STIBO SYSTEMS
MASTER DATA MANAGEMENT

USER GUIDE

Performance Analysis

2024.4 – December 2024

Table of Contents

Table of Contents **2**

Performance Analysis **3**

Setup Requirements 3

Performance Analysis Tools 4

Health checks 4

Activity tree 7

Events 8

Using Performance Analysis Tools 9

Accessing Performance Analysis Tools **11**

Prerequisites 11

Access 11

Healthcheck Test Index **14**

Healthcheck Tests 14

Performance Analysis Workbench Integration **25**

BG Processes 25

Workbench Notifications 25

Performance Analysis

The Performance Analysis tools provide a number of useful screens, enabling administrators to analyze STEP events, activity, and results of the scheduled health checks. Only the health checks that are relevant to customers and provide actionable information to help improve system performance are included in the scheduled health checks. Some functions available within the admin portal are useful only for Stibo Systems Technical Support and/or R&D groups, while others are applicable for any system administrators.

- Monitor performance trends for opportunities to improve system performance.
- Identify whether business rule logic, data objects, and/or event handling are contributing to poor system performance.
- Troubleshoot cases where a performance impact has been observed.

The Performance Analysis tools can be accessed via the Performance Analysis link found on the Resources section of the Start Page. This separate user interface contains three analysis tools: Health checks, Activity tree, and Events. Using the 'Health checks' screen, in conjunction with the 'Activity tree' and Events graph, can help to identify slow-running processes and correlate issues with event queue activity to identify whether business rule logic, data objects, and/or event handling are contributing to poor system performance.

For a list of health checks that are scheduled to run, refer to the **Healthcheck Test Index** topic in the **Admin Portal** documentation.

Setup Requirements

There are two configuration properties associated with the Performance Analysis tools:

- `HealthCheckScheduler.AutoRun`
  - Allows the user to disable, enable, or skip the automated execution until a date in the future.
  - Possible values are '`true`,' '`false`,' and a date specified by the user.

    > **Note:** When a future date is specified, the schedule does not run until that date. Once the specified date has past, the schedule is enabled.

- `HealthCheckScheduler.OperatingDayAndHours`
  - Allows the user to schedule the start day(s) of the week, with possible values of: `sun, mon, tue, wed, thu, fri, sat`
  - Allows the user to schedule the start and end time of scheduled health checks, with possible values for the window of: `[0-9][0-9];[0-9][0-9]`

    Example: `sat 4-8` means that the scheduled health check analysis begins on Saturdays at 4 a.m. with the list of pre-defined health checks and stops new executions at 8 a.m. Ideally, the window size is large enough to run the scheduled health checks during that period of time.

Once the configuration proprieties have been set, the scheduled health checks can be monitored via the Performance Analysis tools. For information on how to access the Performance Analysis tools, refer to the **Accessing Performance Analysis Tools** topic.

Performance Analysis Tools

The 'Health checks', 'Activity tree,' and 'Events' tab give access to the types of performance analysis tools.

Health checks

The 'Health checks' screen displays a table listing the health checks completed for the scheduled runtime, sorted by the most current execution and providing 12 weeks of history to help identify patterns over time. Taking action on the identified data and configuration issues helps prevent issues from growing to a level where the issues are more difficult to resolve.

The following is a list of the columns associated with the 'Health checks' screen:

- **Checkbox** – The user can select none, one, or multiple rows. Clicking on the checkbox in the header selects all or deselects all. No further action can be taken.
- **Name** – The name of the health check that has been run. This column can be filtered to display one name, or select 'All' to display all health check names.
- **Category** – Displays the category of the health check. This column can be filtered for either the performance, data error, or configuration.
- **Issues** – The number of issues associated with that specified health check. This column can be filtered by a minimum and maximum number of issues.
- **Severity** – Displays the severity of the health check. This filterable column can show all, critical, high, or medium severity issues.
- **Status** – Displays whether the scheduled health check succeeded or failed. This column can be filtered to show all, succeeded, or failed.
- **Completed** – Displays the date and time in which the health check was completed.

