Edge Deployment Verdict: YOLO11n

- Based on the comparative analysis of computation complexity, thermal reliability, and real-time safety requirements, YOLO11n is the recommended architecture for roadside edge deployment. While RT-DETRv2-S provides deterministic latency and high precision, its architectural weight makes it unfeasible for low-power embedded hardware at the 30 FPS safety standard.
    - **Real-Time Safety & Hardware Portability**: To maintain safety-critical traffic sensing, the system must process frames in under the 33.3ms hard limit. While RT-DETRv2-S fails this requirement on edge hardware at ~108ms, YOLO11n maintains a massive scaling buffer with a project edge latency of ~25ms
    - **Kinematic Advantage**: On projected edge hardware (~25ms), the YOLO11n system reacts 8x to 10x faster than a typical alert human driver. This ensures a reaction occurs within 39cm (~1.2 feet) of vehicle movement at 35 mph, maintaining a critical safety margin for urban intersections.
    - **Thermal Reliability**: YOLO11n requires only 6.5 GFLOPs of mathematical work per frame, compared to the 60 GFLOPs demanded by RT-DETRv2-S. High GFLOP models utilize a larger surface area of the GPU's processing cores, which spikes the wattage required and acts as a heat source that can lead to thermal throttling in constrained roadside cabinets.
    - **Sustainability & Scale**: The low GFLOP count and parameter size of YOLO11n allow for smaller, more sustainable hardware footprints that can run on low-power edge devices or solar-assisted power grids. Conversely, attempting to throttle a heavier model like RT-DETRv2-S to decrease power consumption would only result in increased latency that violates the 33.3ms real-time safety threshold.

NOTE: See 'Latency and Power Analysis for YOLO11n and RT-DETRv2-S' for more in depth details about this decision.