user to view access records. The access hook would not allow a user to create, edit, or delete an asset record.

Access hooks cannot be directly granted to a user account. To grant the "View Asset" privilege to a user, the access hook must be included in an access key .

- An **access key** is a group of one or more access hooks. When you associate an access key with a user account, the user is granted all the access hooks within the key. For example, you could create an access key called "Manage Asset". That access key could include three hooks: view an asset record, edit an asset record, and create a new asset record.

For details on defining and using access keys, see the manual **Access Permissions**.

Segregation by Organization

ScienceLogic administrators can use organizations to enforce segregating of duties in the platform.

Regardless of access keys, accounts of type "user" can access only pages and actions associated with their organization. For example:

- Suppose your organization includes three regional offices. Suppose you define three organizations: Northeast, Headquarters, and West Coast.
- Suppose each organization includes the devices located at the corresponding office.
- Now suppose the account "JohnDoe" is of type "user" and is a member of the organization "West Coast". User JohnDoe would be able to view and act upon only devices that are included in the organization "West Coast". User JohnDoe would not be able to view or act upon the hardware at the other offices.
- For this reason, SL1 allows you to assign each user a primary organization and an optional additional organization.
- Now suppose that user "JohnDoe" needs to view the status of a device at headquarters. If you add a secondary organization to JohnDoe's account information, that user will now be able to view and act upon all the devices in the "Headquarters" organization.

NOTE: You can still use Access Keys to limit the access of each user, even within his or her own organization.

Credential Management

Credentials are access profiles (usually user name, password, and any additional information required for access) that allow the platform to retrieve information from devices and from software applications on devices. These profiles allow the platform to access external systems while maintaining the security of the access accounts. Users who need the platform to retrieve data from these external systems see only the name of the credential, not the user name, password, and network information. All credentials are encrypted on the SL1 system.

You can use Access Keys to restrict access to the credential definitions.

To further support multi-tenancy, SL1 allows you to align each credential with one, multiple, or all organizations. You can also align a credential with no organizations. When you align an organization with a credential, you control who can view details about the credential, who can view the name of the credential, and who can apply the credential.

NOTE: When you align an organization with a credential, you are restricting only the users who can view and assign the credential. You are not restricting the devices and actions that can be associated with the credential. For example, you can align a credential only with the organization "Operations" but assign the credential to a device in the "Finance" organization.

Credentials that are aligned with an organization have the following behavior:

- For each credential that is aligned with an organization, only administrators and users who are members of the aligned organization will be able to see the credential in the **Credential Management** page.
- In the user interface, in any field or column that **displays the name of the credential**, users who are **not** members of the aligned organization will **not** see the credential name. Instead, these users will see either a dash character (-) or the text "Restricted Credential".
- In the user interface, in any list from which users can **select a credential**, users who are **not** members of the aligned organization will **not** see the credential as an entry in the list.
- In the user interface, in any page where the credential has already been assigned, users who are **not** members of the aligned organization will see only the name "Restricted Credential".
- In the user interface, in any page where the credential has already been assigned, users who are **not** members of the aligned organization can save the page and maintain the credential. The credential will still appear to that user as "Restricted Credential".
- In the user interface, in any page where the credential has already been assigned, users who are **not** members of the aligned organization can change the credential to a credential aligned with their organizations. However, those users cannot change the credential again and re-assign the "Restricted Credential". The entry for "Restricted Credential" is removed from the list of possible credentials.

For details on credentials, see the manual *Discovery and Credentials*.

User Policies

You can easily manage multiple user accounts using **User Policies**. User Policies allow you to define a custom set of account properties and key privileges (from the **Account Permissions** page) and then save them as a policy, for reuse. When you create a user account, you can use the User Policy to quickly apply settings to the new account. User policies can help you to enforce roles and consistent SOPs when creating user accounts in SL1.

User Policies have a dynamic relationship with their member user accounts. You can make a change to a user policy and the platform will automatically update the account settings for each member account. For example, suppose you create a user account called "John Doe" on the first of the month and use the user policy named "NOC users" to create the user account. Suppose you create another user account called "Jane Smith" on the fifth of the month and use the user policy "NOC users". Suppose on the 15th of the month, you add an additional Access Key to the "NOC users" policy. That additional Access Key will appear in the account for John Doe and Jane Smith as soon as the "NOC users" policy is saved.

For details on user policies, see the manual *Organizations and Users*.

Protection Against Injections and Cross-Site Scripting

ScienceLogic uses automated penetration tests that are included in the Quality Assurance process for each release. ScienceLogic uses a variety of scanning tools such as Nessus and Burp to detect vulnerabilities in the user interface and platform. The tools specifically look for the OWASP Top Ten Vulnerabilities such as, various XSS, SQL injection, broken authentication, and session management.

ScienceLogic validates the Top Ten Security Configurations using the most current version of Nessus. Items such as validating functional level access controls and updating modules with known vulnerabilities are currently checked in both manual and automated fashion.

Bi-annually, ScienceLogic contracts with independent, third-party penetration testers to provide independent penetration testing. These third parties conduct bi-annual tests on current and upcoming SL1 releases with the intent of constantly increasing our security posture in response to the expanding set of known vulnerabilities identified by industry experts. Additional corporate and SaaS penetration testing is performed on an annual basis.

In addition, ScienceLogic customers often perform additional penetration tests tailored to their unique requirements. These companies report their findings back to ScienceLogic's security team for remediation.

ScienceLogic responds to the penetration tests performed by ScienceLogic and ScienceLogic customers with the appropriate fixes and updates.