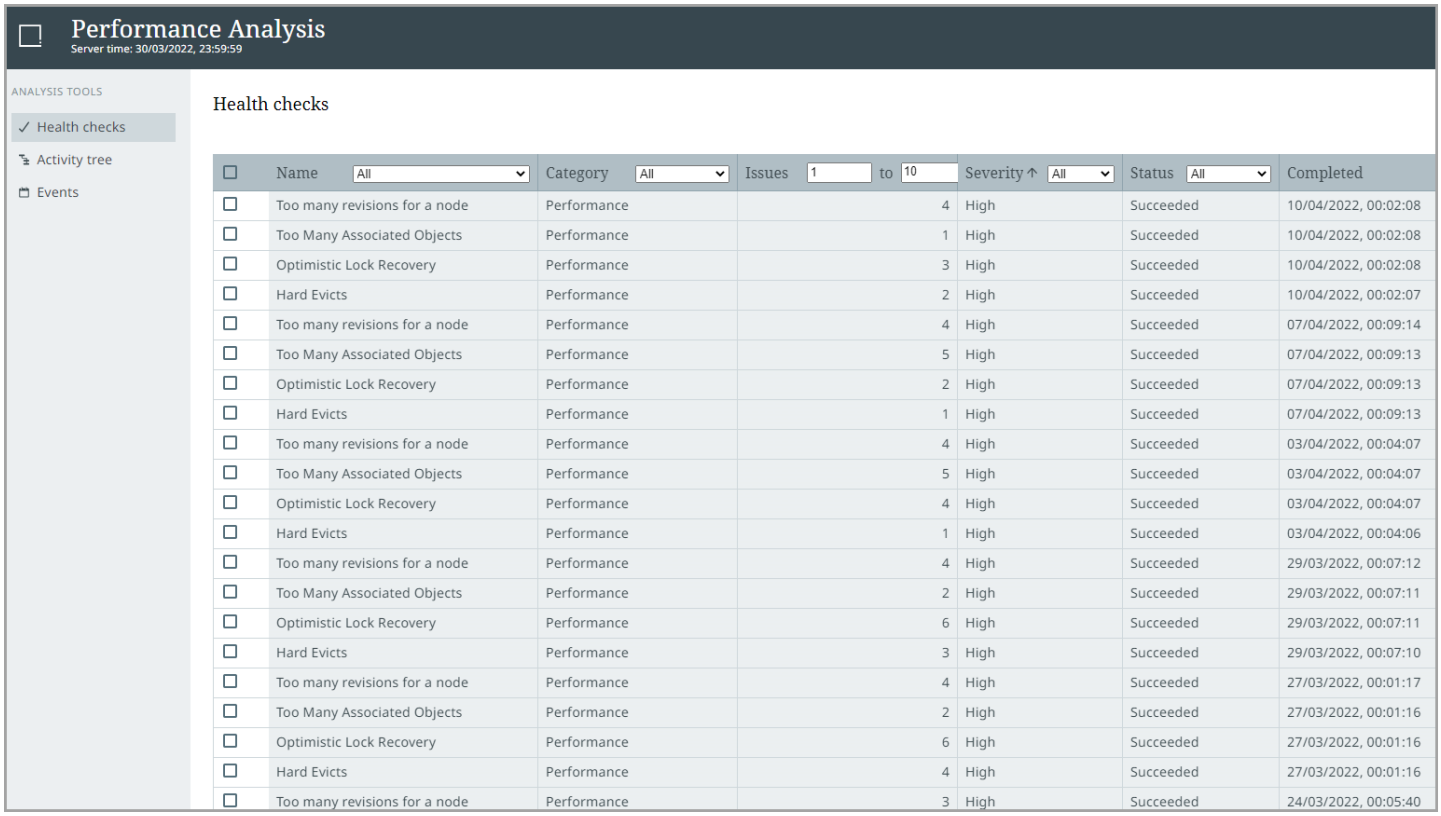

Name	Category	Issues	Severity	Status	Completed
Too many revisions for a node	Performance	4	High	Succeeded	10/04/2022, 00:02:08
Too Many Associated Objects	Performance	1	High	Succeeded	10/04/2022, 00:02:08
Optimistic Lock Recovery	Performance	3	High	Succeeded	10/04/2022, 00:02:08
Hard Evicts	Performance	2	High	Succeeded	10/04/2022, 00:02:07
Too many revisions for a node	Performance	4	High	Succeeded	07/04/2022, 00:09:14
Too Many Associated Objects	Performance	5	High	Succeeded	07/04/2022, 00:09:13
Optimistic Lock Recovery	Performance	2	High	Succeeded	07/04/2022, 00:09:13
Hard Evicts	Performance	1	High	Succeeded	07/04/2022, 00:09:13
Too many revisions for a node	Performance	4	High	Succeeded	03/04/2022, 00:04:07
Too Many Associated Objects	Performance	5	High	Succeeded	03/04/2022, 00:04:07
Optimistic Lock Recovery	Performance	4	High	Succeeded	03/04/2022, 00:04:07
Hard Evicts	Performance	1	High	Succeeded	03/04/2022, 00:04:06
Too many revisions for a node	Performance	4	High	Succeeded	29/03/2022, 00:07:12
Too Many Associated Objects	Performance	2	High	Succeeded	29/03/2022, 00:07:11
Optimistic Lock Recovery	Performance	6	High	Succeeded	29/03/2022, 00:07:11
Hard Evicts	Performance	3	High	Succeeded	29/03/2022, 00:07:10
Too many revisions for a node	Performance	4	High	Succeeded	27/03/2022, 00:01:17
Too Many Associated Objects	Performance	2	High	Succeeded	27/03/2022, 00:01:16
Optimistic Lock Recovery	Performance	6	High	Succeeded	27/03/2022, 00:01:16
Hard Evicts	Performance	4	High	Succeeded	27/03/2022, 00:01:16
Too many revisions for a node	Performance	3	High	Succeeded	24/03/2022, 00:05:40

If the user hovers next to a health check checkbox, a blue expansion arrow is shown, allowing the user to click through to obtain more detail about a specific health check:

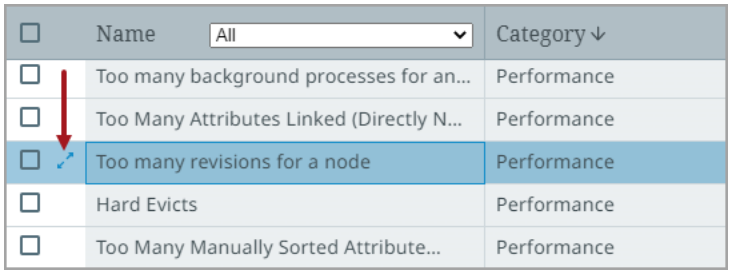

Name	Category
Too many background processes for an...	Performance
Too Many Attributes Linked (Directly N...	Performance
Too many revisions for a node	Performance
Hard Evicts	Performance
Too Many Manually Sorted Attribute...	Performance

In this expanded view, more information is displayed about the health check, including a description of the check.

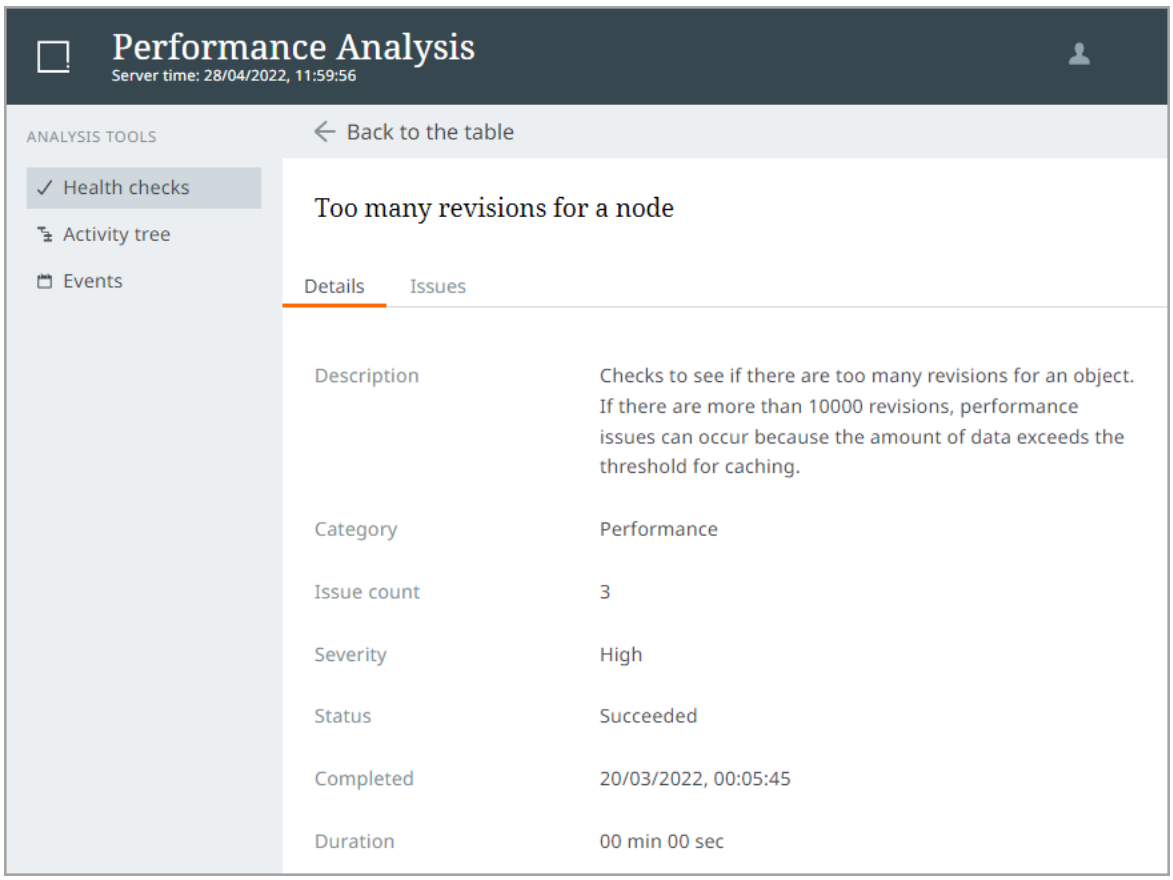

On the Issues tab, additional details about the health check are available, including object information, issue details, whether there is a fix available, if it has been fixed via Stibo Systems Support, and username of who applied the fix (if applicable).

