

5G NR FROM SITE TO TRACE

A Field Engineer's Guide to Numerology, BWP, SSB, PDCCH, Mobility and Throughput



FIELD PROMISE

Connect what is installed, what is configured, what the trace proves and what the service actually delivers.

PHYSICAL SITE	PROTOCOL TRACE	FIELD DECISION
---------------	----------------	----------------

By Takwa Sebai - takwasebai.com - Aligned to Release 18 specifications - Original field diagrams

Telecom engineer and technical manager - 10+ years across mobile networks - 4G/5G RAN validation and OSS

CONFIGURED SET what may be used	ACTIVE STATE what applies now	SCHEDULED RESOURCES what was allocated	DELIVERED SERVICE what the user received
------------------------------------	----------------------------------	---	---

STOP AT THE FIRST MISSING PROOF

Editorial and evidence notice

This guide is an original field-oriented synthesis. It connects the physical site, NR configuration, protocol messages, measurements and service symptoms. Public technical statements are cited to the applicable ETSI-published 3GPP specifications. The diagrams, troubleshooting structures and explanatory wording are original.

No vendor manual, commercial book, private email, customer trace or operator configuration has been copied into this edition.

Evidence status

The current edition uses declared synthetic training cases to make the diagnostic method concrete. They are exercises, not operator benchmarks, customer results or reconstructed incidents.

Chapter	Synthetic case IDs	Purpose
1	C-001	physical-path fault localisation
2	C-002	valid band, bandwidth and SCS context
3	C-003	configured BWP versus active state
4	C-004	cell result, SSB index and traffic behaviour
5	C-005	control allocation before data
6	C-006 and C-007	multi-KPI correlation and allocation
7	C-008	mobility phase localisation
8	C-009 to C-012	RF, scheduler, path and policy branches

Every numerical value, timestamp, route, topology or counter in those cases is synthetic. The cases are designed to close a calculation or expose an evidence boundary. They must not be presented as the author's lived incidents.

About Takwa Sebai

Takwa Sebai is a telecom engineer and technical manager with more than ten years across mobile networks, quality of service and technical-team leadership. Her work spans 4G/5G RAN validation, OSS, network KPIs, industrialisation and automation. She approaches validation from the same direction throughout this guide: start from the physical and configured reality, align the trace and measurements, then stop the conclusion where the evidence stops. Her current work and technical resources are published at takwasebai.com and on [LinkedIn](#).

Rights and trademarks

Copyright Takwa Sebai. All rights reserved. Published at **takwasebai.com**. Every public page carries a discreet Takwa Sebai attribution line. Removing that attribution does not waive the author's rights. Names referenced may be trademarks; descriptive use implies no affiliation or endorsement.

Safety and operational responsibility

This guide is educational material. It is not a substitute for approved method statements, site-access rules, electrical or climbing procedures, spectrum licences, change-management controls, vendor instructions or local law.

Do not access a live site, collect production data, run an intrusive test or change a network parameter without the required authorisation, competence, personal protective equipment and rollback plan.

Contents

The page numbers below are generated from the final continuous layout.

Editorial and evidence notice	2
How to use this guide	4
From the Physical Site to the Protocol Stack	6
Numerology and Bandwidth: Choose for the Scenario	16
Bandwidth Parts: Configured Is Not Active	27
SSB and Beam Sweeping: Coverage Is Not One Beam	38
PDCCH to PDSCH: Control Before Data	49
Read Radio Measurements Together	60
Mobility in NSA and SA	75
Low Throughput: Stop Guessing, Split the Problem	91
Sign off the observed state, not the control-path activity	108
Printable field checklists	109
Continue the work	111

How to use this guide

The chapters follow the same service path an investigation should follow.

The twelve field plates use A4 landscape pages so that their labels remain legible. For duplex printing, print those plates separately or enable automatic page orientation rather than shrinking them into portrait pages.

Chapter	Start from	Finish with
1	the physical site	a site-to-stack evidence map
2	the NR time-frequency grid	a scenario-based numerology decision
3	configured bandwidth parts	the active state at the test instant
4	SSB and cell measurements	a beam-evidence boundary
5	configured control resources	a PDCCH-to-PDSCH evidence chain
6	apparently conflicting KPIs	a synchronised multi-layer interpretation
7	an NSA or SA transition	the last completed mobility phase
8	low application goodput	the RF, allocation, path or policy branch

The field reading rule

For every symptom:

- 1 name the architecture, direction and exact test interval;
- 2 preserve what was configured without calling it active;
- 3 preserve what was measured without calling it the cause;
- 4 identify the last positive evidence;
- 5 identify the first missing or divergent evidence;
- 6 run one controlled test that can remove a branch;
- 7 compare with a known-good interval;
- 8 stop the conclusion where the evidence stops.

Evidence labels

- configured - present in configuration but not necessarily active;
- measured - observed at a defined resource, layer and time;
- reported - signalled by a UE, network function or counter;
- scheduled - selected and allocated by the network;
- delivered - useful data reached the stated measurement boundary;
- inferred - a hypothesis still requiring a discriminating test.

Field glossary

Term	Meaning in this guide
RAN	Radio Access Network
RRC	Radio Resource Control
PHY / MAC / RLC	Physical, Medium Access Control and Radio Link Control layers
PDCP / SDAP	Packet Data Convergence Protocol and Service Data Adaptation Protocol
DU / CU-CP / CU-UP	Distributed Unit, Central Unit Control Plane and Central Unit User Plane
BWP / SCS	Bandwidth Part and subcarrier spacing
SSB / SMTC	SS/PBCH block and SSB-based measurement timing configuration
CORESET / PDCCH / DCI	Control resource set, physical downlink control channel and downlink control information
PDSCH / TBS	Physical downlink shared channel and transport block size
CQI / RI / PMI / CRI	Channel quality, rank, precoding matrix and CSI-RS resource indicators
BLER / HARQ	Block error rate and hybrid automatic repeat request
RACH / BFR	Random-access channel procedure and beam failure recovery
CHO / DAPS	Conditional handover and dual active protocol stack
EN-DC / PSCell	E-UTRA-NR dual connectivity and Primary Secondary Cell
5QI / DNN / AMBR	5G QoS identifier, Data Network Name and aggregate maximum bit rate
UPF	User Plane Function

What this edition does not provide

- universal acceptance thresholds;
- vendor-specific commands or parameter values;
- a promise that one KPI identifies one root cause;
- permission to reproduce customer or operator data;
- a substitute for an authorised troubleshooting procedure.

The purpose is to make field reasoning explicit, repeatable and defensible.

From the Physical Site to the Protocol Stack

FIELD QUESTION

A user reports poor service. Which physical and logical points exist between the antenna and the application, and what evidence would let you isolate the failed part of that chain?

CHAPTER 01 · FIELD ORIENTATION

FROM SITE TO STACK

Physical path · logical functions · field evidence

WALK THE SITE BEFORE THE TRACE

PHYSICAL CHECKPOINTS

- 1 Active antenna / radio
- 2 Passive multiband panel
- 3 GNSS, lightning and aviation safety
- 4 Vertical fibre / DC route
- 5 Sealed cable entry
- 6 Optical distribution
- 7 DU / baseband processing
- 8 Transport, rectifier and battery reserve
- 9 Main earth bar and site earth ring

ROUTES


- fibre
- -48 V DC
- earth
- backhaul

WHAT THE ENGINEER SEES

- Site: connectors, fibre, DC, earth, alarms
- Counters: errors, retransmissions, throughput
- Trace: PHY → MAC → RLC → PDCP → RRC
- Decision: collect evidence before changing parameters

GENERIC FUNCTIONAL VIEW · NOT VENDOR-SPECIFIC

WHERE THE PROTOCOL FUNCTIONS LIVE



- Physical radio-to-DU link shown generically. It does not automatically mean O-RAN or one fixed functional split.
- *Exact function placement depends on implementation and deployment architecture.

BACKHAUL TO AGGREGATION

! COMMON FIELD MISTAKE

Jumping to RF optimisation before checking power, fibre, timing, transport and alarms. Confirm the failed link in the chain first.

Figure V-001. Generic training site. The physical scene is original and synthetic. The DU/CU function split is checked against 3GPP TS 38.401. The radio-to-DU link is shown generically and must not be interpreted automatically as O-RAN.

The first mistake in many investigations is to start with a favourite KPI. An RF engineer opens radio measurements. A transmission engineer checks packet loss. A site technician looks for a hardware alarm. Each view is useful, but none represents the whole path.

A mobile service is delivered by a chain. The chain crosses the air interface, radio equipment, fibre and power routes, baseband processing, transport, centralised RAN functions and the core network. Timing, grounding, configuration and software affect several links at once.

This chapter builds one map for that chain. It is deliberately generic. It does not describe a specific vendor cabinet, functional split or deployment rule.

1. The site is a chain, not a tower

From outside the fence, the tower dominates the site. Operationally, the tower is only one part of the service path.

Use three views during an investigation:

- 1 Physical view - equipment, connectors, fibre, DC power, grounding, timing and transport routes.
- 2 Logical view - RF, PHY, MAC, RLC, PDCP, SDAP, RRC and the interfaces between network functions.
- 3 Evidence view - alarms, measurements, traces, counters, configuration changes and controlled tests.

A diagnosis becomes stronger when all three views point to the same segment. For example, poor throughput plus high retransmissions is not enough to prove an RF problem. If the same time window also contains an unstable optical level or fronthaul error counter on one sector, the investigation should start closer to that physical link.

The practical service path

For field orientation, read the path in this order:

UE ↔ antenna/radio ↔ physical radio-to-DU link ↔ DU ↔ CU ↔ transport/core ↔ service

Then add the cross-cutting dependencies:

- DC power and battery reserve;
- protective earth and bonding;
- timing and synchronisation;
- software and configuration;
- alarms and management connectivity;
- environmental conditions.

These dependencies do not sit neatly in one protocol layer. A timing problem can appear as a radio symptom. A transport limitation can appear as low user throughput with healthy radio conditions. A power problem can be intermittent and leave no obvious service alarm by the time the engineer arrives.

Field principle

Do not ask only, "Which KPI is bad?" Ask, "Which link in the service chain could create this set of observations?"

2. Walk the physical path

The numbered checkpoints in Figure V-001 provide a repeatable site walk. They are not a replacement for the operator's safety procedures or installation documentation.

Only trained and authorised personnel should access towers, live DC systems, batteries or optical equipment. Measurements must follow the applicable electrical, laser, working-at-height and site rules.

Checkpoint 1 - Active antenna or radio

The radio converts between the digital signal path and RF. In an active antenna system, antenna elements and radio functions are integrated more closely than in a conventional passive-antenna installation.

Inspect or collect:

- mounting condition and visible damage;
- connector condition and weather protection where applicable;
- radio alarms and restart history;
- temperature or environmental alarms;
- sector identity and configuration consistency;
- evidence that the symptom is limited to one radio, sector or band.

What this does not prove: a visually correct installation does not prove healthy RF performance, and a clear alarm list does not prove that the link has never been intermittent.

Checkpoint 2 - Passive multiband panel

A passive panel may support other bands or radio chains on the same structure. Its presence matters because the user's service can involve more than the NR carrier being analysed, especially in NSA operation or during inter-RAT mobility.

Before drawing a conclusion, identify:

- which antenna system serves the affected carrier;
- whether the symptom follows one band, one sector or several;
- whether feeders, combiners or remote radio equipment exist in the actual deployment;
- whether the drawing used for analysis matches the as-built site.

Checkpoint 3 - GNSS, timing and site safety equipment

Timing deserves an explicit checkpoint. A synchronisation issue can affect TDD operation and may create symptoms that look like interference, availability loss or unstable performance.

Collect:

- synchronisation source and lock state;
- recent timing alarms, not only current alarms;
- changes to GNSS, PTP or transport timing;
- whether neighbouring cells show a correlated symptom;
- the exact start time of the degradation.

Lightning protection and aviation safety equipment are shown because they exist on the physical site. Protective systems must be inspected under the applicable site and regulatory procedures; this guide does not define those procedures.

Checkpoint 4 - Vertical fibre and DC routes

The route from tower-top equipment to the technical area is a real diagnostic object. It can be exposed to bending, water ingress, poor strain relief, damaged glands, heat, vibration or accidental disturbance during work.

Useful evidence includes:

- optical receive/transmit levels where exposed by the equipment;
- link error, loss-of-signal or flap history;
- DC alarms and voltage-drop evidence;
- inspection of bends, clamps, glands and cable identification;
- correlation between physical intervention time and counter changes.

Do not disconnect or manipulate an in-service link merely to “see what happens.” Define the test, rollback and service risk first.

Checkpoint 5 - Sealed cable entry

The cable entry separates the outside environment from the technical area. It is easy to omit from a logical architecture drawing and easy to disturb during site work.

Check for:

- correct sealing and strain relief;
- cable routing that avoids excessive bend or compression;
- identification that matches both ends;
- water or contamination evidence;
- a recent installation or maintenance change.

Checkpoint 6 - Optical distribution

The optical distribution point provides a place to identify, route and test fibre connections. It is also a place where a clean diagram can hide messy reality: crossed labels, spare fibres, patching changes or contamination.

A useful fibre test needs context:

- expected route and endpoints;
- reference level or known-good comparison;
- measurement method and test equipment;
- connector state before and after intervention;
- counter behaviour during the same time window.

An isolated optical measurement without the equipment’s error history may miss an intermittent fault.

Checkpoint 7 - DU or baseband processing

In the generic 3GPP split used in this chapter, the gNB-DU hosts PHY, MAC and RLC. It terminates the F1 interface toward the gNB-CU.¹

This is where an engineer may correlate:

- cell and carrier availability;
- PHY decoding and retransmission behaviour;
- scheduler and resource usage;
- RLC behaviour;
- radio-link and lower-layer alarms;
- the exact cell, sector, UE and time window.

The physical DU may be at the site or at another edge location. The logical function matters more than the cabinet position.

Checkpoint 8 - Transport, rectifier and battery reserve

These are three different functions grouped in one checkpoint only to keep the overview readable.

Transport

- interface state and flap history;
- packet loss, latency and congestion evidence;
- routing or QoS changes;
- comparison with unaffected cells using the same path;
- location of the first degraded measurement point.

Rectifier and DC distribution

- current and historical alarms;
- supply stability under the applicable test procedure;
- distribution protection state;
- signs of a recent load or cabling change.

Battery reserve

- battery alarms and state information;
- autonomy or capacity tests performed under approved procedures;
- evidence of repeated mains events;
- temperature and maintenance history.

"No outage" does not mean "healthy power." A marginal supply can produce resets or instability without a long site-down event.

Checkpoint 9 - Main earth bar and site earth ring

Grounding and bonding are safety-critical and site-specific. They must be evaluated by qualified personnel against the applicable design and regulation.

For diagnosis, record:

- visible integrity and bonding of cabinets, tower and cable routes;
- evidence of recent civil or electrical work;
- corrosion, damage or loose connections identified through the approved inspection process;
- whether symptoms correlate with weather or electrical events.

Do not turn a safety observation into a speculative RF conclusion. Record it, escalate it correctly and test the service hypothesis with appropriate evidence.

3. Map the physical site to the logical gNB

The 3GPP architecture separates logical functions from physical packaging. A supplier may deploy them in different hardware or locations, but the functional map gives engineers a common language.

gNB-DU

The gNB-DU hosts:

- PHY;
- MAC;
- RLC.

It connects to the gNB-CU-CP through F1-C and to the gNB-CU-UP through F1-U.¹

gNB-CU-CP

The gNB-CU-CP hosts:

- RRC;
- the control-plane part of PDCP.

It connects to the DU over F1-C and to the gNB-CU-UP over E1.¹

gNB-CU-UP

For a gNB, the gNB-CU-UP hosts:

- SDAP;
- the user-plane part of PDCP.

It connects to the DU over F1-U, to the gNB-CU-CP over E1 and toward user-plane core functions over NG-U.¹

Why RU/AAU is marked as generic

TS 38.401 defines the gNB-DU and gNB-CU functions used here. It does not make the physical RU/AAU box in Figure V-001 a universal 3GPP node with one mandatory fronthaul split.

The radio-to-DU path therefore remains intentionally generic. Before applying an O-RAN, CPRI, eCPRI or vendor-specific interpretation, the engineer must identify the actual deployment architecture and its documentation.

Common architecture mistake

Seeing an RU and DU on a drawing does not prove which functional split is deployed.

4. Where evidence appears

The same fault can leave evidence in several places. The table below prevents one source from being treated as definitive on its own.

Area	What the engineer can observe	What it can support	What it cannot prove alone
Physical site	damage, routing, connector state, water, alarms	a physical hypothesis and test location	service impact or root cause
RF measurements	serving cell, level, quality, interference-related indicators	radio condition at a place and time	scheduler, transport or policy health
PHY	decoding, HARQ/BLER-related behaviour, MCS/rank context	lower-layer radio behaviour	the physical cause of that behaviour
MAC	scheduling, resource usage, throughput allocation	capacity or scheduler hypothesis	transport or application performance
RLC	retransmission and buffer behaviour	radio-link or delivery stress	which lower-level cause created it
PDCP/SDAP	user-plane volume and bearer/QoS context	end-to-end bearer analysis	application-server health
RRC	connection, reconfiguration and mobility signalling	control-plane state and failure phase	user-plane performance after setup

Area	What the engineer can observe	What it can support	What it cannot prove alone
Transport	loss, latency, utilisation, interface errors	transport-path hypothesis	RF health
Power/timing	alarms, state changes, test results	cross-cutting infrastructure hypothesis	the affected protocol transaction without correlation

The word correlation matters. Align observations by:

- timestamp and timezone;
- site, sector and cell identity;
- UE or session identity when permitted;
- software/configuration version;
- before/after change window;
- affected and unaffected comparison groups.

A counter accumulated over one hour should not be compared casually with a drive-test sample from ten seconds unless the time and scope are reconciled.

5. Translate symptoms into testable branches

Start with the smallest confirmed scope.

Symptom: the complete site is unavailable

First branches:

- mains/DC and rectifier state;
- transport reachability;
- timing if required for service;
- common DU or site-control function;
- planned work or configuration event.

Do not begin by tuning individual cell parameters when all cells disappeared together.

Symptom: one sector or one band is degraded

First branches:

- affected antenna/radio and its physical link;
- sector-specific alarm and error history;
- carrier and cell configuration difference;
- interference or coverage change;
- comparison with the other sectors on the same site.

Symptom: radio level looks good but throughput is poor

First branches:

- signal quality, BLER and retransmission context;
- scheduler/resource availability;
- UE capability and active configuration;
- transport limitation;
- policy, QoS or core/application limitation.

"Good RSRP" is not a throughput diagnosis.

Symptom: degradation is intermittent

First branches:

- timestamped alarm and interface-flap history;
- temperature, weather, power or timing correlation;
- connector/fibre path disturbed by recent work;
- load-dependent congestion;
- mobility or beam behaviour on a repeatable route.

An intermittent problem is easy to erase by rebooting. Preserve evidence before recovery actions when service procedures allow it.

6. Synthetic training case C-001 - a physical link wearing a radio mask

Declaration

This is a synthetic training scenario. It is not a customer case, measured result or claim about the author's experience.

Trigger

After planned work at a generic three-sector site, users on one sector report unstable throughput. The site remains reachable and the other sectors behave normally.

Initial observations

- no continuing site-down alarm;
- DC and timing states appear stable during the investigation window;
- the transport path is not congested and the unaffected sectors share part of that path;
- the affected sector shows intermittent lower-layer errors and retransmission stress;
- the issue began after physical work near the cable route.

These observations narrow the scope. They do not yet prove a fibre fault.

Weak conclusion

"Retransmissions are high, so increase radio power or change mobility parameters."

This conclusion skips the change history and the fact that only one physical branch is affected.

Better test sequence

- 1 Freeze unrelated parameter changes.
- 2 Export the relevant alarms and interface counters for the same time window.
- 3 Compare the affected sector with the two unaffected sectors.
- 4 Inspect the labelled fibre path and sealed entry under the approved site procedure.
- 5 Measure or compare the optical link using the operator's accepted method.
- 6 If authorised, test with a known-good patch path or replace the suspected passive element with rollback prepared.
- 7 Confirm whether link errors stop and whether the radio symptoms recover in the same window.

Decision

Only if the physical-link evidence remains healthy should the investigation move deeper into RF configuration, scheduling, transport or policy branches.

The learning point is not “fibre always causes retransmissions.” It is:

A logical KPI identifies where stress appears. It does not automatically identify the physical cause.

7. Before changing a parameter

Use this checklist at the start of a field investigation.

Define the problem

- ☐ What exactly fails: access, retainability, mobility, latency or throughput?
- ☐ Which site, sector, band, cell, route, UE type and service are affected?
- ☐ When did it start?
- ☐ Is the problem continuous, intermittent, load-dependent or location-dependent?
- ☐ What changed before the first occurrence?

Build the evidence set

- ☐ current and historical alarms;
- ☐ configuration and software change history;
- ☐ physical inspection evidence where safe and authorised;
- ☐ interface state and error counters;
- ☐ RF/PHY/MAC/RLC/PDCP/RRC observations relevant to the symptom;
- ☐ transport and timing evidence;
- ☐ a healthy comparison site, sector, time window or UE;
- ☐ timestamps normalised to one reference.

Test one hypothesis

- ☐ state the expected observation if the hypothesis is true;
- ☐ state what would falsify it;
- ☐ define service risk and rollback;
- ☐ change one controlled variable where possible;
- ☐ verify recovery with the same evidence used to establish the fault;
- ☐ preserve the before/after record.

Takwa Sebai working criterion

I do not sign a physical-site root cause unless the affected scope is fixed, the first divergent boundary is time-aligned with the symptom, and one controlled test separates that boundary from at least one credible alternative. A clean alarm screen is not that test.

Chapter takeaway

The physical site and the protocol stack are not two separate subjects. They are two views of the same service chain.

When a KPI changes:

- 1 define the affected scope;
- 2 locate the stressed logical function;
- 3 identify the physical links that can influence it;
- 4 align alarms, traces, counters and change history;
- 5 test one hypothesis before tuning the network.

Final evidence gate

Gate	Question before action
Scope	Are the affected site, sector, cell, UE/service and time window identified?
Path	Is the first stressed physical or logical segment located?
Correlation	Are alarms, traces, counters and changes aligned to the same clock and scope?
Comparison	Is there a healthy sector, route, UE or time window for comparison?
Test	Can the next test falsify the hypothesis, and is rollback prepared?

Next chapter

Numerology follows the same method: definition first, then where it appears, what it changes, what it does not prove and how two deployment objectives lead to different choices.

References

¹ ETSI, [TS 138 401 V18.7.0 - NG-RAN Architecture description](#), especially clause 3.1 for gNB-DU, gNB-CU-CP and gNB-CU-UP definitions and clause 6.1.2 for F1-C, F1-U and E1.

Numerology and Bandwidth: Choose for the Scenario

FIELD QUESTION


A carrier uses 30 kHz subcarrier spacing. What does that tell you, what does it not tell you, and which evidence would justify changing the design?

CHOOSE SCS FROM CONSTRAINTS, THEN TEST

A field decision map — not a “faster numerology” ranking

DECISION PATH

Follow the arrows. Do not skip to a parameter change.



WHAT THE ENGINEER ACTUALLY CHECKS

Configured grid + UE context + aligned RF/L2 evidence + load + transport + a known-good comparison.



COMMON FIELD MISTAKE

Changing SCS to fix throughput before checking allocation, MIMO, retransmissions, load, transport and policy.

