Elastic Resource Adaptation in the OpenStack Platform

Pedro Martinez-Julia, Ved P. Kafle, Hiroaki Harai
ペドロ・マルティネスフリア、ベド カフレ、原井 洋明

Network Science and Convergence Device Technology Laboratory, Network System Research Institute
National Institute of Information and Communications Technology (NICT)
{pedro,kafle,harai}@nict.go.jp

IEICE Technical Committee on Network Virtualization (NV)

7 March 2018

Outline

- Problem Statement:
  - Motivation and Research Topic
  - Use Case
- Solution:
  - Proposed Approach
  - Architecture Overview
  - Requirement Anticipation
  - Integration with ETSI-NFV-MANO.
- Conclusions & Future Work

Motivation and Research Topic (I)

Trivia:

“High variation in resource demand” vs “Fixed resource allocation”.

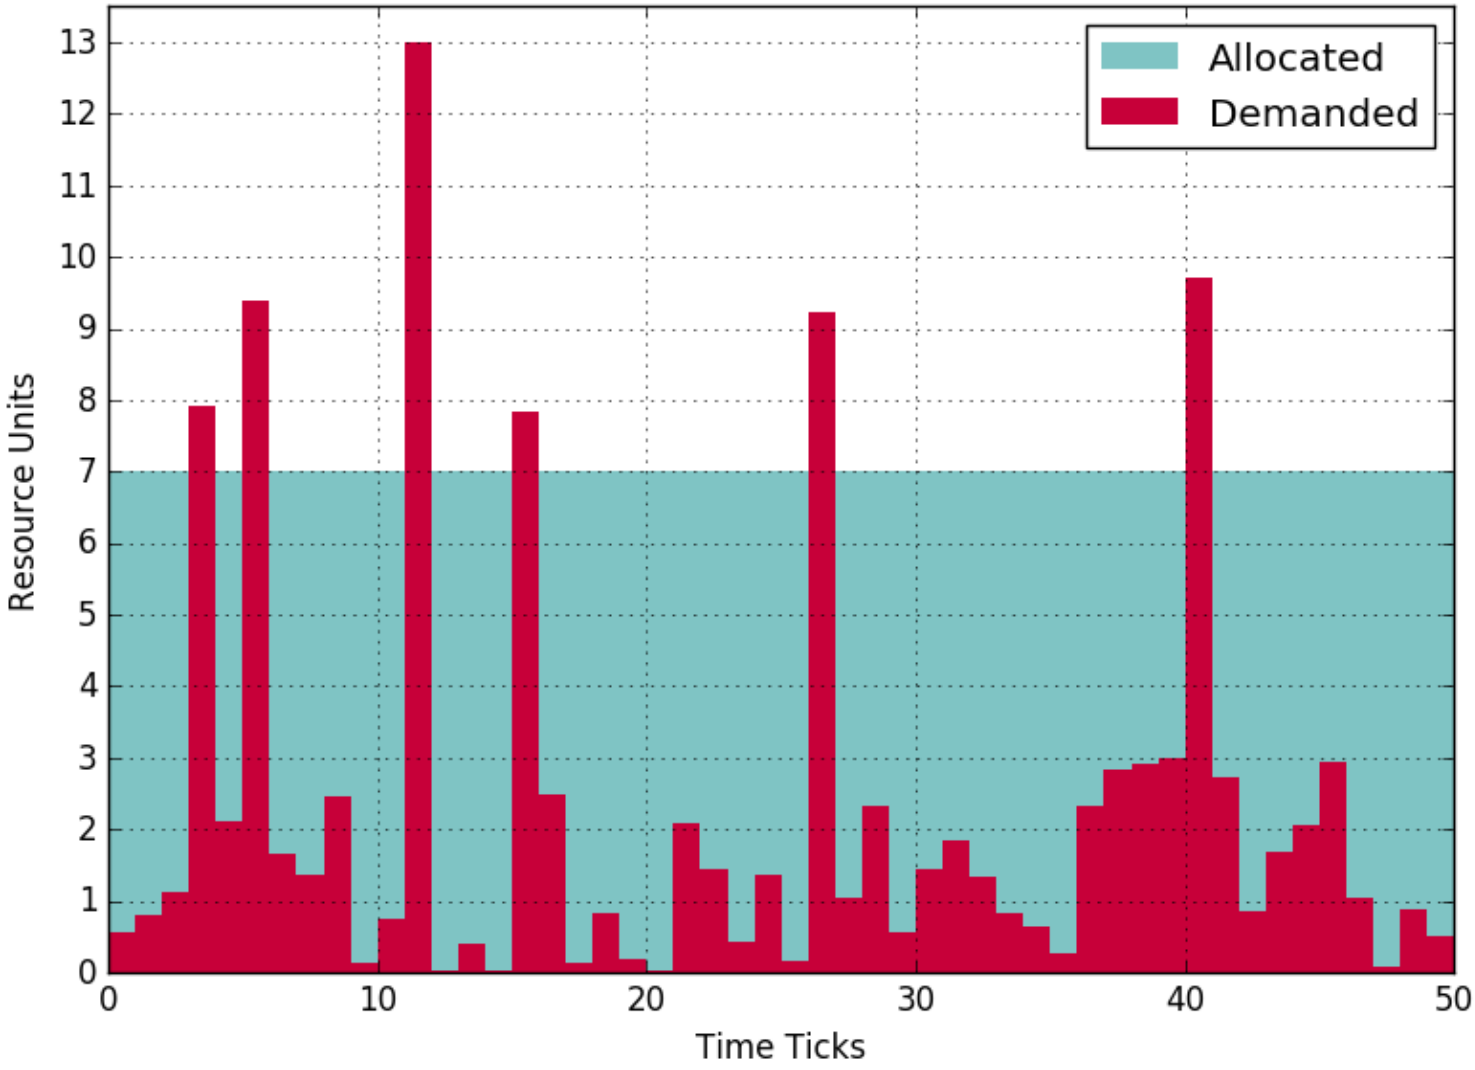

Motivation and Research Topic (II)

Trivia:

“High variation in resource demand” vs “Fixed resource allocation”.

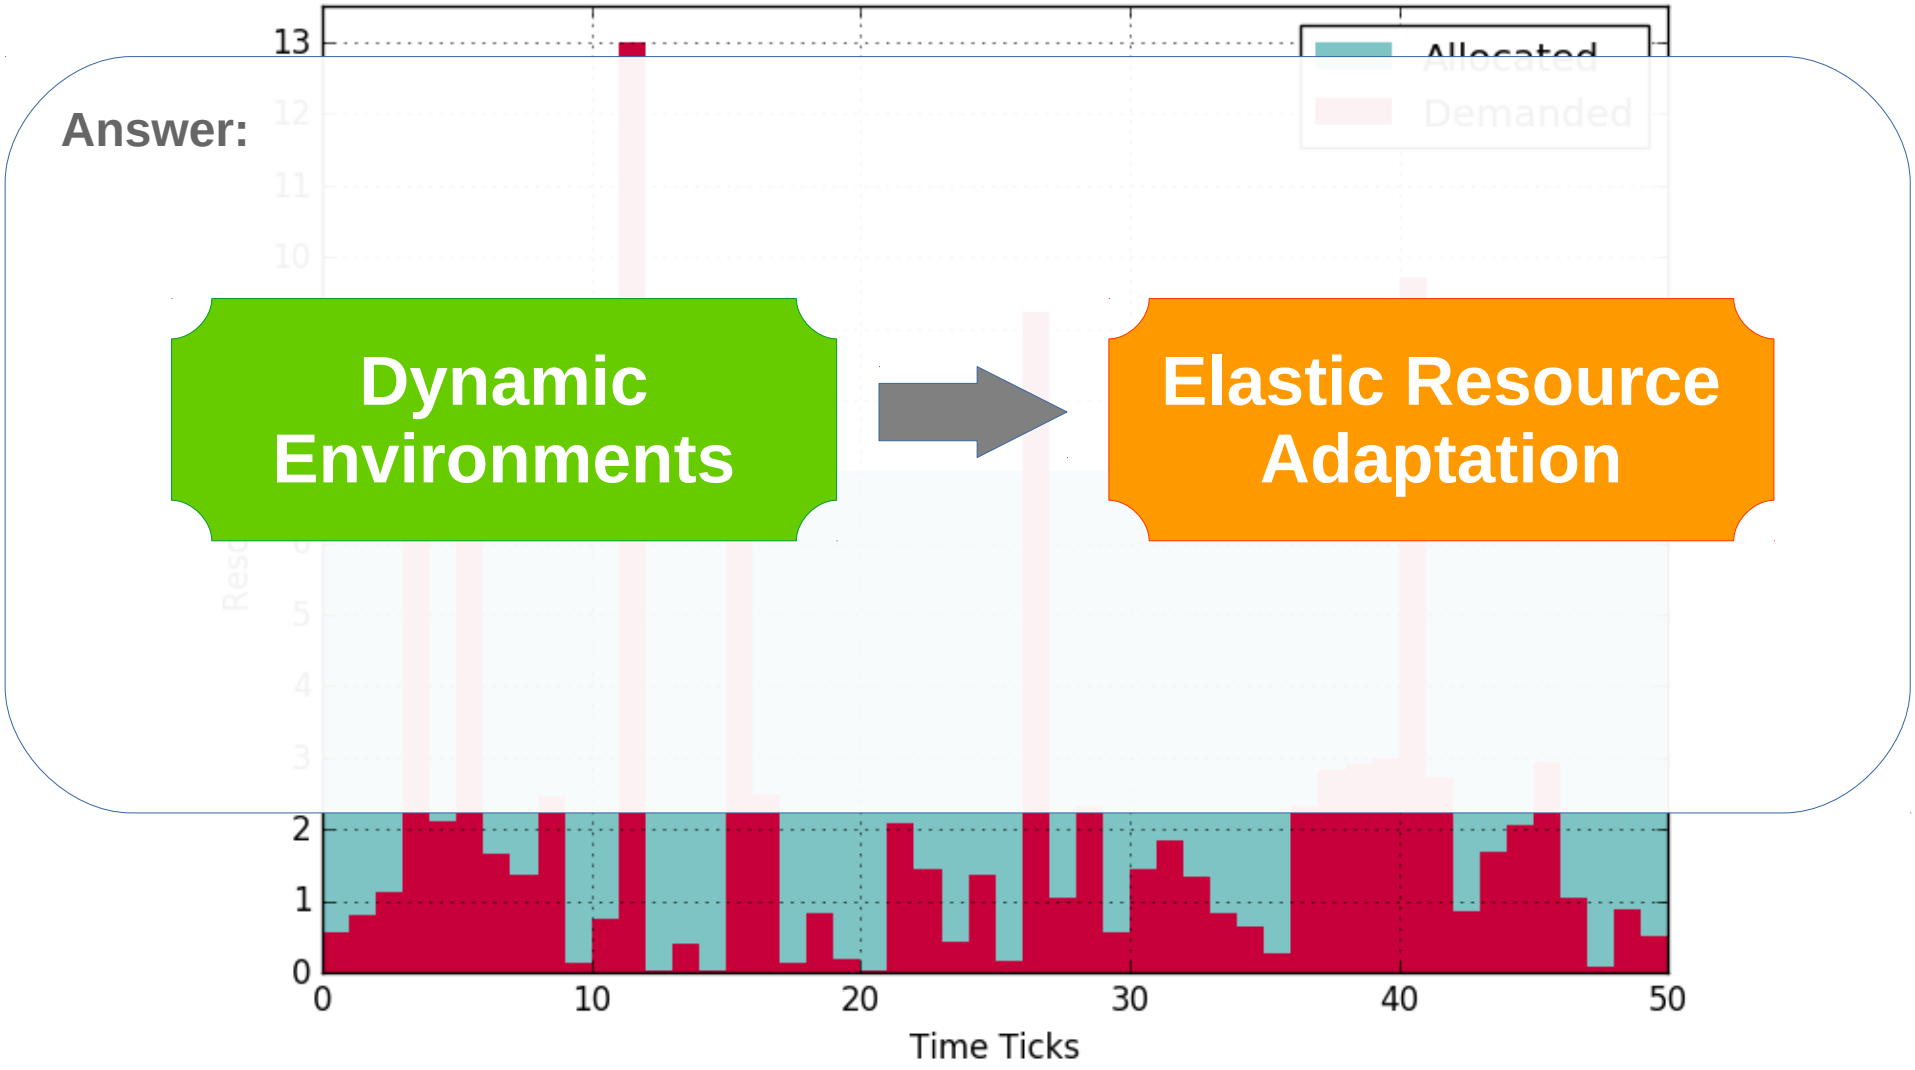

Motivation and Research Topic (III)

Trivia:

vs “Fixed resource allocation”.

Virtual computer and network systems can be dynamically dimensioned to:

- Improve resource utilization.
- Reduce **CAPEX**.

Dynamic Environments → **Elastic Resource Adaptation**

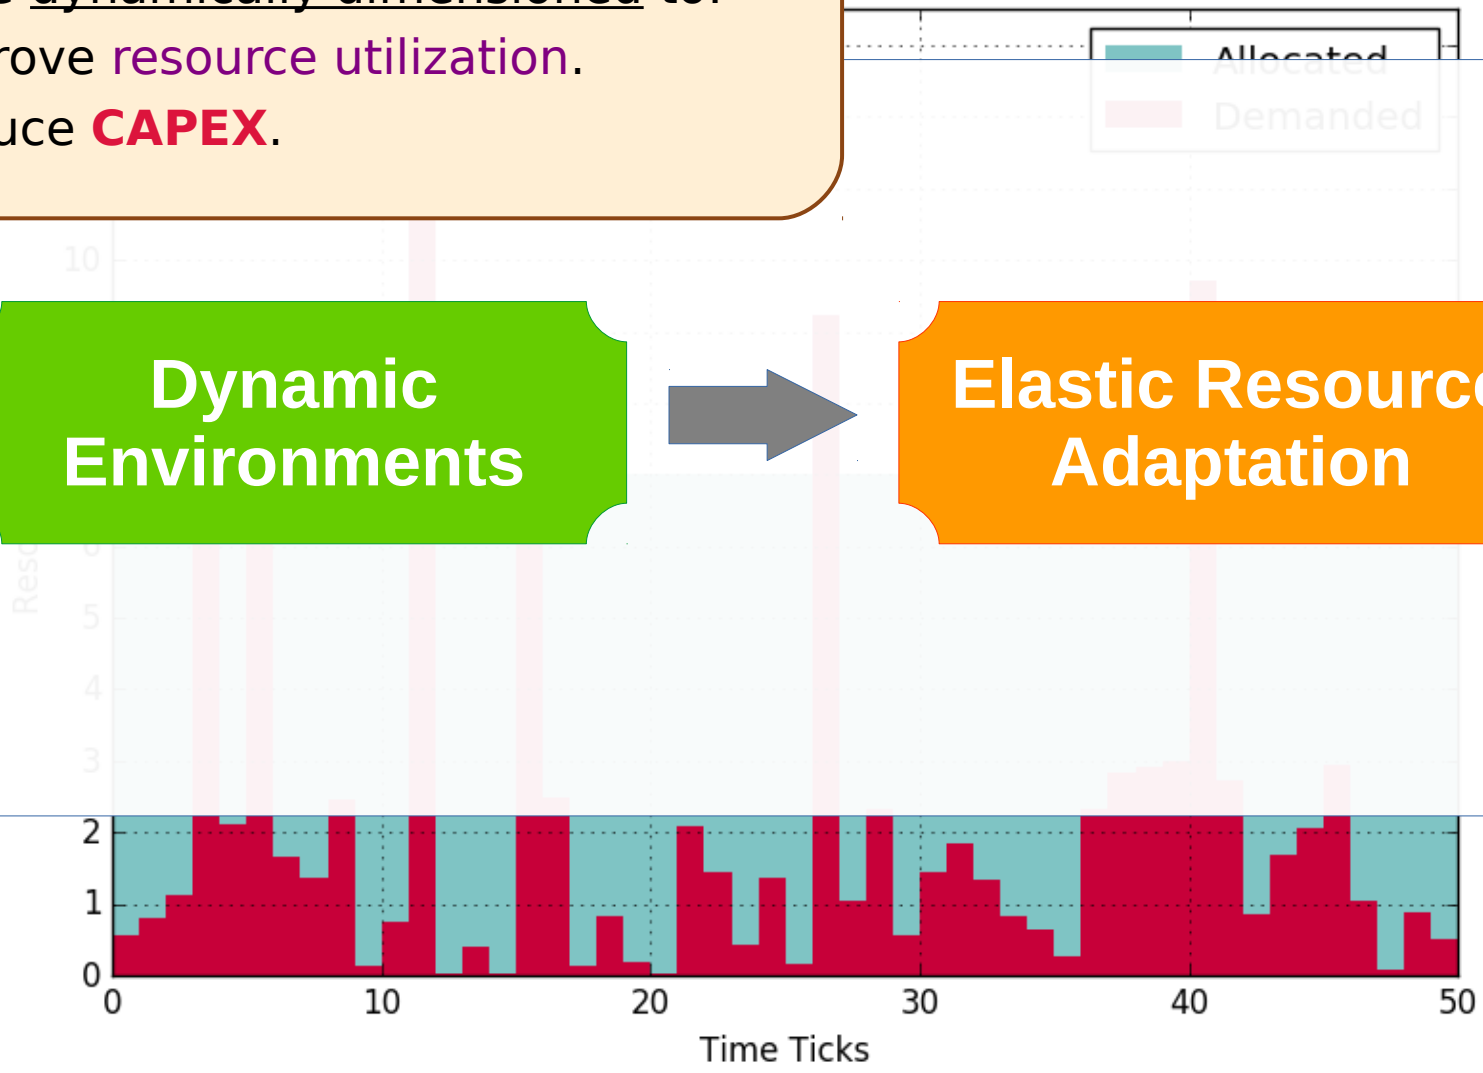

Motivation and Research Topic (IV)

Trivia:

vs “**Fixed resource allocation**”.

Virtual computer and network systems can be dynamically dimensioned to:

- Improve resource utilization.
- Reduce **CAPEX**.

Dynamic Environme

Automated solutions aim to set the **optimum dimension** for every situation:

- Approach increased system complexity with ***intelligent*** and ***intelligence*** methods.
- Reduce both **OPEX** and **CAPEX**.

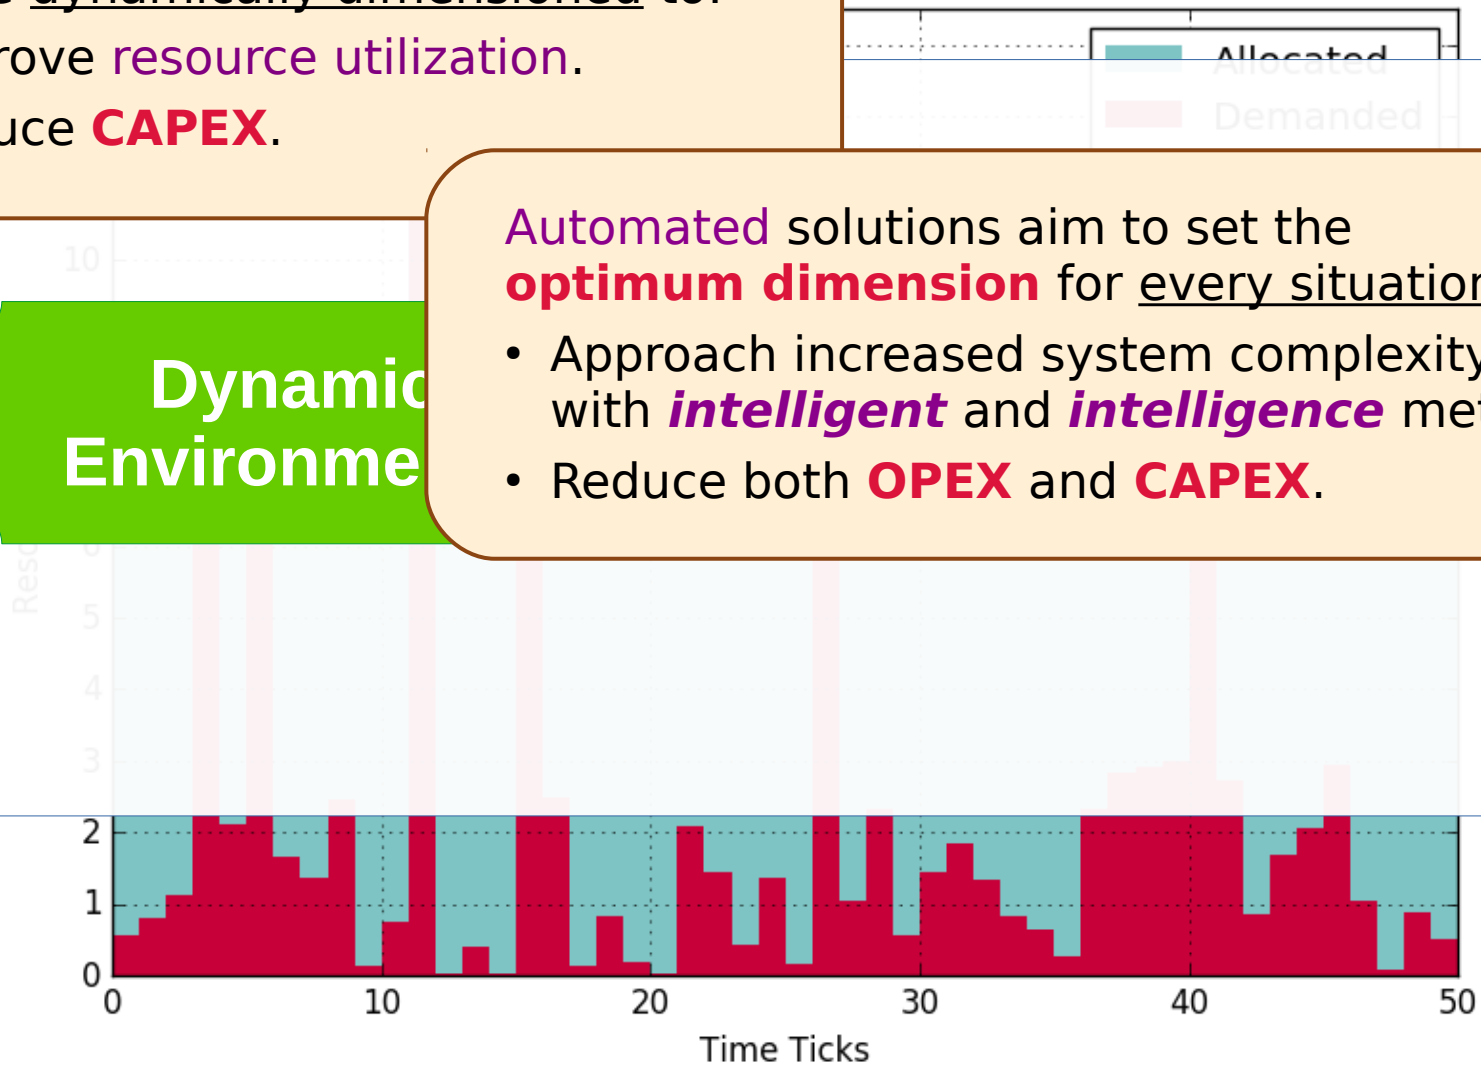

Use Case (I)

