

# Module 1: Solution definition

This toolkit is designed for Adobe Campaign Classic Architect Master Exam Aspirants. There are five modules. Study each module per week to stick to schedule. Technical parts of applications are depicted in videos, about which you can learn more from Experience League. You can visit [Get prep page](#) to understand the contents and anticipate the learning journey.

This is Architect Master Exam, toolkit Module 1. This module contains five sections.

## 1.1 Manage subscriptions

### About information services

An information service comprises:

- Registration and subscription (opt-in),
- Deregistration, voluntary unsubscription (opt out) or automatic unsubscription (limited-time service, for example as a trial offer),
- Subscription and unsubscription confirmation mechanisms (simple mechanisms with confirmation, double opt in, etc.),
- Tracking of subscriber history.

As a standard feature, these services include specific statistical reports: subscriber tracking, loyalty level, unsubscription trends, etc.

For emails, the mandatory unsubscription links are generated automatically, and the entire opt in/opt out process is fully automated, with history tracking to guarantee full compliance with the regulations in force.

There are three service subscription/unsubscription modes:

1. manual
2. by importing (subscription only),
3. via a web form

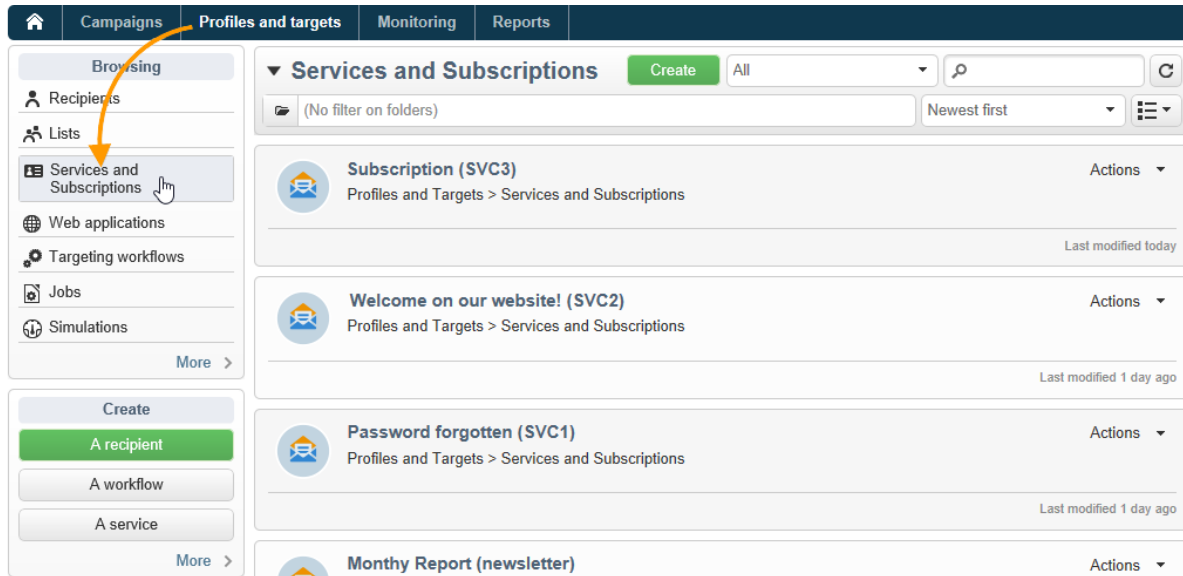
#### **NOTE**

A sample to create a subscription form with double opt-in is detailed in [this section](#).

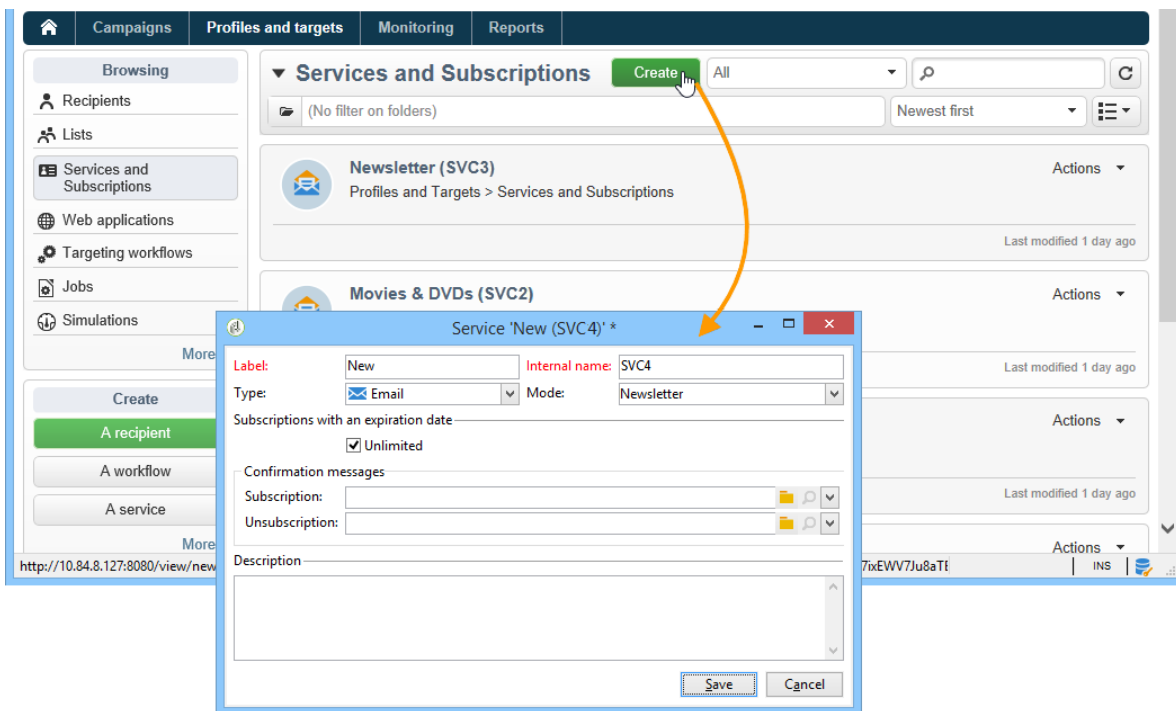
### Creating an information service

You can create and manage subscriptions to information services with associated confirmation messages or automatic deliveries to subscribers.

To access the information services map, open the **Profiles and Targets** tab and click the **Services and Subscriptions** link.



To edit an existing service, click its name. To create a service, click the **Create** button located above the list.



- Enter the name of the service in the **Label** field and select the delivery channel: email, mobile, Facebook, Twitter, or mobile applications.

**NOTE**

Facebook and Twitter subscriptions are detailed in [this section](#). Mobile application subscriptions are detailed in [About mobile app channel](#).

- For an Email type service, select the **Delivery mode**. The possible modes are: **Newsletter** or **Viral**.
- You can send **confirmation messages** for a subscription or unsubscription. To do this, select the delivery templates to be used to create the corresponding deliveries from the **Subscription** and **Unsubscription** fields. These templates must be configured with a **Subscription** type target mapping, without a defined target. See section [About email channel](#).
- By default, subscriptions are unlimited. You can deselect the **Unlimited** option to define a validity duration for the service. The duration can be specified in days (**d**) or months (**m**).

Once the service has been saved, it is added to the Services and Subscriptions list: Click its name to edit it. Several tabs are available. The **Subscriptions** tab lets you look at the list of subscribers to the information service (**Active subscriptions** tab) or the subscription/unsubscription history (**History** tab). You can also add and delete subscribers from this tab. See [Adding and deleting subscribers](#).

Service 'Newsletter (SVC4)'

Dashboard Edit Subscriptions

Recipient subscriptions Filters Add Delete Det

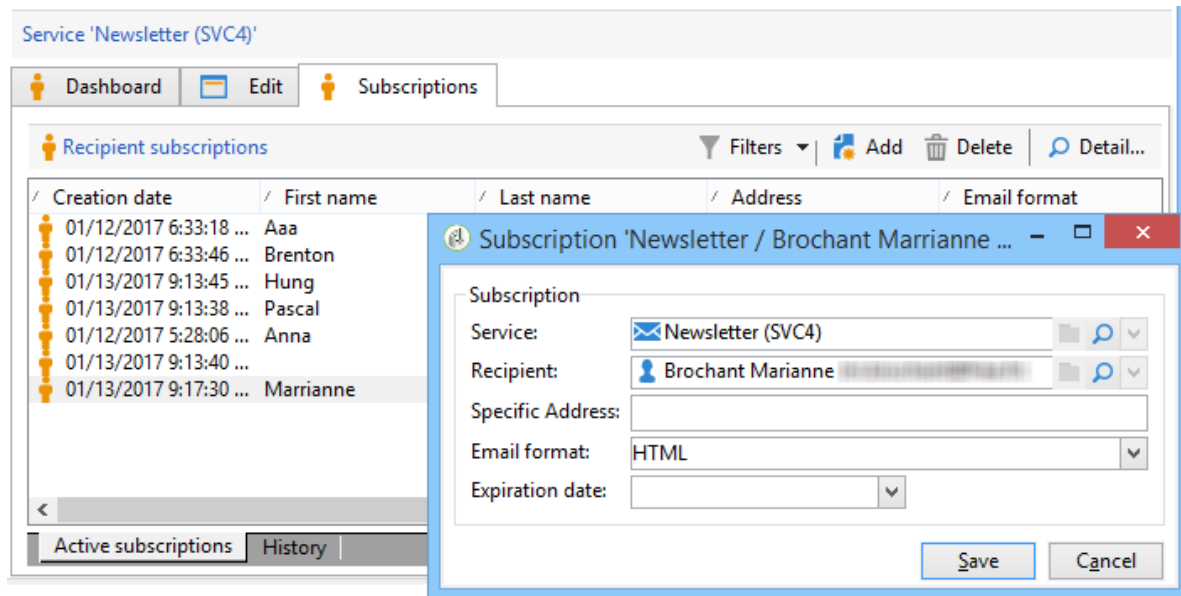
Creation date	First name	Last name	Address
01/12/2017 6:...	Marianne	Brochant	m.brochant@
01/12/2017 6:...	Diego	Casals	dcasals@gma
01/13/2017 9:...	Grégory	Caseneuve	caseneuveg@
01/13/2017 9:...	Michelle	Corette	miche.corette
01/12/2017 5:...	Linn	Damon	linn.damon@
01/13/2017 9:...	Edmond	Dantès	edantes@neo
01/13/2017 9:...	Cecilia	Dapunte	cecilia.dapuer
01/13/2017 9:...	Peter	Braun	peter.braun@

< 8 / 8 >

Active subscriptions History

The **Detail...** button lets you look at the subscription properties for the selected recipient.

You can modify the subscription properties for a recipient.



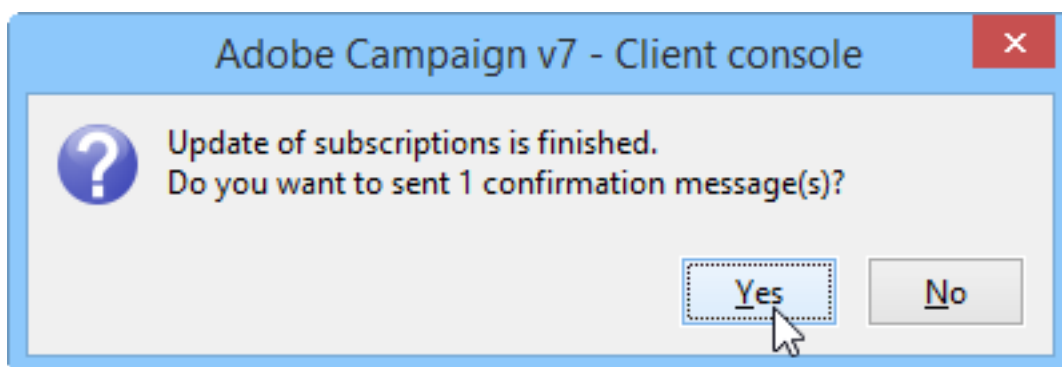
On the dashboard, click the **Reports** tab to track subscriptions: changes in subscription levels, total number of subscribers, etc. You can archive reports and look at histories from this tab.

