A Composable Scientific Data Management Architecture

Yu Ma, Randall Bramley ([yuma,bramley]@cs.indiana.edu)
Indiana University Computer Science

Scientific Data Management

- Applications becoming dominated by moving, finding, archiving large data
- Users are concerned with longevity
  - Robustness and longevity two primary concerns in recent cyberinfrastructure analysis at IU
  - Rightfully or not, they believe open source provides it
- Need more automation of data management and metadata creation

Scientific Data Libraries

- Data model defines structures that that I/O libraries can directly understand/manipulate
- Scientific data libraries manage complex data structures directly and record metadata such as type, size, shape, numerical format, etc.

Some I/O systems and their data models

I/O system	Data models
UNIX	**Sequence of bytes**
netCDF (network Common Data Form)	**Annotated multi-dimensional arrays of typed elements**
HDF5 (Hierarchical Data Format)	**Annotated multi-dimensional arrays of multi-element records Hierarchical groups of objects.**

High Level Data Models

- Pros
  - Data structures closely match what scientific applications use;
  - Self-contained data files with metadata stored in addition to the basic datasets.
- Cons
  - May be too specific to a class of applications, thus not useful in other application areas;
  - Translations between high-level abstraction and low-level storage system model are required, some may be inefficient, and increase the data transfer cost.

Metadata Management Systems

- Metadata makes other data useful
- Separating metadata from described datasets has proven utility
- The earliest and still most commonly used technology: file naming conventions.
- Existing general-purpose metadata management systems include SRB/MCAT, OGSA-DAI, MCS, SAM, Chimera, etc.

Data Grids

- Grid: a set of services for configuring, launching, monitoring, controlling work
- A data grid provides an architecture of managing and analyzing large-scale, shared, and widely distributed datasets.
- Data grids focus on:
  – Secure and efficient data transfer
  – Metadata services
- Data grid efforts include: SDSS, ESG, European DataGrid, FusionGrid, GriPhyN, PPDG, etc.

SRB/MCAT

- SRB is a client-server middleware that provides distributed clients with an uniform interface to access heterogeneous data storage resources.
- MCAT is a metadata repository for SRB resource discovery.
- MCAT system-level metadata categories include Data Object, Resource,Collection, User, Method

SRB/MCAT

- MCAT has introduced several concepts of metadata categorization for apps. However:
  – a universal metadata schema is more useful at abstract level than at practical level;
  – MCAT has rather limited support for application-level metadata;
  – as with any integrated software system, MCAT has implementation limitations;
  – SRB/MCAT is not an open source software

OGSA-DAI

- Provides uniform service interfaces to access data from distributed sources via the Grid
- Promotes standards for grid database services, initially focusing on consistent access to existing autonomously managed databases
- Does not create new data storage/management systems, but makes them more readily usable within a Grid framework
- Can be considered as a grid transport layer of relational and XML database queries, mainly concerned with large scale integration/federation

Obsidian Goals

- Unattended data collection
- Multiple storage levels and hierarchies
- Large size data objects
- Composable modules
  - Each part does one thing, but does it well
- Ability to define user-specified metadata schemas, instead of one size fits all approach
- Open source components with proven robustness
- Support for Unix and Windows platforms

