Rethinking streaming system construction for next-generation collaborative science

Matthew Wolf, Patrick Widener, Greg Eisenhauer
-- and a cast of many more

Streaming to support new science -- Big Data's other 4 V's

- Historically, a great deal of emphasis has been placed on batch processing of data-at-rest
- However, this focus has meant that scientists trying to do interactive or collaborative work have had to work with mismatched tools

- In particular, the **steering/command and control** functions in many scenarios gets short shrift
  - Collaboration is more than sharing repositories
  - Discovery, multi-disciplinary viewpoints on data, verification & gatekeeping on data

STREAMING AT EXASCALE: THE RISE OF IN SITU

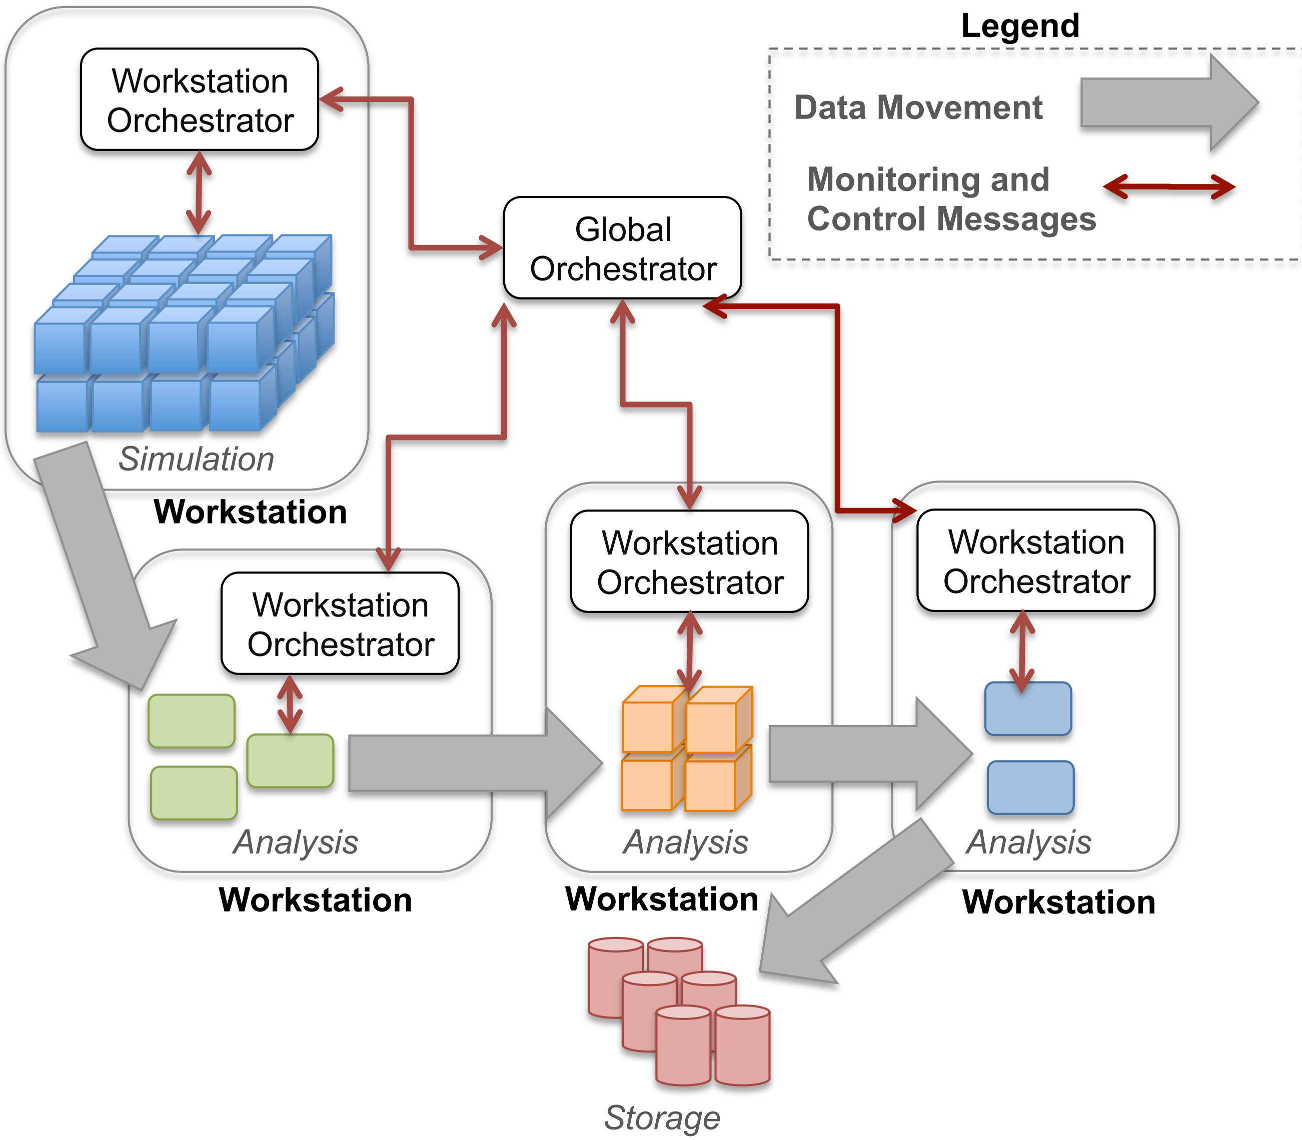

Codes

- GTS
- GTC-P
- LAMMPS
- PIConGPU
- Pixie3D
- S3D
- Einstein Toolkit
- …

Thanks: Jai Dayal, Scott Klasky, Hasan Abbasi, Fang Zheng, Norbert Podhorski, Karsten Schwan, Manish Parashar, Jay Lofstead…

