

IEEE MSST 2016

Workflows

A Procurement/Spec Tool Session

Gary Grider

Division Leader
High Performance Computing Division
Los Alamos National Laboratory

ggrider@lanl.gov

Spring 2016

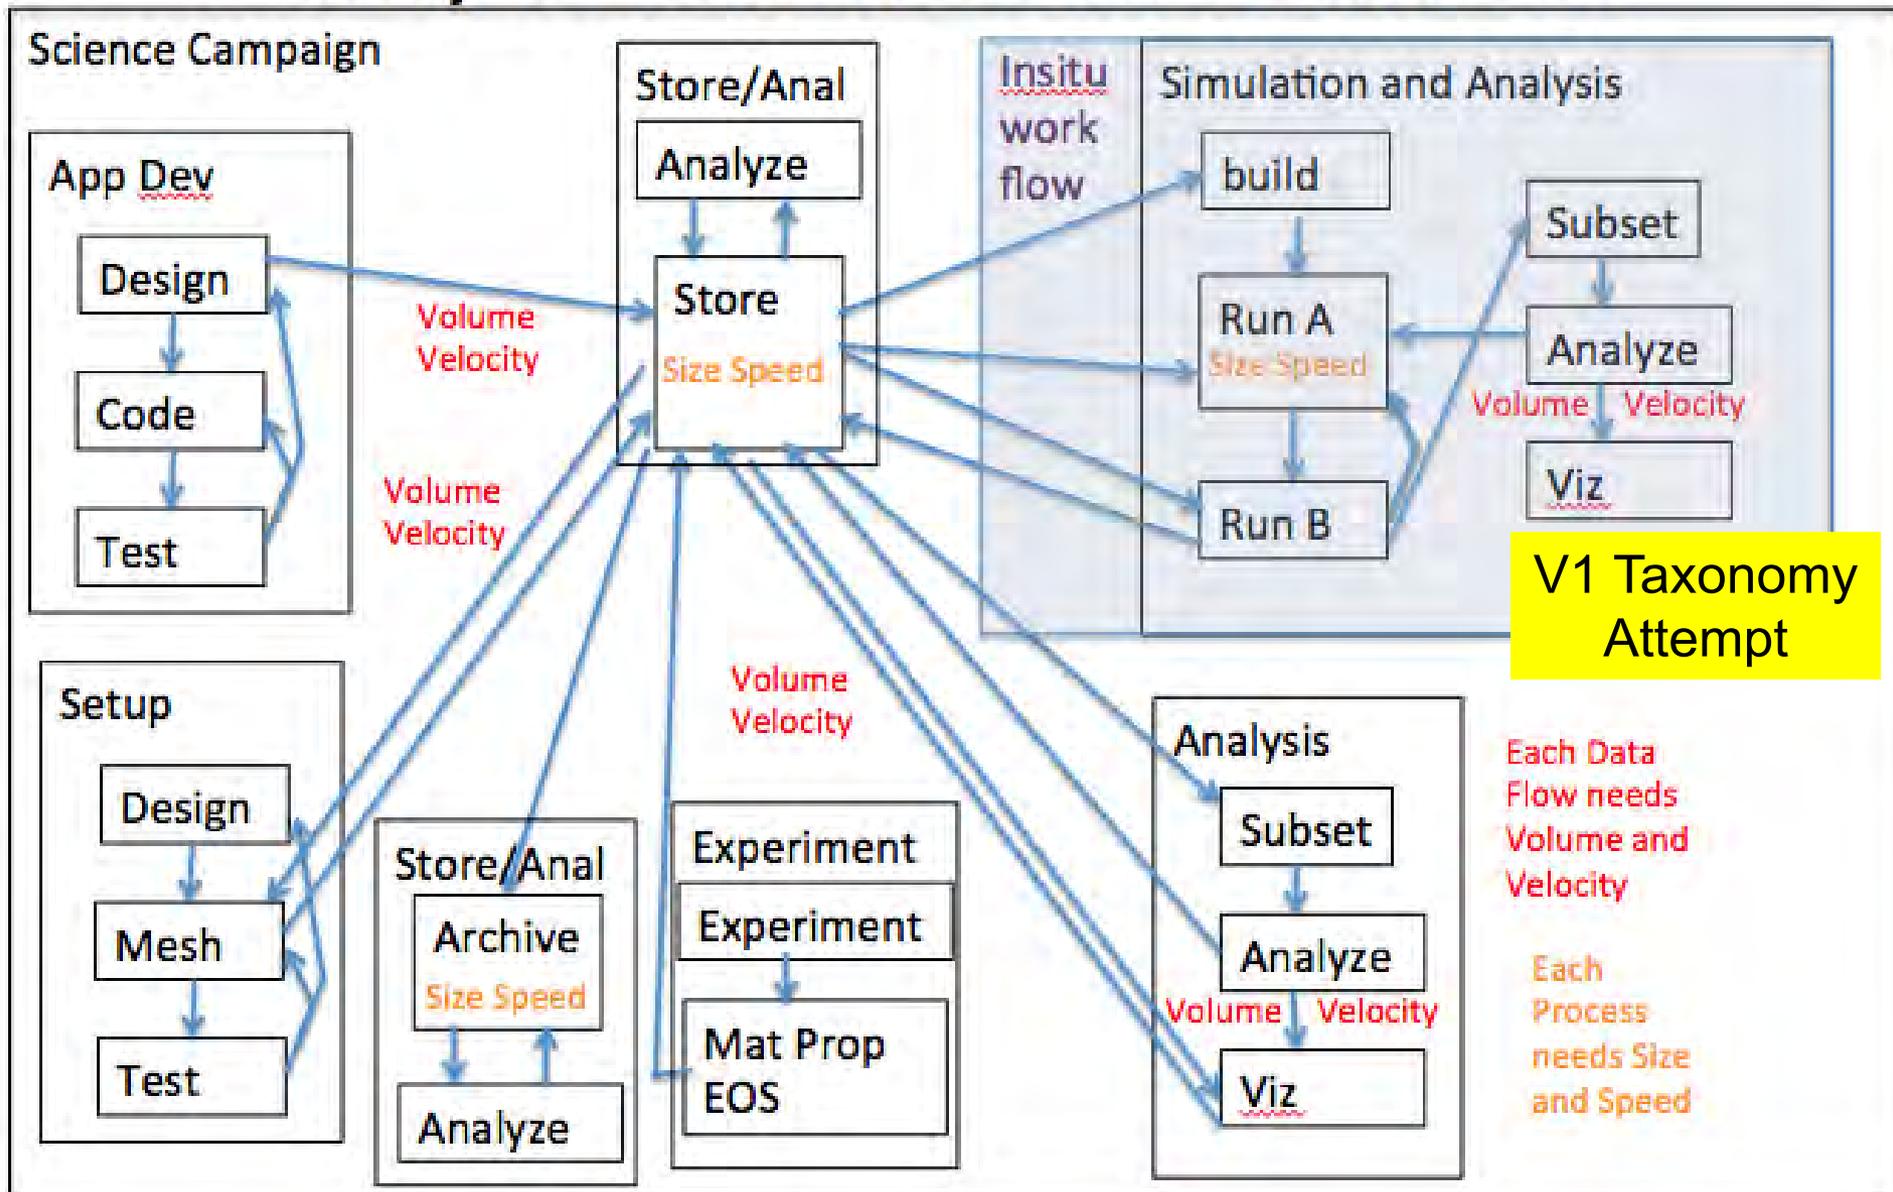
Excerpts from LA-UR-16-20184

Workflows to the rescue?

