AI & APRA Compliance

A Comprehensive Guide for Superannuation Funds

Essential Reading for:

Chief Risk Officers • Chief Operating Officers
Heads of Compliance • Technology Directors
Trustee Board Members

BackPro AI
December 2025

www.backpro.ai

Contents

1 Executive Summary

Australian superannuation funds face increasing operational complexity: APRA reporting requirements grow more detailed, member expectations for instant service escalate, and cost-per-member pressure intensifies. AI automation offers a solution, but implementation must navigate strict APRA prudential standards.

This whitepaper addresses the critical question: *How can super funds deploy AI while maintaining complete APRA compliance?*

1.1 Key Findings

- **Cloud AI services violate SPS 234:** ChatGPT, Claude, and Gemini send member data offshore, creating immediate data sovereignty violations
- **On-premise deployment maintains compliance:** AI deployed within your Australian infrastructure preserves SPS 234 alignment
- **Software licence avoids material outsourcing:** Proper structuring prevents SPS 231 service provider registration requirements
- **70% time savings achievable:** APRA quarterly returns, board papers, and breach notifications automated while maintaining audit trails

2 APRA Regulatory Landscape

2.1 Relevant Prudential Standards

2.1.1 CPS 230: Operational Risk Management (Cross-Industry)

Applies to all APRA-regulated entities. Requires:

- Board accountability for operational risk
- Three lines of defence model
- Risk appetite statements
- Material risk identification and controls

AI Implications: AI deployment creates operational risks (accuracy, availability, dependency) requiring board awareness and documented controls.

2.1.2 SPS 234: Data Sovereignty (Superannuation)

Specific to RSE licensees. Requires:

- Information security controls protecting member data
- Data stored and processed within Australia (or with APRA notification)
- Third-party service provider oversight
- Incident notification within 72 hours

AI Implications: Cloud AI services sending member data to US servers violate data sovereignty requirements. On-premise deployment necessary.

2.1.3 SPS 515: Operational Risk (Superannuation)

Requires operational risk management framework including:

- Risk identification processes
- Control effectiveness monitoring
- Scenario analysis and stress testing
- Risk culture and governance

AI Implications: AI errors, hallucinations, or availability issues constitute operational risks requiring controls and monitoring.

2.1.4 SPS 231: Outsourcing (Superannuation)

Defines material service providers and requires:

- Due diligence before engagement
- Contractual protections (audit rights, termination clauses)
- Ongoing monitoring
- APRA notification for material arrangements

AI Implications: Cloud AI services would be material outsourcing. On-premise software licence avoids this classification.

2.1.5 SPS 232: Business Continuity (Superannuation)

Requires:

- Business continuity plans covering critical processes
- Disaster recovery capabilities
- Regular testing (at least annually)
- Recovery time objectives (RTOs)

AI Implications: AI-dependent processes must have documented recovery procedures and failover capabilities.

2.1.6 CPS 234: Information Security (Cross-Industry)

Requires:

- Information asset register
- Security controls aligned to asset sensitivity
- Penetration testing
- Incident response procedures

AI Implications: AI systems must be included in information security framework with appropriate access controls and vulnerability management.

2.2 APRA's Technology Risk Expectations

APRA's Prudential Practice Guide CPG 234 provides guidance on information security, emphasizing:

1. **Board oversight:** Board understands technology risks and approves risk appetite
2. **Three lines of defence:**
   - First line: Business units using AI
   - Second line: Risk and compliance oversight
   - Third line: Internal audit assurance
3. **Vendor due diligence:** Assess vendor security controls, financial stability, and exit strategies
4. **Change management:** Controlled deployment with testing and rollback procedures

3 The Cloud AI Compliance Gap

3.1 Why ChatGPT Violates SPS 234

Data Flow:

1. Staff member enters member query: "What is John Smith's super balance?"
2. Query transmitted to OpenAI servers in USA
3. OpenAI processes query, potentially logging for model improvement
4. Response returned to staff member

SPS 234 Violations:

- Member data transmitted offshore without APRA notification
- No contractual control over OpenAI's data handling
- Member data potentially used for model training
- No audit trail of which member data was processed

3.2 Why Claude & Gemini Have Same Issues

All major cloud AI services share fundamental problems:

- **Offshore processing:** US or global data centres
- **API architecture:** Data must leave your environment
- **Training data concerns:** Queries may inform future models
- **No Australian residency:** Cannot guarantee data stays onshore

3.3 Material Outsourcing Concerns

Under SPS 231, cloud AI services would likely qualify as material because:

1. **Operational dependence:** If service disrupted, AI-dependent processes stop
2. **Member data access:** Service provider processes member information
3. **Concentration risk:** Single vendor failure impacts operations

APRA Requirements: Would need to notify APRA, conduct extensive due diligence, establish audit rights, and monitor vendor ongoing.

