vAPV Administration Guide

For ArrayOS APV 8.6.0.30 Release

August 3, 2016

1   Server Hardware Requirements

The server hosting the vAPV must meet the following minimal requirements:

- CPU: One or more 64-bit x86 multi-core CPUs with virtualization assist (Intel-VT or AMD-V) enabled
- RAM: 2 GB available for vAPV
- NIC: 1 x 1 Gbps NIC or 2 x 1 Gbps NICs
- Storage: 40 GB of available hard drive space for vAPV

2   Hypervisor Guest Machine Definition

The following hypervisors are supported for administration of the vAPV:

- VMware ESXi 4.1 or later (64-bit only)
- Citrix XenServer 5.6 or later (64-bit only)
- KVM in CentOS 6.0 or later
- Microsoft Hyper-V (64-bit only)

The VMware ESXi virtual machine guest environment for the vAPV must include:

- 2 vCPUs
- 2 GB vRAM
- 4 x 1 virtual e1000 Ethernet adapters (other Ethernet adapter types are not supported, for example, VMXNET)
- 40 GB of available hard drive space

The Citrix XenServer virtual machine guest environment for the vAPV must include:

- 2 vCPUs
- 2 GB vRAM
- 4 x 1 virtual network interfaces
- 40 GB of available hard drive space

The KVM virtual machine guest environment for the vAPV must include:

- 2 vCPUs
- 2 GB vRAM
- 4 x 1 virtual network interfaces
- 40 GB of available hard drive space

The Mircosoft Hyper-V virtual machine guest environment for the vAPV must include:

- 2 vCPUs
- 2 GB vRAM

- 4 x 1 virtual network interfaces
- 40 GB of available hard drive space

3 Deploying the vAPV Virtual Machine

3.1 Deploying the vAPV Virtual Machine on VMware ESXi

The first step in deploying the vAPV virtual machine on VMware ESXi is to download the vAPV OVA file to your local system. Next, run the "**Deploy OVF Template**" wizard from within the VMware vSphere client. Detailed steps are as follows:

1. Please contact Array Networks Customer Support (Email: support@arraynetworks.com) for download instructions.

2. Download the Array vAPV file package for VMware ESXi with the ".ova" extension name.

3. From the "**File**" menu in the vSphere client, select "**Deploy OVF Template**". (This will start the "**OVF Template**" wizard.)

4. In the "**Source**" pane, click "**Deploy from file**" and locate the OVA file using the "**Browse**" button. For example, "\MyDocuments\Work\Virtualization\<vAPV_OVA_filename>".

5. Click "**Next**" to open the "**OVF Template Details**" pane and verify that the OVF template details are correct.

6. Click "**Next**" to open the "**Name and Location**" pane.

7. In the "**Name**" field, type a name for the vAPV virtual machine, for example, "vAPV".

8. In the "**Disk Format**" pane, select "**Thick provisioned format**".

9. Click "**Next**" to open the "**Ready to Complete**" screen.

10. Verify that all deployment settings are correct and click "**Finish**".

11. Once the Virtual Machine is deployed, right-click the deployed Virtual Machine and select **Power** > **Power On** in the pop-up menu to start the vAPV Virtual Machine.

3.2 Deploying the vAPV Virtual Machine on Citrix XenServer

The first step in deploying the vAPV virtual machine on Citrix XenServer is to download the .rar containing the XVA file to your local system. Next, run the "**Import**" from within XenCenter. Detailed steps are as follows:

1. Please contact Array Networks Customer Support (Email: support@arraynetworks.com) for download instructions.

2. Download the Array vAPV .rar file containing the .xva image.

3. Open the .rar file by using the WinRAR archiver and extract the Array vAPV file package ending with .xva.

4. From the "**File**" menu in XenCenter, select "**Import**". (This will start the "**XVA import**" wizard.)

5. In the "**Import Source**" pane, locate the XVA file using the "**Browse**" button. For example, "\MyDocuments\Work\Virtualization\<vAPV_XVA_filename>".

