



[Frame 6335] Downlink burst 1 -- CU-UP hands the DU two voice packets. The CU-UP (user-plane half of the gNB Central Unit) forwards a downlink PDCP-NR PDU to the DU over F1-U. On the wire this is a bare GTP-U packet on TEID 0x06 whose NR RAN Container extension header (type 0x84) marks an NR-U PDU Type 0 -- DL User Data (TS 38.425). The payload T-PDU carries PDCP SN 0 for this voice DRB. Note the NR-U Sequence Number reads 0: this srsRAN build leaves the NR-U SN unused, so the meaningful sequence lives in the PDCP SN inside the T-PDU, which the DU will echo back in its delivery status.

[Frame 6336] Second half of the burst – PDCP SN 1, sent back-to-back with the same capture timestamp. The DU now has two SDUs queued for the air interface. Downlink voice arrives in these small paired bursts, roughly on the codec's 20 ms cadence.

[Frame 6337] DU flow-control feedback -- DL Data Delivery Status. The direction reverses: the DU answers on TEID 0x18 (the uplink half of the same DRB's F1-U tunnel pair) with an NR-U PDU Type 1, DL Data Delivery Status (DDDS). It reports Highest Transmitted NR PDCP SN = 1 -- the DU has now pushed SN 0 and 1 out over the radio -- and advertises a Desired Buffer Size of 6,172,672 bytes (~5.9 MiB), telling the CU-UP how far ahead it may send. Frame is not final: the tunnel stays open. This is the return leg of the NR-U flow-control loop that keeps the CU-UP and DU in step across the higher-layer split.

[Frame 6348] DDDS heartbeat. No new voice data has arrived, so the report is unchanged: Highest Tx SN still 1, buffer still ~5.9 MiB. The DU re-sends DDDS on a timer (~10 ms here), not only when something changes -- flow control on F1-U is a steady keepalive, not purely event-driven.

[Frame 6349] Heartbeat continues -- identical Highest Tx SN = 1. Two DDDS landed within the same millisecond, so the cadence is not perfectly periodic; the DU emits them as its scheduler ticks.

[Frame 6351] Heartbeat -- Highest Tx SN = 1. Nothing to acknowledge yet.

[Frame 6362] Heartbeat -- Highest Tx SN = 1. This repetition is exactly why DDDS dominates the capture: across the full ~82 s trace 36,339 of 37,871 frames (96%) are these upstream status reports, and only 1,513 carry actual downlink data.

[Frame 6365] Last heartbeat at Highest Tx SN = 1, immediately before the next voice burst.





