AppCheck
ACCURACY IS EVERYTHING

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

AppCheck Internal Scan Hub Setup Guide

V2.0.23 (2021-08-11)

Date	Version	Change(s)
(Earlier version changes redacted, see older file versions for full change history)		
2021-03-16	2.0.16	Updating process flow slightly to reflect support team requirements
2021-04-15	2.0.17	Minor layout fix/image placement in Step 11
2021-04-16	2.0.18	Updating broken link for image conversion
2021-06-03	2.0.19	Minor layout fix/image placement in Step 5
2021-06-03	2.0.20	Minor formatting changes to and standardisation of formatting for code block inclusion
2021-06-03	2.0.21	Updating firewall access requirements based upon deployed service changes
2021-06-04	2.0.22	Updating diskspace requirements, adding page numbers
2021-08-11	**2.0.23**	**Removing reference to (legacy) VHD conversion instructions**

This guide will assist you in the necessary steps to deploy a dedicated AppCheck scan hub as a virtual appliance within your hosted (data centre) or cloud estate (AWS) environment. AppCheck provides customer-dedicated scan hubs to allow clients to scan target infrastructure and applications from inside their organisation's perimeter firewall boundary.

Please follow the steps below to deploy your internal scan hub. If you encounter any issues with hub deployment or configuration, please contact our technical support team at tech-support@appcheck-ng.com for assistance.

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Table of Contents

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Hub Setup

Step 1 - Download hub bootstrap image (OVA file)

The hub is deployed via a minimal bootstrap image that is available for download from our website, in the Open Virtual Appliance (OVA) file format.

Hub Download

The hub image is available in the **Ubuntu Linux** distribution format.

https://appcheck-ng.com/get-help/downloads/

The scan hub image is provided in an Open Virtual Appliance (**OVA**) package, which is a *tar* (compressed) archive file with the OVF directory inside. Open Virtualization Format (OVF) is an open standard for packaging and distributing virtual appliances or, more generally, software to be run in virtual machines. The OVF standard is not tied to any particular hypervisor or instruction set architecture but is supported by platforms including VirtualBox, VMWare, RedHat Enterprise Virtualization and Oracle VM.

The same build version is also offered in **VHD** (Virtual Hard Disk) format from the same download location. The VHD format is supported by Microsoft Virtual PC, MicrosoftVirtual Server, Hyper-V (within Windows Server 2008) and Oracle VirtualBox.

Alternative Hub Image Formats & Image Conversion

Conversion to OVF

Should you need to manually extract the OVA file or convert it to Open Virtualization Format (OVF) format for any reason on your chosen hypervisor platform, then you can do so using the free OVFtool from VMWare:
https://code.vmware.com/web/tool/4.3.0/ovf

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 2 (OPTIONAL) – Confirm file checksums

A checksum is a hash calculated from the downloaded file for the purpose of detecting errors that may have been introduced during its transmission or storage via error or file corruption/interference. It allows you to verify that the image you have downloaded is free of errors.

You can optionally compare the checksum for the file you have downloaded with the reference values provided on the download page at https://appcheck-ng.com/get-help/downloads/

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 3 – Ensure Virtualisation is enabled in BIOS / UEFI firmware

NOTE: You can likely skip this step if using a server-grade hypervisor. This step is largely applicable to personal-grade hypervisor users

Modern CPUs include hardware virtualization features that help accelerate virtual machines created in VirtualBox, VMware, Hyper-V, and other apps. But those features aren't always enabled by default.

If not enabled, you may encounter error messages when deploying the internal hub, like the following:

```
VT-x is disabled in the BIOS for all CPU modes (VERR_VMX_MSR_ALL_VMX_DISABLED).
```

If you have an Intel CPU and uninstalling Hyper-V didn't solve your problem—or your virtualization app reported that Intel VT-x was disabled—you'll need to access your computer's BIOS or UEFI settings.

On a BIOS-based system, you'll access BIOS settings by restarting your PC and pressing the appropriate key right when it first boots. The key you press depends on your PC's manufacturer, but it's often the "Delete" or "F2" key. You also will most likely see message during startup that says something like

```
Press {Key} to access setup.
```

Find a setting such as "*Intel VT-x*," "*Intel Virtualization Technology*," "*Virtualization Extensions*," or similar. Often, you'll find the option under a "*Chipset*," or "*Advanced CPU Configuration*" menu.

