



Automation Leveraging NSX REST API

Automate the network via programming languages, vRO, and other tools leveraging NSX REST API

NSBU
UPDATED 6/21/2016

VERSION 1.0

Intended Audience	2
Overview.....	2
Why You Should Care	3
Introduction to Rest API	3
API Specification Formats.....	4
NSX - REST API	5
Using REST API with cURL	5
Using REST API with Postman	7
Example: Create a Universal Distributed Logical Router (UDLR) in NSX 6.2.2	9
Programming Languages with NSX Rest API	12
Why Use Programming Languages with NSX REST API?	12
Basic Concept and Architecture.....	12
Simple Scripts	13
Functions and Modular Programs.....	14
Libraries	15
Python	16
Example: Automating Security Group and Policy Creation.....	18
Perl.....	24
Example: Monitoring VMware NSX SpoofGuard with REST API and Perl	25
Go	32
Example: Creating NSX Services in Go.....	33
Automation Tools with NSX REST API	39
vRealize Orchestrator	39
Example: Create an IP Set in NSX Using vRealize Orchestrator	42
Example: Enable or Disable HA on an NSX Edge Using vRO	51
Ansible	56
Example: Using Ansible Playbook to Manage Logical Switch State	57
Cloud Management Platforms with NSX REST API.....	64
vRealize Automation	66
References	72

Intended Audience

The intended audiences for this paper are devops/automation/system engineers and network engineers looking to automate logical networking and security within their virtualized environment. The goal of this paper is to demonstrate how automation can be accomplished in several ways by leveraging the NSX REST API. This paper explains the fundamental concepts around REST API, programming languages, and automation tools; nevertheless, the reader should already have some general basic understanding of these tools and concepts. The main goal of this document is to enable the reader to easily leverage NSX REST API automation within their environments.

Overview

VMware has helped transform the data center landscape by introducing virtualization technologies for compute, storage, and networking. By replicating traditional hardware constructs in software, VMware allows for not only efficiency in hardware utilization, improved resiliency, lower costs, but also automation at a completely new level that has never been possible before. Gone are the days of hardware specific APIs, screen/CLI scraping and vendor-specific tool sets that would have to be stitched together to create a complete solution.

VMware NSX provides a RESTful API service via NSX Manager that can be consumed in several ways.

The NSX REST API can be consumed directly via a tool/library such as cURL or a REST Client like Postman, via multiple popular programming languages, and via orchestration cloud management tools. Popular programming languages such as Python, PowerShell, Perl, Go, and Java have REST client libraries which can easily be utilized to consume the NSX REST API. This means that elaborate workflows and complete systems/portals can be created to provide custom automation, management, and monitoring capabilities.

Tools such as VMware vRealize Orchestrator (vRO) or configuration frameworks like Ansible can also be used to create advanced workflows for NSX. vRO also provides a HTTP REST API client allowing for NSX REST API to be called directly from the tool; a NSX vRO Plugin is also available.

Further, cloud management systems such as vRealize Automation (vRA) and Openstack incorporate and leverage pre-packaged automation solutions. A vRO plugin for vRA is also available allowing for purpose-built custom automation of NSX objects.

Why You Should Care

By consuming the NSX REST API natively or through an automation tool or cloud management platform, customers can achieve efficiency through automated processes and workflows. There are several important reasons why automation is center-stage in modern data centers:

- Increases productivity by decreasing manual steps for configuration, provisioning, failover/recovery, etc.
- Avoids manual configuration mistakes by automating tasks to follow preset, defined processes
- Allows for an increase in the number of managed systems without a significant opex increase
- Allows the infrastructure to support continuous operations and automated scaling; automatic creation of testing environments can be part of continuous integration workflows
- Repeatable and consistent application environments can be created which minimizes customizations and provides consistency for ease of maintenance and troubleshooting

In a nutshell, by leveraging the NSX automation capabilities, customers can not only increase productivity, but can also avoid increased costs, reduce errors due to manual error-prone processes, simplify processes, and enable ease of troubleshooting and recovery.

Introduction to Rest API

Representational State Transfer (REST) is a software architecture that enables users to query and modify applications states using simple HTTP primitives.

REST defines the use of POST, GET, PUT and DELETE to implement CRUD (create, read, update and delete) operations. The below table describes various HTTP verbs and their corresponding operations; the context used is NSX as the system being utilized.

HTTP “Verbs”	Use	CRUD
POST	Create an NSX object (e.g. logical switch)	Create
GET	Retrieve data about a single NSX object or multiple objects	Read
PUT	Modify the properties of an already existing NSX object	Update
DELETE	Remove an NSX object	Delete

Table 1: HTTP Verbs and CRUD Functions

A system with a northbound REST-based API enables any software that can consume standard-based protocols and formats such as HTTP, JSON, and XML to create clients that are able to interact with its objects and data structures. Simple tools such as **wget**, **cURL**, or even a web browser could be used to interact with an application that supports REST APIs. Programming and scripting languages that support HTTP operations can also be used to interact with REST API based applications.

API Specification Formats

While REST APIs are extremely useful to allow different software components to communicate, they can become complex to understand and to document. In the past, when developers had to write their implementation of an API (for either a client or a server) they had to analyze the documentation and manually build all the required interfaces. API documentation is also often presented in documents that do not follow any specific guideline for formatting and definition.

To solve these issues, several specification formats emerged in order to provide a consistent format to document and consume APIs.

API specification formats provide a number of benefits for developers, including:

- Non-ambiguous ways to define and document APIs
- Dramatic reduction of errors during implementation, testing and troubleshooting
- Diffusion of tools that can automatically generate API implementations in different programming languages using the specification as source

Two of the most common API documentation and specification formats, at the time of writing this document, are RAML and OpenAPI Specification (formerly known as Swagger). Both are widely utilized and have a large open source community. There are also tools that automate the conversion across the different API documentation and specification formats.

RAML is a language based on YAML; YAML is a syntax used to describe an API in a 'human friendly' way. YAML can easily be mapped to data types common to most high-level languages. RAML syntax can be parsed automatically to produce documentation, dynamic clients in different languages such as Python, and collections to be used with tools like Postman.

As an example, a company can automatically generate, from the application's source code, a RAML specification file for a product. This RAML file can then be used by the same company to generate the Java source code to implement the server-side API for the product which can then be shared with third parties. A partner could use the same RAML file to generate Python source code to implement the client-side API. It's somewhat similar to what Web Service Definition Language (WSDL) enables for Simple Object Access Protocol (SOAP) based services: automatically generate client code in multiple

languages for calling the web service specified or server stubs for implementing the service(s) instead. Consistency is a major benefit here: since all the implementations in different languages adhere to the same specification, there is less space for ambiguity.



Additional information on RAML and OpenAPI Specification can be found on the respective project web sites ([RAML](#) and [OpenAPI Specification](#)).

VMware offers a community supported RAML specification of NSX for vSphere API (nsxraml) that can be used to simplify the consumption of NSX services. Some of the examples and tools mentioned in this document use nsxraml. For additional details please check the [nsxraml page on GitHub](#).

NSX - REST API

NSX provides a REST API interface via NSX Manager. Any REST Client can utilize this REST API interface provided by NSX Manager.

The next image shows the typical structure of a REST API call for NSX – using an example that shows scopes.




Figure 1: NSX API Structure (example)

For security reasons, NSX enforces authentication before querying or changing the state using REST APIs.

Using REST API with cURL

cURL is a simple tool used to transfer data from server to client. cURL supports various methods, one being HTTP. While cURL may be used to manage NSX manager, it's not a very scalable approach unless cURL is being used from within another programming language etc., to help automate. As such, the following examples with cURL are primarily to help with the understanding how REST APIs work.



If you are trying this out on a Mac, use **xmllint** to parse the output

The below example uses cURL to call the REST API described in Figure 1 above. Note the usage of various cURL options: **-k** to ignore any SSL related warnings; **-u** to specify username and password.

```
skommu-mac:~ skommu$ curl -k -u admin:password https://nsxmgr-01.corp.local/api/2.0/vdn/scopes | xmllint --format -
```

Following is the output for the above cURL request:

```
1. <?xml version="1.0" encoding="UTF-8"?>
2. <vdnScopes>
3.   <vdnScope>
4.     <objectId>vdnscope-9</objectId>
5.     <objectTypeName>VdnScope</objectTypeName>
6.     <vsmUuid>421DE20E-63C4-914C-785F-C760AB465B3F</vsmUuid>
7.     <nodeId>421DE20E-63C4-914C-785F-C760AB465B3F</nodeId>
8.     <revision>1</revision>
9.     <type>
10.      <typeName>VdnScope</typeName>
11.    </type>
12.    <name>Global NSX Zone</name>
13.    <description/>
14.    <clientHandle/>
15.    <extendedAttributes/>
16.    <isUniversal>false</isUniversal>
17.    <universalRevision>0</universalRevision>
18.    <id>vdnscope-9</id>
19.    <clusters>
20.      <cluster>
21.        <cluster>
22.          <objectId>domain-c332</objectId>
23.          <objectTypeName>ClusterComputeResource</objectTypeName>
24.          <vsmUuid>421DE20E-63C4-914C-785F-C760AB465B3F</vsmUuid>
25.          <nodeId>421DE20E-63C4-914C-785F-C760AB465B3F</nodeId>
26.          <revision>230</revision>
27.          <type>
28.            <typeName>ClusterComputeResource</typeName>
29.          </type>
30.          <name>Compute-POD-02</name>
31.          <scope>
32.            <id>datacenter-2</id>
33.            <objectTypeName>Datacenter</objectTypeName>
34.            <name>DC</name>
35.          </scope>
36.          <clientHandle/>
37.          <extendedAttributes/>
38.          <isUniversal>false</isUniversal>
39.          <universalRevision>0</universalRevision>
40.        </cluster>
41.      </cluster>
42.      <cluster>
43.        <cluster>
44.          <objectId>domain-c219</objectId>
45.          <objectTypeName>ClusterComputeResource</objectTypeName>
46.          <vsmUuid>421DE20E-63C4-914C-785F-C760AB465B3F</vsmUuid>
47.          <nodeId>421DE20E-63C4-914C-785F-C760AB465B3F</nodeId>
```

```

48.      <revision>243</revision>
49.      <type>
50.        <typeName>ClusterComputeResource</typeName>
51.      </type>
52.      <name>Compute-POD-01</name>
53.      <scope>
54.        <id>datacenter-2</id>
55.        <objectTypeName>Datacenter</objectTypeName>
56.        <name>DC</name>
57.      </scope>
58.      <clientHandle/>
59.      <extendedAttributes/>
60.      <isUniversal>false</isUniversal>
61.      <universalRevision>0</universalRevision>
62.    </cluster>
63.  </clusters>
64. </clusters>
65. <virtualWireCount>498</virtualWireCount>
66. <controlPlaneMode>UNICAST_MODE</controlPlaneMode>
67. </vdmScope>
68. </vdmScopes>

```

Figure 2: Result of REST API Call

Using REST API with Postman

Postman is a REST Client tool to execute REST APIs from the Chrome web browser. There are many different REST Client tools available for different browsers. RESTClient on the Firefox browser is another specific example. This example shows how to use Postman with the Chrome browser to call the same REST API used above.

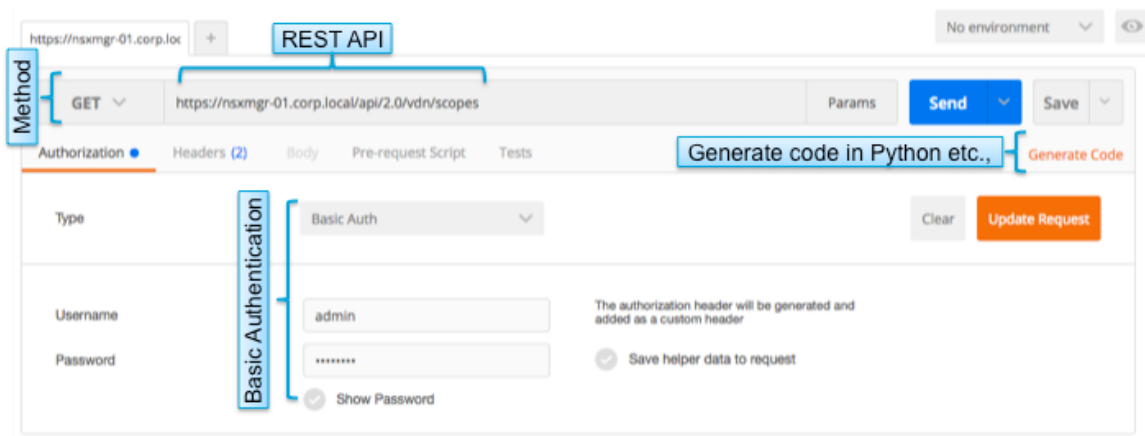


Figure 3: Postman GUI

The below figure shows the output excerpt from Postman, which can be seen is the same result obtained when using cURL.

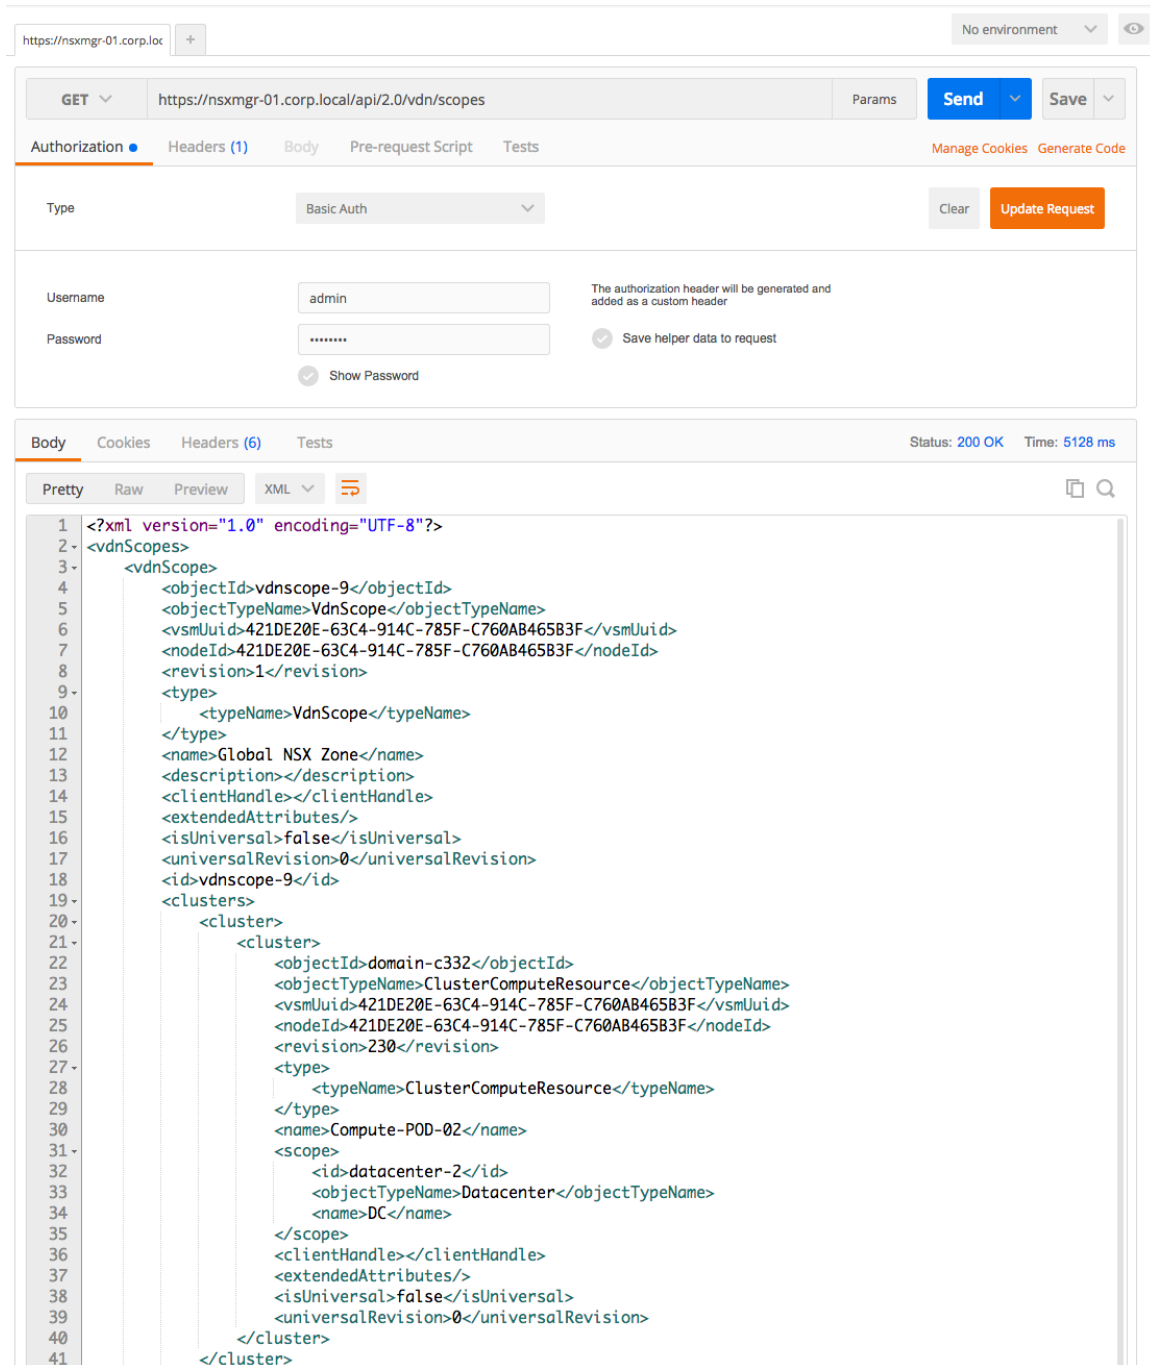


Figure 4: Data Returned via the REST API Call From Postman

Postman is very useful for quickly and automatically generating code snippets for a given API call. To use this feature, one can click on the **Generate Code** option in the upper right corner of the Postman application and select the desired programming language for which the code should be generated for. This auto-generated code will work without any changes; however, if SSL related issues are encountered, the **-k** option can be added to the respective cURL call.



Figure 5: Use Postman to generate cURL code snippet

Below is another example to generate Python code for the same example call using Postman. The procedure is similar to getting code for cURL; just select Python instead of cURL.

