

Exploring NCAR's Campaign Store with Elasticsearch

How much can we know about 120 PiB in a summer?

Anh Nguyen

Mount Holyoke College

Mentors: Nathan Hook, Eric Nienhouse, Jason Cuning

July 30, 2024



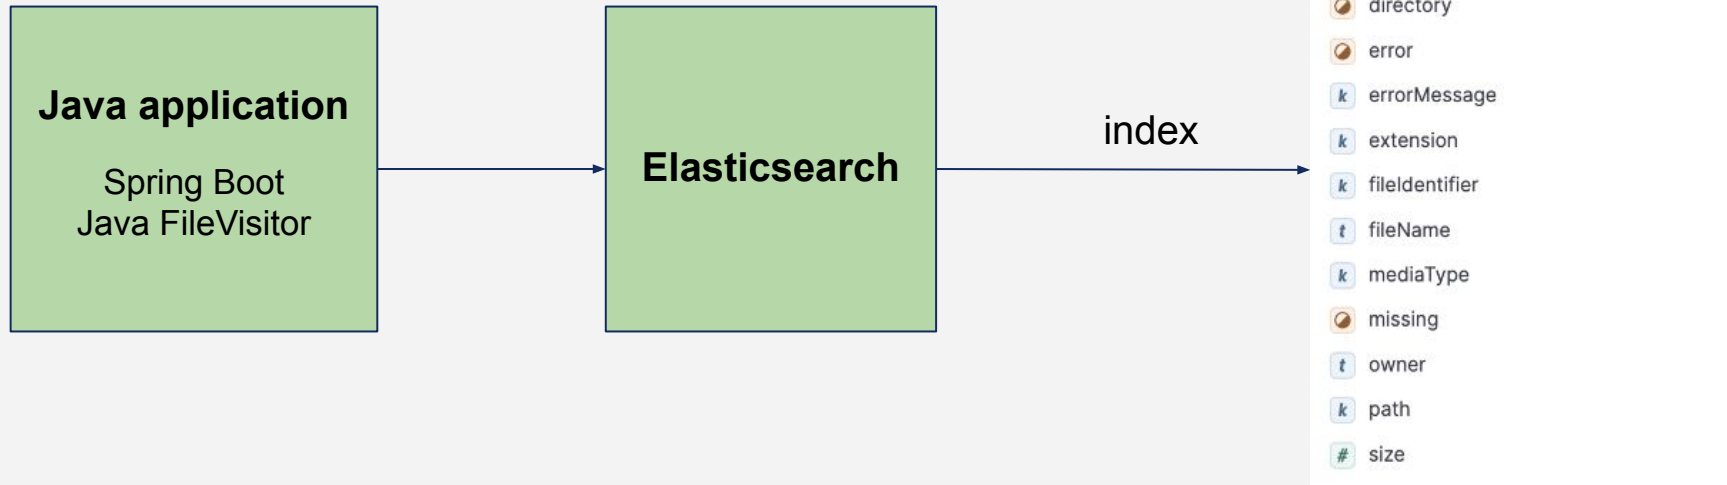
Background

- **NCAR Campaign Store:**
 - Resource for medium-term storage of project data, typically for three to five years, by NCAR labs and universities
 - **Large volume:**
 - ~120 pebibyte capacity (1 PiB ~ 1 million GB)
 - 90% currently used (~4 billion files)
- **Central question:**
 - How do we facilitate data searching within Campaign Store for NCAR scientists?



Background

- **Application:**
 - Java-based web application that interfaces with Elasticsearch to index files and directories



File Walkers State

Start Directory	Ignored Directories	Start Time	Total Time (ms)	Files	Errored Files	Directories	Errored Directories	Other Errors	Running
/glade/campaign/cesm	[/glade/campaign/cesm/development/espwg/SMYLE/archive/]	Mon Jul 22 13:20:23 MDT 2024	596828045	2,622,994	0	360,399	214	1,246	true
/glade/campaign/collections	[/glade/campaign/collections/cdg/data/CMIP6/, /glade/campaign/collections/cdg/data/cmip5/]	Mon Jul 22 13:20:23 MDT 2024	596828046	58,616,412	0	1,303,766	582	344	true
/glade/campaign/mmm	[]	Mon Jul 22 13:20:23 MDT 2024	596828046	196,677,980	0	2,415,254	1,653	2,954	true
/glade/campaign/acom	[]	Mon Jul 22 13:20:23 MDT 2024	596828045	6,176,839	0	162,112	0	0	true