ScienceLogic customers may request an executive summary of ScienceLogic's penetration tests, but ScienceLogic customers must have a fully executed Non-Disclosure Agreement (NDA) that is signed by both parties to have access to this documentation.

Operating System Scan

The SL1 operating system is scanned daily to identify new vulnerabilities. Potential vulnerabilities are evaluated by the security team to determine the likelihood and impact of exploit on SL1. If the security team identifies high and critical vulnerabilities, they are escalated to engineering for remediation within the defined time frame.

Data Integrity

SL1 performs the following to ensure the integrity of data:

- Leverages the xfs journaling file system type for added resiliency
- Automatically detects and self-corrects most sources of database corruption
- Validates data before insert into the database.
- Routinely checks the status of all required platform processes and automatically restarts those that have stopped

- Collects data once, stores the data in an intelligent manner, and uses the stored data in multiple parts of the platform, without requiring the compute burden of multiple collection sessions. For example, SL1 collects bandwidth data once, but performs normalization for different timespans, using the same set of collected data.
- Does not require a dedicated Database Administrator. The Database Server is designed to be self-pruning and self-correcting.

Backups

SL1 includes two levels of built-in, scheduled backups:

- **Configuration Backup.** Stores a local copy of the core database tables that are required to restore an instance of the platform, and optionally transfers the copy to an external system.
 - The platform automatically launches this backup every day, at the time you specify.
 - Optionally, the platform can automatically transfer this backup every day to an external system, at the time you specify. Configuration backups can be saved to a remote server. Multiple secure protocols are supported for file transfers such as, SFTP site, Secure NFS mount, or SMB 3.0 mount.
- **Full Backup.** For instances of the platform in small-to-medium businesses, full backup makes a full backup of the Database Server.
 - Full backup includes all configuration data, performance data, and log data.
 - After backup, the platform automatically copies the compressed file to a remote system. The compressed backup can be copied to an FTP site, SFTP site, NFS mount, or SMB mount.
 - To conserve space, the compressed file is then removed from the Database Server appliance or All-In-One Appliance.

For details on backing up your instance of the platform, see the manual on **System Administration**.

Disaster Recovery and High Availability

SL1 includes multiple options for high-availability, including high-availability for Database Servers and Data Collectors. For details on these configurations, see the manual on **Architecture**.

SL1 also includes a disaster recovery option for Database Servers and All-In-One Appliances. For details on disaster recovery, see the manual on **Disaster Recovery**.

Audit Logs

SL1 includes detailed audit logs of all platform activity, including:

- **System Logs** of all platform activity. For example, starting and stopping key processes, backing up key databases, purging old database records, and other maintenance activities.



- **Audit Logs** of all actions in SL1 that are generated by user or managed entities, including all logins by users; creating, editing, or deleting of any data, reports, or policies; and all events.
- **Access Logs** of all user activity within SL1, including logins, logouts, notifications, and actions executed.
- **Device Logs** that include all messages collected from a device by SL1. These messages include information from syslog, internal alerts, traps, Dynamic Applications, and email messages.

Manage the Security of Your Network

Overview

SL1 easily integrates with existing security hardware, software, and policies and adds additional security features and automation. The following sections will describe the added security that SL1 can add to your network.

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 the menu options, click the Advanced menu icon ().

This chapter includes the following topics:

<i>Monitoring IDS, Firewalls, and Security Hardware</i>	16
<i>Security Events</i>	16
<i>Monitoring Changes to Device Configuration</i>	16
<i>Monitoring for Illicit Behavior</i>	17
<i>Blueprinting Windows Services</i>	17
<i>Blueprinting System Processes</i>	17
<i>Blueprinting DNS</i>	18
<i>Monitoring Open Ports</i>	18
<i>Monitoring Bandwidth Usage</i>	18
<i>Monitoring Hardware Performance</i>	19
<i>Managing Patches and Hot Fixes</i>	21
<i>Using Standard Deviation To Calculate "Normal" Conditions and Abnormal Conditions</i>	21

<i>Using Run Book Automation to Automate Responses to Security Events</i>	21
<i>Reports</i>	22
<i>Proxied Web Services</i>	23

Monitoring IDS, Firewalls, and Security Hardware

As a Manager of Managers, SL1 can monitor security hardware, including intrusion detection systems and firewalls. SL1 can generate alarms (called **events**) about the security hardware. SL1 can also create alarms about network information that the security hardware collects.

For example, SL1 can generate an alarm when a Cisco IPS Sensor indicates a change in security status.

For more information about events, see the manual **Events**.

Security Events

Using policies, Dynamic Applications, and retrieved log messages, SL1 can monitor security on each managed device and generate alerts about security breaches.

For example, SL1 can generate an alert when a user opens a chassis on a Dell server.

SL1 ships with many pre-defined security events. You can also easily customize your system to monitor the security events and generate alerts that are most useful to your operations.

You can define customized events specifically for your operations.

You can define escalation policies based on events.

For more information about events, see the manual **Events**.

Monitoring Changes to Device Configuration

SL1 can automatically create an asset record for each discovered device. SL1 uses information gathered from the device during initial discovery, nightly auto-discovery, and from Dynamic Applications to populate each asset record.

Asset records can contain the following information:

- The name, make, and model of the device
- The serial number of the device
- The function and status of the device
- Networking information, like host ID, IP address, and DNS server for the device
- Hardware information like amount of memory, CPU, and BIOS or EPROM version

- The physical location of the device
- Vendor information for the device, including PO or check number, warranty policy, and service policy
- A description of the network interface
- A description of each hardware component (if applicable)
- A description of installed software (if applicable)

In SL1, you can define a policy for one or more asset fields that you want to monitor. If the value of a monitored asset field changes, the platform will alert you to the change.

