THE LOSS PREVENTION MATURITY MODEL

A Strategic Framework for Evolving LP Systems from Devices to Autonomous Intelligence

Agilence®

TABLE OF CONTENTS

Introduction

A Tipping Point for Loss Prevention Operations

Loss prevention has evolved significantly in the 2020s. Retailers, grocers, and restaurant operators face a broader and more volatile threat landscape than at any point in the past decade. External theft has become more brazen, with shoplifting incidents nearly doubling since 2019 and 73% of retailers reporting increasingly violent behavior during theft events. Organized retail crime (ORC) has evolved into coordinated, multi-location networks that use online marketplaces to rapidly liquidate stolen goods. These developments have been coupled with higher prosecution thresholds, making deterrence and case building more difficult. Self-checkout adoption combined with staff shortages has created an expanded attack surface for both organized and opportunistic theft. Meanwhile, the surge in e-commerce and omnichannel fulfillment has opened new attack surfaces such as BOPIS, curbside pickup, cross-channel returns, and digital payment fraud—each introducing vulnerabilities that traditional LP workflows were never designed to manage.

Rapid expansion of LP complexity has outpaced the capacity of most organizations to respond effectively.

At the same time, long-standing loss vectors still demand attention: internal theft, returns abuse, vendor fraud, administrative errors, loyalty program fraud, gift card schemes, and more. Each requires different detection methods and different data sources. Customer data, transaction records, and inventory information often remain fragmented across platforms that were never designed to communicate with each other. The result is a rapid expansion of LP complexity that has outpaced the capacity of most organizations to respond effectively.

Operational Pressure and the Capacity Gap

These challenges are unfolding against a backdrop of constrained operating environments. Many retailers are managing tight margins, short-staffed stores, high turnover, and increased pressure on earnings. LP teams have to do more with less. While LP funding has generally remained stable, headcount is not keeping up with expanding threat surfaces, as it can be difficult to find talent with LP expertise and data fluency.

Technology is now the primary lever for scaling LP effectiveness, but it introduces a new kind of complexity: with dozens of vendors and hundreds of solutions available, leaders struggle to determine which capabilities matter most, how they fit together, and in what order they should be adopted. Without strategic guidance, technology portfolios risk becoming disconnected collections of tools rather than part of a coherent

LP strategy. With pressure on earnings and retail crime making mainstream news headlines, executives are understanding the value of LP/AP more than ever, and in some organizations, LP/AP is becoming an important data node for the organization as a whole. This presents an opportunity for LP leaders who want to provide greater value to their organizations.

Why Current Approaches Are Falling Short

Today's LP operations are still largely reactive and event-driven. Devices such as cameras, POS terminals, and EAS systems record what happened, but rarely explain why it happened or how to prevent recurrence. Critical insights remain buried in siloed systems: POS data, HR records, video footage, inventory files, and audit results often live on separate platforms that rarely communicate. Analysts must manually piece these fragments together, a process that is often slow and inconsistent.

Even more advanced rule-based alerting systems generate overwhelming noise. Alert fatigue is real and costly, and false positives consume valuable time while true risks blend into the background. As criminals evolve their tactics and omnichannel operations grow more complex, static rules and disconnected systems struggle to keep pace.

Compounding the issue, LP lacks a shared language for modernization. "AI" is used inconsistently across the industry, often referring to basic automation rather than true intelligence. Most organizations know they need to evolve, but few have a roadmap for doing so. Leaders cannot benchmark their current capabilities, prioritize investments, or articulate a clear transformation strategy to cross-functional partners or executive sponsors.

The proliferation of "AI" as a buzzword creates its own challenges. The term is used loosely, often describing little more than rebranded automation or basic machine learning. Distinguishing genuine capability from vendor marketing requires technical sophistication that many LP leaders understandably lack. This uneven, ambiguous landscape makes it difficult for LP leaders to evaluate capabilities, prioritize investments, or chart a coherent modernization path.

Adoption Without Alignment: LP's Modernization Dilemma

Technology adoption in loss prevention is accelerating. Nearly 90% of retailers reported using or evaluating AI solutions in 2024, and adoption of technologies like exception-based reporting, computer vision, and RFID continues to grow (at varying levels of adoption).

> **90% of retailers reported using or evaluating AI solutions in 2024.**

Yet implementation remains fragmented and uneven. One study found that 84% of retail associates are concerned about loss prevention and the lack of technology to spot safety threats or criminal activity.

This lagging adoption creates a widening performance gap between early movers and organizations still operating in reactive modes. It also amplifies confusion: with inconsistent definitions and varied claims from vendors, LP leaders need a structured way to evaluate maturity, avoid hype, and ensure investments align with real business outcomes.

The Case for a Strategic Framework

Other enterprise functions have navigated technological transitions by using maturity models for decades. Marketing, finance, supply chain, and HR all rely on established frameworks to benchmark capability, guide investment, build cross-functional alignment, and measure progress over time. Loss prevention needs an equivalent structure.

A strategic framework provides:

- A common language for describing current LP capability
- A roadmap for progressing from reactive practices to intelligence-driven operations
- Clear investment priorities aligned to business value
- Criteria for evaluating vendor claims and distinguishing true capability from rebranded automation
- A basis for executive sponsorship by tying LP modernization to enterprise outcomes

Without such a model, LP transformation risks becoming a series of disconnected point-solution deployments that improve symptoms in isolated areas but never compound into a coherent, scalable strategy. Having a model can enable honest benchmarking of current state, identify specific requirements for advancement, align technology investments with business value creation, and help build compelling cases for executive sponsorship. LP leaders can better evaluate their own organization, evaluate vendor capabilities, and plan for the future.

Introducing the Loss Prevention Maturity Model

This whitepaper introduces the Loss Prevention Maturity Model, a practical framework grounded in real-world operations. It outlines four stages of LP evolution: Devices, Analytics, AI, and Agentic AI. Each stage has clear characteristics, specific technological requirements, and defined advancement criteria.

The model clarifies how foundational LP technologies build on one another, progressing from physical devices to the use of data generated by those devices. We then explore the potential of AI and how LP technology is in a process of shifting from reactive surveillance to proactive intelligence and ultimately toward autonomous action.

It offers benchmarks for self-assessment, the implementation challenges and prerequisites of each stage, and the steps LP organizations can take to progress through the stages. Most importantly, it provides a coherent roadmap for navigating rapid change, empowering organizations to move from doing more with less to achieving more with technology.