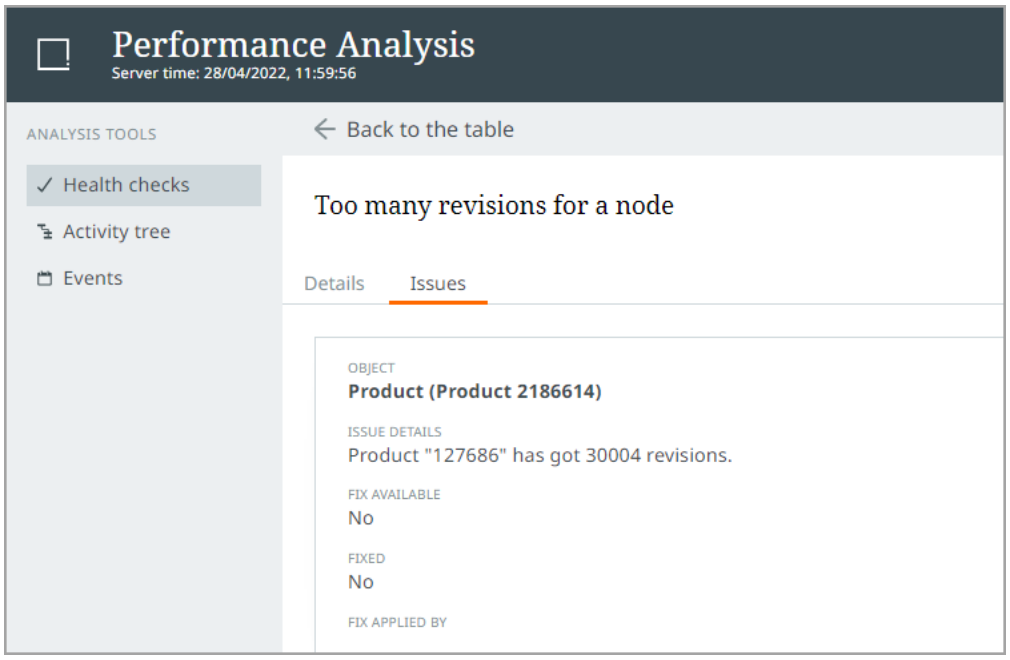

Activity tree

In the 'Activity tree' screen, a user can investigate short-term performance issues when there is a known problem (e.g., overall system slowness, BGP not progressing, etc.) but the cause is unclear. This drill-down tree helps the user discover the problematic activity and/or process, in addition to determining the possible cause. The 'Activity tree' prioritizes items with the most activity at the top, allowing the user to review which processes are using the highest percentage of system resources.

