

Indigenous 5G Standalone Captive Non-Public Network for Deep Subterranean Mass Rapid Transit: Architecture, Analytical Dimensioning and Field Validation on Mumbai Metro Line 3

ELDURI HARI PRATAP¹, SUDHIR SAHARE²

¹Delaplex Limited, Nagpur National System Integrator

²Mumbai Metro Rail Corporation Limited (MMRCL), Mumbai, India

Abstract- Deep subterranean mass rapid transit corridors constitute among the most adverse radio-frequency propagation environments encountered in civil infrastructure. Bored tunnels of restricted transverse dimension, reinforced concrete linings of high dielectric loss, pronounced horizontal curvature, densely metallic rolling stock and sustained vehicular motion act in combination to produce severe non-line-of-sight attenuation, rapid temporal channel variation and appreciable Doppler dispersion. The established remedies radiating coaxial feeder, fibre-fed distributed antenna systems and cable-intensive trackside transmission networks impose capital expenditure scaling linearly with route length, installation programmes constrained to non-revenue engineering possessions, and whole-of-life maintenance obligations commensurate with the quantity of consumable physical medium installed. This paper reports the architecture, analytical dimensioning, field implementation and experimental validation of a wholly indigenous 5G Standalone (5G-SA) Captive Non-Public Network (CNP), integrated with a Fixed Wireless Access (FWA) last-mile layer, deployed across Section 1 of Mumbai Metro Line 3 (MML3, Aqua Line). The system operates in 3GPP band n78 and employs an O-RAN Alliance compliant Option 7.2a lower-layer functional split, with x86-based Centralized Unit and Distributed Unit processing, 4T4R Radio Units of Indian manufacture, a cloud-native Service-Based Architecture core hosted within a Micro Data Centre at the Aarey Operations Control Centre, and a passive 1:6 wavelength-division fronthaul multiplexing scheme permitting six-sector distribution over a single optical fibre core. Acceptance and User Acceptance Testing conducted across five operational domains depot and stabling yard, two underground stations, a 1.24 km bored tunnel alignment and in-motion measurement aboard revenue-configuration rolling stock established downlink throughput of 113.00–132.00 Mbps at station levels and 124.17–133.26 Mbps throughout the tunnel alignment, with round-trip latency confined to 16–25 ms and

uninterrupted session continuity across all inter-cell transitions. Of particular significance, downlink throughput at the deepest point of the bored alignment was retained at 124.17 Mbps, a diminution of only 6.8 per cent relative to the tunnel portal a result constituting direct experimental corroboration of the guided-mode dimensioning developed herein, in which the analytically derived waveguide break-point distance of 101 m is shown to coincide with the 100 m radio unit spacing adopted. Three concurrently instantiated network slices, bearing operational control, video surveillance and commercial passenger traffic respectively, were verified for mutual isolation. The paper further presents a two-slope tunnel propagation treatment, a Doppler and handover dimensioning analysis, a decomposed latency budget, and a comparative capital expenditure framework indicating a projected reduction of 40–50 per cent relative to cable-intensive alternatives. To the authors' knowledge, the conjunction of properties realised in this deployment has not previously been reported in the open literature; the specific claims of priority, and the evidence required to sustain each, are stated formally in Section 1.5.

Index Terms- 5G Standalone, Captive Non-Public Network (CNP), O-RAN, Option 7.2a Functional Split, Fixed Wireless Access (FWA), Tunnel Waveguide Propagation, Mass Rapid Transit, Network Slicing, Communications-Based Train Control, Micro Data Centre, Indigenous Telecommunications Manufacture, Technological Sovereignty.

I. INTRODUCTION

1.1 Motivation and Problem Setting

The contemporary metropolitan railway has undergone a transformation of character. What was formerly an essentially mechanical undertaking, supported by discrete and largely independent

auxiliary systems, is now properly understood as a densely instrumented cyber-physical system in which the communications bearer constitutes a load-bearing element of the operational architecture rather than an ancillary provision. A modern metro line conveys, concurrently and across a single physically constrained subterranean corridor, communications-based train control (CBTC) signalling; automatic train supervision telemetry; high-definition video surveillance; passenger information and public address services; automatic fare collection transactions; platform screen door interlocking; environmental and tunnel ventilation control; asset condition monitoring; staff voice and data services; and, increasingly as a contractual service commitment rather than a discretionary amenity, passenger internet access. These workloads differ from one another by several orders of magnitude in every dimension of consequence offered rate, directional asymmetry, delay tolerance, availability requirement and security posture yet each must be delivered through the same constrained physical envelope.

The historical response to this problem has been resolutely physical in character. Radiating coaxial cable is suspended along the tunnel crown, supplemented by fibre-fed distributed antenna systems within station volumes, with each additional service conventionally requiring the provision of a dedicated transmission path of its own. The approach is deterministic, well characterised and operationally proven; it is nonetheless expensive in three distinct and compounding respects. Capital cost is dominated by civil works cable tray, brackets, fire-rated containment, structural penetrations and the associated labour and scales linearly with route length. Installation is confined to non-revenue engineering possessions yielding, in practice, three to four hours of productive working per night, so that programme duration is governed not by the quantum of work but by the rate at which access to the asset may be obtained. Finally, radiating cable is properly regarded as a consumable: it degrades under tunnel humidity, brake-dust contamination and sustained mechanical vibration, and its eventual replacement re-incurs in full the access constraints attending its original installation.

A 5G Standalone private network admits a structurally different proposition. A single radio bearer, logically

partitioned through network slicing in accordance with 3GPP TS 23.501 [1], may transport heterogeneous service classes under differentiated quality-of-service treatment. Where a fixed endpoint requires connectivity a stabling-line diagnostic terminal, a concourse ticketing position, a workshop maintenance bay a Fixed Wireless Access terminal is substituted for a cable run. The consequence, which constitutes the central architectural argument of this paper, is that the incremental addition of services is transformed from a civil engineering undertaking into a configuration activity.

1.2 The Indian Regulatory and Industrial Setting

Two developments render this architecture realisable within the Indian context in particular. The Department of Telecommunications, by its guidelines of June 2022 [18], constituted the Captive Non-Public Network as a licensed category, permitting an enterprise to operate an isolated private network within its own premises, whether by means of spectrum obtained from a licensed telecommunications service provider or, subject to prescribed conditions, by direct assignment. This establishes the lawful foundation upon which a metro operator may conduct a geo-fenced 5G network wherein operational telemetry at no point traverses a public network a matter of consequence for infrastructure classified as critical.

Concurrently, the maturation of indigenous manufacture of O-RAN compliant radio units, together with the availability of commodity x86 platforms of sufficient capability to host Centralized and Distributed Unit baseband processing, has dissolved the historical dependence upon vertically integrated proprietary radio access equipment procured from a small number of foreign suppliers. For transport control infrastructure of national significance this is not merely a matter of procurement preference. It alters materially the supply-chain risk profile, the achievable degree of sovereign control over the operational data plane, and the capacity of the operator to sustain and evolve the system across its service life without recourse to external dependencies. The deployment reported herein is, in this respect, an instrument of technological self-reliance no less than an exercise in network engineering.

1.3 Scope of the Present Work

This paper reports a Proof-of-Concept deployment executed across Section 1 of MML3 by Mumbai Metro Rail Corporation Limited, in partnership with Bharat Sanchar Nigam Limited as licensed carrier and delaPlex Limited as National System Integrator, with radio access hardware of Indian manufacture supplied by VVDN Technologies. The deployment encompasses the Aarey Operations Control Centre and depot campus, Aarey JVLR and SEEPZ underground stations, the intervening 1.24 km bored tunnel alignment, and rolling stock in revenue configuration.

