ESnet

ENERGY SCIENCES NETWORK

Network Support For Large Scale Science

Eli Dart

Energy Sciences Network (ESnet)

Lawrence Berkeley National Laboratory

SIG-NGN

Prague, CZ

20 April, 2023

Framing and Context

- ESnet is a science network first
  - We serve enterprise needs too (e.g. commodity/cloud transit)
  - Lab business operations important, science mission front of mind
- Multiple operational mechanisms to understand science we serve
  - Interaction with Lab network managers
  - Performance/design/troubleshooting engagements with scientists
  - Learning from our friends and peers
- Formal requirements program – programmatic foundation

ESnet

ESnet Requirements Program

- Part of ESnet's formal governance structure – DOE program effort
- Formal interaction between ESnet and the six DOE/SC program offices
- Regular interaction with each program office – every 3 years
- Collect written information from scientists, then discuss it with them
- Results in formal reports: www.es.net/requirements

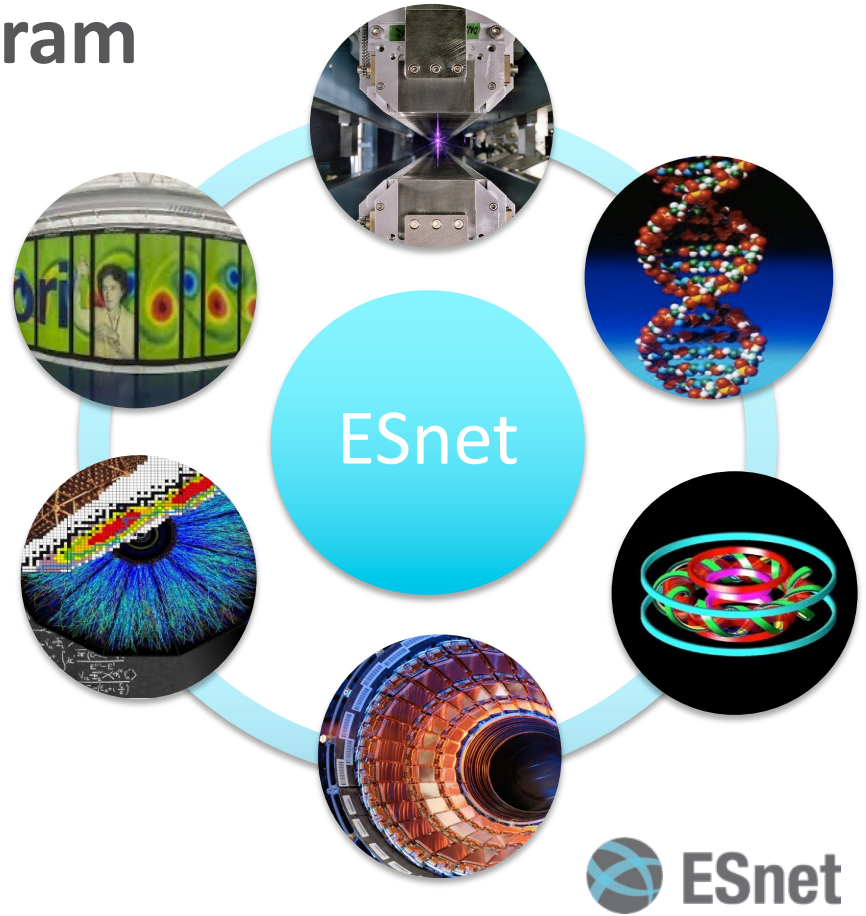

ESnet

ESnet Requirements Program Case Studies

- ESnet attempts to understand the facilities and major projects of each program office, primarily from a networking and data perspective
  - Derive network requirements from the conduct of the science
  - We collect this information in “case study” narratives
- Each science collaboration in the review prepares a case study containing a data-centric narrative describing their science, including quantitative information
  - Three time scales:
    - 0-2 years – current budget envelope
    - 2-5 years – current technology envelope
    - 5+ years – strategic planning

ESnet

Requirements Review Case Studies

Requirements case studies evaluate multiple aspects of science programs

- ***Major science experiments and facilities,*** both in operation and planned.
- ***The process of science*** used for knowledge discovery, and including scientists' interactions with the instruments and facilities.
- ***The volume of data*** produced now, and anticipated in the future, with an emphasis on geographical location of where the data must be shared, computed and/or stored.
- ***The current technology capabilities*** (network, computation, storage, and software stack) used by each science collaboration/facility as well as any planned upgrades, additions or improvements.

