

A Network Design for the Digital Workplace

A NETWORK DESIGN FOR THE DIGITAL WORKPLACE

The Rise of the Digital Device

The next-generation digital workforce uses mobile devices in ways that are beyond conventional. No longer relegated to work email and personal entertainment, devices in today's mobile-first world are becoming the central component for every aspect of daily productivity.

For personal use, mobile devices and apps provide entertainment, social media, shopping, navigation, and financial management. For work, they provide voice communication, file sharing, video conferencing, CRM, unified collaboration, expense reporting, shared calendaring, help desk support, and project tracking. The lines between our social and work lives are becoming thinner every day, as revealed in the results of the latest industry survey conducted by the Economist Intelligence Unit (EIU):

- 30% of respondents say they would never work for a company that did not allow them to use their own mobile devices for work.
- The ability to work anytime, anywhere is regarded as the single greatest impact on employee productivity, with 49% believing remote, flexible working was responsible for increased productivity.
- 31% of companies surveyed use mobile communications apps to improve employee efficiency.

Likewise, the Internet of Things (IoT) is creating new opportunities and considerations for the mobile-first digital workforce as these "headless" devices are coming on-line at a staggering pace. These devices include lighting, temperature control, sound control, appliances, door locks, vehicles, medical equipment, manufacturing machinery, utility meters, parking meters, surveillance cameras, and much more.

Mobile and IoT devices are forcing IT teams and network infrastructures to adapt and support an entirely new set of requirements and challenges.

CHALLENGES FOR IT ORGANIZATIONS IN SUPPORTING THE DIGITAL WORKPLACE

1. Providing network access that meets the demands of digital devices
2. Monitoring and troubleshooting new network issues created by digital devices
3. Securing the network with respect to the new challenges of BYOD, IoT and guest access

4. Improving uptime and access to private and public cloud apps from any device and any location
5. Engaging with customers and employees on mobile devices to improve customer service and operations

Certainly, these are all significant challenges. But with each challenge comes great opportunity. We believe that those organizations that can rise to these challenges and adapt to the digital workplace faster than their competitors can and will recognize substantial benefits.

SOLUTIONS TO THESE CHALLENGES

1. Stable Network Access - Wi-Fi and Wired


Consider a real-world scenario where two users are performing the same tasks on a Wi-Fi network, but one is having a first-rate user experience while the other is not. This scenario is all too common in legacy networks that don't meet the demands of a digital workplace. But why does this occur, and more importantly, what can be done to address it?

To answer this question, we need to look at the device type – the smartphone, the tablet, or the myriad of other devices that the Wi-Fi network is being asked to support. Each of these device types has its own RF radio, its own operating system, its own Wi-Fi network driver firmware, and its own supported 802.11 amendments. As a result, each of these device types will often perform much differently within the same RF network.

Additionally, not all devices should be expected to support every company application. While many companies are successfully deploying VoIP and video on corporate-issued devices or certain models of smartphones, the prospect of delivering these applications flawlessly to any and every possible device in a BYOD environment is daunting.

Addressing the challenge of diverse device types is two-fold. First, devices should be classified by use case. For example, is the device in question meant to be used for location services, voice communication, collaboration, business application access, or IoT? Then, within each classification, each device needs to be tested and the RF environment should be designed to support each use case.

The new normal for network connectivity with mobile and IoT devices



Let's consider another real-world scenario in which a user has great success using gigabit 802.11ac Wi-Fi for all their needs. However, a few times a week, their performance degrades dramatically, even though their device indicates strong signal strength and "4 bars" of connectivity. Upon further investigation, it is determined that the poor performance always occurs while a large nearby conference room is being used. This is a classic example of a Wi-Fi network that was designed for legacy coverage requirements and not for today's capacity requirements.


Historically, Wi-Fi network design has been centered on signal strength. One would determine the required signal strength for an application, and as long as the entire coverage area maintained that pre-determined signal strength, the network performed as desired.

But things have changed, and the increased density of users and devices along with the use of real-time, high-bandwidth applications like voice and video have made focusing only on signal strength inadequate.


