Sendense

TECHNICAL GUIDE

Sendense replication prerequisites

Prepare CloudStack and VMware sources, and a CloudStack recovery site, for Sendense replication.

Companion to the video *Set up Sendense for replication*: https://sendense.com/guides/set-up-sendense-for-replication

Document control	Details
Revision	2.0, 1 November 2026
Applies to	Sendense SHA 4.3.685 and later
Scope	Single-tenant replication from Apache CloudStack (ACS) to ACS, and from VMware to ACS
Audience	Infrastructure, network, platform and application teams
Companion video	*Set up Sendense for replication* (sendense.com/guides/set-up-sendense-for-replication)

Purpose and intended result

Prepare your infrastructure for replication from Apache CloudStack (ACS) to ACS and from VMware to ACS. The intended result is incremental replication into a CloudStack recovery site, proved by a recovery into an isolated test network.

The companion video, *Set up Sendense for replication*, walks through the setup step by step. At each step it names the section of this guide that applies, so prepare the items below before you start.

What to prepare, and when

Before you start	During setup
Reserve appliance addresses and capacity; prepare platform accounts, networks, firewall access, offerings and KVM host access.	Set up the SHA, import the templates, deploy and enrol the SNAs, and install the host agents (SNAgent) on the KVM hosts.
Select the test VMs, and arrange guest access, recovery network mappings and application checks.	Confirm the final replication ports, create the replication patterns and their Controllers, and test the first full sync, incremental syncs, an isolated test failover and its rollback.

Reading order

Sections 1 to 3 cover appliance addresses and network access. Sections 4 to 6 cover CloudStack, VMware and the test workloads. Section 7 is the checklist to complete before you start, and the setup steps in order.

1 Appliance plan and IP addresses

The reference layout is one Sendense Hub Appliance (SHA) at the recovery site, an ACS source Sendense Node Appliance (SNA), a VMware source SNA and an ACS destination SNA. Reserve four stable appliance addresses, one for each, reusing suitable existing allocations where appropriate. Any additional SNA needs its own address.

Appliance address worksheet

For each appliance, record the network or port group, IP address and prefix, gateway, DNS servers and hostname. Use a static address or a DHCP reservation, and confirm it is free. All SNAs and KVM host agents must resolve and reach the SHA name; the SHA certificate must match that name.

Appliance and placement	Network or port group / VLAN	IP / prefix	Gateway	DNS servers	Hostname (FQDN for the SHA)	Static or DHCP reservation
SHA, CloudStack recovery site						
ACS source SNA, source CloudStack site						
VMware source SNA, VMware source site						
ACS destination SNA, CloudStack recovery site						

CloudStack-managed networks must allocate or reserve the selected address on the VM NIC. On an L2 network, use your external IP address management process to keep the address free. Sendense applies the selected addressing during deployment.

Controller DHCP is a separate requirement

Provide working DHCP on the Controller network with at least one free lease per protected VM, plus operational headroom. Sendense creates one Controller VM per protected VM. The four appliance reservations above do not replace this DHCP scope.

Starting capacity for a first deployment

Component	CPU and memory	Storage and capacity to reserve
SHA	16 vCPU, 16 GB RAM	SSD system disk; keep the supplied image disk size or increase it. Allow space to stage the complete VMware OVA, plus at least 4 GiB free reserve.
Each SNA	Recommended start: 4 vCPU and 8 GB RAM	Use the supplied template disk size. Reserve resources in the site the SNA will serve.
Each Controller	Default standby: 2 vCPU and 4 GB RAM	Appliance root disk, replica disks and checkpoint space. Reserve the protected workload's CPU and memory for failover.

Place the controlling SHA at the recovery site, so recovery management remains available if the source site is lost. These are initial allocations for a small test set. Confirm sizing against disk concurrency, change rate, storage performance and the required recovery point objective (RPO).

Have the CloudStack zones, account, network IDs and offering IDs available. For VMware, identify the datacenter, cluster or resource pool, datastore, VM folder and port group for the SNA.