6. Click "**Next**" to open the "**Home server**" pane, and select the home server on which the Virtual Machine should be run.

7. Click "**Next**" to open the "**Storage**" pane, and select the storage repository to store the virtual disks for the new Virtual Machine.

8. Click "**Import >**" to prepare the importing process.

9. In the "**Network**" pane, add, modify or delete virtual network interfaces as required.

10. Click "**Next**" to open the "**Finish**" screen.

11. Select the check box at the bottom of the "**Finish**" pane to start the vAPV Virtual Machine as soon as the import process is over.

12. The vAPV Virtual Machine will be deployed and started automatically.

3.3 Deploying the vAPV Virtual Machine on KVM

The first step in deploying the vAPV virtual machine on KVM is to download the compressed file containing the bootable hard disk of the Virtual Machine and the XML formatted configuration file. Next create a Virtual Machine by using these files. Detailed steps are as follows:

1. Please contact Array Networks Customer Support (Email: support@arraynetworks.com) for download instructions.

2. Download the compressed file (for example, vAPV_image.tar.gz) containing the disk0.raw disk image file and the XML configuration file (for example, vAPV.xml). Please make sure that you have more than 40 GB of free space after the download.

3. Open the compressed file by using the "**tar -xvzf <vAPV_image>.tar.gz**" command. This creates a folder called "vAPV_image" which contains the disk0.raw disk image file and the XML configuration file.

4. Enter the folder "vAPV_image", and edit the XML configuration file to specify the appropriate bridge name that you want the interfaces to be connected to. Specify the disk, CPU, and memory in the vAPV configuration file if required. Please do not change any other fields as it might lead to licensing conflicts.

5. Create a virtual machine from the XML configuration file by using the "**virsh define vAPV.xml**" command.

6. Start the virtual machine by using the "**virsh start vAPV**" command.

Note: The deployment steps of vAPV on KVM are based on the CentOS 7.0 environment. The deployment steps on other Linux environments are similar. For any questions on the deployment, please contact Array Networks Customer Support (Email: support@arraynetworks.com).

3.4 Deploying the vAPV Virtual Machine on Microsoft Hyper-V

The first step in deploying the vAPV virtual machine on Microsoft Hyper-V is to download the .vhd file to your local system installing Microsoft Hyper-V. Detailed steps for deploying vhd file are as follows:

1. Please contact Array Networks Customer Support (Email: support@arraynetworks.com) for download instructions.

2. Download the Array vAPV vhd image file.

3. Access the Microsoft Hyper-V server using the Microsoft Hyper-V Manager. In Hyper-V Manager, click **Action** > **New** > **Virtual Machine** to bring up the New Virtual Machine Wizard.

4. Click "**Next**" on the wizard and input a virtual machine name such as "**vAPV**".

5. In the "**Specify Generation**" pane, select "**Generation 1**" and click "**Next**".

6. In the "**Assign Memory**" pane, input the startup memory of vAPV and click "**Next**".

7. In the "**Configure Networking**" pane, select a network adapter to use as a virtual switch.

8. In the "**Connect Virtual Hard Disk**" pane, select "**use an existing virtual hard disk**", and enter the vAPV vhd image file location. Click "**Next**" when ready.

9. Verify that all deployment settings are correct and click "**Finish**".

10. Once the virtual machine is deployed, it can be started.

4　Post-Installation Configuration

1. Please finalize all Virtual Machine configurations by following the steps in Chapter 5 before requesting a new license.

   Caution: Any Virtual Machine configuration change may invalidate the current license.

2. Please contact the Array Networks Customer Support (Email: support@arraynetworks.com) to obtain a license key after you have finalized the configuration.