When designing a network for mobile and IoT devices, one must also consider the number of devices, the types of applications, and the throughput requirements needed within a coverage area. In the past, a single access point (AP) might have provided adequate coverage. Today, multiple APs are needed to serve the new capacity requirements of the network. In short, it is crucial to examine the need for both capacity and coverage in a Wi-Fi design for the Digital Workplace.

NEW CAPACITY REQUIREMENTS DEMAND A NEW DESIGN


Coverage Requirement (Old)
Coverage Design (Old)



Capacity Requirement (New)
Coverage Design (Old)



Capacity Requirement (New)
Coverage Design (New)




After putting careful thought into designing a network around the devices it will support, the location of the users, and what applications will be used, it is time to examine airtime and channel utilization.

The FCC only permits so much radio frequency (RF) for use in a Wi-Fi network. Subsequently, this RF must be used as efficiently and wisely as possible. There are so many tricks, tips, and factors at play in RF design that an entire white paper could be dedicated to the subject. We will do our best to touch upon the highlights here.

Frequently, device radios use the 2.4 GHz spectrum, although the most bandwidth is found within the 5.0 GHz spectrum. Utilizing mostly the 2.4 GHz spectrum while ignoring the open and unutilized bandwidth available on the 5.0 GHz spectrum can cause needless “congestion” for users. While more and more networks are starting to use the 5.0 GHz spectrum, there are a surprising number of legacy Wi-Fi networks designed with a majority of their clients on 2.4 GHz. The end result is a network that resembles a highway where some traffic sits bumper-to-bumper while other lanes remain illogically empty.

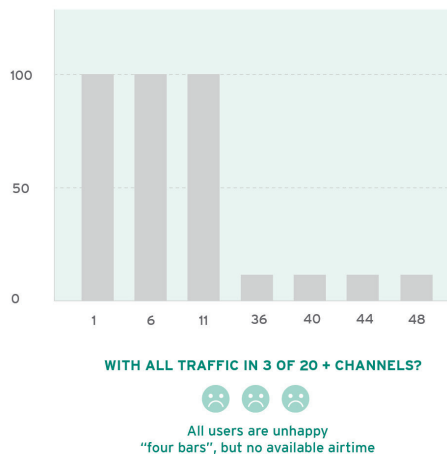
AIRTIME (CHANNEL) UTILIZATION IS EVERYTHING

People don't drive like this...



WITH ALL CARS IN 2 LANES?

Yet they design wi-fi networks like this



WITH ALL TRAFFIC IN 3 OF 20+ CHANNELS?



All users are unhappy
"four bars", but no available airtime

Now let's imagine a different scenario where the IT team has carefully studied the device types they need to support as well as their density and capacity requirements.

This team has also configured their network to reduce inefficiencies in their airtime utilization and has carefully selected the appropriate channel width between 20 and 40 MHz (80 and 160 MHz only have a small use case in the SOHO world, and even then, are generally not recommended).


They have reduced excessive overhead (non-payload traffic) by decreasing SSID count, eliminating low data rates, reducing co-channel inference and network retries, and eliminating rogue APs. They have implemented a quality of service (QoS) policy to optimize the performance of real-time applications like VoIP and video. They use network tools to monitor airtime fairness insuring that a few slow users won't degrade the overall performance and throughput.

Additionally, the team has carefully studied the use of Dynamic Frequency Selection (DFS) channels, which offer considerable bandwidth, but with two risks. The first risk is that not all devices support these channels, particularly IoT devices. The second is that government regulations prohibit the use of these frequencies during radar events, thereby temporarily blocking Wi-Fi on these channels. The IT team has decided to improve capacity in their high-density areas by implementing the use of DFS channels, while also overlaying at least one non-DFS channel in all high-density areas so that both risks are mitigated.

Finally, they have carefully monitored their maximum airtime and channel utilization so that overhead does not exceed 5% on any channel, data channels don't exceed 70% utilization, and VoIP / video channels don't exceed 30% utilization. The end result of the combined measures taken in this scenario is a Wi-Fi network that is capable of supporting the new demands of digital devices.