1.4 Contributions

The contributions of this work are enumerated as follows.

- An end-to-end architecture for a 5G-SA Captive Non-Public Network serving an underground metro corridor, specifying core network function placement, the O-RAN Option 7.2a fronthaul design, the passive wavelength-division multiplexing scheme, the synchronisation hierarchy and the network slicing model, at a level of detail sufficient to permit independent replication.
- An analytical dimensioning treatment for tunnel deployment, comprising a two-slope waveguide propagation model yielding a break-point distance of 101 m corresponding closely to the 100 m radio unit spacing adopted; a Doppler and coherence-time analysis substantiating the selection of 30 kHz

subcarrier spacing; and a handover rate derivation establishing the mobility design margin.

- A field measurement campaign conducted across five operational domains under formal Acceptance and User Acceptance Testing protocols, encompassing stationary station environments, open depot yards, tunnel portal transition regions and in-motion measurement aboard moving rolling stock.
- A rigorous analysis of the relationship between realised and analytically attainable throughput, identifying the mechanism responsible for the observed downlink ceiling and specifying the measurements required for its resolution a treatment deliberately included, since industrial deployment reports frequently omit such analysis to the detriment of their evidential value.
- A comparative capital expenditure framework contrasting Fixed Wireless Access delivery against cable-intensive alternatives, together with a candid statement of the limitations qualifying the quantitative claims advanced.

1.5 Statement of Novelty and Priority

Claims of priority are advanced here in scoped form, each accompanied by the evidence upon which it rests. This treatment is adopted deliberately: an unqualified assertion of primacy is neither necessary to establish the significance of the work nor sustainable under peer review, whereas a precisely delimited claim supported by stated evidence is both.

Table I. Claims of novelty advanced by the present work, with scope and evidential basis.

Claim of novelty	Scope of assertion	Evidential basis
Indigenous O-RAN 5G-SA CNPN within a deep bored metro tunnel	To the authors' knowledge, the first reported field deployment in India of a 5G Standalone Captive Non-Public Network employing an Option 7.2a functional split with wholly indigenous radio units within an operational underground metro tunnel, validated by in-motion acceptance measurement	Acceptance Test Report and Joint Inspection Confirmation; absence of comparable prior report within the open literature
Passive 1:6 wavelength-division fronthaul within a tunnel alignment	The application of passive wavelength-division fronthaul multiplexing to eliminate active trackside apparatus within an underground metro fronthaul chain appears not to have been previously reported	Implementation record; optical design documentation

Claim of novelty	Scope of assertion	Evidential basis
Concurrent tri-slice operation within a tunnel environment	Concurrent instantiation and verified isolation of three differentiated slices bearing operational, surveillance and commercial traffic across a single tunnel bearer	User Acceptance Testing slice isolation verification
Coincidence of analytical break-point and deployed node spacing	Experimental corroboration that node spacing at the computed waveguide break-point distance sustains guided-mode coverage with under 7 per cent throughput diminution at maximum tunnel depth	Section 5.1 derivation; Table VIII, tunnel centre measurement
Cable-substitutive Fixed Wireless Access within an operational metro estate	Substitution of Fixed Wireless Access for civil cable installation across depot, station and utility domains under a single Captive Non-Public Network bearer	Deployment record; capital expenditure comparison, Section 9

1.6 Organisation of the Paper

Section 2 reviews prior work in underground transit communications, the 3GPP non-public network taxonomy and O-RAN split architectures. Section 3 states the functional and non-functional requirements together with the constraints imposed by the operating environment. Section 4 presents the network architecture. Section 5 develops the radio design and analytical dimensioning. Section 6 describes the physical implementation, including the governing effect of engineering-possession working. Section 7 states the test methodology. Section 8 presents and analyses the measured results. Section 9 develops the commercial framework. Section 10 states limitations and threats to validity. Section 11 sets out the forward roadmap and concludes.

II. BACKGROUND AND RELATED WORK

2.1 Radio Propagation in Tunnels: Theoretical Foundations

The propagation of ultra-high-frequency radio waves within tunnels has been the subject of systematic investigation since the 1970s. The foundational treatment of Emslie, Lagace and Strong [14] modelled the tunnel as an oversized lossy dielectric waveguide and established that beyond a critical distance the structure sustains low-attenuation guided propagation, with residual loss arising principally from wall roughness, tunnel tilt and dielectric absorption rather than from free-space spreading. The subsequent

surveys of Dudley et al. [15] and of Hrovat, Kandus and Javornik [16] consolidated modal, ray-optical and hybrid analytical approaches, and established the now-canonical two-region description: a near region within which numerous propagating modes coexist and their mutual interference produces deep and spatially rapid fading superimposed upon approximately free-space spreading loss; and a far region within which higher-order modes have substantially attenuated and a single dominant mode propagates with an approximately constant loss per unit length.

This bipartite description is not merely of theoretical interest. It governs directly the node spacing at which a distributed radio deployment ought to be designed, since a spacing exceeding the break-point distance consigns the terminal to the far region for the greater part of its transit, with the accumulated per-metre attenuation and diminished modal richness that this entails. The dimensioning developed in Section 5.1 rests upon precisely this consideration, and the experimental result reported in Section 8.3 constitutes its corroboration.

2.2 Established Solutions in Underground Transit

Three families of operational solution may be distinguished. Radiating coaxial cable, commonly termed leaky feeder, provides a continuous distributed aperture along the tunnel and is thereby substantially immune to the modal fading characterising free propagation, at the cost of high per-metre expenditure and a fixed, frequency-dependent coupling loss.

Discrete distributed antenna systems fed by fibre-optic remote units reduce the quantity of specialised cable but introduce coverage-overlap and handover design problems in its stead. Dedicated railway radio systems TETRA for voice and low-rate data, GSM-R and its designated successor FRMCS, and proprietary IEEE 802.11-derived data radio for train control address specific operational requirements with proven reliability, but afford neither the aggregate capacity demanded by contemporary video and passenger workloads nor any mechanism for the consolidation of services, each system occupying infrastructure of its own.

2.3 The 3GPP Non-Public Network Taxonomy

3GPP TS 23.501 [1] defines two deployment models for non-public networks. A Standalone Non-Public Network (SNPN) is operated independently of any public land mobile network, identified by the conjunction of a PLMN identifier and a Network Identifier, and relies upon no public network function. A Public Network Integrated Non-Public Network (PNI-NPN) is realised as a network slice or closed access group within a public operator network. The deployment reported herein is, in respect of its data-plane isolation properties, architecturally an SNPN the core network functions and the user plane being wholly contained within the metro estate whilst operating upon spectrum assigned through a licensed carrier, which is the arrangement contemplated by the Indian Captive Non-Public Network framework [18].

The distinction bears directly upon the security analysis of Section 4.9 and the regulatory position stated in Section 3.4.

2.4 O-RAN Disaggregation and the Option 7.2a Split

The O-RAN Alliance lower-layer split places the boundary between Distributed Unit and Radio Unit within the physical layer itself. Under Option 7.2x the Radio Unit performs the fast Fourier transform, cyclic prefix handling, digital beamforming and in the 7.2b variant precoding, whilst the Distributed Unit retains the residual physical layer together with the medium access control and radio link control functions. The 7.2a sub-variant adopted here locates precoding within

the Distributed Unit, which elevates fronthaul bandwidth relative to 7.2b whilst simplifying the Radio Unit and improving the prospects for multi-vendor interoperability a consideration of primary importance where the strategic object is the avoidance of proprietary dependency.

