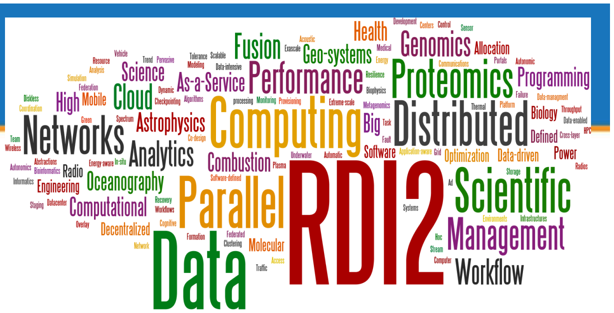

Accelerating Data-driven Science using the Cyberinfrastructure Continuum

Manish Parashar, Ph.D.
Director, Rutgers Discovery Informatics Institute RDI^2
Distinguished Professor, Department of Computer Science

Yubo Qin, Anthony Simonet, Juanjo Villalobos, Ivan Rodero

PPAM 19, Bialystok, Poland, September 10, 2019

VDC VIRTUAL DATA COLLABORATORY

Outline

- Facilities-based, data-driven since: Opportunities & challenges

- The Virtual Data Collaboratory (VDC) project: Leveraging the computing continuum for facilities-based science

  Data discovery, ***data access***, data integration

- Conclusion and next steps

Science & Society Transformed by Data & Computing

- *Nearly every field discovery is transitioning from "data poor" to "**data rich**"*
- *The **scientific process** has evolved to include **computation** & **data***

Science and Engineering in 21st Century

- New paradigms and practices in science and engineering
- Inherently multi-disciplinary
- Extreme scales, data-driven, data and compute-intensive
- Collaborative (university, national, global)

VDC
VIRTUAL DATA COLLABORATORY

Large, Shared-use Facilities can Transform S&E Research

VDC
VIRTUAL DATA
COLLABORATORY

NSF Ocean Observatories Initiative (OOI)

