SC20

Atlanta, GA | more than hpc.

SCCompSC

The SC20 Student Cluster Competition, and How to Join it

November 15th-20th • Atlanta, GA

Housekeeping

- Q&A session at end - please use the link above slides to ask questions
- We're recording this session, the recording and these slides will be posted on the webinars page
- The zoom password is "cluster"

The SC20 Student Cluster Competition

"SC" is The International Conference for High Performance Computing, Networking, Storage, and Analysis.

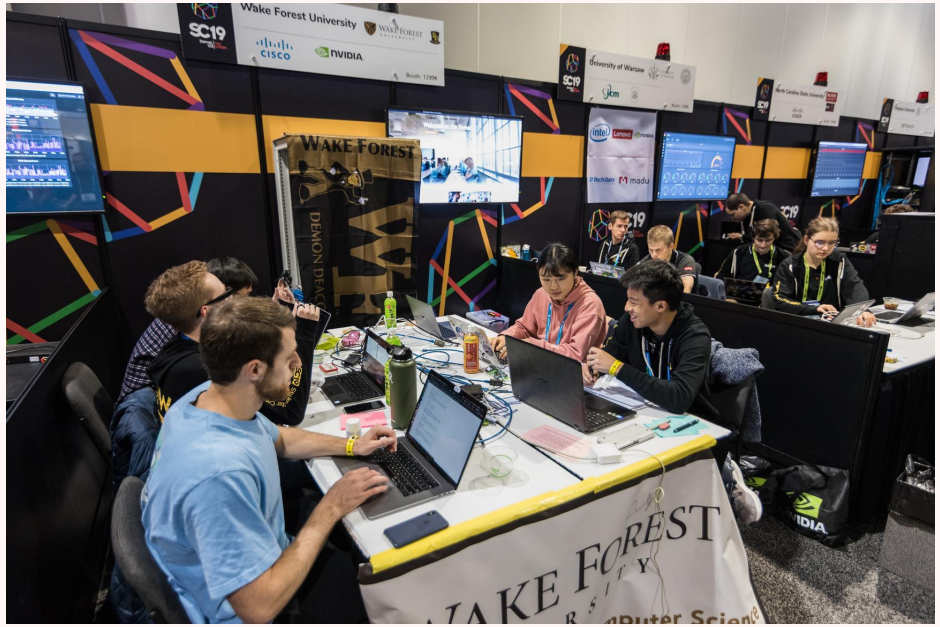

November 15–20, 2020
Georgia World Congress Center - Atlanta, GA

#MoreThanHPC

SC20
Atlanta, GA | more than hpc.

What is a Student Cluster Competition?

HPC is one of the **best tools in existence** for science and engineering

The SCC fosters skills development and social connections to bring new people into HPC

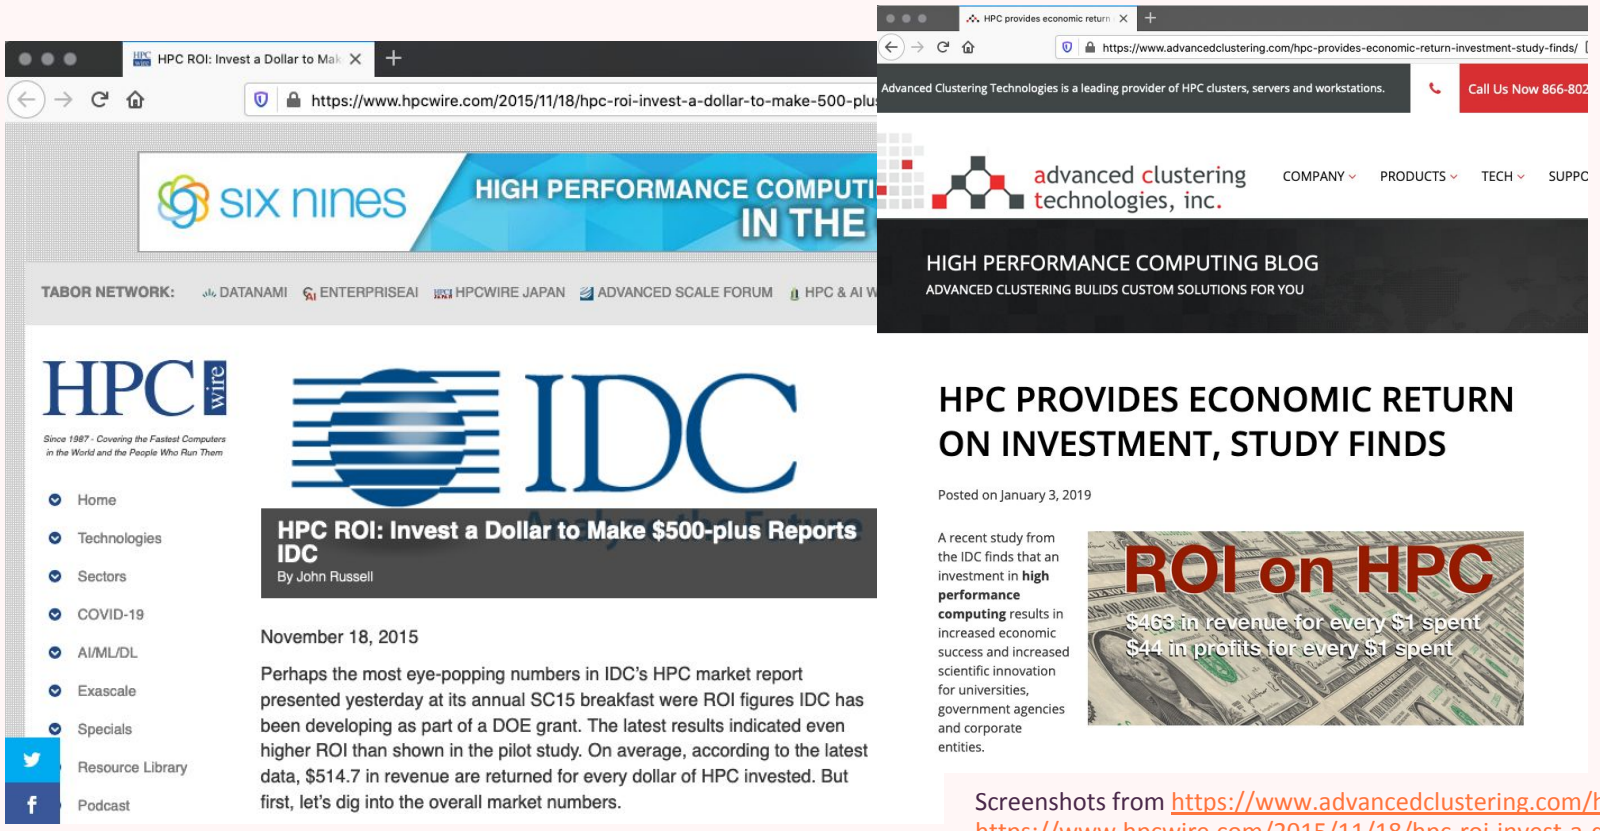

Screenshots from https://www.advancedclustering.com/hpc-provides-economic-return-investment-study-finds/ and https://www.hpcwire.com/2015/11/18/hpc-roi-invest-a-dollar-to-make-500-plus-reports-idc/

Ok, but What is a Student Cluster Competition?