5　Array vAPV Series Installation Guide

Once the virtual machine starts, the login shell will appear on the console tab of the vAPV virtual machine, for example, on the Citrix XenCenter, VMware vSphere client, Virtual Machine Manager or Microsoft Hyper-V Manager. Please note that it may take a couple of minutes for the login shell to appear.

The initial username/password is array/admin.

6　Minimum Initial Configuration

Three ways are available to connect to the vAPV to begin configuration:

1. Console (recommended)

2. SSH

3. Web browser

To set up the vAPV via SSH or a Web browser, first you need to complete the network settings of the vAPV through the Console connection.

It is recommended to access the vAPV system via the virtual machine console tab on the Citrix XenCenter, VMware vSphere Client, Virtual Machine Manager or Microsoft Hyper-V Manager. The vAPV will prompt you for a User Privilege Password. A prompt like "Array Networks Login (AN):" may appear. If this is the first time you have connected to the vAPV, or if you have not changed the default password, then enter the user name "array" and the password "admin".

Now enter "**enable**" and press "**Enter**" to switch to the enable mode. You will be prompted to enter an enable password. The default enable password is null. Therefore, just press "**Enter**". The prompt "AN#" will appear.

Now enter "**config terminal**" and press "**Enter**" to switch to the configuration mode. The prompt "AN(config)#" will appear. At this point, the administrator will have full access to the vAPV CLI.

The CLI commands required for minimum configuration are listed below. It is recommended that you set port1, the default route, and WebUI IP address and port.

ip address *{system_ifname|mnet_ifname|vlan_ifname|bond_ifname} <ip_address> {netmask|prefix} [overlap]*

Allows the user to set each interface IP address and netmask or prefix length.

ip route default *<gateway_ip>*

Allows the user to set the default gateway IP address.

webui ip *<ip_address>*

Allows the user to set the IP address at which that the vAPV will accept Web User Interface commands via the Web browser. It is recommended that a management IP address be used for configuring the WebUI IP address.

webui port *<port>*

Allows the user to set the port from which the vAPV will accept WebUI commands. The port must be designated within the range from 1024 to 65,000. The default port is 8888.

In vAPV WebUI configuration, this command is only used for the new WebUI. For the legacy WebUI, please use "**webui legacy port** *<port>*" command.

webui {on|off}

Enables or disables the WebUI.

In vAPV WebUI configuration, this command is only used for the new WebUI. For the legacy WebUI, please use "**webui legacy {on|off}**" command.

With the above network settings completed for the vAPV, you may set up the vAPV via SSH or a Web browser. You may use a Web browser to connect to the WebUI IP address assigned to the vAPV (with WebUI enabled) or establish an SSH connection to the IP address of the vAPV.

Note: If you change the settings of the vAPV port1 IP address and WebUI IP address, the connection to the vAPV will be terminated and you will have to enter the new WebUI address in the browser to reconnect to the vAPV WebUI.

Configuration Example

```
AN(config)#ip address port1 10.8.6.50 255.255.255.0
AN(config)#ip route default 10.8.6.1
AN(config)#webui ip 10.8.6.50
AN(config)#webui port 8888
AN(config)#webui on
```