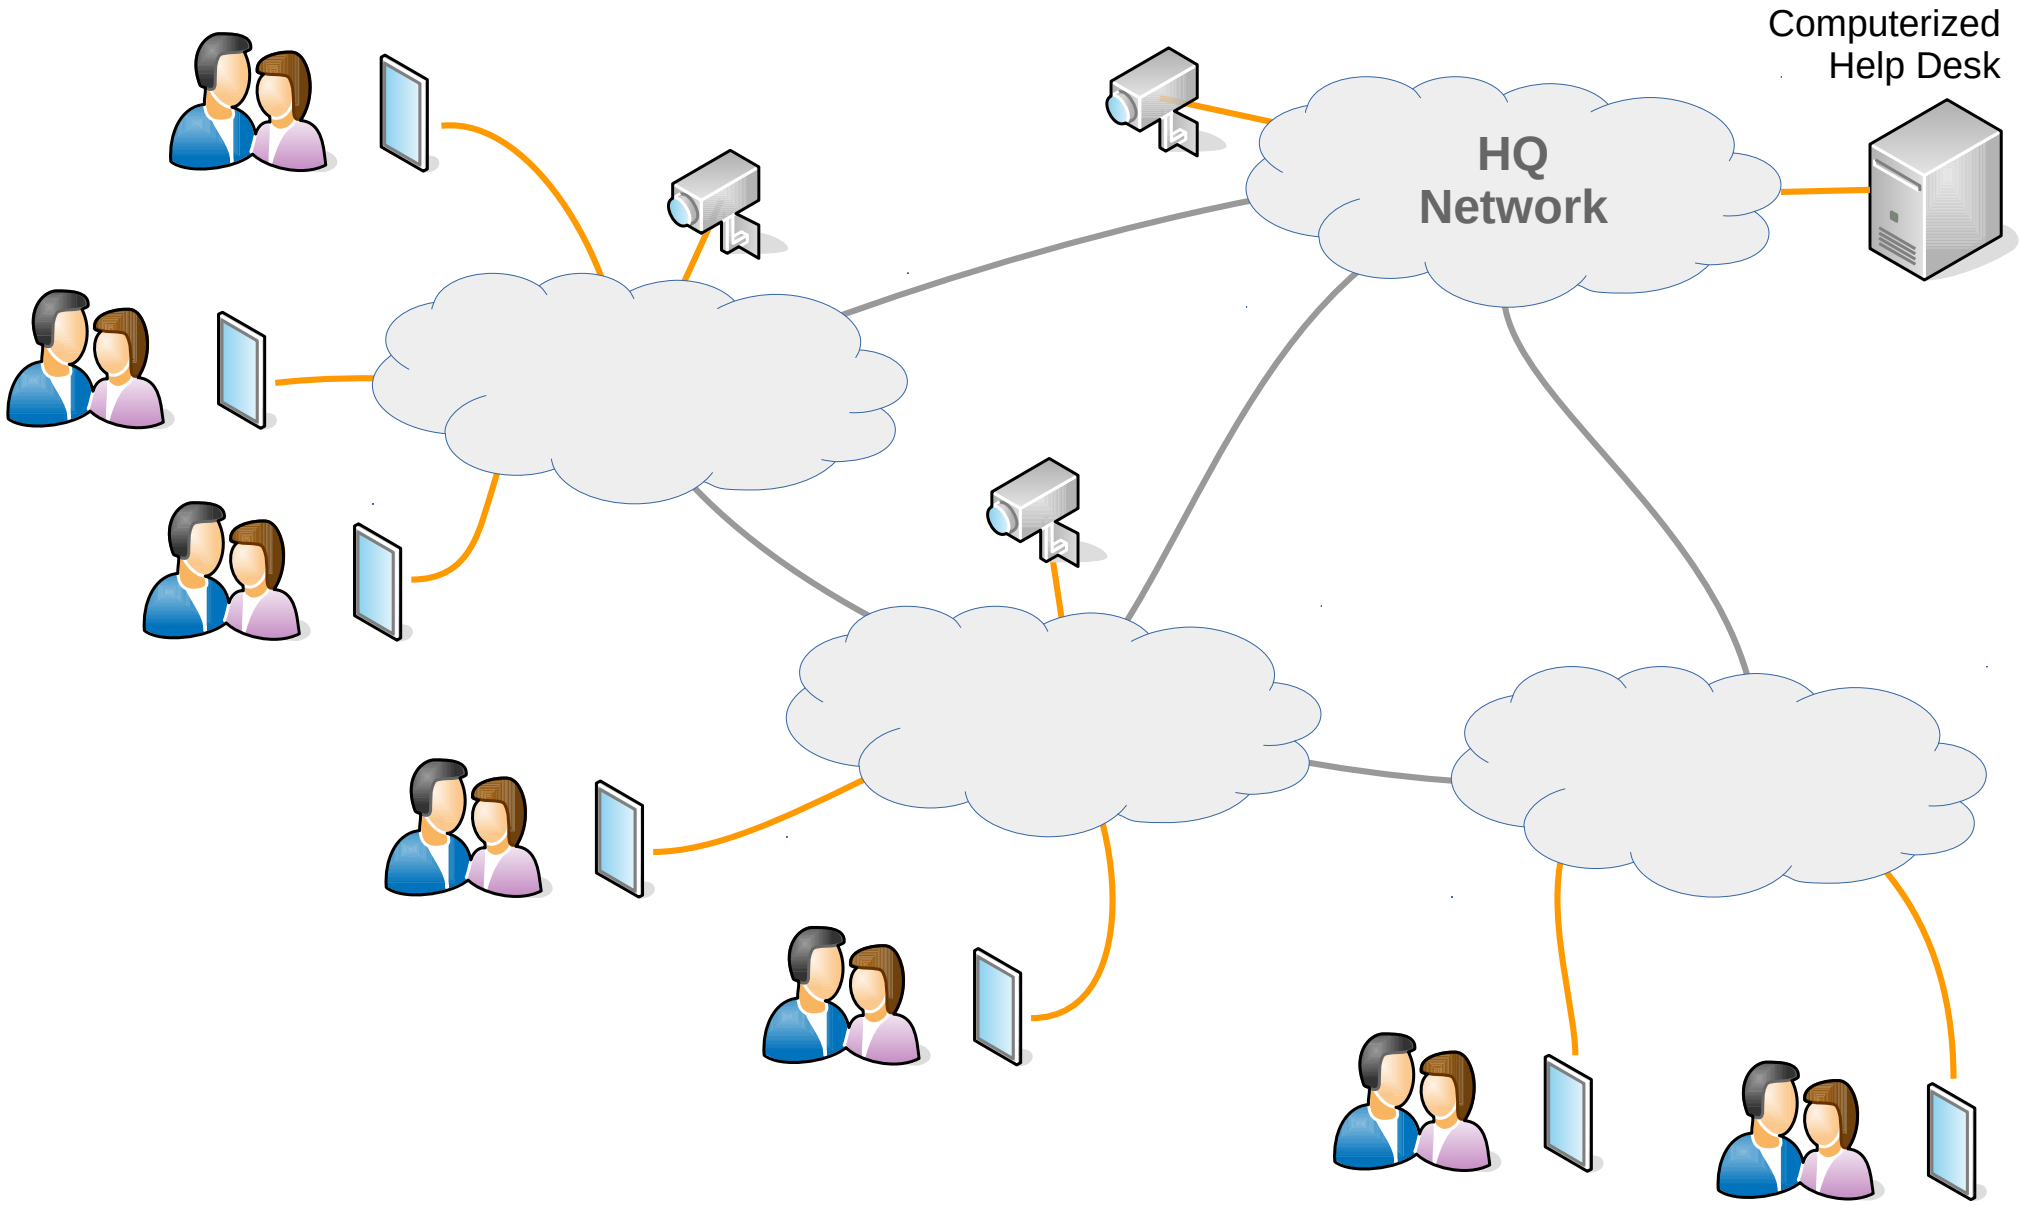

Use Case (II)

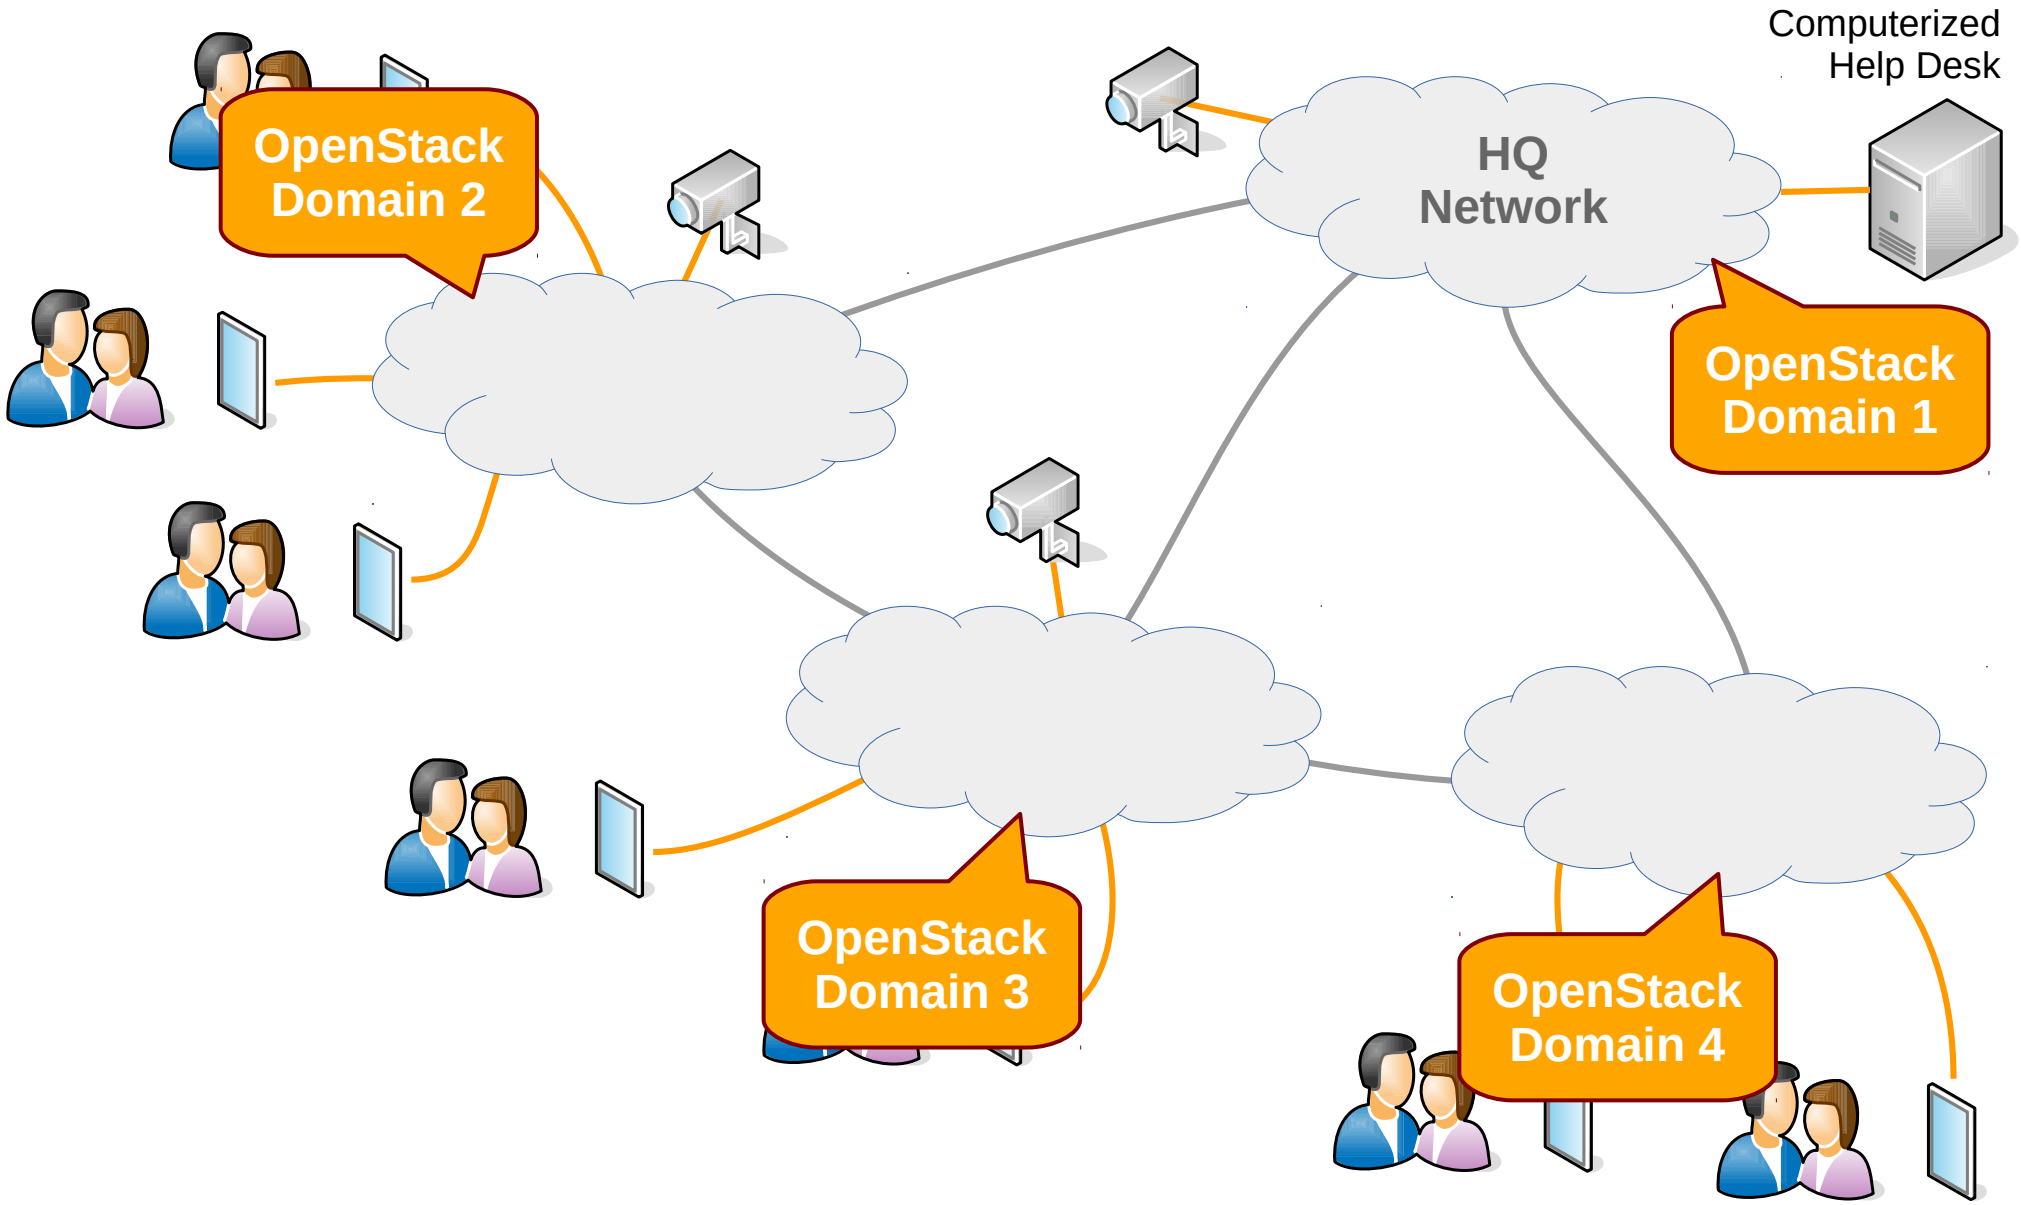

Use Case (III)

OpenStack:

- Facilitates the **construction** of virtual computer and network sytstems.
  - It is widely used to create production-ready *virtualization environments*.
- Enables the **adaptation** of resources:
  - On-demand instantiation or removal of VMs attached to a service.
- Offers application **interfaces**:
  - Monitoring and resource adaptation.
- Supports **NFV**.
- Its operation will be **enhanced** by the results of our research work.

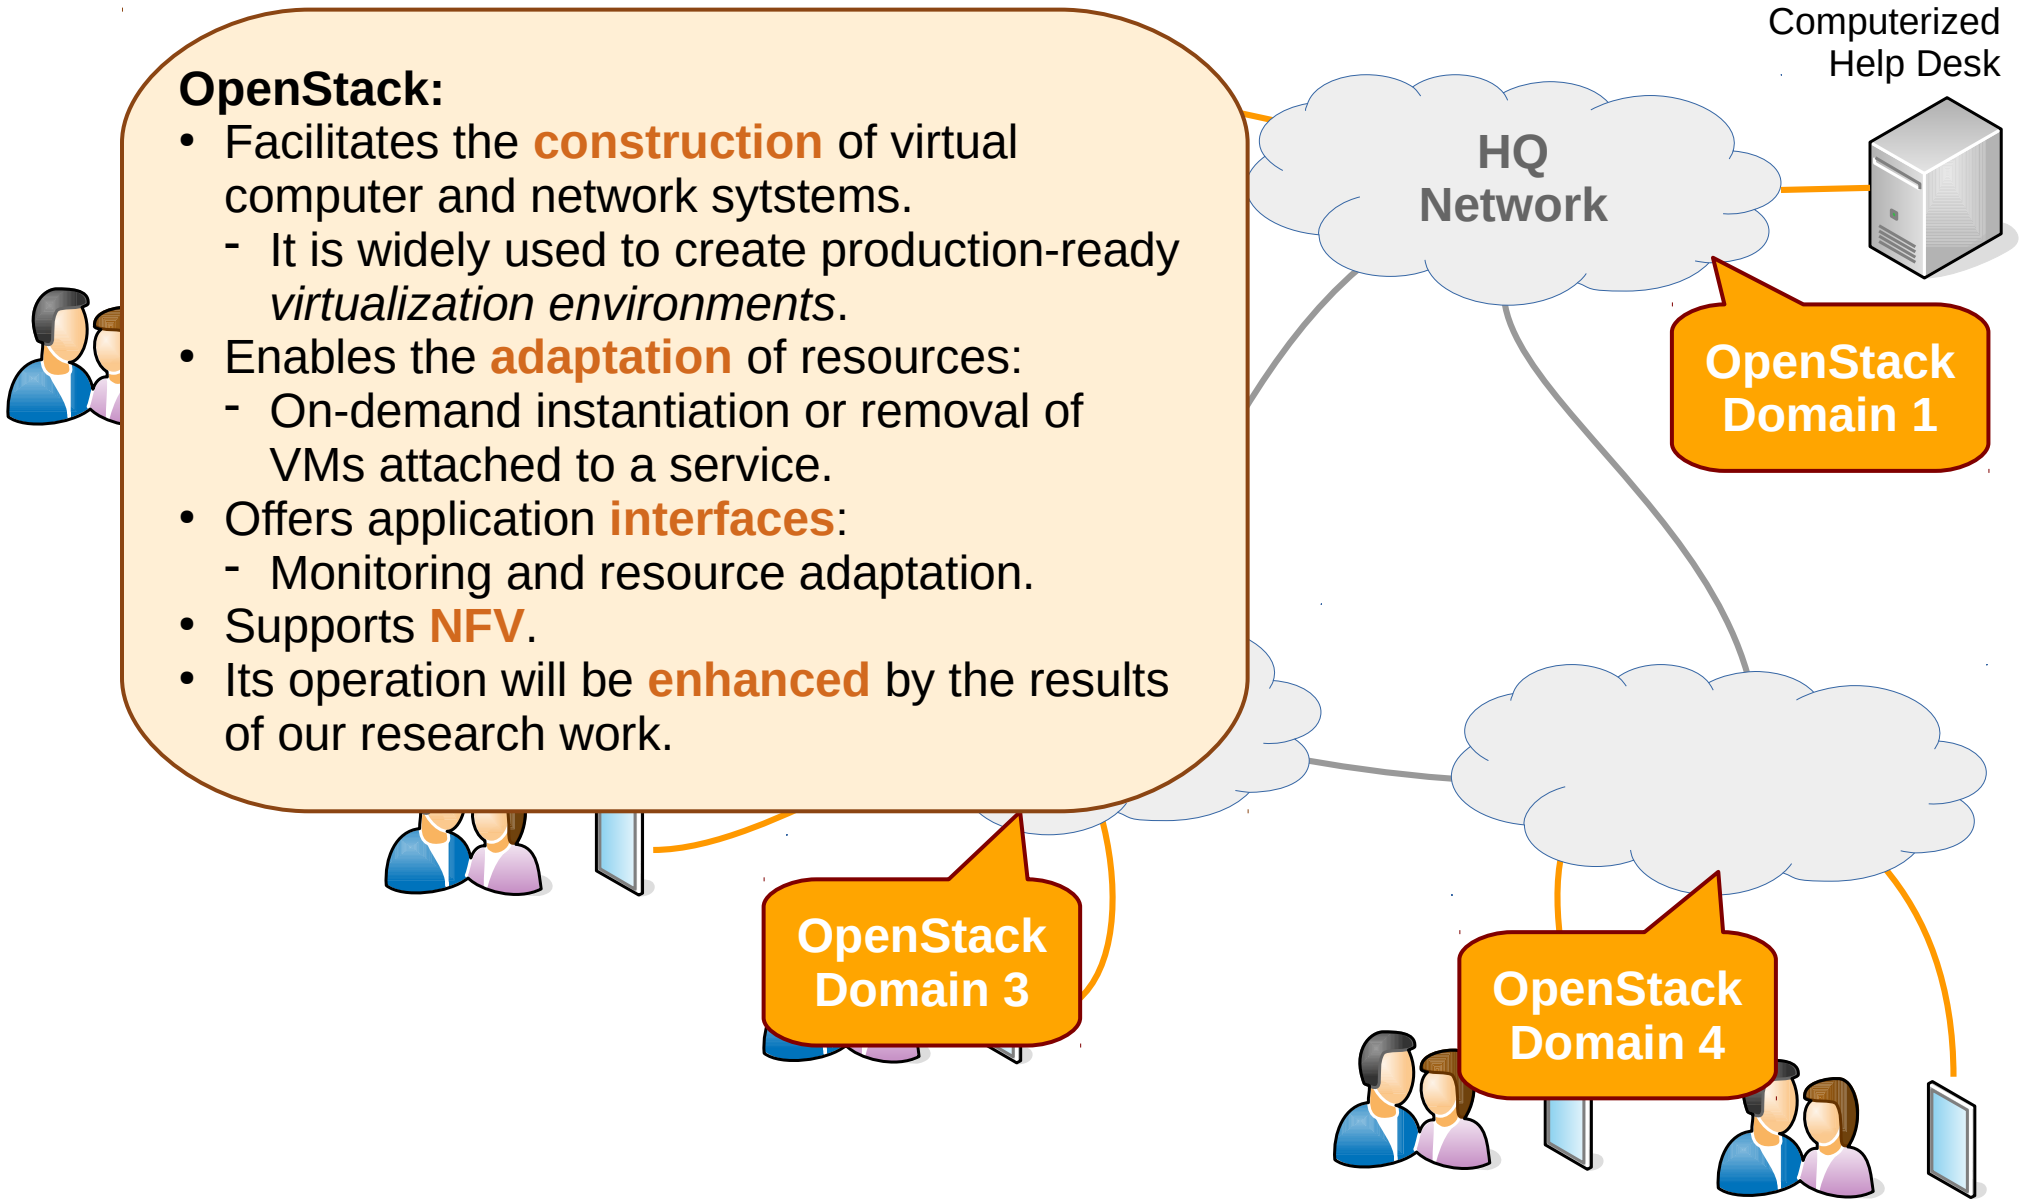

Use Case (IV)

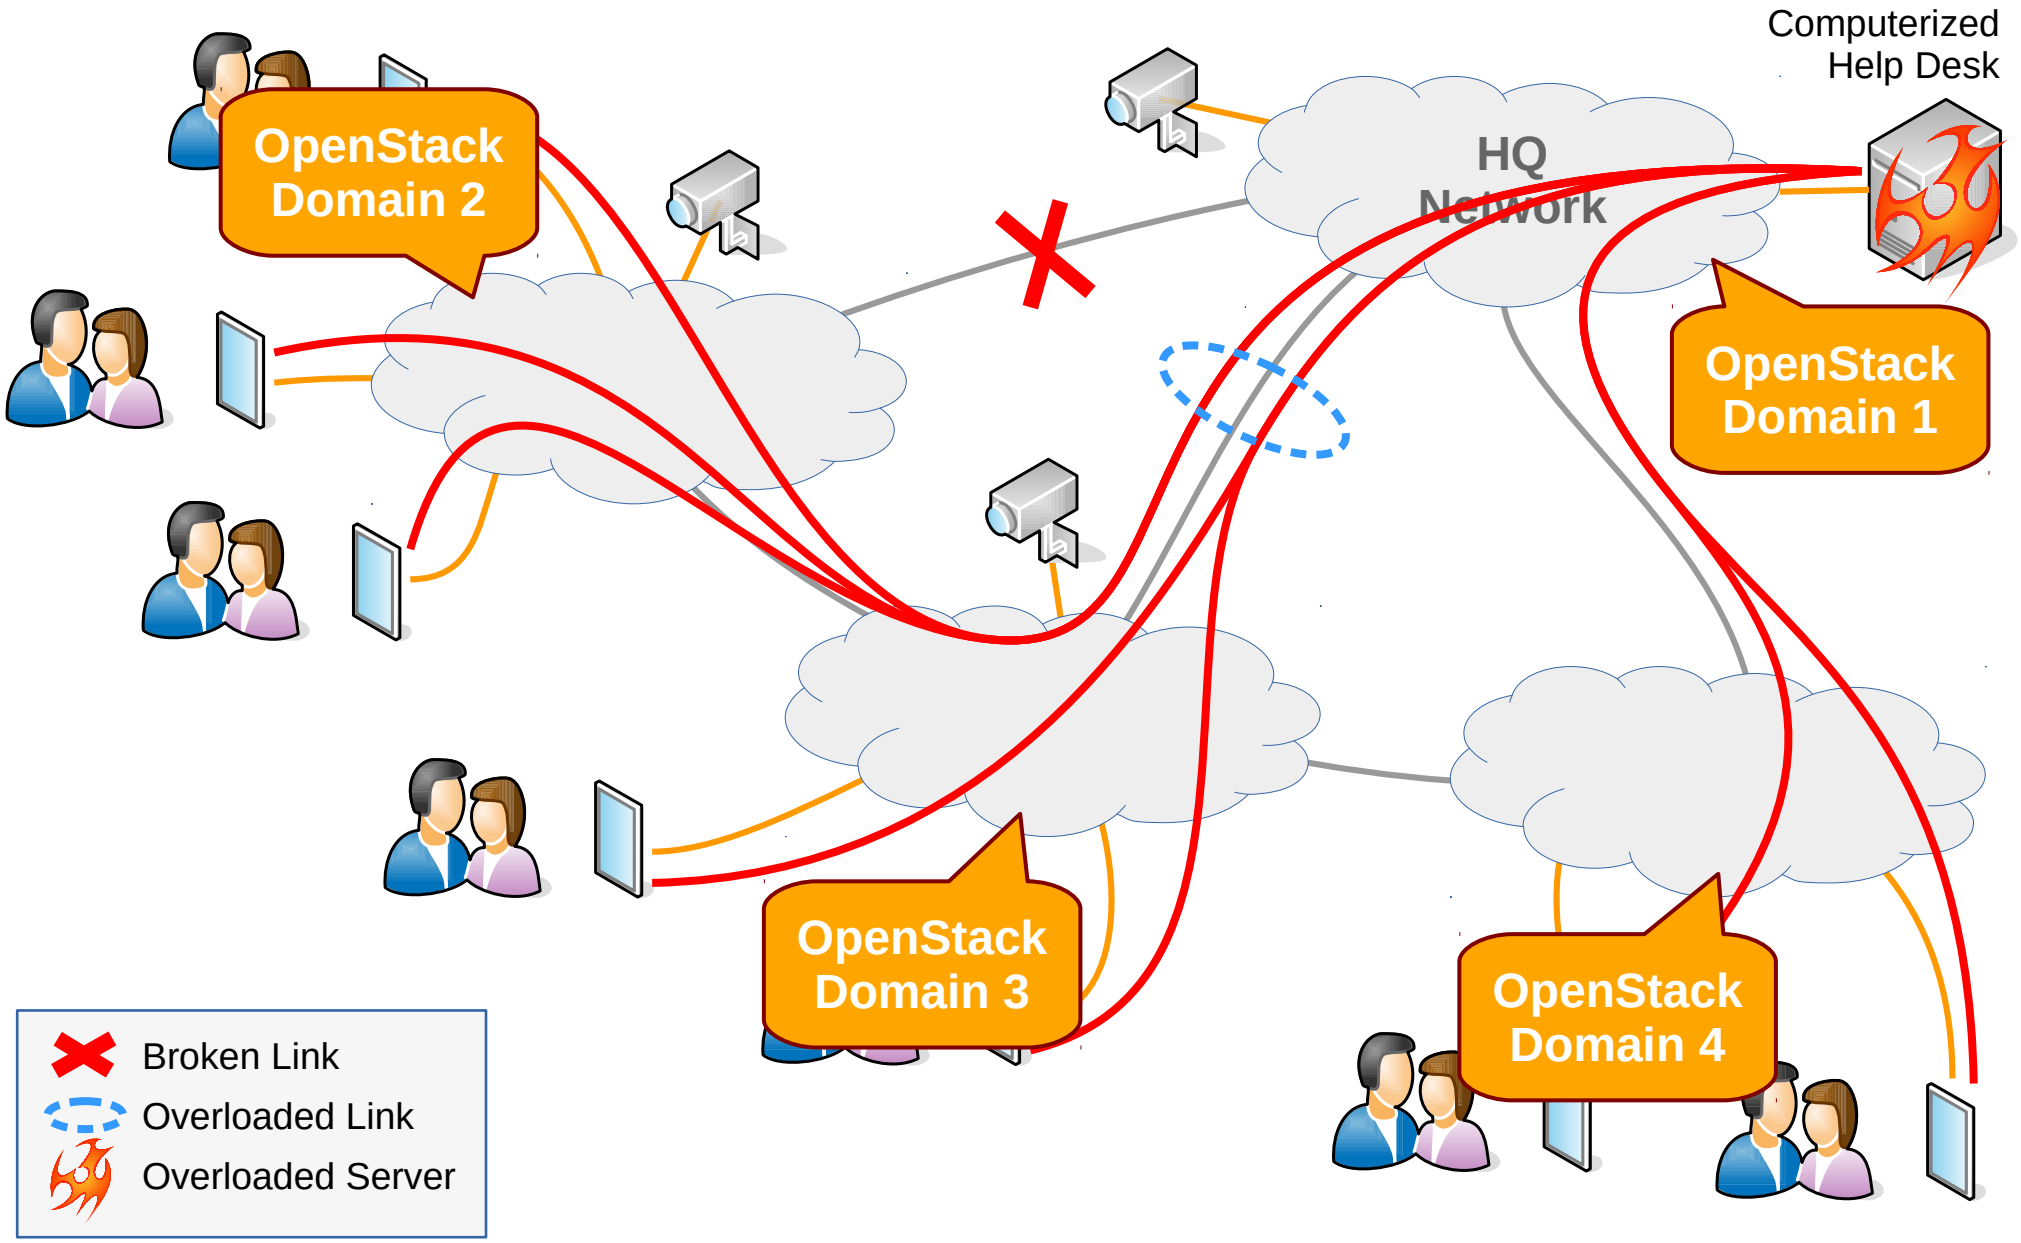

Use Case (V)

