SCTE·ISBE CABLE-TEC EXPO 2019 NEW ORLEANS, LA SEPT. 30-OCT. 3

2019 Fall Technical Forum SCTE·ISBE • NCTA • CABLELABS

Defining the Premise and the Edge

A Managed Wi-Fi Service Application

A Technical Paper prepared for SCTE•ISBE by

John Gammons
Director, Wi-Fi Engineering and Operations
Cox Communications, Inc.
6305 Peachtree Dunwoody Rd. CTECH B09-105, Atlanta, GA 30328
404-269-6992
John.gammons@cox.com

Key Contributors

Michael Hurd, Cox Communications, Inc.

Kevin Klimek, Cox Communications, Inc.

Marco Lin, Cox Communications, Inc.

Jonilson Santos, Cox Communications, Inc.

Acknowledgements

Nick Green, Cox Communications, Inc.

Brady Puckett, Cox Communications, Inc.

Altan Stalker, Cox Communications, Inc.

Drew Stravelli, Cox Communications, Inc.

Table of Contents

Title **Page Number**

Table of Contents 2

Introduction 4
1. Managed Wi-Fi Services Overview 4
   1.1. Use Cases 4

Content 5
2. The Problem 5
3. The Placement of Functions 5
   3.1. Functional Overview 5
      3.1.1. Access 6
      3.1.2. Switching 6
      3.1.3. Wireless LAN Controllers (WLC) 6
      3.1.4. Wireless Access Gateway (WAG) 6
      3.1.5. Routing 7
      3.1.6. Firewall 8
      3.1.7. Other Services 8
   3.2. Considerations 9
      3.2.1. General 9
      3.2.2. Virtualization and Orchestration 9
      3.2.3. Advanced Site to Site Network Requirements 10
   3.3. Analysis 10
      3.3.1. The Case for Centralization at the Edge 10
      3.3.2. The Case for Distribution to the Premise Edge 12
4. Solutions by Vertical and Use Case 14
   4.1. Edge - Public Hotspot and Guest Services 15
   4.2. Premise Edge – Enterprise/School/Medical and Venues 15
   4.3. Hybrid Use Cases 16
      4.3.1. SMB and MDU 16
      4.3.2. Multi-Service Customer 16

Conclusion 17

Abbreviations 18

Bibliography & References 19

List of Figures

Title **Page Number**

Figure 1 – Next Generation Service Provider Network Model 5
Figure 2 - Functional Overview 6
Figure 3 - WAG Functions 7
Figure 4 - Service Overlays by OSI Layer 2/3 9
Figure 5 - WAG Deployed in the Edge 10
Figure 6 - WAG Deployed in the Premise Edge 11
Figure 7 - Distributed Premise Edge Architecture 13
Figure 8 - Customer Segmentation 14
Figure 9 - Advanced Customer Integrations 14
Figure 10 - The Hybrid Approach 16

Introduction

The next generation of Managed Wi-Fi Service offerings have distinct requirements and challenges. These offer unique challenges to both the business and the architects defining the underlying technology solutions. As architects and engineers are beginning to create solutions for these services, there are numerous requirements which need to be considered. The Service Provider will need to find the right balance between deploying these services in the Edge and the Premise Edge. The Service Provider may need to leverage multiple edges depending on the customer requirements. Careful analysis of the offering, the functions and the benefits will be required as Service Provider's aim to solve these fundamental questions surrounding these Managed Services. This analysis will review the offerings, the functions, analyzing the benefits of each architecture, and will provide directional guidance regarding the ideal deployment scenarios by use case, making the case for both Edge, the Premise Edge and a hybrid architecture for Service Providers.

1. Managed Wi-Fi Services Overview

The Managed Wi-Fi Service is a relatively new offering from Service Providers, and an important one. These services are driving new use cases and new demarcations for the Service Provider and the customer. Access speeds are now commonly approaching 1Gbps and beyond, and quickly eyeing 10Gbps on the horizon. The available cloud services which are being provided have made it easy for users to consume complex services with relative ease and reduced cost. Cellular and mobile networks have improved their reliability and general availability. Furthermore, device and network standards have evolved to improve compatibilities. These trends have converged in bringing about the BYOD (Bring Your Own Device) era, where the traditional Internet Service connection has transformed into the ability to easily consume wireless or wired access, in several geographical locations, and in a reliable manner. Consumers are also demanding more features and integration options in an intuitive and easy to use App-like interface where support is available but rarely required. This is the new Service Provider landscape and the new Managed Wi-Fi Service. These trends have continued to evolve across multiple industry verticals including:

- Public Hotspot
- Hospitality/Amenity/Guest Services
- SMB (Small-Medium Business)
- MDUs (Multiple Dwelling Units)
- Enterprise
- School/Campus
- Medical
- Venues

1.1. Use Cases

As we dive deeper into this offering, let's first analyze the use cases Service Providers are being challenged to deliver for these verticals. Typical use cases for these services include the following:

- Guest Wireless Local Area Network (WLAN), Local Area Network (LAN), Internet access service and associated services for end-user devices.
- Private Wireless Local Area Network (WLAN), Local Area Network (LAN), Internet access and associated services for end-user devices.

