



Skylar AI Installation Guide

Version 2.4.0

Table of Contents

System and Sizing Requirements for Skylar AI	1
Cluster Requirements for Skylar Analytics and Skylar Advisor	2
Cluster Requirements for Skylar Analytics	2
Additional Requirements for Skylar Advisor	3
Third-Party Chart Dependencies for Skylar Analytics and Skylar Advisor	4
GPU Deployment Architecture Requirements	4
Cluster Sizing Guidelines for Skylar Analytics	4
Micro Deployment for Customer Proof Of Concept (POC)	5
Small Deployment	5
Medium Deployment	5
Large Deployment	5
Extra Large Deployment	6
Cluster Sizing Guidelines for Skylar Advisor	6
Small Deployment	6
Medium Deployment	6
Large Deployments	7
Required Infrastructure Dependencies	7
Ingress Controller	7
Recommended Infrastructure Components	7
Load Balancing	7
DNS and TLS Requirements	7
Optional Monitoring Integration	8
The Installation Process for Skylar AI	9
Registry Access Setup	10
Step 0: Pre-flight Validation	10
Step 1: Obtain Registry Credentials	10
Step 2: Sign into Helm Registry (Required)	11
Step 3: Install the OpenTelemetry Operator	11
Step 4: Download the Skylar AI Chart	11
Scaling Profiles	12
Availability Profiles	12

Installation Process	13
Step 1: Environment Configuration	13
Ingress Controller-Specific Annotations	14
High Availability Settings	15
Step 2: Deploy the Skylar AI Platform	16
Step 3: Configure TLS Certificate (Optional)	17
Validation and Access	17
Verify Deployment Status	17
Access the Platform	17
Monitoring the Integration (Optional)	18
Troubleshooting and Security for Skylar AI	20
Troubleshooting the Installation Process	21
Common Issues	21
Image Pull Failures	21
TLS Certificate Issues	21
DNS Resolution Issues	21
Ingress Controller Issues	21
Pod Startup Failures	21
Database Connection Issues	21
Large Upload/Download Issues	22
Cloud-Specific Issues	22
AWS	22
GCP	22
Azure	22
Useful Commands	22
Deployment Status Commands	22
Infrastructure Verification Commands	22
Registry Authentication Issues	23
TLS Certificate Issues	23
Ingress Controller Issues	23
Security Considerations	23
Security Best Practices	23

Third-Party Security Responsibilities	24
Backup and Recovery	24
Critical Data Components	24
ClickHouse Data Volumes	24
PostgreSQL Data Volumes	24
Backup Strategy Options	25
Volume-Level Backups	25
Application-Level Backups	25
Hybrid Approach	25
Backup Verification	25
Regular Testing Requirements	25
Monitoring and Alerting	25
Configuration Backup	26
Helm and Kubernetes Configuration	26
Support	26

Chapter

1

System and Sizing Requirements for Skylar AI

Overview

The installation and configuration of Skylar AI and Skylar Advisor uses Harbor for the registry and Helm for deployment. The Skylar AI platform is deployed as a single Helm umbrella chart containing over 20 microservices, databases, and supporting infrastructure components.

This chapter details the different cluster requirements for Skylar Analytics and Skylar Advisor, and it also includes requirements for only Skylar Analytics if you are currently not using Skylar Advisor. This chapter includes the required and recommended infrastructure dependencies for a Skylar AI deployment.

This chapter covers the following topics:

<i>Cluster Requirements for Skylar Analytics and Skylar Advisor</i>	2
<i>Cluster Sizing Guidelines for Skylar Analytics</i>	4
<i>Cluster Sizing Guidelines for Skylar Advisor</i>	6
<i>Required Infrastructure Dependencies</i>	7
<i>Recommended Infrastructure Components</i>	7

Cluster Requirements for Skylar Analytics and Skylar Advisor

IMPORTANT: Before beginning the installation process, verify that your environment has the **required** versions of Kubernetes and the container runtime (where relevant) listed in the following sections.

Cluster Requirements for Skylar Analytics

- **Kubernetes:** Version 1.32 or later (Skylar Analytics only).

You can run the following command to verify your environment:

```
kubectl version
```

NOTE: Kubernetes upgrades must be performed one minor version at a time. If you are planning a cluster upgrade, see the [Kubernetes upgrade documentation](#).