A WELL-DESIGNED NETWORK MAXIMIZES THE ENTIRE RF SPECTRUM

CLEAN WI-FI NETWORK = MAXIMUM THROUGHPUT



Just as we need to adapt the design of the Wi-Fi network to accommodate the needs of the digital workplace, the wired edge must also adapt. The upsurge of Wi-Fi digital devices brings with it a fast decline in the use of wired desktops and wired VoIP phones. With this change, we no longer need to size our edge closets with the expectation of two or three wired drops per office or cubicle. This allows portions of the IT budget to be reallocated towards meeting the digital device demands.

While it is of utmost importance that the wired access edge today supports the APs that are needed to connect wireless digital devices, the need to support wired digital devices has not disappeared. There are still legacy devices such as desktop computers, printers, smart televisions, legacy manufacturing equipment, and MRI machines that necessitate a wired ethernet port. Because of this, today's network will need to continue to support these devices and do so with a unified approach to configuration and management for both wired and wireless.

With programmability and intelligence at the core level that eliminates the need for extensive use of off-box solutions, a complex network consisting of both wired and wireless - of cutting-edge and legacy - becomes much simpler, and cheaper, to manage.

In the past, an Ethernet edge port was given a VLAN and a configuration that hardly ever changed, because the devices rarely ever moved. With the arrival of mobility, the network team now requires a better way to configure the Ethernet access ports. One option is to dynamically assign a configuration to each port as soon as the device type is classified upon connection. In this model, network operations teams do not have to spend hours configuring individual ports or keeping track of which port is assigned for which use.

The continued increase in traffic congestion within the access layer of the wired network requires the implementation of QoS policies on the wired backbone. Business-critical applications like VoIP or live-stream security cameras need to be classified and prioritized to ensure lower levels of latency and improved speed of operation for business processes.

Another important consideration regarding edge Ethernet ports is that they have typically supported only one or two devices at a time, and a speed of 1 Gbps was more than enough. However, an AP will support many devices at a time, raising concerns that 1 Gbps may become a bottleneck. There may be some debate regarding an AP's ability to exceed the 1 Gbps throughput threshold today, but there is little disagreement as to whether APs will exceed 1 Gbps soon. Consequently, multi-gig wired ports at the edge should be a component of today's wired infrastructure purchases.

RECOMMENDATIONS CHECKLIST 1

1. Document the mission-critical applications and device mix
2. Plan for both coverage and capacity
3. Reduce SSID count and disable lower rates whenever possible on Wi-Fi radios
4. Use DFS channels to overlap non-DFS channels to increase Wi-Fi throughput
5. Enforce QoS policies on Wi-Fi and wired to improve the performance of real-time applications
6. Dynamically classify each device on the Wi-Fi and wired portions of the network
7. Track channel utilization and enable airtime fairness on Wi-Fi radios
8. Implement 802.3at PoE+ and multi-gigabit wired ports to future proof

2. Deeper Visibility, Smarter Troubleshooting, Task Automation

Market forces continually challenge businesses to be more agile, asking them respond to change at a breakneck pace. As businesses rely on IT for their day to day operations, IT infrastructure must do the same. As we've discussed at length, the number of digital devices on the network is growing exponentially, as is the number of applications in use. Many organizations are moving their networks from a client / server model to a private / public cloud model. Additionally, more and more network traffic is voice and video, requiring real-time delivery. All of this adds up to a dire need for intelligent monitoring and troubleshooting capabilities. Today's aggregation and core switches absolutely must be able to meet this need. Moreover, when troubleshooting, monitoring and analytical tools are embedded in the switch, the network core's capability is on par with its criticality, and the cost and efficiency benefits of this are immense.

While the core still needs to have support for, HA, L2 / L3 protocols, VSF, and reliable security, it should no longer be based on a closed system. Instead, it should provide easy integration via API. To improve uptime and resiliency, the digital generation core needs to allow for configuration changes in a virtualized environment with easy roll-back.