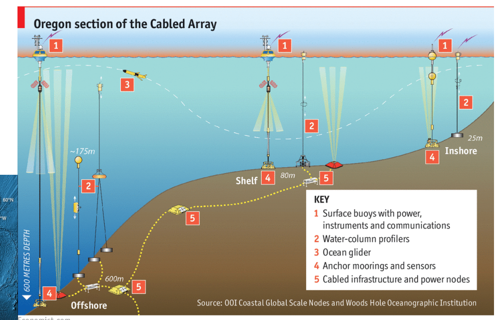

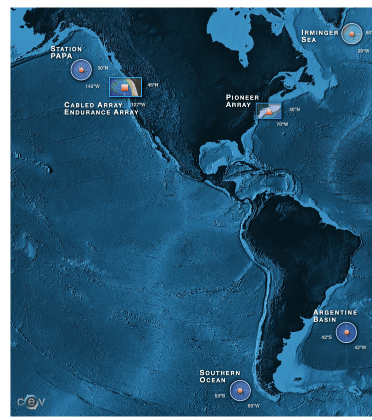

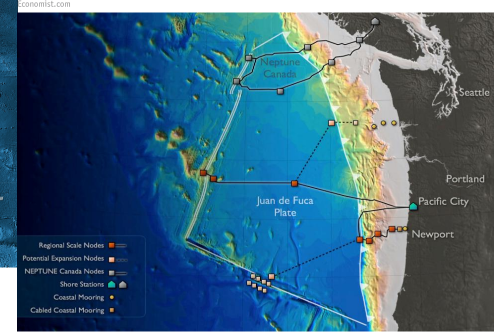

ooinet.oceanobservatories.org

Credit: John Delaney, University of Washington

vatories

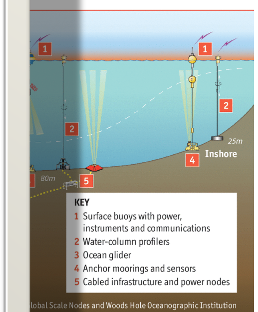

7 Arrays

57 Stable Platforms
Moorings, Profilers, Nodes

31 Mobile Assets
Gliders, AUVs

1227 Instruments (~850 deployed)

>2500 Science Data Products

>100K Science/Engineering Data Products

ooinet.oceanobservatories.org

Credit: John Delaney, University of Washington

Types of Data

Telemetered

Recovered

Cabled

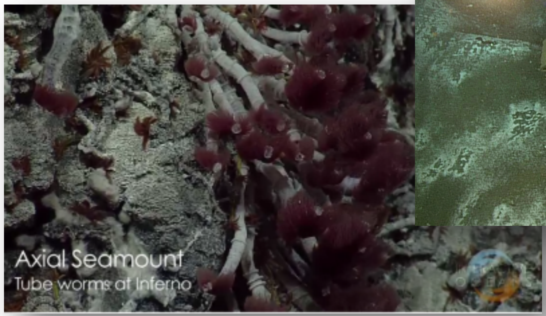

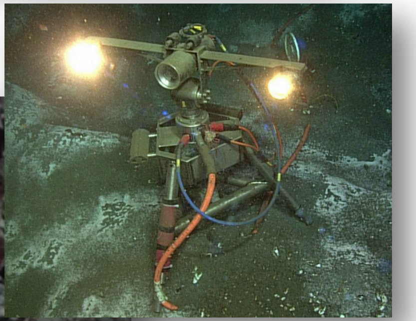

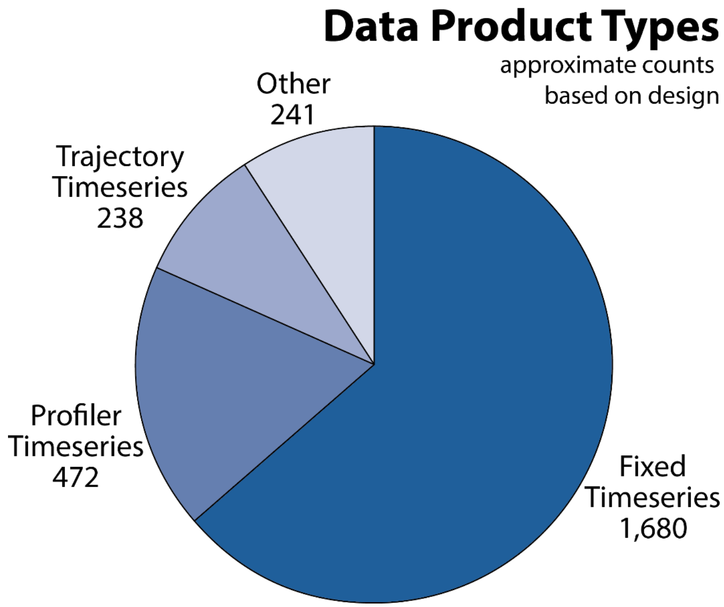

Types of Data

Teleme

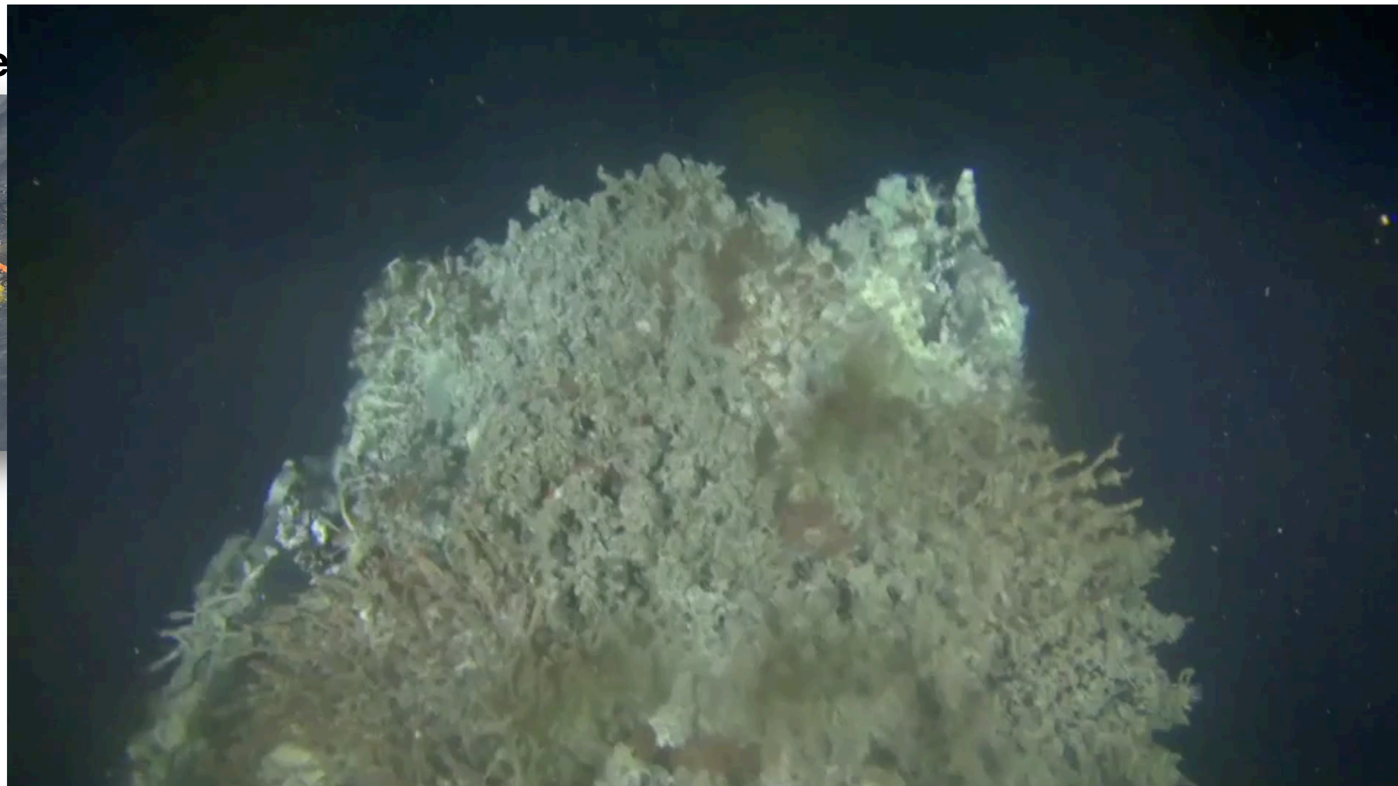

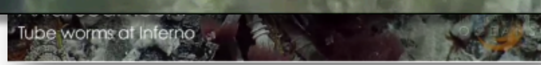

Product Types

approximate counts based on design

Data Download Statistics (Jun’16 – Jun’17)

	OOINet (UI Portal)	THREDDS Server	Raw Data Server
Visits	28,341	3,681	18,829
Distinct countries	104	36	57
Direct entries	22,446 (79%)	3,324 (90%)	17,021 (90%)
Search engines	227(1%)	51 (1%)	(<1%)
From websites	3,228 (26%)	306 (8%)	1,792 (10%)
Distinct websites	131 (540 distinct URLs)	17 (92 distinct URLs)	30 (158 distinct URLs)
Data transferred	**75.31 GB**	**923.3 GB**	**41.85 TB**

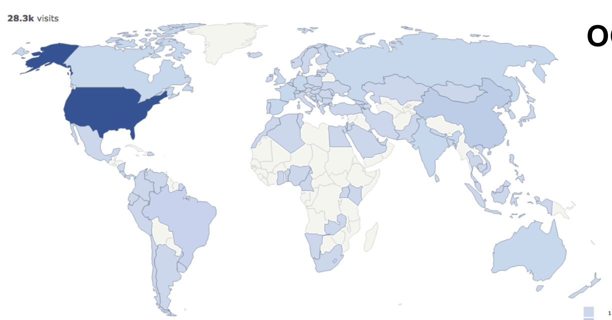

OOINet

OOINet

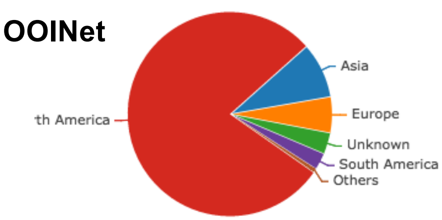

Country	Visits
United States	22,029
China	1,942
Unknown	965
Brazil	659
France	426
United Kingdom	297

Raw Data

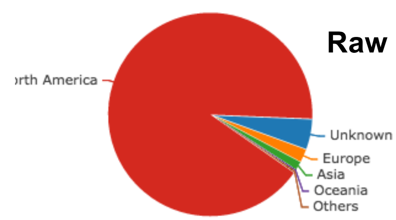

Country	Visits
United States	16,989
Unknown	929
Germany	215
Canada	127
China	116
Israel	84

