

KNOWLEDGE ARCHITECTURE · ENTERPRISE IT

Consolidating 10,000+ engineering knowledge items into one governed SharePoint platform.

10,000+

Content and data items consolidated

4 → 1

Scattered platforms reduced to one environment

18

Months, phased engagement

1

Governed SharePoint knowledge platform

CLIENT TYPE

Large technology organization

ENVIRONMENT

Engineering and infrastructure support

ENGAGEMENT

18-month phased engagement

SCOPE

10,000+ content and data items

STARTING POINT

Knowledge scattered across roughly four platforms

FINAL STATE

One structured SharePoint knowledge environment

CORE WORK

Content inventory, taxonomy design, metadata model, SharePoint calendar views, and governance structure

Summary

A large technology organization had engineering knowledge scattered across wiki pages, Confluence, SharePoint folders, standalone documents, and embedded calendar references. Engineers had documented valuable information, but it was difficult to find, inconsistent in structure, and spread across too many locations.

I was brought in to evaluate the existing documentation environment and create a more usable knowledge structure. The result was a centralized SharePoint platform that consolidated 10,000+ items into one governed environment organized around taxonomy, metadata, structured views, and operational use.

The problem

The issue was not a lack of documentation. The issue was that the documentation did not function as a system.

Engineering and infrastructure support information lived in too many places. Some content was stored in wiki pages. Some lived in Confluence. Some was already in SharePoint, but without a consistent structure. Some existed as standalone files. Time-sensitive operational details were also embedded inside static documents instead of being managed through calendar functionality.

This made knowledge access dependent on familiarity. Engineers and support teams had to know where information lived before they could find it. That created friction for onboarding, audits, support response, release readiness, and day-to-day infrastructure operations.

The approach

I approached the work as an information architecture and knowledge management initiative, not a basic content migration.

The first step was understanding the content landscape: what existed, where it lived, how it was used, and which information needed to be surfaced by audience, system, function, or operational need.

The work included:

- Inventorying 10,000+ content and data items across multiple repositories.
- Identifying duplicate, outdated, misplaced, and inconsistently structured information.
- Grouping content by system, support function, team, document type, and operational use case.
- Designing a taxonomy and metadata model to support filtering, retrieval, and governance.
- Moving calendar-driven information out of static documents and into functional SharePoint calendar views.
- Creating a platform structure that made SharePoint the primary access point for engineering knowledge.

A key friction point was that some operational information was not really “documentation” in the traditional sense. Calendar items, support references, and recurring operational details were buried inside documents and pages. Treating all of that as flat content would have recreated the same problem in a new tool.

The better solution was to separate content types and use SharePoint functionality intentionally: documents for reference material, calendars for time-based information, metadata for retrieval, and views for role- or function-specific access.

The solution

The final platform was designed as a structured SharePoint knowledge environment, not a document dump. Reference materials, calendar-driven information, operational files, and support content were organized according to how teams actually needed to use them.

The design used SharePoint functionality intentionally: documents for reference material, calendars for time-based operational information, metadata for retrieval, and views for audience- or function-specific access. This created a governance foundation that made the platform easier to maintain beyond the initial consolidation effort.

What changed

More than 10,000 content and data items were consolidated into a single SharePoint-based knowledge environment.

Engineering and infrastructure support teams gained one primary starting point instead of having to search across multiple disconnected platforms.

Calendar-driven operational information was moved out of static documents and into calendar views, making time-based support information easier to see and manage.

The new structure improved onboarding, audit readiness, support response, and release readiness by giving teams clearer access to operational knowledge.

The platform reduced dependence on tribal knowledge by applying taxonomy, metadata, and structured navigation to information that had previously been scattered across tools and individual documentation habits.

The organization gained a stronger foundation for knowledge access, operational continuity, and future governance.

TAKEAWAY

When documentation is scattered across tools, teams, and individual habits, the solution is not just to move content into a new repository.

The real value comes from creating structure: inventory, taxonomy, metadata, ownership, access paths, and governance. That is what turns documentation from a collection of files into an operational knowledge system.