

Business Intelligence Platform: Selection, Architecture, and Tradeoffs



- Equip BI leaders, analytics engineers, data/IT teams, and procurement partners to select and govern a BI platform



- Three lenses for every decision: selection criteria, architecture patterns, tradeoffs



- Platform choice is a multi-year commitment: cost, agility, governance, talent, and AI readiness



- Session toolkit: seven-question decision filter, weighted scorecard, POC design, takeaway checklist

Why Platform Choice Matters Now



- 2026 market: hyperscaler suites, independents, and open-source options



- BI now serves analysts, line-of-business users, and AI agents




- Price pressure: Tableau up 9%; Fivetran-dbt merger drives consolidation





- BI is the most swappable layer to optimize TCO



- Poor fit fails quietly: low adoption, shadow IT, metric sprawl




Core Capability Model for a BI Platform



Architecture Patterns: From Warehouse to Semantic Layer




Centralized Warehouse plus BI




- Centralized warehouse plus BI: strong governance, single source of truth, but bottlenecks

Data Mesh and Federated Analytics




- Data mesh and federated analytics: domain ownership, data as a product, federated governance

Semantic Layer-first




- Semantic layer-first: metrics as code (dbt MetricFlow, LookML, Cube, AtScale) reused across tools

Embedded and Composable



- Embedded and composable: headless BI, APIs, customer-facing and operational analytics


Hybrid and Multi-cloud



- Hybrid and multi-cloud: latency, cost, resilience, and data residency implications



Where Should the Semantic Layer Live?



- BI-tool-resident (LookML, DAX): fast for analysts, locks metrics to one tool
- Warehouse-resident (dbt MetricFlow, Snowflake Semantic Views): consistent, emerging default
- Headless/universal (Cube, AtScale, GoodData): portable and AI-agent ready, more infrastructure
- Multiple BI tools, embedded analytics, or AI agents? Weight portability heavily
- Platform-embedded models create lock-in; warehouse-native or independent layers cut switching cost

Selection Criteria and Weighted Evaluation



EVALUATION WEIGHT (ADJUSTABLE)					
CRITERIA					
 - Functional fit: visualization depth, self-service, embedded and customer-facing use cases					
 - Non-functional: petabyte-scale performance, 20–50 concurrent sessions, availability					
 - Governance: SOC 2 Type II, ISO 27001, data residency, RLS enforced at query time					
 - Ecosystem: cloud platform, warehouse depth, SSO/SAML/SCIM, semantic-layer interoperability					
 - Commercial: pricing model, renewal uplift cap, worst-case consumption bill, exit path					



Pricing and TCO: Where BI Deals Quietly Fail



- Sticker spread is extreme: median entry \$55/month, mean \$409 across 201 tools



- 3-year TCO = license + capacity/credits + implementation + training + admin headcount



- Labor dwarfs software: \$9 per software dollar (Tableau), ~\$35 (Power BI)




- Capacity traps: Fabric F64 floor ~\$8,410/month; PPU-to-Fabric break-even ~350 users



- Negotiate 7–10% renewal cap, fixed overage; price at 2x users and data

THREE-YEAR TOTAL COST OF OWNERSHIP (TCO)



CAPACITY FLOOR AND BREAK-EVEN



Fabric F64 floor
~\$8,410/month



PPU-to-Fabric
break-even
~350 users

PRICING MODEL SPLIT



Per-Seat (PPU)
Price by User

VS.