Monitoring and troubleshooting this new dynamic on legacy networks with legacy tools can be extremely slow and time consuming. Today's visibility consists of network flows, logs, and 3rd party tools. Many network issues are not captured in logs, and those that are, can sometimes be difficult to resolve from log data.

Rules based data is critical to streamline troubleshooting. For example, if VoIP quality deteriorates, a rule could be used to collect and deliver QoS queue information, CPU utilization, and the last time a configuration change occurred. This troubleshooting can be done by the network before a human even knows there is a network issue. Taking this a step further, policy-based remediation can be implemented automatically or with one click.

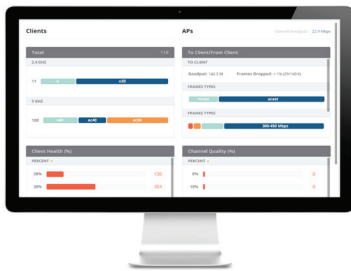
CLI programming (and troubleshooting) of network gear can be time consuming and prone to human error. The explosion of digital devices combined with shrinking network budgets compound this concern. Each new device type or application might necessitate a new VLAN, QoS policy, SSID or VRF. Network engineers need the ability to automate this type of programming, allowing them to focus on the bigger picture issues.

Better visibility and smarter troubleshooting is not only needed at the core, it is needed on the Wi-Fi network. Previously, we explained how identifying and understanding device types, coverage vs. capacity planning, and airtime utilization all have a direct effect on the Wi-Fi user experience. This brings us to questions such as:

- How do we monitor these critical metrics?
- How do we know the thresholds that indicate a poor user experience?
- Outside of Wi-Fi, what other factors could affect the user experience negatively?
- How do we not only collect the data, but fix the problem before the user even knows a problem existed?

To deliver the best possible user experience, IT departments must be able to predict, rather than react. Collecting key RF metrics such as interference levels, device counts, top talkers, retry rates, and signal strength can go a long way in helping to make those predictions.

These metrics can be collected by production APs or by dedicated monitors, and they are normally viewed in a Wi-Fi controller dashboard or monitoring software. Monitors can collect more data on more channels in a shorter timeframe than production APs, which is critical to being proactive rather than reactive. While this data can be used to track the general health status of your Wi-Fi environment, gaining a true snapshot of the user experience will require a deeper dive.



Basic health monitoring for clients and access points

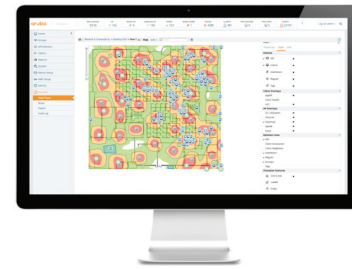
Here is where application-level data collection enters the scene. It is important to have a full understanding of the commonly-used applications and web-traffic on a network, as they are the vital indicators of general usage patterns. From there, we can define mission-critical applications, create QoS policies, monitor QoS tagging, configure QoS queueing, and ultimately guarantee that all mission-critical applications are getting the bandwidth and priority access needed for a flawless user experience.

The application class that is most associated with mobile devices is unified communication and collaboration (UCC), and accordingly, it deserves a special mention. UCC applications allow for voice and video communications as well as document sharing and collaboration and are key for raising productivity in a mobile-first workplace. Because UCC apps are real-time and require video quality, they place the highest burden on the network of any application. An organization that can deliver a UCC application to its employees on their mobile devices successfully meets the standards of a true mobile-first digital workplace.




Visibility into quality of mobile apps, including unified communications

Gone are the days of compiling network metrics in tables, spreadsheets or databases. Today's IT teams require easy-to-interpret visual maps where all the relevant data can be viewed on a single screen. Being proactive, rather than reactive, means being able to see a real-time view of the entire RF environment, as well as the underlying wired network, the location of users, their health and application performance at those locations, and any rogue devices or intrusion events.



