

	Programs
Client	Client (eg XenCenter)
xapi R	Receiving XCP host
xapi T	Transmitting XCP host
SM	Storage Manager
DRBD	Manages block replication

Within pool migrate

Cross-pool migrate

Storage motion

Phase	Action	Notes
1. Connection	<ol style="list-style-type: none"> 1. Client -> xapi R: session.login_with_password(username, password) <- s1 2. Client -> xapi R: VM.receive(s1, SR) <- URI list 3. Client -> xapi T: session.login_with_password(username, password) <- s2 4. Client -> xapi T: VM.migrate(s2, vm, URI list, VDI to SR map) 5. xapi T -> xapi R: HTTP HEAD URI (with empty content) 	<p>URIs encode different addresses, protocols (https/http) URIs encode a capability/token</p> <p>Each URI is tried until one successfully returns 200/OK HTTP HEAD run every 30s as heartbeat XenAPI client can cancel either transmitter or receiver</p>
2. Preparation	<ol style="list-style-type: none"> 1. xapi T -> xapi R: HTTP PUT URI/metadata (with VM metadata export) 2. xapi T -> xapi R: HTTP GET URI/disk/0 3. xapi R: SMAPI VDI.get_replication_target(0) <- http://...../ 4. xapi T -> xapi R: HTTP GET URI/disk/1 5. xapi R: SMAPI VDI.get_replication_target(1) <- http://...../ 6. xapi T -> xapi R: HTTP PUT URI/memory 7. xapi R: SMAPI VDI.attach 	<p>RESTful API = easy to test VM created "hidden" from XenCenter Always 1-1 correspondence between domains and VMs If SRs can't be found, fresh VDIs are created</p> <p>VDIs are attached before receiving replication data</p>
3. Bulk disk transfer	<ol style="list-style-type: none"> 1. xapi T: SMAPI VDI.replicate(0, http://...../, new_ref) 2. SM: does nothing if both arguments identify the same disk 3. SM: detach tapdisks from blktap devices 4. SM: spawn new tapdisks to manage the underlying .vhd/LUNs 5. SM: configure DRBD /dev/drbdN devices 6. SM: attach AIO tapdisks to /dev/drbdN devices 7. xapi T: SMAPI VDI.get_state(new_ref) until the mirror is established 	<p>new_ref allows xapi to query state of the operation</p> <p>guest data path disconnected</p> <p>guest data path reconnected</p>
4. Memory transfer	<ol style="list-style-type: none"> 1. xapi T -> xapi R: Vmops.suspend data 2. xapi T waits for domain to finally suspend 3. xapi T: SMAPI VDI.deactivate(0) 4. xapi T: wait for all VDI.replicate Tasks to enter "finished" state 	<p>VDI.deactivate triggers disk copy to stop iterating</p>
5. Completion	<ol style="list-style-type: none"> 1. xapi T -> xapi R: HTTP PUT URI/control (with "done") 2. xapi R: resynchronises VM powerstate (Halted, Suspended, Running) 3. xapi R: SMAPI VDI.activate 4. xapi T: destroy domain 5. xapi T: SMAPI VDI.detach 	<p>DRBD becomes primary then shuts down</p>