## Adding and deleting subscribers

From the **Subscriptions** tab of an information service click **Add** to add subscribers. You can also right-click the list of subscribers and select **Add**. Select the folder in which the profiles to be subscribed are stored, and then select the profiles to subscribe and click **OK** to validate.

To delete subscribers, select them and click **Delete**. You can also right-click the subscriber list and select **Delete**.

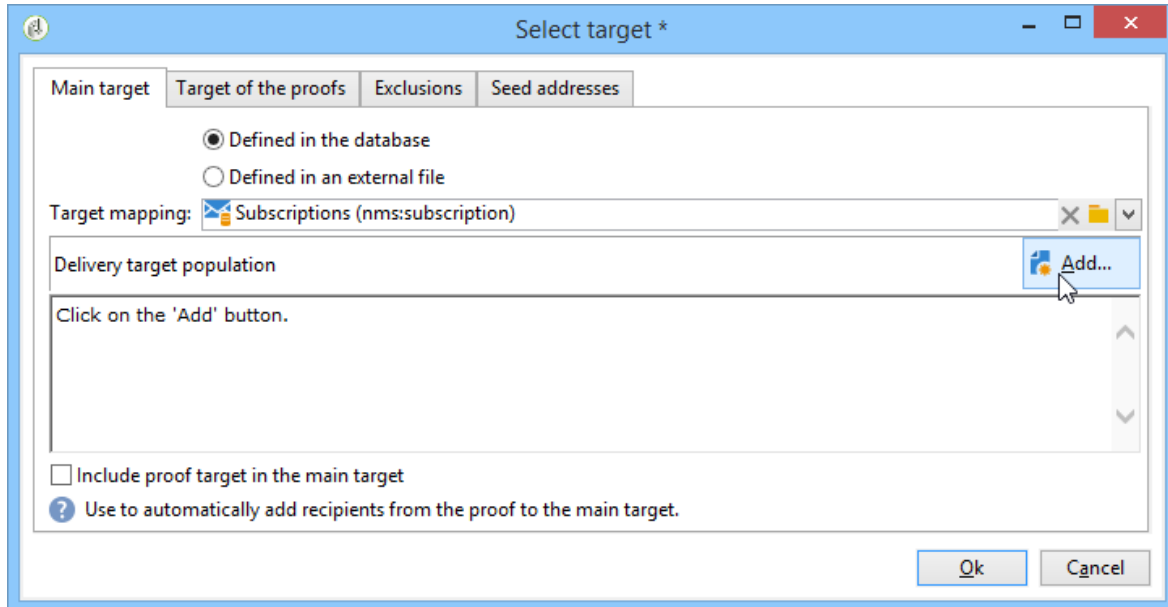
In both cases, you can send a confirmation message to the users concerned if a delivery template for unsubscriptions has been attached to the service (see [Creating an information service](#)). A warning lets you validate or not validate this delivery:



See [Subscription and unsubscription mechanisms](#).

## Delivering to the subscribers of a service

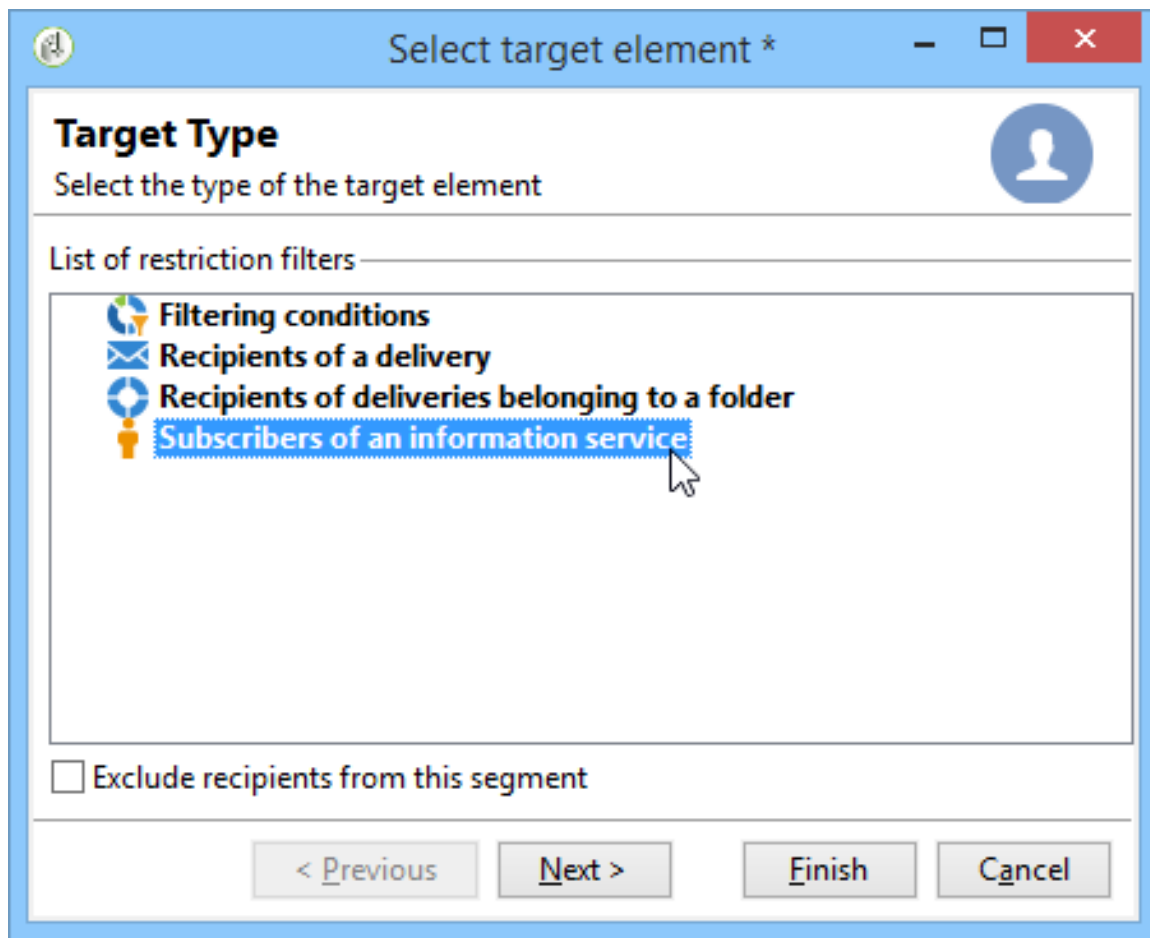
To deliver to the subscribers of an information service, you can target the subscribers to the information service concerned, as in the following example:



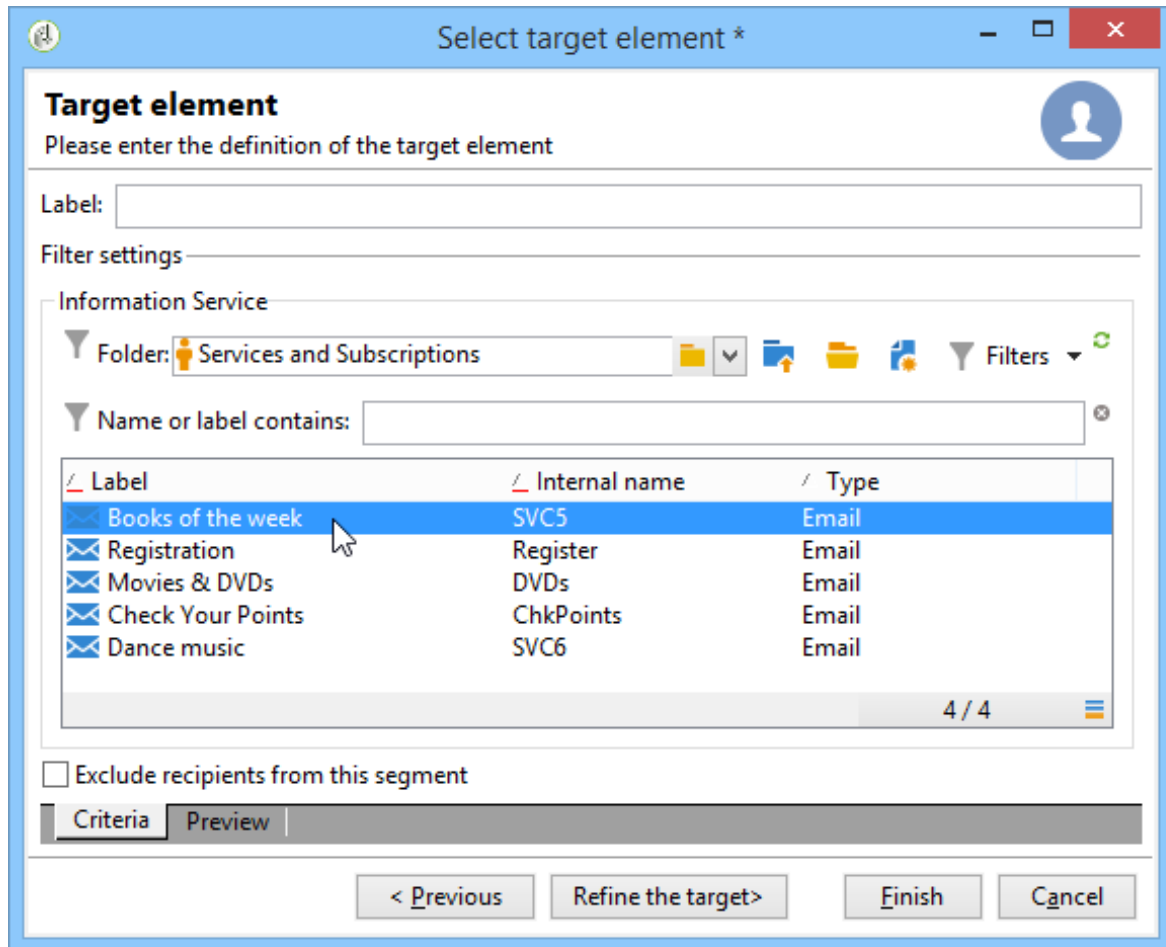
### CAUTION

The target mapping must be **Subscriptions**.

Select **Subscribers of an information service** and click **Next**.



Select the targeted information service and click **Finish**.



The **Preview** tab lets you view the list of subscribers to the selected information service.

## Subscription and unsubscription mechanisms

You can set up subscription and unsubscription mechanisms to automate the processes and subscriber management.

### NOTE

You can send a confirmation message to new subscribers.

The content of this message is defined in the information service configuration via the **Subscription** or **Unsubscription** fields.

The confirmation messages are created via the delivery templates specified in these fields. These target mappings must be **Subscriptions**.

Service 'Newsletter (newsletter)' \* Save ✕ 🔍

Dashboard **Edit** Subscriptions

Label:  Internal name:

Type:  Mode:

Subscriptions with an expiration date

Unlimited

Confirmation messages

Subscription:  📁 🔍

Unsubscription:  📁 🔍

Description

## Subscribing a recipient to a service

To register recipients for an information service, you can:

- Manually add the service: to do this, from the **Subscriptions** tab of their profile, click **Add** and select the information service concerned.

For more on this, refer to the section on profile editing in [this section](#).

- Automatically subscribe a set of recipients to this service. The list of recipients can come from a filtering operation, a group, a folder, an import, or a direct selection using the mouse. To subscribe these recipients, select the profiles and right-click. Select **Actions > Subscribe selection to a service...**, select the service concerned, and launch the operation.
- Import recipients and subscribe them automatically to an information service. To do this, select the service concerned in the last step of the import wizard.

