



# SAFR<sup>®</sup> Large Scale Deployment

Documentation Version = 3.048

Publish Date = August 19, 2022

Copyright © 2022 RealNetworks, Inc. All rights reserved.

SAFR® is a trademark of RealNetworks, Inc. Patents pending.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

**Note:** Documentation pertaining to the macOS platform is no longer being actively maintained.

## Contents

<b>1</b>	<b>Upgrade SAFR Server</b>	<b>3</b>
<b>2</b>	<b>SAFR Server Clusters</b>	<b>4</b>
<b>3</b>	<b>Add Secondary Servers</b>	<b>8</b>
<b>4</b>	<b>Database and Object Storage Redundancy</b>	<b>12</b>

# 1 Upgrade SAFR Server

RealNetworks releases new versions of SAFR almost every month that contain additional features. You will occasionally want to upgrade your SAFR Server(s) to the latest version in order to take advantage of the new features. To upgrade your SAFR Server(s), follow the instructions below.

## 1.1 Upgrading a Single SAFR Server

To upgrade a single SAFR Server, simply download the latest SAFR Platform installer and install it on top of your existing SAFR Server.

**Note:** Your SAFR service will be offline while the new SAFR Server is installing.

## 1.2 Upgrading a SAFR Server Cluster

To upgrade your SAFR Server cluster to the latest version of SAFR, do the following:

1. Prevent new faces from being added during the upgrade by performing one of the following:
  - Stop the CoVi service from running on all secondary SAFR Servers, OR
  - Stop all clients (i.e. VIRGO video feeds, Desktop clients, and Mobile clients) that are adding new faces.
2. Make a backup of your Primary SAFR Server, as described in the SAFR Server Backup documentation.  
**Note:** Backup and restore should only be performed for the primary server. Secondary servers will synchronize automatically with the primary server.
3. Update the primary SAFR Server by running the latest SAFR Platform installer. While the primary server is upgrading:
  - The system will be down unless an external load balancer is being used.
  - Even when an external load balancer is being used, there will be about 2-20 seconds where newly generated events might be dropped.
  - Live video streams will be down.
4. Update the secondary SAFR Servers by running the latest SAFR Platform installer on each of them.
5. Restart the CoVi service on all the secondary servers, if you stopped them in Step #1. Similarly, you can resume submitting faces on all clients.

## 2 SAFR Server Clusters

At some point, your SAFR system's capacity and/or performance may degrade if the number of face recognition requests sent to your SAFR Server overwhelms your server's capacity. (Performance problems may also arise if the number of people in your Person Directory becomes too large.) Fortunately, you can install additional SAFR Servers on other machines in order to increase your SAFR system's capacity, improve performance, and improve resiliency. The first SAFR Server you install is automatically your primary server, while all additional servers are secondary servers.

**Note:** You can change which machine is the primary server by uninstalling the primary server, waiting 24 hours, and then re-installing the SAFR Server on a different machine. The 24 hour wait time can be avoided if you contact your SAFR Account Manager and ask them to manually reset your IP address.

### 2.1 Understand When to Scale

A single SAFR Server that's also running a Desktop Client can handle up to 16 cameras, (assuming each camera view contains just a single face), as long as the host machine meets the recommended hardware requirements. If the machine running the server doesn't have any cameras connected directly to it, then the server's capacity increases to 25 cameras, again assuming that each camera view contains a single face. A higher number of faces per camera or a higher number of cameras requires either vertical scaling of a single server (i.e. more or faster CPUs) or horizontal scaling by installing more SAFR Servers.

Another possible performance bottleneck is the network throughput of the primary server. You may want to monitor its network throughput during maximum concurrency times to make sure the network is not over-saturated.

### 2.2 Load Balancing Configurations

For prescribed deployments, the system requirements of the Desktop Client need to be combined with those of SAFR Server. A single Desktop Client typically handles up to 16 cameras as long as it is equipped with a GPU card (see SAFR System Requirements). In this way, running SAFR Server and the Desktop Client on the same machine using the recommended configuration can host up to 16 cameras, assuming each camera view contains with a single face.

There are three different load balancing configurations you can choose from.

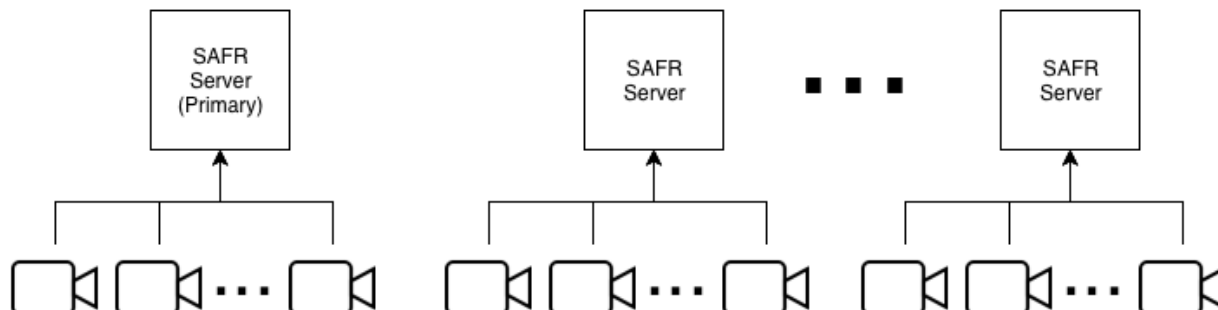
- **Prescribed Load Balancing Configuration:** Cameras are connected to Desktop Clients or Video Recognition Gateway (VIRGO) video feeds running on the same machines that are hosting your SAFR Servers. This gives you tight control over how your face recognition load is distributed, since the video feeds' face recognition requests are processed on the same machine where the video feeds are connected. The system requirements of the Desktop Client need to be combined with those of SAFR Server when calculating the system requirements for a machine hosting the SAFR Server and Desktop Client. A single Desktop Client typically handles up to 16 cameras as long as it is equipped with a GPU card (see SAFR System Requirements). Thus, a machine running SAFR Server and the Desktop Client which meets the recommended system requirements can host up to 16 cameras, assuming each camera view contains just one single face.
- **Software-Based Load Balancing Configuration:** In this configuration the machines hosting SAFR Servers do not also have cameras connected to them. All face recognition requests are initially sent to the primary server, and the primary server acts as the load balancer for the server cluster.
- **External Load Balancing Configuration:** In this configuration all face recognition requests are directed at one or more external load balancer(s), which handle load balancing duties for the SAFR system.