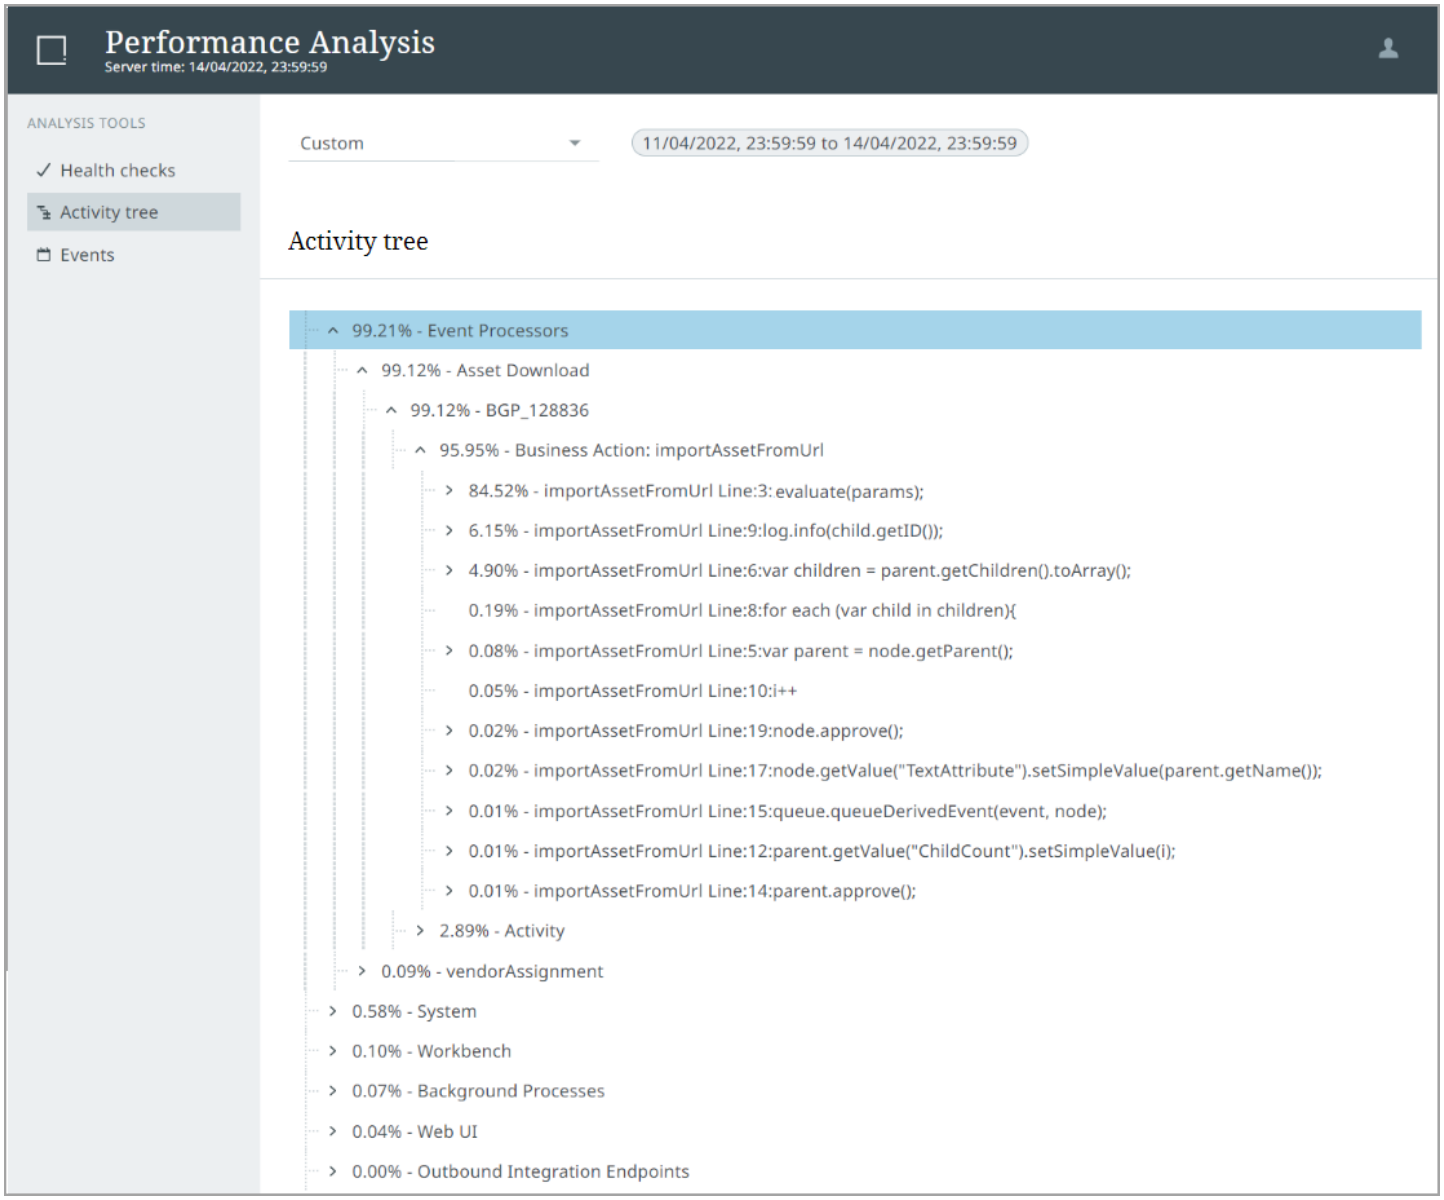

Note: The percentages detailed in the 'Activity tree' are relative to the used activity, not that all of the system resources are being used by these particular activities.

The 'Activity tree' can be run for the following time periods, allowing the user to customize the results:

- Last 5 minutes
- Last 30 minutes
- Last 1 hour
- Last 2 hours
- Last 4 hours
- Last 8 hours
- Last 1 day
- Custom

Note: The selected time period also applies to the Events screen. Alternatively, if the user selects a time period on the Events screen first, that is reflected on the 'Activity tree' screen.

Events

On the Events screen, the 'Events created vs. consumed' graph allows the user to visualize event processing in two ways. The Series tab allows the user to review a line graph of all the events created versus those consumed within the defined time period. Mouse-over text shows the three event queues with the highest count of events for the selected date and time.

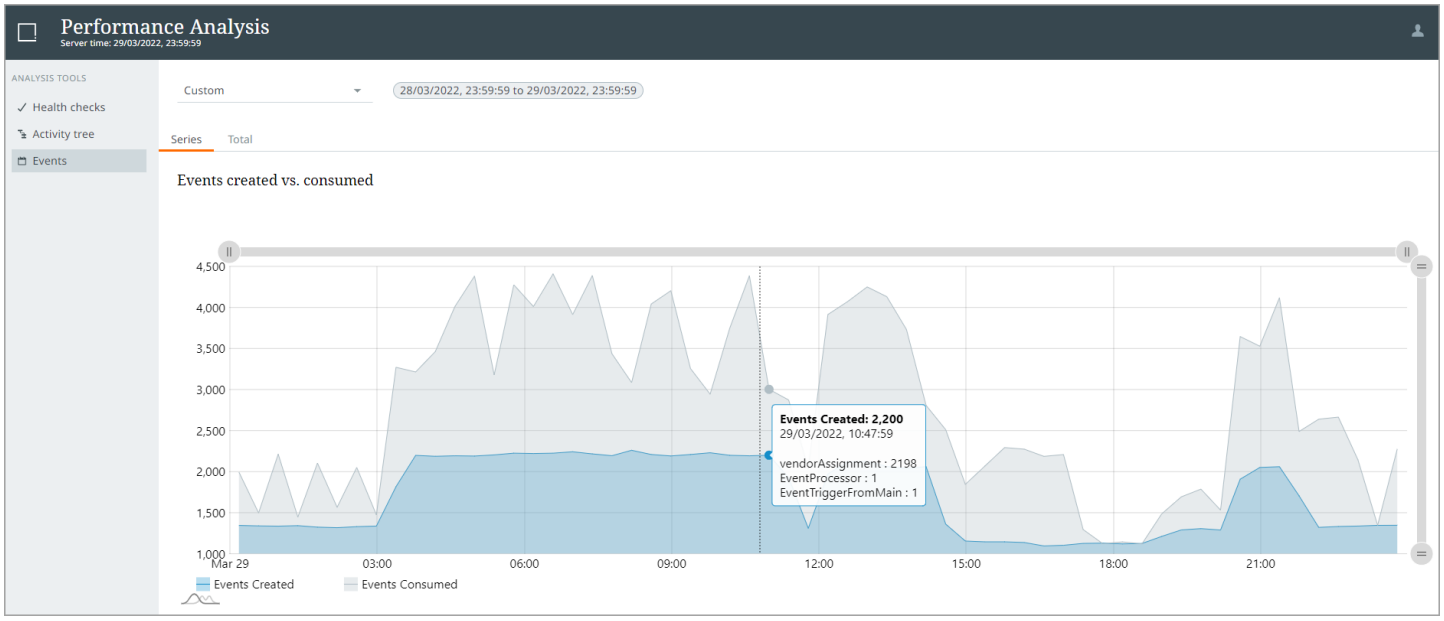

The Total tab displays the events created versus those consumed with individual queue counts in a bar chart.

Like the 'Activity tree', the Events screen can be run for the following time periods:

- Last 5 minutes
- Last 30 minutes
- Last 1 hour
- Last 2 hours
- Last 4 hours
- Last 8 hours
- Last 1 day
- Custom

Note: The selected time period also applies to the 'Activity tree' screen. Alternatively, if the user selects a time period on the 'Activity tree' screen first, that is reflected on the Events screen.

Using Performance Analysis Tools

This is an example of how a user might benefit from Performance Analysis tools.

Consider that a system is configured to trigger an Event Processor executing a business action following the import of products, which triggers an Outbound Integration Endpoint, sending the changes to a downstream system. This flow can quickly create a high volume of events in the system, assuming the import volume is high.

The Events graph shows the rate of event creation on one line and the rate of event consumption on another line of the same graph, which shows how well the system is able to process changes. A minor phase-shift is typical between creation and consumption of events, based on schedules and execution time of the business action.

In this example, we assume the action takes more than a few seconds to execute. This could be because:

- the action is written inefficiently
- the action encounters objects that have a lot of children or references or values or revisions
- there are time delays while concurrent updates to the same object in a different transaction are retried
- or other issues.

In any of these cases, the consumption of events lags behind the creation significantly. When this occurs, the weekly executed health checks indicate issues, detailing which objects, the number of occurrences, timestamps, and the related BGP ID.

Using the information from the health checks and lag in event processing, the 'Activity tree' screen is useful to drill into an identified time frame to help indicate the configuration areas where the system is spending the most time. Exploring the 'Activity tree' for the highest percentages below the Event Processors exposes the ID of the problematic Event Processor, the BGP ID, the ID of the Business Action and the lines of code where the system is spending the most amount of time.

By combining the information in these screens, the administrator is empowered to quickly identify the contributors to performance issues, continue to research, modify and test to resolve the identified issues, which may involve rewriting a business action, adding an Event Processor to manage revisions, restructuring hierarchies, or other configurations changes, in order to optimize the system performance.

Accessing Performance Analysis Tools

This topic describes how to access Performance Analysis tools, including what privileges are required to do so.