Modules Implied by Requirements

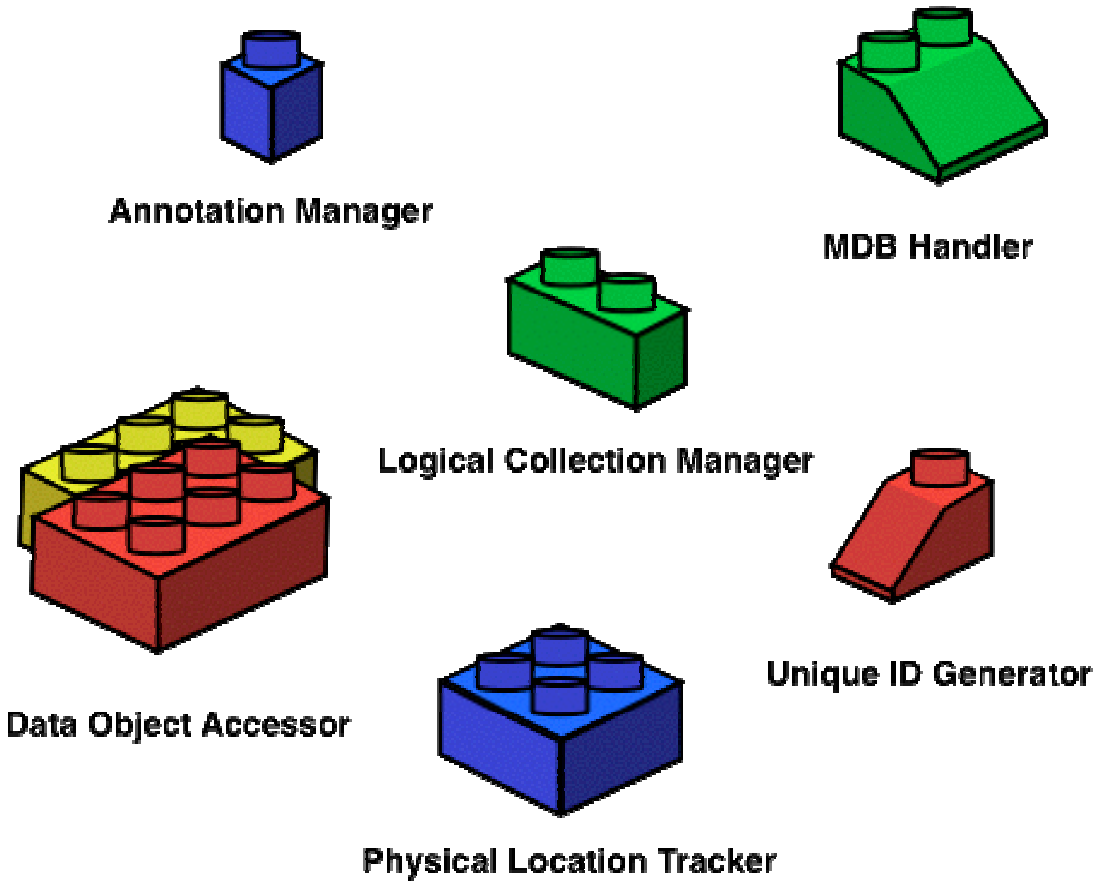

Architecture Overview

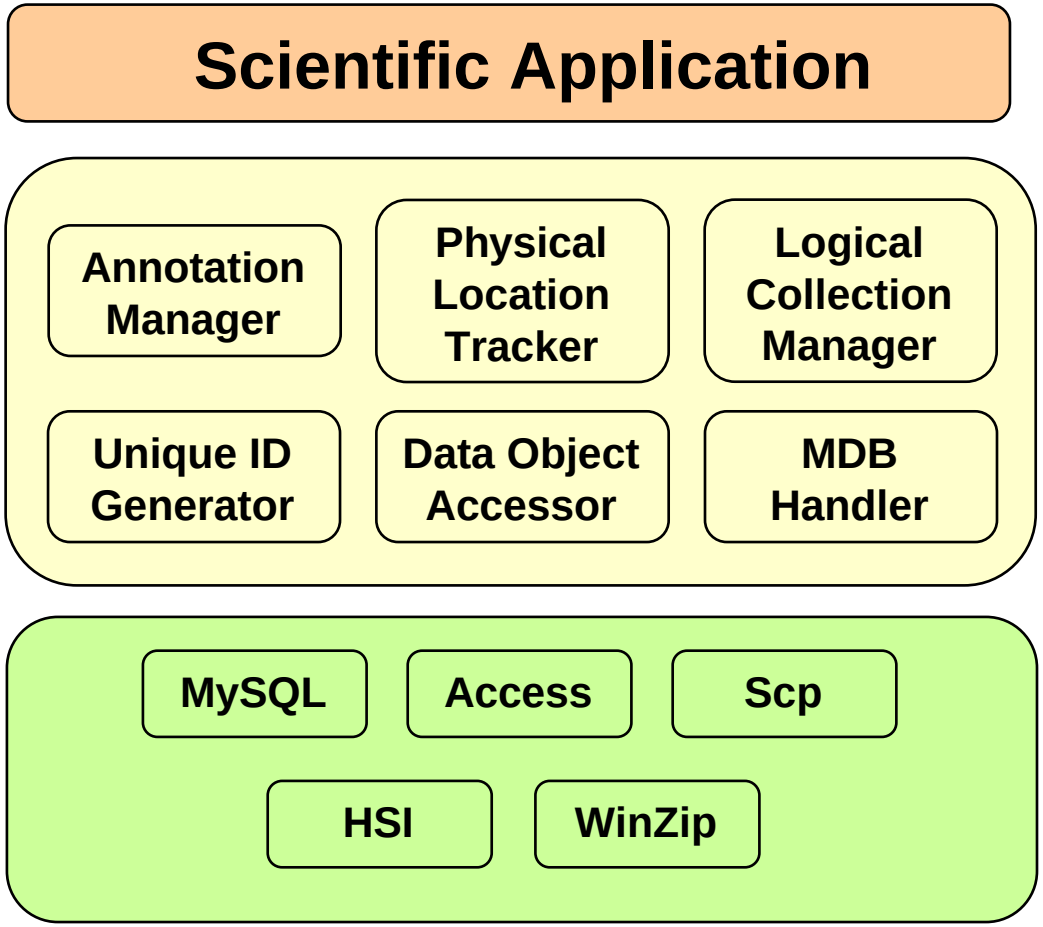

Obsidian

- Supports each needed module with at least one implementation
- Has been used for collaborations in
  - bioinformatics
  - x-ray crystallography
  - astronomy
  - clinical radiation therapy

Application 1: Platcom

- Platcom is an integrated comparative genome analysis system developed at IUB School of Informatics.
- Building and updating a pairwise comparison database require over 48,000 jobs initially, and a few hundred more monthly afterwards.
- Data management needs:
  - interface with batch queuing system PBS;
  - automatic and secure transfer of large amount of computation input/output files;
  - archive history and performance metadata;
  - failed job discovery and re-submission.

