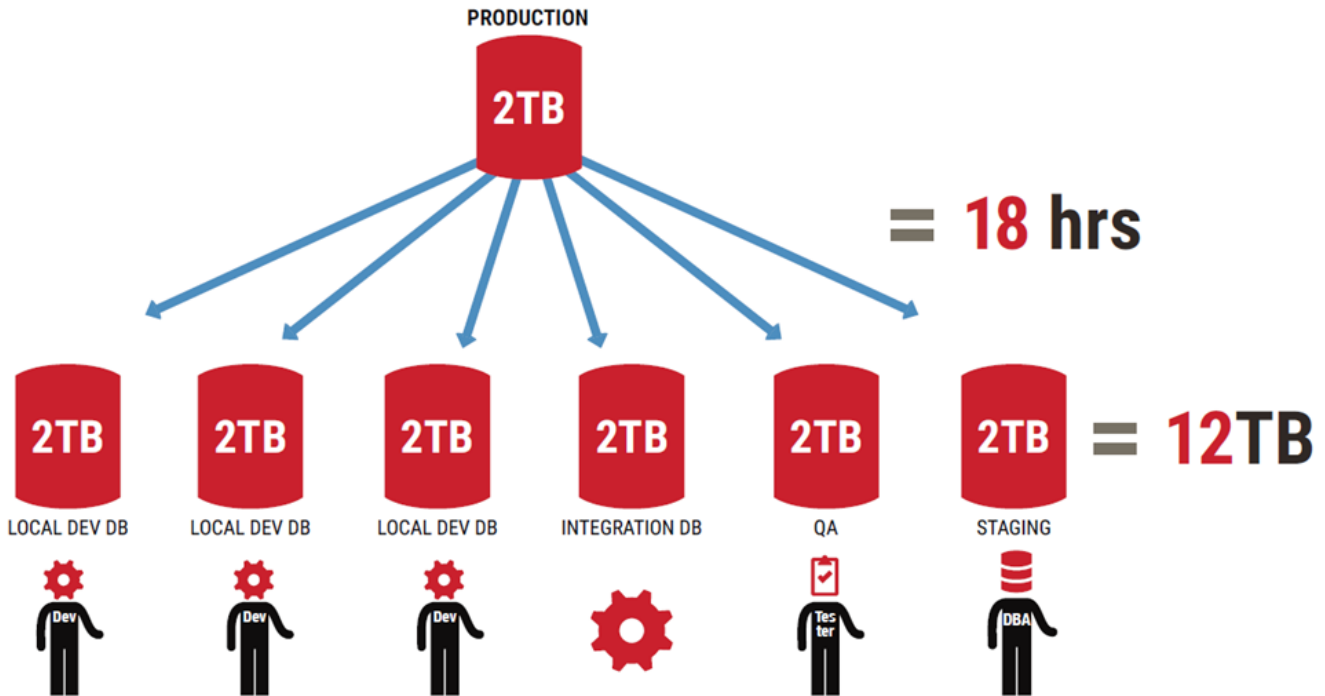


What is SQL Clone?

SQL Clone allows you to make editable copies of databases which can be created within seconds and which use very little additional disk space.

It is often useful to make copies of production databases available for development, QA, reporting, staging and UAT. These databases often contain large amounts of data, which means they can take many hours to restore when provisioning a new environment or performing a refresh, and each copy consumes significant amounts of disk space. Sometimes, an empty schema or sample data generated by SQL Data Generator is sufficient. But access to production data is usually necessary to investigate faults or generate offline reports, and developing and testing against it helps to identify edge cases and performance issues which may not otherwise appear until code is in production. This means that long restore times become inevitable, resulting in refreshes happening infrequently, and the number of copies is constrained by available disk space, making it more difficult to parallelize work.

Before SQL Clone - full disk footprint and long wait times to copy databases for multiple targets



SQL Clone makes it possible to provision copies of databases quickly by using a two-step process in which a snapshot file is created. A snapshot contains the source data from which clones can be derived. It can be created either from a live SQL Server database or from a backup, and is a full copy of the source database at one point in time. This takes about as long as a restoring a backup to create, and consumes as much disk space as a single restore.


Creating a snapshot from a live database requires SQL Clone to be running on the same machine as the SQL Server instance which contains the database. If the snapshot is created from a backup, SQL Clone can run on any machine with a copy of SQL Server which is capable of reading the backup.

Snapshot files can either be stored locally on the machine that created them, or shared with other machines by placing them on an Windows file share.

Once you have a snapshot file, you can create an unlimited number of clones from it on any number of machines which can access the file. Creating clones usually takes only a few seconds and a few megabytes of local disk space. They work just like normal SQL Server databases, and can be connected to and edited using any program. The changes made are specific to each clone and are persisted to a local diff file. The size of this file is proportional to the amount of changes which have been made. The rest of the data is accessed using the snapshot file.

The SQL Clone process

Step One - Create a snapshot of a live SQL Server database or backup



Step Two - Create clones from your snapshot on any machines which can access the file