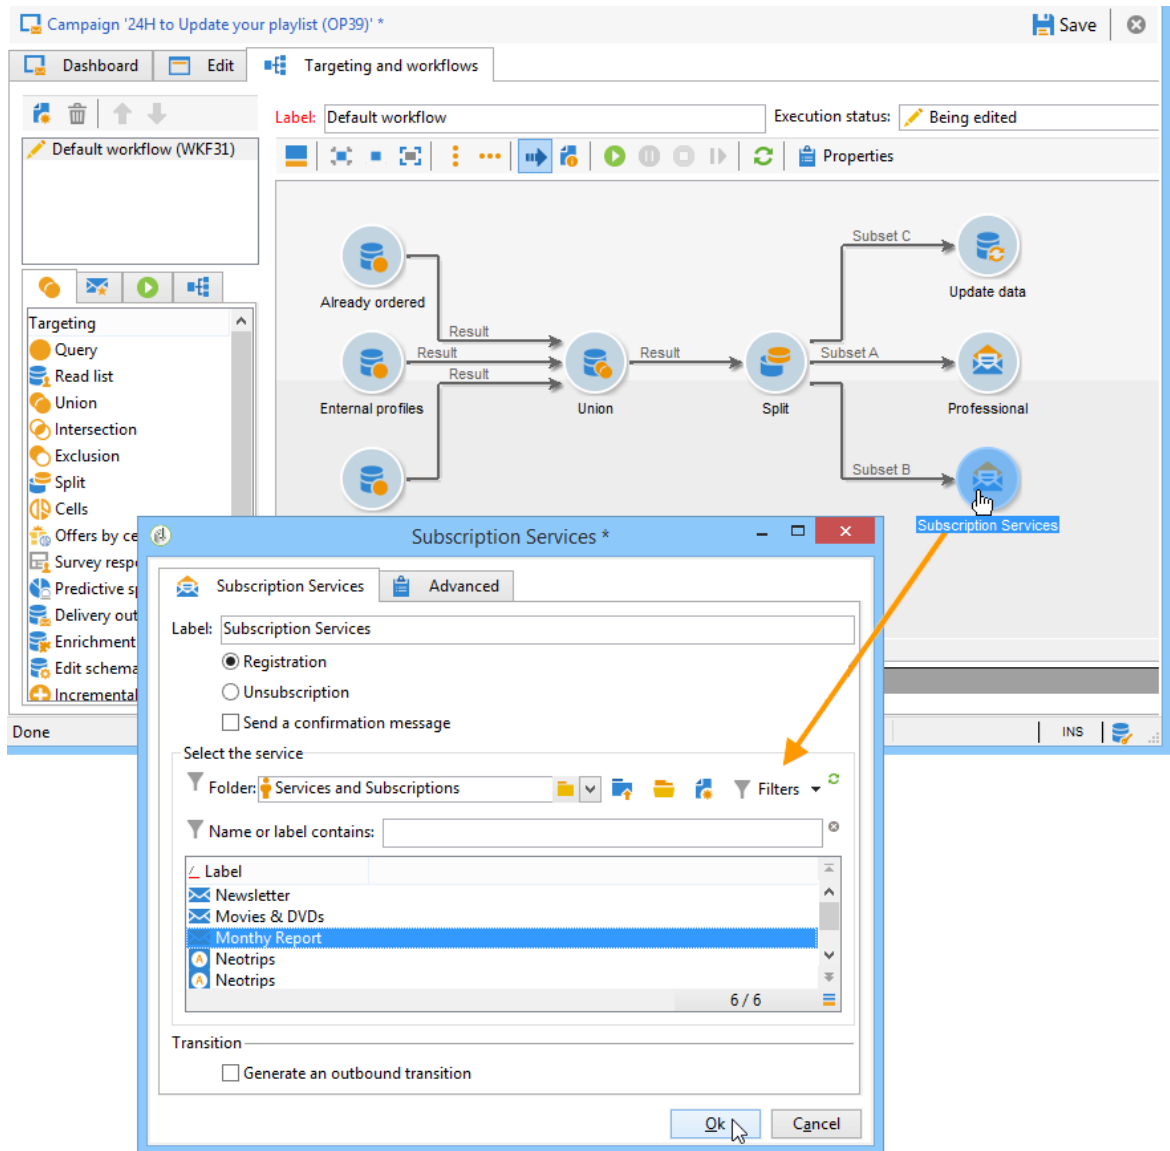
For more on this, refer to [this section](#).

- Use a web form so that recipients can subscribe to a service.

For more on this, refer to [this section](#).

- Creating a targeting workflow and using a **Subscription service** box.





Workflows and how to use them is detailed in [this section](#).

## Unsubscribing a recipient from a service

### Manual unsubscribing

email deliveries must contain an unsubscription link, by law. Recipients can click this link to update their profile and be excluded from the targets of future deliveries.

The default unsubscription link is inserted via the last button in the toolbar of the content editor provided in the delivery wizard (see [About personalization](#)). When the recipient clicks this link, the profile is added to the denylist (opt-out), meaning that this recipient will no longer be targeted by any delivery action.

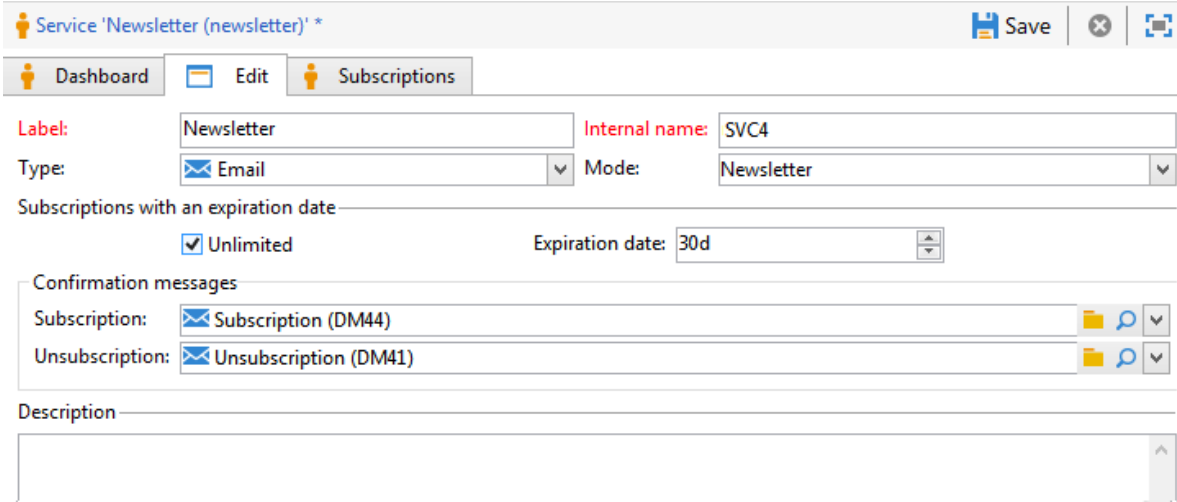
Recipients can, however, choose to unsubscribe from a service without unsubscribing from all services. To allow this, you can use a web form (refer to [this section](#)) or insert a personalized unsubscription link (see [Personalization blocks](#)).

You can also unsubscribe a recipient manually from the recipient profile. To do this, click the **Subscriptions** tab of the recipient concerned, select the information service(s) concerned, and click **Delete**.

You can finally unsubscribe one or more recipients via the information service concerned. To do this, click the **Subscriptions** tab of the service, select the recipients concerned and click **Delete**.

### Automatic unsubscription

An information service can have a limited duration. Recipients will be unsubscribed automatically when the period of validity has expired. This period is specified in the **Edit** tab of the service properties. It is expressed in days.



Service 'Newsletter (newsletter)' \*

Dashboard Edit Subscriptions

Label: Newsletter Internal name: SVC4

Type: Email Mode: Newsletter

Subscriptions with an expiration date

Unlimited Expiration date: 30d

Confirmation messages

Subscription: Subscription (DM44)

Unsubscription: Unsubscription (DM41)

Description

You can also set up an unsubscription workflow for a population. To do this, follow the same procedure as for a subscription workflow, but select the **Unsubscription** option. See [Subscribing a recipient to a service](#).

### Subscriber tracking

You can track the changes in subscriptions to the information services using the **Reports** link on the dashboard.

## 1.2 Get started with Campaign server configuration

This chapter details server-side configurations that can be performed to match your needs and your environment specificities.

## Restrictions

These procedures are restricted to **on-premise/hybrid** deployments and require Administration permissions.

For **hosted** deployments, server-side settings can be configured by Adobe only. However, some settings can be set up within [Campaign Control Panel](#), such as IP allowlist management or URL permissions. [Learn more](#).

For more information, refer to these sections:

- [Control Panel documentation](#)
- [Hosting models](#)
- [Campaign Classic On-premise & Hosted capability matrix](#)

## Configuration files

Campaign Classic configuration files are stored in the **conf** folder of the Adobe Campaign installation folder. The configuration is spread over two files:

- **serverConf.xml**: general configuration for all instances. This file combines the technical parameters of the Adobe Campaign server: these are shared by all instances. The description of some of these parameters is detailed below. The different nodes and parameters are listed in this [section](#).
- **config-<instance>.xml** (where **instance** is the name of the instance): specific configuration of the instance. If you share your server among several instances, please enter the parameters specific to each instance in their relevant file.

## Configuration scope

Configure or adapt Campaign server depending on your needs and configuration. You can:

- Secure the [Internal identifier](#)
- Enable [Campaign processes](#)
- Configure [URL Permissions](#)
- Define [Security Zones](#)
- Configure [Tomcat settings](#)
- Customize [Delivery parameters](#)
- Define [Dynamic page security and relays](#)
- Restrict the list of [Allowed external commands](#)
- Set up [Redundant tracking](#)
- Manage [High availability and workflow affinities](#)
- Configure file management - [Learn more](#)
  - Limit upload files format
  - Enable access to public resources
  - Configure Proxy connection
- [Automatic process restart](#)

## Internal identifier

The **internal** identifier is a technical login to be used for installation, administration and maintenance purposes. This login is not associated with an instance.

Operators connected using this login will have all the rights on all instances. This login won't have a password in the case of a new installation. You must manually define this password.

Use the following command:

```
nlservice config -internalpassword
```

Copy

Toggle Text Wrapping

The following information is then displayed. Enter and confirm the password:

```
17:33:57 > Application server for Adobe Campaign Classic (7.X YY.R build
XXX@SHA1) of DD/MM/YYYY
Enter the current password.
Password:
Enter the new password.
Password: XXXX
Confirmation: XXXX
17:34:02 > Password successfully changed for account 'internal'
(authentication mode 'nl')
```

Copy

Toggle Text Wrapping

## Enable processes

Adobe Campaign processes on the server are enabled (and disabled) via the **config-default.xml** and **config-<instance>.xml** files.

To apply the changes to these files, if the Adobe Campaign service is started, you must run the **nlservice config -reload** command.

There are two types of processes: multi-instance and single instance.

- **multi-instance**: one single process is started for all instances. This is the case for **web**, **syslogd** and **trackinglogd** processes.

Enablement can be configured from the **config-default.xml** file.

Declaring an Adobe Campaign server to access client consoles and for redirection (tracking):

```
vi n16/conf/config-default.xml
<web args="-tomcat" autoStart="true"/>
<!-- to start if the machine is also a redirection server -->
<trackinglogd autoStart="true"/>
```

Copy

Toggle Text Wrapping

In this example, the file is edited using a **vi** command in Linux. It can be edited using any **.txt** or **.xml** editor.

- **mono-instance**: one process is started for each instance (modules: **mta**, **wfserver**, **inMail**, **sms** and **stat**).

Enablement can be configured using the configuration file of the instance:

```
config-<instance>.xml
```

Copy

Toggle Text Wrapping

Declaring a server for delivery, executing workflow instances and recovering bounce mail:

```
<mta autoStart="true" statServerAddress="localhost"/>
<wfserver autoStart="true"/>
<inMail autoStart="true"/>
<stat autoStart="true"/>
```

Copy

Toggle Text Wrapping

## Campaign data storage

You can configure the storage directory (**var** directory) of Adobe Campaign data (logs, downloads, redirections, etc.). To do this, use the **XTK\_VAR\_DIR** system variable:

- In Windows, indicate the following value value in the **XTK\_VAR\_DIR** system variable
- **D:\log\AdobeCampaign**

Copy

Toggle Text Wrapping

- In Linux, go to the **customer.sh** file and indicate: **export XTK\_VAR\_DIR=/app/log/AdobeCampaign**.

For more on this, refer to [Personalize parameters](#).

## Dynamic page security and relays

By default, all dynamic pages are automatically related to the **local** Tomcat server of the machine whose Web module has started. This configuration is entered in the `<url>` section of the query relay configuration for the **ServerConf.xml** file.