ESnet

Strategic Context

- Requirements program facilitates conversations between ESnet, our sites, our program management, and the other DOE program offices.
- ***We learn from each other***
- Not only do we gain insight into the data strategy of programs, we can teach programs and scientists about the value of ESnet
  - Best practice, design patterns (Science DMZ, DTNs, perfSONAR)
  - Benefits of using high-performance networking in effective ways

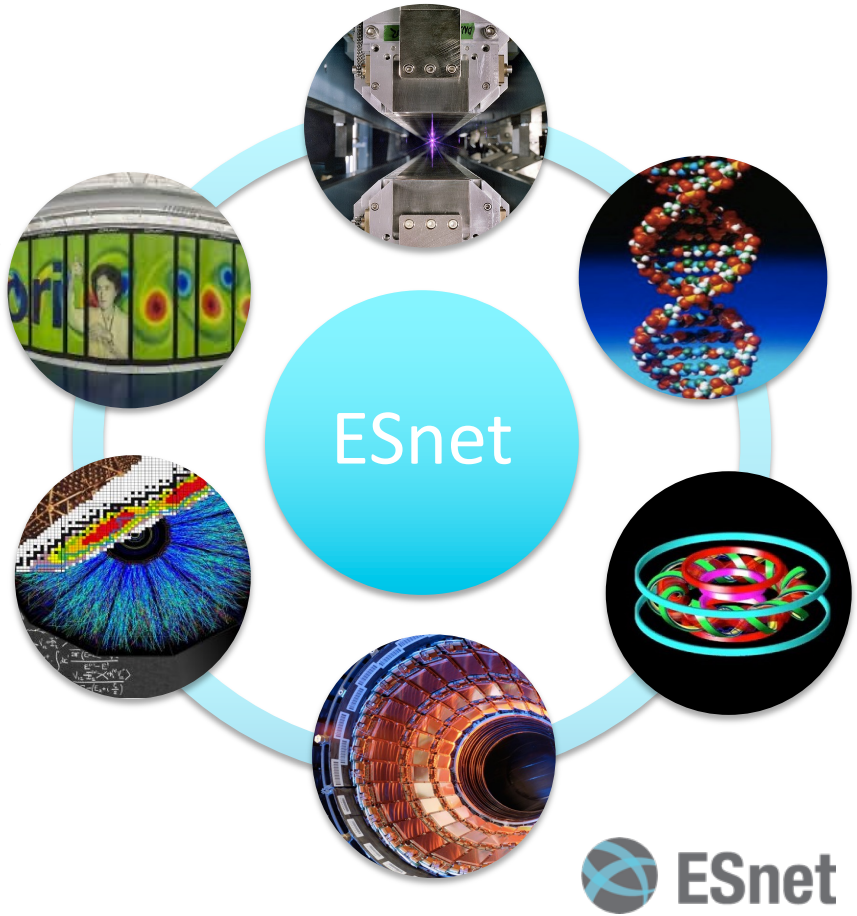

ESnet

Several Benefits To A Formal Program

- ESnet, our program managers, our science constituents, and their program managers are required to sit down and talk to each other
  - It is hard to overstate the value of these interactions
  - Everyone in the same room at the same time: common discussion of needs, and the solutions ESnet undertakes to meet those needs
- The “sociology” of different science collaborations, facilities, and fields is super-important
  - Astronomy is different from Fusion which is different from Genomics, ….
  - For us to be successful we need to meet the scientists where they are, which is different for each field and facility, sometimes for each collaboration