SYNTHETIC DEPLOYMENTS

A · URBAN MID-BAND LAYER

- n78 · 100 MHz · TDD
- 30 kHz SCS
- dense urban macro context
- capacity-oriented test objective

B · LOW-BAND COVERAGE LAYER

- n28 · 20 MHz · FDD
- 15 kHz SCS
- suburban / rural context
- coverage-oriented objective

TRAINING LIMIT

Both combinations are valid examples.
They do not isolate the effect of SCS.
No measured performance is claimed.

Sources: TS 38.104 V18.13.0 §5.3.5; TS 38.211 V18.9.0 §§4.2, 4.3.2; TS 38.214 V18.8.0 §5.1.3.

TAKWA SEBAI · TAKWASEBAI.COM · V-002

Figure V-002. Original decision map. It starts with deployment constraints and ends with observed behaviour. The two example deployments are synthetic and are not a performance comparison.

Numerology is often explained with one formula and then turned into a ranking: a larger subcarrier spacing creates shorter slots, so it must be faster.

That shortcut is technically incomplete and operationally dangerous.

Subcarrier spacing defines part of the NR time-frequency grid. It affects how time and frequency resources are organised. It does not, by itself, establish how many resources a UE receives, which modulation and coding are used, how many layers are realised, how much retransmission occurs, whether the TDD pattern is favourable, whether transport is constrained or whether the test method is valid.

The field method is therefore:

- 1 establish the configured grid;
- 2 verify that the band, channel bandwidth and SCS combination is valid;
- 3 identify the carrier, BWP and UE context actually in use;
- 4 separate the intended design objective from the observed symptom;
- 5 test competing explanations before proposing a numerology change.

1. What numerology defines

For NR, subcarrier spacing follows:

$$\Delta f = 2^{\mu} \times 15 \text{ kHz}$$

where μ is the numerology index.¹

For the normal cyclic-prefix cases used in this chapter:

μ	Subcarrier spacing	Slots in a 1 ms subframe	Reference slot duration	OFDM symbols per slot
0	15 kHz	1	1 ms	14
1	30 kHz	2	0.5 ms	14
2	60 kHz	4	0.25 ms	14
3	120 kHz	8	0.125 ms	14
4	240 kHz	16	0.0625 ms	14
5	480 kHz	32	0.03125 ms	14
6	960 kHz	64	0.015625 ms	14

The slot durations above are derived from the 1 ms subframe divided by the number of slots per subframe. They are useful for reading the grid; they are not end-to-end latency promises. The FR1 examples in this guide use $\mu = 0$ or $\mu = 1$. The higher rows are retained so that the table does not silently stop before the 240 kHz SSB context and the 480/960 kHz FR2-2 grids; their applicability still depends on the band, signal or channel and equipment support.¹

ONE MILLISECOND, FOUR NUMEROLOGIES

The grid changes in time and frequency. No row is a universal “best”.

NORMAL CP

SAME 1 ms SUBFRAME

Δf

0 **15 kHz**
1 slot $\times$ 1 ms




1 **30 kHz**
2 slots $\times$ 0.5 ms



2 **60 kHz**
4 slots $\times$ 0.25 ms



3 **120 kHz**
8 slots $\times$ 0.125 ms



i **Each slot contains 14 OFDM symbols in these normal-CP examples.**
Fine internal lines represent symbols; strong white boundaries represent slots.

WHAT CHANGES

1 **Frequency spacing**
 Δf doubles with each μ step.

2 **Time grid**
Slots per 1 ms double.

3 **Symbol duration**
Useful symbols become shorter.

WHAT THIS DOES NOT PROVE

- user throughput
- end-to-end latency
- coverage or mobility result
- the best deployment choice

Source: 3GPP TS 38.211 V18.9.0, clauses 4.2 and 4.3.2. Slot durations are derived from a 1 ms subframe.

TAKWA SEBAI · TAKWASEBAI.COM · V-003

Figure V-003. Four normal-CP grids aligned to the same 1 ms subframe. More slot boundaries do not mean more user throughput by themselves.

As μ increases:

- the subcarrier spacing doubles;
- the useful OFDM symbol period becomes shorter;
- the number of slots in the same 1 ms window doubles;
- the scheduling grid can be represented at a finer time granularity.

The word can matters. A shorter slot does not prove that the scheduler served a UE sooner, that the packet crossed the complete network faster or that the application responded faster.

2. Read time and frequency together

An engineer should be able to look at a numerology and describe both axes.

Time view

Within the same 1 ms subframe:

- 15 kHz SCS gives one slot;
- 30 kHz gives two;
- 60 kHz gives four;
- 120 kHz gives eight.

This changes the available time structure for scheduling and signal design. It does not show which slots carry downlink, uplink or flexible symbols in a TDD deployment. The actual direction pattern and the resources assigned to a UE must still be checked.

Frequency view

A larger SCS places adjacent subcarriers farther apart. The physical resource grid therefore changes in frequency as well as time.

This is why “shorter slot” is only half the explanation. A design choice occupies a particular band and channel, uses specific signals and bandwidth parts, supports a defined UE population and operates in a real propagation and interference environment.

What the grid does not show

A clean grid diagram does not contain:

- the number of physical resource blocks allocated to the UE;
- modulation order and target code rate;
- realised MIMO layers;
- retransmission pressure;
- cell load and scheduler decisions;
- TDD downlink/uplink opportunity;
- protocol overhead;
- transport or core limitations;
- policy, subscription, server or device limits.

Treat the grid as a coordinate system, not as a performance result.

TDD pattern turns slots into directional opportunity

For a TDD carrier, the slot grid is not all downlink and it is not all uplink. A repeating pattern such as **DDDSU** contains three downlink slots, one special slot and one uplink slot. The special slot must then be decoded into downlink, guard and uplink symbols before it can be used in a throughput budget.

In practice, record:

- the pattern periodicity and the exact slot-format sequence;
- the number of usable PDSCH symbols in full and special slots;
- whether the observed interval crosses a pattern boundary;
- the direction-specific scheduling duty cycle;
- guard, control, reference-signal and other reserved resources.

My first throughput check is deliberately conservative: I translate the configured pattern into an explicit downlink-opportunity factor, state the usable PDSCH symbols assumed, then verify that the observed goodput remains below the resulting physical ceiling. If that arithmetic does not close, the case is wrong before the RCA begins.

One wall-clock factor throughout this guide

Every absolute downlink-goodput ceiling in this edition uses one effective wall-clock downlink factor, **f_{DL}**. For the synthetic **DDDSU** model, the maximum teaching value is **f_{DL} = 0.70** after the special slot is converted to downlink-symbol equivalents rather than counted as a full DL slot.

When a case publishes a scheduled-slot share against all wall-clock slots, that measured share is already the row's **f_{DL}**: it must not exceed **0.70** and must not be multiplied by **0.70** again. When no measured share is given, the ceiling uses the model maximum of **0.70**. Every ceiling also applies the row's stated initial-BLER factor, **1 – BLER**, and the declared number of PDSCH data-symbol equivalents after control and reference-signal overhead.

3. Constraints come before preferences

The first question is not "Which SCS do I prefer?" It is "Which combinations are defined and supported for this deployment?"

TS 38.104 lists supported channel-bandwidth and SCS combinations by operating band.² The relevant release and equipment capability still have to be checked; a training table is not a substitute for the deployed product documentation.

At minimum, establish:

- 1 operating band and frequency range;
- 2 duplex mode;
- 3 configured channel bandwidth;
- 4 carrier SCS and the SCS used by the relevant BWP;
- 5 SSB and common-configuration context;
- 6 UE and network support;
- 7 software and feature baseline.

This order prevents an impossible or irrelevant comparison. A low-band FDD coverage layer and a mid-band TDD capacity layer do not differ only by numerology.

Field principle

If band, bandwidth, duplex mode, load and radio conditions all change, the performance difference cannot be assigned to SCS alone.

4. Why SCS is not a throughput button

TS 38.214 makes the separation visible. The downlink transport-block calculation uses, among other inputs, the modulation order, target code rate, number of layers and allocated physical resource blocks.³

That normative dependency supports a narrow but important conclusion:

SCS alone is insufficient to determine user throughput.

It does not prove that numerology is unimportant. It means that a throughput investigation must include the variables that directly determine how much useful information was scheduled and successfully delivered.

A stronger evidence set

For a low-throughput symptom, align:

- configured carrier and active BWP context;
- UE capability and serving mode;
- allocated resources over the test interval;
- modulation, coding, rank/layers and retransmission evidence;
- radio quality and mobility state;
- cell load and comparison window;
- TDD pattern where applicable;
- transport, server and test-method evidence;
- policy or subscription constraints.

Only then ask whether the numerology design remains a plausible limiting factor.

A common false causal chain

30 kHz SCS → shorter slot → lower latency → higher throughput

Every arrow requires evidence.

- A shorter reference slot is a property of the grid.
- Lower radio scheduling delay depends on the actual configuration and scheduling event.
- Lower end-to-end latency depends on the complete service path.
- Higher throughput depends on successfully delivered resources over time.

The chain may be directionally useful in a design discussion, but it is not a field diagnosis.

5. Synthetic training case C-002 - two deployments, not an A/B test

The following scenarios are pedagogical. They are not measured networks, customer cases or evidence of the author's deployment practice.

Their band, channel-bandwidth and SCS combinations are supported combinations in the cited Release 18 version of TS 38.104.²

Attribute	Scenario A - urban mid-band layer	Scenario B - low-band layer
Band	n78	n28
Channel bandwidth	100 MHz	20 MHz
Duplex mode	TDD	FDD
SCS	30 kHz	15 kHz
Context	Dense urban macro	Suburban or rural
Training objective	Capacity-oriented test	Coverage-oriented test

What this comparison teaches

The two configurations are plausible examples of different deployment roles. They help an engineer ask the right questions:

- What problem is each layer intended to solve?
- Which bandwidth and duplex opportunities exist?
- Which UEs and services must be supported?
- Which measurements demonstrate that the design objective is being met?

What this comparison cannot prove

It cannot prove that:

- 30 kHz is “for speed” and 15 kHz is “for coverage” in every network;
- the SCS caused any difference between the two scenarios;
- one scenario has better latency, mobility or spectral efficiency;
- the same result would occur with a different band, bandwidth, site grid, antenna system, TDD pattern or load.

Those claims would require controlled measurements or a publishable real case.

6. Where to verify the configured reality

Use three evidence layers.

Layer 1 - planned and configured design

Collect the authorised configuration view for:

- band, frequency and duplex mode;
- channel bandwidth;
- carrier-specific SCS;
- common and BWP-related SCS;
- SSB/common configuration;
- TDD pattern where applicable;
- software and feature baseline.

Vendor labels differ. Preserve the original export or screenshot, its timestamp and the object identity instead of rewriting values from memory.

Layer 2 - UE and signalling context

Use an authorised decoded trace or device log to confirm:

- which serving cell and carrier the UE used;
- NSA or SA context;
- which BWP context is relevant;
- whether a reconfiguration occurred during the test;
- whether the test interval includes mobility, fallback or another state change.

This chapter intentionally avoids vendor-specific information-element paths until the exact decoder and release are selected. Do not infer the active configuration from one isolated field name.

Layer 3 - observed performance

Align the configuration with:

- exact test start and stop time;
- radio conditions and location;
- UE model, capability and thermal state;
- resource allocation and load;
- retransmission and MIMO evidence;
- transport and server measurements;
- a known-good comparison.

Configuration says what the network was intended to do. The trace and measurements say what happened for that UE in that interval.

7. A field decision sequence

Symptom: expected throughput is not reached

- 1 Confirm the result is repeatable with a controlled test method.
- 2 Confirm the serving cell, carrier, mode and BWP context.
- 3 Verify band, channel bandwidth and SCS against the approved design.
- 4 Inspect allocation, modulation/coding, layers and retransmissions.
- 5 Check load, TDD opportunity and comparison windows.
- 6 Exclude transport, server, policy and UE constraints.
- 7 Compare with a known-good cell using the same carrier design where possible.
- 8 Only then evaluate whether a numerology-related design change deserves a controlled trial.

Symptom: latency is higher than expected

Do not stop at slot duration. Separate:

- radio scheduling and retransmission delay;
- TDD waiting opportunity;
- RAN processing and queueing;
- transport and core delay;
- server and application response;
- test-tool resolution and clock alignment.

A 0.5 ms reference slot cannot explain a 30 ms end-to-end result without the rest of the path.

Symptom: coverage or mobility is unstable

Verify:

- band and antenna/beam context;
- serving and neighbour measurements;
- SSB and measurement configuration;
- handover or beam-management phase;
- speed and route repeatability;
- timing, interference and retransmission evidence.

Do not use numerology as a substitute name for every RF trade-off.

8. Common field mistake - changing SCS before proving causality**Weak conclusion**

"The speed is low. Move to a wider SCS."

Why it is weak

It does not show:

- that the current combination is invalid;
- that the UE lacks time-frequency opportunities because of SCS;
- that allocation, MIMO and retransmissions are healthy;
- that transport and policy are clear;
- that the proposed configuration is supported;
- what measurable outcome would confirm success.

Better test

Compare the affected case with a known-good cell or interval that has:

- the same carrier design;
- the same UE and traffic method;
- comparable radio conditions;
- a documented load state;
- aligned RAN and transport evidence.

If a configuration trial is still justified, define:

- the hypothesis;
- expected and falsifying observations;
- affected cells and services;
- safety and service risk;
- rollback;
- measurement window;
- acceptance threshold based on real baseline data.

Do not invent a threshold because it looks reasonable in a training document.

9. Minimum evidence package

Configuration

- ☐ release/software baseline;
- ☐ band and duplex mode;
- ☐ channel bandwidth;
- ☐ carrier/common/BWP SCS context;
- ☐ SSB and TDD context where applicable;
- ☐ relevant UE support.

Test context

- ☐ site, sector, cell and time window;
- ☐ UE model, capability and state;
- ☐ NSA/SA and serving-carrier context;
- ☐ location, route, speed and radio conditions;
- ☐ server, protocol, stream count and test duration;
- ☐ cell load and known-good comparison.

Performance evidence

- ☐ allocation over time;
- ☐ modulation/coding context;
- ☐ rank/layers;
- ☐ BLER/retransmission context;
- ☐ downlink/uplink opportunity;
- ☐ transport loss, latency and utilisation;
- ☐ policy or subscription limits checked.

Decision

- ☐ What does the SCS explain directly?
- ☐ What remains only correlated?
- ☐ Which competing cause has been excluded, and with what evidence?
- ☐ What observation would falsify the proposed explanation?
- ☐ Is a controlled change necessary, and is rollback ready?

Takwa Sebai working criterion

I do not recommend an SCS change from a throughput or latency graph alone. The minimum decision record fixes band, bandwidth, duplex mode, TDD opportunity, active BWP, PRBs, MCS table, layers, retransmissions, load and the measurement boundary. If several of those move together, the result is a comparison, not proof that numerology caused it.

Chapter takeaway

Numerology is a grid choice inside a complete radio design.

The engineer's sequence is:

- 1 define the SCS relationship;
- 2 check band and channel constraints;
- 3 identify the carrier and BWP context actually used;
- 4 measure allocation and delivery, not only configuration;
- 5 separate SCS effects from bandwidth, duplex, RF, load, MIMO, transport and policy;
- 6 change the design only when the hypothesis is testable.

Next chapter

A carrier can contain several configured bandwidth parts, while the UE operates in a specific BWP context. Chapter 3 separates configured, default, initial and active BWP states without confusing the container with the resource currently used.

References

¹ ETSI, [TS 138 211 V18.9.0 - NR; Physical channels and modulation](#), especially clause 4.2 for $\Delta f = 2^{\mu} \times 15 \text{ kHz}$ and clause 4.3.2 with table 4.3.2-1 for the normal-CP slot structure.

² ETSI, [TS 138 104 V18.13.0 - NR; Base Station radio transmission and reception](#), especially clause 5.3.5 and table 5.3.5-1 for NR channel-bandwidth and SCS combinations.

³ ETSI, [TS 138 214 V18.9.0 - NR; Physical layer procedures for data](#), especially clause 5.1.3 for MCS-table selection, modulation order, target code rate and transport-block determination.

Bandwidth Parts: Configured Is Not Active

FIELD QUESTION

The cell has a wide carrier. How much spectrum is the UE actually monitoring and using at the instant under investigation?

5G NR FIELD GUIDE • CHAPTER 3 • V-004


CONFIGURED IS NOT ACTIVE

Reconstruct the BWP state at the measurement instant — not only the RRC inventory

COMMON ERROR
List # current state

1 MAP THE CONFIGURED SET

Conceptual DL view • widths and positions not to scale



i All three may be configured.
The list defines possible operating regions.
It does not identify the region active "now".

Read the RRC fields precisely

initialDownlinkBWP


firstActive...Id

default...Id

bwp-InactivityTimer

2 REPLAY THE STATE CHANGES

One active BWP per direction in this serving-cell context



3 PROVE "NOW"

Three evidence layers, one synchronized interval

LAYER 1 • RRC
Proves the configured set and the state imposed by that event.

Does not exclude a later switch.

LAYER 2 • EVENTS
DCI, timer, SCell activation, random access, later RRC.

Build the complete timeline.

LAYER 3 • OBSERVED
UE log, MAC/scheduler trace, resource location and allocation.

Corroborate the active state.

IF AN EVENT IS INVISIBLE
Report UNKNOWN.
Never promote expectation to evidence.

THE ONLY DEFENSIBLE PERFORMANCE CHAIN



TAKWA SEBAI • TAKWASEBAI.COM • V-004 • technical basis: 3GPP TS 38.300 §6.10, TS 38.213 §12, TS 38.321 §5.15 and TS 38.331

Figure V-004. Original state-and-evidence map. The carrier, BWP widths and event sequence are conceptual. They illustrate the method, not a customer configuration.

A carrier bandwidth is not the same object as a bandwidth part. A configured BWP is not necessarily active. A **firstActiveDownlinkBWP-Id** is not a promise that the same BWP remains active later. A **defaultDownlinkBWP-Id** describes a fallback after a specific event; it is not another name for the current state.

These distinctions matter because an isolated RRC configuration can contain several valid BWP definitions while showing only part of the state history. If a switch occurred through PDCCH/DCI, an inactivity timer expired, random access changed the applicable BWP or another reconfiguration arrived later, the configuration list alone cannot prove which BWP was active at the time of the symptom.

The field method is:

- 1 separate the complete carrier from the BWP definitions;
- 2 inventory every configured DL and UL BWP without calling them active;
- 3 identify the event that established or changed the active BWP;
- 4 reconstruct the timeline up to the exact test instant;
- 5 confirm the state with the best available lower-layer or device evidence;
- 6 link performance only to the BWP that was demonstrably active during the same interval.

1. Why bandwidth parts exist

NR bandwidth adaptation allows a UE's receive and transmit bandwidth to be smaller than the cell bandwidth. The configured bandwidth can change in width, frequency location and subcarrier spacing. A BWP is therefore a subset of the cell bandwidth, and the UE is told which configured BWP is active.¹

This creates a useful separation:

- the network can deploy a wide carrier;
- the UE does not have to operate across that entire carrier in every state;
- different BWP configurations can support different operating contexts;
- the active operating region can change without redefining the complete carrier.

The standard explains the mechanism. It does not assign a universal operational motive to each BWP in every network. Energy saving, access, traffic handling, device capability and deployment policy may influence the design, but the actual purpose must be established from the approved configuration and engineering intent.

Field principle

Treat the carrier as the container, configured BWPs as possible operating regions, and the active BWP as a time-dependent state.

2. The terms that must not be mixed

Term	What it establishes	What it does not establish
Carrier bandwidth	The complete NR channel bandwidth deployed for the cell/carrier context	That one UE monitors or uses the complete bandwidth now
Configured BWP	A BWP definition available in the UE configuration	That the BWP is active
Initial BWP / BWP ID 0	The initial common BWP context defined by the configuration; the initial DL BWP includes the common control context described by RRC	That BWP 0 remains active during connected traffic
Additional BWP	Another configured DL or UL BWP with its own identity and parameters	That "additional" means active, wider or faster
Dedicated configuration	UE-specific configuration parameters associated with a BWP	A separate instantaneous state named "dedicated BWP"
First active DL/UL BWP	The BWP selected at a defined transition such as an RRC reconfiguration; for an SCell, the configured BWP used on activation	That it remains active after later switching events
Default DL BWP	The DL BWP selected when the inactivity timer expires; if the field is absent, the initial DL BWP is used as default	The current DL BWP before timer expiry
Currently active BWP	The BWP in which the UE is presently expected to perform the applicable DL or UL operations	Which event made it active unless the timeline is available

TS 38.213 allows a UE to be configured with up to four downlink BWPs and up to four uplink BWPs for a serving cell.² That is a configuration inventory, not four simultaneously active operating regions. MAC BWP operation activates one BWP while deactivating the previously active BWP in the same direction and serving-cell context.³

Why "dedicated BWP" is risky shorthand

RRC defines common and dedicated parameter groups for BWP configuration. Engineers sometimes turn that structure into an operational label: "the dedicated BWP is active."

The first half may describe where parameters were carried. The second half is a state claim and requires separate evidence.

Use explicit wording instead:

- "BWP ID 1 is configured with dedicated parameters."
- "BWP ID 1 became active at the RRC reconfiguration."
- "The available trace does not prove that BWP ID 1 was still active 20 seconds later."

That language preserves the difference between configuration and observation.

3. Downlink and uplink are separate questions

A serving-cell configuration can contain DL and UL BWP definitions. On an active DL BWP, the UE performs operations such as PDCCH monitoring and DL-SCH reception. On an active UL BWP, it performs applicable operations such as UL-SCH, PUCCH or SRS transmission. The UE does not perform those active-BWP operations on an inactive BWP.³

Therefore, "the UE is on BWP 1" is incomplete. Ask:

- DL BWP 1, UL BWP 1, or both?
- paired or unpaired spectrum?
- at which instant?
- established by which event?
- proven by which decoder or measurement?

For unpaired spectrum, DL and UL BWPs with the same ID are treated as a pair, and the specification constrains their centre-frequency relationship.² In other contexts, do not assume that a DL identifier automatically proves the UL state.

The practical reporting format should contain two fields:

Direction	Active BWP at test instant	Evidence
Downlink	record the BWP ID, or write unknown	name the RRC, DCI, timeline or device-log evidence
Uplink	record the BWP ID, or write unknown	name the RRC, DCI, timeline or device-log evidence

“Unknown” is a stronger engineering answer than an unsupported inference.

4. How the active BWP can change

Release 18 MAC procedures describe several controls for BWP switching.³ The exact applicability depends on the serving-cell state and configuration.

RRC reconfiguration

For the SpCell, **firstActiveDownlinkBWP-Id** and **firstActiveUplinkBWP-Id** can establish the active BWP when the relevant reconfiguration is applied. If a first-active field is absent, the RRC reconfiguration does not by itself impose that directional switch.⁴

For an SCell, the first-active values identify the BWP used when the SCell is activated, subject to the applicable configuration.

Evidence needed: the complete decoded reconfiguration, successful application context, serving-cell identity and exact timestamp.

PDCCH/DCI indication

A downlink assignment or uplink grant can include a BWP indicator that changes the active BWP among the configured set.²³

This is the main reason a static RRC snapshot may become stale as evidence of current state. RRC defines the available set and transition parameters; lower-layer control can select within that set later.

Evidence needed: a decoder that exposes the relevant DCI and BWP indicator, aligned with the correct UE, cell, direction and time.

BWP inactivity timer

The **bwp-InactivityTimer** applies to the active DL BWP under the specified conditions. On expiry, the active DL BWP changes to the configured default DL BWP. If **defaultDownlinkBWP-Id** is absent, the initial DL BWP is the fallback.²³⁴

Do not write:

“The default BWP is active.”

