Specification of Cloud-Based Elastic Systems

Srđan Krstić
Politecnico di Milano

POLITECNICO
MILANO

Joint work with Carlo Ghezzi, Domenico Bianculli, Marcello Bersani, Alessio Gambi and Schahram Dustdar

Specification of Cloud-Based Elastic Systems

Srđan Krstić
Politecnico di Milano

Joint work with Carlo Ghezzi, Domenico Bianculli, Marcello Bersani, Alessio Gambi and Schahram Dustdar

Cloud-based Services

- quick & painless access
- provision on-demand
- high availability
- on-the-fly changes
- pay-per-use billing

Cloud-Based Elastic System

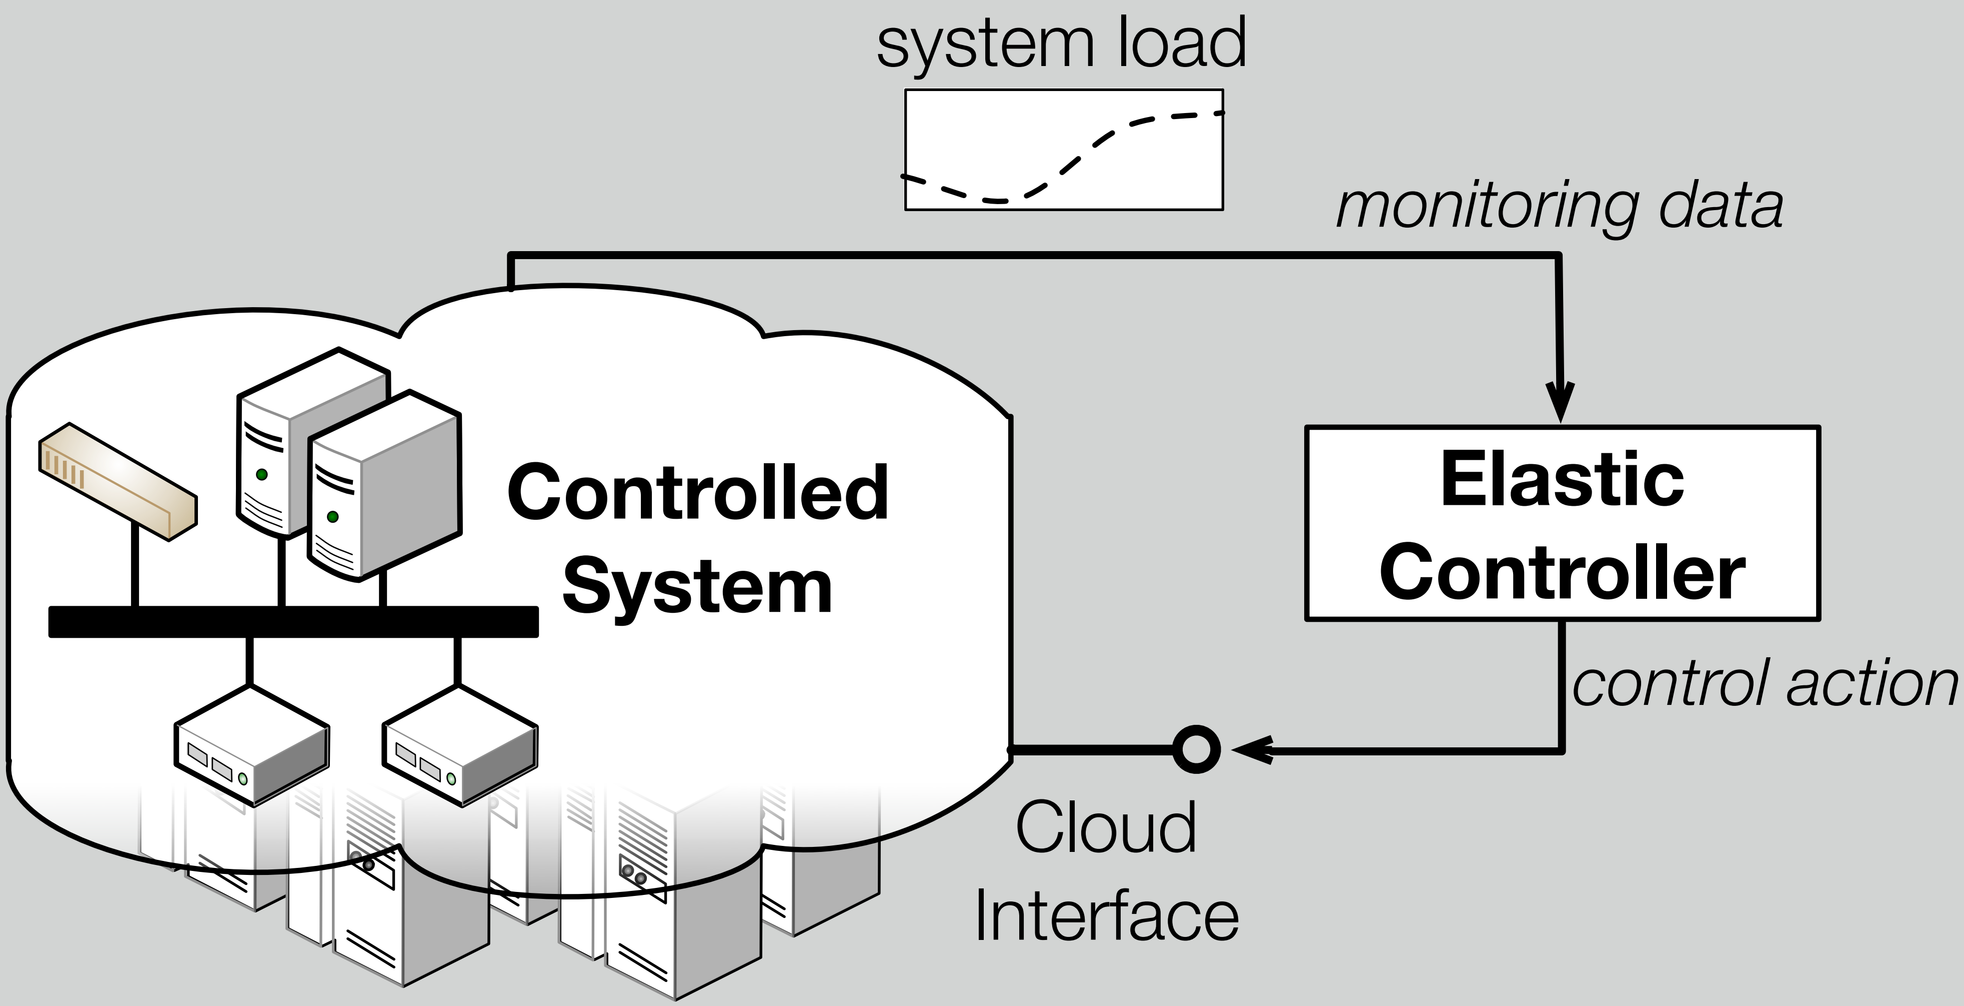

Open issues

