Hewlett Packard
Enterprise

HPE Application Performance Management

Software Version: 9.30

Service Health Administration - Best Practices

Document Release Date: December 2016
Software Release Date: July 2016

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2016 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

AMD and the AMD Arrow symbol are trademarks of Advanced Micro Devices, Inc.

Google™ and Google Maps™ are trademarks of Google Inc.

Intel®, Itanium®, Pentium®, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

iPod is a trademark of Apple Computer, Inc.

Java is a registered trademark of Oracle and/or its affiliates.

Microsoft®, Windows®, Windows NT®, Windows® XP, and Windows Vista® are U.S. registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/group/softwaresupport/search-result?keyword=.

This site requires an HPE Passport account. If you do not have one, click the **Create an account** button on the HPE Passport Sign in page.

Support

Visit the HPE Software Support website at: **https://softwaresupport.hpe.com**

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software Support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract. To register for an HPE Passport ID, go to **https://softwaresupport.hpe.com** and click **Register**.

To find more information about access levels, go to:
https://softwaresupport.hpe.com/web/softwaresupport/access-levels

HPE Software Integrations and Solutions

Visit the Integrations and Solutions Catalog at **https://softwaresupport.hpe.com/km/KM01702731** to explore how the products in the HPE Software catalog work together, exchange information, and solve business needs.

Contents

Introduction 5
Document Overview 5
Chapter 1: APM Service Health Value Proposition 6
Chapter 2: Define Service Health Requirements 7
Chapter 3: Service Health High-Level Design 8
Chapter 4: Customize Indicators and Rules 9
Adding a Data Point/Dimension in Service Health 9
KPI Assignment 14
Changing the KPI Representation 16
Customizing KPI Presentation for a Specific View 16
Changing a KPI's Calculation Rule 16
Changing a Rule Across APM 16
Using a Subset of a CI's KPIs 21
Displaying a KPI Only 22
Example - Customize KPI Display 23
Chapter 5: Views 27
Working With Views 27
Which view type is suitable for your needs? 28
What are some rules of thumb for creating views? 29
What is a local impact view (LIV)? 29
When should you use a Local Impact View? 30
View Permissions 31
Chapter 6: MyBSM Planning 32
Chapter 7: Use Cases / Troubleshooting 33
Service Health Not Reflecting Your True Business Health 33
Exporting Service Health Data 33
Send Documentation Feedback 35

Introduction

APM Service Health has a clear value proposition, which this document will present. If you are configuring, administering, or consuming Service Health, it is essential that you understand this value proposition, and be able to adapt it to your needs.

Service Health was designed to be adjusted and customized by the end user to meet an organization's specific needs. The wide variety of customization and tuning options can be overwhelming to the inexperienced user. The purpose of this guide is to explain some of the options and adjustments available to the end user.

The objective of this document is to help you maximize the value gained from APM Service Health, by helping you:

- Understand the Service Health value proposition.
- Define the correct set of requirements and use cases.
- Configure and tune Service Health to your specific needs.
- Review some practical Service Health use cases.

Document Overview

This document is intended for APM users who are familiar with APM Service Health. This document is not intended to be a comprehensive or introductory guide. For general getting started and background material refer to the Service Health documentation (for example, within the APM Help see **Application Administration > Service Health > Service Health Overview**). In addition, the level of detail in this document may vary from section to section and does not cover the full range of possibilities. Experienced Service Health users should be able to complete the missing parts by using their own knowledge or by using the Service Health documentation.

The content of this document was developed based on R&D's experience with Service Health implementation in the field. Our intent is to capture the major gaps and use cases where Service Health users can benefit from a lab-issued Best Practice document. Naturally, this implies that not all possible solutions and use cases are covered.

The target audience is primarily the following:

- Service Health administrator.
- Application / Service / Line of Business manager.

We encourage the readers of this document to send us your feedback, to help us improve the document as well as identify other gaps where the lab may assist you.

Chapter 1: APM Service Health Value Proposition

Service Health provides a customizable integrated real time view of your business processes and applications, which can be aligned with the infrastructure supporting the business. This means you can determine the health of your business against predetermined objectives at any given time, while being able to link your business performance and health with the relevant supporting infrastructure. This provides you with an improved basis for communication and collaboration between various business and IT teams.

The following are the main elements of the Service Health value proposition:

- Service Health combines data from all data sources available to your APM system, including third party data.
- The processed data is presented over a shared real time service model, alowing you to map the relations between your business and operations or IT.
- Service Health dimensions are derived from data source types.
- You can define and tune the perspectives and portions of your business you are interested in, as well as the type of data presented, threshold values, and status calculation rules. These are all customizable according to the persona type or group consuming this information.

Chapter 2: Define Service Health Requirements

You should start your Service Health implementation by collecting the Service Health consumer requirements. It is important to map the different types of consumers, the data they are interested in, and the manner by which they would like to consume it. It is also important to understand the work processes or flows within your business and IT organizations, where Service Health will be used.

Example 1: The CIO is interested in a high-level online health view of all business-critical services. The desired aspects are availability and performance; drill-down and over-time capabilities are required in case of availability or performance degradation. The data should be presented in a readable and simplified format. The data is used during a daily status call with LOB and IT managers, and presented 24/7 in the CIO's office screen.

Example 2: The Application Owner needs a detailed online view of her application. She would like to be alerted whenever predefined thresholds (availability, performance, and SLA) are breached. As part of her daily Health validation checks, the view will serve as the starting point for triage and discussion with other IT support groups or functions.

With the output of the above you should start planning your Service Health configuration: models, views, customizations (such as HIs, KPIs, and business logic rules), function- or role-based pages, permissions and so on.

Chapter 3: Service Health High-Level Design

Service Health presents its status based on the data collected by APM's data collectors. All APM raw data is the Metrics type.

Metrics: Metrics are data packages sent APM according to a regular predetermined schedule. They contain monitoring information such as measurements and components. APM sources for metrics are BPM, RUM, SiteScope, and Diagnostics.

Service Health status calculation is based on raw data. An additional data layer comprised of Health Indicators (HIs), which will be explained shortly, is calculated on top of the raw data layer. HIs are used to normalize events and metrics on the same scale, and serve as the basic data building block for KPI status calculation.

The following are brief definitions of the vital concepts:

- **Health Indicator (HI).** Provides some indication on the health of a CI. An HI comes with a predefined scope (granularity of health covered by this indicator), and can assume multiple states.

  HI examples:

  - Number of error events (BusinessApplication CI)
  - Database Reads outstanding (SQL Database CI)
  - CPU Load (Computer CI)

- **Key Performance Indicator (KPI).** A KPI is a major characteristic of a CI which pertains to its overall health status. A KPI's scope is greater than an HI ; it is much less granular than an HI.

  KPI examples:

  - Application Availability
  - Software Performance
  - Backlog

- **CI status** is determined by the status of its subordinate KPIs (Worst Status Rule); the worst status out of the group of related KPIs will determine the CI's color.
- **Status calculation**. Status is calculated by a back-end engine running on APM's Data Processing Server, based on raw data and predefined calculation rules.

For a detailed overview of the terms above, refer to the APM documentation (**Application Administration > Service Health > Service Health Overview**).

Chapter 4: Customize Indicators and Rules

There are cases where Service Health out-of-the-box content is not aligned with a specific user's needs. Certain information (data dimensions) may be missing or redundant for certain consumers, and customization and adjustments may be required to fill these needs.

In this section we will review a few use cases, and demonstrate how to customize Service Health HIs, KPIs and rules.

This section includes:

- "Adding a Data Point/Dimension in Service Health" below
- "KPI Assignment" on page 14
- "Changing the KPI Representation " on page 16
- "Customizing KPI Presentation for a Specific View" on page 16
- "Changing a KPI's Calculation Rule" on page 16
- "Using a Subset of a CI's KPIs" on page 21
- "Displaying a KPI Only" on page 22
- "Example - Customize KPI Display" on page 23

Adding a Data Point/Dimension in Service Health

You can create an additional data point/dimension in Service Health, for instance, when a new HI for a business application - Application Availability Status - is needed to report application availability based on a SiteScope URL monitor. (This example assumes we already have the supporting topology configuration and measurement HI mapping).

1. Create a new HI:
   a. Access **Admin > Service Health > Repositories > Indicators**.
   b. In the left pane (CI Types), select the CI for which you would like to create the new HI.
   c. In the center pane (Indicators), click **New** and select **Health Indicator**.

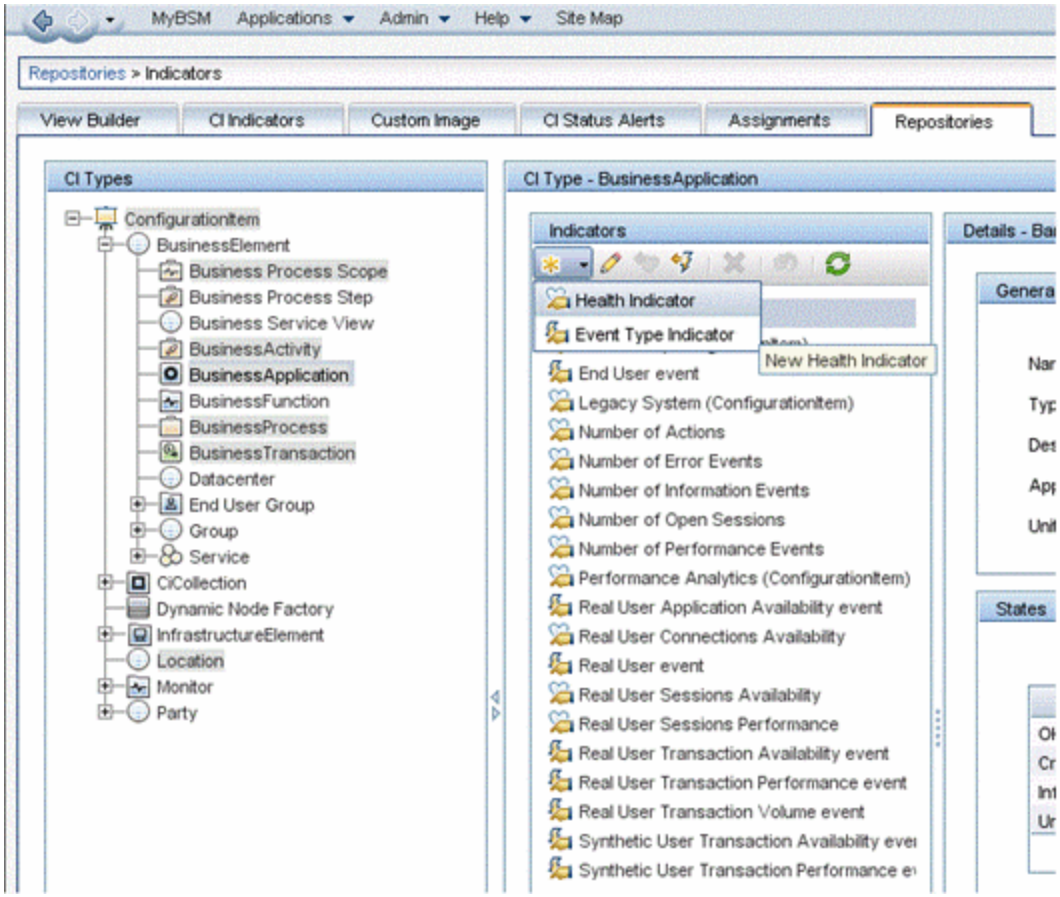

d. Complete the **Display Name** and **Name** fields and click **Save**. In the images below, we created a new HI named Application Availability Status.