ESnet

Rough User Grouping By Data Set Size

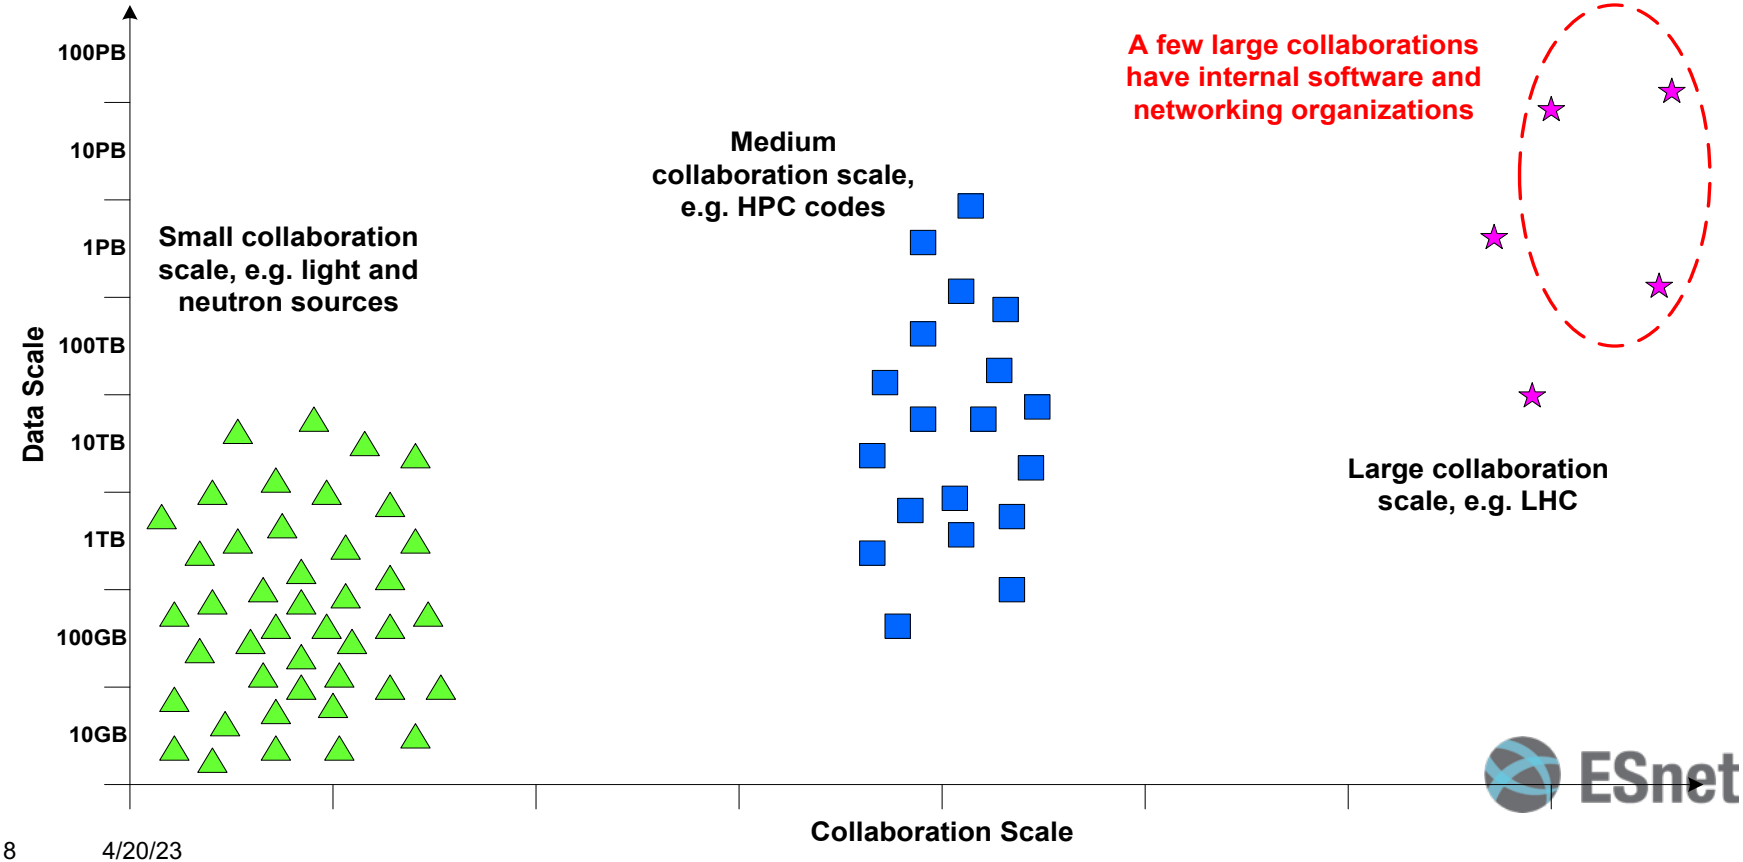