Large, Shared-use Facilities can Transform S&E Research

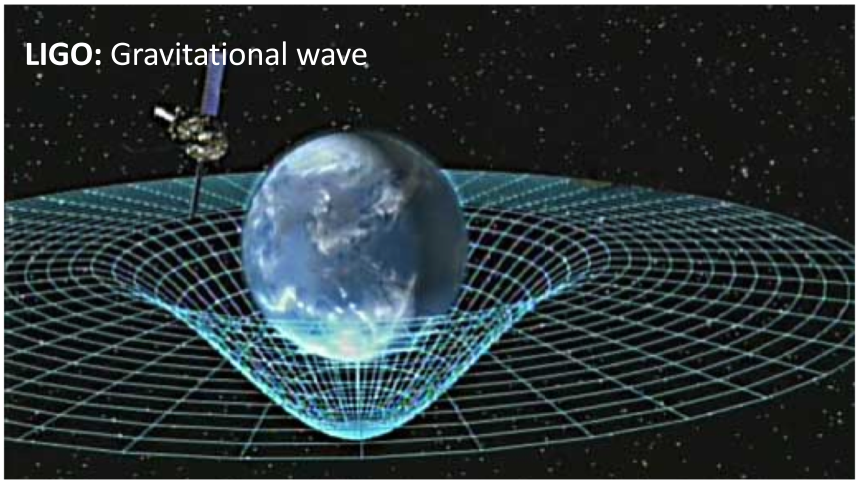

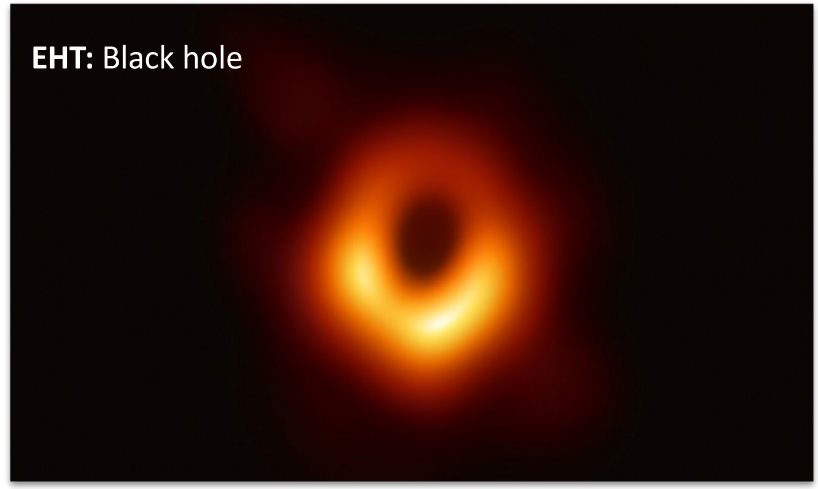

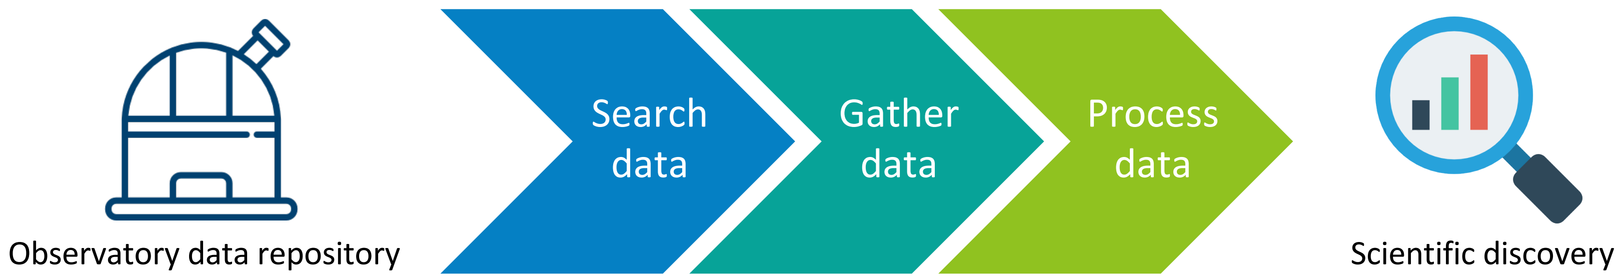

[EHT image] http://news.mit.edu/sites/mit.edu.newsoffice/files/styles/news_article_image_top_slideshow/public/images/2019/01-eso1907a_0.jpg?itok=LNUsoo9M

[LIGO image] https://www.ligo.caltech.edu/system/media_files/binaries/266/original/162571main_GPB_circling_earth3_516.jpg?1446243770

VIRTUAL DATA COLLABORATORY

Virtual Data Collaboratory: Enabling The Large Facilities Science

- Data and services provided by large-scale instruments and observatories have become **important enablers** of scientific discoveries.
- The VDC project explores how the emerging cyberinfrastructures continuum can improve the performance, usability and science impact of data and services provided by facilities.

SYSTEM PLANE	Physical infrastructure (compute, storage, network, FIONAS), operating system and networking. Virtual infrastructure management: VMs, containers, etc.
DATA PLANE	Cross-repository data indexing and discovery, provenance records and other data-related services.
KNOWLEDGE PLANE	Data analytics, cross-repository data fusion, in-transit processing.
USER PLANE	Productivity tools, streaming-based interfaces, advanced caching and prefetching strategies.

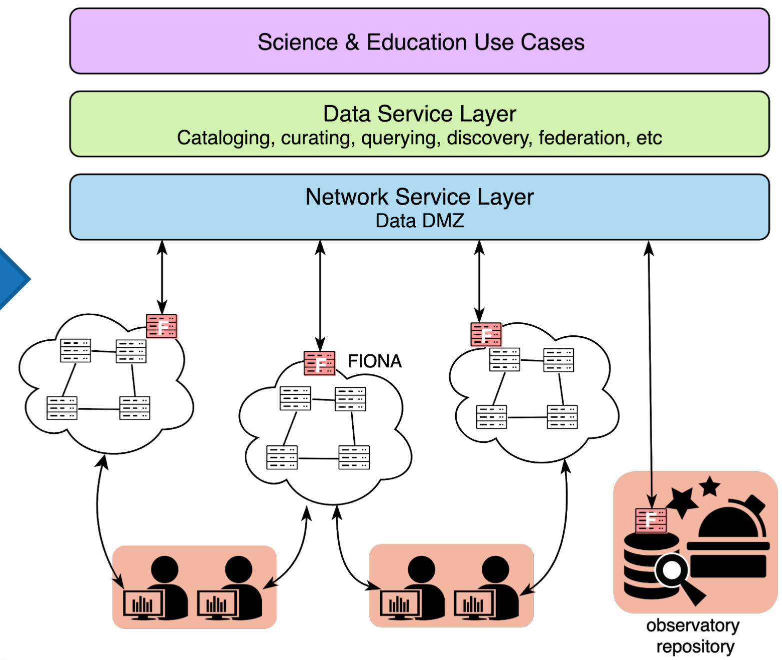

Collaborative data services to enable facilities' data to be **discovered, accessed, integrated** and **analyzed** in a timely manner

Leveraging the Computing Continuum

Emerging computing landscape

- Cloud
  - Hosted in data centers at the core
  - Relatively inexpensive; seemingly infinite
  - Far from data; data access expensive

- Fog/Edge
  - Computation/storage limited and expensive
  - Closer to the data; lower latencies
  - Limited and unreliable connectivity

- In-Transit
  - Distributed along the data path
  - Limited, but can be effective
  - Intermediate latency
  - Fewer guarantees

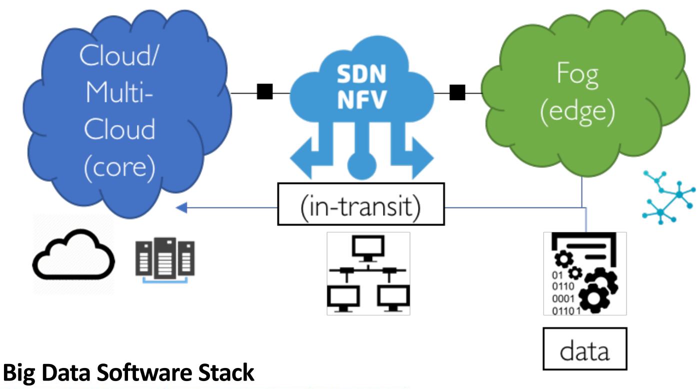

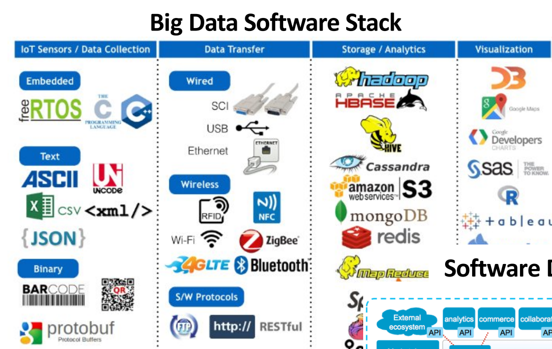

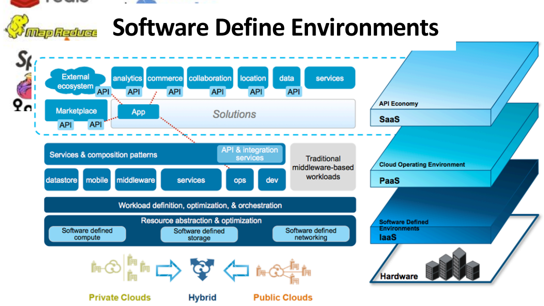

COLLABORATORY

Leveraging the Computing Continuum

Emerging computing landscape

- Clc
- Fog
- In-T
  - Fewer guarantees

Computing across the Continuum

- Leverage resources and services at the logical extreme of the network and along the data path to increase the value of the data while potentially reducing costs
- Exploit the rich ecosystem of data and computation resources at the edge so that data is not moved

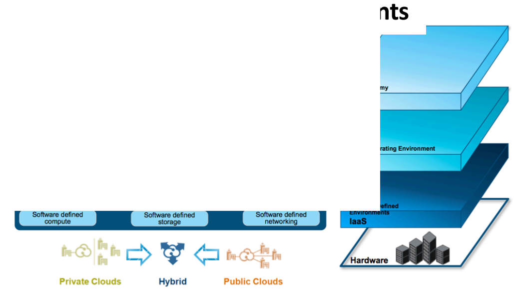

COLLABORATORY

Driving Use-case: Tsunami Early Warning

Increase precision and timeliness of tsunami warning by analyzing multiple geo-graphically-distributed data sources simultaneously,

