



Strategies for movement


Section 3

Why do we move data and computation?


We know how to move data and computation..

Why do we move them?

Case study description



Case study description



Why do we move data and computation?

We know how to move data and computation..

Why do we move them?

To satisfy requirements of users:

- Latency
- Completeness
- Length encryption key
- Location of data
- ...

When do we move data?

Two options

- In advance:
 - Proactive behavior
- When a requirement is not satisfied
 - Reactive behavior

Reactive behavior: violations management

Definition of violation: one or more requirement is not satisfied.

Examples:

Requirements	Violations
Latency < 5ms	Latency 20 ms
Completeness > 80%	Completeness = 70%
Encryption key length > 256	Encryption key length > 128

Tolerance accepted, but have to be agreed with users

Definition of requirements

Requirements is a singular documented physical or functional need (of users) that a particular design, product or process aims to satisfy.

Data utility

- Quality of Service (QoS)
- Quality of Data

Security

Privacy

The elicitation of requirements

It's difficult for users to specify which requirements they needs.

They require a certain level of abstraction that can change based on their expertise

Es:

- Latency < 20 ms
- Responsive system
- Fast system

Goal-based modelling notation


Based on:

- Goals: objectives to be achieved
- A tree-like structure: to organize goals

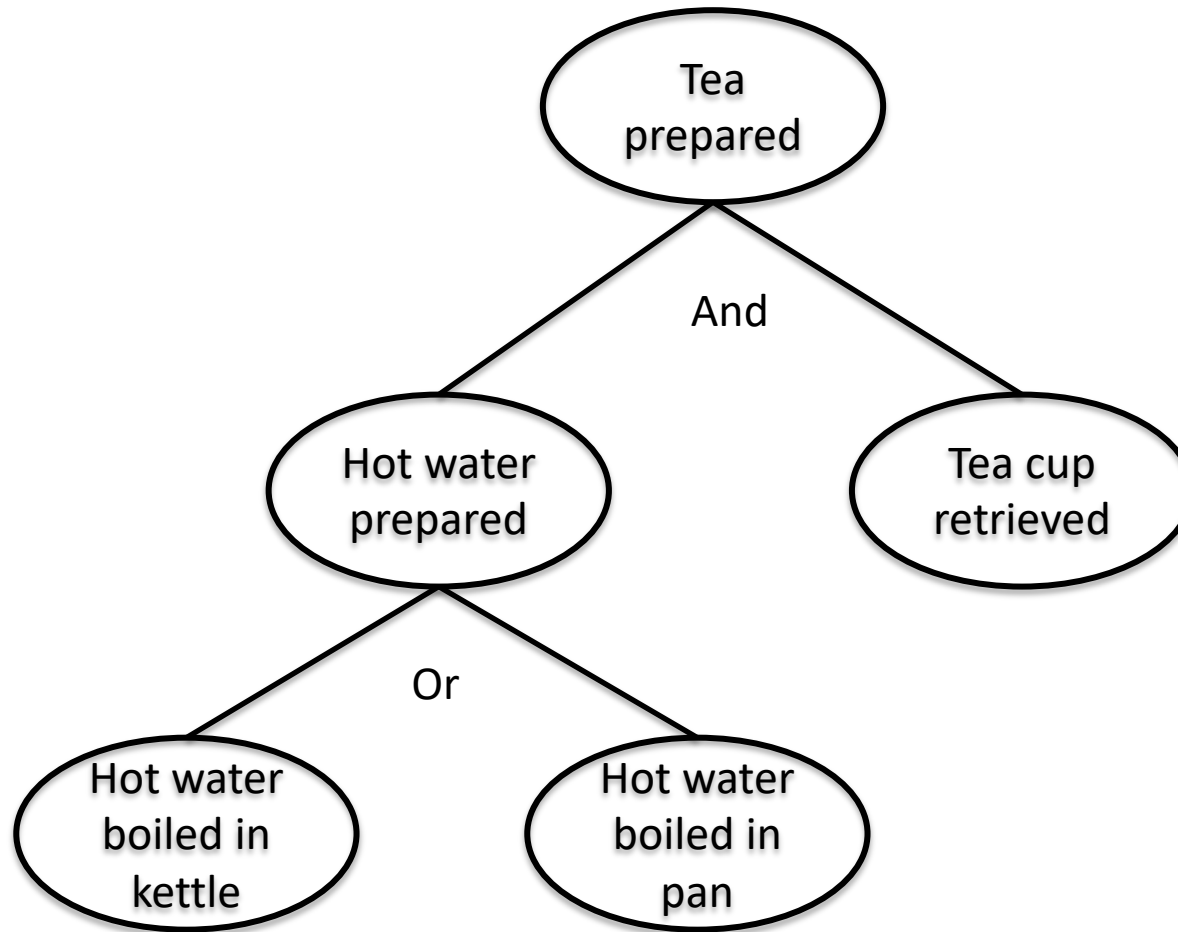
Many goal-based modelling language

- We extended Business Intelligence Model (BIM) [1]

And decomposition



Or decomposition



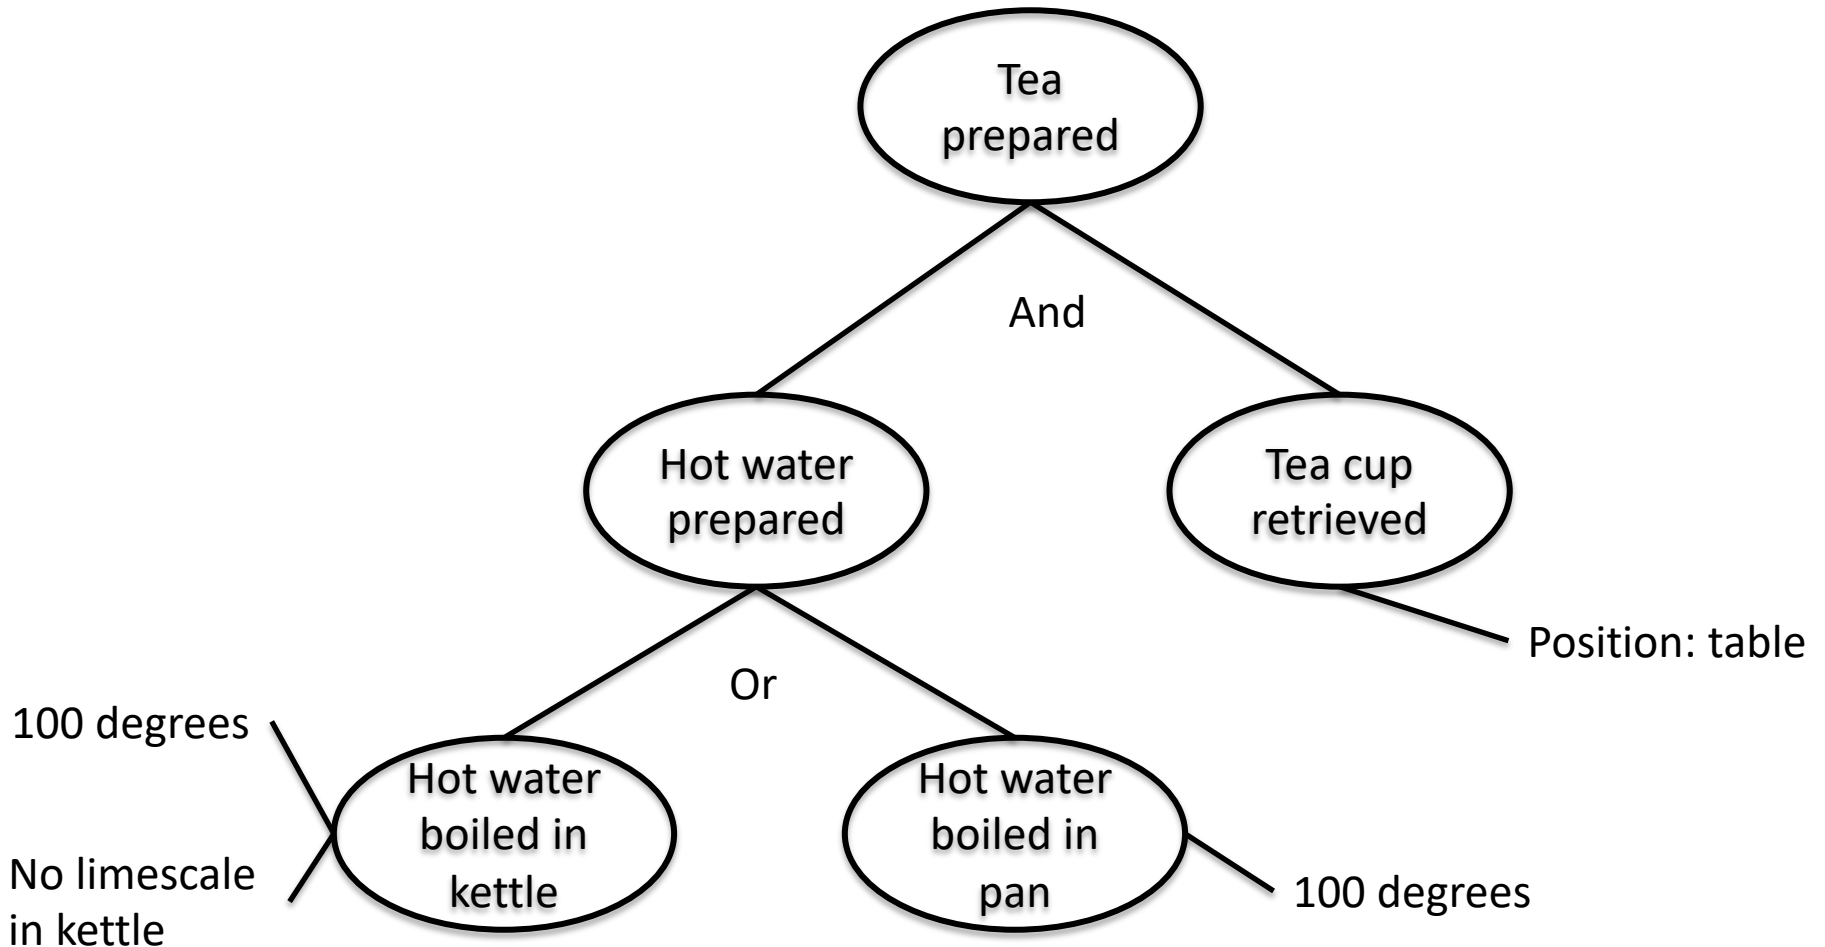
Decompositions rules

- Only one type of decomposition per goal
- Minimum 2 sub-goals
- No upper bound


Key Performance Indicator (KPI)

- Specifies a metric and a threshold
- Attached to a goal
- When all KPIs attached to a goal are satisfied, then the goal is considered reached (satisfied)