Zoom-In Analysis

VMWare, Amazon, DOE

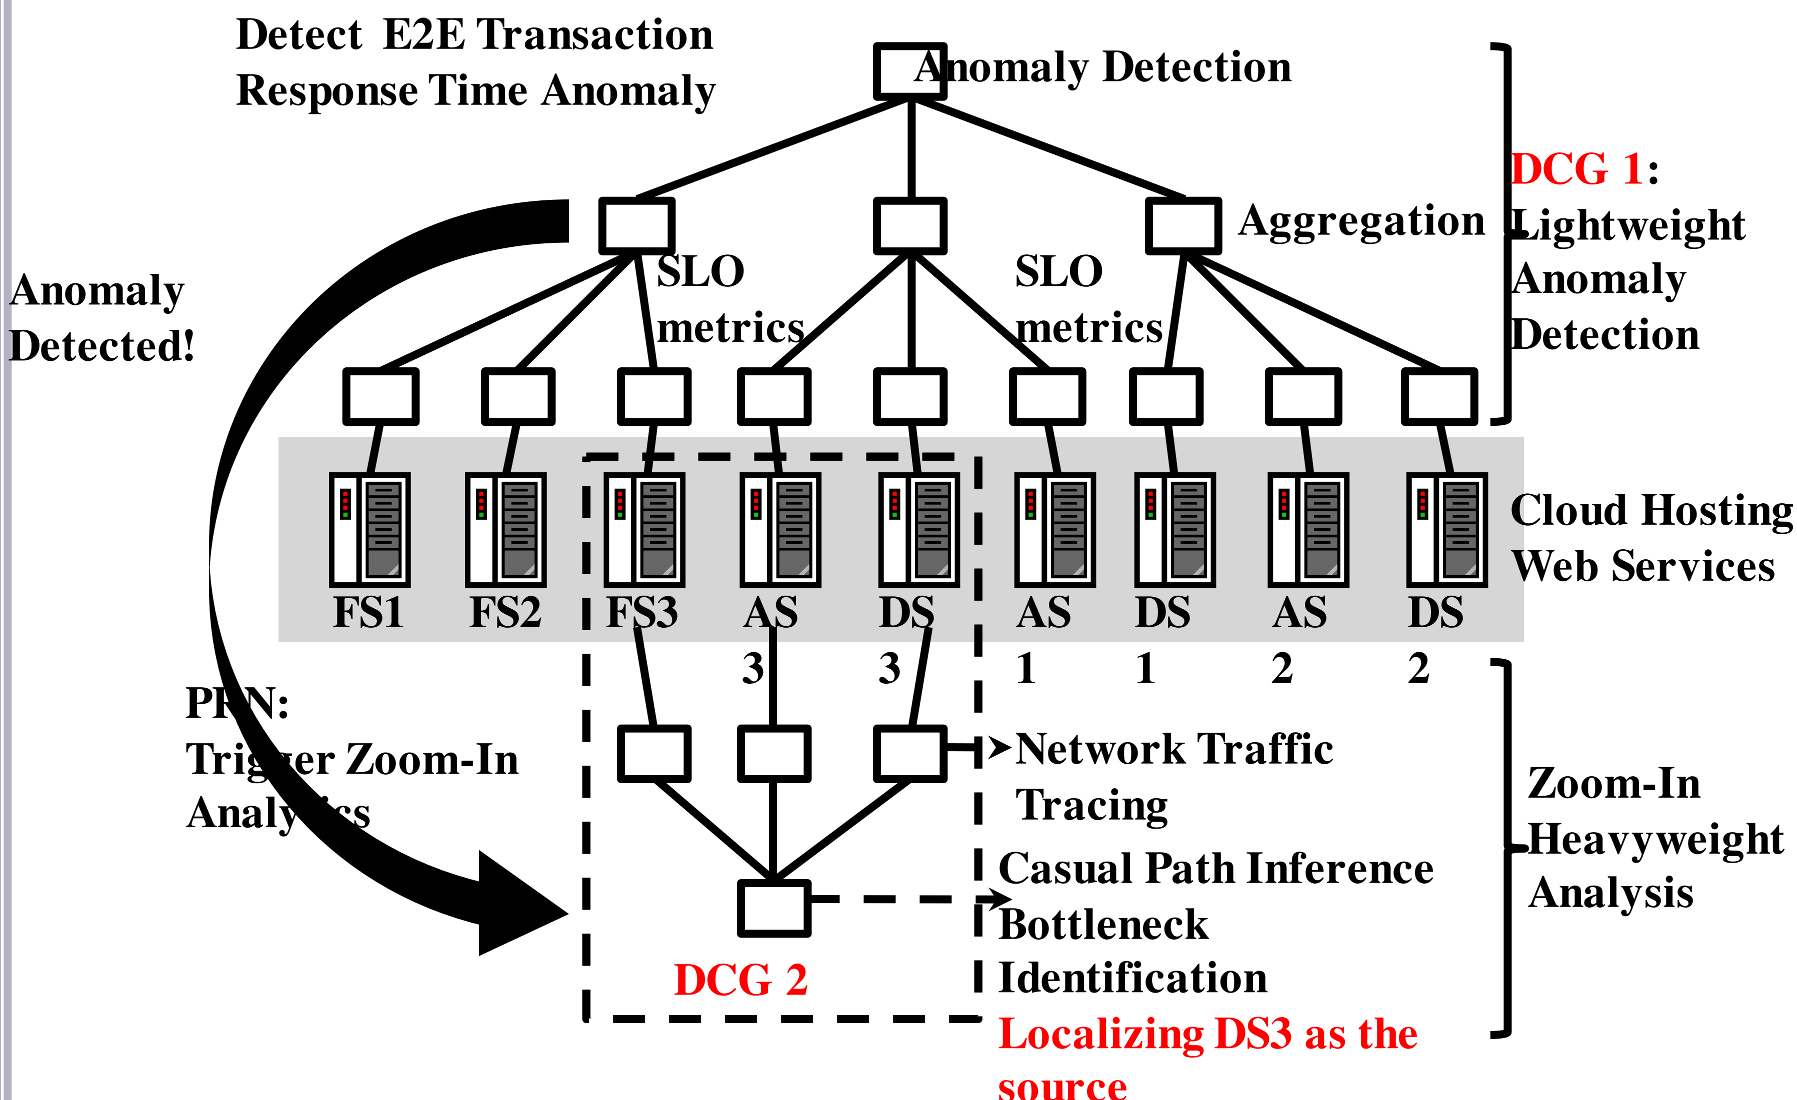

Highly reduced data overhead and focus on problematic area

Thanks: Chengwei Wang, Drew Bratcher, Karsten Schwan, and many more.

