U.S. DEPARTMENT OF ENERGY | Office of Science

Future Lab Computing Working Group (FLC-WG)

Distributed Computing and Data Ecosystem (DCDE)

MAGIC @ SC 19

Ray Bair (ANL), Kuldeep Chawla (LBNL), Luis Corrales (LBNL), David Cowley (PNNL), Zhihua Dong (BNL), Eric Lancon (BNL), Ketan Maheshwari (ORNL), Daniel Murphy-Olson (ANL), Mallikarjun (Arjun) Shankar (ORNL)

ASCR PM: Richard Carlson

Context and Overview

- **DOE/SC Laboratories provide computing/storage resources to lab staff, researchers, and visiting scientists**
- **Demands on these resources are increasing**
- **Labs have the capability to leverage decades of research to create modern Distributed Computing and Data Ecosystems (DCDE) to meet the current and future demands of DOE scientists**
- **ASCR constituted Future Laboratory Computing Working Group (FLC-WG). Met through 2018 and delivered report with findings.**
- **DCDE pilot established for FY2019 fleshes out the key components and documents procedures to establish the infrastructure.**

FLC Working group report (2018): *Background and Roadmap for a Distributed Computing and Data Ecosystem*, https://doi.org/10.2172/1528707

DCDE Components

1. Seamless user access

2. Coordinated Resource allocations and cross-facility workflows

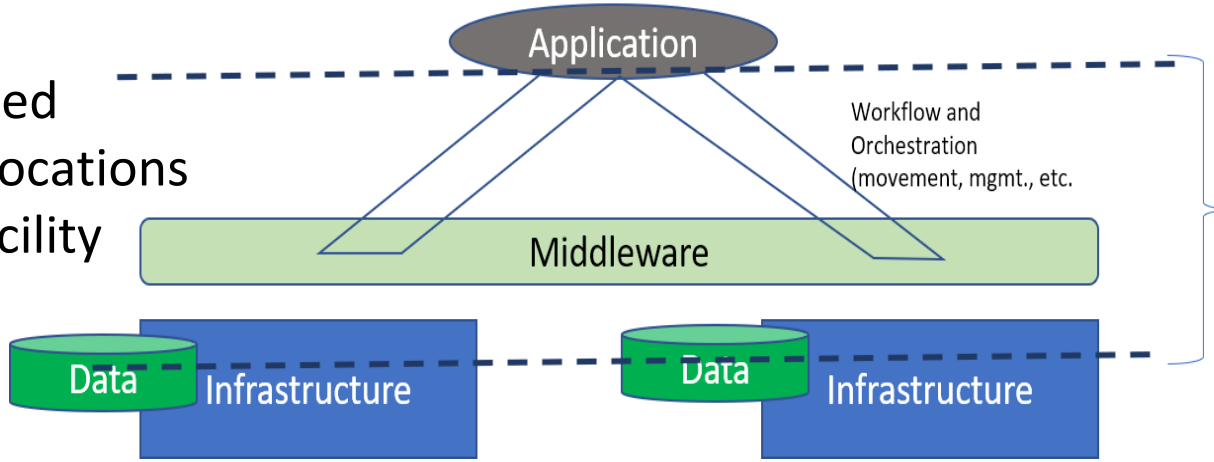

3. Data storage, movement, and dissemination for distributed operations

4. Variety, Portability: through virtualization, and containers, etc.

5. Governance and Policy Structures

Demonstration Components

- Science Driver
- Federated Identity Management
- Portability across laboratories
- Workflow through analytic notebooks
- Data Transfer

Try very hard to not reinvent anything: use available technologies and capabilities!