#### 2.2.1 Prescribed Load Balancing Configuration

In the prescribed configuration, you run multiple SAFR Servers by connecting cameras to Desktop Clients or VIRGO video feeds running on the same machines that are hosting SAFR Servers. In this way, you have tight

control over which servers take the video feed load. This is also a useful configuration for systems with very low video feed count totals where running a Desktop Client on a separate machine from the SAFR Server would take more resources than are required for the given use case.

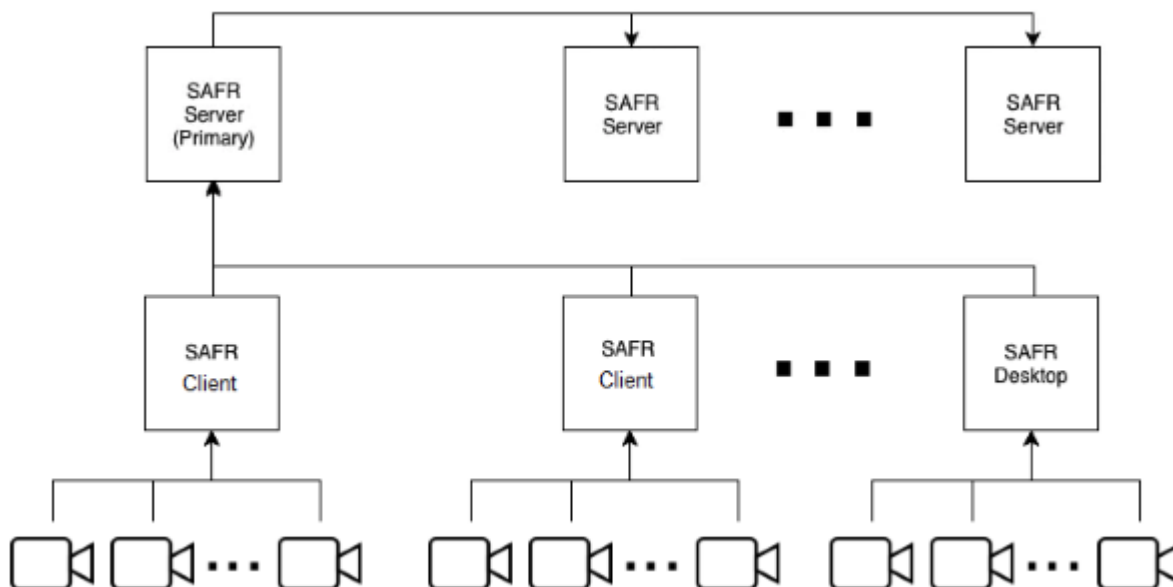
The following diagram illustrates this configuration:



Most services (e.g. face service, events, and reports) are performed on the server where recognition requests are sent.

### 2.2.2 Software-Based Load Balancing Configuration

In the software-based load balancing configuration, cameras aren't connected to machines running SAFR Servers. When newly installed secondary servers are configured, they check in with the primary server and announce that they're ready to receive load-balanced traffic. All recognition requests go through the primary server, which balances the load among itself and all other servers in the SAFR system. The following illustration demonstrates this setup:

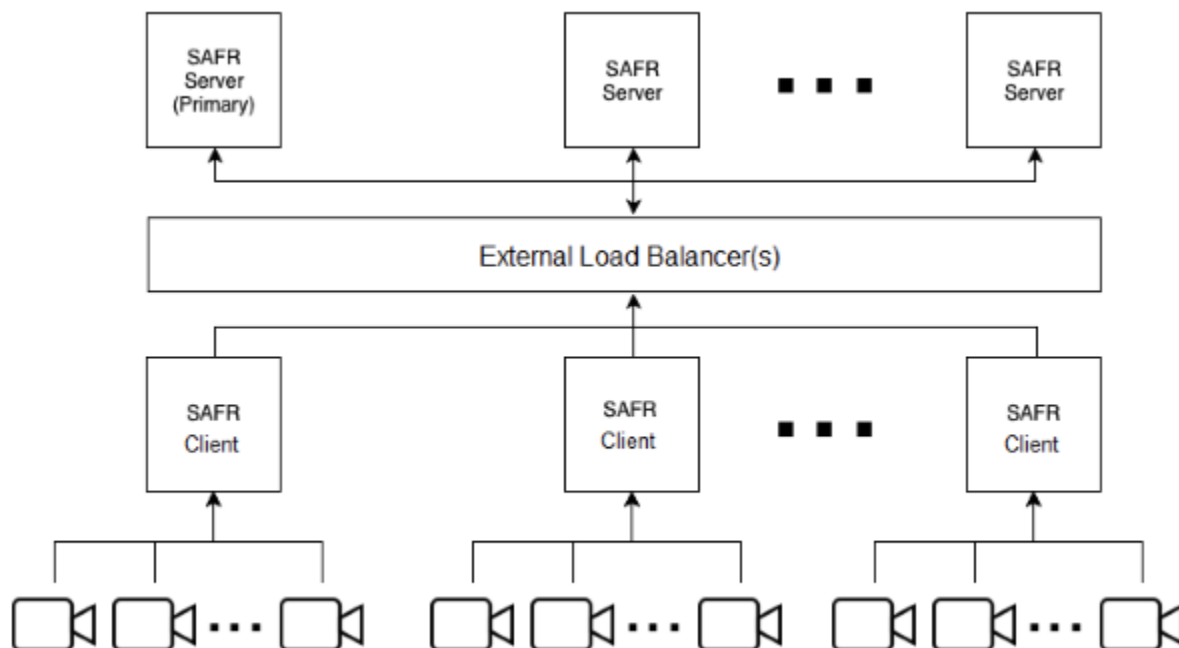


### 2.2.3 External Load Balancing Configuration

The software-based load balancing configuration has the limitation that the primary server is a single point of failure. All traffic is routed through the primary server before any traffic is redirected to the rest of the servers. If the primary server is down, all traffic will stop. External load balancing is an alternate configuration that can be used to provide a more robust setup that can better deal with server failure.