Prerequisites

To access the Performance Analysis tools, the person attempting to log in must have a user created in STEP that is a member of a user group with privileges that include the following setup actions:

- View Administration
- View Context

For example:

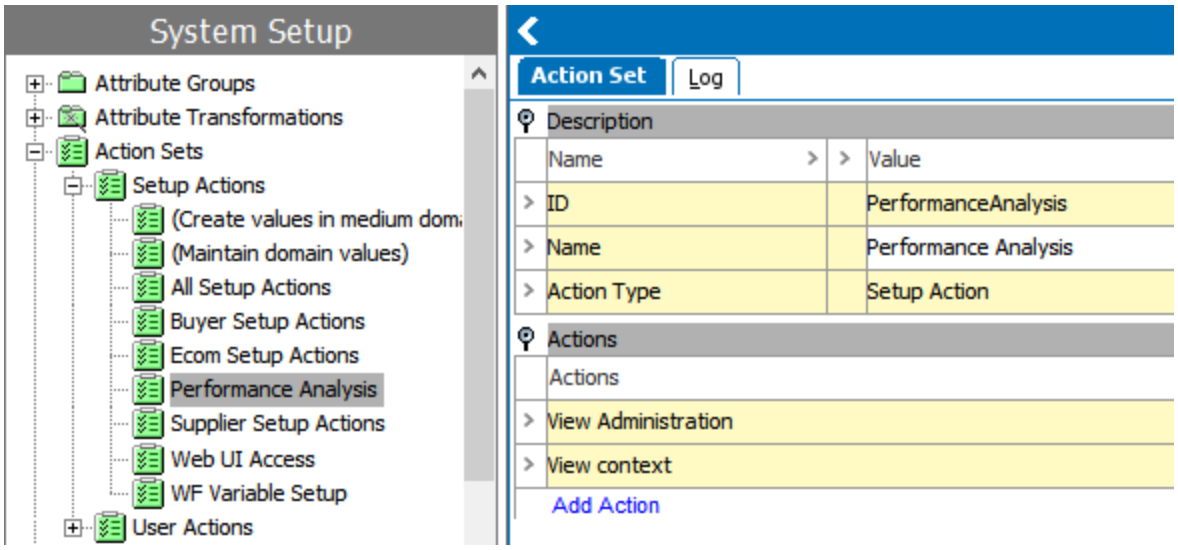

Additional information on Actions, Users, and User Groups can be found in the **System Setup** documentation, under **Action Sets** and **Users and Groups**, respectively.

Access

The Performance Analysis page can be accessed by clicking the Performance Analysis link from the Resources section of the Start Page

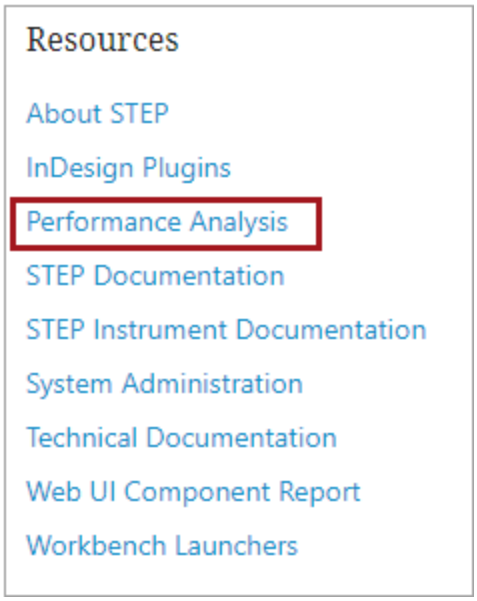

This displays a login prompt:

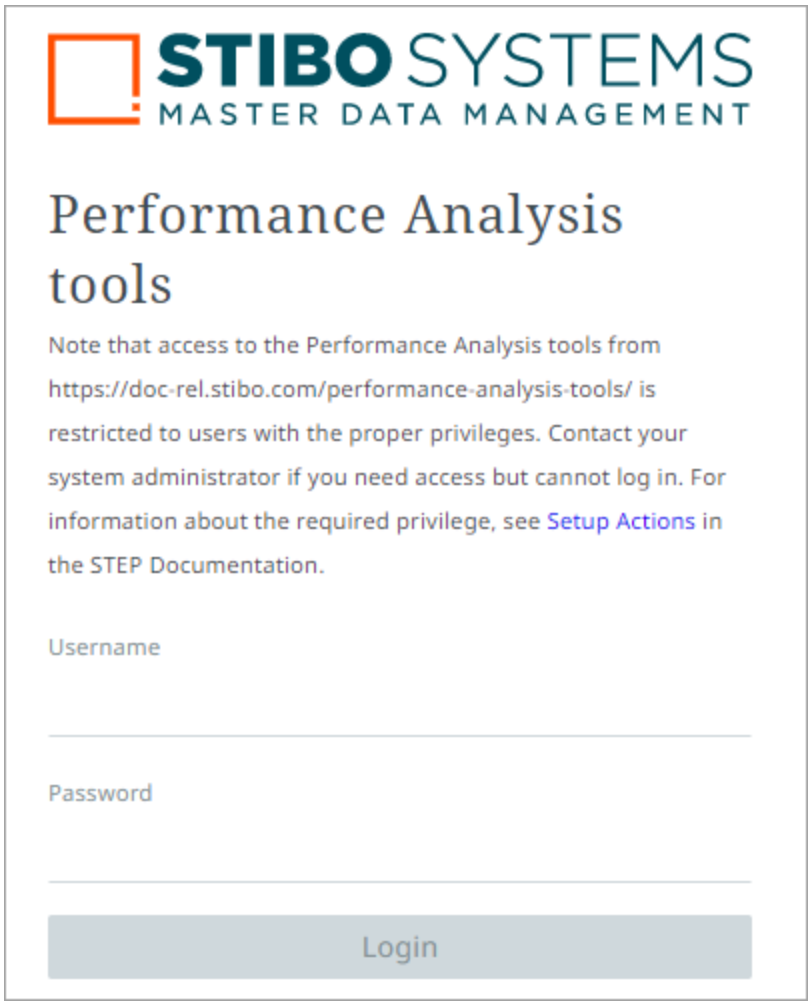

Valid credentials must be entered for login, as indicated in the Prerequisites section of this topic. Once logged in, the Performance Analysis page displays a series of screens, with various functions available on each screen. The remainder of this guide provides details on these tabs and functions.

Healthcheck Test Index

Healthchecks assist users to identify and resolve configuration and data issues that can negatively affect system performance.

Healthchecks are executed or skipped based on the database in use and/or if in-memory is enabled on the STEP system, so not all healthchecks will run on every system. As available, healthchecks can be reviewed and run from the following locations:

- In the Admin Portal on the **Healthcheck** tab, users can run tests and review detected problems as needed. For more information on the Admin Portal Healthcheck tab, refer to the **Healthcheck** topic.
- For on-premise systems, healthcheck information is stored on the application server at [STEPHOME]/diag/healthcheck (for example, opt/stibo/step/diag/healthcheck). This information is automatically included when sending a diagnostics package to Stibo Systems Support.
- From the Start Page, the **Performance Analysis** link displays 12 weeks of results from all scheduled healthcheck results. Some unscheduled healthchecks are long-running or are only useful for Stibo Systems Support, and so they are not available in the Performance Analysis tools.

Healthcheck Tests

The following tables include all available Configuration, Data Error, and Performance healthchecks.

