Silk Echo Installation and User Guide

Contents

About Silk

Silk is more than just a storage solution or a cloud migration tool. We're an intelligent data platform that optimizes cloud resources to significantly improve performance, resiliency, data lifecycle management, and cost efficiency in the cloud. We operate between databases and cloud infrastructure to let our customers manage their data in ways cloud providers can't. Silk gives you the means to shift your mission-critical data to the cloud and operate on par with, or beyond, even the fastest on-prem environments.

Silk lets you use one platform to run multiple databases concurrently, giving you more ways to improve management of your data. Our platform features two-tier virtualization with a fully distributed, cloud-native, multi-cloud framework that decouples your compute and storage and promotes self-optimizing for scalability. Our dynamic database block size alignment uses an adaptive, log-structured datastore to improve performance, while our elastic architecture optimizes latency and throughput. We use a shared multi-tenant datastore so transaction, analytic, and mixed workloads can run together on pooled resources for greater efficiency. Silk is a database-, application-, and cloud-agnostic platform that helps companies harness the full power of the cloud.

The Silk Intelligent Data Platform can give you the tools you need to make your cloud environment run up to 10x faster. Contact our Sales team for a demo to see how Silk can help you make the most of your cloud resources.

Executive Summary

As more organizations move their data to the cloud, Silk has set out on a mission to rewrite the rules of data operations in the cloud. Our platform is architected in a way that makes it easy to migrate and manage your data in any cloud environment. Silk lets customers compose data resources based on application requirements, allocating just the right amount of both performance and capacity while minimizing resource overprovisioning and waste. You can extend your environment's operational flexibility by scaling performance and capacity resources independently of each other in a dynamic cloud framework.

The Silk Intelligent Data Platform operates using industry-leading software and services. Our software stack is built on Silk DataPods, Silk Flex, and Silk Clarity, which work together to provide a rich set of data services, machine learning, analytics, and policy-based automation and orchestration. Our tools and services support a comprehensive cloud optimization strategy through performance, resiliency, data lifecycle management, and cost efficiency.

This guide describes the software features, installation and user experience workflows, and known issues for Silk Echo.

Introduction

Overview

Designed for database management, DevOps, and AI, Silk Echo accelerates business workflows by letting you quickly and efficiently copy databases from one host to another. The service offers a simple, centralized way to manage your databases, with the ability to create DevTest environments, perform backup and restore operations, and generate analytics. Echo resources only use storage if you write to them, providing significant storage savings, instant provisioning, and high scalability.

Echo enables realistic testing, safer infrastructure changes, and faster incident response by giving teams fast, scalable access to production-like environments that won't impact live systems. For DevOps teams, this means streamlined CI/CD pipelines, accelerated troubleshooting and debugging with reduced risk, environment parity, and reduced downtime. For AI enablement, this means data pipeline development, model training on production-level data, feature engineering, and safe data exploration and experimentation.

Echo is accessible through the Silk Flex interface and the Silk API, enabling seamless integration. To use the service, you'll install a host agent on the host to enable cloning operations. The host agent utilizes Microsoft Shadow Copy Service (VSS) in conjunction with the Silk sVSS driver.

The service is supported in Azure and Google Cloud on Windows 2019+ hosts running Microsoft SQL Server 2019+ databases. Echo takes a snapshot of your source database (called a DB-Snapshot), then creates thin clones for all the volumes on the DB-Snapshot and maps them to the target host to create a clone of the database (called an Echo Database) on the target host. The DB-Snapshots are stored in Echo, allowing you to reuse existing DB-Snapshots to create additional Echo Databases.

Supported Capabilities

This list provides our best practice recommendations and the supported conditions for using Silk Echo:

1. Echo is supported in Azure and Google Cloud on hosts installed with Windows OS 2019 and above and with MSSQL 2019 and above.
2. Echo currently works separately with each Silk DataPod, which means you can't copy databases between DataPods using Echo.
3. For MSSQL 2022 and higher, Echo supports hosts mapped to multiple DataPods (you'll select which DataPod to use with Echo during sVSS setup). For MSSQL 2019 and lower, all hosts must be mapped to the same DataPod.
4. Ensure the source database files are on volumes within the DataPod. Likewise, ensure the files for a single database are all located on the same volume group (VG) in the DataPod.
5. Echo Databases (the thin databases created using Echo) cannot be cloned.
6. Cloning is performed serially, which means that during database cloning, other cloning operations are blocked.

