Backup & Restore Guide

2026


# SipLive SBC Manager

Backup & Restore Guide

2026 — Version 12.14

# Contents

# 1 Overview

The SipLive SBC Manager backup system protects your configuration and call-routing data by packaging everything needed to fully recover a node into a single portable ZIP file. Backups can be created on demand, scheduled to run automatically, and optionally pushed to remote S3-compatible object storage for offsite retention.

## 1.1 What Is Included in a Backup

Every backup ZIP contains:

| Item | Description |
|---|---|
| `README.txt` | Generation timestamp, reason, and originating node name |
| `metadata.json` | Node ID, database name, app URL, and generated config path |
| `database.sql` | Full MariaDB dump of the SBC Manager database (trunks, DIDs, routes, HMR sets, fraud rules, CDR records, audit log, settings, users, roles) |
| `asterisk-generated/` | Staged generated Asterisk include files (pre-apply) |
| `asterisk-generated-live/` | Live Asterisk include files currently in use by Asterisk |
| `env` | Copy of the application `.env` file |

A backup captures the complete configuration state of the node at the moment it is created. It does not include call recordings, active call state, or operating system configuration.

## 1.2 Required Permission

Creating backups, uploading restores, and restoring from remote storage all require the **`backups.manage`** permission. Viewing the backup inventory and downloading backup files requires **`backups.view`**. Administrators have both by default.

# 2 Backup Schedule

The backup schedule controls when the system automatically creates backups, how many copies to keep locally, and when to raise a health alert if backups have not run recently. Schedule settings are per-node and take effect immediately after saving.

## 2.1 Accessing Backup Schedule Settings

Navigate to **Settings → Backup Schedule** from the left navigation menu.

## 2.2 Enabling Scheduled Backups

At the top of the page, check **Enable scheduled backups** to activate automatic backup creation. When unchecked, no automatic backups run — only manual backups created from the Backups page are made.

## 2.3 Frequency Options

The **Frequency** dropdown controls how often backups are created:

| Frequency | Behavior |
|---|---|
| **Hourly** | Creates a backup at the top of every hour. No run-time field is shown. |
| **Daily** | Creates one backup per day at the time specified in **Run at**. |
| **Weekly** | Creates one backup per week on the day and time specified. |

For most deployments, **Daily** is the recommended setting. Use **Hourly** only when your configuration changes very frequently and you need tighter recovery point objectives.

## 2.4 Run Time

The **Run at (HH:MM)** field sets the time of day for daily and weekly backups, entered in 24-hour format. Times are in the **server's local timezone**, shown in parentheses next to the field. Choose an off-peak window — typically early morning — to avoid running during high call-volume periods.

### 2.5 Day of Week

When **Weekly** frequency is selected, the **Day of week** dropdown appears. Select the day on which the weekly backup should run. All other days are skipped.

### 2.6 Local Retention Count

The **Retain (local copies)** field sets the maximum number of backup files to keep on the local disk. After each successful backup, the oldest files beyond this count are deleted automatically. The minimum is 1; the maximum is 365.

**Recommendation:** Keep at least 7–14 local copies to allow recovery from a problem discovered days after it occurred. If remote backup storage (S3) is configured, local retention can be set lower since remote copies provide additional history.

### 2.7 Stale Alert Threshold

The **Stale alert threshold (hours)** field sets how old the most recent backup can be before the system raises an operations alert. The default is 30 hours, which gives a daily backup a 6-hour grace window beyond its expected run time before alerting.

When the latest backup is older than this threshold:

- The **Dashboard** health card for Backups shows a warning state.
- The **Operations Alerts** panel (if delivery is configured) sends an alert notification.

Set this value to approximately the backup frequency plus a reasonable grace window. For weekly backups, a value of 180 hours (7.5 days) is appropriate.

### 2.8 Saving Schedule Settings

Click **Save settings** to apply the new schedule. Changes take effect at the next scheduler tick — there is no need to restart any services.