Note: Not all healthchecks are applicable for all STEP systems. On your system, only the healthchecks that are valid are displayed in the Admin Portal, on the application server, and in the Performance Analysis tools.

- The 'Automated Fix' column indicates if a script is available to resolve the reported issue. To apply a scripted fix, contact Stibo Systems Support for assistance. If no automated fix is available, manually update the reported data or configuration.
- The 'Runs on Schedule' column indicates that the test is run on the schedule defined in the sharedconfig.properties file on the application server, as defined in the **Performance Analysis** topic.

The information in the tables below is best viewed online.

Performance Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
Business Rule Execution Time Too Long	High	Performance can be impacted when business rules run too long. Business rules may run too long when there are too many operations combined into one rule, or when accessing too many objects, or when accessing objects with too many revisions, or when configurations call external services with a slow response time, etc. This is not always a problem, for example, if an IEP using the Business Rule Based Message Processor runs too long when processing	No	Yes

Performance Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		batches, performance is not necessarily impacted if large transactions are not being written to the database. Most often, business rules taking longer than one minute require examination of the rule itself for performance improvements or the objects involved for data management. By default, business rules that run longer than five (5) minutes are reported as a healthcheck warning and business rules that run longer than 15 minutes are stopped and generate an error.		
Change Log Entries Per Node	Low	When modifying certain objects in STEP, a change is written to the change log of the object. You can set event queues in STEP to monitor on these events and via, for example, an integration endpoint, information about the change can be exported. An event is put the queue for each interested event queue. Every time you modify an object that generates an event, STEP tries to limit the number of log entries for that object. If more than 20,000 changes are logged for the object, it attempts to delete the old events. However, the attempt only succeeds if there are no event for the change. If you have more than 20,000 change log entries for an object, determine why the events are not being processed.	Yes, in most cases	No
Change Log Total Size	Critical	Checks if the change log has grown too large. This can cause Oracle to perform poorly. The maximum number of rows allowed in the table is 100,000,000.	Yes	No
Check for Common Web UI Configuration Errors	Medium	Checks Web UI for some of the most common configuration errors that can cause performance problems.	No	No
Data cleanup tools	High	To maintain a system that performs well, regularly scheduled data clean-up is highly recommended. This healthcheck detects problems with the configuration of a scheduled background process to empty the recycle bin and/or an event processor to purge old revisions. Configure 'Schedule Empty Recycle Bin' , running at least monthly, and including all contexts in use by your system, to purge items in the Tree recycle bin. Enable an event processor of type 'Revision Management' scheduled to run frequently, with Purge Across Workspaces set to Yes, and Number to Keep set to not much more than 100, to reduce the	No	Yes

Performance Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		number of unnecessary revisions.		
Hard Evicts	High	A hard evict is a forced attempt to remove persistent objects from their cache. Hard evicts can happen when a task is holding many persistent objects for a long time without committing the transaction. In such a case, a hard evict may be executed to make room for caching other persistent objects. This can negatively affect performance since the cache of persistent objects may become less effective. A cause of hard evicts could be one or more business rules accessing too many objects.	No	Yes
Large Commit	High	For Cassandra systems, large commits decrease system performance and increase the risk of concurrency problems. Very large commits may fail.	No	Yes
Leaked Changelog Rows	High	Checks if there are leaked rows in the change log table. If there are many leaked data rows, performance will be negatively affected.	Yes	Yes
Optimistic Lock Recovery	High	Reports that optimistic locking errors were detected when flushing to the data store. This indicates that some objects were concurrently modified in another transaction, or a constraint error occurred. This can negatively impact performance. Repeated occurrences of this may cause the transaction to eventually fail. Resolve this by avoiding and/or minimizing concurrent modifications of the same data.	Yes	Yes
Too Many Associated Objects	High	When there are too many associated objects, degradation of performance is possible because the amount of data exceeds the threshold for caching the given relation. This could be due to too many children, references, referenced by values, or multi-valued attributes.	No	Yes
Too Many Attributes Linked (Directly Not Via Inheritance) to a Product/Classification	Medium	Finds all products / classifications that are directly linked (not inherited) to more than 1,000 valid attributes. More than 1,000 links can cause performance issues when opening the References Editor in workbench.	No	Yes
Too Many Background Processes for an Integration Endpoint	High	Checks if there are too many background processes for an integration endpoint. Too many BGPs for an IEP can degrade	No	Yes

Performance Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		performance. Clean-up of old BGP files and folders is required to resolve this issue.		
Too Many Manually Sorted Attribute Groups	Medium	Checks that no manually sorted attribute group has more than 10,000 children. Only the front revisions are considered and children in all workspaces are counted.	No	Yes
Too Many Manually Sorted Products and Classifications	Medium	Checks that no manually sorted attribute group has more than 10,000 children. Only the front revisions are considered and children in all workspaces are counted.	No	Yes
Too Many Qualifier Relations	Low	Find all qualifiers that are used in too many pseudo qualifiers. Performance problems can result from having a large number of pseudo qualifiers if a real qualifier is linked to large number of pseudo workspaces because, by default, the application cache only caches 10,000. Refer to the property: Install.DataCache.MaxRelationSize=10000. This plugin cannot remove the duplicates, but another plugin can remove the unused pseudo qualifiers.	No	No
Too Many Revisions for a Node	High	Checks if there are too many revisions for an object. More than 10,000 revisions can cause performance issues because the amount of data exceeds the threshold for caching.	No	Yes
Too Many Valid Values for List of Values	Medium	Checks that no list of values has more than 5,000 valid values. Large lists of values (LOVs) make it difficult to find, search, select, and filter on values.	No	No
Too Many Values for a Node	Medium	Checks if there are nodes with too many values, which can cause performance issues.	No	Yes
Too Many Workspace Relations	Low	Finds all workspaces are used in too many pseudo workspaces. If, for example, a node is visible in the Main, Approved, and Staging workspaces, a pseudo workspace representing these three workspaces is created. Performance problems can result from having a large number of pseudo workspaces if a real workspace is linked to large number of pseudo workspaces. The application cache, by default, only caches 10,000 pseudo workspaces. Refer to the property: Install.DataCache.MaxRelationSize=10000.	No	No

Performance Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		While this plugin cannot remove the duplicates, another plugin can remove the unused pseudo workspaces.		
Unused Pseudo Qualifiers	Low	Finds all pseudo qualifiers that are not used. Performance problems can result from having a large number of pseudo qualifiers if a real qualifier is linked to large number of pseudo qualifiers. The application cache, by default, only caches 10,000. Refer to the property: Install.DataCache.MaxRelationSize=10000. Missing qualifiers are only reported when at least 5,000 unused qualifiers exist.	Yes	No
Unused Pseudo Workspaces	Low	Finds all pseudo workspaces that are not used. If, for example, a node is visible in the Main, Approved, and Staging workspaces, a pseudo workspace representing these three workspaces is created. If you create new workspaces, many new pseudo workspaces can display many combinations of data. In this case, the result is a lot of pseudo workspaces, while any of these combinations are not always used. Performance problems result from having a large number of pseudo workspaces if a real workspace is linked to large number of pseudo workspaces. The application cache, by default, only caches 10,000 of these. Refer to the property: Install.DataCache.MaxRelationSize=10000.	Yes	No