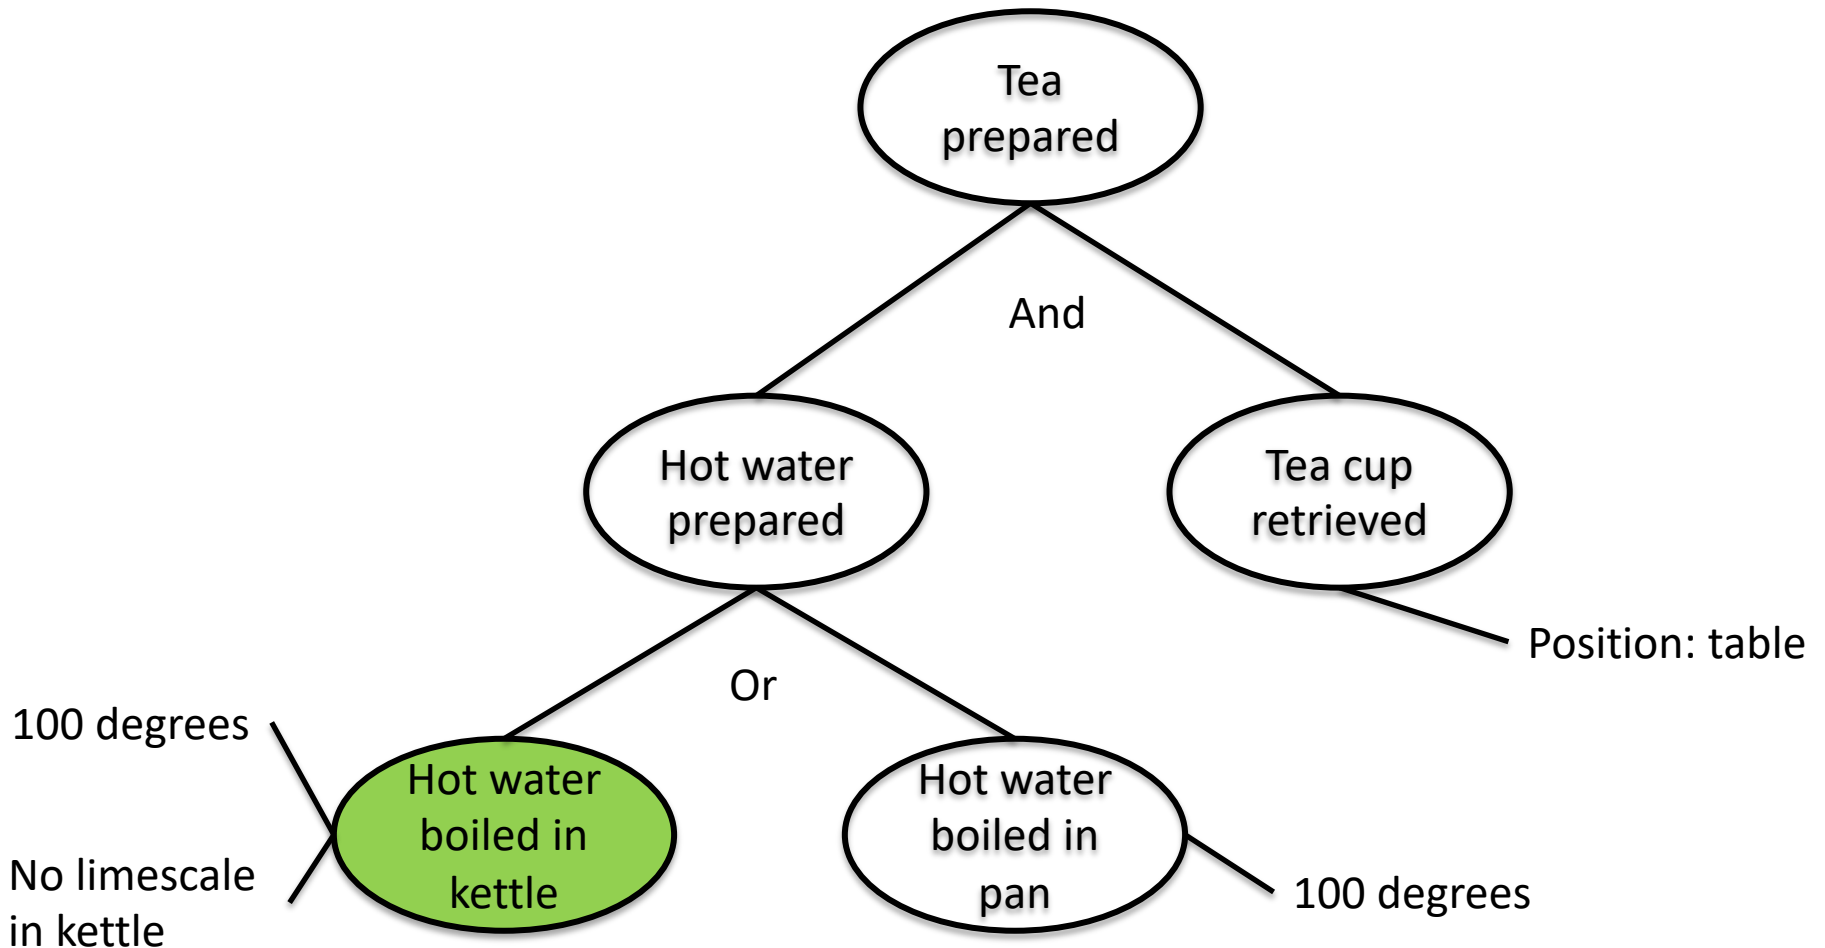
Or decomposition



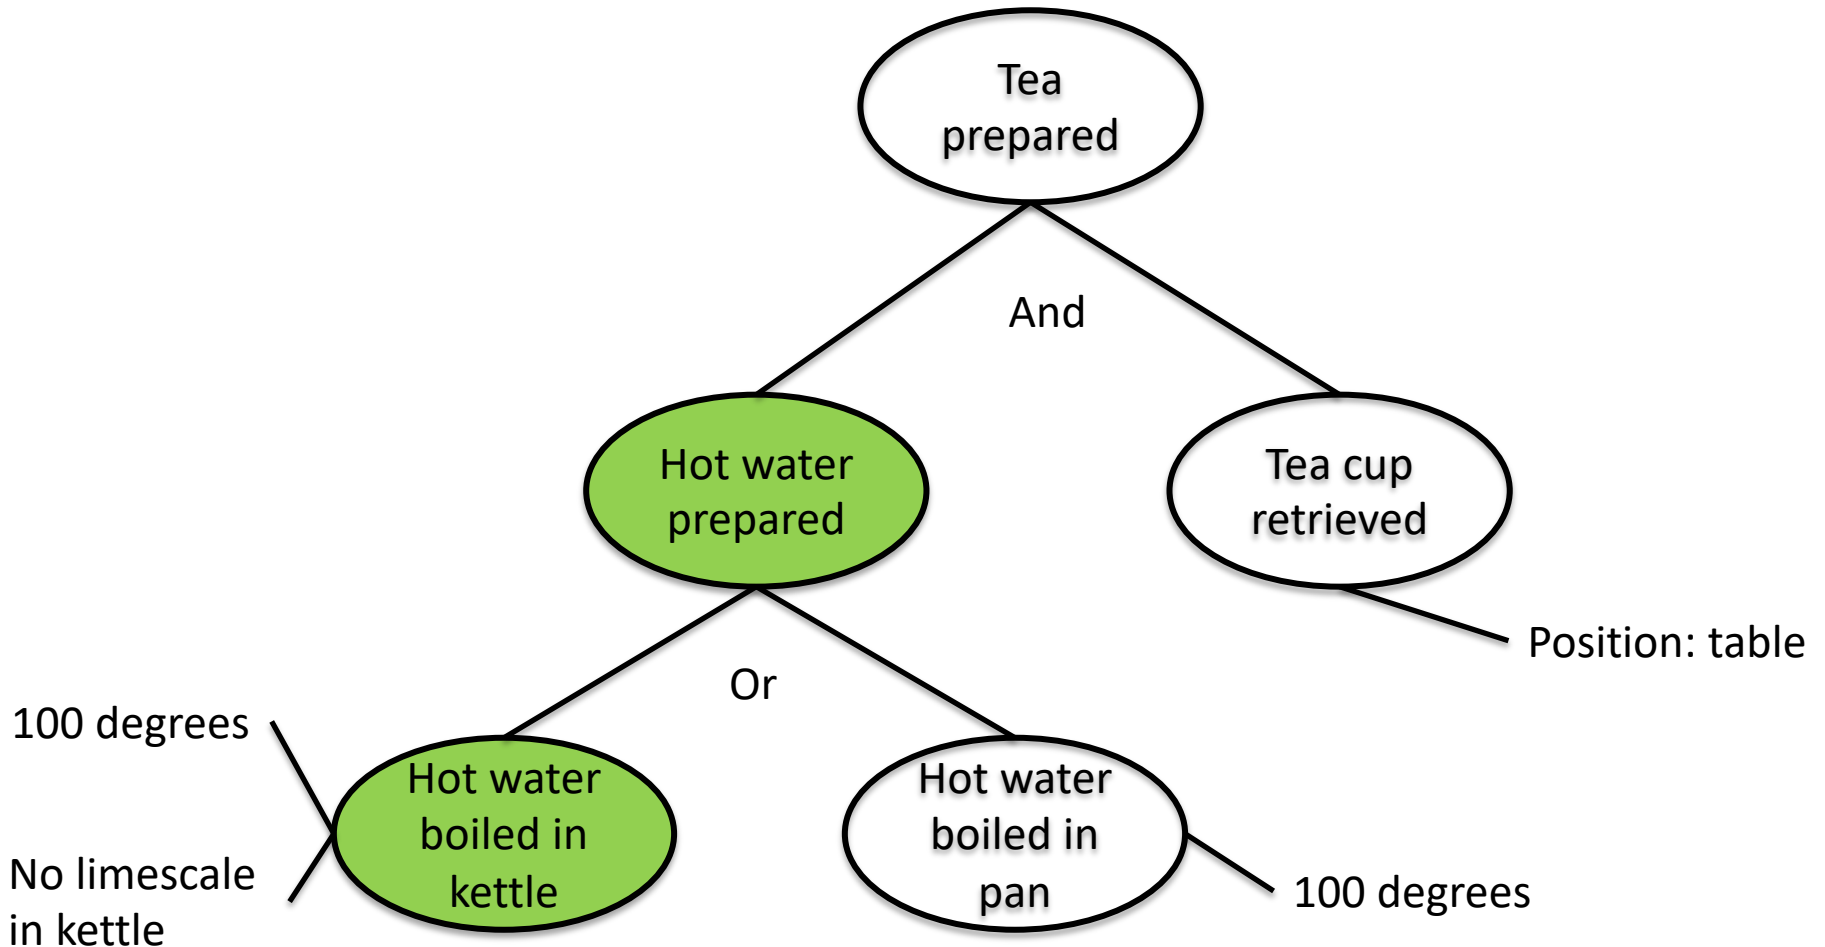
Or decomposition



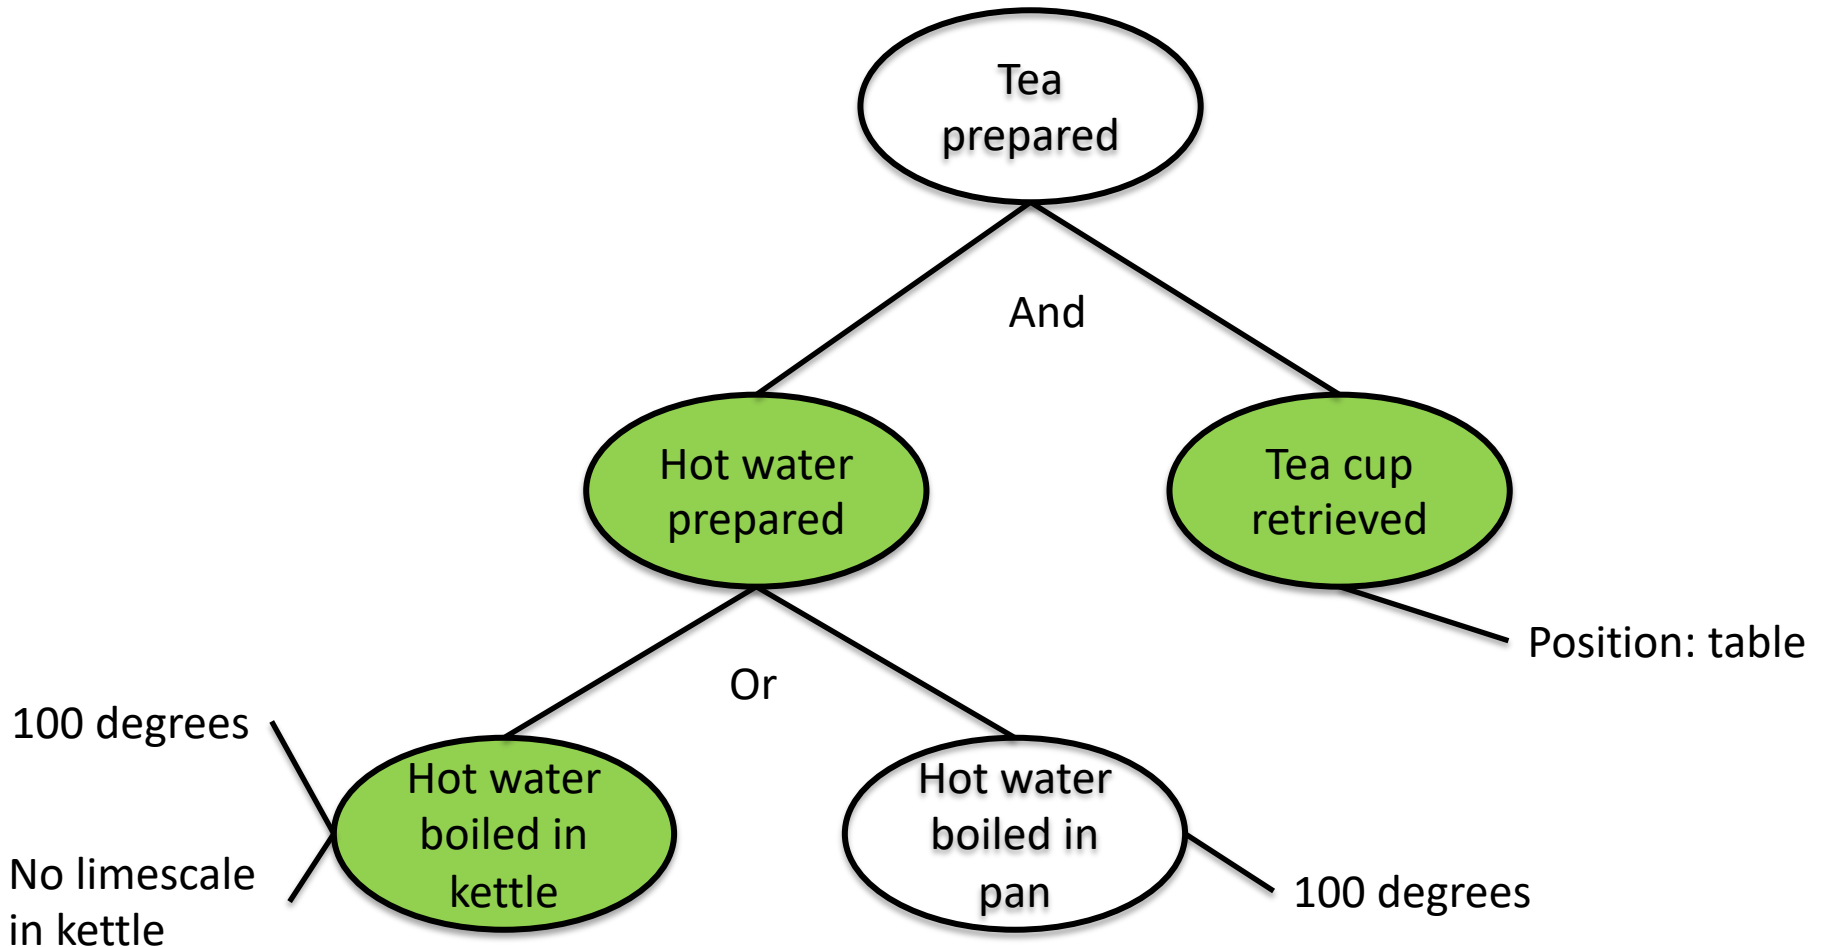
Or decomposition



