Proactive Network Monitoring using Automated Anomaly Detection

Timely identification of deviations in network/AP/client behavior is critical for ensuring uninterrupted connectivity and good quality of experience on enterprise Wi-Fi networks. Anomaly detection enables network administrators to not only to proactively identify and alleviate the immediate issues in the network, but also initiate long term plans for network expansion and its dynamic usage over time.

Arista's CV-CUE offers a robust anomaly detection mechanism to assist network administrators in identifying deviations in performance metrics of APs and clients on a continuous basis. A performance baseline is required in the network for :

- Determining the "normal" operating level of network and devices, thus providing a way to validate the current network status against prescribed performance standards.
- Tracking trends in key performance metrics so that network administrators can take preventive action before the performance degrades to unacceptable levels.

This application note describes the baselining and anomaly detection process in Arista's CV-CUE.

Table of Contents

The Baselining Process

Identifying anomalous behavior against static thresholds across the network is not the recommended practice since the performance of different sections/locations of a Wi-Fi network are affected by their respective deployment scenario, prevalent radio channel conditions, client/AP capabilities and application usage patterns. Instead of static thresholds, CV-CUE maintains a moving average baseline for every metric and highlights anomalies when the metrics deviate from their respective baseline values beyond a dynamically calculated range. Dynamic baselining adjusts or adapts to the steady state behavior of the clients/APs as per the variations in network, traffic and usage scenarios. This reduces, to a great extent, false positives that result in alarm fatigue and false negatives that result in inaction.

CV-CUE captures and stores run-time metrics of both clients and access points at 15 minute resolution. The clients and APs metrics are stored for up to a month in cloud deployments and upto a week in on-prem deployments. This baselining facilitates anomaly detection whereby clients/APs deviating significantly from the baselines of respective metrics are flagged as 'unhappy' or 'bad' and are reflected on the baseline charts as anomalies. For details on the frequency of reporting of the various metrics and duration of retention in CV-CUE, refer Data Reporting and Retention.

Weighted Moving Average

A robust weighted moving average (WMA) baseline is maintained by CV-CUE for all the captured performance metrics. The weighted moving average μ_t is evaluated by capturing sufficient historical context so that it is robust to rapid variations in the new values of the metric. This allows the network administrator to observe trends in the metric over longer time durations. For example, the baseline plot of a client's average data rate over a period of one day is shown in the Figure 1 below. The purple color dots in the plot represent the average value of data rate over the sampling interval of 15 minutes. The light blue color line running through the dots is the weighted moving average of the data rate over time.

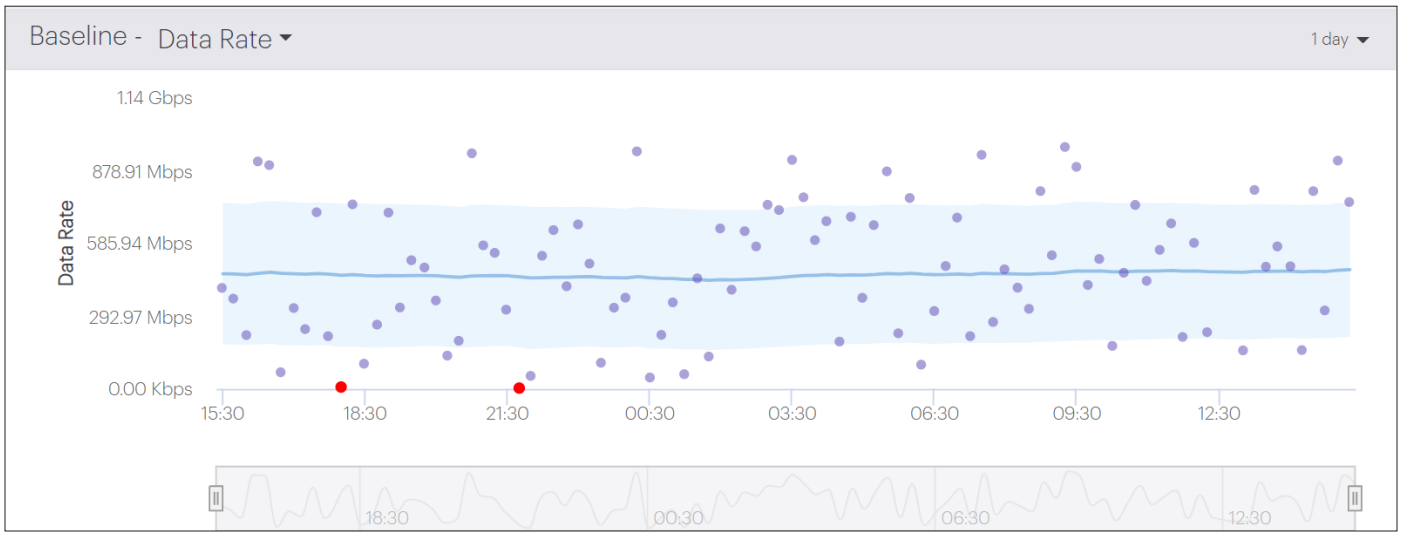

Figure 1: Baseline plot of client data rate

Blue line - Weighted Moving Average; Light blue shaded region - Weighted Moving Standard Deviation; Purple dots - Actual data rate averaged over 15 minute intervals; Red dots - Anomalies in data rate w.r.t baseline

Weighted Moving Standard Deviation:
The key statistic used for identifying anomalies is the weighted moving standard deviation (WMSD), σ_t. While μ_t reflects a metric's normal behavior, a positive or negative deviation by more than $2\sigma_t$ from μ_t is considered as anomalous behavior. The light blue shaded region in the baseline plot represents the range of σ_t. Anomalies are identified by the red dots.

Performance Thresholds vs. Anomalies

CV-CUE maintains baselines for the key RF parameters: data rate, retry rate and signal strength (RSSI) for every client. A client is considered as performing "normal" if the values of these metrics are within the thresholds set in CV-CUE. The respective default thresholds are listed in the table below.