For more information about monitoring asset records, see the manual **Asset Management**.

Monitoring for Illicit Behavior

SL1 can monitor devices for misuse and illicit behavior. SL1 can alert users when:

- illicit DNS domains or DNS records are discovered on a Domain Name server
- illicit processes are found on one or more devices
- processes are running as "root" or "su" on one or more devices
- illicit Windows services are found on one or more devices
- illicit ports are found open on one or more devices

For more information about monitoring DNS, system processes, or Windows services, see the manual **Device Management**.

Blueprinting Windows Services

You can define policies that monitor the Windows services on each Windows server. You can create a "blueprint" of the services and settings that should run on each Windows server. If SL1 detects a change to those services and settings, SL1 can automatically restore the server to its blueprint of services and settings.

For more information about monitoring Windows services, see the manual **Monitoring Windows Systems (WMI) or Monitoring Windows Systems (PowerShell)**.

Blueprinting System Processes

You can define policies that monitor the system process on a server. You can create a "blueprint" of the processes and process settings on each server. If SL1 detects a change to those processes or their settings, SL1 can generate an alert, specifying the exact changes to the blueprint.

For more information about monitoring system process, see the manual **Device Management**.

Blueprinting DNS

You can define policies that monitor DNS servers. You can create a "blueprint" of the records that should exist on the server. If SL1 detects a change to those records, the platform can generate an alert, specifying the exact change to each changed record.

For more information about monitoring DNS, see the manual *Device Management*.

Monitoring Open Ports

SL1 helps you limit open ports on each managed device and reduce your vulnerability to attacks.

During initial discovery, SL1 scans each device to discover all the open ports on the device. SL1 also scans each device every night and looks for any newly opened ports.

- If SL1 discovers an illicit port, SL1 can alert users. System administrators can define which ports to consider "illicit".
- If SL1 discovers a newly opened port, SL1 can alert users of the new port.

For more information about monitoring open ports, see the manual *Device Management*.

Monitoring Bandwidth Usage

SL1 discovers each network interface on each managed device. SL1 can then monitor bandwidth usage and performance for each network interface.

You can define thresholds for bandwidth usage and performance. These thresholds will alert you of unusual or illicit bandwidth usage.

You can define global thresholds for the following bandwidth statistics. If any network interface on any device exceeds these thresholds, the platform will alert users.

You can also define each of these thresholds per interface. The interface will use the custom thresholds instead of the global thresholds. If the interface exceeds these thresholds, the platform will alert users.

- **Inbound Bandwidth.** Rate at which a device is using network bandwidth for inbound traffic, in Mb per second. If a device uses bandwidth for inbound traffic at a greater rate than the specified rate, SL1 generates an event.
- **Outbound Bandwidth.** Rate at which a device is using network bandwidth for outbound traffic, in Mb per second. If a device uses bandwidth for outbound traffic at a greater rate than the specified rate, SL1 generates an event.

- **Inbound Percent.** Percentage of network bandwidth being used by each device for inbound traffic. If a device uses more than the specified percentage of network bandwidth for inbound traffic, SL1 generates an event.
- **Outbound Percent.** Percentage of network bandwidth being used by each device for outbound traffic. If a device uses more than the specified percentage of network bandwidth for outbound traffic, SL1 generates an event.
- **Inbound Errors.** Threshold for number of inbound packet-errors per polling interval per interface. Errors occur when packets are lost due to hardware problems such as breaks in the network or faulty adapter hardware. If the number of inbound packet-errors on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface.
- **Outbound Errors.** Threshold for number of outbound packet-errors per polling interval per interface. If the number of outbound packet-errors on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface.
- **Inbound Error Percent.** Threshold for percentage of inbound packet-errors per polling interval per interface. If the percentage of inbound packet-errors on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface. The percentage is calculated as inbound errors/total inbound packets.
- **Outbound Error Percent.** Threshold for percentage of outbound packet-errors per polling interval per interface. If the percentage of outbound packet-errors on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface. The percentage is calculated as outbound errors/total outbound packets.
- **Inbound Discards.** Threshold for number of inbound packet-discards per polling interval per interface. Discards occur when an interface receives more traffic than it can handle (either very large message or many messages simultaneously). Discards can also occur when an interface has been specifically configured to discard. For example, a user might configure a router's interface to discard packets from a non-authorized IP. If the number of inbound packet-discards on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface.
- **Outbound Discards.** Threshold for number of outbound packet-discards per polling interval per interface. If the number of outbound packet-discards on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface.
- **Inbound Discard Percent.** Threshold for percentage of inbound packet-discards per polling interval per interface. If the percentage of inbound packet-discards on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface. The percentage is calculated as inbound discards/total inbound packets.
- **Outbound Discard Percent.** Threshold for percentage of outbound packet-discards per polling interval per interface. If the percentage of outbound packet-discards on an interface exceeds the threshold specified in this field, SL1 will generate an event for that interface. The percentage is calculated as outbound discards/total outbound packets.

For more information about monitoring bandwidth usage, see the manual **Device Management**.

Monitoring Hardware Performance

SL1 allows you to monitor hardware performance, to detect intrusions or illicit use of hardware resources.

SL1 can monitor each device to determine if:

- network connections to the device are unusually slow
- the operating system or CPU on a device is being over-taxed
- one or more file systems on a device are filling up unexpectedly
- one or more CPUs on a device are being over-taxed
- memory or virtual memory on a device is being over-taxed

You can accept the default values or customize these values.