Underlying virtualization platforms (e.g. OpenStack), require **long time** (~10 s) to be adapted to new requirements:

- **Some client requests could be rejected.**

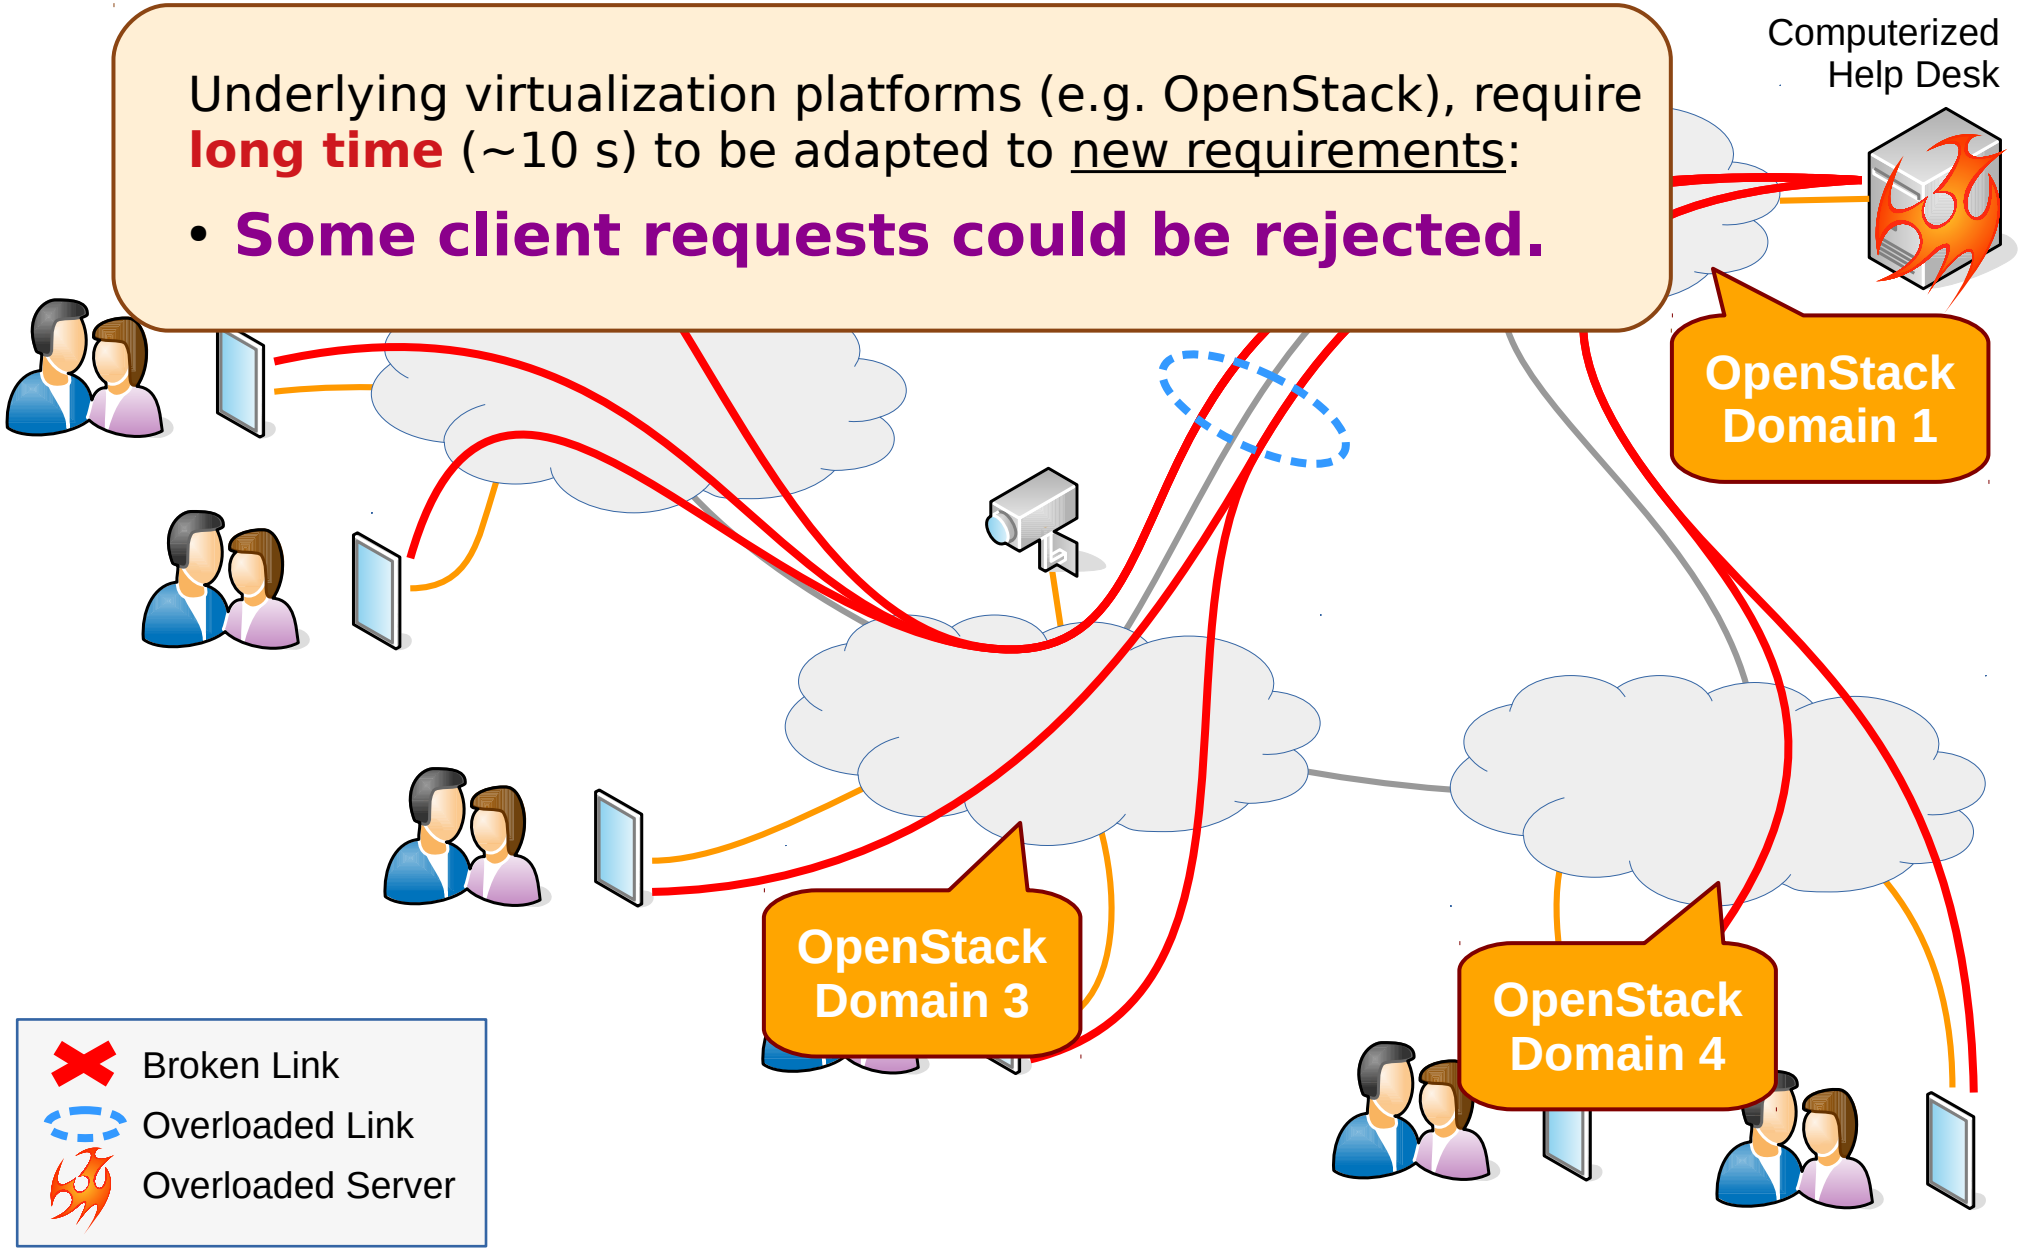

Use Case (VI)

Most changes in requirements are linked to events from outside the system:

- **User response can be derived from event occurrence.**
- **Required resources can be anticipated to reduce adaptation delay by noticing the events as soon as they occur.**
- **The system can be adapted before the client request burst actually reaches the servers.**

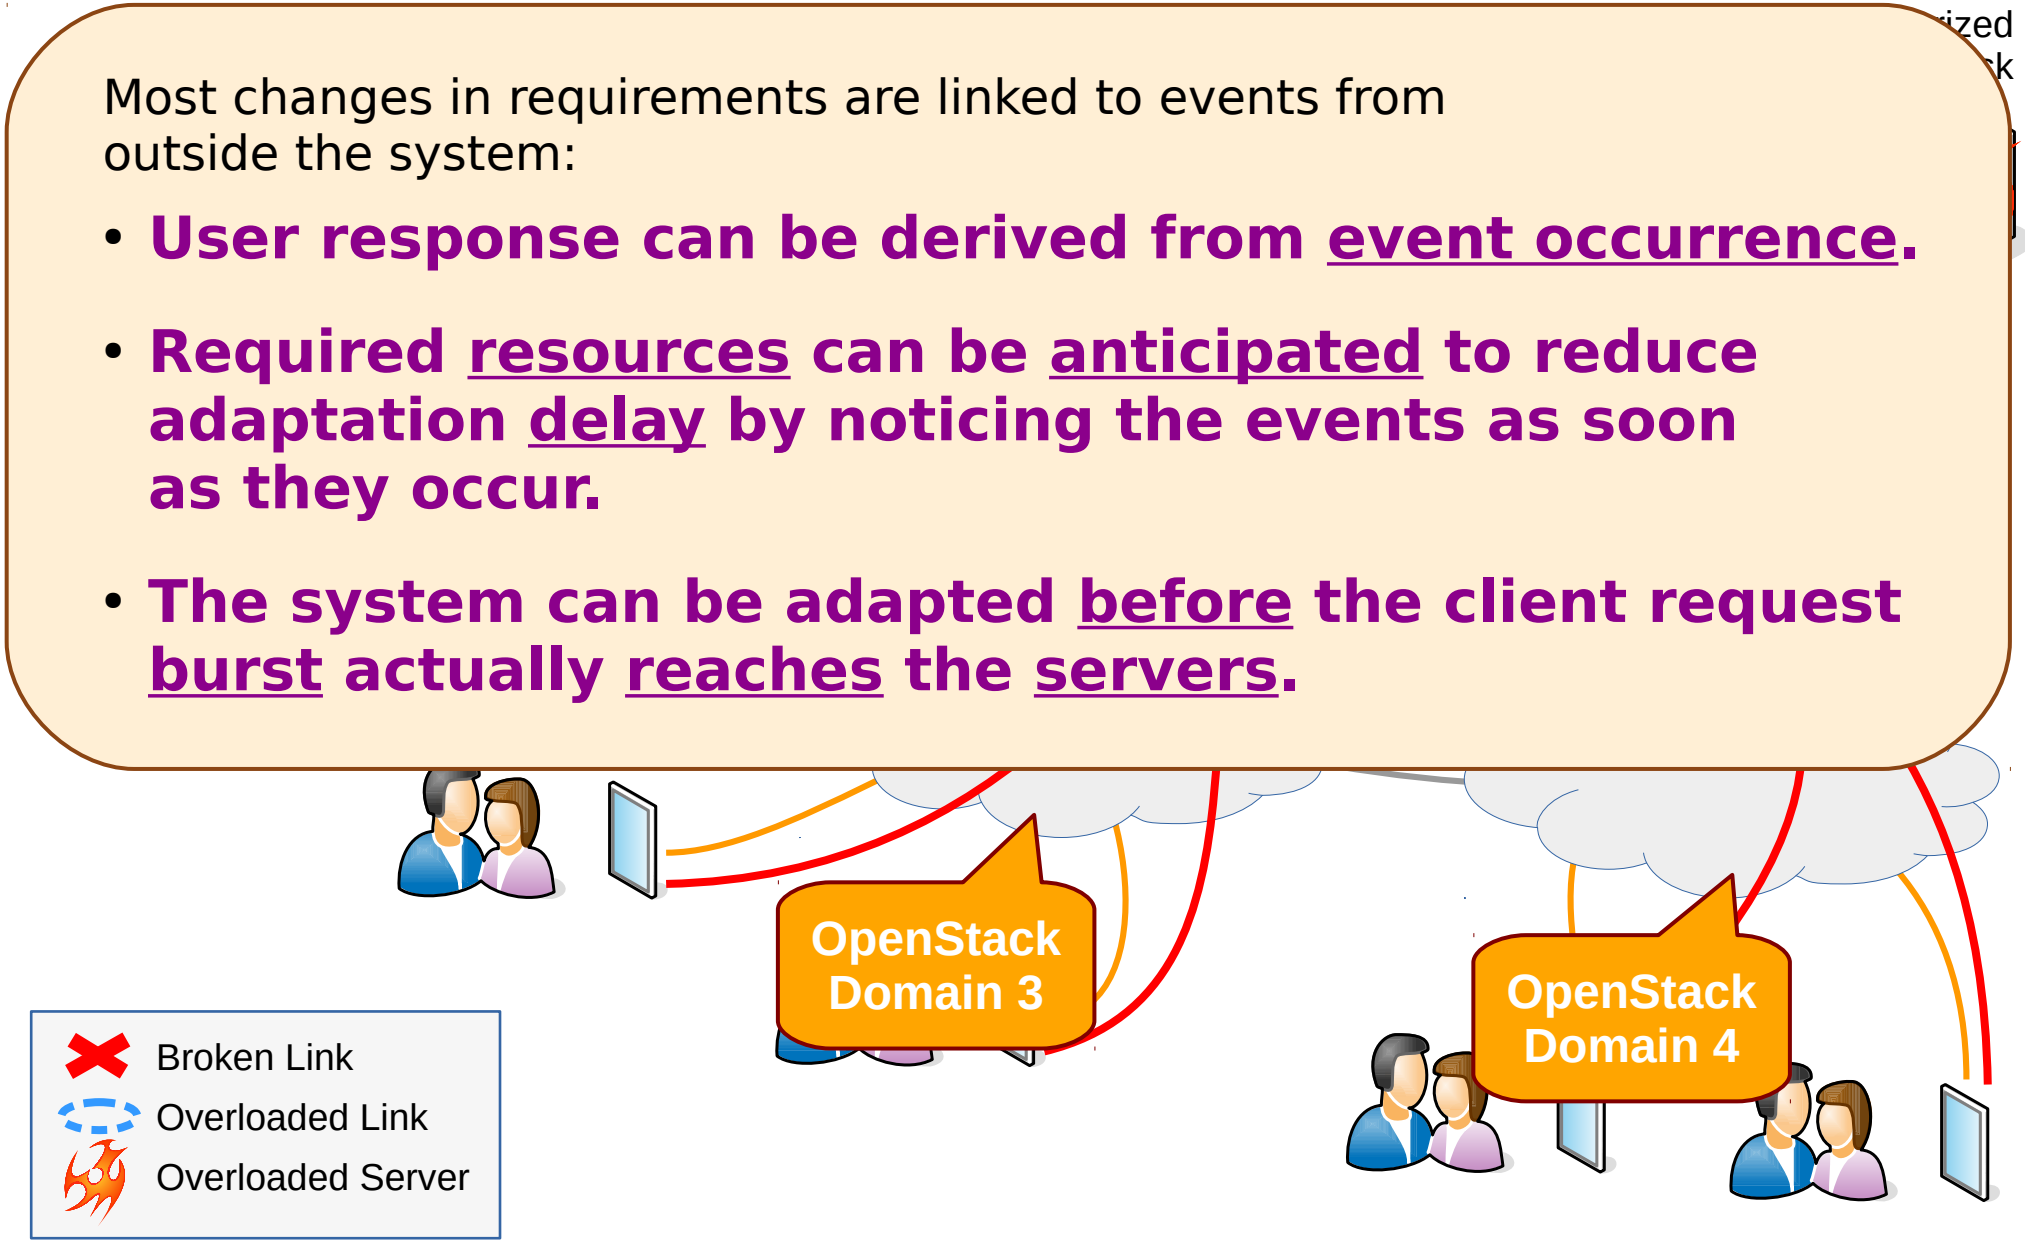

Proposed Approach (I)

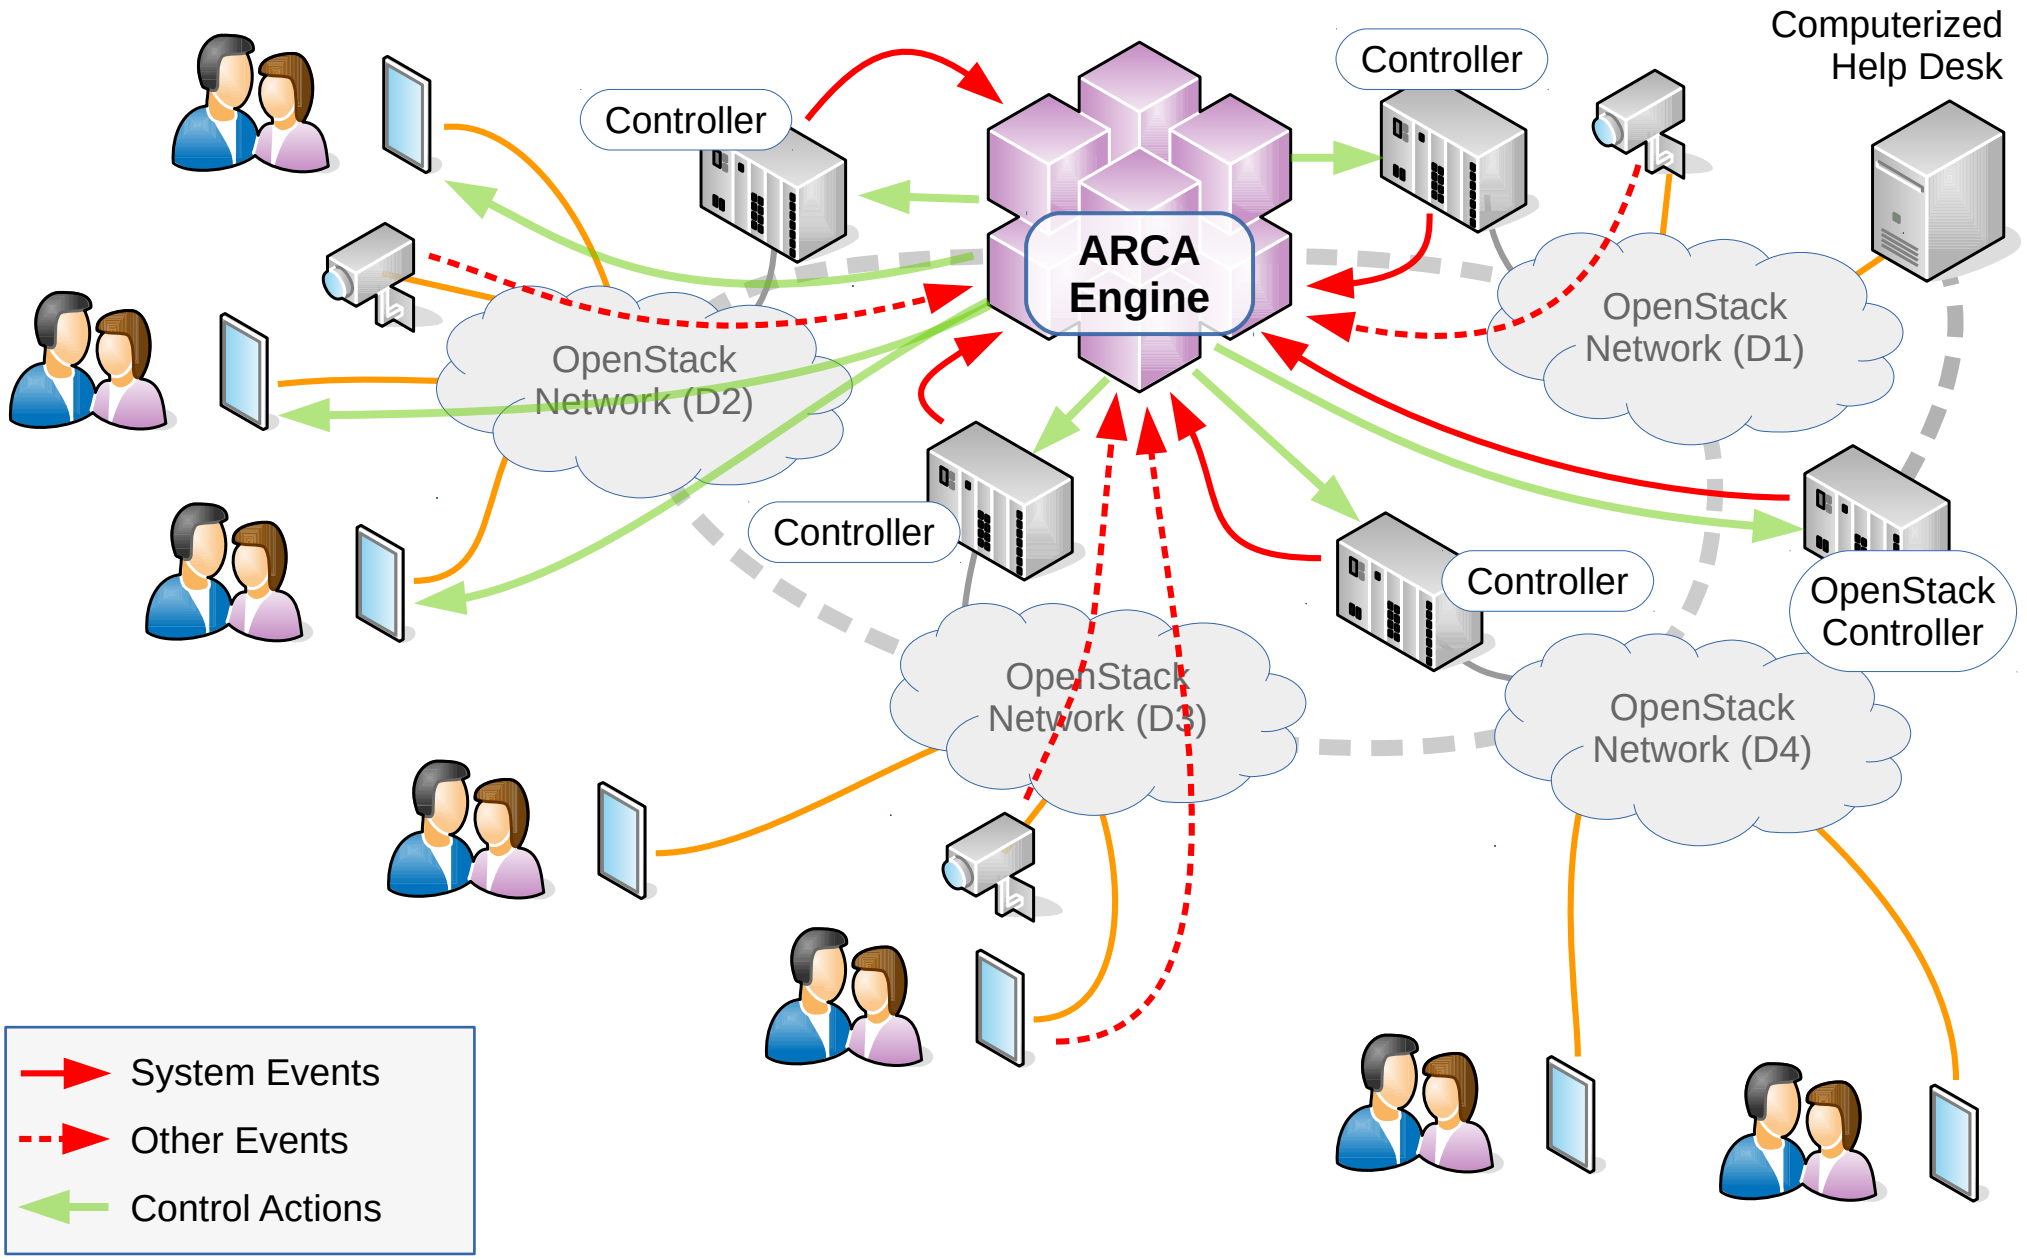

Proposed Approach (II)

