Final Report

Platform Commissioning

Jan 2022

Issue Details

	Name(s)	Signature
Principal author	Marc Ranford	24 Jan 2022
Internal reviewer	Stephen McCartney	24 Jan 2022
External reviewer	Paul Cooper (Zeetta Networks)	21 Jan 2022
Internal authorisation	Sachin Sree Raman	24 Jan 2022
Document number	DETI-ENC-001	
Document issue	1	
Issue date	24 Jan 2022	

Document Issue Log

1	Initial issue

About 5G-Encode

The 5G-ENCODE Project is a £9Million collaborative project aiming to develop clear business cases and value propositions for 5G applications in the manufacturing industry. The project is partially funded by the Department for Digital, Culture, Media, and Sport (DCMS), of the UK government as part of their 5G Testbeds and Trials programme. The project is one of the UK's biggest investments in using 5G to modernise manufacturing.

The key objective of the 5G-ENCODE project is to demonstrate the value of 5G as part of industrial use case delivery within the composites manufacturing industry. It is designed to validate the idea that using private 5G networks in conjunction with new business models can deliver better efficiency, productivity, and a range of new services and opportunities that would help the UK lead the development of advanced manufacturing applications.

The project will play a key role in ensuring that the UK industry exploits the 5G technology and remains a global leader in the development of robust digital engineering capabilities when implementing complex composites manufacturing processes.

The project will highlight how 5G features such as neutral hosting and network slicing can be applied to transform a private 5G network into a dynamically reconfigurable network able to support a wide range of applications (URLL/eMBB/MMTC) including industrial applications of Augmented Reality/Virtual Reality (AR/VR), asset tracking of time sensitive materials and automated industrial control though IoT monitoring and big data analytics. Such a dynamic network would enable new business models and creation of bespoke virtual networks tailored to specific applications or use cases.

A state-of-the-art test bed was deployed across three sites centred around the National Composites Centre in the Southwest of England. In support of the West of England Combined Authority (WECA) industrial strategy, the NCC plans to keep the test bed as an open access facility for the experimentation and development of new products and services for the composites industry after the completion of the 5G-Encode project. The location and nature of NCC's business would ensure the creation of an industrial 5G ecosystem involving multiple industry sectors and SMEs.

The project consortium, led by Zeetta Networks, brings together leading industrial players (e.g., Siemens, Toshiba, Solvay), a Tier 1 operator (Telefonica), disruptive technology SMEs covering all aspects of network design, deployment, and applications (Zeetta Networks, MatiVision, Plataine), a world-leading 5G network research group (High Performance Networks Group in the University of Bristol) and the NCC representing the high value manufacturing industry.

Contents

List of Figures

1. Introduction

This report contains the journey of specification, design, and implementation of the Open Radio Access Network (ORAN), 5G Stand Alone (SA) Private Cellular Network (PCN) installed at the National Composites Centre (NCC), Emersons Green, Bristol, and is targeted at IT professionals. This 5G SA PCN provides the foundation to the government backed Department of Digital, Culture, Media, and Sport (DCMS) 5G-Encode research project which has the overall goal to enable industrial use case demonstrators that springboard 5G technology within a manufacturing environment.

2. Journey

This section contains an insight into the implementation from a System Integrator's perspective (in this case the Infrastructure & Security Lead for the NCC Digital Business Unit) to deploy the 5G SA PCN. For each stage there is a deep dive of the process used to move from conception to vendor selection and deployment. Documented, in this section, are the challenges, workarounds and solutions faced during the integration.

2.1. Requirements

At conception, it was very apparent that a 5G PCN had a more moving components than its 4G PCN predecessor and as such vendor support (in this case Airspan Networks) was critical in ensuring the architecture was correct. This is due to the Open RAN nature of the deployment, which sees collaboration between non-proprietary equipment makers to create an end-end solution based on interoperability and open standards. Provided the equipment meets open RAN standards it should be compatible with equipment from other vendors that share these standards. Open RAN is favourable in private networks as it can create a more cost-effective and innovative solution when distinct functions are sourced from independent specialist suppliers.

Airspan Networks consulted with the lead partner on the Encode Project (Zeetta Networks) to fabricate the High-Level Design (HLD) to identify and plan all hardware and software components needed. The Low-Level Design (LLD) was then created ready for system integration. The following section contains both the HLD and the LLD as well as all components selected for deployment.

2.2. Architecture

Airspan's ORAN architecture supports the split architecture as depicted in Figure 1 where a gNodeB is composed of the following network functions:

- Radio Unit (RU): Handles the digital front end and the parts of the PHY layer
- Distributed Unit (DU): Located in proximity (in sense of timing) to the RU and handles the RLC, MAC, and parts of the PHY layer.
    - DU is a virtual network function that runs over commercial off-the-self servers.
- The Centralised Unit (CU):

- Centralised Unit - Control Plane (CU-CP): A logical node hosting the Radio Resource Control (RRC) and the control plane part of the Packet Data Coverage Protocol (PDCP) of the gNodeB-CU for a gNodeB (5g New Radio base station). The gNodeB-CU-CP terminates the E1 interface connected with the gNodeB-CU-UP (User Plane) and the F1-C interface connected with the gNodeB-DU (Distributed Unit).
- Centralised Unit - User Plane (CU-UP): A logical node hosting the user plane part of the PDCP protocol and the Service Data Adaption Protocol (SDAP) protocol of the gNodeB-CU for a gNodeB. The gNodeB-CU-UP terminates the E1 interface connected with the gNodeB-CU-CP and the F1-U interface connected with the gNodeB-DU.
- Both CU-CP and CU-UP are virtual network functions that run over commercial off-the-self servers.

Figure 1: Open RAN Split Architecture

- The front-haul between the RU and the DU is using Split 7.2a per O-RAN definition.
- The mid-haul between DU and CU is Split 2 (F1 interface) as defined by 3GPP.
- The back-haul between CU and the 5G Core (5GC) includes N2/N3 interfaces as defined by 3GPP.

Coverage requirements of NCC headquarters in Emersons Green, Bristol, drove the scale requirements of the Radio Access Network (RAN). The budget for equipment was considered to minimise investment where possible.

A remedial survey was carried out based on experience with the 4G deployment to indicate the most suitable locations for the 5G radio units (RU) and determine the number of RUs required to cover the workshop spaces with 5G connectivity. Environmental Radio Frequency (RF) interference was also taken into consideration and an overlap was also considered in this process for resilience. The chosen radio locations are shown below in Figure 2.

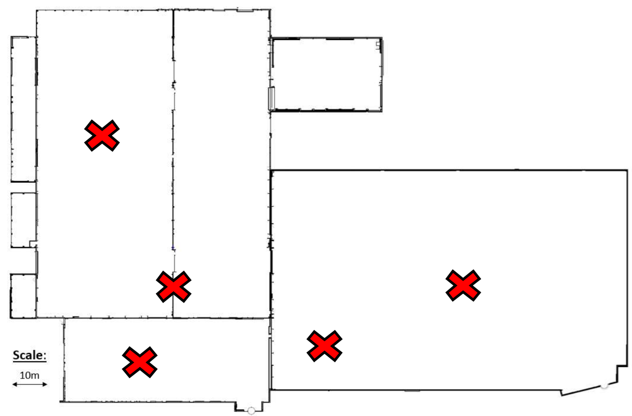

Figure 2: Floor Plan of the Workshop Area at NCC Emersons Green, Bristol

After establishing Radio Unit locations, the system requirements to manage these RUs was designed and equipment procured. The HLD was updated after consultation with Airspan Networks.

2.2.1. High Level RAN Overview

The HLD displayed in Figure 3 below allows for 4 x Radio Units connected to the relevant back-end components that make up the RAN. There is a maximum of a 2 to 1 relationship between the Radio Units (RUs) and the Distributed Units (DUs) and all DUs talk back to a single Centralised Unit (CU) in a tiered format.

A decision was taken to not to provide full coverage in one area of the workshop to reduce RAN infrastructure cost. The RU removed from the coverage plan being planned for a future installation activity as the RAN product matures and this RU can be supported within the infrastructure purchased for the 4 x RUs planned for initial deployment.

NB: The requirement for complete coverage at NCC Emersons Green transpired to be 5 x RUs so the decision to carry 1 x RU for future use meant that coverage in one area was sub-optimal. This would not affect use case execution.

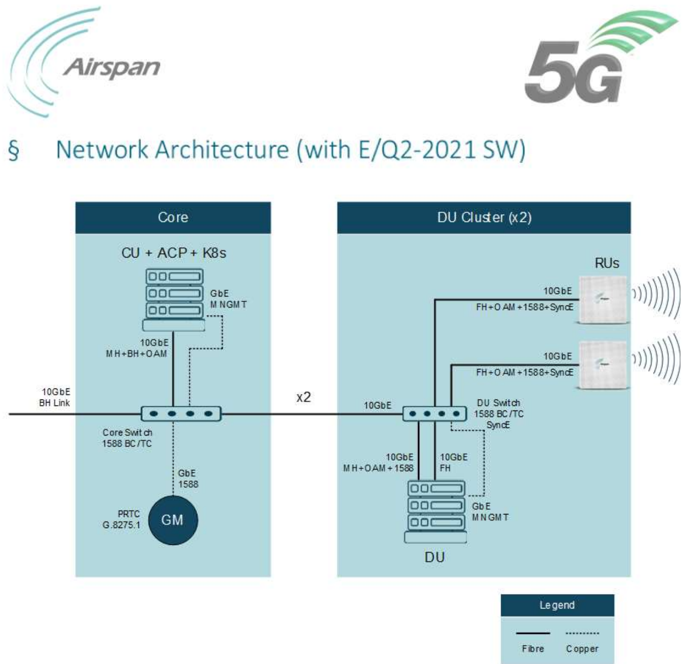

Figure 3: 5G RAN High Level Design Provided by Airspan Networks

A set of hardware requirements accompanied this design to allow for all relevant software to run in an optimised manner, this is discussed in a later section.

2.2.2. Low Level RAN Overview

The LLD iterations in Figure 4 to Figure 8 expand on the previous section and add the related network segments to allow for intercommunication between 4 x RUs and the supporting Distributed Units (DUs), Centralised Unit (CU) and Packet Core. IP Addressing can be any RFC1918 address schema. For the purposes of the project the concept of front, mid and backhaul addressing was implemented, and the IP addresses shown in Figure 5 used. Precision Time Protocol needs are discussed in later sections.

Note: The LLD was created in collaboration between Zeetta Networks and NCC. It is presented in a layered approach to enable easier reading, larger diagrams can be found in the appendices.

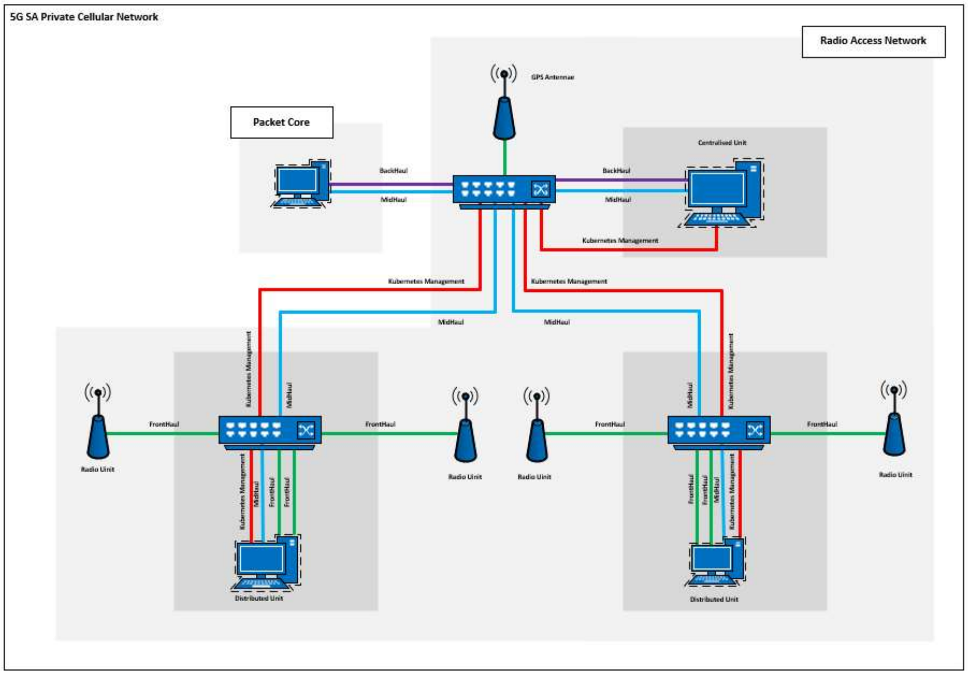

Legend:

- Kubernetes Management:
- Front-Haul:
- Mid-Haul:
- Back-Haul:

Figure 4: RAN LLD Provided by Consortium (Network Segments)

A larger diagram can be found in Appendix A.