System Notes

1. Echo Database names and database suffixes should only include alphanumeric characters, underscores, and dashes. The maximum supported length for Echo Database names is 128 characters.
2. sVSS can only be connected to one DataPod.
3. If the database is rebuilding the table index, attaching the Echo Database may take longer than operation timeout.

User Roles

Echo supports two user roles for customers:

1. **Flex:** The default user role, available through the Silk DataPod.
2. **Echo:** A limited role that allows Echo access without requiring access to a DataPod. Use this role to grant access based on the principle of least privilege. To access Echo with the Echo user role, navigate to https://[Flex-IP] in a browser and log in using Echo credentials.

Note: There is also a separate user role for Flex and Echo used by Silk's Support team for troubleshooting.

Echo Prerequisites

Version Details

Before installing Echo, verify your Silk environment is running the supported versions:

Silk Component	DataPod	Flex	Host Agent	sVSS
Supported Versions	v8.5.95.21 or higher	v1.95.13 or higher	v.1.0.17	2.0.15-setup-64bit

Installation Prerequisites

Prerequisite	Description
DataPod	Prior to installing Echo, ensure Silk Flex and the SDP Management console are configured. For instructions on how to set up Silk Flex and the SDP Management console, refer to the Silk Flex User Guide and the Silk Cloud Platform User Guide.

Host Installation	Echo supports Windows hosts running Microsoft SQL Server. Configure hosts based on the best practices outlined in the Silk SDP Connectivity & Host Best Practices guide.
Networking	The host agent should be able to send HTTPS requests to Flex. Typically, the host agent and Flex are on the same network (but this varies depending on your business needs).
Connection to Silk DataPods	The source database files must be stored on volumes provided by a DataPod. Currently, all Echo source files must be in the same volume group and on the same DataPod.

System Survey

Prior to Echo installation, the Silk team takes a survey of your environment, including:

Component	Description
Cloud	[Azure or Google Cloud]
Windows Server Version	[Insert version (Windows 2019 or above)]
SQL Version	[Insert version (SQL Server 2019 or above)]
Host Connectivity	2 (one for iSCSI and one for management)
DB Sizes	[Insert sizes]
AUTH	[SQL/Windows authentication]
SDP ID	[Insert ID]
Cluster ID	[Insert ID]
SDP Floating IP	[Insert IP] To see the external IP (floating IP) of your DataPod, log in to the SDP Management console and navigate to **System > General**. The floating IP is listed in the **DNS Settings** section as **Primary**.
Flex IP	[Insert IP]

Installing Echo

Echo is available via the Silk Flex interface or the Silk API, and all installation prerequisites and limitations apply to both Flex and the API.

Validate Host Connectivity

1. Log in to the host via remote desktop and then run PowerShell as an administrator.
2. Confirm the host is connected to the DataPod via the iSCSI protocol:

```
Get-IscsiConnection
```

Note: The number of iSCSI portals in the response should correspond to the number of c.nodes on the DataPod (e.g., if you have three c.nodes, you should see three response outputs).

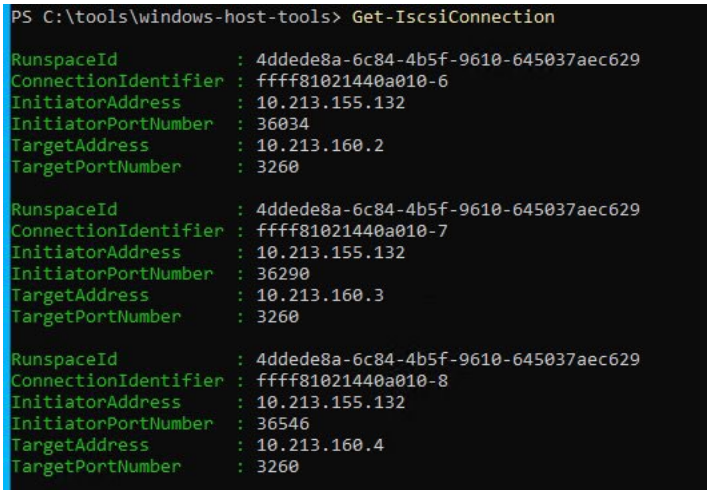

3. Get the Microsoft Mulitpath IO (MPIO) settings and ensure they're identical to the settings in the screenshot below:

```
Get-MPIOSetting
```

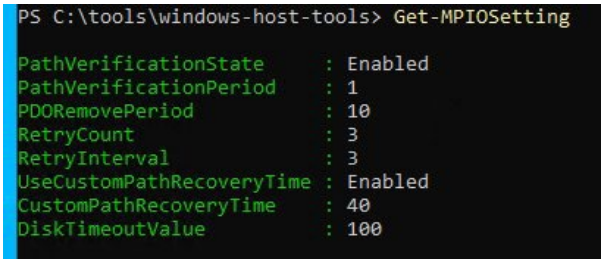

4. Check connectivity to the DataPod and the Flex management IP:

```
Test-NetConnection -ComputerName [Flex IP] -Port 443
Test-NetConnection -ComputerName [SDP Floating IP] -Port 443
```

You should see that `TcpTestSucceeded` is `True` in the response.