Specification

Verification

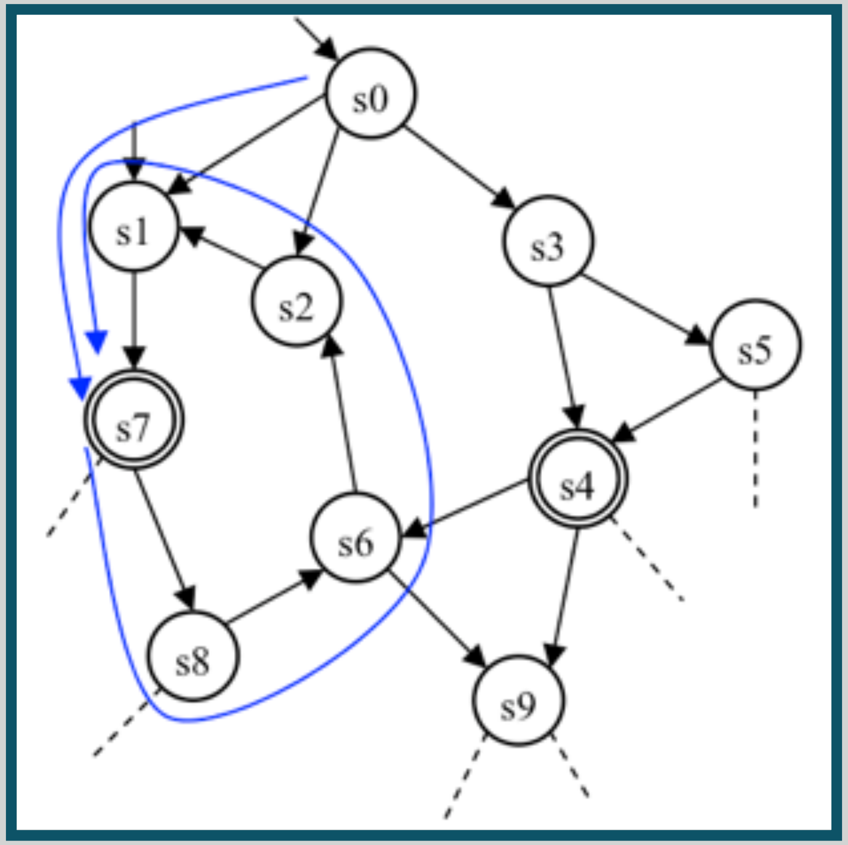

Failure Analysis

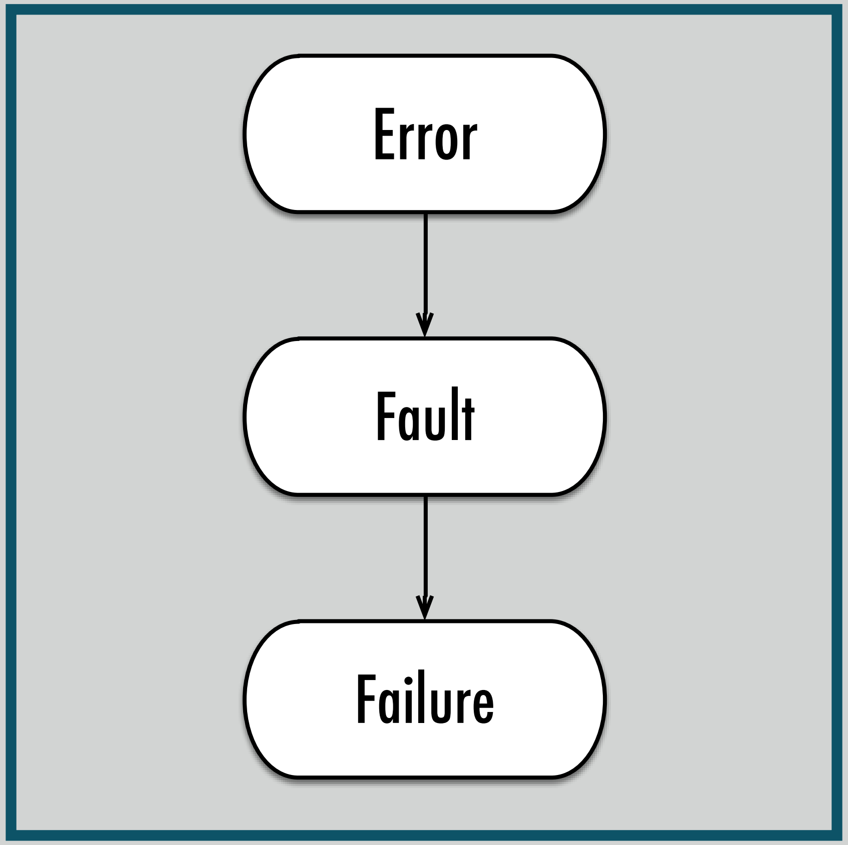

Open issues

Specification

Verification

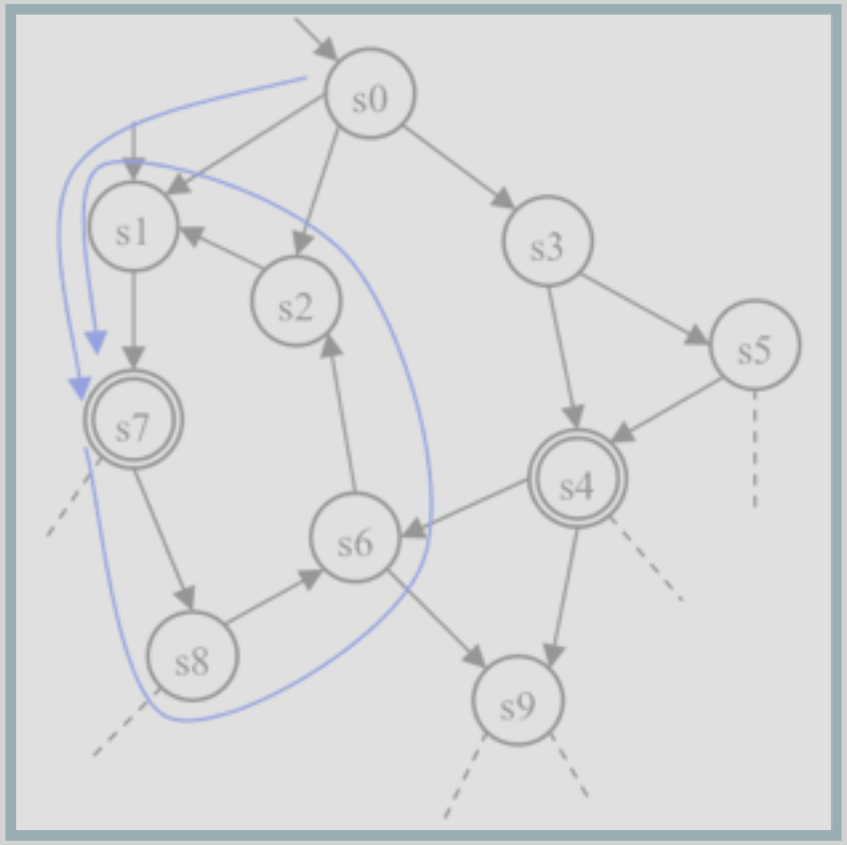

Failure Analysis

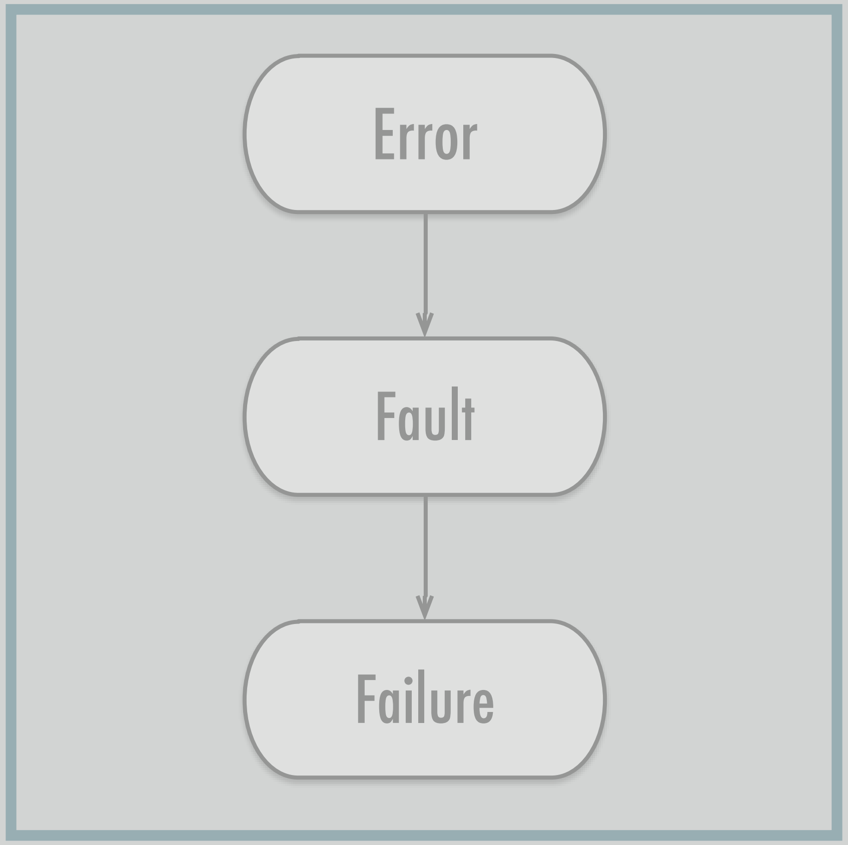

Cloud-Based Elastic System

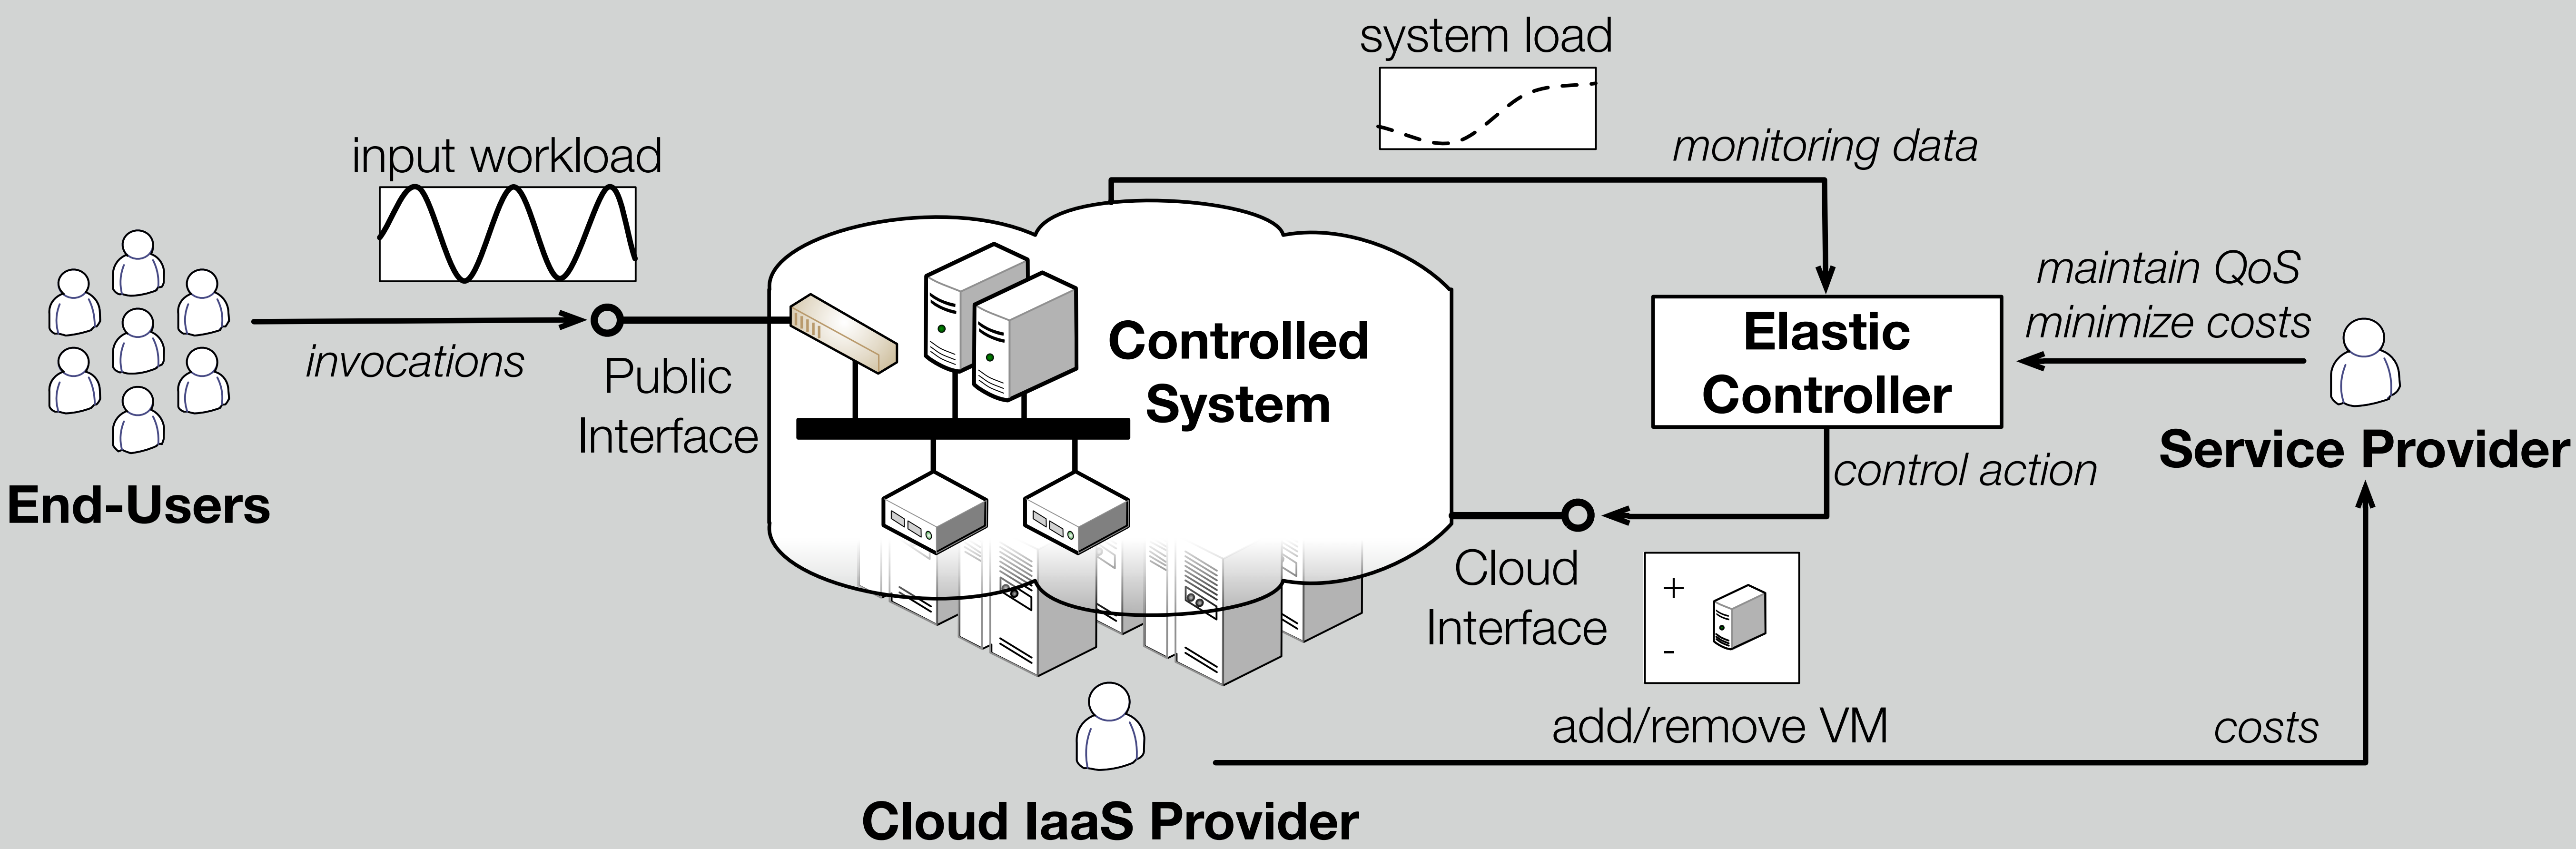

Cloud-Based Elastic System

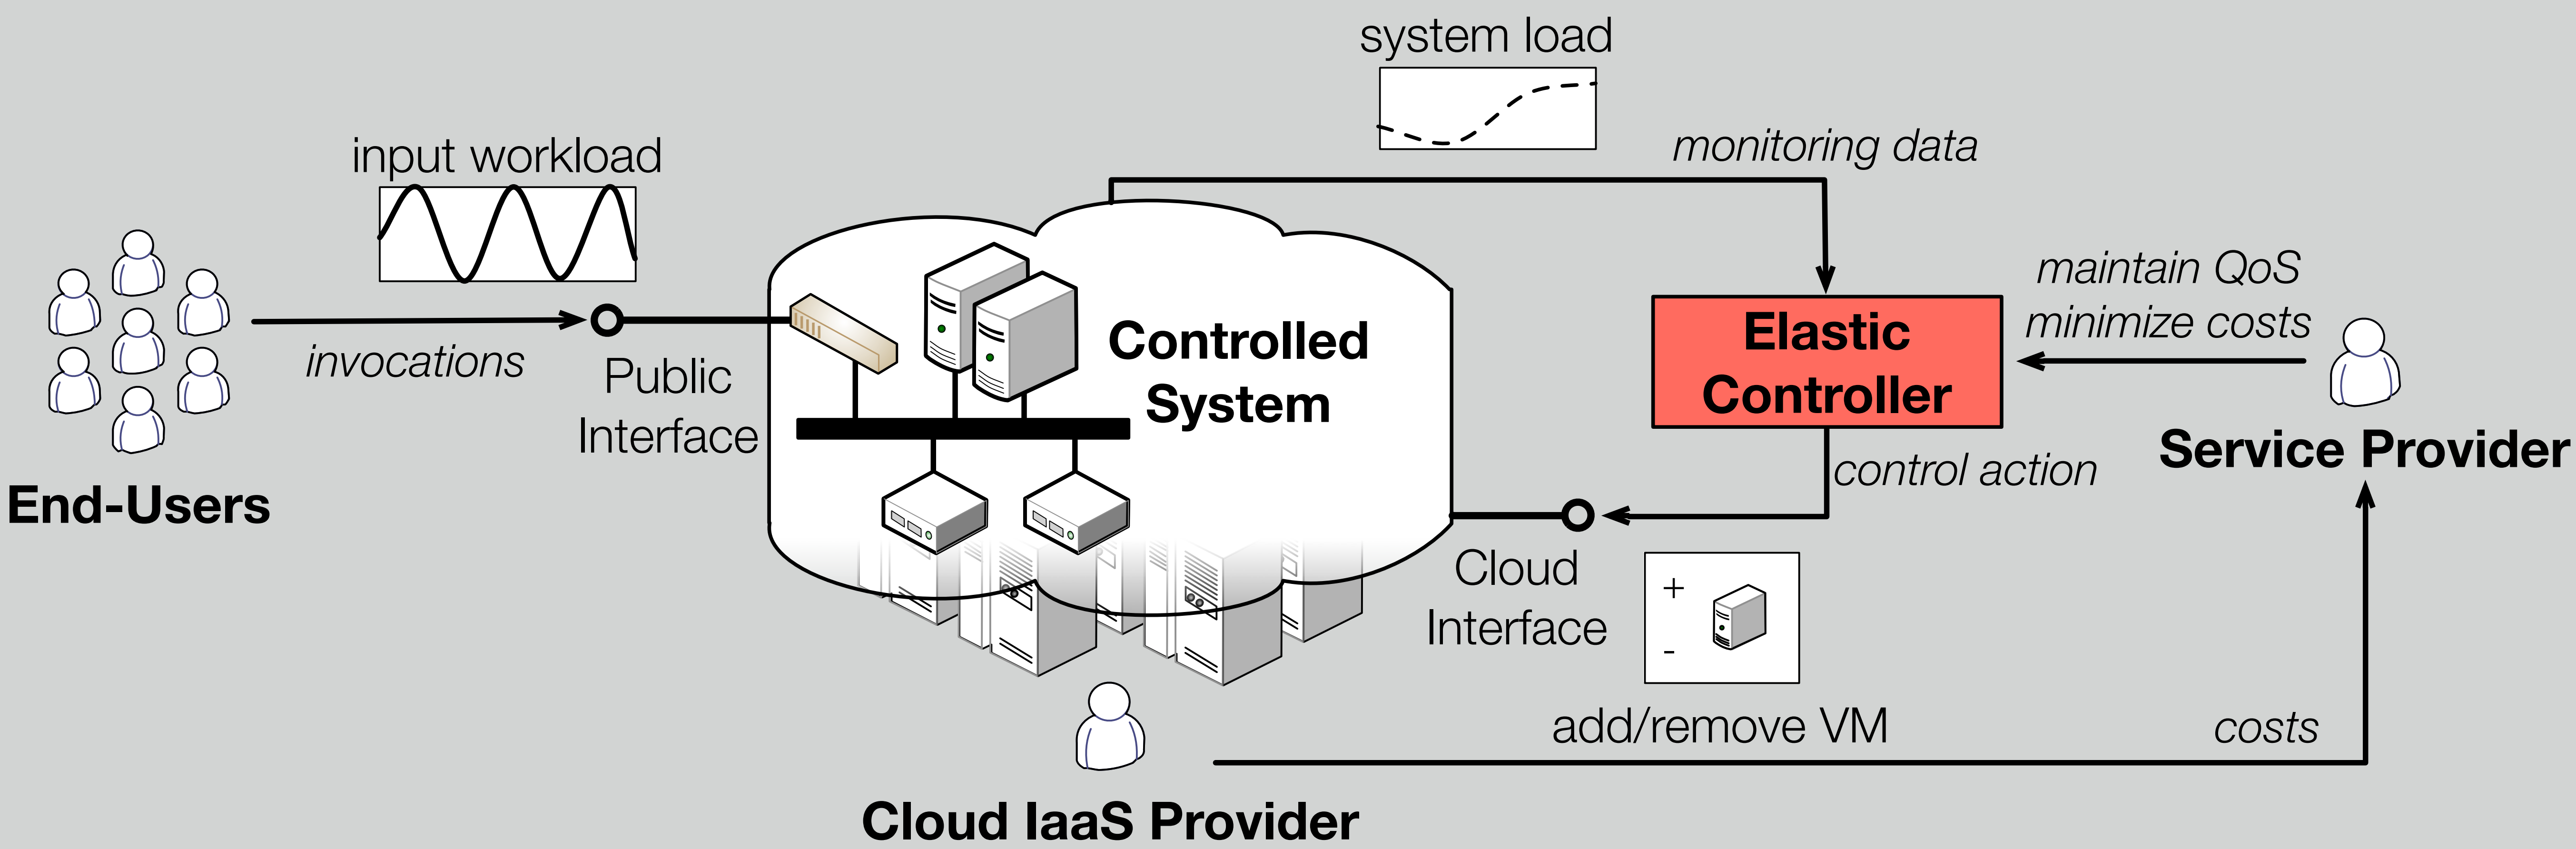

Cloud-Based Elastic System

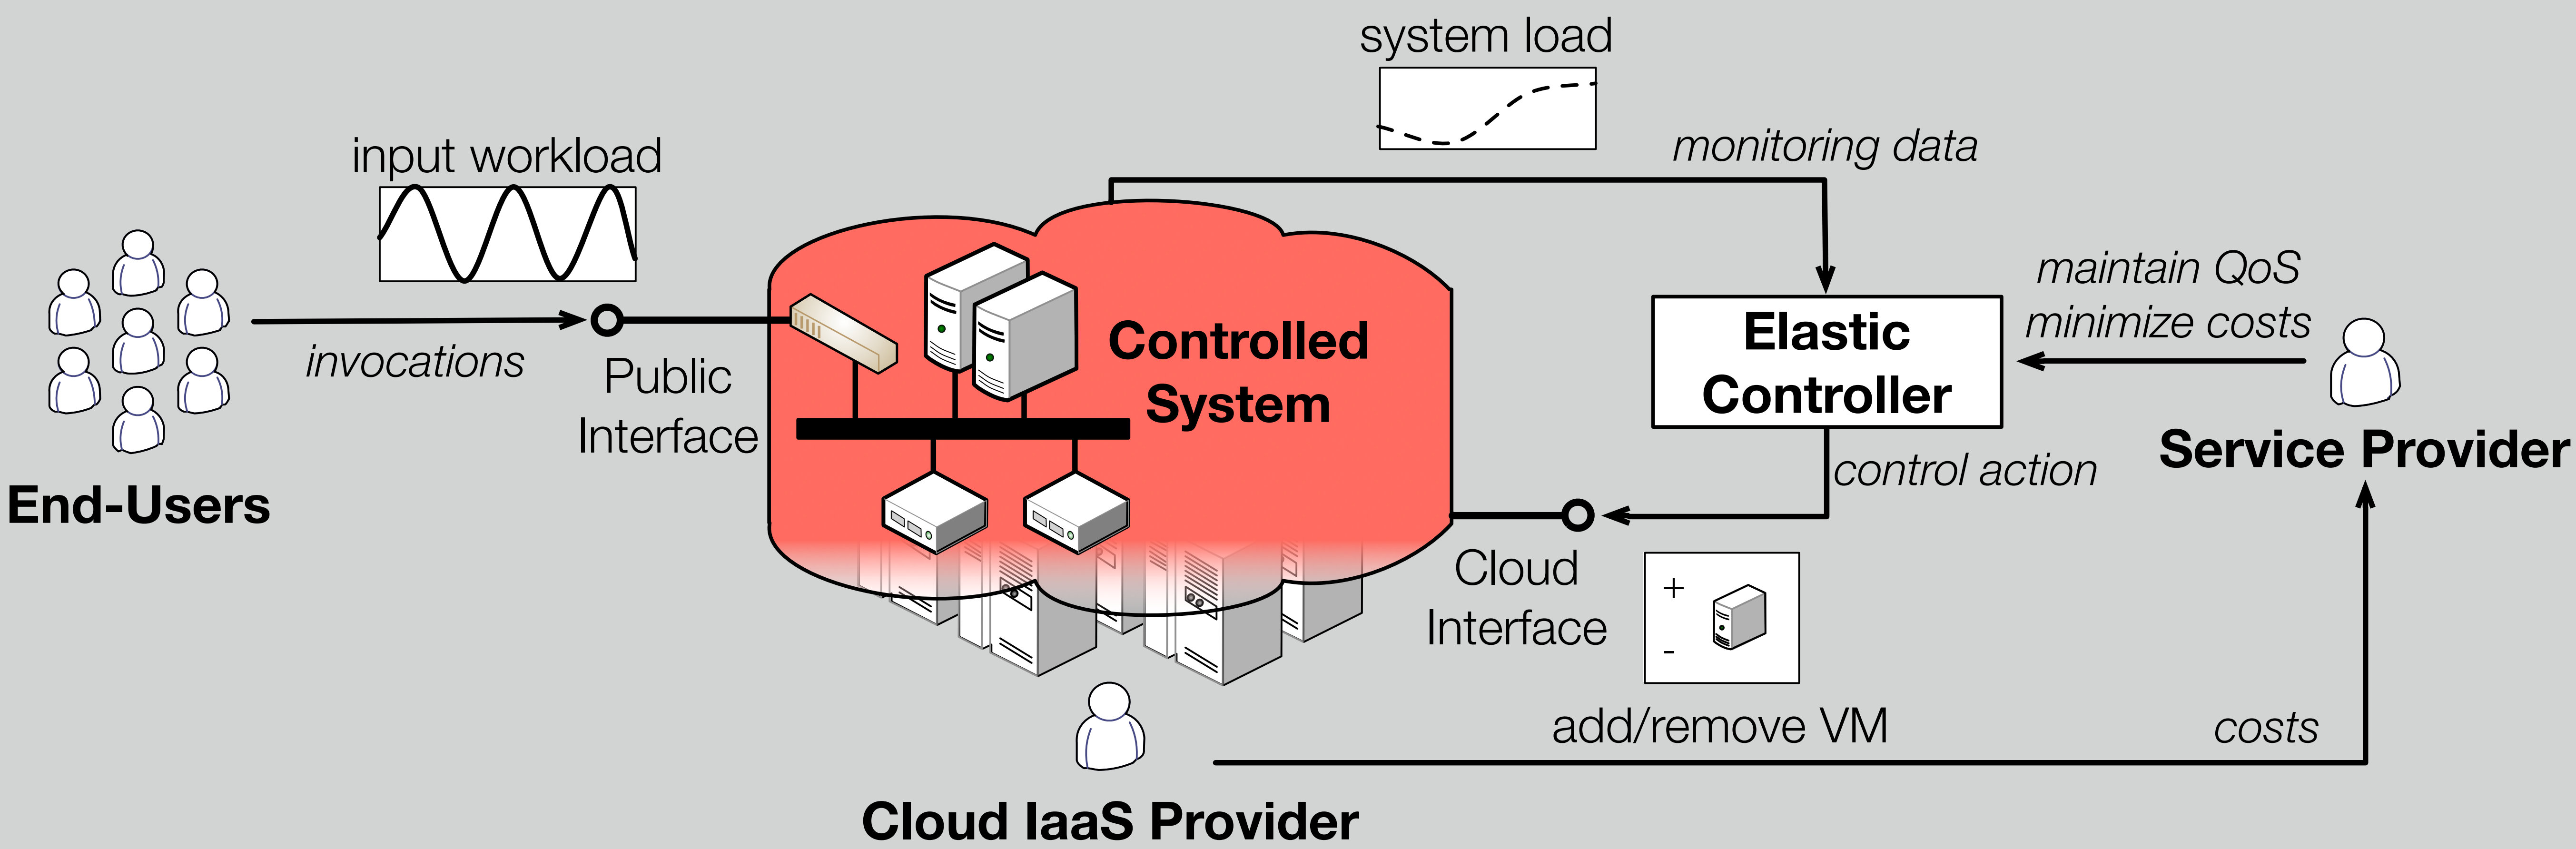

Cloud-Based Elastic System

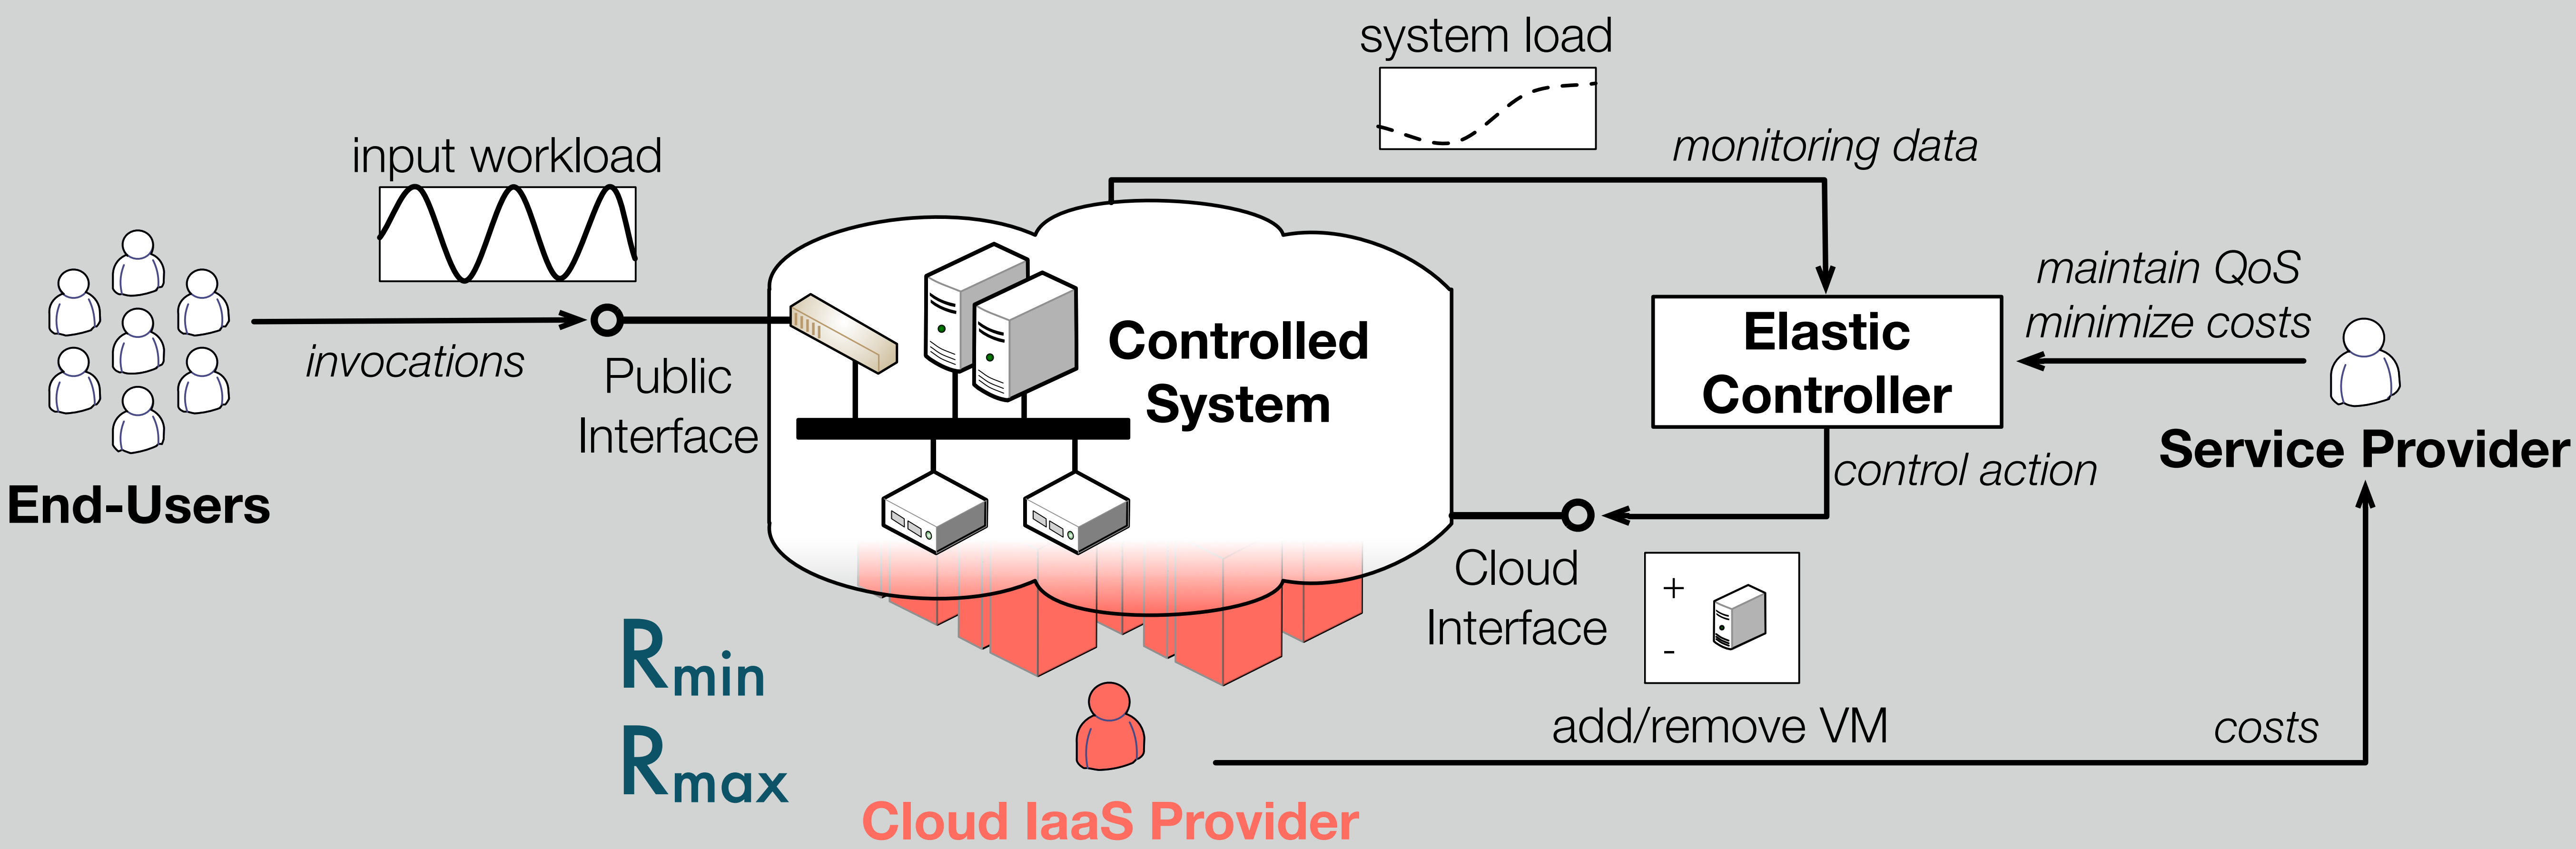

R_{min}

R_{max}

Cloud-Based Elastic System