Rough User Grouping By Data Set Size

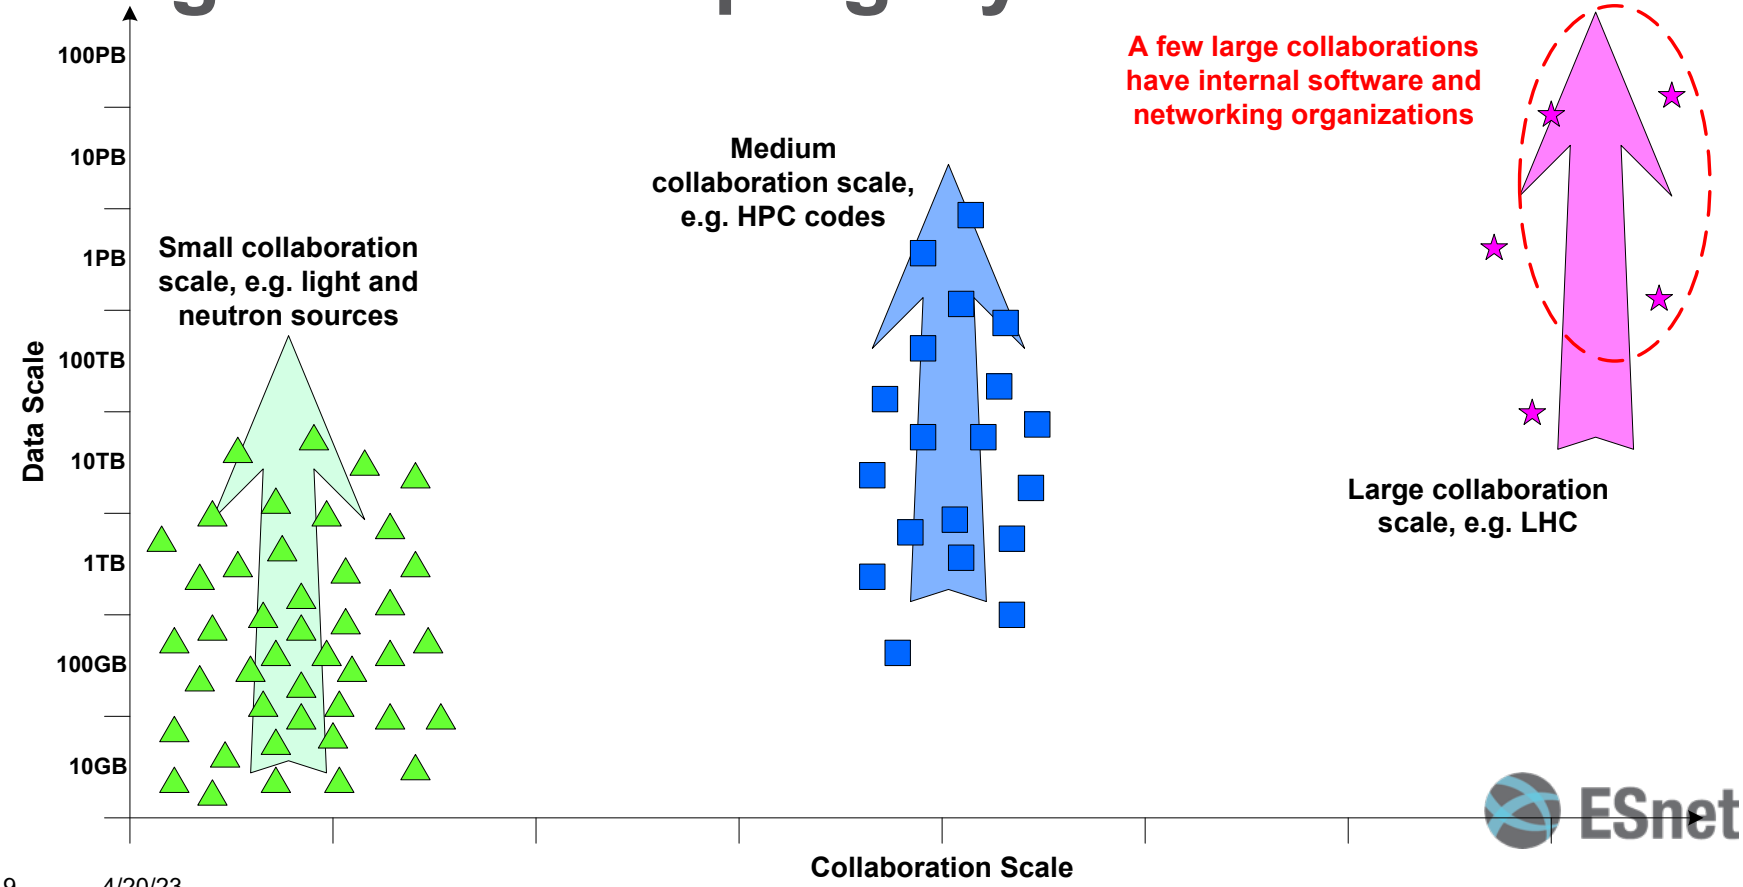

Rough User Grouping Discussion (1)