---

## 3 Remote Backup Storage (S3-Compatible)

Remote backup storage automatically pushes each backup to an S3-compatible object storage bucket immediately after the local copy is saved. This provides offsite protection against local disk failure and allows backup downloads and restores from the remote bucket directly through the Backups page.

Supported providers include **AWS S3**, **MinIO**, **Wasabi**, **Backblaze B2**, **Cloudflare R2**, and any other provider that implements the S3 API.

### 3.1 Accessing Remote Storage Settings

Navigate to **Settings → Remote Backup Storage** from the left navigation menu.

### 3.2 Enabling Remote Sync

Check **Enable remote backup sync** at the top of the form to activate S3 upload. When this checkbox is unchecked, backups are stored locally only, even if endpoint credentials are configured.

### 3.3 Endpoint URL

Enter the S3-compatible API endpoint for your storage provider:

| Provider | Endpoint |
|---|---|
| AWS S3 | `https://s3.amazonaws.com` |
| AWS S3 (regional) | `https://s3.us-east-1.amazonaws.com` |
| Wasabi | `https://s3.wasabisys.com` |
| Backblaze B2 | `https://s3.us-west-004.backblazeb2.com` (region varies) |
| Cloudflare R2 | `https://<account-id>.r2.cloudflarestorage.com` |
| MinIO (self-hosted) | Your MinIO server address, e.g. `https://minio.example.com` |

### 3.4 Region

Enter the storage region for your bucket. For AWS, this matches the bucket's region (e.g. `us-east-1`, `eu-west-1`). For providers that do not use AWS regions (Cloudflare R2, some MinIO deployments), enter `auto` or the value your provider's documentation specifies. The default is `us-east-1`.

### 3.5 Bucket

Enter the name of the S3 bucket where backup files should be stored. The bucket must already exist and the access key must have read/write permission to it.

### 3.6 Object Prefix

The **Object prefix** field sets a folder-like path prefix applied to every backup object uploaded. For example, a prefix of `sbc-backups` stores files as `sbc-backups/sbc-backup-20260621-021500.zip`. A trailing slash is added automatically.

Using a prefix is recommended when the bucket is shared with other applications or when multiple SBC nodes write to the same bucket — use a different prefix per node to keep backups organized.

### 3.7 Access Key ID and Secret Access Key

Enter the access key ID and secret access key for your storage account. The secret access key is stored encrypted in the database and is never displayed after saving — leave the field blank when editing other settings to preserve the existing key.

**Security note:** Create a storage account or IAM role scoped to the minimum permissions needed: `s3:PutObject`, `s3:GetObject`, `s3:ListBucket`, and `s3:DeleteObject` on the target bucket only. Do not use a root or admin-level key.

### 3.8 Saving Remote Storage Settings

Click **Save settings** to store the configuration. The connection is tested on the next backup run. If the upload fails, the local backup is still retained and an error is shown in the operations alert delivery log.



---

## 4 Creating a Manual Backup

A manual backup can be created at any time from the **Backups** page, regardless of the backup schedule.

### 4.1 Steps

1. Navigate to **Backups** from the left navigation menu.
2. In the **Create Backup** section at the top of the page, click **Create backup now**.
3. The page reloads and a success message confirms the backup was created.
4. The new backup appears immediately in the **Local Backup Inventory** table with its filename, creation timestamp, and file size.

If remote backup storage is configured and enabled, the backup is automatically uploaded to S3 after the local copy is saved. There is no separate upload step.

### 4.2 When to Create a Manual Backup

Create a manual backup before any of the following:

- Applying a new configuration to Asterisk (especially large trunk or routing changes)
- Performing a node upgrade
- Importing a large batch of DIDs
- Changing HA node configuration
- Restoring from another backup (back up current state first)

---

## 5 Downloading a Backup

Any backup in the **Local Backup Inventory** can be downloaded to your workstation for archival or transfer to another node.