Distributed Computing and Data Ecosystem (DCDE) Demo Overview

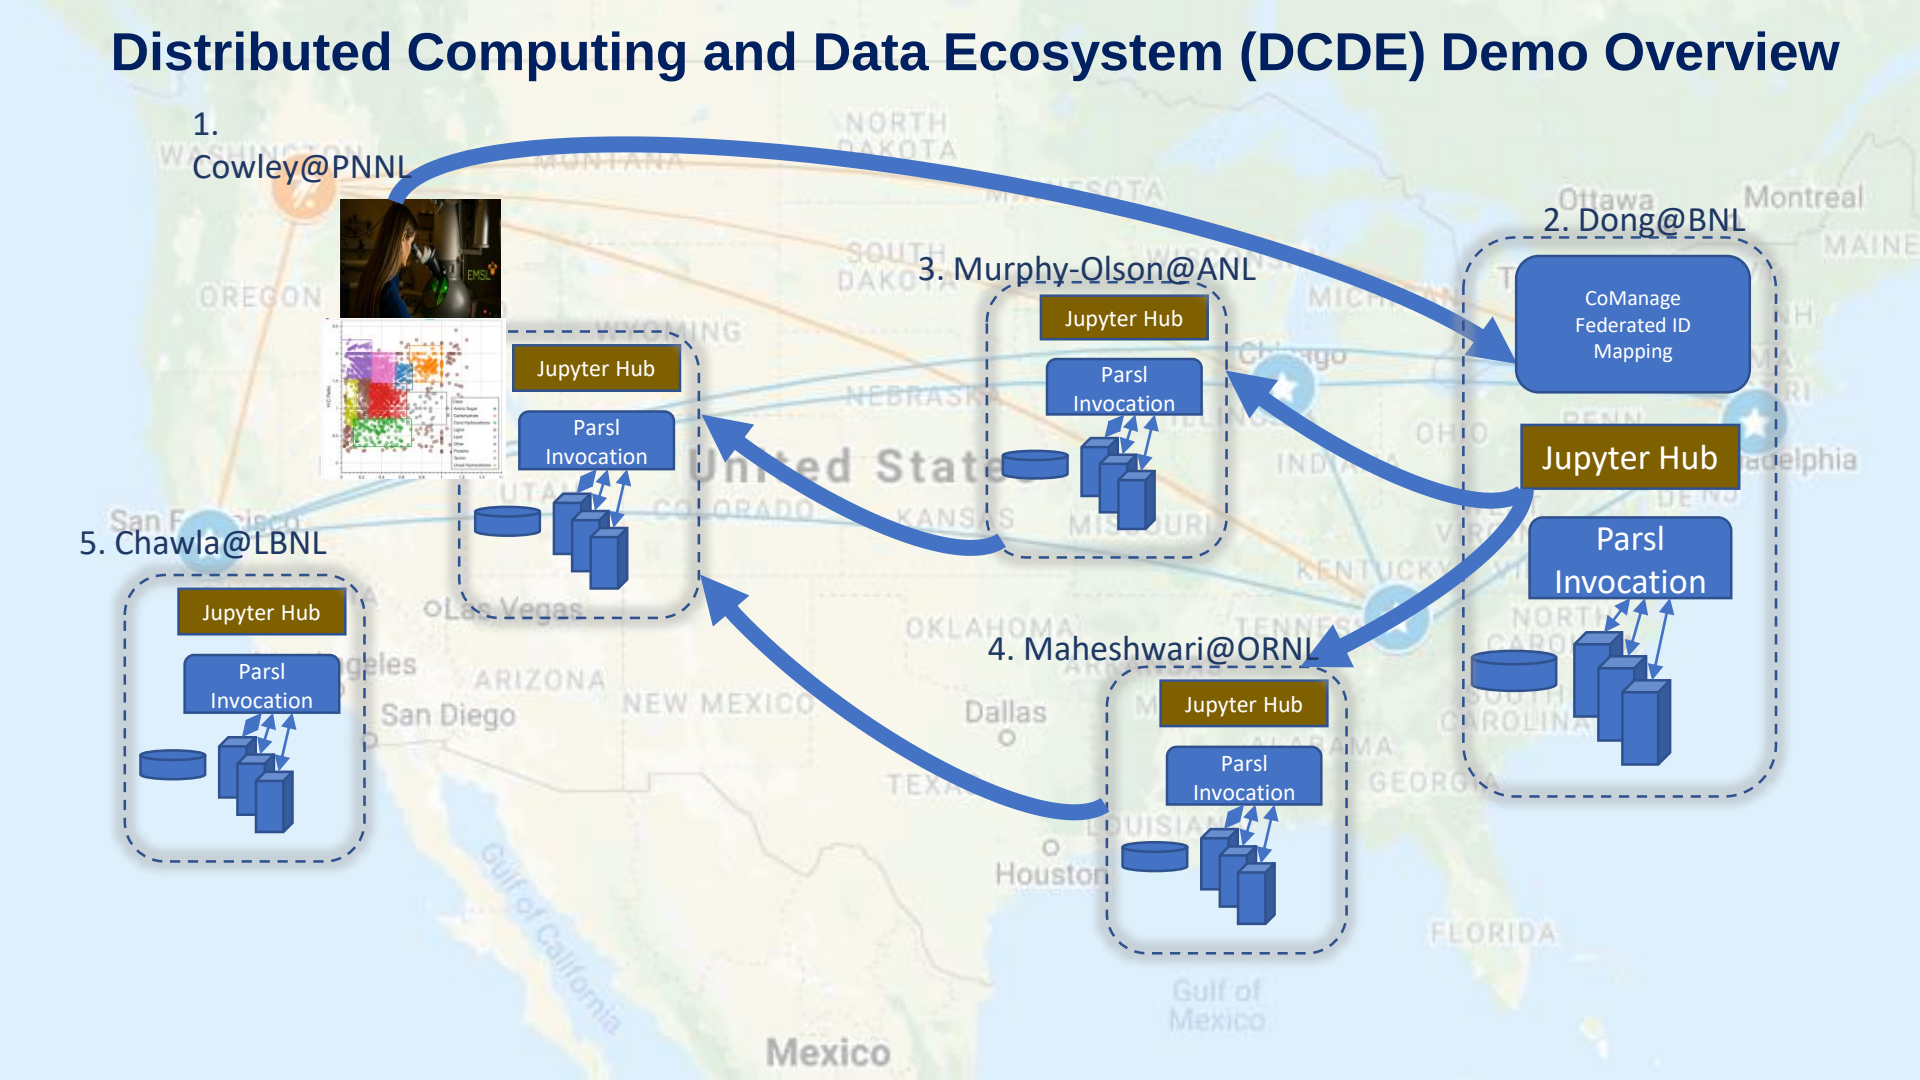

Guidelines for the pilot

- Simplicity (single researcher, small group)
- Lightweight, off the shell components
- Easy local installation
- Works locally and remotely, command line and Web interface
- Portability
- Modern (sustainable) technology

Demonstration Components

- Science Driver
  - DOE-BER EMSL code Relion
- Federated Identity Management
  - CoManage/CILogon
- Portability across laboratories
  - Singularity Containers
- Workflow through analytic notebooks
  - Jupyter Hub and Parsl
- Data Transfer
  - Globus

