The HPC Skill Tree
Design, current content, and how to work with it

Kevin Lüdemann, GWDG

Kevin.luedemann@gwdg.de

Table of contents

The Skill Tree

- A skill defines the contents to be learned and tested
- The skill tree organizes the competences as hierarchical skills
- Certificates bundle several skills into attestable unit

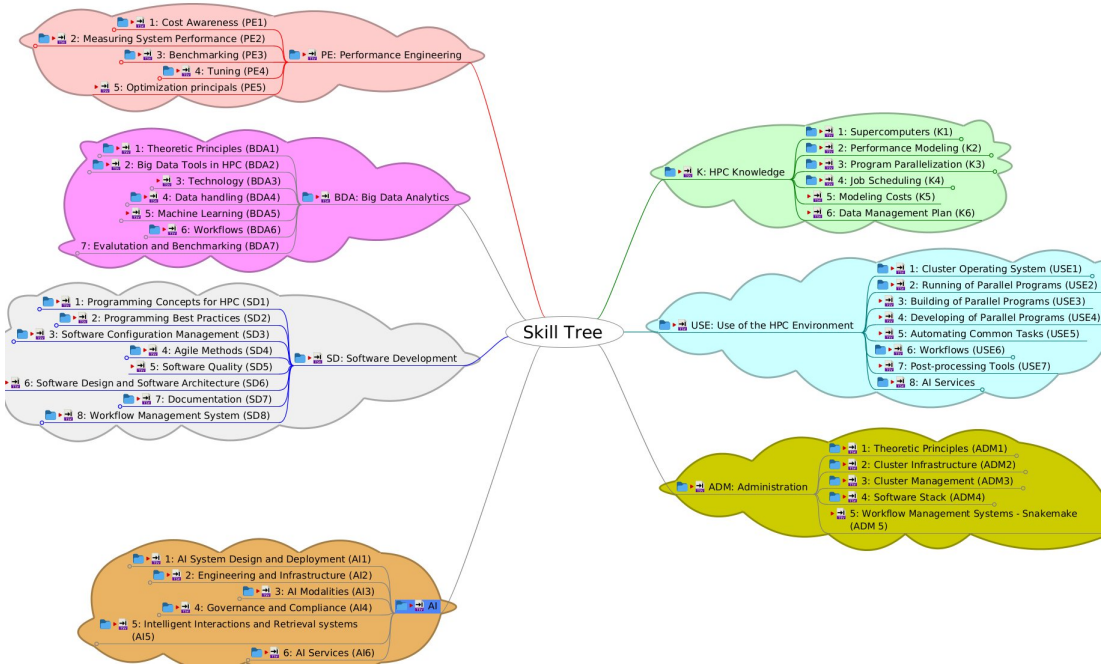

The Skill Tree

- Using an HPC environment including Linux basics
- Organized into Node Skills, further Broken Down into Leaf Skills

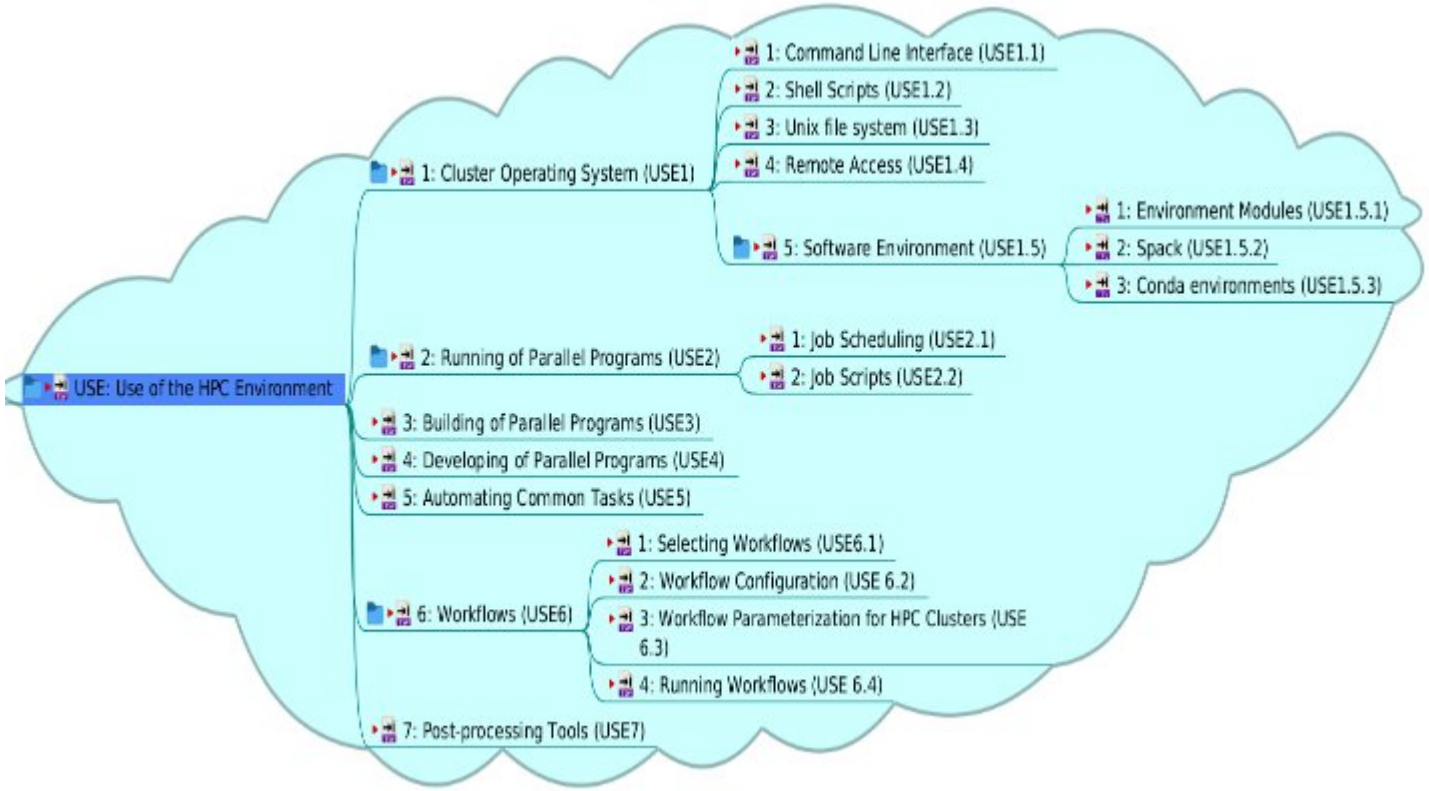

Skill Tree – The Mind Map

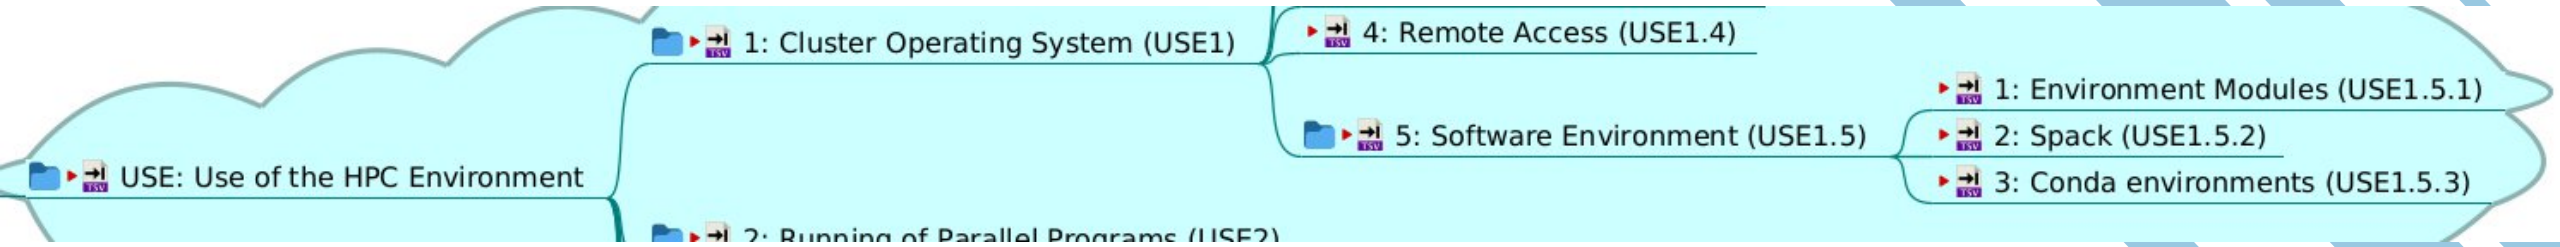

- Each branch contains:
  - Nodes with leaves
  - Folder symbols for node skills
  - Links to files of wiki for all skills, leaf and node skills
- Overview/Node skill
  - Contains the descriptions of skills under it
  - Contains aggregates and distilled learning objectives from skills