- **Storage:** Configured EBS StorageClass capable of dynamic PV/PVC provisioning. Storage cannot be NFS-backed, as this is not supported by ScienceLogic databases. Also, you must have one storage class marked as *default*.
- **Attached Storage Capacity:** Recommended 1 TB or more total storage capacity. This requirement varies by tenant datapoints per minute (DPM) requirements. For more sizing information, see [Cluster Sizing Guidelines for Skylar Analytics](#).
- **Node Storage Capacity:** Around 300 GB per node of internal storage for images and Kubernetes-related functions. Ensure that the root directory has most of this space provisioned. For more sizing information, see [Cluster Sizing Guidelines for Skylar Analytics](#).
- **Networking:** Currently, only IPv4 is supported. IPv6 is not supported at this time.
- **External Networking:**
 - The Skylar AI Kubernetes cluster will need Internet access to <https://registry.scilo.tools/>. This connection can be through a proxy if needed (additional configuration required).
 - 443 connectivity from the Skylar One instance to the Skylar AI Kubernetes cluster.
- **Deployment Options:**
 - Self-hosted Kubernetes clusters, such as bare metal or VMware
 - Cloud-managed Kubernetes services (EKS, GKE, AKS)

- **Availability Options:**
 - By default, Skylar Analytics will deploy the Postgres, Clickhouse-Keeper, OpenTelemetry Collector, and Tenant services in a high-availability pattern.
 - For the smallest resource footprint possible, Skylar deployments can be configured to use the minimum replicas required for each service.
 - Additionally affinity zone and node scheduling requirements for the Postgres and Tenant services may be relaxed from *hard* to *preferred*. For more information, see [Availability Profiles](#).
- **CLI Packages:** Customer will need kubectl access to the cluster along with Helm and Oras packages installed into the command-line interface (CLI).

For more sizing information, see [Cluster Sizing Guidelines for Skylar Analytics](#).

Additional Requirements for Skylar Advisor

IMPORTANT: If you are installing only Skylar Analytics, you can skip this topic.

- **Storage:** For Skylar Advisor, ScienceLogic requires that you use RWX Storage configurations only for Skylar Advisor services. This is in addition to the RW EBS storage class required for Skylar Analytics.
- **Kubernetes:** Version 1.35 with ContainerD 2.1 or cri-o 1.32 or later runtime.

You can run the following commands to verify your environment:

```
kubectl version
```

```
containerd --version
```

- **GPU resources provisioning:** You need to dynamically allocate GPU resources that can be correctly provisioned to Kubernetes nodes under the **[Resources]** tab. Examples would be NVIDIA time-slicing/MIG or VMware vGPU.
- **GPU Driver:**
 - **Minimum driver for MIG support:** 575.x (CUDA 12.9)
 - **Recommended driver (inference only):** 570.195.03 (CUDA 12.8)
 - **Not supported:** 580.x (CUDA 13); this driver configuration causes CUDA exceptions. While 580 is the latest NVIDIA recommended version, 580 requires code changes in Skylar AI for CUDA 13 compatibility.

To verify the drive version, run the following command:

```
nvidia-smi
```

Third-Party Chart Dependencies for Skylar Analytics and Skylar Advisor

Skylar AI includes several third-party Helm charts from Bitnami, Chainguard, and other providers:

- **Bitnami Charts:** ClickHouse, PostgreSQL HA, and associated components
- **Maintenance Notice:** ScienceLogic validates and updates third-party chart versions with each Skylar AI release. However, customers are responsible for:
 - Security patching of third-party components between Skylar AI releases
 - Vulnerability management for non-Skylar AI components
 - Understanding the security posture of included third-party dependencies

GPU Deployment Architecture Requirements

Multi-node appliance (4 or more nodes, 1 or more GPU per node):

- Time slicing per node is sufficient
- Driver 570.195.03 (CUDA 12.8) supported

Single-Node application (1 node, 8 GPUs):

- MIG required for workload isolation
- Driver 575.x minimum (CUDA 12.9) required for MIG on RTX PRO 6000
- Time slicing alone will NOT work – Kubernetes bin packing causes OOM

The XE7740 appliance is a single-node deployment. MIG and driver 575 is the required configuration for this hardware.

For additional Skylar Advisor configuration information, see [Cluster Sizing Guidelines for Skylar Advisor](#).

Cluster Sizing Guidelines for Skylar Analytics

This section lists the recommended configuration for your Skylar Analytics environment based on your expected volume of data, measured in *datapoints per minute* (DPM).

The following sizings are only for the Skylar Analytics application. You will need to consider additional resource requirements for system and cluster level services. In other words, do not size a node to just the requirements listed below.