Install and Configure sVSS

1. Log in to the host and run the sVSS installer provided by Silk's Cloud Infrastructure team. If you're prompted with a warning, click **More info** and then click **Run anyway**.
2. Click **Yes** when prompted to allow the program to make changes to your device.

   The installation wizard opens.
3. Click **Install** to initialize the installation wizard.
4. On the **Silk VSS Provider Configuration** dialog, fill in the configuration details:
   - **SDP External IP:** Copy the floating IP from your Silk system survey and paste it into this field.
   - **Host Name:** Enter the name of your host. To get the host name, log in to the SDP Management console and navigate to **Volumes & Hosts > Hosts**. Any hosts you have configured, and details like the host name, are visible here.
   - **Username:** Enter the username for your DataPod.
   - **Password:** Enter the password for your DataPod.
   - **Snapshot Prefix:** In the text field, enter a prefix for your snapshot names. This prefix is automatically applied when you create DB-Snapshots.
   - **Retention Policy:** In the text field, enter **Best_Effort_Retention**, which retains snapshots for as long as possible.
   - **Check VG Full:** Ensure this checkbox is unchecked. If it's checked, the VSS snapshot only succeeds when all volumes in a VG are used by the database.

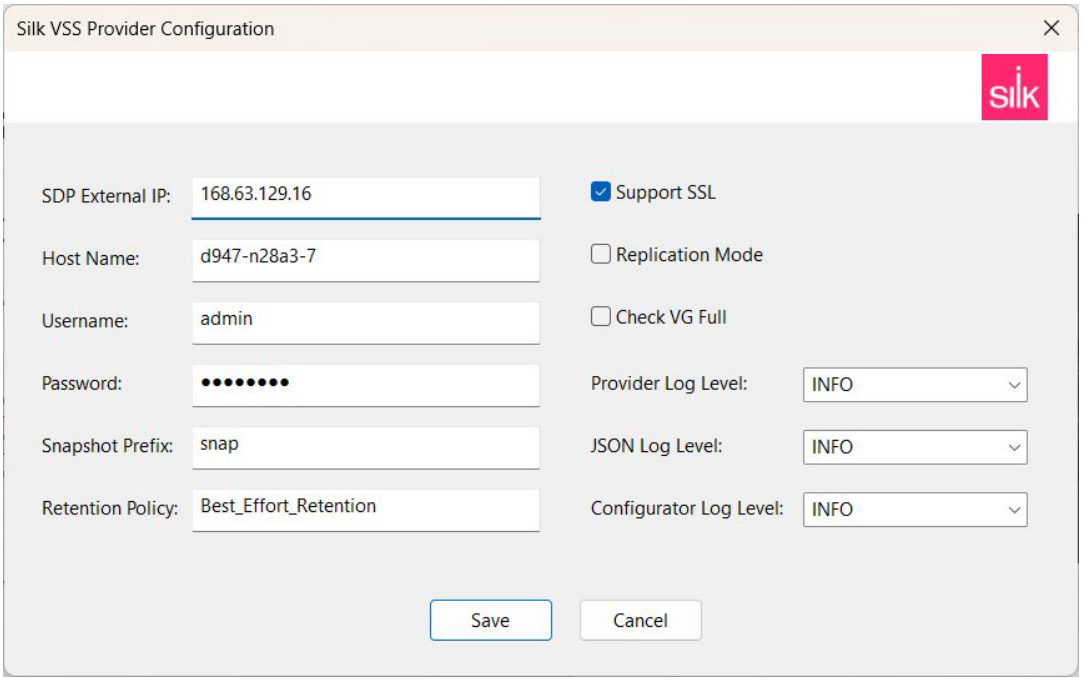

5. Use the dropdowns to update the log levels as needed. By default, all log levels are set to **INFO**.
6. After the configuration details are filled in, click **Save**.

   You should see a confirmation pop-up that your configuration saved successfully.