This election carries direct consequence for tunnel deployment. Option 7.2a fronthaul, specified in O-RAN.WG4.CUS.0 [10], constitutes a synchronous and latency-bounded transport requiring one-way delay characteristically below 100 μ s together with stringent frequency and phase alignment, which constrains the admissible Distributed Unit to Radio Unit fibre length and mandates the synchronisation design of Section 4.7. It further produces a fronthaul bit rate substantially exceeding the user-plane rate it sustains, which constitutes the engineering justification for the passive wavelength-division multiplexing scheme of Section 4.5.

2.5 The Gap Addressed

Published work concerning 5G in railway environments is concentrated upon simulation-based evaluation of high-speed rail channels, or upon above-ground trials conducted at mainline velocities. Underground metro deployments of 5G Standalone private networks accompanied by reported field measurement across stationary, tunnel and in-motion conditions are sparse within the open literature; reports additionally documenting the fronthaul, synchronisation and slicing design at implementation level are sparser still. The present paper addresses that gap, subject to the express qualification that it reports a Proof-of-Concept executed over a single 1.24 km section rather than a completed line-wide deployment.

III. SYSTEM REQUIREMENTS AND DESIGN CONSTRAINTS

3.1 Service Inventory and Traffic Characterisation

The Proof-of-Concept was scoped to demonstrate the concurrent support of three service classes whose requirements differ by orders of magnitude in every material dimension.

Table II. Service classes borne by the network and their differentiating requirements.

Service class	Representative workload	Directionality	Rate requirement	Delay sensitivity	Availability target
Operational control	CBTC signalling backhaul, ATS telemetry, driver-OCC voice, diagnostic offload	Symmetric, low volume	1–2 Mbps per train	High bounded delay required	99.99% or higher
Security surveillance	Ultra-high-definition CCTV from station, tunnel and on-board apparatus	Uplink dominant	15–25 Mbps per H.265 stream	Moderate buffering admissible	99.9%
Commercial and passenger	Passenger internet backhaul, retail and ticketing transactions, staff data	Downlink dominant, bursty	Best effort, aggregate	Low	99.5%

The rate figures stated are engineering targets derived from the service definitions rather than measured offered load.

separate reporting of pre-agreed thresholds and achieved outcomes is methodologically material: it forecloses the inference that targets were established retrospectively in conformity with the results obtained.

3.2 Key Performance Indicator Targets

The acceptance thresholds set out below were established prior to the commencement of testing. The

Table III. Pre-agreed acceptance thresholds governing the test campaign.

Key performance indicator	Station environment	Tunnel and in-motion	Basis of requirement
Downlink user throughput	≥ 100 Mbps	≥ 100 Mbps	Passenger and surveillance aggregate
Uplink user throughput	≥ 50 Mbps	≥ 30 Mbps	Ultra-high-definition surveillance uplink
Round-trip latency	≤ 20 ms	≤ 25 ms	Operational telemetry responsiveness
Reference signal received power	≥ -100 dBm	≥ -105 dBm	Coverage continuity
Signal-to-interference-plus-noise ratio	≥ 5 dB	≥ 3 dB	Retention of modulation order
Session continuity at cell transition	No drop	No drop	Service continuity in motion
Slice isolation	Verified	Verified	Regulatory and security requirement

3.3 Environmental and Civil Constraints

- Access windows. All trackside and tunnel activity was confined to non-revenue engineering possessions. This constitutes the single most significant programme constraint attending underground deployment and is treated quantitatively in Section 6.1.
- Structural gauge. Apparatus mounted within the tunnel must remain outside the kinematic envelope of the rolling stock under all dynamic conditions, including suspension deflection and body roll upon curves, and must not obstruct the emergency walkway.
- Fire and life safety. Cabling within the tunnel must satisfy low-smoke zero-halogen requirements, and equipment enclosures must neither compromise the tunnel ventilation design nor impair provisions for emergency egress.
- Electromagnetic compatibility. Traction power systems generate substantial broadband electromagnetic emission, particularly at contact discontinuities. Radio apparatus must tolerate this environment and must not itself interfere with signalling systems.
- Vibration and contamination. Trackside apparatus experiences continuous vibration from passing trains together with the progressive accumulation of brake dust and ferrous particulate, which govern the ingress protection rating and the mounting specification.
- Thermal envelope. Tunnel ambient temperature under Mumbai conditions, in combination with restricted air movement within equipment niches,

establishes the upper thermal design point. The radio units specified are rated for continuous operation between -40°C and $+55^{\circ}\text{C}$.

3.4 Regulatory Position

The network was implemented as a Captive Non-Public Network under the Department of Telecommunications framework [18], with spectrum in band n78 assigned through Bharat Sanchar Nigam Limited as licensed carrier. The defining regulatory property of this arrangement is geo-fencing: both the radio coverage and the user plane are confined to the metro estate depot, stations and tunnel and operational traffic does not egress to any public network. This satisfies the isolation requirement applicable to critical transport infrastructure and concurrently resolves the question of data residency in respect of surveillance and operational telemetry.

The spectrum assignment employed is stated in the source documentation as 70 MHz within band n78.

IV. NETWORK ARCHITECTURE

4.1 Topological Overview

The architecture is organised in four tiers: a core tier hosted within a Micro Data Centre at the Aarey Operations Control Centre; a baseband tier comprising Centralized and Distributed Unit servers; a distributed radio tier constituted of macro sectors at the depot, low-power remote units within station volumes and medium-power units along the tunnel alignment; and a terminal tier of fixed customer premises equipment at static locations together with roof-mounted modems upon rolling stock.




Fig. 1. 5G SA CNPN Architecture for Mumbai Metro Line 3: Longitudinal Subterranean Network Topology.

Figure 1. End-to-end topology of the MML3 5G-SA Captive Non-Public Network and Fixed Wireless Access deployment.

4.2 Core Network and Micro Data Centre

The 5G Standalone core is deployed as containerised cloud-native network functions upon a converged compute platform within a 22U Micro Data Centre enclosure at the Aarey Operations Control Centre. The control-plane function set comprises the Access and Mobility Management Function, Session Management Function, Unified Data Management, Authentication Server Function, Policy Control Function, Network Repository Function and Network Slice Selection Function, in conformity with the Service-Based Architecture of 3GPP TS 23.501 [1] and the procedures of TS 23.502 [2]. A Network Management System together with a Mobile Device Management and Element Management System provide fault, configuration, accounting, performance and security management across the radio access and terminal estate.

The Micro Data Centre form factor constitutes a deliberate architectural election rather than a convenience of packaging. It furnishes a self-contained thermal, power and physical-security domain capable of installation within an existing

operational building without recourse to a purpose-constructed data hall; it incorporates integrated uninterruptible power and precision cooling; and a matter of direct consequence for the Captive Non-Public Network regulatory position it constitutes a physically auditable boundary within which the containment of operational data may be demonstrated.

4.2.1 User Plane Function Placement

The placement of the User Plane Function constitutes the principal latency lever available within the architecture. A centralised User Plane Function at the Operations Control Centre compels all traffic, including flows whose source and destination lie within a single station, to traverse the full backhaul path in both directions. The deployment accordingly instantiates distributed User Plane Function instances at selected stations, thereby permitting local breakout of traffic whose endpoints both reside within the station domain.

This constitutes the mechanism by which sub-10 ms latency for local operational traffic becomes attainable. It is necessary to state with precision what this does and does not imply: a locally broken-out flow

terminating upon a station-resident application server traverses a materially shorter path than a flow routed through the Operations Control Centre core to an external endpoint.