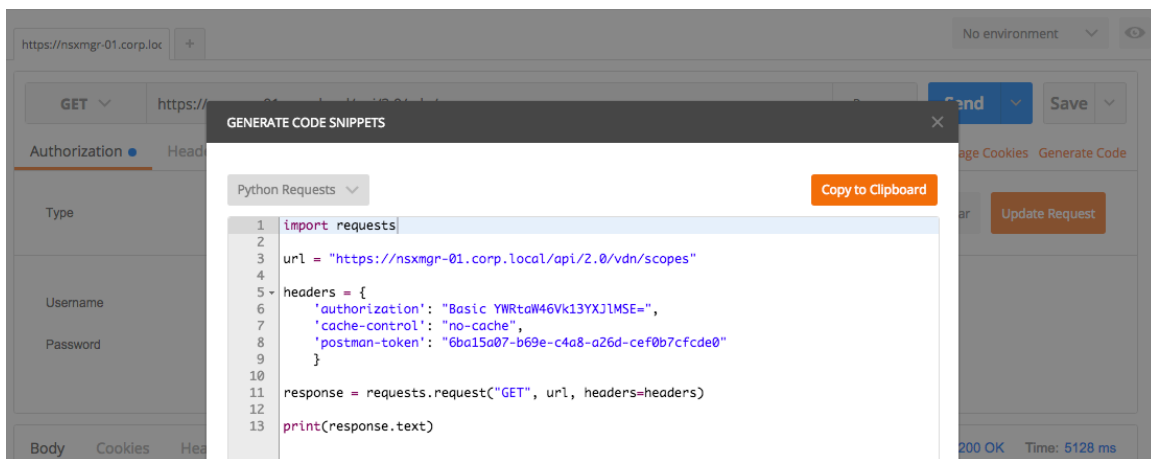


Figure 6: Use Postman to generate Python code snippet



If SSL related issues are encountered with the generated Python code, pass the variable **Verify=False** to the **requests.request** method:

```
response = requests.request("GET", url, headers=headers, verify=False)
```

Example: create a Universal Distributed Logical Router (UDLR) in NSX 6.2.2

This example demonstrates using the RESTClient plugin on the Firefox browser to create a Universal Distributed Logical Router (UDLR) in NSX 6.2.2. The successful creation of the UDLR named **Universal DLR Test** is confirmed from the **Status Code** returned by the NSX Manager. The related XML data structure within the **Body** field shows all the available fields with their respective values for the desired NSX object. If an error is

received when submitting the NSX REST API call, the returned **Status Code** and respective error message will be helpful to understand and troubleshoot the issue.

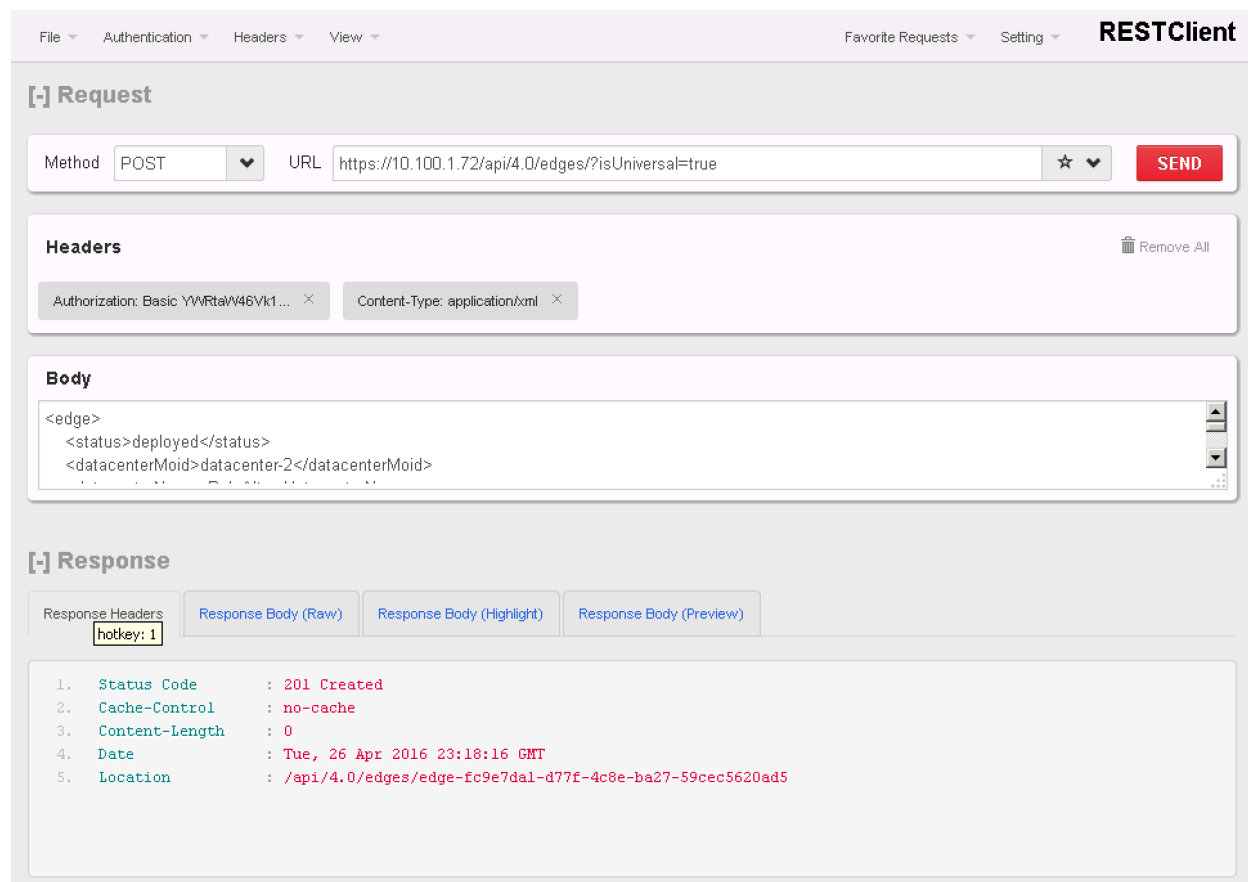


Figure 7: Using RESTClient on Firefox to Create a UDLR in NSX 6.2.2

Below is the body of the REST API call shown in Figure 7 used to create a UDLR. This is entered within the **Body** field of the RESTClient in Figure 7. Note, parameters can be adjusted based on requirements.

```

1. <edge>
2.   <status>deployed</status>
3.   <datacenterMoid>datacenter-2</datacenterMoid>
4.   <datacenterName>PaloAlto</datacenterName>
5.   <tenant>default</tenant>
6.   <name>Universal DLR Test</name>
7.   <fqdn>NSX-edge-Test</fqdn>
8.   <enableAesni>>false</enableAesni>
9.   <enableFips>>false</enableFips>
10.  <vseLogLevel>info</vseLogLevel>
11.  <appliances>
12.    <applianceSize>compact</applianceSize>
13.    <appliance>
14.      <highAvailabilityIndex>0</highAvailabilityIndex>
15.      <vcUuid>5013a4d4-9fa5-d86c-2166-3df1636e7d5a</vcUuid>

```

```

16.      <resourcePoolId>domain-c22</resourcePoolId>
17.      <resourcePoolName>Edge Cluster</resourcePoolName>
18.      <datastoreId>datastore-38</datastoreId>
19.      <datastoreName>EMC_VNX_1-1</datastoreName>
20.      <hostId>host-32</hostId>
21.      <hostName>10.100.1.72</hostName>
22.      <vmFolderId>group-v3</vmFolderId>
23.      <vmFolderName>vm</vmFolderName>
24.      <vmHostname>NSX-edge-Test</vmHostname>
25.      <vmName>Universal DLR</vmName>
26.      <deployed>true</deployed>
27.    </appliance>
28.
29.    <deployAppliances>true</deployAppliances>
30.  </appliances>
31.  <cliSettings>
32.    <remoteAccess>true</remoteAccess>
33.    <userName>admin</userName>
34.    <password>VMwareVMware1!</password>
35.    <passwordExpiry>99999</passwordExpiry>
36.  </cliSettings>
37.
38.  <autoConfiguration>
39.    <enabled>true</enabled>
40.    <rulePriority>high</rulePriority>
41.  </autoConfiguration>
42.  <type>distributedRouter</type>
43.  <isUniversal>true</isUniversal>
44.  <universalVersion>21</universalVersion>
45.  <mgmtInterface>
46.    <label>vNic_0</label>
47.    <name>mgmtInterface</name>
48.    <addressGroups/>
49.    <mtu>1500</mtu>
50.    <index>0</index>
51.    <connectedToId>universalwire-11</connectedToId>
52.    <connectedToName>Universal HA</connectedToName>
53.  </mgmtInterface>
54.
55.  <localEgressEnabled>false</localEgressEnabled>
56. </edge>

```

Figure 8: XML Data Structure Returned After RESTClnet API Call

Once the **Send** button is clicked, in NSX, the new UDLR can be observed as shown below in Figure 9.

Id	Name	Type		
edge-1	ESG1-1	NSX Edge		
edge-2167d8fa-915d-48b7-a7c1-620906...	Universal DLR	Universal Distributed Router		
edge-fc9e7da1-d77f-4c8e-ba27-59cec56...	Universal DLR Test	Universal Distributed Router		

Figure 9: Verification on the UDLR Creation in the NSX Manager GUI

Programming Languages with NSX Rest API

This section demonstrates how the NSX Rest API can be leveraged via different programming languages to further automate tasks and create workflows/programs with advanced logic.

Why Use Programming Languages with NSX REST API?

Leveraging programming languages with the NSX REST API goes a step beyond simply calling NSX REST API calls via REST Client tools. Programming languages such as Python or Perl allow for introducing custom advanced logic and workflows.

Using conditional constructs available in all scripting/programming languages such as if/then statements and control flow logic such as looping mechanisms and boolean constructs, advanced programs and custom solutions can be created.

Further, an organization may already have a custom portal for managing/monitoring their infrastructure. By leveraging the NSX REST API with respective programming languages, organizations can build NSX support into their custom management portals or even create a new custom cloud management platform.

Below, a basic concept and architecture discussion of how programming languages can be utilized with NSX REST API is provided. In later sections, the programming languages Python, Perl, and Go are briefly discussed and examples provided.

Basic Concept and Architecture

As mentioned prior, NSX provides a REST API interface via NSX Manager, and any REST Client can utilize this REST API interface provided by NSX Manager. All popular programming languages have REST Client libraries/methods that can be utilized for REST operations discussed prior such as POST, GET, PUT, and DELETE. The respective REST Client library is typically imported into a program allowing accessibility to its REST Client methods. The below table provides examples of popular REST Client libraries for several common programming languages.

Programming Language	REST Client Library
Python	“requests” is the most popular library and one of the simplest to consume; some other libraries that can also be utilized include “siesta”, “httplib”, “httplib2”, “urllib”
PowerShell	“Invoke-RestMethod” cmdlet
Perl	“REST:Client” is the most popular library; “LWP::UserAgent” can also be used
Java	Several libraries that can be leveraged include: “Unirest”, “Javalite HTTP”, “Apache HttpClient”, “Jersey”
Go	“net/http” is the most popular library; some other libraries that can also be utilized include “sling”, “resty”, “napping”, “grequests”

Table 2: REST API Client Libraries for Common Programming Languages

Simple Scripts

Many simple utility scripts composed of one file may be written as shown in Figure 10 below. The script is short and steps are sequential in nature within the file. This may be adequate for such tasks as returning the NSX Distributed Firewall (DFW) to default state. This script serves its purpose for resetting the DFW to default state in-case one accidentally blocks vCenter communication and can no longer access the GUI to make the change.

In some programming languages this would consist of less than ten lines of code. A generalized example layout below shows what this may look like using any programming language

1. Simple Generalized Program
- 2.
- 3.
4. Program Variables here such as NSX Manager URL, Username, and Password
- 5.
- 6.
7. Main body of script calling NSX REST API to delete DFW section and reset to default

Figure 10: Basic Script/Program Structure

However, as programs get more advanced or additional functionality is added, it soon becomes clear splitting the program up into compartments or smaller functional units becomes helpful not only for management but also for rapid development, troubleshooting, and code reuse. In such cases the desire is to make the programs more modular. The next section discusses how such programs can be made more modular.

Functions and Modular Programs

Assume additional functionality needs to be added to the prior example of returning the NSX Distributed Firewall (DFW) to default state. New requirements now include the following:

- display a warning message that the NSX DFW will be reset to a default state
- user must acknowledge to proceed by clicking **y**
- some new DFW rules must be added automatically to represent custom default state

From the above requirements we can compartmentalize these into three tasks:

1. Display a warning message and confirm acknowledgement
2. Delete DFW rules and put DFW into default state
3. Create new DFW rules to represent custom default state

Each of these identified tasks can be a separate function. A function in programming languages is a specific reusable section of a code that can be called at any time to perform a specific task. If needed, functions can also be written to accept arguments that are used by the logic within the function to achieve needed results. Additionally, functions can also be written to return values once the respective function is complete; a value can consist of perhaps a value representing status or some other specific criteria.

So, instead of writing all logic sequentially within a file, specific functions can be created as separate blocks of code and be called from within the program. These functions can be within the same file or even in a separate file and imported in. Such a workflow would look like the below.

```

1.    Simple Generalized Program
2.
3.
4.    displayWarning() - Function for displaying warning message
5.
6.
7.    deleteDfwSection() - Function for deleting DFW rules (putting DFW into default state)
8.
9.
10.   addCustomDfwRules() - Function for adding custom default DFW rules
11.
12.
13.   Main body of script calling respective functions above in sequence
14.   displayWarning()
15.   deleteDfwSection()
16.   addCustomDfwRules()

```

Figure 11: Basic Script/Program Structure Using Functions

To further organize code, enable code reuse, and allow for faster deployment of applications custom libraries of code can be created. These libraries typically group together specific features/functionality that are related such as file manipulation or Internet connectivity. The next section discusses libraries in more detail.

Libraries

Programming languages will include a standard library out of the box for common tasks such as manipulating files, accessing operating system services, Internet connectivity, etc. Many additional external libraries exist, and custom libraries can be created as needed. Continuing with the example displayed above, below all custom functions prior created for modifying DFW are placed into an external custom library called **CustomDFW** which is then imported into the program to be consumed. **CustomDFW** represents the custom library created providing respective DFW functions.

```

1.    Simple Generalized Program
2.
3.
4.    import CustomDFW - Custom DFW Library
5.
6.
7.    displayWarning() - Function for displaying warning message
8.
9.
10.   deleteDfwSection() - Function for deleting DFW rules (putting DFW into default state)
11.
12.
13.   addCustomDfwRules() - Function for adding custom default DFW rules

```

Figure 12: Basic Script/Program Library Structure

The below generalized program in Figure 13 imports the **CustomDFW** library and calls respective functions. It can be seen how creating libraries enables a clean organized structure which allows for ease of code development, code maintenance, and code reuse.

```
1.      Simple Generalized Program
2.
3.
4.      import CustomDFW
5.
6.
7.      Main body of script calling respective functions from imported library
8.      displayWarning()
9.      deleteDfwSection()
10.     addCustomDfwRules()
```

Figure 13: Basic Script/Program Using Libraries

In the next few sections we'll take a look at some specific programming languages and respective examples.

Python

Python is a popular open-source, high-level programming language. Its popularity is in part due to its ease of use, readability, low learning curve, rapid deployment nature, cross-platform capabilities, and library support. As such, it has become one of the most popular programming languages and a favorite among network admins and devops engineers.

A focus of Python is both rapid and ease of development. For example, to allow for rapid and ease of development, Python is dynamically typed meaning that Python allows the user to assign different object types to the same variable without having to declare that variable of a certain type first. This allows quickly creating data constructs such as lists containing object of different types. Python is also very readable and easy to troubleshoot because it does not require semicolons to terminate statements or use braces to designate blocks of code like other programming languages. Instead indentation is used to designate specific blocks of code.

Figure 14 below demonstrates how indentation is used with Python to structure blocks of code.

```
1.      First line of Python code
2.
3.      Loop 1
4.          Loop line 1
```

```
5.         Loop line 2
6.
7.     Fifth line of Python code
8.
9.     If this is true
10.         Do this
11.     Else
12.         Do this
13.
14.     Last line of Python code
```

Figure 14: Python Code Structure

Additionally, Python is a very expressive language in the sense that the same task can usually be done in Python in less lines of code than other languages. Python also supports object oriented features such as user-defined classes and inheritance where objects can inherit from other parent objects.

Further, Python's a cross-platform language able to run on multiple platforms such as MAC, Unix/Linux, and Windows. Python is an interpreted language, and, as such, can run on any platform that has a Python interpreter; this allows for its cross-platform capabilities.

Python also has a vast assortment of natively included and externally importable libraries; this enables rapid development of solutions as it allows programmers to reuse efficiently developed code within their custom projects. It's important to note that when using the popular REST Client **requests** library with Python as shown in the examples further below, the **requests** library must be manually installed as it is not part of the Python standard library. Example code below is written in Python 2.7.7 and the Python package management system, **pip**, was used to install the **requests** package from the cli with the following command: **pip install requests**.

Another important detail to note about Python is that there are two lines of versioning available: Python 2.x and Python 3.x. Python 2.x is the traditional Python that simply builds upon earlier versions of Python without breaking backward compatibility. Python 3.x is the first version of Python to break backward compatibility with older Python versions; the reason for this was simply to improve some of the traditional Python language syntax and details for consistency and to evolve the language in a way to make it even more easy to understand, use, and build upon.

Whether to use Python 2.x or Python 3.x is a choice based on preference. NSX REST API can be used with either one. However, as Python 2.x has been around longer, there are likely Python 2.x libraries that have not been converted over to Python 3.x. If one of these libraries are needed for respective development purposes or the program needs to be merged into a larger existing Python 2.x program, Python 2.x can be used. Most popular Python libraries have been converted over to Python 3.x and as such Python 3.x would be a good choice for new projects.

To learn more about Python, see the [official Python website](#).

Example: Automating Security Group and Policy Creation

Imagine creating a security posture and having to identify

1. The appropriate workloads
2. What security policies need to be applied to respective identified workloads

With NSX, a user can create what is known as **Security Groups**. Security Groups can easily be defined via the NSX Service Composer or via the Grouping Objects section (**NSX Managers** → **[Select NSX Manager]** → **Manage** → **Grouping Objects** → **Security Groups**) and allows for a unique model in consuming network and security services.

Objects such as virtual machines (VMs) across which similar policies must be applied can be defined inside a Security Group, and a consistent policy can then easily be mapped across all the respective objects. Furthermore, objects can be dynamically included within Security Groups and the security policies leveraging these Security Groups can then be configured based on the built-in VMware NSX features or by those offered by 3rd party solutions.

Assume the tasks at hand are:

1. Identify workload traffic and create security groups to dynamically include the relevant workloads. This will help determine what security policies should be applied to the specific workloads based on requirements
2. Create security policies leveraging created Security Groups

With NSX REST API, we can automate all of these steps and avoid a manual and time intensive process. In this example, the automatic creation of security groups with dynamic matching criteria will be demonstrated. Figure 15 below displays the NSX lab setup for this example.




Figure 15: Example NSX Lab Setup

Figure 16 below displays a screenshot from NSX 6.2.2 where it shows a Web Security Group already created with one dynamically identified VM that's included as part of the Security Group; this correlates to the diagram above in Figure 15. It should be noted here that the Web Security Group is only being used locally on a local logical switch. The screenshot below shows the NSX Manager role as **Primary** because the deployment is part of a Cross-VC NSX deployment, but in this case only local objects and no universal objects are being used. To learn more about Cross-VC NSX and universal objects, see the following blog post on the VMware NSX Network Virtualization Blog : [Cross-VC NSX for Multi-site Solutions](#).

The Security Groups and policies discussed in this example are for local objects within one vCenter domain.