Application 1: Platcom

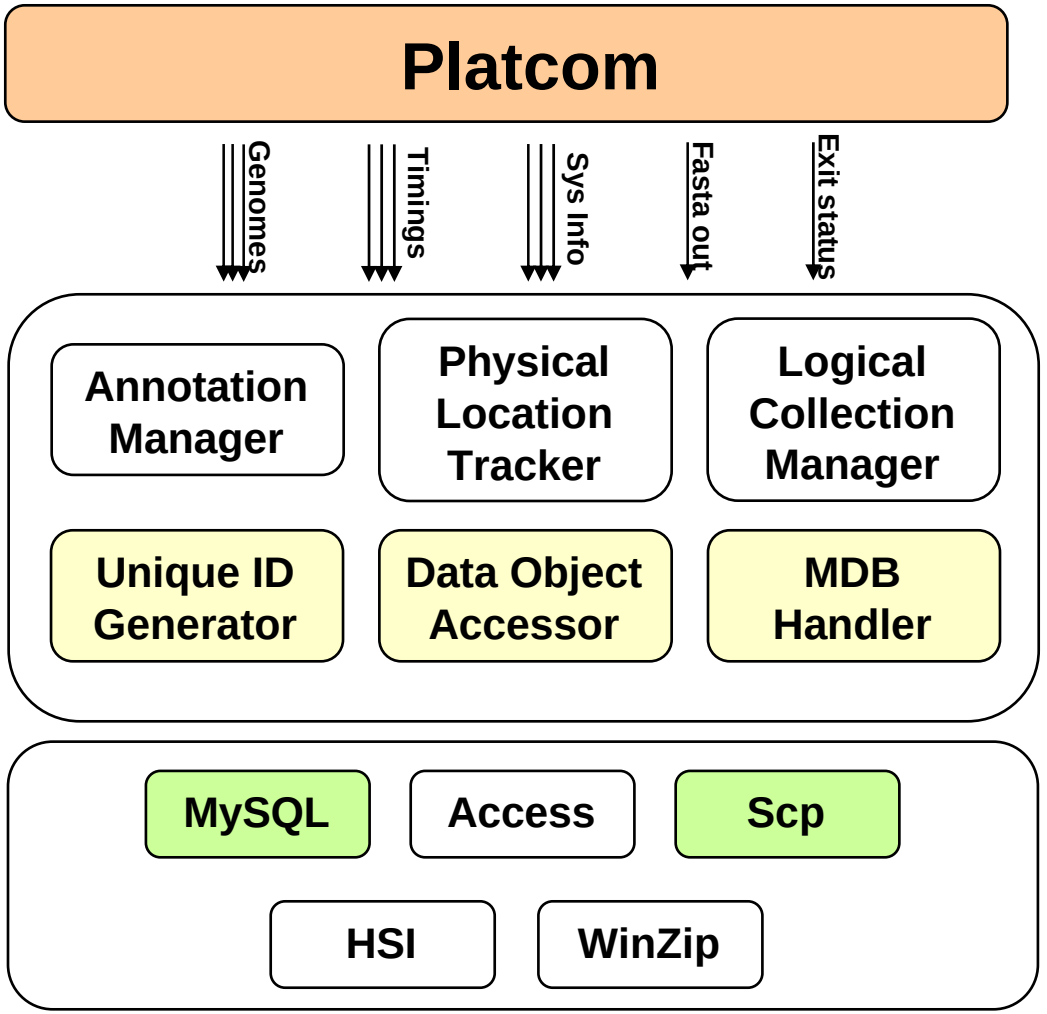

Application 2: X-Ray Crystallography

- IU Molecular Structure Center remotely manages X-ray crystallography data collections of large molecule samples
- To date over 2 terabytes of data; extending to include earlier data from past decade involves multiple petabytes
- Data management needs:
  - track physical locations of different kinds of data objects, including CCD frames and lab camera images;
  - archive metadata about environment under which the datasets are collected;
  - archive metadata about sample providers and instrument operators;
  - move data objects and metadata among front end collectors, staging computers, tape systems like HPSS