- The chart is a crude generalization
  - It is not meant to describe specific collaborations, but to illustrate some common socio-technical aspects of many collaborations
  - Data sets are constantly growing (growth arrows on second slide)
- Small collaborations
  - Light sources, microscopy, nanoscience centers, etc.
  - Typically small number of scientists per collaboration, many many collaborations
  - Individual collaborations typically rely on site support and grad students
  - This group typically has difficulty moving data via the network
  - Science DMZs and Data Transfer Nodes (especially if deployed with Globus) are a big win, multi-facility workflow efforts are building infrastructure

Rough User Grouping Discussion (2)

- Supercomputer simulation science
  - Climate, fusion, bioinformatics, astrophysics simulations, etc.
  - Larger collaborations, often multi-site
  - Reliant on supercomputer center staff for help with network issues, or on grad students
  - This group typically has difficulty transferring data via the network, unless DTNs have been deployed properly (e.g. Science DMZ + Globus)
- Large data instrument science (HEP, NP)
  - Very large collaborations – multi-institution, multi-nation-state
  - Collaborations have their own software and networking shops
  - Typically able to use the network well, in some cases expert (LHC)
  - These groups often deploy their own stack (Rucio, FTS, HTTP-TPC, XRootD, etc) instead of working for interoperability with many different science collaborations

ESnet

So What? What Have We Learned?

- The LHC experiments are big (obviously!)
  - But what does that mean?
- Almost every non-LHC collaboration needs our help, but in different ways
  - Some need is about understanding networks
  - Some need is about understanding systems
  - Some need is other stuff
- Science collaborations will not abstract their workflows for us
- So let's unpack this…

ESnet

The LHC Experiments

- ESnet needs (and has) a multi-faceted, strategic engagement effort with the LHC experiments
  - Started before Run 1, in place today, continuing for the future
  - Ongoing participation in collaborations/groups where it makes sense, proactive action where needed, ability to change/adapt over time
  - Examples: SENSE/Rucio collaboration, Tier2 engagement in the US, strategic transatlantic connectivity program, perfSONAR
- It's important to spend the people cycles to be part of the collaborations. Otherwise, networking doesn't have a seat at the table, and everyone loses

ESnet

Non-LHC Collaborations and Experiments

- Many non-LHC science groups do not have in-house networking expertise
  - Importantly, the senior staff who do strategic planning for experiments or for a field of research do not have networking expertise in the general case
    - It's "IT" or "Cloud" to them
- Many non-LHC groups do not have in-house expertise with high performance data systems (DTNs, etc.)
  - This is why ESnet has the fasterdata knowledgebase (could it be better? Always! We welcome contributions)
- There is a huge opportunity to help our constituents make better use of our networks

ESnet

ITER

- Global collaboration building the world's largest Tokamak in France
  - Member states: China, EU, India, Japan, Korea, Russia, USA
  - Pursuing connectivity to GEANT, so data likely accessible via GEANT
- Data model – different aspects
  - ITER has said there will be no interaction with plasma control via external networks
  - Exabyte data scale by 2035
  - Every member state is entitled to a full copy of all the data
    - 2PB/day capacity per member state
    - 3 months of disk buffer, then data migrated to tape (get it while it's hot!)
    - Transfers via GEANT
    - Transfers starting approx. 2028, fast ramp to 2030
- Data challenge activities beginning later this year
- Schedule – update coming soon

ESnet

Climate Science

- Note well: this is *not* weather! Climate and Weather are different
- Large (many petabytes) of climate model output distributed worldwide
  - Major data repositories in US, UK, France, Germany, Australia, Japan, China
  - 20-30 additional sites
- They need to do periodic transfers to sync their big repositories
- Sometimes they transfer large data sets (e.g. to HPC)

ESnet

Cloud

- Cloud is a capability play in many cases
  - Easy API access
  - Always-on or always-available services (e.g. Globus controller)
  - Burst capacity/elasticity (though expensive)
- Our community needs to think hard about what services and capabilities it wishes to keep in-house, especially in terms of computing
  - Ease of use is a primary driver (compute allocation application delay, API vs. CLI, etc.)
  - Security issues – many orgs find it easier to write a contract for cloud access than defend their own systems against cyberattack: is this what we want?
- Performance is relevant to topology location, but ease of use/access need not be.
- **Our scientists just want to get their work done** – what does that mean for us?
  - Note that the default over time is that science moves to the cloud if it's easier

ESnet

In conclusion – ESnet's mission:

Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.

ESnet

ESnet
ENERGY SCIENCES NETWORK

Thanks!

Eli Dart (dart@es.net)
Energy Sciences Network (ESnet)
Lawrence Berkeley National Laboratory

http://fasterdata.es.net/
http://my.es.net/
http://www.es.net/