Monitoring Business Services

Skylar One version 12.5.22

Table of Contents

Introduction to Business Services **5**

What is a Business Service? 6

The Business Services Page 7

Deleting, Enabling, or Disabling a Business Service 10

Favorite a Business Service 10

Business Service Dashboards 10

Service Insights Page 11

Example: Retail Banking 11

Creating Services and Service Policies **12**

Administrative Processes for Business Services 13

Understanding Health, Availability, and Risk 13

Creating Services 14

Creating Business, IT, or Device Services 16

Creating a Custom Service Model 18

Advanced Search Criteria: Practical Use Cases 26

Using Service Policies 27

Creating a Service Policy 27

Use Cases for Service Policy Creation 31

Selecting a Service Policy 32

Managing Service Policies 33

Deleting a Service Policy 34

Creating a Service Template 34

Creating a Service From a Template 37

Exporting a Service Template 38

Default Service Policy Settings 40

Device Service Default Policy 40

IT Service Default Policy 40

Business Service Default Policy 40

Custom Service Model Default Policy 40

Assigning an Icon to a Service 41

Exporting Service Data with the ScienceLogic API 42

Using the Service Investigator **44**
Viewing the Service Investigator 45
The Sections on the Service Investigator Page 45
The Overview Panel 45
Sunburst View 46
Map View 47
The Timeline Panel 48
Skylar AI Swim Lanes 49
Confirmed Alerts 49
Suggested Alerts 50
Changes Swim Lanes 50
Maintenance 50
ServiceNow 51
Skylar Compliance 51
Status Swim Lanes 51
Health 52
Availability 52
Risk 52
Events Swim Lanes 53
The Events Panel 53
Events Tab 53
Changes Tab 55
Service Analysis Tab 57
Skylar Advisor Tab 57
Using the Service Analysis Feature 59
Enabling Service Analysis 59
Configuring and Enabling the Changes Tab 60
Editing the Run Book Action 61
Syncing Skylar One Devices with Skylar Compliance 62
Permanently Enabling the Changes Tab 62
Temporarily Enabling the Changes Tab 62
Resolving Service Issues with Behavioral Correlation **64**

Understanding Behavioral Correlation in Business Services 64
Analyzing a Service's Behavioral Correlation 65
Troubleshooting Business Services 66
Business Services Have Empty Values 66
All Business Services Have Empty Values 66
Some Business Services Have Empty Values 67
Services Missing Up-to-Date Values 68
Some Services Fail to Generate Health, Availability, or Risk Values 69
Access Log Messages For Warnings or Errors 69
Check if Services Have Over 100 Constituents 73
All Services Fail to Generate Health, Availability, and Risk Values 74
Device Services Fail to Load After an Upgrade 75
502, 503, or 504 Errors: Health, Availability, and Risk Values are All the Same or are Inaccurate 75
Advanced Troubleshooting 77
Customization for Environments with More Than 2,500 Services 78
Update Settings and Increase Default Values 78
Modify NGINX Rate Limit 78

Chapter 1

Introduction to Business Services

Overview

This manual describes how to use Skylar One to create and manage business services for your company. Business services let you gauge the availability, health, and risk of your services and the devices that provide those services.

NOTE: Business services are available as part of a Skylar One Standard solution. To upgrade, contact ScienceLogic Customer Support. For more information, see https://sciencelogic.com/pricing.

NOTE: Business services and IT services created in the classic Skylar One user interface are *not* included in nor related in any way to the current iterations of business services, IT services, and device services that are described in this manual. For more information about the classic versions, see the ***Service Provider Utilities (formerly Business Services)*** and ***IT Services (Classic)*** manuals.

This chapter covers the following topics:

What is a Business Service? 6

The Business Services Page 7

Deleting, Enabling, or Disabling a Business Service 10

Favorite a Business Service 10

Business Service Dashboards *10*

Service Insights Page *11*

Example: Retail Banking *11*

What is a Business Service?

A ***business service*** includes one or more technical services that provide value to internal or external customers. Some examples of business services include verifying Internet access or website hosting, online banking, remote backups, and remote storage. Usually a business service includes an associated Service Level Agreement (SLA) that specifies the terms of the service.

There are two methods by which you can create business services in Skylar One:

Option 1:

You can create the following types of services on the **Business Services** page, in the following order:

1. ***Device service***. Monitors a set of related IT infrastructure components (devices) that deliver a discrete function, such as a DNS or collector group, or all devices in a specific region.
2. ***IT service***. Monitors a service that IT provides to your organization. An IT service provides a way to define how a set of related device services work together to power a given IT service, such as a DNS plus collector group plus a database.
3. ***Business service***. Monitors a service your organization provides to your customers. A business service consists of one or more IT services.

The following figure shows an example of how your business services may be organized.

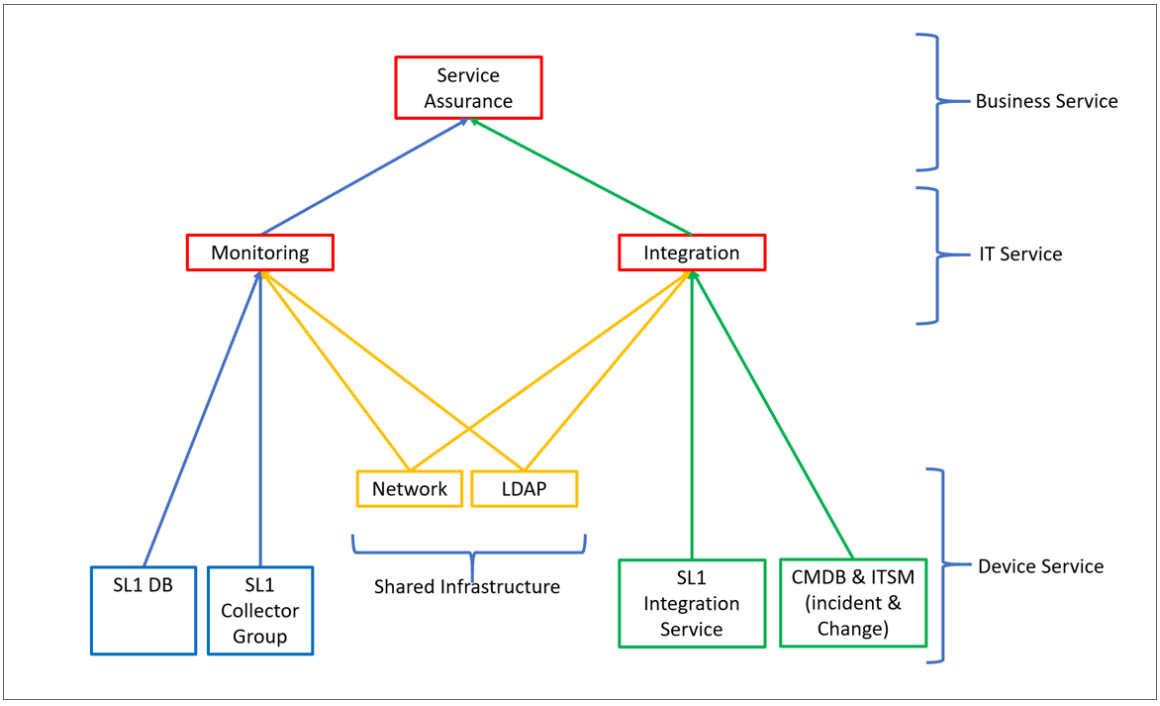

Option 2:

Alternatively, if you require more flexibility in modeling your business service, you can create a custom ***Service Model*** based upon how your organization defines its structure.

This option, which is also called an "N-tier" service model, uses a wizard to walk you through the process of building custom service models with multiple nested or connected service levels, each of which you can label to match the terminology used in your business. This enables you to create service hierarchies with a custom number of tiers that accurately reflect your service structures within your organization, rather than being confined to the three-tier business service/IT service/device service model.

NOTE: Skylar One PowerFlow users can use custom service models and the applications in the "ServiceNow Configuration Management Database (CMDB)"SyncPack to sync business services between Skylar One and ServiceNow.

Using this method, you can create, update, or delete services in ServiceNow and it will be reflected in Skylar One, or vice versa.

However, services that you want to sync between the two systems must either be built entirely in ServiceNow or entirely in Skylar One; you cannot merge services between the two.

ScienceLogic recommends syncing services from ServiceNow into Skylar One rather than building custom service models in Skylar One and syncing them into ServiceNow.

For more information, see the section on "Syncing Business Services" in the ***ServiceNow CMDB SyncPack*** manual.

The Business Services Page

The **Business Services** page displays a list of the business, IT, and device services that you have access to, as well as some basic info and the health, availability, and risk metrics for each service.

To navigate to the **Business Services** page, click the **Business Services** icon ( ):

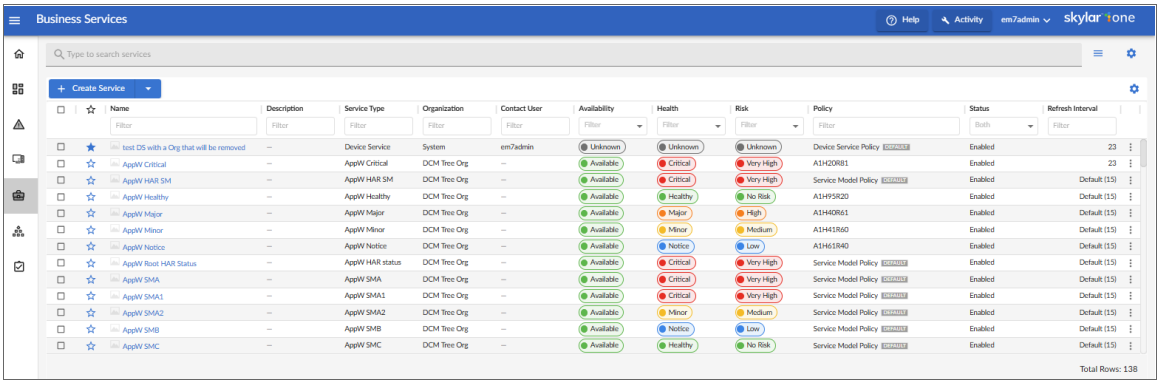

These business services let you gauge the health, availability, and risk of your services or the devices that provide those services. On the **Business Services** page, these values display in the following format and order:

1. ***Availability***: The availability of a device service is derived from the availability rules. This might or might not be linked to device availability. A service or device is considered unavailable if Skylar One is not able to collect data from the device or service, or if a device is usable or not usable. A value of *0* means a device or service is unavailable, and a value of *1* means a device is available. Availability uses the following icons:

2. ***Health***: Indicates the current status of a device service–for example, the rate of processing or throughput for the devices in the device service. Health is represented by a color-coded "severity" icon that corresponds to a numerical value between 0 and 100. In the case of Skylar One Database Servers, the "Rows Behind" presentation objects can provide a good measure of how effectively the Database Server is processing Skylar One Collector data. For example, the health value could indicate when a device is intermittently unavailable because of a power problem, thereby falling below the required level of performance. Health uses the following icons by default:

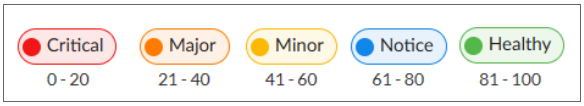

3. ***Risk***: Displays a percentage value between 0 and 100 that indicates how close a service is to being in an undesirable state. The safest possible risk value is 0%, while the worst risk value is 100%. Use risk for data that is known to cause issues if left unchecked, such as critical events, swap usage, or low database logging space. Risk uses the following icons:

These values are computed in this order because Skylar One uses *Availability* values to compute *Health*, and then uses both *Availability* and *Health* values to compute *Risk*.

You can define metrics for *device services* based on:

- availability
- latency
- event count
- event severity
- device state
- Dynamic Application performance data collected by Skylar One
- collection label metrics (for example, CPU, Memory, or Swap)

TIP: You can sort the list by a column's values by clicking on its column heading. You can also filter the items on the page by typing or selecting filter options in one or more of the filters found above the columns. For more information, see the topic on *Filtering Inventory Pages*.

TIP: You can adjust the size of the rows and the size of the row text on this inventory page. For more information, see the topic on *Adjusting the Row Density*.

The **Business Services** page displays the following about each service:

- ***Favorite***. The services that are designated as "favorites" to display at the top of the list.
- ***Name***. The name of the service.
- ***Description***. A description of the service.
- ***Service Type***. The type of service. Values include *Business Service*, *IT Service*, *Device Service*, or a custom service type.
- ***Organization***. The organization that owns the service.
- ***Contact User***. The user who should be contacted with any questions about the service.
- ***Availability***. The service's current availability value.
- ***Health***. The service's current health value.
- ***Risk***. The service's current risk value.
- ***Policy***. The service policy associated with the service.
- ***Last Updated By***. The username of the user who last updated the service.
- ***Date Updated***. The date and time at which the service was last updated.
- ***Service Analysis***. Indicates whether Service Analysis is enabled or disabled for the service.
- ***Contact Organization***. The organization that should be contacted with any questions about the service.
- ***Status***. Indicates whether the health, availability, and risk values are calculated. This field shows "Enabled" or "Disabled".
- ***Refresh Interval***. The frequency at which health, availability, and risk values are calculated. This field is also known as the polling frequency, which is the frequency at which data is communicated between a device and another system.

NOTE: You can update the refresh interval for services in bulk by selecting the checkbox of each service you want to modify, clicking the **[Actions]** button at the top of the page, and selecting *Change Refresh Interval*.

Deleting, Enabling, or Disabling a Business Service

In Skylar One, you can enable, disable, or delete one or more services from the **Business Services** page:

1. Select the checkboxes of the services you want to delete, enable, or disable.
2. Click the **[Actions]** button at the top of the page, and select one of the following actions you want to perform: **[Delete Services]**, **[Enable]**, or **[Disable]**.

> **NOTE:** Alternatively, you can perform any of the actions to a single service by clicking the **Actions** icon (⋮) for that service, and then selecting **[Delete]**, **[Enable]**, or **[Disable]**.

Favorite a Business Service

In Skylar One, you can ***favorite*** one or more services so that they always display at the top of the list on the **Business Services** page.

For example, on the **Business Services** page pictured below, click the ***Favorite Service*** star icon (☆) to add or remove the service from your favorites list. Click the icon (★) again to remove the favorite status.

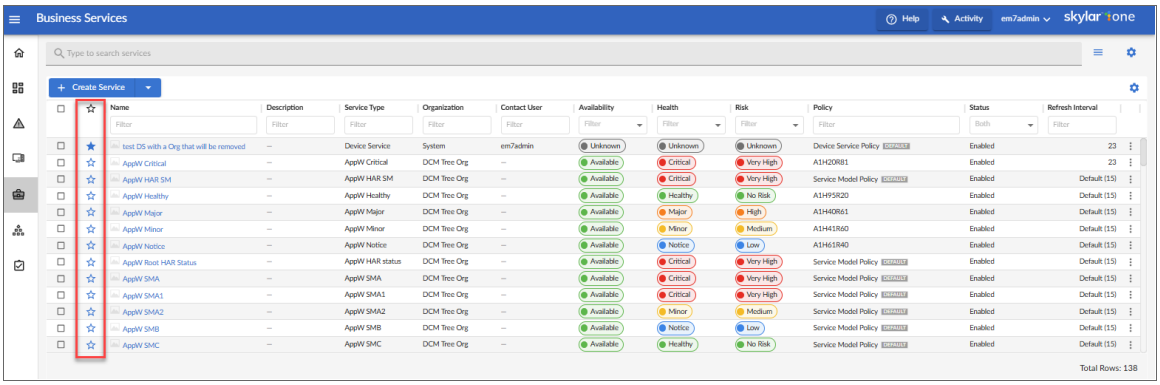

You can then sort your services by their "favorite" status.

With favorite services, you can:

- View your favorite service at the top of the **Business Services** page by default.
- Include favorites in the multi-sort function.
- Filter services by favorite.

Business Service Dashboards

Skylar One includes three default dashboards relating to business services on the **Dashboards** page (▦):

- NOC Overview dashboard
- Business Services dashboard
- Business Service Details dashboard

For more information about these dashboards, see the ***Dashboards*** manual.

In addition to these default dashboards, you can also choose to create your own custom dashboards for business services. For more information, see the ***Dashboards*** manual.

Service Insights Page

The **Service Insights** page (Business Services > Service Insights) provides a comprehensive dashboard for monitoring business service health, usage, and performance. This page displays key metrics and visualizations for all services in your organization at a glance. For more information, see the *Service Insights Page* section.

> **NOTE:** The **Service Insights** page is available only in AP2 Quesito and later releases. If you are running an earlier version of AP2, this feature is not available.

Example: Retail Banking

Using Skylar One to monitor a business service lets you quickly see whether the service is available and working as expected for a customer or end user. For example, a banking company wants to ensure that its retail banking service is available around the world. It would use the following workflow to set up its services in Skylar One:

1. Because the company has offices around the world, it creates multiple ***device services*** that organize devices based on location or region. The company adds all of its devices to the relevant device services.
2. The company then creates multiple ***IT services*** to monitor the device services (from step 1), including separate IT services for online banking, teller systems, and ATM networks.
3. Next, the company creates a ***business service*** for its retail banking business, and this business service includes all of the IT services (from step 2) that deal with retail banking.

> **NOTE:** As needed, the banking company repeats steps 1-3 to create additional business services (made up of IT services and device services) to monitor their commercial banking and investment banking devices and services.

Chapter 2

Creating Services and Service Policies

Overview

This chapter describes how to create and monitor business services, IT services, and device services, as well as custom service models. This chapter also describes how to create and use policies for each service to assist with monitoring those services.

