BDEC Breakout Arch 1

2014/02/28 AM

Pete Beckman (ANL)
Satoshi Matsuoka (Tokyo Tech.)

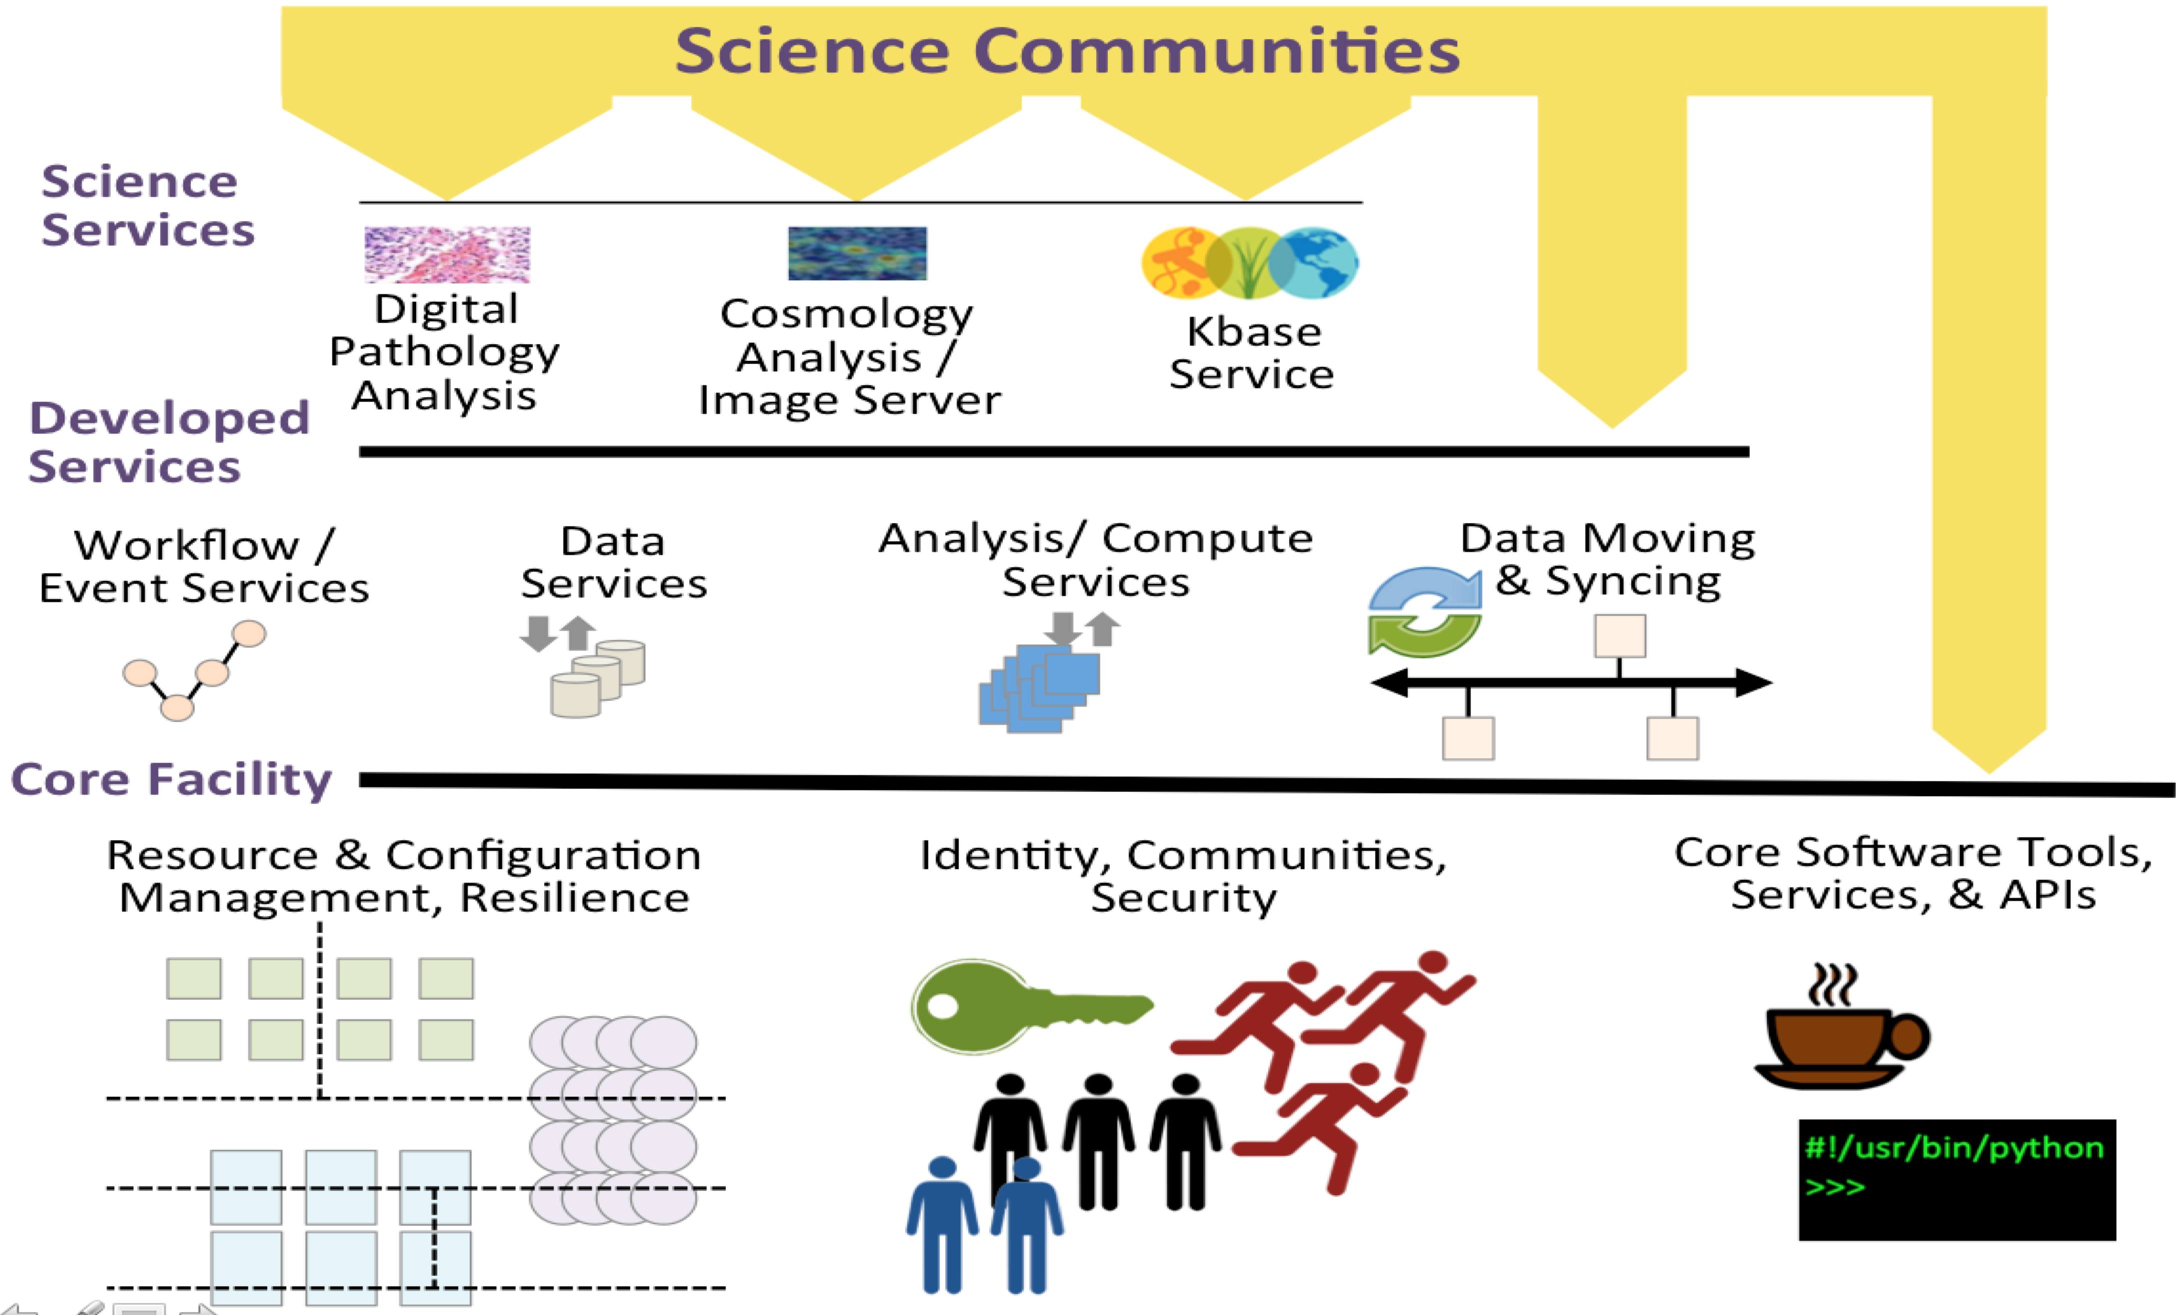

Science Communities
Science Services
Digital Pathology Analysis
Cosmology Analysis / Image Server
Kbase Service
Developed Services
Workflow / Event Services
Data Services
Analysis/ Compute Services
Data Moving & Syncing
Core Facility
Resource & Configuration Management, Resilience
Identity, Communities, Security
Core Software Tools, Services, & APIs
#!/usr/bin/python
>>>

