National Synchrotron Light Source II

Brookhaven National Laboratory

U.S. DEPARTMENT OF ENERGY

Overview of NSLS-II Computational Infrastructure and Plans

Stuart Campbell

NSLS-II Chief Data Scientist & Deputy Program Manager, Data Science and Systems Integration (DSSI)

23rd June 2022

Our Vision

Integration is the key…

Computing technology has transformed many aspects of our lives.

This transformation has largely been enabled by integration and interoperability which hides the complex details of databases, storage, interfaces, file formats from the end user.

In the long term, NSLS-II will use the same concept to transform the user's experiments.

Building a Data Ecosystem

Core to our strategy is to transform NSLS-II to have a “Data Ecosystem”

- “An ecosystem is a community of living organisms in conjunction with the nonliving components of their environment, interacting as a system”.
- We require the same thing: staff and users interacting with the computing environment as a system.

For our users and staff (and NSLS-II) to thrive the data also needs to thrive. That will require:

- A robust and flexible data infrastructure.
- An innovative and supportive development environment.
- A holistic approach to data.
- Staff and collaborator development in coding practices / standards so that work can be transferred and reused.
- Common interfaces and standards.

Data science guiding principals

Development of common platforms where possible and make customization easy to fit the diversity of beamlines and science cases.

Be good stewards; develop where possible to open-source standards and ethos and collaborate where possible.

Develop to accepted engineering practices.

Cyber security is a core value and is considered at all stages of planning, development, deployment and use.