The Loss Prevention Maturity Model

A strategic framework for LP capability evolution

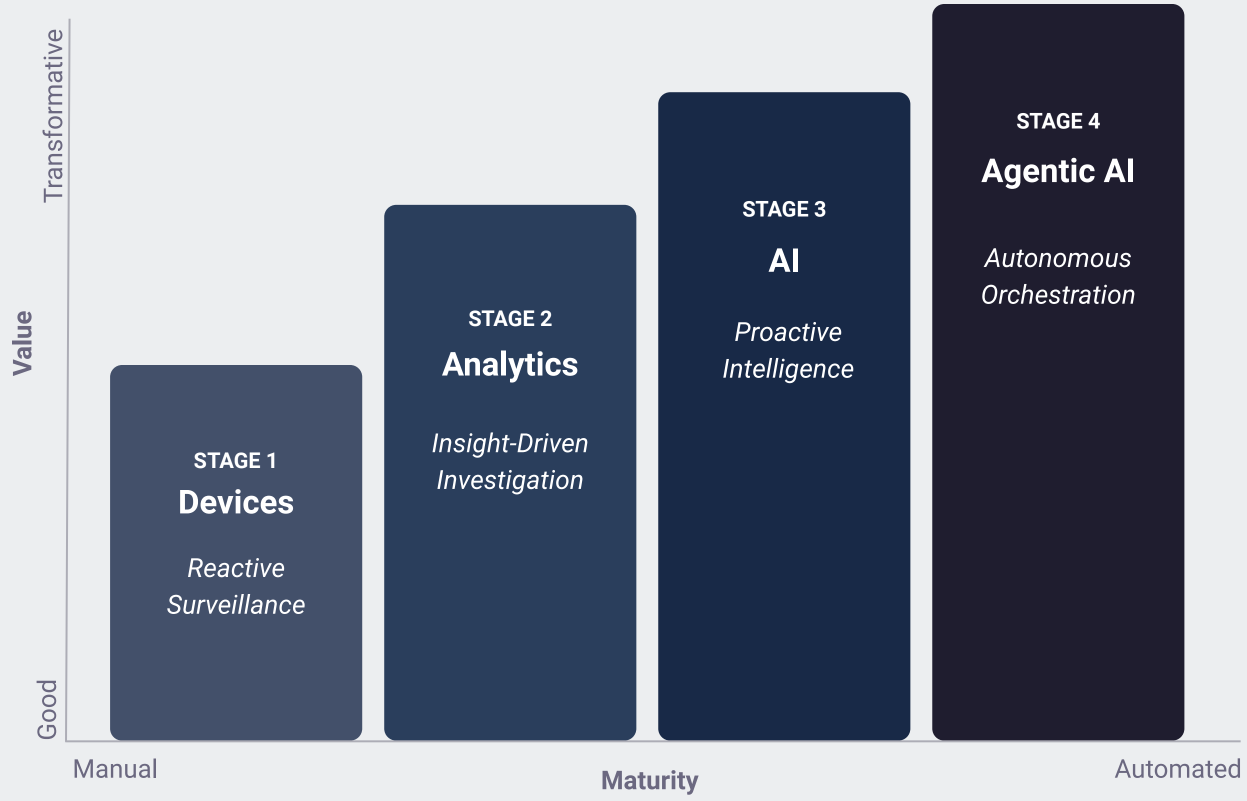

Source: Agilence, 2026

Stage 1 — Devices

Every loss prevention program begins with physical devices. The LP device toolkit includes cameras and sensors that observe and record events, barriers and locks that restrict access to merchandise, tags and alarms that signal theft attempts, point-of-sale (POS) terminals that log transactions, and RFID systems that track inventory movement. These technologies often serve dual purposes: they deter opportunistic theft through visible presence while simultaneously providing visibility into incidents and evidence for after-the-fact investigations.

Yet the core function of these tools is narrow. Devices often tell us what happened but not why it happened. A camera records a shoplifting event. A POS terminal logs a transaction. An RFID reader tracks tagged merchandise. A locked case prevents a theft attempt. Each creates a fragment of the picture, but none provide context or causality. They are tools of monitoring, detection, and prevention, but not explanation or strategy.

The Device Ecosystem

We divide LP devices into two overlapping categories.

Deterrence-focused

Deterrence-focused tools, such as mirrors, visible camera domes, locking display cases, EAS pedestals, spider wraps, keeper boxes, ink tags, and access control systems exist primarily to discourage theft through visible presence or physical barriers. Security personnel, while not devices themselves, function as part of this physical LP toolkit, providing human presence and response capability. Each of these measures adds friction to theft, though at the cost of also adding friction to legitimate shopping and customer service.

Data-generating devices

Data-generating devices, such as video systems, POS terminals, RFID sensors, and alarms capture records that can be reviewed later or, in the more mature LP program stages, fed into analytical systems. Cameras evolved from 1960s closed-circuit television (providing deterrent value even when footage went unreviewed) to 1980s-90s recording systems that made them essential for evidence and prosecution. Electronic Article Surveillance (EAS) systems, commercialized in the 1970s, exemplify dual functionality: tags and pedestals create audible alarms at the moment of theft (deterrence and intervention) while also generating incident data. POS systems matured in the 1980s-90s, creating transaction-level data that would later prove critical for identifying internal theft and fraud. More recently, RFID has emerged as a successor to EAS, offering continuous inventory tracking rather than just exit-point detection. In practice, most devices serve both functions.

A visible camera deters theft and produces video that may become data for analytics. On-site personnel provide presence and file incident reports. Locked cases prevent theft and can generate data on associate unlock frequency and customer wait times. EAS tags prevent theft through deterrence and create records of alarm events.

That distinction matters beyond taxonomy. Deterrence-focused tools reduce theft through presence and friction. Data-generating devices do something else entirely: they create the records that make everything in the following stages possible.

STAGE 1 - DEVICES

The Device Ecosystem

LP devices serve two overlapping functions

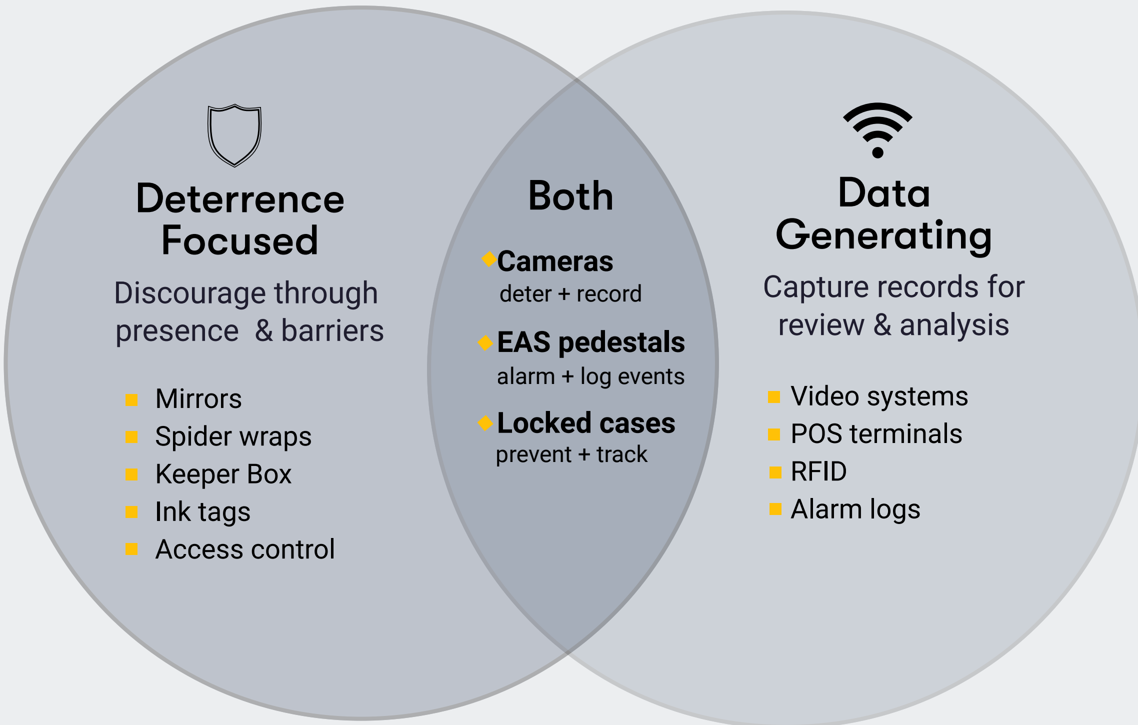

Source: Agilence, 2026

The Operational Reality and Its Limitations

In the maturity model, devices sit at the bottom left: high effort, modest value. Their effectiveness depends almost entirely on human labor. LP staff must watch video feeds, review DVR footage, respond to alarms, and manually correlate information across fragmented systems.

