GRAVITY
DIGITAL SOLUTIONS

PRODUCT DATASHEET

Neural Compress Pro

Compression for the AI era of tolling.

A **software-only sidecar** that compresses toll imagery by **92 to 94%** while maintaining the visual fidelity downstream systems require, from ALPR today to AI workloads tomorrow. Built on patent-pending neural compression architecture by Gravity Digital Solutions.

92–94%	+15.1pp	0
STORAGE REDUCTION	PEAK OCR UPLIFT	HARDWARE REFRESH

THE PROBLEM

Tolling runs on images. Every misread plate is leaked revenue. Every manual review is operational cost. Every compressed archive is the training data for tomorrow's AI workloads, whether agencies plan for it or not.

Conventional compression is scene-blind. The industry has accepted a three-way tradeoff between image quality, storage cost, and recovered revenue for over a decade.

THE TECHNOLOGY

NCP is a learned codec. Rather than applying a uniform compression rule across every region of the image, NCP's encoder learns where the signal lives and allocates bits proportional to information content.

The mechanism is **binary spherical quantization in the latent space**, a patent-pending architecture purpose-built for transportation imagery. Traditional codecs are scene-blind. NCP is scene-aware. *That is the unlock.*

REVENUE

Recover leaked revenue.

Higher OCR confidence on every plate. Fewer manual reviews, fewer misreads, fewer plates falling through to leakage.

AI READINESS

Your archive becomes a strategic asset.

Training-grade imagery preserved at 92 to 94% smaller footprint. Tokens consumed natively by downstream AI workloads.

PRIVACY BY DESIGN

Pseudonymized at rest.

Compressed representation is unreadable without the decoder model. Privacy is structural, not bolted on.

PERFORMANCE

Compression ratio	92–94% reduction
Image quality (PSNR)	~38 dB
Image quality (SSIM)	~0.92
OCR confidence uplift	up to +15.1pp
Visual fidelity	Indistinguishable

Measured under controlled lifecycle testing. Same recognition pipeline, same input. Compression is the only variable.

TWO INSERTION POINTS

01 Recognition tier

Plugs in front of edge inference, alongside any incumbent ALPR. Vendor-additive, not vendor-replacement. Higher confidence reads, fewer manual reviews.

02 Storage tier

Sits between recognition and archive. 92 to 94% storage reduction with no quality loss. Compressed representation is already tokenized for downstream AI.

> Both tiers are independently available. Most agencies start with the storage tier for lower integration friction. Recognition tier pilots are equally available for agencies ready to integrate.

THE PILOT MODEL

30-day shadow deployment with tolling agencies and integrators. Originals retained throughout. Compressed and original run through the same recognition pipeline. **Per-plate confidence deltas measured across the full sample, with regressions reported alongside improvements.** Results published jointly.

Karthik Sivakoti
karthiksivakoti@gravitydigitalsolutions.com
GRAVITY DIGITAL SOLUTIONS · AUSTIN, TX

PATENT-PENDING ARCHITECTURE
Now piloting with agencies & integrators