Data Error Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
Assets Without a History Entry	Medium	Assets cannot be found or viewed in the workbench or in Web UI because there are no visible entries in the history.	No	No
Attributes / Products Which Cannot Be Approved	Medium	These attributes and/or products cannot be approved.	No	Yes
Attributes That Have Both Revised and Not Revised (Externally Maintained / Not	Medium	Finds attribute values where the workspaces are not in agreement. For example, attribute values that are visible in all workspaces and are not externally maintained, or attribute values that are not visible in all workspaces and are externally	Yes	Yes

Data Error Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
Externally Maintained)		maintained.		
Background Processes Incorrectly Linked to Integration Endpoint	High	Verifies if BGPs are correctly linked to an integration endpoint. If not correctly linked, this can cause multiple pollers to be started for IEPs which can cause issues.	Yes	Yes
Check LOV Used for Status by BGPs	High	Checks for duplicate values in the LOVs used for the status of BGPs. Duplicates can cause locking errors when setting the status of the BGP.	Yes	Yes
Check Sequences	High	Identifies when a production database has been incorrectly copied to a Test / QA system, without including Oracle sequences. Other elements like views might also be missing.	No	No
Cycles in a Translation Graph	Medium	Finds nodes where the translations are caught in an infinite loop, like when the source and target of the translation are the same language.	Yes	Yes
Dual Visibility Values With Different Values	Low	Generates a list which can help to decide which values should be deleted / kept. Same as 'Values with Dual Visibility' except only duplicates with different values are listed. This list should be viewed by the customer before deletions.	No	No
Duplicated Contexts	Low	Finds duplicated contexts.	Yes	No
Duplicated Contexts, Their Matching Names and Workspaces	Low	Finds duplicated contexts, including their names / IDs, and workspaces. It can be time-consuming to include the names and workspaces.	No	No
Duplicated Edges	Medium	Searches for duplicated references / links with identical parents and children in the same context.	No	Yes
Duplicated History Entries	Medium	Checks if there are duplicated entries in the history table.	No	No
Duplicated PrivilegeRule Ownership	Critical	Finds Privilege Rules that are erroneously shared between multiple User Groups. Inconsistent behavior can result when editing the Privilege Rule and it prevents In-Memory system from starting up.	Yes	No
Duplicated Workspaces	High	Finds all pseudo workspaces that are duplicated. If, for example, a node is visible in the Main,	No	No

Data Error Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		Approved, and Staging workspaces, a pseudo workspace representing these three workspaces is created. However, if another pseudo workspace represents the same three workspaces, it is a duplicate. Performance problems can result from having a large number of pseudo workspaces if a real workspace is linked to a large number of pseudo workspaces. The application cache, by default, only caches 10,000. Refer to the property: Install.DataCache.MaxRelationSize=10000. While this plugin cannot remove the duplicates, another plugin can remove the unused pseudo workspaces.		
Edges With Invalid Revisability	Medium	Finds link types that are marked as revisable but where unrevised links exist or link types that are marked as unrevised but where revisable links exist.	Yes	No
GDSN Subscription Not Linked to a Datapool	High	Checks if there are subscriptions that are not linked to a datapool.	No	No
GDSN Subscription With Invalid GPC	High	Checks if there are subscriptions with invalid GPC codes. For example, GPC codes for which there is no matching classification in STEP.	No	No
Invalid Previous and Maximum History Revisions	Medium	Finds history rows where previous max of new revision row is not equal max revision of previous revision row. On Cassandra systems, this healthcheck identifies data inconsistency issues. Data inconsistency can be caused by JavaScript that catches exceptions but does not rethrow the exception to handle the exception properly, causing partially committed actions writing to the database instead of being rolled back. Refer to the healthcheck 'Javascript catch without rethrow' for possible identified issues.	No	Yes
Invisible Deleted Products Where the Deletion Cannot be Approved	Medium	Identifies when there is a mismatch between two internal tables in STEP for history and references / links. If there is a mismatch, approval will fail and a message such as 'Cannot subtract Main from PWSpace...local to 203, pseudo does not contain the workspace.' is displayed.	No	Yes
Node Collections From Missing Parents	Medium	Finds collections that are missing a link to the parent. This issue has been found with temporary collections that are used by BGPs. When the user	Yes	Yes

Data Error Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		attempts to delete the BGP, files are out of sync and an error is thrown.		
Nodes Having Multiple Parents	Medium	Identifies nodes with multiple parents.	Yes	No
Packaging Hierarchy Loop	Critical	Identifies when a circular reference is created in the packaging hierarchy, which can cause performance issues or throw an unhandled exception. This report identifies where the user can correct the data issue, preventing future performance issues.	No	Yes
Pollers Started by a Different User Than the One Configured in the IEP	High	Checks if there are pollers started by a different user than the one configured in the IEP. This can cause new revisions being generated each time the IEP is invoked, which can in turn cause performance problems.	Yes	No
Revised Node Missing a Front Revision	Medium	Finds all revised attributes that have unrevised values. This occurs when an attribute is changed and the update to values is interrupted, leaving some values revised and other values unrevised.	Yes	No
Revised Values Should be Unrevised	High	Finds all revised attributes that have unrevised values. This occurs when an attribute is changed and the update to values is interrupted, leaving some values revised and other values unrevised.	Yes	Yes
Search for Duplicated Qualifiers	Critical	Searches for duplicated qualifiers for modified values used for export and/or publication. If the value inconsistency exists for a long period of time, data can become corrupt.	No	Yes
Softvalues With Dual Visibility	Medium	Searches for duplicated softvalues.	No	No
The Object Has One or More (LOV) Attribute Values in a Deleted Context	Medium	Searches for objects with LOV attribute values present in a context that has been deleted. This may cause an approval to fail.	No	Yes
The Object Has One or More Attribute Values in a Deleted Context	Medium	Searches for objects with attribute values present in a context that has been deleted. This may cause an approval to fail.	No	Yes
The Object Has One or More Dimension	Medium	Searches for objects with dimension dependent attribute values present in a context that has been	No	Yes

Data Error Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
Dependent Attribute Values in a Deleted Context		deleted. This may cause an approval to fail.		
The Object is Named in a Deleted Context	Medium	Searches for objects with name(s) present in a context that has been deleted. This may cause an approval to fail.	No	Yes
Unrevised Values should be Revised	High	Finds all unrevised attributes that have revised values. This occurs when an attribute is changed and the update to values is interrupted, leaving some values unrevised and other values revised.	Yes	Yes
Value Link With No Owner	Medium	A data inconsistency for values using either List Of Values validation or older soft values, such as some used for system attributes. Applying the fix deletes value links that have a non-existing node. The fix operation does not support fixing individually selected problems - all fixable problems are fixed if the fix is started regardless of which problems are selected.	Yes	Yes
Value Link With No Value Model	Medium	A data inconsistency for values using either List Of Values validation or older soft values, such as some used for system attributes. Objects with this problem fail to load. Applying the fix recreates the missing value model objects where possible. The fix operation does not support fixing individually selected problems - all fixable problems are fixed if the fix is started regardless of which problems are selected.	Yes	Yes
Value Missing Content	Medium	This check finds and fixes problems with missing entries in BLOB tables.	Yes	Yes
Values Have Not Been Marked Correctly Deleted	Low	Finds all nodes / attribute combinations with values that are marked as deleted but that are still visible.	Yes	No
Values With Dual Visibility	High	Identifies if some of the duplicates have different values which can lead to unexpected behavior in the workbench or Web UI. The two values are randomly displayed and STEP appears unstable.	No	No

