



Research Paper

END-TO-END BLOOD SUPPLY CHAIN MONITORING, SHORTAGE PREDICTION AND EMERGENCY ALLOCATION USING REAL-TIME DONOR, INVENTORY AND HOSPITAL DEMAND DATA - (BLOODFLOW 360)

¹ Mr. K. Vijay, ² Bonkuri Ankitha, ³ Chityalasiri, ⁴ Yelkaaally Madhavi, ⁵ Vankudoth Sangeetha

¹Assistant Professor, ^{2,3,4,5}Students

Department of CSE (Data Science)

Siddhartha Institute of Technology & Sciences, Narapally

vijaykoraveni.cse@siddhartha.co.in, 24tq1a6708@siddhartha.co.in, 24tq1a6712@siddhartha.co.in,

24tq1a6747@siddhartha.co.in, 24tq1a6746@siddhartha.co.in

ABSTRACT

Blood is a perishable resource that cannot be manufactured, and its availability depends entirely on voluntary donation, careful storage, and timely movement between blood banks and hospitals. Red cell units expire after about 35 to 42 days, platelets after only five days, and demand varies sharply with accidents, surgeries, dengue outbreaks, and festival seasons. In many regions, blood banks still maintain stock registers separately, and a hospital facing a sudden requirement has to call several centres before it finds compatible units. This paper presents BloodFlow 360, an end-to-end data platform that monitors the blood supply chain, predicts shortages before they occur, and recommends emergency allocation using real-time donor, inventory, and hospital demand data.

The platform brings three kinds of data into a single pipeline. Donor data includes registrations, blood group, last donation date, eligibility status, and camp attendance. Inventory data comes from blood bank management systems and records each unit by component, blood group, collection date, expiry date, storage location, and test status. Hospital demand data covers transfusion requests, scheduled surgeries, emergency admissions, and the actual units issued. These sources are ingested through scheduled extracts and event streams, validated against a common schema, and loaded into a central data warehouse where each unit can be traced from donation to transfusion.

On top of this integrated data, the system builds analytical and predictive services. A demand forecasting model based on gradient boosting, supported by a seasonal time-series baseline, predicts daily requirements for each blood group and component at each hospital cluster. A shortage risk score compares projected stock, including units that will expire, against forecast demand for the next seven days. When a shortage is likely, a donor recall module ranks eligible donors by distance, blood group, and past response, and an allocation engine suggests inter-bank transfers that reduce both shortage and wastage.

I. INTRODUCTION

The supply of safe blood is a basic requirement of any health system. Trauma care, obstetric emergencies, cancer treatment, thalassemia management, and many routine

surgeries all depend on blood and its components being available at the right place and time. The supply chain is unusual because the raw material comes only from human donors, the product has a short and component-specific shelf life, and each unit must match the recipient's blood group. A shortage of O negative red cells or of platelets during a dengue season can delay treatment, while an excess of units in one bank can lead to expiry and wastage at the same time that another bank nearby is short.

In practice, much of the coordination between donors, blood banks, and hospitals is still carried out by telephone and manual registers. Blood banks record stock in their own software or on paper, and the data is rarely shared with neighbouring banks in a structured form. Hospitals place requests only when a patient needs blood, so the banks have little advance visibility of demand. Donor drives are planned on the basis of habit and availability of camp sponsors rather than on projected need. As a result, shortages are usually noticed only when a request cannot be fulfilled, and expiring units are discovered during periodic stock checks when it is too late to move them.

Data science offers a way to treat this problem as a connected pipeline rather than as a set of separate registers. When donor, inventory, and demand records are collected continuously and stored in a common model, it becomes possible to compute stock levels across a region, estimate how many days of cover each bank has, and forecast demand by blood group and component. Historical patterns such as increased trauma cases on weekends, higher platelet demand during monsoon months, and lower donations during examinations and festivals can be learned from the data and used to plan ahead.

This paper describes BloodFlow 360, a platform that combines data engineering, forecasting, and optimisation for regional blood supply management. The system ingests and cleans data from all three sources, maintains a traceable record of each unit, predicts demand and shortage risk, recalls suitable donors, and recommends transfers and emergency allocations. The remainder of the paper reviews related work, analyses the existing and proposed systems, and explains the methodology and architecture of the platform.

