



Smart City Network Architecture Guide

Table of Contents

- 1 About this document 3
 - 1.1 Purpose 3
 - 1.2 Audience 3
 - 1.3 Scope..... 3
- 2 Introduction and use cases 3
- 3 Solution Overview 5
- 4 Reference Architecture 7
- 5 City Net..... 8
 - 5.1 Introduction..... 8
 - 5.2 Business Drivers and Technical Requirements..... 8
 - 5.3 L2/L3 VPNs and IoT Containers 9
 - 5.4 The Intelligent Fabric in City Nets..... 11
 - 5.5 Solution Highlights 11
 - 5.6 Why ALE’s Intelligent Fabric for Smart City Nets? 12
- 6 Municipal Cloud 12
 - 6.1 Introduction..... 12
 - 6.2 Business Drivers and Technical Requirements..... 12
 - 6.3 The Virtualized Data Centre..... 13
 - 6.4 End-2-End Virtualized Hybrid Multi-Tenant Cloud 13
 - 6.5 Solution Highlights 14
 - 6.6 Why ALE’s Intelligent Fabric for Smart Municipal Clouds?..... 14
- 7 Architecture 15
 - 7.1 Introduction & Overview..... 15
 - 7.2 Core and Border Topology..... 16
 - 7.3 Backbone and Edge Topology 17
 - 7.4 VPN Services..... 18
- 8 Management Service Architecture 26
- 9 Internet Service Architecture 29
- 10 Security Framework..... 32
 - 10.1 Securing IOT Devices..... 33
 - 10.2 Securing the Perimeter 33
 - 10.3 Securing the Network 33
- 11 Conclusion..... 37
- 12 Acronyms 37
- 13 Related documents 38

1 About this document

1.1 Purpose

The purpose of this architecture guide is to present the requirements and considerations relevant to Smart City networks along with design options, best practices and configuration guidelines.

1.2 Audience

This design guide is intended for network architects and network engineers involved in the design, implementation and maintenance of Smart City networks.

To take advantage of this document, it is expected that the reader will be familiar with Shortest Path Bridging and will have a solid understanding of various networking technologies at the ACPS or similar level.

Please refer to [1] for a short introduction to SPB.

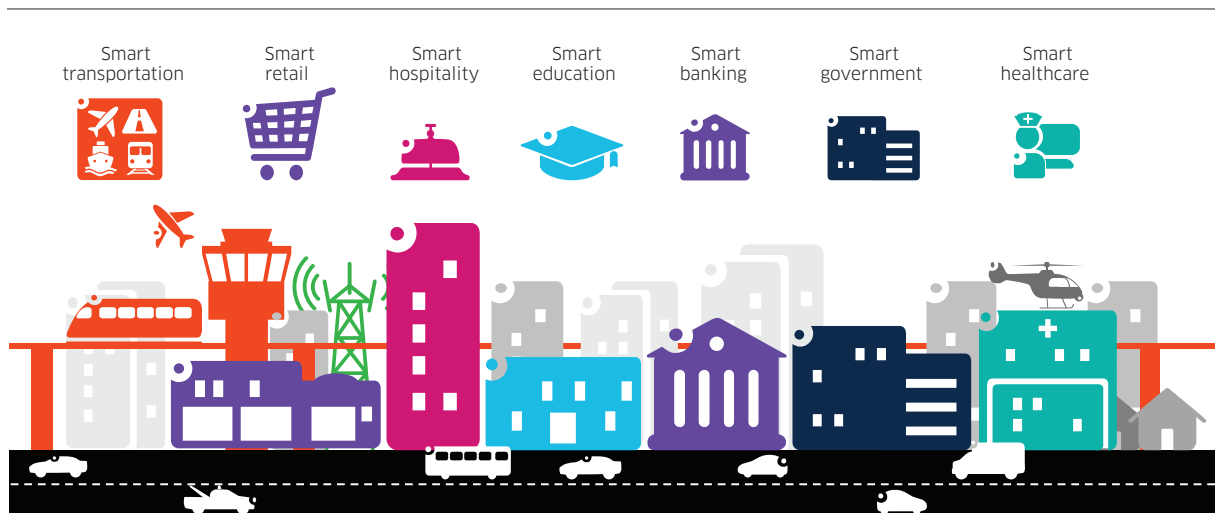
1.3 Scope

This document will focus on Smart City Nets and Municipal Clouds. Municipal Wi-Fi, Smart Transit, Smart Roads & Highways and Smart Civic Venue solutions are out of scope.

2 Introduction and use cases

According to the Smart City Council, “A smart city uses information and communications technology (ICT) to enhance its liveability, workability and sustainability. First, a smart city collects information about itself through sensors, other devices and existing systems. Next, it communicates that data using wired or wireless networks. Third, it analyses that data to understand what’s happening now and what’s likely to happen next.”¹

Figure 1. Smart City



¹ <https://rg.smartcitiescouncil.com/readiness-guide/article/definition-definition-smart-city>

Through this data collection, sharing and analysis, a Smart City can improve on multiple aspects. Let's review some of the most common use cases with practical examples shown in Table 1 below.

Table 1. Use Cases and Examples

Use case	Example
Smart Lighting	<p>Save energy with remotely-controlled on-off timing and dimming depending on time of year, weather conditions and motion detection, such as an approaching cyclist.</p> <p>Example: Public lighting represents 20% of energy consumed by Barcelona City Council.² Replacing old-fashioned sodium lamps with LED ones provides better lighting and reduces energy usage by more than 52%.</p> <p>With the addition of remotely-controlled on-off timing and dimming depending on time of the day, or year, weather or other conditions as well as motion detection, such as approaching cyclists, energy savings of up to 72% are feasible.</p>
Smart Parking	<p>Reduce parking search time, fuel usage, CO2 emissions, and parking violations by detecting available parking spots and guiding drivers to them with an app.</p> <p>Example: City drivers spend a significant amount of time looking for a parking spot. Sensors can detect parking spot availability and a Smart Parking app can guide drivers to it. In the city of San Francisco, parking search time was reduced by 43% which not only saves time but also reduces fuel usage and CO2 emissions by 30%³. In addition, in-app payment makes it easier to pay, resulting in a 23% decrease in the number of parking violations and citations.</p>
Smart Transit	<p>Improve rider safety and satisfaction as well as ridership with video surveillance, on-board Wi-Fi and entertainment and real-time schedule information. Real-time vehicle location tracking provides live updates through an app or smart display at the bus stop.</p> <p>Example: In Bucharest, visually-impaired riders are alerted that their bus is approaching⁴. And bus drivers can be informed that someone may need assistance with boarding at the next stop. With Open Data, real-time location information can be leveraged by 3rd party solutions for wayfinding and more.</p>
Smart Waste Management	<p>Reduce waste collection activities, fleets, fuel usage and CO₂ emissions with smart bins that can compact waste and inform their fill level.</p> <p>Example: Solar-powered smart rubbish bins can compact waste, increasing their capacity by a factor of 5. Bins use sensors to determine their fill status and communicate it to the management platform. In Australia⁵, smart bin deployment has resulted in a 75% reduction in waste collection activities at Bondi Beach. With real-time access to bin fill-level information, waste collection is optimized with smarter routes, resulting in reduced labour costs, fleets, fuel usage and CO₂ emissions.</p>
Smart Metering	<p>Reduce water and electricity waste or theft with smart meters that can measure and report usage in near real-time.</p> <p>Example: Smart Meters measure energy and water usage in short intervals and report this data back to the energy or water company. Utilities can set dynamic usage-based pricing to reduce usage at times of peak demand. Smart metering also helps reduce waste and eliminate the need to dispatch technicians for costly manual meter readings as well as facilitate trading of excess energy. Smart Meters can detect and alert of energy theft in public infrastructure such as public lighting.</p>
Video Surveillance	<p>Improve public safety with license plate recognition, vehicle tracking, facial recognition and analytics.</p> <p>Example: Cities cannot be smart if they are not safe. Smart video surveillance solutions add features such as license plate recognition, vehicle tracking, and analytics. Facial recognition can identify a suspect, or persons of interest, such as people on a watch list, from a crowd. But it's not just about safety. Facial recognition can be up to 50% faster⁶ compared to tap cards when used at metro gates.</p>

2 <http://ajuntament.barcelona.cat/ecologiaurbana/en/services/the-city-works/maintenance-of-public-areas/energy-management/street-lighting-management>
3 <https://www.sfmta.com/press-releases/sfpark-evaluation-shows-parking-easier-cheaper-pilot-areas>
4 <http://www.rfidjournal.com/articles/view?13899>
5 <https://www.iothub.com.au/news/council-introduces-smart-bins-to-bondi-464281>
6 <http://www.straittimes.com/singapore/transport/face-reading-system-can-replace-fare-cards>

3 Solution Overview

If we pay close attention to the use cases presented in Section 2, we may notice that they require a network to function. Some use cases are unfeasible without network connectivity (e.g. smart metering). Other use cases cannot deliver their full benefits in an offline manner (e.g. smart lighting).

Figure 2 illustrates a high-level view of a Smart City infrastructure.

Starting at the bottom of this diagram, we have IoT Devices and Users. IoT devices are a variety of sensors and actuators, the sensory nerves and the muscles. IoT devices gather data and are commanded by the City's brain, software and applications running in a Municipal Cloud and the City's Command and Control Centre.

A Smart City leverages its ICT infrastructure to maintain smart interactions with residents and visitors, the users. Users consume and produce data as they go about their daily lives utilizing the City's resources such as public transport, parking, Wi-Fi and so on. Users interact with City's agencies directly through official websites and applications, or, indirectly, with 3rd party websites and applications that access agency Open Data through APIs. Users connect through fixed (e.g. POL or ADSL) or mobile (e.g. 4G or LTE) broadband connections or Wi-Fi using a personal device, such as mobile phone, or a City asset, such as an information kiosk.

IoT devices can connect using the same technologies available to users. In other cases, however, other technologies are more appropriate. An IoT gateway may be required to connect such devices to IP-based infrastructure.

Let's have a look at some of those technologies.

SCADA: SCADA is a control technology which is prevalent in many industrial and infrastructure processes. As an example, SCADA may be found in power utilities, water, or wastewater collection and treatment. Modern SCADA devices can connect to an IP network directly. Legacy SCADA devices, however, utilize a serial interface and require a gateway to connect to the IP network.

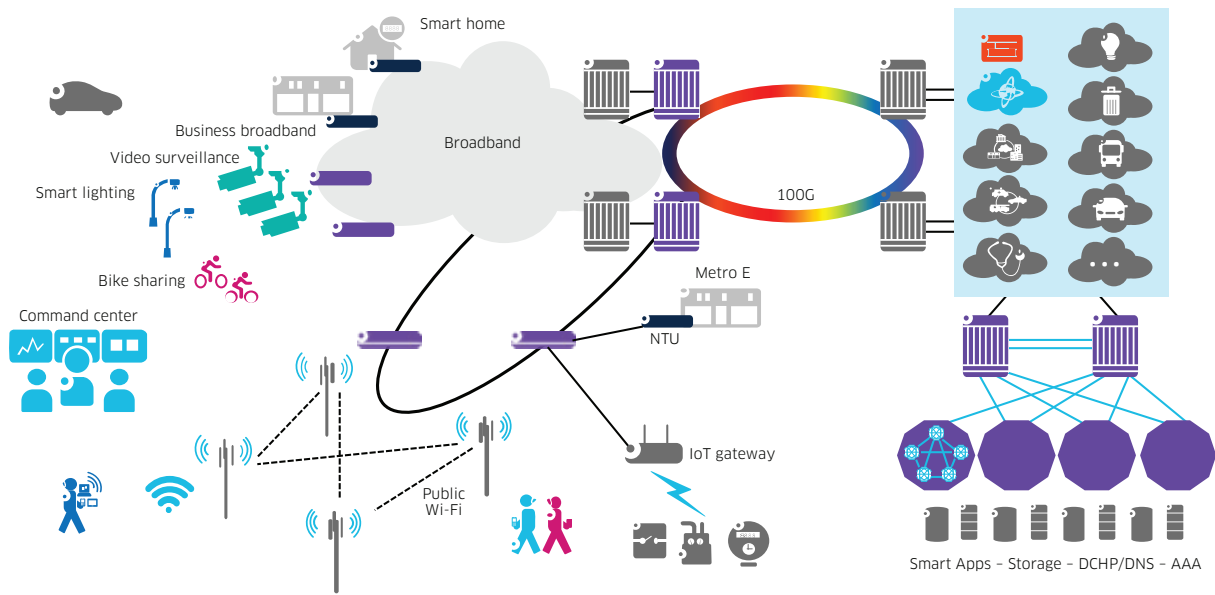
WPAN: WPAN stands for "Wireless Personal Area Network". These technologies are designed to connect medium-bitrate devices operating on a battery over a short distance (i.e. 10 to 20 meters). This distance can be extended by "meshing" to other WPAN devices. Bluetooth and Zigbee are examples of WPAN technologies. WPAN technologies are useful in short-range applications such as those within a "Smart Building".

LPWAN: LPWAN stands for "Low-Power Wide Area Network". These technologies are designed to connect low-bitrate devices operating on a battery, such as sensors, over long distances. Examples of LPWAN technologies include Lora/LoRaWAN, Sigfox and NB-IOT. These technologies are useful in low-density, long-range deployments.

The City's network is the spinal cord linking IoT devices with the software, applications and people that analyse their data and control them. The City's network delivers network services to government agencies, businesses and users over different technologies. Section 5 introduces ALE's solution for Smart City Nets.







The Command Centre is the facility where government personnel manage incident and analyse data for better city planning. The Command Centre also acts as the City's emergency and disaster management centre supporting multiple agencies (e.g. law enforcement, public transport, environmental, etc.).

Figure 2. Smart City Infrastructure



ALE’s network solutions enable Smart City use cases. In this document, we will focus on 6 Smart City solution sets shown in Figure 3 below. This architecture guide focuses on the City Net and Municipal Cloud solution sets. These solution sets are described in Sections 5 and 6.

Figure 3. Smart City Solution Sets

 <p>City Net City-wide fabric to interconnect government agencies, residents, visitors, businesses and IoT devices.</p>	 <p>Municipal Wi-Fi Municipal Wi-Fi provides amenity to visitors and residents and enables multiple smart city use cases.</p>
 <p>Municipal cloud End-to-end virtualized and consolidated cloud enables smart city use cases by breaking siloes whilst at the same time achieving economies of scale.</p>	 <p>Smart Civic venue Location and wayfinding enables various use cases such as visitor guides, geo-fenced alerts and more.</p>
 <p>Smart transit Dedicated solutions for on-board communications and smart shelters.</p>	 <p>Smart roads and highways Reducing congestion, improving safety and keeping drivers informed in real time.</p>

4 Reference Architecture

Implementing these use cases requires breaking organizational siloes: information must be shared and budgets must be pooled across multiple government agencies for this to be technically feasible and cost-effective. A siloed architecture in which each vertical use case relies on its own infrastructure, middleware and applications increases complexity and cost.

