

Autopology Installation & Quick Start Guide

Version 1.0

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition.



You can find the most up-to-date technical documentation on the VMware Web site at:
<http://www.vmware.com/support/>

Copyright © 2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA
94304
www.vmware.com

Contents

Autopology Installation & Quick Start Guide	5
1 Overview of Autopology	6
2 Preparing for Installation	7
System Requirements	7
Installation Overview	8
3 Detailed Installation Procedure	9
4 Designing and Deploying Topologies	11
5 Uninstalling Autopology	13
Stopping Autopology Server	13
Uninstalling Autopology	13
6 Troubleshooting	14

Autopology Installation & Quick Start Guide

The *Autopology Installation and Quick Start Guide*, describes how to install the VMware Autopology and use it. The information includes step-by-step configuration instructions, and suggested best practices.

Intended Audience

This information is intended for anyone who wants to install or use VMware Autopology. This information is written for experienced system administrators who are familiar with virtual machine technology and virtual datacenter operations. This manual assumes familiarity with a virtual machine management service—such as VMware vSphere 5.5 or 6.0, including VMware ESX, vCenter Server, and the vSphere Web Client, VMware OVF Tool—or another virtual machine management service with kernel-based virtual machines (KVMs) and VMware NSX.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Overview of Autopology

NSX is a very powerful network virtualization platform that allows the users to spin up large scale complex networks within seconds.

For new users, the inventory-based interface that NSX exposes might be a little overwhelming.

To ease the network creation process and the workloads' attachment to such networks we present 'Autopology' - a graphical interface aimed to complement NSX and simplify workflows for customers and people stepping into the realm of virtual networking.

What Autopology offers:

Drag and drop interface to create logical networks so that the customers can design topologies customized for their needs.

Single click replication of sections of topologies.

Ability to view the drawn network for future reference.

Ability to easily navigate to NSX Manager interface to further customize the logical entities as and when needed.

Ability to attach multiple VMs to logical entities. VMs could be residing on ESX or KVM hypervisors.

Benefits:

One cohesive view to design logical networks as opposed to filling out tedious forms across multiple views.

Intelligent defaults help bypass advanced configuration for network objects.

Rapid scaling of logical networks.

Inventory of workloads across multiple heterogeneous hypervisors.

Combine compute and network controls in a single view.

Autopology requires the server component to be installed on a Linux machine in your datacenter and you can design / deploy the topologies using the Autopology's web interface which connects to the Autopology server.

Preparing for Installation

Before installing Autopology make sure your environment is prepared.

System Requirements

NSX

Autopology currently supports only NSX-T versions. NSX-T downloads can be found at https://my.vmware.com/web/vmware/info/slug/networking_security/vmware_nsx/1_x

Hypervisors

All hypervisors supported by NSX-T are supported by Autopology.

Operating System and Packages [For Autopology Server]

- Autopology works best when installed on Linux OS. Although there are no restrictions on using any Linux version, we recommend using Ubuntu 16.04 [<http://releases.ubuntu.com/16.04/>] to install Autopology.
- The following packages need to be installed on the Linux machine before proceeding with Autopology installation:
 - python-pip
 - python-libvirt
 - libssl-dev
 - python-dev
 - libffi-dev

You can download the above using: ***apt-get install <package name>***

- Autopology integrates with Linux user management system to authenticate the users on the web interface. Before proceeding with install, please ensure you have added users in Linux using the 'useradd' command.

Detailed Instructions about adding a user on Linux can be found in Detailed installation instructions section below.

Browsers [For Autopology Web Interface]

Supported browsers for Autopology Web interface are:

- Chrome (latest version) <https://www.google.com/chrome/browser/desktop/>
- Firefox (version 35 or higher) <https://www.mozilla.org/en-US/firefox/new/>

We don't support Internet Explorer, Edge or Opera at this point.

Installation Overview

Typically, for the initial installation, the order of procedures is as follows:

- 1 Have your NSX-T environment ready by installing NSX Manager, NSX Controller and NSX Edge as per your datacenter requirements.
- 2 Deploy a virtual machine or have a physical machine with Linux installed on it in your datacenter. Ensure that the Linux machine you use for Autopology Server has connectivity to NSX Manager and all the Hypervisors in your datacenter.
- 3 Install the required packages suggested in the System Requirements- -> Operating System and Packages section above on the Linux machine.
- 4 Add users who need to have access to Autopology Web Interface in the Linux user database using 'useradd' and 'passwd' commands.
- 5 Download the Autopology whl package from VMware Flings web page [\[https://labs.vmware.com/flings/autopology-wysiwyg-designer-for-nsx\]](https://labs.vmware.com/flings/autopology-wysiwyg-designer-for-nsx) on your Linux machine.
- 6 Install the Autopology package using the command: pip install <downloaded package name>.whl
- 7 Run Autopology Server on the Linux machine using the command: autopology.server
- 8 Follow the on-screen prompts to install Autopology Server.

