

BIND DNS Zone Upload and Import

BIND DNS Zone Import

The [BIND DNS Zone Upload / Import](#) tool uses the named.conf configuration file tied to the zones you are uploading, a .zip or .tar file of the zones themselves, and an optional .csv file mapping zones to customers. The following steps are used when importing BIND DNS zones:

- [Preparing your DNS Zones for Import](#)
- [Importing your DNS Zones \(BIND\)](#)
 - [Video Walkthrough](#)
 - [Step 1: Create a new DNS Import Job](#)
 - [Step 2: Map Data Columns \(Optional\)](#)
 - [Step 3: Reviewing Data](#)
 - [Step 4: Execute Import](#)

Preparing your DNS Zones for Import

If your zone data is currently in BIND format - this is very straightforward.

There are three components for the upload process:

1) The named.conf configuration file tied to the zones you are uploading (required)

This tells the importer the Zone Name and where the zone file is written. It could be as simple as a multi-line file:

```
Simple DNS Config File

zone "my-zone.com" { type master; file "my-zone.com.zone"; };
zone "my-other-zone.com" { type master; file "my-other-zone.com.zone"; };
zone "my-third-zone.com" { type master; file "my-third-zone.com.zone"; };
```

or could be more complex like this file structure directory:

```
Complex DNS Config File

zone "my-zone.com" { type master; file "/usr/local/zones/my-zone.com.zone"; };
zone "my-other-zone.com" { type master; file "/usr/local/zones/more/my-other-zone.com.zone"; };
zone "my-third-zone.com" { type master; file "/usr/local/zones/more/even more/my-third-zone.com.zone"; };
```

This configuration file can be taken directly from the DNS server, and can be in either ISC BIND or NSD format. The system auto-detects which one is being supplied.

For a sample Simple Config: [conf.conf](#)

2) A ZIP or TAR file of the DNS zones themselves (required)

This is as it sounds - a file archive where we can find the zones and it should match the configuration file uploaded in Step 1.

Zone Order

These zone files can be in any order, or in sub-directories, so long as the configuration file (Step 1) correctly points to them

For a sample simple ZIP: [zones.zip](#)

3) Match CSV for assigning DNS Zones to Resources (optional)

This file allows the administrator to "assign" zone files to a given Resource. If you have Imported a group of Resources, they have Resource IDs associated with them. You can then import DNS zones and assign them to those Resource IDs. When complete, you will be able to pull up the Resource Record and see the DNS Zones associated to that Resource ID.

Sample CSV File

```
my-zone.com,test-01,fun stuff, 174.23.14.4, 174.23.14.9
my-otherzone.com,test-02,great stuff, dns1.dns.net, dns2.dns.net
even-reverse-zones.arpa,test-03,amazing stuff
```

Note the columns are the "Zone Name", the "Resource ID", "Notes", "Master Server", "Slave Server"

Importing DNS Server Linkages

When importing zones, you can use the "Master Server" and "Slave Server" columns to assign zones to specified DNS Servers. Please note that the IP address or FQDN of the DNS Server is supported in this field.

To successfully map to a DNS server, that server must already exist within Provision.

For a sample CSV: [config.csv](#)

Importing your DNS Zones (BIND)

Video Walkthrough

Note: Video walkthrough is from ProVision version 5.1.x - however, import steps and descriptions still apply to current version. Updated videos are in progress!

Step 1: Create a new DNS Import Job

Navigate to the [Data Import Tab](#) from the [Admin](#) button to import your data. Select "BIND Zone Upload/Import" under "DNS Import".

Import Tools

Information on preparing data for import and using each import tool is available at <https://docs.6connect.com/display/DOC/Importing+Your+Data>.
Sample import templates are available [here](#).

Resource Import:	DNS Import:
Simple Upload/Import from CSV	BIND Zone Upload/Import
Resource Import Tool "Beta"	PowerDNS Zone Import
IP Import:	InfoBlox Zone Import
Upload/Import from CSV	NS One Zone Import
Import from RIR	Dyn DNS Zone Import
Peering Import	DNSMadeEasy Zone Import
Import BGP Sessions	Route53 Zone Import
	IPPlan Zone Import

Under the "New Import" section, select the [DNS Group](#) under which you want the zones to be imported, then create a Job Name and Description for the import. This is especially useful to keep track of progress in cases the data arrives from multiple sources, or will require multiple stages of manual review.

Select the appropriate **Configuration File** (required if the .zip is not a flat archive file), **Archive File** (required), and **CSV File** (optional) that you prepared above by selecting the "Choose File" button(s) under each section, and browsing to the correct file location. Then hit "Start Import".

DNS Zone Import

Clear

Step 1: New Import

The DNS Import accepts an archive file of zones (ZIP or TAR) in both flat and hierarchical formats. You may also submit a CSV file mapping zone names to customer ids and DNS servers. Please make sure the archive file has an appropriate file extension, and that all files are encoded in UTF-8.

DNS Group:

Example Group

Job Name:

Sample Zone Import

Description:

Example

Configuration File:

Choose File

No file chosen

Optional: a configuration file in BIND or NSD format. If it is not provided the ZIP or TAR must be a flat archive from zone files.

