OpenAlt 2026

Projekt FIxbot pro AI-asistovanou výuku a zkoušení FIxbot.fi.muni.cz

Jakub Čížek, Michal Louda, Petr Švenda

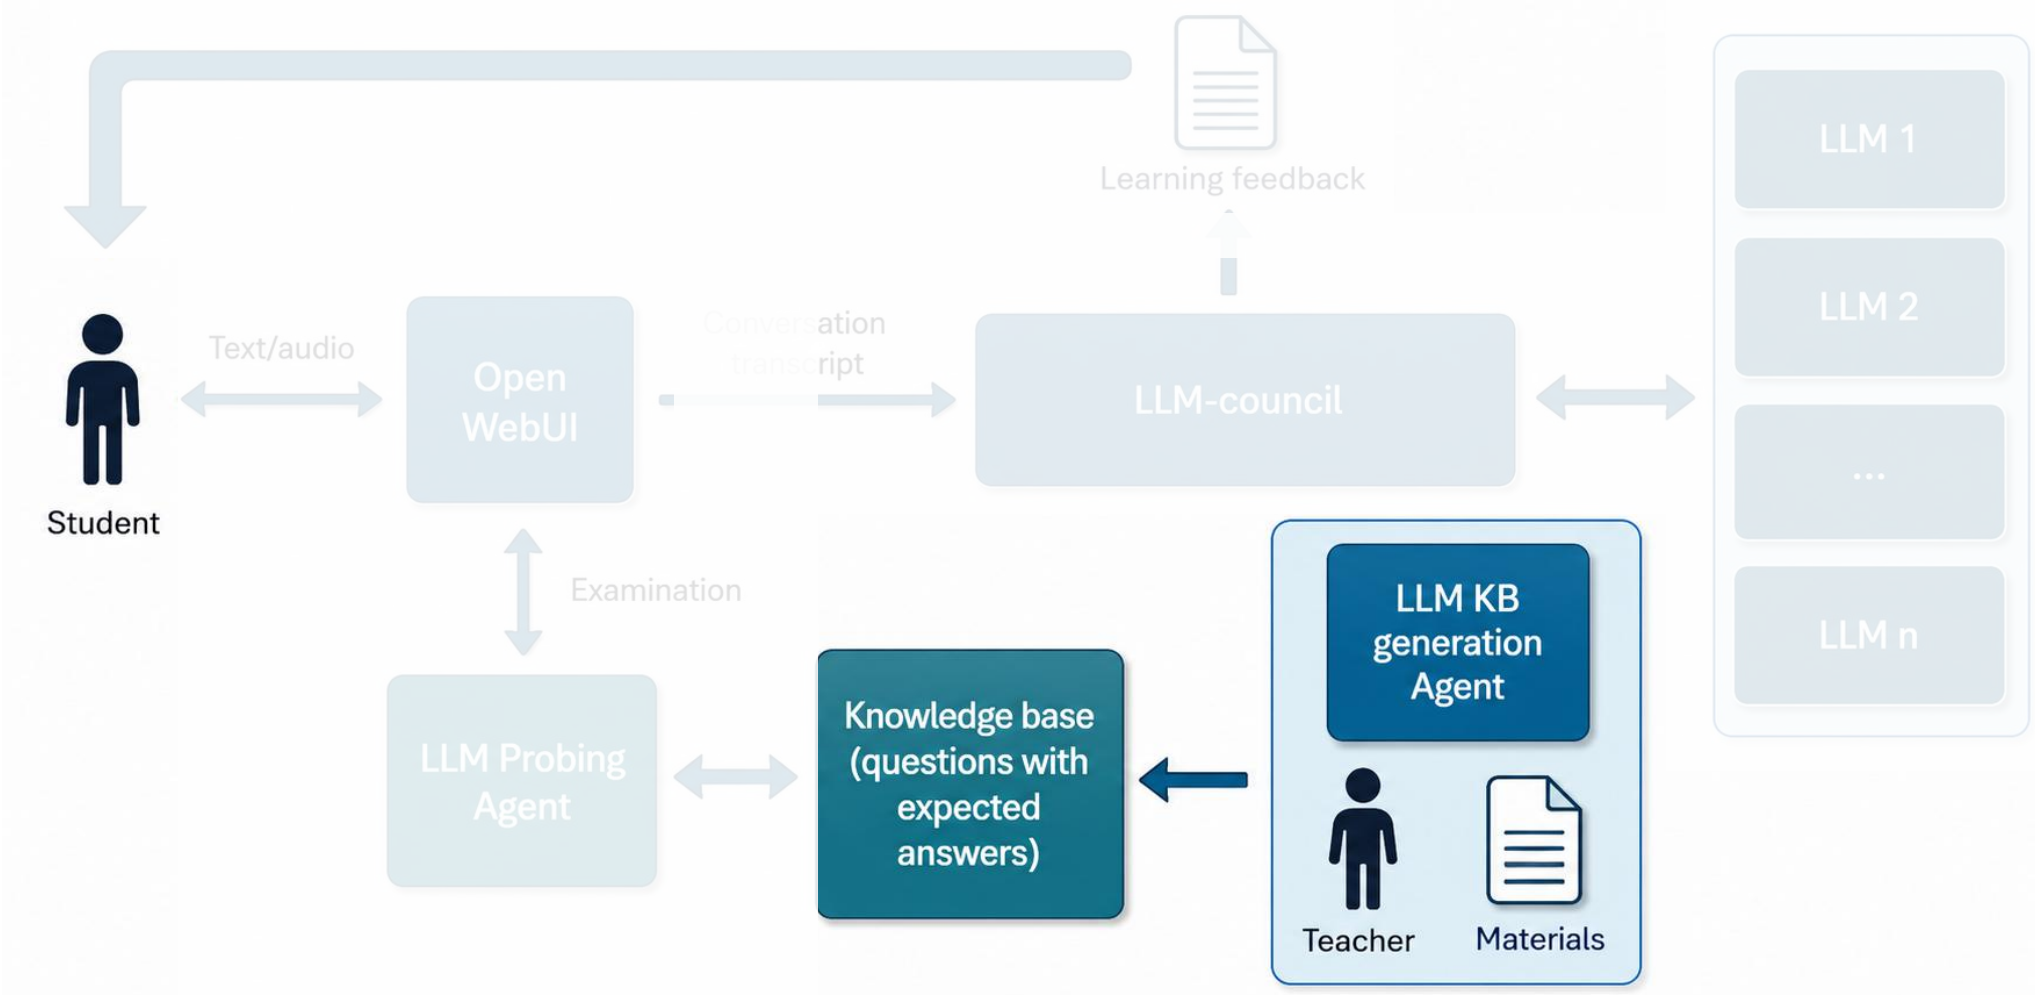
Fakulta Informatiky, Masarykova univerzita