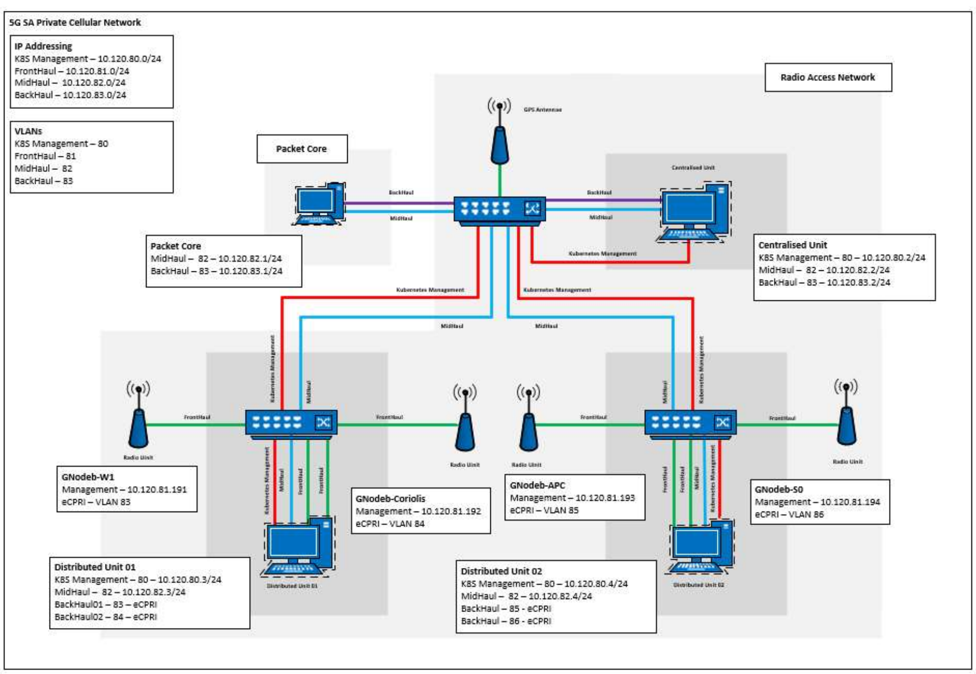

Figure 5: RAN LLD Provided by Consortium (VLANs & IP Addressing)

A larger diagram can be found in Appendix B.

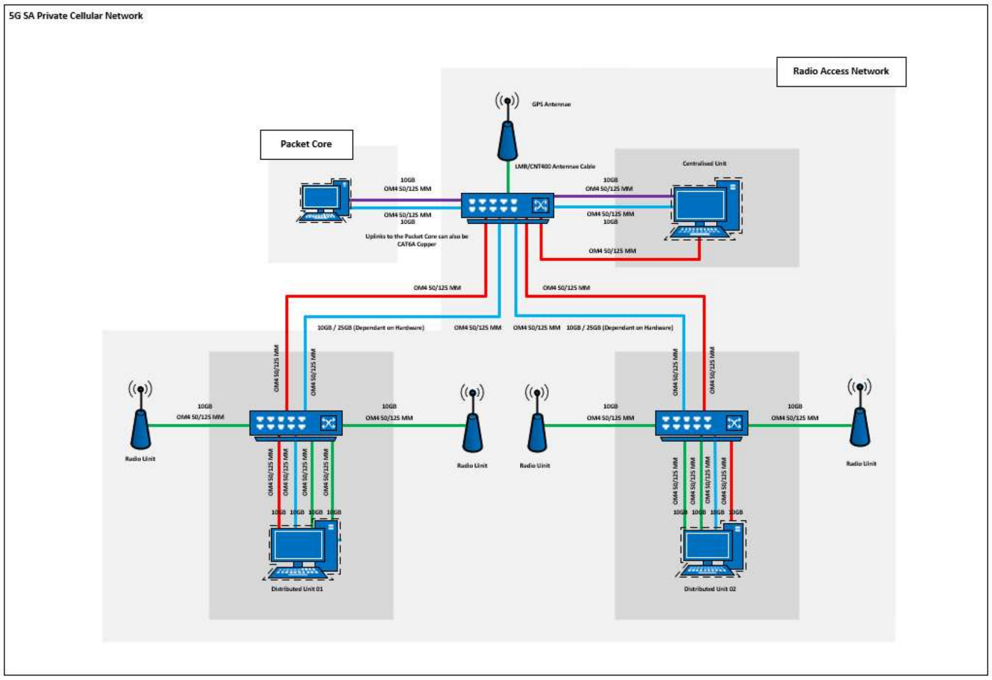

Figure 6: Physical Connectivity

A larger diagram can be found in Appendix C.

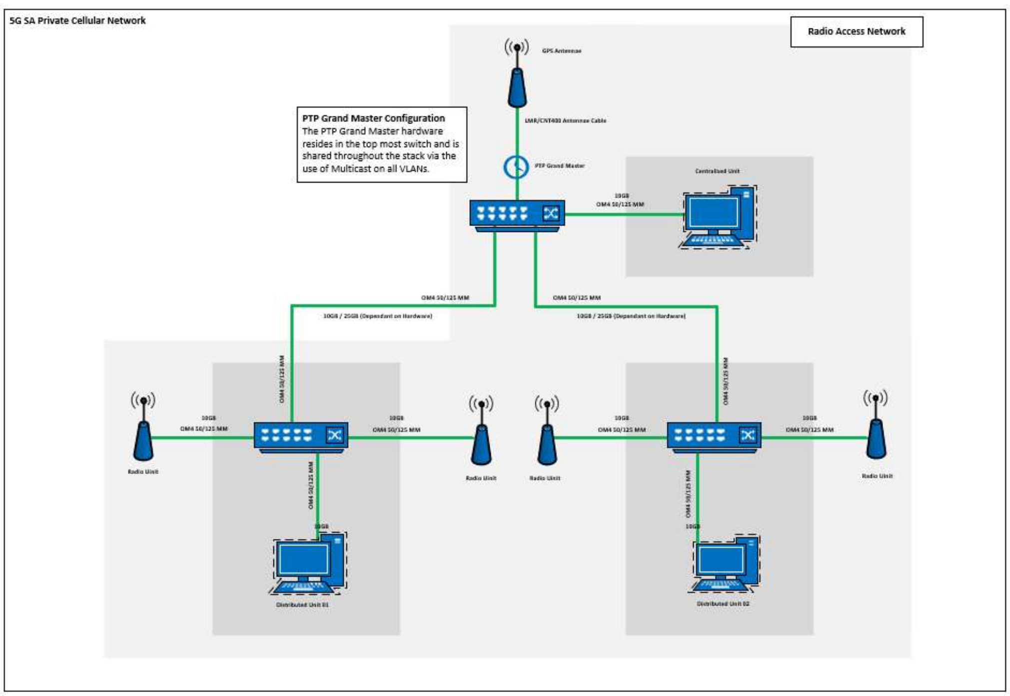

Figure 7: RAN with PTP Clock Integration

A larger diagram can be found in Appendix D.

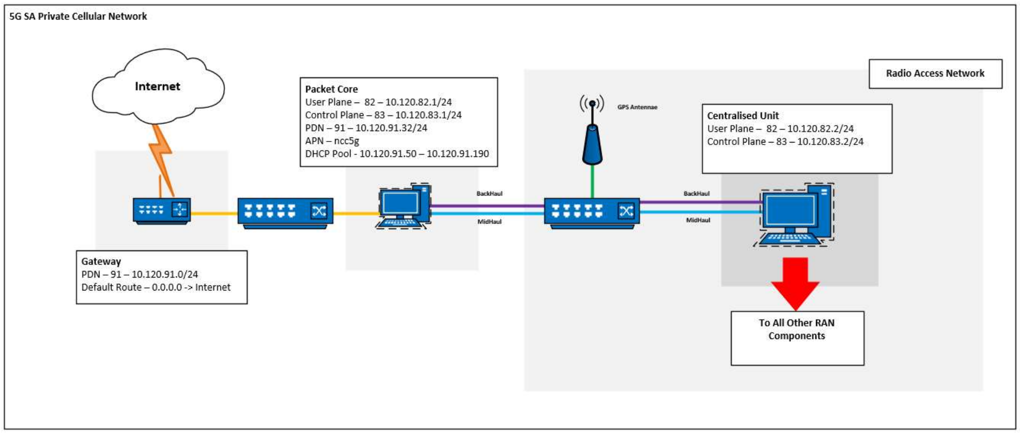

Figure 8: Packet Core Integration

A larger diagram can be found in Appendix E.

Hardware and software selection was limited to vendor interoperability and supported configurations and as such are discussed in later sections.

2.3. Hardware Selection

The following section lists out the hardware specification for each of the components required to deliver the overall solution.

2.3.1. Compute Function

The RAN vendor, Airspan Networks, listed the following server hardware for their supported configuration for the deployment at NCC Emersons Green:

- 1 x Centralised Unit
  - 1 x Dell R740 Chassis
  - 1 x PowerEdge R740/R740XD Motherboard
  - 2 x Intel Xeon Gold 5218R 2.1G, 20C/40T, 10.4GT/s, 27.5 M Cache, Turbo, HT (125W) DDR4-2666
  - 1 x Chassis with up to 8 x 2.5" SAS/SATA Hard Drives for 2CPU Configuration
  - 1 x PowerEdge 2U LCD Bezel
  - 1 x Riser Config 5, 6 x8, 2 x16 slots
  - 12 x 16GB RDIMM, 3200MT/s, Dual Rank
  - 1 x iDRAC9, Enterprise
  - 2 x 960GB SSD SATA Read Intensive 6Gbps 512 2.5in Hot-plug AG Drive, 1 DWPD, 1752 TBW
  - 1 x PERC H330 RAID Controller, Low Profile
  - 1 x Dual, Hot-plug, Redundant Power Supply (1+1), 750W

- 1 x Trusted Platform Module 2.0
- 2 x Intel XXV710 Dual Port 10/25GbE SFP28 Adapter, PCIe Full Height
- 1 x Intel X550 Quad Port 10GbE BASE-T, rNDC

- 2 x Distributed Units
  - 1 x Dell R740 Chassis
  - 1 x PowerEdge R740/R740XD Motherboard
  - 2 x Intel® Xeon® Gold 6240R 2.4G, 24C/48T, 10.4GT/s, 35.75 M Cache, Turbo, HT (165W) DDR4-2933
  - 1 x Chassis with up to 8 x 2.5" SAS/SATA Hard Drives for 2CPU Configuration
  - 1 x PowerEdge 2U LCD Bezel
  - 1 x Riser Config 6, 5 x8, 3 x16 slots
  - 8 x 32GB RDIMM, 3200MT/s, Dual Rank
  - 1 x iDRAC9, Enterprise
  - 1 x 960GB SSD SATA Read Intensive 6Gbps 512 2.5in Hot-plug AG Drive, 1 DWPD, 1752 TBW
  - 1 x PERC H330 RAID Controller, Low Profile
  - 1 x Dual, Hot-plug, Redundant Power Supply (1+1), 750W
  - 1 x Trusted Platform Module 2.0
  - 2 x Intel XXV710 Dual Port 10/25GbE SFP28 Adapter, PCIe Full Height
  - 1 x Intel X550 Quad Port 10GbE BASE-T, rNDC
  - 1 x Silicom eASIC ACC100 FEC Accelerator, PCIe Full Height

Note: Vendor recommended SFP/SFP+ modules are installed in all network cards.