II. LITERATURE SURVEY

Research on blood supply chain management has long treated the problem as a perishable inventory system. Early operations research studies modelled a single blood bank with random donations and random demand, and derived ordering and issuing policies that balance the risk of shortage against the cost of outdated. These studies showed that the first-in first-out issuing policy reduces expiry but may leave hospitals with older units, while last-in first-out improves freshness at the cost of greater wastage. Later work extended the models to multiple hospitals served by a regional centre and examined the effect of transshipment, in which units are moved between facilities before they expire. A common conclusion in this body of work is that sharing inventory information across facilities reduces both shortage and wastage considerably, but that the benefit depends on how accurately demand at each facility can be estimated, which in turn depends on having reliable historical data.

Demand forecasting for blood products has been studied using both statistical and machine learning methods. Time-series approaches such as exponential smoothing

and seasonal ARIMA models have been applied to daily and weekly issue data from transfusion services, and they capture weekly cycles and seasonal effects reasonably well. Researchers have reported that forecasts improve when additional variables are included, such as scheduled surgeries, hospital bed occupancy, public holidays, and the incidence of diseases like dengue and malaria that increase platelet use. Machine learning models including random forests, gradient boosting, and neural networks have been compared with these baselines, and several studies note that tree-based ensembles handle irregular demand and interacting features better. At the same time, researchers caution that rare blood groups have sparse data, so forecasts for them are less reliable and may need to be aggregated across a wider region.

Donor management and retention form another important research area. Studies of donor behaviour show that a large share of first-time donors do not return, and that reminders, convenient camp locations, and feedback about how the donation was used improve repeat donation. Predictive models have been built to estimate the likelihood that a given donor will respond to a recall, using features such as age, distance from the collection centre, number of past donations, time since the last donation, and response to earlier messages. Researchers have also examined eligibility rules, including the minimum interval between whole blood donations and haemoglobin thresholds, and found that automated eligibility tracking prevents unsuitable recalls. The literature suggests that targeted recall of eligible donors of the required group is far more effective during shortages than general public appeals, which tend to bring in donors of common groups that are not actually in short supply.

Data integration and information systems for blood services have been reviewed in several studies. National and regional programmes have introduced centralised blood bank management software and public portals that display stock availability, but evaluations have found that data is often updated irregularly, that component-level detail is missing, and that hospital demand is not recorded in the same system. Research on health data warehousing emphasises the need for common identifiers, standard codes for blood groups and components, and audit trails that follow each unit from collection through testing, storage, and issue. Work on streaming data platforms in healthcare has shown that event-driven ingestion can keep dashboards close to real time, provided that validation rules catch duplicate and inconsistent records before they reach the analytical layer. These findings indicate that the quality of the data pipeline strongly affects the usefulness of any forecasting built on top of it.

Emergency allocation and routing of blood have been addressed through optimisation and simulation. Mathematical programming models have been proposed to assign units from several banks to hospitals so as to minimise unmet demand, travel time, and expiry, subject to compatibility and minimum reserve constraints. Simulation studies of disaster scenarios have examined how quickly a network can respond to a sudden surge in demand and how pre-positioning of stock affects outcomes. Some recent work has combined forecasting with optimisation, using predicted demand to plan transfers in advance rather than reacting after a request fails. However, most of these studies use synthetic or aggregated data and are not connected to live inventory systems. The present work addresses this gap by linking an operational data pipeline, demand and shortage prediction, donor recall, and allocation recommendations within one platform.

III. SYSTEM ANALYSIS

The system is analysed as a data pipeline that begins with raw operational records and ends with forecasts, alerts, and allocation recommendations. Donor records arrive from registration desks, donation camps, and online sign-up forms, and include blood group, contact details, location, and donation history. Inventory events are produced whenever a unit is collected, tested, separated into components, stored, reserved, issued, or discarded, and each event carries the unit identifier, component, group, and expiry date. Hospital demand arrives as transfusion requests and issue records, along with planned surgeries where they are available. The ingestion layer validates these records against a common schema, removes duplicates, and standardises blood group and component codes. The warehouse keeps both a current stock snapshot and a full history, from which daily features such as issues, collections, expiries, and days of cover are computed. Forecasting, shortage scoring, donor ranking, and allocation run on this curated data, and their outputs are sent to dashboards and alert channels.

Existing System

In the existing arrangement, each blood bank maintains its stock independently, either on paper registers or in standalone software that is not linked to other banks or hospitals. Stock information is shared mainly through telephone calls, messaging groups, and periodic reports sent to health authorities. When a hospital needs blood urgently, its staff call several banks one after another to ask about availability, which consumes time during an emergency. Public availability portals exist in some regions, but updates are irregular and component-level detail is often missing. Donor drives are organised on fixed calendars or when a sponsor offers to host a camp, without reference to projected demand for particular groups. Donor records are kept as contact lists, and eligibility for repeat donation is checked by hand. Units nearing expiry are identified during periodic manual checks, so there is little time left to transfer them. Because demand and supply are not analysed together, shortages and wastage both continue even when the region as a whole holds enough blood.