When using an external load balancing configuration, all network traffic is first routed to one or more load balancer(s), and the load balancer(s) proxy requests to the backend servers over either HTTP or HTTPS. HTTP would be OK in situations where network traffic is isolated to a trusted network, or when network sniffing by non-target hosts is impossible.

If HTTPS is used to proxy traffic to SAFR servers, you should manually disable load balancing on all secondary servers as described below so that the primary server isn't double load balancing traffic to them. A valid (i.e. non self-signed) SSL certificate would still need to be installed and configured on the primary server. Secondary servers should be fine with the default (i.e. self-signed) certificate, if your load balancer allows it.



## 2.2.4 Manually Configure Load Balancing Traffic

SAFR Servers can be manually enabled or disabled to accept load balancing traffic.

**Note:** If the server you want to disable is the only one configured to take traffic, you receive a warning and prompt to continue. In this case, should you proceed, your system will most likely go offline.

### Disable Load Balancing Traffic

To stop receiving traffic on a server, log in to a shell on the server and run the appropriate command for your server's OS:

OS	Command
Windows	"C:\Program Files\RealNetworks\SAFR\bin\server-status.py" --disable
macOS	/Library/RealNetworks/SAFR/bin/server-status.py --disable
Linux	sudo /opt/RealNetworks/SAFR/bin/server-status.py --disable

It may take up to one minute for the desired traffic state to change.

### Enable Load Balancing Traffic

To resume receiving traffic on a server, log in to a shell on the server and run the appropriate command for your server's OS:

OS	Command
Windows	"C:\Program Files\RealNetworks\SAFR\bin\server-status.py" --enable
macOS	/Library/RealNetworks/SAFR/bin/server-status.py --enable
Linux	sudo /opt/RealNetworks/SAFR/bin/server-status.py --enable

It may take up to one minute for the desired traffic state to change.

## 2.3 Upgrading Server Clusters

See the Upgrading a SAFR Server Cluster documentation for details on upgrading your SAFR Servers to the latest version.

## 3 Add Secondary Servers

The first SAFR Server you install will automatically become the primary server. All subsequent servers you install will be secondary servers. There are two types of secondary servers:

- **Simple:** Does not replicate the database data.
  - **Redundant:** Replicates database data from the primary server, possibly providing failover functionality. (See the Database and Object Storage Redundancy topic for details.)
- Note:** Only Windows and Linux SAFR Servers can become redundant secondary servers.

### 3.1 Add a Secondary Server While Connected to the Internet

If your system is connected to the Internet, do the following to add a secondary server:

1. Download and install SAFR Platform on the additional machine.
2. Log in to the SAFR auto-discovery process:
  - Connect the Desktop Client to the primary server (for macOS and Windows) as described here.
  - Connect your Web Console to the primary server as described here.
3. During auto-discovery, the following automatically happens:
  1. The secondary server contacts a SAFR Licensing Server in the cloud to acquire a license.
  2. The SAFR Licensing Server authenticates the SAFR account credentials.
  3. The SAFR Licensing Server identifies the license and deployment type.
  4. A suitable license is returned to the secondary server and information about the primary server is returned to the secondary server, including the hostname.
4. If your new secondary server is on a Windows or Linux machine, you will be prompted to choose which kind of secondary server you want: simple or redundant. If your new secondary server is on a macOS machine no prompt will occur; macOS secondary servers are always simple.
5. Auto-discovery will now continue, with the following automatically occurring:
  1. The secondary server re-configures itself to reference the primary server.
  2. The secondary server registers itself with the primary server.
  3. The primary server updates its local database.
  4. If you're using a Software-Based Load Balancing Configuration, the primary server now adds the new secondary server to its load balancer configuration and uses it as an additional node in its cluster.

### 3.2 Add a Secondary Server While Offline

If you are not connected to the Internet, you can still connect your new secondary server to the primary server, but the auto-discovery process is not available. You must instead manually configure the newly installed secondary server to locate the primary server. If your new secondary server is on a Windows or Linux machine, you'll need to choose which kind of secondary server you want: simple or redundant. If your new secondary server is on a macOS machine no such decision is required; macOS secondary servers are always simple.

1. Download and install SAFR Platform on the second machine.
2. Run the *safr-worker* script on your secondary server by doing the following:

**For macOS:**

1. Open Terminal.
2. Run the following command, substituting the primary SAFR hostname for HOSTNAME:

```
sudo /Library/RealNetworks/SAFR/bin/safr-worker HOSTNAME
```

**For Windows**

1. On the primary server record the contents of `C:\ProgramData\RealNetworks\SAFR\mongo\.adminpass` and `C:\ProgramData\RealNetworks\SAFR\mongo\mongod.keyfile`
2. On the new secondary server, open a command prompt by right-clicking on the **Start** menu, selecting **Run**, and entering `cmd`.
3. If you want it to be a simple secondary server, in the new command prompt run the following command, substituting the password from `.adminpass` in Step 1 for `PASSWORD` and the primary server hostname for `HOSTNAME`:

```
python "C:\Program Files\RealNetworks\SAFR\bin\saf-worker.py" -p
    PASSWORD HOSTNAME
```