Write:

“The timer expired at 10:42:18.540; the configured fallback was the initial DL BWP, ID 0; no later DL switch was observed before the 10:42:18.700 test sample.”

Evidence needed: timer configuration, the events that start or restart it, the expiry timeline and any later switch.

Random access and MAC-controlled changes

Under the conditions defined by MAC, random access can require operation on the initial BWP, for example when the currently applicable uplink BWP does not provide the required random-access resources. MAC also contains BWP handling for consistent LBT failure in applicable deployments.³

These mechanisms should not be applied as generic explanations. First confirm that the relevant procedure occurred and that its conditions match the trace.

Evidence needed: random-access or failure procedure timeline, active direction before the event, BWP resources and the resulting state.

5. What each evidence layer can prove

Layer 1 - RRC configuration

Collect the complete serving-cell configuration, including:

- **initialDownlinkBWP;**
- **downlinkBWP-ToAddModList;**
- **firstActiveDownlinkBWP-Id;**
- **defaultDownlinkBWP-Id;**
- **bwp-InactivityTimer;**
- **uplinkConfig;**
- **initialUplinkBWP;**
- **uplinkBWP-ToAddModList;**
- **firstActiveUplinkBWP-Id;**
- subsequent add, modify or release operations.

RRC can prove what was configured and which state a reconfiguration intended to establish. It cannot, by itself, exclude a later DCI switch, timer expiry, random-access transition or newer RRC message.

What the decoded structure looks like

The abbreviated extract below is synthetic. It is an assembled decoder-style teaching view of the applicable common and dedicated serving-cell branches, not a captured operator message and not a verbatim ASN.1 schema.

```
ServingCellConfigCommon
  downlinkConfigCommon
    initialDownlinkBWP
      genericParameters
        locationAndBandwidth: 12925

ServingCellConfig
  downlinkBWP-ToAddModList.item(0)
    bwp-Id: 1
    bwp-Common.genericParameters
      locationAndBandwidth: 1099
  downlinkBWP-ToAddModList.item(1)
    bwp-Id: 2
    bwp-Common.genericParameters
      locationAndBandwidth: 13037
  firstActiveDownlinkBWP-Id: 1
  defaultDownlinkBWP-Id: 0
  bwp-InactivityTimer: ms100
```

For the RRC BWP field, TS 38.331 requires the RIV interpretation with **N_BWP = 275**.⁴ Applying the TS 38.214 formula and encoding the result back again gives:⁵

- **RIV 12925 → RB_start 0, length 48 PRBs** for the initial DL BWP;
- **RIV 1099 → RB_start 0, length 273 PRBs** for BWP 1;
- **RIV 13037 → RB_start 112, length 48 PRBs** for BWP 2.

The widths and locations are teaching values, not a recommended deployment. Here the initial BWP is narrow and BWP 1 is wider, so the decoded values agree with the scenario they illustrate.

Read the structure in two passes. The list proves that BWP 1 and BWP 2 are configured. The first-active field proves the intended DL state when this reconfiguration is applied. The default field identifies the applicable timer fallback as the initial DL BWP. None of those lines excludes a later DCI indication or another state-changing event.

Layer 2 - state-changing events

Build a chronological list:

Time	Direction	Event	Previous state	Resulting state	Evidence quality
10:42:18.000	DL/UL	RRC applied	initial BWP	BWP 1	decoded and timestamped
10:42:18.420	DL	DCI BWP indication	BWP 1	BWP 2	decoder-specific
10:42:18.540	DL	timer expiry	BWP 2	initial BWP, ID 0	reconstructed from the configured timer
10:42:18.610	UL	random access	unknown	initial UL BWP	procedure trace

If the tool does not expose one event, mark the interval as uncertain. Do not fill a trace gap with the configuration you expected to see.

Layer 3 - observed behaviour

Corroborate the reconstructed state with available evidence:

- PDCCH monitoring or scheduling context;
- PDSCH/PUSCH resource location;
- UE diagnostic active-BWP field;
- MAC or scheduler trace;
- configured bandwidth versus observed allocation;
- power-consumption or throughput change, only when measured under controlled conditions.

This layer is tool- and vendor-dependent. Document:

- decoder/tool name and version;
- log source;
- timestamp reference;
- UE and serving-cell identity;
- whether the active-BWP field is directly decoded or inferred;
- known blind spots.

6. Synthetic diagnostic case C-003

Status: pedagogical scenario, not a customer trace and not a measured performance result.

This case uses the same narrow-initial/wide-traffic relationship as the decoded teaching view above, but it is a separate minimal exercise and does not reuse that view's BWP 2 timeline.

Configuration

A conceptual n78 cell has a 100 MHz carrier. The UE configuration contains:

- BWP 0: a narrow initial DL/UL context;
- BWP 1: a wider traffic-oriented DL/UL context;
- **firstActiveDownlinkBWP-Id = 1;**
- **firstActiveUplinkBWP-Id = 1;**
- **defaultDownlinkBWP-Id = 0;**
- a configured DL inactivity timer.

The exact RB locations and timer value are intentionally omitted. They are not needed to teach the state error and must not be mistaken for a recommended production design.

Incomplete analyst conclusion

An RRC decode at **t0** shows BWP 1 configured and selected as first active. A throughput sample is taken later at **t3**.

The report states:

"The UE was using the wide BWP 1 during the low-throughput sample."

That conclusion is not supported by the RRC decode alone.

Missing timeline

Between **t0** and **t3**:

- 1 the reconfiguration makes BWP 1 active;
- 2 traffic becomes inactive;
- 3 the inactivity timer can expire;
- 4 the configured default DL BWP is BWP 0;
- 5 a later DCI or RRC event may switch again.

Without evidence for events 2-5, the active DL BWP at **t3** remains unproven. The UL state must also be analysed separately.

Corrected conclusion

“BWP 1 was configured and became first active when the reconfiguration was applied. The current log set does not prove the DL or UL BWP active at the later throughput sample. We need the intervening DCI/MAC/timer timeline or a synchronised UE active-BWP indication.”

What this case teaches

- A configured wide BWP does not prove wideband operation at the sample time.
- **firstActive...** describes a transition, not a permanent state.
- **defaultDownlinkBWP-Id** matters only when the applicable fallback event occurs.
- DL and UL conclusions need directional evidence.
- Low throughput cannot be attributed to a BWP until the active state and allocated resources are aligned with the same test interval.

The case does not claim that a narrow BWP caused low throughput. It shows that the analyst lacked the evidence required to test that hypothesis.

7. Common field mistakes

Mistake 1 - reading a list as a state

Weak: “The UE has four active BWPs.”

Better: “The UE has four configured DL BWPs. The active DL BWP at the selected instant is BWP 1,” or “the active DL BWP remains unknown for this interval.”

Mistake 2 - turning first active into always active

Weak: “**firstActiveDownlinkBWP-Id** = 1, so BWP 1 was active throughout the test.”

Better: “BWP 1 became active at the qualifying transition. Later switches must be excluded.”

Mistake 3 - treating default as initial in every configuration

Weak: “Timer expiry always returns to BWP 0.”

Better: “Timer expiry returns to the configured default DL BWP; when no default ID is configured, the initial DL BWP is used.”

Mistake 4 - using channel bandwidth as UE bandwidth

Weak: “This is a 100 MHz cell, so the UE was monitoring 100 MHz.”

Better: “The carrier is 100 MHz. The UE operating bandwidth depends on its active BWP at that instant.”

Mistake 5 - attributing throughput before checking allocation

Even a demonstrably active wide BWP does not guarantee that the scheduler allocated the UE the full width, high MCS, multiple layers or sustained time resources.

The correct chain is:

configured set → **active BWP** → **scheduled resources** → **successful delivery** → **application result**

Each arrow needs evidence.

8. Field workflow

Step 1 - freeze identity and time

Record the UE, cell, carrier, direction, timestamp and test interval. BWP state without a time reference is not operational evidence.

Step 2 - map the configured set

Create separate DL and UL tables with:

- BWP ID;
- frequency location and width;
- SCS and cyclic prefix where relevant;
- common and dedicated parameters;
- initial, first-active and default roles;
- add/modify/release history.

Step 3 - establish the first known state

Use the successful configuration procedure or another reliable starting event. State exactly what it proves.

Step 4 - replay every possible state change

Inspect:

- later RRC messages;
- DCI BWP indicators;
- SCell activation/deactivation context;
- inactivity timer behaviour;
- random-access events;
- relevant MAC failure procedures.

Step 5 - align performance

Only compare the BWP state with KPI or throughput evidence from the same interval. Then inspect allocations, MCS/code rate, rank/layers, retransmissions, load, TDD opportunity, transport and policy.

Step 6 - report uncertainty

Use one of three labels:

- confirmed - direct or fully reconstructed evidence;
- probable - partial evidence with stated assumptions;
- unknown - event visibility is insufficient.

Avoid "confirmed" when the decoder cannot expose the switching mechanism.

9. Minimum evidence package

Configuration

- ☐ release and software baseline;
- ☐ band, duplex mode and channel bandwidth;
- ☐ complete DL and UL BWP inventory;
- ☐ initial DL/UL BWP;
- ☐ first-active DL/UL identifiers;
- ☐ default DL identifier or confirmed absence;
- ☐ inactivity timer configuration;
- ☐ SCell and serving-cell context.

Timeline

- ☐ exact RRC application time;
- ☐ later reconfigurations;
- ☐ visible DCI BWP indications;
- ☐ timer start/restart/expiry evidence;
- ☐ random-access events;
- ☐ SCell activation state;
- ☐ identified trace gaps.

Observation

- ☐ active DL BWP at the test instant;
- ☐ active UL BWP at the test instant;
- ☐ proof type and decoder limitation;
- ☐ scheduled resource location and amount;
- ☐ radio quality, retransmission and MIMO context;
- ☐ cell load, TDD opportunity and transport context;
- ☐ known-good comparison.

Decision

- ☐ Is the claim about configuration or current state?
 - ☐ Does the evidence cover the interval between RRC and measurement?
 - ☐ Are DL and UL separated?
 - ☐ What would falsify the BWP hypothesis?
 - ☐ Is any performance effect measured rather than assumed?
-

Takwa Sebai working criterion

I call a BWP active only when the configured set, the last applicable state-changing event and an observed scheduling or device state agree for the same direction and interval. If one event is missing from the trace, the state is reported as unknown, not reconstructed from the desired design.

Chapter takeaway

A BWP configuration is a set of possible operating regions. Active BWP is a state on a timeline.

The engineer's sequence is:

- 1 separate carrier bandwidth from BWP width;
- 2 inventory configured DL and UL BWPs;
- 3 distinguish initial, first active and default roles;
- 4 reconstruct RRC, DCI, timer, activation and random-access events;
- 5 prove the active state at the exact test instant;
- 6 align allocation and performance evidence before assigning a cause.

Next chapter

A correct BWP state still does not explain where the UE detects and measures the cell. Chapter 4 moves to SSB, beam sweeping and the difference between cell coverage, beam measurement and serving-beam behaviour.

References

- ¹ ETSI, [TS 138 300 V18.9.0 - NR and NG-RAN Overall description; Stage-2](#), especially clause 6.10 for bandwidth adaptation and active BWP indication.
- ² ETSI, [TS 138 213 V18.8.0 - NR; Physical layer procedures for control](#), especially clause 12 for DL/UL BWP configuration, DCI indication and default-DL-BWP behaviour.
- ³ ETSI, [TS 138 321 V18.9.0 - NR; Medium Access Control protocol specification](#), especially clause 5.15 for MAC BWP operation and state changes.
- ⁴ ETSI, [TS 138 331 V18.9.0 - NR; Radio Resource Control protocol specification](#), especially clause 5.3.5.5.7 and the **ServingCellConfig**, **BWP-Downlink** and **BWP-Uplink** information elements.
- ⁵ ETSI, [TS 138 214 V18.9.0 - NR; Physical layer procedures for data](#), especially clause 5.1.2.2.2 for RIV encoding and decoding.

SSB and Beam Sweeping: Coverage Is Not One Beam

FIELD QUESTION

The cell-level result still looks strong. Does that prove the UE has a stable, useful beam for the traffic session?

5G NR FROM SITE TO TRACE • FIELD PLATE V-005

STRONG CELL, UNSTABLE BEAM


One route. Three evidence layers. No shortcut from “best SSB” to “serving data beam”.



FIELD SHORTCUT TO REJECT

Best measured SSB = current PDSCH beam

1 WALK THE ROUTE — KEEP LOCATION, TIME AND TRAFFIC SYNCHRONIZED



2 SEPARATE WHAT EACH LAYER PROVES

The field answer gets stronger as the layers are synchronized.



CELL RESULT • resultsSSB-Cell

A configured, filtered cell-level measurement output

Can hide transitions among individual SSB resources.

smooth summary

LAYER 1



PER-INDEX RESULT • resultsSSB-Indexes

What the UE measured on reported SS/PBCH resource

Best measured index is still not a data-beam proof.



LAYER 2



CONTROLLED / USED BEAM + DATA PATH

Lower-layer state, CSI-RS/TCI context, PDCCH/PDSCH evidence, BLER, HARQ, rank, MCS, allocation and throughput.

RRC is not required to know the beam used at every instant.

closest proof

LAYER 3



BEAM FAILURE IS A PROCEDURE, NOT ONE DIP

Configured BFD RS + repeated indications + threshold/timer + recovery evidence

3 SYNCHRONIZED SAMPLE EXTRACT — C-004 (SYNTHETIC TRAINING VALUES)

The cell curve barely moves from A to C; the per-SSB and traffic story does.

CELL SS-RSRP	A OPEN STREET	B APPROACH	C BEHIND CORNER	D INDOOR ENTRANCE
STRONGEST SSBs	SSB 1 -83 dBm SSB 2 -92	SSB 1 -85 dBm SSB 2 -86	SSB 2 -85 dBm SSB 1 -87	SSB 3 -97 dBm SSB 2 -101
SS-SINR	18 dB	10 dB	5-6 dB	1 dB
DL BLER	3%	8%	16-21%	29%
THROUGHPUT	goodput index 100	goodput index 82	goodput index 39-53	goodput index 14

CONCLUSION: cell result = summary • per-SSB result = measured resource • current data beam = separate lower-layer proof

TAKWA SEBAL - TAKWASEBAL.COM - V-005

Figure V-005. Original field interpretation map. The route, beam shapes, SSB-index positions and numerical values are synthetic. They explain the evidence layers; they do not represent an operator design or a universal radio threshold.

No. It proves much less than that.

An NR measurement report can contain a cell-level result and results for individual SS/PBCH block indexes. The lower layers can also manage the beam used at a given instant without RRC being required to know that beam. Meanwhile, a data session depends on more than synchronisation-signal measurements: control reception, scheduling, interference, rank, retransmissions and the beam applied to the relevant channels all matter.

This is where a familiar field shortcut becomes dangerous:

“SS-RSRP is good, so the beam is good.”

Before accepting that sentence, ask four questions:

- 1 Is the value a cell result or a result for one SSB index?
- 2 Which reference signal and reporting configuration produced it?
- 3 Is there evidence that the measured SSB corresponds to the beam used for the relevant control or data transmission?
- 4 Are the measurement, beam event and traffic symptom synchronised?

1. What an SS/PBCH block contains

The SS/PBCH block, usually shortened to SSB, combines four physical components:

- the Primary Synchronisation Signal, or PSS;
- the Secondary Synchronisation Signal, or SSS;
- the Physical Broadcast Channel, or PBCH;
- the demodulation reference signal associated with PBCH.

In the SSB resource mapping, PSS and SSS each occupy 127 subcarriers within their respective OFDM symbols; they do not occupy all 240 subcarriers of those symbols. PBCH spans three symbols, with PBCH and its demodulation reference signal mapped within the 240-subcarrier SS/PBCH block structure.¹

This combination supports essential early functions. The UE uses PSS, SSS and the PBCH demodulation reference signal in the cell-search process, then decodes the broadcast information needed to continue access.² The SSB is therefore both a synchronisation and broadcast structure and a radio reference for several measurement and beam-management procedures.

That does not make the SSB a complete representation of every channel experienced by the UE.

The PDCCH and PDSCH can use other reference signals, scheduling decisions and beamforming arrangements. A UE that can detect a strong SSB may still experience control-channel difficulty, interference, retransmissions or a data-beam problem. Conversely, a temporary fall in one SSB measurement is not by itself proof that the complete link has failed.

Field principle

Treat an SSB result as evidence about the measured synchronisation-signal resource. Do not silently promote it into evidence about every downlink channel.

2. SSB transmission and spatial direction

Several SSBs can occur within the defined half-frame structure. Their time locations depend on the configured subcarrier spacing, and the SSB periodicity is configured by the network.²

Different SSBs in a half-frame may be transmitted in different spatial directions. In a beamformed deployment, this allows SSBs to cover different parts of the cell area. The word “may” matters: the standard enables the behaviour but does not prove the antenna implementation, beam count or footprint of a particular live site.

An SSB index is therefore not automatically:

- a compass bearing;
- a fixed azimuth;
- one physical antenna element;
- the beam used for PDSCH at the same instant;
- a vendor-independent beam identity.

To draw a physical lobe for an SSB index on a customer map, the engineer needs approved design or configuration evidence. Without that mapping, the safe statement is:

“The UE reported SSB index 2 as the strongest configured SSB resource at this sample.”

It is not:

“The UE was served by the 70-degree traffic beam.”

The diagram in this chapter draws four conceptual lobes so that the measurement sequence is easy to understand. Those lobes are not a reconstruction of a real antenna pattern.

Quasi co-location is specific, not universal

The physical-layer specification defines assumptions for SS/PBCH blocks sharing the same index and centre frequency, including quasi co-location properties. Those rules help the UE interpret reference signals over time.¹ They do not justify treating every SSB index, every CSI-RS and every scheduled data transmission as the same physical beam.

Whenever a trace tool labels something simply as “beam,” identify the underlying object:

- SSB index;
- CSI-RS resource;
- transmission configuration indication state;
- beam-failure detection resource;
- vendor or chipset internal identifier;
- antenna design beam.

The label is not the proof. The decoded field and its configuration are.

3. Cell search, RRM measurement and beam management are different jobs

The same family of reference signals can appear in several procedures, but the procedures answer different questions.

Procedure	Practical question	Evidence to inspect
Cell search and synchronisation	Can the UE find the NR cell and obtain the required broadcast context?	PSS/SSS detection, PBCH DM-RS and PBCH decode
Radio resource management measurement	What cell-level and reference-signal-level radio results can the UE derive and report?	Measurement objects, report configuration, SS-RSRP/RSRQ/SINR, SSB indexes, filters
Beam management	Which transmit/receive spatial relationship should the lower layers use or evaluate?	Configured SSB/CSI-RS resources, lower-layer commands and beam reports where available
Beam failure detection and recovery	Has the configured beam-quality process met its failure condition, and what recovery follows?	BFD resources, beam-failure indications, threshold/counter/timer context, random access or BFR signalling
Traffic diagnosis	Did the UE receive useful control and data during the same interval?	PDCCH/PDSCH evidence, MCS, rank, BLER, HARQ, throughput and scheduler context

The mistake is to answer the last row using only the second.

For example, a drive-test plot can show a smooth cell-level SS-RSRP while per-SSB results alternate at a street corner. That pattern may be relevant to beam stability, but it is not enough to declare the root cause of poor throughput. The engineer still needs the traffic and lower-layer evidence.

4. Read SS-RSRP, SS-RSRQ and SS-SINR precisely

These three measurements are related but not interchangeable.

SS-RSRP: reference-signal received power

SS-RSRP is based on the linear average of the power contributions of resource elements carrying the SSS. Under the conditions defined by the measurement configuration, PBCH demodulation reference signals and configured CSI-RS resources may also be used. The measurement is made among signals associated with the same SSB index and physical-cell identity, and the network can restrict the set of SSBs to be measured.³

SS-RSRP is useful for signal-level comparison. It does not contain a complete description of interference, load, scheduling or data-channel performance.

SS-RSRQ: reference quality relative to total received power

SS-RSRQ is defined as:

$$N \times \text{SS-RSRP} / \text{NR carrier RSSI}$$

The numerator and denominator use the same resource-block set. NR carrier RSSI includes total received power from all sources over the relevant symbols, including serving and non-serving cells, adjacent-channel interference and thermal noise.³

Two positions can therefore have similar SS-RSRP and different SS-RSRQ because the total received environment is different.

SS-SINR: reference signal versus interference and noise

SS-SINR compares the average received power on the applicable reference-signal resource elements with the average interference-plus-noise contribution. Its exact measurement resources depend on the configuration; SSS resource elements form the core definition, and other defined resources can be involved under the specified conditions.³

SS-SINR is closer to a quality ratio than SS-RSRP, but it is still a reference-signal measurement. It is not a direct throughput forecast.

Do not manufacture universal thresholds

There is no single valid table that turns every SS-RSRP, SS-RSRQ or SS-SINR value into “good,” “acceptable” and “bad” across all devices, bands, bandwidths, antenna designs and services.

Thresholds used in one project can be useful operational criteria if their provenance is stated. They are not laws of NR. Before applying a threshold, record:

- the band and channel bandwidth;
- the SSB measurement configuration;
- device and logging mode;
- mobility and traffic state;
- indoor/outdoor context;
- project KPI definition;
- service objective;
- time and spatial aggregation.

The values in case C-004 are deliberately not coloured as universal pass/fail thresholds.

5. Three levels of evidence that must remain separate

Level 1 - Cell-level measurement result

RRC measurement results can contain a cell-level SSB result in **resultsSSB-Cell**.⁴ This is the value most likely to appear as a single smooth line on a summary plot.

The configuration and filtering matter. Release 18 measurement procedures allow cell-level results to be derived from beam measurements, and the exact Layer 1 filtering used in that derivation can be implementation-dependent.⁴ A cell curve can therefore hide transitions occurring among individual SSB resources.

Level 2 - Per-SSB-index measurement result

When configured, **resultsSSB-Indexes** can carry beam-level results based on SS/PBCH block resources. The report configuration determines whether those results are included, the quantities reported and the maximum number of reference-signal indexes.⁴

The abbreviated synthetic decode below shows why the two layers must not be collapsed. It is a teaching representation of decoder output, not a captured message or a verbatim ASN.1 schema.

```

MeasResultNR
  physCellId: 321
  measResult.cellResults.resultsSSB-Cell
    rsrp: 98
    rsrq: 51
    sinr: 63
  rsIndexResults.resultsSSB-Indexes.item(0)
    ssb-Index: 7
    ssb-Results: rsrp 101, rsrq 48, sinr 57
  rsIndexResults.resultsSSB-Indexes.item(1)
    ssb-Index: 11
    ssb-Results: rsrp 99, rsrq 46, sinr 54

```

The cell result is one reported object. The indexed results are separate objects. Neither object identifies the beam applied later to every PDCCH or PDSCH transmission.

These results answer a narrower question:

“What did the UE measure on the reported SSB resources?”

They still do not prove which spatial state was applied to every PDCCH or PDSCH transmission.

Level 3 - The controlled or used beam

Beam-level mobility can be handled by the lower layers within a cell. TS 38.300 explicitly notes that RRC is not required to know which beam is being used at a given instant.⁵

That one sentence has an important field consequence: an RRC trace can show the measurement configuration and reports while remaining insufficient to reconstruct the instantaneous data-beam state.

Depending on the network and tool access, stronger evidence may include:

- lower-layer beam commands or state logs;
- configured CSI-RS and transmission-configuration state;
- network-side beam or precoder telemetry;
- chipset diagnostics tied to the same timestamp;
- PDCCH/PDSCH decoding and performance indicators;
- antenna design mapping approved for the site.