Infrastructure

Where we were…

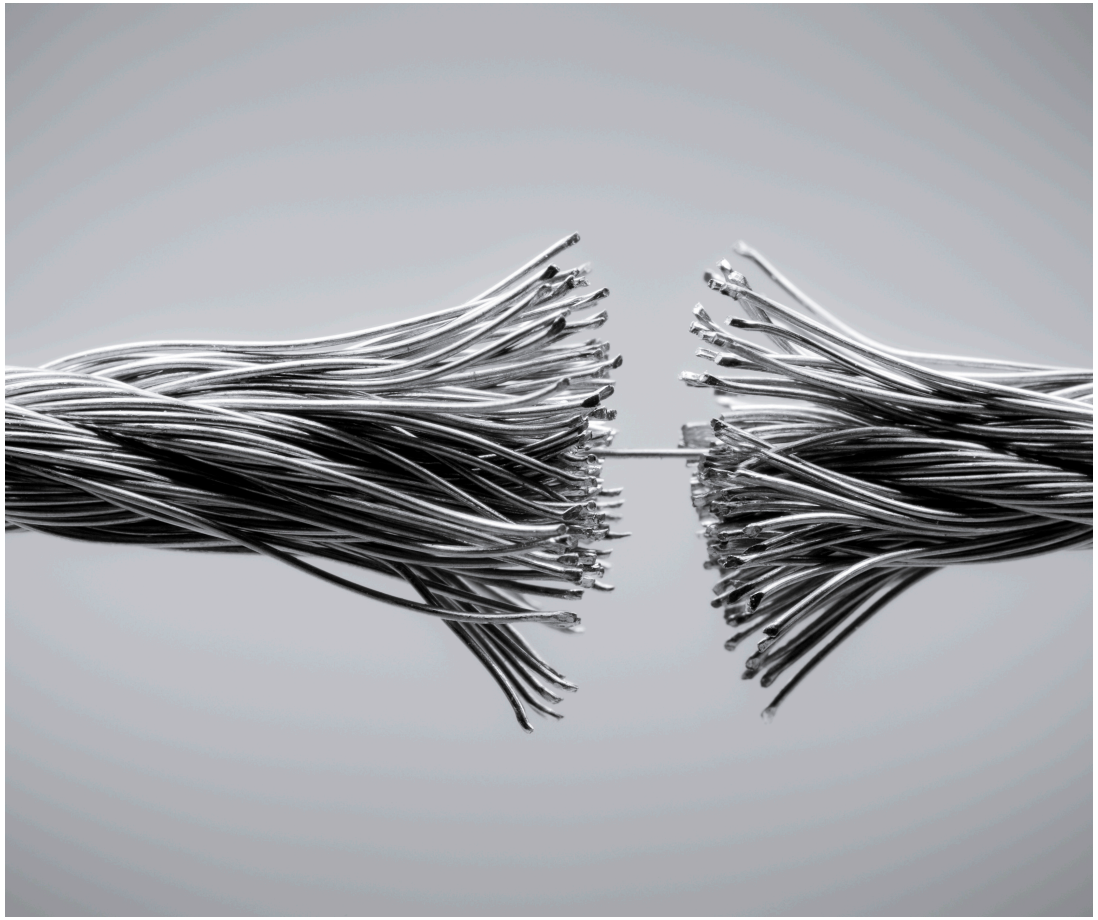

- NSLS-II was built with a non-routable network
- 2 x 1Gbps proxied connections to Campus Network (80 Gbps out of BNL)
- Protocols and Programs that couldn't use the proxy were unable to connect outside NSLS-II
- Resulted in multi-hop connections and indirect methods to copy data out of the facility.

…and where we are now.

- Network is now routable which enables connectivity within BNL and to external sites
- 4 x 100Gbps pipes to High Throughput Science Network (400Gbps out of BNL)
- 4 x 100Gbps backup pipes

Roads are 100x as wide, twice as many and no proxy.

NSLS-II by the numbers

To run 28 beamlines, the NSLS-II has:

- > 2,500 User accounts (automatically generated)

- > 3.5 PB Central Storage
- > 3 PB GPFS Storage distributed across beamlines
- > 3500 network devices
- > 750 Linux Machines
- > 150 Network Switches
- > 85 HPC cluster nodes (split over 4 different clusters)

- 400 Gbps connectivity (4 x 100 Gbps) from NSLS-II to BNL's High Throughput Science Network

vmware®

lustre®

Red Hat

ARISTA

Hewlett Packard Enterprise

NetApp®

Bluesky

bluesky

Design of Bluesky readily accommodates both adaptive and autonomous interfaces.

- Data acquisition system for high-level control and planning
- Collection of co-developed Python libraries (useful a la carte)
- Support both automatic and manual metadata encoding
- Store data/metadata in robust, searchable API called Databroker
- Data emitted in streaming fashion via standard Python structures
- First-class support for adaptive feedback and inline analysis
- Adopted across the DOE Lightsource complex.
- Now an open-source collaborative project between all stakeholders www.blueskyproject.io

