

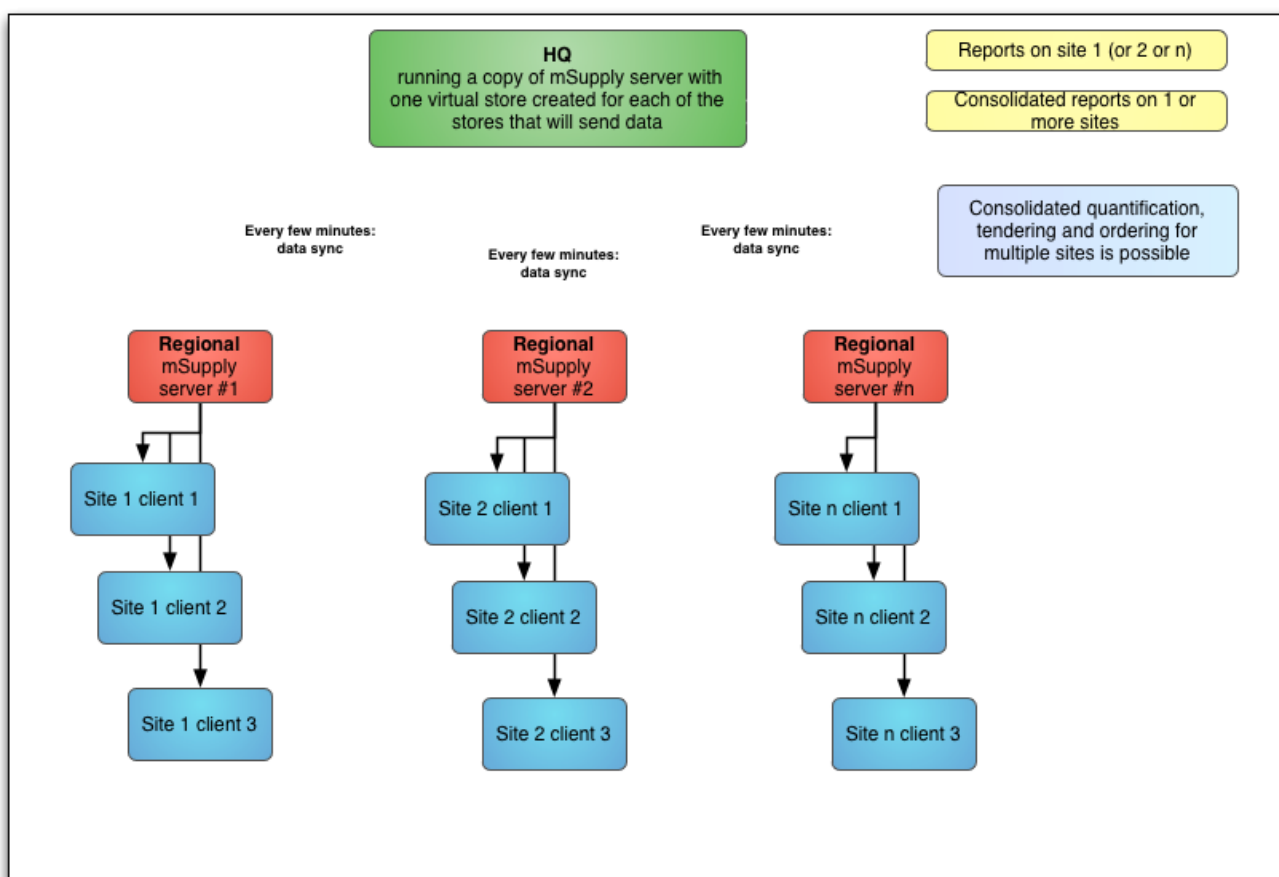
# 29.01. Introduction to Synchronisation

Updated: mSupply Version 7.13

mSupply's remote synchronisation allows mSupply running stores on local servers to send data to a central mSupply server for combined reporting.