Evidence rule

Best measured SSB index $\neq$ proven serving PDSCH beam.

The two may correlate in a particular deployment. Correlation must be demonstrated, not assumed.

6. Synthetic field case C-004 - a street corner that hides the transition

Evidence status: synthetic training case. No customer data, author drive-test file or live-network antenna pattern is represented.

Scenario

A UE follows a short n78 route:

- A - open street: clear exposure to the site;

- B - approaching a building corner: two SSB resources become comparable;
- C - behind the corner: the reported best SSB index alternates;
- D - indoor entrance: the dominant SSB is weaker and quality falls further.

A traffic download runs throughout the route. The logger exports one cell-level measurement, selected per-SSB results and traffic KPIs at synchronised samples.

Synthetic sample extract

Position	Cell SS-RSRP	Reported strongest SSB results	SS-SINR context	DL BLER / normalised goodput index
A	−82 dBm	SSB 1: −83 dBm; SSB 2: −92 dBm	18 dB	3% / 100
B	−83 dBm	SSB 1: −85 dBm; SSB 2: −86 dBm	10 dB	8% / 82
C1	−84 dBm	SSB 2: −85 dBm; SSB 1: −87 dBm	6 dB	16% / 53
C2	−84 dBm	SSB 1: −85 dBm; SSB 2: −86 dBm	5 dB	21% / 39
D	−96 dBm	SSB 3: −97 dBm; SSB 2: −101 dBm	1 dB	29% / 14

The values are synthetic teaching values, not measurements. The goodput index is normalised to position A = 100 because the exercise does not provide the PRBs, symbols, MCS table and layers needed to defend an absolute Mbit/s value. The “cell SS-RSRP” column represents the cell-level output exposed by the imaginary logger; the table does not claim that every implementation derives it by selecting the numeric maximum.

What the sample supports

At B and C:

- the cell-level SS-RSRP changes only slightly;
- two SSB indexes are close;
- the reported strongest index changes between consecutive samples;
- SS-SINR falls;
- BLER rises and throughput falls.

This is enough to justify a focused beam and interference investigation. It is not enough to conclude:

- that the beam switch itself caused the throughput loss;
- that SSB 1 and SSB 2 correspond to fixed azimuths drawn in the figure;
- that PDSCH used the same beam as the strongest reported SSB;
- that a beam-failure procedure occurred;
- that the network configuration is defective.

The next evidence request should be explicit: obtain the lower-layer beam state or the closest available proxy, the measurement configuration, the PDCCH/PDSCH condition, HARQ/BLER and scheduler context over the same samples.

Why the cell curve misleads

If an engineer plots only cell SS-RSRP, A to C appears almost stable: −82 to −84 dBm. The real change in this teaching case is not a 2 dB coverage loss. It is the combination of:

- near-equal candidate SSB results;
- rapid change in the reported strongest index;
- deteriorating signal-to-interference-and-noise context;
- worsening data-plane indicators.

The cell curve is not wrong. It is simply answering a broader and less diagnostic question.

7. Beam change is not beam failure

Normal beam management can change the preferred or used spatial relationship. A beam-failure procedure is a specific configured process.

For beam-failure detection, the gNB configures reference signals such as SSB or CSI-RS resources. The UE receives beam-failure instance indications from the physical layer. Failure is declared only when the configured count or threshold condition is reached within the applicable timer context. Recovery on the special cell can involve random access and selection of a suitable beam; secondary-cell recovery uses its defined MAC signalling path.⁶

Therefore:

- one weak SSB sample is not a beam failure;
- one best-index change is not a beam failure;
- a throughput dip is not a beam failure;
- a random-access event is not automatically beam-failure recovery.

To claim beam failure, collect:

- 1 the configured BFD reference resources;
- 2 the applicable threshold, counter and timer configuration;
- 3 the sequence of beam-failure indications or the closest approved decoded evidence;
- 4 the declaration event;
- 5 the recovery action and outcome;
- 6 the synchronised radio and traffic context.

If those layers are unavailable, report a suspected beam-quality instability, not a decoded beam-failure event.

8. A field workflow that survives trace gaps

Step 1 - Lock the serving context

Record:

- NR-ARFCN and band;
- PCI and serving-cell identity;
- SSB subcarrier spacing;
- SSB periodicity where available;
- measurement object and report configuration;
- configured SSB set;
- device, software and logger version.

Without this context, an index or value can be assigned to the wrong cell or measurement process.

Step 2 - Preserve cell and per-index results

Export both if available:

- cell-level SS-RSRP/SS-RSRQ/SS-SINR;
- per-SSB-index results;
- report timestamps;
- position and motion;
- missing-sample indicators.

Do not keep only "best beam." The second- and third-best resources often explain instability near a boundary.

Step 3 - Reconstruct the measurement rules

Check:

- which SSBs the UE was configured to measure;
- which quantity sorted the reported indexes;
- how many indexes could be reported;
- whether beam measurements were included;
- filtering and reporting intervals;
- whether the tool applied extra averaging.

A missing SSB result can mean "not reported under this configuration," not "not transmitted."

Step 4 - Add lower-layer beam evidence

Use the strongest approved evidence available. If the logger exposes only SSB measurements, say so. Do not rename the best SSB as the serving data beam to fill the gap.

Step 5 - Synchronise the traffic path

Align the beam investigation with:

- PDCCH decoding or control failures;
- MCS, rank and allocated resources;
- DL and UL BLER;
- HARQ retransmissions;
- throughput at a defined layer;
- application or transport limitation;
- network load where available.

The same radio event can have different user impact depending on scheduler and service state.

Step 6 - Classify the conclusion

Use one of four labels:

- measured: directly present in the decoded data;
- configured: present in RRC or network configuration;
- correlated: synchronised variables move together, but causality is not established;
- inferred: best engineering explanation from incomplete evidence.

That vocabulary prevents a good hypothesis from being reported as a proven root cause.

9. Evidence checklist

Before signing off a beam-related diagnosis, verify:

- ☐ Serving frequency, PCI and cell context are fixed.
- ☐ The SSB measurement and reporting configuration is preserved.
- ☐ Cell-level and per-SSB-index results are not mixed.
- ☐ The number and identity of reported indexes are interpreted using the configuration.
- ☐ No SSB index is assigned a physical azimuth without design evidence.
- ☐ "Best measured SSB" is not called "serving PDSCH beam" without lower-layer proof.
- ☐ Location, time, motion and traffic interval are synchronised.
- ☐ SS-RSRP is not used as a substitute for interference and quality.
- ☐ SS-SINR or SS-RSRQ is not used as a direct throughput prediction.
- ☐ Control, data, BLER, HARQ and scheduler evidence are checked.
- ☐ A beam-failure claim includes the configured detection and recovery sequence.
- ☐ Thresholds are labelled as project criteria, not universal NR rules.
- ☐ Synthetic values are visibly declared.
- ☐ Trace gaps remain uncertainties.

Takwa Sebai working criterion

I do not sign a beam root cause from a cell-level curve. The minimum evidence is the configured measurement resource, the per-resource result, the state-changing event or lower-layer beam evidence, and the control/data outcome on one timeline. Without that chain, I report an SSB observation and keep the physical cause open.

Chapter takeaway

Five statements are safe to carry into a field review:

- 1 A cell-level SSB result, a per-SSB-index result and the beam used by lower layers are different objects.
- 2 Different SSBs may use different spatial directions, but an SSB index is not a physical azimuth without deployment evidence.
- 3 Strong SS-RSRP does not prove low interference, reliable control reception or a useful PDSCH beam.
- 4 A best-index change is an observation; a beam-failure claim requires the configured detection and recovery procedure.
- 5 The diagnosis becomes defensible only when configuration, measurements, lower-layer state and traffic KPIs share the same timeline.

Stopping rule

When the evidence stops at SSB measurements, the conclusion must also stop at SSB measurements.

References

¹ ETSI, [TS 138 211 V18.9.0 - NR: Physical channels and modulation](#), especially clause 7.4.3.

² ETSI, [TS 138 300 V18.9.0 - NR and NG-RAN Overall description; Stage-2](#), especially clauses 5.2.4 and 5.2.5.3.

³ ETSI, [TS 138 215 V18.4.0 - NR: Physical layer measurements](#), especially clauses 5.1.1, 5.1.3 and 5.1.5.

⁴ ETSI, [TS 138 331 V18.9.0 - NR; Radio Resource Control protocol specification](#), especially clauses 5.5.3 and 5.5.5 and the **MeasResultNR** information element.

⁵ ETSI, [TS 138 300 V18.9.0](#), clause 9.2.3.1.

⁶ ETSI, [TS 138 300 V18.9.0](#), clause 9.2.8.

PDCCH to PDSCH: Control Before Data

FIELD QUESTION

If the radio conditions look good, can the UE still receive too little scheduling information to sustain the expected throughput?

5G NR FROM SITE TO TRACE • FIELD PLATE V-006

CONTROL BEFORE DATA


A dynamic PDSCH grant must survive five evidence gates before it becomes useful downlink data.



SHORTCUT TO REJECT

No DCI decoded = PDCCH was missed


DYNAMIC DOWNLINK EVIDENCE CHAIN



FROM THE PHYSICAL REGION TO A MONITORED CANDIDATE

Synthetic 48-RB × 2-symbol CORESET • simplified non-interleaved teaching view

ONE DOWNLINK SLOT



CORESET CAPACITY

48 RB × 2 symbols

= 96 REG

96 REG ÷ 6 REG/CCE

= 16 CCE



READ IT CORRECTLY

Candidate = monitoring location


Allocation = scheduler used CCEs

Decode = UE recovered valid DCI

Candidate count ≠ transmitted grants


CCE POOL

One PDCCH candidate occupies 1, 2, 4, 8 or 16 CCEs



WHAT THE ENGINEER MUST ALIGN

Same UE • same serving context • one synchronized timeline



1 RRC CONFIGURATION

active BWP • CORESET • Search Space • monitoring occasions

2 gNB SCHEDULER TRACE

UE selected? • candidate / AL • allocated, deferred or rejected?

3 UE CONTROL LOG

time / slot • DCI format • RNTI • decoded fields • CRC context

4 PDSCH RESULT

RBs • MCS • HARQ / CRC • transport block • delivered bytes

WITHOUT THESE FOUR LAYERS: CONTROL-PATH SUSPICION, NOT PROOF

SYNTHETIC FIELD COMPARISON — SAME STATIC UE, SAME RF CONDITION

The pattern localizes pressure in the control path; it does not by itself prove an over-the-air PDCCH miss.

QUIET WINDOW • 100 ms

6 / 16 CCE used*

12% AL8/16

0.8% deferred

61 DL DCI

4% PDSCH BLER

118 Mbit/s

BUSY WINDOW • 100 ms

15 / 16 CCE used*

56% AL8/16

18% deferred

24 DL DCI

5% PDSCH BLER

44 Mbit/s

RF nearly unchanged

* Synthetic vendor-defined occupancy example; validate the counter definition. Dynamic DL scope. Other CORESETs and scheduling constraints may exist.

TAKWA SEBAI - TAKWASEBAI.COM - V-006

Figure V-006. Original control-to-data evidence map. The CORESET dimensions, Search Space, candidate positions, DCI, counters and performance values are synthetic. The plate illustrates a method for dynamically scheduled downlink traffic; it is not an operator configuration.

Yes. But “the PDCCH is congested” is not a conclusion that can be extracted from RSRP, SINR or a falling throughput curve alone.

For dynamically scheduled downlink data, the UE first needs a valid control path. It must monitor the configured PDCCH candidates, detect a DCI addressed through the applicable RNTI context, interpret the assignment with the higher-layer configuration and then receive the scheduled PDSCH resources. A problem can occur at any point in that chain.

The field method is to separate six questions:

- 1 Configuration: where and when is the UE expected to monitor PDCCH?
- 2 Scheduling decision: did the gNB decide to schedule this UE?
- 3 Control allocation: did the scheduler allocate a PDCCH candidate and aggregation level?
- 4 Air-interface reception: did the UE decode the expected DCI?
- 5 Data allocation: what PDSCH resources and parameters did the DCI indicate?
- 6 Data outcome: did the UE receive useful data on those resources?

1. Why data can fail before the data channel

PDSCH carries downlink user data and other downlink information. PDCCH carries downlink control information. In the common dynamic-scheduling path, the DCI tells the UE which resources and transmission parameters apply to the PDSCH.

That dependency creates a practical asymmetry:

- the network can have PDSCH capacity available;
- the RF power measured from synchronisation signals can look stable;
- the scheduled PDSCH can decode well when it occurs;
- yet the UE can receive fewer useful scheduling assignments than expected.

This does not prove a PDCCH problem. A scheduler may intentionally give the UE fewer grants because of load, QoS, policy, buffer state or another resource constraint. The engineer must distinguish a decision not to schedule from a failed control allocation and a missed control transmission.

Field principle

“No DCI decoded” does not mean “DCI was transmitted and missed.”

This chapter focuses on dynamically scheduled downlink traffic. Semi-persistent scheduling and other configured mechanisms can create PDSCH occasions without a fresh scheduling DCI for every occurrence. The mechanism under investigation must therefore be identified before applying the workflow.

2. CORESET: the physical control container

A Control Resource Set, or CORESET, defines a region in frequency and time in which PDCCH can be mapped.

At the physical layer, a CORESET contains a configured number of resource blocks in frequency and one, two or three OFDM symbols in time.¹ For a dedicated RRC-configured CORESET, **frequencyDomainResources** uses a bitmap in which each bit represents a group of six resource blocks in the applicable BWP resource grid.² Do not reuse that origin rule blindly for CORESET 0: derive CORESET 0 from the applicable Type0-PDCCH common-search-space configuration and TS 38.213 tables.

The basic physical units are:

- REG - Resource Element Group: one resource block during one OFDM symbol;
- CCE - Control Channel Element: six REGs;
- PDCCH candidate: one possible set of CCEs on which the UE may monitor for DCI;
- aggregation level: the number of CCEs used by the candidate.

The supported PDCCH aggregation levels are 1, 2, 4, 8 and 16 CCEs.¹

A concrete calculation

The synthetic case in this chapter uses:

- 48 resource blocks in frequency;
- 2 OFDM symbols in time.

That produces:

$$48 \text{ RB} \times 2 \text{ symbols} = 96 \text{ REG}$$

and therefore:

$$96 \text{ REG} \div 6 \text{ REG per CCE} = 16 \text{ CCE}$$

Those sixteen CCEs are the physical control resource in that one synthetic CORESET and monitoring occasion. They are not the total control capacity of the cell. A UE can be configured with multiple CORESETs, candidates can overlap, different Search Spaces can use different monitoring occasions and an implementation can schedule control through other applicable resources.

The calculation is still useful because it makes aggregation pressure tangible:

- one AL16 candidate occupies all 16 CCEs;
- two non-overlapping AL8 candidates occupy all 16 CCEs;
- eight non-overlapping AL2 candidates also occupy all 16 CCEs.

The actual candidate locations and what a UE monitors are determined by the Search Space procedure and configuration, not by dividing the rectangle visually.

3. Search Space: where and when the UE looks

The CORESET defines physical resources. The Search Space defines the monitoring rules applied to PDCCH candidates associated with those resources.

The Release 18 **SearchSpace** information element can identify:

- the associated **controlResourceSetId**;
- monitoring-slot periodicity and offset;
- a monitoring duration when applicable;
- the first monitoring symbols within a slot;
- the number of candidates at each aggregation level;
- common or UE-specific Search Space type;
- the DCI formats to monitor.²

TS 38.213 defines the UE procedure in terms of Search Space sets. A set can be a Common Search Space, or CSS, or a UE-Specific Search Space, or USS.³

Common and UE-specific do different jobs

Common Search Spaces support control that must be found in common procedures and RNTI contexts, such as system information, paging or random access, depending on the Search Space type.

UE-Specific Search Spaces support DCI monitored for a connected UE using its applicable RNTI and configured formats. The exact combinations depend on the RRC configuration and procedure.

The field interpretation must keep three states separate:

State	What it means
Search Space configured	The UE has rules describing possible monitoring occasions and candidates
Monitoring occasion present	The configured time condition says the UE should monitor the applicable candidates
PDCCH transmitted for this UE	The gNB actually placed an encoded DCI in one candidate for the relevant RNTI

A configured Search Space does not mean that the gNB transmits DCI in every monitoring occasion.

Candidate count is not a grant count

nrofCandidates configures a number of PDCCH candidates per aggregation level.² These are possible decode locations. They are not a promise that the UE receives that many grants.

If a trace shows:

aggregationLevel8 = n2

the safe interpretation is:

“The UE can be required to monitor two configured candidates at AL8 under the applicable Search Space procedure.”

It is not:

“Two AL8 PDCCH transmissions occurred.”

4. PDCCH candidates and aggregation level

A PDCCH consists of one or more CCEs. Aggregation level **L** means that a candidate uses **L** CCEs.¹ The UE monitors the candidate set defined through its Search Spaces and applies the configured DCI and RNTI context.³

Aggregation level creates a field tradeoff:

- a larger aggregation level gives the encoded control message more physical resources;
- that can support a more robust control transmission for the same DCI payload;
- it also consumes more CCEs in the CORESET;
- fewer non-overlapping control messages can then fit in the same physical region.

The network's aggregation-level selection is implementation- and policy-dependent. Do not turn the tradeoff into a universal rule such as:

"Cell-edge UE always uses AL16."

A live selection can depend on control-channel quality estimates, payload size, DCI format, beam state, scheduler policy, available candidates and current control-resource pressure.

Candidate overlap matters

The sum of all configured candidate sizes is not the number of CCEs physically consumed. Candidate sets can overlap. Only transmitted PDCCH allocations consume the corresponding resources at a given occasion.

The UE also has defined limits on the number of PDCCH candidates and non-overlapping CCEs it is expected to monitor, depending on numerology, capability and configuration.³ Therefore, a troubleshooting sheet should record both:

- the network's configured candidate structure;
- the UE capability and active-cell context to which the monitoring limits apply.

5. DCI: the control message inside PDCCH

Downlink Control Information is the encoded payload carried by PDCCH. DCI formats serve different purposes.

For example, DCI format 1_0 is used to schedule PDSCH in one downlink cell. Its applicable fields include frequency-domain and time-domain resource assignment, modulation and coding information, HARQ process information, new-data and redundancy-version information, and feedback-related parameters depending on the RNTI and procedure.⁴

Formats such as 1_1, 1_2 and 1_3 provide other configured scheduling possibilities. A field guide should not list one fixed set of fields and call it "the DCI." The analyst must record:

- DCI format;
- Search Space type;
- RNTI used for CRC scrambling;
- serving and scheduling cell;
- active BWP;
- applicable RRC configuration;
- fields exposed by the decoder;
- decoder limitations.

The RNTI is part of detection

The DCI CRC parity bits are scrambled with the corresponding RNTI.⁴ This helps the UE determine whether a candidate contains control intended for the applicable context.

In a connected-mode synthetic case using C-RNTI, a successfully decoded DCI format 1_1 can be treated as evidence that:

- the UE detected a PDCCH candidate;
- the CRC check succeeded under the C-RNTI context;
- the decoded fields can be interpreted with the current configuration.

It does not prove that every intended PDCCH for the UE was decoded. Successful samples show success; they do not count invisible failures without another source of truth.

6. From DCI to the scheduled PDSCH

The UE combines the decoded DCI with higher-layer configuration to determine the PDSCH allocation.

For time-domain allocation, TS 38.214 defines how the applicable table is selected and how the indicated row provides parameters such as slot offset, start symbol, length and mapping type for the applicable case.⁵

For frequency-domain allocation, Release 18 supports downlink resource-allocation schemes including type 0 and type 1. The DCI field is interpreted using the active BWP and the configured resource-allocation mode.⁵

The control-to-data chain can therefore be represented as:

CORESET → Search Space → candidate → decoded DCI → time/frequency allocation → PDSCH attempt → HARQ outcome

Each arrow needs evidence.

Link	Strong evidence
CORESET → Search Space	Complete RRC configuration and active BWP
Search Space → candidate	Monitoring occasion, candidate index and aggregation level
Candidate → DCI	Network PDCCH allocation plus UE DCI decode when both are available
DCI → PDSCH resources	Decoded assignment fields interpreted with TS 38.214 and the active configuration
PDSCH → outcome	CRC/BLER, HARQ, MCS, rank, allocation and delivered bytes on the same interval

PDSCH performance cannot repair a missing scheduling opportunity. Conversely, a decoded DCI does not guarantee successful PDSCH reception.

7. Four different explanations for “no data arrived”

1. No scheduling decision

The gNB did not select the UE in that interval. Possible contexts include:

- no data waiting in the applicable buffer;
- another UE or service had scheduler priority;
- QoS or policy limited the UE;
- data resources were constrained;
- the UE was not in the expected active state.

There is no missed PDCCH if no DCI was created for the UE.

2. Control allocation deferred or rejected

The scheduler wanted to send DCI but could not allocate an applicable PDCCH candidate under the current CORESET, Search Space, aggregation-level and resource conditions.

This requires network-side scheduler or PDCCH-allocation evidence. A vendor-neutral label such as “control allocation deferral” must be mapped to the real counter definition before publication.

3. PDCCH transmitted, DCI not decoded

The gNB allocated and transmitted the PDCCH, but the UE did not produce the expected valid DCI decode.

This conclusion normally requires synchronised network and UE evidence. A UE-only gap is ambiguous because the UE log does not know about a control message it failed to detect.

4. DCI decoded, PDSCH failed or carried little data

The UE decoded the scheduling assignment, but the PDSCH outcome was poor or the allocation was small. Investigate:

- assigned time and frequency resources;
- MCS, layers and transport-block size;
- PDSCH DM-RS and beam context;
- CRC, BLER and HARQ retransmissions;
- scheduler allocation;
- transport and application path.

These four branches prevent the word “PDCCH” from becoming a generic explanation for low throughput.

8. Synthetic field case C-005 - good RF, fewer grants

Evidence status: synthetic training case. No customer trace, vendor counter or author drive test is represented.

Configuration used by the exercise

- SA connected-mode UE;
- n78, 100 MHz, 30 kHz SCS, TDD;
- active DL BWP containing a synthetic CORESET 1;
- PDSCH MCS table 1, scheduled MCS 24 and two layers in both windows;
- repeating **DDDSU** TDD model;
- one decoded dynamic DL assignment maps to one scheduled PDSCH slot in the exercise;
- CORESET 1: 48 RB × 2 symbols = 16 CCE;
- non-interleaved CCE-to-REG mapping for this exercise;
- UE-Specific Search Space 3, monitored every slot from the first symbol;
- synthetic candidate configuration: AL1 **n4**, AL2 **n4**, AL4 **n2**, AL8 **n2**, AL16 **n1**;
- DCI format 1_1 monitored with C-RNTI;
- continuous downlink traffic.

These values form one plausible teaching configuration. They do not describe an operator design recommendation.