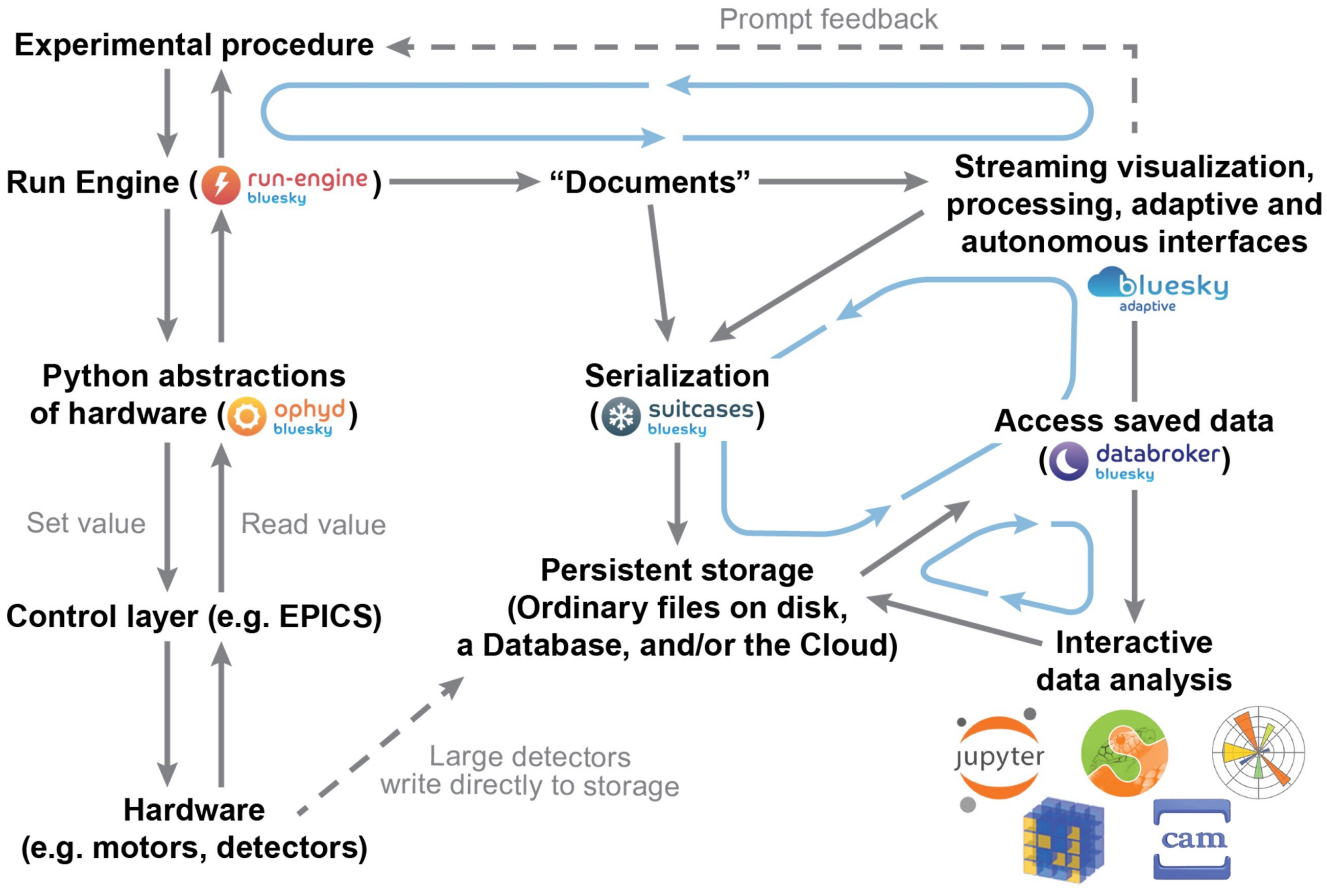

Bluesky is designed in service to data analysis

When analyzing data we want....

To easily find the data we're looking for.

Access to that data, not particularly caring where it's stored or in what file format.

Well-structured data marked up with relevant context, to support easy and sometimes automated batch analysis.

Seamless integration with popular data analysis tools.

Enabling Extensible Workflows

Streaming data and information into central message bus.

Consumers listen to messages on a topic that they are interested in.

Full data and metadata are available to all consumers.