The ALE reference architecture for Smart Cities is shown in Figure 4 below. This horizontal architecture makes common infrastructure and services layers available to use cases and applications.

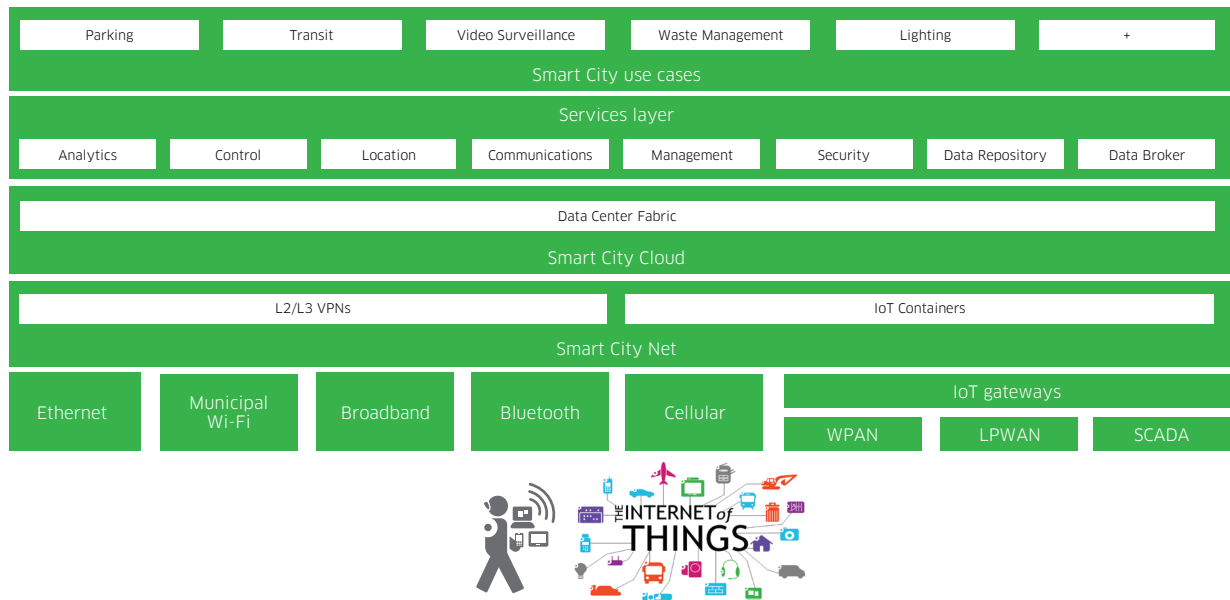
This approach has the following advantages:

- Facilitates sharing of information across different use cases and applications, a key requirement of Smart Cities.
- Consolidates network, cloud and services in shared layers which are abstracted from and consumed by applications, thus simplifying application development.
- Eliminates multiplication of infrastructure and middleware performing similar functions, thus reducing TCO.

The Alcatel-Lucent Enterprise Intelligent Fabric architecture enables this horizontal approach through:

- Automated network node provisioning (ZTP/RCD)
- Automated IoT device provisioning (Access Guardian/UNP)
- Service-oriented architecture (SPB/VXLAN/GRE)
- Northbound APIs (OmniSwitch/OmniVista)

Figure 4. ALE Smart City Reference Architecture

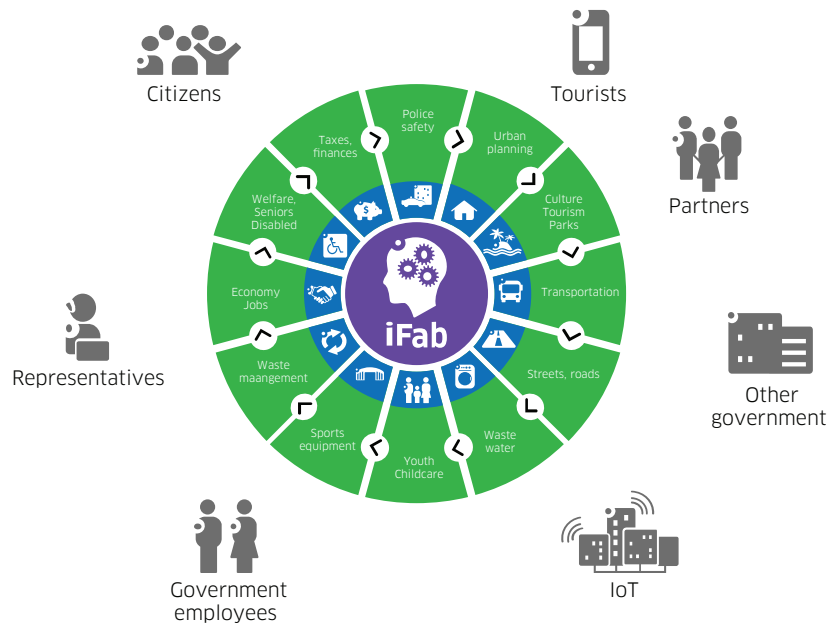


5 City Net

5.1 Introduction

Smart Cities rely on a robust, yet flexible network foundation to interconnect sensors, people and businesses with cloud-based applications. The City Net is the cornerstone of the Smart City since all use cases are reliant on it.


Figure 5. Cornerstone of a Smart City



5.2 Business Drivers and Technical Requirements

Let's review the 3 main business drivers of a City Net in Figure 6 below.

Figure 6. City Net Business Drivers



We can now translate these business drivers into technical requirements in Table 2 below.

Table 2. City Net Technical Requirements

Technical requirements	
Multi-Tenancy	The City Net provides services to multiple tenants. L3 VPN services are needed for most Government agencies, schools and hospitals. L2 VPN services are needed in specific situations such as DC interconnect. IoT devices of different class need to be isolated in their own L2 or L3 container for security.
Availability	Availability is of utmost importance since the City Net provides services to multiple agencies and supports multiple use cases. The City Net must implement redundancy at all layers to protect from node or link failure and to facilitate in-service maintenance. In the event of a failure, the network must reconverge in under 1s.
Scalability	The City Net architecture must be scalable to support the required number of tenants, containers, devices and users as well as bandwidth and multicast flows, etc.
Simple Operations	The City Net can be comprised of thousands of network nodes and IoT devices. Because of its size, template-driven automatic provisioning of nodes and devices is necessary to reduce deployment cost, time and errors.
Security	Because a Smart City Net supports mission-critical infrastructure, security is even more critical than traditional Enterprise security. A multi-layered approach is required including security hardening of nodes and devices, network admission control (NAC) and role-based access, protection of data integrity and confidentiality as well as quarantine and remediation among others.
Environmental	Smart City Net nodes will be deployed in locations such as within a roadside cabinet where they will be subject to extreme temperatures, dust, vibration and shock. The network devices must tolerate such harsh conditions.

5.3 L2/L3 VPNs and IoT Containers

Figure 7 below illustrates the concept of L3 VPNs for government agencies, schools and hospitals. In this diagram, schools, hospitals and other government agencies are tenants of the City Net and each have their own L3 VPN. VPN technology ensures all these different tenants can coexist on the same infrastructure and that tenant data is kept private from other tenants. In a L3 VPN, tenants exchange routes with the City Net and the City Net routes tenant data from site to site over the most optimal route without traffic tromboning back and forth to a central site for routing. In addition, tenant data is isolated from other tenants and tenant routes are not advertised to other tenants.

Figure 7. City Net L3 VPN

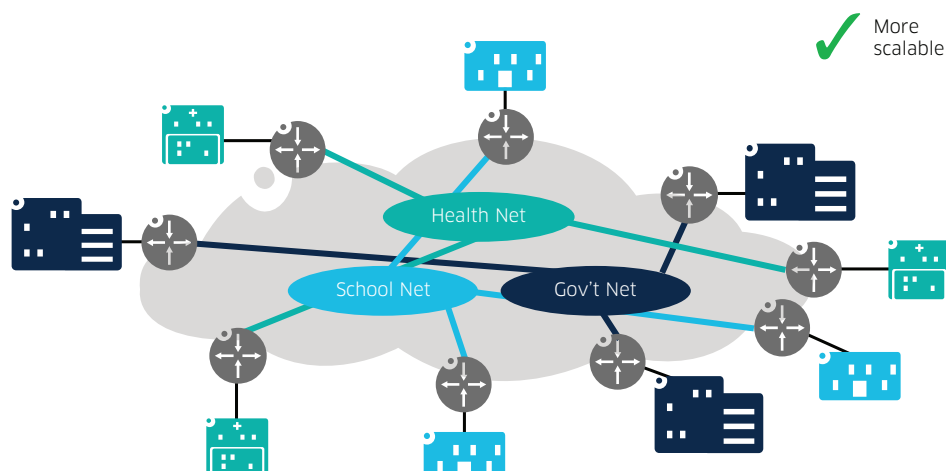



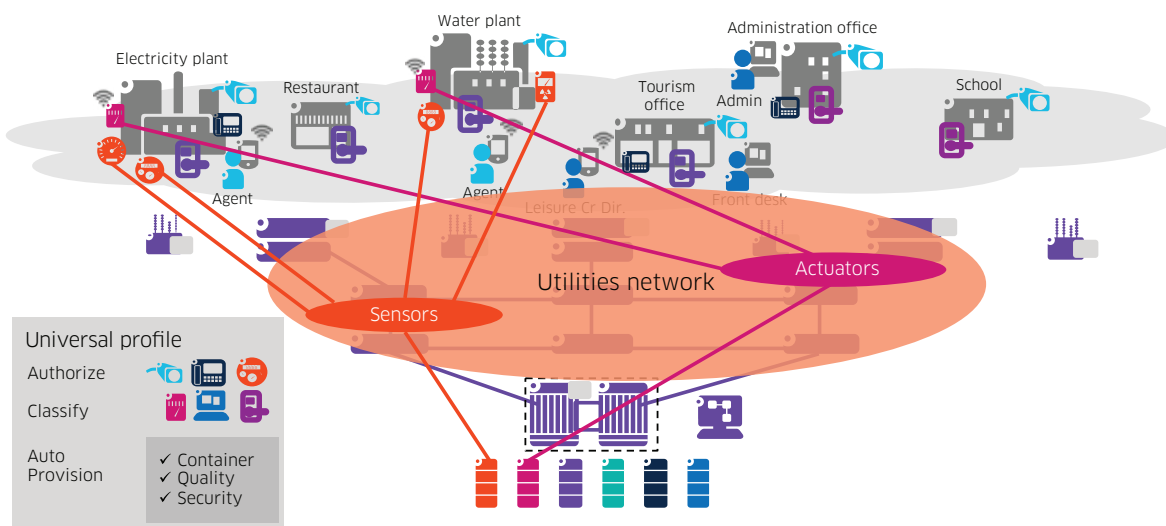
Figure 8 below illustrates the concept of L2 VPNs for DC interconnect. In this diagram, 2 DC locations are interconnected by L2 Services. These DCs provide services to different tenants. Unlike L3 VPN services, in a L2 service there is no exchanging of routes between tenant and City Net, any routing is done by the tenant and data is simply bridged across the City Net. This L2 communication is a requirement in certain scenarios, such as DC interconnect, because it enabled VM mobility an HA across both locations. A L2 VPN also isolates tenant data from other tenant's data.

Figure 8. City Net L2 VPN



Finally, Figure 9 below illustrates the concept of IoT Containment. In this figure, multiple different types of IoT devices such as sensors, actuators, cameras and door locks are connected. IoT devices are grouped into different containers: Utility, Administration, Security, etc. This containerization increases security because a container is isolated from other containers and devices in different containers can only communicate through a Firewall. In addition, devices are mapped to containers according to the device type using authentication (MAC or certificate-based 802.1x). Once the device type and container is determined, the device is bound to a profile which will also restrict communication with other devices, even if on the same container, and apply fine-grained QoS policies to the device.

Figure 9. IoT Containers



5.4 The Intelligent Fabric in City Nets

The Intelligent Fabric is based on the IEEE standard protocol Shortest Path Bridging. Let's compare this technology with others such as legacy (VLANs and Spanning Tree Protocol) and MPLS (Multi-Protocol Label Switching). As we can see, SPB is superior for small to medium-sized City Nets because it is much simpler to deploy and operate, resulting in lower total cost of ownership.






Table 3 - SPB Comparison

	Legacy (VLAN/STP)	MPLS	SPB
Virtualization	VLAN	VPN	VPN
Availability	Slow fault recovery	Fast fault recovery	Fast fault recovery
Performance	Single active path	Multiple active paths	Multiple active paths
Scalability	Small networks	Very large networks	Large networks
Complexity/Operational Costs	Low/High	High/High	Medium/Low
	Limited	High Cost	Best Option

5.5 Solution Highlights

Table 4 below shows the most relevant ALE products for a Smart City Net solution and the typical role that they may fulfill along with the key network and environmental features and market-specific certifications that make them a good fit for the role.

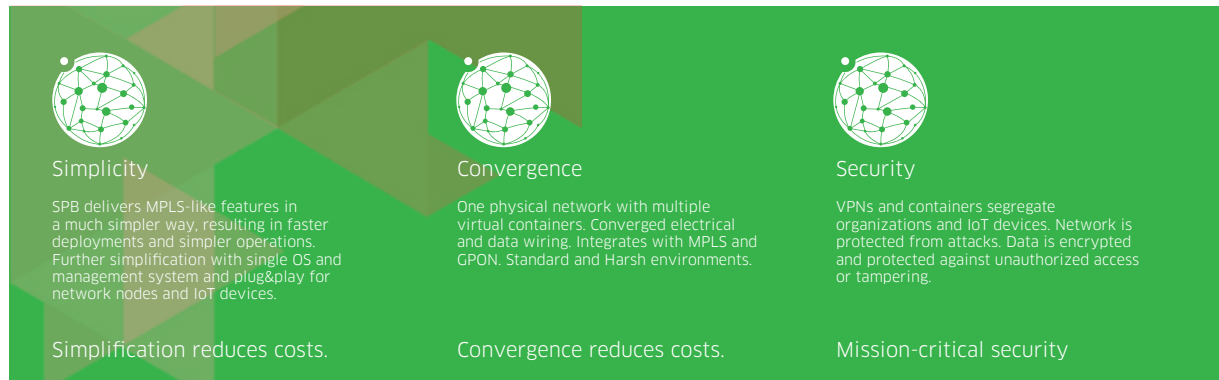
Table 4. City Net Solution Highlights

Product	Role	Key environmental features	Key network features	Market-specific certifications
 OmniSwitch 6865	Advanced hardened Access or collapsed Access + backbone	<ul style="list-style-type: none"> Fan-less Shock, vibration, temperature IP-30 rating* 	<ul style="list-style-type: none"> ERPV2, SPB OAM POE/POE+*/UPOE* Virtual Chassis/ISSU* IEEE 1588v2 	<ul style="list-style-type: none"> NEMA-TS2 (Traffic Controller Assemblies)* CC EAL-2 & NDcPP FIPS-140
 OmniSwitch 6465	Value hardened access	<ul style="list-style-type: none"> Fan-less Shock, vibration, temperature IP-30 rating* 	<ul style="list-style-type: none"> ERPV2 OAM* POE/POE+*/UPOE* Stacking Alarm replay inputs MACSec* IEEE 1588v2 	<ul style="list-style-type: none"> NEMA-TS2 (Traffic Controller Assemblies)* CC EAL-2 & NDcPP FIPS-140
 OmniSwitch 6350/6450/6465T/6560	Value access/smart building	<ul style="list-style-type: none"> Extended temperature range (6465T only)* 	<ul style="list-style-type: none"> ERPV2 OAM* POE/POE+*/UPOE* Stacking 	<ul style="list-style-type: none"> CC EAL-2 & NDcPP FIPS-140
 OmniSwitch 6860/6900/9900	Aggregation, backbone, core	-	<ul style="list-style-type: none"> ERPV2, SPB OAM Virtual chassis/ISSU MACSec* IEEE 1588v2 	<ul style="list-style-type: none"> CC EAL-2 & NDcPP FIPS-140
 OmniVista 2500 NMS	Network Management System	-	<ul style="list-style-type: none"> Network Management Unified access Analytics 	-

5.6 Why ALE's Intelligent Fabric for Smart City Nets?

Let's review the key reasons that make ALE's Intelligent Fabric a great fit for the Smart City's network in Figure 10 below.