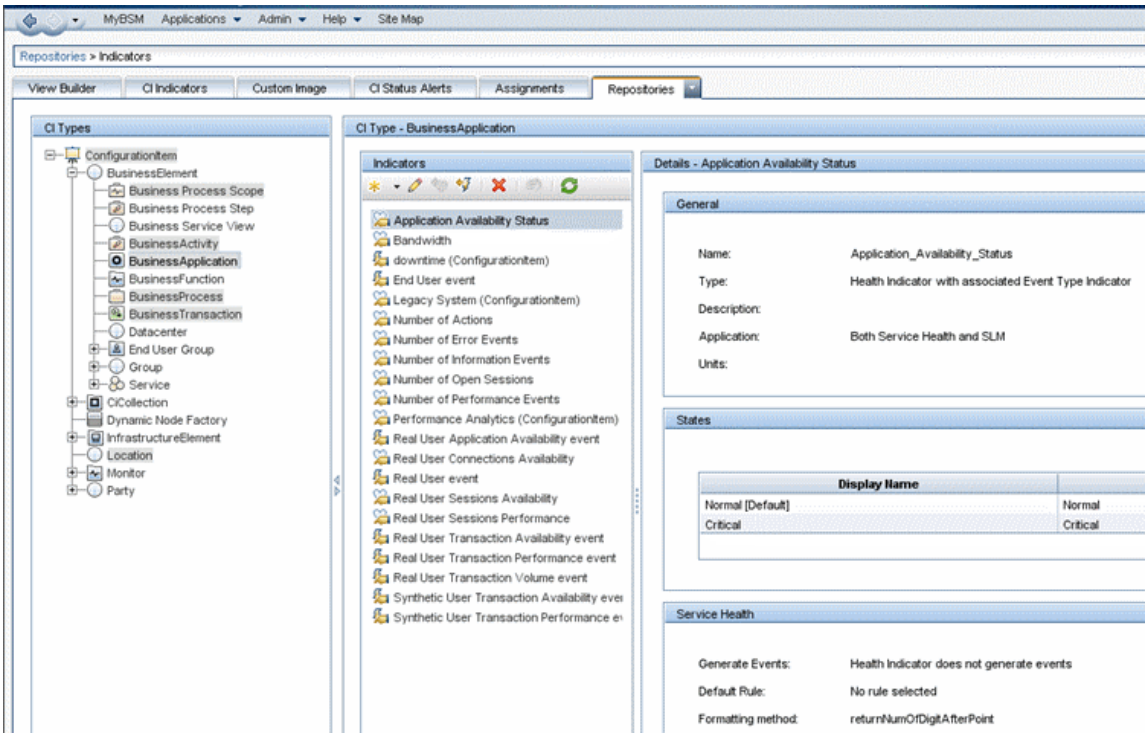

Make sure the new HI's states are aligned with your needs, and use meaningful names. This makes it easier for Service Health consumers such as NOC operators or application support to understand. Referring to the HI in the image above, some examples might be Available, Not Responding, and so on.

Note: In some cases, an HI may have multiple **states** mapped to one **status**. For instance, Memory Usage Level HI (Computer CI type) has nine states mapped to four statuses (see the following image). In such cases, make sure you map statuses to states properly.

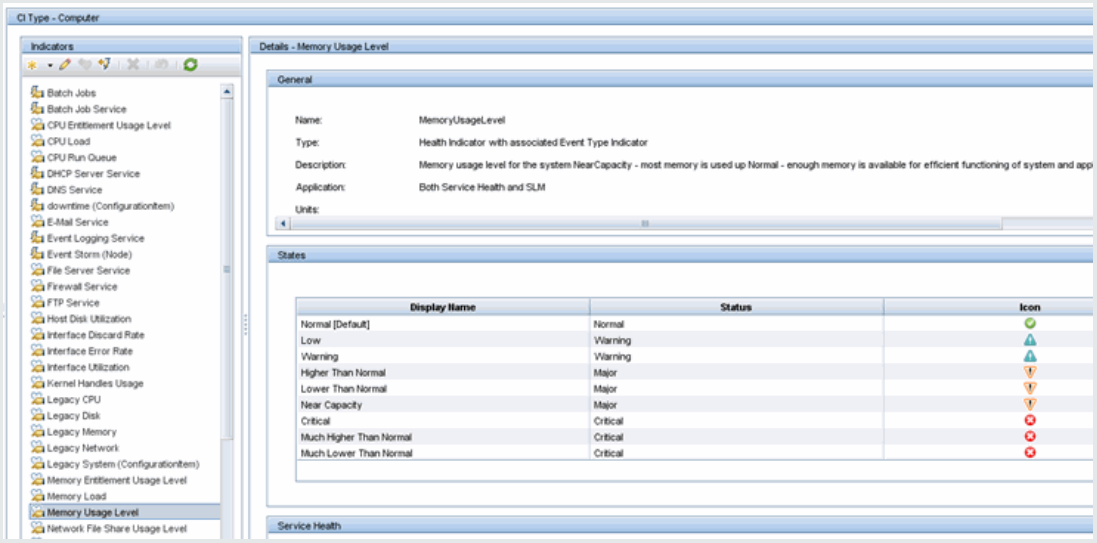

2. Create a new KPI:

   A common use case for creating a new KPI is when you want to present a finer granularity KPI for a high level CI (for example Business Application), so that it is visible in the 360 View or Top View.

   Note: Creating a new KPI is similar to the cloning process. We recommend you create a new KPI only in case the required logic is not similar to any existing KPI. The main difference between Clone and New is in the KPI's selected rules.

   a. Access **Admin > Service Health > Repositories > KPIs**.
   b. Click the KPI to be cloned.
   c. Click **Clone KPI:**

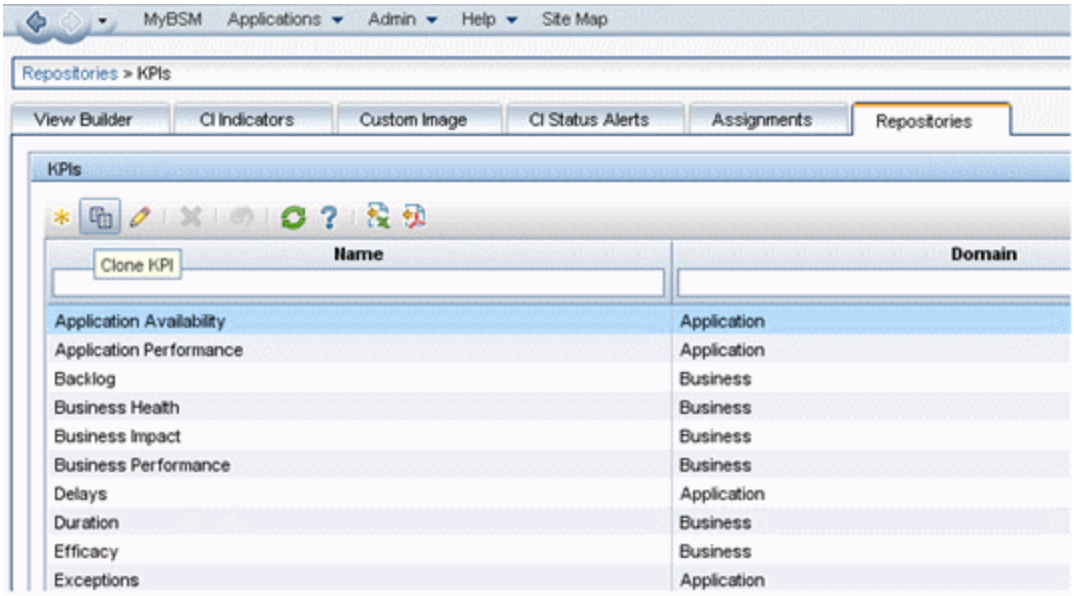

In the example above, we are cloning the Application Availability KPI. The result is as follows:

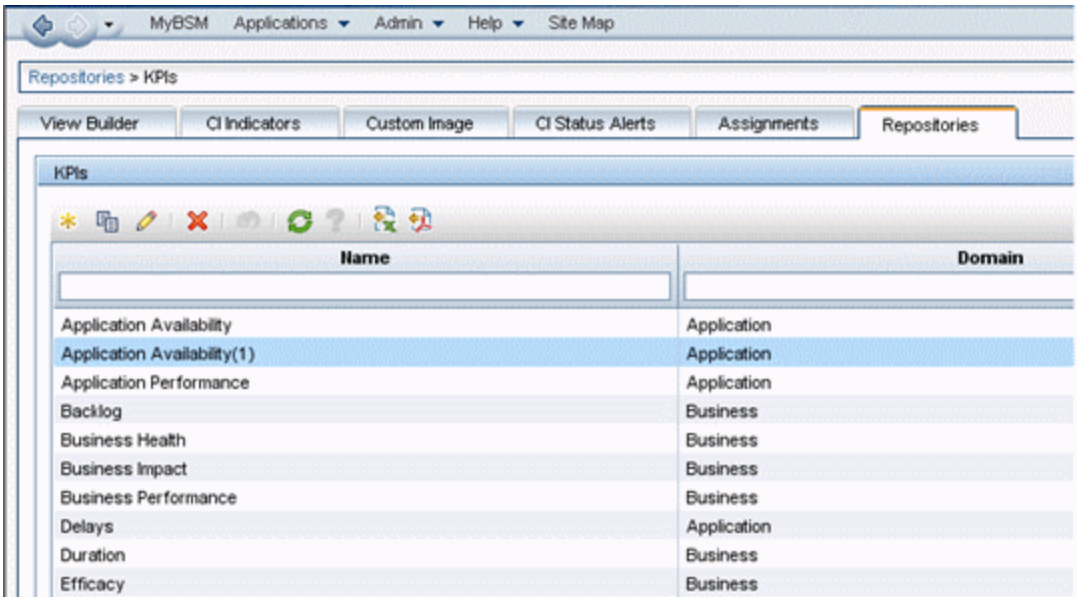

3. Click **Edit** to edit the newly created KPI.
4. In the **Main Settings** section, rename the KPI as relevant, and select **Domain** and **Default Group Rule**, as well as the **Applicable Rules**.