Live monitoring of app usage, device density and air quality

To examine the true end-to-end user experience, we must look at the time it takes for a user to receive the first data packet upon first connect and after each roaming event. What matters most is the total time required to associate to a wireless network, authenticate to a Radius server, obtain an IP using DHCP, resolve host names using DNS, and transmit a first data packet. If we are not able to capture and follow each step of this process, it is somewhat akin to ignoring the weakest link in a chain and then wondering why it keeps breaking. Therefore, the next generation of monitoring tools must assist with evaluating this end-to-end performance. Additionally, if problems are discovered, those monitoring tools should allow us to use our APs as clients to run synthetic on-demand testing of the full user experience.



Predictive analysis of end-to-end user experience on the network

The final component of being proactive versus reactive pertains to notification. Once we can easily review metrics and define what is acceptable for the user and workplace, we need the ability to define thresholds and create alarms that notify the IT team before those limits are reached. In a true digital world, IT administrators are as mobile as the users, and so, notifications should go directly to mobile devices, triggering a response before the user even realizes there was a problem.

The same processes that will enable us to proactively solve network issues will also provide benefits for those occasions when an IT department must react to an issue. With the vast amount of data we can collect, we can send issues to Tier 1 help desk support rather than to higher level, and higher cost, solution teams. And, for those rare situations when a higher level of support is needed, the wealth of metrics we'll be able to provide will give that team a head start.

RECOMMENDATIONS CHECKLIST 2

1. Implement aggregation and core switches that are designed from ground up to support cutting edge and legacy needs.
2. Utilize a network analytics engine to give better network visibility and quicker troubleshooting
3. Automate the deployment of configuration tasks to reduce engineering CLI time
4. Collect information about app performance, in addition to basic Wi-Fi metrics
5. Enable forensics and trend analysis for deeper visibility into the end-to-end user experience
6. Have a method to synthetically simulate network performance and connectivity tests
7. Define threshold metrics and associated alerts for proactive response
8. Create a process for utilizing the tier 1 helpdesk when reactive response is required

3. Securely Support the Many Use Cases of BYOD, Guest and IoT

We have discussed how to design a reliable network, the importance of proactively monitoring the network, and the need to react efficiently when issues arise. The next steps for ensuring a positive user experience is securing the network. This requires three basic steps:

- Carefully author policies and rules
- Inspect all requests for access against the policies
- Deny or grant permissions accordingly

With a wide range of device types, user types, and applications, network access control (NAC) must be designed on a case-by-case basis. What follows are key elements for securing the mobile-first digital workplace.

Let's begin with the most common use case, which is granting network access to employees using a corporate-issued or BYO device. The goal is to provide access to the wired, wireless, or VPN network as easily and seamlessly as possible, without multiple access methods, credentials, or login screens.

To enable this, a single NAC platform that is performing 802.1X-based authentication is becoming the standard, with EAP-TLS using certificates being the most secure methodology. Authentication is often based on a username and password and is tied to existing user stores such as an Active Directory or LDAP. For added security, some companies are implementing multifactor authentication, which might consist of a secret question, fingerprint, voice recognition, photograph, or physical location.

In addition to the user, the device itself can be included as part of the network access policy. For corporate devices, this security policy could be simply based on a MAC address, or it can include the presence of a mobile device management (MDM) software agent installed on the device. For BYO devices, it could be based on OS versions, device model and/or presence of a mobile application management (MAM) solution. Mobile device and app management solutions can segment personal data from corporate apps and data to keep the network secure, or it can be used to help build and enforce blacklists for certain apps or "unhealthy" devices.

Employees, guests, contractors



Another common use case scenario is when network access is requested for a guest who is seeking Internet, printing, and other basic services. These "non-employees" might be visitors, clients, contractors, business partners, shoppers, and so forth.

Providing access to these guest users must be done securely and efficiently. Registration options typically include portal-based solutions such as self-registration or employee sponsor, both of which reduce the IT administration burden. Once registered, users typically obtain credentials via email or text, and are able to login and gain access to a secure VLAN that gives them limited network access.

Yet another example of a guest access use case would be a high-capacity facility such as an airport, sports venue, or shopping mall. In this case a simple "acceptance of use" account could be used, or alternatively, a social media login could be required. In some cases, open authentication could also be an option. In these situations, the guest WLAN must be 100% secure from the corporate network, typically through the use of role-based policies, enforced with firewall security embedded within the network infrastructure.



