Remote Computing Enablement and You

LC User Meeting

Todd Heer
ASC Facilities, Operations, and User Support
Deputy Program Lead
RCE Lead, LLNL
theer@llnl.gov

December 8th, 2020

LLNL-PRES-817469: This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Remote Computing Enablement and You

Agenda

- The *What, Why* and *Who* of RCE
- Foreach thrust area
  - What is/was the starting state?
  - What did/are we achieving?
  - What's on the horizon?

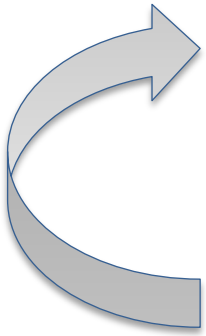

Putting the “C” in RCE

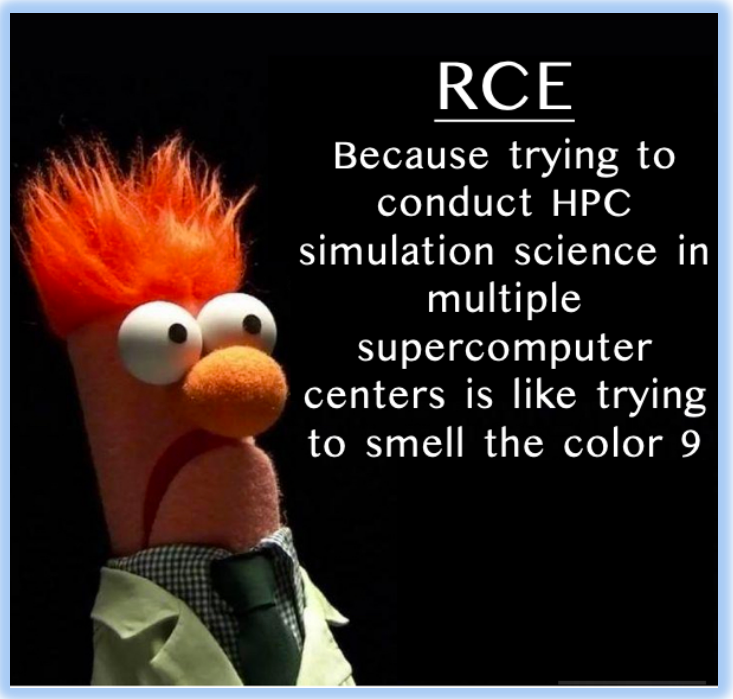

*The single biggest problem in **communication** is the illusion that it has taken place.*

- George Bernard Shaw

***Cooperativeness** is not so much learning how to get along with others, as taking the kinks out of ourselves so that others can get along with us.*

-Thomas S. Monson

*The real art of **conversation** is not only to say the right thing at the right place but to leave unsaid the wrong thing at the tempting moment.*

- Dorothy Nevill

RCE Mission

*Attempt to make the **remote HPC user experience as close as possible to the local user experience** to maximize productive utilization of computing resources across the NNSA HPC simulation complex.*

The RCE team, formed in May of 2019, is comprised of many members from multiple disciplines and management strata across the Sandia, LANL, and LLNL HPC centers.

RCE came about as a confluence of events

- Trinity at LANL, Sierra at LLNL, and of course the eventual arrival of El Capitan at LLNL
- Center of Excellence (CoE) efforts
- Ongoing semi-disparate multi-lab meetings
  - Common Computing Environment (CCE)
  - Tri-Lab Data Movers (TDM) telecon
  - User oriented meetings (e.g. *Tri-Lab Remote Computing Preliminary Planning Meeting* of 5/2019)
- NNSA HQ desire to further multi-lab cooperation to aid remote user bases

ASC FOUS leads at the three sites have programmatic responsibility

RCE focuses on *attainable* deliverables in *measured tempos* that deliver real benefits to the HPC Tri-Lab user community

The Problem Space is Huge

Where to focus effort?

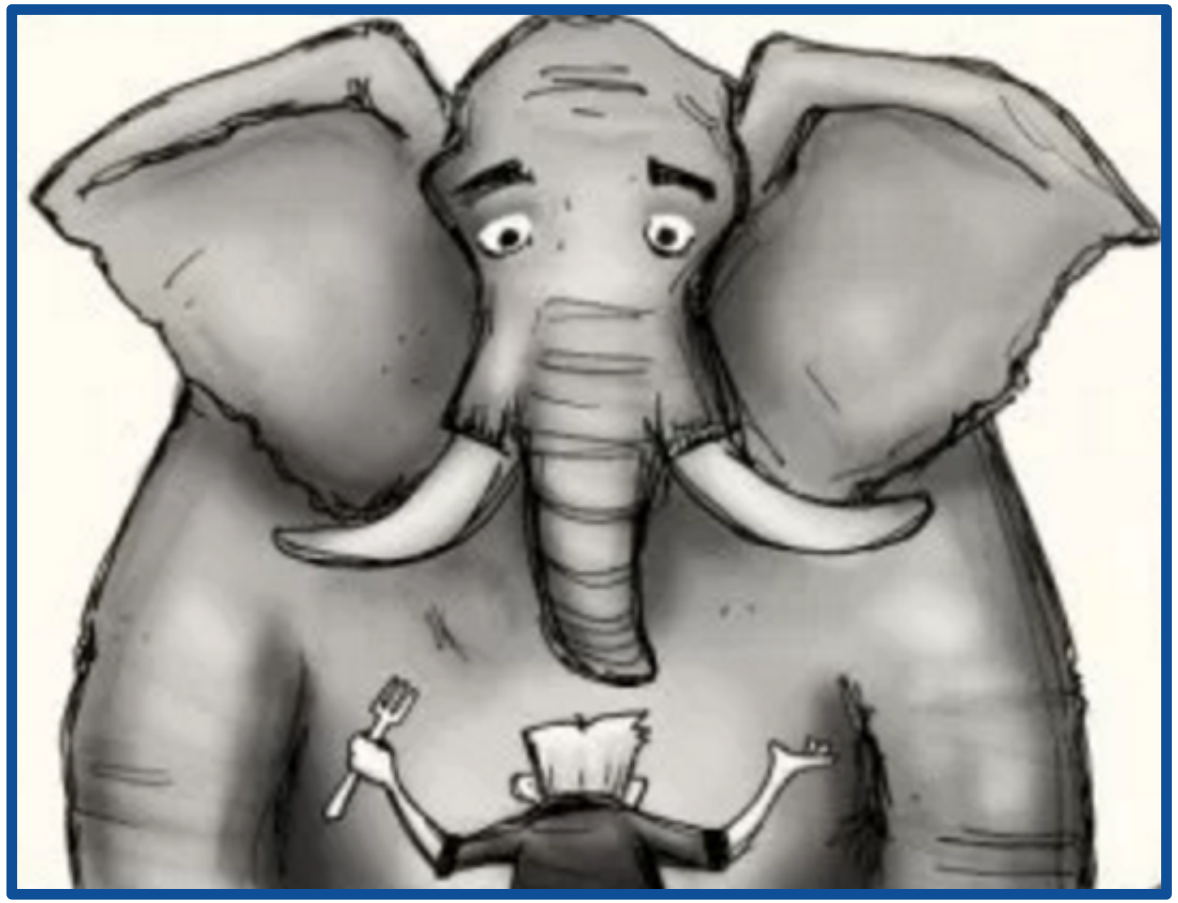

- SMEs are best able to identify problem areas and low-hanging fruit (read: *slam dunks*).
  - Data transfer tools lack of ubiquity
  - Network bottlenecks
  - Gateways and what could be done to *dissolve* them
- Check-ins with user community (either unsolicited or user-driven)
- Identify projects that our in our normal charter (and thus already budgeted at some level)
- Some efforts *do* require an influx of money and/or more strategic planning
- Tri-Lab management (and HQ) help agree upon cooperative efforts and priority

RCE has achieved best success by shedding light in certain areas and enabling the ***conversation***

The Problem Space is Huge

How to focus effort?

- Fluid working groups address items, tasks, projects

Q) How do you eat an elephant?