Further detail: Sendense Controllers · SNA enrolment

2 Platform access for deployment

Set up the SHA first, then import the templates and deploy the SNAs from it. Have the supplied SHA image available; if CloudStack imports it by URL, the SSVM must reach that image URL and port. Prepare the routes below from the reserved SHA address, plus the replication routes in section 3.

Access must cover deployment as well as discovery

Provide accounts that can import or register templates, deploy and configure appliances, attach deployment media and manage their power state. A successful login or inventory test does not prove these write permissions. Platform-specific access is listed in sections 4 and 6.

Connection initiator	Destination	Port and protocol	Why it is needed
SHA	vCenter	TCP 443	Inventory, template import, SNA deployment and lifecycle operations.
SHA	ESXi upload endpoints selected by vCenter	TCP 443 HTTPS	Appliance disk upload during OVA/template import. Include every host eligible for the deployment.
SHA and the site SNA	Each relevant ACS API	TCP 443, or the actual API port	Template registration, SNA deployment, discovery, replication and failover. A directly exposed HTTP API commonly uses 8080.
CloudStack SSVM	SHA template and enrolment-media download address	TCP 80 by default, or configured HTTPS	CloudStack downloads the appliance image and seed ISO. Use the actual URL selected for the import.
Administrator workstation	SHA	TCP 443	GUI and API access.
Appliances, host agents and SSVM	Their DNS servers	UDP and TCP 53	Resolve the platform, appliance, download and update endpoints each system uses.
Appliances and host agents	Time service	UDP 123 (NTP)	Synchronised time for certificates and operations.
Controller network; appliance networks if using DHCP	DHCP service or relay	UDP 67 and 68	Address allocation and renewal. Allow the DHCP/relay path for the selected networks.

Verify the complete image transfer path

The SHA must resolve vCenter and the ESXi hostnames returned for upload, and reach both on TCP 443. VMware replication disk reads use TCP 902 from the source SNA, as shown in section 3. These are different paths.

CloudStack's Secondary Storage VM (SSVM) must resolve and fetch the chosen template and seed URLs. Check its own DNS, routing and egress rules. For HTTPS, its Java trust store must accept the certificate. A

download from a laptop or from the SHA does not prove SSVM access. If a site SNA is selected to serve media, the SSVM needs that SNA's configured download endpoint, normally HTTP TCP 8081.

Outbound services from the SHA

Allow outbound TCP 443 to licensing.sendense.com, registry.sendense.com and downloads.sendense.com. Allow ghcr.io and its required authentication and content-delivery endpoints for updates. Optional remote support uses diag.sendense.com on TCP 443.

SNAs receive updates through the SHA, and Controllers use their local service path. They do not need direct Internet access for a template-based deployment. KVM hosts may need their package mirrors when the SNAgent installer adds dependencies.

Further detail: CloudStack deployment checklist · Firewall ports

3 Replication firewall rules

The initiator column shows which system opens the connection. Allow established return traffic. Scope rules to the reserved appliance addresses, participating hosts and Controller subnet. In this layout the destination SNA handles Controller management and heartbeat relay.

Connection initiator	Destination	TCP port	Purpose
Every SNA	Controlling SHA	443	Enrolment, control and health; also intersite data when the route is Via SHA.
Source SNA	Destination SNA	Assigned ingress range	Direct intersite replication. The range is allocated when the destination SNA enrols.
Destination SNA	Controller subnet	8080	Controller management, disk preparation, verification and guest adaptation.
Destination SNA	Controller subnet	10810 upward	Replica disk writes; one successive port per replica disk.
Controller subnet	Destination SNA	8081	Controller heartbeat relay and local downloads; a separately initiated connection.
VMware source SNA	vCenter	443	Inventory, snapshot, CBT and other source API operations.
VMware source SNA	Every eligible ESXi management address	902	VMware NBD/NFC disk reads; include DRS and vMotion destinations.
KVM hosts with SNAgent	Their SHA	443	Host-agent registration, health, configuration and updates.
KVM hosts with SNAgent	Other participating KVM hosts	8888	Authenticated host-to-host coordination; allow initiation in either direction between participating peers.

