Description

Title

Autonomous Inventory Management Platform for Automated Dispensing Cabinets (ADCs)

Novelty and IP Landscape Analysis

Automated Dispensing Cabinets (ADCs) like **BD Pyxis MedStation** and **Omnicell** systems have long been used to securely store and dispense medications at point-of-care. Various patents and products by BD, Omnicell, **KitCheck/Bluesight**, **Swisslog/Talyst**, etc., cover aspects of ADC technology and medication inventory management. Below we analyze the state of the art for each key feature of the proposed autonomous inventory management platform, in order to identify which aspects are novel and non-obvious:

- **Dynamic PAR-Level Adjustment:** ADCs traditionally rely on fixed PAR (periodic automatic replenishment) levels to trigger restocking when inventory falls below a threshold. Some advanced inventory management algorithms have been proposed in literature to **automatically adjust PAR levels based on usage frequency**[1]. For example, studies have shown that dynamically tuning min/max levels or PAR standards based on recent dispensing patterns can reduce stockouts and excess inventory[2][3]. However, we did not find any *granted patents* from major ADC vendors explicitly claiming automated PAR-level optimization. Omnicell's and Pyxis's core patents mention flagging low stock at static par levels for restock[4], but **no known prior patent dynamically updates the PAR thresholds themselves**. This suggests novelty in a system that continually recalculates optimal inventory levels per item per cabinet using real-time usage data and predictive analytics.

- **Real-Time Trip-Matching of ADC, eMAR, and ADT Events:** Ensuring that every medication removal from an ADC (dispense or waste transaction) is properly documented in the electronic medication administration record (eMAR) and aligns with patient status (ADT – admission, discharge, transfer events) is critical for patient safety and diversion prevention. Current practice often involves manual audits or basic reports to reconcile ADC dispenses with eMAR charting. Some third-party diversion analytics platforms (e.g. KitCheck's **ControlCheck** or Omnicell's acquired **ANiGENT MAAP** platform) consolidate multiple data sources to detect irregularities[5]. **KitCheck's patents** on drug diversion surveillance cover broad anomaly detection rules not limited to a single data source[6], implying that integrating ADC logs with eMAR and other hospital systems is contemplated. **Omnicell's "Medical equipment with diversion mechanism"** (published application US20170109480A1) describes computing a diversion score by comparing users' dispensing behaviors, essentially benchmarking individuals' ADC usage patterns[7]. However, these systems focus on analyzing data post hoc to flag suspicious users, rather than a *real-time trip-matching engine* that immediately correlates an ADC transaction with an eMAR entry and patient ADT status to catch discrepancies on the fly. We did not find prior patents explicitly teaching *trip matching across ADC, eMAR, and ADT in real time*, so this cross-system synchronization and verification appears to be a novel element of the proposed platform.

- **Automated Controlled-Substance Discrepancy Closure with Anomaly Scoring:** In current ADC workflows, controlled substance counts are verified at dispensing and returning, and any count **discrepancies** (e.g. a missing dose) require manual resolution by pharmacy/nursing staff. No automatic closure exists; instead, the discrepancy is investigated and either resolved (explained) or noted as unresolvable. Emerging analytics attempt to identify the causes of discrepancies. For instance, **ANiGENT's MAAP Analytics**, now part of Omnicell, claims to "identify the root cause of discrepancies in minutes" using machine learning across all users and locations[5]. Similarly, KitCheck's **AI-based diversion detection** (multiple patents issued in 2022-2024) monitors ADC transactions for anomalous patterns[8][9] and clusters data to detect possible misappropriation[10]. These suggest that scoring the *likelihood* of a discrepancy being diversion-related vs benign is known in the art. **Omnicell's 2017 application** also proposes generating a **"diversion score"** per user by averaging various suspicious behaviors (like stockout overrides, canceled transactions, etc.) across shifts[7][11], and even locking out users who exceed thresholds[12]. What appears *non-obvious* and novel in our proposed system is the idea of **automatically resolving low-risk discrepancies**: i.e. if the system can confidently reconcile a discrepancy (say, by finding that a nurse removed a medication on one patient but correctly documented it under a different patient), it could auto-close that discrepancy record, alleviating manual work. No prior patent explicitly covers *automated discrepancy resolution*. Our system's combination of an **anomaly scoring engine** with business rules to auto-resolve or escalate discrepancies (closing those within normal variance and flagging high-risk ones) would be a unique feature extending beyond existing discrepancy logs.
- **Expiry-Driven Redeployment of Pharmaceuticals Across Units:** Medication waste due to expiration is a costly issue. Hospitals often manually transfer soon-to-expire meds from low-usage areas to high-usage areas to get them used in time. We found **one prior patent (US20110257991A1)** addressing a similar concept at a network level: it describes maintaining an online inventory across pharmacies and **identifying over-stock or slow-moving products for redistribution** to other pharmacies to reduce expiries[13][14]. That system essentially acts as a broker for inter-pharmacy transfers of unused medications approaching expiry. Our proposed platform brings this concept *inside a single health system*: continuously monitoring each ADC's inventory for upcoming expirations and usage rates, then generating **redeployment tasks** (e.g. transfer 5 vials of Drug X from Unit A where they'll expire next month to Unit B which has higher demand). While the general idea of rebalancing inventory to avoid waste is not entirely new[13], implementing it in an **autonomous hospital-wide ADC management system**

appears novel. None of BD Pyxis or Omnicell's known patents explicitly cover *expiry-driven reallocation* at the cabinet level; this feature would likely provide a non-obvious improvement in inventory optimization.

