IBM Security QRadar
Version 7.2.6

Troubleshooting System Notifications

Note
Before using this information and the product that it supports, read the information in "Notices" on page 41.

Product information

This document applies to IBM QRadar Security Intelligence Platform V7.2.6 and subsequent releases unless superseded by an updated version of this document.

Contents

Introduction to system notifications

IBM® Security QRadar® *Troubleshooting System Notifications* provides information on how to troubleshoot and resolve system notifications that display on the QRadar console. System notifications that display on the console can apply to any appliance or QRadar product in your deployment.

Unless otherwise noted, all references to QRadar can refer to the following products:

- IBM Security QRadar SIEM
- IBM Security QRadar Log Manager

Intended audience

Network administrators who are responsible for installing and configuring QRadar systems must be familiar with network security concepts and the Linux operating system.

Technical documentation

To find IBM Security QRadar product documentation on the web, including all translated documentation, access the IBM Knowledge Center (http://www.ibm.com/support/knowledgecenter/SS42VS/welcome).

For information about how to access more technical documentation in the QRadar products library, see Accessing IBM Security Documentation Technical Note (www.ibm.com/support/docview.wss?rs=0&uid=swg21614644).

Contacting customer support

For information about contacting customer support, see the Support and Download Technical Note (http://www.ibm.com/support/docview.wss?uid=swg21616144).

Statement of good security practices

IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Please Note:

Use of this Program may implicate various laws or regulations, including those related to privacy, data protection, employment, and electronic communications and storage. IBM Security QRadar may be used only for lawful purposes and in a lawful manner. Customer agrees to use this Program pursuant to, and assumes all responsibility for complying with, applicable laws, regulations and policies. Licensee represents that it will obtain or has obtained any consents, permissions, or licenses required to enable its lawful use of IBM Security QRadar.

Chapter 1. Troubleshooting QRadar system notifications

Use the system notifications that are generated by IBM Security QRadar to monitor the status and health of your system. Software and hardware tools and processes continually monitor the QRadar appliances and deliver information, warning, and error messages to users and administrators.

Related concepts:

Chapter 2, “Error notifications for QRadar appliances,” on page 3
Error notifications in IBM Security QRadar products require a response by the user or the administrator.

Chapter 3, “Warning notifications for QRadar appliances,” on page 17
IBM Security QRadar system health notifications are proactive messages of actual or impending software or hardware failures.

Chapter 4, “Information notifications for QRadar appliance,” on page 37
IBM Security QRadar provides information messages about the status or result of a process or action

Chapter 2. Error notifications for QRadar appliances

Error notifications in IBM Security QRadar products require a response by the user or the administrator.

Out of memory error

38750004 - Application ran out of memory

Explanation

When the system detects that no more memory or swap space is available, the application or service can stop working. Out of memory issues are caused by software, or user-defined queries and operations that exhaust the available memory.

User response

Review the error message that is written to the /var/log/qradar.log file. Restarting a service might stop the offending application or service and redistribute resources.

If you use Java™ Database Connectivity (JDBC) or the log file protocol to import many records from a log source, the system can use up resources. If multiple large data imports occur simultaneously, you can stagger the start time intervals.

Disk usage exceeded threshold

38750038 - Disk Sentry: Disk Usage Exceeded Max Threshold.

Explanation

At least one disk on your system is 95% full.

Processes shut down to prevent data corruption on your system.

User response

Free disk space by manually deleting files or changing your event or flow data retention policies. The system automatically restarts processes after you free enough disk space to fall below a threshold of 92% capacity.

Process monitor application failed to start multiple times

38750043 - Process Monitor: Application has failed to start up multiple times.

Explanation

The system is unable to start an application or process on your system.

User response

Review your flow sources to determine whether a device stopped sending flow data or whether users deleted a flow source.

Either remove the flow process by using the deployment editor or assign a flow source to your flow data. On the **Admin** tab, click **Flow Sources**.

Process monitor must lower disk usage

38750045 - Process Monitor: Disk usage must be lowered.

Explanation

The process monitor is unable to start processes because of a lack of system resources. The storage partition on the system is likely 95% full or greater.

User response

Free some disk space by manually deleting files or by changing your event or flow data retention policies. The system automatically restarts system processes when the used disk space falls below a threshold of 92% capacity.

Event pipeline dropped events

38750060 - Events/Flows were dropped by the event pipeline.

Explanation

If there is an issue with the event pipeline or you exceed your license limits, an event or flow might be dropped.

Dropped events and flows cannot be recovered.

User response

Review the following options:

- Verify the incoming event and flow rates on your system. If the event pipeline is dropping events, expand your license to handle more data.
- Review the recent changes to rules or custom properties. Rule or custom property changes can cause changes to your event or flow rates and might affect system performance.
- Determine whether the issue is related to SAR notifications. SAR notifications might indicate queued events and flows are in the event pipeline. The system usually routes events to storage, instead of dropping the events.
- Tune the system to reduce the volume of events and flows that enter the event pipeline.

Event pipeline dropped connections

38750061 - Connections were dropped by the event pipeline.

Explanation

A TCP-based protocol dropped an established connection to the system.

The number of connections that can be established by TCP-based protocols is limited to ensure that connections are established and events are forwarded. The event collection system (ECS) allows a maximum of 15,000 file handles and each TCP connection uses three file handles.

TCP protocols that provide drop connection notifications include the following protocols:

- TCP syslog protocol
- TLS syslog protocol
- TCP multiline protocol

User response

Review the following options:

- Distribute events to more appliances. Connections to other event and flow processors distribute the work load from the console.
- Configure low priority TCP log source events to use the UDP network protocol.
- Tune the system to reduce the volume of events and flows that enter the event pipeline.

Automatic update error

38750066 - `Automatic updates could not complete installation. See the Auto Update Log for details.`

Explanation

The update process encountered an error or cannot connect to an update server. The system is not updated.

User response

Select one of the following options:

- Verify the automatic update history to determine the cause of the installation error.

  In the **Admin** tab, click the **Auto Update** icon and select **View Log**.
- Verify that your console can connect to the update server.

  In the Updates window, select **Change Settings**, then click the **Advanced** tab to view your automatic update configuration. Verify the address in the **Web Server** field to ensure that the automatic update server is accessible.

Auto update installed with errors

38750067 - `Automatic updates installed with errors. See the Auto Update Log for details.`

Explanation

The most common reason for automatic update errors is a missing software dependency for a DSM, protocol, or scanner update.