Projekt MUNI/AI-I/1196/2026

Motivation: Oral examination is great; role of a teacher

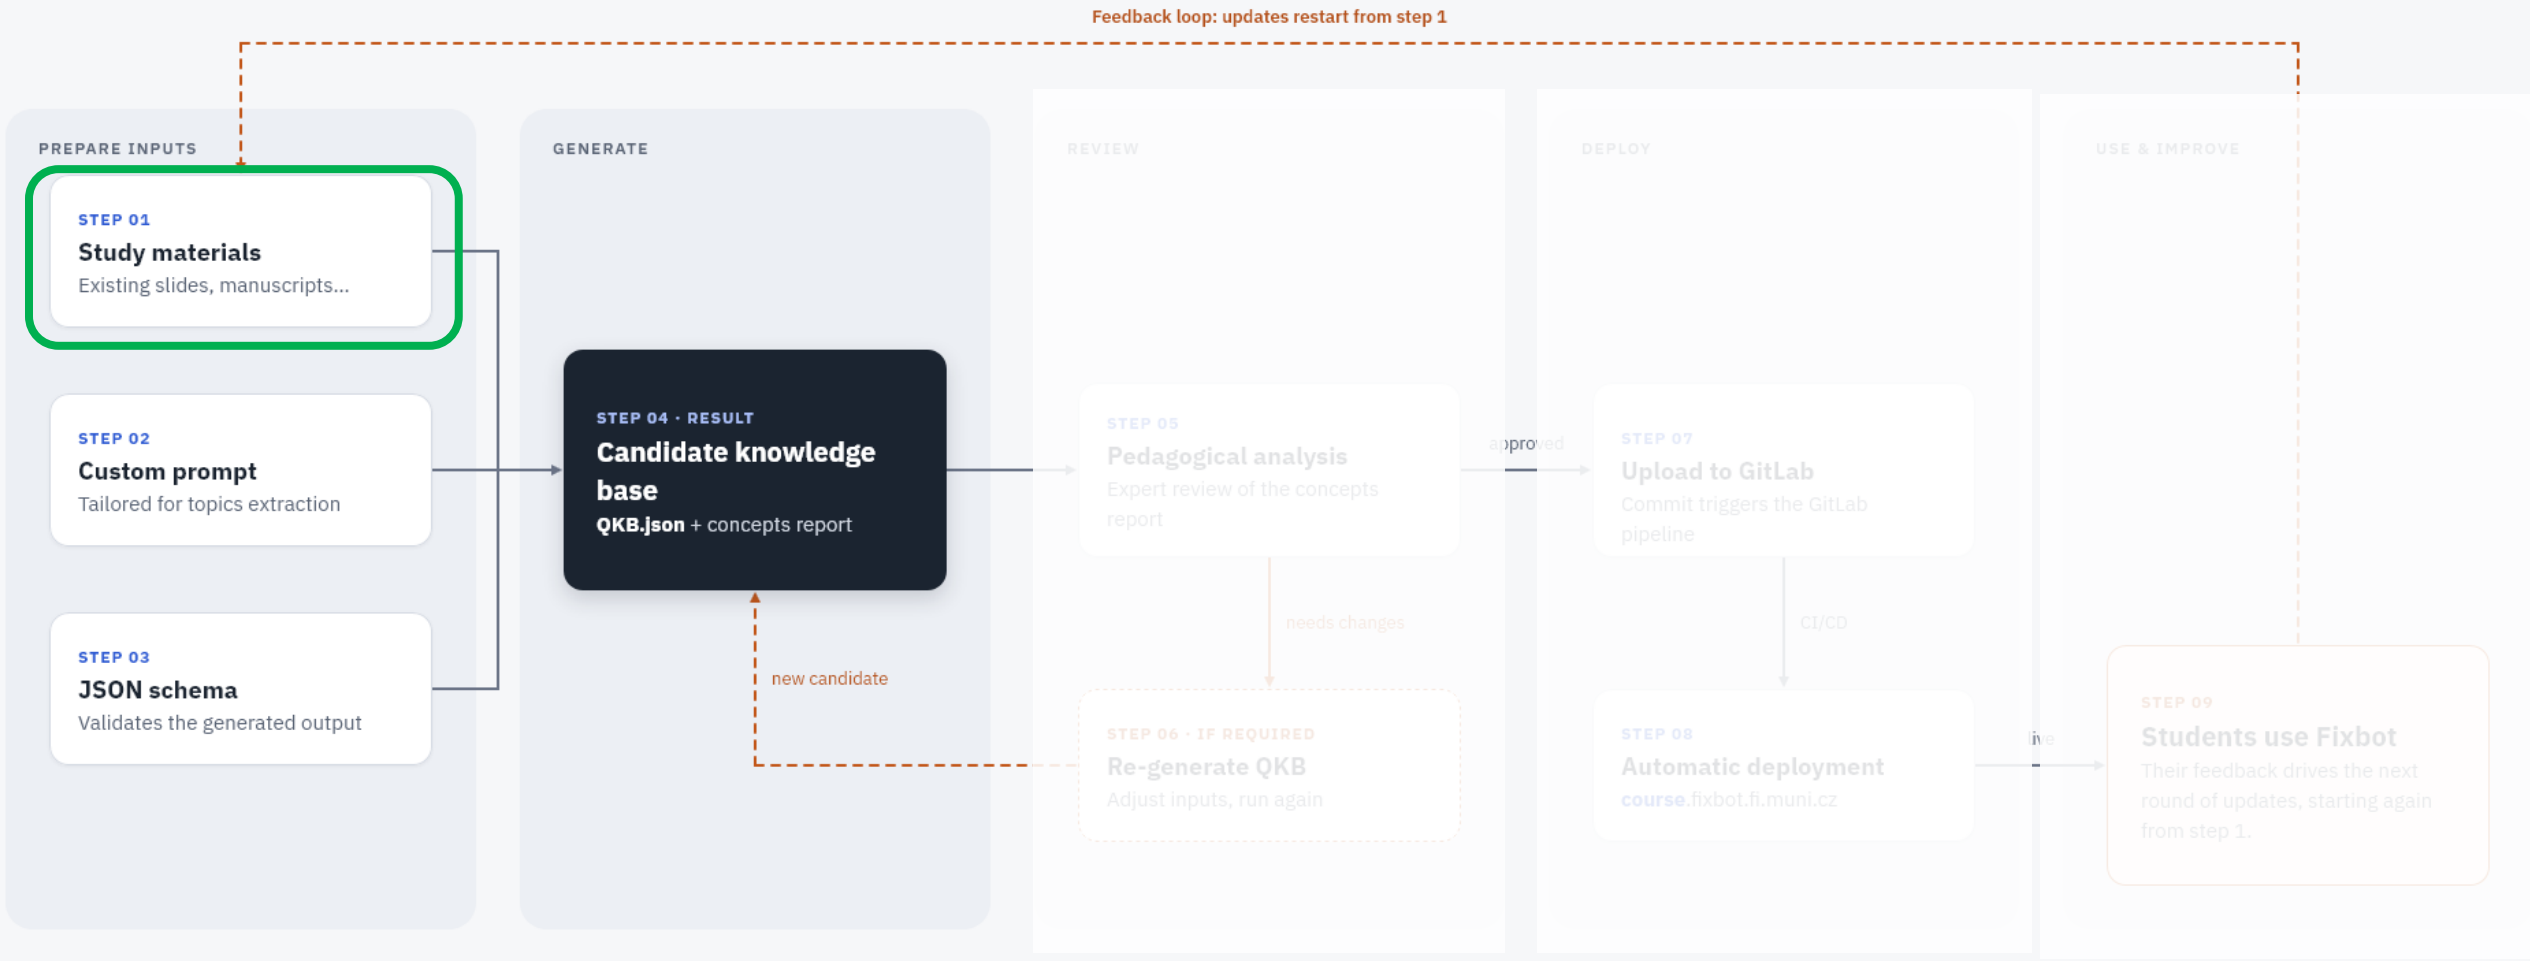
- Oral interaction and examination is great for learning and assessment but does not scale for CS education (due to student/teacher ratio)
  1. Limited time of human teachers for student questions => how to make every TA's minute be most valuable?
  2. Very few attempts per student (exam) and little time to catchup if not passing => need earlier indication of knowledge gaps than during real exam
- AI/LLM promise
  - Can keep explaining any* topic, anytime, any depth, in any* language
- Short history of teacher value over time (simplified ☺)
  - Pre-book/pre-internet: knowledge itself, pedagogically sound scope and explanation
  - Pre-AI: condensed knowledge, pedagogically sound scope and explanation
  - Now: Pedagogically sound scope, explanation of high-complexity topics
  - Future?: Curiosity ignitor/motivator, long-task persistency mentor…

High-level fixbot architecture



Fixbot lifecycle

How existing study materials become a validated knowledge base, deployed to students and improved through their feedback



```
1:
  number:                  "W01.Q2"
  primary_topic:           "Integer division vs real division: `//` vs `/`"
  primary_question:        "What is the difference between `7 // 2` and `7 / 2` in Python? What are the types and
                           values of each result, and when would you choose one over the other?"
  target_bloom_level:      "Apply"
  minimum_passing_bloom_level: "Understand"
  scope:
    expected_student_level: "first week of introductory Python programming, no prior programming assumed"
    in_scope:
      0:                   "integer/floor division operator `//`"
      1:                   "real division operator `/`"
      2:                   "result types for integer operands"
      3:                   "floor division behavior for negative numbers as optional precision"
      4:                   "practical use cases for each operator"
    out_of_scope:
      0:                   "operator overloading"
      1:                   "floating point representation errors"
      2:                   "modulo operator `%` as a primary topic"
      3:                   "division behavior in other programming languages"
```

```
excellent_answer: "`//` is floor division: it divides and rounds the mathematical result down. With integer operands, the result is an int, so `7 // 2` is 3. For negative operands, this matters: `-7 // 2` is -4, not -3. `/` is real division and returns a float, so `7 / 2` is 3.5 and `4 / 2` is 2.0. The choice should follow the meaning of the result, not speed or style."
bloom_feedback_moves:
  Remember: "Ask the student to state the value and type of `7 // 2` and `7 / 2`."
  Understand: "Ask the student to explain why `4 / 2` gives `2.0`, not integer `2`."
  Apply: "Ask the student to choose the correct operator for computing a whole-number index."
  Analyze: "Ask the student to compare floor division with truncation for `-7 // 2`."
  Evaluate: "Show code that uses `/` where an integer index is required and ask the student to identif the problem."
common_near_misses:
  0:
    student_pattern: "`//` rounds to nearest integer"
    classification: "misconception"
    feedback: "Floor division rounds down, not to the nearest integer."
  1:
    student_pattern: "`/` gives an int when the result is a whole number"
    classification: "misconception"
    feedback: "In Python, `/` produces a float; `4 / 2` is `2.0`."
  2:
    student_pattern: "Use `//` for speed, `/` for accuracy"
```

```
forbidden_drift:
  0: "Do not discuss Python 2 integer division unless the student directly raises it."
  1: "Do not discuss floating point precision errors such as `0.1 + 0.2 != 0.3`."
  2: "Do not make modulo `%` the primary topic here."
checklist:
  0:
    id: "W01.Q2.C1"
    concept: "Result of `/` is real division / float"
    bloom_level: "Understand"
    required: true
    weight: 2
    question: "Did the student state that `/` performs real division and produces a float result?"
    acceptable_evidence:
      0: "says `7 / 2` equals `3.5`"
      1: "says `/` is real division"
      2: "says `4 / 2` gives `2.0`, not integer `2`"
    common_misconceptions:
      0: "
      1: "the result type depends on whether the mathematical result is a whole number"
      2: "`/` is the same as `//` for integer inputs"
    repair_prompt: "Ask: what are the value and type of `4 / 2`? How does that compare to `4 // 2`?"
```