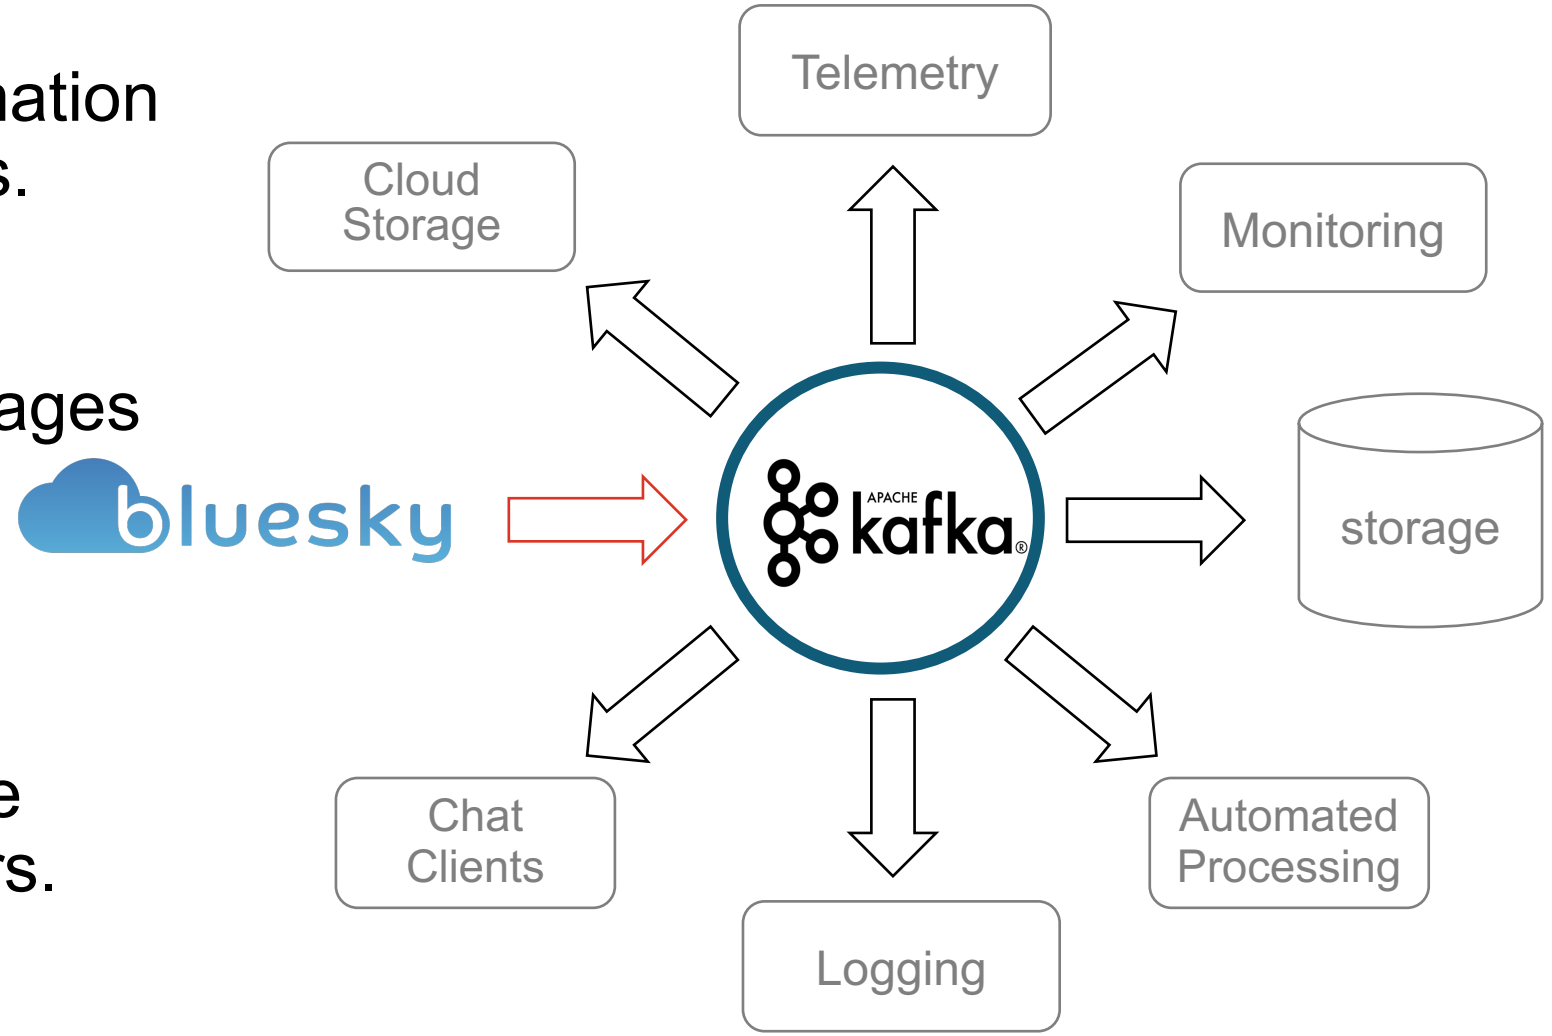

Bluesky is designed for the long term

Make it easy to keep file-reading and -writing code separate from scientific code.