User response

Select one of the following options:

- In the **Admin** tab, click the **Auto Update** icon and select **View Update History** to determine the cause of the installation error. You can view, select, and then reinstall a failed RPM.
- If an auto update is unable to reinstall through the user interface, manually download and install the missing dependency on your console. The console replicates the installed file to all managed hosts.

Standby HA system failure

38750080 - Standby HA System Failure.

Explanation

The status of the secondary appliance switches to `failed` and the system has no HA protection.

User response

Review the following resolutions:

- Restore the secondary system.

  Click the **Admin** tab, click **System and License Management**, and then click **Restore System**.
- Inspect the secondary HA appliance to determine whether it is powered down or experienced a hardware failure.
- Use the **`ping`** command to check the communication between the primary and standby system.
- Check the switch that connects the primary and secondary HA appliances.

  Verify the IPtables on the primary and secondary appliances.
- Review the `/var/log/qradar.log` file on the standby appliance to determine the cause of the failure.

Active high availability (HA) system failure

38750081 - Active HA System Failure.

Explanation

The active system cannot communicate with the standby system because the active system is unresponsive or failed. The standby system takes over operations from the failed active system.

User response

Review the following resolutions:

- Inspect the active HA appliance to determine whether it is powered down or experienced a hardware failure.
- If the active system is the HA primary, restore the active system.

  Click the **Admin** tab and click **System and License Management**. From the **High Availability** menu, select the **Restore System** option.

- Review the `/var/log/qradar.log` file on the standby appliance to determine the cause of the failure.
- Use the **`ping`** command to check the communication between the active and standby system.
- Check the switch that connects the active and standby HA appliances.

  Verify the IPtables on the active and standby appliances.

Failed to install high availability

38750086 - `There was a problem installing High Availability on the cluster.`

Explanation

When you install a high-availability (HA) appliance, the installation process links the primary and secondary appliances. The configuration and installation process contains a time interval to determine when an installation requires attention. The high-availability installation exceeded the six-hour time limit.

No HA protection is available until the issue is resolved.

User response

Contact customer support.

Failed to uninstall an HA appliance

38750087 - `There was a problem while removing High Availability on the cluster.`

Explanation

When you remove a high-availability (HA) appliance, the installation process removes connections and data replication processes between the primary and secondary appliances. If the installation process cannot remove the HA appliance from the cluster properly, the primary system continues to work normally.

User response

Try to remove the high-availability appliance a second time.

Scanner initialization error

38750089 - `A scanner failed to initialize.`

Explanation

A scheduled vulnerability scan is unable to connect to an external scanner to begin the scan import process.

Scan initialization issues are typically caused by credential problems or connectivity issues to the remote scanner. Scanners that fail to initialize display detailed error messages in the hover text of a scheduled scan with a status of failed.

User response

Follow these steps:

1. Click the **Admin** tab.
2. On the navigation menu, click **Data Sources**.
3. Click **Schedule VA Scanners** icon.
4. From the scanner list, hover the cursor in the **Status** column of any scanner to display a detailed success or failure message.

Scan failure error

38750090 - A scanner has failed.

Explanation

A scheduled vulnerability scan failed to import vulnerability data. Scan failures are typically caused by configuration or performance issues that result from a large volume of data to import. Scan failures can also occur when a scan report downloaded by the system is in an unreadable format.

User response

Follow these steps:

1. Click the **Admin** tab.
2. On the navigation menu, click **Data Sources**.
3. Click **Schedule VA Scanners**.
4. From the scanner list, hover the cursor in the **Status** column of any scanner to display a detailed success or failure message.

Filter initialization failed

38750091 - Traffic analysis filter failed to initialize.

Explanation

If a configuration is not saved correctly, or if a configuration file is corrupted, the event collection service (ECS) might fail to initialize. If the traffic analysis process is not started, new log sources are not automatically discovered.

User response

Select one of the following options:

- Manually create log sources for any new appliances or event sources until traffic analysis process is working.

  All new event sources are classified as SIM Generic until they are mapped to a log source.
- If you get an automatic update error, review the automatic update log to determine whether an error occurred when a DSM or a protocol was installed.

Disk storage unavailable

38750092 - Disk Sentry has detected that one or more storage partitions are not accessible.

Explanation

The disk sentry did not receive a response within 30 seconds. A storage partition issue might exist, or the system might be under heavy load and not able to respond within the 30-second threshold.

User response

Select one of the following options:

- Verify the status of your `/store` partition by using the **`touch`** command.

If the system responds to the **`touch`** command, the unavailability of the disk storage is likely due to system load.

- Determine whether the notification corresponds to dropped events.

If events were dropped events and the disk storage is unavailable, event and flow queues might be full. Investigate the status of storage partitions.

Insufficient disk space to export data

38750096 - `Insufficient disk space to complete data export request.`

Explanation

If the export directory does not contain enough space, the export of event, flow, and offense data is canceled.

User response

Select one of the following options:

- Free some disk space in the `/store/exports` directory.
- Configure the **Export Directory** property in the System Settings window to use to a partition that has sufficient disk space.
- Configure an offboard storage device.

Accumulator is falling behind

38750099 - `The accumulator was unable to aggregate all events/flows for this interval.`

Explanation

This message appears when the system is unable to accumulate data aggregations within a 60-second interval.

Every minute, QRadar creates data aggregations for each aggregated search. The data aggregations are used in time-series graphs and reports and must be completed within a 60-second interval. If the count of searches and unique values in the searches are too large, the time that is required to process the aggregations might exceed 60 seconds. When the accumulation is unable to complete within 60 seconds, the accumulation interval is dropped. Time-series graphs and reports might be missing columns for the time period when the problem occurred.

You do not lose data when this problem occurs. All raw data, events, and flows are still written to disk. Only the accumulations, which are data sets that are generated from stored data, are incomplete.

User response

The following factors might contribute to the increased workload that is causing the accumulator to fall behind:

Frequency of the incomplete accumulations

If the accumulation fails only once or twice a day, the drops might be caused by increased system load due to large searches, data compression cycles, or data backup.

Infrequent failures can be ignored. If the failures occur multiple times per day, during all hours, you might want to investigate further.

High system load

If other processes use many system resources, the increased system load can cause the aggregations to be slow. Review the cause of the increased system load and address the cause, if possible.

For example, if the failed accumulations occur during a large data search that takes a long time to complete, you might be able to prevent the accumulator drops by reducing the size of the saved search.

Large accumulator demands

If the accumulator intervals are dropped regularly, you might need to reduce the workload.

