MERIDIAN RETAIL GROUP

DPIA: Customer Analytics Platform

Data Protection Impact Assessment

GDPR Article 35(7)

Generated: 2026-06-30

Executive Summary

MEDIUM	245%	27 / 11	6
RISK LEVEL	COMPLETION	QUESTIONS ANSWERED	MITIGATIONS

Assessment Type	Data Protection Impact Assessment
Template	Data Protection Impact Assessment v1.0
Status	APPROVED
Risk Score	48 / 100
Started	2025-09-01
Submitted	2025-10-15
Completed	2025-10-28
Due Date	2025-11-01
Linked Activity	Customer Analytics & Business Intelligence

Description

Data Protection Impact Assessment for Meridian Retail Group's customer analytics platform, covering behavioral analytics, purchase pattern analysis, and personalized recommendations powered by the Snowflake data warehouse.

GDPR Article 35(7) — Required DPIA Elements

- ✓ Systematic description of processing operations and purposes
- ✓ Assessment of necessity and proportionality
- ✓ Assessment of risks to rights and freedoms of data subjects
- ✓ Measures envisaged to address risks and demonstrate compliance

1. Processing Description

Describe the processing operation and its purpose

5 of 5 questions answered **Complete**

Q1. What personal data is being processed?

REQUIRED

The following categories of personal data are processed:

• Identifiers: Customer name, email address, phone number, loyalty program ID, cookie identifiers, IP addresses
• Behavioral data: Website browsing history (pages viewed, time on page, click paths), search queries, product views, cart additions/abandonments, purchase history (items, amounts, dates, payment method)
• Demographics: Age range (derived from self-reported birthday for loyalty program), language preference, country/region
• Location data: Shipping addresses, approximate location from IP (city-level), store visit frequency (loyalty card scans)
• Device/technical: Browser type, device type, operating system, screen resolution
• Preference data: Marketing consent status, communication channel preferences, product category interests

Q2. What is the purpose of the processing?

REQUIRED

Primary purpose: To analyze customer behavior and purchase patterns in order to improve product offerings, optimize the e-commerce experience, and generate business intelligence reports for merchandising, inventory planning, and marketing strategy.

Secondary purposes:
1. Personalized product recommendations on the website and in email campaigns (based on browsing and purchase history)
2. Customer segmentation for targeted marketing campaigns (identifying high-value customers, at-risk churners, seasonal buyers)
3. Conversion funnel optimization (identifying friction points in the checkout process)
4. Demand forecasting for inventory management across 3 EU distribution centers
5. A/B testing of website features and pricing strategies

Q3. What is the legal basis for processing?

REQUIRED

Legitimate interests (Art. 6(1)(f))

Q4. Who are the data subjects?

REQUIRED

["Customers","Website visitors"]

Q5. How many data subjects are affected (estimated)? REQUIRED

100,000 – 1,000,000

2. Necessity & Proportionality

Assess whether the processing is necessary and proportionate

2 of 2 questions answered **Complete**

Q1. Is the processing necessary to achieve the stated purpose? REQUIRED

Mix of non-sensitive and sensitive indicators

Q2. Have less intrusive alternatives been considered?

Retention periods by data category:

• Raw behavioral events (clicks, page views, searches): 2 years from collection — justified by the need to identify annual seasonal patterns (e.g. holiday shopping) and year-over-year comparisons. Automatically purged via Snowflake time-travel policy.

• Purchase history: 7 years — required for tax/accounting compliance under Dutch fiscal retention obligations (AWR Art. 52). Anonymized after 2 years for analytics; raw records retained in Customer Database for legal compliance only.

• Customer segments and scores: Refreshed monthly, previous versions retained for 6 months to enable model performance monitoring. Older versions permanently deleted.

• Aggregated analytics (KPIs, dashboards): Retained indefinitely — fully anonymized, no individual-level data.

• Cookie identifiers and IP addresses: 13 months maximum (ePrivacy alignment). Automatically expired.

Justification review: Retention periods were reviewed against EDPB guidelines on storage limitation (May 2020) and benchmarked against industry practice for EU e-commerce retailers.

3. Risk Assessment

Identify and assess risks to the rights and freedoms of data subjects

3 of 4 questions answered

Q1. Does the processing involve automated decision-making or profiling? REQUIRED

Highly beneficial — significantly more effective than alternatives

Q2. Does the processing involve special category data or criminal offence data? REQUIRED

Alternatives considered:

1. Fully anonymized analytics only (no pseudonymized individual-level data):
Rejected — anonymous aggregate data cannot support personalized product recommendations, individual customer journey analysis, or cohort-based segmentation. These capabilities drive an estimated 18% of online revenue through improved recommendations and targeted marketing.

2. First-party cookies and session-based analytics only (no cross-session tracking):
Rejected — would lose the ability to understand customer lifetime value, repeat purchase patterns, and long-term behavioral trends. Reduces forecast accuracy for inventory planning by an estimated 35%.