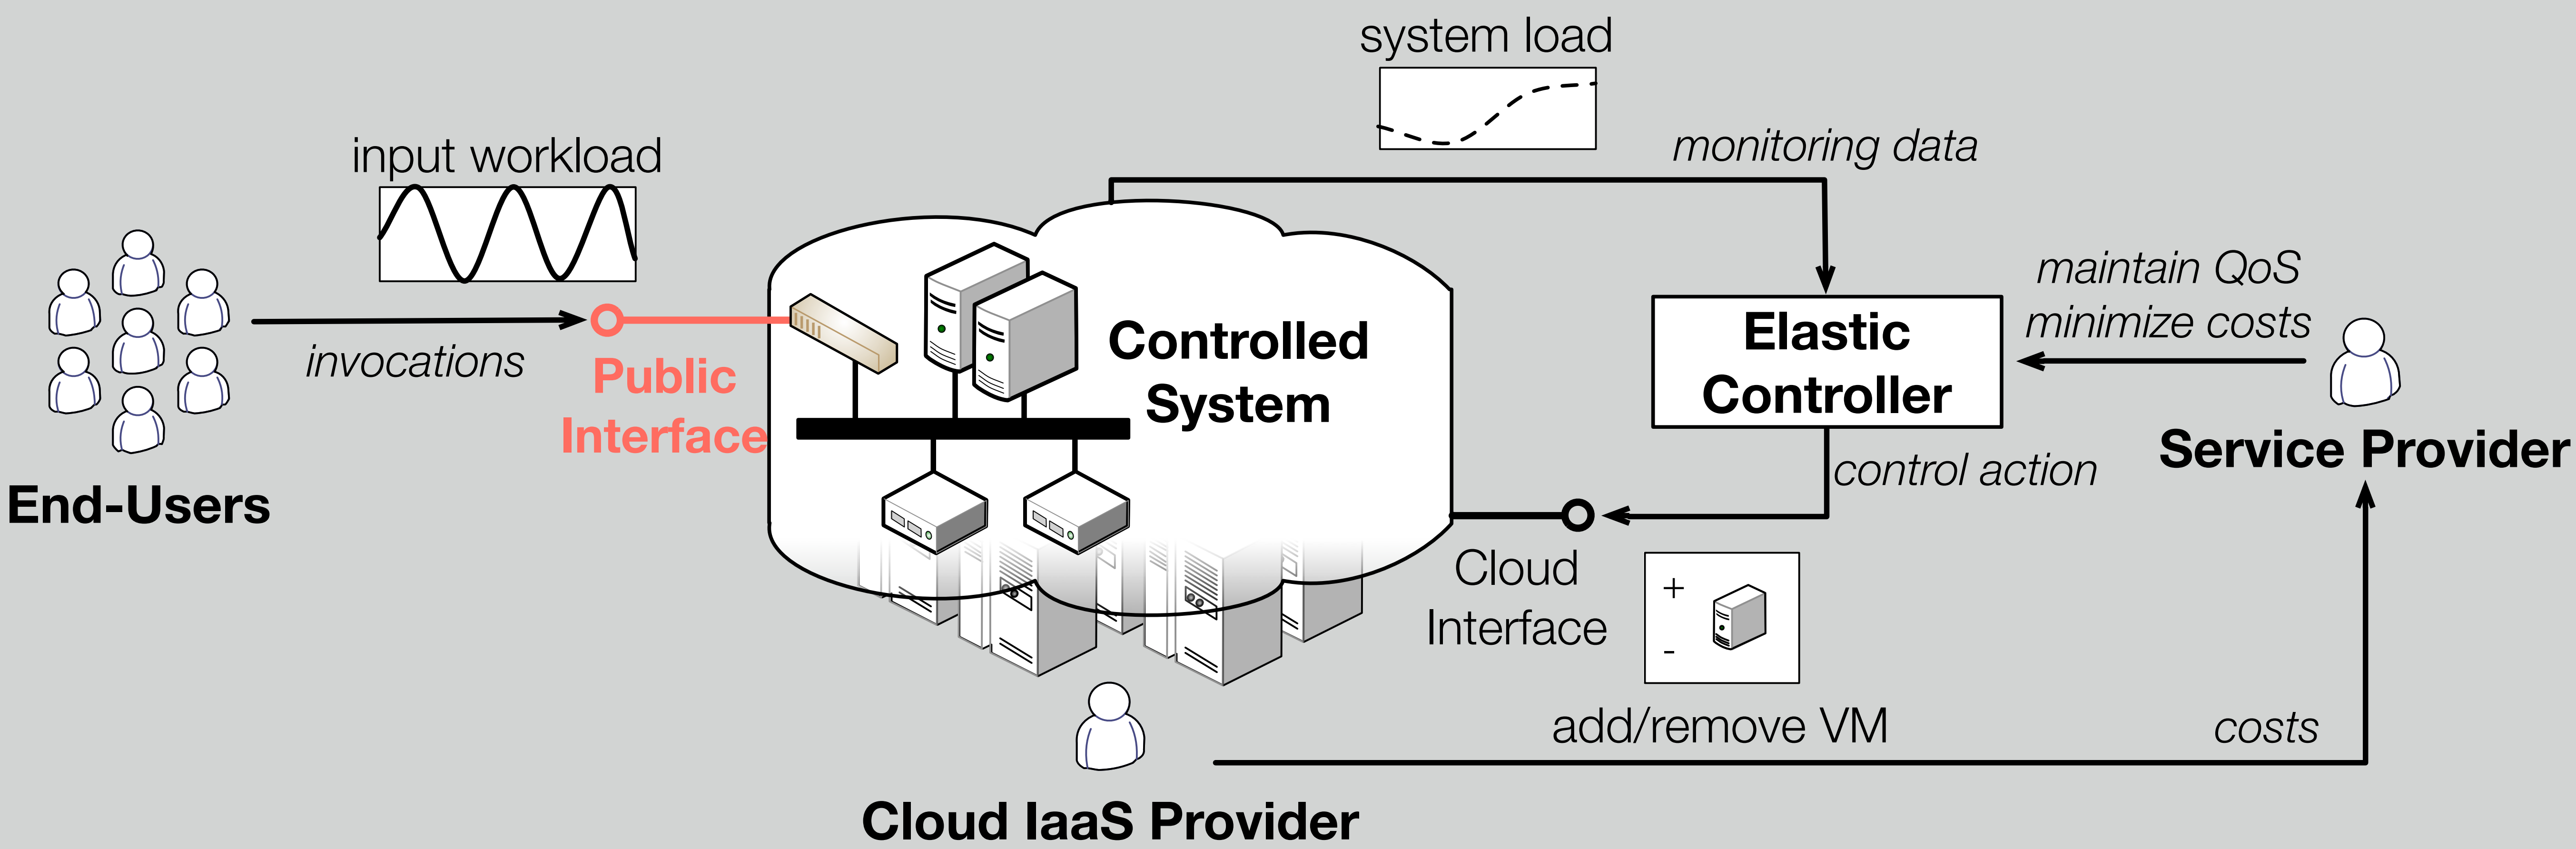

Cloud-Based Elastic System

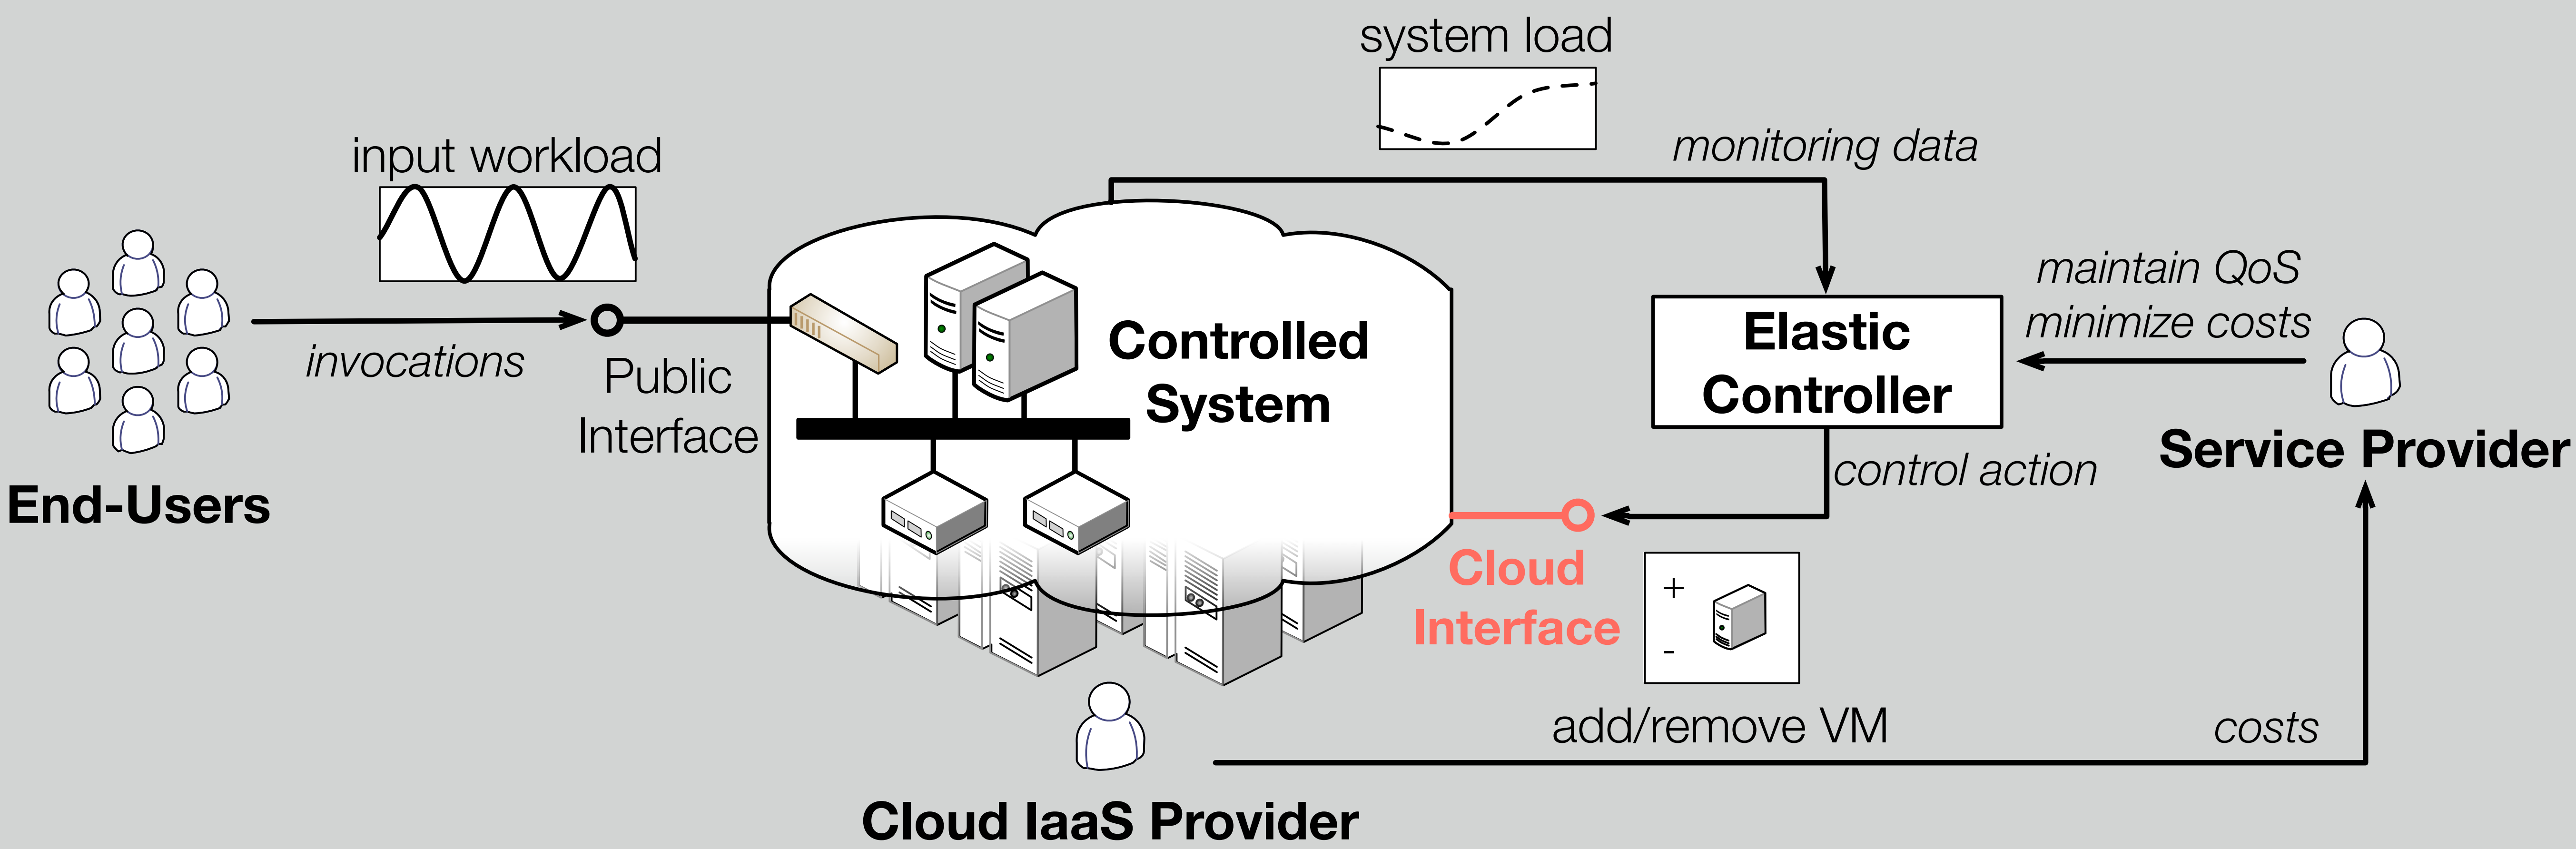

Cloud-Based Elastic System

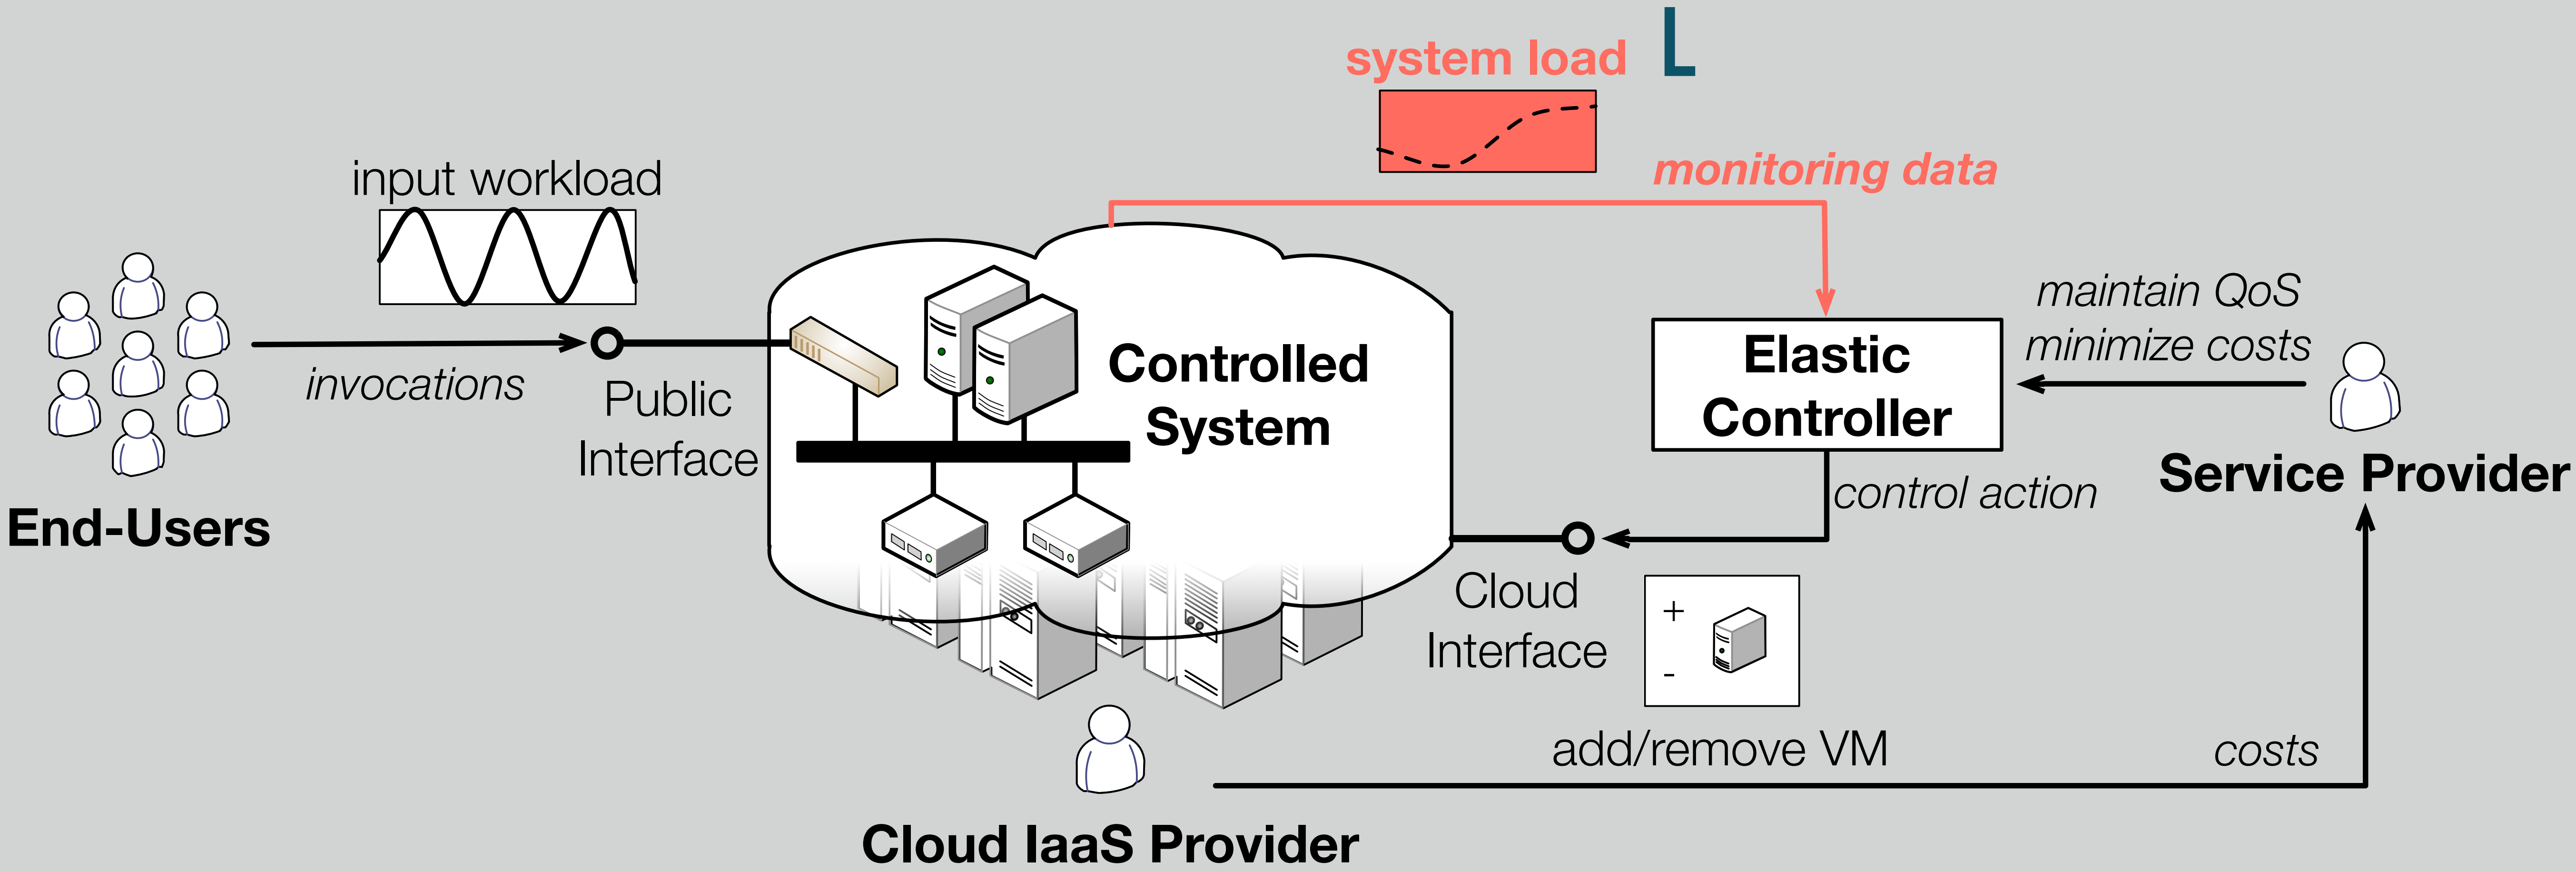

Cloud-Based Elastic System

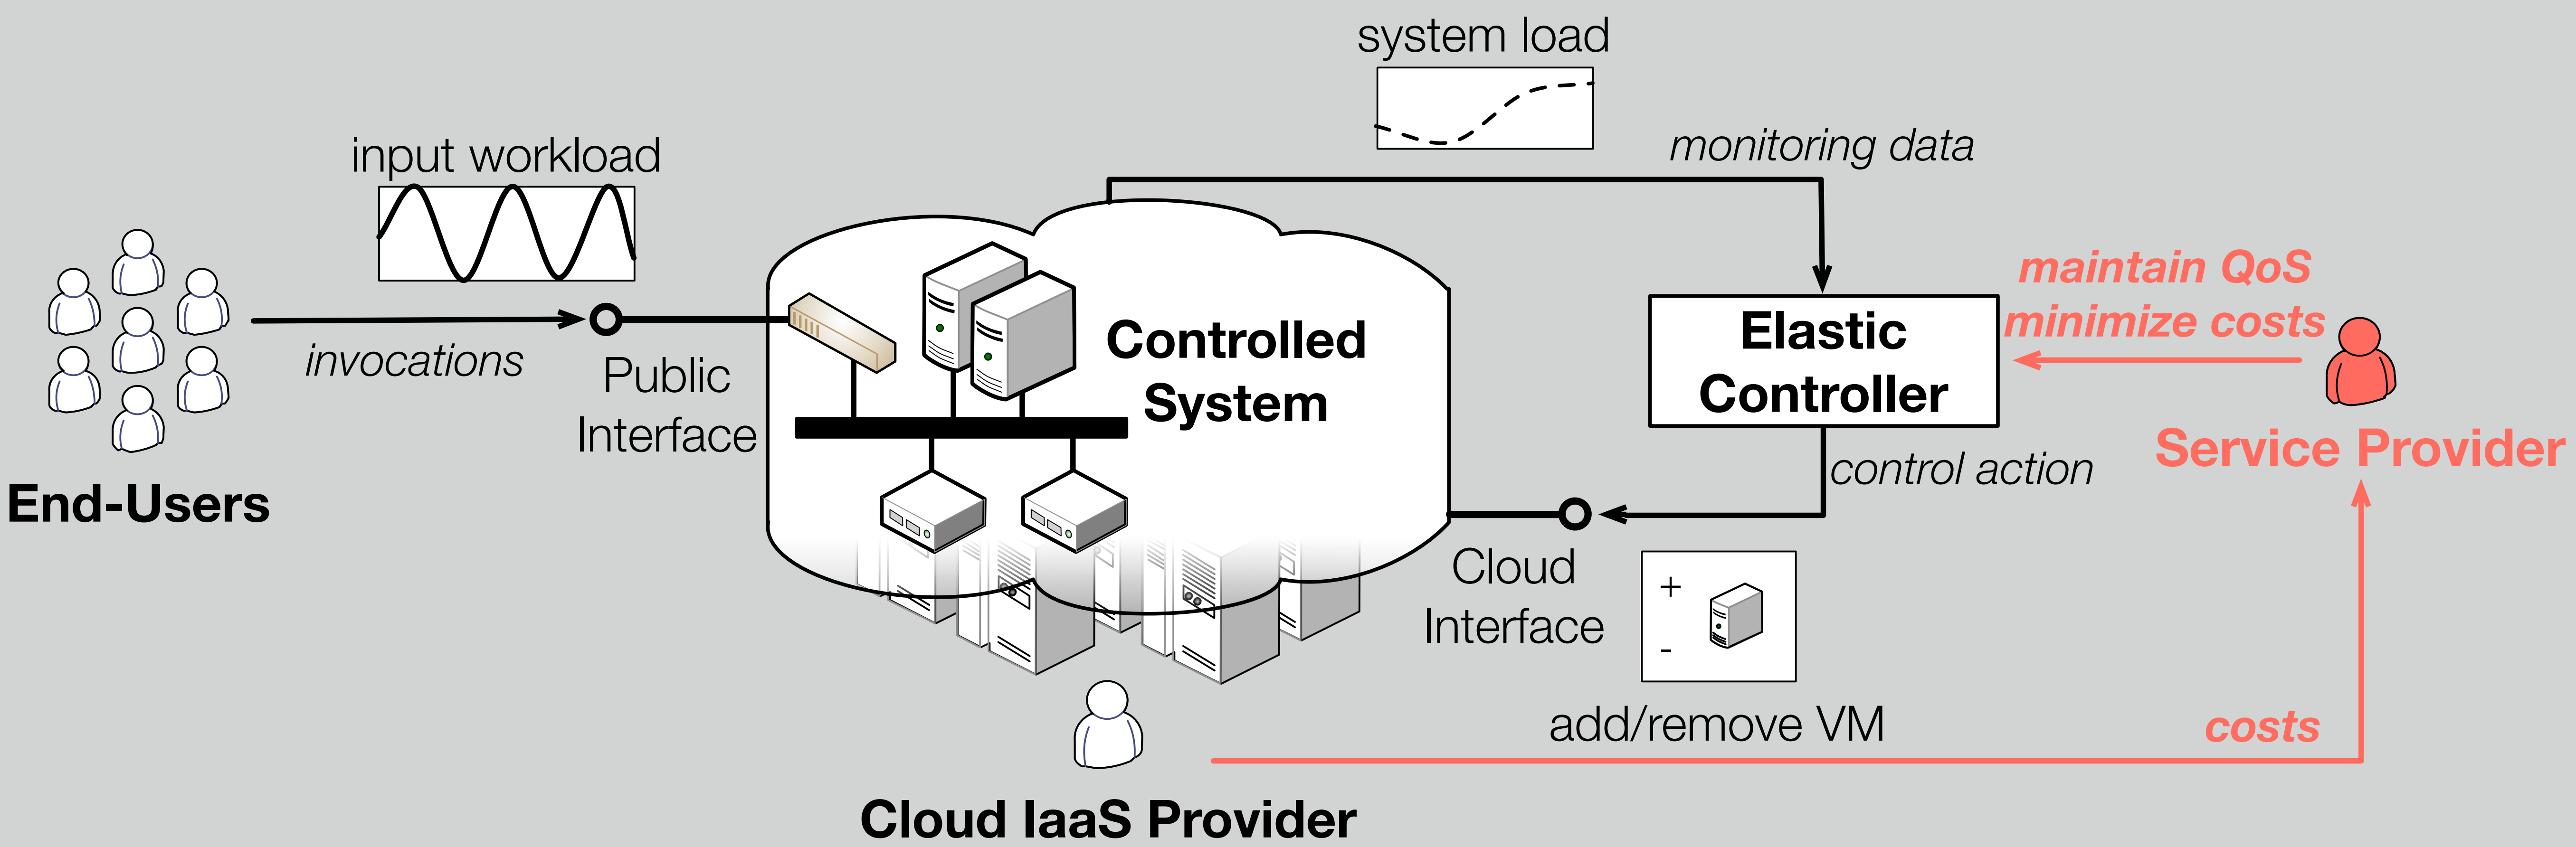

Cloud-Based Elastic System

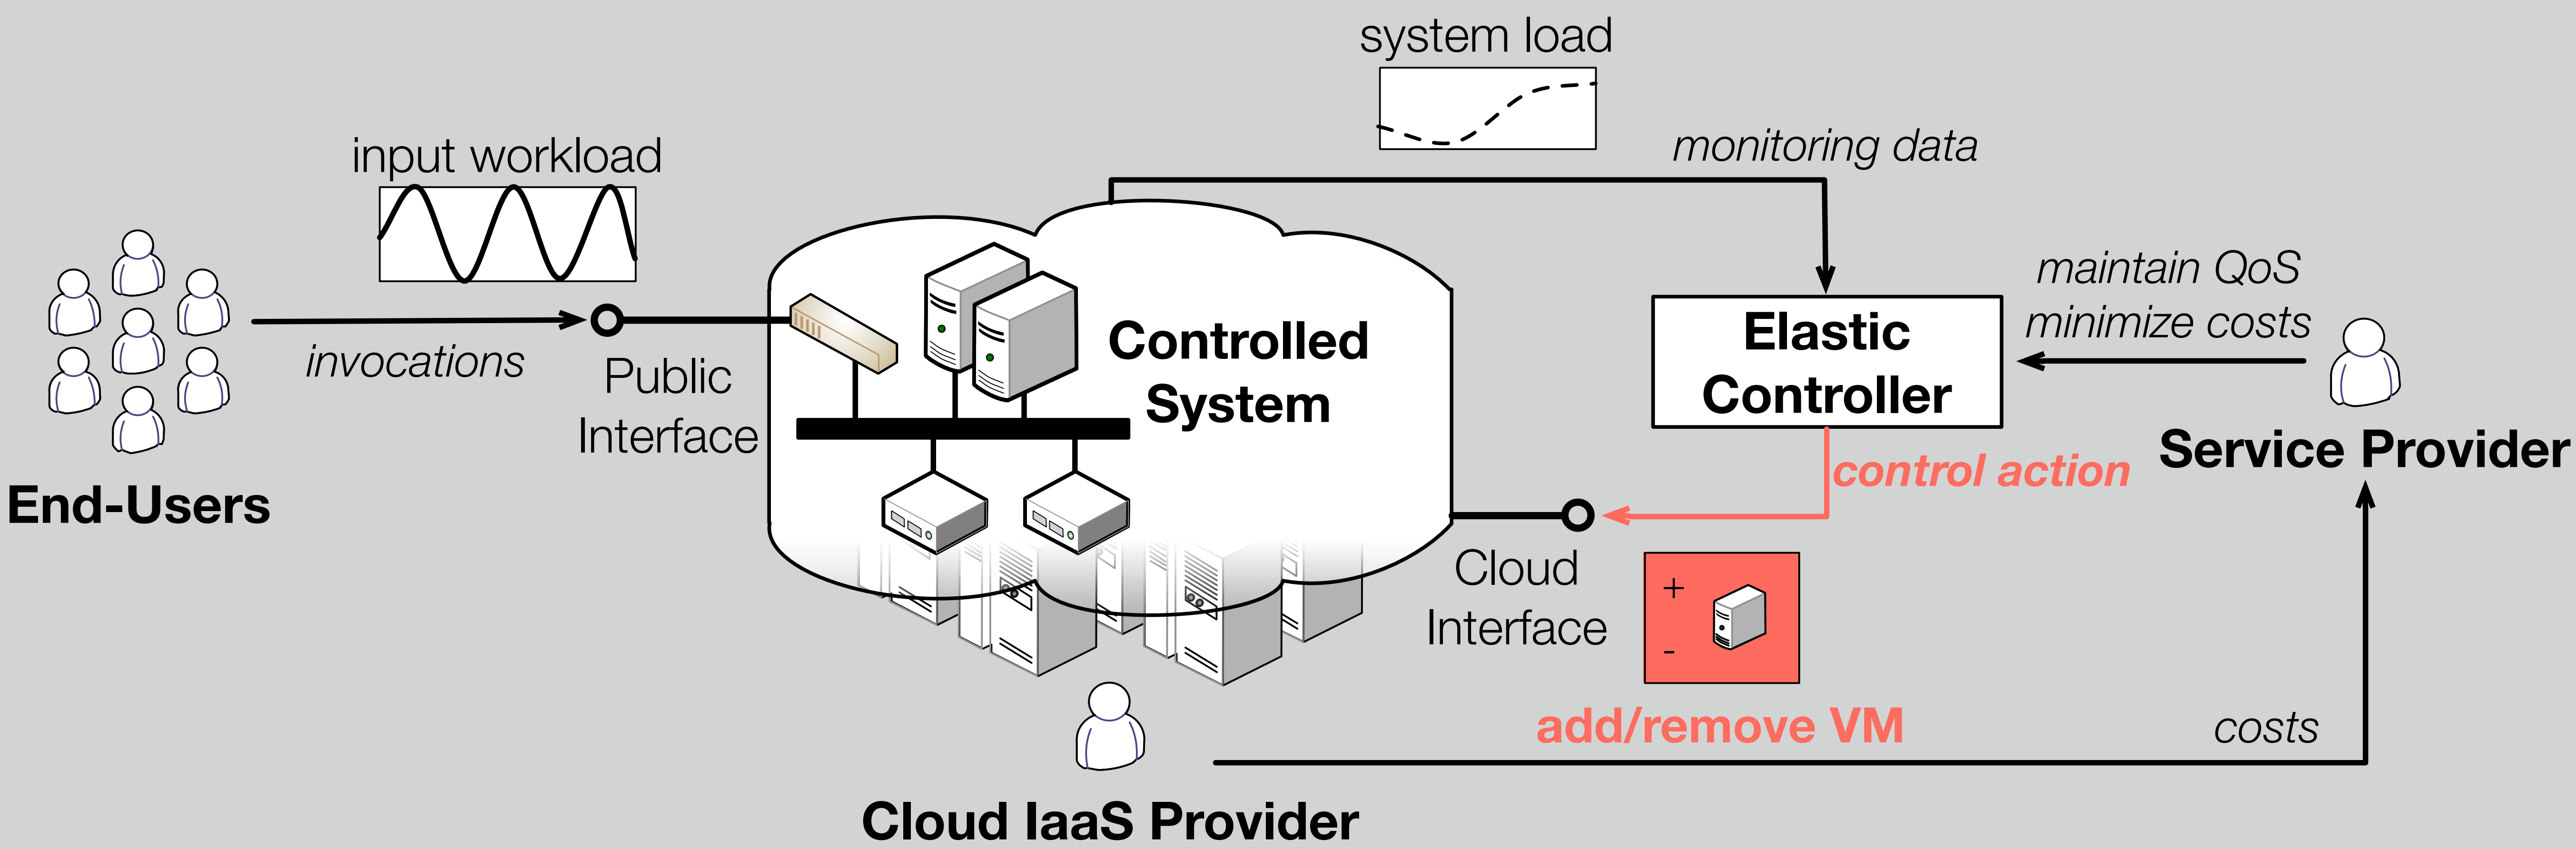

The First Step

Property Groups

Elasticity

Resource Management

Quality of Service

Property Groups

Elasticity

- Eagerness
- Sensitivity
- Plasticity

Resource Management

- Precision
- Oscillation
- Resource thrashing
- Cool-down period
- Bounded concurrent adaptations