A?) One bite at a time

A) Parallelize. Many *bytes* simultaneously

HPC approved

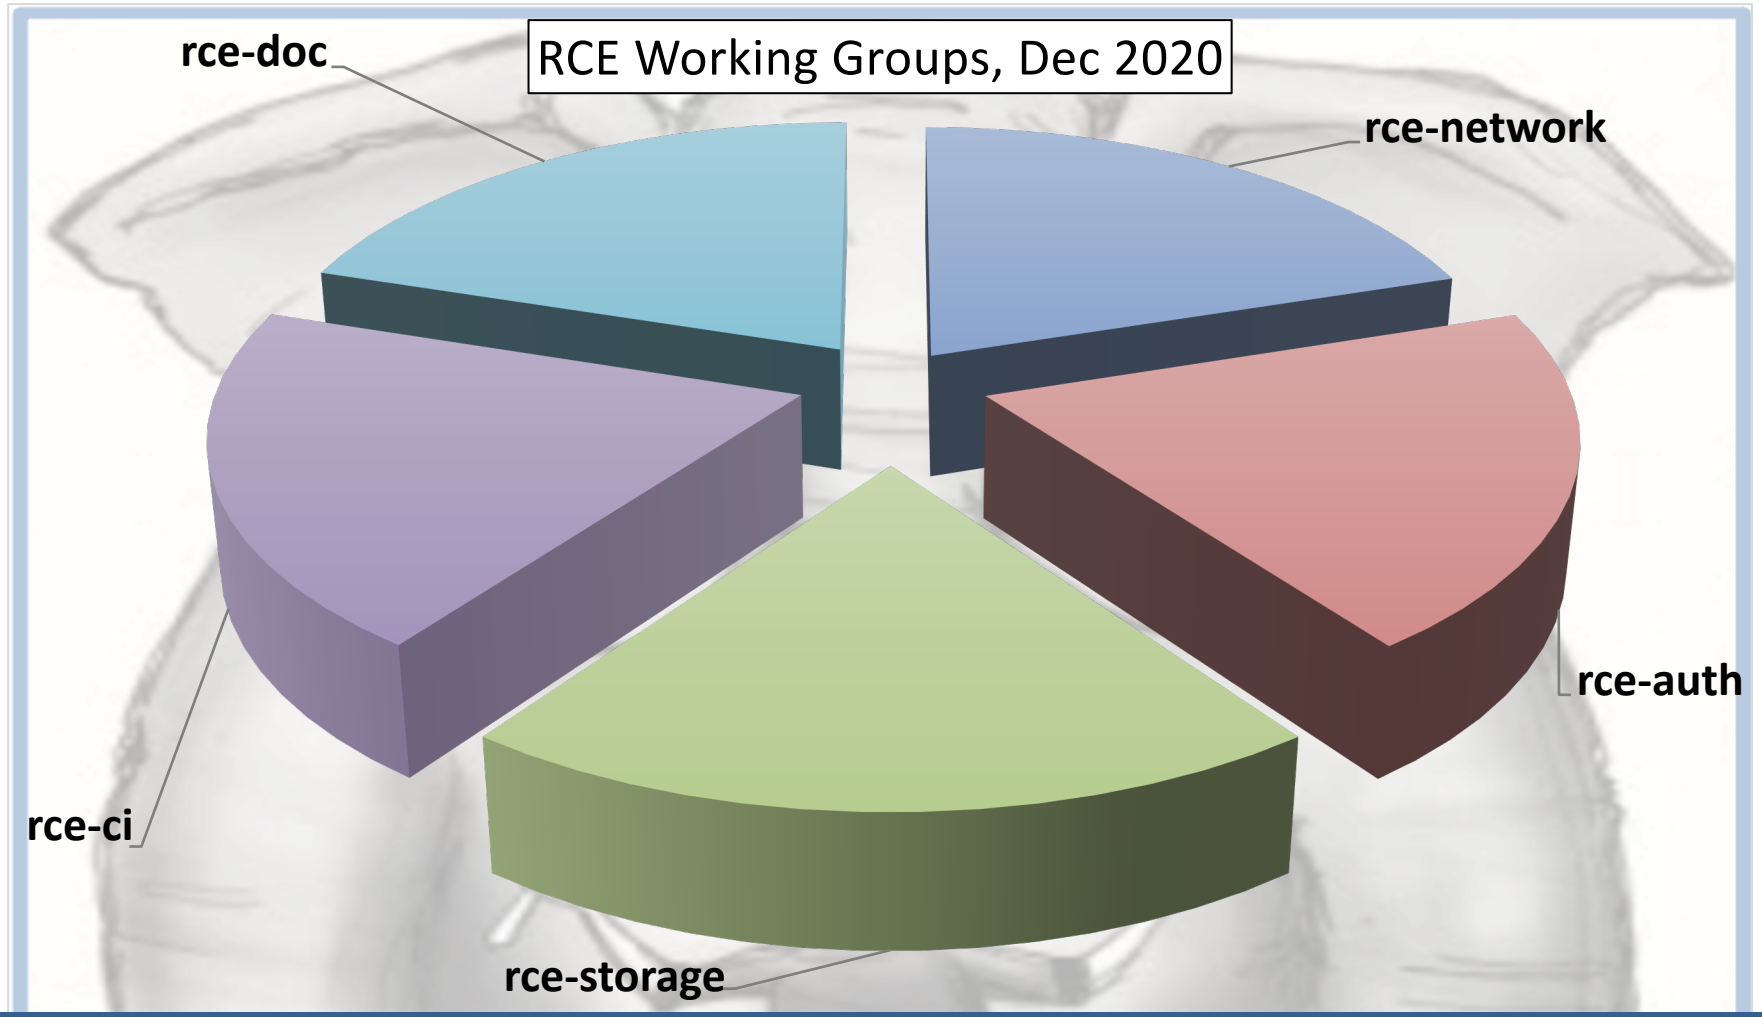

Tri-Lab working groups meet via Webex, use email reflectors, and have Mattermost discussions

NNSA HPC User Community

RCE User Check-Ins

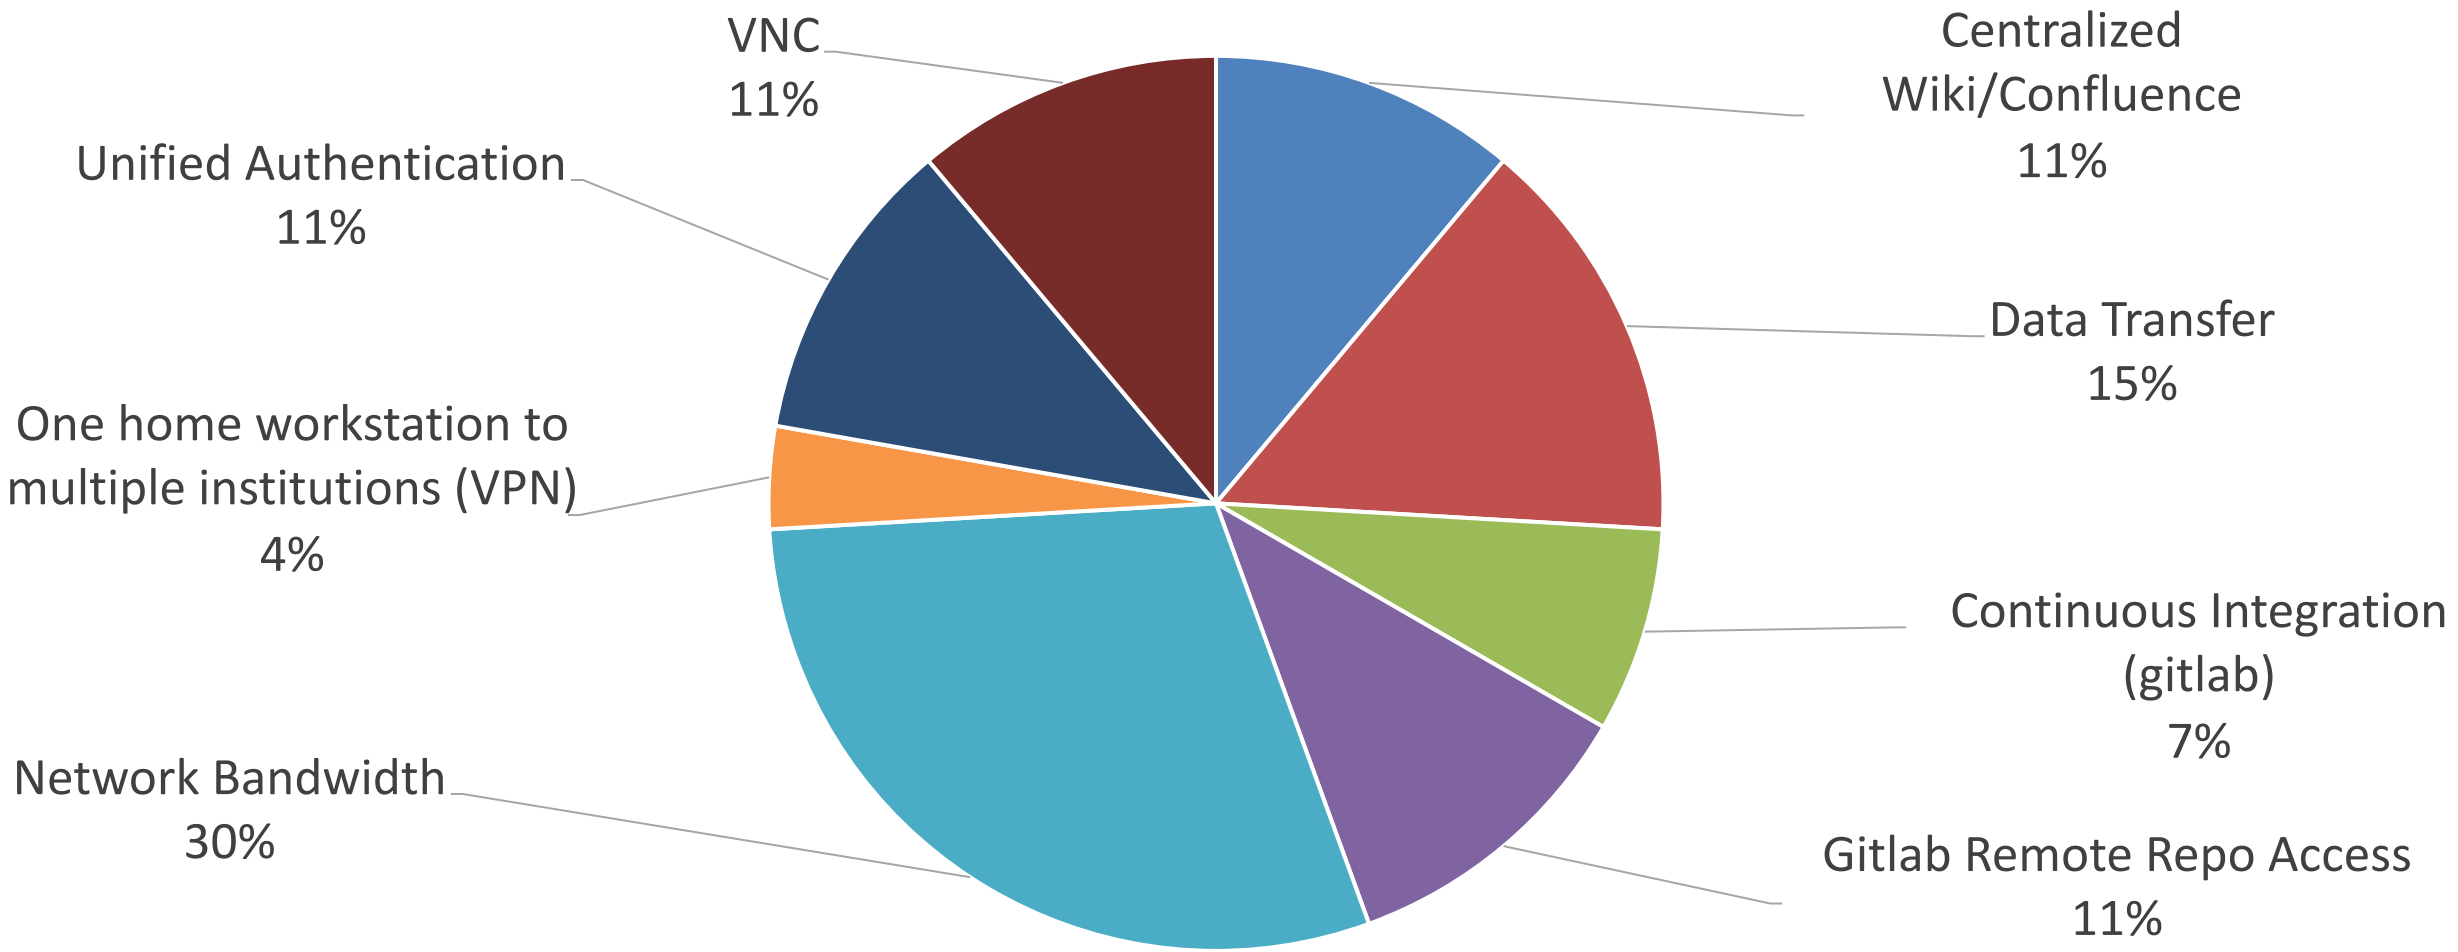

RCE Categories and Examples
Current Grouping of Identified Areas of Interest

- Hardware needs that require multi-lab coordination
  - 100Gb DISCOM encryptors
  - Gitlab server hardware (and licensing)
- Identify SME efforts requiring multi-lab coordination
  - Tuning / chasing network maladies (e.g. dropped packets, buffer overflows, etc.)
- Reduce barriers to effective remote code development
  - Eliminate gateways
  - Find opportunities for unified authentication
- Identify software gaps
  - Data xfer tools
  - VNC
- Multi-lab meetings
  - Hone to reduce repetition and increase SME coverage/involvement
- User support / Documentation
  - Continued efforts on Sarape, hpc.llnl.gov, hpc.sandia.gov, hpc.lanl.gov
- Strategic
  - Gitlab Continuous Integration across the NNSA Tri-Lab HPC complex