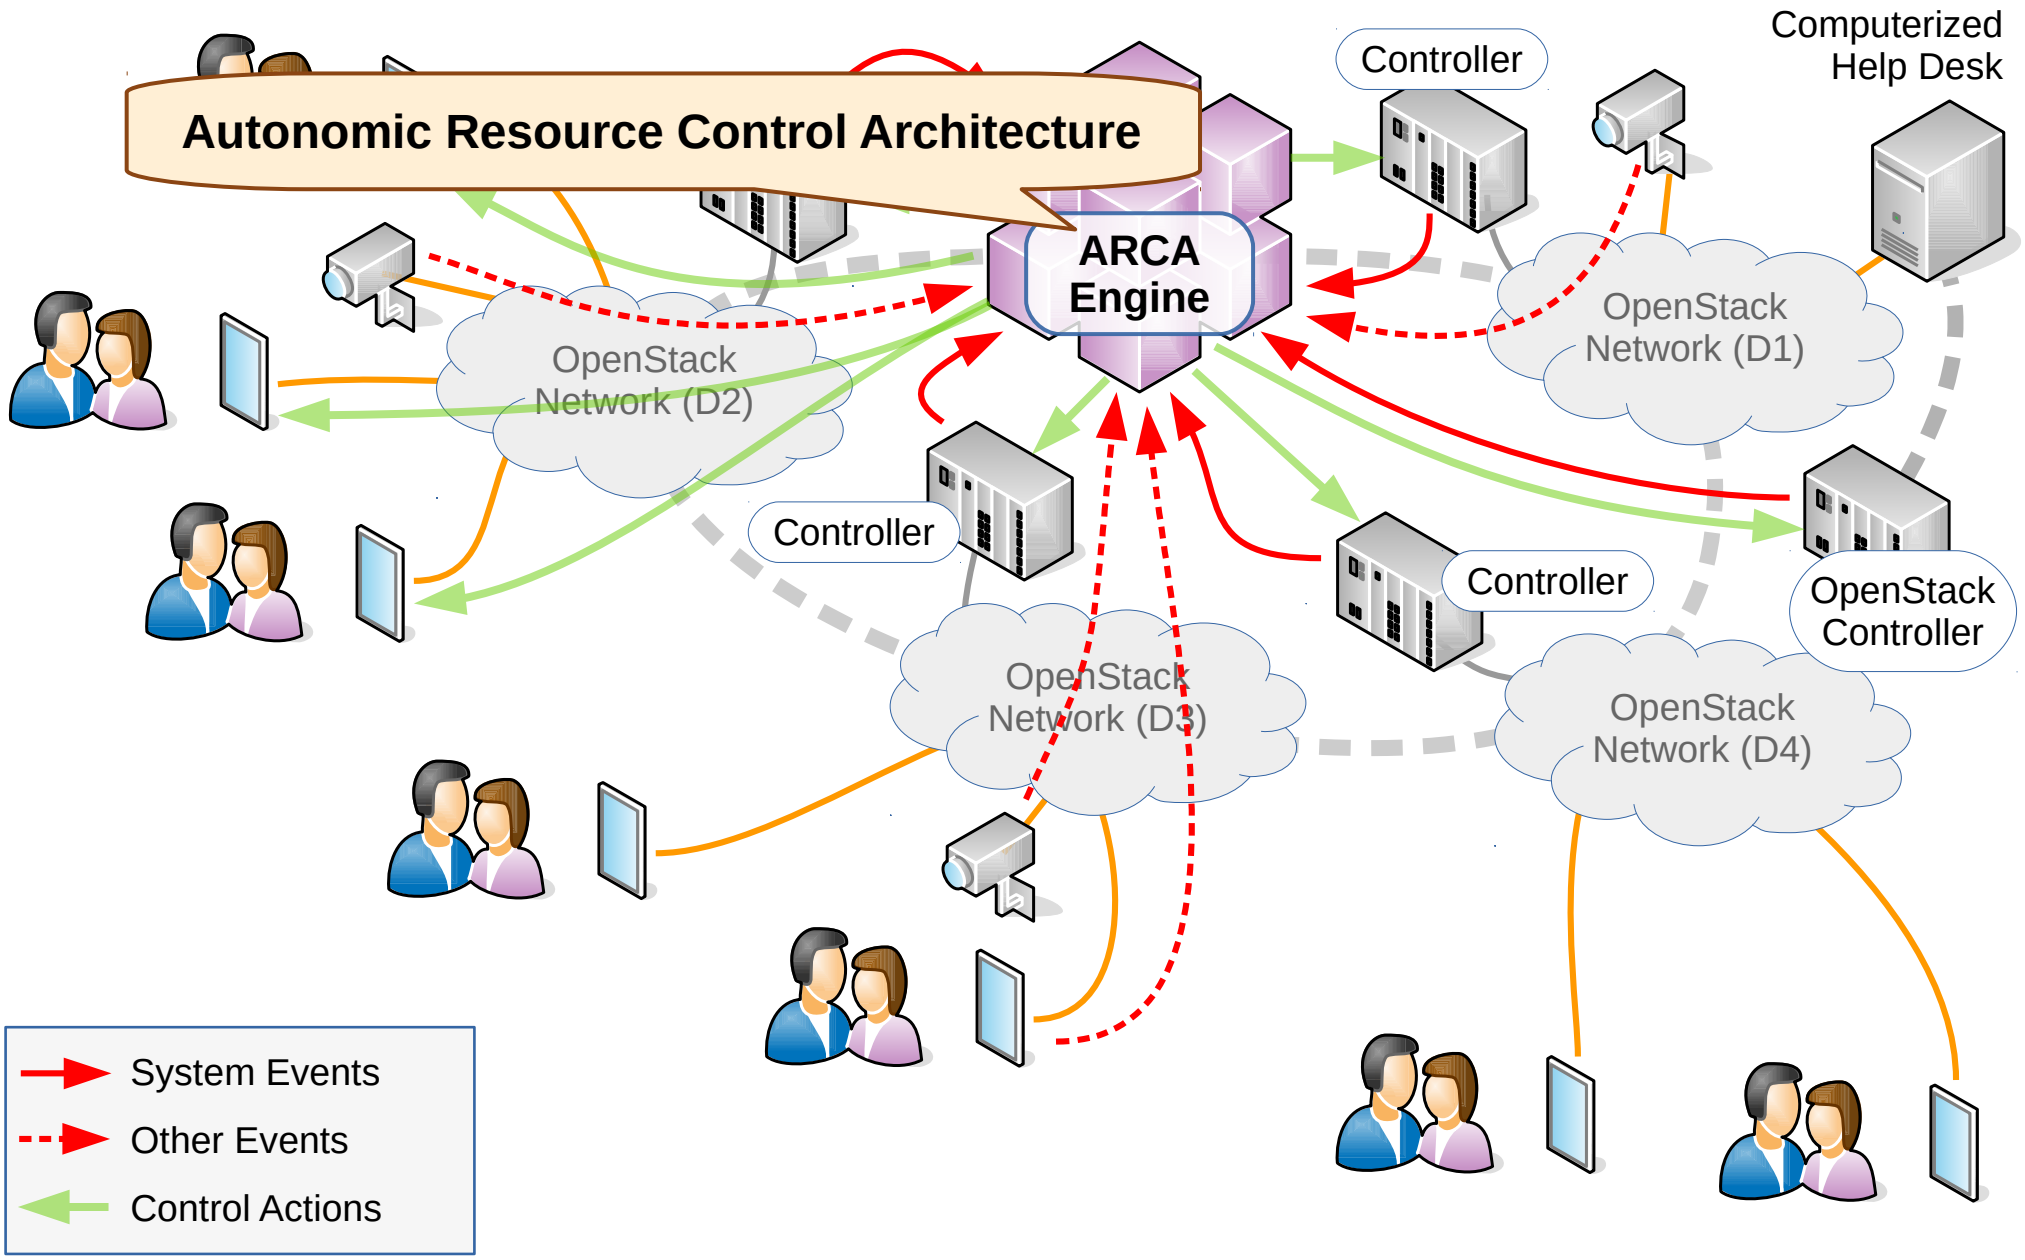

Proposed Approach (III)

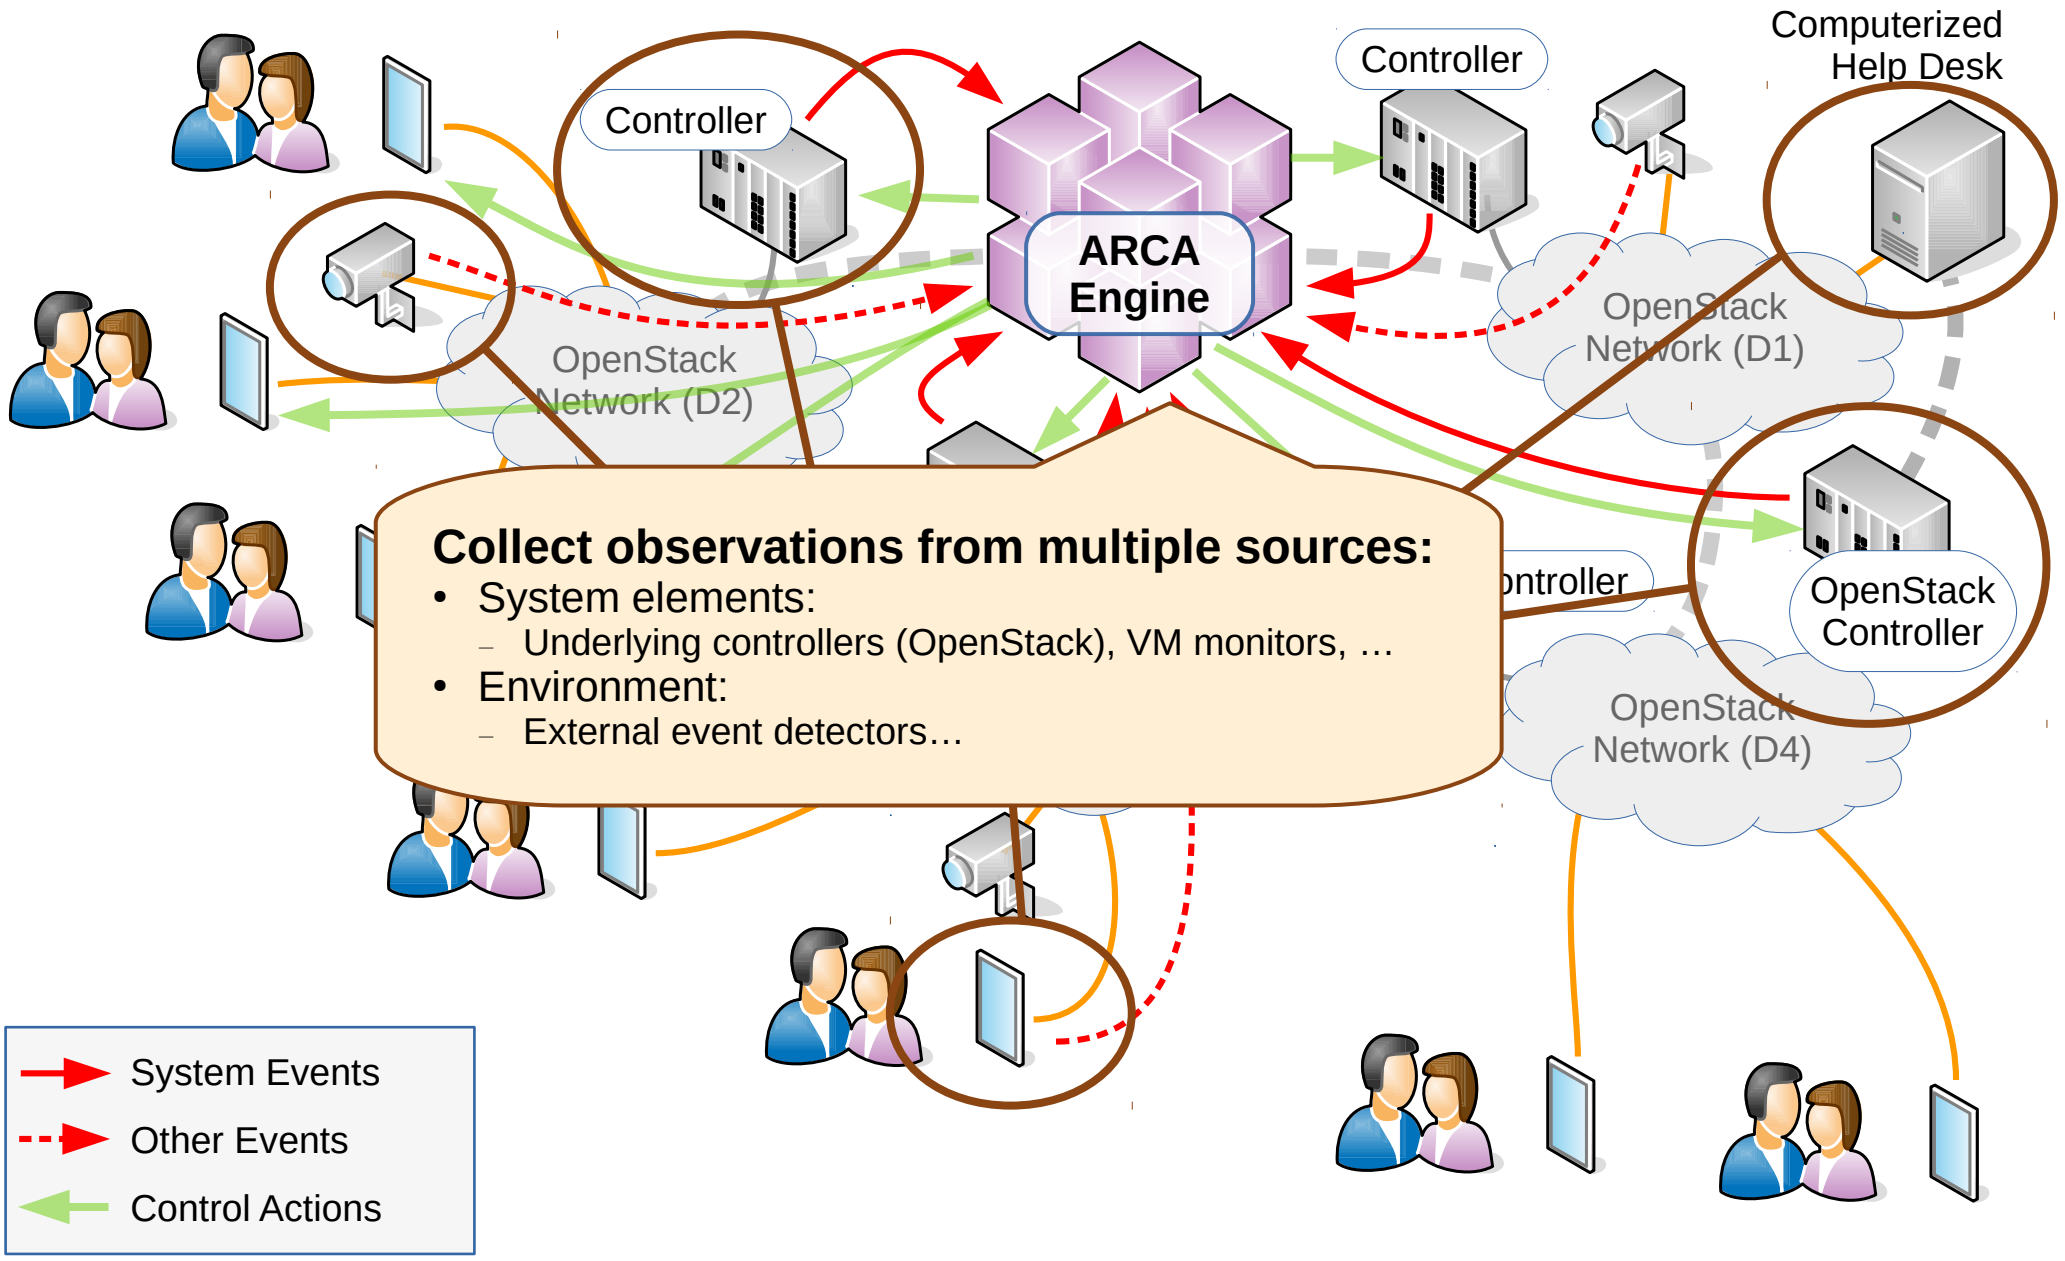

Proposed Approach (IV)

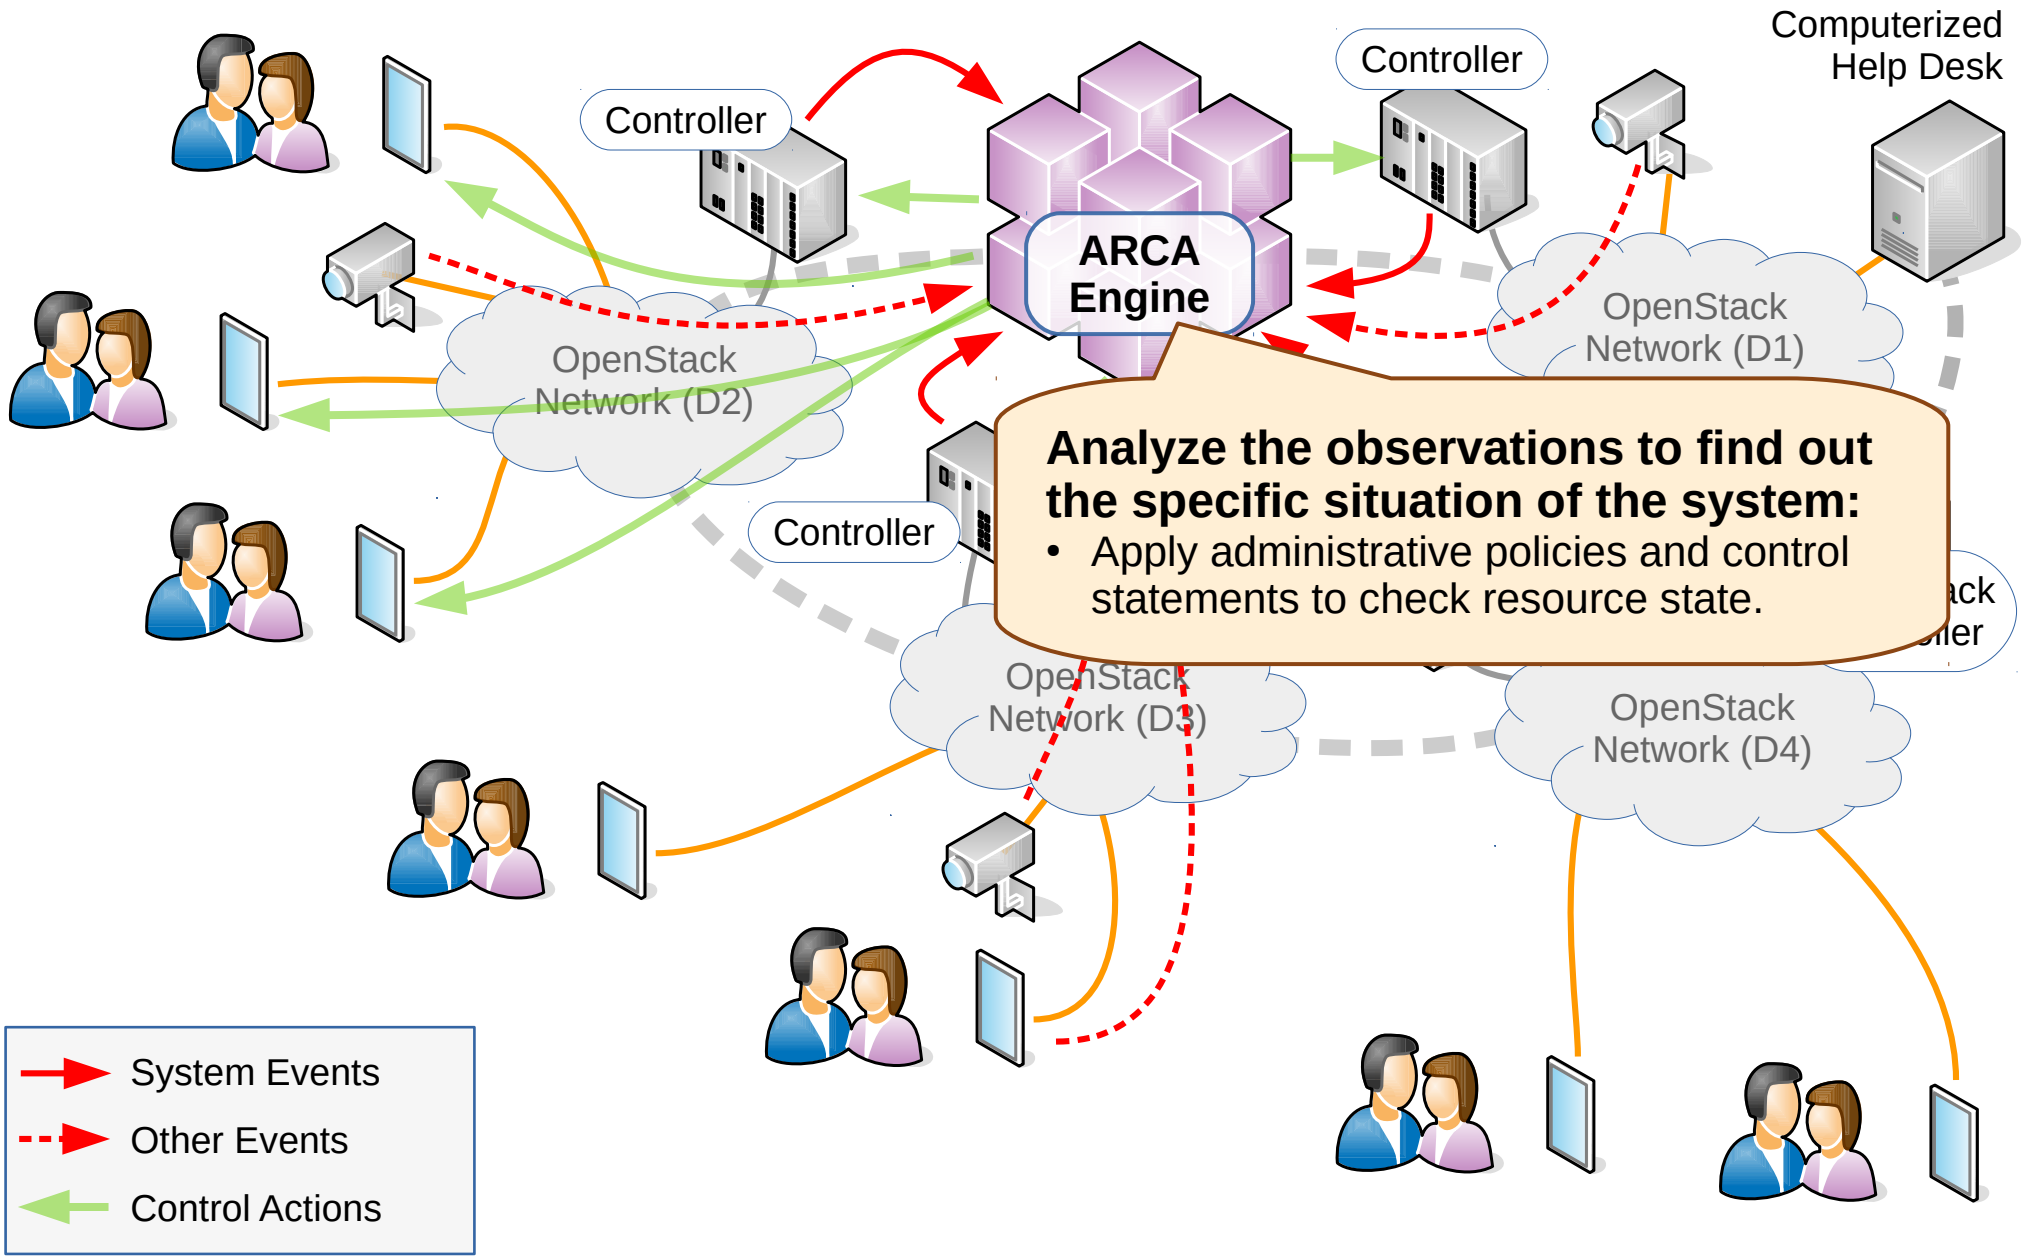

Proposed Approach (V)

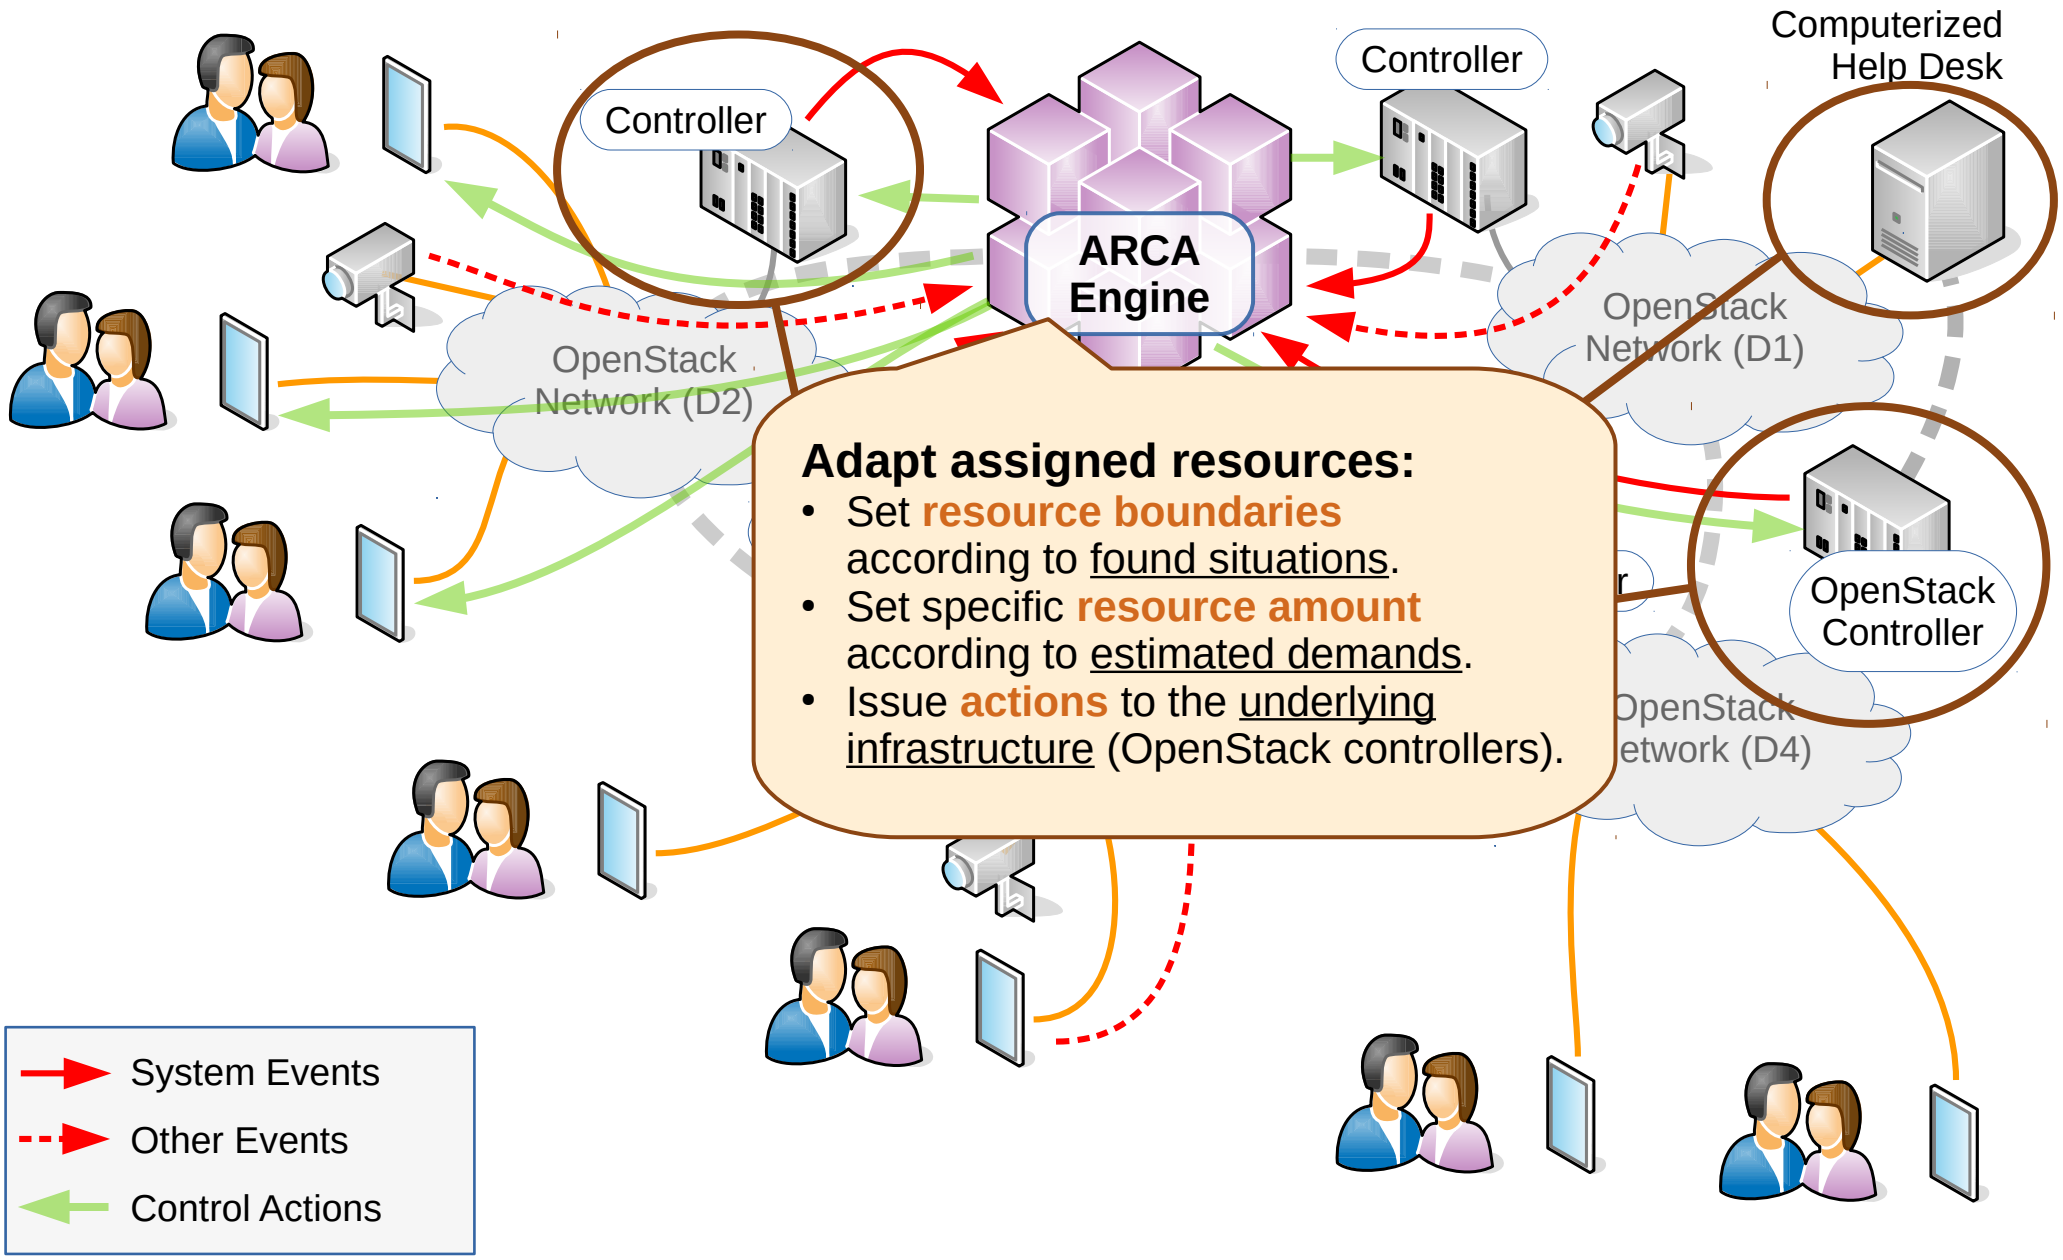

Proposed Approach (VI)