Enable the option and then select "Save and Exit" or the equivalent feature to save your settings changes and reboot your host.

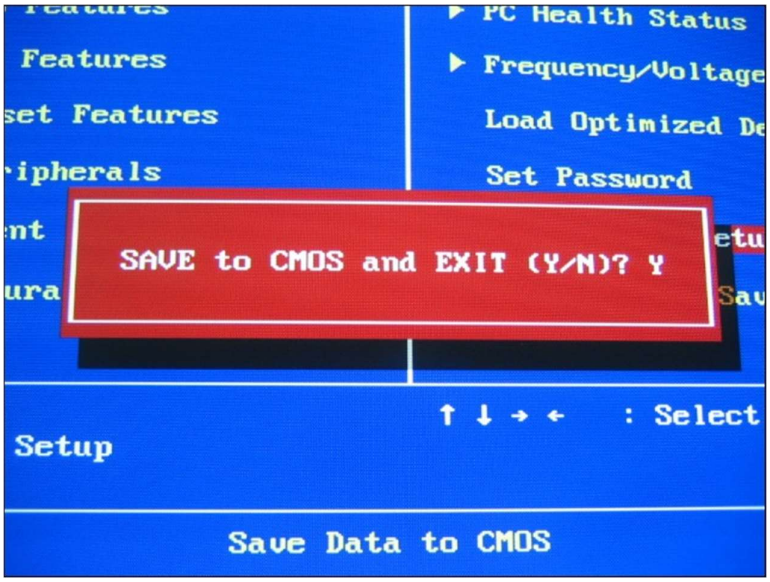

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 4 - Deploy hub to chosen Hypervisor/Cloud platform

The exact mechanism to deploy your new internal hub will depend on your hypervisor, but typically you will need to **Import** the OVA image.

You should use the following settings if your hypervisor requires them:

Machine/OS Type Selection

The virtual machine image is based on **Ubuntu Linux**. When deploying the virtual appliance from the supplied image, most hypervisors will be able to determine the OS automatically. Where your hypervisor requires you to specify the OS, please select the "**general Linux**" setting or similar, dependent on your hypervisor platform.

However, if your hypervisor requires selection of specific OS, please select then select "**Ubuntu (64-bit)**". If your hypervisor further requires selection of the specific OS version, please select "**Ubuntu 18.04 (Linux) OS**".

Virtual Machine (VM) Configuration & Resource Assignment

It is recommended that once you have deployed the virtual appliance, you provide it with at least **16GB RAM.**

The number of assiged CPU cores should match the maximum number of concurrent scans that you want to run, plus one core for other OS function, with a minimum requirement of **4 cores recommended**.

For example, to support **4** concurrent scans, please assign **5** CPU cores.

Disk Capacity

The VM disk image requires a minimum of **60GB** of space on the underlying storage tier or datastore.

AppCheck recommends that the virtual disk provisioning be configured to use dynamically allocated virtual disk/image sizing, managed by the hypervisor, where supported. If only static disk size allocation is supported by your hypervisor or cloud platform, then the image should be statically sized to 60GB.

Cloud Deployment Guidance

NOTE: Please see Appendix A for deploying to Amazon AWS EC2 cloud platform.

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 5 – Power on the internal hub

Ensure that the hub has network provision (in some hypervisors this may be **Bridged Adapter** mode, and then power on the virtual machine in your hypervisor.

After a short boot process, you should be presented with a console screen as per the below screenshot:

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 6 – Accessing the host console (OPTIONAL)

Accessing the VM Console

Once the hub is deployed, access should be possible via the virtual console on your hypervisor platform. See your hypervisor (eg VMWare) documentation for further detail.

Access to the appliance at the command line is not essential, but for some clients, including those needing to set a static IP directly on the device rather than via DHCP assignment (see steps later in this guide), then console access will be needed. Access via console can be gained by logging in with the credentials:

```
Username: appcheck
Password: Flood-Presence-Bus-Hurry-8
```

Root (Sudo) Access

In some cases, if there is some particular reason that you need to configure a host using a system that requires sudo access in order to gain root (superuser) privileges, you can do so using the sudo command as below:

```
sudo -i
```

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 7 - Hostname & IP Assignment

Hostname and IP Address Assignment

You are able to assign a static IP address to the host via either direct static assignment on the host, or a statically assigned address via DHCP. Hostname assignment is not normally necessary for a standard internal hub deployment and should not be necessary in most environments.