Figure 10. Why ALE's Intelligent Fabric in Smart City Nets



6 Municipal Cloud

6.1 Introduction

The Municipal Cloud serves two main purposes:


- Hosting the multiple systems, applications, databases, etc, powering the various Smart City use cases (e.g. Smart Lighting, Smart Waste Management, etc). The Smart City's "brain".
- Hosting systems, applications, databases, etc, for Smart City Net customers, such as government or private organizations, hospitals and schools.

The Municipal Cloud may also provide additional services such as Email. Web, caching, storage, content delivery network (CDN), voice communications and IP TV.

6.2 Business Drivers and Technical Requirements

Let's now review the 3 main business drivers of a Municipal Cloud in Figure 11 below.

Figure 11. Municipal Cloud Business Drivers



We can now translate these business drivers into technical requirements in Table 5 below.

Table 5. Municipal Cloud Technical Requirements

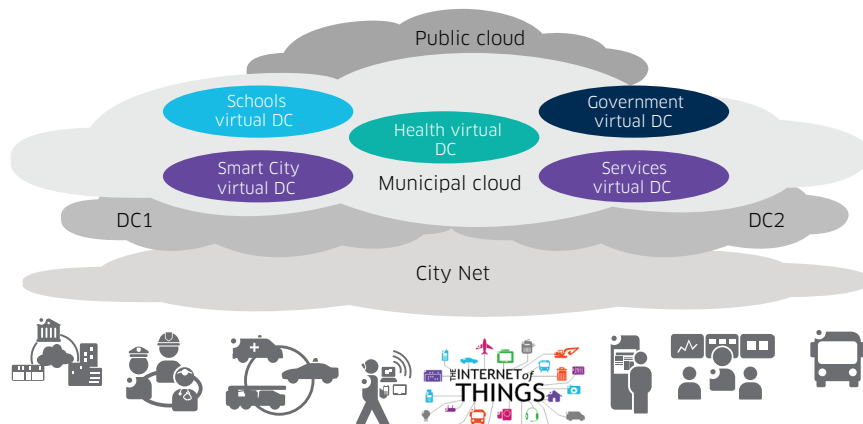
Municipal cloud	Technical requirement
Multi-Tenancy	The Municipal Cloud will provide services to multiple tenants including government and private organizations. Tenants must be isolated from one another for security reasons and to avoid conflicting IP addressing issues. Tenant traffic must be seamlessly stitched from the City Net to Municipal Cloud, across DC sites and between the private and public cloud.
Business Continuity	The Municipal Cloud must recover from multiple failure scenarios without disruption to customers. To accomplish this, Active/Active intra- and inter-DC redundancy with no single point of failure and sub-second convergence is required.
Elasticity	Customers must be able to scale their workloads up and down both on-site at the Municipal Cloud DC sites or off-site across the public cloud.
Fluid Operations	Manual intervention to Moves, Adds and Changes should be minimized. The network must automatically adapt to event such as Virtual Machine Mobility. Operators should be able to perform software upgrades in-service without disruption to customer traffic.
Performance	The Municipal Cloud must offer high throughput and low latency, lossless capability for storage convergence as well as monitoring and optimization of workload performance.
Sustainability	The Municipal Cloud must operate within a compact footprint and with low power consumption and cooling requirements.

6.3 The Virtualized Data Centre

The municipal cloud also adopts a layered horizontal architecture to break siloes and reduce costs by consolidating and sharing common functions across the board.

The Municipal Cloud can extend across multiple physical Active/Active Data Centres for high availability and load sharing. The Municipal Cloud can also extend to Public Clouds for elasticity or to exchange information with those clouds. Containers abstract these aspects such that, to the application and the customer, the container appears to be a virtual DC. Please refer to Figure 12 below.

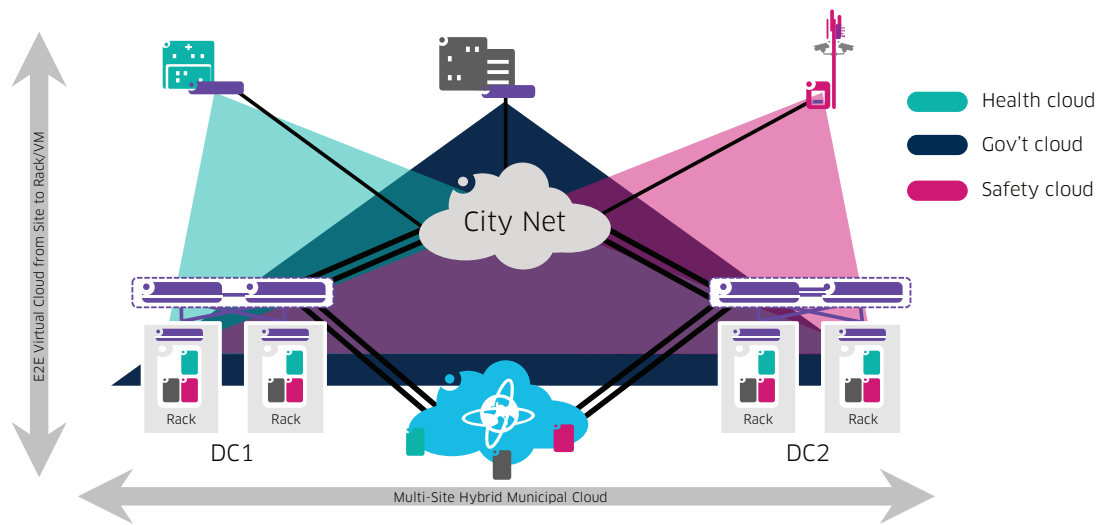
Figure 12. Virtualized Municipal Cloud



6.4 End-2-End Virtualized Hybrid Multi-Tenant Cloud

City Net containers extend deep into the data centres and all the way through the Top-of-Rack (TOR) switch down to the Virtual Machine (VM). This deep containerization keeps different organizations applications and data isolated end-2-end. This not only improves security, but also facilitates migration of hardware and applications from on-premises DC to the Municipal Cloud. Regardless of where the application is hosted (primary or secondary DC or even the public cloud) it is mapped to the appropriate container. Secure containerization is maintained seamlessly across DC sites and the public cloud and from the IoT device at the edge of the network all the way to the application that controls it. Refer to Figure 13 on the following page.

Figure 13. End-2-End Virtualized Hybrid Multi-Tenant Cloud



6.5 Solution Highlights

Table 6 below shows the most relevant ALE products for a Smart Municipal Cloud solution and the typical role that they may fulfill along with the key network and environmental features and market-specific certifications that make them a good fit for the role.

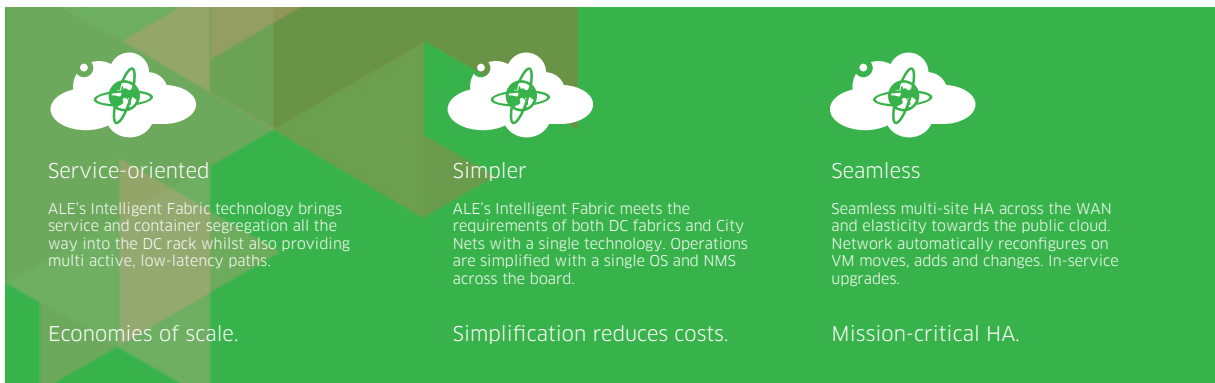
Table 6. Municipal Cloud Solution Highlights

Product	Role	Key network features	Specific features	Ports	
OmniSwitch 9900	Modular Spine Node	<ul style="list-style-type: none"> Virtual chassis SPB VM Snooping VM Mobility VM Profiles Lossless Ethernet FIP Snooping RESTful API OpenFlow Python programmability ZTP App fingerprinting 	In-line routing	1G/10G/25G/40G/50G/100G	
OmniSwitch 690-C32	Stackable 100G Spine Node		VXLAN Gateway	1G/10G/25G/40G/50G/100G	
OmniSwitch 6900-Q32	Stackable 40G Spine Node				1G/10G/40G
OmniSwitch 6900-V72	Stackable 25G Leaf Node				1G/10G/25G/40G/50G/100G
OmniSwitch 6900-X72	Stackable 10G Leaf Node				1G/10G/40G
OmniSwitch 6900-X20/40	Semi-modular 10G Leaf Node			FC Gateway	1G/10G/40G/FC
OmniSwitch 690-T20/40	Semi-modular 10G Copper Leaf Node				1G/10G/40G/10G-T/FC
OmniVista 2500	NMS	<ul style="list-style-type: none"> VM Mobility/Snooping SPB, VXLAN RESTful API Analytics 			

6.6 Why ALE's Intelligent Fabric for Smart Municipal Clouds?

Let's review the key reasons that make ALE's Intelligent Fabric a great fit for the Smart City's Municipal Cloud in Figure 14 on the following page.

Figure 14. Why ALE's Intelligent Fabric for Smart Municipal Clouds



7 Architecture

7.1 Introduction & Overview

A City Net may serve different purposes such as:

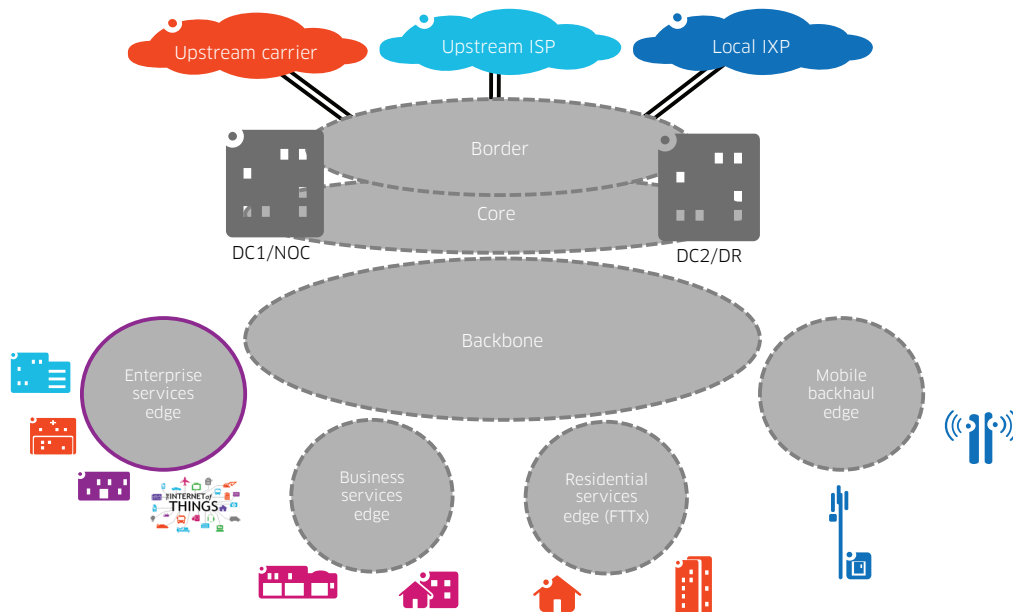
- VPN/Internet services for private and government organizations, hospitals, schools, etc.
- Smart City use cases such as Smart Lighting, Video Surveillance, Public Wi-Fi, etc.
- Residential or Business Internet Access
- Mobile backhaul.

This section will describe the architecture focusing primarily on the first two. The City Net offers the following services:

- L2 VPN: E-LINE (point to point) or E-LAN (multipoint) service.
- L3 VPN: Multipoint IP VPN.
- Internet: Transit to other local or regional/national ISPs.
- Wholesale: Local termination of other carrier or ISP services.

Let's refer to Figure 15. This diagram is showing the architecture for a small city network. A small city can be served with a single POP (Point of Presence). The POP is wholly contained within a single metropolitan area.

