DYNES: Deployment Plan
DRAFT

1 Introduction

This NSF-funded project (grant number 0958998) is developing and deploying the **Dynamic Network System (DYNES)**, a nationwide cyber-instrument spanning multiple US universities and Internet2 connectors. A collaborative team including **Internet2**, **Caltech**, **University of Michigan**, and **Vanderbilt University** are leading this effort in cooperation with regional networks and campuses. The objective is to support large, long-distance scientific data flows in the LHC, other programs in data intensive science (such as LIGO, Virtual Observatory, and other large scale sky surveys), as well as the broader scientific community.

For the latest announcements concerning DYNES, subscribe to the dynes@internet2.edu mailing list.

- http://lists.internet2.edu/sympa/subscribe/dynes.

To ask questions of the DYNES project team, email dynes-questions@internet2.edu.

- http://lists.internet2.edu/sympa/subscribe/dynes-questions

The purpose of this document is to provide an overview of the DYNES deployment plan and schedule. Additional documents are available which provide a DYNES project overview, technical details, and other information at the project web site:

- http://www.internet2.edu/dynes

2 DYNES Infrastructure Overview

This section is intended to provide a high level overview of the DYNES Infrastructure as required to support DYNES program and schedule information presented in the remainder of this document. Additional details on the DYNES system, technologies, and components can be found in the other DYNES documents as referenced on the DYNES project web site.

The DYNES Infrastructure will be composed of Internet2's ION service and extensions over regional and state networks to US campuses. It will connect with transoceanic (IRNC PRONET, USLHCNet), European (GÉANT), Asian (SINET3) and Latin American (RNP and ANSP) R&E networks. It will build on existing key open source software components that have already been deployed and tested including the DCN Software Suite (OSCARS/DRAGON), perfSONAR, the UltraLight Linux kernel (Michigan), and Caltech's Fast Data Transfer (FDT) open source application suite: FDT, FDT/dCache and FDT/Hadoop.

A high-level view of the generic DYNES infrastructure is presented in Figure 1. This topology view will be updated to reflect the specific list of DYNES participating regional networks and campus/sites as they are identified. The DYNES system is based on the combination of the DCN and FDT technologies into a integrated workflow tailored to support science application specific activities. Figure 2 depicts a high level view of the DYNES Data Flow architecture. The following information provides additional detail related to Figure 2.

- Each DYNES Sites will be assigned DYNES Project private address space (10.20/16)
- Each DYNES FDT Server will be assigned a DYNES EndPoint Name (siteZ.fdt1)
- The DYNES FDT Server will include a data storage and reference structure to allow user to identify and indicate the data to be moved via DYNES. This data storage and reference structure

will be project dependent. This will allow users to specify the desired data in the form of a DYNES Data Id.