Synchronised synthetic extract

Indicator	Quiet window	Busy window	What it suggests
UE SS-RSRP	−84 dBm	−83 dBm	Signal level is not the differentiator
UE SS-SINR	18 dB	17 dB	Reference-signal quality is broadly stable
Connected active UEs	9	45	Control demand is much higher
Mean CCE use in scoped CORESET	6 / 16	15 / 16	The physical control region is close to full in the busy window
Share of scoped PDCCH allocations using AL8/16	12%	56%	More CCEs are consumed per affected control message
gNB control-allocation deferrals	0.8%	18%	Network evidence supports control-resource pressure
UE DL DCI decodes per 100 ms	61	24	The observed UE receives fewer dynamic DL assignments
Scheduled PDSCH-slot share	30.5%	12%	The exercise maps one decoded assignment to one PDSCH slot
Mean PRBs when the UE is scheduled	190	180	Individual data grants remain substantial
Scheduled MCS / layers	24 / 2	24 / 2	The data-transmission contract is explicit
PDSCH BLER	4%	5%	Scheduled PDSCH remains comparatively healthy
Application throughput	118 Mbit/s	44 Mbit/s	Fewer scheduling opportunities reduce delivered data

All counter names and values are pedagogical. “Mean CCE use” is scoped to the synthetic CORESET and monitoring occasions defined by the exercise. It is not a standardised 3GPP KPI. The application values sit below a conservative PDSCH ceiling calculated from PRBs, MCS table 1, MCS 24, two layers, eleven PDSCH data-symbol equivalents after overhead and the stated initial BLER. The scheduled PDSCH-slot shares are measured against all wall-clock slots; they are therefore the row-specific **f_{DL}** values and are not multiplied by the **DDDSU** maximum again.

What proves the teaching diagnosis

The RF values alone do not prove anything about control pressure. The useful evidence is:

- 1 the same RRC configuration is active in both windows;
- 2 the gNB trace records much higher CCE use;
- 3 the same trace records attempted control allocations that were deferred;
- 4 the UE decodes fewer DL assignments;
- 5 the PDSCH is still healthy when scheduled;
- 6 throughput follows the drop in scheduling opportunities.

The synthetic conclusion is:

Control-resource pressure in the scoped CORESET contributed to fewer dynamic downlink assignments for the observed UE during the busy window.

It is not:

Good SINR proves the PDCCH was congested.

Why AL8 and AL16 matter here

The synthetic CORESET contains 16 CCE. If a scheduled control message uses AL16, it occupies the entire region for that mapping occasion. Two non-overlapping AL8 allocations can also occupy the region.

This does not mean the network should force lower aggregation. Lowering aggregation without sufficient control robustness can create decode failures and make performance worse. The engineering question is whether the CORESET, Search Spaces, aggregation policy and traffic mix match the intended coverage and capacity objective.

9. A field workflow for control-path diagnosis

Step 1 - Reconstruct the active configuration

Collect:

- active DL BWP;
- every applicable CORESET;
- frequency-domain bitmap and duration;
- CCE-to-REG mapping;
- Search Space IDs and associated CORESET IDs;
- monitoring periodicity, offset, duration and symbols;
- candidate count per aggregation level;
- CSS or USS type and monitored DCI formats;
- UE capability and cross-carrier scheduling context.

Do not analyse a stale RRC message if a later reconfiguration or BWP switch changed the active control context.

Step 2 - Define the scheduler question

Ask whether the gNB:

- had data or control to send;
- selected the UE;
- selected a DCI format and aggregation level;
- found an applicable candidate;
- allocated and transmitted the PDCCH;
- deferred or rejected the allocation.

Use the vendor's exact counter and trace definitions. "PDCCH congestion" is too vague for a final report.

Step 3 - Inspect UE control evidence

Where the tool exposes it, preserve:

- timestamp and slot;
- serving and scheduling cell;
- CORESET and Search Space;
- candidate or CCE information;
- aggregation level;
- DCI format;
- RNTI context;
- decoded scheduling fields;
- decode or CRC status if the tool exposes attempted candidates.

A decoder that logs only successful DCI cannot directly count missed PDCCH.

Step 4 - Follow the DCI into PDSCH

For the same assignment, derive:

- scheduled slot and symbols;
- allocated RBs;
- MCS and redundancy version;
- HARQ process and new-data indicator;
- layers or rank evidence where available;
- transport-block size;
- CRC and HARQ result;
- delivered bytes and throughput.

Step 5 - Compare a known-good interval

Keep constant as much as possible:

- UE and software;
- cell, band, BWP and beam context;
- traffic server and application;
- route or static location;
- service and QoS;
- measurement duration.

Then compare configuration, control demand, aggregation, allocation result, DCI rate and PDSCH outcome.

Step 6 - Label the conclusion

Use one of these:

- not scheduled - network decision is proven;
- control allocation deferred - gNB attempt and failure/deferral are proven;
- PDCCH decode failure - transmission and UE non-detection are synchronised and proven;
- PDSCH failure after valid DCI - control succeeded and data failed;
- control-path suspicion - available evidence is incomplete.

10. Evidence checklist

Before signing off a PDCCH-related diagnosis:

- ❑ The active DL BWP and complete current RRC configuration are known.
- ❑ CORESET frequency resources, duration and CCE-to-REG mapping are decoded.
- ❑ Search Space monitoring occasions and symbols are reconstructed.
- ❑ Candidate counts are not reported as transmitted grants.
- ❑ Common and UE-specific Search Spaces are not mixed.
- ❑ DCI format and RNTI context are identified.
- ❑ Aggregation level is tied to a real allocation or decode, not inferred from coverage.
- ❑ “No UE DCI” is not called a decode failure without network evidence.
- ❑ Scheduler decision, control allocation and UE reception remain separate.
- ❑ DCI fields are interpreted with the active BWP and current configuration.
- ❑ The resulting PDSCH resources and HARQ outcome are synchronised.
- ❑ Load, QoS, buffer, policy and data-resource constraints are checked.
- ❑ UE monitoring capability and multi-cell context are considered.
- ❑ Counter definitions and aggregation intervals are recorded.
- ❑ Synthetic values remain visibly declared.

Takwa Sebai working criterion

I call an interval control-resource limited only when network-side allocation evidence shows the pressure or deferral, the UE receives fewer applicable assignments, and PDSCH remains healthy when it is scheduled. A UE-only gap in decoded DCI proves only that fewer DCI were decoded.

Chapter takeaway

The five objects to keep separate are:

- 1 configured CORESET resources;
- 2 configured Search Space candidates;
- 3 the scheduler’s decision and control allocation;
- 4 the UE’s successful DCI detection;
- 5 the resulting PDSCH allocation and data outcome.

Stopping rule

If the only evidence is fewer DCI decodes at the UE, the conclusion is “fewer DCI decodes observed” - not “PDCCH transmissions were missed.”

References

¹ ETSI, [TS 138 211 V18.9.0 - NR: Physical channels and modulation](#), especially clauses 7.3.2.1 and 7.3.2.2.

² ETSI, [TS 138 331 V18.9.0 - NR: Radio Resource Control protocol specification](#), especially the **ControlResourceSet** and **SearchSpace** information elements.

³ ETSI, [TS 138 213 V18.8.0 - NR: Physical layer procedures for control](#), especially clause 10.1.

⁴ ETSI, [TS 138 212 V18.8.0 - NR: Multiplexing and channel coding](#), especially clauses 7.3.1.2 and 7.3.2.

⁵ ETSI, [TS 138 214 V18.9.0 - NR: Physical layer procedures for data](#), especially clauses 5.1.2 and 5.1.3.

Read Radio Measurements Together

FIELD QUESTION

What does "good coverage" mean when RSRP, SINR, BLER and throughput disagree?

5G NR FROM SITE TO TRACE • FIELD PLATE V-007

ONE ROUTE, ONE CLOCK

Level, quality, CSI, scheduling and outcome must be aligned before the route tells a defensible story.




SYNCHRONIZATION RULE

Same UE • resources • direction • time basis

SYNTHETIC ROUTE C-006

Continuous DL traffic • SA • n78 / 100 MHz / 30 kHz • one-second aligned windows



THE SAME TEN FIELDS AT EVERY POINT

Relative pattern only — values are not acceptance thresholds

A OPEN STREET	-82 dBm	RSRP	20 dB	SINR	CQI 13	BLER 6%	270 Mb/s
MCS 24 • rank-2 78% • 210 PRB							
B CORNER / OVERLAP	-78 dBm	RSRP	4 dB	SINR	CQI 6	BLER 23%	30 Mb/s
MCS 10 • rank-2 6% • 180 PRB							
C BUILDING ENTRANCE	-91 dBm	RSRP	7 dB	SINR	CQI 8	BLER 16%	42 Mb/s
MCS 13 • rank-2 14% • 150 PRB							
D INDOOR WINDOW	-97 dBm	RSRP	12 dB	SINR	CQI 10	BLER 9%	78 Mb/s
MCS 18 • rank-2 55% • 130 PRB							

SHARED TIMELINE — EACH LANE USES THE SAME FOUR WINDOWS

A visual correlation is not a root cause: the red window identifies where the evidence chain must be opened.

ALIGNED SOURCE	A • 00-08 s	B • 08-16 s	C • 16-24 s	D • 24-32 s
SS-RSRP • dBm	-82	-78 ↑	-91	-97
SS-SINR • dB	20	4 ↓	7	12
CQI / MEAN MCS	13 / 24	6 / 10	8 / 13	10 / 18
RANK-2 SHARE	78%	6%	14%	55%
INITIAL DL BLER	6%	23% ↑	16%	9%
MEAN ALLOCATED PRB	210	180	150	130
APP THROUGHPUT	270 Mb/s	30 Mb/s ↓	42 Mb/s	78 Mb/s

Synthetic case C-006 • SS-based values • initial PDSCH BLER • mean PRBs only when scheduled • no real interferer or operator route represented.

TAKWA SEBAI • TAKWASEBAI.COM • V-007

Figure V-007. Original synchronised route plate. The site, street, building, timestamps, route and all measurements are synthetic. Relative changes illustrate the method; values are not acceptance thresholds.

It means the phrase is too imprecise for a diagnosis.

A radio level can be strong while interference or another quality limitation makes data reception inefficient. A clean reference-signal measurement can coexist with low throughput because the UE receives few resources. A high CQI can be followed by a lower scheduled MCS because the report and the transmission do not describe the same instant, resources or assumptions.

The field method is to separate five evidence layers:

- 1 measurement context - UE, direction, cell, BWP, reference resources, time and filtering;
- 2 reference-signal measurements - SS-RSRP, SS-RSRQ, SS-SINR or the corresponding CSI-RS quantities;
- 3 UE channel-state report - CQI, RI, PMI and the report configuration that produced them;
- 4 network transmission decision - scheduled PRBs, MCS, layers, grant frequency and HARQ process;
- 5 outcome - CRC/BLER, retransmissions, delivered radio bytes and application throughput.

ONE KPI DOES NOT CLOSE THE CASE

Configuration, measurement, UE report, network decision and outcome are related — never interchangeable.




SHORTCUT TO REJECT

"RSRP is good, therefore RF is good."

FIVE EVIDENCE LAYERS — READ LEFT TO RIGHT

Each layer answers a different question. A valid conclusion states where the evidence stops.



DEFENSIBLE CHAIN

resource + time → measurement → CSI report → scheduled PDSCH → BLER / HARQ → delivered data

STOPPING RULE

If only one KPI supports the diagnosis, the case is not closed.

FOUR RELATIVE PATTERNS — NO UNIVERSAL “GOOD / BAD” THRESHOLDS

Compare with the same UE, configuration, direction, traffic demand and a known-good interval.

1 • STRONG LEVEL, POOR QUALITY

-78 dBm RSRP • 4 dB SINR

Level exists; quality and data degrade.

DOES NOT PROVE

which interferer, beam or receiver effect
created the interference-plus-noise term.

NEXT CHECK

SS/CSI resources • neighbour activity
beam state • scanner / network evidence

2 • WEAK LEVEL, CLEANER QUALITY

-97 dBm RSRP • 12 dB SINR

The link can be cleaner than a stronger point.

DOES NOT PROVE

unlimited margin, healthy uplink or robust
mobility under load and orientation change.

NEXT CHECK

BLER under load • UL power headroom
mobility • beam / edge stability

3 • GOOD RADIO, POOR DATA

CQI 12 • BLER 7% • 7 Mbit/s

Scheduled link outcome remains healthy.

DOES NOT PROVE

that scheduler load, QoS or transport is
the root cause without allocation evidence.

NEXT CHECK

PRBs • scheduled-slot share • buffer / QoS
active UEs • server and transport path

4 • LEVEL, QUALITY AND DATA FALL

RSRP ↓ • SINR ↓ • BLER ↑ • data ↓

A radio limitation becomes plausible.

DOES NOT PROVE

one unique cause such as coverage, beam,
interference or a physical site fault.

NEXT CHECK

serving/beam change • interference • UL
hardware • geometry • known-good route

Synthetic teaching values • prefixes, direction, resources, sample population and counter definitions must be preserved in a real case.

TAKWA SEBAI • TAKWASEBAI.COM • V-008

Figure V-008. Original multi-layer correlation map. Every value and scenario is synthetic. The plate shows which evidence can support a diagnosis and which shortcuts must be rejected; it does not provide universal KPI thresholds.

1. Power is not quality

The first mistake in many field reports is to use RSRP as a synonym for “radio quality.”

SS-RSRP: reference-signal level

TS 38.215 defines SS-RSRP as the linear average of the power contributions of resource elements carrying the secondary synchronisation signal. Depending on the configured measurement use, PBCH DM-RS and indicated CSI-RS may also contribute under the specified conditions.¹

SS-RSRP therefore answers a scoped question:

How much power from the applicable synchronisation/reference resources was received under this measurement configuration?

It does not directly contain:

- total interference power;
- scheduled bandwidth;
- transport-block size;
- the number of active users;
- the scheduler’s priority decision;
- application-server performance.

A strong SS-RSRP can support coverage, cell detection and beam measurements. It cannot, by itself, prove that a loaded PDSCH will have a good error rate or high throughput.

SS-RSRQ: reference power relative to received wideband power

SS-RSRQ is defined as:

$N \times \text{SS-RSRP} / \text{NR carrier RSSI}$

where N is the number of resource blocks in the NR carrier RSSI measurement bandwidth, and numerator and denominator use the same resource-block set.¹

The RSSI term includes power from all sources observed in the applicable symbols and bandwidth, including serving and non-serving co-channel cells, adjacent-channel interference and thermal noise.¹

That makes RSRQ useful when level alone is misleading. But it still must be read with:

- the exact SS or CSI prefix;
- measurement bandwidth and resources;
- symbols and time window;
- serving-frequency context;
- UE implementation and filtering;
- load and transmission activity during the sample.

SS-SINR: wanted reference-signal contribution versus noise and interference

SS-SINR is defined as the linear average power contribution of resource elements carrying the secondary synchronisation signal divided by the linear average noise and interference contribution.¹

It describes quality more directly than RSRP, but the word “SINR” is not enough. The analyst must identify:

- SS-SINR or CSI-SINR;
- L1 or higher-layer reported value;
- channel-measurement resource;
- interference-measurement resource, if configured;
- report timing and filtering;
- beam or resource index;
- serving cell and BWP.

Field principle

A dB value without its resource, time and reporting context is not a complete measurement.

Do not mix SS and CSI measurements silently

TS 38.215 defines both SS-based and CSI-RS-based RSRP, RSRQ and SINR quantities. Their reference resources and applicable procedures differ.¹

If a UE log displays **RSRP**, **SINR** or **quality** without the prefix, preserve the tool's exact label and verify its documentation. Do not place SS-RSRP from mobility measurement beside CSI-SINR from a CSI report and describe them as one simultaneous physical snapshot unless the timestamps and resources justify that comparison.

2. The same timestamp is not always the same radio instant

Drive-test and trace tools often merge several sources:

- GNSS position;
- UE diagnostic messages;
- decoded RRC;
- L1 measurement reports;
- MAC scheduling records;
- network counters;
- application throughput samples.

Those sources can have different clocks, reporting periods, averaging windows and latencies.

A row labelled **12:05:17** can combine:

- a GNSS fix generated at the start of the second;
- an SS-RSRP value filtered over earlier samples;
- a CSI report based on a defined CSI reference resource;
- MCS averaged only across slots in which the UE was scheduled;
- BLER calculated over a short or long rolling window;
- application throughput averaged over the full second.

This is why chart alignment must be treated as engineering work, not as a cosmetic export.

Build a timing dictionary

Before correlating KPIs, record:

Source	Time basis to establish
GNSS / route	fix time, update rate and interpolation method
RRC	message receive time and activation time of the configuration
SS measurements	SMTC/report timing, L1/L3 filtering and report trigger
CSI	report configuration, periodicity/trigger and CSI reference resource
Scheduler	slot, carrier, BWP, UE, HARQ process and averaging rule
BLER	initial or residual, DL or UL, numerator, denominator and window
Throughput	radio, PDCP/RLC or application layer; active-time or wall-clock denominator

If two sources cannot be synchronised closely enough for the claimed mechanism, mark the relationship as a correlation or suspicion.

3. CQI is a report under assumptions, not the scheduled MCS

Channel Quality Indicator is part of the CSI framework. TS 38.214 allows CSI reporting configurations that can include CQI, RI, PMI, CRI and related quantities.²

The dependencies are explicit:

- CQI is calculated conditioned on reported PMI, RI and CRI;
- PMI is calculated conditioned on reported RI and CRI;
- RI is calculated conditioned on reported CRI.²

The report can be periodic, semi-persistent or aperiodic, using configured channel and interference measurement resources.²

For each reported CQI, the UE selects the highest applicable CQI index satisfying the specified transport-block error probability condition for the configured CQI table and CSI reference resource. The common CQI tables use a target probability of 0.1, while the low-BLER table uses 0.00001.²

That definition has practical consequences.

CQI can be stale relative to a transmission

The CQI may represent an earlier reference resource. Between the CSI reference resource and a later PDSCH:

- interference can change;
- the beam can change;
- the UE can move;
- another set of PRBs can be scheduled;
- the rank or precoder can differ;
- implementation-specific outer-loop adaptation can act.

CQI table matters

CQI index 12 is not a universal modulation-and-code-rate statement without the configured CQI table. TS 38.214 contains several CQI tables with different entries and purposes.²

CQI is input, MCS is a transmission decision

MCS tables map a selected MCS index to modulation order and target code rate for the applicable PDSCH procedure.² The selected MCS is part of the actual transmission decision. A scheduler may use CQI among several inputs, but the specification does not define one universal field rule:

$$\text{MCS} = \text{CQI} + \text{constant}$$

Such a shortcut ignores the configured tables, report age, rank, target error behaviour, scheduler policy, HARQ state, UE history and implementation.

Safe interpretation

CQI describes what the UE reported under a configured CSI hypothesis. MCS describes what the network selected for a particular transmission.

4. Rank report and scheduled layers are related, not identical

RI is the rank indicator in CSI. It is calculated under the configured CSI reporting assumptions and conditioned on the reported channel-resource indicator.²

RI can help answer:

How many spatial layers did the UE report as supportable or preferable under this CSI configuration?

It does not, alone, prove how many layers were used on every PDSCH.

The actual transmission can be affected by:

- UE capability;
- available antenna ports and CSI-RS design;
- reported PMI/RI and their age;
- configured transmission mode;
- scheduler policy;
- traffic-buffer size;
- number of available PRBs;
- control and data resource pressure;
- rank adaptation and link history;
- beam, mobility and interference changes.

Therefore preserve separate fields for:

- reported RI;
- scheduled number of layers;
- rank-one/rank-two sample share;
- codeword count where exposed;
- MCS per codeword where exposed;
- actual transport-block result.

Do not label an RI value as "MIMO throughput" or interpret rank two as a guarantee that throughput doubles.

5. MCS, BLER and HARQ: read the attempt and the outcome

The applicable MCS index and MCS table determine modulation order and target code rate for a scheduled PDSCH. Transport-block size also depends on scheduled resources, layers and other parameters in the TS 38.214 procedure.²

MCS alone therefore cannot explain throughput.

Two intervals can have the same mean MCS but very different:

- scheduled PRBs;
- number of scheduled slots;
- number of layers;
- transport-block size;
- retransmission volume;
- application demand.

Define BLER before using it

Field tools use several BLER-like quantities. Before publishing a number, record:

- DL or UL;
- PDSCH, PUSCH or another channel;
- initial-transmission BLER or residual error after HARQ;
- transport-block or code-block basis;
- first transmission only or all transmissions;
- sample count;
- averaging window;
- behaviour when no transmission occurred.

A BLER percentage from ten transport blocks is not as stable as the same percentage from ten thousand.

HARQ retransmission is evidence of recovery work

HARQ can recover unsuccessful transport-block attempts. More retransmissions can consume time-frequency resources and reduce delivered throughput even when the final application session remains connected.

But retransmissions still need context:

- were they concentrated on one beam or route segment?
- did MCS fall before or after BLER rose?
- did rank collapse at the same time?
- did the scheduler reduce allocation?
- did the traffic buffer remain full?

Throughput is an end result, not a radio measurement

Application throughput can be limited outside the radio link:

- traffic server;

- TCP congestion control;
- latency or packet loss in transport;
- device processing;
- thermal state;
- QoS or policy;
- traffic demand;
- NSA anchor or aggregation state;
- scheduler allocation.

For a radio diagnosis, correlate at least:

scheduled slots × allocated resources × layers × transport efficiency – retransmission cost

with delivered bytes and application throughput. This is an engineering decomposition, not a replacement for the exact transport-block calculation.

6. Synthetic field case C-006 - stronger level, worse data

Evidence status: synthetic training route. No customer trace, author drive test or real interference map is represented.

Exercise configuration

- SA connected-mode UE;
- n78, 100 MHz, 30 kHz SCS, TDD;
- PDSCH MCS table 1;
- repeating **DDDSU** pattern model;
- eleven PDSCH data-symbol equivalents after overhead per eligible slot;
- the common **f_{DL} = 0.70** wall-clock downlink factor and each row's stated initial-BLER factor for the ceiling check;
- continuous downlink traffic;
- one serving cell and fixed test device;
- four route points sampled over the same application session;
- metrics aggregated in one-second windows;
- **rank-2 share** means the share of scheduled PDSCH samples with two layers in this exercise;
- **DL BLER** means synthetic initial PDSCH transport-block BLER.

Synchronised synthetic extract

The extract is split into two tables so that route and delivery fields remain readable in print.

Point	Route context	SS-RSRP	SS-RSRQ	SS-SINR
A	Open street	−82 dBm	−10 dB	20 dB
B	Reflective corner / overlap zone	−78 dBm	−17 dB	4 dB
C	Building entrance	−91 dBm	−15 dB	7 dB
D	Indoor window position	−97 dBm	−12 dB	12 dB

Point	CQI	Mean MCS	Rank-2 share	DL BLER	Mean PRBs	App goodput
A	13	24	78%	6%	210	270 Mbit/s
B	6	10	6%	23%	180	30 Mbit/s
C	8	13	14%	16%	150	42 Mbit/s
D	10	18	55%	9%	130	78 Mbit/s

For the arithmetic check, **rank-2 share** is converted to mean layers as **1 + share**, because all remaining scheduled samples use one layer in this exercise. Each goodput value is below the corresponding conservative physical ceiling. The ceiling is a sanity check, not a prediction of scheduler or application performance.

What the table proves

Within the synthetic exercise:

- the highest SS-RSRP occurs at point B;
- the worst SS-RSRQ and SS-SINR also occur at point B;
- CQI, scheduled MCS and rank-two use fall;
- initial BLER rises;
- throughput falls even though the reference-signal level is stronger than at A.

The defensible conclusion is:

Point B combines strong reference-signal level with poor measured quality and poor data outcome.

The table alone does not prove which transmitter, reflection, beam state or receiver limitation created the interference-plus-noise term. That root cause needs:

- synchronised serving and neighbour activity;
- beam/resource evidence;
- interference measurement configuration;
- scanner or network-side evidence where available;
- a discriminating test such as time, beam, frequency or cell comparison.

Why point D matters

Point D has the weakest SS-RSRP, but better SS-SINR, CQI, MCS, BLER and throughput than B.

This does not make D universally “good.” It shows why a level threshold alone can rank the two points incorrectly for the observed data session.

7. Synthetic field case C-007 - good radio, few resources

Evidence status: synthetic static comparison. No real scheduler counter or operator load data is represented.