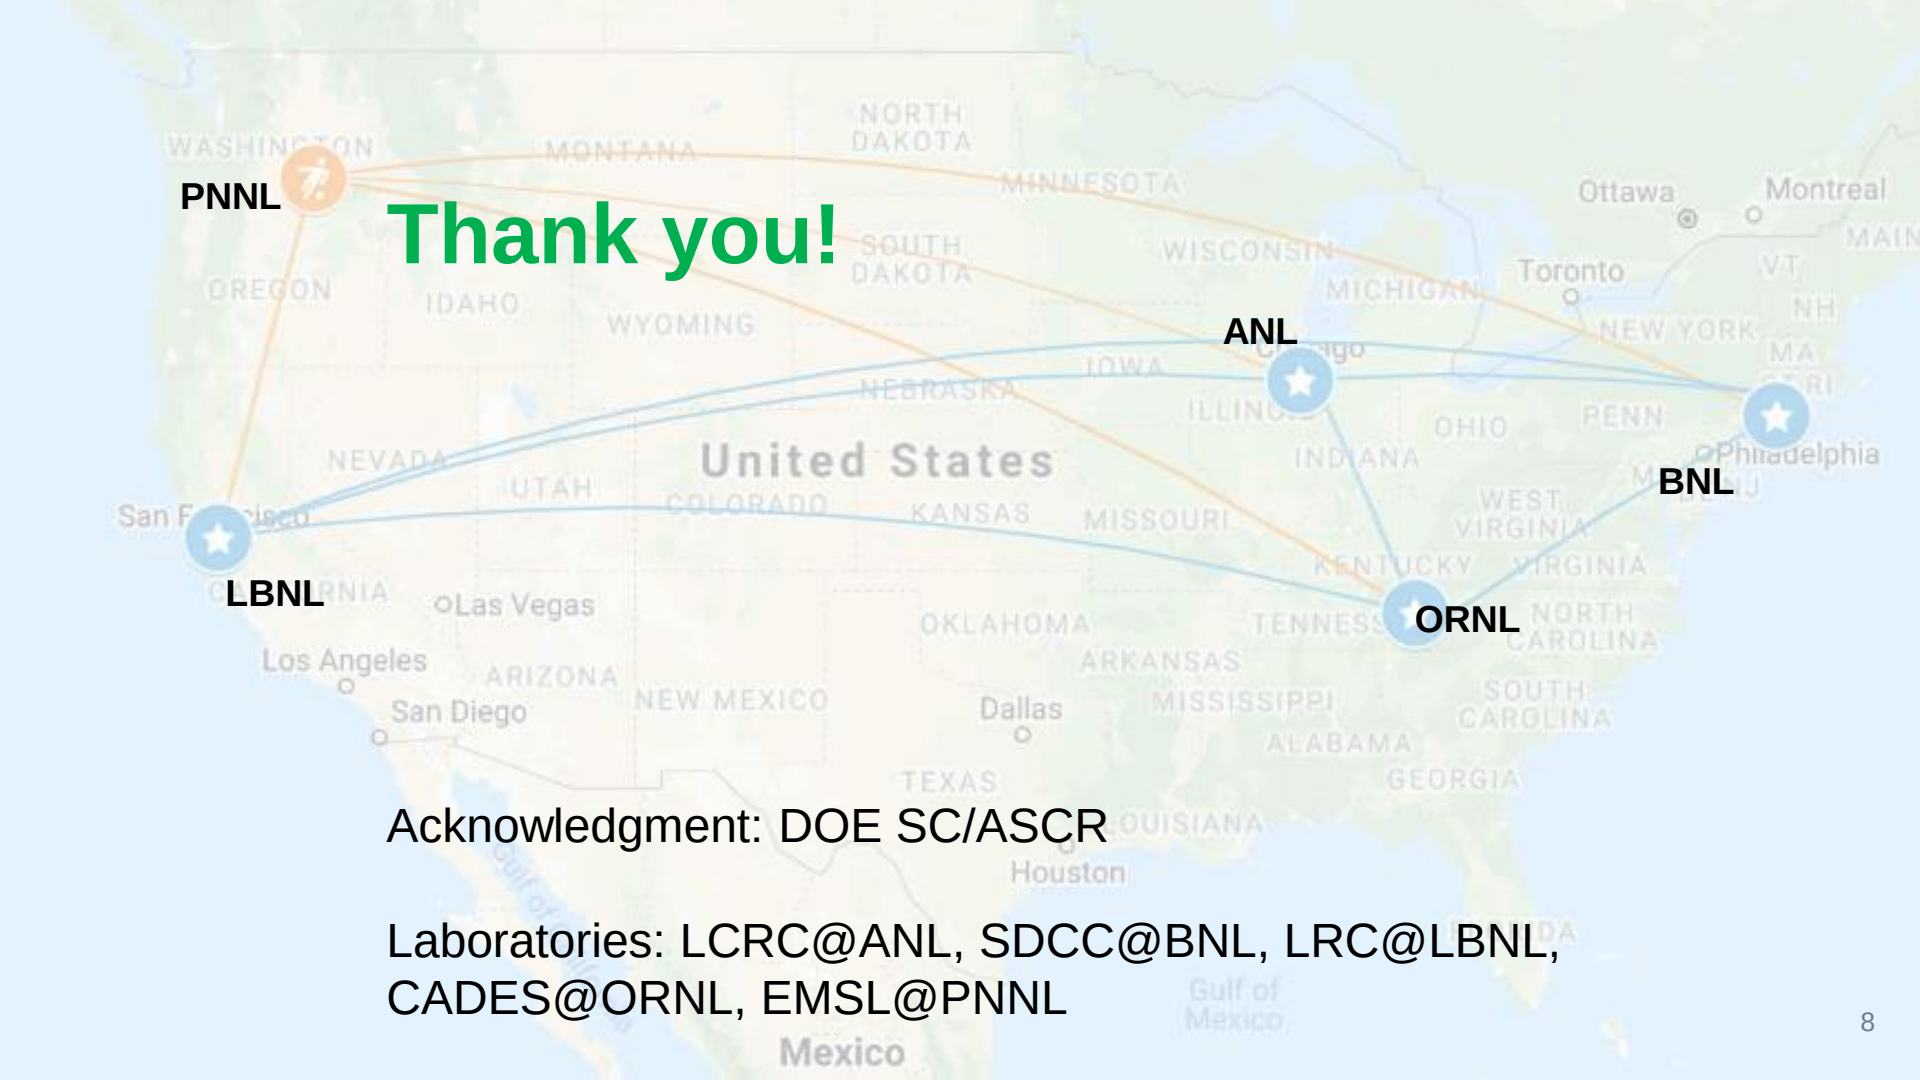

Thank you!

Acknowledgment: DOE SC/ASCR

Laboratories: LCRC@ANL, SDCC@BNL, LRC@LBNL, CADES@ORNL, EMSL@PNNL

U.S. DEPARTMENT OF ENERGY | Office of Science

Future Lab Computing Working Group (FLC-WG)

Distributed Computing and Data Ecosystem (DCDE) Technical elements

MAGIC @ SC 19

Ray Bair (ANL), Kuldeep Chawla (LBNL), Luis Corrales (LBNL), David Cowley (PNNL), Zhihua Dong (BNL), Eric Lancon (BNL), Ketan Maheshwari (ORNL), Daniel Murphy-Olson (ANL), Mallikarjun Shankar (ORNL)

ASCR PM: Richard Carlson

Table of Contents

- Overview
- AuthN/AuthZ: InCommon, CILogon and COManage
- Globus and auth-ssh
- Application and Containers
- Jupyter and Parsl
- Issues, Challenges and Lessons
- Summary and References

Overview

- Distributed Compute and Data Ecosystem (DCDE) project aims to unify and make available to its users the compute and data resources from DOE affiliated National labs.
- This talk discusses the technical components chosen and implemented for the project so far.

InCommon

- InCommon Federation chosen as identity platform
- Participation from multiple DOE labs
  - ANL, BNL, LBNL, JLab, ORNL
- Users from participating labs can authenticate themselves using their lab credentials
- InCommon provides only SAML standard but no Oauth (which is a bit of a problem)
  - DOE HPC and General Purpose compute systems have limited support for SAML

CILogon

- Authentication hub for DCDE
- Serves as a proxy/broker service linked to Incommon
  - Able to translate SAML to Oauth
- Also provides X509 certificate service that is useful for integration with some services (eg. gsissh, gridftp)

COManage

- COManage service is integrated with CILogon

- It provides a web-portal like platform
  - for users to self-register
  - to manage and federate user attributes from multiple sources
  - admins to create groups and sub-groups
  - admins to manage project and user account lifecycle

AuthN/AuthZ: Participating Site Roles

- A site admin for each site
- Provide access to resources for DCDE project
- Provide a site entry point gateway "host" and a batch scheduler
- Obtain registered users' DN from COManage admin
- Create appropriate "mapfiles" mapping the user DNs to local site accounts
- Create local accounts and groups as appropriate to the local site policies

Globus and oauth-ssh

- Globus ssh provides the ssh over oauth
  - users can ssh to sites via oauth-ssh
- Globus transfer is used for data transfer purposes between participating sites
- Technical requirements for oauth-ssh
  - Choose a port (we chose 2222) for each site to be opened
  - Update DNS at site to add a text metadata record for the site gateway host for oauth-ssh to authenticate site using nslookup