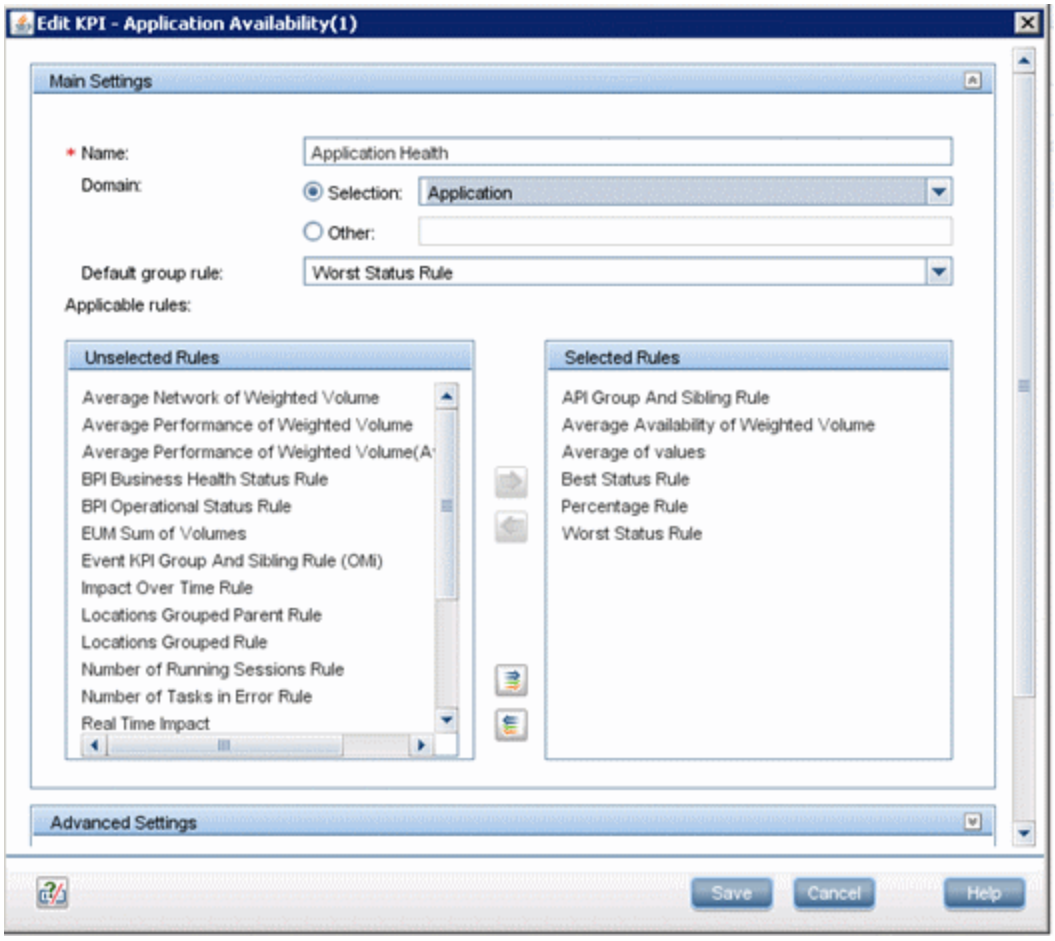

5. In the **Advanced Settings** section, you can change or determine the following:
   - KPI calculation order, compared to other KPIs for the same CI. (This is important if the new KPI will be calculated based on the result of other KPIs.)
   - The location where this KPI will be displayed, relative to other KPIs.
   - Condition for “Critical” state.
   - User type for which this KPI is applicable for.
6. In the **KPI Parameters** section, you can change the presentation type, as well as edit or add keys and colors. For example, you can change KPI presentation from an “Icon” to a “Value” or “Text” presentation type.

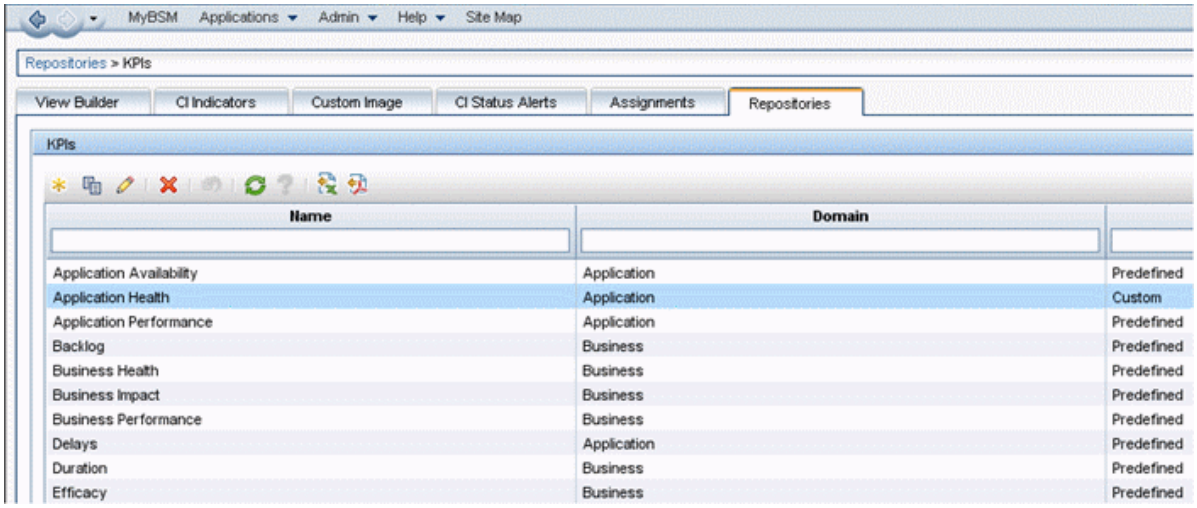

For further details regarding various fields and options, refer to the product documentation **(Application Administration > Service Health > KPI Repository > KPI Repository User Interface)**.

KPI Assignment

KPI assignments allow you to add KPIs to existing assignments or create new ones. It is recommended to create a new assignment only when the condition differs from the OOTB assignment condition.

An example for a case where a new KPI is needed is when Service Health consumers are interested in separating synthetic from real user data. By default, both data types are mapped to the same KPIs (Application Performance and Availability) for Business Transaction CIs.

1. Access **Admin > Service Health > Assignments > KPI Assignments**.
2. Choose the appropriate CI type in the left pane and click **New** or **Edit.**
3. In the **Condition** section, assign a fine-grain condition in addition to the “Monitored By” main condition.
4. In the **KPI Configurations** section, add the KPIs you want assigned by this rule.

   A KPI will be assigned only if the assignment's condition is met, and the CI has at least one of the HIs which are defined in the **Related HIs** section of the KPI configuration.
5. In the **Context Menus** section, add relevant context menus.

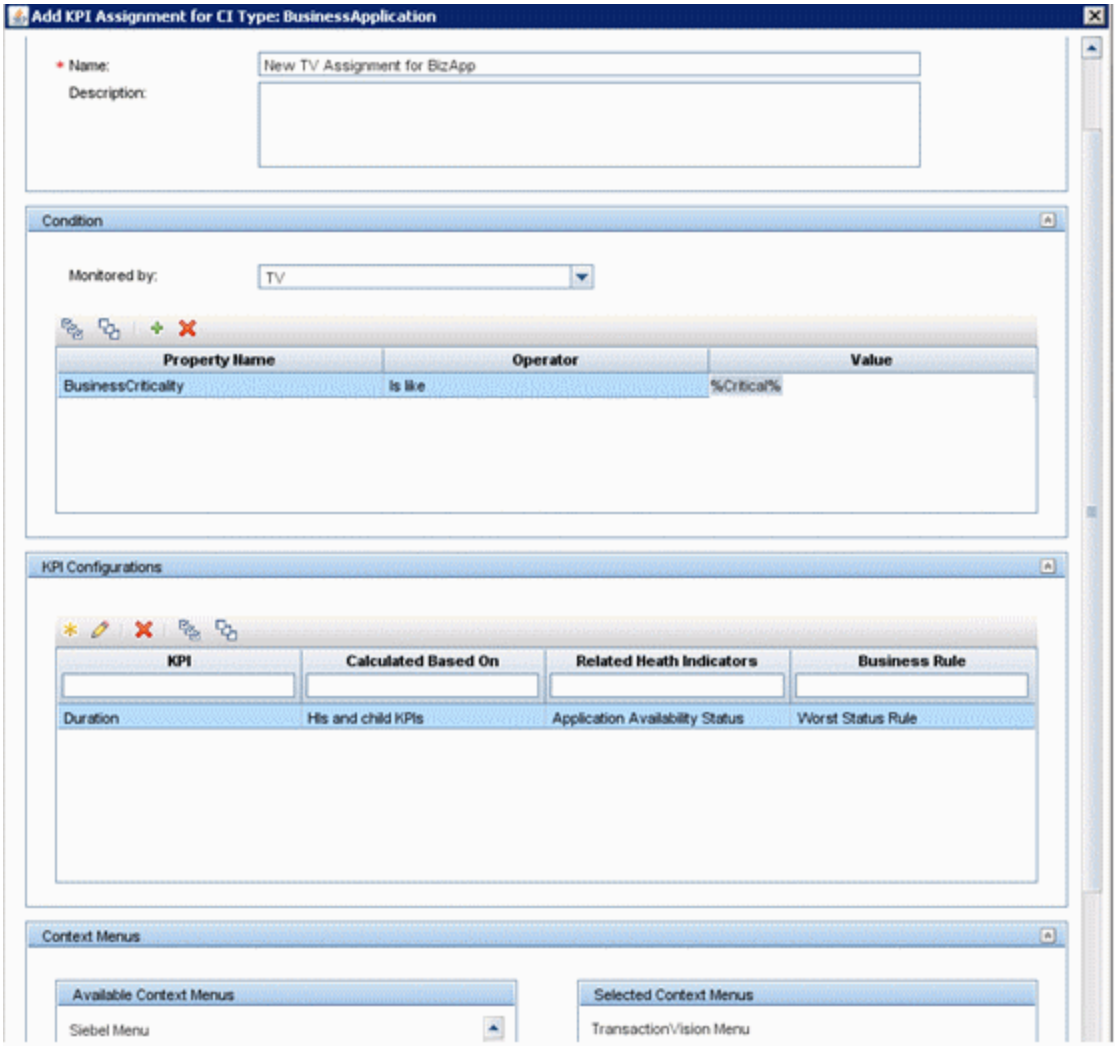

The result should look as follows:

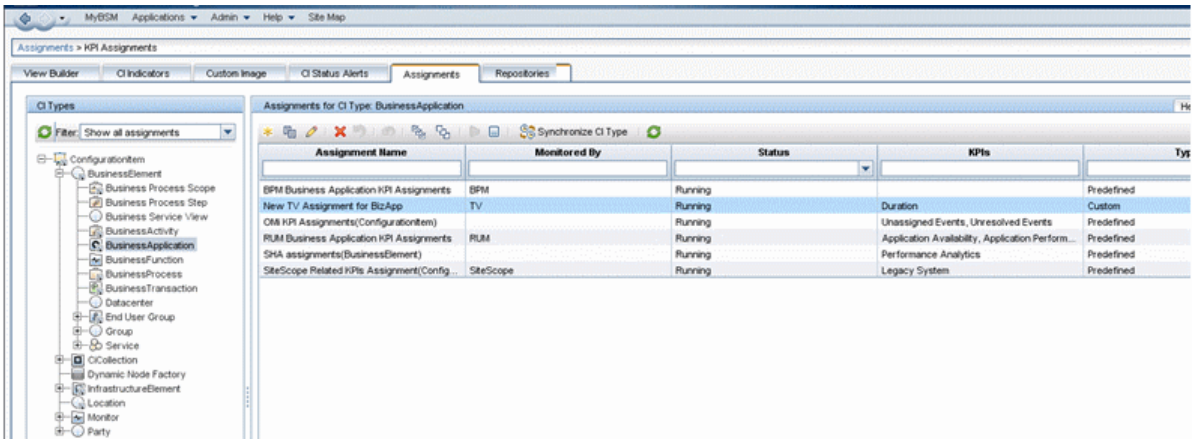

Changing the KPI Representation

You can change KPI representation if you need to display the numeric value of a KPI calculation result, instead of (or in addition to) a status icon. This is relevant if the selected KPI is calculated on a numeric basis.

1. Access **Admin > Service Health > Repositories > KPIs**.
2. Select the KPI and open it for editing.
3. Expand the **KPI Parameters** section.
4. In the **Presentation type** field, select **Value or status**, and save.