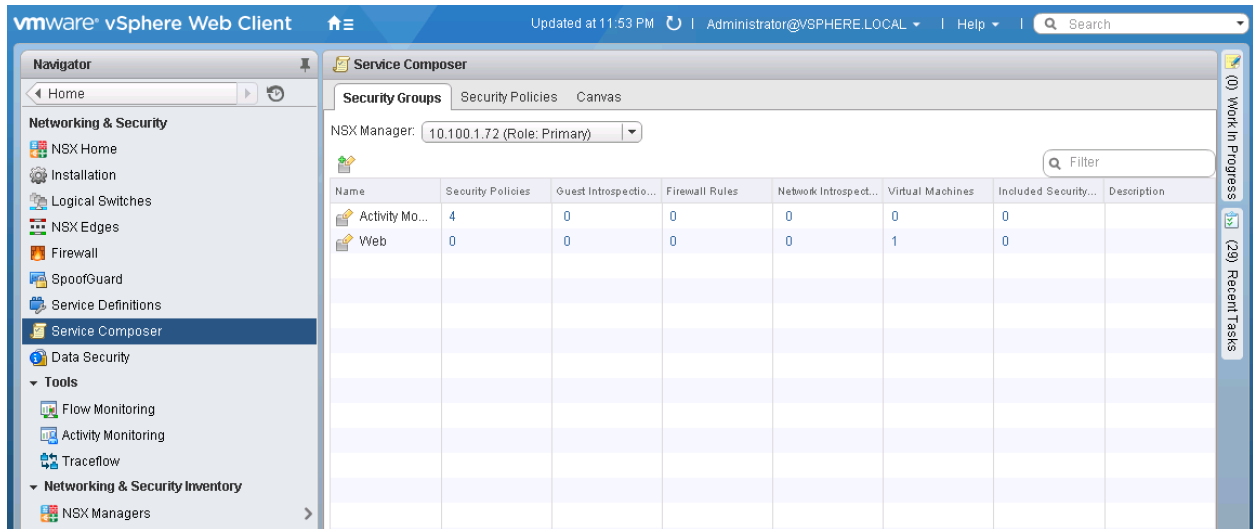


Figure 16: Existing Security Groups in Setup

A Security Group can be manually created as shown below where dynamic matching membership criteria is being defined. Imagine creating 100+ or even a 1000+ Security Groups manually.

Figure 17 below shows the dynamic membership criteria that an object must meet to be part of the **Web** Security Group. In this case the VM just has to contain the word **Web** in the VM Name.

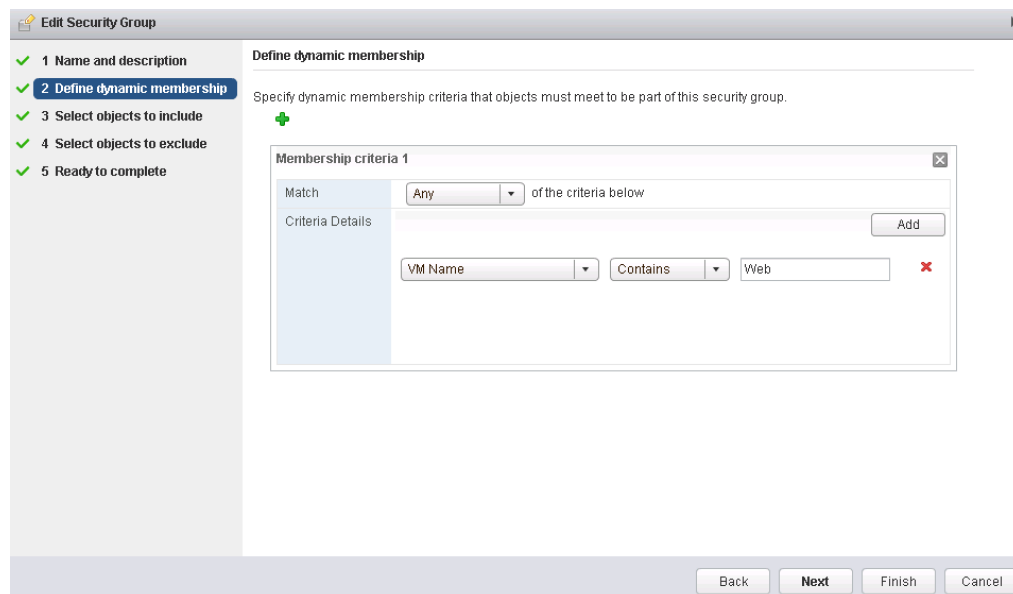


Figure 17: Using NSX Manager GUI to Create a Security Group

Currently, from Figure 16, it can be seen that the **Web** Security Group has been created yet the Security Group is not yet utilized in any security rules/policies. The procedural steps for creating the security posture in this case are:

1. Identify workload traffic
2. Create Security Groups to dynamically include workloads that should be included as part of the Security Group. This will help determine what security policies should be applied to the specific workloads based on requirements.
3. Create security policies leveraging created Security Groups

Let's apply this scenario to a larger scale; depending on specific environment and workloads, there can potentially be many Security Groups and respective security policies that need to be created. Let's assume the end user has 100+ Security Groups that need to be created. It can become tedious if one has to create so many Security Groups manually. One method of automation can be to leverage the NSX REST API and a bit of code. In this case we'll demonstrate with python.

For our example, first, we create a CSV file with one Security Group per row. Each row has two fields; the first field is the Security Group name, and, the second field is the matching criteria, or, in this case, specifically the word to match in the VM name (meaning, if the VM name contains this word, the VM will meet the criteria to dynamically be added to respective Security Group). The fields per row are separated via a comma within a CSV file.

Figure 18 below shows a screenshot of the file used as input to the python script in this example. Note, we demonstrate using just three Security Groups here based on our setup but many more as needed can be added and created with this batch-type Security Group creation script.




Figure 18: CSV File Used as Input to Python Script

By running the below python script, which reads in the CSV file as input and leverages the respective NSX REST API calls, all the Security Groups in the CSV file with the respective dynamic matching criteria are created.

First, the needed python libraries are imported, python variables used for the script are set, and a reader object is created to iterate over the lines in the CSV file. The python requests library is leveraged for the NSX REST API call and the python csv library is leveraged for reading and parsing the CSV input file. It's important to note that the requests library is not part of the Python standard library and must be installed. The below code is written in Python 2.7.7 and the Python package management system, **pip**, was

used to install the requests package from the cli with the following command: **pip install requests**.

Next, the python script iterates through the CSV file data via a loop reading the CSV data line by line and the respective NSX REST API to create the Security Group is called. The **security_group** python list variable contains two items, **security_group[0]**, which contains the Security Group name, and, **security_group[1]**, which contains the Security Group matching criteria for dynamic inclusion as discussed prior. For each iteration through the list, the NSX REST API call is made to create the respective Security Group.

The full python code is displayed below and was run in this example with NSX 6.2.2. The Python script is executed from the cli with the **python create_security_groups.py** command. The script could be run in Unix/Linux or Windows as long as the Python interpreter and respective **requests** library is installed.

```
1.  #Script: create_security_groups.py
2.  #
3.  #Description: Script reads input from a CSV file and creates Security Groups with
4.  #dynamic inclusion criteria. Used with NSX 6.2.2 and Python 2.7.11.
5.
6.
7.  import csv          #library used for reading and parsing data from CSV file
8.  import requests     #library used for making REST API calls
9.
10. #ignore specific warnings
11. from requests.packages.urllib3.exceptions import InsecureRequestWarning
12. requests.packages.urllib3.disable_warnings(InsecureRequestWarning)
13.
14. #initialize variables with needed info for input file and to make NSX REST API call
15. nsx_username = "nsxadmin"
16. nsx_password = "notMyPassword!"
17. nsx_url = "https://10.100.1.72/api/2.0/services/securitygroup/bulk/globalroot-0"
18. csv_filename = "security_groups.csv"
19. myheaders={'content-type':'application/xml'}
20.
21.
22. try:
23.     csv_file = open(csv_filename)          #open CSV file for reading
24. except FileNotFoundError as e:
25.     print "Input CSV file not found!"
26.
27. try:
28.     csv_data = csv.reader(csv_file)        #get reader object to iterate over lines in CSV
29.     file
30. except IOError as e:
31.     print "Error reading CSV file!"
32.
33. for security_group in csv_data:            #loop through parsed lines read in from CSV fil
34.     e
35.     #print security_group[0]              #uncomment for debugging - contains Security Group name
36.     #print security_group[1]              #uncomment for debugging - contains Security Group word
37.     to match in VM name
38.
39.     #create XML payload with Security Group data read in from CSV file
```

```

37.     payload = ''
38.     <securitygroup>
39.         <objectId>securitygroup-1</objectId>
40.         <objectTypeName>SecurityGroup</objectTypeName>
41.         <vsmUuid>422B2F08-25CF-AE4D-D78E-9450BA33618F</vsmUuid>
42.         <nodeId>d4fa9b9f-7973-4971-9ae6-c8f7e693ab30</nodeId>
43.         <revision>1</revision>
44.         <type>
45.             <typeName>SecurityGroup</typeName>
46.         </type>
47.         <name>'' + security_group[0] + ''</name>
48.         <description>
49.         </description>
50.         <scope>
51.             <id>globalroot-0</id>
52.             <objectTypeName>GlobalRoot</objectTypeName>
53.             <name>Global</name>
54.         </scope>
55.         <clientHandle>
56.         </clientHandle>
57.         <extendedAttributes/>
58.         <isUniversal>false</isUniversal>
59.         <universalRevision>0</universalRevision>
60.         <inheritanceAllowed>false</inheritanceAllowed>
61.         <dynamicMemberDefinition>
62.             <dynamicSet>
63.                 <operator>OR</operator>
64.                 <dynamicCriteria>
65.                     <operator>OR</operator>
66.                     <key>VM.NAME</key>
67.                     <criteria>contains</criteria>
68.                     <value>'' + security_group[1] + ''</value>
69.                     <isValid>true</isValid>
70.                 </dynamicCriteria>
71.             </dynamicSet>
72.         </dynamicMemberDefinition>
73.     </securitygroup>''
74.
75.     #print payload      #uncomment this for debugging - payload for REST API request call
76.
77.     #call NSX REST API to create Security Group with XML payload just created
78.     try:
79.         response = requests.post(nsx_url, data=payload, headers=myheaders, auth=(nsx_us
80. ername, nsx_password), verify=False)
81.     except requests.exceptions.ConnectionError as e:
82.         print "Connection error!"
83.
84.     print response.text

```

Figure 19: Python Script Creates Security Groups Based on Input From CSV File

Figure 20 below shows the results within the NSX GUI. As expected, based on the diagram in Figure 15, three Security Groups have been created each with one member VM. Note, the prior **Web** Security Group is still present which displays the same number of dynamically identified VMs as the **Web Security Group** created via python script, thus further validating the results.

Name	Security Policies	Guest Introspection Ser...	Firewall Rules	Network Introspection S...	Virtual Machines	Included Security Groups	Description
Activity Monitoring Data C...	4	0	0	0	0	0	
App Security Group	0	0	0	0	1	0	
DB Security Group	0	0	0	0	1	0	
Web	0	0	0	0	1	0	
Web Security Group	0	0	0	0	1	0	

Figure 20: Three Additional Security Groups Created

Perl

Perl is also quite a popular open-source programming language. One of Perl's key strengths is the vast amount of Perl modules/libraries available via the Comprehensive Perl Archive Network (CPAN). There is also a CPAN module included with Perl which is used to download and install Perl modules automatically from the CPAN repository.

Perl is a much older language than Python and thus has a larger amount of external libraries available. It's important to note that when using the popular REST Client **REST:Client** module with Perl as shown in the examples further below, it must be manually installed as it is not part of the Perl standard library. Example code below is written in Perl 5.24.0.1, and, the Perl CPAN module/command was used to install the **REST:Client** package from the cli with the following command: **cpan install REST:Client**. The **Data::Validate::IP** library, used to validate IPv4 and IPv6 addresses, is also not part of the Perl standard library and was also installed via CPAN.

Many beginning programmers also find Python easier to learn than Perl. Perl is more traditional in the sense that like other programming languages, outside of Python, it uses braces to designate blocks of code like loops, conditional statements, and functions. It also uses semicolons to designate end of statements. Perl is a powerful and very flexible programming language in the sense that there are many ways to accomplish the same task.

Perl also supports object oriented features such as user-defined classes and inheritance where objects can inherit from other parent objects.

Similar to Python, Perl's a cross-platform language able to run on multiple platforms such as MAC, Unix/Linux, and Windows. Perl is an interpreted language, and, as such, can run on any platform that has a Perl interpreter; this allows for its cross-platform capabilities.

Figure 21 below demonstrates how braces and semicolons are used with Perl to structure blocks of code.

```
1.      First line of Perl code;
2.
3.      Loop 1
4.      {
5.          Loop line 1;
6.          Loop line 2;
7.      }
8.      Seventh line of Perl code;
9.
10.     If this is true
11.     {
12.         Do this;
13.     }
14.     Else
15.     {
16.         Do this;
17.     }
18.
19.     Last line of Perl code;
```

Figure 21: Perl Code Structure

To learn more about Perl, see the [Perl Programming Language Website](http://www.perl.org/).

Example: Monitoring VMware NSX SpoofGuard with REST API and Perl

In this example, the program does a query to get Virtual NICs that have the SpoofGuard state of **Active**. SpoofGuard is a feature that can prevent the spoofing of IP addresses by trusting an IP address of a VM upon first use. In this state the IP address will be Active. If the IP address is later changed, communication will be blocked unless the IP address change is manually approved. Figure 22 below displays the NSX lab setup used in this example.




Figure 22: Lab Setup

The complete Perl code is shown further below; the script runs against an environment where each VM has one vNIC with both an IPv4 and IPv6 address. The program returns all **Active** Virtual NICs and associated information in regards to where SpoofGuard is enabled. NSX 6.2.2 and Perl 5.24.0.1 were utilized in this example.

The **REST:Client** module/library is used to provide the REST Client capabilities to communicate with NSX Manager. The **XML:Simple** library is used to easily read the XML output returned from the NSX REST API calls. It can be observed from the code that different tasks within the program are broken into functions and the respective functions are called in sequence based on requirements within the program.

First, the respective needed libraries are imported into the program and respective variables to connect to NSX Manager are initialized. Next a function call is made to obtain a REST Client object. Once the REST Client object is obtained, another function call is made to call the respective NSX REST API to get **Active** IP addresses and associated information in regards to where SpoofGuard is enabled.

Within the **getXMLData** function, once the respective data is returned in XML format from the NSX REST API call, the **XMLin** method of the **XML:Simple** object is used to read and parse the data. The resulting data is returned to the main program where it is looped

through, and, for each entry/record, Nic Name, approved MAC address, approved IPv4 address, and approved IPv6 address is output. The full code is shown below in Figure 24.

Before running the Perl script, we confirm via the NSX GUI for SpoofGuard what the **Active** Virtual NICs are. The **Web Policy** SpoofGuard policy in Figure 23 below is enabled on the web-tier shown as the **Web** logical switch in Figure 22 above. Further, it can be seen that there are two entries, one for the **Web** VM with IP address **172.100.10.1** and another for the **Web 3** VM with IP address **172.100.10.2**; this correlates with the lab diagram in Figure 22. Since the SpoofGuard operation mode is set to **Trust On First Use**, both these respective VMs' IP addresses are automatically initially trusted/approved and considered **Active**.

Virtual NIC	MAC Address	Virtual Machine	Last Approved Date	Approved IP	Detected IP
<input type="checkbox"/> Web 3 - Network adapter 1	00:50:56:93:37:a1	Web 3	6/16/2016	172.100.10.2 Clear	fe80::d097:8e5:7e4a:e6ac 172.100.10.2
<input type="checkbox"/> Web - Network adapter 1	00:50:56:93:75:21	Web	6/16/2016	172.100.10.1 Clear	fe80::4c89:3622:3782:9cc1 172.100.10.1

Figure 23: NSX GUI Displaying SpoofGuard Active Virtual NICs

The Perl script is executed from the cli with the **perl spoofguard_active_ips.pl** command. The full code is shown below. The script could be run in Unix/Linux or Windows as long as the Perl interpreter and respective **REST::Client** and **Data::Validate::IP** Perl modules are installed.

```

1.  #Script: spoofguard_active_ips.pl
2.  #
3.  #Description: Script runs against an environment where each VM has a single vNIC with
4.  #both an IPv4 and IPv6 address. The program returns all "Active" IP addresses in
5.  #relation to SpoofGuard and respective Nic Name, approved MAC Address, approved IPv4
6.  #Address, and approved IPv6 Address. Used with NSX 6.2.2 and Perl 5.24.0.1.
7.
8.
9.  use REST::Client; #REST Client library used to make REST API calls
10. use MIME::Base64; #library used for base64 encoding
11. use XML::Simple;  #library used for ease of reading XML

```

```

12. use Data::Validate::IP qw(is_ipv4 is_ipv6); #library used to validate IPv4 and IPv6
    addresses
13.
14.
15. #uncomment below line if getting SSL connectivity/certificate error
16. $ENV{"PERL_LWP_SSL_VERIFY_HOSTNAME"} = 0; #disable certificate check
17.
18. #function to draw separator between records
19. sub drawSeparator
20. {
21.     print "-----\n";
22.     return;
23. }
24.
25.
26. #function to print heading
27. sub printHeader
28. {
29.     print "SpoonGuard Active Virtual NICs: \n\n";
30.     return;
31. }
32.
33.
34. #function creates and returns REST Client object
35. #takes NSX Manager URL with IP address as argument
36. sub getRestClient
37. {
38.     #create REST Client object
39.     my ($arg) = @_;
40.
41.     my $restclient = REST::Client->new(host => $arg);
42.     return $restclient;
43. }
44.
45.
46. #function calls respective NSX REST API call and returns results
47. #takes REST Client object, NSX REST API call, base64 encoded
48. #username and password, and REST API 'content-type' header variable
49. #as arguments
50. sub getXMLData
51. {
52.     my ($client, $nsx_rest_spoofguard_active_api, $encoded_auth, $nsx_content_type)
    = @_;
53.
54.     #make NSX REST API call to get "Active" Virtual NICs and respective info for
    spoofguard
55.     $client->GET($nsx_rest_spoofguard_active_api,
56.         {"Authorization" => "Basic $encoded_auth",
57.          "Accept" => $nsx_content_type});
58.
59.     $xmlResponse = new XML::Simple; #create object for ease of reading XML
60.
61.     #read in respective XML data from NSX API response using XMLin method
62.     $xmlData = $xmlResponse->XMLin($client->responseContent(), KeyAttr =>
    ['spoofguard']);
63.
64.     return $xmlData
65. }
66.
67.
68. #initialize variables for connecting to NSX Manager and calling NSX REST API

```

```

69. my $nsx_manager = "https://10.10.10.72";
70. my $nsx_username = "admin";
71. my $nsx_password = "notMyPassword!";
72. my $nsx_rest_spoofguard_active_api = "/api/4.0/services/spoofguard/spoofguardpolicy-5?list=ACTIVE";
73. my $nsx_content_type = "application/xml";
74.
75. #use Base64 encoding for transmitting respective data
76. my $encoded_auth = encode_base64("$nsx_username:$nsx_password");
77.
78.
79. my $client = getRestClient($nsx_manager); #get REST Client object
80.
81. #get NSX REST API call result data
82. my $xmlData = getXMLData($client, $nsx_rest_spoofguard_active_api, $encoded_auth,
    $nsx_content_type);
83.
84. printHeader();
85.
86. #loop through respective results and for each entry/record, print Nic Name,
87. #approved MAC Address, approved IPv4 Address, and approved IPv6 Address
88. foreach $xmlRecord (@{$xmlData->{'spoofguard'}})
89. {
90.     drawSeparator();
91.     print "Nic Name: \t\t" . $xmlRecord->{'nicName'} . "\n";
92.     print "Approved Mac Address: \t" . $xmlRecord->{'approvedMacAddress'} . "\n";
93.
94.     $ip_address_1 = $xmlRecord->{'approvedIpAddress'}->{'ipAddress'}->[0];
95.     $ip_address_2 = $xmlRecord->{'approvedIpAddress'}->{'ipAddress'}->[1];
96.
97.     if(is_ipv4($ip_address_1))
98.     {
99.         print "Approved IPv4 Address:\t" . $ip_address_1 . "\n";
100.        print "Approved IPv6 Address:\t" . $ip_address_2 . "\n";
101.    }
102.    else
103.    {
104.        print "Approved IPv4 Address:\t" . $ip_address_2 . "\n";
105.        print "Approved IPv6 Address:\t" . $ip_address_1 . "\n";
106.    }
107. }
108. drawSeparator();
109.
110.
111. #uncomment below lines to see response status and headers
112. #print "Response status: " . $client->responseCode() . "\n";
113. #foreach ($client->responseHeaders())
114. #{
115. # print $_ . "=" . $client->responseHeader($_) . "\n";
116. #}

```

Figure 24: Perl Script For Displaying “Active” IP Addresses for SpoofGuard

Once the above script is run with the **perl spoofguard_active_ips.pl** command at the cli, the below output is generated. As can be seen, the results match what is displayed in the GUI in Figure 23 confirming the results are correct.