Example Area For Improvement

Category: Reduce Barriers to Remote Code Development
Kerberos cross-realm ssh authentication

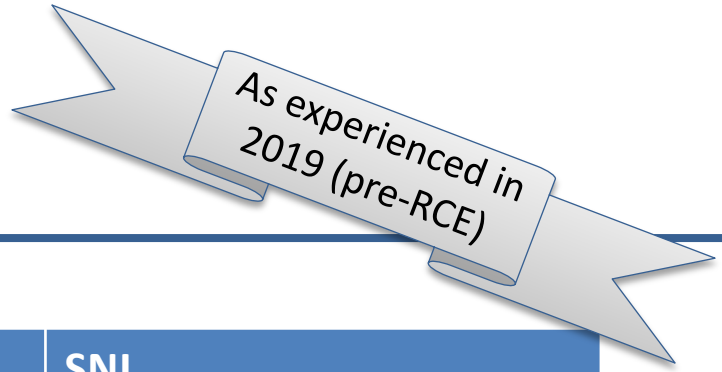

Network Type	LLNL	LANL	SNL
Open "collaborative"	CZ: No gateway	Turquoise: wtrw.lanl.gov	OHPC: No gateway
Open "restrictive"	RZ: rzgw.llnl.gov	Yellow: ihpc-gate.lanl.gov	SRN: srngate.sandia.gov
Closed/classified	SRD: No gateway	Red: red-wtrw.lanl.gov	SCN: No gateway

NOTE: rows across institutions do NOT imply the network types have the same definition/restrictions.

Example Area For Improvement

Category: Reduce Barriers to Remote Code Development
Kerberos cross-realm ssh authentication

As experienced in 2019 (pre-RCE)

Network Type	LLNL	LANL	SNL
Open “collaborative”	CZ: No gateway	Turquoise: wtrw.lanl.gov	OHPC: No gateway
Open “restrictive”	RZ: rzgw.llnl.gov	Yellow: ihpc-gate.lanl.gov	SRN: srngate.sandia.gov
Closed/classified	SRD: No gateway	Red: red-wtrw.lanl.gov	SCN: No gateway

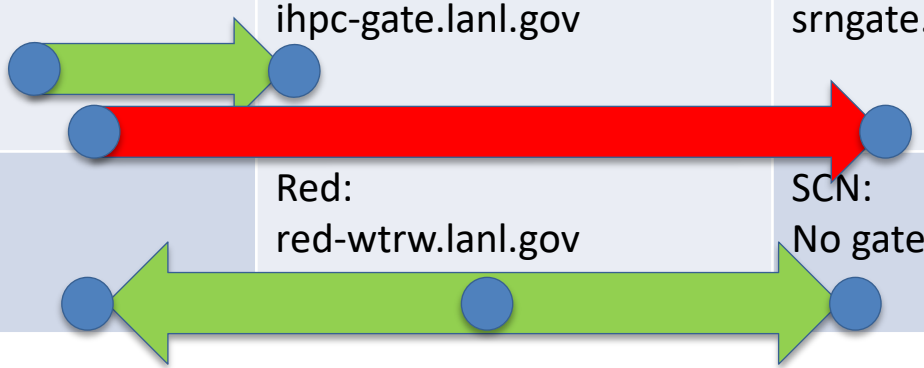

Example Area For Improvement

Category: Reduce Barriers to Remote Code Development
Kerberos cross-realm ssh authentication

As experienced in 2019 (pre-RCE)

Network Type	LLNL	LANL	SNL
Open “collaborative”	CZ: No gateway	urquoise: wtrw.lanl.gov	OHPC: No gateway
Open “restrictive”	RZ: rzgw.llnl.gov	Yellow: ihpc-gate.lanl.gov	SRN: srngate.sandia.gov
Closed/classified	SRD: No gateway	Red: red-wtrw.lanl.gov	SCN: No gateway

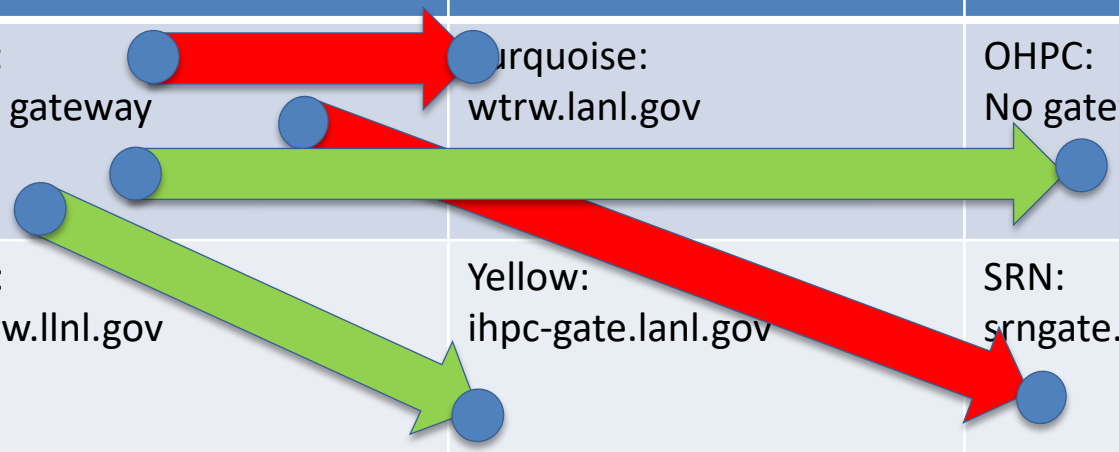

Example Area For Improvement

Category: Reducing Barriers to Code Development

Reduce/Remove gateway hosts

As experienced in 2019 (pre-RCE)

Network Type	LLNL	LANL	SNL
Open "collaborative"	CZ: No gateway	Turquoise: wtrw.lanl.gov	OHPC: No gateway
Open "restrictive"	RZ: rzgw.llnl.gov	Yellow: ihpc-gate.lanl.gov	SRN: srngate.sandia.gov
Closed/classified	SRD: No gateway	Red: red-wtrw.lanl.gov	SCN: No gateway

Example Area For Improvement (Achieved)

Category: Reducing Barriers to Code Development
Reduce/Remove gateway hosts

Network Type	LLNL	LANL	SNL
Open "collaborative"	CZ: No gateway	Turquoise: wtrw.lanl.gov	OHPC: No gateway
Open "restrictive"	RZ: No gateway	Yellow: ihpc-gate.lanl.gov	SRN: srngate.sandia.gov
Closed/classified	SRD: No gateway	Red: No gateway	SCN: No gateway

Example Area For Improvement

Category: Reducing Barriers to Code Development

Reduce/Remove gateway hosts+++

On the RCE Horizon

Network Type	LLNL	LANL	SNL
Open "collaborative"	CZ: No gateway	Turquoise: wtrw.lanl.gov	OHPC: No gateway
Open "restrictive"	RZ: No gateway	**RE** (restricted enclave): No gateway	SRN: srngate.sandia.gov
Closed/classified	SRD: No gateway	Red: No gateway	SCN: No gateway

1) Entirely new HPC network enclave

2) Removal of gateway machine

3) Cross-realm trust model

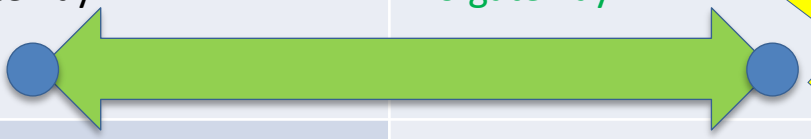

Example Area for Improvement
Category: **Classified** WAN Networking

Frontend network: Used for authentication, ssh, web, application (simulation visualizations, graphical debuggers, etc.).
Managed by NNSA OCIO

Tri-Lab Backend network: Used for data transfer (pftp, hsi, certain protocols under Hopper)
Managed by Tri-Lab HPC Centers

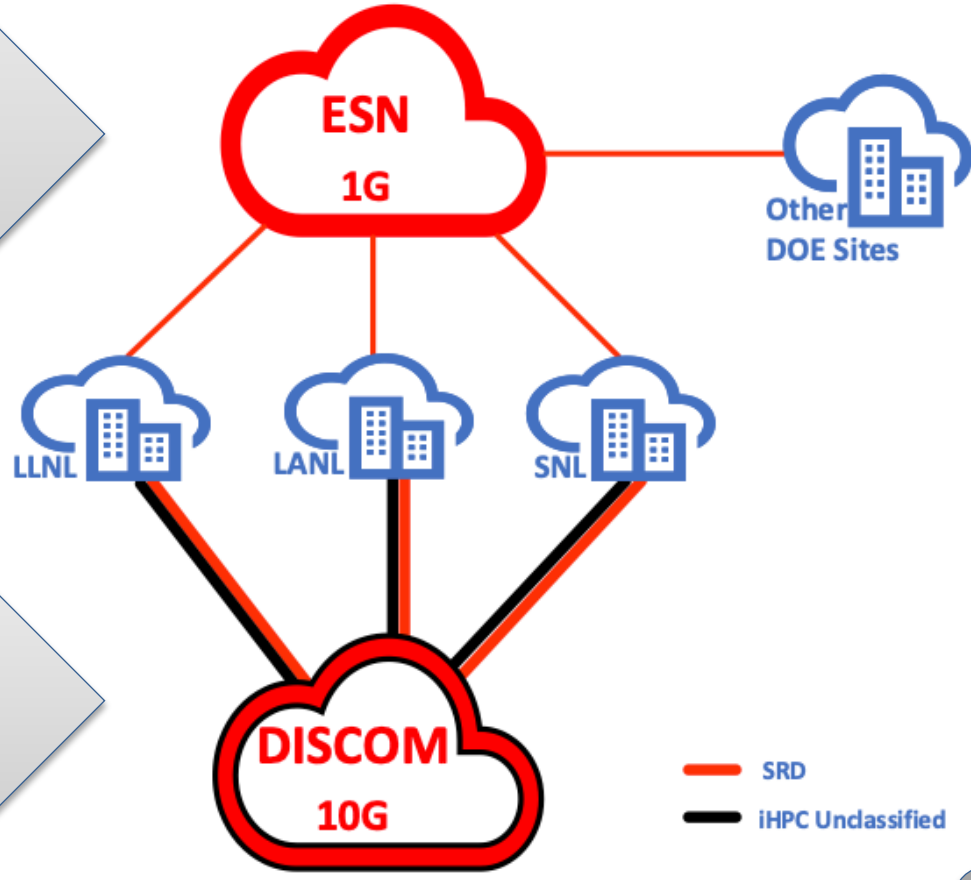

What is needed today is for the ESN managed equipment to be upgraded to support remote computing at higher speeds to absorb pressures from platforms currently in service to the tri-lab community