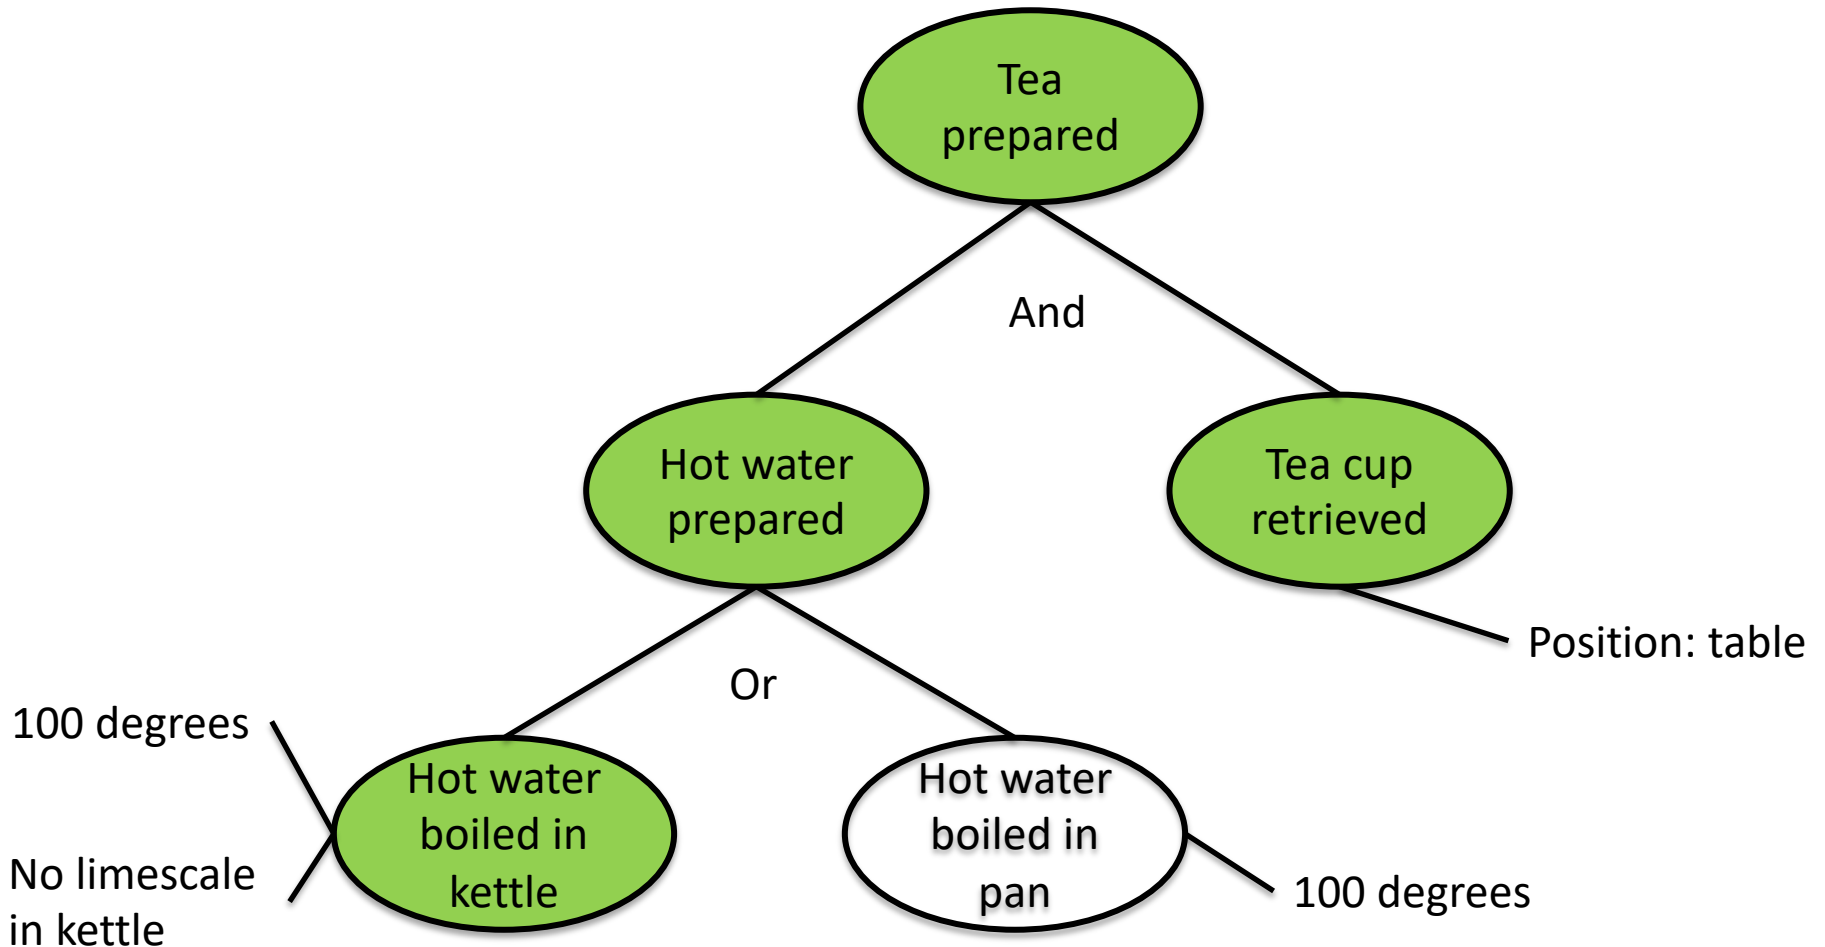
Or decomposition



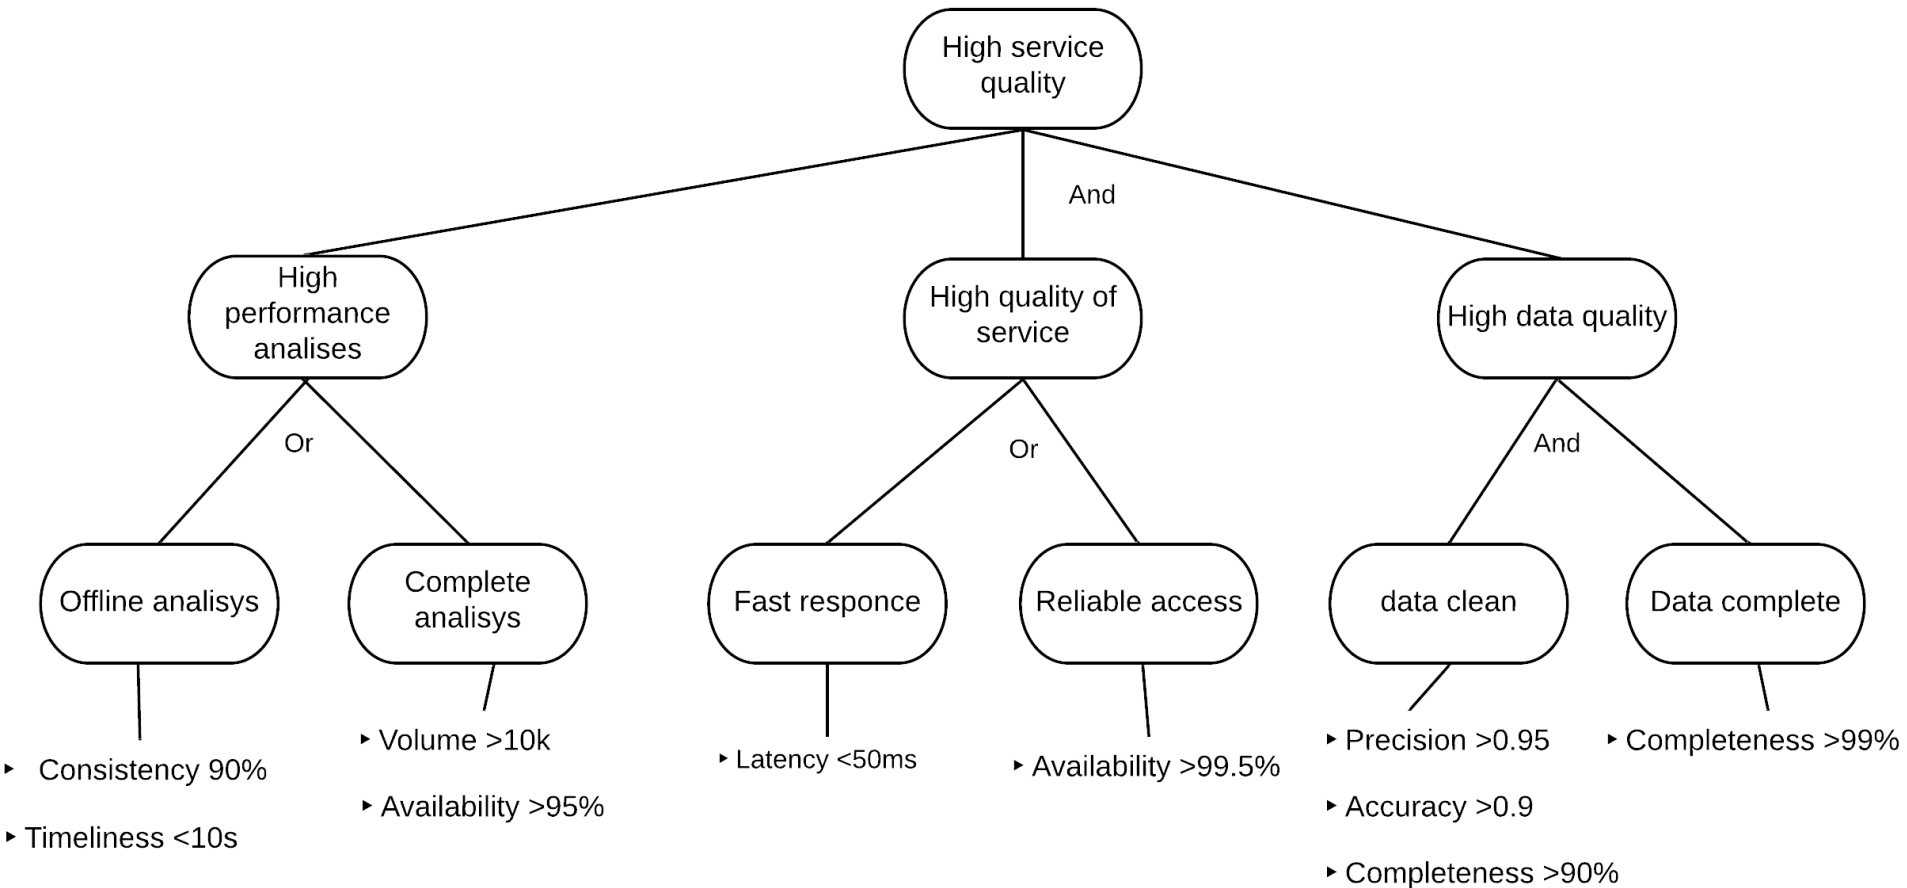
Or decomposition



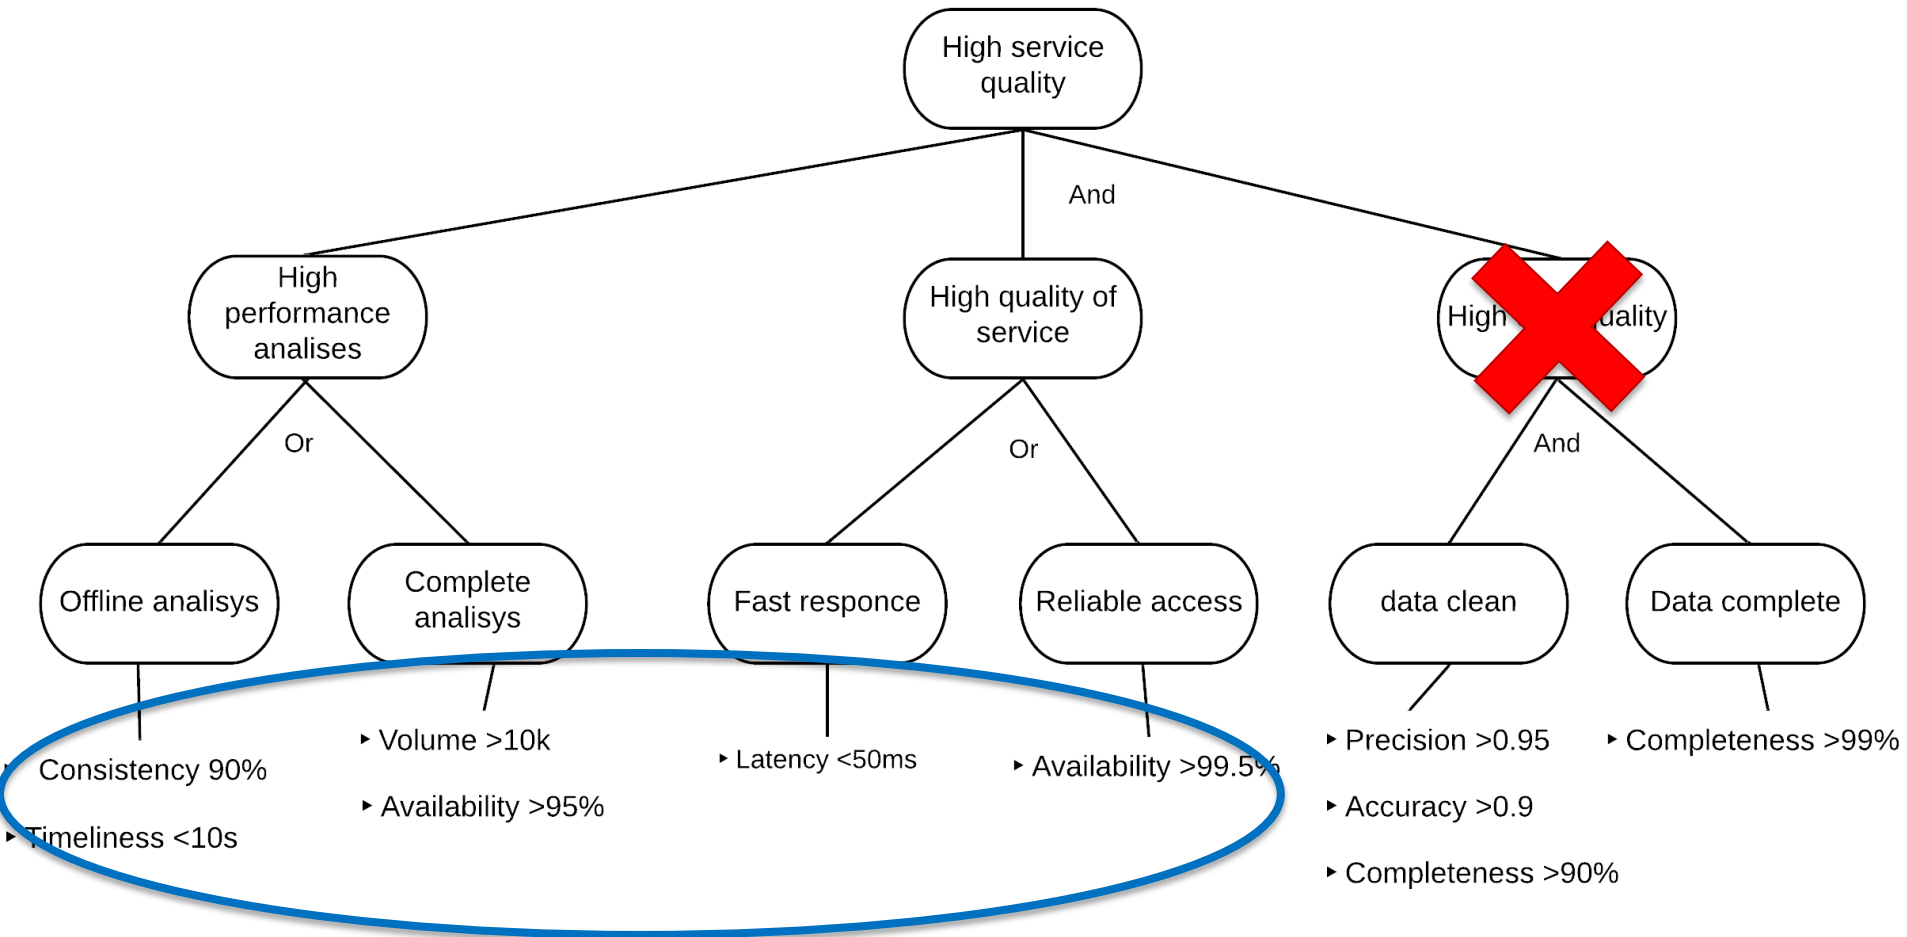
Or decomposition