- **Digitally Signed DEA/State Board Compliance Packet Generation:** Hospitals must maintain compliance records for controlled substances (e.g. dispensing logs, inventory counts, discrepancy resolutions, wastage documentation) and often prepare monthly or quarterly "compliance packets" for internal audit or regulatory inspection. Currently, assembling these reports is a manual or semi-manual process using data from ADC reports, pharmacy systems, etc., and then someone typically signs off on them. We did not find any prior patent or product that automates the **generation of a complete compliance packet with digital signatures** for authenticity. The novelty here lies in the **end-to-end automation of compliance reporting**: our system would compile all required records (transactions, discrepancies, resolutions, audit logs) into a tamper-evident PDF or electronic document, and apply a **digital signature** (e.g. using hospital's private key or a cryptographic certificate) to lock the content. This ensures integrity and non-repudiation, streamlining regulatory reporting. Given that digital signing is a well-known technology, the innovation is in applying it within the ADC context to create a **ready-to-submit compliance bundle**. We have not identified any prior art specifically on this, indicating a likely novel feature.

- **API Resilience (Self-Healing Interface Layer):** Modern ADC platforms interface with other hospital IT systems (EHR, pharmacy information systems, billing, etc.) through APIs or HL7 interfaces. These integrations can fail if an endpoint is down or messages time out, potentially causing data mismatches. A **self-healing API layer** would monitor interface transactions and automatically retry, queue, or reroute in case of failures – essentially ensuring robust, **fault-tolerant connectivity**. While the concept of resilient API middleware exists in general IT (with error handling, circuit breakers, etc.), we did not find any mention in medication management patents. For example, Omnicell's cloud platform patents discuss connecting to EHRs via APIs[15] but not specifically the self-healing aspect. This feature thus appears to be an *obvious engineering best-practice* but **non-obvious to implement and claim in the ADC domain**. Its novelty for patent purposes would stem from the specific context: ensuring ADC-to-EHR data synchronization continuity without manual intervention, which is not addressed in prior ADC systems.

- **Audit Logging via Blockchain or Immutable Ledger:** Using a **blockchain** or distributed ledger for logging medication dispensing events could guarantee an immutable audit trail – once a transaction is recorded, it cannot be altered, which is appealing for compliance. We found conceptual work suggesting blockchain for healthcare data auditing[16] and even specific use-cases like prescription management on blockchain[17], but not a concrete deployment in ADCs. No major ADC vendor appears to have patented the use of blockchain in their systems. (KitCheck/Bluesight's patents on privacy and diversion focus on anomaly detection algorithms, not the storage mechanism[6].) The novelty here is **applying blockchain in a hospital pharmacy automation setting**: e.g. each ADC event (dispense, return, witness, count adjustment) is hashed and appended to a ledger accessible to stakeholders, creating a **tamper-proof log** for audits. This may be non-obvious because of potential throughput and complexity concerns in a high-transaction environment, but if implemented, it significantly strengthens data integrity assurance.

- **Integration with Materials Management for Carousel-Based Restocking:** Both BD and Omnicell offer central pharmacy automation (carousels, robotic storages) that can be used to refill ADCs. Prior art exists for integrated systems – **McKesson's AcuDose-Rx with Connect-Rx** (later Omnicell) had a system where the central carousel would fulfill ADC restock lists and the process was tracked end-to-end[4][18]. **US Patent 8,571,701** and related family describe a system that generates restocking packages (totes) for decentralized ADCs, filled by a carousel or robot, and then verified at the ADC upon replenishment[4][18]. Thus, integrating ADC inventory with a materials management system (like an ERP or warehouse carousel software) has prior art foundation. Our platform likely uses a similar concept (automating restock picks when ADC stock is low). The novelty here may be incremental – e.g. using real-time usage to trigger just-in-time restocks or having the **API-resilient layer** manage the integration. Since patents from the early 2000s already cover much of the basic *carousel-ADC integration*[4], any new patent claims on this aspect need to emphasize improvements (like dynamic prioritization of restocks, or unified inventory visibility across carousel and ADCs).

In summary, the core novelty of the proposed system is in the **combination** of these features into a unified, autonomous platform for hospital medication inventory management. Individually, some elements have prior art (e.g. ADC restocking systems, diversion analytics) but others are fresh (automated PAR tuning, proactive redeployment, self-healing interfaces, blockchain logging). **No single existing solution or patent covers this full spectrum** of capabilities working together. Implementing dynamic inventory optimization alongside real-time data reconciliation, AI-driven anomaly resolution, and advanced compliance automation represents a **non-obvious integration** of technologies. This should provide strong IP defensibility, as competitors would have to replicate the whole chain of innovations to achieve similar functionality. The comprehensive nature of the system – spanning inventory control, compliance, IT integration, and security – creates a *picket fence* of novel improvements that can be protected via multiple claims, deterring design-around and supporting licensing opportunities to ADC manufacturers or healthcare IT providers.