Archive File:

Choose File

No file chosen

Required: a ZIP or TAR of your zones.

CSV File:

Choose File

No file chosen

Optional: a CSV file mapping zones to customers and DNS Servers.

Start Import

Working with Large or Multiple Data Sets

Although you cannot add new files to an existing job, for jobs with multiple sources for data (which may have different formatting), you can simply create separate jobs and descriptions for each source - no need to manually combine the data into one file before importing. The Import tool's mapping and editing functions will allow for the data to be reconciled in ProVision.

For large data sets where multiple stages of manual review might be needed, you can create a new job using the same set of data files in order to work in parallel on a different portion of the data.

After importing, the new job will appear under the "Existing Jobs" section. To continue working with this job, select it from the list and the next step will appear on the page.

Data Import

Job List

Existing Jobs

Sample Zone Import last modified 10-01-2022 9:08 AM

Delete

Step 2: Map Data Columns (Optional)

If you chose to load an optional match CSV file to assign DNS Zones to Resource, a mapping step will be available. Otherwise, proceed to Step 3: Reviewing Data.

For DNS imports, two column definitions are available: **Zone** and **Resource Holder ID**. Using the dropdown menu, select the appropriate column for the zone name and Resource Holder (customer) associated with zone. The Resource Holder ID field will be used to set the Parent Resource of the individual zone once imported.

Server associations will not be imported. Instead, the associated server will be the server attached to the selected DNS Group.

Other columns which do not apply under the available definitions should be left as blank, and will be skipped during the upload process.

When completed, hit "Next".

Step 2: Define Columns

The Import process requires you to enumerate the function of the columns in the provided CSV.

Column Type	Column Type	Column Type	Column Type	Column Type
Zone	Resource Holder ID			
Zone Name	Resource Id	Notes	Master Server	Slave Server
testadd.	test-01	fun stuff	208.39.106.184	
4oakclonelImport.	test-02	great stuff	208.39.106.99	208.39.106.184
example.com.	test-03	amazing stuff	208.39.106.184	208.39.106.82

Next

Step 3: Reviewing Data

After supplying the file set and defining columns (if applicable), a review step is provided. The configuration file is broken into individual jobs, scanned for errors, and shown by row (in batches of 100) to be reviewed. Zones with errors will show as color coded. and can be filtered to be viewed by All, Valid, Warnings, Invalid, or Ignored. From here, the zone can be edited or ignored.

Step 3: Review Data

Please review the data for correctness. Invalid and ignored rows will be skipped.

View: AllValidWarningsInvalidIgnored

Hide

Zone: testadd.	Resource Holder: Example Group	EditIgnore
Zone: 4oakclonelImport.	Resource Holder: Example Group	EditIgnore
Zone: example.com.	Resource Holder: Example Group	EditIgnore

Step 4: Import Data

When you have reviewed the data import job for accuracy, hit the Execute Import button. All rows which are disabled, invalid, have warnings, or were previously successful will be passed over. Successful import rows will be marked as such.

Execute Import

To edit the zone name or select a different Resource Holder, hit "Edit" on the zone record.

After editing, hit "Save", and continue reviewing / editing data as desired.

Step 3: Review Data

Please review the data for correctness. Invalid and ignored rows will be skipped.

View: AllValidWarningsInvalidIgnored

Hide

Zone Name: testadd.

Resource Holder: QA Test Resource

ViewSave

DNS Servers:

Enabled	Server Name	Master	Slave
☐	BIND 6connect QA Server (bind-test.6connect.com)	☐	☐
☐	Test 2 (1.1.1.2)	☐	☐
☐	testaddapproval (1.1.1.1)	☐	☐
☐	Somefakeserver (someserver.com.)	☐	☐
☐	R3-Approvals Test Server (test.com)	☐	☐
☐	Example Server (ns1.dns.example.com)	☐	☐

Zone: 4oakclonelImport.	Resource Holder: Example Group	EditIgnore
Zone: example.com.	Resource Holder: Example Group	EditIgnore

Step 4: Import Data

When you have reviewed the data import job for accuracy, hit the Execute Import button. All rows which are disabled, invalid, have warnings, or were previously successful will be passed over. Successful import rows will be marked as such.

Execute Import

Step 4: Execute Import

When the review step is completed, hit the "Execute Import" button. A progress bar will appear to show progress and note errors if they occur.

Step 3: Review Data

Please review the data for correctness. Invalid and ignored rows will be skipped.

View: All Valid Warnings Invalid Ignored Hide

Zone: testadd.	Resource Holder: Example Group	Edit	Ignore
Zone: 4oakclonelImport.	Resource Holder: Example Group	Edit	Ignore
Zone: example.com.	Resource Holder: Example Group	Edit	Ignore

Step 4: Import Data

When you have reviewed the data import job for accuracy, hit the Execute Import button. All rows which are disabled, invalid, have warnings, or were previously successful will be passed over. Successful import rows will be marked as such.

Execute Import

When the bar reaches 100%, the import is complete.

Step 4: Import Data

When you have reviewed the data import job for accuracy, hit the Execute Import button. All rows which are disabled, invalid, have warnings, or were previously successful will be passed over. Successful import rows will be marked as such.

Execute Import

Current Block: Finished!