# Bio-IT Conference Intelligence Brief

## Personalized Conference Debrief

PREPARED FOR

**Jane Smith, PhD**

VP & Head of AI-ML, Quantitative & Digital Sciences

Pfizer Inc. | New York, NY

| Conference | Sessions Covered | Delivered | Brief Type |
|---|---|---|---|
| BIO-IT World 2025 | DAY 1, 2, and 3 | 48 Hrs Post-Event | Strategic Intelligence Brief |

| 229 | 10 | 3 | 5 |
|---|---|---|---|
| Sessions Analysed | Competitors Tracked | Critical Signals | Actions to Take |

### WHY THIS BRIEF MATTERS FOR YOU

*This brief translates BIO-IT 2025 into structured intelligence — capturing signals across 229 sessions and highlighting the insights, opportunities, and follow-ups most relevant to Pfizer— to support your leadership reporting, programme benchmarking, and near-term decisions.*

### OVERALL CONFERENCE ASSESSMENT

| Dimension | Assessment |
|---|---|
| **Biggest Signal** | Insilico Medicine's Phase 2A success. First public proof that fully AI-designed drugs can pass clinical milestones. More significant than any platform announcement. |
| **Production Reality** | Novartis, BMS, AbbVie, Takeda all disclosed production-grade AI with specific metrics. Industry has moved past experimentation. Velocity is now the competitive variable. |
| **Competitive Surprise** | AbbVie — 6 coordinated sessions spanning molecular GenAI, clinical documents, secondary use of clinical data, data analytics, digital retrospective, and federated learning. |
| **Shared Bottleneck** | Data readiness. Named identically by AstraZeneca, Roche, Takeda, Snowflake: 80% of time on wrangling, very few gigabytes AI-ready despite petabytes stored. |
| **Vendor Standout** | ServiceNow + NVIDIA — disclosed 25% reduction in development costs and 500-day faster time to market from lab standardisation. Most specific ROI claim at the conference. |
| **Open Whitespace** | Responsible AI governance — mentioned as a challenge in many sessions, solved publicly by none. Leadership position available now. |

## CONTENTS

# 01 — EXECUTIVE SUMMARY

## Three Things Leadership Needs to Know

**1. The first fully AI-designed drug has completed Phase 2A clinical trials — meeting its endpoints.**

- Insilico Medicine disclosed at the D2 and D3 plenaries that their lead program completed Phase 2A in China, meeting primary and secondary endpoints.
- AI platform used: Pandomics (target discovery) → Chemistry 42 (molecule design, using multi-agent reinforcement learning) → Cliniko (clinical success prediction).
- Quantified value: saving 18 months in preclinical development improves net present value by ~40%, regardless of where in the pipeline those months are saved.
- This is not a claim — it is a disclosed clinical result. The competitive question has shifted from whether AI can design drugs to who builds the best platform before the field matures.

**2. GenAI has moved from pilot to production — multiple companies disclosed working systems with measured outcomes.**

- Novartis: 4 tools shipped in under 10 months plus NSA (live institutional knowledge search). Disclosed in 7 separate sessions.
- BMS: DocGen AI covering 20+ clinical document types at 40% average time saving, delivered via Word add-in. Built over 18 months with ZS Associates.
- AbbVie: GAIA platform — 100+ person programme, 3 documents in production, company-wide request pipeline, built with Accenture.
- Takeda: Protocol intelligence tool using GenAI to quantify patient burden and benchmark against competitor trials — in production.

**3. FAIR data infrastructure and responsible AI governance are both unclaimed — and both matter within 24 months.**

- Takeda presented 4 coordinated sessions on FAIR data. Their stated reality: 80% of data science time spent on wrangling before any AI can run.
- AstraZeneca, Roche, and Snowflake independently described the same problem — petabytes of data, very few gigabytes AI-ready.
- Across all 229 sessions at BIO-IT 2025, zero companies published a formal responsible AI governance framework.
- Both positions are open. Companies that move on FAIR data now, and publish governance frameworks before regulators require them, establish structural advantages.

