
WHITE PAPER



The Aruba AIOps Advantage

DATA SCIENCE, DOMAIN EXPERTISE
AND PROVEN RESULTS INCREASE THE
EFFICIENCY OF NETWORK OPERATIONS

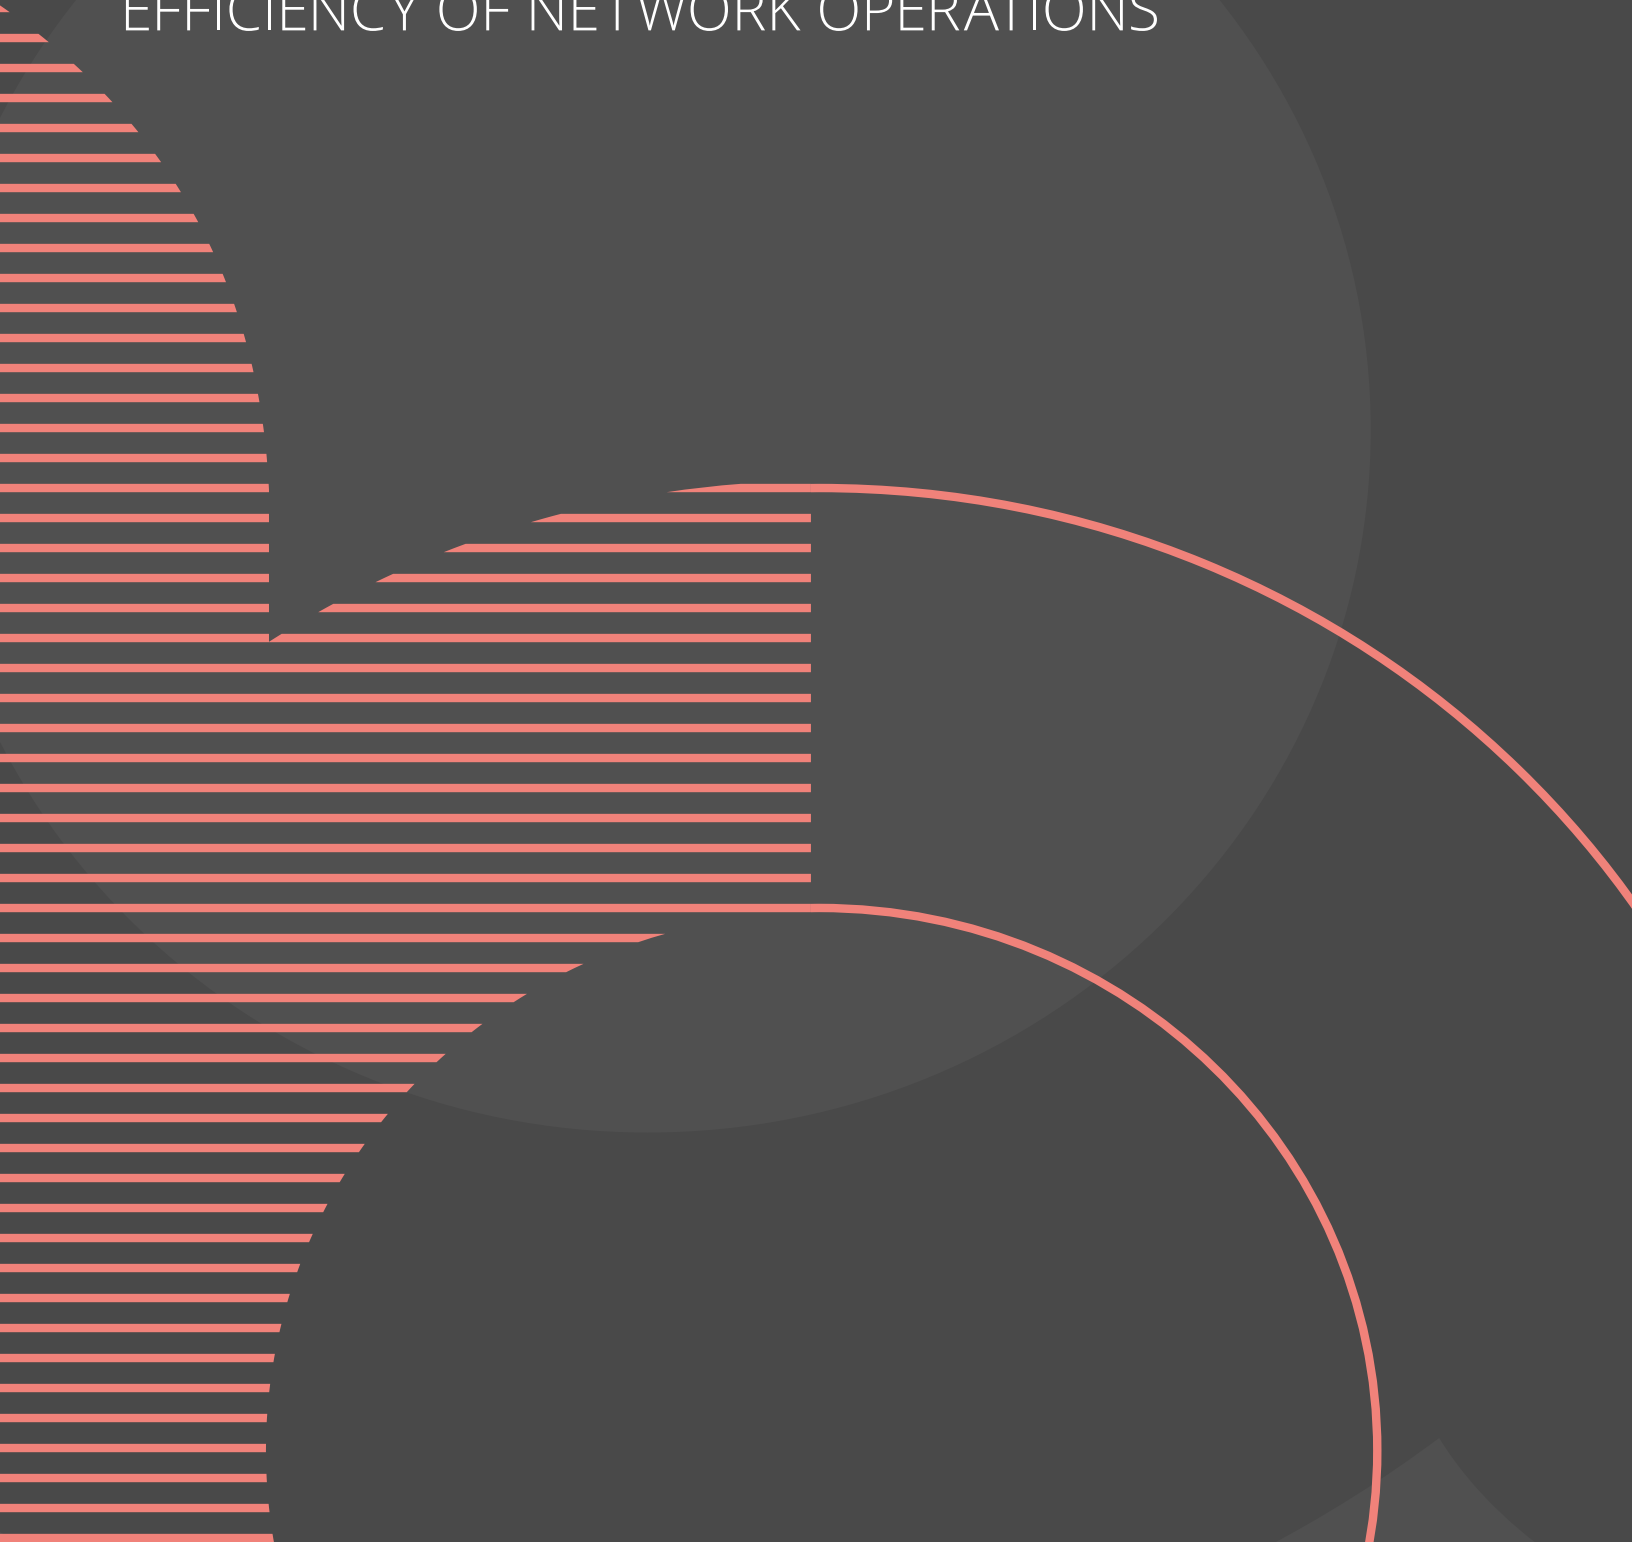


TABLE OF CONTENTS

| | |
|---|----|
| THE ARUBA AI FRAMEWORK | 3 |
| LEVERAGING NETWORK ENVIRONMENT FACTORS TO DELIVER RELIABLE AUTOMATION | 4 |
| USING PEER GROUPS TO OPTIMIZE NETWORK SETTINGS | 6 |
| PUTTING IT ALL TOGETHER: AI INSIGHTS DELIVER ACTIONABLE RECOMMENDATIONS | 8 |
| SUMMARY. AI + AUTOMATION = AIOps YOU CAN RELY ON | 10 |



Modern networking requirements are extremely complex due to increased network size, volume of traffic, and diversity of devices and applications. Manually configuring and operating these networks has become time-consuming, error-prone, and difficult to manage due to the explosion of IoT devices and distributed connectivity from the edge to the branch. Network operators need better insights, and intelligent automation powered by AI that is aware of all aspects needed to maintain a network today.

The Aruba AIOps advantage starts with great AI. We collect over 31 Terabytes of useful data per day, from tens of thousands of installations that range from small stores and offices, to large campuses across all verticals and geographies to form our data lake. Our deep understanding of networking and security technology, and strong team of data scientists then deliver the insights needed to quickly preempt or resolve issues in a fraction of the time required in the past.

This investment in data collection and AI expertise cannot be replicated in months or even years. Vendors with a small customer base are promising insights based on limited data that lacks variety which means the AI results are unreliable. On day one when using Aruba AI, your network is dynamically monitored to form a baseline of operation – from connectivity to the quality of connections. Other vendors require that you manually set service level expectations and reset these values as your network changes.