4 On-Premise AI Architecture

4.1 How On-Premise AI Works

4.1.1 Deployment Model

BackPro AI installs entirely within your infrastructure:

1. **Infrastructure:** Deploy on Azure Australia East, AWS ap-southeast-2 (Sydney), or GCP australia-southeast1
2. **Compute:** Runs on your virtual machines or containers (AKS, ECS, GKE)
3. **Storage:** Uses your database and file storage (Azure Blob, S3, GCS)
4. **Network:** Operates within your VNet/VPC, no internet egress required

4.1.2 Data Access Pattern

1. Staff member asks: "What is John Smith's super balance?"
2. BackPro (running in your infrastructure) queries your admin system API
3. Admin system returns member data within your environment
4. BackPro processes query and generates response locally
5. Response displayed to staff member
6. **Zero external transmission:** No data leaves your infrastructure

4.2 SPS 234 Compliance Maintenance

4.2.1 Data Sovereignty

$$ All member data processed in Australian data centres

No API calls to offshore services

Network logs prove zero external transmission

Can air-gap deployment if required for highest security

4.2.2 Information Security Controls

Your existing controls apply:

- **Access management:** BackPro authenticates via your Azure AD/Entra ID
- **Authorization:** Role-based access controls (RBAC) enforce permissions
- **Encryption:** Data encrypted at rest (your KMS) and in transit (TLS 1.3)
- **Audit logging:** All queries logged to your SIEM (Splunk, Sentinel, etc.)

4.2.3 Third-Party Oversight

BackPro vendor never accesses your member data:

- Software deployed by your IT team
- Configuration managed by your administrators
- Vendor provides software updates (not operational services)
- Support provided via screen sharing (not direct access)

4.3 Avoiding Material Outsourcing

4.3.1 Software Licence vs. Service Provider

BackPro Model: Software license

- You purchase software license
- Deploy in your infrastructure
- Operate and maintain deployment
- Vendor provides updates and support
- No vendor access to your data

Similar to: Microsoft Office, Adobe Acrobat, Salesforce (self-hosted)
Classification: Not a material service provider under SPS 231

4.3.2 Operational Independence

If BackPro (vendor) ceased operations:

- Your deployed instance continues functioning
- No operational dependency on vendor
- You own the deployment and configuration
- Can maintain software internally if needed

5 Operational Risk Framework

5.1 Risk Identification

AI deployment introduces specific operational risks:

5.1.1 Accuracy Risk

Risk: AI generates incorrect information
Example: APRA quarterly return contains wrong member count
Controls:

- Mandatory human review before submission
- Validation checks against source systems
- Error correction workflows
- Accuracy metric tracking (monthly)

5.1.2 Availability Risk

Risk: AI system unavailable when needed
Example: APRA quarterly return deadline approaching, AI offline
Controls:

- Fallback to manual processes documented
- High availability deployment (multi-zone)
- Disaster recovery tested annually
- Service level monitoring

5.1.3 Dependency Risk

Risk: Staff become over-reliant on AI, lose manual process capability
Example: Staff cannot prepare board papers without AI assistance
Controls:

- Maintain documentation of manual procedures
- Quarterly manual process testing
- Staff training on both AI and manual methods
- Knowledge retention programs

5.2 Control Framework

5.2.1 Preventive Controls

1. **Input validation:** Verify queries are appropriate and complete
2. **Output review:** Human checks AI-generated content before use
3. **Access restrictions:** Only authorized staff can use AI
4. **Use case limits:** Define which processes suitable for AI

5.2.2 Detective Controls

1. **Accuracy monitoring:** Track error rates and correction frequency
2. **Audit log review:** Monthly review of AI usage patterns
3. **Escalation tracking:** Monitor cases requiring human judgment
4. **User feedback:** Staff report quality issues