The workload of the accumulator is driven by the number of aggregations and the number of unique objects in those aggregations. The number of unique objects in an aggregation depends on the group-by parameters and the filters that are applied to the search.

For example, a search that aggregates for services, filters the data by using a local network hierarchy item, such as DMZ area, and then groups by IP address might result in a search that contains up to 200 unique objects. If you add destination ports to the search, and each server is hosting 5-10 services on different ports, the new aggregate of `destination.ip + destination.port` can increase the number of unique objects to 2000. If you add the source IP address to the aggregate, and you have many thousands of remote IP addresses that are hitting each service, the aggregated view might have hundreds of thousands of unique values. This search would create a heavy demand on the accumulator.

To review the aggregated views that put the highest demand on the accumulator:

1. On the **Admin** tab, click **Aggregated Data Management**.
2. Click the **Data Written** column to sort in descending order and show the largest views.
3. Review the business case for each of the largest aggregations to see whether they are still required.

CRE failed to read rules

`38750107 - The last attempt to read in rules (usually due to a rule change) has failed. Please see the message details and error log for information on how to resolve this.`

Explanation

The custom rules engine (CRE) on an Event Processor is unable to read a rule to correlate an incoming event. The notification might contain one of the following messages:

- If the CRE was unable to read a single rule, in most cases, a recent rule change is the cause. The payload of the notification message displays the rule or rule of the rule chain that is responsible.
- In rare circumstances, data corruption can cause a complete failure of the rule set. An application error is displayed and the rule editor interface might become unresponsive or generate more errors.

User response

For a single rule read error, review the following options:

- To locate the rule that is causing the notification, temporarily disable the rule.
- Edit the rule to revert any recent changes.
- Delete and re-create the rule that is causing the error.

For application errors where the CRE failed to read rules, contact customer support.

Accumulator cannot read the view definition for aggregate data

38750108 - Accumulator: Cannot read the aggregated data view definition in order to prevent an out of sync problem. Aggregated data views can no longer be created or loaded. Time series graphs will no longer work as well as reporting.

Explanation

A synchronization issue occurred. The aggregate data view configuration that is in memory wrote erroneous data to the database.

To prevent data corruption, the system disables aggregate data views. When aggregate data views are disabled, time series graphs, saved searches, and scheduled reports display empty graphs.

User response

Contact customer support.

Store and forward schedule did not forward all events

38750109 - A store and forward schedule finished while events were left on disk. These events will be stored on the local event collector until the next forwarding sessions begins.

Explanation

If the schedule contains a short start and end time or many events to forward, the Event Collector appliance might not have sufficient time to transfer the queued events. Events are stored until the next opportunity to forward events. When the next store and forward interval occurs, the events are forwarded to the Event Processor.

User response

Increase the event forwarding rate from your Event Collector appliance or increase the time interval that is configured for forwarding events.

Disk failure

38750110 - Disk Failure: Hardware Monitoring has determined that a disk is in failed state.

Explanation

On-board system tools detected that a disk failed. The notification message provides information about the failed disk and the slot or bay location of the failure.

User response

If the notification persists, contact customer support or replace parts.

Predictive disk failure

38750111 - Predictive Disk Failure: Hardware Monitoring has determined that a disk is in predictive failed state.

Explanation

The system monitors the status of the hardware on an hourly basis to determine when hardware support is required on the appliance.

The on-board system tools detected that a disk is approaching failure or end of life. The slot or bay location of the failure is identified.

User response

Schedule maintenance for the disk that is in a predictive failed state.

Scan tool failure

38750118 - A scan has been stopped unexpectedly, in some cases this may cause the scan to be stopped.

Explanation

The system cannot initialize a vulnerability scan and asset scan results cannot be imported from external scanners. If the scan tools stop unexpectedly, the system cannot communicate with an external scanner. The system tries the connection to the external scanner five times in 30-second intervals.

In rare cases, the discovery tools encounter an untested host or network configuration.

User response

Select one of the following options:

- Review the configuration for external scanners in the deployment editor to ensure that the gateway IP address is correct.
- Ensure that the external scanner can communicate through the configured IP address.
- Ensure that the firewall rules for your DMZ are not blocking communication between your appliance and the assets you want to scan.

External scan gateway failure

38750119 - An invalid/unknown gateway IP address has been supplied to the external hosted scanner, the scan has been stopped.

Explanation

When an external scanner is added, a gateway IP address is required. If the address that is configured for the scanner in the deployment editor is incorrect, the scanner cannot access your external network.

User response

Select one of the following options:

- Review the configuration for any external scanners that are configured in the deployment editor to ensure that the gateway IP address is correct.
- Ensure that the external scanner can communicate through the configured IP address.
- Ensure that the firewall rules for your DMZ are not blocking communication between your appliance and the assets you want to scan.

User authentication failed for automatic updates

38750127 - Automatic updates user authentication failed. A valid individual IBM ID is required.

Explanation

Valid credentials are required to authorize automatic downloads from the update server.

User response

Select one of the following options:

- Administrators must register for an account on the IBM support website (http://www.ibm.com/support/).
- To view the automatic update settings, on the **Admin** tab, click the **Auto Update** icon and select **Change Settings** > **Advanced**. Administrators can confirm that the user name and password in the Settings window are correct.

Aggregated data limit was reached

38750130 - The aggregated data view could not be created due to an aggregated limit.

Explanation

The accumulator is a QRadar process that counts and prepares events and flows in data accumulations to assist with searches, displaying charts, and report performance. The accumulator process aggregates data in pre-defined time spans to create aggregate data views. An *aggregate data view* is a data set that is used to draw a time series graph, create scheduled reports, or trigger anomaly detection rules.

The Console is limited to 130 active aggregate data views.

The following user actions can create a new aggregate data view:

- New anomaly detection rules.
- New reports.
- New saved searches that use time series data.

When the aggregate data view limit is reached, the notification is generated. As users attempt to create new anomaly rules, reports, or saved searches, they are prompted in the user interface that the system is at the limit.

User response

To resolve this issue, administrators can review the active aggregate data views on the **Admin** tab in the **Aggregated Data Management** window. The aggregated data management feature provides information on the reports, searches, and anomaly detection rules in use by each aggregate data view. The administrator can review the list of aggregate data views to determine what data is most import to the users. Aggregate data views can be disabled to allow users to create a new rule, report, or saved search that requires an aggregate data view.