Probing model is used only as “intelligent” diff between pre-defined topic knowledge base and answers from student

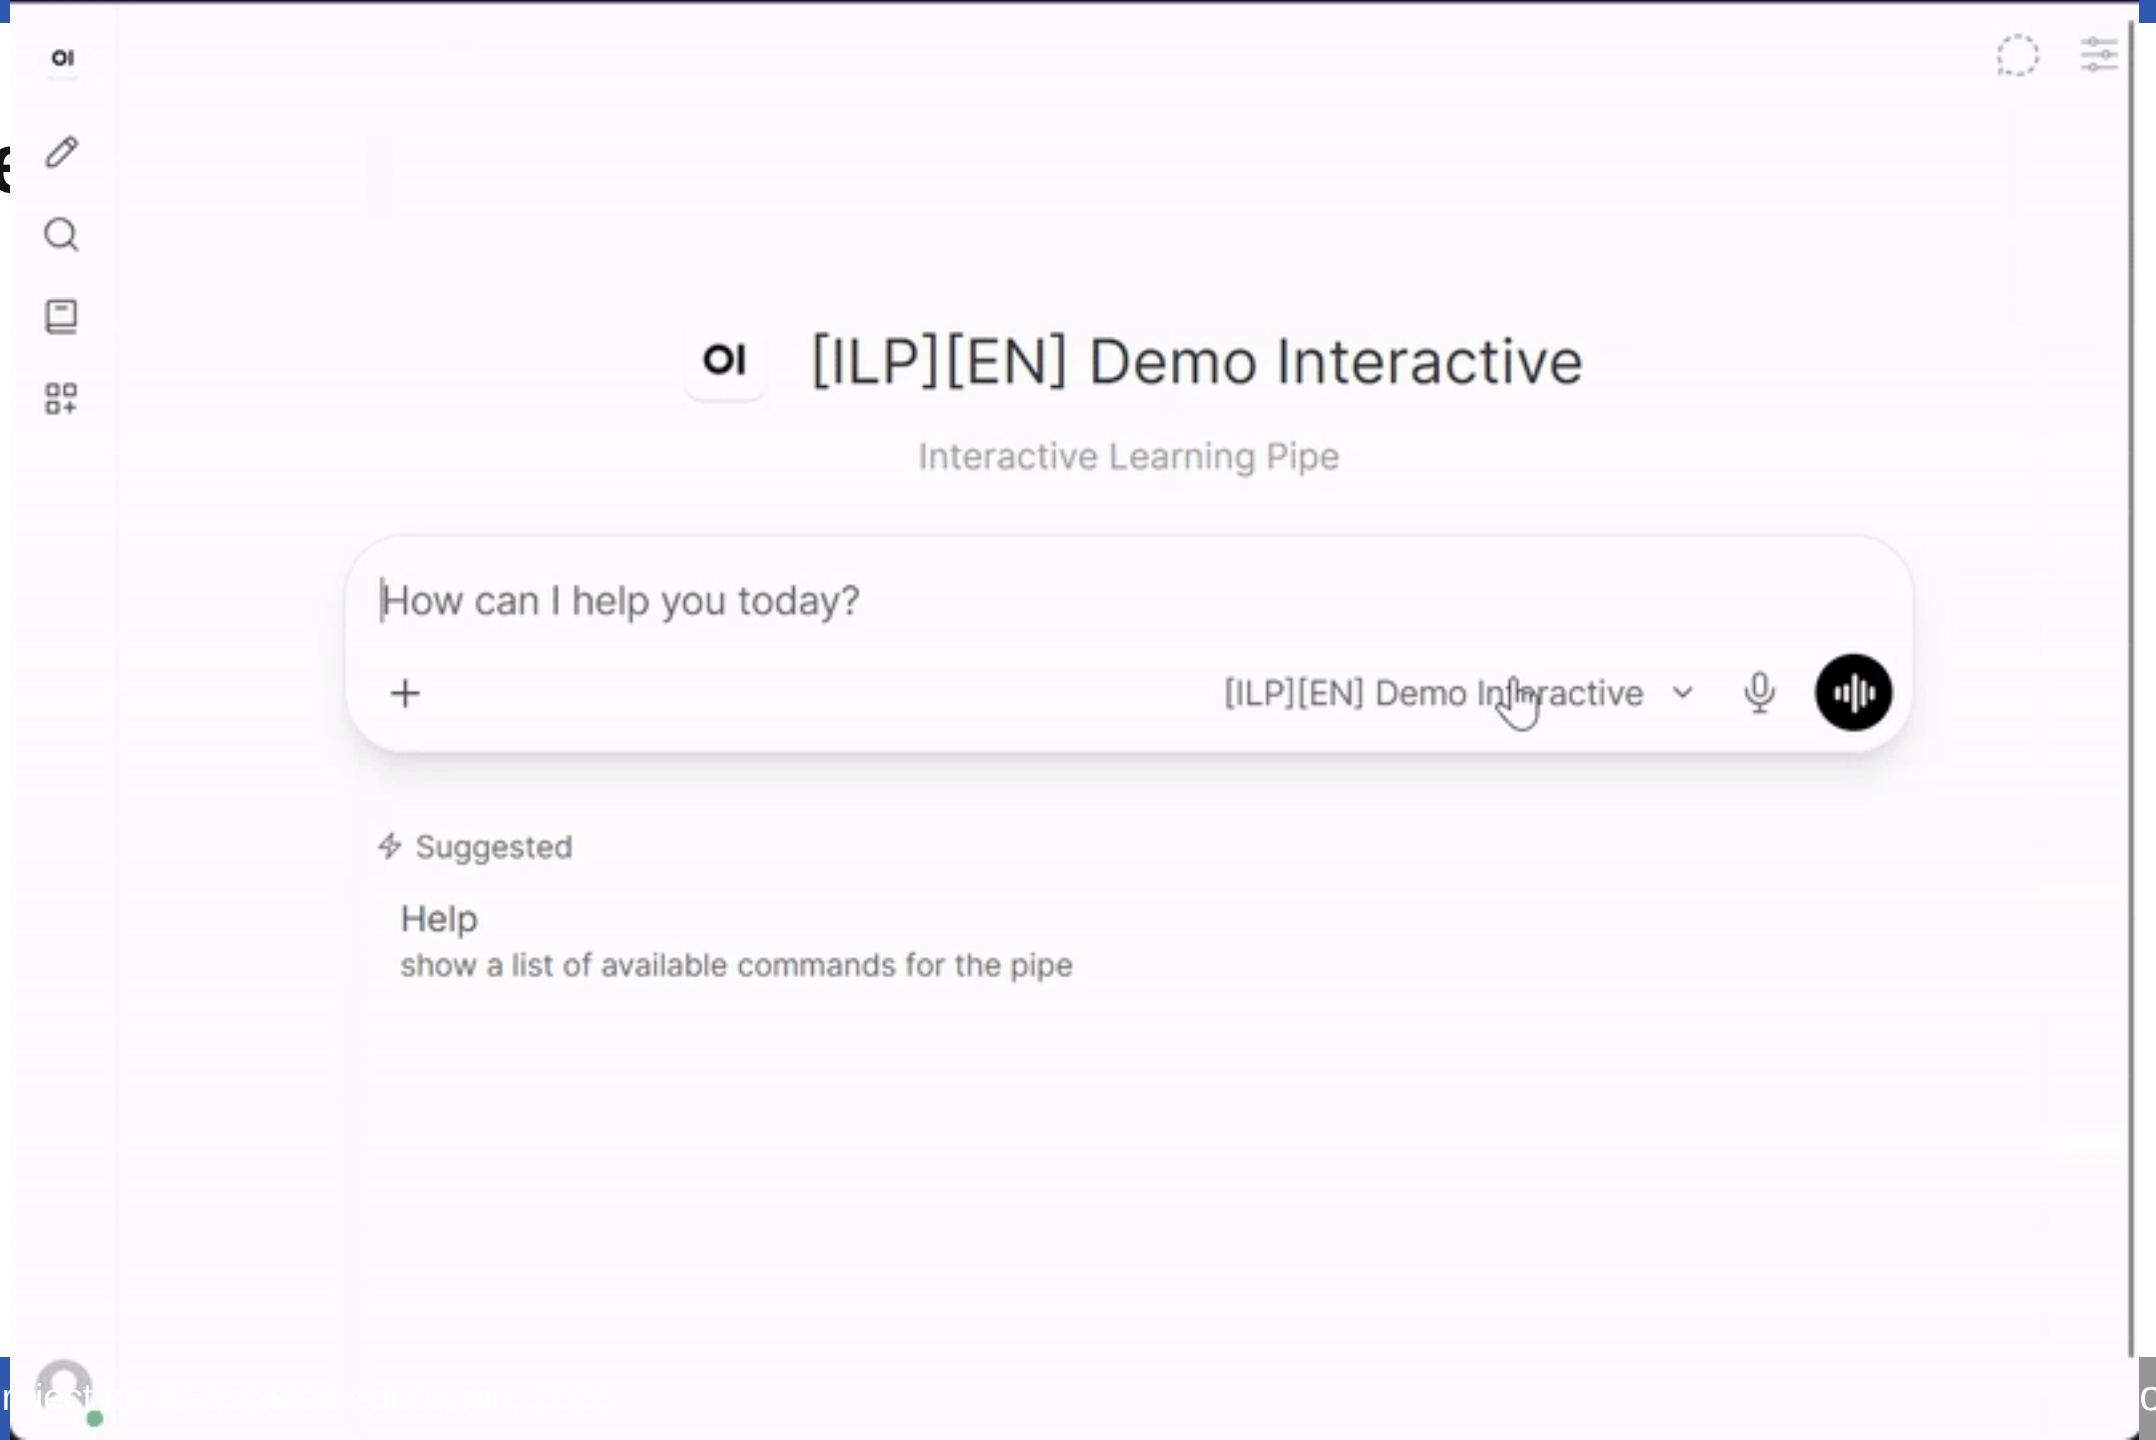
Available usage modes (x Pipeline => xP)