The operational reality of device-centric LP is reactive and event-driven. An incident occurs—merchandise goes missing, a complaint is filed, suspicious activity is noticed—and someone investigates manually. This approach is labor-intensive and doesn't scale. As business grows, the volume of device data expands faster than the human capacity to review it, creating inevitable coverage gaps. Devices are also typically siloed: video systems, POS data, EAS alerts, and access logs rarely communicate with one another in Stage 1 environments, producing fragmentation: plenty of signals, little synthesis.

Organizations operating primarily at Stage 1 face structural constraints:

- Reactive, not proactive: Incidents are reviewed after the fact.
- Labor intensive: Hours spent on video review or manual reconciliation.
- Delayed response: Manual investigations take time, allowing trails to go cold before detection.
- Pattern blindness: Siloed data and lack of cross-location oversight prevents recognizing trends across stores or time periods.
- Resource misallocation: Without broad visibility, resources go to what's most visible rather than what's most impactful.
- Scalability challenges: Human bandwidth cannot keep pace with rising transaction volume or increasingly sophisticated fraud tactics.

Organizations recognize they need more than devices alone when they repeatedly investigate the same issues without understanding root causes, when alert volumes overwhelm response capacity, or when leadership asks questions that devices cannot answer: Which patterns correlate with loss? Where should we focus limited resources for maximum impact?

Stage 1 as Foundation

Despite these limitations, Stage 1 never becomes obsolete. Devices remain essential regardless of analytical sophistication. Organizations don't transcend this stage; they build upon it. LP teams in later stages still rely on cameras, EAS, locked cases, and security personnel. The difference lies in how deployment decisions are made and how device outputs are consumed.

The data generated by physical devices is a prerequisite for moving into the following stages. Without cameras, sensors, and transaction systems creating records, there is nothing to analyze. Stage 1 establishes the data foundation that makes analytics possible.

Later stages use analytics and AI to determine how to arrange the devices of Stage 1: where cameras should be positioned, which products merit physical security measures, and when to increase personnel presence. Intelligence flows back to inform device strategy, creating continuous improvement rather than one-way progression. Learnings from more advanced stages funnel back to create more efficient strategies for using physical devices as both deterrents and data sources.

Stage 2 — Analytics

If devices tell us what happened, analytics tells us why it happened, and where it's likely to happen next. Stage 2 represents the shift from reactive investigation to proactive pattern recognition. Rather than waiting for incidents to surface and reviewing isolated device outputs, LP teams now deploy loss prevention analytics platforms that integrate transactional data with contextual information such as point-of-sale records, employee data, loyalty program data, alarm data, video, and much more in order to understand root causes and identify systemic vulnerabilities.

This is where most of the loss prevention market operates today. Organizations at Stage 2 have embraced data analytics as the core methodology of the discipline. Decisions require evidence. Investigations begin with dashboards. Priorities emerge from reports rather than intuition or anecdote.

The Integration Breakthrough

The defining characteristic of Stage 2 is data convergence, made possible by specialized LP analytics platforms that connect disparate systems into unified views. A suspicious transaction isn't just a line in a POS log, it can be cross-referenced with employee schedules, inventory data, and historical fraud indicators. A theft event captured on video gains context when the platform correlates it with inventory discrepancies, prior incidents at the same location, and regional trends. Data streams that once existed in silos now flow into exception-based reporting tools where analysts can correlate, compare, and investigate across dimensions.

This integration reveals patterns invisible at Stage 1. An isolated refund looks benign; fifty refunds by the same cashier during unstaffed hours becomes a clear signal.

A single case of merchandise walking out the door might be opportunistic; a cluster of similar events across multiple stores on promotion weekends indicates organized activity. Analytics doesn't just answer questions—it surfaces the right questions to ask.

STAGE 2 - ANALYTICS

The Integration Breakthrough

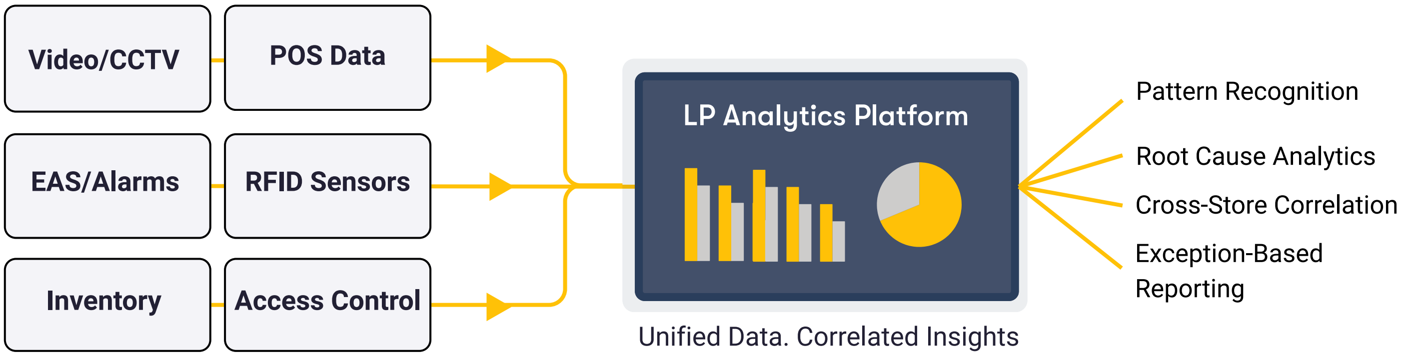

Source: Agilence, 2026

From Manual Reviews to Exception-Based Reporting

The analytical approach to loss prevention emerged in the mid-1990s with Exception-Based Reporting software, a breakthrough that allowed LP teams to mine POS transactional data across entire store chains from one central point rather than investigating locations individually. These EBR tools revolutionized the discipline by enabling teams to search for patterns of internal fraud—refunds, voids, employee discounts, sweethearting—using analytics software rather than manual processes. At the time, database storage was expensive and network speeds were slow, but the productivity gains justified the investment.

As networks accelerated and storage costs fell, EBR platforms expanded their capabilities. POS systems captured richer transaction detail. Video systems integrated directly into analytics suites, allowing analysts to push a button and review footage corresponding to flagged transactions. Cloud computing reduced costs and eliminated IT burden. According to one industry survey, 95% of LP professionals reported that exception-based reporting tools were what they used to accurately identify cases of internal loss. The transformation period roughly spanned 2012 to 2017, as earlier generations of LP professionals—many with law enforcement backgrounds—moved from skepticism about statistical patterns to full adoption of data-driven methods enabled by these platforms.

The Analyst as Detective

Stage 2 elevates the role of the LP analyst from evidence reviewer to pattern detective. Working within data analytics suites, analysts search for anomalies: unusual transaction volumes, statistically improbable refund rates, inventory variances that don't align with sales, time-and-attendance patterns that correlate with shrink. The work remains intensive. Analysts sift through data, test hypotheses, tune alert parameters, and distinguish signal from noise. The analytics platform provides the view; the analyst provides the judgment.