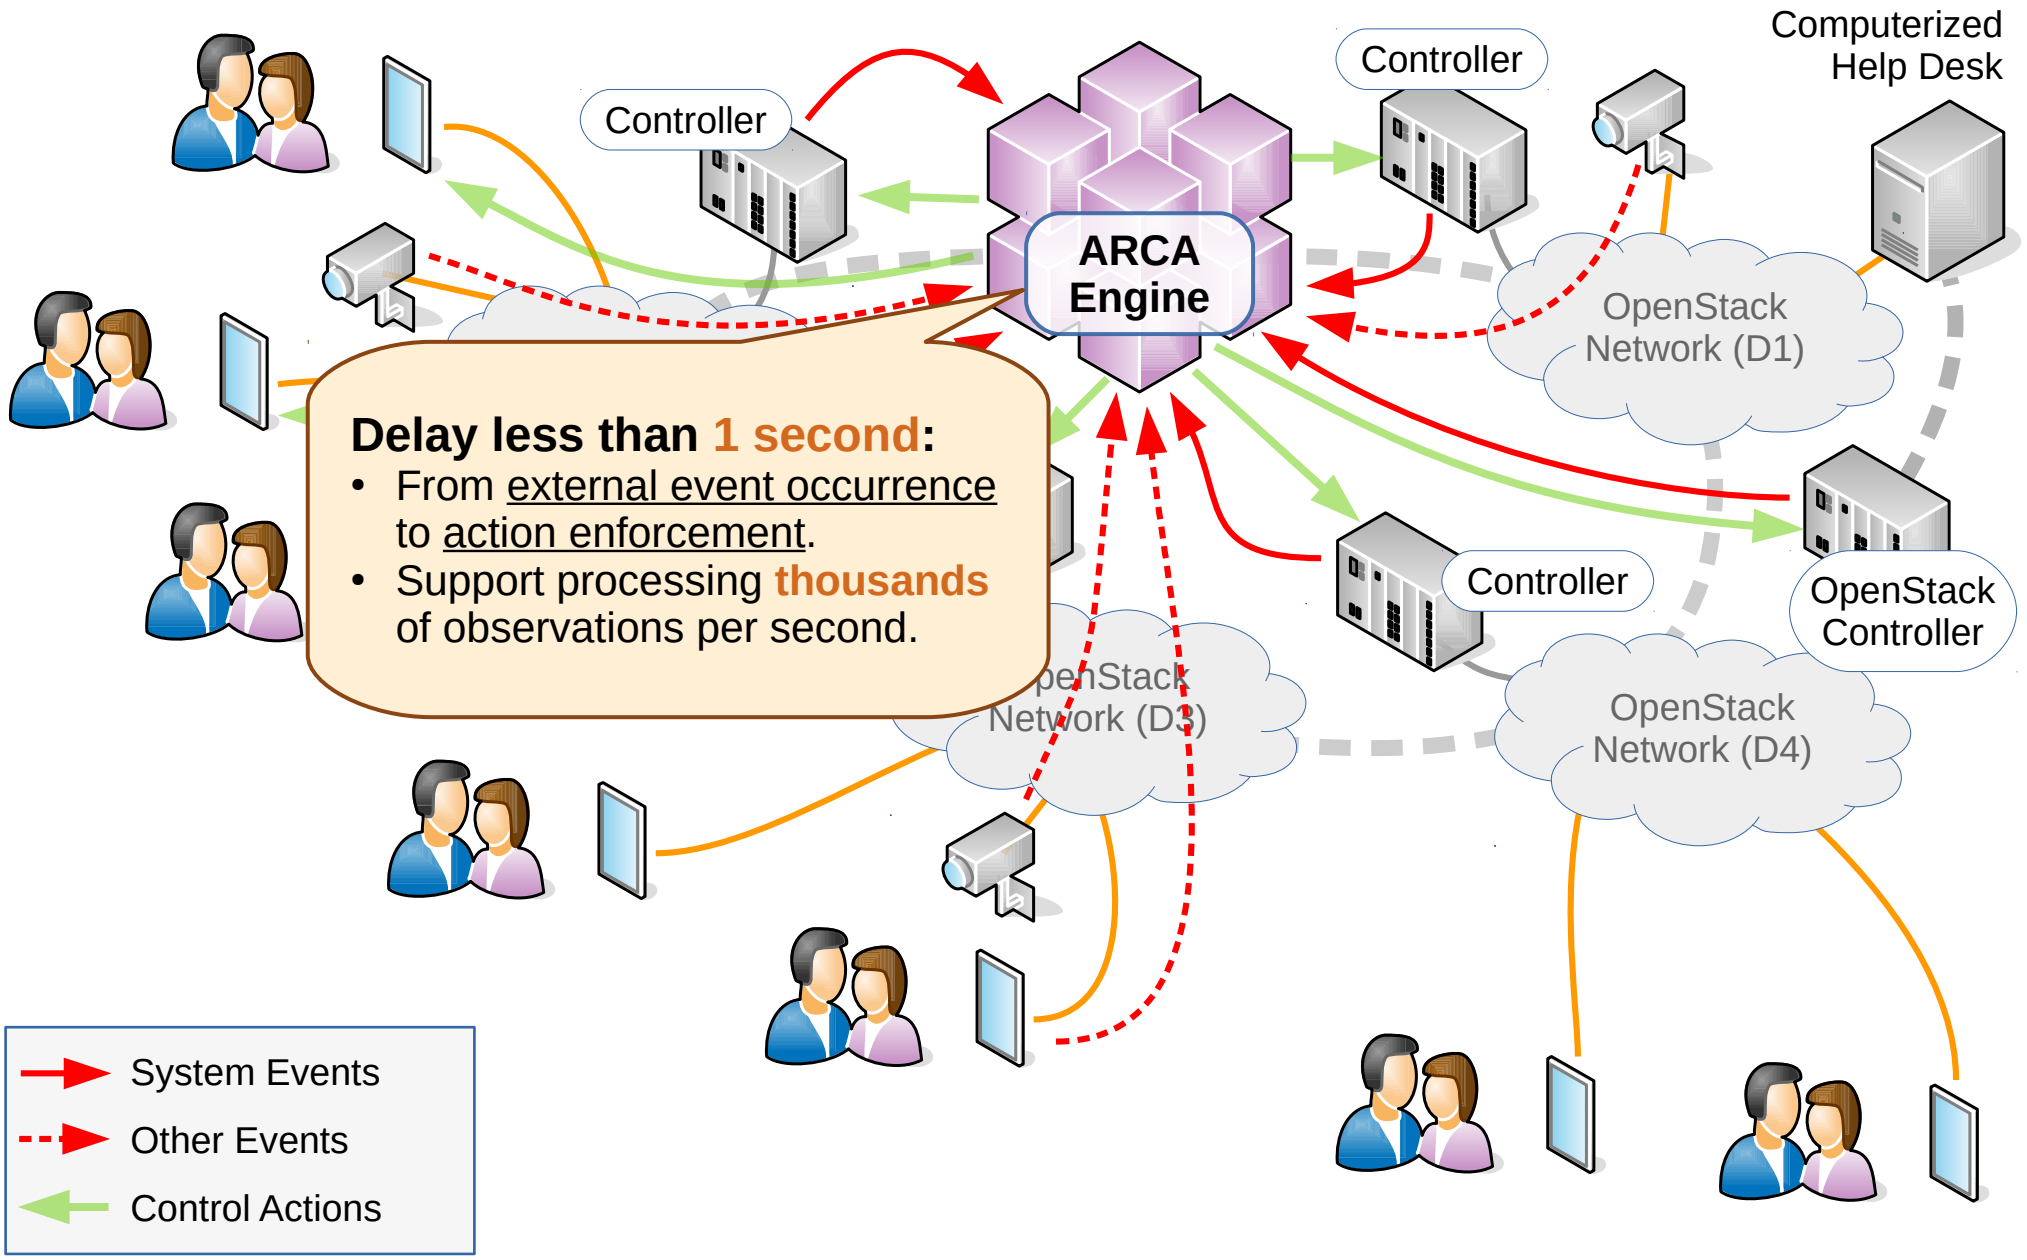

Proposed Approach (VII)

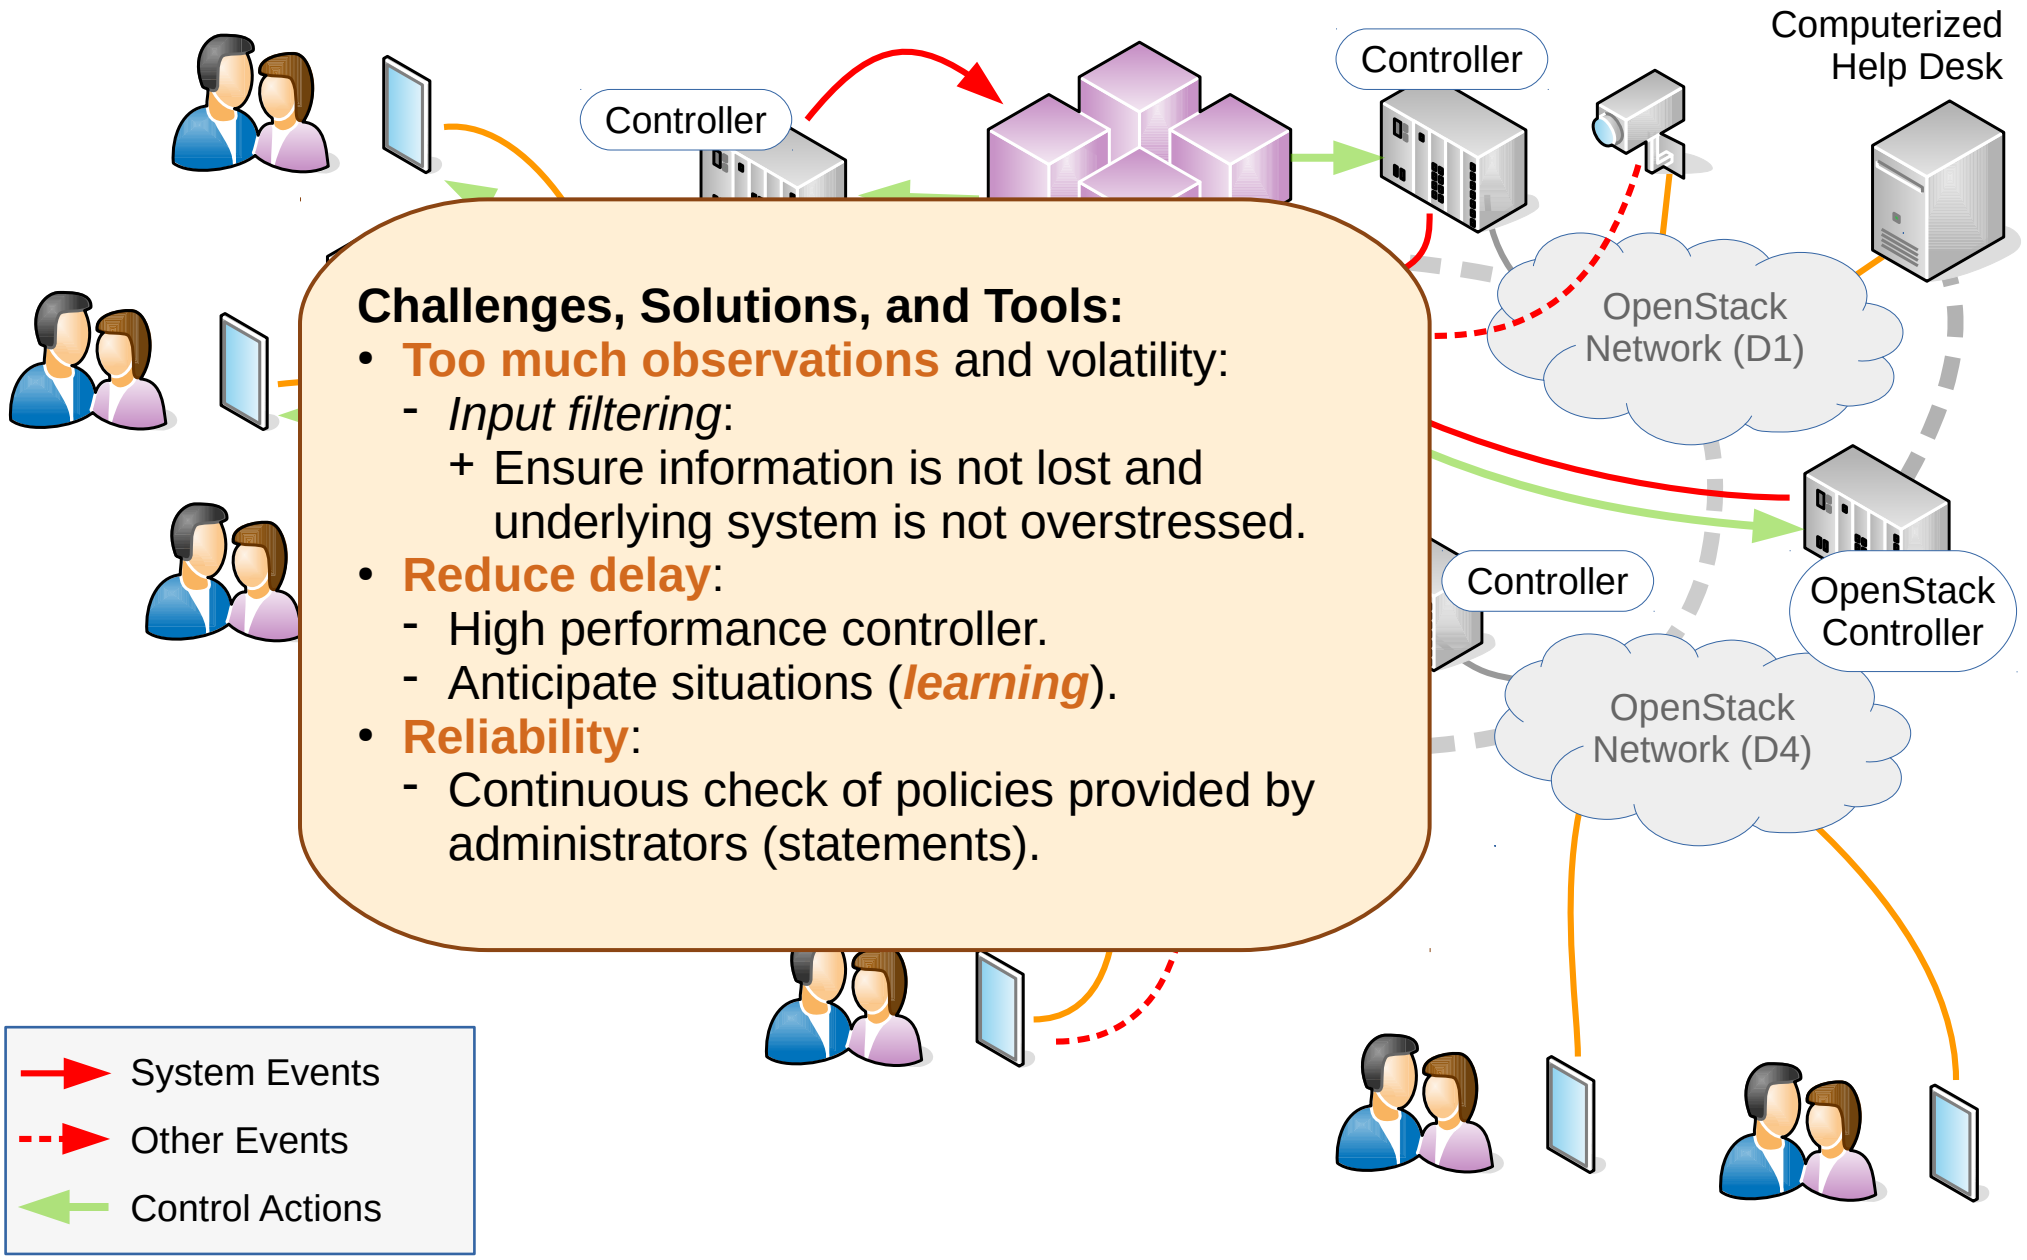

Challenges, Solutions, and Tools:

- **Too much observations** and volatility:
  - *Input filtering*:
    - + Ensure information is not lost and underlying system is not overstressed.
- **Reduce delay**:
  - High performance controller.
  - Anticipate situations (***learning***).
- **Reliability**:
  - Continuous check of policies provided by administrators (statements).

Overview of ARCA (I)

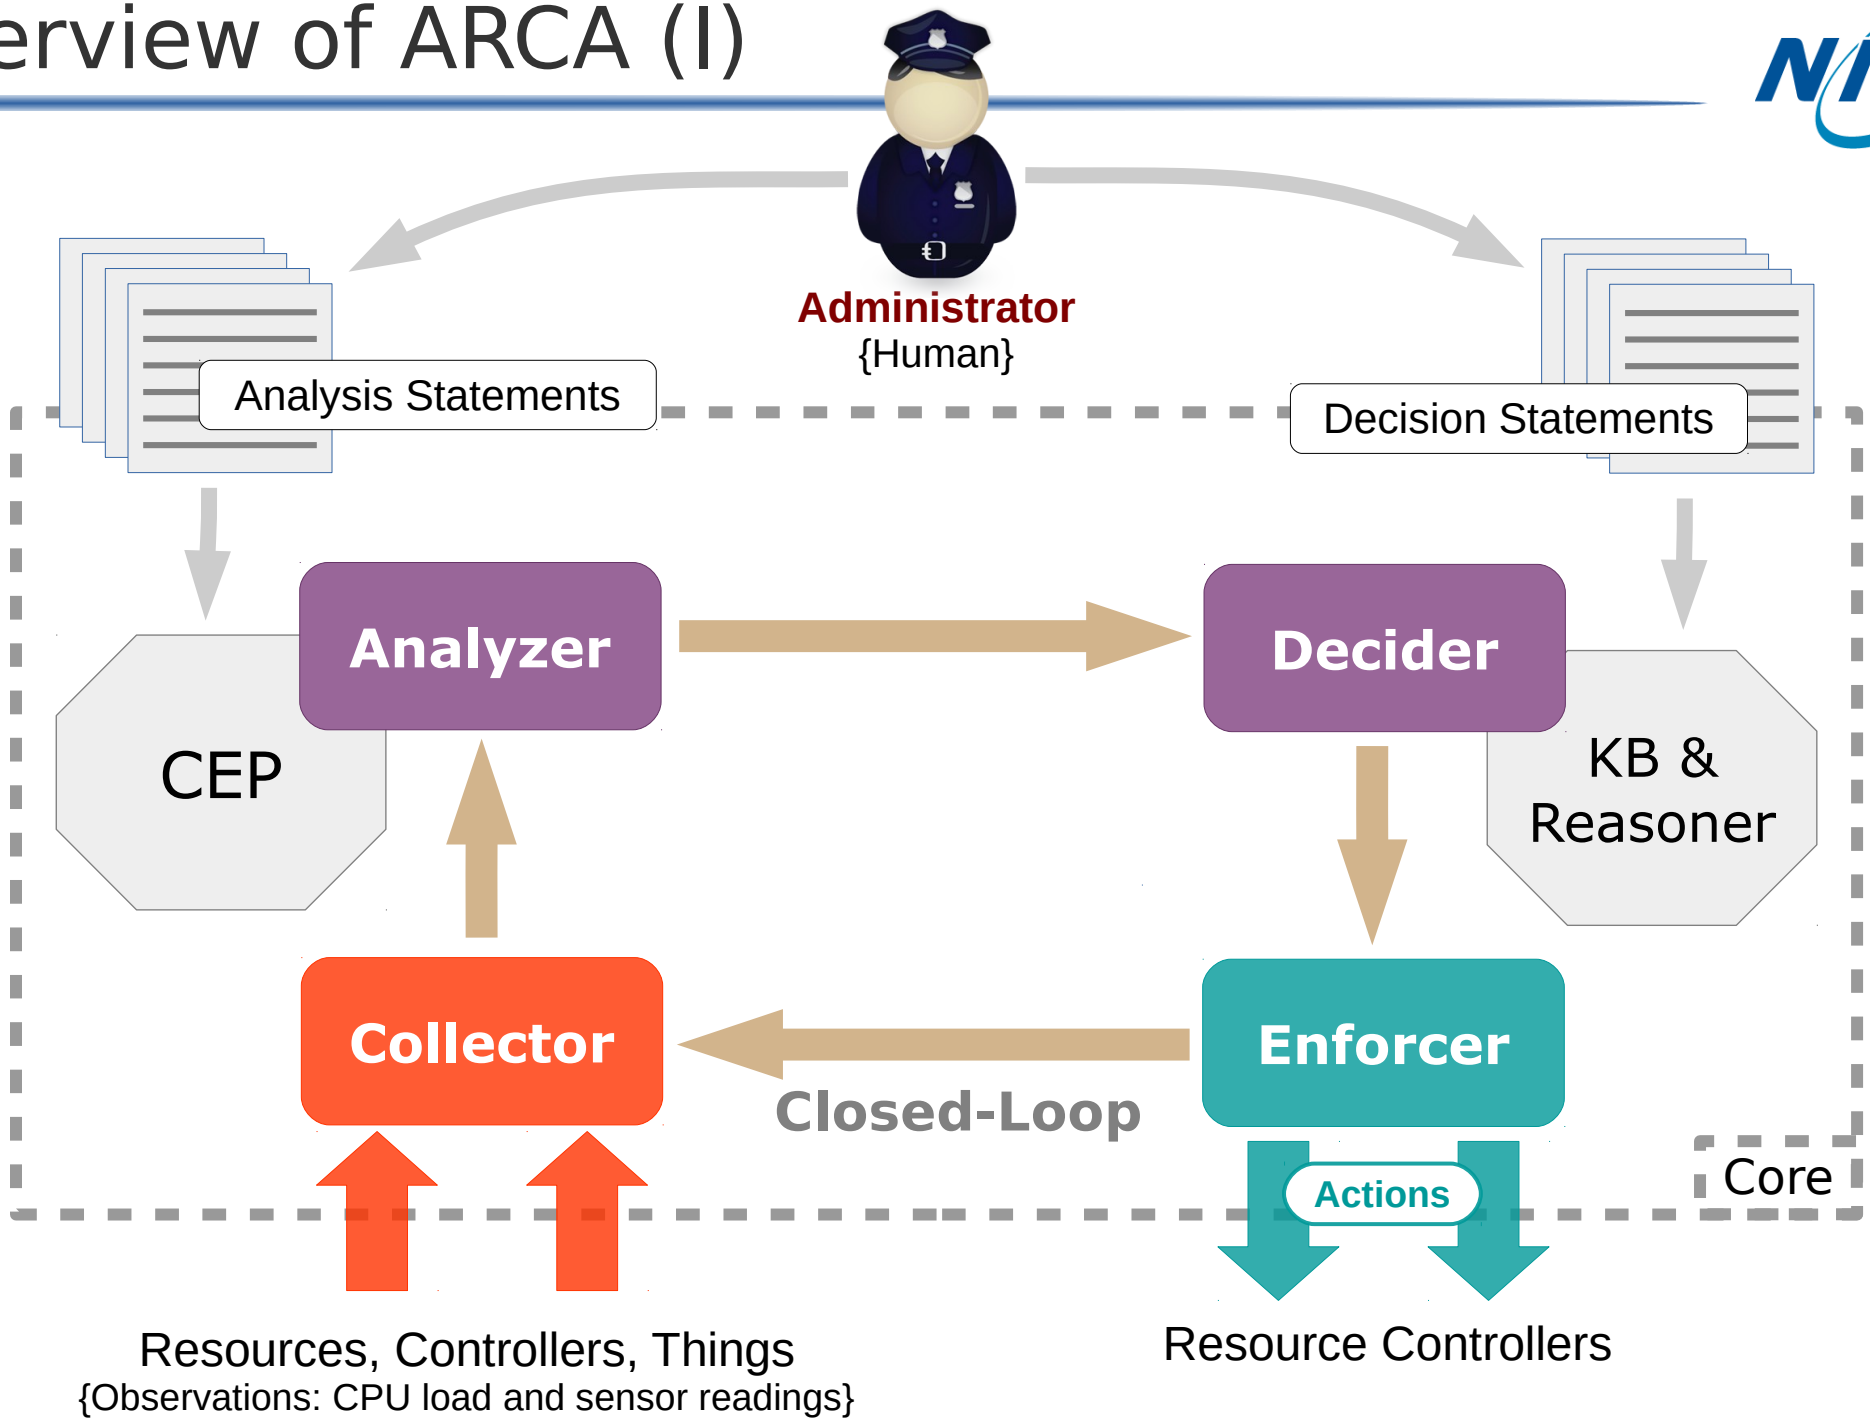

Overview of ARCA (II)