Skill Tree – A Leaf Skill

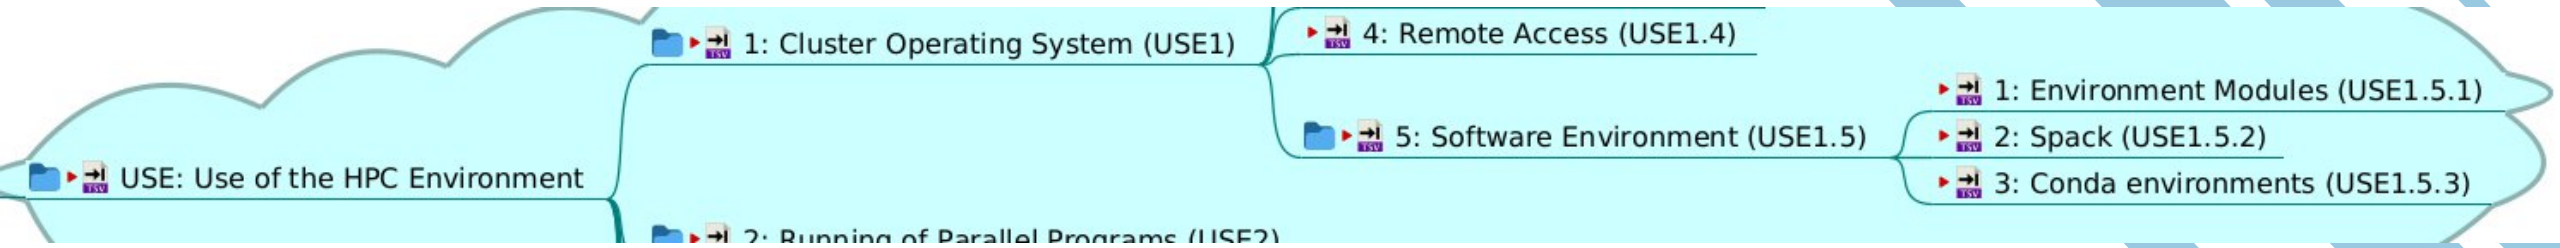

- Single enclosed, universal unit of learning
- Contains several learning objectives
  - A leaf skill has content for 1 to 4 hours (including self-study time)
  - Framing of learning objectives follow: Bloom's taxonomy
- Has a unique identifier: for example USE1.5.1
- Can have requirements of other skills (either node or leaf)

Skill Tree – A Node Skill (Overview)

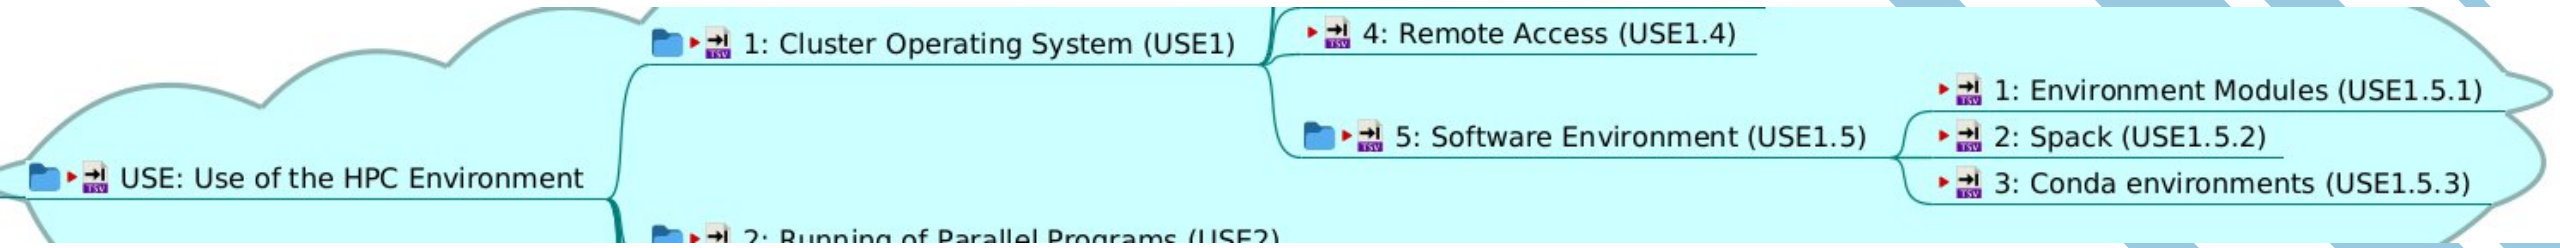

- An aggregate unit of learning
- Contains several learning objectives
  - 1-2 learning objectives distilled from skills immediately bellow it
  - Learning time set by skill below this node
  - Framing of learning objectives follow: Bloom's taxonomy
- Has a unique identifier: for example USE1.5
- Can have requirements of other skills (either node or leaf)