A 48-hour, non-stop contest to build and run a (small) supercomputer

Teams of 6 undergraduate[1,2] students:

- Design and build a HPC Cluster
- Measure and tune its performance
- Run real science workloads on your cluster and the cloud
- Recover from a disaster scenario
- Report on your results

All within a power budget

1. Team members must be enrolled at an educational institution, but not have been granted an undergraduate degree, at Nov 16, 2020.
2. Team members must be at least 18 years old by the start of the competition.

What's the process?

Preparation before the event:

- Teams form partnerships with institutions and vendors
- Design a cluster, practice building and running the applications
- Webinars to help you prepare
- Logistics of getting to the competition
    - SCC provides conference registration and lodging for 6 team members plus 1 advisor
    - We encourage institution and vendor partners to budget for other expenses
    - Some details shortly
- Build and install your cluster, make posters, start on reports

What's the process?

At SC (dedicated webinar to come):

- Teams arrive Friday/Saturday, safety briefing etc on Saturday
- Saturday PM -> Monday AM: Build and test the cluster, social events
- Monday AM: Benchmarking begins! No more advisor help :)
- Monday PM: Competition kickoff!
- 48 non-stop hours of competition including:
    - Running applications
    - Writing reports
    - Daily standups
    - Presenting posters and lightning talks
    - Disaster recovery (power outage)
    - Scoring interviews
    - Conference events
- Wednesday PM: Competition ends! Party night :)
- Thursday: Awards, wrap-up, cluster teardown

Some words from teams GeekPie HPC (ShanghaiTech University) and Wolfpack HPC (North Carolina State University) about the SC19 Student Cluster Competition

Members of team RACKlette from ETH Zürich are here with us today

- What was your most memorable moment in the SC19 competition?
- How much HPC experience did you / your team have before signing up last year?
- How has the competition impacted your career path?
- What advice would you give to someone thinking about joining the SCC?

Introducing the Committee

Verónica - co-chair

Scott - co-chair

Jason - vice-chair

Kathleen - deputy chair

Introducing the Committee

Rigo - infrastructure

Paul - infrastructure

Stephen - reproducibility

Rebecca - student activities, posters

- Abinhav - applications
- Stephen - infrastructure
- Ramin
- Angel
- Sam
- Brian
- Carlos
- Esteban

Andy - cloud

John

Jenna

Steve - webinars

How to apply

1. Bookmark this page:

https://sc20.supercomputing.org/program/studentssc/student-cluster-competition/

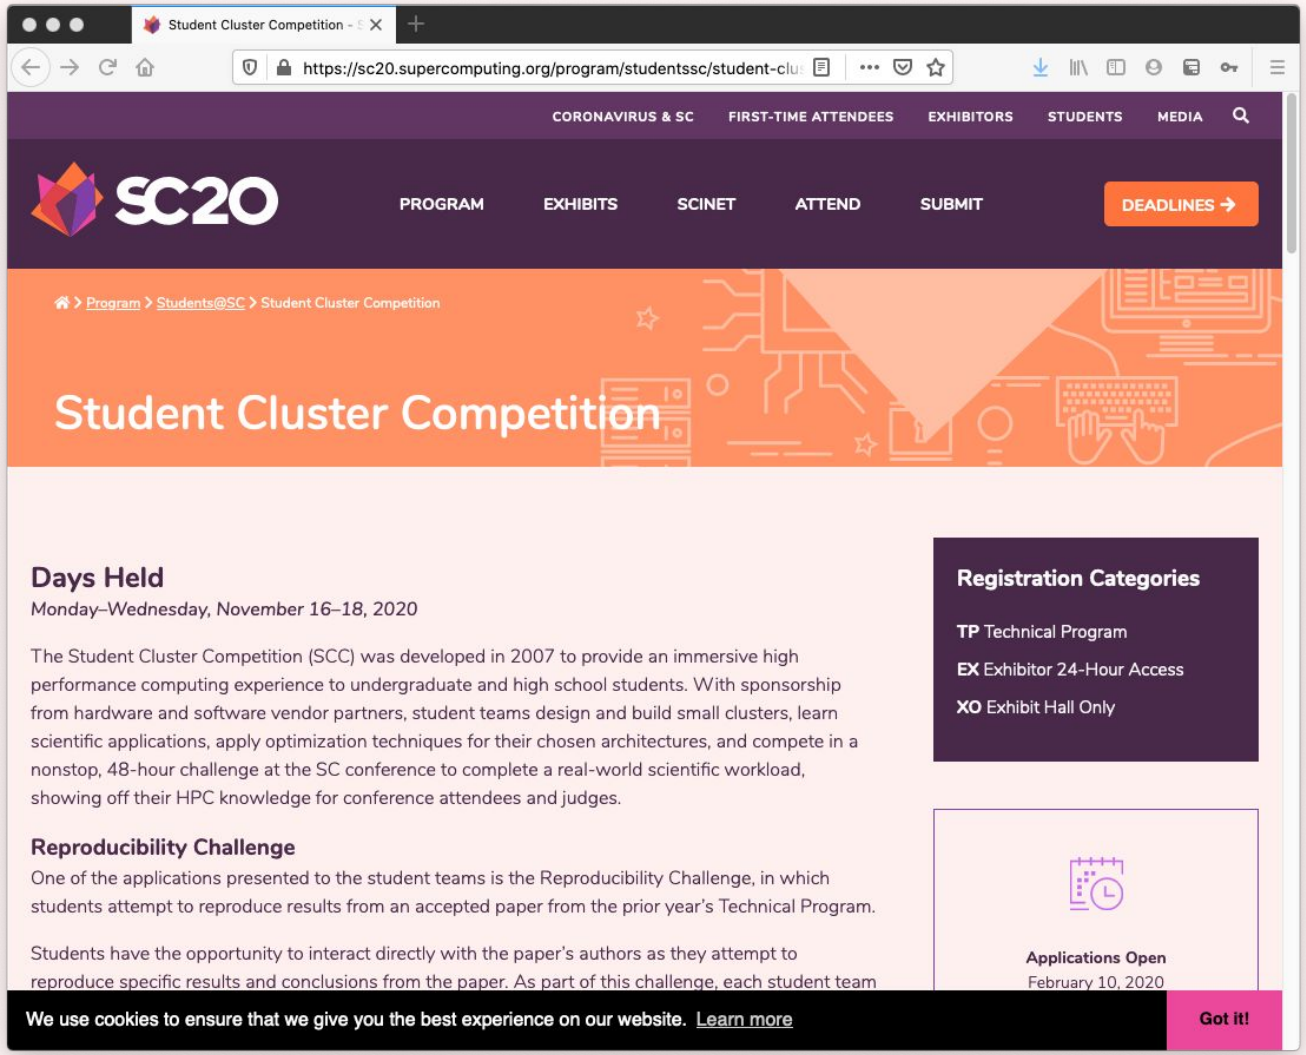

How to apply

2. Read the regulations

Key points:

