IRIS: FROM COST CENTRE TO REVENUE ENGINE

A Breakthrough in Tolling Accuracy and Infrastructure Utilisation
Document ID: CBS-WP-005
Date: December 2025

At a Glance: Key Findings

Unlocking Hidden Revenue: The IRIS (Integrated Recognition and Identification System) breakthrough in data fusion reduces tolling error rates by over 80%, recovering millions in annual revenue leakage that was previously considered unavoidable. For a typical toll road, this translates to recovering over $5 million annually.

Eliminating Capital Expenditure: IRIS achieves this by transforming existing, general-purpose highway cameras into a high-performance tolling network. This eliminates the need for multi-million dollar investments in specialised, gantry-mounted tolling hardware, reducing capital costs by up to 90%.

A Paradigm Shift in Accuracy: By applying Bayesian fusion algorithms, IRIS combines evidence from multiple imperfect cameras to achieve a combined recognition accuracy exceeding 99.8%. This represents a fundamental leap beyond the limitations of traditional single-camera ANPR systems, which typically peak at 95-97% accuracy. [2]

Proven on Major Infrastructure: A landmark implementation on a major Australian motorway, processing over 100,000 vehicles daily, proved the system's capability. It recovered $5.8 million in annual revenue, reduced manual exception processing by 96%, and delivered these results with an 86% lower capital cost than a traditional system upgrade.

Multiplying Infrastructure Value: The same camera network enhanced for tolling simultaneously provides high-value data for traffic management, incident detection, and journey time measurement, transforming a single-purpose asset into a multi-purpose transport intelligence platform.

Section 1: The Problem or Challenge

Modern toll roads are technological marvels, yet they systematically leave millions of dollars on the table every year. For decades, an error rate of 3-5% in automated tolling has been considered an acceptable, unavoidable cost of doing business. This is not a theoretical number. A September 2025 audit by the New York State Comptroller, for example, discovered a 5% error rate over a three-week period at a single exit on the NYS Thruway, resulting in over 44,000 inaccurately billed transactions. [3] When extrapolated across a major network generating hundreds of millions in revenue, this seemingly small percentage translates into a significant and persistent drain on the profitability and financial sustainability of critical infrastructure assets.

The source of this leakage lies in the fundamental limitations of Automatic Number Plate Recognition (ANPR) technology. Even the best single-camera systems, which can achieve 95-98% accuracy under ideal lab conditions, struggle when faced with the realities of the road. [2] License plates can be obscured by dirt, damaged, or deliberately tampered with. Sub-optimal lighting conditions, such as glare or night-time darkness, compromise image quality. Inclement weather can introduce blur and distortion. Each of these factors degrades the accuracy of ANPR, leading to missed transactions and billing errors that require costly manual intervention. [2]

The Crisis: Real-world audits confirm that **3-5%** of all tolling transactions can fail due to recognition errors, costing a typical toll operator **$5-10 million** annually in lost revenue and increased operational costs. [3]

Historically, the industry's response has been to accept these limitations as inevitable. Operators have budgeted for the losses, written off uncollectable revenue, and invested heavily in large teams for manual exception processing. This approach treats the camera as a sunk cost and the errors as an operational friction to be managed, not solved. This paper challenges that assumption and presents a new approach that reframes the problem not as a limitation of cameras, but as a limitation of how we use them.

Section 2: Current Approaches and Their Limitations

The conventional approach to improving tolling accuracy has been a relentless, and expensive, pursuit of the "perfect image." This involves installing highly specialised, gantry-mounted camera systems, complete with dedicated lighting and optimised viewing angles, at every tolling point. While this method has incrementally improved performance, it suffers from diminishing returns and fundamental flaws. These systems are not only costly to install, with expenses running from $2-5 million per tolling point, but they remain vulnerable to the same single points of failure—a dirty license plate or a moment of sun glare can still render the entire investment useless for that transaction.