Performance metric	Default threshold	Threshold type
RSSI	-65 dBm	Lower limit
Data rate	24 Mbps	Lower limit
Retry rate	30%	Upper limit

Even though a performance metric is within the thresholds set in CV-CUE, it will still be flagged as an anomaly if its value lies beyond the $2\sigma_t$ region of the baseline. Since the baseline reflects the steady state behavior of the metric in the prevailing RF conditions and network configuration, CV-CUE flags deviations from the baseline as anomalies, rather than hard thresholds. For example, in the Figure 2 below, the retry rate of the client is 14.3 %, which is well within the default threshold limit of 30% set in CV-CUE. However, it is considered as anomalous behavior since it's value is beyond $2\sigma_t$ from the baseline of 3.2 %.

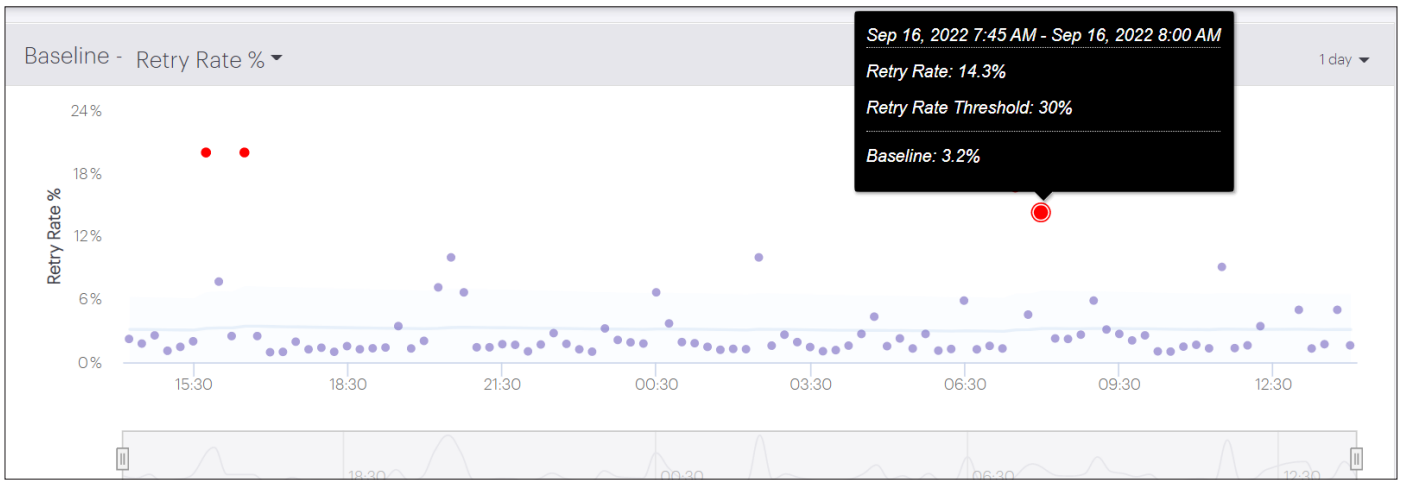

Figure 2: Retry rate -Performance threshold vs Anomaly

Connectivity Anomalies

The CV-CUE connectivity dashboard provides a live summary of connectivity performance as well as the historical baseline of connectivity failures at network/location/floor level. The failures accounted for include authentication, association or network connectivity (DHCP/DNS) failures. Each anomaly can be further drilled down to examine client, AP and network metrics and ascertain appropriate root cause(s). The failures can be analyzed per SSID and per band to gain insights into potential configuration or RF issues. The connectivity failures baseline is maintained at the AP and network levels, and is useful in observing trends in connectivity anomalies for up to a duration of one month.

Use case: Connectivity failures due to high contention

The band-wise baselines depict connectivity failures at the AP/network in the 2.4GHz, 5GHz and 6GHz bands. For example the figure below depicts connectivity failures at an AP in the 2.4GHz and 5GHz bands. The incidence of anomalies is higher in the 2.4GHz band as compared to the 5GHz band, owing to higher contention in that band as can be seen from the Spectrum Occupancy charts (Monitor->WiFi->Access Points-><AP name>) in both the bands in the subsequent Figure 3.1 and Figure 3.2. The Spectrum Occupancy information is stored in CV-CUE for a duration of 12 hours. In this case, on an average, there are 75-90 active nodes in the 2.4GHz band as compared to 30-50 nodes in the 5GHz UNII-2 band. A possible remediation of this anomaly could be configuring the AP's Smart Steering feature to steer all 5GHz capable clients to that band.