USE: Use of the HPC Environment

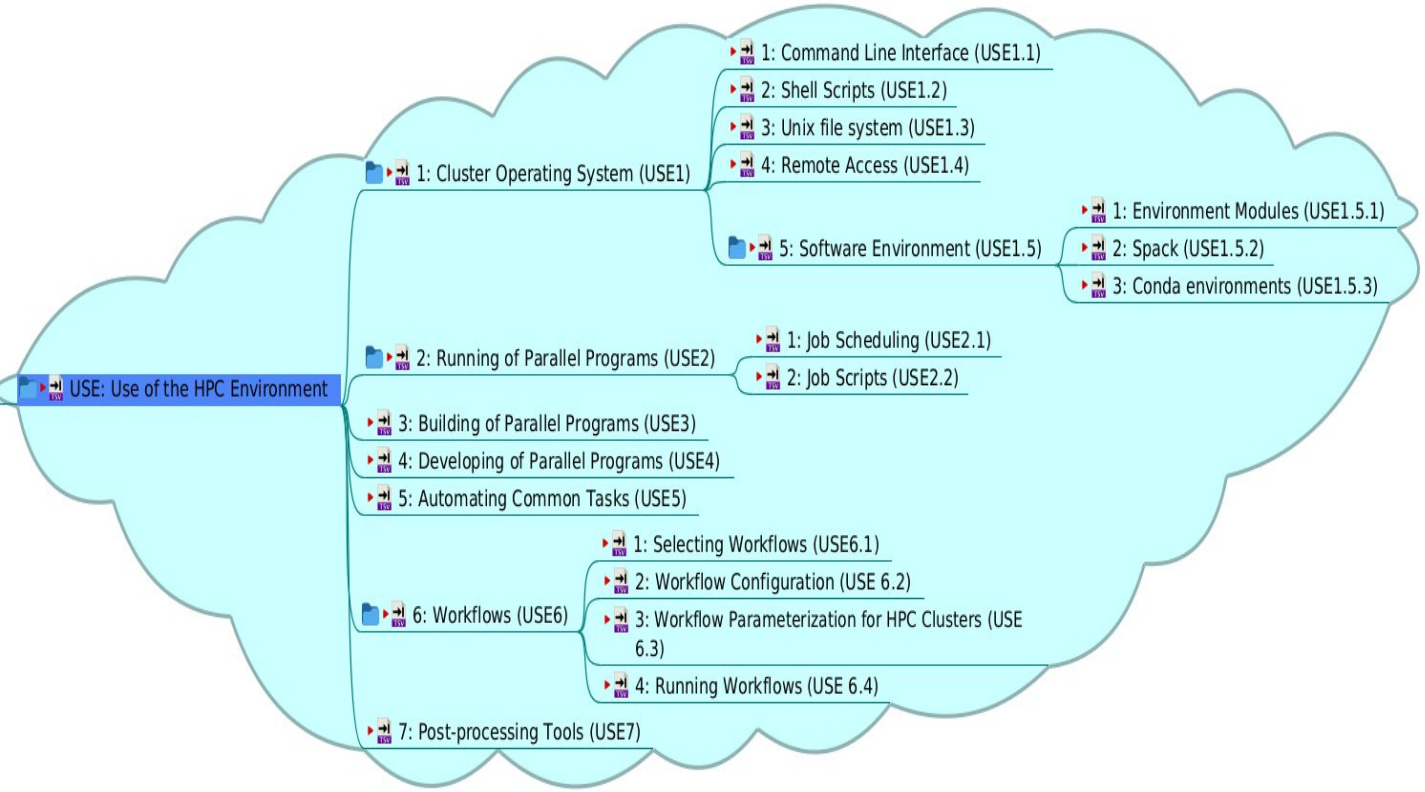

- Using a cluster
- Includes Linux and CLI
- Running jobs without deeper knowledge
  - Deeper knowledge is K
  - Administration is in ADM
- Basis for almost all other branches

K: HPC Knowledge Branch

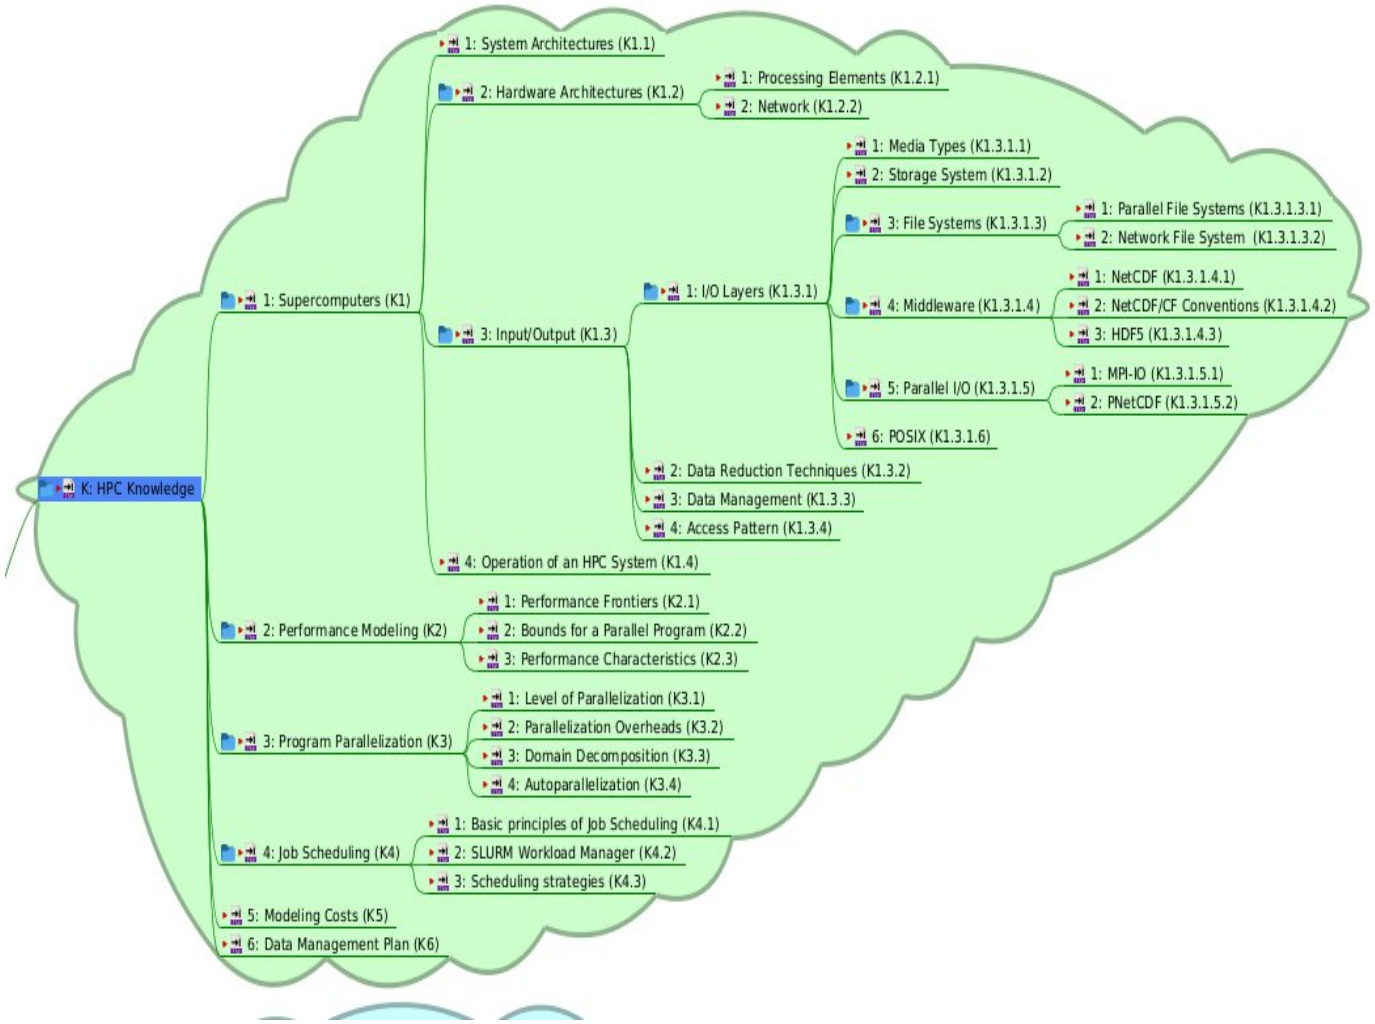

- Understanding of HPC
- Deeper knowledge when compared to USE
- Administration is a different branch
- Example:
  - USE: Using a cluster
  - K: Understanding a specific workload manager i.e. SLURM