# 02 — STRATEGIC RECOMMENDATIONS FOR PFIZER AI-ML

## Five Actions Tied to Conference Disclosures

*Each recommendation is anchored to a specific BIO-IT 2025 session. Each is conditional — only relevant if the stated condition applies to your organisation.*

### 01 Evaluate building an institutional knowledge retrieval tool

- **What was disclosed:** Novartis NIBR disclosed NSA — a RAG system indexing 1,500+ internal documents with molecule, reaction, and text search, returning exact source and page references.
- **How it was built:** A small team during a company recharge period. Elasticsearch, BM25 keyword search, semantic embeddings for deep semantic context. Two engineers, one sprint.
- **Why it matters:** The institutional knowledge risk — expertise walking out when people change roles — applies to every large pharma organisation identically.

| EVIDENCE | RELEVANT IF | TIMELINE / OWNER |
|---|---|---|
| Novartis T6 — NSA disclosure, April 2025 | *you don't already have a comparable institutional knowledge retrieval system* | 30-day feasibility / AI-ML Engineering |

### 02 Confirm or initiate a position on the AISB Consortium

- **What was disclosed:** AbbVie, Columbia University, and Apheris disclosed the AISB Consortium — a live, legally structured federated learning arrangement for protein-ligand structure prediction.
- **How it works:** Each company trains AI inside its own firewall. Only model parameter updates are shared. AbbVie contributes 30 years of protein-ligand crystal structures.
- **The window:** The CEO of Apheris issued an explicit on-stage invitation for additional pharma companies to join. Early members shape the consortium's direction.

| EVIDENCE | RELEVANT IF | TIMELINE / OWNER |
|---|---|---|
| AbbVie T5 — AISB disclosure, April 2025 | *you are not already in discussions with Apheris or a comparable arrangement* | Status confirmed 10 days / BD + AI-ML jointly |

## 03 Benchmark clinical document automation against BMS and AbbVie

- **BMS DocGen AI:** 20+ document types, 40% average time saving, Word add-in delivery. Built over 18 months with ZS Associates. Architecture documented in T6 transcript.
- **AbbVie GAIA:** 100+ person programme with Accenture. 3 documents in production. Company-wide request intake for any department to submit additional documents.
- **Why both:** Two reference points at different scales — BMS for architecture depth, AbbVie for enterprise breadth. Both are now public benchmarks.

| EVIDENCE | RELEVANT IF | TIMELINE / OWNER |
|---|---|---|
| BMS T6 + AbbVie T11 — April 2025 | *your programme has not been benchmarked against a production deployment* | Architecture review 60 days / Clinical AI + Regulatory Affairs |

## 04 Publish your responsible AI governance framework externally

- **The gap:** Zero companies across all 229 BIO-IT 2025 sessions presented a formal, published responsible AI governance framework.
- **The opportunity:** A tier-based, pharma-specific framework — distinguishing patient-facing, regulatory, and internal AI risk tiers — establishes the industry standard before regulators define one.
- **Why now:** This is a communication decision, not a build decision. If your organisation has a governance framework, publishing it now claims the leadership position while it is still available.

| EVIDENCE | RELEVANT IF | TIMELINE / OWNER |
|---|---|---|
| Gap confirmed across all 229 BIO-IT 2025 sessions | *your organisation has a governance framework not yet published externally* | Publishing decision 90 days / AI-ML + Medical Affairs |

## 05 Run an internal FAIR data maturity assessment

- **What Takeda disclosed:** 4 coordinated sessions covering semantic harmonisation, FAIR data lakes, TakOmics multiomics platform (in production), and protocol intelligence.
- **The industry finding:** 80% of data science time spent on data wrangling — stated identically by Takeda, AstraZeneca, and Roche from independent sessions.
- **The action:** A structured internal assessment against Takeda's disclosed standard surfaces where AI-readiness gaps are largest and which datasets need harmonisation prioritisation first.

| EVIDENCE | RELEVANT IF | TIMELINE / OWNER |
|---|---|---|
| Takeda T11, T5, T9 — April 2025 | *you haven't recently assessed FAIR data maturity across key R&D domains* | Assessment scoped 45 days / Data Governance + AI-ML |

## 03 — What This Means for Pfizer AI-ML

# Pfizer's Competitive Position After BIO-IT 2025

*Based on 229 sessions at BIO-IT World 2025 | April 2025 | Prepared for Pfizer AI-ML*

**The in silico first window is open — and the clock is now measurable.**

- Insilico Medicine's Phase 2A success gives every pharma organisation a concrete reference point: AI-designed drugs can reach and pass clinical milestones.
- The companies that build the most capable in silico platforms over the next 2–3 years will have a structural advantage that is very difficult to reverse.
- Pfizer's scale, modality breadth, and proprietary data assets are significant advantages — the question is whether the platform investment is being made at the pace that Insilico's disclosed timeline implies is now competitively necessary.