Post-Installation

Once Autopology is installed, use one of the supported Web Browsers above and navigate to the URL: <https://<linux machine ip>:<port number selected during install>> to open the Autopology Web Interface to Design and Deploy Topologies.

Refer to Section 3 for detailed installation requirements.

Refer to Section 4 for instructions about designing and deploying topologies.

Detailed Installation Procedure

NSX-T Prerequisites Installation

Prior to installing Autopology, please ensure that you have the following set up in your datacenter:

1. NSX-T version NSX Manager, NSX Controller and NSX Edge (if needed in the topology) installed.
2. ESX / KVM hypervisors setup and
3. NSX Controller is registered with NSX Manager. If you have a cluster of NSX Controllers, all controllers are registered with NSX Manager.
4. ESX and KVM Hypervisors are configured as FABRIC NODES on the NSX Manager.
5. Make a note of the Admin credentials for NSX Manager and root credentials for ESX / KVM hypervisors.

Details about performing steps 1 to 4 above can be found here:

https://www.vmware.com/support/pubs/nsxt_pubs.html

Autopology Server Prerequisites Installation

Prior to installing the Autopology package, please ensure that you have the following:

1. Deploy a physical or a virtual machine and install Linux OS on it. Autopology works best with Ubuntu 16.04 which can be downloaded here: <http://releases.ubuntu.com/16.04/>
2. Ensure that your Linux machine has connectivity to NSX Manager and ESX / KVM hypervisors via the Management Network.
3. Autopology uses Linux users for authentication and you will need the credentials for the user to login to the Autopology Web Client. Before proceeding, create users for the members who will have access to Autopology.

For e.g.: if John Doe from your team needs access to Autopology, add a user in Autopology server

Linux machine

as follows:

```
# useradd johndoe
# passwd test123
```

Repeat the above for all the users whom you want to provide access to Autopology.

4. Install the packages below on the Autopology server Linux machine using apt-get:

```
# apt-get update
# apt-get install python-pip
# apt-get install python-libvirt
# apt-get install libssl-dev
# apt-get install python-dev
# apt-get install libffi-dev
```

If you are logged in as a non-root user, you will need to append 'sudo' to all the commands above.

Note: Ensure that the time on NSX components, hypervisors and Linux machine is in sync.


```

0X21xQD3XPGbXi2dwIDbfsPCS0ile31I3CCo+t3mCzA0knH6U+B0XwRdMM43F+Gk
E0Kn/a081sFpENl+lvquNwnV5Q+LpCQFq9DEJKmLM4e0Ce/4lmCcxXUw38yJmyA
X8TAub9cFGC7XX/845N1dQ5uM8Zem9QRYbHS1+K3LxYcyHIggVJ82t1mc/yvkiGB
BYeqPssyPq5tRW/aCIp+CtIC7zt6XyUFR5XZe0J/kP3N8a+gDmj9TWEbkdfCdJ3
P4wXNcDQkx+f0dNLa71GivpxK+aFzJb3VoN3x84LVWRJzdjUZeQI/63MXQIDAQAB
oyEwHzAdBgNVHQ4EFgQU57ho06QvQbHSHXnmXgM180eA6LEwDQYJKoZIhvcNAQEL
BQADggEBAA6kdF/ox2EHpgNSTzQGgq3jo+4v00ymDXF4w76QvE91Cora1KNedlyM
hXDa5N5TqjVz2WdYdkQz00L1qKB4KHKjeqKmQQ3VVSkX2EP8AEPZj7ka8V1lWYPP
DqDnIcqDiUwQr9YHRSpQ/nJJkdcZXhLNQOAAoERACIMJdj96gub2LhZsyXSHnY5u
LARZWwpAza1uPGW4BwvljzzgSctQwq4pYO/HXoGVHidpvmAUsfqc2T/Po75W4BDC
I6fRaJ950+YAc35wZpeS2za9LtdnVfoxQD/cnkN0e4meyDz7MT6RuJ8B3hopzVi
bjPmXXJrDjaFQPTUw2ENaHOieiuDj0s=
-----END CERTIFICATE-----

```

```

# Trust this certificate ? [Y/N]: y
# Enter filename to store the certificate.[Blank for default
(~/nsx_cert.pem)]: ~/my_nsx_cert.pem

```

7. Provide the 'root' user credentials for ESX and KVM hypervisors.

```

# Enter ESX root password:
# Enter KVM root password:

```

Note: Password will not be displayed on the console as you type.

8. Provide a port number on which Autopology server will listen for HTTP requests. Make sure other applications are not using the same port number.

```

# Enter port number for server to listen: 5000

```

9. If all the credentials you provided are valid and Autopology server is able to connect to NSX Manager and all the ESX / KVM hypervisors which are configured as Fabric Nodes, Autopology server will start up and display a message as below:

```

# Successfully started Autopology Server
# Use the URL: https://<your machine ip>:5000 to access Autopology Web
Interface.