This chapter covers the following topics:

Administrative Processes for Business Services 13

Understanding Health, Availability, and Risk 13

Creating Services 14

Using Service Policies 27

Creating a Service Template 34

Creating a Service From a Template 37

Exporting a Service Template 38

Default Service Policy Settings 40

Assigning an Icon to a Service 41

Exporting Service Data with the ScienceLogic API 42

Administrative Processes for Business Services

Two administrative processes (System > Settings > Admin Processes) are used to calculate business service values:

- ***Business Services: Service Management Engine***. This process aggregates all of the metric information from your devices and services to create the Health, Availability, and Risk values. This is a long-running process which typically runs every 15 minutes. You should only run it more frequently if you have a robust Database Server with room to support the more frequent collection.
- ***Business Services: Service Topology Engine***. This process calculates the relationships between your services and your devices. By default, this process runs every 5 minutes. If your services and device relationships are relatively static, you can decrease the frequency to every 15 minutes.

Understanding Health, Availability, and Risk

> **NOTE:** None of the metrics described in these examples actually pinpoint the exact cause of the unavailability, degradation in health, or increase in risk, but they do bring it to your attention quickly and with a minimal level of administration. When you use key performance indicators (KPI) for responsiveness or availability, you may find it much easier than trying to model every way a service can break.

Understanding Availability

Availability assesses whether something is reachable or is performing at a level to be useful. Here are a few examples to help you understand availability:

- **Website**. The URL for a website must be responsive; that is, it must respond either with the expected page or with an error page indicating that the site is unreachable (up/down). The website's response also needs to be fast enough that users will not leave the page due to a slow response time. This should be considered when defining availability.
- **Cluster of Database Servers**. Assume one Database Server can process 1,000 transactions per second with good response times. To maintain those response times with 3,000 transactions per second, four equivalently configured Database Servers are put into a cluster. This method allows for any one Database Server to be down without losing acceptable throughput and responsiveness. If three Database Servers in the cluster become unavailable, the one remaining Database Server will not be able to maintain throughput or responsiveness, so the cluster is effectively unavailable.
- **Processes**. Consider that process A passes work to process B by way of a queue. If the queue depth sits at zero, it indicates that process A is not passing any new work and is considered to be unavailable. If the queue grows to a specified threshold, it indicates that process B is not pulling work from the queue and is considered to be unavailable.

Understanding Health

A decline in health for a given service or device means that one or more KPIs are degrading. Left unchecked, this can be expected to degrade throughput or responsiveness. Here are a few examples of issues that impact health:

- **Database Server**. On a Skylar One Database Server, a key database function is to retrieve and store events and Dynamic Applications data. You can create device service policies that degrade health as the volume of high frequency (HF) rows climbs, as this indicates the Database Server is becoming overloaded or slow to process incoming data. This could lead to delays in events from Skylar One Collectors being presented to automation actions or the **Events** page, and can impact overall system performance.
- **Windows server**. In some cases, the CPU Queue depth on a Windows server starts to increase, indicating the CPU has insufficient bandwidth to process its workload. When this happens, all processes or applications running on the Windows server will run slowly, impacting either responsiveness or throughput. You can build a policy that lets you know if this is happening on any Windows server.
- **Website**. A website that is the face of an application has increasing web URL response times, indicating stress in the delivery of the URL. If it is known that the URL becomes functionally unavailable at 5 seconds, meaning that your customer may give up and goes to another vendor, then setting health to degrade for 1 to 4 seconds will give notice that the service health is degrading and investigations and resolution can be performed before the URL reaches an unavailable state.

Understanding Risk

In considering risk, think of the consequences of a KPI degrading. If a selected KPI is known to indicate situations that, if left unaddressed, will impact health or availability, you will want to create a policy for that. Some examples:

- On a Skylar One Database Server, if the InnoDB table runs out of space, MariaDB will stop, which leads the Database Server to become unavailable. A shrinking level of available InnoDB space will not degrade the responsiveness and throughput of MariaDB, and therefore the Database Server, but it can indicate that your Database Server availability is at risk.
- Another way to measure risk for devices in a service is by monitoring the level of severity for events. This provides a reasonable baseline for risk. For example, many critical events for a device either indicates a false positive that should be suppressed or that monitoring has found a condition that is deemed to be unacceptable.

Creating Services

You can create services in Skylar One using one of two methods:

- You can create a *three-tier service*, consisting of a business service, one or more IT services, and one or more device services.
- You can create a *custom service model* with a user-defined business hierarchy.

NOTE: The "Business Services Sample Services and Policies" PowerPack provides a set of sample device services and service policies that cover most scenarios that users new to business services will want to address. For instructions on how to import and install a PowerPack, see *Installing a PowerPack*.

When designing your service structure, a good design principle is to begin with the end in mind. To create a new business service, you should first determine the following:

- *Stakeholders*. Who is the intended audience for the service?
- *Purpose*. What problem are you trying to solve for your stakeholders?
- *Visibility*. Who needs to see which services?
- *Workflow*. How are your stakeholders currently performing fault isolation?
- *Right-sizing*. What is the right number of services? Consider the following:
    - The devices that impact the business service
    - The IT services that impact the business service
    - The specific conditions that you want to monitor, based on your business processes

If you follow the design flow described above, you will have an outline of which model type to create and which specific services you need to build. For example, if you provide email service, then a failure of your primary SMTP server and backup SMTP server would constitute a Critical status.

The next consideration is to determine which devices share a common description of health, availability, and risk rules. If two devices need different rules, you will need to create two Device Services.

TIP: You can copy an existing service on the **Business Services** page by clicking the **[Actions]** button (⋮) for that service and selecting *Duplicate*.

Creating Business, IT, or Device Services

To create a business, IT, or device service:

1. On the **Business Services** page, click the **[Create Service]** button. The **New Service** page appears.

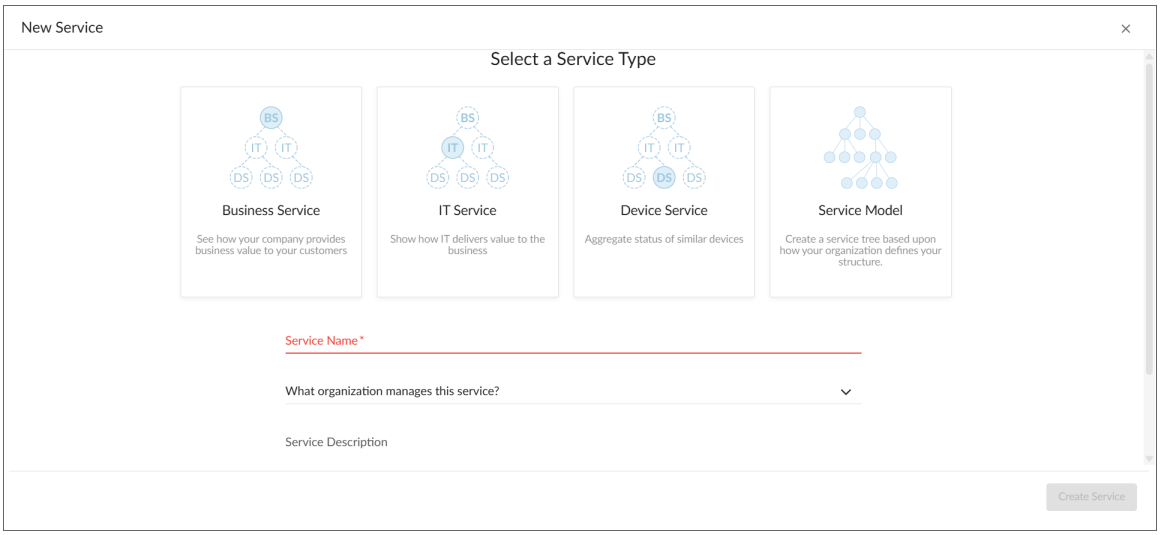

2. Select a service type. You should start by creating your device services, then your IT services, and then finally your business service. Your options include:
    - ***Device Service***. Monitors a set of related devices.
    - ***IT Service***. Monitors a service that IT provides to your organization. An IT service includes one or more device services.
    - ***Business Service***. Monitors a service that your organization provides to your customers. A business service includes one or more IT services.

3. Complete the remaining fields:
    - ***Service Name***. Type a unique name for this service.
    - ***What organization manages this service?***. Select the name of the organization that owns this service.
    - ***Service Description***. (Optional) Type a short description of this service and its purpose. You can use the text in this description to search for this service on the **Business Services** page. For example, if a collection of device services all have a description of "Shared Infrastructure", then an IT service can search to include every device service in the same organization that has a description of "Shared Infrastructure". As you add more "Shared Infrastructure" device services, the IT service will automatically expand to include them. This makes building service trees quick and self-maintaining, without resorting to rigid service names.
    - ***Visible Organizations***. (Optional) Select one or more organizations from which you can select devices to use in the service. For example, if you selected *Acme* for this field, then any service that is aligned with Acme can access devices in the Acme organization.

    > **NOTE:** The ***Visible Organizations*** field allows the selected organizations to view the service and enables these organizations to query the service. For instance, if you want an IT service to have a device service as a child service, the device service will either need to be within the same organization as the IT service, or the device service will need to be included in the visible organizations that are aligned with the IT service.

4. Click the **[Create Service]** button. If you selected *Device Service* in step 2, the **[Devices]** tab appears, with a list of available devices in the ***Preview*** section. If you selected *Business Service* or *IT Service* in step 2, the **[Services]** tab appears, with a list of available services in the ***Preview*** section.
5. In the ***Search*** field, type search criteria for the services or devices you want to monitor. A list of services or devices that match your search criteria appears in the **Preview** pane.

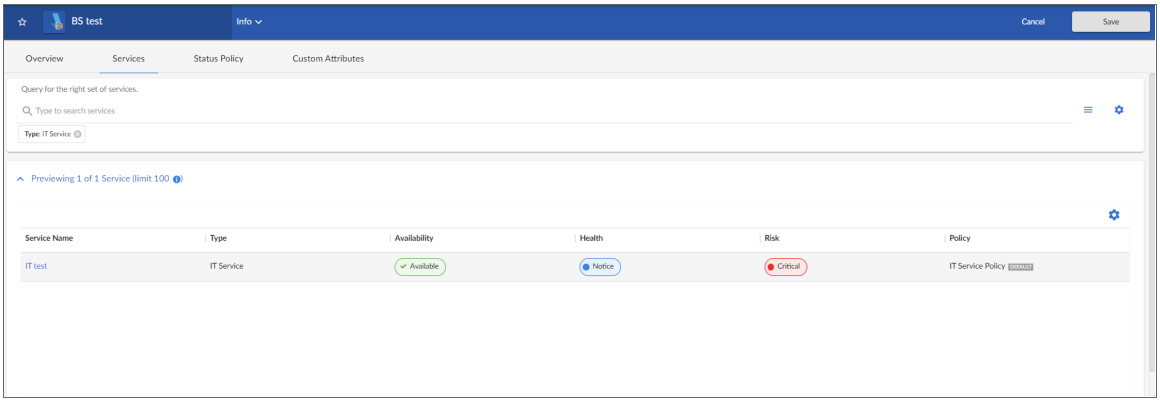

TIP: If you are looking for a very specific set of services or devices, click the gear icon (⚙) to the right of the ***Search*** field and select *Advanced*. In this mode you can create an advanced search using "AND" or "OR" for multiple search criteria. You must click **[Search]** before you can view the results of your search. For more information, see the "Performing an Advanced Search" topic in the ***Introduction to Skylar One*** manual.

TIP: If you want to search for devices that have specific custom attributes, use Advanced Search. Use the following format:

```
attribute has (id == custom attribute and value == value)
```

Note that search cannot process colons (:) in strings. The presence of a colon in service inclusion searches will stop the engine that calculates health, availability, and risk for that service.For more information, see the "Advanced Search" topic in the ***Introduction to Skylar One*** manual.

NOTE: The "ANY" search option is disabled on the **[Services]** or **[Devices]** tab.

TIP: The **Preview** pane indicates the maximum number of constituent services or devices that will be used for computing health, availability, and risk.

NOTE: There is a strict limit of 100 children for each service, whether it involves devices or services. If the rules of a service permit more than 100 children, the results might become unpredictable, as it will be unclear which 100 children will be chosen. You need to modify your design if you exceed this limit.

6. When you have the right combination of services or devices, click **[Save]**. The default policy for the type of service you selected is automatically added to the new service.
7. If you want to use a different business policy with the new service, see *Selecting a Business Service Policy*.
8. If you want to create a *new* business policy to use with the new service, see *Creating a Business Service Policy*.
9. Repeat this process until you have the right combination of device services and IT services in your business service (or business services, if needed).

Creating a Custom Service Model

If you require more flexibility in modeling your service beyond the standard three-tier business service/IT service/device service model, you can instead create a custom ***Service Model***, which is also called an "N-tier" service model.

For example, if you needed to monitor individual conference rooms within your organization from a high level, you could build the following service model to do so:

Country > Region > State > City > Building > Floor > Conference Room

Custom service models enable you to build service hierarchies with a unique number of tiers that accurately reflect the service structures within your organization, using user-defined labels that can match the terminology used in your business.

NOTE: You cannot convert a three-tier business service model to a custom service model.

NOTE: Skylar One PowerFlow users can use custom service models and the applications in the "ServiceNow Configuration Management Database (CMDB)" Synchronization PowerPack to sync business services between Skylar One and ServiceNow.

Using this method, you can create, update, or delete services in ServiceNow and it will be reflected in Skylar One, or vice versa.

However, services that you want to sync between the two systems must either be built entirely in ServiceNow or entirely in Skylar One; you cannot merge services between the two.

ScienceLogic recommends syncing services from ServiceNow into Skylar One rather than building custom service models in Skylar One and syncing them into ServiceNow.

For more information, see the section on "Syncing Business Services" in the ***ServiceNow CMDB SyncPack*** manual.

To create a custom service model:

1. On the **Business Services** page, click the **[Create Service]** button. The **New Service** page appears.

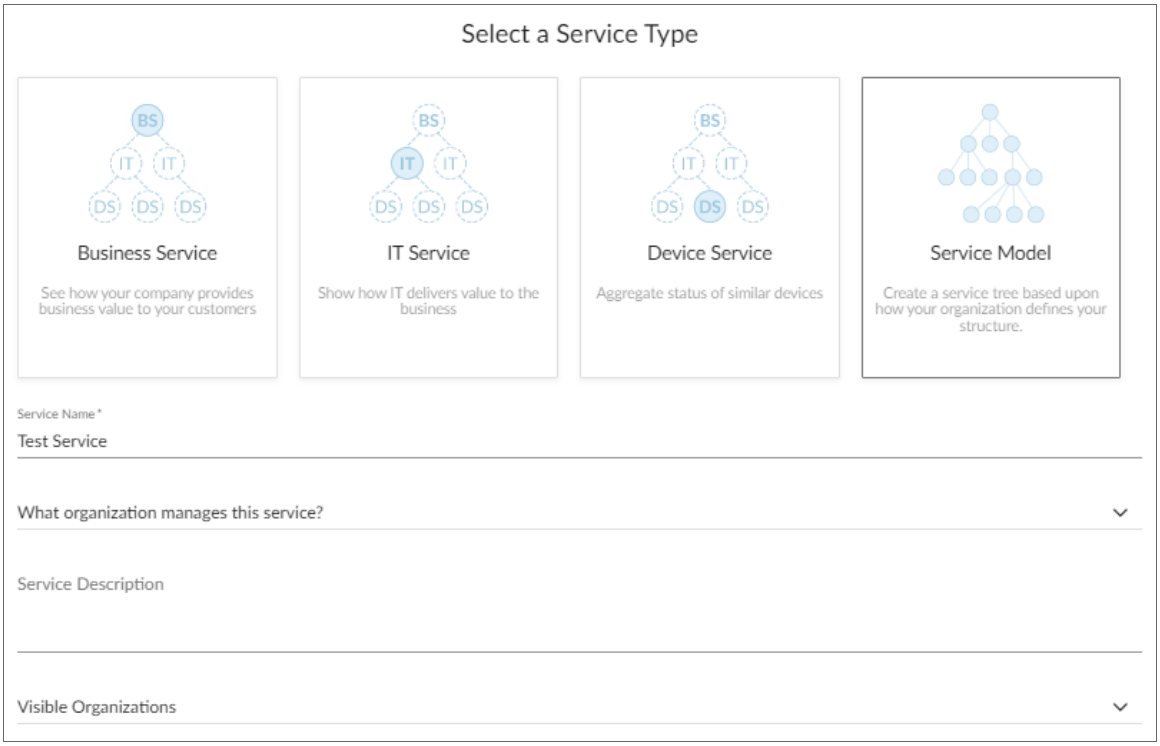

2. Select *Service Model*, and then complete the following fields:
    - ***Service Name***. Type a unique name for this service model.
    - ***What organization manages this service?***. Select the name of the organization that owns this service model.
    - ***Service Description***. (Optional) Type a short description of this service model and its purpose. You can use the text in this description to search for this service on the **Business Services** page. For example, if multiple services all have a description of "Shared Infrastructure", then you could create another service search to include every service in the same organization that has a description of "Shared Infrastructure". As you add more "Shared Infrastructure" services, the other services that include those "Shared Infrastructure" services will automatically expand to include them. This makes building service trees quick and self-maintaining, without resorting to rigid service names.
    - ***Visible Organizations***. (Optional) Select one or more organizations from which you can select devices to use in the service model. For example, if you selected *Acme* for this field, then any service that is aligned with Acme can access devices in the Acme organization.

> **NOTE:** The ***Visible Organizations*** field allows the selected organizations to view the service and enables these organizations to query the service. For instance, if you want an IT service to have a device service as a child service, the device service will either need to be within the same organization as the IT service, or the device service will need to be included in the visible organizations that are aligned with the IT service.