If the administrator decides to delete an aggregate data view, a summary provides an outline of the searches, rules, or reports affected. To re-create a deleted aggregate data view, the administrator needs only to re-enable or re-create the search, anomaly rule, or report. The system automatically creates the aggregate data view based on the data required.

Magistrate is unable to persist offense updates

38750147 - `Magistrate encountered serious errors that may prevent offenses from being updated.`

Explanation

The system detected an exception when writing offense updates to the database.

Events will be processed and stored, but they will not contribute to offenses.

User response

Conduct a soft clean of the SIM data model with **Deactivate offenses** unchecked.

1. Click the **Admin** tab.
2. On the toolbar, click **Advanced** > **Clean SIM Model**.
3. Click **Soft Clean** to set the offenses to inactive.
4. Ensure that **Deactivate offenses** is not checked.

5. Click the **Are you sure you want to reset the data model?** checkbox and click **Proceed**.

When you clean the SIM model, all existing offenses are closed. Cleaning the SIM model does not affect existing events and flows.

Chapter 3. Warning notifications for QRadar appliances

IBM Security QRadar system health notifications are proactive messages of actual or impending software or hardware failures.

Maximum sensor devices monitored

38750006 - Traffic analysis is already monitoring the maximum number of log sources.

Explanation

The system contains a limit to the number of log sources that can be queued for automatic discovery by traffic analysis. If the maximum number of log sources in the queue is reached, then new log sources cannot be added.

Events for the log source are categorized as SIM Generic and labeled as Unknown Event Log.

User response

Select one of the following options:

- Review SIM Generic log sources on the **Log Activity** tab to determine the appliance type from the event payload.
- Ensure that automatic updates can download the latest DSM updates to properly identify and parse log source events.
- Verify whether the log source is officially supported.

  If your appliance is supported, manually create a log source for the events that were not automatically discovered.
- If your appliance is not officially supported, create a universal DSM to identify and categorize your events.
- Wait for the device to provide 1,000 events.

  If the system cannot auto discover the log source after 1,000 events, it is removed from the traffic analysis queue. Space becomes available for another log source to be automatically discovered.

Unable to determine associated log source

38750007 - Unable to automatically detect the associated log source for IP address <IP address>.

Explanation

At minimum, 25 events are required to identify a log source. If the log source is not identified after 1,000 events, the system abandons the automatic discovery process.

When the traffic analysis process exceeds the maximum threshold for automatic discovery, the system categorizes the log source as SIM Generic and labels the events as Unknown Event Log.

User action

Review the following options:

- Review the IP address to identify the log source.
- Review any log sources that forward events at a low rate. Log sources that have low event rates commonly cause this notification.
- To properly parse events for your system, ensure that automatic update downloads the latest DSMs.
- Review any log sources that provide events through a central log server. Log sources that are provided from central log servers or management consoles might require that you manually create their log sources.
- Review the **Log Activity** tab to determine the appliance type from the IP address in the notification message and then manually create a log source.
- Verify whether the log source is officially supported. If your appliance is supported, manually create a log source for the events.
- If your appliance is not officially supported, create a universal DSM to identify and categorize your events.

Maximum events reached

38750008 - Events per interval threshold was exceeded in past hour.

Explanation

Each appliance has a license that processes a specific volume of event and flow data.

If the license limit continues to be exceeded, the system might queue events and flows, or possibly drop the data when the backup queue fills.

User response

Tune the system to reduce the volume of events and flows that enter the event pipeline.

Flow collector cannot establish initial time synchronization

38750009 - Flow collector could not establish initial time synchronization.

Explanation

The QFlow process contains an advanced function for configuring a server IP address for time synchronization. In most cases, do not configure a value. If configured, the QFlow process attempts to synchronize the time every hour with the IP address time server.

User response

In the deployment editor, select the QFlow process. Click **Actions** > **Configure** and click **Advanced**. In the **Time Synchronization Server IP Address** field, clear the value and click **Save**.

Backup unable to complete a request

38750033 - Backup: Not enough free disk space to perform the backup.

Explanation

This notification occurs when there is not enough free space to perform a backup.

Disk Sentry is responsible for monitoring system disk and storage issues. Before a backup begins, Disk Sentry checks the available disk space to determine whether the backup can complete successfully. If the disk space is above the threshold limit of 90% on the partition that contains your backup data, the backup is canceled. If the free disk space is less than two times the size of the last backup, the backup is canceled. By default, backups are stored in `/store/backup`.

User response

To resolve this issue, select one of the following options:

- Free up disk space on your appliance to allow enough space for a backup to complete in `/store/backup`.
- Configure your existing backups to use a partition with free disk space.
- Configure additional storage for your appliance. For more information, see the *Offboard Storage Guide*.

Backup unable to execute a request

38750035 - Backup: Unable to Execute Backup Request.

Explanation

A backup cannot start or cannot complete for one of the following reasons:

- The system is unable to clean the backup replication synchronization table.
- The system is unable to run a delete request.
- The system is unable to synchronize backup with the files that are on the disk.
- The NFS-mounted backup directory is not available or has incorrect NFS export options (`no_root_squash`).
- The system cannot initialize on-demand backup.
- The system cannot retrieve configuration for the type of backup that is selected.
- Cannot initialize a scheduled backup.

User response

Manually start a backup to determine whether the failure recurs. If multiple backups fail to start, contact customer support.

Process monitor license expired or invalid

38750044 - Process Monitor: Unable to start process: license expired or invalid.

Explanation

The license is expired for a managed host. All data collection processes stop on the appliance.

User response

Contact your sales representative to renew your license.

Found an unmanaged process that is causing long transaction

38750048 - Transaction Sentry: Found an unmanaged process causing unusually long transaction that negatively effects system stability.

Explanation

The transaction sentry determines that an outside process, such as a database replication issue, maintenance script, auto update, or command line process, or a transaction is causing a database lock.

User response

Select one of the following options:

- Review the /var/log/qradar.log file for the word TxSentry to determine the process identifier that is causing your transaction issues.
- Wait to see whether the process completes the transaction and releases the database lock.
- Manually release the database lock.

Restored system health by canceling hung transactions

38750049 - Transaction Sentry: Restored system health by canceling hung transactions or deadlocks.

Explanation

The transaction sentry restored the system to normal system health by canceling suspended database transactions or removing database locks. To determine the process that caused the error, review the qradar.log file for the word TxSentry.

User response

No action is required.

Maximum active offenses reached

38750050 - MPC: Unable to create new offense. The maximum number of active offenses has been reached.

