



Protecting data, unlocking insight

PET-enabled architectures for bank supervision

Jakub Demski, *Workshop on data security and privacy*, September 14

The views expressed in this presentation are those of the author and do not necessarily reflect the views of the BIS or its member central banks.

Agenda

- The problem
- Current arrangements
- What PETs offer
- Supervisory use cases
- Selection guide
- Infrastructure implications
- Pilot path


The problem

Analytical demand

- Granular institution-level insights
- Cross-border group and network dependencies
- Timely evidence for supervisory action

Constraints


- Supervisory secrecy and confidentiality
- Data sovereignty and purpose limits
- Commercial, customer and cyber sensitivity



The properties that make supervisory data insightful also make them difficult to share

What works today

- Today's tools: MoUs, colleges, secure channels, institution-specific collections
- International Data Hub (IDH) shows pooling works when scope and rules are tight



- Pooling can work in narrow, governed scopes — but doesn't fit every purpose
- Limiting factors: law, timeliness, compute, output sensitivity
- PETs complement – not replace – governed hubs

PET-enabled architectures

SMPC



Secure multiparty computation (SMPC) for protected aggregation and benchmarking

TEE



Trusted execution environments (TEEs) for controlled in-hardware processing

Federated



Federated analytics for recurring indicators on horizontally split data

Differential privacy



Output governance and differential privacy to protect results

Match architecture to use case, not the other way round

Supervisory use cases

Group risk mapping

Map vulnerabilities across banking groups (intragroup dependencies, concentrations, cross-entity exposures)

Hierarchical + networked

Liquidity and funding contagion

Trace stress spreading through networks

Simulation + network

Peer benchmarking

Compare peer positions across jurisdictions

Horizontal data

AML oversight

Detect suspicious patterns (anomaly scores, clusters)

Transactional + mixed

Macroprudential dashboards

Monitor system-wide trends and distributions

Horizontal + panel

Model reviews

Identify model outliers (cross-firm metrics, consistency checks)

Mostly horizontal

Crisis monitoring is a cross-cutting operating mode

Selection guide: 10 core considerations


- 1 Legal permissibility
- 2 Data topology
- 3 Analytical complexity
- 4 Output sensitivity
- 5 Trust model
- 6 Performance & latency
- 7 Auditability & explainability
- 8 Operational maturity
- 9 Governance burden
- 10 Crisis readiness

Law narrows the feasible set. The other dimensions determine suitability and safeguards.

Infrastructure implications


Output and control

- Output-tiering
- Restricted supervisory reporting
- Disclosure review
- Differential privacy where appropriate
- Audit logs



Metadata, harmonisation and orchestration

- Common definitions
- Entity mapping
- Reporting standards
- Workflow orchestration
- Query / code approval
- Metadata & lineage



Protected computation modes

- Data pooling
- Federated analytics
- SMPC workflows
- TEE / confidential computing
- Hybrid PET workflows



Governance and legal basis

- Supervisory mandates
- Confidentiality rules
- Participation agreements
- Role-based permissions
- Incident-response arrangements



The platform coordinates protected workflows; it need not centralise every input.

Pilot path

- Pick a narrow use case with clear outputs (eg peer benchmarking; G-SIB bucketing)
- Pre-position legal basis, participation rules, access-tiers
- Harmonise inputs; fix specification; pre-approve code; record hash
- Execute protected workflow (SMPC, secure aggregation, or TEE) side-by-side with the established process
- Evaluate: accuracy vs reference; latency under load; resilience; attestation and code-approval; data-lineage capture; output-tiering effectiveness
- Decide: which capabilities become standing infrastructure

Conclusion

- Three takeaways
 - Begin with the supervisory purpose; let it drive architecture choice
 - Combine PETs with governance: purpose control, attestation, audit, output-tiering
 - Orchestrate multi-mode workflows; pre-position for crisis-time use
- Implementation
 - Identify one to two high-value use cases; map legal and operational constraints
 - Select an architecture using the ten considerations; define assurance model
 - Run a bounded pilot alongside the reference process; evaluate and decide
- Thank you