You can relay execution of the dynamic page on a **remote** server; if the Web module is not activated on the computer. To do this, you must replace the **localhost** with the name of the remote computer for JSP and JSSP, Web applications, reports and strings.

For more on the various parameters available, refer to the **serverConf.xml** configuration file.

For JSP pages, the default configuration is:

```
<url relayHost="true" relayPath="true" targetUrl="http://localhost:8080"
urlPath="*.jsp"/>
```

Copy

Toggle Text Wrapping

Adobe Campaign uses the following JSP pages:

- `/nl/jsp/soaprouter.jsp`: client console and Web services connections (SOAP APIs),
- `/nl/jsp/m.jsp`: mirror pages,
- `/nl/jsp/logon.jsp`: Web-based access to reports and to deployment of the client console,
- `/nl/jsp/s.jsp` : Using viral marketing (sponsoring and social networks).

The JSSPs used for the Mobile App Channel are as follows:

- `nms/mobile/1/registerIOS.jssp`
- `nms/mobile/1/registerAndroid.jssp`

### Example:

It's possible to prevent client machine connections from the outside. To do this, simply restrict the execution of `soaprouter.jsp` and only authorize the execution of mirror pages, viral links, web forms and public resources.

The parameters are as follows:

```
<url IPMask="<IP_addresses>" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="*.jsp"/>
<url IPMask="<IP_addresses>" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="*.jssp"/>
<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="m.jsp"/>
```

```

<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="s.jsp"/>
<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="webForm.jsp"/>
<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="/webApp/pub*"/>
<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="/jssp/pub*"/>
<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="/strings/pub*"/>
<url IPMask="" deny="" hostMask="" relayHost="true"
relayPath="true" targetUrl="http://localhost:8080" timeout=""
urlPath="/interaction/pub*"/>
<url IPMask="" deny="true" hostMask="" relayHost="false"
relayPath="false" targetUrl="http://localhost:8080" timeout=""
urlPath="*.jsp"/>
<url IPMask="" deny="true" hostMask="" relayHost="false"
relayPath="false" targetUrl="http://localhost:8080" timeout=""
urlPath="*.jssp"/>

```

Copy

Toggle Text Wrapping

In this example, the `<IP_addresses>` value coincides with the list of IP addresses (separated by comas) authorized to use the relay module for this mask.

## NOTE

Values shall be adapted according to your configuration and your network constraints, especially if specific configurations have been developed for your installation.

## Manage HTTP headers

By default, all HTTP headers are not relayed. You can add specific headers in the replies sent by relay. To do this:

1. Go to the **serverConf.xml** file.
2. In the `<relay>` node, go to the list of relayed HTTP headers.
3. Add a `<responseheader>` element with the following attributes:
  - **name:** header name
  - **value:** value name.

For example:

```
<responseHeader name="Strict-Transport-Security" value="max-age=16070400; includeSubDomains"/>
```

Copy

Toggle Text Wrapping

## Restrict authorized external commands

From build 8780, technical administrators can restrict the list of authorized external commands that can be used in Adobe Campaign.

To do that, you need to create a text file with the list of commands that you want to prevent from using, for example:

```
In
dd
openssl
curl
wget
python
python3
perl
ruby
sh
```

Copy

Toggle Text Wrapping

### IMPORTANT

This list is not exhaustive.

In the **exec** node of the server configuration file, you need to reference the previously created file in the **blacklistFile** attribute.

**For Linux only:** in the server configuration file, we recommend that you specify a user dedicated to executing external commands to enhance your security configuration. This user is set in the **exec** node of the configuration file. All the parameters available in the **serverConf.xml** are listed in this [section](#).

### NOTE

If no user is specified, all commands are executed in the user context of the Adobe Campaign instance. The user must be different than the user running Adobe Campaign.

For example:

```
<serverConf>
  <exec user="theUnixUser" blacklistFile="/pathtothefile/blacklist"/>
```



```
</serverConf>
```

Copy

Toggle Text Wrapping

This user needs to be added to the sudoer list of the 'neolane' Adobe Campaign operator.

### IMPORTANT

You should not use a custom sudo. A standard sudo needs to be installed on the system.

## Redundant tracking

When multiple servers are used for redirection, they must be able to communicate with each other via SOAP calls in order to share information from the URLs to be redirected. At the time of delivery start-up, it is possible that not all the redirection servers will be available; therefore they might not have the same level of information.

### NOTE

When using the standard or enterprise architecture, the main application server must be authorized to upload tracking information on each computer.

The URLs of the redundant servers must be specified in the redirection configuration, via the `serverConf.xml` file.

### Example:

```
<spareserver enabledIf="$(hostname)!='front_srv1'" id="1"
url="http://front_srv1:8080" />
<spareserver enabledIf="$(hostname)!='front_srv2'" id="2"
url="http://front_srv2:8080" />
```

Copy

Toggle Text Wrapping

The `enableIf` property is optional (empty by default) and allows you to enable the connection only if the result is true. This lets you obtain an identical configuration on all redirection servers.

To obtain the hostname of the computer, run the following command: `hostname -s`.

## High availability workflows and affinities

You can configure several workflow servers (`wfserver`) and distribute them on two or more machines. If you choose this type of architecture, configure the connection mode of the load balancers according to the Adobe Campaign access.

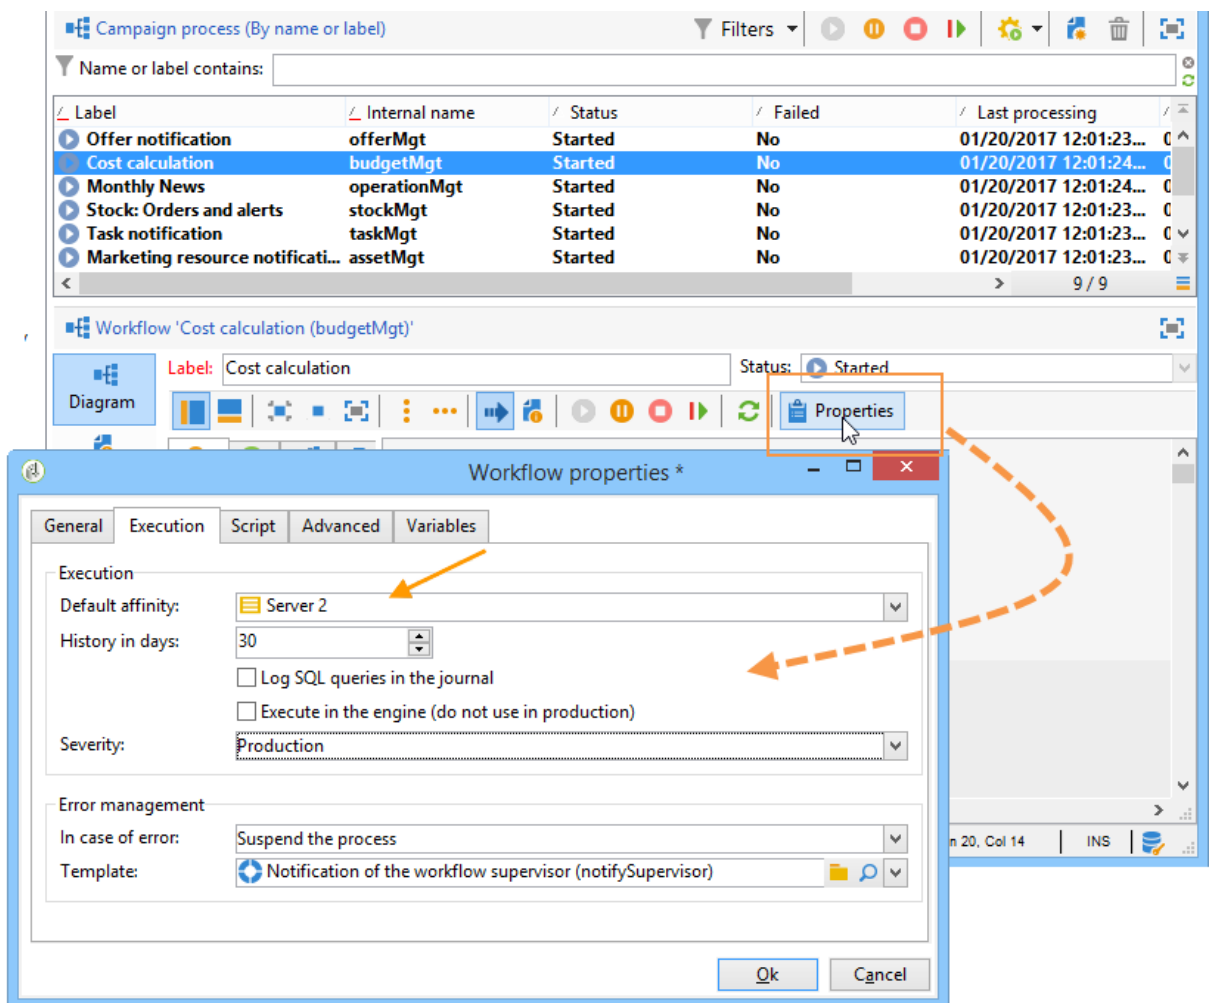
For access from the web, select the **load balancer** mode to limit connection times.

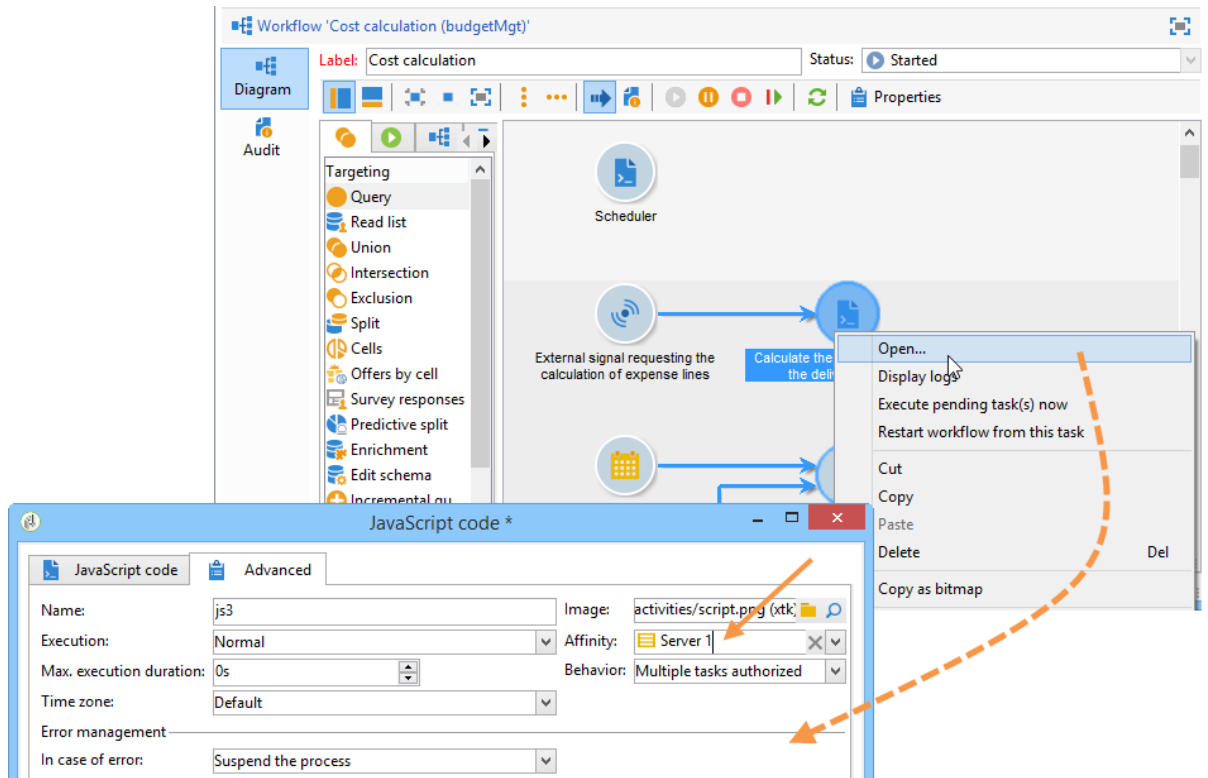
If accessing via the Adobe Campaign console, choose **hash** or **sticky ip** mode. This lets you maintain the connection between the rich client and the server and prevent a user session from being interrupted during an import or export operation, for example.