The UE remains at the same point with the same serving cell, BWP, beam context, device and traffic server. Two 20-second windows are compared. The same PDSCH MCS table 1 and **DDDSU** pattern model are used. Scheduling duty cycle is expressed as scheduled DL slot-equivalents against all 30 kHz slots in

the wall-clock interval. The special slot is not counted as a full DL slot. The quiet-window value therefore remains below the common $f_{DL} = 0.70$ maximum and is not multiplied by 0.70 a second time.

Indicator	Quiet window	Busy window	Interpretation inside the exercise
SS-RSRP	−86 dBm	−85 dBm	Level is stable
SS-SINR	18 dB	19 dB	Reference-signal quality is stable
Wideband CQI	12	12	UE CSI report is stable
Mean scheduled MCS	21	20	Scheduled link adaptation remains similar
Rank-2 share	74%	72%	Layer use is similar
Initial DL BLER	6%	7%	Scheduled PDSCH remains healthy
Mean allocated PRBs when scheduled	204	31	Resource size falls sharply
Scheduling duty cycle	68%	19%	UE is scheduled much less often
Active UEs in scoped cell	8	54	Demand is much higher
Application throughput	205 Mbit/s	7 Mbit/s	Delivered data follows allocation and duty cycle

The synthetic diagnosis is not:

“SINR is wrong because throughput is low.”

It is:

“Radio quality and scheduled-link outcome remain broadly stable, while the UE receives fewer and smaller allocations during the busy window.”

To call this a scheduler/capacity limitation in a real report, preserve the network-side proof:

- full-buffer demand for the test UE;
- cell and BWP utilisation;
- UE selection and scheduling records;
- PRBs and scheduled-slot share;
- QoS/5QI and policy;
- active UE mix;
- control-resource status;
- known-good window;
- transport/server health.

Good SINR removes one hypothesis. It does not identify the remaining cause by itself.

8. Four combinations that prevent one-KPI diagnosis

The labels below are relative to the same configuration and known-good reference. They are not universal ranges.

Relative pattern	What it says	What it does not prove	Next discriminating checks
Strong level, poor quality, poor data	Wanted reference power exists, but quality and outcome degrade	Which source caused interference or whether the beam is wrong	SS/CSI resource context, neighbour activity, beam state, scanner/network evidence
Weak level, cleaner quality, stable data	Low reference power coexists with a cleaner link and usable adaptation	That there is unlimited margin or robust mobility	BLER under load, UL, mobility, power headroom, edge stability
Good level/quality, poor data	Reference measurements do not explain the throughput loss	That the scheduler is definitely at fault	allocation, grant rate, load, QoS, transport, server, buffer
Level, quality and data all degrade	Radio limitation is plausible	One unique RF root cause	serving/beam change, interference, hardware, coverage geometry, UL/DL split

This matrix is useful because it forces the report to state what evidence is still missing.

9. Field workflow: from chart to defensible conclusion

Step 1 - Freeze the measurement context

Record:

- UE model, modem and software;
- capability relevant to bands, MIMO and modulation;
- SA or NSA and the LTE anchor if applicable;
- serving and scheduling cell;
- band, carrier bandwidth, active BWP and SCS;
- SS or CSI reference resource;
- beam/resource index;
- DL or UL direction;
- route, speed and device orientation;
- traffic application, server and demand.

Step 2 - Decode the reporting configuration

For radio measurements and CSI, preserve:

- measurement object and report configuration;
- SMTC where relevant;
- configured SS/PBCH blocks or CSI-RS resources;
- channel and interference measurement resources;
- CQI table;
- report quantity;
- periodic, semi-persistent or aperiodic behaviour;
- report periodicity or trigger;
- filtering and tool aggregation.

Step 3 - Normalise time and sample population

Do not average unlike populations.

Examples:

- RSRP may exist during idle measurement periods;
- MCS exists only when a transmission is scheduled;
- rank share is defined over scheduled samples;
- BLER requires transport-block attempts;
- throughput can include unscheduled wall-clock time.

Record both the value and the number of contributing samples.

Step 4 - Read the chain in order

For each interval ask:

- 1 What reference-signal level and quality were measured?
- 2 What CSI did the UE report, on which resources and when?
- 3 What did the scheduler allocate?
- 4 Which MCS and layers were actually used?
- 5 What was the initial transmission result?
- 6 How much retransmission work occurred?
- 7 How many bytes reached the radio and application layers?

Step 5 - Compare with a known-good interval

Keep constant as much as possible:

- UE and software;
- cell, band, BWP and beam context;
- traffic server and protocol;
- static position or repeatable route;
- service/QoS;
- test duration and sample rule.

Change one useful dimension:

- quiet versus busy;
- open street versus corner;
- one beam versus another;
- before versus after a confirmed change;
- serving cell versus known-good cell.

Step 6 - Label every claim

Use:

- measured - directly recorded by the identified tool;
- configured - decoded from the active configuration;
- scheduled - shown in network or UE scheduling evidence;
- observed outcome - CRC, HARQ, delivered bytes or throughput;
- correlated - synchronised values move together;
- inferred - mechanism is plausible but not directly proven.

Step 7 - Stop at the evidence boundary

Examples of defensible conclusions:

- "SS-RSRP improved while SS-SINR and initial PDSCH BLER degraded."
- "The UE reported CQI 12, while the network scheduled mean MCS 20 under the active table."
- "The PDSCH remained healthy when scheduled, but scheduled-slot share fell."
- "The available evidence supports an interference suspicion; the interferer is not identified."

Examples to reject:

- "RSRP is good, therefore RF is good."
- "CQI 12 means MCS must be 20."
- "Rank two means throughput must double."
- "BLER is low, therefore the server path is healthy."
- "SINR is good, therefore low throughput is a scheduler fault."

10. Evidence checklist

Before publishing a multi-KPI diagnosis:

- ☐ The same UE and software are used in the compared intervals.
- ☐ Serving and scheduling cell, carrier, BWP and beam context are known.
- ☐ DL and UL are not mixed.
- ☐ SS-based and CSI-RS-based quantities are not mixed silently.
- ☐ Measurement resources, timing, filtering and report configuration are recorded.
- ☐ CQI table and CSI report quantity are known.
- ☐ CQI is not presented as the scheduled MCS.
- ☐ Reported RI is not presented as actual layers for every PDSCH.
- ☐ Scheduled MCS, layers, PRBs and slot share use the same interval.
- ☐ BLER direction, channel, initial/residual definition, sample count and window are known.
- ☐ HARQ retransmissions and delivered bytes are synchronised.
- ☐ Traffic demand and buffer state are established.
- ☐ Load, active users, QoS and policy are checked.
- ☐ Application server, transport path and device constraints are checked.
- ☐ A known-good comparison exists.
- ☐ Synthetic values remain visibly declared.
- ☐ The conclusion states what is measured, correlated and inferred.

Takwa Sebai working criterion

I do not correlate two radio KPIs until UE, cell, carrier, BWP, direction, reference resource, clock, window and sample population match. I do not call the correlation causal until a controlled comparison changes one useful variable and the delivery chain moves as predicted.

Chapter takeaway

Read the evidence in this order:

- 1 reference resource and timing;
- 2 received level and quality;
- 3 UE report under CSI assumptions;
- 4 scheduler allocation and transmission parameters;
- 5 error/retransmission outcome;
- 6 delivered radio and application data.

Stopping rule

If only one KPI supports the diagnosis, the diagnosis is not closed.

References

¹ ETSI, [TS 138 215 V18.4.0 - NR; Physical layer measurements](#), especially clauses 5.1.1, 5.1.3 and 5.1.5.

² ETSI, [TS 138 214 V18.9.0 - NR; Physical layer procedures for data](#), especially clauses 5.1.1, 5.1.3, 5.2.1.4 and 5.2.2.1.

Mobility in NSA and SA

FIELD QUESTION

At which phase did mobility fail, and what evidence proves it?

5G NR FROM SITE TO TRACE - FIELD PLATE V-009

NSA MOBILITY: KEEP THE LTE ANCHOR VISIBLE

In EN-DC, prove the Master Node path first, then isolate Secondary Node preparation, PSCell access and NR service.




FIELD RULE

LTE RRC complete does not prove successful NR random access or data.


EN-DC ARCHITECTURE - TWO CELL GROUPS, ONE MASTER ANCHOR

Generic Release 18 topology - no operator identities or proprietary implementation detail



SECONDARY NODE ADDITION - READ SIX EVIDENCE GATES

Signal names are generic teaching labels based on TS 37.340; implementation logs may use different labels.



COMMON REPORTING ERROR

RRC complete != proven PSCell access, NR scheduling or data.

CLASSIFY THE TRANSITION BEFORE YOU NAME THE FAILURE

The same KPI dip can sit behind different procedures. Anchor continuity is the first split.



SN ADDITION

LTE anchor remains; NR resources appear.



prove PSCell access + data

Preparation accepted does not equal active NR service.



SN MODIFICATION

Existing SCG resources are reconfigured.



same SN? decode what changed

Do not report an anchor handover from a config delta.



SN / PSCell CHANGE

NR serving relation changes; LTE may stay.



old PSCell -> new PSCell

Map source, target and Secondary Node transaction.



SN RELEASE

SCG disappears; LTE service can continue.



reason + LTE continuity + traffic impact

Loss of NR is not automatically a dropped session.



MASTER NODE HO

LTE anchor changes; SN handling follows context.



prove source MN -> target MN transition

This is the anchor handover category.

Normative basis: ETSI TS 137 340 V18.7.0. Original vector construction. Generic flow - not a vendor trace.

TAKWA SEBAI - TAKWASEBAI.COM - V-009

Figure V-009. Original EN-DC evidence flow. It separates LTE Master Node continuity, Secondary Node preparation, the nested NR configuration, PSCell access and observed NR data. It is a generic Release 18 teaching flow, not an operator trace.

The question is more useful than “Why did the handover fail?”

Before searching for a cause, establish which procedure was attempted and how far it progressed. In NSA, an NR Secondary Cell Group can be added, modified, changed, deactivated or released while the LTE Master Node remains in place. In SA, an intra-NR handover can complete radio access at the target before the core-network path switch and source-context release are finished.

Those are different state transitions with different messages, timers and failure boundaries.

The field method in this chapter is:

- 1 name the architecture and procedure;
- 2 map the source, target, Master Node and Secondary Node identities;
- 3 separate preparation, UE execution, target access and completion;
- 4 preserve the evidence at every boundary;
- 5 stop the diagnosis where the evidence stops.

No single missing message proves a radio root cause. It only identifies the next boundary to test.

1. Do not call every transition a handover

The word “handover” is often used for any visible serving change. That shortcut hides important differences.

At least five mechanisms can look like mobility in a trace:

Visible change	What can actually be happening
serving beam or beam index changes	lower-layer beam management, possibly without an RRC handover
NR appears while LTE remains	EN-DC Secondary Node or SCG addition
NR PSCell changes while the LTE anchor remains	PSCell change or Secondary Node modification/change
NR disappears while LTE remains connected	SCG deactivation or Secondary Node release
source NR PCell becomes a target NR PCell in SA	intra-NR handover or another configured mobility procedure

The analyst must therefore record:

- network mode: LTE, EN-DC, NGEN-DC, NE-DC, NR-DC or SA;
- core: EPC or 5GC;
- control-plane anchor;
- Master Cell Group PCell;
- Secondary Cell Group PSCell, if configured;
- serving cell before and after;
- procedure name supported by the messages;
- interface carrying each network message.

A throughput interruption during NR release in EN-DC is not automatically an LTE handover failure. A target beam change inside one serving cell is not automatically a cell handover. A source-to-target RRC command in SA does not prove the target was successfully accessed.

2. EN-DC: LTE is the Master Node

TS 37.340 defines EN-DC as E-UTRA-NR Dual Connectivity with the EPC. The UE is connected to an eNB acting as Master Node and an en-gNB acting as Secondary Node. The eNB connects to the EPC over S1 and to the en-gNB over X2.¹

The functional split matters:

- the LTE eNB provides the Master Node role;
- the LTE PCell belongs to the Master Cell Group;
- the NR PSCell belongs to the Secondary Cell Group;
- the en-gNB provides additional NR resources;
- the UE has one RRC state based on the Master Node RRC;
- the initial Secondary Node RRC configuration reaches the UE through the Master Cell Group signalling path.

The Master Node is not merely the cell with the strongest reference-signal level. It is the node providing the control-plane connection to the core network. The Secondary Node provides additional radio resources and has no control-plane connection to the core in that role.¹

This leads to a practical rule:

If LTE remains the same Master Node while NR is added, changed or released, do not label the event as an LTE anchor handover.

3. EN-DC Secondary Node addition: read the gates in order

TS 37.340 clause 10.2.1 describes Secondary Node Addition as a procedure initiated by the Master Node to establish UE context at the Secondary Node and provide Secondary Node resources. Where SCG radio resources are required, the procedure adds at least the first SCG cell.¹

A field-oriented reconstruction contains six gates.

Gate 1 - Master Node decision

The Master Node decides to request Secondary Node resources. It can provide UE capabilities, requested SCG information and recent measurement results that help the Secondary Node choose and configure SCG cells.

Evidence can include:

- active LTE RRC measurement configuration;
- UE MeasurementReport on the expected LTE signalling path;
- Master Node decision or admission record;
- correct UE identity and E-RAB context;
- X2 connectivity and target Secondary Node identity.

A measurement report is an input. It is not proof that the Master Node requested or admitted NR resources.

Gate 2 - Secondary Node preparation

The Master Node sends an SgNB Addition Request. If the Secondary Node admits the request, it chooses the PSCell and any other SCG cells, prepares radio and transport resources, and returns an SgNB Addition Request Acknowledge containing the NR RRC configuration for the UE.¹

At this boundary, ask:

- was the request sent to the intended en-gNB?
- was it accepted or rejected?
- which PSCell was selected?
- which NR frequency, PCI and configuration were prepared?
- did the response contain the expected NR configuration?
- were bearer and user-plane resources admitted?

No UE-side NR access has happened yet.

Gate 3 - Configuration delivery through LTE

The Master Node sends an LTE RRCConnectionReconfiguration to the UE containing the NR RRC configuration. TS 37.340 states that the Master Node transports the Secondary Node configuration without modifying the UE configuration provided by the Secondary Node.¹

Preserve both layers:

- outer LTE RRC transaction and transaction identifier;
- embedded NR configuration;
- target PSCell identity and NR frequency;
- SCG state;
- random-access resources, if present;
- bearer mapping;
- timestamp at the UE and Master Node.

If a decoder exposes only the outer LTE message, the analyst cannot safely claim that the expected NR configuration was present.

Gate 4 - UE configuration response

The UE applies the configuration and replies to the Master Node with RRCConnectionReconfigurationComplete. The Master Node then informs the Secondary Node with SgNB Reconfiguration Complete.¹

This proves that the RRC reconfiguration procedure reached its response boundary. It does not, by itself, prove that the NR PSCell random access or NR user plane is working.

Gate 5 - PSCell synchronisation and random access

When SCG radio resources are configured and the SCG is active, the UE synchronises to the PSCell and performs random access.

A subtle but important point in TS 37.340 is that the order between the UE's RRCConnectionReconfigurationComplete and the random-access procedure toward the SCG is not defined. Successful random access to the SCG is not required for successful completion of the LTE RRC Connection Reconfiguration procedure.¹

Therefore:

RRCConnectionReconfigurationComplete observed is not equivalent to NR radio access complete.

Look for:

- target NR synchronisation;

- PRACH attempt and response evidence;
- target timing advance where exposed;
- PSCell MAC activation;
- NR PDCCH/PDSCH or PUSCH activity;
- SCG failure information if access later fails.

Gate 6 - NR service outcome

Only after the configuration and radio-access evidence should the report evaluate:

- NR bearer state;
- NR scheduling;
- delivered NR bytes;
- split or SCG bearer behaviour;
- user-plane path update where applicable;
- application throughput and interruption.

An NR icon on a device is not a substitute for these observations.

4. Addition, modification, change, release and anchor handover are different

EN-DC exposes several procedures in TS 37.340:

- Secondary Node Addition;
- Secondary Node Modification;
- Secondary Node Release;
- Secondary Node Change;
- PSCell change;
- inter-Master Node handover with or without Secondary Node change.

They should not be merged into one KPI called "5G handover."

Use this classification:

Evidence pattern	Defensible procedure label
SgNB addition preparation, embedded NR configuration, new SCG	Secondary Node or SCG addition
same Master Node, changed SCG configuration	Secondary Node modification
same Master Node, target PSCell replaces source PSCell	PSCell or Secondary Node change, according to messages
same Master Node, SCG removed or deactivated	Secondary Node release or SCG deactivation
source Master Node changes to target Master Node	Master Node handover
both Master Node and Secondary Node relationships change	inter-Master Node handover with Secondary Node handling

Before assigning a failure, identify:

- which node initiated the procedure;
- whether the Master Node changed;
- whether the Secondary Node changed;
- whether only the PSCell changed;
- whether the SCG was active or deactivated;
- whether the bearer termination point changed;
- whether random access was required;
- whether the problem was configuration, target access or delivered data.

5. SA handover: three phases plus a post-check

This chapter uses the basic Xn-based intra-NR handover where neither AMF nor UPF changes. It is the clearest flow for phase isolation.

TS 38.300 states that preparation and execution are performed without 5GC involvement for this intra-NR RAN handover: preparation messages are exchanged directly between gNBs. During completion, the target gNB triggers release of resources at the source gNB.²

The engineering phases are:

- 1 preparation;
- 2 execution and target access;
- 3 completion and path switch;
- 4 post-handover service verification.

SA HANDOVER: STOP AT THE FIRST MISSING EVIDENCE

Preparation, target access and core path switch are separate proof boundaries - one success does not validate the next.




STOPPING RULE

Last positive evidence + first missing evidence = defensible failure phase.

Xn-BASED INTRA-NR HANDOVER - SOURCE, TARGET AND CORE EVIDENCE

Generic Release 18 flow with unchanged AMF and UPF - use N2/NGAP messages when that is the intended mobility path.



THREE PROCEDURE PHASES + ONE SERVICE POST-CHECK

The first missing mandatory evidence defines the phase boundary. Physical root cause still requires RF, RACH, transport or node evidence.

1 PREPARATION

Can the target accept and prepare the UE context?

- measurement report + source decision
- Xn HANDOVER REQUEST
- target ACK + transparent target config

ACK proves preparation
not target radio access

2 EXECUTION + TARGET ACCESS

Did the UE receive the command, synchronize and complete access?

- RRCReconfiguration + reconfigurationWithSync
- target sync + random access evidence
- target RRCReconfigurationComplete

T304 expiry localizes
this phase - it does not
name the physical cause

3 COMPLETION

Did the core path move and the source context release?

- Path Switch Request
- Path Switch Request Ack
- UE Context Release to source

target access can succeed while completion fails later

4 POST-CHECK


Did radio and service remain usable on target?

- target cell + RNTI continuity
- DL/UL scheduling + BLER/HARQ
- bearer and application continuity

protocol success is not service acceptance

SYNTHETIC C-008 - PREPARATION SUCCEEDS, TARGET ACCESS DOES NOT

Relative timestamps are an exercise, not operator data. T304 is configured to 1000 ms in this synthetic comparison.



Normative basis: ETSI TS 138 300 V18.9.0, TS 138 331 V18.9.0, TS 138 401 V18.7.0 and TS 138 413 V18.9.0. C-008 is synthetic.

TAKWA SEBAI - TAKWASEBAL.COM - V-010

CONFIRM THE MOBILITY PATH

Missing Xn messages are not a fault when N2 is intended.




Figure V-010. Original Xn-based SA handover map. The success baseline and C-008 failure timestamps are synthetic. The plate localises evidence boundaries; it does not identify a physical root cause.

Phase 1 - Preparation

Preparation begins before the UE receives a handover command.

The basic evidence chain is:

- 1 the source gNB configures measurements;
- 2 the UE sends MeasurementReport according to that configuration;
- 3 the source gNB makes a handover decision using the report and other RRM information;
- 4 the source gNB sends HANDOVER REQUEST to the target gNB;
- 5 the target performs admission and preparation;
- 6 the target returns HANDOVER REQUEST ACKNOWLEDGE with a transparent RRC container for the UE.²

Preparation succeeds only when the target has accepted and returned usable target configuration.

Possible preparation failure evidence includes:

- no decision despite a valid report;
- wrong target identity;
- HANDOVER REQUEST not sent;
- transport or Xn association failure;
- target admission rejection;
- slice, PDU session or QoS resource rejection;
- HANDOVER REQUEST ACKNOWLEDGE absent or unusable;
- explicit cause in XnAP or network counters.

The UE trace alone normally cannot prove all these network-side states.

Phase 2 - Execution and target access

After successful preparation, the source gNB sends RRCReconfiguration containing the information needed to access the target. TS 38.300 lists the target cell identity, new C-RNTI and security information, with possible dedicated or common RACH resources and target system information.²

In TS 38.331, the target change is represented by **reconfigurationWithSync**. The UE starts T304 for the target SpCell, starts synchronising to the target and applies the target configuration.³

The abbreviated synthetic decode below shows the fields an engineer should expect a tool to expose. It is not a captured operator message and not a verbatim ASN.1 schema.

```
RRCReconfiguration
  rrc-TransactionIdentifier: 2
  secondaryCellGroup
    CellGroupConfig
      spCellConfig
        reconfigurationWithSync
          spCellConfigCommon
            physCellId: 418
            absoluteFrequencySSB: 641280
            newUE-Identity: 0x3A71
            t304: ms1000
            rach-ConfigDedicated: present
```

This proves that the target configuration contains a synchronised change, target identity, timer and dedicated RACH context. It does not prove that the UE decoded the command or completed target access.

The execution evidence chain is:

- 1 source transmitted the RRCReconfiguration;
- 2 UE decoded the intended transaction;
- 3 **reconfigurationWithSync** and target identity are present;
- 4 UE starts target synchronisation;
- 5 target random access is attempted and succeeds, unless a supported procedure removes that need;
- 6 UE sends RRCReconfigurationComplete using the target configuration;
- 7 target gNB observes the UE and accepts the completion.

These distinctions matter:

- source “command sent” is not UE “command decoded”;
- UE “command decoded” is not target “synchronisation achieved”;
- PRACH transmitted is not random access completed;
- random access completed is not path switch completed;
- target RRC completion is not proof of healthy application service.

If T304 expires for the Master Cell Group during reconfiguration with sync, TS 38.331 defines failure handling that can include storing handover-failure information and initiating connection re-establishment.³

T304 expiry localises a reconfiguration-with-sync failure. It does not name the physical cause.

RACH is the target-access proof, not one PRACH line

Random access is a short procedure with several evidence gates. In the common four-step contention-based view, the UE transmits a preamble, receives a Random Access Response with timing and uplink resources, sends the scheduled message, then resolves contention. A contention-free mobility procedure can use dedicated resources and a different evidence path. The trace must tell you which procedure applied.

For mobility RCA, preserve at least:

- 1 selected target cell, beam and RACH occasion;
- 2 preamble transmission and attempt count;
- 3 Random Access Response detection and the grant it carried;
- 4 scheduled uplink transmission and target reception;
- 5 contention resolution or the equivalent completion evidence;
- 6 the first target RRC and user-plane outcome.

PRACH transmitted stops at step 2. It does not prove timing alignment, uplink decoding, contention resolution or usable target service.