- ***Tsunami Early Warnings*** require earthquakes to first be characterized (magnitude, location, speed of displacement, etc.).
  - A **single data source** doesn't able to cover a whole spectrum of events. *Seismometers* are good for the *smaller* earthquakes (< 6.5), *high-precision GPS* are good for *larger* earthquakes.
  - **Centralized data processing** does not support *real-time* and *high volume* of data constraints of such system.
- Goal: Combine multiple data sources to cover the **whole spectrum** of events.
- **Decentralized Early Earthquake Magnitude (DEEM):** A new two-step ensemble ML algorithm leveraging the two types of data for magnitude prediction using in-network resources

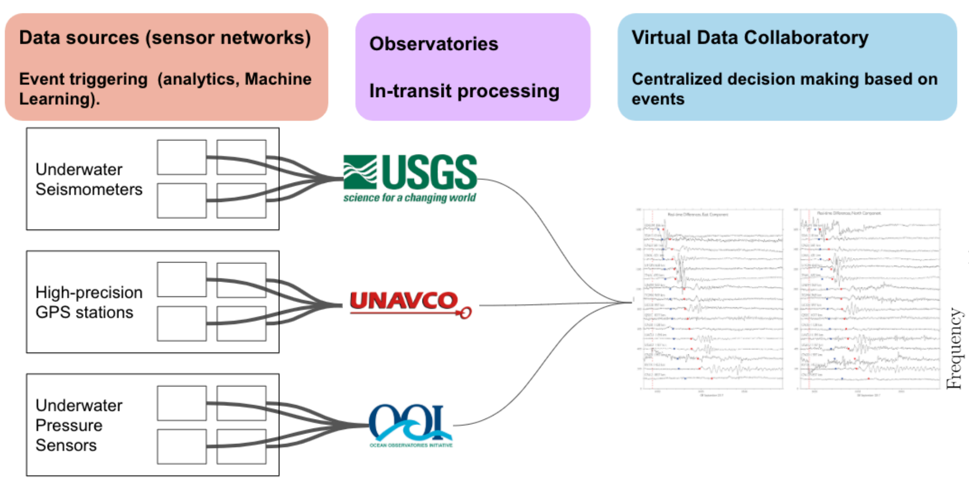

# 60s MTS	GPS (# Events)	Seismic (# Events)
Magnitude < 5	7,718 (170)	1,038 (349)
5 ≤ Magnitude < 6	3,859 (85)	None
6 ≤ Magnitude < 7	991 (4)	266 (4)
7 ≤ Magnitude < 8	432 (6)	249 (6)
Magnitude > 8	265 (4)	133 (4)
Total	13,265 (269)	1,686 (363)

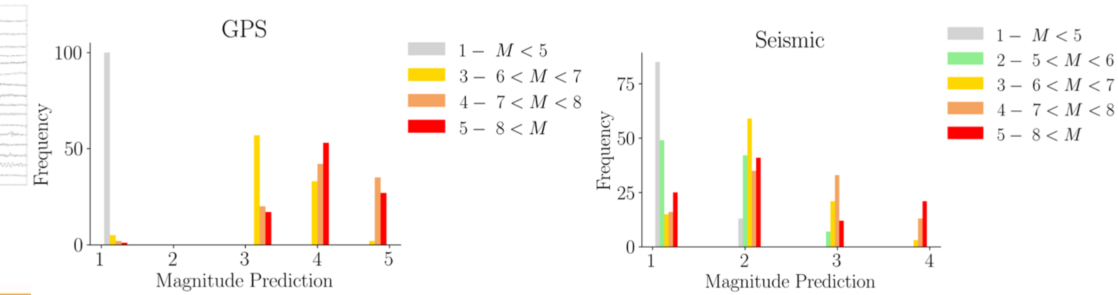

Driving Use-case: Tsunami Early Warning

Increase precision and timeliness of tsunami warning by analyzing multiple geo-graphically-distributed data sources simultaneously,

- ***Tsunami Early Warnings*** r… , speed of displacement, etc.).
  - A **single data source** … are good for the *smaller* earthquakes (< 6.5), *hi*…
  - **Centralized data proc**… traints of such system.
- Goal: Combine multiple data…
- **Decentralized Early Earth**… hm leveraging the two types of data for magnitude predictio…

Key requirements/challenges

- Data discovery

- Data access

- Data integration

… Events)	Seismic (# Events)
(170)	1,038 (349)
) (85)	None
(4)	266 (4)
(6)	249 (6)
(4)	133 (4)
; (269)	1,686 (363)

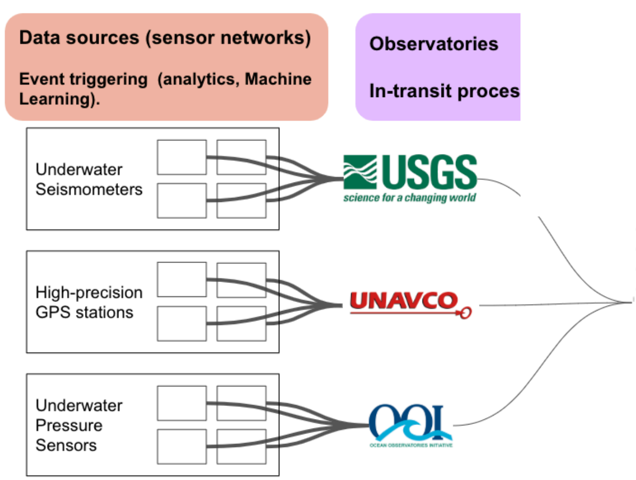

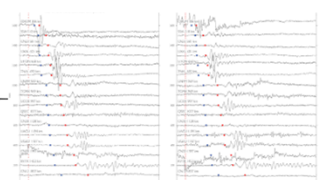

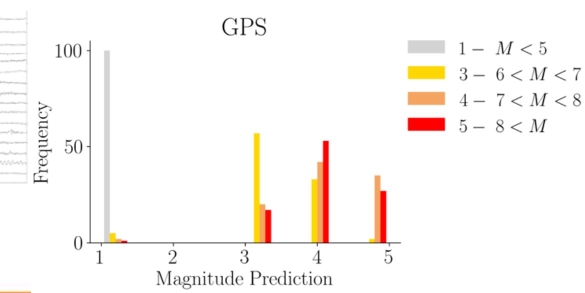

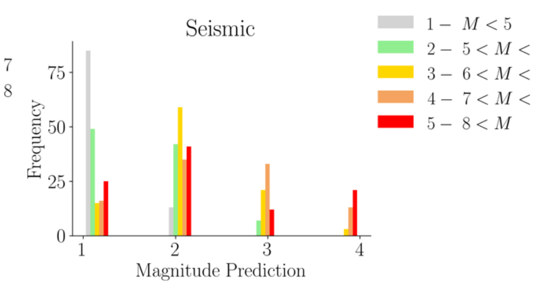

Facilities-based, Data-driven S&E: Data Access Challenge

Large volume, high data-rate, geographically distributed datasets

- Ocean Observatory Initiative
  - 1,227 instruments (~850 deployed)
  - 25,000 science data sets
  - 100,000 scientific data products
- LIGO: Generate TBs data per day, during 'observing' mode
- SKA: Estimate to generate an EB a day of raw data, which could be compressed to around 10 PB
- LSST: Produce 20 TB data per night

Large number of geographically distributed user communities with overlapping interests

