Validate Wi-Fi Throughput : Testing and Monitoring Guide

Contents

Introduction

This document describes how to monitor and troubleshoot throughput issues in large Wi-Fi networks.

Context

In Wi-Fi networks, there are not that many types of end-users perceived problems.

Reported problems can range between:

- clients not being able to connect;
- clients getting disconnected suddenly or;
- the perceived speed of the application on the user device is not satisfactory.

Behind these simple symptoms can lie hundreds of types of problems, most not even involving the actual Wi-Fi networks like DNS problems, Internet connection problems, and so on.

Management servers like Cisco Catalyst Center help the administrator to troubleshoot specific issues and this article does not go in detail over those many types of daily issues that can be easily seen and remediated through Catalyst Center. Instead, this document focuses on the more vague feedback from end users that the network is slow.

How to test that? How to validate the actual throughput throughout your network? How to triage the speed-related problems into actionable items to improve the overall end user experience?

These are all questions this document tries to answer.

Defining the maximum expected throughput

The first question in each network is: what is the maximum speed that could potentially and realistically be reached?

Since Wi-Fi is a shared medium, the speed experienced depends directly on the number of clients and devices using the Wi-Fi at the same moment on the same channel. Therefore, this question of the actual maximum speed that can be achieved directly implies having a single client device and a single access point in a quiet isolated place where no one is using the same Wi-Fi channel. In these conditions, the factors to determine the maximum speed boil down to:

- The Wi-Fi protocol used (Wi-Fi 5, Wi-Fi 6, ...)
- The hardware capabilities of the client and the access point (number of antennas, number of spatial streams, Ethernet connection of the access point, ...)
- The configuration (channel width, ...)

Knowing these factors allows you to have an estimate of what the maximum real-life throughput you could hope to reach in lab conditions.

To get a quick idea, you check at which data rate your client reports to be connected to the access point. This data rate is not the actual throughput you are able to prove in your tests. This is because Wi-Fi is a half-duplex medium that has some management overhead (frames need to be acknowledged, beacons need to be transmitted) and also short silences between frames for better reception and decoding. This means that, when data is sent, it is sent at the documented data rate, but data is not always sent. Management and control frames are sent at a much lower data rate to ensure reception. An estimate is that you can consider achieving 65 to 70% of the data rate used in an actual throughput test. For example, if your client reports being connected and sending data at 866Mbps, actual tests must report a transfer speed around 600Mbps.

If you know the configuration parameters in use as well as the hardware capabilities of the involved devices, you can also figure out which maximum data rate (and therefore throughput, by using the percentage calculation documented in this section) must be achievable.

If there is a mismatch between the reported data rate and the one you were hoping to achieve, you can start the troubleshooting process by the configuration and verify the various parameters to understand where the gap is.

One example: if you have an Access Point model C9120 broadcasting at 20Mhz channel width in the 5Ghz band and a typical 2 spatial streams Wi-Fi 6 client, you can calculate that, in a perfectly clean RF (Radio Frequency) environment, with a single client you could hope to achieve 160 to 200Mbps in a single file transfer.

More information on throughput testing and validation is documented here: https://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/212892-802-11ac-wireless-throughput-testing-and.html.

Establish the baseline experience

It is important to know what can be expected in your venue in typical circumstances. It is often the case that a technician visits the empty site before deployment roll-out, runs speed tests and documents expected numbers.

Then employees or customers come in, the site gets busy, and the actual experience differs a lot.

After a deployment goes live, it is a clever idea to send technicians to measure the actual experience in each area and take a note of how the network looks like on an average good day.

This includes average amount of clients per radio when the network is operating at a satisfactory level as well as the average throughput achieved with a speed test.

Look for deviations of the experience

When operating your network, monitoring for major alerts or devices that suddenly go down is easy. This document focuses on the hard part: how to spot a wireless network that still works but provides a subpar end user experience.

Finding evidence of a problem (passive testing)

You have tested your network yourself; you know it operates fine and you are monitoring your management systems and dashboards. Nothing is reported as down: you can take a step back and relax. Or can you?