SOFTWARE SOLUTION: AN EVENT PROCESSING TOOLKIT

- http://evpath.net & http://korvo.gatech.edu/software
- **EVPath is an Open Source event processing infrastructure designed for high performance**
    - A component of the SDAV SciDAC institute
- **Allows the construction of application-level overlay networks with embedded computation**
    - Fully-typed data flows along the path
    - Very low overhead self-describing binary data
    - Dynamic code generation for on-the-fly processing
    - Flexible network infrastructure allows run-time selection and parameterization of network transport
- **Toolkit that supports construction of CDN-like, DHT-like, aggregation-tree-like, asynchronous, p2p, or other steering infrastructures**

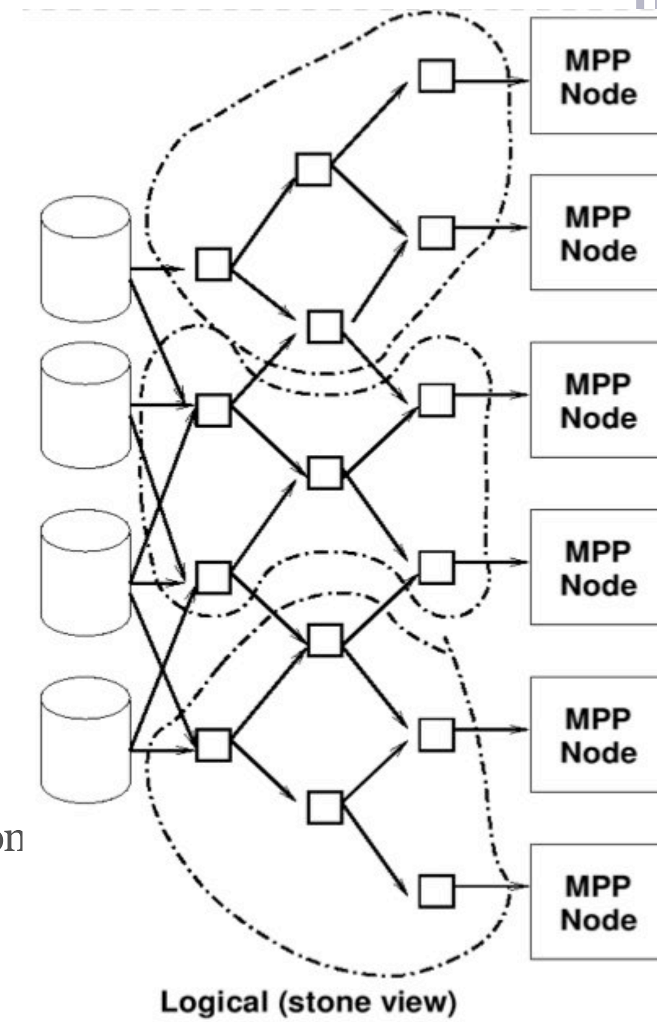

An Illustrative Example: Experimental Combustion Collaboration

- Science goal is to understand the complex dynamics of different fuel mixes, speeds, acoustic interactions, and so on
- Use laser probes and cameras at 10k+ frames per second
- Inject particles so you can trace fuel, flame, and residue in real time.
- Initial process was driven by disk I/O & storage transport