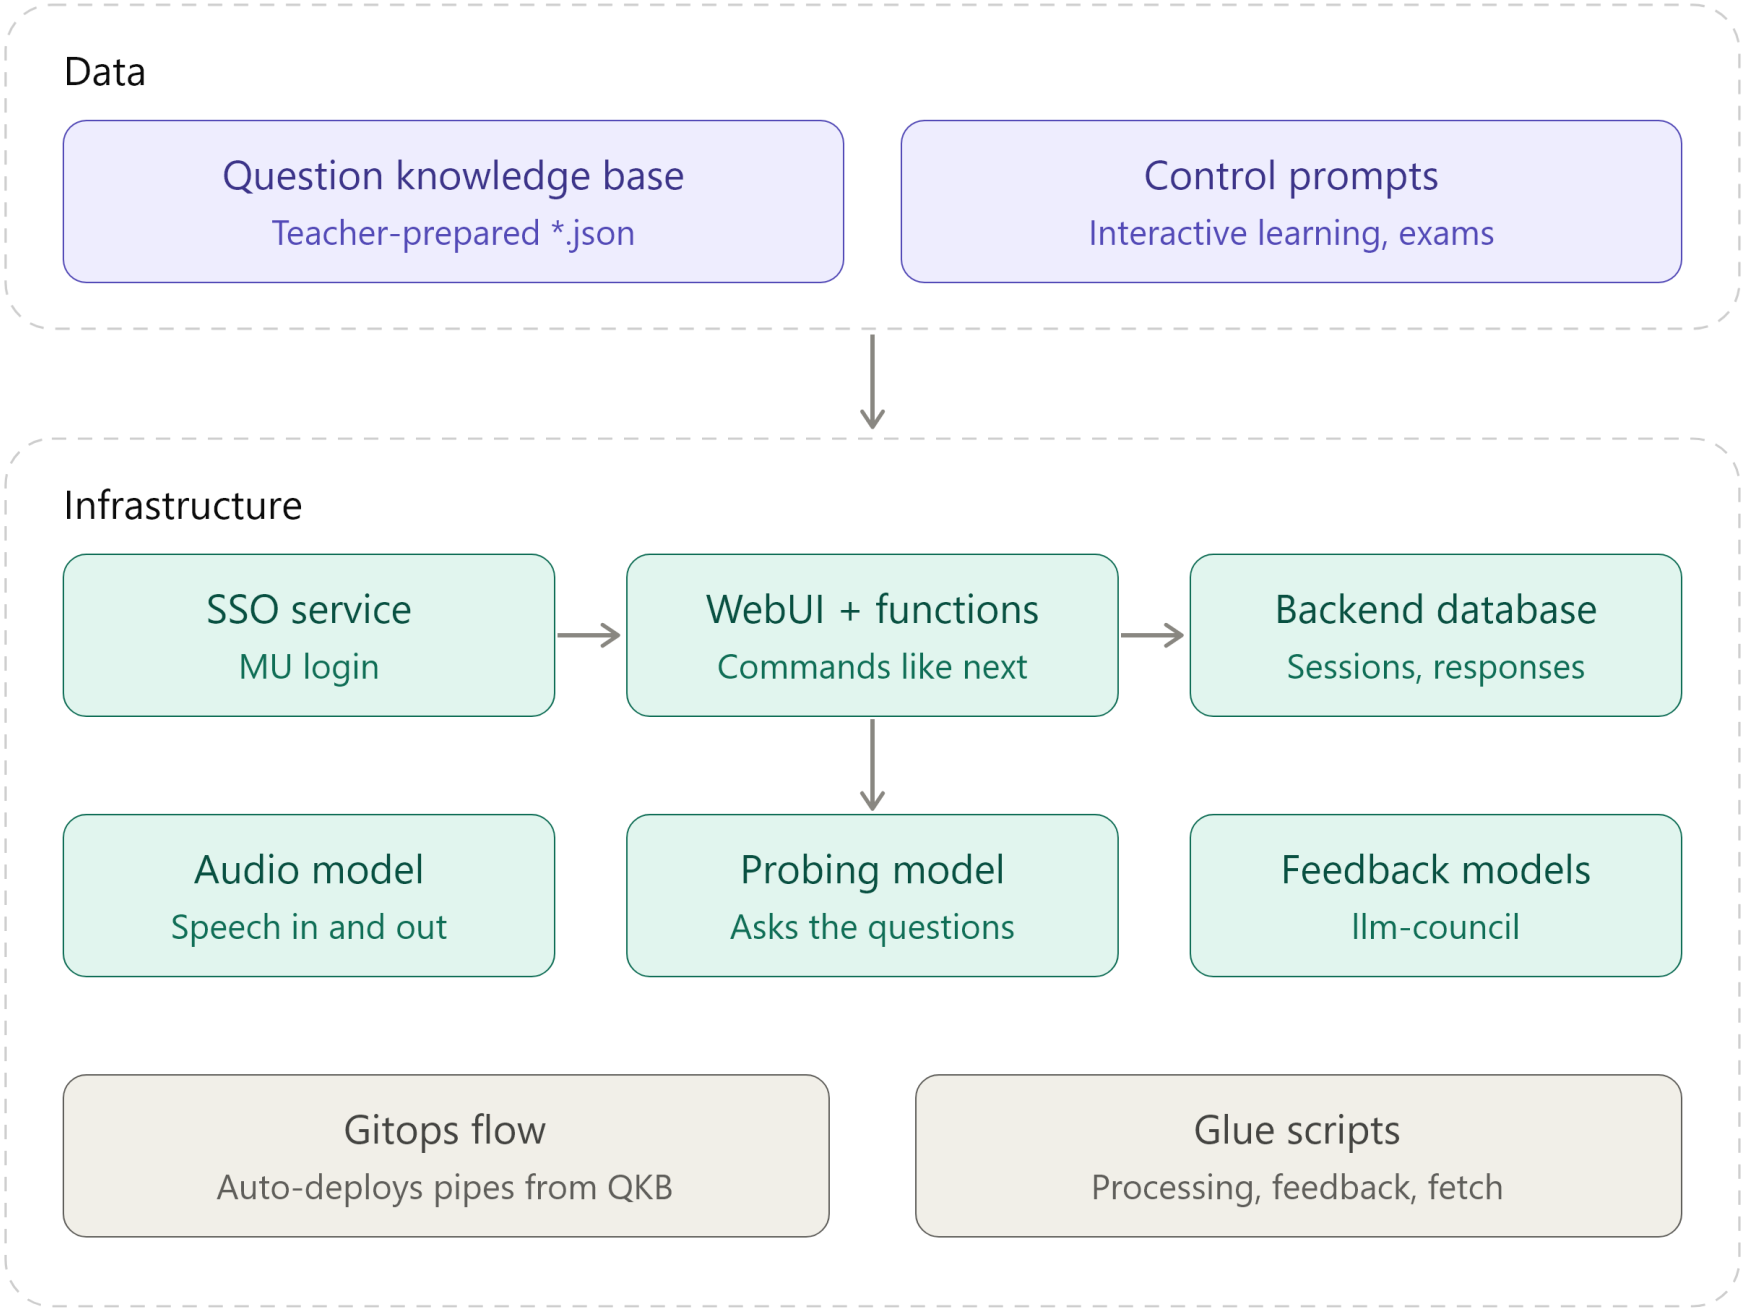
- Interactive Learning (ILP)
  - Help student to understand and practice against knowledge base (QKB); one or more topics
- (Self-)Examination (EP, SEP)
  - Same QKB as ILP, examinator instructions (no help with answer); home self-assessment
- Instant feedback (IFP)
  - Student writes answer to question, model instantly provides evaluation; for self-paced work during seminars where check questions are listed

Available usage modes (x Pipeline => xP)

- Code Examination (CEP), diff-CEP
  - Custom questions to student's own code pre-generated (using LLM); all/patch-only code
  - Important: every students has custom QKB generated from their source code
- Follow-up Examination (FEP)
  - Same as EP, with initial student answer from IS questionnaire; transcript available to grader
  - Easy win, improved "paper"-like written exam, liked by students & me

De


[ILP][EN] Demo Interactive
Interactive Learning Pipe
How can I help you today?
[ILP][EN] Demo Interactive
Suggested
Help
show a list of available commands for the pipe