A more realistic example



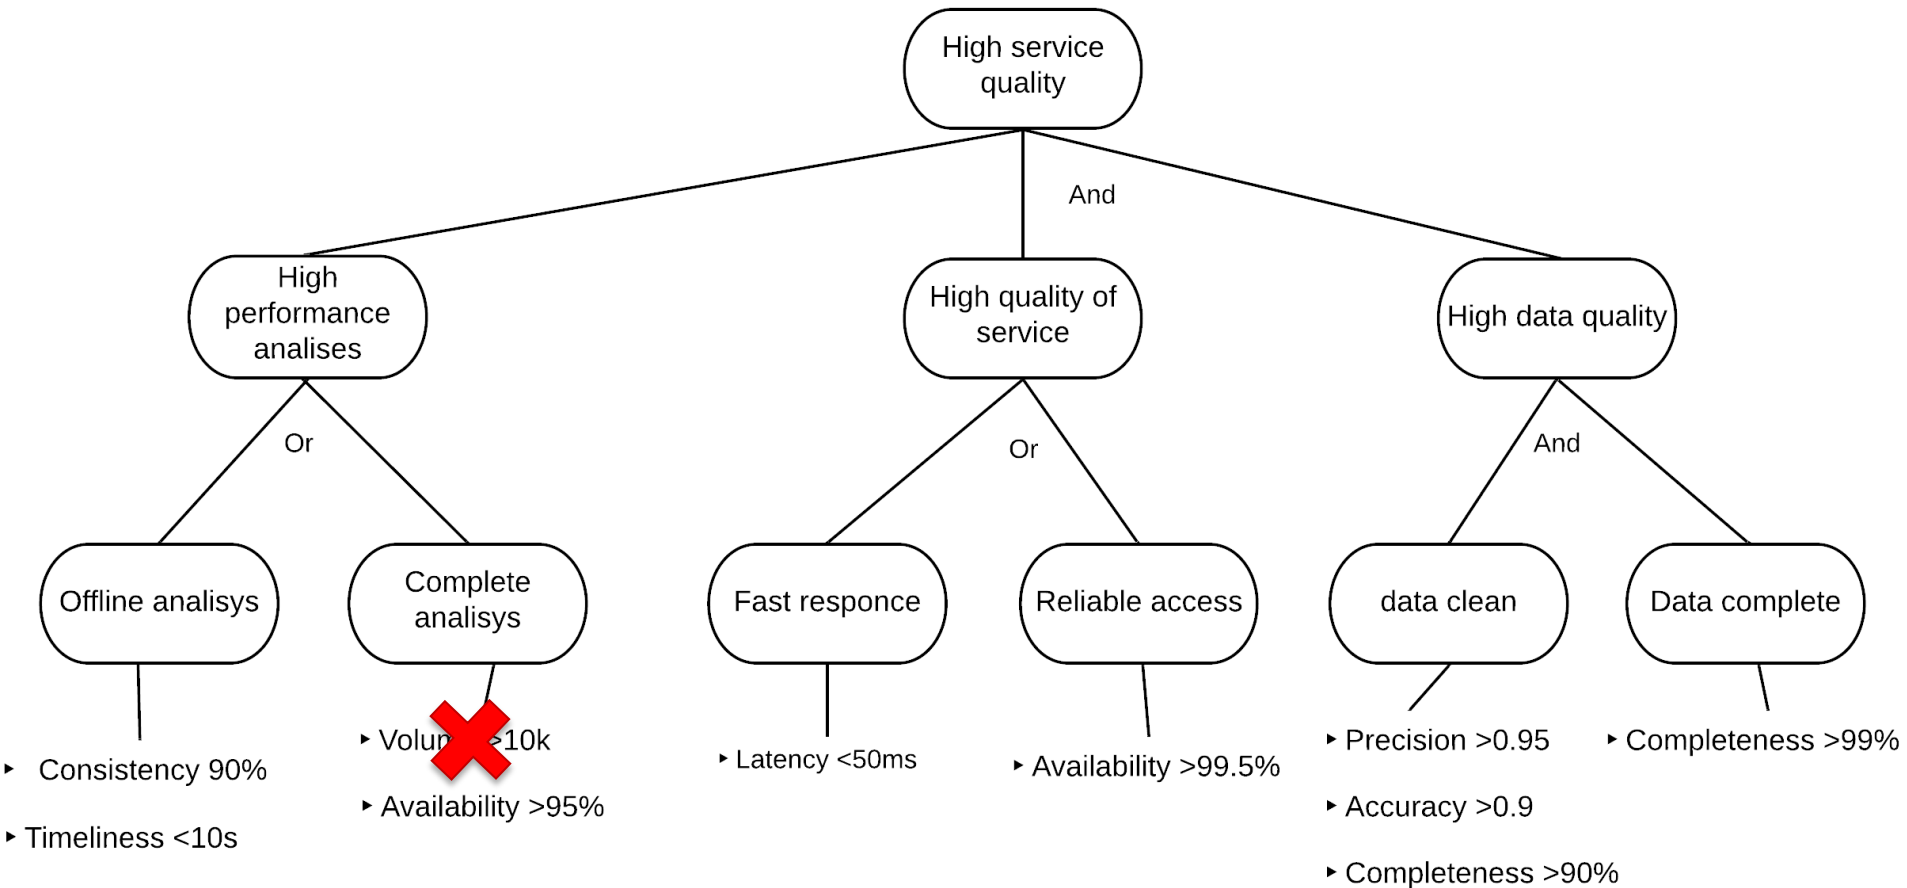
Goal model customization



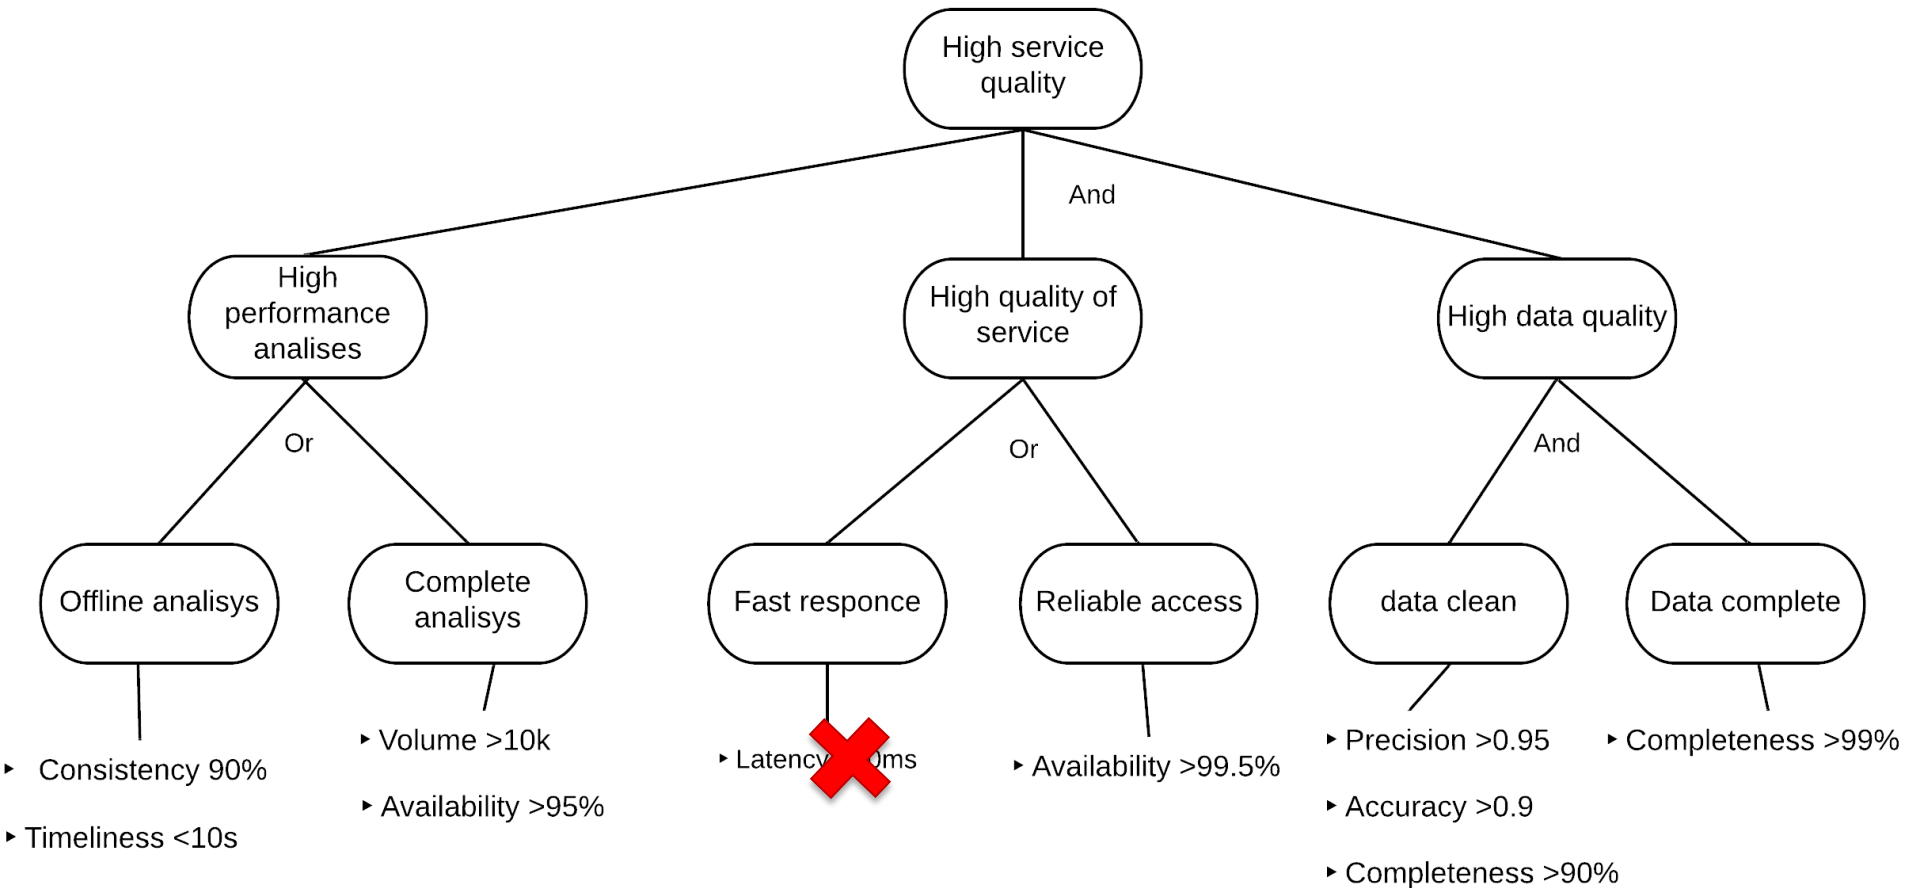
A violation is not a KPI under performing