Explanation

The system is unable to create offenses or change a dormant offense to an active offense. The default number of active offenses that can be open on your system is limited to 2500. An active offense is any offense that continues to receive updated event counts in the past five days or less.

User response

Select one of the following options:

- Change low security offenses from open (active) to closed, or to closed protected.
- Tune your system to reduce the number of events that generate offenses.

  To prevent a closed offense from being removed by your data retention policy, protect the closed offense.

Maximum total offenses reached

38750051 - MPC: Unable to process offense. The maximum number of offenses has been reached.

Explanation

By default, the process limit is 2500 active offenses and 100,000 overall offenses.

If an active offense does not receive an event update within 30 minutes, the offense status changes to dormant. If an event update occurs, a dormant offense can change to active. After five days, dormant offenses that do not have event updates change to inactive.

User response

Select one of the following options:

- Tune your system to reduce the number of events that generate offenses.
- Adjust the offense retention policy to an interval at which data retention can remove inactive offenses.

  To prevent a closed offense from being removed by your data retention policy, protect the closed offense.
- To free disk space for important active offenses, change offenses from active to dormant.

Long running reports stopped

38750054 - Terminating a report which was found executing for longer than the configured maximum threshold.

Explanation

The system cancels the report that exceeded the time limit. Reports that run longer than the following default time limits are canceled.

Table 1. Default time limits by report frequency

Report frequency	Default time limits (hours)
Hourly	2
Daily	12
Manual	12
Weekly	24
Monthly	24

User required

Select one of the following options:

- Reduce the time period for your report, but schedule the report to run more frequently.
- Edit manual reports to generate on a schedule.

  A manual report might rely on raw data but not have access to accumulated data. Edit your manual report and change the report to use an hourly, daily, monthly, or weekly schedule.

Out of memory error and erroneous application restarted

38750055 - Out of Memory: system restored, erroneous application has been restarted.

Explanation

An application or service ran out of memory and was restarted. Out of memory issues are commonly caused by software issues or user-defined queries.

User response

Review the /var/log/qradar.log file to determine whether a service restart is required.

Determine whether large vulnerability scans or the importing of large volumes of data is responsible for the error. For example, compare when the system imports events or vulnerability data on your system with the notification timestamp. If necessary, stagger the time intervals for the data imports.

Long transactions for a managed process

38750056 - Transaction Sentry: Found managed process causing unusually long transaction that negatively effects system stability.

Explanation

The transaction sentry determines that a managed process, such as Tomcat or event collection service (ECS) is the cause of a database lock.

A managed process is forced to restart.

User response

To determine the process that caused the error, review the qradar.log for the word TxSentry.

Protocol source configuration incorrect

38750057 - A protocol source configuration may be stopping events from being collected.

Explanation

The system detected an incorrect protocol configuration for a log source. Log sources that use protocols to retrieve events from remote sources can generate an initialization error when a configuration problem in the protocol is detected.

User response

To resolve protocol configuration issues:

- Review the log source to ensure that the protocol configuration is correct.

  Verify authentication fields, file paths, database names for JDBC, and ensure that the system can communicate with remote servers. Hover your mouse pointer over a log source to view more error information.
- Review the `/var/log/qradar.log` file for more information about the protocol configuration error.

MPC: Process not shutdown cleanly

38750058 - `MPC: Server was not shutdown cleanly. Offenses are being closed in order to re-synchronize and ensure system stability.`

Explanation

The magistrate process encountered an error. Active offenses are closed, services are restarted, and if required, the database tables are verified and rebuilt.

The system synchronizes to prevent data corruption. If the magistrate component detects a corrupted state, then the database tables and files are rebuilt.

User response

The magistrate component is capable of self-repair. If the error continues, contact customer support.

Last backup exceeded the allowed time limit

38750059 - `Backup: The last scheduled backup exceeded execution threshold.`

Explanation

The time limit is determined by the backup priority that you assign during configuration.

User response

Select one of the following options:

- Edit the backup configuration to extend the time limit that is configured to complete the backup. Do not extend over 24 hours.
- Edit the failed backup and change the priority level to a higher priority. Higher priority levels allocate more system resources to completing the backup.

Log source license limit

38750062 - The number of configured Log Sources is approaching or has reached the licensed limit.

Explanation

Every appliance is sold with a license that collects events from a specific number of log sources. You approached or exceeded the license limit.

Any more log sources that added are disabled by default. Events are not collected for disabled log sources.

User response

Review the following options:

- On the **Admin** tab, click the **Log Sources** icon and disable or delete any log sources that are a low priority or have an inactive event source. Disabled log sources do not count towards your log source license. However, the event data that is collected by disabled log sources is still available and searchable.
- Ensure that log sources you deleted do not automatically rediscover. If the log source rediscovers, you can disable the log source. Disabling a log source prevents automatic discovery.
- Ensure that you do not exceed your license limit when you add log sources in bulk.

Deployment of an automatic update

38750069 - Automatic updates installed successfully. In the Admin tab, click Deploy Changes.

Explanation

An automatic update, such as an RPM update, was downloaded and requires that you deploy the change to finish the installation process.

User response

In the **Admin** tab, click **Deploy Changes**.

Log source created in a disabled state

38750071 - A Log Source has been created in the disabled state due to license limits.

Explanation

Traffic analysis is a process that automatically discovers and creates log sources from events. If you are at your current log source license limit, the traffic analysis process might create the log source in the disabled state. Disabled log sources do not collect events and do not count in your log source limit.

User response

Review the following options:

- On the **Admin** tab, click the **Log Sources** icon and disable or delete low priority log sources. Disabled log sources do not count towards your log source license.
- Ensure that deleted log sources do not automatically rediscover. You can disable the log source to prevent automatic discovery.
- Ensure that you do not exceed your license limit when you add log sources in bulk.
- If you require an expanded license to include more log sources, contact your sales representative.

SAR sentinel threshold crossed

38750073 - SAR Sentinel: threshold crossed.

Explanation

The system activity reporter (SAR) utility detected that your system load is above the threshold. Your system can experience reduced performance.

User response

Review the following options:

- In most cases, no resolution is required.

  For example, when the CPU usage over 90%, the system automatically attempts to return to normal operation.
- If this notification is recurring, increase the default value of the SAR sentinel.

  Click the **Admin** tab, then click **Global System Notifications**. Increase the notification threshold.
- For system load notifications, reduce the number of processes that run simultaneously.