After completing these steps, administrators may continue to configure the vAPV via the CLI or WebUI. To access and configure vAPV via WebUI, enter the URL **https://10.8.6.50:8888** in the address bar of a browser. Then follow the prompts to log in the vAPV and enter the password for the "enable" mode. After you enter the "enable" mode, the following window will be displayed.
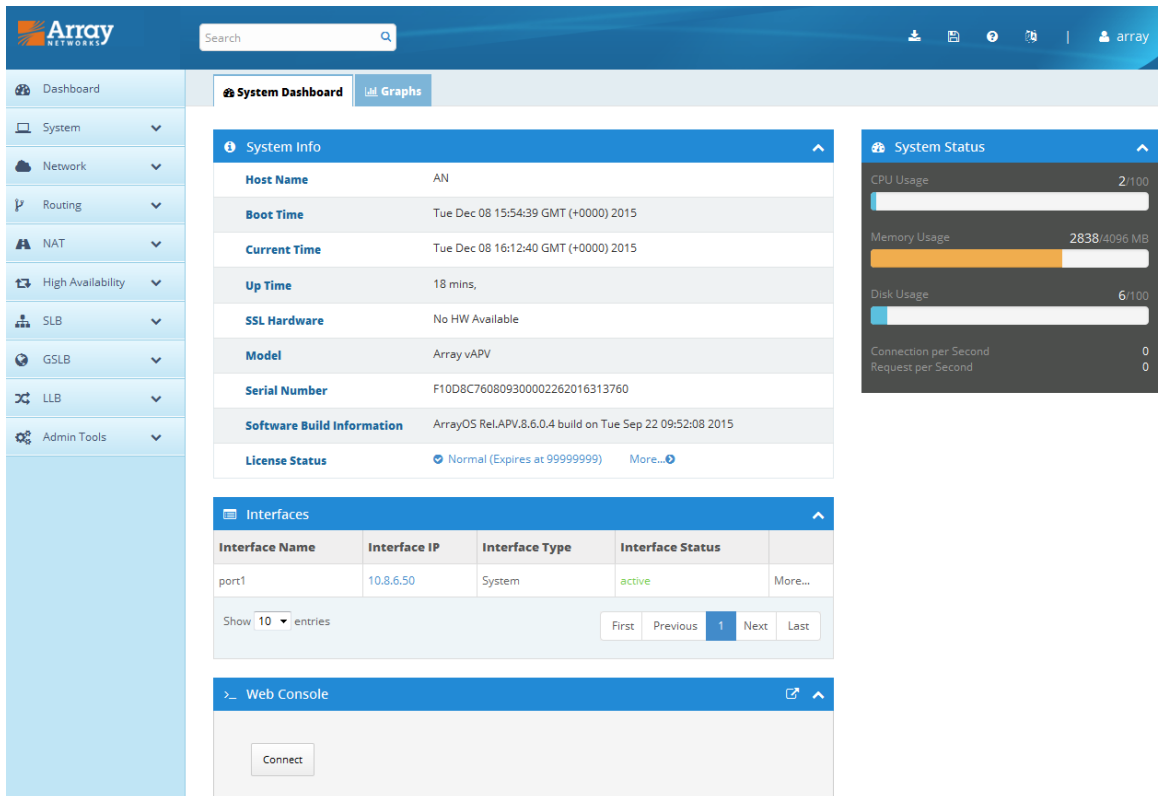
For more detailed configuration information, please refer to the User Guide and CLI Handbook of the APV series products.

7   Resizing a vAPV Virtual Machine on Hypervisors

Resizing the vAPV virtual machine such as adjusting the system resources vCPUs, vNICs, RAM and bandwidth, might cause the license to become invalid, and therefore administrators must be very cautious when resizing a virtual machine. Before any resizing operation, please check the resizing scenario and be aware of the potential effects caused by your operation, as shown in the following table:

Resizing Scenario	Influence
Downgrade the VM size	The license will still be valid.
Upgrade the VM size within the licensed limits	The license will still be valid.
Upgrade the VM size beyond the licensed limits	The license will become invalid. Please repurchase the vAPV license by providing the following information to Array Networks Support (Email: support@arraynetworks.com): • Serial number • New Capacity you wish to purchase

You can use "**show version**" command to check the "**Licensed Limits**" item which displays the system resources authorized by the license.

Note: Before resizing the vAPV virtual machine, please shut it down first.

