

Multi-Phase Plan for the IS&T Data Exploratory

Katherine Nystrom INST-OFF Project Mentors: Gary Grider HPC-DO, Carolyn Connor HPC-5, John Bent HPC-5

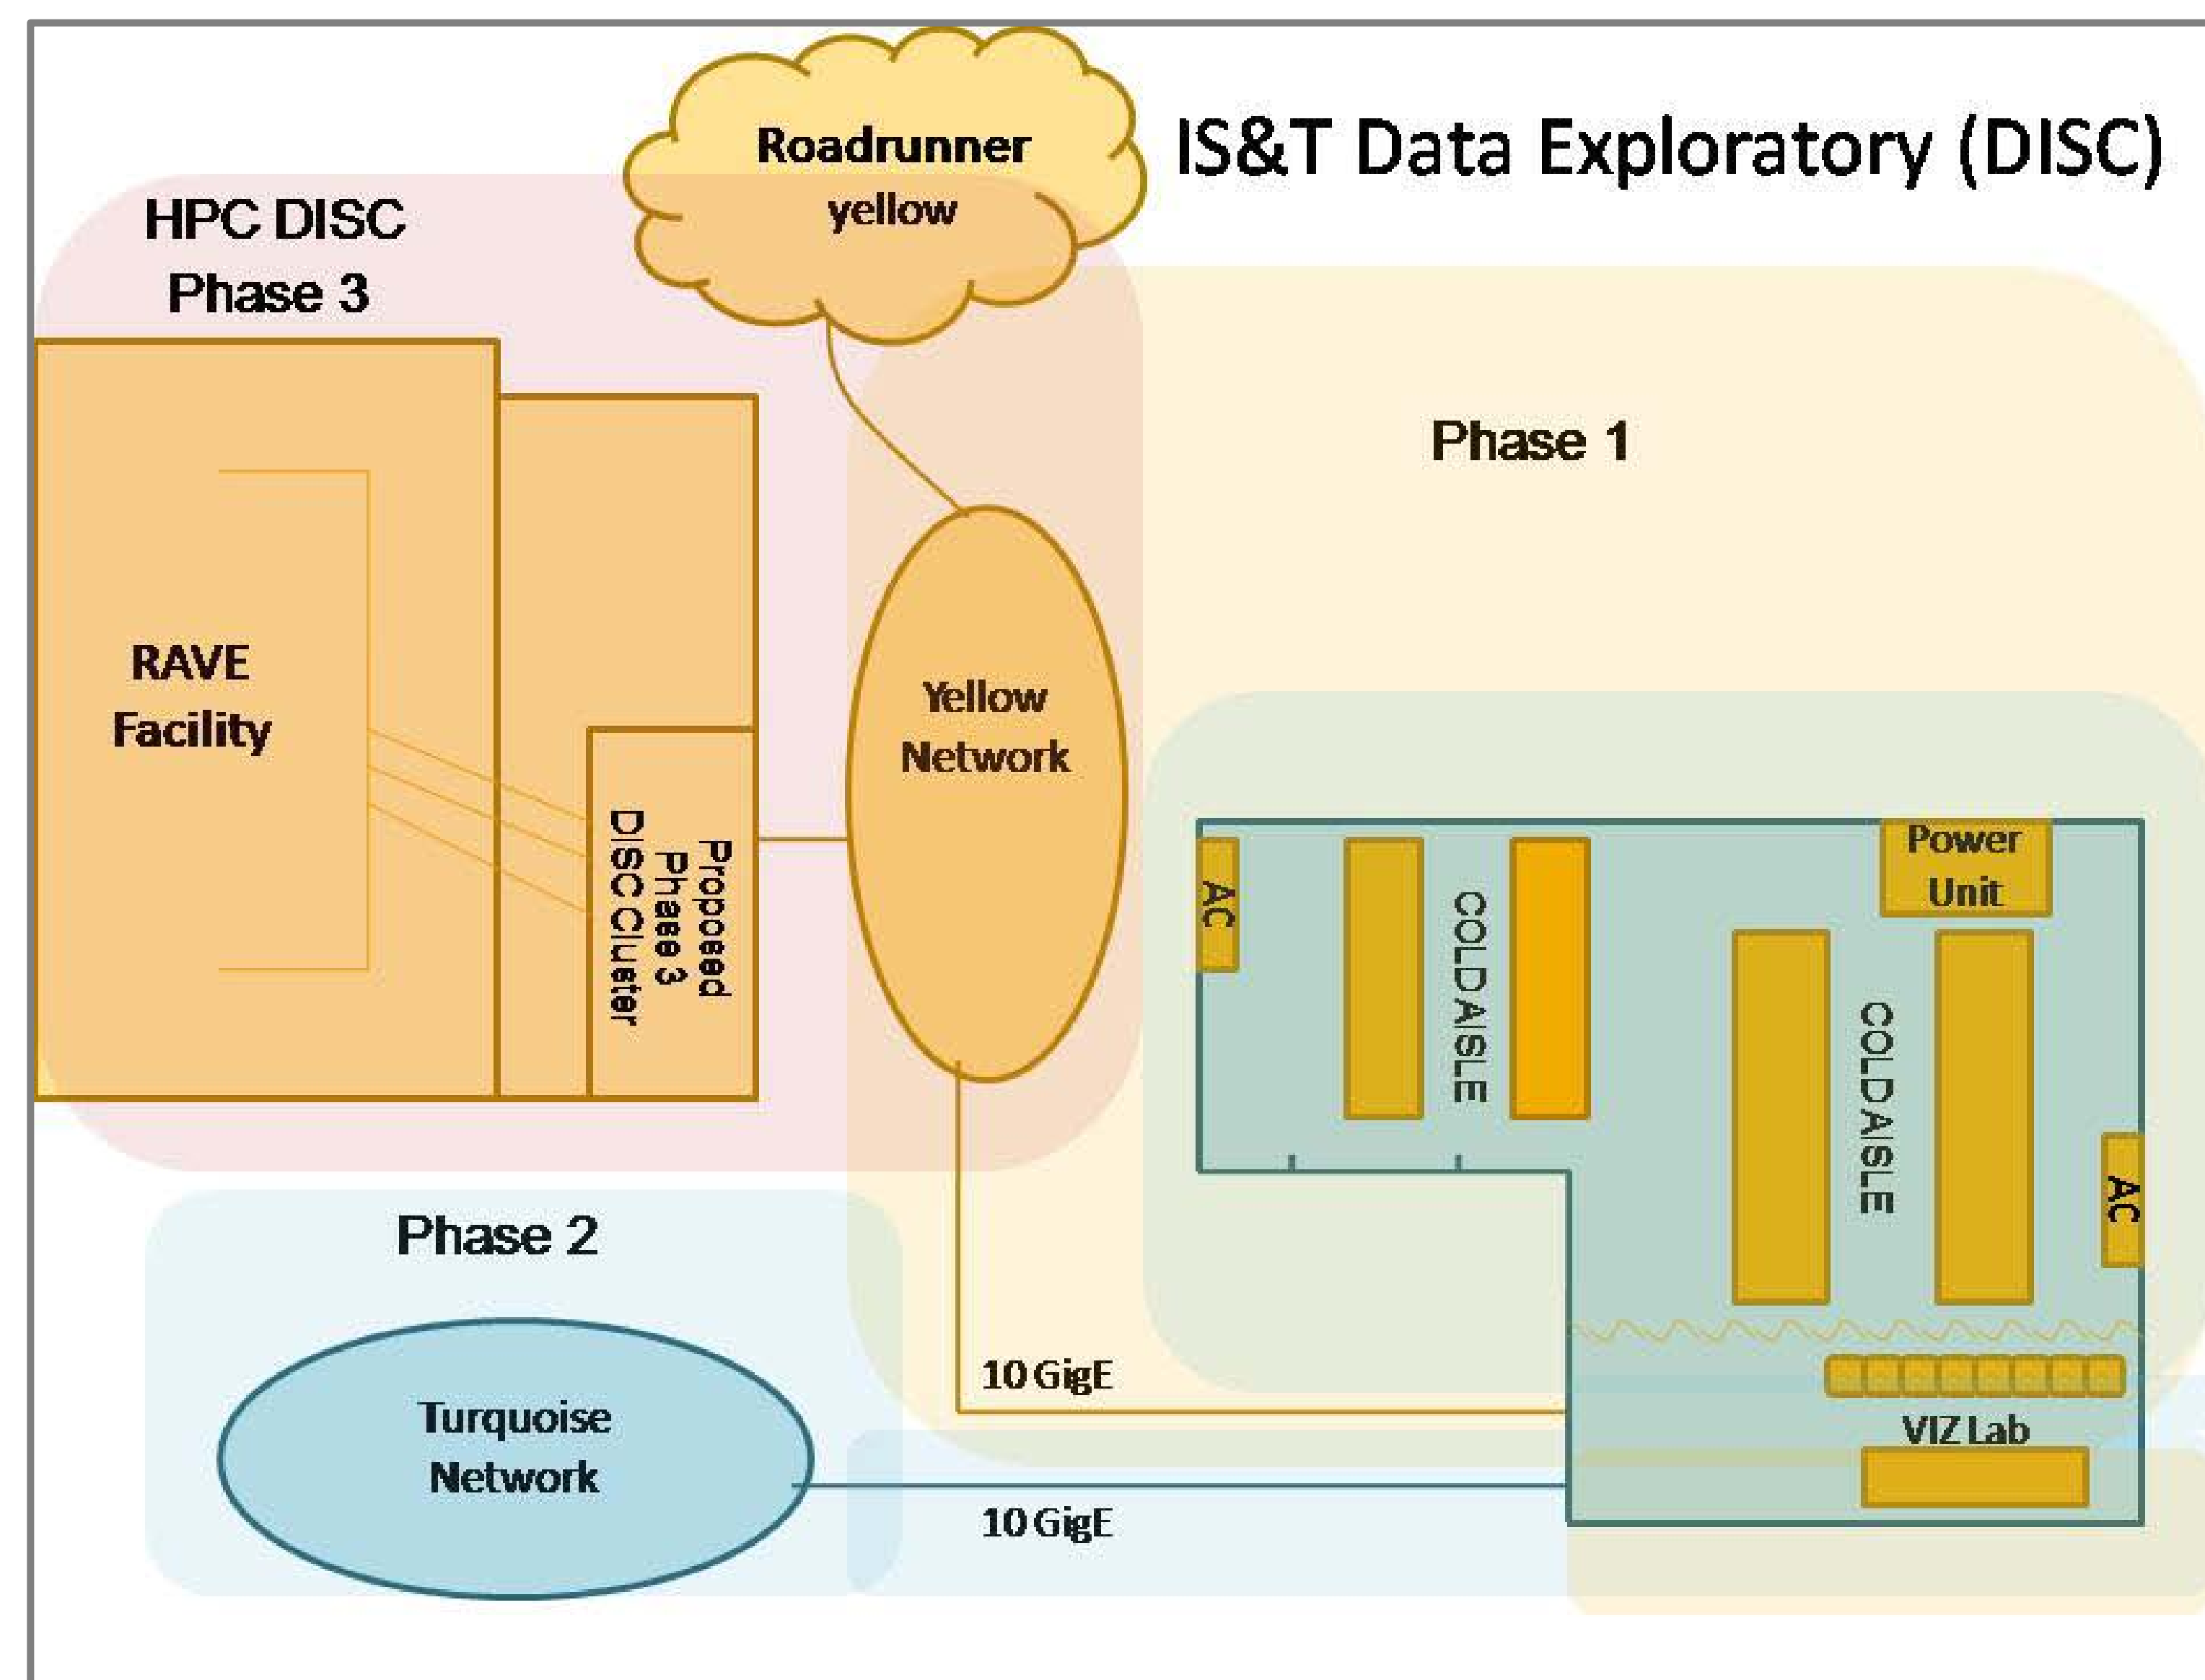
The IS&T Data Exploratory is a unique, data intensive supercomputing facility designed to explore the challenges of working with large datasets. Sensors, Internet packet filters, telescopes and satellites all generate massive amounts of data. To solve problems in areas like energy security, bio-security, and cosmology we need effective ways to analyze this data. The Data Exploratory will provide a trial facility in an open environment where researchers can experiment with data intensive methods.