- Security and content filtering solutions
- Application Visibility, Monitoring and Analytics
- Advanced layer 3 routing and layer 2 transport protocols
- Network Authentication services
- Integration with Customer or 3rd party systems and services

In order to deliver these use cases successfully to the consumer, in a way that satisfies their BYOD (Bring your own device) mindset, the new Managed Wi-Fi Service offering is creating a need for a next-generation network deployment model for the Service Provider. As the demands on the network and the cloud increase, services in the SP Cloud and even the Internet are distributing to take advantage of edge computing. Gartner defines edge computing as "solutions that facilitate data processing at or near the source of data generation". For Service Provider to client communications, this could mean distributing to the edge of their cloud, or the premise itself. This traditional SP premise delivery architecture is being augmented to include one or more access points and switches deployed downstream of x86 based commodity hardware as access and services meet in these new Managed Wi-Fi Service offerings, as illustrated in Figure 1 below:

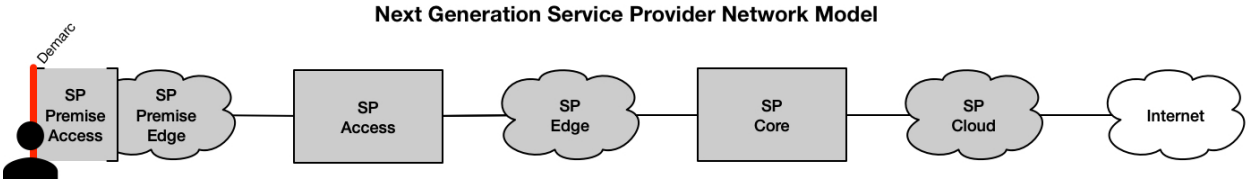

Figure 1 – Next Generation Service Provider Network Model

Content

2. The Problem

While this next generation Service Provider network has advantages, the new capabilities offer new challenges and questions. Most notably, the following questions surface:

- Where does a Service Provider deploy these Managed Services functions?
- When does the Service Provider consider Edge deployments?
- When does the Service Provider consider Premise Edge deployments?

These questions are bound to generate discussions, some of which may even get heated amongst Network, Datacenter and System Architects.

3. The Placement of Functions

3.1. Functional Overview

Before we decide where functions are placed within the network, let's first analyze each function within the typical solution to ensure we understand the functions to be performed. The below illustration in Figure 2 shows the necessary functions to be covered as part of this analysis. The services which will serve as our primary focus are shown in white, while the others are shown for reference in gray to ensure

a proper understanding of the entire landscape of the solution. Each functional area is shown in its logical place within the network, and then subsequently defined below.

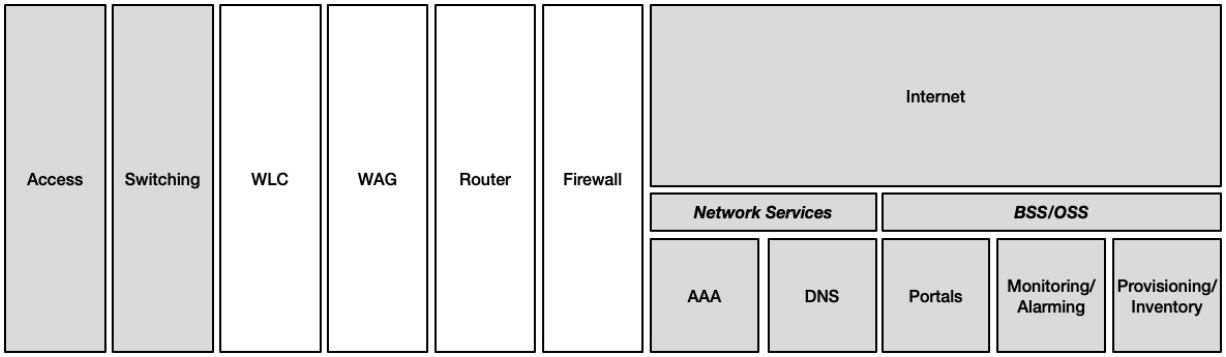

Figure 2 - Functional Overview

3.1.1. Access

The access network is delivered through 802.11 RF based Wi-Fi or wired copper ethernet connections. These are physical components which are not able to be virtualized and are required to be on-site. These are not directly analyzed within this paper.

3.1.2. Switching

The ability to have multiple access mediums connected to common or independent layer 2 broadcast domains. Some aspects of switching could be virtualized through a variety of protocols, and that is assumed to be the case. Since the construct could not be virtualized completely, it is shown to be out of scope.

3.1.3. Wireless LAN Controllers (WLC)

Many enterprise and carrier-grade access points manufacturers incorporate a Wireless LAN Controller function into their offering. These systems provide intelligence in the form of Radio Resource Management (RRM) and Self-Optimizing Network (SON) functionalities for the access points. When WPA2, 802.1x, or 802.11r Fast-Secure Roaming are enabled, the WLC will also serve to ensure keys are cached and exchanged quickly without sacrificing the network integrity. These solutions are often available in virtualized container forms.

3.1.4. Wireless Access Gateway (WAG)

The Wireless Access Gateway (WAG) is a critical component within the typical Service Provider Wi-Fi offering. It may also be referred to as a TWAG (Trusted Wireless Access Gateway) per the 3GPP and Evolved Packet Core (EPC) standards.