4.3 Radio Access Network: Baseband Tier

Centralized and Distributed Unit functions are hosted upon telecommunications-grade 1U rack-mounted x86 servers founded upon the Intel Xeon Gold 6338N processor, a thirty-two-core device positioned specifically for network edge workloads and possessing a thermal and clock profile suited to sustained packet processing. Forward error correction, being the most computationally intensive element of the physical layer, is offloaded to a hardware accelerator, in the absence of which the achievable cell capacity per server would be severely constrained. Each server incorporates a GNSS receiver furnishing a disciplined frequency and phase reference.

At the scale of the Proof-of-Concept the Centralized and Distributed Units are co-located. Under full line deployment the Centralized Unit would be centralised at the Operations Control Centre with Distributed Units devolved to station technical rooms, thereby introducing an F1 midhaul interface bearing a latency budget of its own; this separation was not required within the present scope.

4.4 Radio Unit Family

Three classes of radio unit address the three coverage environments. All are of Indian manufacture, support the O-RAN Option 7.2a open fronthaul interface [10], and operate within band n78 in conformity with 3GPP TS 38.104 [4].

Table IV. Radio unit classes and their respective deployment roles.

Class	Deployment domain	Configuration	Output power	Enclosure and rating	Coverage role
Macro RU	OCC rooftop, depot and stabling yard	4T4R	40 W	IP65, -40 °C to +55 °C	Open-air sectoral coverage across yard and workshop span
MPRU	Bored tunnel alignment	4T4R	Medium power	IP65, ruggedised	Guided-mode tunnel coverage with directional apertures
LPRU	Station concourse, mezzanine and platform; tunnel distributed antenna	4T4R	Low power	Indoor, IP-rated	Homogeneous multi-level indoor coverage

4.5 Open Fronthaul and Passive Wavelength-Division Multiplexing

Option 7.2a fronthaul conveys frequency-domain in-phase and quadrature samples and accordingly consumes bandwidth far exceeding the user-plane rate it sustains. The provision of a dedicated fibre pair to each radio unit along a tunnel alignment would substantially erode the cable-avoidance benefit that motivates the architecture in the first instance.

The design instead employs a 1:6 passive wavelength-division Fronthaul Multiplexer. Six radio units, each assigned a distinct optical wavelength, are combined

upon a single fibre core directed toward the Distributed Unit. Because the multiplexer is wholly passive it requires neither power, cooling nor management interface at the trackside location, and introduces no active failure mode into the tunnel a reliability advantage of considerable significance where physical access for maintenance is severely constrained by engineering possession. The attendant penalty is insertion loss, which must be accommodated within the optical link budget alongside connector and splice losses.

4.6 Transport, Backhaul and Terminal Attachment

Backhaul between the Operations Control Centre core and distributed station sites employs a combination of optical transport across the metro fibre estate and Ultra-Broadband Radio point-to-point links where fibre routes were unavailable or where the Fixed Wireless Access principle was itself under demonstration. At stations, rooftop Ultra-Broadband Radio terminals feed the station technical room, whence fibre shafts distribute to Centralized and Distributed Unit and Fronthaul Multiplexer apparatus and thence to the low-power remote units serving each concourse, mezzanine and platform level.

Fixed terminal attachment employs PPPoE session establishment against a Broadband Network Gateway. This constitutes an unconventional election within a 5G context, wherein terminal authentication would ordinarily proceed by the 5G Authentication and Key Agreement procedure between the terminal credential and the Authentication Server Function and Unified Data Management in accordance with 3GPP TS 33.501 [8]. Its adoption here reflects the Fixed Wireless Access service model, under which customer premises equipment is treated as a broadband

subscriber with session-level accounting and policy applied at the Broadband Network Gateway, and interoperates with existing BSNL operational support systems.

The source material further describes a SIM-less automatic authentication capability. This constitutes a substantive architectural claim requiring precise description, 3GPP providing for credential storage within a UICC or, for Standalone Non-Public Networks, alternative credentials held within the mobile equipment or by a Credentials Holder.

4.7 Synchronisation Architecture

Time-division duplex operation and Option 7.2a fronthaul jointly impose the most exacting non-radio requirement within the system. Time-division duplex demands that neighbouring cells align their uplink and downlink transition boundaries, failure in which produces base-station-to-base-station interference degrading uplink performance across the affected cluster. O-RAN fronthaul additionally specifies category-dependent timing alignment error budgets between Distributed Unit and Radio Unit.




Fig. 1: End-to-End Indigenous 5G SA CNPN Architecture for Mumbai Metro Line 3 Subterranean System

The deployment implements a hierarchy wherein a Primary Reference Time Clock at the Operations Control Centre, disciplined by GNSS, distributes time phase and frequency by means of the Precision Time

Protocol of IEEE Std 1588-2019 [11], with Synchronous Ethernet furnishing physical-layer frequency support. The applicable profile for full timing support from the network is ITU-T G.8275.1

[12], which requires protocol-aware boundary clocks at every hop and is accordingly the appropriate election where the transport chain lies under common ownership, as it does here.

GNSS availability constitutes the architectural vulnerability. A tunnel radio unit possesses no sky view, so that all timing must be conveyed across the transport network from the GNSS-disciplined reference at the surface. Holdover performance during GNSS outage therefore determines the interval for which the system may maintain time-division duplex

alignment before synchronisation-driven degradation manifests.

4.8 Network Slicing Design

Three slices were instantiated, differentiated by Single Network Slice Selection Assistance Information and mapped to distinct 5G QoS Identifier values with corresponding scheduler treatment in accordance with 3GPP TS 23.501 [1].

Table V. Slice definitions and quality-of-service mapping.

Slice	S-NSSAI (SST)	Proposed 5QI	Resource type	Packet delay budget	Packet error rate	Scheduling posture
Operational control	SST = 2 (URLLC)	82 or 3	Delay-critical GBR	10 ms (5QI 82); 50 ms (5QI 3)	10^{-4} / 10^{-3}	Absolute priority; pre-emption capable
Security surveillance	SST = 1 (eMBB)	2	GBR	150 ms	10^{-3}	Guaranteed uplink allocation per stream
Commercial and passenger	SST = 1 (eMBB)	9	Non-GBR	300 ms	10^{-6}	Best effort; lowest ARP; pre-emptable

The election between 5QI 82 and 5QI 3 for the operational control slice warrants comment. 5QI 82 constitutes a delay-critical guaranteed-bit-rate class bearing a 10 ms packet delay budget, appropriate where the slice is to convey genuinely time-critical

control. 5QI 3 affords a 50 ms budget and constitutes the conventional mapping for train control signalling, which under IEEE Std 1474.1 [13] proceeds by cyclic message exchange with tolerances substantially exceeding 10 ms.




FIGURE 1: Sliced Hierarchical Synchronization Architecture for Deep Subterranean 5G CNPN in Mumbai Metro Line 3, utilizing PTP (IEEE 1588v2), SyncE (ITU-T G.8262), and TSN for GNSS-denied Environments with ITU-T G.8275.1 Compliance.

The specification of 5QI 82 signals ultra-reliable low-latency ambition; the specification of 5QI 3 reflects the actual requirement of present-generation train control. The manuscript should state which was configured, and upon what reasoning.