**The competitive gap is not in ambition — it is in deployment velocity and coordination.**

- The companies that stood out at BIO-IT 2025 were not those with the biggest AI strategies — they were the ones that had shipped working tools and could cite specific numbers.
- Novartis had 7 coordinated sessions and a replicable process. AbbVie had 6 sessions and a publicly visible programme spanning molecular design to clinical documents.
- The intelligence question for Pfizer is not 'do we have comparable capabilities' but 'would our comparable capabilities be visible from the outside if we presented at BIO-IT 2025?' Visibility and coordination matter as much as the capability itself for attracting talent, partners, and pre-competitive collaboration.

**The governance whitespace is a disproportionate opportunity given your specific position.**

- You co-presented on the Opening Plenary alongside venture capital leaders (Third Rock Ventures, Bessemer Venture Partners), an AI drug discovery company (BioMap), and Mayo Clinic. The FDA representative was unable to join on the day. Pfizer remained the only large pharma company on the panel.
- Publishing a tier-based responsible AI governance framework — before any regulator requires one — would be cited, referenced, and benchmarked against by the industry.
- That is a different category of impact from any of the five strategic recommendations in this brief. It is available now, and it is yours specifically to take.

# APPENDIX

## Reference Intelligence

BIO-IT World 2025 • Detailed Session Analysis

*Competitive Intelligence • Emerging Themes • Vendor Watch • Follow-Up Contacts*

*Each company card below includes a link to purchase the recorded session.*

## A1 — COMPETITIVE INTELLIGENCE: WHAT COMPETITORS REVEALED

### Key Disclosures by Company

*Every finding is sourced from a named session and a dated transcript. Session links open the full transcript in Airtable. Purchase recording links go to the BIO-IT on-demand portal.*

#### Insilico Medicine • *D2 + D3 Plenaries*

**First publicly disclosed AI-designed drug completes Phase 2A — meets primary and secondary endpoints.**

- **The result:** Phase 2A completed in China for their lead program targeting IPF (idiopathic pulmonary fibrosis) via the TNIK kinase. Met primary and secondary endpoints for safety and efficacy.
- **The platform:** Pandomics (target discovery) → Chemistry 42 (multi-agent reinforcement learning for molecule design) → Cliniko (clinical success prediction). Full pipeline, not point solutions.
- **The timeline:** From target identification to development candidate in under 18 months. Full in vitro and in vivo testing completed — no steps skipped.
- **The value framework:** Saving 6 months in preclinical development improves net present value by 10%. Saving 18 months improves it by nearly 40% — regardless of where in the pipeline those months come from.

QUESTION FOR YOUR CONTEXT

*How does your organisation's in silico drug discovery capability compare to a company that has now demonstrated Phase 2A success with a fully AI-designed compound?*

**Evidence:** D3 — Revolutionizing Drug Discovery D2 — GenAI, Aging & Robotics
**Recording:** Purchase on-demand recording ↗

## Novartis • *7 sessions across S3, T6, T7*

**4 GenAI tools shipped in under 10 months, a live institutional knowledge search system, and active learning now closing the wet-dry lab loop.**

- **Tool suite (S3):** 72 ideas crowdsourced, scored, refined to 5 — 4 in production: Protein Co-pilot (peptide extraction), FAIRchat (NL-to-SQL), ALIVE (animal licence automation), ItKnows (gene research summaries).
- **NSA — institutional knowledge (T6):** RAG system indexing 1,500+ internal SharePoint documents. Elasticsearch + BM25 + semantic embeddings. Returns exact source document and page number. Built by a small team during recharge period.
- **Active learning (T7):** ML now guiding which experiments to run next in drug discovery — closing the wet-dry lab feedback loop. This is beyond document automation into scientific decision-making.
- **What it means:** Two independent teams, two different deployment tracks, all disclosed publicly. The velocity and coordination are the significant signal.