Also, these sizing are for a single tenant cluster. Skylar AI can be multi-tenant, which will change the sizing recommendations as more Data Visualization tenants are added.

Micro Deployment for Customer Proof Of Concept (POC)

- **DPM Range:** 0 - 5,000 DPM (datapoints per minute)
- **Target Use Case:** POC environments
- **Data Visualization:** 5 CPU and 32 GB RAM

Estimated Total: 18 cores, 80 Gi RAM

- About 150 Gi free space for every Kubernetes node for images, logs, etc.
- About 250 Gi attached storage for database mounts.

Small Deployment

- **DPM Range:** 5,000 - 30,000 DPM
- **Target Use Case:** Development, QA, and penetration-testing environments
- **Data Visualization:** 5 CPU and 32 GB RAM

Estimated Total: 25 cores, 70 Gi RAM

- About 150 Gi free space for every Kubernetes node for images, logs, etc.
- About 512 Gi attached storage for database mounts.

Medium Deployment

- **DPM Range:** 30,000 - 215,000 DPM
- **Target Use Case:** Small production environments
- **Data Visualization:** 5 CPU and 64 GB RAM

Estimated Total: 34 cores, 225 Gi RAM

- About 150 Gi free space for every Kubernetes node for images, logs, etc.
- About 1 Ti Storage attached storage for database mounts.

Large Deployment

- **DPM Range:** 215,000 - 300,000 DPM
- **Target Use Case:** Production environments with moderate to high load
- **Data Visualization:** 5 cpu and 64gb RAM

Estimated Total: 53 cores, 432 Gi RAM

- About 150 Gi free space for every Kubernetes node for images, logs, etc.
- About 2 Ti Storage attached storage for database mounts.

Extra Large Deployment

- **DPM Range:** 300,000 - 900,000 DPM
- **Target Use Case:** High-scale production environments with heavy workloads
- **Maximum Single Instance:** 977,000 DPM
- **Data Visualization:** 8 CPU and 128 GB RAM

Estimated Total: 93 cores, 700 Gi RAM

- About 150 Gi free space for every Kubernetes node for images, logs, etc.
- About 4 Ti Storage attached storage for database mounts.

Cluster Sizing Guidelines for Skylar Advisor

The sizing guidelines for Skylar Advisor are independent of Skylar Analytics. For example, you could have a medium-sized deployment of Skylar Analytics with a small-sized deployment of Skylar Advisor. These requirements are in addition to the requirements of Skylar Advisor.

The deployment size is based on the number of expected active concurrent users and the size of the Corpus in Skylar Advisor.

Small Deployment

- **GPU:** 4 NVIDIA L40s or 4 NVIDIA RTX 6000 GPUs
- **CPU:** 80 cores
- **Memory:** 450 GB memory
- **Storage:** 250 GB of node memory
- **Attached Storage:** Depends on the expected Corpus size
- **Number of active concurrent users:** 1-15

Medium Deployment

- **GPU:** 4 NVIDIA H100 or H200s
- **CPU:** 90 cores
- **Storage:** 650 GB of node memory
- **Memory:** 250 GB memory
- **Attached Storage:** Depends on the expected Corpus size
- **Number of active concurrent users:** 16-30

Large Deployments

- **GPU:** 8 NVIDIA H100 or H200 GPUs
- **CPU:** 170 cores
- **Memory:** 1 Tb
- **Storage:** 250 GB of node memory
- **Attached Storage:** Depends on the expected Corpus size
- **Number of active concurrent users:** 31-120

Required Infrastructure Dependencies

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Ingress Controller

HTTPS traffic routing and SSL termination:

- **Required:** An ingress controller must be installed and configured
- **Recommended:** ingress-nginx controller
- **Alternatives:** AWS Load Balancer Controller, GKE Ingress, Azure Application Gateway, Traefik, or HAProxy Ingress

Recommended Infrastructure Components

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Load Balancing

A load balancer solution is needed to distribute traffic across worker nodes:

- **Self-hosted:** F5 BIG-IP, HAProxy with floating IP, or MetalLB
- **Cloud:** Application Load Balancer (ALB) or Network Load Balancer (NLB)

DNS and TLS Requirements

Required:

- **FQDN:** A fully qualified domain name pointing to the load balancer
- **TLS Certificate:** Valid TLS certificate for the FQDN, provided as a Kubernetes secret

Optional Monitoring Integration

Prometheus-based Monitoring

- **Supported:** Integration with existing Prometheus deployments
- **Benefits:** Custom metrics export from Skylar AI services

Chapter

2

The Installation Process for Skylar AI

Overview

The installation and configuration of Skylar AI and Skylar Advisor uses Harbor for the registry and Helm for deployment. The Skylar AI platform is deployed as a single Helm umbrella chart containing over 20 microservices, databases, and supporting infrastructure components.