A violation is the top goals that becomes not achieved

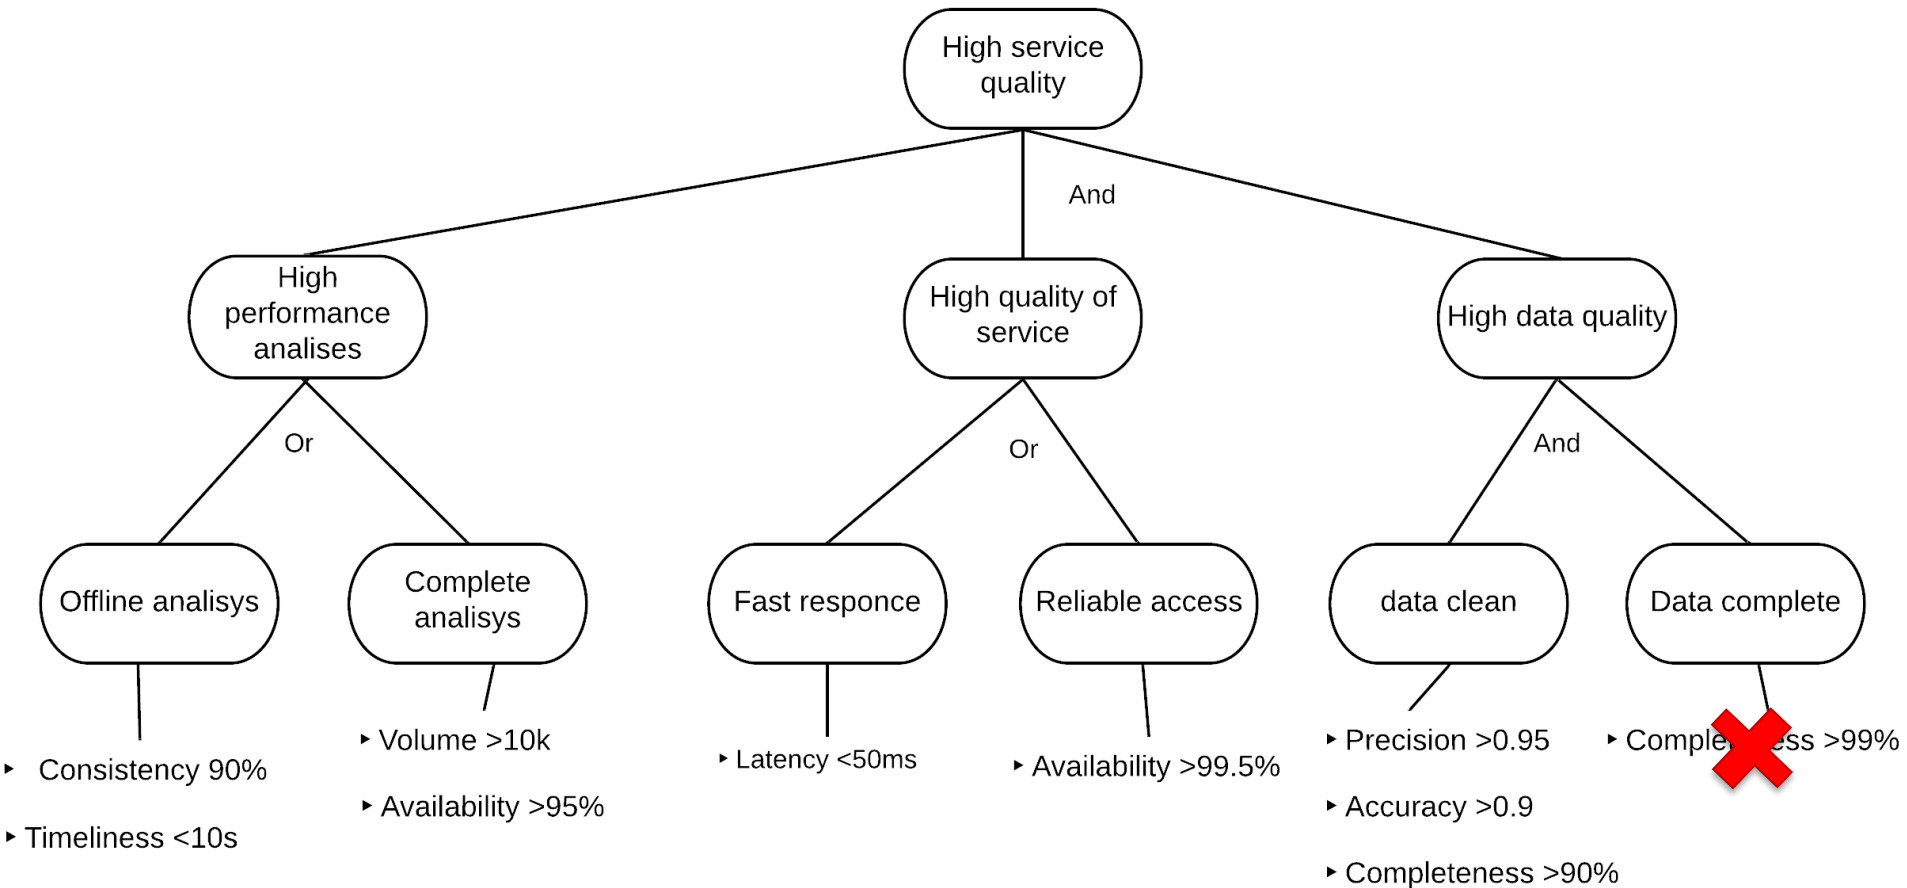
Violation example



Violation example



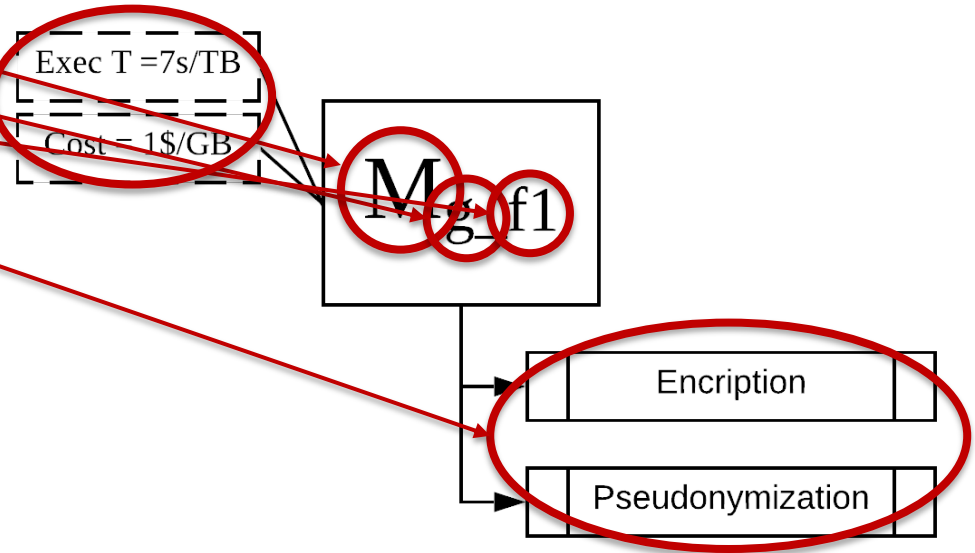
Violation example



How do we represent movement actions?

Information needed

- Type
- Source
- Target
- Transformation
- Costs



Connecting movement actions

When a movement action is enacted the status of the fog environment changes

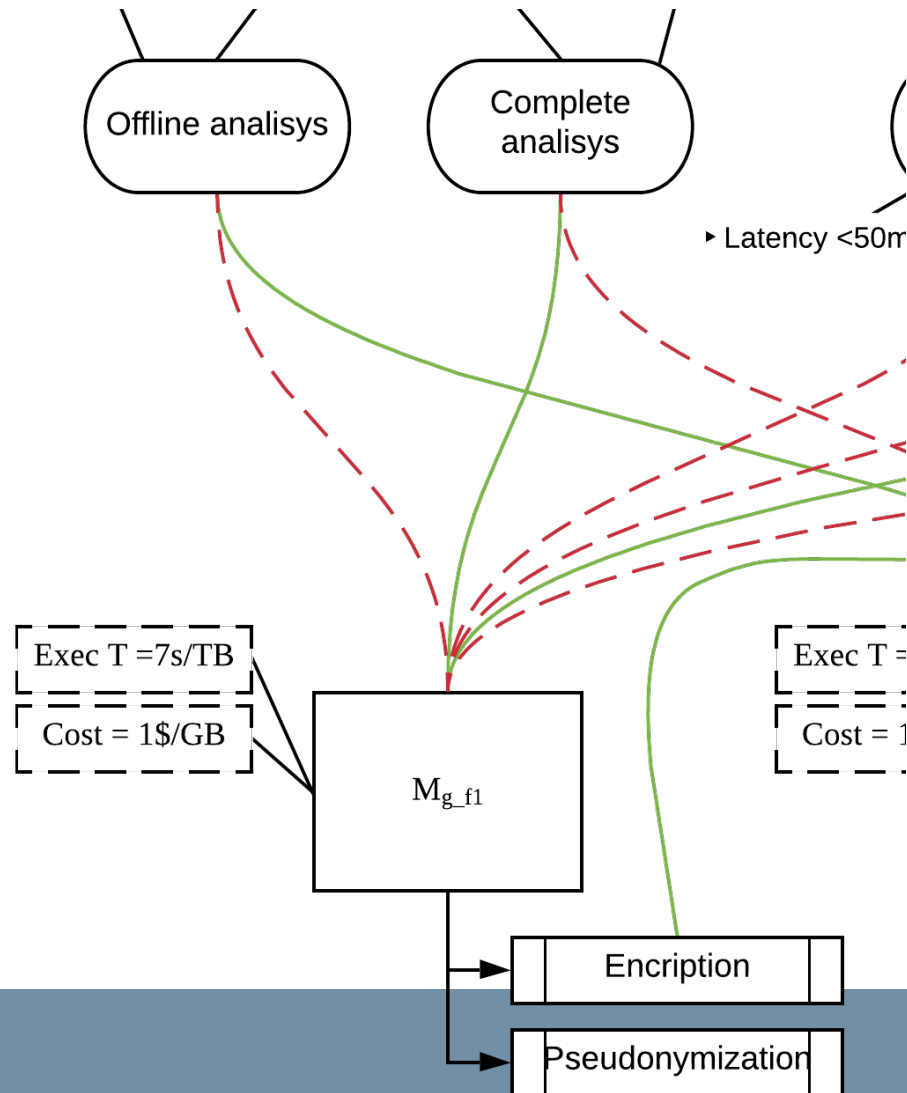


The movement action impacts on the data qualities

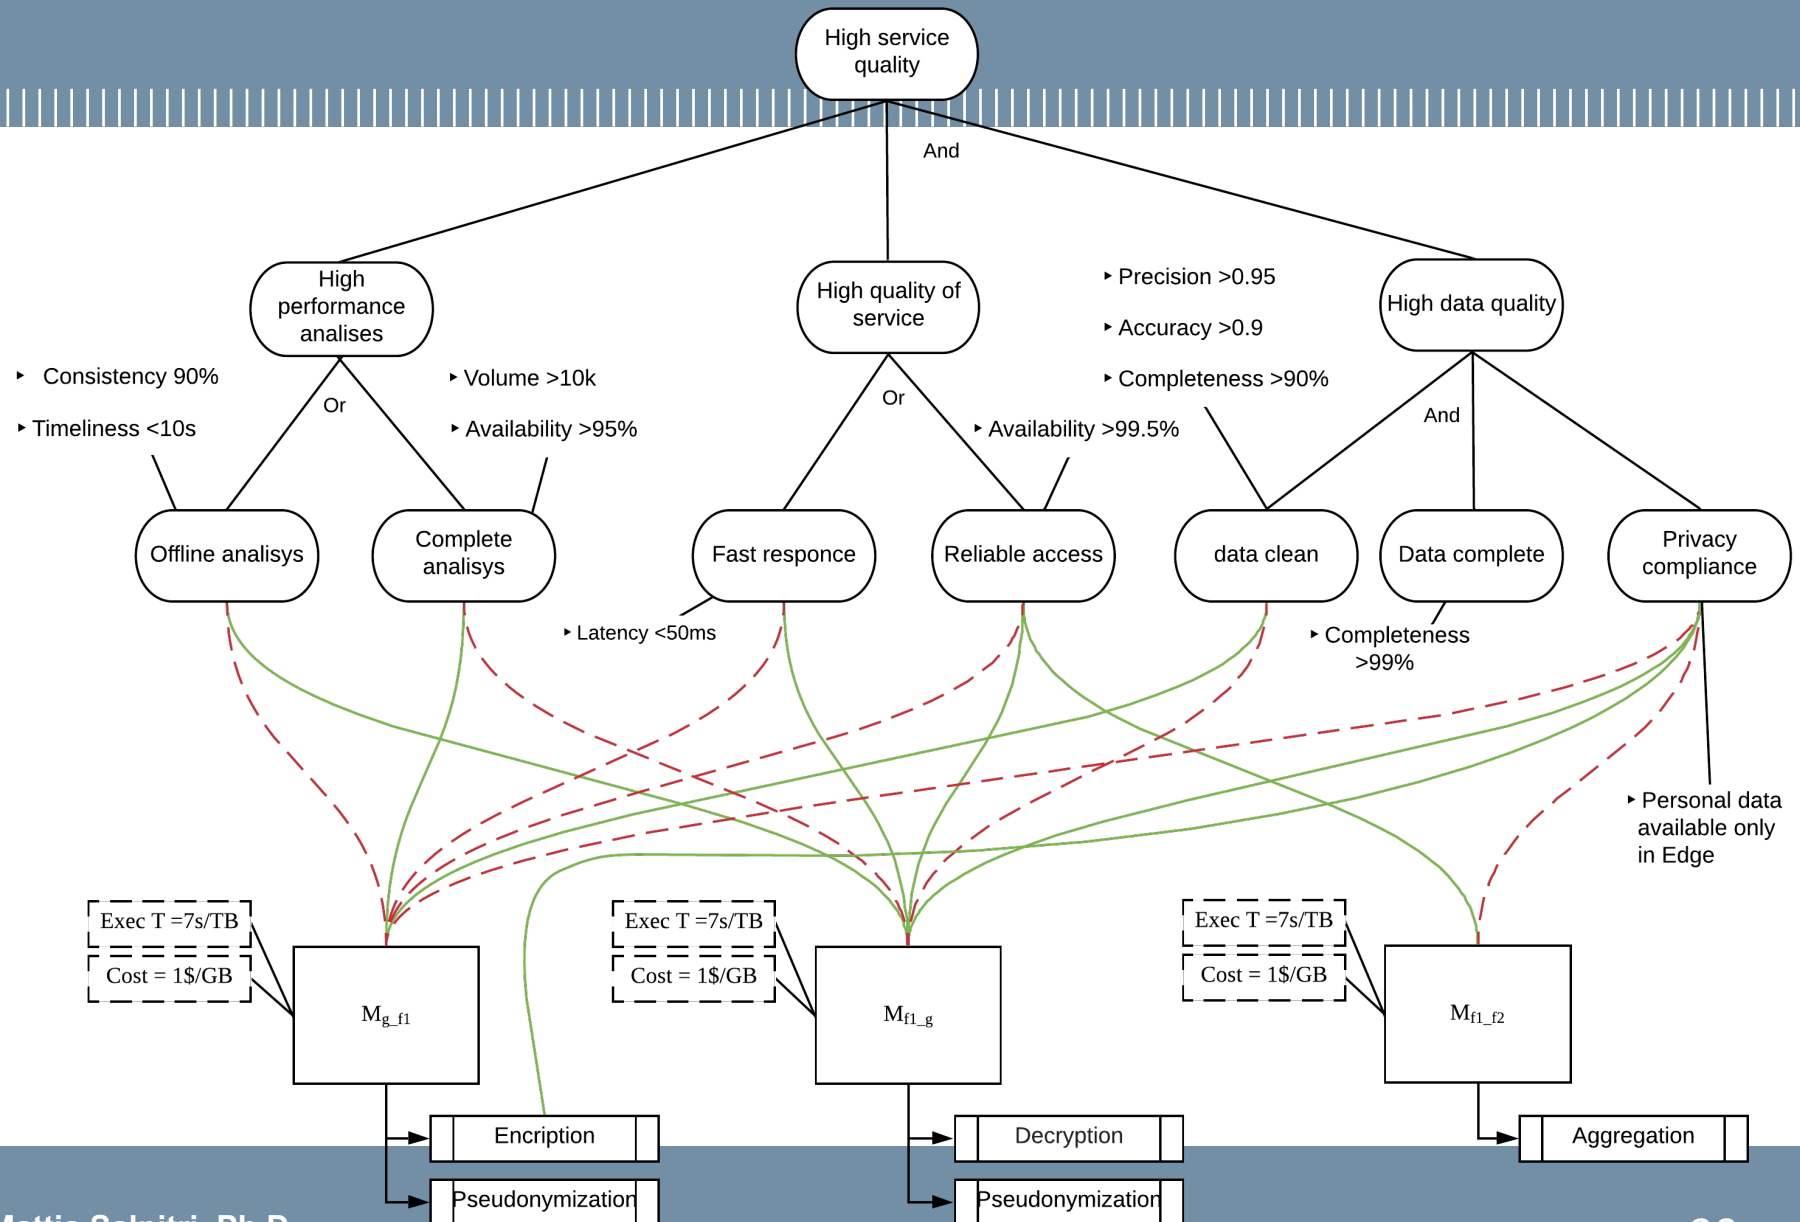
We use two types of impact relations:

- Positive impact
- Negative impact

Example of impact relations



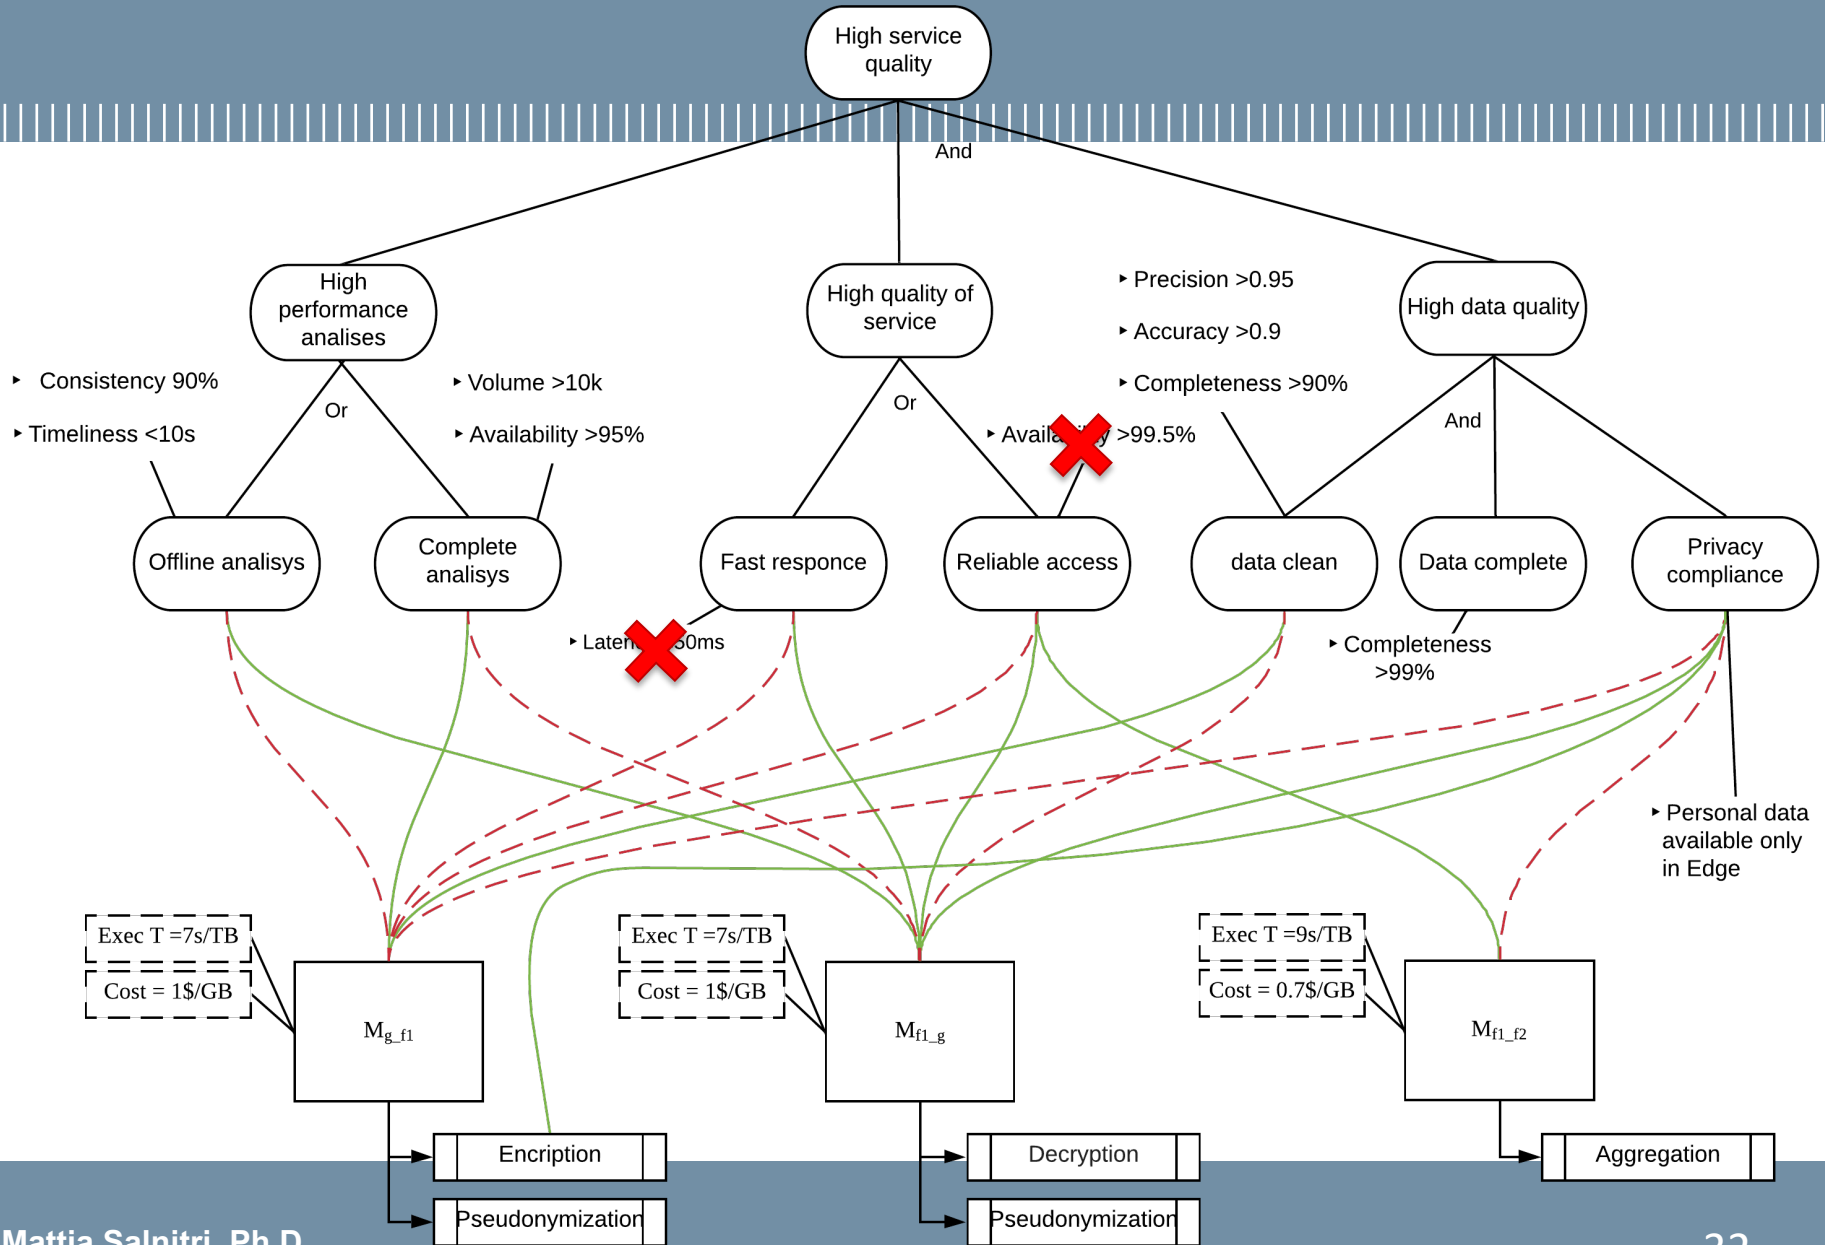
The complete goal model



The selection process

1. Violation is detected
2. Top-goal satisfaction is checked
3. Identification of possible movement actions

The selection of data movement



Two movement selected

- M_{f1_g}
- M_{f1_f2}

Which one do we choose?

Optimization functions define by users

- Prioritize gain
- Prioritize time
- Prioritize positive impact
- ...