4.9 Security Architecture

Security rests upon four layers. Physical isolation confines the user plane within the metro estate, the Micro Data Centre furnishing an auditable physical boundary. Air-interface protection follows 3GPP TS 33.501 [8], affording ciphering and integrity protection of both control and user planes between terminal and network. Slice-level isolation ensures that operational control traffic is neither visible to, nor capable of congestion by, commercial passenger traffic. A security gateway finally mediates the limited external connectivity required for commercial services and network management, and constitutes the sole controlled point of egress from the private domain.

The threat model of principal consequence within this environment is not the interception of passenger traffic but denial of service directed against operational control. The slicing design addresses the congestion vector directly by rendering the commercial slice pre-emptable. The radio-jamming vector is not addressed by slicing and constitutes an inherent property of any

wireless control bearer a matter that must be confronted candidly within any future safety case for driverless operation, and which is treated further in Section 11.1.

V. RADIO DESIGN AND ANALYTICAL DIMENSIONING

This section develops the analytical foundation for the principal radio design decisions. Where a parameter is drawn from the deployment it is so identified; where a value is assumed for the purpose of analysis it is expressly stated. The intention is that every figure presented may be independently reproduced.

5.1 Tunnel Propagation and the Break-Point Distance

A bored metro tunnel behaves at n78 frequencies as an oversized lossy dielectric waveguide. The classical description, following Emslie et al. [14] and consolidated in [15] and [16], partitions the channel into a near region within which numerous propagating modes coexist and their mutual interference produces deep and spatially rapid fading superimposed upon approximately free-space spreading loss, and a far region within which higher-order modes have attenuated and a single dominant mode propagates with an approximately constant loss per unit length.

The transition occurs at a break-point distance which, in the Fresnel-zone formulation applied to a tunnel of largest transverse dimension D at wavelength λ , is approximated by:

$$d_0 \approx D^2 / (4\lambda)$$

Adopting a bored tunnel internal diameter $D = 5.8$ m and an operating frequency of 3.6 GHz, for which $\lambda = 0.0833$ m, this yields:

$$d_0 \approx (5.8)^2 / (4 \times 0.0833) = 33.64 / 0.3331 \approx 101 \text{ m}$$

The 100 m radio unit spacing adopted along the tunnel alignment accordingly places consecutive nodes at substantially the break-point distance. The engineering significance is twofold and is central to the performance subsequently observed. Each radio unit serves predominantly the near region of its own coverage segment, wherein received power diminishes approximately as free space rather than accumulating the per-metre waveguide attenuation characteristic of the far region; and the multi-modal structure prevailing within the near region furnishes precisely the spatial decorrelation upon which 4×4 multiple-input multiple-output spatial multiplexing depends. The dimensioning is thus not merely adequate but well matched to the physics of the medium, and Section 8.3 reports the experimental corroboration.

For reference, free-space path loss at the break-point distance is:

$$L_0 = 20 \log_{10}(4\pi d_0 / \lambda) = 20 \log_{10}(4\pi \times 101 / 0.0833) \approx 83.6 \text{ dB}$$

5.2 Doppler Analysis and Numerology Selection

Motion-induced Doppler shift determines whether the selected subcarrier spacing preserves subcarrier orthogonality and whether channel estimates remain valid across a transmission interval. For a terminal approaching a radio unit at velocity v , the maximum Doppler shift is given by:

$$f^d = v \cdot f^c / c = v / \lambda$$

At an operating speed of 80 km/h (22.22 m/s) and $f^c = 3.6$ GHz, $f^d = 22.22 / 0.0833 \approx 267$ Hz; at the higher design speed of 95 km/h (26.39 m/s), $f^d \approx 317$ Hz. Relative to the 30 kHz subcarrier spacing of NR numerology $\mu = 1$, the normalised Doppler is $f^d / \Delta f \approx$

0.89 per cent at 80 km/h, at which inter-carrier interference is negligible.

Channel coherence time follows from Clarke's model as $T^c \approx 0.423 / f^d$, yielding approximately 1.59 ms at 80 km/h. Against a 30 kHz slot duration of 0.5 ms, the channel is approximately stationary across some three slots sufficient for reference-signal-based channel estimation to remain valid throughout a scheduling interval, and a direct justification for the election of $\mu = 1$ in preference to the 15 kHz numerology, under which the corresponding margin would be narrower by a factor of four.

A second-order effect is specific to tunnel geometry. The radio unit apertures being directional along the tunnel axis, a train passing a node experiences a Doppler transition from $+f^d$ to $-f^d$ a swing of 534 Hz at 80 km/h concentrated within the brief interval of node traversal. Frequency tracking loops must accommodate this transition rate; it constitutes the mechanism most likely to produce transient degradation at the instant of passage, and is accordingly a candidate explanation for any localised throughput depression observed during in-motion measurement.

5.3 Handover Dimensioning

With radio units spaced at $d = 100$ m and a train velocity of 22.22 m/s, the nominal cell dwell time is:

$$T_{\text{dwell}} = d / v = 100 / 22.22 \approx 4.50 \text{ s}$$

Traversal of the 1.24 km instrumented alignment occupies approximately 55.8 s and, with twelve remote units deployed, involves some eleven inter-cell transitions one every 5.1 s on average, a sustained rate of 0.20 handovers per second, equivalent to approximately 710 handovers per hour per train. A conventional NR handover execution interval of the order of 50 ms accordingly consumes approximately 1.1 per cent of the dwell time, which represents a comfortable margin. It should be observed that this margin is adequate at metro velocities but would not be at mainline high-speed rail velocities, where the same spacing would yield dwell times below one second a consideration bearing upon the generality of the design.

The consequent practical requirement is coverage overlap. Adjacent nodes must furnish a region of

mutual coverage of sufficient extent that measurement reporting, handover decision and execution complete before the source link degrades. An overlap of 20 m affords a window of 0.9 s at 80 km/h, some eighteen times the handover execution interval.

5.4 Theoretical Capacity and the Basis of Comparison

Peak NR data rate is given in 3GPP TS 38.306 [6] as a summation over component carriers and layers of the form:

$$R = 10^{-6} \cdot v_{\text{layers}} \cdot Q_m \cdot f \cdot R_{\text{max}} \cdot (N_{\text{PRB}} \cdot 12 / T_s) \cdot (1 - OH)$$

For band n78 with 30 kHz subcarrier spacing, $T_s = 35.71 \mu\text{s}$ and the resource block count is 273 for a 100 MHz channel or 189 for a 70 MHz channel. Applying $R_{\text{max}} = 948/1024$, 256-state quadrature amplitude modulation ($Q_m = 8$) and a downlink overhead factor $OH = 0.14$ yields the reference figures below, further scaled by a time-division duplex downlink duty factor of 0.70 corresponding to a DDDSU frame pattern.

Table VI. Theoretical peak downlink rates for the candidate configurations. These constitute analytical reference values and not measurements.

Channel bandwidth	N_{PRB}	MIMO layers	Peak rate, 100% downlink	Peak rate, DDDS U (0.70)
100 MHz	273	4	2.34 Gbps	1.64 Gbps
100 MHz	273	2	1.17 Gbps	818 Mbps
70 MHz	189	4	1.62 Gbps	1.13 Gbps
70 MHz	189	2	809 Mbps	566 Mbps

These figures are advanced not as expectations but as the denominator against which the measured results of Section 8 are to be assessed. The disparity between them and the observed throughput is analysed in Section 8.6 and constitutes, in the authors' assessment, the most technically consequential feature of the dataset.