This chapter explains how to get access to the registry for the installation, how to run the installation, and how to validate your installation.

This chapter covers the following topics:

<i>Registry Access Setup</i>	10
<i>Scaling Profiles</i>	12
<i>Availability Profiles</i>	12
<i>Installation Process</i>	13
<i>Validation and Access</i>	17

Registry Access Setup

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Step 0: Pre-flight Validation

Perform all of the following validations before continuing to Step 1:

1. Verify the Kubernetes version; the minimum is version 1.35:

```
kubectl version
```

2. Verify the container runtime version; ContainerD 2.1 or later is required:

```
containerd --version
```

3. Verify that the GPU Operator is running:

```
kubectl get pods -n gpu-operator
```

4. Verify that GPUs are visible to Kubernetes:

```
kubectl describe nodes | grep -A5 'nvidia.com/gpu'
```

5. Verify that storage classes are available:

```
kubectl get storageclass
```

6. Test the registry login:

```
helm registry login registry.scilo.tools
```

7. Verify that the ingress controller is running (adjust the namespace for your controller):

```
kubectl get pods -n ingress-nginx
```

8. When all of the checks have passed, continue to Step 1.

Step 1: Obtain Registry Credentials

1. Navigate to the registry at <https://registry.scilo.tools/>. The Harbor landing page appears.
2. Click **[Login via OIDC Provider]**.
3. Click **[Customer Login]**.
4. Enter your Salesforce credentials provided by ScienceLogic.

5. Click your name in the top right corner and select *User Profile*.
6. Click the copy icon in the **CLI Secret** field for use in the next steps.

Step 2: Sign into Helm Registry (Required)

You must authenticate with the Helm registry before proceeding with any chart operations. You authenticate with the registry using your Harbor CLI credentials (from step 6 in the previous procedure):

```
helm registry login registry.scilo.tools --username <your-harbor-username>
```

Step 3: Install the OpenTelemetry Operator

Install the OpenTelemetry Operator with custom image:

```
# Install OpenTelemetry Operator with custom image
helm registry login
helm install opentelemetry-operator oci://registry.scilo.tools/skyilar/opentelemetry-helm-charts/opentelemetry-operator \
  --version 0.114.0 \
  --namespace opentelemetry-operator \
  --create-namespace \
  --set "manager.collectorImage.repository=registry.scilo.tools/skyilar/slo-otelcol" \
  --set "manager.collectorImage.tag=0.16" \
  --set admissionWebhooks.certManager.enabled=false \
  --set admissionWebhooks.autoGenerateCert.enabled=true
```

TIP: If the formatting in the code snippets in this document are not retained when you copy and paste them, use the code snippets in the [online version](#) of this manual.

If you are using Cert-Manager for your certificate management, you can remove the following lines:

```
--set admissionWebhooks.certManager.enabled=false \
--set admissionWebhooks.autoGenerateCert.enabled=true
```

Step 4: Download the Skylar AI Chart

Use Helm to pull the Skylar AI **Chart.yaml** file from the registry, using the registry login from Step 2:

```
helm pull oci://registry.scilo.tools/skyilar/skyilar-charts
```

Decompress the charts:

```
tar -xvf skyilar-charts-x.x.x.tgz
```

Scaling Profiles

Choose the appropriate scaling profile based on your environment. These scaling profiles can be located within the downloaded Helm chart inside the sizing directory. Please work with your ScienceLogic Sales Engineer to understand which scaling profile you will need to deploy. Most Proof of Concepts (PoC) systems can be deployed with a *small* profile.

NOTE: This document will use **small.yaml** in the following steps, though larger deployments will need other files such as **medium.yaml** or **large.yaml**.

Availability Profiles

Starting with version 2.4.0, the Skylar AI platform is deployed by default in a High Availability (HA) pattern that requires each replica of a HA service to be scheduled on a distinct worker node and affinity zone (hard HA). As a result, the cluster must have at least **three** worker nodes.

If you do not want to use the default availability settings, you can choose an availability profile from the **availability** directory.