Below, we present a **draft PCT-style patent application** for this autonomous ADC inventory management suite, including a detailed specification with embodiments and claims.

Draft Patent Application

Title of Invention

AUTONOMOUS INVENTORY MANAGEMENT PLATFORM FOR AUTOMATED DISPENSING CABINETS WITH DYNAMIC PAR OPTIMIZATION, ANOMALY RESOLUTION, AND COMPLIANCE AUTOMATION

Technical Field

This invention relates to medication dispensing systems in healthcare facilities. In particular, it concerns an **autonomous inventory management platform for Automated Dispensing Cabinets (ADCs)**, integrating dynamic stock level optimization, real-time data reconciliation, automated discrepancy resolution, and secure audit/compliance features. The technology touches the fields of hospital pharmacy automation, inventory control, data analytics, and health IT system integration.

Background Art

Automated Dispensing Cabinets are widely used in hospitals to securely store and dispense medications at the point of care. They improve medication availability and safety by controlling access and tracking usage[19]. However, current ADC systems face several challenges:

- **Static Inventory Levels:** Traditional ADCs use fixed PAR levels (minimum/maximum stock thresholds) for each medication. These often fail to account for changing usage patterns, leading to frequent stock-outs or overstock, tying up capital and causing waste. Pharmacy staff must manually review usage reports and adjust inventory levels periodically, an inefficient and error-prone process.

- **Manual Discrepancy Management:** ADCs track inventory counts, but discrepancies (mismatches between expected and actual counts, especially for controlled substances) are common. Resolving these discrepancies typically requires manual investigation by a pharmacist or manager who must comb through transaction logs, nursing records, and patient charts. This reactive process delays resolution and may allow drug diversion (theft or misuse) to go undetected[20][21].

- **Fragmented Data Systems:** Multiple IT systems are involved in medication use: the ADC records what was taken, the eMAR (Electronic Medication Administration Record) records what was given to which patient, and ADT systems record patient admissions, discharges, or transfers. In current practice, these data streams are **not automatically reconciled in real time**. A nurse might remove a medication from an ADC but forget to chart it in the eMAR, or remove it after a patient was discharged. Such inconsistencies might only be caught in after-the-fact audits, if at all.

- **Expiration and Waste:** Medications can expire on the shelf if not used in time. Hospitals attempt to redistribute soon-to-expire drugs (e.g., moving them from a slow unit to a busier unit) but this is done infrequently due to lack of visibility. As a result, valuable drugs are discarded unused, incurring financial loss and potential drug shortages.

- **Integration Downtime:** ADCs rely on integration with other hospital systems (pharmacy inventory systems, EHRs, billing) to function optimally. Interface downtimes or software glitches can disrupt medication availability (e.g., an ADC might not receive a new medication order due to interface failure). Current systems typically require IT intervention to troubleshoot and recover from such issues; there is no self-healing mechanism.

- **Compliance Burden:** Regulatory bodies (e.g., the DEA and state pharmacy boards) require extensive documentation for controlled substances. Pharmacies must produce reports of every dose dispensed, wasted, returned, along with discrepancy resolution reports, often on a monthly basis. Compiling these compliance packets is labor-intensive, and any data integrity issues (missing records, tampering) can have legal implications. Ensuring audit logs are tamper-proof and complete is a growing concern.

In summary, while ADCs have improved medication distribution, **opportunities for automation and intelligence remain**. There is a need for a unified platform that can **optimize inventory dynamically, reconcile data across systems, detect and resolve issues automatically, and produce compliance evidence with minimal human effort**. Such a system would reduce stockouts and waste, enhance diversion control, and ensure regulatory compliance – ultimately improving patient care and operational efficiency.

Disclosure of the Invention (Summary)

The invention provides an **Autonomous Inventory Management Platform** built around Automated Dispensing Cabinets, addressing the above challenges with a suite of integrated, intelligent modules. In one aspect, the system continuously **adjusts PAR levels** and restocking schedules for each medication in each ADC based on real-time usage, predicted demand, and upcoming expiration dates, thus ensuring optimal inventory levels and proactive redistribution of expiring stock.

In another aspect, the platform performs **real-time tripartite matching of ADC transactions, eMAR entries, and ADT patient status**. Every dispense, return, or waste event from an ADC is immediately cross-referenced with the patient's electronic MAR record and the hospital admission/discharge/transfer system. If a mismatch is detected (e.g., medication removed but not charted on a patient, or removed for a patient who has been discharged), the system flags it within seconds for investigation or automatically rectifies documentation if appropriate. This **closed-loop reconciliation** significantly tightens medication administration safety and diversion prevention.