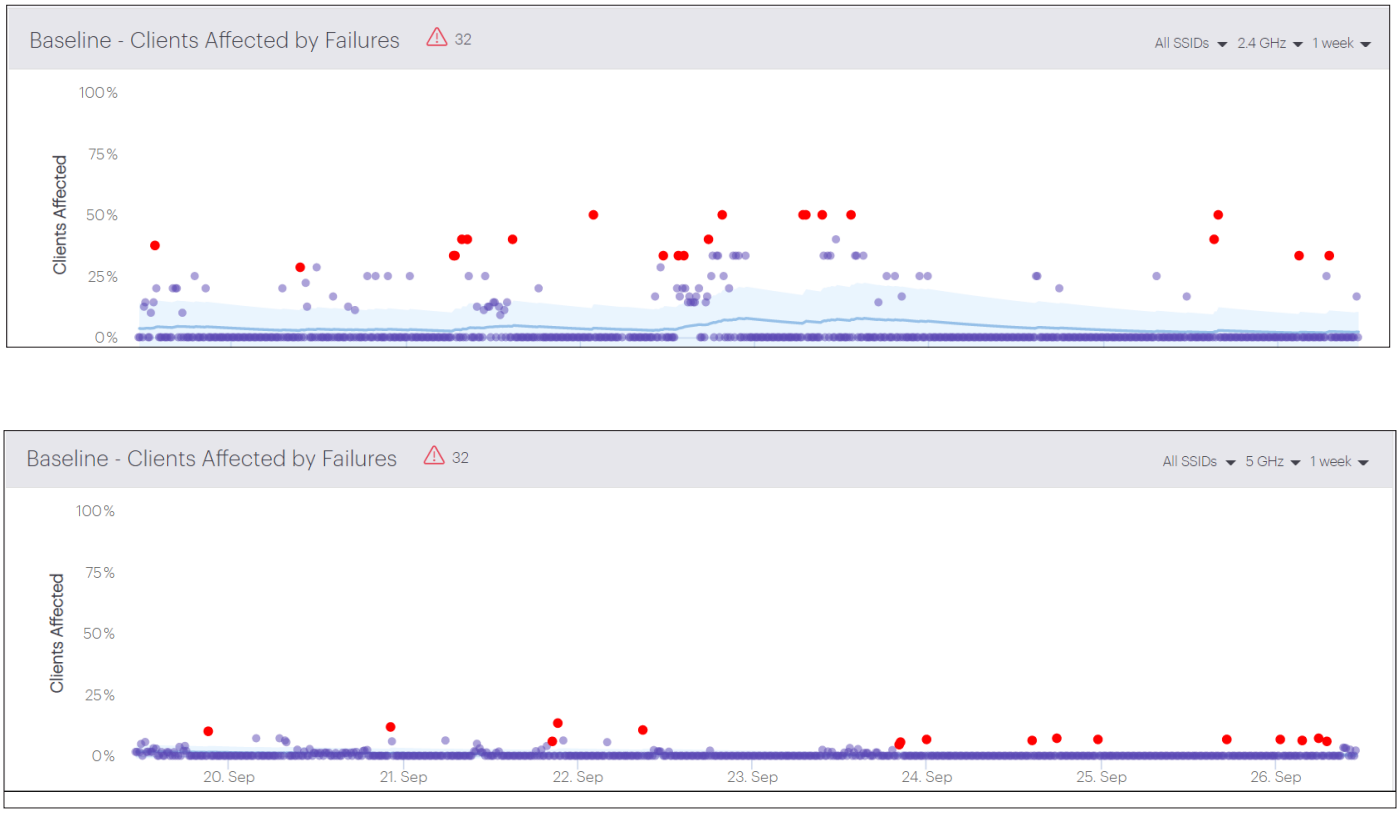

Figure 3.1: Comparison of Connectivity Failures in 2.4GHz and 5GHz bands at an AP