- Use **no-ha.yaml** to deploy the bare minimum number of replicas for eligible services. This option reduces the resources required by the cluster in exchange for availability guarantees.
- Use **soft-ha.yaml** to relax hard affinity zone (AZ) and worker node scheduling requirement to *preferred* for most services. HA services will attempt to distribute replicas across worker nodes and affinity zones, but will allow double scheduling if necessary. Soft HA requires at least **two** worker nodes, as the OTel Collector must schedule its two replicas on separate worker nodes, guaranteeing telemetry can be received from tenants if a worker node fails.

The following table provides additional details about the different types of High Availability (HA):

Profile	File	Minimum Nodes	Minimum Affinity Zones *	Replica Scheduling	Affinity	Best for
Hard HA (default)	N/A	3	3	Each replica on distinct node and AZ (required)	Hard constraints	Production, fault tolerance priority
Soft HA	<code>soft-ha.yaml</code>	2	2	Distributed across nodes and AZs, allows double scheduling if needed	Preferred	Balanced HA & flexibility

Profile	File	Minimum Nodes	Minimum Affinity Zones *	Replica Scheduling	Affinity	Best for
No HA	no-ha.yaml	1	1	Minimal replicas, no distribution	None	Dev/test, resource-constrained clusters

* *Minimum Affinity Zones* applies only to AWS deployments. On-premises deployments should ensure sufficient physical and logical separation across the infrastructure.

Installation Process

IMPORTANT: Unless noted otherwise, the information in this section is relevant for *both* Skylar Analytics and Skylar Advisor deployments.

Step 1: Environment Configuration

You will need to prepare the Helm override files with environment-specific details. Inside the decompressed **skylar-charts** you will find an example directory.

1. If you are deploying Skylar Analytics, populate the **examples/override-analytics.yaml** by creating a copy locally:

```
cp skylar-charts/examples/override-analytics.yaml .
```

```
cp skylar-charts/sizing/analytics-small.yaml .
```

2. Open the **override-analytics.yaml** file and add the values specific to the Kubernetes deployment. The file includes comments with helpful guidelines.

```
vi override-analytics.yaml
```

3. If you are using a proxy for egress Internet traffic, be sure to un-comment the sections that set the `HTTP_PROXY/HTTPS_PROXY/NO_PROXY` environment variables in the **override-analytics.yaml** file.
4. If you are also deploying Skylar Advisor with Skylar Analytics, you will need to complete the **examples/override-advisor.yaml** file. For Skylar Advisor you will need a fourth override that specifies the GPU-specific tolerations and resources. You will need to work directly with your sales engineer to come up with this configuration. You can find examples of this file inside the sizing directory, such as **sizing/l40s-rtx6000-combined.yaml** or **sizing/h100-resources.yaml**:

```
cp skylar-charts/examples/override-advisor.yaml .
```

```
cp skylar-charts/sizing/l40s-rtx6000-combined.yaml .
```

5. Finally, you need to choose a High Availability (HA) profile.

- If you are using the default HA configuration explained in [Availability Profiles](#), then you do not need to do anything else.
- If you want to choose **soft-ha.yaml** or **no-ha.yaml**, you need to apply the correct file. Depending on which configuration you want to use, run *one* of the following commands to copy the correct file out of the **skylar-helm-charts** into your local directory:

```
cp skylar-charts/availability/no-ha.yaml .
```

or

```
cp skylar-charts/availability/soft-ha.yaml .
```

Ingress Controller-Specific Annotations

Inside the **override-analytics.yaml** file, you will need to add annotations specific to your ingress controller. Skylar AI recommends the following ingress annotations. Each controller will be different, so the ingress nginx annotations below are only an example. If you are using Traefik or some other ingress controller, you will need to work with a ScienceLogic Sales Engineer to understand the comparable annotations.

Also, in the `global` section, set the `protocol` to `https`, as using HTTP will result in a redirect loop.