Application and Containers

- We chose a microscopy application called Relion for prototyping purposes
- The application is containerized using Singularity
- Each site runs the same container and same version of Singularity for ease of portability and troubleshooting

Jupyter and Parsl

- Jupyter was integrated into the DCDE project whereby users can sign into the Jupyter web-interface using their DCDE credentials
- Jupyter's capability of custom authentication was linked with the CILogon interface
- CILogin OAuth tokens are available within the Jupyter notebooks to provide seamless authentication across the DCDE resources
- Parsl is chosen as the workflow platform
  - Written in Python -- a python package
  - Natural for Jupyter
  - Well integrated with Globus and oauth

Issues, Challenges and Lessons

- Some learning curve for users -- third-party auth, oauth-ssh, parsl, etc.
  - An approach is to provide a templated solutions to common user issues
- Site admin overhead eg. firewalls management, installation and configuration of oauth stack, jupyterhub etc.
  - Scripted several install steps

Summary

- The DCDE project aims at enabling users at national labs access and use compute and data resources of other national labs.
- We prototype a solution to achieve this objective.
- Challenges such as ease of user on-boarding, tech-learning curve remain.
- Turning DCDE into a turn-key solution for users will open up a broad and efficient inter-lab utilization of resources.

References

- CILogon https://www.cilogon.org/2
- COManage https://www.cilogon.org/comanage
- JupyterHub/Jupyter https://jupyter.org/hub
- oauth-ssh https://github.com/XSEDE/oauth-ssh
- Globus https://www.globus.org
- Parsl http://parsl-project.org
- DCDE Templates https://github.com/bnl-sdcc/dcde-templates
- scripts and config https://github.com/bnl-sdcc/dcde-templates

U.S. DEPARTMENT OF ENERGY | Office of Science

Future Lab Computing Working Group (FLC-WG)

Distributed Computing and Data Ecosystem (DCDE)

Lessons learned

MAGIC @ SC 19

Ray Bair (ANL), Kuldeep Chawla (LBNL), Luis Corrales (LBNL), David Cowley (PNNL), Zhihua Dong (BNL), Eric Lancon (BNL), Ketan Maheshwari (ORNL), Daniel Murphy-Olson (ANL), Mallikarjun Shankar (ORNL)

ASCR PM: Richard Carlson

Lessons Learned

- Federated Identity Management
  - Workarounds are currently required at each site to accept a non-local DCDE member registered in a DCDE CoManage group.
  - Each DCDE member has a unique CoManage ID: Can this be used to provision a unique Unix account at DCDE participating sites? Options: accounts pre-provisioning vs creation on the fly? Work is needed to create good mapping mechanisms between pre-provisioned accounts and local groups and allocation systems.
  - CoManage groups may have various capabilities at each participating site. We need a way to publish these capabilities. (single DCDE group in the pilot).
  - CoManage attributes may provide a path to propagate user policy acceptance and training requirements.

Lessons Learned (contd.)

- Portability across laboratories
  - Compatibility of singularity versions at different labs. Other container technologies?
  - Simple users may not know how to create a container? We may need a service to encapsulate applications.
  - Operating environments, libraries, paths to filesystems may change over time at participating sites. A standard way of communicating this information needs to be developed.

Lessons Learned (contd.)

- Workflow through analytic notebooks
  - Parsl: simple, local & global, Python, integrated with Jupyter. Still in development.
  - Each JupyterHub submission host needs firewall exceptions. Need for special edge services at each participating Lab to interface with other participating labs (similar to existing DTN edge systems but for execution).

Lessons Learned (contd.)

- Data Transfer
  - Globus now a commercial capability? De facto standard but alternatives?
  - Ultimately data transfer should be hidden to basic users, similarly storage locations. Need for a dynamic global storage.

Lessons Learned (contd.)

- Information and policies
  - Each participating site needs to allocate & publish resources to DCDE. How? Accuracy/reliability of information?
  - What are the expectations from the various participants? Service level agreement?
  - How to manage policies and priorities?