If you wait for echoes from end-users complaining about the poor experience, chances are that you are too late. When end-users complain, the issue has been going on for a long time and you only hear from the few users who were vocal enough for you to hear it.

Countless users have been frustrated already, said nothing to you or your helpdesk, but gave a bad reputation to your network.

So, the question is: how can you spot occurrences of poor experience as soon as they occur?

1. The client assurance dashboard on Cisco Catalyst Center

In the Cisco Catalyst Center assurance dashboard, you have an overall graph of your clients health.

There always are some clients who are unable to connect because someone entered the wrong key, or the device is sitting at the very edge of your coverage, so do not hope to reach 100% of healthy clients but be familiar with what is a good percentage of healthy clients for your environment.

Being in the 90s range is typically good news.

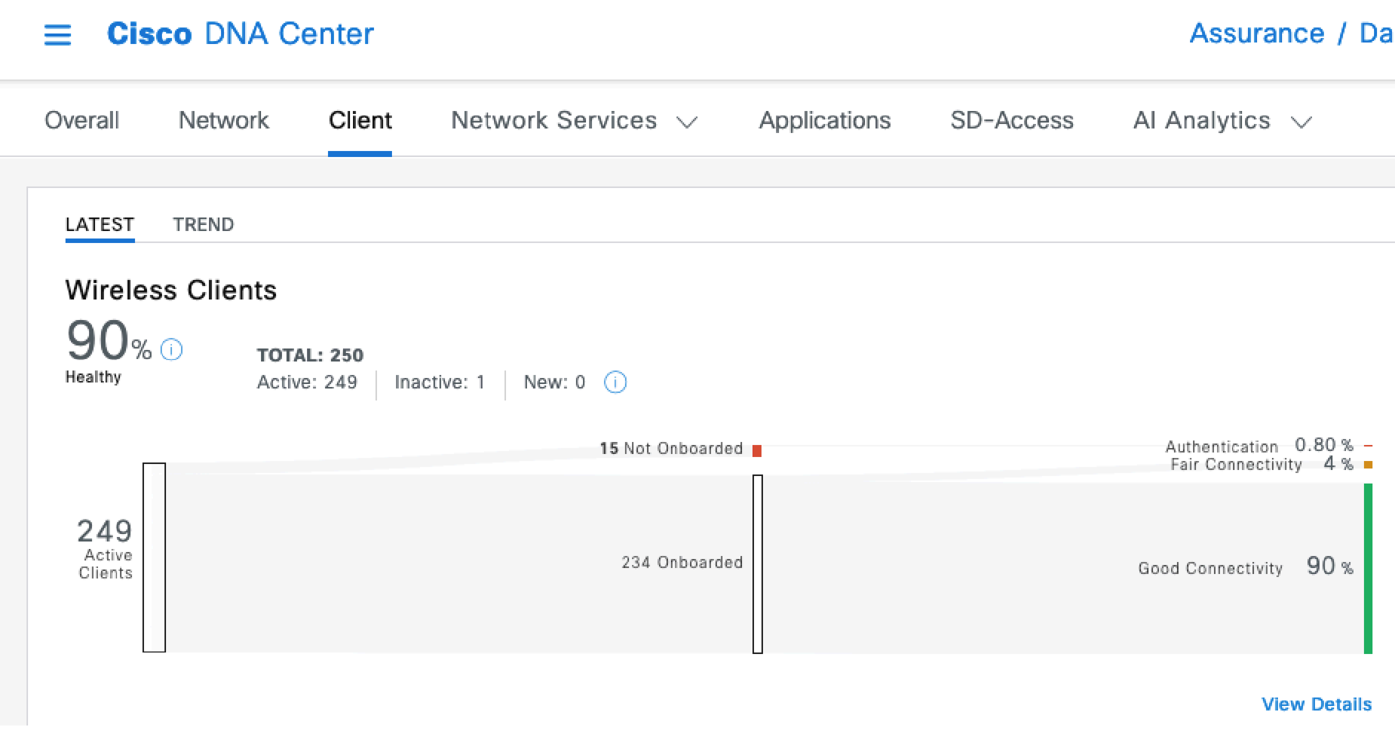

Wireless Clients

90% Healthy

TOTAL: 250

Active: 249 | Inactive: 1 | New: 0

249 Active Clients

15 Not Onboarded

234 Onboarded

Authentication 0.80 %

Fair Connectivity 4 %

Good Connectivity 90 %

With a very quick glance you can see what is happening to the clients who are not healthy:

- Are they far away from the AP (Access Points)?
- Is it an authentication problem?