Focus of discussion---what do we articulate

- What are the key abstractions?
- What are the missing pieces (gap analysis) c.f. extreme computing?
- Pete's initial layer
  - Developed Services
  - Developed services
  - Core Facility
- Single site or multi-site (Franck) – could be location independent?

Gap analysis 1 – core facilities 1

- Resource & configuration management and resilience
- Identity and community management & security
- Core software, services & APIs
  - Persistent service, ad-hoc, interactive service vs. reserved/batch, orchestrated services (throughput vs. response) – are they a dichotomy or coexist
  - Variability of runtime acceptable? – would like being able to predict certain degree of throughput, but not critical (little "scaling" experiments)
  - No elasticity exists in most systems – no ad-hoc addition removal, market-based charging
  - "swarming"
  - Virtual machines – managing complexity of hardware and software stack as well as resource management issues. Do we have a EC-capable VM?
    - Interconnect, I/O, node / heterogeneous processors
    - Nascent state of BD middleware may mandate virtualization (including bare metal)
    - The lack of performance expectations and guarantees out of varying performance due to virtualization hinder

Gap analysis 1 – core facilities 2

- Identity management etc.
  - Central management of identity difficult (c.f. Shibboleth, Dropbox w/email identity)
  - Management of data security etc. curation, etc.
  - Analysis vs. sharing – is there distinction?
  - Access to data services, not just files
  - Different communites behave differently vs. national data thingy service
    - Can we accommodate all the different curation and securitypolicies?
    - Million clientels, international access,
      - Various identity management and authorization
      - Does the technology exist, and/or is it a management issue
      - How do commercial entity deal with the sheer scale and variations of policies?
    - How do we track the access (curation and provanance issues again)
      - Need flexible policies here, too

Gap Analysis 1 - Developed Services 1

- Data moving & synching , workflow
  - How are they different from the past Grid work
    - Schedulable data transfers
    - Complex workflow
  - Lack of adoption – why?
    - Simple data transfers over a TCP/IP network using file staging is a done deal
    - Inability to store data due to Multi-terabit transfer capabilities
      - Lack of Data mover facilities for staging and mid-tier real-time processing
      - Control of terabit scale direct light path to connect the instruments, data archives, compute, etc.
      - Intelligent "SDN" for extreme big data
  - Intelligence of resource control in various service – get human out of the loop as much as possible, only when needed

Gap Analysis 1 - Developed Services 2

- Can existing EC machines (cost) effectively support the BD services?
  - Just a matter of procuring cost configuration heterogeneous tiered machines?
    - Resources at least need to be picked automatically
    - Can we be more aggressive, put more intelligence in data devices for local processing?
- Economics and cost model must become more data centric and sophisticated
- Do we understand the (E)BD community well enough as we understand EC (which is akin to telescopes vs. mass market S&E)
  - Community and number and varieties leading to different policies
  - Data discovery services are also important (as is with web vs. web search)
  - Tools, metadata, workflow etc. also needs to be discovered
  - Journals such as Nature mandate publication of data and procedures for scientific reproducibility- need to support such use cases
- Analyzing, monitoring debugging BD architecture
  - We are getting good tools for EC
  - E.g., anomaly could be glitch in the data processing, not real data
- The need of BDEC testbeds and mini-apps

Morning Breakout: Applications

Spatio-temporal Sensor Integration, Analysis, Classification

- Multi-scale material/tissue structural, molecular, functional characterization. Design of materials with specific structural, energy storage properties, brain, regenerative medicine, cancer
- Integrative multi-scale analyses of the earth, oceans, atmosphere, cities, vegetation etc – cameras and sensors on satellites, aircraft, drones, land vehicles, stationary cameras
- Digital astronomy
- Hydrocarbon exploration, exploitation, pollution remediation
- Aerospace – wind tunnels, acquisition of data during flight
- Solid printing integrative data analyses
- Autonomous vehicles, e.g. self driving cars
- Data generated by numerical simulation codes – PDEs, particle methods
- Fit model with data

Typical Computational/Analysis Tasks

Spatio-temporal Sensor Integration, Analysis, Classification

- Data Cleaning and Low Level Transformations
- Data Subsetting, Filtering, Subsampling
- Spatio-temporal Mapping and Registration
- Object Segmentation
- Feature Extraction
- Object/Region/Feature Classification
- Spatio-temporal Aggregation
- Diffeomorphism type mapping methods (e.g. optimal mass transport)
- Particle filtering/prediction
- Change Detection, Comparison, and Quantification

Coupled data acquisition, data analysis, modeling, prediction and correction – data assimilation, particle filtering etc.