QUESTION FOR YOUR CONTEXT

*How does your organisation's progression from GenAI tools to active learning in drug discovery compare to what Novartis disclosed across these 7 sessions?*

**Evidence:** S3 — GenAI in Action T6 — NSA: Institutional Knowledge T7 — Active Learning
**Recording:** Purchase on-demand recording ↗

## AbbVie • *6 sessions across S2, T4, T5, T6, T10, T11 | April 2025*

**The broadest and most coordinated AI programme at the conference — spanning molecules, clinical documents, secondary use of clinical data, and federated learning.**

- **GAIA — clinical documents (T11):** 100+ person programme built with Accenture. 3 documents in production. Any department can submit a request for document automation. Distinct from BMS — enterprise-scale platform approach.
- **AISB Consortium — federated learning (T5):** Live legal structure with Columbia University and Apheris for protein-ligand structure prediction. AbbVie contributes 30 years of crystal structures. Open to new pharma members — CEO issued on-stage invitation.
- **Secondary use of clinical trial data (T10):** A dedicated group (Convergence) working to make secondary use of completed clinical studies the default company philosophy — democratising access across R&D and commercial.
- **10-year digital retrospective (S2):** Explicitly framed the industry as being in a digitalization-to-digital transformation transition, with in silico-first drug discovery as the endpoint. Now on the public record.

QUESTION FOR YOUR CONTEXT

*Does your organisation have an equivalent secondary use programme for completed clinical trial data? Is the breadth of AbbVie's coordinated AI investment visible in public signals?*

**Evidence:** **T5** — Federated Learning / AISB T11 — GAIA Clinical Docs S2 — 10-Year Retrospective T10 — Secondary Clinical Data
**Recording:** Purchase on-demand recording ↗

## Bristol Myers Squibb • *T6 + S3*

**DocGen AI in production — 20+ clinical document types, 40% average time saving, delivered via Word add-in.**

- **Architecture:** Modular Blueprint (template + prompts) + Reference Sources (prior documents + databases). Flexible per document type. Word add-in delivery for workflow-native adoption.
- **Scale:** 20+ study teams using the platform. 20+ document types enabled. 50+ instances generated. 18 months in development with ZS Associates as the scale-up partner.
- **Metrics:** 40% average time saving (self-reported from early adoption teams, range 20–60%). Not independent audit data — but the most specific productivity disclosure at the conference.
- **Roadmap:** Automated quality review, translation, and self-updating document ecosystems. Architecture designed for extension — not a fixed product.

QUESTION FOR YOUR CONTEXT
*How does your clinical document automation programme compare to BMS DocGen AI on coverage, adoption depth, and architecture? The T6 transcript documents their full architecture.*

**Evidence:** T6 — DocGen AI S3 — LLMs for Breakthroughs
**Recording:** Purchase on-demand recording ↗

## Takeda • *5 sessions across T9, T11, T5, S2*

**The most coordinated FAIR data investment at the conference — TakOmics in production, FAIR data lakes, semantic harmonisation, and a GenAI protocol intelligence tool.**

- **TakOmics (T9):** A FAIR multiomics data management and analytics platform now in production for Cell Therapy and Global Biologics teams. Vendor partners: OntoForce, QuartzBio, AWS, Healthomics.
- **FAIR data lakes (T5):** One Research Digital — ecosystems of FAIR data lakes enabling AI-augmented analysis. Goal: reduce data wrangling from 80% of time to under 20%.
- **Protocol intelligence tool (T11):** GenAI tool that quantifies patient burden from protocol schedules, benchmarks against competitor trials in clinicaltrials.gov, and evaluates inclusion/exclusion criteria using real-world data.
- **Semantic harmonisation (T11):** Standardising ontologies across clinical and preclinical data — addressing the problem that identical data stored by different teams is described with different vocabulary.