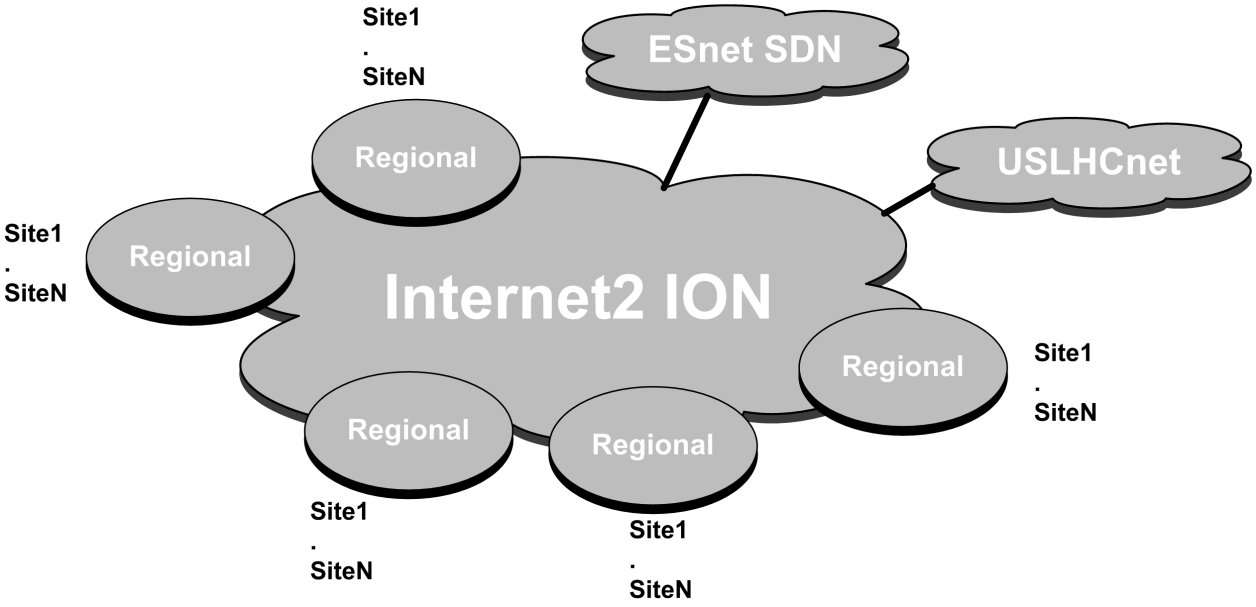

Figure 1 DYNES Infrastructure Overview

- The combination of the DYNES EndPoint Name and DYNES Data ID will form a "DYNES Transfer URL" (siteZ.fdt.1/datalocationref30)
- Users will need to present a "DYNES Transfer URL" to their local DYNES Agent to initiate the data transfer.
- The DYNES Agent (DA) will provide the functionality to request the circuit instantiation, initiate and manage the data transfer, and terminate the dynamically provisioned resources. Specifically the DA will do the following:
    - Accept user request in the form of a DYNES Transfer URLs indicating the data location and ID
    - Locates the remote side DYNES EndPoint Name embedded in the Transfer URL
    - Submits a dynamic circuit request to its home InterDomain Controller (IDC) utilizing its local DYNES EndPoint Name as source and DYNES EndPoint Name from Transfer URL as the destination
    - Wait for confirmation that dynamic circuit has been established
    - Starts and manages Data Transfer using the appropriate DYNES Project IP addresses
    - Initiate release of dynamic circuit upon completion
- The dynamic circuit network infrastructures and control plane will provide for the multi-domain circuit instantiation. The high level workflow is as described below:
    - Upon receipt of a circuit request from a DYNES Agent, the IDC utilized the DYNES LookUp Service to translate the DYNES EndPoint Names into dynamic circuit source and destination URNs.

- The initiating IDC then uses these URNs to set up the multi-domain dynamic circuit and