We call this combination of AI and automation, actionable AIOps. It's our belief that the next generation of networking management solutions must include built-in intelligence that can only be realized with reliable data and proven results that start with thousands of varied customer installations. These new solutions also need to be comprehensive enough to deliver precision insights to quickly solve individual customer issues, and leverage the data lake to provide useful optimization guidance based on the network insights of others.

This white paper describes how Aruba ESP includes AIOps and the essential building blocks needed to deploy a self-optimizing, AI-powered network and how these building blocks work together to reliably solve today's evolving networking challenges

AI and automation quickly deliver ROI

A retail customer was recently seeing what they believed was a significant amount of “passerby” traffic that was dragging down network performance. But they didn't have the data or guidance needed to manually diagnose and fix the problem. When they installed Aruba Central with AI Insights, the root cause of the problem was immediately found. A recommended configuration change eliminated 98% of the unwanted passerby traffic, while maintaining connectivity for legitimate users and devices. The result: 25% more wireless capacity without requiring any additional hardware.

THE ARUBA AI FRAMEWORK

The **objective** of a self-optimizing network is to provide its users with stable connectivity and ability to satisfy any traffic demands so that they can experience the highest quality of service. In practice, this perceived user experience is very difficult to measure as no two sites are identical – from an environmental as well as user behavior perspective. To build, train and evaluate the AIOps we build, we use metrics that are directly measurable, such as throughput, latency, or resource efficiency.

A metric or numeric representation of how well the network is doing and how it compares to other networks is chosen per use case. We also identify the **factors that influence** the chosen metric and separate them into two types: those that the AI used is allowed to control (controllable factors), and those that are assumed as given (environmental factors). For the purpose of optimizing radio transmission parameters for the access points in a Wi-Fi network the controllable factors include RF channel bandwidth and RF transmission power levels. Because the Wi-Fi access point (AP) is physically able to operate at various settings based on these factors, we can choose to let the AI decide which settings are optimal.

Environmental factors include, but are not limited to, the spacing between adjacent access points in the deployment, the propagation characteristics of the RF signals depending on the building materials in the coverage area, and the RF characteristics of the client devices that are connecting. AI is not free to modify these factors (and they may vary naturally over time).



Equipped with the definitions of a measurable objective function and its controllable and environmental factors, we instruct the network to **continuously collect this data**, and store it in the cloud. For each network and each instant in time, the AI engine in the cloud is thus aware of how well the network is running, as well as its current settings of controllable factors and environmental values.

To evaluate performance fairly, we baseline each network against its **peers**, which is a set of other networks with the same environmental factors that belong to the same customer or to others. For example, if a network ranks in the 10th percentile of its peers, that is, it performs worse than 90% of its peers Aruba's AI has the ability to offer optimization guidance. Since the principal difference between the peer networks is the assignment of controllable factors, our AI algorithms can move this network up in the ranking by modifying the controllable factors.

The **optimal setting of controllable factors** is determined in one of two modes:

With supervised learning, the past data of all networks is distilled into an AI model that can predict the objective function value of a network with its current environmental factors with any possible setting of the controllable factors. Using these predictions, it is easy to select the best setting among all possible settings. In this mode, the AI exploits the knowledge embedded in the available data in order to optimize each network.

Alternatively, with active learning which is the second mode, the AI is allowed to acquire new knowledge that is not available in past data. The algorithm purposefully and cautiously sets the controllable factors to combinations of values that have never been tried before. In this mode, AI explores the space of opportunity to identify previously unknown optimization potential for each network.

The **key building blocks** of a self-driving network are a meaningful and measurable objective, its controllable and environmental influencing factors, a large and diverse collection of data, a definition of peer groups to compare against, and AI models and algorithms that automatically select the optimal assignment of the controllable factors in each operating condition.

Why data is so important

It's a well-known AI principle that the mathematical models learn to accomplish a specific task based on continuous training of the data they are intended to process. When Machine Learning models are presented meaningless or too little data to train the models, the results are unusable. For example, if a model has only seen data from small offices, the results are unreliable in larger, more complex environments. Imagine a model designed to recognize pictures of dogs that has only been trained using photos of brown dogs - when it sees a black dog it will fail. The same applies to your network. If the AI used to make configuration changes has not seen enough quality data to match your network, you cannot rely on it to deliver the results you truly need.

LEVERAGING NETWORK ENVIRONMENT FACTORS TO DELIVER RELIABLE AUTOMATION

Environmental factors of a network are a key element of AIOps: they capture characteristics of the network that are not controllable. For the purpose of the AI and the actions it can take, they are externally given parameters. The environmental factors define the extent of the playing field within which the AI is allowed to optimize.

To make the concept concrete, we'll now look at specific environmental factors, what they mean, and how they are measured. The examples are drawn from performance optimization automation for Wi-Fi networks, a core feature of Aruba Central AI Insights.