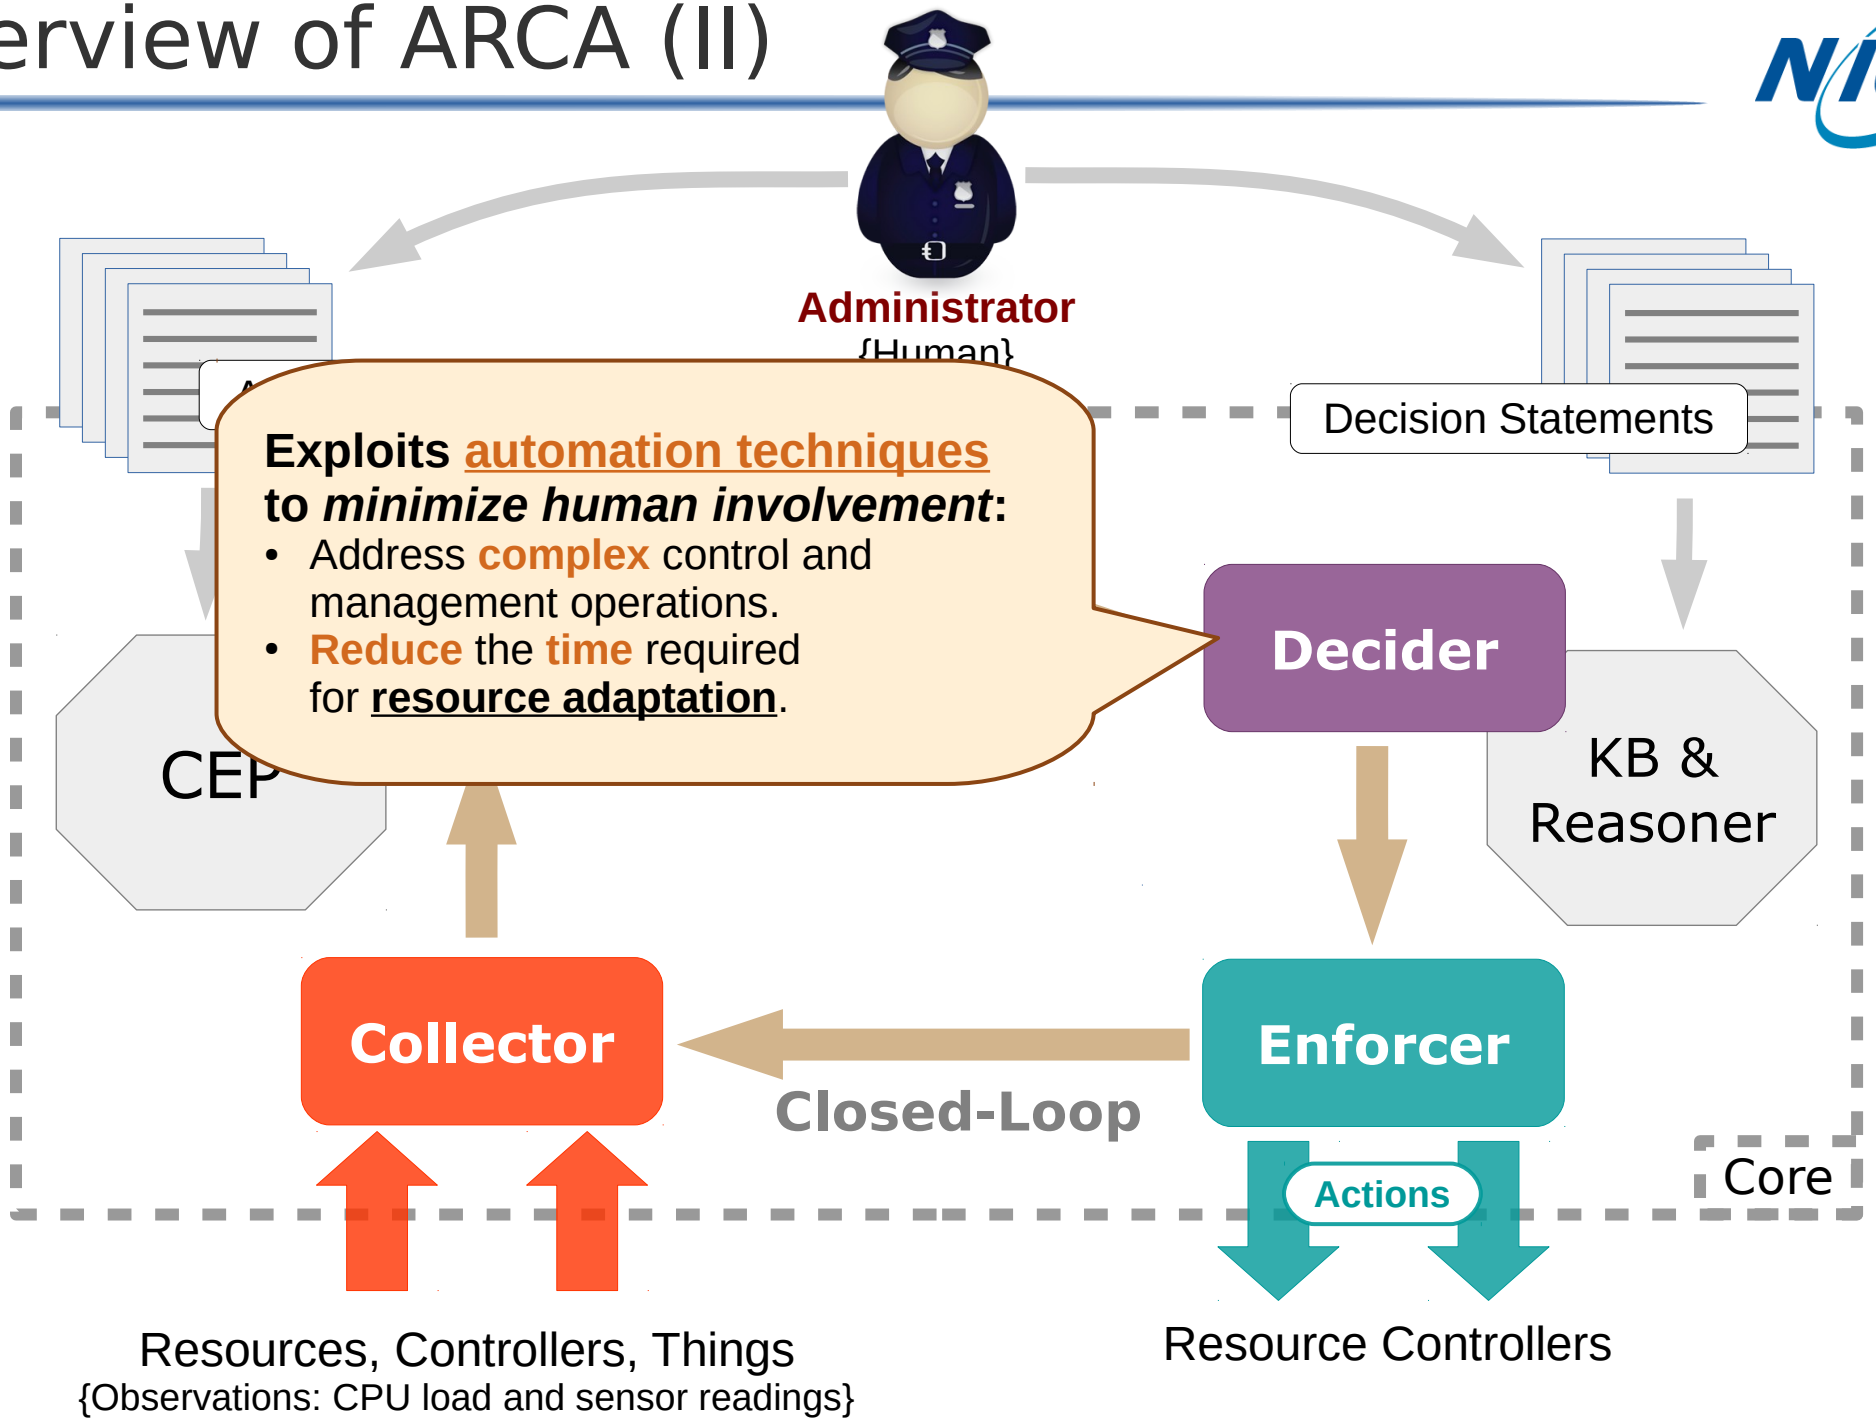

Overview of ARCA (III)

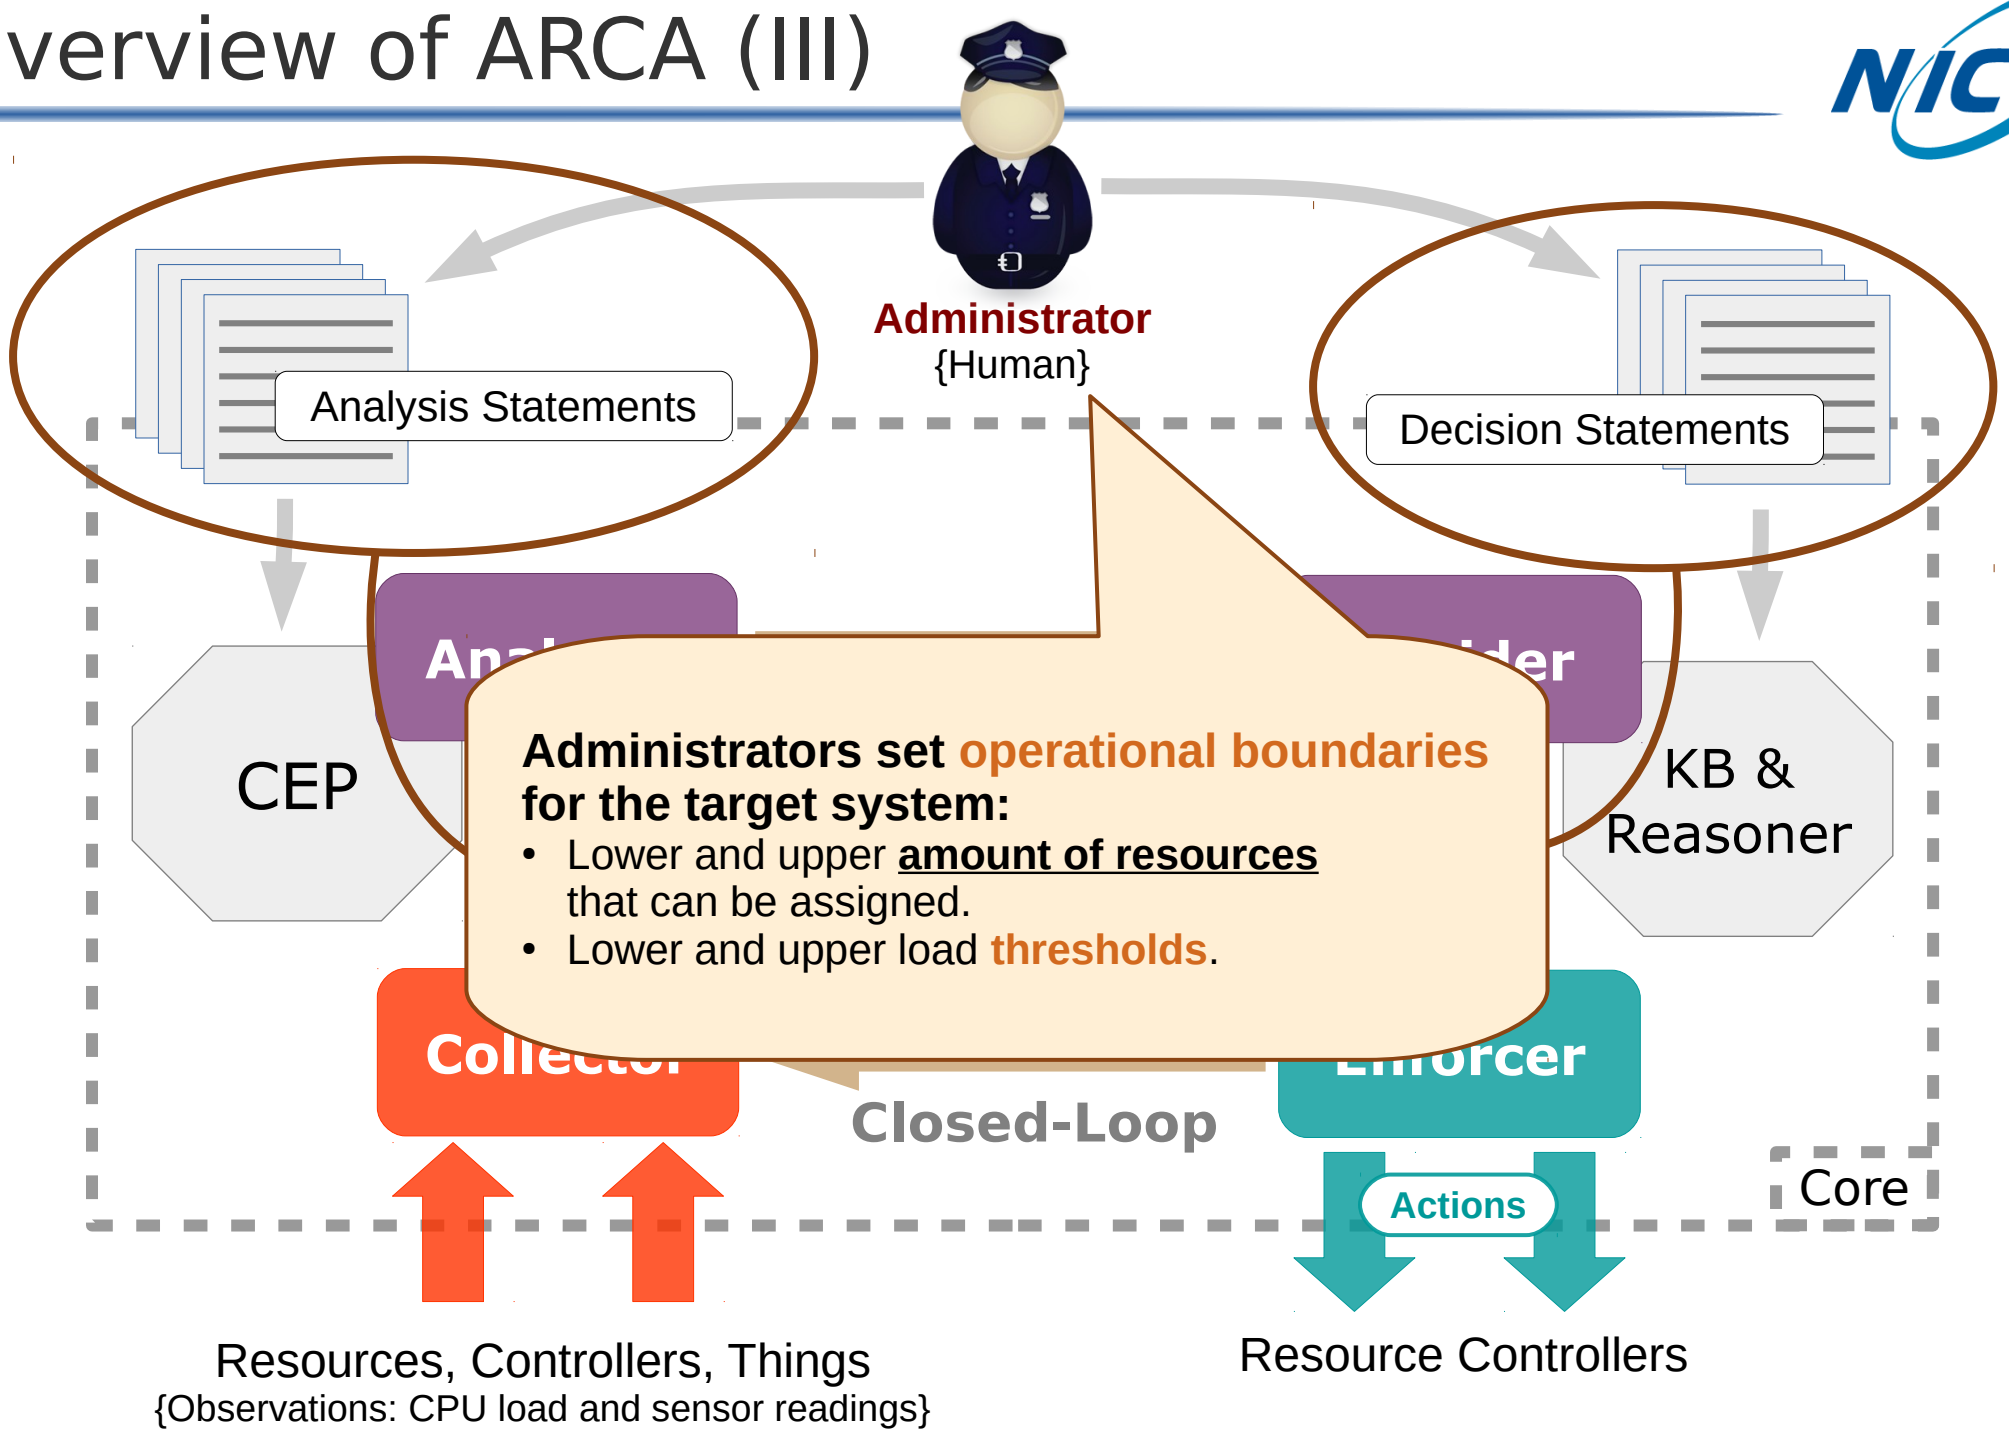

Overview of ARCA (IV)

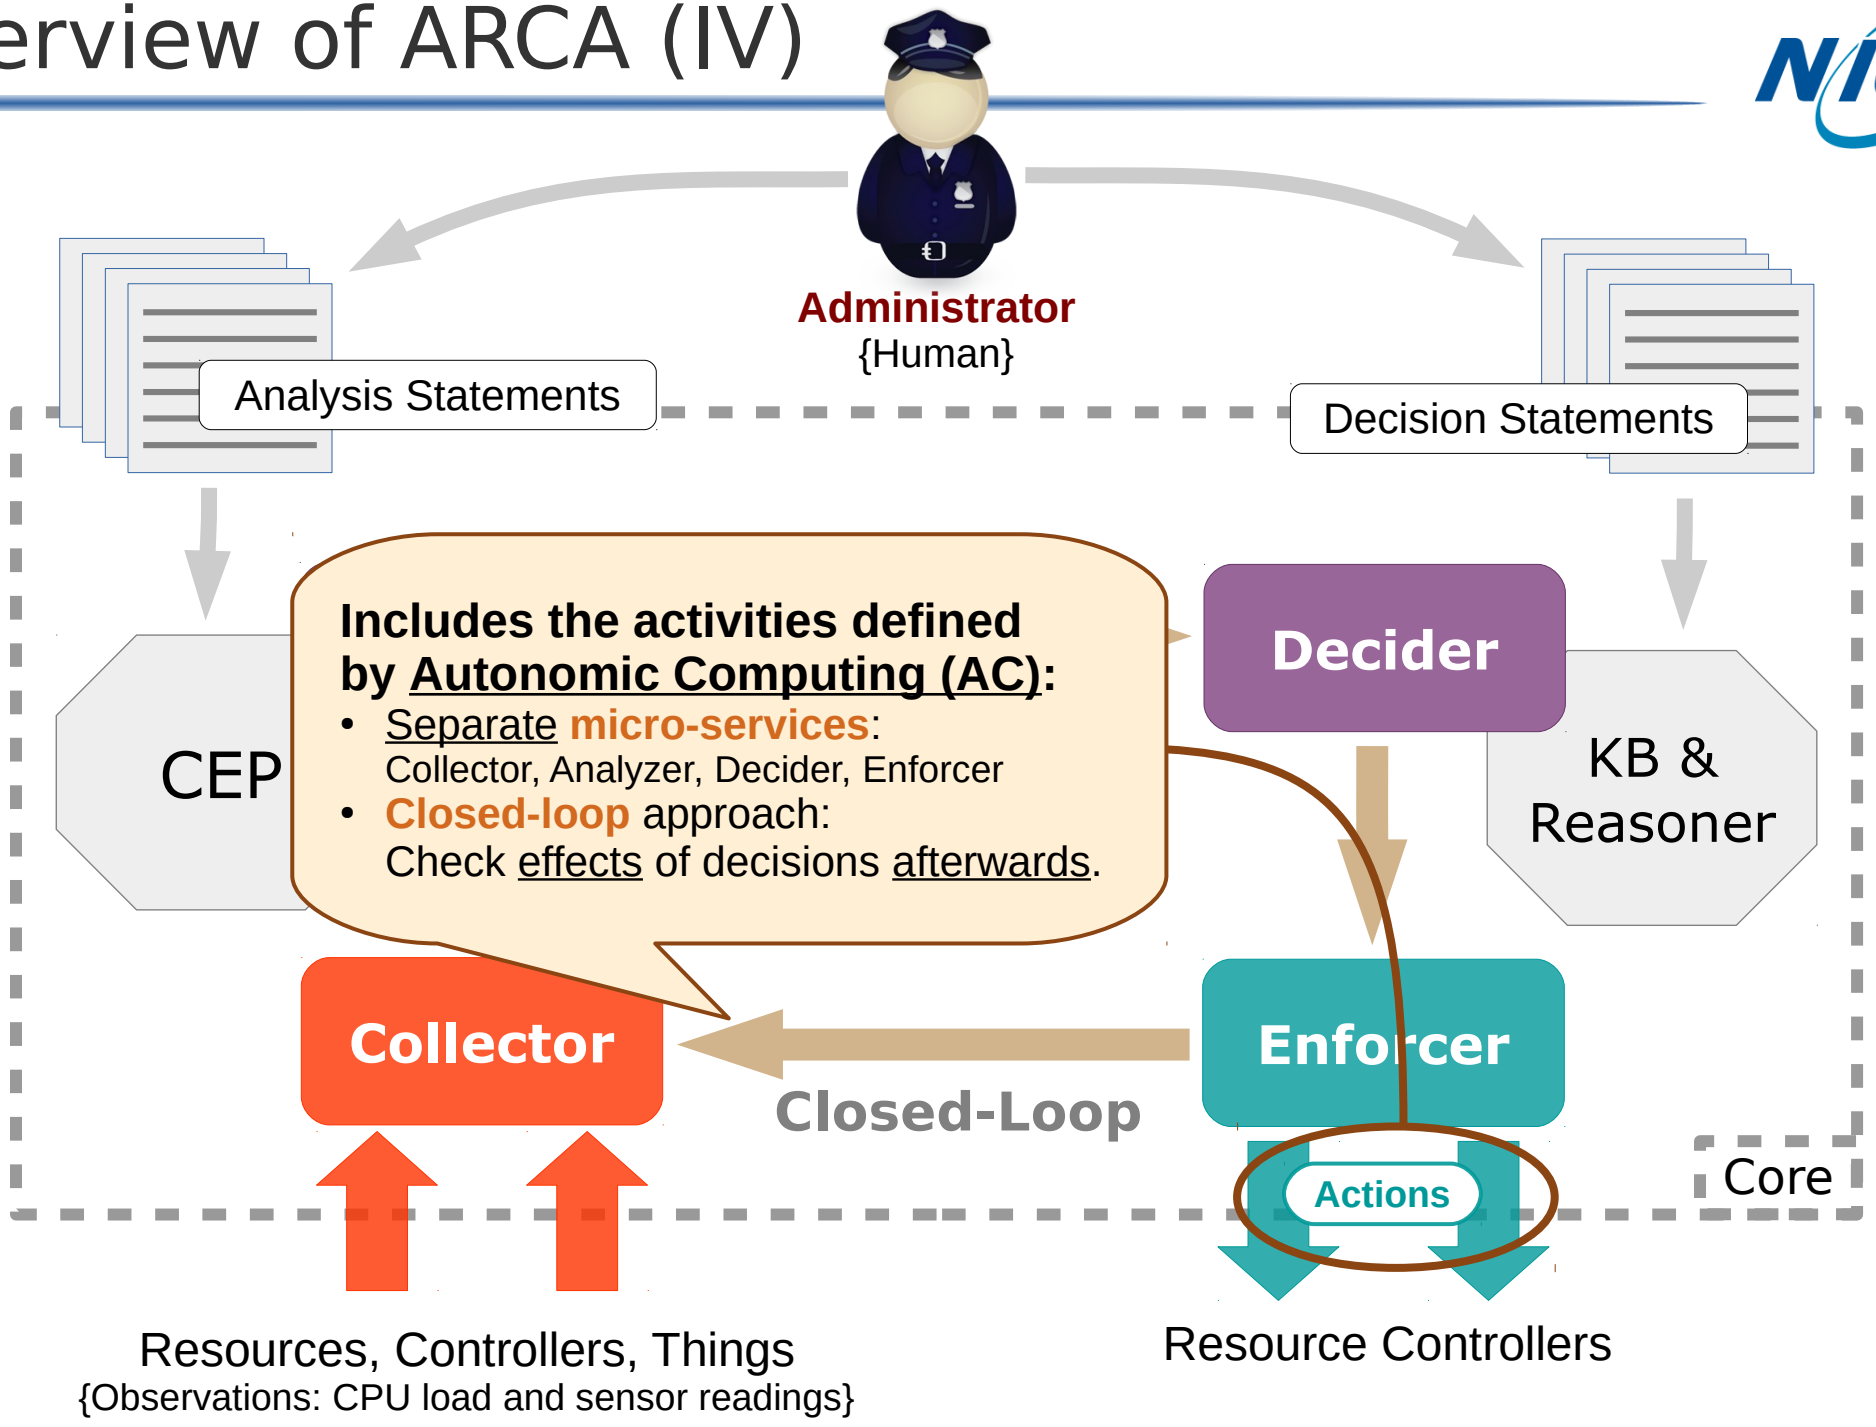

Overview of ARCA (V)