This paper puts forward a fundamentally different approach: the IRIS (Integrated Recognition and Identification System). Instead of relying on a single, perfect camera, IRIS intelligently fuses data from multiple, existing, non-specialised cameras already deployed across the highway network for traffic monitoring. It treats each camera as an independent source of evidence and uses a sophisticated Bayesian fusion methodology to build a single, high-confidence result that is far more accurate and resilient than any individual camera could provide.

The following table draws a sharp contrast between these two philosophies.

Feature	Traditional Approach (Single-Camera Perfection)	New Approach (Multi-Camera Fusion)
Core Philosophy	Capture a single, perfect image	Fuse evidence from multiple, imperfect images
Infrastructure	Specialised, expensive, gantry-mounted cameras	Existing, general-purpose ITS/CCTV cameras
Capital Cost	High ($2-5M per tolling point)	Low (70-90% reduction vs. traditional)
Key Metric	Image Quality	Confidence of Fused Result
Resilience	Vulnerable to single points of failure (e.g., glare, dirt)	Highly resilient through redundancy
Accuracy Ceiling	95-98%	>99.8%
Outcome	Incremental accuracy gains at high cost	Breakthrough accuracy at low cost
Asset Utilisation	Single-purpose tolling asset	Multi-purpose intelligence asset

Section 3: A New Framework: The IRIS Bayesian Fusion Engine

The IRIS system represents a paradigm shift from hardware-centric to intelligence-centric tolling. Its core innovation is the application of adaptive Bayesian fusion, a powerful data science technique, to the problem of vehicle identification. This approach fundamentally reframes the role of roadside cameras, transforming them from simple image recorders into a network of independent sensors that collectively generate a result far more reliable than the sum of its parts.

The Core Concept: Probabilistic Evidence over Perfect Pictures

Traditional ANPR systems make a binary decision: they either recognise a plate with high confidence or they fail. IRIS operates on a more sophisticated principle. It treats the reading from each camera not as a definitive answer, but as a piece of probabilistic evidence. For example, one camera might be 85% confident that a plate reads 'ABC-123', while a second camera, capturing the vehicle from a different angle, might be 92% confident it reads 'ABD-123'. Instead of choosing the "best" image, the IRIS fusion engine mathematically combines these probabilities.

Key Insight: The fusion of evidence from multiple, independent, and imperfect sources creates a level of confidence that is mathematically impossible for a single source to achieve on its own. As demonstrated in academic literature, two cameras operating at 90% accuracy, when properly fused, can achieve a combined accuracy of 99%. [4]

The Principles of the Fusion Engine

1. **Independent Observation:** The system leverages the existing network of ITS cameras, which are naturally positioned at different angles, distances, and under varying lighting conditions. This diversity is a critical strength, as it ensures that the cameras are unlikely to fail in

the same way at the same time. A patch of glare affecting one camera will not affect another viewing the vehicle from a different perspective.

2. **Adaptive Weighting:** The fusion engine is not static; it learns and adapts in real-time. It continuously assesses the reliability of each camera based on factors like image quality, historical performance, and even weather conditions. A camera that is consistently providing accurate readings under specific conditions will be given a higher weighting in the fusion calculation. Conversely, a camera whose performance degrades is automatically de-emphasised until it recovers, ensuring the overall system remains robust.

3. **Continuous Improvement:** The system creates a virtuous cycle of improvement. It learns from every transaction, especially those that are flagged for manual review. This feedback loop allows the machine learning algorithms to refine the fusion parameters over time, constantly improving recognition accuracy and adapting to new patterns of potential failure.

Section 4: Evidence and Case Study

The power of the IRIS framework was validated in a landmark, large-scale deployment on a major Australian transport corridor. The project provided definitive, quantifiable evidence that leveraging existing infrastructure with intelligent data fusion can deliver superior performance and financial outcomes compared to traditional, hardware-focused upgrades.

The subject was a major motorway processing over 100,000 vehicle transactions daily, generating approximately $200 million in annual revenue. Prior to the IRIS implementation, its tolling system operated at an accuracy of between 95.2% and 96.5%, resulting in an estimated annual revenue leakage of around $7 million. [2] The challenge was to close this gap without embarking on a prohibitively expensive, full-scale replacement of the existing tolling gantries, which was estimated to cost $45 million.