3. **[Next]**. The service creation page of the **New Service** wizard appears.
4. Click **[Next]**. The model selection page of the **New Service** wizard appears.
5. On the model selection page, do one of the following:
   - Click the **[Add Model]** button to design a new service model. If you do this, proceed to step 6.
   - Use the search bar if necessary to search for an existing model to which you want to make customizations or changes. If you do this, skip ahead to step 7.
   - Use the search bar if necessary to search for an existing model to which you *do not* want to make any customizations or changes. Select that model's radio button and then click **[Next]**. If you do this, skip ahead to step 9.

> **TIP:** If you are looking for a very specific set of models, click the gear icon (⚙) to the right of the ***Search*** field and select *Advanced*. In this mode you can create an advanced search using "AND" or "OR" for multiple search criteria. For more information, see the "Performing and Advanced Search" topic in the ***Introduction to Skylar One*** manual.

6. On the **New Model** modal, type a name for the service model in the ***Model Name*** field and then click **[Create]**. The newly created service model is added to the model selection page.

7. On the model selection page, select the new service model you just created or an existing model to which you want to make customizations or changes, and then customize the model as needed:

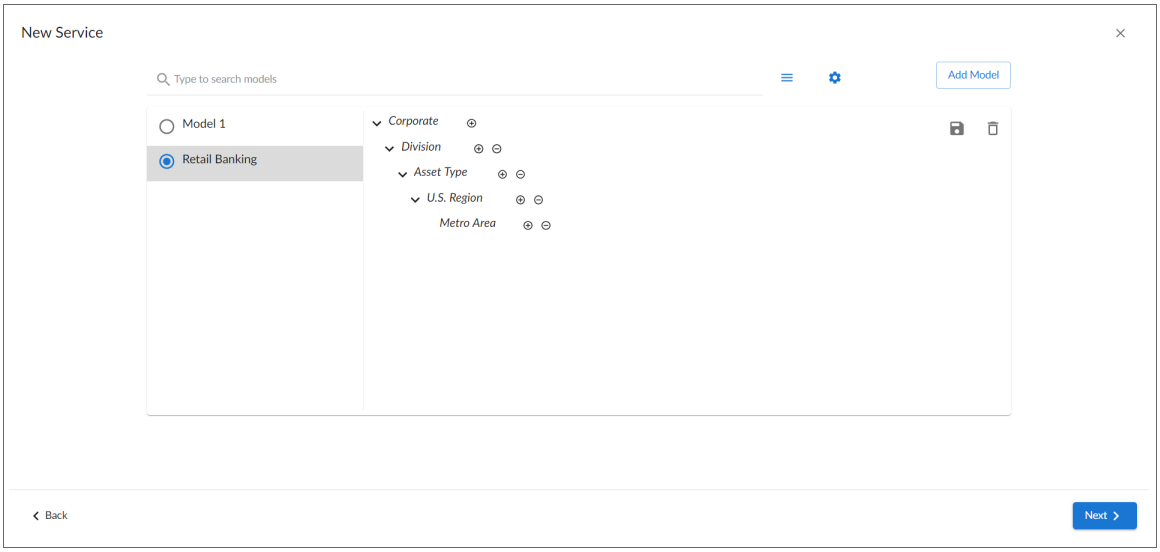

You can customize the service model to fit your business needs in the following ways:

- Click the model name, then type a new model name to replace the existing name.
- Use the plus (⊕) and minus (⊖) icons to build the tiers of your service model. For each tier, click the tier label and then type a new tier label to replace the existing label.
- Click the save icon (💾) to save your custom service model.
- Click the delete icon (🗑) to delete your custom service model.

8. When you are finished customizing the service model, click **[Next]**. The hierarchy creation page of the **New Service** wizard appears.
9. On the hierarchy creation page, use the model you selected to build out the full hierarchy structure of your service:

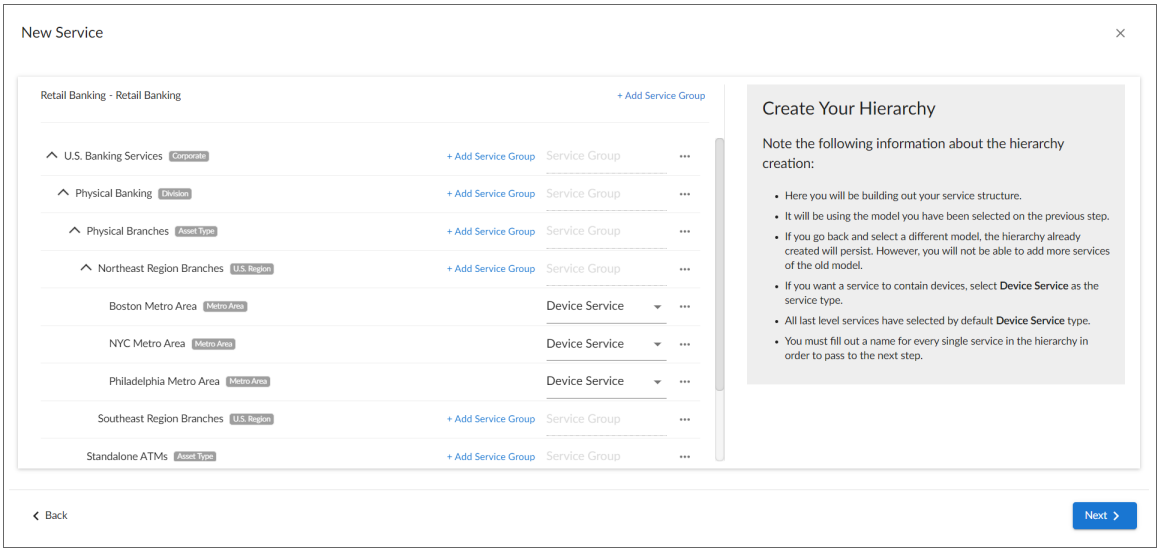

You can customize the service hierarchy to fit your business needs in the following ways:

- Click the name of each tier in your service model, then for each tier, type a new label that is specific to the service you want to monitor. For example, your service model might have a tier labeled "U.S. Region"; you might click that tier and type "Northeast" if your service includes assets in the Northeast.

NOTE: You must fill out a name for every service in the hierarchy before you can proceed to step 10.

- Click the expand (⌄) and contract (⌃) icons to expand and contract parts of the service hierarchy tree.
- Click a hierarchy row to drag and drop that row (and any rows that fall below it on the service hierarchy) into a different location in the service hierarchy. If you are able to drop the row in a particular location, the row will turn solid blue. If you are unable to drop the row in a particular location, the row will be white with a red border.
- Click **[+Add Service Group]** at the top of the service hierarchy to add a copy of the entire service model structure to the hierarchy. For example, if your service model includes five tiers and you click **[+Add Service Group]** at the top of the service hierarchy, Skylar One will add an additional set of all five tiers to the hierarchy.

- Click **[+Add Service Group]** for a particular row to add to a new set of sub-tier rows in the hierarchy under the existing row. For example, if you have a service model tier labeled "Physical Banking" and the next level of the service model under that is a tier labeled "Branch Locations" and under that are several other tiers, then when you click **[+Add Service Group]** on the "Physical Banking" row, Skylar One will add a new "Branch Locations" tier along with all of its sub-tiers under the "Physical Banking" row.
- Click the Actions button (···) for a particular row and then select *Delete* to delete that row from the hierarchy.
- If you want the next-to-last-level service row within the hierarchy to have device services below it, click the Actions button (···) for that row and then select *Add Device Services*. The **Add Device Services** modal appears. On that modal, search for and select any existing device services that you want to add below the selected service row, and then click **[Add Services]**.
- If you want the last-level service row within the hierarchy to contain devices, select *Device Service* from the service type drop-down at the end of the row. If you do not want the row to contain devices, select *Service*. All last-level services are set to *Device Service* by default. All service rows that are not last-level rows have a service type of *Service Group*.

10. When you are finished customizing your service hierarchy, click **[Next]**. The **[Overview]** tab for the new service appears:

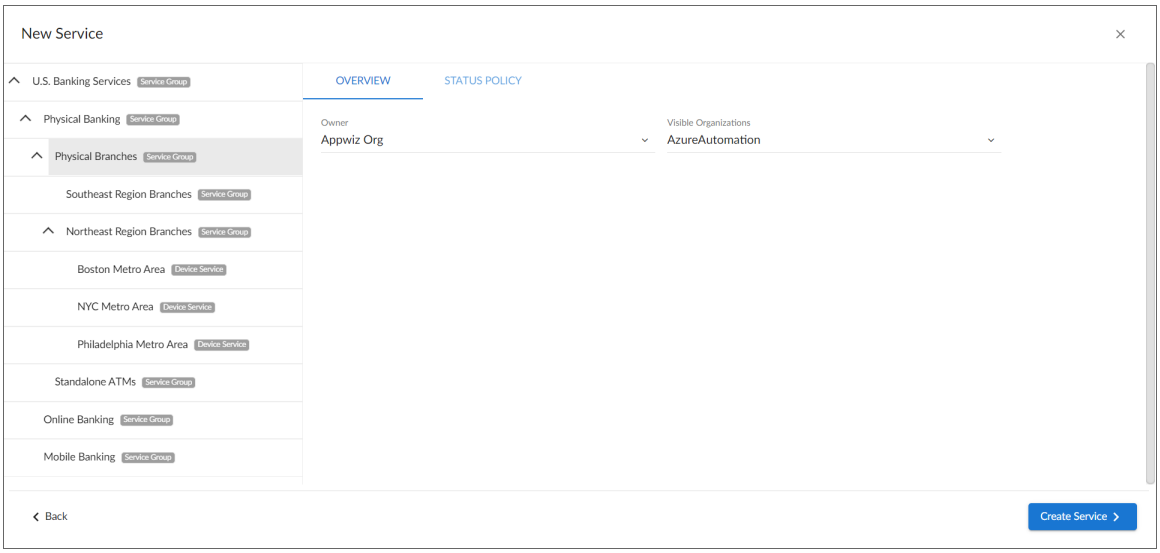

11. On the **[Overview]** tab, you can update the managing organization and visible organizations for the individual levels within your service model hierarchy if they differ from the managing organization and visible organizations you selected for the entire model in step 3. To do so, click the hierarchy level on the left side of the page, and then complete the following fields:
    - ***Owner***. Select the name of the organization that owns the selected service level.
    - ***Visible Organizations***. Select one or more organizations from which you can select devices to use in the selected service level.

12. Click the **[Service Policy]** tab:

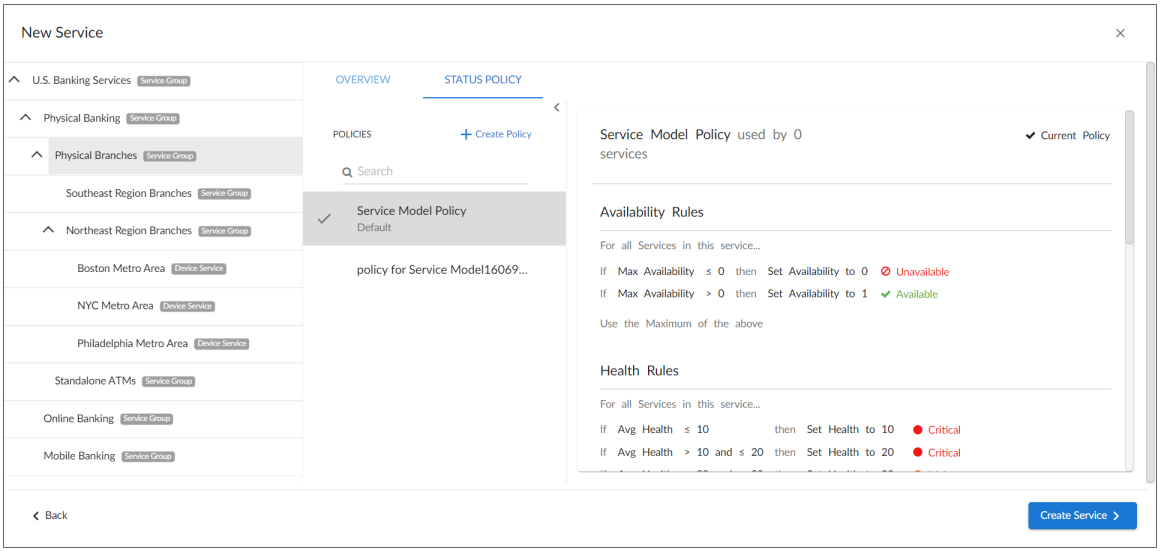

13. On the **[Service Policy]** tab, click on each row within your service hierarchy and then do one of the following for each:
    - To apply an existing policy to the selected service level, search for and select the policy that you want to apply. When you select a policy from the list, the details of that policy appear in the right panel. If a service policy is already applied to a service level, "Current Policy" appears in the top right corner of the right panel; otherwise, a **[Use Policy]** button appears. To apply a different policy, click the **[Use Policy]** button.
    - To create a new policy, click the **[Create Policy]** button. In the **Create Policy** modal, type a ***Policy Name***, and then click **[Create Policy]**. The policy is added to the list. To apply it to the selected service level, select the policy from the list and then click the **[Use Policy]** button in the top right corner of the right panel.
14. Click **[Create Service]**. A confirmation message displays, indicating that your service model and its hierarchy have been created successfully.

> **NOTE:** After you have created the service model, you cannot use the **New Service** wizard to edit the model.

15. Click **[Close]**.
16. If you want to use a different service policy with the new service, see *Selecting a Business Service Policy*.
17. If you want to create a *new* service policy to use with the new service, see *Creating a Business Service Policy*.

Advanced Search Criteria: Practical Use Cases

When creating services, you can type search criteria for the services or devices you want to monitor in the ***Search*** field. If you are looking for a very specific set of models, click the gear icon (⚙) to the right of the ***Search*** field and select *Advanced*. In this mode you can create an advanced search using "AND" or "OR" for multiple search criteria. For more information, see *Performing an Advanced Search*.

- If you want to search for devices that have specific custom attributes, use the following format:

  ```
  attribute has (id == <custom-attribute> and value == <value>)
  ```

  where:

  - `<custom-attribute>` is the unique identifier or name of the custom attribute you want to filter by. This variable is defined in your system or device management platform when you create custom attributes.
  - `<value>` is the specific value assigned to that custom attribute for the devices you want to find.
    For example:

    ```
    attribute has (id == "Environment" and value == "Production")
    ```

- To search for devices aligned with a specific Dynamic Application, use the following format:

  ```
  alignedDynamicApplication has (dynamicApplication has (name contains
  '<dynamic-application-name>') )
  ```

  where:

  - `<dynamic -application-name>` is the name of the Dynamic Application.
    For example:

    ```
    alignedDynamicApplication has (dynamicApplication has (name
    contains "DNS Dynamic Update Performance") )
    ```

- To search for all virtual machines (VM) whose names begin with a specified prefix that are hosted on a specified ESX service, use the following format:

  ```
  name eq '<esx-Service-Name>' and deviceClass has (description eq
  'host server')) and componentAncestor has (name beginsWith '<VM-name-
  prefix>')
  ```

  where:

- `<esx-Service-Name>` is the ESX service you want to filter on
- `<VM-name-prefix>` is the prefix for VM names.
  For example:

```
name eq 'Prod1' and deviceClass has (description eq 'host
server')) and componentAncestor has (name beginsWith 'ABC')
```

NOTE: Depending on the VM installation, the advanced search could be: `name eq '<esx-Service-Name>' and deviceClass has (description eq 'ESC(i)')) and componentAncestor has (name beginsWith '<VM-name-prefix>')`.

Using Service Policies

Each service type requires a ***policy*** that determines what it monitors. A business service policy contains a set of rules and conditions that define the health, availability, and risk (HAR) values for the service, depending on your business needs. Each service requires that one policy be associated with a service at a time.

Creating a Service Policy

When you create a business service of any type, Skylar One automatically creates a policy without any rules for health, availability, or risk. A policy consists of a set of ***rules***, and each rule begins with three ***conditions***. You can also add more conditions to each rule.

NOTE: Service policies for device services must not use Dynamic Application performance objects with polling frequencies longer than 15 minutes. If you select a Dynamic Application performance object with a polling frequency longer than 15 minutes, the health, availability, and risk values of these services will not be calculated.

NOTE: Before you configure your service policy, it is important to understand why each severity can be configured by selecting a severity level for the service policy. For example, you can assign a number to the **Critical** severity level. If the number is exceeded, the health, availability, or risk rules will display the severity that was assigned. Suppose a device service for Linux servers has two risk rules: one for memory utilization and one for swap utilization. A server that has exhausted memory but still has free swap space to expand into will continue running but will slow down. A server that has exhausted swap space is likely to fail. Therefore, while both statuses can be bad, the lack of free swap space is worse than having low memory. When building risk rules, you could set 95% memory utilization as Critical with a score of 85, but set swap at 95% utilization to Critical with a score of 95. This will indicate that swap space is more causal then memory, and that as soon as you fix the swap space issue, you will need to check into the problems with memory.

To create a policy:

1. On the **Business Services** page, select the service for which you want to create a policy. The **Service Investigator** page appears.
2. Click the **[Edit]** button in the upper right corner of the page. The **[Overview]** tab for the service appears.
3. Click the **[Service Policy]** tab, and then click **[Create New Policy]** in the ***Service Policies*** section. A **Create Service Policy** window appears.
4. Enter a policy name in the ***Name*** field. You can also enter a description of the policy in the ***Description*** field.
5. In the ***Health***, ***Availability***, and ***Risk*** columns, click the **[Add Rule]** button to create the rules and conditions for each of the three values that make up this policy. Each column uses the same layout. The **Edit Rule** modal appears.