The invention further includes an **Automated Discrepancy Resolution Engine**. This subsystem aggregates data such as dispensing logs, return/waste records, user activity patterns, and patient data to analyze any count discrepancies or anomalous transactions. Using configurable rules and machine learning anomaly detection, it assigns an **anomaly score** to each discrepancy or suspicious event. Low-risk discrepancies (e.g., a known benign pattern) can be auto-closed by the system, with a note explaining the likely cause, whereas high-risk cases (e.g. potential diversion events) are escalated to managers with a compiled evidence dossier. Over time, the system "learns" from resolved cases to improve its scoring and resolution suggestions. This greatly reduces the manual burden of daily count reconciliation and enhances controlled substance oversight.

Another feature of the invention is **expiry-driven inventory redeployment**. The platform continuously monitors the expiration dates of all medications in all ADCs across the facility. When it identifies medications nearing expiry in low-usage locations, it automatically generates transfer tasks to move those medications to higher-usage locations or consolidates them in the pharmacy, thereby **preventing waste**. This may involve printing instructions for pharmacy technicians or updating inventory systems to reflect new locations. The redeployment suggestions optimize both usage (so that drugs approaching expiry get used) and inventory levels (preventing overstock in some areas).

To ensure **data integrity and compliance**, the system maintains an **immutable audit log** of all critical events (transactions, corrections, user access, etc.) using a secure ledger technology. In one embodiment, a **blockchain-based ledger** is used, wherein each transaction record is cryptographically hashed and linked. This provides a tamper-evident trail ideal for regulatory audits. Building on this, the platform can automatically generate **digitally signed compliance packets** – for example, a monthly report of all controlled substance activity, discrepancy resolutions, and audit log excerpts. The packet is cryptographically signed by the system (or authorized personnel) to certify its accuracy. This one-click compliance reporting ensures that regulators or auditors can be given trustworthy records without painstaking manual compilation.

The platform also includes a **resilient API integration layer**. It interfaces with external systems (EHR, pharmacy information management, inventory procurement systems, robotic warehouses) through well-defined APIs or HL7 messages. If an external system goes down or a network glitch occurs, the layer will queue transactions, retry connections, or failover to a backup interface. It effectively **self-heals** integration faults, ensuring continuous operation of the ADCs. For example, if the connection to the hospital ADT system is temporarily lost, the platform caches patient admission/discharge data and syncs it when the link restores, so that no dispense events go un-checked against patient status. This design minimizes downtime and manual IT interventions.

Finally, the invention provides seamless **integration with materials management systems**, such as pharmacy carousel storage or robotic dispensing machines, for automated restocking. When the dynamic PAR module flags items that need restock, the system can directly interface with an automated storage (carousel) to **pick and dispatch medications for ADC replenishment**. Each restock transaction is tracked from the central pharmacy to the ADC, verifying that the correct medications and quantities are delivered. In embodiments, restocking can be optimized by grouping tasks across multiple ADCs, and the carousel control software can prioritize picks based on urgency (e.g., preventing an ADC stockout). This integration closes the loop from demand sensing to fulfillment, aligning with prior approaches to integrated medication distribution[4]but enhanced with the intelligence of the present invention.

In sum, the invention marries the strengths of automation (speed, accuracy, 24/7 operation) with advanced analytics (pattern recognition, prediction) to create a **self-managing ADC ecosystem**. It is novel in combining *inventory optimization*, *real-time data validation*, *AI-driven anomaly management*, and *secure compliance logging* in one platform. Hospitals deploying this system can expect fewer medication outages, reduced waste, earlier detection of diversion, streamlined compliance, and overall lower operating costs. The invention is also designed with **licensing and integration in mind**: it can overlay on existing ADC hardware via software updates or middleware, making it commercially attractive for vendors like BD, Omnicell, or third-party software providers to license and incorporate into their product offerings. Its comprehensive feature set and patent-protected innovations will allow the owner to enforce exclusivity across a broad range of critical ADC functionalities, establishing a defensible market position and opportunities for partnerships or direct sales to healthcare systems.

Brief Description of the Drawings

FIG. 1 is a **system architecture diagram** illustrating an exemplary deployment of the autonomous inventory management platform in a hospital environment. It shows multiple Automated Dispensing Cabinets (ADCs) on nursing units connected to a central server/cloud service. Interfaces to external systems (EHR, pharmacy system, ADT system, and a pharmacy carousel) are depicted, along with the internal modules of the platform (dynamic inventory optimizer, data reconciliation engine, anomaly detection AI, compliance ledger, etc.). [[FIG. 1]]

FIG. 2 is a **flowchart** showing the process of **dynamic PAR-level adjustment** for a single medication in an ADC. The flowchart details how usage data is collected, how an optimal par level is calculated (taking into account factors like average daily usage, variability, lead time, and expiry), and how restock orders are generated and sent to the pharmacy when inventory drops to the re-order point. [[FIG. 2]]

FIG. 3 is a **sequence diagram** illustrating **real-time trip matching** of an ADC dispense event with eMAR and ADT events. The diagram outlines the sequence: a nurse removes a medication from the ADC, the ADC notifies the central platform, which queries or receives data from the eMAR system (to see if an administration record exists for that patient and drug) and from the ADT system (to verify patient is active). The outcomes (match found, or mismatch leading to alert) are shown. [[FIG. 3]]