You can define global thresholds for the following hardware statistics. If any device exceeds these thresholds, the platform will alert users.

You can also define each of these thresholds per device. The device will use the custom thresholds instead of the global thresholds. If the device exceeds these thresholds, the platform will alert users.

- **System Latency.** During polling, SL1 initially pings monitored devices. The value in this field is the maximum number of milliseconds for the device to respond to the platform's ping (round-trip time/2). When the latency threshold is exceeded, SL1 generates an event for that device.
- **System Availability.** During polling, SL1 monitors devices for availability. Availability means the device's ability to accept connections and data from the network. The value in this field is the percent availability required of each device. When a device falls below this level of availability, SL1 generates an event for that device.
- **File System Warning .** Threshold that will trigger a "low disk space" event. The default threshold is 85%. When a device has used more disk space than the specified percentage, SL1 will generate an alert with a status of "major".
- **File System Critical.** Threshold that will trigger a "low disk space" event. The default threshold is 95%. When a device has used more disk space than the specified percentage, SL1 will generate an alert with a status of "critical".

In addition, SL1 includes Dynamic Applications to monitor CPU performance and memory performance. SL1 includes such Dynamic Applications for each major operating system and for generic SNMP.

- For each device that uses a CPU Dynamic Application:
 - When a device exceeds the "CPU Utilization High" threshold, SL1 will alert users.
- For each device that uses a memory Dynamic Application:
 - When a device exceeds the "Physical Memory Utilization High" threshold, SL1 will alert users.
 - When a device exceeds the "Swap Memory Utilization High" threshold, SL1 will alert users.

For more information about device thresholds and device hardware, see the manual **Device Management**.

Managing Patches and Hot Fixes

During initial discovery and during the nightly update, SL1 searches each device and determines the software installed on each device.

From the list of all discovered software, you can generate Software Exclusion Reports. These reports can help administrators manage patches and software versions. Among other information, a Software Exclusions Report displays the following:

- Name of the software title and the date the report was generated
- List of all devices in SL1 that have the software installed
- List of all devices in SL1 that don't have the software installed. SL1 includes only appropriate devices in this report. For example, Linux servers would not appear in a report for a Windows patch.

For more information about installed software, see the manual *Device Management*.

Using Standard Deviation To Calculate "Normal" Conditions and Abnormal Conditions

SL1 allows you to examine a value retrieved from a device and compare that value to "normal" values for the hour of day on that day of week. SL1 will trigger an alert if the current retrieved value falls outside the range of "normal" values for the hour of day on that day of week.

Standard deviation allows you to monitor security based on the real-life parameters of your environment. For example, if your organization operates Monday–Friday, from 8:00 AM until 5:00 PM, "normal" values during work hours will differ from "normal" values during the weekend or late at night.

Some possible uses for the deviation function are:

- **Determining if an application is functioning properly.** For example, if a log file for an application begins to grow at a rate outside the "normal" range for Sunday at 03:00, you can trigger an alert to determine if there is a problem with the application.
- **Monitoring security.** For example, if bandwidth usage on a Saturday at 23:00 exceeds the normal activity, you can trigger an alert that indicates that your network might have been compromised.

For details on standard deviation, see the manual *Dynamic Application Development*.

Using Run Book Automation to Automate Responses to Security Events

SL1 includes automation features that allow you to specify actions you want the platform to execute automatically when specific security conditions are met.

Automation in SL1 includes two pieces: an automation policy and an action policy. An **automation policy** defines the conditions that can trigger an automatic action. When the criteria in an automation policy are met, one or more actions are executed. These actions are defined in an **action policy**.

For example, an automation policy might specify: if the event "illicit process" occurs on device "mailserver01", and the event is not cleared within five minutes, execute the action policy "email NOC". The action policy "email NOC" could notify all NOC staff about the "illicit process" event.

Automation policies can evaluate one or more of the following criteria:

- whether SL1 has triggered one or more of the specified events (sometimes called "alerts" or "alarms" in other products)
- whether one or more events has occurred on one or more of the specified devices
- whether one or more of the events has the specified severity (critical, major, minor, notice, or healthy)
- whether one or more of the events has the specified status (event is not cleared, event is not acknowledged, ticket is not created for event)
- whether the specified text appeared in the event message

When the criteria in an automation policy are met, an action policy can perform one or more of the following actions:

- Send an Email message to a pre-defined list of users.
- Send an SNMP trap to an external device.
- Create a new ticket.
- Perform an SNMP SET on an external device.
- Execute a custom program written in Python.
- Query a database.

For details on Run Book Automation to automatically respond to security events, see the manual **Run Book Automation**.

Reports

In addition to security policies and their related alerts, the platform generates multiple reports that trend and categorize hardware performance and policy performance, for all devices in SL1 and per device. You can use these reports to quickly find unusual spikes or lulls in usage.

For details on the default reports and graphs included with SL1, see the manual **Reports**.

Proxied Web Services

SL1 allows you to perform requests to an external web service and display the results in a graphical dashboard. The proxy feature allows you to pass authentication information to the external web service. Every time a user views the dashboard, SL1 passes the authentication information to the web service. However, the authentication information is not viewable or available to each user who views the graphical dashboard.

Chapter



4

Security Settings

Overview

This chapter provides a map to some of the security features described in this manual. If a security feature in this manual includes a specific page or field in SL1, this chapter describes that page or field.

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 the menu options, click the Advanced menu icon (.