Custom designed, user-friendly web portals to help users onboard their devices to the network

The need for ‘non-employee’ network access is compounded by the growth of the IoT. While these devices are “headless” – meaning they are not manned by a human – they still require network access. Organizations must develop role-based policies for the IoT that are different than the policies for their employees and guest users. They may even need multiple policies to address the vast array of IoT devices that are connecting to the network.

Device profiling and fingerprinting, already important to network security, are measures that have become mission critical with IoT. Anytime a device attempts to connect to the network, there is a need to automatically profile it. This means knowing what type of device it is and what its function performs within your enterprise, allowing for the correct security policy to be applied. Should the network encounter a device that cannot be automatically profiled, that device should be quarantined until more information is obtained.



Automated classification of devices on the network

In addition to security issues relating to NAC, there are other security issues that must be accounted for in any network. Sitting at the top of that list is intrusion. Together, a wireless intrusion detection system (WIDS) and a wireless intrusion protection system (WIPS) assist in identifying potential threats to the AP infrastructure or associated clients.

A sophisticated, effective, yet easy-to-manage WIDS/WIPS solution should provide a network administrator the means to identify, assess, and defend against attacks while maintaining a seamless, uninterrupted, and secure experience. What’s more, the WIDS/WIPS must be highly customizable since the requirements and regulations that apply to each network can vary significantly, whether it be a conference facility, financial institution, or government agency.

Another prevailing security concern in networking is multi-stage malware or persistent attacks. These are threats where once a person gains access to the network they are able to remain on the network for a long period of time and launch multiple attacks, often stealing data rather than crashing the network. Legacy firewalls and antivirus software are no match for these types of attacks. The solution is next-generation protection with seamless integration between multiple products that perform detection, mitigation, and prevention. This protection must be done on the application layer, as traditional port-based security is insufficient.

These new security systems should integrate with network access policy platforms to enforce real-time policies for both Wi-Fi and wired network access in case any internal and/or external attacks are identified. While traditional NAC solutions only worried about the initial connection for Wi-Fi and wired devices, next-generation policy management solutions need to take action against the increasing risk posture of connected devices as malicious activities are identified.

RECOMMENDATIONS CHECKLIST 3

1. Create network access policies using a rich context of device type, health, user, time and location
2. Use device profiling to on-board IoT devices onto the network, and quarantine as necessary
3. Implement a single platform to handle control for wired, wireless and VPN network access
4. For employees, use 802.1X EAP-TLS coupled with multi-factor authentication
5. For secure guest access, enforce firewall policies within the network
6. Implement a wireless intrusion detection and/or prevention policy, procedure and solution
7. Integrate next-generation threat protection with the policy management platform

4. Deliver Cloud Apps and Collaboration from Any Location

The final piece in the end-to-end user experience is the application itself. Of course, no two networks or organizations have the exact same application needs. However, there are some fundamentals that apply to most scenarios.

The traditional client/server architecture is no longer suitable for the mobile-first network. This model was designed to deliver files and information from a single server to a PC that was connected on the same LAN, or occasionally WAN. While the PC was robust enough to support large applications stored locally, this is no longer the case with mobile devices. Which is why we now demand a streamlined and centralized delivery of applications.

Additionally, users are often not working on the same LAN as the server anymore. On a given workday, a user might be working from home, and therefore needs to access an application on a tablet using a cable internet connection. On the next day, that same user might be traveling, or visiting customers, and needs to access that application on a smartphone using 4G cellular service. And on the next day, that user might be checking on operations at several remote offices, where he needs to use the corporate WAN to connect securely to those applications.

Even though each of these scenarios requires its own unique network infrastructure design, the user requires their experience to remain consistent across all the potential combinations of access methods, device types, or locations. Satisfying that requirement involves moving apps to a centralized location and using virtualized servers -- a simple enough concept that delivers many advantages to the mobile-first network.