There are four distinct types of environmental factors that capture the physical environment, network layout, user mobility behavior, and traffic characteristics.

Factors representing a **physical environment** are based on the architecture, floorplan, and materials of a building,

seen through the lens of the Wi-Fi infrastructure. Firstly, the building size, measured in terms of square footage, is an important factor.

Ceiling types and height on the Wi-Fi signal (through-ceiling signal loss) where each AP is installed can also cause an effect. The combined effect of all building materials on signal propagation can be represented by the path loss exponent, which is the average rate at which the Wi-Fi signal strength decays.

Network layout factors are related to the way the infrastructure is deployed in a space. This includes Wi-Fi *access point density* and the Wi-Fi *hardware capabilities* of the access points. To capture the end user’s perspective, we measure the strength of the client device signals received by the access points, and aggregate these raw measurements over time, space, and clients to compute the typical uplink signal-to-noise ratio. Finally, the level of resource competition caused by other Wi-Fi equipment is captured as the density of *unmanaged* access points.

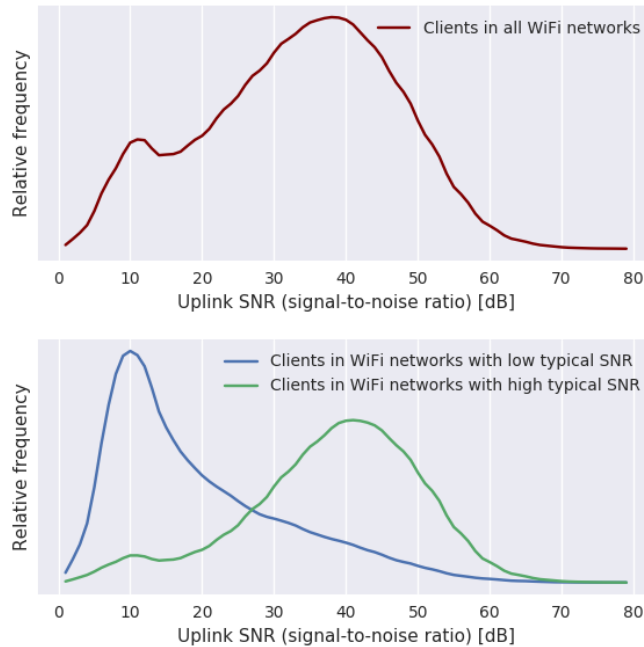
Factors that represent **user behavior** include peak client density, since networks are designed to handle the highest possible load and typical client dwell time, which quantifies the mobility behavior of the users. The population of client

devices can be characterized by the composition of client device types in categories such as laptops, phones, IoT, and legacy devices (that use outdated Wi-Fi protocols).

Finally, **traffic characteristics** capture the kind of applications running on the network. This includes typical *traffic volume per connected client*, as well as the *traffic composition across application types*. These features help our AI allocate network resources preferentially to those applications that directly impact the user experience.

The combination of all environmental factors defines the playing field for the AIOps that powers the self-driving network. Deployments with the same environmental factors have the same playing field, and form a natural peer group for cross-network benchmarks: if the AI plays on the same field, it should be able to achieve the same performance every time.

Aruba’s networking hardware is specially engineered to supply raw data for accurate use of environmental factors.



Across all networks, client to access point signal-to-noise ratio (uplink SNR) varies widely, and there are two distinct modes (upper diagram). Each network can be characterized by its typical uplink SNR. Grouping similar networks together separates the two modes, allowing like-for-like comparison between networks in the same group (lower diagram).

Our cloud architecture supports the design, implementation and validation of these environmental factors across all of our customer networks to ensure they are robustly computed. Aruba's AIOps builds on the data from our hardware and advantages of a cloud infrastructure to compute the environmental factors, translate them into meaningful peer groups, and run every single customer network at its optimal operating point.