5.2.3 Corrective Controls

1. **Error correction:** Documented process to fix AI mistakes
2. **Root cause analysis:** Investigate recurring error patterns
3. **Model refinement:** Improve AI accuracy based on corrections
4. **Process adjustment:** Update workflows if issues persist

5.3 Risk Appetite Alignment

Board must define acceptable AI risk levels:

- **Accuracy tolerance:** What error rate is acceptable? (e.g., ¡2%)
- **Use case boundaries:** Which processes suitable for AI? (e.g., routine queries yes, investment decisions no)
- **Review requirements:** When is human oversight mandatory? (e.g., APRA submissions always)
- **Escalation triggers:** When should AI defer to human? (e.g., complex scenarios, policy gaps)

6 Practical Use Cases

6.1 APRA Quarterly Returns (SRF 600, 700, 800 Series)

6.1.1 Current Manual Process

1. Extract data from administration platform (2 days)
2. Compile member statistics across multiple systems (3 days)
3. Validate data accuracy and reconcile discrepancies (3 days)
4. Format to APRA submission standards (2 days)
5. Executive review and sign-off (2 days)
6. **Total: 10-12 days**

6.1.2 AI-Assisted Process

1. AI queries administration platform and financial systems (30 minutes)
2. AI compiles data and formats to SRF standards (30 minutes)
3. AI runs validation checks against previous quarters (15 minutes)
4. **Human review:** Compliance officer validates output (4 hours)
5. Executive sign-off (1 day)
6. **Total: 1.5-2 days**

Time Saved: 70-80% reduction in preparation time

6.2 Trustee Board Papers

6.2.1 Current Manual Process

1. Gather data from multiple sources (1 day)
2. Draft investment committee report (2 days)
3. Compile risk assessment (1 day)
4. Format and proofread (4 hours)
5. Executive review cycles (1 day)
6. **Total: 5 days**

6.2.2 AI-Assisted Process

1. AI compiles data from investment, risk, and compliance systems (15 minutes)
2. AI drafts board paper using templates and previous papers (30 minutes)
3. **Human review:** Executive edits and refines content (3 hours)
4. Final approval (4 hours)
5. **Total: 1 day**

Time Saved: 60

6.3 Policy & Procedure Management

6.3.1 Current Manual Process

1. Monitor APRA prudential standard updates manually (ongoing, ad-hoc)
2. Assess impact on fund policies (1-2 weeks when changes occur)
3. Draft policy updates (1 week)
4. Review cycle with legal and compliance (1 week)
5. **Total: 3-4 weeks per regulatory change**

6.3.2 AI-Assisted Process

1. AI monitors APRA website and legislative registers automatically
2. AI identifies relevant changes and assesses impact on policies (1 hour)
3. AI drafts policy updates with track changes (2 hours)
4. **Human review:** Legal and compliance validate changes (1 week)
5. **Total: 1-1.5 weeks**

Time Saved: 50

6.4 Breach Notification Preparation

6.4.1 Current Manual Process

1. Compile incident details from multiple sources (4 hours)
2. Draft root cause analysis (1 day)
3. Document remediation plan (1 day)
4. Format to APRA notification standards (4 hours)
5. Executive review (4 hours)
6. **Total: 3 days**

6.4.2 AI-Assisted Process

1. AI compiles incident data from logs and reports (15 minutes)
2. AI drafts breach notification with root cause (30 minutes)
3. **Human review:** Risk officer validates and refines (4 hours)
4. Executive sign-off (2 hours)
5. **Total: 1 day**

Time Saved: 70

7 Implementation Approach

7.1 Phase 1: Technical Deployment (Week 1-2)

7.1.1 Infrastructure Setup

1. Provision Azure/AWS/GCP resources in Australian region
2. Configure networking (VNet, subnets, security groups)
3. Set up storage (databases, file storage)
4. Deploy BackPro software

7.1.2 Integration

1. Connect to administration platform API (Class, Link, Bravura, etc.)
2. Integrate with member records systems
3. Link to financial and investment systems
4. Configure Azure AD/Entra ID authentication

7.1.3 Security Configuration

1. Implement encryption (KMS, TLS)
2. Configure access controls (RBAC)
3. Set up audit logging to SIEM
4. Enable network monitoring