The SNA-to-SNA ports are known after enrolment

The destination SNA receives its own replication ingress range when it is enrolled, so your network administrator applies the Direct P2P rule at that point. It is not a fixed product-wide range and cannot be completed in advance for a new SNA.

The recorded value is shown under DenseView › Appliances › destination SNA › Port allocations › Restore slots. The API equivalent is `GET /api/v1/appliances/{id}/ports`, using `nbd_restore_range.start` through `nbd_restore_range.end`. Recheck it if the appliance is replaced or enrolled again.

Controller data ports depend on disk count

The first Controller data port is TCP 10810. The last is 10809 plus the largest replica-disk count of one protected VM. For example, ten replica disks need TCP 10810–10819. This local range is separate from the destination SNA's assigned ingress range.

Choose the intersite route

Direct P2P uses your private interconnect or VPN. Sendense does not add encryption to the direct stream, so use an encrypted VPN where the link must be encrypted. Via SHA carries intersite data in the encrypted SNA connections on TCP 443 and requires no direct SNA-to-SNA rule. Both routes still require all three local Controller paths above.

Automatic can fall back to Via SHA. Check the actual route during a live sync; success alone does not prove the direct path. SNA data listeners are temporary, so testing an idle listener is not a reliable firewall check.

Further detail: Replication network planning · Replication route selection

4 CloudStack access and KVM source hosts

Prepare the following in each relevant CloudStack environment. The credentials are entered and tested, and the Sendense settings configured, during setup. The host-agent requirements apply to NFS-backed QCOW2 source VMs.

API account and deployment permissions

Create a dedicated Sendense user in the correct account or domain. Clone the Tenant Admin role and explicitly allow `getPathForVolume`. Have the API URL (including `/client/api`), the API key, the secret key and the zone available through a secure handover method.

The account must see the source VMs and the intended zones, networks, templates and offerings. It must be able to register or import templates and ISOs, deploy and configure VMs, attach and detach enrolment media, start VMs and remove deployment-owned resources if cleanup is needed. At the recovery site it also needs the volume, snapshot and VM operations used by replication, failover and rollback.

Use an account network that the deployment user owns or can access. Have a CloudStack administrator available for global configuration, guest OS types, offerings, permissions and quota changes. Each platform credential is stored against its specific Sendense site.

Source storage and appliance settings

Prepare before setup	Configure and verify during setup
Confirm that the source VMs use NFS-backed QCOW2 storage, and identify the eligible source hosts.	Set the CloudStack credential's Default storage backend to QCOW2 (NFS Storage), and check any per-VM overrides. Discovery does not correct a mismatched selection.
Arrange administrator access to create the guest OS type named exactly `STYPE`, in capitals.	Apply `STYPE` to the CloudStack SNAs, and to the SHA if it is hosted in CloudStack, so SNAgent can identify the appliances.
Make the selected account, zone and network available for appliance deployment.	Import the current SNA and Controller templates and verify `rootDiskController=virtio`. Check existing appliances separately, and confirm the SNA QEMU guest-agent channel works.

KVM host access for SNAgent

Identify every KVM host on which a protected source VM or its SNA can run. Have root or sudo access available on those hosts, with working KVM/libvirt, bash, curl, systemd and a package manager. Permit host access to the SHA on TCP 443, and peer-to-peer TCP 8888 between participating hosts.

Generate the SNAgent installer from the SHA once it is ready, install it with your KVM administrator, and verify each host's site, identity, version, health and recent check-in. If SNAgent is already installed, check its current details.

SNAgent runs on the hypervisor; no Sendense protection agent is required inside the protected guest VMs. NFS/QCOW2 does not require one SNA per KVM host. Prepare host agents at every ACS site acting as a source; a destination that later becomes a source needs the same host preparation.