> **NOTE:** The "Availability" vital metric is not populated for component devices. Therefore, any device service that includes component devices will have a null "Availability" value, which is displayed as a hyphen. However, a potential alternative is to change the rule from "Availability" to "Count" and query devices that are enabled to collect data (`isActive = true`). This alternative works because "Count" is the number of devices that matches the filter query, and querying on `isActive` matches all devices that are currently collecting data.

6. In the **Edit Rule** modal, complete the following fields:

- ***Aggregate***. Select an aggregation method for the data for this condition. Your options include *Average*, *Minimum*, *Maximum*, *Count*, *Sum*, and *Percent*.

  For example, suppose you have a web server farm consisting of three web servers. You have created a rule for web response time and are building for ***Health***. By default, the aggregation factor for IT, business, and service model service policies is *Minimum*.

  - *Minimum* will drive health based on the fastest responding web server.
  - *Maximum* will drive health based on the slowest responding web server.
  - *Average* will drive health based on the average between slowest and fastest. This might give false positives. For example, assume that 5 seconds is the ideal target response time. If web server 1 gives a .1-second response time, web server 2 gives a 5-second response time, and web server 3 gives a 10-second response time, then the average will be 5 seconds, masking the fact that one of the response times is grossly unacceptable.
  - *Count* determines how many devices are currently being included in the device service. (The devices must be available as seen on the **Devices** page). This is useful if you need at least 2 out of your 3 web servers to be active at any one time. If the rule does not have a rule filter, this value will display the number of devices in the service. If the rule has a rule filter, this value will display the number of devices and services that match the rule filter.
  - *Sum* is the result of adding up the value of the metric from all devices currently included in the device service. This is useful when you need to know how many devices are available across all the devices in the device service.
  - *Percent* determines the number of devices as a percentage of the total number of devices for the service query that fit the criteria of the filter rules.

    NOTE: For example, you might have a rule filter for devices for a service. If the service has 10 devices, and 5 match the rule filter, the percentage will be 50%.

- ***Metric***. Select the metric you want to monitor for this condition:
  - If this is a business service or an IT service, your options include *Availability*, *Health*, and *Risk* for the services you want to monitor.
  - If this is a device service, select a device metric such as vitals like *Availability* and *Latency*, performance metrics, metrics collected by the Skylar One agent, or Dynamic Application metrics.

> **NOTE:** In the **Edit Rule** modal, when creating or editing rules for specific service policies, and when hourly summarization is available, a preview line graph displays threshold levels for these policies summarized on an hourly basis. Preview data is not available for all devices in a service, and can trigger a warning explaining that the preview data is not complete.

- ***Timespan***. Select a time frame for the data in the graph in the ***Conditions*** section below. Your options include *Day*, *Week*, and *Month*. This field uses summarized data. If the Dynamic Application's performance object is not set to *Summarize*, the preview graph will not display.

7. In the ***Conditions*** section, select a severity level and the inequality sign. Then, add the numeric value you want to set for the severity level. You can add more conditions by clicking the **[Add Condition]** button.

> **NOTE:** When creating or editing a service policy, you can insert new conditions directly into any rule within that policy by clicking the plus icon (+) under which you want to insert a new condition, rather than adding them to the end of the list of conditions.

> **TIP:** To *remove* a condition from a rule, click the "X" icon for that condition in the **Edit Rule** modal. Alternatively, from the **Edit Service Policy** page, you can click the **[Actions]** button (⋮) for that condition and select *Delete*. To *copy* a condition, click the **[Actions]** button (⋮) for that condition and select *Duplicate*.

> **TIP:** You can also set custom alert thresholds for each severity level.

> **NOTE**: The availability value calculates only the minimum (unavailable) and maximum (available) values for rules.

8. Edit any additional conditions or rules on the remaining columns for this policy, and then click **[Save]**. The new policy is added to the **Service Policies** section on the **[Service Policy]** tab. The policy will also display on the **Service Policies** page (Business Services > Service Policies).

NOTE: You can save a service policy only when it has valid rules and conditions for health, availability and risk. If the rules' condition operators and values are not logically consistent, a red warning message displays and the **[Save]** button is disabled.

NOTE: To edit your new service policy from the **[Service Policy]** tab, you can click the **[Actions]** button (⋮) for the new policy and select *Edit* or click the **[Edit]** icon (✎).

Use Cases for Service Policy Creation

This section outlines common scenarios for creating service policies, which define what a service monitors and how it behaves. A service policy is a collection of rules that define the service's health, availability, and risk status. These rules can use any combination of performance metrics, data, events, or device state information provided by the devices within the service.

For example, a policy can include conditions that specify HAR (Health, Availability, and Risk) values for a service, ensuring that critical metrics are tracked consistently. Other use cases could involve applying environment-specific monitoring rules, or enforcing compliance standards across multiple services. By leveraging service policies, you can tailor monitoring to meet operational requirements and maintain service health effectively.

Below are several example use cases for service policies:

- Availability threshold policy:
    - You need at least 50% of the devices in a service to be active for monitoring to classify the service as available. This policy helps maintain service reliability by requiring a minimum level of device availability before monitoring status is marked as healthy.
    - You need at least eight of the ten devices in a service to be active for monitoring to classify the service as available.
- Health Impact Based on MS-SQL Buffer Cache Hit Ratio
    - Configure the service health to decrease as the MS-SQL buffer cache hit ratio drops, indicating potential performance degradation in database operations.
- Risk Adjustment Based on **/var/log** Free Space
    - Increase the service risk level as the percentage of free space in **/var/log** decreases, helping to prevent issues related to log storage exhaustion.

- Aggregate Service Health Based on Constituent Services (Static Count)
    - Reflect the overall health of an aggregate service based on the health of its constituent services:
        - If all constituent services are healthy, the overall health is considered "Healthy".
        - If two to five constituent services have major or critical health, the overall health is considered "Major".
        - If more than five constituent services have major or critical health, the overall health is considered "Critical".
- Aggregate Service Health Based on Constituent Services (Percentage-Based)
    - Reflect the aggregate service health based on the percentage of constituent services with major or critical health:
        - If all constituent services are healthy, the overall health is considered "Healthy".
        - If 25% - 50% of the constituent services have major or critical health, the overall health is considered "Major".
        - If more than 50% of the constituent services have major or critical health, the overall health is considered "Critical".

Selecting a Service Policy

Each service type requires a ***policy*** that determines what it monitors. A business service policy contains a set of rules and conditions that define the health, availability, and risk (HAR) values for the service, depending on your business needs. Each service requires that one policy be associated with a service at a time.

NOTE: The "Business Services Sample Services and Policies" PowerPack provides a set of sample device services and service policies that cover most scenarios that users new to business services will want to address. For instructions on how to import and install a PowerPack,, see *Installing a PowerPack*.

When you create a business service of any type, Skylar One automatically uses the *default* policy for that particular type of business service. You can remove the default policy after you create a new policy. The default policies cannot be edited.

TIP: If a policy contains errors, an error icon (⚠) appears next to the policy name. To view details about what makes the policy invalid, select the policy and hover over the error icon next to the

policy name in the right-hand section. A pop-up window lists the problems with the policy. An error is most likely to occur between the time the policy is saved and the next HAR aggregation cycle. For best results, wait for the next HAR cycle before investigating whether there is a true error.

To select an existing business service policy:

1. On the **Business Services** page, select the service that needs a policy.
2. Click the **[Actions]** button (⋮) for that policy and then select *Edit*. The **[Devices]** or **[Services]** tab appears, depending on the service type.
3. Click the **[Service Policy]** tab.
4. In the **Policies** section on the left, select the policy you want to use.

TIP: You can type basic search criteria in the ***Search*** field to locate a specific policy in the list.

5. To view the details of a selected policy, click the **[Edit]** icon (✎) for that policy. The **Edit Service Policy** page appears.
6. To edit a service policy, make updates to the policy while in edit mode and then click **[Save]**. Otherwise, click the **[Cancel]** button when you are done viewing the details for that policy.

NOTE: Each column shows the number of associated rules in parentheses, provides a list of the rules in both view and edit modes, and includes an **[Add Rule]** button when in edit mode.

7. To add a policy to the service, select the policy in the **Service Policies** section and click the **[Actions]** button (⋮) for that policy in the right-hand section, then select *Use*. A check mark icon (☑) appears next to that policy in the **Policies** section.
8. To make a copy of a policy, click **[Actions]** button (⋮) for that policy and select *Duplicate*.
9. To delete a policy you no longer want to use, click the **[Actions]** button (⋮) for that policy, select *Delete*, and then confirm that you want to delete the policy. If that policy is used by any other services, those services are then assigned the default policy type. You cannot delete a default policy.

Managing Service Policies

The **Service Policies** page (Business Services > Service Policies) lists all of your service policies and their respective statuses and service counts, among other information.

The **Service Policies** page includes the following information:

- ***Name***. The name of the service policy.
- ***Type***. The type of service policy.
- ***Status***. Indicates the service policy's state, showing either "Valid" or "Invalid".
- ***Service Count***. Indicates the number of services assigned to the service policy. To view all services aligned to a service policy, click on the number that displays under the ***Service Count***column. When you do so, the **Business Services** page appears, pre-filtered to only those services aligned to the service policy.

> NOTE: You can delete a service policy only if the ***Service Count*** is zero. If you try to delete a service policy that has services assigned to it, you will receive an error message.

- ***Date Edited***. The date the service policy was last edited.
- ***Last Edited By***. The user to create or last edit the service policy

Deleting a Service Policy

On the **Service Policies** page (Business Services > Service Policies), you can search for and delete one or more service policies.

To delete a single service policy from the **Service Policies** page, click the **[Actions]** button for the service policy you want to delete and then select *Delete*.

To delete multiple service policies from the **Service Policies** page, select the checkboxes of the policies you want to delete and then click **[Delete Policies]** at the top of the page.

> TIP: To select all visible service policies, click the checkbox at the top of the list on the **Service Policies** page. Clear the checkbox at the top of the list to unselect all of the service policies.

> NOTE: When deleting a service policy, the **Delete Policy** modal displays all of the services that are currently using that policy. If the policy is deleted, the associated services will revert to using the default service policy

Creating a Service Template

You can create a ***service template*** from an existing service to simplify the process of replicating an entire service or service hierarchy on another Skylar One system. For example, if you want to create the same service hierarchy, but only change the owner of the service hierarchy, creating a service template from an existing service streamlines this process.

To create a service template:

1. On the **Business Services** page, click the **[Actions]** button (⋮) for the service you want to use as the basis for your template and select ***Create Template***. The **Create Template From Service** window appears. This window contains important information about what you can and cannot do with a service template.
2. After reading the information that appears on the **Create Template From Service** window, click **[Next]**. The next **Create Template From Service** window appears:

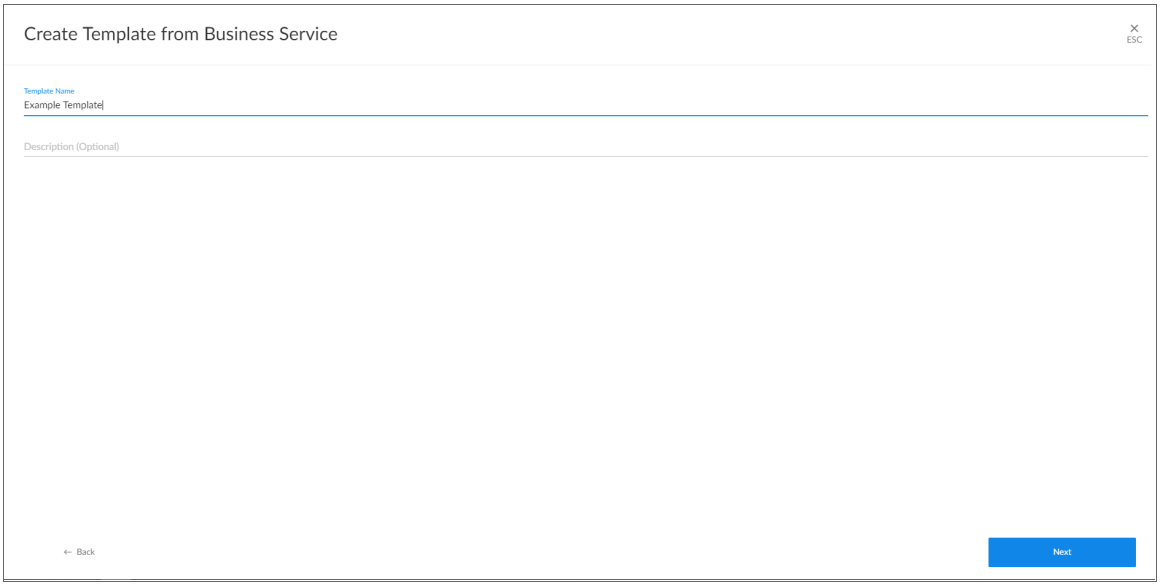

3. Type a name for the template in the ***Template Name*** field and a description of the template in the ***Description*** field, if needed. Click **[Next]**. The next **Create Template From Service** window appears:

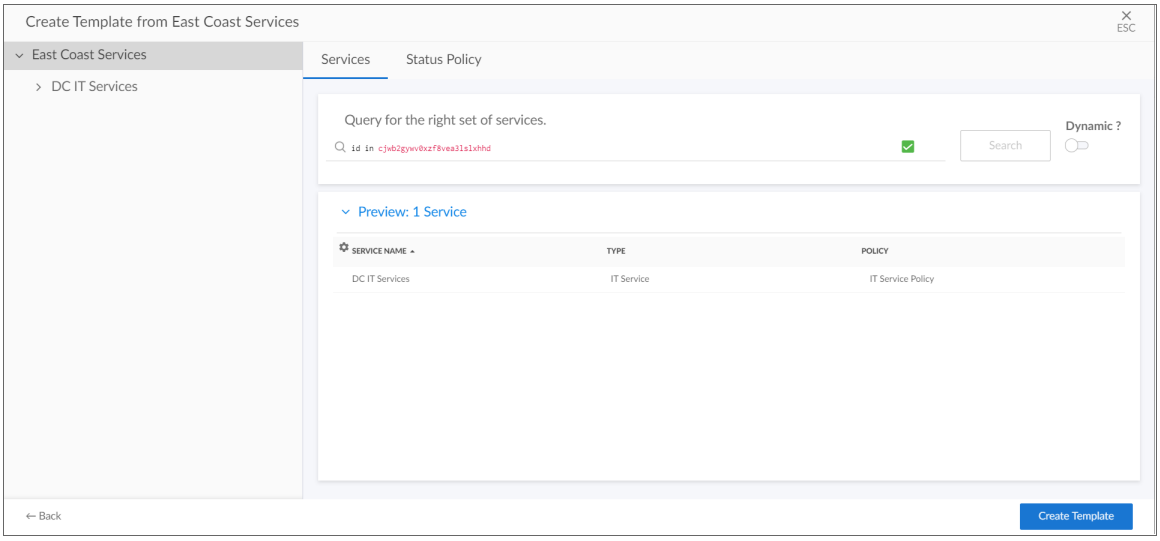

4. The left side of the window displays the tree for the service hierarchy that is being made into a template. You can select each service in the tree to see information related to that service on the right side of the window. For example, if you select a device service, the **Devices** tab displays the search query used for the devices included in that service. If you select a business service or an IT service, the **Services** tab displays the search query for that service. Note the following about the **Dynamic?** slider.

    - If **Dynamic?** is disabled, the template inherits the result of the services inclusion search. This is useful if you want to lock the service tree at the time of template creation. For example, a Managed Service Provider (MSP) might do this to allow end customers to create services from the template but not to modify them. Another use case is if you want to use searches for tags to lock in a set of services that matched the rules at template creation time. By default, **Dynamic?** is disabled.
    - If **Dynamic?** is enabled, the original rule is maintained in the template, so every service tree created from the template will be dynamic based on the services that match the rules.

TIP: The search uses the Advanced Search mode that lets you use "AND" or "OR" for multiple search criteria. For more information, see the "Performing an Advanced Search" topic in the ***Introduction to Skylar One*** manual.

5. Click the **[Service Policy]** tab to view the service policy definition for ***Health***, ***Availability***, and ***Risk*** for that service.
6. On the **[Service Policy]** tab for a device service, you can add annotations for the policies in the template. When a new user uses the template on another system, your annotations can help that user understand the purpose of this policy.

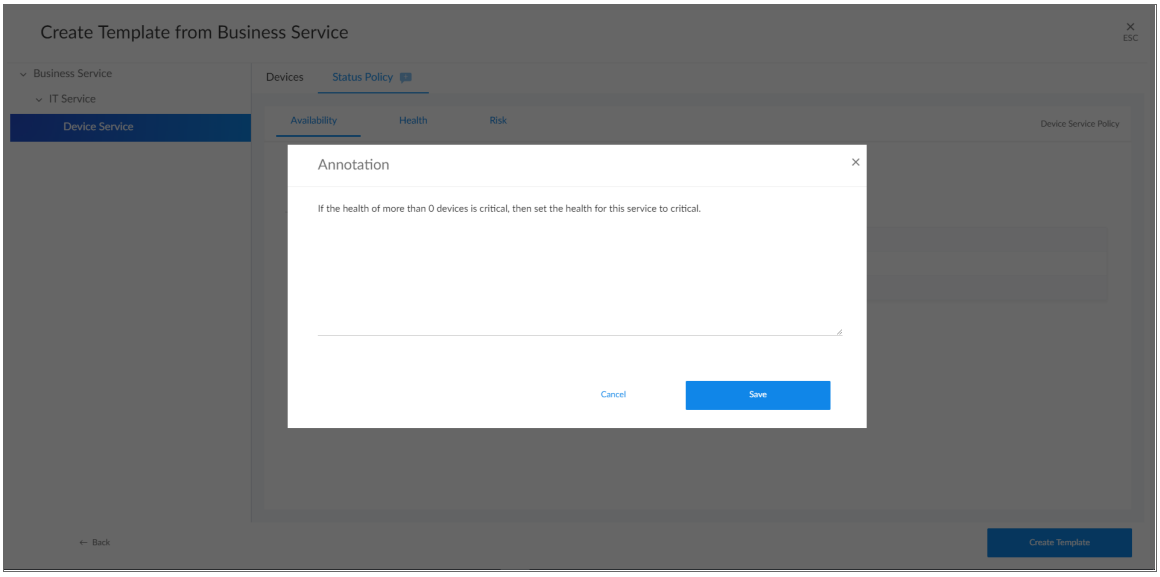

7. To leave an annotation for a service policy or rule, click the annotation icon ( ) next to the rule or tab. Type your annotation text in the **Annotation** window and click **[Save]**. The annotation icon now displays as solid blue, while empty annotation icons contain a plus sign.
8. Click **[Create Template]**. A confirmation window appears stating that you created the template. Click **[Close]**. The template appears on the **Templates** page (Business Services > Templates).