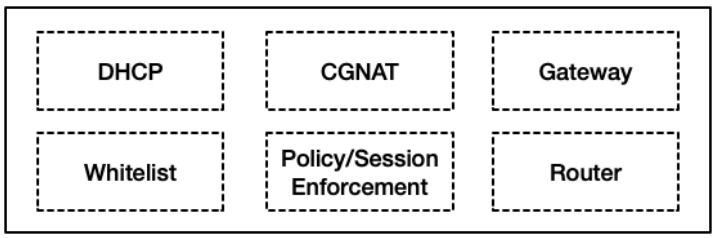

Figure 3 - WAG Functions

3.1.4.1. DHCP

This function assigns IP addresses to end user devices on the network for both wireless and wired mediums across the Guest and Private use cases explained above.

3.1.4.2. Whitelist

This is the ability to allow only certain protocols or destinations through the WAG while a session is in an unauthenticated state. An example of this allowing the ability for a user to render a splash page portal for registration or authentication on the network.

3.1.4.3. Network Address Translation (NAT) or Carrier Grade NAT (CGN/CGNAT)

Network Address Translation as described in RFC 1631 Traditional Network Address Translator is a function providing the ability for clients to exist within RFC 1918 Private Address Space using 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16 or RFC 6598 Shared Address Space using 100.64.0.0/10 but still reach the Internet. NAT is also "often accompanied by application specific gateways (ALGs)" to monitor application payloads and correct as necessary (Srisuresh).

3.1.4.4. Session Policy Enforcement and Accounting

This functionality refers to the ability to create and tear-down sessions as necessary to enforce network policies on individual devices for the purposes of authentication, authorization, and accounting. Sessions may include both wired and/or wireless network devices.

3.1.4.5. Gateway and Router

The WAG serves as a Layer 3 IP Gateway for the end-user devices. It also routes packets in between multiple networks or forwards packets upstream through POPs or Peering locations to the Internet.

3.1.5. Routing

This function implies the ability to have multiple customer networks and the ability to optionally route between multiple networks and the hosts within them using traditional Layer 3 routing techniques. This is listed separately from the WAG in case there are more advanced requirements such as MPLS, VPLS, VXLAN, SD-WAN or other requirements. This is shown upstream of the WAG however it could alternatively exist downstream of the WAG, especially if it is a Layer 2 transport protocol.

3.1.6. Firewall

The firewall function has the ability to analyze and limit both external and internal communications across the physical interfaces. Packet analysis may include source, destination of both the hosts and the ports/protocols. Tracking state or session awareness is optional and not pertinent to this paper and analysis.

3.1.7. Other Services

The following other services may be included in a Managed Wi-Fi Services offering but have limited scope within this specific analysis.

3.1.7.1. Authorization, Authentication and Accounting (AAA)

This service is responsible for authenticating and authorizing devices in accordance with the defined network policies or proxying as appropriate to upstream systems including partners when policy and business rules dictate. The AAA is also responsible for accounting policies and processing near-real time on the network at defined intervals and delivering those records to a data warehouse infrastructure as necessary. For the purposes of this paper, we will abstract any potential upstream authentication, authorization or accounting systems into this function for simplicity. Communications will typically occur over RADIUS or DIAMETER between the Session and Policy Enforcement point and the AAA. This is assumed to reside in the Service Provider's Edge or their Cloud, depending upon their specific requirements.

3.1.7.2. DNS Resolution (Proxy)

This functionality provides the end user device with the ability to resolve Domain Names to IP addresses as described in RFC 1034 and 1035. This is assumed to be provided by the Service Providers traditional network services infrastructure. It is only referenced as consideration for proxy services lower within the network, likely in the WAG as an extension of the Gateway function.

3.1.7.3. Portals

Portal functions include Graphical User Interfaces which are available over HTTP/HTTPS protocols and allow unknown end users to authenticate to the network. They also include the ability to allow customers or Service Provider operational support personnel to manage their end user experience, manage their network configurations and to also view the status and analytics of the platform.

There are extreme cases such as venues where a significant number of authentication transactions may occur within a short amount of time. While these scenarios may warrant distribution of a portal application to the Edge or Premise Edge, that will not be a significant consideration in this analysis.

3.1.7.4. Monitoring and Alarming

Any Managed Service offering delivered by a Service Provider should include the ability to monitor and alarm upon any event that is deemed actionable or any event that either the SP or the Customer have deemed to require awareness. The exact placement of this function could vary, but it is assumed that the SP has existing infrastructure or requirements which would dictate the placement.

3.1.7.5. Provisioning and Inventory Management

Service Providers will require back office integration with both Provisioning and Inventory management systems. These systems are out of scope for this analysis.

3.2. Considerations

3.2.1. General

As we first consider these functions, and their placement, it is important to note the network layers for each function. Consider the following logical diagram of functions, which makes a number of assumptions and simplifications on the WLC, WAG, Router and Firewall placement for the purposes of generalization:

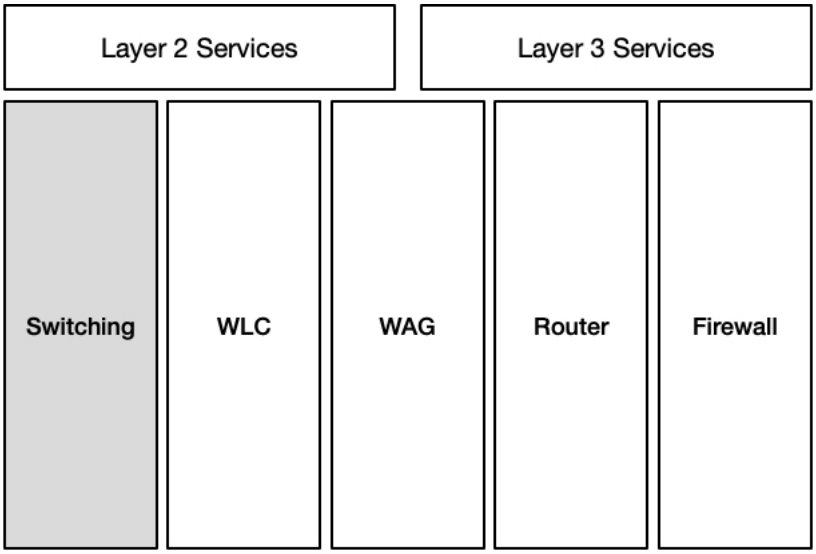

Figure 4 - Service Overlays by OSI Layer 2/3

Depending upon configuration the ordering of these functions can certainly shift. The assumption for this analysis is that these are directionally correct, and more importantly linked directly or through Service Function Chaining (SFC) in either the Premise Edge or Edge to ensure these functions are compatibly linked to one another. The Service Provider would have deployment flexibility to stitch these together in a number of different ways to meet the individual customer requirements. As an abstract, some part or all of these functions are assumed to be required for deployment to fulfill the Managed Wi-Fi Service requirements of the customer. The remainder of the analysis will often times generically refer to an abstract functional representation of the WLC, WAG, Router and Firewall as the Managed Service.

3.2.2. Virtualization and Orchestration

Both the Edge and Premise Edge architectures can take advantage of Virtualization, Orchestration and Service Function Chaining. The ability to deploy container-based services and functions on-the-fly for customers is attractive in all cases. The Premise Edge deployment scenario would require more "clusters", which would drive complexity of those solutions and overall costs to deploy and operate orchestrated virtualization environments such as PaaS (Platform as a Service), IaaS (Infrastructure as a Service) and NFV (Network Function Virtualization) solutions. However, the more distributed Premise Edge model would also likely improve the value proposition of orchestration for the Service Providers.

This is due to the fact that as the number of environments they need to manage increases and those environments become more dispersed, their business would see more overall value in orchestrated provisioning of those services.

3.2.3. Advanced Site to Site Network Requirements

Any Enterprise customers which require advanced site-to-site layer 2 networks such as L2VPN, Metro Ethernet, or L3VPN networks could be ideal candidates for either deployment methodology. While there are benefits to centralize and leverage an SD-WAN or MPLS/VPLS technology to aggregate these networks at the Edge, there is also the opportunity for the Premise Edge to create these overlay networks. Due to this, many of the above factors would have to be considered alongside their exact site-to-site requirements as well in order to best meet the customers individual needs.

3.3. Analysis

3.3.1. The Case for Centralization at the Edge

There are numerous arguments for centralization of the Managed Service functions. The arguments which typically drive centralization are based within the following principles:

- More easily deliver larger mobility domains
- Improved operations and reliability
- Reduced total cost of ownership

We will discuss and analyze these aspects further in the sections that follow.

3.3.1.1. Mobility

The traditional role of a WAG within a wireless architecture is to function as the border of a mobility domain to anchor sessions higher in the network and create session mobility across small or large portions of a network. To accomplish this WAG terminates Layer 2 and serves as the IP gateway for end user networks. In order to have WAG exist in the Edge, either the layer 2 network domains will need to be extended from the Premise to the Edge, or an equivalent layer 2 over layer 3 overlay or tunneling protocol will be required. The end user will have a seamless experience as they roam from Premise A to Premise B, assuming ubiquitous RF coverage as the session is anchored in the SP Edge. This is illustrated in Figure 5 below:

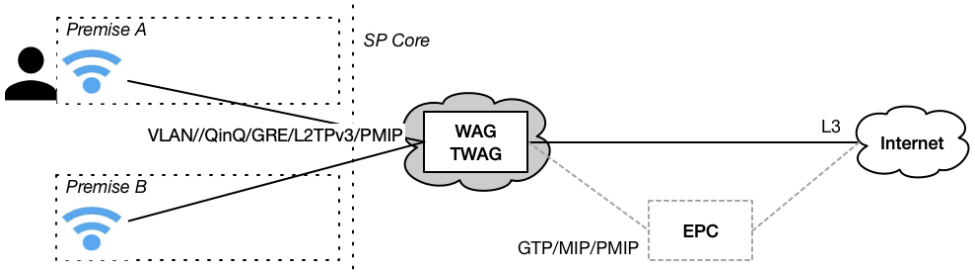

Figure 5 - WAG Deployed in the Edge

As the WAG is distributed to the SP Premise Edge, the mobility domains are also distributed. In order to maintain a seamless experience, sessions will need to be anchored between Premise A and Premise B WAGs either through the use of direct tunnels or another layer of aggregation and anchoring such as is available within an EPC, or Local Mobility Anchor (LMA) in the case of PMIPv6. This adds additional

cost in the form of deploying more WAGs in addition to more layers of complexity to provide a larger mobility domain.