Refresh

Run File Walkers



The challenge of time and volume

- Current number of files on campaign store: 4,088,158,704

Name	Health	Status	Primaries	Replicas	Docs count	Storage size
file-walker-files	● green	open	7	0	499,936,380	140.17gb

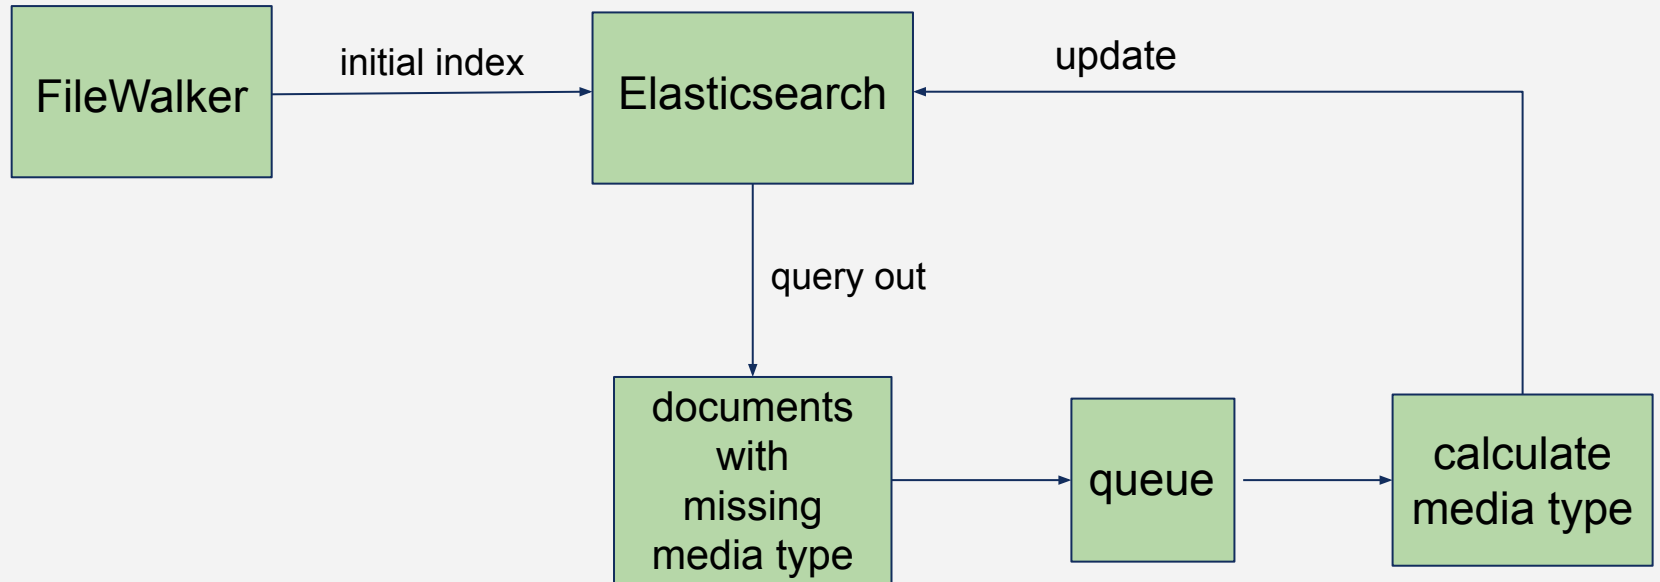
The restart problem

- How do we keep track of the file walker's progress to continue where we stopped?