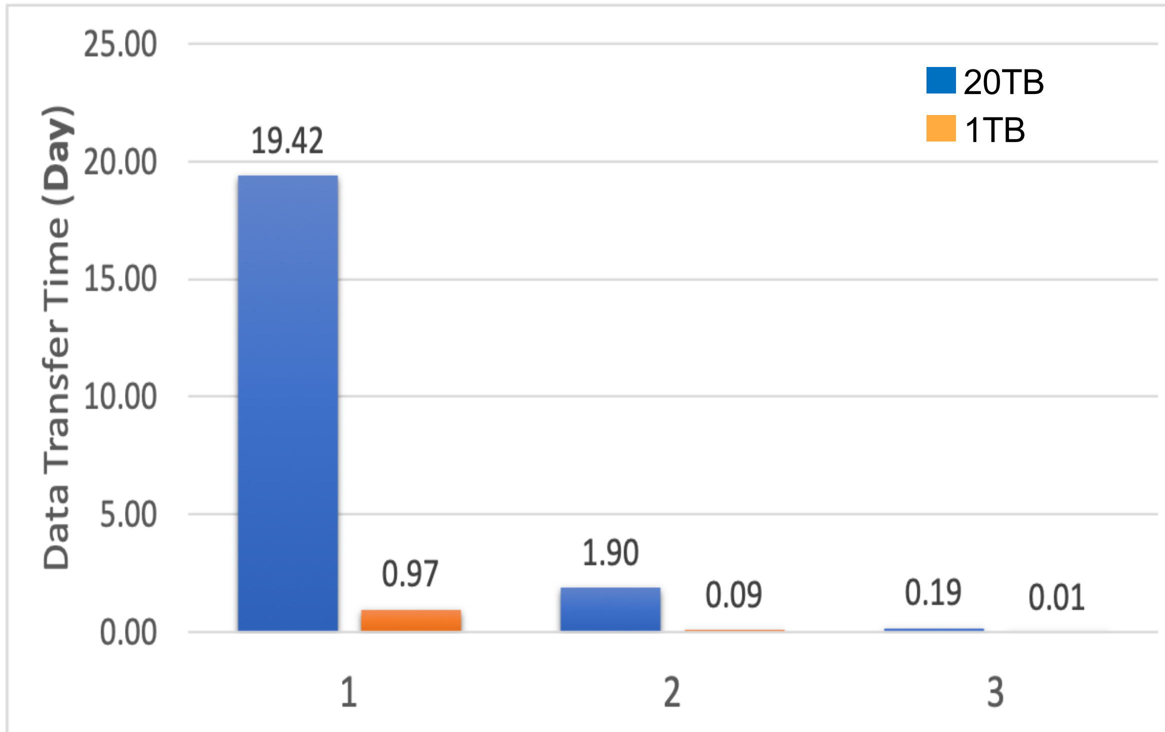

Network bandwidth: 100Mbps, 1Gbps. 10Gbps

VDC VIRTUAL DATA COLLABORATORY

Addressing Data Access: Prefetching & Smart Caching

Objective: A *push-based* data delivery framework that leverages user access patterns and locality to accelerate the data delivery performance

Approach:

- Establish a distributed cache network using in-network DTNs
- Develop a hybrid prefetching model
  - Association-based model
  - Historical record-based model

Data sets:

- Ocean Observatory Initiative
  - Nov. 2018
  - 17.9 million records
- UNAVCO
  - One year of 2018
  - 77.8 million records

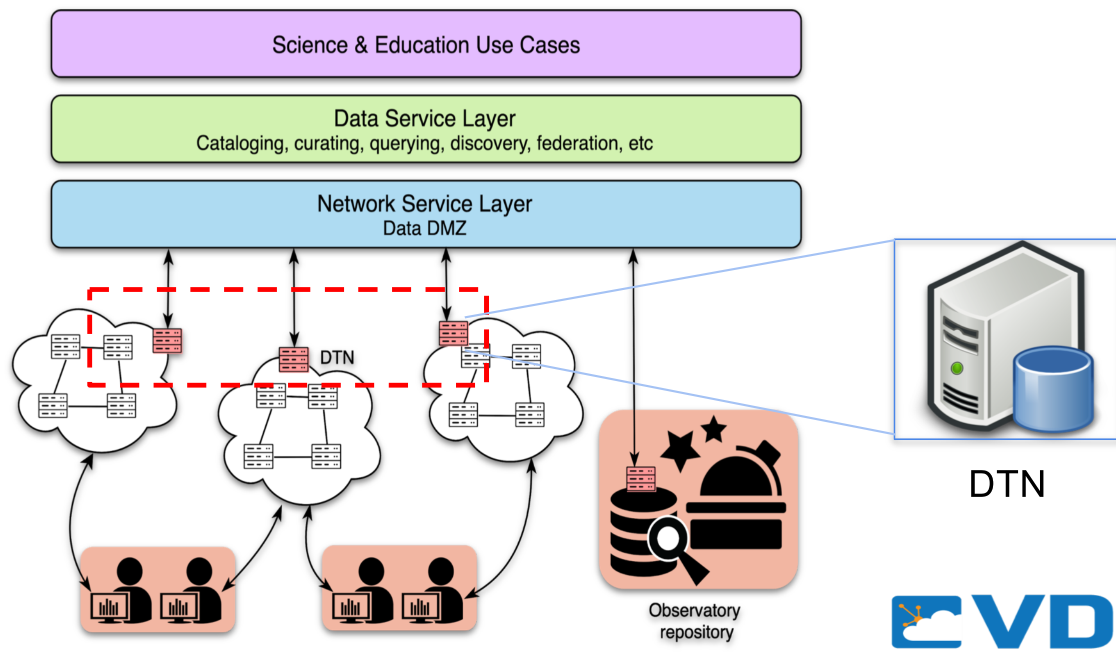

Addressing Data Access: Studying User Access Patterns

User access classification:
- Interactive access: users manually download data
- Program access: scripts automatically download data

Observations:
- Interactive users are the major users (~90%)
- Programs are the major data consumers (~90%)

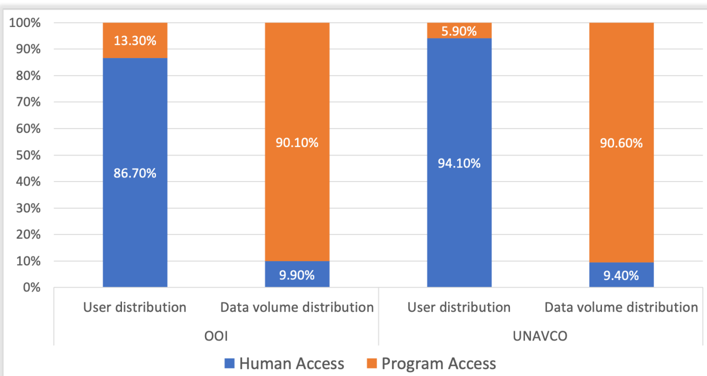

Program access classification:
- Regular data download: at regular intervals; no overlap
- Overlapping access: at regular intervals; large overlap with previous access
- Real-time access: high frequency access (e.g. every 5 sec); no overlp

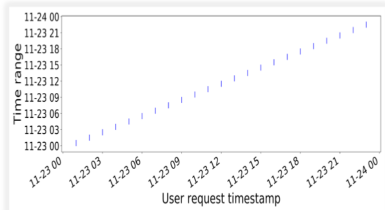

Regular data download

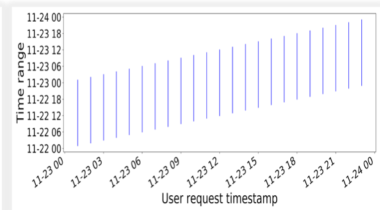

Overlapping access

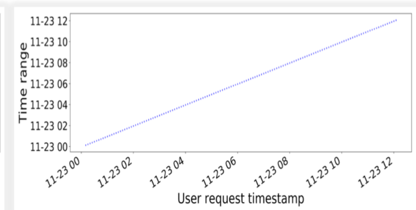

Real-time access

Data usage pattern:
- Duplicated data transfer volume
  - OOI: 60.8%
  - UNAVCO: 17.2%