Quality of Service

- Bounded QoS degradation
- Bounded actuation delay

$\mathrm{CLTL}^{\dagger}(\mathrm{D})$

MTL < $CLTL^{\dagger}(D)$ < MFOTL

$\mathrm{CLTL}^{\dagger}(D)$

CLTL^t(D)

Linear Temporal Logic

$$G(p \rightarrow XX(q))$$

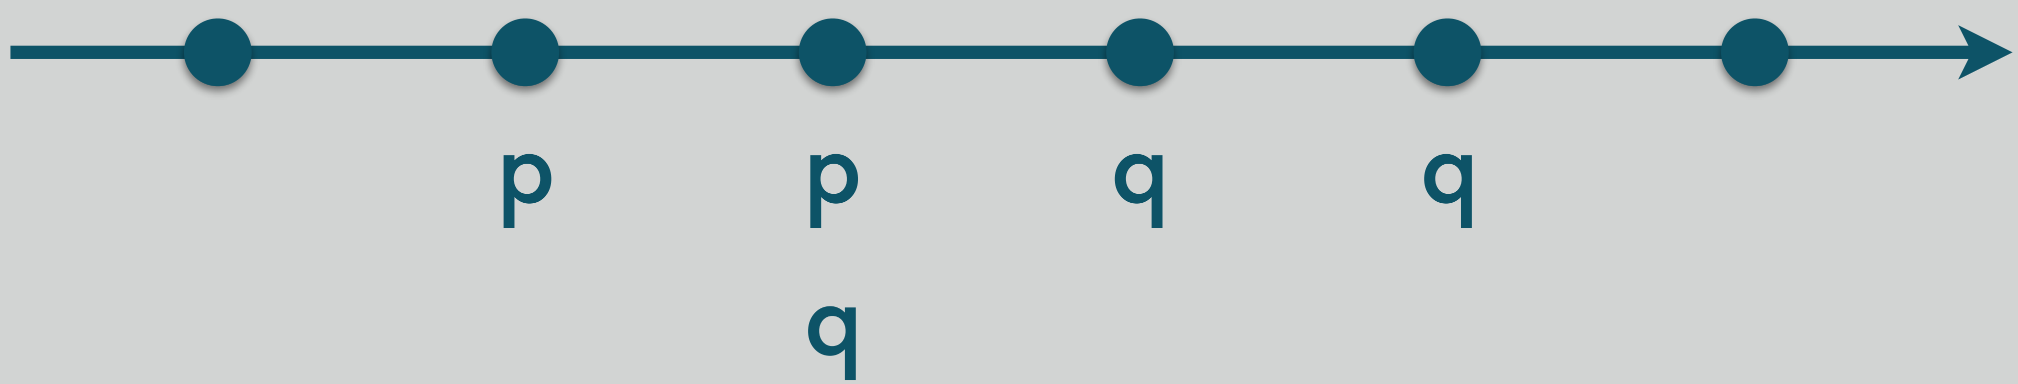

"It is always true that if p occurs then q occurs 2 positions afterwards"

$\mathrm{CLTL}^{t}(\mathrm{D})$

Linear Temporal Logic

$$G(p \rightarrow XX(q))$$

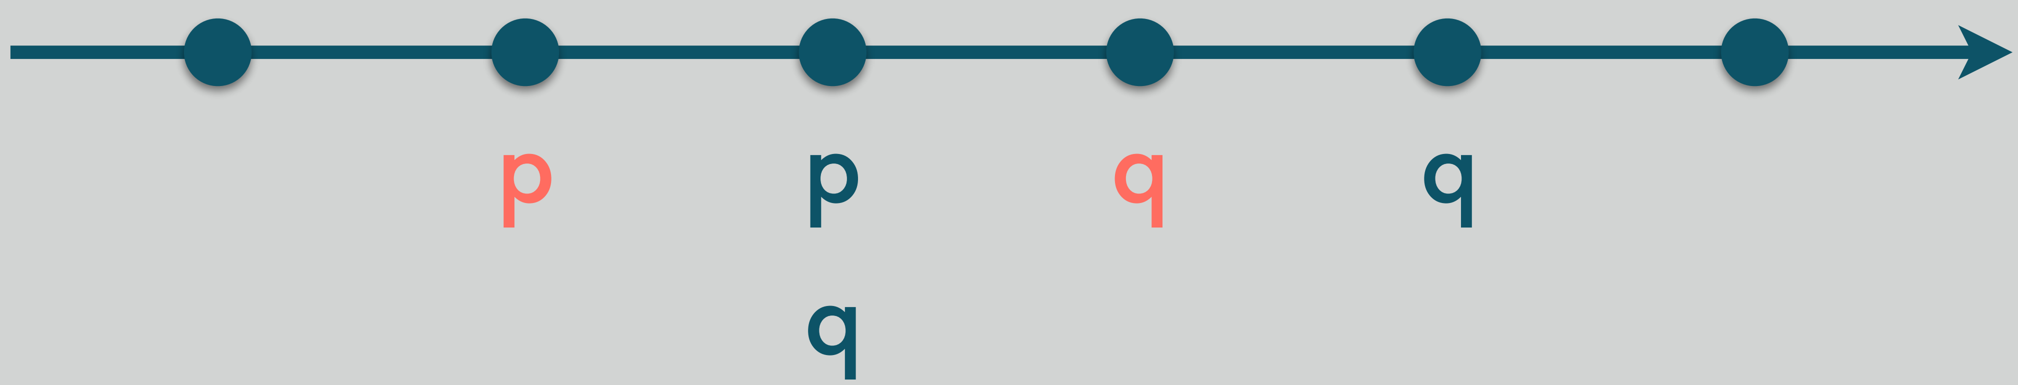

"It is always true that if p occurs then q occurs 2 positions afterwards"

CLTL^{t}(D)

Linear Temporal Logic

$$G(p \rightarrow XX(q))$$

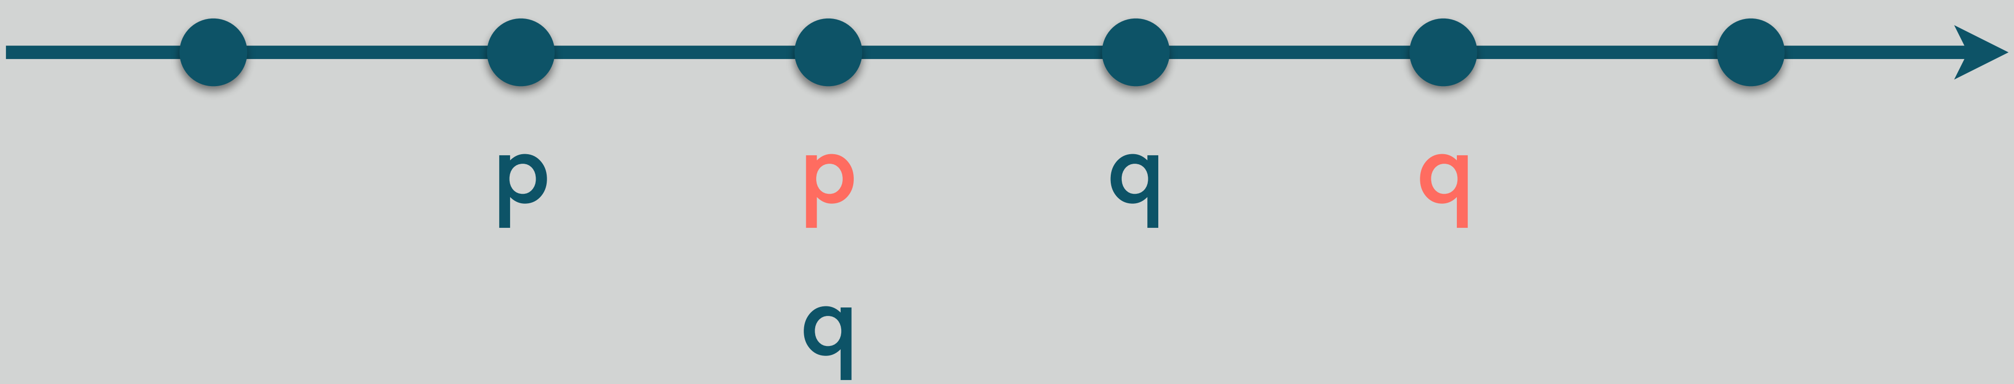

"It is always true that if p occurs then q occurs 2 positions afterwards"

CLTL^{t}(D)

Constraint Linear Temporal Logic (over constraint system D)

$$G(p \leftrightarrow X(z) = z + 1)$$

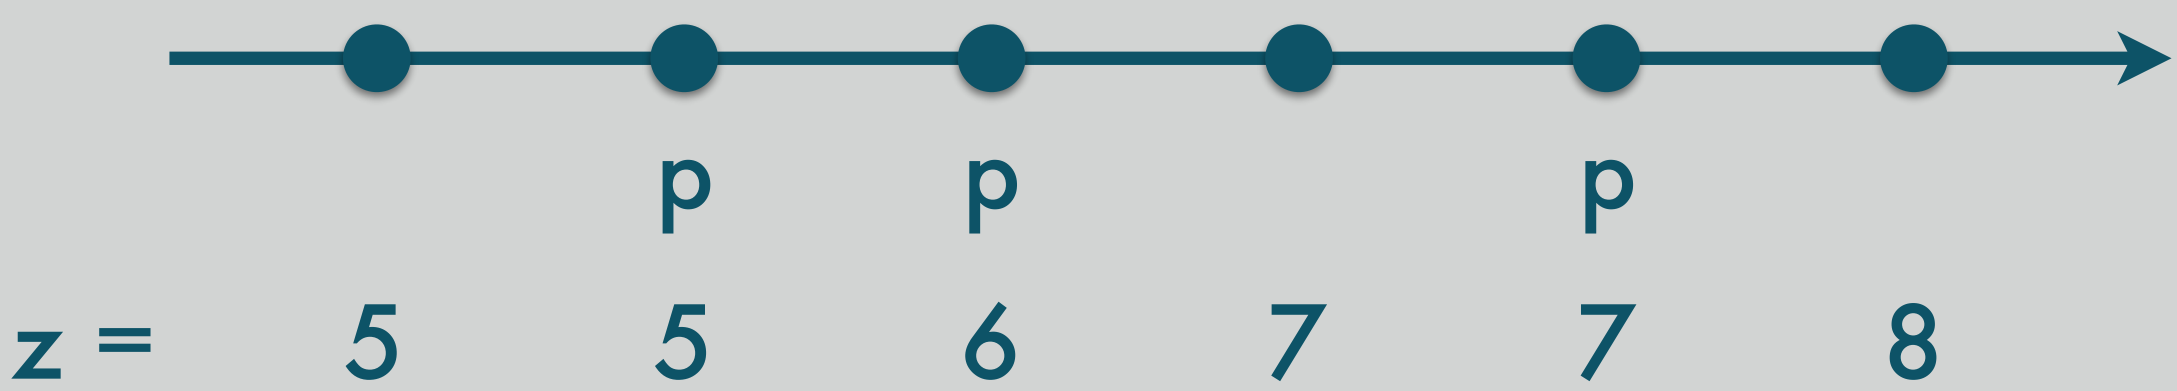

“There is p if and only if variable z is incremented by 1 in the next position”

CLTL^f(D)

Constraint Linear Temporal Logic (over constraint system D)

$$\mathsf{G}(p \leftrightarrow \mathsf{X}(z) = z + 1)$$

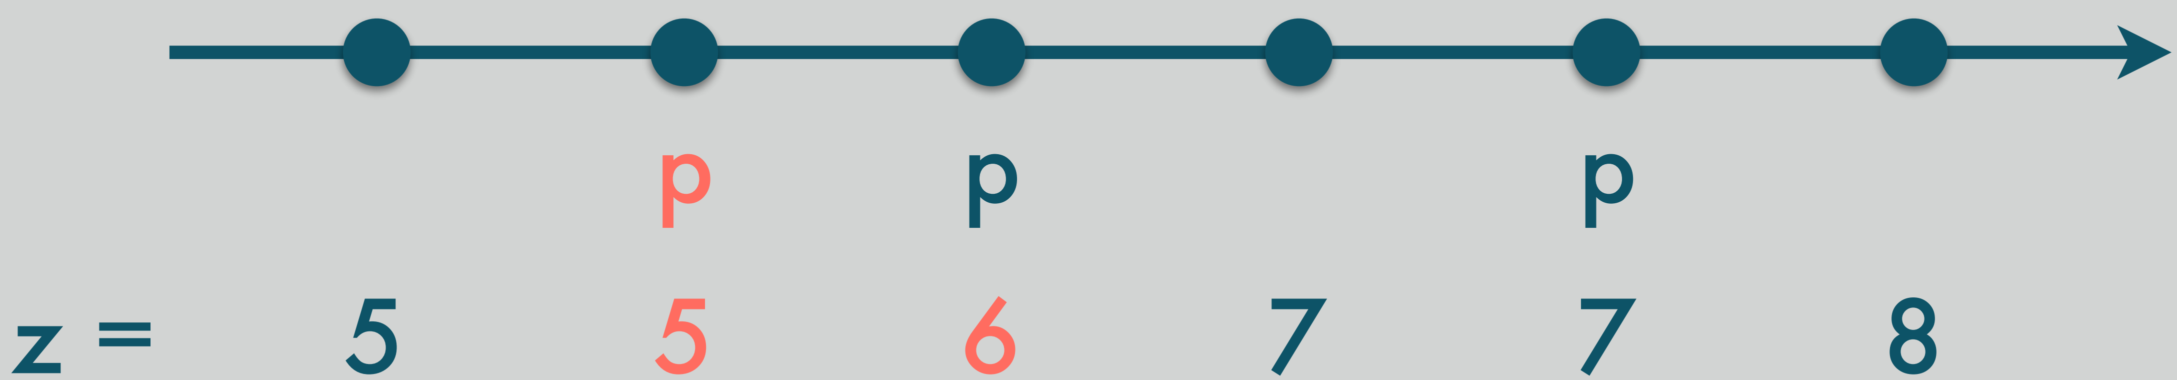

"There is p if and only if variable z is incremented by 1 in the next position"

CLTL^t(D)

Constraint Linear Temporal Logic (over constraint system D)

$$G(p \leftrightarrow X(z) = z + 1)$$

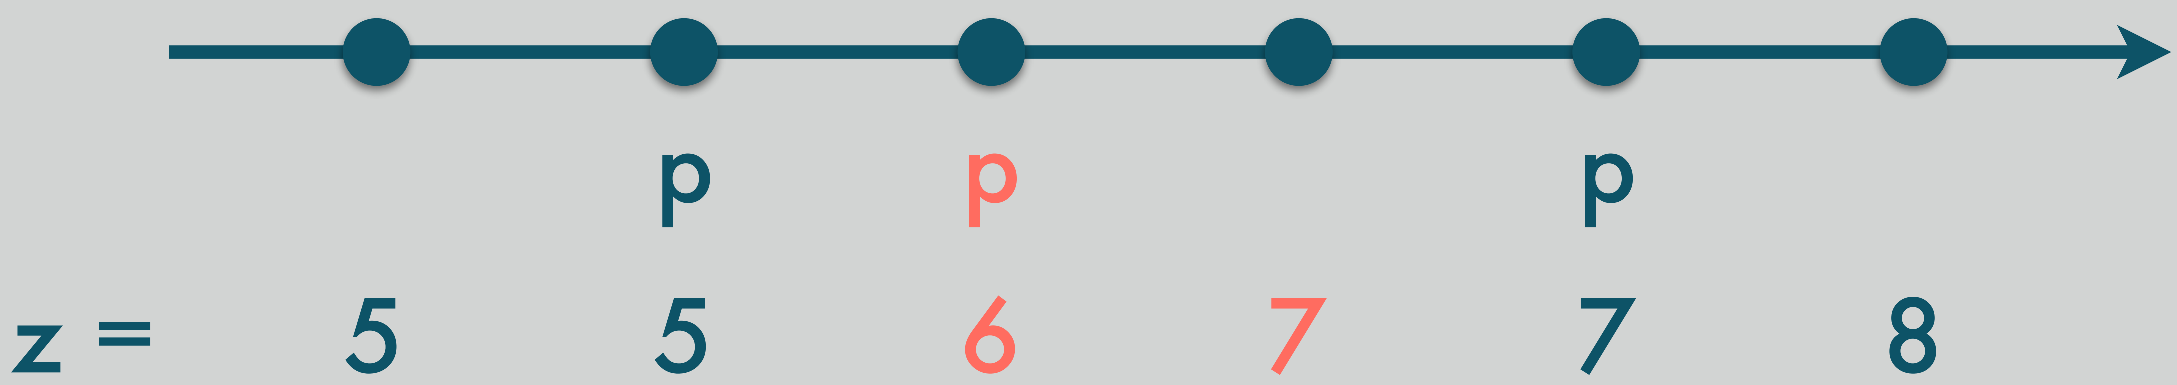

"There is p if and only if variable z is incremented by 1 in the next position"

CLTL^{t}(D)

Constraint Linear Temporal Logic (over constraint system D)

$$G(p \leftrightarrow X(z) = z + 1)$$

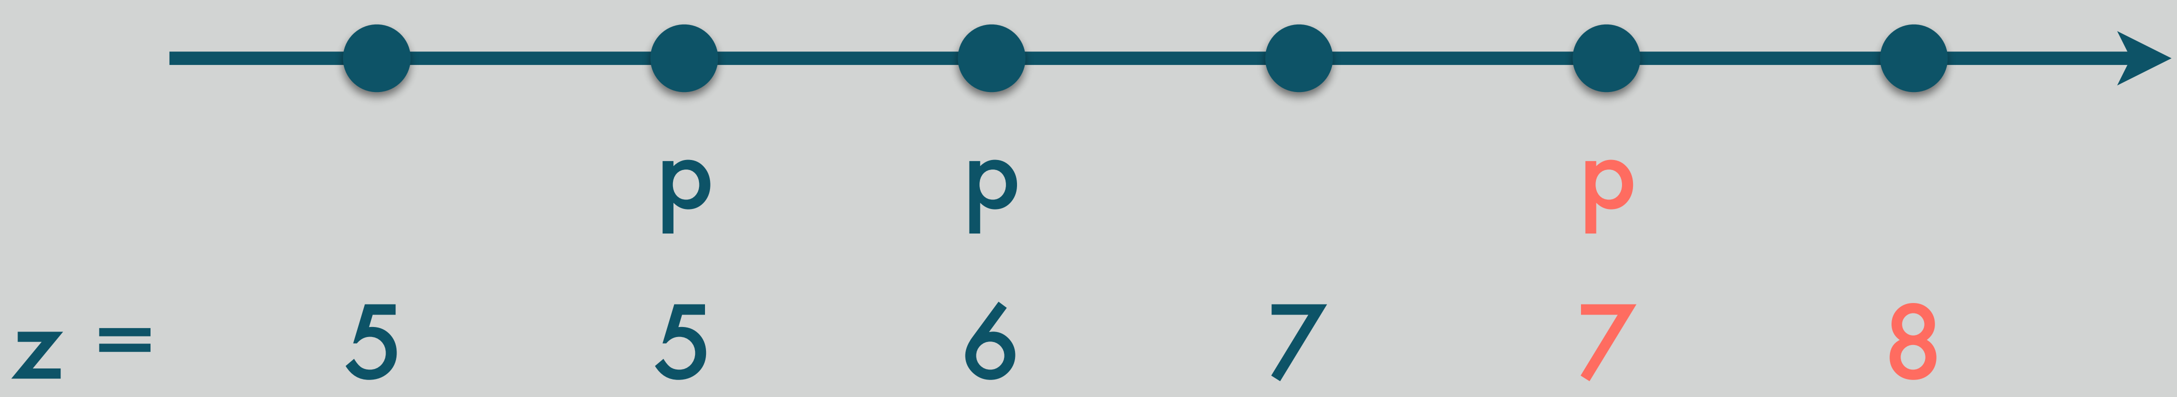

"There is p if and only if variable z is incremented by 1 in the next position"

CLTL^t(D)

Timed Constraint Linear Temporal Logic

$$G(p \leftrightarrow X_{(0,4)}(z = 0))$$

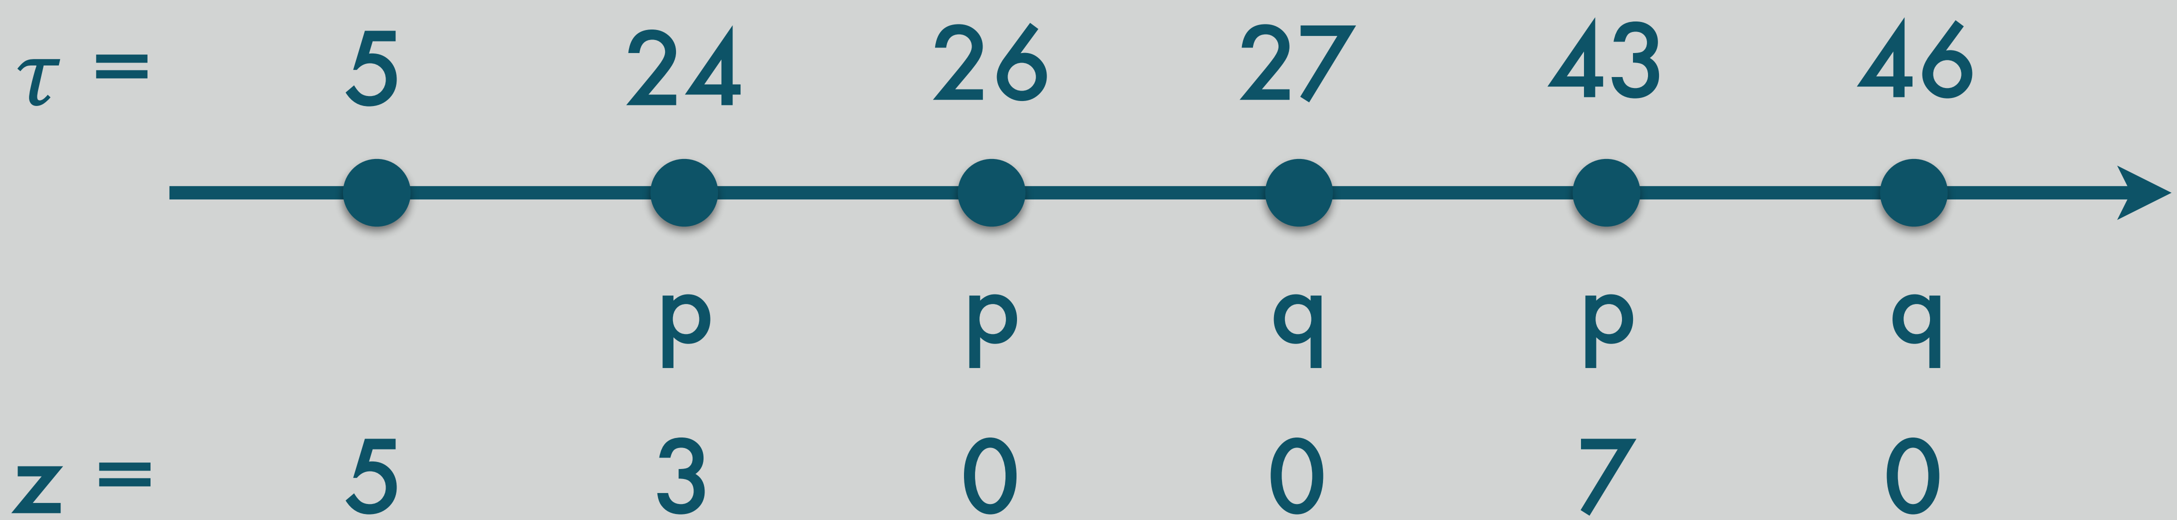

"There is p if and only if z is equal to 0 in the next position which must occur within 4 seconds"

CLTL^t(D)

Timed Constraint Linear Temporal Logic

$$G(p \leftrightarrow X_{(0,4)}(z = 0))$$

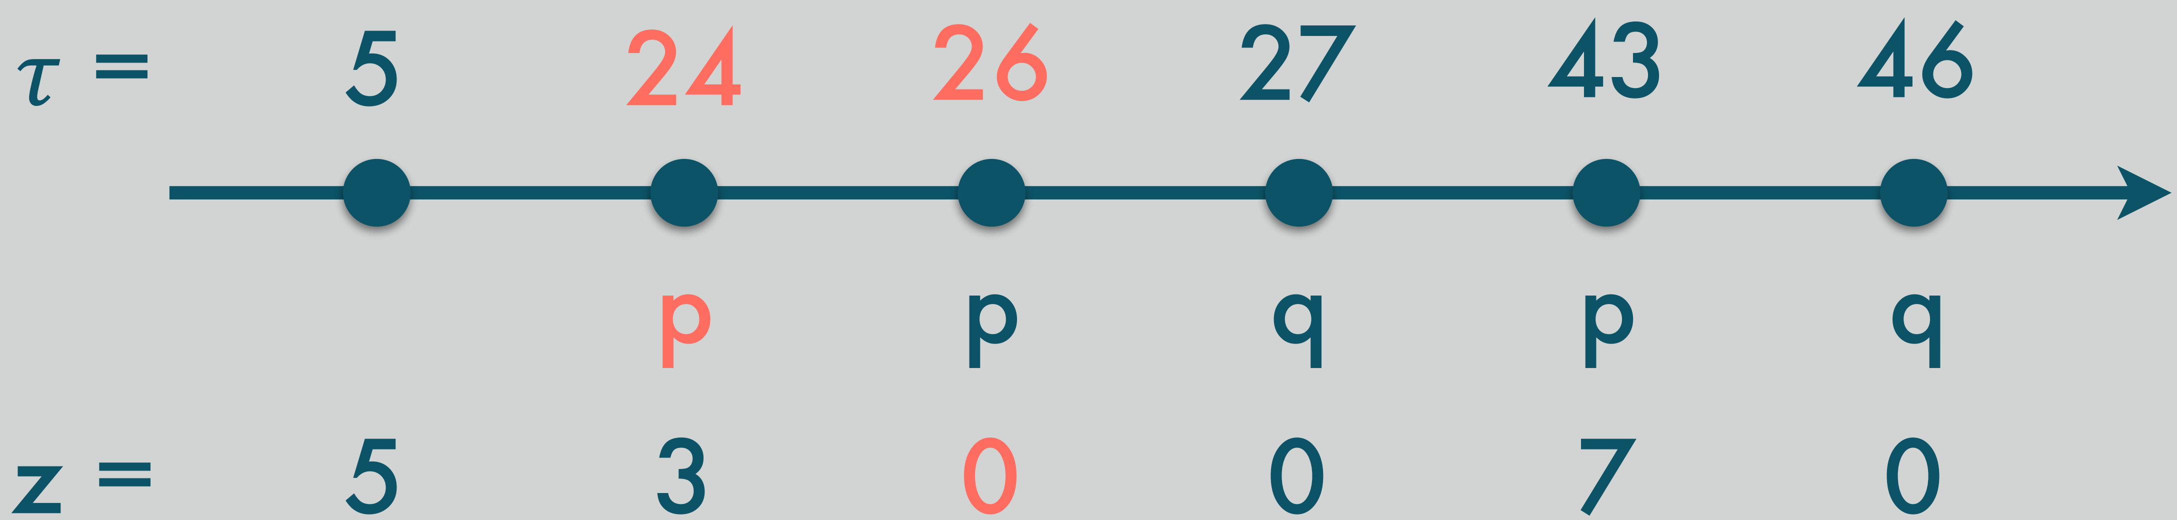

"There is p if and only if z is equal to 0 in the next position which must occur within 4 seconds"