8   Known Limitations

- Fast Failover (FFO) is not supported.
- Hardware compression SR-IOV is not supported.
- Hardware monitoring (on fan speed, system/CPU temperature, etc.) is not supported.
- Software SSL can run only on Intel CPUs.
- If the Network Interface Card (NIC) type is "i82545em_copper" or "i82545em_fiber", the system will not perform hardware checksums even if the command "**system tune hwcksum on**" is configured.
- The Maximum Transmission Unit (MTU) cannot be set greater than 1500 for vAPV deployed on Microsoft Hyper-V.
- Bond is not supported for vAPV deployed on Microsoft Hyper-V.
- Trunk Virtual Local Area Network (VLAN) is not supported for vAPV deployed on Microsoft Hyper-V.
- The 8.6.0.x vAPV image does not support downgrades to non-8.6.0.x versions of vAPV images.

9   vAPV FAQ

9.1   General

- How is the vAPV software upgraded?

  For vAPV running as a virtual appliance, the system update is the same as APV, that is, use the command "**system update** …" for the ArrayOS update.

Please be reminded that upgrading the vAPV virtual appliance from earlier 8.3 builds will require a new license.

- How many vCPUs are supported?

  A maximum of four vCPUs are supported.

- Is SSL offload supported?

  Yes. SSL offload (software only solution) is supported on the Intel CPUs only.

- Can I use a third-party hardware SSL accelerator?

  No. You cannot install any third party hardware accelerator and expect it to work with vAPV.

- Can I put vAPV in a cluster/HA configuration with a physical APV?

  No, vAPV clustering with a physical APV is not supported.

- Why does a vAPV as an HA unit remain active when its port becomes "DOWN"?

  Because vAPV uses virtual NICs, hypervisors do not provide the physical port status to our virtual device. For HA to work, the condition of the gateway health check (using the command "**monitor network gateway**") should be used rather than instead of port status condition, for which the port status is always "UP".

- Why there is no default IP address assigned like the APV appliance offers?

  The default IP address assignment is removed from vAPV to avoid potential IP address conflicts in the case that multiple instances of vAPV are first installed on the same hypervisor.

- Is VMware vMotion supported?

  Yes. Because vMotion operates on VMware using generic interfaces, it requires no GuestOS support. vAPV is proven to work with vMotion. Array Networks will work with customers if any issues are found while vAPV is working with vMotion.

- Is VMware HA supported?

  Yes. vAPV supports VMware HA. Array Networks will work with customers if any issues are found while vAPV is running VMware HA.

- Are the virtio NIC and disk supported on KVM?

  Yes. The virtio NIC and disk are supported on KVM.

- Why can't I ping successfully after changing the MAC address of vAPV port?

  Enabling the "MAC Address Spoofing" option on the hypervisor can help avoid this issue.

9.2 Configuration

- **How do I configure bond interfaces on VMware ESXi?**

  Configuring bond interfaces on vAPV will cause duplicate packets. Thus, it is recommended to use the VMware ESXi NIC teaming function to configure bond interfaces for vAPV. Detailed steps are as follows:

  1. Select the target VMware ESXi host.

  2. Click "**Configure** > **Network** > **Add Network**". Select the physical interfaces for bonding while creating the vSphere standard switch (vSwitch), and complete the operation as prompted.

  3. Select the vAPV that requires interface bonding, and click "**Edit Virtual Machine Setting**". Then, associate the network adapter (vAPV's virtual network interface) with the vSwitch just created.

- **How do I configure the trunk VLAN on VMware ESXi?**

1. Edit the virtual machine portgroup associated with the vAPV by setting the VLAN ID as 4095. Thus, VMware ESXi will not untag the VLAN.

2. Configure VLAN on the vAPV and the peer physical switch respectively.

- **How do I configure bond on Citrix XenServer?**

  Citrix XenServer supports creating bond interfaces for vAPV. Detailed steps are as follows:

  1. Enter the XenCenter, click the Citrix XenServer host that requires interface bonding, and select "**NICs**".

  2. Click "**Create Bond**".

  3. Select the interfaces for bonding, and click "**Create**" to complete the bond creation.

  4. Reboot the Citrix XenServer host.

  5. Associate the bond just created with the interfaces of the vAPV.

- **How do I configure VLAN on Citrix XenServer?**

  Virtual machines running on Citrix XenServer do not see VLAN tags. To support VLAN tagging it must be configured on the Citrix XenServer. Detailed steps are as follows:

  1. Enter the Citrix XenCenter, select the target Citrix XenServer host, and click "**Networking** > **Add Network** > **External Network** > **Next**".