This chapter includes the following topics:

<i>Access Control</i>	25
<i>Authentication</i>	30
<i>Multiple Tenancy and Segregation of Duties</i>	31
<i>Protection of Shared Content</i>	33
<i>Data Integrity</i>	33
<i>Security Events</i>	34
<i>Monitoring Changes to Device Configuration</i>	35
<i>Monitoring for Illicit Behavior</i>	35
<i>Blueprinting DNS, System Processes, and Windows Services</i>	36
<i>Monitoring Open Ports</i>	37
<i>Monitoring Bandwidth Usage</i>	37
<i>Monitoring Hardware Performance</i>	39

<i>Monitoring Patches and Hot Fixes</i>	40
<i>Using Run Book Automation to Automate Responses to Security Events</i>	41
<i>Reports</i>	41
<i>Proxied Web Services</i>	41
<i>Audit Logs</i>	42

Access Control

Description	Click Path	Page/Field	Field Options
Password Expiration. Defines global expiration setting for user passwords. Specifies whether or not the passwords for user account will expire and if so, when the passwords will expire.	System > Settings > Behavior	Behavior Settings / Password Expiration	Disabled, 30 days, 60 days, 90 days, 180 days
Password Hash Method. Specifies how user passwords will be encrypted for storage in the Database Server. You can choose the hashing algorithm that works best for your enterprise.	System > Settings > Behavior	Behavior Settings / Password Hash Method	MD5 (legacy), SHA-512 (FIPS 140-2 Compliant), Automatic (PHP Password API)
Password Minimum Length. Specifies the minimum required number of characters for user passwords. The default value is "8". You can enter any value between 1 and 99. EM7 will enforce this minimum when an administrator creates a new user or edits a user's account properties and when a user changes his/her password.	System > Settings > Behavior	Behavior Settings / Password Minimum Length	1 through 99 characters.
Account Lockout Type. Defines global parameter for lockouts. If a user enters incorrect login information multiple times in a row, that user will be locked out of the system. In this field, you can select how the lockout will be applied.	System > Settings > Behavior	Behavior Settings / Account Lockout Type	Lockout by IP address (default), Lockout by Username and IP address, Lockout by username, Disable

Description	Click Path	Page/Field	Field Options
<p>Account Lockout Attempts. Defines global trigger for lockouts. Specifies the number of times a user can enter incorrect login information before the lockout occurs.</p>	System > Settings > Behavior	Behavior Settings / Account Lockout Attempts	1 through 10 attempts
<p>Account Lockout Duration. Defines global duration for lockouts. Specifies how long a user will be locked out of the system.</p>	System > Settings > Behavior	Behavior Settings / Account Lockout Duration	1 hour to 24 hours in one-hour increments
<p>Login Delay. To prevent unauthorized users from using brute-force login attempts, you can set a login delay in this field. After each failed login, EM7 will not allow another attempt for the number of seconds specified in this field.</p>	System > Settings > Behavior	Behavior Settings / Login Delay	Disabled, 1 second, 2 seconds, 4 seconds, 8 seconds
<p>Single Instance Login (for both Admins and Users). Global settings for how the system will handle multiple instances of the same username. Specifies whether more than one instance of a single username can be logged in to the system at the same time. Separate settings for the default behavior for users of account type "User" and users of account type "Admin".</p>	System > Settings > Behavior	Behavior Settings / Single Instance Login (Admin) and Single Instance Login (Users)	Disabled, session can be transferred instantly, session can be transferred after 1 minute of inactivity, after 5 minutes of inactivity, after 10 minutes of inactivity, after 30 minutes of inactivity, after 1 hour of inactivity, or session can be transferred after a manually specified number
<p>Account Lockout Duration. Defines global duration for lockouts. Specifies how long a user will be locked out of the system.</p>	System > Settings > Behavior	Behavior Settings / Account Lockout Duration	1 hour to 24 hours in one-hour increments

Description	Click Path	Page/Field	Field Options
<p>Lockout Contact Information. This contact information will be displayed when a user is locked out of the system. This information should allow the user to contact his/her administrator to unlock the account.</p>	System > Settings > Behavior	Behavior Settings / Lockout Contact Information	This information should allow the user to contact his/her administrator to unlock the account.
<p>Prevent Browser Saved Credentials. This checkbox specifies whether or not SL1 will allow the browser to cache login credentials and perform auto-complete in the login page. By default, SL1 will allow browsers to cache login credentials.</p>	System > Settings > Behavior	Behavior Settings / Prevent Browser Saved Credentials	Selected, Not Selected
<p>Display Previous Login in Footer. This checkbox specifies whether or not SL1 will display information about the last successful login to the Administration Portal or All-In-One Appliance and the last failed login (if applicable) in the lower right corner of each page.</p> <p>Previous Login: yyyy-mm-dd hh-mm-ss from user's IP address</p> <p>Failed Login: yyyy-mm-dd hh-mm-ss from user's IP address</p>	System > Settings > Behavior	Behavior Settings / Display Previous Login in Footer	Selected, Not Selected
<p>Prevent Loading Interface in External Frames. If you select this checkbox, other pages cannot be loaded in external frames in the same browser session that includes SL1. This option can be used as a security measure to prevent click-jacking attacks.</p>	System > Settings > Behavior	System > Settings > Behavior / Prevent Loading Interface in External Frames	Selected, Not Selected