SpoofGuard Active Virtual NICs:

```
-----  
Nic Name:                Web - Network adapter 1  
Approved Mac Address:    00:50:56:93:75:21  
Approved IPv4 Address:   172.100.10.1  
Approved IPv6 Address:   fe80::4c89:3622:3782:9cc1  
-----  
Nic Name:                Web 3 - Network adapter 1  
Approved Mac Address:    00:50:56:93:37:a1  
Approved IPv4 Address:   172.100.10.2  
Approved IPv6 Address:   fe80::d097:8e5:7e4a:e6ac  
-----
```

Figure 25: Output of Perl Script Displaying "Active" IP Addresses for SpoofGuard

In this example, entries with the SpoofGuard state of **ACTIVE** are shown. To see entries with other SpoofGuard states, we can simply modify the key value pair at the end of the query. For example, to see entries with **REVIEW_PENDING** state, the **list=ACTIVE** key value pair at the end of the NSX REST API call can be changed to **list=REVIEW_PENDING**. In the code above, the `$nsx_rest_spoofguard_active_api` variable can be changed to `/api/4.0/services/spoofguard/spoofguardpolicy-5?list=REVIEW_PENDING`.

Note, for purposes of calling the NSX REST API, the default policy ID is always **spoofguardpolicy-1**. If custom policies have been created, the policy ID can be retrieved via REST API Get call with the following URL: **`https://[NSX Manager IP Address]/api/4.0/services/spoofguard/policies/`**.

In Figure 23, it can be seen there are no entries that need review for the Web Policy SpoofGuard policy as shown by the **0** under the **Need Review** column. For demonstration, the IP address of the **Web** VM with IP address **172.100.10.1** is changed to **172.100.10.3** and the IP address of **Web 3** VM with IP address **172.100.10.2** is changed to **172.100.10.4**. These IP address changes are detected by SpoofGuard. As shown in Figure 26 below, before the VMs are allowed to communicate using the new respective IP addresses, the changes must be approved. The column **Need Review** has also changed to **2** as expected. Selecting the **Virtual Nics IP Required Approval** option from the drop down box displays the entries needing review/approval. Note, the **Approved IPv4 Address** and **Detected IPv4 Address** for each of the two entries respectively is different.

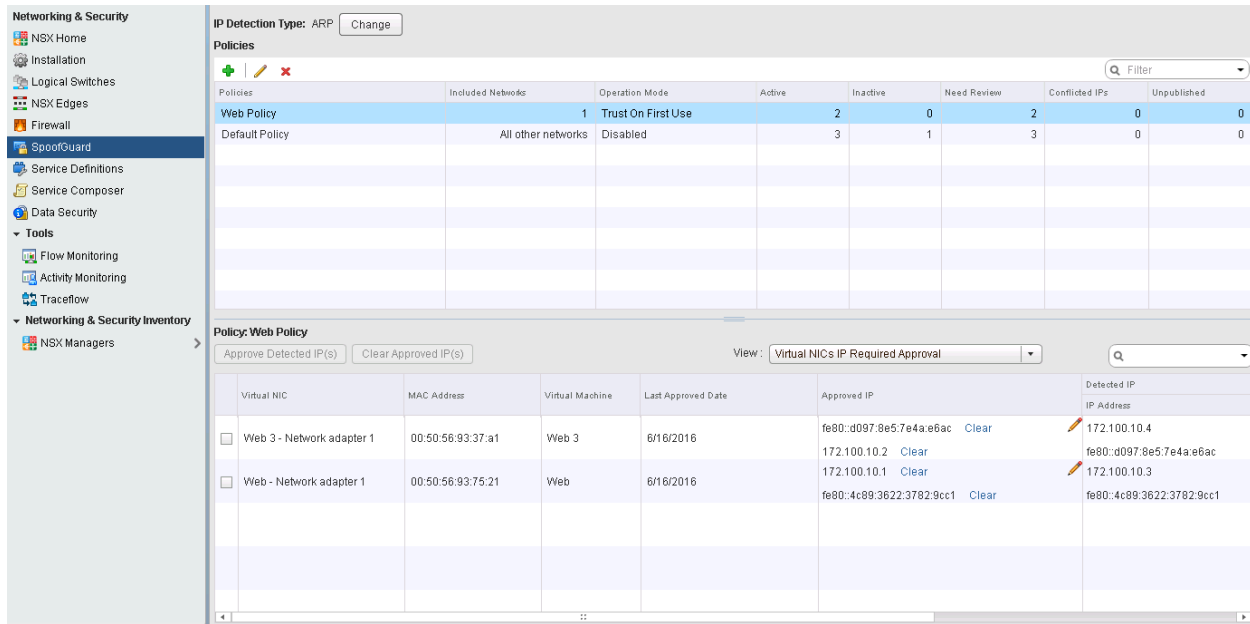


Figure 26: Virtual NIC IP Address Changes Requiring Review/Approval

If in the Perl code, the `$nsx_rest_spoofguard_active_api` variable is changed to `/api/4.0/services/spoofguard/spoofguardpolicy-1?list=REVIEW_PENDING`, information can be displayed about IP Addresses that SpooGuard has blocked and need review/approval. As can be seen in Figure 27 below, the correct results are displayed matching the results shown by the GUI in Figure 26. Note, additional fields and respective code was added to show the corresponding detected MAC and IP Addresses. These new fields correspond to the **detectedMacAddress** and **detectedIPAddress** data fields that can be obtained from the **\$xmlRecord** variable within the script. The respective title/header for the results was also updated.

SpoofGuard Virtual NIC IP Addresses Pending Review:

```
-----  
Nic Name:                Web - Network adapter 1  
Approved Mac Address:    00:50:56:93:75:21  
Approved IPv4 Address:   172.100.10.1  
Approved IPv6 Address:   fe80::4c89:3622:3782:9cc1  
  
Detected Mac Address:    00:50:56:93:75:21  
Detected IPv4 Address:   172.100.10.3  
Detected IPv6 Address:   fe80::4c89:3622:3782:9cc1  
-----  
Nic Name:                Web 3 - Network adapter 1  
Approved Mac Address:    00:50:56:93:37:a1  
Approved IPv4 Address:   172.100.10.2  
Approved IPv6 Address:   fe80::d097:8e5:7e4a:e6ac  
  
Detected Mac Address:    00:50:56:93:37:a1  
Detected IPv4 Address:   172.100.10.4  
Detected IPv6 Address:   fe80::d097:8e5:7e4a:e6ac  
-----
```

Figure 17: Output of Perl Script – SpoofGuard IP Addresses Pending Review

Go

Go (also often referred to as Golang) is an open source, compiled, programming language developed by Google and many community contributors since 2007. It's distributed under a BSD-style license.

Go's mission is to “make easy to build simple, reliable and efficient software.” It provides several packages with built-in functions that allow programmers to develop software quickly. For example, it provides HTTP and XML packages that can be used to interact with REST APIs very effectively. The Go language is rapidly gaining traction in the open source community and several popular projects (Docker, Terraform) are written in Go. Additional information about Go, as well as examples and tutorials can be found on the [Go project's web site](#).

Go provides REST Client libraries that can be used to directly interact with the NSX REST API. There are also open source libraries that allow RAML to be parsed and used. This document is not meant to be a guide for the Go programming language, and a simple example is provided below that does not leverage RAML parsing.

Example: Creating NSX Services in Go

Services (or Applications) are used in NSX to define port numbers used in the Distributed Firewall rules. Once NSX is installed, a large number of built-in services for the most common use cases is already provided; however, users might be required to define new custom services.

Below, a simple program is shown written in Go that will create custom NSX Services using the NSX REST API, allowing the user to specify name, description, protocol (TCP, UDP, etc.) and value (a comma-separated list of ports and port ranges).

The program consists of code in a single **main** function that performs the following:

- Checks initial arguments to retrieve host name of the NSX Manager, credentials and an input file
- Parses the input file (semicolon separated list of entries) to retrieve the list of Services to create
- For each Service, an XML document is generated and sent over the NSX Manager API using the HTTP POST method
- Parses the return code, and, if successful, returns the NSX object ID of the created Service

The NSX REST API Guide, under section **Add Service**, specifies that all services must be created by executing an HTTP POST to the following URL:

`https://nsxmgr-ip/api/2.0/services/application/scopeId`

The guide also mentions that for all Services the **scopeId** parameter in the URL corresponds to the global scope (**globalroot-0**). The NSX API documentation also provides a detailed description of the XML structure that must be included in the HTTP request body to create the new Service:

```
1. <application>
2.   <objectId/>
3.   <type>
4.     <typeName/>
5.   </type>
6.   <description>(application description)</description>
7.   <name>(application name)</name>
8.   <revision>0</revision>
9.   <objectTypeName/>
10.  <element>
11.    <applicationProtocol>(application protocol)</applicationProtocol>
12.    <value>(port value(s))</value>
13.  </element>
14. </application>
```

Figure 28: XML Structure to Insert in the Request Body

Some of the parameters are empty or fixed, and the relevant ones are highlighted in red:

description, **name**, **applicationProtocol** and **value**. All these parameters are strings of text. As can be seen, this program does not do any validation of the inputs: if some illegal argument is specified on the command line or in the input file, it will be passed directly to the NSX API. If errors are encountered, NSX will return respective errors and the program will exit. Best practice is to add validation routines to such programs.

The following figure shows how the semicolon separated input file looks (comments begin with the **#** character):

```
1. # name, description, applicationProtocol, value
2. Test1;Test TCP Protocol;TCP;444
3. Test2;Test UDP Protocol;UDP;4871
4. Test3;Test TCP Protocol Range;TCP;181,1000-2000
5. Test4;Test Protocol Range;UDP;391,1821-3291
6. Test5;Test FTP;FTP;2121
```

Figure 29: Protocols Definition Input File

The program's Go code is presented below, and is compiled from the **nsx-importservices.go** source file. The reader is encouraged to refer to the comments in the code to better understand the program flow:

```
1. package main
2.
3. // Import required packages
4. import (
5.     "encoding/xml"
6.     "fmt"
7.     "os"
8.     "path/filepath"
9.     "net/http"
10.    "io/ioutil"
11.    "strings"
12.    "crypto/tls"
13.    "encoding/csv"
14.    "bytes"
15. )
16.
17. // Define the NSX structure for the NSX service ("application") as a set of structs
18. // The desired XML format is as follows:
19. //<application>
20. // <objectId/>
21. // <type>
22. // <typeName/>
23. // </type>
24. // <description>(application description)</description>
25. // <name>(application name)</name>
26. // <revision>0</revision>
27. // <objectTypeName/>
28. // <element>
29. // <applicationProtocol>(application protocol)</applicationProtocol>
30. // <value>(port value(s))</value>
31. // </element>
32. //</application>
33.
34. type Application struct {
```

```

35.     XMLName xml.Name    `xml:"application"`
36.     ObjectId string `xml:"objectId"`
37.     Revision    int    `xml:"revision"`
38.     Type        Type    `xml:"type"`
39.     Name        string  `xml:"name"`
40.     Description string  `xml:"description"`
41.     Element []Element  `xml:"element"`
42. }
43.
44. type Type struct {
45.     XMLName xml.Name    `xml:"type"`
46.     TypeName string  `xml:"typeName"`
47. }
48.
49. type Scope struct {
50.     XMLName xml.Name    `xml:"scope"`
51.     Id      string  `xml:"id"`
52.     ObjectTypeName string `xml:"objectTypeName"`
53.     Name    string  `xml:"name"`
54. }
55.
56. type Element struct {
57.     XMLName xml.Name    `xml:"element"`
58.     ApplicationProtocol string `xml:"applicationProtocol"`
59.     Value      string  `xml:"value"`
60. }
61.
62. func main() {
63.
64.     // Check arguments and return error if incorrect
65.     // Note: there is no formal validation of the arguments in this example
66.     if(len(os.Args) != 5) {
67.         fmt.Printf("Syntax error\nUsage: %s [NSX Manager Address] [Username] [Password]
[Input File]\n\n", os.Args[0])
68.         os.Exit(1)
69.     }
70.     nsxManager := os.Args[1]
71.     nsxUser    := os.Args[2]
72.     nsxPassword := os.Args[3]
73.     inputFileName := os.Args[4]
74.
75.     // Check the path and open the input file for read
76.     readFilePath, err := filepath.Abs(inputFileName)
77.     if err != nil {
78.         fmt.Println(err)
79.         os.Exit(1)
80.     }
81.     csvfile, err := os.Open(readFilePath)
82.     if err != nil {
83.         fmt.Println(err)
84.
85.         os.Exit(1)
86.     }
87.     // make sure the descriptor will be closed later
88.     defer csvfile.Close()
89.
90.     // Read and parse the CSV file
91.     reader := csv.NewReader(csvfile)
92.     // Four fields are expected for each record, comments start with '#' and separator
is semicolon
93.     reader.FieldsPerRecord = 4

```

```

94.     reader.Comment = '#'
95.     reader.Comma = ';'
96.     rawCSVdata, err := reader.ReadAll()
97.     if err != nil {
98.         fmt.Println(err)
99.         os.Exit(1)
100.    }
101.
102.
103.    // Parse the CSV data and add it a slice of Application types
104.    // Warning: there is no formal validation of the input in this example, all argumen
    ts are passed to the NSX API
105.    var nsxServices []Application
106.    for _, chunk := range rawCSVdata {
107.        var x Application
108.        var e Element
109.        x.Name = chunk[0]
110.        x.Description = chunk[1]
111.        e.ApplicationProtocol = chunk[2]
112.        e.Value = chunk[3]
113.        x.Element = append(x.Element, e)
114.        x.Revision = 0
115.
116.        fmt.Printf("Read application from CSV file:\n\tName: %s\n\tDescription: %s\n\tP
    rotocol: %s\n\tValue: %s\n", chunk[0], chunk[1], chunk[2], chunk[3])
117.        nsxServices = append(nsxServices, x)
118.    }
119.
120.    // Initialize the HTTPS client to skip SSL certificate verification
121.    tr := &http.Transport{ TLSClientConfig: &tls.Config{InsecureSkipVerify: true}}
122.    client := &http.Client{Transport: tr}
123.
124.    // For each application read from the file, create an HTTPS POST request to the NSX
    Manager
125.    for _, srv := range nsxServices {
126.
127.        // Encode the XML representation of this object in a buffer
128.        var buf bytes.Buffer
129.        xmlbuf := xml.NewEncoder(&buf)
130.        err = xmlbuf.Encode(&srv)
131.        if err != nil {
132.            os.Exit(1)
133.        }
134.        // Uncomment the following line to output the created XML body
135.        //fmt.Println("XML Body: " + buf.String())
136.
137.        // Prepare the HTTPS POST request with the XML body
138.        fmt.Println("Importing service: " + srv.Name)
139.        req, err := http.NewRequest("POST", "https://" + nsxManager + "/api/2.0/service
    s/application/globalroot-0", &buf )
140.        if err != nil {
141.            fmt.Println(err)
142.            os.Exit(1)
143.        }
144.
145.        // Configure Basic authentication and Content-Type
146.        req.SetBasicAuth(nsxUser, nsxPassword)
147.        req.Header.Set("Content-Type", "application/xml")
148.
149.        // Execute the HTTPS request
150.        resp, err := client.Do(req)

```

```

151.     if err != nil {
152.         fmt.Println(err)
153.         os.Exit(1)
154.     }
155.     defer resp.Body.Close()
156.
157.     // Check the response type (expected 201 - Created)
158.     fmt.Println("HTTP Response is: " + resp.Status)
159.     if(resp.StatusCode != 201) {
160.         b, err := ioutil.ReadAll(resp.Body)
161.         if err != nil {
162.             fmt.Println(err)
163.             os.Exit(1)
164.         }
165.         fmt.Println("Error in executing query - got the following error:")
166.         fmt.Println(string(b))
167.         os.Exit(1)
168.     }
169.
170.     // NSX returns the created application id as the last field of the Location header URL
171.     // Parse it and print it out
172.     loc, err := resp.Location()
173.     if err != nil {
174.         fmt.Println(err)
175.         os.Exit(1)
176.     }
177.     l := strings.Split(loc.String(), "/")
178.     fmt.Println("Successfully created application Object ID: " + l[len(l)-1])
179. }
180. }

```

Figure 30: Go Program to Create Custom NSX Services Leveraging the Relevant NSX REST API

In this example, we used XML structures to define the document hierarchy and leveraged the built-in XML encoder to generate the HTTP body. Users can also choose to build the HTTP body by simply concatenating strings with variables. There is no preferred option: the choice depends on the desired flexibility and reusability of the code.

The following picture shows the output of the program when run with the correct parameters:

```

root@dev:~# ./nsx-importservices nsxmgr-01a admin Password1 input.csv
Read application from CSV file:
    Name: Test1
    Description: Test TCP Protocol
    Protocol: TCP
    Value: 444
Read application from CSV file:
    Name: Test2
    Description: Test UDP Protocol
    Protocol: UDP
    Value: 4871
Read application from CSV file:
    Name: Test3
    Description: Test TCP Protocol Range
    Protocol: TCP
    Value: 181,1000-2000
Read application from CSV file:
    Name: Test4
    Description: Test Protocol Range
    Protocol: UDP
    Value: 391,1821-3291
Read application from CSV file:
    Name: Test5
    Description: Test FTP
    Protocol: FTP
    Value: 2121
Importing service: Test1
HTTP Response is: 201 Created
Successfully created application Object ID: application-401
Importing service: Test2
HTTP Response is: 201 Created
Successfully created application Object ID: application-402
Importing service: Test3
HTTP Response is: 201 Created
Successfully created application Object ID: application-403
Importing service: Test4
HTTP Response is: 201 Created
Successfully created application Object ID: application-404
Importing service: Test5
HTTP Response is: 201 Created
Successfully created application Object ID: application-405
root@dev:~# █

```

Figure 31: Output After a Successful Run of the Go Program

After a successful run, the program outputs that the services were correctly created (HTTP Post returned **201 Created**). In the example shown above, the object IDs of the created services are **application-401**, **application-402**, **application-403**, **application-404** and **application-405**: these handlers can be used for subsequent CRUD operations.

Finally, it's possible to connect to the NSX Manager UI and observe that the new services were created according to the information specified in the input file; this is shown below.

