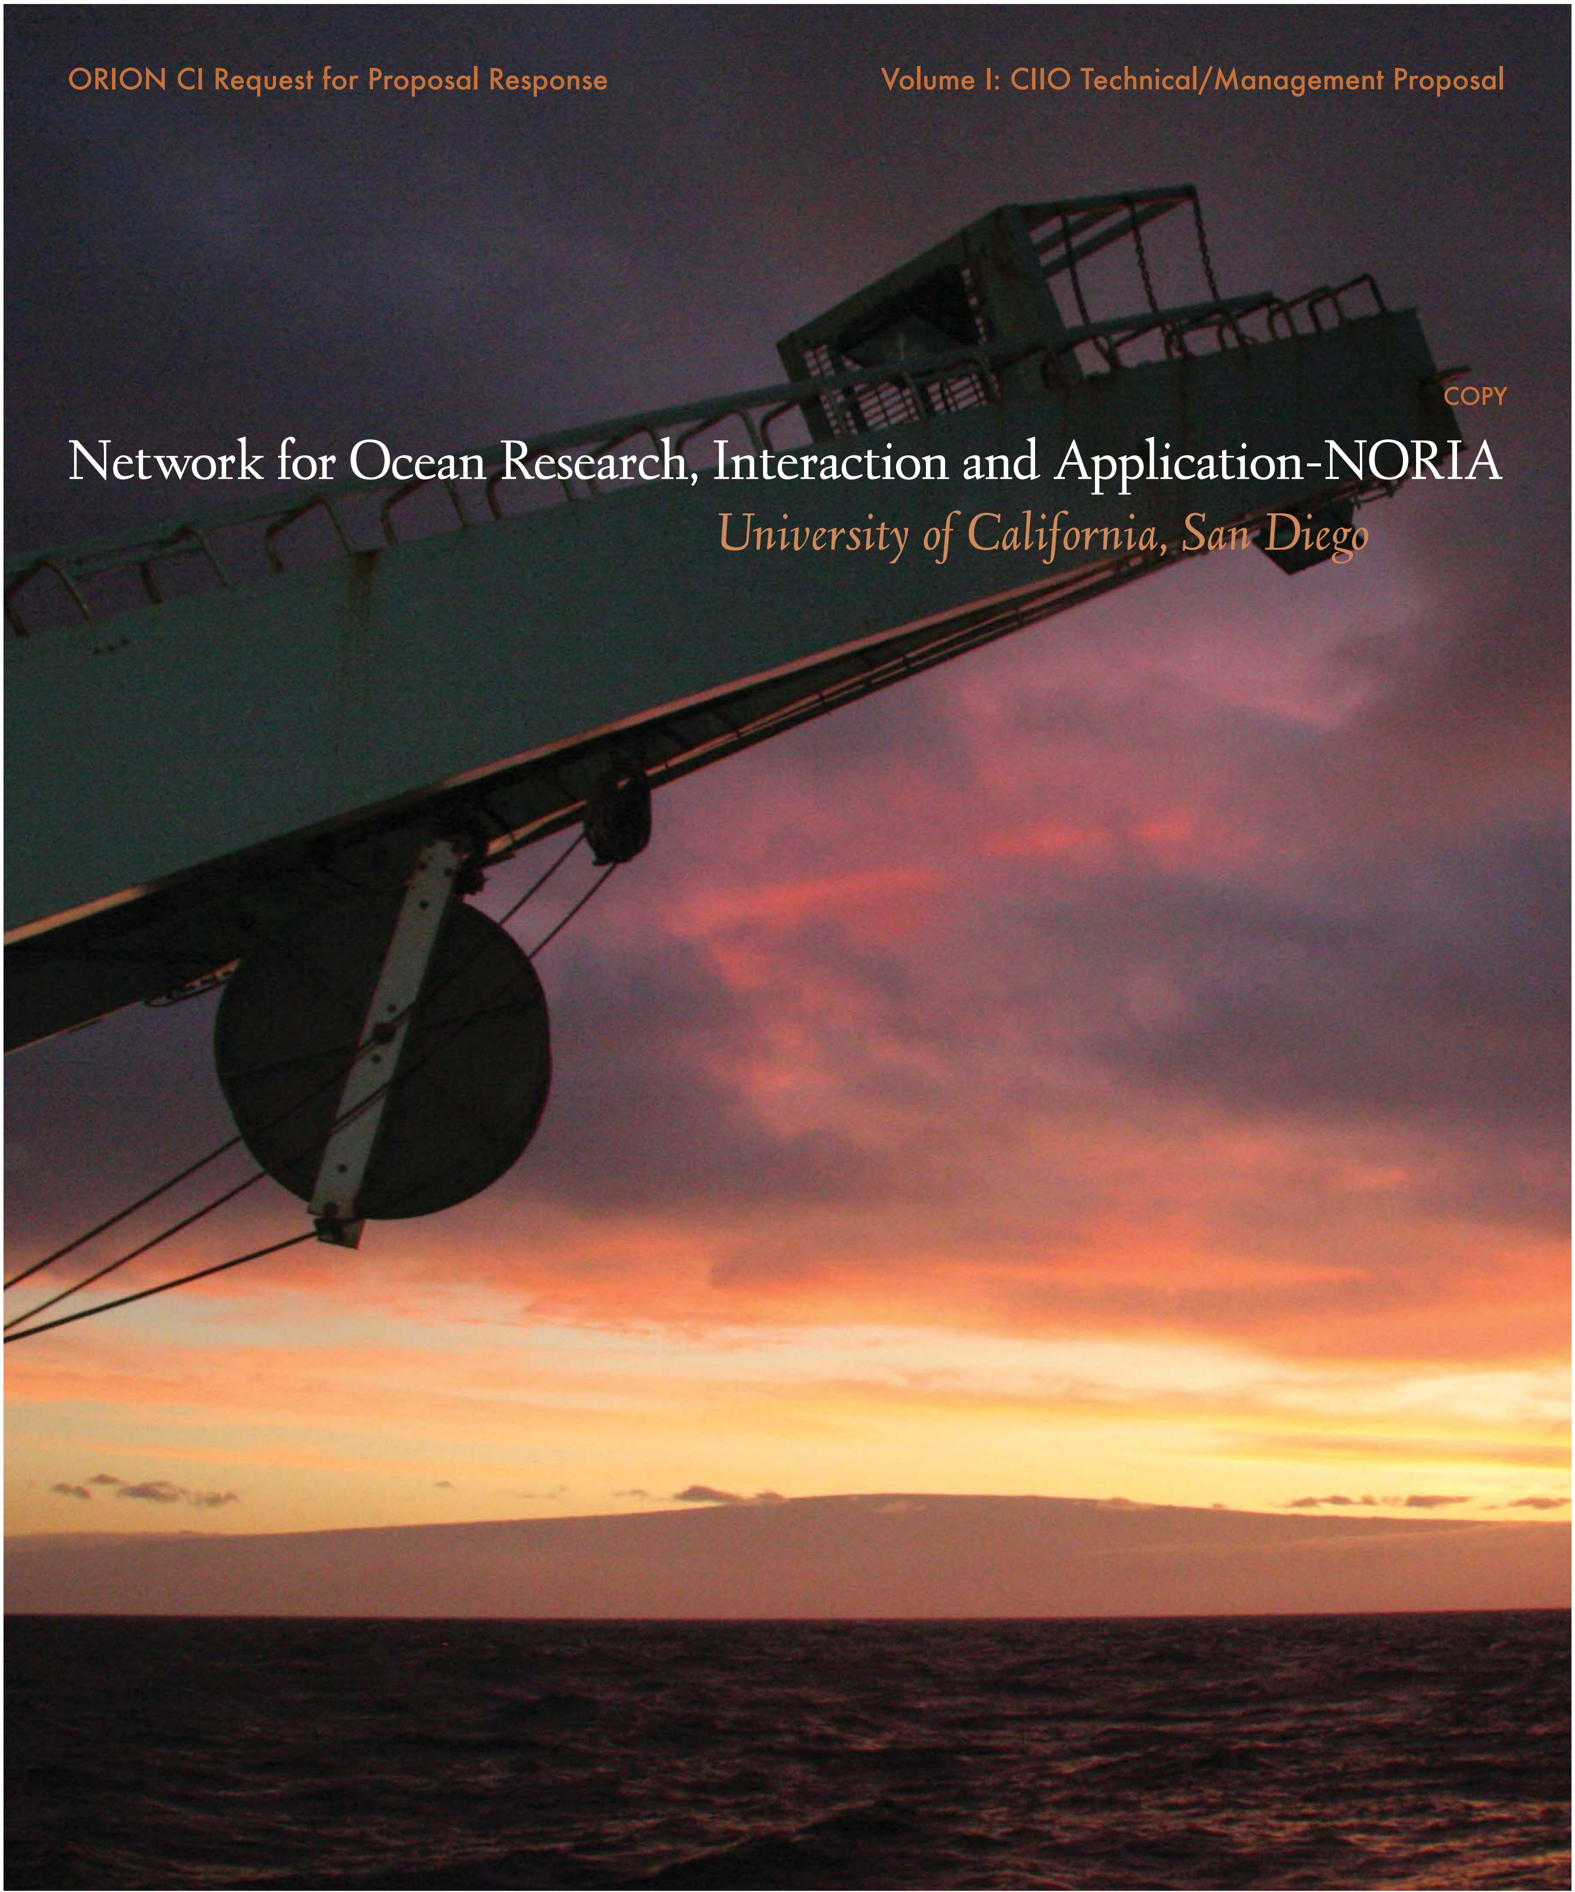

COPY

Network for Ocean Research, Interaction and Application-NORIA

University of California, San Diego

JPL/NASA MIT MBARI NCSA NCSU Rutgers UCSD Univ Chicago USC/ISI WHOI

22 December 2006

image courtesy of Dan Cartamil

The
University of California, San Diego

VOLUME I - TECHNICAL/MANAGEMENT PROPOSAL

UCSD# 20071735

Principal Investigator:

Last Name	**Orcutt**
First Name	**John**
Title	**Professor of Geophysics**
Department/ORU	**Institute of Geophysics & Planetary Physics**
Mail Code	**0225**
Phone #	**858-534-2887**
Fax #	**858-534-6354**
Email	**jorcutt@ucsd.edu**

Associate Investigators/Key Personnel:

Frank Vernon, Research Geophysicist, Institute of Geophysics & Planetary Physics, Scripps Institution of Oceanography, University of California, San Diego (858) 534-5537, flvernon@ucsd.edu

Matt Arrott, Project Manager, California Institute for Telecommunications & Information Technology, University of California, San Diego (885) 822-4633, marrott@ucsd.edu

Alan Chave, Sr. Scientist, Deep Submergence Laboratory, Woods Hole Oceanographic Institution, (508) 289-2833, achave@whoi.edu

Oscar Schofield, Associate Professor, Institute of Marine & Coastal Sciences, Rutgers University, (732) 932-6555 x 548, oscar@marine.rutgers.edu

Scott Glenn, Professor, Institute of Marine and Coastal Sciences, Rutgers University, (732) 932-6555 x 506, glenn@marine.rutgers.edu

Agency Information:

Agency Name	**Joint Oceanographic Institutions, Inc.**	If applicable, include the following information:
Contact Name	**Strat Cavros, Contracts Dept**	Award #
Contact Phone	**202-787-1631**	PA/RFA/RFP, etc # **Cyberinfrastructure/ORION/OOI**
Street Address	**1201 New York Ave., NW, Suite 400**	Other required agency information, such as DUNS Number, etc.
City, State, Zip	**Washington DC 20005**	**DUNS 17-510-4595 Cage code 7H459 TIN: 95-6006-144**
Proposal Title	**Proposal for the ORION Cyberinfrastructure Implementation Organization by the Network for Ocean Research, Interaction and Application (NORIA) Consortium**	

Project Begin Date **03/01/07** Project End Date **02/28/18**

Principal Investigator

John A. Orcutt
Principal Investigator Signature

12/08/06
Date

OCGA Official Authorized to Sign on Behalf of The Regents

Nancy F. Wilson
Signature

12/21/06
Date

Nancy A.F. Wilson, Director, Contract and Grant Administration, SIO/UCSD
Print or Type Name and Title

Please send award documents as indicated below:

Mailing Address:
University of California-San Diego
SIO Contract and Grant Office
9500 Gilman Drive, Dept 0210
La Jolla, California 92093-0210
Phone # 858-534-4570
Fax # 858-534-9642

Courier Address:
University of California-San Diego
SIO Contract and Grant Office
8622 Discovery Way, Room 116
La Jolla, California 92037
Phone # 858-534-4570
Fax # 858-534-9642

Contents

Volume One: CIIO Technical/Management Proposal

6.0 Appendices

6.1 Cyberinfrastructure Project Execution Plan (PEP)

6.2 Task Area Experience

6.3 ORION CI Facilities

6.4 Work Breakdown Structure (WBS)

7.0 Exhibits

8.0 References

9.0 Vitae

10.0 Letters of Support

Volume I: CIIO Technical/Management Proposal

ORION CI Proposal

1.0 STATEMENT OF WORK

1.1 Introduction

Oceanography is augmenting the ship-based expeditionary science of the last two centuries with a distributed, observatory-based approach in which scientists continuously interact with instruments, facilities, and other scientists to explore the earth-ocean-atmosphere system remotely **(Exhibit 1: Ocean Observatory)**. Routine, long-term measurement of episodic oceanic processes on a wide range of spatial and temporal scales is crucial to resolving scientific questions related to Earth's climate, geodynamics, and marine ecosystems. Innovative ocean observatories providing unprecedented levels of power and communication and access to real-time sensor networks will drive scientific innovation and provide education and outreach capabilities that will dramatically impact the general understanding of, and public attitude toward, the ocean sciences.

The Ocean Sciences Division of the National Science Foundation implemented the Ocean Research Interactive Observatory Networks (ORION) program to focus on science, technology, education, and public awareness activities needed to develop and deploy a network of science-driven ocean observing systems. The ORION infrastructure, which initially will be funded by the Ocean Observatories Initiative (OOI) from the Major Research Equipment and Facilities Construction account, will provide users with the means to characterize the oceans for decades.

The OOI comprises three types of interconnected observatories spanning global, regional and coastal scales **(Exhibit 2: Observatory Scales)**. The global component addresses planetary-scale problems via a network of moored buoys linked to shore via satellite. A regional cabled observatory will 'wire' a single region in the Northeast Pacific Ocean with a high speed optical and power grid. The coastal component of the OOI will expand existing coastal observing assets, providing extended opportunities to characterize the effects of high frequency forcing on the coastal environment. The OOI CyberInfrastructure (CI) constitutes the integrating element that links and binds the three types of marine observatories and associated sensors into a coherent system-of-systems. Indeed, it is most appropriate to view the OOI as a whole, which will allow scientists and citizens to view particular phenomena irrespective of the observing elements (e.g. coastal, global, regional, ships, satellites, IOOS...) to which the observations belong. The proposed work will create a CI Implementing Organization (CIIO) to construct the CI.

The core capabilities and the principal objectives of ocean observatories are collecting real-time data, analyzing data and modeling the ocean on multiple scales, and enabling adaptive experimentation within the ocean. A traditional data-centric CI, in which a central data management system ingests data and serves them to users on a query basis, is not sufficient to accomplish the range of tasks ocean scientists will engage in when the OOI is implemented. Instead, a highly distributed set of capabilities are required that allow **(see Exhibit 3: Glossary)**:

- end-to-end data preservation and access,
- end-to-end, human-to-machine and machine-to-machine control of how data are collected and analyzed,
- direct, closed loop interaction of models with the data acquisition process,
- virtual collaborations created on demand to drive data-model coupling and share ocean observatory resources (*e.g.*, instruments, networks, computing, storage and workflows),
- end-to-end preservation of the ocean observatory process and its outcomes, and
- automation of the planning and prosecution of observational programs.

In addition to these features, the CI must provide the background messaging, governance and service frameworks that facilitate interaction in a shared environment, similar to the role of the operating system on a computer.

All of this CI functionality either exists today or is in an advanced state of development. The proposed work will achieve our vision by:

- Assembling an interdisciplinary team of

- o Top oceanographers and ocean engineers from leading institutions (Scripps Institution of Oceanography, Woods Hole Oceanographic Institution, MIT, Rutgers University and Caltech's Jet Propulsion Laboratory),
  - o CI specialists and information architects from cutting-edge organizations (UCSD's Calit2 and NCMIR, San Diego Supercomputing Center, National Center for Supercomputing Applications, MIT, Caltech's Jet Propulsion Laboratory and Globus).
- Integrating "best of breed" technologies from proven, active systems in environmental sensing, data acquisition, data analysis, ocean modeling and adaptive, behavior-based observing into a comprehensive CI.
- Implementing an integration strategy that builds on experience over the past decade with Grid computing and Web development, and uses an incremental build and deploy approach to minimize risk.

1.2 Proposal Overview

The remaining sections of Volume One are devoted to technical and management aspects of this CI vision.

Volume One Organization

Section	Subject	Content
1.3	Science User Requirements	Describes our policy and proposed plan to ensure that the CI architecture is constructed around stakeholder needs. Stakeholder involvement will be facilitated throughout the project life cycle.
1.4	Motivation and Background	Details the pathway that led to the proposed architecture through an evaluation of relevant CI projects. A key objective is to capitalize on these efforts and produce original ORION software and middleware only as necessary.
1.5	High-level Overview of Architecture	Provides an overview of the proposed CI architecture that implements the ocean observatory scientific process.
1.6	Implementation	Provides specific detail on how the proposed CI will be implemented. The subsections cover the System Life Cycle, Software Development, Hardware Development, Implementation Plan, and the Operations and Maintenance Phase of ORION.
2.0	Software Release Schedule	Details the CI deployment capabilities and milestones at 18, 30, 42, 54 and 66 months.
3.0	Project Management	Describes implementation of the PEP, and includes a description of the project life cycle, the management structure, education and public awareness activities, key management personnel, a labor projection table, and subcontractor responsibilities.
4.0	Task Area Experience	Details past experience with projects of comparable size and scope to the ORION CI.
5.0	Key Resources	Gives an overview of the facilities available at USCD and its partner institutions, including infrastructure, equipment, labs and instrumentation. Information is also provided on ORION office, lab and workspace requirements and on the dovetailing of ORION and institutional objectives.
6.0	Appendices	
7.0	Exhibits	
8.0	References	
9.0	Vitae	
10.0	Letters of Support	

1.3 Science User Requirements

Design of the CI architecture must be driven by the needs of the diverse domain stakeholders who will use it. We view the delineation of user requirements as an initial phase of an ongoing stakeholder oversight process, a process that must continue through the validation and acceptance phases of the system life cycle. Extracting requirements from diverse stakeholder communities presents a challenge, as the typical user may not be familiar with many of the relevant information technologies, and will not be able to readily quantify present and future needs in a manner that will lead to a formal design. The science user requirements team led by the Project Scientist (with assistance from the System Engineer, System Architect, their support team [see Section 1.6.1, System Life Cycle], and the Education and Outreach Manager) will work with domain scientists and educators to this end.

The requirements team must construct a wide range of use scenarios and concepts of operations incorporating representative suites of sensors and platforms in close collaboration with a representative group of potential domain users. Three candidate use scenarios are described in **Exhibits 4.1 to 4.4: Use Scenarios**. A complementary process involves preparation of a Concept of Operations Document by the architecture team. A preliminary version of this document is part of the ORION CI Conceptual Architecture Document (the outcome of a conceptual design effort carried out by the ORION Cyberinfrastructure SubCommittee), and serves as a starting point for further development. The ConOps document describes what the system will do and why it will do it from the user perspective. It serves the important function of seeding stakeholder thinking with feasible capabilities and applications of the CI, and serves to bootstrap the use case scenario development effort. **Exhibit 5: ORION CI Conceptual Architecture–Concepts of Operations** contains an example adapted from the ORION CI ConOps document.

We envision an iterative process leading to increasingly more elaborate use case scenarios and concepts of operations, and hence more refined science user requirements. The use scenarios will, at a minimum, include the six examples defined in the RFP:

- Monitoring and control of a single observatory;
- Monitoring and control of multiple observatories;
- Event detection and response involving single and multiple observatories;
- Fusion of data from an observatory with an OOI-supplied or user-supplied ocean model;
- Design of field experiments in the vicinity of an observatory;
- The capability to compose virtual observatories from components of multiple physical observatories;

but can be expected to go beyond this point, as the three candidate use scenarios in **Exhibits 4.1 to 4.4: Use Scenarios** already incorporate many aspects of these examples.

The final version of the user requirements will be released as the CI Science User Requirements Document (CI SURD) under configuration control, with the Project Scientist as document custodian. The Concept of Operations Document will also be released with the System Engineer as document custodian. The SURD and ConOps Documents are subject to review by the stakeholder community and approval by the Program Office. We expect that both documents will be defined within the initial six-month planning phase of the project (see Section 2.1, Project Life Cycle and **Exhibit 6: Project Master Schedule**), although they can be expected to evolve over time as stakeholder experience with the CI grows. It is the responsibility of the Project Scientist to communicate user-initiated proposed changes to the CI design teams and the ORION advisory structure on an ongoing basis. It is also the responsibility of the Project Scientist to communicate changes in CI capabilities to the stakeholder community.

The SURD and ConOps Document are key inputs to a definition of system requirements, as described in Section 1.6.1, System Life Cycle. A preliminary System Requirements Document (SRD) was produced as part of the ORION CI Concept Architecture, and is largely derived from, and traceable to, the requirements of related observatory projects. Further evolution of the SRD will occur iteratively with the CI SURD.

1.4 Motivation and Background

ORION's construction will occur during the confluence of several significant technology innovations in Web and distributed processing: Semantic Webs, Social Networks, Grid Computing, Sensor Nets, Service Oriented Architectures (SOA), Event-Driven Architectures, Policy-Based Security and Machine Virtualization. Each offers different capabilities, and each may increase the scope and reliability of ORION while lowering its complexity and cost. The challenge to building the CI at this time of convergence is finding an appropriate integration architecture and roadmap to deliver a functioning system as early as possible, while maintaining the ability to refine and extend operating characteristics as technology evolves. To this end, the proposal team has gone to great lengths to review a wide range of related projects and technologies and include participants from projects that are a good fit with ORION.

Detailed architecture reviews were completed for a representative set of Grid projects (see **Exhibit 7: Hyperlinks to Cited Projects,** Standards and Technologies for the URLS for blue-highlighted phrases), including Teragrid, Open Science Grid (OSG), National Virtual Observatory (NVO), Biomedical Informatics Research Network (BIRN), Science Environment for Ecological Knowledge (SEEK), Geosciences Network (GEON), the Telescience Project, Linked Environments for Atmospheric Discovery (LEAD) and National Ecological Observatory Network (NEON). Digital library and preservation environments, including DSpace and Fedora were also examined. Architectures that introduced networking and instrumentation such as Optical networking, Internet Protocol Computer (OptIPuter), Canada's advanced Internet development organization (CANARIE), Common Instrument Middleware Architecture (CIMA), Java Distributed Data Acquisition and Control (JDDAC) and Grid Enabled Remote Instrumentation with Distributed Control and Computation (GRIDCC) were given particular focus.

To understand the current methodologies and future trends in real-time interactive observatory based research, key Cyber-Physical coupled systems were also examined: SIO's Real-time Observatories, Applications, and Data management NETwork (ROADNet), MBARI's Monterey Ocean Observing System (MOOS), Canada's North-East Pacific Time-series Undersea Networked Experiments (NEPTUNE), JPL's OurOcean Portal, MIT's Laboratory for Autonomous Marine Sensing Systems, Ocean.US's Integrated Ocean Observatory System, and NASA's Mars Exploratory Rover. In addition, the IEEE 1451 Standard for a Smart Transducer (sensor and actuators) Interface, IEEE 1588 Standard for Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Open Geospatial Consortium's Sensor Web Enablement Suite of Standards and the Marine Metadata Interoperability Project standards were evaluated. The collective results of these efforts are contained in the NSF Laboratory for the Ocean Observatory Knowledge Integration Grid (LOOKING) program's Year-1 and Year-2 annual reports, and are reflected in the ORION CI Conceptual Architecture.

In aggregate, these projects represent a sequence of extensions on the original concept of the Grid as a means of seamlessly linking distributed computers into a coherent whole [*e.g.*, Foster *et al.*, 2003]. The projects can be mapped onto three classes: Computational Grids that emphasize the shared use of computational resources, Data Grids that externalize data to assemble shared collections, and Service Grids that extend this functionality to externalize the behaviors needed to manipulate data (*e.g.*, instrument or simulation models). The Open Science Grid and Teragrid are Computational Grids. Data Grids include NVO, BIRN, SEEK, ROADNet and GEON. The Telescience Project, LEAD, NEON and ORION are both Service Grids and Data Grids.

These evaluations strongly influenced the ORION CI Conceptual Architecture. The proposed work subscribes to the scope and framework defined in that effort. However, the proposed architecture builds on this foundation through more complete consideration of related projects and relevant technologies. Further refinement is derived from the scope, structure and activities defined by the ORION Conceptual Network Design documents and the ORION Concept Design Review report.

The project reviews also provide information for determining when to leverage existing designs and approaches, and when to invest in innovation. Grid service approaches for distributed presentation, management, and control of networked instruments pioneered by ROADNet, CIMA, GRIDCC, JDDAC and LOOKING provide the basis for an interactive instrument model for ORION. The maturing technologies of the Grid trust services (Grid Security Infrastructure and its harmonization with Shibboleth and GridShib) provide the ability to effectively scale research collaborations between ORION participants while assuring the secure operation of their respective resources. Most importantly, to ensure universal access and long term preservation of ORION's most valuable products, the data and knowledge derived from them, the project will leverage the federated data management, archive, and mediation designs already established in other initiatives, notably BIRN, GEON, MMI and the Storage Resource Broker (SRB). The policy mechanisms supported by the NSF-funded Integrated Rule-Oriented Data System (iRODS) will be used to automate resource management operations. In addition, technology reviews suggest that the core CI architecture should leverage the integrative principles of a Service-Oriented Architecture (SOA) approach built on a new breed of light weight, agile Enterprise Service Bus (ESB) containers, such as the PicoContainer and Spring. SOA and the agile service container are further discussed in Section 1.6.2, Software Development.

1.5 Overview of the Architecture

The ORION CI Conceptual Architecture articulates an ocean observing system that supports the design, deployment, and prosecution of programs that investigate natural phenomena and obtain objective and reproducible knowledge **(Exhibit 8: ORION CI Conceptual Architecture-Observatory Network Model)**. The final CI architecture must be capable of supporting scientific methods used implicitly or explicitly by scientists on a daily basis. In order to serve science, observational programs must be interactive, easily coupled with other programs, and continually operational.

The OOI comprises four interoperating infrastructure elements run by four Implementing Organizations (IOs; see **Exhibit 9: The OOI Construction Program for ORION**). Three of the IOs (global, coastal and regional cabled) are responsible for the construction and operation of marine facilities providing the physical, power and communications infrastructure needed to sustain a long-term ocean presence. The CyberInfrastructure Implementing Organization (CIIO) is responsible for the interaction and interoperation of the three marine facilities with the research community and external entities, such as the Integrated Ocean Observing System (IOOS) and the international Global Earth Observing System of Systems (GEOSS) program.

1.5.1 *Capabilities of the Cyberinfrastructure*

Throughout the observatory network, strategically placed computation and storage infrastructure will be provided at Cyberinfrastructure Points of Presence (CyberPoPs; see **Exhibit 10: Cyberinfrastructure Point of Presence**). These include integrated real-time data processing and archive sites located at a few central facilities and at marine observatory shore stations or control centers. *In situ* computation and storage resources will be located within science instrument interface modules (SIIMs) and selectively within the marine networks. Large computational models and the Virtual Ocean Simulator will be run on the national Grid infrastructure (*i.e.*, the Teragrid and the Open Science Grid). Finally, observatory participants will be able to securely incorporate computation and storage capabilities within their instruments, instrument platforms (*e.g.*, AUVs and gliders) and research facilities.

Use of CyberPoPs will be prioritized based on latency, bandwidth, computing capacity and user specifications. The bulk of the real-time processing and all long-term data cataloging and archiving will occur at the central CyberPoPs. Data quality control, and short-term (less than six months) archiving will, by default, occur at the marine observatory CyberPoPs. All short-term archives will be replicated at the central CyberPoPs with the exception of HD video, which will be kept at the originating site until such time that cost of storage and transit are less prohibitive. Lower resolution representations of the HD data (*e.g.*, mpg, Quicktime®) can be made available more generally. Instrument control, data acquisition, and resource scheduling will be located where continuous connectivity can be maintained. In the case of buoys, this may mean on the wet side of the communications link, while for cabled installations it means at the marine observatory shore station or operations center. Low latency automated instrument control and/or peer-to-peer interactions will tend towards the distal ends of the marine network.

The CI, layered atop the physical computing, storage, and network infrastructure, will provide a scalable framework for the dynamic acquisition and distribution of data streams that can be coupled with immediate-mode, on-demand processing across the integrated resource network. The architecture supports the composition and deployment of user-specified processes and functionality, giving users the ability to arrange system components and construct a wide variety of data collection and modeling scenarios. A representative set includes:

- Data discovery and acquisition from real-time data streams and historical data archives.
- Controlled ocean observing using:
  - Single-task sensor systems with configuration and operating state control.
  - Multi-task sensor systems with control of task.
  - Articulated sensor systems with control of movement.
  - Mobile sensor platforms with control of geospatial location.
  - Coordinated sensor system clusters with control over scheduling.
  - Autonomous behavior-based sensor systems with collaborative collection, analysis, and control capabilities directed by defining behaviors and setting objectives.
- Hind-, now- and fore-casting of the ocean through computational modeling by assimilating observations.
- Automated ocean surveillance using coupled data collection and modeling components with event detection and classification capabilities.
- Direct environmental participation using nested clusters of behavior-based adaptive sensor systems (*e.g.*, coordinated fleets of AUVs).
- Collaborative investigation that coordinates research campaigns across teams of researchers.

The proposed architecture makes a clear distinction between the activities that users engage in (collection, assimilation, surveillance and adaptive response) and the resources (data, instruments, networks, analysis processes, computational models and behavior systems) they employ to complete the activities. The interactive nature of the activities and the constraints imposed on the resources represent a substantial extension to data analysis systems as typified by the Data Grid architectures of NVO, BIRN or GEON (see Section 1.4). While data management and analysis are essential elements of the proposed CI, they are not sufficient to support the interactive observation and response activities in cyber-physical coupled systems. The proposed architecture extends the capabilities of Data and Computational Grid architectures by incorporating the ability to employ, couple, and control shared resources (computing power, bandwidth, dynamic data sources, data collection/assimilation systems, and on-demand processing) operating across the ORION enterprise in real-time over extended periods. Representative additional capabilities include:

1. Standardized data, measurement, and process semantics with a common metadata model.
2. Interactive instrument network providing secure operation and coordinated utilization of resources.
3. Automated real-time data acquisition, preservation, and distribution with associated metadata.
4. Immediate-mode execution of user- and system-supplied processes.
5. Data processing coupled with data acquisition on time schedules and/or event triggers.
6. Extensible repository of processes for data calibration, validation, analysis, visualization, and assimilation based on established ORION semantics.
7. Interactive ocean modeling system providing access to multiple community-based numerical ocean models for parameter estimation/optimization and data assimilation.
8. Virtual ocean simulator that produces model ocean fields on a wide range of scales.
9. Planning and control framework to design, test, and prosecute observational programs.
10. Extensible knowledge management system able to deliver, exploit, and annotate relevant data from the observatory based on a localized semantic frame of reference (ontology).

These ten capabilities require six resource elements for their implementation that inform core aspects of the architecture. At the top level, the Instrument (2,3) and Modeling (5,6,7,8) elements circumscribe the key purposes of a research ocean observatory. At the next level, Data elements (1,3) support the ability to persist and catalog data and data products. Workflows and analysis are implemented through Processing elements (4,6). Interactive operation in a shared environment requires Control (2,9) elements. Finally, Knowledge (10) elements extend the concept of Data elements.

1.5.2 Architectural Organization and Abstraction

Activities represent user objectives and methodologies that establish the requirements and ultimately the prioritization of the architecture's development. *Activities* are the verbs of the system. *Resource* elements represent the actors and information needed to support the portfolio of activities. The capabilities of, and constraints on, resources define the scope of the system, and are its nouns. The infrastructure represents the organizing principles and operating frameworks that provide coherence, reliability and scale by normalizing the interactions exerted on the resources by different activities. The infrastructure is the syntax or sentence structure of the architecture.

1.5.2.1 Organization of User Activities.

User activities within ORION are organized into three categories: scientific investigation of natural phenomena, management of resources throughout their life cycle, and management of collaborations between participants. Support for scientific investigation is the central purpose of the observatory. The other two activities are a natural consequence of working within a network of resources and operating within a community.

The Observatory Network Model **(Exhibit 8: ORION CI Conceptual Architecture-Observatory Network Model)** proposed in the ORION CI Conceptual Architecture depicts a combination of the investigative process and some of the resources it employs. The proposed work expands the set of activities and provides a process model that structures their interaction and compartmentalization **(Exhibit 11: Scientific Investigation Process)**. The model supports the real-time coupling and aggregation of the main *observe, model* and *exploit* activities as well as their respective sub-activities. These can operate in isolation or be coupled with or without direct human interaction. Continuously relying upon a human in the loop is at once expensive, exhausting, and non-scalable. As discussed more thoroughly in

Section 2.0, Software Release Schedule, the order of CI development follows from the core activity of *observing* through its coupling to *modeling* and finally to the *exploitation* of the aggregate capability of Observing and Modeling.

The ORION CI Conceptual Architecture proposed an overarching set of activities for provisioning, managing, and utilizing resources. In chronological order, these represent the life cycle for the operational state of any resource in the system, starting with its inception and ending when the resource ceases to be an operational entity **(Exhibit 12: Resource Life Cycle)**. The proposed architecture will utilize this model, including the delineation of the process steps within each activity, and supports its application across all resource elements.

1.5.2.2 Organization of Resources.

The resource elements extending beyond storage, data and computation needed to realize the high level CI capabilities can be generalized as:

- Dynamic data sources (*e.g.*, data, product and event streams).
- Taskable elements (*e.g.*, instruments, AUVs, and ocean modeling systems).
- Executable processes (*e.g.*, behaviors, workflows, scripts, and compiled code).

By leveraging and structuring the shared characteristics of the complete set of resource elements, the architecture provides an extensible model for the operation, presentation, and composition of resources. The architecture then provides a standard basis for provisioning, managing, and sharing resources across the physical, technical, and organizational domains of the observatory network. This resource network model provides a consistent and scalable approach to the federation, governance, and operation of all types of resources; one that enables pervasive and universal access without compromising control and security.

The resource network concept facilitates the federation, governance, and management of interactions between system activities and a base set of resource types. It includes information as well as managed and taskable resources. The information resource supplies the interaction semantics for accessing (acquisition and presentation of) data and its associated declarations of identity along with structural and semantic contexts. Information resources can represent immutable (unchangeable), mutable (changeable), and dynamic (continually changing) content with respect to the activity using them. The managed resource extends the information resource to include declaration of operational state and governance context. The taskable resource provides the interaction semantics for the control (configure, command and monitor) of a managed resource that has a dynamic, command-driven operational state. As a simple example, a 1-D time series (*e.g.*, temperature) can be acquired including relevant metadata. The sensor can also facilitate a change in sample rate through an easily understood interface.

The organization and structuring of resources taken in the proposed architecture recognizes the process-driven nature of an ocean observatory, ranging from simple data collection protocols to the complete Investigation Model. **Exhibit 13: Activity Resource Model** illustrates the relationships between the investigation model and the six core resource networks it employs: the Control, Data, Processing, Instrument, Modeling, and Knowledge Networks. The relationships between activities, resources, and outcomes are defined on the left side of the figure. The purposes of the six resource networks are:

- The Control Network establishes standard models for the management of stateful and taskable resources. It provides the semantics to monitor and control the operating state of an active resource as well as to initiate, monitor, and amend the tasks of a taskable resource.
- The Data Network provides an automated data distribution and preservation network with pervasive and universal access subject to ORION Data Policy. It provisions a federated system of data streams, repositories, and catalogs that supports the distributed organization of resources.
- The Processing Network provides immediate-mode scheduling of processes at specified locations within the integrated network based on explicit time requirements and/or event triggers.
- The Instrument Network provides interactive and coordinated relations with real and/or synthetic environments through the use of transducers (sensors or actuators). It ensures the safe and secure operation of individual sensing platforms, and provides reliable delivery of acquired data with their associated metadata. These capabilities must be integrated with network-wide resource allocation and observation planning.
- The Modeling Network establishes baseline processes and tools comprising a coherent framework for the analysis and assimilation of data. The capability to network and interact with multiple community-based numerical ocean models for parameter estimation/optimization and data assimilation is integrated into the framework.

- The Knowledge Network provides the capability to deliver, exploit and annotate data from across the observatory based on a localized semantic frame of reference (ontology). The Knowledge Network will provide mediation between different protocols, content standards, vocabularies and concepts, and will capture and exploit the characterization and correspondence of data to concepts as well as concepts to concepts.

The resource networks are further described in Section 1.6.2. Their uses are defined in Section 2.0, Software Release Schedule.

1.5.2.3 Organization of the Infrastructure.

A key innovation of the proposed CI lies in a comprehensive strategy to support technology integration without depending on any particular software or hardware component. The CI will be built around existing and proven "best of breed" technologies. The heart of this integrative strategy lies in a core infrastructure that frames provisioning, interoperation, and management of activities and resource elements circumscribing the CI. The proposed work decomposes the infrastructure into two main elements: the Common Operating Infrastructure (COI) and the Common Execution Infrastructure (CEI). The COI focuses on operation of the system, while the CEI supports provisioning and management of the system.

The Common Operating Infrastructure provides a cross-cutting common platform that supports collaboration between individuals, groups, and institutions through the controlled sharing and aggregation of their capabilities. The COI is the core set of frameworks that establish consistent semantics for communication, state, execution, governance, resources, services, and presentation across the ORION enterprise. Each framework implements the domain model for one of these operational aspects of the system. The frameworks establish the interaction interfaces for and between their elements of the system. Using the capabilities of dependency injection and process orchestration provided by an Enterprise Service Bus (see Section 1.6.2.3), new technology solutions can be incorporated by implementing an interaction pattern without impacting the remaining infrastructure. This approach is the basis for all modern operating systems, and is a core principle of object-oriented programming. It will be extended to a distributed resource and authority environment in the proposed work. The same integrative capability is provided to applications through the Service framework.

The Common Execution Infrastructure provides configuration management and demand-driven provisioning of capability at selected locations (CyberPoPs) within the CI network. The CEI is elastic, having the ability to expand and contract the configuration of computing resources as the need rises and falls. This subsystem extends the Amazon® service model for the "Elastic Computing Cloud®" by incorporating a security framework and a demand-based scheduler to enable dynamic configuration.

Both of these infrastructure elements are further described in Sections 1.6.2. Their uses are defined in Section 2.0, Software Release Schedule.

1.6. Implementation

1.6.1 *System Life Cycle*

Our team will provide an integrated system engineering and architecture team to carry out the system engineering and architecting processes. The key planning, requirements definition, and system-level design, prototyping, integration, and verification activities are carried out with oversight from the System Engineer and System Architect assisted by their direct reports, with operational details contained in the Project Execution Plan. Use of the spiral project life cycle (see Section 2.1) does not alter the function of system engineering, and in fact it becomes the key activity that binds the cyclically-growing system into a coherent whole. The system engineering framework used by the project will be a tailored version of that defined in the System Engineering Handbook Version 3 issued by the International Council on System Engineering (www.incose.org).

1.6.1.1 Requirements Definition.

Under the spiral model, requirements definition constantly evolves throughout the project life cycle. Requirements definition is the dominant activity in each inception phase **(Exhibit 14: Spiral Release Cycle)**, but continues at reduced levels throughout each development spiral. The science user requirements definition process is outlined in Section 1.3. The System Engineer and System Architect are participants in this process, and are responsible for extracting system requirements from the science user requirements and other sources. Initial system requirements are part of the ORION CI Conceptual Architecture, and were developed by the JOI Program Office from an examination of existing observatory projects (LEAD, SIAM, SSDS, IOOS DMAC, and VENUS/NEPTUNE Canada) along with input from the ORION Advisory Structure and the OOI Science Plan. Further requirements definition will involve the stakeholder communities organized through the Project Scientist and E&O Manager, the entire CI development team, and the ORION Advisory Structure. Additional elements of the system requirements are the policies that govern use of observatory resources. Their specification will follow a parallel process, but may involve further negotiation with the observatory IOs, analogous to the process used for Interoperability Agreements (see Project Execution Plan, Appendix 6.1, Section 4.4.2), and may require Advisory Structure and Program Office approval.

The System Engineer and System Architect are responsible for ensuring that requirements are stated in an atomic, clear, validatable manner with source traceability. They are also responsible for trade studies, constraint evaluation, and cost-benefit analyses that may be required to refine them. The system requirements will be divided into four major categories (functional requirements, performance requirements, design principles, and interface requirements), and then further sorted into categories that are consistent with the CI subsystems. The system requirements and their provenance will be captured using a requirements data base tool to be selected at project inception.

The outcome of these efforts is the System Requirements Document (SRD), which will serve as the top-level description of the desired CI capabilities. Evaluation of the SRD is a key element of each Initial Operating Capability milestone, which is analogous to the waterfall model Concept Design Review.

1.6.1.2 System Architecture.

System architecture definition is the dominant activity during each elaboration cycle, although it continues at a reduced level in the other spiral phases. With input from the Subsystem Development Teams, the System Architect is responsible for definition and refinement of the system architecture with guidance from the evolving SRD. Architecture specification activities may include prototyping and trade studies as needed, and will be documented using DoDAF. Evaluation of the system architecture is a key purpose of each Life Cycle Architecture review, which is analogous to the waterfall model Preliminary Design Review.

The DoDAF documents replace the System and Subsystem Specification Documents, and guide the CI development throughout the system development life cycle. They are used to partition the system into subsystems, define the subsystems, and task each subsystem lead. The System Development Manager supervises the Subsystem Development Teams with carrying out design and construction of the subsystems. The System Architect provides oversight and coordination during the construction phase of each development spiral. To ensure stakeholder commitment to the outcome, selected ocean scientists will be asked to participate as Subsystem Development Team members.

1.6.1.3 System Integration and Verification.

During the construction phase of each development spiral, the newly created architectural elements are integrated into the existing CI release using the roadmap defined in the Integration and Verification Plan. The System Development Manager is responsible for delivery of a quality, integrated system to the System Engineer. The System Architect is responsible for verification of the integrated CI, while the System Engineer is responsible for integration of each CI release with the marine and external observatory elements and verification of the result. Each CI release will be evaluated for correctness, completeness, security, and quality using the criteria established in the Integration and Verification, Security, and Quality Plans as well as for consistency with the SRD.

1.6.1.4 System Validation.

Near the end of the construction phase, the newly integrated and verified CI will be validated using the roadmap given in the Validation Plan to show that the result meets stakeholder needs. This activity is the responsibility of the Project Scientist assisted by the System Engineer and selected domain stakeholders. The integrated, verified, and validated system is reviewed at an Initial Operating Capabilities milestone (analogous to the waterfall model Critical or Final Design Review) at the end of each construction phase.

1.6.1.5 System Deployment and Acceptance.

During each transition phase, the integrated, verified, and validated system is delivered to the Operations Team for deployment. As part of this activity, the system will be accepted by the ORION Program Office using specified criteria. The Integration and Verification Report and the Validation Reports provide key input to this process. After deployment, the CI is available for operations.

1.6.1.6 Milestone Reviews.

The LCA and IOC milestone reviews will be organized by the Project Manager as internal project reviews, with selected participants from the marine IOs and the Program Office as required. Except for the first instance, the LCO milestones coincide with completion of the deployment of a CI release, and will be treated as a comprehensive review of the previous development spiral as well as a forward looking review of the next development spiral. To ensure ongoing stakeholder commitment to the CI development process, ten ocean scientists and educators will be asked to attend this review and provide independent evaluation of the CI project. Funds for this purpose have been budgeted. The Project Manager is responsible for documenting each of the reviews, and all review reports will be publicly available.

1.6.2 *Software Development*

1.6.2.1 Introduction and Overview.

This section maps the conceptual architecture laid out in Section 1.5 onto an implementation strategy and deployment architecture. The task and challenge for the Software Development Team (SDT) is translating this vision into an operational system that can be deployed to deliver the required capabilities throughout the ORION life cycle. The SDT will leverage and further develop existing technologies to ensure both the rapid availability of an initial CI implementation and the systematic, iterative, and incremental implementation of the full set of CI capabilities in accordance with a release schedule. At its core, the ORION software development effort is a system-of-systems integration challenge, with software development primarily focused on subsystem adaptation and integration of best-of-breed, proven technologies.

The conceptual architecture rests on a rigorous service-oriented *design* approach to project subsystem capabilities and integration into system capabilities. The SDT will leverage service-oriented *implementation* techniques to yield a seamless software and system engineering project framework. Intuitively, every ORION entity (examples range from instruments; to laboratories; to data repositories; to coastal, regional, and global observatories; to the computational Grid; to observatory management) will represent itself as a set of services that can be located and accessed via the CI. Web services and related technologies enable rapid implementation, provisioning, and integration, along with flexible configuration of these services to yield the CI.

This section is structured as follows: first, for illustrative purposes, a concrete ORION deployment architecture scenario is described. This scenario captures the data acquisition activity for an ORION global buoy installation. Second, the core software abstractions and their implementation patterns are defined; this discussion includes an exposition of how the models for resources, services, communications, and presentation are provided as patterns for the overall integration architecture. Third, the architectural elements are mapped to the software development projects that implement them. In Section 2.0, Software Release Schedule, the deliverables for each project are defined and placed in a software development roadmap.

1.6.2.2 ORION CI Deployment Architecture Scenario: Data Acquisition.

Section 1.5 provides a high-level conceptual view of the ORION CI, its subsystems and major domain models and processes. The CI is broken down into six resource networks (see Section 1.5.2.2) that contain different classes of resources and the capabilities needed to manage them. These resources are tied together by two cross-cutting infrastructure elements:

1. The Common Operating Infrastructure (COI) serves as the integration platform and communication conduit, and supports the activity, resource, service, identity, communication, and presentation models that must be used across ORION.
2. The Common Execution Infrastructure (CEI) is the computational substrate that supports provisioning, configuration and direction of collection and analysis protocols.

There are two steps that illustrate the deployment approach planned for the ORION CI. First, we describe the *subset* of the networks and infrastructure elements that support the *data collection* activity of a *global scale, buoyed observatory* as one relevant deployment scenario. This spans the entire range of deployed systems and networks from ocean-based instruments to CyberPoPs and user applications. Second, we establish the mapping of cross-cutting COI infrastructure capabilities to deployment sites in this scenario.

In **Exhibit 15: Global Obsevatory Deployment Model**, the Instrument Point represents the interface between the physical world and the CI; it comprises proxies that provide a programming interface to the instruments. The Acquisition Point provides instrument control and data acquisition and transmission functions. It comprises a process and instrument controller and a data acquisition subsystem. Researcher-supplied triggers initiate data acquisition processes that the process controller translates into commands for the instrument controller. Data from the instruments transit from the instrument controller to the acquisition component, where researcher-supplied filters result in either new data acquisition or transmission of the data to the Ingest Point. Data are sectioned appropriately by a segmentation component for transmission and handed over to the transport broker, which uses the Object Ring Buffer (ORB) and the communications controller to locally store and transmit the data, respectively, based on network availability and Quality-of-Service (QoS) constraints. At the Ingest Point, data arrive via the local communications controller and transport broker. The latter feeds data correction and ingestion components. The ingested data, along with their metadata, are buffered via the local storage broker. The storage broker interacts with the Storage Point that offers repositories for data and services, as well as data and metadata cataloging. The researcher has multiple ways to view and further process data. First, an Access Portal supports data access and presentation via web- and other browsers. It interfaces with the Storage Point via a local storage broker and offers components for search, navigation, and presentation. Second, an Application Integration Point supports data access, transformation and analysis via a set of analysis tools. It also connects to the Storage Point via a local storage broker, and offers programming interfaces so that researcher-supplied analysis and transformation processes can access and manipulate the data.

In this deployment scenario, the Instrument Network comprising the Instrument Point, the process and instrument controller, acquisition and segmentation components, transport broker, ORB and communications controller of the Acquisition Point, as well as the communications controller and transport broker of the Ingest Point are built from BRTT Antelope® components. The Data Network comprising the storage provider with its repositories and catalogs at the Storage Point and the federated storage brokers of the Ingest Point, Access Portal, and Application Integration Point are built from UCSD Storage Resource Broker (SRB) components. Presently, v1.0 of the ROADNet PoP (RPoP) integrates Antelope and SRB in a small, low-power, low-cost LINUX box using Intel XScale Processors. The next operational release will include web services support. The Processing Network is implemented using researcher-provided filter and trigger processes in the Acquisition Point, a data correction process in the Ingest Point, a presentation process in the Access Portal, and the transformation and analysis processes in the Application Integration Point, where the visualization and modeling capabilities of, for example MatLab, are also provisioned. The data management functionality of the Data Network and its interface to the Knowledge Network, comprising ingest and metadata cataloger components at the Ingest Point, the metadata-based search and navigate components of the Access Portal, and the navigate component of the Application Integration point, are implemented using components of the MBARI Shore Side Data System (SSDS). The repositories and catalogs at the Storage Point are implemented using ORION-specific adaptations of SRB repositories and catalogs currently deployed for BIRN/Telescience, with necessary extensions to house service repositories.

As elaborated in Section 1.6.2.3 below, each deployment site implements CI capabilities organized as CI *capability containers* (Cap Containers), see **Exhibit 16: Cyber Capability Container**. All Cap Containers provide a rich set of infrastructure services, including communication, governance and policy, process execution, state management, and repository services. The pervasive provisioning of these infrastructure services across the deployment architecture constitutes the core of the Common Operating Infrastructure (COI). A Cap Container provides access to the resource networks via service interfaces. Each Cap Container also has a presentation capability to project its services to its environment in various forms. **Exhibit 17: Global Observatory Capability Model** illustrates the layout of these containers for the data collection activity scenario. Because of its resource-constrained nature, the Instrument Point Cap Container provides only general CI infrastructure and instrument proxy services. The Acquisition Point Cap Container illustrates the use of infrastructure services to implement the processes introduced in **Exhibit 15: Global Observatory Deployment Model**: the process execution infrastructure service provides the filtering and triggering processes at this Cap Container. An additional function supported by the Cap Container is the presentation capability implemented in the Access Portal Cap Container. Building on the BIRN/Telescience GridSphere-based portal framework, portlets for session management and Google Earth presentation, as well as an http container, can be built.

1.6.2.3 Core Software Abstractions, Infrastructure Technologies and Deployment Patterns.

Because of their pervasive nature, the Cap Containers are ideally suited to addressing cross-cutting infrastructure concerns, including security, reliability, governance, and scalability. In this section, we introduce the two fundamental technology and deployment patterns that will be used in the proposed work to implement the Cap Container concept: web services and the Enterprise Service Bus (ESB). Furthermore, we will show how the ORION collaboration and policy framework can easily be deployed using Cap Containers.

Web services provide a lightweight, well-established deployment pattern for capabilities in distributed systems. They are particularly attractive as a deployment technology because 1) they support a seamless transition from the service-oriented conceptual architecture to a service-oriented deployment architecture, and 2) data and computing Grid technologies come equipped with web services interfaces, as exemplified by the ATOMIC interface of BIRN/Telescience to rationalize access to the national Grid Computing infrastructure. This results in low integration overhead.

ESB technologies are rapidly emerging as the standard approach to system-of-systems integration. ESBs provide a highly scalable integration platform built on web and other open standards that combine access to heterogeneous data sources with messaging and web-service technologies to produce coherent, event-driven, service-oriented architectures that can rest on practically any transport and network access protocol. In essence, an ESB consists of four major components: a service/data interface, a messaging component, a router/interceptor component, and a set of infrastructure plug-ins. The service/data interface acts as a gateway and adapter, connecting the ESB with other services and heterogeneous data and message sources, but also offering access to the services provided by components connected directly to the ESB. The service/data interface also serves as a gateway to other ESBs; enabling a wide range of integration topologies. The messaging component provides reliable messaging via freely configurable point-to-point and publish-subscribe communication channels. The router/interceptor component captures messages intended for the messaging component and subjects them to a suite of infrastructure plug-ins according to a freely configurable routing scheme. Infrastructure plug-ins process messages they receive via the router/interceptor. Examples of plug-ins are data transformers, encryption engines, authentication, policy enactment, and failure management components. This combination of router/interceptor mechanism and infrastructure plug-ins is known as a *dependency injection* mechanism or *aspect-oriented* infrastructure. The message resulting from the routing/processing combination is then made available for consumption via the appropriate channel of the messaging component.

The proposed work will exploit these ESB characteristics to implement capability blocks, as shown in **Exhibit 18: Fractal ESB Design Pattern**. A capability block can be either simple or composite. An example of a simple capability block is an instrument proxy, as shown by the Interface Point in **Exhibit 15: Global Observatory Deployment Model;** it displays no further decomposition, but rather exposes its capability as a web service. A composite capability block is comprises the ESB parts mentioned above, bundled with a specific set of plug-ins and additional hierarchically-composed capability blocks that are connected, via their service/data interfaces, to the messaging component of the composite capability block. An example of a composite capability block is the overall ORION CI. Most importantly, the scale-independent logical structure supports a very broad suite of capabilities and will allow users well into the future to compose capabilities that are not now obvious. This abstraction is critical to future success and acceptance of the ORION CI.

At each deployment site, a Cap Container implements the capability block pattern as described above. In particular, the infrastructure services will be provided as ESB-plug-ins and the resource networks will be integrated via web services, or as capability blocks in their own right.

ORION *activities* emerge from the interplay of multiple capability blocks, as shown by the data collection activity example in Section 1.6.2.2. Activities will be enabled by the precise specification of the collaboration pattern, including the required subsystems, their interaction protocols, and a description of the information exchanges over time with each observatory activity, through an *interaction interface*. Furthermore, the proposed work will associate cross-cutting authentication, security, governance, and policy requirements with each interaction interface. The interaction interfaces will be provisioned via the COI so they can be bound to actual resources either at the time of deployment or at runtime to provide the required degree of flexibility in system configuration. In effect, the ORION activity model is mapped to a service-oriented process model that is supported by appropriate configuration of the orchestration plug-in of each Cap Container.

Exhibit 19: Collaboration and Policy Framework illustrates this pattern for the base case of a single service provider and consumer. The pattern generalizes to arbitrary numbers of participants in a service orchestration. Conceptually, the example captures the establishment of a service agreement between two parties; for example, this could unfold between a regional cabled observatory (service provider) and a buoy-based global observatory (service consumer). Each one of the parties has established contractual commitments with their respective user communities, including membership agreements. Upon establishing mutual commitments, a contract between the two parties is in place. Furthermore, each party operates under a set of policies. The negotiation and contracting process, as well as the actual service usage, leads to an interaction pattern between the two parties that is constrained by the contractual commitments and policy subscriptions of both parties.

Because each Cap Container is equipped with plug-ins for orchestration, governance, policy enforcement, and monitoring/audit, the deployment mapping for the collaboration and policy framework is straightforward: the corresponding interaction interface is stored and accessed COI-wide at the Storage Point (see Section 1.6.2.2). Each party's Cap Container orchestration component executes the projection of the interaction pattern on the respective role to participate in the overall collaboration. The governance and policy constraints are extracted from the interaction interface and provided to the corresponding Cap Container plug-ins for monitoring and enforcement.

The Data Acquisition example provided above shows how the COI facilitates the interoperability and on demand mobility of the capabilities of the Instrument, Processing, Data, and Knowledge resource networks. The core software abstraction illustrates how the COI, through the use of the CI capability container, factors out the common aspects of communication, state management, execution, governance, and service presentation to provide a very scalable, secure and extensible model for managing user-defined collections of information and taskable resources. This ability to integrated resources of different types implemented by different technologies is the central proposition of the proposed architectural approach. It provides the basis for an integrated observatory network that will remain viable and pertinent over multiple decades.

To appreciate the versatility of this model and the strength of the proposed implementation, it is important to understand the specific capabilities, data models, interface points, technologies and systems behind each of the six resource networks and the two infrastructure elements. Given the rather detailed and technical nature of such a discourse, a set of exhibits discussing each is provided in **Exhibits 20.1-20.8: Architectural Elements**.

1.6.2.4 Implementation of the Capabilities.

In this section, we turn our attention to the implementation of the architecture in the context of an integrated set of subsystem development projects. The development effort is structured into seven projects based on the major activities of scientific investigation as well as data and knowledge management and the two infrastructure elements. The projects are: 1) Sensing & Acquisition, 2) Analysis & Synthesis, 3) Planning & Prosecution, 4) Data Management, 5) Knowledge Management, 6) Common Operating Infrastructure, and 7) Common Execution Environment.

Each project has a duration of either two or three 16-month development cycles. They are staggered across the five release cycles of the overall proposed project. Their ordering is based on a prioritization of their value to the ORION community and their interdependency. The projects deliver complete subsystems and have a prescribed set of deliverables. These deliverables, starting with the domain models and ending with implementation code, are essential for the long-term viability of the system. **Exhibit 21: Engineering Life Cycle** provides the correspondence between the deliverables (on the right) and the community of participants (on the left) engaged in shaping the outcome of these projects through the engineering process (in the center).

Close attention has been paid to the shared architectural elements that can be assembled by one team, enhanced by another and used by most. This strategy of "implement, enhance, use" across teams is employed to ensure sustainability of the interaction interfaces across the duration of ORION and drive the quality of the implementations behind them. **Exhibit 22: Development of Architectural Elements** lists the projects and shows their implementation responsibilities and dependencies based on the architecture introduced in Section 1.5 and detailed in this section.

In addition to the teams identified for each project, architectural leadership, domain modeling, software design, and user interface resources have been allocated from the System Architect's engineering team to work alongside each project. This will ensure coherence of the design and products generated across the projects while giving the proposed work the ability to engage the expertise of the Ocean Sciences and CI communities in a very targeted and cost effective manner to the ORION vision.

A comprehensive exposition of the full set of deliverables in terms of end user services, resources, and infrastructure for each of the projects appears in Section 2.2. The following is a summary review of the projects, detailing the team, objectives, and major technologies that will be used to realize the deliverables.

The Sensing & Acquisition project involves software engineers from two of the top oceanographic institutions, Scripps Institution of Oceanography and Monterey Bay Aquarium Research Institute, deploying real-time cabled and moored instrument management and data acquisition systems. The team is lead by ROADNet's lead architect and designer Dr. Kent Lindquist. The project will utilize and further develop the Antelope Object Ring Buffer (ORB) server and MBARI's PUCK standard. The ORB is a robust, content-neutral, packetized buffering and event-driven distribution system that can be configured for wide area (global) data acquisition, sharing, and processing. Agilent Technologies, Inc. and the IEEE 1451 standards group are interested in participating with ORION in building on the PUCK and IEEE 1451 standards to establish a network standard for instrument control and measurement processing.

The Analysis & Synthesis project merges the observing and modeling communities. Dr. Yi Chao, PI for JPL's OurOcean Portal project, will lead a team that incorporates his work with the distributed workflow execution of Pegasus from the USC Information Science Institute and the Harvard Ocean Prediction System (HOPS), which recently relocated to MIT. The project will use, adapt, and further develop community-based numerical ocean models such as the Regional Ocean Modeling System (ROMS) and HOPS, combined with a suite of integrated applications, including a standard Web portal interface and Matlab, Kepler and WS-BPEL workflow editors that will support process and model specification, simulation, analysis, and visualization.

The Planning & Prosecution project leverages the consistent nested and autonomous capabilities of the integrated network of sensing, modeling and control resources. Dr. Henrik Schmidt, PI of MIT's Laboratory for Autonomous Marine Sensing Systems, leads this JPL and MIT team of engineers to integrate Dr. Schmidt's work on PLUSNet and Dr. Steve Chien of JPL's work on autonomous Earth observing sensor webs to develop a generalized design and control framework for ORION. The objective is to plan, schedule, and prosecute multi-objective observational programs. The project will use, and further develop, the behavior-based autonomous control software MOOS-IvP for fully autonomous event capture and characterization. The Mission Oriented Operating Suite (MOOS) is open source middleware for connecting software components on an autonomous platform. MOOS-IvP extends MOOS via Interval programming (IvP), a unique, new mathematical model for representing and solving multi-objective optimization problems for reconciling vehicle behaviors during missions.

The Data Management project capitalizes on two significant opportunities: the next generation of the most influential data Grid technology, SDSC's Storage Resource Broker (SRB), and the open source release of MBARI's Shore Side Data System (SSDS). Led by Dr. Arcot Rajasekar, designer of SRB and iRODS, and Dr. Reagan Moore, inventor of the SRB, the team of MBARI, MMI and UCSD will establish the metadata and semantic foundations of the ORION data repository, a foundation that is required to operate for decades. The SDSC SRB system provides all of the features that are needed to build a production-level data Grid, including facilities for collection building, managing, querying, accessing, and preserving data in a federated, distributed framework with uniform access to diverse, heterogeneous storage resources across administrative domains. The SRB's Metadata Catalog, MCAT, holds system and application- or domain-dependent metadata about the resources and datasets, and methods and users that are being brokered by SRB. Together, the SRB and the MCAT provide a scalable information discovery and data access system for publishing and computing with scientific data and metadata.

The Knowledge Management project brings together a design team from the leading institutions working on semantic-based data integration in environmental sciences. Under the leadership of SDSC's Dr. Amarnath Gupta, the semantic architect for BIRN, the design team from MBARI, NCSA and SDSC will collaborate with the Marine Metadata Interoperability project to develop methodologies and tools for disseminating the use of vocabularies and ontologies in the pursuit of ocean research. The project will utilize and further develop SDSC's Data Mediation technology that is currently in use in the BIRN/GEON projects. The tool chains for vocabulary authoring and negotiation being developed jointly between MBARI and MMI will play a large role in the adaptation of a semantic approach to data interoperability. This work will enable mediation of science domain-specific queries across the broad heterogeneous set of oceanographic data sources.

The Common Operating Infrastructure project brings together the agile integration model of SOA with maturing distributed Grid security to provide a secure and scalable federated operating infrastructure. Dr. Ingolf Krueger, the PI of the Calit2 Software and Systems Architecture and Integration Team, will lead a group with considerable experience in the security- and policy-driven governance of Grid and Web infrastructure from NCSA, U of Chicago (Globus) and North Carolina State University. The project will engage a number of significant technologies. At its core, it will configure and utilize the Enterprise Service Bus suite of technologies, comprising of scalable, reliable, and secure messaging technology for data streams, a service and data presentation and subscription mechanism, a routing & filtering framework for flexible plug-in of governance and policy management components, dependency injection, and a service framework. Furthermore, it will utilize the time-tested BIRN/Telescience cyberinfrastructure stack with its ATOMIC set of web-service-enabled interfaces to data and computational Grids, as well as user authentication, authorization, and security. This configured technology suite yields the communication and integration substrate for the proposed work.

The Common Execution Environment project lays the foundation for scalable, immediate mode, reconfigurable, elastic computing with a process integration strategy that will allow ORION to adapt to its participants manner of processing data anywhere in the network. The project draws on the extensive experience in provisioning, scheduling, and management of Grid and data workflow infrastructure. Dr. Kate Keahey from the University of Chicago and Ilkay Altintas of SDSC lead a team currently developing distributed analysis environments for science. The project will incorporate the scheduling advances made by NCSA, USC-ISI and SDSC with the Enterprise Service Bus suite of technologies to provide process execution environments for scripts and binary programs. This will use Globus's Virtual Workspace with SDSC's Rocks optimized configuration deployment technology to establish an on-demand computing capability that can operate across the integrated observatory network.

The proposed work will use the following tool for design, development and management of all system artifacts:

- Telelogic's System Architect for authoring and management of all design artifacts,
- Eclipse, the industry standard open development platform, for software development,
- Perforce for software configuration management and version control,
- Atlassian's JIRA for system bug and issue tracking, and
- Apache's Ant for the software build environment.

