Pre-Class Discovery Handout: Privacy-Tech Business Models

Activity 1: Business Model Canvas Detective — OneTrust

Scenario: Pick the privacy-tech vendor OneTrust (or, if you already have direct experience with a different one from the mini-lecture slate — BigID, Privitar, Anonos, or Sherpa — use that instead). Fill in the Business Model Canvas below by investigating how that vendor actually creates value. Focus on the mechanics of pricing an absence, not on marketing slogans.

Canvas Element	Your Analysis
Value Proposition *Which avoided downside does the buyer purchase?*	
Customer Segments *Which officer or team holds the budget?*	
Channels *How does the vendor reach privacy-officer buyers?*	
Revenue Streams *What does the pricing scale with (footprint, jurisdictions, modules)?*	
Key Activities *What does the vendor do every day on the customer's behalf?*	

Q1: What specific downside is the buyer trying to avoid by paying for this product?

Q2: How does the vendor reach new customers, given that the buying centre is a legal or compliance function rather than a revenue-generating business unit?

Q3: If this vendor disappeared tomorrow, what would the customer lose that a generic cloud platform cannot replace?

Activity 2: Unbundling Map — The Privacy-Tech Stack

Scenario: A single enterprise used to buy one bundled compliance suite from its data-centre vendor. Today, specialist privacy-tech vendors attack individual links of that bundle. Match each vendor below to the link it unbundled, then answer the questions.

Vendor		Link Unbundled
OneTrust	____	Consent management + vendor-risk hub
BigID	____	Data discovery + classification
Privitar	____	Masking + pseudonymisation in-pipeline
Anonos	____	Regulated pseudonymisation + patent overlay
Sherpa	____	On-device privacy assistant

Q1: For each pair, describe in one sentence what avoided downside the vendor delivers to its buyer.

Q2: Which of these vendors has started adding products beyond its original wedge? What did it add?

Q3: Why might a privacy-tech vendor that starts with one narrow product eventually want to offer a suite?

Activity 3: The Privacy-Hub Puzzle

Scenario: A privacy-operations hub connects two sides of a market — enterprise customers who hold personal data and third-party vendors whose processing those enterprises must audit. Neither side finds the hub useful without the other.

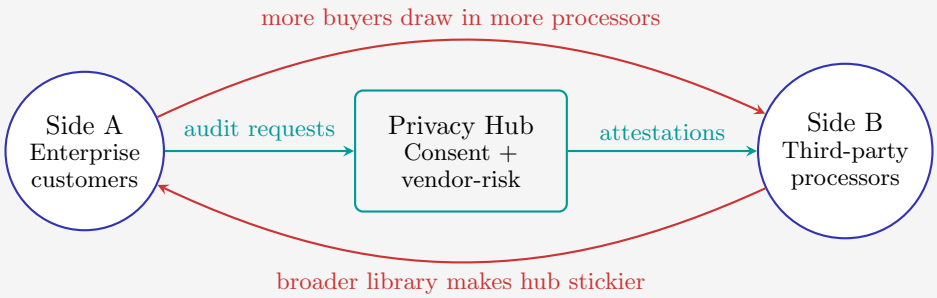

Q1: Why does a privacy hub with more enterprise customers attract more third-party processors, and vice versa?

Q2: The chicken-and-egg problem: which side should the hub attract first, and why?

Q3: Once the hub reaches critical mass, why is it hard for a competing hub to enter?

Solutions

Activity 1: Business Model Canvas Detective

A1: Model answer for OneTrust: The buyer avoids three downsides simultaneously — a cookie-consent enforcement action, a missed subject-access-request deadline, and an unflagged third-party processor that breaches and takes the enterprise's name into the headline. None of those avoided outcomes is visible on a dashboard; the value proposition is the specific shape of exposure that stops accumulating.

A2: OneTrust reaches privacy-officer buyers through direct enterprise sales, industry communities, peer-referenced benchmarking, and compliance conference presence. The buying centre is the Data Protection Officer, General Counsel, or Chief Information Security Officer; acquisition runs through those professional communities rather than through end-user self-service.

A3: The customer would lose the consolidated consent record, the mapped processing inventory, the vendor-risk library, and the DSAR workflow audit trail. A generic cloud platform provides storage and compute, not the privacy-specific orchestration that turns raw personal data into a defensible audit position.

Canvas elements (OneTrust):

- **Value Proposition:** Consolidated privacy-operations hub spanning consent, data-mapping, subject-access, and vendor-risk — priced against the downside of a fragmented-tool failure.
- **Customer Segments:** Primary — large enterprises with Data Protection Officers; secondary — mid-market firms whose growth triggered statutory thresholds.
- **Channels:** Direct enterprise sales, compliance-officer communities, analyst-report placement, regulator-conference visibility.
- **Revenue Streams:** Subscription tiers scaled by data footprint and regulated-jurisdiction count, with module add-ons (consent, mapping, DSAR, vendor-risk, breach ops).
- **Key Resources:** Vendor-risk assessment library, privacy-regulation content team, integration surface with identity providers and cloud stores.

Activity 2: Unbundling Map

A1: OneTrust → Consent + vendor-risk hub (removes the fragmented-tool risk of inconsistent consent records and unmonitored processors). BigID → Discovery + classification (removes the risk of holding personal data whose existence the enterprise cannot even enumerate). Privitar → Masking in-pipeline (removes the analytics-team pressure to reach for raw identifiers). Anonos → Pseudonymisation under a privileged regulatory reading (removes the interpretive uncertainty about whether a downstream use is in scope). Sherpa → On-device processing (removes the central-server footprint that creates server-side exposure).

A2: OneTrust began with a cookie-consent wedge and added data-mapping, subject-access, vendor-risk, and policy modules. BigID started with discovery and expanded into classification, access auditing, subject-access workflows, and breach operations. Both illustrate **rebundling** — the wedge earns trust, adjacent modules attach, and the suite re-forms around the original entry point.

A3: A single-product privacy vendor has high acquisition costs because the buyer is a specialist with a full procurement cycle. Once the vendor is deployed and trusted, the marginal cost of offering an adjacent module is low while the marginal revenue is high. Rebundling raises lifetime value per customer and creates switching costs, because privacy data flows that span multiple modules are hard to reconstitute elsewhere.

Activity 3: The Privacy-Hub Puzzle

A1: This is a **cross-side network effect**: more enterprise customers mean the hub's vendor-risk library covers a larger share of the third-party processors any new enterprise needs to audit. Simultaneously, more processors publishing attestations through the hub reduce audit-time for every enterprise. Each side's growth reinforces the other's.

A2: Most successful privacy hubs attract the **enterprise-customer side** first, because that side writes the cheques. Processors follow because enterprise customers demand it; cold-starting the processor side is hard when no enterprise is asking for attestations through the hub.

A3: Once critical mass is reached, the hub enjoys a self-reinforcing loop that creates a **structural moat**. A new entrant would need to simultaneously attract enterprises (who have no reason to leave a hub where their processors already publish) and processors (who have no reason to add an extra publication surface). The incumbent's library grows with every audit cycle, widening the gap. Competing hubs must find a vertical niche or offer dramatically superior attestation tooling to pry either side loose.