Organizations build libraries of reports and dashboards within their LP analytics tools tailored to known risk scenarios, gradually automating the discovery of familiar patterns while still requiring human investigation for novel or complex cases. Rule-based alerts built into these platforms handle routine monitoring—flagging when conditions exceed thresholds—but effectiveness depends on how well humans configure those rules and evaluate their outputs. This represents a fundamental shift from Stage 1, where LP teams reacted to incidents after they occurred. At Stage 2, analysts actively hunt for patterns before losses compound, using exception-based reporting tools to prioritize investigations and allocate resources to the highest-impact areas.

The Analytics Hierarchy

Not all analytics are created equal. Stage 2 loss prevention platforms typically offer capabilities across a spectrum of analytical sophistication:

Descriptive

Descriptive Analytics form the foundation. These basic calculations show what happened: sales increased, refunds spiked, inventory variance occurred. Dashboards and reports within the analytics suite deliver visibility into historical performance but offer limited insight into causality.

Diagnostic

Diagnostic Analytics go deeper, examining why it happened. By combining transactional data with contextual variables—employee schedules, weather patterns, promotional calendars, product attributes—these analyses identify root causes. A spike in voids might correlate with undertrained staff, system errors during a promotion, or deliberate theft. This is where exception-based reporting tools demonstrate their value, automatically surfacing correlations that would be invisible in manual reviews.

Predictive

Predictive Analytics shift focus to what could happen. Using historical patterns and statistical models, advanced LP analytics platforms allow teams to test scenarios: What happens if we change staffing patterns at high-risk locations? How will a new promotion structure affect exception volumes? Predictive approaches remain uncommon in Stage 2 LP operations, though adoption is growing.

Prescriptive

Prescriptive Analytics represent the frontier of Stage 2, recommending what should be done and, in advanced implementations, executing changes automatically. Few loss prevention analytics platforms operate at this level today—it marks the threshold between Stage 2 and Stage 3, where machine learning begins to replace rule-based logic.

Most Stage 2 organizations concentrate on descriptive and diagnostic analytics. This is where the analyst's detective work operates: investigating the patterns surfaced by reports, tuning alert thresholds within the platform to reduce false positives and translating statistical anomalies into actionable cases. The intelligence resides in human judgment applied to software-generated insights.

The Analytics Hierarchy

Not all analytics are created equal

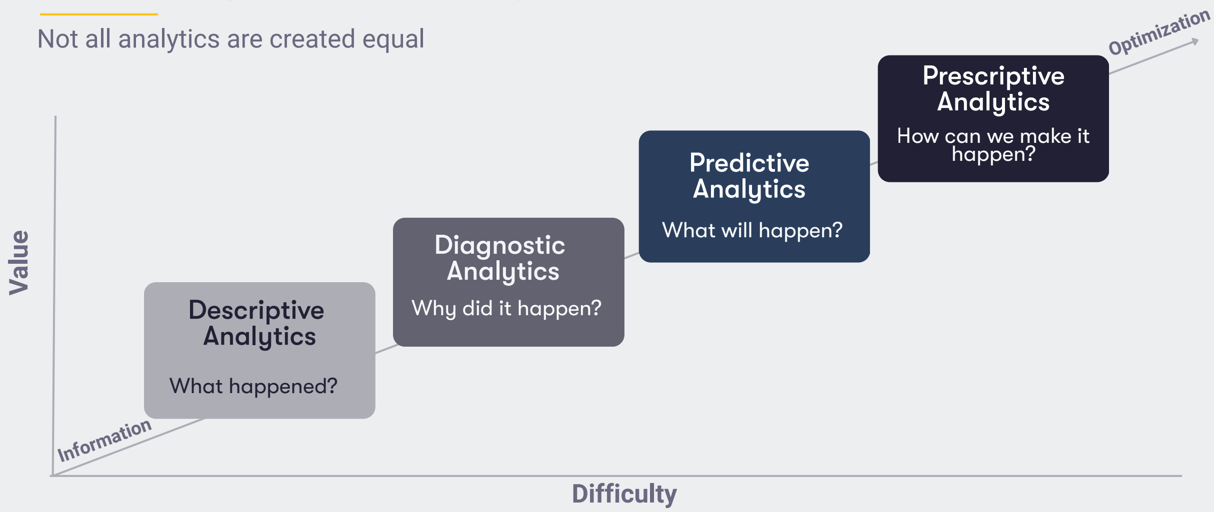

Source: Agilence, 2026, based on Gartner, 2012

Limitations and Readiness Signals

Despite measurable gains, Stage 2 organizations face structural constraints. Analysts remain the bottleneck—their capacity to investigate limits how much insight the organization can extract from its data. Alert volumes can overwhelm response capacity, creating false positive fatigue. Effective monitoring requires continuous tuning—rules that catch this quarter's schemes may miss next quarter's variations. The time spent configuring, refining, and managing alert systems leaves less time for investigation and action.

Organizations recognize readiness for Stage 3 when analysts spend more time managing alerts than investigating cases, when leadership asks for predictive insights that static reports can't provide, or when data volume exceeds human capacity to process it effectively.

These signals indicate that rule-based analytics, while valuable, no longer match the complexity of the problem—and that machine learning may offer the next step forward.

Stage 2 never becomes obsolete, just as Stage 1 devices remain essential. Organizations continue to rely on descriptive and diagnostic analytics even as they adopt more advanced capabilities. But as teams master the analytical approach and face its limits, they increasingly ask: Can the software figure it out and just tell us where to look? That question points toward Stage 3.

Stage 3 — AI: From Analysis to Action

While analytics platforms reveal patterns, LP analysts must still hunt for them. LP/AP teams spend hours building queries, reviewing dashboards, and investigating hunches that may lead nowhere. Stage 3 is about changing this fundamental dynamic: Artificial Intelligence (AI) surfaces what matters, and LP teams focus on taking action. The LP Analyst workflow transforms from "searching for problems" to "responding to prioritized, confidence-scored recommendations."

This shift allows a new level of scalability for LP/AP teams. Machine learning models now approximate the pattern recognition of expert investigators, but at a scale and speed no human team could match. While most organizations still operate at Stage 2, leading LP teams are embracing new AI tools to multiply their effectiveness. Stage 3 is the current state of the art in Loss Prevention, but adoption remains uneven across the industry, and most AI features are new and maturation and enhancements until they become reliable bedrocks of LP workflows.

The AI Transformation

Artificial Intelligence tools often act like an assistant for LP/AP professionals. The transition from Stage 2 to Stage 3 mirrors a detective getting an intelligent assistant who pre-screens all leads, prioritizes the most promising cases, and gathers preliminary evidence—all before the detective arrives at work. AI doesn't replace human analysis; it accelerates and amplifies it.

AI changes LP workflows. At Stage 2, an analyst might spend their morning querying transaction databases, building custom reports, and manually scanning for anomalies. They rely on predetermined rules and thresholds that generate alerts (many of them false positives). By lunch, they might identify one or two issues worth investigating.

In a mature Stage 3 environment, that same analyst arrives to find AI has already analyzed thousands or even millions of transactions, considered hundreds of variables simultaneously, and surfaced a prioritized list of likely incidents, each with a confidence score and impact assessment. The AI has pre-gathered relevant evidence: correlated transactions, video segments, employee schedules, historical patterns. The analyst's expertise shifts from finding problems to validating AI recommendations and determining appropriate responses. Investigation time drops dramatically, allowing teams to address more cases without additional headcount.