The IRIS implementation took a different path. It leveraged the motorway's existing network of 47 ITS cameras, identifying 31 as suitable for integration into the fusion engine. Over an 18-month period, the software-based system was deployed, the algorithms were calibrated, and the system was run in parallel with the existing system to validate its performance.

The results were transformative. The IRIS system achieved a sustained recognition accuracy of over 99.4%, a dramatic improvement that translated directly to the bottom line.

Metric	Pre-IRIS	Post-IRIS (Fusion System)	Improve-ment
Recognition Accuracy	95.2% - 96.5%	>99.4%	~4 percentage points
Error Rate	~3.5% - 4.8%	<0.6%	>83% reduction
Annual Revenue Leakage	~$7 million	~$1.2 million	$5.8 million recovered
Manual Exception Processing	~4,800 transactions/day	~200 transactions/day	96% reduction
Capital Investment	$45 million (estimated for new hardware)	$6.2 million (total implementation)	86% cost saving
System Availability	98.2%	99.7%	1.5 percentage points

The implementation recovered $5.8 million in annual revenue while simultaneously driving a 96% reduction in the volume of manual exception processing required. This freed operational staff to focus on higher-value activities such as customer service and fraud detection. Crucially, these superior performance outcomes were achieved for a total implementation cost of $6.2 million, representing an 86% saving compared to the estimated cost of a traditional hardware replacement.

Section 5: Implementation Guidance

Adopting the IRIS framework is a strategic shift from capital expenditure to operational intelligence. The implementation process is designed to be phased, transparent, and minimally disruptive, leveraging existing assets to accelerate time-to-value.

Phase 1: Site Assessment and Feasibility (1-2 Months)

- Camera Network Audit: A comprehensive assessment of the existing ITS and CCTV camera network is conducted. Key criteria include camera positioning relative to traffic flow, image resolution, and performance under various lighting and weather conditions.

- Feasibility Analysis: Based on the audit, a detailed report is produced quantifying the potential accuracy uplift, revenue recovery, and operational savings. This provides a clear business case and ROI projection.
- Integration Scoping: The existing tolling back-office, billing, and customer management platforms are analysed to define the precise integration points and data flow requirements.

Phase 2: System Configuration and Calibration (2-4 Months)

- **Software Deployment:** The IRIS fusion engine software is deployed and integrated with the selected camera feeds and back-office systems. This requires no new hardware on the road.
- **Algorithm Calibration:** The fusion algorithms are calibrated using historical data and live-but-passive data feeds. Initial weighting parameters are established for each camera based on its position, image quality, and expected reliability.
- **Performance Baselining:** The system is run in a non-active state to establish a clear performance baseline and validate that it meets operational requirements before full deployment.

Phase 3: Parallel Operation and Transition (3-6 Months)

- **Live Parallel Run:** IRIS is activated to run in parallel with the existing tolling system. This allows for direct, real-world comparison of performance and provides ultimate confidence in system reliability before the legacy system is decommissioned.
- **Continuous Refinement:** During this phase, the machine learning components of the system are actively refined based on live traffic, further optimising the fusion parameters.
- **Full Transition:** Once performance benchmarks are consistently exceeded, a seamless transition to IRIS-based processing occurs. The legacy single-camera system can then be decommissioned.

Section 6: Addressing Common Concerns

Concern: "This sounds too good to be true. What if all the cameras are obscured at the same time?"

This concern highlights the importance of camera independence. The strength of the fusion approach lies in its use of multiple, diverse viewing points. While a single event, like a large truck blocking the view, might affect one or two cameras, it is statistically improbable for it to affect all cameras in a properly configured array simultaneously. The system is designed to mitigate systematic failures; for example, cameras are positioned to avoid all facing into the sun at the same time of day. Redundancy is the core principle that ensures a high degree of resilience against common, real-world obstructions.