"Applications delivered from a centralized location" is the long way of saying cloud applications. Cloud applications are made possible by virtualized server farms and data centers, and can be delivered by the public or the private cloud. Private cloud applications are served from a data center that is owned and managed internally, and only users that have gained network access can use these applications. The clear advantage of using a private data center is the ability to maintain QoS and security protocols all the way from the client to application servers. The distinct benefit of public cloud applications is cost-efficiency and easy access. With all these options at hand, each organization must decide what combination of public and private cloud application use fits their specific needs and budget accordingly.

When organizations want to explore new geographies for their business, the faster they move into their new locations and are connected back to corporate resources, the faster they can realize a RoI. Network infrastructure should not hinder the way folks do business at remote sites.

RECOMMENDATIONS CHECKLIST 4

1. Design a network template for home, small and branch office connectivity
2. Implement an easy-to-use system for road-warrior connectivity on 4G, LTE and Wi-Fi hotspots
3. Create a list of apps that will be delivered by public and private cloud for remote users
4. Clearly define network metrics that will satisfy user expectations and monitor them in real-time
5. Implement end-to-end QoS for your private cloud applications, primarily for unified communications
6. Carefully negotiate and understand service level agreements (SLAs) for public cloud apps

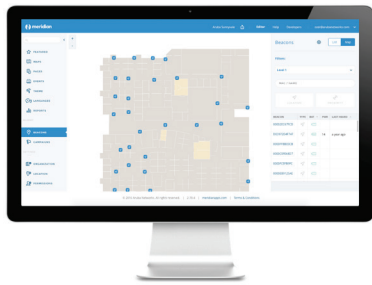
5. Leverage Bluetooth Low Energy (BLE) for Location Services

Location-based services and mobile app engagement are an increasingly important aspect of the digital experience. The ability to engage with customers and employees in a contextual manner via capabilities like targeted push notifications or turn-by-turn directions delivers a wide range of advantages for both the user experience and an organization's bottom line.

Bluetooth Low Energy (BLE) is the ideal technology for indoor location-based services. Operating on the 2.4 GHz frequency, BLE beacons have a typical range of about 200 feet and can provide location accuracy within a foot. When combined with a well-designed mobile app, a tablet or smartphone can be used as the client-side device within a BLE beacon network. The value of location-based information will become virtually limitless as the interactions between smartphones and IoT continues to grow.

BLE is not solely used for human location. It is also a great technology for asset tracking. A low-cost and small BLE tag can be placed on any asset, whether it is a mobile MRI machine in a hospital, a pull-out bed in a hotel, or a valuable tool in a manufacturing plant, and that asset can be tracked and located via an app on a mobile device. Asset tracking of this kind saves time and money.

BLE infrastructure can be deployed as stand-alone BLE beacons, as BLE enabled APs, or a combination of both. BLE beacons are simple to deploy and use, but large deployments can introduce operational challenges in terms of management, such as battery changes and adjustments to power levels. To overcome these challenges, and eliminate the time and expense of manual management, organizations interested in BLE beacons should look for solutions that include sensors and other remote management capabilities.