The 4 Main Categories of AI Applications in LP

It's important to understand that "AI" encompasses multiple distinct technologies at different maturity levels, each delivering different business outcomes. Computer vision algorithms that detect theft are fundamentally different from natural language processing that links cases in both their use case and underlying technology, though both fall under the AI umbrella. To help clarify, we've broken AI into four main categories based on their main application areas.

Seeing

Seeing (Computer Vision) represents the most mature AI application in LP today. Computer vision is a field of artificial intelligence that enables computers to interpret and understand the visual world from digital images or videos. Self-checkout monitoring systems detect scan avoidance and ticket switching in real-time. Facial recognition platforms alert staff when known offenders enter stores. License plate recognition (LPR) systems track vehicles associated with organized retail crime. These technologies are delivering measurable results today and adoption is growing.

Thinking

Thinking (Risk Detection and Prioritization) encompasses machine learning models that analyze datasets to detect anomalies, identify patterns, and help prioritize what matters most. Risk-scoring algorithms evaluate every transaction against dozens of variables—time of day, SKU mix, operator patterns, discount usage, tender types, and historical behavior—to flag events with high fraud likelihood. Anomaly detection models uncover emerging patterns that rules would miss, such as unusual refund sequences or outlier inventory movements.

Alert scoring ranks the flood of exceptions to highlight the incidents most likely to represent actual loss, reducing noise and improving analyst efficiency. Alert scoring systems, like those from Agilence and others, prioritize which alerts deserve immediate attention versus those likely to be false positives. Cross-location correlation can reveal regional patterns and ORC activity that no individual store could see. These systems accelerate investigation speed, save time, and improve accuracy for analysts.

Understanding

Understanding (Natural Language Processing) is still in early adoption, but the rise of Large Language Models (LLMs), such as those underlying applications like ChatGPT, has accelerated progress dramatically. NLP features or tools can analyze unstructured written text (such as incident reports, case notes, emails, audit findings) to surface new insights. For example, case linking algorithms compare the narrative descriptions in case reports to identify related incidents and ORC patterns. Natural language interfaces can simplify the use of analytics tools by allowing analysts to query data conversationally ("show me all cases involving razors this month") rather than building filters. LLM-powered models may be able to summarize lengthy incident reports, extract key details automatically (suspect descriptors, loss amounts, etc.), classify cases, and highlight emerging trends across thousands of reports. While adoption is early and less mature than Computer Vision, NLP features may become important analysis tools and interface options.

Advising

Advising (Prescriptive AI) is the most aspirational branch of AI in LP. While vendors frequently use terms like "predictive" or "preventive," true prescriptive systems that recommend specific operational actions before incidents occur are not yet deployed. Current tools can forecast trends and highlight high-risk stores or categories, but they cannot reliably predict when or where theft will occur or advise stores on preventive actions. Still, the potential is significant. As more data sources are integrated and machine learning models mature, prescriptive AI features could eventually guide LP teams with actionable recommendations—such as adjusting SCO staffing, relocating high-risk merchandise, or modifying return workflows—laying the foundation for Stage 4's autonomous agents. But today, prescriptive AI remains an ambition rather than a reality, and honest practitioners treat predictive claims with caution.

STAGE 3 - AI

Four Categories of AI in LP

AI is not monolithic - each uses different technology at different maturity

Seeing
MOST MATURE

Computer Vision

- SCO scan avoidance
- Facial recognition alerts
- License plate recognition
- Behavior analysis

Thinking
GROWING ADOPTION

Risk Detection & Prioritization

- Transaction risk scoring
- Anomaly detection
- Alert prioritization
- Cross-location correlation

Understanding
EARLY ADOPTION

Natural Language Processing

- Case linking from text
- Conversational queries
- Report summarization
- Trend extraction

Advising
ASPIRATIONAL

Prescriptive AI

- Staffing adjustments
- Merchandise placement
- Workflow optimization
- Proactive risk mitigation

Source: Agilence, 2026

Operational Reality and Model Evolution

Stage 3 changes the day-to-day experience of LP analysts not by eliminating work, but by reshaping it. AI-driven alerts, risk scores, and anomaly detections surface far more potential issues than manual methods—but they are not perfect. As of 2025, AI features are still new and developing. Early deployments can often generate noise: false positives, missed incidents, or alerts that don't yet reflect the nuances of a retailer's processes.

Over time, these systems can be refined, particularly the machine-learning models that score transactions and prioritize alerts. As analysts mark alerts as useful or irrelevant, the system learns which patterns matter for their organization. What appears suspicious in one context may be normal in another; feedback helps the models become better calibrated to local realities. Computer vision tools tend to improve more slowly and globally, with enhancements delivered by the vendor rather than learned automatically from store feedback.

The practical impact is that Stage 3 analysts spend less time digging for issues and more time evaluating and acting on them. Workflows shift from building queries or manually reviewing video to triaging prioritized alerts, validating findings, escalating cases, and closing the loop with store teams. Organizations that invest in training, feedback, and iterative refinement see AI become more reliable over time, supporting analysts with higher-quality insights rather than replacing human judgment.

The Human-AI Partnership

Discussions of AI bring with them concerns about AI replacing human workers. Will AI take the jobs of LP teams? The truth is that for now and the foreseeable future, AI augments human analyst capability rather than replaces it. Machine learning models excel at pattern recognition across large data sets, but humans provide context, judgment, and strategic thinking that no algorithm can match.

When AI flags an unusual return pattern, experienced LP professionals determine whether it's fraud or a legitimate edge case. When computer vision identifies suspicious behavior, humans decide if intervention is appropriate or if there's reasonable explanation. AI brings immense processing power, but as any LP practitioner knows, there are too many aspects of the job technology cannot do.

Paradoxically, AI makes experienced LP professionals more valuable, not less. Their expertise trains the models, their judgment validates recommendations, and their strategic thinking directs how AI capabilities get deployed. Rather than eliminating jobs, Stage 3 AI elevates the LP role from data processor to strategic decision maker.

Stage 4 — Agentic AI

The final stage of our model, Stage 4, is about the foreseeable future of loss prevention: Agentic AI. No organizations currently operate in this stage. Agentic AI is the currently hypothetical frontier of LP technology, where AI-powered systems don't merely detect anomalies or surface insights, but take autonomous action to address them. In this stage, AI tools move from "tell what to do" to "here's what I've done, please approve." While this is currently only aspirational, the foundational technologies and organizational capabilities required for this transformation are beginning to emerge across the industry.

The progression follows a natural evolution: Devices capture what happened, Analytics reveal why it happened, AI determines what to focus on, and Agentic AI completes the loop by executing what needs to be done. This represents the ultimate maturity in loss prevention: autonomous systems that don't just inform human decision-making but actively participate in operational response.

Understanding Agentic AI in Loss Prevention

Agentic AI refers to autonomous systems capable of planning, reasoning, and taking action with minimal human oversight. These systems interpret data from multiple feeds, make decisions, pursue goals, and orchestrate multi-step workflows within defined parameters. While traditional AI tools respond to queries or surface anomalies, agentic AI carries out multi-step workflows. For example, rather than simply alerting an analyst on an ORC pattern, an agentic system could compile evidence across stores, check it against external data sources, generate incident reports, and initiate a response protocol — delivering a completed package for review rather than waiting for human direction.

