In-Class Exercise: Privacy-Tech Business Models

Exercise 1: Structured Debate — "Is BigID a Privacy Tool or a Data-Discovery Platform?"

Format: Split into two teams. Each team prepares arguments for its assigned position, then presents. After both sides speak, the class votes — but first, read the debrief questions.

Team A — "BigID Is a Privacy Tool"

Anchoring evidence: BigID is sold into Data Protection Officers and Chief Privacy Officers. Its product is purchased under GDPR, CCPA, and similar statutes. Its reference customers describe it as a privacy-operations suite. Its marketing language is privacy-first.

	Team A: BigID Is a Privacy Tool
Argument I	
Argument II	
Argument III	
Concession	*Strongest argument AGAINST your position:*
Closing	*How you address the concession:*

Team B — "BigID Is a Data-Discovery Platform"

Anchoring evidence: BigID's core engine indexes enterprise data stores and produces a classified inventory. That inventory drives privacy workflows, but the same inventory can drive security, governance, and analytics use cases. Many customers deploy it alongside or instead of data-governance platforms.

	Team B: BigID Is a Data-Discovery Platform
Argument I	
Argument II	
Argument III	
Concession	*Strongest argument AGAINST your position:*
Closing	*How you address the concession:*

Debrief Questions

Q1: Does the answer — privacy tool or data-discovery platform — matter for how BigID's competitors should position against it? Why or why not?

Q2: Could the answer genuinely be "both"? If so, what does that imply about how category analysts should classify privacy-tech vendors?

© Joerg Osterrieder 2026

Q3: Name another software vendor (in any sector) that blurs an established category boundary in a similar way. What tensions does that blurring create?

Exercise 2: Value Chain Mapping — The Data-Processing Pipeline

Scenario: The enterprise data pipeline can be broken into six links. Privacy-tech vendors attack individual links with specialised solutions. Your task: for each link, identify a privacy-tech vendor from the mini-lecture slate, describe the friction it removes, and predict the long-term outcome.

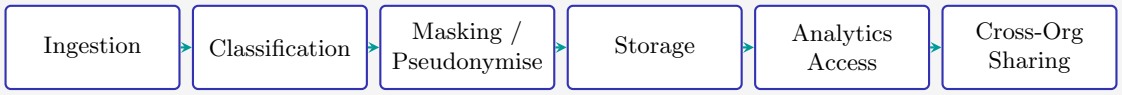

Pipeline Link	**Vendor Attacking It**	**Friction Removed**	**Replaces or Improves?**	**Incumbent Loses or Adapts?**
Ingestion				
Classification				
Masking				
Storage				
Analytics Access				
Cross-Org Sharing				

Synthesis Question

Q1: Which link in the pipeline is *most vulnerable* to privacy-tech attack? Which is *most resistant*? Defend your reasoning with reference to switching costs, regulatory triggers, and integration depth.

© Joerg Osterrieder 2026

Facilitator Solutions

Sample answers for instructor reference. These are illustrative; student reasoning may diverge and still be valid.

Exercise 1: Debate Sample Answers

Team A (BigID Is a Privacy Tool) — sample arguments

Argument I. BigID's buying centre is the Data Protection Officer, General Counsel, or Chief Privacy Officer. The budget it is sold into is the privacy-programme budget, not the data-platform budget. A product whose buyer, sponsor, and procurement pathway are all defined by privacy regulation is, for practical commercial purposes, a privacy tool.

Argument II. The product's differentiating features are all privacy-shaped: automated data-subject-access-request workflows, consent-state auditing, privacy-impact-assessment templates. A pure data-governance platform would not need any of those capabilities. The engineering investment concentrates on privacy-specific workflows that generic governance platforms deliberately do not build.

Argument III. The vendor's go-to-market is structured around privacy events — regulatory deadlines, supervisory guidance, and breach headlines. That cadence only makes sense for a privacy tool. A data-discovery platform sold into engineering leadership would market against use cases like data-quality, cataloguing, or analytics enablement; those are not the use cases BigID leads with.

Concession. The strongest argument against Team A is that the underlying discovery engine is general-purpose — it can (and at some customers does) serve security and governance use cases that have nothing to do with privacy.

Closing. Category follows buying centre. A tool purchased by privacy officers under privacy statutes, with privacy-specific workflows, in response to privacy triggers, is a privacy tool regardless of how general-purpose the underlying indexing engine may be.

Team B (BigID Is a Data-Discovery Platform) — sample arguments

Argument I. The core product is a crawler and classifier that indexes enterprise data stores. That capability is orthogonal to privacy — it is equally useful for security-operations, cloud-cost optimisation, data-governance, and analytics-enablement use cases. A product whose engine serves many use cases, even if one is dominant today, is a platform.