At least three copies of all design and development artifacts will be maintained at all times with an additional offline storage strategy. The system artifacts will be stored in the same manner as all of the scientific data to ensure long-term preservation.

1.6.3 Hardware Development

The proposed work does not anticipate the need for any specific hardware development. The use of Kinemetrics Marmot marine-deployable computing platform with embedded BRTT Antelope software obviates the development of any bespoke hardware solutions.

The proposal team anticipates working directly with the Observatory IOs in their development of the Science Instrument Interface Module (SIIM) to identify the most effective microprocessor that meets power and communication requirements while affording the CI a viable compute platform directly associated with each instrument.

All expenses associated with hardware purchases are either accounted for as Software Development or Implementation expenses.

1.6.4 Implementation Plan

In Sections 1.5 and 1.6.2, the proposed CI architecture was described and its constituent software elements were defined. Section 1.6.1 outlines the system engineering process that will drive the CI from requirements definition through integration and deployment in the marine observatories. This section defines the compute and storage systems that will be provided throughout the CI and the additional, external resources and services that will be integrated to produce an ocean observatory system-of-systems. It also contains a timeline for their deployment.

The integration of the CI with marine observatories is constrained by the installation schedule for their constituent wet and dry elements. The ORION Conceptual Network Design and associated planning documents project that coastal, global, and regional cabled observatory elements will respectively be ready for service 12-24, 18-30, and 54-66 months after start of the OOI. The CI development schedule has been broadly defined around these dates. As described in Section 2.1, the CI deployment will be phased, with major software releases beginning at 18 months after initiation of the OOI. Successive releases will provide new or improved capabilities while maintaining backward compatibility with earlier releases, hence minimizing the impact on the user communities.

1.6.4.1 Provisioning.

The primary compute and storage building block for the CI is the Cyberinfrastructure Point of Presence (CyberPoP) introduced in Section 1.5.1 (see Exhibit 10: Cyberinfrastructure Points of Presence). From the perspective of users, observatory operators, and the CI, CyberPoPs are virtual resources that can be provided in ways ranging from actual hardware installations to an elastic service analogous to Amazon®'s Elastic Compute Cloud®.

There are four classes of CyberPoP provisioned by the proposed work; the on-demand highly scalable (*i.e.*, elastic) CI CyberPoP; the high availability, shore side, Observatory CyberPoP; the ocean deployable Marine CyberPoP; and the high bandwidth stream processor (HBSP) CyberPoP.

The highly scalable CI CyberPoPs will be accessible at two central facilities (one at SDSC based on Teragrid, and one whose location is yet to be determined). These two installations will serve as the central repository and core processing centers for ORION. They will also support the CI Test Facility. The CI CyberPoPs and the CI System Test Facility will be elastic installations, with compute cycles and storage provided as a per unit service. In this way, costs will be proportional to actual rather than anticipated use. Capacity, availability (up time), physical operations and maintenance will all be ensured by the service provider, as governed by a service agreement.

This model also provides for the migration from service provider to service provider over the duration of the ORION program. The only binding asset of a physical computing and storage facility is access to the tapes associated with off-line storage. The proposed model for storage is to maintain tapes for only an initial period until NASA's Earth Observing System Data and Information System (EOSDIS) strategy of moving all data online has been verified. The proposal has accounted for three copies of all data being on-line at geographically distributed repositories, with an additional copy on tape. We intend that data persistence be provided at these locations, including the maintenance of metadata catalogs. Maintaining full data provenance implies that modifications to data sets are controlled through a versioning process such as that implemented on SRB and iRODS. The O&M budget provides funding for data persistence through the next 11 years, and anticipates that this will be continued as the ORION project continues into the future to the projected 30-year lifetime of the observatories. The commoditization of computing and storage will continue to expand throughout this period, and exponential growth of capacity will reduce overall costs.

The CI CyberPoPs also support the CI Test Facility. It has three major functions: 1) as a test platform to carry software releases from alpha, to beta, to production status, 2) as a test harness for software release conformance and performance characterization, and most importantly 3) as a test facility to verify and qualify all services and resources introduced into the integrated network (*i.e.*, instruments, processes, models and missions).

The shore side Observatory CyberPoPs will be located at each Observatory IO operations center, and will support all of the capabilities of the CI CyberPoP. The Observatory CyberPoPs will comprise installed hardware scaled to provide storage for six months of anticipated data production and the compute capacity to manage its associated processing, product development, and presentation. The core infrastructure will be based on the high availability Sun Microsystems server architecture. This selection will allow for phased upgrades of hardware from a proven provider that has demonstrated clear direction over many years. We will select the hardware platforms from products that are not nearing the end of the manufacturer's production cycle. The Observatory CyberPoPs will operate as independent entities working in concert with the other Observatory and CI CyberPoPs to form the nucleus of an integrated observatory network. The operational independence ensures that the failure of one component won't impact the overall reliability and continuous operations of the others.

In addition to an Observatory CyberPoP, the Regional Cabled Observatory will be provided with the High Bandwidth Stream Processor (HBSP) CyberPoP. This recognizes the extreme capabilities of a 10Gbps cabled system to produce unprecedented volumes of observed data on a continuous basis. At the current time and for the first five years of the ORION program at least, it is more cost effective to deploy enough computing capability onsite to characterize, classify, and buffer the data as it arrives than to provision enough bandwidth to backhaul it to one of the CI facilities. The proposed work will provide enough storage capacity to support short term (order of days) buffering. Networking and storage costs at the current time preclude permanent archival storage of all high bandwidth data streams (*i.e.*, HD Video) by ORION, and responsibility for storage will devolve to the user until costs diminish or additional storage is made available. To illustrate, the current cost to store a single lossless compressed HD Video stream at 30 fps for a year online is $1.7M; to transport the same will cost ~$500k/yr for a 10Gpbs circuit. The system being used as a model for the HBSP CyberPoP (8 node, 64 processor Sun cluster w/ 11.5TB SAN storage system) is less than half the annual circuit cost. As exponential growth in network speed and data storage capacity occur, this model will change.

Each of the CI facilities and Observatory IO operations centers will initially be provisioned with Internet2 connectivity. It is anticipated this will be more than sufficient for the Coastal and Global observatory operations. When the RCO comes online in Year 5 of the program, the national connectivity will be augmented, opportunistically, by using a number of under subscribed academically-owned 10 Gb/s lambda circuits. It is anticipated that by the end of the OOI program, the RCO will be producing enough data to warrant provisioning a national 1Gbps circuit connecting all of the ORION facilities and most of the major research institutions. This can be accomplished and is currently planned, by purchasing such capacity from either National Lambda Rail or Internet2.

The Operations Manager, assisted by the Operations Team, is responsible for delivery of a tested and verified instance of each CyberPoP to the System Engineer. The System Engineer, assisted by the System Integration Team and the service provider, is responsible for installation and verification of each CyberPoP. Observatory installations will involve interfacing the shore side CyberPOP to the observatory's power, data communications and time service infrastructure in accordance with previously negotiated Interface Agreements (see Section 1.6.1). Following installation, the Operations Manager is responsible for ongoing operation and maintenance of the CyberPoPs, which will be managed remotely from UCSD ("lights out management").

The CI Test Facility will be provisioned and operating by 12 months after initiation of the project. The SDSC and coastal CyberPoPs will be operational at the time of the first CI software release (18 months). The global CyberPoP will be on-line a year later. The regional cabled observatory CyberPoP will be provisioned around 54 months, or just prior to the RCO coming online. The HBSP CyberPoP is scheduled to be installed in the last year of the program to ensure the best purchasing power for the allocation and to be coincident with the anticipated ramp-up of activity at the observatory. The second central CyberPoP will be completed toward the end of the CI development phase.

The wet hardware end of the CI contains numerous instances of an embedded CyberPoP located in instruments, in SIIMs, and possibly in observatory physical nodes. As a result of differences in connectivity continuity, available bandwidth, and the clear need for *in situ* buffering at buoys, these implementations will be distinct for different marine observatories. Because the embedded CyberPoPs are simultaneously an integral element of a given marine observatory and of the CI, their design must be carried out jointly by the marine and cyberinfrastructure IOs through a process that remains to be defined. It is anticipated that the marine observatory IO will be responsible for hardware implementation and maintenance, while the CIIO will be responsible for software development and maintenance. Clarification of the embedded CyberPoP engineering process is a priority issue that must be negotiated between the ORION Program Office, the marine IOs, and the CIIO early in the OOI life cycle. The first embedded CyberPoP must be ready for installation on a coastal observatory at 18 months after OOI initiation, and different versions will be required as subsequent observatory types come on-line. The standard Marine CyberPoP provided by the CI leverages the proven technology of Kinemetrics Marmot® data acquisition appliance built on Intel® XScale microprocessor technology with a Linux OS. A Marine CyberPoP will have access to the full capability of the other types of CyberPoP with obvious constraints.

1.6.4.2 Integration.

The process for internal integration of the CI elements and final integration of the CI with the marine observatories has been described in Section 1.6.1. In this section, additional resource, service, model, compute and external observatory elements that will be integrated with the CI are briefly described, along with timelines for their completion.

The ORION Conceptual Network Design documents define about 30 instrument types that will be deployed with the initial OOI infrastructure. A selected set of two or three instrument types will be chosen for pilot integration with the CI, both to characterize the range of integration issues and to set the pattern for future integration efforts, hence simplifying and standardizing the process. The Project Scientist will be responsible for selecting the two to three initial instrument types, in consultation with the ORION science user community and CI System Engineer. Instrument integration will be carried out by the Sensing and Acquisition project team. Instrument integration for the 30 basic types will be completed by 30 months after initiation of the OOI, with the schedule determined through community input collected by the Project Scientist.

Section 1.6.2 describes the major technologies and technology families that will comprise the internal elements of the proposed CI. In addition, there are three or four key processing engine applications that are either widely used by the ocean sciences community (*e.g.*, Matlab), or in increasing use by the sensor Grid communities (*e.g.*, Kepler, BPEL, and Riverglass), that will be integrated with the ORION CI to increase user functionality. **Matlab** is an interactive, GUI-based data visualization and analysis application with a diverse set of toolboxes. **Kepler** is an open-source scientific workflow system that allows scientists to design scientific workflows and execute them efficiently using emerging Grid-based approaches to distributed computation. **Business Process Execution Language** (BPEL) orchestrates peer-to-peer interactions of processes, and is widely used in business web applications. **Riverglass** is a powerful emerging application that combines real-time streaming data with analytic and visualization processes. Seamless interfaces to Matlab and Kepler will be available at 30 months, while BPEL and possibly Riverglass (depending on demand) will be integrated by 42 months.

The CI modeling framework will rest on two community-based numerical ocean models, the **Regional Ocean Modeling System** (ROMS) and the **Harvard Ocean Prediction System** (HOPS). ROMS is a flexible and extensible framework for a free-surface, terrain-following, primitive equation ocean model that is widely used by the scientific community for a diverse range of applications. HOPS is a complementary modeling package with significant nowcast and forecast capabilities. Both models will be integrated with the Analysis and Synthesis design element, and will be available at 42 months.

Connection of the ORION CI to the national computing Grid infrastructure is essential for high capability modeling and data assimilation. **Teragrid** is the most extensive existing structure for distributed high performance computing, and will be linked to the ORION CI by 30 months after project inception. The **Open Science Grid** is a second viable candidate for integration, depending on its evolution.

The ORION CI will also be integrated with two selected classes of external services to enhance its functionality: data standardization and mediation services, and geo-referenced access and visualization services.

Three widely used data standardization and mediation service frameworks will support the creation, access, and sharing of scientific data within the ORION CI: **Network Common Data Form** (NetCDF), **Open-source Project for a Network Data Access Protocol** (OPeNDAP) and **Thematic Realtime Environmental Distributed Data Services** (THREDDS). NetCDF is an interface for array-oriented data access and a library that provides an implementation of the interface. The NetCDF library also defines a machine-independent format for representing scientific data. OPeNDAP makes local data accessible to remote locations regardless of local storage format, and provides tools for transforming existing applications into OPeNDAP clients. THREDDS provides middleware to link data providers and users, and simplifies the discovery and use of scientific data. These frameworks will be part of the initial CI software release at 18 months.

Two flexible geo-referenced access and visualization services are the **Live Access Server** (LAS) and the **Open Geospatial Consortium** (OGC). LAS is a highly configurable web server that provides flexible access to geo-referenced scientific data. It can present distributed data sets as a unified virtual database, and it is integrated with a wide range of visualization packages including Matlab. OGC provides standards for open source, web services-based tools for geo-referenced data access and visualization. The specifications currently include the **Web Map Service** (WMS), **Web Feature Service** (WFS), and **Web Coverage Service** (WCS). WMS provides three operations to support the creation and display of registered and superimposed map-like views of information that come simultaneously from multiple remote and heterogeneous sources. WFS allows a client to retrieve and update geospatial data from multiple WFSs. WCS extends the WMS interface to access raster data. These two services will be integrated to the CI at 30 months.

Metadata standardization and cataloging represents an important element of the ORION Data Management subsystem that will require ongoing attention throughout the development lifecycle. Four organizations have been identified to provide standardization liaison or alternatively, the ORION CI will provide interoperability with the standards and catalogs. The **Marine Metadata Interoperability** (MMI) Project is the most relevant standards-coordinating organization to ORION. The **Federal Geographic Data Committee** (FGDC) promotes the coordinated development, sharing, and dissemination of geospatial data through the **National Spatial Data Infrastructure** (NSDI). NSDI is a physical, organizational, and virtual network that enables the sharing of digital geographic information resources. ORION must be compliant with this *de facto* standard. NASA's **Global Change Master Directory** (GCMD) is a comprehensive catalog of environmental data and ORION interoperability is desirable. Finally, **Ocean Biogeographic Information System** (OBIS) is an emerging database of global marine animal and plant distributions that may become an important interoperability candidate in the future. Liaison with MMI will be maintained on an ongoing basis. Interoperability with FGDC and GCMD will be achieved at 30 months, and possible interoperability with OBIS will be evaluated later in the project.

Finally, the **Integrated and Sustainable Ocean Observing System** (IOOS) remains an important potential development partner for the future. Since IOOS has not been funded, and hence neither their ultimate product nor their development timeline can be fully defined, meaningful interoperability planning cannot be carried out at this time. However, we will maintain contact with IOOS through John Orcutt and Matthew Arrott, who serve on Ocean.US DMAC Committee, and a CIIO development partner (Raytheon). We will make the link with IOOS more explicit as it solidifies and as funding allows. John Orcutt and Yi Chao participated with Raytheon in a NOAA-funded **IOOS Conceptual Design Study** during 2006. The design proposed a service-oriented architecture that would be simple to couple to the proposed architecture using the resource adapter concept described in the Raytheon Conceptual Design.

1.6.5 Operations and Maintenance Phase

In the operation and maintenance phase of the OOI, the CIIO will be receiving sensor data streams from each of the observatory subsystems, ensuring that those data streams are made available through the CI, providing the computational and storage resources to support data processing, analysis and other aspects of the CI, and provide access to these resources for the user community.

The Operations Manager, assisted by the CIIO Operations Team, will be responsible for:

- Data communications to and from sensors and actuators;
- Sensor command and control;

- Presentation of real-time data and information through the ORION portal;
- Presentation of non-real-time data and information through the ORION portal;
- Archiving of data;
- Quality control of data and metadata;
- Maintenance of all metadata and station information;
- Maintenance of the federated metadata catalog;
- Maintenance of the federated resource catalog;
- Regular system calibration;
- Monitoring sensor and actuator state-of-health;
- Presentation of sensor and actuator state-of-health;
- Monitoring the ORION system state-of-health;
- Presentation of the ORION system state-of-health;
- Alarm notification to appropriate operational personnel;
- Delivery of processing services for the ORION infrastructure.

1.6.5.1 Operations and Maintenance.

The approach we propose for operations in the CIIO will be the scalable computing service as defined in the Common Execution Infrastructure with capability for on-demand capacity provisioning and support for immediate-mode execution of processes across a federation of computing service providers. To minimize cost, we will use elastic infrastructure to be deployed at SDSC that will operate a standard CI software stack. These costs are projected as a base level capability and can be scaled by any increase in the computational and network bandwidth requirements of ORION. This strategy will minimize computing and networking hardware life cycle costs, system administration costs, and network administration costs.

Operations and maintenance (O&M) management will be governed by the Operations and Maintenance (O&M) plan developed by project management (see Appendix 6.1, Section 4.10), and will be used for the post-deployment phase extending to the end of the ORION life cycle. It is envisioned that CI O&M will require 9 FTEs. The CI project office will require a Project Manager (0.5 FTE) along with a full time administrator. Observatory services will require 4 total FTEs, including 2 FTEs supporting Data and Knowledge Management and one FTE each for Planning and Prosecution and training and support services. Cyber Operations will need 3.5 FTE to support the engineering operational requirements for marine observatory liaison, software development, and system build and deployment. The CIIO will develop, perform, and document CI tests prior to their acceptance by the Program Office. The CIIO will support and document operational readiness tests that may be required.

Finally, we will implement a state-of-health monitoring system using standard tools such as IPERF, Nagios or RRDtool, or their equivalent, to provide system and networking information as well as notification of system or component failures. The CIIO will provide a 24-7 on-call point-of-contact person who is a member of the CIIO O&M staff. The CIIO will provide training to maintain the technical and process competency of the operating staff. We will also design a set of metrics to determine how well the system is operating. Some examples of the metrics include information on data return rates, data latency, bandwidth usage, and time quality.

A summary of these metrics will be included in the Annual Report. The core of the processing and networking infrastructure will be subcontracted to facilities such as SDSC and NLR that will include 24-7-365 service as part of their contracts. These entities will be required to provide state-of-health monitoring and reporting for their processing systems and network connectivity, respectively.

1.6.5.2 Data Preservation.

The current ORION data policy acknowledges that all data will be rapidly disseminated (near-real-time) in an open (unrestricted except for reasons of national security) and freely available (without charge when accessing via the Internet) environment. The CI architecture was designed with this in mind, and the system described in this proposal meets these requirements. The ORION policy also requires that providers of data establish an archiving and metadata plan that will be approved by the ORION management structure. The provider must associate their data with metadata that meets or exceeds national requirements. The proposed work implements this criterion, although the ORION CI committee has yet to examine any specific plan for consistency with the policy or to establish data and metadata standards for the wide variety of data in the ORION collection. It must be noted that other data centers and archives will not provide the full functionality of the ORION CI for metadata, maintenance of data provenance, discovery, and real-time management. Both the Ocean Sciences Data and Sample Policy and the ORION policy require deposition of data in the National Data Centers. In the case of ORION, these centers include:

- National Oceanographic Data Center (NOAA NODC)
- National Climatic Data Center (NOAA NCDC)
- National Geophysical Data Center (NOAA NGDC)

along with several discipline-specific centers (*e.g.*, UNAVCO, IRIS, and RIDGE).

Formal policies must be established for archival services from the three National Data Centers that will include transfer of data with or without latency, the appropriate ontologies for the specific data models, and potential costs. The integration of the OOI with IOOS, as proposed for a specific instance in this proposal, will provide a mechanism for transferring OOI data to the appropriate National Data Centers. We have considerable experience through EarthScope's ANF, the IDA network, and the Ocean Bottom Seismograph Instrument Pool in transferring data to IRIS. Likewise, through the Scripps Orbit and Permanent Array Center (SOPAC), we transfer geodetic GPS data to/from UNAVCO. Establishing service agreements with these consortia for these specific scientific domains will be straightforward. However, the ORION Advisory Structure as well as the NSF must set specific goals for these services given the cost implications. The maintenance of the unique capabilities of the ORION CI is an argument for long-term NSF funding for an intramural data system, as has been the case with the IRIS DMC.

1.6.5.3 Virtual Ocean Simulator.

In summary, the virtual ocean models for the three observatories will have a size on the order of 12200*10400*100 grid points that need to run on a daily (24-hour) basis. Most of the codes are written in MPI. Some of the data processing and assimilation codes will be initially written in OpenMP and tested on the shared-memory computers, and will be converted to MPI. Thus, we plan to run the production codes using MPI on any distributed-memory computers. A test has been conducted on the SGI Altix computer using the Intel Itanium2 processors (900 MHz, 1.5 MB cache, 4 GFLOPS peak speed, and 1 GB main memory, gigabit network). We assume that the production run for the virtual ocean will start in FY2009. If we assume that the processor speed will double every 18 months (this is a conservative estimate because it roughly doubles every year), the 2009 processor will be faster than the 2006 Intel Itanium 2 processor by a factor of four. To integrate the virtual ocean with a dimension of 12200*10400*100 grid points with 20 variables over a period of 24 hours requires about 20 hours on a 2009 1024-processor cluster computer.

By 2009 at the start of the daily update of the Virtual Ocean, the proposal team will request 20,480 single processor hours per day on the supercomputer facility provided by NSF (*e.g.*, via TeraGrid) or other agencies. An expanded justification is in Exhibit 23: Virtual Ocean Allocation Justification.

1.6.5.4 Mission Scheduling and Coordination.

An important component of the CIIO is the scheduling and coordination of observatory hardware and software resources. It is anticipated that there will be several simultaneous users of individual observatory nodes, yielding the potential for damaging interference. For example, a whale biologist may be tracking whale vocalization, while the operator of an underwater vehicle is attempting to command the vehicle acoustically to respond to an oceanographic event, with obvious interference being the result. The Planning & Prosecution component of the CI will be designed to minimize such interference while maintaining efficiency. The coordination will, in part, be centralized through CIIO scheduling and planning support, but it will also be incorporated in the distributed network control framework and inherently enabled by the behavior-based MOOS-IvP command and control architecture. Thus, the 'behavior algebra' will allow the mission planning to automatically incorporate safeguards in the plan. For the acoustics example, the observation plan for the AUV will automatically be amended to restrict acoustic communication unless a specific 'clear' command is given by the mission tracking the whales. This is achieved by adding a Boolean condition for the communication behavior to the MOOS-IvP mission plan. An additional layer of safety will be provided by the requirement that all individual and contemporary measurement campaigns in the observatory be thoroughly tested by executing the entire mission plan in the Virtual Ocean Simulator before the actual field deployment will be approved.

2.0 SOFTWARE RELEASE SCHEDULE

2.1 Project Life Cycle

As described in the Project Execution Plan (Appendix 6.1, Section 2.2), the CI project will be organized around the spiral project life cycle model that has become the *de facto* standard for complex technology projects, especially those that are software-intensive. The spiral project life cycle differs from the traditional, sequential, waterfall project life cycle that proceeds linearly in discrete phases with limited iterations between them. **Exhibit 24: Comparison of Waterfall & Spiral Management Models** contrasts the waterfall and spiral approaches. In the spiral project life cycle, the system is defined, refined, and developed in multiple, risk-driven spiral iterations comprising of four phases (inception, elaboration, construction, and transition; see **Exhibit 14: Spiral Release Cycle**) bounded by anchor point milestones (Life Cycle Objectives (LCO), Life Cycle Architecture (LCA) and Initial Operating Capability (IOC); see **Exhibit 25: Spiral Management Milestones**) that constitute project reviews. The details and impact of the risks (whether technical, management, operational, or stakeholder) drive the number of spirals, the level of detail, and effort within each phase of a spiral. The riskiest elements are brought forward as early in the development process as possible. Each spiral includes management, engineering, and support activities in proportion to the risks. Each spiral expands system definition and results in a deployed representation of the CI.

The spiral project life cycle and annual budget limitations that preclude concurrent development of all of the subsystems leads to a project schedule consisting of five full development spirals **(Exhibit 6: Project Master Schedule)**. The first spiral has an inception phase that is six months long, but this phase is reduced to four months in subsequent cycles. Each development cycle has a notional four-month elaboration phase, a six-month construction phase, and a two-month transition phase, excepting that the last transition phase, which is eight months and leads to the final deployed CI. Beginning with the second development cycle, the initiation phase starts two months prior to the end of the preceding construction phase, so that a full cycle and CI release is completed every twelve months beginning eighteen months after project inception. A LCO, LCA, and IOC anchor point milestone occurs at the end of each inception, elaboration, and construction phase, respectively.

2.2 Software Capabilities and Schedule

As illustrated by **Exhibit 14: Spiral Release Cycle** and **Exhibit 6: Project Master Schedule**, the software is designed, developed, iterated, and promoted on an annual basis. The proposed work will have an operating system in place within 18 months in order to support the first observatory nodes being deployed. The release cycles illustrated by **Exhibit 6: Project Master Schedule** and the Software Release Plan, **Exhibit 26: Software Release Plan** represent the broad demarcation of the major planning and design milestones. In reality, different capabilities of the system will require different development cycles. All capabilities of the system will be broken down into development cycles much shorter than a year. As practitioners of the Agile Development Process, the proposal team advocates having the development code base operating and available continuously with its migration to the production system on a regular basis. We propose to do this migration on a 6 to 8 week schedule. This tight coupling between the development efforts and the operational system proves to be very effective for bringing users into a closer relationship with the system that is being built to provide the users with increasing functionality. At the same time, it provides the development team much finer granularity of control over the direction and quality of the product being produced.

The software release strategy is to plan, design, and reassess on an annual basis, yet develop, test, and deploy on a bi-monthly basis. **Exhibit 26: Software Release Plan** provides the overview of how the subsystem development projects are scheduled across the five spiral Release Cycles. The next seven sections define the major deliverables of each of the subsystem development projects and when they will be implemented (*i.e.*, which Spiral Release Cycle). The deliverables are described as services; each potentially has multiple representations, but all will have at least one user interface and one application interface, the "Interaction Interface." The last section looks at the deliverables from the release perspective and provides a characterization of each release with a listing of the deliverable that will be implemented.

2.2.1 *Sensing and Acquisition*

Duration: R-1 through R-3

Service Deliverables

Observatory Services: R-1 & R-2

Provides services to task, coordinate, and manage the observatory resources and their interdependencies. The coordination services are the primary means for allocating and scheduling instrument use of communications and power, but will extend to the coordination of environmental interactions (*i.e.* sound, chemical, light). The management services provide oversight to ensure safe and secure operations and to maximize the total data return from all instruments.

Instrument Direct Access Service: R-1 & R-2

This service provides direct IP connectivity between the research team and their instrumentation from anywhere within the integrated network. The service is designed to support instrument connects to telnet and/or proprietary instrument software. Such a channel has a higher-level security requirement, and initiation will require a separate and more stringent authentication process.

Instrument and Process Repository: R-2 & R-3

Maintains informational representations of instruments and their configuration and calibration, along with references to their acquired data. It also maintains copies of all processes applied to data from acquisition through product delivery. All are associated with their respective metadata.

Instrument Services: R-2 & R-3

Provides the command, control, and monitoring services to operate and manage an instrument. Operating an instrument has a higher-level security requirement, and engagement will require a separate and more stringent authentication process. This service also supports instrument development and deployment through test and validation services.

Instrument Integration Services: R-3

Provides testing and validation services to ensure conformity with different operational requirements in the network.

Data Acquisition Services: R-2 & R-3

Provides services to configure data acquisition and apply specific data processing steps at the acquisition site and/or at the ingest site.

Data Calibration & Validation Services: R-2

Enables configuration of the data calibration and validation processes and the application of custom automated data processing steps. The service supports the flagging and sequestering of derived data until reviewed by responsible participants. Derived data are automatically associated with their data source. The service supports automated revisions of the derived data on a partial or complete basis.

Data Product Development Service: R-3

Provides services to produce and publish data products and apply processes for generating products from data and/or derived data. Data products are automatically persisted and published based on the configuration set for individual product development.

Resource Network Deliverable

Instrument Network: R-3

Provides interactive and coordinated interaction with real and/or synthetic environments through the use of transducers (sensors or actuators). Provides the command and control semantics for interacting with an Instrument resource.

2.2.2 *Analysis & Synthesis*

Duration: R-3 through R-5

Service Deliverables

Laboratory Services: R-3 & R-4

Provides services to organize, manage, and control research activities, the resources they use, and the participants involved. It is the virtual home where research teams gather their resources, carry out their objectives, and collect their results. It belongs to an individual or a group. It provides the group management tools to facilitate membership and collaborations and to assign roles and responsibilities.

Analysis & Model Repository: R-4

Maintains analysis processes with their associated metadata. A representative suite of analysis and visualization processes that adhere to the ORION measurement semantics will be introduced.

Data Analysis & Visualization Service: R-3

Provides a generalized analysis and synthesis framework for transforming, analyzing, and visualizing data through the application of user- and community- developed processes.

Model Repository: R-4

Maintains a hierarchy of evolving interdisciplinary models (e.g. from 'reduced' process-oriented models to operational forecast systems). Supports the registration and dissemination of model data sets. An initial set of community-based numerical ocean models, such as Regional Ocean Modeling System (ROMS) and the Harvard Ocean Prediction System (HOPS), will be introduced.

Modeling Services: R-3 & R-4

Provides ocean modeling network services for access to multiple community-based numerical ocean models for parameter estimation/optimization and data assimilation. Provides the services to construct, modify, and execute numerical ocean models with command and control services for their operation and management. It provides a Virtual Measurement Sampling service to drive virtual instruments and/or virtual data acquisition processes. Services support: multiple models used in ensemble techniques, uncertainty and error estimation, and adaptive multi-domain 2-way nested configurations for generating dynamical interpolation of data sets, data assimilation, re-analyses (hindcasts), nowcasts, and forecasts.

Model Integration Services: R-4

Provides testing and validation services to ensure conformity with the different operational requirements in the network.

Virtual Ocean Simulator: R-4

Provides services to interact with the ocean through a simulator producing virtual ocean fields updated on a daily basis covering all three observatory types. The simulator involves on the order of twenty tracers including four physical variables (temperature, salinity, zonal and meridional current), a dozen biogeochemical variables (silicate, nitrate, phytoplankton, ammonium, two phytoplankton groups, two zooplankton grazers, two detrital pools, DIC, and oxygen), and four more tracers of interests (*e.g.*, tracers from hydrothermal event plumes).

Event Detection Services: R-5

Provides services to register processes to detect and publish events from data streams. Events are automatically persisted and distributed based on the configuration set for the detector.

Ocean System Simulation Experiments: R-5

Provides services to help select observatory sites and sensor network designs, as well as generate trade studies, data impact investigations, and data and information management exercises.

Resource Network Deliverable

Model Network: R-3

Establishes baseline processes and tools comprising a coherent framework for the analysis and assimilation of data. Provides the command and control semantics for interacting with a Model resource.

2.3 Planning & Prosecution

Duration: R-4 & R-5

Service Deliverables

Virtual Observatory: R-4

Provides the services to design, assemble, and operate configurations of resources from across ORION into unique systems for planning, testing,and prosecuting observation requests, leveraging the nested and autonomous capabilities of the fully integrated network of sensing, modeling, and control resources. Provide experimentalists with services to define, compose, and schedule multi-instrument observations that can execute across the observatory. As an example of a simple observation statement: on event "X" provide a CTD and a current profile of region "Y" using gliders "A, B, C" in configuration "Z" using behavior scenario "W".

Event Response Services: R-4

Provides services for policy- and behavior-based reconfiguration of tasks and observational programs. Provides a nested communication, command, and control architecture which enables and supports the deployment and prosecution, fully autonomously or under operator control, of new missions, processes and behaviors, in parallel to and without interruption of prior platform objectives.

Portable Control Software: R-4

Provides a portable, platform-generic higher-level control software package based on the public-domain MOOS mission control software that can run natively on fixed observatory assets, and for download and implementation into platforms such as gliders and AUV's operated in the observatory. The software provides a standard communication, command, and control connectivity with the overall ORION CI, and a standard NMEA interface to native control software on the platforms.

Planning Services: R-4

Provides software tools and user interfaces for the scientist defining a set of states for each fixed or mobile node involved in a planned experiment or observation campaign, and to design the associated, conditional state transitions, forming the basis for defining the behavior algebra necessary to complete a pre-determined, as well as autonomously adaptive sensing task.

Multi-objective Mission Plans & Behavior Repository: R-5

Provides and maintains platform specification, planning element, and plan and behavior modules for a variety of ocean sensing missions, such as the capture of a coastal upwelling event. A representative set of plan and behavior modules that adhere to a full Boolean logic pre-condition language for generically-conditioned autonomy actions will be introduced.

Mission Services: R-5

Provides standard safety procedures protecting the fixed or mobile assets that could be damaged through improper use by inexperienced operators, such as collision control for multiple AUVs and assurances of depth-limits for sensor packages.

Mission Simulator: R-5

Provides a complete mission simulation capability for pre-deployment planning and testing of specific measurement campaigns. Seamlessly linked to the ORION Virtual Ocean Simulator, this enables comprehensive testing of pre-determined as well as adaptive missions, such as the capture and measurement of a rapidly developing coastal front or a subsea volcanic eruption.

Resource Network Deliverable

Control Network: R-5

Enhances the control network to incorporate multi-objective optimization for behavior reconciliation, with a Boolean logic behavior calculus.

2.2.4 Data Management

Duration: R-1 through R-3

Service Deliverables

Data Repository: R-1

Maintains data and data products with their associated metadata.

Archive Services: R-1 & R-2

Provides cataloging, preservation, and curation services to index, maintain, and present the data holdings of an individual, group or community.

Aggregation Service: R-2

Provides for the classification, categorization, and general grouping of resources into collections.

Attribution Service: R-3

Associates and retrieves attributes to resources. The attributes can be associated within a semantic context (ontology). The service facilitates the characterization, qualification, and general commentary about the elements that participants interact with.

Metadata Search & Navigation Services: R-2

Provides query and browsing services by content and structure of the metadata.

Dynamic Data Distribution Services: R-2 & R-3

Provides publication, subscription, and query services associated with variant and dynamic data resources. Used in combination with the Processing Service to drive the policy decision to execute a process.

Data Access Services: R-2 & R-3

Provides an extensible suite of access interfaces and data formats for interoperability with external communities and applications (*i.e.*, OPeNDAP, THREDDS, LAS and the suite of OGC Web services).

Resource Network Deliverable

Data Network: R-1

Provides an automated data distribution and preservation network with pervasive and universal access subject to ORION Data Policy.

2.2.5 *Knowledge Management*

Duration: R-4 & R-5

Service Deliverables

Classroom Facility Services: R-4

Provides services to design and present curricula and demonstrations. Maintains a repository of frequently asked questions in the form of queries.

Data Integration Service: R-4

Provides querying and the integration of results from a heterogeneous class of complex data sources (*e.g.*, geospatial, temporal, spatiotemporal, relational, XML, and ontological data structures).

Semantic Context Service: R-4

Provides for the development, negotiation, and provisioning of vocabularies, thesauruses, and ontologies to be employed by individuals, groups, and communities.

Vocabulary & Ontology Repository: R-5

Maintains vocabularies and ontologies with their associated mappings. A representative set of the major vocabularies and ontologies in use within the oceanographic research community will be provided. A set of relevant reference resources such as regulatory boundaries and environmental climatology will be provided as the initial basis for characterization. An initial set of ontologies for instrument types, habitats, etc. will be included.

Semantic Search & Navigation Services: R-4 & R-5

Provides query and browsing services by concept or associated attributes (e.g. instrument or data product discovery by quality of measurement within an ecological habitat). A vocabulary and/or ontology can be used to provide greater scope to the query.

Resource Network Deliverable

Knowledge Network: R-4

Provides the capability to query, deliver, exploit, and annotate data from across the observatory based on a localized semantic frame of reference (ontology).

2.2.6 *Common Operating Infrastructure*

Duration: R-1 through R-3

Service Deliverables

Facility Services: R-1 & R-2

Provides the management and governance services for a collection of resources on behalf of a group or individual. It represents the domain of authority for the set of resources managed by the facility. The governance services provide for the following set of collaboration agreements: membership, partnership, federation, and delegation. Delegation, for example, is used to give a marine observatory the rights to operate/manage a research team's instrument on their behalf.

Resource Repository: R-1

Maintains all relevant information associated with resources registered with the system.

Resource Lifecycle Service: R-1

Provides resource management services to transition a resource from cradle to grave.

Resource Integration Services: R-2

Provides testing and validation services to ensure conformity with the different operational requirements in the network.

Resource Collaboration Services: R-2

Facilitates the negotiations between participants for sharing of resources (*e.g.*, instruments, processes, and models). Agreements are captured and associated with all parties materially involved.

Identity Management Services: R-1 & R-3

These services provision and securely manage information about participants used in the governance (*i.e.* authentication, authorization) of their activities across the network. The services ensure that personal information is owned and its exposure to other participants is controlled by the participant.

Repository & Catalog Services: R-1 & R-2

These services provide for persistence, preservation, and retrieval of information elements.

Resource Network Deliverable

Resource Network: R-1

Provides the standard semantics and behaviors for the management, federation, delegation and governance of a resource. This is the base implementation of the resource network model. From this all other resource networks, control, processing, data, etc., are derived.

Infrastructure Frameworks Deliverables

Presentation: R-1

Provides the web service & browser presentation containers as well as the web user interface "portlet" building blocks.

Service: R-1

System for provisioning, federating, delegating, and binding service interactions between resources.

Resource: R-1

System for provisioning, managing, and tracking the use of resources.

Governance: R-2

System for Identity Management, governing the use of resources by participants through policy enforcement & decision services.

State: R-1 & R-2

System for managing active and persisted distributed state.

Communications: R-1 &R-2

System for messaging, bulk data transfer, guaranteed data transfer and provisioning stream media channels.

2.2.7 *Common Execution Infrastructure*

Duration: R-2 & R-3

Service Deliverables

Deployment Repository: R-2

Maintains references to registered execution sites and Virtual Compute Node configuration packages.

Elastic Computing Services: R-2 & R-3

Provides the scheduling, provisioning, and monitoring services to maintain a balanced deployment of Virtual Compute Nodes to the computational engines.

Process Repository: R-3

Maintains process itineraries and references to registered process engine configurations and execution sites.

Processing Services: R-2 & R-3

Provides the validation, scheduling, and management services for policy-based process execution at specified execution sites. The service supports the coupling of the dynamic data distribution service with the process and its triggering. Provenance and citation annotation are registered associating the input and output products with the execution process and its operating context.

Resource Network Deliverables

Control Network: R-2

Establishes standard models for the management of stateful and taskable resources.

Process Network: R-2

Provides immediate mode scheduling of processes at specified locations within the integrated network based on explicit time requirements and/or event triggers.

Infrastructure Frameworks Deliverable

Execution: R-2

System for distributed orchestration and dispatch.

2.2.8 Software Release Roadmap

The five releases are specifically designed to build from the foundation of data collection and preservation up to advanced concepts of interactive ocean science supporting:

- real-time modeling and data assimilation.
- adaptive sensing and platform control.
- rapid response and event capture.
- closed loop, integrated sensing, modeling, and distributed network control.

Even though multiple subsystems and infrastructure elements are being developed simultaneously, each release has a specific theme targeted at providing value to the end users and building on value previously delivered:

R-1 End-to-End automated Data preservation and distribution.
R-2 End-to-End control of how Data are collected.
R-3 End-to-End control of how Data are processed.
R-4 Control of Models driven by the Collection Process.
R-5 Control of Data, Processes, and Models to drive the Collection Process.

One the great advantages of laying down the data system in concert with the interoperation of resources based on shared semantics early in the project is that while all of the advances in interactive capability are being developed, the data and knowledge management teams can work behind the scenes to capture, organize and associate the structure and semantics of each of these activities. This yields one of the great values of ORION, the End-to-End capability to preserve and present the entire collection process and its outcomes as knowledge.

Release 1: End-to-End automated Data preservation and distribution

Sensing & Acquisition	Observatory Services Instrument Direct Access Service
Data Management	Data Network & Repository Archive Services
Common Operating Infrastructure	Resource Network & Repository Resource Lifecycle Services Facility Services Repository & Catalog Services Infrastructure Frameworks w/out Policy-based Governance, Distributed Execution
Implementation & Integration	Primary CI CyberPoP

Release 2: End-to-End control of how Data are collected

Sensing & Acquisition	Instrument Network & Repository Instrument Services Data Acquisition Services Data Calibration & Validation Services
Data Management	Aggregation Service Metadata Search & Navigation Services Dynamic Data Distribution Services THREDDS/OPeNDAP Data Access Live Access Server Data Access
Common Operating Infrastructure	Resource Integration Services Resource Collaboration Services Infrastructure Frameworks w/ Policy-based Governance

Common Execution Infrastructure	Control Network Deployment Repository Elastic Computing Services Processing Services Infrastructure Frameworks w/ Distributed Execution
Implementation & Integration	GSO CyberPoP

Release 3: End-to-End control of how Data are processed

Sensing & Acquisition	Data Process Repository Instrument Integration Services product development Services
Analysis & Synthesis	Laboratory Services Data Analysis & Visualization Services & Repository Modeling Services
Data Management	Attribution Service OGC: Web Coverage, Mapping & Feature Services OGC: Sensor Web Enablement Suite (*i.e.*, Sensor Observation & Sensor Planning Services)
Common Execution Infrastructure	Process Network & Repository Processing Services coupled w/ Dynamic Data Distribution (Streaming Data & Events)
Implementation & Integration	GSO CyberPoP w/ Observatory Suite Services IOOS Prototype Integration

Release 4: Control of Models driven by the Collection Process

Analysis & Synthesis	Model Network & Repository w/ ROMS & HOPS Model Integration Services Virtual Ocean Simulator
Planning & Prosecution	Virtual Observatory Event Response Services Planning Services
Knowledge Management	Knowledge Network Classroom Facility Services Data Integration Service Semantic Context Service Semantic Search & Navigation Services
Implementation & Integration	RCO CyberPoP w/ Observatory Suite Services Teragrid Integration

Release 5: Control of Data, Processes, and Models to drive the Collection Process

Analysis & Synthesis	Event Detection Services Ocean System Simulation Experiments
Planning & Prosecution	Mission Plans & Behavior Repository Mission Services Mission Simulator
Knowledge Management	Vocabulary & Ontology Repository
Implementation & Integration	Secondary CI CyberPoP RCO High Bandwidth Stream Processor CyberPoP w/ HD Video Services Secondary Grid Infrastructure - Open Science Grid (candidate)

3.0 PROJECT MANAGEMENT

3.1 Project Management

The Project Execution Plan (Appendix 6.1, Section 2.1) fully defines the project organizational structure **(Exhibit 27: Organizational Structure)**. Beginning at the top, the Project Director (John Orcutt) serves as Principal Investigator for the CI with overall authority and responsibility for the project. He serves as the principal point of contact with the ORION Program Office, and appoints the Deputy Project Director (Frank Vernon), Project Manager (Matthew Arrott), Project Scientist (Oscar Schofield & Scott Glenn), and Education and Outreach (E&O) Manager (Cheryl Peach). Reporting to the Project Director, the Deputy Project Director is responsible for oversight of internal project functions, and will serve as the external point of contact in the absence of the Project Director. He approves plans and reports produced by the Project Manager, Project Scientist, and E&O Manager.

The Project Manager has day-to-day responsibility for all project life cycle activities. He carries out project planning, holds regular meetings of the entire project team, assesses cost and work progress against plans and schedules, and maintains up-to-date projections of performance. The Project Controller and administrative staff support the Project Manager. The Project Scientists have primary responsibility for the scientific integrity of the CI. They are responsible for the CI Science User Requirements and for validation of the system prior to each deployment. The E&O Manager guides the development of education and outreach plans for the CI and maximizes the opportunities for relevant communities at system level reviews. Each of these positions reports to the Deputy Project Director.

The System Engineer (Alan Chave) is responsible for management of the system life cycle, integration of the CI with the observatory elements of the OOI and external observatories, and verification of the result during each development spiral. In this capacity, he develops, verifies, and maintains all system level engineering policies and plans. The System Architect (Ingolf Krueger) is responsible for design, synthesis, and documentation of the system software architecture and oversight of its implementation by Subsystem Development Teams. The System Engineer and System Architect are jointly responsible for requirements definition and internal interface control. The System Engineer is responsible for negotiating interface agreements with the observatory IO system engineers or external entities like IOOS, subject to Program Office oversight and approval. Both the System Engineer and System Architect report to the Project Manager.

The System Development Manager oversees the activities of all of the Subsystem Development Teams, and is responsible for delivering a quality, integrated CI to the System Engineer during each development spiral. The Operations Manager is responsible for the planning and execution of all deployments of the CI, and for post-deployment operations and maintenance. The Quality Assurance Engineer is responsible for quality control during each development spiral. These positions report to the Project Manager.

Seven Subsystem Development Teams will be constituted to construct and deliver subsystems for integration, and to assist in integration with other subsystems. The leads of each team report to the System Development Manager, and the System Architect provides coordination for their activities. The System Integration Team supports all integration and verification activities and reports to the System Engineer.

The Project Execution Plan (Appendix 6.1, Sections 3.2-3.9 and 4.1-4.10) defines a series of management tasks for which one of the Project Manager, Project Scientist, E&O Manager, System Engineer, System Architect and Operations Manager positions is responsible. Each of these is governed by a plan that defines scope, procedures, responsibilities and authorities. Collectively and with the project organizational structure, these define the management model for the project. The management tasks include:

- Business management, including Earned Value Management (EVM).
- Work Breakdown Structure (WBS) and Dictionary, and Master Project Schedule.
- Risk opportunity and mitigation management.
- Communications, configuration and change control management.
- Integrated logistics support.
- Education and outreach activities.
- Work and operations plans and reports.
- Detailed Project Schedule and performance baseline.

- System life cycle management.
- Integration and verification management.
- Validation, deployment and acceptance management.
- Quality and security management.
- Operations and maintenance management.

Further details may be found in the Project Execution Plan (Appendix 6.1).

The ORION project is a highly distributed observatory comprising coastal, global and regional observatories at a wide variety of spatial scales. At present, none of the other IOs have been selected, although announcement of the regional cabled observatory IO is imminent. A letter of support from the PI of a proposed RCO IO is contained in Section 10.0. Close collaboration between the CIIO and the other OIs is essential for success. The coordination of the OOI through the ORION Office constitutes a challenge for managing dispersed independent operators, especially when the operators themselves are likely to integrate experts from around the country. This is certainly the case with the proposed CIIO, where it is beneficial to take advantage of distributed expertise in a wide range of computer sciences and engineering specialties. Insofar as possible, the key personnel are located at UCSD and Scripps, but there are exceptions. We do have considerable experience in overcoming the distance barriers effectively, and believe these can be applied to interacting with the other IOs.

As an example, the NSF OptIPuter project significantly furthers collaboration by expanding modern videoconferencing. Through the use of 10Gbps point-to-point connections for HD videoconferencing and large display walls coupled with HDTV, scientists are encouraged to collaborate by sharing large data sets. High-speed Internet (or more accurately National Lambda Rail) connections between many of the CIIO participants will greatly aid the management of the proposed work. Examples of projects where frequent video meetings have been effective include (1) the OptIPuter; (2) the Moore Foundation sponsored Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis (CAMERA) integrating the Venter Institute (Craig Venter), Calit2(Larry Smarr) and the Center for Earth Observations and Applications (CEOA—John Orcutt); and (3) the proposed UCSD Energy Biosciences Institute (EBI) between UCSD, Iowa State University, the Venter Institute, Battelle Memorial Institute (PNNL), The Scripps Research Institute, and the Salk Institute. In the latter example, John Orcutt led the partners in developing the proposal, an activity made possible by daily videoconferencing.

While technology will lead to many new approaches to managing large projects, the current management plan clearly defines roles within the CIIO and is classically hierarchical. Furthermore, we have spent considerable time building trust relationships between participants through an exhaustive travel/meeting schedule and anticipate this will continue throughout the life of the program.

3.2 Key Personnel

The five key personnel for the CI project are:

- Project Director/Principal Investigator **John A. Orcutt** (UCSD)

Orcutt has been at Scripps for 33 years and has spent 10% of his lifetime at sea. He has been involved in and led the development and use of autonomous seafloor instruments (largely seismographs) since 1974. His administrative responsibilities have included Director of the Cecil H. & Ida M. Green Institute of Geophysics and Planetary Physics (IGPP) for 18 years, Deputy Director of Scripps for four years, and, currently, Director of the UCSD Center for Earth Observations and Applications (CEOA) and Associate Vice Chancellor for Research with responsibility for research across School (e.g. Medicine), Division (e.g. Biology), and Departmental boundaries. His is Past President of the American Geophysical Union (AGU) with 48,000 members. His current research interests include broadband seismology in the oceans as applied to understanding mantle structure and dynamics. He has worked for the past six years in extending visualization and information technology methodologies to global observations in real-time.

- Deputy Project Director **Frank L. Vernon** (SIO/UCSD)

Vernon has over 25 years experience designing, deploying, operating, and maintaining real-time seismic sensor networks throughout the world. His current research interests are focused on developing distributed real-time sensor networks in terrestrial and marine environments. Specific programs include the ROADNet project (roadnet.ucsd.edu) that is developing the real-time software systems to acquire, process, distribute, and archive data from environmental, oceanographic, geophysical, and structural monitoring sensor nets; the USArray

Array Network Facility (anf.ucsd.edu) collecting real-time data as part of the NSF EarthScope MRE; and HPWREN (hpwren.ucsd.edu), a large-scale wireless high-performance data network used for interdisciplinary research and education applications.

- Project Manager **Matthew Arrott** (Calit2/UCSD)

Arrott has over 20 years experience in project management, design leadership, and engineering management for software systems focused on information design and financial services. His current projects are focused on remote interactive instrument networks and the policy-based architectures to govern their secure and safe operations. Specific positions held include Currenex, a leading independent global electronic multi-bank foreign currency exchange linking the Fortune 1000 worldwide. As VP of Product Development, he designed, architected, and delivered the industry's first multi-bank "Executable Streaming Price" product, distributing over 2 million real-time pricing events daily from financial institutions globally. At Dreamworks SKG, he served as Head of Software responsible for 3 software development departments, R&D, production software and asset management systems groups. At Autodesk, as Software Development Manager & Systems Architect, he designed and executed a 2-year project to replace the entire AutoCAD graphics subsystem with the HEIDI Graphics System, a hierarchical multi-tiered rendering architecture.

- Project Scientist **Oscar Schofield & Scott Glenn** (Rutgers University)

Schofield and Glenn have developed and operated the coastal Ocean observation laboratory for over a decade, hosting large science campaigns. They have been actively involved in the inception of the NSF ORION program, the NOAA Integrated Ocean observing System (IOOS), and the DOD Littoral BattleSpace Fusion and Integration (LBSF&I) program. Their research group has pioneered the use of electro-optic cables, shore-based radars, numerical data assimilative forecast models, and underwater autonomous vehicles.

- System Engineer **Alan D. Chave** (Woods Hole Oceanographic Institution)

Chave has over 25 years of experience with oceanographic technology development. Recent efforts include the design of an underwater optical modem system capable of Internet speeds over distances of 200 m, and the design of an extensible data communications infrastructure for ocean observatories that includes a high reliability, out-of-band communication channel capable of delivering synoptic high accuracy time to the seafloor. Both of these are patent pending. Chave also serves as the Project Scientist for the LOOKING project conducting research into ocean observatory cyberinfrastructure. He was trained in system engineering as an undergraduate at Harvey Mudd College, and has taken post-graduate courses in the topic.

These personnel represent national/international expertise in oceanography and IT, and they are responsible for the internal oversight of all project functions, day-to-day management of the project life cycle, representation of the science user community perspective, and day-to-day management of the system development life cycle, respectively. This collection of activities represents the principal management functions required to implement the ORION CI. Resumes are contained in Section 9.0.

In addition to the five key personnel, the management team includes:
- Education and Outreach Manager Cheryl Peach (SIO/UCSD)
- System Architect Ingolf Krueger (Calit2/UCSD)
- System Development Manager (TBD)
- Operations Manager (TBD)

The roles, responsibilities, and authorities of all these positions are defined in the Project Execution Plan (Appendix 6.1, Section 2.1). Persons to fill the last two positions will be recruited from within UCSD if possible, and from outside if necessary.

3.3 Education and Public Awareness

A well executed education and public awareness (EPA) program is critical to the long-term success of ORION. For ORION EPA to reach its full potential, the OOI EPA efforts must focus on building the educational infrastructure and prototypes essential for successful EPA plan implementation. As ORION's integrating IO, the CIIO serves as a natural nexus for OOI EPA activities. By providing the breadth of services and interactivity required for a robust ORION EPA program and accommodating the technical capabilities and needs of ORION target audiences, the CIIO will play a pivotal role in creating the bridge between observatory resources and a diverse set of observatory users.

In addition to scientific users skilled in the acquisition and manipulation of large digital data sets, other potential users of observatory data include scientists less familiar with these processes, observatory operators, graduate and undergraduate students, K-12 and informal educators and students, the public and policy makers. To effectively reach this range of audiences, a multi-faceted approach is envisioned in which the CIIO EPA team will work collaboratively with the ORION program office and other IO EPA teams to develop:

CI training programs for scientists, observatory operators, graduate and undergraduate students and educators: The training will be conducted using a phased approach in which training priorities evolve as ORION moves from planning to operation. The initial focus will be on the research community (two to three years) and subsequently, in consultation with the ORION project office, will expand to include other target audiences. Training will be conducted both on site and online, and approaches and materials will be designed to empower those initially taught to subsequently instruct colleagues, students, and other users.

Educational Prototypes: The CIIO will be the entity where education systems are developed, tested and refined in advance of the observatory systems going live. Examples of possible prototype products include graphing and mapping tools, simple modeling tools for non-scientists, and online investigative learning activities for secondary and college students. The CI EPA team will work with the ORION Program Office to identify and prioritize key audiences, assess user needs and design, develop, and test educational prototypes.

Visualizations and Animations: A critical component of the ORION EPA program, including a far-reaching media and public relations campaign, is the capacity to visualize observatories and observatory data in ways that are accessible to non-scientists. Working collaboratively with the ORION project office and IO EPA teams, visualizations and animations will be designed to serve as central components of the ORION EPA web portal (see below) and to support ORION scientists and the ORION Project Office in their EPA activities. Initially, visualizations may be developed to elucidate three to five ORION science themes. Over time, a visualization resource center will be established for use by the entire ORION community.

ORION EPA Portal: The CIIO EPA team will develop a Google-aware web interface that will serve as a gateway to ORION CI EPA activities, products and programs. The portal will serve as a key source for educational products, images, visualizations, Maya animations and other public relations information for the ORION community. The CI EPA team will work closely with the ORION Program Office to integrate the CI EPA web portal into the ORION EPA web site.

The four thrusts described above constitute a ORION EPA cyberinfrastructure that will be necessary to support the wide array of ORION EPA activities and users envisioned in the ORION EPA strategic plan.

3.4 Labor Projection Table

This section defines the allocation of labor in FTEs per year to a core set of tasks:

Project Office: management of the project life cycle by the Deputy Project Director, Project Manager, Project Scientist, Project Controller and administrative support staff.

System Engineering: management of the system life cycle by the System Engineer, System Architect and their support staff.

	Y-1	Y-2	Y-3	Y-4	Y-5	Y-6	Total	
	15.90	22.35	22.90	19.13	16.30	14.30	110.9	
Project Office	3.50	3.50	3.25	2.80	2.80	2.80	18.7	17%
System Engineering	3.85	4.00	3.55	2.63	1.70	1.00	16.7	15%
Software Development	8.05	13.10	13.60	11.20	9.05	7.50	62.5	56%
Management	1.00	1.00	1.00	1.00	1.00	1.00	6.00	
Subsystem Projects	6.55	11.10	11.60	9.20	7.05	5.50	51.00	
Sensing & Acquisition	1.30	2.20	2.10	1.00	-	-	6.60	6%
Analysis & Synthesis	0.15	0.45	1.20	1.90	2.20	1.10	7.00	6%
Planning & Prosecution	0.15	0.10	0.35	1.40	2.20	2.20	6.40	6%
Data Management	2.25	2.90	2.05	1.10	0.25	-	8.55	8%
Knowledge Management	0.15	0.10	0.45	1.70	2.40	2.20	7.00	6%
Common Operating Infrastructure	2.20	3.50	2.25	0.50	-	-	8.45	8%
Common Execution Infrastructure	0.35	1.85	3.20	1.60	-	-	7.00	6%
Quality Assurance	0.50	1.00	1.00	1.00	1.00	1.00	5.50	
Hardware Development	-	-	-	-	-	-	-	0%
Implementation	0.50	1.75	2.50	2.50	2.75	3.00	13.0	12%

Software Development: All subsystem design and integration activities. This category is broken down by subsystem.

Hardware: The CIIO does not propose to develop hardware. All hardware used is accounted for in Software Development or Implementation.

Implementation: Provisioning and integration with the three ORION Observatory IOs with their respective CyberPoP configuration, integration with external data and computational resources (*i.e.*, TeraGrid and IOOS), and the integration of the core science instruments and specified suite of analytical applications.

3.5 Participant Responsibilities

Exhibit 28: Project Participant Labor presents the level of effort in FTE of each participating organization by task area (as defined in Section 3.4, Labor Projection Table) and calendar year. In most instances, the subcontractors are academic institutions or federally-funded research and development centers (FFRDC) chosen for the expertise of specific individuals that they employ. The only industrial subcontractors are Triad Project Management and Raytheon. Triad has considerable experience with MREFC projects, and will provide financial management consulting to the Project Office. Raytheon will provide project management and system engineering support for the Project Manager and System Engineer, respectively, and also provide a key link into the emerging IOOS cyberinfrastructure.

4.0 TASK AREA EXPERIENCE

Our experience with projects of comparable size and scope to ORION CI is documented in Appendix 6.2, Task Area Experience at the end of this volume. Choosing among numerous potential examples, we opted to include:

- *Southern California Coastal Ocean Observing System (SCCOOS/COCMP):* illustrating experience in all six task areas;
- *US Array: Array Network Facility (ANF):* illustrating experience in project management;
- *Persistent Littoral Undersea Surveillance Network (PLUSNET/ASAP):* illustrating experience in system engineering;
- *Storage Resource Broker (SRB)/Integrated Rule-Oriented Data System (IRODs):* illustrating experience in software development;
- *MARS Data Communication Subsystem:* illustrating experience in hardware development;
- *Biomedical Informatics Research Network (BIRN):* illustrating project implementation;
- *CI Coordination of Science Campaigns (SW06):* illustrating experience with operations and maintenance.

Appendix 6.2, Task Area Experience provides project descriptions, dates, and funding levels as well as information on the participation of key personnel.

5.0 KEY RESOURCES

5.1 Facilities

The strength of our collaboration is reflected in the lengthy compilation of facilities, technical capabilities, and instrumentation found in Appendix 6.3, ORION CI Facilities at the end of this volume.

5.2 Space Requirements

The ORION CI project office, lab and workspaces will be housed in UCSD's Division of the California Institute for Telecommunications and Information Technology (Calit2), adjacent to many of our project collaborators. We have received office and workspace to accommodate an average of 12 people, and sufficient lab space to support hardware system assembly and instrument integration.

5.3 ORION and Institutional Goals

5.3.1 Integration with Institutional Research.

The CIIO contract will dovetail substantially with many research programs at UCSD and Scripps. Related programs are shown in the table.

Major Related Research at UCSD & Relationship to personnel involved in this proposal

Program	Sponsor	Related Personnel
Southern California Coastal Ocean Observing System (SCCOOS)	State of California & NOAA	Orcutt (PI), Terrill (PI), Yi, Vernon
California Ocean Current Monitoring Program (COCMP)	California	Orcutt
Community Cyberinfrastructure for Advanced marine Microbial Ecology Research & Analysis (CAMERA)	Moore Foundation	Orcutt, Arrott
UCSD Energy Biosciences Institute (EBI)	BP-proposed	Orcutt (lead), Vernon, Arrott
Laboratory for Ocean Observatory Knowledge Integration Grid (LOOKING)	NSF	Orcutt (PI), Chave, Arrott, Vernon, Delaney
EarthScope USArray Array National Facility (ANF)	IRIS/NSF	Vernon (PI)
Real-time Observatories, Applications, and Data management Network (ROADNet)	NSF	Orcutt (PI), Rajasekar (PI), Vernon (PI)
HiSeasNet	NSF	Orcutt, Berger, Vernon
Center for Earth Observations & Applications (CEOA)	Multiple	Orcutt (Director)
Biomedical Informatics Research Network (BIRN)	NIH	Arrott
National Surface Currents Mapping Initiative	NOAA	Terrill

UCSD is a major leader in research, development, and applications in computer sciences and engineering. The projects listed above represent a strategic direction for UCSD as a whole, and computer/observational sciences in particular, including the Scripps Institution of Oceanography. The mission of CEOA, directed by John Orcutt, is to stimulate, support, and coordinate sustained research and applications in Earth observations at UCSD. The Computer Science & Engineering Department in the Jacobs School of Engineering conducts basic research in computer science and engineering. The Jacobs School of Engineering is named after Prof. Irwin Jacobs who was a faculty member in the Department and best known as the founder and original CEO of Qualcomm, one of the world's largest wireless companies. The California Institute of Telecommunications & Information Technology (Calit2) is a two-campus institute including UCI that dates to 2002. The Director, Larry Smarr, characterizes Calit2 as "a new mechanism to address large-scale societal issues by bringing together multidisciplinary teams of the best minds." The San Diego Supercomputer Center (SDSC) was founded in 1985 and its current activities range from participation in the NSF's TeraGrid, to hosting data for the Protein Data Bank, to GEON. The storage capacity at SDSC has recently increased to 25 PetaBytes and SDSC is competing for the TeraScale computing facility in the most recent NSF Cyberinfrastructure Office competition. Scripps Institution of Oceanography (SIO) partners with all of these entities, but most closely with Calit2 for research and development, and SDSC for applications and operations.

The projects in the table above are all related to the proposed work. For example, for the EBI, a sensor data and control network, similar to that proposed here, will be used for command and control of analytical instruments in a distributed, virtual laboratory and, notably, a modern, prototype biofuels plant in the UCSD EBI laboratory itself. This will allow our partners at the Venter Institute and Iowa State University to participate directly in experiments. HiSeasNet presently uses ROADNet to return data from Scripps' research vessels, the

R/V Melville and Roger Revelle, in real-time via satellite. HiSeasNet is now installed on nearly all large and intermediate ships in the UNOLS fleet. Making the fleet an element of the global observatory would be straightforward and would provide an excellent testbed for ORION CI. Finally, the marine genome database with environmental metadata will go on line in January 2007, more than doubling the number of available marine microbe genomes. Future extensions of the OOI to include microbial ecosystems can build upon this growing database for identifying species and understanding ecosystem responses to changing ocean climate.

5.3.2 Integration with Institutional Education and Public Awareness Activities

A central component of Scripps Institution of Oceanography's mission is to train the next generation of leaders in ocean sciences. Recognizing that technological advances have played a critical role in many, if not most, of the advances in our understanding of the ocean, SIO has a long tradition of providing students with training not only in research methods, but in the skills required to excel in an arena that is increasingly technical, cross-disciplinary and dependent on large data sets. SIO/UCSD and its partners are committed to preparing the next generation of researchers to capitalize on the unprecedented volume of data from the observatories to transform the conduct of ocean sciences research. Graduate student participation in the CIIO will be a central element of that preparation.

SIO is increasingly engaged in initiatives addressing workforce development and student recruitment into academic and technical careers in ocean sciences. To that end, SIO is developing a new undergraduate marine sciences minor program designed to attract majors in physics, chemistry, biology and engineering into academic careers in ocean sciences. SIO is also a key player in the Centers for Ocean Sciences Excellence (COSEE), and has led the community in engaging researchers and graduate students in reaching out to students and the public to heighten awareness of ocean sciences and ocean science careers. As part of its COSEE activities, SIO was recently funded as a partner in a National Ocean Partnership Program project investigating the ocean observatory workforce needs of the future. Finally, keeping the students in the pipeline for academic and technical careers in ocean sciences requires effort at the pre-college level. SIO has recently established a collaboration with San Diego Unified School district (2nd largest in California) to provide online educational programming that supports their new 9th grade Earth sciences program. All of these initiatives provide excellent opportunities to pilot CIIO educational prototypes and training programs while simultaneously engaging undergraduate and high school students in educational experiences that are truly on the cutting edge of ocean science research.