Description	Click Path	Page/Field	Field Options
<p>CAC/ClientCert Auth. This page allows you to define an SSL certificate check that controls whether the login page is displayed to the end user. This feature is primarily used to authenticate Common Access Card (CAC) users against a Department of Defense (DoD)-issued server-side certificate; however, based on your business needs, this feature can also be used with your own client/server certificates.</p>	<p>System > Settings > CAC/ClientCert Auth</p>	<p>Client Certificate & CAC Authentication</p>	
<p>Login Alert Message. This page defines a customizable click-through banner at login. This banner prevents further activity on the SL1 until the user agrees to the terms by clicking on the [OK] button.</p>	<p>System > Settings > Login Alert Message</p>	<p>Login Alert Editor</p>	
<p>Password Reset Email Editor. This page allows system administrators to define the Email message that is sent to users who select the "I forgot my password" option from the Login page. If the user enters a valid username in the Login page and then selects the "I forgot my password" option, the system will check the account information for that user. If the user's account information includes an Email address, the system will send the user an Email message. The Email message will include a link that allows the user to redefine their password.</p>	<p>System > Settings > Password Reset Email</p>	<p>Password Reset Email Editor</p>	
<p>Change Password. The user's new password.</p>	<p>Registry > Accounts > User Accounts > Create/Edit</p>	<p>Account Permissions / Change Password</p>	<p>Enter the new password.</p>

Description	Click Path	Page/Field	Field Options
Login State. Default login state for the user.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Login State	Active, Suspended, Vacation
Password Strength. When defining or editing a user account, the administrator can define the required password strength. The user must then always use a password that meets or exceeds that specified password strength. The system will not allow the user to save changes to his/her password that do not meet the password-strength requirement.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Password Strength	Good, Strong, Very Strong
Password Expiration. Specifies whether or not the password for this account will expire and if so, when the password will expire.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Password Expiration	Disabled, 30, 60, 90, 180 days
Password Shadowing. Specifies requirements for password reuse. By default, when a user defines a new password, he/she cannot reuse any passwords that he/she has used in the last 12 months.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Password Shadowing	Default - cannot reuse passwords from the past year, 1 - cannot reuse current password, 2 - cannot reuse last two 2 passwords, 3 - cannot reuse last 3 passwords, 4 - cannot reuse last 4 passwords, 5 - cannot reuse last 5 passwords
Require Password Reset. If selected, the user will be prompted to change his/her password at the next login. When creating a new user account, this option is selected by default. After the user's first login, when he/she is prompted to change his/her password, this option is then deselected.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Require Password Reset	Selected, Not Selected

Description	Click Path	Page/Field	Field Options
Restrict to IP. The user will be allowed to access the system only from the specified IP address. Specify the IP address in standard dotted-decimal notation.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Restrict to IP	Blank or enter an IP address

Authentication

Description	Click Path	Page/Field	Field Options
Authentication Method. Specifies how the user's username and password will be authenticated.	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Restrict to IP	For configuration details, see the manual on Active Directory and LDAP Integration .
Force Secure HTTPS. If enabled, requires the Administration Portal appliance, the All-In-One Appliance, and/or the combination Database and Administration Portal appliance to use HTTPS (secure HTTP) instead of HTTP.	System > Settings > Behavior	Behavior Settings / Force Secure HTTPS	Enabled or Disabled
Client Certificate & CAC Authentication. The Client Certificate & CAC Authentication page allows you to define an SSL certificate check that controls whether the login page is displayed to the end user. This feature is primarily used to authenticate Common Access Card (CAC) users against a Department of Defense (DoD)-issued server-side certificate; however, based on your business needs, this feature can also be used with your own client/server certificates.	System > Settings > Authentication > CAC/ClientCert Auth	Client Certificate & CAC Authentication	

Multiple Tenancy and Segregation of Duties

Description	Click Path	Page/Field	Field Options
<p>Primary Organization. Specifies the primary organization to align with a user account.</p>	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Primary Organization	You can select a single organization from a list of all organizations in the system.
<p>Additional Organization Memberships. Specifies additional organizations for the user. This allows the user to view and access elements from multiple organizations.</p>	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Additional Organization Memberships	You can select one or more organizations from a list of all organizations in the system.
<p>Account Type. For each user account, specifies the standard account type. These account types affect the list of Privilege Keys for the user.</p>	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Account Type	User, Administrator
<p>Access Key. The Privilege Keys pane displays a list of access keys that can be assigned to the user's account. Access Keys define the tabs and pages users have access to and the actions that a user may perform. These access keys are defined by the system administrator from the Access Keys page (System > Manage > Access Keys).</p>	Registry > Accounts > User Accounts > Create/Edit	Account Permissions / Privilege Keys	Accounts of type "Administrator" are automatically assigned all access keys. For accounts of type "User", you can manually assign one, multiple, or all access keys. For details on defining and using access keys, see the manual Access Permissions .

Description	Click Path	Page/Field	Field Options
<p>Credential Management. To support multi-tenancy, the system allows you to align each credential with one, multiple, or all organizations in the system. You can also align a credential with no organizations. When you align an organization with a credential, you control who can view details about the credential, who can view the name of the credential, and who can apply the credential in the system.</p>	<p>System > Manage > Credentials > Organization icon </p>	<p>Align Organizations</p>	<p>Enabled alignment, select one or more organizations to align with a credential.</p>
<p>Cloaked Comments. Ticket notes that are created with the Cloaked checkbox enabled can be viewed only by: the user who created the note; all users of type "administrator"; and users in the same organization as the user who create the note, who also have Access Hooks that allow them to view the ticket where the cloaked note resides, and who also have the Access Hook "Ticket:Notes:Cloaked".</p>	<p>Tickets > Create/Edit > Notepad > Cloak</p>	<p>Notepad / Cloak</p>	<p>Enabled, Disabled</p>
<p>User Policies. User Policies allow you to define a custom set of account properties and key privileges (from the Account Permissions page) and then save them as a policy, for reuse. When you create a user account, you can use the User Policy to quickly apply settings to the new account.</p>	<p>Registry > Accounts > User Policies > Create/Edit</p>	<p>User Policies</p>	<p>For details on defining and applying user policies, see the manual Organizations and Users.</p>