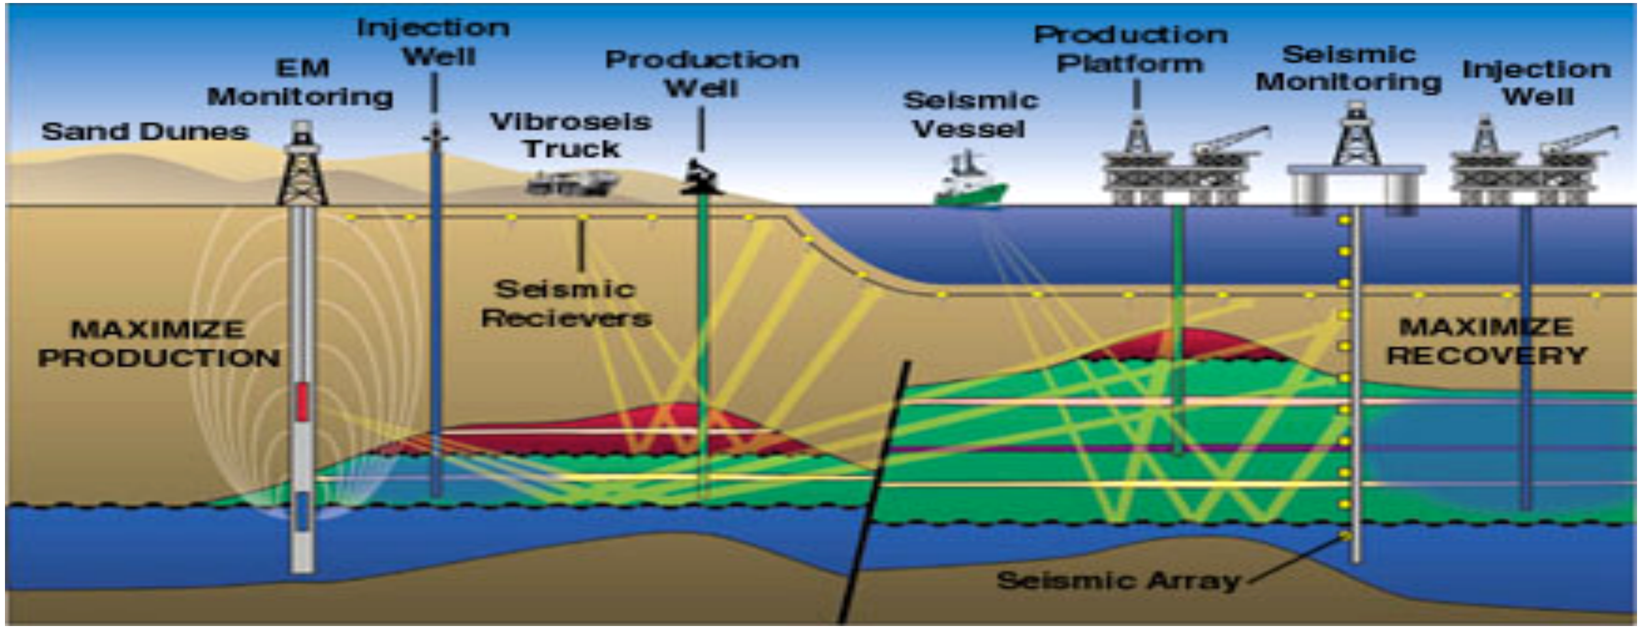

Detect and track changes in data during production
Invert data for reservoir properties
Detect and track reservoir changes

Assimilate data & reservoir properties into the evolving reservoir model
Use simulation and optimization to guide future production

Soft real time and streaming Sensor Data Analysis, Event Detection, Decision Support

- Integrated analyses of patient data – physiological streams, labs, mediations, notes, Radiology, Pathology images, mobile health data feeds
- High frequency trading, arbitrage
- Real time monitoring earthquakes, control of oilfields
- Control of industrial plants, aircraft engines
- Fusion – data capture, control, prediction of disruptions
- Internet of things
- Twitter feeds
- Intensive care alarms