7. Click **Finish** to close out of the sVSS installation wizard.

Install the Host Agent

1. Download the Silk host agent (**SilkAgent**) package provided by Silk's Cloud Infrastructure team.
2. Open and run the installer. If you're prompted with a warning, click **More info** and then click **Run anyway**.
3. Click **Yes** when prompted to allow the program to make changes to your device.

   The **SilkAgent** installation wizard opens.
4. Click **Next** to proceed, then review the **Silk Host Agent License Agreement** and click **I Agree** to continue.
5. Choose the install location for the **SilkAgent** package and then click **Next**. By default, **C:\Program Files\SilkAgent** is populated as the destination folder.
6. In the **Enter Flex/Host IP** field, enter the IP address of your Silk Flex instance.
7. In the **Enter Token** field, create and then paste in your token by registering the host**:**
   - Log in to Silk Flex and use the cogwheel menu in the top right to select **Echo**.
   - Along the top of the **Manage Hosts** page, click **New Host Registration**.
   - In the **Host ID** field, enter an alphanumeric ID for the host, starting with a letter. This is a unique ID within the Echo management page and does not have to be the same as the host name.
   - In the **DB vendor** field, leave **mssql** as the default selection. Currently, Microsoft SQL Server is the only supported database vendor.
   - Click **REGISTER HOST**.

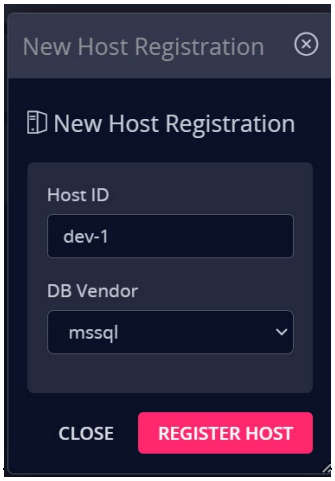

   You should see a **New Host Was Registered** confirmation along with your automatically generated token.
   - Copy the token to your clipboard and paste it into the **Enter Token** field in the host agent installation wizard.
8. (Optional) Check the **Use mount points instead of drives** checkbox to use mount point instead of drive letters. In the corresponding text field, specify a locally attached disk to serve as the mount point base location (by default, **c:\silk_mounts** is populated).

Leave the checkbox unchecked to use drive letters. Each volume will be mapped to a disk with a different drive letter, so the maximum number of supported volumes is limited to the number of free drive letters available.

9. Click **Test Connection** to verify connectivity, and then click **Next**.
10. Use the **Authentication** dropdown to choose an authentication method:
    - **SQL Server Authentication:** Enter the username and password for your SQL Server application. Use this option if your server listens for `localhost: 1443`. Otherwise, we recommend using **Connection String**.
    - **Windows Integrated Authentication:** Enter the username and password for the Windows user on the SQL Server host. This user should have permission to run as a service.
    - **Connection String (Recommended):** In the corresponding text field, enter this syntax, using your SQP Server address, username, and password:

      ```
      Server=myServerAddress;User Id=myUsername;Password=myPassword;
      ```

    - **Preserve existing:** On subsequent installations or upgrades on the same host, this option lets you reuse your previous credentials (from the first installation) without re-entering them.

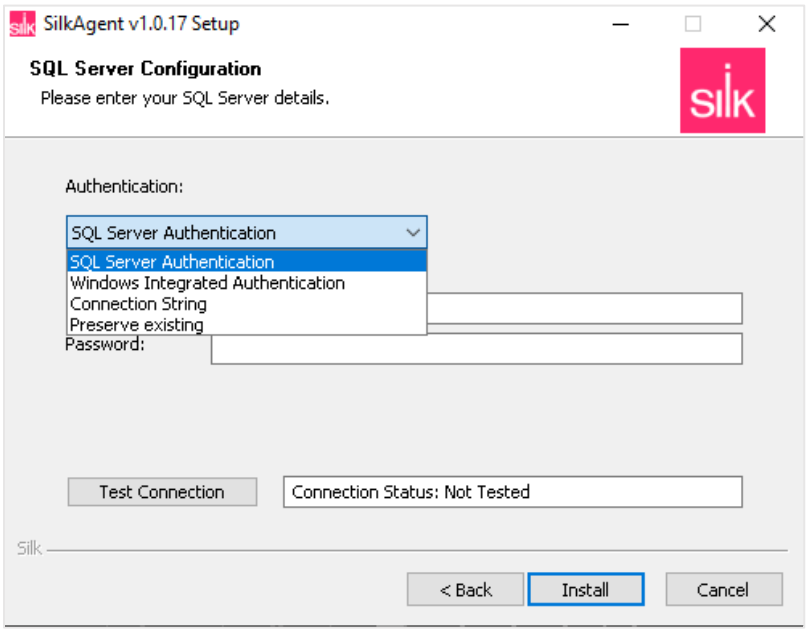

11. After filling in the Echo details, click **Install** to complete the installation.
12. After installation, click **Finish** to close out of the installer.