- 6 undergrads + 1 advisor
- Age 18+ at time of event
- Sponsoring institution and/or vendor supplies the hardware and software
  - (can be a loan)
- Also shipping costs, team transportation, etc
- Hardware must be commercially available at time of event
- No NDAs

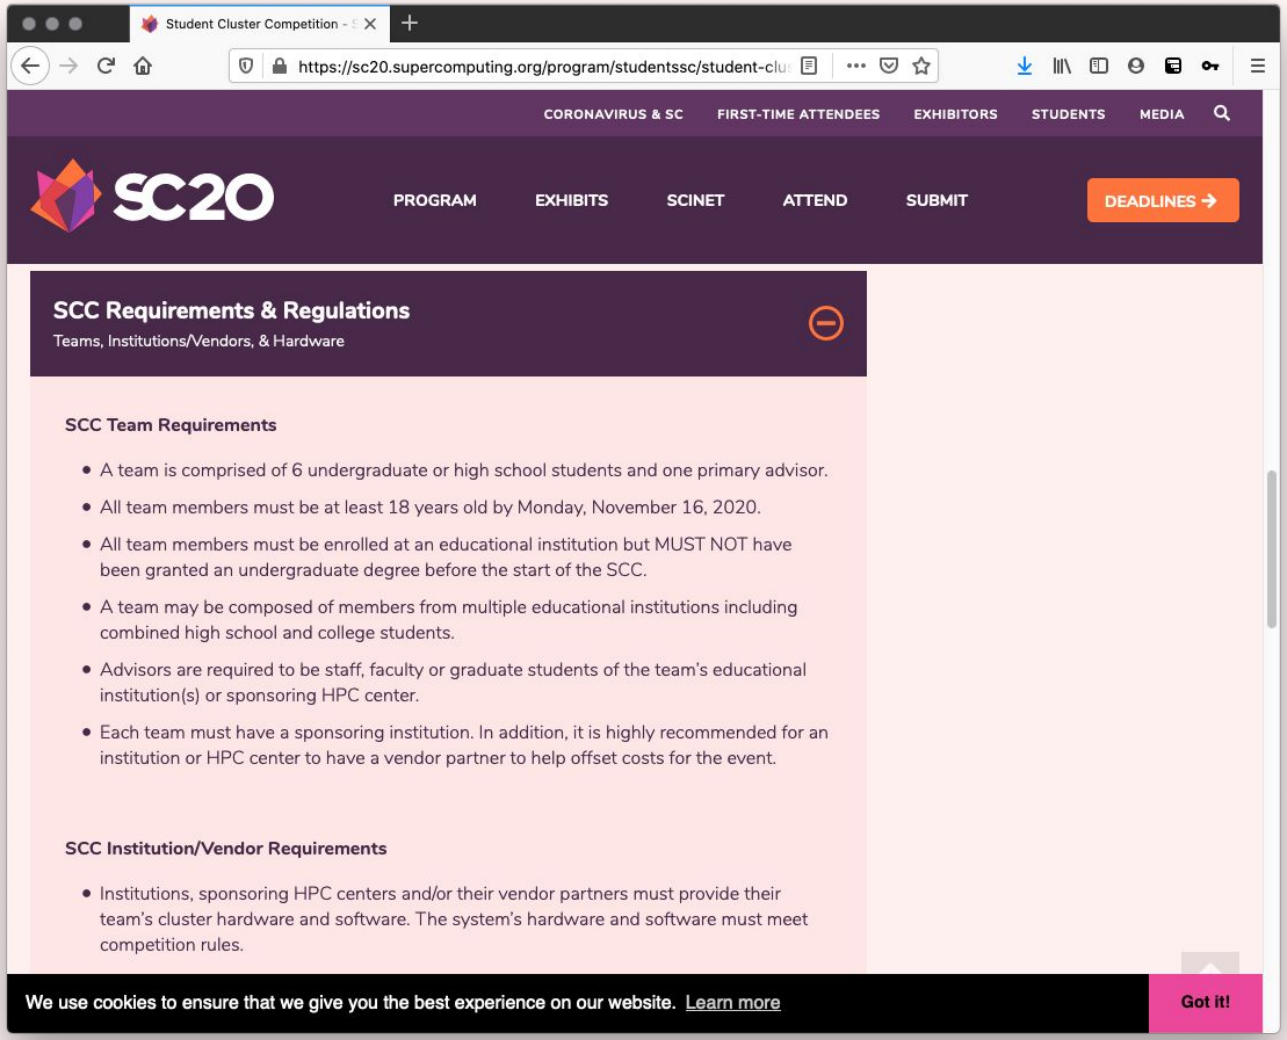

How to apply

3. Checkout these tips

.. we'll walk through them in a moment

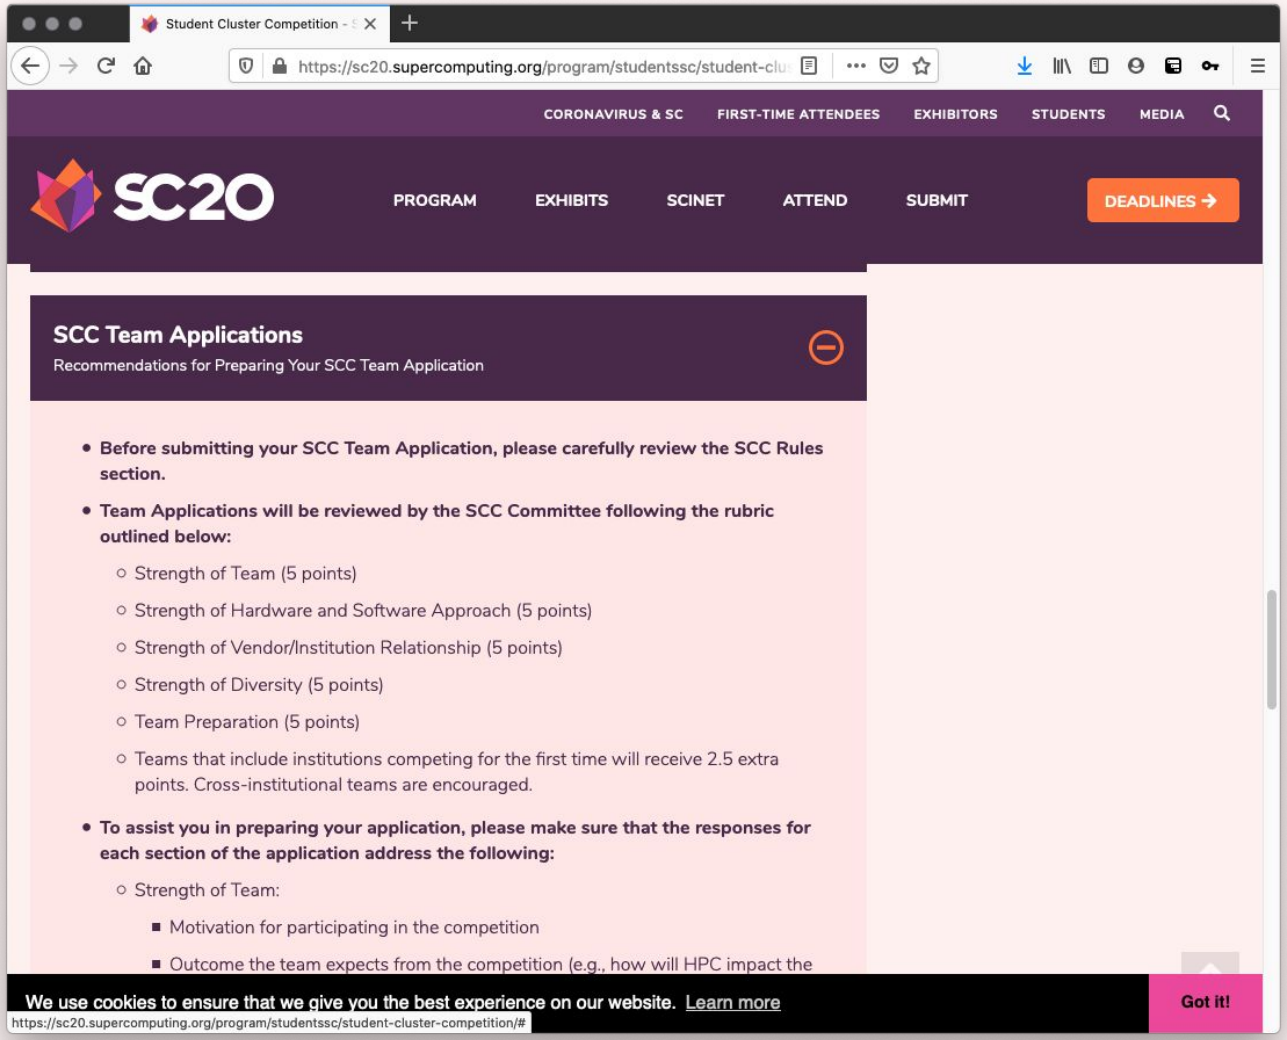

How to apply

4. Submit via this link

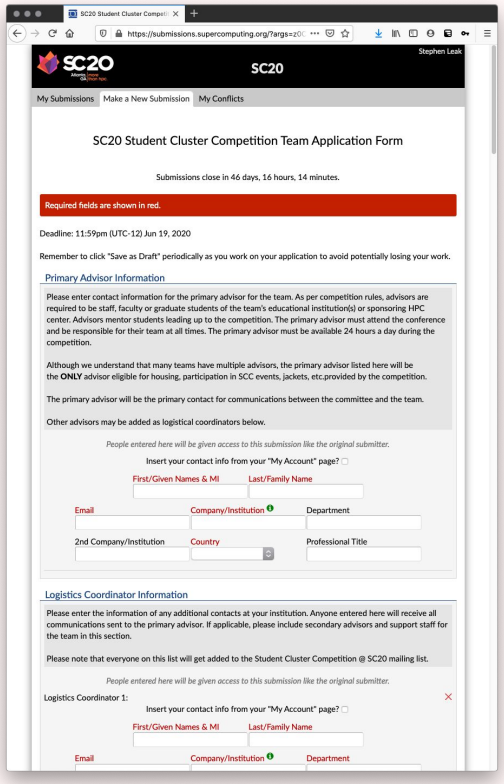

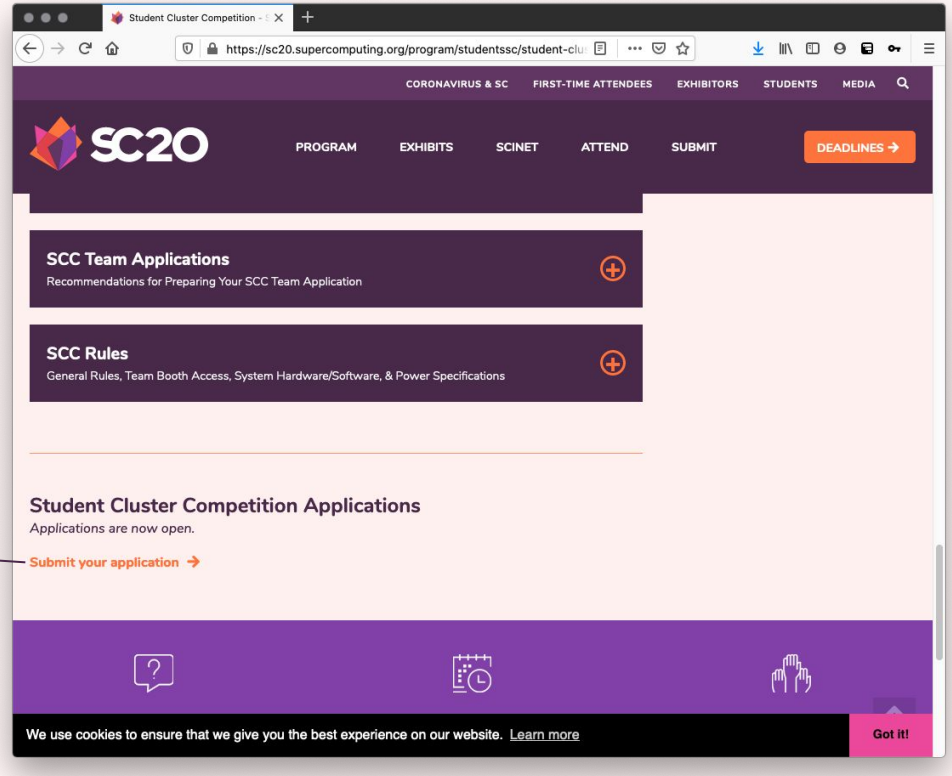

(or this one)

