A P R I T Y

Aprity — Service Reliability & Availability

Service Reliability & Availability

PUBLIC

Version 1.0 | March 2026 | aprity.ai

1. Executive Summary

Aprity is designed for **enterprise reliability**. The platform's architecture ensures that scans can survive failures, recover automatically, and produce consistent results even under adverse conditions. This document describes the reliability mechanisms built into the platform.

2. Platform Availability

2.1 Infrastructure

Component	SLA	Provider
Azure Container Apps	99.95%	Microsoft Azure
Azure Cosmos DB	99.99%	Microsoft Azure
Azure Blob Storage	99.9% (Hot tier)	Microsoft Azure
Azure Key Vault	99.99%	Microsoft Azure

2.2 Aprity Platform

Metric	Target
API availability	99.9%
Scan completion rate	> 99% (with auto-retry)
Planned maintenance windows	Communicated 48h in advance
Incident notification	Within 4 hours

3. Scan Reliability

3.1 Idempotent Processing

Every step of the scan pipeline is **idempotent** — running the same operation twice produces the same result without side effects. This means:

- **Safe to retry**: If a step fails, it can be re-executed without corrupting data
- **Safe to resume**: A failed scan can restart from the last successful step
- **No duplicates**: Deterministic IDs prevent duplicate work items

3.2 Automatic Retry

Failure Type	Behavior
Network timeout	Automatic retry with exponential backoff
AI provider error	Retry with backoff; failover to secondary provider

Salesforce API limit	Pause and retry after cooldown period
Transient database error	Automatic retry (Cosmos DB built-in)
Permanent error	Marked as terminal; scan continues with partial results

3.3 Distributed Lease Management

To prevent duplicate processing in a multi-worker environment:

- Each work item acquires a **time-limited lease** before execution
- If a worker crashes, the lease expires and another worker picks up the item
- **Lease reclaim cycle**: 45 seconds (orphaned items detected quickly)
- No zombie processes — expired leases are automatically reclaimed

3.4 Circuit Breaker

When an AI provider experiences repeated failures:

Condition	Action
5 consecutive failures	Pause for 60 seconds
Provider still unavailable	Failover to secondary provider
All providers unavailable	Scan paused; auto-resume when available

4. Data Durability

4.1 Database

Feature	Guarantee
Storage	Azure Cosmos DB with automatic replication
Consistency	Session consistency (strong within a session)
Backup	Continuous backup with point-in-time restore
Retention	Configurable (7-30 days for backups)

4.2 File Storage

Feature	Guarantee
Storage	Azure Blob Storage with versioning enabled
Durability	99.999999999% (11 nines)
Redundancy	Locally redundant (LRS) minimum; Geo-redundant (GRS) available
Versioning	Immutable — previous versions retained for audit

4.3 Secrets

Feature	Guarantee
Storage	Azure Key Vault (HSM-backed)
Soft delete	Enabled (accidental deletion recoverable)
Purge protection	Enabled (mandatory retention period)

5. Scan Performance

5.1 Typical Processing Times

Org Size	Metadata Items	Estimated Scan Time
Small (< 100 objects)	< 500	10-20 minutes
Medium (100-500 objects)	500-5,000	30-90 minutes
Large (500+ objects)	5,000-50,000	2-6 hours
Very Large (1,000+ objects)	50,000+	6-12 hours

> *Times depend on: Salesforce API response times, AI provider throughput, and metadata complexity.*

5.2 Optimization Features

- **AI analysis cache**: Unchanged metadata reuses previous analysis (cross-scan)
- **Batch processing**: Metadata items analyzed in optimized batches
- **Parallel extraction**: Multiple metadata types fetched concurrently
- **Incremental analysis**: Only changed items are re-analyzed (when supported)

6. Monitoring & Alerting

6.1 Real-Time Monitoring

Metric	Monitoring
API response time	Azure Application Insights
Worker queue depth	KEDA metrics
Database request units	Cosmos DB metrics
AI provider latency	Application telemetry
Error rates	Structured logging + alerts

6.2 Health Checks

- **Liveness probe**: Verifies the application is running
- **Readiness probe**: Verifies the application can process requests
- **Dependency checks**: Database, Key Vault, and storage connectivity
- **Endpoint**: `/actuator/health` (standard Spring Boot)

6.3 Scan Progress

Customers can monitor scan progress in real-time through:

- **Salesforce UI**: Progress bar and phase indicators in the Aprity component
- **Admin Console**: Detailed work item status for each scan
- **API**: Programmatic access to scan status (`GET /v1/scans/{id}`)

---

7. Disaster Recovery

7.1 Recovery Objectives

Metric	Target
RTO (Recovery Time Objective)	< 4 hours
RPO (Recovery Point Objective)	Near-zero (continuous backup)

7.2 Scenarios

Scenario	Recovery
Single container failure	Auto-restart by Container Apps (seconds)
Worker crash during scan	Lease expiration → another worker resumes
AI provider outage	Automatic failover to secondary provider
Database failure	Azure Cosmos DB automatic failover
Region failure	Re-deployment to secondary region (manual)

7.3 Scan Recovery

Because all work items are idempotent and resumable:

1. Failed scans can be **restarted** from the admin console
2. The scan resumes from the **last successful work item** — no work is repeated
3. Partial results are preserved — the scan does not start from scratch
4. Customers are notified of any scan interruption via the Salesforce UI

---

8. Maintenance & Updates

8.1 Platform Updates

Update Type	Impact	Notification
Security patches	Zero downtime (rolling deployment)	Changelog published
Feature releases	Zero downtime (rolling deployment)	Release notes + changelog
Infrastructure changes	Planned window (if needed)	48h advance notice
Salesforce package updates	Admin-initiated install	Release notes provided

8.2 Backward Compatibility

- API versioning ensures existing integrations continue to work
- Salesforce package updates are backward compatible
- Database schema migrations are non-breaking

9. Support

Channel	Availability
Email	support@aprity.ai (business hours)
Admin Console	Self-service scan management
Documentation	Online knowledge base
Incident status	Status page (URL provided to customers)

Contact

For reliability or SLA questions:

Email: support@aprity.ai

Website: https://aprity.ai

This document describes the reliability architecture of the Aprity platform as of March 2026. Specific SLA terms are defined in individual customer contracts.