



# Things You Didn't Know You Needed to Know

The Data Dynamics Sales Engineering Team Responds to Frequently Asked Questions About StorageX

With any enterprise-class software purchase, there's bound to be complexity. That's just the nature of the beast when working with a product or product suite that works across your global environment, and that you rely on to function just how you expect. You and your team are going to have questions and concerns, and you may not always know what those questions are the first time you take a look at the software.

With that in mind, we reached out to the Data Dynamics Sales Engineering team to get their take on questions for which you're likely to want to know the answers, even if you haven't realized it yet. Use these questions as a stepping stone when starting your journey with **Data Dynamics StorageX**.

## | Setting Up StorageX

1

### **Is Microsoft SQL Server bundled with StorageX?**

The full version of Microsoft SQL Server is not bundled with StorageX. Smaller projects can leverage the free-of-charge Microsoft SQL Server Express deployment, but larger projects will require Microsoft SQL Server Standard/Enterprise. Licenses for SQL are provided by the customer.

2

### **Can SQL Server Express be upgraded?**

Yes, SQL Server Express can be upgraded to SQL Server Standard or Enterprise.

3

### **Is SQL Server installed locally on the StorageX server?**

SQL Server can be installed locally on the StorageX server, or you can configure your StorageX installation to connect to a managed SQL Server instance.

4

### **What are the SQL Server database requirements? How many tables?**

The StorageX installation will create the database and all its associated tables. Our product cannot use a pre-created database.

5 **What is required to allow StorageX to create the database?**

Our service account user must have database creator privileges, as with the **dbcreator** role, on the SQL instance that is provided.

6 **Can database creator privileges be removed?**

Yes, database creator privileges can be removed after installation. However, database creator privileges may be required for a future upgrade.

## | Configuring & Scaling Storage

1 **Is the StorageX product able to control the use of network bandwidth?**

Yes. Our Universal Data Engine (UDE) can be throttled to use only a specified amount of bandwidth.

2 **How does the StorageX product control policy threading?**

The StorageX UDE automatically creates a threadpool based on the number of processor cores that are provided on the UDE server.

## | Discovering Your Data

1 **What are the privilege requirements for scanning file systems?**

For best results, our product services run as a service account user.

- **For CIFS/SMB shares:**

This user account must be in the local **Administrators** and **Backup Operators** groups of the server hosting the filesystems to be scanned or migrated. For migrations, this requirement applies to both source and destination.

- **For NFS exports:**

The IP addresses of our core server and Universal Data Engines must have **Read/Write** and **Root** client privileges to the exports that are in scope for the scan or migration.

2 **Can a discovery scan be stopped and restarted?**

Yes, a discovery scan can be stopped and restarted.

3 **Does the discovery scan restart from a checkpoint?**

No; a restarted discovery scan will rescan the entire filesystem.

4 **How long does it take to scan a filesystem?**

This depends on many factors, including the performance of the storage platform and BAU workload that is currently running against that storage platform. Performance also depends on the specification of the Universal Data Engine that is used to run the policy. We've seen customers scanning upwards of **10-12 million files per hour per engine**.

5 **Does it take longer to scan larger files?**  
StorageX collects metadata, which is the same for every file, regardless of the file size. We collect **4.5 KB of metadata per file**. It's very fast.

6 **What is the overhead that is applied to the storage when scanning?**  
We've seen less than **1% CPU overhead** in our lab. This is dependent upon the number of policies being executed against the storage device and how much overhead is currently applied by BAU workloads.

7 **How does the product scale out to multiple PB?**  
Once the core services have been installed and configured, the product is scaled by adding Universal Data Engines and Metadata servers in a ratio of **3:1**. Universal Data Engines are typically installed at a ratio of 1 per 16 policies to be run, with a minimum of 1 per site with data to be scanned. Multiple engines deployed to a single site can be grouped for round-robin policy distributions. Every 3 UDEs usually requires 1 StorageX Metadata server.

## | Archiving or Replicating Your Data

1 **Which object storage platforms are supported for StorageX Archive and Replication?**  
StorageX supports archiving and replicating data to AWS S3, Azure Blob, Google Cloud Platform (GCP) S3, Hitachi Content Platform (HCP), NetApp StorageGrid, Dell Elastic Cloud Storage (ECS), IBM Cloud Object Storage (COS), and most other S3-compliant platforms.

2 **Instead of Archiving my data, can StorageX delete or cleanup old, cold, or ROT data?**  
StorageX does not directly delete old/cold data, but an Archive policy can be set up to "move" those files to an object storage bucket or container that has an ILM policy configured to delete everything after a certain period of time.

3 **Can StorageX report the list of archived files?**  
Running an Analysis query on the files to be archived will show a file list that can be exported in comma-separated value (CSV) format. Additionally, you can configure the Archive policy to save a manifest in CSV format each time the policy is run.

## | Retrieving Your Data

1 **When a file is archived, what does the end user see in the file system?**  
StorageX is a true archive. Files are hard-moved to the archive destination. We do not leave stubs or links.

2 **How are files retrieved?**  
Admin access to our **Retrieval Portal** can retrieve all files that have been previously archived to object storage. End user access to the Retrieval Portal can retrieve only the files for which they are the owner.

3 **Are files retrieved back to original location?**  
The administrator or end user can specify the retrieval destination.

4

#### Is all metadata retrieved?

The administrator or end user can specify ownership of the retrieved files.

## | Auditing & Reporting

1

#### Does your product report filesystem security?

StorageX will report on the users and groups that are owners of files on the access control lists and will also report where permissions are inherited or where inheritance is broken. You can also search for specific privileges of a user or group and can report on open shares (Everyone, auth user, Domain Users on the ACL).

2

#### Can StorageX report audit information?

StorageX can ingest the audit logs of NetApp and Isilon products. The audit information is then merged with metadata to show who/what/when for each file.