Protection of Shared Content

Description	Click Path	Page/Field	Field Options
<p>Intellectual Property Protection. After adding content to a PowerPack, you can add Intellectual Property Protection to included Dynamic Applications and event policies. Intellectual Property Protection prevents users from viewing or editing advanced implementation details about the Dynamic Application or the event policy after it has been installed on another system.</p>	<p>System > Manage > PowerPacks > Create/Edit > Dynamic Applications or Event Policies</p>	<p>Embedded/Available page / IPP</p>	<p>Enabled, Disabled</p>

Data Integrity

Description	Click Path	Page/Field	Field Options
<p>Backups. The Backup Management page allows you to define two types of backups for your system: configuration backup that stores a local copy of the core database tables that are required to restore a system, and full backups that make a full backup of Database Server.</p>	<p>System > Settings > Backup</p>		

Description	Click Path	Page/Field	Field Options
<p>Collector Groups. For distributed systems, a collector group is a group of Data Collectors. ScienceLogic Data Collectors retrieve data from managed devices and applications. Grouping multiple Data Collectors allows you to create one of the following configurations: load-balanced collection system, where you can manage more devices without loss of performance; or 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.</p>	System > Settings > Collector Groups	Collector Group Management / Collector Failover	Off (Maximize Management Devices), On (Maximize Reliability)

Security Events

Description	Click Path	Page/Field	Field Options
<p>Events console. The Event Console page displays a list of currently active events. One of the easiest ways to monitor the health of your network is to look at events. Events are messages that are triggered when a specific condition is met. For example, an event can signal that a server has gone down, that a device's hard-drives are getting too full, or simply display the status of a device.</p>	[Events] tab	Event Console	

Description	Click Path	Page/Field	Field Options
Event policies. The Event Policy Editor page allows you to define a new event or edit the properties of an existing event definition.	Registry > Events > Event Manager > create or edit	Event Policy Editor	

Monitoring Changes to Device Configuration

Description	Click Path	Page/Field	Field Options
Alert when asset record changes. The Asset Automation page allows you to define the default behavior for all asset records. For each standard asset field, you can specify how the field should be populated and whether or not the system should generate an event if the field's value changes.	System > Settings > Assets	Asset Automation / Alert on change	Yes, No

Monitoring for Illicit Behavior

Description	Click Path	Page/Field	Field Options
Generate an event if the system discovers an illicit domain record	Registry > Monitors > Domain Name > Create / Edit	Domain Name Policy / Alert if Found	Yes. Use this setting to look for an illicit domain record. If the system finds the specified, illicit domain record, the system will generate an event.
Generate an event if the system discovers an illicit process	Registry > Monitors > System Processes > Create / Edit	System Process Policy / Alert if Found	Yes. Use this setting to look for an illicit system process. If the system finds the specified, illicit system process, the system will generate an event.

Description	Click Path	Page/Field	Field Options
Generate an event if the system discovers an illicit Windows service	Registry > Monitors > Windows Services > Create / Edit	Windows Service Policy / Alert if Found	Yes. Use this setting to look for an illicit Windows service. If the system finds the specified, illicit Windows service, the system will generate an event.
Generate an event if the system discovers a specified, illicit port open	System > Customize > TCP-IP Ports	TCP/IP Port Editor / Illicit Port Alarm	On, off

Blueprinting DNS, System Processes, and Windows Services

Description	Click Path	Page/Field	Field Options
Generate an event if the system discovers a change to a domain record .	Registry > Monitors > Domain Name > Create / Edit	Domain Name Policy / Alert if Found	No. Use this setting to ensure that a required domain record is running. If the system does not find the specified domain record, the system generates an event.
Generate an event if the system discovers a change to a process .	Registry > Monitors > System Processes > Create / Edit	System Process Policy / Alert if Found	No. Use this setting to ensure that a required system process is running. If the system does not find the system process, the system generates an event.
Generate an event if the system discovers a change to a Windows service .	Registry > Monitors > Windows Services > Create / Edit	Windows Service Policy / Alert if Found	No. Use this setting to ensure that a required Windows service is running. If the system does not find the specified Windows service, the system generates an event.

Monitoring Open Ports

Description	Click Path	Page/Field	Field Options
Generate an event if the system discovers a specified, illicit port open .	System > Customize > TCP-IP Ports	TCP/IP Port Editor / Illicit Port Alarm	On, off

Monitoring Bandwidth Usage

Description	Click Path	Page/Field	Field Options
Define a global threshold and generate an event if the counter rolls over .	System > Settings > Thresholds	Global Threshold Settings / Rollover Percent	0% - 100%
Define a global threshold and generate an event if packets are sent out-of-order .	System > Settings > Thresholds	Global Threshold Settings / Out-of-order Percent	0% - 100%
Define a global threshold and generate an event if inbound bandwidth exceeds the specified percentage .	System > Settings > Thresholds	Global Threshold Settings / Inbound Percent	0% - 100%
Define a global threshold and generate an event if outbound bandwidth exceeds the specified percentage .	System > Settings > Thresholds	Global Threshold Settings / Outbound Percent	0% - 100%
Define a global threshold and generate an event if inbound bandwidth exceeds the specified Mbps .	System > Settings > Thresholds	Global Threshold Settings / Inbound Bandwidth	0 - 1,000,000 Mbps
Define a global threshold and generate an event if outbound bandwidth exceeds the specified Mbps .	System > Settings > Thresholds	Global Threshold Settings / Outbound Bandwidth	0 - 1,000,000 Mbps

Description	Click Path	Page/Field	Field Options
Define a global threshold and generate an event if <i>inbound errors exceed the specified number of packets.</i>	System > Settings > Thresholds	Global Threshold Settings / Inbound Errors	0 - 10,000 packets
Define a global threshold and generate an event if <i>outbound errors exceed the specified number of packets.</i>	System > Settings > Thresholds	Global Threshold Settings / Outbound Errors	0 - 10,000 packets
Define a global threshold and generate an event if <i>inbound discards exceed the specified number of packets.</i>	System > Settings > Thresholds	Global Threshold Settings / Inbound discards	0 - 10,000 packets
Define a global threshold and generate an event if <i>outbound discards exceed the specified number of packets.</i>	System > Settings > Thresholds	Global Threshold Settings / Outbound Discards	0 - 10,000 packets
Define a global threshold and generate an event if <i>inbound errors exceed the specified percentage.</i>	System > Settings > Thresholds	Global Threshold Settings / Inbound Error Percent	0% - 100%
Define a global threshold and generate an event if <i>outbound errors exceed the specified percentage.</i>	System > Settings > Thresholds	Global Threshold Settings / Outbound Error Percent	0% - 100%
Define a global threshold and generate an event if <i>inbound discards exceed the specified percentage.</i>	System > Settings > Thresholds	Global Threshold Settings / Inbound Discard Percent	0% - 100%
Define a global threshold and generate an event if <i>outbound discards exceed the specified percentage.</i>	System > Settings > Thresholds	Global Threshold Settings / Outbound Discard Percent	0% - 100%
Define a device-specific threshold and generate an event if the <i>counter rolls over.</i>	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / Rollover Percent	0% - 100%

Description	Click Path	Page/Field	Field Options
Define a device-specific threshold and generate an event if packets are sent out-of-order .	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / Rollover Percent	0%-100%

Monitoring Hardware Performance

Description	Click Path	Page/Field	Field Options
Define a global threshold and generate an event if system latency exceeds the specified number of milliseconds.	System > Settings > Thresholds	Global Threshold Settings / System Latency	0 ms - 5,000 ms
Define a global threshold and generate an event if system availability falls below the specified percentage.	System > Settings > Thresholds	Global Threshold Settings / System Availability	0% - 100%
Define a global threshold and generate a warning event if filesystem usage exceeds the specified percentage .	System > Settings > Thresholds	Global Threshold Settings / Filesystem Warning	0% - 100%
Define a global threshold and generate a critical event if filesystem usage exceeds the specified percentage .	System > Settings > Thresholds	Global Threshold Settings / Filesystem Critical	0% - 100%
Define a global threshold and generate an event if ICMP availability falls below the specified percentage .	System > Settings > Thresholds	Global Threshold Settings / Avail Required Ping Percent	0% - 100%
Define a device-specific threshold and generate an event if system latency exceeds the specified number of milliseconds.	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / System Latency	0 ms - 5,000 ms

Description	Click Path	Page/Field	Field Options
Define a device-specific threshold and generate an event if system availability falls below the specified percentage.	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / System Availability	0% - 100%
Define a device-specific threshold and generate a warning event if the filesystem usage exceeds the specified percentage .	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / File System Thresholds (Warning)	0% - 100%
Define a device-specific threshold and generate a critical event if the filesystem usage exceeds the specified percentage-order .	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / Files System Thresholds (Critical)	0% - 100%
Define a device-specific threshold and generate an event if ICMP availability falls below the specified percentage .	Registry > Devices > Device Manager > wrench icon > Thresholds	Device Thresholds / Avail Required Ping Percent	0% - 100%

Monitoring Patches and Hot Fixes

Description	Click Path	Page/Field	Field Options
Generate an exclusion report for a selected software title. This report displays devices where the software is installed and devices where the software is not installed .	Registry > Devices > Software > printer icon (🖨️)	Software Titles	printer icon (🖨️)

Using Run Book Automation to Automate Responses to Security Events

Description	Click Path	Page/Field	Field Options
Create an automation policy that defines the conditions during which you want the system to execute automated actions.	Registry > Run Book > Automation > Create button	Automation Policy Manager	Create button
Create an action that you want the system to execute automatically when specific conditions occur.	Registry > Run Book > Actions > Create button	Action Policy Manager	Create button

Reports

Too numerous to list in a table. For details on all the reports in SL1, see the manual on **Reports**.

Proxied Web Services

Description	Click Path	Page/Field	Field Options
Define how the system will make requests to an external web service.	Registry > Web Proxies > Proxied Web Services > Create button	Proxied Web Services	Create button
Create a graphical report that displays the results from a request to an external web service.	Dashboards > Actions > Add Widget	New Widget Configuration / Widget Type	Proxied Web Services

Audit Logs

Description	Click Path	Page/Field	Field Options
View all messages about the system's standard operations , like starting and stopping key processes, backing up data, purging old data, and other maintenance activities.	System > Monitor > System Logs	System Logs	
View a complete audit trail for all actions in the system that are performed by users or related to managed devices .	System > Monitor > Audit Logs	Audit Logs..	
Monitor and manage user logins and logouts to the system.	System > Monitor > Access Logs	Access Sessions.	
View all the messages and log entries generated for a device .	Registry > Devices > Device Manager > wrench icon > Logs	Device Logs & Messages.	

© 2003 - 2021, 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



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

International: +1-703-354-1010