- Tri-Lab ASC Program Directors plus ASC Program Mgr NNSA
Sept. 2019

Example Area for Improvement

Category: WAN Networking

Secure Tri-Lab DISCOM Network Performance

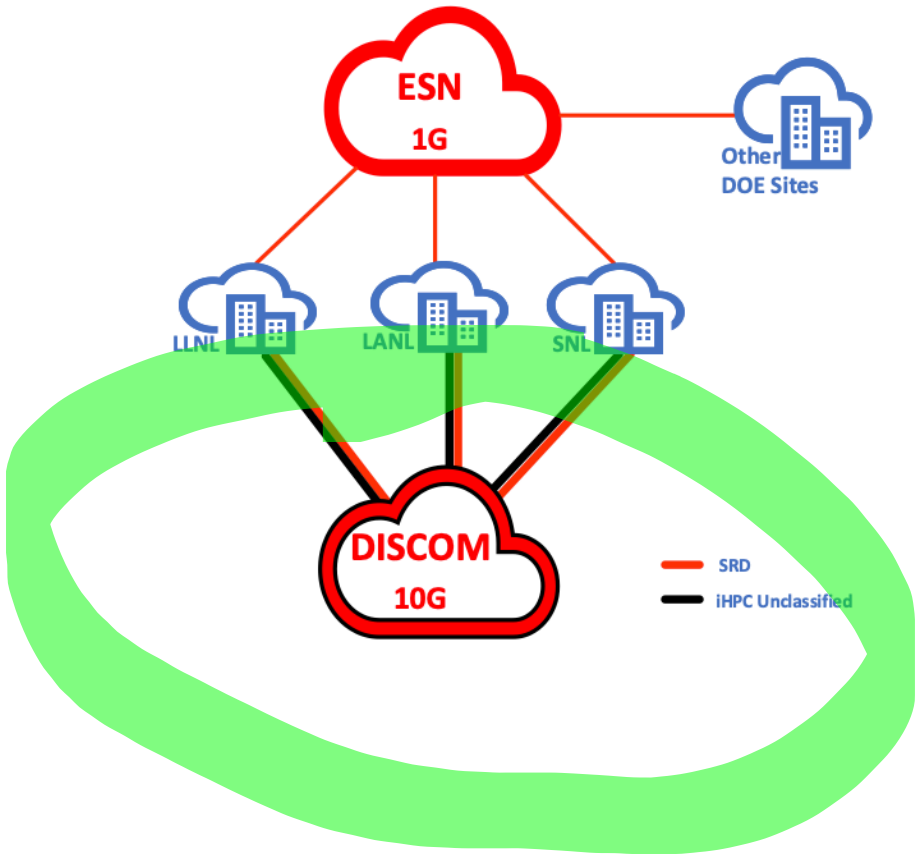

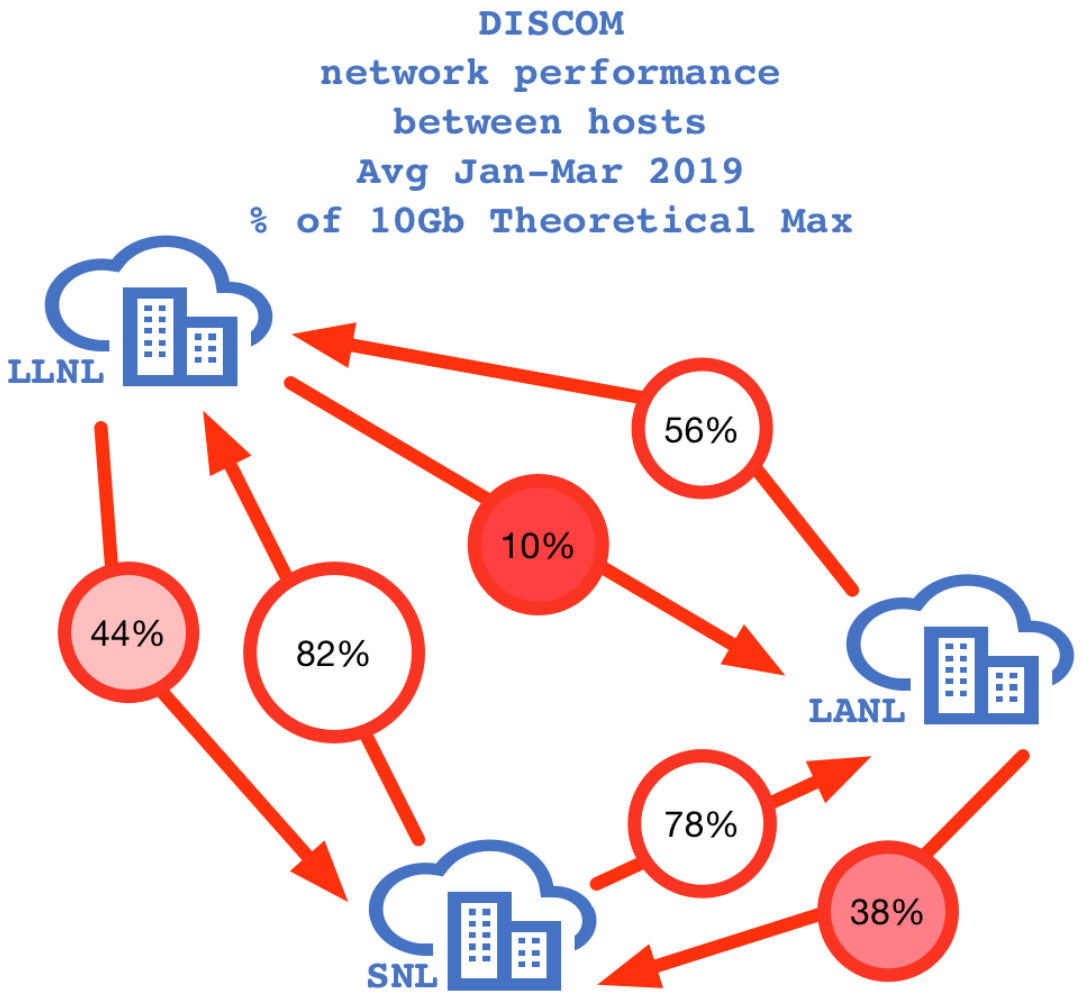

RCE driven network improvements

On the RCE Horizon

- These two projects are currently budgeted and underway
- 10x improvement in bandwidth for both frontend and data transfer networks
- Red lines are classified networks
- Note importance of leveraging iHPC <u>unclassified</u> footprint