completed	id
<input type="checkbox"/> [file:///glade/campaign/acom/archive/otserver/jamesw/mlo/cof2/A/20060531.033854/, file:///glade/campaign/acom/archive/otserver/jamesw/mlo/cof2/A/20170418.174420/, file:///glade/campaign/acom/archive/otserver/jamesw/mlo/cof2/A/20170407.002501/, ...	acom
<input type="checkbox"/> [file:///glade/campaign/cesm/development/cross-wg/S2S/jaye/SPPT/archive/cesm2cam6.2015-11-09.09/, file:///glade/campaign/cesm/development/cross-wg/S2S/jaye/SPPT/archive/cesm2cam6_spptv2.2004-01- 12.10/, file:///glade/campaign/cesm/development/cross-wg/S2S/jaye/SPPT/archive/cesm2cam6.2015-11-...	cesm
<input type="checkbox"/> [file:///glade/campaign/cisl/fs1_p_relocation/nwc/model_testing_env/, file:///glade/campaign/cisl/fs1_p_relocation/nwc/jamesmcc/, file:///glade/campaign/cisl/aiml/, file:///glade/campaign/cisl/fs1_p_relocation/nwc/katelynw/nwm_v3/groundwater/nwm_groundwater_los...	cisl
<input type="checkbox"/> [file:///glade/campaign/collections/cmip/CMIP6/timeseries- cmip6/f.e11.FAMIPC5CN_RCP85.f09_f09.AERBMBemis_1920sst.009/, file:///glade/campaign/collections/cmip/CMIP6/timeseries-...	collections
<input type="checkbox"/> [file:///glade/campaign/ral/nsap/rkumar/NASA_HMA/GSI_24mar/2008020209/, file:///glade/campaign/ral/nsap/rkumar/NASA_HMA/GSI_24mar/2008020206/, file:///glade/campaign/ral/nsap/rkumar/NASA_HMA/GSI_24mar/2005032712/, ...	ral
<input type="checkbox"/> [file:///glade/campaign/mmm/parc/schwartz/JTTI/precip_verif_MRMS_daily_pcent/, file:///glade/campaign/mmm/parc/schwartz/JTTI/exp_hybrid_single_res_3km_static_BECS/, file:///glade/campaign/mmm/parc/schwartz/JTTI/be/, ...	mmm

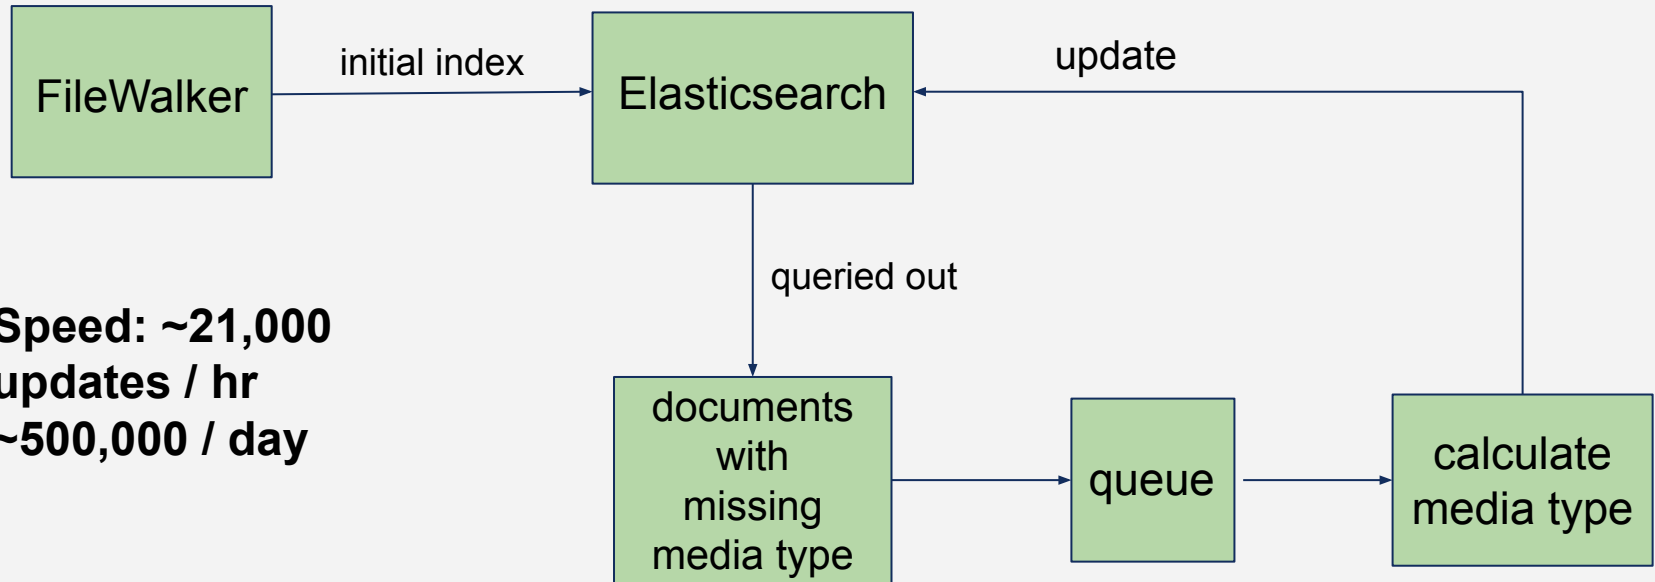
The metadata calculation problem

- Tika: used to calculate media type
- Current Workflow:



The metadata calculation problem

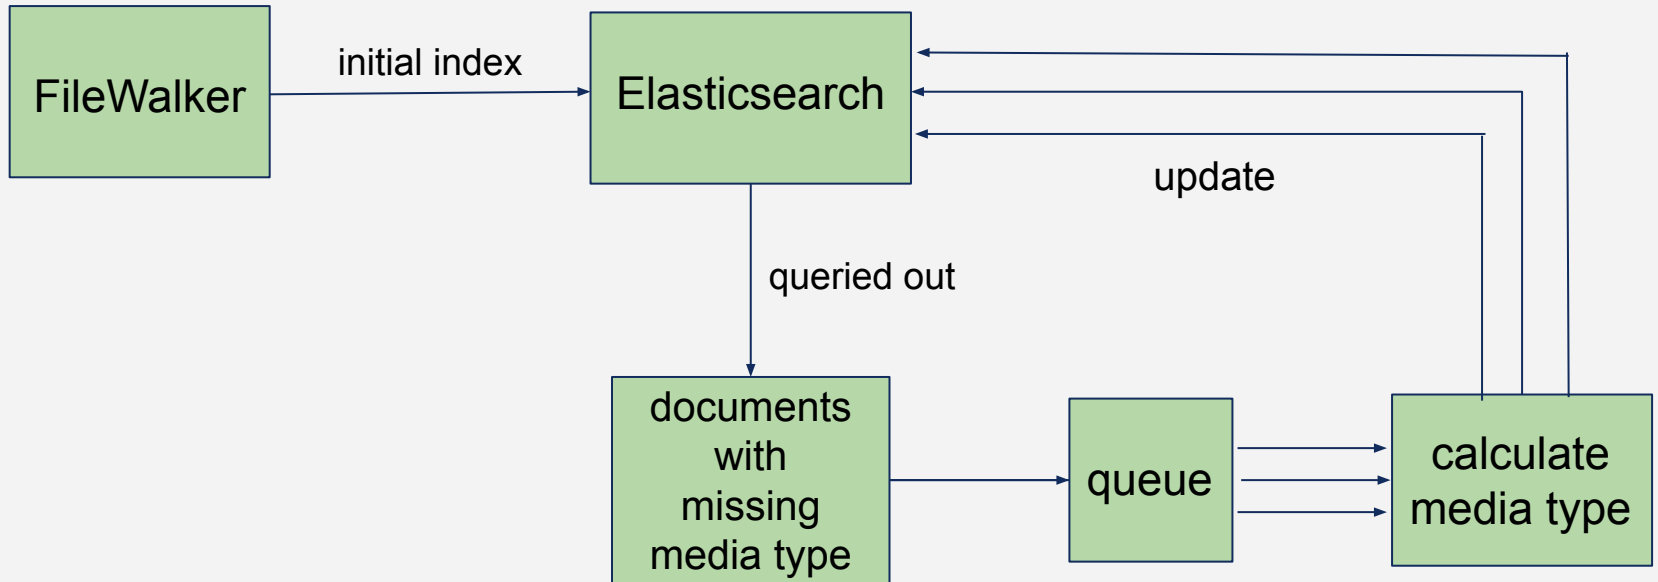
- Current Workflow:



- **Speed: ~21,000 updates / hr**
- **~500,000 / day**

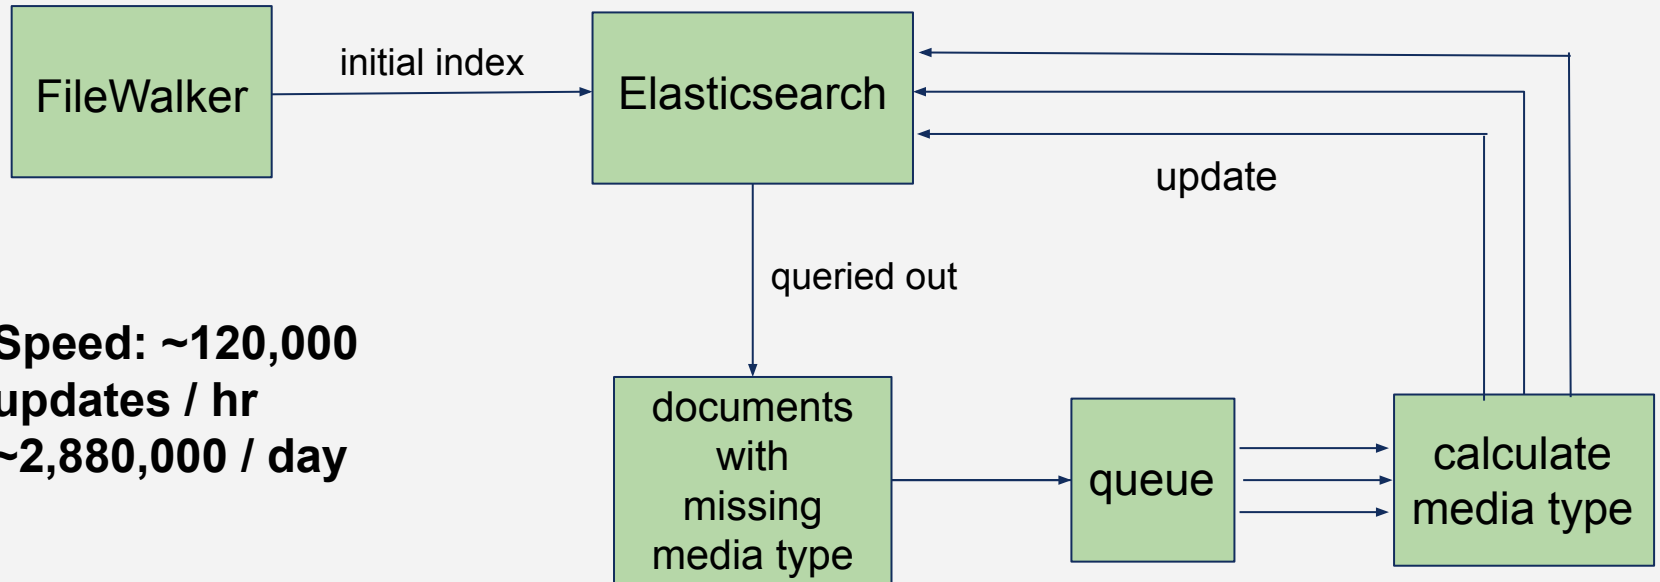
The metadata calculation problem

- Improvement: multithreading and tuning variables
- Workflow:



The metadata calculation problem

- Improvement: multithreading and tuning variables
- Workflow:



- Speed: ~120,000 updates / hr
- ~2,880,000 / day

Exploring scientific metadata

- **We know things about the files, what's actually in them?**
 - Harvest data inside netcdf, hdf, and grib files using Unidata NetCDF library
 - Ex: variable names (standard_name, long_name, short_name), contact, author




Results

- Indexed **486,899,650** files and **13,075,066** directories (~**37.1** pebibytes)
- Calculated media type for **36,262,238** files
- Checked **2,657,132** files for scientific metadata