FIG. 4 is a **schematic view** of the **anomaly scoring and discrepancy resolution interface**. It shows an example dashboard listing open discrepancies or flagged events, each with an anomaly score. One entry is expanded to illustrate data the system has compiled: the user involved, medication, timestamps, related eMAR documentation, and an AI-generated explanation or recommendation (e.g., “Likely documentation delay – auto-resolved” or “Potential diversion – investigation required”). This figure also shows how the user can provide feedback or additional input, which the system uses to update its machine learning model. [[FIG. 4]]

FIG. 5 is a **block diagram** focusing on the **compliance and audit subsystem**. It depicts how transactions are recorded on an immutable ledger (for example, blocks chaining together), and how a compliance report generator assembles data from the ledger and the hospital’s records into a digitally signed PDF. A notional depiction of a blockchain and a generated compliance report (with digital signature stamp) is included. [[FIG. 5]]

(In the drawings, like reference numerals denote like components. For instance, 100-series numbers in FIG. 1 refer to components of the ADC units, 200-series in FIG. 2 refer to steps in the flowchart, etc. These figures are exemplary and not to scale, intended to illustrate possible embodiments of the invention.)

Description of Embodiments

Embodiment 1: System Overview (FIG. 1).

Referencing FIG. 1, the autonomous inventory management system (**100**) comprises multiple network-connected Automated Dispensing Cabinets (**101**), a central server or cloud-based control system (**120**), and integration points to external systems (**130**). Each ADC unit 101 includes a local computer (**102**) that controls its drawers/dispensing mechanisms (**103**) and manages local inventory data (medication types and counts stored). The ADCs 101 are deployed in various hospital units (e.g., ICU, Med-Surg, ER) and are connected via a secure hospital network (**110**) to the central server 120.

The central server 120 (which could be on-premise or hosted in the cloud) hosts the intelligent modules of the platform. In this embodiment, these modules include:

- a **Dynamic Inventory Optimizer** module (**121**),
- a **Data Reconciliation Engine** (**122**),
- an **Anomaly Detection & Resolution AI** (**123**),
- a **Compliance Ledger & Reporting** module (**124**),
- and an **Integration Manager** (**125**) for external interfaces.

External systems 130 shown here include the hospital's **EHR system** 131 (which encompasses eMAR functionality), the **ADT system** 132 (patient administration system), the **Pharmacy Information Management System** 133 (for inventory and purchasing), and an **Automated Pharmacy Carousel** 134 located in the central pharmacy.

Communication between the central server 120 and ADCs 101 is bidirectional. ADCs regularly send usage data (transactions, inventory levels, alerts) to the server, and receive back commands or updates (e.g., new par values, restock orders, configuration changes). The Integration Manager 125 handles communication with external systems 130, for example using APIs or HL7 messaging. This includes sending queries (e.g., asking EHR 131 if a dose was charted) and receiving event feeds (e.g., real-time ADT admission/discharge messages).

Embodiment 2: Dynamic PAR-Level Adjustment (FIG. 2).

The Dynamic Inventory Optimizer 121 runs a continuous process for each medication in each ADC. FIG. 2 outlines the logic. At step **201**, the module retrieves recent usage data for the medication: how many doses were dispensed, days with zero usage, etc., over a configurable lookback period (say 30 days). It also pulls relevant context: current stock level, last restock date, upcoming patient census (from ADT 132, if it can correlate that a certain unit will have more patients requiring that med), and the medication's expiration date if applicable.

At step **202**, the module forecasts demand using a predictive algorithm. This could be a simple moving average or a more complex machine learning model considering day-of-week trends, seasonality, or even specific patient orders. The forecast yields an expected usage rate (e.g., 5 doses/week with a peak of 2 on Mondays).

Using this forecast, step **203** computes an **optimal PAR level range**: a maximum level (target stock to have on hand) and a minimum threshold (trigger point for restock). It does so by balancing the costs of stockout vs. overstock. For instance, it might use a service level target (e.g., 99% probability of no stockout) to determine safety stock. The algorithm also factors in the lead time to replenish from central pharmacy, and whether some stock will expire soon (effectively reducing usable inventory). The output could be, for example: "Keep 20 doses on hand, reorder when 5 left."

Decision block **204** checks if the current configured PAR levels differ significantly from the optimal values. If so, at **205** the ADC's configuration is automatically updated: the new min/max values are sent to the ADC 101. This might also trigger a one-time restock if current inventory is below the new minimum. The ADC's user interface can indicate the updated par to pharmacy staff (e.g., via a message on the cabinet or central dashboard).

If the levels are already optimal, the module simply loops or waits until next evaluation cycle. The evaluation can be continuous or periodic (e.g., once every 24 hours or when significant usage changes occur).