CLTL^t(D)

Timed Constraint Linear Temporal Logic

$$G(p \leftrightarrow X_{(0,4)}(z = 0))$$

τ =	5	24	26	27	43	46
		p	p	q	p	q
z =	5	3	0	0	7	0

“There is p if and only if z is equal to 0 in the next position which must occur within 4 seconds”

CLTL^t(D)

Timed Constraint Linear Temporal Logic

$$G(p \leftrightarrow X_{(0,4)}(z = 0))$$

τ =	5	24	26	27	43	46
		p	p	q	p	q
z =	5	3	0	0	7	0

"There is p if and only if z is equal to 0 in the next position which must occur within 4 seconds"

CLTL†(D)

Virtual Machines

VM Lifecycle

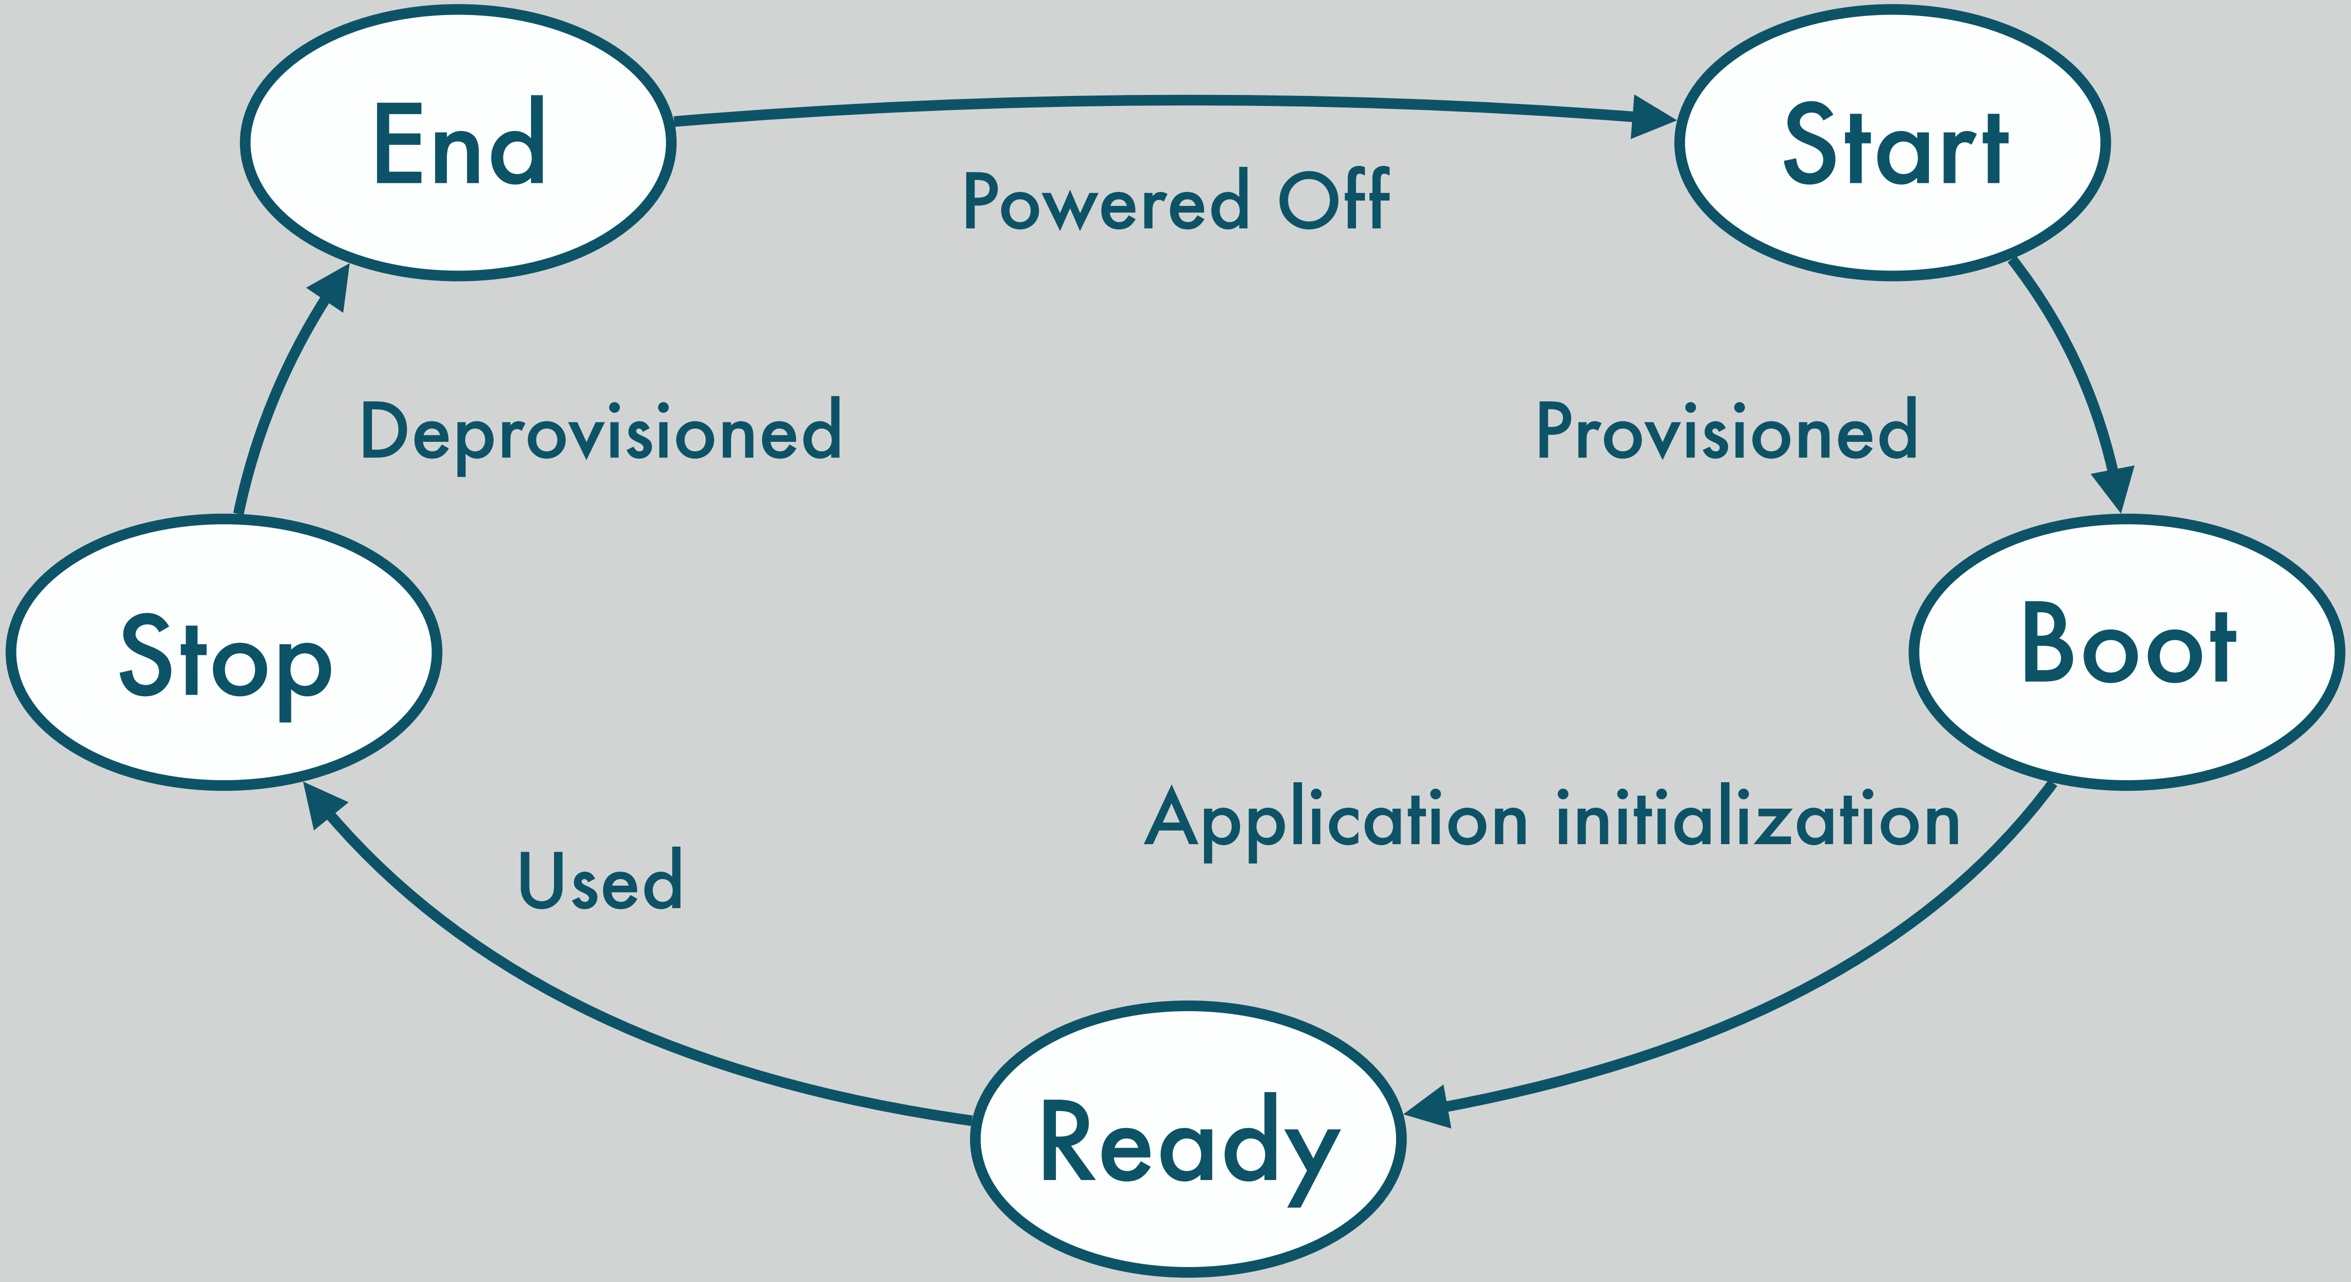

VM Lifecycle

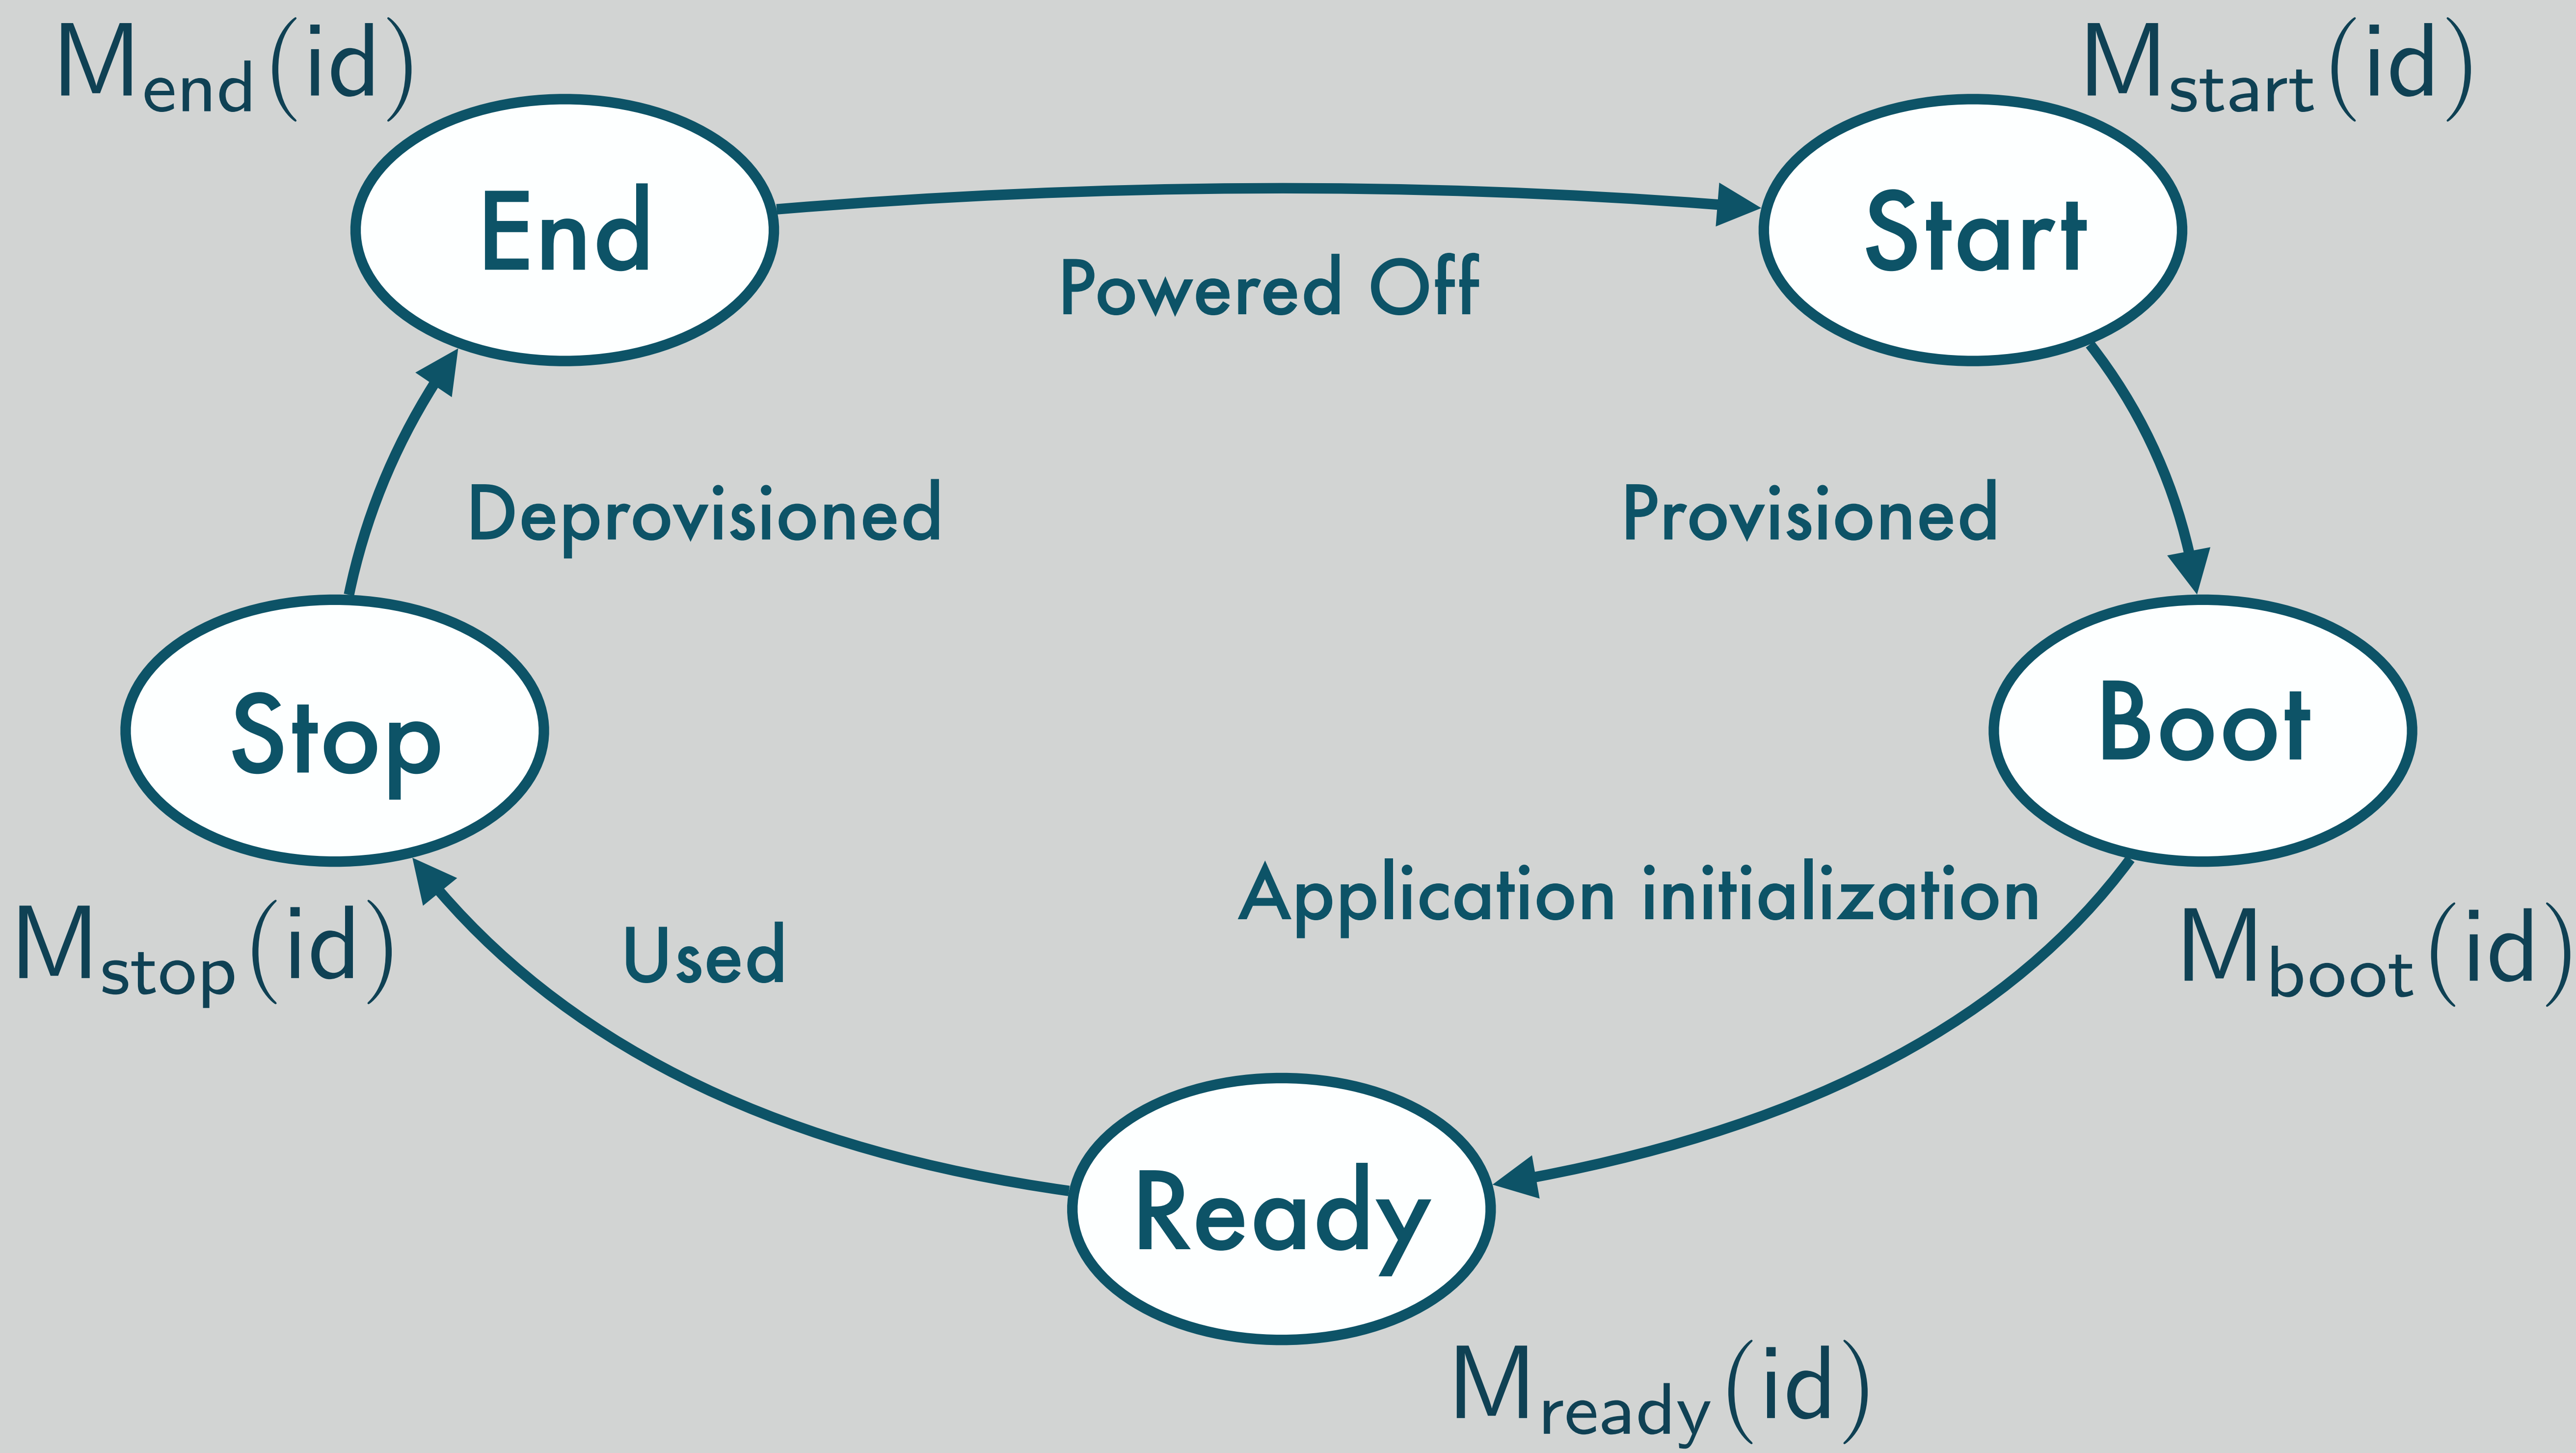

Example monitored execution

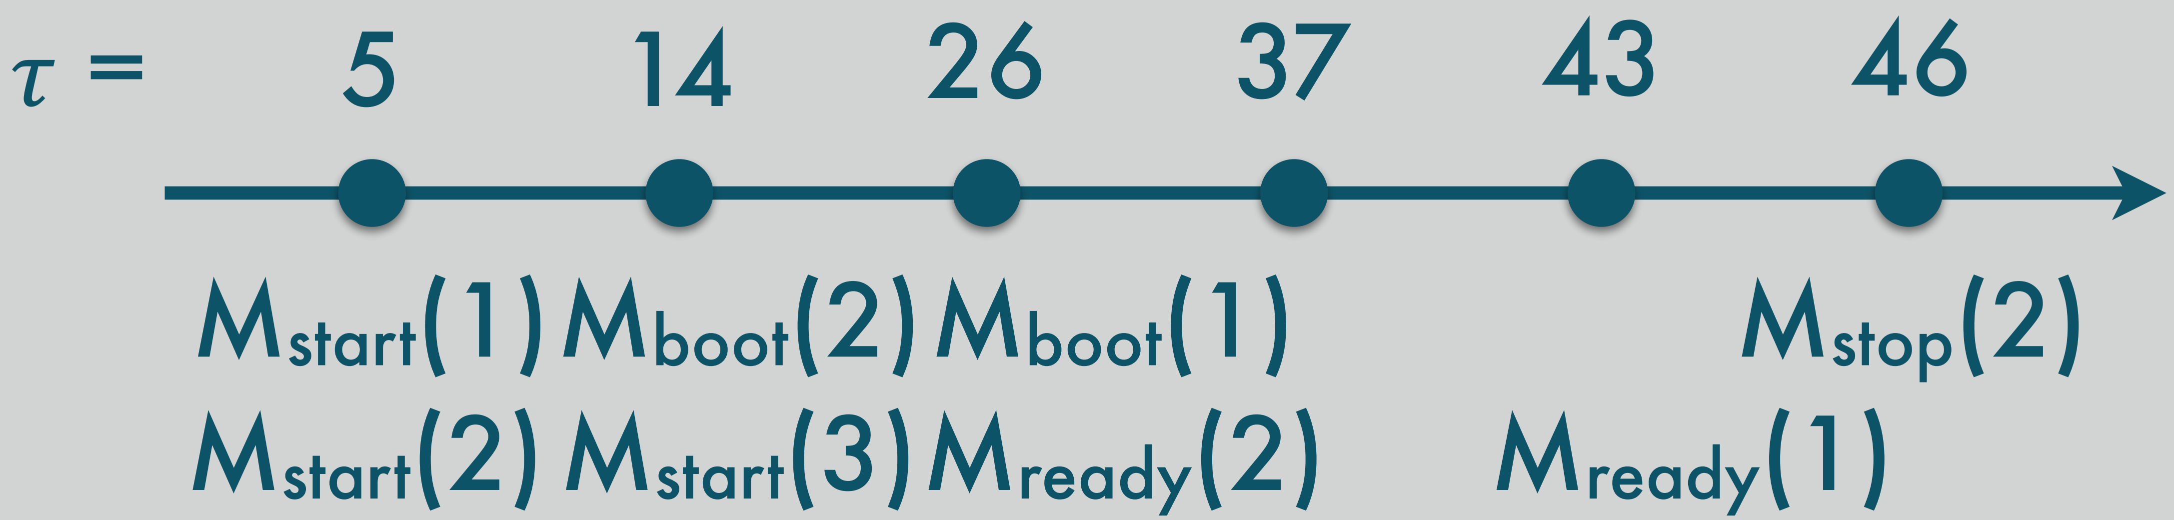

Allocated Resources

$R = R_{min}$

$\forall id : \mathrm{G}(M_{start}(id) \rightarrow R = \mathrm{Y}(R) + 1)$

$\exists id : \mathrm{G}(R = \mathrm{Y}(R) + 1 \rightarrow M_{start}(id))$

$\mathrm{G}((\forall id : \neg M_{start}(id) \wedge \neg M_{stop}(id)) \leftrightarrow R = \mathrm{Y}(R))$