MAGIC Meeting Discussion

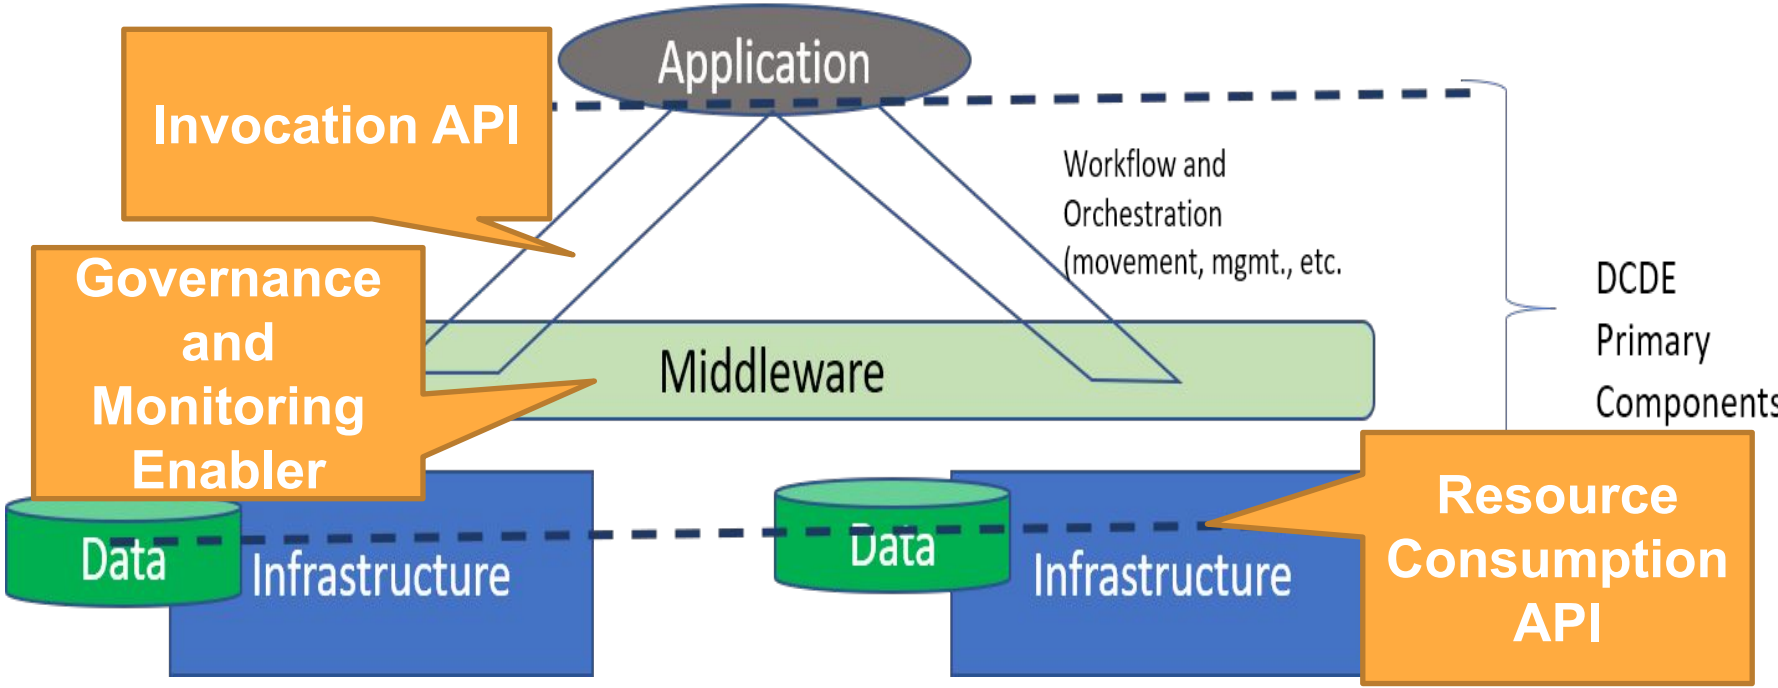

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov