

Structured Knowledge as Shared Infrastructure

Experience from CMS Computing Operations

Jason Mohoney · MIT CSAIL

On behalf of the Archi team

Team: MIT Data Systems Group and CMS

Principal investigators



Christoph Paus



Tim Kraska



Funding: TWG Global and Intersystems,
through the MIT Generative AI
Impact Consortium



Mariarosaria D'Alfonso · Julius Heitkoetter · Luca Lavezzo · Pietro Lugato
Dmytro Kovalskyi · Warren Nam · Christoph Paus

Tim Kraska · Michael Stonebraker · Samuel Madden · Jason Mohoney · Xinjing Zhou

Muhammad Hassan Ahmed · Antonio Battaglia · Hasan Ozturk · Viphava Ohm · Krittin Phornsiricharoenphant

Juan Pablo Salas Galindo · Aron Soha

Samantha Sunnarborg · Gabriele Benelli

Austin Swinney · Manasvita Joshi

CMS / Physics @ MIT

Data Systems Group, MIT CSAIL

CMS @ CERN

CMS @ Fermilab

CMS @ Brown

Computing @ Harvard



The problem

The data of a large organization or experiment is scattered across many systems. It is heterogeneous and multi-modal, it keeps changing, parts of it are stale or conflicting, and it is **hard to stitch together.**

This talk

A **single structured knowledge base** can serve many applications in an organization.

1. An assistant for CMS computing operations
2. Why it needs structured knowledge
3. Knowledge bases as shared infrastructure

CMS computing operations

*“turning the raw output of the CMS detector into usable scientific data,
and making that data available to physicists worldwide”*

Tier-0


repack, two tape copies, prompt reconstruction

Production & Reprocessing

Monte Carlo and reprocessing campaigns

Data Management


transfers and placement across the sites



Most questions touch several of these systems.

Archi

Our assistant for CMS Computing Operations. In production since February 2026.





What are the most severe issues that occurred in the past year that affected multiple sites or involved multiple systems



CMS comp-ops graph

✗ Explored 13 okg-cms_inspect, 6 okg-cms_search, okg-cms_filter, okg-cms_expand

For the year ending 2026-09-30 (2025-09-30 through 2026-09-30), the strongest candidates are the incidents with broad service impact, long outages, or substantial potential data impact. The graph labels these downtime records as **outages**, but does not provide a common severity ranking, so this is a qualitative assessment.

- **KIT-T2 and FZK-LCG2 storage/network outage — 2026-01-22, 11:30–13:33 UTC.** The matching downtime records describe the same network-related storage outage at two sites. Both affected XRootD and WebDAV; FZK-LCG2 also lists the XRootD redirector and tape WebDAV. This is the clearest multi-site event found. (OKG: `downtime:37475` , `downtime:37476`)
- **LRZ-LMU computing-centre shutdown — 2025-10-16 to 2025-10-17.** Water ingress led to an emergency shutdown of the whole computing centre for about 28 hours. The record lists ten affected services, including compute, storage, WebDAV, and monitoring/accounting services. (OKG: `downtime:37199`)
- **IEPSAS-Kosice power cut — 2025-10-25.** An unexpected electrical cut caused an approximately 12-hour outage spanning 13 services, including compute, storage, XRootD, and WebDAV. (OKG: `downtime:37239`)
- **CSCS-LCG2 cluster outage — 2025-10-22.** "Cluster down" for about 10 hours; the record lists 11 services affected, including compute, storage, XRootD, and WebDAV. (OKG: `downtime:37215`)
- **MIT disk data-loss incident — January 2026.** `CMSTRANSF-1298` reports **799,428 missing files** following accidental deletions at T2_US_MIT. The issue says subsequent analysis suggested the files were recoverable, so this is a high-potential-impact incident rather than confirmed permanent loss; JIRA priority was Medium. (OKG: `jira:CMSTRANSF-1298`)

A further cross-system issue was WMCore's Rucio client failing to authenticate after a service-certificate renewal in July 2026, affecting the Global WorkQueue LocationUpdateTask and other WMCore services. It is marked Minor, so it ranks below the outages and data-loss concern above. (OKG: `jira:CMSTRANSF-1365`)

Evidence generation: `oka:cms:branch:default:aen:20260930T140718024468Z:81f8c9831fb0` .


Send a Message



CMS comp-ops graph




Privacy



Ticket text is anonymised at ingest: names, emails, and usernames are removed before anything is stored. Open-weight models run on our own GPUs, so nothing has to leave CERN. An API model can be used instead.

Evaluation by CMS operators

63 questions from real traffic, graded blind by CMS operators on a 1-5 scale for usefulness and correctness



A Archi's responses are useful

GPT-5.5 scores 4.07 for usefulness and 4.25 for correctness.

B Open-source models are competitive

Qwen trails GPT-5.5 by 0.35-0.38 points.

C Non-agentic retrieval is obsolete

The same model gains 1.29 in correctness and 1.39 in usefulness with the agent loop.

Pieces that are needed for a good retrieval system

Coverage, structure, and attribution with time.

Coverage

All of the relevant records

ticket history

workflow configuration

dataset placement

site topology

Structure

Explicit links among those records

workflow → requires → dataset

dataset → available at → site

workflow → allowed at → site

Attribution + temporal history

Where a claim came from and when it was true

source record + revision


observed at

valid from / valid to


supersedes

Data is represented as a temporal knowledge graph

Data items are represented as nodes, and the relations between them as edges




Incremental Ingestion Pipeline



Knowledge graph vs. vector database

42 CMS operator questions · % of required facts found · langgraph agent + GPT-6 Luna



CMS Knowledge Base

To be released to CMS on 19 October 2026, behind CERN SSO.

35

sources

3.2 M

records

76,444

tickets

43,746

wiki pages

1.8 M

datasets

5,314


analysis notes

dev build, 1 September 2026

Contains analysis notes, sites, datasets, workflows, tickets, code, software releases, and the relations between them


Each application rebuilds, replicates and maintains the same knowledge

We observed this problem when building agents for CMS



This is inefficient and difficult to maintain

Shared knowledge base for downstream applications



A single shared knowledge base makes downstream applications consistent and easier to maintain

Knowledge base infrastructure across application domains

We have ten deployments that use the same underlying infrastructure for building and maintaining knowledge bases

HEP related applications at CMS and MIT

CMS computing operations

CMS data quality monitoring

with Fermilab

CMS physics analysis

CMS CRAB

MIT SubMIT cluster

Other domains

statistical genetics

UCLA · published

healthcare records

InterSystems

astronomy literature

University of Arizona

financial records and documents

FinTechAI at CSAIL

research computing

Harvard FASRC

Conclusions

Problem The data of a large organization or experiment is scattered across many systems, heterogeneous, changing, partly stale or conflicting, and hard to stitch together.

Proposed solution A single structured knowledge base maintained and used by multiple downstream applications.

Knowledge Bases and Agents: workshop at CERN, 19 to 23 October 2026

Workshop indico.cern.ch/event/1716863

Code github.com/archi-physics/archi

Slack join.slack.com/t/a2rchi/shared_invite/zt-3ski3n4bb-24LFO~5adtdzmGNUa2vTXg