



Data Cyberinfrastructure for End-toend Science: Experiences from the Ocean Observatories Initiative

Manish Parashar, Ph.D.

Director, Rutgers Discovery Informatics Institute RDI² Distinguished Professor, Department of Computer Science

> PPAM'17, Lublin, Poland September 12, 2017

Outline

RUTGERS

- Science transformed by computing and data
- Ocean Observatories Initiative (OOI)
- Towards end-to-end data-driven science
- Conclusion

The Era of Big Data and Extreme Compute

- Data and computing data are pervasive
 - Extreme scales; extreme data volumes and rates
 - Novel paradigms: cloud services everywhere, cloud/fog/edge, intransit, SDN/NFV, IoT, ...
 - New technologies: accelerators, storage, communication, ...
 - New concerns: correctness, energy, fault tolerance, security, etc.
- New paradigms and practices are essential
 - Multi-disciplinary
 - Collaborative



Entire printed collection of the US library of Congress is 10 Terabytes I Exabyte is 100,000 US Libraries of Congress



Science/Society Transformed by Data & Compute

- The scientific process has evolved to include computation & data
 - New **data rich** reality in virtually every domain
- **Convergence** of opportunities and challenges



RDI²



NSF Ocean Observatories Initiative (OOI)









ooinet.oceanobservatories.org

Credit: John Delaney, University of Washington





Select critical locations, of high science impact yet data sparse due to challenging nature

Deploy cutting edge platforms

Equip with multidisciplinary sensors

Provide as much real time data as possible

Provide as much interactive capability as possible

Provide additional bandwidth and carrying capacity

Launch crucial long-term measurements that resolve high frequencies and episodic events

Make all data available









Select critical locations, of high science impact yet data sparse due to challenging nature

RMINGER EA

60°N

39°W

Deploy cutting edge platforms

Equip with multidisciplinary sensors

Provide as much real time data

SOUTHERN OCEAN

90°W



ooinet.oceanobservatories.org



OOI CI Overall Architecture



- Enterprise-level system and network architecture
- CI best practices
- Information lifecycle management
- 24/7 delivery of data & services
- Integrated software stack
- Configuration management
- Comprehensive monitoring and analytics
- Mirroring and fail-over
- Integrated cyber-security















Types of Data

















OOI CI Overall Architecture





Raytheon







Integrated Software Stack





uFrame-based OOINet – Data Ingest





uFrame-based OOINet – Plotting and download













Cyber-security Overview

- Implementation
 - Redundant perimeter firewall appliances in all sites
 - VPN IPSec tunnels across all sites
 - Two-factor authentication
 - Federated Identity Management
 - 1,250 attacks mitigated in first half of 2017
 - Cyber-security Program
 - Engagement with NSF Center for Trustworthy Scientific Cyberinfrastructure (CTSC)
 - Maintains policies and procedures (12+ documents)
 - Cyber-security scans/audits
 - Leverage Rutgers' policies and protocols
 - Training activities













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)	
Γ	Data transferred	75.31 GB	9		
k visits		OOINet th America	North America South America Others	Asia Europe Unknown South America Others	ania ers
		United China	Country Visits Actions per 1,942 Av25	Actions per A Visitions Visiting Vi 342,523 Unknown Product Control of the states of t	vg. Time on Swebsite n ⁹⁸⁹ 5s 251s
» 	OOINet ed State	es 22,029 🖓 ຢາ23,523 (,942 💭 ឆ្លៃទោល ខ ,942 💭 ឆ្លៃទោល ខ ())))))))))))))))))	Visit 965 3Me 60.9 6519 min123 8.2 4126min 15.12 Kiongdom 2397min 31.22	by te Country te Visit 5982 Upp te 5% tates B6, mg 8 5982 Upp te 25% tates B6, mg 8 50 2 Upp te 25% tates B6, mg 8 50 2 Upp to 5% 3295 1 370 German % 2295 1 370 German % 21 m 1 57 a 21 m 1 57	512s 48s 24s
	blas Decan Leadership Decan Leadership Decan Leadership Decan Leadership Decan Leadership Decan Leadership Decan Leadership	dom 297 250 1426 1,9783005 sia 4 1,9783005 sia 4 1,97830005 sia 4 1,978300000000000000000000000000	18.3 3 min 78.5 4.1 13/3 s 5 28 10.3 13/3 s 5 28 208 epa 13/7 min 35.2 3.7 14/0 s 753	1/8 1/2 1/2 1/2 1/2 1/2 1/2 1/2 3/2 1/2 3/2 3/2 1/2 3/2 3/2 1/2 1	525 105 1285 145
	Germany	133 🛁 Syonain 🤅	21.8 1305min 11.03	/st4 🔜 Riller striðiðiðiðiðiðiðiðiðiðiðiðiðiðiðiðiðiðið	1 325