Application 2: X-Ray Crystallography

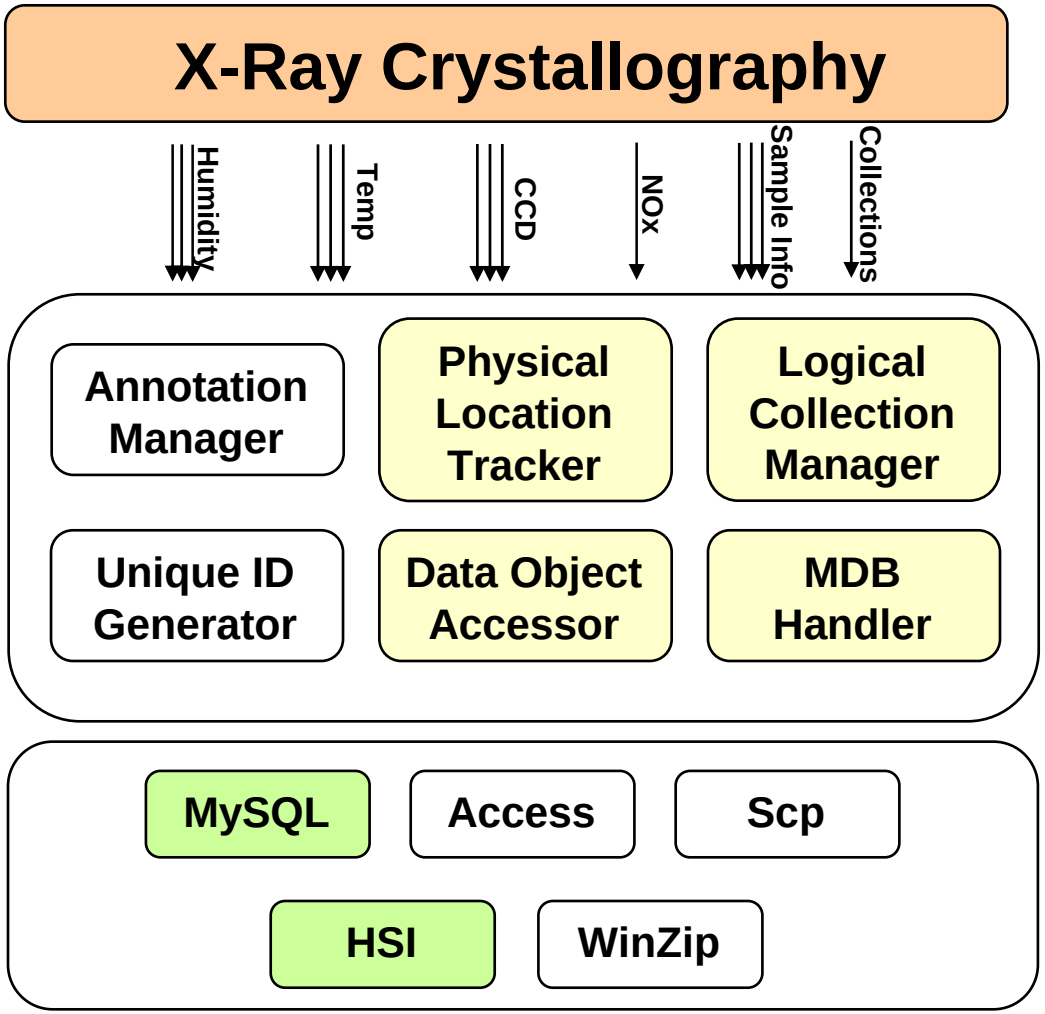

Application 3: Radiation Therapy

- In radiation therapy of cancers, large amount of data are produced from patient pretreatment imaging scans
- Collaboration with one of the clinical practices at Radiation Oncology Department, University of Maryland School of Medicine
- Data management needs:
  - Automatically locate a patient's treatment files, whether on hard disk or tape, through a few variations of queries;
  - Automatically build a MS Access database of specified treatment parameter combinations and populate it with existing patient data to facilitate clinical trial designs.