NOTE: To delete one or more service templates, select the check boxes of the templates you want to delete from the **Service Templates** page and then click **[Delete Templates]**. You can delete a single service template by clicking the **Actions** button ( ) for that template and then selecting *Delete*. To select all visible service templates, click the checkbox at the top of the list on the **Templates** page. Clear the checkbox at the top of the list to unselect all of the service templates.

Creating a Service From a Template

To create a service from a template:

1. Go to the **Templates** page (Business Services > Templates) and click the **[Actions]** button (⋮) for the template you want to use, then select *Create Service*. The **Create Service from Template** window appears.

> TIP: You can also go to the **Business Services** page, click the down arrow on the **[Create Service]** button, and select *Create Service from Template*. When you do so, a list of templates you can use to create a service displays. Select the template that you want to use and then click **[Next]**.

2. On the **Create Service from Template** page, complete the following fields:
   - ***Service Name***. Type a name for the new service.
   - ***Description***. Type a description of the new service. (Optional.)
   - ***What organization manages this service?*** Select the organization that will manage the new service from the drop-down field.
3. Click **[Next]**.
4. To edit the names of the services in the hierarchy at the left, click the service name and update the name. Updating the service names is recommended if you are creating the new service on the same system from which the template was created.
5. Any annotations for a device service that were added when the template was created will be present, and you can edit them and add new annotations.
6. On the **[Service Policy]** tab, you can view the rules for ***Health***, ***Availability***, and ***Risk*** for a device service in the template.
7. Click the **[Create Service from Template]** button to save your service. A confirmation window appears.
8. Click the **[Close]** button. The new services appear on the **Business Services** page.

Exporting a Service Template

If you want to use a business service template on another Skylar One system, you can package that template into a PowerPack and then export it to the other system.

To package and export a service template:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks.
2. Click the **[Actions]** button and select *Create New PowerPack*. The **PowerPack Properties** page appears.
3. On the **PowerPack Properties** page , type a name for the PowerPack in the ***Name*** field and click **[Save]**.

> **NOTE:** You can also complete the other fields on the **PowerPack Properties** page. For more information, see the section on "Creating a PowerPack" in the ***PowerPacks*** manual.

4. On the left side of the **PowerPack Properties** page, click **[AP Content Objects]**. Your template appears in the **Available AP Content Objects** pane.

> **TIP:** You can use the filter fields at the top of the columns in the **Available AP Content Objects** pane to narrow down the list of potential templates.

5. Click the lightning bolt icon ( ) next to the template you want to add to the PowerPack. The template moves up to the **Embedded AP Content Objects** pane.
6. On the left side of the **PowerPack Properties** page, click **[Build/Export]**. The **Compiled PowerPacks** page appears.
7. Click the *Create a new build* link at the top of the page.
8. In the **Configure New Export File** window, select *Administrative (including export & license)* from the ***Embedded license key*** drop-down list, and then click **[Build]**.
9. When the PowerPack finishes building, you can click the download icon ( ) to download the build and use that file to upload the template to a new Skylar One system by importing and installing the PowerPack on that system. For instructions on how to import and install a PowerPack, see the section on "Installing a PowerPack" in the ***PowerPacks*** manual. After you install the PowerPack, you can access the template on the **Templates** page (Business Services > Templates).

> **NOTE:** Service templates include every service policy assigned to the services in the template. Templates can back up a service hierarchy and its status policies at the current state. They also allow you to migrate service hierarchies between Skylar One systems.

Default Service Policy Settings

The following sections describe how the three default service policies calculate health, availability, and risk:

Device Service Default Policy

Health: Based upon the worst device severity, then uses the following settings:

- Critical = 0-20
- Major = 21-40
- Minor = 41-60
- Notice = 61-80
- Healthy = 81-100

Availability: Maximum available: if one device is available, then all are available

Risk: Based upon the worst device severity, then uses the following percentages:

- Healthy= 0-20%
- Notice = 21-40%
- Minor = 41-60%
- Major = 61-80%
- Critical = 81-100%

IT Service Default Policy

Health: Average health value of all services

Availability: Maximum available: If one service is available, then all are available

Risk: Maximum risk value of any service

Business Service Default Policy

Health: Average health value of all services

Availability: Maximum available: If one service is available, then all are available

Risk: Maximum risk value of any service

Custom Service Model Default Policy

Health: Average health value of all services

Availability: Maximum available: If one service is available, then all are available

Risk: Maximum risk value of any service

> **NOTE:** Unlike IT services and business services, which use "0" as the lowest possible health and risk values in their default service policies, custom service models use "10" as the lowest health and risk values in their default service policies.

Assigning an Icon to a Service

To assign an icon to a service:

1. On the **Business Services** page, locate the service to which you want to add an icon.
2. Click the **[Actions]** button (⋮) for that service and select *Assign Icon*. The **Select an Icon** window appears:

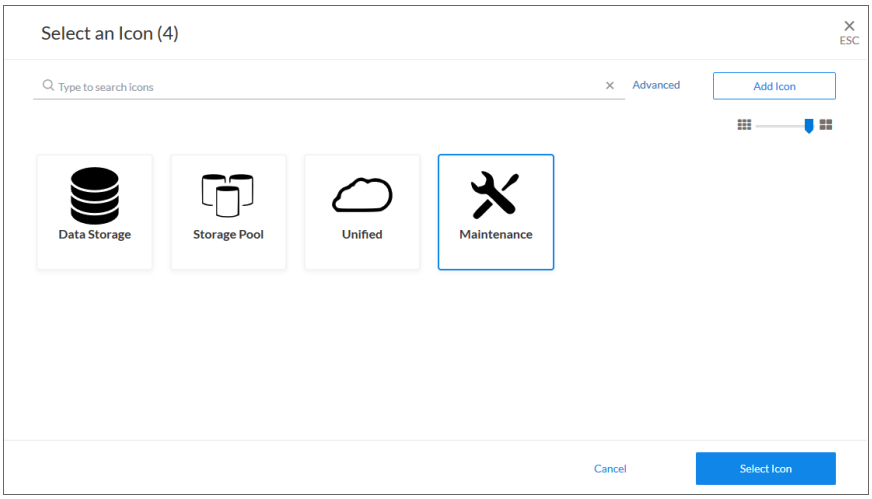

3. To use an existing icon, select that icon from the list of icons and click the **[Select Icon]** button.

> **TIP:** If an icon includes a tag, you can search for that icon by typing some or all of the tag text in the ***Search*** field.

4. To upload an icon from your local drive, make sure that the image file meets the following criteria:
    - The image file should be in .SVG format.
    - The file should not be larger than 40 KB.
    - The file should not be animated.
    - The file should not contain bitmaps

5. To start the upload process, click the **[Add Icon]** button. The **Add an Icon** window appears:

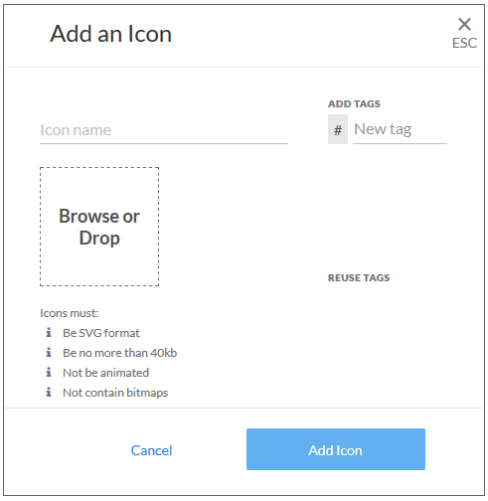

6. In the ***Icon name*** field, type a name for the icon you want to upload.
7. In the ***Add Tags*** field, type a short descriptor for the icon, without spaces. You can use this tag for searching later.
8. You can click the ***Browse or Drop*** area to browse for and select the icon, or you can drag and drop the icon file onto the **Add an Icon** window.
9. Click the **[Add Icon]** button. The icon is added to the **Select an Icon** window.
10. Click the **[Select Icon]** button to add the icon to the service.

Exporting Service Data with the ScienceLogic API

By navigating to the **GQL Browser**, you can export business service data with the ScienceLogic API. The **GQL Browser** is a user interface for interactively exploring the capabilities of, and executing queries against, the Skylar One GraphQL API.

To export service data using the **GQL Browser**:

1. Go to the **GQL Browser** (System > Tools > GQL Browser).
2. In Skylar One, make a note of the URL that displays for the service you want to export. For example, if you have a service named "East Coast Tech," and its URL in Skylar One is **http://sl1.sciencelogic.com/inventory/services/cjumt2se20p3izg6lmiqool5b/overview**. Make a note of the unique value between **/services** and **/overview**. In this example, the value you need is ***cjumt2se20p3izg6lmiqool5b***.
3. In the **GQL Browser**, create a *harProvider* query for the service you want to export, using the following format:

```
query {harProvider (id:"<Service_URI>") { name} }
```

where `<Service_URI>` is the value found in the URL for the Service you want to export.

4. Click the **[Execute Query]** (Play) button to tell the **GQL Browser** to send the query to the GraphQL server and get the results.
5. To export additional data, use the filter-while-you-type capabilities of the **GQL Browser** to gather other information, such as the collection timestamp, health, availability, and risk.
6. After you finish updating your query, click the **[Execute Query]** button.
7.

> **TIP:** For more information, ask your ScienceLogic contact for access to the ScienceLogic GraphQL documentation.

Chapter 2

Using the Service Investigator

Overview

This chapter describes how to use the **Service Investigator** page for a particular business service, IT service, device service, or custom service model.

This chapter covers the following topics:

Viewing the Service Investigator 45

The Sections on the Service Investigator Page 45

Using the Service Analysis Feature 59

Configuring and Enabling the Changes Tab 60

Viewing the Service Investigator

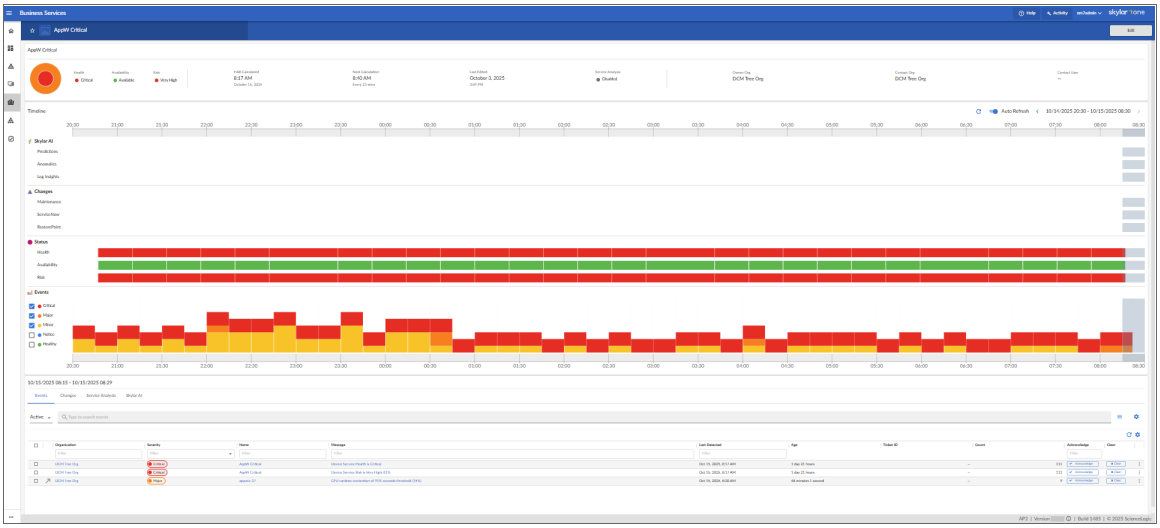

The Sections on the Service Investigator Page

The **Service Investigator** page contains the following sections:

- *Overview*. A summary panel at the top of the page that displays organization and system information such as ***Contact Organization***, ***Visible Organization***, and ***Owner***. This information bar also displays a preview of a sunburst chart, which you can click to see a more detailed breakdown of the **Health**, **Availability**, and **Risk** statuses of your devices.
- *Timeline*. A panel that displays swim lanes and bar graphs to show **Historical**, **Change**, **Health**, **Availability**, **Risk**, and **Skylar AI** events. Swim lanes are visual flowcharts that show a process from start to finish for an event.
- *Events* . An interactive pane that displays the ***Events***, ***Changes***, ***Service Analysis***, and ***Skylar Advisor*** related to the service you have selected.

Each of these sections is described in the following sections.

The Overview Panel

The overview panel on the **Service Investigator** page displays the following information:

- The current ***Health***, ***Availability***, and ***Risk*** values for the service.
- ***HAR Calculated***. The date and timestamp of the most recent system refresh.
- ***Next Calculation***. The timestamp of the next system refresh and its frequency.

- ***Last Edited***. The date and timestamp of the most recent change made to this service and the username of the last user to edit this service.
- ***Service Analysis***. Indicates whether Service Analysis is enabled or disabled for the service.
- ***Owner Org***. The organization that owns the service.
- ***Contact Org***. The organization that should be contacted with any questions about the service.
- ***Contact User***. The user who should be contacted with any questions about the service.
- ***Location***. The geographic location aligned to the service.

> **NOTE:** The ***Location*** field is available only in AP2 Quesito and later releases. If you are running an earlier version of AP2, this feature is not available.

You can click the sunburst chart to see a more detailed breakdown of the ***Health***, ***Availability***, and ***Risk*** statuses of your services. Doing so opens a larger modal that provides different ways to view and filter the statuses of your devices on this particular service.

This diagram window consists of the following viewing options for your services:

- *Sunburst*
- *Map*

Each of these viewing options is described in the following sections.

Sunburst View

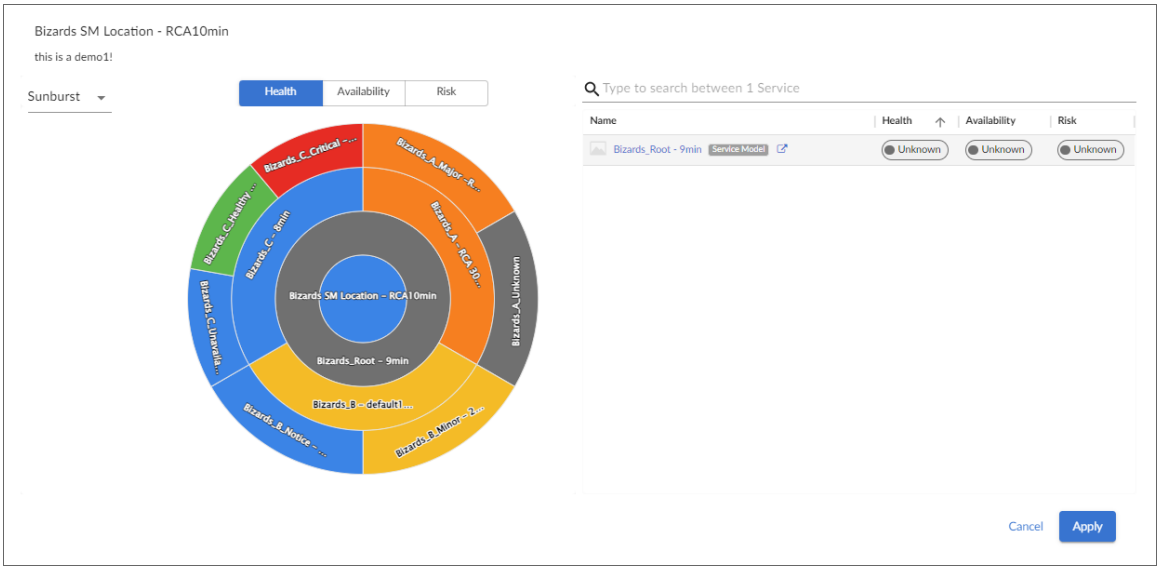

The overview summary panel displays either a *Sunburst* chart view or a *Map* view of your services. Use the drop-down menu in the top left corner of the window to select which view you want to use.

When you select the *Sunburst* view:

- The sunburst chart displays the current ***Health***, ***Availability***, and ***Risk*** values for the services, as well as for any constituent services or device services that belong to that top-level service. For device services, the sunburst chart includes the device name and health values for any devices that belong to the service. Additionally, this chart indicates the maximum number of constituent services or devices that are used for computing health, availability, and risk.
- The right panel includes a list of constituent services or devices. Each service in this panel includes icons that represent that service's ***Health***, ***Availability***, and ***Risk*** metrics; devices include icons that represent each device's health value. The right panel also includes a search bar at the top of the panel that enables you to search for specific constituent services or devices.

In the sunburst chart view, the center circle represents the selected service. The selected service drives the context for the page title and **Info** drawer, as well as all the other panels and widgets in the **Overview** section. This means that the right panel and other elements on the page all reflect the metrics for the service in the center circle of the sunburst.