UCSD has emerged as a national leader in outreach programs and facilities that raise public awareness and support for innovation in science and technology. Equipped with a unique suite of capabilities in multimedia production, visualization technology, broadcasting, web casting and distance education, UCSD and its campus partners (CALIT2, SDSC) deliver science content, cyberinfrastructure training, and educational programming to millions globally via programs such as UCSD-TV/SIO's Perspectives on Ocean Sciences lecture series, SDSC's CI Channel (a K-20 cyberinfrastructure online training program), and UCSD's Science Matters to name a few. Combining these global outreach programs with campus-based EPA activities, including ocean science exhibit and educational program development at UCSD/SIO's world-class public science center, the Birch Aquarium, the CIIO partners are well positioned to leverage ongoing activities to promote CIIO EPA activities. This broad spectrum of on site, on air and online capabilities provide not only key foundational pieces of the educational cyberinfrastructure for the OOI, but programmatic elements that can enhance the overall ORION EPA program.

Appendix 6.1-Cyberinfrastructure Project Execution Plan (PEP)

ORION CI Proposal

1.0 PROJECT PURPOSE

Next generation studies of dynamic, interacting processes in the Earth-ocean-climate system require new *in situ* approaches to complement the more traditional ship-based, expeditionary science that has dominated oceanographic research for the past century or more. Routine, long-term measurement of episodic oceanic processes is crucial to continued growth in our understanding and predictive modeling of complex natural phenomena that are highly variable and span enormous scales in space and time. This access will be enabled by innovative ocean observatory facilities providing unprecedented levels of power and communication to access and manipulate real-time sensor networks deployed within the ocean. These facilities will empower entirely new approaches to science and enable education and outreach capabilities that will dramatically impact the general understanding of, and public attitude toward, the ocean sciences.

To accomplish this paradigm shift, ocean scientists require at least seven infrastructural capabilities they do not now have. They must be able to:

- fully and quantitatively characterize selected volumes of the ocean, the atmosphere overhead and the lithosphere beneath;
- receive information about all interrelated components of the system simultaneously, in real-time;
- recognize departures from the norm and observe emergent phenomena;
- conduct interactive experiments within the environment;
- reconfigure observational-sampling systems in response to events;
- assimilate *in situ* data efficiently into models that expand the space/time view of the data and feed back onto the measurement protocols; and,
- continue and expand this real-time interaction within the oceans for decades.

These functions can only be realized through the development of state-of-the-art cyberinfrastructure (CI). The ORION CI Conceptual Architecture contains further details on the information technology capabilities required to bind the three observatory elements of the OOI into a coherent system-of-systems. These documents are incorporated into this PEP by reference.

2.0 PROJECT STRUCTURE

2.1 Organizational Structure

2.1.1 *Management Team Roles and Responsibilities*

The eight person management team comprises the Deputy Project Director (chair), Project Manager, Project Scientist, Education and Outreach (E&O) Manager, System Engineer, System Architect, System Development Manager, and Operations Manager. The management team reports to the Project Director. The group will meet regularly to keep abreast of developments at the top levels of the project, changes initiated by the ORION Program Office or NSF, and the status of the three observatory IOs. The Project Director will attend the meetings as necessary to gather advice and gauge progress. The management team will seek consensus on strategic decisions that crosscut the project. In the event that consensus is not reached, the Project Director will make a final and binding decision.

The Project Director serves as Principal Investigator for the CI with overall authority and responsibility for the project. He serves as principal point of contact with the ORION Program Office, and appoints the Deputy Project Director, Project Manager, Project Scientist, and E&O Manager. The Project Director is the final and binding arbiter of all internal project conflicts that cannot be resolved satisfactorily at lower levels.

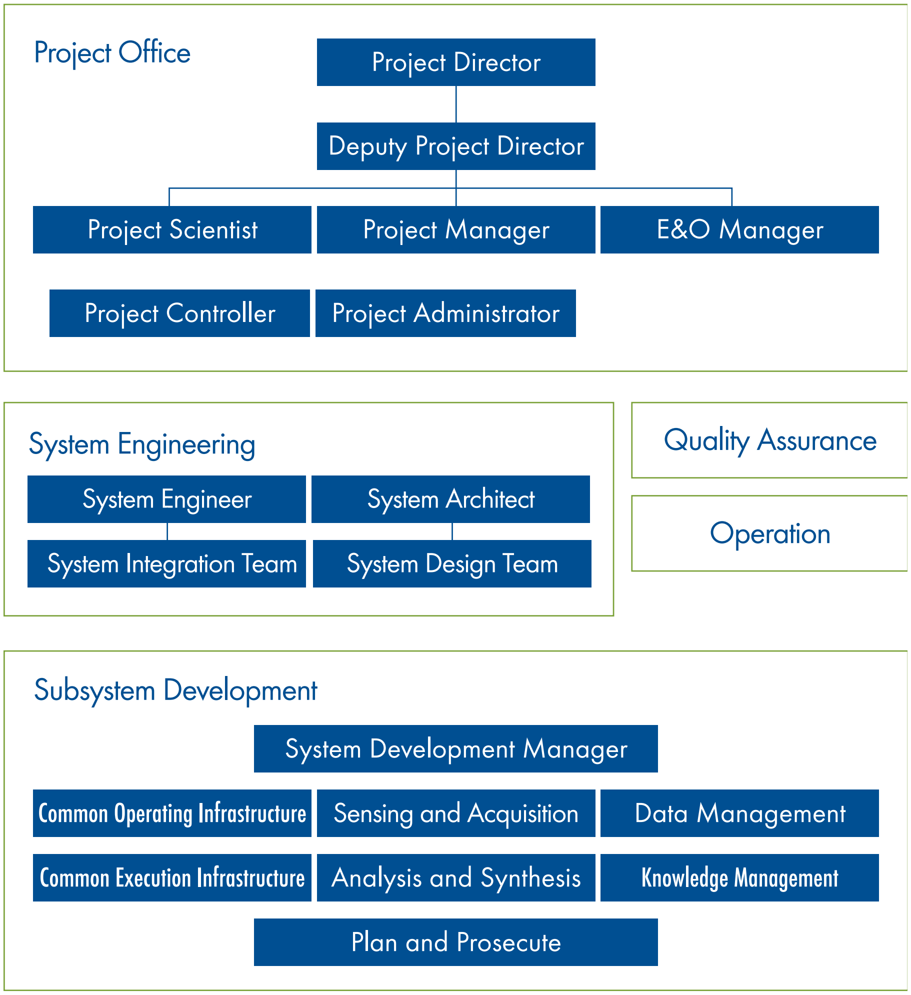

Reporting to the Project Director, the Deputy Project Director is responsible for oversight of internal project operations, and will serve as the external point of contact in the absence of the Project Director. He approves plans and reports produced by the Project Manager, Project Scientist, and Education and Outreach (E&O) Manager.

The Project Manager reports to the Deputy Project Director and has day-to-day responsibility for managing the project life cycle. He is responsible for key project planning actions, including generation of all project-level plans and approval of system-level plans. Other tasks include oversight of project activities to ensure timely correction of problems, convening regular meetings of the entire project team, assessing cost and work progress against plans and schedules including Earned Value Management, maintaining up-to-date projections of the project schedule and cost-to-complete/life-cycle costs, conducting design reviews and ensuring that the results of such reviews are incorporated into the project plans. The Deputy Project Director is the first point of contact for the Project Manager in resolving conflicts regarding resources not under his control and obtaining decisions beyond his authority. If these cannot be adjudicated by the Deputy Project Director, such issues will be referred to the Project Director, and, if necessary, the Program Office for resolution.

The Project Manager is also a member of the OOI Management Team chaired by the OOI Project Director, as established in the OOI PEP.

The Project Scientist reports to the Deputy Project Director and has responsibility for the scientific integrity of the CI and communication with the scientific community on CI issues. The Project Scientist will organize a stakeholder team comprising representatives of interested user groups to develop use case scenarios. He will be responsible for the CI science user requirements in consultation with working groups, the System Engineer and Architect, the ORION advisory committees and the Program Office as appropriate, and will be responsible for validation of the CI at the end of each development cycle.

The E&O Manager reports to the Deputy Project Director and is responsible for the development of education and outreach plans and for maximizing E&O opportunities for relevant communities at system-level reviews. She assists the Project Scientist with the development of user requirements, representing the interests of the E&O community. The E&O Manager will coordinate and integrate activities with the overall ORION E&O effort.

The System Engineer reports to the Project Manager and is responsible for management of the system life cycle, integration of the CI with the observatory elements of the OOI and external observatories, and verification of the result. He/she develops, verifies, and maintains all system-level engineering policies and plans. Together with the System Architect, the System Engineer assists with the definition of science user requirements, defines the system requirements and specifies internal system hardware/software interfaces in consultation with subsystem design teams. The System Engineer is a member of the OOI System Engineering Team as defined in the OOI PEP, and will negotiate external system interfaces with the observatory IO system engineers and/or non-OOI observatories.

The System Architect reports to the Project Manager and is responsible for the design, synthesis, and documentation of the system software, and the oversight and coordination of its implementation by distributed Subsystem Development Teams. He is also responsible for verification of the integrated CI.

The System Development Manager reports to the Project Manager and oversees the activities of all of the Subsystem Development Teams, and is responsible for delivering a quality integrated CI to the System Engineer during each development cycle.

The Operations Manager reports to the Project Manager and leads the Operations Team. The Operations Manager is responsible for planning and execution of all deployments of the CI, and for post-deployment operations and maintenance. The Operations Manager also provides critical input to the System Engineer and System Architect with the goal of minimizing the life cycle cost of the CI.

2.1.2 *Subsystem Development Teams*

Seven Subsystem Development Teams will construct and deliver subsystems for integration and promote integration with other subsystems. A Development Team Lead reporting to the System Development Manager is responsible for delivery of a quality subsystem. A Subsystem Development Team comprises the Lead, a Subsystem Architect, Expert Users, Design Participants, Development Participants, and Technology Providers. A single individual may play multiple roles. The Subsystem Architect provides the architectural vision. To ensure the delivery of an end-user focused product, Expert Users will work with each development team throughout the development life cycle. Design Participants assist the System Architect to produce the architecture documents relevant to their subsystem. Development Participants construct the subsystem. Technology Providers bring CI capabilities to the Development Team.

2.1.3 *Project Management Support*

The project office staff comprises the Project Controller, Project Administrator, and other support personnel as required and reports to the Project Manager. The Project Controller assists with project financial management activities, including Earned Value Management. The Project Administrator assists with general administrative activities.

The Quality Assurance Team oversees quality control during each development cycle. The key activities include monitoring the implementation of the project plans, controlling the acceptance of project deliverables, and providing input to the Project Manager on risk identification and mitigation. The Quality Assurance Engineer reports to the Project Manager and is responsible for system development quality assurance activities.

The System Integration Team reports to the System Engineer and assists with integration and verification of the CI and subsequent integration and verification of the CI with the OOI observatory elements and external observatory systems such as IOOS. The System Architecture Team reports to the System Architect, and assists in architecture definition.

Since the project and its personnel will be housed at existing institutions, it is expected that their human resources, property management, facility management, physical security, and supply chain management capabilities will support project activities.

2.1.4 Relation to the ORION Program Office

The project will adhere to the policies and constraints laid out in the OOI PEP that supersedes this document in the event of conflicts. This includes, but is not necessarily limited to:

- Participation in the cross-organizational structure defined by the Program Office.
- Compliance with international and interagency partnership agreements.
- Compliance with the accounting system, including the Earned Value Management component, defined by the Program Office.
- Adherence to a document control system consistent with document control at the ORION Project Office.
- Compliance with the ORION Data Policy.
- Submission of plans and reports for approval as required.
- Providing ex officio members of ORION advisory committees as required.

2.2 Project Life Cycle

Due to increasing problems (e.g. high costs) with the conventional waterfall development model, the spiral development model was introduced in the mid-1980s for the management of complex technology projects, especially those that are software-intensive. The spiral model was created by Boehm [1988], and a comprehensive treatment appears in Royce [1998]. The table here compares the waterfall and spiral development approaches. The spiral model addresses several fundamental waterfall model flaws:

- Critical risks (including requirement and design flaws) are often identified and resolved at advanced stages in a project, leading to unanticipated labor and costs.
- User needs and requirements evolve as the community better understands the design, leading to requirements changes.
- Late stage integration frequently uncovers problems with performance, quality, or user satisfaction, leading to significant unplanned work.

Waterfall Development Model	Spiral Development Model
Requirements may be completely defined in advance of development	Requirements are discovered and defined during development
Requirements contain no high risk implications that are unresolved	Risks are continually discovered and the reduction of high risk elements drives the development process
The requirements will change very little during development and operations	The requirements will evolve throughout the life cycle
The requirements are compatible with expectations of all key stakeholders	Ongoing negotiation with key stakeholders will be required to meet expectations
The best architecture to implement the requirements can be defined in advance	The best architecture to implement the requirements cannot be defined in advance
There is sufficient calendar time for sequential system development	Plans (including cost and schedule) are continually refined as the requirements and solutions become better defined

The spiral approach has become the *de facto* standard in the software industry, and is finding growing use beyond.

Boehm [2000] describes spiral development as a risk-driven process model generator for guiding multi-stakeholder concurrent engineering of complex systems. Its distinguishing features are

- A cyclic approach that incrementally grows the definition and implementation of a system while pulling risk forward.
- A set of anchor point milestones that ensure ongoing stakeholder commitment, and demonstrate the feasibility of the incremental definitions and implementations.

The figure on page 5 depicts a set of successive development spirals that are used to increase system definition from a concept to deployed products over time. The increasing definition typically occurs within four phases called inception, elaboration, construction, and transition [Royce 1998]. Anchor point milestones, called Life Cycle Objectives (LCO), Life Cycle Architecture (LCA), and Initial Operating Capability (IOC), provide and manage criteria for progressing from one phase to the next (see milestones table p5). In practice, more than one development cycle will be completed; for the proposed work, a total of five full cycles are planned.

The key activities during the inception phase are requirements discovery and conceptual architecture definition based on negotiation with and among stakeholders. This culminates in the LCO anchor point milestone that produces stakeholder commitment to building the architecture. The elaboration phase is focused on defining a feasible system architecture and addressing its riskiest elements. It ends with the LCA anchor point milestone that commits the stake-

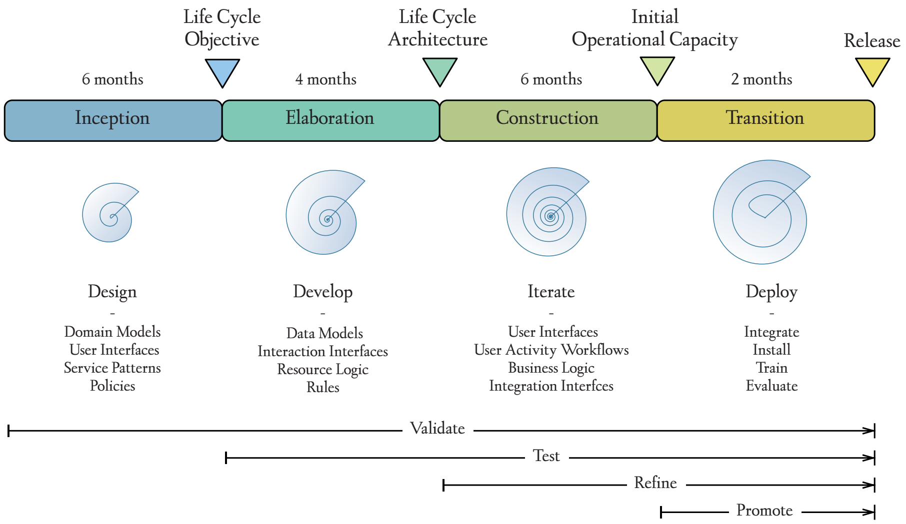

Life Cycle Objectives (LCO)

Goal: Definition of what the cyberinfrastructure design will accomplish.

Focus: Ensuring that at least one architecture choice is acceptable to the stakeholders

Stakeholder Commitment: Building the identified architecture.

Completion: End of planning phase

Life Cycle Architecture (LCA)

Goal: Definition of the software architecture and technologies needed to implement the cyberinfrastructure design.

Focus: Committing the project to a viable CI design.

Stakeholder Commitment: Supporting deployment of the cyberinfrastructure design.

Iterations: Five (every 12 mo starting at 18 mo)

Initial Operating Capability (IOC)

Goal: Approval of final production release.

Focus: Validation of a viable, maintainable system.

Stakeholder Commitment: Going into operations phase.

Completion: End of project.

holders to construction of the system. The construction phase is centered on building alpha and beta releases of the system. It terminates with the IOC anchor point milestone that commits the stakeholders to deployment of the system. It is followed by transition to a deployed system.

As the spiral model has evolved, variants (notably the Rational Unified Process; Kruchten [2004]) have been introduced, but all share several key properties. The system is defined, refined, and developed in multiple, risk-driven spiral iterations bounded by anchor point milestones. The details and impact of the risks (whether technical, management, operational, or stakeholder) drive the number of spirals and the level of detail and effort within each phase of a spiral. The riskiest elements are brought forward as early in the project as possible. Each spiral includes management, engineering, and support activities in proportion to the risks. Each spiral expands system definition and results in a deployed representation.

A tailored version of the spiral development model as described by Royce [1998] will be used in the project. This will require some modification of the documentation and schedule requirements outlined in the CI RFP, as described throughout this PEP.

2.3 Key Deliverables

The key project deliverable is an integrated, verified, validated, and deployed CI. The CI will be delivered in an increasingly advanced form every 12 months beginning 18 months after project inception. Interim releases that provide incremental capability improvement may also be released at the discretion of the Project Manager.

The project will also deliver a range of plans, reports, and manuals. The table on page 6 maps the project documents onto those specified in the CI RFP, and includes additional documents. Planning documents including this PEP will be provided in initial form within six months of project inception, and will be updated at least annually. Reports will be provided on a schedule specified by the Program Office and consistent with the budget.

2.4 Decision Making Process

The project is organized in teams that will operate under standard team rules. Insofar as possible, decisions will be made by consensus. However, in the event that consensus is not reached, a person at the next level in the organizational structure will always have the authority to make a final and binding decision. In that event, team members will abide by the decision without further discussion.

Document Map

RFP Document Name	PEP Document Name	Responsible
Annual Work Plan	Annual Work Plan	PM
Annual Operations Plan	Annual Operations Plan	PM
Annual Report	Annual Report	PM
Monthly Status Report	Monthly Status Report	PM
Quarterly Status Report	Quarterly Status Report	PM
Earned Value Baseline and Budget	Earned Value Baseline and Budget	PM
Commissioning and Acceptance Test Plan	Deployment and Acceptance Plan	SE
Commissioning and Acceptance Test Report	Deployment and Acceptance Report	OM
Project Execution Plan	Project Execution Plan	PM
Performance Requirements for the CI	System Requirements Document	SE/SA
Requirements Traceability Matrix	System Requirements Document	SE/SA
Signed Interface Agreements with CSO	Interoperability Plan	SE
Signed Interface Agreements with GSO	Interoperability Plan	SE
Signed Interface Agreements with RCO	Interoperability Plan	SE
Requirements Verification Compliance Matrix	Integration and Verification Report	SE
Site Surveys	Not Applicable	
Deficiency List	Deficiency List	SE
OA&M Plan	Operations and Maintenance Plan	PM
Network Verification Test Plan	Integration and Verification Plan	SE
Network Verification Test Report	Integration and Verification Report	SE
Operational Readiness Plan	Operations and Maintenance Plan	PM
Operational Readiness Report	Monthly Status Report	PM
OA&M Procedures	Operations and Maintenance Plan	PM
Requirements for CI Facility Use	Security Management Plan	SE
Integrated Logistic Support Plan	Integrated Logistic Support Plan	SE
Configuration Management Plan	Configuration Management and Change Control Plan	PM
Quality Plan	Quality Plan	PM
Reliability, Maintainability and Availability Plan	System Engineering Plan	SE
Software Development Plan	System Engineering Plan	SE
Interoperability Plan	Interoperability Plan	SE
Interface Control Drawings	Interoperability Plan	SE
Security Plan	Security Management Plan	SE
Science User Requirements	CI Science User Requirements	PS
Risk Management Plan	Risk and Opportunity Management Plan	PM
Transition to Operations Plan	Operations and Maintenance Plan	PM
	Communications Management Plan	PM
	Subsystem Requirements Documents	SA
	DoDAF Documents *(System &Subsystem Specification Documents)*	SA
	Validation Plan	PS
	Validation Report	PS
	E&O Plan	E&OM

A formal decision making process will be invoked when

- the outcome may be of medium to high risk,
- work products under configuration control are changed,
- schedule delays occur,
- costs exceed threshold values set by the Project Manager, and
- the ability to meet project objectives is impacted.

In response, a formal change control process will be implemented (See section 3.7).

2.5 Initial Risks and Constraints

The central risk to the project is failure to deliver a CI with the functionality required for oceanographic research. The areas of highest vulnerability are the elaboration and construction phases, as they are where the major resources are expended, and where missteps may take months to detect. Therefore, it is necessary for the project to be conservative with control over the development process to ensure the ability to execute with precision and agility. Partnership risk will be reduced through use of design and implementation teams that are dominantly based at the prime institution, with partnerships being formed only when the necessary expertise cannot be found in-house. Development partners have been chosen based on a demonstrated ability to perform. Technology providers will be chosen based on the project's ability to capitalize on their proven functionality and for alignment of their technology roadmap with that of the project.

3.0 PLANNING AND CONTROL

3.1 Establishment of Project Planning Parameters

Project planning parameters include all information needed to perform planning, organization, staffing, directing, coordinating, reporting, and budgeting functions. These include requirements imposed by NSF and/or the ORION Program Office; the scope of the project as defined by science user, system requirements, and the system architecture; the spiral project life cycle; the deployment schedule; and build-to-cost constraints. Task and work product identification and their conversion to costs are based on prior experience in related projects.

3.2 Business Management

3.2.1 Earned Value Management System

In accordance with requirements imposed by the ORION Program Office and the Large Facility Project office at NSF, a formal Earned Value Management System will be implemented for the project. This entails development of a performance baseline (PB) that represents the work to be performed along with the required resources and schedule. The project will link its PB into the integrated PB at the ORION Program Office that will be used to generate EVMS reports to NSF. It will also adopt the EVMS software specified by the Program Office. The Project Manager is responsible for implementation of EVMS with the assistance of the project controller.

Because EVMS is not widely used in the academic environment, we have contracted with Triad Project Management to provide project controller services for the initial two and a half years of the effort. Triad has prior experience with project management support of MREFC projects. During the third year of the project, this capability will be transferred to the Project Office with the assistance of Triad personnel.

EVMS was designed for, and has primarily been used with, projects following a waterfall development life cycle. In particular, it is based on the premise that the requirements and high-level architecture can be fully defined prior to construction, and the process to define, implement, and deploy the system can be fully specified. None of these principles are valid under a spiral model. However, the integration of EVMS into the spiral development process has been thoroughly examined by Brownsword and Smith [2005]. They recognized that, prior to LCA at the end of each elaboration phase, most of the effort is oriented at identifying and reducing the system risk. A rational measure of earned value during the inception and elaboration phases must account for the degree to which risk is reduced. In contrast, during the post-LCA construction and transition phases, sufficient design detail exists to permit normal planning of product-oriented work packages, and the standard measures of earned value apply. Brownsword and Smith introduced a risk-oriented performance baseline for use during the early phases of a development cycle. The Project Manager will evaluate this approach at project initiation, and work with the ORION Program Office and the project controller to define an EVMS system that meets both project and program level needs.

3.2.2 Financial Management

The Project Manager is responsible for implementing an accounting system that complies with the requirements for EVMS. The accounting system and its products will be made available for audit as required by the Program Office or NSF.

The Project Manager is also responsible for negotiation of subcontracts with design partners at other academic institutions, federally funded research and development centers (FFRDC) and industry as appropriate, and for procurement of items that exceed the prime institution's policies for a no-bid process. Guidelines will be instituted that delegate procurement below a threshold value to the System Engineer, System Architect, System Development Manager and Operations Manager, and require Deputy Project Director approval for procurement above another threshold level. The Project Manager is responsible for monitoring and controlling all subcontractors, and has the authority to modify or revoke subcontracts as needed.

The Project Manager is responsible for obtaining necessary permits and insurances. Given the scope of the project, it is anticipated that standard prime institution and subcontractor insurance provisioning will suffice, and that no permits are going to be necessary.

3.2.3 Contingency Management

With Deputy Project Director approval, the Project Manager is responsible for managing contingency funds held by the project, applying contingency to mitigate project risk, and replenishing contingency when opportunities are realized. The Project Manager is also responsible for coordination with the ORION Program Office regarding contingency funds held by the OOI Project Director. Contingency funds will be debited or credited only after formal approval of scope or design changes by the project and/or program Change Control Boards, and by the OOI Management Team or NSF if required.

3.3 Work Breakdown Structure and WBS Dictionary

The Work Breakdown Structure (WBS) provides the central scheduling, budgeting, and resource tracking framework for organizing the CI project. It is a hierarchy of elements that decomposes the work plan into a set of work tasks. Conventional WBSs are typically structured around the product design, with top-level headings of requirements, system, subsystem, integration and verification, etc., requiring that the system architecture and its breakdown be defined at the outset. In contrast, an evolutionary WBS organizes the planning elements around a process rather than a product framework, and is better suited for a project whose design and definition proceeds iteratively. A number of structures have been proposed, but the phase-oriented approach suggested by Brownsword and Smith [2005] best meshes with the spiral development life cycle. A phase-oriented WBS has development cycle at the top level, development phase (inception, elaboration, construction, and transition) at the second level, the product design process at the third level, product design activities at the fourth level, and so on. In essence, the full WBS for the six-year effort becomes a set of development spiral WBSs appended to each other. Each development cycle in the WBS is 16-18 months long and overlapped so that a transition phase is completed every 12 months beginning 18 months after project inception. This approach explicitly recognizes that the WBS in successive years is increasingly notional, as the requirements, architecture, and its implementation are in a constant state of evolution. The WBS will be maintained using Microsoft Project. A phase-oriented WBS for the CI project is contained in the WBS Appendix. The WBS Dictionary is an accompanying document that defines the scope of each element in the WBS. Both of these documents will be continuously refined as the project moves forward, and are incorporated by reference into this PEP. With Deputy Project Director approval, the Project Manager is responsible for the WBS and WBS Dictionary.

3.4 Master Project Schedule

The Master Schedule identifies the major engineering phases and milestones along with their success criteria. It time sequences the WBS elements at the top level (development spiral), the second level (spiral phase), and the third level (product design processes). The Master Schedule is established at project inception and is updated only as required. With Deputy Project Director approval, the Project Manager is responsible for preparing the Master Project Schedule.

3.5 Risk and Opportunity Management

Risk and opportunity management is an organized process to identify and categorize situations so that undesirable risks may be mitigated and advantageous situations may be exploited throughout the project life cycle. Risk and opportunity management are tightly coupled with contingency management. Contingency is debited to mitigate high-risk items, and funds saved by opportunity realization are credited back to the contingency pool.

The OOI Risk Management Plan specifies a program-wide process for controlling risk with which the project will be compliant. With Deputy Project Director approval, the Project Manager is responsible for preparation and implementation of the project Risk and Opportunity Management Plan. He will serve as project risk and opportunity manager, and will be a member of the OOI Risk Management Team. Risk and opportunity management will follow the standard process of ranking identified and categorized risks and opportunities, determining their schedule and cost impacts, and developing strategies to mitigate or take advantage of them.

Page 9 of this document contains a preliminary project risk matrix. The Quality Assurance Team plays a pivotal role in identifying and mitigating risks throughout the system life cycle.

3.6 Communications Management

Communications management encompasses the timely and appropriate generation, collection, dissemination, storage, maintenance, and ultimate disposition of program information to relevant stakeholders. It also includes the monitoring of stakeholders' involvement and eliciting their needs, expectations, and constraints for all phases of the project life cycle. With Deputy Project Director approval, the Project Manager is responsible for developing and implementing a Communications Management Plan (CMP) that addresses these items. The CMP contains a list of stakeholders, the information the stakeholders need, the times when the information will be received, who is responsible for information dissemination, the method by which information will flow, and procedures by which commitments by the project team to the communication requirements will be obtained.

Risk Matrix

Risk ID	Title/Assessment	Description	Lead	Status/Plan	Updated
1	Internal partnerships L=1, I=5	Risk from a distributed project team	PM	Mitigate through formal communication plan	12/14/06
2	ORION partnerships L=4, I=5	Risk from a distributed program	PD	Mitigate by maintaining good communication with Program Office and other IOs	12/14/06
3	External partnerships L=3, I=3	Risk from partnerships with non-ORION entities	PM	Mitigate by maintaining close liaison with key entities (e.g., IOOS)	12/14/06
4	Standards organizations L=2, I=4	Risk from changing standards	PM	Mitigate by watching evolving standards	12/14/06
5	Project understaffing L=5, I=5	Risk from underfunding of project	PM	Mitigate by maintaining tight control of cost, schedule and scope	12/14/06
6	Community engagement L=5, I=5	Risk from inadequate community buy-in during the project life cycle	PS	Mitigate by maintaining ongoing community outreach and involving representatives in the system life cycle	12/14/06
7	Requirements redundancy L=1, I=1	Risk from redundant requirements from different domains	SE	Reject	12/14/06
8	Requirements prioritization L=1, I=4	Risk from inadequate prioritization of requirements	SE	Mitigate by maintaining formal requirements review process	12/14/06
9	Subsystem interactions L=1, I=5	Design risk from inadequate management of subsystem interactions	SE	Mitigate by maintaining formal interface control process	12/14/06
10	Policy-based resource management L=2, I=5	Design risk of policy-based resource management element	SA	Mitigate by utilizing proven technologies	12/14/06
11	Instrument network L=2, I=5	Design risk in realizing instrument network capabilities	SA	Mitigate by utilizing proven technologies and early design	12/14/06
12	Streaming data L=2, I=5	Design risk in realizing streaming data capabilities	SA	Mitigate by utilizing proven technologies and early design	12/14/06
13	Absence of instrument standards L=2, I=2	Design risk from lack of instrument standards	SE	Watch	12/14/06
14	Stream processing/storage L=1, I=5	Design risk from stream processing/storage requirement	SA	Watch; utilize proven technologies	12/14/06
15	Converting advanced concepts into reality L=3, I=3	Design risk from advanced concepts in architecture	PM	Watch; use proven development team	12/14/06
16	Technology choice L=3, I=3	Technology risk when multiple technologies have the same functionality	SA	Watch	12/14/06

Likelihood					
5					5, 6
4		4			2
3			3, 15, 16		
2		13			10, 11, 12
1	7			8	1, 9, 14
Risk Summary Matrix	1	2	3	4	5
	Impact				

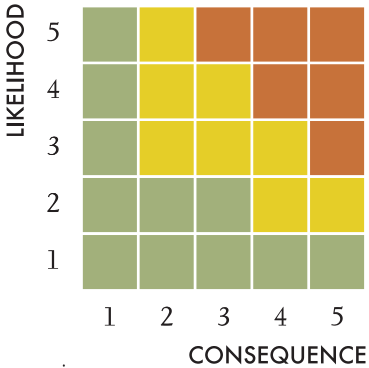

LIKELIHOOD					
5	🟩	🟨	🟧	🟧	🟧
4	🟩	🟨	🟨	🟧	🟧
3	🟩	🟨	🟨	🟨	🟧
2	🟩	🟩	🟩	🟨	🟨
1	🟩	🟩	🟩	🟩	🟩
	1	2	3	4	5
	CONSEQUENCE				

Likelihood (L): *L=1, very unlikely, < 1%; L=2, unlikely, 1-30%; L=3, moderate, 30-50%; L=4, high, 50-70%; L=5, almost certain*

Impact (I): *I=1, minimal or no impact to project success; I=2, minor impact to project success, but can be accommodated with established reserves; I=3, moderate impact to project success, but can be handled within established reserves; I=4, major impact to project success, but threatens established reserves; I=5, cannot achieve project success, no alternatives exist, exceeds available reserves*

Classifications of risks: *Green, not a project threat; Yellow, potential threat to project success; Orange, significant and likely threat to project success*

Status: *Watch; Research; Accept or reject; Mitigate; Retired*

3.7 Configuration and Change Control Management

Proper configuration and change control management is a key element in maintaining the reliability and quality of complex systems. A formal project-level configuration management and change control process will be implemented for all project, system, and sub-system level plans and requirements, interface control, and design documents. This process will be compliant with the program level configuration and change control management procedures defined in the OOI Configuration and Change Control Management Plan. As an element of configuration management, the System Engineer administers the Deficiency List, a list of all identified defects and information about their resolution.

With Deputy Project Director approval, the Project Manager is responsible for the preparation and implementation of a project Configuration Management and Change Control Plan (CMCCP). The CMCCP provides security assurance through control of changes made to documents, hardware, software, and firmware throughout the project life cycle. Source code management (revision control) is an integral part of configuration management. The CMCCP describes the documents and software components that are maintained under configuration control, specifies the custodian for each element through whom proposed changes must be routed, and defines the change approval process.

The CMCCP establishes a Change Control Board that makes decisions on whether proposed changes may be implemented. The Change Control Board will comprise the Project Scientist and E&O Manager (who represent the stakeholders), the System Engineer, System Architect and System Development Manager (who represent the system), the Operations Manager (who ensures that changes have minimal impact on operations and maintenance), and the Project Manager, who chairs the Board and will make a final and binding decision on subsystem level changes in the event that consensus is not reached. System-level changes are approved by the Deputy Project Director, who will make a final and binding decision in the event that the Board cannot reach consensus. Proposed changes may also be submitted to the OOI Change Control Board and NSF as specified in the OOI Configuration and Change Control Management Plan, depending on the impact of the change.

3.8 Integrated Logistics Support

Integrated Logistics Support (ILS) defines all of the elements required to support the system throughout its life cycle. It is usually divided into ten components:

- Maintenance planning.
- Supply support.
- Test equipment/equipment support.
- Manpower and personnel.
- Training and training support.
- Technical data.
- Computer resources support.
- Facilities.
- Packaging, handling, storage, and transportation (PHS&T).
- Design interface.

All components of ILS must be developed in coordination with each other and with system engineering. Tradeoffs may be required between them to minimize life cycle cost, which is a principal design goal.

With Project Manager approval, the System Engineer will prepare the Integrated Logistics Support Plan (ILSP) that defines the ILS process. The ILS process must begin at project inception to ensure that life cycle costs are minimized. The System Engineer and System Architect are responsible for implementing the plan.

3.9 Education and Outreach

With Deputy Project Director approval, the E&O Manager will develop and implement the E&O Plan that provides a roadmap for E&O activities during the project life cycle. The E&O Plan will be compliant with the OOI E&O Plan.

4.0 PROJECT EXECUTION

4.1 Work Plan and Reports

With the assistance of other project team members as required, the Project Manager is responsible for completion of an Annual Work Plan that:

- defines the next year's engineering goals and activities;
- provides schedules for design, construction, integration verification, validation, and deployment;
- states the required budgets and resources to accomplish the goals; and
- identifies major planning activities and milestones.

The Annual Work Plan is used to modify the scope, schedule, and cost baselines, and hence define the annual performance baseline. The Project Manager is also responsible for preparation of an Annual Report that:

- gives the key accomplishments in the prior year;
- provides a comprehensive financial report;
- states project changes such as adjustments to schedule, contingency usage, or cost variance; and
- identifies risk status.

Both the Annual Work Plan and Annual Report are approved by the Project Director and submitted to the Program Office for final approval.

The Project Manager will submit Monthly and Quarterly Status Reports to the Program Office. Both of these reports will document major accomplishments and project changes, and the quarterly report will also include a financial report. These reports are approved by the Deputy Project Director.

4.2 Detailed Project Schedule

The Detailed Project Schedule provides a calendar view to support the events in the Master Schedule (see exhibits) and Annual Work Plan, and hence is produced before beginning a development spiral. It expands the Master Schedule from the fourth WBS level (product design activities) downward. With Deputy Project Director approval, the Project Manager is responsible for preparing the Detailed Project Schedule.

4.3 Performance Baseline

The performance baseline defines the performance capabilities required to meet the mission, and is used for Earned Value Management. It comprises three elements: scope, schedule, and cost baselines. With Deputy Project Director approval, the Project Manager is responsible for preparing and maintaining the performance baseline, which is a key input to the EVMS.

The Science User Requirements and the System Requirements circumscribe the intended purpose of the CI. The System Architecture defines the functionality of the CI. Additional scope modifications are contained in the Annual Work Plan. Together, these elements provide the scope baseline.

The schedule baseline is the Detailed Project Schedule.

The cost baseline defines the total cost of providing the necessary capabilities in the Detailed Project Schedule. The Annual Work Plan defines the annual cost baseline. The initial cost baseline is defined in Volume Two of the proposal.

4.4 System Engineering

A comprehensive system engineering effort lies at the heart of any complex system, especially one that is software-intensive. The OOI certainly fits that description, and the respondents will provide an integrated System Engineering and System Architecture team to carry out the system engineering and architecting processes throughout the system life cycle.

The system engineering framework used by the project will be a tailored version of that defined in the System Engineering Handbook, Version 3 (SEH) issued by the International Council on System Engineering (www.incose.org). DoDAF www.defenselink.mil/nii/doc/DoDAF_v1_Volume_I.pdf) defines a common approach for software architecture description, development, presentation, and integration that is especially suitable for systems that are implemented in stages. DoDAF was used to describe the ORION Conceptual Architecture, and will continue to be used throughout the system life cycle.

Fulfilling the anchor point milestones in the spiral model is a key system engineering responsibility. Use of the spiral management model does not alter the function of system engineering, and in fact it becomes the key activity that binds the cyclically-growing system into a coherent whole.

4.4.1 System Engineering Plan

The System Engineering Plan (SEP) encompasses the Software Development Plan (SDP), and serves as the roadmap for developing and delivering the CI, including its hardware elements. The SEP addresses a series of questions regarding the system deliverables:

- What system will be delivered?

- What tasks must be accomplished to deliver it?
- When must each task be started and finished?
- What is the order in which the tasks must be completed?
- What are the task dependencies?
- What are the final acceptance criteria?
- Who will be responsible for each task?
- How will each task be carried out?

The SEP describes all stages in the system life cycle from requirements definition through integration to deployment. The SEP will include reliability, maintainability and availability criteria, and subsumes the Reliability, Maintainability and Availability Plan. The SEP is incorporated into this PEP by reference. A comprehensive template for the SEP is given in SEH, and will be tailored to include all elements of the SDP. With Project Manager approval, the System Engineer is responsible for developing and implementing the SEP. The SEP will be updated at least annually, and can be expected to evolve through successive development spirals.

4.4.2 Interoperability Management

The Interoperability Plan describes procedures that will ensure the interoperability of the CI both internally and with the hardware and software elements produced by the three observatory IOs and key external entities, notably IOOS. It includes Interface Control Documents (ICDs) that describe the interfaces between all system and subsystem elements internal to the CI, between the CI and the three OOI observatories, and between the CI and external entities. ICDs internal to the CI are negotiated between the System Development Manager and the System Engineer/System Architect, who have approval authority. The Interoperability Plan also includes Interface Agreements (IAs) negotiated between the project and the three observatory IOs that establish interface requirements, constraints, and milestones. The IAs incorporate relevant ICDs, and are negotiated between, and signed by, the cognizant IO system engineers subject to approval by the cognizant IO project managers. The Program Office has final approval authority, and resolves any conflicts that may arise. Finally, IAs may be negotiated between the project and external entities under similar conditions. With Project Manager approval, the System Engineer is responsible for developing and implementing the Interoperability Plan.

4.4.3 Integration and Verification Management

Integration and Verification Plan (IVP) establishes sequences and schedules for integration of the subsystems with each other, with the existing CI, and with the observatory and external elements at successive development spirals. It also establishes criteria to verify the system by asking "was the system built right?" through establishing that the system requirements have been met. ISO 9126 serves as a framework of verification attributes and criteria. The IVP subsumes the Test Plan (with the exception of the elements in the Validation Plan), and is the responsibility of the System Engineer with Project Manager approval. The IVP is implemented by the Software Development Manager and System Architect for internal CI elements and by the System Engineer for final integration and verification. At the end of each integration and verification phase, the System Engineer will submit an Integration and Verification Report to the Project Manager that includes a Requirements Verification Compliance Matrix.

4.4.4 Concept of Operations and System Requirements Documents

With Project Manager approval, the System Engineer will be responsible for the Concept of Operations and System Requirements Documents (SRD). A Requirements Traceability Matrix will link the SRD to the science user requirements. The system requirements are divided into four major categories (functional requirements, performance requirements, design principles, and interface requirements), and then further sorted by the System Architect into categories that are consistent with the CI architecture and its subsystems to yield Subsystem Requirements Documents. The SRD serves as the top level description of desired CI capabilities, and is a key guide to the software developers as the project moves forward.

4.4.5 DoDAF System Architecture Documents

The system architecture will be specified using the DoDAF framework. With Project Manager approval, all system and subsystem architecture documents will be the responsibility of the System Architect. This document set replaces the usual System and Subsystem Specifications Documents.

4.4.6 User Documentation

The System Development Manager is responsible to the System Architect and System Engineer for the production of user documentation for the CI. He may be assisted by selected members of the Subsystem Development Teams for this activity.

4.5 Validation Management

The Validation Plan establishes criteria to validate the system by asking "was the right system built?" It must include evaluation of the system in the context of the use scenarios that help define the science user requirements, and serves as final stakeholder acceptance of the CI at each deployment. With Deputy Project Director approval, the Project Scientist is responsible for development and implementation of the Verification Plan with support from the System Engineer. At the completion of each validation, the Project Scientist will submit a Validation Report to the Deputy Project Director.

4.6 Deployment and Acceptance Management

With Project Manager approval and Operations Manager assistance, the System Engineer is responsible for developing a Deployment and Acceptance Plan that defines the process for CI deployment at the end of each development spiral and the criteria for its acceptance by the OOI Program Office. The Plan will also describe the documentation that allows the system to make the transition to operations and specify the training required for operations personnel. The acceptance process will comply with requirements imposed by, and will be overseen by, the ORION Program Office, which has ultimate responsibility for accepting the CI. The deployment and acceptance process follows on the integration, verification, and validation processes described in the Integration and Verification Plan and Validation Plan, respectively. Deployment will be carried out by the Operations Team and the System Integration Team, with oversight by the System Engineer and Project Manager. The Operations Manager is responsible for delivering a deployed system to the Project Manager at the end of each development spiral, and will prepare a Deployment and Acceptance Report for submission to the Program Office with Project Manager approval after each development spiral.

4.7 Quality Management

With Deputy Project Director approval and in coordination with the OOI Project Director, the Project Manager is responsible for preparing a Quality Management Plan. He is also responsible for staffing a qualified Quality Assurance Team that will audit the engineering deliverables and oversee the quality assurance and quality control process throughout the system life cycle. The Quality Assurance Engineer is responsible for implementing the Quality Management Plan.

4.8 Security Management

With Project Manager approval, the System Engineer and System Architect are responsible for preparing and implementing the Security Management Plan that covers all aspects of operational and CI security for the system, including defining the software and hardware "best practices" (e.g., firewalls, one-time passwords, anti-virus software) that will be used to protect against intrusion on a real-time basis and the processes used to define and manage reportable incidents both within ORION and at the federal level. It will also describe the authorization and auditing policies for the CI at different levels of access and the ongoing process for ensuring that repositories remain free from external aggression. Compliance with national security requirements will also be described. The Security Management Plan will incorporate any additional requirements imposed by the ORION Program Office and NSF.

4.9 Reviews

The Project Manager is responsible for organizing, documenting, and reporting all internal project reviews (e.g., LCO, LCA, and IOC during each development spiral), and for incorporating the findings from the reviews into the ORION development process. With Deputy Project Director approval, all reports will be submitted to the Program Office. The project will also support program-level design reviews. The Project Manager is the principal point of contact for that purpose.

4.10 Operations and Maintenance Management

Operations and maintenance (O&M) management is the process that governs post-deployment operations and maintenance of the CI. It is governed by the O&M Plan that is effectively a project execution plan for the post-deployment phase extending to the end of the ORION life cycle. Some of its elements are constrained by the Integrated Logistics Supply Plan. The Annual Operations Plan is an annual work plan for operations and maintenance. With Deputy Project Director approval and support from the Operations Manager, the Project Manager will prepare the O&M Plan and the Annual Operations Plan. The Operations Manager is responsible for implementing both plans.

Appendix 6.2-Task Area Experience

ORION CI Proposal

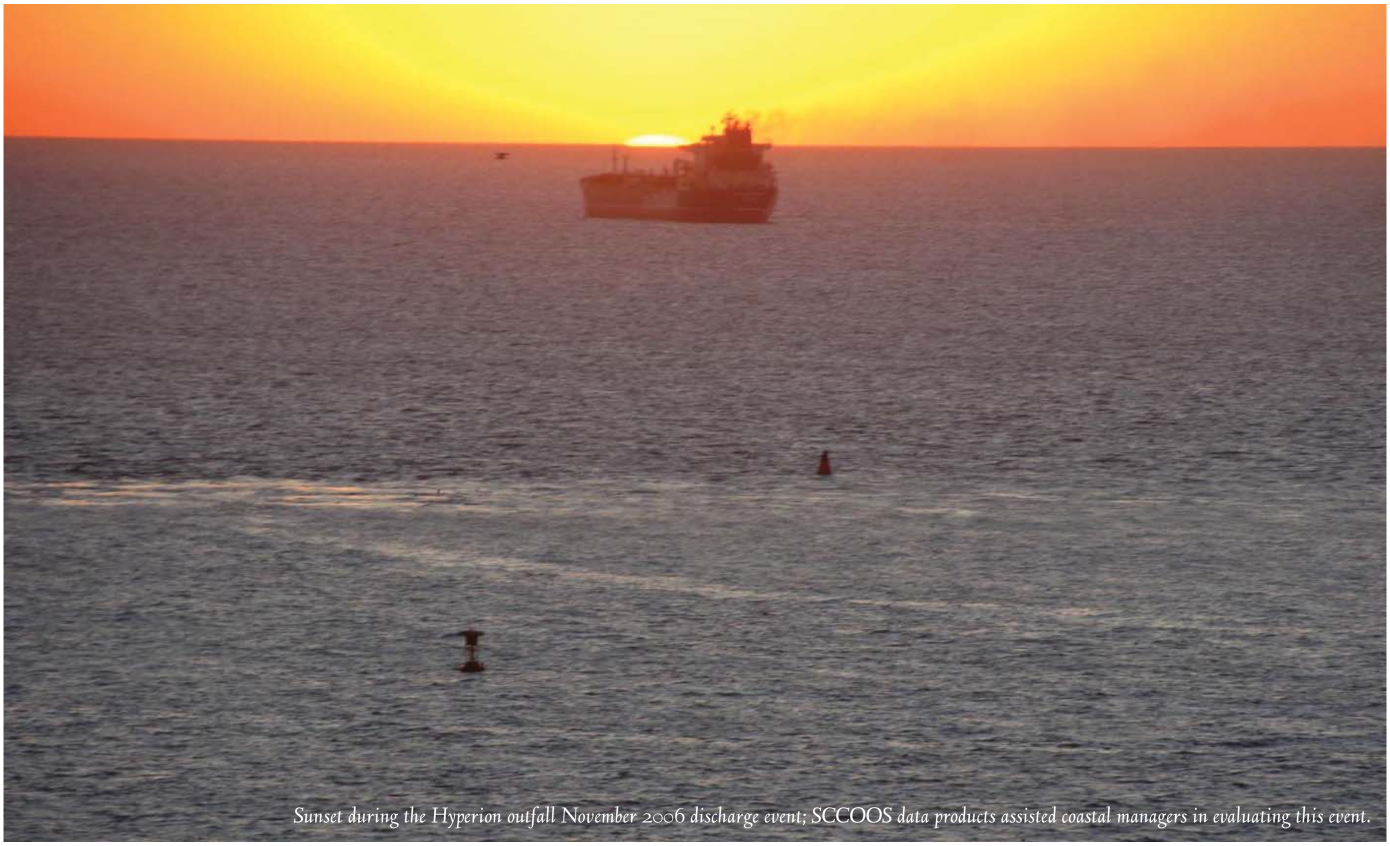

Sunset during the Hyperion outfall November 2006 discharge event; SCCOOS data products assisted coastal managers in evaluating this event.

6.2.1 SOUTHERN CALIFORNIA COASTAL OCEAN OBSERVING SYSTEM - EXPERIENCE IN ALL SIX TASK AREAS

COTS (Coastal Observation Technology Systems)

PIs: Orcutt, Davis and Terrill

Title: Southern California Coastal Ocean Observing System (SCCOOS): Shelf to Shoreline Observatory Development

Award Number: NA17RJ1231

Amount: $1,405,394

Period: 07/01/06-06/30/07

STATE CONSERVANCY

PIs: Terrill, Orcutt and Davis

Title: Coastal Ocean Currents Monitoring Program (COCMP)

Award Number: 04-078

Amount: $11,443,566 (Year Two 11/05 to 11/06-$2,970,873)

Period: 11/15/04-11/15/07

Project Management

SCCOOS is a consortium of eleven Southern California and Mexican universities, including California Polytechnic State University, San Luis Obispo (CalPoly); the University of California campuses at Santa Barbara (UCSB), Los Angeles (UCLA) and Irvine (UCI); the Jet Propulsion Laboratory; the University of Southern California; Cal State Los Angeles; the Southern California Coastal Water Research Project (SCCWRP); Scripps Institution of Oceanography; the Universidad Autonoma Baja California (UABC); and Centro de Investigacion Cientifica y de Educacion Superior de Ensenada (CICESE). Each institution has a history of coastal observing, monitoring, and modeling, and a reputation for developing novel environmental sensors, platforms, and data management techniques. The consortium (www.sccoos.org), which extends from Northern Baja CA in Mexico to Morro Bay at the southern edge of central California, aims to streamline, coordinate, and further develop individual institutional efforts to create an integrated, multidisciplinary, coastal observatory for the Southern California

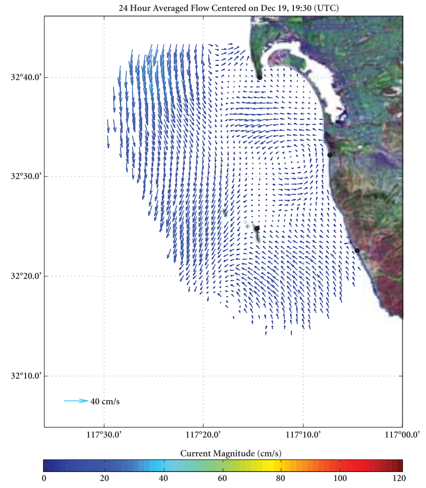

24-hr average of surface current vectors.

Bight. SCCOOS is committed to leveraging current infrastructure, partnerships, and other resources to develop a fully operational coastal observing system to address a variety of policy, monitoring, and management needs. SCCOOS is a leading Integrated Ocean Observing System (IOOS) Regional Association (RA).

The SCCOOS organization structure is based on a Memorandum of Understanding between consortium members, and is directed by a Board of Governors comprising an institution head or other person of administrative stature from each member institution. John Orcutt is the Chair of the Board of Governors. The fiscal function of the corporation is provided by the business office of the Marine Physical Laboratory and Joint Institute for Marine Observations (JIMO) located at Scripps Institution of Oceanography. Consortium members collaborate to pursue the interests of SCCOOS through contracts, grants, or other mutual agreements between their home organizations. Terms, liability, and management structures are defined by written agreement.

An Executive Steering Committee, appointed by the Board of Governors, works closely with a Chief Operations Officer (currently Eric Terrill) to provide advice on technical issues and strategic planning. The Chief Operations Officer provides day-to-day management of SCCOOS and acts as general manager for all grant and technical matters. An external Senior Advisory Committee provides external insight and perspective on technical, market, legislative, and political matters that may affect the observing system. They also serve as an external source of information and reference that links SCCOOS with broad stakeholder interests and knowledge within the region.

Systems Engineering

The technical elements of SCCOOS operate as a system-of-systems, with common architectures and protocols developed to link together existing data gathering or data management activities. SCCOOS has incorporated legacy observing systems, and is providing incremental improvement to those systems in a manner that does not degrade the ongoing collection of data. In many ways, this reflects the IOOS challenge of linking disparate Regional Association observatory instruments into a common enterprise "bus" capable of providing a Service Oriented Architecture to eternal users.

Due to the stakeholder/customer engagement approach used in SCCOOS activities, new system engineering activities are typically approached in a spiral development mode that allows build-test-build involving end users of the technical components of the observing system, and in particular, the data management/information technology back end. This allows a cyclical approach in which customers evaluate early results, and in-house engineers identify potential trouble spots at an early stage.

Examples of engineered systems include HF radar networks using ROADNet and Antelope/Datascope, mooring networks, automated pier sensor networks, glider networks, and data assimilating models.

Software Development

Software development takes place using

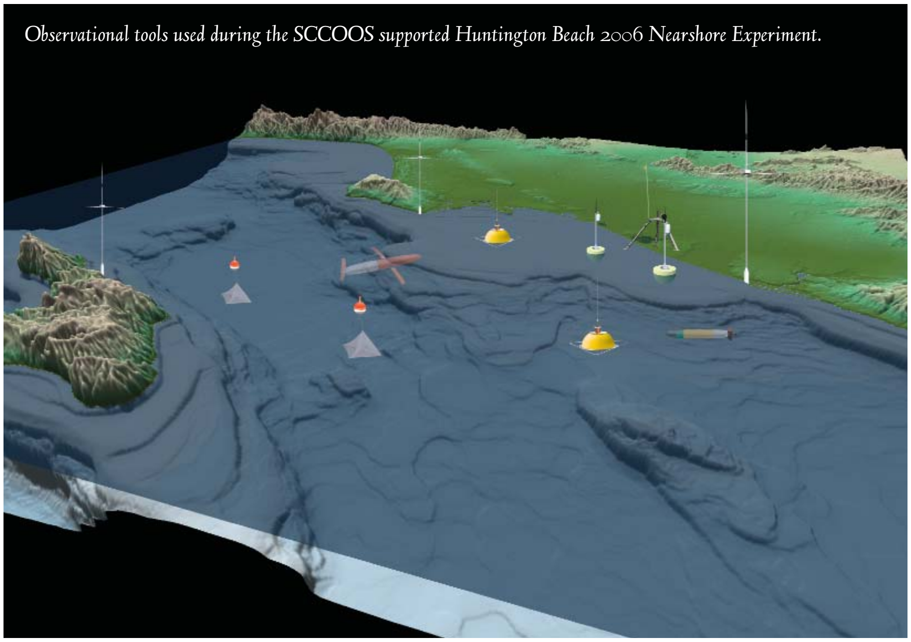

Observational tools used during the SCCOOS supported Huntington Beach 2006 Nearshore Experiment.

community best practices. Whenever possible, open sources are used to allow the greatest flexibility in programming. Data bases employed in the SCCOOS data management system include SQL/MySQL, Oracle, and Antelope/Datascope. Storage systems are both redundant RAID arrays and the SDSC Storage Resource Broker. The programming languages utilized for data manipulation and transport include C/C++, Java/Java Script, PhP, Perl, Matlab, JpGraph, and Google Earth/Maps. Programming teams are engaged on a per-project basis. Successful integration is exemplified by an excellent web portal, which includes ready access to archived data.

Examples of developed software include the data storage, communications, and delivery system for a real-time network of HF radar, telemetering ocean buoys, pier sensors, and ocean glider systems. SCCOOS has also established data ingestion, archiving, and visualization tools for a network of over 400 coastal meteorology sensors, shoreline bacteria water quality data, and CTD cast data gathered throughout Southern California by a network of small boats. SCCOOS works closely with operational Navy meteorology groups for access, delivery, and evaluation of coupled ocean and atmospheric forecasting models (COAMPS). SCCOOS has also established a near real-time data delivery system for remote sensing data gathered by both U.S. and international satellites that have ocean color and SST sensors.

Hardware Development

Hardware development is supported by engineers and technicians employed by SCCOOS. Examples include ocean mooring hardware and associated *in situ* sensors, pier-based sensor systems, fabrication of gliders, and wireless networking hardware for providing coastal connectivity to real-time sensors. SCCOOS also maintains the back-end computing systems necessary for the management, storage, and delivery of data.

Implementation, Operations and Maintenance

SCCOOS manages and supports the engineering and technical staff that implement, operate, and maintain the sensor networks that comprise the system. A description of data gathering activities can be found at www.sccoos.org/interactive-map. Operations and maintenance typically take place during a Monday to Friday, nine-to-five work period, principally due to funding constraints. On a case-by-case basis (e.g., this year's large Huntington Beach experiment—details may be accessed online: www.sccoos.org/projects/hb06/instrumentMap.php), after hour maintenance and operations do occur. All software is designed to work on an automated basis, with real-time data operations occurring on a round-the-clock basis. SCCOOS operations have achieved very high reliability while supported on a work-day schedule. This experience translates directly into ORION's need for essentially 100% data return.

The Coastal Observing R&D Center at Scripps Institution of Oceanography (www.cordc.ucsd.edu) staffs the principal information technology functions for SCCOOS. The center consists of more than a dozen highly trained engineers, programmers, staff research associates, and technicians, who develop, test, operate, and implement all facets of a coastal observing system. The Center provides these services to agencies including NOAA/IOOS, Navy/ONR, National Science Foundation, and the State of California. This experience in providing deliverables to both operational and scientific communities has allowed SCCOOS to effectively deliver data and informational products to its end-users. SCCOOS can serve as a conduit between ORION's research network and the more operationally focused Integrated Ocean Observing System, which will be principally managed by NOAA. A demonstrated example of this conduit is the data management system designed by CORDC to manage a national network of HF radars. The system, which manages real-time data from close to 60 different sites, is used by SCCOOS scientists, local coastal managers, and the NOAA National Data Buoy Center.


NSF Earthscope USArray telemetry station currently deployed in northeastern Oregon

6.2.2 US ARRAY: ARRAY NETWORK FACILITY–EXPERIENCE IN PROJECT MANAGEMENT

PI: Vernon

Title: UCSD Proposal for the USArray Array Network Facility

Award Number: IRIS Subaward 472 (MRE) and 479 (O&M)

Amount of Award: $3,805,009 (MRE) and $1,435,763 (O&M)

Period of Award: 10/1/04-9/30/09

The USArray Array Network Facility (ANF, anf.ucsd.edu) is one of the critical elements that has made the NSF MRE EarthScope USArray a success. The ANF provides the cyberinfrastructure to guarantee that data from the USArray Transportable Array (TA) and telemetered Flexible Array (FA) stations are delivered to the IRIS Data Management Center (DMC) for permanent archiving and distribution to the user community. The ANF provides measures of quality control for all data and ensures that the proper calibration and metadata are always up to date and available. There are many types of monitoring that need to be accomplished, including seismic data quality, IP network communications, and higher-level data communications. The ANF interacts closely with the Array Operations Facility (AOF) and the TA Field Operations Contractor providing immediate feedback on station data quality.

The ANF is a major undertaking given that there will be 400 TA stations and up to 200 FA stations telemetering data in real-time when USArray completes its build-out. This system is larger in scale than any existing digital telemetry network. As of 1 November 2006, there are 276 TA stations on line, and the TA will be fully deployed by 1 October 2007. The main activity of the ANF for its first three years was development of the infrastructure to manage data and QA obligations for full implementation. There currently are enough computing resources at the IGPP Broadband Data Collection Center to manage USArray operations well into Year Five. The ANF has demonstrated the system is scalable; a result important to the OOI.

The PI, Dr. Frank Vernon, manages and is responsible for all aspects of the ANF. He has extensive experience developing, deploying, and operating real-time seismic networks and arrays. He is responsible for interactions with IRIS management and for interactions with UCSD. He is assisted by the following key personnel:

- Specialist Dr. Luciana Astiz is a seismologist with significant field and network experience (she was responsible for deploying much of the UN International Monitoring System while an employee of the UN in Vienna), and assists the PI in ANF operations and troubleshooting station problems. Dr. Astiz is responsible for final data review as well as evaluation of seismic noise and calibration data, and supervises the Data Analysts. Dr. Astiz is the primary point of contact to the AOF.

- Staff Research Associate Jennifer Eakins is in charge of network quality assurance, data distribution, and metadata maintenance and delivery. Ms. Eakins is the primary point of contact with the IRIS DMC.
- Programmer/Analyst Rob Newman is responsible for the WWW views of the project.
- Staff Research Associate Dr. Vladik Martynov is the lead data analyst.
- System Administrators Steve Foley and Geoff Davis maintain the Sun computers and networking hardware/software.
- Principal Scientist Arcot Rajasekar is the PI for the SDSC Storage Resource Broker (SRB), and assists in implementing the SRB for data backup.
- Principal Publications Coordinator Jennifer Matthews assists in document and figure preparation for internal and external information dissemination, and responds to requests from IRIS for the generation of figures for reports and presentations.

All project personnel attend a weekly meeting, and all personnel involved in the IGPP Broadband Data Collection Center contribute to prognostic weekly e-mail summaries that are distributed to the entire project. The PI and Key Personnel meet monthly to review program status and carry out strategic planning.

6.2.3 PLUSNET: PERSISTENT LITTORAL UNDERSEA SURVEILLANCE NETWORK– EXPERIENCE IN SYSTEM ENGINEERING

PI: Schmidt, Co-PIs: Baggeroer, Leonard, Rus, Makris, Battle, Stojanovic

Title: PlusNet; Persistent Littoral Undersea Surveillance Network

Award Number: S05-06

Amount of Award: $5,110,000

Period of Award: 1/1/05-12/31/07

MIT is leading a multi-institutional, multi-disciplinary research effort funded by the Office of Naval Research aimed at developing PLUSNet, a new autonomous, distributed acoustic sensing network concept for persistent undersea surveillance in the littoral ocean.

The PLUSNet concept is based on a synergy of new acoustic sensing technology, advanced signal processing, ocean modeling and assimilation, marine robotics, underwater communication and navigation, and artificial intelligence and automation, all integrated into an operating system aimed at providing persistent surveillance capability in the ocean for periods of several months, without being dependent on centralized, human control. By seamlessly integrating a wide range of new and developing technologies, PLUSNet provides a unique example of the system engineering challenges facing the development of a net-centric, distributed and interactive ocean observation system, taking advantage of the latest technological advances in a wide range of traditionally independent disciplines.

A major obstacle to the development of distributed sensing networks in a littoral ocean environment is the inherent limitation of underwater acoustic communications in terms of bandwidth, intermittency and latency, typically limiting communication to a few tens of bytes per minute, with intermittency and latency of order tens of minutes. To enable robust operation of the network, MIT has developed and recently implemented a comprehensive Autonomous Communication, Command and Control (AC3) architecture, inherently capable of supporting nested autonomy, where individual nodes and clusters of nodes can complete adaptive and collaborative sensing missions without operator control and intervention. Based on the open architecture of the MOOS-IvP robotic network control framework developed in a collaborative effort by MIT, Oxford University, and Naval Undersea Warfare Center, the PLUSNet architecture provides a generic framework for nested and autonomous, fully integrated sensing, modeling and control, which is directly applicable to adaptive and collaborative sensing in other ocean observation networks with fixed and mobile assets, such as ORION.

The MOOS-IvP framework (Mission-Oriented Operation Suite—MOOS, not to be confused with the MBARI Monterey Ocean Observing System; Interval Programming—IvP) was successfully implemented and demonstrated in the recent MB'06 experiment in Monterey Bay. The goals were to (a) demonstrate network connectivity and the autonomous, adaptive tracking of an acoustic target by an autonomous underwater vehicle (AUV) equipped with a 100 m long towed acoustic array, and (b) conduct adaptive environmental sampling using a CTD. With the principal objective being detection, localization, and tracking of a fast moving target, particular considerations were given to proper control while towing the array. Thus, mission-related objectives were tempered by the need to consider the effect of a sequence of maneuvers on the motion of the towed array, which cannot tolerate sharp bends or twists. A key to the robust satisfaction of such multiple objectives—characteristic of most autonomous operations - is the behavior-based control model of MOOS augmented with a novel interval programming (IvP) approach for performing behavior coordination using multi-objective optimization.