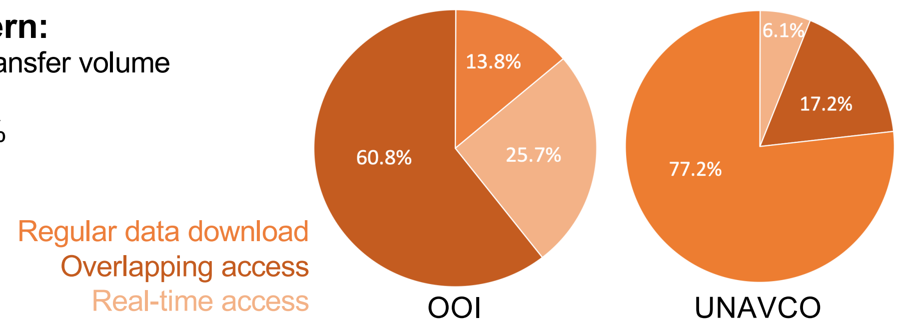

Summary:
- Focus on the **program access**: ~90% of total data transfer
- Add a **cache layer** to alleviate redundant data transfers
- Implement **prefetching:** program accesses are predictable (~90%)

VDC
VIRTUAL DATA COLLABORATORY

Addressing Data Access: Prefetching

A hybrid prefetching model:

- Association-based prediction model
  - Spatial correlation
- Historical record-based prediction model
  - Temporal correlation

Example: OOI Instruments distributed across 7 platforms

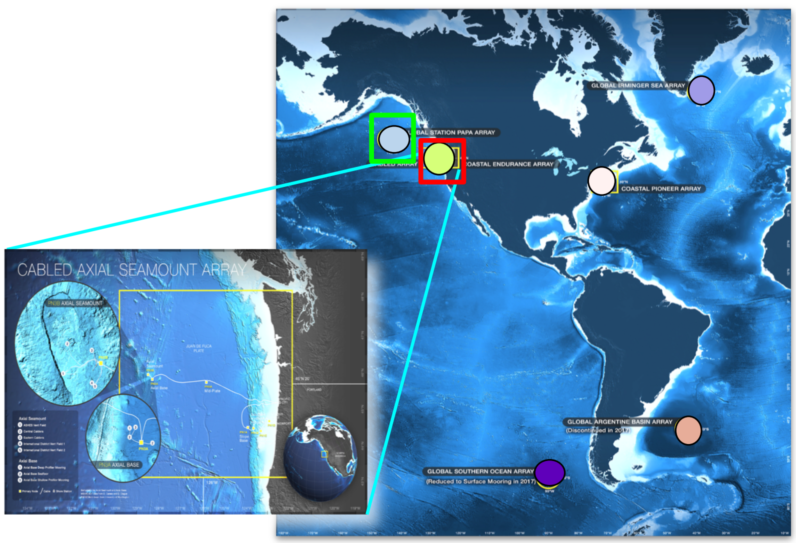

Spatial locality:

- User requests target specific locations
- Multiple users request the same data

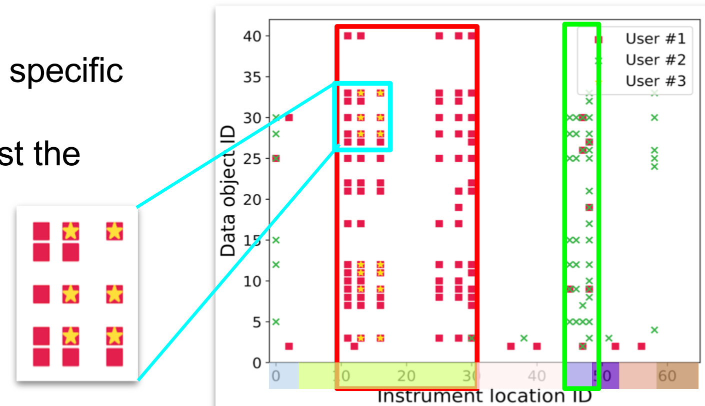

Temporal locality:

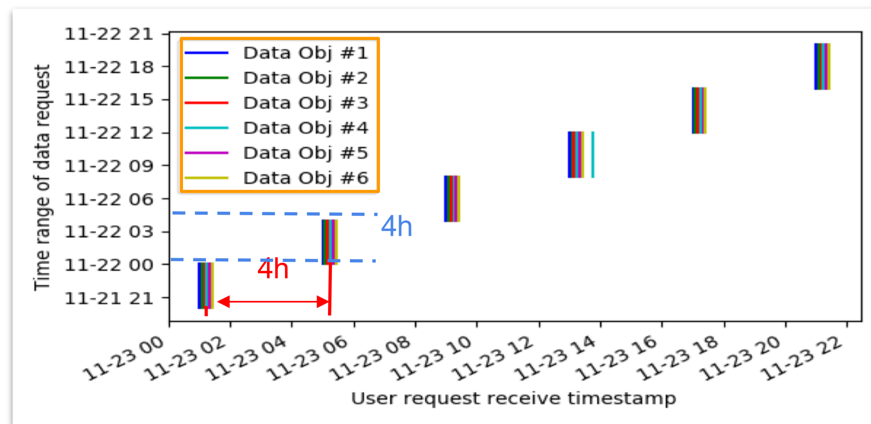

Prefetching Approach

Modeling the request sequence

Request: $R = \langle TS_n, D_n, I_n, TR_n \rangle$
Timestamp: $TS_n = \langle ts_1, ts_2, \ldots, ts_n \rangle$
Data stream ID: $D_n = \langle d_1, d_2, \ldots, d_n \rangle$
Instrument ID: $I_n = \langle i_1, i_2, \ldots, i_n \rangle$
Time range: $TR_n = \langle tr_1, tr_2, \ldots, tr_n \rangle$

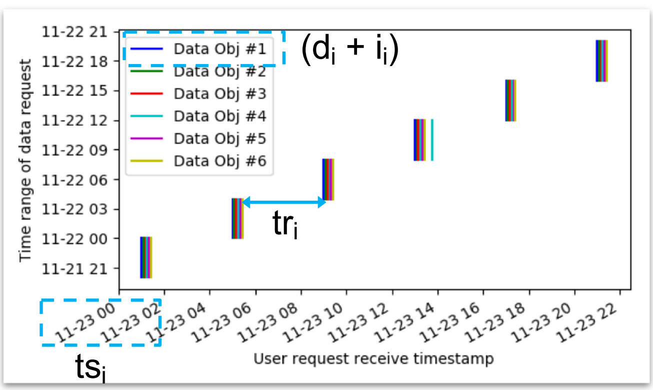

Association-based prediction model

- Exploit spatiotemporal correlation
- Use the Frequent Pattern growth algorithm (FP-growth) to predict next request
  - Construct the frequent-pattern tree (FP-tree)
  - Prefetch the data with confidence $\emptyset_j$ > threshold
    - $(d_1, d_2, \ldots, d_{i+m-1}) \rightarrow \langle d_{i+m}, \emptyset_j \rangle$, $\emptyset_j$ is confidence value

Historical record-based prediction model

- Identifying program-based access (PA)
  - Maintain a detection time window (e.g. 2 weeks)
  - Check repetition patterns for a user's request *R*
- Once PA is identified
  - Determine $\langle D_n, I_n, TR_n \rangle$
- Use ARIMA to predict TS_n

Beer and Diapers

Image: http://myshopdiscountsblog.com/wp-content/uploads/2014/06/dad-shopping.jpg