Step **206** handles **restock order generation**. Whenever an ADC's inventory falls below the minimum (whether old or newly updated threshold), the optimizer flags the need for restock. Instead of just firing off an alert, the system integrates with the central pharmacy. It creates a restocking task listing the drug and quantity needed (e.g., "ADC in ICU needs 15 of Morphine 2mg vial"). This task is sent to the **Automated Pharmacy Carousel 134** via Integration Manager 125.

In embodiments, as shown in FIG. 2, this interaction is automated: the carousel (or pharmacy tech using it) receives an electronic **pick list** (step **207**). The carousel 134 lights up or directs the operator to pick the required medication (**208**). The items are then placed in a restock tote with a barcode or RFID tag identifying the destination ADC and slot. At step **209**, when the tote is delivered to the ADC, the ADC verifies the contents (via scanning) and updates its counts accordingly. This closes the loop with minimal human decision-making; pharmacy staff just follow optimized restock instructions.

Through this embodiment, the system continuously right-sizes inventory. **Dynamic PAR adjustment** ensures fast-moving meds don't run out (PAR goes up if usage increases) and slow movers don't hog space (PAR goes down if usage drops). Unlike static systems, this adapts to changes like a new protocol that increases usage or a ward closure that decreases demand, all without requiring a manual review by pharmacy.

Embodiment 3: Real-Time Trip-Matching & Anomaly Alerts (FIG. 3).

FIG. 3 illustrates how the **Data Reconciliation Engine 122** links ADC events with eMAR and ADT data in real time. The timeline begins with a nurse performing a dispense transaction at the ADC (step **301**). For example, Nurse Jane authenticates at the ADC 101, selects patient John Doe (who is in Room 100, Unit ICU), and removes 1 vial of Medication X at 10:00 AM. The ADC's local system logs this event with details: {timestamp, user, patient ID, med ID, quantity, transaction type=dispense}. The ADC 101 immediately sends this event to the central server 120 (message **302**).

Upon receiving the event, at step **303**, the Data Reconciliation Engine 122 queries the eMAR system 131 (which could be a part of the EHR) to find if there is a corresponding medication administration record. It looks for an order for Medication X for patient John Doe around that time. This can be done via an API call or by having subscribed to eMAR events. Simultaneously, in step **304**, it checks the ADT system 132 to confirm the patient's status (e.g., John Doe is indeed admitted in ICU, not discharged or transferred at that time).

Two outcomes are possible:

- **Case A: Everything matches.** Suppose the nurse documented the administration in the eMAR at 10:05 AM for John Doe, and ADT confirms John Doe is an inpatient in ICU.

The engine finds the match (step **305**). It could log an "all good" verification internally. Optionally, it could augment the ADC record with the eMAR reference for future auditing (i.e., mark the transaction as reconciled). No alert is needed. This entire check might occur within a few seconds without any user visible action.

- **Case B: Mismatch or missing data.** If by a certain short timeout (say 5 minutes) no matching eMAR entry is found, or ADT says John Doe was discharged at 9:00 AM (before the dispense), then an anomaly is detected. In step **306**, the system raises a **real-time alert**. This alert can manifest in multiple ways: a notification on the pharmacy dashboard, an email/text to a diversion specialist, or even a popup at the ADC for the nurse (in some implementations). The alert contains details: "Medication X removed for Patient John Doe by Nurse Jane at 10:00, but no administration documented (Patient possibly not eligible)." The nurse or pharmacist can then investigate immediately—perhaps catching a documentation omission while it's fresh or preventing an improper use.

If a documentation was just delayed, the system can keep checking for a late eMAR entry for a configurable window. If it finds one (e.g., nurse charts at 10:15), it can auto-clear the alert with a note "Late charting reconciled."

This tripartite matching addresses a significant gap in today's systems by creating an **immediate feedback loop**. It essentially enforces proper practice: meds taken must be meds given (or returned/wasted with documentation). Additionally, by involving ADT data, it can catch scenarios where a nurse might withdraw medication for a patient who isn't currently admitted (which is a red flag for diversion, e.g., using a discharged patient's profile to obtain a drug). Conventional ADC reports might list transactions by patient, but without cross-checking ADT they might not realize the patient was not there. This embodiment therefore greatly enhances security and compliance by **preventing and detecting errors or misuse in real time**.

Embodiment 4: Anomaly Scoring and Discrepancy Resolution (FIG. 4).

The **Anomaly Detection & Resolution AI 123** works in conjunction with the reconciliation engine. While FIG. 3 dealt with a single transaction instantaneously, FIG. 4 shows the higher-level dashboard and process for ongoing monitoring and resolution of anomalies and discrepancies.

In FIG. 4, a user interface (**400**) is depicted, likely accessible to pharmacy managers or a diversion team. The system aggregates various anomaly types into this dashboard:

- **Discrepancies:** These are inventory count mismatches in the ADC (e.g., drawer says 5 were expected but only 4 found). Each discrepancy entry (**410**) lists the ADC, medication, when it was detected, and involved users if known.
- **Documentation Mismatches:** These come from the trip-matching process above where an ADC dispense lacked matching eMAR/ADT. Each such event (**420**) is listed with nurse, patient, drug, time.
- **Other Anomalies:** The system can also flag patterns (like frequent transaction cancellations by a user, excessive wasting of a particular drug by one nurse, etc.). Each pattern or threshold-based alert (**430**) would be an entry.