```

10. If autopology server is unable to start-up due to any reason (bad password, connection issues etc.) the following message will be displayed

```

# Failed to start Autopology Server
# Please check /tmp/autopology.log for more information.

```

What to do next

Use any machine which has one of the supported web browser installed and has connectivity to autopology server to access the Autopology Web Interface.

Login to the web interface using Linux credentials and start designing and deploying topologies. Refer to section 4 for more details about designing and deploying topologies using Autopology.

Designing and Deploying Topologies

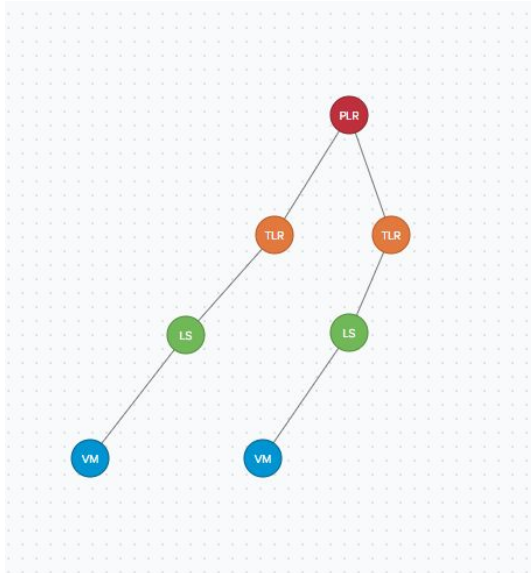
4

To login to the Autopology UI, please use the same credentials used for your Ubuntu machine. Once you are successfully logged in you may start creating your logical topologies.

The UI is a collection of 3 core pieces -

1. Object inventory - lists all the objects that may be added to the canvas to form a topology.
2. Action bar - lists controls like Topology listings, settings and deploy control
3. Canvas - this is where user constructs the logical topology user drag-drop-describe interactions. It also has additional controls on the top-left toolbar to clone/delete/reset the topologies

Following is a step-by-step guide to creating a 3-tier logical topology as seen in the illustration -



1. Click on the PLR (Provider Logical Router) icon in the Object Inventory. This adds a PLR object to the canvas which is the root of your topology.

2. Add a TLR (Tenant Logical Router) by clicking the appropriate icon from the inventory. To link the TLR to PLR, simply right click to open the context menu and click the “Build Link” option. Drop the link on the PLR object to set the connection.
3. You can also use the context menu to configure additional properties on each of the objects.
4. Similarly add a LS (Logical Switch) and link it to the TLR like above. This completes the subtree on the left.
5. To create the subtree on the right, one can use the clone feature on the canvas.
 - a. Simply select the node that you intend to copy (we use the TLR in this example). The selected node is annotated with a dark border around the object.
 - b. Click the clone option and enter the number of copies that are needed.
 - c. Once you hit “GO”, the node along with all its children are copied automatically with all the configurations intact.
6. To add VM (Virtual Machine) to your LS, simply click the VM icon in the Object Inventory and follow earlier steps to link the VM and LS.
7. At this point, the topology is ready for deployment. To deploy you’d need to execute the steps below -
 - a. Click on the deploy button which then prompts you to optionally configure (TN) Transport Nodes for your topology
 - b. You can add multiple ESX/KVM/Edge nodes along with their IP addresses by using the form fields in the “Configure Transport Nodes” Dialog.
 - c. Once you “Save Nodes”, the system prompts you to enter a name for your topology. The name is used to differentiate between different topologies on the UI and is also repurposed as a tag for the resources created on NSX Manager for a particular topology.
 - d. Click “Deploy” to initiate the deployment.
8. The state of the deployment is reflected on the top right corner of the canvas as it goes through different stages.
9. Once successfully deployed, each resource’s realized name may be found on the tooltip as you hover over each object on the canvas.
10. Additionally, one may right click on the object and select “Show in MP” link to open the realized object on the NSX Manager UI.

Uninstalling VMware Autopology

Stopping Autopology Server

You can stop Autopology server running on the Linux machine as below:

1. Use the command below to find the autopology server's process id:

```
# ps -aux | grep autopology
root  14245  5.5  0.6 1292124 52804 pts/1  Sl  18:38  0:00 python
/usr/local/lib/python2.7/dist-packages/autopology/server.py
root  14351  0.0  0.0 14228 1092 pts/1  S+  18:38  0:00 grep --color=auto autopology
```

2. Use the process id shown in the output, 14245 in this example, and stop the server using the command below:

```
# kill 14245
```

Uninstalling Autopology

You can remove elements of VMware Autopology using the command below on the Linux machine which hosts the Autopology Server: `pip uninstall autopology`

What to do next

After removing Autopology, make sure the log files are removed from `/tmp` directory on the Linux machine which hosted Autopology Server to free up space on the hard disk.

Troubleshooting

Use the autopology.log file for troubleshooting the system. Autopology logs are collected in /tmp folder and have prefix autopology for all the files.