

A Connectivity-Resilient Blockchain Architecture for Digital Twin in Bangladesh RMG

Alvi Nahian Ishti¹, Sadikur Rahman¹, S.M.Fahamid Rahman¹, M M Dian¹, Md Mohiuddin sikder¹

¹Department of Industrial and Production Engineering(IPE)

Jashore University of Science and Technology, Jashore-7408, Bangladesh

alvinahianishti@gmail.com, shakkar300@gmail.com, fahamidrahmansiam2020@gmail.com,
mohiuddinsikder123@gmail.com, mmdiananan23@gmail.com

ARTICLE INFO

Keywords:

Blockchain,
Connectivity resilience,
digital shadow,
Edge computing,
Hyperledger fabric,
Production monitoring,
Ready-made garment.

ABSTRACT

Bangladesh's ready-made garment (RMG) sector generates more than 80% of national export earnings, yet real-time production monitoring has gained almost no foothold on its factory floors. This paper combines two bodies of evidence. Documentary analysis of the Bangladesh Garment Manufacturers and Exporters Association's digital programme, covering DigiProd Pass, the Digital Factory Passport and AWARE, shows that these are compliance instruments built around European Digital Product Passport requirements rather than around the cost, connectivity and skill constraints that keep factories from adopting production-monitoring technology. A targeted survey of 61 practitioners then tested whether those constraints are recognised at the production level. Respondents rated skilled-personnel shortage, initial cost and interruption of network and electricity supply as substantial obstacles, and a Friedman test found no reliable severity ordering among the five barrier items, so they are treated as a joint constraint set rather than ranked. The strongest inferential result is that the skills barrier was rated well above respondents' assessment of their own factories' capacity to absorb minimal-training systems. Perceived monitoring value and system acceptability were positively associated. From this evidence the paper derives design requirements and specifies a connectivity-resilient edge-gateway architecture using RFID or QR capture, a permissioned Hyperledger Fabric ledger with private data collections, and an offline-first store-and-forward layer. At the level specified the system is a digital shadow rather than a closed-loop digital twin, and it has not been prototyped.

1. Introduction

1.1 Introduce the Problem

Bangladesh is the world's second-largest exporter of apparel. The ready-made garment (RMG) sector generates over 80% of national export earnings, contributes close to a tenth of GDP, and has in recent years approached USD 47 billion in annual exports (Salman et al., 2024). Digital maturity, however, remains elusive: factories still fall short of rival economies in connectivity, capital and technical skills (Salman et al., 2024), and a survey of 142 textile and apparel factories placed overall Industry 4.0 maturity at 1.91 on a five-point scale, with digital integration the weakest of the measured dimensions (Mim et al., 2024).

This is the gap digital twin technology is positioned against. A digital twin is a virtual model of a physical production system, continuously updated by live data and used for monitoring, simulation and optimisation (Grieves & Vickers, 2017; Tao et al., 2019); in apparel manufacturing it has been associated with improved traceability and faster response to buyer requirements (Jayalakshmi et al., 2024). Adoption within Bangladeshi RMG nonetheless remains close to negligible. Four constraints plausibly explain why, and each is tested against practitioner evidence in Section 4: interruption of network and electricity supply, initial cost, recurrent cost, and a shortage of skilled technical personnel.

1.2 Explore Importance of the Problem

Over the past two years BGMEA has committed to three separate digital compliance programmes. In May 2025 it signed a memorandum of understanding with DigiProd Pass Ltd. and Digital Architect Ltd. for a 24-month pilot of blockchain-anchored Digital Product Passports (DPPs) (DigiProd Pass, 2025). In October 2025 it announced a strategic cooperation with Digital Bridge Partners for an AI-enabled Digital Factory Passport covering environmental, social and governance reporting (BGMEA, 2025). In May 2026 it signed a further agreement with the Dutch traceability platform AWARE, covering fibre-to-garment records for roughly 4,000 member factories, with factories retaining ownership of the data they generate (Fibre2Fashion, 2026; Textile Focus, 2026). The driver in each case is regulatory: the Ecodesign for Sustainable Products Regulation has been in force since July 2024, and the delegated act fixing DPP requirements for textiles carries an indicative adoption date of 2027, followed by a transition period (European Union, 2024).

This is a different problem from the one a digital twin addresses. A DPP records where a garment has come from; a production-monitoring system indicates what a sewing line is doing at a given moment. Knowing where a bundle is and how long it has stayed at a station is what allows a production manager to identify a delay before it cascades into a missed shipment date. The distinction matters because, as Section 5.3 argues, BGMEA's programme does not appear to have been designed around the limits that keep factories from adopting production-monitoring technology in the first place. The industry therefore finds itself in an unusual position, in which its apex body is rolling out several advanced compliance applications while the obstacles to a closely related one have not been mapped at the level of the production floor.

1.3 Describe Relevant Scholarship

Research on Industry 4.0 and digital twin adoption in emerging economies returns repeatedly to cost, skills and infrastructure (Raj et al., 2020; Kamble et al., 2018; Kumar et al., 2021). Bangladesh-specific work has ranked such barriers using Delphi and DEMATEL (Salman et al., 2024) and fuzzy interpretive structural modelling (Karmaker et al., 2023), and has measured adoption maturity across 142 factories (Mim et al., 2024). A separate body of work proposes Hyperledger Fabric traceability frameworks for the Bangladeshi RMG supply chain, though these remain conceptual or simulated (Niloy et al., 2024; Al Amin et al., 2025). A third stream, developed largely outside Bangladesh, has built and benchmarked blockchain traceability for apparel production and for edge-constrained hardware (Faridi et al., 2023; Honar Pajoooh et al., 2021). Section 2 examines all three.

1.4 State Objectives and Their Correspondence to Research Design

The main objective is to develop a practitioner-informed specification for a low-cost, connectivity-resilient production-monitoring architecture for Bangladesh RMG. Five research questions structure the empirical work.

RQ1 asks which operational and infrastructure conditions practitioners perceive as obstacles. RQ2 asks whether any barrier is perceived as more severe than the others. RQ3 asks how strongly practitioners value real-time visibility, work-in-progress tracking and bundle dwell-time monitoring. RQ4 asks whether perceived value is associated with acceptability of the proposed system. RQ5 asks what further barriers practitioners raise unprompted. All five are answered in Section 4. A sixth objective, assessing whether BGMEA's programme is directed at the barriers identified, is addressed in Section 5.3 through documentary analysis; the seventh, deriving requirements and specifying an architecture, in Section 6.

The correspondence between objective and method is deliberately asymmetric. The first six objectives are evidential; the last is constructive. No hypothesis about system performance is tested here, because no system has been built. The survey establishes which constraints a design must respect, not whether any design that respects them works.

1.5 Contribution

The contribution is the evidentiary chain the paper documents rather than any single component of the architecture: practitioner experience generates a perceived operational problem; the survey produces evidence about it; the evidence implies a system requirement; the requirement becomes a design decision and then a specific architectural feature. Prior Bangladeshi barrier studies stop at ranking, and prior Fabric-for-RMG proposals are not grounded in practitioner evidence. A second, narrower contribution is the connectivity-resilience mechanism itself, an offline-first, sequence-aware synchronisation pattern for RMG production data on constrained hardware. Figure 1 sets out the chain and marks the point at which validation stops.




Figure 1. *Research Design and the Evidentiary Chain From Evidence to Architecture*

Note. The dashed transition marks the boundary of validation in this study. The architecture is assessed against published benchmarks and derived targets; it has not been prototyped or deployed.

2. Literature Review

2.1 The RMG Industry's Digital Transformation Imperative

The apparel industry operates as a buyer-driven value chain: design, branding and retail sit in developed countries while manufacturing concentrates in economies with lower production costs. Bangladesh occupies the manufacturing segment under increasing pressure, as buyers demand stricter compliance and improved sustainability, demand patterns become less predictable, and lead times shorten (Arnob et al., 2020), all of which push factories towards digitalization and cleaner production (Chaity et al., 2025; Fahim et al., 2024). The sector's readiness for that transformation is limited by insufficient capital, inconsistent connectivity and a shortage of

technical expertise (Salman et al., 2024), and persistent social-compliance concerns sharpen the pressure further (Sharfuddin Ashbin et al., 2024).

2.2 Digital Twin, Digital Shadow and Digital Model

A digital twin maintains synchronization in both directions through a continuous data connection, which is what distinguishes it from a static simulation (Grieves & Vickers, 2017). In manufacturing the technology now covers condition monitoring, predictive maintenance and process optimization (Tao et al., 2019; Qi & Tao, 2018); in apparel it has been applied to virtual sampling, real-time line monitoring, inventory optimization and traceability (Jayalakshmi et al., 2024).

Precision about what counts as a twin matters for the claim this paper makes. Kritzinger et al. (2018) distinguish three levels by the degree of automation in the data connection. In a digital model both data flows are manual. In a digital shadow the physical-to-digital flow is automated but the return path is not. Only where both directions are automated, so that a change in the digital object propagates to the physical object without human mediation, does the arrangement qualify as a digital twin proper. Much of the applied literature uses the term more loosely, and systematic reviews have noted the resulting inconsistency in what the label denotes (Jones et al., 2020; Hu et al., 2021). Section 6.6 scopes the present contribution against this taxonomy.