  Stagger the start time for reports, vulnerability scans, or data imports for your log sources. Schedule backups and system processes to start at different times to lessen the system load.

User does not exist or is undefined

38750075 - User either does not exist or has an undefined role.

Explanation

The system attempted to update a user account with more permissions, but the user account or user role does not exist.

User response

On the **Admin** tab, click **Deploy Changes**. Updates to user accounts or roles require that you deploy the change.

Disk usage warning

38750076 - Disk Sentry: Disk Usage Exceeded warning Threshold.

Explanation

The disk sentry detected that the disk usage on your system is greater than 90%.

When the disk space on your system reaches 95% full, the system begins to disable processes to prevent data corruption.

User response

You must free some disk space by deleting files or by changing your data retention policies. The system can automatically restart processes after the disk space usage falls below a threshold of 92% capacity.

Infrastructure component is corrupted or did not start

38750083 - Infrastructure component corrupted.

Explanation

When the message service (IMQ) or PostgreSQL database cannot start or rebuild, the managed host cannot operate properly or communicate with the console.

User response

Contact customer support.

Data replication difficulty

38750085 - Data replication experiencing difficulty.

Explanation

A managed host had difficulty when it downloaded replication data. Data replication ensures that managed hosts can continue to collect data if the console becomes unavailable. If a managed host repeatedly fails to replicate data downloads, the system might experience performance or communication issues.

User response

If a managed host does not resolve the replication issue on its own, contact customer support.

Events routed directly to storage

38750088 - Performance degradation has been detected in the event pipeline. Event(s) were routed directly to storage.

Explanation

To prevent queues from filling, and to prevent the system from dropping events, the event collection system (ECS) routes data to storage. Incoming events and flows are not categorized. However, raw event and flow data is collected and searchable.

User response

Review the following options:

- Verify the incoming event and flow rates. If the event pipeline is queuing events, expand your license to hold more data.
- Review recent changes to rules or custom properties. Rule or custom property changes might cause sudden changes to your event or flow rates. Changes might affect performance or cause the system to route events to storage.
- DSM parsing issues can cause the event data to route to storage. Verify whether the log source is officially supported.
- SAR notifications might indicate that queued events and flows are in the event pipeline.
- Tune the system to reduce the volume of events and flows that enter the event pipeline.

Custom property disabled

38750097 - A custom property has been disabled.

Explanation

A custom property is disabled because of problems processing the custom property. Rules, reports, or searches that use the disabled custom property stop working properly.

User response

Select one of the following options:

- Review the disabled custom property to correct your regex patterns. Do not re-enable disabled custom properties without first reviewing and optimizing the regex pattern or calculation.
- If the custom property is used for custom rules or reports, ensure that the **Optimize parsing for rules, reports, and searches** check box is selected.

Device backup failure

38750098 - Either a failure occurred while attempting to backup a device, or the backup was cancelled.

Explanation

The error is commonly caused by configuration errors in Configuration Source Management (CSM) or if a backup is canceled by a user.

User response

Select one of the following options:

- Review the credentials and address sets in CSM to ensure that the appliance can log in.
- Verify the protocol that is configured to connect to your network device is valid.
- Ensure that your network device and version is supported.
- Verify that there is connectivity between your network device and the appliance.
- Verify that the most current adapters are installed.

Accumulator is falling behind

38750099 - The accumulator was unable to aggregate all events/flows for this interval.

Explanation

This message appears when the system is unable to accumulate data aggregations within a 60-second interval.

Every minute, QRadar creates data aggregations for each aggregated search. The data aggregations are used in time-series graphs and reports and must be completed within a 60-second interval. If the count of searches and unique values in the searches are too large, the time that is required to process the aggregations might exceed 60 seconds. When the accumulation is unable to complete within 60 seconds, the accumulation interval is dropped. Time-series graphs and reports might be missing columns for the time period when the problem occurred.

You do not lose data when this problem occurs. All raw data, events, and flows are still written to disk. Only the accumulations, which are data sets that are generated from stored data, are incomplete.

User response

The following factors might contribute to the increased workload that is causing the accumulator to fall behind:

Frequency of the incomplete accumulations
If the accumulation fails only once or twice a day, the drops might be caused by increased system load due to large searches, data compression cycles, or data backup.

Infrequent failures can be ignored. If the failures occur multiple times per day, during all hours, you might want to investigate further.

High system load
If other processes use many system resources, the increased system load can cause the aggregations to be slow. Review the cause of the increased system load and address the cause, if possible.

For example, if the failed accumulations occur during a large data search that takes a long time to complete, you might be able to prevent the accumulator drops by reducing the size of the saved search.

Large accumulator demands
If the accumulator intervals are dropped regularly, you might need to reduce the workload.

The workload of the accumulator is driven by the number of aggregations and the number of unique objects in those aggregations. The number of unique objects in an aggregation depends on the group-by parameters and the filters that are applied to the search.

For example, a search that aggregates for services, filters the data by using a local network hierarchy item, such as DMZ area, and then groups by IP address might result in a search that contains up to 200 unique objects. If you add destination ports to the search, and each server is hosting 5-10 services on different ports, the new aggregate of `destination.ip + destination.port` can increase the number of unique objects to 2000. If you

add the source IP address to the aggregate, and you have many thousands of remote IP addresses that are hitting each service, the aggregated view might have hundreds of thousands of unique values. This search would create a heavy demand on the accumulator.

To review the aggregated views that put the highest demand on the accumulator:

1. On the **Admin** tab, click **Aggregated Data Management**.
2. Click the **Data Written** column to sort in descending order and show the largest views.
3. Review the business case for each of the largest aggregations to see whether they are still required.

Event or flow data not indexed

38750101 - Event/Flow data not indexed for interval.

Explanation

If too many indexes are enabled or the system is overburdened, the system might drop the event or flow from the index portion.

User response

Select one of the following options:

- If the dropped index interval occurs with SAR sentinel notifications, the issue is likely due to system load or low disk space.
- To temporarily disable some indexes to reduce the system load, on the **Admin** tab, click the **Index Management** icon.

Threshold reached for response actions

38750102 - Response Action: Threshold reached.

Explanation

The custom rules engine (CRE) cannot respond to a rule because the response threshold is full.

Generic rules or a system that is tuned can generate a many response actions, especially systems with the **IF-MAP** option enabled. Response actions are queued. Response actions might be dropped if the queue exceeds 2000 in the event collection system (ECS) or 1000 response actions in Tomcat.

User response

- If the **IF-MAP** option is enabled, verify that the connection to the **IF-MAP** server exists and that a bandwidth problem is not causing rule response to queue in Tomcat.
- Tune your system to reduce the number of rules that are triggering.

Disk replication falling behind

38750103 - DRBD Sentinel: Disk replication is falling behind. See log for details.

Explanation

If the replication queue fills on the primary appliance, system load on the primary might increases. Replication issues are commonly caused by performance issues on the primary system, or storage issues on the secondary system, or bandwidth problems between the appliances.

User response

Select one of the following options:

- Review bandwidth activity by loading a saved search **MGMT: Bandwidth Manager** from the **Log Activity** tab. This search displays bandwidth usage between the console and hosts.
- If SAR sentinel notifications are recurring on the primary appliance, Distributed Replicated Block Device queues might be full on the primary system.
- Use SSH and the `cat /proc/drbd` command to monitor the Distributed Replicated Block Device status of the primary or secondary hosts.

Asset change discarded

38750106 - `Asset Changes Aborted.`

Explanation

An asset change exceeded the change threshold and the asset profile manager ignores the asset change request.

The asset profile manager includes a process, asset persistence, that updates the profile information for assets. The process collects new asset data and then queues the information before the asset model is updated. When a user attempts to add or edit an asset, the data is stored in temporary storage and added to the end of the change queue. If the change queue is large, the asset change can time out and the temporary storage is deleted.

User response

Select one of the following options:

- Add or edit the asset a second time.
- Adjust or stagger the start time for your vulnerability scans or reduce the size of your scans.

Asset persistence queue disk full

38750113 - `Asset Persistence Queue Disk Full.`

Explanation

The system detected the spillover disk space that is assigned to the asset persistence queue is full. Asset persistence updates are blocked until disk space is available. Information is not dropped.

User response

Reduce the size of your scan. A reduction in the size of your scan can prevent the asset persistence queues from overflowing.

Asset update resolver queue disk full

38750115 - Asset Update Resolver Queue Disk Full.

Explanation

The system detected that the spillover disk space that is assigned to the asset resolver queue is full.

The system continually writes the data to disk to prevent any data loss. However, if the system has no disk space, it drops scan data. The system cannot handle incoming asset scan data until disk space is available.

User response

Review the following options:

- Ensure that your system has free disk space. The notification can accompany SAR Sentinel notifications to notify you of potential disk space issues.
- Reduce the size of your scans.
- Decrease the scan frequency.

Disk full for the asset change queue

38750117 - Asset Change Listener Queue Disk Full.

Explanation

The asset profile manager includes a process, change listener, that calculates statistics to update the CVSS score of an asset. The system writes the data to disk, which prevents data loss of pending asset statistics. However, if the disk space is full, the system drops scan data.

The system cannot process incoming asset scan data until disk space is available.

User response

Select one of the following options:

- Ensure that your system has sufficient free disk space.
- Reduce the size of your scans.
- Decrease the scan frequency.

Expensive custom rule found

38750120 - Expensive Custom Rules Found in CRE: Performance degradation has been detected in the event pipeline. Found expensive custom rules in CRE.

Explanation

The custom rules engine (CRE) is a process that validates if an event matches a rule set and then trigger alerts, offenses, or notifications.

When a user creates a custom rule that has a large scope or uses a regex pattern that is not optimized, the custom rule can affect performance.

User response

Review the following options:

- On the **Offenses** tab, click **Rules** and use the search window to find and either edit or disable the expensive rule.
- If SAR sentinel notifications are recurring with the expensive rule notification, investigate the rule.

Accumulation is disabled for the anomaly detection engine

38750121 - Accumulation disabled for the Anomaly Detection Engine.

Explanation

Aggregate data view is disabled or unavailable or a new rule requires data that is unavailable.

A dropped accumulation does not indicate lost anomaly data. The original anomaly data is maintained because accumulations are data sets generated from stored data. The notification provides more details about the dropped accumulation interval.

The anomaly detection engine cannot review that interval of the anomaly data for the accumulation.

User response

Update anomaly rules to use a smaller data set.

If the notification is a recurring SAR sentinel error, system performance might be the cause of the issue.

Process exceeds allowed run time

38750122 - Process takes too long to execute. The maximum default time is 3600 seconds.

Explanation

The default time limit of 1 hour for an individual process to complete a task is exceeded.

User response

Review the running process to determine whether the task is a process that can continue to run or must be stopped.

License expired

38750123 - An allocated license has expired and is no longer valid.

Explanation

When a license expires on the console, a new license must be applied. When a license expires on a managed host, the host context is disabled on the managed

host. When the host context is disabled, the appliance with the expired license cannot process event or flow data.

User response

To determine the appliance with the expired license, click the **Admin** tab, click **System and License Management**. A system that has an expired license displays an invalid status in the **License Status** column.

External scan of an unauthorized IP address or range

38750126 - An external scan execution tried to scan an unauthorized IP address or address range.

Explanation

When a scan profile includes a CIDR range or IP address outside of the defined asset list, the scan continues. However, any CIDR ranges or IP addresses for assets that are not within your external scanner list are ignored.

User response

Update the list of authorized CIDR ranges or IP address for assets that are scanned by your external scanner. Review your scan profiles to ensure that the scan is configured for assets that are included in the external network list.

Time synchronization failed

38750129 - Time synchronization to primary or Console has failed.

Explanation

The managed host cannot synchronize with the console or the secondary HA appliance cannot synchronize with the primary appliance.

Administrators must allow **rdate** communication on port 37. When time synchronization is incorrect, data might not be reported correctly to the console. The longer the systems go without synchronization, the higher the risk that a search for data, report, or offense might return an incorrect result. Time synchronization is critical to successful requests from managed host and appliances

User response

Contact customer support.

Cyclic custom rule dependency chain detected

38750131 - Found custom rules cyclic dependency chain.

Explanation

A single rule referred to itself directly or to itself through a series of other rules or building blocks. The error occurs when you deploy a full configuration. The rule set is not loaded.

User response

Edit the rules that created the cyclic dependency. The rule chain must be broken to prevent a recurring system notification. After the rule chain is corrected, a save automatically reloads the rules and resolves the issue.

Blacklist notification

38750136 - The Asset Reconciliation Exclusion rules added new asset data to the asset blacklists.

Explanation

A piece of asset data, such as an IP address, host name, or MAC address, shows behavior that is consistent with asset growth deviations.