  2. In the "**New Network**" pane, select the NIC on which you want to set VLAN.

  3. Set the VLAN ID, which should be the same as the VLAN ID set on the switch connected with the Citrix XenServer, and click "**Finish**".

  4. Select the desired vAPV, and click "**Networking** > **Add Interface**".

  5. In the "**Add Virtual Interface**" pane, select the network on which to set the VLAN. Then, click "**Add**" to complete the operation.

- **How do I limit the interface speed on vAPV?**

  On vAPV, the interface speed setting only supports the "auto" mode, indicating that the interface automatically negotiates the speed. Even if you set the interface speed manually on vAPV, the setting will not work.

9.3 License

- How do I import a valid license to vAPV?

  Please perform the following steps to generate a serial number and import a new valid license:

  1. Generate a new serial number by executing the command "**system serialnumber**".

  2. Contact the Array Networks Customer Support (Email: support@arraynetworks.com) to apply for a new license based on the serial number in the output of the command "**show version**" and the new capacity you wish to purchase.

  3. Import the license obtained from the Array Networks Customer Support (Email: support@arraynetworks.com) by executing the command "**system license**".

  For example:

  AN(config)#**system license** "**be99af00-ef53d5fd-e21501ad-9eee8179-3dd843b3-1108c600-1505fdab-20130108-99999999**"

Note: You cannot copy and paste the command out when accessing the system via the server's console. It is recommended to access the system via SSH in order to generate the serial number and import a valid license.

9.4   Performance

- What is the expected performance on my server?

    The performance you should expect largely depends on the available CPU power of your server and the licensed amount of bandwidth of vAPV. Below is a reference table of performance data collected from QA lab (using Intel Xeon E5-2690 CPU and Intel 10G NIC).

vAPV (8.6)				
Hypervisor	Citrix XenServer 6.1.0	VMware ESXi 5.0	KVM-1.1.1-1.8.1 (on OpenSUSE 12.3)	Microsoft Hyper-V
vCPU	2	2	2	2
vNIC	4	4	4	4
vRAM	2 GB	2 GB	2 GB	2 GB
L4 Throughput	3.7 Gbps	5.6 Gbps	3.8 Gbps	3.3 Gbps
L4 CPS	135K	130K	80K	85K
L4 Concurrent Connection	1M	1M	250K	1M
L7 Throughput	3.7 Gbps	5.6 Gbps	3.8 Gbps	3.3 Gbps
L7 RPS	310K	390K	220K	255K
L7 Concurrent Connection	260K	250K	62.5K	250K
SSL (1024-bit) TPS	2800	2950	950	3000
SSL (2048-bit) TPS	1100	1050	940	1150
SSL Concurrent Connection	10K	10K	10K	10K
SSL Bulk Crypto	1.43 Gbps	1 Gbps	680 Mbps	1.2 Gbps

Note: All L4 performance data, including L4 Throughput, L4 CPS, and L4 Concurrent Connection, are measured with the DirectFWD function enabled (by the "**slb directfwd on**" command).

- Can more than one vAPV instance run on the same hypervisor?

    Yes.

- Can I add more vCPUs to boost the performance?

    You can only license up to 2 vCPUs on the vAPV. If more performance is needed, please move the vAPV instance to a more powerul server running a faster processor.

- Why do I observe CPU usage while no traffic is running through my vAPV?

    This is the same behavior as found on a physical APV appliance. An APV appliance is constantly monitoring the system and looking for work to do. This CPU usage is already accounted for when measuring the performance. In another word, this "extra" CPU usage does not affect the overall system performance.