EasyBuild @ ITER

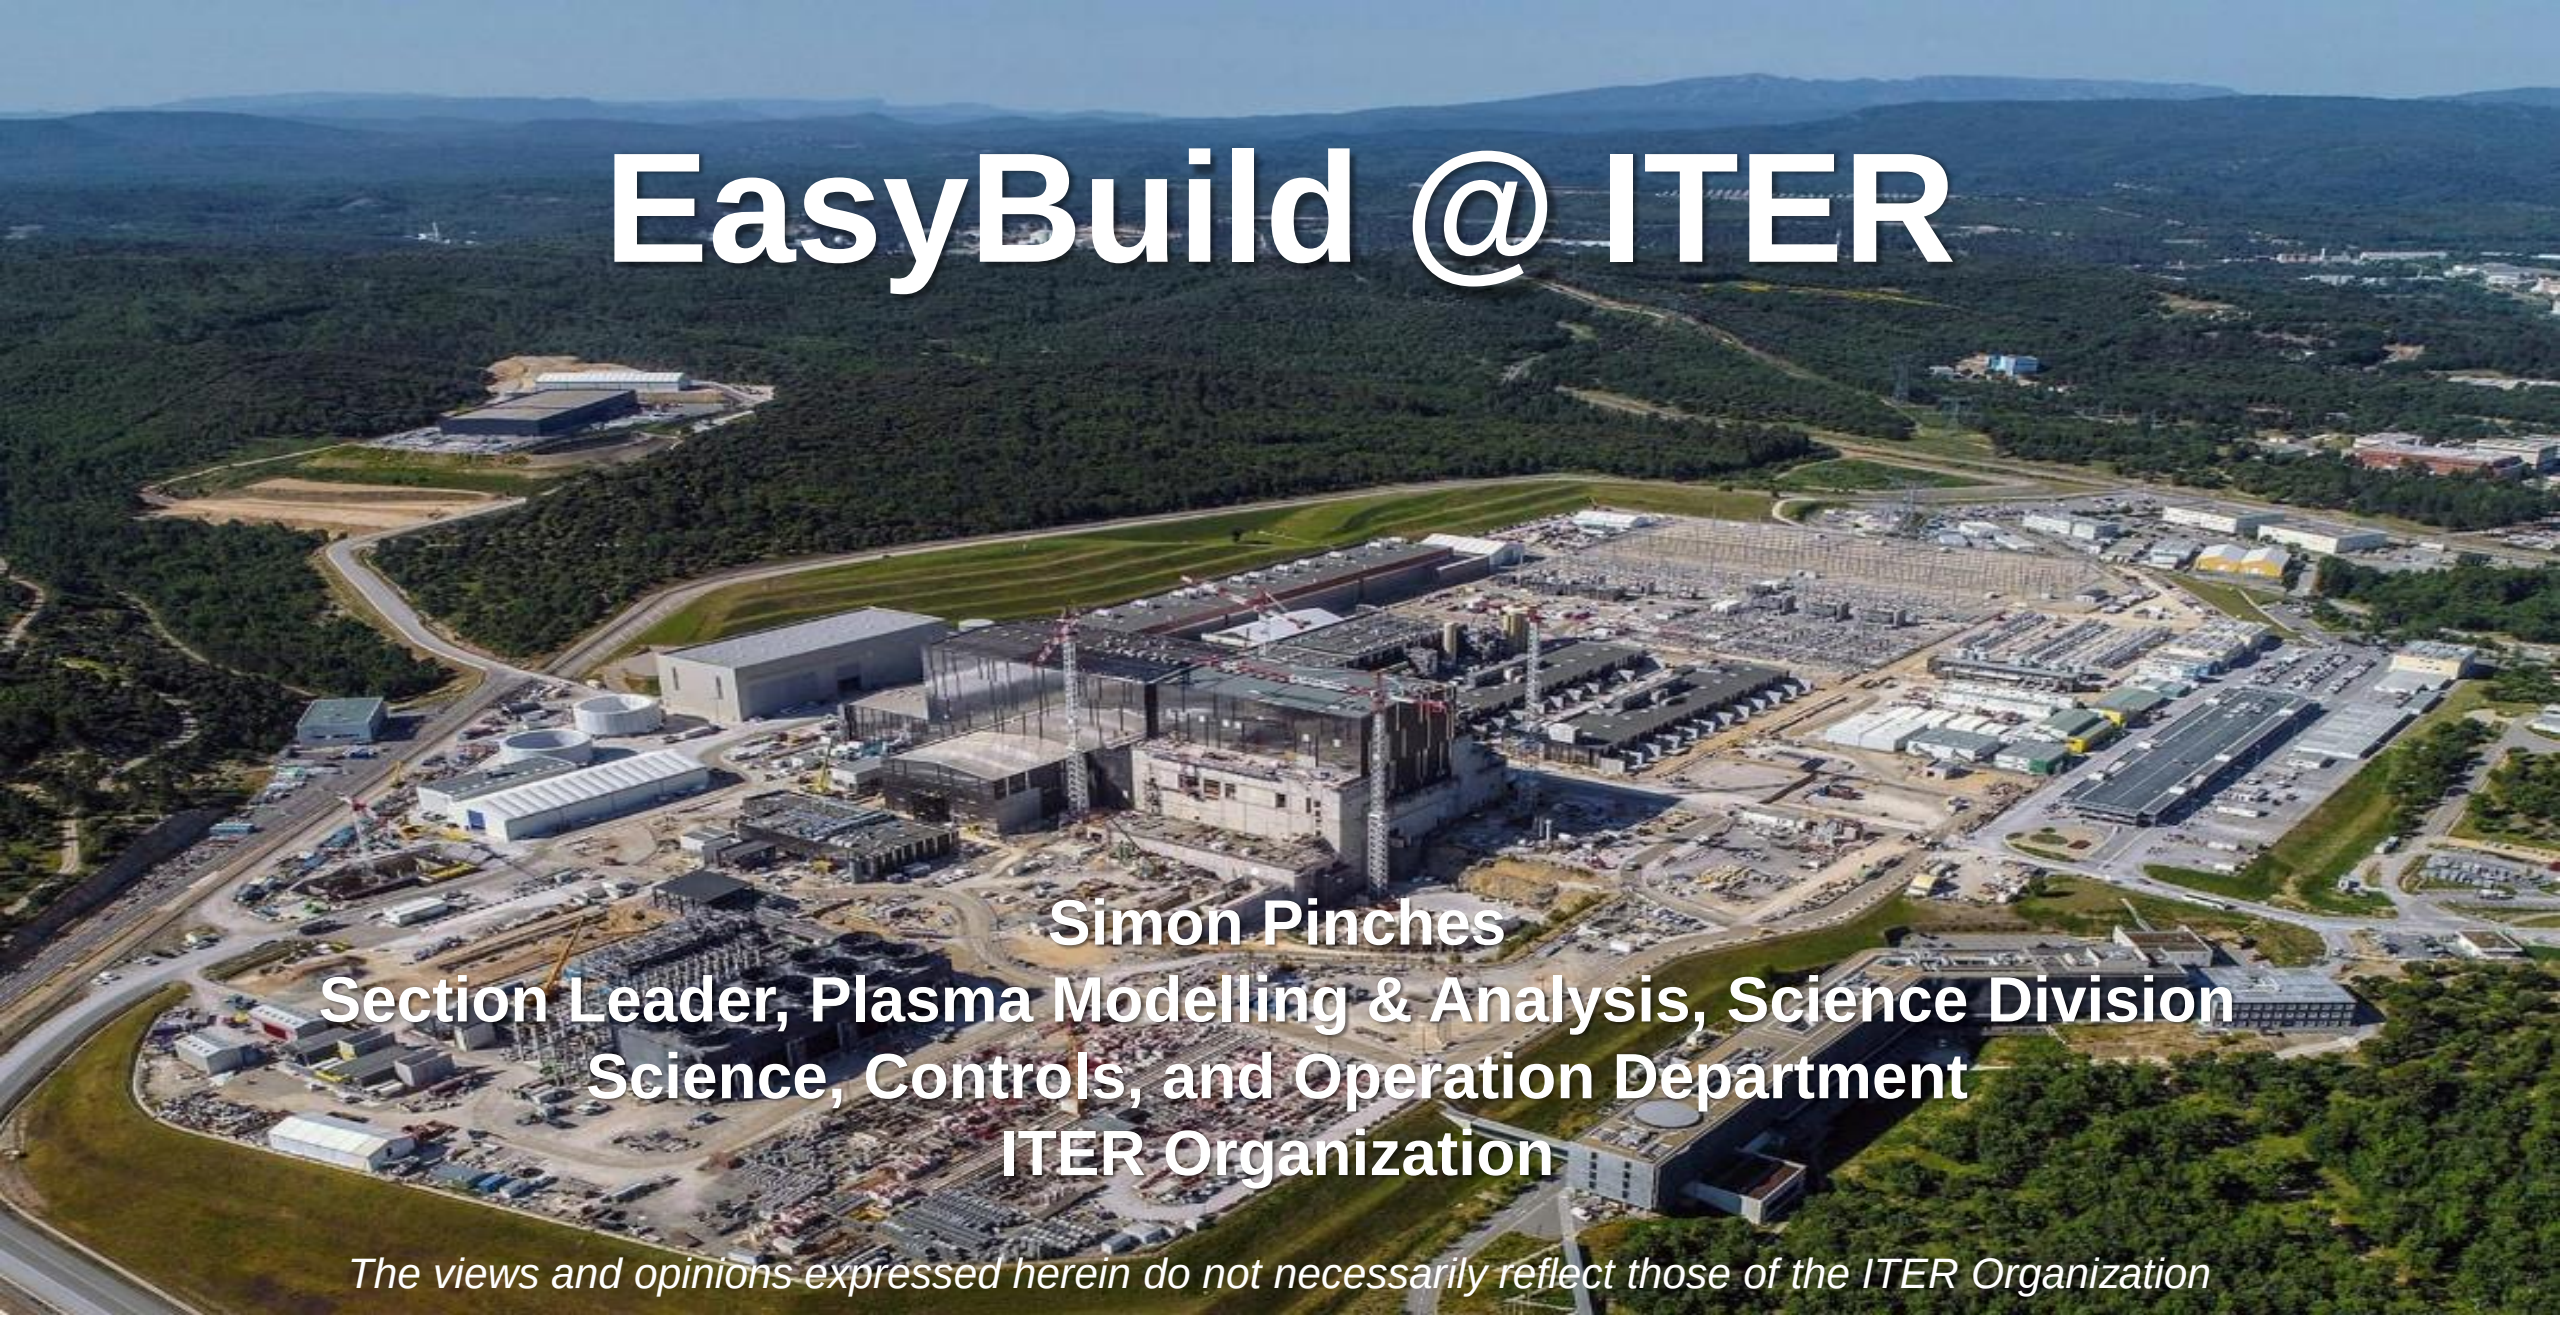

Simon Pinches
Section Leader, Plasma Modelling & Analysis, Science Division
Science, Controls, and Operation Department
ITER Organization

The views and opinions expressed herein do not necessarily reflect those of the ITER Organization

ITER

- ITER Organization
  - One of worlds largest scientific projects
  - Seven ITER Members who contribute in-kind
    - China, EU, India, Japan, Korea, Russian Federation and USA
  - Aims to demonstrate the scientific and technological feasibility of using magnetically confined fusion as a future energy source
  - Mission
    - Fusion energy gain, Q ≥ 10 for 300 – 500 s
    - Long pulse and steady-state operation, Q ≥ 5 for 1000s and 3000s
  - Machine assembly on-going, first plasma mid-2020s