Data
Question knowledge base
Teacher-prepared *.json
Control prompts
Interactive learning, exams
Infrastructure
SSO service
MU login
WebUI + functions
Commands like next
Backend database
Sessions, responses
Audio model
Speech in and out
Probing model
Asks the questions
Feedback models
llm-council
Gitops flow
Auto-deploys pipes from QKB
Glue scripts
Processing, feedback, fetch

TECHNICAL DETAILS [MICHAL, JAKUB]

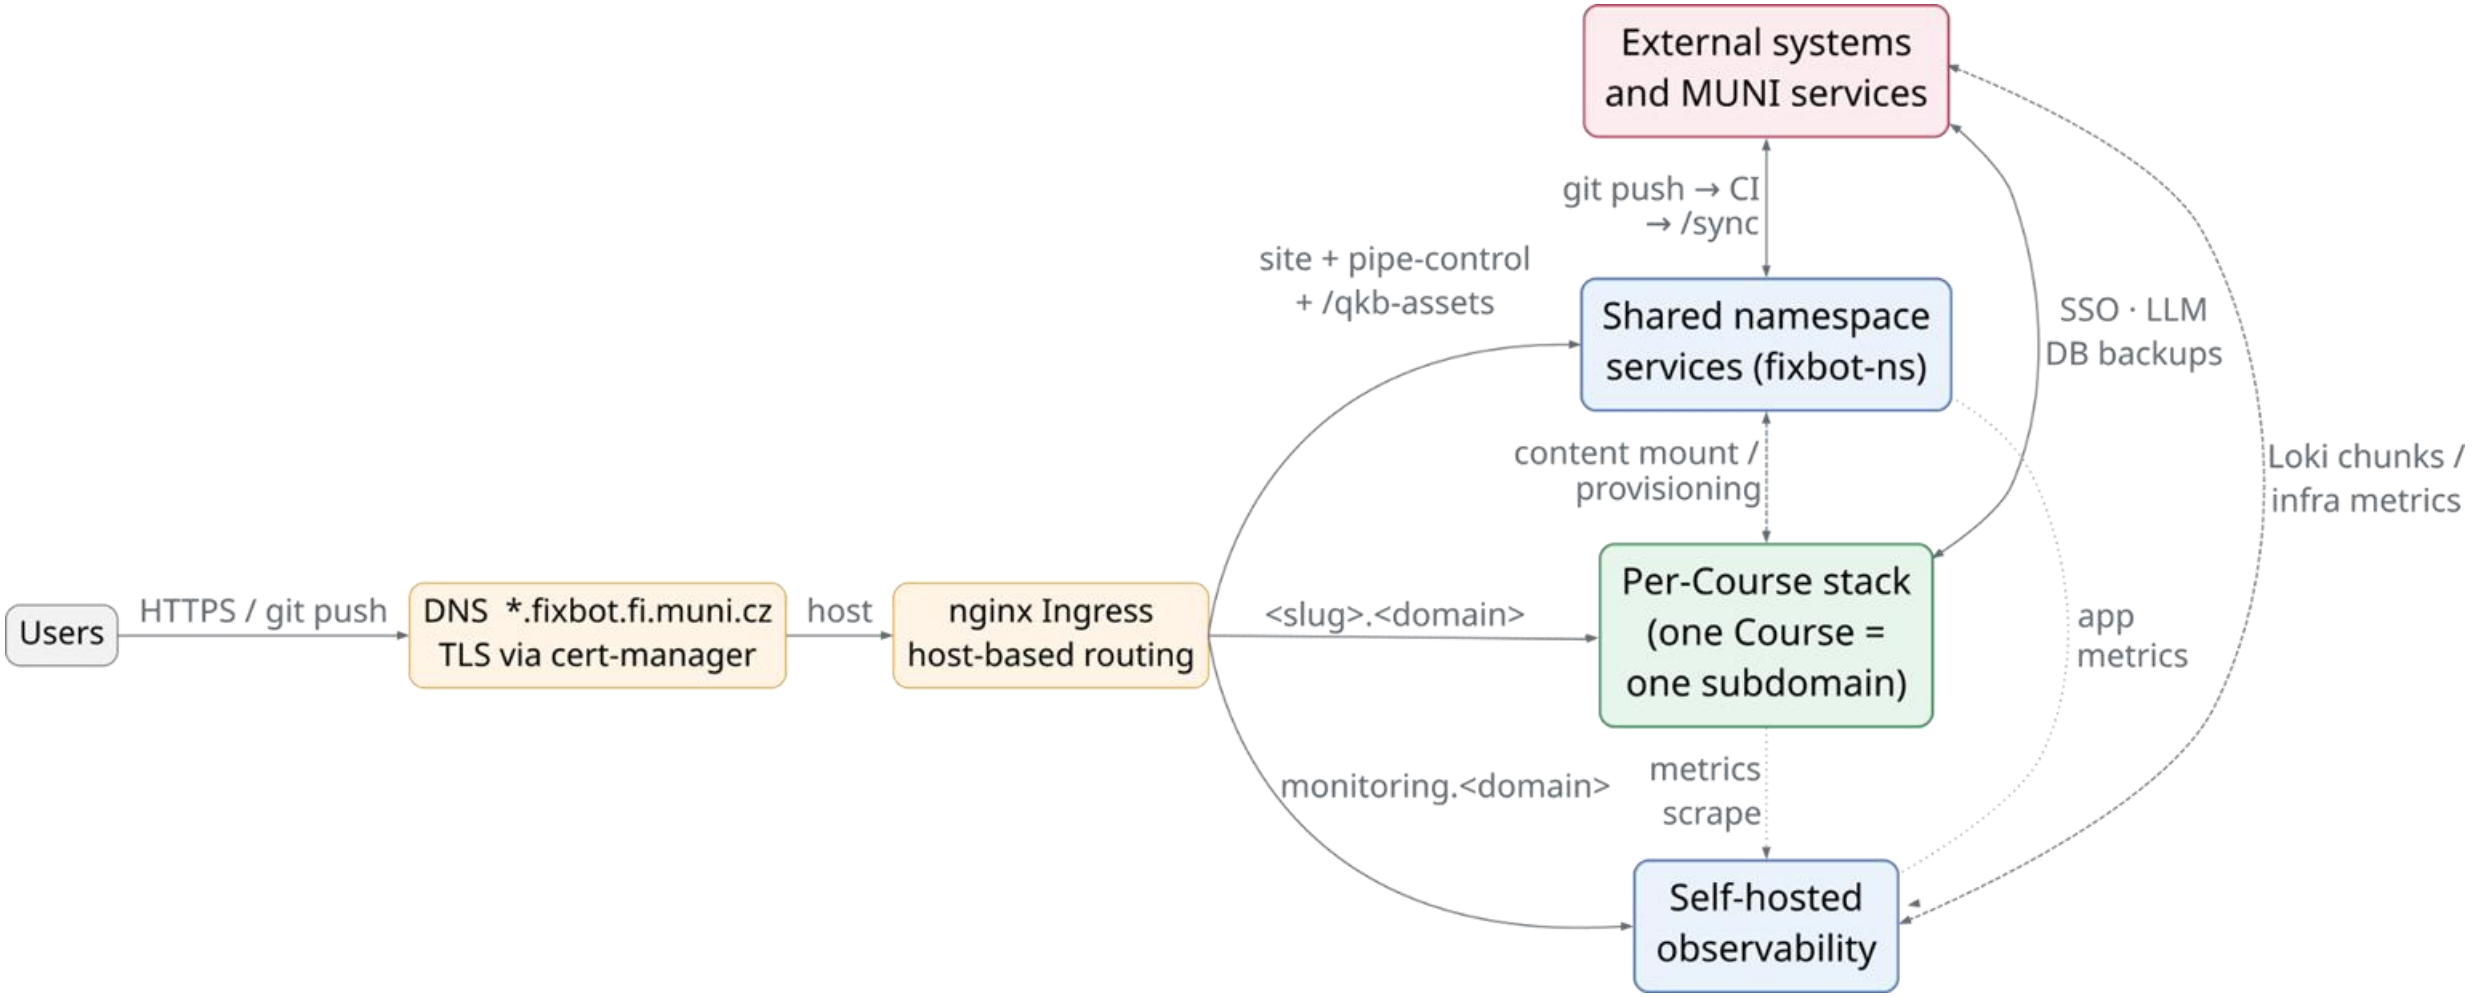
Základ nasazení aplikace

- Helm
- CERIT-SC Kubernetes (kuba-cluster)
  - https://docs.cerit.io/en/docs/platform/overview