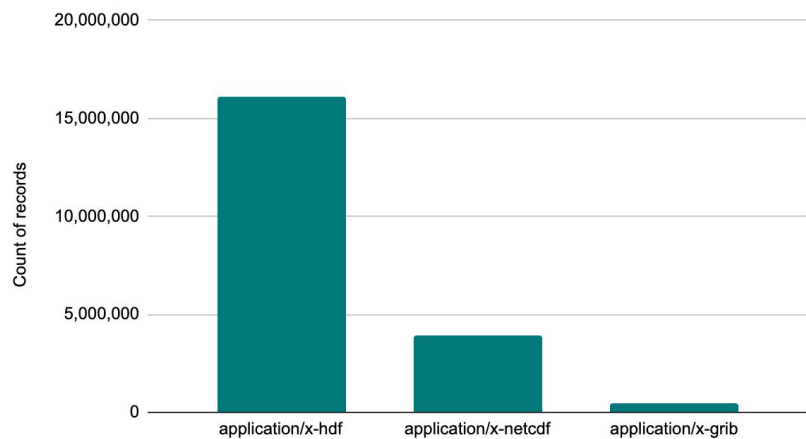
Results

- **152** unique media types
- **21,722,204 (59.8%)** are hdf, netcdf, grib




Results

Count of scientific files




Median size of scientific files




Results

- **70%** has standard_names, **9.1%** has contact, **0.1%** has author

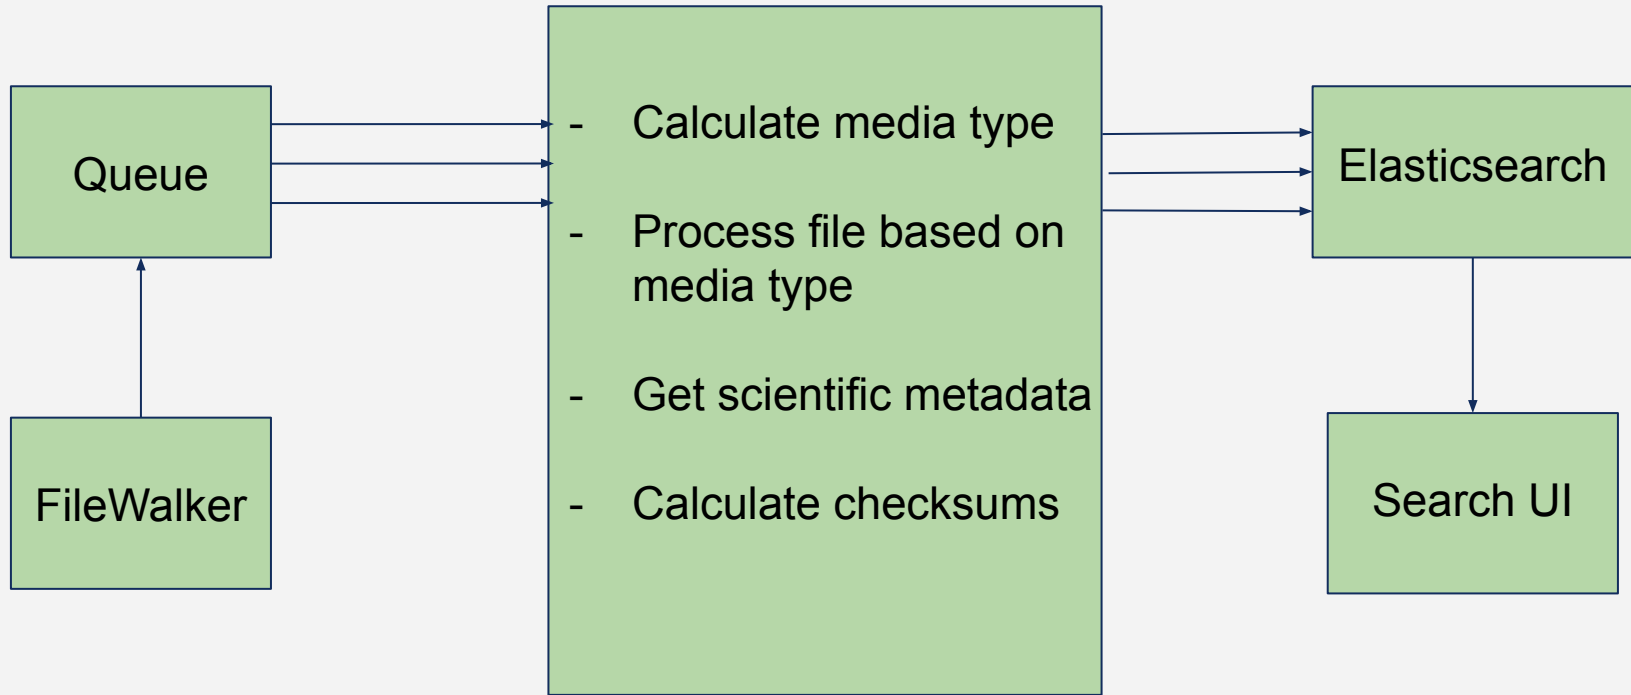
Top 10 values of long_name



Top 10 values of standard_name



Future Work: adding functionalities and increasing efficiency



Acknowledgements

NSF, NCAR, CISL, Sage Team

Mentors: Nathan Hook, Jason Cunning, Eric Nienhouse

SIParCS Program: Virginia Do, Jerry Cycone, all administrative staffs, fellow SIParCS interns