Figure 15. City Net Architecture Overview

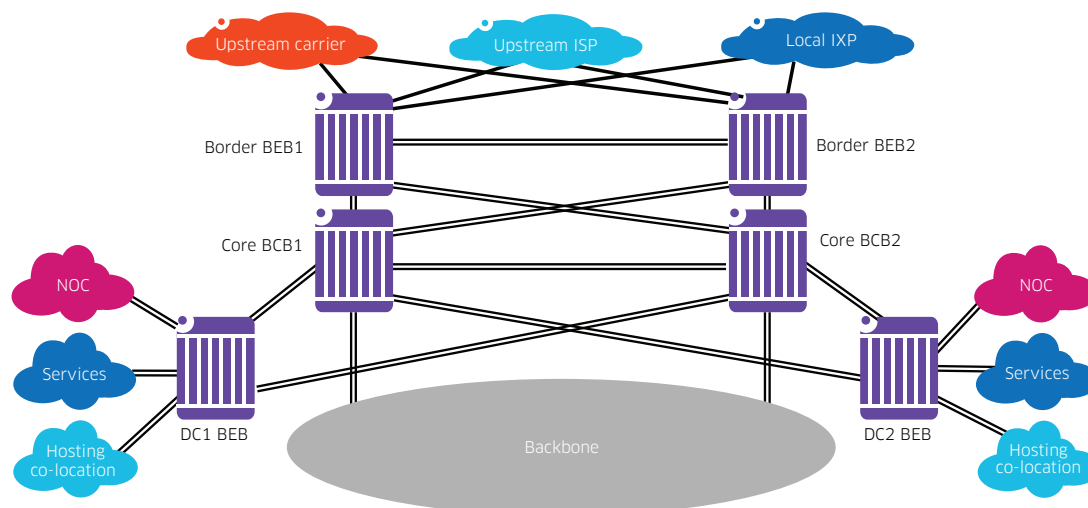


- **Core:** This is the City Net's Core and the rest of the City Net connects to it through the Backbone. The Core may be physically split between two sites for redundancy. These two sites are usually co-located with redundant DC1 and DC2 as well as the NOC and DR sites as shown in the diagram.
- **Border:** The Border block is the point of demarcation between the City Net and 3rd party upstream and local ISPs. Border nodes are also usually co-located with DC1 and DC2.
- **DC1 and DC2:** These two Data Centres operate in an Active/Active redundancy scheme. DC1 and DC2 provide essential network services such as DHCP/DNS and value-added services such as Email, WWW, firewalling, web-caching or CDN (Content Delivery Network), IP Telephony and IPTV. DC1 and DC2 also offer hosting, co-location and "cloud" services to the City Net's customers such as public or private organizations and businesses. DC1 and DC2 also host the applications and services implementing Smart City use cases such as Smart Lighting and Video Surveillance.
- **NOC and DR:** The Network Operations Centre is the facility where the City Net is monitored and managed from. A DR (Disaster Recovery) site may also replicate the NOC at a separate location such that NOC activity can quickly resume in the event of a major disruption.
- **Backbone:** The Backbone block is an aggregation layer linking Edge blocks with the Core. In smaller cities or communities, the Backbone and Edge blocks may be collapsed in a single block.
- **Edge:** Edge blocks are aggregation layers linking customer connections with the backbone. Different Edge blocks such as Enterprise, Business, Residential or Mobile Backhaul provide different services. For instance, the Enterprise edge provides Enterprise-grade L2/L3 VPN and Internet services whilst the Residential block provides Internet services over a FTTH (Fibre to the Home) connection.
- **Upstream ISP:** The City Net needs to connect to at least one ISP, such as a regional ISP, to communicate with the rest of the Internet.
- **Local IXP:** The City Net may connect with other local ISPs at an IXP (Internet Exchange Point). The benefit in this is that, unlike traffic transiting through the upstream ISP, peering with other local ISP is free of charge.
- **Upstream Carrier:** City Net customers may require services in locations not within the City Net's footprint (off-net). Conversely, other carrier customers may require network services within the City Net's footprint (on-net). Interconnection with other carriers at enables cross-selling of network services to cater to these situations.

7.2 Core and Border Topology

Let's refer to Figure 16 below.

Figure 16. Core and Border Topology



The Core and Border blocks serve the following functions:

- Backbone aggregation
- Interfacing with other ISPs or carriers
- NOC connection
- Services connection
- Hosting and co-location connection

The Core and Border blocks are comprised of three different node types according to the functions that they perform. In a smaller network however, more of these functions may be performed by the same node.

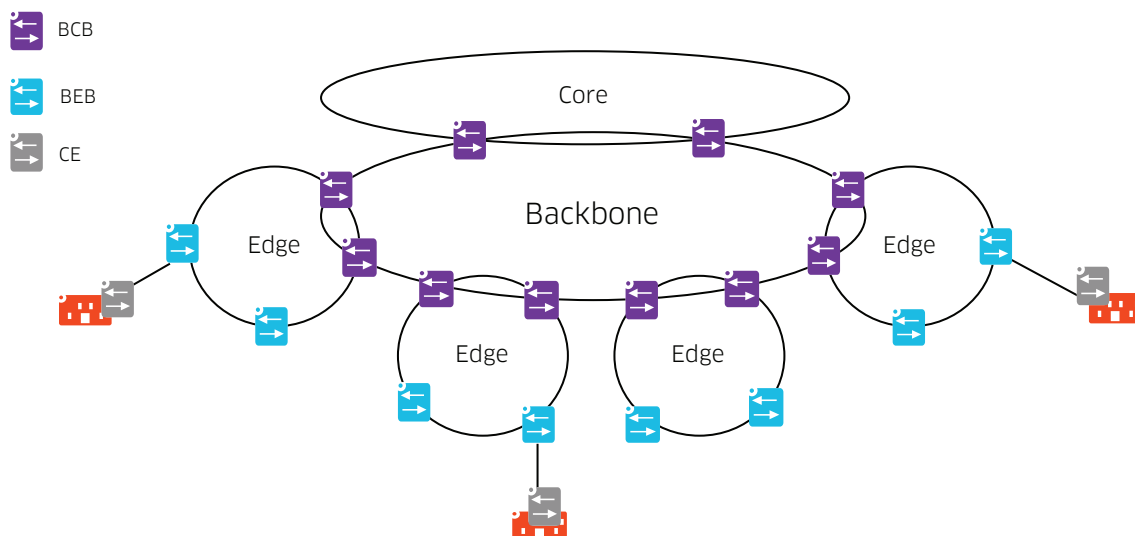
- **Core BCBs:** These nodes are pure Core nodes which means they do not directly connect to any external party (e.g. ISP or customer) and do not terminate any service, except for Internet and Management services which, as we will see later one, are implemented as separate services. For this reason, these nodes require little touch.
- **Border BEBs:** These nodes are a point of demarcation between the City Net and other ISPs or carriers. Routing and security policies are enforced on these nodes.
- **DC BEBs:** These nodes connect the NOC, essential services blocks and other customer services hosted or co-located at the DC sites.

Note that nodes in Figure 16 are illustrated as modular chassis but they can be standalone stackable devices or a Virtual Chassis. Also note that hairpins or s-hooks required for routing or service mapping are not shown.

7.3 Backbone and Edge Topology

Let's refer to the topology shown on Figure 17 below.

Figure 17. Backbone and Edge Topology



In this diagram, we are using ring topologies as an example. Ring topologies support redundant physically-diverse paths and minimize fibre runs when compared to other topologies such as star or mesh. For these reasons, ring topologies are very common when a large area needs to be covered. It should be noted however that the same design principles are applicable to any topology because SPB is powered by a link-state protocol (IS-IS) and therefore supports any topology. We are also showing one backbone block and several edge blocks. Larger networks may have several backbone blocks and smaller networks may collapse backbone and edge blocks in a single block.

Backbone nodes are SPB BCB nodes because they do not terminate any customer service. The only exception is, as we will see later, Internet and Management services.

Edge blocks are comprised of Edge BEBs. These nodes terminate customer services such as L2/L3 VPN. Edge blocks redundantly connect to diverse Backbone BCBs.

Customer Edge (CE) devices connect to one or more BEBs. Please refer to Section 7.4.1 for the various CE connection redundancy options.

Note that nodes in Figure 17 may be standalone stackable, VC or modular chassis. Also note that hairpins or s-hooks required for routing or service mapping are not shown.

7.4 VPN Services

Figure 18 below illustrates a L2 service interconnecting three customer sites. This is an example of multipoint (E-LAN) L2 service. From the customer's point of view, this would be equivalent to having all three sites connected to an Ethernet switch. Within the City Net, customer traffic is isolated in its own ISID.

Figure 18. L2 Service


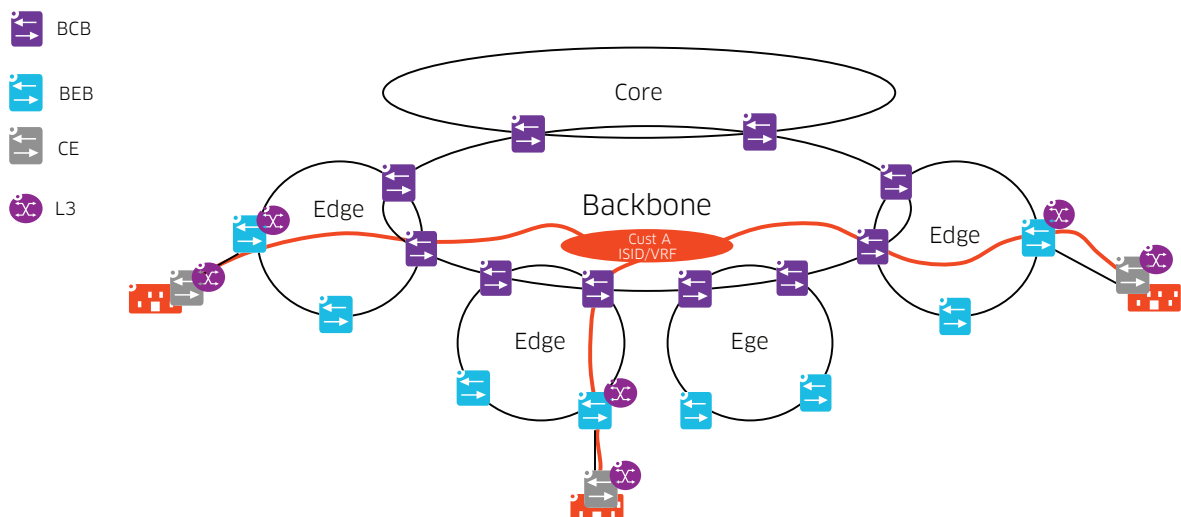


Figure 19 below illustrates a L3 service interconnecting three customer sites. This is an example of multipoint IP-VPN. From the customer's point of view, this would be equivalent to having all three sites connected to an Ethernet router. Customer sites can exchange routing information using standard protocols such as OSPF or simply use static or default routes. Within the City Net, customer traffic is isolated in its own ISID and VRF.

Figure 19. L3 Service



Let's define the elements in these diagrams.

- **CE:** Customer Edge. This device is located on customer premises but is owned and managed by the City Net. The CE device connects to and peers with the BEB. In a L2 Service, the CE device is a L2 device. In a L3 service, the CE device may be a L3 device particularly when each site has multiple subnets. For smaller sites, the CE device can be a L2 device and routing can be done by the BEBs.
- **BEB:** Backbone Edge Bridge. An SPB switch located at the edge of the City Net. The BEB connects CE devices to the City Net. A BEB may also connect to third party PE (Provider Edge) nodes.
- **BCB:** Backbone Core Bridge. An SPB switch located within the City Net network. A BCB connects to other BCB or BEB nodes but does not connect to CE devices or third-party PE nodes.

The Management and Internet services will be described in Sections 8 and 9.

7.4.1 Service Access

A given customer site may require a combination of L2, L3 and Internet services. In addition, CE devices need access to the Management service so they can be provisioned, configured and monitored from the NOC. When defining service access architecture, a few properties are desired:

- Support for multiple different services (L2 VPN, L3 VPN, Internet, Management) over the same CE-BEB interface and link
- Support for CE device auto-provisioning (Remote Configuration Download)
- Support for E-LAN or E-LINE L2 services
- Support for overlapping C-VLAN numbers on same BEB
- Single hairpin for all VPN services on a BEB
- No hairpins on BCB nodes

We have defined the CE-BEB interface and the BEB's hairpin interfaces to accomplish all the above.

7.4.1.1 L2 Service Access


Let's refer to Figure 20 below. A L2 CE device is installed on customer premises. The CE is owned and managed by the City Net. The CE's LAN or User ports constitute the UNI (User-Network Interface). The UNI is the point of demarcation between the customer's responsibility and City Net's responsibility. A customer may utilize one or more LAN ports on the CE device. For example, one port may be mapped to one service and another port to a different service. CE UNI ports will be tagged (802.1q) or untagged (native VLAN) with C-VLANs (Customer VLANs).

In a L2 Service, the CE device performs VLAN Stacking, also known as Q-in-Q. The CE's WAN port is the NNI (Network-Network Interface). An extra 802.1q tag known as the S-VLAN (Service VLAN) tag is added to customer traffic before forwarding out the NNI port. S-VLANs are assigned and managed by City Net and are BEB-significant only. This means different customers connected to the same BEBs will use different S-VLANs. But the same customer service may use different S-VLANs on different BEBs. VLAN translation is needed when near-end and far-end S-VLAN tags are different, which is normally the case because S-VLANs are only BEB-significant. Therefore, there is no need to coordinate near-end and far-end S-VLANs between BEBs.

This VLAN-stacked traffic is received at the BEB on a standard VLAN-tagged NNI and forwarded to another standard VLAN-tagged port which is physically wired to another port (hairpin). This second port in the hairpin is the SPB UNI or SAP (Service Access Point) port.

In point-to-point services (two BEBs), SAPs may simply match on the outer S-VLAN tag and map to a single ISID. C-VLANs will be transparently carried to the other end. VLAN translation can be enabled in case near-end and far-end S-VLANs don't match (usually the case).

Figure 20. L2 Service Access



Multipoint services (more than two BEBs) can use the same approach as point-to-point services. However, in multi-point services, there is a risk that duplicate MAC addresses may cause a “mac-move” issue. Normally, there should be no duplicate MAC addresses but, in reality, it can happen particularly in virtualized environments. Duplicate MAC addresses in different C-VLANs do not collide, however, if these C-VLANs are mapped to the same SPB service and the client devices are connected to different CEs, they will be constantly learned, re-learned and flushed on different BEBs. This is known as mac-move and should be avoided to maintain stability. To avoid mac-move, it is recommended that each C-VLAN is mapped to a different SPB service (ISID) when offering multipoint L2 Services. This will require one SAP and ISID per S-VLAN:C-VLAN combination.

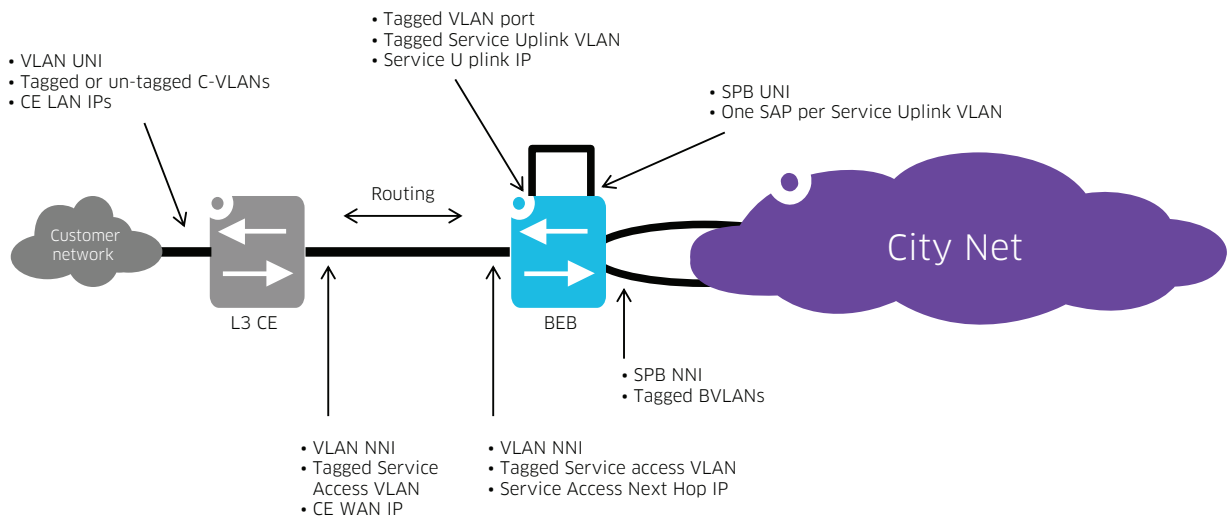
7.4.1.2 L3 Service Access

There are two possibilities for L3 Service Access: L3 CE and L2 CE. L3 CE is typically applicable to larger sites where there is a need to have multiple local routes and subnets, such as in a Hospital. L2 CE is typically applicable to smaller sites and where all traffic or most traffic will flow towards the City Net and/or Municipal Cloud and therefore it is easier, or more cost-effective, to deploy a L2 CE device and perform routing at the BEB. A Smart Light Pole or a Wi-Fi Hotspot is an example of such situation. Let’s review both cases.