Argument II. Customers who start with a privacy use case routinely extend the deployment into security and governance. That extension pattern only works if the engine is general enough to serve adjacent use cases without re-platforming. Privacy tools that lack a discovery layer cannot be extended that way — the extension fact proves the platform nature.

Argument III. Competing categorisations from industry analysts increasingly place BigID alongside data-catalogue and data-governance vendors, not purely alongside privacy-operations tools. Capital markets value discovery platforms at higher multiples than privacy tools because the addressable market is larger and the product's use cases are less tied to any single regulatory cycle.

Concession. The strongest argument against Team B is that most revenue today still comes from privacy-programme budgets, and most sales motions still cite privacy statutes as the trigger.

Closing. Current revenue concentration reflects go-to-market history, not product architecture. A discovery platform whose engine, integrations, and expansion pattern are general-purpose is a platform; the privacy wedge is the entry point, not the ceiling.

© Joerg Osterrieder 2026

Debrief Q1 — Competitive positioning

The answer matters because a competitor targeting BigID must decide whether to match a privacy tool or a discovery platform. Matching a privacy tool means investing in consent, data-subject-access, and privacy-impact-assessment workflows; the addressable buyer is the privacy officer, and the displacement strategy is feature parity. Matching a discovery platform means investing in the indexing engine, integrations, and classification taxonomies; the addressable buyer is broader, and the displacement strategy is horizontal expansion. Choosing the wrong frame wastes engineering effort on features the buyer does not care about or leaves the competitor with a better privacy tool that still loses deals because BigID is displacing it as a platform.

Debrief Q2 — "Both" as an answer

The answer genuinely can be "both": BigID is a general-purpose discovery engine sold through a privacy-shaped go-to-market. That duality reveals that traditional category boundaries, drawn when each tool served a single buyer with a single use case, no longer capture modern enterprise software where a single engine crosses several buying centres. If "both" is the right answer, it implies that category analysts need functional classifications — based on what the engine can do — rather than institutional classifications based on which department currently pays for it. A purely privacy classification understates the engine's reach; a purely discovery classification understates the commercial reality that privacy budgets are underwriting today's growth.

Debrief Q3 — Cross-sector blurring example

Snowflake blurs the boundary between data warehouse and application platform. It sells compute and storage but derives a growing share of its value from the marketplace of data-apps and shared data-sets built on top of its platform. Traditional warehouse vendors are valued on workload multiples; platform vendors on ecosystem multiples. The tension this creates is acute for competitors (do you benchmark against Oracle or against a platform-as-a-service vendor?), for customers (do you buy it as infrastructure or as an application layer?), and for analysts (which comparable set should price the company?). The parallel to BigID is direct: the blurring is not a marketing claim but a structural consequence of a general-purpose engine shipped through a narrower go-to-market.

© Joerg Osterrieder 2026

Exercise 2: Pipeline Value-Chain Mapping Sample Answers

Pipeline Link	Vendor Attacking It	Friction Removed	Replaces or Improves?	Incumbent Loses or Adapts?
Ingestion	Sherpa (on-device capture)	Central-server footprint that creates server-side exposure	Replaces	Incumbent Adapts
Classification	BigID (discovery + classification)	Unknown personal-data inventory across enterprise stores	Replaces	Incumbent Loses
Masking	Privitar (in-pipeline pseudonymisation)	Downstream teams reaching for raw identifiers	Improves	Incumbent Adapts
Storage	Anonos (regulated pseudonymised stores)	Interpretive uncertainty about whether a downstream use is in scope	Improves	Incumbent Adapts
Analytics Access	OneTrust (consent + purpose gating)	Audit-time reconstitution of lawful-basis per query	Improves	Incumbent Adapts
Cross-Org Sharing	OneTrust (vendor-risk hub for processor attestations)	Fragmented vendor-attestation workflows across processors	Improves	Incumbent Adapts

Synthesis Question Sample Answer

The most vulnerable link is Classification. A new entrant can deploy a discovery engine against an existing data estate without any re-architecture of upstream ingestion or downstream consumers; the vendor ships crawlers and classifiers that sit outside the hot path. Switching costs at the classification layer are modest because the output is an inventory artefact that any downstream workflow can ingest. Network effects compound this: each new customer extends the classification taxonomy with more edge cases, raising accuracy for every other customer. The most resistant link is Masking. Once a masking vendor is inserted into the pipeline, every downstream dataset depends on the vendor's pseudonym mapping and tokenisation keys; switching away requires re-processing every downstream store and regenerating historical joins. Regulatory triggers entrench the incumbent further because changing the masking layer can invalidate prior consent records and audit evidence. A new entrant can improve the interface, but rarely displaces the in-pipeline chokepoint without a full migration programme that customers are reluctant to underwrite.

© Joerg Osterrieder 2026