Name	Protocol	Destination ports	Source ports	Scope	Inheritance
Test1	TCP	444	any	Global	✗
Test2	UDP	4871	any	Global	✗
Test3	TCP	181,1000-2000	any	Global	✗
Test4	UDP	391,1821-3291	any	Global	✗
Test5	FTP	2121	any	Global	✗

Figure 32: NSX Services Created Using a Go program – NSX GUI verification

Automation Tools with NSX REST API

The previous sections documented how to consume the NSX REST API using clients such as a web browser, cURL, or via different programming languages. Administrators and engineers use such methods to build scripts that avoid the manual execution of repeatable tasks, and these methods are very effective. However, sometimes these methods can produce challenges in terms of management and maintenance.

To streamline the automation of tasks in a more structured manner, products, such as orchestrators and configuration management tools can be leveraged. With these solutions, users can define workflows or playbooks that include a list of tasks which will be automatically executed when some event occurs. Orchestrators and configuration management tools are often modular and highly configurable, allowing customers and partners to add additional capabilities to the product.

This document is not focused on discussing how these technologies work or how they differ from each other, but presents examples on how to leverage the NSX REST API using the respective tools. Two example products discussed in more detail below are VMware vRealize Orchestrator, a solution adopted by many VMware customers, and Ansible, an open source platform written in Python.

vRealize Orchestrator

VMware vRealize Orchestrator (vRO) is a powerful solution that simplifies the automation of complex IT tasks and integrates with several VMware and 3rd party components to extend service delivery and operational management. Given its flexibility and integration capabilities, many VMware customers use it for their automation strategy. It's packaged with vCenter Server, and, at the time of writing, does not require additional licenses.

This document is not meant to be a comprehensive guide on vRealize Orchestrator, but a few concepts below are important to mention:

- vRO is modular software and can be extended through Java plugins. VMware provides plugins to integrate it with many solutions, including vCenter and NSX
- vRO also includes a number of general purpose plugins that allow users to integrate with other solutions through several technologies, including AMQP, SSH, SQL, JDBC, PowerShell and HTTP REST APIs
- VMware Solution Partners ship plugins that can be imported in vRO and used to create workflows that involve the integration with 3rd party components. Available plugins can be downloaded from [VMware's Solution Exchange web site](#).
- Users can build simple or complex workflows using vRO UI by simply concatenating different components. vRO also ships with an integrated Javascript parser that can be used to build the required logic and is easy to learn
- vRO workflows can be invoked manually, through APIs, by vRA, and AMQP message bus. The workflows can also be scheduled to run at specific times.

For additional details on vRealize Orchestrator please refer to the [VMware vRealize Orchestrator web site](#). Figure 33 below displays a vRO workflow being built and leveraging NSX REST API.

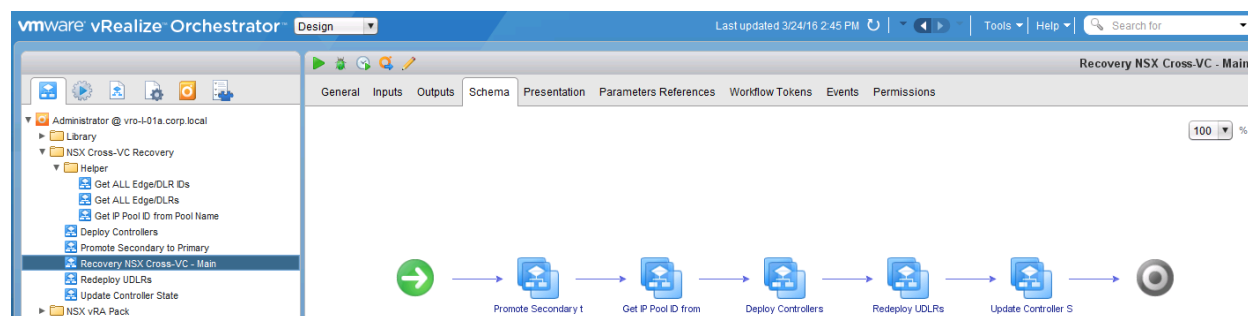


Figure 33: Building a vRO Workflow

vRO also provides a HTTP REST API client allowing for NSX REST API to be called directly from the tool; a NSX vRO Plugin is also available.

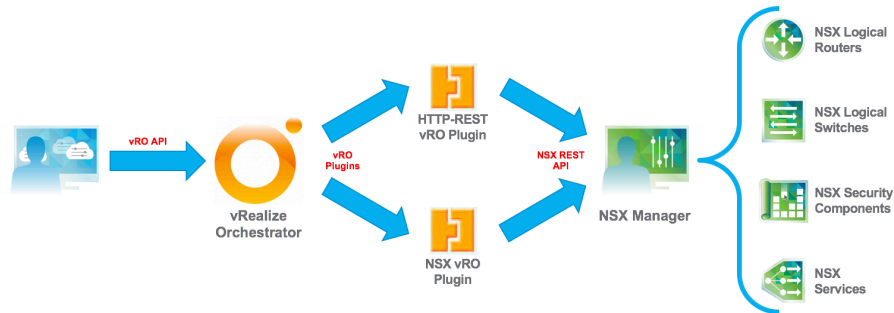


Figure 34: vRO Leveraging NSX REST API

As mentioned, VMware ships an NSX plugin that includes a number of built-in workflows that allow basic consumption of NSX features without requiring knowledge of the NSX API. The below screenshot shows the existing workflows available in the NSX vRO plugin at the time of writing.

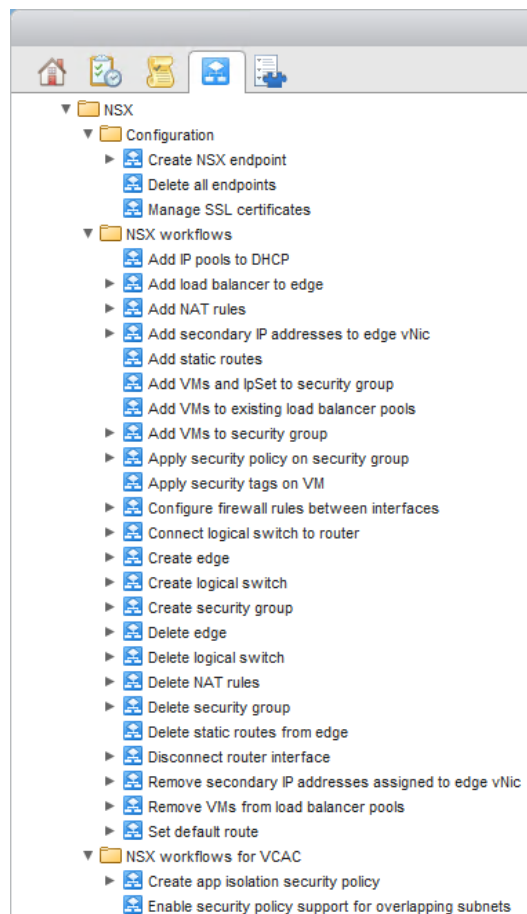


Figure 35: NSX vRO Plugin Workflows

More capabilities will be added over time as the initial focus of the NSX plugin is meant to enable the NSX integration with vRealize Automation. More information on this topic can be found in the vRealize Automation section of this document.

While several basic NSX operations are already included in the plugin, users might be required to automate tasks that are not natively present: extensibility can easily be done as vRO ships with a powerful HTTP REST API plugin that administrators can leverage to build custom integrations with NSX.

Two examples are provided here on how to create vRealize Orchestrator workflows using the HTTP REST API plugin and JavaScript to achieve the following:

- Create an IP Set in NSX
- Enable or Disable HA on an NSX Edge

These examples require basic knowledge of vRO, but are easy enough to follow even for beginners. Please refer to the [vRealize Orchestrator page](#) on the VMware Technology Network website for additional information and resources.

Example: Create an IP Set in NSX Using vRealize Orchestrator

IP Sets can be included within NSX Security Groups or used natively in the Distributed Firewall (DFW). IP Sets can be used to describe objects on the network that are not present within the vCenter environment (i.e. an external physical server).

We'll create a vRealize Orchestrator workflow that will create a new IP Set on the NSX Manager, allowing the user to specify name, description and value (a comma-separated list of IP addresses).

Before creating the workflow, we'll register our NSX Manager as an endpoint in our vRealize Orchestrator inventory: this is a common practice that will allow users to not have to specify the NSX Manager IP address and credentials every time, but just refer to the NSX Manager endpoint as an object (of type **REST:RestHost**).

To create the endpoint just run the **Add a REST host** workflow in the **Library → HTTP-REST → Configuration** folder and specify the required NSX Manager information (URL, credentials, timeout and certificate), as shown in the below figures.

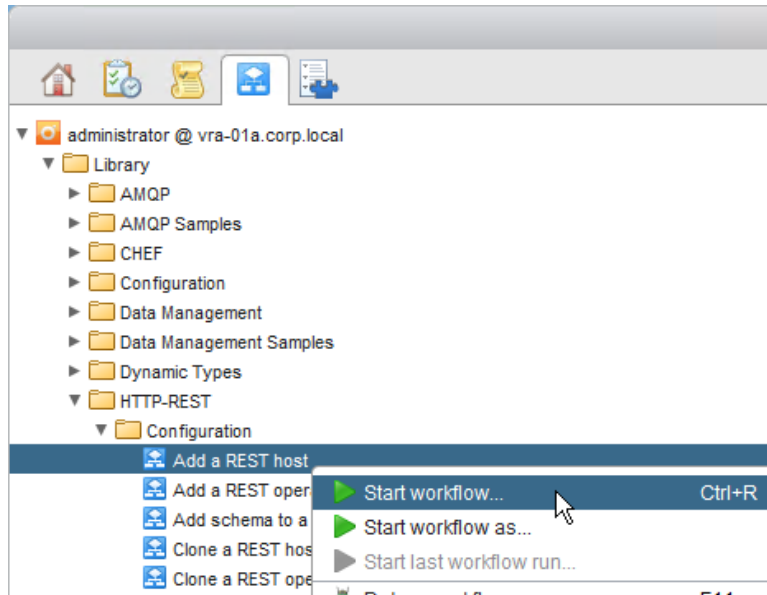


Figure 36: Run the "Add a REST host" Workflow in vRO

 A screenshot of the 'Start Workflow : Add a REST host' configuration window. The window has a blue header and a sidebar on the left with a tree view showing the workflow steps: '1 Add a REST host' (selected), '1a Host Properties', '2 Host Authentication', '2a', '2b User credentials', '3 Proxy Settings', '3a Proxy Settings', '4 Advanced', and '4a SSL'. The main area contains the following fields and options:

- Properties to create a new host. The name is the host's unique identifier.**
- * Name:** A text field containing 'NSX Manager'.
- * URL:** A text field containing 'https://nsxmgr-01a/api'.
- Connection timeout (seconds):** A text field containing '30.0'.
- Operation timeout (seconds):** A text field containing '180.0'.
- If set to true, the certificate is accepted silently and the certificate is added to the trusted store.**
- Yes/No radio buttons:** The 'Yes' radio button is selected.

Figure 37:: Specify NSX Manager IP Address, Timeout and Certificate in vRO

Start Workflow : Add a REST host

- 1 Add a REST host
 - 1a Host Properties
 - 2 Host Authentication
 - 2a
 - 2b User credentials
 - 3 Proxy Settings
 - 3a Proxy Settings
 - 4 Advanced
 - 4a SSL

* Host's authentication type
Basic

Figure 38: Specify NSX Manager Basic Authentication in vRO

Start Workflow : Add a REST host

- 1 Add a REST host
 - 1a Host Properties
 - 2 Host Authentication
 - 2a
 - 2b User credentials
 - 3 Proxy Settings
 - 3a Proxy Settings
 - 4 Advanced
 - 4a SSL

If Shared Session is selected, the provided credentials are used. If Per User Session option is used, the credentials are retrieved from the currently logged-in user.

Session mode
Shared Session

* Authentication user name
admin

* Authentication password

Figure 39: Specify NSX Manager Credentials in vRO

Please note that the **Operation timeout (seconds)** parameter is set to 180 seconds as some tasks on NSX are blocking (they do not return until the operation in the backend is completed) and the default value (60 seconds) could lead to timeout errors.

Check whether the workflow has been executed correctly and if the endpoint is shown in the HTTP-REST section of the vRO inventory:




Figure 40: Verify vRO Workflow Completion

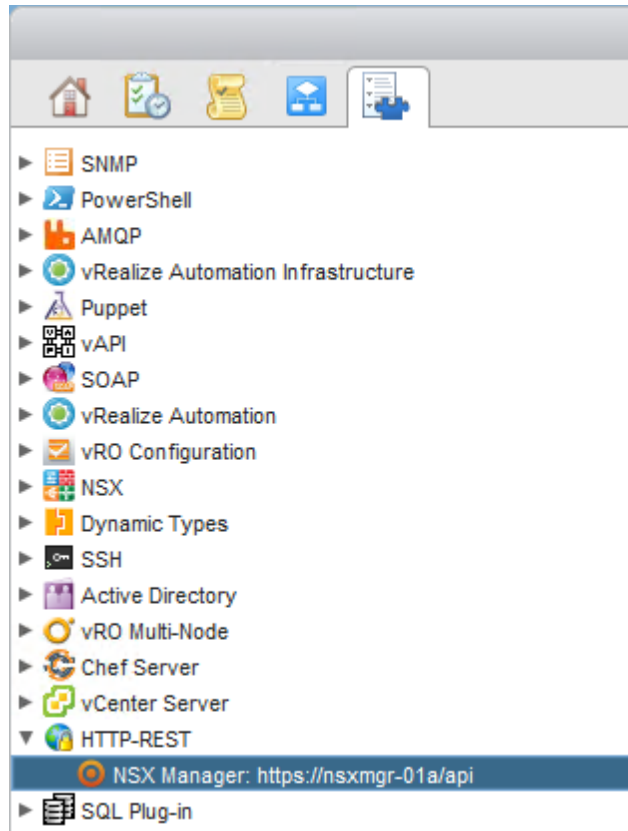


Figure 41: vRO Inventory with NSX as an HTTP-REST Endpoint

We can now prepare a vRealize Orchestrator workflow that creates a new IP Set in NSX. The workflow contains a single scriptable task, which includes JavaScript code that:

- Builds an XML document according to the NSX API specification (refer to the NSX API guide)
- Sends an HTTP POST request to the NSX Manager with the XML document as the body
- Parses the return code and, if successful, returns the NSX object ID of the created IP Set

The NSX API documentation, under section **Create an IPset**, specifies that all IP Sets must be created by executing an HTTP POST to the following URL:

`https://nsxmgr-ip/api/2.0/services/ipset/scopeId`

The guide also mentions that for all IP Sets the **scopeId** parameter in the URL corresponds to the global scope (**globalroot-0**).

The NSX API documentation also provides a detailed description of the XML document that must be included in the HTTP request body to create the new IP Set:

```

1. <ipset>
2. <objectId />
3. <type>
4. <typeName />
5. </type>
6. <description>
7. (IP Set description)
8. </description>
9. <name>(IP Set name)</name>
10. <revision>0</revision>
11. <objectTypeName />
12. <value>(comma separate list of IPs)</value>
13. </ipset>

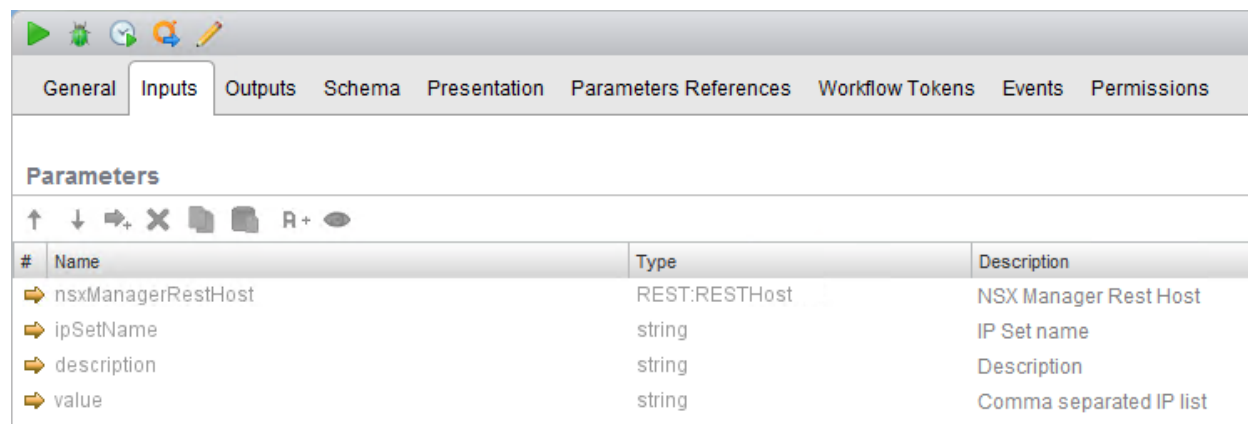
```

Figure 42: XML Schema

Some of the parameters are empty or fixed, and, the relevant ones are highlighted in red: **description**, **name** and **value**. All these parameters are strings of text.

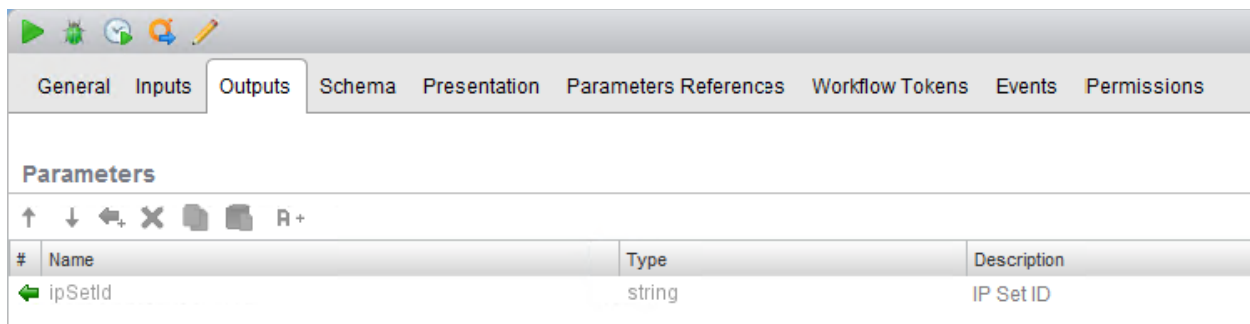
The vRealize Orchestrator workflow can be created interactively using the UI by dragging and dropping components in the canvas as well as binding input and output parameters. This document assumes that the reader has the basic knowledge to use vRO: please refer to the respective documentation for additional information.

The following figures show how the vRO workflow looks like (inputs, outputs and schema):



#	Name	Type	Description
➡	nsxManagerRestHost	REST:RESTHost	NSX Manager Rest Host
➡	ipSetName	string	IP Set name
➡	description	string	Description
➡	value	string	Comma separated IP list

Figure 43: "Create NSX IP Set" vRO Workflow Inputs



The screenshot shows the 'Outputs' tab of a vRO workflow. The 'Parameters' section contains a table with one output parameter.