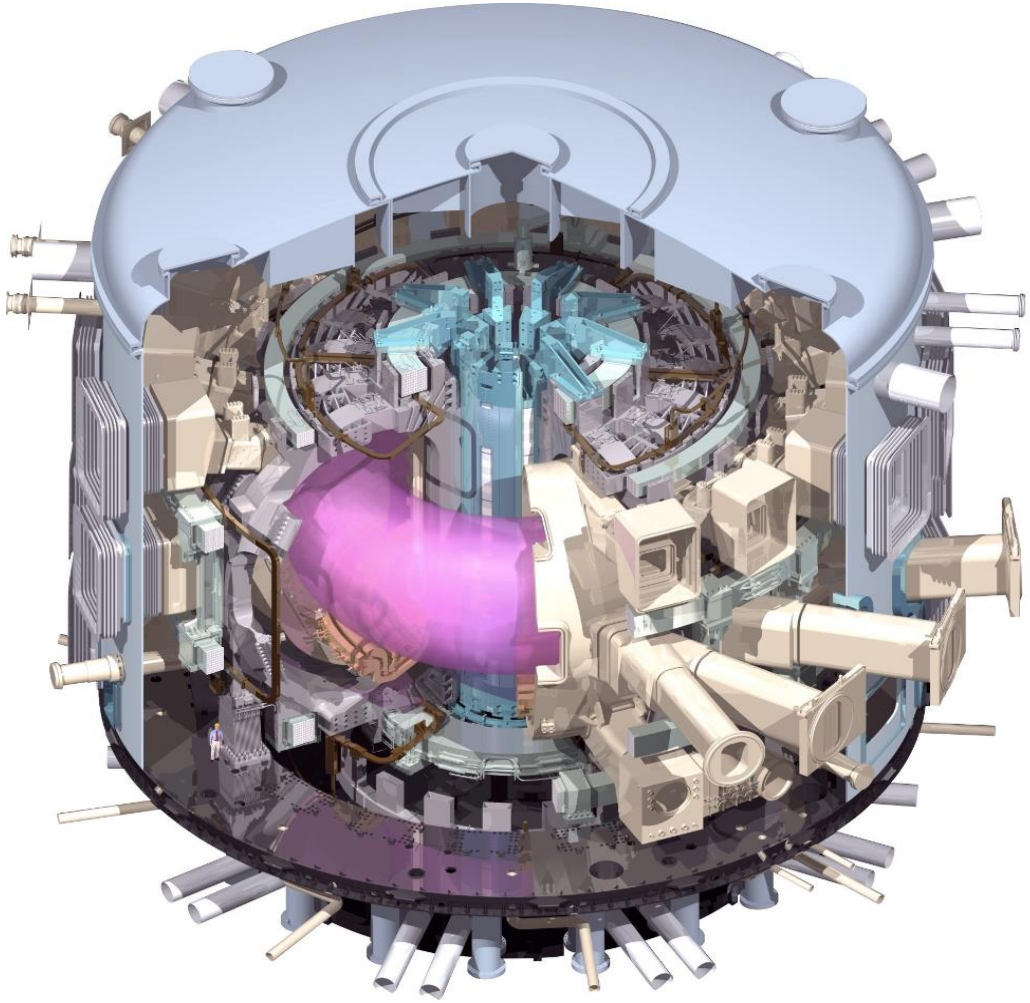

Accompanying Integrated Modelling Programme

- To prepare all scientific software to predict ITER performance, support on-going design refinement and prepare for data processing and interpretation
- Builds upon expertise and knowledge (often captured in software!) from within ITER Members
- Initial focus has been on preparing a High-Fidelity Plasma Simulator
  - Couples codes by enforcing data exchange based around a standard Data Model
  - An infrastructure (part of Integrated Modelling & Analysis Suite, IMAS) provides an access layer and a component generator for use with workflow engine (e.g. Kepler) or development environment (e.g. Python)
  - Up until now, the entire IMAS infrastructure has been provided to users through a single installer (supporting all languages) and giving rise to one IMAS environment module
- Development of data processing and analysis pipelines starting

Present Data and Computing Capabilities

- Cluster
  - 235 nodes with ~7000 cores
  - Infiniband, SLURM
  - Modules-Tcl v3.2.10
    - May move to v4 or Lmod
- GPGPU server
  - 8 x P100 NVIDIA cards
- Remote Access
  - 3 x NoMachine servers
- Shared storage
  - GPFS
  - Total storage currently ~2 PB
  - Total experimental data:
    - 0.1 – 1 PB for First Plasma → Exabyte scale during Fusion Power Operation (mid-2030s)
    - 30 – 50+ GB/sec
    - Up to 2.2 PB/day

Heterogeneous environment

- Cluster has already gone through various expansions (and will continue to do so)
- EB builds are currently optimized for lowest common instruction set:

```
optarch=Intel:axAVX,CORE-AVX2;GCC:march=sandybridge
```

  - Looking into creating separate mounts providing same software but optimized for each architecture (and perhaps also OS as we migrate from CentOS 7 to 8)

GPU server

- Primarily used for Monte-Carlo power load calculations
- Written for PGI CUDA Fortran compiler
  - Last year, PGI compilers rebranded as part of NVIDIA HPC SDK (and now free to download)
  - New EB files would be good… ☺

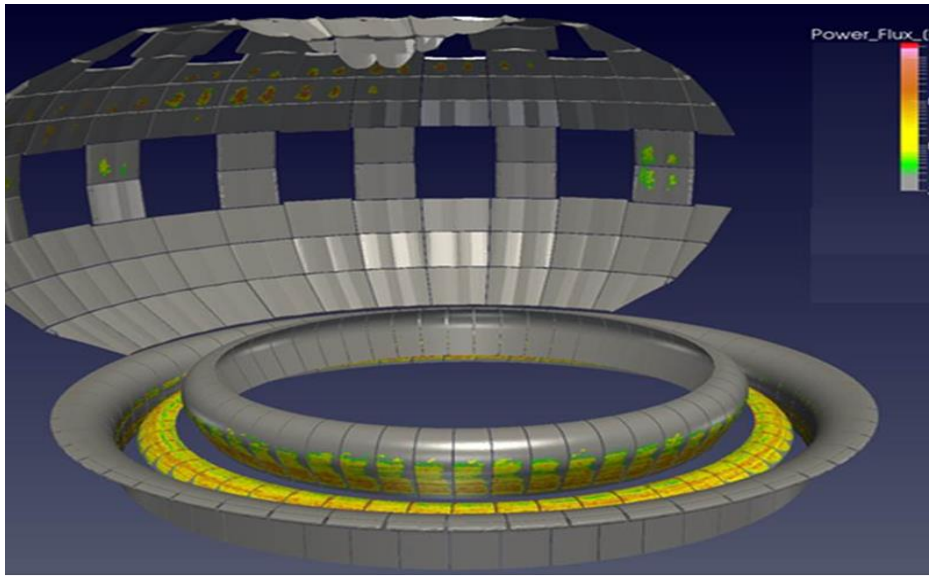

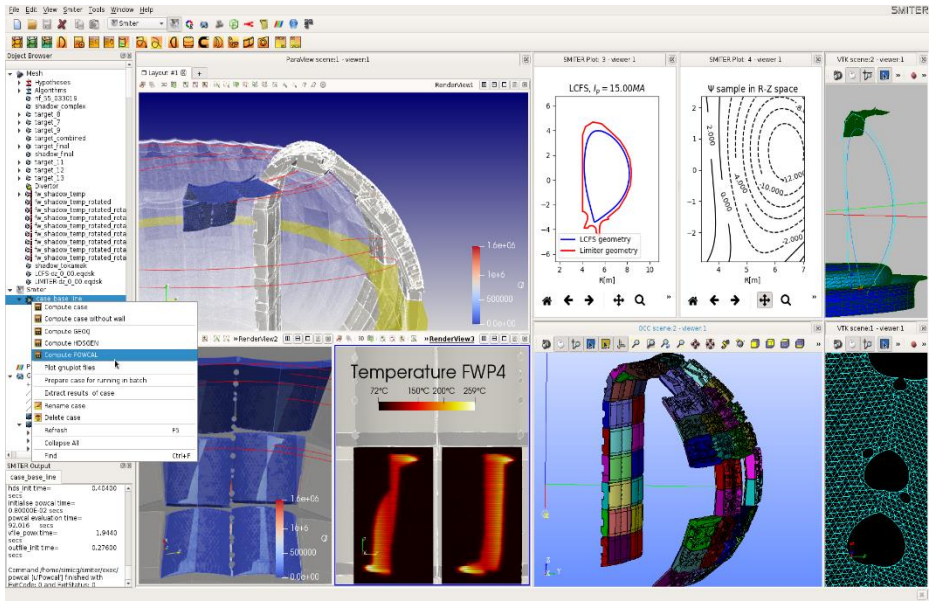

IMAS Software Management Tools

- Issue Tracking (JIRA)
  - Integrated across management tools
  - >3500 issues categorised across ~100 components
  - Issues are associated with components and are auto-assigned to the relevant ROs
- Distributed Revision Control (Git)
  - 70+ repositories grouped across 17 projects
  - Branch-based access control
- Automatic building, unit & regression testing, and deployment (CI server)
  - Continuous Integration to support agile development
  - Changes pushed to ITER repositories trigger automatic builds and testing to prevent regression (deny Pull Requests to merge changes)