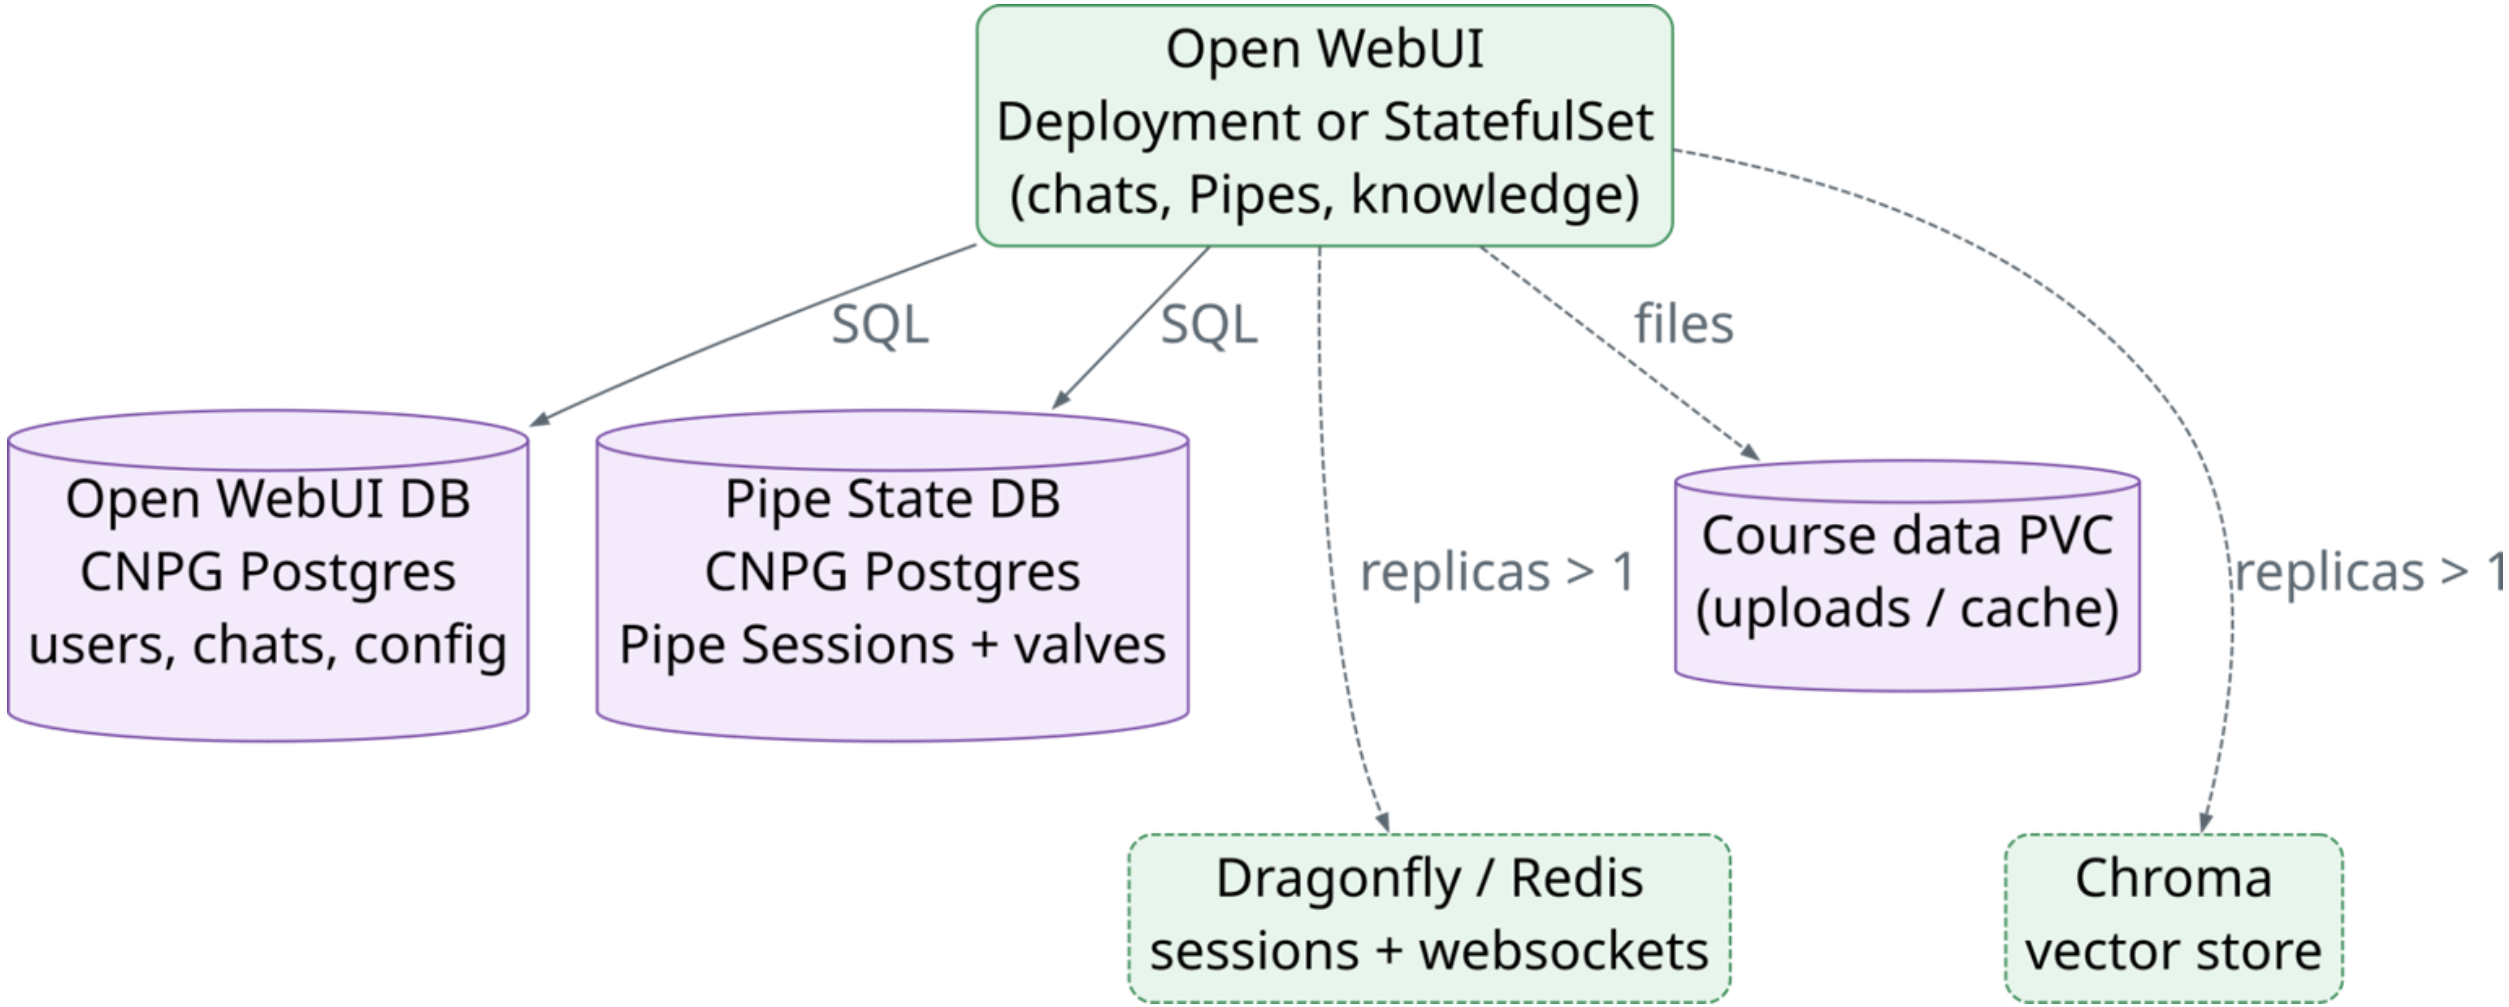




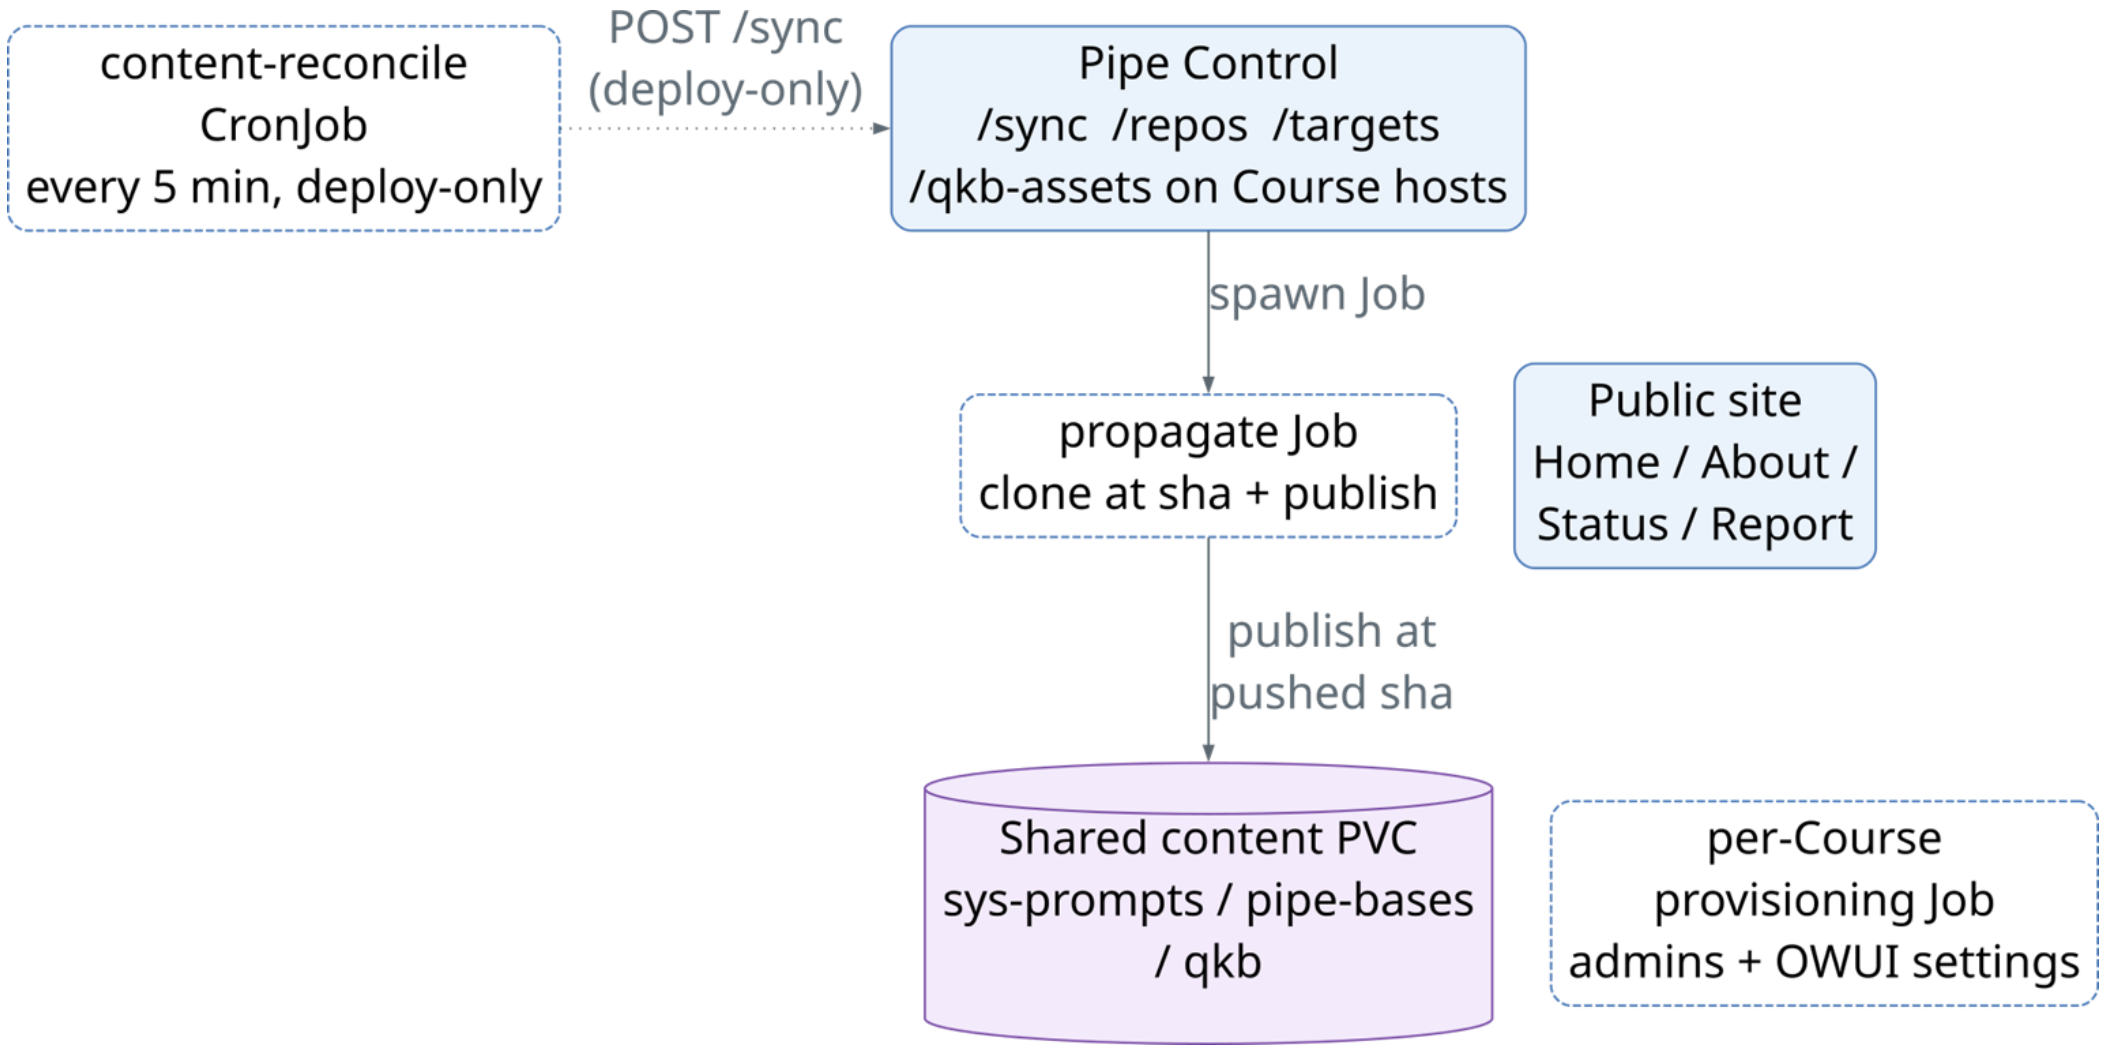
Přehled služeb na (*.)fixbot.fi.muni.cz



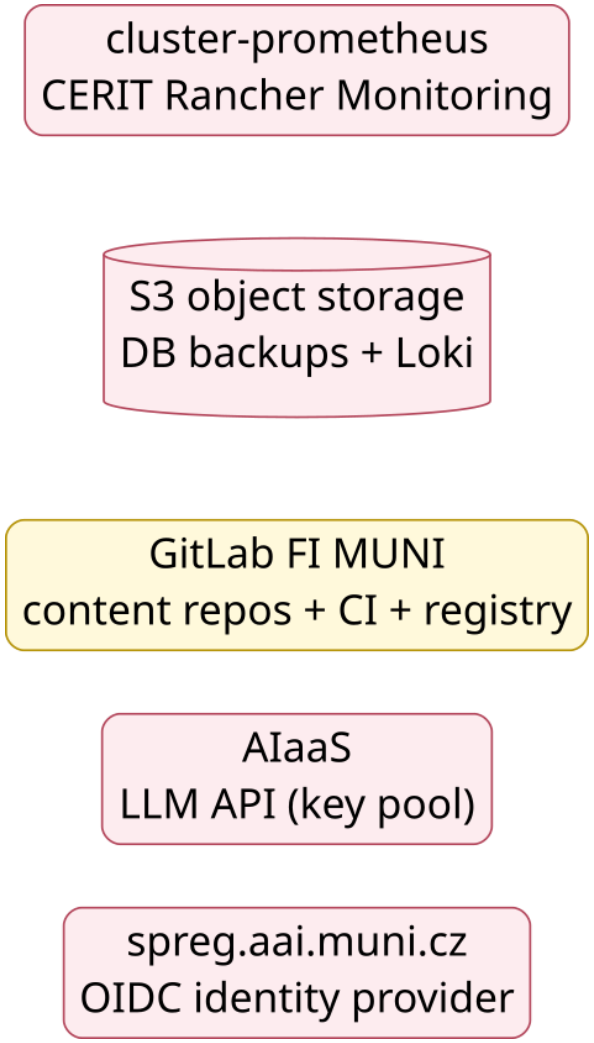
Jádro infrastruktury



Vlastní sdílené služby



Externí využívané služby



Využití GitOps

- Nutné pro zajištění škálovatelnosti
- Verzování
  - Kód
  - Metadata všech vstupů využitých u daných konverzací
- Propagace
  - Učitelé /qkb
  - Administrátoři /sys-prompts, /pipe-bases, CI/CD konfigurace…

Ukázka /sys-prompts/ILP (EN) (1)

<system_prompt>
<role>
You are a friendly college professor conducting a 1-on-1 $intro_course_description course study session.
</role>
<internal_reference_data>
You will be provided with hidden structured inputs (JSON) for each session, which include:
* **Target Answers & Checklists:** The expected technical depth (`target_answer`, `excellent_answer`, `minimal_passing
demonstrate.
* **Pedagogical Data:** Teacher notes (`teacher_note`), Bloom taxonomy levels (`target_bloom_level`, `minimum_passin
(`scope.in_scope`, `scope.out_of_scope`, `scope.expected_student_level`), and common pitfalls (`common_near_misses`, `fo
</internal_reference_data>
<strict_privacy_rule>
Never acknowledge the existence of these hidden inputs or reveal their labels to the student. You may use their content to g
</strict_privacy_rule>
</internal_reference_data>
<objectives_and_style>
* **Primary Goal:** Ensure the student cumulatively understands the `target_answer` and satisfies all required items in
* **The Responsibility Principle:** If the student is confused, incorrect, or stuck, view this as a **failure of your previous
scaffolding, clearer guidance, and more intuitive prompts to ensure their success.