However, "agentic AI" has also become a marketing buzzword applied to everything from chatbots to workflow automation. In practice, LP remains years away from tools with true autonomy, and even bounded agents (systems that automate multi-step processes within strict guardrails) are not yet meaningfully deployed. Current LP offerings are sophisticated automation, not autonomous agents. When vendors advertise "AI agents," they generally refer to predefined workflows triggered by conditions, not systems capable of independent reasoning or goal pursuit. Understanding this gap is essential for separating today's reality from tomorrow's potential.

The Spectrum of Agency

What does "agency" for a software system really mean? In the broadest sense, for a person or a piece of software, agency is the ability to operate independently. Agency exists on a continuum, not a binary. Understanding this spectrum helps LP teams realistically evaluate progress and opportunities.

No Agency
Traditional rule-based tools such as EAS pedestals, threshold alerts, and scheduled reports. They execute fixed "if-then" logic with no adaptation.

Low Agency
AI-assisted tools that improve human capability but do not initiate action. Examples include current AI capabilities such as natural-language interfaces, anomaly detection alerts, and recommendation engines. These tools provide insight, but humans drive all action.

Bounded Agency
The emerging frontier for AI agents. Bounded agents can run multi-step workflows within strict constraints—for example, detecting a refund anomaly, creating a case, attaching relevant data, and routing it for review. Human oversight remains required, and boundaries are narrow. Most agentic systems in LP will likely operate at this level for the foreseeable future due to governance, legal, and safety requirements.

High Agency
Systems capable of launching complex investigations, adjusting protocols dynamically, and coordinating multi-system responses. These agents would proactively identify and mitigate risks, escalating only ambiguous or sensitive situations. Such capabilities remain technically out of reach today.

Full Autonomy
A theoretical stage where agents operate independently with minimal human intervention and make strategic decisions. This level raises fundamental questions about accountability and governance and may never be adopted in LP or most enterprise environments.

Across industries, most organizations using AI today operate somewhere between low agency and early bounded agency. A small number of more advanced sectors (such as cybersecurity, financial operations, customer support, and IT service management) are beginning to implement workflow automation that resembles multi-step, goal-oriented behavior, but always within strict parameters and with human oversight. These early patterns suggest how bounded agents may evolve in LP once data integration, governance, and operational trust mature.

STAGE 4 - AGENTIC AI

The Spectrum of Agency

Agency exists on a continuum, not a binary

No Agency	Low Agency	Bounded Agency	High Agency	Full Autonomy
Fixed if-then logic with no adaptation	**AI-assisted tools that improve human capability**	**Multi-step workflows within strict constraints**	**Complex investigations and dynamic protocols**	**Independent strategic decisions, minimal oversight**
• EAS pedestals • Threshold alerts • Scheduled reports	• NL interfaces • Anomaly detection • Recommendation engine	• Auto case creation • Evidence attachment • Workflow routing	• Multi-system coordination • Proactive risk mitigation • Adaptive escalation	• Strategic LP Planning • Cross-domain decisions • Self-governing systems

Source: Agilence, 2026

Operational Vision

What might LP operations look like in a Stage 4 organization? Most likely, a growing share of today's manual workflows will become automated, more processes will be software-initiated rather than human-initiated, and AI systems with greater agency and deeper integration across various enterprise tools will proactively execute cross-departmental workflows.

Early Stage 4 environments will likely begin with bounded agents that automate predictable, low-risk tasks. For example, an agent might generate a case from an anomaly detection event, attach relevant transaction data and video clips, and assign it to the right investigator based on workload. These systems operate within strict parameters and include human validation at key decision points.

Over time, organizations will gradually expand agent boundaries as reliability improves and trust grows. Bounded agents may standardize incident documentation, automate compliance checks, orchestrate multi-system data gathering, or prepare investigation packets. As systems demonstrate consistency, some actions move from "AI prepares, humans approve" to "AI acts, humans audit."

In more advanced Stage 4 environments, agents autonomously coordinate workflows across systems. Instead of analysts manually searching for patterns and creating cases, agents continuously monitor data streams and trigger appropriate responses. A fraud pattern detected across multiple locations could automatically create investigation cases, notify store management, adjust staffing schedules to increase oversight, and generate compliance or risk reports — all with complete audit trails for human review.

Organizational Shifts Beyond the Technical

At maturity, autonomous cross-system orchestration becomes the defining feature of Stage 4. LP systems may trigger inventory reorders, notify HR of personnel-linked anomalies, or alert operations teams to process breakdowns. LP becomes embedded in the nervous system of enterprise decision-making rather than siloed.

As this occurs, the human role shifts dramatically. LP professionals spend less time executing tactical tasks and more time providing strategic oversight, handling ambiguous judgment calls, validating agent decisions, and managing edge cases that exceed the system's contextual understanding. Investigators and analysts evolve into orchestrators and strategists who guide, govern, and audit digital agents.

Reaching Stage 4 is as much an organizational transformation as a technological one. Moving from Stage 3's human-directed AI to Stage 4's autonomous agents require not only system integration and API connectivity, but also redesigned workflows, new governance structures, clear accountability models, and a redefinition of team roles and skills.

Current Reality and Evolution Path

Expert timelines for agentic AI development vary dramatically, from within a year to a decade or more. This uncertainty reflects both the inherent unpredictability of technological progress and the tendency toward over-optimism that has characterized AI predictions since the 1960s. Understanding where loss prevention sits within this landscape helps separate realistic planning from marketing hype.

Most expert predictions focus on the arrival of artificial general intelligence (AGI), a system with human-level cognitive capability. However, AGI measures the intelligence of a system, not its autonomy. A system can be extremely intelligent yet have very low agency (e.g., today's LLMs). Agentic AI measures the ability to take action toward goals, not simply reason about them.

AI research labs project aggressive timelines. Anthropic CEO Dario Amodei predicted "powerful AI systems" matching Nobel Prize-level intellect by early 2027. OpenAI's Sam Altman suggested AI agents would meaningfully enter workplaces in 2025. Google DeepMind's Demis Hassabis forecasts AGI within 5-10 years. These predictions have proven overly optimistic, yet debate continues; some argue current AI capabilities already constitute AGI, while others believe we remain decades away.

Industry analysts offer more optimistic but tempered expectations. Firms such as Gartner, Forrester, and PwC see 2026 as a pivotal year for AI agent adoption but also warn that many projects will fail due to unclear ROI, integration challenges, and governance risks. Gartner projects that by 2028 roughly 15% of day-to-day work decisions could be made autonomously by agentic AI—but also expects over 40% of agentic AI initiatives to be canceled by 2027.

Recently, some technical leaders have offered more sobering forecasts. Andrej Karpathy, OpenAI founding member and former Tesla AI director, recently said this would be "the decade of agents, not the year of agents," given current systems' limitations in reasoning ability, reliability, multimodal capability, interface use, and lack of continual learning. Realistically, we are still at the dawn of AI agent adoption, and reliable systems may take a decade to mature.

What This Means for LP

For LP leaders, the implications are clear: autonomous and semi-autonomous systems are coming, but they will take time, and LP will naturally adopt them later than faster-moving sectors. Retail operates on long technology refresh cycles, LP depends on upstream systems it does not control, and errors carry uniquely high costs. As a result, Stage 4 adoption will appear gradually—first through bounded agents in tightly scoped workflows—long before broader autonomy becomes practical. This evolution will unfold over years, not months.