The host OS uses `netplan` for its network configuration. To configure a static IP address on your Ubuntu 18.04 server or to set the IP to be assigned by DHCP, you need to modify a relevant netplan network configuration file within `/etc/netplan/` directory.

To set up DHCP, ensure that the default netplan configuration file called `01-netcfg.yaml` has the following content, where `enp0s3` is your adapter interface found listed in `/sys/class/net/:`

Please note that netplan files use YAML formatting. Whitespace indentation is used for denoting structure. It is very important that you maintain indentation of two space characters at each "level" as in the below

```
# This file describes the network interfaces available on your system
# For more information, see netplan(5).
network:
  version: 2
  renderer: networkd
  ethernets:
    enp0s3:
      dhcp4: yes
```

To set your network interface `enp0s3` to static IP address `192.168.1.222` with gateway `192.168.1.1` and DNS server as `8.8.8.8` and `8.8.4.4` replace the above configuration with the one below:

```
# This file describes the network interfaces available on your system
# For more information, see netplan(5).
network:
  version: 2
  renderer: networkd
  ethernets:
    enp0s3:
      dhcp4: no
      addresses: [192.168.1.222/24]
      gateway4: 192.168.1.1
      nameservers:
        addresses: [8.8.8.8,8.8.4.4]
```

Once ready apply changes with:

```
$ sudo netplan apply
```

In case you run into some issues execute:

```
$ sudo netplan --debug apply
```

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 8 – Change default credentials (RECOMMENDED)

The credentials listed in the step above are common to all hubs downloaded. From a security point of view is is therefore highly recommended to change these default credentials to a unique set. However please be aware that if you do so, AppCheck will not know these credentials, and you may need to provide credentials to AppCheck support staff in the event of an issue with your scan hub appliance requiring remote (shadow) support assistance.

To change the credentials enter the following command when logged in as the **appcheck** user:

```
appcheck@host:~$ passwd
```

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 9 - Open up firewall access (OUTBOUND)

After your host has had a static IP configured, it is important to ensure that you have then add the permitted outbound access for your internal hub to necessary services.

All access granted is **outbound** from your network to AppCheck, and it is not necessary to open up any inbound connectivity from the public internet.

Once the hub is deployed and an IP address assigned, the next step is to set up outbound (hub-initiated) access to some remote endpoints. The internal hubs are designed to work behind NAT so do not need external IP addresses allocated to them; they do however require access outbound to the following hosts and services:

Source	Destination Host	Destination (IP)	Port(s)	Protocol
Purpose: Hub system and OS updates and provisioning				
(internal hub)	assets.appcheck-ng.com	167.99.85.223	80, 443	TCP
(internal hub)	*.archive.ubuntu.com	-	80, 443	TCP
(internal hub)	docker.appcheck-ng.com	68.183.33.54	80, 443	TCP
Purpose: Hub command & control communication with AppCheck cloud platform				
(internal hub)	wire1.appcheck-ng.com wire2.appcheck-ng.com	178.128.173.89	80, 443, 5671	TCP
(internal hub)	wire3.appcheck-ng.com	178.128.163.167	80, 443, 5671	TCP
(internal hub)	lograbbit.appcheck-ng.com	178.62.17.110	80, 443	TCP
Purpose: DNS / hostname resolution				
(internal hub)	dns.google	8.8.8.8 8.8.4.4	53	TCP, UDP
Purpose: Scan hub software licence activation and renewal				
(internal hub)	licensing.appcheck-ng.com	104.248.173.23	80, 443	TCP
(internal hub)	licensing-master.appcheck-ng.com	142.93.43.105	80, 443, 4505, 4506	TCP

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 10 – Configure HTTP Proxy

The deployed internal hub will, during normal operation, call out to the AppCheck cloud platform on ports 80 and 443 in order to retrieve Command & Control (C&C) tasking, and to report back results. These connections are initiated outbound from your network to the AppCheck cloud.