    You should now see the status of your new host is active.

Configuring Echo Authentication

Windows Integrated Authentication Permissions

Run Windows Integrated Authentication as a Service

Windows Integrated Authentication requires user permissions to run as a service. To configure these permissions:

1. Open the **Local Group Policy Editor**:
   - Press **Windows + R** to open the **Run** dialog.
   - In the **Open** text field, enter **gpedit.msc** and click **OK**.
2. In the lefthand menu, navigate to **Computer Configuration > Windows Settings > Security Settings > Local Policies**.
3. Double-click on the **User Rights Assignment** folder and then double-click the **Log on as a service** policy.
4. Click **Add User or Group...**
5. In the **Enter the object names to select** field, enter the username or group name you want to grant this permission to, and then click **OK**.
6. In the **Log on as a service** dialog, click **Apply** and then click **OK.**
7. Close out of the **Local Group Policy Editor**.

Grant Windows Integrated Authentication Admin Permissions

Windows Integrated Authentication also requires admin privileges to run Microsoft Volume Shadow Copy Service (VSS). To configure these permissions:

1. Open the **Local User Management** settings:
   - Press **Windows + R** to open the **Run** dialog.
   - In the **Open** text field, enter **lusrmgr.msc** and click **OK**.
2. In the left pane, select the **Users** folder.
3. In The middle pane, right-click on the user you want to add to the local **Administrators** group, and then select **Properties**.
4. In the **Properties** dialog, select the **Member of** tab and then click **Add**.
5. In the **Enter the object names to select** field, enter **Administrators** and then click **Check names**.

   This finds, highlights, and underlines the **Administrators** group.
6. Review the group details and then click **OK**.
7. In the **Properties** dialog, click **OK** to save your changes.

SQL User Permissions

These permissions are required on the SQL user and for Integrated Windows Authentication:

```
GRANT VIEW ANY DEFINITION TO "silk_agent";
GRANT VIEW SERVER STATE TO "silk_agent";
ALTER SERVER ROLE dbcreator ADD MEMBER "silk_agent";
```

Getting Started with Echo

Access to Echo depends on your user role. **Flex** users can access Echo through Silk Flex. **Echo** users have limited access and use the Flex IP to log in directly to Echo.

To access Echo via Silk Flex:

After you log in to Flex, use the cogwheel menu in the top right corner to select **Echo**. This opens the **Manage Hosts** page, which allows you to register/add new hosts, manage existing hosts, and perform Echo operations.