You can navigate through services on the widget in the following ways:

- On the sunburst chart, you can click any constituent IT services or device services in the sunburst to select that service. To return to the parent IT service or business service, click the center circle or click the breadcrumb links that appear in the top-left corner of the window.
- In the right panel, you can click the service name of any of the constituent IT services or device services to select that service. To return to the parent IT service or business service, click the breadcrumb links that appear in the top-left corner of the window.

By default, the sunburst displays the health value for the selected service and its constituent services or devices. To view the current availability or risk value for the selected service, click either the **[Availability]** or **[Risk]** tab above the sunburst chart.

Map View

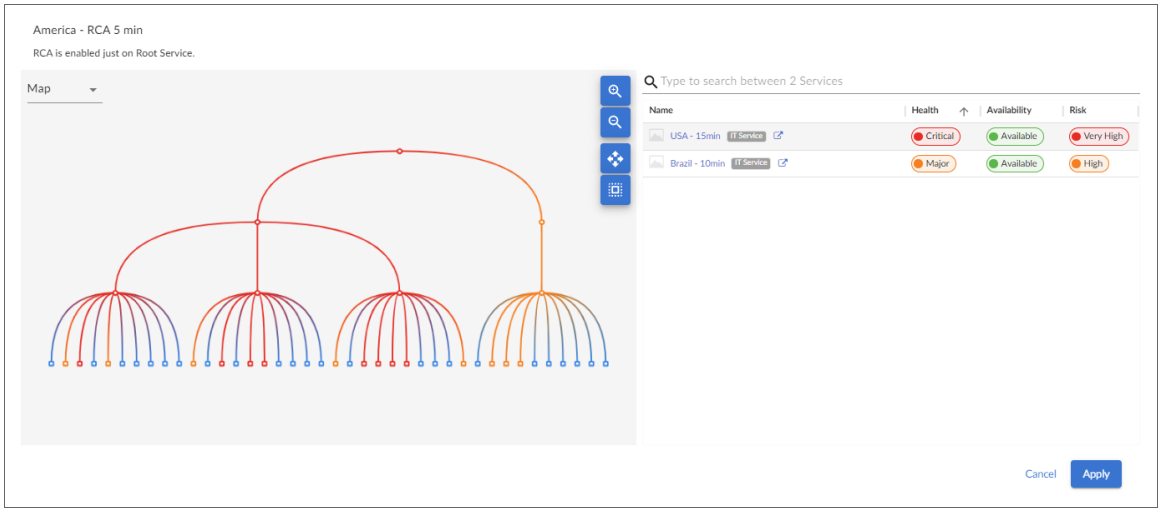

The overview summary panel displays either a *Sunburst* chart view or a *Map* view of your services. Use the drop-down menu in the top left corner of the window to select which view you want to use.

When you select the *Map* view:

- The map view displays a map of the service and any constituent services and devices that belong to that top-level service.
- The right panel includes a list of constituent services or devices. Each service in this panel includes icons that represent that service's ***Health***, ***Availability***, and ***Risk*** metrics; devices include icons that represent each device's health value. The right panel also includes a search bar at the top of the panel that enables you to search for specific constituent services or devices.

In the map view, you can click the top-level service or any of its constituent services or devices. The selected service drives the context for the page title and Info drawer, as well as all the other panels and widgets on the overview summary pane. This means that the right panel and other elements on the page all reflect the metrics for the service that you have clicked in the map view.

In the map view, use the following buttons to manipulate the map in the left pane:

 : Zoom in on the map.

 : Zoom out on the map.

 : Fit all elements of a map into the viewing pane.

 : Center all selected elements of a map in the viewing pane.

The viewing pane displays the following two types of graphical elements:

1. ***Nodes*** that represent devices, topology elements, and business services defined in Skylar One. The shape of the node represents its type: *Services*, such as business services, IT services, or device services, are represented by hexagons, while devices are represented by squares. The color of the outline specifies the current state of the node.
2. ***Edges*** are lines that represent the relationships and hierarchies between nodes.

TIP: When you hover over a node, a pop-up **Properties** pane appears with the metadata for that node. Click the **[Go to service]** or **[Go to device]** link at the top of the pane to open the **Investigator** page for that service or device in a new browser window.

The Timeline Panel

The **Timeline** panel displays an interactive graphic that combines swim lanes and bar graphs to visualize the health, availability, and risk of your events. Swim lanes act as visual flowcharts, showing the progression of an event from start to finish. You can select any time range on the timeline to view **Changes**, **Health**, **Availability**, and **Risk** details specific to your device or service for that period. Timeline data is presented in 12-hour segments, allowing you to navigate backward and review up to seven days of historical activity.

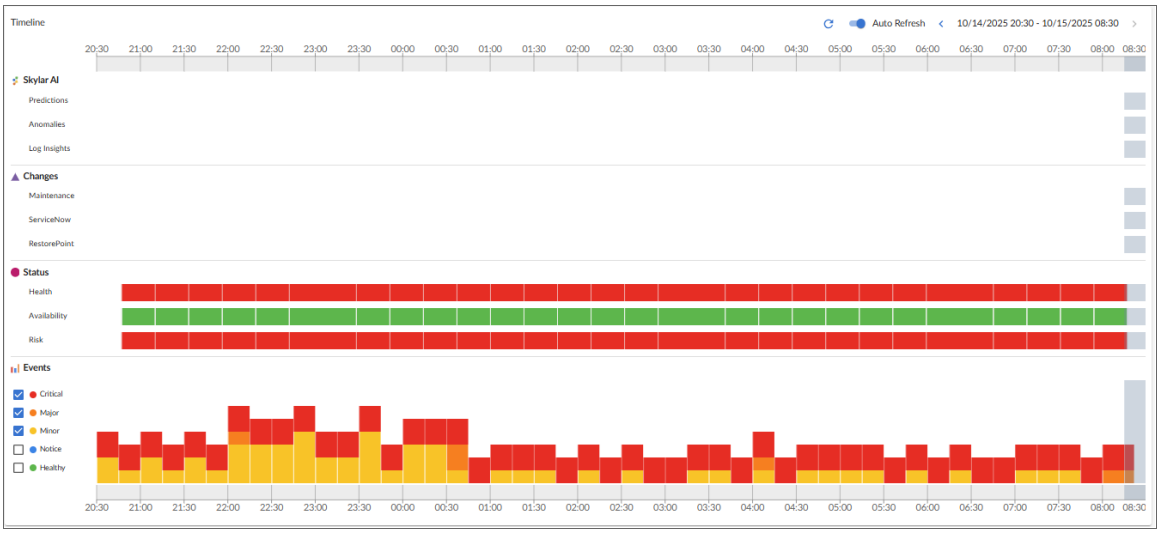

The **Timeline** panel consists of the following sections:

- *Skylar AI*
- *Changes*
- *Status*
- *Events*

NOTE: The swim lanes under the **Skylar AI** and **Changes** sections automatically expand or collapse based on configuration, subscription, and event availability. Each section collapses when it is not applicable or when no relevant events exist in the selected time frame, and expands automatically when events are present. The **Skylar AI** and **Changes** sections can also be manually collapsed or expanded.

Skylar AI Swim Lanes

You can view Skylar AI suggestions and alerts in the **Skylar AI** swim lanes of the **Timeline** panel on the **Service Investigator** page. To use this feature, you will need to set up the connection between Skylar AI and Skylar One. By default, suggestions and alerts cover the last 12 hours.

The **Skylar AI** section includes three swim lanes:

- ***Predictions***, which shows predicted future events from start to finish.
- ***Anomalies***, which shows events marked as anomalies from start to finish.
- ***Log Insights***, which shows Skylar AI suggestions and alerts. The alerts are either *Confirmed* or *Suggested*.

Confirmed Alerts

Confirmed alerts in the ***Log Insights*** swim lane of the **Skylar AI** section represent log events that Skylar AI has collected for your service or device. In Skylar AI, they appear as categorized events under

Accepted/Custom. These logs typically contain metadata such as title, description, and so forth and use the **Skylar AI** to analyze and collect event logs that show abnormalities.

Suggested Alerts

Suggested alerts in the ***Log Insights*** swim lane of the **Skylar AI** section contain metadata such as title, description, and so forth; a root cause report, which is a set of correlated log lines that help to explain a problem; and a suggested alert rule consisting of one or two log event types that form the signature for this type of alert.

As logs are ingested, Skylar AI analyzes your system or device logs for event patterns such as rare events and error events that are abnormally correlated across all log streams. When it detects one of these "abnormal" clusters, it generates a suggested alert, which allows you to choose if you want these events to be recorded as event logs in the future.

You can choose to either accept or reject a suggested alert.

- If you accept a suggested alert, you can edit the metadata and alert rule. You can also decide on the action to take if the same kind of alert occurs again, such as sending a notification to Slack, email, or another communications platform.
- If you reject a suggested alert, the same kind of alert is not recorded in the future, and you will not be offered the option to either accept or reject it.

Changes Swim Lanes

The **[Changes]** section on the **Timeline** panel is available to customers who have purchased Configuration and Change Management as part of their Skylar One Standard or Advanced subscription. This section displays a list of events that are created when Skylar Automation pulls change data from ServiceNow or Skylar Compliance, including both active and cleared change events. By default, the events cover the last 12 hours.

This **[Changes]** section contains the following swim lanes:

- *Maintenance*
- *ServiceNow*
- *Skylar Compliance*

> **NOTE:** For swim lanes under the **Skylar AI** and **Changes** section, the total number of events is displayed in parentheses next to the swim lane title.

Maintenance

The ***Maintenance*** swim lane indicates whether a device is in maintenance mode. Devices in maintenance have start and end times, which are determined by a start event and an end event for each change data. The device's ***Collection State*** field will change from *Active* to *Scheduled Maintenance*, and then back to *Active* when maintenance reaches its end time.

You can see a device's collection state in the ***Collection State*** column on the **Events** page. To see a list of all devices in maintenance mode and their start times, end times, and duration, go to the **Schedules** tab of the **Device Investigator** page.

NOTE: For more information, see the section on *Viewing the Schedule Manager*.

ServiceNow

The ***ServiceNow*** swim lane shows a visual representation of all Skylar One events that are synced with ServiceNow incidents in a specific time range.

NOTE: You can permanently enable or disable ***ServiceNow*** swim lane diagrams on the **Timeline** panel through the nextui.conf file, or temporarily enable or disable them through GraphQL mutations. To do so, follow the instructions outlined in the *Configuring and Enabling the Changes Tab* section.

NOTE: For more information on how to monitor a ServiceNow instance with Skylar One events, see the section on *Using the ServiceNow Base Pack PowerPack*.

Skylar Compliance

The ***Skylar Compliance*** swim lane shows a visual representation of all Skylar One events on devices that are synced between Skylar One and Skylar Compliance. The events are created when Skylar Automation pulls change data from Skylar Compliance.

NOTE: You can permanently enable or disable ***Skylar Compliance*** swim lane diagrams on the **Timeline** panel through the **nextui.conf** file, or temporarily enable or disable them through GraphQL mutations. To do so, follow the instructions outlined in the *Configuring and Enabling the Changes Tab* section.

NOTE: For more information, see the section on the *Skylar Compliance Syncpack*.

Status Swim Lanes

The **[Status]** section provides an overview of your service or device health, availability, and risk. Selecting a time range in the **Timeline** panel opens a pop-over modal with ***Health***, ***Availability***, and ***Risk*** details for the selected device or service. By default, the status shows the health, availability, and risk statuses of events from the past 12 hours, which is the maximum supported time range.

- *Health*
- *Availability*
- *Risk*

Health

The ***Health*** swim lane displays device or service health over a period of time.

This swim lane indicates the current status of a device service–for example, the rate of processing or throughput for the devices in the service. In the case of Skylar One Database Servers, the "Rows Behind" value can provide a good measure of how effectively the Database Server is processing data from Data Collectors. ***Health*** is represented by a color-coded "severity" icon that corresponds to a numerical value between 0 and 100. For example, the ***Health*** value could indicate when a device is intermittently unavailable because of a power problem, thereby falling below the required level of performance. ***Health*** uses the following icons:

Availability

The ***Availability*** swim lane shows the availability status of your device or service over a period of time.

The availability of a device service is derived from the availability rules. This might or might not be linked to device availability. A service or device is considered unavailable if Skylar One is not able to collect data from the device or service, or if a device is usable or not usable. A value of *0* means a device or service is unavailable, and a value of *1* means a device is available. Availability uses the following icons:

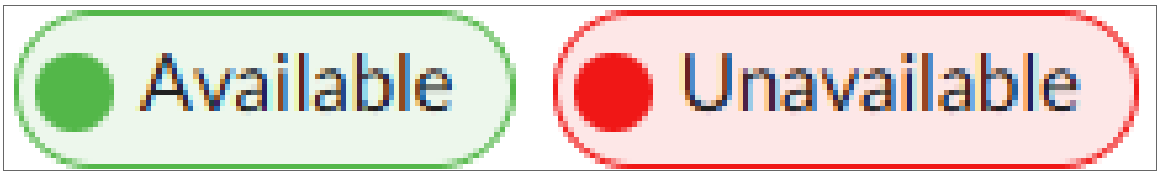

Risk

The ***Risk*** swim lane displays the risk status of your device or service over a period of time.

This swim lane measures the risk status of your device or service using a percentage value between 0 and 100 that indicates how close a service is to being in an undesirable state. Use risk for data that is known to cause issues if left unchecked, such as critical events, swap usage, or low database logging space. The safest possible risk value is 0%, while the worst risk value is 100%.

These values are computed in this order because Skylar One uses ***Availability*** values to compute ***Health***, and then uses both ***Availability*** and ***Health*** values to compute ***Risk***.

Events Swim Lanes

The **[Events]** section contains five swim lanes that display the number of unique events of each severity level that impacted your service or device over a period of time. By default, the events cover the last 12 hours.

The Events Panel

The **Events** panel contains the following tabs:

- *Events*
- *Changes*
- *Service Analysis*
- *Skylar Advisor*

Events Tab

The **[Events]** tab displays events associated with the currently viewed service, as well as events associated with that service's current services and devices. The list of events shown in the table is defined by the time range selected in the **Timeline** panel.

The **[Events]** tab uses much of the same functionality as the **Events** page, but the results shown here are limited to the current service context and filtered based on the selected timeline range.

> **NOTE:** By default, the **[Events]** tab summarizes events by severity and reflects events that meet the current table criteria and timeline context. Whether masked events are shown or hidden is determined by the user's event table preferences. Based on those preferences, masked events may be included in or excluded from the table.

When opened, the **[Events]** tab lists events that qualify under the current view (*Active* or *Cleared*) and the selected timeline range. Inclusion in the table is not limited to events that are currently impacting the service. Instead, events are included based on their current state (*Active* or *Cleared*) and how their activity relates to the selected time range, as described below.

You can perform the following actions from the **[Events]** tab:

- Use the drop-down menu to choose which type of change events display in the widget: *Active* or *Cleared*.
    - *Active* . Shows events that are currently active (not cleared) and meet specific criteria relative to the selected time range, as defined by one of the following criteria:
        - Became active in one of the selected time ranges
        - Was already active and saw a recurrence that updated the existing event record in one of the selected time ranges
        - Became active prior to the selected time range and has occurred again and updated the existing event record after the selected time range, but did not become active or see recurrence within the selected time range

        > **NOTE:** The *Active* view lists only events that are still active. An event that becomes active and then clears within the same selected time range does not appear in the *Active* view. Although such an event was active during part of the selected range, it is excluded because it is no longer active. Instead, it appears only in the *Cleared* table, where events are listed based on when they transition from *Active* to *Cleared*.

        For example, consider a selected interval from 12:15 to 12:30. An event that becomes active at 12:20 appears in the *Active* view, unless it clears within that same interval. An event that was already active before 12:15 and has a recurrence at 12:22 that updates the existing event record also appears in the *Active* view. An event that became active prior to 12:15, does not recur between 12:15 and 12:30, but is updated again at 12:45 may still appear in the *Active* view, because the system can determine that the event remained active across the selected interval. However, if an event becomes active at 12:18 and then clears at 12:27, it does not appear in the *Active* view and instead appears only in the *Cleared* view, even though it was active during part of the selected interval.
    - *Cleared*. Shows events that transitioned from *Active* to *Cleared* within the selected time range.

    > **NOTE:** The *Cleared* view is specifically tied to the transition from *Active* to *Cleared* within the selected time range, not to when the event originally became active.

- Use the search field to search for specific events.
- For events that are aligned to devices, click the arrow icon (↗) next to the event to open the **Device Summary** window, which displays the following panes:
    - ***Tools***. A set of network diagnostic tools or user-initiated actions that you can run on the device associated with the event. Click the search bar to search for a tool or action to run, or click one of the default tools or actions that are available based on the device type and your user permissions.

    - ***Vitals***. A widget displaying the past four hours of CPU and memory usage for the device related to the event. You can zoom in on a shorter time frame by clicking and dragging, and you can go back to the original timespan by clicking the **[Reset zoom]** button.
    - ***Logs***. A list of the log entries from the device's log file, sorted from newest to oldest by default.
- View the **Organizational Summary** page for the organization aligned with the event by clicking the link in the ***Organization*** column.
- View the **Service Investigator** or **Device Investigator** page for the service or device aligned with the event by clicking the link in the ***Name*** column.
- View the **Event Investigator** page for the event by clicking the link in the ***Message*** column.
- View or edit event notes by clicking the **Note** icon ( ) in the ***Event Note*** column or by clicking the actions icon ( ) and selecting *Edit Event Note*. Event notes contain event definitions, probable causes, and resolutions for the event, along with a text field where you can add more information about the event or the service or device you are monitoring.
- View more information about masked events by clicking the masked events icon ( ) in the ***Masked Events*** column. Masked events are related events that occur in quick succession on a single device or service that are rolled up and posted together under one event description, with only the highest severity event displayed.