#	Name	Type	Description
1	ipSetId	string	IP Set ID

Figure 44: “Create NSX IP Set” vRO Workflow Outputs

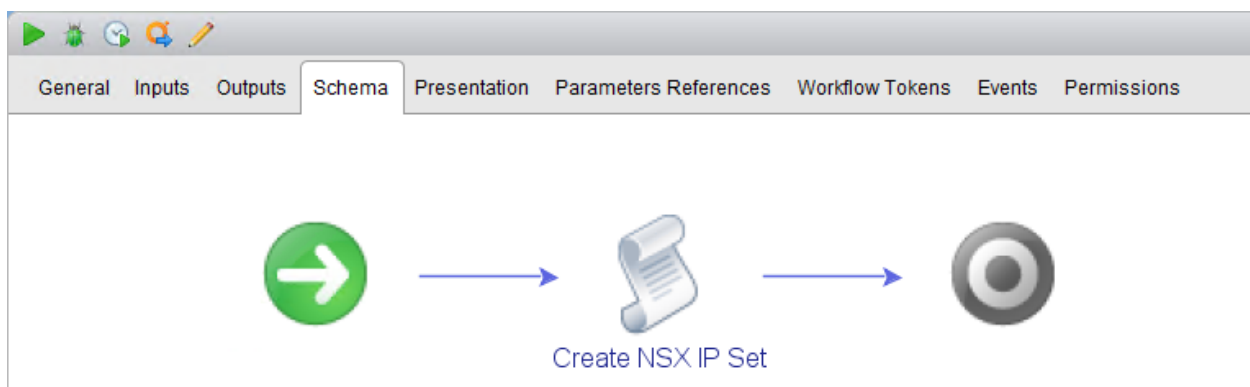


Figure 45: “Create NSX IP Set” vRO Workflow Schema

The scriptable task is bound to all the workflow inputs and outputs so that all variables are available in the JavaScript code. The visual binding of the task is shown in the below figure.

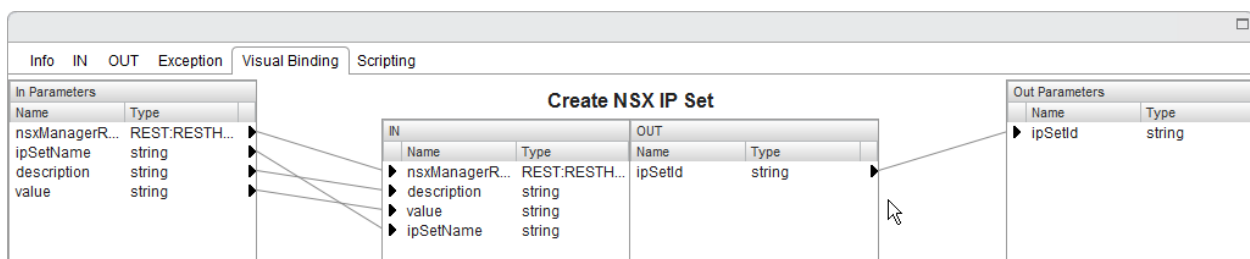


Figure 46: “Create NSX IP Set” vRO Workflow’s Scriptable Task Visual Binding

The following JavaScript code is used in the **Scripting** section of the task. Please refer to the comments in the code to understand the logic/flow better.

```
1. // vRealize Orchestrator script: Create NSX IP Set
2. // This script uses the REST Plugin for vRO to create an IP Set in NSX
3. // Inputs are:
```



```

4. // - nsxManagerRestHost (type REST:RESTHost): specifies the NSX Endpoint to be used as
   REST API target
5. // - ipSetName (type string): specifies the name of the IP Set to be created
6. // - description (type string): specifies the description of the IP Set to be created
7. // - value (type string): specifies the value (comma separate IP list) of the IP Set to
   be created
8. // Outputs:
9. // - ipSetId (type string): NSX ID of the created IP set
10.
11.
12. // XML Schema for IP Set creation (from API guide) is:
13. //
14. //<ipset>
15. //<objectId />
16. //<type>
17. //<typeName />
18. //</type>
19. //<description>(IP Set description)</description>
20. //<name>(IP Set name)</name>
21. //<revision>0</revision>
22. //<objectTypeName />
23. //<value>(comma separate list of IPs)</value>
24. //</ipset>
25. //
26.
27. // Prepare the HTTP body as an XML document with the required parameters
28. var xmlbody = new XML('<ipset />');
29. xmlbody.objectId = "";
30. xmlbody.type.typeName = "";
31. xmlbody.description = description;
32. xmlbody.name = ipSetName;
33. xmlbody.revision = 0;
34. xmlbody.objectTypeName = "";
35. xmlbody.value = value;
36. // Uncomment the next line to dump the generated XML to the debug
37. //System.debug("Generated XML is: " + xmlbody);
38.
39.
40. // HTTP POST request is prepared using the URL: /2.0/services/ipset/scopeId
41. // Note that "/api" is not shown in the request URL as it's already set in the endpoint
   configuration
42. // scopeId is globalroot-0
43. var request = nsxManagerRestHost.createRequest("POST", "/2.0/services/ipset/globalroot-
   0", xmlbody.toString());
44. request.contentType = "application/xml";
45.
46.
47. // Execute the HTTP request
48. System.debug("Creating IP Set " + ipSetName);
49. System.debug("POST Request URL: " + request.fullUrl);
50. var response = request.execute();
51.
52. // Evaluate the HTTP response
53. // We expect a HTTP 201 (Created), otherwise throw an exception
54. System.debug("GET Response Status Code: " + response.statusCode);
55. if (response.statusCode == 201) {
56.     System.debug("Successfully created IP Set " + ipSetName);
57. }
58. else {
59.     throw("Failed to create IP Set " + ipSetName);
60. }

```

```

61.
62. // Return created IP Set id (it's the last element in a returned URI in the "Location"
    HTTP header
63. // (it is also contained in the body of the response)
64. ipSetId = response.getAllHeaders().get("Location").split('/').pop();
65. System.debug("Created IP Set ID is " + ipSetId);

```

Figure 47: Javascript Code for vRO Script

The way the XML schema is created follows the ECMAScript for XML (E4X), which is an extension that adds native XML to JavaScript. This simplifies how XML information can be prepared and parsed to better interact with RESTful APIs.

Users can also choose to build the HTTP body by simply leveraging JavaScript strings, as in the below code example (that achieves the same result as in the previous example, but without using the E4X syntax to prepare the XML document) – note that the string is now used directly in the **createRequest** function, without the **.toString()** conversion.

```

1. // Prepare the HTTP body as a string with the required parameters
2. stringbody='<ipset><objectId /><type><typeName /></type><description>' + description +
    '</description><name>' + ipSetName + '</name><revision>0</revision><objectTypeName /><v
    alue>' + value + '</value></ipset>';
3. // Uncomment the next line to dump the generated string to the debug
4. //System.debug("Generated string is: " + stringbody);
5.
6. // HTTP POST request is prepared using the URL: /api/2.0/services/ipset/scopeId
7. // Note that "/api" is not shown in the request URL as it's already set in the endpoint
    configuration
8. // scopeId is globalroot-0
9. var request = nsxManagerRestHost.createRequest("POST", "/2.0/services/ipset/globalroot-
    0", stringbody);
10. request.contentType = "application/xml";

```

Figure 48: Alternative Usage of Javascript to Build HTTP Body

As an alternative, vRO also provides an XML library to implement the document structure if required.

Now that the workflow is created, users can run it in vRealize Orchestrator to create a new IP Set in NSX, specifying the parameter strings and the NSX Manager endpoint (selected from the available HTTP-REST object list in the inventory). A screenshot displaying this is shown below in Figure 49.

Figure 49: Run "Create NSX IP Set" vRO Workflow

Once the workflow run is completed successfully, it's possible to verify in the vRealize Orchestrator debug logs that the IP Set was successfully created (HTTP Post returned **201 – Created**). In the example shown in the following figure, the new object ID is **ipset-6**; this handler can be used for subsequent CRUD operations.



Figure 50: "Create NSX IP Set" vRO Workflow Logs

Using the NSX GUI, it's possible to verify that the new IP Set is created and it contains the expected addresses; this is validated below in Figure 51.

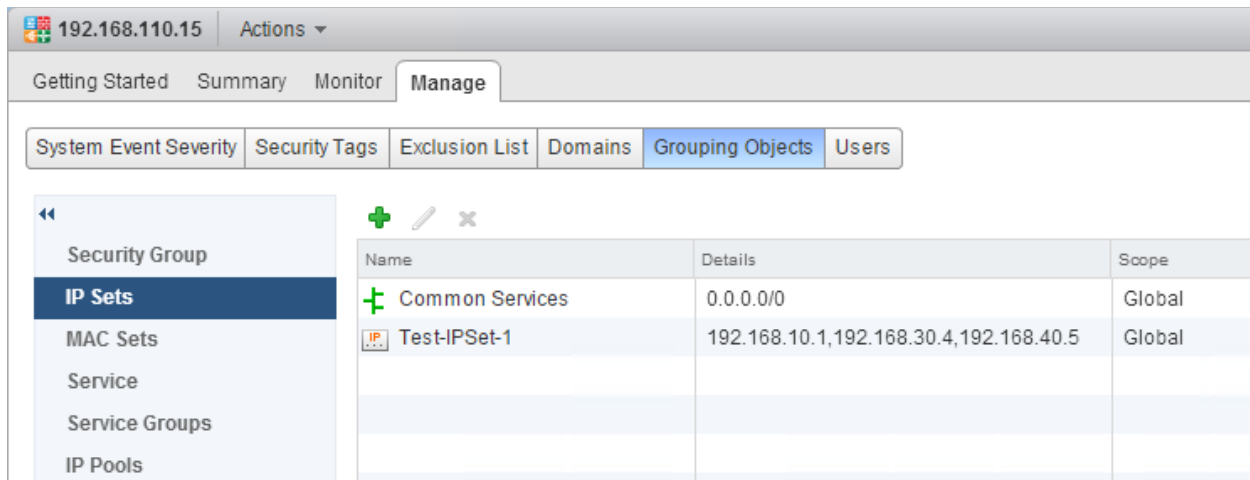


Figure 51: IP Sets as They Appear in NSX GUI

This example demonstrated how to create vRealize Orchestrator workflows that leverage the NSX REST API via the HTTP-REST plugin.

Example: Enable or Disable HA on an NSX Edge Using vRO

NSX Edges can be configured to run in Active/Standby high availability (HA) mode; the feature can be enabled at deployment time or later. The vRealize Orchestrator plugin for NSX does not implement such methods, but, similar to the previous example, in this example, it will be implemented using the HTTP-REST module.

To avoid redundancy, this example omits the general initial details presented in the previous section, including the configuration of an HTTP-REST endpoint for the NSX Manager.

This workflow requires the user to specify the NSX Manager HTTP-REST endpoint, the Edge object ID as an input string, and a boolean that determines whether HA must be enabled or not. The difference from the previous example is that now the user prompts for a desired state (i.e. HA enabled or disabled), so we first need to query the current configuration to determine whether an action is needed.

The workflow contains a single scriptable task, which includes JavaScript code which performs the following:

- Queries the NSX API to retrieve the HA configuration of a specific Edge
- Parses the returned XML document to discover whether HA is currently enabled
- Determines if HA is already in the desired state (if so, nothing to do)

- If HA is not in the desired state, modifies the XML document to change configuration and sends an HTTP POST request to the NSX Manager with the updated body. Note that only a single field of the original XML document is modified, while all other HA parameters are preserved
- Checks the return code to determine if the operation was successful. The workflow does not return any output; an exception will be thrown in case of error

The NSX API documentation, under the section **Working with High Availability (HA)**, specifies that HA configuration must be managed through the following URL:

`https://nsxmgr-ip/api/4.0/edges/edgeId/highavailability/config`

The same URL is used for retrieving the existing configuration (via HTTP GET) or modifying it (via HTTP PUT).

Another difference from the previous example is that now an Edge object ID is required; this information is specified in the URL with the parameter **edgeId**. The workflow must complete the URL with the ID provided by the user as an input parameter.

The NSX API documentation also provides a detailed description of the XML document that is returned by the HTTP GET request as shown below.

```

1. <highAvailability>
2. <vnic>(vNic index for HA)</vnic>
3. <ipAddresses>
4. <ipAddress>(First peer IP address)</ipAddress>
5. <ipAddress>(Second peer IP address)</ipAddress>
6. </ipAddresses>
7. <declareDeadTime>(Dead timer in seconds)</declareDeadTime>
8. <enabled>(enabled - true or false)<enabled>
9. </highAvailability>

```

Figure 52: XML Schema

For this specific scenario we just want to enable or disable HA, so we ignore all the parameters except the **<enabled> field**.

As this workflow has no outputs and the schema and scriptable task visual bindings are similar to the previous example, only the inputs are shown in the below figure.

General Inputs Outputs Schema Presentation Parameters References Workflow Tokens Events Permissions			
Parameters			
#	Name	Type	Description
➡	nsxManagerRestHost	REST:RESTHost	NSX endpoint for REST API
➡	edgeID	string	Edge ID
➡	enabled	boolean	Enable or Disable HA

Figure 53: "Configure Edge HA" vRO Workflow Inputs

The following JavaScript code is used in the **Scripting** section of the task. Please refer to the comments in the code to understand the logic/flow better.

```

1. // vRealize Orchestrator script: Configure HA on NSX Edge
2. // This script uses the REST Plugin for vRO to enable HA on an NSX Edge
3. // It first retrieves the configuration of the NSX Edge and evaluates whether it's required to enable or disable HA
4. // If a change is required, perform the action
5. // Inputs are:
6. // - nsxManagerRestHost (type REST:RESTHost): specifies the NSX Endpoint to be used as REST API target
7. // - edgeID (type string): specifies the ID of the Edge (from NSX Manager)
8. // - enabled (type bool): specifies whether Edge HA should be enabled or disabled
9. // No outputs
10.
11. // Retrieve the existing configuration
12. // HTTP GET request is prepared using the URL: /api/4.0/edges/edge-id/highavailability/config
13. // Note that "/api" is not shown in the request URL as it's already set in the endpoint configuration
14. var request = nsxManagerRestHost.createRequest("GET", "/4.0/edges/" + edgeID + "/highavailability/config");
15. request.contentType = "application/xml";
16.
17. // Execute the HTTP request
18. System.debug("Querying current HA configuration for Edge " + edgeID);
19. System.debug("GET Request URL: " + request.fullUrl);
20. var response = request.execute();
21.
22. // Evaluate the HTTP response
23. // We expect a HTTP 200 (OK), otherwise throw an exception
24. System.debug("GET Response Status Code: " + response.statusCode);
25. if (response.statusCode == 200) {
26.     System.debug("Response is success (200)!");
27.     // We can optionally output in the debug the entire XML document output (uncomment the next line)
28.     //System.debug("Response content is " + response.contentAsString);
29. }
30. else {
31.     throw("Failed to get HA configuration for Edge " + edgeID);
32. }
33.
34. // We were able to retrieve the Edge configuration, let's now parse the returned XML

```

```

35. // Use the XML parser to parse the HTTP XML in the response
36. var document = XML(response.contentAsString);
37. // Uncomment the next line to dump the parsed XML to the debug
38. //System.debug("Edge HA configuration is " + document.toString());
39.
40.
41.
42. // Determine if we need to do change something or if the HA state is already the desired one
43. // Initialize a boolean variable to false: if we need to do something will change to true later
44. var dosomething = false;
45. if(enabled) { // HA desired state is enabled
46.     if(document.enabled == "true") { // HA already enabled, nothing to do!
47.         System.debug("HA is already enabled on Edge " + edgeID);
48.     }
49.     else { // HA was not enabled, we must do something
50.         // Directly modify the XML document with the syntax below
51.         document.enabled = "true";
52.         dosomething = true;
53.     }
54. }
55. else { // HA desired state is disabled
56.     if(document.enabled == "true") { // HA was enabled, we must do something.
57.         // Directly modify the XML document with the syntax below
58.         document.enabled = "false";
59.         dosomething = true;
60.     }
61.     else { // is already disabled, nothing to do!
62.         System.debug("HA is already not enabled on Edge " + edgeID);
63.     }
64. }
65.
66. // It's been determined that we need to do something, we need to submit the change to the NSX API
67. if(dosomething) {
68.     // Uncomment the next line to dump the modified XML to the debug
69.     //System.debug("Modified HA configuration is " + document.toString());
70.     // Convert the XML to a string that we will send in the HTTP body
71.     xmlbody=document.toString();
72.     // Uncomment the next line to dump the HTTP body we are about to send
73.     //System.debug("HTTP body being sent: " + xmlbody);
74.     // HTTP PUT request is prepared using the URL: /api/4.0/edges/edge-id/highavailability/config with the required HTTP body
75.     // Note that "/api" is not shown in the request URL as it's already set in the endpoint configuration
76.     var request = nsxManagerRestHost.createRequest("PUT", "/4.0/edges/" + edgeID + "/highavailability/config",xmlbody);
77.     request.contentType = "application/xml";
78.     System.debug("Setting HA to " + enabled + " on Edge " + edgeID);
79.
80.     // Execute the HTTP request
81.     System.debug("PUT Request URL: " + request.fullUrl);
82.     var response = request.execute();
83.
84.     // Evaluate the HTTP response
85.     // We expect a HTTP 204 (No Content), otherwise throw an exception
86.     System.debug("PUT Response Status Code: " + response.statusCode);
87.     if (response.statusCode == 204) { // success!
88.         System.debug("HA successfully set to " + enabled + " on Edge " + edgeID);
89.     }

```

```

90.     else {
91.         throw("Failed to set HA to " + enabled + " on Edge " + edgeID);
92.     }
93. }

```

Figure 54: Javascript code for vRO Workflow

Now that the workflow is created, users can run it in vRealize Orchestrator, specifying the Edge object ID and its desired HA state; a screenshot is shown below in Figure 55.

Figure 55: Run "Configure Edge HA" vRO Workflow

Once the workflow run is completed successfully, it's possible to verify in the vRealize Orchestrator debug logs that the configuration is changed as shown in Figure 56 below.

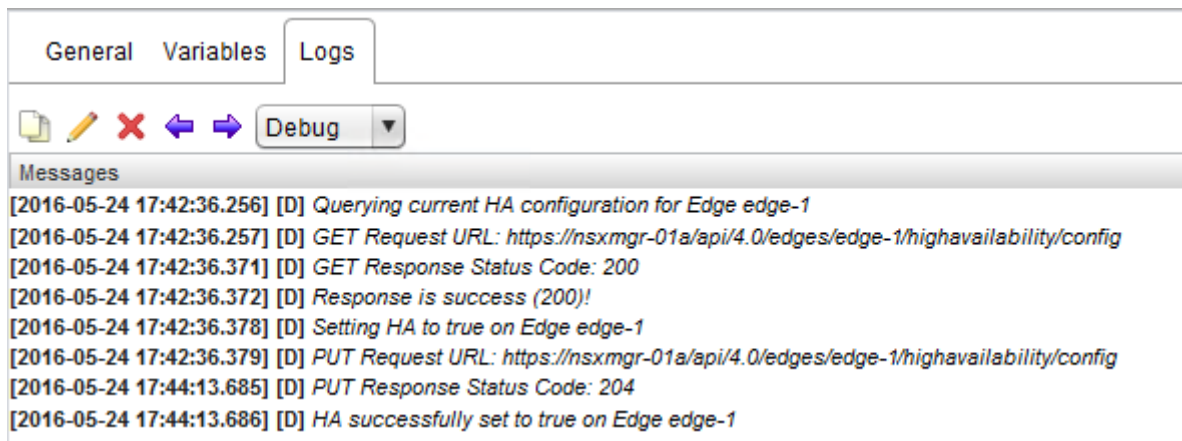


Figure 56: "Configure Edge HA" vRO Workflow Logs When an Action is Performed

On the NSX UI, it can be verified that the HA for **edge-1** is now enabled as shown below.

HA Configuration:		Change
HA Status:	Enabled	
Connected To:	vds-mgt_Management Network	

Figure 57: NSX Edge HA

As HA is already enabled, running the workflow again with the same parameters will return the below messages in the debug log and no action is performed.