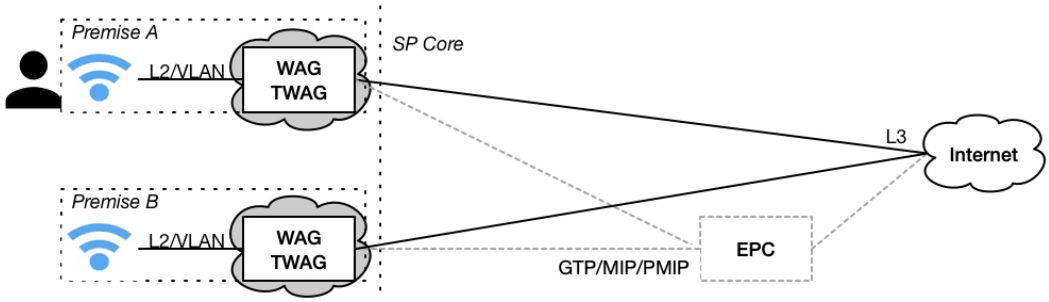

Figure 6 - WAG Deployed in the Premise Edge

In both cases, the architecture is very similar, only potentially introducing different levels of overhead to transport Layer 2 packets to the WAG and altering the location of the deployed equipment. As the WAG location is determined, and more importantly where end-user layer 2 is terminated, one can determine how overlay managed services are applied for the customer. It is important to note that in both cases there are standardized means to enable session mobility across multiple WAGs, however these technologies invoke additional cost as well as additional complexity and overhead which will need to be balanced.

It is important to understand the implications of mobility. A customer or service which has requirements for a large mobility domain will typically require some centralized Edge or Cloud elements in order to accomplish mobility across numerous locations. This should be one of the primary considerations when determining if a function or service should be deployed on the Edge or Premise Edge. Another consideration in mobility, is security of the RF layer. WPA2/3, 802.1x (and all EAP types therein), as well as 802.11r all have implications to mobility as well. These requirements will need to be well understood and the interconnections between the functions will need to be designed to meet the exact customer requirements. For all of these reasons, use cases where a mobility domain is required over a larger area prefer deployments within the Edge.

3.3.1.2. Improved Operations and Reliability

As solutions centralize, the Service Provider can create larger clusters that serve more customers. Even if they are supporting the same number of virtual instances, having them coexist on less physical hardware will utilize less operational resources. Additionally, having them in fewer geographical locations would enhances their supportability and overall reliability of the service for the customers. Centralization should therefore directionally improve the overall operation of the solution. Quantification of the improvements would vary greatly based on the individual Service Provider circumstances and would need to be analyzed for each individual SP and for each service offered.

In general, for the use cases and customers that this paper is referring to, the Service Provider network is assumed to have improved datacenter facilities when compared to the customer premise. This will likely include more highly available and consistent connectivity, power, cooling, and physical security at the SP facilities. As the size of a customer's network declines so will their environmental protections. As an example, consider the following potential deployment examples of customer premise equipment:

- on a table or desk in the corner of a common area such as a breakroom where coffee is spilt
- setup under an employee's desk where the cables and equipment are stepped on multiple times a day

- plugged into an extension cord that gets unplugged in favor of a vacuum each evening

These are all too common occurrences for Service Provider support personnel. Inversely, consider the Service Provider facilities, where typically there are varying degrees of criticality assigned to datacenters and the greater the impact, the greater the environmental protections which are put in place to protect the facility and therefore improve the services they provide. With this logic, centralization will correlate to improved environment availability. This may also imply that smaller customers may see additional benefit in more centralization than larger customers, as they are less likely to have taken any precautions to improve the provisions at their facility.

3.3.1.3. Reduced Total Cost of Ownership

As resources including compute, memory and storage are centralized there are direct financial gains to the business. There are a number of reasons for this including:

- Economies of Scale
- Reduced Operating Expenses
- Improved abilities to utilize/repurpose available resources

Oversubscription of the network is an assumed part of the business case. Centralization allows what would otherwise be Premise Edge hardware to be centralized and therefore oversubscribed based on true peak utilization. Additionally, as the solution centralizes, there are opportunities for multiple customers to share the same virtual instances within the Edge as well. In addition to the hardware, this would lead to even greater utilization of the software and subsequent licensing. There are a number of factors to consider including the different market segments and what that may mean for peak utilization of the Edge. As an example, the same Managed Service delivered to a dinner restaurant and a coffee shop will likely have distinctly different peak times when customers and employees are utilizing the services. Allowing the same hardware/software to serve both segments will allow those resources to be maximized. If this is done across the scale of the entire Service Provider, this could mean significant reductions of resources and overall cost to the business.

While there are technologies which could take advantage of underutilized resources on the customer premise, these would result in additional utilization on the network which is less than ideal. Similarly, local traffic which does not need to traverse the backhaul or the core should not as this creates undue network utilization and drives network costs. Therefore, centralization is only beneficial under the appropriate use cases.

3.3.2. The Case for Distribution to the Premise Edge

There are many considerations and requirements which would drive the Managed Service to be distributed to the Premise Edge. In every case, individual customer requirements must be analyzed to ensure that all considerations are taken into account. Generally speaking, the requirements which drive the Managed Service functions to be deployed on the Premise Edge fall into the following categories:

- Intra-premise Communications
- Customer Security and Segmentation
- Advanced Customer Integrations
- Performance

We will discuss each of these in greater detail in the following sections. Please reference the diagram below in Figure 7.