## Synchronisation explained

It's like this:



## Definitions

- The **Central server** (a.k.a. **sync server**) is a central mSupply server. In any given setup, there will only be *one* central server, which must be running a *web server*, and it is responsible for:
  - aggregating the store data from all stores ⇒ enables centralised reporting, and a real-time backup
  - controlling how all of the stores sync with each other
  - routing transfers from one store to another
  - controlling the store records themselves, their preferences and **name** visibility in stores ⇒ all stores exist on the central server

- running the mSupply dashboard
- The **primary server** is a central mSupply server. In any given setup, there will only be **one** primary server, and it is responsible for:
  - controlling the central data which is common across all stores e.g. names and items
  - controlling the item visibility in stores ⇒ all stores exist on the primary
- A **remote site** is any mSupply server, single-user standalone mSupply, or mSupply mobile instance which is *not the central server*:
  - the syncing process is initiated from each remote site on a schedule, and pushes its own updated records to the central server first, before pulling any updated records from the central server
- A **satellite site** is any mSupply server, single-user standalone mSupply, or mSupply mobile instance which is *not the primary server*:

## Primary vs Central vs Mirror servers

- For most installations, the **Central server** and the primary server are one and the same, but for **mirrored** systems (see [Mirrored sync](#) below for details) they can be different - in that case the primary server is just a special case of a remote site.
- The **Central server**, and all of the remote sites are collectively referred to as **sync sites**.
- Each sync site has a unique ID and connection parameters, along with a list of the other sync sites with which it can send/receive synchronisation data. These are largely determined by inter-store visibility i.e. if store A (active on site A) is visible in store B (active on site B), then store B is automatically visible in store A, and store A on site A can send/receive transfers to/from store B on site B.

## Active vs Collector

- Each store must be **Active** on a *single* sync site, and that is where its own store data can be added, edited, and deleted.
- **Collector** copies of stores are *full read-only* copies of the store.
- Each store that is **Active** on a **remote site** also has a **Collector** copy of the store on the **Central server**. These can be logged into and reported on.
- The **Central server** can optionally also have **Active** stores on it.

## Sync and data integrity

- To preserve data integrity, and to avoid potential clashes where more than one sync site tries to modify the same record, *only one sync site can edit/update any specific type of data*. There are some special cases (see [Data types](#) below for details), but most data falls into one of the following main types:
  - **Central data** (a.k.a. **system data**) is common to all stores (e.g. suppliers, customers, items, master lists) - this can only be edited or imported on the primary server, and changes are synced to all sites
  - **Central store data** is data related directly to a store record e.g. store preferences, store tags and **name** visibility - this can only be edited or imported on the central server (the primary server for versions of mSupply up to v7.13), but changes are only synced to the sites where the store is *visible*
  - **Store data** (e.g. transactions, stock, locations) can only be edited or imported on the

sync site where the store is *active*, and changes are only synced to any site where the store exists as a *collector copy* (e.g. the central server)

- **Patient data** (e.g. patients, PMR records, insurance policies) can only be edited on the sync site where its *home store* is *active*, and changes are synced to any site where the patient is visible
- **Name data** is data related directly to a name (e.g. name tags, master lists) - this can only be edited or imported on the primary server, and changes are synced to all sites where the name is *visible*

## Mirrored sync

In a mirrored synchronisation system, the primary server becomes a remote site with responsibility for editing the central data:

- this remote site doesn't need to be online all the time, so that it can run operations for its own active stores and still be able to administer centrally-controlled data, without being a bottleneck for sync
- the central server (especially if it's in the cloud) can be always online and available for remote access to reports and the dashboard, and for processing and routing of sync data

The typical use case is where a site in a country that has the responsibility for administering items has poor internet (typically the Central Medical Stores), so that it either can't run the central server locally and/or can't reliably connect to a cloud-based central server. By using mirrored sync, this site retains authority, but a reliable central server elsewhere ensures that synchronisation with other sites is not hindered and that external users can still access the dashboard.