7.4.1.2.1 L3 Service Access with L3 CE

Please refer to Figure 21 on the following page. There is no difference in the UNI when compared to an L2 Service: UNI ports accept tagged or un-tagged C-VLAN traffic. But there is a change at the NNI: customer traffic is no longer “stacked” with an S-VLAN but rather it is routed to a different Service Access VLAN by the CE device. The CE’s WAN IP address resides on this Service Access VLAN and the Service Access VLAN is tagged on the NNI interface. Note that the Service Access VLAN used for L3 services can co-exist with S-VLANs used for L2 Services on the same port. What this means is that this design supports mixing L2 and L3 services on the same interface. It should be noted that the Service Access VLAN is BEB-significant only and will be assigned by the City Net.

Figure 21. L3 Service Access with L3 CE



Customer traffic is received at the BEB on a standard tagged VLAN port. The same Service Access VLAN hosts the next-hop IP address at the BEB. This IP address is mapped to the customer's VRF. CE and BEB exchange routing information using dynamic routing protocols such as OSPF.

From the Service Access VLAN, traffic is routed to a different Service Uplink VLAN where a Service Uplink IP address resides. This IP address is also mapped to the customer's VRF.

Similarly, this Service Uplink VLAN traffic is forwarded to another standard VLAN port and, from there, is sent through a hairpin to a SAP port where it is mapped to an SPB service.

Using a dynamic routing protocol such as OSPF between CE and BEB, local OSPF routes are exported from the customer's VRF into the customer's ISID and far-end routes are imported from the customer's ISID into the customer's VRF. Imported routes are then re-distributed into OSPF. Tags are set and matched with route-maps to avoid circular route re-distribution.

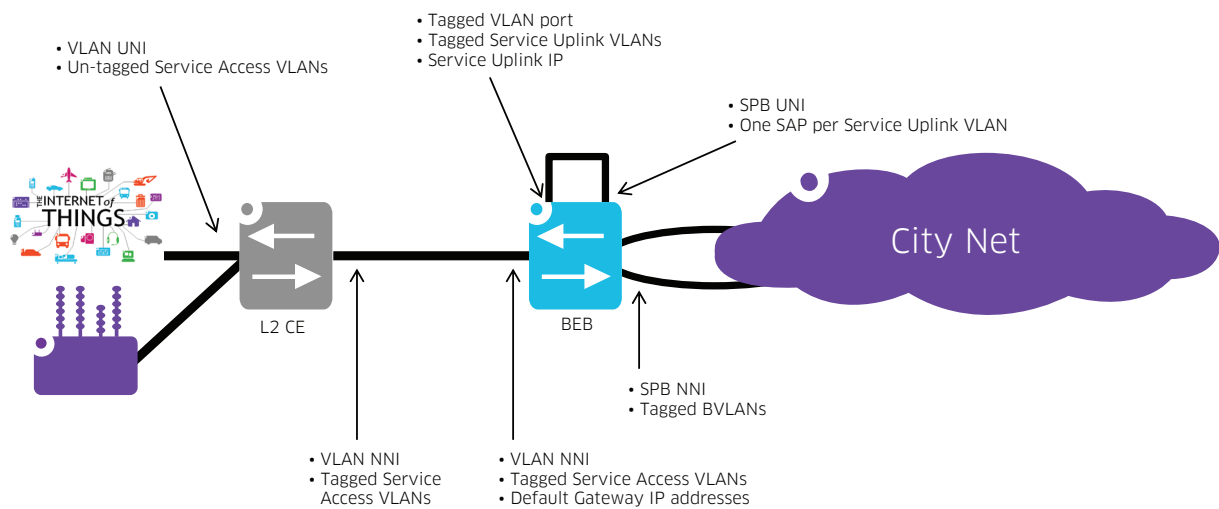
At the far-end BEB, this process reverses and traffic is delivered to the customer's far-end site.

All Service Uplink IP addresses on different BEBs belong in the same subnet. Traffic between near-end and far-end Service Uplink IP addresses is bridged across the SPB backbone.

7.4.1.2.2 L3 Service Access with L2 CE

Let's refer to Figure 22 below. Unlike the previous case, the CE device does not perform routing between local Customer VLANs. Local devices connect on standard, un-tagged VLAN ports on the L2 CE device. Note that "standard" does not mean "static". These ports can be mobile ports and dynamically assign Network Profiles based on authentication (802.1x, MAC, Captive Portal), rules (MAC OUI, range, IP, etc.) or LLDP (IP Phone or Wi-Fi AP).

Figure 22. L3 Service Access with L2 CE



Whether statically or dynamically, end devices are mapped to Service Access VLANs. The link between the L2 CE and the BEB is a VLAN NNI where these Service Access VLANs are tagged. At the BEB, there is a default gateway IP address for each of these Service Access VLANs. And traffic is routed from the Service Access VLAN to the Service Uplink VLAN. From the Service Uplink VLAN, traffic is mapped to an SPB SAP through a hairpin.

All Service Access default gateway and Service Uplink IP addresses belonging in the same L3 VPN are mapped to the same VRF. Unlike the previous case, there is no need to run a dynamic routing protocol. Local connected routes are exported from the VRF to the VPN's ISID and, conversely, imported from the VPN's ISID into the VRF.

At the far-end BEB, this process reverses and traffic is delivered to the far-end site.

All Service Uplink IP addresses on different BEBs belong in the same subnet. Traffic between near-end and far-end Service Uplink IP addresses is bridged across the SPB backbone.

7.4.1.3 Internet Service Access

In this section, we will explain access to the Internet Service. This refers to Enterprise-grade Internet services as opposed to residential or “business” Internet services.

Conceptually, Internet access is a special case of a L3 service. There are, however, some differences between a L3 VPN Service and an Internet Service:

- Unless Internet is offered in conjunction with other VPN services, a CE device may not be necessary.
- Even if using a CE device, the CE device does not require an Internet IP address, unless BGP is required (refer to Section 9). Customers are assigned an Internet IP address block (subnet, mask) and next-hop (gateway) IP address.
- A different Customer Internet VLAN exists at the BEB for each customer and this is where the customer's next-hop Internet IP address resides. These Customer Internet VLANs are BEB-significant only and all next-hop IP addresses belong in the same Internet VRF.
- Internet traffic is not mapped to SPB Services because all nodes participate in the Internet Service.
- The Internet Service traffic is routed across the backbone.

In summary, the Internet Service is not an SPB Service but it does run on a VRF of its own.

Please refer to Section 9 for further details on the Internet Service.

Figure 23 illustrates access to the Internet Service with a CE device. Internet access with CE device would normally be applicable to customer sites where other L2 or L3 VPN services are delivered and eliminates the need for a separate connection at the site.

Figure 23. Internet Access with CE Device

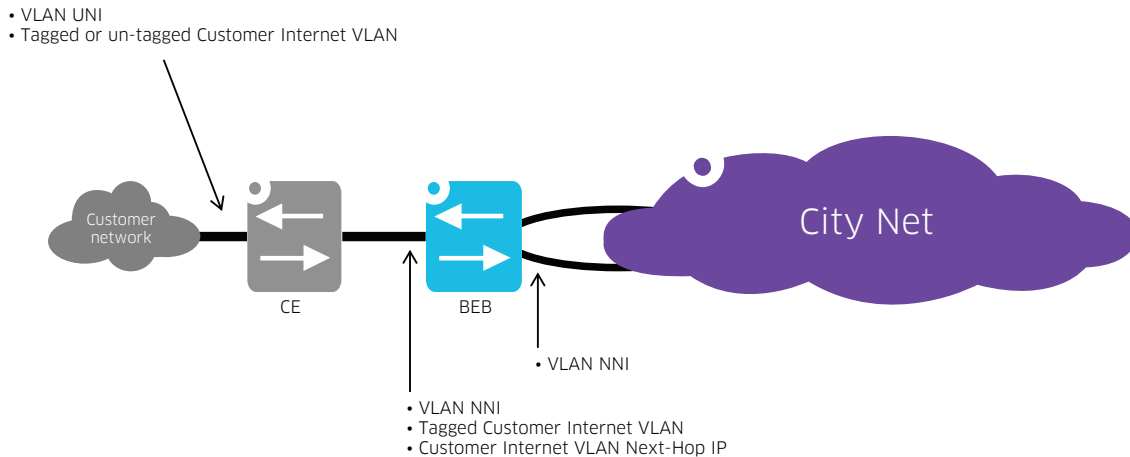



Figure 24 illustrates access to the Internet Service without a CE device. Internet access without CE device would normally be applicable to customer sites where no other L2 or L3 VPN services are needed.

Figure 24. Internet Access without CE Device



7.4.1.4 Management Service Access and CE Auto-Provisioning

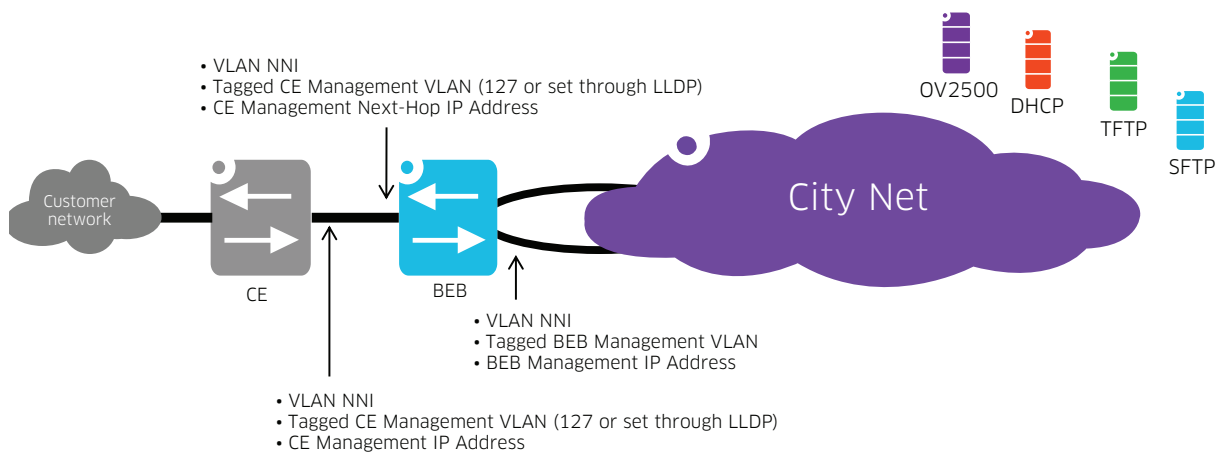
Conceptually, the Management Service is a special case of a L3 service. There are, however, some differences between a L3 VPN Service and the Management Service:

- The purpose of the Management service is to manage and monitor the CE device. Therefore, there is no UNI for this service.
- A dedicated CE Management VLAN and IP address are defined at the CE. This CE Management VLAN is tagged on the interface between CE and BEB.
- The CE Management VLAN is BEB-significant and is the same for all CE devices connected to the same BEB. The next-hop CE Management IP address on the BEB resides on this VLAN and belongs in the Management VRF.
- The Management Service is not mapped to an SPB Service.
- CE nodes do not route management traffic.

This definition of the CE Management VLAN and Management Service facilitates use of the Remote Configuration Download (RCD) feature to auto-provision CE devices. When a blank, out-of-the-box OmniSwitch boots up, it goes through the RCD sequence. It requests an IP address through DHCP on the un-tagged (native) VLAN and tagged VLAN 127 or, on a specific VLAN that the BEB can inform through LLDP (Link Layer Discovery Protocol). The OmniSwitch will receive an IP address lease and, with it, certain attributes specifying the IP address of a TFTP server and the name of an instruction file. With that, the OmniSwitch will download the instruction file from the TFTP server. The instruction file contains details of firmware and configuration file to be downloaded from an FTP or SFTP server. The instruction file can also contain details of a script file. The script file is used to add commands that are not saved in the configuration file (e.g. changing the default password). After the firmware and configuration file are downloaded, they are saved to the “working” directory. Finally, the OmniSwitch will reload from its working directory and can be managed from the NOC. Please refer to the “OmniSwitch AOS 8 Switch Management Guide” for further details.

Therefore, VLAN 127 is the recommended and default CE Management VLAN.

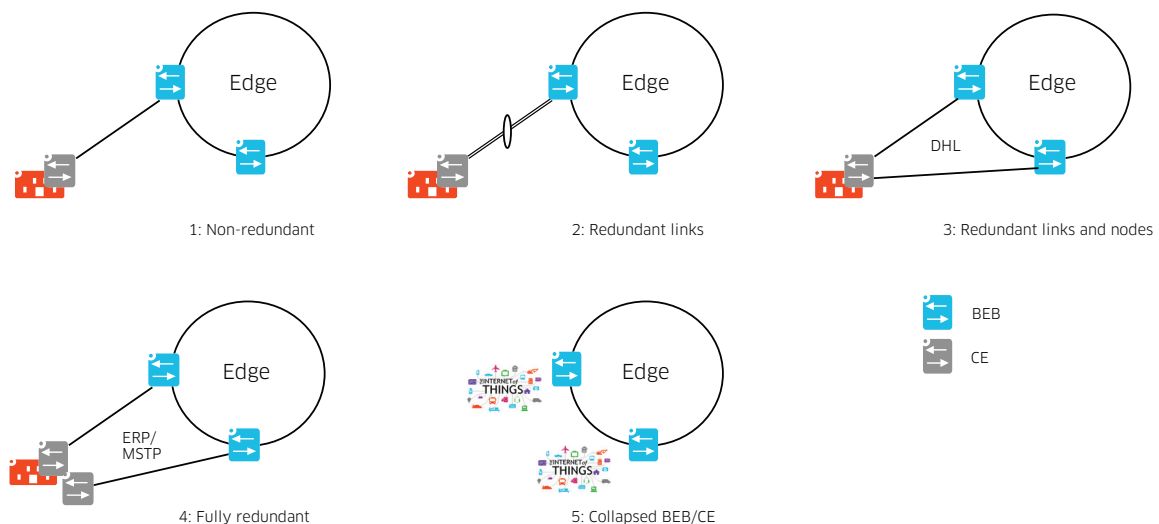
Figure 25. Management Service Access and CE Auto-Configuration



7.4.2 Service Access Redundancy

Service access redundancy increases service availability by protecting from link or device failures. Figure 26 below illustrates different options for access to L2 Services.