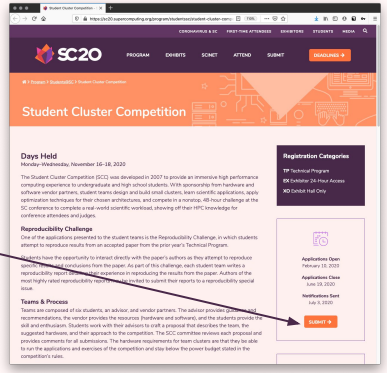

Preparing your application

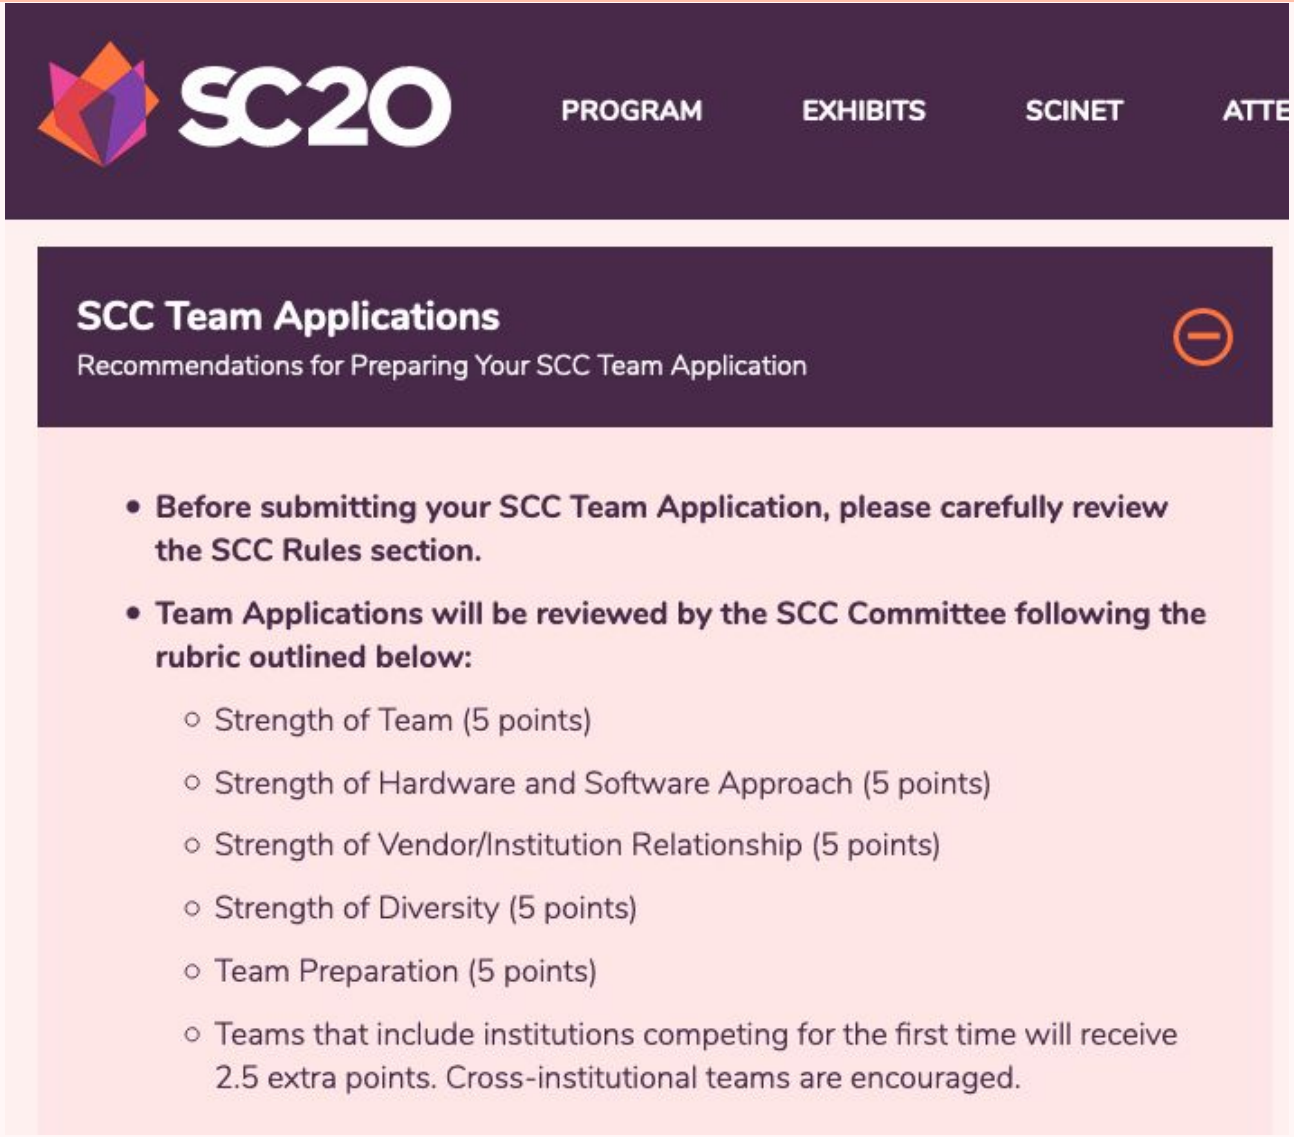

We get many great proposals each year, but we have limited booth space

Knowing what we're looking for can help you to write a winning Team Application

Tip: First-time teams get a bonus!