You can choose to force the execution of a workflow or a workflow activity on a particular machine. To do this, you must define one or more affinities for the concerned workflow or activity.

1. Create the affinities of the workflow or activity by entering them in the **Affinity** field.

You can choose any affinity name, but make sure you do not use spaces or punctuation marks. If you use different servers, specify different names.





The drop-down list contains formerly used affinities. It is completed over time with the different entered values.

2. Open the `nl6/conf/config-<instance>.xml` file.
3. Modify the line which matches the **wfserver** module as follows:

4. `<wfserver autoStart="true" affinity="XXX," />`

Copy

Toggle Text Wrapping

If you define several affinities, they must be separated by commas without any spaces:

`<wfserver autoStart="true" affinity="XXX,YYY," />`

Copy

Toggle Text Wrapping

The comma following the name of the affinity is necessary for the execution of workflows for which no affinity is defined.

If you wish to execute only workflows for which an affinity is defined, do not add a comma at the end of the list of your affinities. For example, modify the line as follows:

`<wfserver autoStart="true" affinity="XXX"/>`

Copy

Toggle Text Wrapping

## Automatic restart

By default, the different Adobe Campaign processes restart automatically at 6am (server time) every day.

However, you can change this configuration.

To do this, go to the `serverConf.xml` file, located in the `conf` repository of your installation.

Each process configured in this file has a `processRestartTime` attribute. You can modify the value of this attribute to adapt the restart time of each process according to your needs.

### IMPORTANT

Do not delete this attribute. All processes must be restarted every day.

## 1.3 Get started with delivery monitoring

Monitoring your deliveries after they have been sent is a key step to ensure your marketing campaigns are efficient and reach out to your customers.

In this section, you will learn more about the information you can monitor after sending a delivery, as well as understand how delivery failures and quarantines are managed.

### Monitor your deliveries

The list of deliveries allows you to see all created deliveries into one single location.

For each delivery, a dedicated dashboard is available. It allows you to monitor eventual issues encountered during the sending as well as various types of information regarding the delivery: reports, mirror pages, exclusions, tracking logs, rendering, etc.

- [Accessing the list of deliveries](#)
- [Delivery dashboard](#)

### Secure delivery performances

Several guidelines should be followed in order to ensure your deliveries perform well. Common issues you may encounter when sending deliveries are also available to help you send deliveries efficiently.

- [Delivery performances and best practices](#)
- [Delivery sending troubleshooting](#)

## Understand delivery failures

When a message cannot be sent to a profile, the remote server automatically sends an error message, which is picked up by the Adobe Campaign platform and qualified to determine whether or not the email address or phone number should be quarantined.

[Understanding delivery failures](#) is a key step to help you improve your marketing campaigns.

## Understand quarantine management

Adobe Campaign manages a list of quarantined addresses. Recipients whose address is quarantined are excluded by default during delivery analysis, and will not be targeted.

In [this section](#), you will find information on how to identify and manage quarantined addresses, and learn more about the conditions for sending an address to quarantine.

# 1.4 Hardware sizing recommendations

## Overview

### CAUTION

This article is provided as a general example guide only. You must engage with your Adobe Campaign Customer Success Manager to measure the exact sizing of your deployment before starting your Campaign project. **Do not** acquire or deploy any infrastructure or hardware until this is done.

This document provides general recommendations for Adobe Campaign Classic v7 deployment at your on-premise data center or virtualized cloud environment. This type of deployment, referred to as **hybrid** or **mid-sourcing**, places the Campaign marketing instance and marketing database under your operational control, while using Adobe Cloud Messaging services to send emails, SMS or SMPP messages, and collect email open, bounce, and click tracking data.

The marketing instance is the portion of the Adobe Campaign architecture that drives all marketing activity, and stores all recipient data and analytic data returned by campaigns. The marketing instance is a set of on-premise servers running Adobe Campaign services, and a relational database.

## CAUTION

The information in this document do not apply if you are using a fully hosted Adobe Campaign instance (deployed in Adobe Cloud Services).

Software compatibility is detailed in the [Compatibility Matrix](#).

## Scenarios

Deployment diagrams and hardware sizing recommendations are provided for three representative scenarios:

1. [Moderate-Size](#) – 5 millions active recipients in the system
2. [Large-Size](#) – 20 million active recipients in the system
3. [Enterprise](#) – 50 million active recipients, with transactional messaging

## Assumptions

This document also assumes the following types of usage for all three scenarios:

- Large email campaigns are sent twice a week, to approximately 50% of your active recipients
- Direct mailings are generated once per month for each recipient in the system
- SMS messages are sent to approximately 10% of your active recipients each month
- The database schema defining each recipient has been extended with one additional table, containing about 200 bytes of data for each recipient
- Adobe Campaign Interaction module is used to add offers to outgoing email
- Email tracking data is retained in the Campaign system for 90 days

## General guidelines

Campaign is a database-centric application, and database server performance is critical. Running workflows, segmentation, tracking data uploads, inbound Interactions, analytics and other activities all generate database activity. In general, the size and frequency of these operations determine the size of your database servers.

The application servers in your marketing instance require enough CPU and memory to run workflows and respond to SOAP API calls, including requests from Campaign Console users. CPU requirements can be significant for workflows that use outbound interactions with complex offer rules, workflows that execute custom Javascript, and web applications with high traffic levels.

Campaign web applications can also be deployed on the marketing instance App servers, or on separate web server systems. Since web application workloads conflict with critical workflows and Campaign Console users, web applications and inbound interactions can be deployed to separate servers, to ensure that core Campaign functionality runs reliably with good performance.

For security and availability, Adobe recommends separating the traffic of the Internet from the traffic generated by the business users. For that reason, the diagrams contain two groups of servers: the Web server (Internet facing Web1 and Web2), and the App servers (business processes App1 and App2).

It is a legal requirement for commercially email senders to have a functional opt-out web page. Adobe recommends having redundant machine in each group server for fail-over scenarios. It is especially true if Adobe Campaign hosts the opt-out pages.

## Reverse proxies

The Campaign architecture enforces high security by using SSL over HTTP (HTTPS) to communicate between your marketing instance and Adobe Cloud Messaging. Security, reliability, and availability are enforced by using reverse proxies in a “demilitarized zone” (DMZ) subnet to isolate and secure the marketing instance servers and database.

## Load balancer

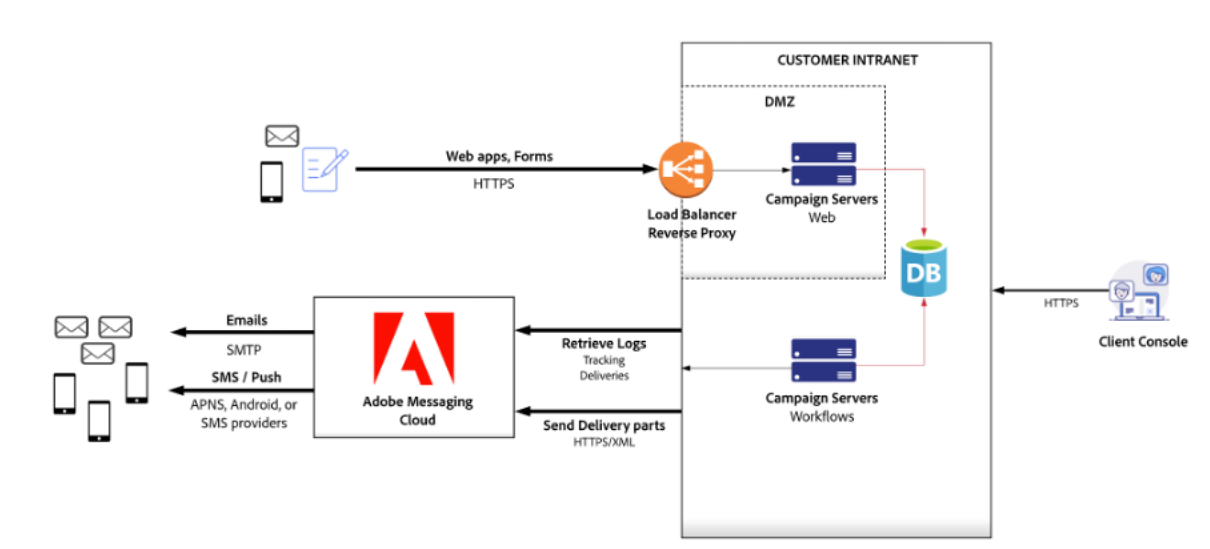
The load balancer for the App servers is set up in an active/passive configuration, with HTTPS terminated at the proxy. The load balancer for the Web servers is set up in an active/active configuration, with HTTPS terminated at the proxy.

Adobe provides you with the exclusive list of URL paths that can be relayed to the Adobe Campaign server in your deployment environment.

## Architecture

The general architecture is almost identical regardless of the volumes. The security and high availability requirements dictate a minimum of four servers; two servers if no WebApps are used. The difference in the configuration mainly varies in the hardware configuration like CPU core and memory.

## Scenario 1: Moderate-Size Deployment



Estimated volume:

Channel	Volume
Active recipients	5 millions
Email	4.2 million/month
Direct mail	1 million/month
Mobile SMS	100,000/month
Peak daily email volume	500

For these volumes, a pair of Adobe Campaign application server systems provides all functionality for Adobe Campaign Client users and workflow execution. For 5 millions active recipients and this email volume, application server workloads are not CPU or I/O intensive; most of the stress is on the database.

The Adobe Campaign Web servers are shown in the secure zone.

### **Web & Application Servers**

This scenario recommends installing Adobe Campaign on four machines, with the following specification:

**3Ghz+ quad-core CPU, 8-GB RAM, RAID 1 or 10, 2 x 80-GB SSD**

These systems create the marketing instance Application Server, which directly supports your Campaign Console users and executes the campaign workflows.

Reverse proxies in a DMZ load balance traffic to the Adobe Campaign Web servers. There is no requirement to install the Adobe Campaign software stack on proxy machines; any reverse proxy software or network equipment can be used.

Subscription opt-in/opt-out and preference center features can be provided by Campaign or by your own website. If you choose to implement this functionality on your website, you must ensure that preference and subscription information is propagated to the Campaign marketing database. It is normally done by creating extraction files which are uploaded by Campaign workflows automatically.

Disk space consumption on Application Servers depends on the retention period of files exchanged with third-party service providers (for example, print vendors for Direct Mail), and on the size and retention of imported flat files, like subscription or preference updates from your website, or extracts from your own CRM or marketing systems.

### **Database**

Hardware recommendations for the database server are as follows:

**3Ghz+ 4-core CPU, 16-GB RAM, RAID 1 or 10, 128GB SSD minimum**



The memory estimate assumes full caching of approximately 500,000 recipients for a large campaign launch, plus RDBMS buffer space for executing workflows, importing tracking data, and other concurrent activities.

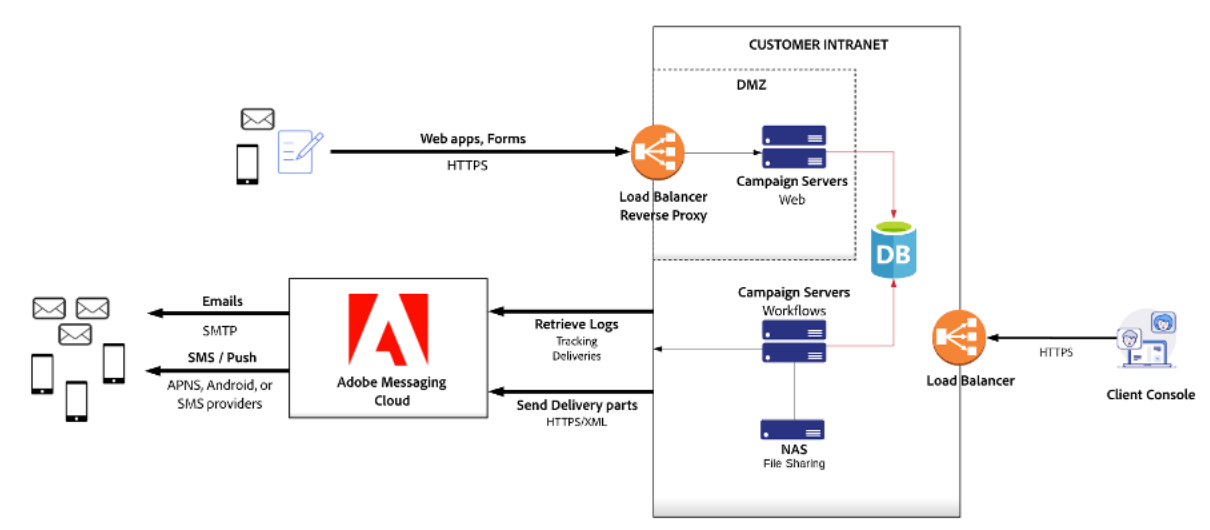
It is estimated that the disk space required on the database to store all Adobe Campaign technical data (campaigns, tracking, working tables, and so on) is approximately 35 GB based on a retention period of three months. If you choose to retain tracking data for 6 months, database size increases to approximately 40 GB, and 12-month retention increases database size to approximately 45 GB. Recipient data consumes about 5 GB for this environment.