Because of the use of ports 80 and 443, some customers may find that the traffic is intercepted (and blocked) by operated HTTP proxies. Despite using port 80 and 443 (among others), the traffic is **not** HTTP traffic (it uses a custom protocol), so it is necessary to **bypass the HTTP proxy** (if you use one) or **add a whitelist/exception** for each of the below endpoints:

Source	Destination Host	Destination (IP)	Protocol
(internal hub)	wire1.appcheck-ng.com wire2.appcheck-ng.com	178.128.173.89	TCP
(internal hub)	wire3.appcheck-ng.com	178.128.163.167	TCP
(internal hub)	licensing.appcheck-ng.com	104.248.173.23	TCP
(internal hub)	licensing-master.appcheck-ng.com	142.93.43.105	TCP
(internal hub)	docker.appcheck-ng.com	68.183.33.54	TCP
(internal hub)	lograbbit.appcheck-ng.com	178.62.17.110	TCP
(internal hub)	assets.appcheck-ng.com	167.99.85.223	TCP

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 11 – Access the hub setup GUI

Once you are happy that your hub's IP has been assigned outbound internet access as per the above firewall configuration guidance, then you should confirm outbound access and finalise hub setup.

To do this, access the login GUI (web interface) that the hub presents on a network port. The address for this is given in the console output when opening a virtual console to your host (see **Step 5**).

Access this interface and login using the credentials below:

```
Username: admin@appcheck-ng.com
Password: 48-Fear-Some-Shoulder-77
```

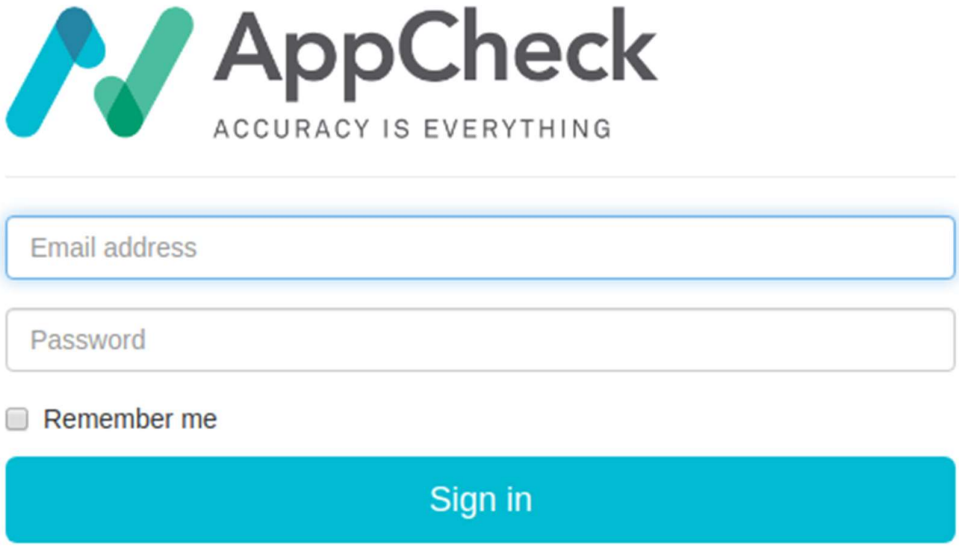

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 12 – Confirm outbound connectivity.

Once authenticated to the hub's main dashboard, review the connectivity, and confirm that all required connectivity is in place. When everything shows green, click **Next** to continue.