Digital production monitoring, in the narrower sense used here, means continuously capturing what a line is doing without necessarily constituting a full twin. Within apparel manufacturing this appears mainly as one application inside digital twin implementations (Jayalakshmi et al., 2024) rather than as a research stream of its own, and no dedicated literature on lightweight, RMG-specific production-monitoring systems was identified in the sources reviewed. The practical entry point for a sewing line is modest: bundles carrying RFID or QR tags move through stations and generate a work-in-progress record that feeds a lightweight digital representation. This is bounded and physically instrumentable and falls well short of simulating an entire production system.

2.3 Barriers to Digital Twin and Industry 4.0 Adoption

Deploying a digital twin requires capital, connectivity, technical talent and organisational readiness, and manufacturers in developing countries frequently possess none of the four (Pathri & Ganduri, 2025). Cross-country and India-focused studies name cost, skills and infrastructure with some regularity (Raj et al., 2020; Kamble et al., 2018; Kumar et al., 2021); sustainability-oriented supply-chain work adds data security and policy uncertainty (Fuller et al., 2020; Luthra & Mangla, 2018; Pasi et al., 2020).

Bangladeshi studies exist, but share two features limiting what can be drawn from them for design purposes. Salman et al. (2024), Karmaker et al. (2023) and Sayem et al. (2022) all reveal relative importance and interdependence among barriers, using Delphi and DEMATEL, fuzzy interpretive structural modelling with MICMAC, and a catalogue of critical barriers respectively. All three, however, draw on small expert panels or firm-level respondents rather than practitioners working at the line, and all three stop at ranking without carrying the result into a system specification. Mim et al. (2024) survey 142 factories but measure adoption maturity rather than the operating conditions a system would have to tolerate. The survey in Section 4 addresses the first limitation; the second is a design problem, and Section 6 treats it accordingly.

One case study indicates the practical scale of these constraints. DBS Group, a medium-sized Bangladeshi exporter that has implemented IoT, blockchain and ERP, reports an IoT implementation costing approximately USD 150,000, around 30% of its workforce lacking the necessary skills, and power outages affecting operations roughly 30% of the time (Hosain, 2025). This is a single firm and generalises to nothing, but it is a firm well ahead of the sector average, which makes the cost and outage figures informative about what the median supplier faces.

2.4 Blockchain and Hyperledger Fabric in Garment Supply Chains

Blockchain work on Bangladeshi RMG has so far produced frameworks rather than deployments. Al Amin et al. (2025) and Niloy et al. (2024) both propose Fabric-based supply-chain systems for the sector; the former states plainly that it is untested in local factories and does not address connectivity gaps, and the latter is evaluated by simulation with no offline or edge capability. Trials outside RMG report the same constraints binding in practice, in e-commerce (Haque et al., 2022) and FMCG logistics (Hasan et al., 2022), while Ahmed and MacCarthy (2021) document a lead-actor-driven rollout at the fibre producer Lenzing.

Work conducted outside Bangladesh has taken the technical questions considerably further, and converges on a common omission. Faridi et al. (2023) implement Chain Apparel, a Fabric-based traceability framework with IoT capture at each workstation, reporting write throughput saturating near 400 transactions per second and noting that no disconnected operating mode is addressed. Hou and Sun (2021) propose a permissioned-ledger security architecture for garment production using conventional RFID readers at each station, without a full prototype evaluation. Agrawal et al. (2021) develop a traceability framework for textiles and clothing using unique identifiers assigned at source, assuming stable connectivity between all tiers. Honar Pajoo et al. (2021) run Fabric on Raspberry Pi 4 hardware and identify the validation phase as the binding throughput constraint. Reviews of the edge and blockchain intersection confirm the pattern: most architectures assume at least periodic connectivity to a cloud or core network, and fully disconnected operation is not treated as a design case (Nguyen et al., 2024).

The performance literature is more sobering. Zulkarnain et al. (2025) benchmark Fabric on Raspberry Pi-class boards against server hardware and report that the small boards sustain only tens of transactions per second under load, roughly 50 to 100 TPS against more than 1,100 TPS on server-class machines, with CPU contention rather than network bandwidth binding first once operating-system and container overhead are counted; Lee et al. (2023) model Fabric latency analytically. None of this work, however, tests Fabric under the combination this paper is concerned with: intermittent connectivity, low-cost hardware, and RMG production data on a single sewing line.

2.5 BGMEA's Current Digital and Blockchain Initiatives

BGMEA is not standing still, and any fair account of the sector has to acknowledge as much. The DigiProd Pass memorandum established a 24-month DPP pilot framed explicitly around EU compliance, with DigiProd Pass Ltd. supplying the platform and Digital Architect Ltd. handling local deployment, lifecycle assessment and workforce training (DigiProd Pass, 2025). The Digital Factory Passport extends the same logic from product to facility, using AI-driven analytics to consolidate factory-level ESG reporting (BGMEA, 2025). The AWARE agreement covers fibre-to-garment records for roughly 4,000 factories, with data sovereignty retained by the factories themselves (Fibre2Fashion, 2026; Textile Focus, 2026).

All three are compliance instruments. They establish what a product is made of, where it has been, and how the facility performs against external standards, for the benefit of an external audience; none models or optimises how a line is running for the factory's own benefit. All three also assume that a factory can already generate, timestamp and transmit data reliably, and the results in Section 4 suggest this is precisely where smaller suppliers are weakest. What can be observed is real institutional momentum directed at a problem other than the one respondents identify as most severe.

2.6 Research Gap and Positioning of This Study

A 40-source literature matrix compiled for this study, covering 2019 to 2026 across IEEE Xplore, ScienceDirect, SpringerLink, MDPI, Emerald, the ACM Digital Library, arXiv and PMC, supports a narrow gap claim. Every individual component of the proposed architecture has clear prior art: Fabric-based apparel traceability has been implemented and benchmarked, Fabric has been run on Raspberry Pi-class hardware, RFID and QR capture for garment production has been specified, and Bangladeshi barrier studies have documented cost, skill and infrastructure constraints. What the matrix did not identify was any study combining these under the condition that distinguishes the present case, namely an environment in which loss of connectivity is a routine operating state rather than an exceptional fault. Table 1 positions this study accordingly.

Two further gaps follow. Existing Bangladeshi barrier studies rank obstacles at industry level but do not connect severity to specific system requirements through practitioner-level evidence. And BGMEA's programme, which represents substantial institutional momentum in an adjacent technical domain, has not been assessed against practitioner-perceived priorities. This study addresses the second and third gaps directly and situates its architecture against the first. The contribution is therefore an integration and feasibility claim rather than a first-of-kind technology claim: neither Hyperledger Fabric nor private data collections are novel, and what the review found unclaimed is the narrower combination of offline-first synchronisation, sequence-aware conflict handling and local plausibility screening on low-cost hardware for sewing-line data.

Table 1. *Positioning of This Study Against the Closest Prior Art*

Study	Setting	Ledger	Edge hardware	RFID / QR	Offline operation	Evaluation
Faridi et al. (2023)	Apparel manufacturing	Fabric	Server class	Yes	Not addressed	Prototype; benchmarked
Hou and Sun (2021)	Garment production	Permissioned	Not specified	Yes	Not addressed	Conceptual architecture
Agrawal et al. (2021)	Textile and clothing	Permissioned	Not specified	Yes	Stable link assumed	Case study
Honar Pajoooh et al. (2021)	Generic edge IoT	Fabric	Raspberry Pi 4	No	Not addressed	Prototype; throughput
Niloy et al. (2024)	Bangladesh RMG chain	Fabric	None	Not specified	Not addressed	Simulated
Al Amin et al. (2025)	Bangladesh RMG chain	Fabric	None	Not specified	Acknowledged gap	Conceptual
Nguyen et al. (2024)	Edge/blockchain review	Multiple	Various	Not specified	Identified as a gap	Systematic review
This study	Bangladesh RMG sewing line	Fabric	Raspberry Pi class	Yes	Central design case	Design proposal only

Note. Cells record what each source addresses, not the quality of its contribution. “Not addressed” indicates that the dimension falls outside the source's stated scope. Entries are drawn from a 40-source literature matrix covering 2019 to 2026.

3. Method

3.1 Research Design

The study combines two evidential components with one constructive component: a cross-sectional practitioner survey, establishing which operating constraints are recognised at the production level and with what relative severity; a documentary analysis of BGMEA's digital programme, establishing what the sector's apex body is currently building; and a derivation of design requirements from both, together with the literature matrix, specifying an architecture intended to satisfy them. The documentary strand is methodologically separate from the survey and is not treated as survey evidence at any point.

3.2 Survey Instrument and Measures

The instrument comprised nineteen items: three collecting respondent and factory characteristics, one collecting factory size (excluded, see Section 3.3), fourteen agreement statements scored from 1 (strongly disagree) to 5 (strongly agree), and one open-ended item. Table 2 sets out the block structure; full item wording appears in Appendix A.

