

Aruba AirWave Connectivity Issues on Hyper-V

Created by Travis G. Kench – travis@tkcomputersolutions.com – 4/2/2018

Issue: AirWave on Hyper-V randomly loses connectivity when multiple vCPUs are assigned to the VM. I have verified that this is an issue that exists on Windows Server 2012 R2 Datacenter (VM Configuration Version 5.0) and also on Windows Server 2016 Datacenter (VM Configuration Version 8.0).

Summary


In most circumstances more than one vCPU will be assigned to the AirWave VM to help balance the processor load. In this case when you assign more than one vCPU to a Hyper-V VM running AirWave (CentOS 6.9) you will encounter one of two things from my experience. It can cause the VM to randomly loose connectivity which can sometimes be resolved by a restart of the VM or it can cause the VM to completely lose connectivity in which a restart will not resolve the issue. This documentation explains how you should appropriately assign resources for the AirWave VM and discusses how I developed an ideal configuration for my environment. The information below will help you develop a configuration that should be beneficial for your environment. If you haven't already purchased AirWave please make sure you run through the calculations below before you buy it so that you can assess your network infrastructure's currently available resources. You may not have a server in production with enough resources available to efficiently use the application so you may have to do some upgrades or potentially buy another server that can handle the applications load.

Current Setup Experiencing Issues

- Physical Server Hyper-V Host: Dell PowerEdge R520
 - OS: Windows Server 2016 Datacenter
 - RAM: 160GB
 - CPU: (2) Xeon E5-2430 @ 2.20GHz
- Hyper-V VM Configuration
 - AirWave Version: 8.2.6 (CentOS 6.9 Final)
 - VM Generation: 1
 - CPU: 4 vCPU
 - RAM: Startup 68GB & Dynamic 8GB-68GB w/20% memory buffer
 - Disk: 300GB VHDx on iSCSI SAN w/10k SAS drives
 - Network Adapter: Legacy
 - Note: AirWave will not install with a standard Hyper-V default network adapter.
 - Linux Integration Tools: 4.2.2
 - Hyper-V VM Configuration Version: 8.0
- AirWave 8.2.6 Average Stats:
 - Access Points UP: 299
 - Rogues Detected: 918
 - Clients: 900

AirWave Min-Max Wireless Client Sizing Analysis

According to my AirWave server we are peaking at around 900 clients as show below.



AirWave Server Sizing

As of this tutorial's creation date there isn't any formal documentation informing us of the exact AirWave Hyper-V hardware system requirements that we should be assigning the VM as a best practice scenario. We have to make estimates based upon the [AirWave 8.2 Server Sizing Guide](#) table as shown below. Please note that the server sizing guide also states that "Aruba recommends a 20% increase buffer for virtual environments" so please make sure you add that buffer.

Model	AW-HW630-PRO	AW-HW630-ENT
Maximum Devices	1500	4000
CPU	1 x Intel Xeon E5-2640, 2.5 GHz	Dual Xeon e5-2640 v3 2.6GHz
Physical Cores	6	12
Passmark Score	9761	14846
RAM	48 GB	96 GB
Disks (RAID10 only)	6 x 146 GB 15k rpm SAS	8 x 146 GB 15k rpm SAS
Storage Capacity	800GB	1.1 TB
Maximum Input/Output Operations Per Second (IOPs)	2132	2842

If I followed the math behind the calculations for the “AW-HW630-PRO” I would assign the following resources.

AW-HW630-PRO Calculations with 1500 Maximum Devices

1500 (Maximum Devices) / 6 (Physical Cores) = 250 devices per core

1500 (Maximum Devices) / 48 (GB RAM) = 31.25 clients per GB RAM

Applying their logic to my Hyper-V VM with 1000 Maximum Devices Estimate

1000 (Maximum Devices) / 250 devices per core = 4 vCPUs

1000 (Maximum Devices) / 31.25 clients per GB RAM = 32 GB RAM

According to the [AirWave Server Sizing Calculator](#) online which is based upon the following information below I should be assigning the following resources to the VM in my environment:

Device Count: 1000

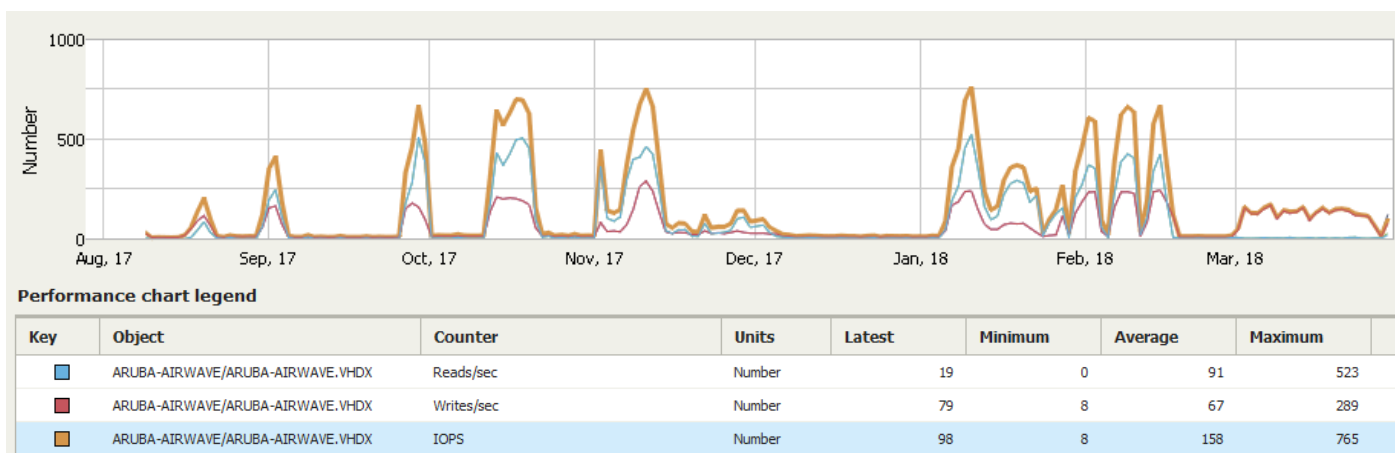
Elements	Value
High mobility environment?	TRUE
Number of Clients per AP?	30
Many small/single IAP clusters?	FALSE
Mostly new clients?	FALSE
More than 20% monitored devices are switches?	FALSE
High SNMP trap volume?	FALSE
Large demand for reports or fast report generation?	TRUE
Higher than default retention of monitoring data?	TRUE
High volume of rogues?	TRUE
High VisualRF utilization?	TRUE
Greater than default SNMP polling periods?	FALSE
Using AMON preferred?	FALSE
Virtualized environment?	TRUE

According to this sizing calculator I should assign the following system requirements:

- Memory: 68GB
 - According to [VeeamOne](#) this is pretty accurate for my environment. I am currently using 66GB of 68GB. I have watched it release dynamic memory at times so this is working for my setup currently and I will continue to monitor it.
 - Note: If you do not assign a reasonable amount of memory to the VM it will page to the virtual disk and can potentially fill up the disk resulting in boot and backup issues.
- Disk Space: 1479GB
 - My VHDx file is 300GB with 115GB being used at this time so I cannot comment as to whether or not it will ever need a VHDx drive that big. I haven't been able to find any information on why I would need such a large hard drive. Possibly the drive may need to be expanded in the future depending on log file sizes, reports, or that the database size has grown so I will be monitoring those variables over time. I am also thinking that there

may be some reasoning behind the large disk size recommendation being that not all environments have enough physical memory to allocate to a VM so it does in fact page a lot to the hard disk.

- CPU: PassMark Score => 19956
 - The AirWave online calculator doesn't provide you with a particular vCPU quantity instead it goes off of **PassMark** benchmark scores.
 - The CPU that I have documented above in my current setup has a PassMark score of **13637** and is more than sufficient based upon analyzation and testing with VeeamOne as you will see below.
- Disk IOPS: 6254
 - My implementation was done back in August 2017, VeeamOne reports a maximum of 765 IOPS and an average of 158 IOPS. So their recommended value is extremely high for my environment at this time.



With the help of VeeamOne, I have tweaked my VM system resources so that I wasn't over assigning vCPUs and memory. I ended up assigning the resources to my VM as documented in the specifications on the first page of this document.



Testing and Resolution of Connectivity Problems:

1. Test connectivity with ping inside the AirWave host as shown below. Let the ping run for 30-60 seconds so that it creates a group of errors within the log and then press Ctrl-Z to stop the ping.

```
[root@ARUBA-AIRWAVE mercury]# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
From      icmp_seq=10 Destination Host Unreachable
From      icmp_seq=11 Destination Host Unreachable
From      icmp_seq=12 Destination Host Unreachable
From      icmp_seq=14 Destination Host Unreachable
From      icmp_seq=15 Destination Host Unreachable
From      icmp_seq=16 Destination Host Unreachable
From      icmp_seq=18 Destination Host Unreachable
```

2. Run the following command to view the kernel message log. The `less` command option allows the output to be moved forward one screen at a time by pressing the space bar.

```
[root@ARUBA-AIRWAVE mercury]# dmesg | less
```

3. Scroll through the log with the spacebar to see if you are generating the error messages below:

```
eth0: 21140 transmit timed out
eth0: 21140 transmit timed out
eth0: 21140 transmit timed out
eth0: 21140 transmit timed out
eth0: 21140 transmit timed out
```

4. Check to see if irqbalance is turned on.

```
[root@ARUBA-AIRWAVE mercury]# chkconfig | grep irqb
irqbalance    0:off    1:off    2:on     3:on     4:on     5:on     6:off
```

5. If it is reporting as ON turn it off with the command below.

```
[root@ARUBA-AIRWAVE mercury]# chkconfig irqbalance off
```

6. Restart the AirWave VM and your connectivity problem will be solved.

```
[root@ARUBA-AIRWAVE mercury]# reboot
```

Bug Background

This issue has been publicized via the following Red Hat links and at this time there isn't a fix other than to turn irqbalance off. The current version of AirWave that I have been troubleshooting is based on CentOS 6.9 which is derived from Red Hat Enterprise Linux sources.

Why is irqbalance not balancing interrupts?

[Bug 655855 - RHEL network stops functioning on Hyper-V hosted VMs with irqbalance daemon started](#)