## Store setup

Stores to be synced will exist as separate instances of the same store on more than one sync site:

- on the central server - the store itself and its related preferences, along with name visibility in the store are editable ⇒ all stores and their transaction data must exist on the central server.
- on the primary server - item visibility in the store is editable ⇒ all stores must exist on the primary server



On a mirrored sync system, the store must be created first on the central server. It will then sync to the primary server, where item visibility and/or master lists can be configured.

- store data (e.g. store transactions, stock) can only be edited on a single sync site - the site on which the store is defined as *active* - and this store data is synced to the central server, as a read-only copy. This is controlled on the central server by assigning which site the store belongs to - see [sync sites](#).

Once a store has been set up (see the relevant parameters below), item and name visibility for the

store needs to be set up - see [Virtual stores](#), [Controlling item visibility](#):

- To do that, once you log in to the central/primary server, you need to switch to the new store - see [Switching store](#)
- To do that, you need to have permission to log in to it - see [Managing users](#) (It's a good thing

this doesn't happen often! 😬 )

## Data types

The table below is a high-level summary of the different data types:

- **Central data** can only be edited or imported on the primary server (or sometimes the central server), and changes are synced to all sites
- **Central store data** can only be edited or imported on the central server (or sometimes the primary server), but changes are only synced to the site where the store is visible
- **Store data** can only be edited or imported on the sync site where the store is *active*, and changes are only synced to any site where the store exists as a *collector* copy (e.g. the central server)
- **Patient data** can only be edited on the sync site where its *home store* is *active*, and changes are synced to any site where the patient is visible
- **Name data** can only be edited or imported on the primary server, and changes are synced to all sites where the name is visible
- **Local data** can be edited or imported on any site but doesn't sync anywhere
- **Sync data** can only be edited on the central server but doesn't directly sync anywhere
- **Message data** can be created on any site, but can't be edited anywhere, and syncs according to the sending/receiving store:
  - if the receiving store is set, then the message will be processed on the site where that store is active
  - if the receiving store is blank and the sending store is set, then the message will be processed only on the central server
  - if both are blank, then the message will be processed on all remote sites (i.e. all except the central server)
- Some data can fall into more than one type, depending on the situation.

Data	Type	Editable	Notes
Items	Central	Primary	Including item-related data e.g. item categories, units, BOM masters
Names (except patients)	Central	Primary	Including name-related data e.g. name categories, contacts
Merging of items, units and names (except patients)	Central	Primary	
Groups and departments	Central	Primary	
Item master lists and programmes	Central	Primary	
Budgets, periods and accounts	Central	Primary	
Transaction categories and payment types	Central	Primary	
Purchase order categories	Central	Primary	

<b>Data</b>	<b>Type</b>	<b>Editable</b>	<b>Notes</b>
Custom data	Central	Primary	
Barcodes	Central	Primary	Can be added on any site; duplicates are automatically merged when synced to the primary
Currencies	Central	Primary	
Options and properties	Central	Primary	
Location types	Central	Primary	
Regimens and indicators	Central	Primary	
Drug interactions and warnings	Central	Primary	
Vaccinators and vaccine settings	Central	Primary	
Custom reports	Central	Primary	Standard reports are regenerated on each upgrade
Asset settings	Central	Primary	
Regions	Central	Primary	
Incoterms and tender conditions	Central	Primary	
Abbreviations and item directions	Central	Primary	Dispensary data
Diagnoses	Central	Primary	Dispensary data
Insurance providers	Central	Primary	Dispensary data
Patient event types	Central	Primary	Dispensary data
Purchase orders (centralised)	Central store	Primary	
Tenders and quotes (centralised)	Central store	Primary	
Payments (centralised)	Central store	Primary	
Visibility of items	Central store	Primary	
Visibility of names (including existing patients)	Central store	Central	
Visibility of existing prescribers	Central store	Central	
Stores and non sync-related store preferences	Central store	Central	
Sites and sync-related preferences	Central	Central	Changes on the central server indirectly update related records on remote sites
Dashboard reports	Central	Central	
Authorisers and authorisation	Central	Central	
Messages	Message	Store	Depends on sending and/or receiving store
Visibility of new patients and prescribers	Patient	Store	New visibility records sent to central server
Patients and prescribers	Patient	Store	Including patient-related data e.g. PMR, insurance policies
Merging of patients and prescribers	Patient	Store	Both records must have the same home store
Patient events	Patient	Store	Dispensary data
Repeats	Patient	Store	Dispensary data; preference can be set to allow processing on all sites where the patient is visible
Prescriptions	Patient	Store	Preference can be set to sync to all sites where the patient is visible