Figure 26. L2 Service Access Redundancy



- **Non-redundant:** the site is attached to a single BEB through a single link. Link, BEB or CE failure will result in loss of service to the site.
- **Redundant links:** the site is attached to a single BEB through a link aggregate (LAG). This adds protection from single-link failure. Note that fibre runs should use diverse physical paths to protect against fibre cuts which would typically interrupt both links otherwise.
- **Redundant links and nodes:** the site is attached to two different BEBs through two different links. This adds protection from BEB failure. When possible, both links should use physically diverse paths such that link failure events are not correlated. Dual-Home Link (DHL) is a high availability feature that provides fast failover without implementing Spanning Tree or Link Aggregation. Please refer to the “AOS 8 Network Configuration Guide” for further details. Note that this redundant access option requires using the same S-VLAN at both BEB nodes and this S-VLAN must be tagged on the link between both BEBs.
- **Fully redundant:** this option adds CE device redundancy. MSTP (Multiple Spanning Tree Protocol) or ERP (Ethernet Ring Protection) can be used to avoid loops in this redundant connection. As in the previous case, this option requires using the same S-VLAN at both BEB nodes and this S-VLAN must be tagged on the link between both BEBs. L2 control protocol frames must not be tunneled at the SPB SAP to avoid joining multiple sites in the same STP or ERP domain.
- **Collapsed BEB/CE:** this is another example of a non-redundant connection. In this case, there is no CE device and this role is implemented in the BEB itself. End devices can connect on a VLAN or SAP UNI port. When connecting on a SAP port, the null tag is matched and traffic is mapped to the SPB service. When connecting on a VLAN port, traffic is forwarded to the VLAN-port side of a hairpin loop and from there forwarded onto a SAP port where the VLAN tag is matched and traffic is mapped to the SPB service.

Note that Virtual Chassis (VC) can be combined with the options above to increase resiliency.

Figure 27. L3 Service Access Redundancy with L3 CE

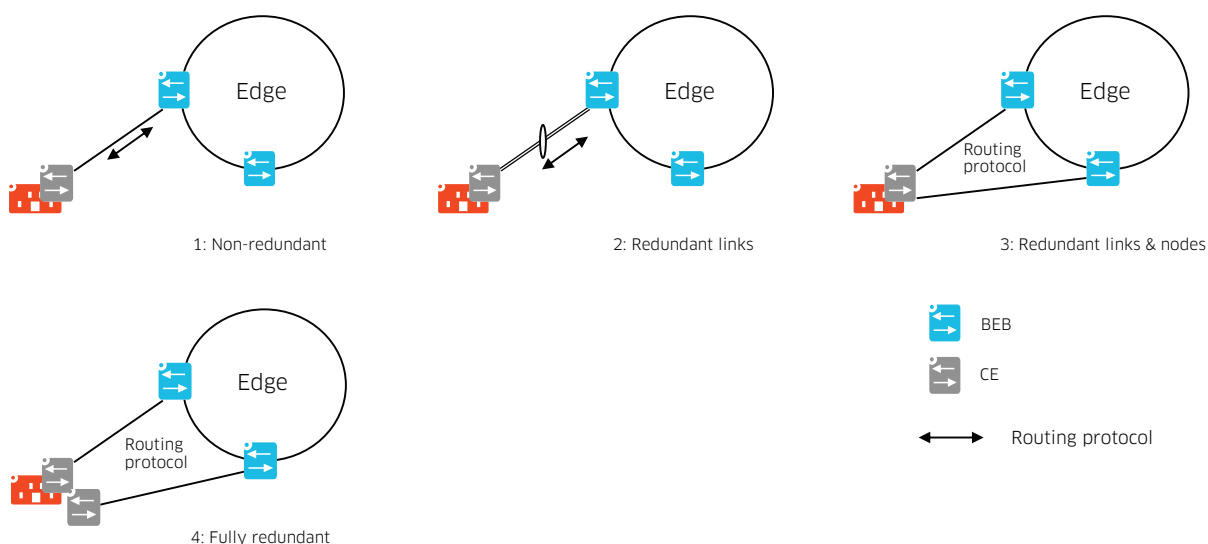


Figure 27 illustrates different options for access to L3 Services using a L3 CE.

- **Non-redundant:** the site is attached to a single BEB through a single link. Link, BEB or CE failure will result in loss of service to the site.
- **Redundant links:** the site is attached to a single BEB through a link aggregate (LAG). This adds protection from single-link failure. Note that fibre runs should use diverse physical paths to protect against fibre cuts which would typically interrupt both links otherwise.

- **Redundant links and nodes:** the site is attached to two different BEBs through two different links. This adds protection from BEB failure. When possible, both links should use physically diverse paths such that link failure events are not correlated. A dynamic routing protocol such as OSPF is used between BEBs and CEs to exchange routing information. Import/Export and re-distribution of routes must be carefully planned to avoid circular re-distribution of routes. This is accomplished with route maps and tags.
- **Fully redundant:** this option adds CE device redundancy.

Figure 28. L3 Service Access Redundancy with L2 or no CE

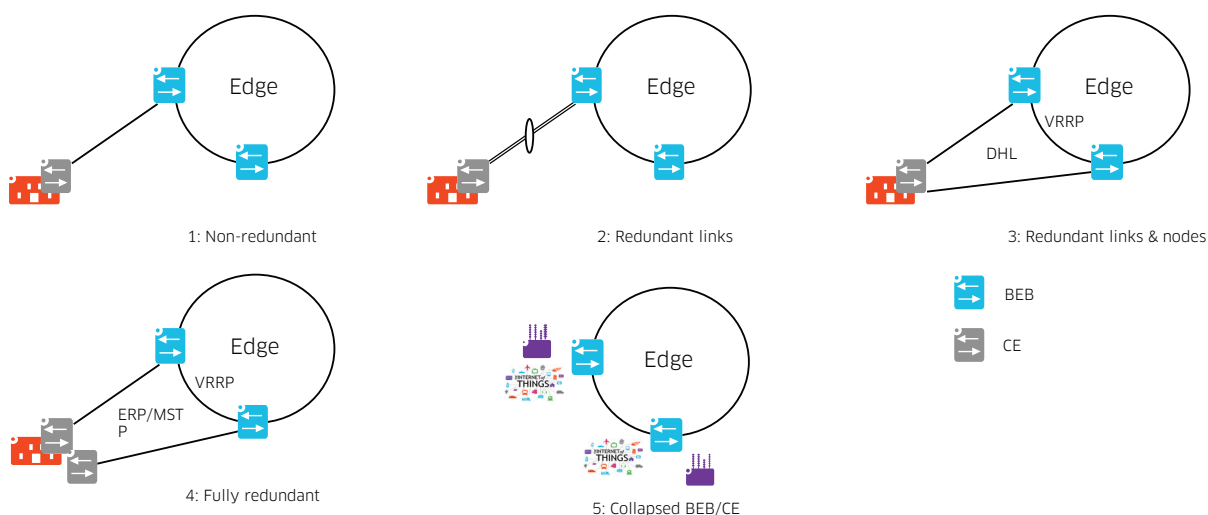


Figure 28 illustrates different options for access to L3 Services using a L2 or no CE. You may notice that these options are almost identical to the ones presented in Figure 26 with the exception that VRRP is added when sites redundantly connect to two BEBs.

8 Management Service Architecture

Although it can be used to manage City Net nodes, the purpose of the Management Service is managing CE devices. A separate out-of-band (OOB) management network connected to console or EMP ports is recommended for City Net node management. This OOB management network will not be described in this document.

As explained previously, the Management Service is not an SPB service. The Management Service is routed over dedicated VLANs and VRF. On SPB backbone interfaces, these Management VLANs run alongside BVLANS. This design eliminates the need for hairpins on BCB nodes.

Let's refer to Figure 29 on the following page. The Management Service's architecture is modular and hierarchical and can use either OSPF or IS-IS. All routing interfaces reside on different VLANs and the type is point-to-point. This is illustrated on Figure 30 for a small Edge Block with 3 BEBs and on Figure 31 for Backbone and Core BCBs.

Figure 29. Management Service Architecture

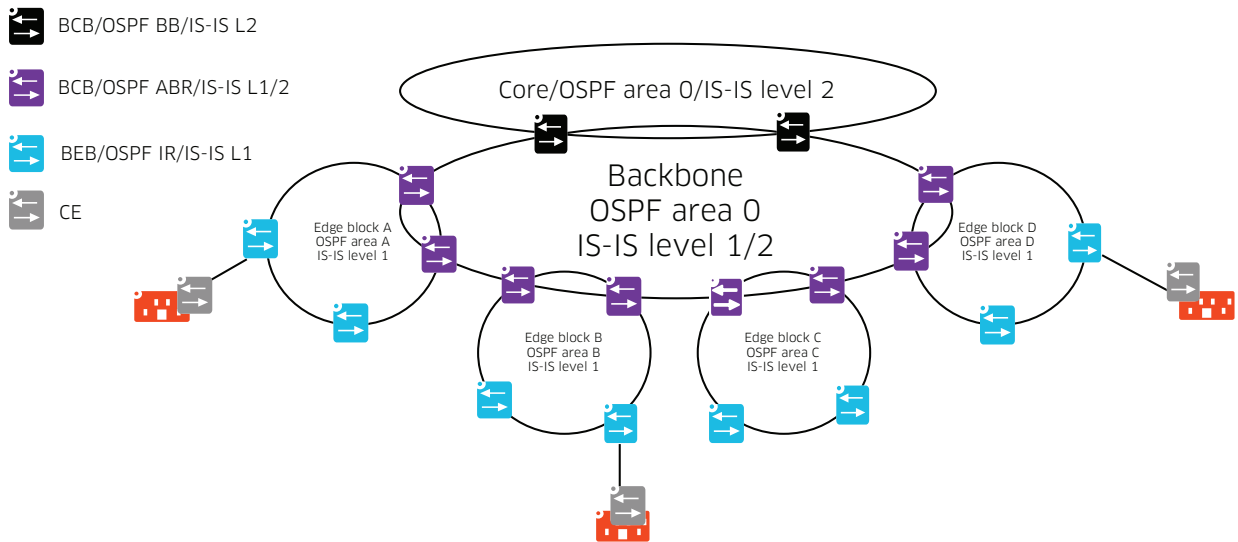


Figure 30. BEB Management Interfaces

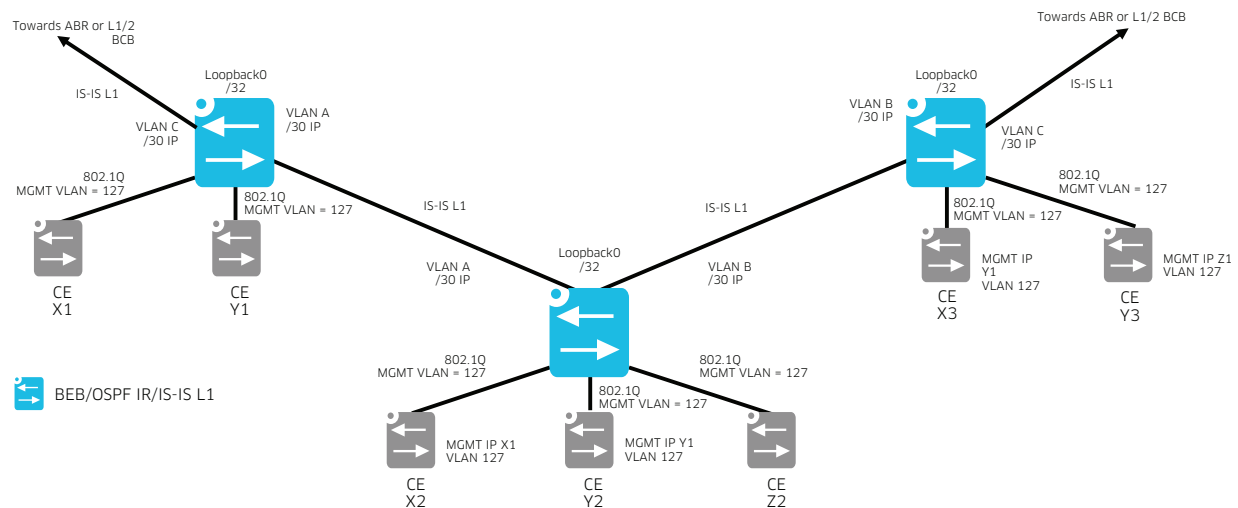
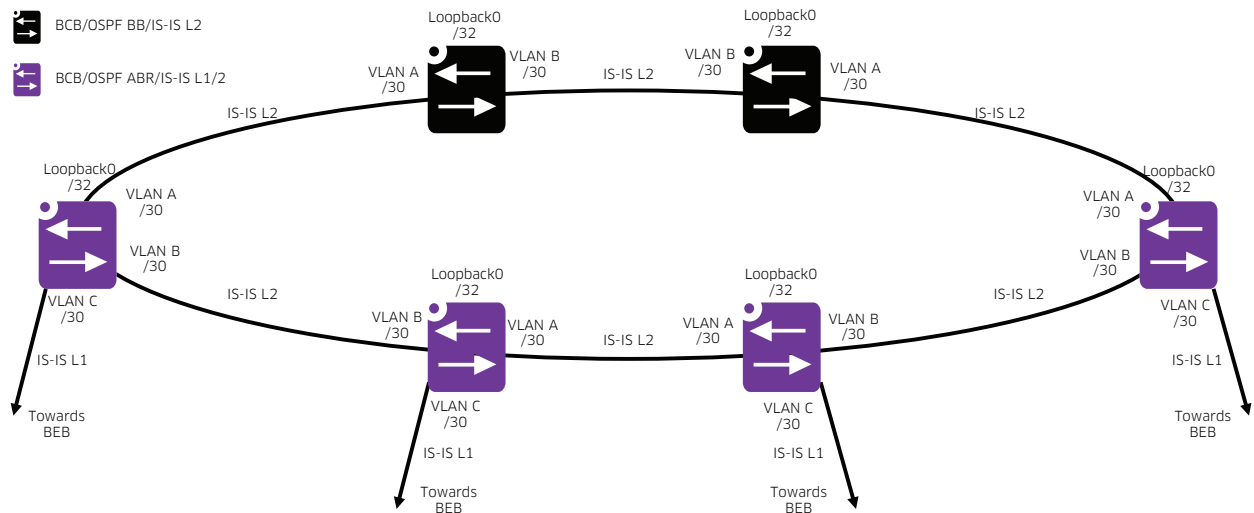


Figure 31. BCB Management Interfaces



These diagrams show both IS-IS and OSPF terminology and roles but only one routing protocol is needed. The choice between these two protocols does not make a big difference for the management service. It does make a bigger difference for the Internet Service as we will see in the next section. The designer may prefer using the same protocol for both services.