QUESTION FOR YOUR CONTEXT
*How does your FAIR data maturity compare to Takeda's 4-session coordinated investment? Is patient burden analysis at protocol design phase something your clinical teams currently do systematically?*

**Evidence:** T9 — TakOmics T11 — FAIR R&D Data T5 — FAIR Data Lakes T11 — Protocol Intelligence
**Recording:** Purchase on-demand recording ↗

## Roche

• *8 sessions across T1, T5, T10, S3*

**259+ applications, 60% custom-built — and scientists say they can't find their data. Roche's external AI positioning vs. internal reality is a material gap.**

- **Scale of the problem (T10):** Chapter Lead for Data Products (pRED) disclosed 259+ applications — estimated real number 300–350, as their own inventory is incomplete. 60% custom-built. 80+ vendors. Mostly on-premise.
- **The candid quote:** Their new head of research, after 100 days, called it 'too many clunky systems.' Scientists report difficulty knowing where to store data and finding what colleagues have already created.
- **Data mesh (T1):** Roche is actively building data productisation in a data mesh — mid-process, not complete. They disclosed the true cost of drug development at $2.8B median per approval as the business case.
- **Navify (Diagnostics):** Roche Diagnostics is separately using GenAI for cancer patient journey tracking and clinical decision support — a different and more advanced maturity than pRED's R&D infrastructure.

QUESTION FOR YOUR CONTEXT

*Does Roche's disclosed internal data landscape resemble your environment? Note: their external AI positioning and internal infrastructure reality, as disclosed by their own teams, are materially different.*

**Evidence:** T10 — Too Many Clunky Systems T1 — Data Mesh Platform
**Recording:** Purchase on-demand recording ↗

## Johnson & Johnson

• *8 sessions across D3 Plenary, T3, T10, T11*

**Med.ai HealthLink tokenisation at 70–80% patient consent, CAR-T analytics, and a clear philosophy: AI is an architecture problem, not a model problem.**

- **Med.ai HealthLink (T11):** Multi-tenant tokenisation enabling real-world data linkage pre and post trial — up to 15 years of historical RWE linked to RCT data. 70–80% patient consent rate disclosed.
- **Architecture philosophy (T10):** J&J Technology's Head of Research Technology argued that the future is not a chat client — it is AI embedded into all scientific tool architecture. A differentiated position in the industry.
- **Modality-agnostic AI (D3):** In silico predictive models now integrated across the full discovery pipeline — small molecules, antibodies, RNA therapies, gene therapies, CAR-T. Not siloed by modality.
- **CAR-T analytics (S2):** Med.ai CAR-T: integrated manufacturing and analytics platform connecting lab data directly to predictive models for cell therapy programmes.

QUESTION FOR YOUR CONTEXT

*How does your clinical trial tokenisation and real-world data linkage compare to J&J's disclosed infrastructure? For cell and gene therapy: is there a comparable integrated data analytics platform?*

**Evidence:** D3 — ML & Data at Scale T11 — Med.ai HealthLink T10 — Architecture as Driver
**Recording:** Purchase on-demand recording ↗

## AstraZeneca • *6 sessions across T3, T2, S3*

**Foundation model data centre (1.1M public samples), data-centric regulatory submissions, and an enterprise-wide AI Accelerator with dedicated prompt upskilling.**

- **AIDC — AI Data Centre (T3):** Harmonises 1.1M samples from 19,000+ public studies via ETL pipelines resolving gene identifier differences across genome versions. Foundation models for oncology biomarker discovery.
- **The honest admission:** Presenters directly disclosed that organisations 'typically have petabytes of data but very few gigabytes that are AI-ready.' The AIDC is their attempt to close that gap — for public data only. Clinical trial data not included.
- **Regulatory submissions (T2):** Targeting elimination of manual data entry for Module 3 CTD — scientists still manually typing data from LIMS into Word. Active programme to change this.
- **AI Accelerator (S3):** Enterprise-wide initiative with a dedicated prompt upskilling lead conducting training across all teams. Privacy, IT, data security, and business functions involved.