<b>Data</b>	<b>Type</b>	<b>Editable</b>	<b>Notes</b>
Adverse drug reactions	Patient	Store	Dispensary data
Name-related tags, budgets and master-lists	Name	Primary	
Name notes	Store	Store	
Customer stock history and requisitions	Store	Store	
Locations	Store	Store	
Merging locations	Store	Store	
Stock and replenishments	Store	Store	
Stocktakes and inventory adjustments	Store	Store	
AMC projections	Store	Store	
Transactions (but not prescriptions)	Store	Store	Including other transaction-related data e.g. backorders, builds
Transaction notes	Store	Store	
Item notes	Store	Store	
Boxes	Store	Store	
Goods received	Store	Store	
Indicator values	Store	Store	
Vaccine monitors/sensors	Store	Store	
Assets	Store	Store	
Store credentials	Store	Store	
Tenders and quotes	Store	Store	Except for centralised procurement
Purchase orders	Store	Store	Except for centralised procurement or supervisor-mode ordering
Payments	Store	Store	Except for centralised payments
New users	Store	Store	New user records sent to central server
User licenses and existing users	Local	Local	
User permissions	Local	Local	
Preferences (non-store)	Local	Local	Except for a few special cases which are explicitly synced
Reference documents	Local	Local	
HIS	Local	Local	
Drug registration	Local	Local	
Labels	Local	Local	
Logs	Local	Local	
Reminders	Local	Local	

## Centralised procurement

If this preference is switched on in the primary site, then purchase orders can be entered for any store on the primary, regardless of its sync type (i.e. not just for active stores).

If it is switched on in a satellite, then purchase orders cannot be entered for any store on the satellite, even if it is active.

When this preference is changed on the primary, the change will propagate to all the satellites as well.



It is still possible to edit the preference on the remote (e.g. to allow local procurement for stores active on that site), but if that is the case, then the related permissions for purchase orders for those stores on the primary will need to be disabled manually in order to prevent purchase orders for those stores being editable on both sites.

## Dispensary data

### Stock transactions

By default, prescriptions and any other operations in dispensary mode which affect stock levels are not synced to the sync server, unless the *Include transactions in sync* option is enabled in the store sync preferences, as these can potentially generate a lot of sync traffic and there is usually no need for the central users to have this level of detail - the same site will usually have another store which supplies the dispensary and its transactions will be synced, and that's usually enough to infer usage for the dispensary store.

If this preference is switched off in a dispensary store on a remote site, the following data will not be synced back to the central sync server:

- Transactions (including backorders, builds etc.)
- Prescriptions
- Stock lines
- Stocktakes

If the preference is switched on, all of the store's transaction & stock data will be synced to the sync server.

### Patients

Patients in mSupply are a special kind of customer, but for the purposes of synchronisation, we need to treat them differently - more like store-specific data rather than system data. This is because patient data is often locally controlled rather than centrally.

Since mSupply v4.09, patients have a *home store*, initialised either according to the store they were created in, or according to their most recent prescription. Patients and their related data (patient medication records, repeats, insurance policies) can only be edited in an active dispensary store on the same *home site* where their home store is active.

Patients and their related data records will be synced to the primary server, and forwarded on to any other site which has an active/collector store in which the patient is visible. Prescriptions will be allowed in dispensary stores on any of those sites, except for repeats which can only be processed in dispensary stores on the patient's home site.