Each entry has an associated **Anomaly Score (401)** – a numeric or categorical rating of how suspicious or severe the issue is. For example, a minor count discrepancy of a non-controlled med might score 20/100 (low), whereas multiple undocumented opioid removals by the same nurse might score 90/100 (high). The scoring can be AI-driven: the model compares against typical behavior patterns and known risk factors. For instance, an entry showing "Nurse Jane removed 5 vials morphine for 5 different patients, none of which have matching eMAR, and those patients were discharged – Score 95" would combine multiple red flags.

The **Resolution Assistant** panel (**440**) on the UI shows details and suggestions when an entry is selected. For a discrepancy entry 410 example: "ADC on Unit 3 reported -1 fentanyl missing." The system might correlate that just before the count, a nurse had a transaction that was interrupted or an emergency code occurred – context that suggests a likely cause. The AI might suggest: "Possible cause: transaction #12345 by Nurse A was not properly completed. Recommend marking as resolved with reason 'Inadvertent removal not administered.' Score 15 (low)." The user can accept this suggestion, and the system will automatically log the discrepancy as resolved with that reason, and even update the ADC's count if appropriate (with a phantom adjustment record for audit).

For a more suspicious case (like 420 or 430 entries with high score), the assistant 440 might say: "Pattern of undocumented removals by Nurse B. 3 occurrences in 2 weeks involving controlled substances. Score 88 (high). Recommended action: Investigate for diversion, interview Nurse B. Do you want to open a case?" The interface might allow the manager to click "Open Case", which then moves this to a case management system (as noted, integration or built-in case tracking can be present). Supporting evidence – such as a list of all transactions by that nurse, timing relative to shifts, any witness co-sign data for wastes – can be bundled automatically for the investigation.

Notably, the system stores feedback. If the manager marks an alert as "False positive" or enters the actual resolution (e.g. "Medication found elsewhere, false discrepancy"), the AI model can learn from that. Over time, this reduces noise by tuning the anomaly detection to site-specific patterns.

This embodiment automates much of what would otherwise be a retrospective, manual audit process. By assigning scores and even auto-resolving trivial issues, it **frees up human experts to focus on the truly problematic cases**. It also provides a clear digital trail of what happened and how it was addressed, which is invaluable for compliance and legal protection (showing due diligence in monitoring).

Embodiment 5: Immutable Audit Logging and Compliance Reporting (FIG. 5).

FIG. 5 illustrates the **Compliance Ledger & Reporting module 124**. At its core is an **Immutable Ledger 500**, conceptually represented as a blockchain or sequential log. Every significant event in the system – a dispense, return, discrepancy, resolution action, user login, system configuration change – can be recorded as a transaction on this ledger. Each entry might include a timestamp, event type, involved identifiers (user ID, med ID, patient ID if relevant), and metadata (quantity, old/new values, etc.). A cryptographic hash of each entry is computed, and each new entry's hash incorporates the previous hash (forming a chain). This means if anyone tried to alter past records, the hashes would mismatch, flagging tampering[16].

The ledger 500 can be implemented using blockchain technology (either a private blockchain for the hospital or even a public one for maximum trust, though probably private/permissioned for privacy reasons). Alternatively, a simpler append-only file with secure checksums could suffice – the key is that once data is written, it's effectively read-only and verifiable.

On top of this ledger, the **Compliance Report Generator 510** runs. Periodically or on-demand, it compiles a **Compliance Packet 520**. This could be, for example, a PDF or XML file that contains: - A cover page with hospital info and reporting period. - Controlled substance summary (e.g., total dispensed per drug, total wasted, etc.). - Detailed transaction logs for the period (or references to the ledger entries). - List of all discrepancies and their resolution. - Sign-off sections.

Once compiled, at step **511**, the system applies a **Digital Signature 530** to the packet. This uses a certificate (e.g., the hospital pharmacy director's digital certificate, or a system certificate) to produce a signature that can be validated by regulators. The digital signature ensures the document hasn't been altered after generation.

The figure shows a stamp 530 (conceptually) on the report 520 indicating it's signed and perhaps the signer name/time. The system can either automatically use an administrator's private key (with permission), or more simply, a cryptographic hash of the report can be generated and stored such that any change is detectable.

This one-click generation means that at the end of the month, the pharmacy can produce the exact set of records the DEA or state board requires, with confidence in completeness and integrity. **Conventional methods** might involve printing multiple reports from the ADC software, getting managers to sign each page, etc. – those are prone to omissions and are very time-consuming. In contrast, this embodiment leverages the comprehensive data already captured (especially given modules above) to output a thorough compliance document. Because it's digitally signed and based on an immutable ledger, it provides **strong evidence of accuracy**, addressing regulators' concerns about record tampering.

Additionally, since the ledger is machine-readable, future audits can be performed by external parties using software – for example, an inspector could be given a read-only access to a blockchain explorer that shows all transactions. This kind of transparency can become a selling point for hospitals under strict oversight.