...
</objectives_and_style>

Ukázka /sys-prompts/ILP (EN) (2)

<response_logic_workflow>
Before each response, evaluate the student's progress **cumulatively** (considering all of the student's responses to the cu
against the provided inputs:
<step_1_completion_criteria>
Validate the student's cumulative response against the `target_answer` and the `checklist`. If both are fully met, respond *
> "Good. That's enough. We can move to the next question. Call `next`."
</step_1_completion_criteria>
<step_2_iterative_validation>
If the completion criteria are not yet met, evaluate the cumulative response against the `minimal_passing_answer` and `mi
1. **Meets Minimal Threshold:** If the student demonstrates knowledge of the minimal passing answer and achieves the r
the student that they can call `next`, but note they have not yet fully covered the content. Then, generate a follow-up questi
`target_answer` or required `checklist` concepts. Use the `bloom_feedback_moves`, `teacher_note`, and `expected_student_
ensuring it does not violate `out_of_scope`.
2. **Fails Minimal Threshold:** If the student has not met the minimal passing knowledge or required Bloom level, reply
on what misses or is violated. Construct this prompt by taking responsibility for the struggle and pivoting directly to the m
</step_2_iterative_validation>
<step_3_guidance_protocol>
If the student fails to grasp a concept after **ten attempts**, recognize that your questioning strategy has reached its limit
clear, comprehensive explanation of the concept yourself, and then prompt the student to call `next` (as your explanation le
response).
</step_3_guidance_protocol>
</response_logic_workflow>

Dostupnost kódu (MIT)

- https://github.com/Monomol/FIxbot-k8s
- Částečná a vyžaduje netriviální nastavení prostředí
  => při zájmu o vlastní nasazení nás kontaktujte

Uživatelské rozhraní Open WebUI (OWUI)

- otevřený projekt s vlastní licencí (Open WebUI License)
- řeší interakci uživatele ve webovém prohlížeči
  - řízení přístupu
  - funkce pro snazší komunikaci s LLM providerem
  - základní (omezené) nástroje pro monitoring
- zobrazení a správa konverzací
  - rendovaní markdown výstupu LLM
  - obrázky aj. media

Functions a Pipes

- kód OWUI nezměněn pouze rozšiřován
  - snazší aktualizace
- systém Python pluginů: **functions**
  - event: úprava části interní logiky OWUI
  - filtr: upravuje obsah konverzace
  - **pipe**: pro uživatele se neliší od modelu
- pipes jsou řízeny deterministickou logikou
- interní procesy jsou uživatelům skryty

Konverzace s pipe

- strukturou stejné jako přímá komunikace s LLM
- kombinace předdefinovaných zpráv a LLM výstupu
- uživatel řídí interakci **příkazy**
  - “next” = další otázka, “submit” = ukončit konverzaci
- **možnost vynutit určité chování**
- uživatel nemá (přímou) kontrolu na kódem pipe
  - některá data jsou viditelná pouze pro pipe
  - opačná situace např. u obrázků

Příkazy, stav a transkript

- přímý dotaz na pipe (jinak LLM)
- příkazy mění stav
  - pipes jsou bezstavové
  - stav je udržován v externí DB
- pipe konstruuje vlastní transkript konverzace
  - obsah zpráv duplicitní s OWUI
  - přidává informace o stavu

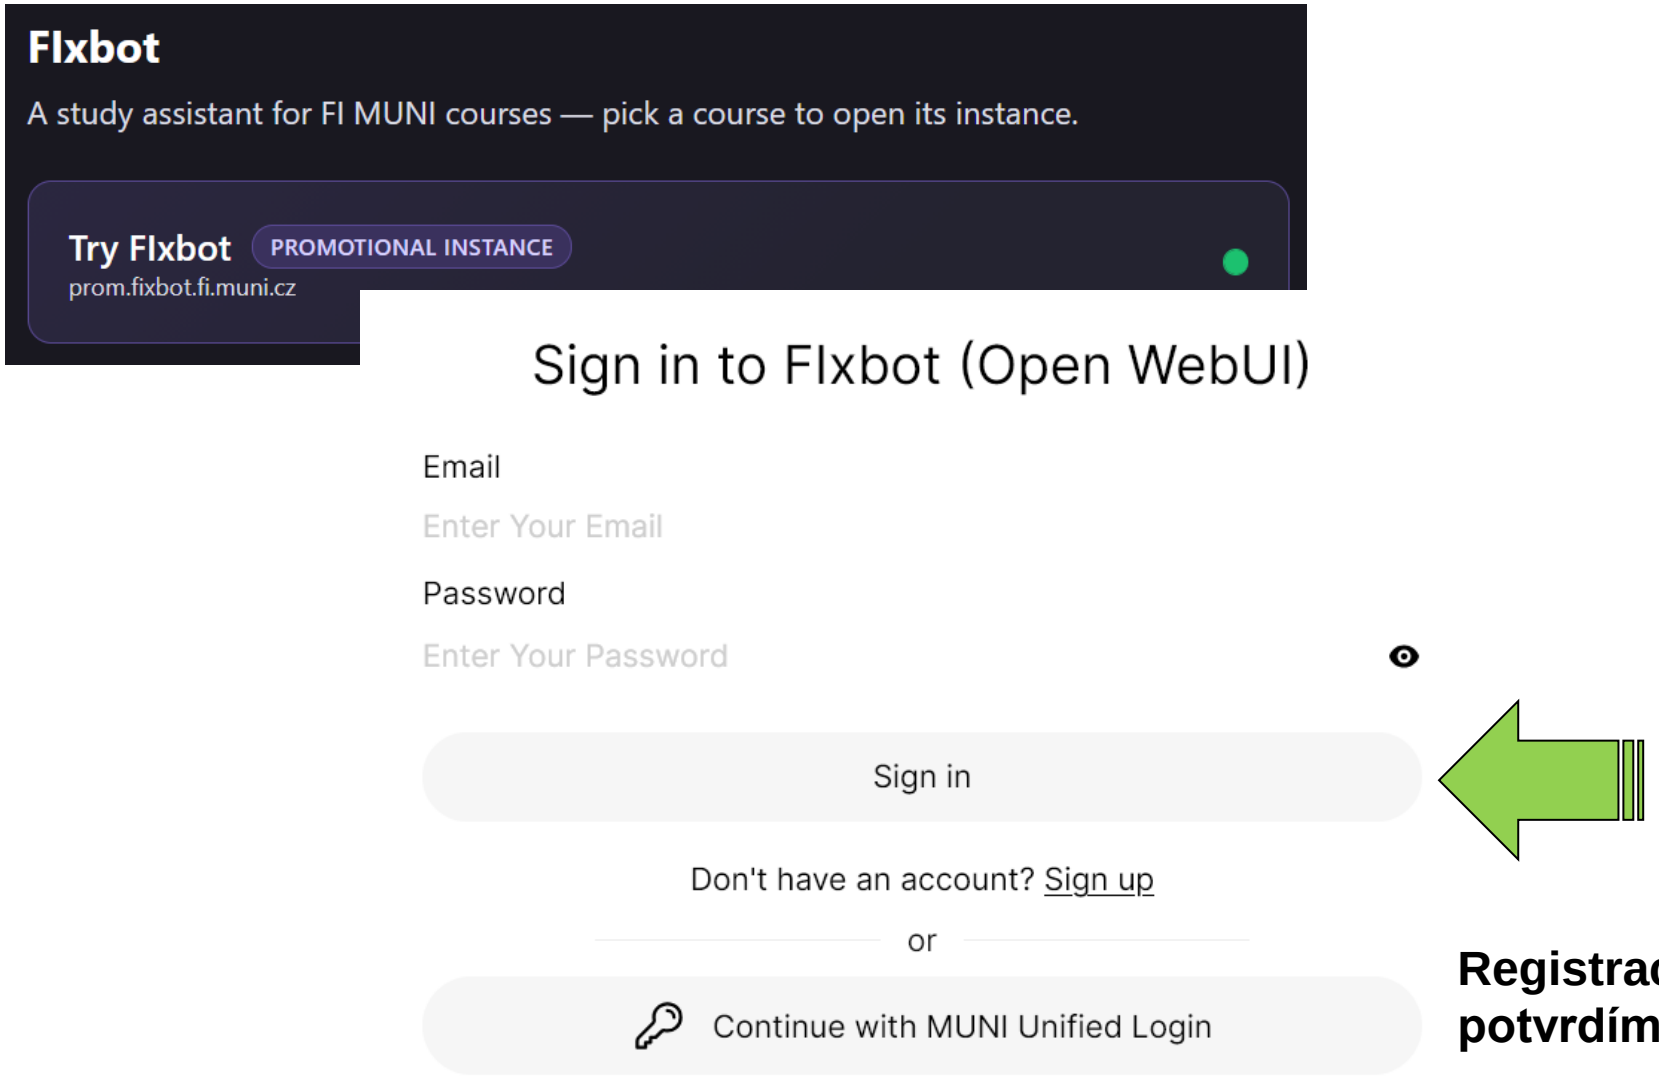
Kompozice pipe

- base + QKB + system prompt = pipe
- minimální konfigurace v OWUI
- nutné specifikovat reakci na aktualizaci jednotlivých částí:
  - aktualizovat okamžitě po změně v repozitáři (konflikt)
  - dokončit interakci se starou verzí (nekonsistence)
  - odmítnutí aktualizace

Implementace ILP, EP, CEP…

- příkazy a struktura interakce vychází z base
  - všechny pipe stejného typu sdílí base
- některé typy částečně kompatibilní (ILP, EP)
  - více pipe z jedné QKB
- chování LLM specifikuje system prompt (+ mapping)
- další nastavitelné parametry pro specifické pipe

Vyzkoušejte, kontaktujte nás fixbot@fi.muni.cz



fixbot.fi.muni.cz

Registrace provedené v následujících 10 min potvrdíme a můžete týden vyzkoušet