You can easily see on this graph the ratio of each category.

In the same range of ideas, you can scroll to the bottom of that page and filter to display the client devices that are reported as having poor health. You can then try to spot if there is any pattern:

- They potentially are all connected on 2.4Ghz band (which is known to give a poorer experience in many cases);
- They are potentially all reported at a low signal strength;
- They are potentially all in the same area physically.

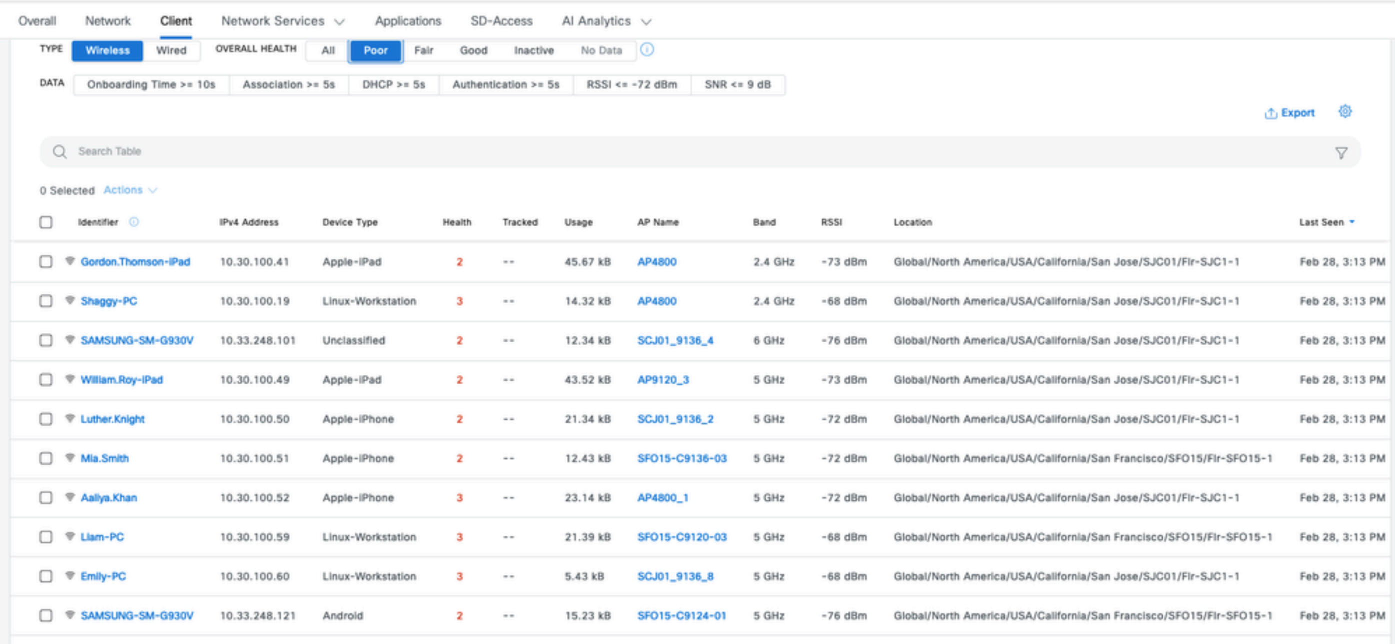

Identifier	IPv4 Address	Device Type	Health	Tracked	Usage	AP Name	Band	RSSI	Location	Last Seen
Gordon.Thomson-iPad	10.30.100.41	Apple-iPad	2	--	45.67 kB	AP4800	2.4 GHz	-73 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
Shaggy-PC	10.30.100.19	Linux-Workstation	3	--	14.32 kB	AP4800	2.4 GHz	-68 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
SAMSUNG-SM-G930V	10.33.248.101	Unclassified	2	--	12.34 kB	SCJ01_9136_4	6 GHz	-76 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
William.Roy-iPad	10.30.100.49	Apple-iPad	2	--	43.52 kB	AP9120_3	5 GHz	-73 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
Luther.Knight	10.30.100.50	Apple-iPhone	2	--	21.34 kB	SCJ01_9136_2	5 GHz	-72 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
Mia.Smith	10.30.100.51	Apple-iPhone	2	--	12.43 kB	SFO15-C9136-03	5 GHz	-72 dBm	Global/North America/USA/California/San Francisco/SFO15/Flr-SFO15-1	Feb 28, 3:13 PM
Aaliya.Khan	10.30.100.52	Apple-iPhone	3	--	23.14 kB	AP4800_1	5 GHz	-72 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
Liam-PC	10.30.100.59	Linux-Workstation	3	--	21.39 kB	SFO15-C9120-03	5 GHz	-68 dBm	Global/North America/USA/California/San Francisco/SFO15/Flr-SFO15-1	Feb 28, 3:13 PM
Emily-PC	10.30.100.60	Linux-Workstation	3	--	5.43 kB	SCJ01_9136_8	5 GHz	-68 dBm	Global/North America/USA/California/San Jose/SJC01/Flr-SJC1-1	Feb 28, 3:13 PM
SAMSUNG-SM-G930V	10.33.248.121	Android	2	--	15.23 kB	SFO15-C9124-01	5 GHz	-76 dBm	Global/North America/USA/California/San Francisco/SFO15/Flr-SFO15-1	Feb 28, 3:13 PM

2. The network assurance dashboard and device 360 on Cisco Catalyst Center

A particularly good metric to spot a specific potential area of problems is to go to the Network Assurance page of Cisco Catalyst Center. You have a widget showing the top access points by client count:

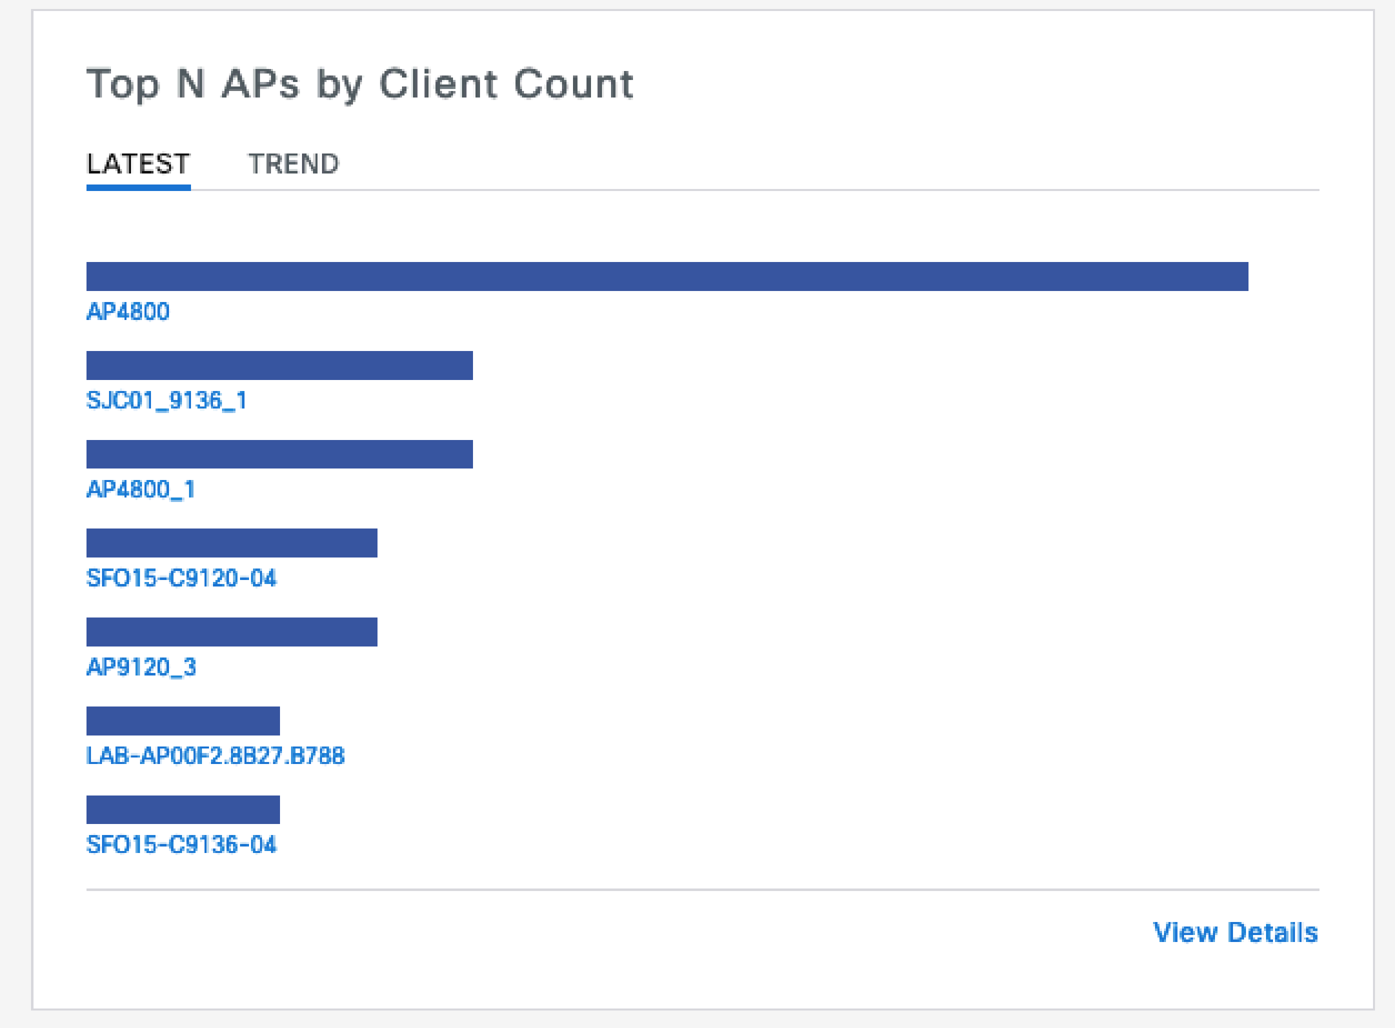

If the top access point in your network has 40 clients connected, you are good. This implies that all the other APs (Access Points) have a lower client count.

On the other hand, if you find the top AP(s) having an unusually high number of clients, you can make a guess that the client experience there is particularly poor (unless most clients are sleeping and not active on the network).

You can then move to a “per AP” investigation where you zoom in on the specific top APs reported in this widget to understand their current health.

Another method of looking at client count is to go to the maps in the Network Hierarchy page of your Catalyst Center. Once in the floor view page, click on “View Options” and in the Access Points section, change the display to “Assoc. Clients" to display the client count per AP:

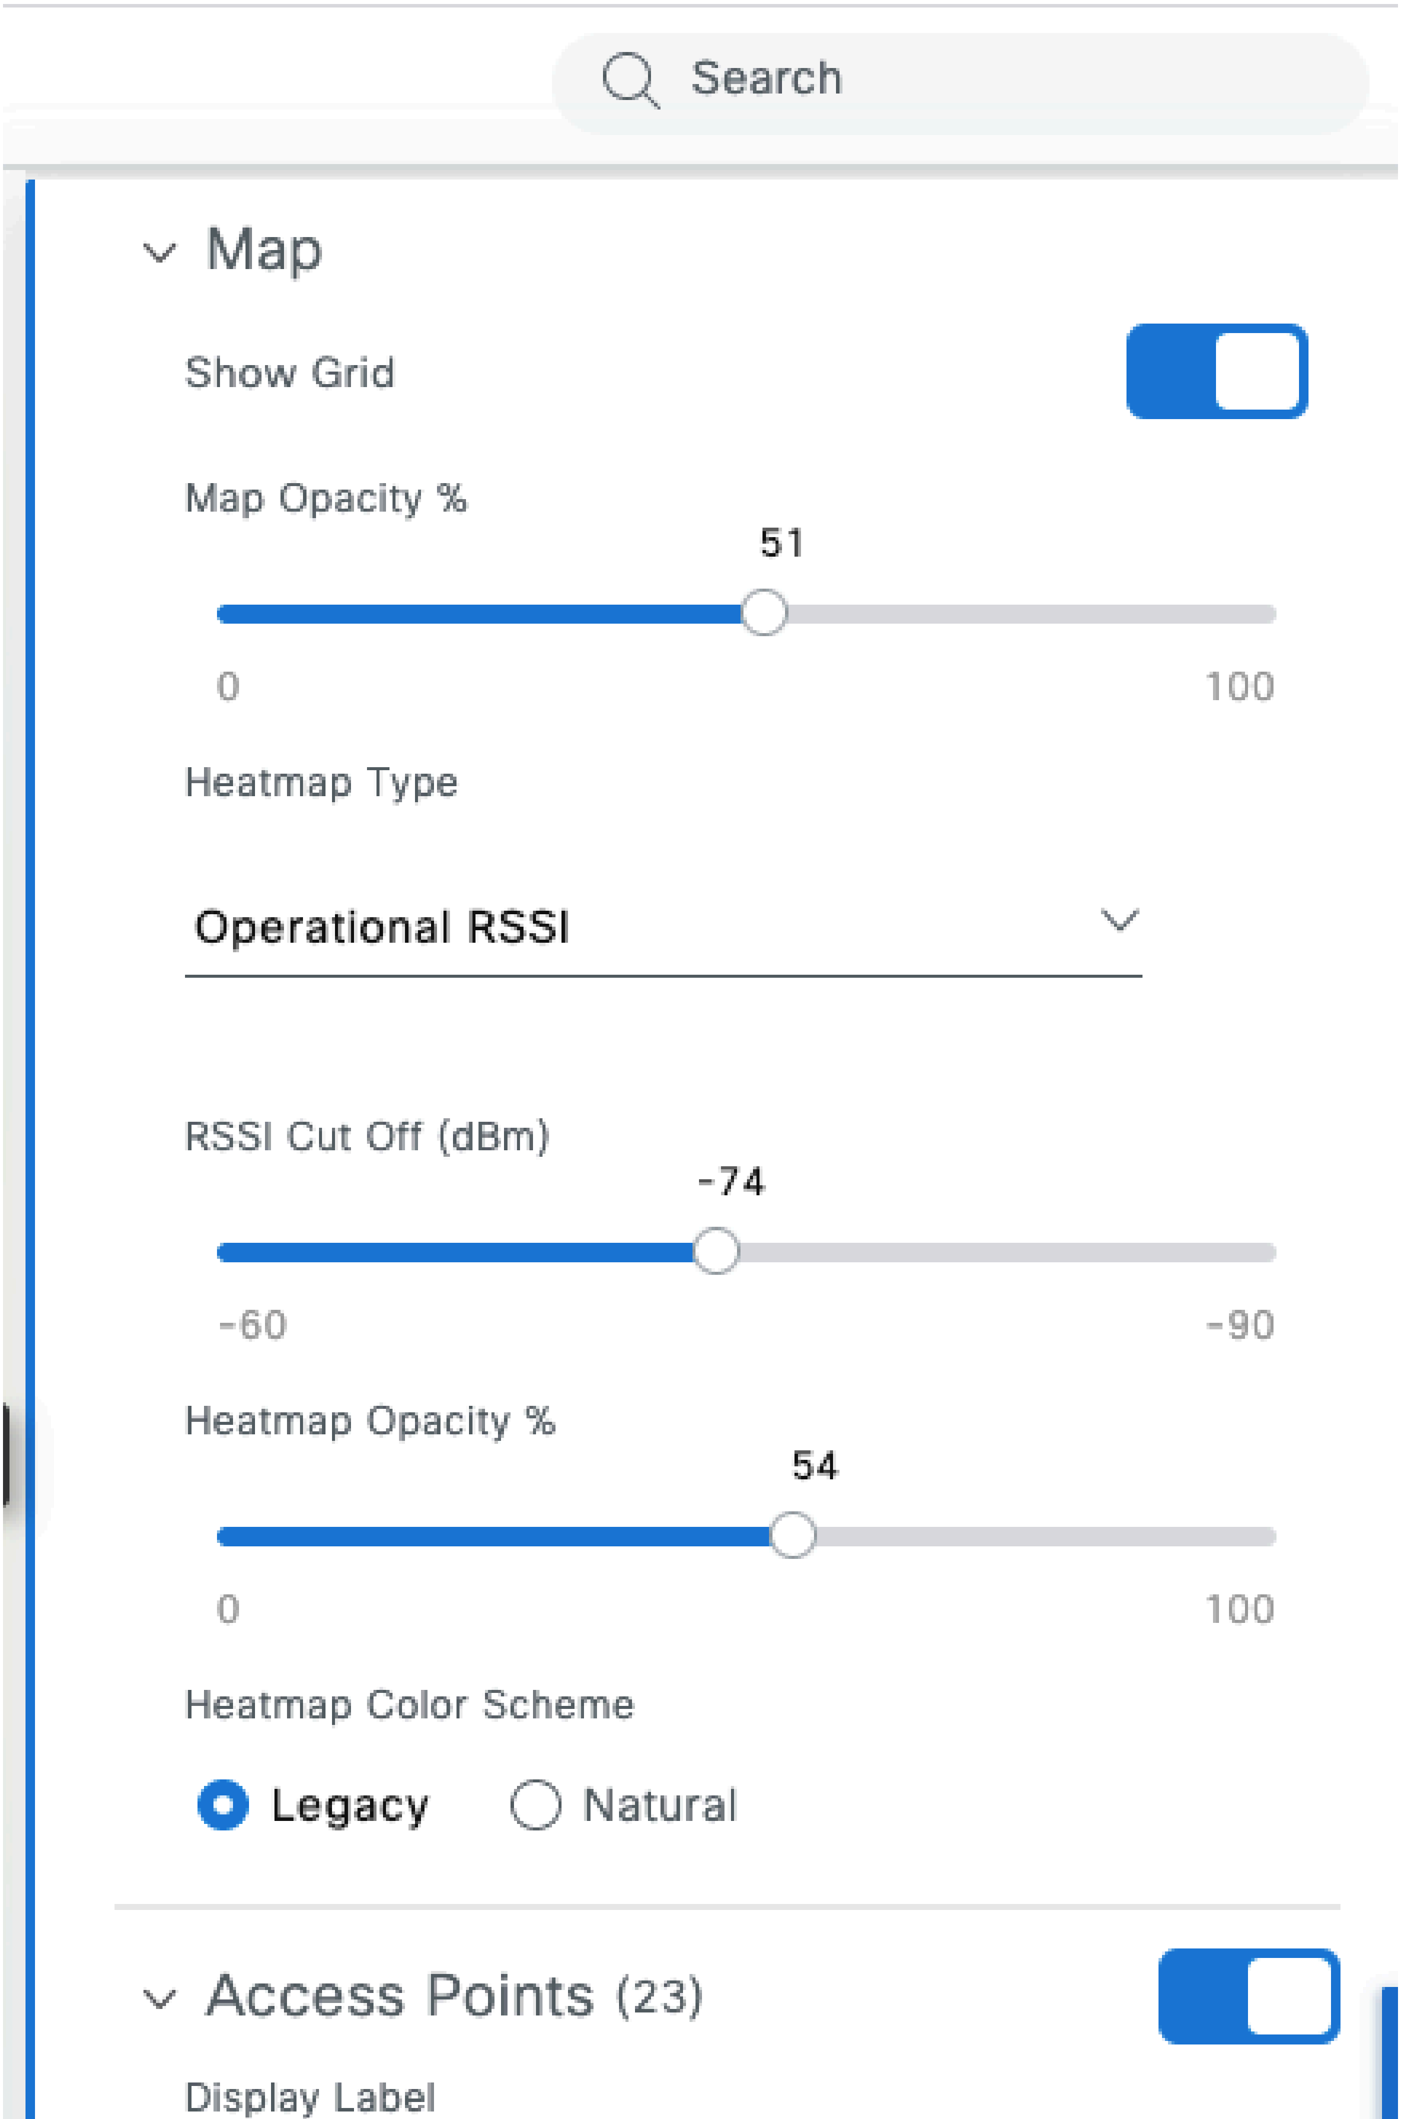

Search

Map

Show Grid

Map Opacity %

51

0 100

Heatmap Type

Operational RSSI

RSSI Cut Off (dBm)