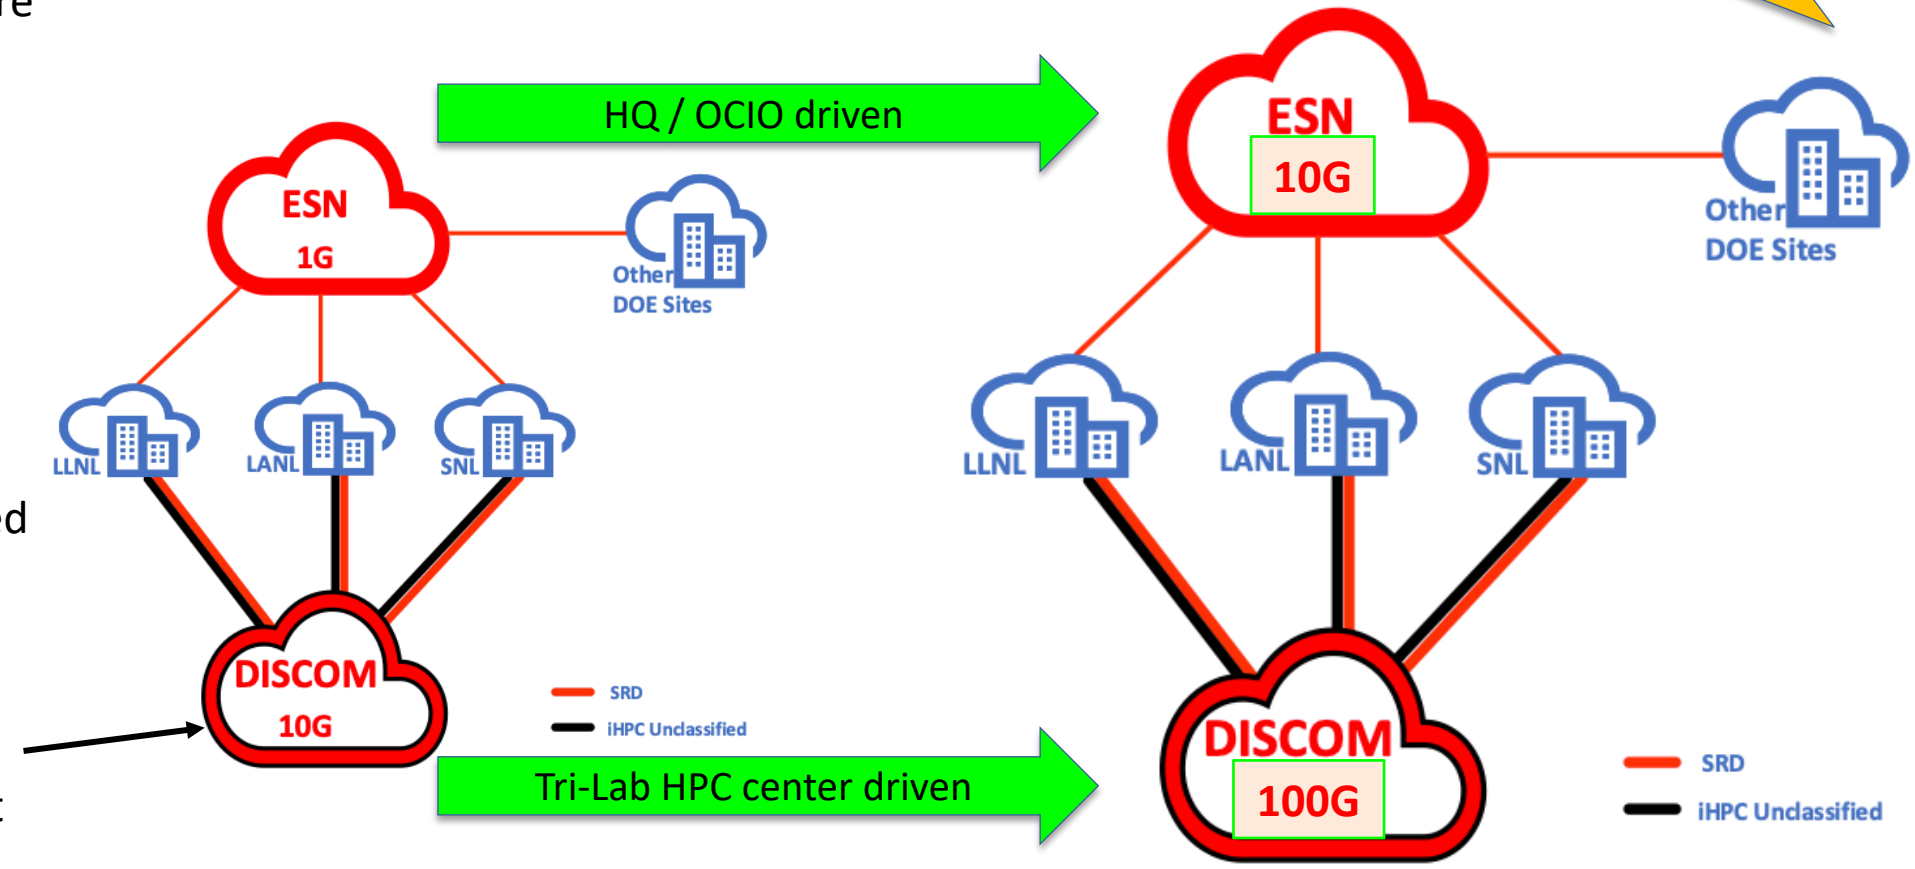

Example Area for Improvement

Category: Identify Software Gaps

Unify endpoint choices and data transfer mechanisms

As experienced in 2019 (pre-RCE)

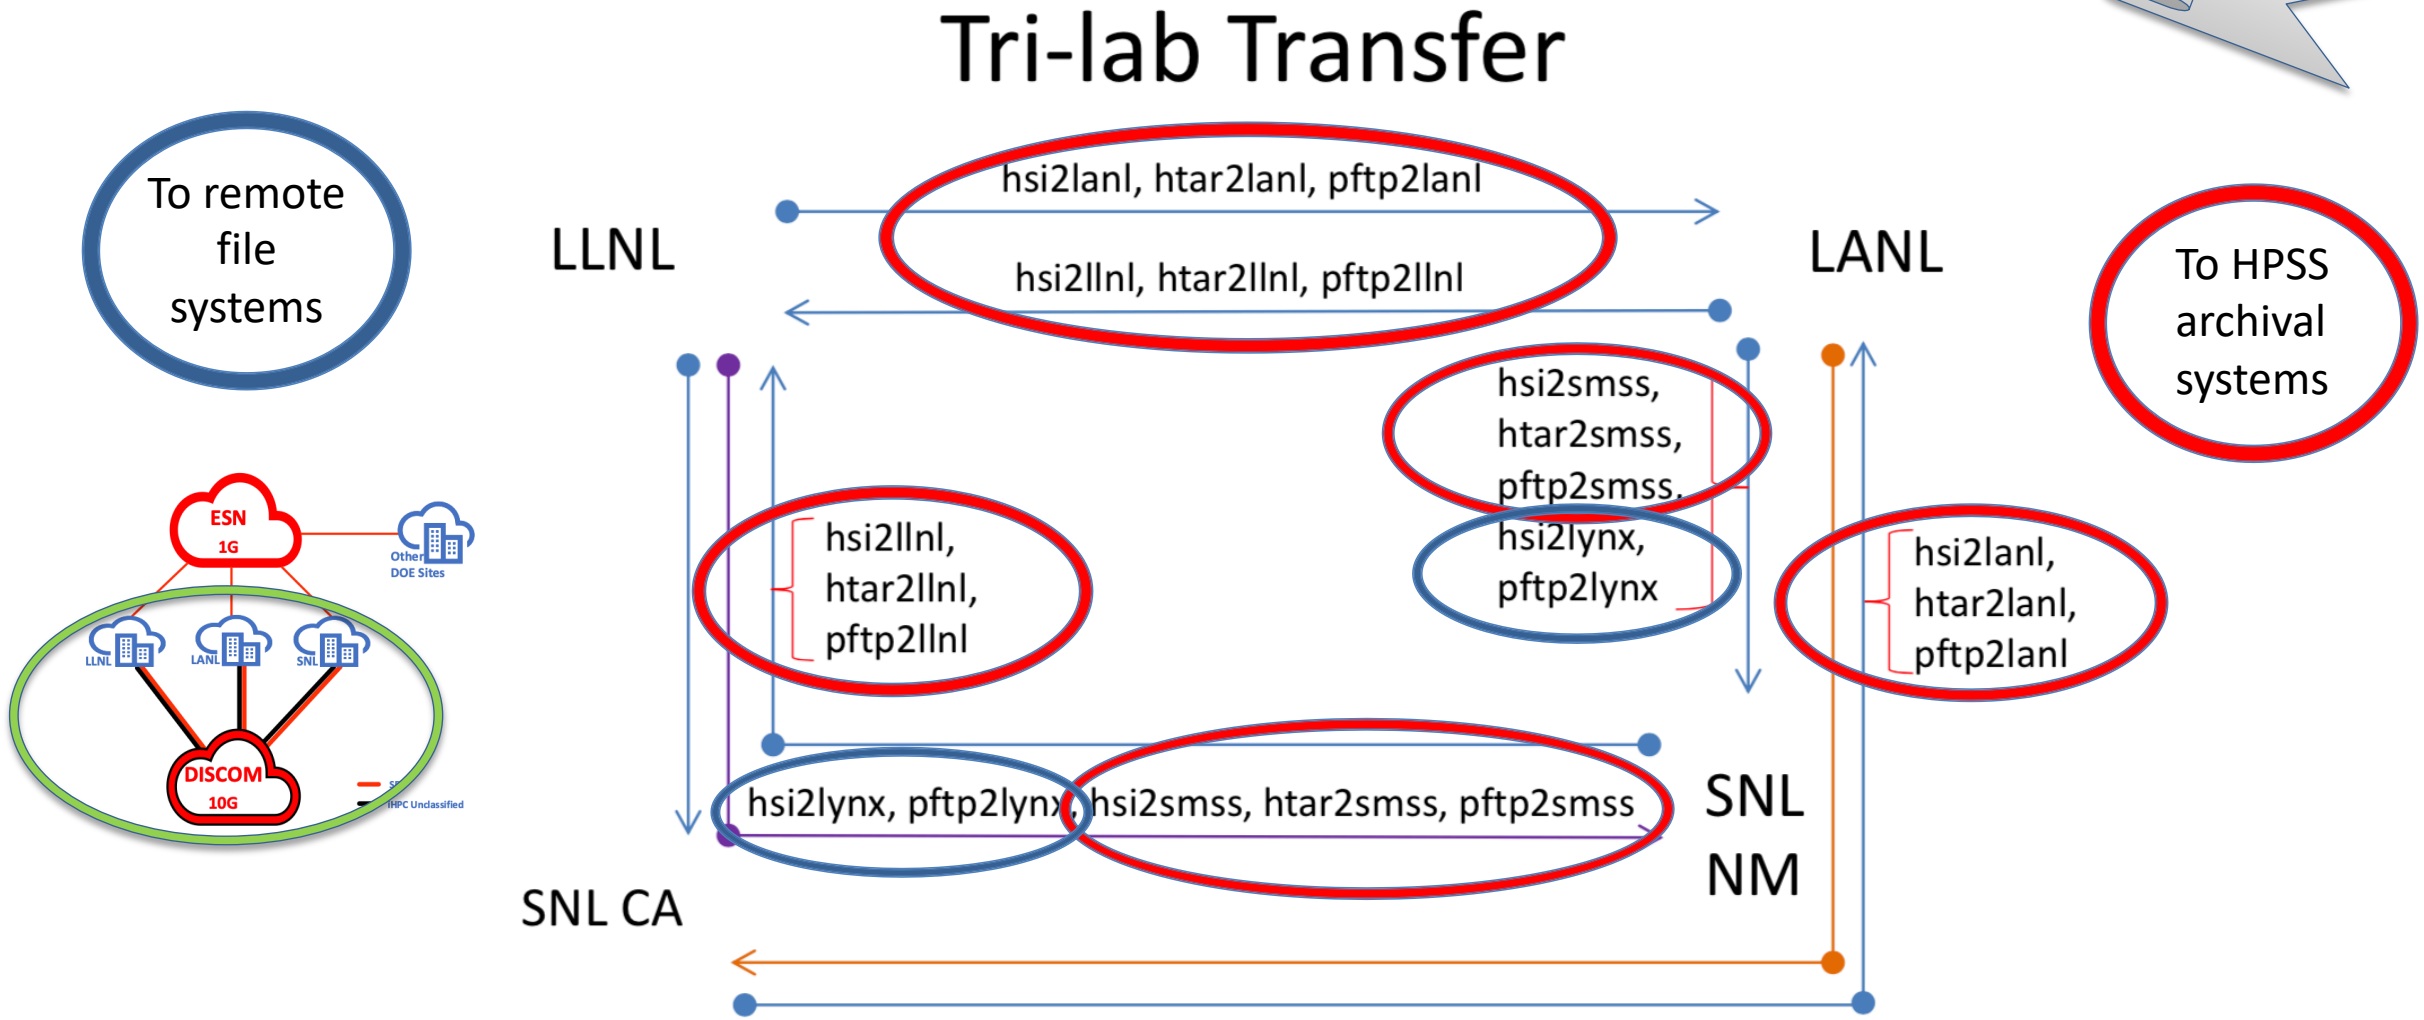

Example Area for Improvement

Category: Identify Software Gaps

Classified Tri-Lab DISCOM WAN Data Transfer, **file system to file system**

As experienced in 2019 (pre-RCE)