System Architecture

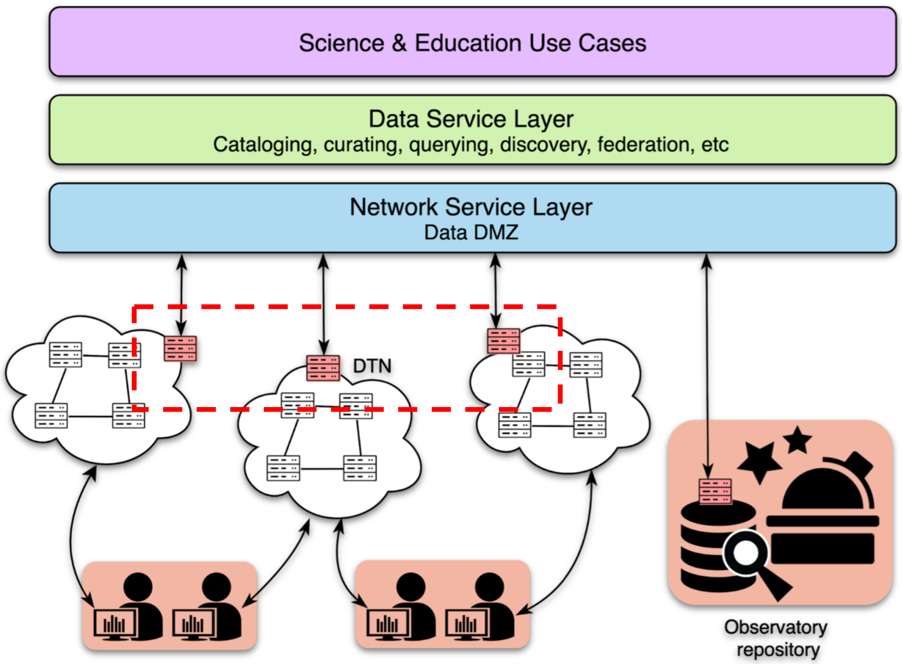

Virtual Data Collaboratory

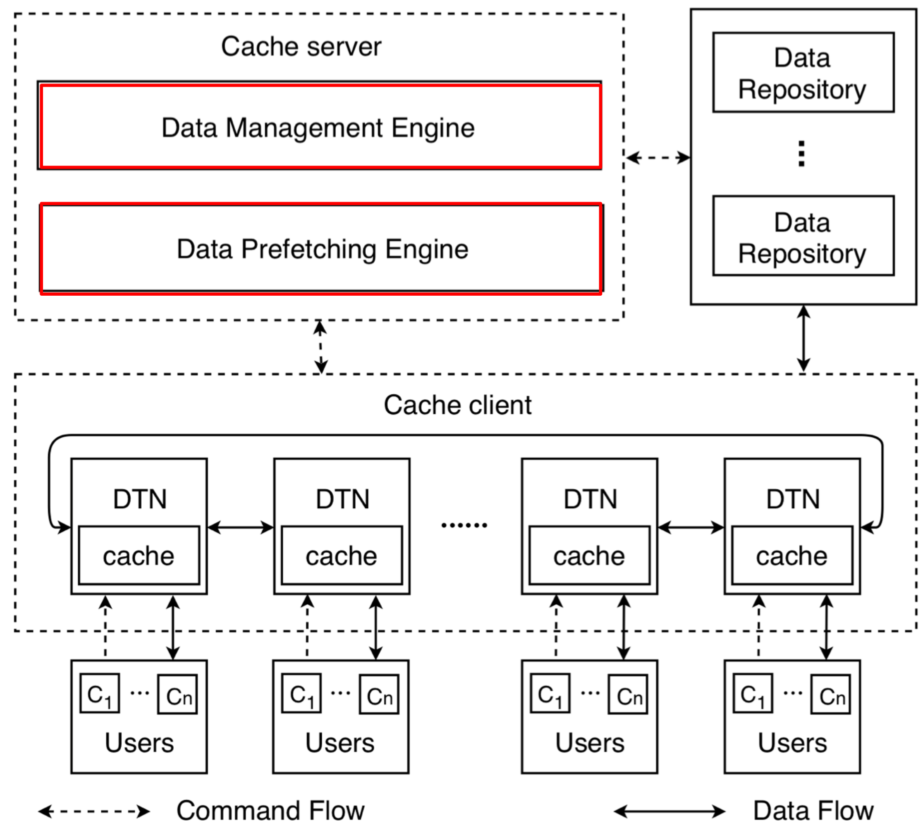

Experimental Evaluation

- System setup emulated:
  - 8 DTNs based on parameters obtained from the PRP Dashboard (Feb. 13, 2019 at 15:39:00)
  - Cache sizes: 128GB, 256GB, 512GB, 1TB, 2TB; LRU cache eviction
- Data source: OOI access log from November 2018 (17 Million records)
- Four scenarios:
  - W/O Cache
  - Simple Cache (LRU)
  - Simple Cache W/ Virtual Group
  - Smart Cache (including prefetch)

The **Smart Cache** enables users get more than **56% data** from their **local DTN cache**

The volume of data retrieved from the repository

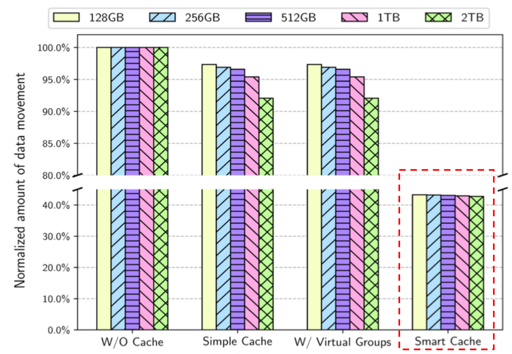

OOI

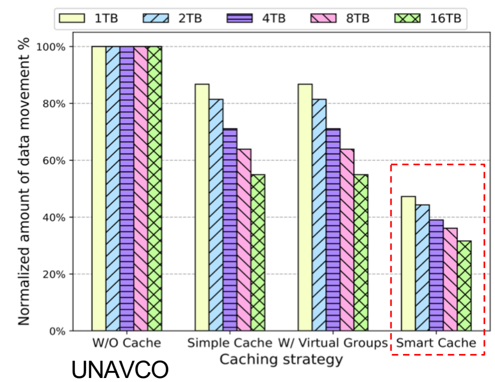

UNAVCO

The volume of data retrieved from the local DTN cache

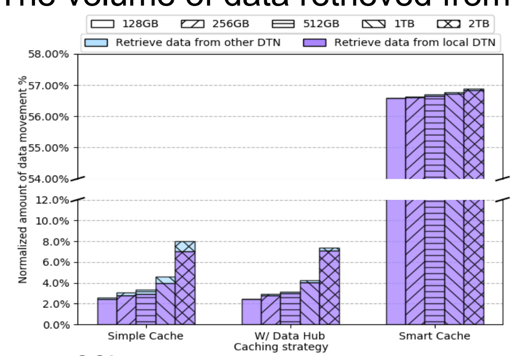

OOI

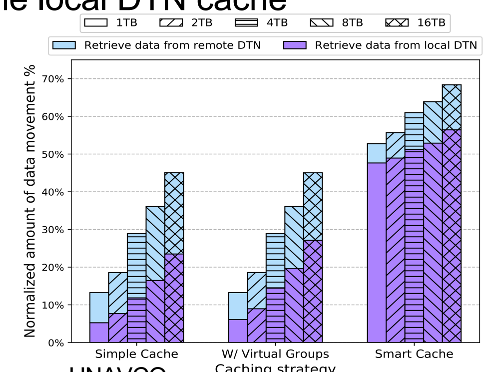

UNAVCO

Facilities-based Data-driven S&E: Challenges

#1 Data access

#2 Data discovery

- Users typically manually explore data via a web portal

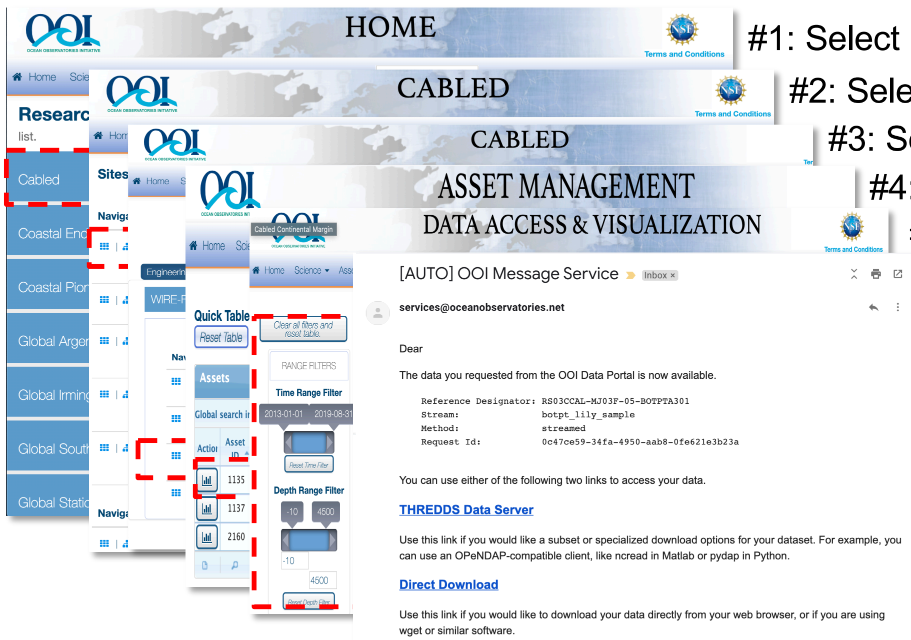

#1: Select research arrays

#2: Select the facility site

#3: Select the instrument

#4: Select the data stream

#5: Specify the data time range and click download

#6: Receive the data download link and Done!