### 5.1 Steps

1. Navigate to **Backups** from the left navigation menu.
2. Locate the backup you want in the **Local Backup Inventory** table.
3. Click **Download** in the rightmost column of that row.
4. Your browser downloads the backup as a `.zip` file.

Downloaded backup files can be used to restore any SBC Manager node — they are not locked to the originating node.

### 5.2 Remote Backup Downloads

If remote backup storage is configured, the **Remote Backups (S3)** section on the Backups page lists all backup files found in the configured bucket and prefix. Remote backups can be restored directly from this list (see Section 7) without downloading the file to your workstation first.

---

# 6 Restoring from a Local Backup

A restore replaces the current database, Asterisk generated configuration, and application settings with the contents of the chosen backup file. The operation is immediate and irreversible — always create a manual backup of the current state before restoring.

## 6.1 What Restore Replaces

| Component | Restore Behavior |
|---|---|
| **Database** | Fully replaced with the `database.sql` from the backup. All trunks, DIDs, routes, HMR sets, fraud rules, users, CDR records, and settings in the current database are overwritten. |
| **Asterisk generated config** | Both the staged (`asterisk-generated/`) and live (`asterisk-generated-live/`) config directories are replaced. Asterisk is automatically reloaded after the live directory is restored. |
| **Application settings (.env)** | Restored from the backup, with the current node's `APP_KEY`, `APP_URL`, `DB_HOST`, `DB_PORT`, `DB_DATABASE`, `DB_USERNAME`, and `DB_PASSWORD` preserved automatically so the instance keeps working after restore. |

## 6.2 What Restore Does Not Change

- Operating system configuration
- Network settings
- Asterisk binary installation
- SSH keys or system users
- Current node's database connection credentials (always preserved)
- Current node's encryption key (`APP_KEY`, always preserved)

### 6.3 Steps

1. **Create a current-state backup first.** On the Backups page, click **Create backup now** before proceeding.
2. Navigate to **Backups** from the left navigation menu.
3. In the **Restore from External Backup** section, click **Choose file** and select the `.zip` backup file from your workstation.
4. Click **Upload & Restore**.
5. A confirmation dialog appears: *"This will overwrite the current database and Asterisk configuration. Continue?"* Click **OK** to proceed.
6. The system uploads the file, extracts and validates its contents, restores each component, and then reloads Asterisk.
7. On completion, a summary message is displayed showing the restore result for each component (database, Asterisk config, application settings).

### 6.4 After Restore

After a successful restore:

- Log out and log back in to refresh your session against the restored database.
- Navigate to **Operations → Apply Status** to confirm the Asterisk configuration state matches expectations.
- If the restoration is to a different node than the one that created the backup, review **Settings → Network** to confirm the node's SIP IP and hostname are correct for the new environment.

---

## 7 Restoring from a Remote (S3) Backup

When S3-compatible remote backup storage is configured, backups stored in the remote bucket can be restored without downloading the file manually.

### 7.1 Steps

1. **Create a current-state backup first.** On the Backups page, click **Create backup now** before proceeding.
2. Navigate to **Backups** from the left navigation menu.

3. Scroll to the **Remote Backups (S3)** section. If S3 is not configured, a link to the settings page is shown instead.
4. Locate the backup you want to restore in the remote backups table.
5. Click **Restore** in the rightmost column of that row.
6. A confirmation dialog appears: *"Restore from <filename>? This will overwrite the current database and Asterisk configuration."* Click **OK** to proceed.
7. The system downloads the backup from S3 to a temporary location, then runs the same restore process as a local restore.
8. On completion, a summary message is displayed and the temporary file is deleted.

The remote restore process is identical to a local restore in every respect — the same components are replaced and the same credentials are preserved.

---

# 8 Backup Health Monitoring

The SBC Manager continuously monitors backup freshness and surfaces warnings when backups fall outside the expected schedule.

## 8.1 Dashboard Health Card

The **Dashboard** displays a **Backups** health card in the system status panel. The card shows:

- **Healthy (green):** The most recent backup is within the stale alert threshold configured in **Settings → Backup Schedule**.
- **Warning (amber/red):** The most recent backup is older than the stale alert threshold, or no backup has ever been created.

Click the card to navigate directly to the Backups page.

## 8.2 Operations Alerts

If operations alert delivery is configured (via **Settings → Operations Alerts**), the system sends an alert notification when the latest backup exceeds the stale threshold. This allows NOC teams to be notified by email or webhook without having to monitor the dashboard continuously.

The alert clears automatically after the next successful backup is created.

### 8.3 Backup Inventory Timestamps

Every entry in the **Local Backup Inventory** table shows its exact creation timestamp. Use this to quickly verify that scheduled backups are running at the expected times and that the retention policy is pruning old files correctly.

---

## 9 High Availability Considerations

In a two-node HA deployment, each SBC node operates its own independent backup schedule and local backup storage. Backups are not automatically shared between nodes.

### 9.1 Per-Node Backup Independence

- Each node runs its own scheduler and creates backups on its own schedule.
- The **Local Backup Inventory** on each node shows only that node's local backups.
- A backup created on `SBC-Node-1` must be manually downloaded and uploaded to `SBC-Node-2` if you want to restore it there, or pulled from S3 if both nodes share a remote bucket.

### 9.2 Recommended HA Backup Strategy

For HA deployments, configure the same S3 bucket on both nodes with different object prefixes (e.g. `sbc-node1/` and `sbc-node2/`). This ensures that:

- Both nodes independently push backups offsite.
- Either node can restore a backup from the other node by navigating to **Settings → Remote Backup Storage**, temporarily pointing to the other node's prefix, and using the remote restore workflow.

### 9.3 Restoring One Node While the Other Stays Active

A restore on one node does not affect the other node's Asterisk process or active calls. However, after restoring the database on one node, the two nodes will have diverging configuration states until a new Apply is run. After restore:

1. Verify the restored node's configuration is correct.
2. Navigate to **Operations → Apply Status** and apply the configuration to the restored node to bring its live Asterisk files in sync with the database.

3. Check **Operations → HA Nodes** to confirm both nodes show the same configuration hash.

---

# 10 Best Practices

## 10.1 Before Any Significant Change

Always create a manual backup immediately before applying large configuration changes, node upgrades, or DID bulk imports. This ensures you have a known-good restore point with a timestamp that maps directly to the moment before the change.

## 10.2 Test Your Restore Process

Periodically restore a backup to a test node or VM to confirm the process works end-to-end. A backup you have never tested restoring from is an untested hypothesis, not a recovery plan.

## 10.3 Use Remote Storage for Offsite Protection

Local backups are lost if the node's disk fails. Configure S3-compatible remote storage to ensure backup copies survive node hardware failures. Use a geographically separate storage region when possible.

## 10.4 Monitor the Health Card

Check the Dashboard backup health card after any planned maintenance window to confirm scheduled backups resumed. After a long maintenance period, the scheduler may have missed one or more scheduled runs — create a manual backup to reset the freshness counter.

## 10.5 Backup Before Restore

Always create a backup of the current state before performing any restore operation. If the restore fails or produces unexpected results, this gives you a recovery path back to the state the node was in before the restore began.

## 10.6 Retention and Storage Budget

At the default schedule (daily backups, 14 retained), each node stores approximately 14 backup ZIPs locally. Backup file sizes depend on CDR volume and the number of routes and trunks

but are typically between 500 KB and a few MB per node. Monitor local disk usage if you increase retention significantly or switch to hourly backups.

---

## 11 Permissions Reference

| Permission | What It Allows |
|---|---|
| `backups.view` | View the backup inventory list and download backup files |
| `backups.manage` | Create manual backups, restore from local or remote backup, configure backup schedule, configure remote S3 storage |

Users without `backups.manage` see the backup inventory and download links but do not see the **Create Backup** or **Restore** controls. Users without `backups.view` cannot access the Backups page at all.

Permissions are assigned per role under **Security → Roles**. The default **Administrator** role includes both permissions.