> **NOTE:** For more information on masked events, see *Viewing Events*.

- Acknowledge the event by clicking the **[Acknowledge]** button. When you acknowledge an event, you let other users know that you are aware of that event, and you are working on a response.
- Clear an event by clicking the **[Clear]** button. When you clear an event, you let other users know that the event has been addressed.
- Create a ticket from the event.
- View the event policy.
- View a log of automations that have occurred for the event by clicking the actions icon ( ) and selecting *View Automation Actions*.
- Select multiple events for action using the check boxes next to the events.

> **NOTE:** For more information about events, see the section on *Events*.

Changes Tab

The **[Changes]** tab displays the number of active change events that are impacting the service. Events on this tab automatically clear after 30 minutes.

NOTE: The **[Changes]** tab is available if you have purchased Configuration and Change Management as a part of your Skylar One Standard or Advanced subscription. This tab displays a list of events that are created when PowerFlow pulls change data from ServiceNow or Skylar Compliance, including both active and cleared change events. For more information on how to configure, enable, or disable the **[Changes]** tab on the **Events** pane, see the section on *Configuring and Enabling the Changes Tab*.

You can perform the following actions on the **[Changes]** tab:

- Use the drop-down menu to choose which type of change events display in the widget: *Active Events* or *Cleared Events*. *Active Events* are events that are currently active and whose most recent detection occurred within the selected time range. *Cleared Events* are events that were cleared within the selected time range, up to a maximum of 12 hours.
- Filter and search for events by date: 5, 7, 14, 30, or more than 30 days.
- Use the ***Search*** field to search for specific change events.
- For active events that are aligned to devices, click the arrow icon ( ) next to the event to display the following information:
    - ***Tools***. A set of network diagnostic tools or user-initiated actions that you can run on the device associated with the event. Click the search bar to search for a tool or action to run, or click one of the default tools or actions that are available based on the device type and your user permissions.
    - ***Vitals***. A widget displaying the past 24 hours of CPU and memory usage for the device related to the event. You can zoom in on a shorter time frame by clicking and dragging, and you can go back to the original timespan by clicking the **[Reset zoom]** button.
    - ***Logs***. A list of the log entries from the device's log file, sorted from newest to oldest by default.
- View the **Organizational Summary** page for the organization aligned with an active event by clicking the link in the ***Organization*** column.
- View the **Service Investigator** or **Device Investigator** page for the service or device aligned with an active event by clicking the link in the ***Name*** column.
- View the **Event Investigator** page for an active event by clicking the link in the ***Message*** column or the ***Event ID*** column.
- For ServiceNow integrations, view the ServiceNow ticket associated with an event by clicking the link in the ***External Ticket*** column.
- Acknowledge an active event by clicking the **[Acknowledge]** button. When you acknowledge an event, you let other users know that you are aware of that event, and you are working on a response.
- Clear an active event by clicking the **[Clear]** button. When you clear an event, you let other users know that the event has been addressed.
- Create a ticket from an active event.
- Align an event to an existing ticket.

- View the event policy for an active event.
- Select multiple active events for action using the check boxes next to the events.

Service Analysis Tab

The **[Service Analysis]** tab displays the Service Analysis of a device or service and shows what is causing either one to be unhealthy based on the **Service Policy**.

> **NOTE:** For more information about enabling Service Analysis for a service, see the section on *Using the Service Analysis Feature*.

You can use the drop-down menu to choose whether you want to view the service's *Health*, *Availability*, or *Risk* analysis.

> **NOTE:** If you select *Active*, Skylar One displays events that are currently active and whose most recent detection occurred within the selected time range. If you select *Cleared*, Skylar One displays events that were cleared within the selected time range. The selected time range spans a maximum of 12 hours.

The following columns appear on the **[Service Analysis]** tab:

- ***Service/Device Name***. The name of the service or device that contributed to the health, availability, or risk status for the selected time period.
- ***Current State***. The current health, availability, or risk status for the service or device.
- ***Condition***. The equation that is used to determine the health, availability, or risk status for the service or device.
- ***Current Value***. The current health, availability, or risk value for the service or device, as determined by the value of the equation used in the ***Condition*** column.
- ***Historical Value***. The historical health, availability, or risk value for the service or device for the selected time period, as determined by the value of the equation used in the ***Condition*** column.
- ***Timestamp***. The date and time the service analysis was collected from the service or device.

> **TIP:** You can click any of the column heading labels to sort the **[Service Analysis]** tab by the values in that column.

Skylar Advisor Tab

The **[Skylar Advisor]** tab displays a list of Skylar Automated RCA events.

You can view Skylar Automated RCA suggestions and alerts in the **[Skylar Advisor]** tab of the **Events** pane on the **Service Investigator** page. To use this feature, you will need to set up the connection

between Skylar Automated RCA and Skylar One. For more information, see *Configuring the Skylar Automated RCA Connector for Skylar One*.

You can perform the following actions on the **[Skylar Advisor]** tab:

- Use the drop-down menu to choose which type of change events display in the tab: *Active Events* or *Cleared Events*.

> **NOTE:** If you select *Active*, Skylar One displays Skylar AI sourced events that are currently active and whose most recent detection occurred within the selected time range. If you select *Cleared*, Skylar One displays events that were cleared within the selected time range. The selected time range spans a maximum of 12 hours.

- Filter and search for Skylar Automated RCA events by their date: either by 5, 7, 14, 30, or more than 30 days.
- Use the ***Search*** field to search for specific change events.
- For active events that are aligned to devices, click the arrow icon ( ) next to the event to open the **Device Summary** window, which displays the following panes:
  - ***Tools***. A set of network diagnostic tools or user-initiated actions that you can run on the device associated with the event. Click the search bar to search for a tool or action to run, or click one of the default tools or actions that are available based on the device type and your user permissions.
  - ***Vitals***. A widget displaying the past 24 hours of CPU and memory usage for the device related to the event. You can zoom in on a shorter time frame by clicking and dragging, and you can go back to the original timespan by clicking the **[Reset zoom]** button.
  - ***Logs***. A list of the log entries from the device's log file, sorted from newest to oldest by default.
- View the **Organizational Summary** page for the organization aligned with an active Skylar Automated RCA event by clicking the link in the ***Organization*** column.
- View the **Service Investigator** or **Device Investigator** page for the service or device aligned with an active Skylar Automated RCA event by clicking the link in the ***Name*** column.
- View the **Event Investigator** page for an active Skylar Automated RCA event by clicking the link in the ***Message*** column.
- For ServiceNow integrations, view the ServiceNow ticket associated with an active or cleared event by clicking the link in the ***External Ticket*** column.
- Acknowledge an active event by clicking the **[Acknowledge]** button. When you acknowledge a Skylar Automated RCA event, you let other users know that you are aware of that event, and you are working on a response.
- Clear an active Skylar Automated RCA event by clicking the **[Clear]** button. When you clear a Skylar Automated RCA event, you let other users know that the event has been addressed.
- Create a ticket from an active Skylar Automated RCA event.
- View the event policy for an active Skylar Automated RCA event.

- Select multiple active Skylar Automated RCA events for action using the check boxes next to the events.

Using the Service Analysis Feature

Skylar One users can use the ***Service Analysis*** feature to determine what is causing a service to be unhealthy, troubleshoot that service, and refine its policies.

> **NOTE:** When you enable Service Analysis on a business service or IT service, it also implicitly enables Service Analysis on any child IT services or device services.

Enabling Service Analysis

To enable Service Analysis:

1. Click the **Business Services** icon ( ) to go to the **Business Services** page.
2. Click the ***Name*** of an existing service. The **Service Investigator** page for that service displays.
3. On the **Service Investigator** page, click **[Edit]**.

4. Select one of the following options from the ***Service Analysis*** drop-down field:
    - ***Disabled***. The Service Analysis feature is disabled.
    - ***Enabled (contributors only)***. The Service Analysis feature is continuously enabled only for contributing rules and devices. When you select this option, a full analysis is generated and saved in the time series chart, but it excludes results from non-contributing rules and devices.
    - ***Enabled (next run only)***. The Service Analysis feature is enabled only for the next data collection.
    - ***Enabled***. The Service Analysis feature is continuously enabled for all rules and devices. When you select this option, a full analysis is generated and saved in the time series chart, and it includes results from non-contributing rules and devices.

    NOTE: If a service has ***RCA Options*** set to *Enabled* and has a child service removed, Skylar One does not compute the health, availability, and risk values until the Service Topology Engine returns an updated topology, which occurs every 5 minutes by default.

    IMPORTANT: Before deleting child services in a 3-tier hierarchy, check if the parent service has the ***RCA Options*** field set to *Enabled*. If it does, then set the field to *Disabled* before deleting the child services.

    CAUTION: You might experience performance slowdown if Service Analysis is continuously enabled.

5. Click **[Save]**.

Configuring and Enabling the Changes Tab

To use the **[Changes]** tab on the **Events** panel, you must first configure and enable the tab. To do so, you must meet the following prerequisites:

- You are running Skylar One version 12.1.0 or later.
- You have "Business Services Base Pack" PowerPack version 2.2.0 or later installed in Skylar One.
- You are running Skylar Automation Platform version 2.2.2 or later.
- You have the following SyncPacks and PowerPacks installed, depending on your integration:
    - ***For a ServiceNow integration***:
        - "ServiceNow CMDB" SyncPack version 3.2.0 or later installed in Skylar Automation
        - "ServiceNow Change Management" SyncPack version 3.2.1 or later installed in Skylar Automation.

- ***For a Skylar Compliance integration***:
  - "Skylar Compliance" (formerly "Restorepoint") SyncPack version 1.2.0 or later installed in Skylar Automation.
  - "Skylar Compliance" (formerly "Restorepoint") PowerPack version 102 or later installed in Skylar One.
  - "Skylar Compliance Automation" (formerly "Restorepoint Automation") PowerPack version 102 or later installed in Skylar One.

If you have met those prerequisites, you must then do the following to configure the **[Changes]** tab:

1. In Skylar One, *create a SOAP/XML credential* to connect with Skylar Automation. Make note of its credential ID. This number is found at the top of the **Edit SOAP/XML Credential** modal or in the ***ID*** column on the **Credentials** page (Manage > Credentials) or **Credential Management** page (System > Manage > Credentials).
2. ***For a ServiceNow integration***:
   - In Skylar Automation, *sync Skylar One devices with ServiceNow*. Make note of the ***Configuration*** field value in the "Sync Devices from Skylar One to ServiceNow" application.
   - In Skylar One, open the "ServiceNow: Send Change Request Event to PowerFlow" run book action and *edit the input parameters* to include the credential ID and the ***Configuration*** field values that you previously noted.
3. ***For a Skylar Compliance integration***, follow the steps in the section on *syncing Skylar One devices with Skylar Compliance* in Skylar Automation.
4. Finally, do one of the following:
   - Permanently enable the **[Changes]** tab by *editing the NextUI configuration file* on your Skylar One system.
   - Temporarily enable the **[Changes]** tab by *running a GraphQL mutation* on your Skylar One system.

Editing the Run Book Action

To edit the input parameters in the "ServiceNow: Send Change Request Event to PowerFlow" run book action:

1. Go to the **PowerPack Manager** page (System > Manage > PowerPacks).
2. Locate the "Business Services Base Pack" PowerPack and click its wrench icon (🔧). The **Editing PowerPack** modal appears.
3. In the **Editing PowerPack** modal, click ***Run Book Actions*** in the left Navbar. The **Embedded Run Book Actions** page appears in the modal.
4. Click the wrench icon (🔧) for the "ServiceNow: Send Change Request Event to PowerFlow" run book action. The **Policy Editor** modal appears.
5. In the **Policy Editor** modal, make the following edits to the ***Input Parameters*** field:

- Replace `<sl1 credential id for powerflow>` with the credential ID of the *SOAP/XML credential you created for Skylar Automation*.
- Replace `<pf config id>` with the ***Configuration*** field value from the *"Sync Devices from Skylar One to ServiceNow" application in Skylar Automation*.

6. Click **[Save]**, then exit the **Policy Editor** modal.
7. Exit the **Editing PowerPack** modal.

Syncing Skylar One Devices with Skylar Compliance

To sync Skylar One devices with Skylar Compliance:

1. Follow the steps in the section *Configuring the "Skylar Compliance: Sync Devices" application*.
2. In Skylar Automation, open the **Configuration** pane for the "Skylar Compliance: Sync Devices" application and select *Enable* for the ***restorepoint_config*** field to allow device change detection.
3. Make a note of the ***restorepoint_id*** value on the **Configuration** pane for the "Skylar Compliance: Sync Devices" application.
4. In Skylar One, go to the **Device Investigator** page for the devices synced from Skylar Compliance, click on the **[Attributes]** tab, and ensure that the same ***restorepoint_id*** value was added to the ***Values*** column.

Permanently Enabling the Changes Tab

To permanently enable the **Changes** widget or tab using the NextUI configuration file, run the following steps on all appliances, including the Administration Portal, the Data Collector, the Database Server, the Data Engine, and the All-In-One Appliance.

To permanently enable the **Changes** widget or tab:

1. Start an SSH session into one of the Skylar One appliances.
2. Using vi or another text editor, edit the `/opt/em7/nextui/nextui.conf` file. To do so, enter the following at the shell prompt:

```
sudo vi /opt/em7/nextui/nextui.conf
```

3. Add the following line at the bottom of the NextUI configuration file:

```
BUSINESS_SERVICES_CHANGE_EVENTS_TAB=enabled
```

4. Save your changes, and then restart the NextUI service by running the following command:

```
sudo systemctl restart nextui
```

5. Repeat steps 1-4 for the remaining Skylar One appliances.

Temporarily Enabling the Changes Tab

To temporarily enable the **Changes** widget or tab using GraphQL:

1. To access the GraphiQL interface, type the URL or IP address for Skylar One in a browser, add `/gql` to the end of the URL or IP address, and press **[Enter]**. The GraphiQL interface appears.
2. In the main query pane, type the following mutation:

```
mutation updateChangeEventsTab {
        updateFeatureToggle(
                id: "system:BUSINESS_SERVICES_CHANGE_EVENTS_TAB"
                value: "enabled"
        ) {
                id
                value
        }
}
```

> **TIP:** Click the **[Prettify]** button to format the mutation and to add syntax highlighting to make the mutation easier to read. Note that the *Prettify* process removes the `query` syntax if only one query is present in the main query pane.

3. Click the **[Execute Query]** (Play) button. The mutation executes, and the results appear in the pane on the right side.

> **NOTE:** If the **Changes** widget or tab does not appear in Skylar One after executing the mutation, refresh the page using the **[F5]** key or by clicking the refresh button in your web browser.

> **NOTE:** For more information about GraphQL, see the GraphQL documentation. For more information about the GraphiQL user interface, see the GraphiQL user interface documentation.

Chapter

3

Resolving Service Issues with Behavioral Correlation

Overview

This chapter describes how to identify and diagnose service issues using behavioral correlation in Skylar One.

This chapter covers the following topics:

Understanding Behavioral Correlation in Business Services *64*

Analyzing a Service's Behavioral Correlation *65*

Understanding Behavioral Correlation in Business Services

Skylar One offers an elevated enterprise visibility experience that allows you to monitor and view the health of your services, from an individual device to a larger business service view. Skylar One has the ability to show a service's behavior and its interactions; in other words, it enables you to analyze a service's ***behavioral correlation***. This behavioral correlation analysis enables you to implement changes with the hopes of reducing complexity and noise.

To determine a business service's health, you might first look at the **NOC Overview** dashboard, where the health of a Business Services's components is displayed.

Using the color-coded health widgets in Skylar One, you can determine which areas of your enterprise are considered healthy and which are not.

Analyzing a Service's Behavioral Correlation

You can determine the behavioral correlation for a business service by further analyzing the health data of specific areas.

To analyze the health of a specific area:

1. On the **NOC Overview** dashboard, click the widget of the service area that you want to further analyze. The **Service Investigator** displays.
2. Identify problem areas by reviewing the relevant widgets on the **Service Investigator**. This page allows you to sift through any problem sources and pinpoint their various severities.
3. By clicking into a widget's source data, you can view the health, availability, risk, events, changes, and Skylar AI data for a specific area. The machine learning capability of Skylar One lets you view predictions, anomalies, and insights for your service's problem source areas.

See *Using the Service Investigator* for more information on this page and its functionality.

Appendix

A

Troubleshooting Business Services

Overview

This chapter covers some of the issues you might encounter while working with services and policies on the **Business Services** page, and how to resolve those issues.

This chapter covers the following topics:

Business Services Have Empty Values 66

Some Services Fail to Generate Health, Availability, or Risk Values 69

All Services Fail to Generate Health, Availability, and Risk Values 74

Device Services Fail to Load After an Upgrade 75

502, 503, or 504 Errors: Health, Availability, and Risk Values are All the Same or are Inaccurate 75

Advanced Troubleshooting 77

Business Services Have Empty Values

All Business Services Have Empty Values

If all of your business services show empty values, ensure that you have given your admin processes adequate time to complete. To populate these values, both the "Business Services: Service Management

Engine" and "Business Services: Service Topology Engine" processes must run once. With default settings, it could take up to 30 minutes to see your first results.