Newly created patients will also be made visible in any other dispensary stores on the home site, depending on the store's visibility preferences. Subsequently, patient visibility is controlled from the

primary server in the same way as for other name (customer & supplier) records.

## Other dispensary data

Abbreviations, item directions, insurance providers and patient event types are treated as a special kind of system-specific data i.e. they are editable on the primary, but only synced to sites having at least one active dispensary store.

Prescribers are also treated as a special kind of store-specific data, like patients i.e. they are editable in any active dispensary store on their home site, synced back to the primary server, and shared with dispensary stores on other sites where they've previously prescribed.

## Transferring patients/prescribers to a different home store

Watch this space...

## Transfers

Transfers occur when there are two stores involved in a transaction, and includes stock transfers, requisitions (from a mobile store) and internal requisitions (from another desktop store). In a syncing system, very often these two stores are not active on the same site, so there has to be extra processing to ensure that both halves of the transaction are synced to both the initiating store/site and the responding store/site.

Both stores need to exist on both sites, and usually the initiating store is a transfer store on the responding site, and vice-versa. As of mSupply v3.83, unless both stores are active on the same site (in which case, everything can be done locally), the bulk of the processing is done on the central sync server when it detects the initiating half of a transfer transaction. In simple terms, the logic is something like this:

1. when the sync server detects the initiating half of a transfer transaction
  1. it creates the responding half of the transaction, but with a dummy invoice/serial number of -1
  2. it ensures that both halves of the transaction are synced to both initiating and responding sites
2. when the responding site receives the responding half of the transaction
  1. it assigns the next invoice/serial number for the store and sends that back to the sync server
  2. it also creates a log message for the initiating half of the transaction and sends that back to the initiating site via the sync server
3. any subsequent changes to the initiating half of the transaction (usually very limited in scope) on the initiating store/site will be synced back to the sync server (according to the normal sync rules)
4. any subsequent changes to the responding half of the transaction on the responding store/site will be synced back to the sync server (according to the normal sync rules)

## Stock transfers

- This is where the customer in a customer invoice is another store
- The initiating half of the transaction is where the customer invoice is finalised in the initiating store/site
- The responding half of the transaction is a supplier invoice (on hold by default) in the responding store/site

## Mobile requisitions

- This is where the supplier in a supplier (aka request) requisition is another store
- The initiating half of the transaction is where the supplier requisition is finalised on mobile
- The responding half of the transaction is a response requisition in the responding store/site
  - this shows up as a customer requisition on mobile if the responding store is active on another mobile site
  - or as a response requisition if the responding store is active on a desktop site
    - the user can create one or more customer invoices to fulfil the requisition
    - these customer invoices (when finalised) will generate corresponding supplier invoices back in the initiating store as stock transfers

## Internal requisitions

- This is where the supplier in a purchase order is another store, and is a two-stage process
- The initiating half of the first transaction is where the purchase order is confirmed in the purchase order store/site (i.e. where the purchase order is editable)
  - if centralised procurement is enabled, this will be the primary site, otherwise it will be the normal initiating store/site (i.e. where the purchase order's store is active)
- The responding half of the first transaction is a customer invoice in the responding store/site
- The initiating half of the second transaction is where the customer invoice is finalised in the responding store/site
  - note that adding extra customer invoice lines before the invoice is finalised will create corresponding new goods received lines
- The responding half of the second transaction is a goods received record in the initiating store/site (i.e. where the purchase order's store is active)
- If centralised procurement is enabled:
  - the received quantities for the original purchase order lines will be updated on the primary whenever the corresponding goods received lines are received there (and forwarded from there to the initiating store/site)
  - when the primary receives any subsequent updates to goods received lines from the initiating store/site, it will update the quantities in the corresponding purchase order lines
- If centralised procurement is not enabled, the received quantities for the original purchase order lines will be updated in the initiating store/site when it receives the goods received lines from the primary

## Reporting

- On the sync server if you login in supervisor mode, you can then run reports on one or more

stores. The reports can answer questions such as:

- How much stock on hand of item X (or all items) are there at each location?
- What is the value of stock on hand across the whole system?
- How many of item X is being used each month at each location?
- When the synchronisation system is turned on, a new **Special > Users report... > Sync report...** menu item is available. Choosing this will create a report showing the last time each of your sync stores connected to the central sync server. Those that connected more than a month ago will be highlighted in red.