Core BCB nodes reside within OSPF area 0 or use Level-2 IS-IS adjacencies only. BCB nodes are OSPF Backbone nodes or IS-IS Level-2 nodes.

Backbone BCB nodes are OSPF ABRs (Area Border Routers) because they are attached to area 0 as well as other areas. If using IS-IS, Backbone BCB nodes are Level-1/2 nodes and have both L1 and L2 adjacencies: L1 adjacencies with BEBs and L2 adjacencies with other BCBs.

BEB nodes are OSPF IR (Interior Routers) because all their interfaces belong in the same area. If using IS-IS, they are Level-1 routers (all their adjacencies are Level-1).

Addressing suggestions:

- All BEB and BCB nodes use Loopback0 for management
- CE nodes do not use Loopback0 for management
- CE nodes have a Primary MGMT IP address on a Primary CE MGMT VLAN. We will use VLAN 127 as an example as it is the default.
- CE nodes redundantly connected to diverse BEBs also have a Secondary MGMT IP address on a different VLAN. We will use VLAN 128 as an example.
- All CE Primary MGMT IP Addresses of interfaces connected to the same BEB belong in the same subnet
- All CE Secondary MGMT IP Addresses of interfaces connected to the same BEB belong in the same subnet
- Therefore, each BEB has two VLANs for CE MGMT: 127 and 128
- And each BEB also has two different subnets for Primary and Secondary CE MGMT IP Addresses, mapped VLANs 127 and 128, respectively.
- This eliminates the need for L2 protection mechanisms such as DHL/MSTP/ERP between CE and BEBs.
- It also eliminates the need for static or dynamic routing towards the CEs.
- CE nodes will have static routes to the NOC pointing to each BEB that they are connected to.
- On L3 CE nodes, MGMT IP Addresses use a separate MGMT VRF.
- CE and BEB MGMT subnets in each Edge Block should be allocated in such a way that they can be summarized at the ABR or L1/2 BCBs before injection into the Backbone to minimize routing table entries.
- CE Management IP addresses are assigned through DHCP.

Additional guidelines:

- For stability, always configure Loopback0 IP address or Router ID on all nodes running OSPF or IS-IS.
- BEB nodes perform DHCP relay for MGMT VLANs such that CE nodes can receive DHCP-assigned MGMT IP addresses.
- If using OSPF, Edge Block areas should be defined as totally stubby areas because they only need a default route. If using IS-IS, BEBs should be configured as L1 only to accomplish the same outcome.
- Routing (OSPF or IS-IS) should be enabled on all BEB management interfaces.
- CE Management IP interfaces (e.g. in VLANs 127 and 128) should be configured as passive interfaces because no adjacencies should form on these interfaces.
- For stability, do not redistribute static or connected routes.
- Configure link costs as inversely proportional to bandwidth.
- Use authentication between routing nodes.

9 Internet Service Architecture

In this section, we will describe the Internet Service Architecture as it relates to Enterprise customers in greater detail. We will not describe the Internet Service Architecture for Residential or “Business” customers.

Let’s start with some definitions:

- City Net is an ISP (Internet Service Provider).
- City Net has its own ASN (Autonomous System Number).
- The City Net’s ASN is obtained from the relevant RIR (Regional Internet Registry) such as APNIC in Asia Pacific or ARIN in North America.
- City Net has its own IPv4 and IPv6 address space.
- City Net may apply for IP address space with the relevant RIR or lease it from the Upstream ISP (particularly in the case of IPv4 as this address space is depleted).
- City Net may implement dual IPv4 and IPv6 stacks for Enterprise customers.
- City Net customers (whether public or private organizations) may also have their own ASN and IPv4/IPv6 address space.
- City Net may allow address space portability to facilitate migration from incumbent to City Net services.
- City Net will connect to one or two Upstream ISPs and receive a default route, but not the full Internet table.
- City Net may connect to an IXP or exchange domestic routes privately with other local ISPs. The term “domestic routes” refers to direct customers of those local ISPs, not default or full routes.
- City Net will not be a transit AS between other ISPs.
- City Net will advertise prefixes owned by City Net and its customers. City Net will neither receive nor advertise the full Internet routing table.

The Internet Service requires two routing protocols:

- An Interior Gateway Protocol (IGP) to exchange routing information between routers across the City Net’s Autonomous System (AS).
- BGP (Border Gateway Protocol) to exchange routing information with routers in other AS, such as City Net’s customers and other ISPs.

IGP options come down to OSPF and IS-IS. OSPF and IS-IS are largely equivalent: they both implement Dijkstra’s Shortest Path First algorithm, have similar convergence times, etc. There are however a few strong points in IS-IS’s favour:

- IS-IS natively supports multiple address families. This means both IPv4 and IPv6 can be supported with a single protocol. If using OSPF, two protocols are required: OSPFv2 for IPv4 and OSPFv3 for IPv6.
- IS-IS does not run on IP. Therefore, IS-IS cannot be attacked using IP packets.
- IS-IS is the same protocol used by SPB. Although SPB’s IS-IS runs on a separate instance, consolidating on a single protocol is simpler than using multiple protocols.

Figure 32 below provides an IGP view of the Internet Service architecture. This architecture is very similar to the Management Service architecture. For this reason, we will not describe the IGP architecture again in this section. You can refer to Section 8 for details. Note in Figure 32 that the chosen IGP is used to exchange routing information amongst City Net nodes only. Interfaces connected to external parties such as customers or ISPs are passive interfaces for the IGP.

Figure 32. Internet Service - IGP View

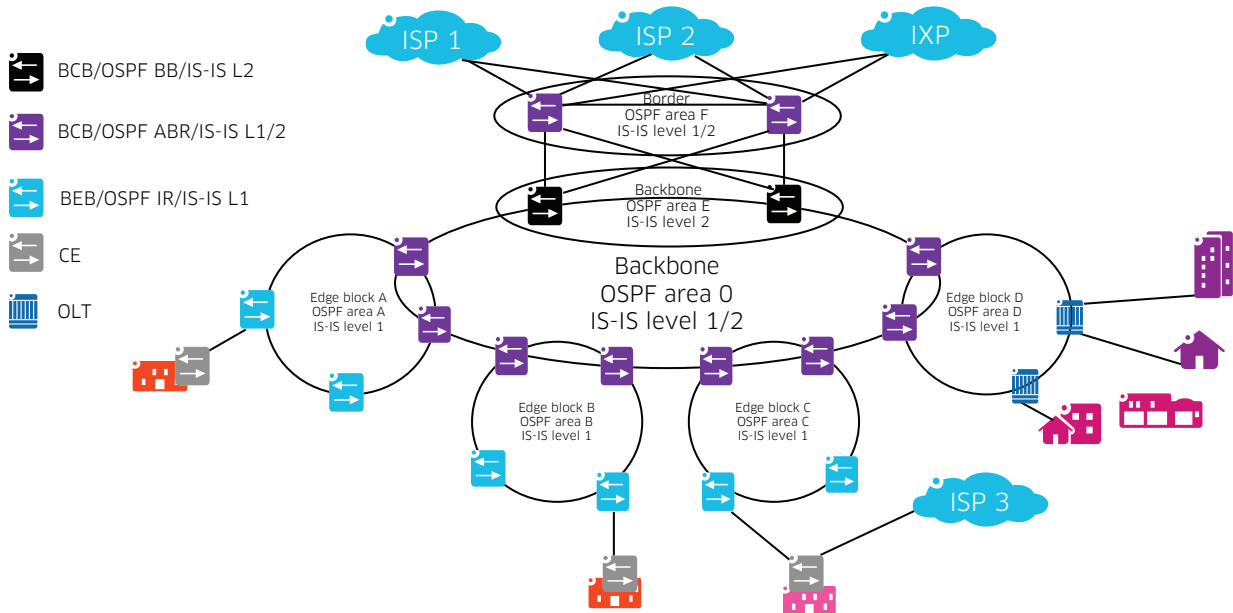


Figure 33 provides a BGP view of the Internet Service architecture. In this figure, iBGP is used between City Net nodes whilst eBGP is used between City Net and other autonomous systems. It should be noted that not all City Net customers need to use BGP. Default routes are sufficient for many or most customers. Customers with more sophisticated needs may require BGP. One such example is shown at the bottom of this diagram. This customer is connected to two different ISPs for redundancy. The customer has its own IP address space which is advertised through BGP.

The iBGP architecture uses hierarchical route reflector clusters for scalability. Cluster 0 is comprised of Backbone and Core nodes as well as the Border nodes connecting to upstream ISPs and IXP. Both Core nodes are Route Reflectors (RR). Border nodes are Route Clients (RC). Other Backbone nodes are RCs of the Core RR and, in turn, RR of their own cluster. Finally, BEB nodes are simple RCs.

Residential or Commercial customers shown in Figure 33 normally use FTTH connections. These customers do not use BGP or any routing protocol and their IP addresses are assigned by City Net.

Addressing guidelines:

- City Net needs to allocate public IP addresses to:
 - Customers
 - Point-to-point interfaces
 - Loopback interfaces
- City Net can partition its allocated public IP address space in three contiguous blocks for each purpose.
- Addresses for point to point interfaces should be allocated in such a way that they can be summarized at the ABR or L1/2 BCBs before injection into the Backbone to reduce routing table size.

- Loopback0 addresses should be allocated from a contiguous block and not be summarized across areas or levels.
- Customers are assigned contiguous blocks of different size according to their needs.

BGP Guidelines:

- Customer IP address space is only advertised through BGP, not through the IGP.
- For customers that do not run BGP, their address space is injected into the BGP routing table with explicit network statements at the BEB matching the required prefixes.
- Before customer prefixes can be advertised through BGP with a network statement, they must exist in the BEB's routing table. At the BEB, this can be accomplished with static routes pointing to the customer's CE router. In some cases, for instance, when the customer is using City Net-assigned public address space which is not subnetted, the BEB will already have an IP interface with an address and mask that will match the required prefix.
- Redistributing IGP routes into BGP is highly discouraged because flapping IGP routes will lead to instability. In fact, running an IGP between the City Net and the customer's CE is not recommended and static routes, if required, should be used instead.
- To minimize routing table entries, "ip bgp neighbour next-hop-self" is recommended at BEB nodes connected to customer eBGP nodes. This eliminates the need to add an IGP network statement or configuring the BEB-CE interface as an IGP passive interface which will create an additional routing table entry.
- For stability, always use Loopback0 for BGP sessions. This requires enabling multi-hop for eBGP neighbours.
- For scalability and stability, customer prefixes should be aggregated at the BEB when possible.

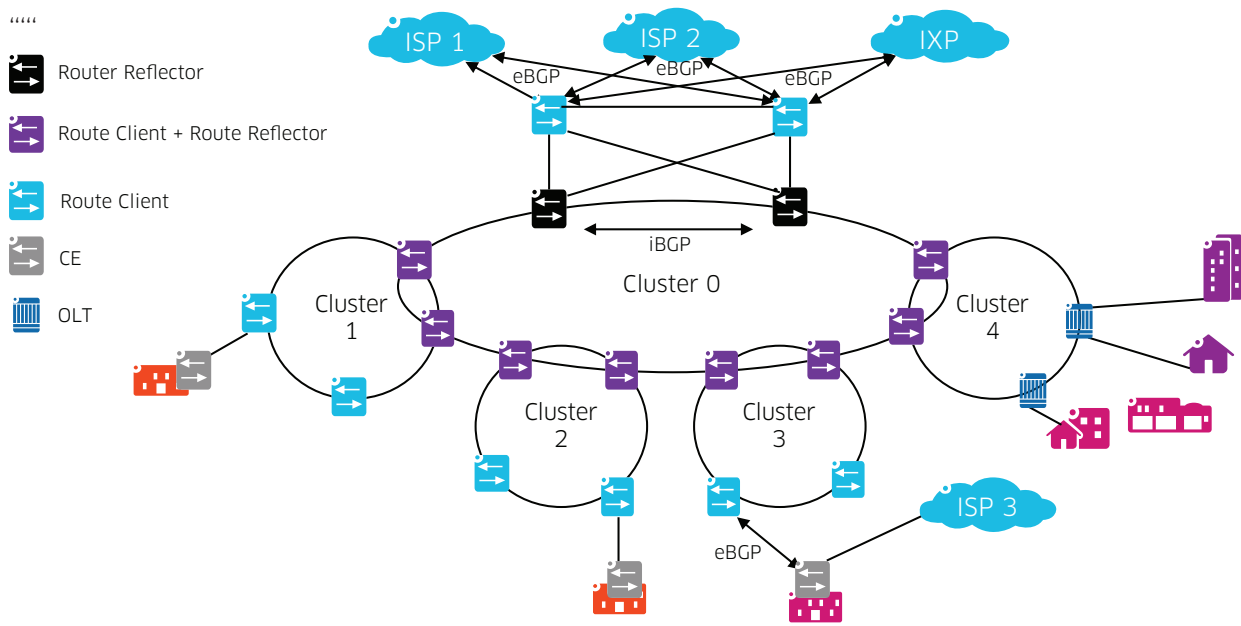
Peering Guidelines:

- eBGP connections between BEBs and City Net customer CE routers will only accept prefixes belonging to those customers. This can be accomplished with prefix lists.
- eBGP connections between Border BEBs and other local ISPs at the local IXP will only accept domestic prefixes (i.e. prefixes belonging to direct customers of those ISPs). This can be accomplished with as-path list policies.
- eBGP connections between Border BEBs and upstream regional or national ISPs will only accept default routes and their direct customer prefixes. This can be accomplished with prefix lists and as-path policies.
- eBGP connections between BEBs and City Net customer CE routers will only advertise prefixes belonging to City Net customers and a default route. This can be accomplished with prefix lists and as-path list policies.
- eBGP connections between Border BEBs and other local ISPs at the local IXP will only advertise City Net domestic prefixes (i.e. City Net and its direct customer's). This can be accomplished with as-path list policies.
- eBGP connections between Border BEBs and upstream regional or national ISPs will only advertise City Net domestic prefixes (i.e. City Net and its direct customer's). This can be accomplished with as-path list policies.
- When connecting to multiple upstream ISPs, route preferences can be implemented with BGP attributes such as local-preference, MED and as-path pre-pending based on policy.

Other guidelines:

- To avoid IP spoofing, BEBs connected to City Net customer CE routers should only accept IP packets with source IP addresses matching the customer's public IP address space. Traffic with different IP source addresses, including private IP address space, should be dropped. This can be accomplished with ACLs.
- A number of other invalid IP packets can be filtered with the DoS Filtering feature in AOS. Please refer to Section 10.3.3 for details.