Preparing your application

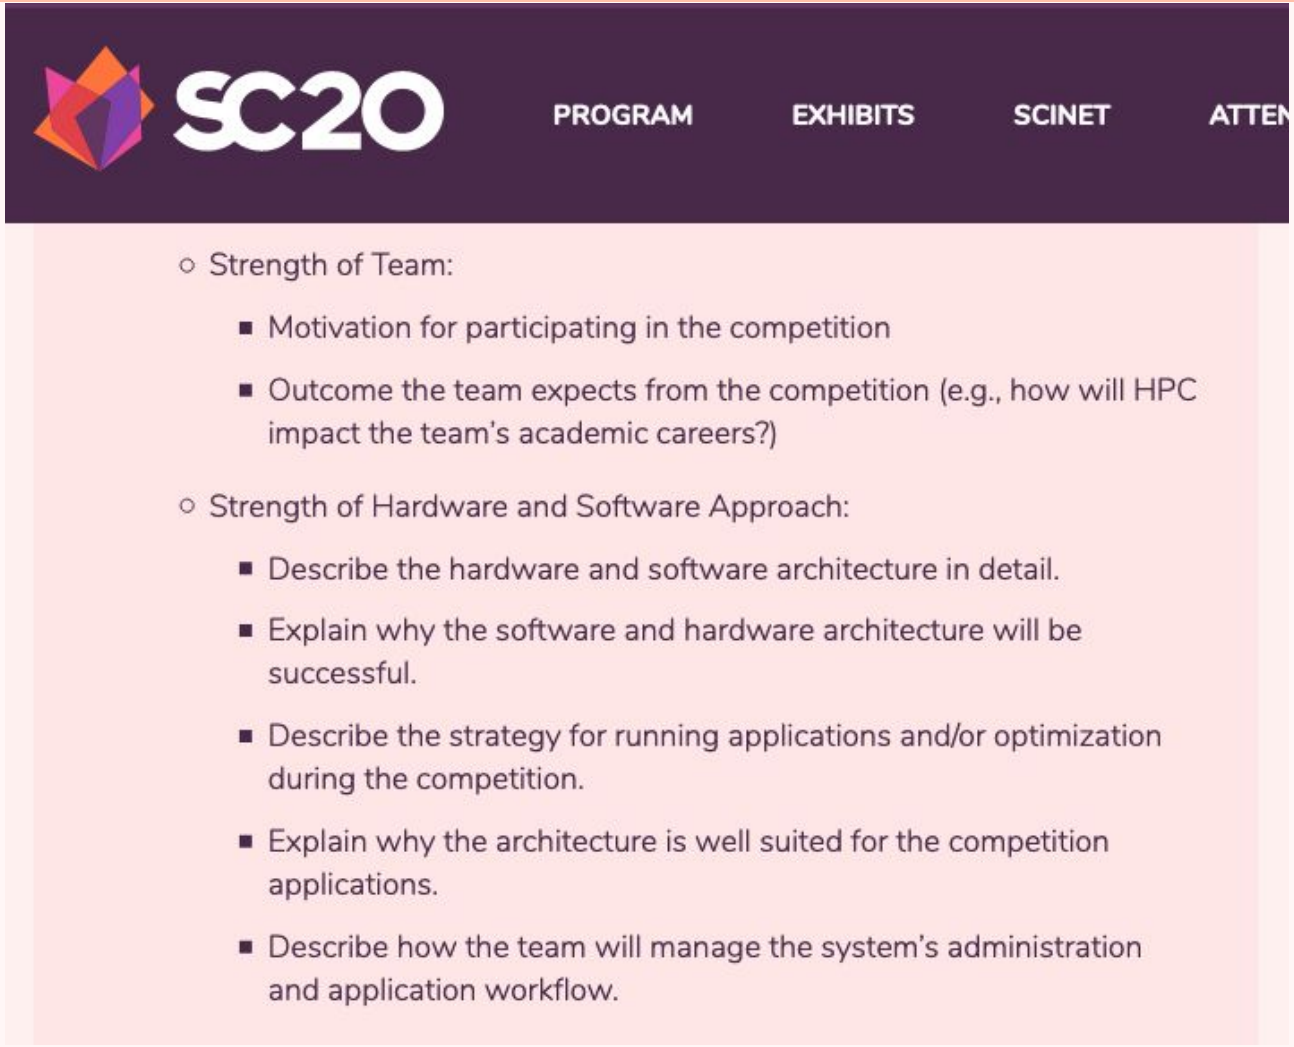

Strength of team

- We're looking for teams who will develop into great HPC professionals
- You don't need to be experts already
- Curious, collaborative, motivated

Hardware and Software

- Why are you choosing this hardware/software?
- Do some research on the applications to get a sense of what each one needs
- What tools will you learn to use? (Tip: ask your vendor partner!)

Preparing your application

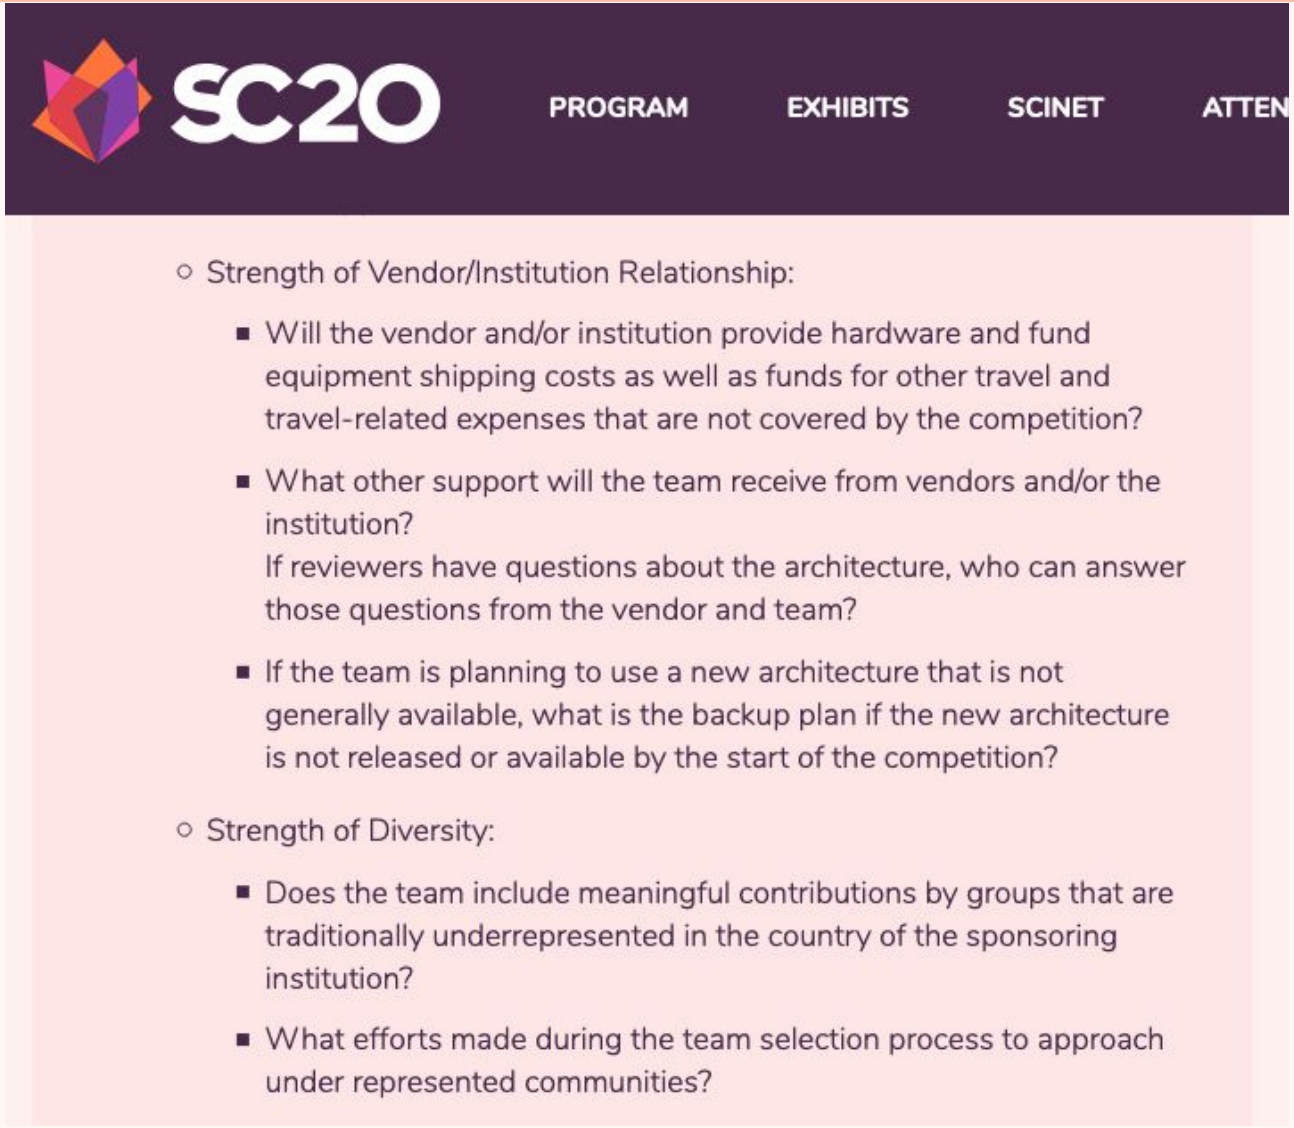

Vendor/Institution

- We want to know that you'll have sufficient support!
- And that your hardware will be available by the time of the competition
  - If you are using new/upcoming hardware, **have a plan B**

Diversity

- Many HPC professionals look and sound a lot like me
  - And we can contribute a lot .. **BUT**:
- We need more voices informed by different cultural, language, gender, career, etc backgrounds
  - So tell us what your team brings!