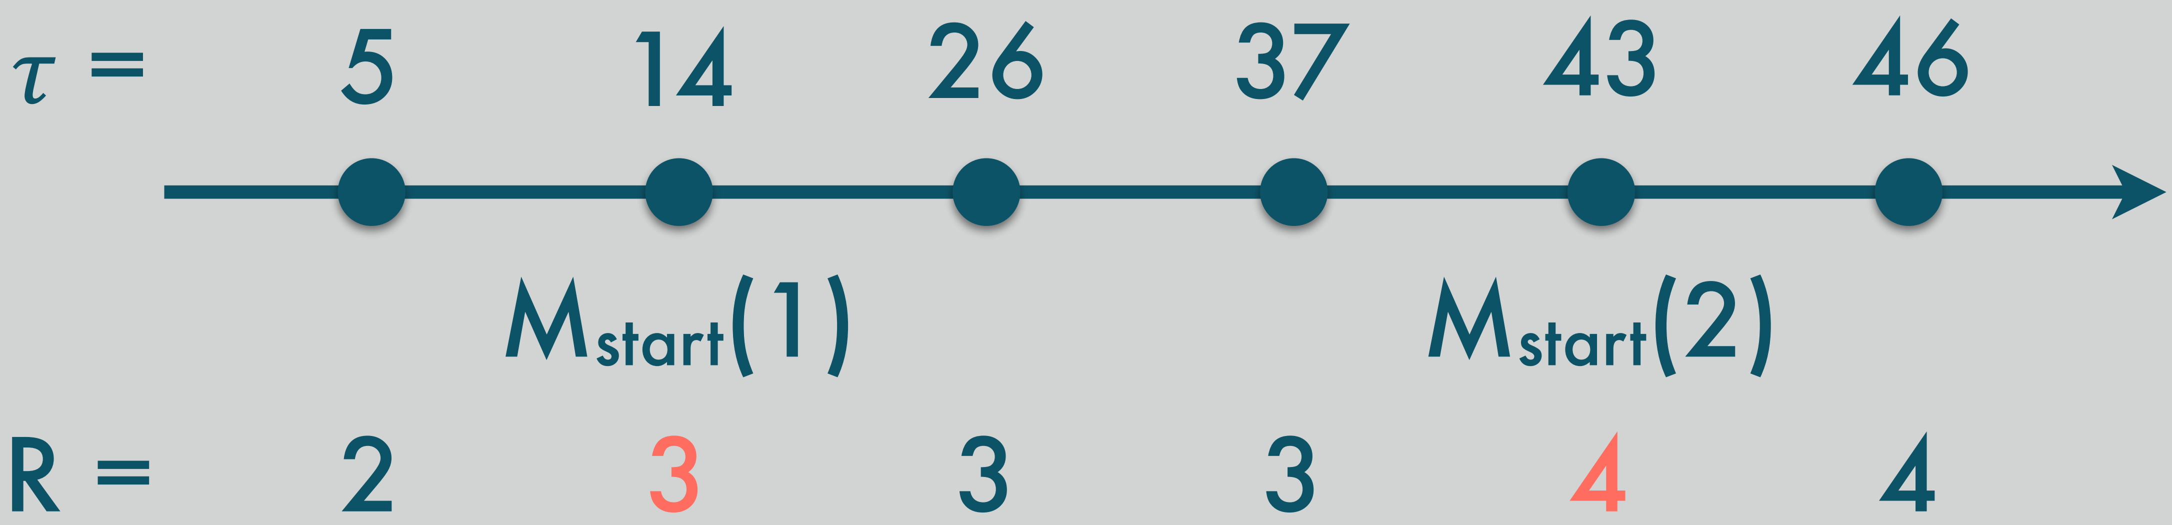

Allocated Resources

$R = R_{min}$

$\forall id : G(M_{stop}(id) \rightarrow R = Y(R) - 1)$

$\exists id : G(R = Y(R) - 1 \rightarrow M_{stop}(id))$

$G((\forall id : \neg M_{start}(id) \wedge \neg M_{stop}(id)) \leftrightarrow R = Y(R))$

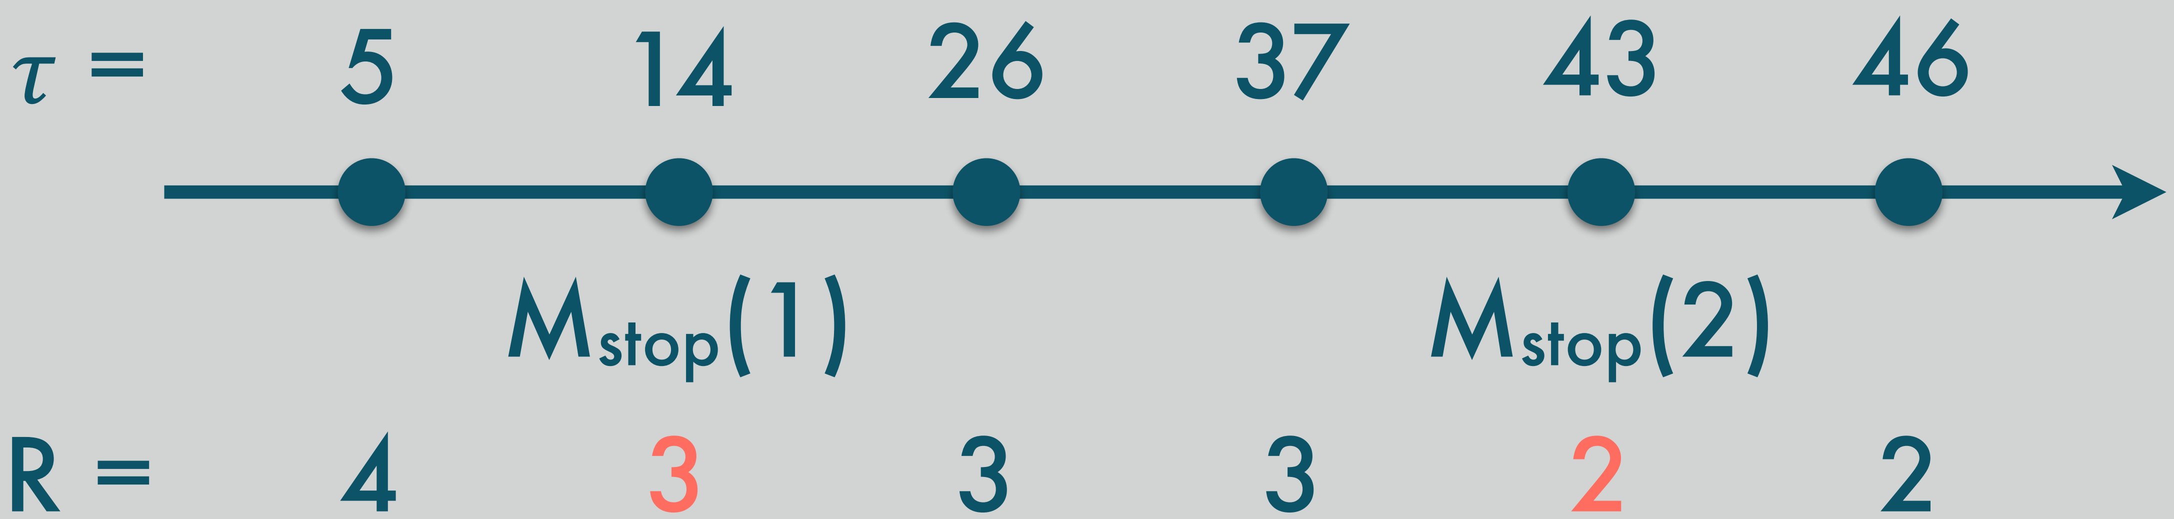

Elasticity

"Capabilities can be rapidly and elastically provisioned to quickly scale out, and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time."

–National Institute of Standards and Technology (NIST)

Elasticity

Eagerness

Sensitivity

Plasticity

Eagerness

"Eagerness captures responsiveness of a system to the changes in the workload."

Sensitivity

"Sensitivity captures robustness of a system to changes in the load which are below a certain threshold."

Eagerness and Sensitivity

$L_a = 0$

$\mathsf{G}((-\Delta \leq L_a \leq \Delta) \rightarrow \mathsf{X}(L_a) = L_a + \mathsf{X}(L) - L)$

$\mathsf{G}((L_a > \Delta) \rightarrow (\mathsf{X}(L_a) = \mathsf{X}(L) - L \wedge \mathsf{F}_{(0,T_e]}(\mathsf{X}(R) > R)))$

$\mathsf{G}((L_a < -\Delta) \rightarrow (\mathsf{X}(L_a) = \mathsf{X}(L) - L \wedge \mathsf{F}_{(0,T_e]}(\mathsf{X}(R) < R)))$

Te, Δ=2

L	5	6	7	8	9	9	4	3	4	5	6	6	6
L_a	0	1	2	3	1	1	-4	-1	0	1	2	2	2
					R↑			R↓					

Plasticity

"When the load drops to zero an elastic system must be able to deallocate all its resources within a reasonable time and return to its minimal configuration."

Plasticity

$$G(G_{(0,T_{p_1}]}(L = 0) \rightarrow F_{(0,T_{p_2}]}(R = R_{min}))$$

"If the load remains zero for T_{p1} time units, the system needs to deallocate all additional resources within T_{p2} time units"

Plasticity

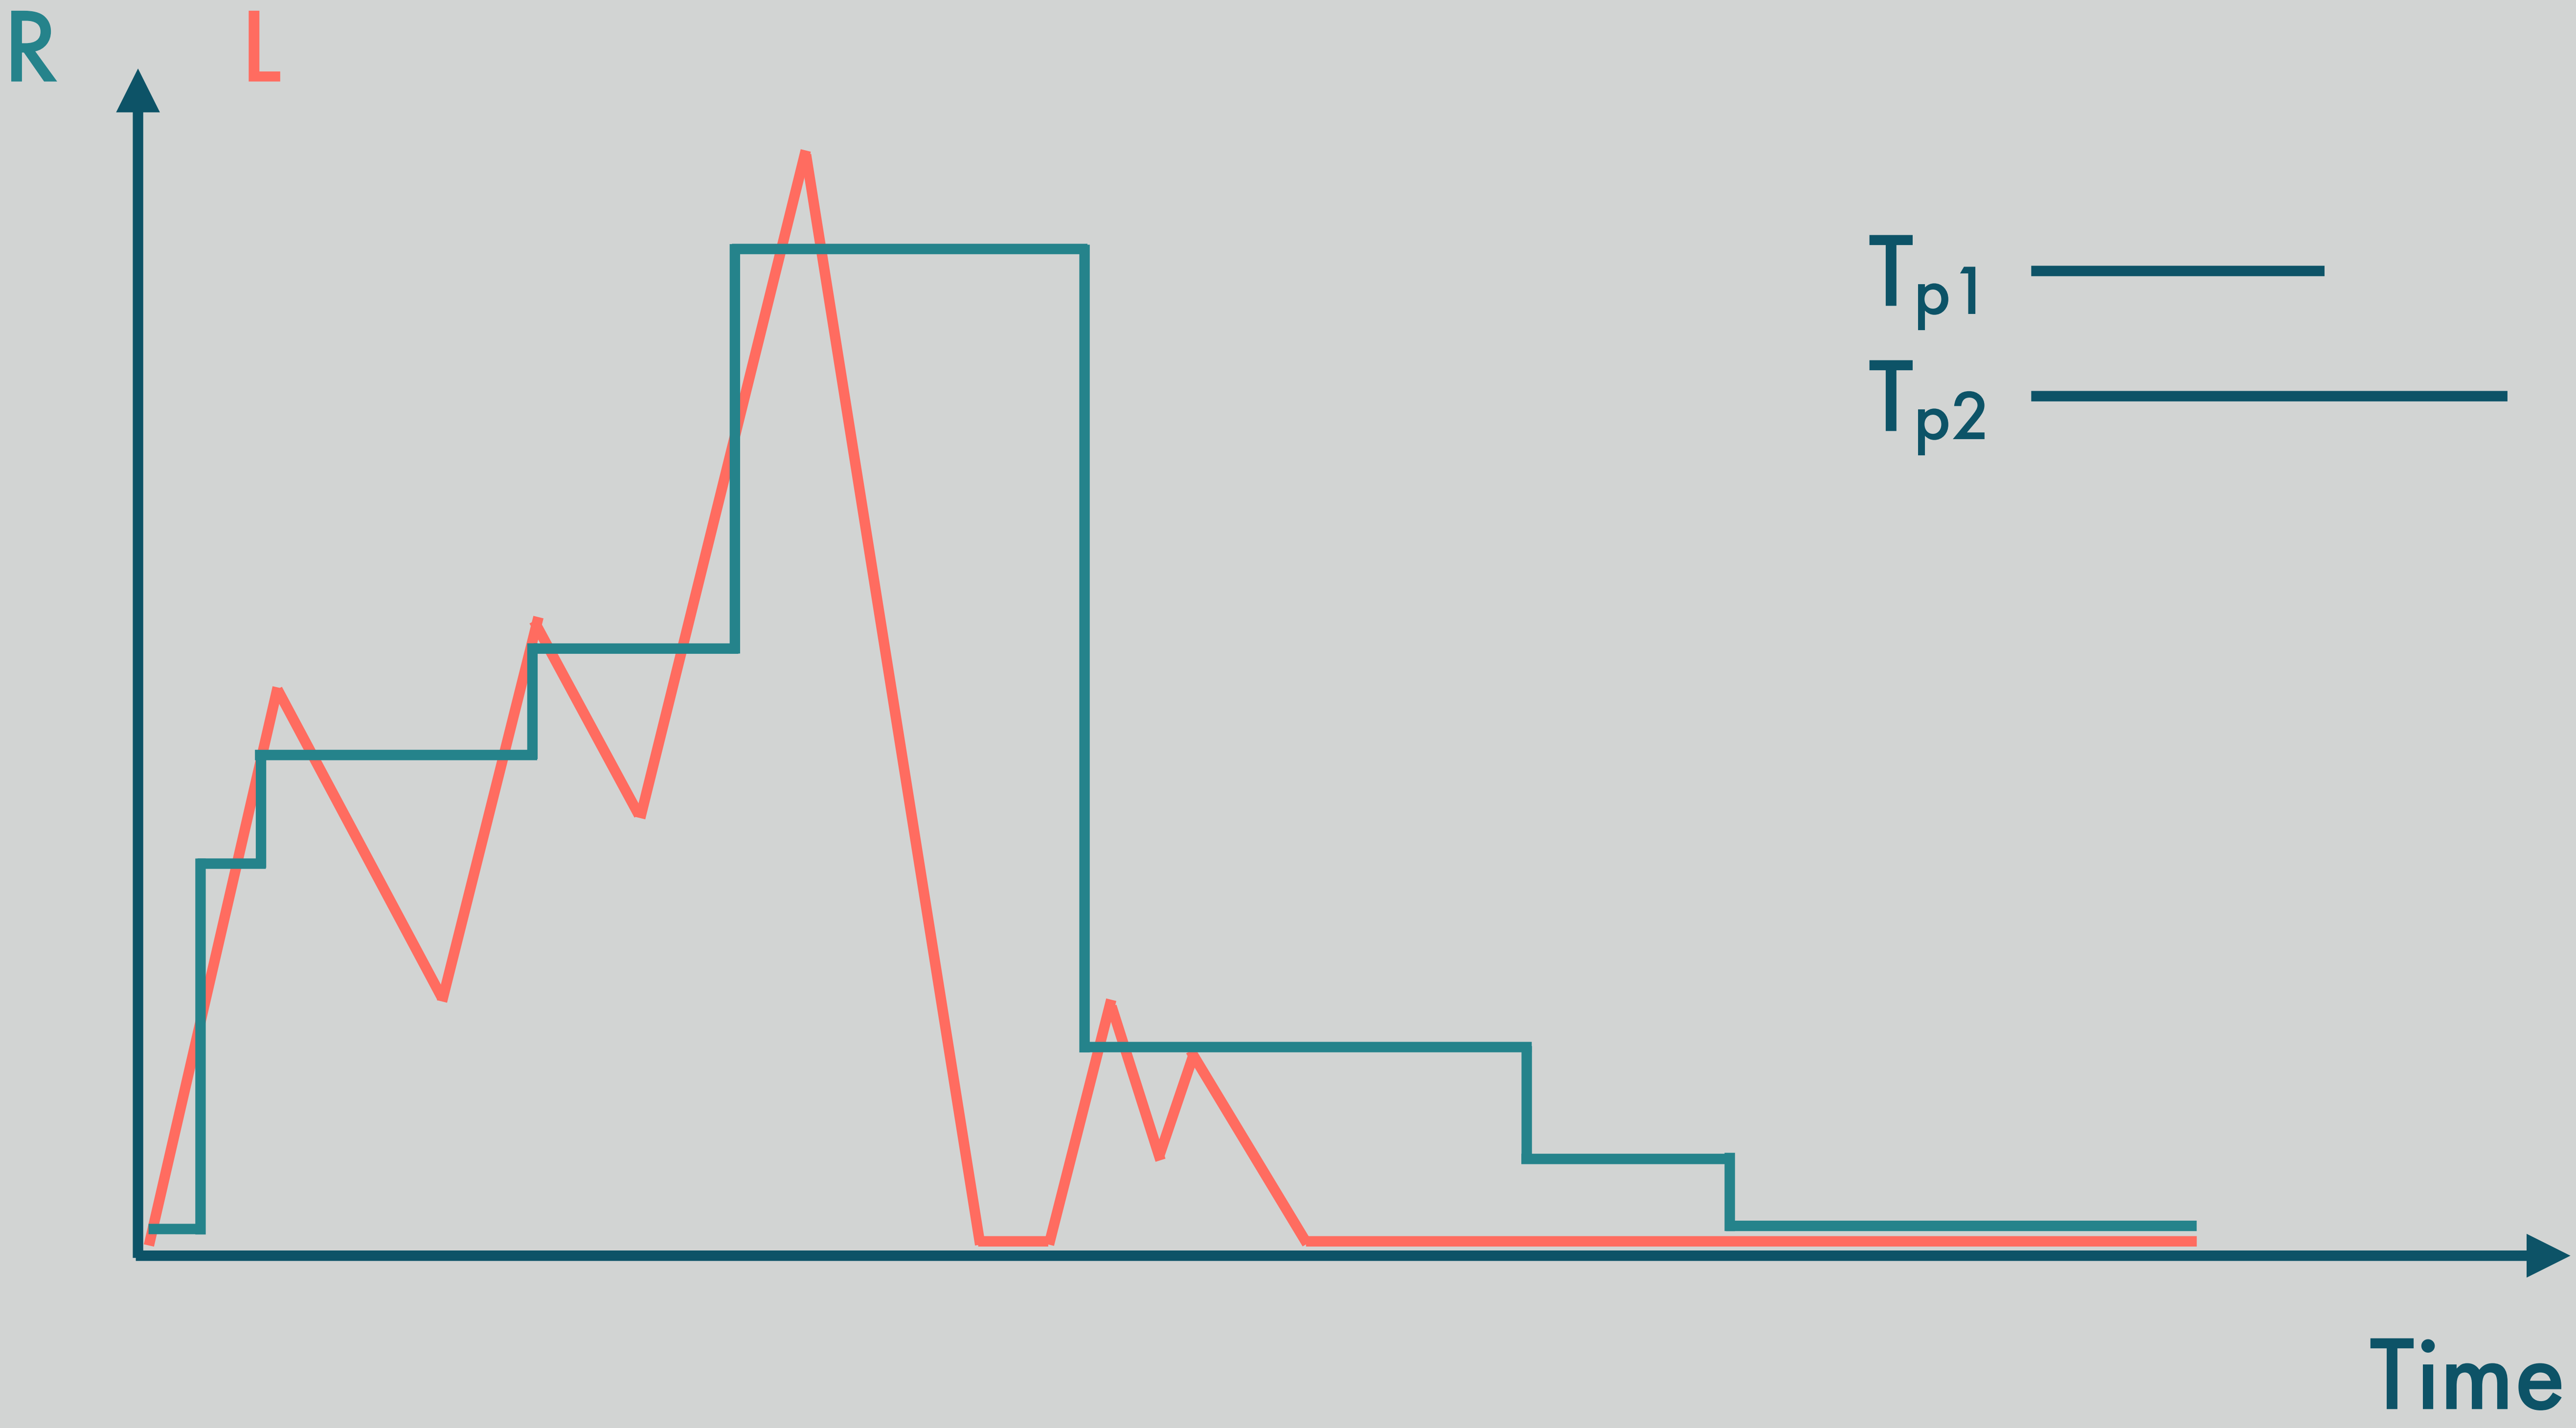

Resource Management

Precision

Oscillation

Resource Thrashing

Cool-down Period

Bounded Concurrent Adaptations

Precision

"Precision constrains the amount of resources that system is allowed to over- or under-provision."