Phase 3 - Completion and path switch

After successful target access and RRC completion, the target gNB requests the core path switch. In the basic flow:

- 1 target gNB sends PATH SWITCH REQUEST to the AMF;
- 2 the 5GC switches the downlink path;
- 3 the AMF sends PATH SWITCH REQUEST ACKNOWLEDGE;
- 4 target gNB sends UE CONTEXT RELEASE to the source;
- 5 source radio and control-plane resources can be released.²

Completion failures therefore need evidence beyond the UE’s target RRC message:

- PATH SWITCH REQUEST sent or absent;
- AMF response and cause;
- accepted and failed PDU sessions;
- new NG-U path state;
- end marker and forwarding state where exposed;
- source context release;
- duplicate or stale context;
- post-handover data continuity.

An RRCReconfigurationComplete at the target with no path switch evidence is not a fully closed handover record.

Phase 4 - Post-handover verification

The protocol flow can complete while the service remains degraded.

Check:

- target serving cell and beam stability;
- RLF or re-establishment immediately after completion;
- target PDCCH/PDSCH and PUSCH;
- BLER, HARQ, rank and allocation;
- RLC/PDCP continuity;
- delivered bytes and packet loss;
- voice or application interruption;
- source resources actually released;
- repeated ping-pong or immediate fallback.

This is the difference between protocol success and service success.

Release alignment is not a Release 18 feature catalogue

This guide cites Release 18 versions of the specifications for the baseline procedures it uses. It does not claim to teach every Release 18 mobility feature. Conditional handover, DAPS and L1/L2-triggered mobility require their own capability, configuration, trigger and execution evidence; they must not be inferred from the basic Xn flow above.

6. Xn-based and N2-based flows do not expose the same messages

Do not search for XnAP messages before confirming the handover path.

For the Xn-based flow described above:

- source and target gNBs exchange preparation directly;
- 5GC is involved for path switch during completion;
- source and target XnAP UE identifiers must be mapped.

For an NG-based or N2 handover, TS 38.413 defines NGAP mobility procedures:

- source NG-RAN node sends HANDOVER REQUIRED toward the AMF;
- AMF initiates target resource allocation with HANDOVER REQUEST;
- target returns HANDOVER REQUEST ACKNOWLEDGE or a failure;
- AMF returns the handover command toward the source;

- target sends HANDOVER NOTIFY after the UE is identified in the target cell and the NG-based handover completes.⁵

Therefore the absence of an Xn HANDOVER REQUEST is not a preparation failure if N2 was the intended path.

Preserve the identifier mapping:

- source and target global node IDs;
- source and target cell identities;
- PCI and ARFCN;
- old and new C-RNTI;
- source and target XnAP UE IDs;
- AMF UE NGAP ID;
- RAN UE NGAP ID;
- RRC transaction identifier;
- UE trace identity and clock.

TS 38.401 distinguishes the UE-associated logical NG and Xn connections and their identifiers.⁴ Without that mapping, messages from two UEs or two attempts can be merged into one false flow.

7. Synthetic field case C-008 - preparation succeeds, target access does not

Evidence status: synthetic SA Xn comparison. No operator trace, customer counter, real cell identity or author incident is represented.

Exercise scope

- SA connected-mode UE;
- intra-frequency NR handover;
- Xn-based flow;
- same AMF and UPF;
- no DAPS and no conditional handover;
- source and target gNB clocks normalised to one timeline;
- UE, XnAP and NGAP records mapped to one synthetic attempt;
- configured T304 in the exercise: 1000 ms;
- all timestamps below are teaching values, not performance targets.

Known-good baseline

Time	Evidence
0 ms	UE MeasurementReport received at source gNB
8 ms	source handover decision recorded
16 ms	Xn HANDOVER REQUEST sent
31 ms	HANDOVER REQUEST ACKNOWLEDGE received
36 ms	RRCReconfiguration with reconfigurationWithSync decoded by UE
82 ms	target random access completed
88 ms	target received RRCReconfigurationComplete
96 ms	target sent PATH SWITCH REQUEST
113 ms	PATH SWITCH REQUEST ACKNOWLEDGE received
121 ms	source received UE CONTEXT RELEASE

The baseline proves that the full preparation, execution and completion chain can succeed for the synthetic configuration.

Failed attempt

Time	Evidence
0 ms	UE MeasurementReport received at source gNB
9 ms	source handover decision recorded
18 ms	Xn HANDOVER REQUEST sent
35 ms	HANDOVER REQUEST ACKNOWLEDGE received
40 ms	intended RRCReconfiguration with reconfigurationWithSync decoded by UE
50 ms	target access attempt begins
214 ms	further target-access activity visible; no completed target RRC transaction
1040 ms	synthetic T304 expiry recorded
1083 ms	RRCReestablishmentRequest observed on the recoverable source cell

Absent in the failed attempt:

- target RRCReconfigurationComplete;
- PATH SWITCH REQUEST;
- PATH SWITCH REQUEST ACKNOWLEDGE;
- UE CONTEXT RELEASE at the source;
- target user-plane delivery.

What the evidence proves

Preparation succeeded:

- the source made a decision;
- the target received the request;
- the target admitted the attempt;

- the source received target configuration.

Command delivery succeeded:

- the UE decoded the intended RRCReconfiguration;
- the target identity and **reconfigurationWithSync** were present.

Execution did not reach target completion:

- target access activity began;
- no target RRCReconfigurationComplete was observed;
- T304 expired;
- the UE initiated recovery.

The defensible classification is:

Synthetic C-008 failed during execution or target access after successful preparation and command delivery, before target RRC completion and core path switch.

What the evidence does not prove

The timeline does not identify why target access failed.

Possible hypotheses still requiring discriminating evidence include:

- inadequate target downlink synchronisation;
- target uplink coverage or power limitation;
- incorrect or unavailable RACH resource;
- contention or response failure;
- target beam mismatch;
- target configuration inconsistency;
- security or lower-layer application problem;
- UE implementation behaviour;
- missing network-side target reception evidence.

No one of these is the synthetic root cause. The case intentionally stops at phase localisation.

8. Failure matrix: state what is known and what is missing

Last positive evidence	Missing next evidence	Phase classification	Next check
MeasurementReport	no source decision	pre-preparation decision gap	RRM policy, target eligibility, event validity
HANDOVER REQUEST	no ACK or explicit reject	preparation	Xn/N2 path, target admission and cause
HANDOVER REQUEST ACKNOWLEDGE	no command at UE	source command delivery	source scheduling, RRC transaction, UE decode
RRC command decoded	no target access completion	execution/target access	T304, sync, RACH, target UL/DL evidence
target RRC complete	no path switch	completion	NGAP, PDU sessions, AMF/UPF and transport
path switch ACK	no source context release	completion cleanup	target release trigger, Xn state, source context
protocol flow complete	service degraded	post-handover outcome	RF, scheduler, PDCP/RLC, transport, application

Use the narrowest classification supported by the last positive and first missing evidence.

Do not write “handover failed at random access” merely because the UE log lacks RRCReconfigurationComplete. Random-access failure requires lower-layer or target-side proof.

9. Compare with a known-good flow

A mobility flow is easier to diagnose when it is compared with a successful attempt under controlled conditions.

Keep constant where possible:

- UE model, modem and software;
- source and target cells;
- frequency and bandwidth;
- measurement configuration;
- service and QoS;
- traffic direction;
- route direction and speed;
- time synchronisation;
- Xn or N2 procedure type;
- network software and configuration revision.

Compare by phase:

Phase	Comparison fields
measurement	event, measId, object, reportConfig, quantities, neighbour identity
decision	target selected, decision timestamp, reason/policy where exposed
preparation	request, acknowledgement/failure, cause, admitted resources, target container

Phase	Comparison fields
command	RRC transaction, target identity, new C-RNTI, T304, RACH resources, security
access	DL sync, SSB/beam, PRACH, RAR, contention result, timing advance, target MAC
RRC completion	target receives RRCReconfigurationComplete, transaction matches
core completion	path switch, PDU session result, source context release
service	interruption, loss, BLER/HARQ, PDCP/RLC continuity, delivered data

Compare elapsed times only after the same start and stop events are defined. A vendor counter named “handover execution time” may not use the same boundary as the UE trace.

10. Evidence checklist

Before publishing a mobility diagnosis:

- ☐ The architecture is named: EN-DC, another MR-DC mode or SA.
- ☐ EPC and 5GC procedures are not mixed.
- ☐ The Master Node, Secondary Node, source and target roles are explicit.
- ☐ PCell, PSCell and serving-cell identities are mapped.
- ☐ The procedure is named from messages, not from an icon or KPI label.
- ☐ Xn-based and N2-based handovers are distinguished.
- ☐ UE, X2/Xn, NGAP and counter clocks are normalised.
- ☐ UE-associated identifiers are mapped across interfaces.
- ☐ Measurement object, reportConfig, measId, event and quantities are preserved.
- ☐ A measurement report is not presented as a handover decision.
- ☐ Preparation request, response and cause are recorded.
- ☐ The target RRC container and UE command are associated with the same attempt.
- ☐ **reconfigurationWithSync**, target identity, C-RNTI, T304 and RACH configuration are decoded where applicable.
- ☐ Command sent, command decoded, target synchronisation and random-access completion are separate states.
- ☐ Target RRCReconfigurationComplete is confirmed on the target side where possible.
- ☐ Path switch and PDU session results are checked.
- ☐ Source-context release is checked.
- ☐ RLF, re-establishment and immediate fallback after mobility are checked.
- ☐ Protocol completion and service outcome are evaluated separately.
- ☐ The conclusion identifies the last positive and first missing evidence.
- ☐ Synthetic values remain visibly declared.

Takwa Sebai working criterion

I report a mobility failure by phase: last positive evidence, first missing or divergent evidence, and the identifiers that tie both to one attempt. A timer expiry can close the phase. It cannot name RF, transport or target configuration as the physical cause without a separate discriminating test.

Chapter takeaway

Mobility diagnosis is a boundary problem.

For EN-DC:

- 1 confirm the LTE Master Node;
- 2 identify Secondary Node addition, change, deactivation or release;
- 3 separate RRC response from PSCell access;
- 4 verify NR scheduling and delivered data.

For SA:

- 1 prove target preparation;
- 2 prove command delivery;
- 3 prove target synchronisation and access;
- 4 prove target RRC completion;
- 5 prove path switch and source release;
- 6 verify post-handover service.

The last positive evidence locates progress. The first missing evidence defines the next test. Neither one identifies a root cause without supporting proof.

References

- ¹ ETSI, [TS 137 340 V18.7.0 - Multi-connectivity: Overall Description; Stage-2](#), especially clauses 3.1, 4.1.2, 4.2.1 and 10.2 to 10.7.
- ² ETSI, [TS 138 300 V18.9.0 - NR and NG-RAN Overall description; Stage-2](#), especially clause 9.2.3.2.1.
- ³ ETSI, [TS 138 331 V18.9.0 - NR RRC protocol specification](#), especially clauses 5.3.5.3, 5.3.5.5.2 and 5.3.5.8.3.
- ⁴ ETSI, [TS 138 401 V18.7.0 - NG-RAN Architecture description](#), especially clause 6.4.
- ⁵ ETSI, [TS 138 413 V18.9.0 - NG Application Protocol](#), especially clauses 8.4.1 to 8.4.4.

Low Throughput: Stop Guessing, Split the Problem

FIELD QUESTION

Throughput is low. Which test separates RF, scheduling, transport and policy?

LOW THROUGHPUT: SPLIT THE PROBLEM

A controlled field test should force the ceiling to follow one evidence branch.

FIELD METHOD
TEST - PROVE - STOP

GATE 0 - VALIDATE BEFORE DIAGNOSING

Freeze the traffic contract and prove demand

- DL or UL
- server + path
- protocol + streams
- duration + clock
- UE thermal/interface



OBSERVED: LOW APP GOODPUT

Name its point, numerator, denominator and exact window.

WHICH CONTROLLED TEST MOVES THE CEILING?



01 RF EFFICIENCY

PROVE TOGETHER

Quality/interference + CQI/RI
+ scheduled MCS/layers + initial
BLER/HARQ + allocation.

CONTROLLED TEST

Same UE + server + load.
Move to a known-good radio
position or time-aligned window.

DEFENSIBLE REPORT

"Efficiency deteriorates with the
radio window." Do not invent the
physical source of degradation.

C-009 SYNTHETIC FINGERPRINT

same PRB; SINR/rank/BLER and rate diverge



02 SCHEDULER / CAPACITY

PROVE TOGETHER

Traffic demand + grant frequency
+ PRB/symbols/layers/TBS + load
+ healthy radio outcome.

CONTROLLED TEST

Same UE + position + server.
Compare busy and quiet windows
with RF and demand established.

DEFENSIBLE REPORT

"The ceiling follows allocation."
This does not yet explain why the
scheduler allocated less.

C-010 SYNTHETIC FINGERPRINT

healthy RF; 34 vs 210 PRB; rate follows



03 TRANSPORT / TEST PATH

PROVE TOGETHER

Local/remote endpoint + N3/S1-U
+ UPF/PGW-U + loss/RTT/TCP
+ ports/queues/server state.

CONTROLLED TEST

Same UE + cell + test settings.
Change one endpoint or measured
boundary; keep the radio stable.

DEFENSIBLE REPORT

"The ceiling follows this path."
Do not name a hop without port,
queue or packet evidence.

C-011 SYNTHETIC FINGERPRINT

local 250-255; remote repeatable near 94



04 POLICY / CONFIG / UE

PROVE TOGETHER

SIM/session matrix + policy record
+ capability vs configured state
+ UE thermal/local interface.

CONTROLLED TEST

Cross two SIMs and two validated
UEs, or repeat after a documented
session/policy change.

DEFENSIBLE REPORT

"The ceiling follows SIM/session."
A flat graph alone never proves
policy enforcement.

C-012 SYNTHETIC FINGERPRINT

2 UE x 2 SIM; independent policy record

FIELD STOP RULE: report the last positive boundary and the first divergence. Do not turn branch localisation into an unsupported root cause.

TAKWA SEBAI · TAKWASEBAI.COM · V-011

Figure V-011. Original four-branch diagnostic tree. Start with the test contract and traffic demand. Each branch contains a discriminating test and a limit on what can be concluded. C-009 to C-012 are synthetic training fingerprints.

“Low throughput” is a symptom observed at one point in an end-to-end path. It is not a root cause and it is not one counter.

The fastest useful investigation is not the one that opens the most dashboards. It is the one that changes one controlled variable, keeps the clocks aligned and proves which branch the ceiling follows.

The field method in this chapter is:

- 1 validate the traffic test before diagnosing the network;
- 2 define the measurement point, interval and denominator of every throughput value;
- 3 split downlink from uplink;
- 4 prove traffic demand before judging allocation;
- 5 separate RF efficiency from scheduled resources;
- 6 move the test boundary to challenge transport and server limits;
- 7 cross UE and subscription only when configuration or policy is suspected;
- 8 stop the conclusion at the last point supported by evidence.

Four diagnostic branches are used throughout the chapter:

- RF quality and radio efficiency;
- scheduler and capacity;
- transport and test path;
- policy, configuration and UE capability.

The branches are an original troubleshooting framework. They are not vendor thresholds, and a real incident may involve more than one branch.

1. Validate the test before touching the network

A throughput value is meaningful only when the test contract is known.

Record the following before opening a radio trace:

Test field	What must be recorded	Why it matters
direction	DL, UL or simultaneous	DL and UL use different scheduling, buffer and power evidence
server	name or anonymised ID, location and ownership	the server can be the ceiling
path	local edge, operator network, public internet, VPN	changing the path changes more than the radio
protocol	TCP, UDP or application-specific	congestion control, loss and acknowledgements behave differently
streams	number of parallel sessions	one stream may not fill the available path
duration	warm-up, measurement and stop interval	short tests can measure ramp-up instead of steady service
payload	application bytes, IP bytes or another layer	headers and retransmissions change the number
UE state	model, software, temperature, power and interface	the terminal or local interface can throttle
timestamp	common timezone and clock offset	unsynchronised windows create false correlations
demand	proof that the source had data to send	no demand can look like no scheduling

Do not compare two tests when one used a local multi-stream server for 60 seconds and the other used a remote single-stream service for 10 seconds. The radio may be identical while the test is not.

A practical pre-check

Run three short controls before escalating:

- 1 repeat the same test without changing location;
- 2 repeat against a known-good server or path;
- 3 repeat after checking UE temperature, power state, VPN and local tethering or Wi-Fi limits.

If the result moves when only the server or path changes, keep the radio branch open only if separate radio evidence also changed. If the result moves after the terminal cools, investigate the UE before optimising the cell.

Do not diagnose an idle test

The network can allocate few resources because the application has little data waiting. That is not a scheduler failure.

For DL, establish that the network-side buffer or traffic source offered enough data during the interval. For UL, preserve SR, BSR and grant evidence where available. TS 38.321 separates scheduling requests, buffer status reporting, uplink grants, logical-channel prioritisation and power headroom procedures.²

The basic question is:

Was the UE limited by the network, or was the network simply responding to the amount of traffic offered?

My discriminating-test rule

The practical rule I use before touching a network is:

A controlled test is discriminating only when the expected state transition is observable and the competing explanations predict different results.

Repeating the same UE, cell, endpoint and path and obtaining the same ceiling proves repeatability, not the branch. Before the test, write:

- 1 the state expected to change;
- 2 the counter or trace that will prove the change;
- 3 the result expected if each competing hypothesis is true;
- 4 the safety stop and rollback;
- 5 the post-check that confirms service recovery.

If none of the branches predicts a different observable result, redesign the test before touching the network.

2. Theoretical, expected and observed throughput are different numbers

Three numbers are often placed in one table as if they described the same thing.

Theoretical peak

A theoretical peak uses an assumed radio configuration and idealised resource use. It can include bandwidth, usable resource blocks, scheduled symbols, modulation order, target code rate, number of layers and an assumed overhead.

TS 38.214 defines how an MCS index selects a modulation order and target code rate in the applicable table, then how a transport block size is determined from scheduled resources and other inputs.¹

Therefore:

MCS alone is not throughput.

Two samples with the same MCS can have very different throughput because they use different numbers of PRB, symbols, layers, slots or successful transport blocks.

Expected field range

An expected range is an engineering baseline for a stated configuration and test method. It should be built from:

- a known-good test on the same network;
- the same UE category and capabilities;
- the same carrier aggregation or dual-connectivity state;
- a comparable radio position and load;
- the same server, protocol, streams and duration.

It is more useful than a marketing peak because it preserves real overhead, implementation and traffic-path constraints.

Observed result

An observed result is valid only for its measurement point and time base.

TS 28.552 defines RAN UE throughput measurements using data volume and an active-time denominator. For downlink, the measurement can use RLC SDU data successfully transmitted and acknowledged while excluding parts of the interval when the buffer is empty. It can also be split by attributes such as QoS, slice, PLMN or BWP.⁵

An application speed test normally reports delivered application data over its own wall-clock interval. It may include:

- TCP ramp-up;
- application warm-up;
- periods with an empty buffer;
- packet loss and end-to-end retransmission effects;
- server processing;
- public-network path changes;
- headers and encryption overhead;
- a different start and stop time.

The two values do not have to be identical. A difference is not automatically a counter fault.

ONE FLOW, MULTIPLE THROUGHPUT NUMBERS

A rate only becomes evidence when its layer, numerator, denominator and time window are known.

END-TO-END FIELD MAP
DL + UL EVIDENCE

APPLICATION TEST


Application goodput = useful application bytes / test wall-clock time

Includes its own warm-up, idle periods, endpoint behaviour and end-to-end loss.

RAN UE THROUGHPUT

RAN throughput = defined data volume / defined active transfer time

The measurement can remove idle gaps and use a lower-layer delivery boundary. It need not equal app goodput.



Generic SA/NSA user-plane teaching chain. No vendor topology, operator counter name, customer address or production value is represented.

MEASURED - DEFINED - SYNCHRONISED

Figure V-012. Original end-to-end measurement map. It shows why application goodput, transport throughput, N3/S1-U traffic, PDCP delivery, RLC volume and scheduled radio data can legitimately produce different numbers. The chain is generic and contains no operator topology.

3. One flow crosses several accounting boundaries

Treat the user flow as a chain of measurement points.

Boundary	Useful evidence	What it cannot prove alone
application	useful bytes, duration, server, streams	where below the application the limit occurred
TCP/UDP/IP	rate, RTT, loss, retransmissions, congestion window	radio scheduling state
Data Network/N6	server and external path traffic	RAN air-interface efficiency
UPF or PGW-U	per-session or tunnel volume, drops, shaping	which radio mechanism limited delivery
N3 or S1-U	interface throughput, loss, delay, drops	UE capability or cell allocation by itself
PDCP	delivered volume, discard, duplication context	exact MAC grant or RF condition
RLC	SDU/PDU volume and AM recovery where applicable	lower-layer HARQ result by itself
MAC	buffers, grants, PRB, TBS, HARQ process	application goodput over another interval
PHY	CQI context, MCS, layers, BLER/CRC, power	server, policy or transport ceiling
UE	capability, aggregation, thermal and local interface	network-side admission or path state without corroboration

MAC, RLC and PDCP are not interchangeable

At MAC, distinguish:

- whether data was waiting;
- whether the scheduler selected the UE;
- the frequency and size of grants;
- PRB, symbols, layers, MCS and transport block size;
- initial transmission and HARQ retransmission outcome.

At RLC, first identify the mode. TS 38.322 defines TM, UM and AM entities. RLC AM can provide retransmission and recovery; that mechanism must not be confused with MAC HARQ, and it must not be claimed for a bearer using another RLC mode.³

At PDCP, useful context can include delivery, discard, status, recovery, data volume and duplication. These functions depend on the configured bearer and procedure.⁴

The investigation should ask where useful data stopped progressing, not which layer owns the largest counter.

Align clocks before aligning values

Use one incident timeline:

- traffic test start, steady interval and stop;
- UE trace clock;
- gNB or eNB clock;
- PM counter collection window;
- transport device clock;
- UPF or PGW-U clock;
- policy or session-event clock.

A 15-minute PM bin cannot safely explain a 20-second test unless the test dominates the bin or a more granular counter supports the claim.

4. Split downlink and uplink

Do not write “the throughput is low” without a direction.

Downlink evidence chain

For DL, preserve:

- 1 application or network-side demand;
- 2 data waiting for the UE;
- 3 scheduling opportunities and actual grants;
- 4 PRB, symbols, layers, MCS and TBS;
- 5 HARQ and BLER outcome;
- 6 RLC/PDCP progression;
- 7 N3/S1-U and application delivery.

A strong SS-RSRP with few DL grants can indicate low demand, capacity pressure, QoS priority, policy, a control-channel limitation or another scheduler constraint. Strong level alone does not choose between them.

Uplink evidence chain

For UL, add:

- Scheduling Request state;
- Buffer Status Report;
- uplink grant size and frequency;
- Power Headroom Report or other transmit-power evidence;
- configured maximum power and power-limited state where available;
- PUSCH MCS, layers and HARQ result;
- RLC/PDCP progression toward the network.

An uplink UE can have data buffered but receive small grants. It can receive grants but fail to use the intended modulation or bandwidth because of power limitation. It can deliver radio data while an upstream transport or server path limits the application.

The DL and UL branches may therefore end at different causes in the same location.

5. The four-branch diagnostic tree

Use the tree in this order.

Branch 1 - RF quality and radio efficiency

Ask whether the scheduled transmission is efficient under the observed radio conditions.

Check together:

- SS-RSRP or CSI-RS level with the correct prefix and resource;
- SS-RSRQ and SS-SINR or applicable CSI quantities;
- CQI and RI context;
- scheduled MCS and layers;

- initial BLER, CRC and HARQ retransmissions;
- PRB and scheduled symbols;
- UL power headroom where relevant;
- location, beam, movement and interference context.