-74

-60 -90

Heatmap Opacity %

54

0 100

Heatmap Color Scheme

● Legacy ○ Natural

Access Points (23)

Display Label

    - It is very tolerant to jitter as it buffers a few seconds or minutes of video beforehand. The pattern looks like a large file transfer for a brief period and then silence while the video plays from the buffer until the next pre-loading occurs.
- Voice call: this consumes a negligible amount of bandwidth but is extremely sensitive to latency and jitter.
    - This can potentially use QoS (Quality of Service) tagging and therefore face a different (prioritized) experience from best-effort traffic.
- Data: a social media application downloads data by bursts.
    - The amount varies based on the content and how fast the user scrolls.

A typical throughput testing application maximizes the protocol to achieve the highest transfer speed possible: it tries to book the medium and send as many data frames concatenated as possible. This does not represent the same usage type as real-life applications (other than file transfers) who are very bursty by nature.

Testing real-life applications mimics the user behaviors but makes it impossible to get actual metrics and numbers to compare. You only get a subjective feeling if the network is smooth or not.

For throughput testing, many websites are popular, and they give a decent picture of the end user experience as they test the whole bandwidth between the client and the internet. However, if you want to validate your wireless network separately from the Internet connection and routing and firewalling issues, it is recommended to use a dedicated throughput testing tool such as Iperf: https://community.cisco.com/t5/wireless-mobility-knowledge-base/iperf-test-for-measuring-the-throughput-speed-of-a-wlan-client/ta-p/3142047.

This tool allows specific testing between a client and a server that you place in your network. This allows you to move the server to specific places in the network and test the throughput over longer and longer network sections to validate each section. Start by placing the Iperf server on the same switch as the AP where your wireless client is in case of local switching or fabric-enabled wireless or on the same switch as the WLC (Wireless LAN Controller) (and in the client VLAN if possible) in case of central switching.

If you are using an anchor WLC, you must place the Iperf server on the same switch as the anchor WLC as that is where the traffic is terminated. It can be interesting sometimes to create a non-anchored WLAN (Wireless LAN) to see if the potentially disappointing throughput results are caused by the anchoring itself versus a non-anchored WLAN.

It does not really make sense to use several clients to do throughput testing at the same time. During throughput testing, it is expected that this single client use the entirety of the available channel airtime. Therefore, if two clients do a throughput test at the same time, they each see a result divided at least in half. If more clients are used, collisions start to occur in numbers and the results are not representative anymore.

There are multiple 3rd party tools to automate the network testing. Be aware that while you are testing the throughput in one area, you are effectively using all the airtime for the duration of the test, and it is then a bad idea to test the network too often as it is disruptive to other clients.

Troubleshooting a throughput problem

When you identify a throughput problem, there are several things that can be looked at to isolate the problem:

- Isolate if, before you start the test, the RF environment is already busy. The higher the Channel Utilization is (outside of the test), the lower the throughput test result becomes. If a Channel Utilization problem is identified, check if other APs are present in the same area on the same channel and reconsider your RF design. Reducing channel width, eliminating rogues, using different antennas with more focused coverage are all good options. Adding more APs is not always the best idea.

- Get an Over-The-Air capture of the throughput test and see if there are a lot of data retries at an 802.11 layer (in percentage of all data frames). A high number of retries means the RF environment is potentially the problem. Check also what data rates are used, potentially sub-optimal protocol or number of spatial streams are being used. A large data transfer is very characteristic in an over-the-air capture, you see dozens of data frames with the same source and destination and with an exceedingly small delta time between each other followed by a block ACK. If the transfer is characterized by a regular ACK after every data frame, or a lot of request-to-send/clear-to-send, the low throughput can be easily explained.
- Verify if the throughput problem happens with all security types on the WLAN. Sometimes, a specific security incompatibility between the client and the AP can lead to poor throughput.