Customizing KPI Presentation for a Specific View

In some use cases, you may want to customize KPI presentation to a specific view. For example:

- A consumer is interested only in a subset of a CI's KPIs. For instance, an application owner is interested only in End User experience data, but the OOTB view contains Application, Software, Business, and System KPIs. Assume we want to remove the Legacy System KPI from the view's presentation. For details, see "Using a Subset of a CI's KPIs" on page 21.
- A variation of the above is when a consumer would like a certain KPI to be displayed, but not considered when calculating CI statuses. For instance, a business user would like the Legacy System KPI to be presented in a view, but not taken into consideration in CI status calculation. For details, see "Displaying a KPI Only" on page 22.
- A user would like to customize the Application Availability KPI. For details, see "Example - Customize KPI Display" on page 23.

Changing a KPI's Calculation Rule

In certain cases, you may want to change the KPI rule or calculation method. This may be needed if:

- OOTB thresholds do not suit your organization's standard.
- You would like to apply a different rule for a certain KPI.
- The rule logic you need is not part of the OOTB rules set. (This document does not address new rule implementation; for details refer to the APM Extensibility Guide.)

There are three levels where a KPI's rule can be modified, as described in the following sections:

Changing a Rule Across APM

The following will apply to every instance of this KPI; any CI assigned with the selected KPI will use the changed rule:

1. Access **Admin > Service Health > Repositories > Business Rules**.

1. Open a rule for editing.
2. In the **Edit Rule** dialog box, the **Advanced Rule Settings** section defines the set of CIs on which this rule can be applied. In order to extend the set of eligible CIs, select one or more CIs in the **CI Types** area, and move them to the **Selected CI Types (Includes Descendants** area.

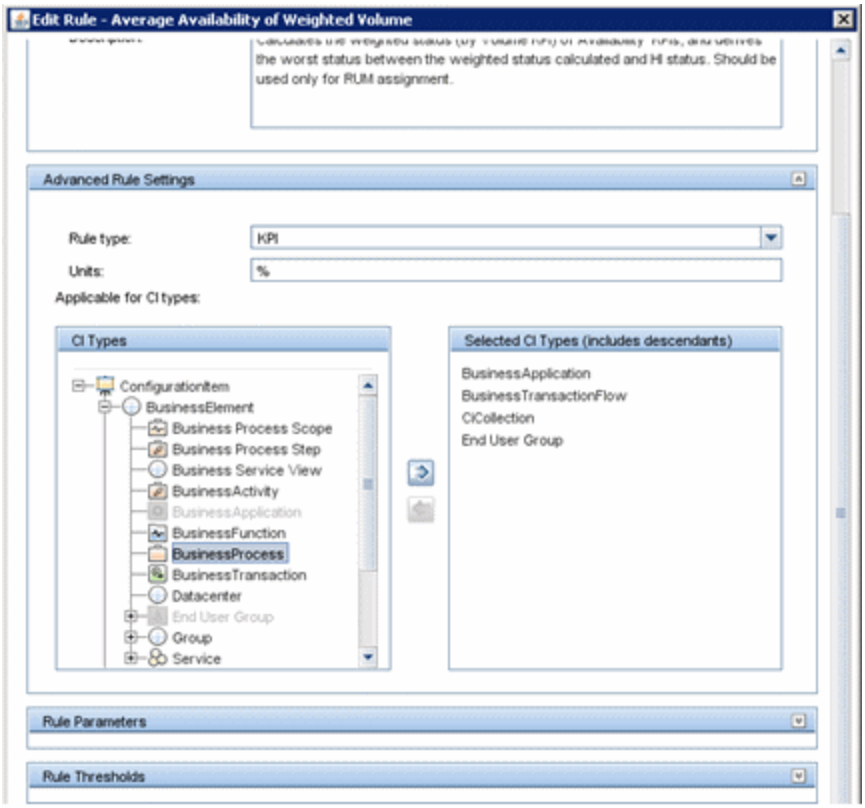

3. The **Rule Thresholds** section allows you to customize status condition values that are applied to the result of the rule calculation, to your organization's standard. In the image below, we changed the Major state threshold from 30 to 15.

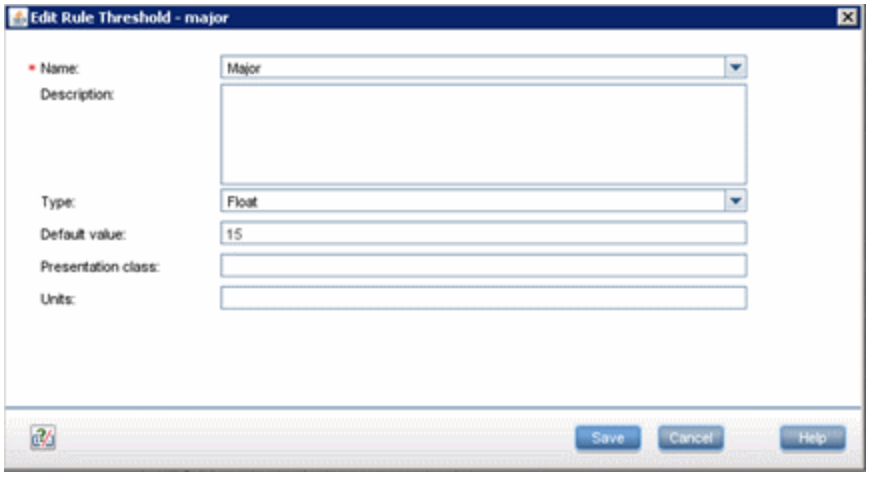

Changing a Rule for a CI Type

The following rule modification will apply to CI instances of the selected CI type (class):

1. Access **Admin > Service Health > Assignments > KPI Assignments**.
2. Select the CI you want to modify in the **CI Types** pane on the left.
3. The **KPI Assignments** tab (on the right) displays the list of assignments associated with this CI type.
4. Select the Assignment you want to modify and click **Edit**.
5. In the **KPI Assignment for CI Type** dialog box:
   - The **KPI Configurations** allows you to add additional KPIs as well as edit existing KPI rules.
   - The **Propagation Rule** tab in the **Assignment for CI Types** pane allows you to change or add propagation rules for this CI type.
6. In the **Propagation Rule for CI Type** edit box:
   - The **Condition** section allows you to change the parent CI type and the KPI type that will be propagated.
   - The **Task** section allows you to restrict propagations, change KPI calculation rules, or use a custom KPI for propagation.

Changing a Rule for a CI Instance

The most granular form of rule modification impacts one CI only; this type of change is rare. If the motivation for such a change is valid, we would expect it to be implemented on the CI type level.

Note that this kind of customization disconnects the CI instance from its OOTB assignment. The only way to link it with its original assignment is to delete the changed KPI, and run KPI Assignment.

1. Access **Admin > Service Health > CI Indicators.**
2. Select the view containing the CI whose KPI you want to change (in the upper section).
3. In the **CI Data** section (the lower part of the page), the **KPIs** tab displays all the KPIs assigned to the selected CI instance. You can either edit existing KPIs, or add new KPIs. The procedure is similar to

what is described in the previous sections.

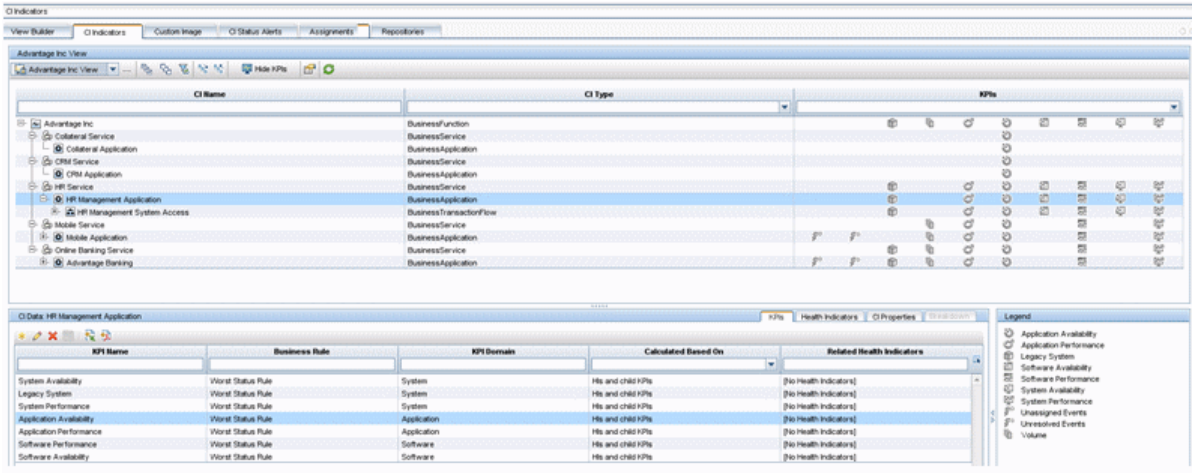

Bulk Operations

Various operations mentioned in this chapter can be performed in bulk. If you have a view containing several CIs and you want to update their KPIs in the same manner:

1. Access **Admin > Service Health > CI Indicators**.
2. Select the relevant view.
3. Select the (multiple) CIs you want to change. The **Common CI Data** (lower pane) presents the group of shared KPIs for all the selected CIs.

4. Perform the desired changes.

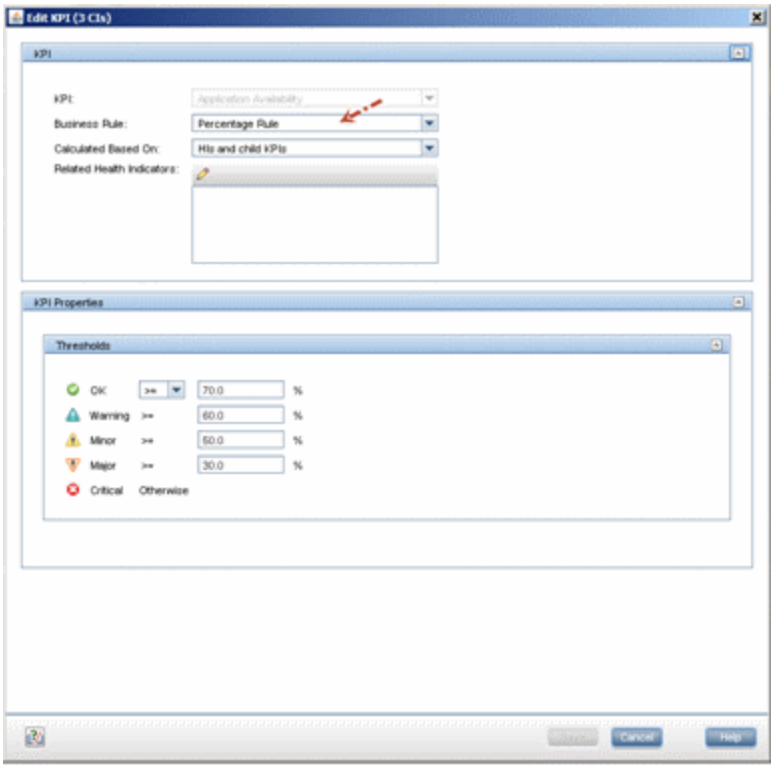

Note: If you have created new assignments or propagation rules, or if you have updated existing assignments or propagation rules, new KPIs should be created in order for those assignments and rules to take effect. You must therefore delete existing KPIs (which were created or assigned based on the previous rules). After that, you must refresh KPI assignments (in the **Admin > Service Health > Assignments > KPI Assignments** tab) thus triggering the creation of the new KPIs.

If the CIs which are impacted by the new assignments or propagation rules are part of one view, or grouped in a form that allows bulk operation, this KPI deletion can be performed in bulk (see the previous section).

