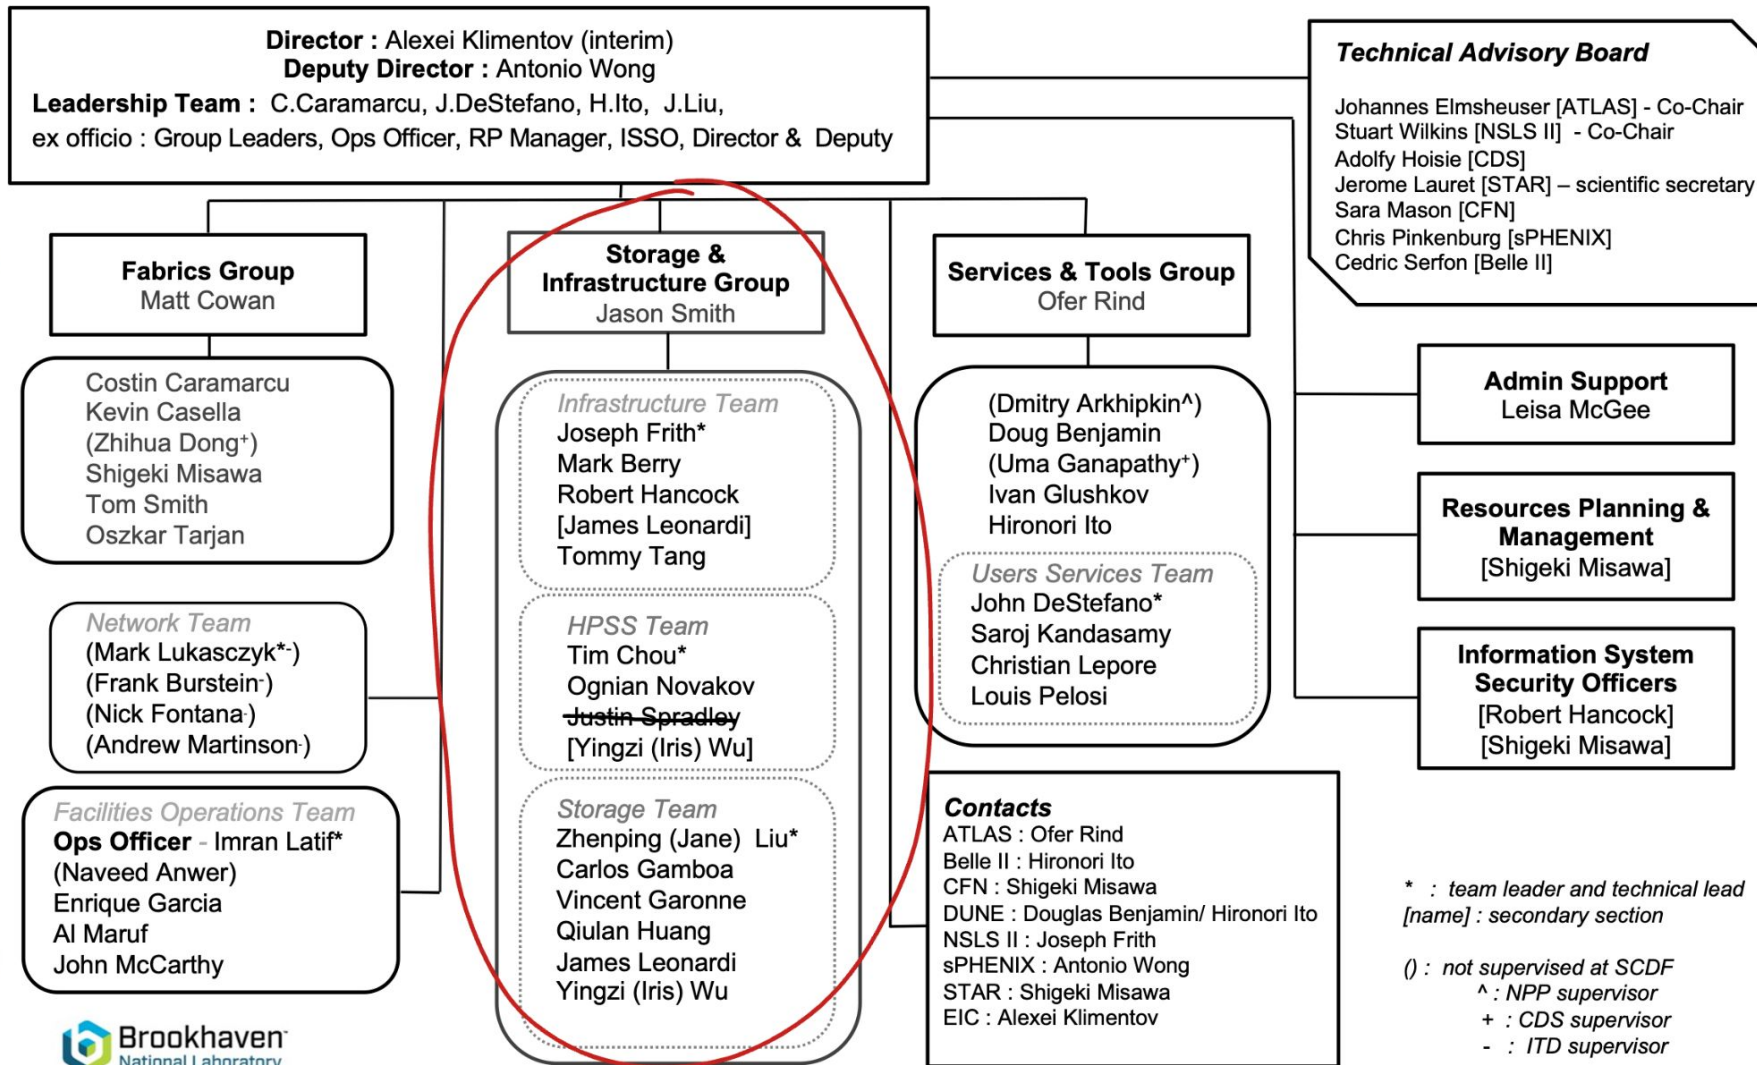# Storage & Infrastructure Group

SCDF Stakeholders and Users Technical Meeting

June 12, 2025

Jason Smith (Group Leader)

@BrookhavenLab

Storage & Infrastructure Group

Director : Alexei Klimentov (interim)
Deputy Director : Antonio Wong
Leadership Team : C.Caramarcu, J.DeStefano, H.Ito, J.Liu,
ex officio : Group Leaders, Ops Officer, RP Manager, ISSO, Director & Deputy

Technical Advisory Board

Johannes Elmsheuser [ATLAS] - Co-Chair
Stuart Wilkins [NSLS II] - Co-Chair
Adolfy Hoisie [CDS]
Jerome Lauret [STAR] – scientific secretary
Sara Mason [CFN]
Chris Pinkenburg [sPHENIX]
Cedric Serfon [Belle II]

Fabrics Group
Matt Cowan

Costin Caramarcu
Kevin Casella
(Zhihua Dong+)
Shigeki Misawa
Tom Smith
Oszkar Tarjan

Network Team
(Mark Lukasczyk*-)
(Frank Burstein-)
(Nick Fontana-)
(Andrew Martinson-)

Facilities Operations Team
Ops Officer - Imran Latif*
(Naveed Anwer)
Enrique Garcia
Al Maruf
John McCarthy

Storage & Infrastructure Group
Jason Smith

Infrastructure Team
Joseph Frith*
Mark Berry
Robert Hancock
[James Leonardi]
Tommy Tang

HPSS Team
Tim Chou*
Ognian Novakov
~~Justin Spradley~~
[Yingzi (Iris) Wu]

Storage Team
Zhenping (Jane) Liu*
Carlos Gamboa
Vincent Garonne
Qiulan Huang
James Leonardi
Yingzi (Iris) Wu