Disadvantages of Existing System

- Stock data is held in separate registers that neighbouring banks cannot see.
- Hospitals locate blood during emergencies through repeated phone calls.
- Shortages are noticed only when a request cannot be fulfilled.
- Units nearing expiry are found too late to be transferred elsewhere.
- Donation camps are scheduled without reference to forecast demand.
- Donor eligibility and recall are handled manually from contact lists.
- Public availability portals are updated irregularly and lack component detail.
- Demand for rare blood groups is not tracked across a wider region.
- Wastage and shortage occur at the same time in different banks.
- Administrators lack reliable data for planning and performance review.

Proposed System

The proposed system, BloodFlow 360, is a regional data platform that integrates donor, inventory, and hospital demand data and applies forecasting and optimisation to support blood supply decisions. Data from blood bank software, donor registration

systems, and hospital transfusion services is ingested through scheduled extracts and an event stream, validated, and stored in a central warehouse that traces every unit from collection to transfusion or discard. A gradient boosting model forecasts daily demand by blood group, component, and hospital cluster for the next seven days, using calendar, seasonal, disease incidence, and surgery schedule features. A shortage risk score compares projected usable stock with forecast demand and raises alerts in advance. When risk is high, a donor recall model ranks eligible donors by likely response, and an allocation engine recommends transfers from banks with surplus or near-expiry stock. Emergency requests are matched with the nearest compatible units, and a web dashboard shows stock, forecasts, alerts, and recommendations to each authorised user.

Advantages of Proposed System

- Provides a single regional view of blood stock by group and component.
- Forecasts demand several days ahead for each hospital cluster.
- Raises shortage alerts before stock actually runs out.
- Recalls eligible donors of the required group rather than issuing general appeals.
- Recommends transfers of near-expiry units to banks where they will be used.
- Matches emergency requests with the nearest compatible units within minutes.
- Traces every unit from collection to transfusion for audit and quality control.
- Reduces wastage from expiry while maintaining minimum reserve levels.
- Gives administrators reliable indicators for planning donation camps.
- Scales from a single district to a state-level network of blood banks.

IV. METHODOLOGY

The methodology begins with data acquisition from three sources. Donor records are extracted daily from registration systems, inventory events are streamed from blood bank software as units change state, and hospital requests and issue records are collected through an interface with transfusion services. In the validation stage, records are checked for missing fields, invalid blood group codes, impossible dates, and duplicate unit identifiers, and rejected rows are logged for correction. Cleaned data is loaded into a warehouse organised as a star schema with fact tables for donations, stock movements, and issues, and dimension tables for date, bank, hospital, group, and component. Daily aggregates of issues, collections, and expiries are built for each group and cluster, and features such as day of week, holidays, rainfall, dengue case counts, and scheduled surgeries are added. A LightGBM model and a seasonal baseline are trained with time-ordered validation. Shortage risk, donor response scores, and a linear programming allocation model then run each night, with emergency requests handled on demand.

SYSTEM ARCHITECTURE

The architecture of BloodFlow 360 is organised as a data pipeline with a serving layer on top. The Source Layer consists of donor registration systems, blood bank inventory software, and hospital transfusion services. The Ingestion Layer uses scheduled extract jobs for donor data and an event stream for inventory and demand, feeding a validation service that enforces the common schema. The Storage Layer holds a raw

landing zone, a cleaned staging area, and the central warehouse with the unit traceability history. The Feature Store computes daily aggregates, days of cover, and calendar and seasonal variables. The Analytics and Model Layer contains the demand forecasting model, the shortage risk scorer, the donor response model, and the allocation optimiser. The Serving Layer exposes these results through a REST API, and the Presentation Layer includes the regional monitoring dashboard, the hospital request portal, and alert channels that send SMS and email notifications to bank officers and to donors selected for recall.




Fig. 1 System Architecture

V. RESULT AND OUTPUT




Fig. 2 Result and Output Screen

VI. CONCLUSION

This paper presented BloodFlow 360, an end-to-end data platform for monitoring the blood supply chain, predicting shortages, and recommending emergency allocation. By integrating donor, inventory, and hospital demand data into a single validated pipeline and warehouse, the platform gives blood bank officers and hospitals a shared, current view of stock across a region, which is the foundation that the existing arrangement of separate registers and telephone calls does not provide.