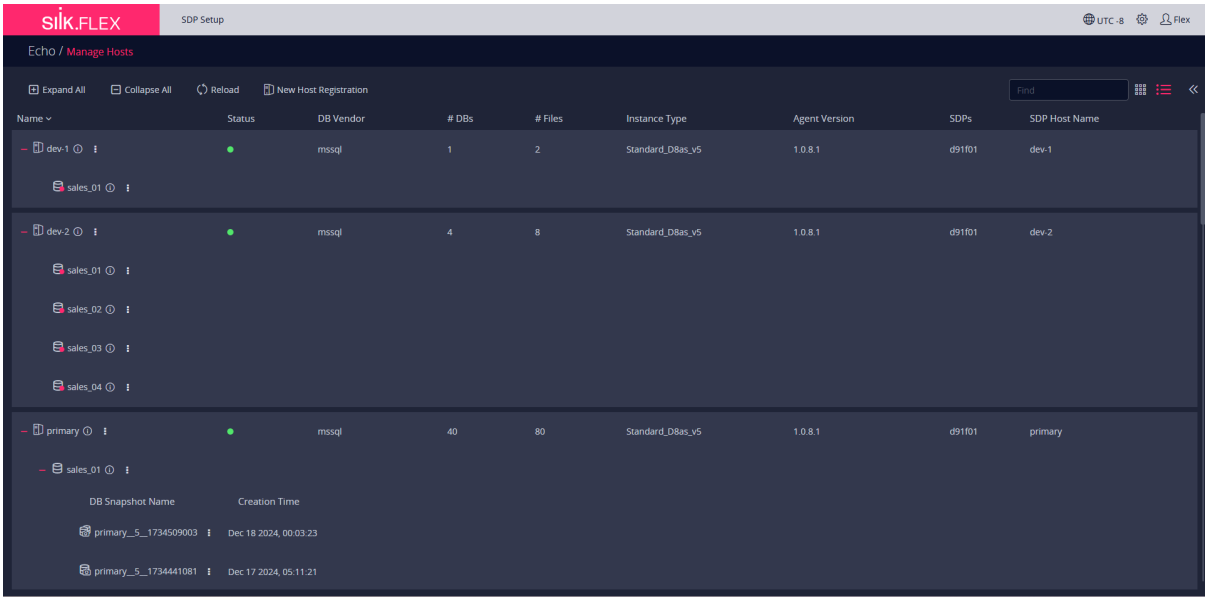

View Hosts and Databases

On login to Echo, you'll see the **Manage Hosts** dashboard with all your host details. The top navigation allows you to expand all details, collapse all details, reload Echo data, and add a new host registration. Click **Expand All** to show all your Echo hosts, databases, and snapshots. Click **Collapse All** to show only your Echo hosts.

Silk offers two ways to visualize and manage your Echo data: using Cards view or Grid view. Use the icons toward the top right corner of the **Manage Hosts** page to toggle between these views:

- **Cards view:** This default view displays all your Echo details at the host level in an expandable list. Click the plus icon (**+**) to the left of a host name to expand the host details. The expanded view shows all databases associated with the host.

  This view shows more details about your hosts than Grid view, but cloning in this view is a bit more manual.

- **Grid view:** This view displays all your Echo details in dynamic tiles that let you drag and drop databases from one host to another for even faster cloning. To clone a database in Grid view, simply drag it from the source to the target, or select the three-dot menu to the right of the database name.

Cards view shows these details for each host:

Column Name	Column Description
Name	The name assigned to the host during host registration.
Status	Whether the host is installed. A green dot indicates that the host is installed, while an orange dot indicates that the host is not installed.
DB Vendor	The vendor used to create the database. This field is only populated if there are databases associated with your host. Currently, Microsoft SQL Server (**mssql)** is the only supported database vendor.
# Databases	The number of databases on the host.
# Files	The number of database files on the host.
Instance Type	The host's instance type (e.g., **Standard_D4as_v5**).
Agent Version	The Silk agent version installed on the host.
SDPs	The DataPod the host is connected to.
SDP Host Name	The name of the host as it appears in the SDP Management console. This can be the same as the host name or it can be different.

Select the information icon to the right of a host name to see the host details, including the host name, host ID, OS type, instance type, location, zone, and host IPs.

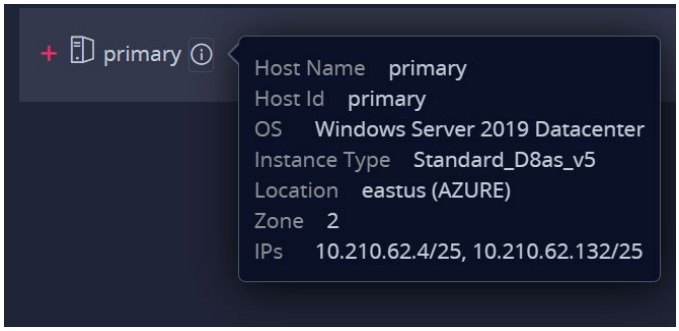

Select the three-dot action menu to the right of a host name to clone databases at the host level, create a DB-Snapshot, or unregister the host.

Source databases are represented with a database icon, while Echo Databases include a pink dot on the database icon to distinguish them. You can also hover over a database name to see whether it's a source database or an Echo Database.

Select the three-dot menu to the right of a source database name to clone the database, create a DB-Snapshot, or delete the database if it was created using Echo. If the database wasn't created in Echo, you can't delete it in Echo.

Select the three-dot menu to the right of an Echo Database to delete the database (you can't clone an Echo Database or create a DB-Snapshot from an Echo Database).

View DB-Snapshots

Snapshots created in Echo are called DB-Snapshots. To view your DB-Snapshots, navigate to your host, expand the host details if you're in **Cards** view, and then click the plus icon (+) to the left of the database name to expand the database resources.

Any DB-Snapshots on the database are listed in a table that shows the DB-Snapshot name and the creation time. Select the information icon to the right of a snapshot name to see the Db-Snapshot details, including snapshot name, creation time, database vendor, consistency type (application or crash) and all the databases the snapshot is included in. Click the three-dot menu to the right of the DB-Snapshot name to clone the database from the DB-Snapshot or delete the DB-Snapshot.

Creating and Managing Echo Databases

You can create an Echo Database at the host level, database level, or at the DB-Snapshot level. Databases created using snapshots are created almost instantly. You can also refresh an Echo Database to use the latest DB-Snapshot data.