**CAUTION**

This estimate does not include any additional customer data. If you are planning to replicate additional columns or tables of customer data into the Adobe Campaign database, then you have to estimate the additional disk space requirement for it. Uploaded segments/lists also require more storage, depending on their size, frequency, and retention period.

Also consider that due to the volume of information processed daily, the IOPS of the database server is critical. For example, on a peak day, you can deploy campaigns targeting a total of 500,000 recipients. To execute each campaign, Adobe Campaign inserts 500,000 records into a table containing roughly 12 million records (the delivery log table). To provide acceptable performance during the campaign deployment, Adobe recommends a minimum of 60,000 4-KB Random read/write IOPS for this scenario.

**Scenario 2: Large-Size Deployment**



Estimated volume:

Channel	Volume
Active recipients	20 millions
Email	42 millions/month
Direct mail	10 millions/month

Mobile SMS	1,000,000/month
Peak daily email volume	5,000,000

## Web & Application Servers

In this scenario, Adobe recommends installing Adobe Campaign on four machines, two application servers and two web servers, with the following specification:

### **3Ghz+ quad-core CPU, 8-GB RAM, RAID 1 or 10, 80-GB SSD**

The Application servers directly support Campaign Console users and the execution of campaign workflows. This functionality is deployed on two identical servers for high-availability, sharing a Network-Attached Storage (NAS) file system to enable failover.

The Web servers host Campaign web applications supporting the 10-million active recipients in the system.

Refer to [Scenario 1: Moderate-Size Deployment](#) for more comments on proxies, preference centers/subscription handling, and disk space usage.

## Database

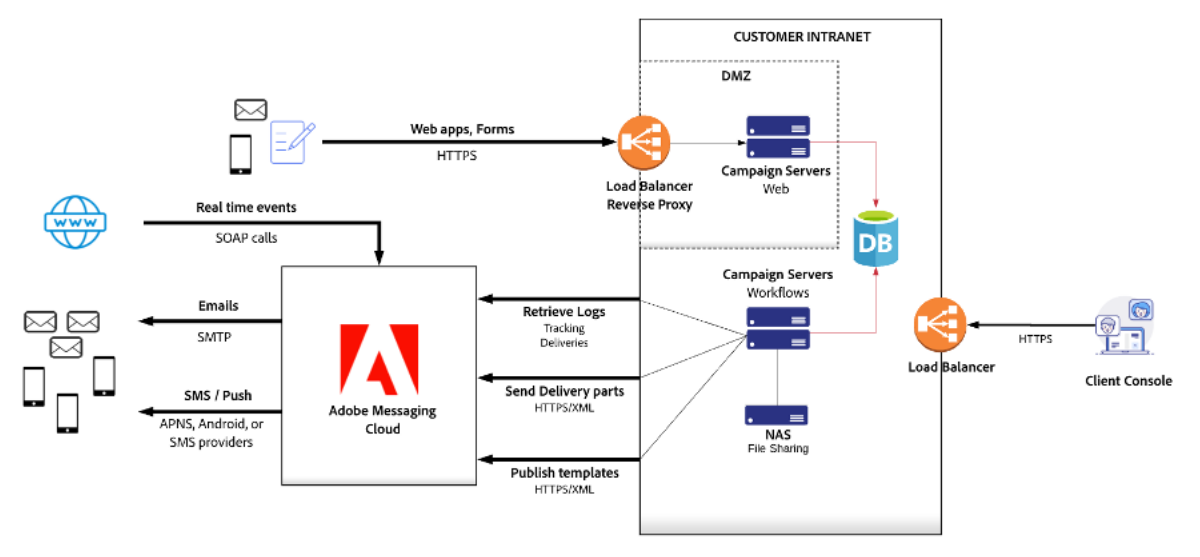
Hardware recommendations for the database server are as follows:

### **3Ghz+ 8-core CPU, 64-GB RAM, RAID 1 or 10, 2 x 320-GB SSD or RAID 10, 640 GB SSD minimum**

The memory estimate assumes full caching of approximately 5,000,000 recipients for a large campaign launch, plus RDBMS buffer space for executing workflows, importing tracking data, and other concurrent activities.

It is estimated that the disk space required on the database to store all Adobe Campaign technical data (campaigns, tracking, working tables, and so on) is approximately 280 GB based on a retention period of 3 months. If you choose to retain tracking data for 6 months, database size increases to approximately 450 GB, and 12-month retention increases database size to approximately 900 GB. Recipient data consumes about 15 GB for this environment.

## Scenario 3: Enterprise Deployment with Message Center



Estimated volume:

Channel	Volume
Active recipients	50 millions
Email	108 millions/month
Direct mail	25 millions/month
Mobile SMS	2,5 millions/month
Transactional messages	2,5 millions/month
Peak daily email volume	2,5 millions

The deployment supporting 50 million recipients is essentially the same as shown in [Scenario 2](#): Campaign web application traffic is routed to Campaign web servers, so the bursts of web traffic after large campaign launches does not impact Campaign workflows and Client Console users.

This deployment also includes Message Center calls, driven from your own websites and applications.

## Web & Application Servers

In this scenario, Adobe recommends installing Adobe Campaign on four machines, as follows:

- Application servers  
**Two systems, 3Ghz+ quad-core CPU, 8-GB RAM, RAID 1 or 10, 80-GB SSD**
- Web servers  
**Two systems, 3Ghz+ quad-core CPU, 16-GB RAM, RAID 1 or 10, 80-GB SSD**

The Application servers directly support Campaign Console users and the execution of campaign workflows. This functionality is deployed on two identical servers for high-availability, sharing a Network-Attached Storage (NAS) file system to enable failover.

The Web servers host Campaign web applications supporting the 10-million active recipients in the system.

Refer to [Scenario 1: Moderate-Size Deployment](#) for more comments on proxies, preference centers/subscription handling, and disk space usage.

## Database

Hardware recommendations for the database server are as follows:

### **3Ghz+ 8-core CPU, 96-GB RAM, RAID 1 or 10, 1.5TB SSD minimum**

The memory estimate assumes full caching of approximately 12,500,000 recipients for a large campaign launch, plus RDBMS buffer space for executing workflows, importing tracking data, and other concurrent activities.

It is estimated that the disk space required on the database to store all Adobe Campaign technical data (campaigns, tracking, working tables, and so on) is approximately 700 GB based on a retention period of 3 months. If you choose to retain tracking data for 6 months, database size increases to about approximately 1.2TB, and 12-month retention increases database size to approximately 2TB. Recipient data consumes about 50 GB for this environment.

## Guidelines for Changing Assumptions

The assumptions made for these scenarios all have a significant impact on the hardware recommendations and deployment architecture. This section discusses guidelines around different assumptions. Contact the Adobe Campaign Consulting team for specific recommendations to meet your requirements.

- **Number of Recipients**  
Active recipients require both storage space and database buffer space, so more recipients generally require more memory and CPU capacity on the database server. Storage increases are relatively small for the recipients themselves, but can be significant for the event tracking data kept for email campaigns.
- **Email Campaign Size**  
The frequency of campaign launches has an impact on database server CPU requirements. Combined with direct mailing, inbound Interactions, and other workflows, segmentation operations for email campaigns put a significant load on the database server.
- **Direct Mail Frequency**  
The frequency of direct mailings can impact database server CPU requirements. Combined with campaign launches and other workflows, segmentation operations for direct mailings put a significant load on the database server.
- **SMS Message Volume**  
Like email campaign size, SMS message volume does not place large loads on Campaign servers located on-premise; the load is mostly on Adobe Cloud Messaging servers on the cloud. Segmentation for SMS campaigns, like email and direct mail,

can place a significant load on the marketing database. Therefore the frequency of SMS campaign launches and complexity of segmentation are more relevant than the volume of SMS messages.

- **Database Schema Complexity**

The amount of data for each active recipient requires both storage space and database buffer space, so more recipients generally require more memory and CPU on the database server. Complex schemas also require more tables to be joined for segmentation, so segmentation operations can run much slower, and require more database CPU and memory when data is spread across multiple tables.

Database server memory is estimated by ensuring that the database buffer pool can be large enough to contain all recipient data, plus temporary tables for running workflows, plus a margin for other database operations.

- **Outbound Interaction Usage**

Rules for Interaction in batch mode are evaluated in workflows that hand over all the calculation complexity to the database. The main factor of effort on the database is the total number of eligible offers computed during one engine call (target size X average number of offers per recipient before keeping the N best offers). The database server CPU speed is the first factor of performance.

- **Inbound Interactions or SOAP API Usage**

Inbound Interaction rules and offers are evaluated in the marketing database, requiring significant database server resources, especially CPU. Heavy usage of inbound Interactions or SOAP APIs requires separate web servers to separate the work load from running Campaign workflows.

- **Tracking Data Retention Period**

Increasing the retention of tracking data beyond 90 days requires more database storage, and can slow down the system because inserting new tracking data goes into large tables. Tracking data is not useful for campaign segmentation after 90 days, so the shorter retention period is recommended.

Tracking data should be moved into Adobe Analytics or another analytics system if you require long-term analysis of recipient marketing experience.

## Virtualization

All Campaign servers are good candidates for virtualization. Several issues must be addressed to ensure appropriate availability and performance.

- **Fail-Over Configuration**

Clustered servers, for example, redundant application servers under a load-balanced proxy, must be deployed on separate hardware to ensure that both VMs do not go down if there is hardware failure.

- **I/O Configuration**

Any recommended RAID configuration must be maintained for database security, to ensure that loss of a storage device does not cause loss of data.

- **I/O Performance**

The recommended IOPS rating for database storage must be respected. Cloud services like Amazon EC2 may not provide the required performance, and must be carefully evaluated. For example, Amazon EC2 provisioned SSD volumes are currently rated at 20,000 IOPS each. Learn more in [Amazon documentation](#), so a 4-volume RAID configuration would be rated at 80,000 IOPS, which may not be sufficient.