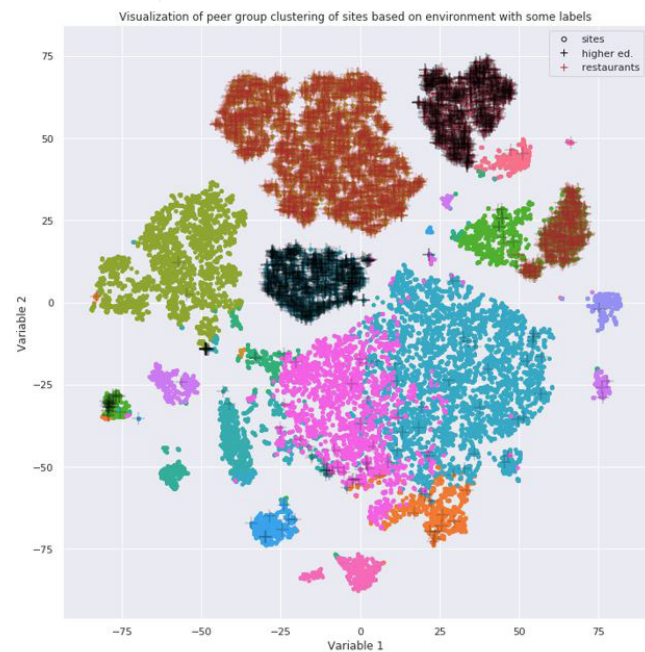
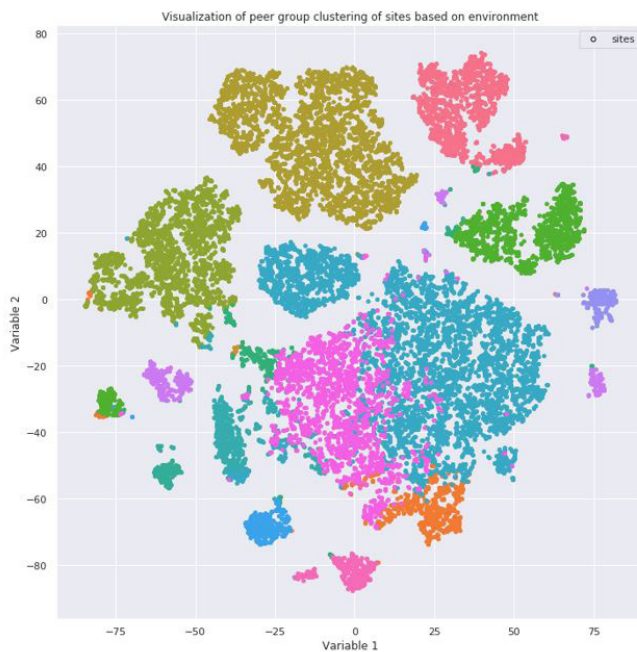
USING PEER GROUPS TO OPTIMIZE NETWORK SETTINGS

Now that we've described environmental factors that can be used to quantify the behavior and performance of different sites, we'll explain how and why we use those factors to cluster sites into peer groups.

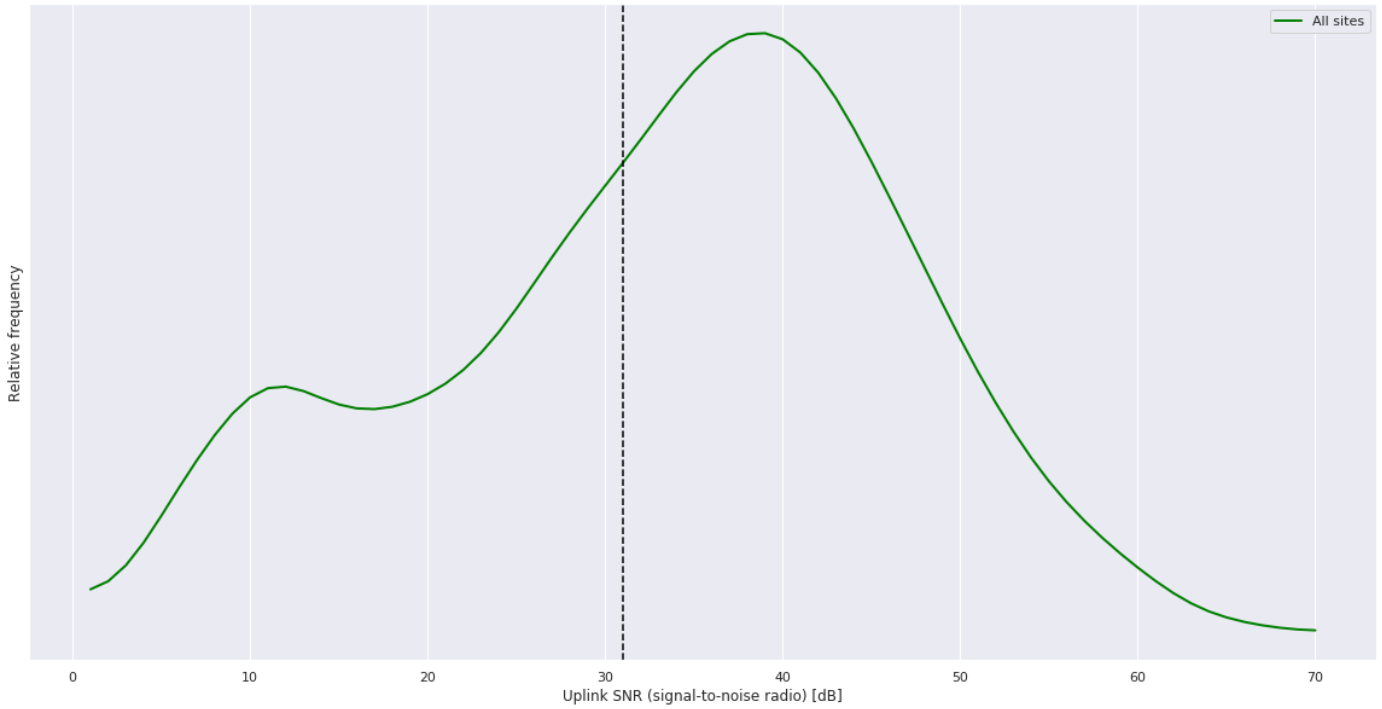
Because there are more mobile and IoT devices now than ever before and data traffic continues to grow it is no longer good enough to have just an ok network; it must also be optimized for specific use-cases. We'll talk about optimization in more detail in the next section. For this section let's discuss an important prerequisite called peer group definitions that helps our AI improve our customers' networking performance.