Create an Echo Database at the Host Level

Clone at the host level to create one or more Echo Databases on one or more hosts.

Create an Echo Database in List View

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Echo **Manage Hosts** page, click the three-dot menu to the right of the host you want to create a database on, and then select **Create Echo Database**.

   The **Create Echo Database** pop-up opens.
3. Use the **Select source Databases** checkboxes to select which database(s) to clone, or click **Select all** to choose all the databases.
4. Click **NEXT** and then use the **Select Destination Hosts** checkboxes to select the destination host(s).
5. Click **NEXT** and then review the clone details to verify the source database and target host details are accurate.

   To rename a database (so it's easier to identify), click the edit icon to the right of the target host database name and make your updates inline.
6. Click **CREATE ECHO DATABASE**.

   Your database(s) is cloned, which takes about a minute. You can track cloning progress using the **Activity** panel on the right.

   After the Echo Database is created, a snapshot is automatically created on the source database.

Create an Echo Database in Cards View

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Echo **Manage Hosts** page, drag the database you want to clone from its source host to the target host.

   The **Create Echo Database** pop-up opens.
3. Review the clone details to verify the source database and target host details are accurate.

   To rename a database (so it's easier to identify), click the edit icon to the right of the target host database name and make your updates inline.
4. Click **CREATE ECHO DATABASE**.

   Your database(s) is cloned, which takes about a minute. You can track cloning progress using the **Activity** panel on the right.

   After the Echo Database is created, a snapshot is automatically created on the source database.

Create an Echo Database at the Database Level

Clone at the database level to replicate an existing database on one or more hosts. This feature only works for databases created outside of Echo (i.e., you can't clone an Echo Database).

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Echo **Manage Hosts** page, click the three-dot menu to the right of the database you want to clone, and then select **Create Echo Database**.

   The **Create Echo Database** pop-up opens.
3. Use the **Select Destination Hosts** checkboxes to select which host(s) you want to clone your database to.
4. Click **NEXT** and then verify the source and target details are correct.

   To rename a database (so it's easier to identify), click the edit icon to the right of the target host database name and make your updates inline.
5. Click **CREATE ECHO DATABASE**.

   Your database(s) is cloned, which takes about a minute. You can track cloning progress using the **Activity** panel on the right.

   After the Echo Database is created, a snapshot is automatically created on the source database.

Create an Echo Database at the Snapshot Level

Create an Echo Database on one or more of your hosts using an existing DB-Snapshot:

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Echo **Manage Hosts** page, click the three-dot menu to the right of the DB-Snapshot you want to clone, and then select **Create Echo Database from DB-Snapshot**.

   The **Create Echo Database from DB-Snapshot** pop-up opens.
3. Use the **Select source Databases** checkboxes to select the database(s) to clone your snapshot from.

4. Click **NEXT** and then use the **Select Destination Hosts** checkboxes to select the database(s) to clone your snapshot to.
5. Click **NEXT** and then verify the source and target details are correct.

   To rename a database (so it's easier to identify), click the edit icon to the right of the target host database name and make your updates inline.
6. Click **CREATE ECHO DATABASE**.

   Your snapshot is cloned as a new Echo Database. You can track cloning progress using the **Activity** panel on the right.

Refresh an Echo Database

You can refresh an Echo Database at the database level from an existing DB-Snapshot to use the latest snapshot data at any time. This feature lets you refresh data with minimal disruption to your database operations and doesn't require you to re-map hosts or applications. You can perform a refresh easily via the Silk Echo GUI or you can schedule automatic refreshes using the Silk Echo API.

To refresh an Echo Database using the Echo GUI:

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Echo **Manage Hosts** page, click the three-dot menu to the right of the Echo Database you want to refresh, and then select **Refresh Echo Database**.

   The **Refresh Echo Database from DB-Snapshot** pop-up opens.
3. Use the **Choose DB-Snapshot** screen to select which snapshot to use for the refresh. This screen shows any snapshots associated with the Echo Database.
4. Click **NEXT** and then click **CONTINUE** to confirm the refresh.

   The Echo Database is refreshed with the latest data from the DB-Snapshot. You can track refresh progress using the **Activity** panel on the right.

Creating Echo DB-Snapshots

Create a DB-Snapshot at the host level or the database level. You can only create DB-Snapshots on source databases (those created outside of Echo).