Application 3: Radiation Therapy

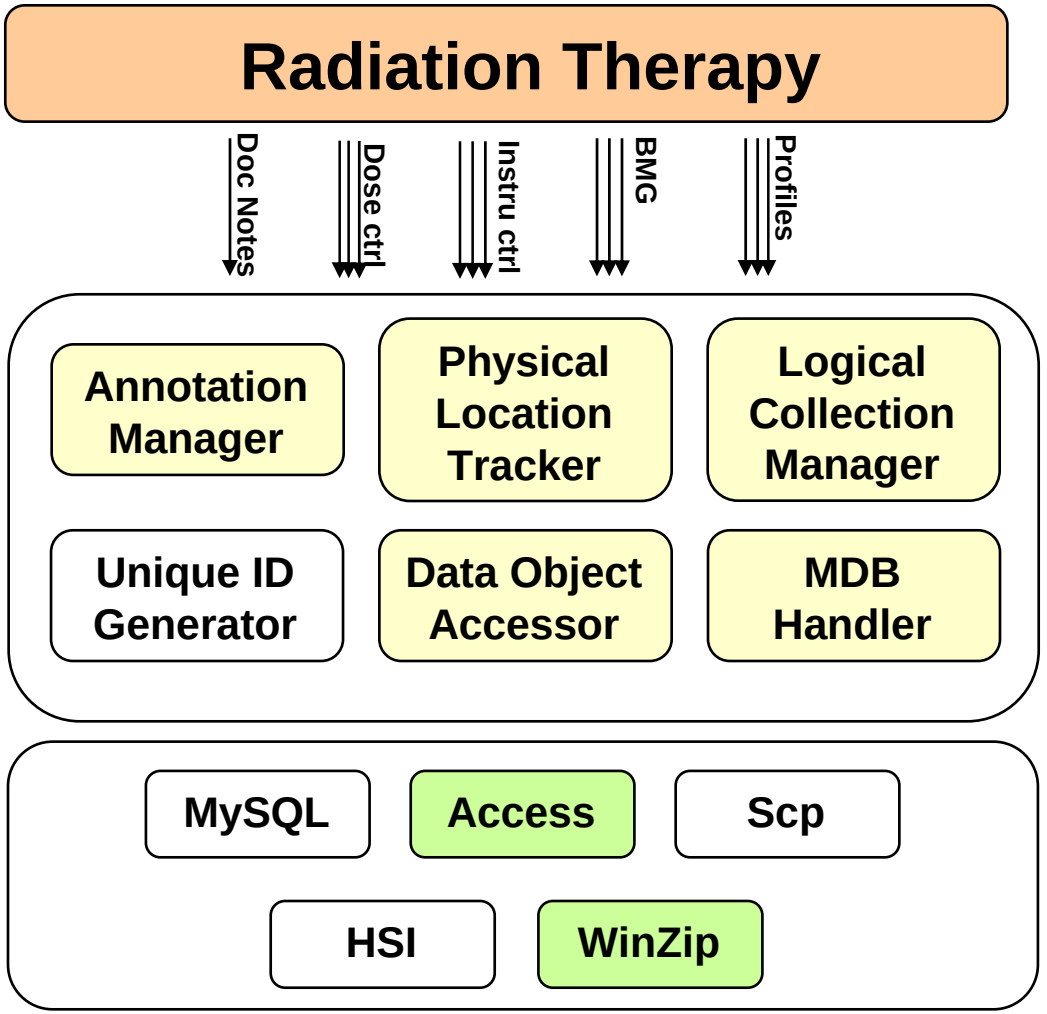

Application 4: Automated Photometry

- Facilitate managing astronomical data collected from Morgan-Monroe Station of the Goethe Link Observatory telescopes.
- Data are produced from a wide range of instruments including:
  - CCD images in FITS format
  - Thermometer readings and precipitation records
  - Weighing scale readings for liquid NOx
  - Telescope and dome control parameters
- Data management needs:
  - Data direct from A2D cards on instruments
  - Transfer across low-bandwidth connections
  - Support scientific queries involving large amounts of complex computation