The selection process

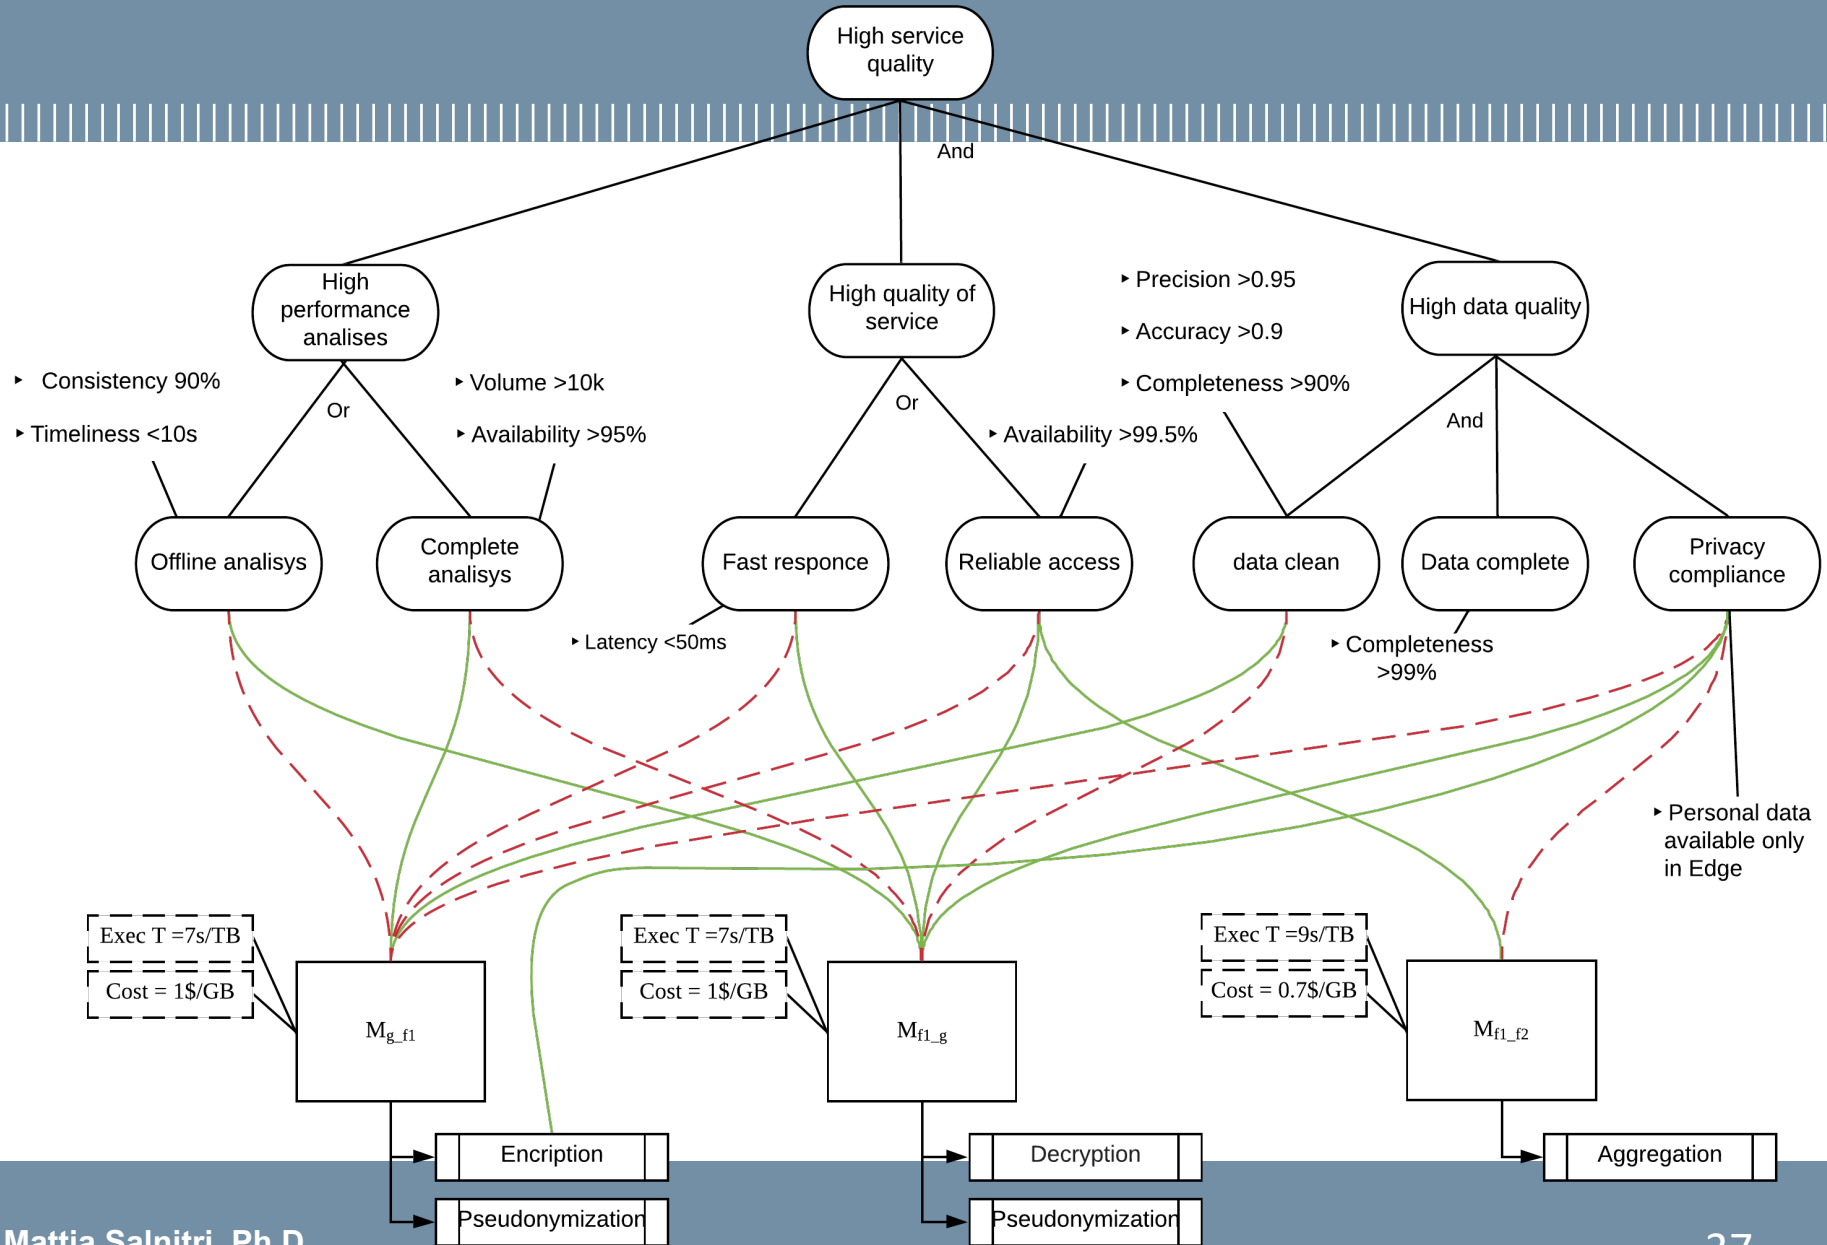
1. Violation is detected
2. Top-goal satisfaction is checked
3. Identification of possible movement actions
4. Selection of movement action based on movement strategy
5. Enactment of the movement action

Every enactment of a movement action influences the fog environment

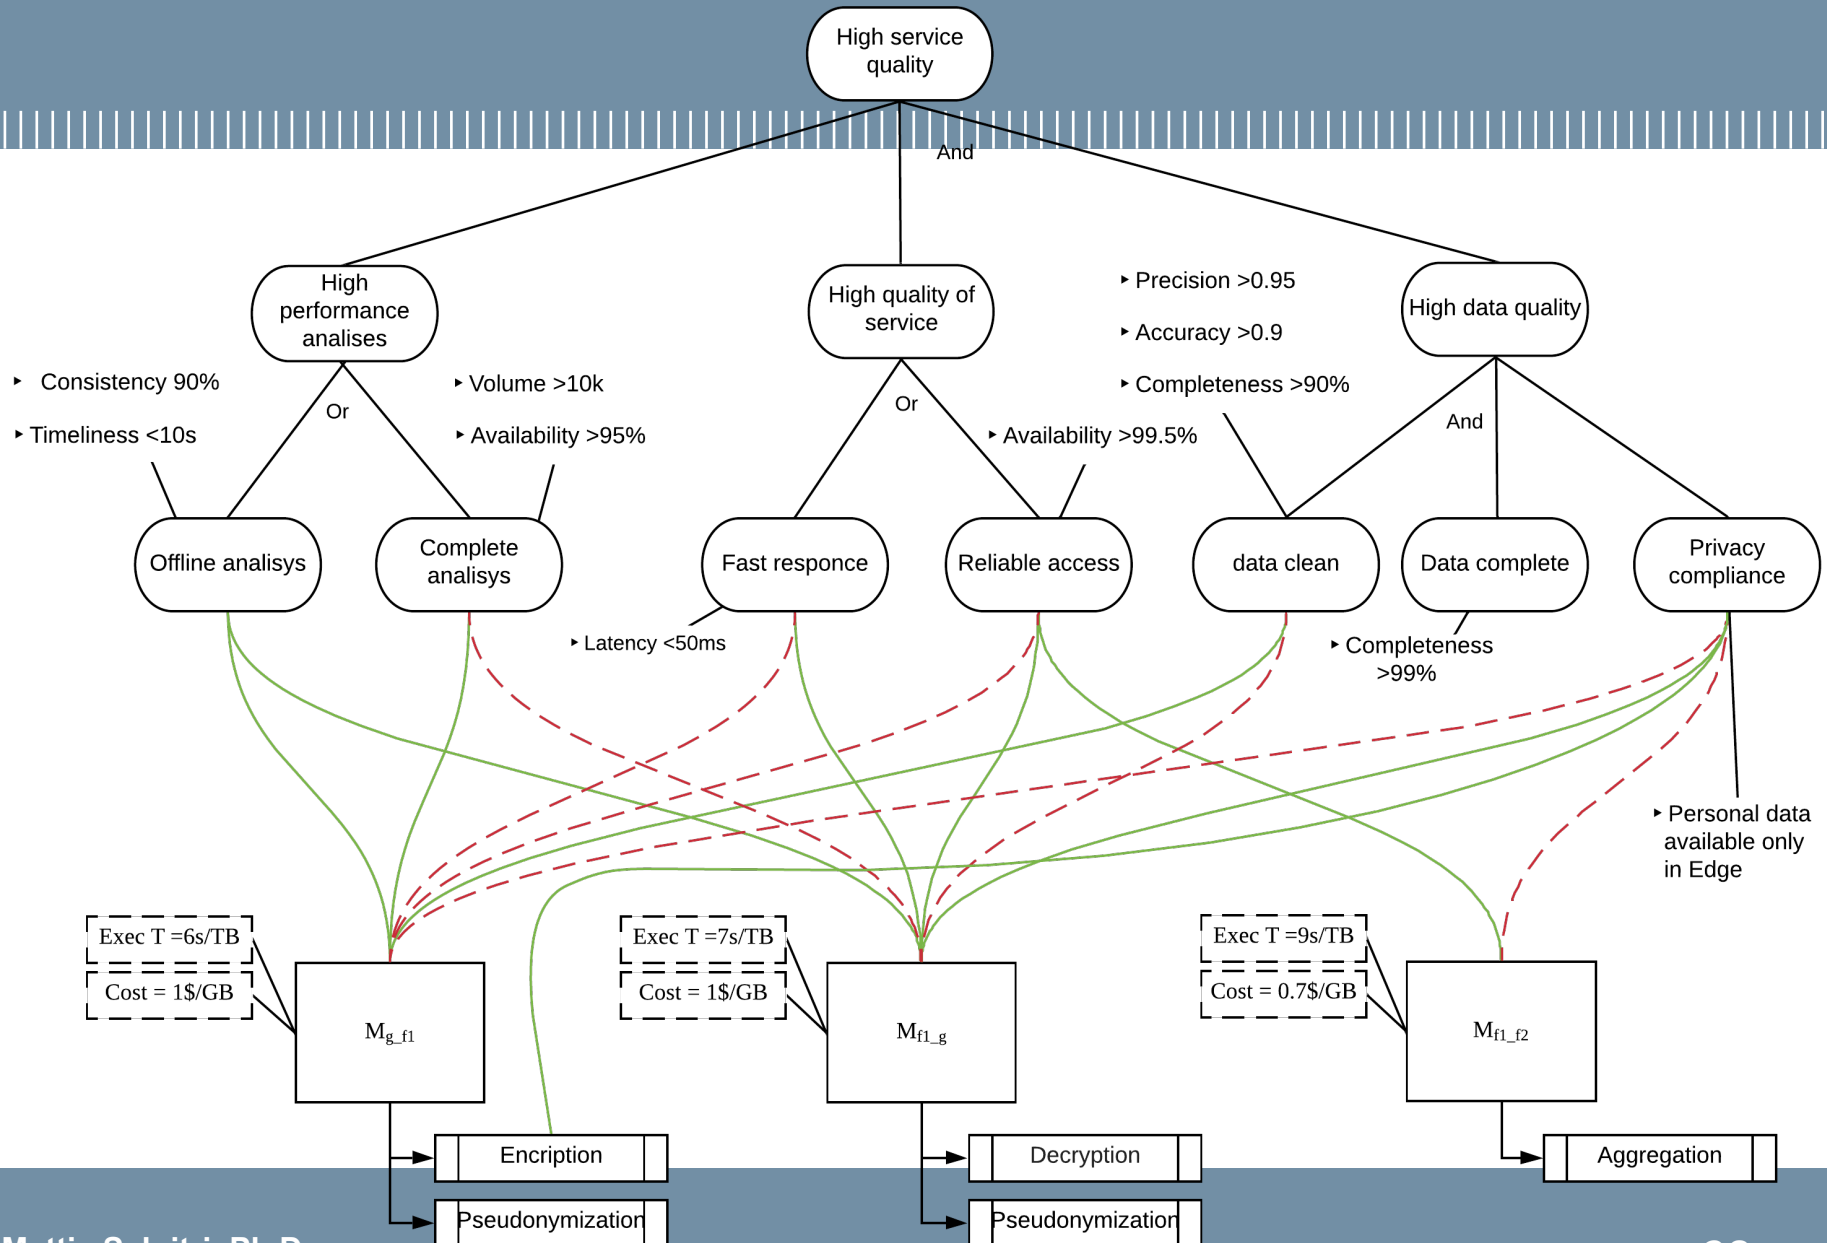
If we monitor the requirements before and after the enactment...

we can refine the contribution links of a goal model

Contribution link update



Contribution link update



Multiple data consumers

Frequently multiple users use data/perform computation

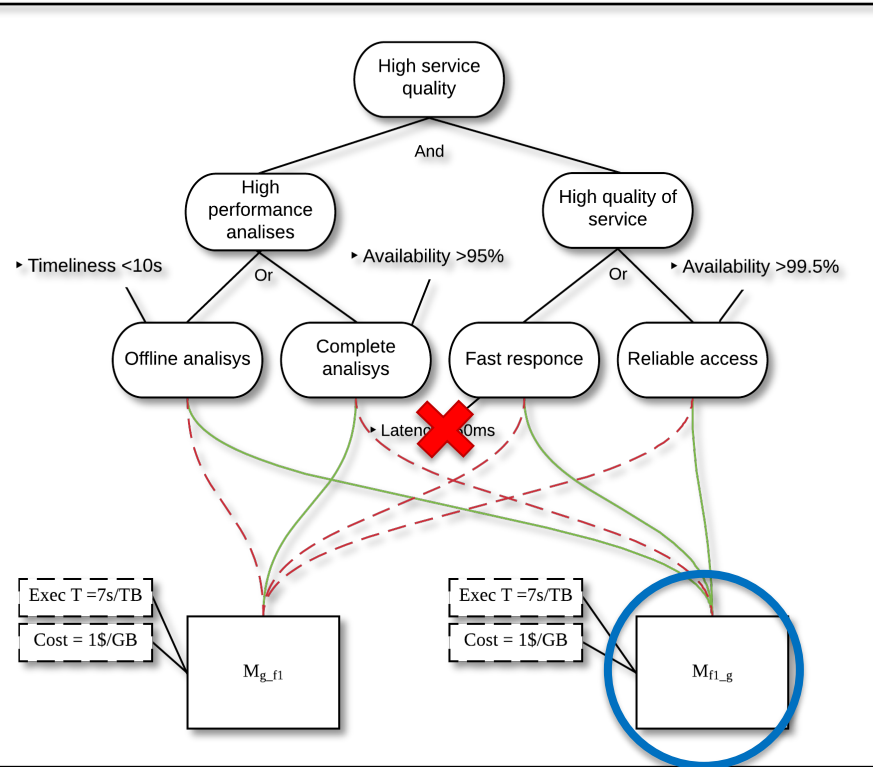
- For different purposes
 - With different requirements
- From different locations
 - Different impact of data movement