A low MCS can be a response to radio conditions, but only after the MCS table, grant, layers and timing are known. A high PRB allocation with low useful throughput and poor radio efficiency is different from a low PRB allocation with healthy efficiency.

Discriminating tests include:

- move to a known-good radio position without changing server or UE;
- compare a good and bad interval with the same test contract;
- compare direction-specific RF and HARQ evidence;
- check whether the ceiling follows beam, position or interference change.

Defensible conclusion:

The degraded interval is consistent with an RF-efficiency limitation because quality, scheduled adaptation and transmission outcome deteriorate together while the test path is held constant.

Do not claim the physical source of interference or coverage loss without separate evidence.

Branch 2 - Scheduler and capacity

Ask whether sufficient traffic was offered but the UE received too few usable resources.

Check:

- DL buffer or UL BSR/demand;
- grant frequency;
- mean and distribution of PRB;
- scheduled symbols and slots;
- layers and transport block size;
- active UE count and cell/BWP load;
- QoS priority and admission state;
- PDCCH/CORESET constraints if supported by evidence;
- radio efficiency during low and high allocation intervals.

The useful comparison is not “busy hour versus midnight” by itself. Keep the UE, position, server and radio quality comparable. Then test whether the throughput follows allocated resources.

Discriminating tests include:

- repeat in a quieter window with the same test method;
- compare per-UE allocation while demand and RF remain established;
- use a controlled load test where authorised;
- compare the same UE with a known-good cell or sector without changing the path unnecessarily.

Defensible conclusion:

The observed ceiling follows reduced scheduling allocation in the compared interval.

That statement does not identify why the scheduler allocated less. Capacity pressure, QoS, control resources, fairness, admission, implementation and configuration remain separate hypotheses until proven.

Branch 3 - Transport and test path

Ask whether the radio can deliver more when the traffic endpoint or path is moved.

Check:

- local and remote server results;
- N3 or S1-U interface volume;
- UPF or PGW-U throughput;
- N6 or external path;
- port errors, queue drops and utilisation;
- RTT, jitter, packet loss and TCP retransmissions;
- MTU and fragmentation;
- load balancer, firewall, VPN and encryption limits;
- server CPU, NIC and per-session constraints.

Discriminating tests include:

- same UE and cell, local known-good server versus remote server;
- same remote server through another validated access path;
- simultaneous interface counters during the exact test window;
- UDP and TCP comparison where authorised and correctly controlled;
- packet capture at two boundaries with aligned clocks.

Defensible conclusion:

The ceiling follows the selected path beyond the radio in the controlled comparison.

Do not name a specific cable, router or negotiation rate unless interface evidence identifies it.

Branch 4 - Policy, configuration and UE

Ask whether the limit follows the subscription, session, bearer, configuration or terminal.

Policy evidence can include:

- subscription and session policy;
- UE-AMBR, Session-AMBR or other applicable rate limit;
- QoS Flow and DRB mapping;
- slice and DNN context;
- charging or policy-control decision;
- UPF enforcement record;
- APN/DNN or PDU Session differences.

TS 23.501 describes Session-AMBR for the aggregate Non-GBR QoS flows of a PDU Session and UE-AMBR for aggregate Non-GBR traffic of a UE, with enforcement at defined functions.⁶

UE/configuration evidence can include:

- supported band combinations;
- feature-set and maximum-layer capability;

- actual carrier aggregation or MR-DC configuration;
- configured and active BWP;
- modem software;
- thermal state;
- local USB, Ethernet or Wi-Fi ceiling.

TS 38.306 provides capability signalling for supported band combinations and associated feature sets. An advertised capability is not proof that the network configured or realised it during the test.⁷

Discriminating tests include:

- same SIM in two validated UEs;
- two SIMs in the same validated UE;
- same session before and after a documented policy change;
- capability versus configuration versus scheduled realisation;
- direct-on-device test versus tethered test.

Defensible conclusion:

The ceiling follows the subscription, session or terminal in the controlled matrix.

A flat graph alone does not prove shaping.

6. NSA and SA change the path you must prove

NSA or EN-DC

In EN-DC, identify:

- LTE Master Node and NR Secondary Node;
- PCell and PSCell;
- split, MCG or SCG bearer;
- which node schedules the observed data;
- whether NR was configured, accessed and actually carrying data;
- EPC user-plane termination and transport path.

A 5G icon or an NR measurement does not prove that the tested bytes used NR. Preserve scheduled bytes or bearer evidence for the radio contribution.

If throughput falls when NR disappears, separate:

- loss of NR resources;
- LTE anchor limitation;
- bearer reconfiguration;
- transport-path change;
- policy change associated with access state.

Carrier aggregation and EN-DC contribution

An aggregate result must be decomposed before it is used as evidence. For carrier aggregation, record the serving-cell set, active BWPs, scheduling state, PRBs, layers and delivered bytes per component carrier. For EN-DC, also record the LTE Master Cell Group, NR Secondary Cell Group and the bearer path.

Three different statements must not be merged:

- the UE and network support a band combination;
- the combination is configured and activated;
- each component actually carries scheduled user data in the measured window.

If the total falls when one component disappears, quantify that component's previous contribution. Do not assign the complete delta to it unless the other carriers, load, RF, policy and test path remained controlled.

SA

In SA, map:

- PDU Session;
- QoS Flow and 5QI;
- DRB;
- gNB user plane;
- N3 tunnel;
- UPF;
- N6 and Data Network;
- application server.

If the RAN reports healthy delivery but the application is low, move the measurement boundary toward N3, UPF, N6 and the server. If N3 traffic is already low, return toward PDCP, RLC, MAC, RF and demand.

Arithmetic contract for C-009 to C-012

Every Mbit/s value in the four exercises is checked against a declared single-carrier envelope: n78, 100 MHz, 30 kHz SCS, PDSCH MCS table 1, a repeating **DDDSU** model, eleven PDSCH data-symbol equivalents after overhead per eligible slot and the common wall-clock factor **f_{DL} = 0.70**. Every row also applies its stated initial-BLER factor, **1 – BLER**; C-011 and C-012 use the known-good radio state of 180 mean PRBs, MCS 24, two layers and 6% initial BLER. No case assumes carrier aggregation or EN-DC.

The resulting ceiling is deliberately an upper sanity bound. Application goodput must remain below it; passing the check does not predict a real scheduler or protocol result.

7. Synthetic case C-009 - the RF branch

Evidence status: synthetic DL training comparison. No customer data, operator threshold or author incident is represented.

Test contract held constant:

- n78, 100 MHz;
- same synthetic UE and software;
- same cell, server, multi-stream DL test and duration;
- comparable cell load;
- two radio positions or windows.
- PDSCH MCS table 1;
- repeating **DDDSU** TDD model;
- eleven PDSCH data-symbol equivalents after overhead per eligible slot;
- the common **f_{DL} = 0.70** wall-clock factor and each row's stated initial-BLER factor;
- one NR carrier, with no CA or EN-DC contribution.

Observation	Degraded window	Known-good window
SS-RSRP	−92 dBm	−79 dBm
SS-SINR	3 dB	21 dB
CQI	6	13
scheduled MCS	9	24
scheduled rank	1	2
initial DL BLER	24%	6%
mean scheduled PRB	180	180
application throughput	27 Mbit/s	255 Mbit/s

What the exercise proves:

- the test path and allocation magnitude are held comparable;
- quality, adaptation, rank and radio outcome degrade together;
- the ceiling follows the degraded radio window in this synthetic comparison.

Defensible classification:

C-009 is consistent with an RF-efficiency limitation.

What it does not prove:

- a universal SS-SINR or BLER threshold;
- the physical origin of the degraded quality;
- whether interference, geometry, beam, hardware or another RF condition is responsible.

8. Synthetic case C-010 - the scheduler/capacity branch

Evidence status: synthetic DL training comparison. The numbers are illustrative, not a vendor benchmark.

The same UE, position, server and test method are compared in two windows. SS-SINR, CQI, MCS, rank and initial BLER remain in a similar healthy range. Demand is established in both windows. The case uses PDSCH MCS table 1, scheduled MCS 24, two layers, the same **DDDSU** pattern model and the same eleven-symbol ceiling assumption in both windows.

Observation	Busy window	Quiet window
active UEs in comparison scope	31	5
mean scheduled PRB for test UE	34	210
scheduled MCS / table	24 / table 1	24 / table 1
scheduled rank	2	2
initial DL BLER	7%	6%
application throughput	45 Mbit/s	270 Mbit/s

What the exercise proves:

- RF efficiency did not deteriorate in proportion to the throughput;

- the test UE received materially fewer resources in the busy window;
- throughput follows allocation in this comparison.

Defensible classification:

C-010 is an allocation-limited interval in the synthetic exercise.

It does not prove that the cell has a design-capacity defect. Active users, offered traffic, QoS, control resources and scheduler policy still need their own evidence.

9. Synthetic case C-011 - the transport/test-path branch

Evidence status: synthetic controlled path comparison.

Held constant:

- same UE;
- same cell and radio position;
- same direction, protocol, stream count and duration;
- stable radio allocation and outcome.

Changed:

- server and the path beyond the operator edge.

Repetition	Local known-good server	Remote path
1	250 Mbit/s	93 Mbit/s
2	255 Mbit/s	95 Mbit/s
3	252 Mbit/s	94 Mbit/s

What the exercise proves:

- the UE and radio can deliver substantially more than 94 Mbit/s in the controlled window;
- the low ceiling follows the remote path or endpoint.

Defensible classification:

C-011 localises the next investigation beyond the radio path used by the local test.

It does not identify a 100 Mbit/s interface, a router, a firewall or the remote server as root cause. That requires interface counters, topology and endpoint evidence.

10. Synthetic case C-012 - the policy/configuration branch

Evidence status: synthetic two-UE/two-SIM matrix with a synthetic policy record.

Test	UE-A	UE-B
SIM-X	98 Mbit/s	101 Mbit/s
SIM-Y	250 Mbit/s	255 Mbit/s

Held constant:

- radio position and cell;

- server and test method;
- software state and thermal pre-check;
- supported band combination and realised aggregation;
- time window selected for comparable load.

Additional synthetic evidence:

- SIM-X session record contains an applicable aggregate rate limit consistent with approximately 100 Mbit/s;
- SIM-Y session record does not contain that limit;
- no corresponding radio allocation or error ceiling appears at 100 Mbit/s.

What the exercise proves:

- the ceiling follows SIM-X across two validated UEs;
- an independent policy record supports the same limit.

Defensible classification:

C-012 is consistent with a subscription or session policy limit in the synthetic exercise.

Without the policy record, the matrix would only localise the issue to the subscription/session branch. The repeated 100 Mbit/s plateau alone would not prove enforcement.

11. Build an evidence matrix before writing the root cause

Use one row per hypothesis.

Hypothesis	Evidence that supports it	Test that can reject it	Limit of conclusion
RF efficiency	quality, MCS/layers and HARQ outcome degrade together	same path at known-good radio position	does not name physical RF source
scheduler/allocation	demand healthy, RF healthy, grants/PRB fall with throughput	same setup in controlled quiet/load window	does not explain scheduler decision
transport/path	radio can deliver more to local endpoint, remote path remains low	move one endpoint or boundary at a time	does not identify the exact hop
policy/session	ceiling follows SIM/session and policy record agrees	UE/SIM cross-matrix or controlled policy change	plateau alone is insufficient
UE capability/config	ceiling follows UE; capability or realised configuration differs	same SIM in validated UE and capability-state comparison	capability support is not actual scheduling
application/server	endpoint CPU, session or protocol limits result	alternate known-good endpoint with same path	does not prove network health outside test window

The useful output of the matrix is not a long list of possible causes. It is the next test with the highest power to remove branches.

Last positive, first divergence

Write two sentences:

- last positive: the last boundary where the expected volume or state is established;
- first divergence: the next boundary where the observed value or state no longer matches the controlled comparison.

Examples:

- “PDCP and N3 volumes remain consistent with the known-good radio interval; application goodput diverges only on the remote endpoint path.”
- “Traffic demand remains saturated and RF efficiency is comparable; scheduled PRB diverge in the busy interval.”

These statements are stronger than “transport issue” or “capacity issue” because they expose the evidence boundary.

12. Field workflow

Step 1 - Freeze the test contract

Record direction, endpoint, path, protocol, streams, duration, UE state and clock.

Step 2 - Prove demand

Show that data was available to send during the analysed interval.

Step 3 - Compare throughput definitions

Label application goodput, RAN UE throughput, interface rate and any PM value with its layer, numerator, denominator and window.

Step 4 - Read radio efficiency and allocation separately

Radio efficiency asks how successfully scheduled resources carried data. Allocation asks how many resources the UE received.

Step 5 - Move one boundary

Change only the radio position, load window, endpoint, SIM or UE according to the branch being tested.

Step 6 - Preserve the known-good comparison

Keep the successful trace. A “good” screenshot without the same contract is not a baseline.

Step 7 - Stop at the proven boundary

If the ceiling follows a remote path, report a path-localised problem. Do not name a router without router evidence.

Step 8 - Measure the post-check

After an authorised correction, repeat the original contract and preserve both before and after windows.

13. Evidence checklist

Before publishing a low-throughput diagnosis:

- ☐ DL and UL are reported separately.
- ☐ The server, path, protocol, streams and duration are recorded.
- ☐ UE software, thermal state, power and local interface are checked.
- ☐ Traffic demand is proven for the measurement interval.
- ☐ Every throughput value has a layer, numerator, denominator and time window.
- ☐ Application goodput is not treated as identical to RAN UE throughput.
- ☐ Clocks are normalised across UE, RAN, transport, core and application.
- ☐ Carrier, BWP, aggregation, layers and actual scheduling state are preserved.
- ☐ CQI is not presented as the scheduled MCS.

- ❑ MCS is not presented as throughput.
- ❑ PRB, symbols, layers, TBS and grant frequency are included where available.
- ❑ Initial BLER/HARQ is distinguished from upper-layer recovery.
- ❑ RLC mode is known before claiming RLC retransmission.
- ❑ PDCP delivery, discard or duplication is interpreted only in its configured context.
- ❑ UL SR, BSR, grants and power evidence are checked for UL cases.
- ❑ NSA bearer and NR contribution are proven rather than inferred from an icon.
- ❑ SA PDU Session, QoS Flow/DRB, N3, UPF and Data Network path are mapped.
- ❑ A known-good comparison uses the same test contract.
- ❑ Only one useful variable changes in each discriminating test.
- ❑ A flat plateau is not called policy without independent policy/session evidence.
- ❑ A path-localised result is not turned into an exact-hop root cause.
- ❑ Synthetic cases and values remain visibly declared.

Takwa Sebai working criterion

I do not sign a throughput root cause until traffic demand is proven and the ceiling follows one controlled boundary: radio position, allocation window, endpoint path, SIM/session or validated UE. The reported cause is that boundary. Naming a specific interferer, router or policy requires its own evidence.

Chapter takeaway

Low throughput becomes manageable when the ceiling is forced to choose a branch.

- 1 validate the traffic test;
- 2 prove demand;
- 3 name every measurement point and denominator;
- 4 separate radio efficiency from allocation;
- 5 split DL from UL;
- 6 move one boundary at a time;
- 7 compare with a known-good interval;
- 8 report only the boundary the evidence proves.

Do not ask which counter is low. Ask where useful data stops progressing, and which controlled test proves it.

References

- ¹ ETSI, [TS 138 214 V18.9.0 - NR; Physical layer procedures for data](#), especially clauses 5.1.3.1, 5.1.3.2 and 6.1.4.1.
- ² ETSI, [TS 138 321 V18.9.0 - NR; Medium Access Control protocol specification](#), especially clauses 5.3 and 5.4.
- ³ ETSI, [TS 138 322 V18.2.0 - NR; Radio Link Control protocol specification](#), especially clauses 4.2.1 and 5.2.
- ⁴ ETSI, [TS 138 323 V18.5.0 - NR; Packet Data Convergence Protocol specification](#), especially clauses 5.2 to 5.6 and 5.11.
- ⁵ ETSI, [TS 128 552 V18.12.0 - Management and orchestration; 5G performance measurements](#), especially clauses 5.1.1.3.1 and 5.1.1.3.3.
- ⁶ ETSI, [TS 123 501 V18.12.0 - System architecture for the 5G System](#), especially clauses 5.7.1.8 and 5.7.2.6.
- ⁷ ETSI, [TS 138 306 V18.9.0 - NR; UE radio capabilities](#).

REAL CASE · ILLUSTRATIVE LAB SCENE · EVIDENCE-BOUNDED

REQUESTED STATE IS NOT ACTIVE STATE

FIELD DECISION RULE
VERIFY THE ACTIVE LAYER

LAB CONTEXT · ILLUSTRATIVE


Observable target before control

Same carrier · different PCI




END-TO-END EVIDENCE CHAIN

Follow the state transition, not the command narrative



THE DEFENSIBLE ACCEPTANCE CHAIN



STOP AT THE FIRST MISSING PROOF.

© Takwa Sebai · takwasebai.com · Reuse with attribution

Figure V-013. Real LTE lab case reconstructed and analysed by Takwa Sebai. The laboratory scene is illustrative, not a raw evidence image. LTE is used here because the configured-versus-active validation rule applies across LTE and NR.

Sign off the observed state, not the control-path activity

I reconstructed this case from timestamped traces and contemporaneous engineering notes. Exact operator, device, carrier, PCI and proprietary command values are withheld. The positive path is supported by contemporaneous engineering notes; the failed path also has timestamped raw evidence.

Checkpoint	Observed evidence	Decision
1. Initial control path	Transaction matched, but the requested function was rejected before radio mutation.	Stop. No lock claim and no rollback required.
2. Natural recamp	Serving PCI changed on the same carrier while registration remained available and recorded RSRP was stable.	Do not attribute the PCI change to the rejected request.
3. Alternate control	The restriction was written and refreshed, but the active serving state varied while two diagnostic processes remained active.	Configured state existed; active enforcement was not yet repeatable enough to accept.
4. Controlled refresh	The two processes were paused for the refresh window, then resumed. Target serving state, registration and unrestricted selection after rollback were recorded.	Accept the bounded procedure on the tested stack. Internal mechanism remains unproven.

WHAT I WOULD SIGN OFF The requested LTE target state was observed after the controlled refresh, and unrestricted selection was observed after rollback. Registration remained available in the recorded checks.	WHAT REMAINS UNKNOWN The internal modem mechanism, the individual effect of either diagnostic process, portability to other modem families and application performance were not proved. IN_SERVICE establishes registration only.
---	--

THE FIVE-PROOF ACCEPTANCE CHAIN

1. CONFIGURED What was requested?	2. ACTIVE What serving state changed?	3. DISCRIMINATING Could another outcome occur?	4. SERVICE What remained available?	5. ROLLBACK What proved release?
--------------------------------------	--	---	--	-------------------------------------

Field stop rule: if one proof is missing, report the first missing boundary. Do not replace it with a plausible mechanism. Rollback means unrestricted selection, not a forced return to the identical starting PCI.

Printable field checklists

These two-sided checklists compress the eight chapter gates. Print them, write the test interval at the top, and do not tick an item that is merely assumed.

1 - Site to stack

- ☐ Affected site, sector, cell, UE/service and interval are fixed.
- ☐ Physical path, power, earth, timing, fibre and transport checks are scoped.
- ☐ Current and historical alarms are aligned with the first symptom.
- ☐ The first divergent physical or logical boundary is named.
- ☐ A known-good sector, route, UE or interval exists.
- ☐ The next test can reject at least one credible branch.

2 - Numerology and TDD

- ☐ Band, bandwidth, duplex mode, SCS and active BWP are recorded.
- ☐ The TDD pattern and usable DL/UL opportunity are decoded.
- ☐ PRBs, usable symbols, MCS table, layers and retransmissions are known.
- ☐ Load, UE capability and aggregation state are controlled.
- ☐ The observed goodput is below a declared physical ceiling.
- ☐ An SCS claim is separated from bandwidth, RF, load and path changes.

3 - Bandwidth parts

- ☐ Configured DL and UL BWP sets are inventoried separately.
- ☐ Initial, first-active and default roles are not confused.
- ☐ The last applicable RRC, DCI, timer or RACH event is timestamped.
- ☐ Active state is proved for the same direction and test instant.
- ☐ Scheduling or device evidence corroborates the reconstructed state.
- ☐ Any trace gap is reported as unknown.

4 - SSB and beam evidence

- ☐ Cell-level and per-SSB-index results are kept separate.
- ☐ SMTC, report configuration, filtering and reference resources are recorded.
- ☐ SSB index is not presented as a physical azimuth without site evidence.
- ☐ Lower-layer beam state is proved or explicitly unavailable.
- ☐ PDCCH/PDSCH outcome shares the measurement timeline.
- ☐ The report stops at an SSB observation when deeper evidence is absent.

5 - PDCCH to PDSCH

- ☐ Active BWP, CORESET and Search Space configuration are complete.
- ☐ Monitoring occasions, candidates, aggregation level and RNTI are known.
- ☐ Network scheduling decision is separated from UE DCI detection.
- ☐ DCI fields are followed into PRBs, symbols, MCS, layers and TBS.
- ☐ PDSCH CRC, BLER, HARQ and delivered bytes use the same assignment.
- ☐ A UE-only decode gap is not called a missed transmission.

6 - Multi-KPI reading

- ☐ UE, cell, carrier, BWP, direction and reference resource match.
- ☐ Clocks, windows, filters and sample populations are reconciled.
- ☐ CQI report assumptions are separated from scheduled MCS.
- ☐ Reported RI is separated from actual PDSCH layers.
- ☐ PRBs, scheduling duty cycle, BLER and delivered bytes align.
- ☐ Correlation is not called causality without a controlled comparison.

7 - Mobility

- ☐ NSA/SA architecture and the exact procedure are named.
- ☐ Source, target and UE identifiers map to one attempt.
- ☐ Last positive and first missing preparation evidence are identified.
- ☐ Command delivery, target synchronisation and RACH are separate gates.
- ☐ Completion, path switch and post-handover service are verified.
- ☐ Timer expiry localises a phase but is not promoted to a physical cause.

8 - Low throughput

- ☐ Direction, server, path, protocol, streams and duration are fixed.
- ☐ Traffic demand and terminal thermal/interface state are proved.
- ☐ Every throughput value has a layer, numerator, denominator and window.
- ☐ Radio efficiency and scheduler allocation are analysed separately.
- ☐ CA/EN-DC contribution is proved per component or bearer.
- ☐ One boundary changes per test and the cause stops at that boundary.

Continue the work

Use the printable checklists on the preceding pages during a real review. Keep the completed sheet with the trace timestamps and the known-good comparison; that record is more useful than an unscoped screenshot.

Challenge a case with Takwa Sebai

If you have a difficult RAN validation or RCA case, prepare an anonymised one-page evidence brief:

- architecture, direction and service symptom;
- exact interval and common clock;
- last positive evidence and first divergence;
- one known-good comparison;
- hypotheses already tested;
- the decision that still needs review.

Do not send customer, operator, subscriber or proprietary identifiers without the required permission.

Find the next resource

- Portfolio and field guides:

takwasebai.com

- Professional contact:

takwasebai.com/#contact

- LinkedIn:

linkedin.com/in/takwa-sebai

- Substack:

substack.com/@takwasebaihicelltek

Author

Takwa Sebai is a telecom engineer and technical manager with more than ten years across mobile networks, specialising in 4G/5G RAN validation, OSS, network KPIs, industrialisation and automation. Her working rule is simple: configuration is not state, correlation is not cause, and an RCA must stop where the evidence stops.

Copyright Takwa Sebai. All rights reserved. Personal reading and internal evaluation are permitted. Redistribution, resale, removal of attribution or reuse in another training product requires written permission.