TIP: If the formatting in the code snippets in this document are not retained when you copy and paste them, use the code snippets in the [online version](#) of this manual.

```
#ingress-nginx

global:
  ingress:
    className: external
    annotations:
      nginx.ingress.kubernetes.io/proxy-body-size: 1024m
      nginx.ingress.kubernetes.io/proxy-buffer-size: 512k
      nginx.ingress.kubernetes.io/proxy-buffering: "on"
      nginx.ingress.kubernetes.io/proxy-buffers-number: "4"
      nginx.ingress.kubernetes.io/proxy-max-temp-file-size: 1024m
      nginx.ingress.kubernetes.io/proxy-read-timeout: "300"
    protocol: 'https'
    hostname: 'skylar.<yourdomain>.com'
```

```

skylar-auth:
  ui:
    ingress:
      className: external
      sslEmbed:
        annotations:
          nginx.ingress.kubernetes.io/rewrite-target: "/skylar-auth/sllembed"

          nginx.ingress.kubernetes.io/configuration-snippet: |
            more_clear_headers "Content-Security-Policy";
            more_set_headers "Content-Security-Policy: frame-ancestors:
*;"

#If deploying advisor add in the below
skylar-advisor-ui:
  ingress:
    annotations:
      # enables launching content residing outside our domain (eg:
storylane)
      nginx.ingress.kubernetes.io/configuration-snippet: 'more_clear_headers "Content-Security-Policy"; more_set_headers "Content-Security-Policy: frame-ancestors: *;"'

```

High Availability Settings

Skylar AI version 2.4.0 improved the stability of Skylar AI and allowed for zero-downtime upgrades by deploying Postgres and PgPool in High Availability (HA) mode by default, each with three replicas. Please note that you will need to account for the increase in resources required in Kubernetes before you upgrade to this release, as this update could triple the size of the Postgres and PgPool footprint, depending on the current deployment.

If you do **not** want to use HA, you will need to turn it off in an **override.yaml** file supplied to the Skylar Helm chart at deploy time, with the following values:

```

postgresql-ha:
  postgresql:
    replicaCount: 1
    pdb:
      create: false
  pgpool:
    replicaCount: 1

```

```
pdb:
  create: false
```

Step 2: Deploy the Skylar AI Platform

Run the following steps to deploy the Skylar AI platform with the scaling profile and customer overrides:

TIP: If the formatting in the code snippets in this document are not retained when you copy and paste them, use the code snippets in the [online version](#) of this manual.

```
helm upgrade --install skylar-production \
  oci://registry.scilo.tools/skylar/skylar-charts \
  --namespace skylar-production \
  --create-namespace \
  --values analytics-small.yaml \
  --values override-analytics.yaml
```

If you are enabling Skylar Advisor, add the following lines:

```
--value override-advisor.yaml
--value 140s-rtx6000-combined.yaml
```

NOTE: The second `.yaml` file name above is just an example. You can name this file based on the GPU being deployed.

If you are enabling *soft or no High Availability (HA) options*, add the following line:

```
--value no-ha.yaml
```

When running the Helm install command, you can open a second and third terminal to monitor pod creation and events in parallel while the Helm install runs:

1. Open a second terminal, full pod state, that refreshes every two seconds:

```
watch kubectl get pods -n skylar-production
```

NOTE: The `watch` command is preferred over `kubectl get pods -w` because `-w` only shows incremental changes, while `watch` shows the full current state.

2. Open a third terminal and watch for error events:

```
kubectl get events -n skylar-production --sort-by='.lastTimestamp' -w
```

Step 3: Configure TLS Certificate (Optional)

NOTE: You can skip this step if you have automated certificate management (such as [cert-manager](#)) or deployment processes that handle TLS configuration.

Create a TLS secret with your provided certificate and private key:

```
kubectl create secret tls skylar-tls-secret \  
  --cert=path/to/your/certificate.crt \  
  --key=path/to/your/private.key \  
  --namespace=skylar-production
```

Validation and Access

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Verify Deployment Status

Run the following commands to verify your deployment was successful:

1. Check to see if all pods are running:

```
kubectl get pods -n skylar-production
```

2. Verify the ingress configuration:

```
kubectl get ingress -n skylar-production
```

3. Check to see if the TLS certificate is properly configured:

```
kubectl describe ingress -n skylar-production
```

4. Check to see if image pull secrets are working:

```
kubectl describe pod <any-skylar-pod> -n skylar-production | grep -A5  
"Events:"
```

Access the Platform

To access the Skylar AI platform for the first time:

1. Verify DNS Resolution:

```
nslookup skylar.<yourdomain>.com
```

2. Navigate to the Skylar AI user interface using your provided FQDN, such as `https://skylar.<yourdomain>.com`.

IMPORTANT: Skylar AI requires HTTPS for access, and using HTTP will result in a redirect loop. Also, you will need to accept the self-signed certificate warning if you need to [configure TLS](#).

3. Log in for the first time with the default email of `skylar@sciencelogic.com`.
4. Set a password for your first login. The password must be at least 15 characters, with at least one special character required.
5. To set up a new Skylar AI account, see [Provisioning a New Account](#).
6. For more information about other Skylar AI settings, see [Configuring Skylar AI System Settings](#).