MOOS is an open source software suite for robotic control, in which all communication between processes and tasks is handled via a central database, the MOOSDB. Processes for sensing, processing, and control can be operated at different cycle frequencies. PHelmIvP is the core process managing adaptive and collaborative, multi-objective, behavior-based control. It uses computationally efficient IvP for choosing between conflicting behaviors, based on a set of cost functions (IvPFunction) mapped over the action parameter space by each individual behavior. MOOS-IvP is ideally suited for highly autonomous systems with limited, latent, and intermittent communication, and with specific application to autonomous detection, localization, and tracking of oceanographic events such as upwelling or internal wave packets.

Although the principal objective of PLUSNet is the detection, classification, localization, and tracking of submerged targets such as submarines, in a coastal environment, the MOOS-IvP AC3 framework lends itself directly to the autonomous and adaptive capture of the episodic oceanographic events that are among the core new capabilities provided by ORION. Thus, PLUSNet itself is exploiting the oceanographic environment to enhance performance. For example, several of the standard PLUSNet behaviors are directly aimed at measuring the sound speed distribution, with particular focus on episodic events such as internal wave solitons, upwelling, and fronts, all of which are critical to the acoustic performance of both sensing and acoustic communication.

It is envisioned that ORION nodes in the Pioneer Arrays, the cabled regional and coastal observatories and large global buoys will at some point be equipped with docking units for one or more AUVs, significantly extending their spatial coverage. Through adaptive sampling and behaviors, they will enable rapid response to episodic oceanographic events similar to those addressed in PLUSNet. We envision a single docking unit serving as a re-charging station for several AUVs deployed around the node in a hibernation mode. Once an event is detected or forecast, scientists using the vehicles can activate one or more of the dormant AUVs using simple acoustically-communicated commands similar to those used in PLUSNet.

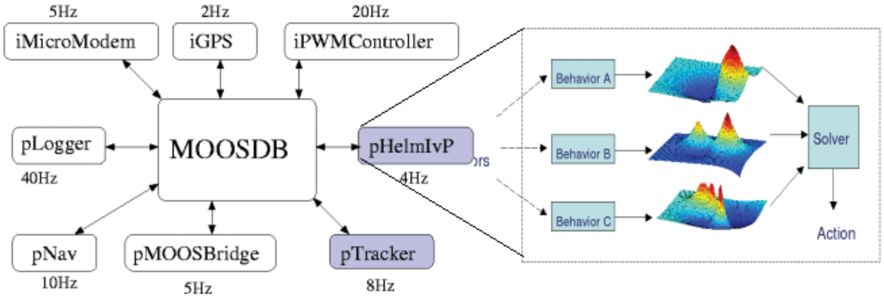

MOOS-IvP behavior-based command and control framework

Using the MOOS-IvP architecture, this control can be utilized by the scientific community without extensive training through a dedicated GUI built on the same principles as the PLUSNet field control and Command and Control Language (CCL). This has ample flexibility to incorporate 10-20 different classes of adaptive and non-adaptive oceanographic mapping behaviors and event capture missions. At the same time, more advanced users will be able to develop their own defined behaviors using the standard MOOS-IvP toolset. In support of this, the ORION CI team will work with the scientific community in identifying and developing the advanced behavior suite for oceanographic event capture, as well as a suite of safety behaviors which will eliminate or reduce the risk of damaging the mobile assets and the fixed node facilities. Further, the implementation team, and later the O&M team, will work with all manufacturers and operators of mobile ORION assets on integrating with the MOOS-IvP command and control infrastructure.

With its modular structure, MOOS-IvP lends itself directly to simulating all missions prior to field deployment. As is currently the case for PLUSNet, all ORION researchers intending to use the interactive, mobile assets must demonstrate the safe execution of their mission plans in a MOOS-IvP simulation environment, where the assets are ‘operated’ in ORION’s ‘virtual ocean.’

6.2.4 SRB/IROD: EXPERIENCE IN SOFTWARE DEVELOPMENT

PIs: Moore, Rajasekar and Wan

Title: Constraint-based Knowledge Systems for Grids, Digital Libraries, and Persistent Archives

Award Number: NSF IIS-0427196

Amount of Award: $714,000

Period of Award: 09/01/04-08/31/06

PIs: Berman, Moore

Title: Supplement to Delivering Cyberinfrastructure: NARA Intercontinental Archive Testbed

Award Number: NSF NPACI ACI-9619020 (NARA supplement)

Amount of Award: $1,850,000

Period of Award: 10/1/05-9/31/07

PIs: Orcutt, Vernon, Rajasekar, Braun, Ludaescher

Title: NSF/ITR: Exploring the Environment in Time: Wireless Networks & Real-Time Management

Award Number: NSF OCE-0121726

Amount of Award: $1,758,078

Period of Award: 10/1/01-9/30/04

PIs: Orcutt, Vernon, Rajasekar, Ludaescher

Title: NSF/ITR: Real-Time Data Aware System for Earth, Oceanographic, and Environmental Applications

Award Number: NSF ATM 03-25963

Amount of Award: $2,344,407

Period of Award: 9/1/03-8/30/06

UCSD has extensive experience in the management of real-time sensor data streams (ROADNet), the implementation of data grids for managing distributed data collections (Storage Resource Broker), and the processing of massive data sets (Teragrid). These systems integrate software technology from multiple communities. The technology on which the ORION cyberinfrastructure will be initially assembled is well understood, has been successfully integrated in multiple production systems, and provides a platform for the archiving of sensor data. Specific instances of the application of these technologies include:

Virtual Object Ring Buffer:

The NSF-funded ROADNet project has developed a Real-time Data Grid (RTDG) based on the concept of a Virtual Object Ring Buffer (VORB) to ease the discovery, use, and access of distributed multi-sensor data. The ROADNet system was based on the concepts of infrastructure independence provided by data Grids and sensor protocol abstraction provided by object ring buffers (ORB). VORB provides a virtual sensor name space with attribute-based discovery and access. It also provides an easy way to perform server-side sensor processing. Currently, the ROADnet system is operational with more than 4000 data streams from multiple-disciplines, including seismic sensors (over 1000 locations from all seven continents and most islands), environmental sensors (temperature, pressure, wind), image streams (on ships, from ecological reserves, etc), oceanographic data, ocean wave monitoring using HFCR, etc. These data streams are accessed from more than 100 ORB servers running on more than 70 distributed hosts. ROADNet's VORB system is also being used to manage the NSF EarthScope USArray Transportable and Flexible Array real-time data systems and the Surface Currents Mapping Initiative.

Storage Resource Broker

Under multi-agency funding support, SDSC has developed leading-edge data Grid middleware, the Storage Resource Broker (SRB), and applied the technology in support of scientific discipline collections, education digital libraries, and preservation environments. The SRB is federated client-server middleware that implements data virtualization through the insertion of a shared collection management layer between the user and distributed storage resources. Logical name spaces are used to identify users (global user identity), storage resources, and dig-

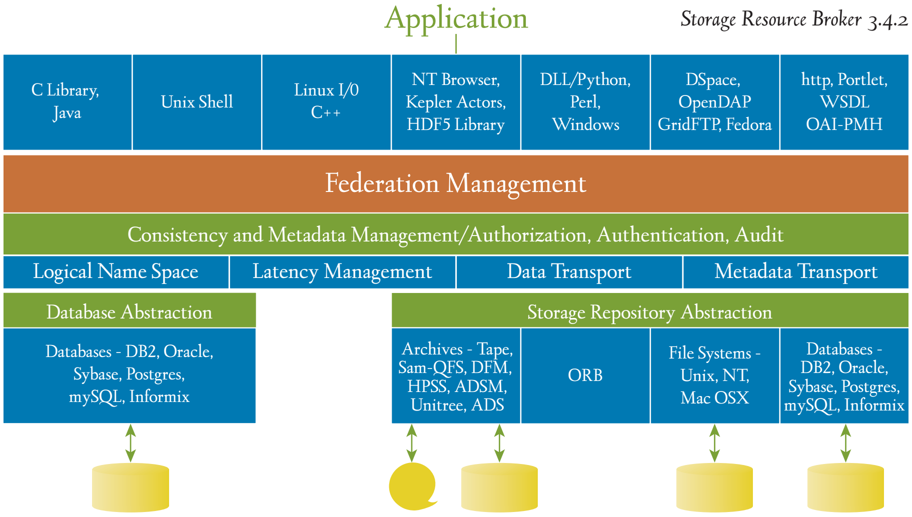

Storage Resource Broker 3.4.2

ital entities independent of the remote storage location. Digital entities (sensor streams, files, database tables) are logically organized into collections. The SRB manages the logical context (both state information and descriptive metadata) in a Metadata Catalog (MCAT), and also provides facilities to associate user-defined metadata (both free-form attribute-based and schema-based) to enable data and resource discovery and access. The SRB has been used in national and international projects ranging from biomedical informatics to environmental data to high-energy physics data. A list of significant projects using SRB is shown in the Table below. These projects have contributed design requirements, for which extensions have been developed and up to four releases of SRB versions made per year. SRB is used currently at SDSC to manage more than 877 terabytes of data stored in more than 131 million files.

Integrated Rule-Oriented Data System

The next generation of distributed data management technology is being developed with NARA and NSF funding. The iRODS system integrates rule engines with data Grids to automate application of management policies. For each management policy, a set of rules is defined that control the execution of remote micro-services. Persistent state information is saved to track the outcome of the application of each policy. The iRODS system is extensible through the use of logical name spaces to identify rules, micro-services, and metadata. This means it is possible to add a new micro-service, a new controlling rule set, and new metadata without disrupting the execution of the original micro-services. The iRODS environment is designed to manage its own evolution, enabling interaction between new and old technologies transparently. Since the ORION cyberinfrastructure will evolve to manage interactive data stream manipulation, the ability to add new services without disrupting the file-based archiving technology is essential.

Example National and International Projects that have downloaded the SRB Data Grid Technology:

Affiliation	ProjectName
Academia Sinica, Taiwan	ASCC, Computing Centre
Advanced Computing in Victoria, Australia	Australian Partnership for Advanced Computation (APAC)
Alfred Wegener Institute, Germany	Collaborative Climate Data Grid
Bioinformatics Institute, Singapore	Bioinformatics Institute
British Antarctic Survey, UK	Data management project
Caltech	National Virtual Observatory
Cambridge e-Science Center, UK	eMinerals data grid
Cardiff University, UK	Welsh e-Science Centre
CINECA, Bologna, Italy	High Performance Computing - EUROPA project
City University of New York	National Science Digital Library (SRMA)
Colorado University	Cires/Cism
Cornell University	Fedora project
CRS4, Italy	Biomedical data grid
CSIRO, Australia	Bureau of Meteorology
Data Storage Institute, Singapore	Quality of Storage service
Drexel University	Digital library project
Environmental Protection Agency	EPA Data Grid Initiative
French National Center	Enabling Grids for E-science
GeoForschungsZentrum, Germany	Telegrafenberg
Griffith University, Australia	Research Computing Services

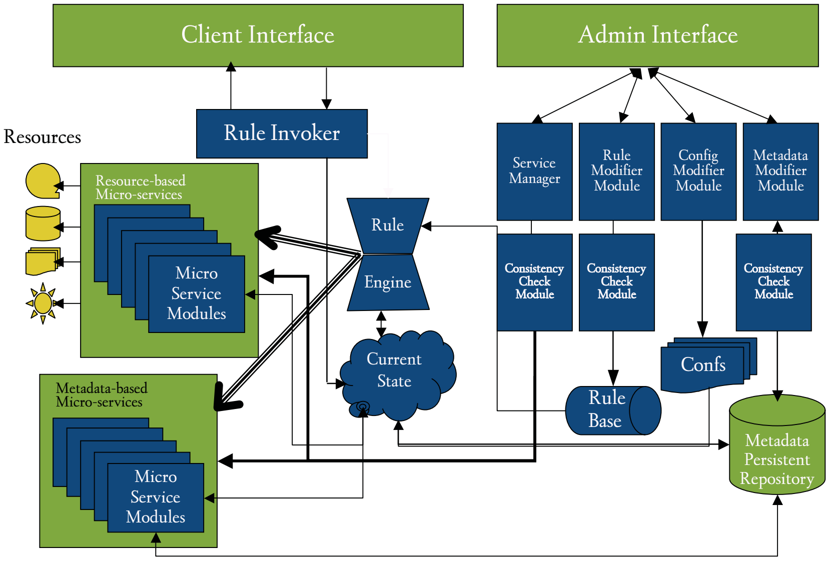

iRODS - integrated Rule-Oriented Data System

Indiana University	Digital Library Program
ISREC, Switzerland	Swiss Institute for Experimental Cancer Research
KEK, Japan	BELLE High Energy Physics Data Grid
Konkuk University, Korea	Korean national grid
Ky Dept. Libraries & Archives	Persistent Archive Testbed
LLNL	Digital Library for scientific collections
Melbourne, Australia	APAC Grid Project
MIT	Integration of Dspace , SRB
Monash University, Austalia	Microsoft SQL Server (MCAT port)
Nanyang Centre for Supercomputing, Singapore	Monash E-Research Grid
NASA Goddard DAAC, Code 902	EOSDIS Distributed Active
National University of Mexico	UNAM Grid
National University, Singapore	Bio data grid
NAVY	SPAWAR data grid
NIH	National Cancer Institute Center
NOAO	NOAO data grid
NYU Libraries	Web-at-Risk NDIIPP (CDL)
Osaka University, Japan	Virtual Tissue Bank
Pacific Northwest National Laboratory	BioPilot
Penn State University	CiteSeer
Pittsburgh Supercomputing Center	ETF project
Purdue University	ITAP data grid, TeraGrid project
Queen's University, UK	Belfast e-Science Centre
SARA	Computing and Network Services, Netherlands
Sickkids Hospital, Toronto	Medical data grid
SLAC/Stanford	BaBar High Energy Physics data grid
Space Telescope Science Inst.	National Virtual Observatory
Taiwan University, Taipei Taiwan	National Center for High Performance Systems
Texas A & M	Multiview Storage System
Tokyo Institute of Technology	NEES project
Trinity College, Ireland	TCHPC (HPC-Europa)
UC Merced	CUAHSI/ DLS
UCAR	NCAR Visualization
UCD	DBIS Lab
UCLA	Digital library project
UCSD	Swartz Center Neuroscience
UCSD	NCMIR, Telescience
UCSF	VA Medical Center, Workflow Project
University of Amsterdam	Virtual Laboratory for eScience
University of Bergen	Parallab (HPC-EUROPA project)
University of Bologna	Grid for Logistic Optimization, CS Dept.
University of Bristol, UK	Physics Labs
University of Buffalo	NEES project
University of Calgary	Research Repository with DSpace
University of Cambridge	UK e-Science
University of Edinburgh	University of Edinburgh
University of Florida	UF Research Grid (HPS)
University of Genoa, Italy	Laboratory for Bioimages and Bioengineering
University of Hawaii	Institute for Astronomy
University of Hong Kong	Computer Centre (Data Grid)
University of Kansas	Bioinformatics
University of Leeds, UK	School Computing
University of Liverpool	Dept. of Computer Science
University of Manchester, UK	WUN data grid
University of Maryland	Department of Computer Science, DataCutter
University of Michigan	CAC department
University of Minnesota	NEES project
University of New Mexico	Long Term Ecological Reserve
University of Oslo	Archiving scientific data (WUN)
University of Oxford	Large Hadron Collider Computing Grid
University of Pittsburgh	Library archive
University of Queensland, Australia	The Earth Systems Science Center
University of Sao Paulo, Brazil	Instituto do Coracao
University of Sheffield, UK	White Rose Grid
University of Southampton, UK	GRIA and SRB
University of Technology, Sydney	APAC Data grid
University of Texas	NEES project
University of the West Indies	Jgrass data grid
University of Washington	Streaming Technologies (WUN)
University of Wisconsin	Condor Project
University of Zürich	Computational Chemistry environment
US Army Research Laboratory	Rapid Unified Generation of Urban Databases
USC	Southern California Earthquake Center
USGS	Bedford Oceanography,Canada
Washington University	Department of Anatomy and Neurobiology
Woods Hole Oceanographic Institution	Multi-institution Testbed for Scalable Digital Archiving
Yale University Library	Digital Library
York Univ, UK	Worldwide Universities Network (WUN) data grid
ZIB, Germany	German Data Grid

6.2.5 MARS DATA COMMUNICATION SUBSYSTEM: EXPERIENCE IN HARDWARE DEVELOPMENT

PIs: Chave, Maffei, Yoerger, and Wooding

Title: Development of a Seafloor Cabled Observatory Communications System

Award Number: NSF OCE-0079720

Amount of Award: $1,400,000 plus $422,780 Cisco Systems match

Period of Award: 10/1/2000-9/30/06

PIs: Chave and Maffei

Title: MARS: Monterey Accelerated Research System

Award Number: subcontract from MBARI

Amount of Award: $1,303,177

Period of Award: 10/1/02-9/30/07

The Woods Hole Oceanographic Institution has extensive experience in the design, construction, deployment,and operation of hardware systems for oceanographic research. As an example, WHOI has carried out research and development on data communications subsystems for ocean observatories over the past six years. Goals for the work included research, design, construction, testing, and operating modular data communications subsystems for uses ranging from single buoys to the regional cabled observatory. This design has also provided the data communications subsystem for the Monterey Accelerated Research System (MARS) testbed.

The work was divided into three phases — 1) research and interface definitions, 2) sub-component identification, reliability analysis, development, and testing, and 3) underwater subsystem fabrication, integration, testing, and operations. An innovative aspect of this work has been the use of commercial-off-the-shelf, field-replaceable components. This has significantly lowered the cost of building submarine networks while taking advantage of the inherent reliability of COTS products rather than absorbing the costs required to validate the reliability of custom hardware solutions. The first two phases of the work have been completed and evaluated at standard design review stages. The last phase of the work has been the construction of a prototype (1-node) RCO observatory to be installed off Monterey, California, in 2007. The WHOI project team worked closely with its industry partner (Cisco Systems, Inc.) through the conceptual, preliminary, and detailed design stages for the seafloor and shore station communications subsystem.

A key design decision for cabled observatories was the use of eight-channel, full duplex, dense wavelength division multiplexing (DWDM) to achieve higher bandwidth using a small number of optical fibers. The optical design is a Gigabit Ethernet over DWDM solution. Using commercially available optical transponders, optical amplifiers and multiplexing equipment, distances in excess of 200 km between nodes are achievable. The equipment works within thermal, electrical, and physical size parameters acceptable for oceanic deployment. The RCO design uses eight wavelengths carrying 1 Gb/s each. This could easily be expanded in multiples of eight channels by using additional fibers. For buoyed installations, the same design, but without the DWDM and long haul optical equipment, is sufficient.

A novel aspect of the design is provision of a low-speed, highly reliable out-of-band communications system capable of operating over a variety of topologies independently of the Gigabit Ethernet backbone. This is essential for any seafloor network, as it replaces the service person that visits the router closet. As a significant bonus, this out-of-band system also supports the distribution of high accuracy time pulses (order of 1 microsecond, which is approximately 1000 times better than can be achieved using Internet protocols over the in-band network) that incorporates correction for time-dependent latency on long runs of optical fiber. This effort required at least four major innovations, and comparable commercial technology does not exist. A patent application has been filed for the out-of-band system.

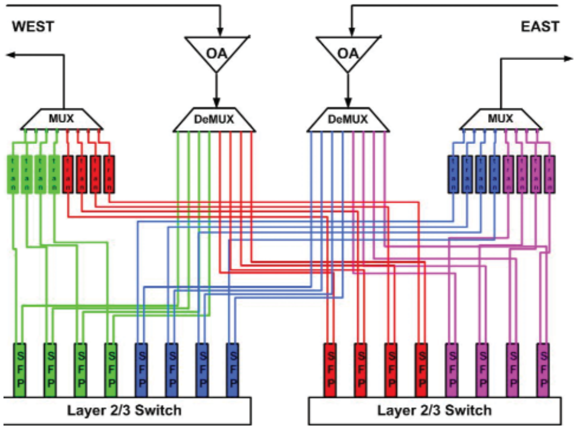

The optical and packet switching architecture for a cabled ocean observatory. A slightly modified version of this design is used in the shore station. Starting from the top, the optical amplifier (OA) units are terrestrial-grade erbium-doped fiber amplifers (EDFAs). Eight-channel optical multiplexers and demultiplexers provide the ability to run 8 optical channels over both an east and a west pair of fibers. The east and west fibers provide redundancy. Optical transponders convert the input signal from the layer 2/3 routers to individual DWDM wavelengths in the 1500 nm band.

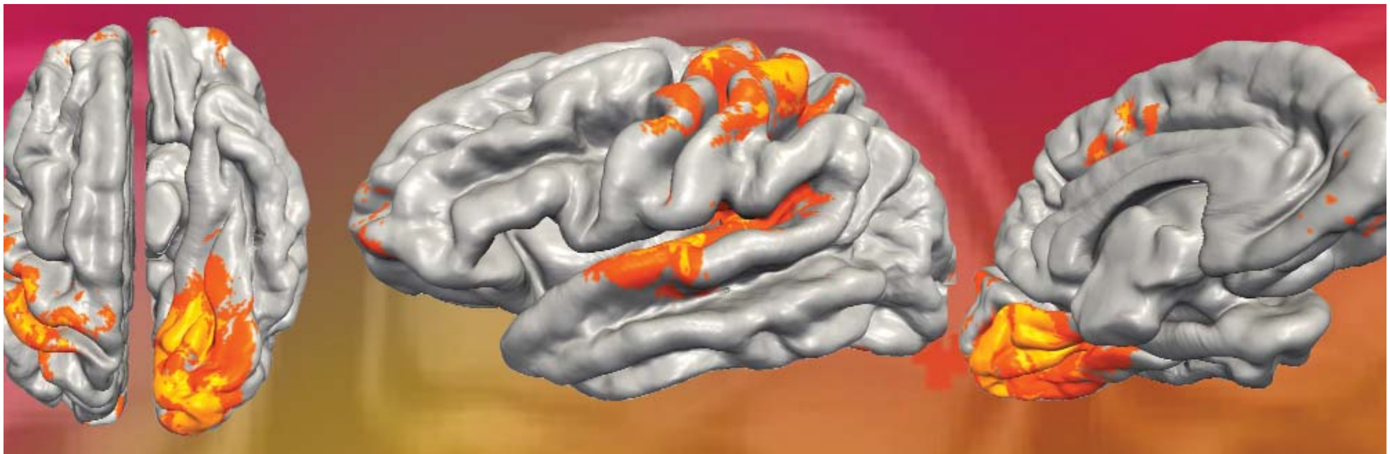

Brain areas activated during task involving finger tapping while viewing a flashing checkerboard and listening to a sequence of tones. Yellow and orange regions show areas reliably activated by this task, including motor, visual, and auditory cortices. Images courtesy of Steve Pieper, Harvard University/Brigham & Women's Hospital.

6.2.6 BIRN: EXPERIENCE IN PROJECT IMPLEMENTATION

PI: Ellisman (with Rajasekar, Moore, Gupta)

Title: BIRN; Biomedical informatics Research Network

Award Number: 1R24RR019701-01

Amount of Award: $18,579,826

Period of Award: 5/1/04-4/30/09

The proposed ORION knowledge management framework is based upon earlier SDSC work in the context of several other projects. An NIH-sponsored project, the Biomedical Informatics Research Network (BIRN), is a geographically distributed virtual community of shared resources offering tremendous potential to advance the diagnosis and treatment of disease. BIRN has changed how biomedical scientists and clinical researchers make discoveries by enhancing communication and collaboration across research disciplines—an objective shared by the ORION program.

BIRN's cyberinfrastructure consists of a cohesive implementation of key information technologies and applications specifically designed to support biomedical scientists in conducting their research. The widespread adoption of BIRN's cyberinfrastructure is allowing investigators to virtually pool their data and share common resources. In order to enable enhanced data sharing, a community-wide effort to develop biomedical ontologies is underway. Ontologies include computer-usable definitions of basic biological concepts and the relationships among them, and enhance the re-use of knowledge by providing a systematic framework for encoding knowledge within and across domains. The Biomedical research communities deal with large and often highly heterogeneous data sets, spanning the scale from whole organs to subcellular structures and traversing multiple species, conditions, and imaging modalities.

BIRN provides the means to integrate these data in ways that can generate new insights. To achieve this, a data integration engine called Metropolis has been developed that is based on a global-as-view mediator wrapper architecture. This system was designed to integrate information over a number of heterogeneous relational databases and currently operates on Oracle, SQL Server, MySQL, PostGreSQL, and other JDBC databases, as well as the Metadata Catalog (MCAT) of the Storage Resource Broker. Further, the engine allows one to use the external ontologies in defining integrated views, thus implementing "semantic information integration."

A number of specialized tools for schema registration, view definition, and query building, and a number of domain-specific clients have been developed to enable the practicing biomedical scientist to browse and query integrated information without the help of computer scientists. The system is currently being used in three different BIRN test beds, with about five to six data sources in each.

In BIRN, the engine sits within a complete cyberinfrastructure framework developed within the BIRN Coordinating Center, and thus is tightly linked with the user portal, the authentication/authorization mechanism, the ROCKS-based remote software deployment system, and the application clients developed by test beds and the Coordinating Center that access the mediator through a web-service mechanism.

The Metropolis subsystem was designed in a modular manner to work in any cyberinfrastructure project. It has been modified for the OceanLife Project funded by the Moore Foundation. In this project, the task is integration of physical oceanography data with biological observations on marine life, as well as ontologies such as habitat classifications. Wrappers were developed for OpenDAP, HDF/NetCDF and Spatial Analyst from ESRI data sources, and the mediator was modified to manage spatial data as first class objects within the mediator. This extended mediator has the ability to treat OWL-specified ontologies as a first class data object. In effect, it serves as a multi-model integration engine over a wider variety of data sources. The data integration engine is being shipped to OBIS (Ocean Biogeographic Information System) at Rutgers University with the goal of incorporating it into a production system in a phased manner.

6.2.7 CI COORDINATION OF SCIENCE CAMPAIGNS (SW06): EXPERIENCE IN PROGRAM OPERATIONS AND MAINTENANCE

PIs: Schofield, Glenn, Fennel, Wilkin, McGillicuddy, He, Gawarkiewicz, Moline

Title: Rapid Environmental Assessment Using an Integrated Coastal Ocean Observation and Modeling System

Award Number: Department of Defense, Major University Research Initiative Program (MURI)

Amount of Award: $4,916,133

Period of Award: 9/1/05-8/30/11

PIs: Glenn, Schofield, Kohut

Title: Adaptive Sampling in a Research Observatory During the Shallow Water 2006 Acoustics Experiment

Award Number: Department of Defense-Office of Naval Research

Amount of Award: $600,000

Period of Award: 10/1/05-12/30/06

Schofield and Glenn have extensive experience leading large interdisciplinary observatory-modeling science programs. These programs have been built around the concept of using observatory assets to coordinate and support many physical, chemical and biological experiments conducted by multiple science teams in real-time. As such, these experiments serve as prototype demonstrations of ORION. The real-time data provided by the Coastal Ocean Observation Lab (COOL) have contributed to the development of large "unplanned" science campaigns where distributed researchers joined originally funded teams to create a critical-mass data collection and exploration exercise.

As an example, the original 1998-2001 HyCODE/COMOP experiment called for 100 scientists, three ships, and one aircraft to provide optical closure in optically complex coastal waters. However, the promise of real-time data attracted additional scientists from NOAA, NASA, NRL, industry, and other universities. By the end of the science experiment, over 300 scientists, 10 ships, and four aircraft participated in the field/modeling effort. Some of the topics studied by these new partners included the impact of the neustonic layer on surface capillary waves, the first ocean deployment of a Webb Slocum glider, calibration/validation of HF radar networks, the migration of larval fish into bays and estuaries, the microphysics of aerosols in the atmosphere, and the operational detection range of swimmers at night. Each of these diverse projects leveraged the backbone of the COOL experiment where data and imagery were provided in real-time, without restriction to anyone. In fact, many of the experiments never developed a formal linkage to existing projects, and the support provided by COOL was often provided without Schofield or Glenn being aware of the new partners until after the experiment. The transparent web delivery of the data also galvanized local community involvement. For example, 75% of the over 250,000 daily hits on the COOL web site are from the general public (as ascertained by the web domain name).

ORION aims to provide science campaigns, and this kind of experiment is impossible to conduct without critical cyberinfrastructure capabilities. Schofield and Glenn have been conducting these science campaigns for a decade. Their most recent experiment provides a relevant example of the scale and potential for future ORION science campaigns. The Office of Naval Research (ONR) Shallow Water 2006 (SW06) Joint Experiment was the field component of three combined ONR programs: Littoral Environmental Acoustics Research (LEAR), the Non-Linear Internal Wave Initiative (NLIWI), and Autonomous Wide Aperture Cluster for Surveillance (AWACS). 50 PIs and co-PIs representing research labs assembled from government agencies, academic institutions, and companies (totaling an estimated 300 scientists and students) conducted independent, but linked, science experiments that ranged from the nonlinear dynamics of internal wave generation, propagation, and dissipation; the sub-mesoscale dynamics of the shelf-slope front and its effect on cross-shore exchange; the interaction of underwater acoustics with complex oceanography and simple bathymetry; to understanding the limits of predictability of shelf processes and their effect on acoustic uncertainty using state-of-the-art data assimilative physical-acoustic models. Deployed platforms included 58 oceanographic moorings, six UNOLs research vessels, a fleet of autonomous underwater vehicles made up of seven Slocum Gliders and four REMUS vehicles, and two research aircraft. All assets were linked through an extensive near real-time web-based data distribution system.

Rutgers scientists served as one of two NLIWI coordinating co-PIs with the responsibility of (a) maintaining a formation-flying fleet of Slocum Gliders at the shelf break for three months, (b) transferring real-time quality-controlled glider data to several modeling groups running predictive data-assimilative models, and (c) using the full resources of the Rutgers observatory to provide daily environmental summaries to scientists. The CI backbone was designed to promote collaboration between the scientists distributed between ships and labs around the country to enable adaptive sampling of all assets. The virtual collaboratory required hardware to sustain broadband connectivity and software to enable both continuous and asynchronous information sharing between ship- and shore-based researchers. Continuous ship-to-shore communication was maintained by antennas on the larger UNOLs vessels, which provided intermittent high-speed line-of-sight connectivity to the other vessels. Commercial WiFi connections enabled shore-based scientists to provide daily coordination reports to the research fleet from a wide variety of locations. The reports were distributed using the WHOI ExView software that enabled the sharing of products and discoveries, the archiving of discussions for later viewing, and the distribution of science and severe weather alerts. In addition, the multiple vehicle control and project coordination software enabled Rutgers to participate in simultaneous experiments in Hawaii and California during the extensive New Jersey operations.

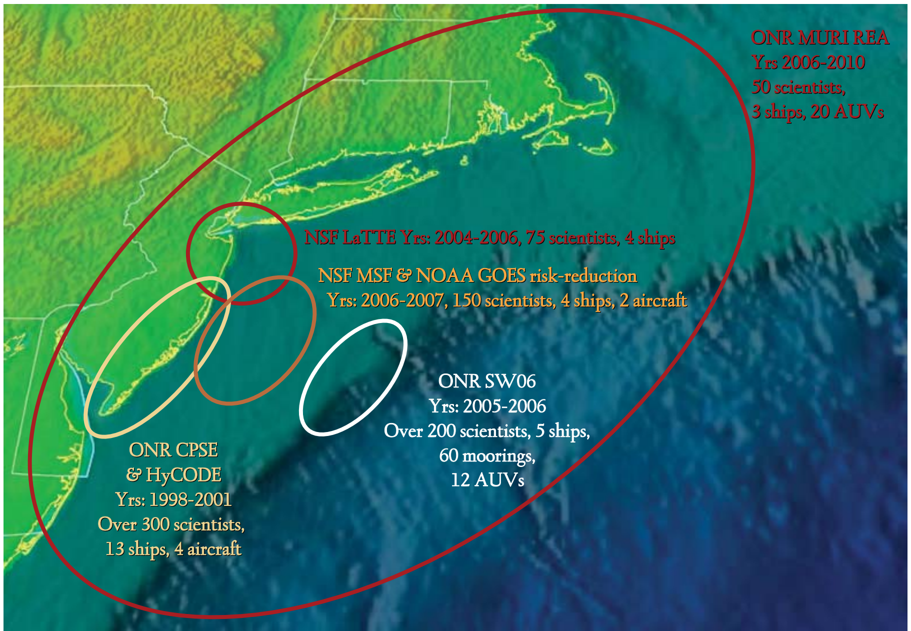

Observatory-Enabled Collaborative Research Campaigns in the Mid-Atlantic Bight with the scales of study increasing over time.

The transformational aspect of the SW06 Joint Experiment was the ability of the distributed cyber-physical observation network to provide unprecedented data coverage that inspired real-time collaboration between scientists independent of their physical location. For CI enabled science, data were acquired from multiple sources, some of which were not even associated with the main experiment. Daily scientific interpretations based on extensive 3-D visualization of the datasets proved critical. For CI-enabled technology, multiple leveraged testbed activities enabled iterative development and ultimate success of new platforms, sensors and software. Sustained and coordinated operation of glider fleets in multiple locations was demonstrated, but communication bandwidth was found to be the limiting factor for collaboration. For CI-enabled human interactions, the individual comfort level with real-time collaboration was wide ranging, and scientists were won over by new data products that emphasized discovery and enabled decision making. Over time the group of scientists actively contributing to the real-time collaboration grew.

Appendix 6.3-ORION CI Facilities

ORION CI Proposal

Scripps Institution of Oceanography Nimitz Marine Facility

JET PROPULSION LABORATORY, NASA

Image (Above): SGI Altix 350 with 16 processors

JPL is a federally funded research and development facility managed by the California Institute of Technology for the National Aeronautics and Space Administration.

JPL Supercomputing Center

Additional computing resources will be provided by the JPL Supercomputing Project (sc.jpl.nasa.gov) and consists of a 64-processor SGI Altix computer and 1024-processor Dell cluster computer system. The data storage will be provided by the JPL's StorageTek tape system with 360 TBytes of space.

NASA's Columbia supercomputer is a 10,240 processor system composed of twenty 512-processor nodes, twelve of which are SGI® Altix™ 3700 nodes, and eight of which are SGI® Altix™ 3700 Bx2 nodes. Each node is a shared memory, single-system-image (SSI) environment, running a Linux® based operating system. Four of the Bx2 nodes are linked to form a 2048-processor shared memory environment (2048-PE).

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Image (Above): Bluefin'21 AUVs and autonomous surface crafts operated by the MIT Laboratory for Autonomous Marine Sensing from the NATO R/V Leonardo in July 2005 at Elba Island, Italy

Laboratory for Autonomous Marine Sensing

The MIT Laboratory for Autonomous Marine Sensing in the Department of Mechanical Engineering is specializing in the development of new distributed ocean sensing concepts for oceanographic science, national defense and coastal management and protection. The laboratory was established in 2005 following the merger of the Mechanical and Ocean Engineering departments. It continues two decades of multi-disciplinary research and development into such systems by the Department of Ocean Engineering, and the MIT Sea Grant AUV Laboratory. In addition to the Laboratory Director, Prof. Henrik Schmidt; four faculty, Prof. John Leonard, Prof. Daniela Rus, Prof. Arthur Baggeroer, and Prof. Pierre Lermusiaux; and four Research Engineers, Dr. David Battle, Dr. Arjuna Balasuriya, Mr. Joseph Curcio and Dr. Michael Benjamin contribute to the lab's research. The faculty and staff, together with a significant number of students and post-doctoral associates provides a strongly multidisciplinary team with expertise in oceanographic sensing and modeling, sonar system technology, computational underwater acoustics, and marine robotics and communication networking.

Major Equipment

The Laboratory owns and operates two state-of-the-art Bluefin'21 autonomous underwater vehicles (AUV), and a fleet of ten autonomous surface craft (ASC), as well as support equipment for their operation, including an acoustic communication bouy, and a Sonadyne long baseline navigation system. The laboratory is operating the AUV and ASC from research vessels, and have been deployed and operated on average once a year during major field experiments. Through joint research agreements with the NATO Undersea Research Centre (NURC) in La Spezia, Italy, the Laboratory has regular access to their two research vessels, and their engineering research staff. The Laboratory operates several payloads for the AUV, including three acoustic sensing payloads for oceanographic mapping and bottom/sub-bottom object detection. Currently, a new payload is being built under ONR DURIP funding, including a highly capable acoustic towed array, intended for littoral undersea surveillance, and other underwater acoustic research.

Computing

The laboratory operates a cluster of ten workstations, in addition to a large number of laptop computers used for the field efforts. Most of the computers are operated using the Linux operating system, with the laptops in general having a dual-boot capability in Windows. The computer network includes a central CVS server for the open-source MOOS-IvP autonomous platform control software applied exclusively for the robotic platforms in the Laboratory. It also supports a state-of-the-art underwater acoustic modeling capability available to faculty, staff and students. A specific capability is a comprehensive MOOS-IvP simulation capability, which is used for simulation experiments, which the Laboratory uses extensively in the planning and preparation of field experiments, and for development and testing of new autonomous behaviors and processes. A unique capability of the MOOS-IvP simulator is the full integration of a high-fidelity, real-time underwater acoustic modeling capability, which supports simulation of fully autonomous acoustic sensing missions with AUVs towing hydrophone arrays. This capability is a crucial component of the development of the fully integrated,

Five MIT Autonomous Surface Craft in use in the Charles River Test Bed

autonomous sensing, modeling and control concepts being required for distributed underwater surveillance systems with no or extremely limited possibilities for operator involvement and intervention.

Charles River ASC Test-bed

Through an agreement with the MIT Sailing Pavilion the Laboratory has established a test-bed for the autonomous surface craft on the Charles River, adjacent to MIT. This Charles River ASC Test-bed is used for smaller scale testing of new adaptive and collaborative sensing concepts. As such it provides a useful intermediate link between the simulation environment and ocean field experiments. The test-bed is used routinely for development and testing of collaborative and adaptive, autonomous control of up to ten ASC using the MOOS-IvP control architecture.

Harvard Ocean Prediction System

With Prof. Pierre Lermusiaux joining MIT in Jan. 2007, and becoming the newest faculty member in the group, the Laboratory will add the Harvard Ocean Prediction System (HOPS) to its in-house capabilities. The Laboratory faculty has a decade long history of collaborating with the HOPS group, and most recently, as part of the ONR PLUSNet effort, HOPS is being closely integrated with the MOOS-IvP simulation environment, further enhancing the Laboratory's capabilities for simulating adaptive and collaborative sensing missions with distributed, fixed and mobile platforms.

MONTEREY BAY AQUARIUM RESEARCH INSTITUTE

Image (Above): Moss Landing Marine Laboratories

Laboratory

The Monterey Bay Aquarium Research Institute (MBARI) is a privately funded, nonprofit oceanographic research Institution. The Institute maintains three research vessels and associated remotely operated vehicles (ROVs) and autonomous underwater vehicles (AUVs) for deep-sea research, as well as several offshore moorings for monitoring physical, chemical, and biooptical parameters in Monterey Bay. A 15-foot by 30-foot electrical engineering/software development lab space has been allocated to this project and is available for use by project staff for the duration of the requested funding.

Computer

Computer resources allocated to this project by MBARI:

Beowulf compute cluster, Racksaver Systems

- 8 2-CPU Xeon 2.4 GHz, 512MB Ram, 40GB disk, 1GB Ethernet nodes (16 CPU nodes)
- 1 2-CPU Xeon 2.4 GHz, 2GB Ram, 240GB disk, 1 GB Ethernet master
- 24 port 1GB Ethernet switch
- dedicated keyboard/mouse/monitor tray
- rackmount chassis with dedicated UPS, Power distribution, and cooling fans
- 1 TB RAID disk array for capturing streaming video
- Fiber channel interface to RAID disk array
- HD/SD direct video capture at full frame rate and full resolution

Video Capture workstation

- 2 GHz Pentium IV, 512MB Ram, 80GB Disk
- Pinnacle Targus Digital capture card
- IEEE 1394 interface
- Fiber channel interface to RAID disk array

2 Linux workstations

- 1.6-2.4 GHz Pentium IV, 512 MB Ram, 200GB Disk

Network disk storage (1 TB) for storing processed results

Office

Office space for the P.I. and technician has been allocated by MBARI to this project.

Temporary office and lab space for visiting collaborators and project team members and summer interns has also been allocated by MBARI to this project.

Other

Other resources allocated to this project by MBARI:

- MiniDV digital video Deck (JVC BR-DV600UA)
- Video Monitor (Sony PVM14L5/1)
- Sony Digital BetaCam playback deck
- Panasonic High Definition (HD) digital video Deck

Access to Sony Digital BetaCam decks and MBARI's library of ROV dive video (on Digital BetaCam tape and on Digital HD tape) is provided

NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS

Facilities, Equipment and Other Resources

NCSA continues to support user communities by offering the resources that are the foundations of advanced cyberinfrastructure. The total computational resources exceed 43 TF supported by over 1 PB of disk storage as part of the infrastructure. The systems are on an internal 10GbE network. Below is a summary of those resources.

NCSA Compute Resources

Copper (Cu)

Power4 IBM p690 systems 384 processors, 8 with 64 GB/system, 4 with 256 GB/system Peak performance: 2 TF 35 TB SAN GPFS filesystem

Tungsten (W)

Xeon 3.2 GHz Dell cluster 2,560 processors, 3 GB memory/node, Myrinet Peak performance: 16.4 TF 140 TB Lustre Filesystem

Mercury, Phase 1 (Hg 1)

Itanium 2 1.3 GHz IBM Linux cluster 512 processors, 4 GB and 12 GB memory/node, Myrinet Peak performance: 2.662 TF 230 TB SAN GPFS filesystem

Mercury 1 and 2 comprise the largest computational resource for the TeraGrid

Mercury, Phase 2 (Hg 2)

Itanium 2 1.5 GHz IBM Linux cluster 1334 processors, 4 GB memory/node, Myrinet Peak performance: 8 TF 50 TB NSD GPFS filesystem

Mercury 1 and 2 comprise the largest computational resource for the TeraGrid

Cobalt (Co)

SGI Altix systems, 2x512 processors Itanium 2 1.6 GHz systems, Linux 1,024 processors, 3 TB total memory Peak performance: 6.5 TF 370 TB SAN storage with SGI CxFS filesystem 8 x 8p SGI Prism visualization systems with Infiniband interconnects to the 512p SMPs 30 TB SGI-based Oracle server

Tungsten 2 (T2)

Intel EM64T 3.6 GHz Dell Linux cluster 1024 processors, 6GB of memory per node, Infiniband interconnect Peak Performance 7.4 TF 4 TB IBRIX filesystem Primarily used by NCSA Industrial Partners

Mass Storage

The environment currently consists of 2 SGI Origin 3900 servers running EMC/Legato DiskExtender (UniTree) with 35 TB of SAN disk cache, 38 LTO2 tape drives, 6 IBM3590 tape drives, and 2 ADIC libraries. The total archival storage capacity of this environment is 3 PB.

Infrastructure SAN

284 TB of SAN connected storage for infrastructure and special projects.

High-Performance Network

All computing platforms are interconnected to a multi-10gigabit network core. The NCSA high-performance computing environment has access to the Abilene high-performance network through a shared 10-gigabit-per-second connection. NCSA also is one of the leading sites for I-WIRE, an optical networking project funded by the state of Illinois. I-WIRE provides lambda services for several projects, including NCSA's 30-gigabit-per-second connection to the TeraGrid network.

Display Systems

Tiled Display Wall: This environment consists of 40 NEC VT540 projectors, arranged in a matrix 5 high and 8 across. The output of the NEC VT540s is rear-projected towards a single screen, creating a large-format, high-resolution image space that is 8192 x 3840 pixels. A 40-node PC Linux cluster is used to drive the display wall. The machines are dual-processor Intel Xeons, running at 2.4 GHz, with Nvidia FX 5800 Ultra graphics accelerator cards, and communicating over Myrinet.

High Definition Passive Stereo Theater: The NCSA High Definition Passive Stereo Theater is a 1920x1080 display on an 6' x 3'5" screen. The projectors used are JVCD-1LA. The display is driven by a dual AMD Opteron 242 processor running at 1.6 GHZ. Graphics hardware consists of a Nvidia Quadro FX3000.

Applications Software

NCSA offers a variety of third-party applications and community codes that are installed on the high-performance systems at NCSA. These applications cover a wide range of science and engineering domains, data analytics and visualization, mathematics and statistics. Complete information on the packages available and detailed descriptions of them are available at: http://hpcsoftware.ncsa.uiuc.edu/Software/user/index.php?view=NCSA.

NORTH CAROLINA STATE UNIVERSITY

Image (Above): NCSU Engineering Building II exterior in Spring

The NCSU component of this project will be performed at the Department of Computer Science's Service-Oriented Computing Laboratory. The Department recently moved to a new building (inaugurated April 2006) on the NCSU Centennial Campus.

The Service-Oriented Computing Laboratory houses several Graduate Research Assistants with commercial-quality work facilities (work cubes and furniture). It includes several computers including three dual processor workstations, and is networked into the Department of Computer Science research network. The lab conducts research on services and agents for information management, IT architectures, Web semantics, business processes, trust, organizations, and technical aspects of policy management.

To facilitate the dissemination of research results, the Department provides Web servers running PHP for use by laboratory faculty and staff. It also provides an online Technical Report repository and links to the Ancestral International Index of Computer Science Technical Reports (NCSTRL).

The Department provides fully backed up storage for research and education use.

Additional facilities that are accessible to us are described below.

Networking Facilities

North Carolina and North Carolina State University boast among of the most extensive and sophisticated advanced high-performance communications infrastructures available for broad-based use today. The facilities include

- The three Research Triangle research intensive universities (Duke, UNC-Chapel Hill and NC State) have a long and fruitful history of partnership and collaboration in leading edge research. This includes networking and a joint high-performance production network and testbed (NCNI GigaPoP) spanning Raleigh, Durham, Chapel Hill and Research Triangle park. This network is operating at 2.4 to 10 Gbps speeds with drops at NC State, UNC-CH, Duke University, MCNC/NCSC and several industrial research sites.
- Abilene (Internet2): NC State is a member of Internet2 and has Abilene connectivity at 2.4 Gbps. Its Centaur Laboratory [Centaur] hosts the North Carolina Internet2 Technology Evaluation Center [ITEC]
- National Lambda Rail (NLR): NC State is a member of NLR—a 10 Gbps national research network [NLR]
- NC State's high-performance production network has capabilities that include a 4 and 10 Gbps backbones with redundant 1 Gbps (and in some cases 10 Gbps) drops into buildings and research laboratories that will participate in this project.
- The research network consists of over a number of switches and edge-devices from many vendors. The extent of the research network implementation is comparable to that of the production network and provides similar service coverage. We operate a number of networking protocols within both networks, and we experiment with a variety of them in the research portion of the network. This includes IPv6.
- NC State has a campus-area wireless networking solutions (802.11b based).

The above networking facilities are available for use in research projects. In addition, facilities in our specialized research and teaching laboratories including the

NCSU Engineering Building II Atrium interior

Computer Science Undergraduate Networking Laboratory, the Centennial Networking Laboratories, the Electrical and Computer Engineering Networking Laboratory and the Graduate Networking Laboratory are being used by courses in hands-on networking, as well as by researchers. These facilities can be accessed through NCSU's Virtual Computer Laboratory (VCL) solution.

NC State Networking Technology Institute (NTI), which includes Center for Advanced Computing and Communication (CACC) and the Centennial Networking Laboratories (CNL) provides on a cost-recovery basis networking technical support, and any additional networking related equipment, that may be needed for any project. It also provides advice to the researchers and students.

Computing Facilities

The College of Engineering and **NC State** maintain a state-of-the-art general-purpose academic computing environment known as Eos/Unity, a large-scale distributed system that consists of literally thousands of Unix and Windows-based workstations and servers all over campus. Eos/Unity uses robust, centrally managed storage and application system that features a number of software packages and tools, including simulation and analysis software, and mathematical software. The academic computing environment is operated by a professional support group that provides consultation and basic system and software services.

NC State University High-Performance Computing (HPC) operations provides NC State students and faculty with entry and medium level high-performance research and education computing facilities and consulting support. This service is complementary, and now joint with the NC State Grid operations which build on the NCBiogrid project.

Special Laboratory Facilities

SOC: Service-Oriented Computing Lab—research on advanced services

SDM: Scientific Data Management Center facilities (sdm.ncsu.edu)—workflow support

VCL: Virtual Computing Laboratory

CDL: Cyber Defense Lab—Security research

PrivacyPlace.org: Privacy and policy (theprivacyplace.org)

CNL: Centennial Networking Laboratory —testing and evaluation of production-level solutions

Centaur: Centaur Laboratory—Internet2 and National Lambda Rail

Software Engineering Laboratory: software and system reliability, fault-tolerance, testing, software processes

Space

Office space is provided for faculty, graduate students, and technicians in the Department of Computer Science. Routine technical support of the faculty and students is also provided by the Department of Computer Science.

RUTGERS UNIVERSITY COASTAL OCEAN OBSERVATION LAB FACILITIES

Image (Above): Shore-based CODAR receive antenna

The Rutgers University (R.U.) Coastal Ocean Observation Lab (COOL) run by Drs. Scott Glenn and Oscar Schofield and includes an interdisciplinary scientific research group (marine.rutgers.edu/COOL), an education outreach group (www.COOLclassroom.org), and an Operations Center (www.theCOOLroom.org). Faculty and students comprising the scientific teams participate in collaborative research programs in which academic, industry and government partnerships are forged between physicists and biologists, between scientists and engineers, and between observationalists and modelers. The education group is the focal point for outreach activities to the K-12 community and to non-science majors within Rutgers. The Operations Center maintains a sustained coastal ocean observatory that provides real-time ocean data to the research and education groups and also serves as the training ground for Operational Oceanography students.

COOL Operations Center

The COOL Operations Center maintains the world's most advanced coastal ocean observatory. State-of-the-art sampling capabilities are continuously upgraded as new technologies developed and demonstrated by the research group are immediately transitioned into the operational setting of the Center. Cost-effective sustained spatial sampling of the coastal ocean is accomplished with a variety of new platforms and sensors that include: (1) the local acquisition of satellite imagery from the international constellation of thermal infrared and ocean color sensors, (2) a triple-nested multi-static HF radar network for surface current mapping and waves, (3) a fleet of long-duration autonomous underwater gliders equipped with physical and optical sensors, and (4) a cabled observatory for water column time series. Raw datasets are shared with a variety of super-users throughout the U.S. for real-time backups, data archiving, and advanced product generation. Operational data products are produced in real time and displayed on the World Wide Web for use by scientists, educators, decision-makers and the general public. Website access peaks in the summer, averaging over 250,000 hits/day in summer 2006.

SeaSpace Satellite Acquisition Systems

COOL has continuously operated an L-Band satellite tracking and data acquisition system since 1992 and a larger X-Band system since 2003. Both systems enable local real-time access to the full resolution direct-broadcast imagery from an international constellation of polar orbiting satellites. The L-Band system currently tracks the NOAA Polar Orbiting Environmental Satellites (POES) and China's FY1-D. Products include the operational Sea Surface Temperature (SST), visible and simple ocean color. The X-Band system is used to acquire data from more recent satellites featuring higher spatial and spectral resolution. This currently includes the NASA MODIS (both Terra and Aqua satellites) and India's OceanSat. The increased spectral resolution enables more advanced ocean color products to be generated in optically-complex coastal zones. Tracking multiple satellites, including those operated by other countries, increases revisit intervals, providing multiple overflights of rapidly evolving coastal features at different times of day. Missed data due to groundstation downtime is minimized through an automated real-time backup system with the University of Maine where either system can write recently acquired raw data to the other's pass disk if it senses a disruption, enabling the downstream data flow to continue uninterrupted.

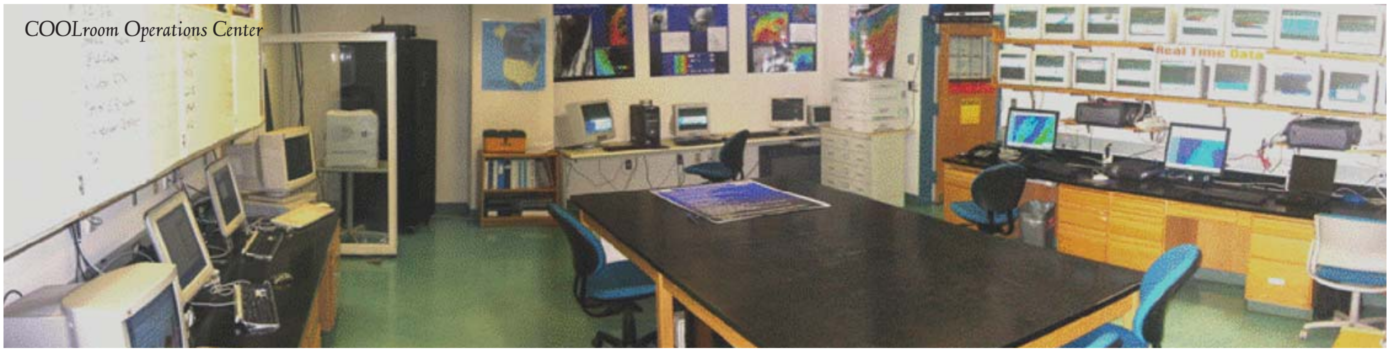

COOLroom Operations Center

CODAR HF Radar Network

CODARs are compact HF radar systems that provide a current mapping, wave monitoring and ship tracking capability. COOL has continuously operated CODAR HF radars since 1999. COOL currently owns and operates 13 individual CODAR HF radars deployed in three nested multi-static networks in the New York Bight. Traditional HF radars operate in a mono-static backscatter mode, where the transmitter and receiver are collocated. Multi-static operation, enabled by GPS-based synchronization, allows a radar receiver to acquire signals from any radar transmitter within range. This transforms N individual mono-static radars into a network of N2 multi-static radars, increasing both the coverage area and the accuracy of the derived current fields. Nesting is achieved by operating at different frequencies, in our case 5 MHz, 13 MHz and 25 MHz. Higher frequencies result in higher resolution but over shorter ranges. The long range 5 MHz network is deployed on the New Jersey coast and the island of Nantucket, providing coverage of the continental shelf out beyond the shelf break. The intermediate 13 MHz network is deployed around the entrance to New York Harbor. The high-resolution 25 MHz network is deployed at the entrance to and within New York Harbor. In addition to the usual shore based systems, COOL operates the only two buoy-based bi-static transmitters, a larger spar buoy for 5 MHz and a smaller surface buoy for 25 MHz. The transmitter is bi-statically paired with an onshore receiver, extending coverage offshore and improving the accuracy of total vector currents nearshore. In 2005 a compact super-directive receiver at 13 MHz system which increases range and directivity was added.

Webb Slocum Glider Fleet

Slocum Gliders are autonomous underwater vehicles that propel themselves through the water by changing their buoyancy and using their wings to glide in a sawtooth pattern through the water column along a subsurface transect. At user specified intervals, the glider surfaces, transmits its data to shore via the Iridium satellite system, and checks its email boxes for new directions or missions. The Slocum Gliders have been operated jointly by COOL scientists and Webb Research Corporation engineers in science experiments since 1999, transitioning to sustained deployments by the COOL Operations Center in 2003. Since then, the Gliders have logged over 25,527 km of underwater sampling in the New York Bight, offshore Massachusetts, Virginia, California, Hawaii, Florida, as well as the waters offshore Australia, England, France, Asia, and Germany. Sensors on the gliders currently include a SeaBird CTD and a payload bay capable of carrying one of several optical sensors, including Scattering Attenuation Meter (SAM), hyperspectral absorption sensor, and the ECO-VSF pucks. A mission control center monitors glider progress on current missions and alerts operators of any problems. Artificial intelligence is being added to the mission control center using an Agent oriented programming approach similar to NASA's approach for intelligent spacecraft. Reactive Agents are currently used to make many of the yes/no control decisions while Planning Agents are being developed to adjust flight paths to optimize sampling for specific goals.

Shore-based CODAR buoy-based bi-static transmitter.

Cabled Observatory

The Mid-Atlantic Bight (MAB) National Undersea Research Center (NURC) has operated the LEO-15 cabled observatory located 10 km offshore of Tuckerton, New Jersey since 1997. The buried electro-optical cable connects the Tuckerton shore lab with a pair of instrumented bottom nodes, providing power and two-way communications. In partnership with WetSat, COOL scientists are leading the upgrade of the original equipment on the shore side and at the nodes to include higher bandwidth communications, more reliable power, and standardized instrument interfaces compatible with other operating or planned cabled observatories. Once the upgrade is complete, long-term sampling at the WetSat nodes will be controlled in the COOL Operations Center.

Shipboard systems and sensors

COOL researchers maintain an extensive set of towed platforms, winched profiling systems, and hand-deployed profilers for shipboard sampling. Towed systems include a SWATH downward facing ADCP and a Guildline Inc. MiniBat undulating towbody containing a FSI CTD, a WetLabs Fuorometer and a DTA OBS. Profiling sensors include standard physical CTDs and numerous bio-optical sensors to characterize both the apparent and inherent optical properties. Optical sensors include a Satlantic downward and upward facing ultraviolet/visible radiometer, a Wetlabs ac-9, a Wetlabs spectral fluorescence excitation/emission fluorometer (SaFire), 2 3-wavelength Wetlabs EcoVSF systems, 2 Sequoia LISST systems, a Satlantic hyperspectral radiometric buoy (TSRB), a HOBI labs 6-wavelength backscatter sensor (HS-6), a bioluminescence bathyphotometer (BIOLUM) and an above water hyperspectral HOBI labs reflectance meter (HYDRORAD).

Research Vessels

Rutgers maintains the 30 foot R/V Caleta in New York Harbor and the 50 foot R/V Arabella in Tuckerton for scientific sampling and diving. COOL scientists also make constant use of the New Jersey Marine Science Consortium's fleet at Sandy Hook, the New York Oil Spill Response Team's fleet at Miller's Launch on Staten Island, and the SeaTow fleet in Atlantic City.

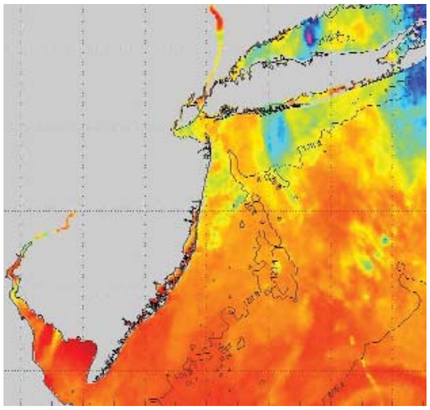

*Sample satellite product-*NOAA AVHRR SST

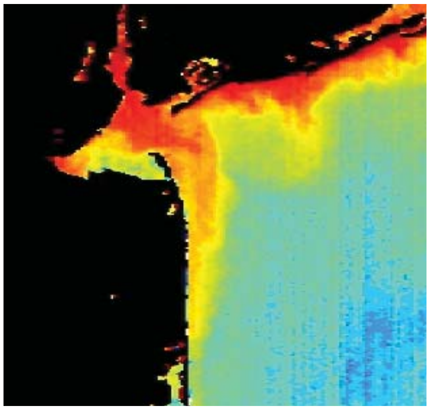

*Sample satellite product-*OceanSat *ocean color*

Mooring and Tripod Systems

COOL researchers have deployed numerous surface moorings, bottom tripods and bottom sensor packages on the New Jersey shelf and estuaries since 1992. Surface moorings include two CODAR bistatic transmitter buoys, a MetOcean meteorological buoy, and several high resolution thermistor and low resolution CT strings. A Benthic Acoustic Stress Sensor (BASS) Tripod has been deployed both as a self-contained system and as a plug-in device to a guest port in the cabled observatory. Numerous diver or ship-deployed bottom packages (ADCP, CTD, optical sensors) have been constructed for deployments on the shelf and in the estuaries.

Meteorological Systems

COOL researchers maintain modern meteorological sensor suites in three locations. A meteorological tower with instrument platforms at 8 m, 16 m, 32 m and 64 m is located at the waters edge in Tuckerton. A 10 m meteorological tower, weather station, and atmospheric profiling SODAR is located 2 km inland within the national estuarine reserve. A trailer-based mobile weather station consisting of a 3 m tower, and a portable atmospheric profiling SODAR is operated at the Pinelands research station 50 km inland.

Microbiology Lab

The Rutgers COOL lab shares one of the world's largest phytoplankton analysis facilities built in collaboration with Dr. Paul Falkowski's Environmental Biophysics and Molecular Biology Lab. The facility is equally shared between the 2 groups The facility is outfitted with a full suite of microbiological supplies, two mass spectrometers, two scanning UV/vis Aminco DW-2000 spectrophotometers, 3 Fast Repetition Rate Fluorometers, an SLM Luminscence UV/vis spectrofluorometer, a Coulter Multisizer, a full genomics laboratory with full sequence capability, a Pulse Amplitude Modulated Fluorometer, 3 Photosynthetrons, 1 Shimadzu High Performance Liquid Chromatograph, a Perkin Elmer C/N Analyzer, a scintillation counter and a full image analysis system connected to a phase contrast microscope.

Rutgers University (R.U.) Coastal Ocean Observation Lab (COOL) Facilities

Satellite Tracking and Data Acquisition Systems

SeaSpace L-Band System with 1 m Tracking Antenna

SeaSpace X-Band System with 3 m Tracking Antenna

Linux Back-up system

CODAR HF Radar Network

Six (9) 5 MHz Long-range SeaSonde CODAR Systems

One (1) 13 MHz Medium-range SeaSonde CODAR System

Two (2) 25 MHz High-resolution SeaSonde CODAR Systems

Two (2) CODAR bistatic transmitter buoys (5 MHz and 25 MHz)

One (1) CODAR Superdirective receiver (13 MHz)

Ten (15) Macintosh CODAR data processing computers

One (1) CODAR Transponder for antenna pattern measurements.

CODAR servicing van equipped with tools and spare parts

Slocum Glider Fleet

Six (16) Webb Research Corp. Slocum Gliders
Optimized for shallow water
Each with standard CTD and payload bay

One (6) Eco-VSf Wetlabs sensors (for glider payload bay)

Five (5) Scattering Attenuation Meter (SAM) (for glider payload bay)

One (1) hyperspectral spectrophotometer (Breve Buster) (for glider payload bay),

Two (2) Slocum Glider simulators

Two Linux computer based control systems

Glider service van equipped with tools and spare parts

Two (2) glider shipping crates for air transport

Shipboard Sensors and Systems

Guildline Minibat Undualating Towbody with Falmouth Scientific CTD, DTA, OBS, WetLabs Fluorometer

Surface Towed SWATH platform for RDI ADCPs

One (1) 1200 KHz RDI Broadband ADCP

Seabird SBE-25 CTD

Seabird SBE-19 CTD

Other CTDs

Other Fluorometers

Satlantic hyperspectral radiometric buoy (TSRB)

Optical profiling cage

Satlantic ultraviolet/visible radiometer

Wetlabs ac-9

Wetlabs spectral fluorescence exitation/emission fluorometer (SaFire)

Two (2) 3-Wavelength Wetlabs EcoVSF systems

Two (2) Sequoia Scientific LISSTS

Hobi labs 6-Wavelength Backscatter sensor (HS-6)

Bioloumiescence bathyphotometer (BIOLUM)

Hobi labs above-water hyperspectral reflectance meter (HYDRORAD)

Moored sensors and systems

Two (2) 1200 KHz RDI Workhorse ADCPs

Two (2) Sontek ADPs

Two (2) Nortek Aquadop Acoustic Doppler Current Meters

Five (5) Sequoia Scientific Laser In Situ Scaterometers & Transmissometers (LISST)

Many Seabird CT sensors

Several Buoys

Meteorological Systems

64 m Meteorological Tower - Tuckerton

10 m Meteorological Tower - Tuckerton

Mobile Meteorological Trailer - Pinelands

Two (2) Remtec SODAR Atmospheric Profilers

MetOcean Meteorological Buoy

Communication Systems

Freewave Line of Sight Radio Modems

Satellite Broadband Systems

Iridium Phones

Cell phone Modems

Broadband line-of-sight antennas

Computational Facilities

Sun Webserver

RAID backup infrastructure

Disc backup systems

Three (3) 9-Monitor computer display systems

Atmospheric forecasting machine

Laboratory Facilities

Three (3) salt water test tanks

Guildline Salinometer

Two (2) Mass Spectrometers

Two (2) Scanning UV/vis Aminco Dw-2000 Spectrophotometers

Three (3) Fast Repetition Rate Fluorometers

SLM Luminscence UV/vis Spectrofluorometer

Coulter Multisizer

Full genomics lab with full sequence capability

Pulse Amplitude Modulated Fluorometer

Three (3) photosynthetrons

Shimadzu High Performance Liquid Chromatograph

Perkin Elmer C/N Analyzer

Scintillation Counter

Full image analysis system with a phase contrast microscope

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Image (Above): Geisel Library

Atkinson Hall

Facility Overview - UCSD Division

Richard C. Atkinson Hall

Floors or Levels:	7
Number of Rooms:	418
Assignable Area:	150,891 sq. ft.
Outside Gross Area:	245,173 sq. ft.