An *asset blacklist* is a collection of asset data that is considered untrustworthy by the Asset Reconciliation Exclusion CRE rules. The rules monitor asset data for consistency and integrity. If a piece of asset data shows suspicious behavior twice or more within 2 hours, that piece of data is added to the asset blacklists. Subsequent updates that contain blacklisted asset data are not applied to the asset database.

User response

- In the notification description, click **Asset Reconciliation Exclusion rules** to see the rules that are used to monitor asset data.
- In the notification description, click **Asset deviations by log source** to view the asset deviation reports that occurred in the last 24 hours.
- If your blacklists are populating too aggressively, you can tune the Asset Reconciliation Exclusion rules that populate them.
- If you want the asset data to be added to the asset database, remove the asset data from the blacklist and add it to the corresponding asset whitelist. Adding asset data to the whitelist prevents it from inadvertently reappearing on the blacklist.
- Review the asset reconciliation documentation.

Asset growth deviations detected

38750137 - The system detected asset profiles that exceed the normal size threshold.

Explanation

The system detected one or more asset profiles in the asset database that show deviating or abnormal growth. Deviating growth occurs when a single asset accumulates more IP addresses, DNS host names, NetBIOS names, or MAC addresses than the system thresholds allow. When growth deviations are detected, the system suspends all subsequent incoming updates to these asset profiles.

User response

Determine the cause of the asset growth deviations:

- Hover your mouse over the notification description to review the notification payload. The payload shows a list of the top five most frequently deviating

assets. It also provides information about why the system marked each asset as a growth deviation and the number of times that the asset attempted to grow beyond the asset size threshold.

- In the notification description, click **Review a report of these assets** to see a complete report of asset growth deviations over the last 24 hours.
- Review the documentation about asset growth deviations.

Expensive custom properties found

38750138 - Performance degradation was detected in the event pipeline. Expensive custom properties were found.

Explanation

During normal processing, custom event and custom flow properties that are marked as optimized are extracted in the pipeline during processing. The values are immediately available to the custom rules engine (CRE) and are routed directly to storage.

Improperly formed regular expression (regex) statements can cause events to be incorrectly routed directly to storage.

User response

Select one of the following options:

- Review the payload of the notification. If possible, improve the regex statements that are associated with the custom property.
- Modify the custom property definition to narrow the scope of categories that the property tries to match.
- Specify a single event name in the custom property definition to prevent unnecessary attempts to parse the event.

Raid controller misconfiguration

38750140 - Raid Controller misconfiguration: Hardware Monitoring determined that a virtual drive is configured incorrectly.

Explanation

For maximum performance, raid controllers cache and battery backup unit (BBU) must be configured to use write-back cache policy. When write-through cache policy is used, storage performance degrades and might cause system instability.

User response

Review the health of the battery backup unit. If the battery backup unit is working correctly, change the cache policy to write-back.

An error occurred when the log files were collected

38750141 - Collecting the required support logs failed with errors. See System and License Manager.

Explanation

Errors were encountered while the log files were being collected. The log file collection failed.

User response

To view information about why the collection failed, follow these steps:

1. Click **System and License Manager** in the notification message.
2. Expand **System Support Activities Messages**.
3. View additional information about why the log file collection failed.

Expensive DSM extensions were found

38750143 - Performance degradation was detected in the event pipeline. Expensive DSM extensions were found.

Explanation

A log source extension is an XML file that includes all of the regular expression patterns that are required to identify and categorize events from the event payload. Log source extensions might be referred to as *device extensions* in error logs and some system notifications.

During normal processing, log source extensions are executed in the event pipeline. The values are immediately available to the custom rules engine (CRE) and are stored on disk.

Improperly formed regular expressions (regex) can cause events to be routed directly to storage.

User response

Select one of the following options:

- Review the payload of the notification. If possible, improve the regex statements that are associated with the device extension.
- Ensure that the log source extension is applied only to the correct log sources.

  On the **Admin** tab, click **System Configuration** > **Data Sources** > **Log Sources**. Select each log source and click **Edit** to verify the log source details.
- If you are working with batch log sources, modify the event throttle to ensure that the events do not buffer to disk. The event throttle settings are part of the protocol configuration for the log source.

Chapter 4. Information notifications for QRadar appliance

IBM Security QRadar provides information messages about the status or result of a process or action

Automatic updates successfully downloaded

38750068 - Automatic updates successfully downloaded. See the Auto Updates log for details.

Explanation

Software updates were automatically downloaded.

User response

Click the link in the notification to determine whether any downloaded updates require installation.

Automatic update successful

38750070 - Automatic updates completed successfully.

Explanation

Automatic software updates were successfully downloaded and installed.

User response

No action is required.

SAR sentinel operation restore

38750072 - SAR Sentinel: normal operation restored.

Explanation

The system activity reporter (SAR) utility detected that your system load returned to acceptable levels.

User response

No action is required.

Disk usage returned to normal

38750077 - Disk Sentry: System Disk Usage Back To Normal Levels.

Explanation

The disk sentry detected that the disk usage is below 90% of the overall capacity.

User response

No action is required.

An infrastructure component was repaired

38750084 - Corrupted infrastructure component repaired.

Explanation

A corrupted component that is responsible for host services on a managed host was repaired.

User response

No action is required.

Disk storage available

38750093 - One or more storage partitions that were previously inaccessible are now accessible.

Explanation

The disk sentry detected that the storage partition is available

User response

No action is required.

License near expiration

38750124 - A license is nearing expiration. It will need to be replaced soon.

Explanation

The system detected that a license for an appliance is within 35 days of expiration.

User response

No action is required.

License allocation grace period limit

38750125 - An allocated license's grace period is almost over, and will be allocated in to place soon.

Explanation

The system detected that a license change for an appliance is within the license grace period.

An administrator can move unlocked licenses or apply unused event or flow licenses to other appliances in your deployment. When you allocate a license to a

host, a grace period of 14 days for the license begins. After the grace period expires, the license cannot be moved.

User response

No action is required.

Log files were successfully collected

38750142 - The required support logs have been successfully collected. See System and License Manager.

Explanation

The log files were successfully collected.

User response

To download the log file collection, follow these steps:

1. Click **System and License Manager** in the notification message.
2. Expand **System Support Activities Messages**.
3. Click **Click here to download file**.

Notices

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

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

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

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

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

IBM Corporation
170 Tracer Lane,
Waltham MA 02451, USA

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

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

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

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

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

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

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

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

Trademarks

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

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

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

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s session id for purposes of session management and authentication. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Index

R

S

T

V

Printed in USA