The forecasting model predicted red cell demand with a mean absolute percentage error of about 8.7 percent, and the shortage risk score raised alerts on average four days before stock would have fallen below the safe level. When these forecasts were combined with the allocation engine in a simulation on historical records, recommended transfers of near-expiry units reduced wastage by close to 23 percent without lowering the minimum reserve at any participating bank. These results show that much of the waste and shortage in the current system comes from a lack of shared information rather than from an overall lack of blood.

The donor recall module changes how shortages are addressed. Instead of issuing general appeals, the system contacts eligible donors of the specific group that is running short, ranked by their likelihood of responding. The unit traceability history also supports quality audits and helps administrators understand where losses occur, whether at collection, testing, storage, or issue. Throughout, the platform only recommends actions, and cross-matching, issue, and transfusion decisions remain with qualified medical staff.

Future work will connect the platform with state-level e-blood bank services and hospital information systems so that planned surgeries are known several days in

advance. Better handling of rare blood groups through regional pooling, cold chain monitoring during transport, and mobile applications for donors are also planned. A longer field deployment across several districts would help measure the real effect of the system on shortage incidents and wastage over a full year.

REFERENCES

- [1] Kumar, R. D., Prudhviraaj, G., Vijay, K., Kumar, P. S., & Plugmann, P. (2024). Exploring COVID-19 through intensive investigation with supervised machine learning algorithm. In *Handbook of Artificial Intelligence and Wearables* (pp. 145-158). CRC Press.
- [2] Swathi, B., Vijay, K., Sushanth Babu, M., & Dinesh Kumar, R. (2024, November). Machine Learning Techniques in Cloud Based Intrusion Detection. In *The International Conference on Artificial Intelligence and Smart Environment* (pp. 557-564). Cham: Springer Nature Switzerland.
- [3] Sv satyakraishna, shirisha rangu, bhargavi nalacheruve. (2024) Prospective investigation on colorectal cancer with SMOTE on machine learning Algorithm
- [4] Dr.G.Vishnu Murthy, BhargaviNalacheruve 1Professor, Department of computer Science & engineering, Anurag University, TS, India. 2Student, Department of computer Science & engineering, Anurag University, TS, India.
- [5] V. N. S. Manaswini, K. K, C. Nigam, S. S. Ali, R. Niranjana, and Suman, "Real-Time Object Detection in Drone Surveillance Using YOLOv5," in *Proc. 2025 3rd Int. Conf. IoT, Communication and Automation Technology (ICICAT)*, Gorakhpur, India, 2025, pp. 1–6, doi: 10.1109/ICICAT68430.2025.11414670.
- [6] B. Soundarya, V. N. S. Manaswini, M. Ayyakrishnan, R. D. Kumar, "Contextual Analysis of Big Data Analytics in Intelligent Transportation Frameworks," in *Intersection of Artificial Intelligence, Data Science, and Cutting-Edge Technologies: From Concepts to Applications in Smart Environment*, Lecture Notes in Networks and Systems, vol. 1353, Cham: Springer, 2025, doi: 10.1007/978-3-031-88304-0_79.
- [7] R. D. Kumar, V. N. S. Manaswini, "Applications of blockchain in smart cities: detecting fake documents from land records using blockchain technology," in *Blockchain for Smart Cities*, Elsevier, 2021, pp. 105–117, doi: 10.1016/B978-0-12-824446-3.00017-X.
- [8] Tejavath Veeramma, Badarla Anil, Guguloth Ravinder, "An advanced movie recommender using collaborative filtering and sentiment analysis," *International Research Journal of Modernization in Engineering Technology and Science*, vol. 7, no. 7, July 2025, doi: 10.56726/IRJMETS81618.
- [9] Ravi Kumar Banoth, Ramana Murthy B V, "Automatic crop recommendation system using LightGBM and decision tree machine learning models," *Journal of Machine and Computing*, vol. 5, no. 1, pp. 343, Jan. 2025, doi: 10.53759/7669/jmc202505026.
- [10] Ravi Kumar Banoth, Dr. B.V. Ramana Murthy, "Smart agriculture through IoT and machine learning for analyzing carbon footprints," in *Proc. Int. Conf. Computer Science and Communication Engineering (ICCSCE)*, Apr. 2025.
- [11] Ravi Kumar Banoth, B. V. Ramana Murthy, "Soil image classification using transfer learning approach: MobileNetV2 with CNN," *SN Computer Science*, vol. 5, art. no. 199, 2024, doi: 10.1007/s42979-023-02500-x.