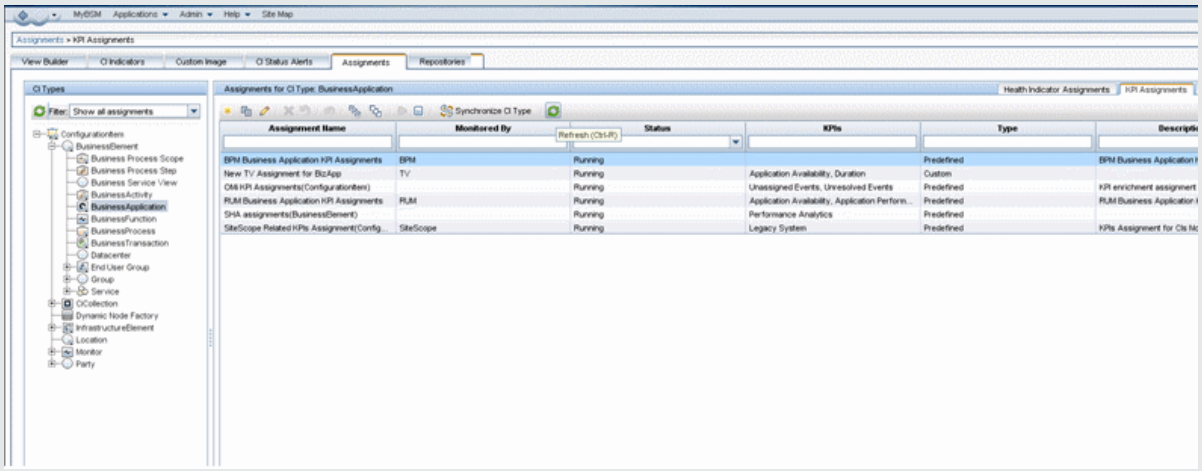

Using a Subset of a CI's KPIs

In some cases, a consumer may be interested only in a subset of a CI's KPIs. For example, an Application owner might be interested only in end user experience data, and we want to remove the **Legacy System** KPI from a view's presentation.

1. Access **Admin > Service Health > View Builder**.
2. Right-click the view you want to change in the **View Selector** (left pane), and select **Properties**.
3. Select the relevant KPI.

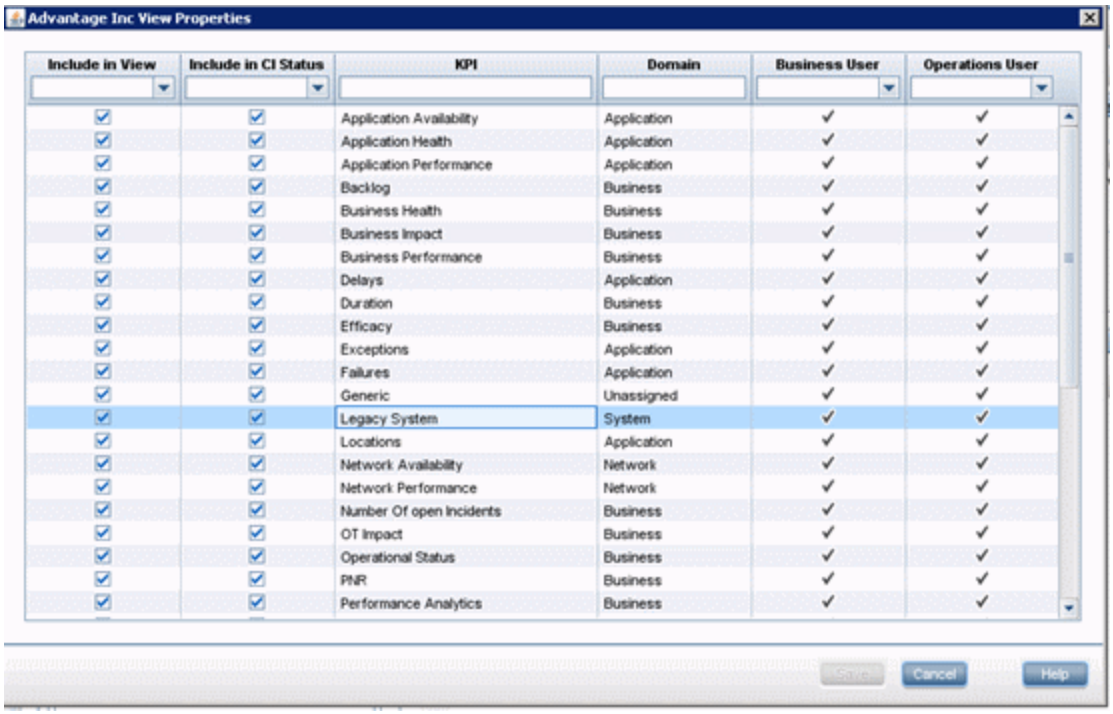

4. Clear the **Include in view** check box.

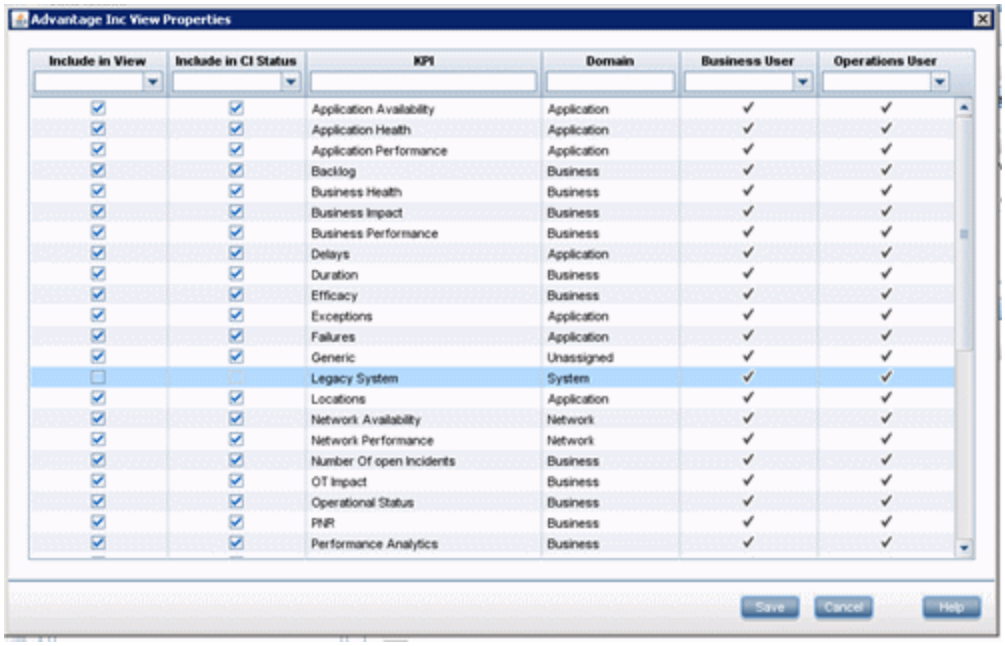

5. Click **Save**.

   This action automatically removes this KPI from CI status calculation as well.

Displaying a KPI Only

A variation of the above scenario is when a consumer would like a certain KPI to be displayed, but not considered when calculating the CI's status. For example suppose a business user would like the **Legacy System** KPI to be presented in a view, but not taken into consideration in the CI's status calculation.

1. Access **Admin > Service Health > View Builder**.
2. Right-click the view you want to change in the **View Selector** (left pane), and select **Properties**.
3. Select the relevant KPI.
4. Clear the **Include in CI Status** check box.
5. Click **Save**.

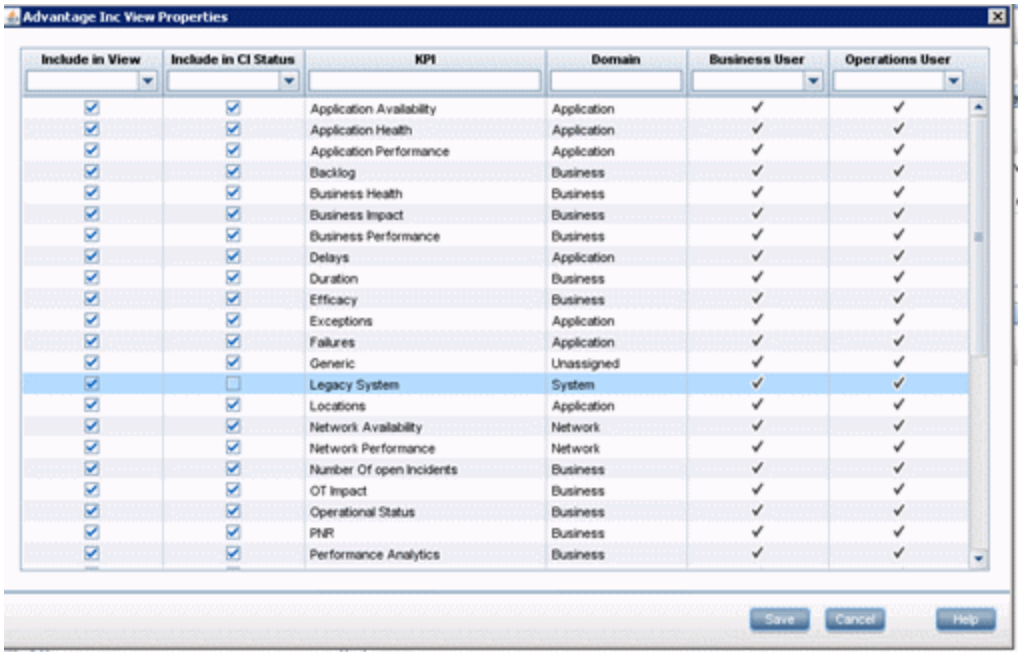

Example - Customize KPI Display

Suppose we want to customize the **Application Availability** KPI display.

Here's how the view appears in the Service Health 360 View; note that the Advantage Inc CI status is **Critical**, which is derived from the **Application Availability** KPI's status:

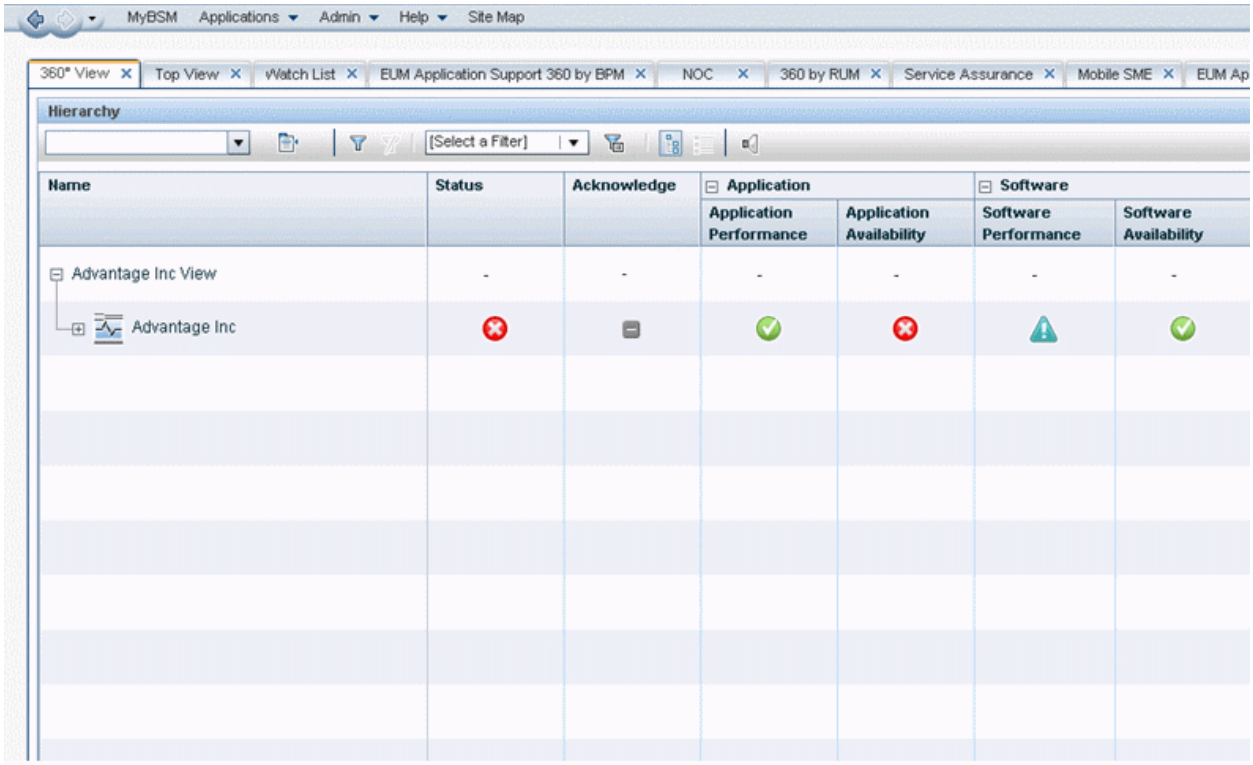

We clear the **Include in CI Status** check box:

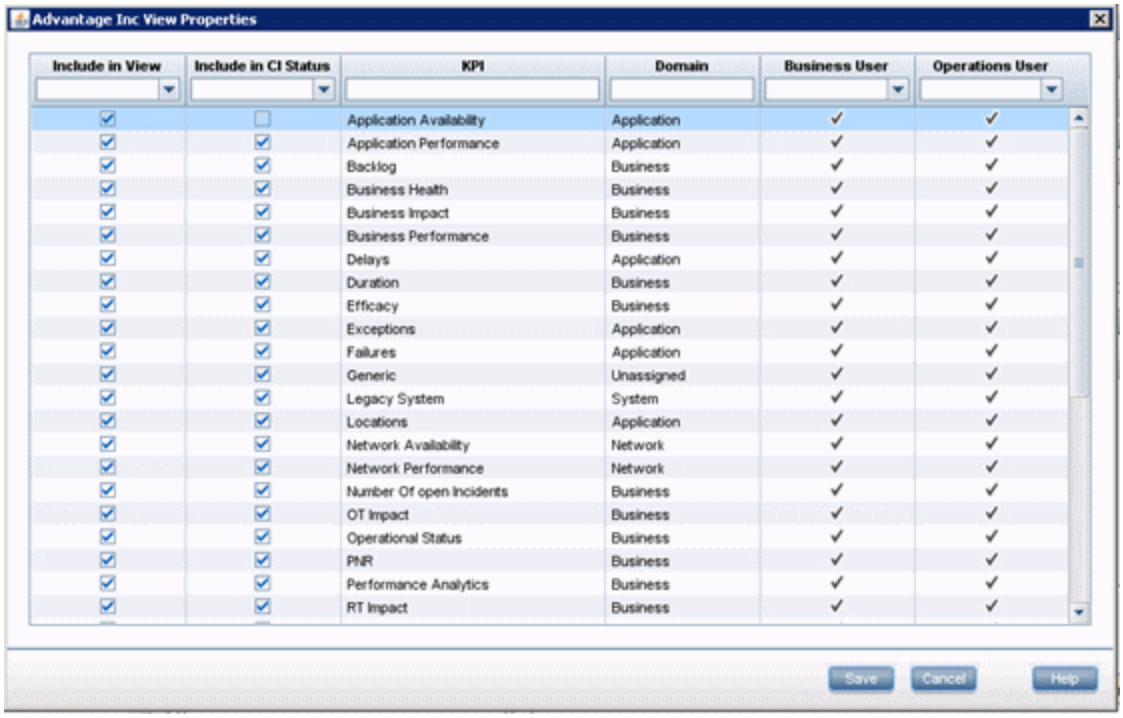

Now the view looks like this; the Advantage Inc CI's status changes to **Warning** despite the fact that **Application Availability** KPI is still **Critical**:

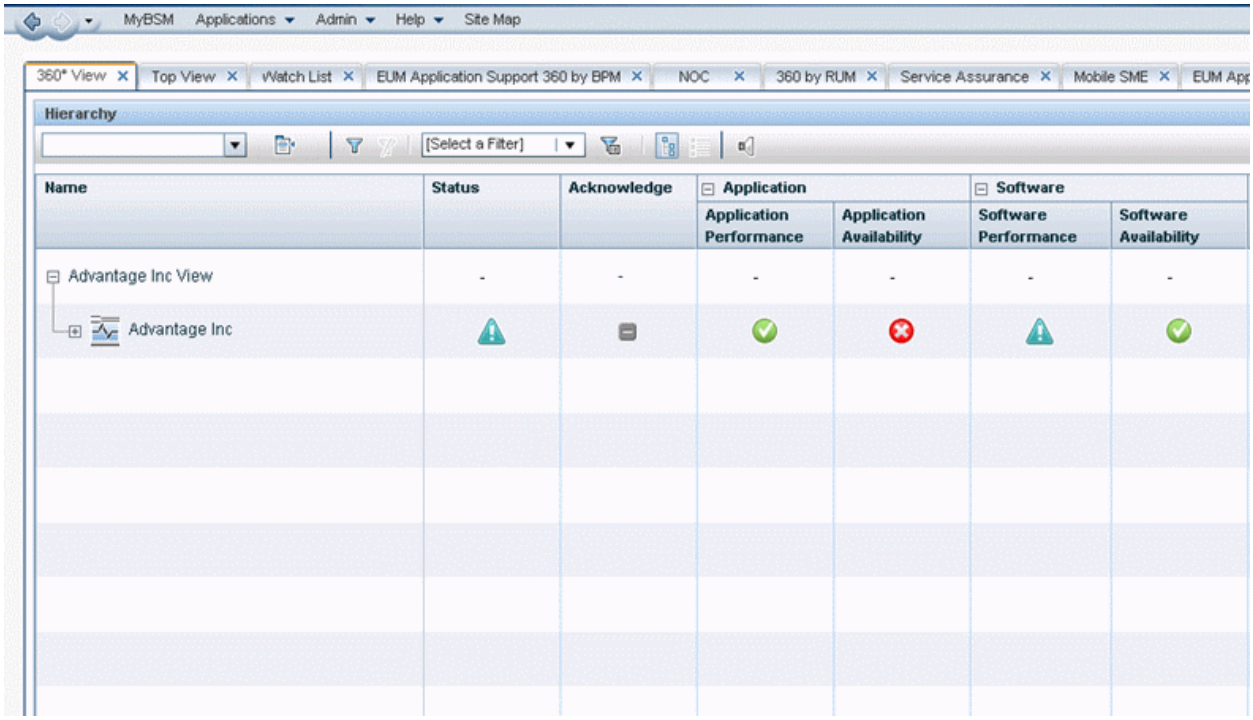

Now we will clear the **Include in View** check box as well:

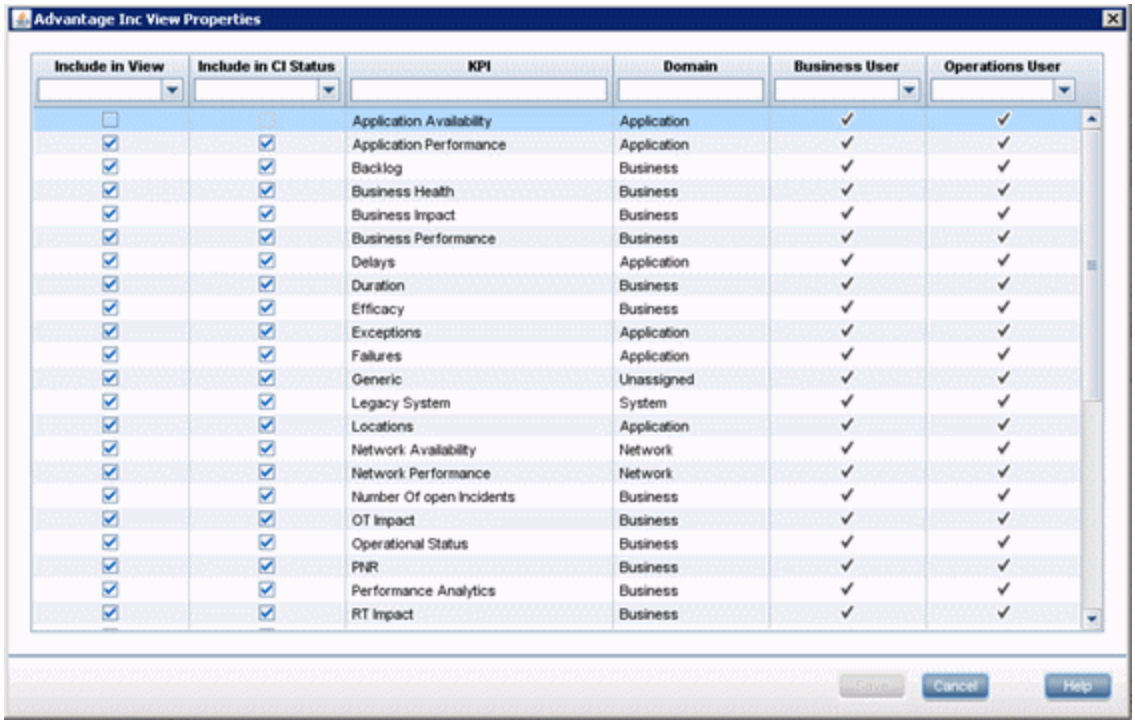

The **Application Availability** KPI is removed from the view entirely:

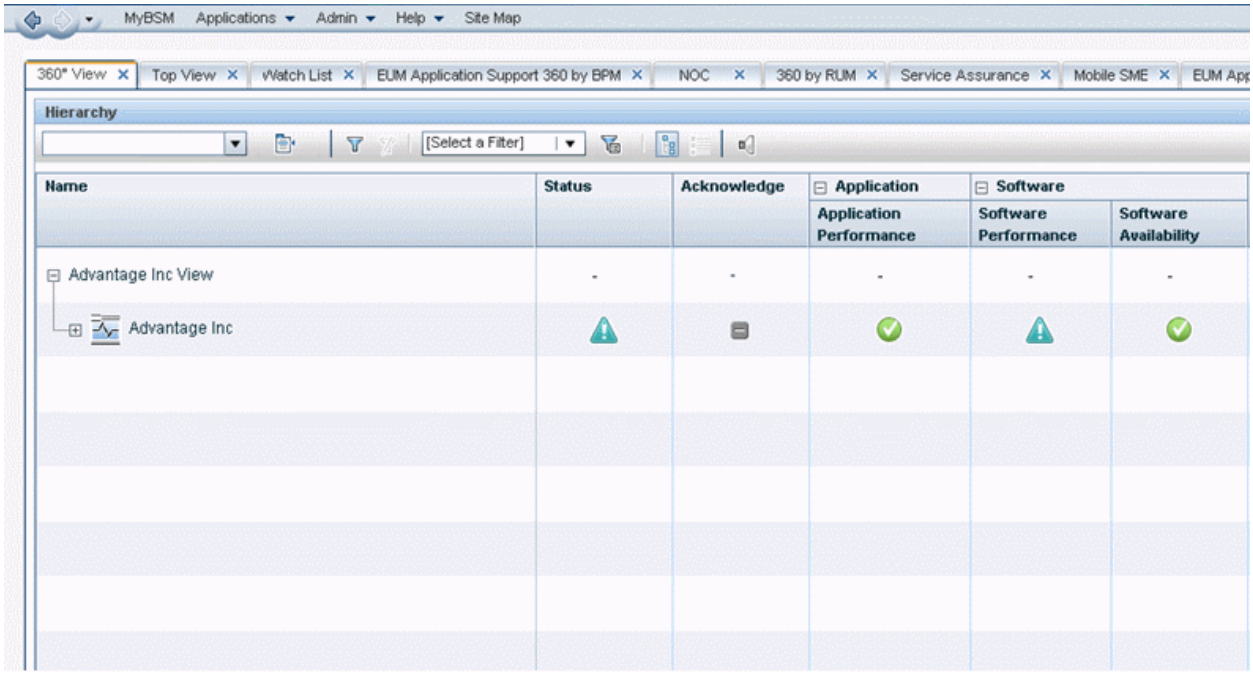

Chapter 5: Views

Views are not the main focus of this best practice document; however in order to maximize Service Health's value, it is important to understand the concept behind views, and the motivation behind the different types of APM views.