When creating a DB-Snapshot, you can choose whether it is application-consistent or crash-consistent:

- **Application-consistent snapshots** freeze and flush your IO to ensure all data and transaction logs are fully synchronized with Echo so they can act as a clean, consistent backup. These snapshots are useful for datasets that require no data loss during restore operations.
- **Crash-consistent snapshots** are essentially taken while IO is running, which means they might include uncommitted or in-progress transactions and partial logs. These snapshots are useful if you don't want to pause database operations while the snapshot is taken and for data that isn't mission-critical in case it requires a rebuild. In Echo, crash-consistent snapshots require all databases to be on the same VG in the DataPod.

By default, Echo Databases created in one step (using the drag-and-drop feature in Grid view) automatically use application-consistent snapshots.

DB-Snapshots are indicated by a camera icon with either a lightning bolt for crash-consistent snapshots or a check mark for application-consistent snapshots.

Create a DB-Snapshot

You can create a DB-Snapshot for any host or database created in Echo:

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. To the right of the host or database you want to create a DB-Snapshot on, select the three-dot menu and then select **Create DB-Snapshot**.

   The **Create DB-Snapshot** pop-up opens.

3. On the **Select DB-Snapshot Type** screen, select **Application Consistent Snapshot** or **Crash Consistent Snapshot**.
4. (For crash-consistent snapshots) Click **NEXT** and then use the **Select source Databases** screen to filter by VG and select one or more databases.
5. (For application-consistent snapshots) Click **NEXT** and then use the **Select source Databases** checkboxes to select which database(s) to clone, or click **Select all** to choose all the databases.
6. Click **NEXT** and then verify the source details are correct.
7. Click **CREATE DB-SNAPSHOT**.

   Your DB-Snapshot is created. You can track creation progress using the **Tasks** panel on the right.

   After the snapshot is created, it automatically displays under your database, along with any other snapshots on the database.

Deleting Echo Databases and DB-Snapshots

You can delete an Echo Database using either Silk Flex or the Silk API. However, source databases (the databases used to create Echo Databases) cannot be deleted using Echo.

DB-Snapshots can be deleted as long as they don't have any dependencies (i.e., as long as they weren't used to create any Echo Databases). If a DB-Snapshot has dependencies, the deletion option is grayed out to prevent you from deleting it.

Note: Resources created with Echo should **only** be deleted using Echo. Be aware that:

- DB-Snapshots (snapshots created with Echo) use a naming convention of **<snap_prefix>_<timestamp>**. To ensure Echo works as expected, these snapshots should **not** be deleted at the DataPod level.
- Echo Databases can be deleted via Echo but should **not** be deleted at the DataPod level, otherwise the system may experience unexpected behavior.

Delete an Echo Database

Deleting an Echo database removes the database mapping, deletes the database clones, and detaches the database from the hosts:

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Silk Echo **Manage Hosts** page, navigate to the database you want to delete.
3. To the right of the database, select the three-dot menu, and then select **Delete Echo Database**.
4. On the confirmation pop-up, click **DELETE**.

The database is removed from the host.

Delete a DB-Snapshot

You can delete any DB-Snapshots that don't have Echo Database dependencies:

1. Navigate to Silk Echo using the Silk Flex GUI or the Flex IP.
2. From the Silk Echo **Manage Hosts** page, navigate to the DB-Snapshot you want to delete.
3. To the right of the DB-Snapshot, select the three-dot menu, and then select **Delete DB-Snapshot**.
4. On the confirmation pop-up, click **DELETE**.

   The DB-Snapshot is removed from the database.

Echo Known Limitations

Limitation ID	Description
7	For crash consistent DB-Snapshots, all source database files must be in a single VG on the DataPod.
6	Creating crash-consistent DB-Snapshots across multiple VGs is unsupported. When using the API, ensure snapshots are within a single VG, as there is no validation in the API.
5	Databases created under the default DataPod VGs are unsupported by Echo.
4	Echo cannot clone a database from a host with a newer MSSQL version to an older one, (i.e., from MSSQL 2022 to 2019), but it can clone a database from a host with an older version to a newer one.
3	When creating an Echo Database from a source database in one operation, (e.g., via drag and drop from one host to another), there is no way to choose the **DB-Snapshot** consistency. These snapshots are automatically application-consistent. To create a crash-consistent DB-Snapshot, use the Echo Database creation wizard instead of a single operation.
2	In Grid view, it is possible to start the operation of creating an Echo Database from one DataPod to another, although this functionality is unsupported.
1	The system cannot validate naming rules, so if you use unsupported characters, you can't delete the Echo Database. Echo supports alphanumeric characters, underscores, and dashes. The maximum supported length for Echo Database names is 128 characters.