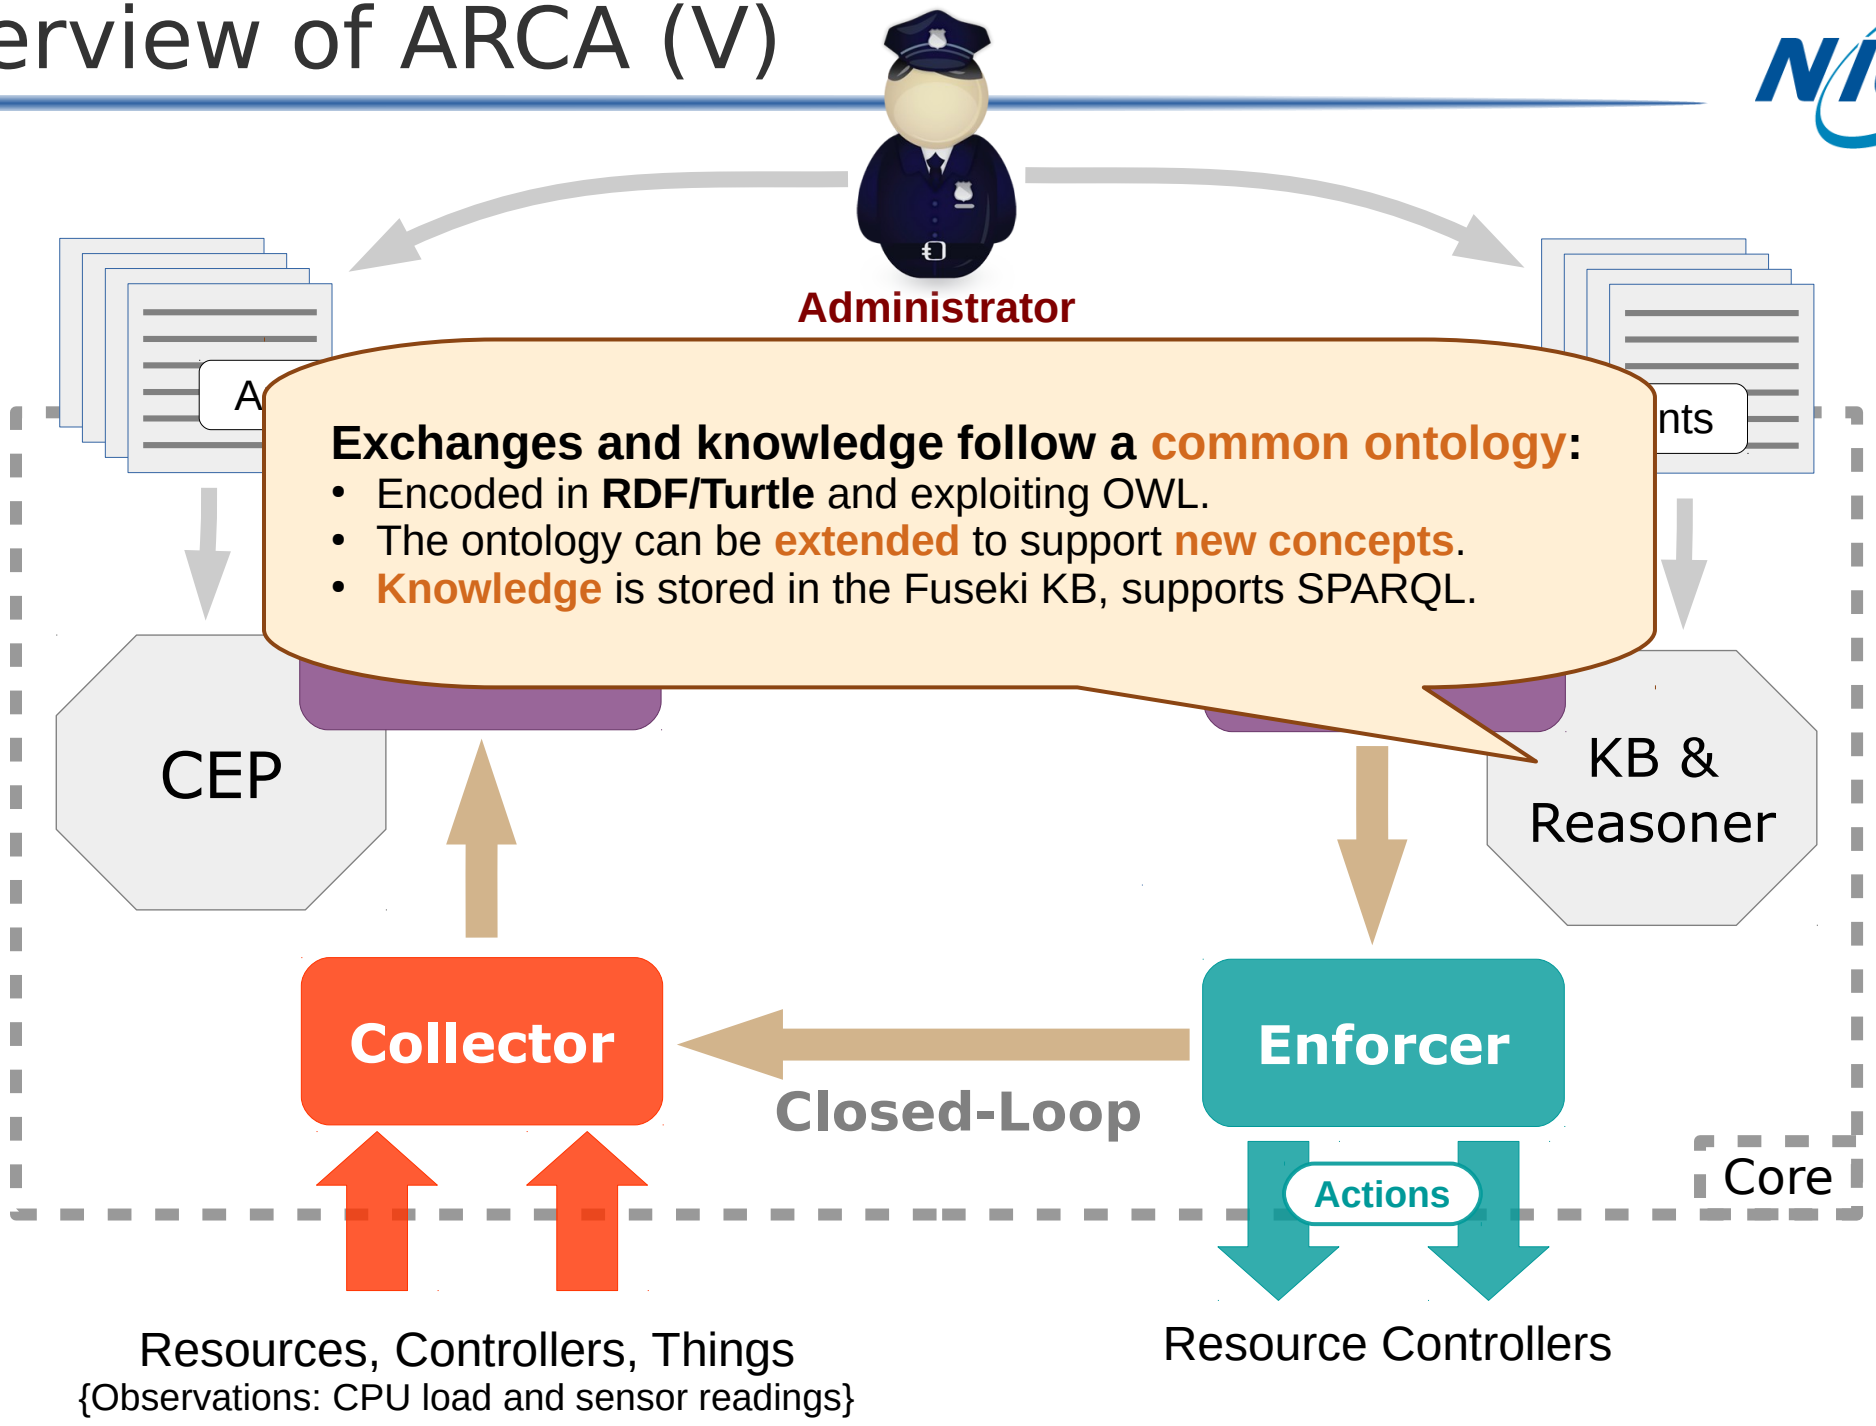

Resource Anticipation Strategy

- Functional and performance target:
  - **Anticipate** the amount of resources that a controlled system will require before it becomes effective.
- Involve **external event detectors**:
  - Physical: Things (IoT)
  - BigData
- Learn the event/reaction correlation:
  - Predict user behavior.
  - Correct mistaken predictions:
    - Improve and optimize learned model…
- Limit the memory used by the learning algorithm:
  - Keep only the most relevant vectors.
- **Fast adaptation** to big changes:
  - Discard old vectors when resizing.

Control Flow

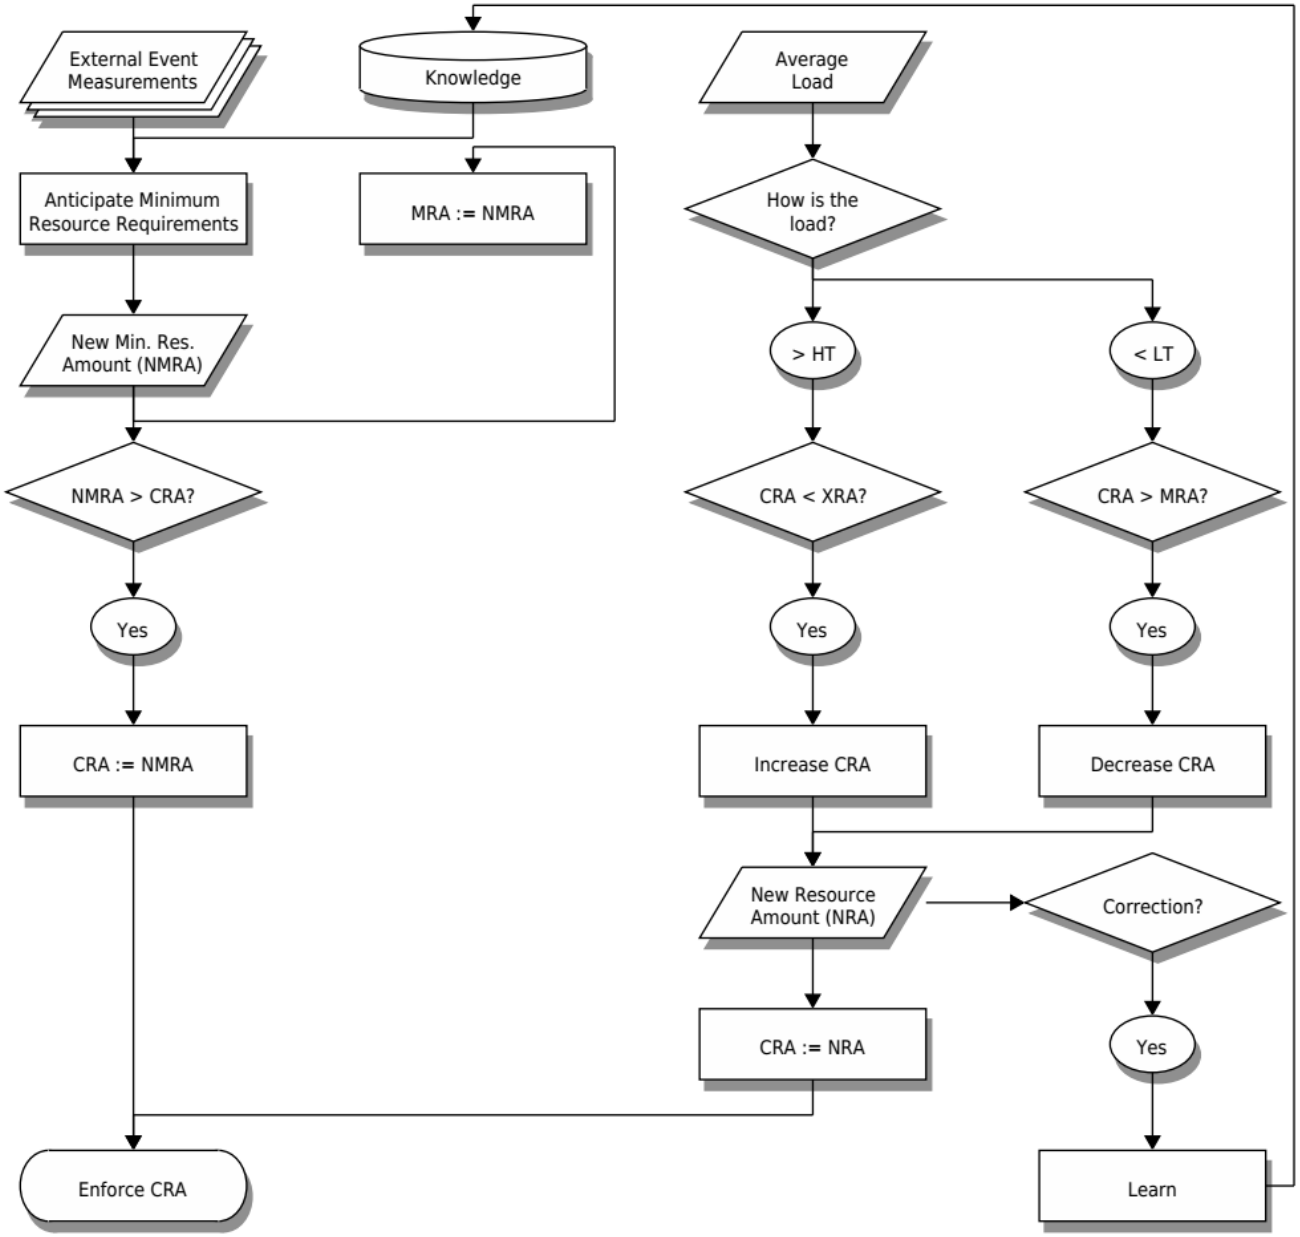

- Two key controlled parameters:
  - Current Resource Amount (**CRA**).
  - Minimum Resource Amount (**MRA**).
- Two concurrent sub-routines:
  - **Anticipation**.
  - Threshold checking and **correction**.
- **Self-assessed** learning process:
  - Correcting learned data when finding mistakes...

Algorithm (I)

1: **procedure** CONTROL(detectors, resources)
2: $mra \leftarrow MIN\_RESOURCES$
3: $cra \leftarrow mra$
4: $anticipator \leftarrow$ LEARNERCREATE$(MIN_RESOURCES, MAX_RESOURCES)$
5: $ant\_severity, ant\_time, ant\_peak, ant\_peak\_rel\_time \leftarrow 0, 0, 0, 0$
6: $severity, pseverity, load, drate \leftarrow 0, 0, 1, 0$
7: $attl\_model \leftarrow$ LEARNERCREATE(MIN_ATTL, MAX_ATTL)
8: **while** TRUE **do**
9: $severity \leftarrow$ COLLATESENSORREADINGS(COLLECT$(detectors, SEVERITY)$)
10: $load \leftarrow$ CALCULATEAVGLOAD(COLLECT$(resources, LOAD)$)
11: $drate \leftarrow$ CALCULATEAVGDROPRATE(COLLECT$(resources, DROP_RATE)$)
12: **if** $ant\_severity \neq 0$ **then**
13: $demand \leftarrow cra * \frac{load + drate}{SERVER\_WORK\_QCAP}$
14: **if** $demand > ant\_peak$ **then**
15: $ant\_peak \leftarrow demand$
16: $ant\_peak\_rel\_time \leftarrow (NOW - ant\_time) * 1.25$
17: **end if**
18: **if** $NOW - ant\_time >$ LEARNERGET$(attl_model, ant_severity)$ **and** $load < LT$ **then**
19: LEARNERSET$(anticipator, ant_severity, \frac{ant_peak}{HT})$
20: LEARNERSET$(attl_model, ant_severity, ant_peak_rel_time)$
21: $ant\_severity \leftarrow 0$
22: $ant\_time \leftarrow 0$
23: $mra \leftarrow MIN\_RESOURCES$
24: **end if**
25: **end if**

Algorithm (II)

26: **if** $severity \neq 0$ **and** $severity \neq pseverity$ **then**
27: $ant\_severity \leftarrow severity$
28: $ant\_time \leftarrow NOW$
29: $ant\_peak \leftarrow 0$
30: $mra \leftarrow \text{LEARNERGET}(anticipator, severity)$
31: **end if**
32: $pseverity \leftarrow severity$
33: $nra \leftarrow cra$
34: **if** $load > HT$ **or** $cra < mra$ **then**
35: $nra \leftarrow \text{MIN}(\text{MAX}(cra + INC, \frac{load * cra}{HT}, mra), MAX\_RESOURCES)$
36: **end if**
37: **if** $load < LT$ **then**
38: $nra \leftarrow \text{MAX}(nra - DEC, mra, MIN\_RESOURCES)$
39: **end if**
40: **if** $nra \neq cra$ **and** $\text{NOSIDEEFFECT}(nra)$ **then**
41: $cra \leftarrow nra$
42: $\text{ENFORCE}(cra)$
43: **end if**
44: **end while**
45: **end procedure**

Alignment With ETSI-NFV-MANO (I)

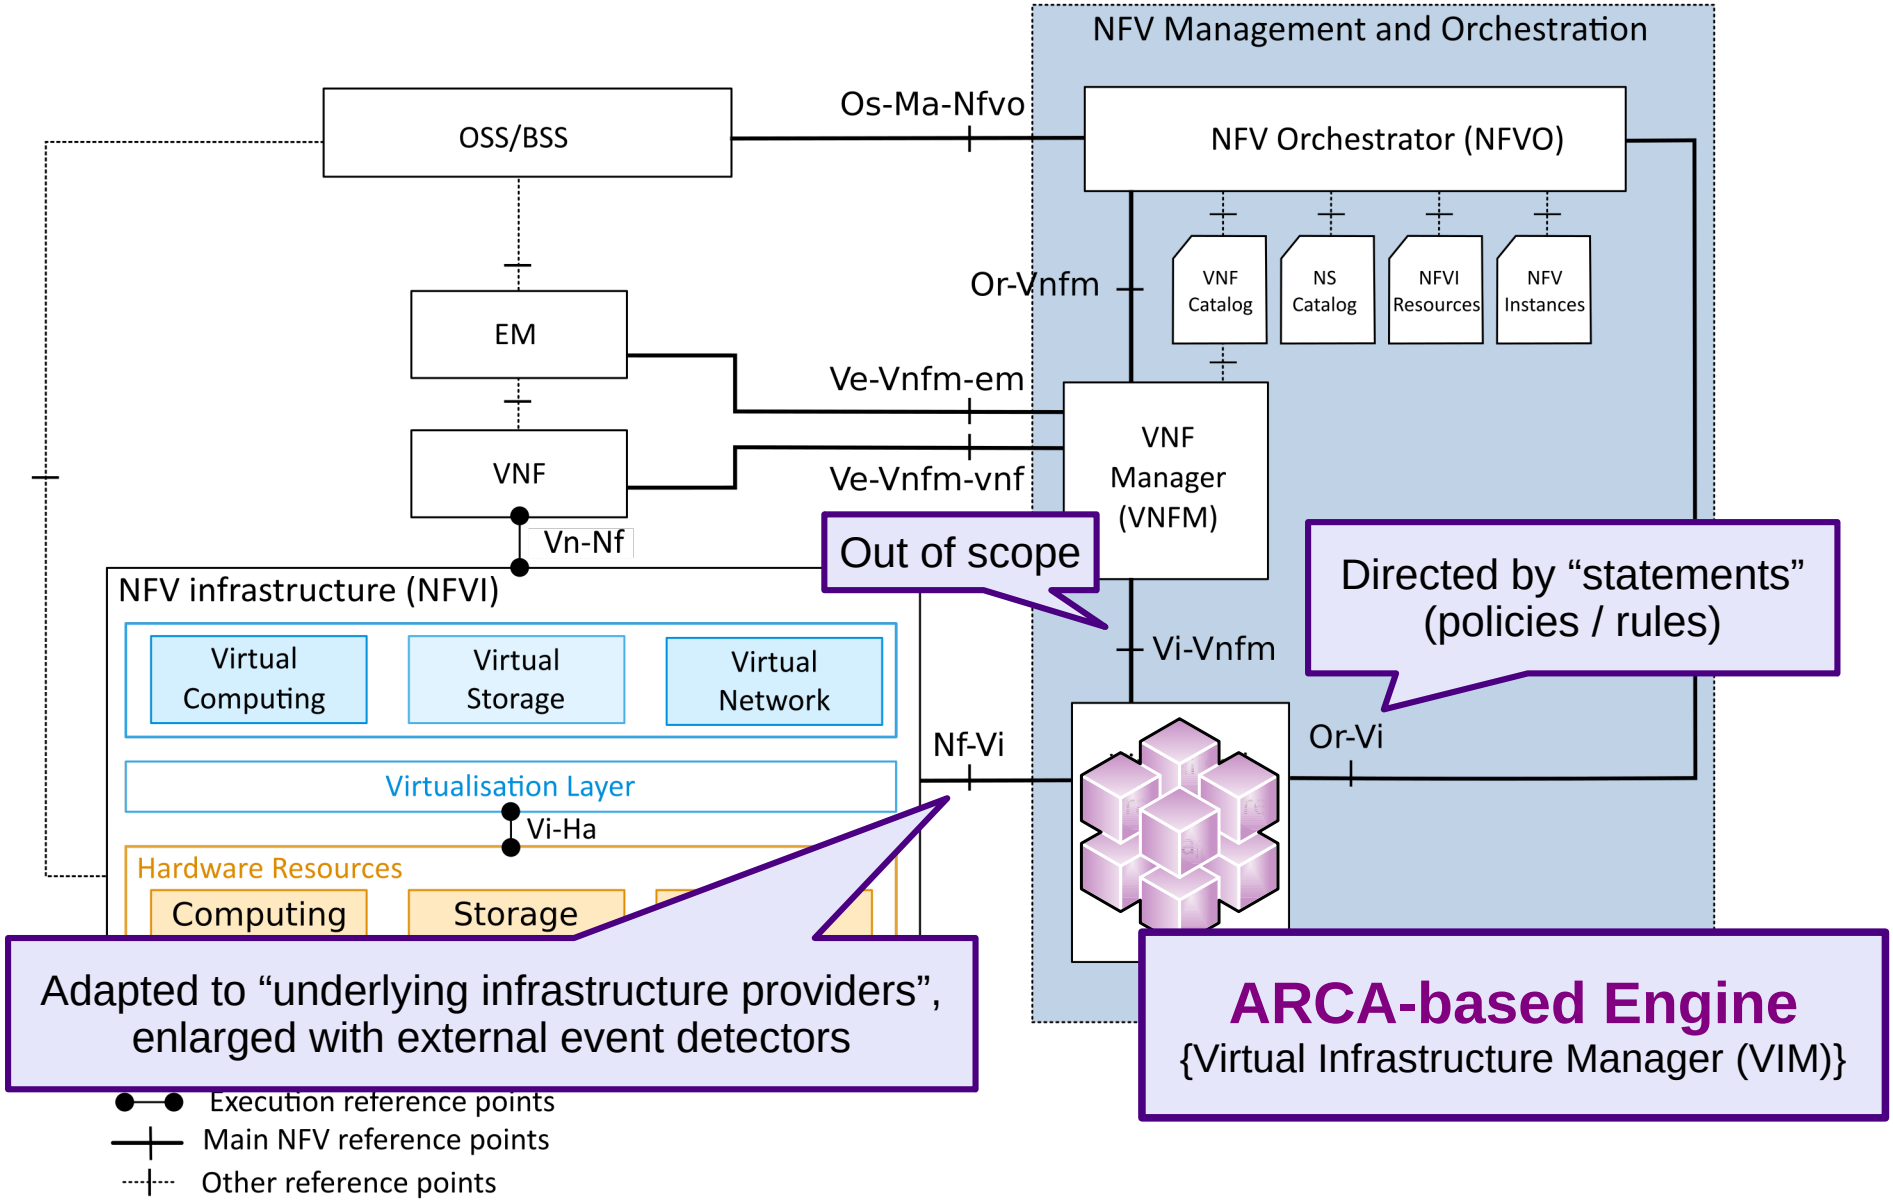

Alignment With ETSI-NFV-MANO (II)

ARCA is fitted as the Virtual Infrastructure Manager (**VIM**):

- Discharges responsibilities from **VNFM** and **NFVO**.
- Improves the **scalability** and **resiliency**...
  ...in case of disconnection from the orchestrator.
- Meets requirements of Virtual Network Operators (**VNO**s).