Configuration Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
Hidden Oracle Parameters With Non-Default Values	Medium	Lists hidden Oracle parameters with a changed default value. The default value of a hidden parameter should only be changed when recommended by Oracle or Stibo Systems Support.	No	No
JavaScript Catch Without Rethrow	High	Identifies business rules that do not correctly handle exceptions in try-catch statements. When catching an exception in JavaScript business rules using try-catch, only checked exceptions that have been declared in the Stibo Systems Scripting API are safe to catch without a rethrow of the same or another exception. All runtime exceptions should be rethrown. For some runtime exceptions, this will be strictly enforced so that if the business rule completes successfully, the exception will be rethrown by the framework when omitted in JavaScript. This protects against possible database inconsistencies that occur when the rethrow is omitted. If an API method partially completed a change when the exception occurred, the database transaction needs to roll back by letting the exception fall through the execution scope of the transaction. When issues are reported in this healthcheck, the system-detected missing rethrow(s) and the reported business rule(s) need to be revised to include a rethrow of the same or another exception.	No	Yes
Non-Compacted Attributes	Medium	Identifies attributes that are not using the compact storage model. Compacted attributes (excluding LOVs) take up less storage space than non-compacted attributes, which results in reduced I/O during read and write, improving response-time of the system. Additionally, for customers migrating to Stibo Systems SaaS (Cassandra), it is a prerequisite and may take multiple	No	Yes

Configuration Healthcheck Name	Severity	Description	Automated Fix	Runs on Schedule
		days to complete. When issues are detected in this healthcheck, review the attributes reported and start the migration to compact soft values, or convert the attributes to LOVs, where feasible (many usages and few distinct values). Refer to the **Attribute Value Migration** topic for prerequisites and technical migration details.		
Reflection usage in business rules	High	This healthcheck detects the use of reflection in business rules when they are executed and a warning with the text 'Attempted to call reflection API...' is written to the step.0.log. A future STEP release will block the use of reflection. For any necessary methods or functions that are not covered in the Public JavaScript API, enter an enhancement request in the Stibo Systems Service Portal for review and approval. Also, plan to rewrite the business rules identified by this healthcheck when preparing for a future upgrade.	No	Yes
Residual Events for a Queue	Medium	Identifies event queues with events not being processed. When a queue-based event processor or outbound integration endpoint is set to 'Read events' but is not enabled, or is stopped in error, or is enabled without a schedule being run, large numbers of events can build up, which can negatively affect performance. If this issue is detected, inspect the objects in the report and make configuration changes to either consume the events or discard them.	No	Yes

Performance Analysis Workbench Integration

This topic details how Performance Analysis tools and scheduled health checks integrate with the workbench. Because the scheduled health checks run automatically, workbench integration is limited to only being able to view certain information about the health checks.

BG Processes

Through the BG Processes 'Health Checks' node, the user can view the scheduled health checks that have been run on a previous execution. The BG Processes display which health checks are being run, which have been skipped, and the approximate time that each test begins. These read-only fields allow the user

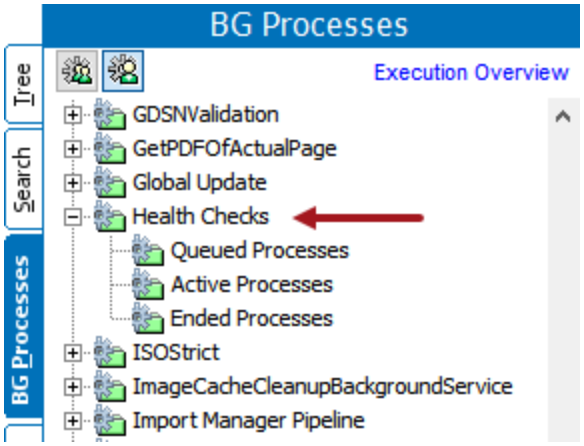

Workbench Notifications

If any of the scheduled health checks detect one or more of the following issue(s), users who have the View Administration privilege get a popup upon opening the workbench, alerting them to review the issue(s) in the Performance Analysis tools:

- One or more scheduled health checks detected an issue with either critical or high severity.
  - The popup displays the sum of the issues found by each health check and the total count of health checks by severity, as shown in the example below:

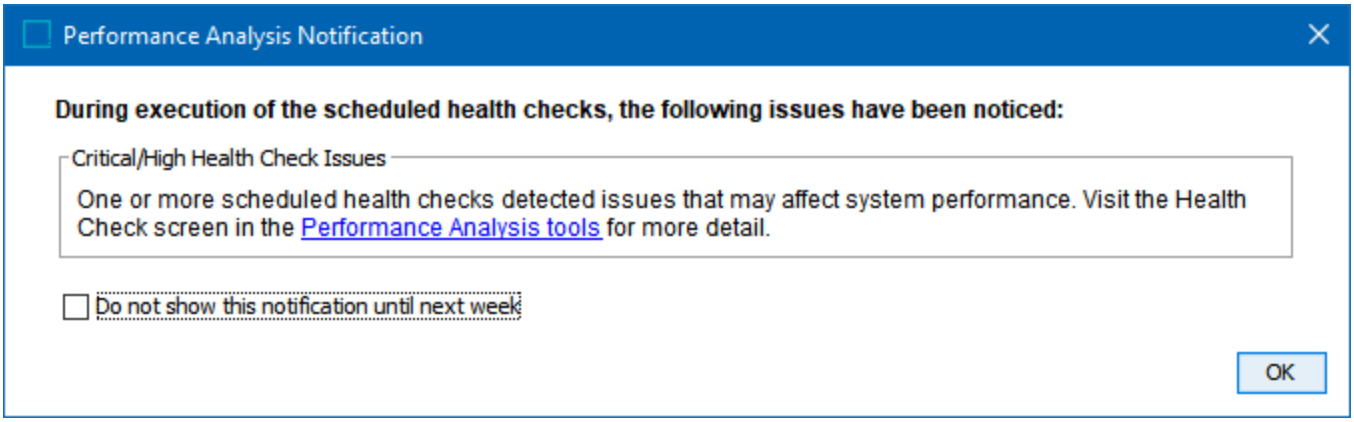

- The schedule to automatically run health checks has been disabled.
  - If this popup occurs, update the `HealthCheckScheduler.AutoRun` configuration property to enable the schedule.
- The operation duration is inadequate to complete all scheduled health checks.
  - When this popup occurs, not all of the scheduled health checks could be completed during the operation window. The oldest health check will be run first during the next execution. Update the `HealthCheckScheduler.OperatingDayAndHours` configuration property to extend the duration of operation to allow for enough time to start all scheduled health checks on a weekly basis.

Note: If the 'Do not show this notification until next week' is checked, the notification is silenced until the popup time resets on Monday at midnight, allowing the popup to be shown the next time the workbench is opened.

To learn more, refer to the **Performance Analysis** documentation.