- At either LLNL or LANL connecting to SNL:
  hsi2lynx “cd sync_dir; put -R source_dir”
  OR
  pftp2lynx <...>

- At either LANL or SNL connecting to LLNL:
  rsync -av source_dir/ cslic.llnl.gov:target_dir
  OR
  scp -r source_dir cslic.llnl.gov:target_dir
  OR
  sftp

- At either LLNL or SNL connecting to LANL:
  rsync -av **--rsh='ssh red-wtrw.lanl.gov ssh'** source_dir/ redcap.lanl.gov:target_dir
  OR
  scp -r source_dir **red-wtrw.lanl.gov:**redcap:target_dir

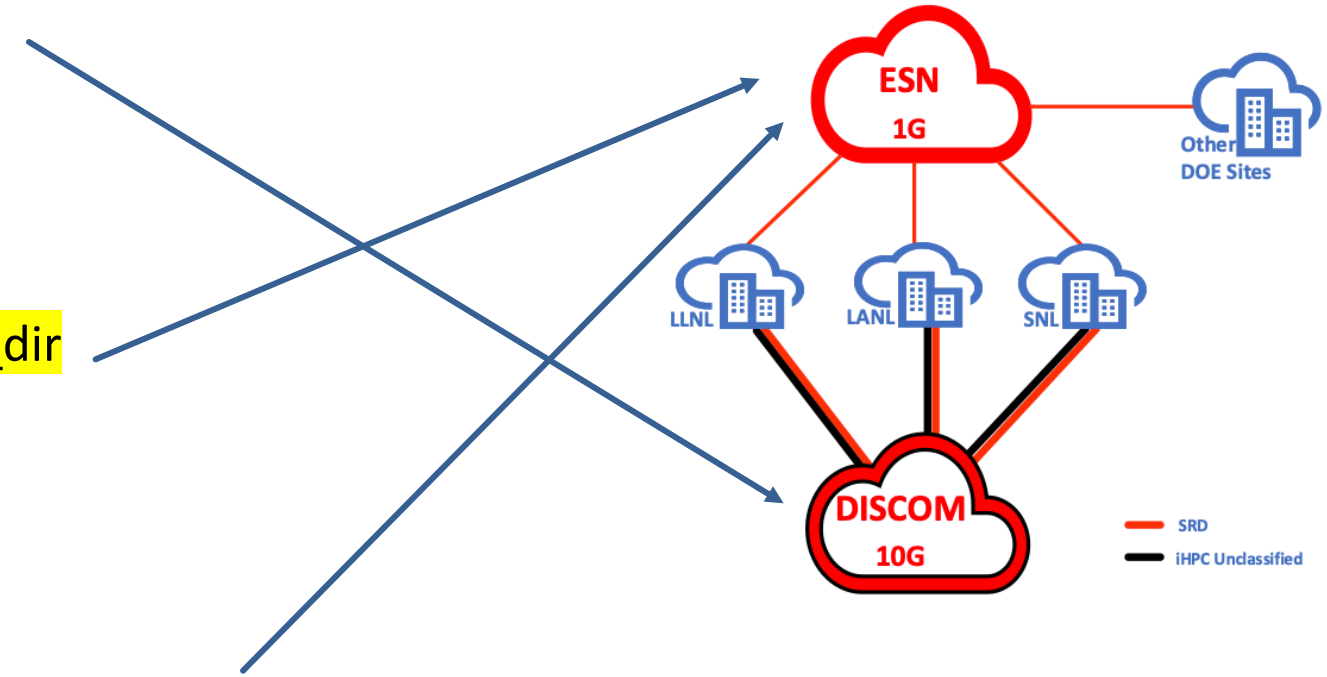

Example Area for Improvement

Category: Identify Software Gaps

Classified Tri-Lab DISCOM WAN Data Transfer, file system to file system

RCE Achieved Today

Classified High Bandwidth Data Movement

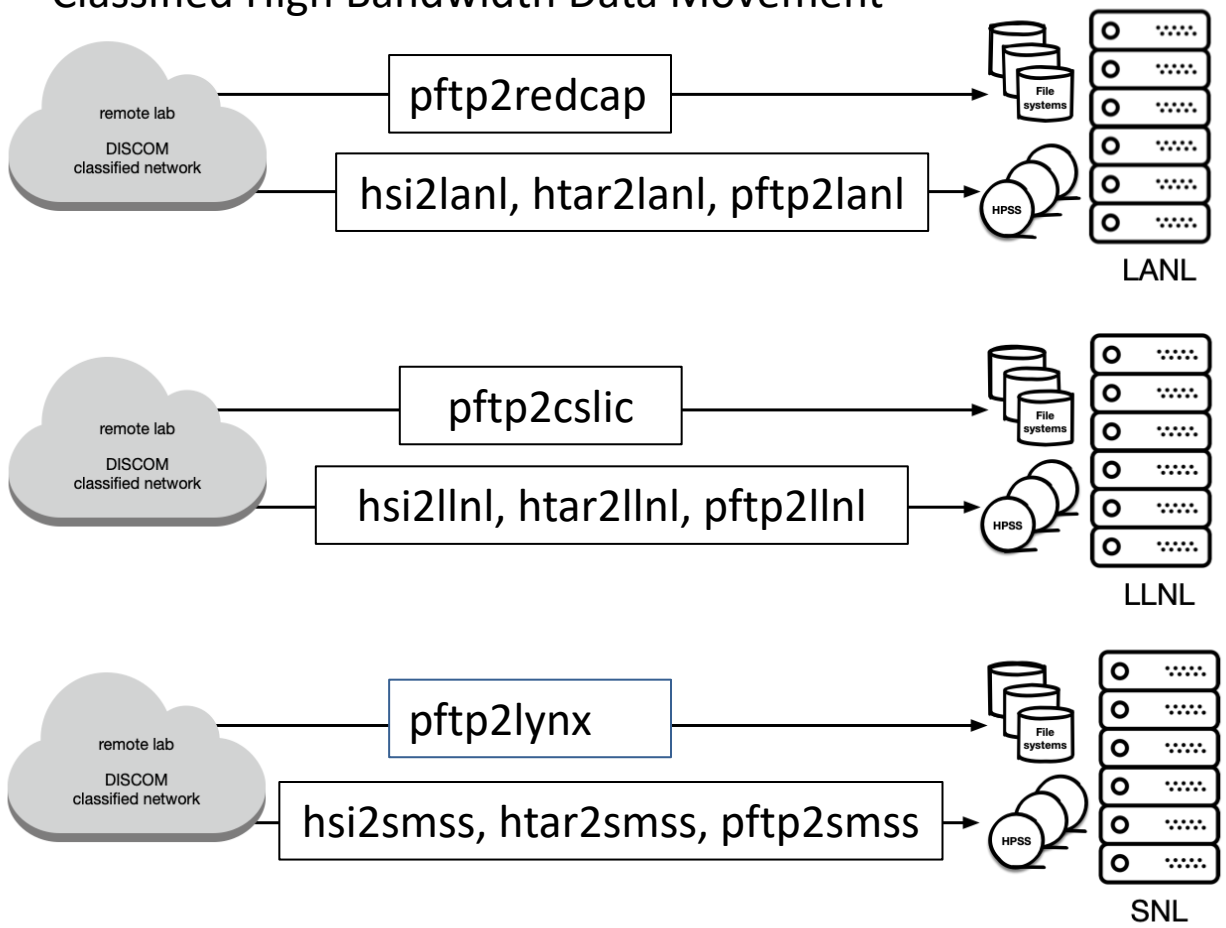

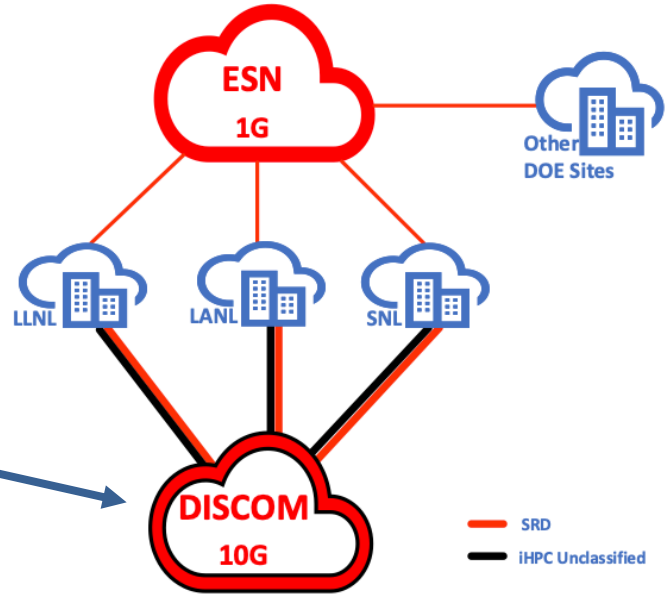

Documented on hpc.llnl.gov

Example Area for Improvement

Category: Identify Software Gaps

Unclassified Tri-Lab iHPC WAN Data Transfer, file system to file system

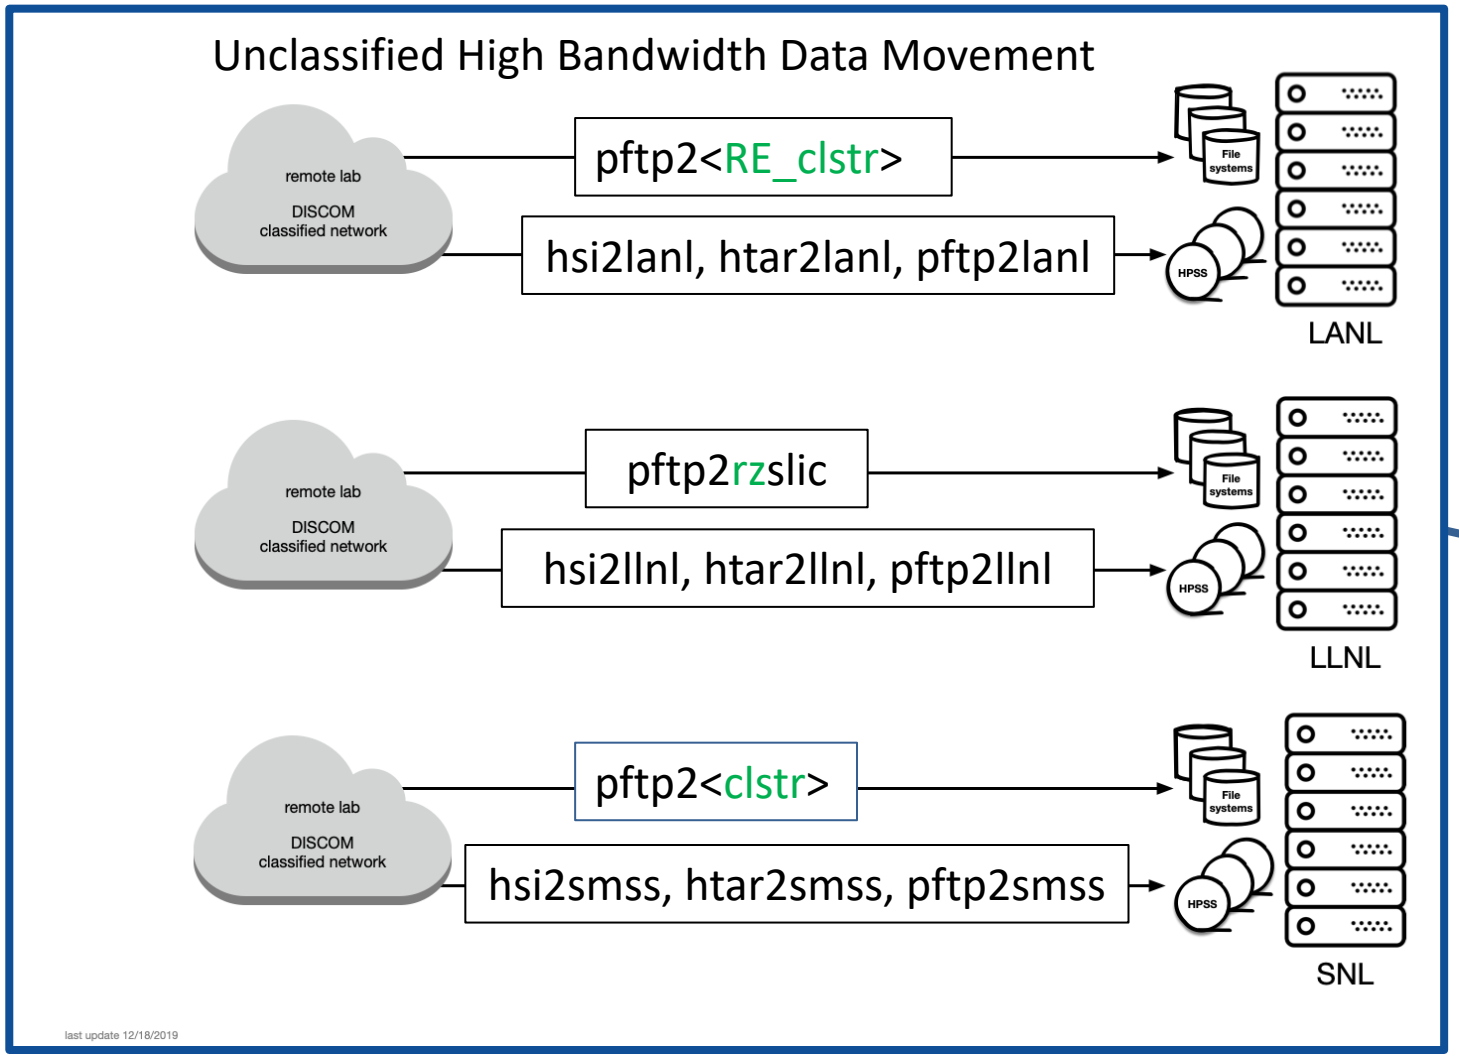

On the RCE Horizon

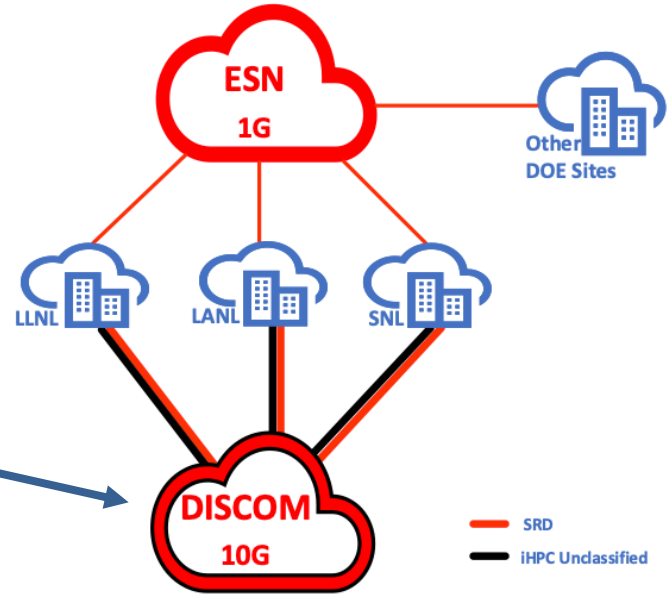

- Not decided at this time
- May require special budget allocations (e.g. DTNs)