ADM: Administration

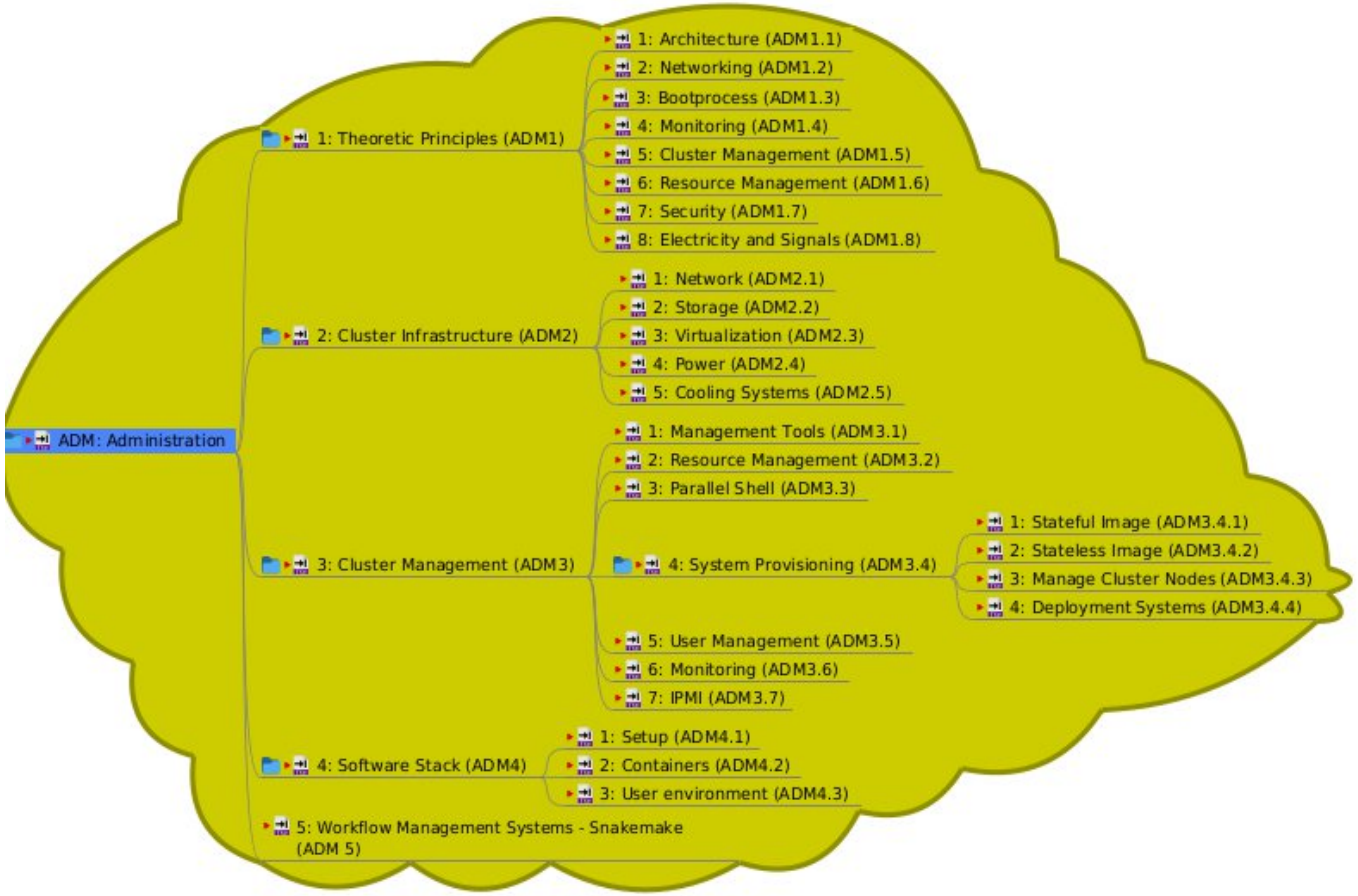

- Administrating an HPC system (sysadmin)
- Setting up and managing a cluster
- Requires K knowledge
- Example:
  - USE: Using a cluster
  - K: Understanding workload manager
  - ADM: Setting up a workload manager

SD: Software Development

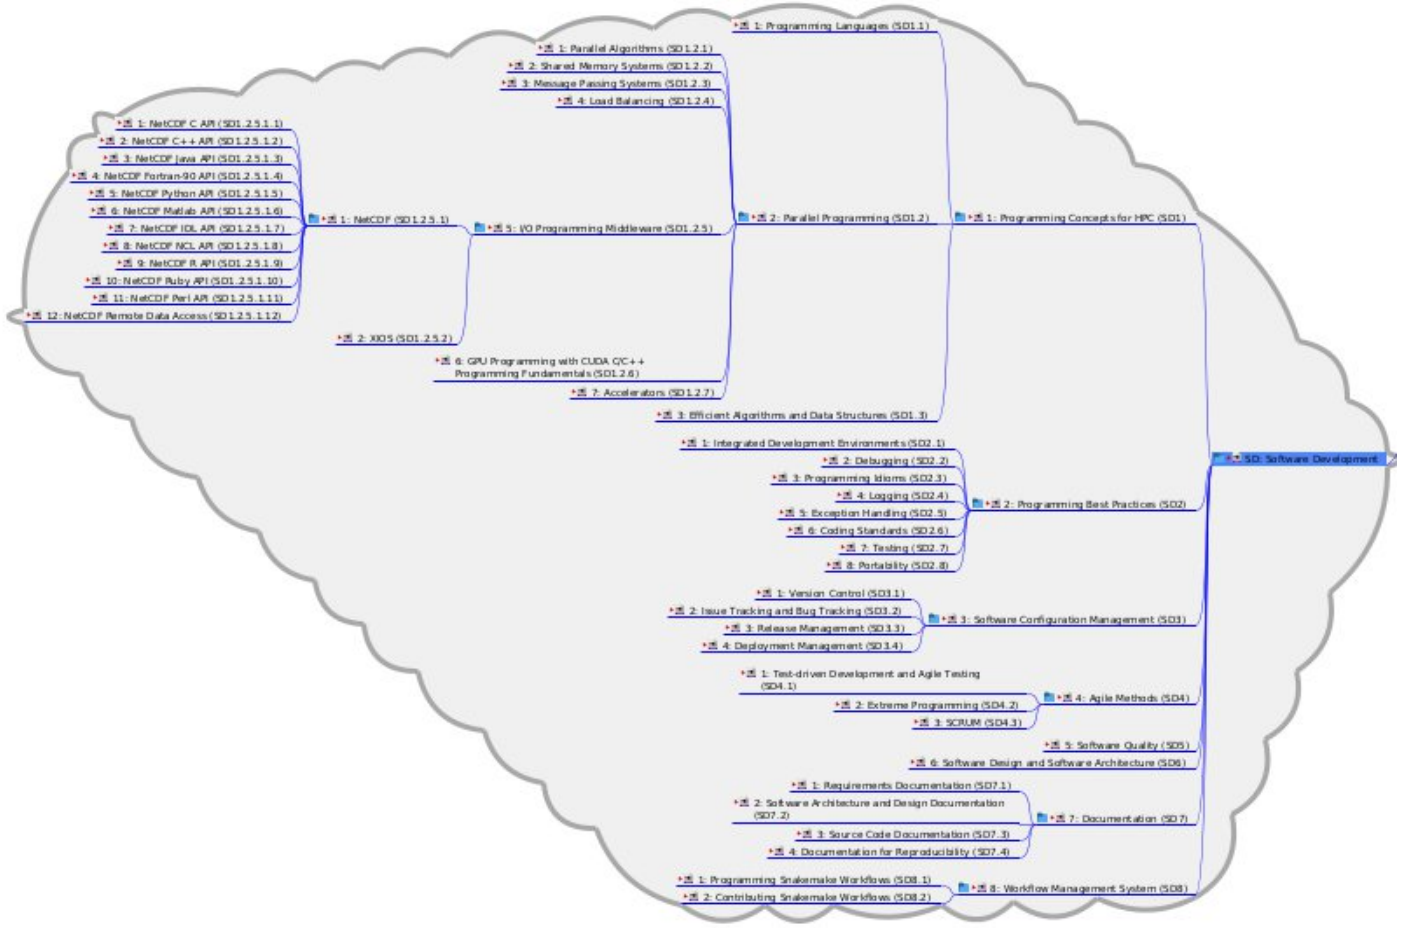

- Parallel Programming,. I.e. MPI
- Best practices
- Documentation
- Understanding hardware accelerators
- Programming languages are external for now