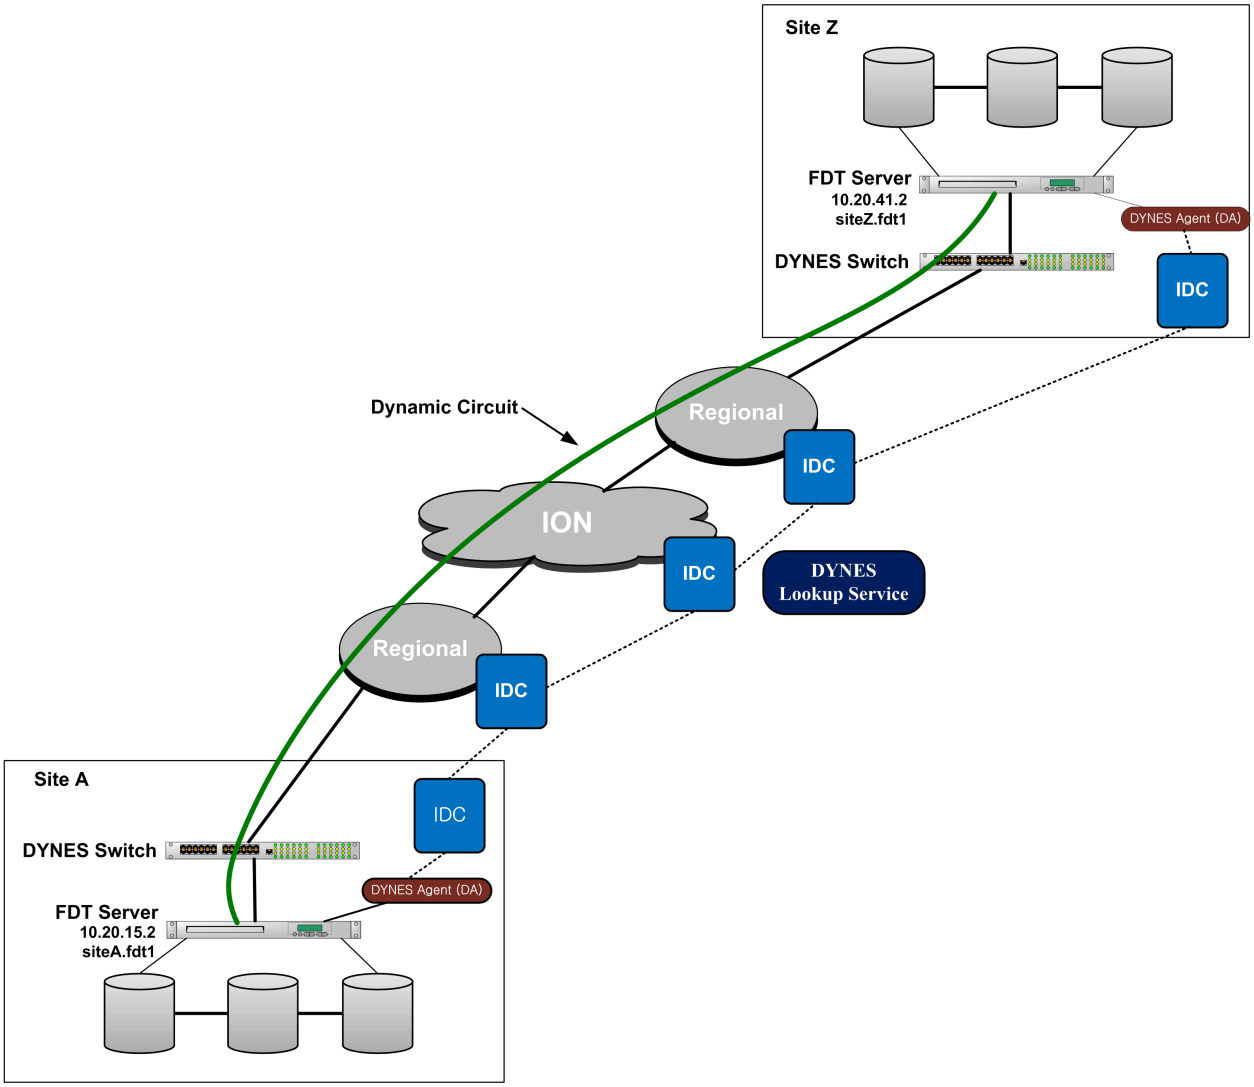

Figure 2 DYNES Data Flow View

notify the DYNES Agent when circuit is ready for use.

- The IDC will also accept requests from the local DYNES Agent to tear down the dynamic circuit after data transfer is complete.

- The FDT Servers can also have public IP addresses which can also be utilized for data transfers when dedicated circuits are not instantiated.
- In this basic scenario, only the provided FDT server will be integrated into the DYNES instrument. However, other site servers can also be integrated into the DYNES infrastructure.