General Spaces

Atkinson Hall is a physical manifestation of Calit2's multidisciplinary agenda. A defining feature of the UCSD Division facility is the shared facility, including clean rooms for nanofabrication, digital theaters for new media arts and scientific visualization, test and measurement labs for circuit design, smart spaces for experiments in augmented reality, transmission and networking testbeds for wireless and optical communications, and labs for designing systems on a chip. In addition to such highly specialized research facilities, floors 1-6 of Atkinson Hall include reconfigurable open research spaces to accommodate hundreds of personnel, standard offices, sixteen conference rooms, and public spaces for events as well as informal collaboration and impromptu gatherings.

Specialized Facilities

First Floor

Digital Cinema This 200-seat theater/concert hall includes ultra-high resolution digital video/cinema projection (4K, 10,000 lumen and dual 1600x1200 computer/HDV projection, 7,000 lumen each) and 21 Terabytes of ultrafast disk playback and real-time computer graphics capability at 4K or HDTV. It is networked via 1 Gigabit and 10 Gigabit Ethernet to Calit2/SDSC servers and the CineGrid network, and has 22-channel 8.2 stereo sound. Capabilities also include high definition video, H.323 teleconferencing, and webcasting.

Immersive Visualization Lab A multi-screen, multi-user virtual reality environment has been created with 1980x2160 resolution in stereo, 60 Terabytes of data/visualization servers, and 30-unit dual-Opteron cluster w/120 Gigabytes of RAM. Networked via two10 Gigabit Ethernet connections to Calit2/SDSC servers and the CAVEwave network, it includes a 100-Megapixel panel display with HDTV input and HDTV uncompressed video conference equipment.

Multipurpose Room Designed for experiments exploring the audience's relationship to the media and the physical environment, the MPR has four 3,500 lumen 1600x1200 projectors, stereo video on 16'x12' screen, seating for up to 100, and bundles of single-mode and multi-mode optical fiber to the Calit2/SDSC server rooms for networking.

HD Production Studio This studio will be an advanced, high-definition video studio for production as well as experimental research. The facility will produce programming on the arts, sciences and other fields; create content for display on devices ranging from iPods to Calit2's 4K Super-HD system; originate faculty presentations to international conferences in high-definition through Calit2's CalViz system; and feature a Calit2-developed OptiPortal, which permits interactive discussion and high-resolution visualization for collaborators across campus—or across the world.

Audio/Video editing suites The suites are optimized to support high-definition video post-production of content for science programming as well as production of Super-HD content and scientific visualizations to be displayed throughout the


Atkinson Hall

Calit2 building. Equipment will include an 'encoder farm' to facilitate the export of edited videos into a variety of online, tape, and hard copy (DVD, CD)formats to support research projects and "tell" the Calit2 story to its constituencies and audiences.

Art Gallery This 800-square-foot networked exhibition space will showcase world-class experimental art and prototype technology.

Audio Spatialization Lab The Audio Spatialization Lab includes a reconfigurable, multi-channel audio system, which allows for the production of audio content in a variety of standard as well as custom multi-channel audio formats. A base setup of 16 audio channels provides for content production occurring alongside the development of software tools which facilitate multi-channel composition in addition to the development of imaging algorithms. The lab is configured to allow for synchronization and staging of complex audio along with multi-media content.

Performative Computing Lab This computer vision/motion capture lab supports research into new techniques for integrating image based data into computer environments. Hybrid approaches are provided to allow for specialized developments, as well as for the integration of previously disparate approaches for the capturing of complex spatial and motion data.

3-D Fabrication Lab Manipulating physical material with the same facility expected from manipulation of virtual material will be supported by the 3D fabrication lab. CNC machines such as a mill, router and lathe provide for subtractive fabrication, while rapid prototyping devices are an additive platform. A range of activities will be supported from the development of new machine control methodologies, to the development of new sculptural forms and the ability to design and implement designs for new types of devices such as antennas, encasement's and robotic parts.

Server Room Atkinson Hall initially had 1,000 sf of server space on the first floor, with room for approximately 32 racks (19x84x30) of equipment. The space provides 558,000 BTU (48 Ton) of cooling capacity and 4 3 phase 225 Amp power panels. The server room is currently being expanded to add 1000 square feet of additional server space.

Nanofabrication Facility An approximately 10,000 asf Materials and Devices laboratory supports nanoengineering, nanoscience, and nanomedicine research with a clean room (approximately 7,000 asf of clean space, class 100 and class 1000), a growth and processing facility, analysis facilities, and a nanomedicine laboratory. This state-of-the-art research and demonstration facility's capabilities will support research in photonics, electronic devices, semiconductor materials engineering, heterogeneous integration/packaging, and biomedical electronics. The process bays house functions including photo and electron beam lithography, nanofabrication, thermal processing, wet processing, metrology, and metallization/thin film deposition. The analysis facilities provide laboratories for the characterization of materials and devices developed in clean room facilities.

Wet Etch Lab This lab includes capabilities for fabricating RF circuit boards using chemical etching techniques and photograph development. Additional capabilities include tin plating and gold plating.

Fourth Floor

High-Definition Studio A dedicated broadcasting studio, the HD studio is capable of producing and transmitting high-definition quality audio and video over optical fiber to distant locations around the world-giving Calit2 "telepresence" at international conferences, and providing a venue for Calit2 and UCSD experts to be interviewed by national and international media.

Fifth Floor

Smart Room Lab Instrumented with displays, cameras, microphone arrays and other sensors to experiment with the concepts of augmented reality, telepresence and collaboration, this lab will develop a prototype of "Super-Studio", where the resident of this futuristic space will be able to experience virtual reality through a Smart-Window.

Circuit Assembly and Sub-System Integration Lab This lab supports preproduction assembly and testing of prototype circuits for large-scale demonstrations of wireless networks for telecommunications, telematics, sensors, safety and disasters. Production capabilities are available for building dozens of modules of various types required for the system. The lab also enables electrical schematic capture and Printed Circuit Board (PCB) layout tools.

MicroWave and High Power Circuits Lab The lab focuses on high and low-power microwave amplifier experiments for high efficiency and high linearity for wireless communication. The lab also includes capabilities for digital signal processing experiments with power amplifiers and low noise amplifiers to mitigate impairments and improve quality of wireless services. DC/DC 'super' converter experiments and wireless protocol-aware battery management experiments can be performed.

MillimeterWave and MicroProbe Circuits Lab Research capabilities of this facility include state of the art research on Radio Frequency (RF) transistors and RFICs. On-wafer probing of devices and circuits up to 110GHz for full S-Parameter Characterization, IV Curves, Noise Figure, Phase Noise is possible. The lab also includes a mmWave antenna range for testing intelligent multiple-antenna systems (MIMO) and high gain antennae.

Sixth Floor

Wireless Platforms Lab and System Integration Current platforms under development and test include portable software defined radio, bench-top software defined radio, wireless pulse oximeter, wireless paperless patient tag, ground-based wireless mesh network nodes, wireless smart-door control, wireless patient drug dosage monitor, and mushroom networks. These sub systems are also integrated in the lab.

Photonics Lab This facility houses a high capacity optical network testbed and associated physical layer research. The high capacity testbed includes up to 200 independent channels and advanced modulation facilities for 10Gb/s (OC-192) and 40Gb/s (OC-768) rates, complemented with real-time RS forward error correction (FEC) system. The testbed is equipped with a complement of fiber and EDFA/Raman amplifier plants for up to 1500km in-line transmission and the recirculating loop for ultralong haul (>5000km) experiments. An optical parametric facility provides ultrawideband fiber amplification, band mapping and 320Gb/s signal processing, and, combined with the conventional high-capacity capability, remains unmatched by present industrial or academic laboratories. A commercial 40Gb/s terrestrial system provides interoperability with in-ground and experimental networks, both on campus and nationwide. High capacity fiber research is complemented by high-speed free-space optical capability for next-generation access networks.

Systems-on-Chip Lab Calit2's SoC Lab is dedicated to the design and development of applications, architectures and system software for state-of-the-art microelectronic integrated systems. Capabilities will support new applications, system-chip architectures, system-chip (microelectronic and microfluidic) platforms and packaging, semiconductor intellectual property, and embedded software: mobile code, middleware and infrastructure software, and cross-cutting thrusts on energy, bandwidth, usability, availability, mobility, security and standards. It does so utilizing four core capabilities: (1) to architect, prototype and build hardware and software platforms that provide capabilities for a variety of sensor, embedded computing and wireless networking functions

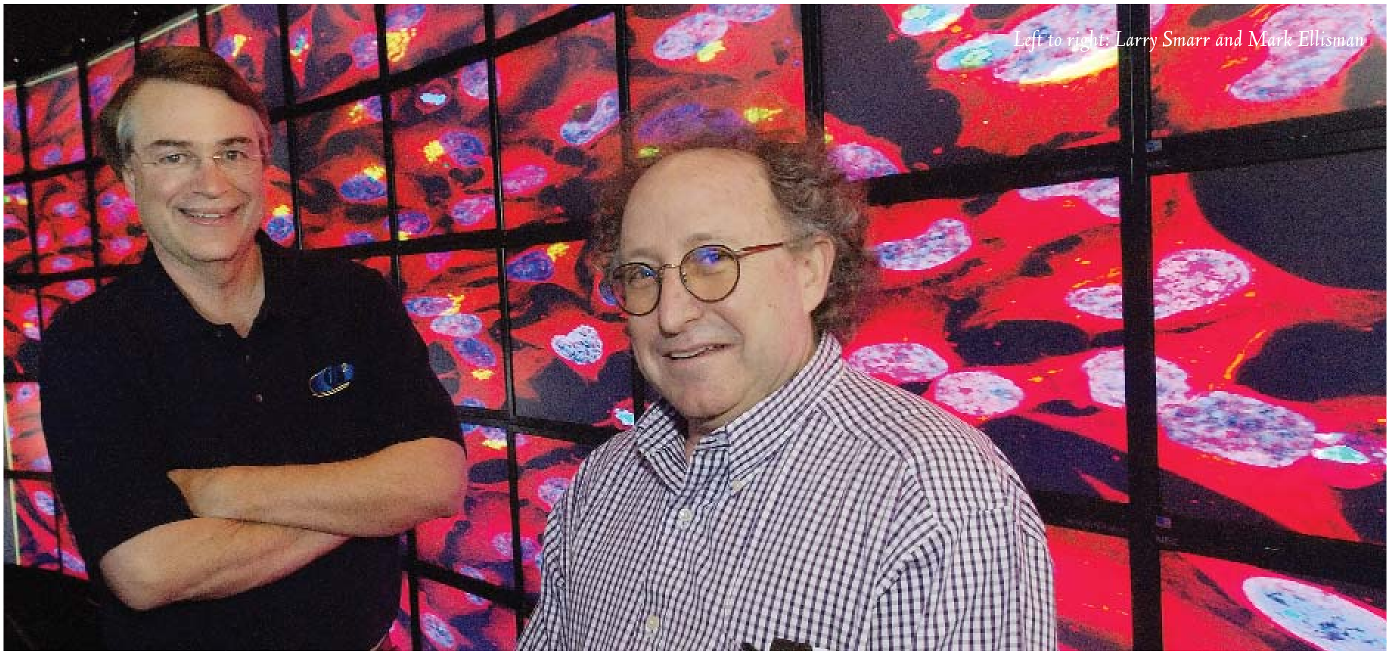

Left to right: Larry Smarr and Mark Ellisman

under domain-specific physical and performance constraints; (2) to build models of the target platforms and associated software development environments, (3) to test board- and chip-level parts operating up to 500 MHz, and (4) to design up to RTL implementations using IP blocks.

Undergraduate Research Laboratory This lab supports undergraduate research as part of senior-level, project-based courses as well as the Calit2 Scholars summer research program.

Roof

Antenna Garden Lab The roof of Atkinson Hall provides 13 antenna pedestals with AC power, GigiBit Ethernet, and RF cables for experimental wireless communication systems. Additional locations are available for cameras associated with wide area surveillance systems for experiments in artificial intelligence detection and estimation of security related events.

Radio Base Station (RBS) Lab An Ericsson experimental CDMA base station allows experimentation with live, on the air CDMA systems. Researchers can use the base station for experiments at the physical, MAC, and network layers.

UCSD Center for Research in Biological Systems (CRBS)

Laboratory

CRBS, directed by Mark H. Ellisman, is a UCSD organized research unit (ORU) that exists to provide an integrative framework to facilitate multiscale studies of biological systems by bringing together innovative, interdisciplinary teams, providing a vehicle for these teams to organize and manage resources, and providing a collaborative core or "glue" to sustain the energy and vision necessary to manage team science and lead novel scientific inquiry. CRBS is organized to uniquely propel integrative, multi-scale research through a marriage of leading edge science and technology. The activities of this project will be principally conducted within one of CRBS's cornerstone projects, the NIH/NCRR National Center for Microscopy and Imaging Research, where Dr. Martone is co-Director. This project will further leverage the expertise and resources of another CRBS cornerstone activity, NIH's Biomedical Informatics Research Network Coordinating Center (BIRN-CC) as well as the the California Institute for Telecommunications and Information Technology, which is a principal partner in CRBS technology development and houses the CRBS front office and many of its leading edge IT activities.

As described in detail below, it is possible that there is no other organization in the world with the unique convergence of domain expertise and resources available for software development to integrate multi-scale biomedical imaging, informatics, and advanced information technology as is assembled within CRBS.

Computer

The NCMIR provides a vast array of supporting high performance computational resources as well as unique access and connectivity to large-scale national and international resources such as the TeraGrid, NSF's OptIPuter, and an assortment of clusters associated with CRBS and NIH's National Biomedical Computational

Resource (NBCR, PI - Peter Arzberger). NCMIR's in-house computing and image-processing facilities consist of approximately 30 networked workstations: SGI, Sun, Intel Unix boxes and PCs (Wintel and Power Mac systems). These are networked, administered by professional staff and backed up nightly. This part of the laboratory is being upgraded continuously and is an official satellite facility of the San Diego Supercomputer Center (SDSC).

NCMIR also houses four clusters, two of which are dedicated for computing and two that have dual purposes as a computational cluster and visualization cluster. Cluster1 comprises 52 CPUs and at least 512MB of RAM per CPU, interconnected with high speed Myrinet. Vogon is a 34 CPU Opteron 244 cluster with 1GB of RAM per CPU and an internal PVFS high-speed storage area. Vogon is also connected to the OptIPuter network via a 10Gbit/s Extreme switch. Brainywall is a visualization cluster composed of 10 CPUs with 1GB of RAM per CPU and Quadro 900XGL graphics powering 19 million pixels on two IBM T221 displays. Brainywall is also connected to both commodity Internet as well as the OptIPuter network. Raster, another visualization cluster, is a 42 CPU Opteron 244 cluster with 4GB per CPU and QuadroFX 3000G genlocked graphics cards powering approximately 40 million pixels on 20 UXGA flat panels in a tiled display configuration connected to the OptIPuter 10Gbit/s network.

Throughout the NCMIR laboratory are immersive visualization and data interaction workspaces specialized for different image modalities. In addition to the 40 million pixel tiled display, NCMIR houses a GeoWall passive stereo 3D system, multiple video teleconferencing resources, and a 72-inch touch screen Smartboard. Within the lab is a two panel IBM T221 18 million pixel display, driven by a 32bit five-node cluster, which is currently being converted into a stereo display.

The NCMIR facility also has dedicated, high bandwidth access to resources within SDSC and the California Institute of Information Technology and Telecommunications (Calit2).

Through SDSC, numerous conventional and high-performance resources are available to NCMIR, including supercomputers, an IBM SP2, several Sun HPC Enterprise servers, and a TeraGrid IA-64 PC cluster. SDSC also houses the HPSS, a 240 TB archival storage system, a visualization laboratory, a Fakespace BOOM virtual-reality device, a tele-manufacturing facility for producing solid 3D models of geometric datasets, a variety of hard-copy and film output devices, and an audio/video suite for producing professional-quality videotapes of visualized research results. The NCMIR core facility machine room is connected to SDSC via a Cisco 6509 catalyst Ethernet switch to the campus backbone and SDSC's border gateway (Juniper T320). SDSC's border gateway connects over private fiber utilizing DWDM at 10Gbit/sec to the main CENIC Calren2 (www.cenic.org) HPR router in Los Angeles. This same CENIC router is 10Gbit/sec connected to the Internet2 Abilene LAX-NG gateway located at the Qwest Qwave facility in Los Angeles. (www.internet2.edu).

This project will also leverage the computational infrastructure of the BIRN-CC. The BIRN-CC supports 96 development, staging and production servers to support the following services (see Figure 1): BIRN web site, portal interfaces, manage security certificates, application development, database and data integration support, test bed applications, data management, and data storage. A new 33-node computational cluster is being deployed to the BIRN infrastructure to support large-scale computational applications like JHU's Large Deformation Diffeomorphic Metric Mapping (LDDMM) tool. Later this year, a 30 terabyte data cluster will be added to assist BIRN collaborators with short-term supplemental storage.

The BIRN-CC central compute and networking resources reside in the main SDSC computer room at UCSD (www.sdsc.edu). All BIRN compute/network devices are utilizing IP and IPv6 protocols and are 1Gbit/sec Ethernet attached into a redundant set of core Cisco 6509 catalyst Ethernet switches that comprise the SDSC backbone. These switches use HSRP to give redundancy for the network. These switches are connected via 10Gbit/sec Ethernet to each other as well as to the SDSC border gateway and the TeraGrid (www.teragrid.org) switch fabric. The TeraGrid switch fabric utilizes 30Gbit/sec connection into the TeraGrid backbone via a Juniper T640 router.

The BIRN-CC also occupies ~100 square feet of the machine floor of the SDSC to support the operation of the BIRN Network Operations Center and to provide an area for assembly and testing of the BIRN racks before shipment to the test bed sites. Hotline services available 24 hours a day, seven days a week allow BIRN participants to ask technical support questions, report problems, or request new services using an 800 telephone number. Often problems are resolved with a single telephone call. A problem ticket is opened if the item cannot be immediately resolved, and a triage system assigns the appropriate priority to the ticket.

Office

As previously described, the NCMIR occupies roughly 3500 square feet of the Calit2 building (UCSD Atkinson Hall). All of the project principals named in this proposal currently reside within this CRBS controlled space. All also have full access and privileges within the NCMIR core facility and within the office space of the BIRN-CC. The NCMIR core facility consists of over 7500 square feet of instrumentation laboratories, wet lab, general lab, offices, interactive conference room, computer lab, and machine room space. The suite is located on the first floor in the Basic Sciences Building (BSB) of the UCSD School of Medicine. The BIRN-CC facility consists of an additional 4300 square feet (also within the UCSD School of Medicine) consisting of office space, a conference room, and two supplemental meeting rooms for focused discussions and/or small impromptu meetings. All CRBS spaces applicable to this project are equipped with telephones, printers, plotters, and reception/administration facilities to support daily operations.

Other

Extending the video teleconferencing capabilities of the NCMIR, this project will further leverage the advanced VTC capabilities of the BIRN-CC which features a Multi-Channel Unit (MCU) device that provides support for up to 24 concurrent VTC sessions. Remote administrators can make reservations on the MCU via a web interface. Conference Meeting templates and Participants templates make it easy to add and create new meetings.

Image Analysis: The laboratory maintains an array of custom and commercially available software packages for confocal and electron microscopy image processing and analysis. A variety of commercial packages are available for use including ANALYZE software package (from Dr. Robb at the Mayo clinic), BITPLANE IMARIS, NEUROLUCIDA, AMIRA, and AUTODEBLUR (from AUTOQUANT for image deblurring and deconvolution).

In addition to these commercially software packages, several Java based applications have been developed at NCMIR. These applications include:

JFido: an interactive client-server application created for the Telescience environment. JFido allows a user to interactively view the raw series of projections, place fiducial markers on the images for alignment, and perform several image processing operations (cropping, gain normalization, designation of the tilt axis, etc). The application centralizes several software components developed locally at NCMIR and provides a simple graphical user interface, where previously a collection of tools with heterogeneous interfaces (command line, X Windows, etc.) were used.

JViewer: a web/portal visualization tool for generalized viewing of three dimensional geometry. Using JViewer, users can load SRB data from their own collections, open files from their own local file systems, or open volumes from remote URLs. Currently, JViewer can read three types of file formats: NCMIR's Xvoxtrace 3D contours, Microbrightfield's Neurolucida V3 ASCII contour and tree structures, and 3D Synu (polygon-mesh, contour, stack). This program is adaptable and will extended to handle additional data formats as necessary.

ImageJ: ImageJ (rsb.info.nih.gov/ij/) is an open source 2D image processing and analysis tool developed by the NIH. With the open source aspect of ImageJ, we will continue to develop custom image processing plug-ins to serve in the processing and assembly of 3D and 4D light microscopy montages.

All NCMIR Java applications are also available from the NCMIR Multi-scale Imaging Grid Portal (formerly Telescience Portal). The latter utilizes NCMIR developed Telescience technologies to consolidate access for controlling instruments remotely, managing data, and controlling batch jobs with a single login and password. The Portal walks the user through the complex process of remote data acquisition via Telemicroscopy; Globus-enabled parallel tomographic reconstruction; advanced visualization, segmentation, and data processing tools; and transparent deposition of data products into federated libraries of cellular structure. Key features of the Portal include personalized user information, collaboration tools such as chat and shared white boards, and automatic storage of data and job tracking tools.

HiSeasNet antenna installation on the SDSC roof

San Diego Supercomputer Center

The resources available through the San Diego Supercomputer Center include supercomputers, archival storage systems, data-handling platforms, high-bandwidth networking, and advanced visuali-

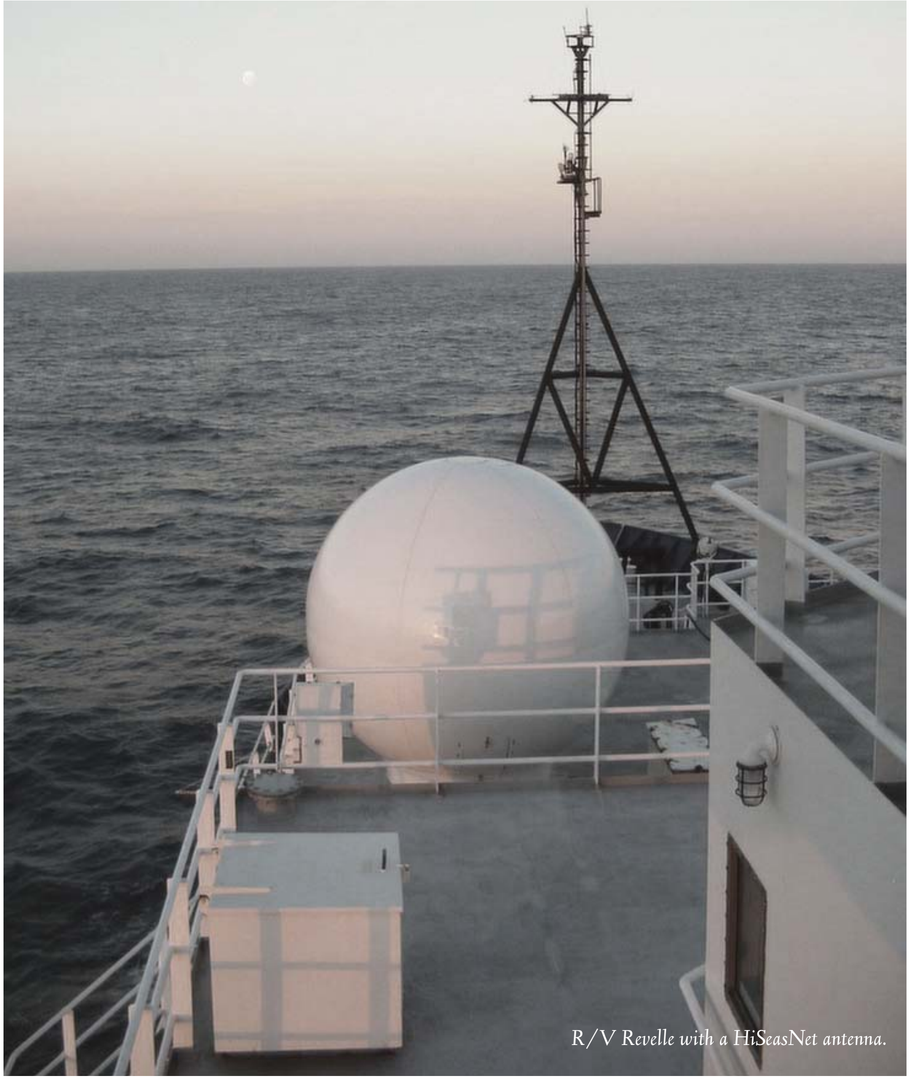

R/V *Revelle with a HiSeasNet antenna.*

zation systems. The capabilities of the center are being upgraded continually to include higher-capability systems that provide a robust environment for cyberinfrastructure research, development and deployment.

Among the hardware resources at SDSC, the foremost is DataStar, an IBM system with a peak performance of 15.6 teraflops. DataStar has 2,518 Power4+ processors in 283 nodes connected to the same high-speed Federation switch and parallel file system, giving DataStar communication and I/O performance far in excess of conventional clusters. SDSC also hosts an IBM/Intel cluster associated with the TeraGrid containing 512 compute processors with a peak performance of 3.1 teraflops. Most recently, SDSC has deployed the first IBM Blue Gene/L system at an academic institution. This unique architecture boasts 2,048 compute processors plus 128 nodes for the maximum I/O performance possible.

Data-handling resources include a storage-area network (SAN) of 1.4 petabytes (1,400 terabytes) of disk and a 25-petabyte tape-storage archive. Managed by a powerful Sun Fire 15K server, with 72 processors and 288 GB of shared memory, SDSC's data-handling environment provides support for databases, data management, and data mining. Associated data-intensive computing software includes the Storage Resource Broker, a distributed data management system developed at SDSC, digital library technology acquired through collaborations with MIT and Cornell, parallel object-relational database technology acquired in collaboration with IBM, and the High-Performance Storage System (HPSS) archival storage software that is being developed and tested in conjunction with IBM and LLNL. SDSC also has available and continues to work with Sun on the SAM-QFS online/archival storage environment. SDSC has integrated these systems to provide support for massive data collections. The archival storage systems at SDSC has 32 tape drives, and sustains up to 10 terabytes of data movement per day per tape drive.

SDSC's core program supports scientific data collections for disciplines including oceanography, seismology, neuroscience, molecular science, Earth systems science, and astronomy. Access to these data collections is provided through the SDSC Storage Resource Broker. The combination of information management technology, scientific data collections, and the data-handling platforms that support rapid access to the data provides an excellent testbed for evaluating new infrastructure for managing scientific data and scientific algorithms.

The SDSC Synthesis Center supports collaborative viewing of scientific data and advanced scientific visualization capabilities. A complete video and audio production suite is used to produce publication quality animations. The video lab is network accessible and can be used to render scientific images.

Scripps Institution of Oceanography

Keck Center

The Keck Center is located next to Nierenberg Hall, and has a bridge-style industrial crane comprising a horizontal load-bearing beam, which is attached at its end to wall columns. This crane is rated to 15 tons, and is used primarily for hoisting equipment (buoys, transducers, cages, etc.) into the test pool. The Center includes a 10 meter test pool with fresh/seawater capability and a pressure test facility.

HiSeasNet

HiSeasNet is a satellite communications network designed specifically to provide continuous Internet connectivity for oceanographic research ships and platforms. Access to the Internet is an integral part of nearly every research lab and office on land; extending this access to oceanographic ships—our seagoing laboratories—will broadly impact seagoing research activities. For the ships, HiSeasNet provides:

- Transmission of hot data in real-time to shore-side collaborators;
- Basic communications-email, voice and video teleconferencing for scientists, engineers and crew at sea;

Tools for real-time educational interactions between shipboard scientists, teachers and the classroom, as well as informal science and other education and outreach activities.

HiSeasNet provides connectivity to R/Vs Atlantis, Endeavor, Kilo Moana, Knorr, Langseth, Melville, New Horizon, Revelle, Pelican, Seward Johnson and the Thompson.

Marine Facility Shop (Marfac)

The marine facility shop building is 265 feet long and 70 feet deep on the main floor. This 18,500 sq. ft is divided into a carpenter shop, welding shop, mechanical shop, machine shop, electric shop, and the office of the shop superintendent and his assistant. The second floor is office space, store rooms, and a lounge, shower, wash room. There is a second building that houses the electronics repair shop and our radio station, WWD, and has a small indoor storage space for shop overflow. An outdoor storage area is available for large objects such as containers and winches.

The shop has four employees trained in pipefitting, general steel and aluminum fabrication, woodworking, hydraulic and mechanical repair, and the operation of material handling equipment. When the need arises, a pool of retired employees and ship's personnel on leave is available and can be called on.

The shop material handling equipment consists of five forklifts ranging in capacity from 15,000 pounds to 500 pounds. Two mobile cranes, one 10 ton and one 12 ton, are equipped with all the slings and bridles needed to lift anything in their capacity. The shop also has four cable spooling devices, one of them truck mounted and powered.

The portable winch pool has four winches that can be deck mounted to the standard bolt pattern. Three of them need only an electrical hook up. The fourth, which can carry 1/2" dredge or .680 cable, needs a hydraulic hook up. The shop also maintains eight portable marine cranes with either a 1000 or 2,000 pound capacity. Six plug in electrically and two need a hydraulic hook up. The MarFac shop is also the west coast depository for new UNOLS wire. A large supply of used, standard-sized oceanographic wire is maintained at the facility. A small supply of stock and hardware is kept in the shop, but a vast network of suppliers is just a phone call away.

Nimitz Marine Facility

The Nimitz Marine Facility, which is the support and management center for the Scripps fleet of research vessels and platforms, is one of the largest and most completely outfitted operating bases at any oceanographic institution. Located on 5.7 acres of land on Point Loma, the 110-

Nimitz Marine Facility

meter finger pier and 85-meter quay wall can accommodate 5 ships and the platform FLIP, and as many as 7 ships doubled up. The piers afford a complete suite of utilities connections for vessels. The marine facility serves as homeport for the NOAA ship David Starr Jordan, and it hosts visiting research vessels from US and foreign institutions as time and space permit. The facility houses R/V's Melville, Revelle, New Horizon, Sproul and FLIP.

Buildings adjacent to the pier and quay wall house shops, the control room of marine radio station WWD, scientific staging and storage areas, administrative offices, and shipboard technical support spaces and offices.

The marine facility is capable of carrying out a variety of ship maintenance, repair, and modification work "in house." Scientific equipment of every description can be loaded and unloaded, or prepared and sent to ports around the world for scientists from Scripps and from many other institutions.

Shipboard Technical Support (STS) provides specialized expertise, personnel, instrumentation and support services to science groups that use oceanographic research vessels.

SIO Hydraulics Laboratory

The SIO Hydraulics Laboratory provides space, facilities, and equipment for physical and biological oceanographic research. Building dimensions are 31 m by 46 m. An asphalt-surfaced work yard 31 m by 40 m is adjacent to the lab. The lab provides access to fresh and salt water (which can be chilled and filtered), drainage trenches, electric power (120 v single phase and 208 v and 460 v three phase), and compressed air. Also available are oil hydraulic power supplies of a combined total power of 160 kW for servo systems and variable speed drives. The lab includes the following equipment and facilities:

- Deep tank
- Granular fluid basin
- Pressure test chambers
- OAR pressure facility
- Oscillatory flow tunnel
- Rotating table
- Stratified flow channel
- Temperature/Pressure Calibration Facility
- Glass walled wave channel
- Field instrument pool

SIO Visualization Center

The Visualization Center at the Cecil and Ida Green Institute of Geophysics and Planetary Physics, Scripps Institution of Oceanography, University of California, San Diego is a state-of-the-art laboratory for the display, integration, analysis and exploration of various geophysical datasets supporting interdisciplinary research. It provides a diverse array of research tools supporting data sharing

iCluster at SIO

among distributed locations, live field reports, real-time data acquisition and presentation, real-time video teleconferencing and lectures. Its resources are also being applied to wider community needs, such as providing heightened response to natural disasters such as earthquakes, education, and outreach. The center comprises high-resolution projection and tiled display systems supporting geosciences visualization that are driven by multiprocessor SGI machines and Linux and MacOS X clusters. It is a node on the 10 Gbps OptIPuter grid.

The Highly Immersive Visualization Environment (HIVE) is a cylindrical wall display 28 feet wide and 8 feet tall installed in the Revelle Conference Room to provide immersive exploration capabilities to large groups of researchers (up to 40 people can be accommodated in the room). A 16 processor SGI Onyx 3400 drives three front mounted projectors that display on the Panoram GVR120E for a fully immersive environment.

Besides the SGI Onyx, the HIVE can also use a Windows PC (with the 3 channel Matrox graphics card) to display on the Panoram screen. The facility also has a DVD player, a VCR and the capability to connect any laptop using a VGA cable for presentation purposes.

The iCluster is a 50 megapixel tiled display visualization system built using Apple G5s and 30" monitors. This system is used to visualize real time data from the USArray network of sensors (part of the NSF funded Earthscope project) and other global observing systems. The iCluster is housed in the Earthscope Array Network Facility (ANF) office at IGPP and receives funding from the ANF, CEOA and Calit2. The iCluster has been built using a 7 node PowerMac G5 cluster and 12 displays arranged in a 4 wide x 3 tall array. Each display supports a resolution of up to 2560 x 1600. Each cluster node is a dual 2.5 GHz PowerPC with 8 GB RAM and the powerful NVIDIA GeForce 6800 Ultra graphics cards.

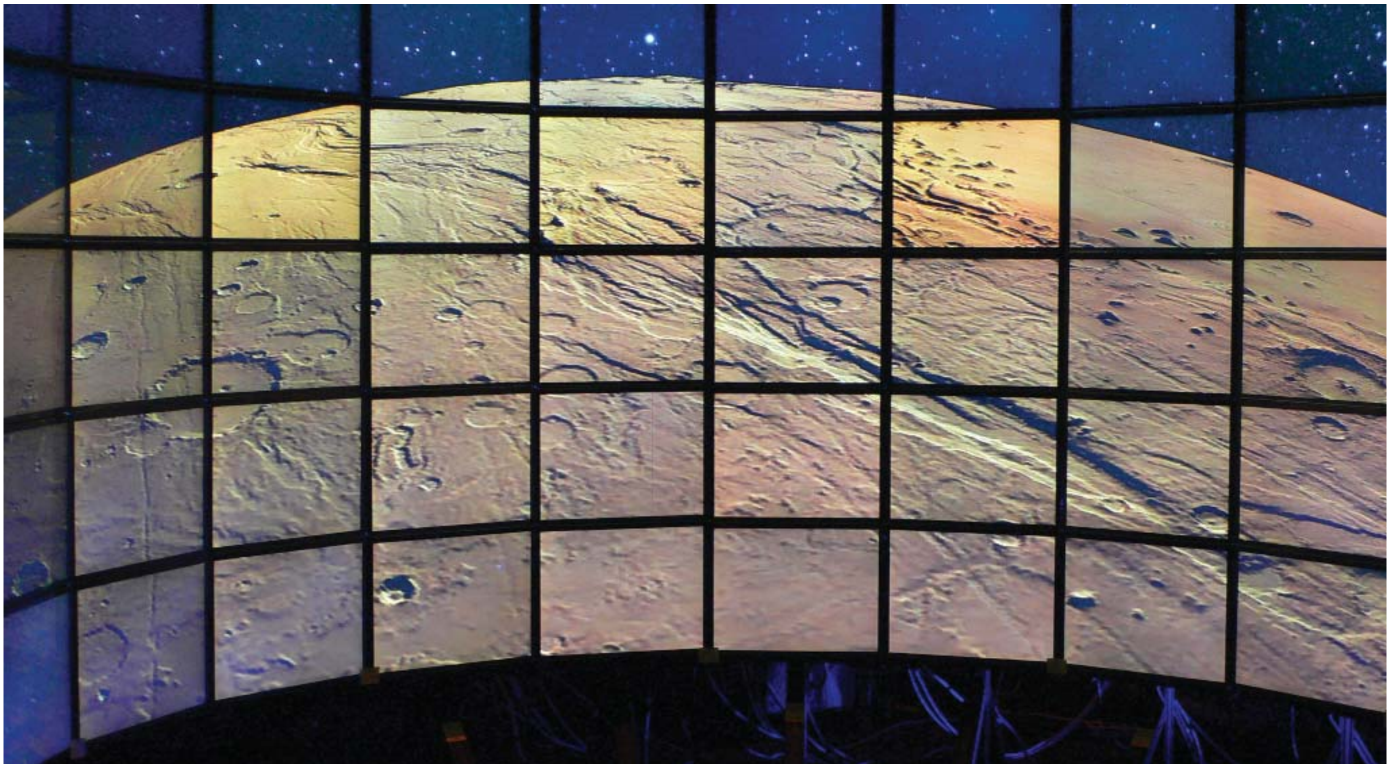

Image (Above): *NASA Mars terrain data set on Calit2's 65-panel Varrier. Varrier developed by EVL/UIC. Image courtesy of Dan Sandin, EVL/UIC.*

UNIVERSITY OF CHICAGO

The Computation Institute at the University of Chicago operates (a) TeraPort, a 256-processor compute cluster acquired under NSF MRI funding that serves as a facility for the University of Chicago community; and (b) the TeraGrid's visualization facility. The entire TeraGrid is distributed across NCSA, SDSC, the Pittsburgh Supercomputer Center, Purdue, Indiana University, the Texas Advanced Computing Center, ORNL, and University of Chicago. The individual clusters are connected by a dedicated 40 Gb/s link that acts as the backbone for the machine. The University of Chicago component of the machine consists of 16 dual IA-64 nodes for computation, a 96 dual Pentium IV nodes with G Force Ti 4600 graphics accelerators for visualization, and 20 TB of storage.

In addition, through the auspices of the Computation Institute, personnel associated with this proposal will have access to facilities associated with the Mathematics and Computer Science Division at Argonne, including major parallel computing clusters, visualization systems, advanced display environments, collaborative environments, and high-capacity network links. These resources include (c) "Jazz," a Linux system that has achieved a sustained teraflop, with 350 compute nodes, each with a 2.4 GHz Pentium Xeon with 1.5GB of RAM; Myrinet 2000 and Ethernet for interconnect; and 20 TB of on-line storage in PVFS and GFS file systems; and (d) a one-rack, 2048-processor IBM Blue Gene/L system with a peak performance of 5.7 teraflops.

University of Chicago and Argonne are both participants in the I-WIRE project, which links them at 10 Gb/s to TeraGrid and to the StarLight international interconnection point in downtown Chicago, and to other research institutions in Illinois.

UNIVERSITY OF SOUTHERN CALIFORNIA/ INFORMATION PROCESSING CENTER

USC/ISI has a distinguished history of producing exceptional research contributions and successful prototype systems under support from DARPA and other government agencies. USC/ISI has built a reputation for excellence and efficiency in both experimental computer services and production services. USC/ISI originally developed and provided and/or now helps to support many mature software packages for the entire Internet community.

Computer

USC/ISI's computer center has been an integral part of ISI since its founding in 1972. Today's Information Processing Center (IPC) maintains a state-of-the-art computing environment and staff to provide the technical effort required to support the performance of research. Resources include client platform and server hardware support, distributed print services, network and remote access support, operating systems and application software support, computer center operations, and help desk coverage. The IPC also acts as a technical liaison to the ISI community on issues of acquisition and integration of computing equipment and software. In addition, research staff have access to grid-enabled cluster computing, and to USC's 5,400-CPU compute cluster with low latency Myrinet interconnect that is the largest academic supercomputing resource in Southern California.

The ISI Center for Grid Technologies has a 96 node Linux (Debian) cluster located at ISI with Gigabit connectivity to the local LAN. Each node in the cluster is a dual-processor node with 1 GHz processors, 1 GB RAM, and with SCSI and IDE connectivity. The cluster is used on a daily basis for regression testing of nightly Globus builds, and general purpose high-performance computing.

WORK BREAKDOWN STRUCTURE

	WBS	Activity Name	Duration (Work Days)	Start Date	Finish Date
1	1	**Development Cycle 1**	**381.00**	**3/1/07**	**8/29/08**
2	1.1	Inception Phase...	**130.00**	**3/1/07**	**8/31/07**
3	1.2	Elaboration Phase...	**82.00**	**9/4/07**	**12/31/07**
4	1.3	Construction Phase...	**126.00**	**1/2/08**	**6/30/08**
5	1.4	Transition Phase...	**43.00**	**7/1/08**	**8/29/08**
6	2	**Development Cycle 2**	**337.00**	**5/1/08**	**8/31/09**
7	2.1	Inception Phase...	**85.00**	**5/1/08**	**8/29/08**
8	2.2	Elaboration Phase...	**84.00**	**9/2/08**	**12/31/08**
9	2.3	Construction Phase...	**125.00**	**1/2/09**	**6/30/09**
10	2.4	Transition Phase...	**43.00**	**7/1/09**	**8/31/09**
11	3	**Development Cycle 3**	**337.00**	**5/1/09**	**8/31/10**
12	3.1	Inception Phase...	**85.00**	**5/1/09**	**8/31/09**
13	3.2	Elaboration Phase...	**84.00**	**9/1/09**	**12/31/09**
14	3.3	Construction Phase...	**125.00**	**1/4/10**	**6/30/10**
15	3.4	Transition Phase...	**43.00**	**7/1/10**	**8/31/10**
16	4	**Development Cycle 4**	**337.00**	**5/3/10**	**8/31/11**
17	4.1	Inception Phase...	**85.00**	**5/3/10**	**8/31/10**
18	4.2	Elaboration Phase...	**84.00**	**9/1/10**	**12/31/10**
19	4.3	Construction Phase...	**125.00**	**1/4/11**	**6/30/11**
20	4.4	Transition Phase...	**43.00**	**7/1/11**	**8/31/11**
21	5	**Development Cycle 5**	**461.00**	**5/2/11**	**2/28/13**
22	5.1	Inception Phase...	**86.00**	**5/2/11**	**8/31/11**
23	5.2	Elaboration Phase...	**83.00**	**9/1/11**	**12/30/11**
24	5.3	Construction Phase...	**126.00**	**1/3/12**	**6/29/12**
25	5.4	Transition Phase...	**166.00**	**7/2/12**	**2/28/13**

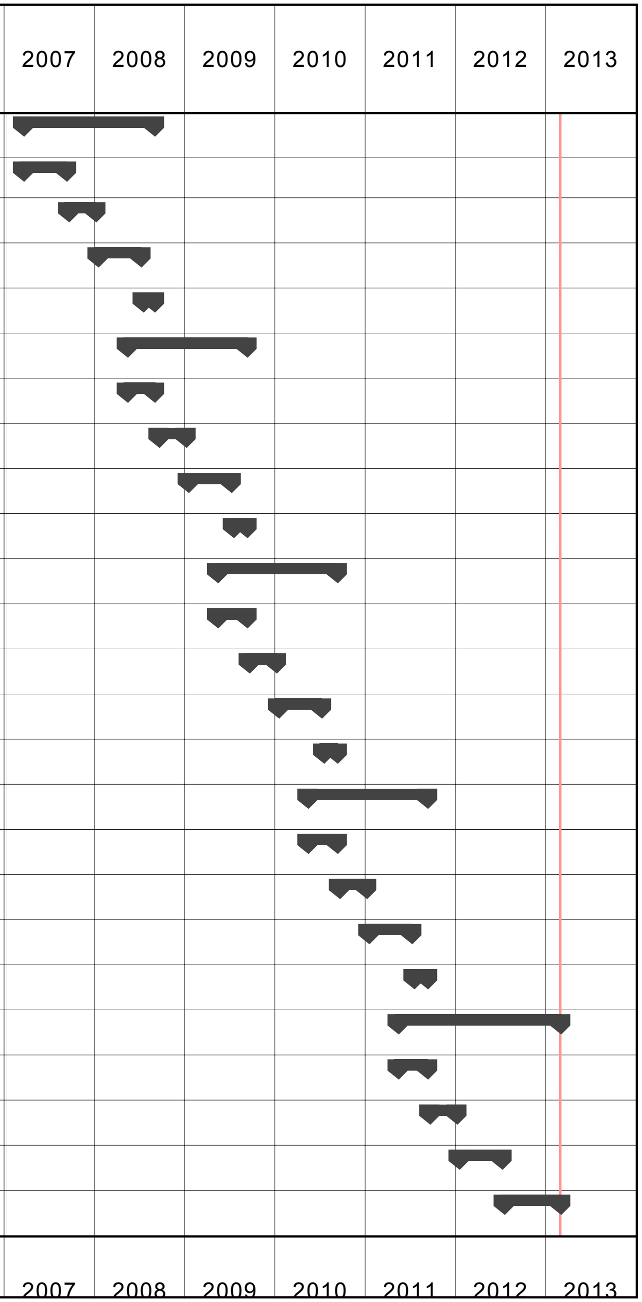

Please se CIIO-WBS.mpp on CD-ROM for full WBS (above is top two levels of a four level break down structure)

WOODS HOLE OCEANOGRAPHIC INSTITUTION

Image (Above): Chave-Blake building

Chave, his engineering, and administrative staff have office space in the recently renovated Blake building on the village campus of Woods Hole Oceanographic Institution. Chave's electronics and instrument fabrication laboratory is also conveniently located within the building.

Computer

All analysis will be undertaken on a Macintosh computer maintained by Chave. Two PC's are dedicated to instrument support and a third PC is dedicated to project administration.

Other

Project dedicated administrative support is available for tasks that are directly related to the project.

WHOI maintains a wide variety of services for use by members of the staff. Services provided by overhead funds that will be important to this project include:

The Computer and Information Services department supports all computer-related applications including networking and system upgrading. The computer infrastructure is in place to support data archival and analysis. CIS services are also available as a technical resource for project specific advice and expertise on emerging technologies.

Grant and Contract Services facilitates contracts, subcontracts and awards and provides guidance regarding agency policy and procedures.

Office of Finance & Administration oversees the Institution's compliance with regulations of state and federal agencies with the assistance of an independent auditor, PricewaterhouseCoopers; is responsible for the Institution's insurance risk management, and insurance provisioning related to specific contracts and awards.

The Controller's Office provides general accounting, accounts receivable, and property administration services and maintenance of project financial status reports used by all projects.

The Procurement department supports the acquisition and payment of goods and services conforming with the policies and procedures of most federal and private agencies.

Exhibits

ORION CI Proposal

Exhibit 1: Ocean Observatory

ORION CI Proposal

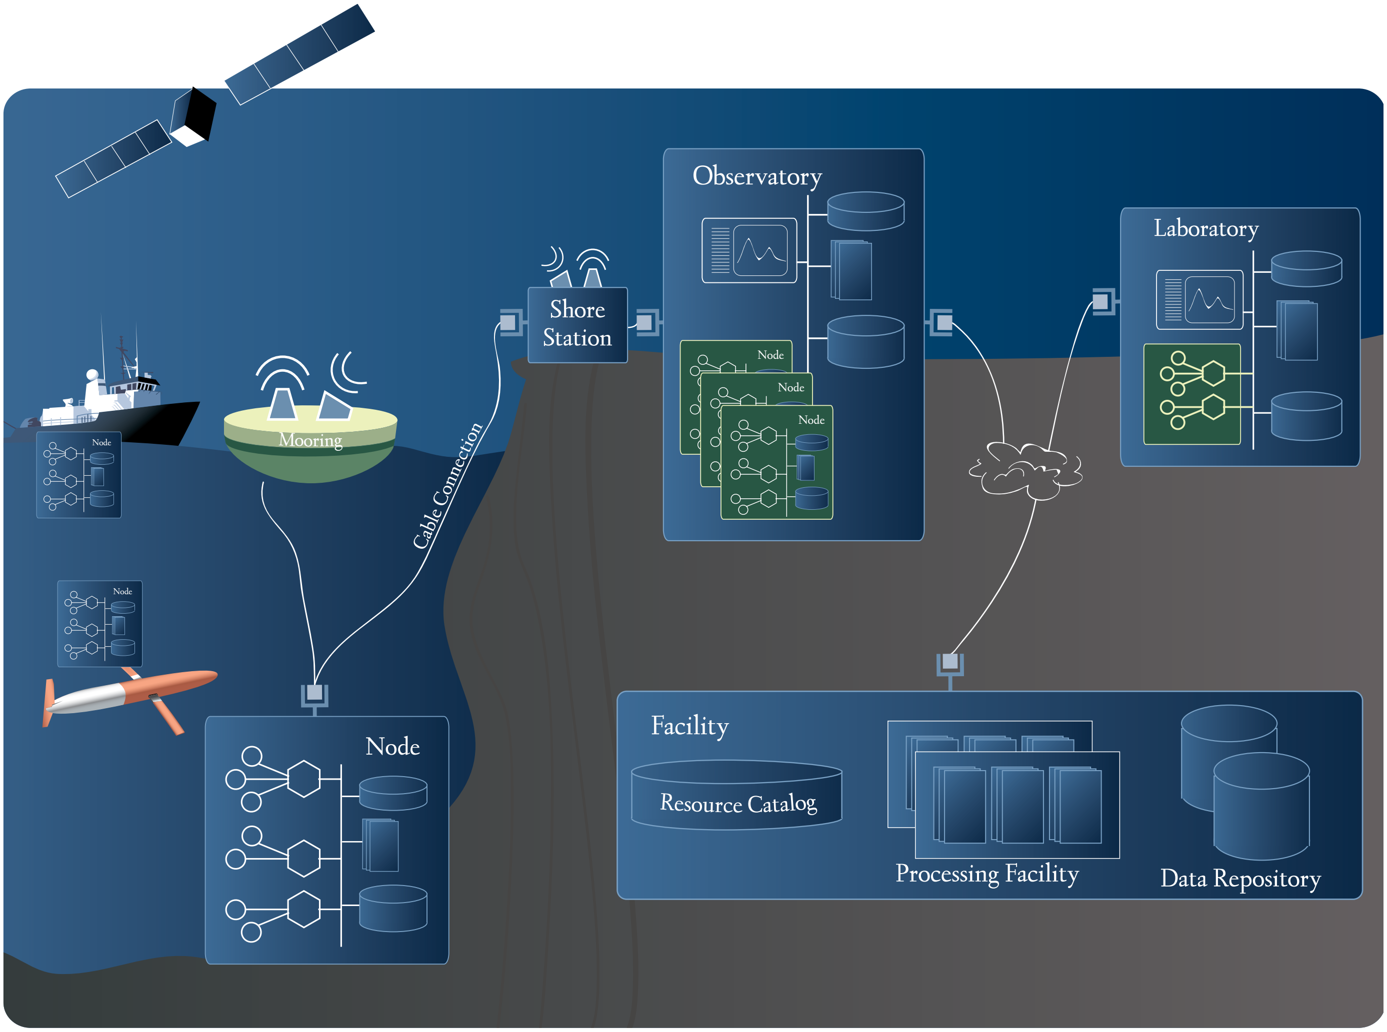

Exhibit 2: Observatory Scales

ORION CI Proposal

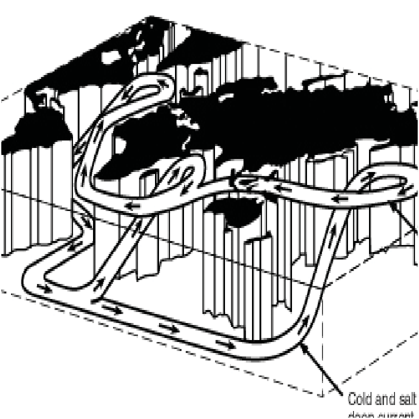

1000-km, decades

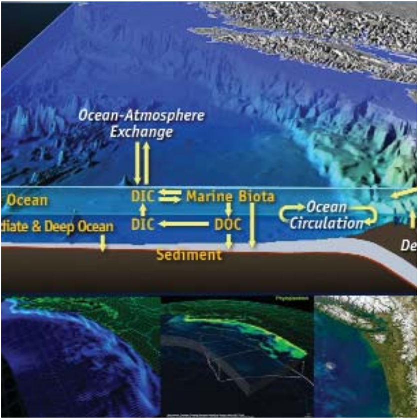

100-km; years

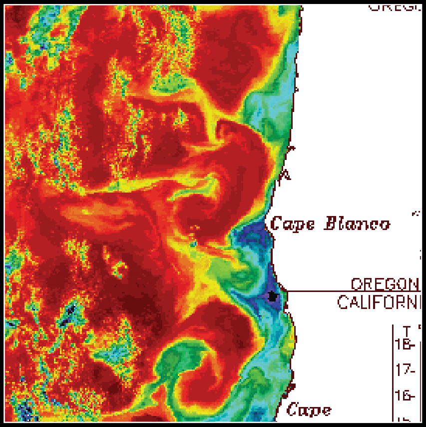

10-km; days

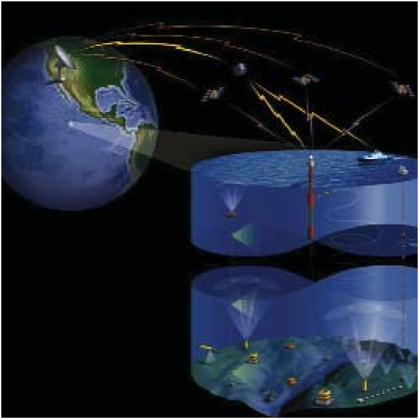

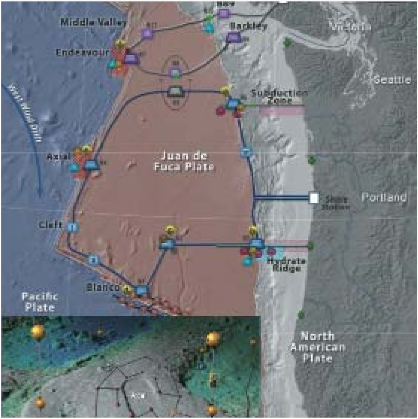

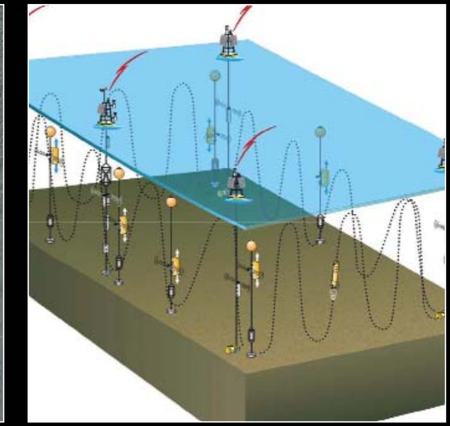

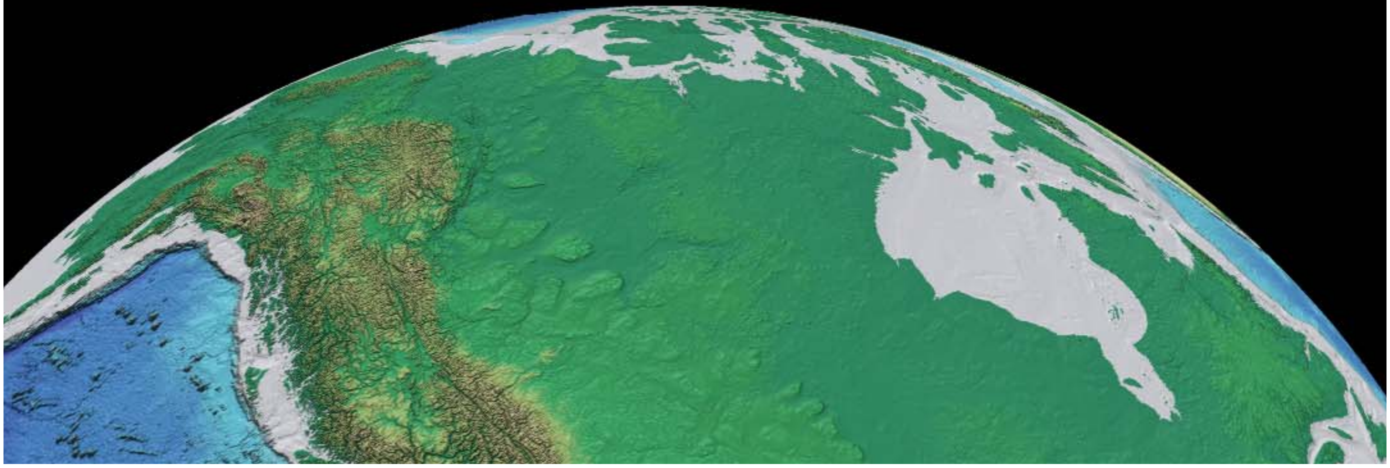

Exhibit 3: Glossary

ORION CI Proposal

actor: an entity (such as a human or a device) external to a system (such as an autonomous underwater vehicle) that plays a role in (interacts with) that system

application: self-contained functionality that performs a specific function directly for a user. An application utilizes resources to achieve an outcome.

catalog: a searchable repository of resource descriptors (e.g., metadata)

discover: the act of locating exposed resources

domain: a part of a system over which rule or control is exercised

enrichment: the process of improving the quality of something by adding something new. In a CI context, this includes declarations of provenance, citation, correspondence or association.

federated: a set of distinct entities joined together into a unified entity

governance: the act of controlling or exercising authority

inform: the process of communicating information or knowledge to an entity

messaging: the process of sending information from a source to a receiver consisting of mediation (i.e., semantic standardization) and a communication (i.e., transmission) steps

ontology: a data model that represents a domain that can be used to reason about objects in the domain and their relationships

persist: to continue in a fixed state

policy: a guiding principle containing the terms of an agreement between a resource provider and a resource user that is described in a language that makes it computationally enforceable

presentation: the process of delivering and formatting information to applications

process: a sequence of activities

protocol: the formal design or plan of an experiment or research activity

publish: the process of making information available, including notification of its availability

query: a user's (or agent's) request for information, generally as a formal request to a data catalog or repository

repository: a container for resources

resource: an entity that provides capability. Resources have a use, an identity, an owner, and their use is subject to policy. Examples of resources include instruments, data, networks, and workflows.

semantic: a conceptualization of the implied meaning of information, that requires words and/or symbols within a usage context.

service: the durable means by which the needs of a resource user are brought together with the capabilities of a resource provider

service oriented architecture: defines standard interfaces and protocols that enable the encapsulation of information tools as services that clients can access without knowledge of, or control over, their internal working

subscribe: the process of receiving information that has been exposed through publication

task: a particular piece of work that is assigned

Exhibit 4.1: Use Scenario

ORION CI Proposal

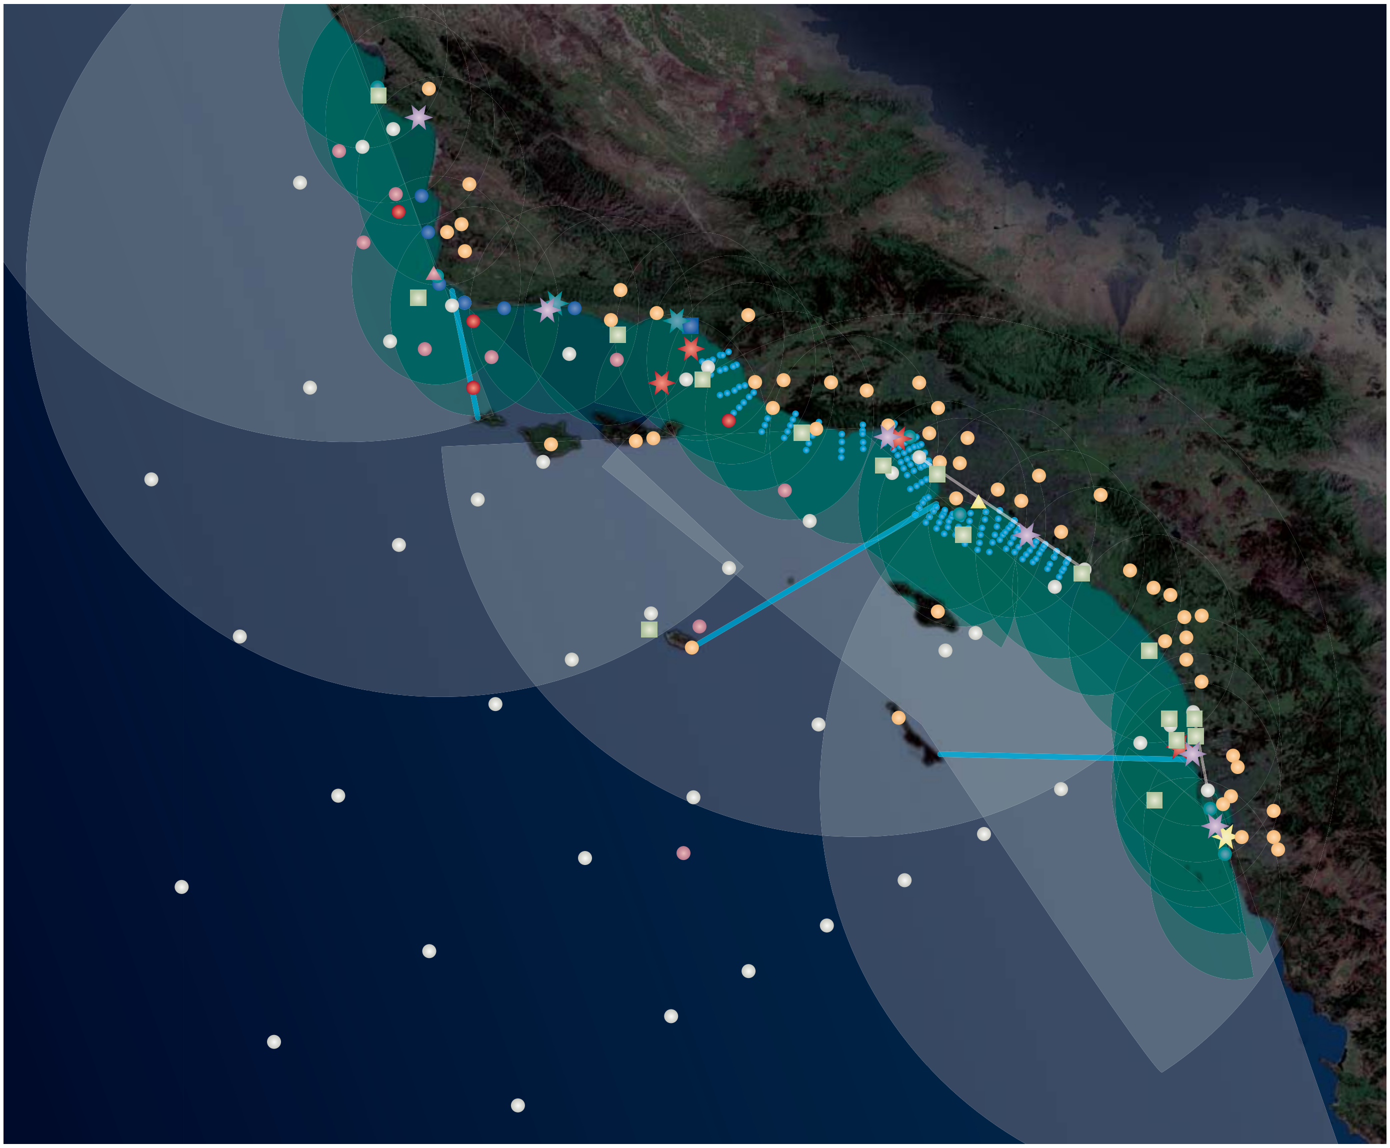

This figure shows a large coastal observatory comprised of long and short range coastal radar (CODAR) nodes and a mix of buoys and glider tracks covering most of offshore southern California. This constitutes a regional framework for coastal sciences processes and events composed of semiautonomous resource nexuses (e.g., discrete buoys). At the node level, resource allocation (e.g., power or bandwidth) is comparatively simple and can be implemented in local hardware or autonomous software. However, coordinating large numbers of nodes into a coherent scientific whole that is larger than the sum of the individual parts is a significant challenge. For example, linking the functionality of CODARs up and down the coast without human intervention is a major science requirement. Management of diverse types of data and their associated metadata is another. CI is needed to provide automatic control of these and other aspects of the overall observatory. In a very real way, the concept of a regional framework is important at the operational as well as the scientific level. One of the major operations and maintenance challenges for a distributed ocean observatory is tracking and coordinating the state of observatory resources. Thus, the science use case is also the operations use case.