Solution:

- Use one goal model for each user

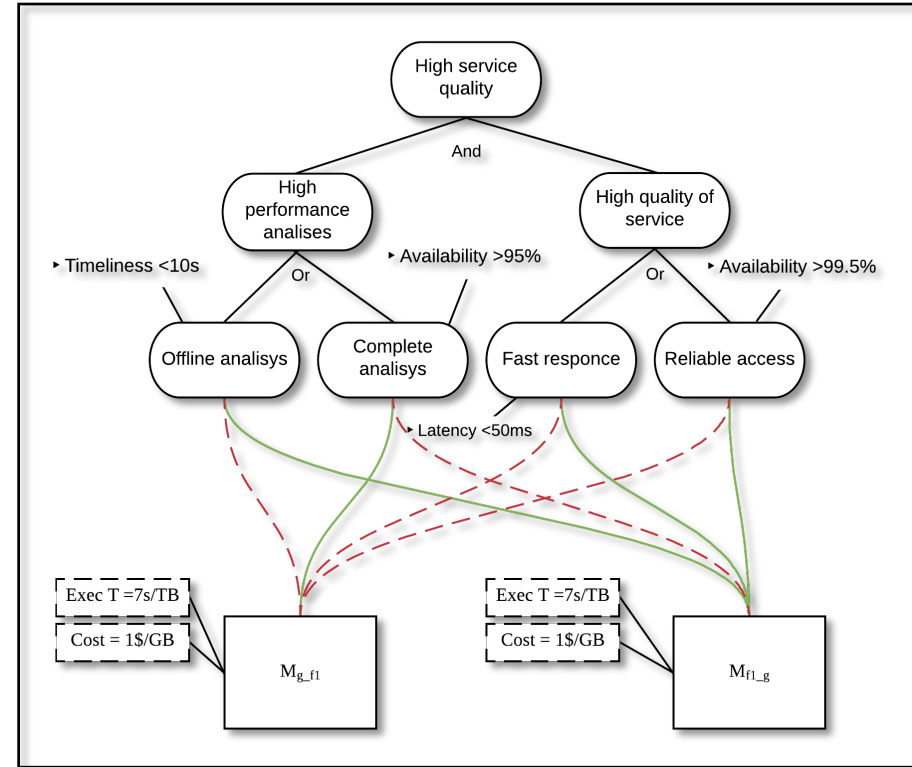
Cycles on movement actions

User A




A: M_{f1_g} ,

User B

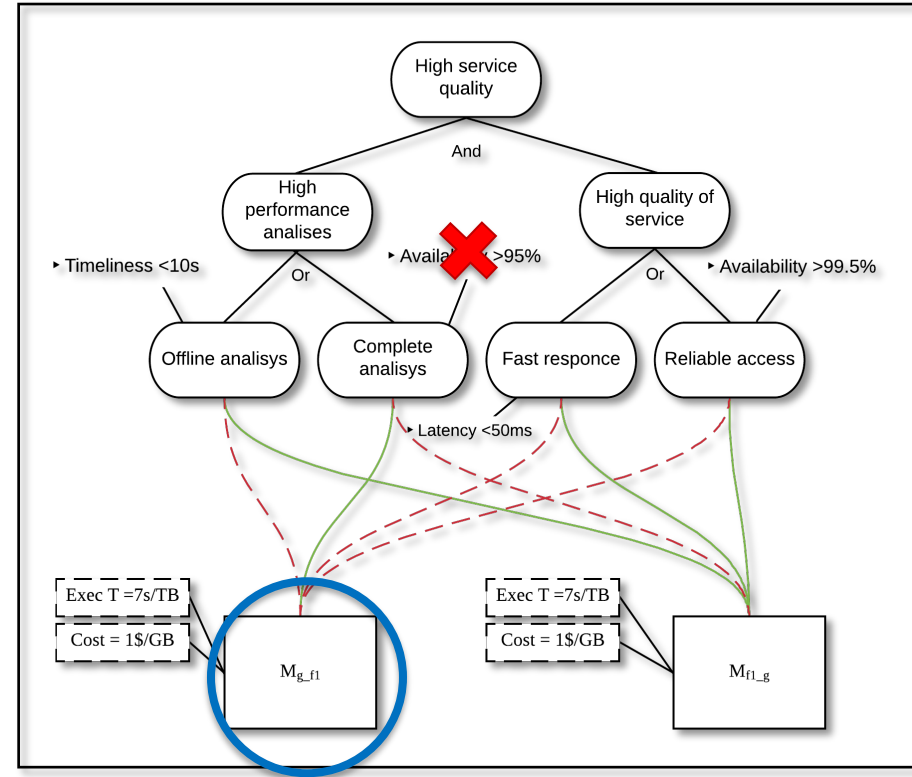


Cycles on movement actions

User A




User B




A: M_{f1_g} , B: M_{g_f1} ,

Cycles on movement actions

User A



User B



A: M_{f1_g} , B: M_{g_f1} , A: M_{f1_g}


Cycles on movement actions

Cycles must be avoided

- Waste of resources
- Hard to detect (multiple parties involved)

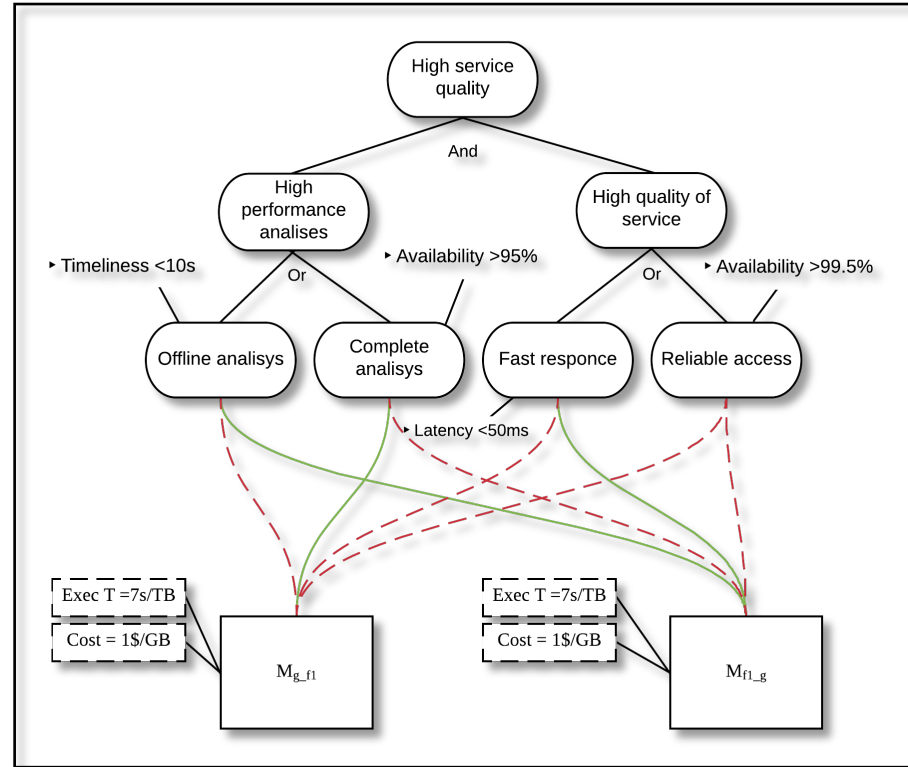
Conflicts on multiple goal models

User A



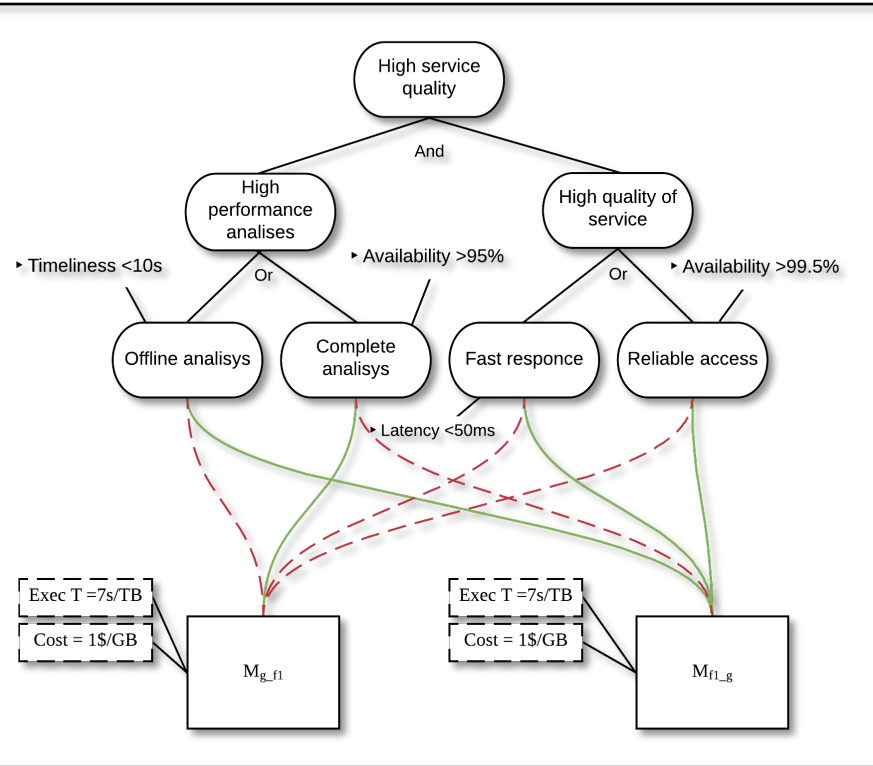
A: M_{f1_g}

User B



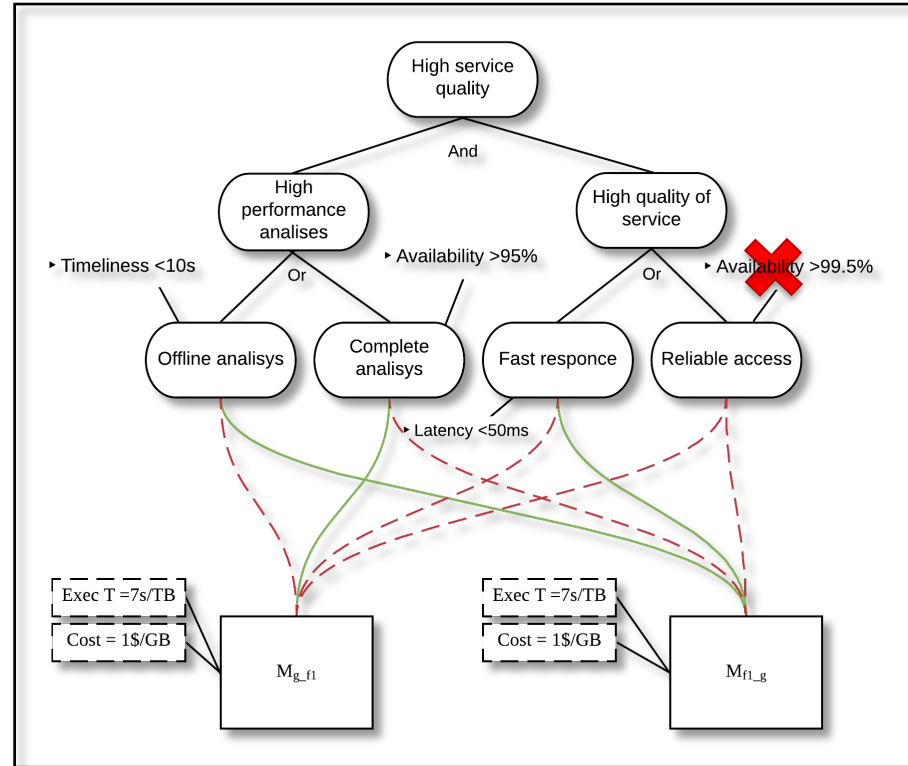
Conflicts on multiple goal models

User A



A: M_{f1_g}

User B



Possible solutions

- Centralized manager
 - DITAS approach
- Decentralized solution
 - Reinforcement learning

Virtual Data Container (VDC)

- Access point for consumer
- Computation tasks performed by the VDC (that can be moved)
- Data provided by the VDC
- One VDC per user -> 1 goal model per VDC


Virtual Data Manager (VDM)

- Coordinates VDC
- Blocks movement action cycles

Centralized manager

Computation
movement

Data
Movement



Decentralized solution


Each VDC is independent

Every time a movement action is enacted, feedbacks arrive from other VDC who uses the resources moved.

Decentralized solution

Computation
movement

Data
Movement





The DITAS framework uses a centralized approach

- Monitoring component in VDCs detect violations,
- A decision system in the VDM enacts the movement actions considering all goal models.




POLITECNICO
MILANO 1863

Thank you!

Pierluigi Plebani, Mattia Salnitri, Monica Vitali
name.surname@polimi.it

Jeju, 23rd September 2019




DITAS
Data-intensive applications
Improvement by moving data
and computation in mixed
cloud/fog environments

Acknowledgements

This research work have been fouded by Ditas Project.

DITAS project is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement RIA 731945.



1. Horkoff, J., Borgida, A., Mylopoulos, J., Barone, D., Jiang, L., Yu, E., Amyot, D.: Making data meaningful: The business intelligence model and its formal semantics in description logics. In: Proc. of On the Move to Meaningful Internet Systems. pp. 700–717 (2012)