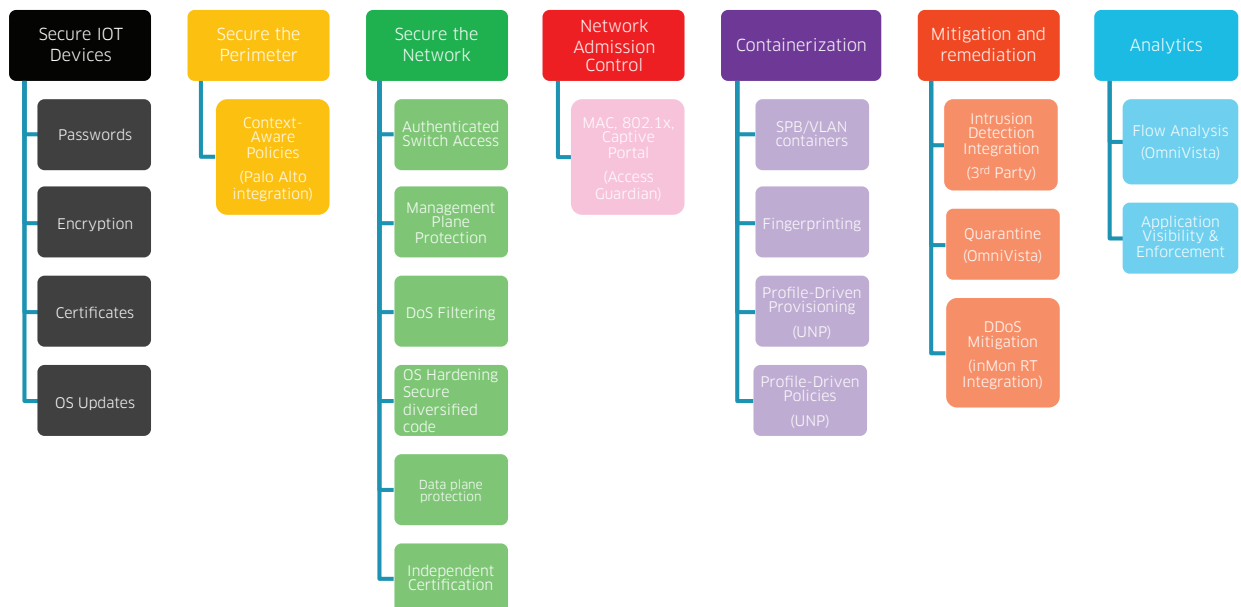
Figure 33. Internet Service - BGP View



10 Security Framework

Securing mission-critical city infrastructure requires a layered “defense-in-depth” approach that combines both proactive and reactive defense mechanisms. Please refer to Figure 24 for snapshot of the various applicable mechanisms.

Figure 34. Layered Security Framework



Let's examine these mechanisms in more detail.

10.1 Securing IOT Devices

IoT devices are vulnerable to security threats just like other IT assets. A compromised IoT device can also become an attack vector into other devices and systems. Because of the high number of IoT devices, the impact of such attacks can be very high.

These devices must be securely configured and managed. The exact measures will depend on the device capabilities and manufacturer recommendations, but some of these are listed below.

- Passwords: Password complexity, renewal and authentication against a central database.
- Certificates: X.509 certificates can be installed on IOT devices allowing for mutual authentication between the IOT device and the Server. Certificates can also be used for Network Admission Control (NAC).
- Encryption: Secure protocols such as Transport Layer Security (TLS) must be used when managing these devices and any unsecure protocol must be disabled.
- OS Updates: IOT devices must be updated and patched according to manufacturer specifications to prevent exploitation of known vulnerabilities.

10.2 Securing the Perimeter

When communication between different containers is required, it must only be allowed through Firewalls and controlled by fine and specific policies. Different systems may have physical or virtual Firewalls of their own if operated by different organizations. Please refer to [4] for an explanation of how integration between Alcatel-Lucent Enterprise's Unified Access solution and Palo Alto Networks firewall can be used to create dynamic fine-grained policies that take the user's identity, device, application, location and other situational factors into consideration.

10.3 Securing the Network

Several steps must be taken to secure the network infrastructure itself. An overview of network security mechanisms is given below.

10.3.1 Authenticated Switch Access & Logging

Switch security features increase the security of the basic switch login process by allowing management only through particular interfaces for users with particular privileges. Login information and privileges should be stored on an external server such as Radius or LDAP. External servers should also be used for accounting, which includes logging statistics about user sessions. Admin authentication against the local database may be allowed on the console port only as a failover mechanism in case the external servers become unavailable or in case of misconfiguration.

10.3.2 Management Plane Protection

Management protocols must be secured, as detailed below:

- Insecure protocols such as Telnet, FTP/TFTP and SNMP must be disabled
- SSHv2 should be used with keys larger than 2048 bits
- HTTPS can be used if required for RESTful API/WebServices access and if so, it should be setup with X.509 certificates for mutual authentication between network node and Server.
- SNMPv3 with authentication and privacy options should be used for monitoring
- TLS 1.1. or 1.2 (preferred) should be used when connecting to remote RADIUS, LDAP or Syslog servers.

10.3.3 Denial of Service (DoS) Filtering

By default, an OmniSwitch filters denial of service (DoS) attacks. Some DoS attacks aim at system bugs or vulnerability, while other types of attacks involve generating large volumes of traffic so that network service is denied to legitimate network users.

10.3.4 OS Hardening – Secure diversified code

Secure diversified code is a technology which mitigates risks at the source, enabling an enhanced security profile through:

- Independent verification and validation of OmniSwitch source code
- Software diversification of OmniSwitch object code to prevent exploitation
- Secure delivery of OmniSwitch software

The OmniSwitch AOS secure diversified code protects networks from intrinsic vulnerabilities, code exploits, embedded malware, and potential back doors that could compromise mission-critical operations. The secure diversified code technology is continuously applied on every new AOS release, therefore it will address both current and future threats.

10.3.5 Data Plane Protection - MACSec

Data integrity and confidentiality must be protected whilst in transit through the network. MACSec is an IEEE standard (802.1AE) which provides point-to-point authentication and optional encryption between MACSec-capable devices such as switches. MACSec can prevent various threats such as man-in-the-middle, sniffing, spoofing and playback attacks.

Because MACSec operates at the MAC layer, it transparently secures all upper layer traffic transiting through MACSec-enabled links. This includes both application-layer data as well as control-plane and management-plane communication. In addition, unlike IPSec, MACSec is implemented in hardware at wirespeed and does not introduce additional latency or bandwidth limitations.

In a Smart City, the main application for MACSec is protecting data integrity and confidentiality whilst transiting over public spaces outside of the physical security perimeter, where it can be subject to tapping and other malicious activity.

10.3.6 Independent Certification

OmniSwitch products have been independently certified to comply with rigorous security standards. Compliance with these standards is mandated in public sector, but is valuable beyond public sector. Independent certification is an objective benchmark to compare security features in different products.

OmniSwitch products are certified to comply with the following standards.

Figure 35. Security Certifications



10.3.6.1 Common Criteria

The Common Criteria for Information Technology Security Evaluation, normally referred to as Common Criteria for short, is a set of specifications which serves as a framework in the evaluation of security products.

There are two important aspects in the Common Criteria certification:

- The Protection Profile: This defines the specific security requirements which are relevant to a particular device class such as a Network Device or a Firewall.
- The Evaluation Assurance Level (EAL): This is a number from 1 to 7 representing how thoroughly the product has been tested. Higher EAL levels do not necessarily mean more secure products but rather more thorough testing.

OmniSwitch products 6250, 6350 and 6450 with AOS 6.7.1.79R04 and 6860, 6865, 6900, 9900 and 10K with AOS 8.3.1.348.R01 are EAL-2 (Structurally Tested) against the Network Device Collaborative Protection Profile (NDcPP) set of requirements. Please refer to [8] and [9] for the EAL-2 and NDcPP certificates, respectively.

10.3.6.2 FIPS 140-2

The Federal Information Processing Standard 140-2 (FIPS 140-2) is a US and Canadian security standard focusing on hardware and software solutions using cryptography.

FIPS 140-2 certificates for AOS 6.7.1R04 (used in OmniSwitch 6350 and 6450) and AOS 8.3.1R01 (used in OmniSwitch 6860, 6865, 6900 and 9900) can be found in [10] and [11] respectively.

10.3.6.3 JITC

The Joint Interoperability Test Command (JITC) is the US Department of Defense's Joint Interoperability Certifier and only non-Service Operational Test Agency for Information Technology (IT)/National Security Systems. JITC provides risk based Test, Evaluation & Certification services, tools, and environments to ensure Joint Warfighting IT capabilities are interoperable and support mission needs. OmniSwitch 6860, 6865, 6900 and 9900 are certified, please refer to [12].

10.3.7 Network Admission Control - Access Guardian

Physical devices attached to a LAN port on the switch can be authenticated using port-based network access control. The following options for authentication are available:

- 802.1X authentication for supplicants.

Uses Extensible Authentication Protocol (EAP) between an end device and a network device (NAS) to authenticate the supplicant through a RADIUS server.

- MAC-based authentication for non-supplicants.

MAC-based authentication does not require any agent or special protocol on the non-suppliant device; the source MAC address of the device is verified through a RADIUS server. The switch sends RADIUS frames to the server with the source MAC address embedded in the username and password attributes.

- Internal or External Captive Portal

Captive Portal authentication enables Web-based authentication for Guest and BYOD users.

The authentication server may return a User Network Profile (UNP). UNPs map devices to VLANs or Services and may contain security and QoS policies. If no UNP is returned, a default UNP is applied.

10.3.8 Containerization

Containers are virtual network segments where IoT devices and applications that control them are isolated from other devices and applications. This segmentation facilitates enforcement of security policies and limits the damage in the event of a security breach.

SPB intrinsically segments IOT devices into containers by way of MAC-in-MAC tunneling and the ISID field designates the container. Any communication outside of the container will be controlled by a Firewall.

10.3.9 Threat Mitigation and Remediation

So far we have focused on best practices and preventive measures that can proactively stop a security incident from happening in the first place. But due to the fast evolving nature of security threats, no security strategy is complete without reactive mechanisms to thwart or dampen the effect of those threats that cannot be proactively avoided.

OmniVista 2500 NMS can be integrated with external Intrusion Detection Systems (IDS) through Syslog. When intrusion or malware attacks are detected by the IDS, a Syslog message is sent to OmniVista. The Syslog message includes the intruder's or attacker's address. OmniVista can use this information to locate the switch and port or AP that the device is connected to and quarantine it by shutting down the port or applying a quarantine profile (restrictive VLAN and ACLs) such that the malicious activity is stopped and remediation activities (e.g. OS patching and cleanup) can be performed.

An additional DDoS mitigation method is described in [5] InMON sFlow-RT collects sFlow data from switches and detects DDoS attacks in real time. InMON notifies the DDoS Mitigation SDN Application which in turn instructs the SDN Controller to push the necessary rules to drop traffic associated to the attack.

10.3.10 Analytics

Smart Analytics in OmniVista 2500 brings unprecedented visibility into the network status and usage patterns up to the application level. Understanding patterns assists in fine tuning and enforcing security policies to drive compliance.

OmniVista 2500 collects data from all OmniSwitch products through SNMP and SFlow. Various reports transform this data into valuable information.

Application Visibility extends this view to the Application Layer through Deep Packet Inspection. This means applications can be detected even if running on standard Web ports such as TCP ports 80 and 443.

When using the OmniSwitch 6860E as a Site Access Node, Application Visibility and Policy Enforcement can be enabled on standard (VLAN) access ports.

11 Conclusion

Alcatel-Lucent Enterprise's offers a comprehensive set of solutions that allow Smart Cities to optimize investment while bringing new services that enhance the city's liveability, sustainability and competitiveness. Alcatel-Lucent Enterprise's multi-layered security protects the city's mission-critical infrastructure from cyber-threats.

12 Acronyms

ABR	Area Border Router
ADSL	Asymmetric Digital Subscriber Line
API	Application Programming Interface
AS	Autonomous System
ASN	Autonomous System Number
BB	Backbone Router
BCB	Backbone Core Bridge
BEB	Backbone Edge Bridge
BLE	Bluetooth Low Energy
CDN	Content Delivery Network
C-VLAN	Customer VLAN
DC	Data Centre
DHCP	Dynamic Host Configuration Protocol
DR	Disaster Recovery
DHL	Dual-Homed Link
EMP	Ethernet Management Port
FTP	File Transfer Protocol
FTTH	Fibre to the Home
GRE	Generic Routing Encapsulation
HA	High Availability
IoT	Internet of Things
IR	Internal Router
ISID	Backbone Service Identifier
IS-IS	Intermediate System to Intermediate System
ISP	Internet Service Provider
IXP	Internet Exchange Point
LLDP	Link Layer Discovery Protocol

LPWAN	Low-Power Wireless Area Network
LTE	Long Term Evolution
NNI	Network-Network Interface
NOC	Network Operations Centre
OOB	Out of Band
POL	Passive Optical LAN
POP	Point of Presence
RCD	Remote Configuration Download
RESTful	Representational State Transfer
RIR	Regional Internet Registry
SAP	Service Access Point
SDK	Software Development Kit
SFTP	Secure File Transfer Protocol
SPB	Shortest Path Bridging
S-VLAN	Service VLAN
TFTP	Trivial File Transfer Protocol
TOR	Top of Rack
UNI	User-Network Interface
UNP	Universal Network Profile
VM	Virtual Machine
VXLAN	Virtual Extensible LAN
WPAN	Wireless Personal Area Network
WWAN	Wireless WAN
ZTP	Zero-Touch Provisioning

13 Related documents

- [1] SPB-Based Transportation Networks Design Guide, Alcatel-Lucent Enterprise
- [2] SPB-Based Transportation Network Technical Case Study, Alcatel-Lucent Enterprise
- [3] Mobile Campus 2.0 Solution Design Guide, Alcatel-Lucent Enterprise
- [4] “Context-aware security for the mobile enterprise” application note, Alcatel-Lucent Enterprise
- [5] “SDN Analytics for DDoS Mitigation”, Application Note, Alcatel-Lucent Enterprise.
- [6] Alcatel-Lucent Enterprise Datasheet URL: <http://enterprise.alcatel-lucent.com/?content=ResourceLibrary&page=datasheets>

- [7] Alcatel-Lucent Enterprise User Guide URL: <http://enterprise.alcatel-lucent.com/?content=ResourceLibrary&page=userguides>
- [8] OmniSwitch Common Criteria EAL-2 Certificate: <https://www.commoncriteriaportal.org/files/epfiles/Certification%20Report%20Omniswitch.pdf>
- [9] OmniSwitch Common Criteria NDcPP Certificate: <https://www.commoncriteriaportal.org/files/epfiles/Certification%20Report%20ALE%20NDcPP.pdf>
- [10] OmniSwitch AOS 6.7.1R04 FIPS 140-2 Certificate: <https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/3071>
- [11] OmniSwitch AOS 8.3.1R01 FIPS 140-2 Certificate: <https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/2996>
- [12] OmniSwitch 6860/6865/6900/9900 JITC Certificate: http://jitc.fhu.disa.mil/tssi/cert_pdfs/ALE_OS9900_OS6900_OS6860_OS6860E_OS6865_8.4.1_TN1628601_DTR1_13NOV2017.pdf