## Requirements

### Synchronisation Server

The Sync server module is the component that controls all the logic of the sync system described in this chapter. This module is priced separately - refer [Pricing](#)

### Web Server

Any communication through the web to an mSupply central server requires the [Web Server Module](#). This module is priced separately - refer [Pricing](#)

### Internet connection

Each sync site requires an internet connection. This doesn't have to be on all the time for each satellite server, but at least an hour or so per day or per week, depending on the transaction volume, how often primary server data needs to be updated, and the speed and latency of the internet.




Obviously, the internet needs to be on at the central server at the same time as any other sync site that it will communicate with. For this reason, the central server needs to have high availability, and so in most cases, will be cloud hosted.

## How to tell if synchronisation is happening

**Note:** the synchronisation system can be disabled completely or paused in the preferences (see the [16.01. General preferences](#) page for details).

### The Manual Sync Button



If you click the “2 circular arrows” icon at the bottom it will initiate an immediate synchronisation and update the statistics on the number of records remaining as sync progresses.

## On a remote site

If there are queued sync records, there will be a message at the bottom left of the navigator which shows the number of sync records in the queue. If this is not going down, there may be an entry in the log.

## On the sync server

- Choose **Help > About mSupply**
- At the bottom left of the window there is a list of tables and the number of records in each table.
- Scroll the list to near the bottom, and you will see the number of records in the **sync\_out** table.
  - If the number is zero, your copy of mSupply is up to date
  - If the number is growing from day to day, there are possible reasons:
    - you need to provide more internet time or faster internet
    - there may be a problem that needs the attention of Sustainable Solutions. Before contacting us, make sure the internet is connected for an hour, and see if the number is decreasing or not.



- administer system data?
  - otherwise, it's better to have a separate local primary server and cloud sync server
- 2. Do a site survey for all the proposed stores
  - how is the internet access and local network?
  - how reliable is the power supply?
  - is there a suitable location for mSupply?
- 3. Decide which version of mSupply is best for each store - don't worry, this can be changed later and licenses can be transferred
  - central primary server: mSupply server with local clients (depending on the number of users)
  - central sync server: mSupply server with local clients (depending on the number of users) and web server
  - multi-user central/district warehouse/busy dispensary: mSupply server with local clients (depending on the number of users)
  - single-user store/quiet dispensary (desktop): mSupply single-user desktop
  - small store/clinic: mSupply mobile
- 4. Decide on the heirarchy of the stores you want to include
  - which stores supply other stores?
  - can stores which have the same supplying store transfer to each other, or just with the supplying store?
  - for dispensary stores, do you need to see prescription details on the central server
- 5. Decide on the users for each site and their roles/permissions

## Moving stores between sync sites



Misconfiguring of store or site settings can corrupt data across the synchronisation system. This should only be done by someone trained to do so and with guidance from the mSupply team



How mSupply handles sync sites was 'refactored' and released in [mSupply v5.07 \(2022-03-22\)](#). This refactoring was significant and has resulted in a simpler and more strictly controlled way of managing sync sites. Rather than use the methods documented in this section, it is **highly recommended** to use the new tool as documented in [Synchronisation Sites](#).

## Moving a store from the central server to a new sync site

If you are using [mSupply v5.07 \(2022-03-22\)](#) or later, it is **highly recommended** to use the new tool as documented in [Synchronisation Sites](#), but if you insist on going 'old school' you can use the [site wizard](#).