Preparing your application

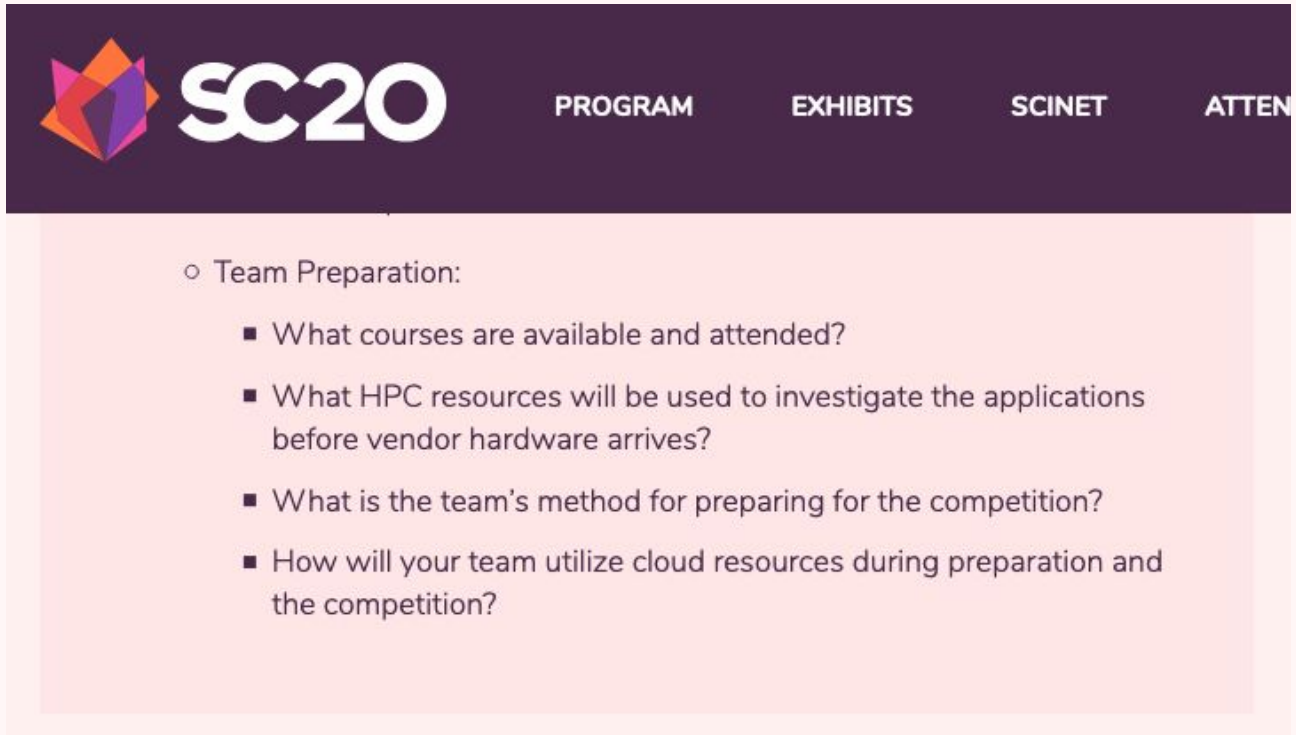

Preparation

- You can get from completely new to HPC rising star in time for the competition
- BUT it takes dedication and organization
- So tell us how you're preparing
  - What challenges do you foresee?
  - What is your plan for meeting those challenges?
  - Tip: collaboration with your institutions, vendors and each other will greatly increase your resources

Hardware and Software: Things to know

- New power budget: **4500W**
- Power specifications (including PDU model) are under "SCC Rules" on the SCC main page

- There are 3 benchmarks:
  - Linpack (HPL) (Shows peak performance)
  - HPCG (More realistic, tougher performance benchmark)
  - IO500 (You will need storage, and a parallel filesystem)
- We will have a webinar about benchmarks, watch the webinars page for announcements

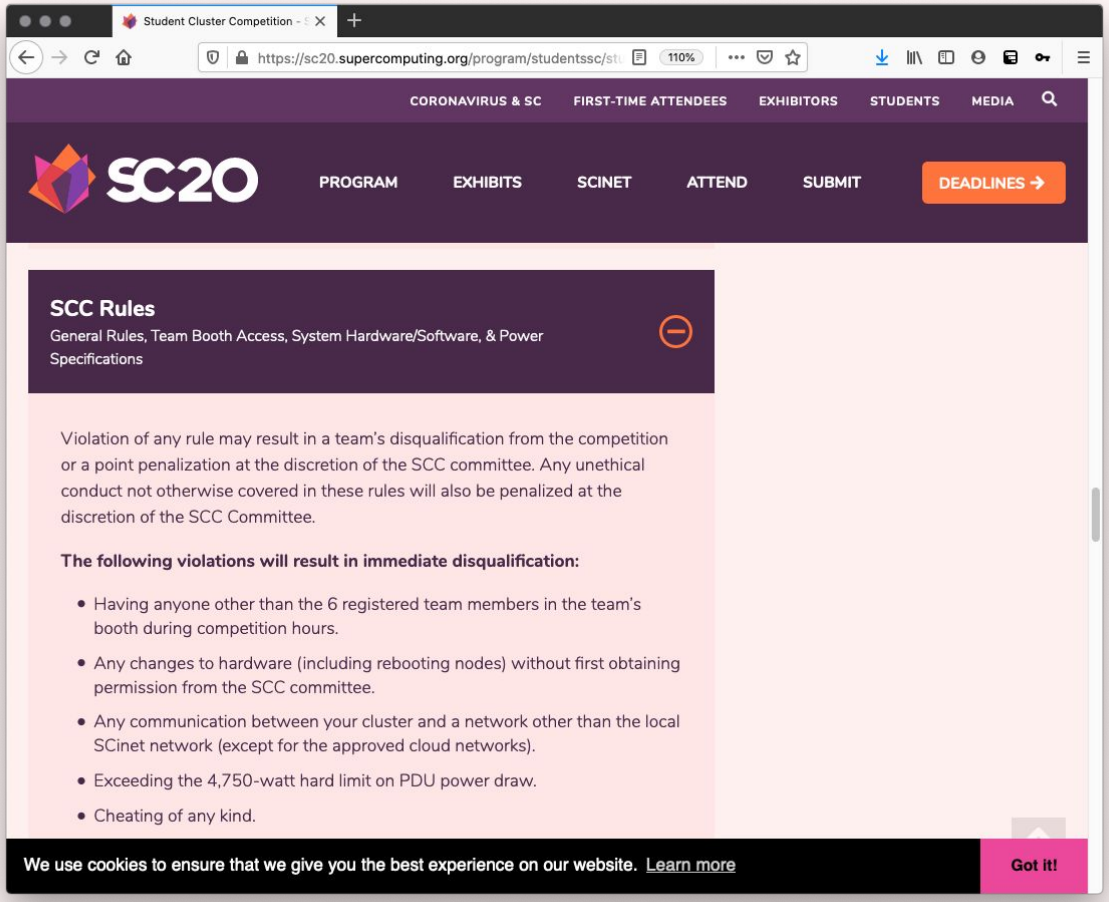

Hardware and Software: Things to know

- This year's applications are posted on on SCC page
- We aim to have webinars about each
- CESM - Parallel climate model.
  - Complex and challenging to optimize
  - Used in IPCC reports
- GROMACS - Molecular dynamics
  - Fast, vectorizes well.
  - Used in COVID-19 research, as well as materials science
- Reproducibility Challenge
  - Reproduce some results from the SC19 paper MemXCT: memory-centric X-ray CT reconstruction with massive parallelization
- Mystery Application
  - Announced when the event begins!