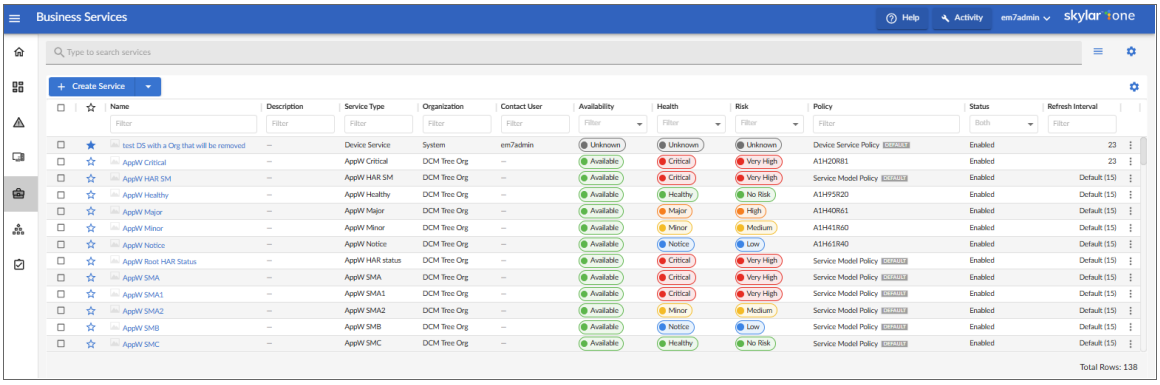

Services are not evaluated if they have an empty filter. For more information on using a filter, see the section on *Creating a Service*. The figure below shows the results of using a filter to find all devices for which the IP address begins with "10.2.":

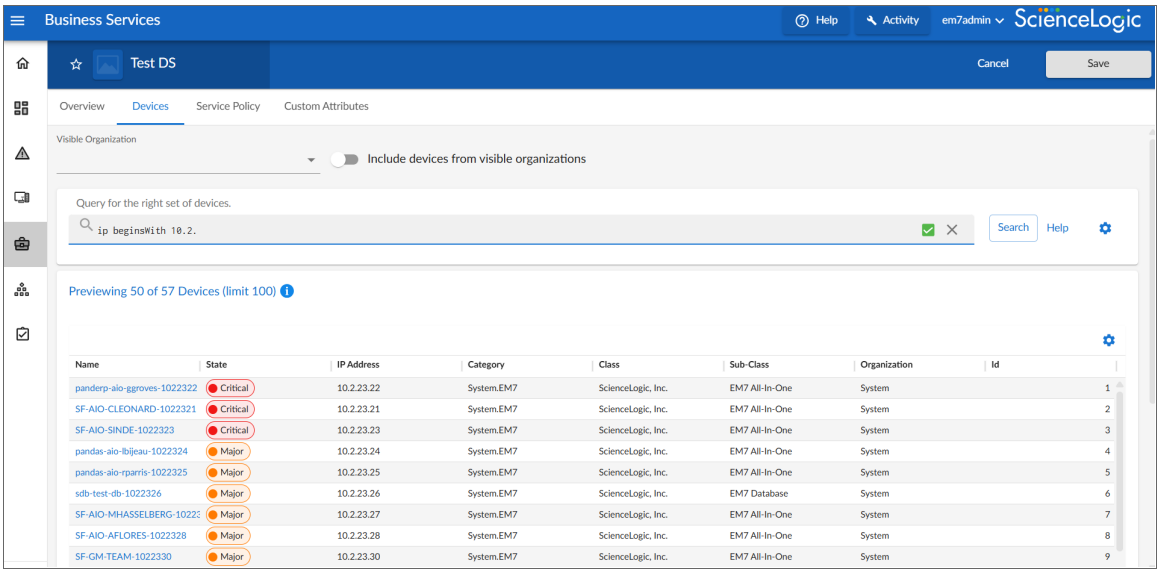

Some Business Services Have Empty Values

If only some of your business services are missing values, troubleshoot using the following procedure:

1. Ensure that your business service has some constituents:
   a. Go to the **Business Services** page ( ).
   b. Select the service that is missing values. The **Service Investigator** appears.
   c. Click **[Edit]**.

d. Click the **[Devices]** or **[Services]** tab and review the devices or services listed. Modify your query as needed.

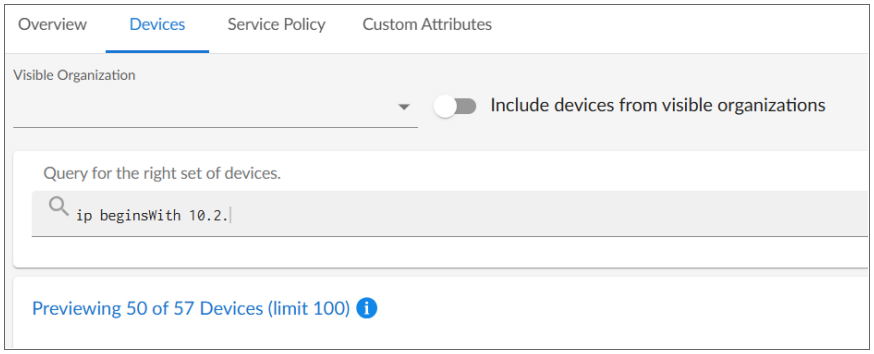

2. Ensure that your service policy results in some constituents. Click on the **[Service Policy]** tab and modify your service policy as needed.
3. Click **[Save]**.

Services Missing Up-to-Date Values

If you have disabled the default administrator account ("em7admin"), you will need to identify another account to use for running business services and run a database query to change the account used for internal communication in Skylar One.

To change the internal account:

1. Go to the **User Accounts** page (Registry > Accounts > User Accounts).
2. Identify the account you want Skylar One to use for internal communication. Take note of the value that appears in the ***User ID*** column for that user account.
3. Update the internal account.
   a. Go to the **Database Tool** page (System > Tools > DB Tool).

> **NOTE:** The **Database Tool** page is available only in versions of Skylar One prior to 12.2.1 and displays only for users that have sufficient permissions to access the page. Alternatively, if you are not a SaaS user, you can use SSH or a separate software to access the Database Server.

   b. Select "master" as the database.

c. Enter the following SQL Query and then click **[Go]**:

```
UPDATE
   master.system_settings_core
SET
   api_internal_account =<account_id>
```

Where `<account_id>` is the ID number of the account you want to use.

Some Services Fail to Generate Health, Availability, or Risk Values

Access Log Messages For Warnings or Errors

In this situation, some services in Skylar One do not generate any values for ***Health***, ***Availability***, or ***Risk***. For example, "Unknown" might appear instead of a value in one of the widgets on the **Service Investigator** page:

To address this issue, review the following settings and suggestions:

Step 1: Turn up the log level to trace:

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.

> **NOTE:** If you are a SaaS user, you cannot access the Skylar One appliance with SSH. Instead, go to the **Developer Logs** page (System > Tools > Skylar One Developer Logs) to obtain trace-level logs.

2. Log in as user **em7admin**.
3. Open the `nextui.env` file with vi or another text editor:

```
sudo vi /usr/local/silo/nextui/nextui.env
```

4. Change the log setting to the following: **NEXT_UI_LOG_LEVEL=all:trace**
5. Restart Skylar One and GraphQL with the following command:

```
sudo systemctl restart nextui
```

6. Tail the log with the following command:

```
sudo journalctl -u nextui -f
```

Step 2: Ensure that your service policy is valid:

1. In Skylar One, go to the **Business Services** page.
2. Select your service. The **Service Investigator** appears.
3. Click **[Edit]**.
4. Click the **[Service Policy]** tab.
5. Review the policy used by that service for any validation errors.
6. Address any errors in the service policy, and then click **[Save]**.

Step 3: Ensure that your service contains at least one service or device:

1. Go to the **Business Services** page.
2. Select your service. The **Service Investigator** appears.
3. Click **[Edit]**.
4. Click the **[Devices]** or **[Services]** tab for the service or services that are not displaying values.
5. Ensure that at least one device or service appears in the **Preview** section. If not, create a new filter to search for devices or services.
6. When you are finished, click `Save`.

Step 4: Generate audit data by running onDemandProcessing with GraphQL:

> **NOTE:** The `onDemandProcessing` GraphQL query currently retrieves data by calculating the "lastValue" in the last three intervals. If the value of the two most recent intervals is null, the system will log the calculation as incomplete because the metrics did not exist during that time.

1. Go to the **GQL Browser** page (System > Tools > GQL Browser).
2. In the left pane, type the following query:

```
query onDemand {
  harProviderOnDemandProcessing(ids: []) {
    results {
      serviceId
      timestamp
      health
      availability
      risk
    }
    auditHistory {
      serviceId
      ruleSetId
      ruleId
      timestamp
      sequence
      message
    }
  }
}
```

3. Click the **[Execute Query]** (Play) button.
4. Review the resulting audit information in the right pane.
5. If you know the service ID you are looking for, search for it by clicking inside the right pane and pressing **Ctrl+F**. The GQL Browser highlights the services that match the ID you searched for:

6. Scroll down to see the audit information for this service (look for the highlighted information):

> **NOTE:** The `onDemandProcessing` GraphQL query displays timing data. This includes the total number of services examined, the maximum time required to calculate these values for all included services, and the individual time taken for each service.

7. After running the `onDemandProcessing` query and updating the log settings on the server to do *all:trace*, you can now see trace-level log messages in the terminal where you ran `sudo journalctl -u nextui -f` in *Step 1*.

> **NOTE:** If you are a SaaS user, you can download the log files from the **Developer Logs** page (System > Tools > Skylar One Developer Logs).

8. Review the log messages for errors and warnings:

Check if Services Have Over 100 Constituents

Services will overload if they have over 100 constituents. Ensure that all of your business, IT, device, and custom (N-tier) services have less than or equal to 100 constituents.

To check if a service has over 100 constituents:

1. Go to the **Business ServicesBusiness Services** page.
2. Select the service you are reviewing. The **Service Investigator** appears.
3. Click **[Edit]**.
4. Click the **[Services]** or **[Devices]** tab.
5. Make sure that the chosen service does not have more than 100 constituents.

You can also check if your services have over 100 constituents by running SQL commands to search and return your results.

1. Either go to the console of the Skylar One Database Server or use SSH to access the Skylar One appliance.

> **NOTE:** If you are a SaaS user, you cannot access the Skylar One appliance via SSH. Instead, go to the **Developer Logs** page (System > Tools > Skylar One Developer Logs) to obtain trace-level logs.

2. Run the following SQL queries to check if the types of services you are searching for have over 100 constituents:
    - **For Device services**: `SELECT parent_id, count(child_id) FROM master_ap2.topo_service_relationships where link_type = 200 group by parent_id having count(*) > 100;`
    - **For IT or Business services**: `SELECT parent_id, count(child_id) FROM master_ap2.topo_service_relationships where link_type = 201 group by parent_id having count(*) > 100;`
    - **For N-tier services**: `SELECT parent_id, count(child_id) FROM master_ap2.topo_aggregate_services group by parent_id having count(*) > 100;`

> **NOTE:** ScienceLogic recommends searching for services that have over 80 constituents to identify those services approaching the 100-constituent limit.

All Services Fail to Generate Health, Availability, and Risk Values

> **NOTE:** If you are a SaaS user experiencing this issue, contact ScienceLogic Support: https://support.sciencelogic.com/s/contactsupport

In this situation, *all* of your services in Skylar One fail to generate any values for health, availability, or risk.

To address this issue, review the following settings and suggestions.

Step 1: Confirm that the Business Services processes exist:

1. Go to the **Process Manager** page (System > Settings > Admin Processes) and start typing "Business" in the ***Process Name*** filter.
2. Ensure that the "Business Services: Service Management Engine" and "Business Services: Service Topology Engine" processes appear and are enabled.

Step 2: Follow the steps in *Generate audit data using GraphQL*, above. If the process times out, then the processing has taken more than two minutes to complete, and no computed results are stored.

Step 3: Look for logs from the Python process:

1. The Python process calls the `onDemandProcessing` GraphQL query. If Python is having trouble connecting to GraphQL, it could be an authentication problem or some other code-related issue.
2. Look in `/var/log/em7` for newly created logs, and use the command `ls -lrt` to see if any new error logs were created with "business" in the file name.
3. Also check the `silo.log` for messages related to the business_service_management process by using the following command:

```
grep service /var/log/em7/silo.log
```

Device Services Fail to Load After an Upgrade

If you have upgraded Skylar One from an earlier version and your device services are not loading on the **Business Services** page, you might have outdated device class filters in your user preferences.

To clear the older device class filters:

1. Go to the **GQL Browser** page (System > Tools > GQL Browser).
2. In the left pane, type the following mutation:

```
mutation deletePreference{
  deletePreference(preferenceId:
"services.detaildevices.table.sort.order") {
    id
    preferenceValue
  }
}
```

3. Click the **[Execute Query]** (Play) button.

502, 503, or 504 Errors: Health, Availability, and Risk Values are All the Same or are Inaccurate

NOTE: If you are a SaaS user experiencing this issue, contact ScienceLogic Support: https://support.sciencelogic.com/s/contactsupport

Step 1: Check the number of services you have configured. If you are seeing 503 errors in the nextui log or within the Skylar One user interface, use the following procedure to check the number of services you have configured on your Skylar One system.

To determine the number of services you have:

1. Go to the **GQL Browser** page (System > Tools > GQL Browser).

2. In the left pane, type the following query:

```
query harProviders {
  harProviders {
    pageInfo {
      matchCount
    }
  }
}
```

3. Click the **[Execute Query]** (Play) button to see the number of services. In this example, the results shows that 10 services are configured.

```
"data": {
  harProviders {
    pageInfo {
      matchCount: 10
    }
  }
}
```

Step 2: (503 Errors) Confirm that the nginx configuration has an appropriate limit set. In some cases, the `limit_conn` value might be set to 20. Increase the value to 200.

To address this issue:

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.
2. Log in as user **em7admin**.
3. Confirm that the nginx config file has the `limit_conn perip` value set to 200 instead of 20:

   `sudo vi /etc/nginx/conf.d/em7_limits.conf`
4. If needed, update the line to say:

   `limit_conn perip 200;`
5. Run the following command:

   `sudo systemctl restart nginx`

Step 3: (503 Errors) Check to see if the nginx server is rate-limiting you.

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.
2. Log in as user **em7admin**.
3. Enter the following command:

   `sudo grep excess /var/log/em7/ngx.log`

4. If you see any results from the above command, then the nginx proxy is rate-limiting requests to your database. In that case, you should increase the rate limit to 100 requests per second. Edit the **em7_limits.conf** file:

   ```
   sudo vi /etc/nginx/conf.d/em7_limits.conf
   ```

5. Change the following line to **100r/s** from the default *5 r/s*:

   ```
   limit_req_zone $binary_remote_addr zone=addr_req:10m rate=100r/s;
   ```

6. Restart your Skylar One system:

   ```
   sudo systemctl restart nextui
   ```

Step 4: (502 Errors) Check node memory usage.

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.
2. Log in as user **em7admin**.
3. Enter the following command:

   ```
   sudo journalctl -u nextui|grep "JavaScript heap out of memory"
   ```

4. If you see any results form the above command, the node.js process is running out of memory. In that case, you should increase the space limit allocated. Edit the **nextui.service** to increase memory to 4096 or 8192 MB, depending on how much memory you have at your disposal:

   ```
   ExecStart=/usr/bin/node --max-old-space-size=4096
   /usr/local/silo/nextui/index.js
   ```

5. Restart your Skylar One system:

   ```
   sudo systemctl restart nextui
   ```

Step 5: (504 Errors) Check Nginx timeout.

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.
2. Log in as user **em7admin**.
3. Edit the `nextui.fragment` file:

   ```
   sudo vi /opt/em7/share/config/nginx.d/nextui.fragment
   ```

4. Change the `proxy_read_timeout` under "location /gql" to **900** as follows:

   ```
   proxy_read_timeout 900;
   ```

5. Restart your Skylar One system:

   ```
   sudo systemctl restart nextui
   ```

Advanced Troubleshooting

NOTE: If you are a SaaS user experiencing any of the following issues, contact ScienceLogic Support: https://support.sciencelogic.com/s/contactsupport.

Customization for Environments with More Than 2,500 Services

If you have an environment that has more than 2,500 services, you might need to modify some default settings in Skylar One, as described in this section.

Update Settings and Increase Default Values

To update your settings and increase your default values:

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.
2. Log in as user **em7admin**.
3. Increase the maximum service count variable. (The default value is 2500.)
   a. At the command line, enter:

      ```
      sudo vi /opt/em7/nextui/nextui.env
      ```

   b. Add the following line (or modify it, if it already exists), where `<new_service_limit>` is the maximum number of services you need in your environment:

      ```
      BUSINESS_SERVICES_MAX_SERVICES=<new_service_limit>
      ```

4. Increase the Node.js memory limit.
   a. At the command line, enter:

      ```
      sudo vi /etc/systemd/system/multi-user.target.wants/nextui.service
      ```

   b. Change the `ExecStart` line to the following, where the size is either 4096 or 8192, depending on how much memory you have available:

      ```
      ExecStart=/usr/bin/node --max-old-space-size=size
      /usr/local/silo/nextui/index.js
      ```

5. Restart `nextui` by entering the following at the command line:

   ```
   sudo systemctl restart nextui
   ```

Modify NGINX Rate Limit

If you have a large number of services in your environment and are seeing 503 errors, you might need to increase your NGINX rate limit.

To increase your NGINX rate limit:

1. Either go to the console of the Skylar One server or use SSH to access the Skylar One appliance.
2. Log in as user **em7admin**.
3. At the command line, enter the following:

   ```
   sudo grep excess /var/log/em7/ngx.log
   ```

4. If you see any results from this command, consider increasing your NGINX rate limit to 100 requests per second.

   a. Enter the following at the command line to edit the limit file:

      ```
      sudo vi /etc/nginx/conf.d/em7_limits.conf
      ```

   b. Change the `<value>` in the following line to "300r/s" from the default value of "100r/s":

      ```
      limit_req_zone $binary_remote_addr zone=addr_req:10m rate=<value>
      ```

   WARNING: If this value is set too high, the database will begin experiencing errors for too many connections.

5. Restart NGINX:

   ```
   sudo systemctl restart nginx
   ```

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

International: +1-703-354-1010