PE: Performance Engineering

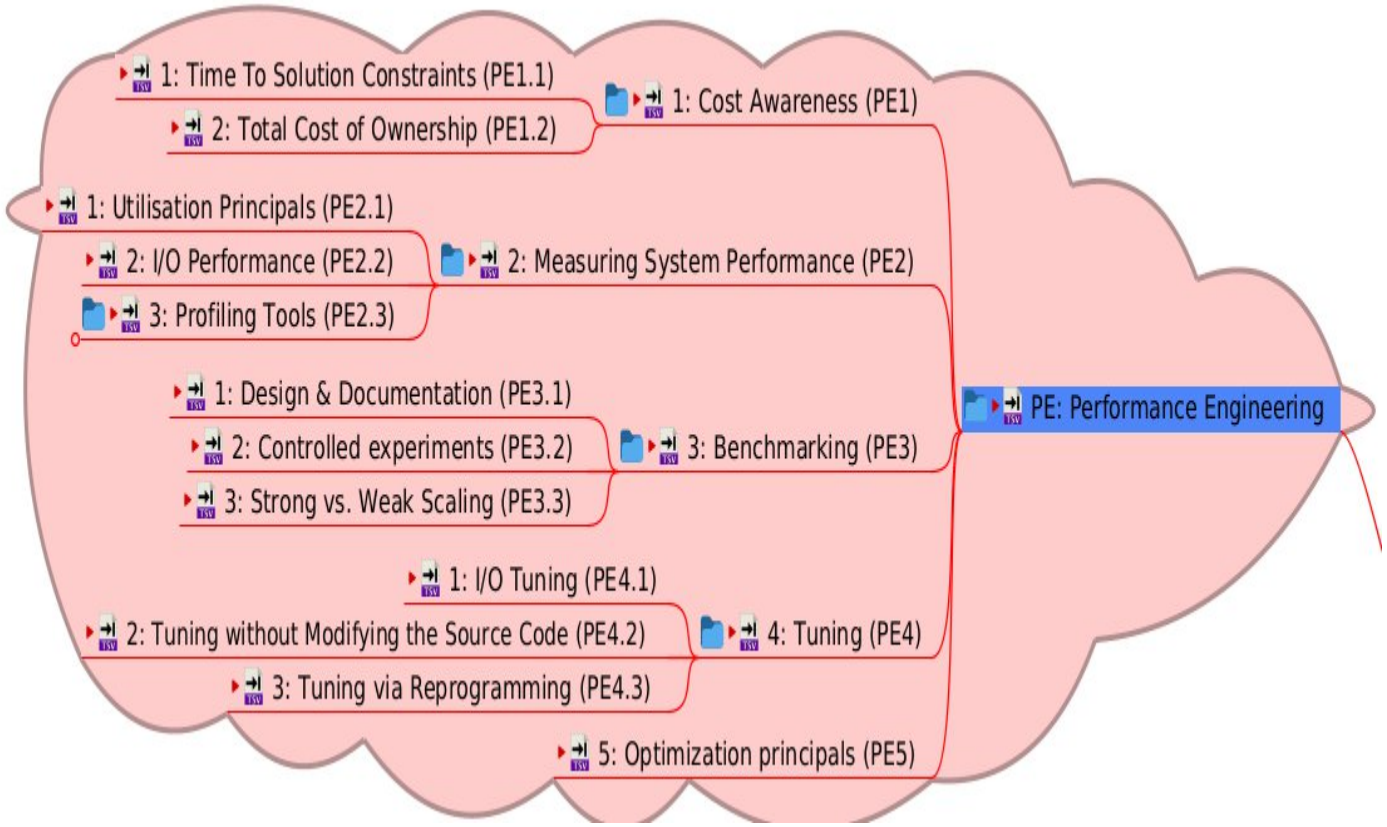

- Understanding costs
- Tools for performance analysis and measuring
- Benchmarking and specific benchmarks
- Principles of optimization/tuning

BDA: Big Data Analytics

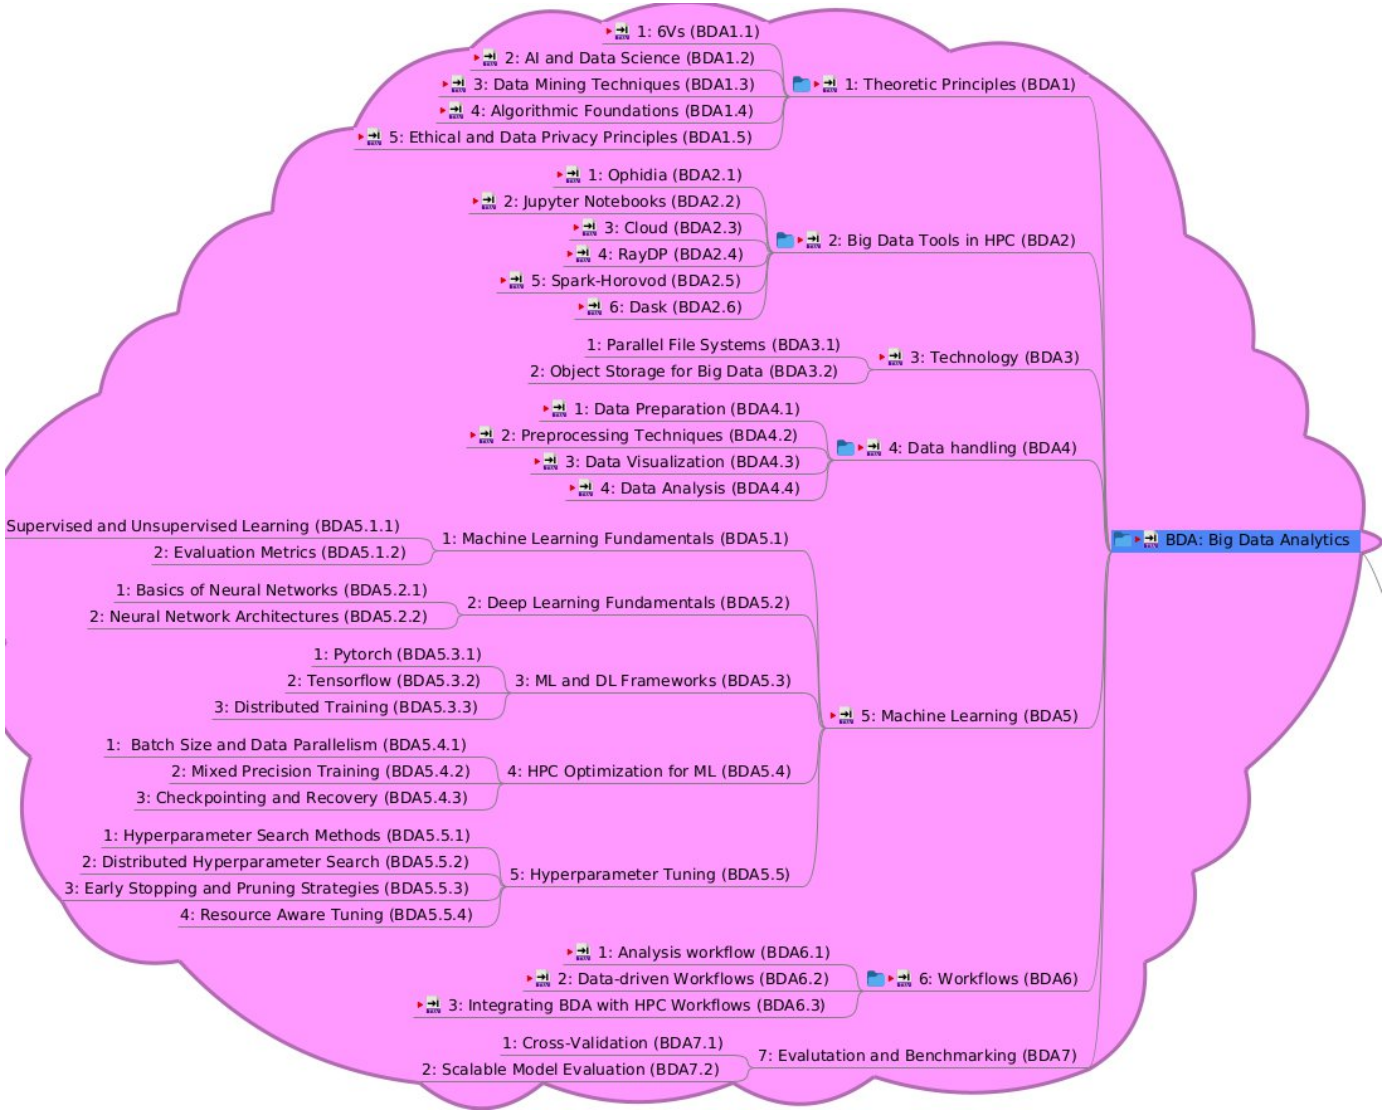

- Everything that related to big data
  - Tools
  - Processing best practices
  - Currently Machine Learning
- AI will become its own branch