Exhibit 4.2: Use Scenario

ORION CI Proposal

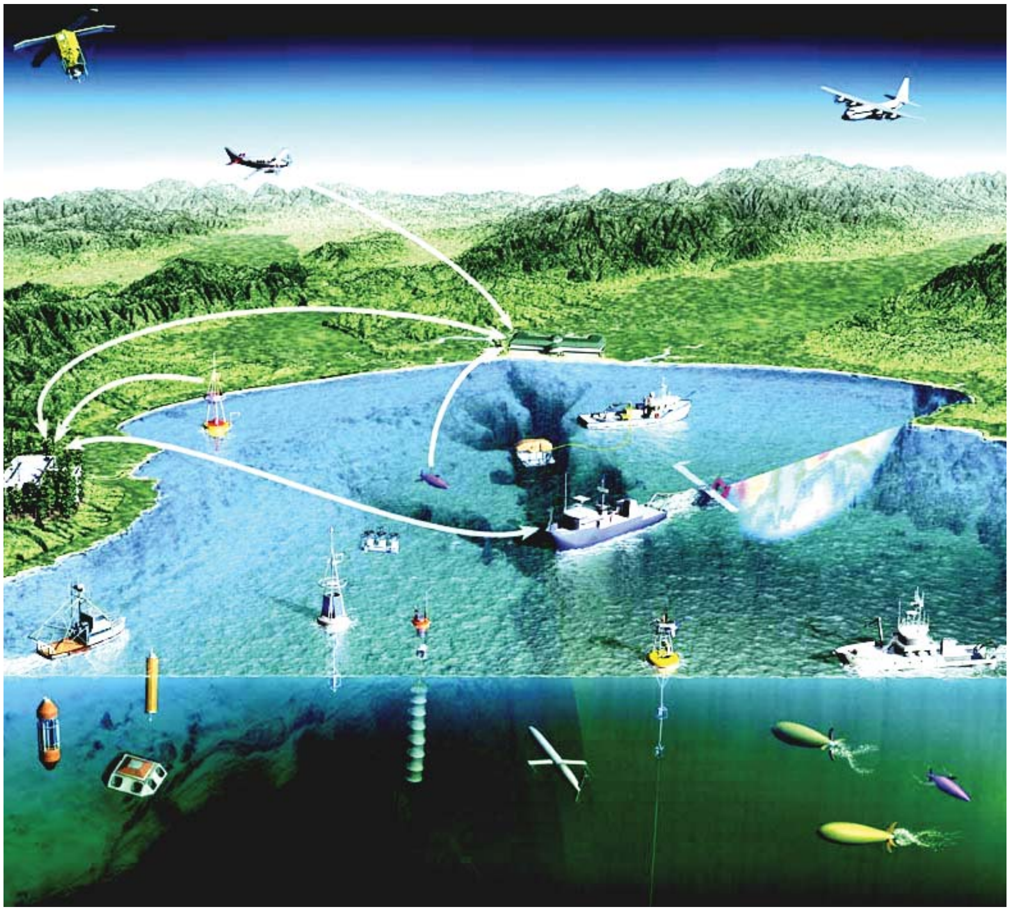

Traditional data assimilation models operate in open loop form, incorporating retrospective or real-time data into the model run without altering the measurement protocols. Dynamic data-driven application systems (DDDAS; Darema, 2005) close the loop by allowing modification of sampling by the assimilation model. In a simple scenario, the assimilation model may change sample rates for selected instruments in response to an event. A more complex scenario has the assimilation model steering instruments on a mobile platform (such as a ship) to locations where property gradients are largest in the simulation. An even more complex scenario shown in the figure might incorporate the addition or removal of fixed or mobile instruments from the domain of interest in response to model output.

Accomplishing a DDDAS scenario with fixed instruments requires a wide range of resource allocation, instrument control, and instrument communication services to coordinate the functionality of the assimilation model, the instrument suite, and the ocean observatory infrastructure. If some of the instruments are mobile or the sensor mix changes with time, then additional services for discovery and localization or tracking may be needed. Cross-cutting requirements for time synchronization and security services also exist. However, the primary communication path in this scenario is between dispersed instruments and terrestrial assimilation models, resulting in a comparatively simple network topology.

Exhibit 4.3: Use Scenario

ORION CI Proposal

A more elaborate use case encompasses many heavily instrumented sites distributed around a regional cabled observatory (e.g., ten or more multidisciplinary moorings extending through the water column). This adds additional complexity through shared use of instruments and resources by multiple users and the difficulty of remote coordination of resources over large distances.

The figure depicts a single science site in this use case, where a diverse suite of sensors and actuators are deployed over a small area (for example, on the scale of a hydrothermal vent field) to accomplish multidisciplinary science. The sensor suite may include physical, chemical, and biological types, and the science mission may require frequent changes in their location or mix. Heavy use of stereo HDTV and high resolution acoustic imaging are anticipated, with concomitant demands on bandwidth and power resources. Acquisition and storage of physical samples for later retrieval and onshore analysis may be needed. Accurate repeat positioning of actuators for sampling may also be required, which imposes closed loop control constraints on the hardware and software infrastructure. This use case imposes stringent demands on the shared use of instruments and resources by many users. Quality of service, latency, and jitter requirements implied by real-time stereo HDTV and closed loop control of sampling actuators are stringent. From a CI perspective, a diverse set of services for resource allocation, time synchronization, instrument monitoring and control, bi-directional instrument communication, cross-calibration, coordination of sensing regimes (e.g., optical or acoustic), localization, tracking, and security are required. Closed loop control may not be feasible in the presence of high seafloor-to-shore latency without CI assistance, such as that used in remote surgery applications.

Exhibit 4.4: Use Scenario

ORION CI Proposal

Looking a decade into the future, the sensor suite at ocean observatory sites of interest will consist of a mix of large numbers of low capability, low cost fixed sensors (e.g., for the measurement of temperature over an area) and small numbers of high capability, high cost sensors (e.g., in situ spectrometers) in mobile platforms. This combination simultaneously accomplishes continuous areal-scale, high resolution and directed, local-scale resolution measurements in an economical fashion. The enabling technology which makes this approach feasible is a network of high bandwidth optical modems [Farr et al., 2006] that provide a wireless extension of the observatory infrastructure, both making it possible to accommodate large numbers of sensors without physically attaching them to the observatory and allowing real-time access to fixed sensors and mobile platforms. The mobile platforms may operate continuously to accomplish pre-programmed sampling missions or under human control for exploratory sampling. Arrays of sensors that fuse into coherent sensor networks are a rapidly evolving application in terrestrial monitoring. This can be accomplished by either linking all sensors to an optical modem network or through pervasive, direct peer-to-peer interconnection Since the characteristics of the terrestrial wireless and seafloor optical environments are similar, it is reasonable to expect both methods to be widely utilized on the seafloor in the future.

This use case aggregates all of the requirements of the previous three scenarios, involving both resource intensive applications and an ever-changing mix of mobile sensors that are complex in their own right and whose operation must be coordinated in real-time. Additional services to provide for discovery of topology and location-aware routing in a time-varying network may be necessary. Sensor networks may also require group management and collaborative information processing applications. A cross-cutting requirement is one of simplicity; for example, low cost sensors with wireless links may not have the capability to process complex time services.

Exhibit 5: ORION CI Conceptual Architecture-Concepts of Operations

ORION CI Proposal

Dr. Adrian Chu of the University of Nebraska Oceanography Center has been working on an analysis tool for some time. The tool integrates several sources of oceanographic data into a small model, and produces a prediction when certain conditions are met. Data from multiple ORION observatories are blended together in the model. Dr. Chu will work with several international researchers from Canada (Dr. Nicole Jones), France (Mlle. Jeanne Fleuris) and Russia (Dr. Dmitri Istantov). He has previously used the ORION cyberinfrastructure to set up his collaborative work group and construct a virtual workspace for them to use. The group has then interactively modified and updated Dr Chu's model and added new features to it. The model has been tested by subscribing to data streams from the ORION observatories. This resulted in further model changes, and Dr Chu and his team are now ready for an operational run. As he starts running the software, although he doesn't fully appreciate it, the ORION infrastructure is performing a lot of steps to make sure the products show up where they are expected.

Configure ORION system to accept data products

Just as Dr. Chu received pointers to subscribe to specific resources, his colleagues Dr. Jones and Dr. Istantov received resource descriptors when configuring the code to publish its observational events and summary data. To obtain this publication resource, the collaborators had to enter information into a publication metadata form that describes the source and nature of their publications. These metadata descriptions help users learn more about data products, assist administrators to troubleshoot any problems, and allows the CI to create a processing history for each of the data products. They are also critical to supporting search functions for the products created by ORION. Because the forms use dropdown menus with controlled vocabularies to fill out most of the fields, and auto-population of subfields based on user selections, all of the members of the team fill out the metadata form consistently and relatively quickly.

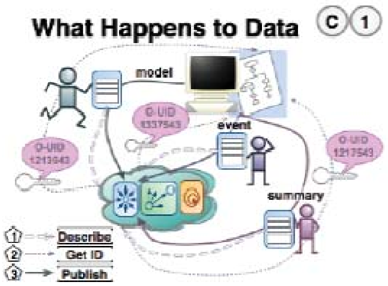

When the software runs, it uses the publication resources to announce to ORION that it is the source of this particular observational event, data stream, or data set. ORION can then connect the people or systems who have sought out and requested these observational events or data.

System performs publication when software runs

As modified by Dr. Jones, Predicitve Ocean Integration Model (POIM) publishes a prediction whenever it detects an observational event. Although she marked this output as an 'observational event', it also has the characteristics of a data stream: it arrives repeatedly over time (not necessarily at a consistent interval), the same type of information is in every record, and it is associated with a single data source, in this case a software process. The additional identification of this record as an observational event serves several purposes: it lets people find the item by searching within a list of publishable observational events, it helps describe the nature of the item (specifically, that arrival of the publication constitutes a message of significance), and it enables general-purpose event-oriented tools (event counters and summarizers, news bulletin generators) to be developed by ORION or other organizations.

Now that the software is executing, observational events will be published on an occasional basis. Each publication is logged by the ORION infrastructure, so that it can be reviewed later in the context of other activities. As described earlier, each publication can be obtained by ORION members in one of several forms: as a subscription, as an email or other notification, upon request (''show me the last observational event of this type''), or in archived form. People who have not registered with ORION can see data products (e.g., the archived logs of observational events), but not the more complicated services.

Publishing data as a stream

Just as the observational events are published (and accessed) as a resource, so too can the data summaries from the model. In fact, this same publication technique can publish any ORION data stream, including those generated by ORION instruments. The key characteristics necessary

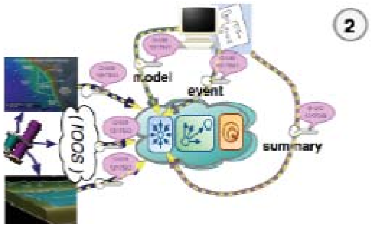

to publish data as a stream are that the data be described in advance, that the data creator (the software or instrument which generates them) use the ORION APIs to submit the data to ORION, and that the resource identifier for the data stream be associated with every data record that is output as part of that data stream. If developers writing software that creates data generation want to take full advantage of ORION's capabilities to integrate, display, and process data—and most developers on ORION are either strongly urged, or required, to do so—they must describe their data in a consistent format, and output it in a way that the format can describe. If a data source like a GPS (or modeling software) actually generates multiple types of data records (for example, one data record, one summary record, and one error record), then the developer must create a separate description for each record, get a separate resource identifier for each record, and publish each record type along with the appropriate resource identification. While this seems like a lot of work up front, it usually is fairly straightforward and saves a lot of time in post-processing the data streams.

In this case, Dr. Chu's colleagues have used these features well, and Mlle. Fleuris in particular quickly understood the process of describing her outputs from the model. She created a metadata description for the model summaries she produced, defining the meaning of each item in the summary and the data source used to present it. Unfortunately, Dr. Chu's model output, which is the data source for her summary, is itself unpublished, since he is keeping it private for now. She plans to suggest to Dr. Chu that the model itself be published as an ORION resource, so that users can trace the sources for these summaries and predictions back through the entire chain of operations in the ORION workflow system. For now, she has referenced the unpublished data by description, as well as pointing back to the observational data streams that Dr. Chu's model uses.

Data Stream Archival

Mlle. Fleuris set up the publication of the model to occur once every hundred times the model runs, as well as every time the model generates an observational event prediction. This allows the team to review the operation of the system over time and contrast its operation in predictive and non-predictive cases. Since the model runs hundreds or even thousands of times a day, this technique should limit the output to only a few outputs each day. This output volume is not very large, and the ORION infrastructure will respond accordingly by archiving them for an extended period. The holder of any reference to an ORION data stream can ask to view the data's historical records, as Dr. Chu did for the other data he wanted to review. If the reference holder has permission to view the data, they can be obtained from the ORION operational data archive. At this point, the events and model summaries can be viewed on-line or from the archive by the collaborators on the team.

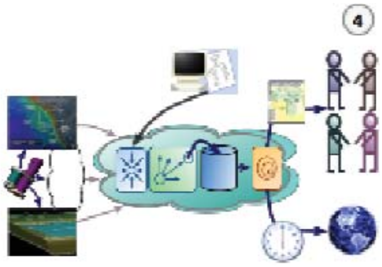

When the verification period (a period set by ORION policy, during which only proprietary access is allowed, so that the data can be evaluated and tested) expires, the data will be available to the public. At first Dr. Chu found this idea to be disturbing, but he has gotten used to it since he wants to use the full capabilities of the ORION.

Publishing as publicly available data

In fact, Dr. Chu expects he will make these data products—the events and the summaries, at least—publicly accessible well before the validation period expires. This takes a minor effort on his part, and he knows a lot of colleagues will want to take advantage of the resulting predictions for their own studies. As an enlightened act of self-promotion, he intends to make the results available with a request to acknowledge him on any papers that ensue. While he knows he may only be acknowledged on half of the papers that use the work, his name will still become widely known as the originator of the information. From his previous experience in publishing data from an instrument, Dr. Chu knows there are several steps required to make his results publicly available, including certification and verification. First, he must certify that the data source meets the standards described in the ORION service agreements. For software, this is little more than has already been specified in the metadata, along with running the software on an ORION test bed system. Obviously, standards for instruments to be deployed at ocean depths are somewhat more demanding. The observatory on which the data source is deployed will confirm that the interface specifications have been met. This is done automatically for software, and with some manual confirmation for hardware interfaces. A further step required before releasing software is the verification step. This consists of evaluating the results from the data source to confirm that it is operating as expected. As Dr. Chu has already accomplished this step to his own satisfaction, reviewing the ORION products from his system should be simple. He is prepared to quickly go to the trouble of releasing his data to a wider audience, and establishing its verified status on ORION. For core instruments on ORION observatories, more detailed criteria must be met, including verification that the metadata describing the data source are correct and QA and QC procedures are in place.

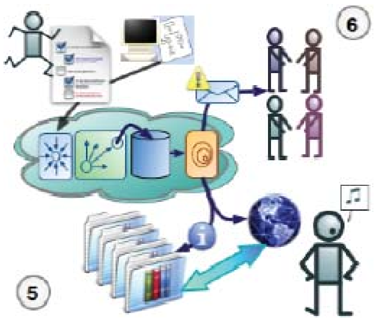

There are several advantages to publishing the data—in this case, events, notification of an event, and model summaries—to a wider audience. First, it makes the data immediately available to the public through the data products that ORION produces. It also makes the services easier to reference and use within ORION—while this could also be achieved by changing the access permissions on his data sets, making the data public automatically changes those access permissions. Further, it makes it clear that Dr. Chu has reviewed the data sources and believes they are functional. Finally, making his data public advertises his products to a wider audience, since the ORION data product registries will only replicate complete metadata descriptions for a data product if that product is in fact publicly accessible. Once the ORION data product registries announce their availability, Dr. Chu's results will become visible in four other data publication registries (three of which are internationally well known), and he will get extra credit and attention for his work.

Dr. Chu has been following some interesting developments related to the publication of metadata in external registries. Some scientists have been quoting as part of their "publication rate" the number of entries they have in data product registries, and some search tools have begun indexing the registries as a way to provide more contextual information about data sources, data owners, and data systems. As a result, the "free registration" that ORION provides will likely have benefits for Dr. Chu's work.

Publishing notifications

Dr. Istantov wants to email a notification to each member of the team whenever the software detects an event, and he has used an almost identical mechanism as the others. Some of the metadata for his "data stream" are different, but much of the cyberinfrastructure used for publishing the notifications is the same as for events and other data streams. In fact, although he didn't realize it, Dr. Istantov's metadata form was made easier to fill out because Dr. Jones and Mlle. Fleuris had filled out almost identical ones earlier that was used to pre-populate some of the fields on Dr. Istantov's form. While he was testing his code, he sent the notifications to himself, but after completing his code changes, he updated the distribution list. Because the notification message is published via email, Dr. Istantov can select the destinations from a number of email address lists, including a list of aliases, of actual users, and of virtual laboratories of which he is a member. He configures the email destination for this published message to be Dr. Chu's newly created virtual laboratory, and awaits further word.

Exhibit 6: Project Master Schedule

ORION CI Proposal

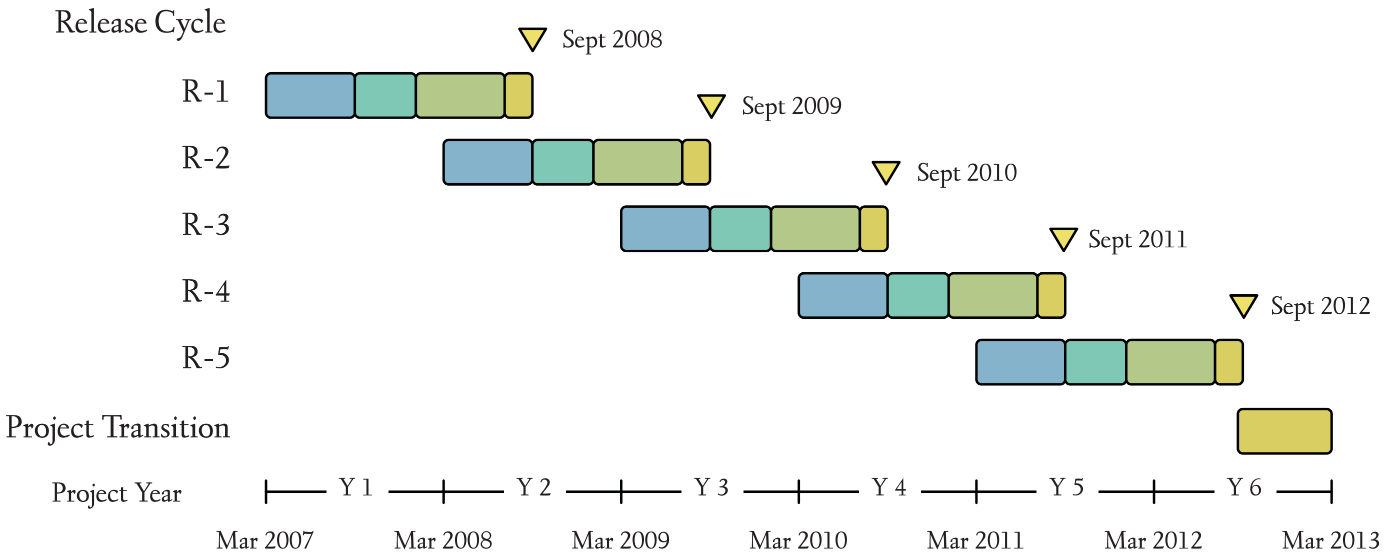

Exhibit 7: Hyperlinks to Cited Projects, Standards and Technologies

ORION CI Proposal

ActiveMQ	http://www.activemq.org/site/home.html
AMQP (Advanced Message Queuing Protocol)	http://www.amqp.org/
BIRN (Biomedical Informatics Research Network)	http://www.nbirn.net/
BPEL	http://en.wikipedia.org/wiki/BPEL
CANARIE (Canada's advanced Internet development organization)	http://www.canarie.ca/
CIMA (Common Instrument Middleware Architecture)	http://www.instrumentmiddleware.org/metadot/index.pl
Concept Design Review	http://www.orionprogram.org/capabilities/cdr/Final_OOI_CDR_Report.pdf
Conceptual Architecture	http://www.orionprogram.org/organization/committees/ciarch/default.html
Conceptual Architecture	http://www.orionprogram.org/organization/committees/ciarch/default.html
Conceptual Network Design	http://www.orionprogram.org/capabilities/cnd.html
Cyber-Physical	http://varma.ece.cmu.edu/cps/
DoDAF	http://www.defenselink.mil/nii/doc/DoDAF_v1_Volume_I.pdf
DSpace	http://www.dspace.org/
Enterprise Service Bus (ESB)	http://en.wikipedia.org/wiki/Enterprise_service_bus
Fedora	http://www.fedora.info/
FGDC (Federal Geographic Data Committee)	http://www.fgdc.gov/
GCMD (Global Change Master Directory)	http://gcmd.nasa.gov/
GEON (Geosciences Network)	http://www.geongrid.org/
Grid Security Infrastructure	http://www.globus.org/security/overview.html
GRIDCC (Grid Enabled Remote Instrumentation with Distributed Control and Computation)	http://www.gridcc.org/
GridShib	http://gridshib.globus.org/
Groovy	http://groovy.codehaus.org/
Hibernate	http://www.hibernate.org/
HOPS (Harvard Ocean Prediction System)	http://oceans.deas.harvard.edu/HOPS/
HPWREN	http://hpwren.ucsd.edu/
Integrated Ocean Observatory System	http://www.ocean.us/
IOOS (Integrated and Sustainable Ocean Observing System)	http://www.ocean.us/
IPERF	http://dast.nlanr.net/Projects/Iperf/
JDDAC (Java Distributed Data Acquisition and Control)	https://jddac.dev.java.net/
JMX	http://java.sun.com/javase/technologies/core/mntr-mgmt/javamanagement/
Kepler	http://kepler-project.com
Laboratory for Autonomous Marine Sensing Systems	http://acoustics.mit.edu/faculty/henrik/LAMSS/laboratory_for_autonomous_marine_sensing_systems.html
LAS (Live Access Server)	http://ferret.pmel.noaa.gov/Ferret/LAS/
LEAD (Linked Environments for Atmospheric Discovery)	http://lead.ou.edu
LOOKING	http://lookingtosea.ucsd.edu/
LOOKING Annual Report Year One	http://lookingtosea.ucsd.edu/library/LOOKING-Y1-AnnualReport-200507.pdf/view
LOOKING Annual Report Year Two	http://lookingtosea.ucsd.edu/library/LOOKING-2006-annualRpt.pdf/view
Marine Metadata Interoperability	http://marinemetadata.org/
Matlab	http://www.mathworks.com/
MMI (Marine Metadata Interoperability)	http://www.marinemetadata.org/

Exhibit 7: Hyperlinks to Cited Projects, Standards and Technologies

ORION CI Proposal

Monterey Ocean Observing System (MOOS)	http://www.mbari.org/moos/
Nagios	http://www.nagios.org/
NEON (National Ecological Observatory Network)	http://www.neoninc.org/
NetCDF (Network Common Data Form)	http://www.unidata.ucar.edu/software/netcdf/
North-East Pacific Time-series Undersea Networked Experiments (NEPTUNE)	http://www.neptunecanada.ca/
NVO (National Virtual Observatory)	http://www.us-vo.org/
OBIS (Ocean Biogeographic Information System)	http://www.iobis.org/
OGC	http://www.opengeospatial.org/
Open Science Grid	http://www.opensciencegrid.org/
OPeNDAP (Open-source Project for a Network Data Access Protocol)	http://www.opendap.org/
Optiputer (Optical networking, Internet Protocol Computer)	http://www.optiputer.net/
OSG (Open Science Grid)	http://www.opensciencegrid.org/
OurOcean Portal	sa.gov/
PicoContainer	http://www.picocontainer.org/
Real-time Observatories, Applications, and Data management NETwork (ROADNet)	http://roadnet.ucsd.edu/
Riverglass	http://www.riverglassinc.com/
ROCKS	http://www.rocksclusters.org/wordpress/
ROMS (Regional Ocean Modeling System)	http://marine.rutgers.edu/po/index.php?model=roms
RRDtool	http://oss.oetiker.ch/rrdtool/
SEEK (Science Environment for Ecological Knowledge)	http://seek.ecoinformatics.org/
Sensor Web Enablement Suite of Standards	http://www.opengeospatial.org/projects/groups/sensorweb
Service-Oriented Architecture	http://en.wikipedia.org/wiki/Service-orientated_architecture
Shibboleth	http://shibboleth.internet2.edu/
Spring	http://www.springframework.org/
Standard for a Smart Transducer (sensor and actuators) Interface	http://ieee1451.nist.gov/
Standard for Precision Clock Synchronization Protocol for Networked Measurement and Control Systems	http://ieee1588.nist.gov/
Storage Resource Broker	http://www.sdsc.edu/srb/index.php/Main_Page
System Engineering Handbook Version 3 issued by the International Council on System Engineering	http://www.incose.org
Telescience Project	http://telescience.ucsd.edu/
Teragrid	http://www.teragrid.org/
THREDDS (Thematic Realtime Environmental Distributed Data Services)	http://www.unidata.ucar.edu/projects/THREDDS/
USArray Array Network Facility	http://anf.ucsd.edu/