The standard set of DYNES provided equipment will include the following:

- Inter-domain Controller (IDC)

- Ethernet switch
- Fast Data Transfer (FDT) server with an attached disk array

Some sites will want to receive a set of equipment different from they DYNES standard configuration. This will be accommodated to the maximum extent possible This topic is discussed in more detail in Section 4 of this document. The remainder of this document presents the program deployment plan and schedule for the DYNES Instrument.

3 DYNES Deployment Plan and Schedule

The DYNES deployment plan and schedule as proposed to the National Science Foundation (NSF) identified four project phases which are summarized below.

- Phase 1: Site Selection and Planning (4 months) (Sep-Dec 2010)
- Phase 2: Initial Development and Deployment (6 months) (Jan 1-Jun 30, 2011)
    - development of DYNES at a limited number of sites
    - sites will be a minimum of Caltech, University of Michigan and Vanderbilt and their regional networks, plus 2-3 additional campuses (such as the University of Nebraska) and at least one of their respective regional networks.
- Phase 3: Scale Up to Full-scale System Development (14 months) (July 1, 2011-August 31, 2012)
    - full-scale deployment and development at all selected sites
- Phase 4: Full-Scale Integration At-Scale; Transition to Routine O&M (12 months) (September 1, 2012-August 31, 2013)
    - DYNES will be operated, tested, integrated and optimized at scale, transitioning to routine operations and maintenance as soon as this phase is completed

As part of Phase 1 work, the DYNES Project has developed additional details for the deployment plan and schedule as presented below.

<u>Phase 1: Site Selection and Planning (Sep-Dec 2010)</u>
i) Participant Applications Due: December 15, 2010
ii) Application Reviews: December 15 2010-January 31 2011
iii) Participant Selection Announcement: February 1, 2011

<u>Phase 2: Initial Development and Deployment (Jan 1-Jun 30, 2011)</u>

i) Initial Site Deployment Complete - February 28, 2011

- Caltech, Vanderbilt, University of Michigan, Internet2, USLHCnet

ii) Initial Site Systems Testing and Evaluation Complete: April 29, 2011

<u>Phase 3: Scale Up to Full-scale System Development (14 months) (July 1, 2011-August 31, 2012)</u>
i) Phase 3-Group A Deployment (10 Sites) (March 28 -July 1, 2011) (~13 weeks)

- Teleconferences and Planning with individual participants: March 28-May 2 2011(~5 weeks)
- Finalize Phase 3-Group A Equipment Order List: May 2-9, 2011(~1 weeks)
- Place Equipment Order: May 10, 2011
- Receive DYNES Equipment: May 24, 2011
- Configure and Test Individual Participant Configurations: May 24 - June 7 (~2 weeks)
- Ship Phase 3-Group A Equipment to sites: June 14 2011

- Deploy and Test at Phase 3-Group A Sites: June 21 - June 28, 2011
- Begin Phase 3-Group A: July 1, 2011

ii) Phase 3-Group B Deployment (15 Sites): July 18-August 26, 2011(~7 weeks)

- Telecons and Planning with Individual Participants
- Finalize Phase 3-Group B Equipment Order List
- Place Equipment Order
- Receive DYNES Equipment
- Configure and Test Individual Participant Configurations
- Ship Phase 3-Group B Equipment to sites
- Deploy and Test at Phase 3-Group B Sites
- Begin Phase 3-Group B: September 1, 2011

iii) Phase 3-Group C Deployment (15 Sites): September 5-October 14, 2011(~6 weeks)

- Telecons and Planning with Individual Participants
- Finalize Phase 3-Group C Equipment Order List
- Place Equipment Order
- Receive DYNES Equipment
- Configure and Test Individual Participant Configurations
- Ship Phase 3-Group C Equipment to sites
- Deploy and Test at Phase 3-Group C Sites
- Begin Phase 3-Group C: October 17, 2011

iv) Full-scale System Development, Testing, and Evaluation (October 17 2011- August 31, 2012)