OR

4. If you want it to be a redundant secondary server, in the new command prompt run the following command, substituting the `mongod.keyfile` contents from Step 1 for `KEYFILE`, the password from `.adminpass` in Step 1 for `PASSWORD`, and the primary server hostname for `HOSTNAME`:

```
python "C:\Program Files\RealNetworks\SAFR\bin\saf-worker.py" -s KEYFILE
    -p PASSWORD HOSTNAME
```

#### For Linux

1. On the primary server, record the contents of `/opt/RealNetworks/SAFR/mongo/.adminpass` and `/opt/RealNetworks/SAFR/mongo/mongod.keyfile`
2. On the primary server, open `/opt/RealNetworks/SAFR/virgo/config/virgo-factory.conf`
3. Within that file, look for a section that looks something like:

```
"global":{
    "environment": "CUSTOM",
    "user-id": "ubuntu18int2tst",
    "user-encrypted-password":
        "%qy4Effq2cxYUrEopuIFS3LE22hDBMOG6NdeqTWfok4=",
    "status-interval":5000,
    "remote-control-enabled":true
},
```

4. Record the "user-id" and "user-encrypted-password" values. These will be your `SAFRUSER` and `SAFRPASSWORD` values in the steps below.
5. If you want it to be a simple secondary server, on the new secondary server run the following command.

```
sudo python /opt/RealNetworks/SAFR/bin/saf-worker.py -p PASSWORD -u
    SAFRUSER -x SAFRPASSWORD HOSTNAME
```

where:

- **PASSWORD** = the password from `.adminpass` in Step 1
- **SAFRUSER** = "user-id" from Step 4. Don't include the enclosing double quotation marks.
- **SAFRPASSWORD** = "user-encrypted-password" from Step 4. Don't include the enclosing double quotation marks.
- **HOSTNAME** = the primary server hostname

OR

6. If you want it to be a redundant secondary server, on the new secondary server run the following command.

```
sudo python /opt/RealNetworks/SAFR/bin/safr-worker.py -s KEYFILE -p
PASSWORD -u SAFRUSER -x SAFRPASSWORD HOSTNAME
```

where:

- **KEYFILE** = the contents of `mongod.keyfile` from Step 1
- **PASSWORD** = the password from `.adminpass` in Step 1
- **SAFRUSER** = "user-id" from Step 4. Don't include the enclosing double quotation marks.
- **SAFRPASSWORD** = "user-encrypted-password" from Step 4. Don't include the enclosing double quotation marks.
- **HOSTNAME** = the primary server hostname

### 3.3 Error Messages

When attempting to join a new secondary server, you might encounter the following error messages:

Error Message	Description
System is offline	Network or system connectivity issue. Attempt to access the system at a later time.
SAFR master host is not reachable	Ensure all servers are connected to the same network and try again.
Improperly configured SSL certificate	Installing non self-signed SSL certificates is recommended when setting up multiple servers. See the SSL Certificates page for information about how to install an SSL certificate.
Secure connection error. Check server for valid SSL certificate	Installing non self-signed SSL certificates is recommended when setting up multiple servers. See the SSL Certificates page for information about how to install an SSL certificate.
Incomplete server connection	Attempt to join again; a persistent issue may require either uninstalling and reinstalling SAFR Platform on your servers or contacting your SAFR support representative.

### 3.4 Secondary Server Health Checks

- At startup each server, both primary and secondary, registers itself by posting its status to the database on the primary server.
- The primary server directs requests to all secondary servers in a *least connection method* that keeps the load evenly balanced among all secondary servers.
- As long as a given secondary server remains healthy, the primary server keeps that secondary server in its load balance rotation.
- Status information about all secondary servers is stored in the database on the primary server.
- Every minute all servers (the primary server as well as the secondary servers) send a status update to the database on the primary server.
- Every five seconds, the primary server attempts to ping all servers (the primary server as well as all secondary servers) via the SAFR health check API.
- If the health check fails for a given secondary server for 15 seconds (i.e. for 3 health check API calls in a row), that secondary server is removed from load balancing rotation and face recognition requests are no longer routed to it. If the health check succeeds for the removed secondary server for ten seconds (i.e. for 2 health check API calls in a row), the secondary server is returned to the load balancing rotation and resumes accepting face recognition requests.
- If a secondary server's status has not been reported for over five minutes, it is removed from the load balancer configuration. In this case, it is no longer sent face recognition requests or health check API

calls.

- If a secondary server has been pulled out of rotation for not responding to health checks, or is removed from the load balancer configuration for not reporting status for more than five minutes, it can still be put back in rotation through any of the following:
  - If a network interruption prevents the secondary server from sending a request, the secondary server continues to send a status update at its regularly scheduled interval after it goes back online and its status is updated in the primary server.
  - If the secondary server is restarted, it sends a status update after all services are started and ready.
  - If the secondary server IP address is changed, the secondary server must be manually restarted to force it to send a status update to the primary server with the new IP address.

## 4 Database and Object Storage Redundancy

### 4.1 Database Redundancy

The first SAFR Server you install will automatically become the primary server. All subsequent servers you install will be secondary servers. There are two types of secondary servers:

- **Simple:** Does not replicate database data.
  - **Redundant:** Replicates database data from the primary server, possibly providing failover functionality. (See the Failover Functionality section below for details.)
- Note:** Only Windows and Linux SAFR Servers can become redundant secondary servers.

With both types of secondary servers services such as feed management, reports, and the Web Console are not load-balanced and are always served from the primary server.

#### 4.1.1 Failover Functionality

If there are at least two redundant secondary servers (three servers total), failover functionality is enabled. This means that if the primary server goes offline and both of the first two installed redundant secondary servers are still online, one of the redundant secondary servers will become the new primary server and the server cluster will continue to function as normal.

If additional redundant secondary servers are installed beyond the first two, database data will be replicated on them, but they don't count for the purpose of failover functionality.

### 4.2 Object Storage Redundancy

**Note:** Object Storage Redundancy is only available on Windows and Linux.

The Object Storage Service is used for storing objects, such as profile and event images, as well as ephemeral data, such as event reply messages.

The service can operate in a redundant configuration when you have multiple SAFR servers running. All redundant secondary servers are load-balanced by the primary server for all Object Storage Service requests it receives.

#### 4.2.1 Shared Object Storage (Network Storage)

Using shared object storage provides a shared location for each server to save and retrieve objects from. This provides each Object Storage Server with access to all of the objects, rather than just objects saved to their local storage.

Shared storage also provides an easier backup process, as you only have to run it from the primary server.

#### 4.2.2 Local Object Storage (Not Recommended)

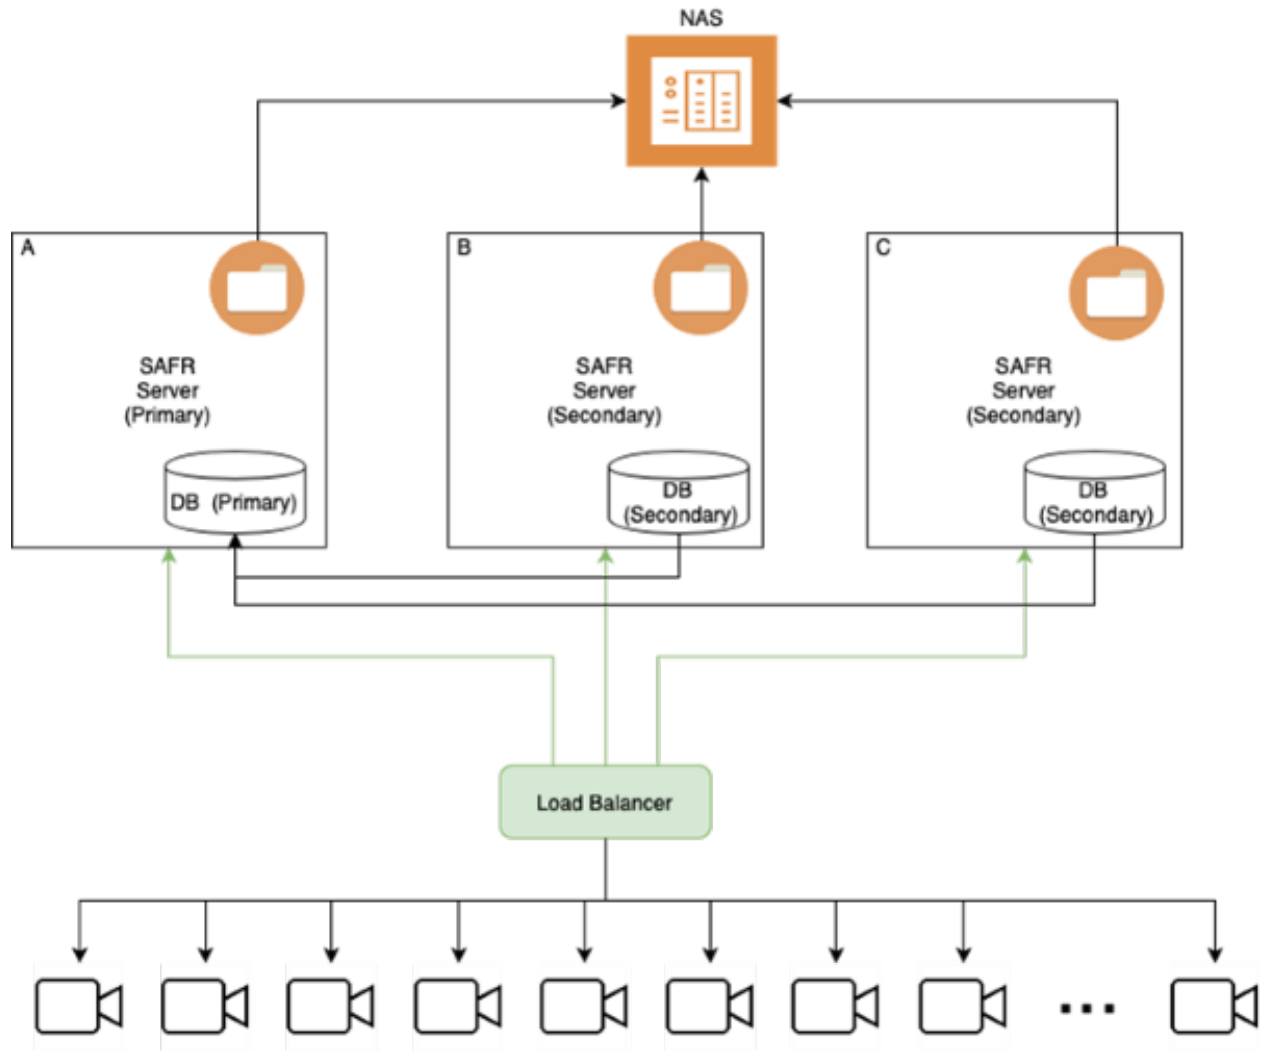
By default all redundant servers will save objects locally, and ask other Object Storage Servers for objects it does not have locally.

When you're using local object storage, you will lose access to all objects that are only stored by an offline Object Storage Server until the server becomes healthy again. If that server's objects are lost, and you do not have backups, they will be unrecoverable.

Backups must be run on every redundant server that has Object Storage enabled.

### 4.3 External Load Balancing (ELB) Walkthrough

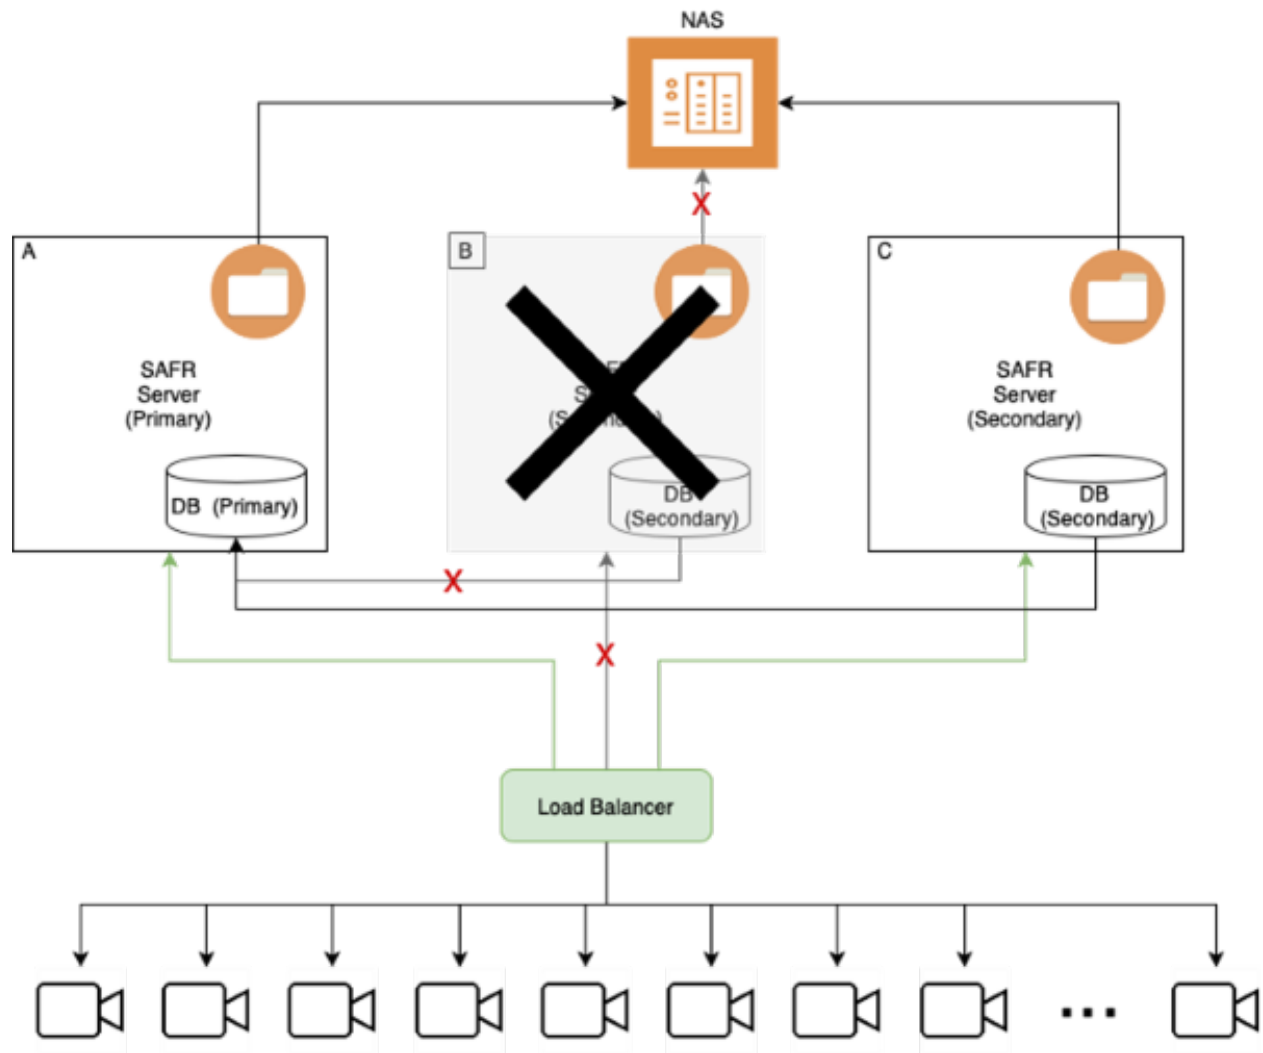
This section describes the functionality of a common large scale deployment configuration: an external load balanced system with 2 redundant secondary servers and a network-attached storage (NAS), a type of shared object storage.



When all the servers are working, this configuration has the following properties:

- Face recognition requests are distributed across all three servers.
- Server A's database is the primary database (i.e. performs database writes).
- All three server share in the database read load.
- All three servers read and write directly to the NAS.
- There isn't a single point of failure.
- Data is redundantly protected across all three servers.