- A View is a mechanism for consuming sub-sets of CIs from APM's configuration management database (RTSM). It enables the user to extract a selected group of CIs that meet specific interest or needs. For example, a view might include all infrastructure components that support a specific business application or service (application support). Another example could be all DB servers in the organization and their storage devices (DBA infra team). Service Health, as well as other APM components, consumes views as the basic building block.
- Note that not all RTSM views are available to Service Health, but only those which belong to the Service Health bundle. Bundles are a method of grouping TQL queries or views by category. In order for a view to be accessible to a certain APM application or module it should belong to its bundle. For further details, refer to the APM documentation (**Modeling > Modeling > Modeling Studio > Building a Business View > Making Views Available in APM**).

Tip: If you want to reduce the number of views in Service Health, for example if there are too many OOTB discovery views, then you can remove their Service Health bundle.

Working With Views

1. Access **Admin > RTSM Administration > Modeling > Modeling Studio**.
2. Select the view in the **Resources** tab.

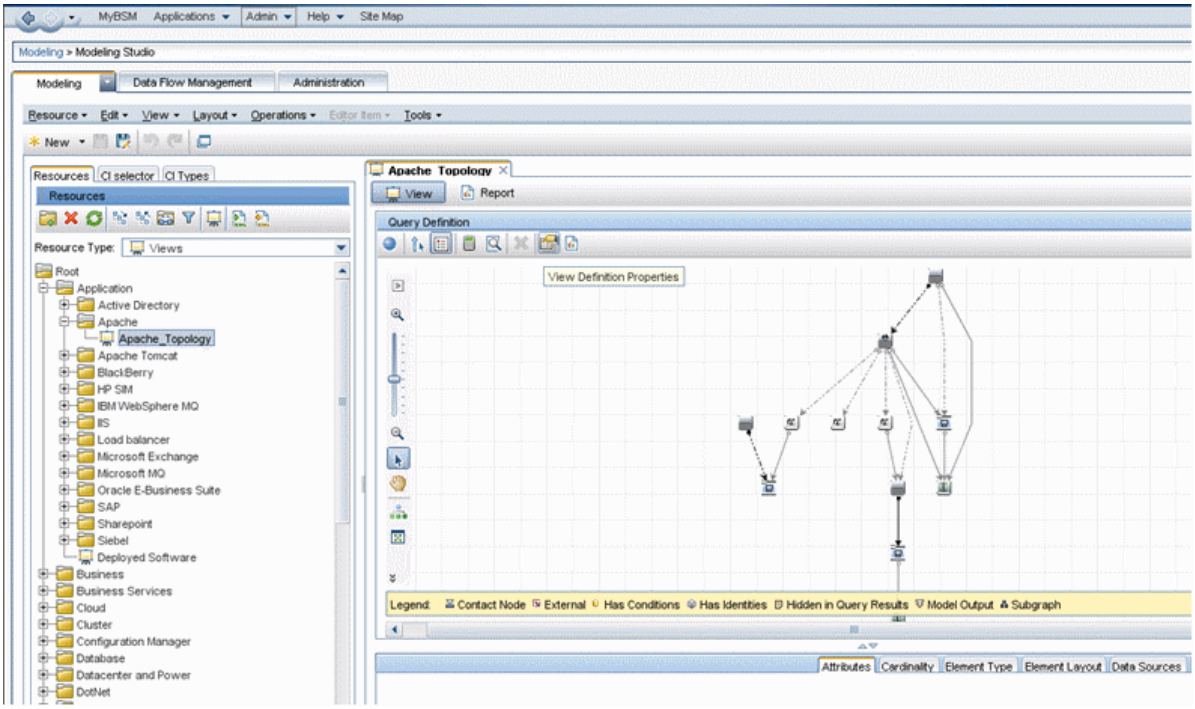

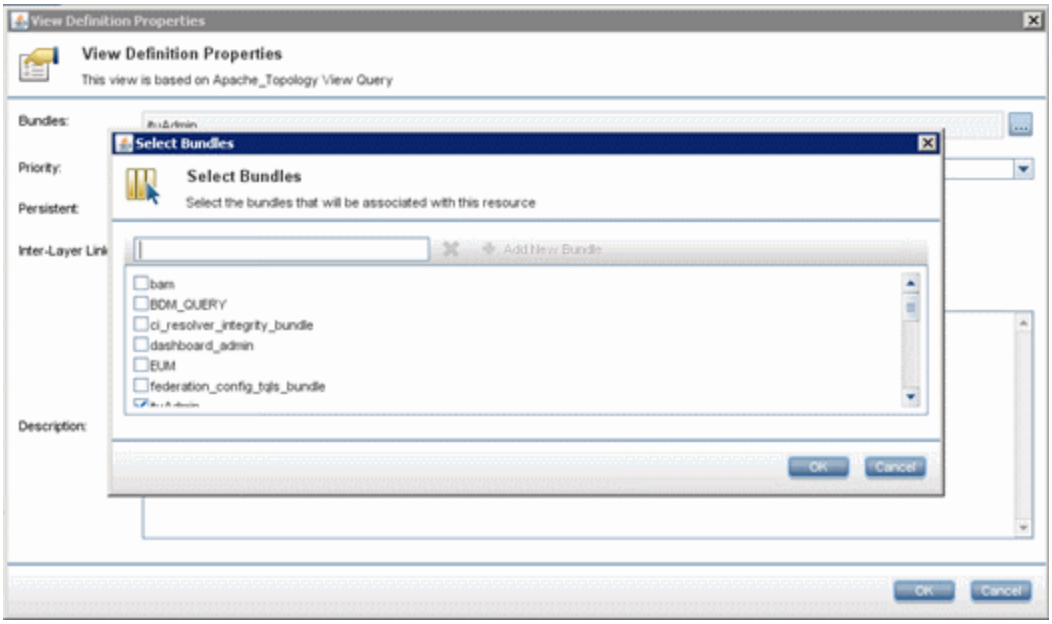

Which view type is suitable for your needs?

- Choose a **pattern view** when searching for a specific signature in RTSM (for example, all nodes with 2 IP addresses). The advantage to using a pattern view is that it is dynamic; a view's result is updated behind the scenes without any user intervention. Note that building pattern views requires an understanding of TQL and RTSM topology.
- **Template-based views** allow the creation of multiple views based on the same template. If you need to create multiple views with the same characteristics, consider whether those can be incorporated into one

template. Parameterization lets you define which attributes of the template can be assigned with changing value or content, thus saving repetitive manual work. Naturally, any change to the base template will impact all template-based views.

- **Perspective-based views** allow superimposing certain vantage points (or perspectives) on a collection of CIs. You start with a "root" CI and apply the desired perspectives (whether an OOTB perspective or one you've created) on that CI. Perspective-based views are easy and quick to generate, and are recommended when the scope (root) of the view is known in advance, and static.

  Note that APM has a limit of 1000 active TQLs, and each perspective-based view generates 3 or more active TQLs (depending on the scope of applied perspectives). Hence, capacity should be taken into consideration when using perspective-based views.

What are some rules of thumb for creating views?

- Limit the number of CIs to no more than 1000. This will allow your view to remain manageable and usable; it is also effective from a capacity perspective, keeping in-memory footprint as small as possible.
- Use View Hierarchy (previously called Folding) to minimize the number of CIs in each layer. This will help you consume the view in Service Health components such as Top View or 360.
- Use "**vertical**" views where appropriate; this refers to the approach of dividing your model and business into "silos" (for example a business application and all its related CIs in one view). A vertical view helps you focus on those elements relevant to your specific needs. Applying this approach allows you to use the Service Health **Watch List** very effectively. You can incorporate multiple vertical views into a high level overview of related business applications, business services, and so on. You can then select a specific watch list element to change context to the vertical view, providing a higher degree of freedom and flexibility when compared to a single model.
- Use "**horizontal**" views where appropriate; horizontal views serve best for users who are interested in a wide view (focused on a specific layer) across the organization. For example, DBAs (all DB servers and related storage devices in the organization), Windows administrators (all Windows servers in the organization), and so on.
- **High level views**. When creating high level views, such as overall business services view for CIO, or critical cross organization infrastructure services view, it is recommended to apply the "less is more" approach. Include the minimal set of items or CIs that will provide the required information or status, gaining simplicity and focus as well as an effective performance footprint.

What is a local impact view (LIV)?

In a LIV, only CIs that are displayed in the view participate in the view's status calculations (KPIs, HIs and the underlying rules); in a LIV what you see is what you get.

Examples:

- Green node: part of LIV (returned by view TQL), participating in view's status calculation
- Blue node: not part of the view and not participating in view's status calculation
- Pink node: not part of LIV (not returned in TQL result set), participating in view's status calculation

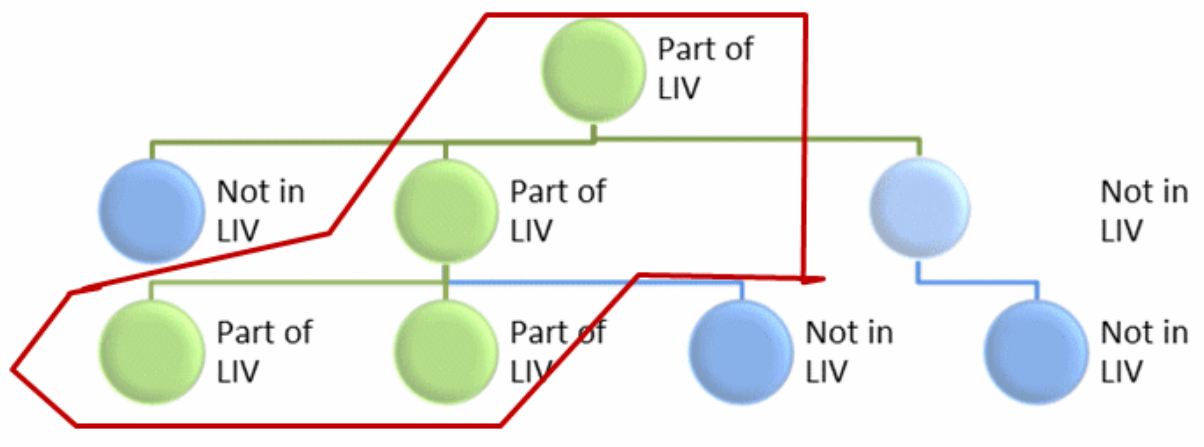

The only exception to the above is where a CI's status calculation relies on KPIs on its child CIs‘ which are not members of the LIV. In such a case, in order to maintain data integrity, the child CIs‘ status will be taken into account:

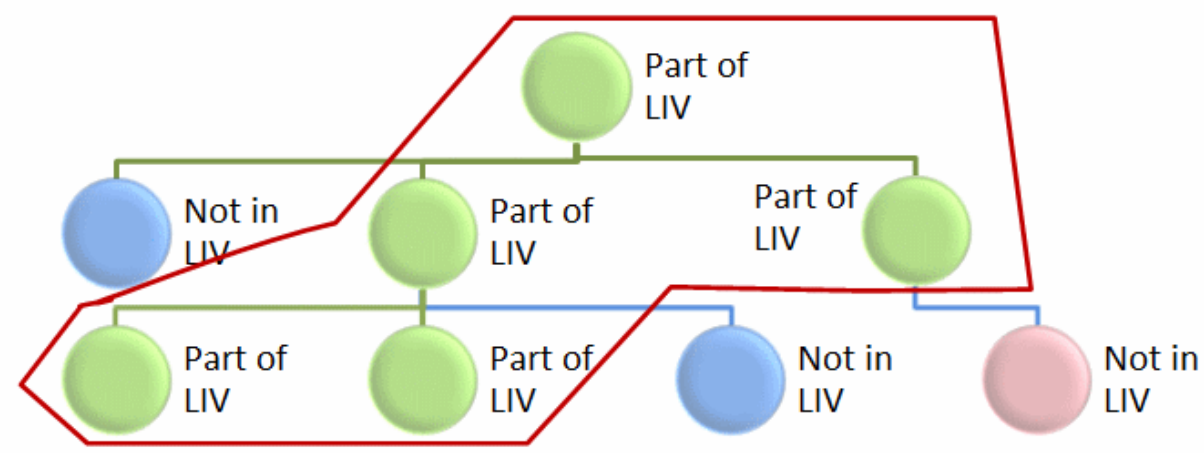

When should you use a Local Impact View?

- You wish to observe the status of only a subset of the Impact model (for example, if you are not interested in Login transaction data or impact as part of Business Application status).
- If you need customized or ad-hoc status calculation (such as a different calculation rule, thresholds, or rule parameters) in addition to the OOTB calculation.
- Breakdown. From BSM v9.0, EUM breakdown was removed from the model (Breakdown CI is no longer part of the class model). LIVs provide you with various breakdown capabilities, such as:
  - More breakdown options than previous versions (see image below).
  - When breakdown is configured for high-level CIs, new child CIs are assigned with breakdown as well.
  - Breakdown can be applied over metric-based HIs only (not events).
  - Breakdown calculation cannot be applied retroactively using historical data prior to when the breakdown was enabled; it is applied on new samples only.

Note: LIV bears a relatively high memory footprint, and should be used as needed (it should be the exception rather than the rule).

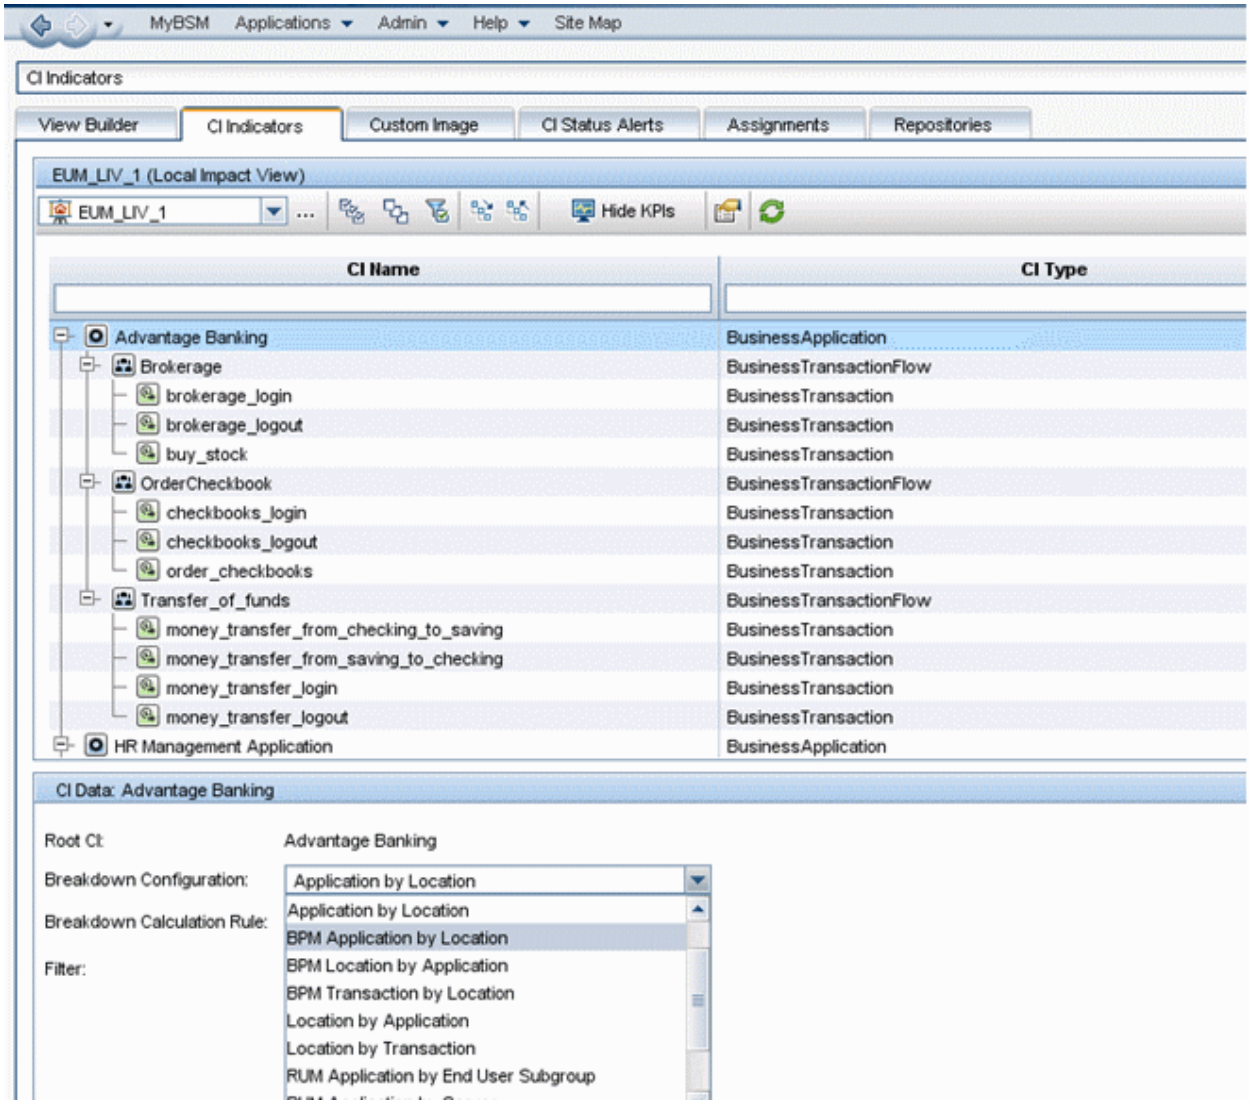

View Permissions

Use user groups (**Admin > Platform > Users and Permissions > User Management > Permissions** tab) to assign view permissions. This gives you more control over information access. For example, NOC operators will have access to high level operational views and the ability to drill down for initial triage purposes, while application support will have access to views relevant for their application and the supporting infrastructure.

Chapter 6: MyBSM Planning

Note: This section is a supplement to the detailed guidelines in the MyBSM documentation (see **User Guide > MyBSM > MyBSM Overview**).

Building a MyBSM page requires research and planning, especially in the following aspects:

- Layout
  - How many components you need in your MyBSM page?
  - How much space (real estate) does each component need?
  - How should you arrange the different components in the page in order to be consumed easily? Horizontally, vertically, and/or in tabs?
- Content
  - Determine which components are relevant for the purpose of that particular page. Keep the overall number of components to the minimum that is necessary for the page to remain usable and effective. Note that every component is generating an action in APM, thus, loading APM's Gateway servers.
- Component wiring
  - Wiring enables the user to design a flow or a role-specific "work desk". Configure the components' order, relation and context to reflect the flow of actions and information needed by the function using this page.
  - Use components from different APM modules to encompass as broad a range as possible of relevant use cases and processes, as implemented in your organization.
- Permissions. Configure access and action permissions to reflect your organization's needs and policies.

Chapter 7: Use Cases / Troubleshooting

The purpose of this section is to discuss various use cases or issues we find common based on R&D's field engagements. Our goal if for this section to grow as more popular use cases are gathered by the lab and field.

The corrective technical steps mentioned below are given in more detail throughout this document.

Service Health Not Reflecting Your True Business Health

The underlying assumption for effective use of Service Health is that only true cases of service disruption, or risk for potential disruption, will be reflected in Service Health's components.

If there is a discrepancy between the actual health of your business and the way it is reflected in Service Health, try the following to identify the cause for discrepancy:

- Validate that relevant indicators are contributing to status calculation. For instance, there could be redundant KPIs which are not needed in order to reflect application or service status. For example, if an application owner is interested in his application's availability only, the performance data presented OOTB is not needed, and may add redundant clutter.

  The same applies for Health Indicators (HIs).
- Verify that the Service Health rules being used reflect your business logic and need. Service Health OOTB rules aim at a common use case; however you may need to apply a different logic. For example, in a High Availability architecture application, the OOTB Availability KPI is calculated using the Worst Status Rule. However, from the perspective of the application owner and your organization, it may be perfectly acceptable if one of the application interfaces is down, as long as the application is available via other interfaces. In such a case, you would replace the OOTB rule for this KPI with the Best Status Rule.
- Confirm that only relevant CIs are participating in the view. Consider using pattern-based or perspective-based views in order to filter out irrelevant CIs and unneeded content.

Exporting Service Health Data

If you need to export Service Health's data (usually for the purpose of integration with external tools), Service Health exposes an API to query its database for view information and data. For technical details and examples, refer to the product documentation (**Extensibility Guide > Service Health > Service Health External APIs > Service Health Database Query API**) . The API will return requested data in XML format which can easily be formatted into any desired format.

The following example will return data on the End User Monitors view, including CIs' KPIs attributes including status, business rule, and last status change:

```
http://GWServer/topaz/bam/BAMOpenApi?customerId=1&userName=<usr>&password=<pwd>&viewName
=End User
Monitors&command=getNodes&addRecallAtts=y&depth=20&forcedPropogationDepth=20&layout
=hierarchical&showTooltip=true
```


- <dimension name="Application Performance" status="15" id="6">
null
- <tooltip>
- <body>
- <entry key="Status">
<![CDATA[Warning]]>
</entry>
- <entry key="Business Rule">
<![CDATA[Worst Status Rule]]>
</entry>
- <entry key="Last status changed">
<![CDATA[8/27/2013 10:54 AM]]>
</entry>
</body>
- <trend>
- <entry key="Trend">
<![CDATA[Stable]]>
</entry>
</trend>
- <history>
- <entry key="History type">
<![CDATA[Worst]]>
</entry>
- <entry key="Historical worst">
<![CDATA[Warning]]>
</entry>
</history>
</tooltip>
</dimension>

Send Documentation Feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Service Health Administration - Best Practices (Application Performance Management 9.30)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to SW-doc@hpe.com..

We appreciate your feedback!