Please note that multi-factor authentication (MFA) is enabled by default when you first log in. Complete the following steps if you need to disable MFA (for test or development environments, for example):

1. To disable MFA at the account level, go to [Skylar Settings](#), navigate to the **Authentication** page ( > Account Settings > Account Access > Authentication), and click the **[Not Required]** button for **MFA Status**.
2. To reset MFA at the user level, navigate to the **Users** page ( > Account Settings > Account Access > Users). Click the Actions icon () and select **Reset MFA**.

If your first login results in a timeout loop after password change:

1. Press **[Ctrl]+[Shift]+[Delete]** and clear the browser cache and cookies.
2. If the issue persists, try a different browser.
3. Wait 2-3 minutes and retry the login.

NOTE: The timeout issues might differ between NGINX and Traefik ingress controllers. Report occurrences to ScienceLogic support for investigation.

Monitoring the Integration (Optional)

If you have an existing Prometheus setup, you can configure it to scrape metrics from Skylar AI services.

In the **override.yaml** file, enable the metrics endpoints:

```
global:
  enableMonitoring: true
```

This setting exposes Prometheus metrics endpoints on Skylar AI services that can be scraped by your existing Prometheus deployment. Configure your Prometheus to discover and scrape these endpoints based on your service discovery method, such as Kubernetes service discovery or static configurations.

You can ingest your scraped metrics from Prometheus into Skylar One leveraging the "SL1 Prometheus" PowerPack.

Chapter

3

Troubleshooting and Security for Skylar AI

Overview

This chapter includes a list of common issues with the installation process for Skylar AI. This chapter also include security considerations, backup and recovery procedures, and a list of potentials topics where you might need to contact Skylar AI Enablement.

This chapter covers the following topics:

<i>Troubleshooting the Installation Process</i>	21
<i>Security Considerations</i>	23
<i>Backup and Recovery</i>	24
<i>Support</i>	26

Troubleshooting the Installation Process

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Common Issues

Image Pull Failures

- Verify that the registry secret is correctly configured.
- Check that the CLI secret is still valid in Harbor.
- Ensure the namespace has access to the registry secret.
- Verify that the Helm registry login was successful.

TLS Certificate Issues

- Verify that the TLS secret contains a valid certificate and key.
- Check that certificate matches the FQDN.
- Ensure the certificate is not expired.

DNS Resolution Issues

- Verify that the FQDN points to the correct load balancer IP.
- Check the DNS propagation if recently updated.

Ingress Controller Issues

- Verify that the ingress controller is running and healthy.
- Check the ingress controller logs for errors.
- Ensure that the ingress class name matches your controller.
- Verify that the ingress annotations are compatible with your controller.

Pod Startup Failures

- Check the resource constraints and storage availability.

Database Connection Issues

- Ensure that the ClickHouse and PostgreSQL pods are healthy.

Large Upload/Download Issues

- Verify that the ingress controller supports large request bodies (256 MB and up).
- Check the proxy timeout configurations.
- Ensure that the proper buffering settings are applied.

Cloud-Specific Issues

AWS

- Check the IAM permissions for EBS/EFS access.

GCP

- Verify the service account permissions for persistent disks.

Azure

- Ensure the proper RBAC for storage resources.

Useful Commands

Deployment Status Commands

To check the deployment status:

```
helm status skylar-prod -n skylar-production
```

To view pod logs for troubleshooting:

```
kubectl logs <pod-name> -n skylar-production
```

To describe the pod for detailed information

```
kubectl describe pod <pod-name> -n skylar-production
```

Infrastructure Verification Commands

To check the ingress controller status (adjust the namespace based on your setup):

```
kubectl get pods -n ingress-nginx
```

To check persistent volume claims:

```
kubectl get pvc -n skylar-production
```

To test registry connectivity:

```
kubectl run test-registry --image=registry.scilo.tools/skyelar/skyelar-charts:latest \
--dry-run=client -o yaml | kubectl apply -f -
```

Registry Authentication Issues

If you encounter image pull errors:

- Make sure the CLI secret has not expired in Harbor.
- Re-authenticate with the Helm registry login.
- Verify that the Kubernetes secret was created correctly.
- Make sure the secret exists in the correct namespace.

TLS Certificate Issues

If you encounter TLS-related problems:

- Verify that the certificate is valid for your FQDN.
- Make sure the certificate is in PEM format.
- Include intermediate certificates if required.
- Verify that the private key matches the certificate.