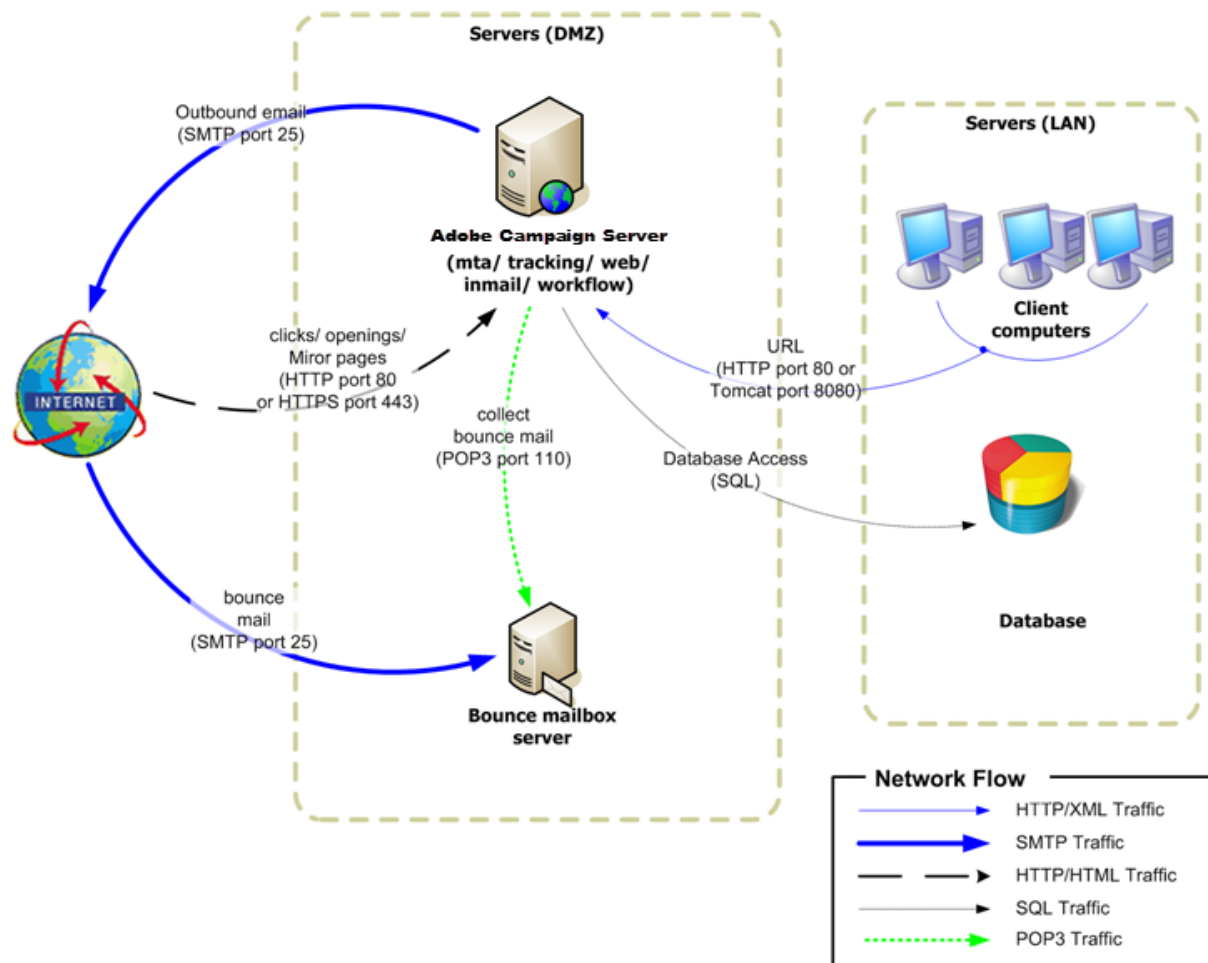
Adobe recommends performance testing for any virtualized deployment of Adobe Campaign before putting the system into production.

## 1.5 Standalone deployment

This configuration includes all components on the same computer:

- application process (web),
- delivery process (mta),
- redirection process (tracking),
- workflow process and scheduled tasks (wfserver),
- bounce mail process (inMail),
- statistics process (stat).

Overall communication between the processes is carried out according to the following schema:



This type of configuration can be run when managing lists of fewer than 100,000 recipients and with, for example, the following software layers:

- Linux,
- Apache,
- PostgreSQL,
- Qmail.

As the volume grows, a variant of this architecture moves the database server to another computer for better performance.

#### NOTE

An existing database server may also be used if it has sufficient resources.

## Features

### Advantages

- Fully standalone and low configuration cost (no billable licenses required if the open-source software listed below is used).
- Simplified installation and network configuration.

## Disadvantages

- A critical computer in case of incident.
- Limited bandwidth when broadcasting messages (in our experience, around several tens of thousands of mails per hour).
- Potential slowing of the application when broadcasting.
- The application server must be available from the outside (whilst being located in the DMZ, for example) since it hosts the redirection server.

## Installation and configuration steps

### Prerequisites

- JDK,
- Web server (IIS, Apache),
- Access to a database server,
- Bounce mailbox accessible via POP3,
- Creation of two DNS aliases:
  - the first exposed to the public for tracking and pointing to the computer on its public IP;
  - the second alias exposed to internal users for console access and pointing to the same computer.
- Firewall configured to open SMTP (25), DNS (53), HTTP (80), HTTPS (443), SQL (1521 for Oracle, 5432 for PostgreSQL, etc.) ports. For further information, refer to [Network configuration](#).

In the following examples, the parameters of the instance are:

- Name of the instance: **demo**
- DNS mask: **console.campaign.net\*** (only for client console connections and for reports)
- Database: **campaign:demo@dbsrv**

### Installing and configuring (single machine)

Apply the following steps:

1. Follow the installation procedure for the Adobe Campaign server: **nserver** package on Linux or **setup.exe** on Windows.

For more on this, refer to [Prerequisites of Campaign installation in Linux](#) (Linux) and [Prerequisites of Campaign installation in Windows](#) (Windows).

2. Once the Adobe Campaign server is installed, start the application server (web) using the command **nserver web -tomcat** (the Web module enables you to start Tomcat in standalone Web server mode listening on port 8080) and to make sure Tomcat starts correctly:



```
3. 12:08:18 > Application server for Adobe Campaign Classic (7.X YY.R
build XXX@SHA1) of DD/MM/YYYY
4. 12:08:18 > Starting Web server module (pid=28505, tid=-
1225184768)...
5. 12:08:18 > Tomcat started
6. 12:08:18 > Server started
```

Copy

Toggle Text Wrapping

## NOTE

The first time the Web module is executed it creates the **config-default.xml** and **serverConf.xml** files in the **conf** directory under the installation folder. All the parameters available in the **serverConf.xml** are listed in this [section](#).

Press **Ctrl+C** to stop the server.

For more on this, refer to the following sections:

- For Linux: [First start-up of the server](#),
- For Windows: [First start-up of the server](#).

7. Change the **internal** password using the command:

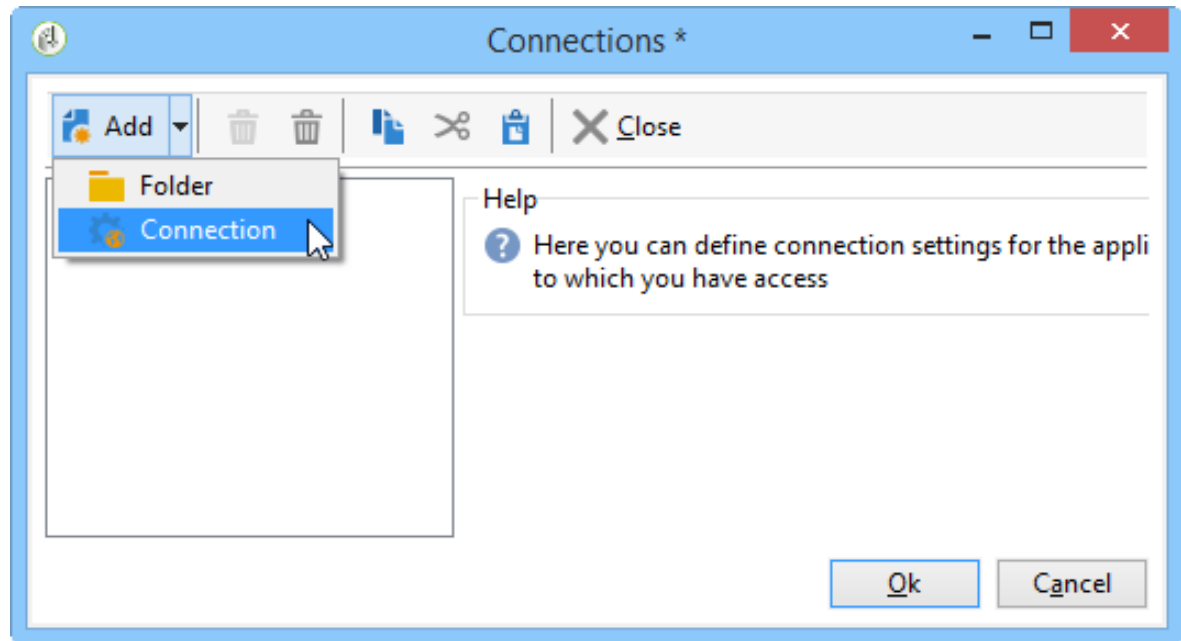
```
8. nlserver config -internalpassword
```

Copy

Toggle Text Wrapping

For more on this, refer to [this section](#).

9. Create the **demo** instance with the DNS masks for tracking (in this case, **tracking.campaign.net**) and access to client consoles (in this case, **console.campaign.net**). There are two ways of doing this:
- Create the instance via the console:



For more on this, refer to [Create an instance and log on.](#)

or

- Create the instance using command lines:

```
nlservice config -  
addinstance:demo/tracking.campaign.net*,console.campaign.net*
```

Copy

Toggle Text Wrapping

For more on this, refer to [Creating an instance.](#)

10. Edit the **config-demo.xml** file (created in the previous step next to **config-default.xml**) and make sure the **mta** (delivery), **wfserver** (workflow), **inMail** (bounce mails) and **stat** (statistics) processes are enabled. Then configure the address of the statistics server:

```
11.<?xml version='1.0'?>  
12.<serverconf>  
13.  <shared>  
14.    <!-- add lang="eng" to datastore to force English for the  
15.     instance -->  
16.    <dataStore hosts="tracking.campaign.net*,console.campaign.net*">  
17.      <mapping logical="*" physical="default"/>  
18.    </dataStore>  </shared>  
19.    <mta autoStart="true" statServerAddress="localhost"/>  
20.    <wfserver autoStart="true"/>  
21.    <inMail autoStart="true"/>
```

```
21. <sms autoStart="false"/>
22. <listProtect autoStart="false"/>
23.</serverconf>
```

Copy

Toggle Text Wrapping

For more on this, refer to [this section](#).

24. Edit the **serverConf.xml** file and specify the delivery domain, then specify the IP (or host) addresses of the DNS servers used by the MTA module to answer MX type DNS queries.

```
25.<dnsConfig localDomain="campaign.com" nameServers="192.0.0.1,
192.0.0.2"/>
```

Copy

Toggle Text Wrapping

#### NOTE

The **nameServers** parameter is only used in Windows.

For more on this, refer to [Campaign server configuration](#).

26. Copy the client console setup program **setup-client-7.XXX.exe** to the **/datakit/nl/eng/jsp** folder. [Learn more](#).
27. Follow the Web server integration procedure (IIS, Apache) described in the following sections:
  - For Linux: [Integration into a Web server for Linux](#)
  - For Windows: [Integration into a Web server for Windows](#)
28. Start the website and test redirection using the URL:  
<https://tracking.campaign.net/r/test>.

The browser must display the following message:

```
<redir status="OK" date="AAAA/MM/JJ HH:MM:SS" build="XXXX"
host="tracking.campaign.net" localhost="localhost"/>
```

Copy

Toggle Text Wrapping

For more on this, refer to the following sections:

- For Linux: [Launching the Web server and testing the configuration](#)
- For Windows: [Launching the Web server and testing the configuration](#)

29. Start the Adobe Campaign server (**net start nlserver6** in Windows, **/etc/init.d/nlserver6 start** in Linux) and run the command **nlserver pdump** once more to check for presence of all enabled modules.