2.3.2. Switching

The lead partner, Zeetta Networks, listed the following LAN hardware for the interconnectivity between RAN components:

- ADVA FSP150-XG118PRO

The above device is classed more as an Ethernet Demarcation Device as it can manage traffic flows at a lower level than your "out-of-the-box" switch. The device has 8 ports and is designed to provide connectivity specifically designed for 5G implementations with the additional facility to provide edge compute capacity where required.

Note: All upstream connectivity from the 5G stack uses standard switching and vendor recommended SFP/SFP+ modules installed in all network hardware.

2.3.3. Timing

The lead partner, Zeetta Networks, listed the following time clock hardware for the 5G stack:

- Oscilloquartz OSA5401

The above device is a Small Form-Factor Pluggable GNSS receiver and plugs into the above referenced network hardware to provide GPS based precision timing compliant with G8275.1 on the LAN to synchronize the RAN devices. The features embedded in the ADVA FSP150-XG118PRO then share this time synchronisation across other units via multicast on all VLANs. The antenna for the device is attached to a 120m cable via risers inside the fabric of the building to the roof for reliable GPS connectivity.

2.3.4. Radio Units

The RAN vendor, Airspan Networks, listed the following RU hardware for their supported configuration for the deployment at NCC Emersons Green:

- Airspan AirVelocity 2700

Airspan Networks AirVelocity 2700 has the following features required to facilitate 5G connectivity:

- 4T4R Tx/Rx Paths
- Up to 100 MHz Bandwidth
- 2+ Gbps
- CBRS Supported
- Split 7.2x

All connectivity provisioned to the radios is 10GB leveraging specific LC SFP+ Modules and are listed as the following:

- Multi-Mode - Finisar FTLX8574D3BCL
- Single-Mode - Finisar FTLX1475D3BCL

All RU SFPs at NCC Emersons Green use Finisar FTLX8574D3BCL (Multi-Mode) modules for connectivity.

2.4. Software Selection

The following section details the Operating Systems, pre-requisites and operational software required to enable the RAN and Packet Core to function as a complete Private Cellular Network. All Operating Systems are Linux based with a Kubernetes (K8S)/Kernel-based Virtual Machine (KVM) overlay.

2.4.1. Operating Systems

Operating Systems are installed on the hardware specified in section 2.3.1 and are part of the build document supplied by the RAN vendor. All Operating Systems are Linux based and are a CentOS version. CentOS 7 is the long term supported distribution that runs as the underlying Operating System to a Kubernetes cluster and host Operating System for KVM hypervisor.

CentOS 8 is used to underpin Airspan's ACP software that controls all containers relating to the Virtual Radio Access Network (vRAN) deployment (CU and DU software). This part of the deployment is a virtual machine configuration.

2.4.1.1. Challenges

Specific Operating System versions were quoted by the RAN vendor (Airspan Networks) and not the current release. As such a legacy version was obtained and installed. Special considerations were required to ensure these operating systems are deployed on a security compliant network with minimal attack vectors e.g. the installations were restricted from accessing the internet and other parts of the wider LAN.

2.4.2. Packet Core

The Packet Core is installed on a virtual machine upstream from the Radio Access Network. The Packet Core selected by the lead partner, Zeetta Networks. was a Druid Raemis 5G SA system with HSS+ enabled. HSS+ allows for the SIM database to be referenced on an the already existing 4G Packet Core as opposed to having its own integrated SIM database. This means that devices can use a common SIM to authenticate to and access either the 4G or 5G RAN. The Packet Core uses Back-haul and Mid-haul connectivity to manage the user control plane and route the user data plane through the RAN to the internet and edge compute resources.

The packet core is installed on the latest CentOS 7 Operating System on a virtual machine on separate hardware to the previously discussed 5G stack.

2.4.2.1. Challenges

Druid's support team were engaged at through both the Packet Core integration and subsequence RAN integration phases of the project to diagnose and resolve encountered issues.

SIMs using the Public Land Mobile Network (PLMN) ID of 001 01 were imported into the 5G packet core and used on some CPE/UE devices to enable the successful attach and data flows for devices with 5G SA limitations that included being limited to that PLMN ID. This configuration issue is discussed further in section 2.7.

Note: A PLMN ID is comprised of two parts; a 3-digit Mobile Country Code (MCC) and a 2-digit Mobile Network Code (MNC)

2.4.3. RAN

RAN software provided by Airspan was selected for the project by Zeetta Networks considering availability, capability and cost. This software is installed on hardware platforms with Operating Systems as discussed in 2.4.1. The construct of this software is a combination of physical installations sitting on bare metal hardware, Kubernetes containers and virtual machine deployments. The physical installation on both CU and DU hardware was preceded by a substantial pre-requisite list with specific versions of software needed to prevent impact on the performance of the solution when in use, a significant element was the Kubernetes configuration.

Kubernetes is installed with network plugins (specifically Flannel & Multus) to allow for clustering of the three servers listed in section 2.3.1 and to run containers that are pre-configured for specific functions throughout the deployment. All containers are monitored and configured from the Airspan Control Platform (ACP).

ACP was installed to run on a virtual machine located on the Centralised Unit hardware. It is a tiered web application with a Linux derivative of a Microsoft SQL back end containing all configuration data. The version of Microsoft SQL will depend on the scale of the deployment and the retention requirements of logs and metrics. This project utilised Microsoft SQL Express.

2.4.3.1. Challenges

The software is a constantly evolving product due to the immaturity of the Open RAN technology. The installation and configuration consistently encountered multiple challenges in both stability and performance of the product. The pre-requisite software configuration required modification on multiple occasions to enable optimisation of the hardware and new releases of the vRAN to be installed. During the project the product evolved into significantly more stable and reliable solution that is installed today. These issues are described in more detail in section 2.5.4.

2.4.4. Monitoring & Statistical Analysis

Monitoring & statistical analysis was performed via two functions:

1. Cellular statistics were generated by the ACP platform and visualised by the engineering team at Zeetta Networks
2. Live network metrics were captured and provided by Accedian using their Skylight product

Accedian used a hybrid approach that allowed for a mixture of in-line hardware probes, port mirroring configuration, docker based software probes and Android applications used on UE devices. The probes enabled an end-to-end transparent picture of network metrics and availability to be visualised.

To visualise the network, in-line hardware probes were implemented on every uplink between the back-end hardware. Probe detection was configured on each link to create full transparency. A port mirror was configured on the uplink to the gateway for application detection on ingress/egress traffic on the 5G VLANs and the associated metrics for uplink/downlink speeds. Software probes were located inside each use case VLAN allowing for part of the end-to-end visualisation and the Android software allowed for metrics to be generated on the end device (in this case a UE) to visualise real-time 5G metrics.

All data to be analysed was sent to a centralised Accedian software installation that processed receiver data and sent to an online portal for visualisation and presentation to the end user.

2.4.4.1. Challenges

Due to the mainstream hardware vendor support of the software, there were compatibility issues with the upstream hardware from the 5G stack to detect some of the metrics. Specific models of EdgeCore hardware running PICOS (open networking stack) could not negotiate a connection with the SFP modules resulting in the inability to monitor traffic outside of the core. As such an alternative piece of equipment (an Accedian ANT) was installed in-line to mitigate the issue and capture the required metrics.

Further hurdles were faced when deploying some elements of the core server that all probes talk back to due to available hypervisor technology, other issues such as nested virtualisation also reared their head and are discussed in 2.5.5.

2.4.5. Software Defined Networking

The Software Defined Network solution was provided by lead partner, Zeetta Networks. The solution encompassed all aspects of the network allowing for VLAN creation, visualisation, and service control functionality. This developed as the project progressed to encompass routing and the ability to create APNs within the Packet Core. Solution upgrades were delivered as functions were completed. A separate instance of this solution was integrated into the remote NCCi site, at Filton. This instance along with a small 5G network was used to demonstrate a new SDN capability Multi-Domain Orchestration (Management of multiple sites from a centralised cloud-based platform).

2.4.5.1. Challenges

The routing between NCC Emersons Green and NCCi Filton had to be provisioned on an already "in-service" link so due diligence was necessary to enable connectivity. This was a timely process due to resource allocation within the NCC IT department.

Access during periods where COVID infection rates were high, and some employees were off work due to infection hindered the deployment timescales.

Features available at the time of implementation only covered Open Systems Interconnect (OSI) model Layer 2 configurations.

2.5. Deployment

The following section provides deployment information throughout the journey, specific configuration of components that overcome hurdles and the final configuration that is currently operational.

2.5.1. Computer Hardware

The DU servers suffered repeated stability issues. Investigation revealed the hardware specification of the DU server recommended by Airspan conflicted with hardware capabilities of the R740 platform provided by Dell. Airspan Networks specify a set number of cores within their requirements documentation for the DU server and later iterations of the documentation state that the deployment requires a single processor configuration, this was not possible due to the number of slots required to allow for network connectivity and an (Field Programmable Gate Array) FPGA card.

2.5.1.1. Challenges

The system was installed in a resilient configuration and subsequently altered to the Airspan recommended configuration. The consequences of this alteration created some constraints in system robustness that need documenting. These are as follows: with the requirement to leverage multiple PCIe risers within the specified chassis it was necessary to have multiple processors. This did not align with the Airspan single processor configuration as per the documentation as the removal of one of these processors would result in loss of functionality of a Peripheral Component Interconnect Express (PCIe) riser causing a network card installed in the riser to move to an offline state. An increased number of PCIe risers within an R740 platform requires an additional CPU to be brought on-line. Loss of connectivity to a card would result in loss of either the Front-haul, Mid-haul, or Kubernetes Management networks. The Airspan configuration required that all ports carrying the enhanced Common Public Radio Interface (eCPRI) protocol to arrive onto the riser hosting the FPGA Card. To enable other links to be connected to the platform these were positioned on the other riser hosted by the second CPU.

DU Processor upgrades were required to accommodate the overall number of cores on a single processor as opposed to the correct number of cores split across both processors. The processor configuration also had to accommodate the number of PCIe risers required to support 4 x SFP modules, a storage controller, and an FPGA card. Airspan Networks supported a dual processor configuration for a virtual “all in one” deployment that accommodated the previously mentioned core requirement, so these chips were specified (upgraded from Intel Xeon Silver in a dual socket configuration to Intel Xeon Gold dual socket configuration) and installed.

Post-processor upgrades the configuration also allowed for redundancy across network interface cards for the radio units (split eCPRI interfaces per network card). When the radio eCPRI interfaces were brought online the radios were showing what appeared to be an out-of-sync time source. Further investigation proved that when a single radio was online the issue was not present. The issue was escalated to Airspan Networks to which it was discovered that both eCPRI interfaces needed to be configured on the network card on the same riser as the FPGA card and separating these will cause a slight time delay. Figure 9 below displays the layout of an R740, the final layout of the hardware configuration for the 5G vRAN DU application is also displayed.