Table 2. *Structure of the Survey Instrument*

Block	Items	Purpose
Respondent profile	Q1–Q3	Position, sector experience, factory type

Block	Items	Purpose
Excluded	Q4	Factory size; not used in any analysis (Section 3.3)
Infrastructure	Q5–Q7	Perceived disruption from network and electricity interruption; offline recording
Cost	Q8–Q9	Initial and recurrent cost as perceived barriers
Skills and training	Q10–Q11	Technical personnel shortage and capacity to absorb low-training systems
Perceived value	Q12–Q15	Value attributed to line visibility, WIP tracking and dwell-time monitoring
Acceptability	Q16–Q18	Perceived usefulness, outage-tolerance preference and conditional adoption intent
Open-ended	Q19	Unprompted account of the biggest practical adoption difficulty

Note. Items Q5 to Q18 were scored 1 (strongly disagree) to 5 (strongly agree). Full item wording appears in Appendix A. Q7 and Q17 are directionally leading; see Section 3.2.

Two features of the instrument govern how the results should be read. First, every item from Q5 to Q18 is an agreement statement, so the data record perceptions and stated preferences rather than operating conditions. A respondent agreeing that network interruptions create problems is reporting a perception of disruption, not an observed outage rate, and no claim about outage frequency rests on these responses anywhere in this paper. Second, Q7 and Q17 are worded in a direction that invites assent, asserting respectively that data should continue to be recorded during connectivity loss and that an outage-tolerant system would suit the respondent's factory better than a continuously connected one. Both are close to tautological for anyone who values the underlying activity, so their agreement levels are uninformative in isolation; Section 4.6 tests whether including them changes any substantive conclusion.

3.3 Sampling and Data Collection

Responses were collected through a structured online questionnaire distributed to practitioners in production, industrial engineering, planning, maintenance, information systems and management roles within Bangladeshi RMG factories. Sampling was purposive rather than probabilistic, since the target population is defined by occupational role rather than by an enumerable frame (Etikan et al., 2016). Sixty-one usable responses were obtained, with no missing data on Q1 to Q3 or Q5 to Q18 and no case excluded for missingness. The open-ended item was answered substantively by 17 respondents (27.9%). One respondent returned the neutral midpoint on all fourteen agreement items; the response was retained, and its removal alters no result reported below.

Factory size was excluded from analysis in its entirety, and is used neither as a descriptor nor as a predictor nor as a stratification variable. Its response options were evidently edited during the collection period, producing overlapping categories in which a factory of 2,200 employees could fall into either of two bins, so any mapping onto a consistent scheme would be a reconstruction rather than a measurement.

3.4 Analytical Procedure

The analysis plan was fixed before the results were inspected. Descriptive statistics are the primary evidence: median and interquartile range are reported alongside mean and standard deviation, because Likert responses are ordinal, together with the proportions selecting 4 or 5 and 1 or 2.

Four inferential comparisons were pre-specified. A Friedman test compares the five barrier items (Q5, Q6, Q8, Q9, Q10) as related repeated ordinal ratings, with Kendall's W as the effect size and Holm-corrected post-hoc Wilcoxon tests planned only if the omnibus test proved significant. Three Wilcoxon signed-rank tests then compare network against electricity disruption, initial against maintenance cost, and the skills barrier against the claimed ease of adopting minimal-training systems; these are paired ordinal comparisons within the same respondents. Effect sizes are reported as r over non-tied difference pairs. Associations between block composites use Spearman's rank correlation. Tests are two-tailed at the 5% threshold, with exact probabilities reported.

Latent-variable methods were considered and rejected: partial least squares structural equation modelling, confirmatory factor analysis and exploratory factor analysis all require a larger case-to-indicator ratio or more indicators per construct than this instrument provides, and two of the five blocks contain only two items. Cronbach's alpha is reported for each block, but the composites are treated as descriptive averages rather than estimates of latent constructs. Open-ended responses were coded thematically using deductive codes from the barrier literature plus inductive codes where responses did not fit; the frame appears in Appendix B.

3.5 Secondary Research for Industry Context and Architecture Design

BGMEA's blockchain activity was established from dated secondary sources, with every claim cross-checked against at least two independent outlets. The architecture in Section 6 draws on the same 40-source matrix, covering Fabric benchmarking, edge-computing blockchain deployments, offline synchronisation patterns and prior Bangladeshi RMG studies. The design targets in Table 8 were derived by pairing published hardware performance figures with the constraints the survey and literature jointly identify; none is presented as an empirical finding of this study.

4. Results

4.1 Respondent Profile

Table 3 and Figure 2 describe the achieved sample. Production and industrial engineering roles account for the largest share and, with planning and supply-chain roles, make up rather more than half the sample, consistent with the intention of reaching practitioners close to the line rather than executives. Tenure is skewed towards earlier career stages, with roughly seven respondents in ten reporting five years of sector experience or less. Denim is the largest single factory type.

Two characteristics bear on interpretation. The sample is not a probability sample, and multiple respondents may work at the same factory, so responses cannot be treated as independent factory-level observations. Several cells in the position and factory-type distributions fall below ten respondents, so no subgroup comparison by role or product type is attempted.

Table 3. *Respondent and Factory Profile*

Characteristic	Category	n	%
Current position (Q1)	Production / industrial engineering	20	32.8
	Planning / supply chain	13	21.3
	Management	10	16.4
	Maintenance / technical	7	11.5
	IT / MIS	7	11.5
	Other	4	6.6

Characteristic	Category	n	%
Sector experience (Q2)	Less than 2 years	22	36.1
	2 to 5 years	21	34.4
	6 to 10 years	12	19.7
	More than 10 years	6	9.8
Factory type (Q3)	Denim	19	31.1
	Knit	9	14.8
	Woven	9	14.8
	Sweater	7	11.5
	Other	17	27.9

Note. N = 61. Percentages are of the achieved sample and may not sum to 100.0 because of rounding. Factory size (Q4) is excluded and appears in no table. Several cells fall below ten respondents, so no subgroup analysis by position or factory type is reported.




Figure 2. Respondent Profile by Position, Sector Experience and Factory Type

Note. N = 61. Bar labels give counts; the horizontal axis gives the percentage of the achieved sample. Factory size (Q4) is excluded (Section 3.3).

4.2 Descriptive Findings

Table 4 reports full descriptive statistics and Figure 3 the response distributions. Agreement is broadly positive: twelve of the fourteen items return a median of 4, and the proportion agreeing exceeds 60% on eleven of them. Disagreement is generally low, exceeding a quarter of the sample on only one item. Two items stand apart. The claim that a factory can easily adopt systems requiring minimal training returns the lowest mean in the instrument and the only median below 4. The assertion that recording should continue during connectivity loss returns the widest dispersion, which is more plausibly read as variation in how respondents interpreted a normative statement than as disagreement about operating practice.

Table 4. Descriptive Statistics for the Fourteen Agreement Items

Item	Statement (abbreviated)	M	SD	Mdn	IQR	% agree	% disagree
Q5	Network interruptions create problems for digital production	3.84	1.17	4	3–5	68.9	13.1
Q6	Electricity interruptions disrupt production-data collection	3.62	1.17	4	3–4	67.2	18.0
Q7	Production data should continue to be recorded when offline	3.43	1.42	4	3–5	50.8	23.0
Q8	Initial cost of digital tools is a significant barrier	3.89	1.00	4	3–5	72.1	6.6
Q9	Hardware and maintenance costs limit adoption	3.79	1.17	4	3–5	67.2	11.5
Q10	Lack of skilled technical personnel is a barrier	3.92	1.16	4	3–5	70.5	13.1
Q11	Our factory can easily adopt minimal-training systems	3.11	1.08	3	2–4	42.6	29.5
Q12	Real-time sewing-line visibility would be valuable	3.61	1.08	4	3–4	60.7	13.1
Q13	Better WIP tracking would improve production management	3.74	1.15	4	3–5	68.9	18.0
Q14	Bundle dwell time at each station would identify delays	3.77	1.10	4	3–5	62.3	13.1
Q15	Continuous line status updating would benefit our factory	3.70	1.20	4	3–4	73.8	14.8
Q16	Such a system would be useful for our factory	3.90	0.93	4	3–5	72.1	4.9
Q17	An outage-tolerant system would be more suitable	3.85	1.11	4	3–5	70.5	9.8
Q18	Our factory would consider adopting it if affordable	3.84	1.13	4	3–5	73.8	13.1

Note. N = 61. Items scored 1 (strongly disagree) to 5 (strongly agree). “% agree” is the proportion selecting 4 or 5; “% disagree” the proportion selecting 1 or 2. Block composites: infrastructure (Q5, Q6) M = 3.73, SD = 1.01, $\alpha = .67$; cost (Q8, Q9) M = 3.84, SD = 0.90, $\alpha = .53$; perceived value (Q12–Q15) M = 3.70, SD = 0.80, $\alpha = .66$; acceptability (Q16–Q18) M = 3.86, SD = 0.90, $\alpha = .81$. Composites are descriptive averages, not validated scales. Q7 and Q17 are directionally leading (Section 3.2).




Figure 3. *Distribution of Agreement Across the Fourteen Items*

Note. N = 61. Bars are centred on the neutral category, so the horizontal position of each bar reflects the balance of agreement and disagreement rather than an absolute percentage. Q7 and Q17 are directionally leading items and are shown for completeness only.

4.3 Barrier Comparisons

The five barrier items cluster tightly. Network and electricity interruption, initial and recurrent cost, and skilled-personnel shortage all attract agreement between 67% and 72%, with the skills item returning the highest mean in the instrument. The pre-specified Friedman test found no reliable difference in perceived severity among them, and mean ranks fell within a narrow band from 2.85 to 3.17. Because the omnibus test was not significant, the planned Holm-corrected post-hoc comparisons were not performed. Two of the three pre-specified paired comparisons were likewise non-significant. Table 5 reports all four tests.

This pattern of null results is informative rather than disappointing, and was specified in advance as a test that could have told against the proposed design. Had electricity disruption been rated materially more severe than network disruption, the case for a connectivity-focused architecture would have weakened, since a gateway without power records nothing regardless of how its synchronisation logic is arranged. Had one barrier dominated, the sensible response would have been to attack it first. Neither happened. The design implication is that these constraints must be addressed jointly rather than sequentially, and that the resilience requirement should be framed around operational interruption generally rather than internet outages specifically.

The one comparison that did reach significance is the most consequential result in the study. The skills barrier was rated substantially higher than respondents' assessment of their own factories' capacity to absorb minimal-training systems ($W = 147$, $p < .001$, $r = .40$), with a median within-respondent difference of a full scale point; only 32.8% of respondents agreed with both propositions. This constrains a common design response. The obvious answer to a skills barrier is to reduce the training a system demands, yet the same people who rate that barrier highest decline to agree that minimal-training systems would be readily absorbed at their own factories. What they describe is therefore not solely a training-hours problem but an organisational absorption problem, in which the constraint lies in supervision, maintenance capability and workflow change. Minimising interface complexity is necessary but, on this evidence, not sufficient.

Table 5. *Pre-Specified Inferential Comparisons*

Comparison	Test	Statistic	p	Effect size
Severity ordering across five barrier items (Q5, Q6, Q8, Q9, Q10)	Friedman	$\chi^2(4) = 2.36$	.670	W = .010
Network vs. electricity disruption (Q5 vs. Q6)	Wilcoxon	W = 192.0	.164	r = .17
Initial vs. maintenance cost (Q8 vs. Q9)	Wilcoxon	W = 222.0	.415	r = .10
Skills barrier vs. minimal-training ease (Q10 vs. Q11)	Wilcoxon	W = 147.0	< .001	r = .40

Note. N = 61; two-tailed throughout. Kendall's W is reported for the Friedman test and rank-based r for the Wilcoxon tests, computed over non-tied difference pairs (n = 32, 32 and 40 respectively). Holm-corrected post-hoc comparisons were planned but not performed, because the omnibus test was not significant. All four comparisons were specified before the results were inspected.

4.4 Perceived Value and System Acceptability

The four value items returned agreement between 60.7% and 73.8%, with a median of 4 throughout. Continuous status updating drew the highest agreement, followed by work-in-progress tracking and bundle dwell-time visibility; real-time line visibility, the most abstract formulation and the capability most directly targeted by the architecture, drew the lowest. One reading is that practitioners respond more readily to a described operational benefit than to a general technological capability, which would suggest the case for such a system is better made in terms of delay identification than visibility as such. The ordering was not tested and lies within the range sampling variation alone could produce, so this is an observation rather than a finding.

The three acceptability items were the most positively rated block, with agreement between 70.5% and 73.8% on perceived usefulness, outage tolerance and conditional adoption intent. These figures should not be read as endorsement of the architecture specified in Section 6. Respondents saw a short general description of a production-monitoring system with offline capability, not the design in Sections 6.3 to 6.5, and the adoption-intent item is explicitly conditional on affordability and compatibility, neither of which this study establishes. What the block supports is the weaker claim that a system of this general description is not perceived as unwanted.

4.5 Associations Between Measures and Reliability

Internal consistency was modest to acceptable: acceptability returned the strongest value across its three items, the infrastructure pair and perceived value fell just below or around the conventional .70 threshold, and the cost pair fell well below it. The cost composite should therefore be read as the average of two related but distinct concerns rather than as a measure of one underlying construct, and the two-item blocks would not support a reflective measurement model in any case. All composites are reported as descriptive averages.

Perceived monitoring value and system acceptability are positively associated (Spearman's $\rho = .535$, $p < .001$), as is cost concern with acceptability, more strongly still ($\rho = .631$, $p < .001$); both are shown in Figure 4. The skills-barrier item is also associated with acceptability ($\rho = .406$, $p = .001$), while infrastructure concern is associated only weakly and marginally ($\rho = .284$, $p = .026$). Because Q17 is a leading item, the value-acceptability association was recomputed with acceptability restricted to the two non-leading items; the coefficient is essentially unchanged ($\rho = .543$, $p < .001$), so the conclusion is insensitive to its inclusion. The two leading items correlate at only $\rho = .310$ ($p = .015$), low enough to suggest they are not simply capturing a uniform tendency to agree.

No causal reading is available from any of these coefficients. They are cross-sectional associations among items answered by the same respondents in a single sitting, and a respondent generally favourable towards digital production monitoring would score high across several items at once, which alone would generate correlations of

this magnitude. The cost-acceptability association illustrates the difficulty plainly: it is at least as consistent with general engagement with the topic as with any specific mechanism linking cost sensitivity to willingness to adopt.




Figure 4. Association Between Block Composites and System Acceptability

Note. $N = 61$. Points are jittered slightly to reveal overlapping ordinal values; lines are least-squares fits shown to indicate direction only. Both associations are cross-sectional and admit no causal reading (Section 4.6).

4.6 Open-Ended Findings

Seventeen respondents answered the open-ended item, of which sixteen were codeable. Table 6 reports code frequencies. Two respondents submitted near-identical wording four minutes apart while differing on role, tenure and factory type; these were retained separately, but the possibility that they reflect a single discussed answer should be borne in mind.

Skills and digital literacy dominate, appearing in twelve of sixteen responses and described specifically as a shop-floor and middle-management problem rather than a shortage of specialist engineers. Resistance to change appears in eight, frequently paired with skills in the same sentence and anchored to entrenched paper-based record-keeping. Cost appears in six, and physical or power infrastructure in only three.

That last figure complicates the closed-item picture in a way worth preserving rather than smoothing over. Where the Likert items place infrastructure, cost and skills at comparable severity, the unprompted responses weight skills and organisational resistance far more heavily. Two readings are available and the data cannot separate them: either the closed items overstate infrastructure concern by prompting for it, or respondents treat infrastructure as a background condition rather than an adoption barrier they would volunteer. Either way, the finding tempers any suggestion that connectivity resilience is what practitioners themselves regard as the central obstacle, and Section 5.5 addresses that tension rather than resolving it in the architecture's favour.

Three concerns emerged that the closed items did not anticipate: the difficulty of training a large workforce without disrupting tight schedules, which is a delivery constraint distinct from the skills gap; trust in the data, with one respondent naming fear of surveillance explicitly; and integration with older, non-instrumented machinery. All three carry into the design requirements in Section 6.1.

Table 6. Thematic Codes in the Open-Ended Responses

Code	Description	Count
Skills and digital literacy	Shop-floor and middle-management capability to operate and rely on the system	12

Code	Description	Count
Resistance to change	Attachment to established manual and paper-based practice	8
Capital and operating cost	Upfront investment and ongoing cost of hardware and software	6
Physical and power infrastructure	Load shedding and inadequate plant-level infrastructure	3
Training delivery constraint	Difficulty releasing staff for training without disrupting schedules	2
Trust and surveillance concern	Doubt about data accuracy or fear of individual monitoring	2
Legacy equipment integration	Difficulty interfacing with older, non-instrumented machinery	1

Note. Sixteen codable responses of 17 received; one restated the question. Counts exceed 16 because responses could receive more than one code. The final four codes were added inductively and were not anticipated by the closed items. The coding frame appears in Appendix B.

5. Discussion

5.1 Interpretation of the Practitioner Evidence

Three claims are supported by these data. Practitioners in the surveyed sample regard the shortage of technical skills as the most severe single constraint, and as more severe than their factories' capacity to absorb low-training systems would suggest is manageable. They regard infrastructure, cost and skill barriers as comparably serious, with no reliable ordering among them. And they regard a low-cost, outage-tolerant production-monitoring system as useful in principle, conditional on affordability and compatibility.

Three claims are not supported and are not made. Nothing here establishes how often Bangladeshi factories actually lose connectivity or power; that would require operational logs, not agreement scales. Nothing establishes that perceived value causes willingness to adopt. Nothing generalises to Bangladeshi RMG factories as a population, since the sample is purposive, respondents may share employers, and the profile is weighted towards shorter tenures.

5.2 Comparison With Previous Research

The survey aligns with the barrier literature on substance while differing on level of observation. Cost and skills emerge as the most strongly endorsed constraints, which is what Salman et al. (2024) and Mim et al. (2024) report from expert panels and firm-level respondents. The convergence matters because the present respondents occupy different organisational positions: constraints identified at board and expert level are recognised, in similar rank order, by people working at or near the line.

Where this study departs from its predecessors is in the null results. Salman et al. (2024) and Karmaker et al. (2023) both produce ranked or structured barrier hierarchies, which is what DEMATEL and interpretive structural modelling are designed to yield. The Friedman test reported here finds no reliable severity ordering. The two findings are not in conflict, because the methods answer different questions: DEMATEL elicits expert judgements about influence and interdependence among barriers, whereas the present test asks whether practitioners rate them at different levels of severity. But the difference is consequential for design. A ranked hierarchy invites a sequenced intervention that tackles the dominant barrier first. The present evidence does not support that inference, and points instead towards a design that must satisfy cost, skill and infrastructure constraints simultaneously, since relaxing any one of them will not by itself make a system adoptable.

A second departure concerns what practitioners volunteer. Barrier studies working from fixed instruments cannot detect concerns their instruments do not name. The open responses surfaced training-delivery scheduling, data

trust and surveillance, and legacy machine integration, none of which appears prominently in the Bangladeshi barrier literature reviewed in Section 2.3, and two of which translate directly into design requirements.

5.3 What BGMEA's Current Response Addresses, and What Remains

The DigiProd Pass pilot, the Digital Factory Passport and the AWARE agreement all produce verifiable records for external audiences. Set against the survey findings, three observations follow.

The first concerns an assumption. All three programmes take for granted that a factory can already generate, timestamp and transmit data reliably, which implies a working baseline of infrastructure and connectivity. A compliance system with no offline-tolerant capture layer will work best at the factories needing help least; smaller and mid-sized suppliers, who make up most of BGMEA's membership, are the ones likely to struggle. The DBS Group case in Section 2.3 illustrates the resource asymmetry: a factory able to commit six figures to technology does not match the profile the barrier data describe.

The second is structural. These are traceability and reporting systems, not production-monitoring systems. They record what has happened, for verification by an outside party; they do not tell a production manager what is happening on the line at present. That is a legitimate goal, and nothing here suggests BGMEA has chosen badly on its own terms, given that European market access is not optional for most exporters.

The third is that the two are not in competition for the same data. The records a sewing line generates are, in principle, inputs to the compliance record rather than an alternative to it. A factory that can attest what happened at each station has the raw material for a Digital Product Passport claim, which suggests these investments are complementary rather than rival. The criticism is therefore narrow: not that BGMEA has invested in the wrong thing, but that the current scoping leaves production-floor barriers unaddressed and risks widening the distance between compliance-ready exporters and the rest.

5.4 Evidence-to-Design Implications

Four findings shape the architecture. The absence of a severity ordering means the design must satisfy cost, skill and infrastructure constraints jointly. The comparability of network and electricity disruption widens the resilience requirement from connectivity to operational interruption generally, which is why local power provision is treated as in scope. The gap between the skills barrier and the claimed ease of adopting minimal-training systems indicates that interface simplicity alone will not resolve what respondents describe; the design responds by removing configuration from the floor entirely rather than making it easier. And any system recording operator-attributable timings will be read by some as a surveillance instrument, which is a reason to draw the privacy boundary deliberately. Table 7 sets out the full chain from evidence through interpretation to architectural response.

5.5 Why Connectivity Resilience Matters, and What Qualifies It

The case for treating connectivity resilience as a basic design requirement rests on three independent strands, and the survey supplies the weakest of them. The strongest is documentary: outages are a recorded operating condition of Bangladeshi manufacturing, reported at roughly 30% of operating time even at a technologically advanced exporter (Hosain, 2025). The second is technical: the edge-blockchain literature consistently assumes at least periodic connectivity, so a system deployed where that assumption fails has no prior art to rely on (Nguyen et al., 2024; Faridi et al., 2023). The third is perceptual, and is the survey's contribution: roughly two-thirds of respondents recognise interruption as a problem for digital production activity.

Two qualifications belong here. Infrastructure appeared in only three of sixteen open-ended responses, well below skills and resistance to change, so practitioners asked unprompted to name their biggest difficulty do not reach for connectivity first. And the survey cannot establish outage frequency at all. The honest position is that connectivity resilience is justified primarily by documented operating conditions and by the absence of prior art addressing them, with practitioner perception offering corroboration rather than foundation. Presenting it as a practitioner-driven requirement would overstate what the evidence carries.

5.6 Contribution and Boundaries

What this study contributes is a documented chain from practitioner-reported constraint to architectural feature, in a setting where prior work has established the constraints and prior work has established the technologies, but no prior work has connected them. What it does not contribute is evidence that the resulting architecture works.

6. Proposed Connectivity-Resilient Architecture

6.1 Problem and Design Constraints

Most digital twin and blockchain traceability proposals for garment manufacturing assume reliable internet, a genuine IT budget, and machines with sensors already fitted. Small and mid-sized Bangladeshi factories have none of these. No validated architecture currently maintains a live digital representation together with verifiable production records for a sewing line under those conditions. Table 7 converts the evidence assembled in Sections 2, 4 and 5 into requirements, recording for each the interpretive step linking evidence to design.

Table 7. *Evidentiary Chain From Practitioner Evidence to Architectural Response*

Evidence	Source	Interpretation	Design requirement	Architectural response
No reliable severity ordering among five barrier items ($p = .670$)	Section 4.3	Constraints bind jointly; relaxing one will not make a system adoptable	Cost, skill and interruption requirements must be satisfied together	Single design evaluated against all targets in Table 8, not sequenced
Network and electricity interruption rated comparably disruptive	Section 4.3; Hosain (2025)	Resilience must cover operational interruption, not internet loss alone	Capture survives interruption of any duration	Durable local event queue; gateway-scoped power provision
Cost is a barrier, with no distinction between capital and recurrent	Section 4.3	Cheap to buy but costly to run addresses only half the concern	Low total cost of ownership	Commodity single-board hardware; open-source stack; no per-line cloud tenancy
Skills barrier rated above capacity to absorb low-training systems ($r = .40$)	Section 4.3	Absorption problem, not only a training-hours problem	No configuration, tuning or maintenance by floor staff	Single-scan interface; pre-configured gateway; heavy roles hosted centrally
Monitoring value endorsed at 60.7–73.8% agreement	Section 4.4	Line-level visibility is wanted, framed operationally	Bundle position and station dwell time must be visible in near real time	Line-level digital shadow updated from station scans
Legacy machines are not instrumented	Section 4.6; Section 2.1	Retrofit-dependent designs are not deployable here	Capture independent of machine-level sensors	Bundle-level RFID or QR identification

Evidence	Source	Interpretation	Design requirement	Architectural response
Surveillance concern raised unprompted	Section 4.6	Operator-attributable timing is a trust risk, not only a privacy one	Individual performance data must not leave the factory	Timings aggregated at the line before transmission (Section 6.5)
Records must be relied on across an organisational boundary	Sections 2.4, 5.3	Factory-controlled databases cannot underwrite external claims	Tamper-evidence and cryptographic identity for the writer	Permissioned ledger anchoring of verified event batches

Note. Requirements derived from survey evidence describe constraints reported by respondents, not measured operating conditions. The interpretation column records the inferential step between evidence and requirement, so that each step can be contested independently.

Expressing these requirements as numerical targets makes the design testable but risks presenting assumptions as findings, so Table 8 records each target with its derivation and evidential status. Two points deserve emphasis in the text. The cost ceiling is an indicative bill-of-materials estimate summing to roughly USD 220 to 350 per line at retail prices; no procurement quotation has been obtained, and it is emphatically not survey-derived, since the cost items establish that cost is perceived as a barrier, not what threshold would count as affordable. The throughput target is derived arithmetically from line rate, and at roughly an order of magnitude below reported single-board Fabric throughput (Zulkarnain et al., 2025) it suggests throughput is unlikely to be the binding constraint. The latency tolerance, by contrast, is derived from nothing and is labelled an assumption.

Table 8. *Design Targets, Derivation and Evidential Status*

Target	Value	Derivation	Status
Cost ceiling per line	≤ USD 300	Indicative bill of materials: single-board computer, power supply, enclosure and storage, small gateway-scoped UPS, two to four handheld readers, local networking. Components sum to approximately USD 220 to 350 at retail.	Design target. No procurement quotation obtained; excludes duty, installation and spares. Not survey-derived.
Sustained throughput	≥ 5 events per second	A 15 to 30 station line at roughly 30 bundles per hour, with entry and exit scans per station, yields about 0.5 events per second in steady state; a tenfold burst allowance gives 5.	Derived arithmetically. An order of magnitude below reported single-board Fabric throughput (Zulkarnain et al., 2025).
Online latency	≤ 5 seconds, scan to local ledger write	Judgement that floor operations do not require the sub-second confirmation expected of financial applications.	Assumption. Not derived and not empirically grounded.

Target	Value	Derivation	Status
Interruption tolerance	No data loss at any interruption duration; visibility may be delayed	Follows from the requirement in Table 7; bounded in practice by local storage capacity.	Requirement. Storage bound not yet quantified.
Interface complexity	Single-scan operation; no configuration by floor staff	Follows from the Q10 and Q11 contrast reported in Section 4.3.	Requirement supported by survey evidence on constraints, not by usability testing.
Privacy boundary	Commercial terms invisible to floor operators; operator-attributable timing not exported	Follows from sector data-sovereignty practice and the surveillance concern in Section 4.6.	Requirement. Enforcement mechanism specified (Section 6.5); not tested.

Note. No value in this table is an empirical finding of the present study. The survey establishes that cost is perceived as a barrier; it does not establish what threshold counts as affordable. The status column records what each figure is, so that later work can test the derived targets and revise the assumed ones.

6.2 Candidate Architectures

Three patterns were compared against those requirements, each trading cost, throughput and interruption resilience differently.

Table 9. *Candidate Architecture Comparison*

Architecture	Description	Key limitation
A. Cloud-anchored Fabric	Edge devices act as thin RFID or QR relays. The full Hyperledger Fabric network, including peers, orderer and certificate authorities, runs centrally and is reachable only over the internet.	There is no offline path. The system stops recording the moment connectivity drops, and cloud hosting costs are unbounded against a USD 300 ceiling.
B. Edge-gateway hybrid (recommended)	One Raspberry Pi-class gateway per line takes RFID or QR scans, holds local state, buffers events on disk, and periodically batches verified transactions to a centrally hosted Fabric network using private data collections and oracle-style verification.	Resource budgeting is delicate. Full Fabric roles are heavier than the peer-only role tested in published benchmarks, so the gateway has to remain a lightweight client with local buffering.

Architecture	Description	Key limitation
C. Off-chain hash-anchored	The digital representation runs on a central or on-premises server. The ledger stores only hashes of batched records, with edge devices relaying raw data.	The weakest link to verifiability at the edge. It falls back to Architecture A's connectivity dependence for the underlying capture.

Note. Compared against the requirements in Table 7 and the targets in Table 8. Architecture C is retained as a fallback should field testing show the buffering and verification logic to be too heavy for the target hardware.

Architecture B fits best. It is the only pattern that keeps recording during an interruption; it meets the cost ceiling on consumer single-board hardware; and, because the heavy Fabric roles remain off the gateway, it should stay inside the CPU headroom that independent benchmarking suggests exists on this class of machine (Zulkarnain et al., 2025; Honar Pajoo et al., 2021). Architecture C is retained as a fallback should field testing show the buffering and verification logic too heavy for the target hardware.

6.3 Recommended Design: Edge-Gateway Hybrid

Each sewing station is fitted with a low-cost RFID or QR reader that registers bundles as they pass, and each line with a Raspberry Pi-class gateway that collects scan data and updates the local representation: bundle location, dwell time at each station, and defect alerts raised during inspection. Before an event is queued for the ledger the gateway runs a local plausibility check, an oracle-style step intended to catch obviously impossible entries, such as a bundle clearing a station faster than the line's minimum cycle time. Verified batches are then sent to a centrally hosted, lightweight Fabric network. Figure 5 shows the layered arrangement and marks the functions that remain available to the supervisor dashboard when the network is unavailable.

The reader choice is deliberately left open between RFID and printed QR, because the two trade differently against the constraints in Table 7. Printed QR costs almost nothing per bundle and can be produced on existing office equipment, but requires line-of-sight and a deliberate scan action. Passive RFID removes that action and tolerates dust and handling better, but raises per-bundle and per-reader cost against a tight ceiling. Under the cost constraint the QR variant is the more defensible starting point, with RFID as the upgrade path once per-line cost is established in practice.




Figure 5. *Proposed Edge-Gateway Architecture*

Note. Heavy Hyperledger Fabric roles remain centrally hosted; the per-line gateway runs only a lightweight client with local buffering. The local path indicates the functions that remain available to the supervisor dashboard during an interruption, when the ledger write is deferred. At this level of automation the digital representation is a digital shadow, as Section 6.6 sets out.

6.4 Connectivity-Resilience Mechanism

The one real departure from a conventional Fabric deployment is the gateway's offline-first event log. Every scan is written to a durable local queue the instant it occurs, whether or not the internet is up. Each event carries an idempotent identifier built from factory, line, bundle, operation type and a sequence number, so events replayed after reconnection are either accepted once or recognised as duplicates and discarded. Because a sewing line follows a fixed processing sequence, a bundle must pass station N minus one before reaching station N, which allows conflicting sequences emerging after a long interruption to be flagged for human intervention instead of merged automatically, a common failure of naive snapshot-replay and a frequent source of silent corruption. On reconnection the gateway replays its buffered queue in order; where the local view of ledger state has become outdated, a chunked, Merkle-verified catch-up procedure should bring it current without resending the entire history. Figure 6 sets out the three connectivity states.

The mechanism introduces costs that should be stated alongside its benefits. Local buffering means the gateway holds unreplicated production data for the duration of an interruption, so a gateway lost or damaged during an outage takes its unsynchronised events with it. The queue is bounded by local storage, and that bound has not

been quantified. And deferred ledger writes mean tamper-evidence applies from the moment of synchronisation, not of capture: events resident only in the local queue are protected by ordinary filesystem controls rather than by the ledger. For an internal monitoring record this is an acceptable trade; for a record intended to substantiate an external compliance claim it may not be.




Figure 6. *Offline-First Store-and-Forward Synchronisation Logic*

Note. Every scan is durably queued locally before any network transmission, so panels A and B differ only after the local state update. Sequence-numbered idempotent identifiers make post-outage replay safe; conflicts that cannot be resolved by sequence are referred for human review rather than merged automatically.

6.5 Privacy and Data Governance

Buyer identity, unit pricing and other commercially sensitive fields are segregated into a private data collection visible only on the buyer-facing channel, following the data-sovereignty logic to which the AWARE partnership has already committed (Fibre2Fashion, 2026). Enforcing confidentiality through platform configuration rather than application logic removes a class of error: a field held in a private data collection is invisible to an unauthorised peer even if the application layer is misconfigured.

The surveillance concern raised in the open responses requires a boundary drawn in the other direction, and is the harder problem. A system recording which bundle passed which station and when necessarily produces operator-attributable timing data. The design response is to hold those timings at the line, aggregate before transmission, and expose station-level rather than person-level statistics on the buyer-facing channel, so the record supports delay diagnosis without exporting individual performance data. This narrows exposure but does not eliminate it, since a supervisor with dashboard access can still observe individual pace. That residual is a governance question rather than a technical one, and should be settled with worker representatives before deployment.

6.6 Digital-Shadow Scope

Applying the taxonomy from Section 2.2 gives an unambiguous answer. The physical-to-digital flow is automated: scans propagate to the local representation and onward to the ledger without manual transcription. The digital-to-physical flow is not. Nothing here writes back to the sewing line; a delay flagged by the system reaches the floor through a supervisor reading a dashboard. On Kritzing et al.'s (2018) criteria the system as specified is a digital shadow, not a digital twin.

Retaining the digital twin framing is defensible provided the scope is stated rather than implied. The shadow is the layer on which a twin would be built: instrumentation, state representation and a verified event history are the prerequisites for simulation, for what-if analysis on line balancing, and eventually for automated feedback into scheduling. The corollary is that this design does not optimise, does not predict and does not close a control loop, and any of those would require both additional modelling and a degree of process automation the target factories do not have.

6.7 Status and Limitations of the Proposed Architecture

No prototype exists. The throughput, latency and CPU figures referenced above come from independent published benchmarks (Zulkarnain et al., 2025; Lee et al., 2023; Honar Pajoo et al., 2021), not from measurements on this configuration, and full Fabric roles are heavier than the roles those benchmarks tested, so real headroom is probably tighter than the published numbers imply. The survey supports the requirements but does not validate the design: respondents were not shown this architecture, so it would be a misuse of the data to read the adoption-intent figure as endorsement of the proposal.

The use of Hyperledger Fabric is not the novel element. Fabric-based RMG supply-chain concepts already exist (Niloy et al., 2024), Fabric-based apparel traceability has been built and benchmarked (Faridi et al., 2023), and private data collections are a standard platform feature. What Section 2.6 found genuinely unclaimed is narrower: offline-first, store-and-forward synchronisation with sequence-numbered conflict handling and oracle-style plausibility verification, on low-cost hardware, for RMG production data. Future work should lead with that combination and not with Fabric. Several components also remain underdetermined: the plausibility check is described by purpose rather than rule set, the catch-up procedure is named rather than designed, the conflict-flagging threshold is not fixed, the storage bound is not quantified, and the power provision is costed but not engineered.

7. Recommendations

7.1 Factory Owners and Management

The finding in Section 4.3 argues against treating adoption as a procurement decision followed by a training event. Because the binding constraint includes supervision and maintenance capability rather than initial instruction alone, management would do better to identify, before committing to any system, who will own it once the vendor has left, and to scope a first deployment to a single line where that ownership can be established before scaling. Two respondents separately identified the difficulty of releasing staff for training without disrupting delivery schedules, which is a planning problem for management rather than for the technology.

7.2 BGMEA and BKMEA

Given Section 5.3, the trade bodies are well placed to extend their Digital Product Passport and Digital Factory Passport work with a low-cost, interruption-tolerant production layer, or to commission a parallel initiative aimed at production monitoring. Because line-level records are inputs to a compliance claim rather than an alternative to one, the two can be scoped as a single investment rather than two demands on the same budget. Cluster and shared-service models, in which several small factories share infrastructure and expertise, would cut the effective per-firm cost and sit naturally with the associations' convening role.

7.3 Government and Regulators

Low-interest technology loans, tax relief on digital equipment and matching grants aimed at SME garment firms attack the cost barrier directly, and the finding that respondents distinguish little between capital and recurrent cost suggests instruments covering only acquisition will be less effective than those covering the operating period. Continued investment in reliable electricity and affordable industrial broadband would reduce the infrastructure barrier, on which both this study's respondents and the sector's own compliance rollout quietly depend.

7.4 Technology and Blockchain Providers

Vendors should build simple, affordable, modular tools rather than heavy enterprise systems. Factories run older, non-digital sewing machines, so new software must work alongside them; one respondent named legacy integration explicitly. Vendors need local presence, payment structures that do not require the full cost at

acquisition, and interruption tolerance designed in rather than added later. That configuration should not be performed by floor staff is a product decision, not a deployment detail.

7.5 Educational Institutions

Teaching industrial data analytics, sensor systems and the fundamentals of distributed ledgers within existing programmes would send graduates into the sector with the skills the barrier data identify. Short courses aimed at incumbent supervisors, rather than only at students, would address the part of the constraint graduate supply alone cannot reach, since the gap the open responses describe sits at supervisory and middle-management level.

8. Limitations and Future Research

8.1 Limitations

The survey carries the limitations of a purposive cross-sectional design with a modest sample. Sixty-one responses drawn without a sampling frame cannot support generalisation to Bangladeshi RMG factories as a population; respondents may share employers; and the profile is weighted towards earlier-career practitioners. Distribution through professional contacts makes self-selection likely, and practitioners already interested in digitalization are plausibly over-represented, which would inflate the acceptability figures. All items are self-reported and were collected in a single sitting, so common-source variance offers a plausible partial account of the associations in Section 4.5. Two blocks return alpha below the conventional threshold and two contain only two items, so the composites are descriptive rather than validated scales. The factory-size item was unusable, and cell sizes preclude subgroup comparison. No causal inference is available from any analysis reported here. The architecture component is limited more fundamentally: there is no prototype, no deployment, no measured throughput or latency, no controlled outage testing, no security evaluation, no usability testing, and no validated cost.

8.2 Future Research

The immediate priority is to build the Section 6 design and evaluate it against the targets in Table 8. A first phase would assemble the gateway on Raspberry Pi-class hardware with RFID or QR capture and establish the actual per-line bill of materials. A second would benchmark Fabric throughput and scan-to-write latency on that configuration under a realistic line event profile, testing whether the published headroom survives the addition of buffering and verification logic. A third would test the resilience mechanism specifically, through controlled network-outage and power-failure experiments: injecting interruptions of increasing duration, verifying that no events are lost, measuring replay time, quantifying the storage bound, and deliberately constructing sequence conflicts to confirm they are flagged rather than silently merged. A fourth would submit the design to security and privacy evaluation, covering the local queue's exposure during an outage and the operator-attributable data of Section 6.5.

Beyond the prototype, a pilot deployment on a real sewing line, ideally at a factory already in a compliance pilot, would allow production-monitoring and DPP requirements to be evaluated together. Logging connectivity and power availability at the gateway would replace the perceptual evidence of Section 4.3 with measurement. Operator usability evaluation would test the single-scan assumption, and a longitudinal study would establish whether the barrier profile reported here is stable or a function of current maturity. Replication in comparable apparel-exporting economies would establish how far these constraints are specific to Bangladesh.

9. Conclusion

This study asked which operating constraints RMG practitioners recognise, whether the sector's current blockchain investment is directed at them, and what an architecture built around them would look like.

Respondents identified skilled-personnel shortage as the most severe single constraint, but a pre-specified omnibus test found no reliable severity ordering among infrastructure, cost and skill barriers, which argues for treating them as a joint constraint set rather than a ranked list. The most informative result is the gap between the severity attributed to the skills barrier and respondents' reluctance to agree that minimal-training systems would be easily absorbed, indicating an organisational absorption problem alongside a training one. The unprompted

responses weight skills and resistance to change considerably more heavily than infrastructure, which qualifies rather than confirms the connectivity premise on which the architecture rests.

BGMEA is certainly moving, and its three agreements represent real institutional momentum. That momentum is pointed at European compliance rather than at production-floor barriers, and left as it is, it risks widening the distance between compliance-ready exporters and the smaller factories that cannot reach the baseline these platforms assume. Since line-level records are inputs to a compliance claim rather than a competing use of the same budget, the two could be scoped together.

The architecture proposed here has not been validated: a Fabric-anchored, edge-gateway production-monitoring system with an indicative target cost below USD 300 per line, and offline-first capture designed to keep recording through interruptions that are a normal condition of Bangladeshi manufacturing rather than a rare failure. At the level specified it is a digital shadow, and a foundation for a twin rather than a twin itself. Its claim to novelty is narrow: it lies in the offline-first, sequence-aware synchronisation pattern for RMG production data on constrained hardware, and not in the use of Hyperledger Fabric.

Acknowledgments

The authors thank the practitioners who gave their time to complete the survey instrument. No external funding was received for this work.

Appendix A. Survey Instrument

Q1. Current position (categorical). Q2. Years worked in the RMG sector (ordinal). Q3. Factory type (categorical). Q4. Factory size (categorical; excluded from all analysis, see Section 3.3).

Items Q5 to Q18 were scored on a five-point agreement scale (1 = strongly disagree, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree).

Q5. Internet or network interruptions create problems for digital production activities. Q6. Electricity interruptions can disrupt digital production-data collection. Q7. Production data should continue to be recorded when internet connectivity is unavailable. Q8. The initial cost of advanced digital technologies is a significant barrier. Q9. Hardware and maintenance costs limit adoption. Q10. A lack of skilled technical personnel is a barrier. Q11. Our factory can easily adopt systems requiring minimal staff training. Q12. Real-time visibility of sewing-line production would be valuable. Q13. Better work-in-progress tracking would improve production management. Q14. Knowing how long a bundle stays at each station would help identify delays. Q15. Our factory would benefit from continuously updated sewing-line status. Q16. A system of the kind described would be useful for our factory. Q17. A system that continues operating during connectivity interruptions would be more suitable for our factory than one requiring continuous internet access. Q18. Our factory would consider adopting such a system if it were affordable and compatible with existing equipment.

Q19 (open-ended, optional). What do you think would be the biggest practical difficulty for Bangladeshi factories in adopting a digital production-monitoring system?

Appendix B. Coding Frame for the Open-Ended Item

Codes derived deductively from the barrier literature: capital and operating cost; skills and digital literacy; physical and power infrastructure. Codes added inductively from the responses: resistance to change; training delivery constraint; trust and surveillance concern; legacy equipment integration; non-substantive. Multiple codes were applied to a single response where the response raised more than one concern, which is why the counts in Table 6 exceed the number of coded responses. Of 17 responses received, 16 were codeable and one restated the question.

Appendix C. Analysis and Reproducibility

All descriptive and inferential statistics reported in Section 4 were computed in Python (pandas, scipy.stats) directly from the 61-response cleaned dataset. Wilcoxon signed-rank effect sizes were computed as $r = |Z| / \sqrt{2 \times n_{\text{nonzero}}}$, where n_{nonzero} is the number of non-tied difference pairs. Kendall's W for the Friedman test was

computed as $W = \text{chi-squared} / [n(k - 1)]$. Cronbach's alpha was computed from item and total-score variances. The cleaned dataset and analysis script are available from the authors on request.

References

- Agrawal, T. K., Kumar, V., Pal, R., Wang, L., & Chen, Y. (2021). Blockchain-based framework for supply chain traceability: A case example of textile and clothing industry. *Computers & Industrial Engineering*, 154, 107109. <https://doi.org/10.1016/j.cie.2021.107109>
- Ahmed, W. A. H., & MacCarthy, B. L. (2021). Blockchain-enabled supply chain traceability in the textile and apparel supply chain: A case study of the fiber producer Lenzing. *Sustainability*, 13(19), 10496. <https://doi.org/10.3390/su131910496>
- Al Amin, M., Baldacci, R., & Kerbache, L. (2025). Blockchain-based green supply chain management framework for sustainable practices in Bangladeshi RMG industries. *Discover Sustainability*, 6, 508. <https://doi.org/10.1007/s43621-025-01196-8>
- Arnob, S. H., Atik, R. I., Ahmed, M. A., & Hassan, S. M. (2020). Lead time optimization of RMG sector supply chain in Bangladesh. *Supply Chain Insider*, 4(1). <https://supplychaininsider.org/ojs/index.php/home/article/view/35>
- Bangladesh Garment Manufacturers and Exporters Association. (2025, October 5). BGMEA signs strategic cooperation with Digital Bridge Partners to launch AI-enabled ESG compliance framework for RMG sector. https://www.bgmea.com.bd/page/BGMEA_Signs_Strategic_Cooperation_with_Digital_Bridge_Partners_to_Launch_AI-Enabled_ESG_Compliance_Framework_for_RMG_Sector
- Chaity, N., Islam, K., Masud, A., Islam, M., Rahman, M., & Kabir, S. (2025). Sustainable pathways to circular economy adoption: Insights from the RMG sector of Bangladesh. *Business Strategy and the Environment*, 34, 5857-5868. <https://doi.org/10.1002/bse.4277>
- DigiProd Pass. (2025, September). DigiProd Pass partners with BGMEA to equip Bangladesh's RMG industry for EU compliance. <https://digiprodpass.com/news/digiprod-pass-partners-with-bgmea>
- Etikan, I., Musa, S. A., & Alkassim, R. S. (2016). Comparison of convenience sampling and purposive sampling. *American Journal of Theoretical and Applied Statistics*, 5(1), 1-4. <https://doi.org/10.11648/j.ajtas.20160501.11>
- European Union. (2024). Regulation (EU) 2024/1781 of the European Parliament and of the Council establishing a framework for the setting of ecodesign requirements for sustainable products. *Official Journal of the European Union*.
- Fahim, M. F. H., Hossen, M. L., Sakib, N. A., Shakil, M. S., & Haque, M. (2024). Transforming Bangladesh's traditional textile and apparel industry: A sustainable future with green supply chain management. *Supply Chain Insider*, 14(1). <https://supplychaininsider.org/ojs/index.php/home/article/view/120>
- Faridi, M. S., Ali, S., & Wang, G. (2023). ChainApparel: A trustworthy blockchain and IoT-based traceability framework for apparel Industry 4.0. *Computers, Materials & Continua*, 77(2), 1837-1854.
- Fibre2Fashion. (2026, May). Bangladesh RMG sector to adopt blockchain-based transparency & DPP. <https://www.fibre2fashion.com/news/apparel-news/bangladesh-rmg-sector-to-adopt-blockchain-based-transparency-dpp-310229-newsdetails.htm>
- Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952-108971. <https://doi.org/10.1109/ACCESS.2020.2998358>
- Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85-113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
- Haque, A. Z., Shuvo, A. A., Hossain, M. M., & Ebrahim, J. A. (2022). A scenario of adopting blockchain technology in supply chain: A study of e-commerce in Bangladesh. *Supply Chain Insider*, 8(1). <https://supplychaininsider.org/ojs/index.php/home/article/view/20>

- Hasan, N., Ferdouse, N., Khan, T., Adiba, A. F., & Shams, K. A. I. (2022). Implementing blockchain in the supply chain of FMCG companies for better demand forecasting and logistics operation. *Supply Chain Insider*, 8(1). <https://supplychaininsider.org/ojs/index.php/home/article/view/18>
- Honar Pajoo, H., Rashid, M., Alam, F., & Demidenko, S. (2021). Hyperledger Fabric blockchain for securing the edge Internet of Things. *Sensors*, 21(2), 359. <https://doi.org/10.3390/s21020359>
- Hosain, R. (2025). Digital transformation in Bangladesh's RMG supply chains: A case study on DBS Group. *International Journal of Research and Innovation in Social Science*, 9(6), 1696-1703. <https://doi.org/10.47772/IJRISS.2025.906000133>
- Hou, L., & Sun, J. (2021). A security architecture for garment production based on blockchain technology and traditional RFID. In *Proceedings of the 2nd International Conference on Computing and Data Science (CDS 2021)*. Association for Computing Machinery.
- Hu, W., Zhang, T., Deng, X., Liu, Z., & Tan, J. (2021). Digital twin: A state-of-the-art review of its enabling technologies, applications and challenges. *Journal of Intelligent Manufacturing and Special Equipment*, 2(1), 1-34. <https://doi.org/10.1108/JIMSE-12-2020-0010>
- Jayalakshmi, I., Vasanthi, D., & Perumal, V. V. (2024). Digital twin contribution in integrated processes of fashion and textile supply chains. In *Illustrating digital innovations towards intelligent fashion* (pp. 573-599). Springer. https://doi.org/10.1007/978-3-031-71052-0_24
- Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29, 36-52. <https://doi.org/10.1016/j.cirpj.2020.02.002>
- Kamble, S. S., Gunasekaran, A., & Sharma, R. (2018). Analysis of the driving and dependence power of barriers to adopt Industry 4.0 in Indian manufacturing industry. *Computers in Industry*, 101, 107-119. <https://doi.org/10.1016/j.compind.2018.06.004>
- Karmaker, C. L., Al Aziz, R., Palit, T., & Bari, A. B. M. M. (2023). Analyzing supply chain risk factors in small and medium enterprises under a fuzzy environment: Implications towards sustainability for emerging economies. *Sustainable Operations and Computers*, 4, 1-12. <https://doi.org/10.1016/j.susoc.2022.10.005>
- Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016-1022. <https://doi.org/10.1016/j.ifacol.2018.08.474>
- Kumar, P., Bhamu, J., & Sangwan, K. S. (2021). Analysis of barriers to Industry 4.0 adoption in manufacturing organizations: An ISM approach. *Procedia CIRP*, 98, 85-90. <https://doi.org/10.1016/j.procir.2021.01.010>
- Lee, S., Kim, M., Lee, J., Hsu, R.-H., Kim, M.-S., & Quek, T. Q. S. (2023). Facing to latency of Hyperledger Fabric for blockchain-enabled IoT: Modeling and analysis. *IEEE Network*, 37(5), 48-54. <https://doi.org/10.1109/MNET.001.2200558>
- Luthra, S., & Mangla, S. K. (2018). Evaluating challenges to Industry 4.0 initiatives for supply chain sustainability in emerging economies. *Process Safety and Environmental Protection*, 117, 168-179. <https://doi.org/10.1016/j.psep.2018.04.018>
- Mim, I. Z., Rayhan, M. G. S., & Syduzzaman, M. (2024). Prospects and current scenario of industry 4.0 in Bangladeshi textile and apparel industry. *Heliyon*, 10(11), e32044. <https://doi.org/10.1016/j.heliyon.2024.e32044>
- Nguyen, T., Tran, N., Balador, A., & Loke, S. W. (2024). Exploring the integration of edge computing and blockchain into IoT systems. *Journal of Network and Computer Applications*, 226, 103884. <https://doi.org/10.1016/j.jnca.2024.103884>

- Niloy, S. J. A., Alam, M. I., Karim, S., & Arfi, N. (2024). Implementing blockchain technology in the supply chain of RMG industries in Bangladesh. In Proceedings of the 7th Bangladesh Conference on Industrial Engineering and Operations Management (IEOM 2024), Dhaka, Bangladesh.
- Pasi, B. N., Mahajan, S. K., & Rane, S. B. (2020). The current sustainability scenario of Industry 4.0 enabling technologies in Indian manufacturing industries. *International Journal of Productivity and Performance Management*, 70(5), 1017-1048. <https://doi.org/10.1108/IJPPM-04-2020-0196>
- Pathri, B. P., & Ganduri, K. V. (2025). From smart cities to manufacturing: A comprehensive study on digital twin applications and barriers. *Digital Twins and Applications*. <https://doi.org/10.1049/dgt2.70014>
- Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360-degree comparison. *IEEE Access*, 6, 3585-3593. <https://doi.org/10.1109/ACCESS.2018.2793265>
- Raj, A., Dwivedi, G., Sharma, A., Lopes de Sousa Jabbour, A. B., & Rajak, S. (2020). Barriers to the adoption of Industry 4.0 technologies in the manufacturing sector: An inter-country comparative perspective. *International Journal of Production Economics*, 224, 107546. <https://doi.org/10.1016/j.ijpe.2019.107546>
- Salman, S., Hasanat, S., Rahman, R., & Moon, M. (2024). Analyzing the key barriers of adopting Industry 4.0 in Bangladesh's ready-made garment industry: An emerging economy example. *International Journal of Industrial Engineering and Operations Management*, 6(3), 232-255. <https://doi.org/10.1108/IJIEOM-04-2023-0034>
- Sayem, A., Kumar, P., & Muhsin, M. A. (2022). Critical barriers to Industry 4.0 adoption in manufacturing organizations and their mitigation strategies. *Journal of Manufacturing and Materials Processing*, 6(6), 136. <https://doi.org/10.3390/jmmp6060136>
- Sharfuddin Ashbin, M., Ahnaf Hamim, M., Rehan, S. M., Didarul Islam, M., & Nazmul Alam Gazi, M. (2024). Social violations in the RMG sector of Bangladesh. *Supply Chain Insider*, 13. <https://doi.org/10.5281/zenodo.13865700>
- Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405-2415. <https://doi.org/10.1109/TII.2018.2873186>
- Textile Focus. (2026, May). BGMEA, AWARE sign deal to meet EU's 'digital product passport' for Bangladesh RMG. <https://textilefocus.com/bgmea-aware-sign-deal-to-meet-eus-digital-product-passport-for-bangladesh-rmg/>
- Zulkarnain, M. M., Ramli, N., & Nordin, A. N. (2025). Performance benchmarking of Hyperledger Fabric on heterogeneous hardware for IoT applications. *IIUM Engineering Journal*, 26(3), 156-170. <https://doi.org/10.31436/iiumej.v26i3.3610>