#### NOTE

Starting 20.1, we recommend using the following command instead (for Linux): **systemctl start nlserver**

```
12:09:54 > Application server for Adobe Campaign Classic (7.X YY.R
build XXX@SHA1) of DD/MM/YYYY
syslogd@default (7611) - 9.2 MB
stat@demo (5988) - 1.5 MB
inMail@demo (7830) - 11.9 MB
watchdog (27369) - 3.1 MB
mta@demo (7831) - 15.6 MB
wfserver@demo (7832) - 11.5 MB
web@default (28671) - 40.5 MB
```

Copy

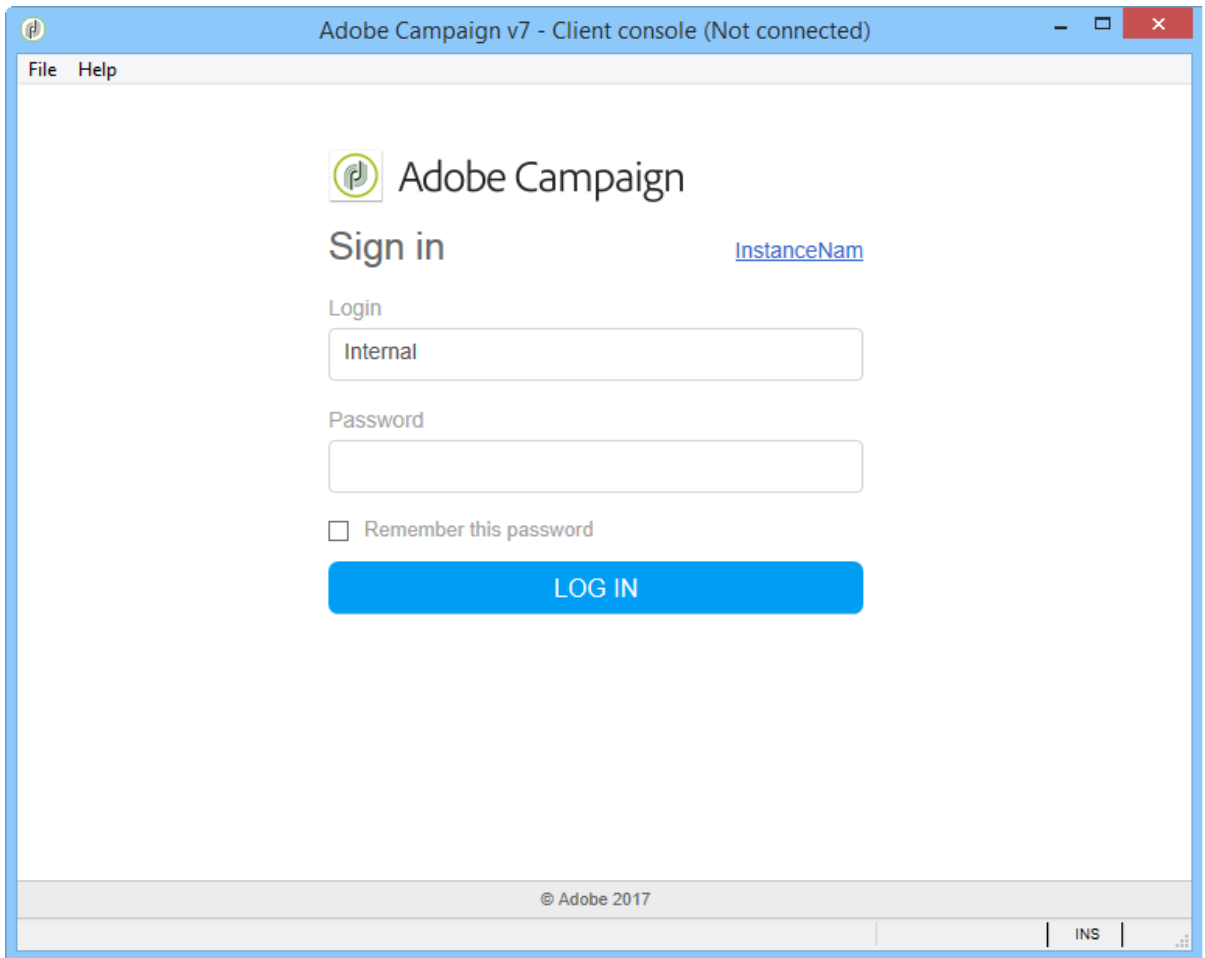
Toggle Text Wrapping

This command also lets you know the version and build number of the Adobe Campaign server installed on the computer.

30. Test the **nlserver web** module using the URL:  
<https://console.campaign.net/nl/jsp/logon.jsp>

This URL enables you to access the download page for the client setup program.

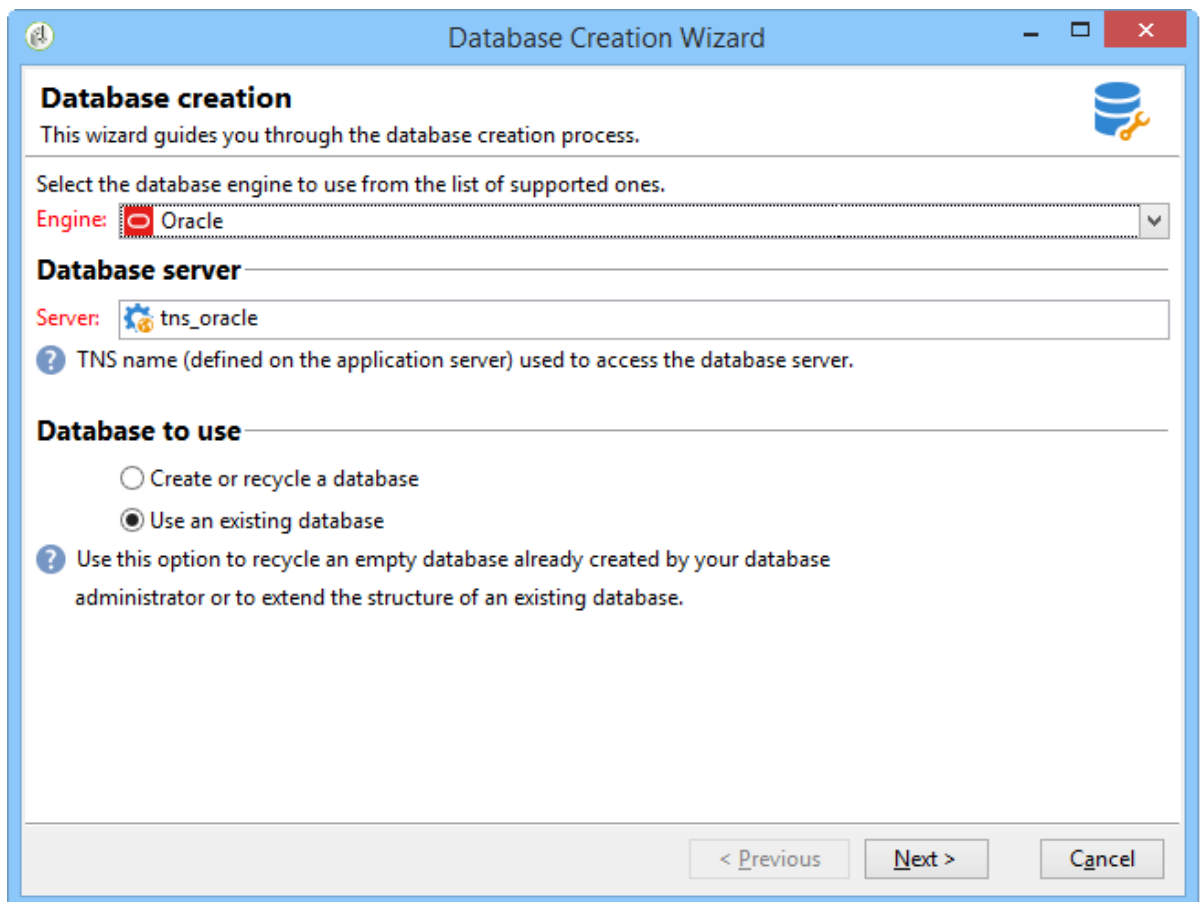
Enter the **internal** login and associated password when you reach the access control page. [Learn more](#).



31. Start the Adobe Campaign client console (from the previous download page or launched directly on the server for a Windows installation), set the server connection URL to <https://console.campaign.net> and connect using the **internal** login.

Refer to [this page](#) and [this section](#).

The database creation wizard appears when you log in for the first time:



Follow the steps in the wizard and create the database associated with the connection instance.

For more on this, refer to [Creating and configuring the database](#).

Once the database is created, log off.

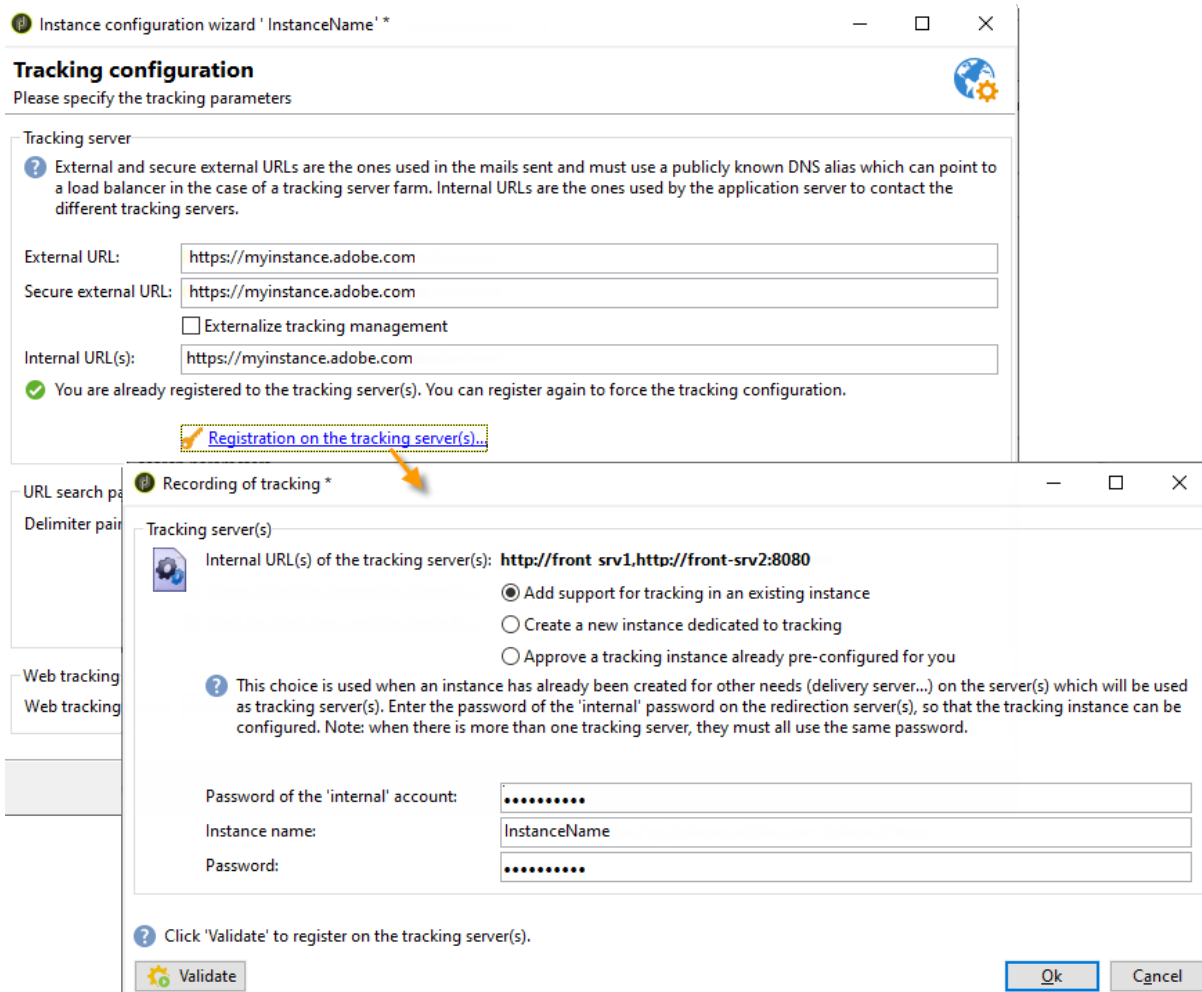
32. Log back on to the client console using the **admin** login without a password and start the deployment wizard ( **Tools > Advanced** menu) to finish configuring the instance.

For more on this, refer to [Deploying an instance](#).

The main parameters to set are the following:

- Email delivery: sender and reply addresses and the error mailbox for bounce mail.
- Tracking: Populate the external URL used for redirection and the internal URL, click **Registration on the tracking server(s)** and then validate it on the **demo** instance of the tracking server.

For more on this, refer to [Tracking configuration](#).



As the Adobe Campaign server is used both as the application server and the redirection server, the internal URL used to collect tracking logs and transfer URLs is a direct internal connection to Tomcat (https://localhost:8080).

- Bounce management: Enter the parameters to handle bounce mail (do not take the **Unprocessed bounce mails** section into account).
- Access from: Provide the two URLs for reports, Web forms and mirror pages.