Peer groups at the most basic level are groups of networks, sites, or even APs (access points or switches), against which a given network is compared and benchmarked. It's possible to compare all networks against each other, but this is not a good idea. The operating environment of a local coffee shop is very different than that of a large grocery, or big box store. It makes sense to compare these environments, and, eventually, allow the AIOps solution to learn from them separately. Optimization guidance based on how a peer operates makes logical sense. The following images are a 2D projection of one instance of AIOps using environmental peer-group clustering for a subset of our sites.

The left image shows a grouping of clusters, which means that our sites do display distinct characteristics. In the right image, we've introduced market and customer verticality by adding black and brown plus symbols to denote some labelled colleges and restaurants, respectively. As clustering gathers more data the fact that the black and brown symbols group together into a few clusters bears that different types of sites do exhibit different environmental factors that are important. Colleges and restaurants do not perform the same even though they may be using the same infrastructure.

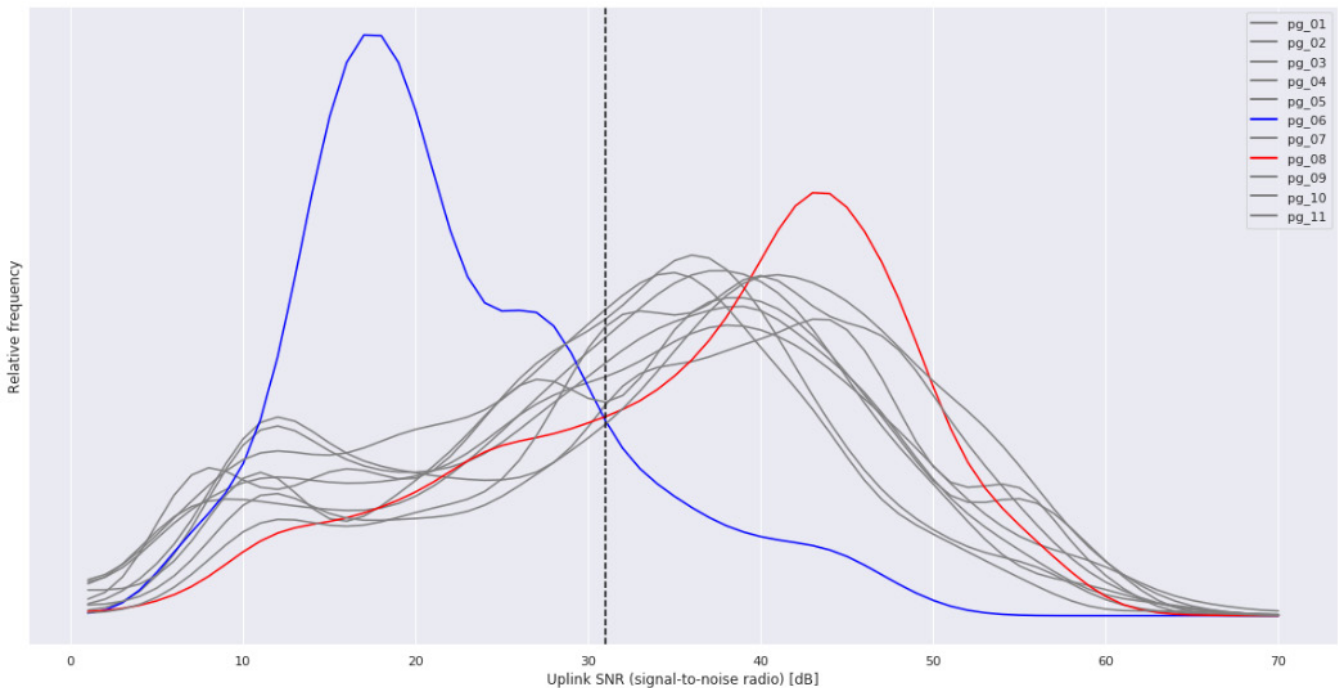


Being able to group similar environments together is great, but what's the broader impact? Let's revisit the example of uplink SNR (signal-to-noise ratio) from the previous section:



This is the average uplink SNR of client stations connected to APs at many different sites for a specific customer. As previously discussed, there are multiple modes, with some client stations connected with low uplink SNR and others connected with high uplink SNR. If one of this customer's sites has an average client station uplink SNR of 31 dB (the

black dashed line), is that good or bad? It's actually right in the middle of the entire distribution, based on the total distribution. However, by using the peer groups created by our AIOps workflow we can recommend a change would help that site perform better.