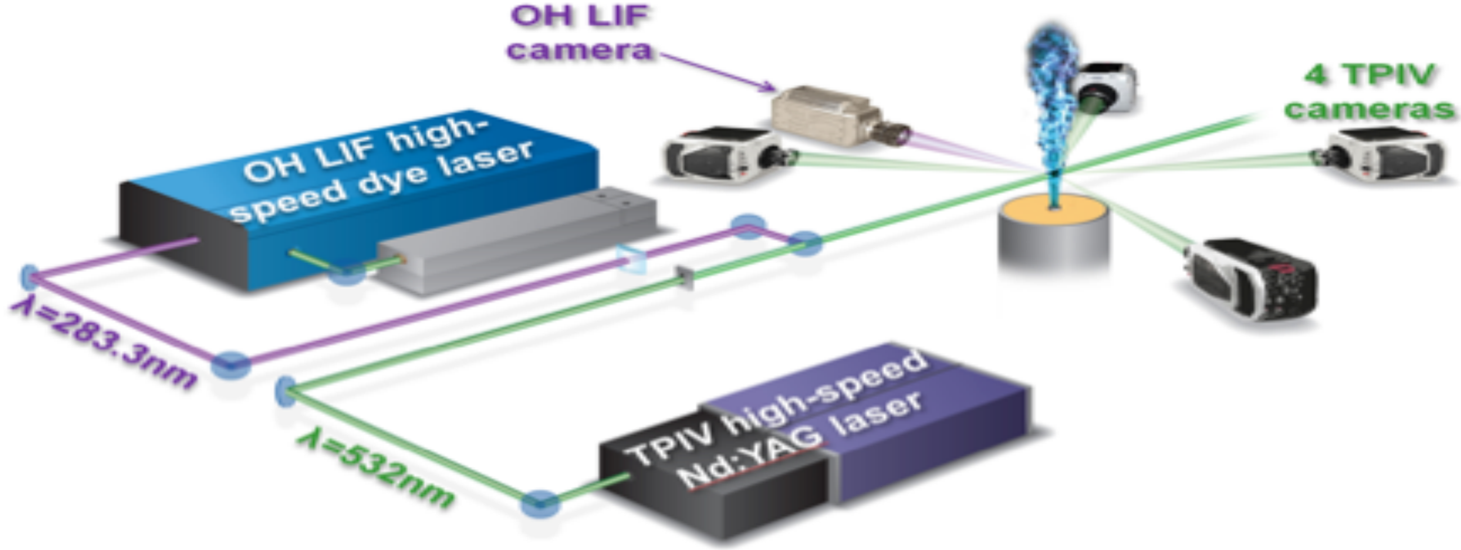

Thanks: Tim Lieuwen, Ben Emerson, Vishal Acharya, Jonathan Frank, Akash Gagnil, Drew Bratcher

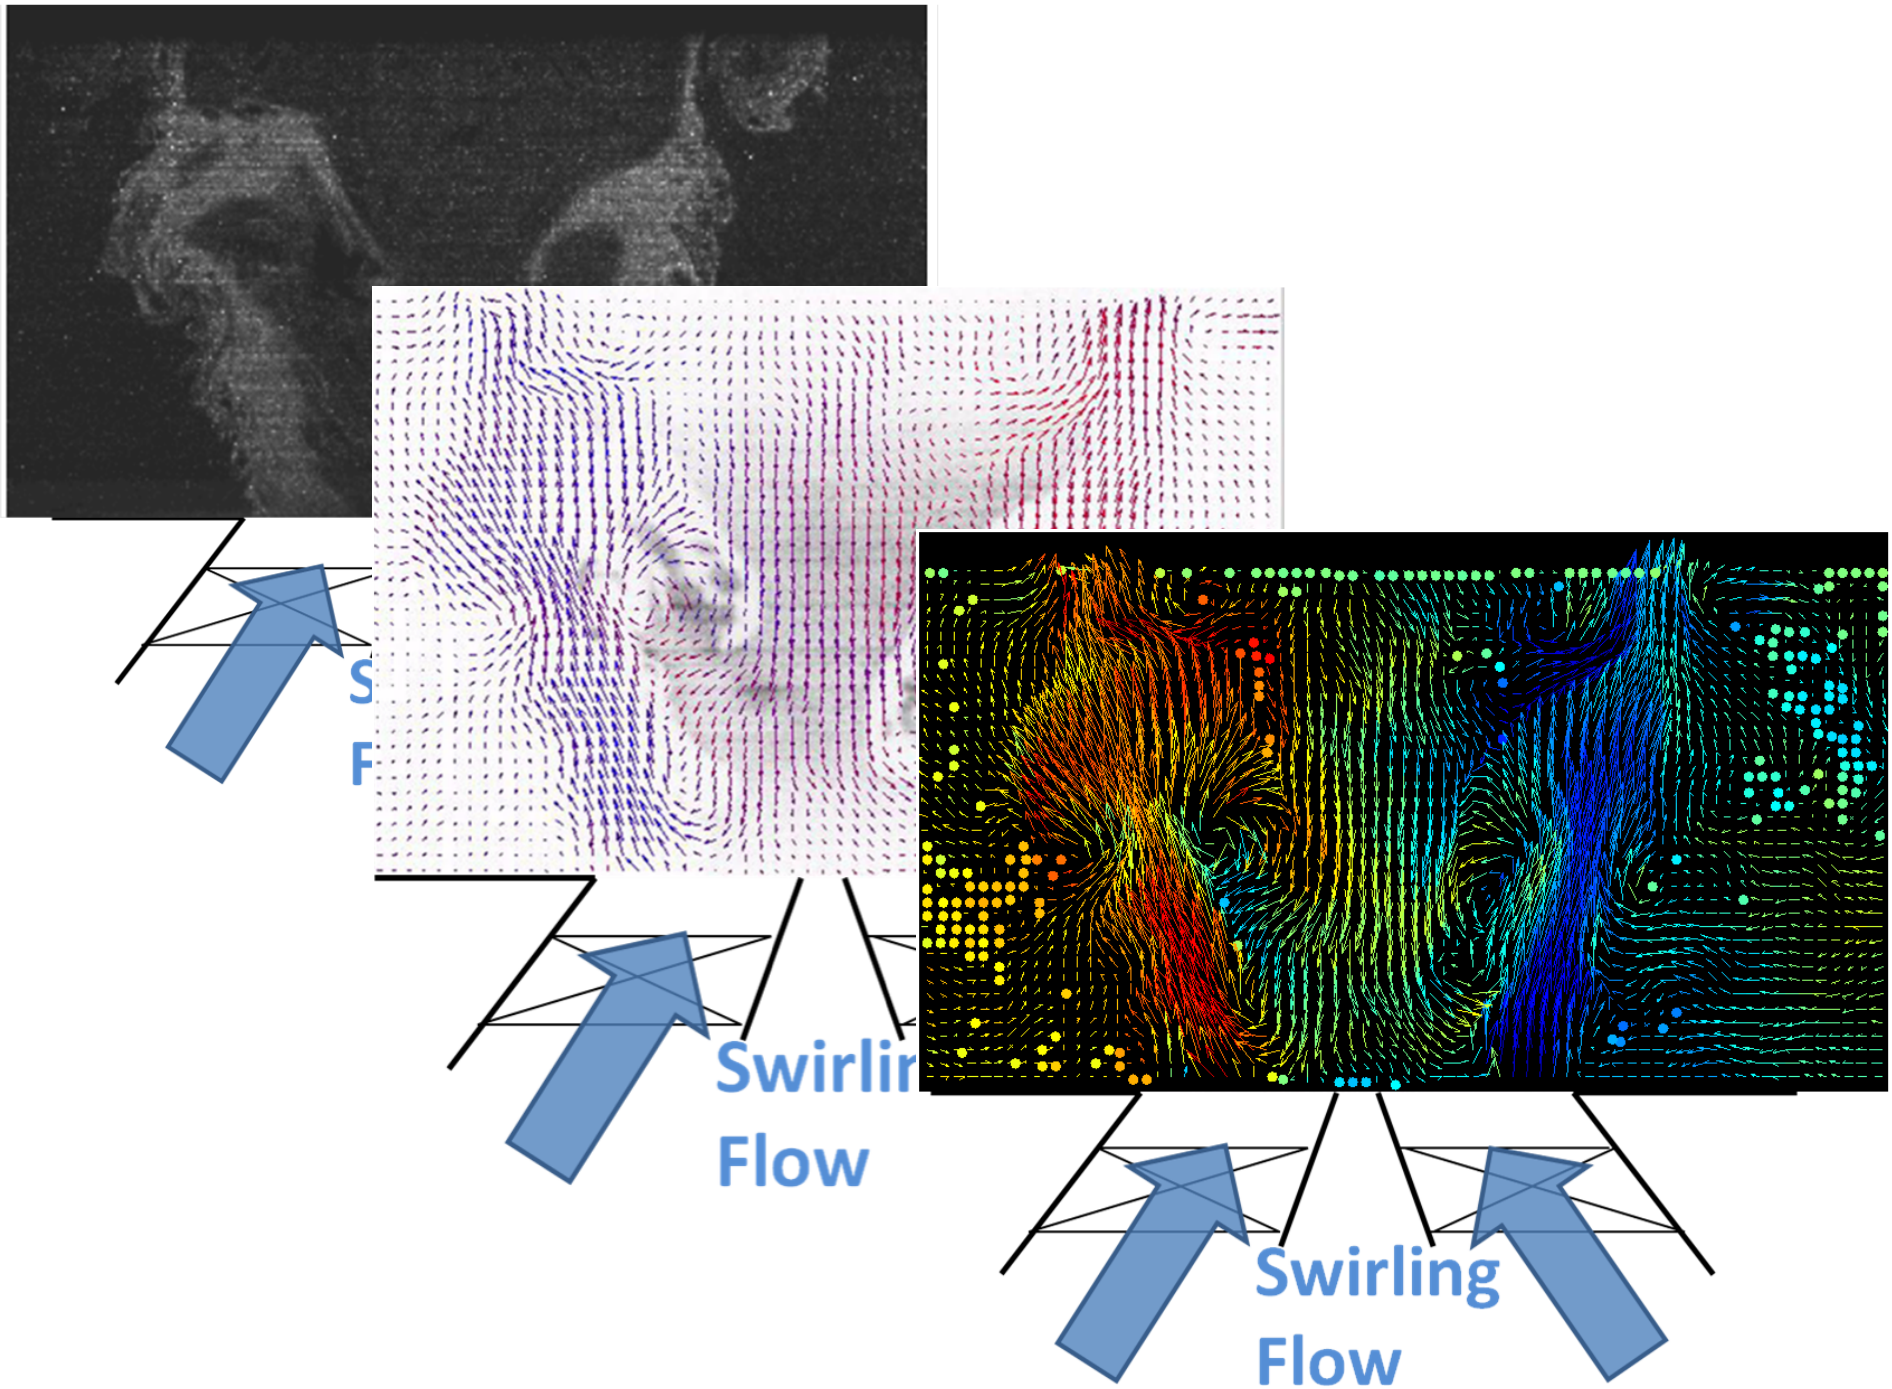

Flow
Swirling
Flow

- Stream processing lets us address a number of critical issues:
  - Are the lasers properly aligned? Did someone bump something?
  - Are the particle injectors working correctly?
  - Are there any obvious experimental defects in the data (i.e. chunks of foam)?
  - Does this look approximately right for the input parameters (i.e. did someone leave a wrench in the inlet)?
  - Has the effect we're looking at saturated? Should we change the next parameter test in the campaign?
  - **Does this line up with what we know from simulation? Should I adapt the campaign to better probe the difference?**
  - **Are the Physical Chemists right?**

SciKhan – an Initial demonstration

- The interactions between data-in-motion and data-at-rest (thanks, IBM!) can be complicated.
- Scientists wanted the stream-based capabilities, but they were used to a file system interface.

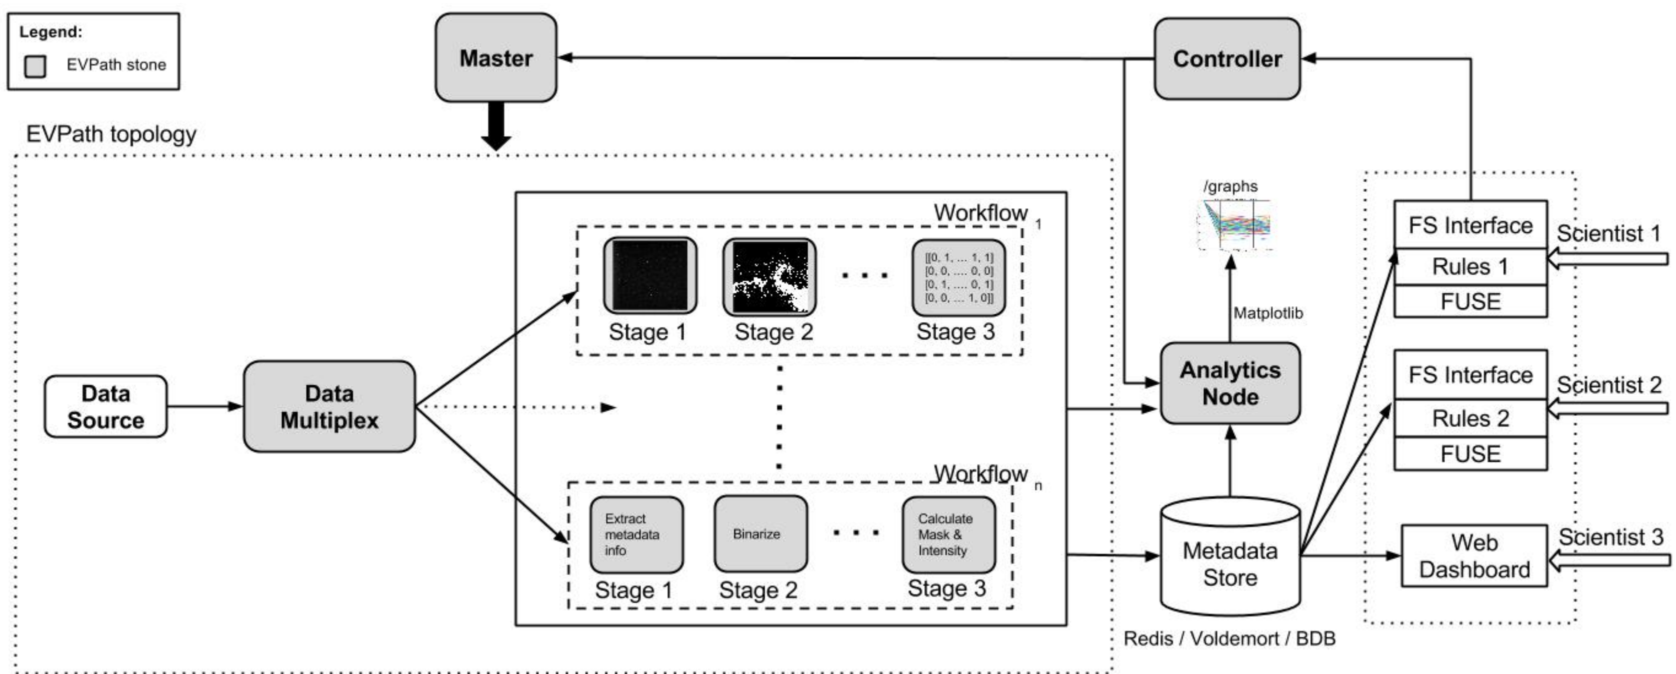

Conclusion

- The data management problem is beyond just large Volume.
  - Streaming has been treated as a corner case for a long time
  - Critical gap when all 5 V's (volume, velocity, variety, value and veracity) are in play
- Steering and/or control requires highly specialized designs for each of the users
  - Use a **toolkit** that allows that **customization**
    - Human-in-the-loop, delegated control, etc.
- There is a change management problem
  - The science questions and the way science is conducted can change as the technology shifts