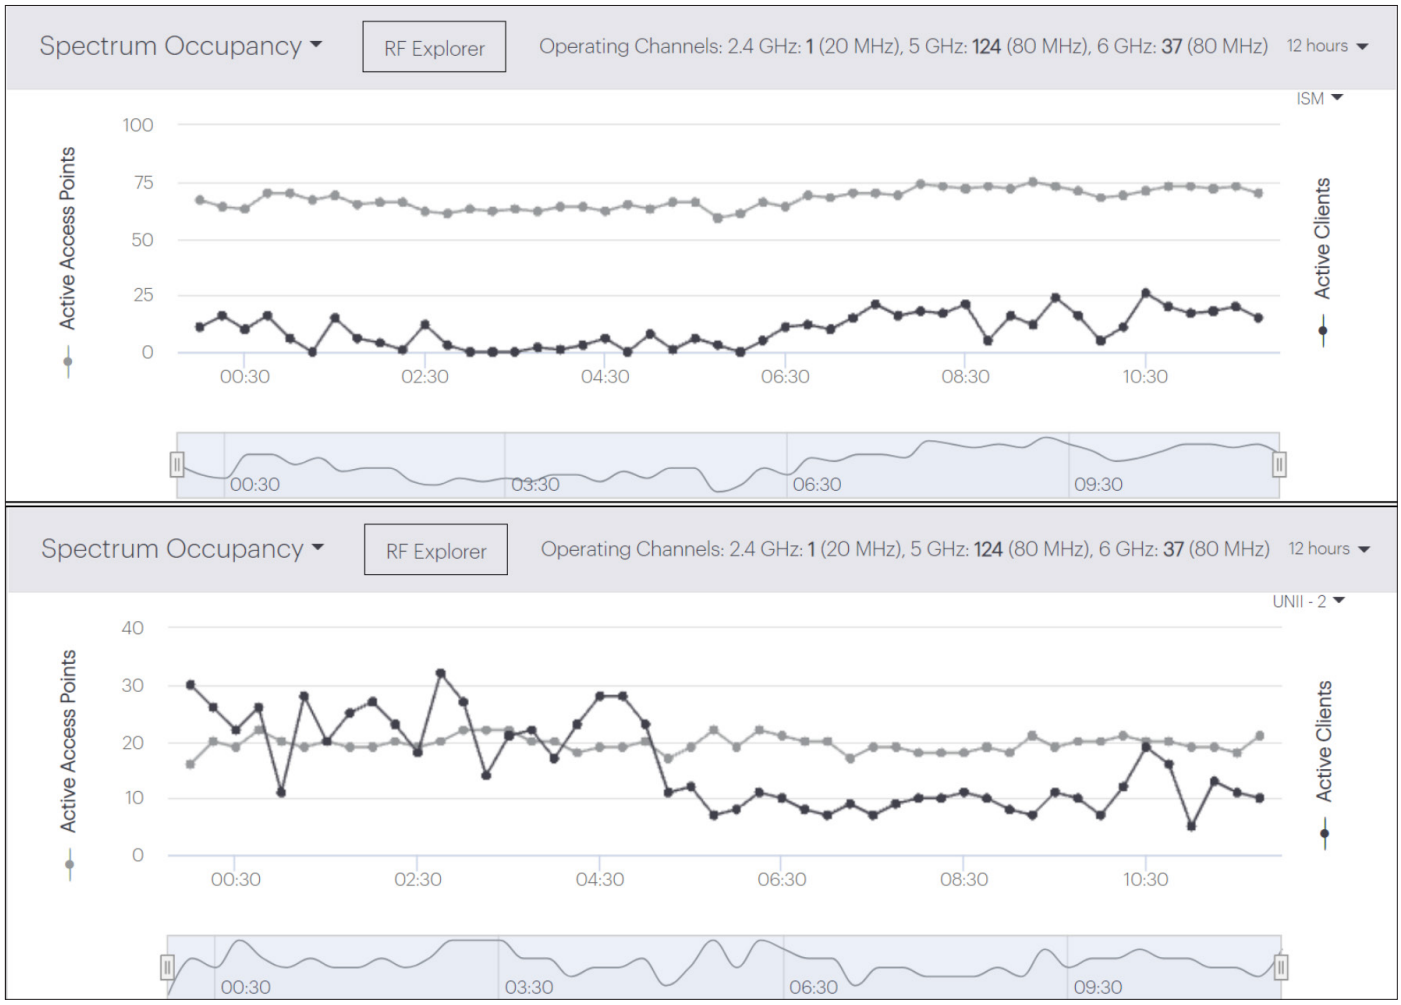

Figure 3.2: Spectrum Occupancy Top: 2.4GHz and Bottom: 5GHz bands

Use case: Fast roaming failure

The list of all the clients affected by an anomaly (red dot) can be seen by clicking on the dot. The reason for connectivity failure for any specific client can be viewed from the Client Events table, accessible by clicking on the respective client from the list of Client Connections. An example scenario is depicted in Figure 4 below where the second client in the list of affected clients is seen to have encountered a fast roaming failure due to PMKR1 mismatch. CV-CUE stores all client events up to a week, which can be used to drill down the root causes of connectivity anomalies.

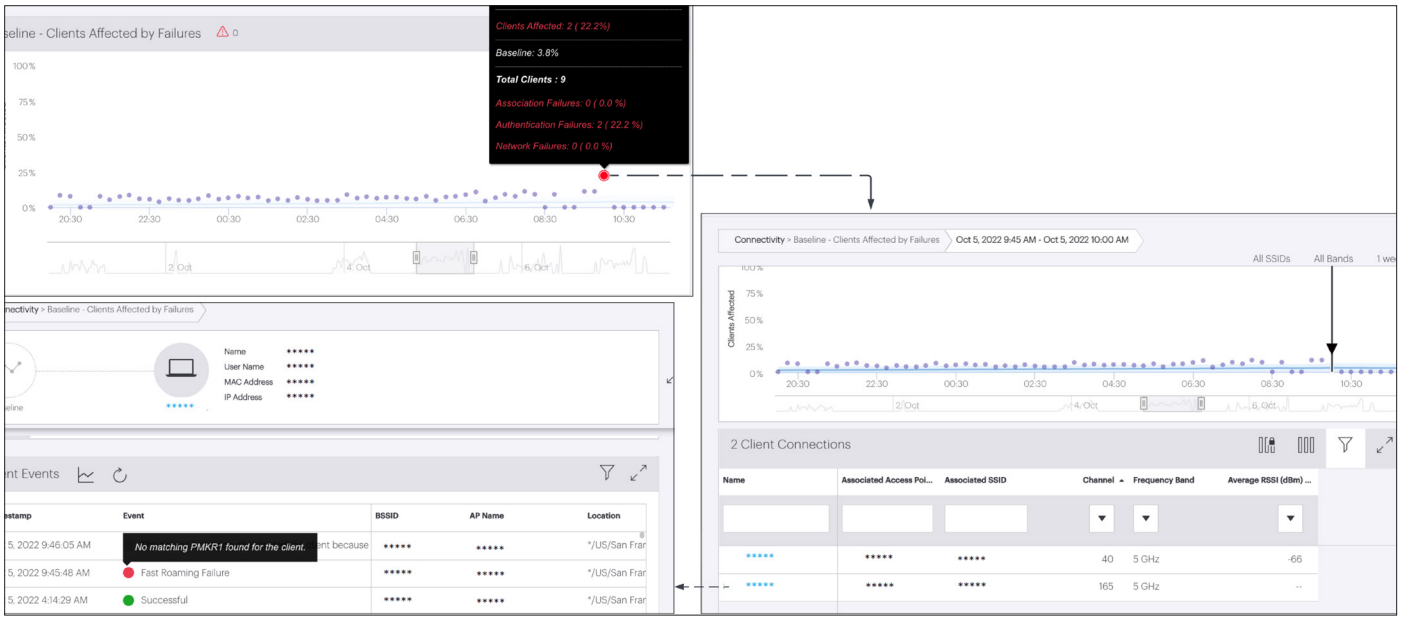

Figure 4: Connectivity Failure Anomaly - Fast Roaming Failure

Performance Anomalies:

Performance anomalies are flagged on the 'Clients Affected by Poor Performance' baseline on the Performance dashboard, whenever the total percentage of clients affected by low RSSI, high retry rate and low data rate issues deviate from the baseline.