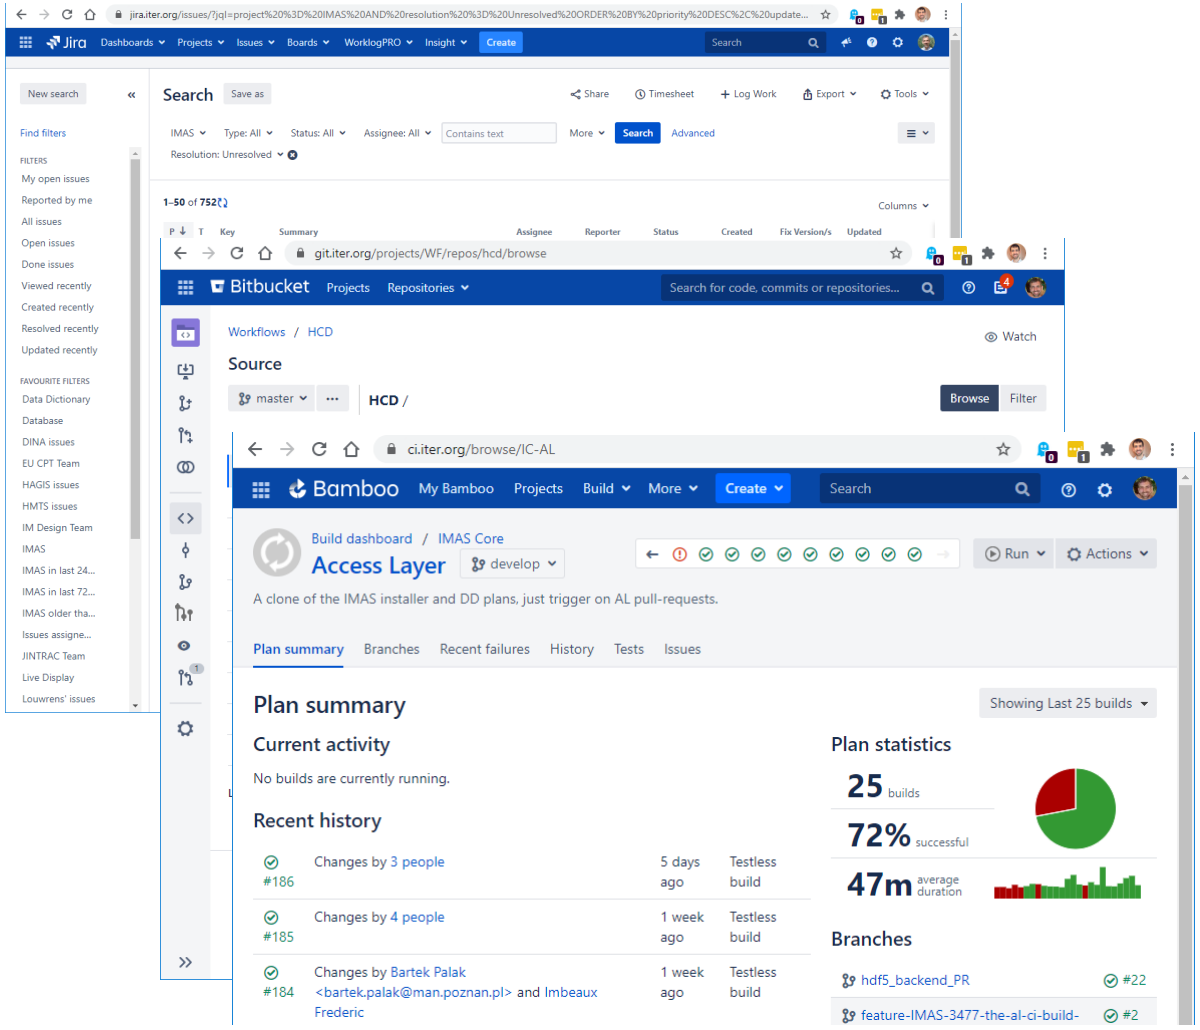

Use of EasyBuild at ITER

- EasyBuild originally used just to manage installation of physics codes' dependencies
  - Compilers and libraries, etc., with new easyconfigs created when not available
- The entire IMAS infrastructure being provided to users through a bespoke installer (supporting C, Fortran, Python, Java, MATLAB) and giving rise to one monolithic IMAS environment module with many dependencies
  - Compilers and libraries: GCC, intel, PGI, Java, Python, Blitz++,…
  - Installers for Linux, Windows and Macs

Using EasyBuild to install fusion physics software

- Started to install physics codes and physics workflows using EasyBuild
  - Includes compilation of native code (adapted to understand IMAS data structures)
  - And creation of workflow components (by running component generator) which for Python components are then added to $PYTHONPATH
  - Treats IMAS as an external module:
    - `builddependencies = [('IMAS/3.31.0-4.8.7', EXTERNAL_MODULE)]`
  - Physics workflows (depending upon physics codes) can also be installed using EB
    - E.g. HCD workflow which using different physics models to synergistically describe all the ITER plasma heating and current drive systems
  - Still mainly using 2018a toolchains but now moving to 2020a/b and hoping to do a better job at keeping up

Using EasyBuild to install IMAS infrastructure

- Fusion physics software not always open source
  - Different ITER Members have different attitudes
  - ITER is bound by ITER Agreement and Annex on Information and Intellectual Property
  - Patch EB framework to allow downloading source from sites requiring authentication (PR-3472) → *Essential for ITER software to be installed within Members using EB*

Using EasyBuild to install IMAS infrastructure