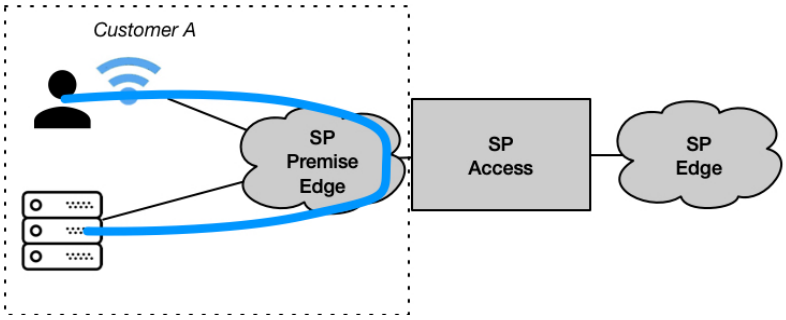

Figure 7 - Distributed Premise Edge Architecture

3.3.2.1. Intra-Premise Communications

The primary driver for distribution is the flow of customer traffic. Anytime there are requirements to route between multiple networks, the packets will have to traverse the Managed Service. The Premise Edge deployment is able to keep all premise communications local to the premise, whether they are layer 2 or across distinct layer 3 networks. These would include customers with local file shares, client to client communications or sharing, or even printing.

Key customer requirements that may drive this include:

- Substantial Client to Client communications
- Multiple premise networks with interconnectivity requirements
- Mission critical Client to Client communications even through an Internet disruption/maintenance activity

3.3.2.2. Customer Security and Segmentation

Certain aspects of the solution may have specific requirements which require additional network segmentation. These situations may create considerations with regards to which functions are deployed on-premise in order to improve the segmentation between customers. As an example, consider a customer with locally hosted servers/services which are managed by the customer or an independent 3rd party but reside on the customer premise. These could range from an on-premise video camera termination appliance, to a server which terminates network authentication requests to 802.1x/LDAP/Active Directory databases. As these network, firewall or proxy functions move upstream into the Service Provider's Edge there will inevitably be multiple customer networks overlapping on the common Edge. While these can be virtualized to segment customers from one another, there is still a risk that these hooks into the customer's private network could be compromised by another customer.

As these functions are distributed to the Service Provider Premise Edge, the amount of separation between the customers increases. From a security perspective, physical separation is always preferred over virtual segmentation. The improved segmentation is certainly preferred by not only the Service Provider, but more importantly the customer. In the centralized model, the connections would have to be secured against both the public internet, as well as other customers. For extremely secure applications, when leveraging the Premise Edge, the Service Provider could recommend that the customer keeps certain networks completely isolated from the Internet and from other customers, to further limit the potential impacts. Consider the figure below which shows these potential use cases:

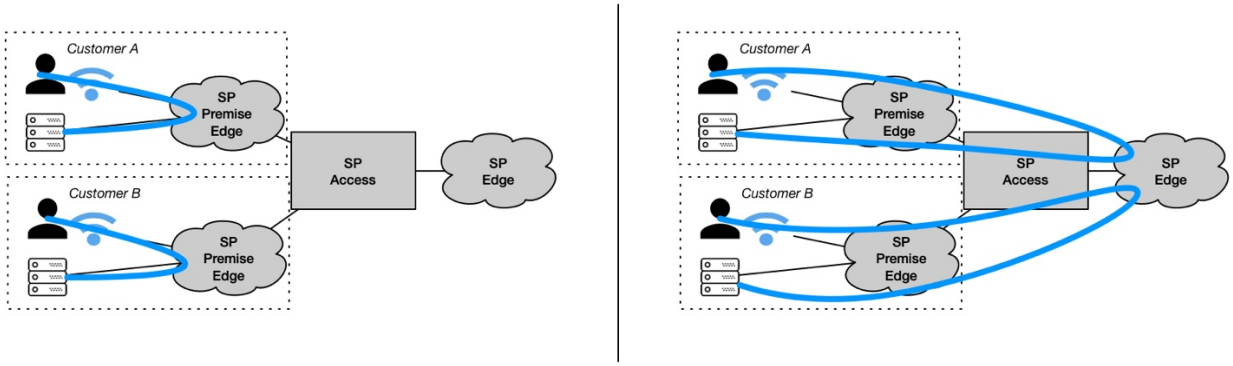

Figure 8 - Customer Segmentation

3.3.2.3. Advanced Customer Integrations

Customers may require that the Managed Service integrates with other premise-based solutions or services. These 3rd party solutions may be owned by the customer or may be provided by 3rd party services that the customer has chosen. These could include enterprise or venue customers which have special requirements for visibility to outbound traffic. These may exist out of line and simply require a tap, or a port devoted to the monitoring and replication of traffic such as a SPAN port. Other examples may include more intrusive integration such as a traffic shaping appliance or firewall as shown in Figure 9 below. It is important to acknowledge that the Service Provider may choose not to support these types of integrations. However, these complex integrations may also be a necessary challenge that is associated with deploying Managed Services.