Precision

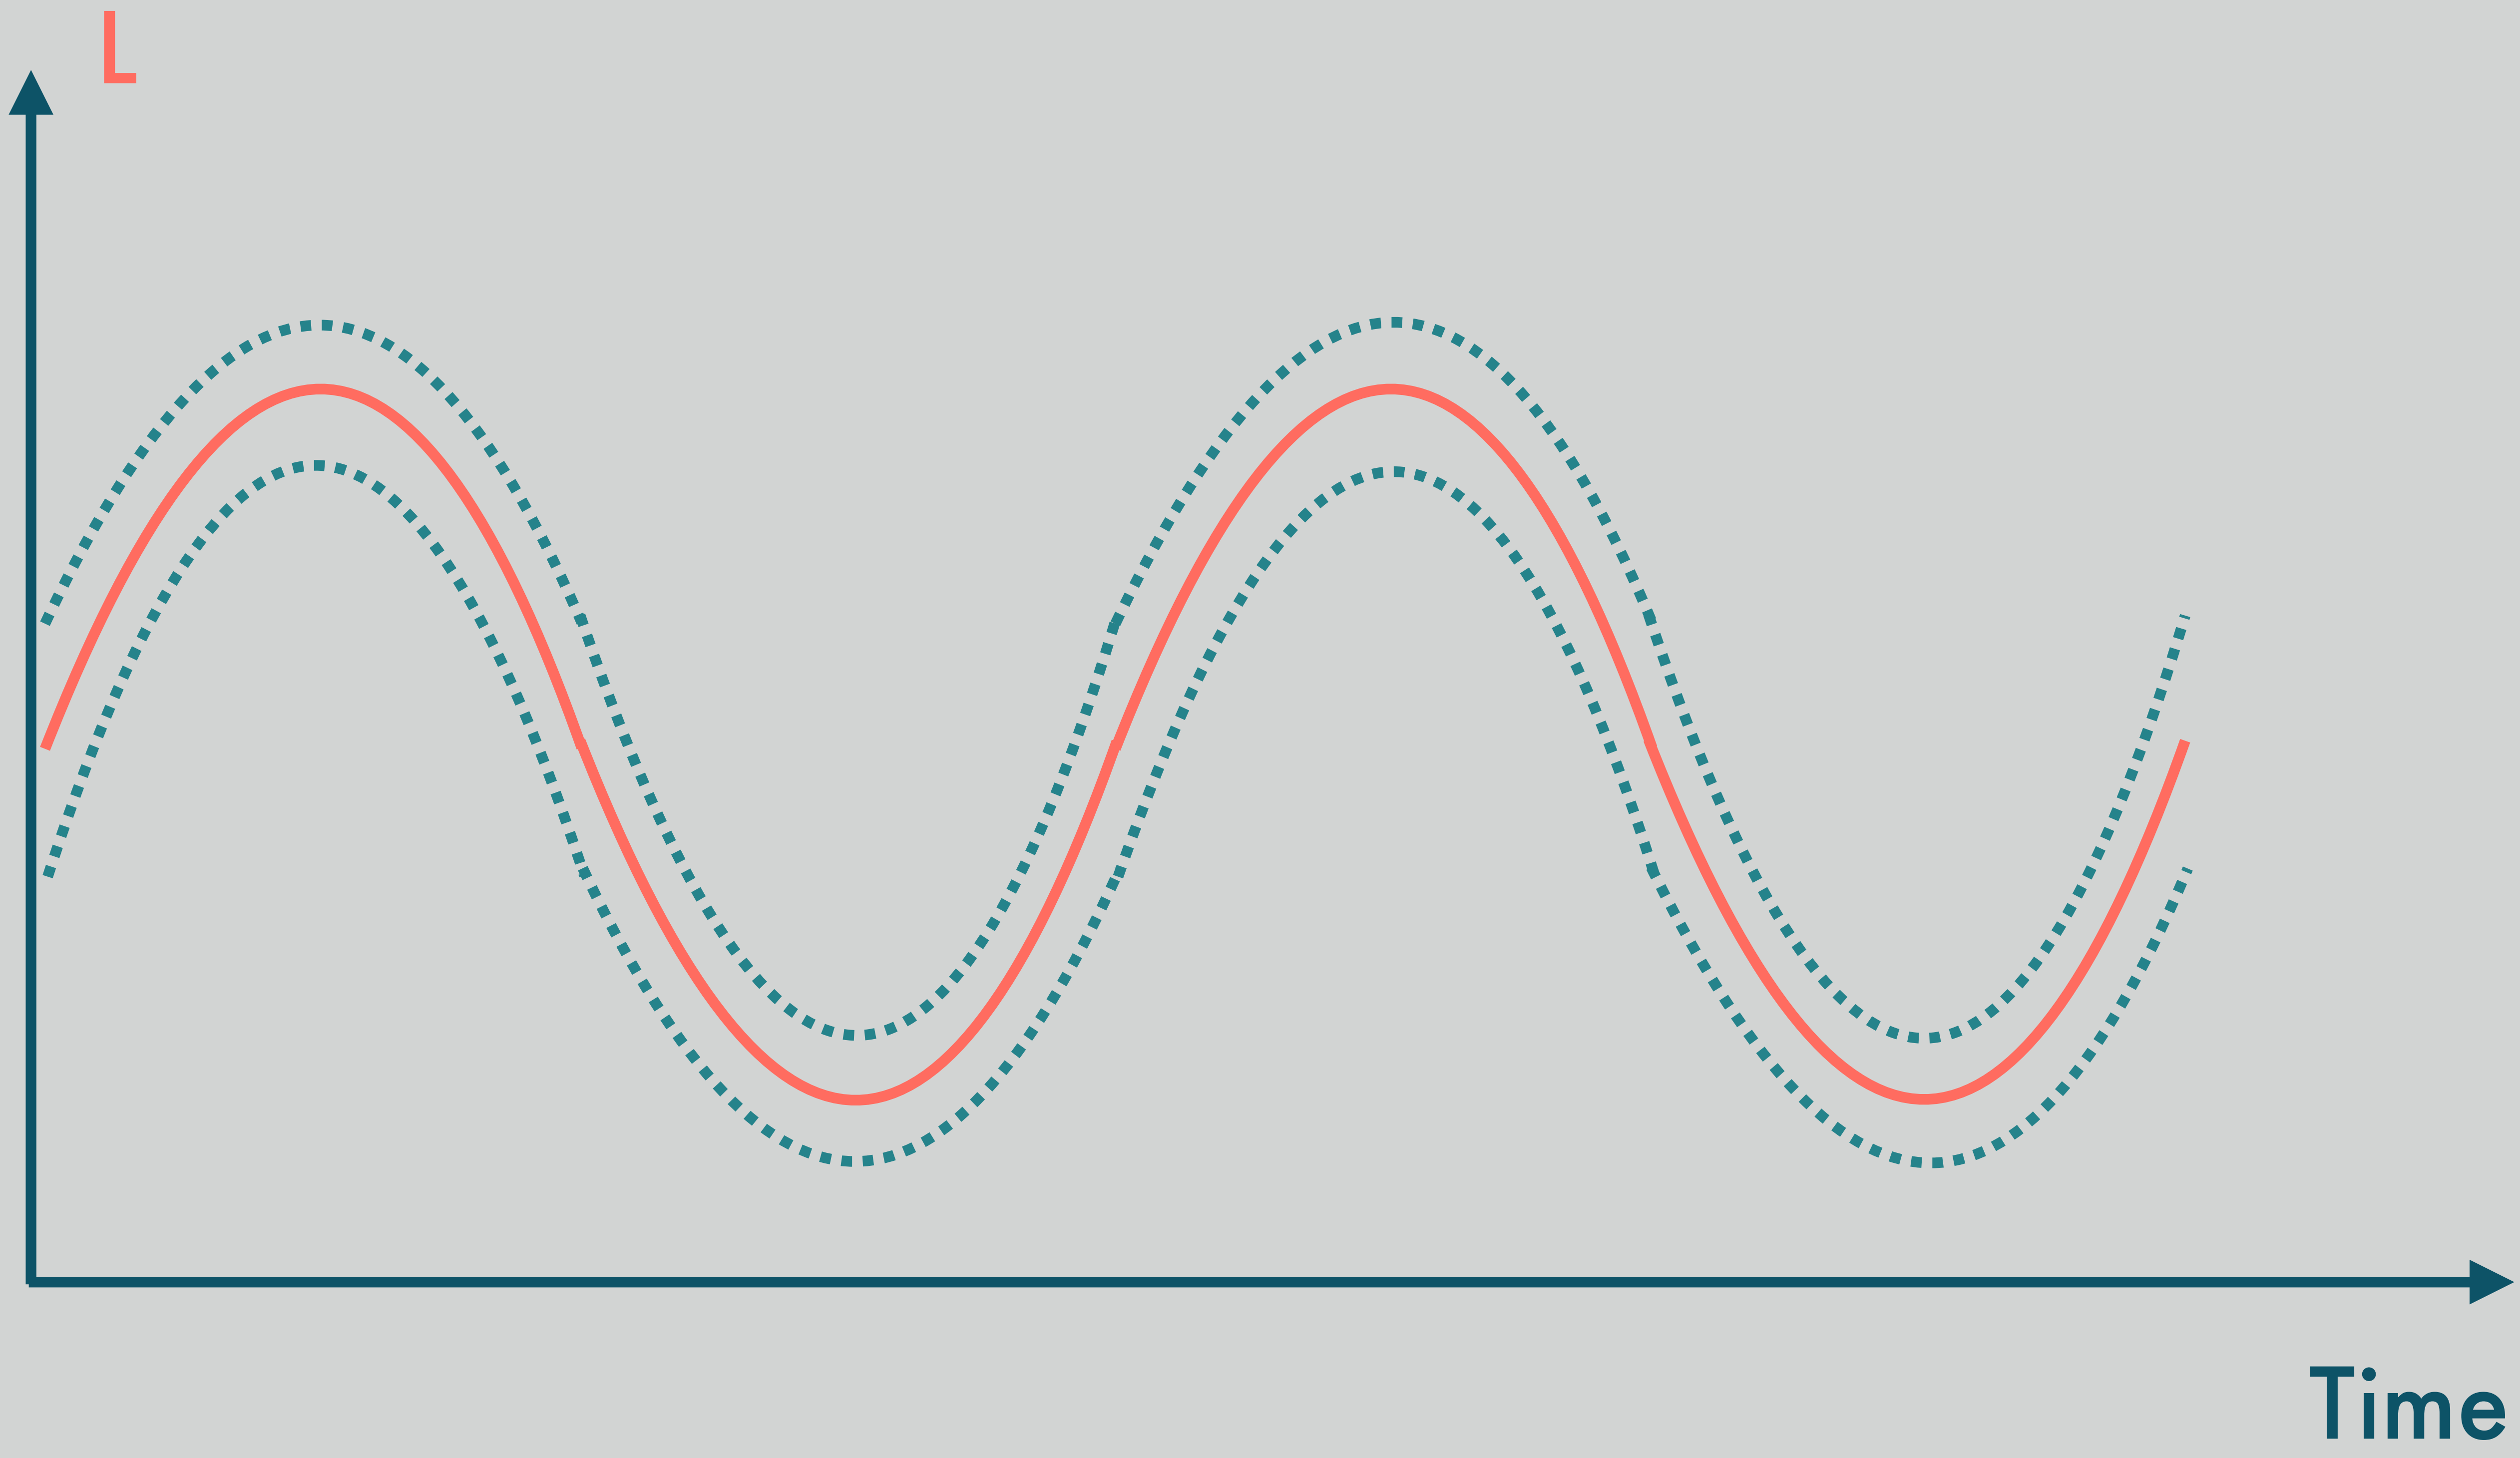

Precision

$$G(|\boldsymbol{R} - \boldsymbol{L}| < \epsilon)$$

"At any time, the absolute difference between provisioned resources and needed resources must be less than a certain threshold "

Oscillation

"Elastic system must not allocate or deallocate resources when the load is stable."

Oscillation

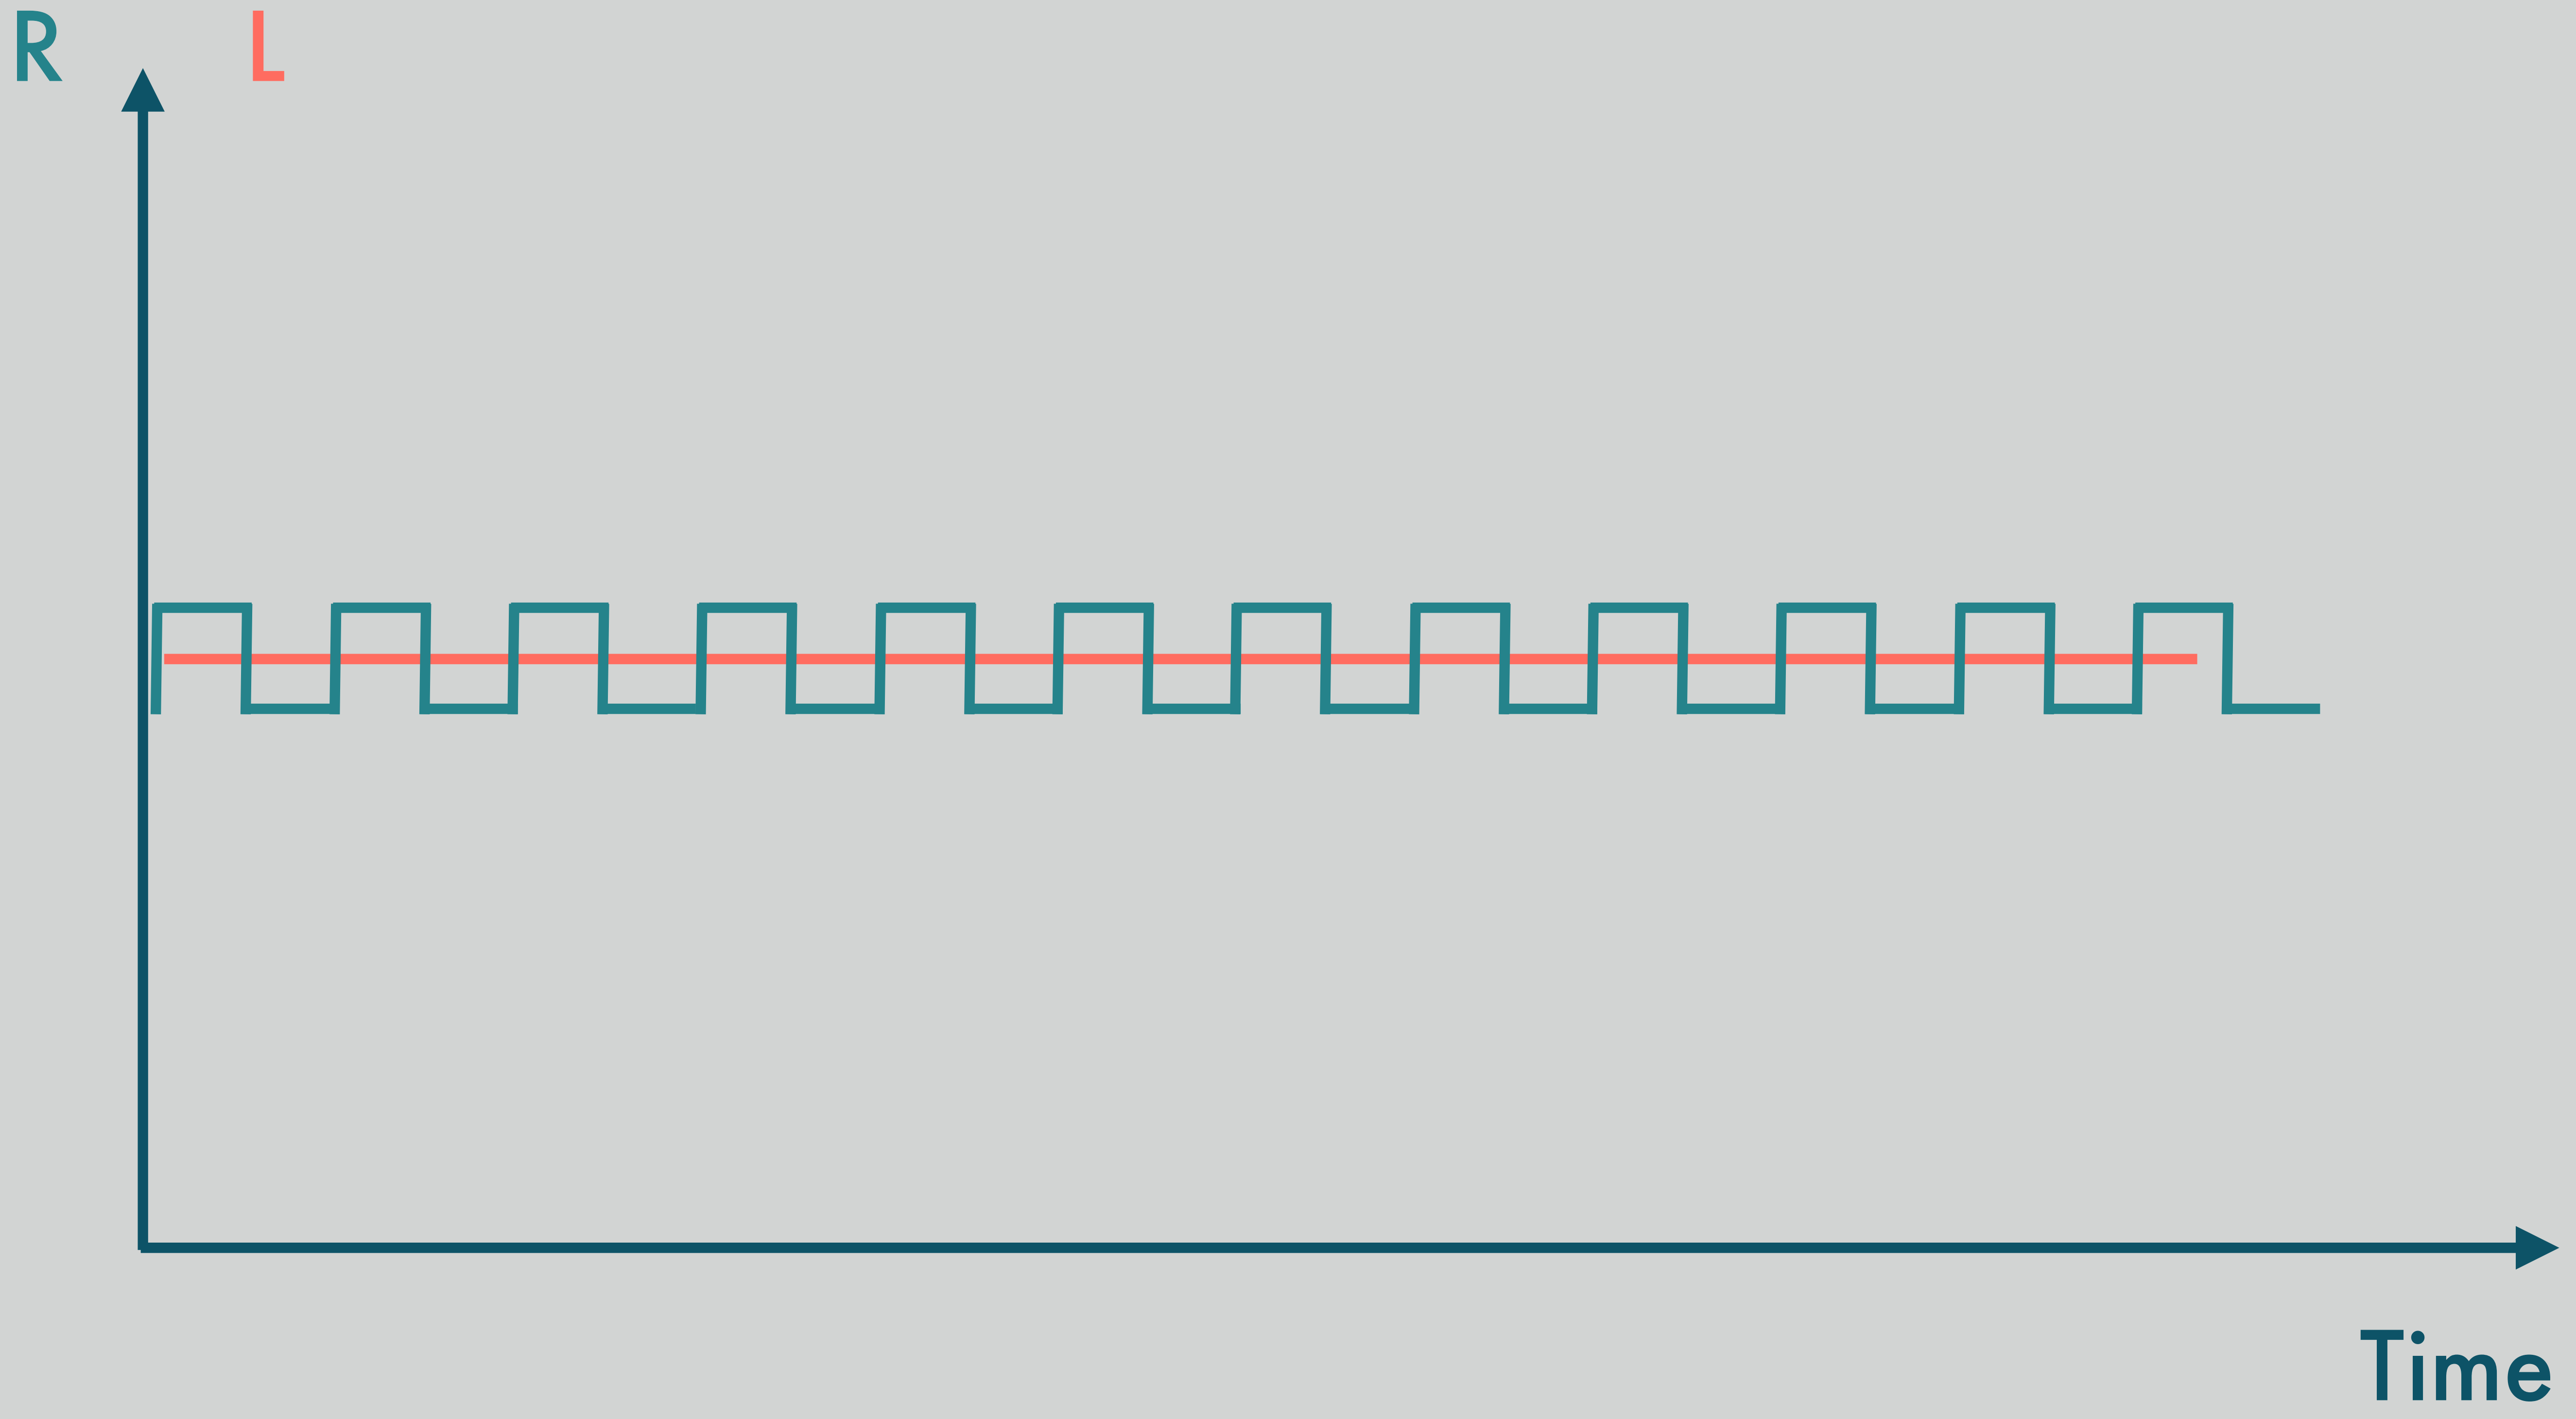

Oscillation

$$G(X(\boldsymbol{R}) > \boldsymbol{R} \rightarrow P_{(0,T_e]}(X(\boldsymbol{L}) > \boldsymbol{L}))$$

$$G(X(\boldsymbol{R}) < \boldsymbol{R} \rightarrow P_{(0,T_e]}(X(\boldsymbol{L}) < \boldsymbol{L}))$$

"The formulae constrain the increase (decrease) of the number of resources only in correspondence with an increase(decrease) of the load that happened some Te time before."

Resource trashing

"Elastic system must not deallocate resources shortly after allocating them and vice versa."

Resource trashing

$$G(R < X(R) \rightarrow \neg F_{(0,T_{rtx}]}(R > X(R)))$$

$$G(R > X(R) \rightarrow \neg F_{(0,T_{rtx}]}(R < X(R)))$$

“The occurrence of opposite adaptations can happen only after a minimum amount of time Trtx has passed"