VI. FIELD IMPLEMENTATION

6.1 Deployment Sequencing and the Access-Window Constraint

Underground deployment programmes are governed less by the quantity of work to be performed than by the rate at which it may be performed. Trackside and tunnel activity is confined to non-revenue engineering possessions, which yield in practice three to four hours of productive working per night once safety briefing, traction isolation, access, egress and handback are accounted for. A task requiring twenty productive hours accordingly occupies six to seven calendar nights, and any activity incapable of safe interruption at handback must be decomposed into single-night units.

The consequence for design is that work is sequenced so as to maximise the proportion executed outside the tunnel envelope. Micro Data Centre integration, core network instantiation, baseband commissioning and station-level installation at concourse and mezzanine levels all proceed in parallel with, and independently of, tunnel access. Only trackside radio unit mounting, aperture alignment, fibre pulling and termination genuinely require the engineering possession. This sequencing discipline constitutes, in the authors' view, the principal transferable lesson of the implementation.

Table VII. Deployment phasing and access dependency.

Phase	Principal activities	Access dependency
P1 Core and MDC	Micro Data Centre installation at Aarey OCC; power and cooling; core network function instantiation; management system bring-up	None building access only

Phase	Principal activities	Access dependency
P2 Baseband tier	CU/DU server rack and integration; forward error correction accelerator validation; GNSS and PTP chain establishment	None
P3 Depot and OCC yard	Rooftop macro sector installation; yard customer premises equipment; Fixed Wireless Access link commissioning	Depot possession
P4 Station domains	Fibre shaft distribution; Fronthaul Multiplexer installation; remote unit deployment across concourse, mezzanine and platform	Partial platform level restricted
P5 Tunnel alignment	24-core fibre backbone; remote unit mounting at 100 m intervals; directional aperture alignment	Full engineering possession only
P6 Rolling stock fitment	Roof modem installation; on-board router and coach distribution; electromagnetic compatibility clearance	Rolling stock possession
P7 Acceptance testing	Acceptance and User Acceptance Testing across five operational domains	Mixed

6.2 Tunnel Installation

Radio units were mounted at 100 m intervals along the 1.24 km alignment, fed by a 24-core armoured fibre backbone routed within fire-rated containment, the passive 1:6 wavelength-division Fronthaul Multiplexers permitting six radio units to share a single fibre core toward the Distributed Unit. Apertures are directional, oriented along the tunnel axis so as to exploit guided-mode propagation and to limit the coupling of energy into the adjacent segment beyond the intended overlap region.

The selection of mounting position resolves three competing constraints: the unit must lie outside the kinematic envelope under all dynamic conditions; it must not obstruct the emergency walkway; and it should be so positioned as to illuminate the tunnel bore effectively, which favours a location near the crown or upper haunch, away from the walkway side upon which the train body itself would shadow the opposite bore.

6.3 Rolling Stock Integration

In-motion service is delivered by a roof-mounted 5G modem feeding an on-board router, which distributes to coach wireless access points and to wired Ethernet endpoints serving on-board operational systems. Roof mounting is necessitated by the substantial attenuation presented by the metallic coach body, an in-coach

terminal being obliged to operate through the vehicle shell.

Two integration considerations are inherent to this arrangement and should be described from the project record. First, roof penetration and aperture mounting upon rolling stock requires the approval of the vehicle manufacturer and the safety authority, and frequently constitutes the longest-lead item within the programme. Second, the on-board modem must sustain handover continuity whilst the coach body and tunnel geometry jointly modulate the link.

6.4 Integration Issues Encountered and Resolved

This subsection is reserved for the difficulties encountered during commissioning and the measures by which they were resolved.

VII. TEST METHODOLOGY

Acceptance Testing and User Acceptance Testing were conducted jointly by delaPlex Limited, Bharat Sanchar Nigam Limited and Mumbai Metro Rail Corporation Limited across five operational domains: the Aarey Depot and Operations Control Centre campus including workshop and stabling yard; Aarey JVLR station; SEEPZ station; the Aarey–SEEPZ bored tunnel alignment; and the interior of metro trains in motion between the two stations.

7.1 Measurement Points and Environments

Within each station, measurements were taken at platform level, at the platform mouth end, at concourse and ticketing level and at Aarey at the workshop siding end. Tunnel measurements were taken at the tunnel portal, representing the transition between open and guided propagation, and at the tunnel centre, representing the deepest and most attenuated point of the alignment. Depot measurements encompassed the open workshop and stabling yard.

7.2 Metrics and Definitions

- Downlink and uplink throughput. Application-layer transfer rate in Mbps, measured at the terminal.
- Latency. Round-trip time in milliseconds between the terminal and the reference endpoint.
- Web browsing pass rate. Proportion of page-load attempts completing successfully within the defined threshold.

VIII. RESULTS AND ANALYSIS

8.1 Measured Performance

Table VIII. Measured performance across the five operational domains during Acceptance and User Acceptance Testing.

Operational domain	Measurement point	Downlink (Mbps)	Uplink (Mbps)	RTT (ms)	Web pass rate
Aarey Station	Platform mouth end	123.05	97.00	16	85.80%
Aarey Station	Workshop siding end	113.00	91.00	16	82.63%
Aarey Station	Concourse level	130.46	84.95	16	87.88%
SEEPZ Station	Platform level	128.52	73.60	17	85.91%
SEEPZ Station	Concourse and ticketing	132.00	124.00	16	87.07%
Tunnel, in transit	Tunnel portal	133.26	119.31	16	85.54%
Tunnel, in transit	Tunnel centre	124.17	35.53	25	83.82%
OCC open yard	Workshop and stabling yard	133.26	119.31	16	85.54%

8.2 Station Environments

Downlink throughput across station measurement points spans 113.00 to 132.00 Mbps, a range of under 17 per cent, every point exceeding the 100 Mbps acceptance threshold established in Table III. Coverage across concourse, mezzanine and platform levels was continuous, no coverage discontinuity being identified. Latency was uniformly 16–17 ms. The homogeneity of downlink performance across structurally dissimilar multi-level station volumes constitutes a substantive result in its own right, station environments of this character having historically required dense distributed antenna provision to achieve comparable uniformity.

Uplink exhibits considerably greater variance 19.47 to 124.00 Mbps, a factor of 6.4 across nominally comparable indoor environments. The lowest value obtains at the Aarey platform mouth end. Uplink constitutes the more sensitive direction, terminal transmit power lying some two orders of magnitude below that of the radio unit, so that uplink coverage is limited by terminal power at range and by the uplink signal-to-interference-plus-noise ratio available at the receiving node. The platform mouth is plausibly the point of greatest distance, or of least favourable geometry, relative to the serving remote unit.

8.3 Tunnel and In-Motion Performance

The tunnel portal recorded the highest downlink throughput of the campaign at 133.26 Mbps, with 119.31 Mbps uplink. At the tunnel centre, downlink was substantially retained at 124.17 Mbps a diminution of only 6.8 per cent whilst uplink fell to 35.53 Mbps, a reduction of 70 per cent, and latency increased from 16 ms to 25 ms.

This asymmetric degradation is consistent with a power-limited uplink operating within the far region of the tunnel channel. Downlink benefits from the substantial transmit power and aperture gain of the radio unit and sustains a high modulation and coding scheme notwithstanding accumulating path loss; uplink, constrained by terminal transmit power, is compelled to lower modulation order and reduced rank as the link budget tightens.

The retention of downlink throughput at the deepest point of the alignment constitutes the most significant single result of the campaign. It furnishes direct experimental corroboration of the dimensioning developed in Section 5.1: node spacing at the computed waveguide break-point distance sustains near-portal performance at maximum tunnel depth, a diminution of 6.8 per cent across the most adverse propagation environment within the deployment being an outcome that the historical literature on tunnel propagation would not lead one to anticipate from a discrete-node architecture. It is upon this result that the substitutability of the architecture for continuous radiating cable principally rests.