> **By 2028, 15% of day-to-day work decisions could be made autonomously by agentic AI**

How to Advance Through the LP Maturity Stages

The preceding sections established a framework for understanding loss prevention maturity. They defined the four stages of our model (Devices, Analytics, AI, and Agentic AI) and characterized the capabilities, operational patterns, and business value associated with each. But knowing what defines each stage is only half the equation. LP leaders also need a practical roadmap for how to move forward. This section provides that roadmap.

Advancing through the maturity model is not simply a matter of purchasing new products. It requires coordinated progress across four foundational dimensions: Data, People, Processes, and Technology. We walk through each stage transition, identifying prerequisites, key actions, and common implementation challenges. The goal is to enable LP leaders to assess their current position, understand what needs to be true to move forward, and build realistic transformation plans that align with business priorities and organizational capacity.

The Four Foundations of Advancement: Data, People, Processes, Technology

Every organization progressing through the maturity model encounters the same bottlenecks and the same enablers. These four foundations provide a simple, universal framework for understanding what must evolve to support each stage transition.

Data

Data is the foundation that enables every subsequent stage of capability. Stage 1 generates abundant device data (video, POS transactions, RFID reads, alarm events) but little of it is connected or structured in ways that support analysis. Stage 2 begins the process of integrating these streams. Stage 3 depends on consistent and well-tagged data to train AI systems. Stage 4 requires data governance strong enough to support automated actions. Three aspects matter most:

Readiness Data must be sufficiently clean, consistent, and granular for its intended use. Even AI systems that can begin learning on day one require stable identifiers (store, associate, SKU, transaction) and basic reliability in timestamps, event logs, and case outcomes.

Integration Mature LP organizations link data from across systems into one or more platforms. EBR tools demonstrate the power of viewing data in context; modern AI (particularly "Thinking" category risk scoring and anomaly detection) depends on it.

Governance and permissible use As organizations adopt computer vision ("Seeing"), NLP ("Understanding"), and eventually prescriptive tools ("Advising"), questions of privacy, compliance, and HR acceptability become central. Not all data can be used for all purposes. LP, IT, Legal, and HR must align on what is permissible, what is restricted, and what requires explicit documentation or oversight.

People

Technology cannot compensate for a workforce that is unprepared or unwilling to adopt new tools. Advancing through the maturity stages requires shifts in skills, culture, and roles.

At Stage 1, LP staff perform manual reviews and rely on intuition and field familiarity. Stage 2 demands analysts who are comfortable using dashboards, correlating data, and interpreting statistical patterns. Stage 3 requires teams who can work with probabilistic outputs such as risk scores, confidence ratings, and anomaly detections, rather than deterministic rules. Stage 4 shifts LP roles again, toward supervising automated workflows, governing agent behavior, reviewing edge cases, and acting as stewards of accuracy and fairness.

Change management is as important as technical skill. Organizations must prepare teams for the psychological shift from "the system helps me find problems" to "the system decides which problems I should focus on." Trust, transparency, and role clarity are essential. AI does not eliminate LP work; it elevates it. Analysts must see themselves not as displaced investigators but as strategic decision-makers guiding increasingly capable systems.

Processes

Processes turn insights into outcomes. Without well-defined workflows, even the most advanced systems produce little value.

Stage 1 organizations often lack standardized investigative workflows; incidents are reviewed ad hoc, outcomes are inconsistently documented, and recurrence analysis is limited. Stage 2 brings structure: reports lead to cases, cases follow escalation paths, and dispositions generate learnings that inform future action. Stage 3 requires disciplined feedback loops so AI systems can learn which alerts are useful and which are noise. Stage 4 demands documented, auditable workflows that agents can partially or fully automate within defined guardrails.

This includes governance processes: approval matrices, risk thresholds, exception handling, and procedures for halting or overriding automation. Agentic AI, particularly bounded agents, depends on processes that are explicit, structured, and consistently applied.

Technology

Technology is the visible layer of maturity, but it is the last to fully develop because it depends on the other three foundations.

Stage 1 provides the device layer: cameras, POS, RFID, EAS, alarms, and physical deterrents. Stage 2 centralizes data through analytics and EBR platforms, finally giving LP a cockpit for investigation. Stage 3 introduces AI capabilities that align with the four categories described earlier: computer vision ("Seeing"), risk scoring and anomaly detection ("Thinking"), natural language interfaces ("Understanding"), and early prescriptive guidance ("Advising"). Stage 4 introduces workflow automation and bounded agents capable of orchestrating multi-step tasks. The architecture behind these technologies matters: systems must interoperate. The LP analytics hub, paired with case management, becomes the integration point through which insights flow and actions originate. Modern LP programs succeed not by collecting tools, but by designing a coherent ecosystem.

HOW TO ADVANCE

The Four Foundations of Advancement

What must evolve at each stage across Data, People, Processes, and Technology

	STAGE 1 Devices	STAGE 2 Analytics	STAGE 3 AI	STAGE 4 Agentic AI
Data	Fragmented device outputs	Integrated into analytics platform	Structure, tagged AI-ready	Governed for automated decisions
People	Manual reviewers Intuition-based	Dashboard analysts Data-driven methods	Interpret AI scores Feedback loops	Supervise automation AI stewards
Processes	Ad hoc investigation Inconsistent docs	Structured case workflows	AI feedback loops Model tuning	Auditable workflows Approval matrices
Technology	Camera, POS, EAS, RFID	LP analytics / EBR platform	Computer vision risk scoring, NLP	Workflow automation Bounded agents

Source: Agilence, 2026

Stage-by-Stage Advancement Playbooks

The following playbooks describe what organizations need to move from one stage to the next. These recommendations reflect lessons from the history of LP analytics and incorporate both vendor-provided capabilities and in-house model development paths.

Moving from Stage 1 → Stage 2 (Devices → Analytics)

As we discussed earlier, organizations operating primarily at Stage 1 rely heavily on devices for prevention and to document events but lack the context and analytical structure to interpret them. Cameras, POS terminals, and alarms each capture a fragment of reality, but analysts must manually assemble these fragments into a coherent picture. The result is reactive, labor-intensive, and difficult to scale.

Prerequisites

Transitioning to Stage 2 requires only modest maturity in each foundational dimension. Data must be reasonably consistent. POS data must link correctly to store and associate IDs, and basic event timestamps must align across systems. Teams must be willing to adopt dashboards and embrace data-driven investigation. Processes need enough structure to record and track cases. Technologically, organizations must adopt an LP analytics/EBR platform to serve as the investigative cockpit.

Actions to Advance

The first step is integrating major systems (typically POS, video, and inventory) into a unified analytics environment. Modern LP platforms make this manageable; full integration is not required up front. What matters most is establishing workflows that help analysts transition from scanning video to interpreting dashboards and following repeatable investigative paths. Organizations should focus early analytics efforts on a few high-value use cases: internal theft, refund abuse, void manipulation, high-risk SKUs, or stores with unusual shrink trends.

Challenges

LP teams accustomed to manual investigation may initially resist shifting to data-first workflows. Early dashboards may surface inconsistencies that require operational cleanup. Leaders should expect several months of adjustment as analysts learn to trust analytical insights and as systems begin linking previously disconnected data streams. The payoff is significant: Stage 2 organizations gain broad visibility into patterns across stores, improved consistency in case handling, and the ability to identify systemic issues.

Moving from Stage 2 → Stage 3 (Analytics → AI)

Stage 2 organizations often experience a growing pain: analysts spend increasing amounts of time building queries, tuning rules, and managing alert volumes. They find more anomalies than they can meaningfully investigate, and patterns grow too complex for manual rule systems to detect. These are the signals that an organization is ready for Stage 3.