Further detail: CloudStack deployment checklist · SNAgent · Credential vault

5 CloudStack recovery resources

Custom offerings and template import access

Prepare the offerings below in the recovery zone and make them available to the Sendense API account. The Controller template is imported, and checked to be Ready, during setup. Controllers are created automatically when protected VMs are added to a replication pattern.

Resource	Required configuration
Compute or service offering	Customisable CPU and memory, with root disk size 0. Available to the target account and zone.
Replica disk offering	Custom disk size enabled (`IsCustomized=true`). Set storage tags and provisioning policy for the intended recovery primary storage.
Controller template access	Permission and download connectivity for Sendense to import the current image. Keep its supplied root disk and verify `rootDiskController=virtio`.

The custom disk offering lets Sendense create each replica disk at the required size; a fixed-size disk offering is unsuitable. Source and destination storage types can differ, and shared storage between the

sites is not required.

Controller network and DHCP

Provide a DHCP-enabled guest network with a working gateway, DNS and enough free addresses for one Controller per protected VM, plus headroom. Record the network ID, subnet and DHCP scope. The destination SNA must reach the Controllers on TCP 8080 and the required TCP 10810 onward range; Controllers must reach the SNA on TCP 8081.

Production recovery and isolated test networks

Prepare the production network for a real failover, and a separate isolated network for testing. Both must exist in the recovery zone and be available to the API account. Decide the destination network for every source NIC, and whether each recovered guest keeps its address or receives a new one.

Controller DHCP is separate from the recovered workload's IP policy. The isolated test network must support the intended guest addressing and application dependencies without connecting a duplicate production identity to the live network.

Snapshot and rollback capability

Enable CloudStack KVM volume snapshots with `kvm.snapshot.enabled=true` and provide the required permissions. Have the CloudStack administrator available for any version-specific configuration and restart requirements. Snapshot creation must work while Controllers are running; rollback may stop and restart a Controller.

During setup, confirm the checkpoint method for the installed release and test create, revert and delete through Sendense. For an NFS recovery site, verify the backing storage before setting File-based primary storage in the site details. This is separate from the source credential's QCOW2 (NFS Storage) setting.

Capacity and quotas

Check account and domain quotas for VM count, vCPU, RAM, volumes, primary storage and snapshots. Each protected workload requires a Controller VM, its appliance root disk, one replica disk per source disk and temporary checkpoint space. Reserve workload compute for failover as well as standby resources.

Allow time and bandwidth for the first complete seed. Ongoing bandwidth and storage performance must allow subsequent changes and checkpoint work to finish within the chosen replication interval. Measure this during validation.

Further detail: Sendense Controllers · CloudStack replication checkpoints · Replication patterns

6 VMware access and test workloads

vCenter access for appliance deployment

Provide an approved service account, and have a VMware administrator available. The account must access the deployment datacenter, cluster or resource pool, datastore, VM folder and port group, as well as the protected VM inventory and eligible hosts.

Deployment permissions must cover OVF/OVA import, marking the imported VM as a template, deploying a VM from that template, assigning storage, network and resource-pool placement, configuring devices,

attaching and detaching enrolment media, and powering on. Include removal of deployment-owned failed VMs, templates and temporary media where cleanup is needed.

The VMware SNA is deployed from the SHA using these resources. The SHA needs vCenter and ESXi HTTPS upload access on TCP 443. VMware does not require SNAgent on ESXi or a separate Broadcom VDDK installation.

Ongoing VMware replication access

- The service account must support discovery, snapshot creation and removal, changed block tracking (CBT) and disk reads for the selected VMs. Include source power control if planned failover will be tested.
- The source SNA must resolve and reach vCenter on TCP 443, and every eligible ESXi management address on TCP 902, including DRS or vMotion destinations.
- Confirm the source VMs can take and consolidate snapshots, with working space on their datastores. Resolve existing snapshot or consolidation problems before the first sync.