#### 4.3.1 Secondary Server Failure in an ELB Configuration



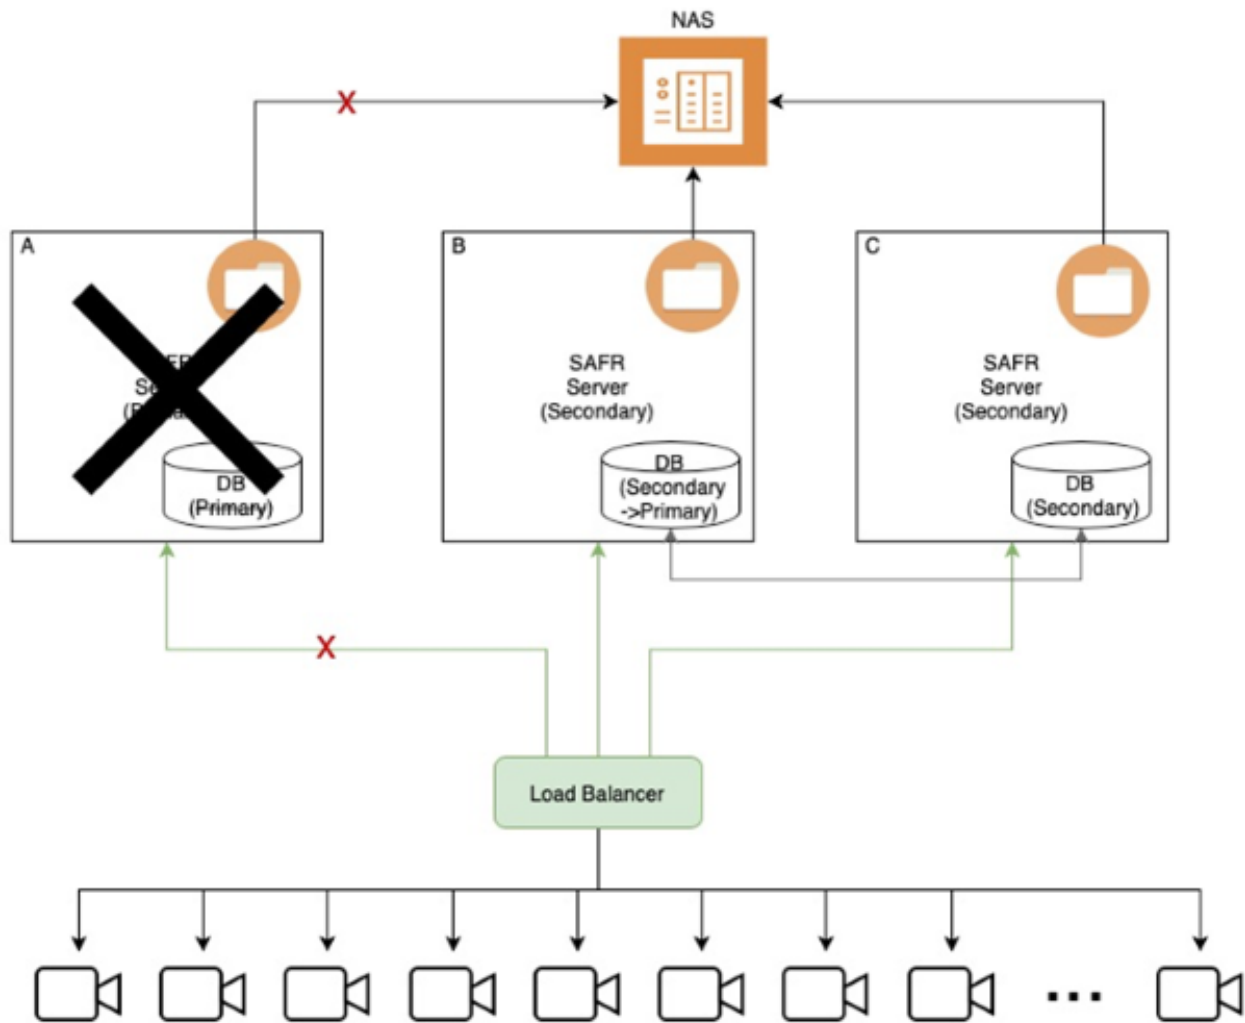
When one of the redundant secondary servers fail, the following occurs:

- Face recognition requests are handled by servers A (primary) and C (secondary).
- Server B is removed from rotation by the load balancer.
- Server B is removed from the database replica set.
- Server A's database continues operating as the primary database (i.e. it performs database writes).
- Servers A and C share in the database read load.
- Both remaining servers read and write directly to the NAS.

##### Impact:

- No service outage occurs, but longer latency may result.
- Data remains redundantly protected.
- There isn't any impact to object storage.

### 4.3.2 Primary Server Failure in an ELB Configuration



When the primary server fails, the following occurs:

- Face recognition requests are handled by servers B (secondary) and C (secondary).
- Server A is removed from rotation by the load balancer.
- Server A is removed from the database replica set.
- The server B database takes over as the primary database (i.e. it performs database writes).
- Servers B and C share in the database read load.
- Both remaining servers read/write directly to the NAS.

**Impact:**

- No service outage occurs, but longer latency may result.
- Data remains redundantly protected.
- There isn't any impact to object storage.

## 4.4 Migrate from Local to Shared Storage

If you start with local storage but later decide to move to shared storage, you will need to consolidate all of your objects to the new shared storage location, delete the local copies, and then mount the shared storage to the correct location. To do this, do the following:

1. Back up both the primary and redundant secondary servers to ensure you have a full backup of all SAFR content.
  - **On Linux:**
    - **Primary:** `python /opt/RealNetworks/SAFR/bin/backup.py`
    - **Redundant Secondaries:** `python /opt/RealNetworks/SAFR/bin/backup.py -o`
  - **On Windows:**
    - **Primary:** `python "C:\Program Files\RealNetworks\SAFR\bin\backup.py"`
    - **Redundant Secondaries:** `python "C:\Program Files\RealNetworks\SAFR\bin\backup.py" -o`
2. Stop all primary and redundant secondary servers by using the **stop** command. This can be done by doing the following on each server:
  - **On Linux:** `/opt/RealNetworks/SAFR/bin/stop`
  - **On Windows:** `"C:\Program Files\RealNetworks\SAFR\bin\stop.bat"`
3. Mount the new shared storage to a temporary location on primary and redundant secondary servers.
4. Copy all files from the primary server and every redundant secondary server(s) to the temporary location of the shared storage. from within the following paths:
  - **On Linux:** `/opt/RealNetworks/SAFR/cv-storage`
  - **On Windows:** `C:\ProgramData\RealNetworks\SAFR\cv-storage`
5. Delete or move the contents of the CV Storage folder on each primary and redundant secondary server as specified below.
  - **On Linux:** `/opt/RealNetworks/SAFR/cv-storage`
  - **On Windows:** `C:\ProgramData\RealNetworks\SAFR\cv-storage`
6. Unmount the temporary location of the new shared storage.
7. Mount the shared storage to the correct CV Storage location, or create a symlink to the shared storage location.
8. Start the primary and redundant secondary servers by using the **start** command. On each server, do the following:
  - **On Linux** `/opt/RealNetworks/SAFR/bin/start`
  - **On Windows** `"C:\Program Files\RealNetworks\SAFR\bin\start.bat"`
9. Disable any automatic backups on redundant secondary servers.
  - Now that you're using shared storage, only the primary server needs to be backed up. Disable any automatic backups you may have configured on your secondary servers.