## Moving a store from one sync site to another

To move a store from one site to another means to make it *not* **Active** on the 'From' site, and **Active** on the 'To' site. By default, mSupply will leave the store configured to maintain a full **Collector** copy on the 'From' site with all the associated sync traffic. This is not normally desired. Except for this process of determining which site a store is **Active** on, all sync settings are governed by [store visibility settings](#).

If you are using [mSupply v5.07 \(2022-03-22\)](#) or later, the methods described here will not work. You will have to use the new tool as documented in [Synchronisation Sites](#).



This process involves editing store sync settings which is dangerous. Editing store sync settings must happen on the central server and is restricted (by password) to Sussol technicians. Please consult with us to coordinate this. The full process is documented here for your information.

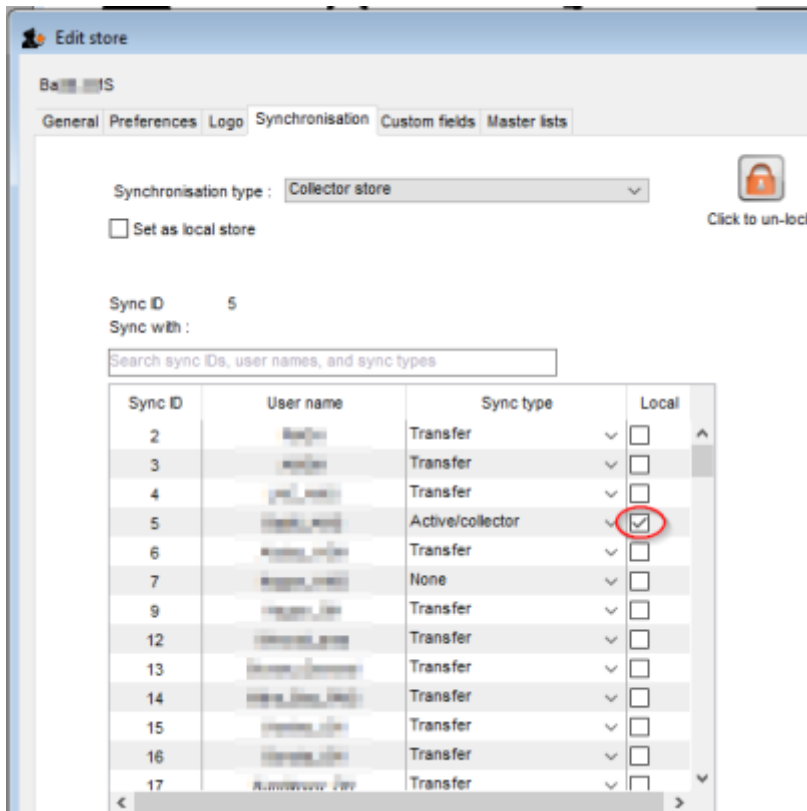


- A mobile site can only have one **Active** store.
- Therefore, only one of the stores on a 'From' desktop site can be moved to a single mobile site. If another **Active** store on the 'From' desktop site needs to be moved to a mobile site, then it will need to be moved to another sync site, mobile or desktop.