Centralized management of BLE beacons

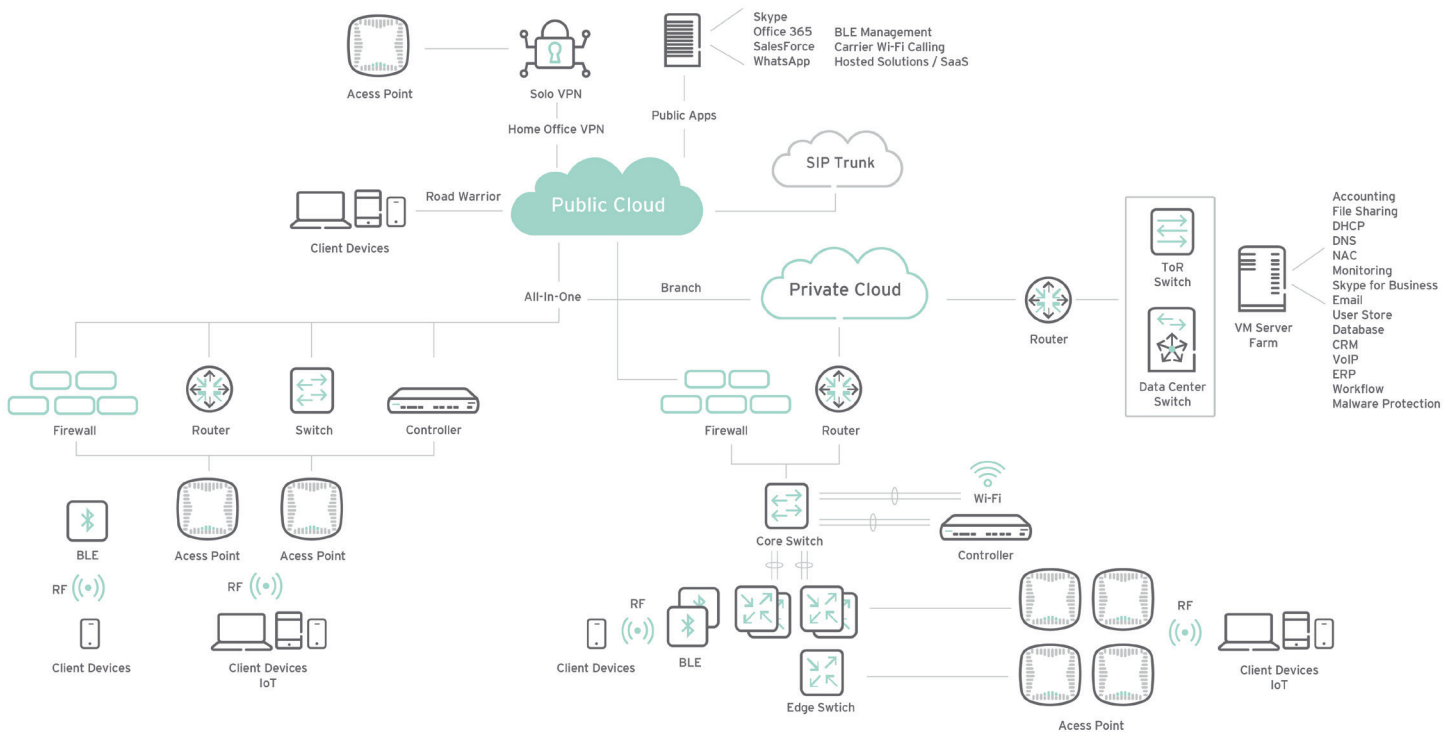
For organizations that lack the resources to develop their own mobile apps from scratch, there are a wide range of templates that can be used. For organizations that have an app but want to add location-based services to it, there are SDK add-ons for functionality such as navigation and context-aware marketing. Some of the common use cases for location-based services include:

- Deliver a “blue dot” experience on a map to help users navigate your facility with turn-by-turn directions
- Allow users to find friends and co-workers quickly and easily
- Asset tracking of anything valuable with the use of BLE tag
- Allow customers to find employees easily when they need assistance
- Deliver push notifications to customers based on their current location and past shopping habits
- Provide automated check-out, room access, and temperature control for hotel guests
- Automatically turn lights and other utilities on and off as employees enter and exit rooms

RECOMMENDATIONS CHECKLIST 5

1. Clearly define use cases for BLE beacons to enable location-based services
2. Start working with an app developer partner, with a laser focus on improving the end user experience
3. Create a plan and strategy for how your mobile app will interact with IoT devices
4. Use centralized management for BLE beacons to reduce operational cost

What does a digital device network look like?



CONCLUSION

No matter what type of business you are in-financial services, government, healthcare, education, retail, hospitality, or manufacturing -digital devices are impacting the way you conduct business and the type of network infrastructure you need.

This document has outlined some of the challenges that digital devices are presenting to organizations across

all vertical markets and it has reviewed some of the solutions to those issues. Forward-looking enterprises will take that information and will work with mobility specialists, network architects, and engineers to find detailed, customized approaches for their specific use case. In doing so, they will enjoy benefits ranging from improved communication and collaboration to enhanced productivity and reduced capital and operating expenses.