Manually searching for data/data products is not efficient (feasible)

Remember, OOI has:

- 1,227 instruments (~850 deployed)
- > 25,000 science data sets
- > 100,000 scientific data products

VDC
VIRTUAL DATA COLLABORATORY

Facilities-based Data-driven S&E: Challenges

#1 Data access

#2 Data discovery

- Users have to manually explore data from the observatory data web portal
- Gets even more tedious when exploring data across multiple repositories

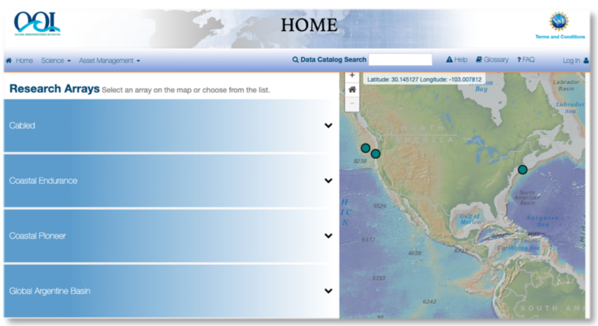

OOI

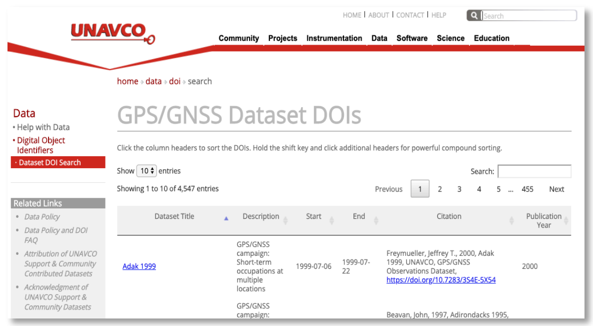

UNAVCO

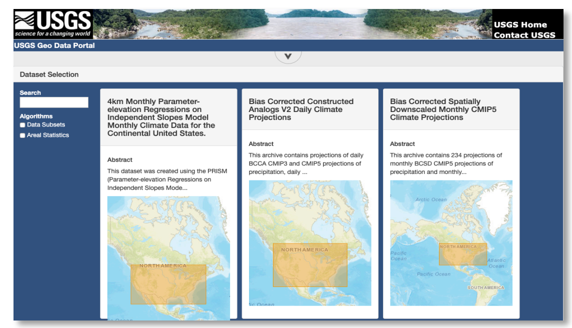

USGS

VIRTUAL DATA COLLABORATORY

Facilities-based Data-driven S&E: Challenges

#1 Data access

#2 Data discovery

#3 Data integration

Enable data from multiple data sources to be dynamically (opportunistically) integrated as part of support data-driven workflows

Example: Tsunami early warning system

- Integrate three data sources:
  - GPS (UNAVCO)
  - Seismograph (USGS)
  - Underwater pressure (OOI)

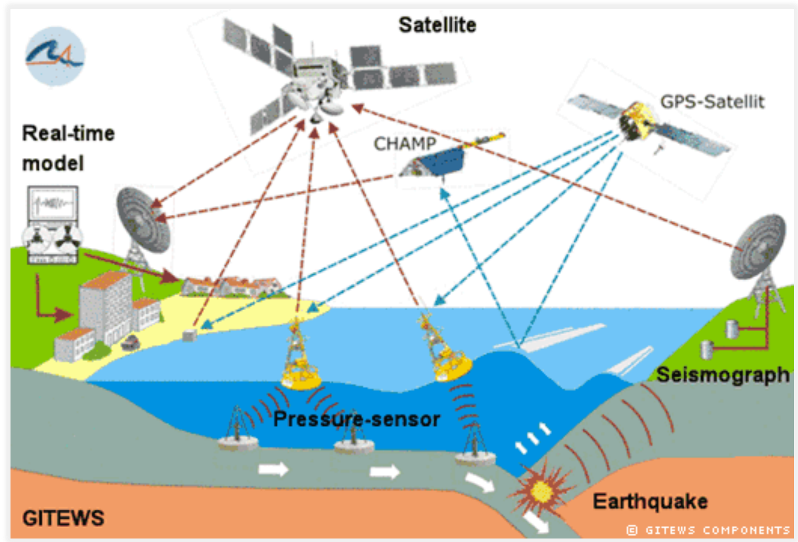

Tsunamic early warning system [1]

[1] https://www.dhigroup.com/presences/nala/usa/news/2006/9/27/~/media/Images/News/2006/20060927_GITEWS.ashx

VIRTUAL DATA COLLABORATORY

Summary

Data-driven science and engineering research enabled by large-scale, shared-use experimental and observational facilities presents new opportunities for discovery

Data distribution, size, heterogeneity presents discovery, access, integration, processing challenges

- Large scale, heterogeneous in nature and geographic location
- Data needs to be processed by complex application workflows in a timely manner

The VDC project aims to provide data services that can leverage the computing continuum to address the needs for facilities-based data-driven science

VDC
VIRTUAL DATA
COLLABORATORY

Thank you!

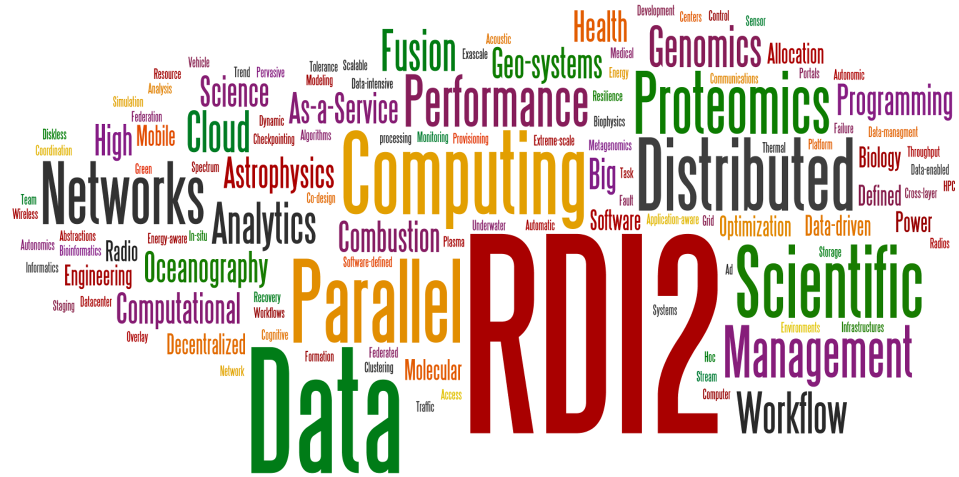

Manish Parashar
Rutgers Discovery Informatics Institute (RDI^2)
Rutgers, The State University of New Jersey

Email: parashar@rutgers.edu
WWW: http://parashar.rutgers.edu/