Exhibit 8: ORION CI Conceptual Architecture Observatory Network Model

ORION CI Proposal

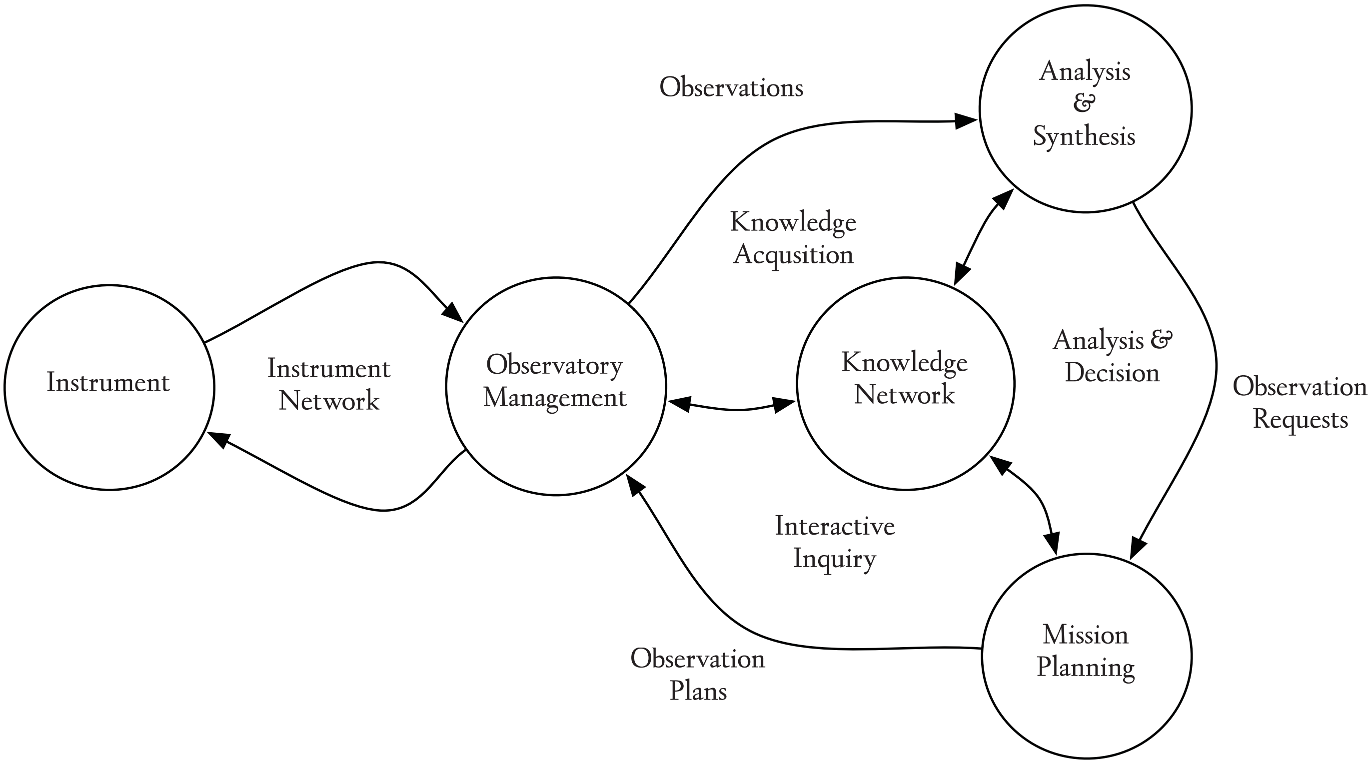

Exhibit 9: The OOI Construction Program for ORION

ORION CI Proposal

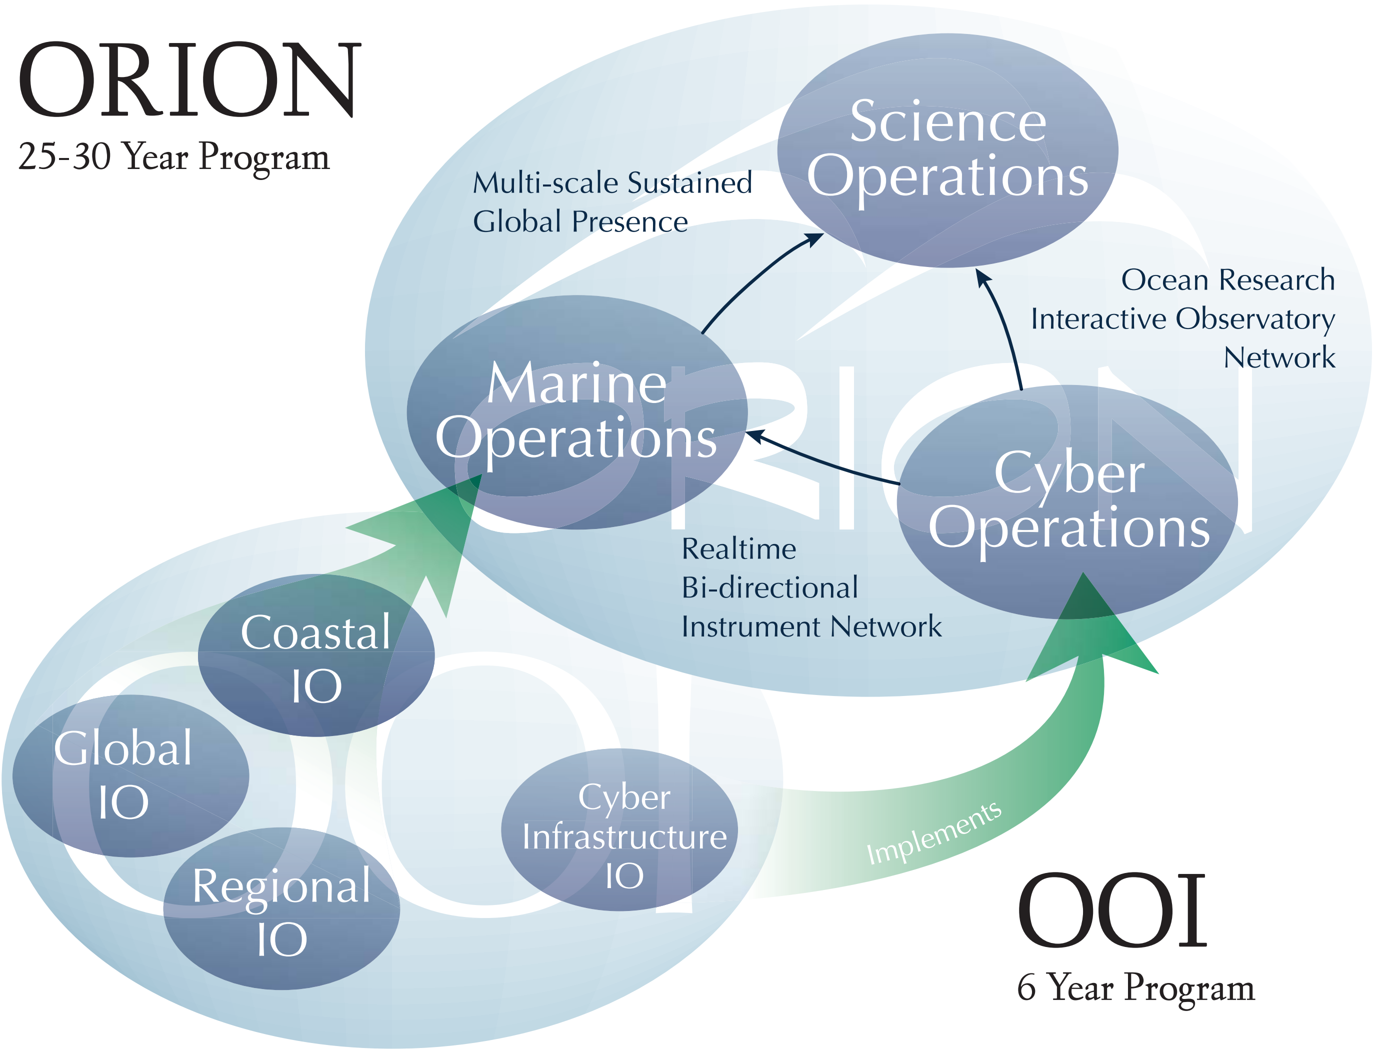

Exhibit 10: Cyberinfrastrucure Points of Presence

ORION CI Proposal

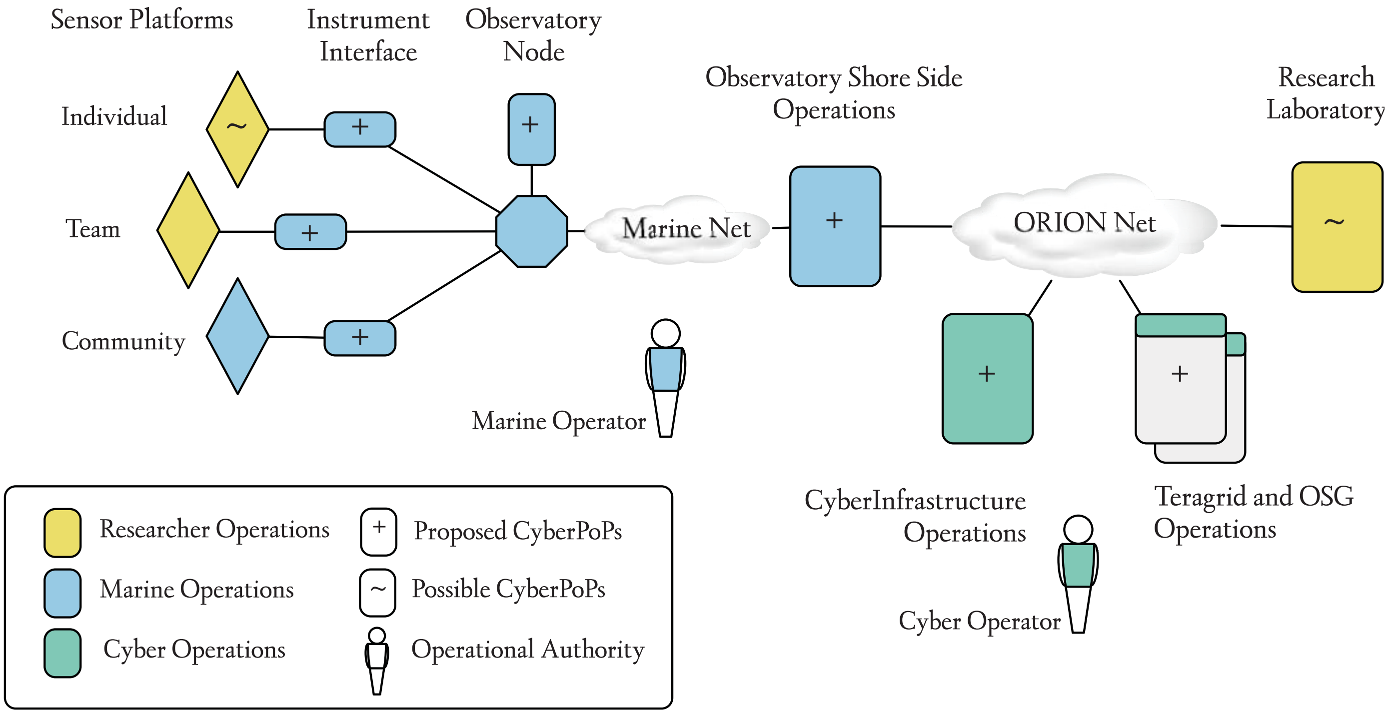

Exhibit 11: Scientific Investigation Process

ORION CI Proposal

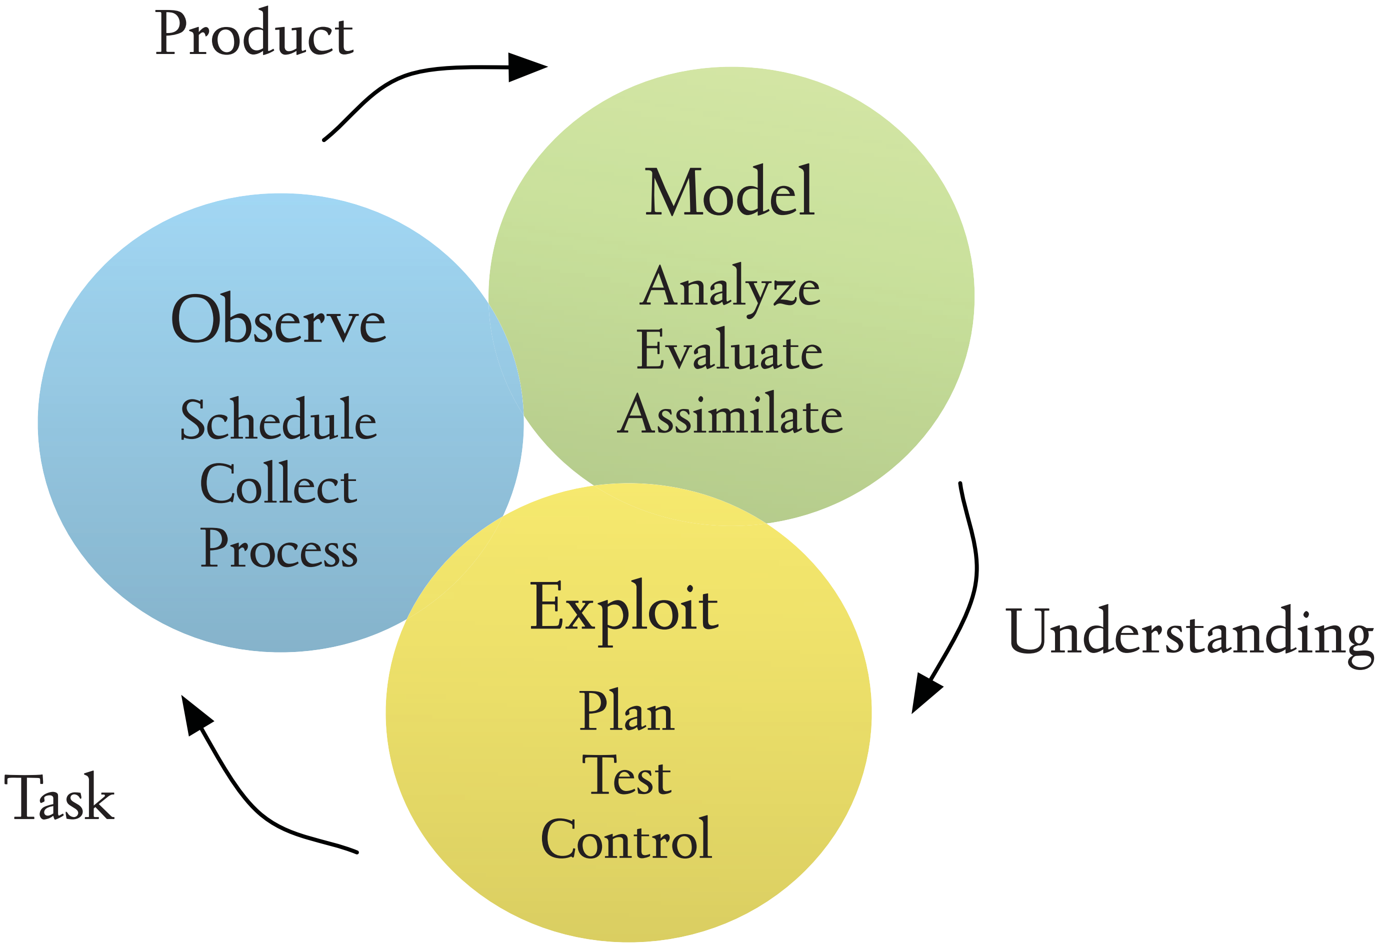

Exhibit 12: Resource Life Cycle

ORION CI Proposal

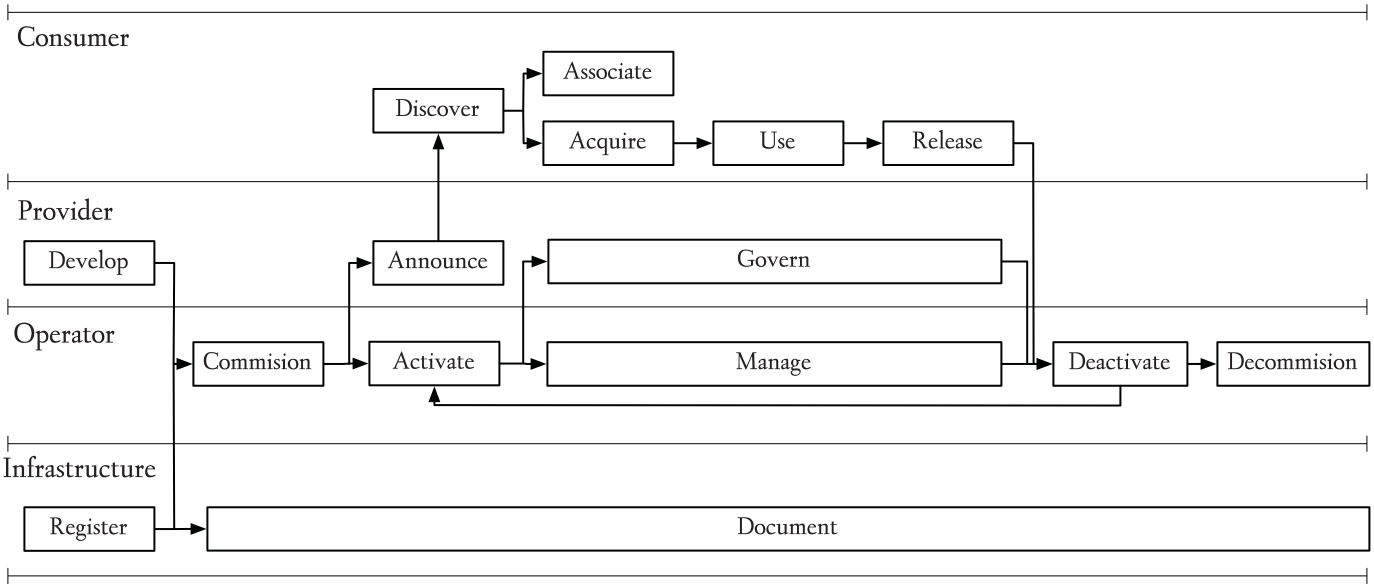

Exhibit 13: Activity Resource Model

ORION CI Proposal

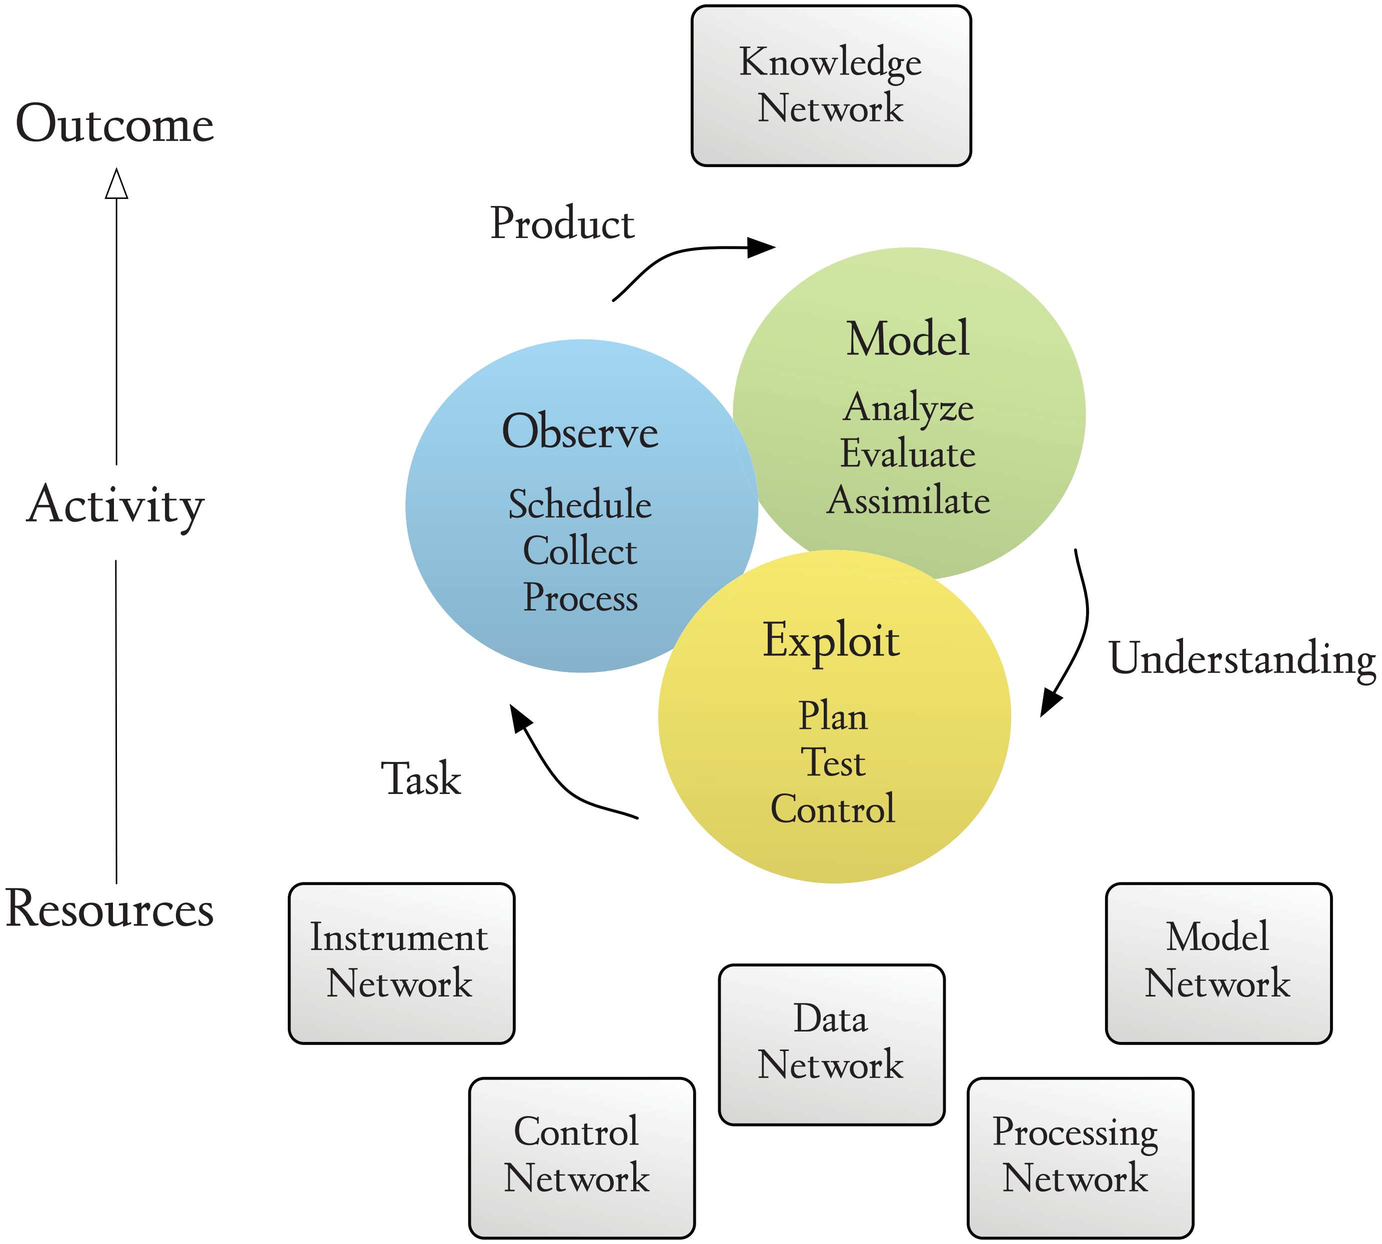

Exhibit 14: Spiral Release Cycle

ORION CI Proposal

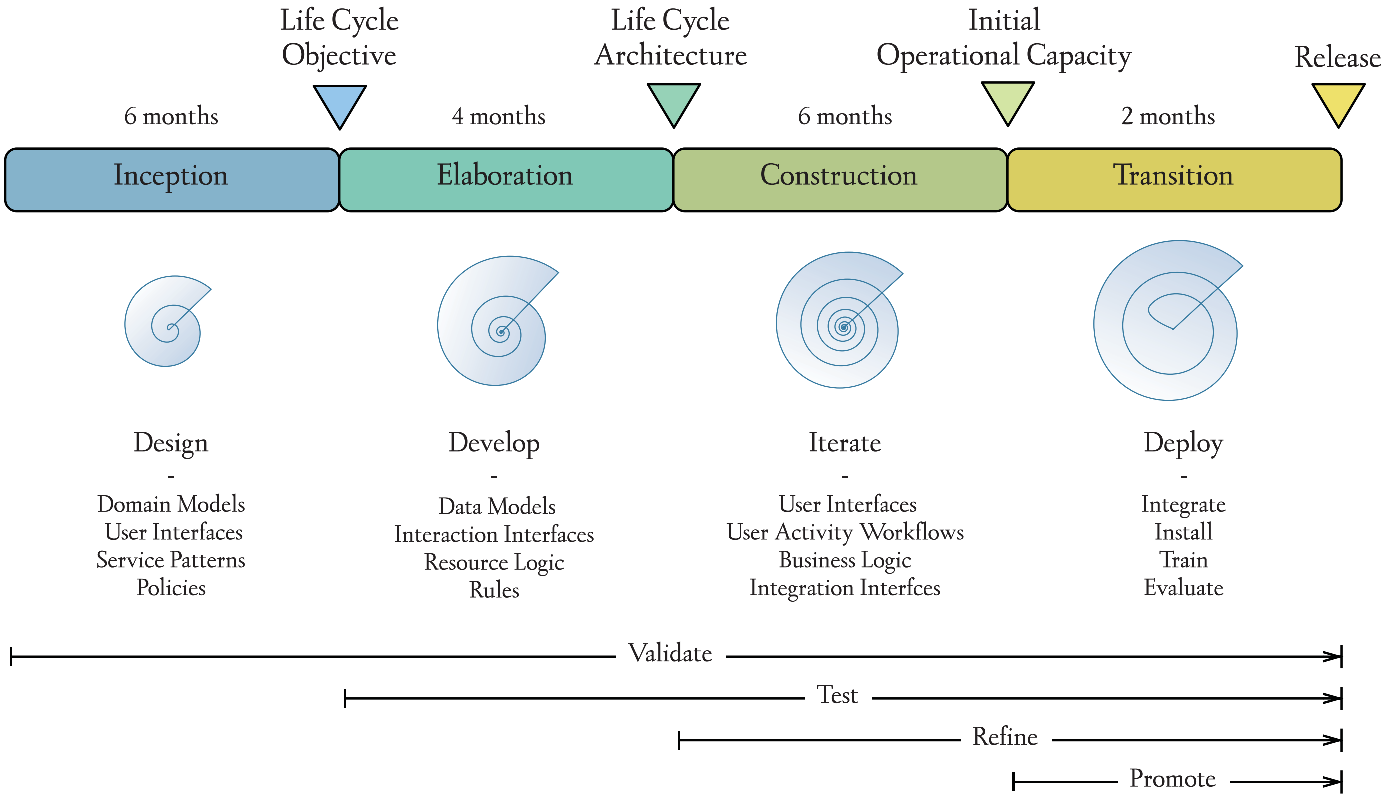

Exhibit 15: Global Observatory Deployment Model

ORION CI Proposal

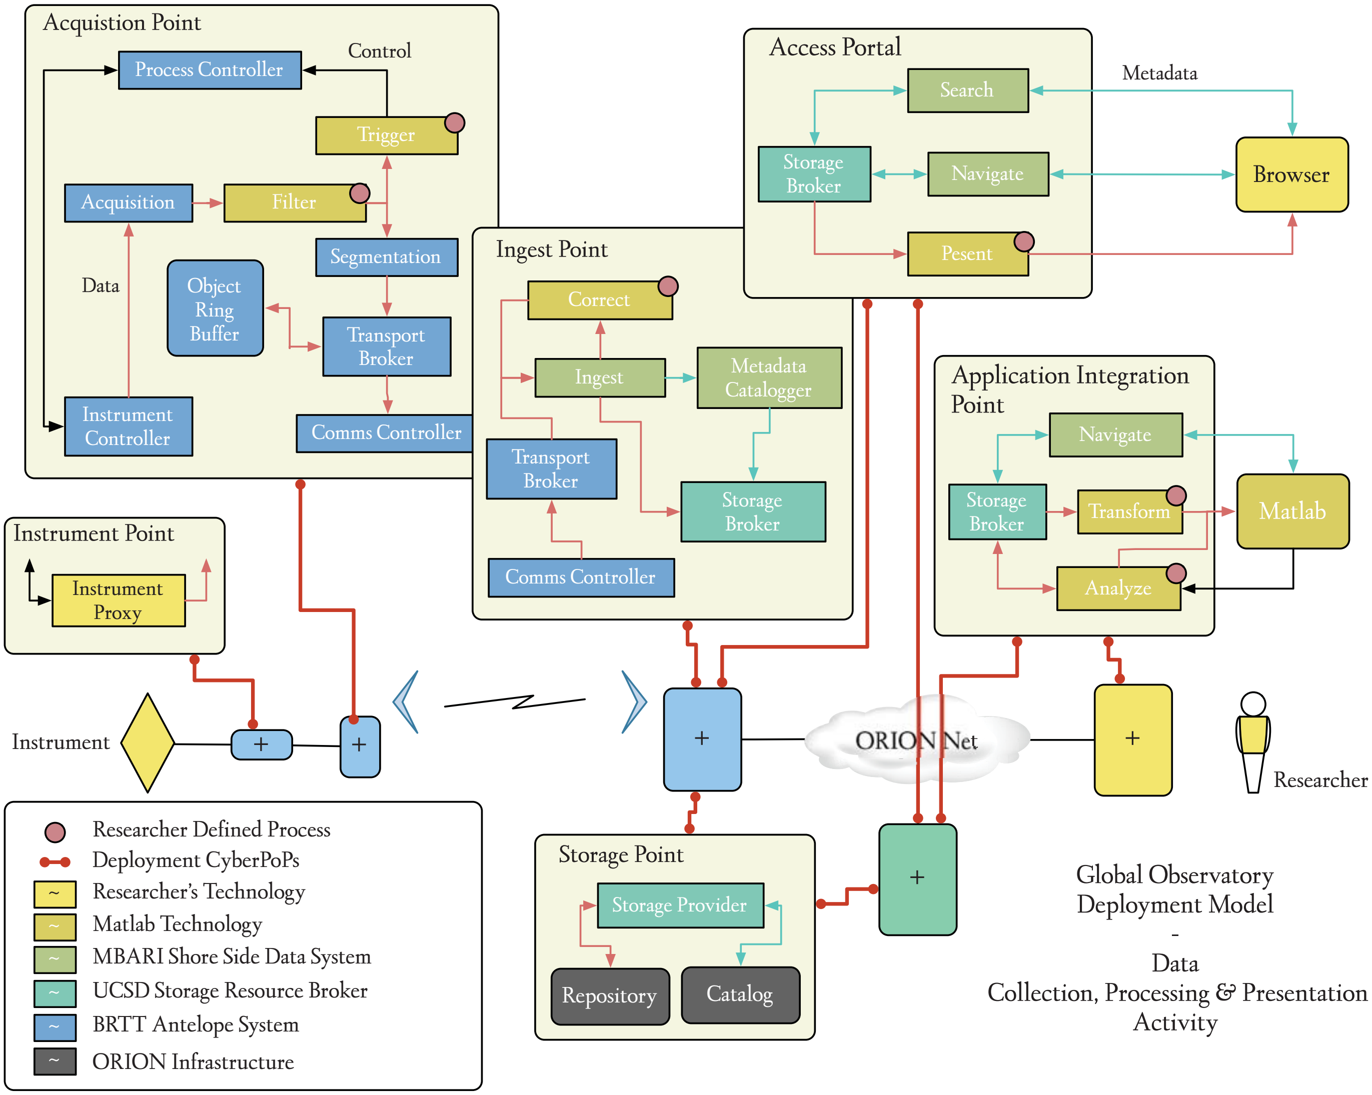

Exhibit 16: Cyber Capability Container

ORION CI Proposal

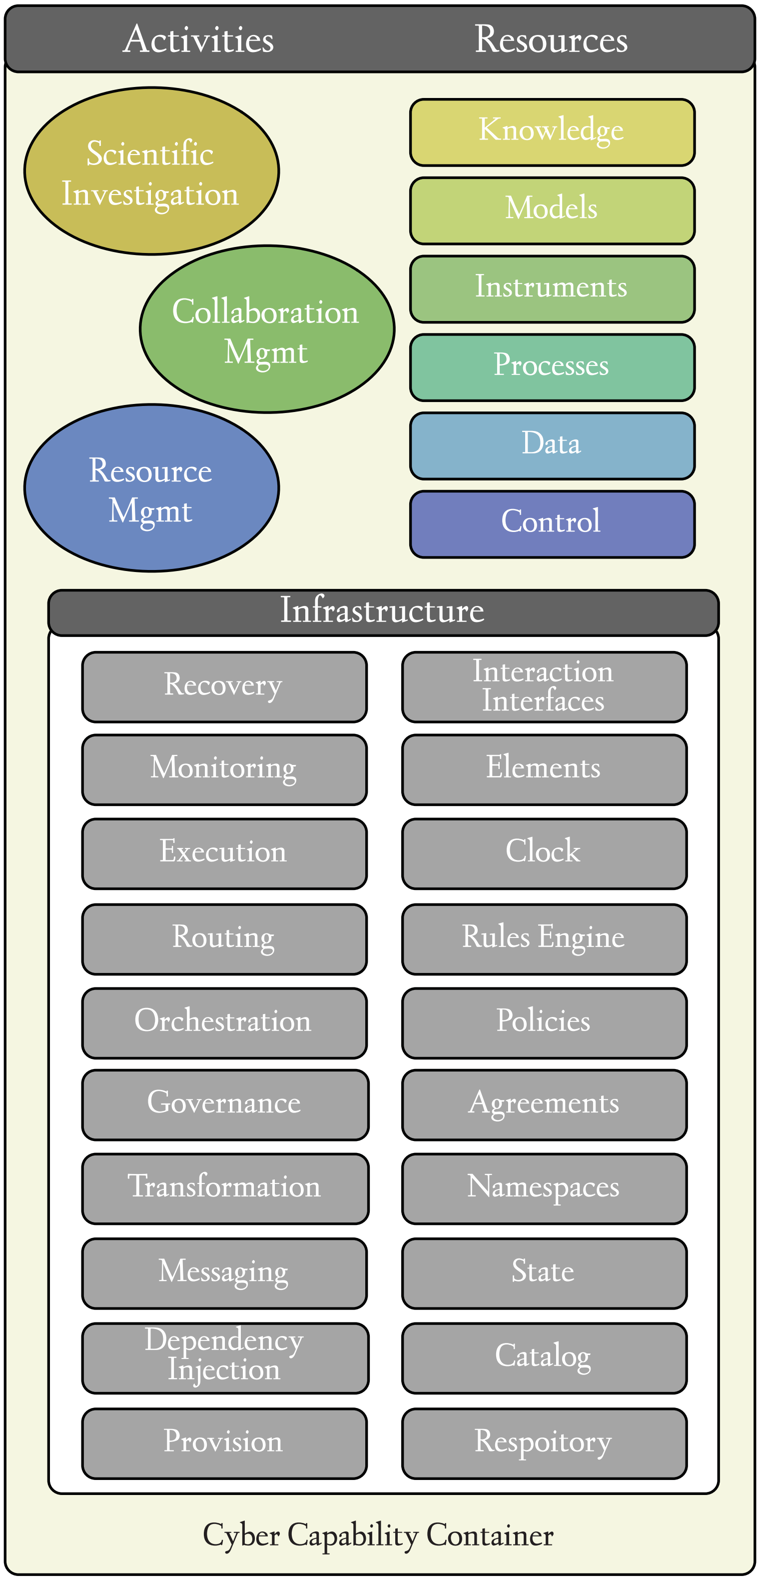

Exhibit 17: Global Observatory Capability Model

ORION CI Proposal

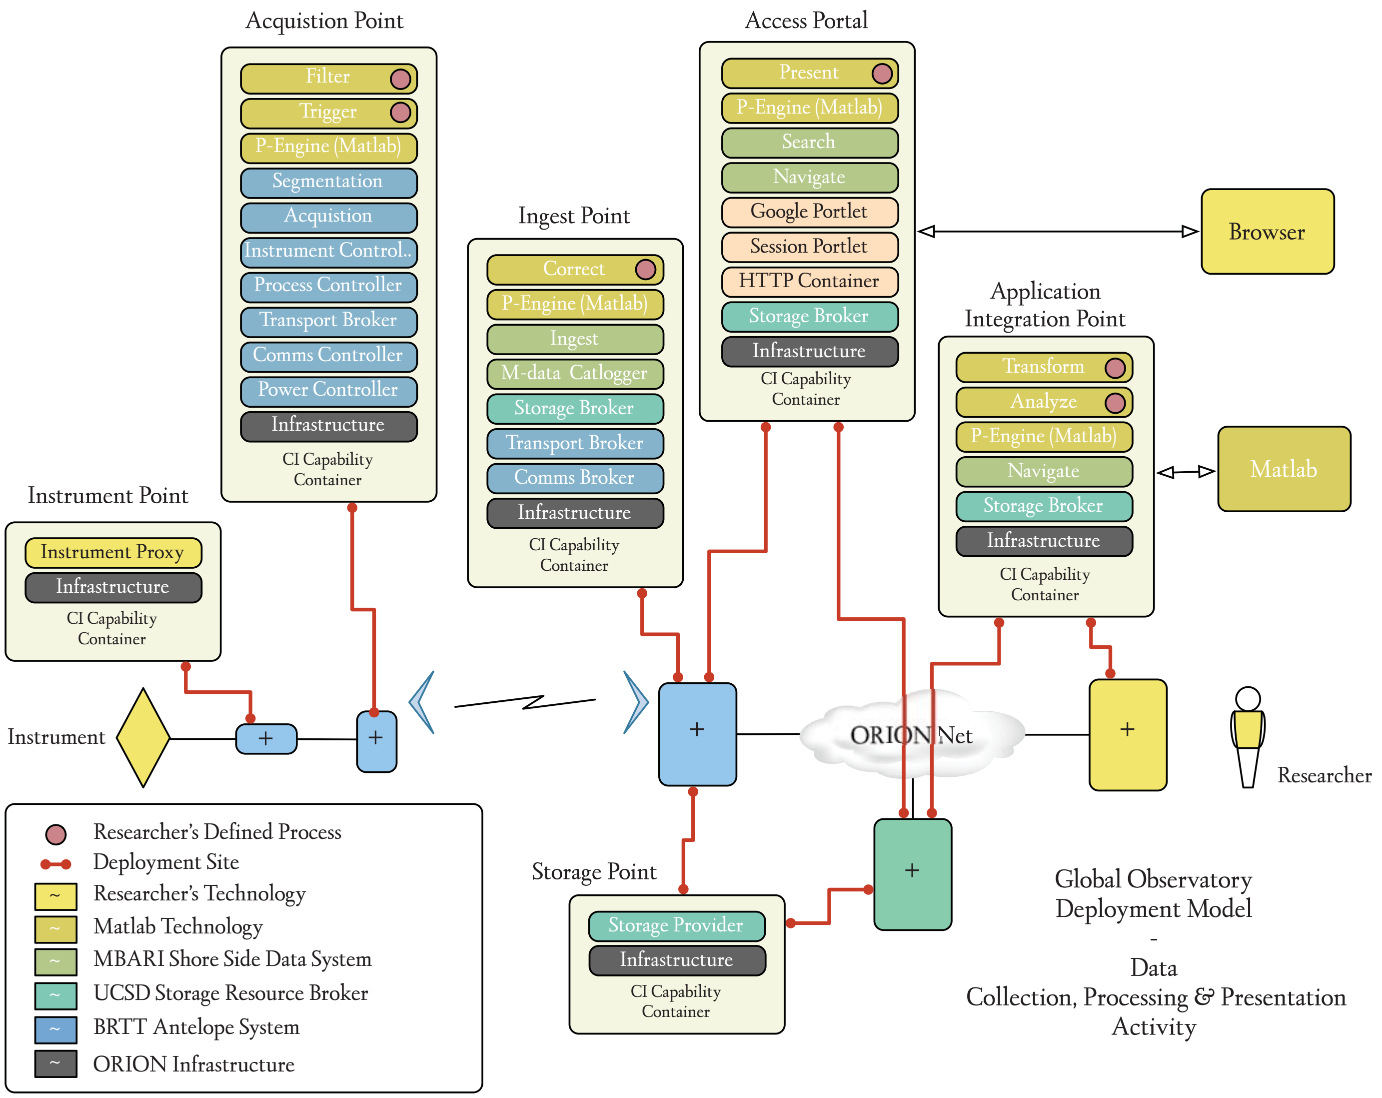

Exhibit 18: Fractal ESB Design Pattern

ORION CI Proposal

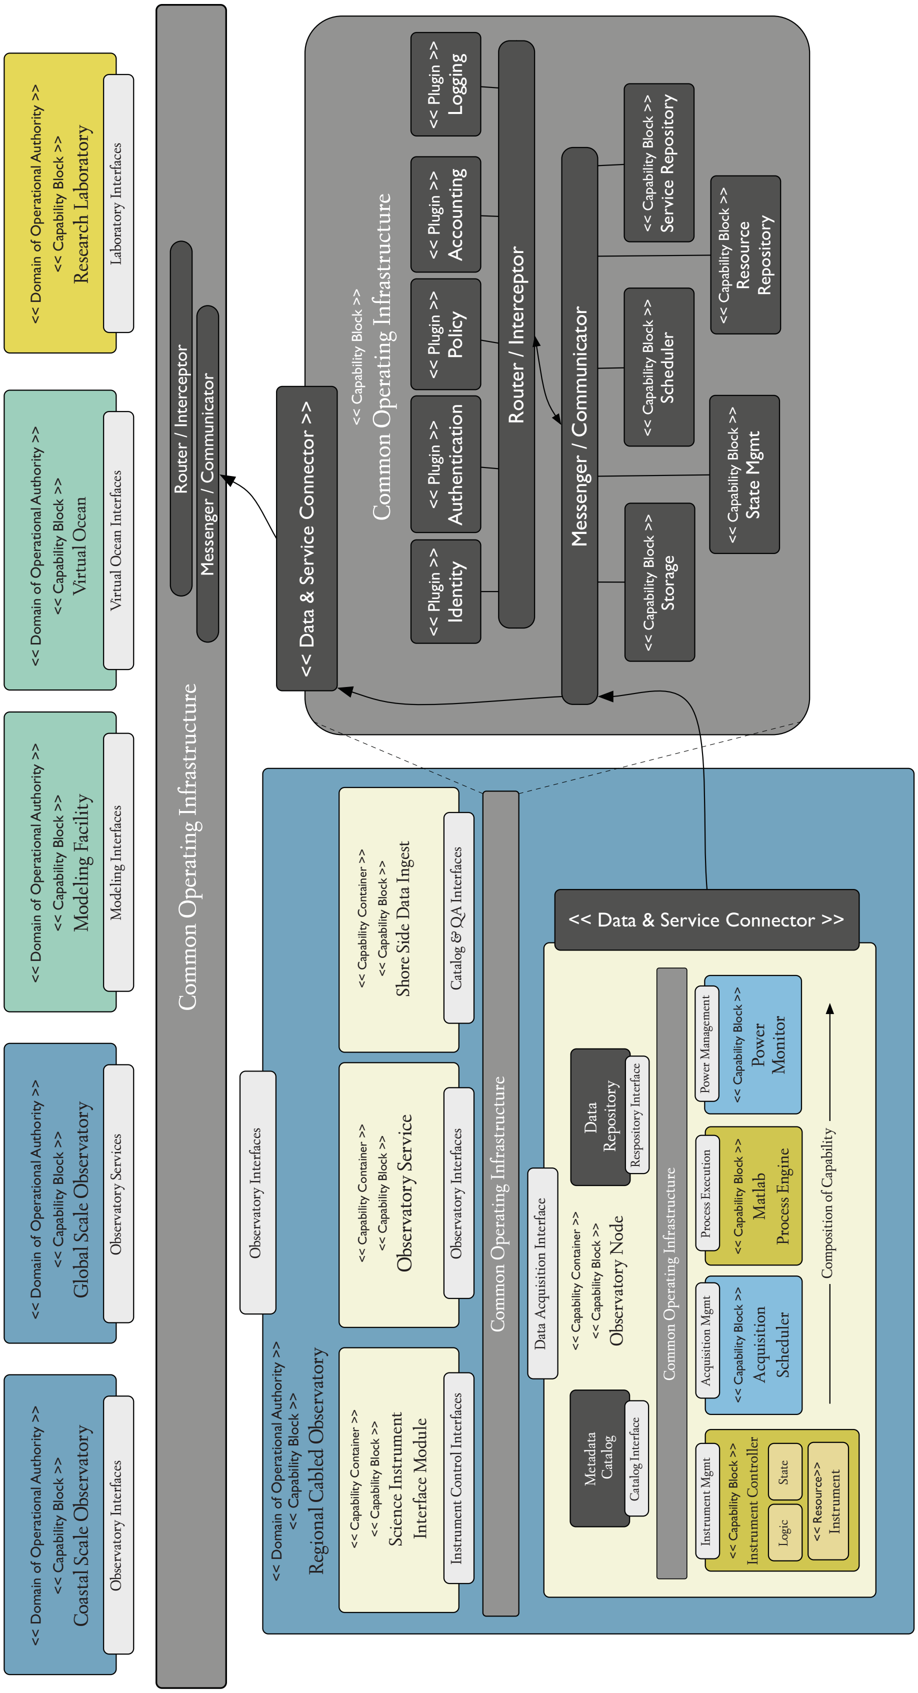

Exhibit 19: Collaboration and Policy Framework

ORION CI Proposal

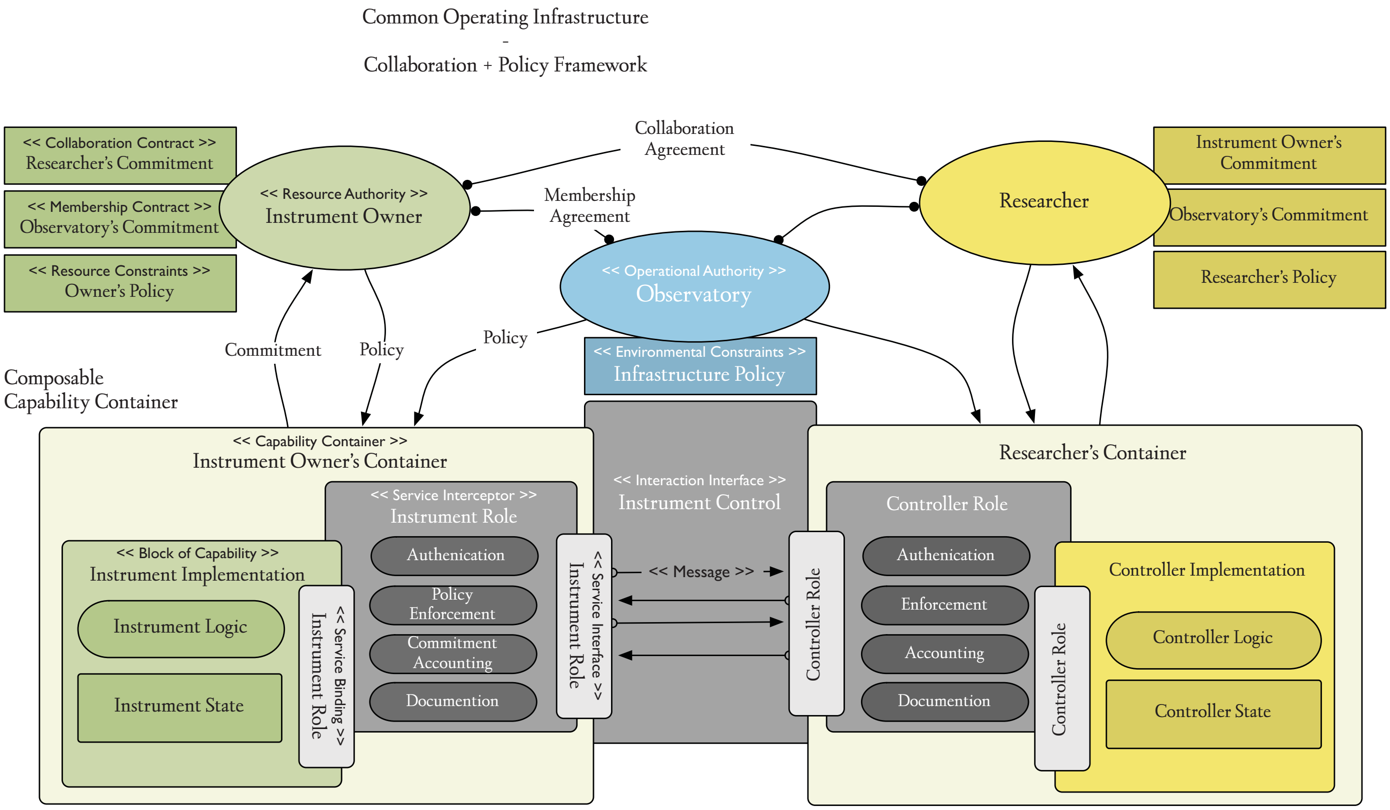

Exhibit 20.1: Architectural Elements
Common Operating Infrastructure (COI)
ORION CI Proposal

Common Operating Infrastructure (COI)

The COI provides the technologies and services to play the role of 1) a unifying information conduit, enabling data and control streams to be published and consumed by all of the subsystems; 2) a platform to execute the core elements of the activity model by allowing the subsystem services to be combined as workflows; and 3) the implementation location for cross-cutting aspects of the CICC model.

Capabilities

1. Collaboration provisioning and agreement management,
2. facility provisioning and rights management,
3. identity validation and verification,
4. Service provisioning and interchange management,
5. federation and delegation of service presentation and fulfillment,
6. resource collection management, navigation and search,
7. resource lifecycle management,
8. policy enactment and enforcement management, and
9. communication provisioning and interchange management.

Technologies & Systems

Following on Section 1.6.4.3, the COI will be realized as a composite capability block providing 1) an integration platform for data and control channels (streams), block data transfer and streaming media, 2) a rich set of options for integrating heterogeneous data sources and applications using a variety of data transports, and 3) an interface for injecting infrastructure services, such as policy monitoring and enforcement services, as plug-ins to effect federated authentication and security policies. This is the basis for provisioning the service (registry, brokering, binding and execution of services), facility and resource networks within the COI. The services and their corresponding data models provide a uniform mechanism to detect and exploit the capabilities of ORION entities. Governance will be supported through templates for collaboration agreements that define partnership, and delegation and federation policies for participants and their assets. In particular, we will establish an integrated policy/security framework that is directly coupled to the interaction interfaces capturing the activity models. Design partners NCSA and NCSU have developed a comprehensive architectural framework for policy and security management that ties in directly with the notion of interaction interface. Policy and security properties are stored together with the interaction protocol that defines a service using the zone federation architecture of the SRB. The proposed work will inject this framework into the capability block pattern such that all interactions with the COI directly fall under the governance of this policy framework. Specifically, the ESB dependency injection mechanism and GridShib, GridGrouper, and myVocs integrated with SRB's Zones will be used as the core technology for implementing and enforcing the COI data structures and models for governance, security and other policies.

In the initial deployment, the capability blocks will be populated as follows: the messaging component will be instantiated to ActiveMQ/AMQP and the router/interceptor and service/data interfaces will be instantiated to Mule. Mule's dependency injection mechanism provides immediate access to persistence, application/transaction, workflow, configuration and monitoring frameworks that will be instantiated to Hibernate, Spring, Groovy and JMX, respectively. Furthermore, the COI will leverage the successful CI software stacks of BIRN/TeleScience with their web-service-based ATOMIC interfaces to the national Grid computing and security infrastructure. All of these web services can easily be configured as capability block plug-ins. The flexibility of this ESB-based approach allows the development team to rapidly integrate new capabilities as they become available.

The transport-transparent messaging component of the ESB will be exploited to implement data and control streams among CI subsystems, and provision and broker any service, data source, data transport and delivery mechanism, as well as any policy that is injected into the routing/interceptor facility of the ESB. Such a messaging system supports secure, durable, fault tolerant and high availability connections.

Exhibit 20.1: Architectural Elements

Common Operating Infrastructure (COI)

ORION CI Proposal

Data Structures & Models

The proposed architecture will encompass authoring, implementing and maintaining the key conceptual model underlying the patterns of governance that reflect security considerations as well as the ORION activity model. These patterns will be stored and made accessible via a Data Network repository for high-quality governance management.

Interface Points

The COI is the communication and integration substrate for the proposed system; it, therefore, interfaces with all subsystems, as well as with the environment via the data/service interface of the ESB.

Exhibit 20.2: Architectural Elements

Common Execution Infrastructure (CEI)

ORION CI Proposal

Common Execution Infrastructure (CEI)

The CEI will provision the services required to implement an elastic compute network together with a corresponding management UI module. This constitutes the computation and execution substrate for the entire CI.

Capabilities

1. Virtualized computing resource provisioning, operations and maintenance,
2. provision parameterized configurations of service and application modules into compute node deployment packages,
3. monitor and provision compute nodes based on compute resource utilization and latency of provisioning,
4. on-demand scheduling of processes,
5. optimized scheduling of stream process subscriptions,
6. extendable process execution environment that supports multiple execution formats,
7. federation of process execution service providers, and
8. process control interface.

Technologies & Systems

The technical components of the CEI include a virtual compute node configuration repository, a virtual compute node loader, a virtual compute node deployment scheduler, and a base virtual compute node; the implementation will leverage Globus' Virtual Workspaces. Furthermore, the CEI will provide virtual compute node test and certification services to support its monitoring and management using an adaptation of the ROCKS roll retrieval and system configuration mechanisms. A system for administering virtual compute node configuration repositories and retrieving both the contents and their catalog will be defined based on the Xen technology, and connected with the Data Network. The virtual compute node loader will retrieve elements from one or more local or system repositories and use them to configure a virtual compute node.

Data Structures & Models

Globus' Virtual Workspaces and Amazon's Elastic Compute Cloud (EC2) provide the base models for the Common Execution Infrastructure.

Interface Points

The CEI provides its services via the COI to all CI subsystems. However, major interaction points exist between the COI and the Control, Processing, Modeling and Data Networks.

Exhibit 20.3: Architectural Elements

Control Network (CN)

ORION CI Proposal

Control Network (CN)

The Control Network will provision the services required to establish the standard models for the management of stateful and taskable resources. It provides controller processes the semantics to monitor and control the operating state of an active resource as well as to initiate, monitor and amend tasks being carried out by a taskable resource. The managed resource is required to present their operational state and declare their governance context.

Capabilities

1. Provisioning the command, control and monitor semantics to operate and manage a resource,
2. Time structured concurrent coordination and prioritization of shared resources that are distributed and constrained,
3. Provisioning of a behavior-based architecture for rapidly reconfigurable autonomous task reconfiguration,
4. Unique multi-objective optimization of behavior coordination, allowing for effective compromise to be attained between periodically competing task objectives for a collection of resources,
5. Provisioning of a behavior calculus, allowing sequences of task states to be structured for long-term, persistent plans while remaining highly reactive to events and in-situ control requests.

Technologies & Systems

The deployment of the technical components of the Control Network is predicated on the ESB implementation that is the basis of the capability container concept of the COI—this allows the control network to have a federated presence across the CI. In particular, the management of state, execution scheduling and orchestration of taskable resources can be provisioned as state management and orchestration/process execution plug-ins of the ESB. Further techniques and implementation technologies on which the Control Network is predicated include: Interval programming (IvP), a unique, new mathematical model for representing and solving multi-objective optimization problems for reconciling vehicle behaviors active during a mission; MOOS (Mission Oriented Operating Suite), an open source middleware for connecting software components on an autonomous platform; IvP Helm, a behavior-based autonomy package using multi-objective optimization for behavior reconciliation, with a full Boolean logic behavior calculus and an interface to the MOOS middleware.

Data Structures & Models

The taskable or executive resource is at the center of a set of resources comprised of the Task, Plan, Scheduler and Executive. The Task is a specification that declares the execution statement as well as its pre- and post-conditions. An execution statement is either a basic command (such as a Data Network query) or a composite (such as a batch script or binary executable). The Plan resource couples Tasks to time and/or information events using Event-Condition-Action rules. The Scheduler coordinates the capacity, and constraints of the execution environment with the demand and priority of tasks. The Executive resource provides the controlling process with the ability to assess, prosecute and monitor a Task.

Interface Points

The Control Network is used as the basis for and thus has major interaction points with: the Processing Network, the Instrument Network, the Modeling Network; the provisioning model for the CEI's Elastic Compute Cloud.

Exhibit 20.4: Architectural Elements
Data Network (DN)
ORION CI Proposal

Data Network (DN)

The Data Network provides services for the secure preservation, management and presentation of scientific data associated with their structural and semantic context. It supports enactment and enforcement of the ORION Data Policy. The repository and its rule-based policy system will enable and support system-wide registration, persistence and presentation of resources. Each resource is mapped onto a logical name space, and persistent attributes are then managed for each object within that logical name space. Finally, the Data Network will integrate with digital library technologies and preservation environments for publishing and archiving data streams and derived data products.

Capabilities

1. Provision, manage and present data and metadata supporting the ORION domain and data models,
2. provide policy-governed data access,
3. provide user-defined data presentation,
4. provision, manage and present data repositories, collections and streams,
5. negotiate and manage federations of data repositories, collections and streams,
6. negotiate and manage delegation of data preservation and presentation responsibilities, and
7. maintain and ensure the integrity of data in perpetuity.

Technologies & Systems

The core element of the Data Network is the Storage Resource Broker (SRB) data Grid, which is in production use within the ROADNet project to manage federation of multiple independent data streams. SRB data Grids have also been integrated with digital library and knowledge management technology (DSpace and Fedora), and have been used to build preservation environments. Furthermore, the SRB is in production use on the BIRN/Telescience projects, where it acts as a federated data store for both metadata catalogs and vast volumes of imaging and other data. The integrated Rule-Oriented Data System (iRODS) adds policy-based data management capability to the SRB environment. MBARI's Shore Side Data System (SSDS) will be integrated with SRB/iRODS to provide domain-specific data cataloging, search and navigation capabilities.

Data Structures & Models

The project will leverage and adapt the successful SRB implementations, featuring data, catalog and repository models for ROADNet and Telescience/BIRN, integrating IRIS's Standard for the Exchange of Earthquake Data (SEED), and MBARI's SIAM/SSDS.

Interface Points

The major interaction points of the Data Network are with the Instrument Network through receipt of data streams packetized into files that are stored on repository systems specified by management policies. The Data Network interacts with the CEI by submission of long-running tasks such as checksum-based integrity validation on a collection to a workflow environment for execution. The DN also interacts with the COI to observe the ORION policies managed within that element, as well as to provide data access for presentation and transformation purposes. External data access services will be provisioned leveraging UCAR's THREDDS, OPeNDAP & IDD, as well as OGC Web.

Exhibit 20.5: Architectural Elements

Processing Network (PN)

ORION CI Proposal

Processing Network (PN)

The Processing Network provides immediate-mode scheduling of processes at specified locations based on explicit time requirements and event triggers. The processes to be scheduled and executed come in a variety of forms, including compiled code, scripting languages, and workflow graphs, and are placed within the common time and event semantics of the COI.

Capabilities

1. Immediate-mode scheduling of processes at specified locations,
2. coupling of processes to the streaming environment of the Data Network,
3. coordinated and/or chained scheduling of processes,
4. an extendable set of process execution engines,
5. standard process execution planning and control, and
6. standard providence capture and reporting. Process authoring and monitoring applications that will be integrated as a user interface to the Processing Network include MatLab and Kepler, as well as NCSA's Community Ensemble Service.

Technologies & Systems

The scheduler will be built based on NCSA's Community Ensemble Service. The core execution engines will leverage and adapt the existing Kepler framework, as well as the execution engines of BRTT Antelope, JJDAC, ISI Pegasus, and Matlab.

Data Structures & Models

The process and scheduling models will be based on the UCSD's Kepler scientific workflow system and ISI's Pegasus Grid workflow execution framework.

Interface Points

The major interaction point for the Processing Network is the streaming environment of the Data Network. The coupling of processes is primarily accomplished through publication and subscription to the Data Network. The COI, especially its CICCs, provision the execution infrastructure, as well as access to the CEI elastic node provisioning service and the Knowledge Network resource association and tagging services.

Exhibit 20.6: Architectural Elements

Instrument Network (IN)

ORION CI Proposal

Instrument Network (IN)

The Instrument Network will implement services falling into the categories of network-wide sensor Grid configuration and control, system support capabilities for individual instruments on the Grid, and support for various aspects of the entire mission (i.e., domain models). The Instrument Network will provide facility components, such as instrument management and presentation modules, as well as services for taskable interfacing with instruments, and instrument and metadata registration.

Capabilities

1. Provide data acquisition, buffering, and transport mechanisms,
2. provide command and control systems,
3. maximize total data return from all instruments,
4. provide instrument test and certification,
5. provide instrument registration with associated metadata,
6. place processing capability into the instrument network,
7. manage and allocate resources to instruments,
8. manage and allocate resources from instruments,
9. acquire, marshal, buffer, and transport data to the Data Network,
10. prioritize data delivery, and
11. provision storage and processing throughout the Instrument Network.

Technologies & Systems

The Instrument Network will initially use the ORBserver technology as the backbone of the transport system. This is a reliable, content-neutral, packetized buffering and event-driven distribution system that can be configured for gridded data acquisition, sharing, and processing. Extensive measures have been taken in its implementation to assure robust delivery. Network state-of-health monitoring issues have already been addressed in part by land-based sensor grids such as ROADNet, based in part on the Nagios open-source network monitoring package (http://www.nagios.com) together with local modifications.

Data Structures & Models

The Instrument Network will implement and contribute domain models via the COI for instruments, observations, plans, schedules, marine resources, allocation, and transducers, leveraging standards such as OGC Sensor Web Enablement (SWE) and IEEE 1451. The plug-and-play instrument support and remote control capabilities of the Software Infrastructure and Applications for Monterey Ocean Observing System SIAM are expected to provide models from which the ORION CI can draw.

Interface Points

The Instrument Network provides its services via the COI to all subsystems. However, major interaction points exist with the COI, the Control Network and the Data Network. While the COI should largely mitigate end-chain communication technology limitations, it is clear that the pathways used for data acquisition and state-of-health monitoring for instruments must often share features with network-wide resource allocation, network state-of-health monitoring, and sometimes data Grid transmission capabilities. Thus, the Instrument Network must at the very least peacefully coexist with features of network-wide communication from central observatory acquisition and control sites out to those of sensors. The Instrument Network interacts heavily with the front-end of the Data Network, and hence must partially drive the interface engineering.

Exhibit 20.7: Architectural Elements

Modeling Network (MN)

ORION CI Proposal

Modeling Network (MN)

The Modeling Network will provision ocean models and provide service implementations for their exploitation. To that end, it will incorporate 1) a virtual ocean model covering the full range of observatory scales; 2) a model control and virtual sampling interface; 3) data access and assimilation interface; 4) model coupling interface; 5) a taskable resource interface supporting construction, modification and execution of numerical ocean models, as well as Observing System Simulation Experiments (OSSEs) for network design and trade studies, data impact investigations, and data and information management exercises.

Capabilities

1. Data calibration and validation,
2. derived data product generation,
3. a stream process network, including stream subscription with process and execution location, stream process scheduling, and stream process execution,
4. a measurement processing network, including a measurement calculus and measurement semantic model, and
5. a modeling network, including on-demand modeling, assimilative modeling, and an observing system simulator.

Technologies & Systems

The development effort for the Modeling Network will rest on community-based numerical ocean models such as the Regional Ocean Modeling System (ROMS) and the Harvard Ocean Prediction System (HOPS). Its modeling and simulation capabilities will leverage existing and emerging data assimilation modules based on the variational method (e.g., 3DVAR or 4DVAR) or Kalman Filter (KF). The MN will be configured as a set of web services with a portlet implementation for model configuration, data delivery and visualization.

Data Structures & Models

JPL's OurOcean Portal, ROMS and HOPS will provide the core data models for the Modeling Network.

Interface Points

The Modeling Network will interface with the COI to interact with the other resource networks and infrastructure elements. In particular, it will interact with the Instrument, Control, Processing, and Knowledge Networks, and the CEI, to inform the direction of observations and to execute simulation tasks on the Grid. Furthermore, the Modeling Network will integrate into the user-configurable virtual observatory, and provide interface control with the observatory data systems and sensor simulator.

Exhibit 20.8: Architectural Elements
Knowledge Network (KN)
ORION CI Proposal

Knowledge Network (KN)

The Knowledge Network provisions the capabilities, resources, tools and presentation solutions for ontology management and mediation. For instance, the Knowledge Network will handle queries such as "what data exist in anoxic regions off Washington" and "show me sensors that can measure temperature within 0.1ºC and are available between March 10 and 20th in my observatory of interest". Users will be able to add annotations to a data stream, a derived data product, or a group of data streams or derived data products, and, in particular, will be able to annotate data with respect to one or more ontologies. Introduction and sharing of these ontologies will also be enabled. The Knowledge Network will hold data related to multiple ontologies and translate data returns into the ontology of interest to the user.

Capabilities

1. Data mediation: complex querying across and integration of geospatial, temporal, spatiotemporal, relational, XML, and ontological (tree and graph structures) resources,
2. present, find, exploit and annotate data based on a semantic frame of reference,
3. provision and exploit sharable semantic frames of reference, and
4. provision and exploit sharable mappings between different semantic frames of reference (i.e. crosswalks between multiple ontologies).

Technologies & Systems

Ontology and mediation research carried out in the following projects will be leveraged: GEON, BIRN, SWEET and MMI. The following existing technologies will be extended and customized to interact with the resource and activity modules of the system through the COI:

1. the Data Integration Engine used by the BIRN and OceanLife projects,
2. the VINE ontology mapping tool developed by MBARI, and
3. a TBD ontology definition tool.

Data Structures & Models

The Data Integration Engine used by the BIRN and OceanLife and the VINE ontology mapping tool provide the core data structures and models for the Knowledge Network.

Interface Points

The Knowledge Network will interface with the COI to interact with the other resource networks. In particular, it will interact with the Instrument, Modeling and Processing Networks using the schema for resource holding sensor capabilities and scheduling information, and with the Data Network using SRB/iRODS schemas and other details needed to interact with data resources. Furthermore, the development team will depend on direct interaction with the scientific user community to facilitate ontology creation and maintenance.

Exhibit 21: Engineering Life Cycle

ORION CI Proposal

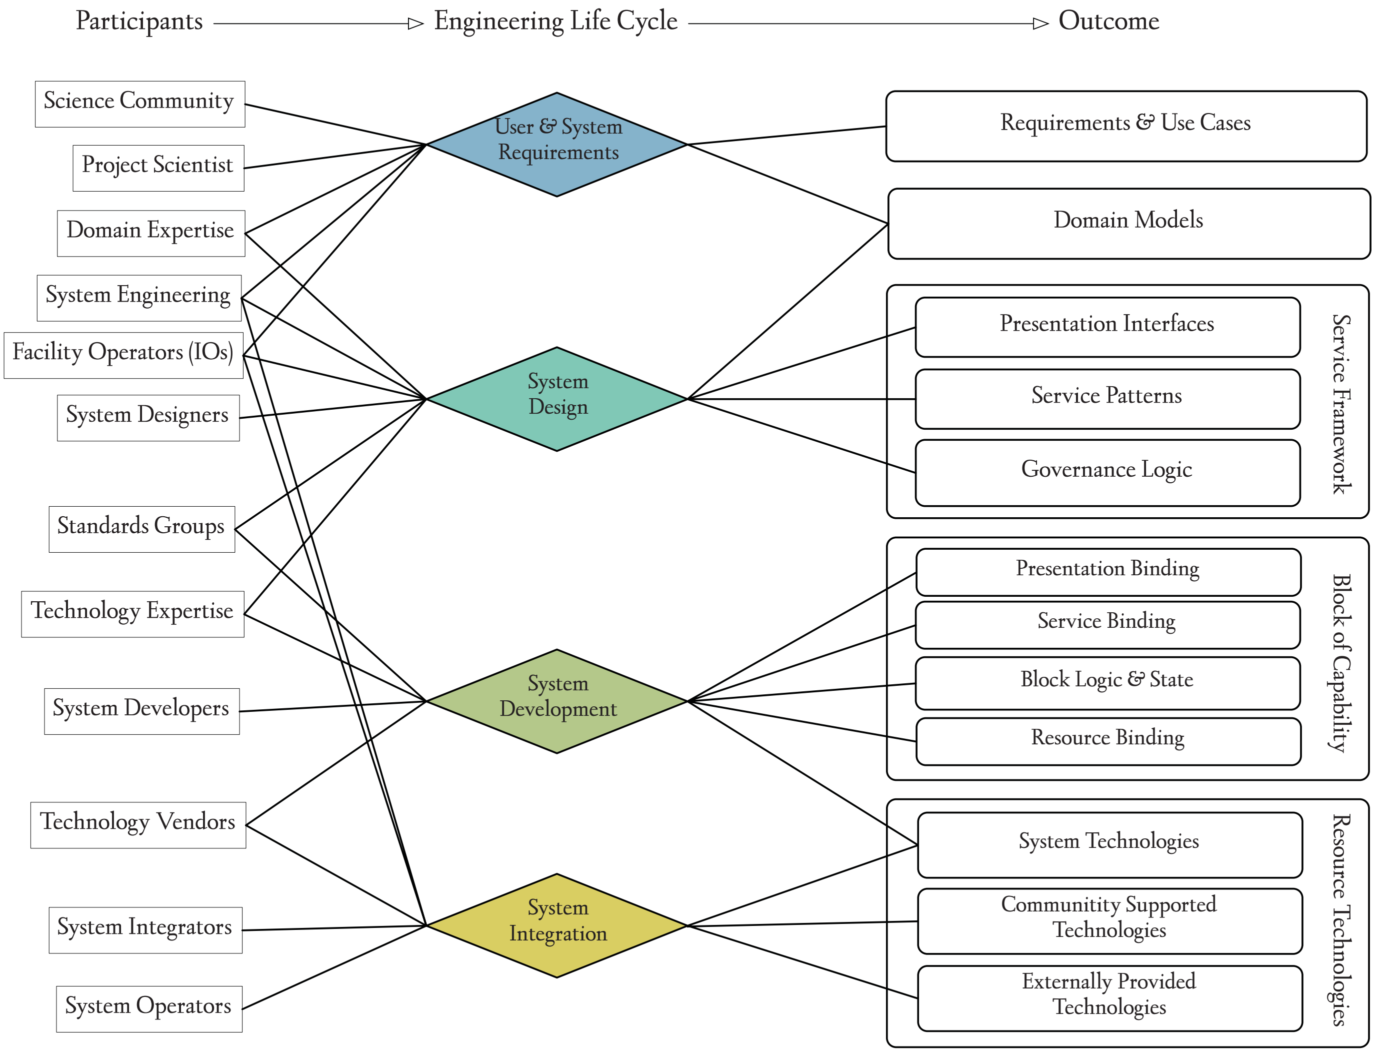

Exhibit 22: Development of Architecture Elements

ORION CI Proposal

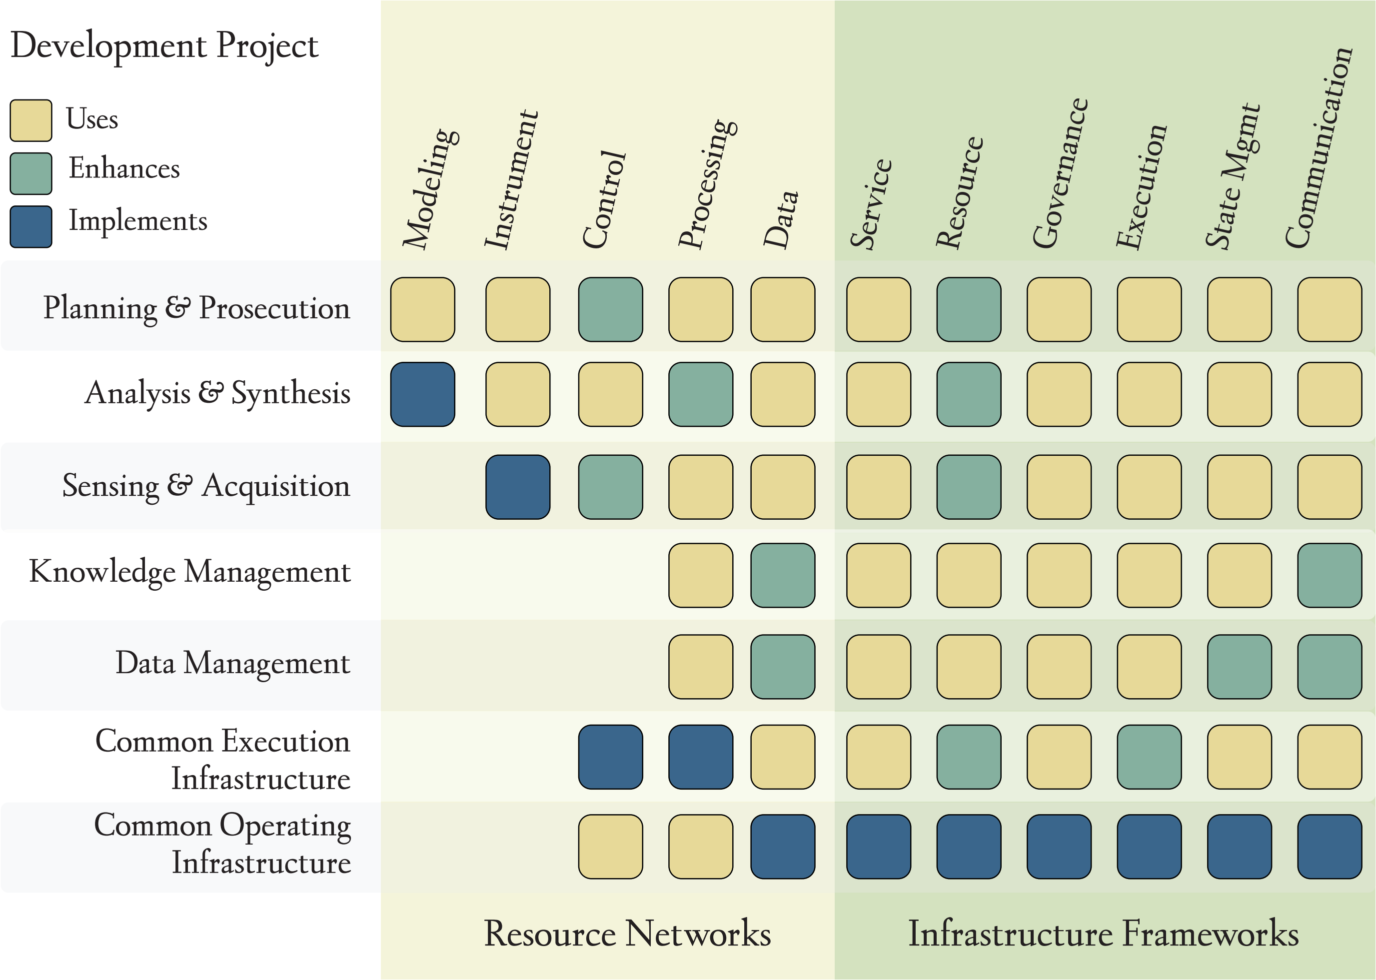

Development Project

🟨 Uses
🟩 Enhances
🟦 Implements

	Resource Networks					Infrastructure Frameworks					
	Modeling	Instrument	Control	Processing	Data	Service	Resource	Governance	Execution	State Mgmt	Communication
Planning & Prosecution	🟨	🟨	🟩	🟨	🟨	🟨	🟩	🟨	🟨	🟨	🟨
Analysis & Synthesis	🟦	🟨	🟨	🟩	🟨	🟨	🟩	🟨	🟨	🟨	🟨
Sensing & Acquisition		🟦	🟩	🟨	🟨	🟨	🟩	🟨	🟨	🟨	🟨
Knowledge Management				🟨	🟩	🟨	🟨	🟨	🟨	🟨	🟩
Data Management				🟨	🟩	🟨	🟨	🟨	🟨	🟩	🟩
Common Execution Infrastructure			🟦	🟦	🟨	🟨	🟩	🟨	🟩	🟨	🟨
Common Operating Infrastructure			🟨	🟨	🟦	🟦	🟦	🟦	🟦	🟦	🟦

Exhibit 23: Virtual Ocean Allocation Justification

ORION CI Proposal

Estimate of Computing resources required for the virtual ocean task

We assume that the virtual ocean will cover all the three observatories on the global, regional and coastal scales. The virtual ocean will be updated on the daily basis and include a data processing, modeling and data assimilation components.

The global virtual ocean model (360º in longitude by 180º in latitude) covers the global domain with a horizontal resolution on the order of 10 km with 100 vertical layers. Thus, the total grid points are 3600*1800*100.

The regional virtual ocean model (10º in longitude by 10º in latitude) covers the Juan de Fuca tectonic plate with a horizontal resolution on the order of 2-km with 100 vertical layers. Thus, the total grid points are 5000*5000*100.

It is expected that there will be multiple coastal observatories covering the U.S. coastal waters. The West Coast Endurance Array includes four sites: Central Washington, Central Oregon, Central California, and Southern California. The East Coast Endurance Array will be centered at the South Atlantic Bight. The Pioneer Array will be located at the Mid-Atlantic Bight. Thus, there will be at least six coastal observatories. A typical coastal observatory would be on the order of 3º in longitude and 3º in latitude. With a typical horizontal resolution at 500-meter and 100 vertical layers, the total grid points for the six observatories would be 3600*3600*100.

In summary, the virtual ocean models for the three observatories will have a size on the order of 12200*10400*100 grid points that need to run on the daily (24-hour) basis. There are usually no more than twenty tracers including four physical variables (temperature, salinity, zonal and meridional current), a dozen biogeochemical variables (silicate, nitrate, phytoplankton, ammonium, two phytoplankton groups, two zooplankton grazers, two detrital pools, DIC, and oxygen), and four more tracers of interests (e.g., tracers from the Hydrothermal event plumes). Most of the codes are written in MPI. Some of the data processing and assimilation codes will be initially written in OpenMP and tested on the shared-memory computers, and will be converted to MPI. Thus, we plan to run the production codes using MPI on any distributed-memory computers.

Our initial test has been conducted on the SGI Altix computer using the Intel Itanium2 processors (900 MHz, 1.5 MB cache, 4 GFLOPS peak speed, and 1 GB main memory, gigabit network). We assume that the production run for the virtual ocean will start in FY2009. If we assume that the processor speed will double every 18 months (this is a conservative estimate because it is now roughly double every year), we will have the 2009 processor a factor of four faster than the 2006 Intel Itanium 2 processor. To integrate the virtual ocean with a dimension of 12200*10400*100 grid points with 20 variables over a period of 24 hours requires about 20 hours on a "2009" 1024-processor cluster computer.

By 2009 when we start the daily update of the virtual ocean for the Orion CI project, we will request 20,480 single processor hours per day on the supercomputer facility provided by NSF (e.g., via TeraGrid) or other agencies.

Exhibit 24: Comparison of Waterfall & Spiral Management Models

ORION CI Proposal

Waterfall Development Model	Spiral Development Model
Requirements may be completely defined in advance of development	Requirements are discovered and defined during development
Requirements contain no high risk implications that are unresolved	Risks are continually discovered and the reduction of high risk elements drives the development process
The requirements will change very little during development and operations	The requirements will evolve throughout the life cycle
The requirements are compatible with expectations of all key stakeholders	Ongoing negotiation with key stakeholders will be required to meet expectations
The best architecture to implement the requirements can be defined in advance	The best architecture to implement the requirements cannot be defined in advance
There is sufficient calendar time for sequential system development	Plans (including cost and schedule) are continually refined as the requirements and solutions become better defined

Exhibit 25: Spiral Management Milestones

ORION CI Proposal

Life Cycle Objectives (LCO)

Goal: Definition of what the cyberinfrastructure design will accomplish.

Focus: Ensuring that at least one architecture choice is acceptable to the stakeholders

Stakeholder Commitment: Building the identified architecture.

Completion: End of planning phase

Life Cycle Architecture (LCA)

Goal: Definition of the software architecture and technologies needed to implement the cyberinfrastructure design.

Focus: Committing the project to a viable CI design.

Stakeholder Commitment: Supporting deployment of the cyberinfrastructure design.

Iterations: Five (every 12 mo starting at 18 mo)

Initial Operating Capability (IOC)

Goal: Approval of final production release.

Focus: Validation of a viable, maintainable system.

Stakeholder Commitment: Going into operations phase.

Completion: End of project.

Exhibit 26: Software Release Cycle

ORION CI Proposal

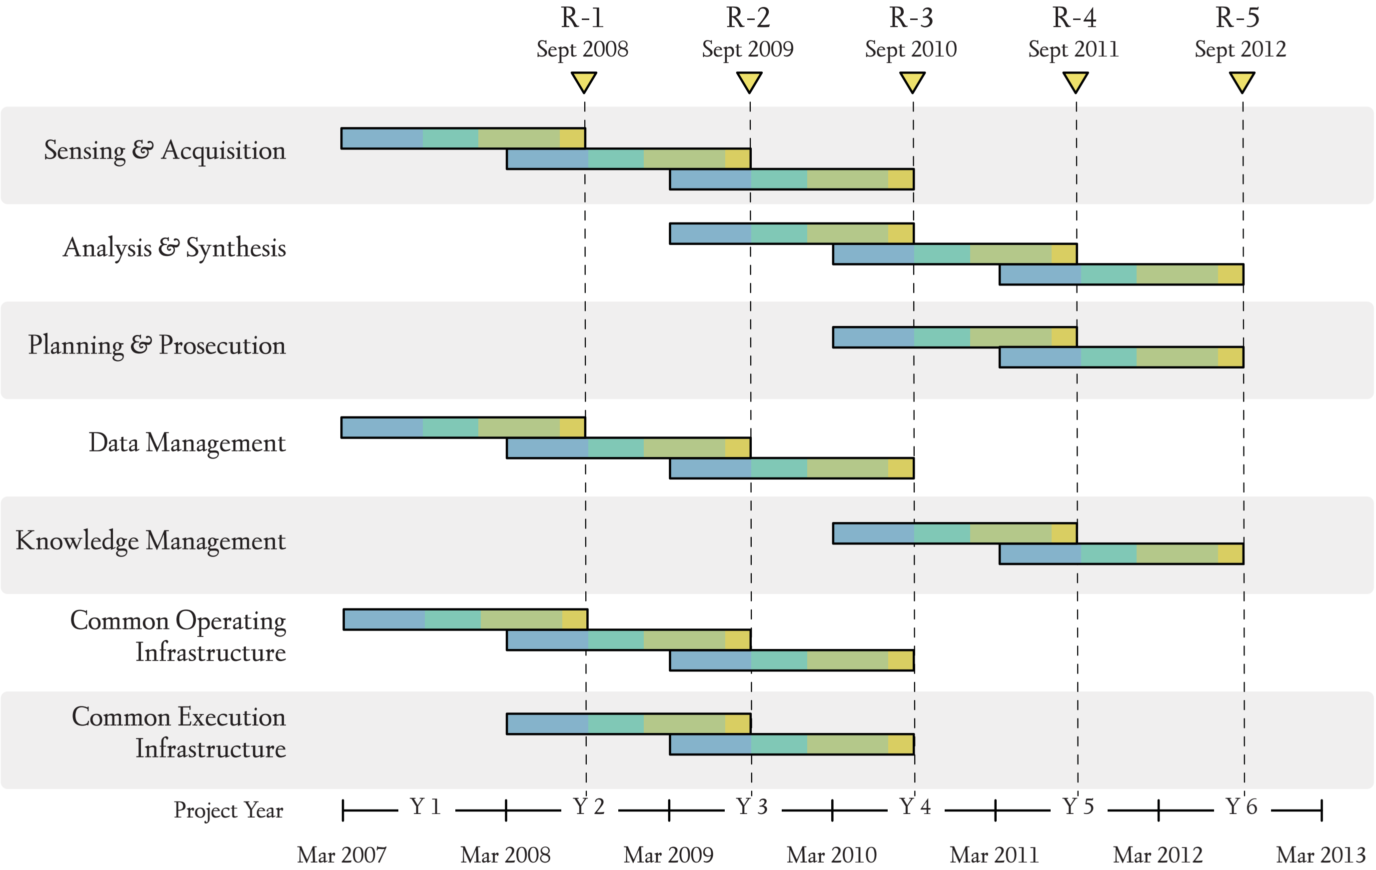

Exhibit 27: Organizational Structure

ORION CI Proposal

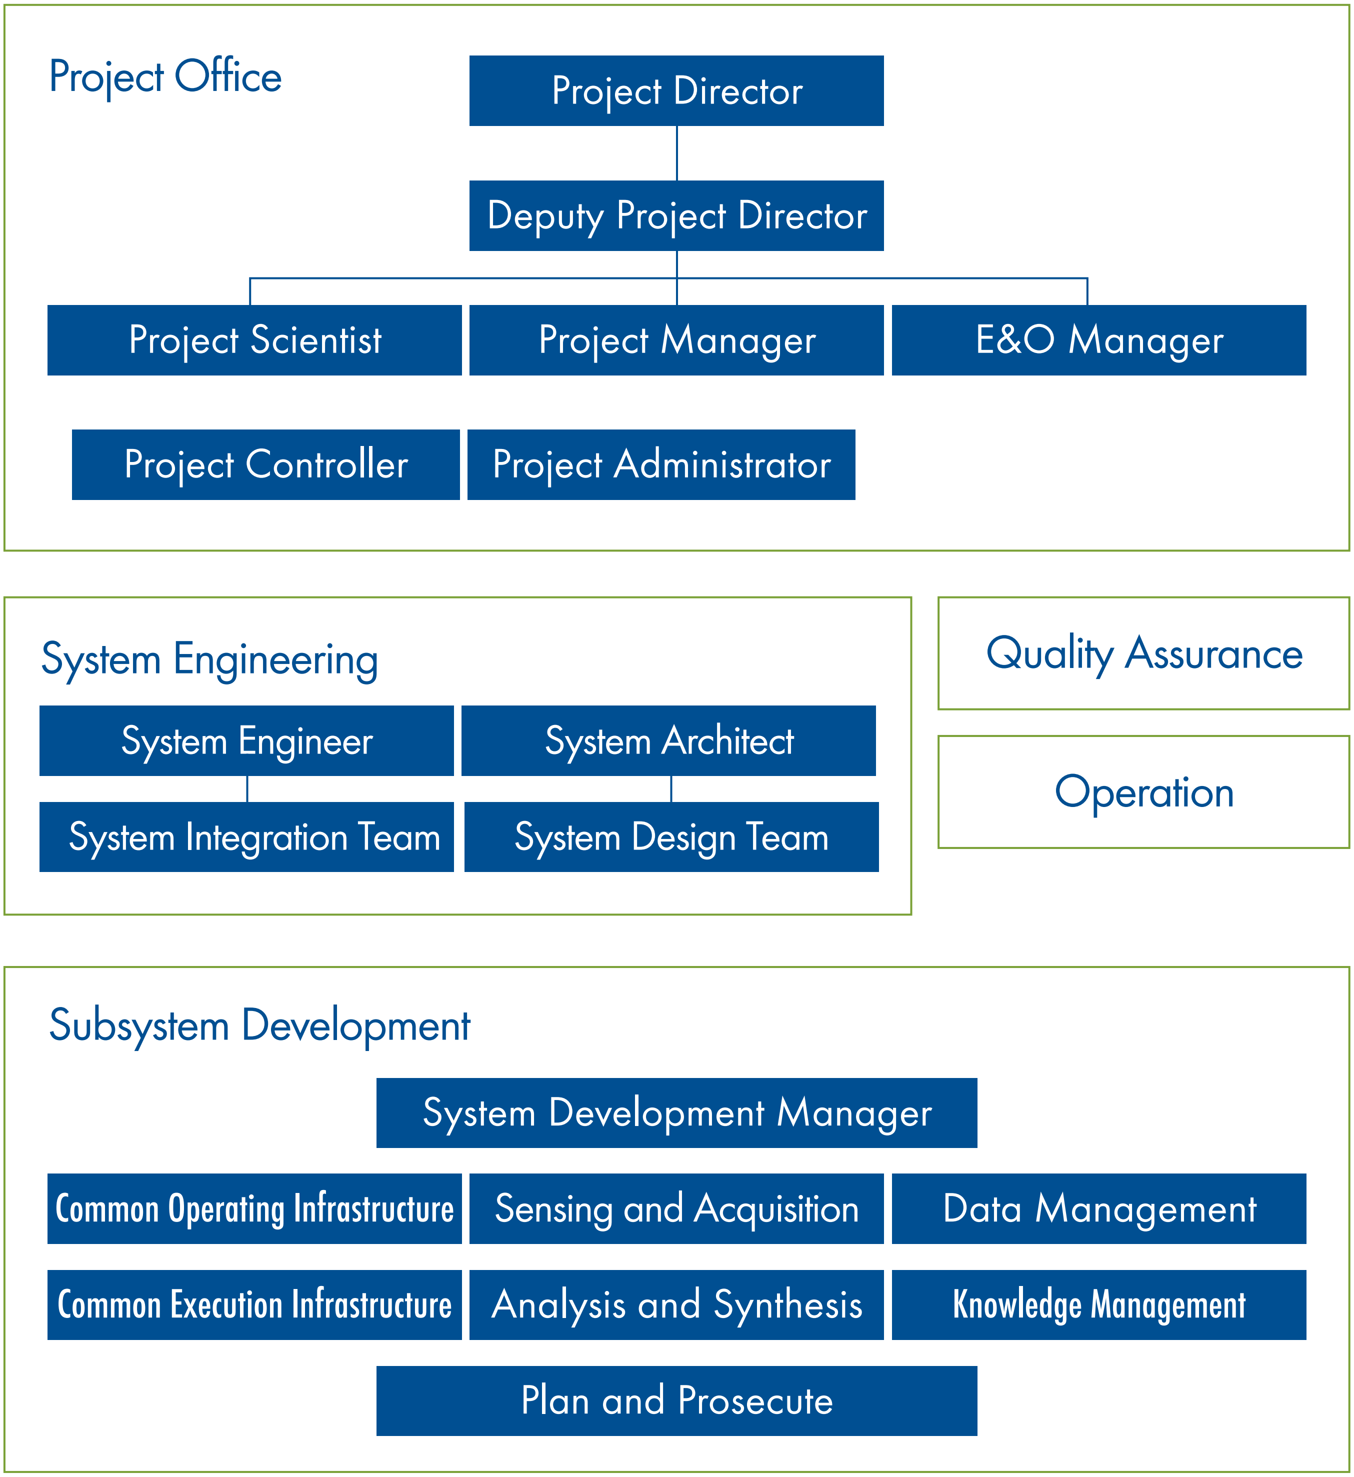

Exhibit 28: Project Participant Labor

ORION CI Proposal

	Y-1	Y-2	Y-3	Y-4	Y-5	Y-6	Total	
	15.90	22.35	22.90	19.13	16.30	14.30	110.9	
Project Office	3.50	3.50	3.25	2.80	2.80	2.80	18.65	17%
UCSD-SIO	0.55	0.55	0.55	0.55	0.55	0.55	3.30	
UCSD-Calit2	2.00	2.00	2.00	2.00	2.00	2.00	12.00	
Rutgers	0.25	0.25	0.25	0.25	0.25	0.25	1.50	
Raytheon	0.20	0.20	0.20	-	-	-	0.60	
Triad PM	0.50	0.50	0.25	-	-	-	1.25	
System Engineering	3.85	4.00	3.55	2.63	1.70	1.00	16.73	15%
WHOI	0.75	0.75	0.75	0.63	0.50	0.50	3.88	
UCSD-Calit2	2.60	2.75	2.30	2.00	1.20	0.50	11.35	
Raytheon	0.50	0.50	0.50	-	-	-	1.50	
Software Development	8.05	13.10	13.60	11.20	9.05	7.50	62.50	56%
Management	1.00	1.00	1.00	1.00	1.00	1.00	6.00	
UCSD-Calit2	1.00	1.00	1.00	1.00	1.00	1.00	6.00	
Subsystem Projects	6.55	11.10	11.60	9.20	7.05	5.50	51.00	
Sensing & Acquisition	1.30	2.20	2.10	1.00	-	-	6.60	6%
Lindquist Consulting	0.35	0.50	0.50	0.25	-	-	1.60	
UCSD-SIO	0.50	1.00	1.00	0.50	-	-	3.00	
UCSD-Calit2	0.25	0.50	0.50	0.25	-	-	1.50	
MBARI	0.20	0.20	0.10	-	-	-	0.50	
Analysis & Synthesis	0.15	0.45	1.20	1.90	2.20	1.10	7.00	6%
JPL	0.15	0.25	0.80	1.70	2.20	1.10	6.20	
MIT	-	0.10	0.20	0.10	-	-	0.40	
USC-ISI	-	0.10	0.20	0.10	-	-	0.40	
Planning & Prosecution	0.15	0.10	0.35	1.40	2.20	2.20	6.40	6%
MIT	0.15	0.10	0.25	0.80	1.20	1.20	3.70	
JPL	-	-	0.10	0.60	1.00	1.00	2.70	
Data Management	2.25	2.90	2.05	1.10	0.25	-	8.55	8%
UCSD-Calit2	0.75	1.00	1.00	0.75	0.25	-	3.75	
UCSD-SDSC	0.70	1.20	0.95	0.35	-	-	3.20	
UCSD-NCMIR	0.60	0.50	-	-	-	-	1.10	
MBARI	0.20	0.20	0.10	-	-	-	0.50	
Knowledge Management	0.15	0.10	0.45	1.70	2.40	2.20	7.00	6%
UCSD-SDSC	0.15	0.10	0.25	0.80	1.20	1.20	3.70	
UCSD-Calit2	-	-	-	0.50	1.00	1.00	2.50	
MBARI	-	-	0.10	0.20	0.10	-	0.40	
NCSA	-	-	0.10	0.20	0.10	-	0.40	
Common Operating Infrastructure	2.20	3.50	2.25	0.50	-	-	8.45	8%
UCSD-Calit2	0.85	1.50	1.25	0.50	-	-	4.10	
NCSA	0.60	1.00	0.50	-	-	-	2.10	
NCSU	0.75	1.00	0.50	-	-	-	2.25	
Common Execution Infrastructure	0.35	1.85	3.20	1.60	-	-	7.00	6%
U of Chicago	0.05	0.55	1.00	0.50	-	-	2.10	
NCSA	0.05	0.55	1.00	0.50	-	-	2.10	
UCSD-SDSC	0.25	0.75	1.20	0.60	-	-	2.80	
Quality Assurance	0.50	1.00	1.00	1.00	1.00	1.00	5.50	
UCSD-Calit2	0.50	1.00	1.00	1.00	1.00	1.00	5.50	
Hardware Development	-	-	-	-	-	-	-	0%
	-	-	-	-	-	-	-	
Implementation	0.50	1.75	2.50	2.50	2.75	3.00	13.00	12%
UCSD-Calit2	0.50	1.75	2.50	2.50	2.75	3.00	13.00	

8.0 References

ORION CI Proposal

Arrott, Matthew, Alan Chave, John Graybeal, Eric Guillemot, Ingolf Krüger, Benoît Pirenne, "ORION CI Conceptual Architecture Team Progress Report," Mar 2006.

Arzberger, Peter (ed.), "Sensors for Environmental Observatories." Report of the NSF-Sponsored Workshop held December 2004, (2005).

AV1: Overview and Summary Information, containing an explanation of the project and design products. (ver. 7/26/06) www.orionprogram.org/organization/committees/ciarch/CI_AV-1_07262006.pdf

AV2: Integrated Dictionary, containing definitions of special terms used in the design. (ver. 5/30/06), www.orionprogram.org/organization/committees/ciarch/CI_AV-2_05302006.pdf

Baru, Chaitanya, R. Frost, J. Lopez, R. Marciano, R. Moore, A. Rajasekar, M. Wan, "Metadata Design for a Massive Data Analysis System." *Proc.* CASCON '96. November 1996.

Baru, Chaitanya, R. Moore, A. Rajasekar, W. Schroeder, M. Wan, R. Klobuchar, D. Wade, R. Sharpe, J. Terstriep, "A Data Handling Architecture for a Prototype Federal Application." *Sixth Goddard Conference on Mass Storage Systems and Technologies*, March 1998.

Benjamin, Michael, "Multi-Objective Helming with Interval Programming on Autonomous Marine Vehicles." IEEE/RJS IROS 2006 Workshop on Multi-Objective Robotics (IROS-MOR 2006), October 2006.

Benjamin, Michael, David Battle, Donald Eickstedt, Henrik Schmidt, Arjuna Balasuriya, "Control of an Autonomous Underwater Vehicle Towing a Vector Sensor Array." International Conference on Robotics and Automation (ICRA), April 2007. Submitted for peer review September 15th, 2006.

Berriman, G. B., D. Curkendall, J. Good, J. Jacob, D. S. Katz, M. Kong, S. Monkewitz, R. Moore, T. Prince, R. Williams, "An Architecture for Access to a Compute Intensive Image Mosaic Service in the NVO." *SPIE Conference* 4686 "*Virtual Observatories,*" Hawaii, August 2002.

Blue Ribbon Review of the "OOI Scientific Objectives and Network Design," August 2006.

Boehm, B.W., (1988) "A spiral model of software development and enhancement." *Computer*, 21, 61-72.

Boehm, B.W., (2000) "Spiral Development: Experiences, Principles, and Refinements." *Technical Note* CMU/SEI-00-SR-008, Carnegie Mellon University, 49 pp.

Boisvert, Ronald, P. Tang, "The Architecture of Scientific Software." *Data Management Systems for Scientific Applications.* Kluwer Academic Publishers. (2001) pp. 273-284.

Brownsword, L. and J. Smith, (2005) "Using earned value management (EVM) in spiral development." *Technical Note* CMU/SEI-2005-TN-016, Carnegie Mellon University, 35pp.

Brunner, R., S. Djorgovski, A. Szalay, "Virtual Observatories of the Future." *Astronomical Society of the Pacific Conference Series*, Vol. 225, pp. 257-264, June 2000.

Burns, Brad, Steve Castelin, Dave DeMartino, Steve Ebner, Barbara Fletcher, Julia Gazagnaire, Ray Harnois, Sam Hester, Martha Head, George Kindel, John Lisiewicz, Pat Madden, Steve Wells, "The Navy Unmanned Undersea Vehicle (UUV) Master Plan," November 9, 2004.

Chave, Alan D., A. Bowen, S. Glenn, W. Hill, M. Kosro, E. Massion, L. Mayer, D. Schwartz, K. Smith, B. Wall, F. Wooding, P. Worcester, "Ocean Observatories Initiative Facilities Needs from UNOLS." Report of the UNOLS Working Group on ocean Observatory Facility Needs, December 2003.

Clark, H. Larry, "New Seafloor Observatory Networks in Support of Ocean Science Research." IEEE Conference Publishing, (2001).

Consortium for Oceanographic Research and Education, "A National Initiative to Observe the Oceans: A CORE Perspective."

Daly, Kendra (Chair), Richard Jahnke (Co-Chair), Mark Moline (Co-Chair), Robert Detrick, Doug Luther, George Matsumoto, Larry Mayer, Keith Raybould, "Report of the D&I Workshop, 27-30 March 2006," May 2006.

DEOS Moored Buoy Ocean Observatory Design Study. obslab.whoi.edu/buoy.html

Devadithya, Tharaka, K. Chiu, K. Huffman, D.F. McMullen, "The Common Instrument Middleware Architecture: Overview of Goals and Implementation." *IEEE International Conference on e-Science and Grid Computing*, Melbourne, Australia. December 5-8, 2005.

Edgington, Duane R., and Davis, Daniel (eds.), (2004) "SENSORS: Ocean Observing System Instrument Network Infrastructure," NSF SENSORS Project Workshop Report, MBARI, Moss Landing, CA.

Eickstedt, D.P., Benjamin M.R., "Cooperative Target Tracking in a Distributed Autonomous Sensor Network." *In Proceedings of the IEEE/MTS Oceans'06 Conference*, Boston, MA, September 18-21, 2006.

Eickstedt, Donald, Michael Benjamin, Ding Wang, Henrik Schmidt, "Behavior-Based Adaptive Control for Autonomous Oceanographic Sampling." International Conference on Robotics and Automation (ICRA), April 2007. Submitted for peer review September 15, 2006.

Eickstedt, Donald, Michael Benjamin, Henrik Schmidt, John Leoanrd, ”Adaptive Tracking of Underwater Targets with Autonomous Sensor Networks.” *Journal of Underwater Acoustics*, Accepted for publication, to appear 2007.

Eickstedt, Donald, Michael Benjamin, Henrik Schmidt, John Leonard, ”Adaptive Control of Heterogeneous Marine Sensor Platforms in an Autonomous Sensor Network.” IEEE/RJS International Conference on Intelligent Robots and Systems (IROS), October 2006.

Embedded Network Sensing. www.cens.ucla.edu

Enabling Grids for E-Science. public.eu-egee.org

Enabling Ocean Research in the 21st Century. www.nap.edu/catalog/10775.html

European Data Grid. eu-datagrid.web.cern.ch/eu%2datagrid/

First Annual IOOS Development Plan. www.ocean.us/ioosplan.jsp

First IEEE International Conference on e-Science and Grid Computing (e-Science 2005), Melbourne, Australia, Dec. 5-8, 2005.

Foster, I., C. Kesselman, and S. Tuecke, “The Anatomy of the Grid.” in F. Berman, G. Fox, and T. Hey (eds.), *Grid Computing: Making the Global Infrastructure*

Glenn, S.M. and T. D. Dickey (eds.), (2003) “SCOTS: Scientific Cabled Observatories for Time Series.” NSF Ocean Observatories Initiative Workshop Report, Portsmouth, VA, 80 pp., www.geoprose.com/projects/scots_rpt.html.