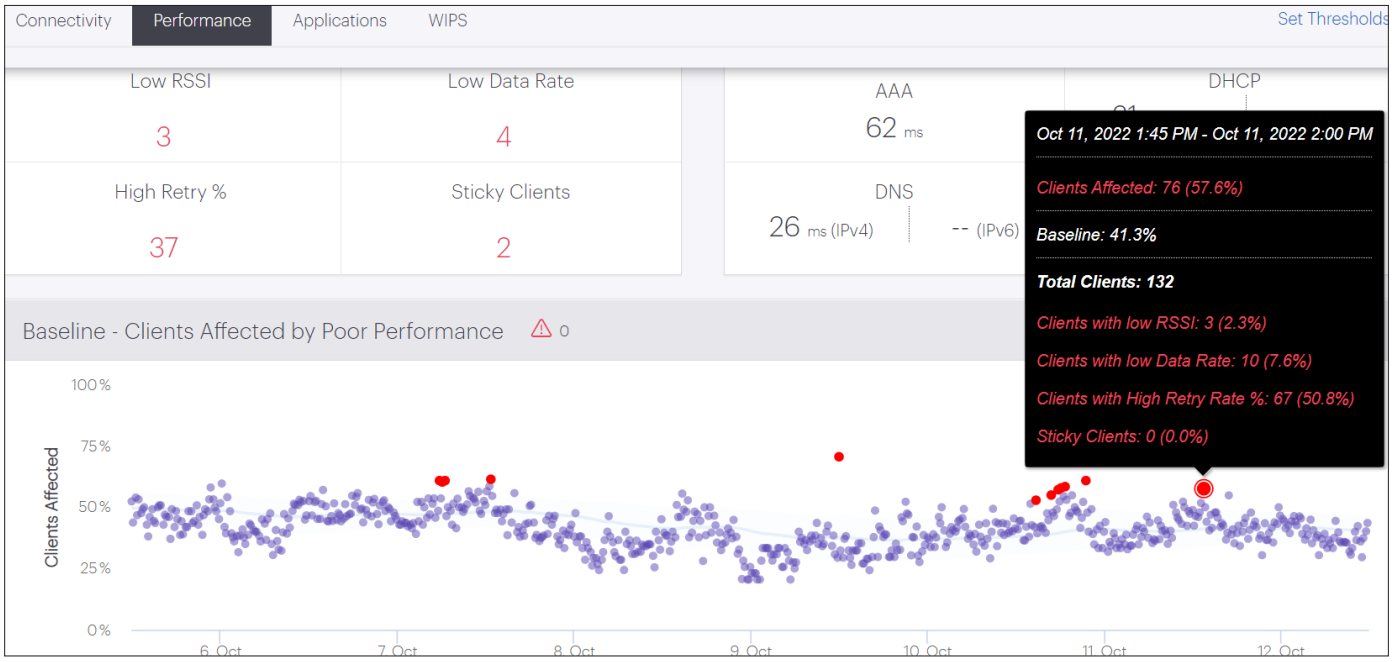

Figure 5: Connectivity Failure Anomaly - Fast Roaming Failure

Top: Real time - Number of clients performing poorly & network latencies

Bottom: Baseline of clients affected by poor performance

A useful feature provided in CV-CUE is the real time categorization of poorly performing clients according to their most common characteristics viz., associated SSID, frequency band, capability, OS, vendor, stickiness etc. This categorization helps in drilling down the root causes of the clients/APs anomalous behavior.

Use case: High retry rate in 5GHz band

An example scenario is shown below where all clients with high retry rate are in the 5GHz band. A drill down to the baseline of the worst performing client in this category shows that the client's retry rate is 41.54%, much higher than its baseline of 13.1%.

The associated AP is on channel 157 (UNII-3 band), and the Spectrum Occupancy chart of the UNII-3 band shows 47 nodes using that band (see below Figure 6.1 and Figure 6.2). A further drill down to the channel wise occupancy of the UNII-3 band using the RF Explorer view shows that most of these nodes are on channel 157, leading to high contention and consequently high retry rate at the AP. This anomaly can be resolved by enabling Dynamic Channel Selection feature at the AP location so that it chooses to move to a less congested channel upon detecting high interference levels in its current channel. The channel information in the RF Explorer is valid for a duration of 15 minutes, while the Spectrum Occupancy information is available for a duration of 12 hours.

