Xtensyon.
WHITE PAPER

Identity Sync for Permission-Aware Retrieval: Getting ACLs Right in Practice

December 3, 2025 • Xtensyon Labs • 9 min read

> *Permission-aware retrieval fails when group memberships drift or metadata is inconsistent. This paper covers a practical identity sync and ACL strategy that keeps retrieval correct without slowing teams down.*

TL;DR

- Sync identity and groups on a schedule you can defend in audits.
- Store ACLs with documents at ingestion, and re-check at query time.
- Use stable group IDs, not display names, to avoid silent mismatches.
- Treat access bugs as security incidents, not search issues.

Executive Summary

Enterprises want assistants that know who can see what. The hard part is not the UI. It is identity data: groups, role changes, exceptions, and legacy systems. We outline an ACL approach that works with real identity providers. The system ingests ACL metadata, enforces it during retrieval, and keeps a clear sync policy. It also has a testing strategy that catches permission regressions before release.

Why It Matters

If a user can retrieve and see content they should not access, the incident is already serious. Even a single leak can block a rollout. Permission-aware retrieval is a security requirement and a trust requirement. Good ACL handling also improves UX because results do not include documents that look relevant but are inaccessible.

What We Built

- A group sync service that stores stable IDs and maps them to document ACLs.
- Index-time ACL tagging plus query-time enforcement based on the requesting user.
- A permission test suite that covers role changes, group removals, and edge cases.
- Operational dashboards for sync freshness, failed syncs, and retrieval deny rates.

Observed Outcomes

- Fewer permission incidents after adding both index-time and query-time checks.
- Cleaner debugging when deny decisions were logged with codes and reasons.
- Faster rollouts because security sign-off had concrete evidence and tests.

Implementation Notes

- Do not sync only daily if offboarding needs to be effective within hours.
- Validate identity inputs. Directory data is messy in the real world.
- Prefer deny-by-default for unknown groups or missing ACL metadata.
- Keep audit logs separate from analytics logs so retention rules are easier.

Governance & Risk

- Define a sync SLA and monitor it. Stale identity data is a risk.
- Use least privilege for sync connectors and store credentials securely.
- Document exception handling for legacy folders and shared mailboxes.

Release Checklist

- Is there a defined identity sync schedule and SLA?
- Are ACLs stored with documents and enforced at retrieval?
- Do we log deny decisions with reasons?
- Are permission regressions covered by tests?

- Is unknown or missing ACL metadata handled safely?

Conclusion

Permission-aware retrieval is not optional in enterprise settings. When identity sync is treated as infrastructure, retrieval becomes both safer and more predictable for users.

Keywords

access control acl identity sync rag security enterprise search