Support streaming (live-updating) visualization and processing, and adaptive experiment steering.

Integrate well with web technologies and be cloud-friendly.

But also meet users where they are!

Tiled incorporates years of user feedback on *Databroker*

- Users of Python Databroker, now rebuit on *Tiled*, will find...
  - Speed!
  - More useful interactive output
  - More obvious, easy-to-guess usage
  - More powerful searches
  - Easy export-to-file
  - Security
  - Backward-compatibility with existing scripts, notebooks, etc.

- And we aren't locked into just Python anymore...

Be unopinionated about file formats

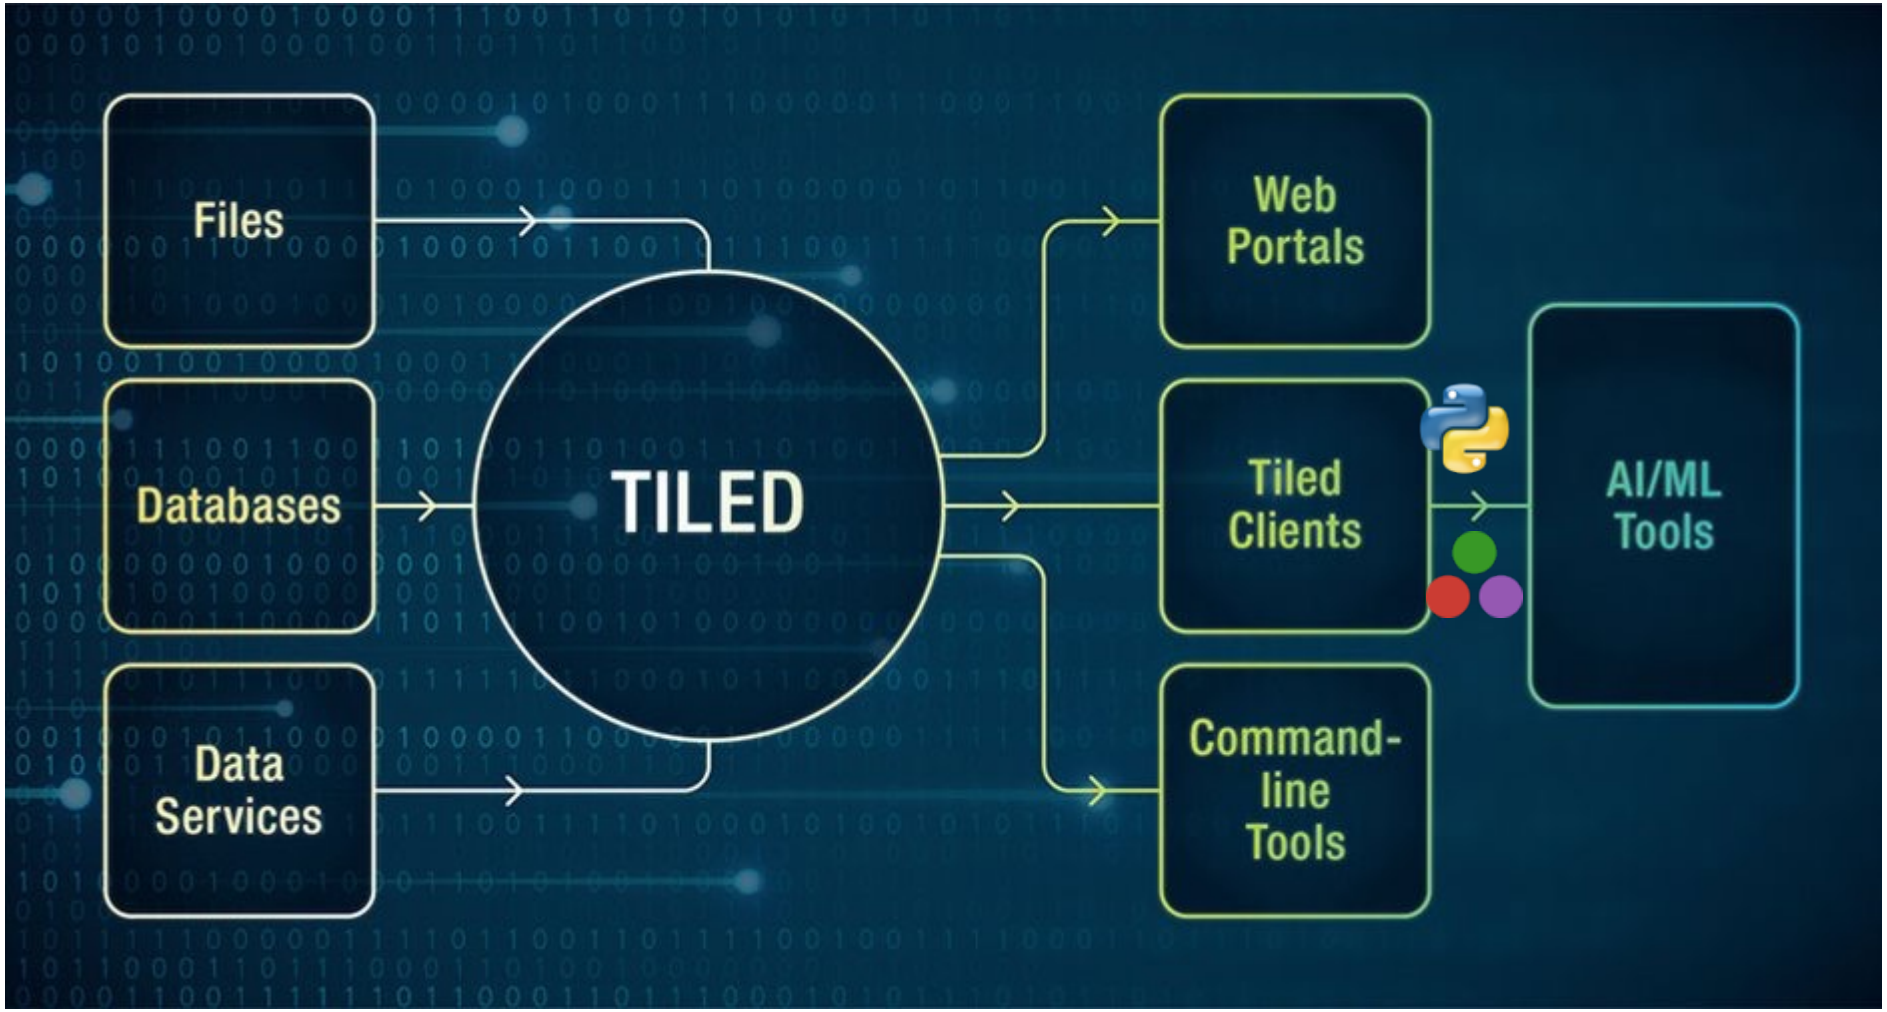

Tiled can provided slices of data as...

- Custom, one-off text format designed to be parsed by a **30-year-old bash script** that still works ("if it ain't broke...")
- Traditional formats like CSV, TIFF, or HDF5 to be opened by **Igor, Origin, ImageJ, PyMCA**, ….
- **Web-friendly** image or data formats to be directly displayed by a web browser or web application
- Chunks of compressed **C** or **Arrow**-encoded buffers to be fetched on demand and fed zero-copy to **Tensorflow**

We can meet all users where they are!

In summary: separate how data is *stored* from how it is *accessed*

- Easy access from any Internet-connected software on any Operating System
- Let analysis code request it in **any appropriate format**
  - Many natively supported
  - Extensible registry of formats for each data structure
- Users doesn't need alphabet soup of technologies (Lustre, Globus, Samba, Mongo, …) to get their data.
- Makes it practical to practice good data management and stewardship

Thank you