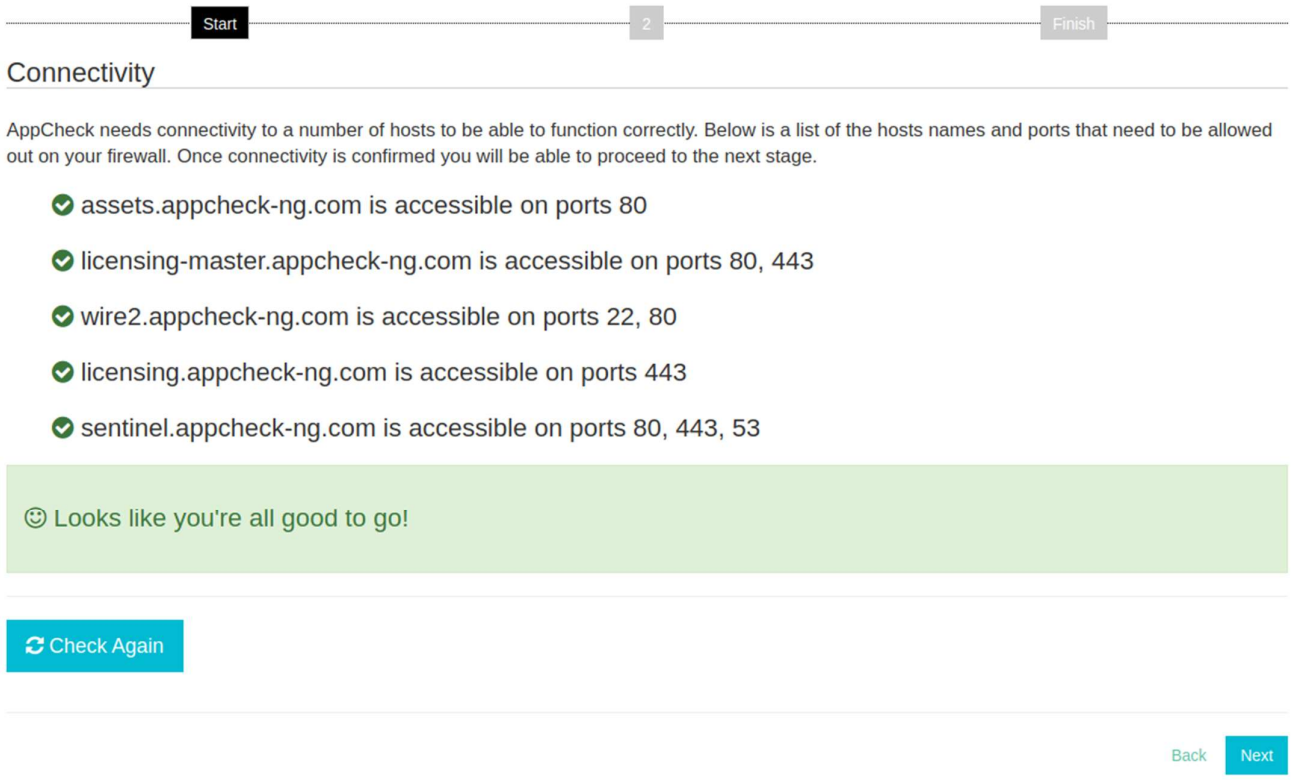

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 13 – Hub licence activation

The final configuration step is to activate the hub using the licence key provided by your account manager. Please contact your account manager if you are unsure what your licence key is.

Once the licence key has been entered click **Next** to continue:

Click **Finish**

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 14 – Package update

When you click **Finish** the hub will perform a full package update.

You can refresh the web interface to see progress of provisioning and wait until it is complete, however this process may take up to 1 hour.

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 15 – Contact AppCheck Support

Your hub is now provisioned and is ready to be linked to your scan portal account. **Please contact the AppCheck technical support team and request that they link the hub to your customer account**.

If you do not already have the contact details for the support team, please see our website at https://appcheck-ng.com/get-help/ or ask your AppCheck account manager for the team's contact details.

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 16 – Configure FQDN resolution

When the hub is licenced and ready to perform scanning, it is important that you consider target address resolution. Internal network hostname and web application targets often do not appear in public DNS records, so manual configuration is required so that the scan hub is able to resolve these and target them for scanning.

Determining whether you need to configure FQDN Resolution

- For **infrastructure targets designated by IP** (eg "*192.168.0.1*", "*10.0.0.2*"), there is no further configuration required.
- However, for both (a) **infrastructure targets designated by FQDN** (eg "*SCBG221.network.co*") and (b) **web application targets** (eg "*https://internal.example.com*") then you will need to configure the internal scan hub so that it is able to resolve the hostnames and Fully Qualified Domain Names (FQDNs) to IP addresses for scan targeting.

Configuring FQDN Resolution

In order to set a list of FQDN to IP address mappings , you can add entries in the Linux `/etc/hosts` format to your scan hub page at https://scanner.appcheck-ng.com/scan_hubs – these settings will be collected by your internal scan hub(s) and used to resolve scan targets specified in FQDN format.