Typical Computational Analysis Tasks

Streaming Sensor Data Analysis, Event Detection, Decision Support

- Prediction algorithms – Kalman, particle filtering
- Machine learning algorithms on aggregated data to develop model, use of model on streaming data for decision support
- Searching for rare events
- Statistical algorithms to distinguish signal from noise
- On the fly integration of multiple complementary data streams

“omics”

- Sequence assembly
- Metagenomics – identification/characterization of populations of organisms based on DNA/RNA sequencing
- Phylogenetics, genetic based population biology, cancer mutation landscaping
- Pathway modeling using integrated sequence, expression, epigenetics, protein, glycans
- Genetic/genomic based design of organisms with specific properties

Typical Computational/Analysis Tasks
"omics"

- Discrete algorithms – hashing, searching, sorting, comparisons, dynamic programming, indexing, similarity search
- Compression
- Statistical algorithms to distinguish biological signal from experimental artifacts/ noise
- Graph construction, traversal, partial/sub graph matching, graph partitioning
- Statistical methods on graphs e.g. Bayesian networks

Population and Social Network Analyses

- Aggregated electronic health data to predict likelihood of disease onset, treatment response, likelihood of re-hospitalization etc
- Predict demand for products, target advertising, store shelf placement
- Characterize influence in social networks

Typical Computational/Analysis Tasks
Population and Social Network Analyses

- Structural properties of graphs – PageRank, diameter, radius, connected component
- Graph spectral analysis
- Graph mining
- Machine learning, cluster analysis
- Statistical modeling and analyses
- Natural language processing

Approach

- Detailed example workflows led by application experts
- Key cases involve interplay between simulation and data acquisition – "data assimilation"
- Scenarios involving current and future state associated with origin and movement of data between workflow stages
- Definition of workflow components

Follow up

- Yutaka Akiyama– Metagenomics
- Geoffrey Fox – Components
- Jean-Claude Andre – Climate
- Philippe Ricoux – Oil exploration/reservoir management
- Joel Saltz – Medical imaging

Data Services

James Ahrens, Chaitan Baru

Attendees

- James Ahrens, LANL
- Chaitan Baru, SDSC
- Volker Markl, TU Berlin
- Ian Foster, ANL/Chicago
- Jergi Girona, PRACE/BSC
- Yusong Tan, NUDT, China
- Kenji Ono, RIKEN, Japan
- Robert Triendl, DataDirect
- Maryline Lengert, ESA
- Bodin Francois, IRISA/Univ of Rennes
- Dries Kimpe, ANL
- Kate Keahey, ANL
- Shinji Shimojo, Osaka U
- Yoshio Tanaka, AIST
- Alok Choudary, Northwestern

Describe what is needed in Common/Basic data services

Services

- Store
- Store and Compute
- Querying
- Transfer
- Archive
- Metadata
- Sharing
- Workflow
- Provenance
- Curation
- Visualization/Analytics

General Property of Services

- Performance, Quality of Service (Cost/Access Time), Length of Time of Storage,
- Location, Quality, Resilience, Online/batch, Access latency, Streaming

Who invokes the service, When

What are the most accepted deployed services – on HPC platforms?

- Services
  - Store - yes
  - Store and Compute – maybe
  - Querying - no
  - Transfer - yes
  - Archive – could be better
  - Metadata - maybe
  - Sharing – little bit
  - Workflow - some
  - Provenance – no

What are the largest gaps for BDEC science communities

- Rigid scheduling and processing HPC model
  - Permanent fast data “not allowed”
  - Solution - to provide HPC as a service
- Inflexible cost model
  - Need measure and flexible cost model
    - power, buffer sizes, bandwidth
- Too low level of service abstractions and composition
  - Lack of automatic optimization of composite services
- Type of machine
  - High utilization single user – HPC
  - High utilization multi-user – cloud
  - Online high throughput pipeline processing
    - Real-time streaming analytics
- No platform at scale to test data services
  - Software and hardware
- Need data benchmarks to measure
  - Not LINPACK, maybe TPH, other benchmarks

What kinds of optimizations /specialization are required for science communities?

- Domain-specific scientific data analytics
  - To build on top of generic capability
  - Example: Library services for data analytics
    - Idea: Idealized statistics package
  - Knowledge processing

- Custom allocation, hardware, networking
  - Heterogeneity
  - Based on specific workflow
    - Relationship between HPC and Experimental facilities
  - Procurement based on specific user workload