Select representative test VMs

Start with one ACS NFS/QCOW2 VM and one VMware VM on the storage type you intend to protect. Include Windows, Linux and a VM with multiple disks where these represent the workloads you intend to protect.

Record for each VM	What the information supports
Name, source site and guest OS/version	Correct discovery and guest preparation.
Disk count, provisioned sizes and source storage	Replica capacity, Controller data ports and multi-disk validation.
CPU, memory, BIOS/UEFI and special boot requirements	Recovery resource selection and operating system compatibility.
Every NIC, source network and destination IP policy	Test and production network mapping, gateway and DNS configuration.
Guest access, application checks and dependencies	Validation beyond a successful VM boot.

Guest checks and other backup products

Have guest administration access and application owners available. Confirm the destination guest CPU model exposes the instructions required by the chosen operating systems, including SSE4.2 and POPCNT where required. Check the virtual CPU presented to the guest as well as the physical host.

After a test failover, validate the required drivers, every disk, NIC, IP setting and application. Windows adaptation can involve a reboot. Flag guest encryption, unusual disks and firmware requirements when selecting the test set.

If another backup product also snapshots these VMs, record its backup and snapshot windows. Run the first Sendense tests away from its active snapshot operations, then perform a controlled coexistence check.

Further detail: Failover VM remediation

7 Readiness checklist and setup steps

Have these ready before you start

Ready	Preparation
☐	CloudStack, VMware, network, KVM and guest/application administrators are available to provide access and make the required changes.
☐	Four initial appliance addresses are reserved; networks, prefixes, gateways, DNS and the SHA hostname and certificate details are recorded.
☐	SHA and SNA capacity and placement are agreed; VMware resources and CloudStack zones, accounts and network IDs are available.
☐	Platform accounts and a secure credential handover are ready, including template import and appliance deployment permissions.
☐	SHA routes to vCenter, the ESXi HTTPS upload endpoints and both ACS APIs are prepared; SSVM image and seed download access is available.
☐	Known replication and host-agent firewall rules are prepared; a network administrator can add the allocated Direct P2P ingress range during setup.
☐	Eligible KVM source hosts and root or sudo access are available for SNAgent installation or review.
☐	Custom compute and custom-size disk offerings, Controller DHCP, and the production recovery and isolated test networks are ready.
☐	Quotas, replica and checkpoint space, failover compute and CloudStack snapshot capability have been checked.
☐	Test VMs, disk counts, the required RPO, guest/application checks and any other backup product's windows are agreed.

The setup steps, and where the video shows them

Step	Setup and validation	Video chapter
1	Set up the SHA at the recovery site, activate its licence, give it a web certificate for its address, and create the source and recovery sites and their credentials.	1
2	Create the `STYPE` guest OS type in CloudStack.	2
3	Import the SNA and Controller templates; deploy, address and enrol the SNAs.	3, 4
4	Install or verify SNAgent with the KVM administrator, and check the storage backend, `STYPE`, VirtIO settings and appliance health.	5
5	Test the credentials and discover the VMs; record the destination SNA's ingress range, apply any remaining firewall rule and verify the selected route.	6

Step	Setup and validation	Video chapter
6	Create separate ACS-to-ACS and VMware-to-ACS replication patterns. Select the offerings and networks, add the VMs, and verify Controller DHCP and health. Run the first full sync and incremental syncs, and record the actual route and timings.	7
7	Keep the SHA, SNAs, host agents and Controllers up to date: the SHA first, then the site SNAs and host agents, then the Controllers soon after.	8
8	Run an isolated test failover with application checks; roll it back, resume replication and verify the next sync.	—

During setup, record the installed component versions, Sendense site IDs, imported template IDs, final appliance addresses and assigned replication port ranges. These are setup results; they do not need to exist before you start.

Hand over passwords and API secrets through a secure method; they are stored in the Sendense credential vault during setup. The completed worksheet should contain addresses and resource identifiers, not passwords or secret keys.

Further detail: CloudStack deployment checklist · DR failover · Updates and version management