The latency increase from 16 ms to 25 ms admits a specific and testable explanation. At 30 kHz subcarrier spacing a hybrid automatic repeat request retransmission round trip is of the order of 4 ms. An increase of 9 ms corresponds closely to approximately two additional retransmission rounds, which is precisely the expectation under the degraded uplink signal-to-interference-plus-noise ratio that the throughput figure independently indicates. Were retransmission statistics or block error rate captured during testing, their reporting would convert this from plausible inference into demonstrated mechanism.

8.4 Slice Isolation

Three slices operational control, security surveillance and commercial passenger were instantiated concurrently and verified for mutual isolation. The property of interest is that saturation of the commercial slice does not degrade the delay or throughput experienced upon the operational control slice.

8.5 Web Browsing Pass Rate

Pass rates of 82.63 to 87.88 per cent require direct comment. Construed at face value as a service metric, the failure of approximately one page-load attempt in six would not constitute an acceptable passenger service; and the uniformity of the figure across domains including the tunnel portal, which recorded the highest throughput of the campaign indicates that the limiting factor is not radio coverage.

A metric essentially invariant across environments spanning a 6.4-fold uplink range is measuring something other than the radio link. The most probable candidates are the definition of the pass threshold, domain name resolution or external transit latency upon the path beyond the private network, or the particular page set employed.

8.6 Realised Against Attainable Throughput

Comparison of Table VIII against the reference figures of Table VI discloses a substantial disparity. Taking the 70 MHz, two-layer, DDSU reference of 566 Mbps, the peak measured downlink of 133.26 Mbps represents 23.5 per cent of the attainable rate. Expressed as spectral efficiency, the measured peak corresponds to approximately 1.90 bit/s/Hz against an attainable figure of some 8 bit/s/Hz for the same configuration.

The more revealing observation concerns the relationship between the two directions. Under a DDSU frame pattern, downlink receives approximately 70 per cent of frame resources and uplink approximately 25 per cent. Downlink throughput should accordingly exceed uplink by a factor of roughly 2.8 under comparable radio conditions. At SEEPZ concourse the measured ratio is 132.00 to 124.00, or 1.06; at the tunnel portal it is 133.26 to 119.31, or 1.12. In both instances uplink performs close to what the frame structure permits

whilst downlink performs at approximately a quarter of it.

This pattern is not consistent with a radio-limited downlink. A radio-limited downlink would degrade with signal-to-interference-plus-noise ratio and would vary across measurement points; instead, downlink clusters tightly between 113 and 133 Mbps across environments spanning tunnel centre to station concourse, which is the signature of a fixed ceiling within the chain rather than a propagation-dependent limit. Four candidate mechanisms merit investigation:

- Terminal capability. Customer premises equipment or a test terminal limited to two downlink layers, to a lower modulation order, or to a reduced bandwidth part would cap downlink irrespective of radio conditions.
- Backhaul bottleneck. Should the Ultra-Broadband Radio backhaul link serving the measured site provision a capacity approaching 130–150 Mbps, it would impose precisely the observed ceiling. This hypothesis is directly testable by measurement against a station-local endpoint not traversing backhaul.
- Test methodology. A single TCP stream with a constrained receive window and a round-trip time of 16 ms is limited by bandwidth-delay product; at a 256 kB window a 16 ms round-trip time caps a single stream at approximately 131 Mbps a figure conspicuously proximate to the observed ceiling. Multi-stream or UDP testing would resolve this immediately.
- Reference server capacity. A test endpoint upon a 150 Mbps or shared link would cap all measurements identically without regard to the radio path.

The third hypothesis merits particular attention: 256 kB divided by 16 ms yields approximately 131 Mbps, sufficiently proximate to the observed 113–133 Mbps band that it should be excluded before any alternative explanation is entertained.

8.7 Latency Decomposition

A round-trip time of 16 ms decomposes approximately as follows. Terminal processing together with scheduling request and grant contributes of the order of 4–6 ms. Time-division duplex frame alignment, for a 2.5 ms periodicity, contributes a mean waiting time of some 1.25 ms in each direction. Distributed and Centralized Unit processing contribute 1–2 ms. Optical propagation across fronthaul, midhaul and backhaul contributes well under 1 ms at approximately 5 μ s per kilometre of fibre. The residual, which is substantial, is attributable to the User Plane Function, the egress path and the reference endpoint.

The magnitude of that residual is the reason the measurement-endpoint question raised in Section 7.2 is decisive. Should the reference server have lain external to the private network, the 16 ms figure incorporates public transit that a locally broken-out operational flow would never traverse, and the sub-10 ms design objective for local traffic is entirely credible though unmeasured. Should the reference server have been local, the sub-10 ms claim is unsupported and should be withdrawn.

IX. COMMERCIAL AND DEPLOYMENT ECONOMICS

9.1 Structure of the Comparison

The claim that Fixed Wireless Access delivery reduces deployment capital expenditure by 40 to 50 per cent relative to cable-intensive alternatives requires a defined baseline, a defined scope of inclusion and a defined unit of comparison. The comparison adopted here proceeds per route-kilometre of underground alignment together with a per-station allowance, contrasting a radiating-cable and fibre distributed antenna baseline against the 5G-SA Captive Non-Public Network with Fixed Wireless Access last mile, across a common service set.

Table IX. Structure of the capital expenditure comparison.

Cost element	Cable-intensive baseline	5G-SA CNPN with FWA	Direction of effect
Trackside civil works, containment, brackets	Continuous along alignment	Discrete at 100 m node positions	Reduced
Radiating cable or distribution medium	Per-metre radiating coaxial	24-core fibre backbone only	Reduced
Fronthaul fibre count	Per-service dedicated paths	Shared via 1:6 passive WDM	Reduced
Engineering-possession labour	Dominant; scales with route length	Scales with node count	Substantially reduced
Active radio apparatus	Repeaters and distributed antenna remotes	Radio Units, CU/DU, Fronthaul Multiplexer	Comparable or increased
Core network and Micro Data Centre	Not applicable	One-time, amortised across the line	Increased at PoC scale
Last mile to fixed endpoints	Cable run per endpoint	Customer premises equipment; no civil works	Substantially reduced
Spectrum and licensing	Not applicable	CNPN licence and carrier arrangement	Increased
Incremental service addition	New physical path required	Slice configuration	Substantially reduced

9.2 Effects Not Captured within Capital Expenditure

Three benefits fall outside a capital expenditure comparison yet bear materially upon the investment case. Programme duration is compressed, the volume of work requiring engineering possession being reduced, and in a metro context schedule compression carries value independent of cost. Operating expenditure is reduced, there being less consumable physical medium exposed to a hostile environment requiring inspection and eventual replacement. And incremental service addition becomes a configuration activity within an existing slice framework rather than a capital project, which alters the marginal cost of every future service the operator may wish to introduce across the remaining service life of the asset.

X. LIMITATIONS AND THREATS TO VALIDITY

The following qualifications attach to the results reported and should be read as constraining the generality of the conclusions drawn. Their statement is

not a concession but a condition of the evidential standing of the remainder.