The exploratory will have two 10GigE connections to bring in data: one to the yellow network and one to the turquoise network. The data will be stored on a distributed file system of hard drives local to compute nodes. The focus of the facility will be to provide an infrastructure to effectively support data intensive applications.

One of the challenges of large datasets is representing the data in a manageable format that scientists can learn from. The IS&T Data Exploratory will include a visualization environment to provide an alternative method for looking at results from data intensive applications. The goal is to find new ways of using visualization for information applications, to provide an abstraction of the data rather than a simulation of it.



- Hadoop is an open source distributed computing system for working with tera- and petabyte size datasets.
- Hadoop Distributed File System (HDFS) stores data on hard drives in the nodes and is designed to be highly fault-tolerant.
- Map/Reduce is a software framework for writing applications to run in parallel on Hadoop clusters.
- Plan is to experiment using Hadoop and Map/Reduce to go beyond text-based mapping applications.



The IS&T Data Exploratory is a data intensive supercomputing facility located on the first floor of the Los Alamos Research Park. The focus of the facility is to provide an infrastructure to effectively support data intensive applications.

Phase 1 Fast Facts

- ~100-200 node mixed architecture cluster
- Perceus for cluster management and node provisioning
- Fedora 9 operating system
- Hadoop software framework for providing data intensive support
- Initial user applications include cosmology, cyber security and image processing

Data Intensive vs. Compute Intensive

Compute Intensive – maximizes system resources for processing large computations for simulation

Data Intensive – simplifies the challenges of working with large datasets generated by sensors and simulation results. Goal is to move computation to the data.

Phase 1: Initially the cluster will be connected to the yellow network via 10gigE. This will enable researchers to access the data generated from open science projects on Roadrunner.

Phase 2: A second 10gigE connection will be added to the turquoise network. This is the network available for guest scientists and visiting universities. The Roadrunner archives will be moved to the turquoise network after Roadrunner moves to the red network.

Phase 3 (proposed): A similar facility will be developed in Bldg. 132 after several older clusters are decommissioned. This facility will be near the Rave Visualization Center and will use the Rave for a visualization environment. This cluster will consist of ~2000 nodes.



Current configuration of the DISC server room with seven racks in place.

Goals

- Provide an open facility for researchers to experiment with data intensive methods
- Develop new capabilities to address problems involving large datasets
- Use data intensive methods to assist with processing large datasets generated by compute intensive applications
- Discover new ways to use visualization tools such as fusing data from simulation and sensors
- Expand LANL's expertise in data fusion and data intensive computing for science