AI: Artificial intelligence

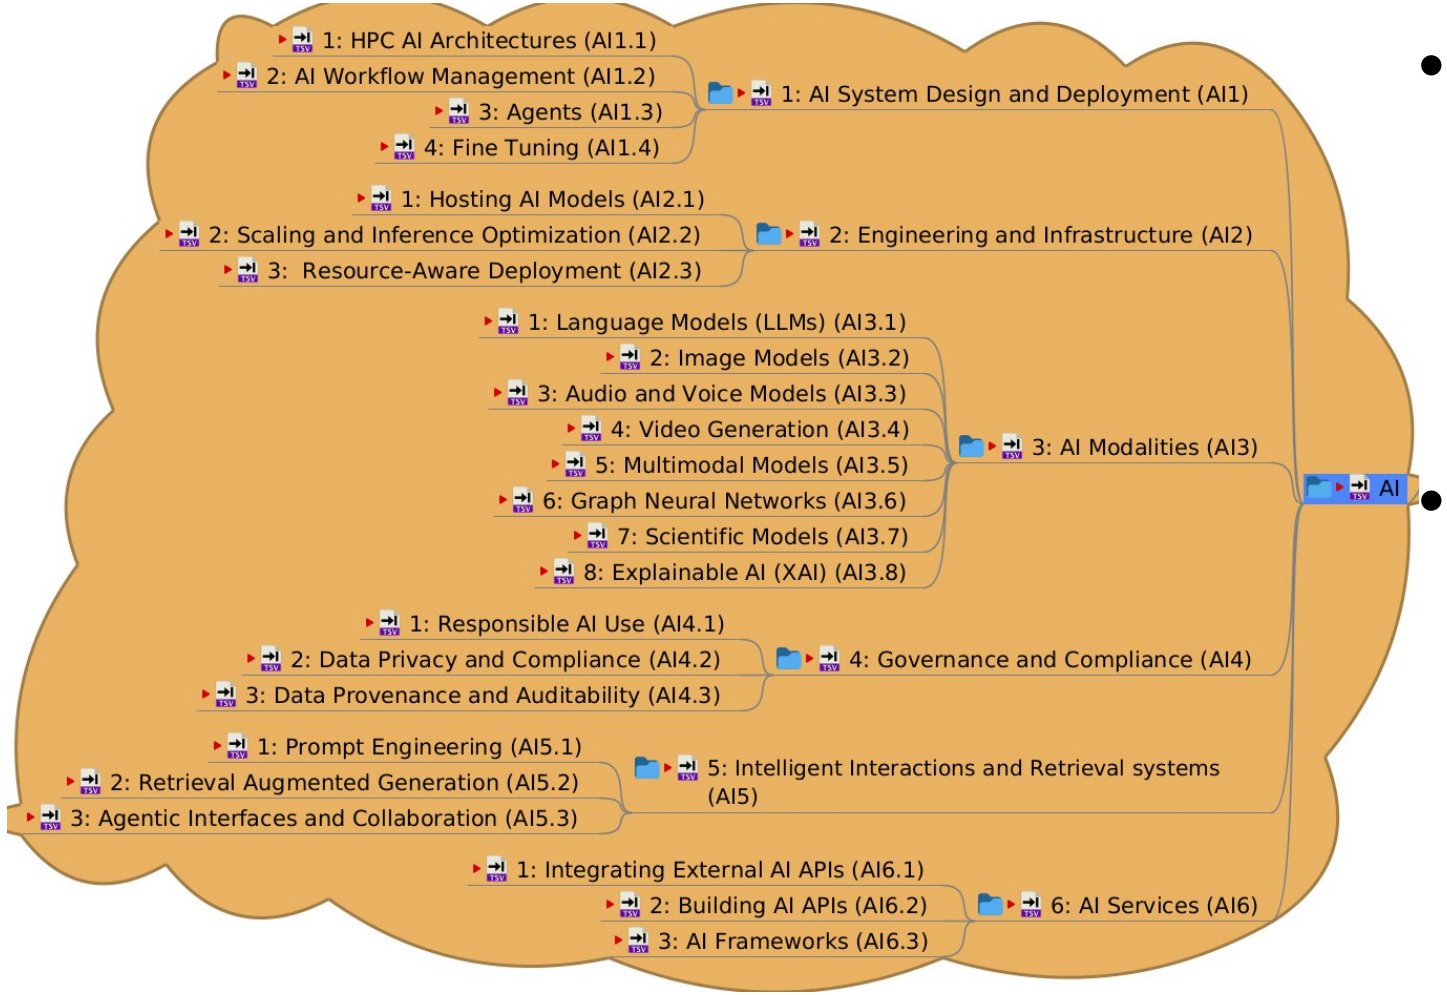

- Everything that related to big data
  - Tools
  - Processing best practices
  - Currently Machine Learning
- AI will become its own branch

Table of contents

Beyond the Mind Map

- Requirements also allow a network representation
- Mind Map only contains relations and names
- Reminder, what is a leaf Skill
  - Single enclosed, universal unit of learning
  - Contains several learning objectives with content for 1 to 4 hours (including some self-study time)
- Reminder, what is a node skill
  - An aggregate unit of learning
  - Contains several learning objectives
    - 1-2 learning objectives distilled from skills immediately bellow it
    - Learning time set by skill below this node
    - Framing of learning objectives follow: Bloom’s taxonomy