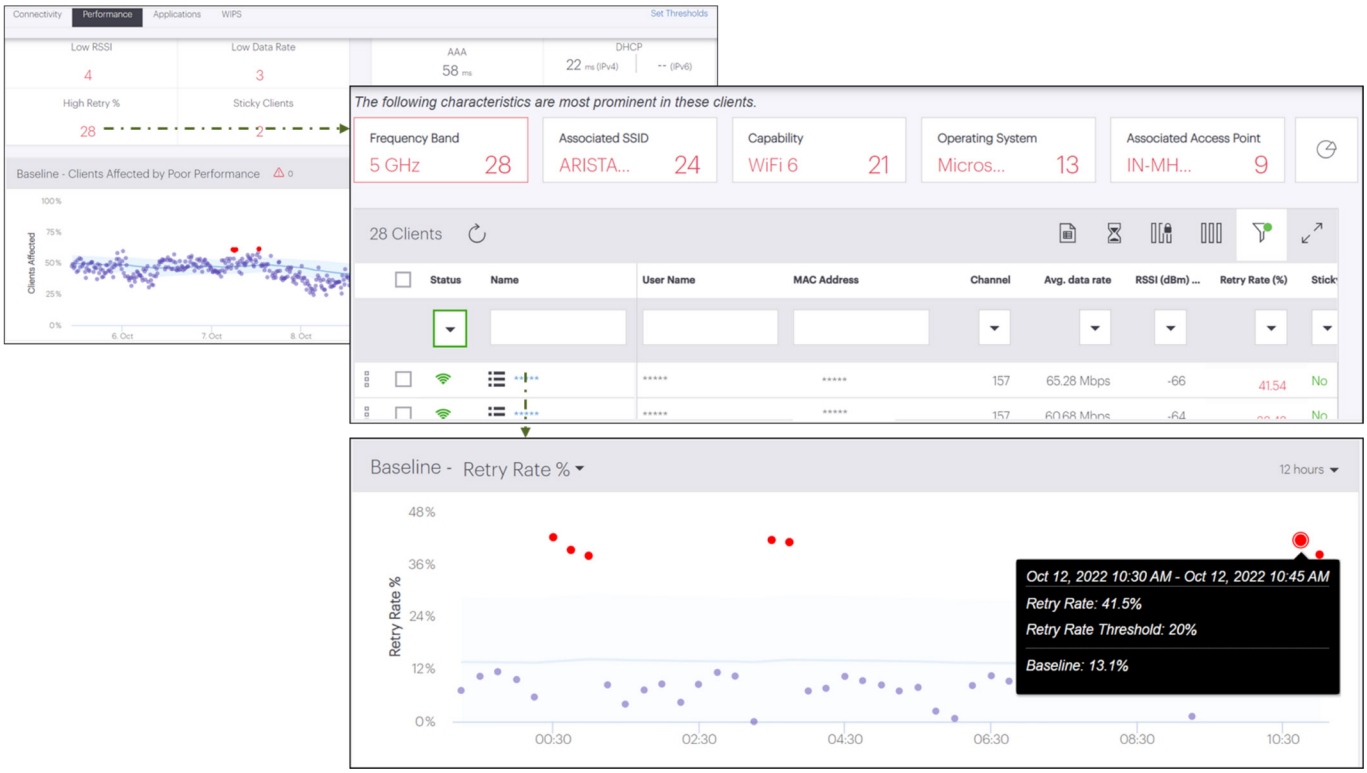

Figure 6.1: High Retry Rate in 5GHz band - Client drill down

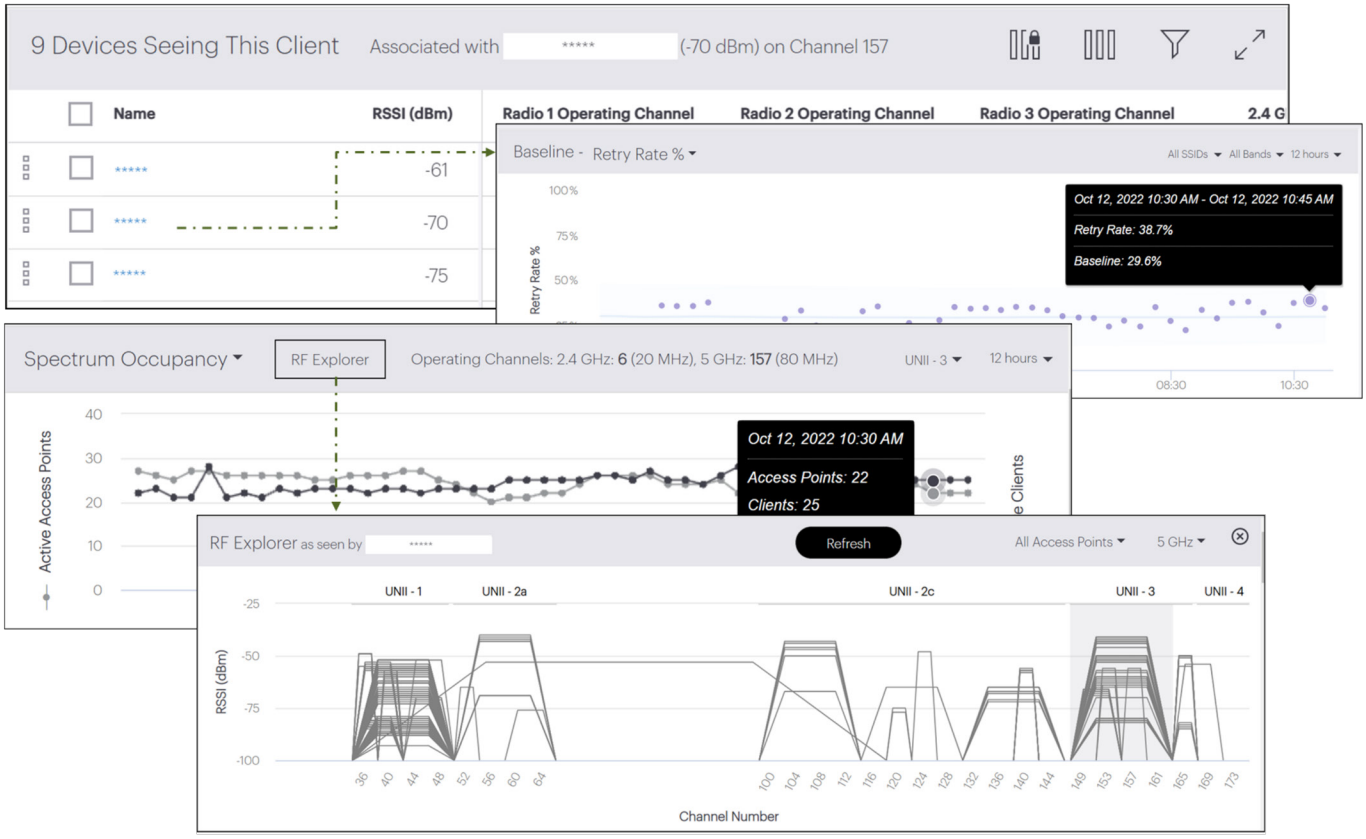

Figure 6.2: High Retry Rate in 5GHz band - AP drill down

Use case: Low data rate due to multi-band client operating in 2.4 GHz

The peak data rate experienced by a client depends on the client's capabilities. For example, a client will be able to experience higher data rates in the 5 GHz band, if it is capable of operating in that band. However, if the client is connected to an AP on the 2.4 GHz channel, its peak data rate will be limited by that offered in the 2.4 GHz band. A scenario reflecting this situation is depicted in the figure below. The drill down of the client with low data rate from the Performance dashboard shows anomalous behavior with an average data rate of approximately 17 Mbps. The properties of this client examined from CV-CUE 's Monitor-> WiFi-> Clients show that it is capable of operating in the 5 GHz band also. Configuring the Smart Steering feature at the AP will steer this client to 5 GHz band and resolve the low data rate issue.