FY21 Goal: Tri-Lab HPC Center GitLab Continuous Integration

Mirror Repositories – Open Restricted Enclave

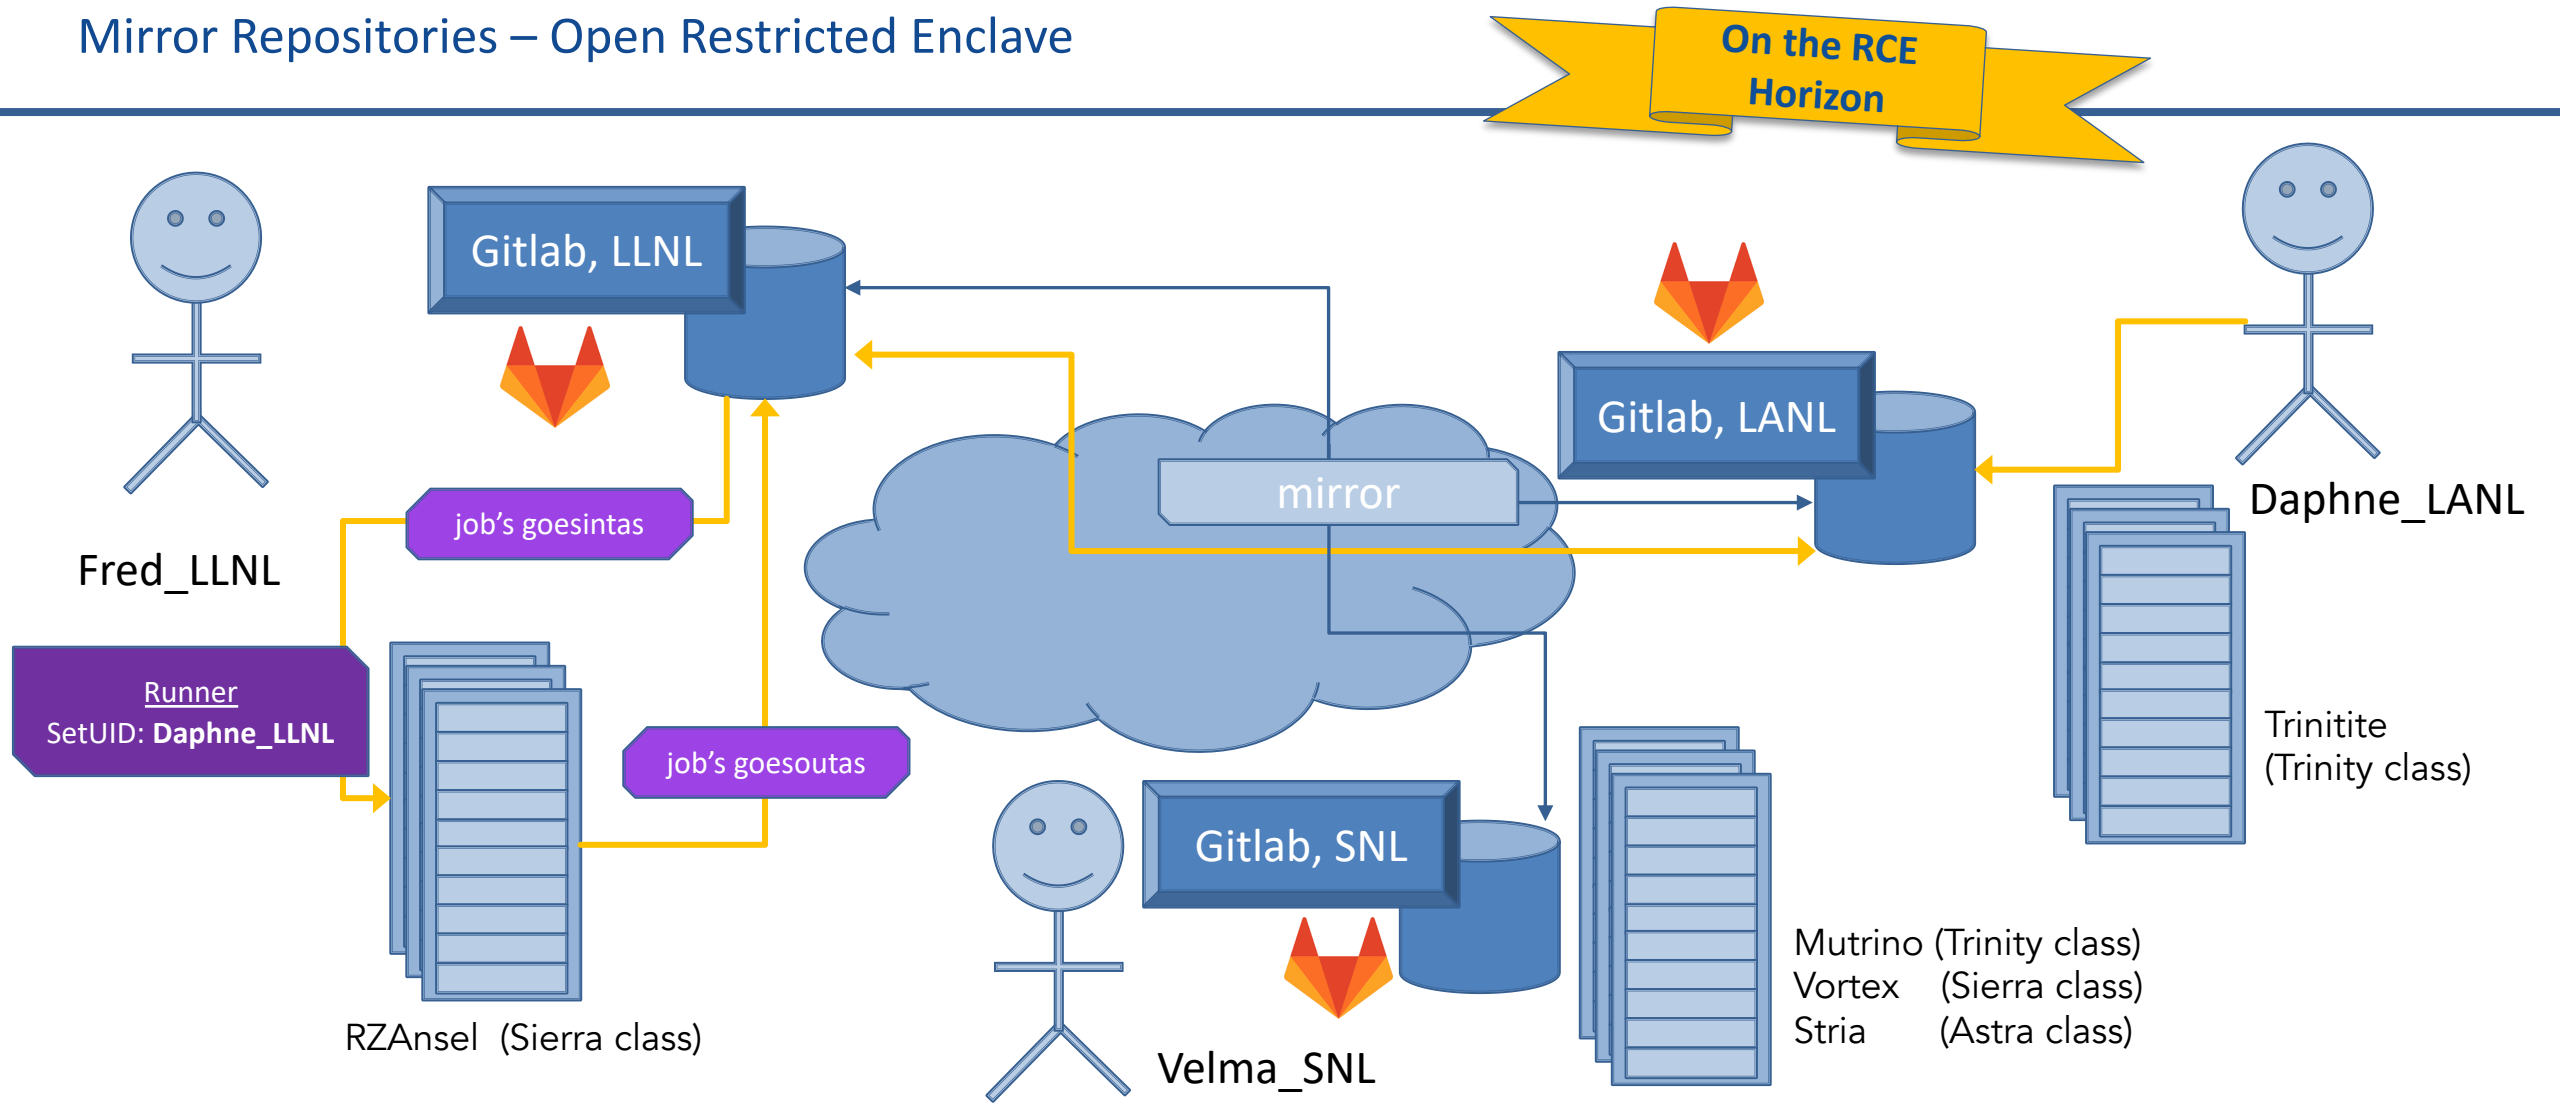

More local Gitlab CI info and How-To at *lc.llnl.gov/confluence*

You are not alone

Huge thanks to all the RCE participants at the labs!

Reference:

- sarape.sandia.gov (remote accounts)
- hpc.llnl.gov (remote access How-To's)
- lc.llnl.gov/confluence (gitlab CI)
- **LC Hotline** – *any* question
- RCE issues/feedback:

Todd Heer theer@llnl.gov

Livermore Computing Hotline
Hours: M-F: 8A-12P,1-4:45P
Email: lc-hotline@llnl.gov
Phone: 925-422-4531

SSH Gateways		Destination									
		LLNL CZ	LLNL RZ	LLNL SCF	LANL Turquiose	LANL Yellow	LANL Red	SNL OHPC	SNL IHPC	SNL SRN	SNL SCN
Source	LLNL EN*	no gw	no gw		wtrw	ihpc-gate via LLNL RZ/CZ		no gw	via LLNL RZ/CZ	srngate	
	LLNL CZ	-			wtrw	ihpc-gate		no gw		srngate	
	LLNL RZ	no gw	-		wtrw	ihpc-gate		no gw	no gw	srngate	
	LLNL SCF			-			no gw				no gw
	LANL Turquiose				-						
	LANL Yellow*	no gw	ihpc-gate		wtrw	-		no gw	ihpc-gate	srngate	
	LANL Red			no gw			-				no gw
	SNL OHPC	no gw			wtrw	ihpc-gate		-	no access	srngate	
	SNL IHPC	no gw	no gw		wtrw	ihpc-gate		no gw	-	srngate	
	SNL SRN*	no gw	via SNL IHPC		wtrw	ihpc-gate		no gw	no gw	-	
	SNL SCN			no gw			no gw				-

*desktops live here

NOTES: o when filling out above, make sure to include the possibility of TWO or more gateways in the single cell if the source site has an outbound gateway and the destination site has a

o the above is intended to illustrate only SSH gateway hosts and does not consider authentication types or requirements, and may yet serve as a springboard for those more inv

HPC Cluster Resource Examples								
LLNL CZ	lassen (ATS)	oslic	catalyst	corona				
LLNL RZ	rzansel (ATS)	rzslic	rztopaz	rztrona	rzhasgpu			
LLNL SFC	sierra (ATS)	cslic	jade (CTS)	zin				
LANL Turquoise	badger (ba-fe, CTS)	grizzly (gr-fe, CTS)	woodchuck	kodiak (ko-fe)	fog (CTS)	snow (CTS)	dtn (xfer nodes)	
LANL Yellow	trinitite (tt-fey)	capulin (Cray xc50)						
LANL Red	trinity	fire (CTS)	ice (CTS)	cyclone (CTS)			redcap (xfer)	