- **What did I learn from the workflow-fest circa 04/2015?**
 - There are 57 ways to interpret in situ 😊
 - There are more workflow tools than requirements documents
 - There is no common taxonomy that can be used to reason about data flow/workflow for architects or programmers ☹️
- **What did I learn from FY15 Apex Vendor meetings**
 - Where do you want your flash, how big, how fast, how durable
 - Where do you want your SCM, how big, how fast
 - Do you want it near the node or near the disk or in the network
 - --- YOU REALLY DON'T WANT ME TO TELL YOU WHERE TO PUT YOUR NAND/SCM ---

Can workflows help us beyond some automation tools?

Insitu/Post/Active Storage/Active Archive Analysis Work Flows In Work Flows



Workflows: Potential Taxonomy

Led to Work by Dave Montoya

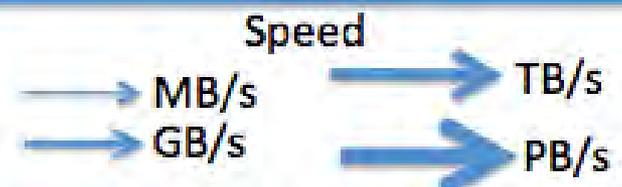
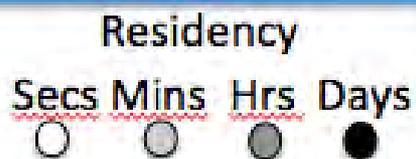
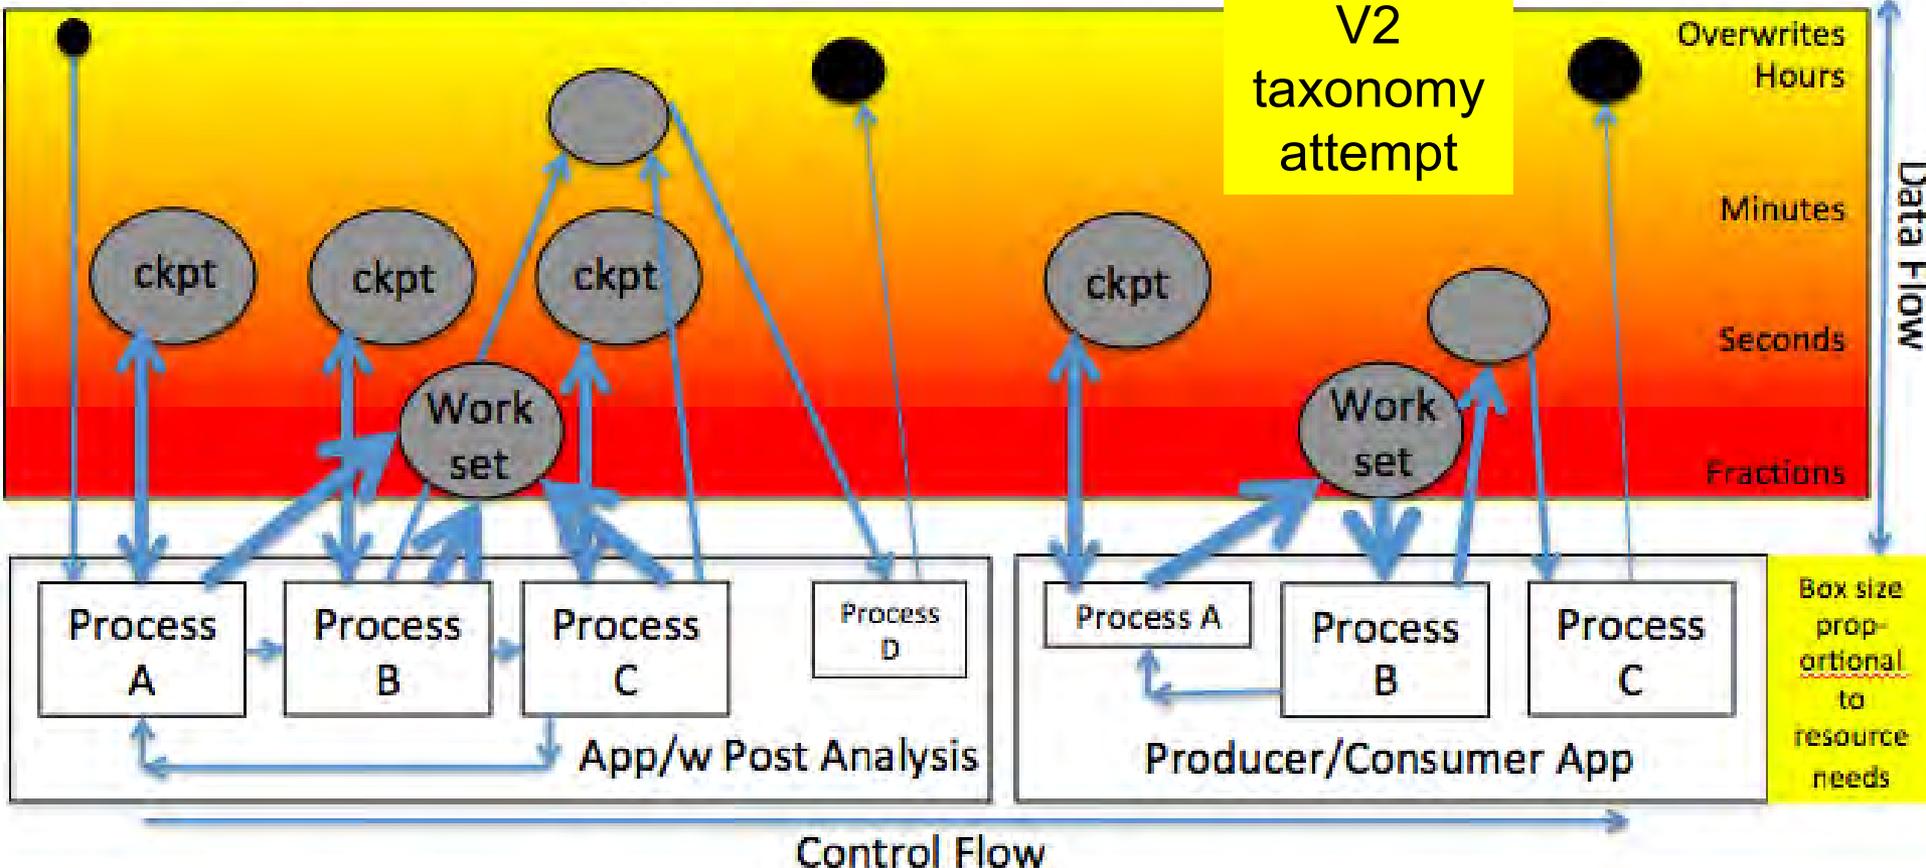
Where to load/store based on size, speed, frequency

Capacity of tiers of memory/storage can be derived

Speed of tiers of memory/storage can be derived

Durability of tiers of memory/storage can be derived

Circa 05/2015



Workflows can help us beyond some automation tools: Workflows enter the realm of RFP/Procurement

■ Trinity/Cori

- We didn't specify flops, we specified running bigger app faster
- We wanted it to go forward 90% of the time
- We didn't specify how much burst buffer, or speeds/feeds
- Vendors didn't like this at first but realized it was degrees of freedom we were giving them

■ Apex Crossroads/NERSC+1

- Still no flops 😊
- Still want it to go forward a large % of the time
- Vendors ask: where and how much flash/nvram/pixy dust do we put on node, in network, in ionode, near storage, blah blah
- We don't care we want to get the most of these work flows through the system in 6 months

The Panel

- **Workflow as a procurement/specification tool**
 - Examples of how workflow specifications have been used to shape RFP responses
 - Examples of how solid state storage technology has been used to accelerate scientific or other significant workflows
- **Panelists**
 - Dave Montoya - DOE/LANL (previous talk)
 - Yoonho Park – IBM
 - Bret Weber – DDN
 - Lance Evans – Cray
 - Gary Grider – DOE/LANL (arm waving moderator)