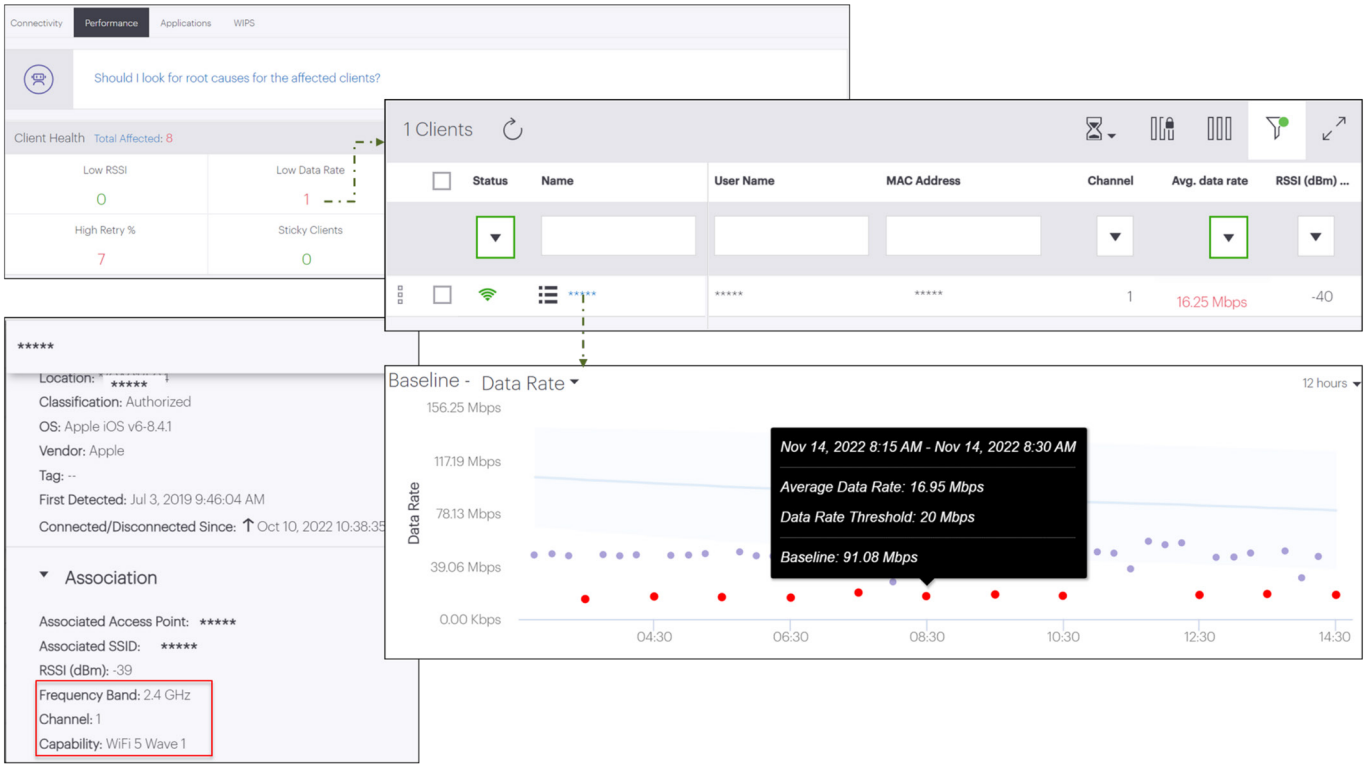

Figure 7: Low Data Rate due to multi-band capable client operating in 2.4 GHz band

Application Performance Anomalies

The quality of end user experience at the application level is often difficult to predict due to the differences in the characteristics of packet traffic generated by each of the applications and the dynamics of the TCP protocol. The end user experience is also impacted by the number of concurrent users and the application server performance. CV-CUE evaluates the end user quality of experience for AVC (Google Meet, Zoom, Skype) and non-AVC (Gmail, web browsing, Social media) applications using Arista's Application QoE solution powered by machine learning algorithms.

Use case: Time-of-day trends

One of the important uses of baselines is trending of the metrics over time. The Figure 8 below captures the application experience measured over a period of one week. We can observe that the application experience is poor during a specific time interval every day, which is probably the busy hour period in the enterprise. The list of active application sessions at the time of occurrence of the anomaly shows that SSL and BugsBy are the most severely impacted applications.

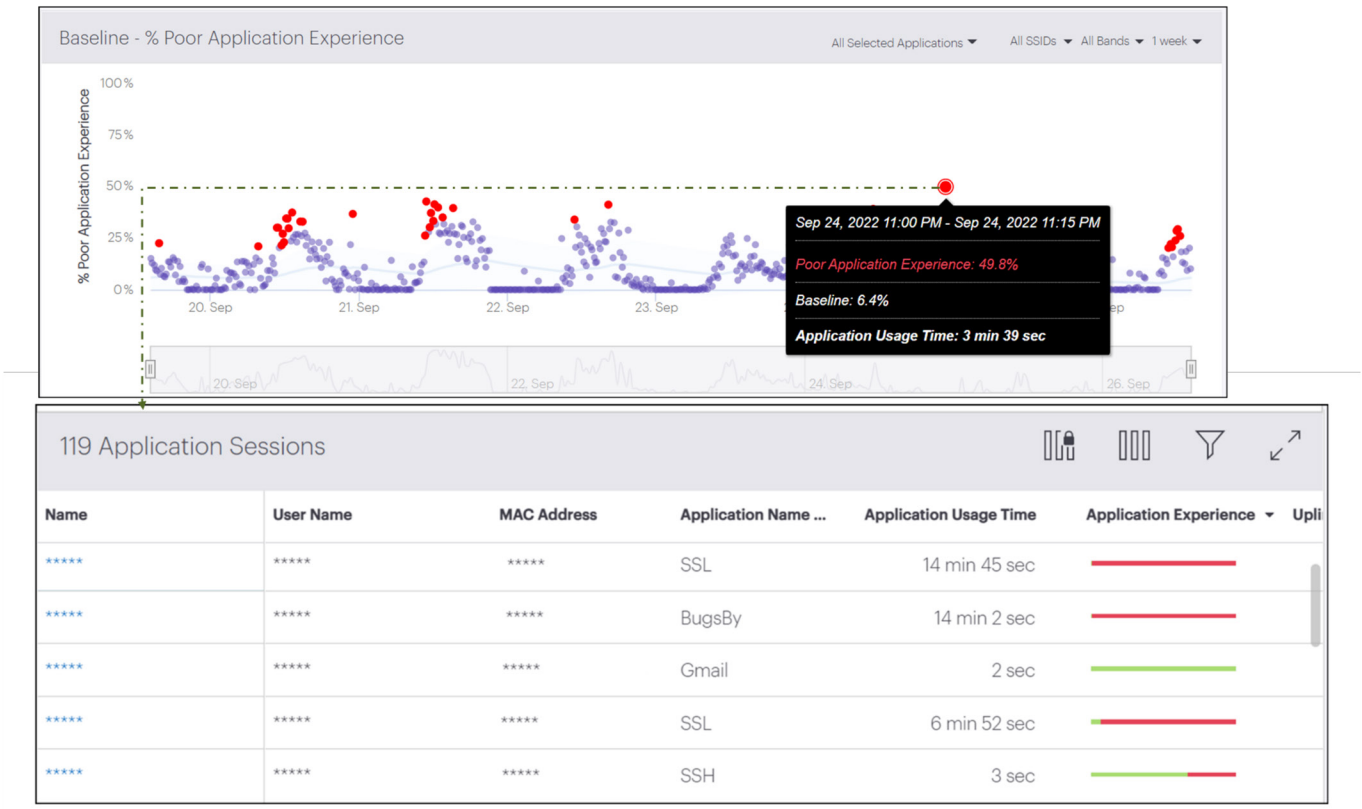

Figure 8: Time-of-day trend in poor application experience by clients

Use case: Poor application experience due to high retry rate

The figure below depicts the drill down of a poor application performance anomaly. CV-CUE lists all 101 active application sessions at the time of occurrence of the anomaly upon clicking on the anomaly.