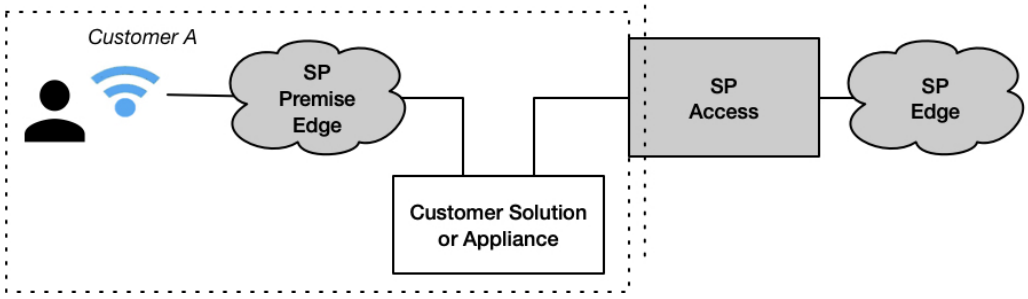

Figure 9 - Advanced Customer Integrations

3.3.2.4. Performance

In any network architecture and across any medium, shortening the path always equates to less latency and a higher performing network. Inversely, adding to a path, especially with additional unnecessary hops or processing, degrades the network. Higher performance applications will almost always require additional dedicated resources devoted to making the best use of the network. In this case, the additional overhead and wrappers of tunneling are not preferred, and in these cases, neither would the shared infrastructure of the Edge. Therefore, higher performance applications will typically prefer distribution.

4. Solutions by Vertical and Use Case

We have a clear understanding of the problem, a solid grasp on the functions, and have explored the considerations which would drive the Managed Service placement. The analysis of the architecture benefits thus far has shown that centralization at the Edge is necessary when mobility requirements

demand it and is preferred whenever possible due to improved operations and reliability while also reducing cost and improved utilization of resources. Distribution to the Premise Edge is required as the solution increases in complexity either due to site integrations or demands for increased security or performance. Let us next consider the offerings of the Managed Service we are offering to our customers and apply those use cases to our analysis thus far.

4.1. Edge - Public Hotspot and Guest Services

Public hotspot services typically will require a Service Provider to provide basic Internet Access services to a high number of temporary mobile or nomadic users. These use cases will typically be highly susceptible to peaks and valleys of utilization based on their individual traffic patterns. These networks are typically openly accessible as well, which tends to drive significant walk-by user connections as there are a number of devices which are configured to attach to nearby open networks by default. These networks will rarely have local switching and routing requirements such as premise file-sharing or printing.

The demands for high capacity and limited accessibility requirements make these use cases highly attractive for centralized deployments within the Service Provider Edge or even higher within the network. Large aggregate CIDR blocks of IP addresses can be allocated in these central locations and reused across a high number of end-user subscriber devices. The highly mobile and nomadic nature also allows these IPs to be configured to expire and be relinquished often. All of these traits combine to limit the strain of hotspot users on a Service Provider's network. Also as previously discussed, multiple verticals can also be balanced against one another. For instance, there may be a large public hotspot which is available in a park with lots of daytime foot traffic, and another hotspot a few blocks away along a corridor of night clubs. The same hardware, software and network resources that feed the daytime hotspot can then be repurposed to serve the night club hotspot at night with no perceived impact to the end-users whom are consuming the service. Hotspot services are the ideal use case for an Edge deployment model which will ease operational support, improve the service availability, while also providing a lower cost service for both the Service Provider and their Customers.

4.2. Premise Edge – Enterprise/School/Medical and Venues

Private Enterprise services will typically have more requirements for accessing local resources for transfers including printing, servers or client to client communications. They will also tend to have more requirements for segmentation between specific departments or areas of the building, which will drive their design to be more complex than a single flat network while possibly still requiring interconnectivity. They are also more likely to have premise-based 3rd party solutions which they require to be integrated. This will especially be the case at Enterprise customers where they are in the process of transitioning to a Managed Service. These customers have a lot to risk and may not feel comfortable swiftly transitioning to a Managed Service model for the entire offering. They may also have legacy network components which are tightly integrated to their business or just require significant efforts to successfully migrate. Some of these legacy services may also have security segmentation requirements as previously discussed. As larger business customers realize the Managed Services value proposition, they are more likely to implement these strategies in a cap and grow model. This will require the Service Provider to be prepared to aid the customer during this transition, which will likely drive complex integration requirements on the premise. Similarly, large venues tend to have very custom requirements depending on the exact venue type such as a Convention Center, Stadium, or Amphitheatre, and sometimes even depending upon the event that they are hosting. In both cases, these types of business customers will likely require the Premise Edge architecture to be deployed to ensure the necessary flexibility to meet their requirements.

4.3. Hybrid Use Cases

The above scenarios reflect the extremes, which are beneficial in determining the outer bounds of the solution. Our architecture analysis confirmed that both the Edge and the Premise Edge architecture offers benefits in certain use cases. Customers whom have solution requirements somewhere between those of a guest/public hotspot service and the large enterprise or venue service are another challenge. These hybrid customers likely include SMB and MDU segments. We will also need to discuss customers who require multiple services including both a guest/public hotspot and a private enterprise service.

4.3.1. SMB and MDU

The SMB and MDU use cases are quite challenging architecture propositions. In most cases, these have slightly more requirements than a guest/hotspot service, but less requirements than a private enterprise. In short, these verticals have requirements for elements of both solutions. Drivers for the Premise Edge architecture likely include some premise local communications, a desire for the benefits of customer security and segmentation, and also include the occasional advanced integration on premise. Similarly, they also would benefit from the Edge architecture with improved operations and reliability, while also reducing the cost and resources driving a lower cost product for the customer. Arguably, both segments will have users that likely fall on either end of the spectrum – pun intended. This will likely result in some customers falling into both architectures, and the SP having reason to support both architectures, and possibly a business process which helps direct the customer to the right solution early in the process. This could be in the form of sales questionnaires or surveys. Effectively driving towards a Premise Edge architecture for the advanced customer, and an Edge architecture for the basic customer, with the flexibility to upgrade the basic customer should they suddenly require a more advanced integration or premise-based communication.