Services & Tools Group
Ofer Rind

(Dmitry Arkhipkin^)
Doug Benjamin
(Uma Ganapathy+)
Ivan Glushkov
Hironori Ito

Users Services Team
John DeStefano*
Saroj Kandasamy
Christian Lepore
Louis Pelosi

Admin Support
Leisa McGee

Resources Planning & Management
[Shigeki Misawa]

Information System Security Officers
[Robert Hancock]
[Shigeki Misawa]

Contacts
ATLAS : Ofer Rind
Belle II : Hironori Ito
CFN : Shigeki Misawa
DUNE : Douglas Benjamin/ Hironori Ito
NSLS II : Joseph Frith
sPHENIX : Antonio Wong
STAR : Shigeki Misawa
EIC : Alexei Klimentov

* : team leader and technical lead
[name] : secondary section

() : not supervised at SCDF
^ : NPP supervisor
+ : CDS supervisor
- : ITD supervisor

May 1, 2025. v 1.42

Brookhaven National Laboratory

Infrastructure Team

- Joe Frith - Technical Lead, NSLS2 PoC, GPFS, NFS, RHEV/OpenShift , Globus & various services
- Mark Berry - GPFS, NFS, RHEV/OpenShift & various services
- Robert Hancock - ISSO, OpenShift, XRootD, COManage & IAA, Storage hardware support, ZFS & many more things…
- James Leonardi - Config Mgmt. (Puppet, Foreman), Gitea, etc.
- Tommy Tang - OpenShift, Storage hardware support, ZFS, Globus, etc.

Brookhaven
National Laboratory

Infrastructure Services

- RHEV/OpenShift for virtualization & container management
  - RHEV (to be retired ~ 1 year) - 3 Hypervisors hosting ~90 VMs
  - OpenShift: 13 Workers hosting ~480 VMS & 39 containers (early testing)
- Central Storage
  - NFS (home dirs, RHEV/OpenShift, CVMFS, NSLS2, OKD)
  - GPFS (3.6PB STAR, 3.9PB PHENIX, 1.1PB USATLAS, 2.4PB HPC)
  - XRootD (STAR): 20PB on 18 servers
  - AFS (RHIC & ATLAS) *to be retired - hopefully
- Central Config Management
  - Puppet, Foreman, Gitea, Cobbler (to be replaced), GLPI
- Various Services
  - IPA (ldap, kerberos), syslog, dns, ntp, ssh, sftp, RT, Globus, Satellite RACF Org, etc…

Brookhaven
National Laboratory

Storage Team

- ## HPSS
  - Tim Chou - Technical Lead, HPSS, TSM
  - Ognian Novakov - HPSS, Libraries & Robots
  - Iris Wu - HPSS, Software & Monitoring
- ## Storage (distributed)
  - Jane Liu - Technical Lead, Lustre, Storage Monitoring
  - Carlos Gamboa - USATLAS Tier1 Storage Manager, dCache, Storage Monitoring & Performance, DB Management,
  - Vincent Garonne - dCache, Storage Monitoring
  - Qiulan Huang - dCache, Lustre, ARLAS Analysis Facility, Storage Monitoring
  - James Leonardi - dCache
  - Iris Wu - dCache, Lustre

Brookhaven National Laboratory

Storage Services

- HPSS (High Performance Storage System)
  - Tape libraries, robots, data movers, gateways, etc…
  - Over 389PB of data in 16 libraries on more than 100k tapes
- TSM (Tivoli Storage Manager): tape backup system
- Lustre
  - sPHENIX: 83PB on 93 servers
  - NSLS2: 13PB on 24 servers
  - EIC/STAR: 7.6PB on 8 servers
  - CFN: 1.8PB on 2 servers
  - LQCD: 217TB on 3 servers
- dCache
  - ATLAS: 47.7PB data on disk (57 servers) & 104PB tape
  - BelleII: 3.7PB data on disk (4 servers) & 2.7PB tape
  - PHENIX: 7.8PB data on disk (13 servers) & 23PB tape
  - Dune: 1.3PB data on disk (4 servers)
  - Integration/Test: 700TB

Brookhaven National Laboratory