Phase 4: Full-Scale Integration At-Scale; Transition to Routine O&M (12 months) (September 1, 2012-August 31, 2013)

-DYNES will be operated, tested, integrated and optimized at scale, transitioning to routine operations and maintenance as soon as this phase is completed

As indicated in the above plan, Phase 3 has been divided into three Groups. Each of the groups are expected to contain approximately ten to fifteen participants which will be a mixture of sites and regional networks. The reason for this phased approach to Phase 3 is to find a good balance between deployment, testing, and incorporating lessons learned as we scale up the DYNES instrument. The exact participant groupings will be decided as part of consultation with the participants and project management. It should also be noted that Phase 3-Group A includes a more ambitious schedule then the original proposal plan. This will allow for some buffer in case unexpected issues are uncovered as part of the initial deployment and testing.

4 DYNES Provided Equipment - Installation, Configuration, and Deployment

The standard set of DYNES provided equipment will be as described below.

Inter-domain Controller (IDC) Server and Software

The IDC creates virtual LANs (VLANs) dynamically between the FDT server, local campus, and wide area network. The IDC software is based on the OSCARS [1] and DRAGON [2] software which is packaged together as the DCN Software Suite (DCNSS) [3]. The DCNSS version correlates to stable tested versions of OSCARS. The current version of DCNSS is v0.5.3. It expected that DCNSSv0.6 will be utilized for Phase 3-Group B deployments and beyond. DCNSSv0.6 will be fully backward compatible with v0.5.3. This will allow us to have a mixed

environment as may result depending on actual deployment schedules. The IDC server will be a Dell R610 1U machine.

Ethernet switch

The standard DYNES Ethernet switch options are:

- Dell PC6248 (48 1GE ports, 4 10GE capable ports (Select SFP+, CX4 or optical)
- Dell PC8024F (24 10GE SFP+ ports, 4 "combo" ports supporting CX4 or optical)

Fast Data Transfer (FDT) server, Attached Disk Array, and DYNES Agent Software

The Fast Data Transfer (FDT) [4] server connects to the disk array via the SAS controller and runs the FDT software developed by Caltech. FDT is based on an asynchronous, flexible multithreaded system that automatically adjusts I/O and network buffers so that the network achieves maximum utilization. The disk array stores the datasets to be transferred among the sites in some cases. The FDT server serves as an aggregator/throughput optimizer in this case, feeding smooth flows over the WAN and LAN directly to the Tier2 or Tier3 clusters. The standard DYNES FDT Server will be a DELL 510.

The standards FDT server will be a DELL 510 server with a dual port Intel X520 DA. This server will include a PCIe Gen2.0 card x8 card along with 12 disks for storage.

The DYNES project will provide this equipment along with detailed installation and configuration steps for individual site deployments. It is recognized that some sites will want to receive a set of equipment different from the DYNES standard configuration. This will be accommodated to the maximum extent possible. The deployment plan includes planning sessions with individual sites in advance of equipment purchase and system configuration. These planning sessions will be utilized to document in detail the equipment list and the associated installing documentation will be modified to reflect each sites configuration. It should be noted that the standard DYNES configuration will be pre-configured and tested by the DYNES project team prior to shipping to the sites. Configurations which are different from the DYNES standard configuration will likely require more support from local site personnel for configuration, testing, and maintenance.

5 DYNES Operational Model - Monitoring and Troubleshooting

The DYNES equipment will be monitored by a centralized DYNES monitoring system. This will require SNMP access to the site located DYNES equipment. This system will not require root access to DYNES equipment at the sites. As our experience using the DYNES infrastructure increases, we expect that the exact methods for identifying and correcting for failures will evolve.

6 References

[1] OSCARS, www.es.net/oscars

[2] DRAGON, dragon.east.isi.edu

[3] DCN Software Suite (DCNSS), wiki.internet2.edu/confluence/display/DCNSS/

[4] FDT System, monalisa.cern.ch/FDT/disk2disk.html