HPC CF Wiki – The Markdown files

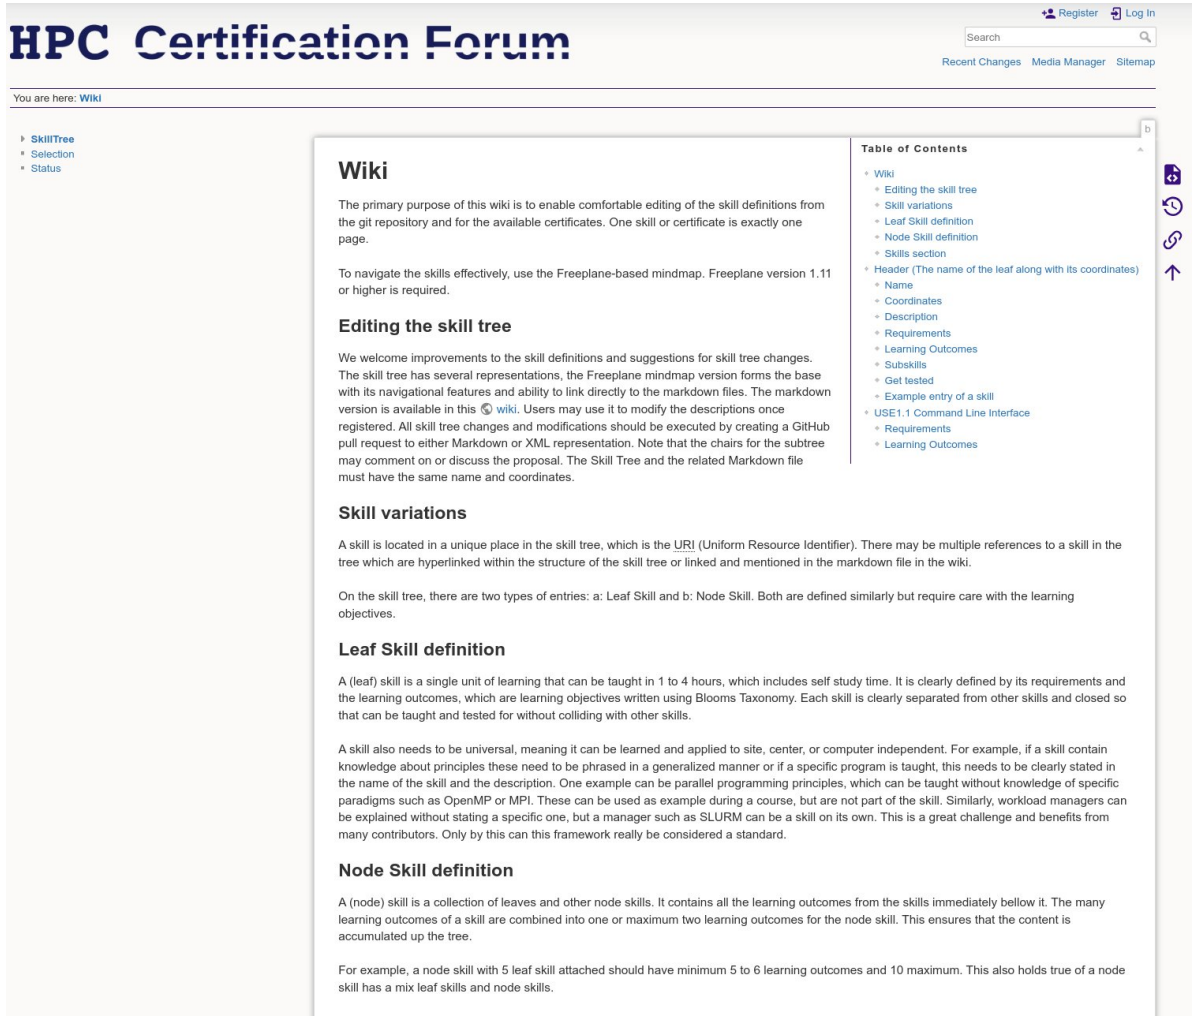

- Wiki main page
- Interaction point for each skill
- Each skill has associated Markdown file
- File contains learning objectives
- Edit the Wiki if anything
- There is assistance for the Mind Map file
- Definition for both leaf skill and node skill

The Wiki – A Skill

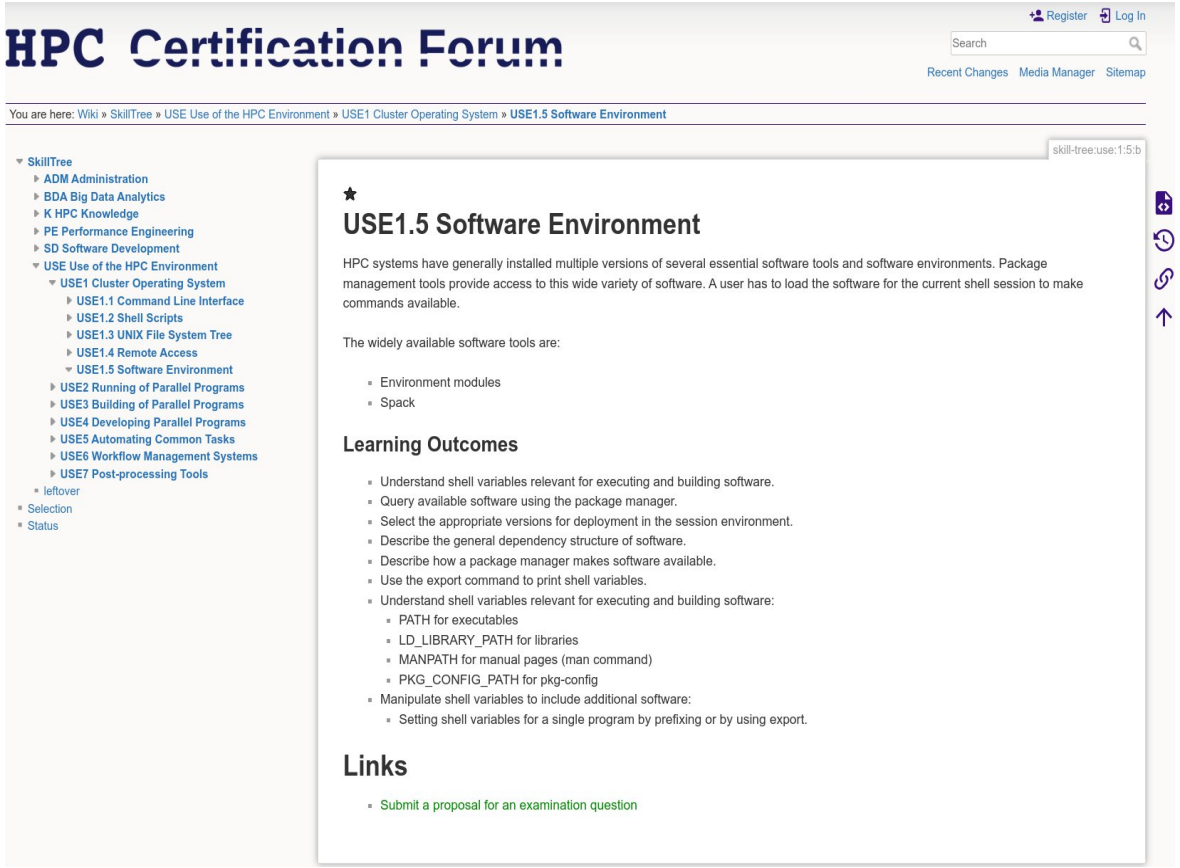

- This is a (leaf) skill and it has
  - A unique coordinate
  - A description
  - Some learning objectives
  - Maybe requirement skill(s)
- Learning objectives cover
  - Module environment
  - SPACK system
  - Are written out using Blooms taxonomy

Selecting skills for courses

- Each course has specific learning objectives
- Write them all out in the style of Blooms taxonomy
- Find skills that contain your learning objectives
- A course/module
  - will cover more content than can be expressed in skills
  - should be more specific to a site or cluster
  - needs to fully cover a skill if it is tagged