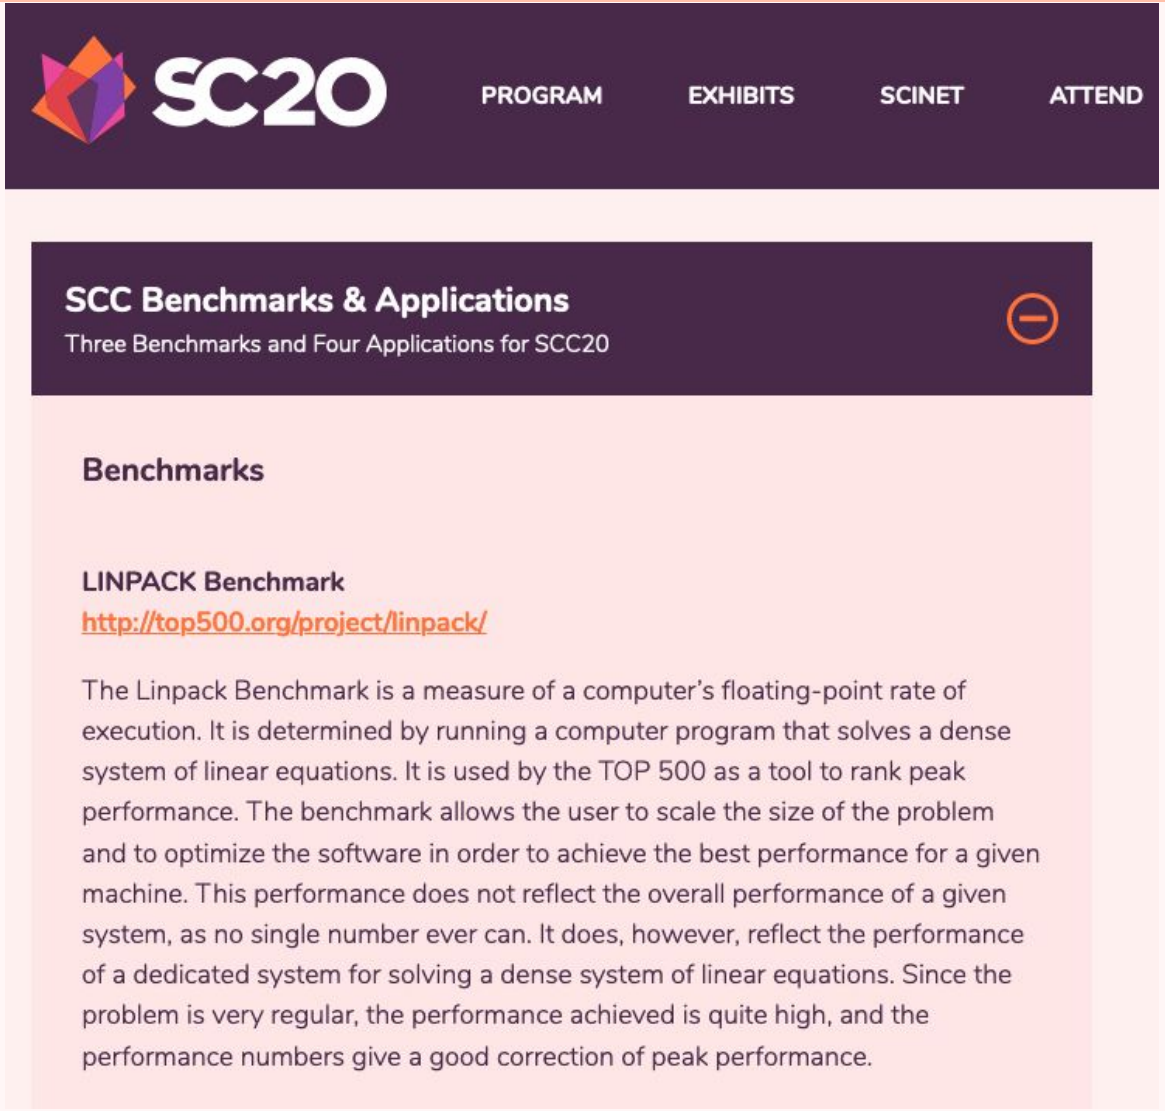

Getting help

- Teams will get access to a Google group
  - Committee, application experts and other teams will be on the group - great place to get help and help each other
- Team liaisons
  - Each team will be assigned a liaison from the committee
  - Your liaison will be able to guide you for questions about the ev logistics, etc
- Webinars:
  - https://sc20.supercomputing.org/program/studentssc/studentssc-webinars/
  - We'll have a series of webinars like this, about once per month, on topics to help you prepare
- Blog: https://sc20.supercomputing.org/attend/blog/
  - https://sc20.supercomputing.org/2020/05/01/how-to-prepare-for-a-student-cluster-competition/
- Websites
  - https://studentclustercompetition.us/index.html
  - https://sc20.supercomputing.org/program/studentssc/student-cluster-competition/

Coronavirus & SC

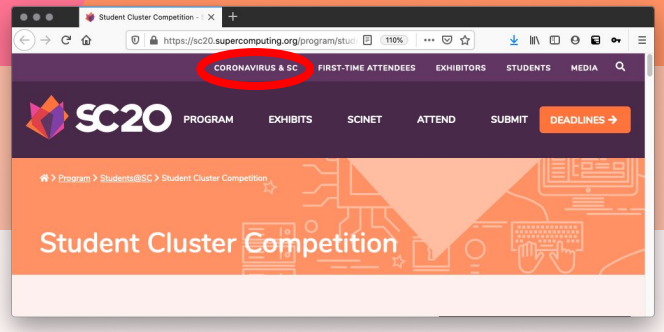

The health and safety of the SC20 family—participants and volunteers alike—are our first priority. We recognize the impact that COVID-19 is having, and will have, on international travel in particular. *We are closely monitoring any specific travel advisories through official channels, which include the World Health Organization (WHO) and the U.S. State Department.*

The SC20 committee is investigating options for limited remote participation for the conference in November 2020 should the need arise due to travel restrictions associated with COVID-19.

We will make a broad announcement if remote participation becomes available for the conference. We continue to monitor the situation closely, but at this point are hopeful for normal conference operations.

Important Dates

Team applications close: June 19, 2020

Notifications will be sent July 3, 2020

Competition: Monday–Wednesday, November 16–18, 2020

Teams will arrive on the weekend to set up

Webinar schedule:
https://sc20.supercomputing.org/program/studentssc/studentssc-webinars/

Apply Now & Follow Us!

@SCCompSC

SC20 SCC Program details and team application information:
https://sc20.supercomputing.org/program/studentssc/student-cluster-competition/

SC20 SCC Mystery Application:
https://sc20.supercomputing.org/program/studentssc/student-cluster-competition/

SCC History:
https://studentclustercompetition.us/

StudentsSC

SC20

Atlanta, GA | more than hpc.

Questions?