QUESTION FOR YOUR CONTEXT
*How does your approach to making data AI-ready compare to AstraZeneca's AIDC? Does Pfizer have an equivalent structured GenAI upskilling initiative across functions?*

**Evidence:** T3 — Foundation Models / AIDC T2 — Data-Centric Submissions S3 — Prompting Mastery
**Recording:** Purchase on-demand recording ↗

## Sanofi • *S2*

**Publicly committed to in silico first biologics design by 2030, via BioAIM — while candidly acknowledging data fragmentation as their primary blocker.**

- **BioAIM platform:** An end-to-end AI biologics discovery platform for multi-specific antibodies. Three already in clinical phase. 2030 target: all biologics design begins computationally before any wet lab work.
- **The candid disclosure:** An Executive Director acknowledged directly that building a unified data foundation across an acquisition-heavy organisation remains their biggest challenge. Data fragmentation from M&A is the blocker.
- **The significance:** The 2030 in silico first commitment is now on the public record. This is a direction, not a current state — but it is a public benchmark against which Sanofi will be measured.

QUESTION FOR YOUR CONTEXT

*How does your biologics AI timeline compare to Sanofi's publicly committed 2030 target? Is the acquisition-driven data fragmentation challenge they described familiar from your own environment?*

**Evidence:** S2 — Biologics AI Moonshot / BioAIM
**Recording:** Purchase on-demand recording ↗

## Alnylam Pharmaceuticals • *T5 + T6*

**Cloud genomics platform for precision medicines at scale — plus Deloitte industry data showing R&D lags all functions in GenAI ROI despite being the largest long-term opportunity.**

- **Cloud genomics / REVEAL (T5):** Platform built with Paradigm4 for large-scale human genetics supporting RNAi drug design. Accessing datasets including Our Future Health — targeting 5M+ individuals, 10% of the UK population.
- **The Deloitte industry finding (T6):** 2,700 respondents across 6 industries: C-suite leaders consistently rate their GenAI programmes as more advanced than the implementation teams closest to the work.
- **R&D vs. other functions:** R&D is the function with the least current GenAI ROI despite being the largest long-term opportunity. IT and cybersecurity are seeing the most immediate returns.
- **The value projection:** $5–7B in GenAI value potential for large biopharma over 5 years (Deloitte modelled projection, not observed results).

QUESTION FOR YOUR CONTEXT
*For RNA therapeutics programmes: how does your cloud genomics infrastructure compare to what Alnylam disclosed? Does the C-suite vs. implementation team perception gap resonate internally?*

**Evidence:** T5 — Cloud Genetics / REVEAL T6 — AI Transformation Blueprint
**Recording:** Purchase on-demand recording ↗

# A2 — EMERGING THEMES THAT WILL SHAPE 2025–2026

## Conference-Wide Patterns — Multiple Independent Sources

*Each theme emerged from multiple independent sessions — not from a single company.*

| Theme | Strategic implication | Evidence from BIO-IT 2025 |
|---|---|---|
| **AI-designed drugs are in the clinic** | The Insilico Phase 2A disclosure moves AI drug design from hypothesis to clinical evidence. The competitive question is who builds the most capable platform now. | *Insilico Medicine — D2 + D3 Plenaries* |
| **GenAI has crossed from pilot to production** | Velocity is now the competitive differentiator — not the fact of having AI. Companies presenting at BIO-IT 2025 were reporting systems, not plans. | *Novartis S3+T6; BMS T6; AbbVie T11; Takeda T11* |
| **FAIR data is the AI infrastructure layer** | The 80% data wrangling problem was cited identically by multiple independent presenters. Companies that solve this first will have a structural advantage within 2–3 years. | *Takeda T11+T5; AstraZeneca T3; Roche T10; Snowflake D2* |
| **Pre-competitive data collaboration is formalising** | AbbVie's AISB Consortium demonstrates that IP barriers to pre-competitive AI collaboration are solvable through federated learning. The shift from informal to formal legal structures is the significant move. | *AbbVie / Columbia / Apheris T5* |
| **Knowledge graphs + LLMs converging into agent architecture** | BRAG, SLAG, MAG are the next architecture layer beyond RAG. AI infrastructure built today without accounting for agent-native design will need rebuilding within 18–24 months. | *NVIDIA / AbbVie / BMS T10+S1* |
| **Active learning is closing the wet-dry lab loop** | Novartis and Dyno Therapeutics both disclosed ML-guided experiment selection in production — the AI is now deciding which experiments to run next. | *Novartis T7; Dyno Therapeutics T1* |
| **Responsible AI governance is unclaimed** | Across all 229 sessions, zero companies presented a formal published governance framework. An open leadership position. | *Gap confirmed across all 229 sessions* |