Embodiment 6: API Resilience and Self-Healing Integration.

While not explicitly illustrated in a figure, a further embodiment covers the **Integration Manager 125** in FIG. 1 with its self-healing capabilities. In practice, the Integration Manager runs numerous background services handling connections to EHR, ADT, pharmacy systems, etc. For each external interface, it maintains a heartbeat and error log.

For example, if the EHR API (for eMAR queries) becomes unresponsive, the manager will detect this (no response, or an HTTP error) at time **T1**. Immediately, it can switch to a backup strategy: maybe querying a replicated database or queueing the requests. It also notifies the anomaly engine that eMAR data is delayed (to not falsely flag every transaction during downtime). The Integration Manager will then periodically retry the EHR connection. Once it succeeds at time **T2**, it will flush any queued requests and back-fill any missed checks. Similarly, if an HL7 ADT feed stops (no messages in X minutes), it can attempt a reconnect.

This module can also use redundancies: e.g., two network routes or a local cache of the last known patient list if ADT is down. In extreme cases, if external systems are offline for long, the platform can enter a **graceful degradation mode** where it continues core ADC functions locally and marks any unverifiable data for later reconciliation when systems recover.

Such resilience ensures the overall solution is robust in real hospital conditions, where IT downtime happens. It **minimizes manual IT support**, as the system "heals" itself and keeps track of what needs reconciling later. This embodiment complements the others by ensuring that the advanced features (like trip-matching) don't themselves become points of failure.

Embodiment 7: Licensing and Deployment Considerations.

The architecture of this platform is modular, which facilitates **licensing and integration** into existing products. For instance, an ADC manufacturer could license the Dynamic Inventory Optimizer and embed it into their cabinet software to offer an "auto PAR" feature. Or a hospital could adopt the anomaly detection module as an add-on to their current ADC management software. The system can run on standard server hardware or cloud environments, and it uses standard protocols (REST APIs, HL7, FHIR, etc.) to communicate, making it interoperable.

From an enforcement perspective, each functional component provides a point of differentiation that can be protected. For example, the specific method of using both eMAR and ADT data to cross-verify ADC transactions is a novel process that competitors would have to replicate to match efficacy, which this patent would preclude. The combination of blockchain logging with digital report signing is another distinct inventive step that is patentable and enforceable.

In commercialization, the platform could be offered as a whole or in parts. Hospitals could implement it to achieve an "autonomous pharmacy" vision where inventory management and diversion surveillance are largely automated (aligning with industry goals of reducing human workload in medication distribution). The return on investment comes from labor savings (pharmacy and nursing time saved), reduced drug waste, and improved compliance (avoiding fines or patient harm). These economic drivers support adoption. The invention's features also position it well against competing offerings, providing a strong **licensing opportunity** for companies looking to offer next-generation medication management solutions.

By covering a broad but interrelated set of innovations, this patent filing effectively creates a **picket fence** around the concept of an AI-driven, self-managing ADC ecosystem. It would be difficult for a competitor to implement a similarly comprehensive solution without infringing at least some claims, thereby protecting the commercial value and giving the patent holder leverage for cross-licensing or strategic partnerships in the healthcare technology domain.

Footnotes:

[1] [PDF] Optimization of Automated Dispensing Cabinets at a Large ...

https://d-scholarship.pitt.edu/44044/1/Robb_Masters%20Essay_Dec22.pdf

[2] Optimizing Medication Distribution in Automated Dispensing Cabinets

https://www.researchgate.net/publication/362649131_Optimizing_Medication_Distribution_in_Automated_Dispensing_Cabinets_Dashboard_Implementation_and_Evaluation

[3] Best practices for using Automated Dispensing Cabinets (ADC ...

https://www.capsahealthcare.com/blog/medication-management/automated_dispensing_cabinets_adc_best_practices/

[4] [18] [19] US20030105552A1 - Carousel product for use in integrated restocking and dispensing system - Google Patents

https://patents.google.com/patent/US20030105552A1/en

[5] Omnicell Acquires ANiGENT | Omnicell

https://www.omnicell.com/resources/news/omnicell-acquires-anigent/

[6] [8] [9] [10] Patents - Bluesight

https://bluesight.com/patents/

[7] [11] [12] [20] [21] US20170109480A1 - Medical equipment with diversion mechanism - Google Patents

https://patents.google.com/patent/US20170109480

[13] [14] US20110257991A1 - Pharmacy Product Inventory Control or Redistribution - Google Patents

https://patents.google.com/patent/US20110257991A1/en

[15] US12248826B2 - Cloud-based healthcare platform - Google Patents

https://patents.google.com/patent/US12248826B2/en

[16] Pharmaceutical Supply Chain with Blockchain: Drug Tracking ...

https://gkc.himss.org/resources/pharmaceutical-supply-chain-blockchain-drug-tracking-tracing-and-verification-last-mile

[17] Blockchain prescription management system - Google Patents

https://patents.google.com/patent/US20190198144A1/en