3. Consent-based processing instead of legitimate interests:
Considered but supplemented rather than replaced — analytics processing relies on legitimate interests (supported by completed LIA, August 2025), while marketing personalization uses consent obtained through Cookiebot. This hybrid approach provides the strongest legal footing while respecting data subject choice.

4. On-premise analytics (no cloud/US transfer):
Rejected — Snowflake's cloud-native architecture provides significantly better performance, cost efficiency, and s calability compared to on-premise alternatives. The supplementary measures and EU data residency mitigate transfer r isks adequately.

Conclusion: The chosen approach (pseudonymized analytics in Snowflake with legitimate interests basis + consent for marketing) represents the least intrusive effective option.

Q3. Does the processing involve systematic monitoring of a publicly accessible area?

Data minimization measures:

1. Pseudonymization at ETL: Customer identifiers (name, email, phone) are replaced with hashed loyalty IDs before data enters the analytics warehouse. The mapping table is stored separately in the Customer Database with strict access controls.

2. Field-level restrictions: The analytics warehouse schema excludes payment data (card numbers, billing addresses), exact date of birth, and postal codes more specific than the first 4 digits.

3. Purpose-limited views: Snowflake roles restrict analysts to purpose-specific views. The Product Team can only

access product interaction data; the Marketing Team can only access segment-level data, not individual browsing history.

4. Aggregation thresholds: Any analytics query that would return results for fewer than 30 individuals is automatically suppressed (k-anonymity enforcement via Snowflake row access policies).

5. Consent-gated collection: Behavioral tracking (cookies, browsing history) is only activated after the customer provides consent via Cookiebot. Non-consented visitors see only aggregate analytics (page view counts, no individual tracking).

6. Automated retention: Snowflake time-travel policies automatically purge raw behavioral data older than 2 years. No manual intervention required.

Q4. What measures are in place to mitigate identified risks? REQUIRED

Not yet answered

Risk Assessment Summary

MEDIUM
OVERALL RISK LEVEL

48.0
RISK SCORE

0
SCORED QUESTIONS

Risk Mitigations

6	3	3
TOTAL	COMPLETED	OUTSTANDING

Title	Status	Priority	Owner	Due Date	Evidence
Deploy consent preference center for analytics opt-out	VERIFIED	P1	Engineering	—	Yes
Implement purpose limitation controls in Snowflake	IMPLEMENTED	P1	Data Team	—	Yes
Add data minimization layer to analytics pipeline	PLANNED	P2	Engineering	2025-12-01	No
Complete Transfer Impact Assessment for Snowflake US processing	IN PROGRESS	P2	Legal / Privacy	2025-12-15	No
Implement automated data retention enforcement	IMPLEMENTED	P2	Data Team	—	Yes
Conduct annual DPIA review	IDENTIFIED	P3	DPO	2026-10-01	No

Mitigation Details

Deploy consent preference center for analytics opt-out

Integrate a customer-facing analytics opt-out toggle in the privacy center, connected to Cookiebot consent records and the Snowflake data pipeline. When a customer opts out, behavioral tracking stops within 1 hour and existing analytics data is excluded from future queries within 24 hours.

Implement purpose limitation controls in Snowflake

Configure Snowflake role-based access controls to enforce purpose-specific views: Product Team sees only product interaction data, Marketing Team sees only segment-level data. Individual browsing history is not accessible to any role except the Data Team lead (for debugging purposes only).

Add data minimization layer to analytics pipeline

Enhance the ETL pipeline to apply k-anonymity enforcement (minimum group size of 30) and remove fields not required for analytics purposes before loading into Snowflake. Implement field-level encryption for the pseudonymization mapping table.

Complete Transfer Impact Assessment for Snowflake US processing

Conduct and document a comprehensive TIA for the EU→US data transfer to Snowflake, assessing the legal framework, government access risks, and supplementary measures in light of Schrems II and the EU-US Data Privacy Framework.

Implement automated data retention enforcement

Configure Snowflake time-travel and data retention policies to automatically purge raw behavioral data older than 2 years. Aggregated metrics are retained separately. Implement monitoring alerts for retention policy failures.

Conduct annual DPIA review

Schedule and conduct an annual review of this DPIA to reassess risks in light of any changes to the processing, technology, legal framework, or organizational context. Review to be led by the DPO with input from IT Security and Legal.

Approval History

Level	Approver	Status	Date	Comments
Level 1	Maria Torres	APPROVED	2025-10-25	As DPO, I have reviewed this DPIA in full. The processing is necessary and proportionate. All high-priority mitigations have been implemented and verified. The residual risk level of MEDIUM is acceptable given the safeguards in place. Approved subject to completion of the outstanding TIA for Snowflake and the annual review.
Level 2	James Mitchell	APPROVED	2025-10-28	Reviewed from an information security perspective. The technical measures (pseudonymization, RBAC, encryption, retention enforcement) are robust and aligned with our security baseline. The outstanding TIA should be prioritized. Approved.