7.2 Phase 2: Testing & Validation (Week 3-4)

7.2.1 Functional Testing

1. Test APRA reporting data extraction
2. Validate board paper generation
3. Test policy update workflows
4. Verify breach notification formatting

7.2.2 Accuracy Validation

1. Compare AI outputs to manually prepared versions
2. Identify and correct accuracy issues
3. Establish baseline accuracy metrics
4. Define acceptable error thresholds

7.2.3 Security Testing

1. Verify zero external network calls
2. Test access controls and permissions
3. Review audit logs for completeness
4. Conduct basic penetration testing

7.3 Phase 3: Pilot Deployment (Week 5-8)

7.3.1 Limited Rollout

1. Start with one use case (e.g., board papers)
2. Limit to 2-3 experienced staff members
3. Mandatory human review of all outputs
4. Daily monitoring of accuracy and issues

7.3.2 Refinement

1. Collect staff feedback on usability
2. Adjust AI prompts and templates
3. Improve integration with source systems
4. Optimize performance

7.4 Phase 4: Full Production (Week 9+)

7.4.1 Expanded Rollout

1. Add additional use cases (APRA reporting, policy management)
2. Train all relevant staff
3. Implement ongoing monitoring dashboards
4. Establish reporting to risk committee

7.4.2 Governance Framework

1. Document AI governance policy
2. Define roles and responsibilities
3. Establish escalation procedures
4. Schedule periodic audits

8 Ongoing Monitoring & Governance

8.1 Monthly Metrics

Track these key performance indicators:

- **Accuracy rate:** % of AI outputs requiring no corrections
- **Escalation volume:** Number of queries AI cannot handle
- **Time savings:** Hours saved vs. manual processes
- **Error corrections:** Types and frequency of mistakes
- **User satisfaction:** Staff feedback scores

8.2 Quarterly Risk Committee Reporting

Report to board risk committee:

1. AI usage statistics (queries, documents, time saved)
2. Accuracy trends and error analysis
3. Control effectiveness assessment
4. Incidents or near-misses
5. Emerging risks or concerns

8.3 Annual Internal Audit

Include in audit plan:

- Review governance documentation
- Test control effectiveness
- Validate audit trail completeness
- Assess compliance with APRA standards
- Interview staff on AI usage and concerns

9 APRA Examination Preparedness

9.1 Documentation to Maintain

Keep readily accessible for APRA prudential reviews:

1. **Governance framework:** AI policy, roles, approval processes
2. **Risk assessment:** Operational risks identified, controls implemented
3. **Technical architecture:** Deployment diagram showing data flows
4. **Compliance mapping:** How deployment aligns with SPS 234, SPS 515, CPS 234
5. **Audit trails:** Sample exports showing source attribution and approvals
6. **Testing results:** Accuracy validation, security testing, DR testing
7. **Monitoring reports:** Monthly metrics, quarterly risk reports
8. **Incident register:** Any AI errors or security incidents

9.2 Demonstrating Compliance

Be prepared to show APRA:

- **Data sovereignty:** Network logs proving zero offshore transmission
- **Access controls:** Authentication logs, permission matrices
- **Human oversight:** Approval workflows, review documentation
- **Control effectiveness:** Error rates, correction cycles, accuracy trends
- **Board awareness:** Risk committee minutes, board reports

10 Conclusion

APRA compliance does not preclude AI adoption in Australian superannuation funds. With proper architecture (on-premise deployment), governance (operational risk controls), and oversight (audit trails and human review), funds can achieve 70% time savings on compliance and operational processes while maintaining regulatory alignment.

The key is rejecting cloud AI services that violate SPS 234 data sovereignty in favor of on-premise deployments that keep member data within your controlled environment.

10.1 Next Steps

1. **Technical assessment:** Review deployment architecture with your infrastructure team (45 minutes)
2. **Risk evaluation:** Document operational risks and controls with risk team (1-2 hours)
3. **Security validation:** Walk through access controls and encryption with CISO (45 minutes)
4. **Pilot deployment:** 4-week pilot in non-production environment to validate compliance
5. **Board approval:** Present to board risk committee for deployment authorization

Contact

BackPro AI
Email: backpro@backpro.ai
Web: www.backpro.ai

Schedule a 30-minute compliance walkthrough to demonstrate SPS 234 controls, operational risk framework, and APRA examination readiness.