# A3 — VENDOR WATCH: POST-CONFERENCE ASSESSMENT

## Vendors That Made Substantive Disclosures at BIO-IT 2025

*Assessed on what was disclosed on stage — not on vendor positioning or marketing materials.*

| Vendor | Action | Assessment based on BIO-IT 2025 disclosures |
|---|---|---|
| **Snowflake** | **Evaluate** | • Delivered the D2 Plenary on survivorship bias in AI model training. A second T6 session with Trinity Life Sciences disclosed Snowflake Cortex AI and Cortex-I with RAG for natural language patient cohort development and protocol acceleration.<br>• Their 8-pillar enterprise AI framework was referenced independently by multiple attendees across sessions. |
| **Apheris** | **Evaluate** | • The federated learning infrastructure behind the AISB Consortium — training inside each company's firewall, only model updates shared. AbbVie's on-stage endorsement is the strongest third-party validation at the conference.<br>• Relevant for any pre-competitive molecular or structural biology data collaboration evaluation. |
| **ServiceNow + NVIDIA** | **Evaluate** | • Disclosed Lab in a Loop — joint laboratory operations automation. Specific ROI: 25% reduction in development costs and 500-day faster time to market from lab standardisation.<br>• NVIDIA has partnered with ServiceNow for 5+ years on domain-specific language models for lab workflows. |
| **ZS Associates** | **Assess** | • Scaling partner for BMS DocGen AI (T6) and data infrastructure partner for AstraZeneca AIDC (T3).<br>• Two independent large pharma engagements at the same conference. Their approach to pharma AI delivery is now documented in both session transcripts. |
| **Insilico Medicine** | **Watch** | • Phase 2A success for a fully AI-designed drug. Their platform (Pandomics, Chemistry 42, Cliniko) spans target discovery to clinical prediction.<br>• Both a competitive intelligence priority and a potential partnership candidate — CEO explicitly discussed licensing and collaboration models on stage. |
| **Domino Data Lab** | **Monitor** | • Their T3 session on agentic AI framed the concept as a junior data scientist that must be trained with organisational knowledge, tools, and examples — a practical and well-received framework.<br>• Domino is the dominant pharma data science platform; their agentic AI roadmap is relevant for any data science infrastructure evaluation. |

# A4 — RECOMMENDED FOLLOW-UPS

## Priority Contacts — Within Two Weeks of the Conference

*All roles below presented at BIO-IT 2025. Personal names are not included — roles and companies are sufficient for identification through conference connections.*

| Role | Company | Why reach out — and what to ask | Urgency |
|---|---|---|---|
| **CEO** | Apheris | • Issued an on-stage invitation for new pharma members to join the AISB Consortium.<br>• Ask about governance and IP terms for new members, what data contribution looks like for Pfizer's scale, and the AISB-1 results timeline. | Priority |
| **Head of Computational Drug Discovery** | AbbVie | • Co-founder of the AISB Consortium.<br>• Ask how the legal structure was established, what the early challenges were, and whether the model is genuinely open to new members. Peer conversation — not a vendor engagement. | Priority |
| **Head of AI Data & Analytics** | BMS | • Presented DocGen AI architecture in detail.<br>• Ask about the ZS Associates partnership model, Word add-in adoption strategy, which document types were most complex to configure, and what they would do differently. | Seek Out |
| **Director of Data Science** | Novartis NIBR | • Built NSA and explicitly invited peer feedback.<br>• Ask about access control design for sensitive project data, BM25 vs. semantic search balance in practice, and adoption curve in the first 3 months. | Seek Out |
| **CEO** | Insilico Medicine | • Presented Phase 2A success on stage. Both a competitive intelligence conversation and a potential partnership exploration.<br>• Ask about platform licensing model and what pharma data they would need to improve model performance. | Seek Out |
| **Director & Research Professor** | Tufts CSDD | • Leads the pharma AI adoption benchmark study (302 respondents, 2024).<br>• Ask about Pfizer-sponsored custom benchmarking engagement. Follow-up study with case examples noted for later in 2025. | Connect |