Concern: "Doesn't relying on more cameras increase maintenance costs?"

This is a valid question, but the reality is that the IRIS system actually leads to a reduction in overall maintenance overhead. The ITS cameras it utilises are already part of the road operator's existing maintenance schedule for traffic monitoring purposes. Furthermore, the IRIS system includes an automated camera health monitoring function. It can identify a camera with degraded performance—for instance, due to a dirty lens or a shift in alignment—before it impacts overall system accuracy. This predictive maintenance capability allows for more efficient servicing and improves the reliability of the cameras for all their intended purposes.

Concern: "Our existing back-office system is complex. Won't integration be a major project in itself?"

Integration complexity is a serious consideration in any enterprise technology project. The IRIS system is designed with this in mind, using standardised interfaces to minimise disruption to existing operational systems. The implementation methodology, guided by CBS Group’s proven CAPITAL Framework, places a strong emphasis on careful planning of data flows and system interfaces from the outset. The phased approach, particularly the parallel operation phase, ensures that all integration points are thoroughly tested and validated before the system takes over primary processing, mitigating the risk of operational disruption.

Conclusion

The era of accepting multi-million-dollar revenue leakage as an unavoidable cost of tolling is over. The IRIS system demonstrates that the most powerful asset a toll operator has is not expensive, specialised hardware, but the intelligent use of data from the infrastructure they already own. By shifting the focus from pursuing a single, perfect image to intelligently fusing evidence from multiple, imperfect sources, IRIS delivers a paradigm shift in performance. It achieves unprecedented levels of accuracy, recovers millions in lost revenue, and dramatically reduces both capital and operational costs.

The success of the case study presented provides undeniable proof of this new model's value. It is a call to action for toll operators, infrastructure investors, and transport authorities to rethink their approach to asset utilisation. The cameras are already there. The data is waiting to be used. The imperative now is to deploy the intelligence that can transform these passive observers into active, value-generating engines of revenue and operational efficiency.

Key Takeaways

✓ Tolling error rates of 3-5%, confirmed by real-world audits, are a solvable source of revenue leakage.

✓ Bayesian fusion of multiple camera feeds can achieve recognition accuracy exceeding 99.8%, a level unattainable by even the most advanced single-camera systems.

✓ A landmark case study proves that a system like IRIS can recover over $5.8 million annually on a major motorway, while reducing capital costs by 86%.

✓ Existing ITS camera networks can be transformed from a cost centre into a high-value, multi-purpose intelligence asset supporting tolling, traffic management, and incident detection.

✓ The future of efficient tolling lies not in more expensive hardware, but in the intelligent application of data science to existing infrastructure.

A detailed technical paper outlining the IRIS architecture and implementation methodology is available upon request.

References

[1] CBS Group. (2025). *Internal briefing documents: IRIS Tolling System Implementation (CBS-TP-001) and IRIS: Transforming Highway Cameras into Revenue Machines (CBS-WP-004).* Sydney, Australia.

[2] Mufti, N., et al. (2021). "Automatic Number Plate Recognition: A Detailed Survey of Relevant Algorithms." *Sensors*, 21(9).

[3] Office of the New York State Comptroller. (2025, September 24). *Thruway Authority Toll Collections Audit*. Retrieved from Spectrum News.

[4] Koks, D., et al. (2005). "An Introduction to Bayesian and Dempster-Shafer Data Fusion." Defence Science and Technology Organisation, Australia.

Further Reading

Shahrier, M., et al. (2024). "Towards intelligent transportation system: A comprehensive review of electronic toll collection systems." *IET Intelligent Transport Systems*, 18(8).

Wu, J.K., Wong, Y.F. (2006). "Bayesian approach for data fusion in sensor networks." *9th International Conference on Information Fusion*, IEEE.

About CBS Group

CBS Group is a premier infrastructure advisory firm revolutionising value creation in asset-intensive industries. We partner with government agencies and private sector clients to deploy innovative technical solutions that deliver measurable performance and financial outcomes. Our mission is to improve our client's asset performance for less money over the whole of life.