Application 4: Automated Photometry

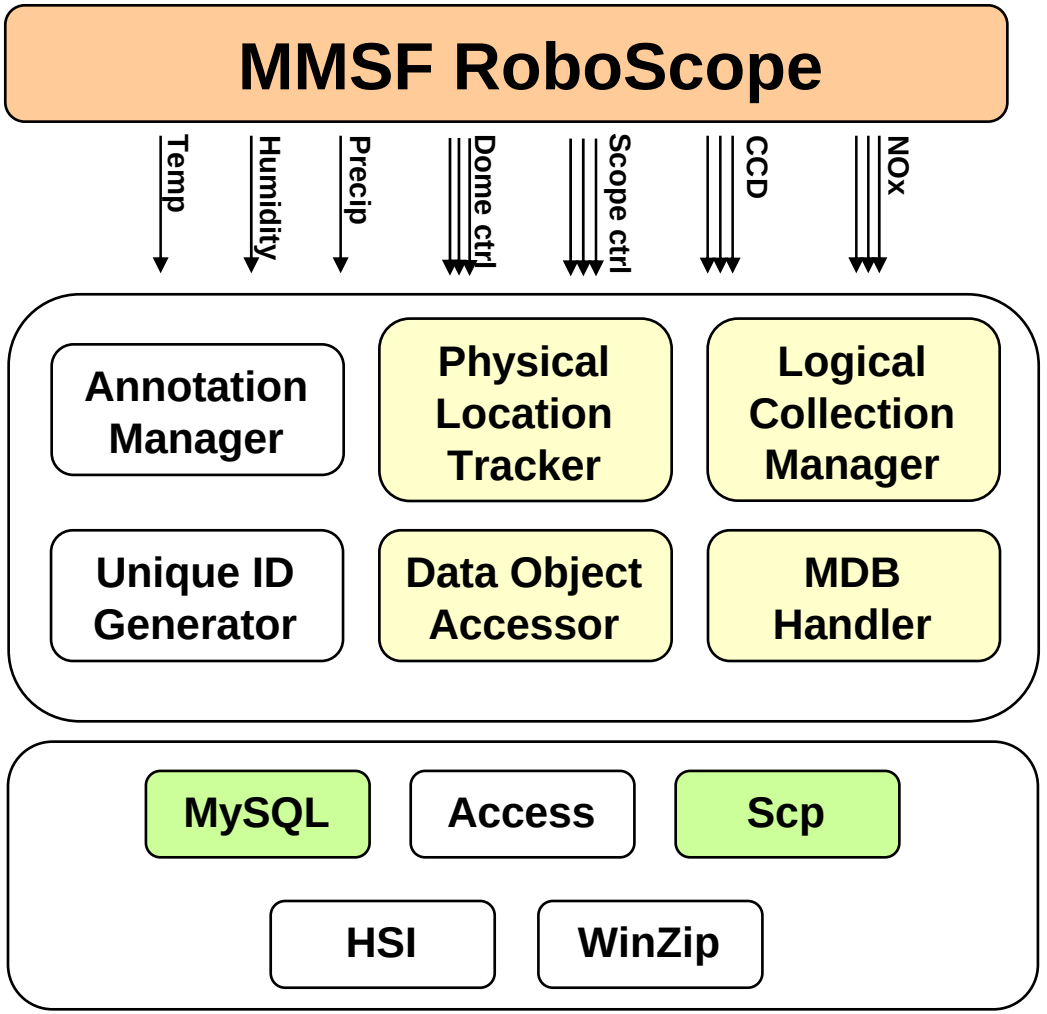

Summary

- Obsidian’s modular approach to creating data management systems allows using only what is needed for particular application
- End-user defined schema are superior to putting user metadata into DB text fields
- Does not replace need for application communities to define ontologies, agree upon terms, definitions, and interfaces