Prerequisites

To adopt AI effectively, organizations do not need perfect data or years of historical records. They do need structured, consistently tagged transaction data; reliable case outcomes; and stable identifiers linking events across systems. The key readiness factor is not volume, but structure. AI models—whether vendor-provided or developed in-house—require enough signal to begin learning and enough consistency to avoid misinterpretation.

People readiness is equally important. Analysts must be willing to interpret confidence scores and risk ratings, not just binary alerts. They must also participate in feedback loops that help models improve. Processes must ensure that investigators record dispositions accurately so AI systems can distinguish true incidents from false positives. Technologically, organizations may either rely on vendor-provided AI features embedded in LP analytics platforms or build models internally.

Vendor models abstract much of the underlying complexity; they deliver capabilities such as anomaly detection, risk scoring, and correlated evidence gathering without requiring data science expertise. In-house models provide greater control but require LP–IT–Legal alignment on training data, permissible use, explainability, and model governance.

Actions to Advance

The shift to Stage 3 begins by selecting a small number of high-value use cases for AI assistance, such as refund fraud scoring, SCO anomaly detection, or analysis of historically high-risk SKUs. These tasks reflect the “Thinking” category of AI: models that analyze many variables simultaneously and prioritize cases that warrant attention. Some organizations may also adopt “Seeing” (computer vision) or “Understanding” (NLP-driven case linking) capabilities as part of this phase.

Human-in-the-loop workflows are essential. AI systems surface prioritized alerts; analysts validate them; and both successes and misses are fed back into the system. Over time, the model becomes better attuned to the organization’s environment, operational rhythms, and risk patterns. Transparency is critical to adoption. Whether using vendor AI or internal models, analysts must see the factors that led to an elevated risk score. Explainability builds trust, accelerates learning, and helps organizations tune models to reduce noise.

Challenges

Stage 3 introduces new complexities. Data inconsistencies become more visible as models encounter edge cases. Privacy and legal considerations become more prominent as organizations consider using employee data, video analytics, or sensitive operational metrics. False positives can erode trust if not carefully managed. Analysts may initially be skeptical of AI recommendations, particularly when outputs contradict established investigative instincts.
Success hinges on maintaining a tight feedback loop, being realistic about model maturity, and positioning AI as an assistant and not a replacement for LP expertise.

Moving from Stage 3 → Stage 4 (AI → Agentic AI)

Stage 3 organizations have reached a new equilibrium: AI surfaces and prioritizes most or much of the investigative work, and analysts focus primarily on validation and action. But the workflow is still human-initiated. Cases are opened manually; tasks are assigned manually; escalation processes depend on human triggers.

Stage 4 introduces bounded agency (the early form of Agentic AI) where systems can initiate multi-step workflows autonomously within strict parameters. As described earlier in the agency spectrum, LP technology is far from high autonomy, but the industry may be beginning to see the emergence of low- to mid-level agency, where software can execute structured tasks such as case creation, evidence gathering, documentation, or cross-store pattern correlation.

Prerequisites

Advancing to Stage 4 requires maturity across all foundations. Data must be clean and consistent enough to support automated decisions. People must be prepared to oversee automation rather than perform every step manually. Processes must be defined with enough precision for agents to follow them reliably. And technology must support interoperability across systems—case management, HR, workforce management, POS, video, and audit platforms.

Cross-functional alignment becomes critical. Automated actions can affect HR policy, operational procedures, customer experience, and even legal exposure. Governance structures must define which actions can be automated, which require human approval, and which must remain manual.

Actions to Advance

Organizations should begin with narrow, low-risk workflows suitable for bounded agents. Examples include automatically creating cases based on high-confidence anomalies, standardizing incident documentation, attaching relevant video clips, linking related incidents across locations, or automatically routing tasks to investigators based on workload. Automation should always begin with guardrails: human approval for sensitive steps, transparent audit logs for all agent actions, and explicit thresholds for when automation should pause or escalate to a human.
As reliability improves and trust grows, organizations can gradually expand agent responsibilities, eventually moving some workflows from “AI suggests, human approves” to “AI acts, human audits.”

Challenges

The challenges of Stage 4 are less technical than organizational. Automated actions can expose gaps in policy, create new forms of risk, or surface disagreements among LP, HR, Legal, and Operations. Model drift and workflow drift must be monitored continuously. Errors can carry reputational or regulatory consequences. Organizations must implement fairness checks, governance reviews, and documented escalation paths. Despite these challenges, early bounded agents offer significant value: standardized workflows, consistent documentation, faster investigations, and reduced manual workload—without requiring full autonomy. The limits to how much autonomy software agents will be given will be a decision for each organization to navigate based on their own systems, risk tolerance, needs, governance, and more.

Charting the Path Forward

Loss prevention is not a cost center; it is a margin multiplier. Every dollar recovered flows directly to the bottom line. Every fraud pattern identified prevents future losses. Every operational inefficiency surfaced by LP analytics creates value beyond shrink reduction, improving inventory accuracy, labor allocation, and process discipline across the enterprise. As LP systems mature, they become strategic data hubs capable of informing decisions far beyond the traditional scope of asset protection – from merchandising and workforce management to supply chain optimization and customer experience.

The Loss Prevention Maturity Model provides a framework for capturing this value systematically. It maps the progression from reactive, device-centric operations to intelligence-driven, and eventually autonomous, LP capabilities. Devices establish visibility into what happened. Analytics reveal why it happened. AI determines what matters most. And Agentic AI in LP, when it arrives, promises to complete the loop by executing what needs to be done.

Each stage builds on the previous. Organizations don't abandon devices when they adopt analytics, or discard analytics when they embrace AI. Instead, intelligence from later stages flows back to inform earlier ones: AI-driven insights guide where cameras are positioned, which products merit physical security, and how analyst time is allocated. The model is cumulative, not sequential.

Advancement requires more than technology purchases. The Four Foundations (Data, People, Processes, and Technology) must evolve in coordination. Data must be integrated and governed. People must develop new skills and trust new tools. Processes must be structured enough to support automation. Technology must interoperate across systems. Organizations that invest in only one dimension will stall.

Most LP organizations today operate at Stage 2. Stage 3 adoption is growing but uneven, and Stage 4 remains aspirational. Progress will be measured in years, not months, particularly for an industry that operates on long technology cycles. But the direction is clear, and organizations that begin building foundational capabilities now will be positioned to capture value as the technology matures.

The $112 billion shrink problem will not be solved by just working harder. It will be solved by working smarter: by replacing manual investigation with intelligent prioritization, by transforming siloed data into connected insight, and by elevating LP professionals from data processors to strategic decision-makers. This model provides the roadmap. The journey begins with an honest assessment of where you stand today.

Agilence is the leader in loss prevention and operations analytics, helping prominent retail, restaurant, grocery, and hospitality companies increase their profit margins by reducing preventable loss in all its forms. Every day, Agilence analyzes over 45 million transactions from over 200+ sources, including point-of-sale (POS), eCommerce, HR, labor, inventory, product, third-party delivery platforms, alarms, RFID, loyalty, access control, video surveillance, and more.

With integrated Case Management for streamlined incident investigations and Audit Management for compliance, risk, and safety standards, Agilence provides an end-to-end solution for loss prevention and operational excellence. Companies have saved millions of dollars by optimizing operations, identifying sources of margin erosion, and reducing shrink using Agilence.