Here we see that our customer’s sites actually fall into many different environmental peer groups. **Note:** this is just one slice from a high-dimensional space, so we expect many of the distributions to look similar. Two of the distributions stand out from the others – peer groups 6 and 8 – shown in blue and red. Peer group 6 looks to be older APs in outdoor locations while peer group 8 seems to consist of mixed models and mid-size deployments.

Our same test site from above with an average client station uplink SNR of 31dB would sit in the 73rd percentile for peer group 6 but only the 21st percentile within peer group 8. These insights allow our customers to hone in on their specific problem sites. We’ll see in the next section that we can also use the relative performance of a site versus its peers to recommend changes that will optimize how the site is performing.

The addition of peer groups allows us to more accurately compare sites to one another and is a key building block in our ability to deliver more than just solving day-to-day issues. It’s unique to Aruba and the networking industry.

PUTTING IT ALL TOGETHER: AI AND ACTIONABLE INSIGHTS

With a strong AI framework, available environmental data, and the use of peer clustering, we will now illustrate how we put these techniques together to provide continuous and reliable Insights to each customer. Aruba AI Insights provide the analysis, data, and recommendations necessary for IT teams to rapidly respond to, prevent, and continuously optimize the network across a broad range of wired, wireless and WAN issues.

To illustrate the power of AI Insights, we’ll look at Airtime Efficiency which is a metric that Aruba uses to measure the overall Wi-Fi performance for an access point (AP), an entire building, or even a customer’s complete deployment. If an AP is providing high SNR links with high speeds on bands accessible to most of the client devices present, it will have high airtime efficiency.

For a better understanding, let’s take a look at two different APs – “AP A” in red and “AP B” in blue below in Figure 1. These APs are the same model, in similar environments, and have approximately the same number of clients connecting each day. For any given day, we can look at a couple of metrics related to airtime efficiency to see how they’re performing: client uplink SNR (signal-to-noise ratio) and client uplink speed.

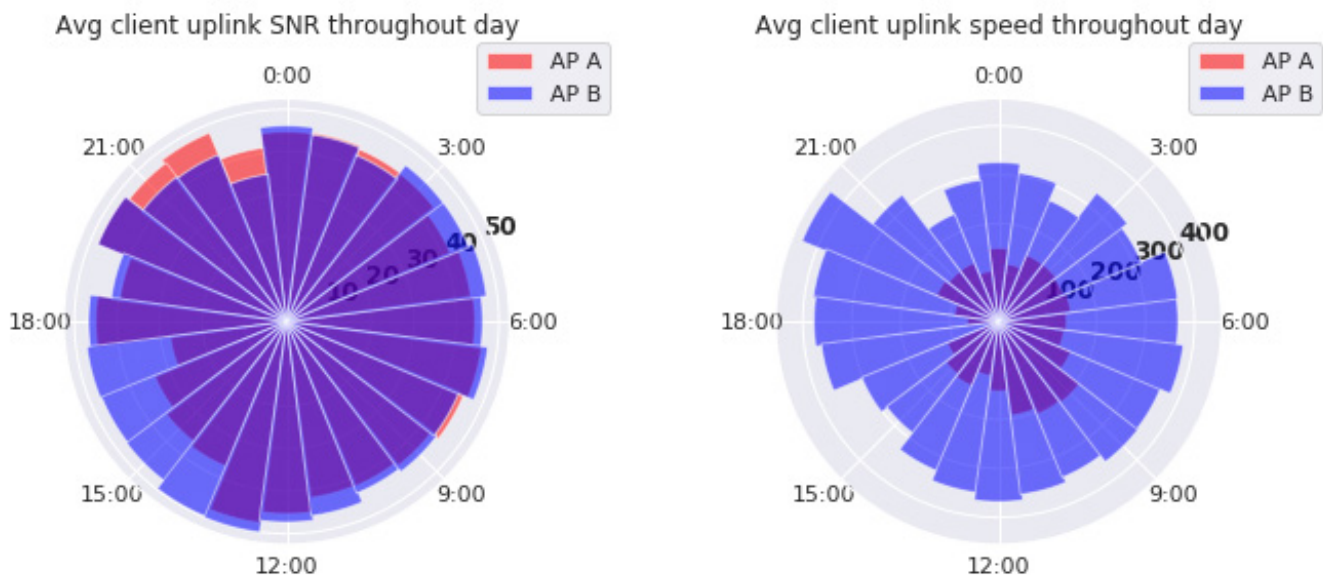


Figure 1 (Left) The average client uplink SNR across all clients connected for each hour of a given day. (Right) The average client uplink speed across all clients connected for each hour of a given day.

The two APs have a fairly close SNR, averaging 30-40 dB throughout most of the day. There are a few periods where AP B does much better than AP A, but the difference over the day isn't significant. However, there is a big difference in the client uplink speeds attained by AP B and AP A, with AP B averaging 2-3 times greater speeds. This could be due to a hardware issue, in which we have mechanisms in place to check for that, but in this case the only real difference between these two APs is their configuration settings. Now let's take a step back from the APs and look at the building's performance.