Ingress Controller Issues

If ingress resources are not working:

- Verify that the ingress controller is running.
- Make sure the ingress class name matches your controller.
- Check the ingress controller logs for configuration errors.
- Make sure the backend services are healthy and have endpoints.
- Make sure the ingress annotations are supported by your controller.

Security Considerations

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Security Best Practices

- **Registry Credentials:** Store CLI secrets securely and rotate them regularly.
- **TLS Certificates:** Ensure certificates are from trusted CAs and renewed before expiration.

- **Network Policies:** Consider implementing Kubernetes network policies to restrict inter-pod communication.
- **RBAC:** Configure appropriate role-based access controls for the Skylar AI namespace.
- **Ingress Security:** Configure appropriate security headers and rate limiting on your ingress controller.
- **Secrets Management:** Use external secret management solutions, such as HashiCorp Vault, or AWS Secrets Manager for production.
- **Cloud Security:** Follow cloud provider security best practices.
- **Third-Party Components:** Monitor security advisories for included third-party components.

Third-Party Security Responsibilities

While ScienceLogic validates and updates third-party chart versions with each Skylar AI release, customers should:

- **Monitor Security Advisories:** Stay informed about security issues in included third-party components.
- **Plan for Updates:** Be prepared to upgrade Skylar AI when security patches are available.
- **Vulnerability Assessment:** Include third-party components in security scanning and assessment processes.
- **Risk Management:** Understand the security posture of all included dependencies.

Backup and Recovery

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

Critical Data Components

Skylar AI stores critical data in the following persistent volumes, which must be backed up regularly.

ClickHouse Data Volumes

- Contains analytics data, metrics, and time-series information.
- Recommended backup frequency: Daily with 30-day retention.

PostgreSQL Data Volumes

- Contains application metadata, user data, and configuration.
- Recommended backup frequency: Daily with 30-day retention.

Backup Strategy Options

This section covers the different types of backup strategies you can use with Skylar AI.

Volume-Level Backups

- Use cloud provider snapshot capabilities, such as AWS EBS, GCP Persistent Disks, or Azure Disks.
- Implement Kubernetes volume snapshots using CSI drivers.
- Consider third-party backup solutions like Velero for comprehensive cluster backup.

Application-Level Backups

- Database-native backup tools for ClickHouse and PostgreSQL.
- Export application configuration and secrets.
- Backup Helm chart values and deployment configurations.

Hybrid Approach

- Combine volume snapshots for fast recovery with application-level backups for granular restore options.
- Implement both local and off-site backup storage for disaster recovery.

Backup Verification

Regular Testing Requirements

- Monthly: Test backup restoration procedures in non-production environment.
- Quarterly: Full disaster recovery simulation.
- Document and validate recovery time objectives (RTO) and recovery point objectives (RPO).

Monitoring and Alerting

- Monitor backup job completion and success rates.
- Alert on backup failures or missing backups.
- Verify backup accessibility and integrity.

Configuration Backup

Helm and Kubernetes Configuration

- Export and store Helm values files.
- Backup Kubernetes secrets and ConfigMaps.
- Maintain version-controlled infrastructure as code.

CAUTION: Store configuration backups containing secrets in encrypted storage with restricted access. Never commit secrets to version control systems.

For detailed backup implementation guidance, consult your cloud provider's backup documentation and the Kubernetes Backup Best Practices.

Support

IMPORTANT: The information in this section is relevant for both Skylar Analytics and Skylar Advisor deployments.

For deployment assistance, configuration guidance, or troubleshooting support, contact Skylar AI Enablement. They can provide:

- Environment-specific override file templates.
- Scaling recommendations based on your requirements.
- Custom configuration for enterprise integrations.
- Cloud-specific deployment guidance.
- Registry access troubleshooting.
- TLS certificate configuration assistance.
- Ingress controller configuration guidance.
- Post-deployment optimization.
- Third-party component guidance (limited to integration aspects only).

NOTE: Support for third-party components, such as Bitnami charts, is limited to integration and configuration guidance. For issues specific to these components, consult their respective documentation and support channels.

© 2003 - 2026, ScienceLogic, Inc.

All rights reserved.

ScienceLogic™, the ScienceLogic logo, and ScienceLogic's product and service names are trademarks or service marks of ScienceLogic, Inc. and its affiliates. Use of ScienceLogic's trademarks or service marks without permission is prohibited.

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information herein, the information provided in this document may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described herein at any time without notice.



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

International: +1-703-354-1010