Grid Physics Network. www.griphyn.org

Guiding National Ocean Research Investment: Public Workshop for the Development of the Ocean Research Priorities Plan, ocean.ceq.gov/about/jsost_workshop/welcome.html

H.L. Clark, (2001) “New Seafloor Observatory Networks in Support of Ocean Science Research,” IEEE conference Publishing.

H.R. 1552, the Marine Research and Related Environmental Research and Development Programs Authorization Act of 1999, www.core-ocean.org/Dev2Go.web?id=213518

Hansen, Todd, K. Lindquist, F. Vernon, J. Orcutt, “VORBrouter: A Dynamic Data Routing System for Real-Time Seismic Networks.” *EOS Trans.* AGU 84(46) Fall Meeting Supplement, Abstract U21A-06. (2004).

International Virtual Data Grid. www.ivdgl.org

Ingolf H. Krueger, Michael Meisinger, Massimiliano Menarini, Stephen Pasco: Rapid Systems of Systems integration - Combining an Architecture-Centric Approach with Enterprise Service Bus Infrastructure. *In: Proceedings of the 2006 IEEE International Conference on Information Reuse and Integration, IRI* - 2006: Heuristic Systems Engineering, September 16-18, 2006, Waikoloa, Hawaii, USA. IEEE Systems, Man, and Cybernetics Society.

Jahnke, Richard, John Bane, Andrew Barnard, John Barth, Francisco Chavez, Hans Dam, Ed Dever, Paul DiGiacomo, James Edson, Rocky Geyer, Scott Glenn, Ken Johnson, Mark Moline, James O’Donnell, Joan Oltman-Shay, Ola Persson, Oscar Schofield, Heidi Sosik, Eric Terrill, “Coastal Observatory Research Arrays: A Framework for Implementation.” Planning Report on the CoOP CORA Workshop, 12-13 November 2003, Chicago, Illinois. *Coastal Ocean Processes (CoOP) Program Report Number 9*, December 2003.

Jahnke, Richard, Larry Atkinson, John Barth, Francisco Chavez, Kendra Daly, James Edson, Peter Franks, James O’Donnell and Oscar Schofield, “Coastal Ocean Processes and Observatories: Advancing Coastal Research.” Report on the CoOP Observatory Science Workshop, May 7-9, 2002 Savannah, Georgia US, *CoOP Report Number 8*, November 2002.

Kruchten, P., (2004) “The Rational Unified Process: An Introduction,” 3rd ed. New York, NY: Addison-Wesley.

Lim, Hock Beng, Y. Teo, P. Mukherjee, V. Lam, W. Wong, S. See, “Sensor Grid: Integration of Wireless Sensor Networks and the Grid.” *Local Computer Networks* (2005).

Ludaescher, B., R. Marciano, R. Moore, “Towards Self-Validating Knowledge-Based Archives.” *SDSC TR 2001-1*, January 2001.

Ludäscher, B., R. Marciano, R. Moore, “Preservation of Digital Data with Self-Validating, Self-Instantiating Knowledge-Based Archives.” *ACM SIGMOD Record*, 30(3), 54-63, (2001).

Marciano, R., B. Ludaescher, R. Moore, “The Senate Legislative Activities Collection (SLA).” *SDSC TR 2001-05*, January 2001.

Moore, R., “Data Grids, Persistent Archives, and Digital Libraries.” published in *Business Briefing: Data Management and Storage Technology*, (2003).

Moore, R., “High Performance Computing for Preservation of Electronic Records.” *63rd Annual Meeting of The Society of American Archivists*, Pittsburgh, August 1999.

Moore, R., “Information Based Computing.” Invited talk for the Supercomputing *’98 Conference*, Orlando, Florida, November 1998.

Moore, R., “Knowledge-based Grid.” *SDSC TR 2001-02*, January 2001.

Moore, R., “Knowledge-based Persistent Archives.” *Proceedings of La Conservazione Dei Documenti Informatici Aspetti Organizzativi E Tecnici*, in Rome, Italy, October 2000.

Moore, R., “Operations for Access, Management, and Transport at Remote Sites.” *Global Grid Forum*, December 2003.

Moore, R., “Preservation of Data.” *e-Science Issue on Intelligent Systems, SDSC TR 2003-006*, San Diego, California, September 2003.

Moore, R., "The San Diego Project: Persistent Objects." *Archivi & Computer, Automazione E Beni Culturali, l'Archivio Storico Comunale di San Miniato*, Pisa, Italy, February 2003.

Moore, R., C. Baru, "Virtualization Services for Data Grids." Book chapter in *Grid Computing: Making the Global Infrastructure a Reality*, John Wiley & Sons Ltd, (2003).

Moore, R., C. Baru, P. Bourne, M. Ellisman, S. Karin, A. Rajasekar, S. Young, "Information Based Computing." *Proceedings of the Workshop on Research Directions for the Next Generation Internet*, May 1997.

Moore, R., C. Baru, S. Karin, A. Rajasekar, "Information Based Computing." *Proceedings of the Workshop on Research and Development Opportunities in Federal Information Services*, March 1997.

Moore, R., J. Lopez, C. Lofton, W. Schroeder, G. Kremenek, M. Gleicher, "Configuring and Tuning Archival Storage System." *Proceedings of the 16th IEEE Symposium on Mass Storage Systems*, March 1999.

Moore, R., R. Marciano, M. Wan, T. Sherwin, R. Frost, "Towards the Interoperability of Web, Database, and Mass Storage Technologies for Petabyte Archives." *Proc. Fifth Conference on Mass Storage Systems and Technologies*, (1996).

Moore, R., T. Prince, M. Ellisman, "Data-Intensive Computing and Digital Libraries." CACM 41(11): 56-62 (1998).

Moore, Reagan W., "High Performance Data Assimilation." *Proceedings of the Committee on Information and Communications R&D (CIC) of the National Science and Technology Council*, July 1995.

Moore, Reagan W., "Enabling Petabyte Computing." *The Unpredictable Certainty, Information Infrastructure through 2000*, National Academy Press, (1997).

National Oceanographic Partnership Program, "Toward a U.S. Plan for an Integrated, Sustained Ocean Observing System." A report prepared on behalf of the National Ocean Research Leadership Council, (1999).

National Research Council, "Illuminating the Hidden Planet: The Future of Seafloor Observatory Science." Washington, D.C.: National Academy Press, (2000).

NEPTUNE Documents and Reports, www.neptune.washington.edu/pub/documents/documents.html

NEPTUNE Phase 1 Partners (University of Washington, Woods Hole Oceanographic Institution, Jet Propulsion Laboratory, Pacific Marine Environmental Laboratory), "Real-time, Long-term Ocean and Earth Studies at the Scale of a Tectonic Plate." NEPTUNE Feasibility Study (prepared for the National Oceanographic Partnership Program), University of Washington, Seattle, (2000).

Network for Earthquake Engineering Simulation, www.nees.org

NPACI Data Intensive Computing Environment Thrust Area. www.npaci.edu/DICE

NSF Geosciences Beyond 2000: Understanding and Predicting Earth's Environment and Habitability. www.nsf.gov/pubs/2000/nsf0028/nsf0028.htm

NSF Ocean Sciences Workshop on Information Technology Infrastructure. www.geo-prose.com/oceans_iti_trends/index.html

Ocean Information Technology Infrastructure Committee Report — An Information Technology Infrastructure Plan to Advance Ocean Sciences. www.geo-prose.com/oiti/report.html

Ocean Observatories Initiative Facilities Needs from UNOLS, Report of the UNOLS Working Group on Ocean Observatory Facility Needs, December 2003.

Ocean Sciences Decadal Committee, P. Brewer and T. Moore (Co-Chairs), (2001) "Ocean Sciences at the New Millennium," University Corporation for Atmospheric Research.

One Hundred Sixth Congress Of the United States of America at the Second Session, "The Oceans Act of 2000." December 2002.

OOI Brochure, Geosciences Professional Services, Inc. (editors), June 2003.

OOI Scientific Objectives and Network Design. July 2006.

Open Science Grid. www.opensciencegrid.org/

Orcutt, John (Co-Chair), Adam Schultz (Co-Chair), Jeremy Bloxham, Rhett Butler, John Collins, Robert Detrick, Kang Ding, Adam Dziewonski, Gary Egbert, Marcia McNutt, Barbara Romanowicz, Sean Solomon, Mark Zumberge, "DEOS Global Working Group Report," December 1999.

Orion CI Committee, ORION Data Policy Draft, September 2005 (updated March 2006).

Orion Cyberinfrastructure: Concept of Operations. Version Or08 2006.05.11

ORION Executive Steering Committee. (2005) "Ocean Observatories Initiative Science Plan." Washington, DC, 102 pp.

ORION, "Advisory Committees—Cyberinfrastructure." www.orionprogram.org/organization/committees/cyberinfrastructure.html

ORION, "Cyberinfrastructure Conceptual Architecture." www.orionprogram.org/organization/committees/ciarch/default.html

ORION, Conceptual Architecture Materials Concept of Operations (version 5/11/06), www.orionprogram.org/organization/committees/ciarch/ConceptOfOperations.pdf

ORION, Conceptual Architecture Materials Executive Summary, www.orionprogram.org/organization/committees/ciarch/CI_ExecSumm.pdf

ORION, Plenary Presentation (Conceptual Design Review, Aug. 14-18, 2006), www.orionprogram.org/organization/committees/ciarch/Plenary_CI.pdf

ORION, Request for Assistance?Conceptual Science Experiments, www.orionprogram.org/news/default.html

ORION,"ORION Coastal Scale Observatory Science Requirements." Draft Report of the ORION Coastal Scale Observatory Refinement Committee, May, 2005.

ORION,"ORION Global Buoy Array Science Requirements," June 2005.

ORION,"ORION Regional Cabled Observatory Science Requirements." Draft Report of the ORION Regional Cabled Observatory Refinement Committee, May, 2005.

OV1: High-Level Operational Concept Graphic (ver. 5/30/06), www.orionprogram.org/organization/committees/ciarch/CI_OperConcpt_05302006.pdf

OV2: Operational Node Connectivity Description (ver. 7/26/06), www.orionprogram.org/organization/committees/ciarch/CI_OperNode_07262006.pdf

OV4: Organizational Relationships Chart (ver. 8/3/06), www.orionprogram.org/organization/committees/ciarch/CI_OrgRelat_08032006.pdf

OV5: Operational Activity Model (ver. 5/30/06), www.orionprogram.org/organization/committees/ciarch/CI_OperActiv_05302006.pdf

Particle Physics Data Grid. www.ppdg.net/

Poulsen, A.J., Eickstedt, D.P., "Bearing Stabilization and Tracking for an AUV with an Acoustic Line Array." In Proceedings of the IEEE/MTS Oceans'06 Conference, Boston, MA, September 18-21, 2006.

Putman, J.R., *Architecting with RM-ODP*, Prentice-Hall, 834 pp., (2000).

Rajasekar, A., M. Wan, R. Moore, W. Schroeder, G. Kremenek, A. Jagatheesan, C. Cowart, B. Zhu, S.-Y. Chen, R. Olschanowsky, "Storage Resource Broker - Managing Distributed Data in a Grid." *Computer Society of India Journal*, special issue on SAN, (2003).

Rajasekar, A., R. Moore, B. Ludäscher, I. Zaslavsky, "The Grid Adventures: SDSC's Storage Resource Broker and Web Services in Digital Library Applications." *4th Russian Conference on Digital Libraries*, Dubna, Russia, October 2002.

Rajasekar, Arcot, F. Vernon, K. Lindquist, S. Lu, J. Orcutt, R. Moore, "Accessing Sensor Data Using Meta Data: A Virtual Object Ring Buffer Framework." *VLDB Workshop on Data Management for Sensor*. Trondheim, Norway. (2005).

Rajasekar, Arcot, M. Wan, R. Moore, "mySRB and SRB, Components of a Data Grid." *11th High Performance Distributed Computing Conference*,. Edinburgh, Scotland, July 2002.

Reality, John Wiley, pp. 171-198, (2003).

RECONN, "Report of the Cabled Regional Observatory Workshop, October 7-10, 2003," San Francisco, CA. March, 2004.

Review Report of the NSF Conceptual Design Review Panel for the Ocean Observatories Initiative. Monterey Bay Aquarium Research Institute. September 2006.

ROADNet, Real-time Observatories, Applications, and Data management Network. roadnet.ucsd.edu/

Royce, W., (1998) "Software Project Management: A Unified Framework." New York: Addison-Wesley, 406 pp.

Rudnick, D.L. and M.J. Perry (eds.), (2003) "ALPS: Autonomous and Lagrangian Platforms and Sensors," Workshop Report, 64pp., www.geo-prose.com/ALPS

Schofield, Oscar and Margaret K. Tivey, (Co-Chairs), (2004) "Orion Ocean Research Interactive Observatory Networks." A report of the Workshop held January 4-8, 2004, San Juan Puerto Rico.

Sea Change in Ocean Observing for the Ocean Planet: Toward Implementation of an Integrated and Sustained Ocean Observing System, www.nopp.org/Dev2Go.web?id=220628&rnd=1869

SRB, Storage Resource Broker, Version 3.4.2, (www.sdsc.edu/srb), (2006).

Stockinger, Heinz, O. Rana, R. Moore, A. Merzky, "Data Management for Grid Environments." *European High Performance Computing and Networks Conference*, Amsterdam, Holland. June 2001.

Surface Currents Mapping Initiative.cordc.ucsd.edu/projects/mapping/

Trends in Information Technology Infrastructure in the Ocean Sciences, Report of the National Science Foundation Ocean Sciences Workshop on Information Technology Infrastructure held May 21-23, 2003 at the Skamania Lodge, Stevenson, WA.

Vernon, Frank, T. Hansen, K. Lindquist, B. Ludäscher, J. Orcutt, A. Rajasekar, "ROADNet: A Real-time Data Aware System for Earth, Oceanographic, and Environmental Applications." *Eos Trans. AGU* 84(46), Fall Meeting Supplement, Abstract U21A-06. (2003).

MATTHEW ARROTT

650-814-4250 marrott@ucsd.edu

PRODUCT DEVELOPMENT & ENGINEERING MANAGEMENT

Financial Services ◆ eCommerce Applications ◆ Information Technologies

Results-driven Product Development and Engineering manager with a proven track record of delivering products that exceeded business and technical objectives. Demonstrated ability to transform business goals and requirements into vision and final product by delivering on the complete product lifecycle from conceptualization through market and product development to customer acceptance.

Skilled at focusing product and engineering teams to execute on business strategies and achieving financial objectives within complex requirements and time constraints, grounded by years of experience in demanding product development and engineering environments. Delivered ground-breaking products across a wide range of disciplines for such companies as *Currenex, DreamWorks, Autodesk* and *NCSA*.

ENGINEERING EXPERIENCE

An accomplished technologist skilled in architecture and software development practices with extensive experience in Trading & Settlement services, as well as, Information & Graphics applications. In-depth senior manager setting strategies, roadmaps, resources and budgets, with years of experience managing rapidly growing software development teams backed by a solid background in Product Management and Business Development.

Software Development

- 10+ years experience in requirements analysis, design, coding and unit testing of scalable, distributed, fault-tolerant applications in UNIX and NT environments with a thorough understanding of synchronous and asynchronous communication architectures
- 15+ years design experience in object-oriented design methodology and application/service development in C, C++, Java and J2EE environments
- Project and Development experience with WebObjects, Weblogic, Oracle, Veritas and VeriSign
- Working knowledge of HTTP/HTML, XML, SOAP, XSLT, as well as, the ebXML, Web Services and Grid stacks, with a particular interest in the transformation of the Grid architecture into the Web Service Resource Framework, WSRF.
- An accomplished User Interface and Information architect with MS WFC, X/Motif, and Swing grounded by 8 years in Information design and presentation.

Network Operations

- 6+ years design and management experience building and operating 24x7 geographically diverse Network Operation Centers with a facile architectural understanding of all the major components from IP, BGP, and VLANs through Switches, Routers, Firewalls, IDSs, and Load Balancing to monitoring and maintaining 99.999 SLAs.

Standards Participation

- Active member TWIST, FIX and ISO 20022 WG4. Past member of WS-I as liaison between Technical and Marketing working groups
- Collaborator with FpML, OASIS ebBP (formerly BPSS), and W3C CDL

PROFESSIONAL EXPERIENCE

CALIFORNIA INSTITUTE OF TELECOMMUNICATION & INFORMATION TECHNOLOGY, UNIVERSITY of CALIFORNIA SAN DIEGO Oct. 2004 - present

E-Science Program Manager and System Architect

- Program development of the E-Science initiative focused on the federation of loosely-coupled distributed resources in support of large-data E-Sciences applications. E-Science refers to large scale science carried out through distributed global collaborations
- Leading the NSF ITR grant to research the design and development of the national cyberinfrastructure to support a global and permanent integrated interactive instrument capability; facilitating the simultaneous and continuous measurements of physical, chemical, biological and geological properties through major portions of the earth-ocean-atmosphere system

TWIST PROCESS INNOVATIONS, LIMITED, London Jun 2001 - present

Chief Technical Officer

- Executive member and CTO for TWIST a not-for-profit industry group delivering non-proprietary XML-based financial process standards for Wholesale financial market, Commercial payments and Working capital finance, and Cash management

CURRENEX, INC., London, New York, Singapore, Menlo Park Dec. 1999 – Jul. 2003

The leading independent global electronic multi-bank foreign currency exchange linking Fortune 1000 worldwide

Vice President, Product Development

- Member of executive management directing a global team located in London, New York, Menlo Park, and Singapore with a budget responsibility of $10+m, managing 4 departments, Product Management, Engineering, Operations and Member Services
- Designed, architected, and deployed the industry's first multi-bank "Executable Streaming Price" product; distributing over 2 million pricing events daily from tier one banks globally.

Vice President, Engineering & Operations

- Redesigned and deployed the service infrastructure from a traditional 3 tier web architecture to an asynchronous event-driven service oriented architecture
- Redesigned and deployed the client interface from DHTML/JavaScript based architecture to a lightweight, "Zero Administration" Java Swing applet with bidirectional asynchronous messaging to the server

Director, Operations and Integration Services

- Designed and deployed the first Internet based Foreign Exchange transaction network; meeting the strict security, reliability, and availability requirements of the top 50 global banks

DREAMWORKS, SKG, Glendale, CA & **SILICON GRAPHIC, INC.**, Mountain View, CA Jun. 1995 – Oct. 1998

Head of Software R&D

- Senior member of the technology department directing all software development for the DreamWorks Feature Animation division. Responsible for 3 departments, research, production software as well as production and asset management software supporting 2 feature productions running concurrently

Software Development Manager & Application Architect

- Development Manager and Application Architect developed for *DreamWorks SKG's* Feature Animation a scene design and planning application used to integrated 2D traditional animation with 3D computer graphic imagery for their films *Prince of Egypt* (1998) and *Road to El Dorado* (2000) .

AUTODESK, INC., San Raphael, CA Jan. 1993 – Jun. 1995

Software Development Manager & Systems Architect

- Architected and executed 2-year project to replace entire AutoCAD graphics subsystem and delivered the *HEIDI* Graphics System, also used as the underpinnings of 3D Studio Max
- Directed all AutoCAD Graphics, Render, and Driver development groups

NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS Feb. 1987 – Jan. 1993
UNIVERSITY OF ILLINOIS, URBANA-CHAMPAIGN

Project Manager & Systems Architect

- Founding member of the Scientific Visualization program that pioneered the definition and promotion of computer graphic imagery for the visual exploration and presentation of complex scientific numerical modeling.
- Design and deployment one of the first distributed rendering farm environments coupled with an automated digital recording facility

WAVEFRONT TECHNOLOGIES, INC., Santa Barbara, CA Jan. 1986 – Jan. 1987

One of the original 3D graphics software vendors providing tools for the film & video production and industrial design industries

Graphics Programmer

- Developed the texture mapping module for the renderer, image to mipmap texture conversion, and the mapping of texture coordinates to geometry for the Advanced Visualizer Product.
- Designed and developed the 2nd generation geometry subsystem for the Advanced Visualizer Product
- Co-designed the user interface management system deployed in the Personal and Data Visualizer as well as in the 2nd generation Advanced Visualizer

CONTRACT EXPERIENCE

NOVEMBER GROUP, LLC Oct. 1998 – present

Principal

Principal and senior consultant for the management and design consultancy focused on product and service development. Client engagements include:

COMMERCENET CONSORTIUM, INC., Palo Alto Aug. 2003 - Sept. 2004

Vice President, Engineering and Operations

- Developed service vision for a distributed multi-industry interoperability platform based on a Self-Provisioning Services framework and shared registry based on Web Services and Grid infrastructure.
- Design of the California Clinical Data Exchange in collaboration with the California Institute for Telecommunications and Information Technologies and the Mayo Clinic
- Designed, staffed and deployed their Network Operation Center from ground zero to operation, in 45 days, allowing CommerceNet to recognize their first revenue from their first operating subsidiary on schedule and under budget

BLUE MARTINI SOFTWARE, INC., San Mateo, CA Jun. 1999 - Jan. 2000

Director of Integration

- Worked directly with executive management to develop their Enterprise Application Integration, EAI, strategy and build their integration team for their enterprise web-based retail and manufacturing application.
- Led negotiations and development efforts for all EAI projects with third party enterprise application vendors, e.g. SAP, JDA, Retek, Remedy, Apropos, Facetime, utilizing EAI solutions from Tibco, Vitria, and WebMethods.

FASTV, INC., Los Angeles, CA Jan. – Jul. 1999

Director of Operations

- Worked directly with executive management to design, staff, and deploy their Network Operations Center. The facility supported video indexing, archiving, search and streaming services for news and sport content over the Internet.

EDUCATION

HARVARD UNIVERSITY, Cambridge, MA
Class ’81, Bachelor of Arts
Varsity Rowing ‘80, ‘81

ALAN D. CHAVE

Senior Scientist
Deep Submergence Laboratory, Department of Applied Ocean Physics and Engineering
Woods Hole Oceanographic Institution, Woods Hole, MA 02543
508-333-4711 alan@whoi.edu http://www.whoi.edu/science/AOPE/people/achave/

PROFESSIONAL PREPARATION

B.S. (Physics)	1975	Harvey Mudd College
Ph.D. (Oceanography)	1980	MIT/WHOI Joint Program in Oceanography
Postdoctoral Investigator	1980-2	Scripps Institution of Oceanography
Chartered Statistician #1773	2003	

APPOINTMENTS HELD

Senior Scientist	1993-	Dept of AOPE (2000-present), Dept of Geology and Geophysics (1993-2000), Woods Hole Oceanographic Institution
Visiting Professor	1998-99	Earthquake Research Institute, University of Tokyo, Tokyo, Japan
Associate Scientist with Tenure	1992-93	Dept. of Geology and Geophysics, Woods Hole Oceanographic Institution
Member of the Technical Staff	1986-91	Physics Research Division, AT&T Bell Laboratories
Asst. and Assoc. Res. Geophysicist	1982-86	Scripps Institution of Oceanography, University of California, San Diego

RECENT PROFESSIONAL ACTIVITIES

Member	2006	ORION Cyberinfrastructure Design Team
Member	2005-6	ORION Engineering Committee
Chair	2003-4	UNOLS Working Group on Ocean Observatory Facility Needs
Editor-in-chief	1991-7	*Reviews of Geophysics*

CRUISE EXPERIENCE: Participant in 32 research cruises, 21 as chief scientist

PUBLICATIONS: Author or co-author of 102 refereed papers, 2 patents, 77 technical reports or extended abstracts, and 242 conference presentations

FIVE MOST RELEVANT PUBLICATIONS

1. Chave, A.D., E. Massion, and H. Mikada, Science requirements and the design of cabled ocean observatories, *Ann. Geophys.*, *49*, 569-579, 2006.

2. Chave, A.D., G. Waterworth, and A. Maffei, Cabled ocean observatory systems, *MTS Journal*, 38, 31-43, 2004.

3. St. Arnaud, B., A.D. Chave, A. Maffei, E. Laszowka, L. Smarr, and G. Gopalan, An integrated approach to ocean observatory data acquisition/management and infrastructure control using web

services, *Mar. Tech. Soc. J.*, *38*, 155-163, 2004.

4. Chave, A.D., and D.J. Thomson, A bounded influence regression estimator based on the statistics of the hat matrix, *J. Roy. Stat. Soc., Series C (Appl. Statist.)*, 52, 307-322, 2003.

5. Petitt, R.A., D. Harris, F.B. Wooding, J.W. Bailey, J. Jolly, E. Hobart, A.D. Chave, F.K. Duennebier, R. Butler, A.D. Bowen, and D.R. Yoerger, The Hawaii-2 Observatory, *IEEE J. Ocean Eng.*, *27,* 245-253, 2002.

FIVE OTHER SIGNIFICANT PUBLICATIONS

1. Baba, K., A.D. Chave, R.L. Evans, G. Hirth, and R.L. Mackie, Mantle dynamics beneath the East Pacific Rise at 17S: Insights from the Mantle Electromagnetic and Tomography (MELT) experiment, *J. Geophys. Res., 111*, B02101, doi: 10.1029/2004JB003598, 2006.

2. Kanzow, T., U. Send, W. Zenk, A.D. Chave, and M. Rhein, Monitoring the integrated deep meridional flow in the tropical North Atlantic: Long-term performance of a geostrophic array, *Deep Sea Res. I, 53*, 528-546, 2006.

3. Baba, K., and A.D. Chave, Correction of seafloor magnetotelluric data for topographic effects during inversion, *J. Geophys. Res.*, *110*, B12105, doi: 10.1029/2004JB003463, 2005.

4. Chave, A.D., D.S. Luther, and C.S. Meinen, Correction of motional electric field measurements for galvanic distortion, *J. Atm. Oceanic Tech.*, *21*, 317-330, 2004.

5. White, S.N., A.D. Chave, and G.T. Reynolds, Investigations of ambient light emission at deep-sea hydrothermal vents, *J. Geophys. Res.*, *108* (1), doi: 10.1029/2000JB000015, 2002 (Editor's Choice, *Science, 295*, 1427, 2002).

SCIENTIFIC COLLABORATORS/CO-EDITORS (PAST 4 Y OUTSIDE WHOI)

Mark Abbott (OSU)
S. Michael Angel (U S. Car.)
Matthew Arrott (UCSD)
Kiyoshi Baba (U Tokyo)
James Bellingham (MBARI)
Dale Chayes (LDEO)
Bernard Coakley (UAF)
John Delaney (UW)
Fred Duennebier (UH)
J.H. Filloux (SIO)
Frank Flechtner (Postsdam)
George Jiracek (SDSU)
Alan G. Jones (DIAS)
Torsten Kanzow (SOC)
Uwe Send (SIO)
Ed Lazowska (UW)
Douglas Luther (UH)
Randy Mackie (GSY-USA)
Gene Massion (MBARI)
Chris Meinen (NOAA)
Hitoshi Mikada (JAMSTEC)
John Orcutt (SIO)
Bill St. Arnaud (Canarie)
Nobukazu Seama (Kobe)
Larry Smarr (UCSD)
Ken Smith (SIO)
Pascal Tarits (Brest)
David Thomson (Queens)
Hisashi Utada (U Tokyo)
Frank Vernon (SIO)
Gary Waterworth (Alcatel)

GRADUATE ADVISEES: Sheri White (Ph.D. 2000), Anna Michel (Ph.D. candidate)

POSTDOCTORAL ADVISEES: Sheri White, Pamela Lezaeta, Kiyoshi Baba

SCOTT M. GLENN

Coastal Ocean Observation Lab
Institute of Marine and Coastal Sciences
Rutgers, The State University of New Jersey
71 Dudley Road, New Brunswick, NJ 08901-8521
(Tel.) 732-932-6555, x506; (Fax) 732-932-8578, E-mail: glenn@marine.rutgers.edu
http://marine.rutgers.edu/cool; http://www.thecoolroom.org; http://www.coolclassroom.org

EDUCATION:

University of Rochester, Rochester, New York, B.S. with High Honors, Geomechanics, 1978.
MIT/WHOI Joint Program, Cambridge & Woods Hole, Massachusetts, Sc.D., Ocean Engineering, 1983.

APPOINTMENTS:

Professor, Rutgers University, New Brunswick, New Jersey, 1998 - present.
Associate Professor, Rutgers University, New Brunswick, New Jersey, 1990 – 1998.
Project Scientist, Harvard University, Cambridge, Massachusetts, 1986 – 1990.
Research Engineer, Shell Development Company, Houston, Texas, 1983 – 1986.

ADDITIONAL APPOINTMENTS:

Adjunct Scientist, Mote Marine Laboratory, Sarasota, Florida, 2001-Present.
Vice-Chair, Department of Marine and Coastal Sciences, Rutgers University, 1999-present.

FIVE SELECTED RELATED PUBLICATIONS (Over 100 total):

Glenn, S.M. and O.M.E. Schofield, 2004. Observing the oceans from the COOLroom: Our history, experience, and opinions, *Oceanography*, V16, N4, pp. 37-52.

Glenn, S.M., O.M.E. Schofield, T. Dickey, R. Chant, J. Kohut, J. Bosch, L. Bowers, L. Creed, C. Haldemann, E. Hunter, J. Kerfoot, C. Mudgal , M. Oliver, H. Roarty, E. Romana, M. Crowley, D. Barrick and C. Jones, 2004. The Expanding Role of Ocean Color & Optics in the Changing Field of Operational Oceanography, *Oceanography*, June, pp 86-95.

Oliver, M., Kohut, J., Irwin, A., Schofield, O., **Glenn, S.M., Bissett, W.P., and Moline, M. A., 2004. Bioinformatic Approaches for Objective Detection of Water Masses, *Journal Geophysical Research,* 109, C12S05, doi: 10.1029/2003JC001985.

** Durski, S. M., **S. M. Glenn**, and D. B. Haidvogel, 2004. Vertical mixing schemes in the coastal ocean: Comparison of the level 2.5 Mellor-Yamada scheme with an enhanced version of the K profile parameterization, *Journal Geophysical Research*, *109*, C01015, doi:10.1029/2002JC001702.

Wilkin, J.L., Arango, H.G., Haidvogel, D.B., Hedstrom, K.S., **Lichtenwalner, C.S., and **Glenn, S.M**., 2005. A regional ocean modeling system for the Long-term Ecosystem Observatory. *Journal Geophysical Research,110,* C06S91, doi:10.1029/2003JC002218.

FIVE SELECTED OTHER PUBLICATIONS: (** = Student)

Glenn, S. M., Arnone, R., Bergmann, T., Bissett, W. P., Crowley, M., Cullen, J., Gryzmski, J., Haidvogel, D., **Kohut, J., Moline, M. A., Oliver, M., Orrico, C., Sherrell, R., Song, T., Weidemann, A., Chant, R., and Schofield, O., 2004. The biogeochemical consequences of summer upwelling off the New Jersey coast. *Journal Geophysical Research, 109*, C12S02, doi:10.1029/2003JC002265.

Chant, R, **Glenn, S. M.**, and **Kohut, J., 2004. Flow reversals during upwelling conditions on the New Jersey inner shelf. *Journal Geophysical Research, 109,* 12S03, DOI:101029/2003JC001941.

Kohut, J., **Glenn, S.M., and *Chant, R., 2004. Seasonal Current Variability on the New Jersey Inner Shelf. *Journal Geophysical Research, 109*, C07S07, doi:10.1029/2003JC001963.

Styles, R., and **Glenn, S. M., 2005. Long-term sediment mobilization at LEO-15. *Journal Geophysical Research,* 110, C04S90, doi:10.1029/2003JC002175.

Kohut, J., **Glenn, S. M., and Paduan, J., 2005. The inner-shelf response to tropical storm Floyd. *Journal Geophysical Research,* in press.

FIVE SYNERGISTIC ACTIVITIES:

1) Member, NSF ORION International Ocean Observatory Committee.
2) Steering Committee, Ocean.US Surface Current Mapping Initiative.
3) Co-PI, NOAA Middle Atlantic Coastal Ocean Observing Regional Association (MACOORA)
4) Member, ONR Glider Consortium
5) Advisory Committee, NSF Center for Ocean Science Education Excellence, Mid Atlantic (COSEE-MA)

COLLABORATORS & OTHER AFFILIATIONS

(a) Collaborators and Co-Editors.

Ken Able (Rutgers), Hernan Arango (Rutgers), Bob Arnone (NRL), Roni Avissar (Duke), John Bane (UNC), Andrew Barnard, Don Barrick (CODAR Ocean Sensors), Jack Barth (Oregon State), Jim Bellingham (MBARI), Tom Berger (SAIC), Trisha Bergmann (U. Maine), Paul Bissett (FERI), Shelly Blackwell (CalPoly), Allan Blumberg (Stevens), William Boicourt (U. Maryland), Philip Bogden (GoMOOS), Emanuel Boss (U. Maine), Louis Bowers (Rutgers), Tod Bowers (NRL), Jack Burbank (JHU/APL), Pat Burke (Stevens), William Browning (Applied Mathematics Inc), Mike Bruno (Stevens), Brad Butman (USGS), Tom Campbell (Webb Research), John Case (UCSB), Rob Cermak (U. Alaska), Grace Chang (UCSB), Bob Chant (Rutgers), Francisco Chavez (MBARI), Bob Chen (UMass), Yi Chao (JPL), Jim Churchill (WHOI), Peter Cornillon (URI), Bob Cowen (RSMAS), Liz Creed (Rutgers), Mike Crowley (SeaSpace), Jay Cullen (WHOI), Carla Curran, Hans Dam, Meredith Dermarest, Ed Dever, Mike DeLuca (Rutgers), Tommy Dickey (UCSB), Paul DiGiacomo, Paul Dragos (Battelle), Fred Duennebier (U. Hawaii), Rich Dunk (Rutgers), Scott Durski (Oregon State), Jim Edson (WHOI), John Fracassi (Rutgers), Tom Frazier (U. Florida), Brian Fullerton (Stevens), Ann Gargett (Old Dominion), Rocky Geyer (WHOI), Lou Goodman (UMass), John Govoni (NOAA), Fred Grassle (Rutgers), Tom Gross (NOAA), Joe Gryzmski (Rockerfeller U.), Dale Haidvogel (Rutgers), J Hamrick, John Hare (NOAA), Courtney Harris (VIMS), Kate Hedstrom (U. Alaska), Tom Herrington (Sea Grant), Christy Herron (MBARI), Jim Hillier (Mote), Dick Hires (Stevens), Bob Houghton (Lamont), Debra Iglesias-Rodriguez (U. Nottingham), Andrew Irwin (CUNY), Rick Jahnke (Skidaway), Ken Johnson, Clayton Jones (Webb Research), Alex Kahl (Rutgers), Dave Karl (U. Hawaii), William Kasch (JHU/APL), Tim Keen (NRL), John Kerfoot (Rutgers), Gary Kirkpatrick (Mote Marine Lab), Josh Kohut (Rutgers), Pierre Lermusiaux (Harvard), Ed Levine (NUWC), Sage Lichtenwalnar (U. South Florida), Pete Lilleboe (CODAR Ocean Sensors), Belinda Lipa (CODAR Ocean Sensors), Tom Lee (RSMAS), Janice McDonnell (Rutgers), W Miller, Ralph Milliff (CORA), Curt Mobley (Sequoia Scientific), Mark Moline (CalPoly), Andrew Moore (U. Colorado), Chhaya Mudgal (Rutgers), Lauren Mullineaux (WHOI), Dave Musgrave (U. Alaska), Robert Nichols (JHU/APL), Jim O'Donnell (UConn), Matt Oliver (Rutgers), Joan Oltman-Shay (NWRA), Chris Orrico (UCSB), Jefff Paduan (Naval PG School), Hai Pan (Rutgers), Bruce Parker (NOAA), A Pence, Ola Person, E Peters, Allan Pluddeman (WHOI), K Prasad, David Porter (JHU/APL), Mike Purcell (WHOI), Kelly Rankin (Bigelow), Clare Reimers (Oregon State), John Reinfelder (Rutgers), Hugh Roarty (Rutgers), Allan Robinson (Harvard), Oscar Schofield (Rutgers), Rob Sherrell (Rutgers), Chris Sherwood (USGS), P Shrestha, Rich Signell (USGS), Tony Song (JPL), Heidi Sosik (WHOI), Rich Styles (U. South Carolina), Eric Terrill (Scripps), Carolyn Thoroughgood (U. Delaware), Sasha Tozzi (VIMS), Mike Twardowski (URI), Chris von Alt (WHOI), Doug Webb (Webb Research), Alan Weidemann (NRL), Bob Weller (WHOI), John Wiggins (Rutgers), John Wilkin (Rutgers), Don Wright (VIMS), Joel Young (RD Instruments), Phoebe Zhang (Rutgers), Mark Zumberge (UCSD), Mung Zuo (UMass)

(b) Graduate Advisor – William D. Grant (WHOI, deceased)

(c) Thesis Advisor and Postgraduate-Scholar Sponsor

Students – Won Chul Cho, Michael Crowley , Tom Fikslin, Raju Datla, Richard Styles, Hai Pan, Hoyle Lee, Cheng-Chin Yang, Scott Durski, Majid Yavary , Justin Sharp, Hongguang Ma, Hugh Roarty , Josh Kohut, Trisha Bergmann, Kristie Andresen, Matt Oliver, Roy Messaros, Louis Bowers, Donglai Gong, YongChul Lee.
Post-docs – Tim Keen (NRL), Anna Mateoda (U. Delaware), Robert Chant (Rutgers), Richard Styles (U. South Carolina).

FIVE HONORS AND AWARDS:

1) Special Edition of *The Bulletin, New Jersey Academy of Science,* dedicated to the Rutgers University Graduate Course in Remote Sensing of the Ocean and Atmosphere, Scott Glenn and Jim Miller, 1996.
2) Named a Teaching Fellow by the Rutgers University Teaching Excellence Center, 1997.
3) First recipient of The Rutgers University President's Scholar-Teacher Award for using research initiatives to enhance teaching, 2000.
4) New Jersey Assembly Resolution No. 209, Commending Rutgers University Coastal Ocean Observation Lab for its Research and Education Outreach Programs, 2002.
5) Rutgers Cook College Team Award, Rutgers University Coastal Ocean Observation Lab, 2006.

John A. Orcutt
Associate Vice Chancellor for Research
University of California, San Diego
La Jolla, CA 92093-0210

Birth	29 AUG 43
Telephone	(858) 822-6378
FAX	(858) 822-3372
Email:	jorcutt@ucsd.edu

Place of Birth Holyoke, CO **Nationality** USA

Education

B.S.	1966	U.S. Naval Academy, Annapolis, MD
M.Sc.	1968	University of Liverpool, UK
Ph.D.	1976	University of California, San Diego

Research and/or Professional Experience

2006 - present	Associate Vice Chancellor – Research Affairs; Director, UCSD Center for Earth Observations & Applications
2002-2006	Deputy Director, Scripps
1984-present	Professor of Geophysics, Scripps Institution of Oceanography, UCSD
1984-2002	Director, Cecil H. & Ida M. Green Institute of Geophysics and Planetary Physics, UCSD
1982-1984	Associate Professor of Geophysics, Scripps Institution of Oceanography, UCSD
1980-1982	Associate Research Geophysicist, Scripps Institution of Oceanography, UCSD
1977-present	Visiting Associate in Geophysics, California Institute of Technology
1977-1980	Assistant Research Geophysicist, Scripps Institution of Oceanography, UCSD
1976-1977	Postgraduate Research Geophysicist, Scripps Inst. of Oceanography, UCSD
1967-1973	Chief Engineer, the nuclear submarine USS Kamehameha, U.S. Navy

Honors and Awards

Trident Scholar, U.S. Naval Academy	1965-1966
Graduate, 3rd in Class, U.S. Naval Academy	1966
Summer College Intern Program, U.S. Dept. of State	1966
Fulbright Scholar, United Kingdom	1966
Woods Hole Visiting Scholar	1980
Newcomb-Cleveland Prize from American Assoc. for Advancement of Science	1980
Fellow, American Geophysical Union	1989
Maurice Ewing Medal, American Geophysical Union	1994
Secretary of the Navy/Chief of Naval Operations Oceanography Chair	1996-2002
Member, American Philosophical Society	2002-
President, American Geophysical Union	2004-2006
Past-President, American Geophysical Union	2006-2008

Research Interests

Applications of information technology to integrating global observations
Ocean bottom seismology and the structure of mid-ocean ridges & hotspots
Wireless networking and real-time data management
Global and crustal seismic tomography

Relevant Publications

Suyehiro, K., J.-P. Montagner, R.A. Stephen, E. Araki, T., Kanazawa, J. Orcutt, B. Romanowicz, S.Sacks, and M. Shinohara (2006), "Ocean Seismic Observatories.", Oceanography **19**(4),

Taesombut, N., F. Uyeda, AA. Chien, L. Smarr, T. DeFanti, P. Papadopoulos, J. Leigh, M. Ellisman, and J. Orcutt (2006), "The OptIPuter: High-Performance, Q0S-Guaranteed Network Service for Emerging E-Science Applications." IEEE Communications Magazine **44**(5), 38-45, 10.1109/MCOM.2006.1637945.

Berger, J., J.A. Orcutt, and F. Vernon (2005). "HighSeasNet: Providing Internet to the UNOLS fleet: A model for real-time Internet data collection from the ORION platforms." Sea Technology **46**(6), 17-20.

Sutherland, F.H., F.L. Vernon, and J.A. Orcutt (2004). "Results from OSNPE: Improved teleseismic earthquake detection at the seafloor." Bull. Seismological Soc. of Am. **94**(5) 1868.

Orcutt, J. (2003) "The Ocean Research Interactive Observatory Networks (ORION) program." EOS Trans AGU **46**(40), 44, 10.1029/2003EO400005.

Sutherland, F.H., F.L. Vernon, and J.A. Orcutt (2004). "Results from OSNPE: Improved teleseismic earthquake detection at the seafloor." Bull. Seismological Soc. of Am. **94**(5) 1868-1878.

Five Other Significant Publications

Orcutt, J. (2005) "Global scale sensor networks – opportunities and challenges" Information Processing in Sensor Networks. 434, 10.1109/IPSN.2005.1440965.

Blackman, D.K., C. de Groot-Hedlin, P. Harben, A. Sauter, and J. Orcutt, (2004). "Testing low/very low frequency acoustic sources for basin-wide propagation in the Indian Ocean." Jour. Acoustical Soc. Am. **116**(4), 2057-2066, 10.1121/1.1786711.

Stephen, R.A., F N. Spiess, J.A. Collins, J.A. Hildebrand, J.A. Orcutt, K.R. Peal, F.L. Vernon, and F.B. Wooding (2003). "Ocean seismic network pilot experiment." Geochemistry Geophysics Geosystems **4**(910), 1092, 10.1029/2002GC000485,2003.

Sandwell, D.T., S. Gille, and J. Orcutt (2003). "Bathymetry from space is now possible." EOS **84**(5), 37, 44.

Newman, H.H., M.H. Ellisman, and J.A. Orcutt (2003), "Data-intensive e-science frontier research." Communications of the ACM, **46**(11), 68-77.

Synergistic Activities (Last 12 Months)

Member, Ocean Research Integrated Observatories Networks Executive Steering Committee (2004-present)
Member, Board of Governors, Joint Oceanographic Inst, Inc (1984 – present)
Member, Board of Governors, Integrated Ocean Drilling Program (2002-present)
Member, JOI/CORE Ocean Council (2005-present)
Member, Board of Governors, Consortium for Ocean Research & Ed (2004-present)

Recent Collaborators:

A. Baggeroer (MIT), D.K. Blackman (SIO), A. Chave (WHOIC. Collins (NPS), J. Collins (SIO), C. deGroot-Hedlin (SIO), R. Detrick (WHOI), A.J. Harding (Scripps), M.A.H. Hedlin (SIO),G.M. Kent (SIO), S. Singh (IPGP-Paris), M. Sinha (Southampton), S. Solomon (DTM/CIW), S. R. Stephen (WHOI), S. Webb (SIO), C. Wolfe (DTM./NSF), J. Mutter (LDEO), F.L. Vernon (SIO), T. Wallace (UA), Arcot Rajasekar (UCSD/SDSC)

PhD Thesis Advisors: LeRoy Dorman, Freeman Gilbert (Scripps)

Graduate Students Supervised (Last 5 Years):

Fiona Sutherland, Sara Bazin, Renee Bulow

OSCAR M. SCHOFIELD

Coastal Ocean Observation Lab, Institute of Marine & Coastal Sciences
Rutgers University, New Brunswick, NJ 08901
(TEL) 732.932.6555, x. 548 (FAX) 732.932.8578
oscar@marine.rutgers.edu · http://www.marine.rutgers.edu/cool

A. PROFESSIONAL PREPARATION

1983-1987	B.A. in Aquatic Biology, Department of Biology, University of California, Santa Barbara
1989-1993	Ph.D. in Biology, Department of Biology, University of California, Santa Barbara
1994	Postdoctoral Researcher, Center for Remote Sensing and Environmental Optics, University of California, Santa Barbara
1994-1995	Postdoctoral Researcher, Southern Regional Research Center, Agriculture Research Service

B. APPOINTMENTS

2001-Present	Associate Professor, Institute of Marine and Coastal Sciences, Rutgers University
2001-Present	Adjunct Professor, California Polytechnic State University, San Luis Obispo, CA
2000-Present	Member of Rutgers Ocean Systems Engineering Center
1999-Present	Member of Rutgers Environmental Biophysics and Molecular Biology Program
1999-2005	Co-Director of the Coastal Ocean Observation Laboratory
1995-2001	Assistant Professor, Institute of Marine and Coastal Sciences, Rutgers University
1995-Present	Adjunct Research Scientist, Mote Marine Laboratory, Sarasota, FL
1995	Adjunct Professor of Biological Sciences, Loyola University, New Orleans, LA
1989-1990	Curator, Algal Culture Collection, Department of Biology, UCA, Santa Barbar

C. PUBLICATIONS (5 RELATED, 5 OTHER*) (** PAPER BY GRADUATE STUDENT)

Schofield, **O**., Bosch, J., Glenn, S. M., Kirkpatrick, G., Kerfoot, J., Moline, M., Oliver, M., Bissett, W. P. Harmful algal blooms in a dynamic environment: How can optics help the field-going and sample poor biologist? In Real Time Coastal Observing systems for ecosystems dynamics and harmful algal blooms. Babin, M. And Cullen, J. J. (Eds) UNESCO, Paris. (in press).

Irwin, A., Finkel, Z., **Schofield**, **O**., Falkowski P. 2006. Scaling-up from nutrient physiology to the size-structure of phytoplankton communities. Journal of Plankton Research 28: 1-13.

Iglesias-Rodriguez, D., **Schofield**, **O**., Batley, J., Probert, I., Medlin, L.K., Hayes, P.K. 2006. Extensive intraspecific genetic diversity in the marine coccolithophorid *Emiliania huxleyi*: The use of microsatellite analysis in marine phytoplankton populations studies. Journal of Phycology doi: 10.1111/j.1529-8817.2006.00231. 1-12.

Schofield, **O**., J. Kerfoot, K. Mahoney, M. Moline, M. Oliver, S. Lohrenz, and G. Kirkpatrick (2006), Vertical migration of the toxic dinoflagellate *Karenia brevis* and the impact on ocean optical properties. Journal of Geophysical Research, 111, C06009, doi:10.1029/2005JC003115

**Finkel, Z., V., Katz, M. E., Wright, J. D., Schofield, O., Falkwoski, P. G. 2005. Climatically-driven macroevolutionary patterns in the size of marine diatoms over the Cenozoic. Proceedings of National Academy of Sciences. 102(25): 8927-2932.

Wolfe, F., Grzebyk, D., **Schofield, **O**., Falkowski, P. G. 2005. The role and evolution of superoxide dismutases in algae. Journal of Phycology DOI: 10.1111/j.-1529-8817.2005.00086: 1-13.

**Jiang, L., Schofield, O., Falkowski, P. G. 2005. Adaptive evolution of phytoplankton cell size. American Naturalist. 166(4): 496-505.

Falkowski, P.G., M. Katz, A. Knoll, J. Raven, O. Schofield, M. Taylor (2004) The consequences of the evolution of eukaryotic phytoplankton. *Science*, 305: 354-360.

Oliver, M.J., S. Glenn, J.T. Kohut, A.J. Irwin, **O. **Schofield**, M.A. Moline, and W.P. Bisset (2004), Bioinformatic Approaches for Objective Detection of Water Masses on Continental Shelves. *J.Geophy. Res, 109*, C07S04, doi: 10.1029/2003JC002072

Schofield, **O**., R. Arnone, W.P. Bissett, T. Dickey, C. Davis, Z. Finkel, M. Oliver, M. A. Moline, (2004) Watercolors in the coastal zone: What can we see? *Oceanography*. 107: 28-37.

Schofield O., T. Bergmann, M. J. Oliver, A. Irwin, G. Kirkpatrick, W. P. Bissett, M. A. Moline, C. Orrico (2004), Inversion of spectral absorption in the optically complex coastal waters of the Mid-Atlantic Bight, J. Geophys. Res., 109, C12S04, doi:10.1029/2003JC002071.

D. SYNERGISTIC ACTIVITIES

2004	Science and Risk Reduction Team GOES-R Hyperspectral Environmental Suite (HES) Coastal Water (CW) Imager
2004	Steering Committee for Alliance of Coastal Technolgoies Autonomous Underwater Vehicle Workshop
2004-2006	Editorial Advisory Board Continental Shelf Research and Journal of Geophysical Research
2004-2006	North American Chair for the Oceanography Society Meetings in Paris, France, Spring 2005
2004-2006	Executive Steering Committee for the ORION Program

E. COLLABORATORS & OTHER AFFILIATIONS (LAST 48 MONTHS): B. Arnone (NRL), K. Benoit-Bird (Oregon), P. Bissett (FERI), S. Blackwell (CalPoly), A. Blumberg (Stevens), W. Boicourt (U. MD), P. Bogden (GoMOOS), E. Boss (U. Maine), M. Bruno (Stevens), T. Campbell (Webb Research), J. Case (UCSB), G. Chang (UCSB), B. Chant (Rutgers), B. Chen (UMass), M. Crowley (SeaSpace), T. Dickey (UCSB), R. Dunk (Rutgers), P. Falkowski (Rutgers)K. Fennel (Rutgers), T. Frazer (U. Florida), R Geyer (WHOI), F. Grassle (Rutgers), J. Gryzmski (Rockerfeller U.), D. Haidvogel (Rutgers), B. Houghton (Lamont), D. Iglesias-Rodriguez (U. Southhampton), A. Irwin (CUNY), C. Jones (Webb Research), William Kasch (JHU/APL), L. Kerkhoff (Rutgers), G. Kirkpatrick (Mote Marine Lab), J. Kohut (Rutgers), M. Moline (CalPoly), R. Nichols (JHU/APL), D. Porter (JHU/APL), M. Purcell (WHOI), J. Reinfelder (Rutgers), C. Thoroughgood (U. Del.), S. Tozzi (VIMS), J. Trowbridge (WHOI), M. Twardowski (URI), C.s von Alt (WHOI), D. Webb (Webb Research), A. Weidemann (NRL), J. Wiggins (Princeton), J. Wilkin (Rutgers), P. Zhang (Rutgers), M. Zuo (UMass).

Graduate and Postdoctoral Advisors: Barbara Prézelin (Ph.D.) – UCSB, David Millie (Post-doctoral) – Agricultural Research Service ***Thesis Advisor and Postgraduate-Scholar Sponsor***: Students – Joe Gyzymski (Ph.D.), Felisa Wolfe (Ph.D), Trisha Bergmann (Ph.D.), Zoe Finkel (Ph.D.), Mathew Oliver (Ph.D.) – Rachael Sipler (Ph.D.) – Alex Kahl (Ph. D.) – Jessie Sebbo (MS) Post-docs – Mark Moline, Yu Gao, Antionetta Quigg, Elena Litchman, Lin Jhang.

Honors and Awards:

Antarctic Service Medal (1988), University Research SCUBA certification (1988), University of California at Santa Barbara Travel Award (1992), University of CA Regents Fellowship Award (1992), Invited Scientist National Academy of Sciences and Max Planck for the German-American Frontiers of Science, Münich Germany (1997), Invited Participant National Academy of Sciences and Japan Science & Technology Corporation (JAMSTEC), Japanese-American Frontiers of Science Symposium (1999), Rutgers University Faculty Academic Service Increment Program (FASIP) Award (1998-2002), NJ State Legislation Resolution Assembly Resolution No. 209 recognizing RU COOL as a state resource (2003)

BIOGRAPHICAL SKETCH

Name	**Birthdate**
Frank L. Vernon III	06-Apr-54

Place of Birth	**Present Nationality**
Pasadena, California, USA	U.S.A.

Education

Degree	Year	Institution and Location
B.A.	1977	University of California, San Diego
Ph.D.	1989	University of California, San Diego

Major Research interests

Sensor Networks and Real Time Data Acquisition Systems
Digital Telemetry Seismic and Environmental Networks and Arrays
Time Series and Array Analysis Techniques
Seismicity and Seismic Structure of Central Asia and the Middle East
Local Earthquake Source Properties

Research and/or Professional Experience

2002-present Research Geophysicist, IGPP SIO, UCSD
1997-present Lecturer, IGPP, SIO, UCSD
1996-2002 Associate Research Geophysicist, IGPP SIO, UCSD
1990-1996 Assistant Research Geophysicist, IGPP, SIO, UCSD
1989-1990 Post Graduate Research Geophysicist, IGPP, SIO, UCSD

Relevant Publications:

T. Hansen, S. Tilak, S. Foley, K. Lindquist, F. Vernon, J. Orcutt (2006). ROADNet: A network of SensorNets. Proceedings of the First IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp 2006), in conjunction with LCN 2006.

C. Cotofana, L. Ding, P. Shin, S. Tilak, T. Fountain, J. Eakins, F. Vernon (2006). An SOA-based Framework for Instrument Management for Large-scale Observing Systems (USArray Case Study). IEEE International Conference on Web Services (ICWS) 2006.

Chave, A.D., J.W. Bailey, S. Beaulieu, R. Butler, F.K. Duennebier, J.H. Filloux, D. Harris, M. Mandea, J.A. Orcutt, K. Smith, R. Stephen, P. Tarits, F.L. Vernon, and F.B. Wooding (2003). 2003-2004 Upgrades and Additions to the Hawaii-2 Observatory. *Proc. 3rd Int.Workshop on Scientific Use of Submarine Cables and Related Technologies, IEEE, 14-19.*

Schulte-Pelkum, V., P. S. Earle, and F. L. Vernon (2004), Strong directivity of ocean-generated seismic noise, *Geochem. Geophys. Geosyst.*, **5**, Q03004, doi:10.1029/2003GC000520.

Sutherland, F. H., F. L. Vernon, J. A. Orcutt, J. A. Collins, and R. A. Stephen (2004). Results from OSNPE: Improved Teleseismic Earthquake Detection at the Seafloor. *Bull. Seismol. Soc. Amer.*, **94** (5), 1868-1878.

Five Other Significant Publications:

Prieto, G. A., F.L. Vernon, G. Masters, and D. J. Thomson (2005), Multitaper Wigner-Ville Spectrum for Detecting Dispersive Signals from Earthquake Records, *Proceedings of the Thirty-Ninth Asilomar*

Conference on Signals, Systems, and Computers, IEEE, 938-941.

Sohn, R. A., F. Vernon, J. A, Hildebrand, S. C. Webb (2000). Field measurements of sonic boom penetrations into the ocean. *J. Acoust. Soc. Am*., **107**, 3073-3083.

Stephen, R. A. , F. N. Spiess, J. A. Collins, J. A. Hildebrand, J. A. Orcutt, K. R. Peal, F. L. Vernon, and F. B. Wooding (2003). Ocean Seismic Network Pilot Experiment. *Geochem. Geophys. Geosyst.*, **4**, 1092, doi: 10.1029/2002GC000485.

Collins, J.A., F.L. Vernon, J.A. Orcutt, R.A. Stephen (2002). Upper Mantle Structure Beneath the Hawaiian Swell: Constraints from the Ocean Seismic Network Pilot Experiment. *Geophy. Res. Lett*. 29

Collins, J.A., F.L. Vernon, J.A. Orcutt, R.A. Stephen, K.R. Peal, F. B. Wooding, F.N. Spiess, and J.A. Hildebrand (2001). Broadband seismology in the oceans: lessons from the Ocean Seismic Network Pilot Experiment. *Geophy. Res. Lett*.., **28**, 49-52.

Doctoral Advisors:
James Brune, University of Nevada, Reno
Jonathan Berger, University of California, San Diego

Postdoctoral Advisor:
Jonathan Berger, University of California, San Diego

Post Graduate Scholar Sponsor
Robert Mellors San Diego State University
Catherine DeGroot-Hedlin, University of California, San Diego
Deborah Kilb, University of California, San Diego

Recent Collaborators:

John Collins	Terry Wallace	
Ralph Stevens	Gary Pavlis	Alan Levander

Synergistic Activities:

1. Presented Invited Seminar to National Insurance Conference on Seismic Hazards in Southern California
2. Created Internet Web sites for Hector Mines, Turkey, Taiwan, Mexico earthquakes in 1999 for public outreach. These sites are linked by Education and Outreach Programs such as the Incorporated Research Institutions for Seismology
3. Participated in the development of the Datascope relational database and associated application software which is used in a wide range of educational and research environments.
4. Scientific writings for general public interest such as the IRIS newsletter which reaches 98 IRIS institutions and many foreign affiliates
5. Presentations to US Congressional Members on Comprehensive Test Ban Treaty Issues.

UNIVERSITY OF WASHINGTON

SCHOOL OF OCEANOGRAPHY
Box 357940

19 December 2006

Dr. John A. Orcutt
Assoc. Vice Chancellor, Research Affairs
Director, UCSD Center for Earth Observations & Applications
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0225

Dear John:

This letter offers strong support for the Scripps Institution of Oceanography, UCSD proposal to become the Implementing Organization for the Cyberinfrastructure component within the ORION Program. As you know, the University of Washington responded to the RFP from the Joint Oceanographic Institutions, Inc. with a proposal to become the Implementing Organization for the Regional Cabled Observatory within ORION. Should the University of Washington be selected as the IO for that portion of the Program, we are strongly committed to working closely with the group selected as the IO for Cyberinfrastructure.

Having collaborated with you for many years--beginning in 1983, when we worked together on the formation of the U.S. Science Advisory Committee for the ocean drilling program, to today, when we are co-PIs on the LOOKING program within NSF's Information Technology Research program--I am well aware of your long-standing interest and expertise in the cyberinfrastructure aspects of the ocean sciences enterprise.

As we both know, the ocean-science community is in the midst of a transformation in the ways we study the ocean basins. Your leadership of successful programs such as ROADNet, which delivers seismic, oceanographic, hydrological, ecological, and physical data to a variety of end users in real time, is a clear indication of the approaches that must be adopted for this transformation to be successful and productive. Your work on the Southern California Coastal Ocean Observing System (SCCOOS) is at the forefront of such activities as it synthesizes observations into products that will provide a scientific basis for evaluating and improving management and guardianship of, and response to the ocean environment and its resources. Clearly, the support that the state of California provides coastal sciences will be an important leveraging capability for the activities planned within ORION.

Your initiation of the Dynamics of Earth and Ocean Systems (DEOS) effort in 1996 -7 was one of the early and important elements contributing to the current successes of ORION. One of the key results of DEOS was support to provide community input for efforts within OCE-NSF to successfully put forward the Ocean Observatories Initiative (OOI) as an MREFC Program within NSF.

The University of Washington is hopeful that our response to the ORION RFP for the Implementing Organization of the Regional Cabled Ocean Observatory will be successful. If we are granted the award we will work closely with the winner of the CI IO competition.

Our interest is in defining and developing the appropriate interfaces with the ORION CI group to ensure the successful completion of the construction phase and a smooth transition into an operational phase of the RCO component of ORION. The University of Washington has considerable assets in terms of connectivity and networking expertise, with high-speed communication links associated with the National Lambda Rail, Internet 2, and beyond. We look forward to bringing these resources to the ORION effort.

Many of our colleagues recognize the vision of the ORION CI in providing all users with a system that enables simple and direct use of ORION resources to accomplish their scientific objectives. This vision includes direct access to instrument data, control, and operational activities, and the opportunity to seamlessly collaborate with other scientists, institutions, projects, and disciplines. There are important dialogues to be had among the all the ORION IOs and NSF to establish the optimum management and use of a ocean observatory network within the U.S., and with our international partners such as NEPTUNE Canada at the University of Victoria.

Scripps has a long history of commitment to education and public awareness programs related to the ocean sciences. Technology, particularly cyberinfrastructure, now offers unparalleled opportunities for bringing the excitement of our scientific explorations and discoveries in real time to students, educations, and the public. Just as we so successfully collaborated on the VISIONS '05 effort to bring first-ever live feeds of high definition video from the seafloor to distribution sites around the U.S., including iGrid05 in La Jolla, I feel we can continue to work together on ORION efforts if each of our institutions is awarded an IO.

We wish you success in your quest to become the CI-IO.

Sincerely yours,

John R. Delaney
Director, NEPTUNE Program
Professor of Oceanography, and
Jerome M. Paros Endowed Chair in Sensor Networks
University of Washington

<MMI />

Mailstop 46
Woods Hole Oceanographic Institution
Woods Hole, MA. 02543

December 6, 2006

Professor John Orcutt
Institute of Geophysics and Planetary Physics
Scripps Institution of Oceanography
University of California, San Diego
9500 Gilman Drive # 0225
La Jolla, CA 92093-0225

John,

The Marine Metadata Initiative (MMI) Executive Committee met on Friday, December 1 2006 to discuss MMI's possible role in the UCSD-led ORION CI proposal to JOI and NSF. This was in response to an email from Matthew Arrott. The committee understands that if the proposal effort is successful MMI's role will be to provide guidance to a software engineer employed by ORION CI and that MMI will provide this guidance to this individual strictly on a pro bono basis.

The goals outlined in your correspondence to John Graybeal are well aligned with MMI current and future work and goals. Therefore, the Executive Committee endorses MMI's participation in the manner suggested. However, since it would be inappropriate for us to explicitly endorse one ORION CI proposal over another, this endorsement of our participation in your proposal should not be construed by anyone that we endorse your proposal. We note here that we have not seen any part of your proposal.

Details, based on the description of the agreement in Matthew Arrott's previous correspondence, are shown below:

We approve a joint venture between the UCSD-led ORION CI proposal team and the Marine Metadata Interoperability (MMI) project as a part of the your team's proposal to JOI and NSF.

The basic premise of the joint venture is that the partners will contribute time and know how towards the development of a set of shared objectives. The proposed contributions are:

- The UCSD-led ORION CI team will provide the core development resources and a community/platform of opportunity on which to realize mutually beneficial outcomes.

- MMI project member(s) will provide community and design leadership as well as participate in development on a best effort basis.

The UCSD-led ORION CI team will contribute at minimum 4 months of development effort each year for the duration of the OOI project (6 years) towards the advancement of the following outcomes. The work will be performed by a software engineer employed by the ORION CI Project under the direction of MMI.

1. A set of semantic domain models (informational and procedural schema) that will form the basis of the ORION CI semantic framework, in particular, models for vocabulary, thesaurus, and ontology with their respective model for mapping between instances of the a model.

2. A tool chain for the development and negotiation of vocabularies, thesauri and ontologies to be employed with in the ORION Ocean Observatory by individuals, ad hoc groups of participants and the ORION community as a whole.

3. A representative set of the major vocabularies and ontologies in use within the oceanographic research community. These are to be provided with a thesaurus that can be localized and modified.

4. A methodology based on “best practices” for disseminating the use of vocabularies and ontologies along with the associated tool chain in the pursuit of scientific investigation within the ORION Ocean Observatory system.

If there are any questions related to the contents of this letter please reply directly to Andrew Maffei who is acting as point-of-contact for MMI. We also wish to inform you that various members of the MMI project are in “conflict of interest” in this matter. Restricting your correspondence to my direct attention is intended to help to protect any sensitive correspondence.

Sincerely Yours,

Andrew M Maffei

Andrew Maffei
amaffei@whoi.edu
Senior Information Systems Specialist
Woods Hole Oceanographic Institution