Setting up Wi-Fi for a building involves a lot of specialized knowledge and work. The average number of simultaneous clients as well as peak demand must be estimated. The right number of APs must be installed so that there are no coverage holes, and not too many such that they interfere with each other (and balloon the cost of the project.) There are many other factors that must be considered, but even so, the initial deployment is only a small part of the overall picture.

Changing client behavior and demands can tax even the most well-designed network. With roughly a dozen different knobs that can be tuned to optimize the performance of a wireless network, that include transmission power and number of channels we can use AI to take the guesswork out of optimizing a network. A feature within AI Insights called **configuration recommendations** lets us use our large dataset of Wi-Fi performance insights to recommend optimal configuration settings.

By calculating the airtime efficiency for each customer network and saving it along with the environmental features and configuration settings for each building, customers benefit from other sites experiences. For example, if a candidate building is not performing as well as peers, we use AI to consider possible performance gains by looking at configuration settings used at the top performing peer buildings. If we find a significant improvement, and our AI model has high confidence in it, we suggest the new settings for the underperforming customer's site. Figure 2 shows the result of such a process.

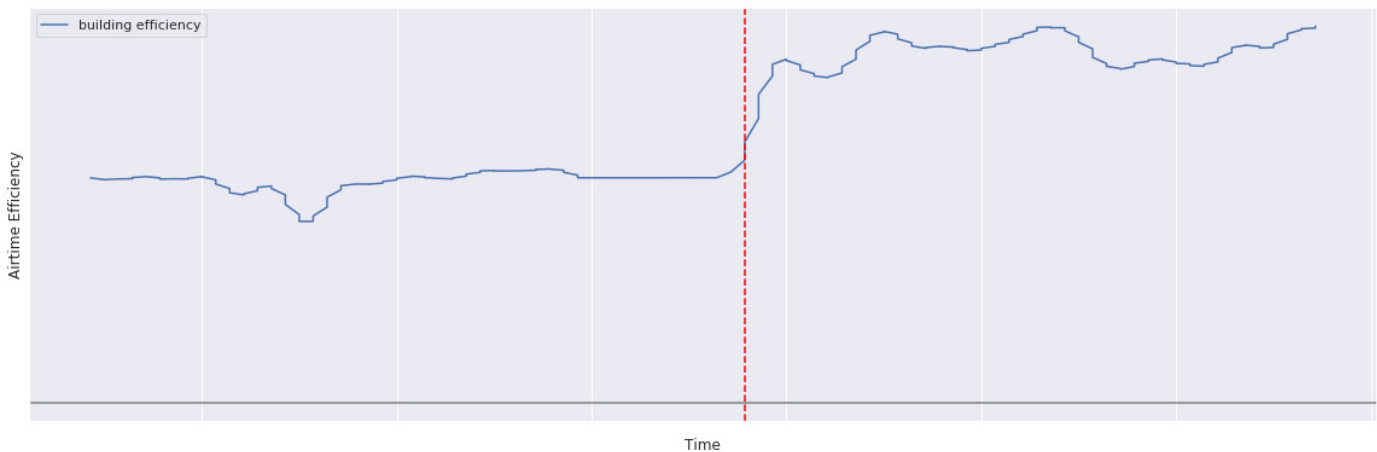


Figure 2. Plot of airtime efficiency vs time. The grey line is the origin. The red line when the recommended configuration changes were applied.



The AI Insight focused on airtime efficiency found a candidate configuration setting for this building from the tens of thousands of customer environments in our dataset to improve this building's Wi-Fi performance. The recommended change was implemented at the time indicated by the red-dashed line and the performance improvement was immediate. In this instance, an improvement of about 50%. What's better is that over time, the AI Insights will continue learning from our experts and other sites to continuously offer optimization guidance for the future of your Aruba deployment. What's more, the next step is to make this even easier for our customers by applying AI Insights automatically in a closed-loop fashion.

As a final thought, let's revisit those two APs from Figure 1. In actuality, they're not two different APs, but rather the same AP. The left image (AP A) and the right image (AP B) are the exact AP a week later and after the configuration changes recommended by our algorithm shown in Figure 2. By using our AIOps solution, our customer was able to improve the speeds provided by a factor of 3 while also slightly improving the signal quality (SNR) of the connections. This is in effect was possible by our measurement of the building's airtime efficiency.

SUMMARY: AI + AUTOMATION = AIOps YOU CAN RELY ON

Artificial Intelligence with Natural Language Processing for fast access to information, as well as supervised and unsupervised Machine Learning for network diagnosis, remediation and optimization can be a critical ally for any IT team – big or small.

Aruba provides the five most critical elements for AI success:

- Access to the volume and variety of data needed to train our AI models
- Domain expertise to know which problems to attack
- Built-in automation that helps IT teams identify issues more quickly
- The ability to scale and offer insights for any size organization
- Years of AI operating in live environments for problem solving and optimization guidance

You don't have to be a data scientist to leverage the power of AI, but it is important to understand what separates real AI that you can rely on from empty promises.