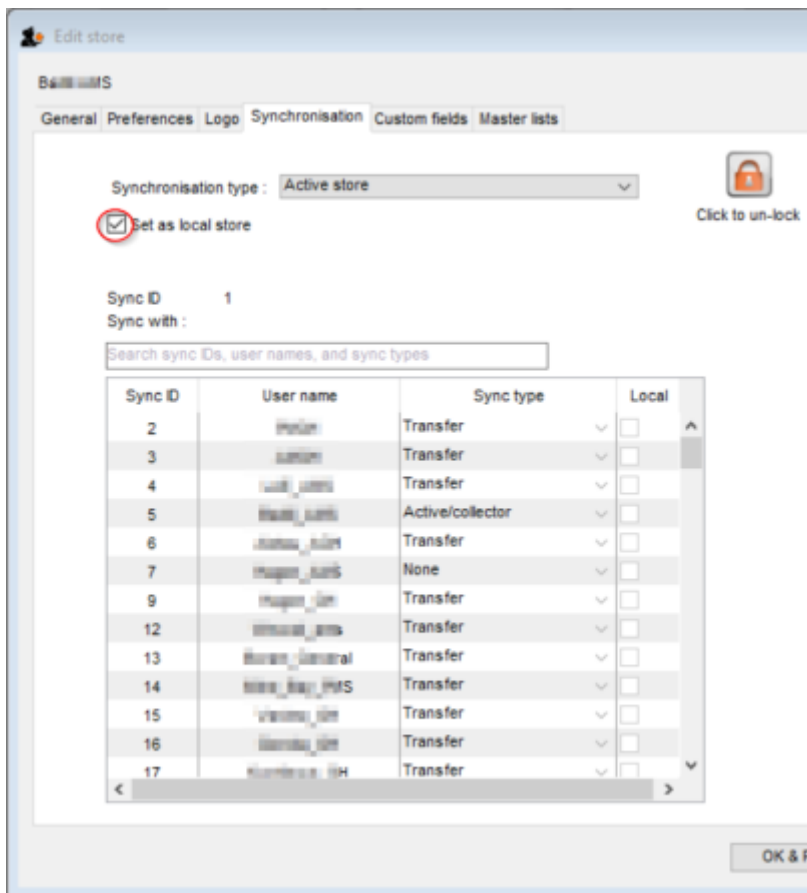


If the 'To' store is the central server, then this same procedure applies

1. Ensure that both the 'From' site and the 'To' site (only relevant if it is a desktop site) are fully synchronised.
2. Ensure that no new data is entered on either site until the move is complete. Therefore, best to plan for this process to happen overnight or over the weekend...
3. On the central server, **Special** > **Stores** > double-click the store that is being moved > **Synchronisation** tab > **Click to un-lock**, enter code to edit (**Sussol technicians only**)
4. If the 'To' site is a desktop site that already exists, then make the store **Active** on the 'To' site by finding the 'To' site in the list and click on the **Local** checkbox.



5. If the 'To' site is the central server, or doesn't yet exist, make the store **Active** on the central server by clicking on the **Set as local store** checkbox



6. Set the sync type for this store on the 'From' site to whatever you want:
1. **None** if the 'From' site only had this store on it. You will likely also want to delete the sync site - see step 12 below.
  2. **Transfer** if this store will likely need to transfer to / from another store on the 'From' site
  3. **Collector only** if a full copy of the store is to be maintained on the 'From' site. This is not

normally desired and is not recommended!

7. Click **OK** to save changes
8. Click **OK** to exit the stores list
9. If the 'To' site is the central server, skip to step 12.
10. If the 'To' site doesn't yet exist, use the [site wizard](#) to create the new sync site and move the store to it.
  1. If the 'To' site is a *new* desktop site, create the data file using [Sync site export](#) and [Sync site import](#).
  2. If the 'To' site is a *mobile* site, then get it initiated as discussed in [Installation of mSupply Mobile APK and initiation](#).  
The store should now be visible on the on the new site, in which case it is safe to operate it.
11. If the 'To' site is an *existing* desktop site, then code needs to be run on the Central server to generate sync out records for that store - refer [Sussol internal reference](#).
12. If the 'From' site now has no stores on it (it was a mobile site or a desktop site with only one store on it), then the sync site can be safely deleted. (**Sussol technicians only**)



If you are moving the store from **Mobile** to **Desktop** or vice versa, this process will not change the [site type](#). As of 2020-10-02 this must be done by Sussol technicians only by use of the record browser.

Previous: [28. Frequently Asked Questions](#) | | Next: [29.02. Synchronisation Sites](#)

From:  
<https://docs.msupply.org.nz/> - **mSupply documentation wiki**

Permanent link:  
<https://docs.msupply.org.nz/synchronisation:introduction?rev=1743522943>

Last update: **2025/04/01 15:55**