Course selection

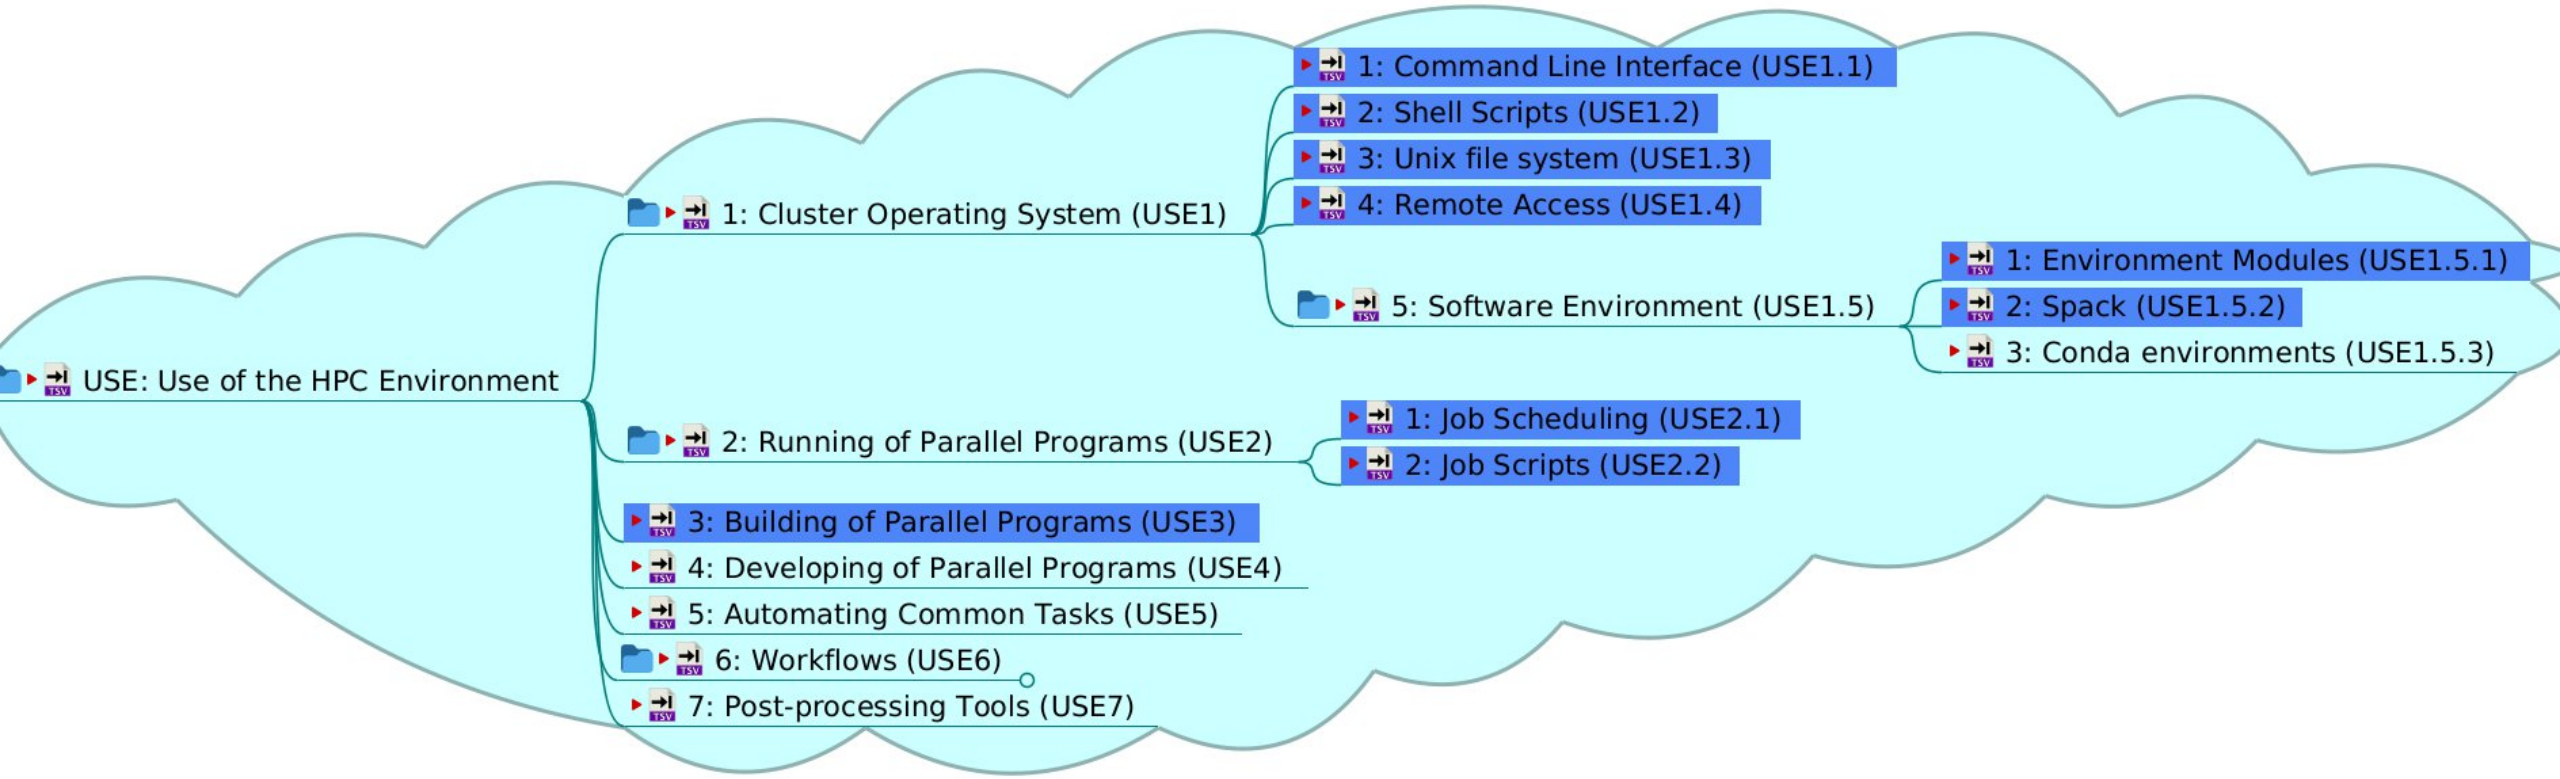

A skill does not fit correctly

- Skill has more learning objectives:
  - Investigate if you really need to link to this skill
  - Check the course/module and update the content to match
  - Check if the skill needs to be split if the difference is too large
- Skill has less/fewer learning objectives:
  - Check if you need two skills
  - Update the skill
    - either with your additional learning objectives
    - or split the skill and put your additional objectives into a new skill
- In both cases update node skills as well

Style of courses

- Overview courses/modules
  - Can have many skills, maybe linking node skills is better
  - Will require a lot of self study time from participants
  - Learning objectives might be mostly low in Taxonomy pyramid
- Deep dive courses/modules
  - Might just cover one or two skills for full day course
  - Contains already a large amount of study time with exercises
  - Learning objectives might be mostly high in Taxonomy pyramid
- Longer courses/modules
  - Usually a mix of both covering several skills
  - Exercises included in study time, i.e Semester courses

Table of contents

The HPC Certification Forum

- Goals: Fine-grained HPC knowledge representation ⇒ Competence Standard
  - What competences exist, how are they defined?
  - Puzzle of competences for everyone (practitioners, students, admins)
  - Supporting navigation and role-specific knowledge maps -> Learning pathways/baselines
- Establishing international certificates attesting knowledge
- Scope: Central authority for competence representation, certification, and support with purposeful limitations:
  - We do not compete with content providers
  - We do not create a curriculum (university/centers responsibility)

The HPC Certification Forum

- Organization Details
  - An independent international body
  - Organized into
    - Steering board (elected)
    - Full members (with voting rights) - Contributors
    - Associate members (anyone and any institution)
    - Collaboration with e.g. EuroHPC JU projects HPC SPECTRA & EViTA
- Activities
  - Curating and maintaining the Competence Standard
  - Providing tools and ecosystem around the competences

The HPC Certification Forum

- Organization of the members:
  - Webpage is the central hub (https://www.hpc-certification.org)
  - Mailinglists (news, members, board)
  - Monthly public meetings on our Slack channel
  - Monthly meeting of the board
  - Annual general assembly (BoF at ISC or independent workshop)
- Data handling
  - Everything is developed/available in the open GitHub (https://github.com/HPC-certification-forum)
  - Exception are examination questions

AI Tools

- Learning Objectives in Blooms Taxonomy
  - https://chatgpt.com/g/g-67768fbf69dc819199207ef9e3d4d697-blooms-taxonomy-for-hpccf
- Sorting through the Skill tree using AI
  - https://chatgpt.com/g/g-67eab1af25908191a09bfc33eeb3811d-hpc-cf-skill-tree-sorter
  - Contains the current skill tree and all Descriptions and Learning objectives
  - Will be constantly updated once updates happen

Table of contents

HPC SPECTRA

- Funded by EuroHPC JU: 101136267
- Main deliverable: Training platform + Supporting IHPCSS

HPC SPECTRA

- International HPC Summer School (IHPCSS)
  - Organizing Summer School
  - Funding students
- Training platform: https://events.hpc-portal.eu/
  - Central database for courses across Europe
  - Browsable and bookable trainings
  - Interfacing with skill framework from HPC CF

Table of contents

HPC SPECTRA EVITA

- Funded by EuroHPC JU: 101196394
- 4 year project and successor to HPC SPECTRA with 1 year overlap

EVITA

- Progressing the Training platform to include a material catalog
- Funding module creation with skill framework as basis
- Interfaces and advances the HPC CF Skill tree
- Establishing certificates for modules through examination framework
- Connecting certificates to ECTS and Micro-credentials

Table of contents

Summary

- The Skill Tree is a hierarchical approach
- Every skill contains learning outcomes split into learnable and attestable usnits
- A course tagging a skill covers all of the learning outcomes
- The HPC Certification Forum (HPC CF) maintains the standard
- HPC SPECTRA and EVITA both use this Skill Tree
- The Skill Tree is a community effort and free to use

Thank you!

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101136267. The JU receives support from the European Union's Digital Europe Programme.

Funded by
the European Union