A more detailed analysis of the first client experiencing poor quality of BugsBy application reveals that the client's retry rate has been consistently high over a 12 hour period, which might be the reason its application experience is poor. The root cause of the client's high retry rate can be drilled down further as described in the Performance Anomalies section.

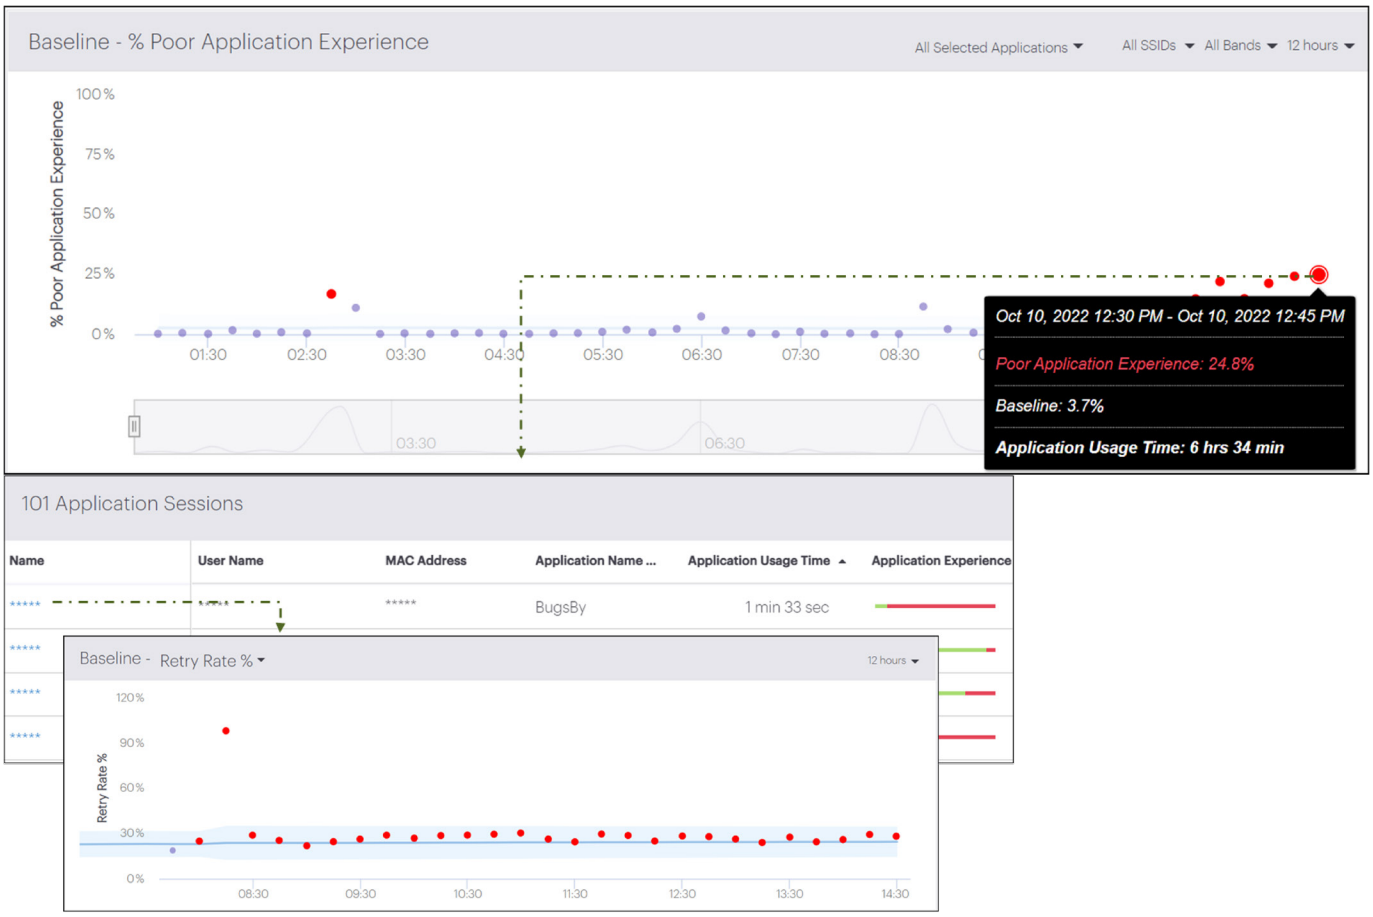

Figure 9: Poor application experience - High retry rate

What is not an anomaly?

Though anomalies are deviations from the baseline of a metric, they do not always necessarily reflect poor performance and have to be analyzed in the context of the network scenario. One such example scenario is discussed below.

Use case: False positives of average data rate at AP

In the first plot in the Figure 11 below, the average data rate of the AP during the time interval 06:30-08:30 is much below the baseline and is depicted as an anomaly. However, a look at the number of clients served by the AP (second plot) reveals that there is just one connected client generating low traffic volume during that period. The average data rate of the AP is low because it is serving a single client, and not because of any performance issues. This case demonstrates that anomalies should not be taken *prima facie*, but rather should be deciphered in the context of the network scenario.

Figure 11: Example of nom-anomalous behaviour.

Average data rate at AP is low between 06:30-08:30 because the number of connected clients is also low during this period

Use case: Seasonal/Time-of-day phenomena

One of the key factors that influences network performance is user behavior. For example, in a University campus deployment, the number of disconnections was observed to be high at around 8:00 a.m and 4:00 p.m during any given day. This can be attributed to the fact that a large number of students disconnect their mobiles before boarding lifts while entering/ leaving the campus at these specific times. The disconnections are "normal" owing to the user behavior and not the result of any connectivity issues in the network.

References

Data Reporting and Retention

Application QoE - Monitoring end user experience of enterprise applications

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062