- Currently splitting-up software stack into separate development repositories (language APIs, data operations, storage, etc) which can each be installed using EasyBuild
  - IMAS/3.31.0-4.8.7 → AL-Fortran/4.8.7-intel-2020a-DD-3.31.0, etc.
- Once ready, it's hoped this will provide an easier way for collaborators to install (and thus validate and use) ITER tools at their sites
  - Physics applications, IMAS infrastructure and all dependencies

Using EasyBuild to install IMAS infrastructure

- To further simplify the installation and use of ITER physics applications within the ITER Members, would like to use EB to help create containers

Summary of EasyBuild at ITER

- EasyBuild started off as a way to simplify installing scientific software dependencies on our computing systems
  - With new easyconfigs created as necessary
- Use has grown to include installing physics applications that externally depend upon ITER's own IMAS software stack
- IMAS software stack being further modularized with aim to being installed using EasyBuild
  - Aim is to provide an easy and robust way for our partners to build and install ITER scientific software on their systems
- Hope this will provide an easy route to creating containers for ITER physics applications

…and finally

- Travel restrictions permitting, feel free to come and pay us a visit and see ITER up close
  - See https://www.iter.org/visiting
- We have student internships, postdoctoral fellowships (incl. with Principality of Monaco) and jobs for those who may be interested
  - See https://www.iter.org/jobs
- We also currently have a Call for Nomination on-going for anyone who may be interested in helping us further develop IMAS and prepare for ITER operations
  - See https://www.iter.org/org/team/adm/proc/overview#CFT/789