## 4.5 Simple Secondary Server Behavior with Local Object Storage

On simple secondary servers, the Object Storage Service will operate in proxy mode.

Object Storage Servers operating in proxy mode will not attempt to use their own storage for objects, but will instead proxy the request to Object Storage Services that are running on either the primary server or on a redundant secondary server. If the redundant server it contacts doesn't have the object, the contacted redundant server will ask all other redundant servers for the object.

The list of servers that run the Object Storage Service is stored in the database and updated every minute. If a host does not respond within a timeout, it is de-prioritized.

## 4.6 Redundant Secondary Server Behavior with Local Object Storage

On both the primary server and on redundant secondary servers the Object Storage Service stores new objects in storage.

When a server receives a request for a file it does not find in its storage, it will request the object from other Object Storage Servers via HTTPS, and return the object if found. (The same applies for DELETes.) This allows multiple Object Storage Servers to operate without using shared network storage, with each server saving a subset of the total objects, and relaying requests for other objects to its neighbors.

Even when using shared network storage, sometimes a request will come in for a new object before it is

visible to all systems on the shared storage. The Object Storage Service will ask all the other Object Storage Servers for the object until it finds one that has the object.

## 4.7 Backup and Restore

The SAFR backup and restore process when using shared network storage is straightforward - you just need to back up the primary server. This will back up all configs, database content, and Object Service Storage objects.

When using local storage, however, the objects are distributed to multiple servers, so the backup must be run on the primary server as well as all redundant secondary servers.

The primary server should run a regular backup, while the redundant secondary servers run an '*objects only*' backup. The difference is just the addition of the **"-o"** flag to the backup script.

When restoring multiple backups, you can restore them all to the primary server, or you can restore the '*object only*' backups back to the same servers that they were backed up from.

### 4.7.1 Backup for Local Storage

- **On Linux**
  - **Primary:** `python /opt/RealNetworks/SAFR/bin/backup.py`
  - **Redundant Secondaries:** `python /opt/RealNetworks/SAFR/bin/backup.py -o`
- **On Windows**
  - **Primary:** `python "C:\Program Files\RealNetworks\SAFR\bin\backup.py"`
  - **Redundant Secondaries:** `python "C:\Program Files\RealNetworks\SAFR\bin\backup.py" -o`

### 4.7.2 Restore for Local Storage

- **On Linux**
  - **Primary:** `python /opt/RealNetworks/SAFR/bin/restore.py BACKUPFILENAME`
  - **Redundant Secondaries:** `python /opt/RealNetworks/SAFR/bin/restore.py -o BACKUPFILENAME`
- **On Windows**
  - **Primary:** `python "C:\Program Files\RealNetworks\SAFR\bin\restore.py" BACKUPFILENAME`
  - **Redundant Secondaries:** `python "C:\Program Files\RealNetworks\SAFR\bin\restore.py" -o BACKUPFILENAME`

## 4.8 Example Shared Storage Configurations

Below are two example shared storage configurations.

### 4.8.1 Linux

Shared storage on Linux is very straightforward. Simply mount your shared storage to the `/opt/RealNetworks/SAFR/cv-storage` location.

1. Stop SAFR.

```
/opt/RealNetworks/SAFR/bin/stop
```

2. Create a shared storage location. The example below uses Amazon's Elastic File System (EFS).

## Create file system

Step 1: Configure file system access

Step 2: Configure optional settings

Step 3: Review and create

### Review and create

Review the configuration below before proceeding to create your file system.

#### File system access

VPC	Availability Zone	Subnet	IP address	Security groups
vpc-71169419 - vpc-mcv-gsprod	us-west-2a	subnet-	- public 2a	Automatic sg- - default
	us-west-2b	subnet-	- private 2b	Automatic sg- - default
	us-west-2c	subnet-	- public 2c	Automatic sg- - default
	us-west-2d	Not configured		

#### Optional settings

Tags No tags added  
 Performance mode General Purpose  
 Throughput mode Bursting  
 Encrypted No  
 Lifecycle policy None

Cancel Previous Create File System

- Edit `/etc/fstab` to create a mount point of `/opt/RealNetworks/SAFR/cv-storage` for your shared storage. The specific mount options should be provided by your specific storage service or device.

```
fs-12345678.efs.us-west-2.amazonaws.com:/
/opt/RealNetworks/SAFR/cv-storage nfs4
nfsvers=4.1,rsiz=1048576,wsiz=1048576,hard,timeo=600,retrans=2,_netdev
0 0
```

- Mount the remote share.

```
sudo mount -a
```

- Start SAFR.

```
/opt/RealNetworks/SAFR/bin/start
```

## 4.8.2 Windows

Windows cannot mount a shared storage location directly to `C:\ProgramData\RealNetworks\SAFR\cv-storage`. It must instead create a symbolic link by doing the following:

- Stop SAFR.

```
"C:\Program Files\RealNetworks\SAFR\bin\stop.bat"
```

- Create a shared storage location.

- Delete the existing `C:\ProgramData\RealNetworks\SAFR\cv-storage` by running `rmdir /q /s C:\ProgramData\RealNetworks\SAFR\cv-storage` in an administrative command prompt. Deleting the existing `cv-storage` allows you to create a symbolic link from the `cv-storage` location to your shared storage location.

**Note:** Be sure you either followed the migration steps above to consolidate your objects onto the new shared storage location, or that you're doing this on a new system without any data.

- Create the symbolic link from `C:\ProgramData\RealNetworks\SAFR\cv-storage` to your shared storage location. To do this, run the appropriate command in an administrative command prompt:

- If you're using a mapped network drive, run

```
mklink /d C:\ProgramData\RealNetworks\SAFR\cv-storage Z:\
```

- If you're using an SMB share, run

```
mklink /d C:\ProgramData\RealNetworks\SAFR\cv-storage \\servername\share
```

5. Start SAFR.

```
"C:\Program Files\RealNetworks\SAFR\bin\start.bat"
```