4.3.2. Multi-Service Customer

The simplest solution for the customer which requires both a guest/hotspot service while in tandem delivering a private enterprise service, would be to deliver this based on the most stringent requirements. Since the large business private network would drive a Premise Edge hardware deployment, the hardware would already be deployed on-site, and dedicated to the customer, so leveraging this same hardware for the guest/hotspot service would be a plausible solution. An additional consideration is to tunnel guest traffic to reduce the size of the hardware deployed on-site and better leverage resources which are centrally deployed, or if tunneling an increase is needed to the size of the mobility domain. A hybrid approach may be a viable consideration, especially if the SP decides to deploy guest services for basic tiers in the Edge. If those services exist already, there may be significant costs reduced at scale by leveraging the Edge resources for Guest services. This architecture is shown below:

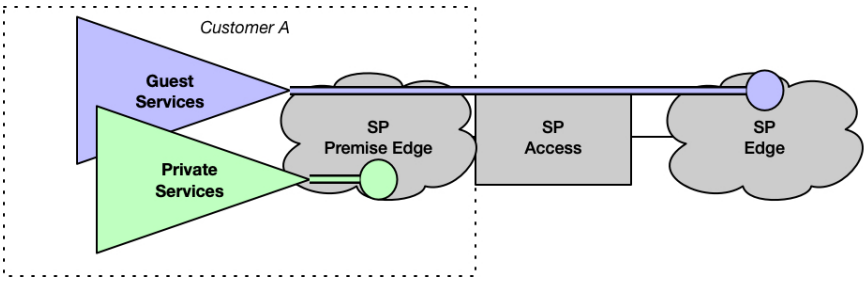

Figure 10 - The Hybrid Approach

Conclusion

The Managed Wi-Fi Service and all associated features are an important offering for Service Providers. These service offerings are ripe with challenges and opportunities for both the business and the solution architects. There are numerous functions that are required to be delivered in the solution to successfully manage an access service for the end-users. Elegance in the delivery and the solution architecture will be critical to the Service Provider's overall success. While either methodology could be leveraged to deploy to any of these verticals, there are benefits to both the Service Provider and the customer in both the Edge and the Premise Edge architectures. The Premise Edge architecture favors premise-based local communications, customer security and segmentation, advanced customer integrations, and high-performance applications, making it ideal for the increased demands of private enterprise or venue deployments and similar verticals. The Edge architecture improves operations and reliability, while also reducing cost and better leveraging available resources, improving both the availability of the service and reducing the cost to both the Service Provider and the customer. This makes the Edge scenario a preferred architecture for guest or hotspot services.

Between these extremes are the SMB and the MDU, which may be best served by creating a tiered offering which supports both deployment methodologies. The final consideration is the customer whom requires both a private network with complex local requirements for traffic and integrations, while also requiring a guest or hotspot service. These customers would likely require the Premise Edge which could be utilized to meet the needs of both services. Ideally the Service Provider could deploy elements at both the Premise Edge and the Edge architectures for these customers, creating a hybrid delivery architecture. This hybrid architecture carries the burden of providing infrastructure at both the Premise Edge and the Edge, but which places the Service Provider in the best position to deliver the Managed Services of today, while being prepared for the Managed Services requirements of the future.

Abbreviations

AAA	Authentication, Authorization and Accounting
AP	access point
BSS	Business Support Systems
BYOD	Bring Your Own Device
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
EAP	Extensible Authentication Protocol
EPC	Evolved Packet Core
IaaS	Infrastructure as a Service
LAN	Local Area Network
LMA	Local Mobility Anchor
OSS	Operational Support System
NAT	Network Address Translation
NFV	Network Function Virtualization
PaaS	Platform as a Service
PMIPv6	Proxy Mobile IPv6
RF	Radio Frequency
RFC	Request for Comment
RRM	Radio Resource Management
SFC	Service Function Chaining
SMB	Small Medium Business
SON	Self-Optimizing Networks
SP	Service Provider
TWAG	Trusted Wireless Access Gateway
WAG	Wireless Access Gateway
WLAN	Wireless Local Area Network
WPA2/3	Wi-Fi Protected Access 2/3

Bibliography & References

Firmin, Frédéric. "The Evolved Packet Core". 3GPP. 2019. https://www.3gpp.org/technologies/keywords-acronyms/100-the-evolved-packet-core. Accessed July 15, 2019.

Rekhter, Y., Moskowitz, B., Karrenberg, D., J. de Groot, G., Lear, E. "Address Allocation for Private Internets". RFC 1918. February 1996. https://tools.ietf.org/html/rfc1918. Accessed July 10, 2019.

Srisuresh, P., Egevang, K., "Traditional IP Network Address Translator (Traditional NAT)". RFC 1631. January 2001. https://tools.ietf.org/html/rfc3022. Accessed July 10, 2019.

van der Meulen, Rob. "What Edge Computing Means for Infrastructure and Operations Leaders". Gartner, Inc. October 3, 2018. https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders/. Accessed July 1, 2019.

Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., Azinger, A. "IANA-Reserved IPv4 Prefix for Shared Address Space". RFC 6598. April 2012. https://tools.ietf.org/html/rfc6598. Accessed July 10, 2019.