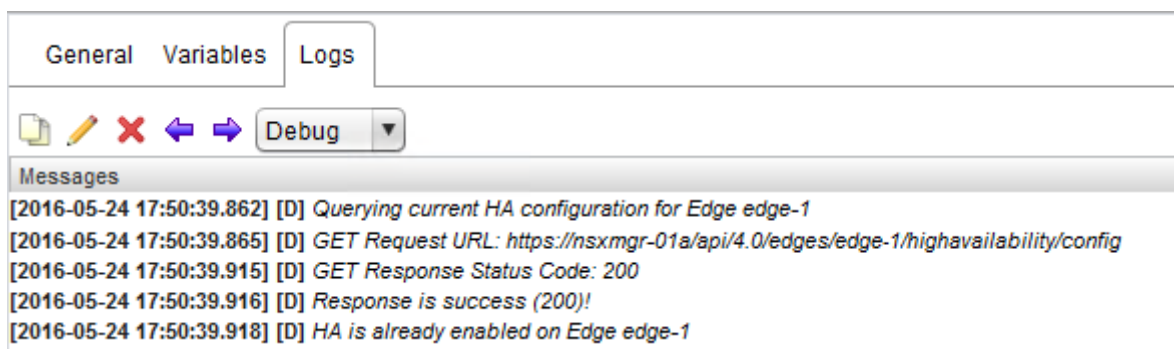


Figure 58:: "Configure Edge HA" vRO Workflow Logs When No Action is Performed

This example again demonstrated how to create vRealize Orchestrator workflows that leverage the NSX REST API via the HTTP-REST plugin.

Ansible

Together with Puppet, Chef and Salt, Ansible is a popular configuration management tool; it can be used to describe the state of a system and provides tools to manage its lifecycle. It can also be used for provisioning new systems; this operation can really be seen as a (complex) change of state. Ansible can also provide and orchestrate these functions on a vast number of systems at the same time.

Given that state management is at the core of Ansible, idempotency is a critical characteristic: if the current state and the desired state are the same, no actions are carried out. So running twice the same Ansible playbook on the same system will result in the second run not producing any change.

To end on a light note, the name Ansible is taken from science fiction where it stands for a device that can transfer information with a speed greater than lightspeed.

Example: Using Ansible Playbook to Manage Logical Switch State

The goal of this example is to show how an Ansible playbook could be used to manage the state of a logical switch; it will be done leveraging the work done with RAML (<http://raml.org/>).

As always, the first thing to do is check the start state of the testing environment, specifically the number of NSX Logical Switches already available. We use Postman and leverage a collection automatically generated from the RAML file defining the NSX-v API.

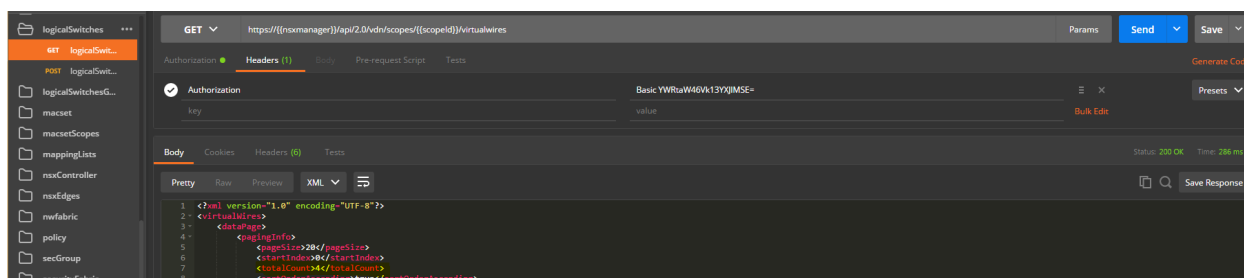


Figure 59: Using Postman to Verify the Number of Logical Switches Present

There are a total of four logical switches, which is confirmed by looking at the vSphere web client as shown in Figure 60 below.

NSX Manager: 192.168.110.15













Name	Status	Transport Zone	Scope	Segment ID	Control Plane Mode	Description	Tenant
 App_Tier_01	 Normal	 Local-Transport-Zone-A	Global	5002	Unicast		virtual wire tenant
 DB_Tier_01	 Normal	 Local-Transport-Zone-A	Global	5003	Unicast		virtual wire tenant
 Edge_Transit_01	 Normal	 Local-Transport-Zone-A	Global	5000	Unicast		virtual wire tenant
 Web_Tier_01	 Normal	 Local-Transport-Zone-A	Global	5001	Unicast		virtual wire tenant

Figure 60: Using NSX Manager GUI to Verify the Number of Logical Switches Present

Let's now have a look at the Ansible playbook that will be used to create a new NSX logical switch: **create_logicalswitch.yml**

```
1. -emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ more create_logicalswitch.yml
2. ---
3. - hosts: 127.0.0.1
4.   connection: local
5.   gather_facts: False
6.   vars_files:
7.     - answerfile.yml
8.   tasks:
9.     - name: nsx_logical_switch Operation
10.      nsx_logical_switch:
11.        nsxmanager_spec: "{ nsxmanager_spec }"
```

```

12.     state: present
13.     transportzone: "Local-Transport-Zone-A"
14.     name: "Ansible-LS"
15.     controlplanemode: "UNICAST_MODE"
16.     description: "Ansible Logical Switch"
17.     register: create_logical_switch

```

Figure 61: Ansible Playbook (create_logicswitch.yml)

The first lines specify that this playbook will be run on the same machine where Ansible is installed, that facts gathering will not be performed, and that some parameters needed by the playbook are stored in the file named **answerfile.yml** (see below). Note that the path of the file specifying the RAML format of the NSX-v API is explicitly called out, together with the NSX Manager login credentials

```

1.  emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ more answerfile.yml
2.  nsxmanager_spec:
3.    raml_file: '/home/emazza/Documents/AnsiblePlaybook/nsxraml/nsxvapi.raml'
4.    host: '10.152.65.157:21443'
5.    user: 'admin'
6.    password: 'VMware1!'

```

Figure 62: Parameter Specification (answerfile.yml)

Looking again at the Ansible playbook, it's important to note the following line

```

1.  nsx_logical_switch:
2.    nsxmanager_spec: "{{ nsxmanager_spec }}"
3.    state: present
4.    transportzone: "Local-Transport-Zone-A"
5.    name: "Ansible-LS"

```

Figure 63: "state: present" Line in Ansible Playbook

The code in Figure 63 highlights the specification of the desired state the system must be at the end of playbook run: a logical switch named **Ansible-LS** must be present in the NSX transport zone named **Local-Transport-Zone-A**. This is linked with the well-known idempotent characteristic of Ansible: it compares the current and desired state of the system and only executes changes when needed. So when the playbook will be called twice in a row, the second run will not produce any change to the system as the starting state already match the desired final state.

Below it is shown how this works at a high level. The playbook calls the function **nsx_logical_switch.py** (line 10 in Figure 61); this script is written in Python and stored under the library folder of the machine where Ansible is installed.

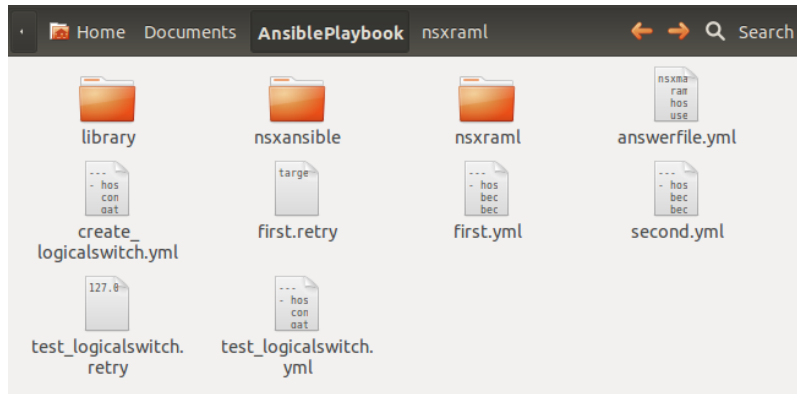


Figure 64: File Organization

Looking at the code, we can see the definition of the parameters that have been passed via the Ansible playbook (line 02 to line 11). At line 15 the **NsxClient** class is made available via importing the respective module via the **client.py** script; **NsxClient** is used to create a NSX client session object (line 15) which will be used to interact with the system programmatically.

```

1. def main():
2.     module = AnsibleModule(
3.         argument_spec=dict(
4.             state=dict(default='present', choices=['present', 'absent']),
5.             nsxmanager_spec=dict(required=True, no_log=True),
6.             name=dict(required=True),
7.             description=dict(),
8.             transportzone=dict(required=True),
9.             controlplanemode=dict(default='UNICAST_MODE', choices=['UNICAST_MODE',
10. 'MULTICAST_MODE', 'HYBRID_MODE'])
11.         ),
12.         supports_check_mode=False
13.     )
14.
15.     from nsxramlclient.client import NsxClient
16.     client_session=NsxClient(module.params['nsxmanager_spec']['raml_file'],
17. module.params['nsxmanager_spec']['host'], module.params['nsxmanager_spec']['user'],
18. module.params['nsxmanager_spec']['password'])

```

Figure 65: nsx_logical_switch.py (snippet)

Looking now at the rest of the code, the further below steps can clearly be identified

```

1. from nsxramlclient.client import NsxClient
2. client_session=NsxClient(module.params['nsxmanager_spec']['raml_file'], module.params['
nsxmanager_spec']['host'], module.params['nsxmanager_spec']['user'],
3. module.params['nsxmanager_spec']['password'])
4.
5. vdn_scope=retrieve_scope(client_session, module.params['transportzone'])
6. lswitch_id=get_lswitch_id(client_session, module.params['name'], vdn_scope)
7.
8. if len(lswitch_id) is 0 and 'present' in module.params['state']:

```

```

9.     ls_ops_response=create_lswitch(client_session, module.params['name'],
10.     module.params['description'], module.params['controlplanemode'], vdn_scope)
11.
12.     module.exit_json(changed=True, argument_spec=module.params,
13.     ls_ops_response=ls_ops_response)

```

Figure 66: nsx_logical_switch.py (snippet)

- Line 05: retrieval of the vdn scope corresponding to the transport zone passed as input
- Line 06: retrieval of the id of the logical switch belonging to the transport zone with name equal to the parameter **name**: defined in the playbook (**Ansible-LS** in this example)
- Line 08 to Line 13: if there is no such logical switch and the requested state is **present**, then the **create_lswitch** function will be called

Below the **create_lswitch** function that is defined in the same **nsx_logical_switch.py** file is looked at in more detail. This function basically creates a python dictionary representing a logical switch, fills it with the appropriate data, and passes it to the create method of the NSX client session object previously defined.

```

1. def create_lswitch(session, lswitchname, lswitchdesc, lswitchcpmode, scope):
2.     lswitch_create_dict = session.extract_resource_body_schema('logicalSwitches',
3.     'create')
4.     lswitch_create_dict['virtualWireCreateSpec']['controlPlaneMode'] = lswitchcpmode
5.     lswitch_create_dict['virtualWireCreateSpec']['name'] = lswitchname
6.     lswitch_create_dict['virtualWireCreateSpec']['description'] = lswitchdesc
7.     lswitch_create_dict['virtualWireCreateSpec']['tenantId'] = 'Unused'
8.     return session.create('logicalSwitches', uri_parameters={'scopeId': scope},
9.     request_body_dict=lswitch_create_dict)

```

Figure 67: "create_lswitch" function

The below screenshot shows the respective Ansible playbook run output.

```

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$
emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ ansible-playbook create_logicswitch.yml

PLAY *****

TASK [nsx_logical_switch Operation] *****
changed: [127.0.0.1]

PLAY RECAP *****
127.0.0.1          : ok=1    changed=1    unreachable=0    failed=0

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$

```

Figure 68: Ansible Playbook Run Output

The output is showing that the operation completed correctly and that one change to the state of the system has been made (remember that this logical switch did not exist before). Another look at Postman output will confirm this: there are now five logical switches, one of which has been recently created by the Ansible playbook

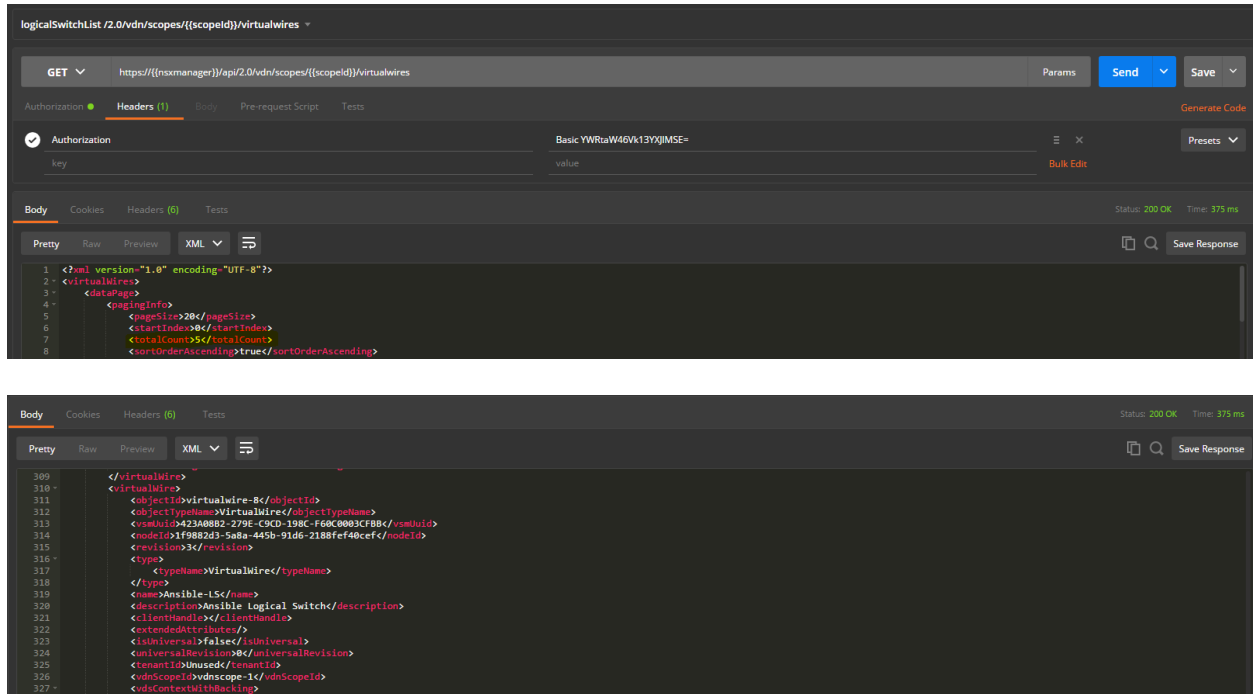


Figure 69: Verification via Postman NSX REST API Call Output

NSX Manager: 192.168.110.15							
Name	1	Status	Transport Zone	Scope	Segment ID	Control Plane Mode	Description
Ansible-LS		Normal	Local-Transport-Zone-A	Global	5004	Unicast	Ansible Logical Switch
App_Tier_01		Normal	Local-Transport-Zone-A	Global	5002	Unicast	
DB_Tier_01		Normal	Local-Transport-Zone-A	Global	5003	Unicast	
Edge_Transit_01		Normal	Local-Transport-Zone-A	Global	5000	Unicast	
Web_Tier_01		Normal	Local-Transport-Zone-A	Global	5001	Unicast	

Figure 70: Verification of Ansible Playbook's Run Results

Running the same playbook again will not cause any change in the system as the current state is already the final desired state. Figure 71 below shows this in practice: the playbook runs correctly (**ok=1**) but no changes have been made (**changed=0**).

```

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ ansible-playbook create_logicalswitch.yml

PLAY *****

TASK [nsx_logical_switch Operation] *****
ok: [127.0.0.1]

PLAY RECAP *****
127.0.0.1          : ok=1    changed=0    unreachable=0    failed=0

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$

```

Figure 71: Ansible Playbook Run Output

We can look at how this is managed at the code level in the **nsx_logical_switch.py** file. At line 07 the Ansible playbook is required to create a logical switch with the name of a logical switch that already exists. From line 08 to line 25 we cycle through the other attributes of this logical switch to understand if their values match the ones set in the playbook and mark accordingly a change request flag. Finally, from line 26 to line 32 we deal with the appropriate course of action based on the final value of the change request flag.

```

1.  if len(lswitch_id) is 0 and 'present' in module.params['state']:
2.      ls_ops_response=create_lswitch(client_session, module.params['name'],
3.      module.params['description'], module.params['controlplanemode'], vdn_scope)
4.
5.      module.exit_json(changed=True, argument_spec=module.params,
6.      ls_ops_response=ls_ops_response)
7.  elif len(lswitch_id) is not 0 and 'present' in module.params['state']:
8.      lswitch_details=get_lswitch_details(client_session,lswitch_id[0])
9.      change_required=False
10.     for lswitch_detail_key, lswitch_detail_value in
11.         lswitch_details['virtualWire'].iteritems():
12.             if lswitch_detail_key == 'name'
13.                 and lswitch_detail_value != module.params['name']:
14.                     lswitch_details['virtualWire']['name']=
15.                     module.params['nsxmanager_spec']['name']
16.                     change_required=True
17.             elif lswitch_detail_key == 'description'
18.                 and lswitch_detail_value != module.params['description']:
19.                     lswitch_details['virtualWire']['description']=module.params['description']
20.                     change_required=True
21.             elif lswitch_detail_key == 'controlPlaneMode'
22.                 and lswitch_detail_value != module.params['controlplanemode']:
23.                     lswitch_details['virtualWire']['controlPlaneMode']=
24.                     module.params['controlplanemode']
25.                     change_required=True
26.     if change_required:
27.         ls_ops_response=change_lswitch_details(client_session,lswitch_id[0],
28.         lswitch_details)
29.         module.exit_json(changed=True, argument_spec=module.params,
30.         ls_ops_response=ls_ops_response)
31.     else:
32.         module.exit_json(changed=False, argument_spec=module.params)

```

Figure 72: Idempotency in nsx_logical_switch.py

Let's see this in action and request a final state which will result in a change of the description of the already existing logical switch.

```
emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ more modify_logicalswitch.yml
---
- hosts: 127.0.0.1
  connection: local
  gather_facts: False
  vars_files:
    - answerfile.yml
  tasks:
    - name: nsx_logical_switch Operation
      nsx_logical_switch:
        nsxmanager_spec: "{{ nsxmanager_spec }}"
        state: present
        transportzone: "Local-Transport-Zone-A"
        name: "Ansible-LS"
        controlplanemode: "UNICAST_MODE"
        description: "Rebranded Ansible Logical Switch"
      register: create_logical_switch
emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$
emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$
emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ ansible-playbook modify_logicalswitch.yml

PLAY *****

TASK [nsx_logical_switch Operation] *****
changed: [127.0.0.1]

PLAY RECAP *****
127.0.0.1 : ok=1 changed=1 unreachable=0 failed=0

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$
```

Figure 73: Ansible Playbook Run Output

NSX Manager: 192.168.110.15						
<div> <div>+</div> <div>✗</div> <div>✕</div> <div>🔍</div> <div>⚙️</div> <div>Actions</div> </div>						
Name	Status	Transport Zone	Scope	Segment ID	Control Plane Mode	Description
Ansible-LS	✓ Normal	Local-Transport-Zone-A	Global	5004	Unicast	Rebranded Ansible Logical Swit..
App_Tier_01	✓ Normal	Local-Transport-Zone-A	Global	5002	Unicast	

Figure 74: Checking name change using NSX Manager GUI

Finally, let's request a final state such that this logical switch must not exist anymore in the system.

```
1. emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ more destroy_logicalswitch.yml
2. ---
3. - hosts: 127.0.0.1
4.   connection: local
5.   gather_facts: False
6.   vars_files:
7.     - answerfile.yml
8.   tasks:
9.     - name: nsx_logical_switch Operation
10.       nsx_logical_switch:
```



```

11.     nsxmanager_spec: "{{ nsxmanager_spec }}"
12.     state: absent
13.     transportzone: "Local-Transport-Zone-A"
14.     name: "Ansible-LS"
15.     controlplanemode: "UNICAST_MODE"
16.     description: "Rebranded Ansible Logical Switch"
17.     register: create_logical_switch

```

Figure 75: Ansible Playbook (destroy_logicalswitch.yml)

```

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$ ansible-playbook destroy_logicalswitch.yml

PLAY *****

TASK [nsx_logical_switch Operation] *****
changed: [127.0.0.1]

PLAY RECAP *****
127.0.0.1      : ok=1    changed=1    unreachable=0    failed=0

emazza@emazza-ubuntu-vm:~/Documents/AnsiblePlaybook$

```

Figure 76: Ansible Playbook Run Output




Figure 77: Using Postman to Verify the Number of Logical Switches Present

NSX Manager: 192.168.110.15						
Name	Status	Transport Zone	Scope	Segment ID	Control Plane Mode	
App_Tier_01	✓ Normal	Local-Transport-Zone-A	Global	5002	Unicast	
DB_Tier_01	✓ Normal	Local-Transport-Zone-A	Global	5003	Unicast	
Edge_Transit_01	✓ Normal	Local-Transport-Zone-A	Global	5000	Unicast	
Web_Tier_01	✓ Normal	Local-Transport-Zone-A	Global	5001	Unicast	

Figure 78: Using NSX Manager GUI to Verify the Number of Logical Switches Present

Cloud Management Platforms with NSX REST API

Many organizations are building private cloud solutions to provide their users with tools that automatically deploy infrastructure services. Cloud architects have quickly realized that automating the provisioning and destruction of Virtual Machines (VMs) is not sufficient to satisfy business requirements: too often configuration changes in network

and security are also required to provide the desired service, and, with manual provisioning of such components, the benefits of the entire cloud initiative are limited.