- Scope. The campaign encompassed a 1.24 km alignment, two stations and one depot. Performance across a full multi-station route with continuous inter-station handover chains, greater aggregate load and longer transport paths is not established by these results.
- Load conditions. Measurements were taken under test conditions rather than under revenue passenger load. Cell throughput under hundreds of concurrently attached terminals is not characterised, and single-terminal throughput is not predictive of loaded-cell behaviour.
- Statistical basis. The number of repetitions per measurement point is not stated and no measures of dispersion are reported. Single-observation results cannot sustain claims concerning reliability or availability.
- Duration. The campaign does not establish long-term stability, seasonal thermal behaviour, or the

effect of accumulated tunnel contamination upon aperture and enclosure performance.

- Reliability metrics absent. Availability, handover success rate, block error rate and radio link failure statistics were not reported. These, and not peak throughput, are the metrics governing suitability for operational traffic.
- Measurement endpoint. As discussed in Sections 7.2 and 8.7, the location of the reference endpoint is unstated, which limits the interpretability of both the latency and the throughput figures.
- Operational traffic. The paper advances no claim that live train control signalling was conveyed. Any inference that the network is qualified to bear safety-related train control is unsupported by this work.

XI. FORWARD ROADMAP AND CONCLUSION

11.1 The Path to Grade of Automation 4

Driverless operation at Grade of Automation 4 places the communications bearer within the safety envelope, which alters the evidential requirement entirely. Peak throughput becomes largely immaterial; what governs is bounded worst-case latency at a specified percentile, quantified availability against a defined outage budget, deterministic failure behaviour, and a defence against deliberate radio interference. The attainment of these properties requires redundancy within the radio path, formal safety-case treatment under the applicable CENELEC standards, and most demandingly latency characterisation at the 99.999th percentile rather than at the mean. The results presented herein establish that the bearer is capable; they do not establish that it is qualified, and the distinction should not be elided.

11.2 Recommended Subsequent Measurement Campaign

- Downlink re-measurement employing parallel TCP streams and UDP against a station-local endpoint, so as to resolve the ceiling identified in Section 8.6.
- Latency measurement against a locally broken-out endpoint, reported as a distribution bearing 95th,

99th and 99.9th percentiles rather than as a single value.

- Quantified handover success rate across a statistically meaningful number of in-motion runs, with radio link failure and re-establishment events logged.
- Slice isolation under controlled load, reporting operational-slice latency with the commercial slice both idle and saturated.
- Reference signal received power and signal-to-interference-plus-noise ratio logging at every measurement point, so enabling attribution of the uplink variance observed in Table VIII.
- Loaded-cell testing with multiple concurrent terminals, so characterising the capacity envelope rather than single-terminal peak rate.

11.3 Conclusion

This paper has presented the architecture, analytical dimensioning, field implementation and experimental validation of an indigenous 5G Standalone Captive Non-Public Network with Fixed Wireless Access last mile, deployed across a 1.24 km underground section of Mumbai Metro Line 3. The system employs an O-RAN Option 7.2a functional split with radio units of Indian manufacture, a cloud-native Service-Based Architecture core within a Micro Data Centre at the Operations Control Centre, and passive wavelength-division fronthaul multiplexing permitting six radio units to share a single optical fibre core.

Field measurement across five operational domains established downlink throughput between 113 and 133 Mbps in every environment examined, including the deepest point of the bored tunnel, with round-trip latency between 16 and 25 ms and verified session continuity across inter-cell transitions aboard moving rolling stock. The retention of 124 Mbps downlink at tunnel centre a diminution of under 7 per cent from the portal constitutes experimental corroboration that node spacing of 100 m, which the propagation analysis of Section 5.1 places at substantially the tunnel break-point distance, delivers effective guided-mode coverage. The correspondence between an analytically derived break-point of 101 m and the deployed spacing of 100 m, and the field performance

consequent upon it, is advanced as the principal technical finding of this work.

Analysis of the results further indicates that measured downlink throughput was constrained by a mechanism lying outside the radio link, and the paper specifies the measurements required for its resolution. Taken together, the work demonstrates that an indigenous, open-standards 5G Standalone private network constitutes a technically viable substitute for cable-intensive communications infrastructure within underground transit, whilst stating expressly the reliability characterisation that remains outstanding before such a network may be qualified for safety-related train control. The architecture is offered as a replicable blueprint for metro operators pursuing both operational consolidation and sovereign control of critical transport infrastructure.

REFERENCES

- [1] 3GPP TS 23.501, "System architecture for the 5G System (5GS)," 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects.
- [2] 3GPP TS 23.502, "Procedures for the 5G System (5GS)," 3rd Generation Partnership Project.
- [3] 3GPP TS 38.300, "NR; NR and NG-RAN Overall Description; Stage 2," 3rd Generation Partnership Project.
- [4] 3GPP TS 38.104, "NR; Base Station (BS) radio transmission and reception," 3rd Generation Partnership Project.
- [5] 3GPP TS 38.214, "NR; Physical layer procedures for data," 3rd Generation Partnership Project.
- [6] 3GPP TS 38.306, "NR; User Equipment (UE) radio access capabilities," 3rd Generation Partnership Project.
- [7] 3GPP TS 22.261, "Service requirements for the 5G system; Stage 1," 3rd Generation Partnership Project.
- [8] 3GPP TS 33.501, "Security architecture and procedures for 5G System," 3rd Generation Partnership Project.
- [9] 3GPP TR 38.901, "Study on channel model for frequencies from 0.5 to 100 GHz," 3rd Generation Partnership Project.
- [10] O-RAN Alliance, "O-RAN Control, User and Synchronization Plane Specification," O-RAN.WG4.CUS.0.
- [11] IEEE Std 1588-2019, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems."
- [12] ITU-T Recommendation G.8275.1, "Precision time protocol telecom profile for phase/time synchronization with full timing support from the network."
- [13] IEEE Std 1474.1, "IEEE Standard for Communications-Based Train Control (CBTC) Performance and Functional Requirements."
- [14] A. G. Emslie, R. L. Lagace, and P. F. Strong, "Theory of the propagation of UHF radio waves in coal mine tunnels," *IEEE Transactions on Antennas and Propagation*, vol. AP-23, no. 2, pp. 192–205, March 1975.
- [15] D. G. Dudley, M. Lienard, S. F. Mahmoud, and P. Degauque, "Wireless propagation in tunnels," *IEEE Antennas and Propagation Magazine*, vol. 49, no. 2, pp. 11–26, April 2007.
- [16] K. Hrovat, G. Kandus, and T. Javornik, "A survey of radio propagation modeling for tunnels," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 2, pp. 658–669, 2014.
- [17] EN 50121-4, "Railway applications Electromagnetic compatibility Part 4: Emission and immunity of the signalling and telecommunications apparatus," CENELEC.
- [18] Department of Telecommunications, Government of India, "Guidelines for obtaining licence for establishing Captive Non-Public Network (CNPN)," June 2022.
- [19] ITU-R Recommendation P.1238, "Propagation data and prediction methods for the planning of indoor radiocommunication systems and radio local area networks."
- [20] Mumbai Metro Rail Corporation Ltd. and Bharat Sanchar Nigam Limited, "Quantum 5G Wireless: SpectraX Metro Rail Communications," Technical Report, 2026.
- [21] delaPlex Limited and BSNL, "Acceptance Test Report: MMRCL Project Data Connectivity using Quantum 5G Spectrum (CNPN/FWA)," Inspection Document v1.0, March 2026.

- [22] BSNL Telangana Circle, "High Level Solution Design for MML3 5G SA Network," Technical Architecture Document, 2026.
- [23] MMRCL, "Joint Inspection Confirmation for BSNL Quantum 5G SIM-less FWA and Micro Data Center Solution," Project Acceptance Letter, March 2026.