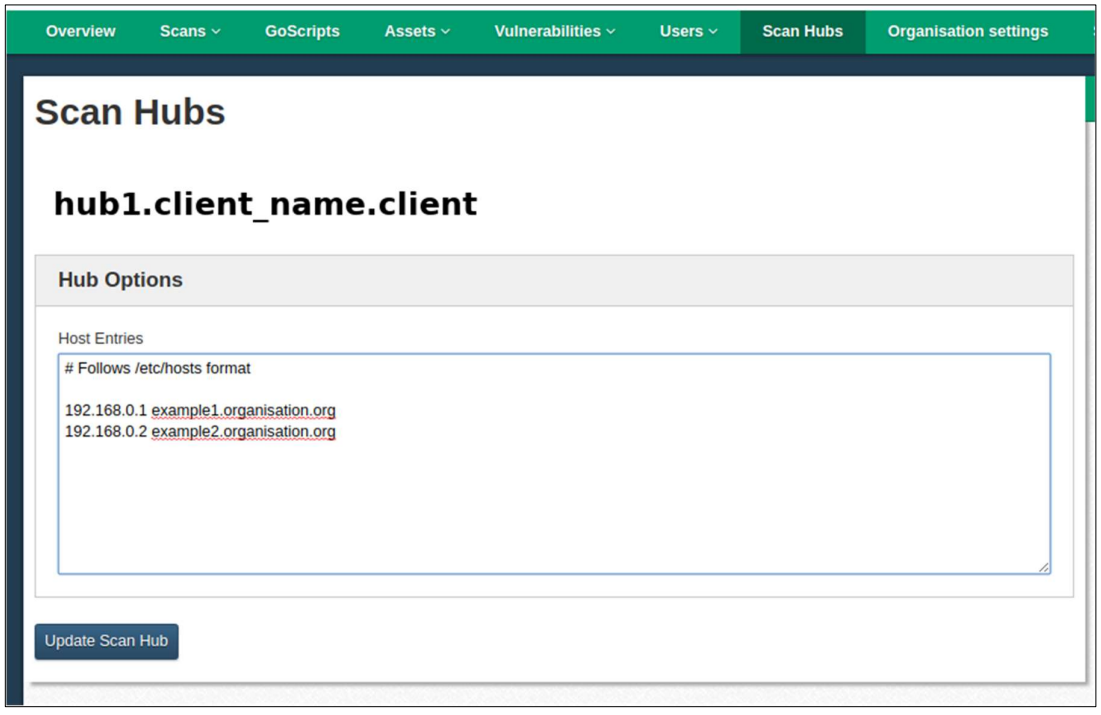

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Step 17 – Perform a test scan

You are now ready to perform a test scan using your new hub. We would recommend testing against a single target initially.

You will need to assign your new scan hub to any scans that you wish it perform by setting the scan hub in the "**Advanced Config Settings**" options at the foot of your AppCheck scan portal scan configuration page:

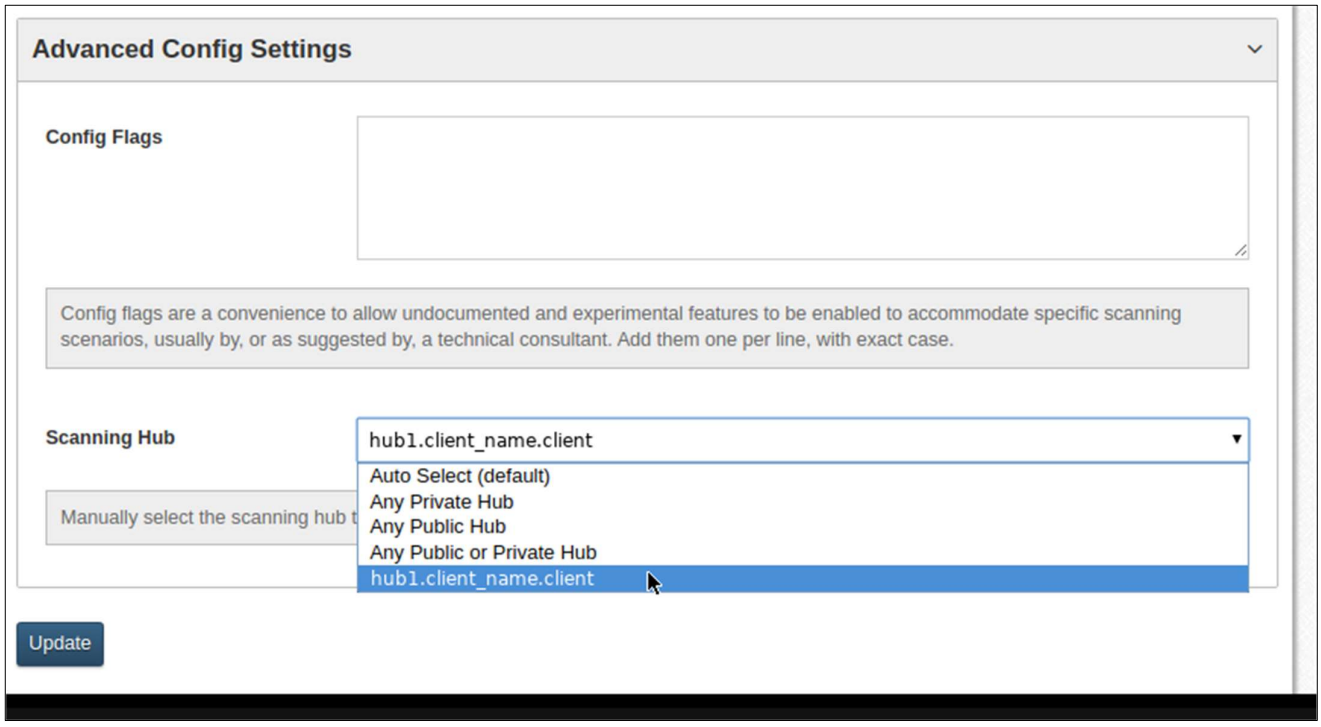

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Target Definitions

Targeting by IP

Targeting of hosts by IP address will work out of the box and no additional configuration is required.

Targeting by FQDN (entry in public DNS)

If scanning a hostname or service name by FQDN (fully qualified domain name) that exists in public DNS then this should just work out of the box. By default, the scanner will ship with Google's 8.8.8.8 DNS resolver set, and use this for DNS resolution. If your organisation blocks external DNS then you will need to ensure this has been whitelisted for the internal hub in Step 5 above.

Targeting by FQDN (no entry in public DNS)

For some truly internal services that are not exposed publicly, it is likely that it will not exist in DNS. You may therefore need to add internal host entries for any FQDNs (fully qualified domain names / hostnames) that you wish to target with scans. This can be done using the `/etc/hosts` format entry on the Scan Hubs page of your scan portal at https://scanner.appcheck-ng.com. If you do not see this option, please contact AppCheck support for assistance.

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

FAQs

What is the course of action in the event of an issue developing with the internal hub appliance or a failed appliance?

AppCheck support staff proactively monitor the scan hubs using heartbeat checks and synthetic monitoring and are proactively alerted to scan hub failures. In the first instance, a scan hub can be rebooted remotely, or escalated to a senior team member for further debugging and investigation. In the event of a total scan hub failure, a new scan appliance would need to be redeployed. Other than currently executing scans, scan hubs preserve no state locally (state preserved in cloud) so there is little or no impact to redeploying a failed scan hub, and it is also highly rare occurrence.

Are there regular updates issued for the scanner?

Yes, updates are regularly and automatically applied. Customers will be notified if it is necessary to completely redeploy appliances, but that would be a highly rare occurrence. Updates are normally applied seamlessly in the background without client intervention. Ubuntu is presently the base OS and from time to time the hub needs to update its OS packages. It is important to permit the access to the update servers as per Step 5 of this guide.

E: support@appcheck-ng.com
W: www.appcheck-ng.com
T: 0113 887 8380

Appendices

Appendix A - Deploying to Amazon AWS EC2 Cloud

We do not offer an AMI image or specific AWS deployment guidance at this time. However, Amazon produce a guide to converting a virtual machine into an image suitable for import to AWS EC2 at the following URL:

https://aws.amazon.com/ec2/vm-import/