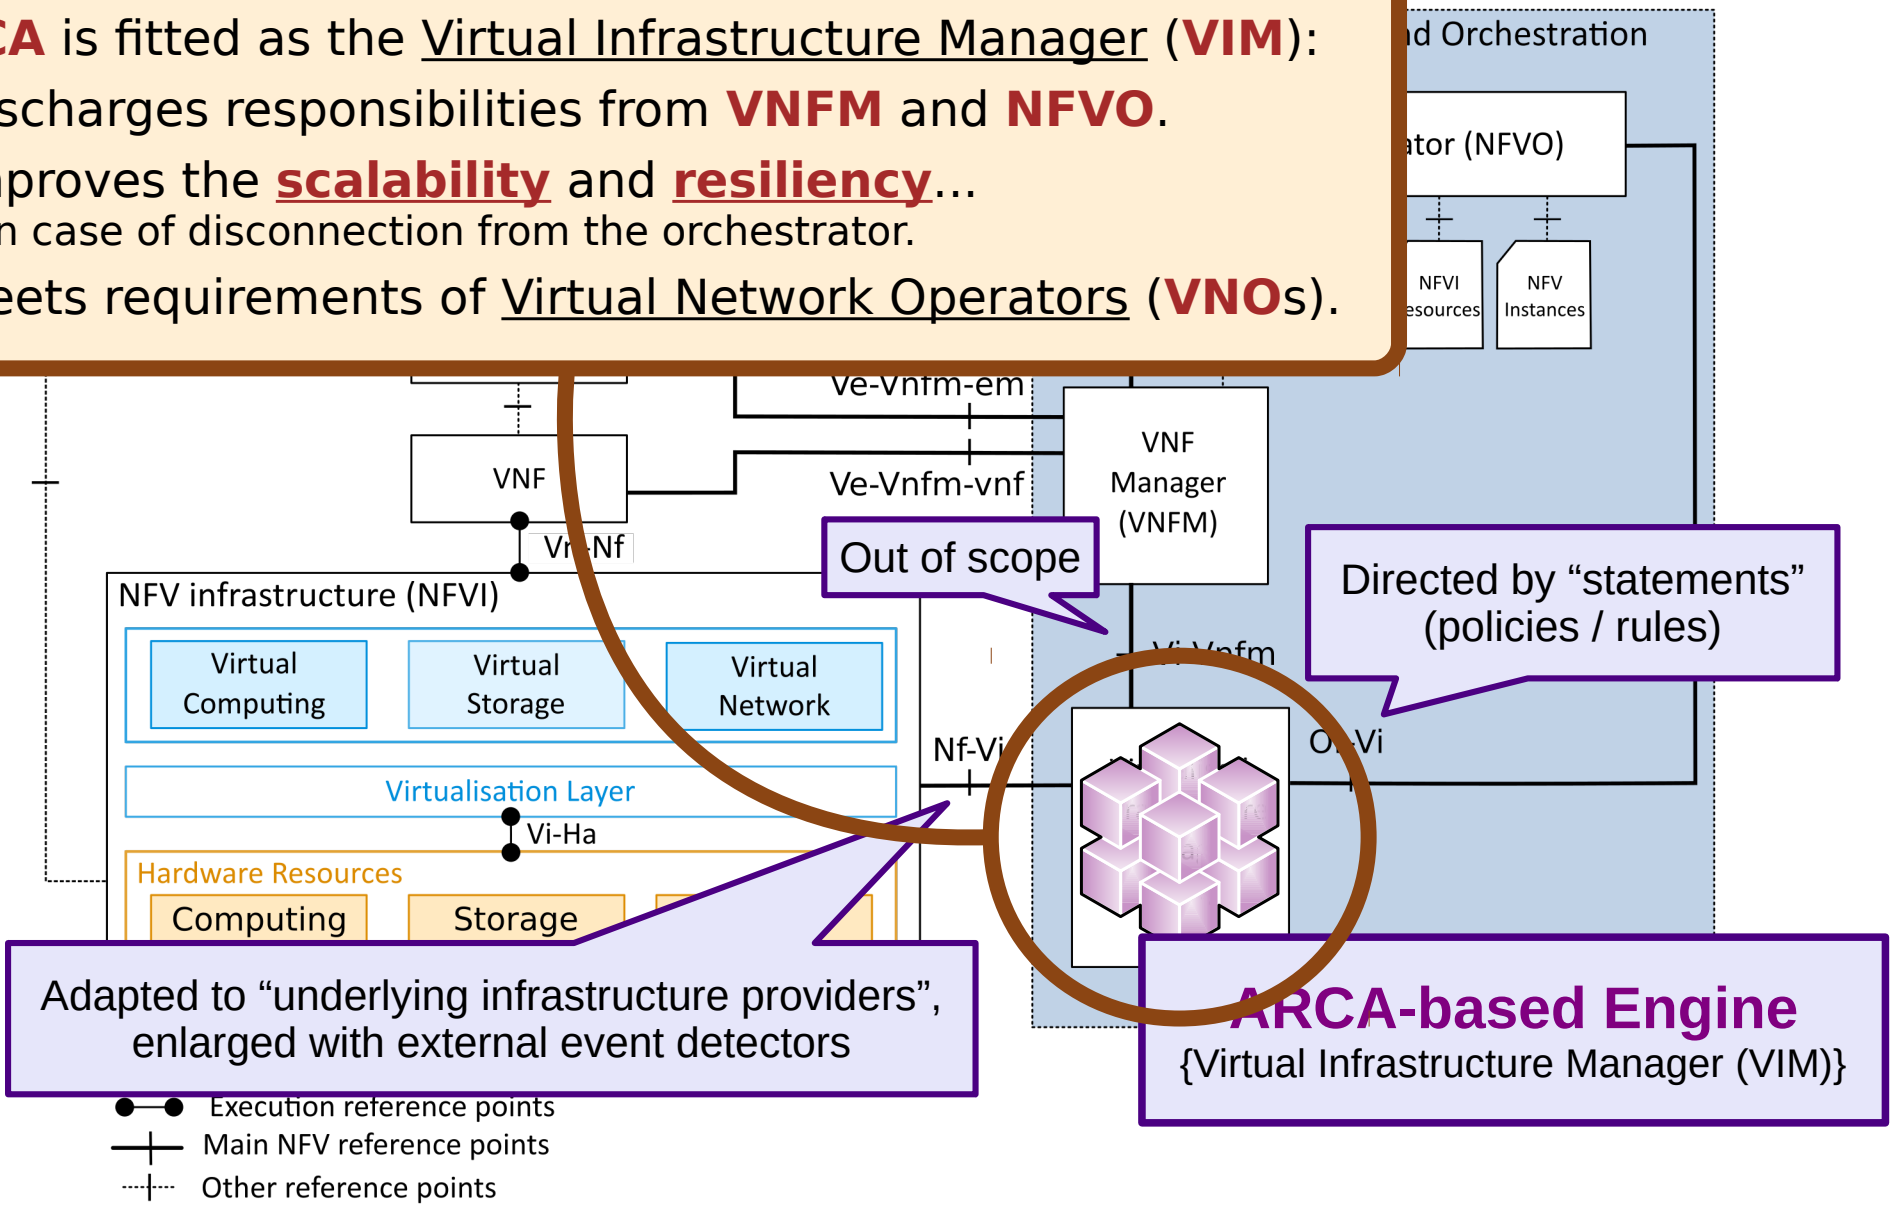

Alignment With ETSI-NFV-MANO (III)

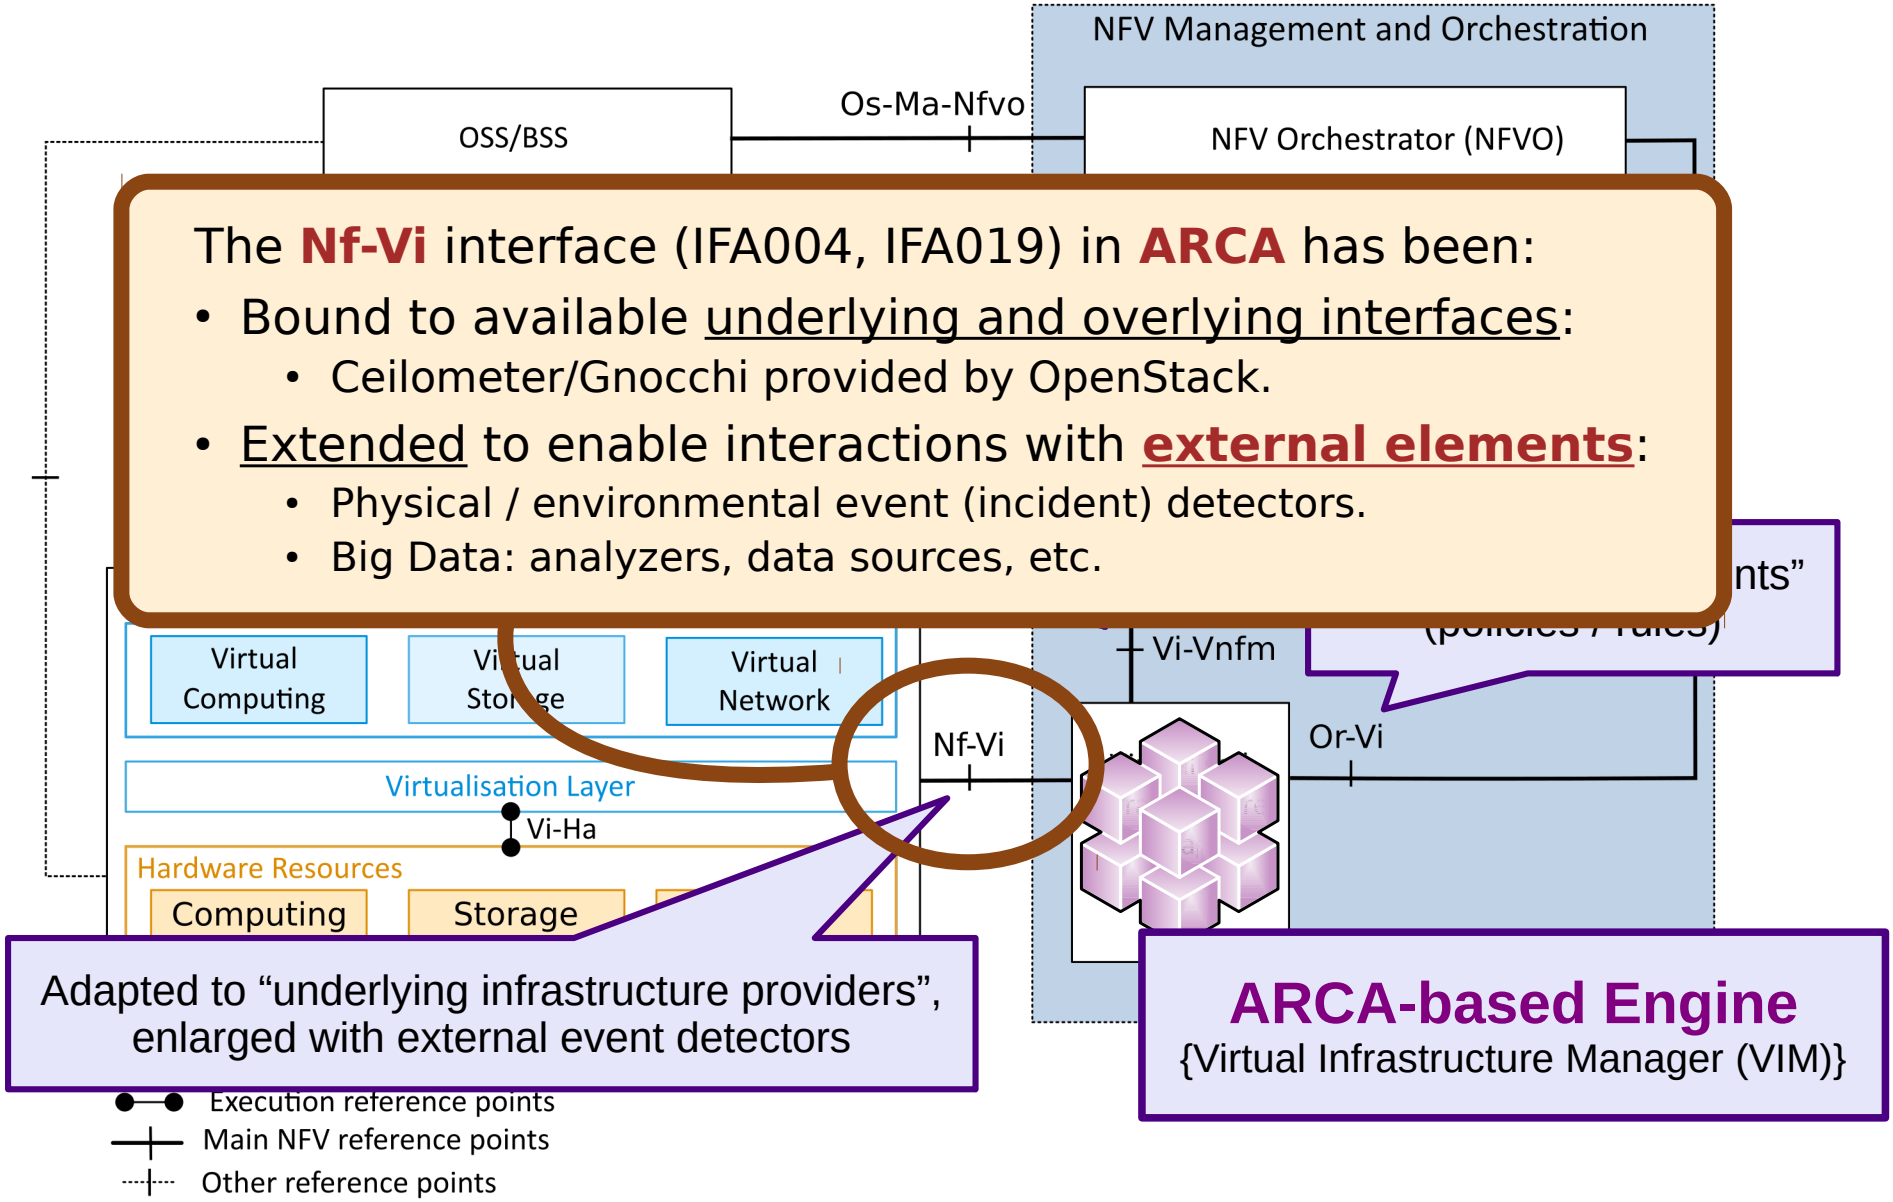

Alignment With ETSI-NFV-MANO (IV)

The **Or-Vi** interface (IFA005) is provided by:

- The specification of **control/mngmt targets** (statements):
  - Represent the rules and policies that ARCA must enforce.
  - Provided by system administrators and/or external orchestrators.
- **ARCA** will enforce the statements in response to changes in the environment and/or user requirements:
  - Requirements are communicated with additional statements.

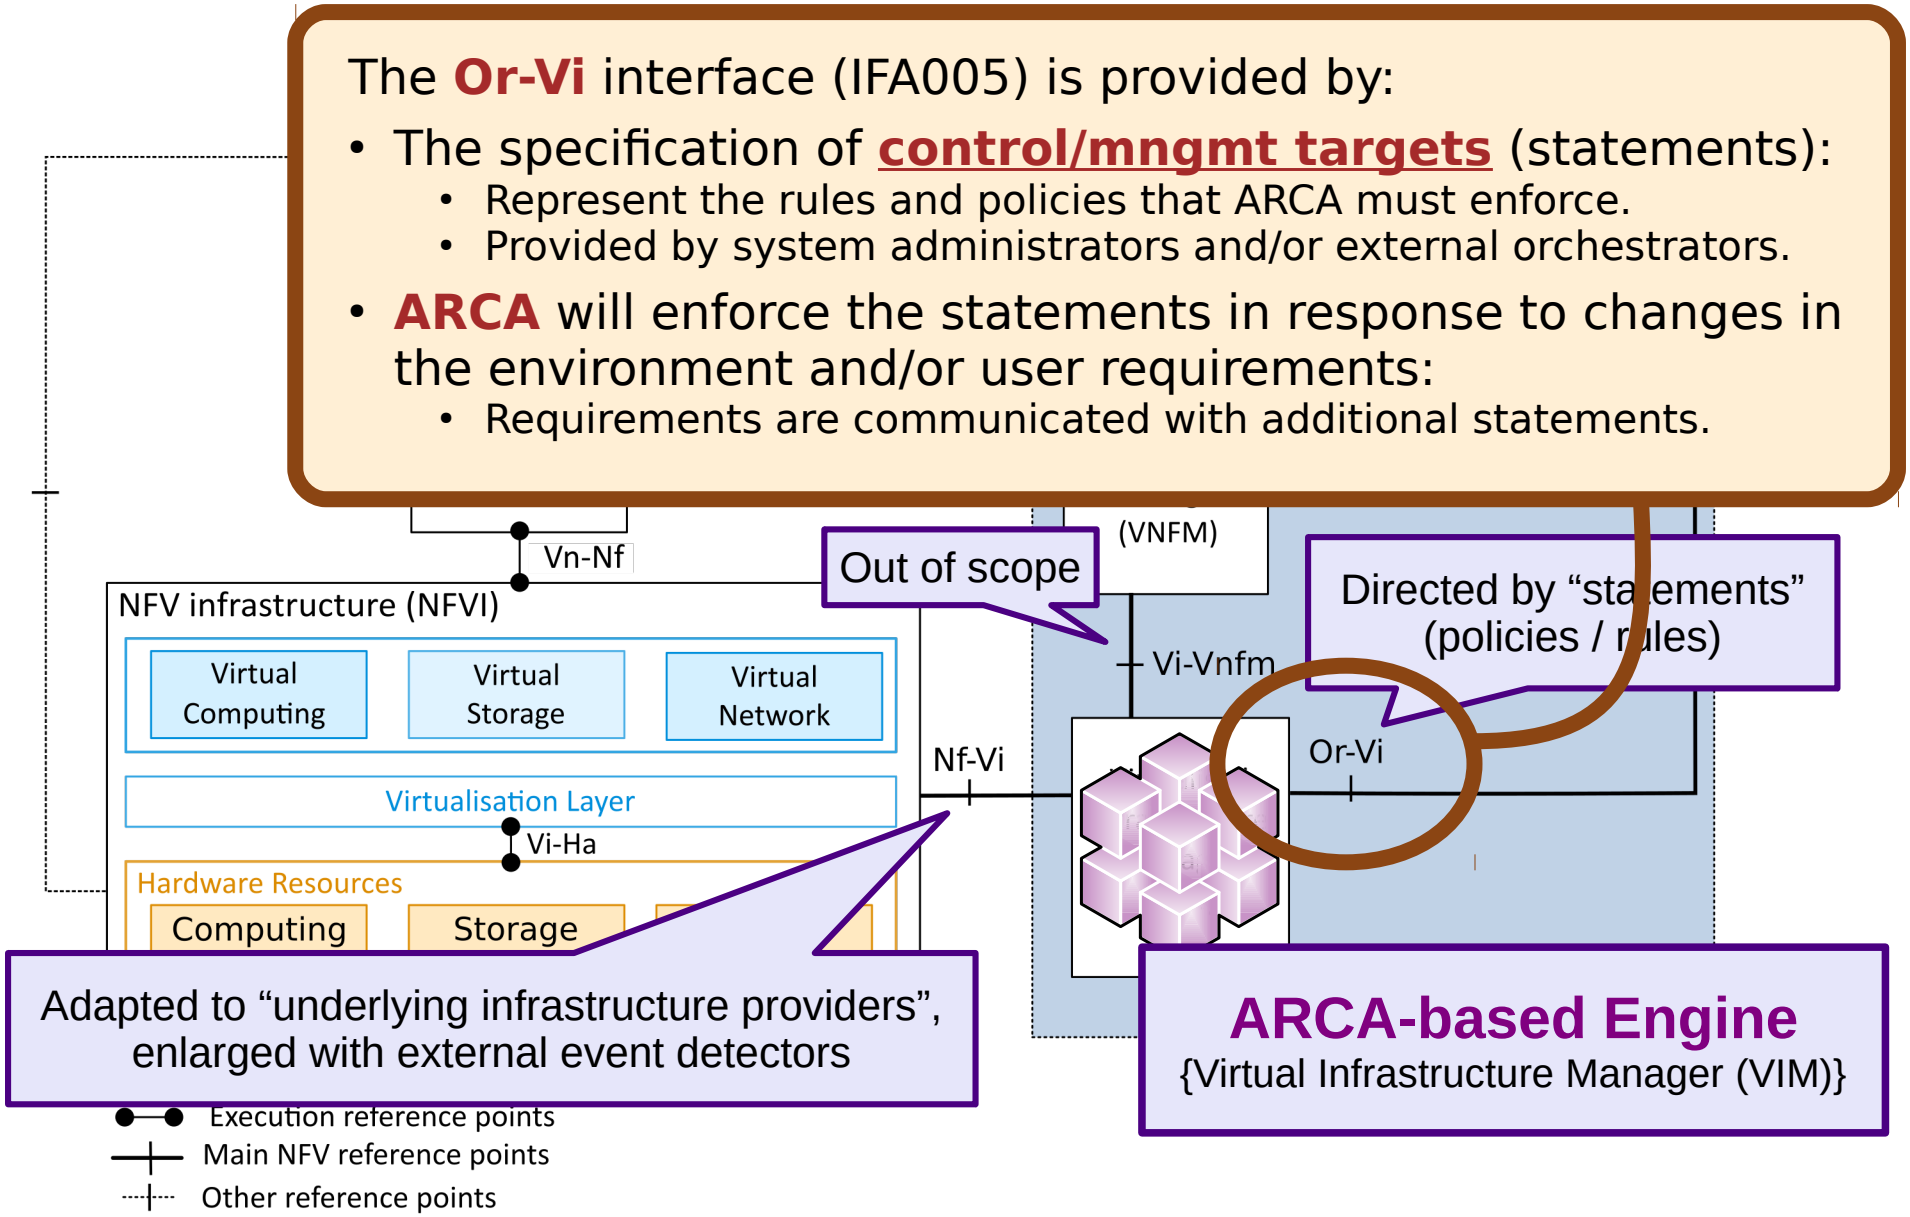

Alignment With ETSI-NFV-MANO (V)

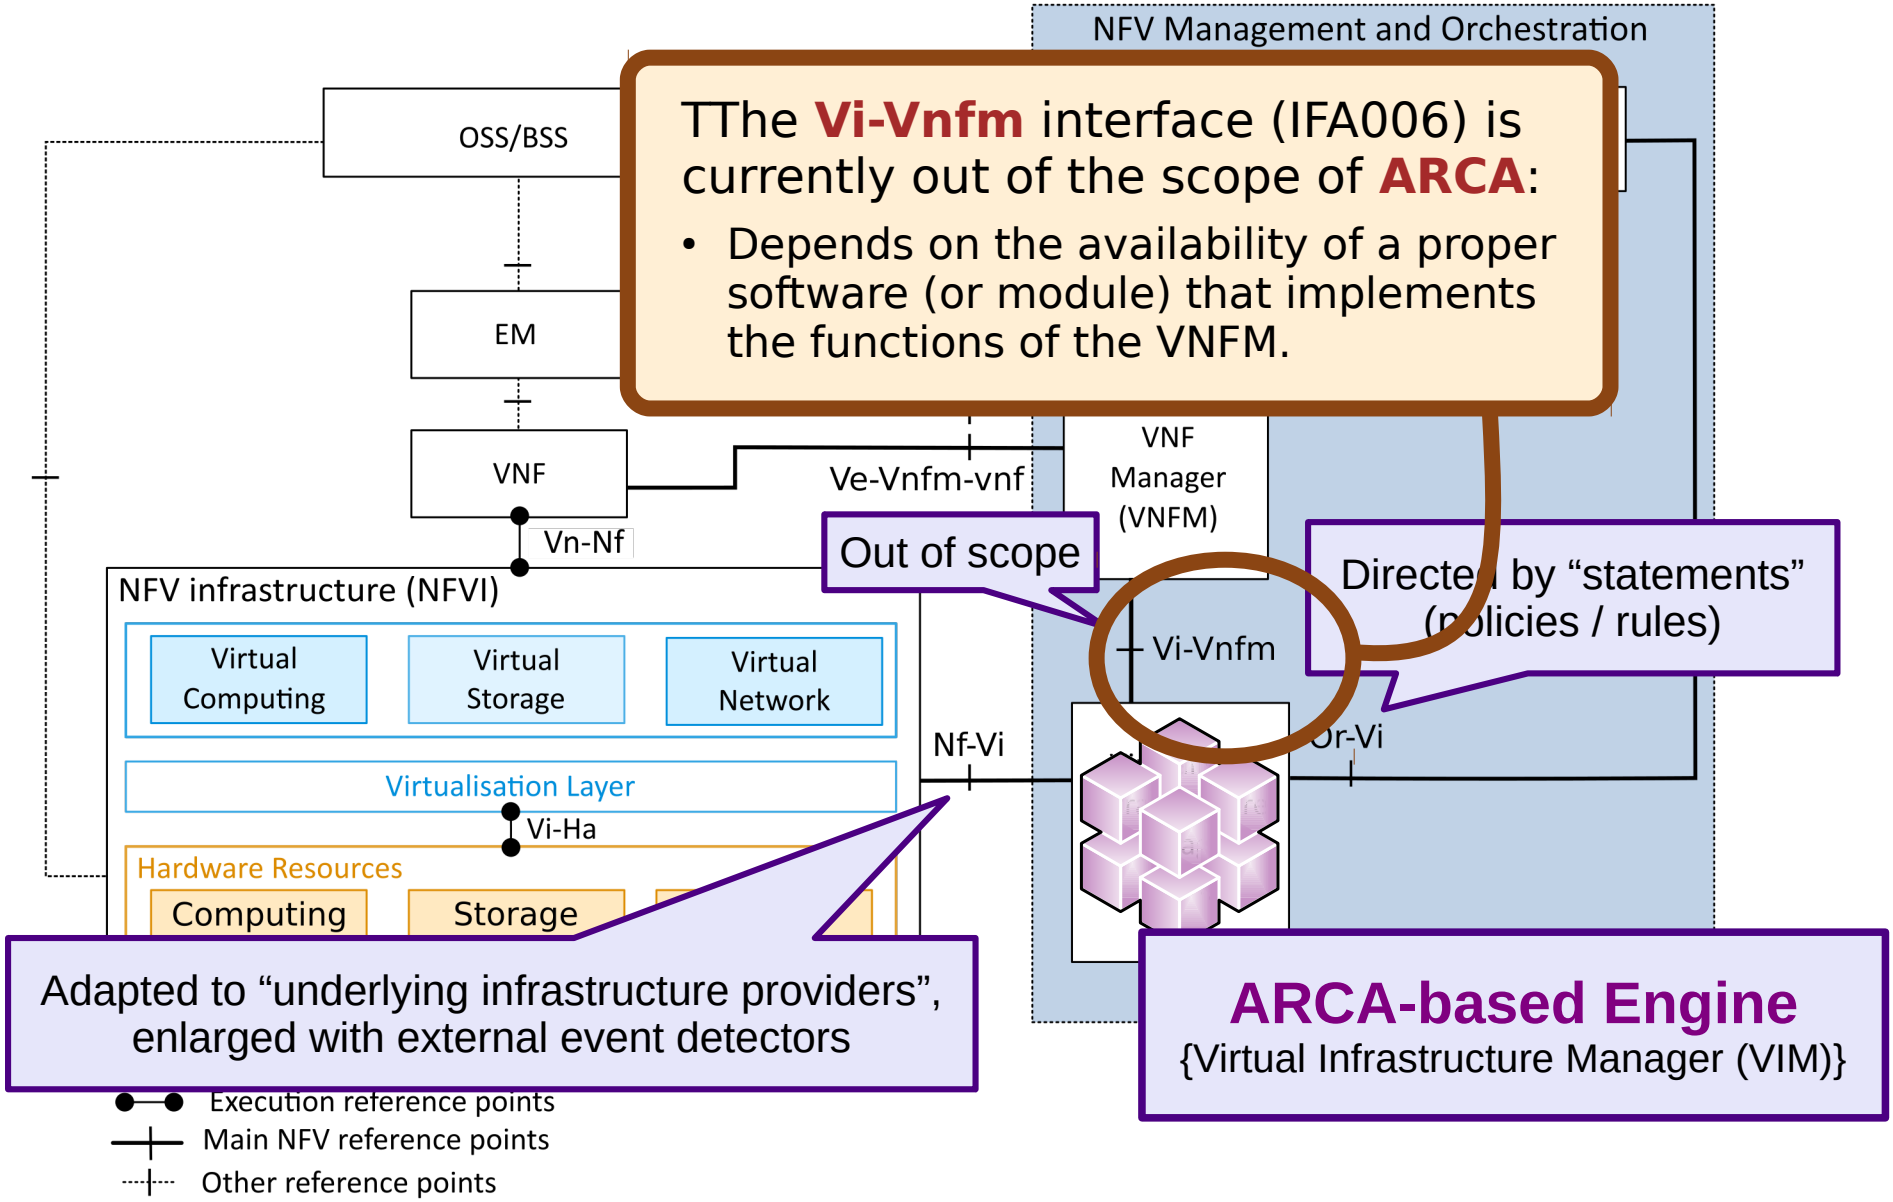

Conclusions & Future Work

- Designed **ARCA**:
  - To provide functions of the Virtual Infrastructure Manager (**VIM**) of **NFV-MANO**.
  - Extended **VIM interfaces** to meet requirements of the **real world**:
    - Sport events, TV shows, emergency scenarios...
  - Achieved good perfomance within an OpenStack-based deployment:
    - Detailed **overlying** and **underlying** infrastructures.
    - Reproduction of production-like environments to ensure transferable research results.
- SDN/NFV and OpenStack stakeholders will benefit from ARCA features:
  - **Efficient** use of **resources**:
    - Further reduce CAPEX and OPEX:
  - Benefit to both infrastructure providers and consumers.
- Next steps:
  - Keep **reducing** ARCA **response time**.
  - Increase complexity of the validation scenario:
    - Mix clients and servants in the same domains.
  - Align **ARCA**-based VNC to **additional equirements** from **NFV**/**SDN**.

Thanks for Your Attention

Q & A

– EOF –