Capacity-Based (Fabric)
Price by Capacity




Negotiate 7–10% renewal cap, fixed overage;
price at 2x users and data

Tradeoffs That Decide the Outcome




Governed Self-Service: The Hybrid Model

- **Centralized governance:** 57% fewer security incidents, 83% fewer compliance violations
- **Up to 60% less dashboard duplication** through consolidated definitions and ownership
- **Decentralized execution:** 25–40% faster reporting cycles, but fragments without guardrails
- **Governed hybrid:** central semantic layer plus controlled self-service at the edge
- **Maturity path:** centralize, decentralize gradually, then run distributed ownership under guardrails
- **Watch overload:** shadow BI and backlog vs. metric drift and duplicate tools






Governance, Security, and Cost Management




- **Access control:** RBAC, ABAC, row/object-level security, and data residency enforced at the semantic layer
- **Metadata and lineage:** source-to-dashboard-to-AI visibility, certification, and named ownership
- **Usage monitoring and cost allocation:** chargeback/showback, license optimization, CU/token tracking
- **Compliance readiness:** retention, eDiscovery, SIEM-exportable audit logs, regulated-industry evidence
- **AI governance:** sensitivity labels, DLP, RLS validation, approved use cases, human-in-the-loop

AI and the Semantic Layer: Evidence and Limits

Natural-language query accuracy



A machine-queryable semantic layer



- **Raw schemas fail AI:** 10.1% GPT-4o baseline on 1,000+ column enterprise schemas



- **Semantic context lifts accuracy:** 45.5–50.5% to 67.7–68.7%, a 17–23 point gain



- **Even with oracle hints,** ~70% error rate remains on the hardest benchmarks




- **A real semantic layer** is machine-queryable, governed, reusable — not a glossary



- **Trust needs** named ownership, certification, provenance, and RLS regression tests

Migration and Consolidation Playbook



- ✓ Inventory first: retire the 40–60% of legacy reports unused or redundant
- ✓ Logic extraction finds 30–50% more hidden complexity than expected
- ✓ Migrate in waves: pilot, parallel run two cycles, validated cutover
- ✓ Announce 90 days ahead; champions drive training on daily workflows
- ✓ Validate metric parity and RLS; align decommission dates with renewals

Decision Framework: Seven Questions and a Weighted Scorecard



1



Ask the seven questions before any demo; answers eliminate most of the market

2



Scorecard weights: cost, adoption, capacity economics carry ~39 points

3



Score all candidates against identical criteria; demand written evidence, not claims

4



Map the buying committee: CFO, data lead, security, IT

5



Bring one page: decision, three-year cost, risk, payback, security, rejected option

6



Pre-signing checklist: pricing floor, renewal cap, worst-case bill, exit path, named owner

7



WEIGHTED SCORECARD

CRITERIA	WEIGHT	CANDIDATE A	CANDIDATE B	CANDIDATE C
 Cost	Highest Weight			
 Adoption	Highest Weight			
 Capacity Economics	Highest Weight			
 Security	Other Factors			
 Operational Fit	Other Factors			
 Risk	Other Factors			
 Strategic Value	Other Factors			

BUYING COMMITTEE



CFO



DATA LEAD



SECURITY



IT

Pilot Design and the 10-Question Takeaway Checklist

30-DAY PILOT FIELD CARD

TESTER ROLES



ANALYST

Represents power data consumer



NON-ANALYST

Represents casual data consumer

WEIGHTED SUCCESS CRITERIA (7 TOTAL)

WEIGHT



Query coverage

20%

☐

Performance and responsiveness

15%

☐

Security and governance

15%

☐

Usability and adoption

15%

☐

Integration and ecosystem

15%

☐

Total cost of ownership (TCO)

10%

☐

Extensibility and roadmap fit

10%

☐

FAILURE SIGNALS



Excel exports



Metric drift above 0.5%



Undocumented consumption costs

ONE-PAGE PILOT REPORT TEMPLATE



- ✓ Run a 30-day POC on real data, with one analyst and one non-analyst
- ✓ Set written success criteria before contacting any vendor
- ✓ Weighted rubric: query coverage 20%, five criteria at 15%, TCO and extensibility 10%
- ✓ Failure signals: Excel exports, metric drift above 0.5%, undocumented consumption costs
- ✓ Output: one-page report plus 10-question checklist before signing

PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/725e488e/public