SNL OHPC	testbeds	mayer (testbed Astra)	solo					
SNL IHPC	ihpc.sandia.gov							
SNL SRN	mutrino, vortex	lynx	Sky Bridge	Chama	Uno	Ghost	Serrrano	stria (unclass Astra)
SNL SCN		lynx-s	Pecos	Jemez	Cayenne	Astra		

Gateway Name	Owning Institution
wtrw	LANL
ihpc-gate	LANL
red-wtrw	LANL
ihpc.sandia.gov	Sandia
srngate	Sandia

RCE Authentication Working Group

- realvnc functionality was tested and validated from SNL to LLNL
- LANL configured testing for access without red-wtrw in the secure
  - SNL to LANL had no issues
  - LLNL to LANL has some issues (solved during 10.17 testing)!
- SNL is standing up a git repo with CI in iHPC
  - we will need common auth for this!
- LANL is working on plan for common-auth by reorganization of resources and working to coordinate between non-HPC groups to provide the needed mechanisms.
  - Many of the resources (such as the KDC) currently live outside of the LANL HPC network making this a challenge to change/update.
- LLNL is looking at multiple options for common auth with SNL/LANL
  - currently have 6 options - trying to narrow down to best one(s)
  - will provide us with updates
- We all agreed on some sort of MOU Tri-Lab Agreement document
  - Catherine Hinton will lead drafting this document

Communication

- RCE combines SMEs in the know, people with budget authority, various levels of interested management, and people who know people across the three institutions
- SMEs are best able to identify problem areas and low-hanging fruit
  - Data transfer tools lack of ubiquity
  - Network bottlenecks
  - Gateways and what could be done to *dissolve* them
- Seek to identify projects that our in our normal charter (and thus already budgeted)
- Some efforts *do* require an influx of money
- Management creates priorities and supports these cooperative efforts
- Check-ins with users (either unsolicited on our part or user-driven)
- RCE ultimately achieves its best success by shedding light in certain areas and enabling the conversation

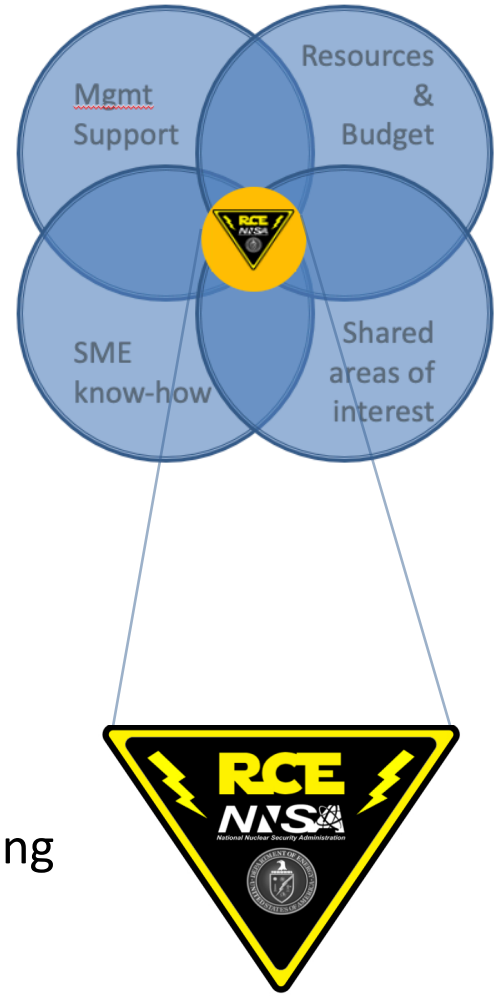

GitLab CI

https://lc.llnl.gov/confluence/display/GITLAB/GitLab+CI