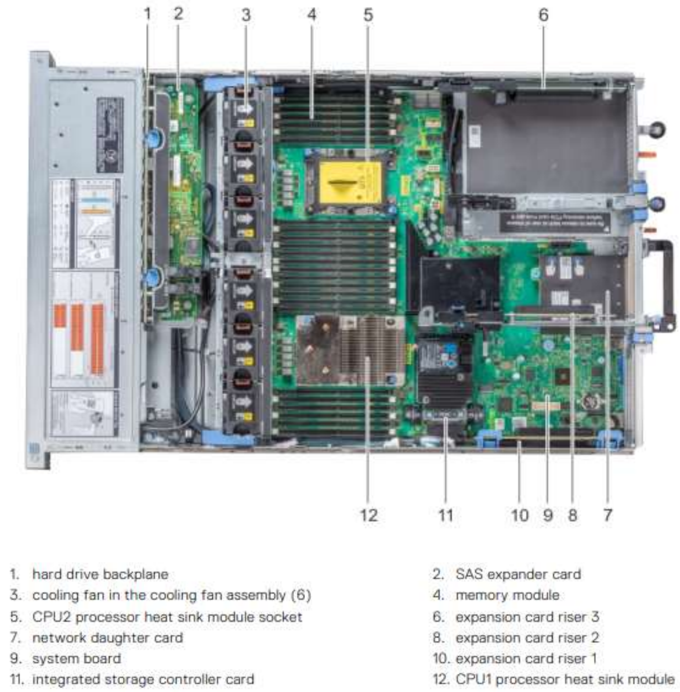

Figure 9: Dell R740 DU Physical Layout

Airspan Specific Hardware Configuration:

Hardware Location Reference	Item Description	Use
12	CPU1 Location	Processing power for container software
5	CPU2 Location	Processing power for container software
6	Expansion Card Riser 3 – Dual Port SFP+ Intel NIC	Mid-Haul & Kubernetes Management networks
10	Expansion Card Riser 1 – Dual Port SFP+ Intel NIC	eCPRI to Radio Units (1 per Radio Unit)
8	Expansion Card Riser 2 – PERC H330	RAID controller for hard disk redundancy

NB: All other RAN computer hardware specification is as recommended by Airspan Networks.

2.5.2. Network Hardware

The ADVA FSP150-XG118PRO is a Telco grade switch designed to work with 5G environments, as such this style of switching is new to most enterprise network engineers. Specialist configuration support and engineer training was required.

2.5.2.1. Challenges

Due to the niche nature of the above product, the lack of understanding made troubleshooting issues (namely the Precision Time Protocol (PTP) time synchronisation issue) in a timely manner challenging and requiring engagement of professionals from all suppliers.

2.5.3. Operating Systems with Kubernetes & KVM Hypervisor

2.5.3.1. Centralised Unit (CU)

As stated previously, the CU server is built on a CentOS 7 platform with KVM Hypervisor installed on it. It also acts as a worker node within a Kubernetes cluster. The installation is based on CentOS7 1908 with “Infrastructure Server” and “Development Tools” selected on installation. The pre-requisite installation files provided by Airspan Networks make up the rest of the bare metal installation.

NB: all versions of the software and pre-requisites are to be kept “as is” until further mainstream releases of containers are provided by Airspan Networks.

2.5.3.1.1. Kernel Virtual Machine (KVM) Hypervisor

KVM is installed as part of the pre-requisite installation and configuration. This component allows for virtual machines to run on the bare metal and function as a hypervisor. Airspan Networks recommend that the Kubernetes Master (the unit that controls all containers and workers within Kubernetes) run on a virtual machine on the CU hardware.

This virtual machine requires network connectivity to the Kubernetes Management network and due to the selection of operating system it is necessary to create a network bridge for interaction between the host and the virtual machine.

2.5.3.1.2. Kubernetes (K8S)

Kubernetes is required to underpin the containers running on the network that form the vRAN. The Kubernetes Control Plane virtual machine can talk to the worker nodes via a management network and communication between the pods are facilitated via the use of the Flannel and Multus add-ins. These plug-ins allow expansion and interactivity of virtual networking functions inside Kubernetes and take advantage of SR-IOV capabilities within the network interface cards leveraged as part of the hardware specification.

NB: The virtual network capabilities manifest themselves to the Operating System as additional network adapters with specific class A subnets associated with them.

Figure 10 below displays the architecture of the Kubernetes Cluster:

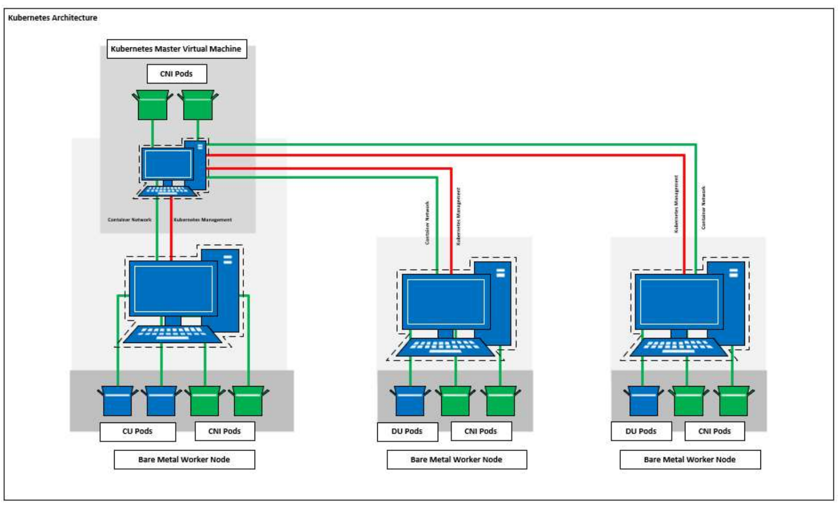

Figure 10: Kubernetes Cluster

Blue: Containers provided by Airspan Networks

Green: Container Network Interfaces (CNI) via Flannel & Multus Add-ins

NB: A larger diagram can be found in Appendix F.

After Kubernetes is installed to the Kubernetes Master virtual machine the add-ins are installed and the pods containing CNI configuration start. Each worker node added to the Kubernetes Master will then have pods deployed to it supporting the CNI and will start as soon as the node becomes available on the Kubernetes Management interface (this is always the interface that adds the node to the cluster).

When all nodes are added and pods containing CNI functionality start, the node switches to a "Ready" status and is available for the addition of pods containing third party applications.

2.5.3.2. Distributed Unit (DU)

The DU server is also built on a CentOS 7 platform. The installation is based on CentOS7 1908 with "Infrastructure Server" and "Development Tools" selected on installation. The pre-requisite installation files provided by Airspan Networks make up the rest of the bare metal installation.

NB: all versions of the software and pre-requisites need to be kept "as is" until further mainstream releases of containers are provided by Airspan Networks.

The DU servers have the most intensive workload of all servers in the deployment and post installation, are required to be added to the Kubernetes Cluster as worker nodes. When the third-party pods are deployed that make up part of the Virtual Radio Access Network (vRAN), they commandeer physical interfaces defined inside configuration files to send eCPRI traffic to the Radio Units. These physical interfaces are hidden from the base OS when active. The container software also leverages the FPGA card installed in the hardware via a direct access mapping (physical hardware presented to a virtual component) to deliver optimum performance on the Front-Haul networks.

2.5.3.3. Challenges

Throughout the entire deployment the underlying hardware required configuration to allow for single processor functionality in a unit that required dual processors to accommodate the number of expansion cards to support the correct number of attached radios. During deployment unexpected behaviour was experienced from multiple components. This unexpected behaviour includes instability of Kubernetes pods, resource spikes that eventually caused Operating System failure and intermittent connectivity to Radio Units. Investigation of the stability issues resulted in processor upgrade to the DU servers to be a match to an alternative, more powerful Airspan configuration. After processor upgrades the same behaviour was observed until a pre-requisite component, named "tuned", was configured to emulate a single slot processor for Kubernetes components whilst leaving all other available cores for the Operating System to leverage.

This resolved most of the issues although container restarts still happened intermittently, and Operating System freezes caused by resource consumption also occurred from time-to-time. The workaround for this was to enable Watchdog (a method of physically restarting a server via the use of an IPMI interface if a response from the host Operating System is not detected over a set period) to allow for recovery in the event of an issue being present. Containers were also restarted in a controlled manner via a cron job running on the Kubernetes Master VM that re-deployed DU containers every morning.

NB: As the software matures these issues will eventually evolve into a stable solution.

2.5.4. Timing

A key difference between 4G and 5G networks is timing. For the radio to achieve the higher capacity and lower latency offered by 5G the network timing precision is increased. In 4G network timing protocol (NTP) was sufficient to support the system. For 5G NTP does not meet the timing precision required, thus Precision Timing Protocol (PTP) needs implementing. To implement PTP all devices in the radio subsystem need the capability to support the protocol. The OSA5401 Grand Master Clock required configuration inside the ADVA FSP150-XG118PRO to allow for a time source to be made available across all 3 pieces of hardware in a Master/Slave architecture.

2.5.4.1. Challenges

The PTP Grand Master Clock was not plug and play, it required extensive configuration when leveraged across multiple pieces of hardware, as such ADVA Professional Services were engaged to configure the OSA5401. The method used to broadcast the PTP in accordance with Airspan documentation was multicast and this was configured to broadcast on all VLANs. Due to the accuracy requirement of timings for 5G deployments it was also necessary to take into consideration the length of cable (120m at NCC Emersons Green) and the material of the core of the cable. Both factors require inputting into an equation that calculates an offset to guarantee accuracy of the time source. When operational, Airspan DUs and RUs verified that timing was correct via the CLI.

2.5.5. Airspan ACP

Airspan ACP is installed on CentOS 8.4.2150 and is a minimal deployment on a virtual machine within the CU hardware. The network configuration needs to reflect the Kubernetes Management network so that ACP can talk to all nodes inside the Kubernetes Cluster. After patching, the Linux version of Microsoft SQL Server is installed on the OS and configured with an instance to which the System Administrator (SA) password needs to be configured. Post installation, the Airspan ACP software is installed using the SA credentials to build the database back-end to the web front-end upon installation.

As soon as this process has been completed ACP needs to be licensed and configured. Each node in the Kubernetes cluster is added to ACP so that the containers can be configured and leveraged to form the complete end-to-end Radio Access Network.

Figure 11 below displays where Airspan ACP sits in the architecture of the Kubernetes Cluster.

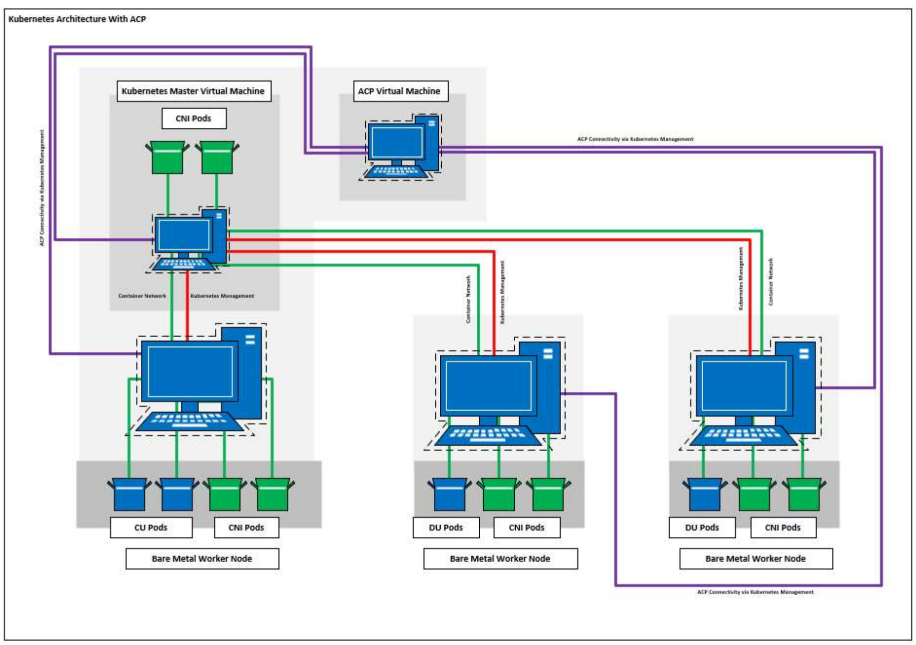

Figure 11: Kubernetes Architecture with ACP

Purple: ACP Connectivity via Kubernetes Management Network for ACP to control all pods contained within worker nodes.

NB: A larger diagram can be found in Appendix G

2.5.5.1. Challenges

No challenges were observed when installing ACP.

2.5.6. Packet Core

The packet core leveraged on both the 4G and 5G deployments is developed by Druid and the product is Raemis. The Raemis packet core operated on the up-to-date release of CentOS 7 and can be patched regularly without issue. The installation is script based and installs by pulling information from the internet.

Post installation configuration is required to allow the packet core to see the RAN. This is in the form of a Control Plane and User Plane L2 connection made available between the packet core and the CU containers (diagrammatical representation in Figure 8; earlier in the document). Connectivity is provisioned via 2 physical interfaces present on the host OS and selected during initial configuration of the Raemis software.

Note: The RAN User Plane and RAN Control Plane Network Interface Cards need to be configured as purely L2 connectivity, so default routes need to be removed during initial configuration.

The CU server is then detected as a gNodeB, screenshot displayed in Figure 12 below:

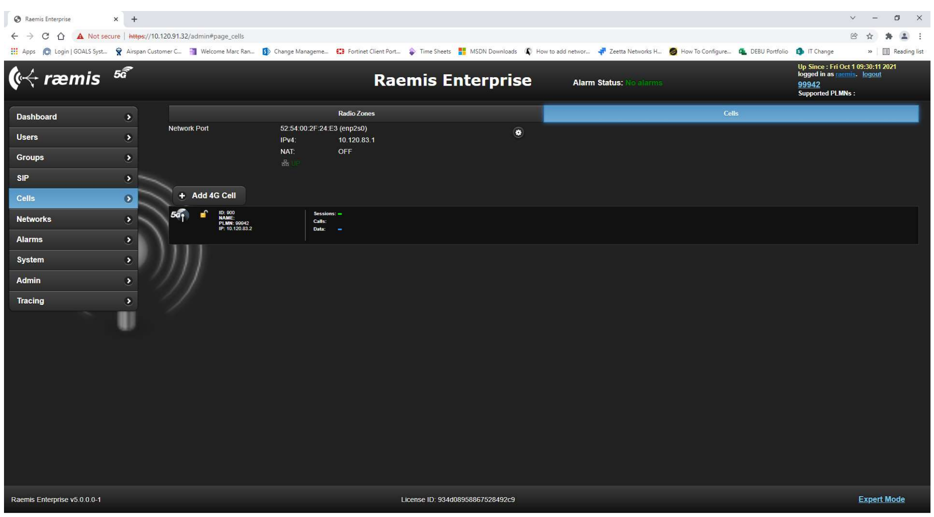

Figure 12: gNodeB Inside Raemis

Note: If the control plane is connected but the user plane is not then the gNodeB will be present but there will be no data flow from the remote devices that are attached to the packet core.

After the RAN is provisioned, it was necessary to configure an APN and attach it to the third and final physical network interface within the software. A default route is required on this Network Interface Card to be able to access resources both LAN and WAN side.

Figure 13 below displays the configured Network Interface Cards inside the Raemis software

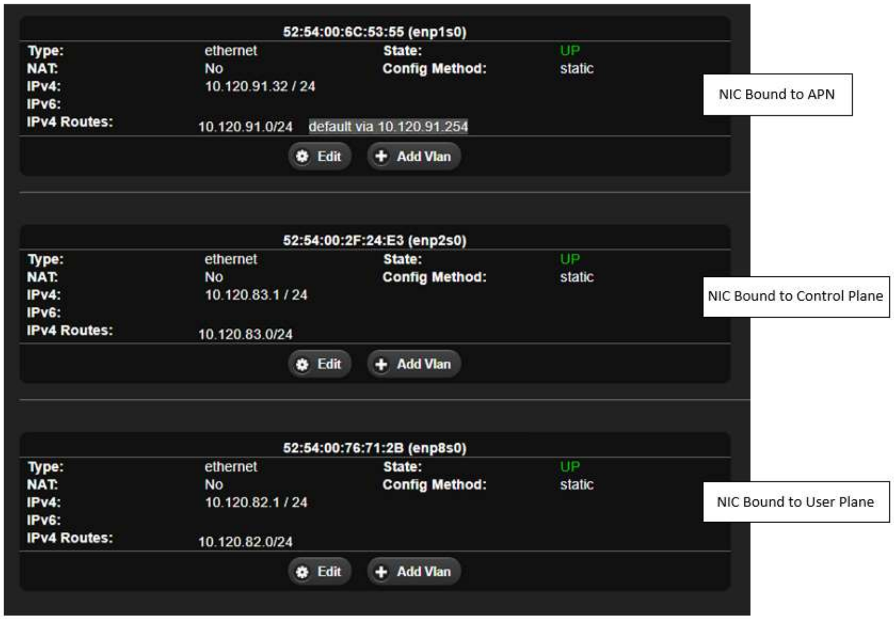

Figure 13: Raemis NIC Configuration

As observed above, the Network Interface Card associated with upstream connectivity has a default route. This card can be seen in Figure 14 bound to the APN “ncc5g” with associated core services connectivity.

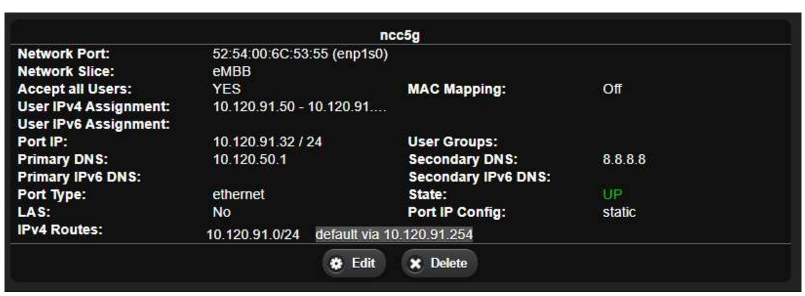

Figure 14: APN Configuration

When a device attaches with the configured APN and its associated registered SIM card it is allocated an address within the designated range and connects to the network.

Note: Licensing dictates the number of APNs, Networks, gNodeBs and Registered SIM Cards available within the software.

Additional features within the software enable a shared SIM card database (HSS+) located on the 4G core (PLMN ID is also 999 42) used to maintain a single set of SIM cards for network access.

2.5.6.1. Challenges

No issues were observed whilst deploying the Packet Core although as the project progressed it was apparent that support was required to confirm connectivity issues with the 5G RAN vendor. Collaboration between both vendors ironed out issues with User Plane and Control Plane configuration to achieve a fully functional system.

Further issues were observed with Customer Premises Equipment (CPE) and User Equipment (UE) devices attaching to the RAN and Packet Core. These issues were also external to the RAN and Packet Core software and related to 5G cellular interfaces specific to these products not recognising the test PLMN ID of 999 42. This was test PLMN used in the NCC 4G that was re-used for NCC 5G. To devices with constraints to attach to the 5G RAN it was necessary to use a SIM card with the PLMN ID of 001 01. Many of the devices also required specific configuration to enable 5G SA network detection and attach protocols. Once configure for 5G SA attach and with a recognised PLMN i.e. 001 01 available on the installed SIM these devices attached to the 5G RAN.

Note: not all devices once configured would attach to the 5G RAN thus it is recommended that care is taken in the selection of devices for use on a 5G SA networks at this time.

2.5.7. Monitoring & Statistical Analysis

Network Probe

The deployment of the Accedian platform involved installation of in-line multimode SFP+ modules to a multitude of third-party vendor hardware. EdgeCore 10GB Enterprise Fibre Optic switches accepted the modules but the 1GB Copper switches with 6 x SFP+ expansion modules would not accept the modules. ADVA switching and Dell server hardware observed no issues and the SFP+ modules functioned as expected.

The construct of the system was extensive involving 5 x virtual machine deployments for components to manage hardware/buffer/analyse and send the metrics to a cloud-based platform for visualisation. These virtual appliances were supplied in specific formats and leverage a proprietary configuration shell. Configuration of the in-line SFP+ modules, the ANT traffic interception devices and port mirroring statistical analysis (for application aware traffic logging) are completed via the web portals on the individual appliances after post setup tasks have been completed.

Furthermore, the base system configuration, is that each probe that sits in-line on the network required an IP address to be able to provide an end point within a required network segment. A Single probe can sit in multiple VLANs and have multiple IP addresses associated with it for complete transparency. In the case of the Encode project, each use case has a separate network segment associated with it meaning probes can be positioned in a full end-to-end configuration giving full transparency on traffic flows and LAN side metrics.

The port mirroring configuration (PVX Sensor) for application detection was deployed on the uplink between the core switch and default gateway.

Figure 15 below displays the topology of the Accedian Networks monitoring and statistical analysis platform:

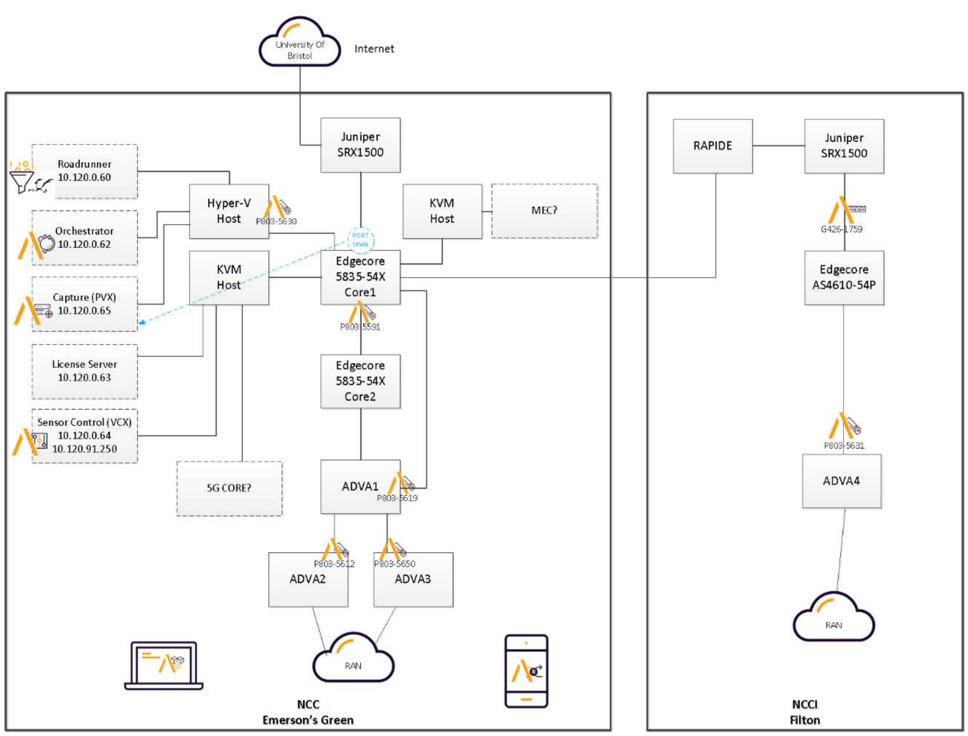

Figure 15: Accedian Topology

Note: All virtual appliances can be found on the left-hand side of the diagram.

SFP+ In-line probe

ANT Traffic Interception Device

- Android Probe
- Docker Based Software Probe

All cellular specific statistics were generated within the Airspan ACP platform and processed by Zeetta Networks. These metrics are created from statistical counts points included in the software of the CU and DU elements. These metrics are collated and sent to ACP as a single XML file at the end of each collection period. The default collection period of every 15 minutes was used for network metric processing and collection. The collected metrics are stored in ACP as XML files, loaded by ACP and visualised on screen as tables of text and numeric data. These tables were exported in CSV format and transferred to Zeetta for chart preparation and initial analysis. Charts and pre-analysis for each use case were returned to each NCC use case lead for further analysis and use in understanding the benefits of 5G.

Note: Similar 4G metric files were generated every 60 minutes. 5G metrics can be set to be generated every 60 minutes, the team decided to retain the 15-minute interval for metrics generation.

Cellular

When considering cellular network level metrics, it must be remembered that the metrics reflect all activity on the network in the collection period. Therefore, distinguishing metric data that were created as a result of executing a use case from all other network related metrics data is not possible. To mitigate this use cases were planned such that the only network activity during the use case was generated by the use case itself. This created network metrics that were as close as could be to expected results from the use case. In busier, or mature, 5G networks, the network metrics created by the RAN will represent overall network performance and not metrics specific to a use case.

2.5.7.1. Challenges

Network Probe

Due to all use case demonstrators utilising Virtual Machine deployments on a multitude of Hypervisor platforms it was necessary to enable nested virtualisation to allow for the software probes to install and function within a containerisation environment. Some of the use case demonstrators leveraged VDI style shared graphics devices for 3D rendering which caused the nested virtualisation to malfunction when certain Microsoft Windows features were enabled. As such a probe running directly on a use case demonstrator server or workstation was not possible so a workaround was required. The Accedian container software was deployed like the hardware probes in the sense that additional interfaces within the host Operating System facilitated connectivity tests within use case network segments ultimately producing the required LAN side network metrics.

Cellular

When considering cellular network level metrics, it must be remembered that the metrics reflect all activity on the network in the collection period. Therefore, distinguishing metric data that were created as a result of executing a use case from all other network related metrics data is not possible. To mitigate this use cases were planned such that the only network activity during the use case was generated by the use case itself. This created network metrics that were as close as could be to expected results from the use case. In busier, or mature, 5G networks, the network metrics created by the RAN will represent overall network performance and not metrics specific to a use case.

The above approach used to network metrics as close as would be expected to the results from the use case had a consequence that without other network traffic the capacity of the 5G network in a multi-use environment was not measured e.g. interruption as a result of other traffic and capacity limits of the network resulting from multiple users were not observed.

2.6. Testing

Throughout the entire deployment the metrics varied depending on what version of the RAN software was deployed. This affected stability, latency, and throughput. As the software evolved it was apparent that a stable baseline would be reached with the Q3 2021 release of the vRAN from Airspan Networks. Stability was achieved from the Radio Units whereby packet loss was consistently less than 1% in what could be considered as “close to optimal” RF conditions. Latency was observed as sub ~10ms and throughput was observed at ~410Mbps downlink and ~57Mbps uplink.

Testing was performed on both CPE and UE devices compatible with 5G SA connectivity. These values were consistently observed throughout testing for use case demonstrators throughout December 2021.

Note: Please refer to the Airspan 5G SA roadmap for 2022/23 as the performance of the system is significantly enhanced throughout this period (roadmap subject to change).

3. Conclusion

After the hardware platform was built and handed over to the Airspan Networks professional services team, it was necessary to modify and upgrade components within the configuration throughout the deployment to achieve a stable and performant solution. Intermittent issues were observed through the project that were resolved with software releases. These releases were installed quarterly throughout the project, which created delay in multiple use case delivery timescales. This highlighted that the RAN software was not initially production ready.

Performance and stability of the technology increased with each software release. The RAN vendors roadmap outlines incremental releases that a ~1GB downlink throughput would be achieved within the next 12 months. Uplink speeds remain at ~10% of the downlink throughput, although there is potential to modify the uplink and downlink throughput, which was discussed in the concluding phases. Latency is currently ~10ms and this is will also be improved to a lower value as part of the RAN vendor roadmap.

Note: Performance is dependent on radio conditions and, consequently, subject to environmental variation.

5G is still in the early adoption phase by most third-party hardware suppliers, limiting end user devices that are available. The uptake of the technology is expected to continue to grow and as devices mature the need for advanced configuration to force the hardware to connect to the technology be retired. For this project the challenge that every device is different and needed specific engineering or SIM configuration did add delays to the timescale for deliverables.

In conclusion, the software from the selected RAN vendor should perform at gigabit speeds during 2022 based on their roadmap. This will help to achieve what is expected downlink from a technical perspective within a 5G SA network. Uplink RAN performance will improve through 2022 as per the RAN vendor roadmap, however, the hardware limitations of the currently deployed radio units will constrain the uplink performance. The technology shows promise to far exceed what is currently available with its 4G derivative and will benefit from maturing over the next 12 to 24 months.

Appendix A – Open RAN Architecture

Appendix B – RAN LLD Provided by Consortium (Network Segments).

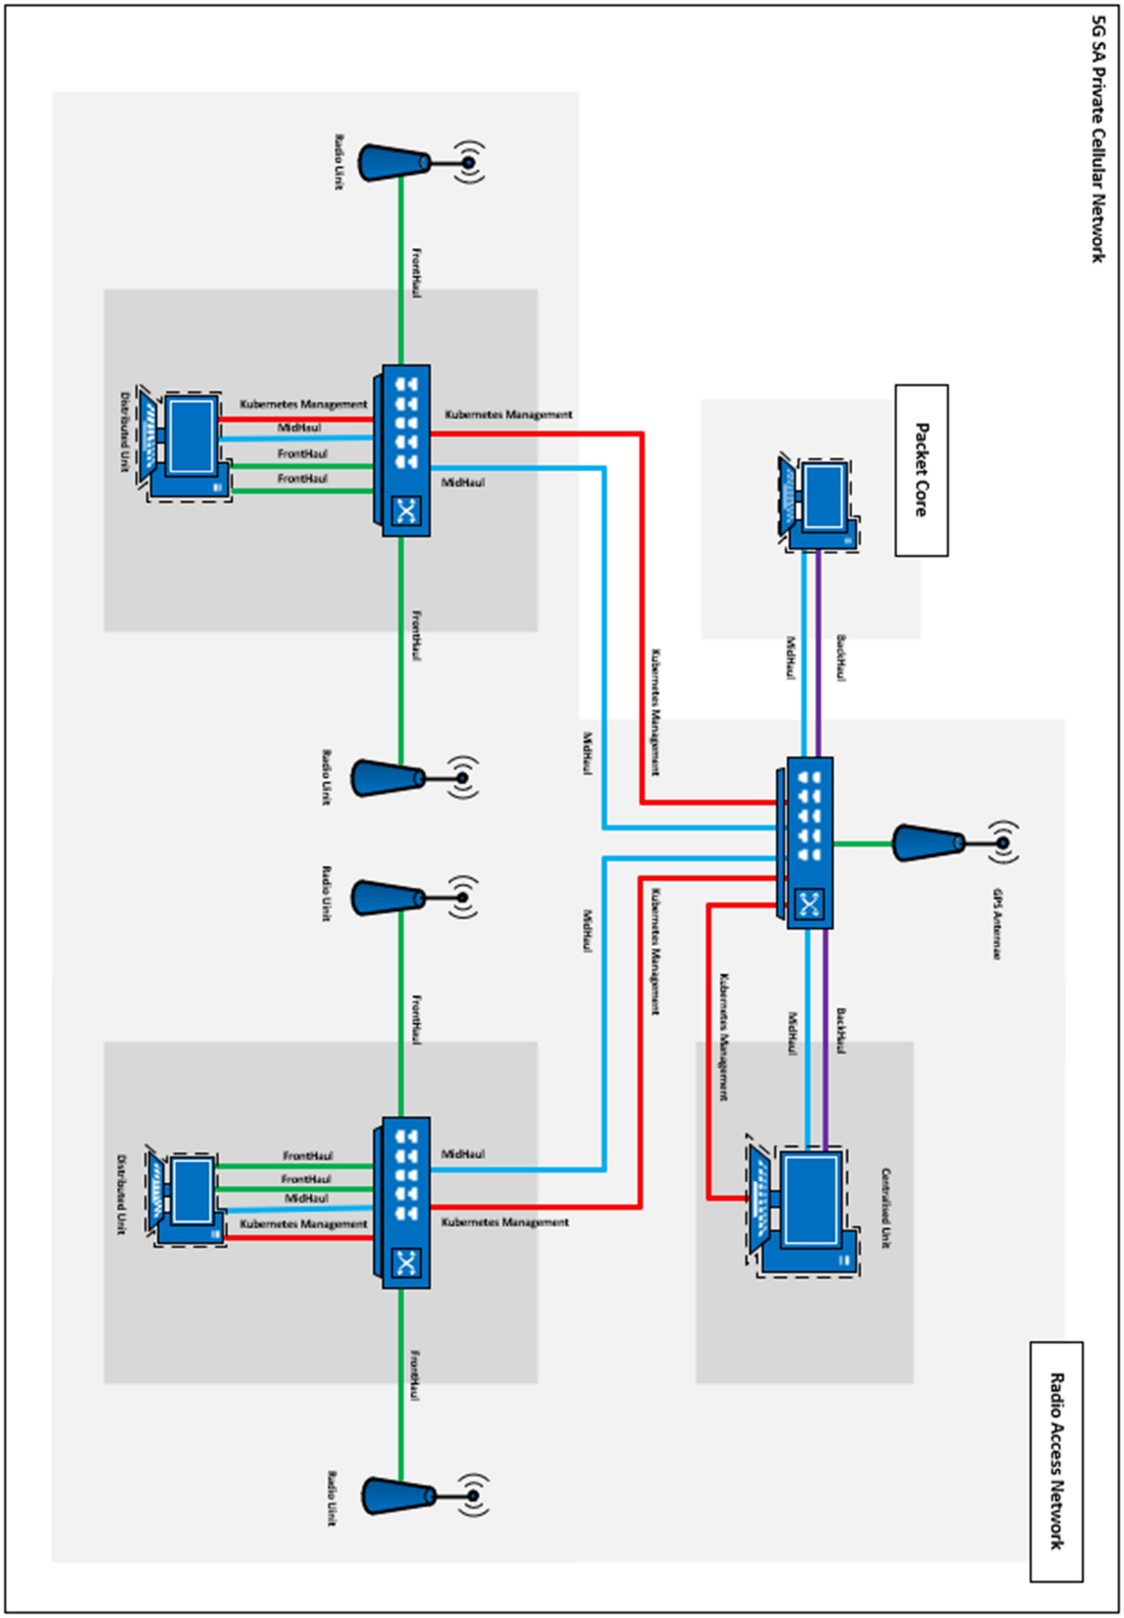

Appendix C – RAN LLD Provided by Consortium (VLANs & IP Addressing).

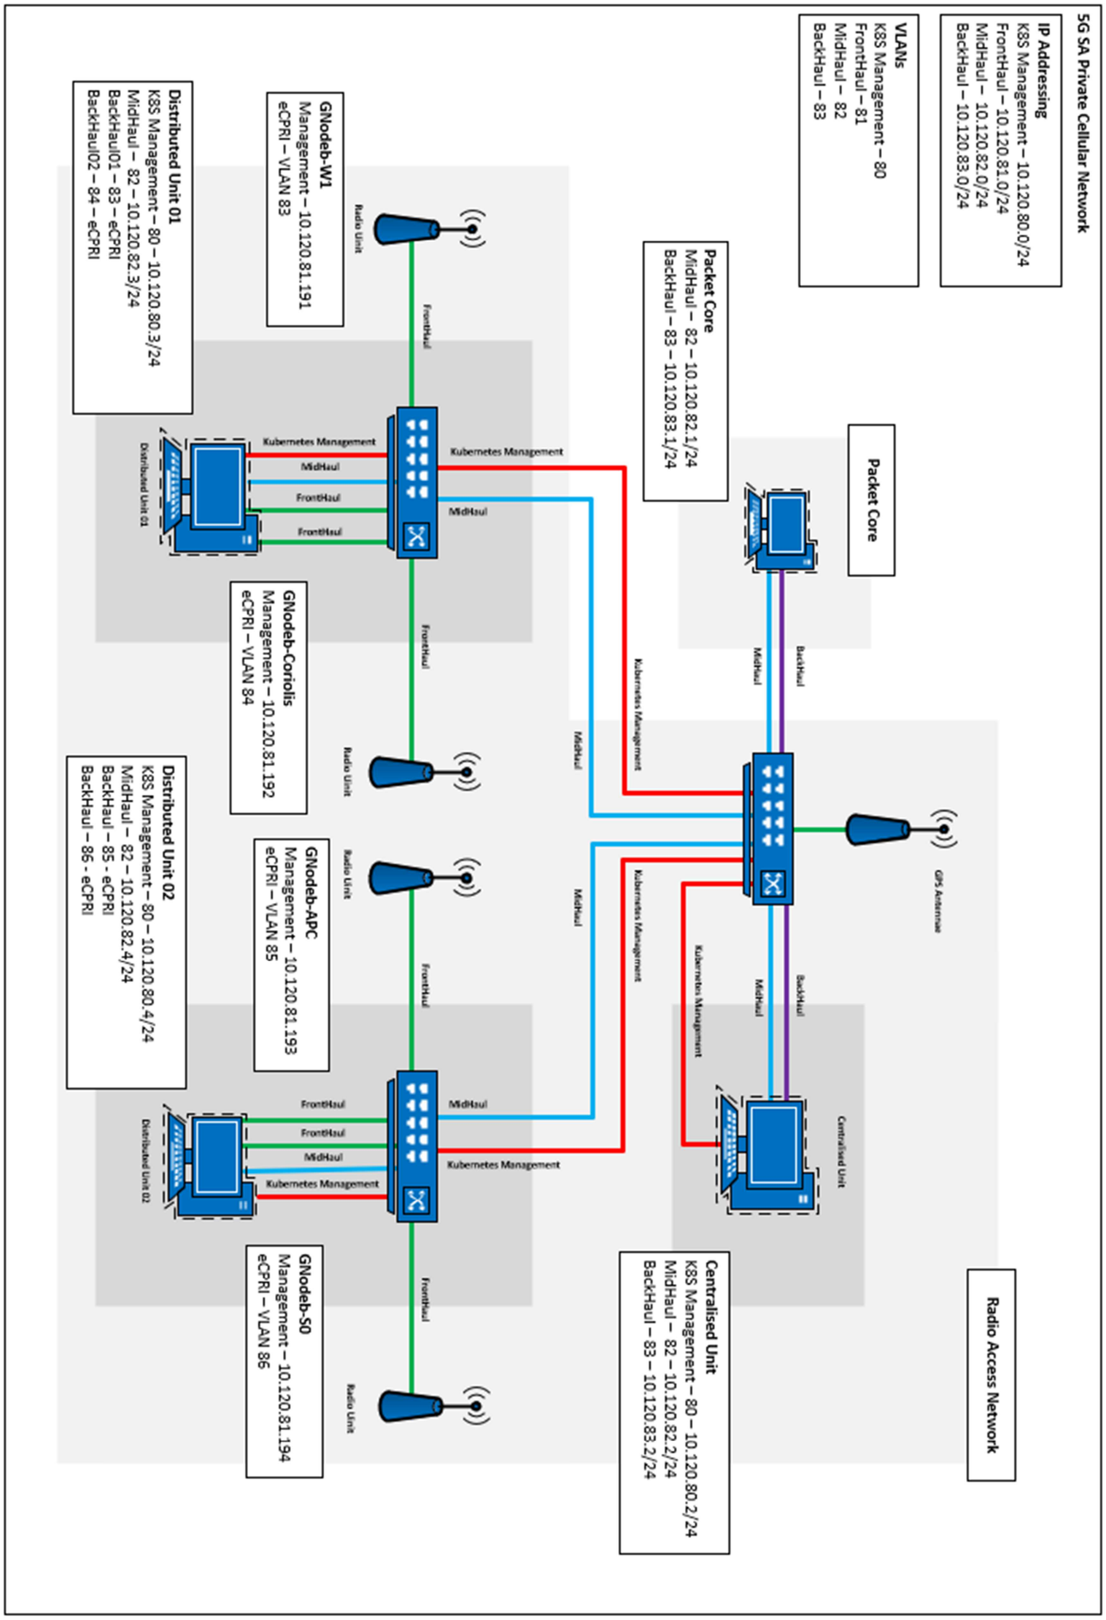

Appendix D – Physical Connectivity.

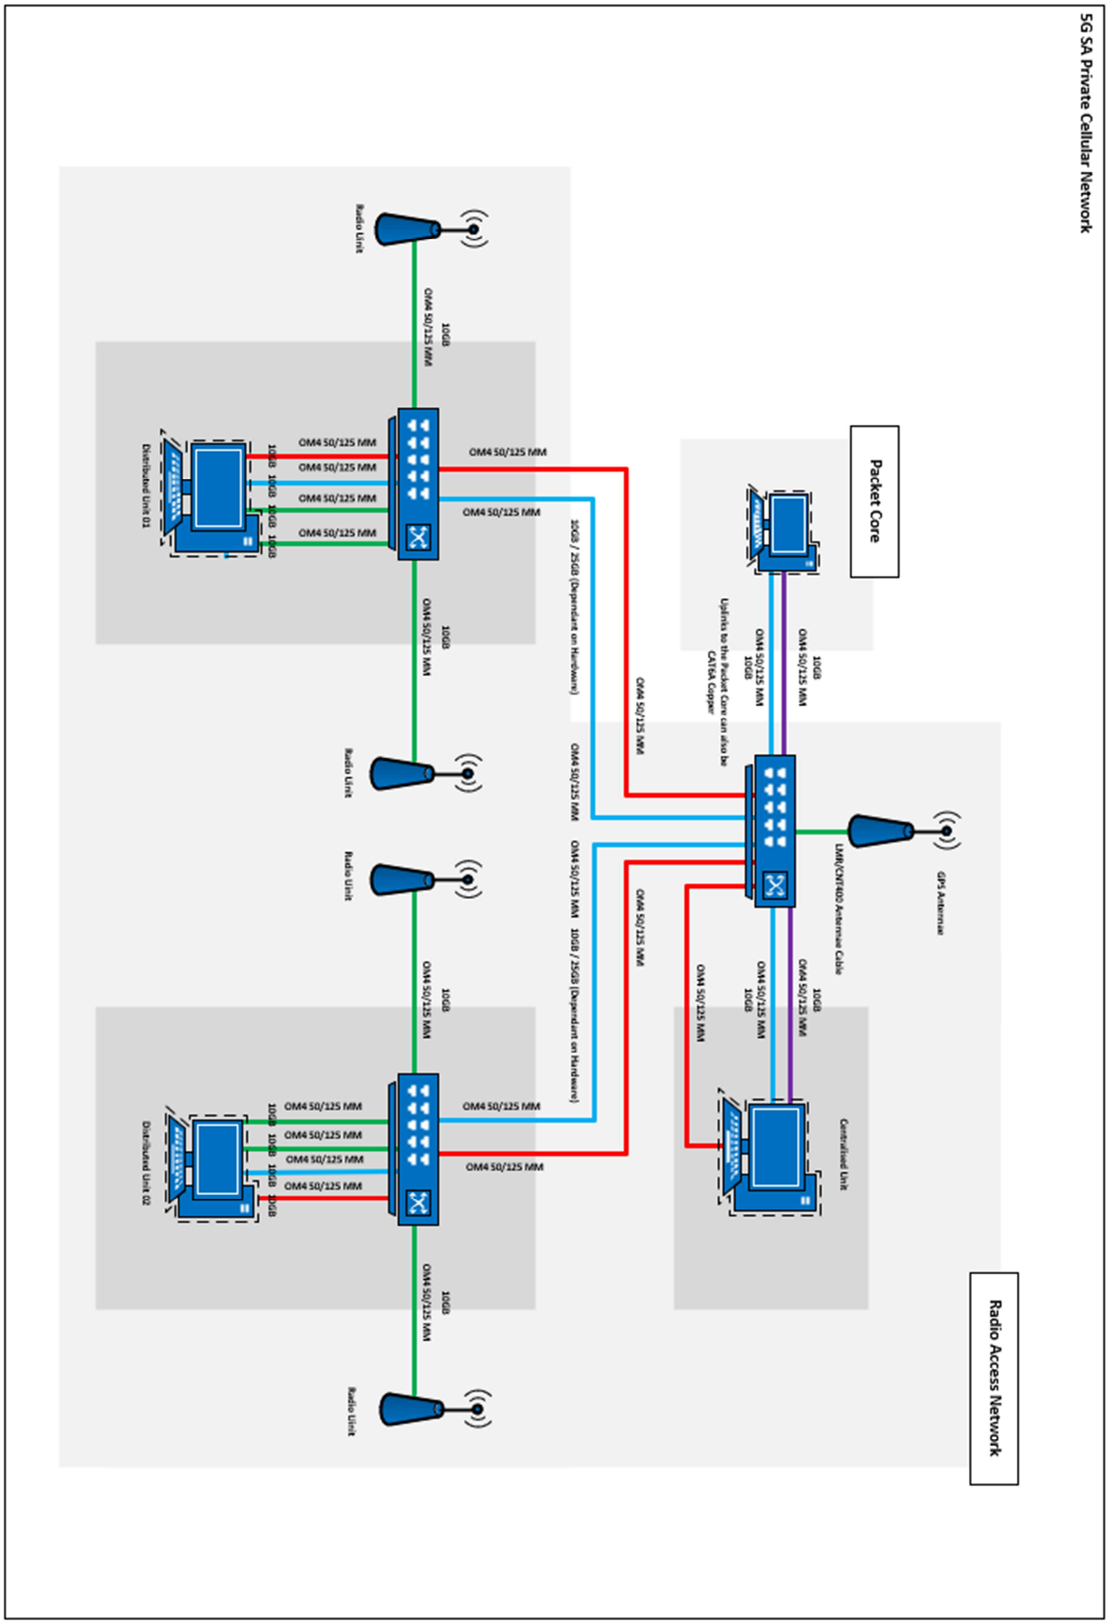

Appendix E – RAN with PTP Clock Integration.

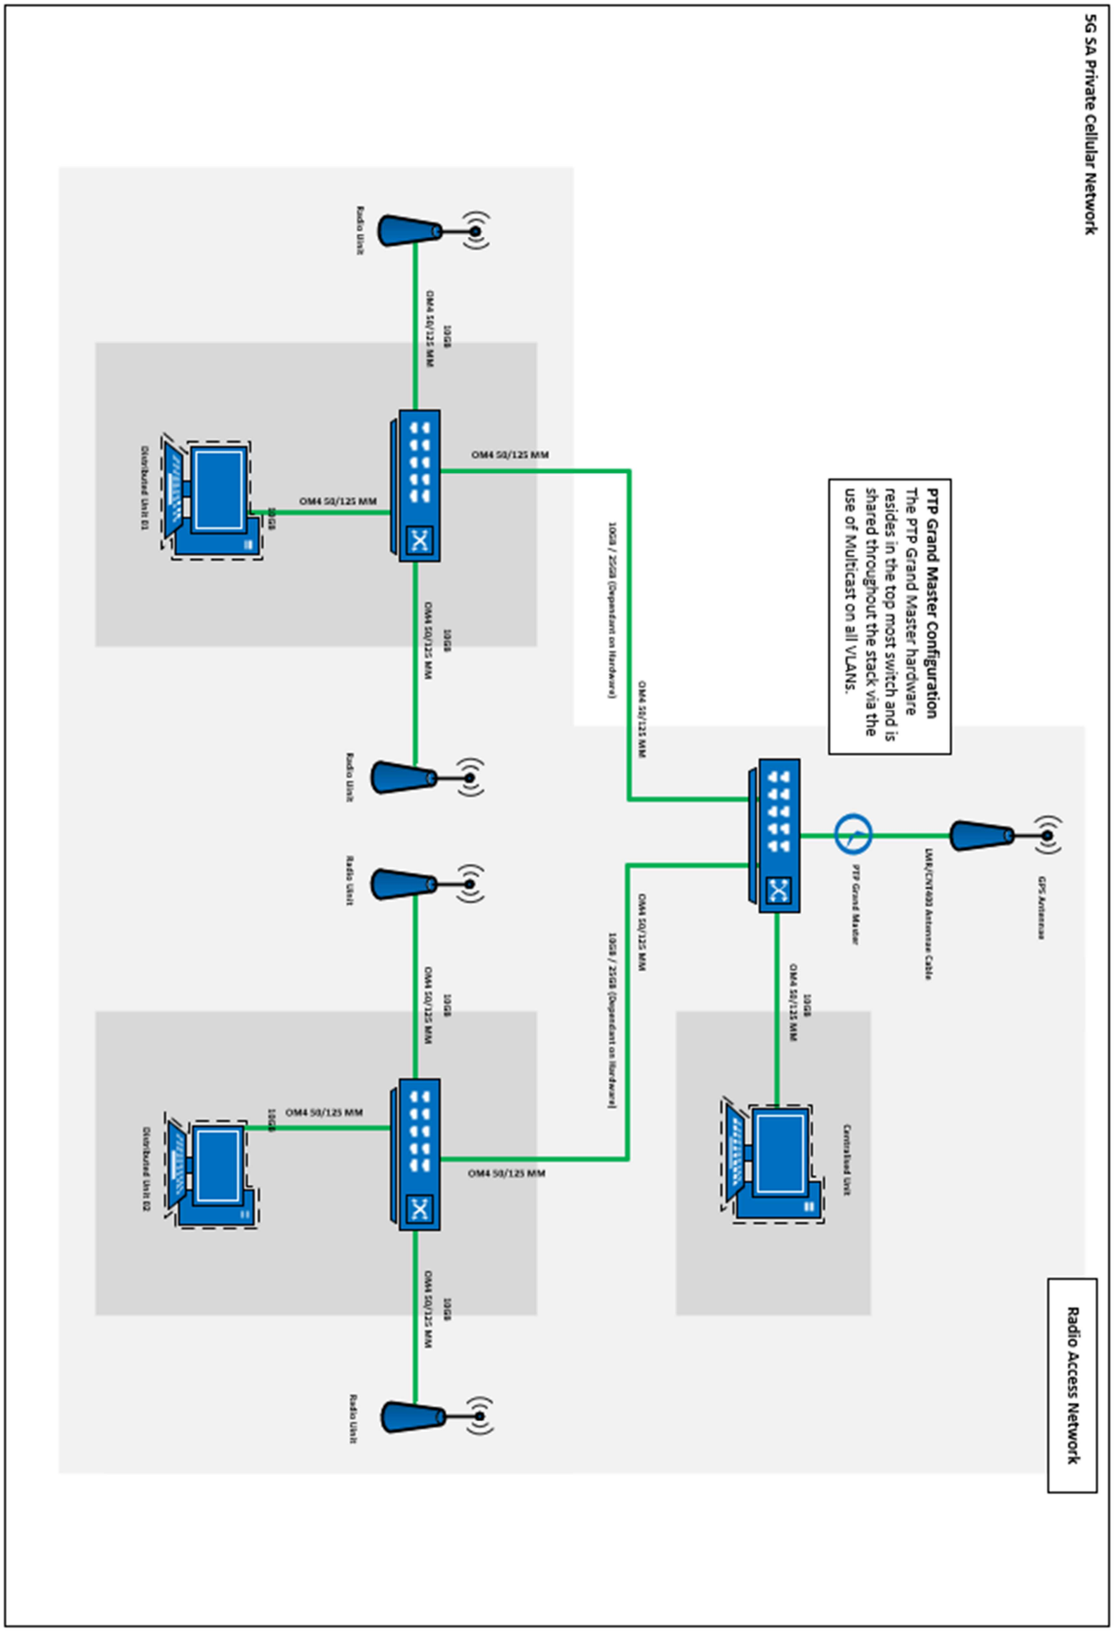

Appendix F – Packet Core Integration

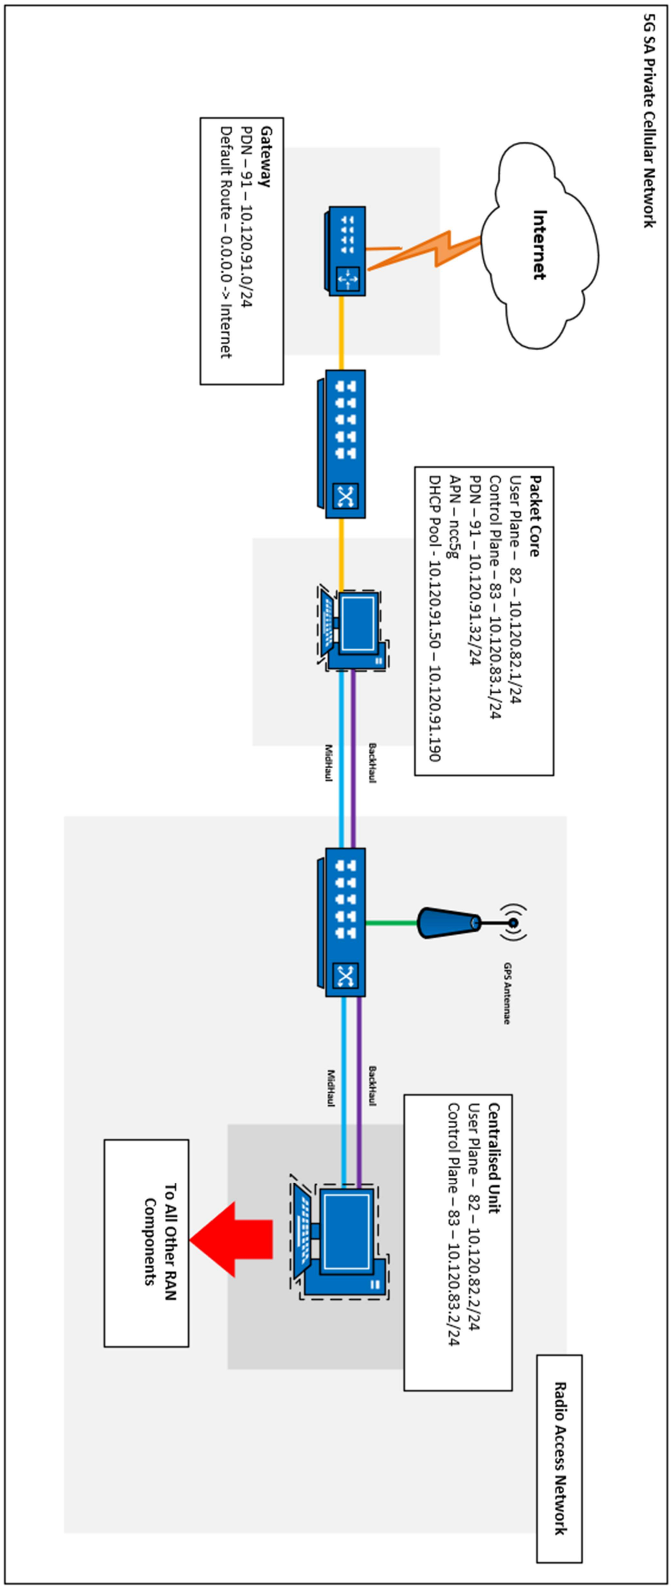

Appendix G – Kubernetes Cluster

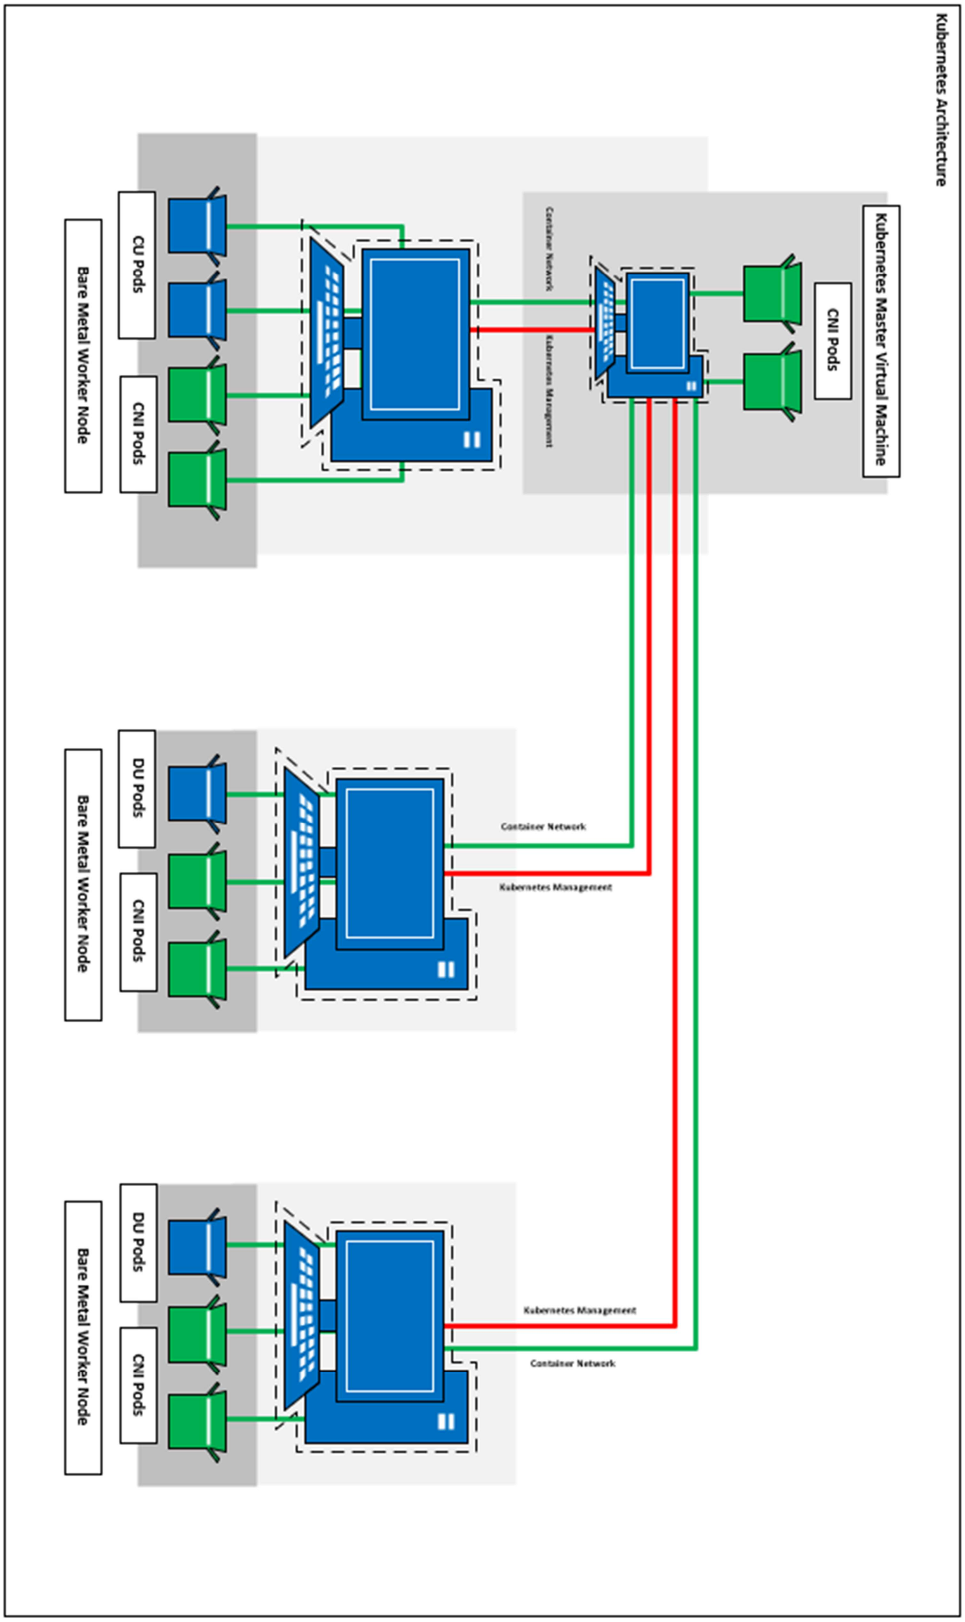

Appendix H – Kubernetes Architecture with ACP