End-to-end Data-driven Workflows

- Large-scale scientific facilities are an essential part of science and engineering enterprise
 - Provide open, real time access to data from geographically distributed sensors and instruments
- Increasing scale, heterogeneity, and richness of data
 - Data downloads and local processing no longer feasible
- Commercial/academic computing services remain largely disconnected from scientific facilities
- New delivery modes for data and data-products necessary

Re-examining Data Delivery

JTGERS

- Distributed content-delivery networks (CDNs) publish/subscribe semantics for data and data products
- Data-driven (content, location, quality) workflows that seamlessly orchestrate the entire data-to-discovery pipeline
 - Automated processing of data
 - Users issue query for data to be transferred from large-scale scientific facilities and processed at national/commercial resources
 - Data-driven query and processing of data
 - Users register queries and workflows that automatically trigger computation when certain types of data are available or when certain results from previous queries are observed
 - Data-driven query, aggregation, and processing across multiple data-stores
 - Allow users to define and register complex queries to couple multiple sources of data, which can potentially be processed at different ACI facilities

Edge/In-transit Data Processing

- Edge/In-transit processing
 - Integrates non-trivial resources and services along the data path
 - Transform/process data as it moves toward consumers
- Approximation and Content Delivery
 - Explore approximate computing and tradeoffs between quality, time, and resources
 - Disseminate data and data to multiple consumers more effectively
 - Using edge and in-transit resources to stage and process data
- Reducing Number of Streams
 - Streams combined based on content of data, users interests and workflow
 - In-transit resources leveraged for data aggregation/disaggregation



System Architecture

- Distributed streaming engine
 - Subscription based data delivery
- Data processing
 - Commercial/academic services (AWS, Azure, Jetstream)
- Content delivery network
 - Leverage network appliances close to users
- Broker
 - Provision services / map workflow to resources (intransit, and at the ends)
- Builds on Kafka, Storm, ...



Usecase: Digital Still Camera Image Analysis

- Processing images from an ocean floor high resolution digital camera
- Total of 50 images

RUTGERS

- Rate: image every 7 seconds
- Large images sizes
 - Image slicing at source
 - Image reconstruction at destination
- Image processing algorithms applied to detect objects





Implementations



Case Study – Base Data Streaming

• Latency: Time to deliver all slices of an image to a consumer

RUTGERS





Case Study – Early Detection



Case Study – Content Approximation

RUTGERS



Image Sequence

Case Study – Hybrid (Early Detection & Approximation)

Consumer Site



Image Sequence

Bottom Pressure and Tilt (BOTPT) – System Performance

- 20 S/sec rate, millimeter resolution of water depth
- Real-time data access, e.g., volcanic eruption detection
- Proof-of-concept using AWS, NSF CloudLab, NSF JetStream
- System implementation based on Apache Kafka/Zookeeper
- Experiments
 - # producers
 - # partitions
 - Sync/Async replication



Concluding Remarks

- Large-scale scientific facilities are an essential part of science and engineering enterprise
 - Provide open, real time access to data from geographically distributed sensors and instruments
 - E.g., Ocean Observatories Initiative
- Increasing data volumes and heterogeneity, complexity of workflows provide new challenges
 - Data streaming and workflow automation mechanisms are essential
 - Publish/subscribe data delivery, in-transit data processing, approximation, etc.





Thank you

Manish Parashar Email: parashar@rutgers.edu WWW: parashar.rutgers.edu