For example, when manually deploying production environments, the process often follows a chain of events that involve operations by several teams (VI, network, security, storage, load balancing, applications, etc.). This model not only slows down the deployment, but is also prone to human errors that can lead longer wait times or misconfigurations in security policies thus resulting in insecure environments. A common example is the removal of firewall policies related to decommissioned applications, which is a rarely performed task.

Moreover, as organizations are transitioning towards Micro-segmentation, the requirement of granular security policies deployed alongside the applications is becoming important, which again is not easy to perform via manual operations. Lastly, many companies are developing applications in-house and their dev and test environments, while not requiring the same levels of availability and security of production, are critical to deliver new features and are therefore considered revenue generating. In such environments, network and security configurations (such as NAT, Load Balancer and Firewall) are often required and must be automated.

For the above reasons, modern IaaS clouds provide automated and integrated consumption of virtual machines, network constructs (such as Logical Switches and Load Balancers), security policies, middleware and applications. Depending on the scope of the initiative (production, dev, test/QA, etc.) and the organizational structure, private cloud solutions can offer different levels of governance and consumption models (i.e. whether giving users access to service catalogs or directly to CMP APIs, implementing approval workflows and strict resource management, etc.).

While describing the differences among several Cloud Management Platform (CMP) solutions is out of the scope of this document, it is important to mention that using NSX as the underlying network and security platform is extremely beneficial as all possible governance and consumption models can apply to it. With NSX it's possible to automate all the configuration actions by leveraging the NSX REST API: several Cloud Management Platforms (CMPs) already provide out-of-the-box integrations to consume the services provided by NSX, such as Logical Switching, Logical Routing, Distributed Firewall and Logical Load Balancing.

This section describes how one of the most widely used CMP solutions, VMware vRealize Automation, leverages the NSX API to automate the consumption of services.

Other solutions can already integrate with NSX out-of-the-box and it's important to understand that, since most CMPs are modular, it's possible to extend their capabilities to consume NSX services either via the REST API or through vRealize Orchestrator. The other examples provided in this document should be useful to understand how to implement such extensions using common programming languages.

vRealize Automation

VMware vRealize Automation (vRA) is a cloud automation solution that automates the delivery of custom infrastructure, application, and IT services across a multi-vendor hybrid cloud infrastructure. Details on vRA are available on [VMware's vRA web site](#), but, in short, it provides an automated way of deploying and managing applications through a service catalog and natively integrates with NSX.

The latest version of vRA at the time of writing is 7.0. It allows native consumption of both pre-created and on-demand NSX network and security services through the concept of a blueprint: cloud architects can design entire application stacks (that include VMs, pre-existing or on-demand networks, load balancers, security groups, policies and tags) and publish them in a service catalog. Cloud consumers can then deploy, manage and destroy applications by leveraging the service catalog, as shown in the picture below.

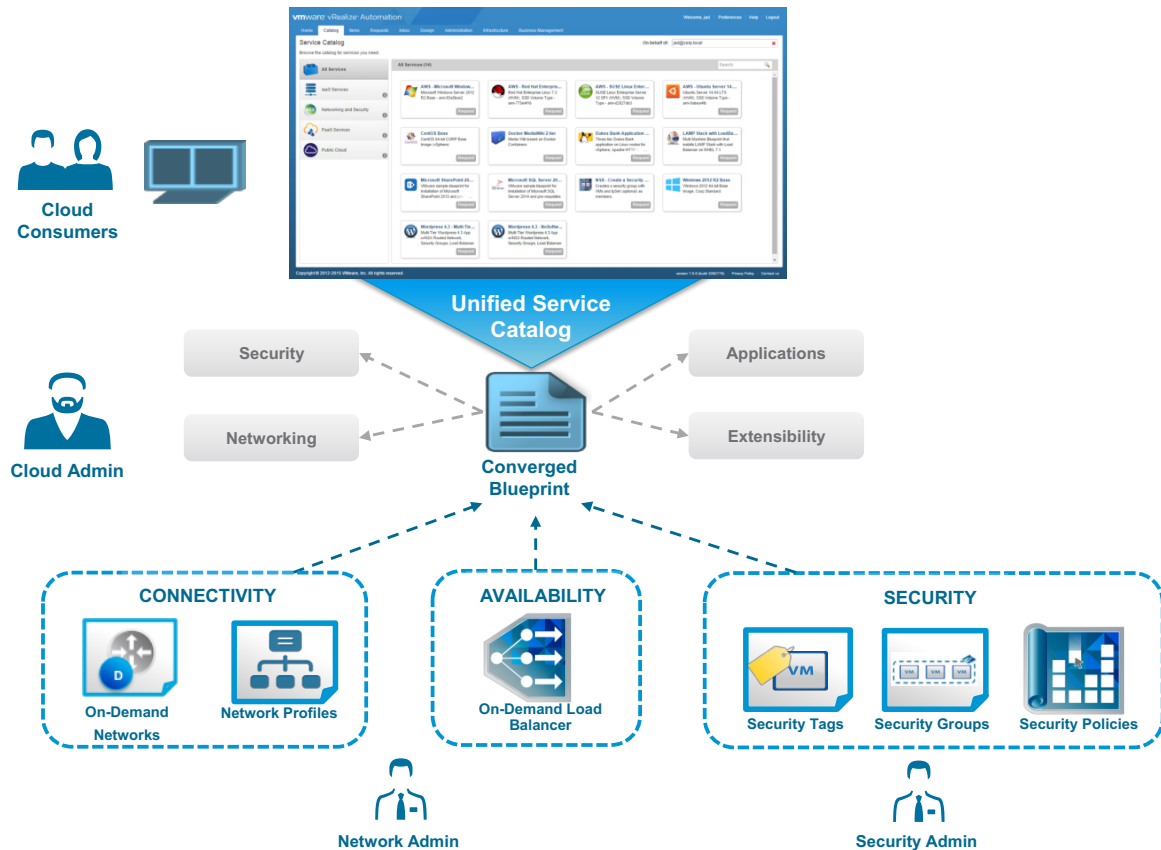


Figure 79: Cloud Consumption Model with vRA

This document is not meant to be a comprehensive guide on how vRA can consume NSX services, but the most significant aspects of the consumption layer from the network and security perspective are presented below.

In regards to Networking, vRA can natively consume the following services based on NSX:

- External networks: vRA can deploy VMs on to existing NSX Logical Switches. This is very useful for production applications that get deployed on pre-existing, shared networks.
- NAT networks: when overlapping addressing is required (very often for dev and test environments), vRA can create new NSX Logical Switches and NSX Edges and define NAT rules to allow reusability of address space across different deployments. It's possible to define 1:1 and 1:Many network profiles, depending on the required external access. Once the application is decommissioned, vRA removes the created network objects alongside VMs.
- Routed networks: some organizations require cloud users to create on-demand routed networks. As part of a blueprint, cloud architects can configure vRA to create new NSX Logical Switches and attach them to an existing NSX Distributed Logical Router during the provisioning process. vRA also manages the IP addressing accordingly to avoid overlapping. As in the NAT case, once the application is decommissioned, dynamically created NSX Logical Switches are also removed.
- Load balancing: vRA can create an NSX Edge at application deployment time and configure it as a load balancer for the required services. Both in-line and one-arm topologies are supported, depending on the network profile.
- DHCP: if required, vRA can configure DHCP services on an NSX Edge when it configures 1:Many NAT networks.

vRA can also natively consume NSX Service Composer capabilities, allowing security admins to define a security posture that is consumed by the cloud platform. vRA leverages the following NSX security capabilities:

- Pre-created NSX Security Groups: vRA can discover existing Security Groups that exist in NSX and add VMs to them during the provisioning of the application. Once the application owner deletes the VMs, they disappear from the Security Groups, reducing the risk of stale firewall rules.
- On-demand NSX Security Groups: vRA is also capable of creating new NSX Security Groups for each VM tier (i.e. web, app or DB), and attaching existing Security Policies to them. When users decommission the applications, the on-demand Security Groups are also removed by vRA.
- Pre-created NSX Security Tags: vRA can discover and attach existing NSX Security Tags to the VMs it creates, allowing a tag-based consumption of security.

- App Isolation: to simplify security automation, vRA also provides an option (App Isolation) that can be enabled at the blueprint level. When selected, vRA will automatically generate NSX Security Groups and Security Policies to block all communications across different application deployments, unless explicitly permitted by other rules. This is extremely useful to provide Micro-segmentation without the need of complex application scoping.

The below picture shows a graphical representation of two applications deployed by vRA that include security, load balancing and App Isolation with NSX.




Figure 80: vRealize Automation Topology Example with App Isolation

All the NSX capabilities that vRA can consume are configured using the Converged Blueprint Designer (CBD): it provides a canvas in the UI where users can drag and drop different components (VMs, networks, security groups, load balancers, etc.) and build the required application topology. Blueprints can also be exported in a YAML file that users can be modify and re-upload (blueprint-as-code).

The below picture shows an example of the Converged Blueprint designer, where a 2-tier application (Web and DB) that includes two on-demand routed networks, a load balancer and an on-demand security group for the web tier are being utilized.

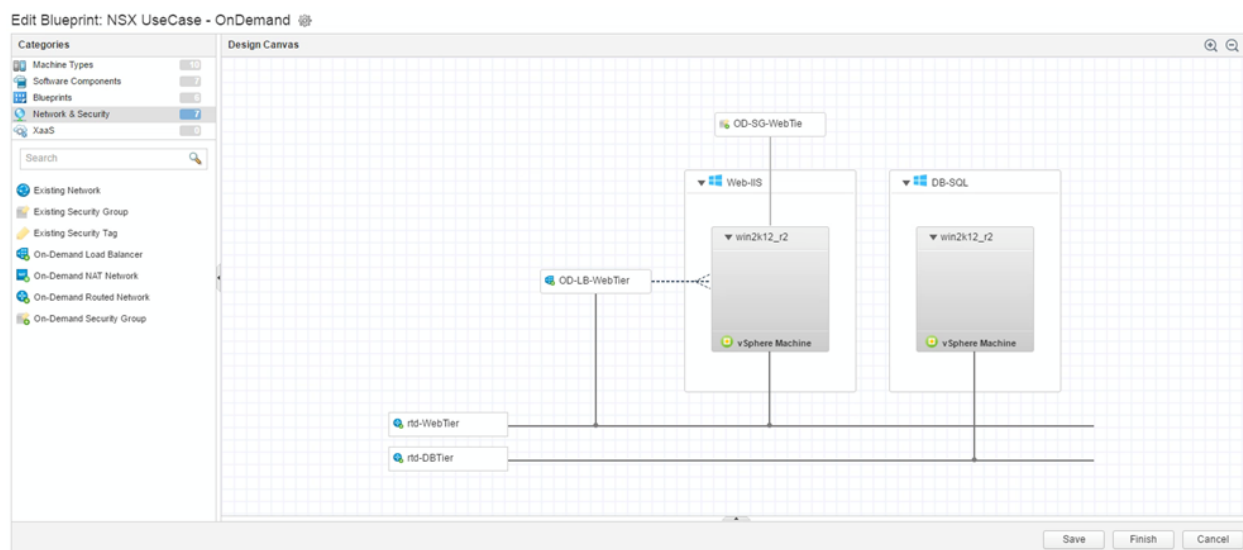


Figure 81: vRealize Automation Converged Blueprint Designer

As mentioned, vRealize Automation can leverage the NSX REST API to achieve the required configurations. To add some detail about the implementation, it is relevant to mention that the actual consumption model is implemented through vRealize Orchestrator (vRO), using the NSX vRO Plugin, as shown in the Figure 82 below.

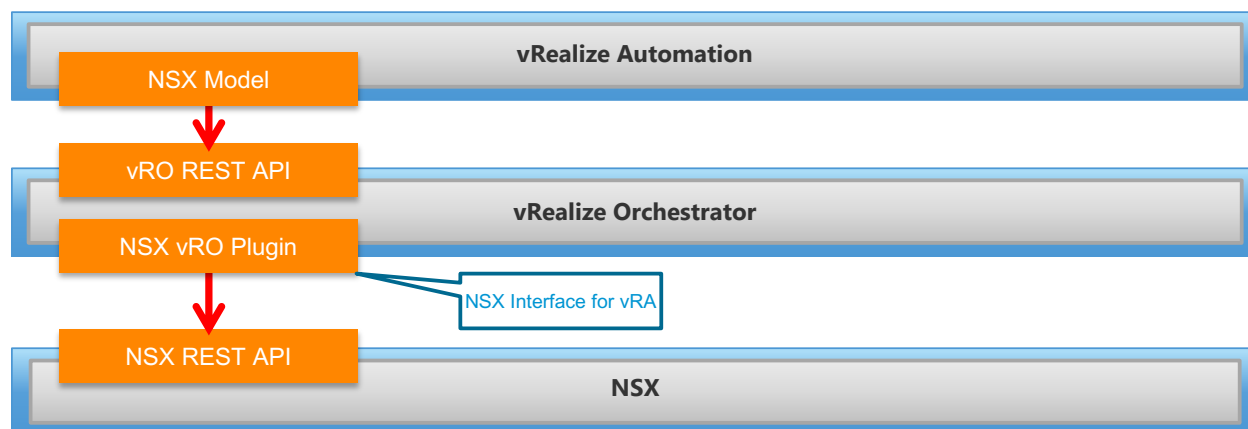


Figure 82: vRealize Automation Consumption of NSX Through vRealize Orchestrator

Among the several design reasons that drive this choice, the most relevant is the ability of extending the native consumption model. NSX is a platform that provides a comprehensive set of capabilities and not all of them are automated by the native vRealize Automation integration, as the out-of-the-box focus is set on the most common use cases. The ability to extend the default behavior and enable additional use cases is a great capability that can be leveraged through the different vRA extensibility options, that all leverage vRO as the underlying orchestration layer.

A vRO plugin for vRA allows for purpose-built custom automation of NSX objects.

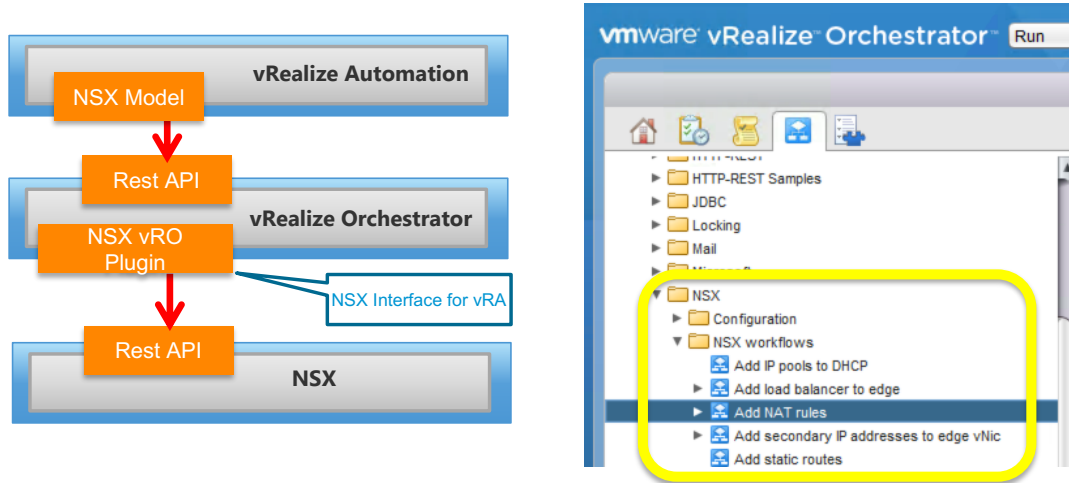


Figure 83: NSX and vRA – Extensibility using vRO

While more features are being added in the native integration, some examples on use cases that require extensibility and are simple to enable are the following:

- Change the size or the Syslog targets of an NSX Edge deployed by vRA
- Enable a persistence algorithm on an NSX Load Balancer configured by vRA
- Define Day 2 operations to add additional NAT rules for an application

These and many other use cases can be enabled by leveraging two vRA extensibility features:

- **EBS (Event Broker Service):** allows cloud administrators to subscribe vRO workflows to specific events, effectively modifying the behavior of the native integration. For example, one could subscribe a workflow that will change the size of the NSX Edge to the event that is generated upon deployment: this way, as soon as the end user triggers the deployment of an application, vRA will create an NSX Edge and call the workflow to change its size.
- **XaaS (anything as-a-service):** vRO workflows can be published in vRA and invoked directly from the service catalog by the end users, allowing direct consumption. XaaS can also be used to embed vRO workflows directly within a converged blueprint, allowing potentially unlimited extensibility use cases. For example, a Day 2 operation that adds additional NAT rules on an existing NSX Edge could be defined using XaaS and published in the service catalog, allowing end users to invoke it against a previously deployed application.

As both extensibility approaches leverage vRealize Orchestrator as the underlying layer, the same content (vRO workflows and actions) can be reused across both options.

Additional details about the NSX plugin for vRO and how it's possible to consume the NSX API via the REST plugin are provided in the section related to vRealize Orchestrator.

References

[NSX REST API](#)

[VMware vRealize Orchestrator Documentation](#)