Resource trashing

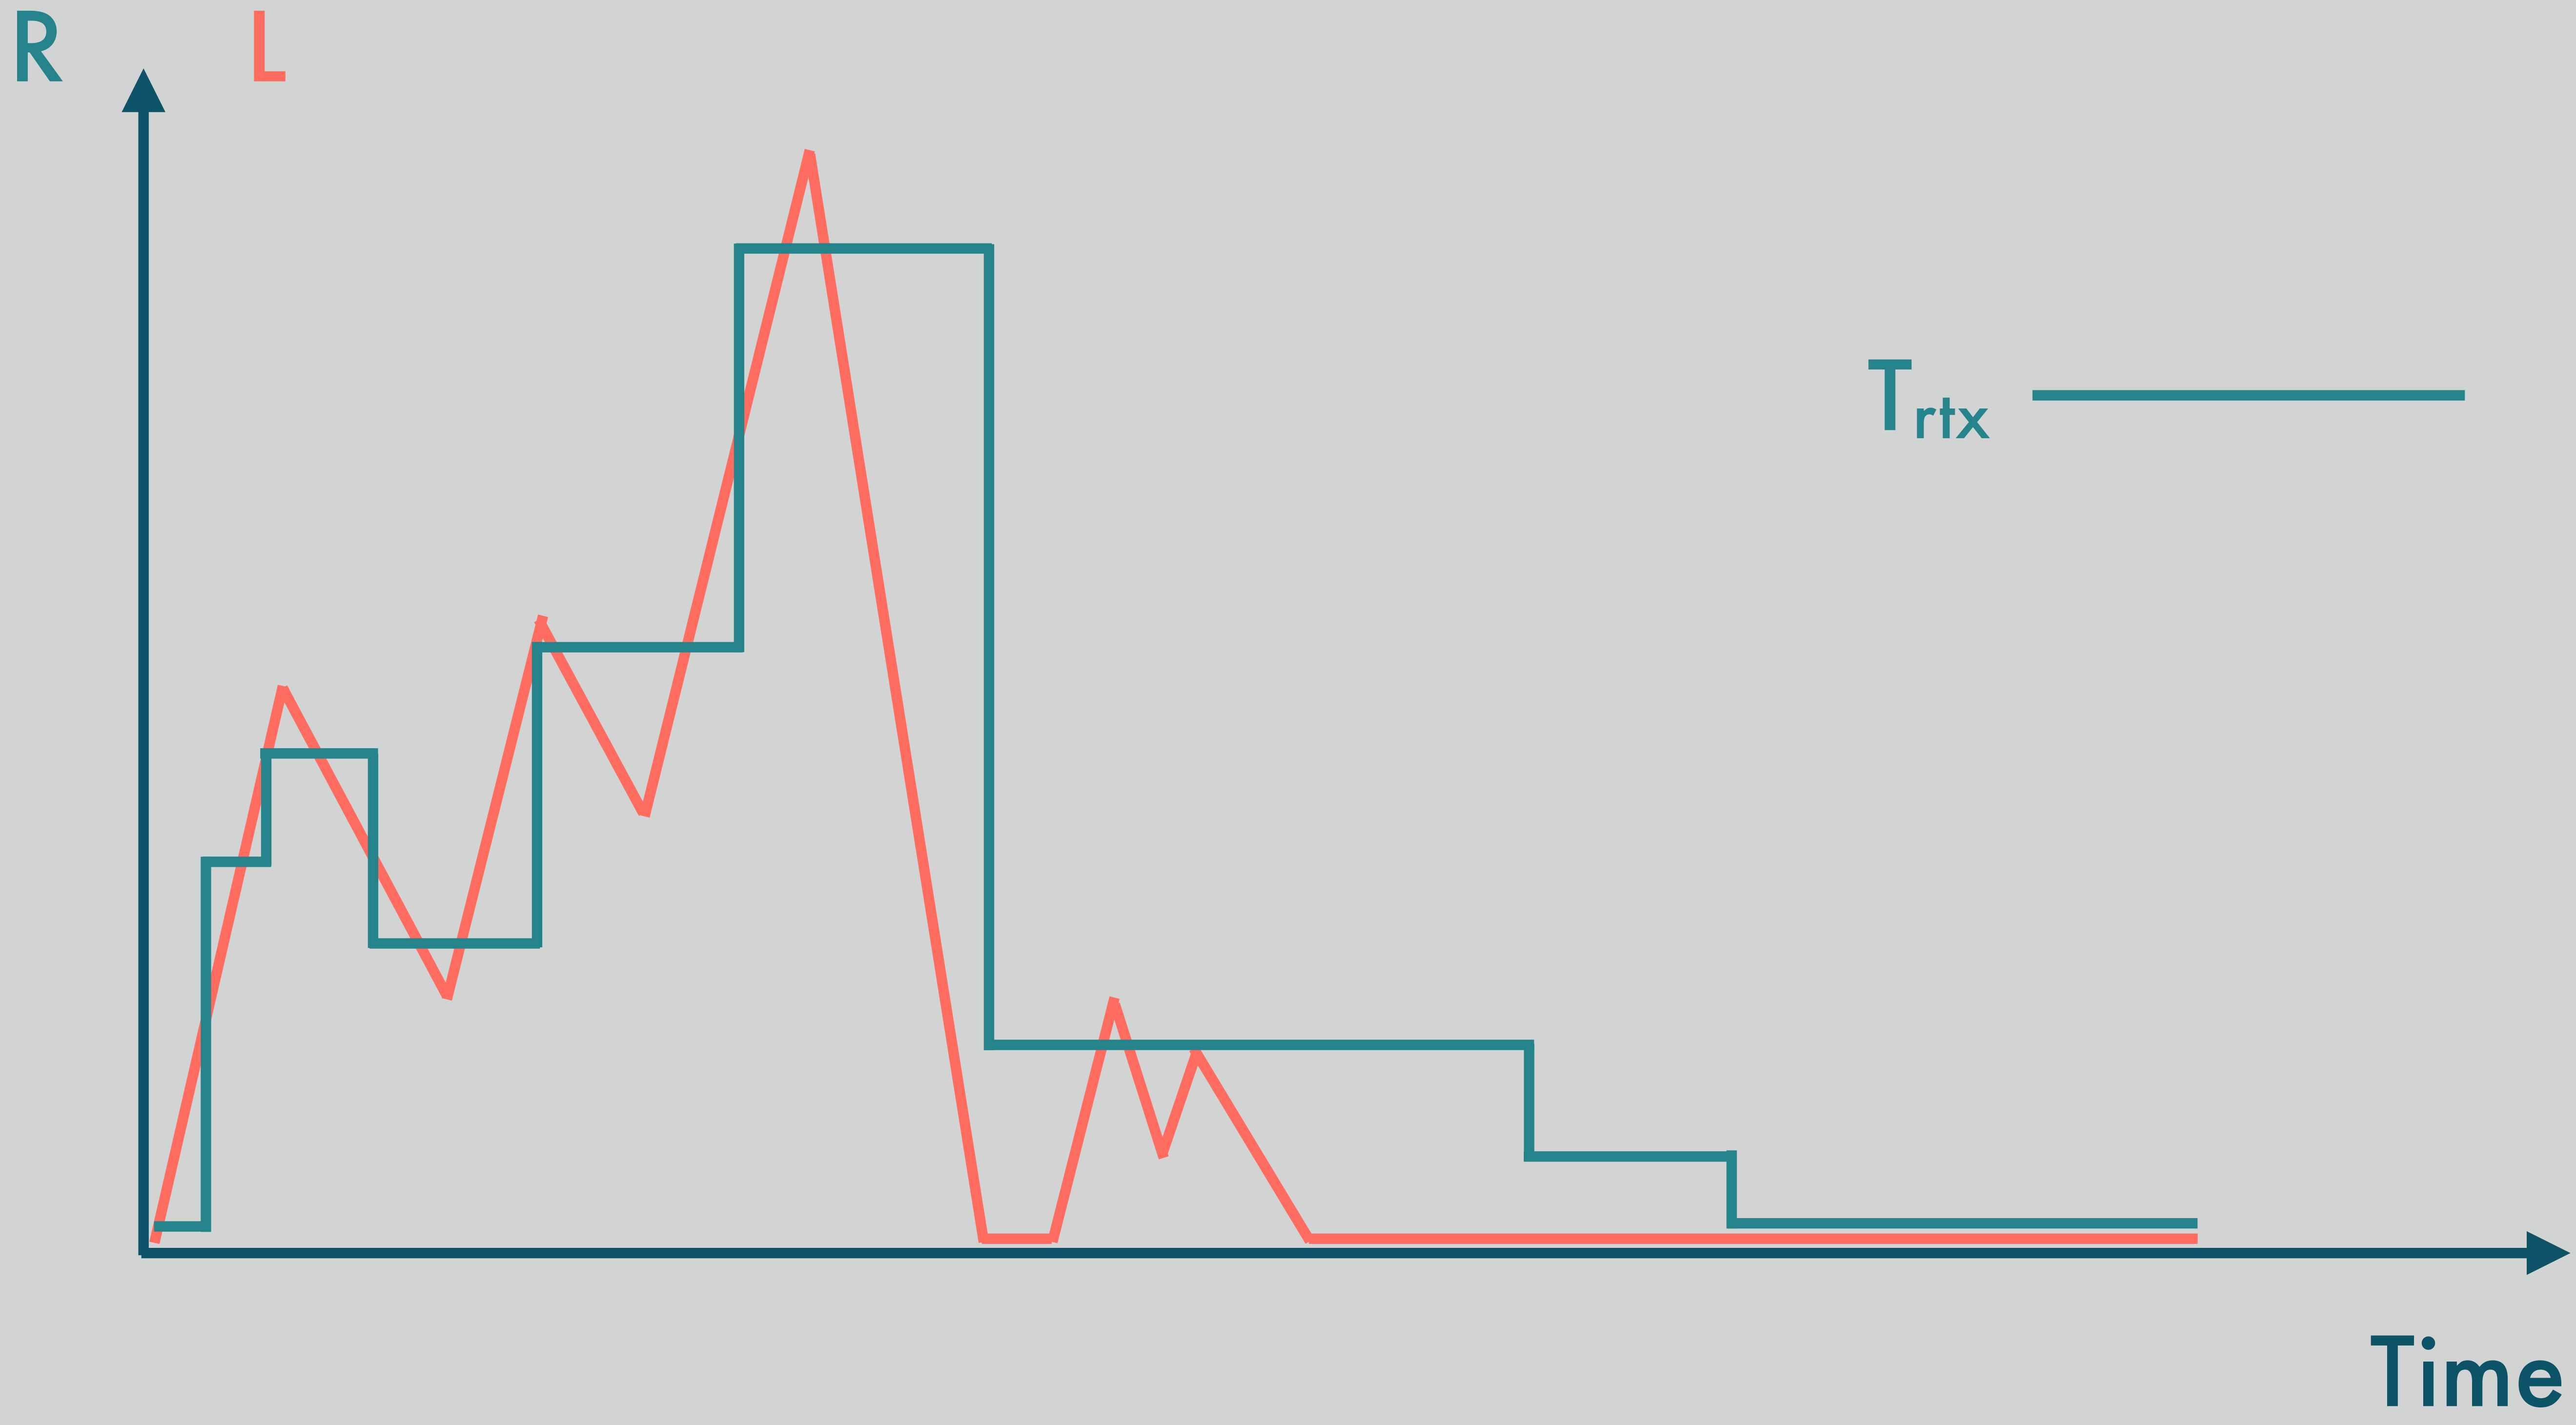

Resource trashing

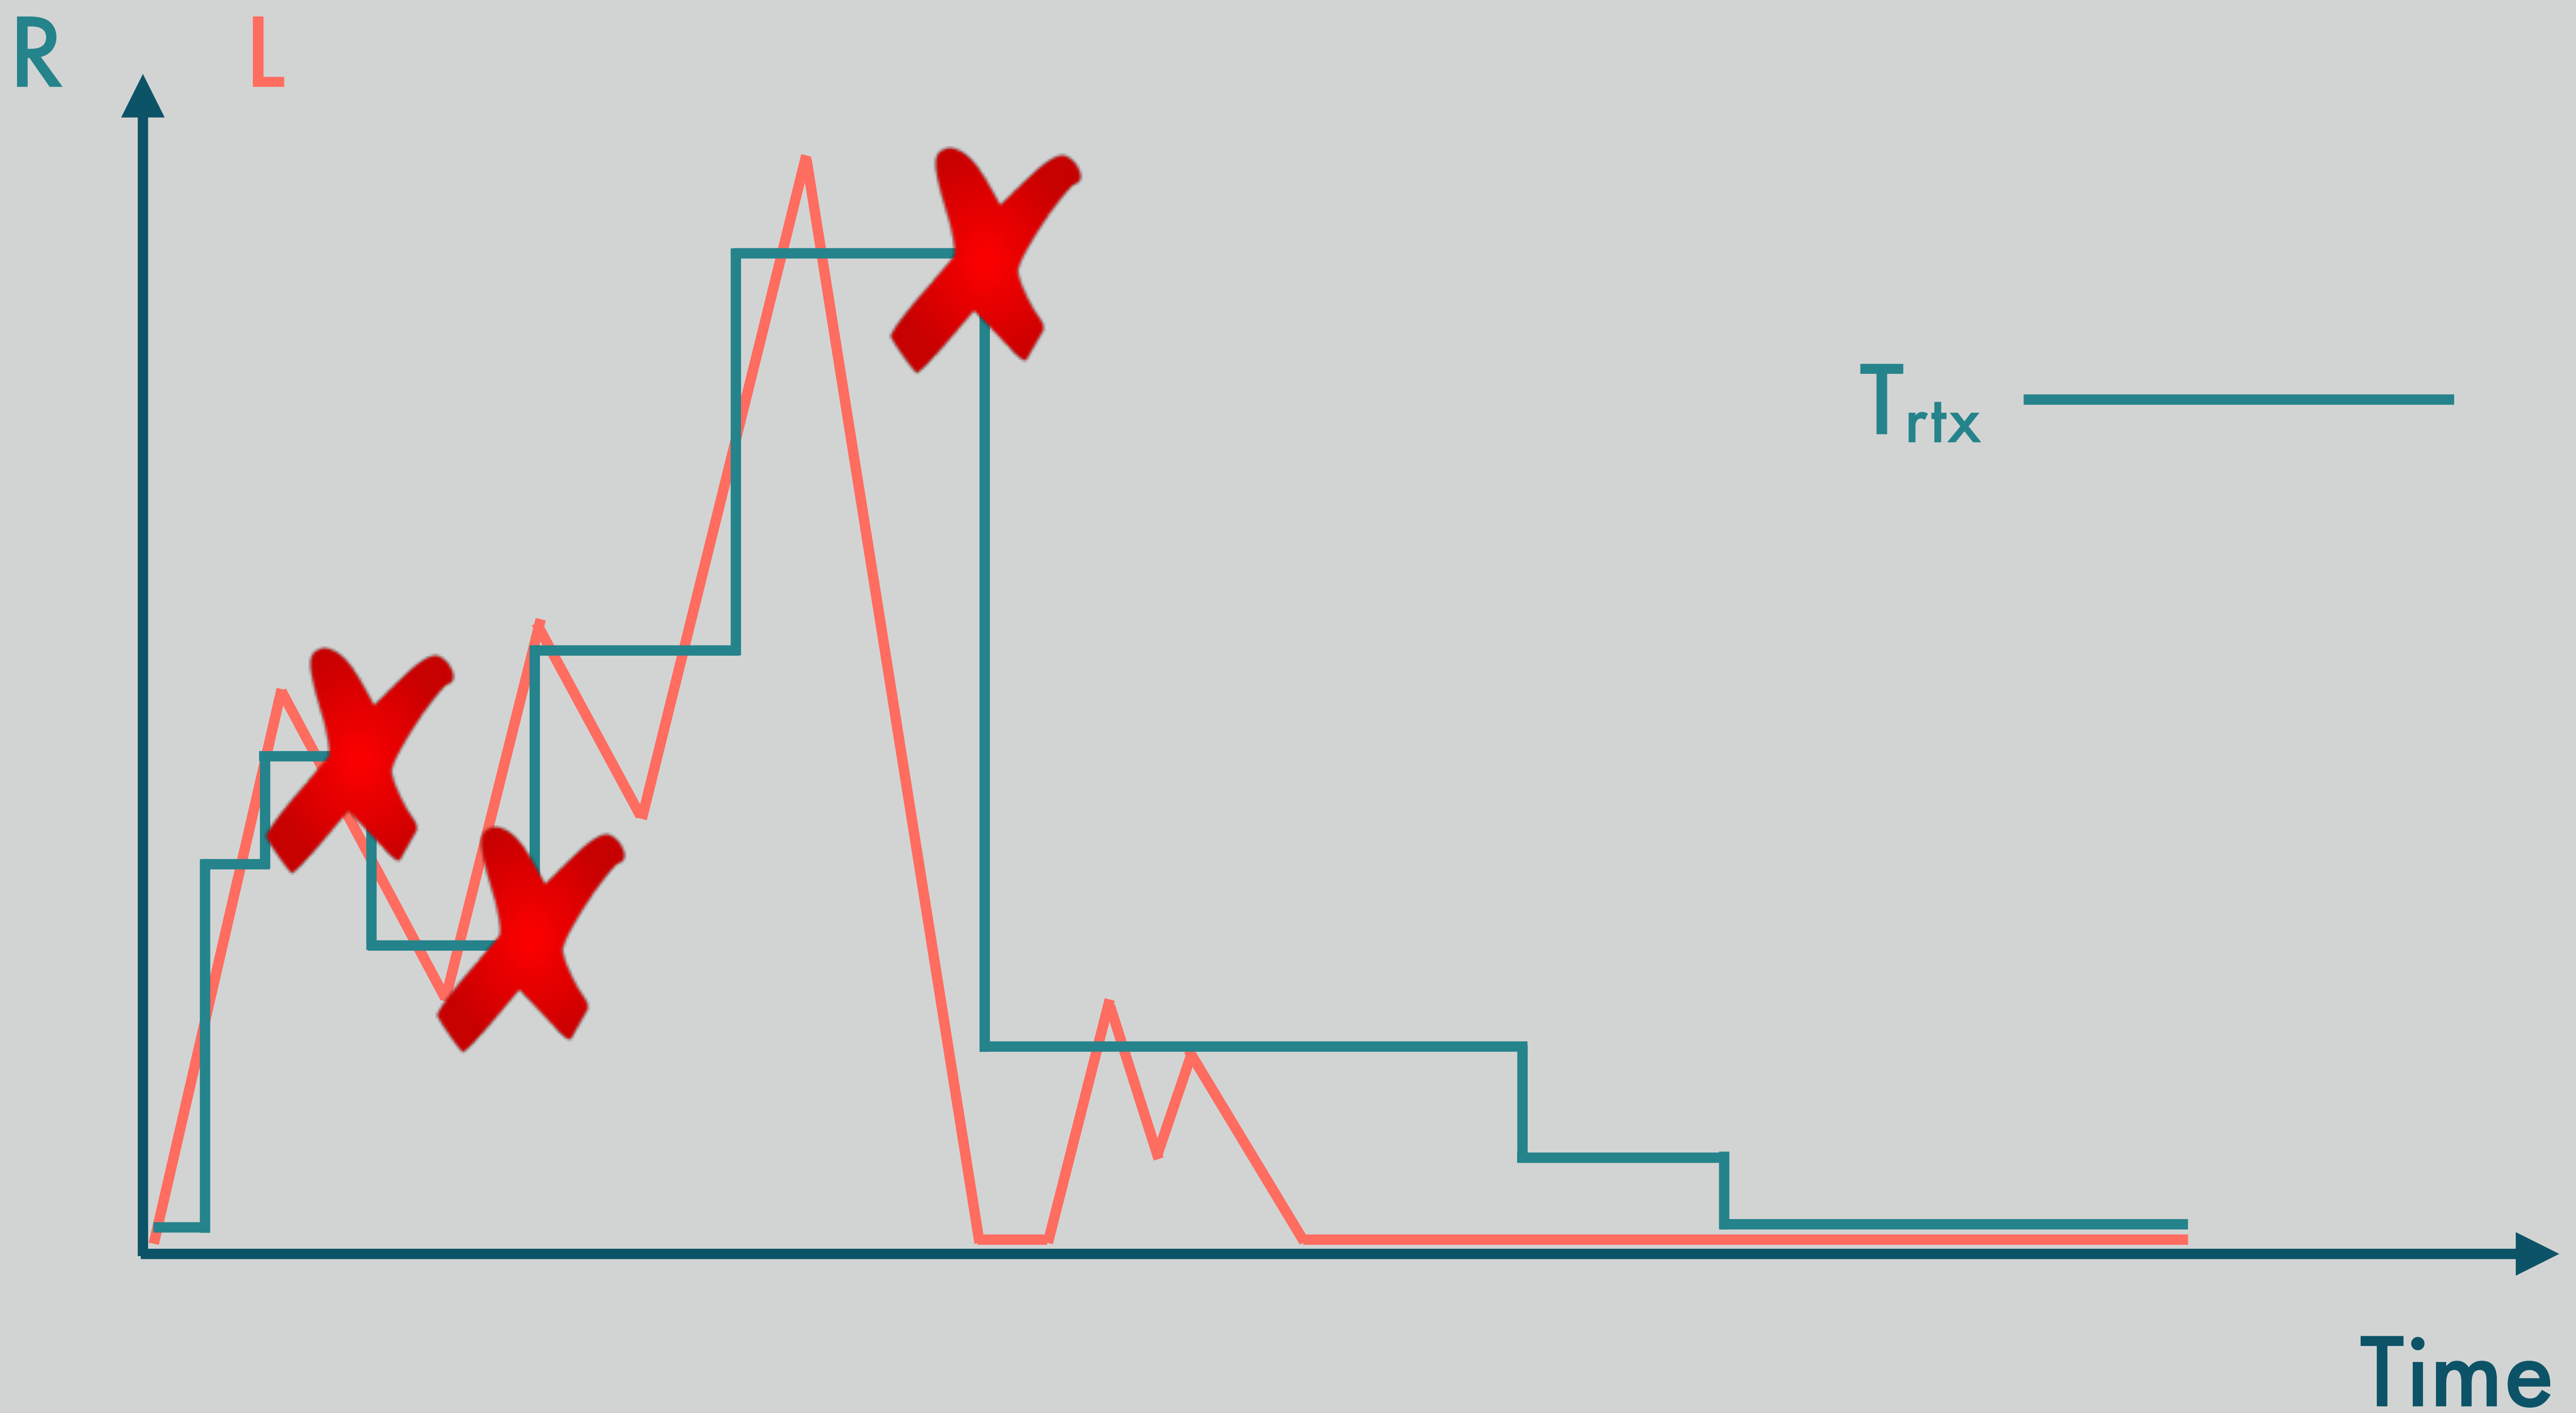

Cool-down period

"Elastic controller must not change resources during the period of VM initialization."

Bounded Concurrent Adaptations

"Elastic controller must not perform more than N changes during the period of VM initialization."

Quality of Service

Bounded QoS Degradation

Bounded Actuation Delay

Bounded QoS Degradation

"During the adaptation elastic systems may relax the QoS requirements up to a certain value."

Bounded QoS Degradation

$$\mathsf{G}(A \rightarrow \textit{some-lower-level-of-QoS})$$

$$\mathsf{G}(\neg A \rightarrow \textit{standard-level-of-QoS})$$

Bounded Actuation Delay

"It expresses a bound on the actuation time of the controller, i.e., time it takes to provision/ deprovision a VM."

Bounded Actuation Delay

$$\forall id : G(M_{boot}(id) \rightarrow F_{(0,T_{ad})}(M_{ready}(id)))$$

$$\forall id : G(M_{stop}(id) \rightarrow F_{(0,T_{ad})}(M_{end}(id)))$$

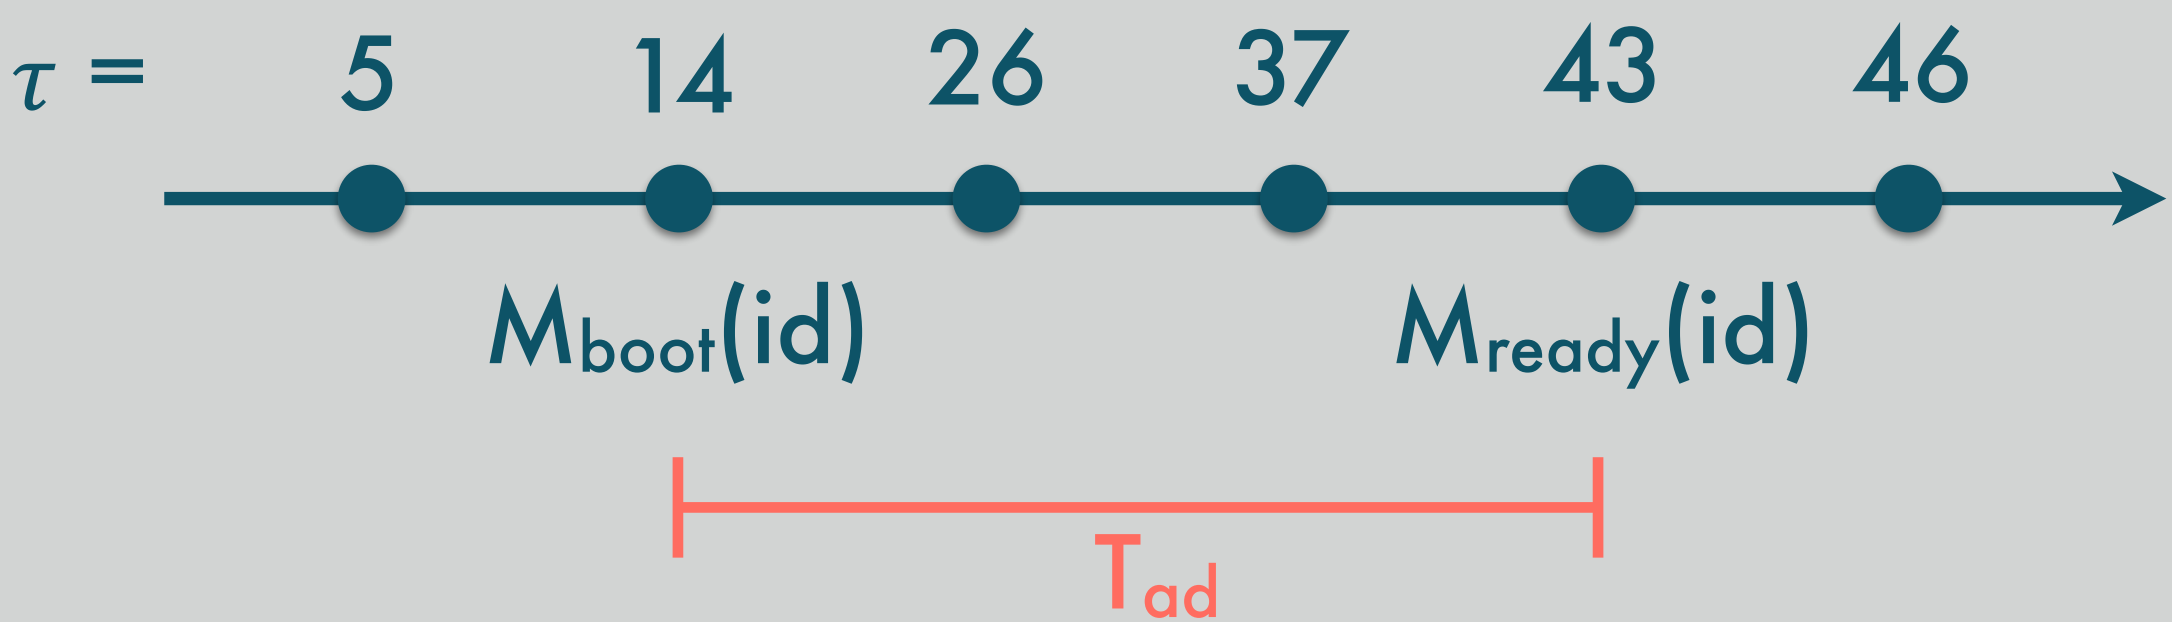

Open issues

Specification

Verification

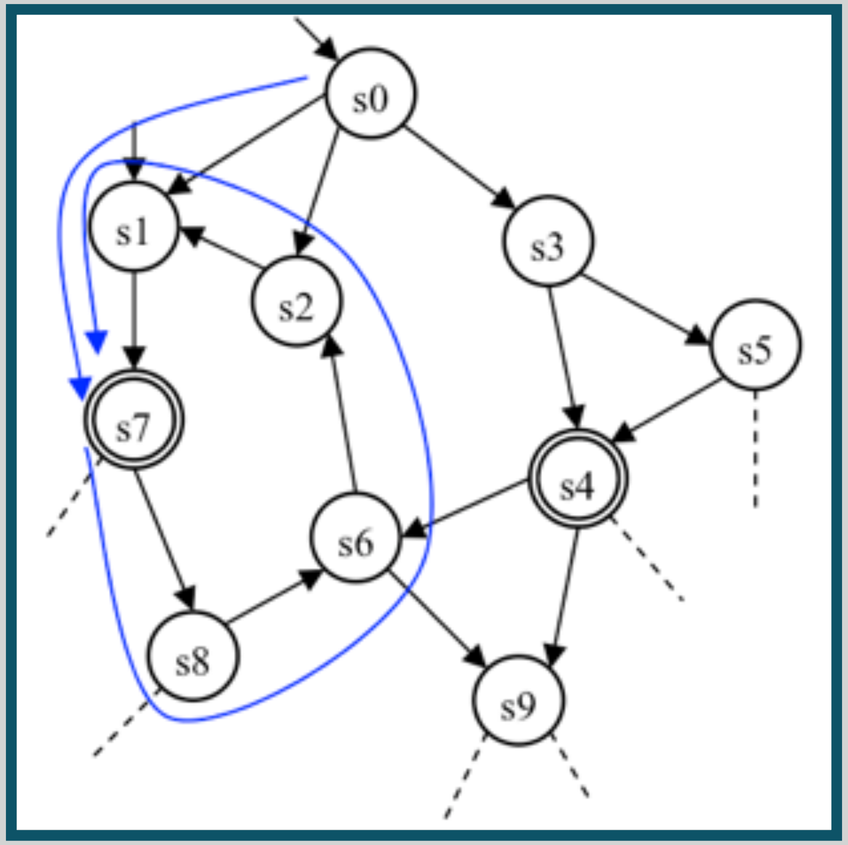

Failure Analysis

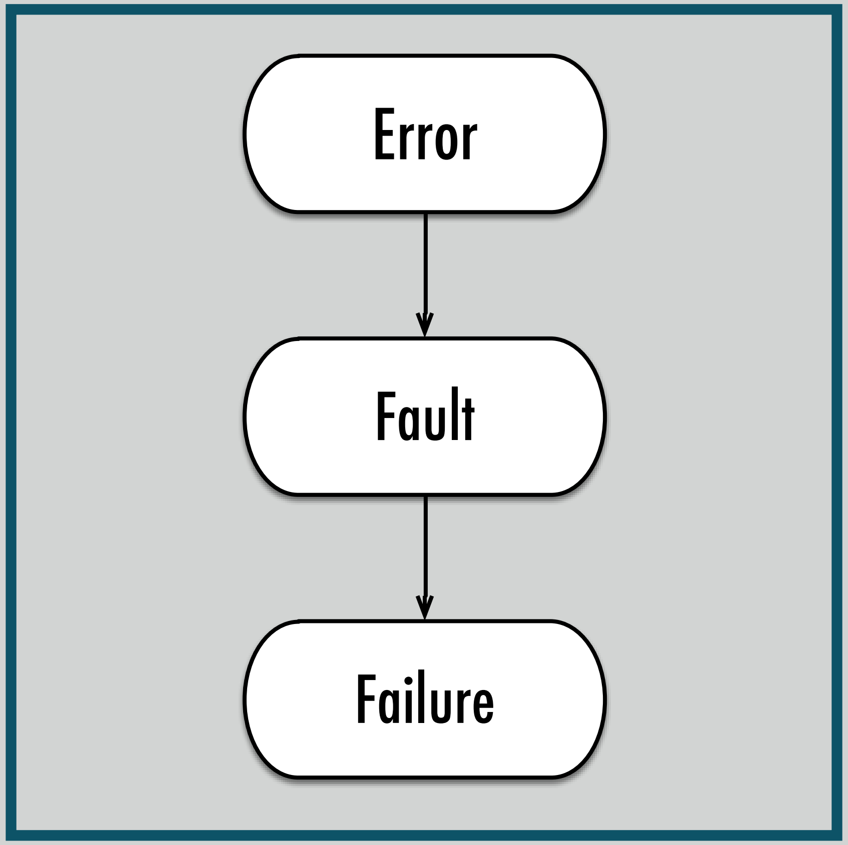

Open issues

Specification

Verification

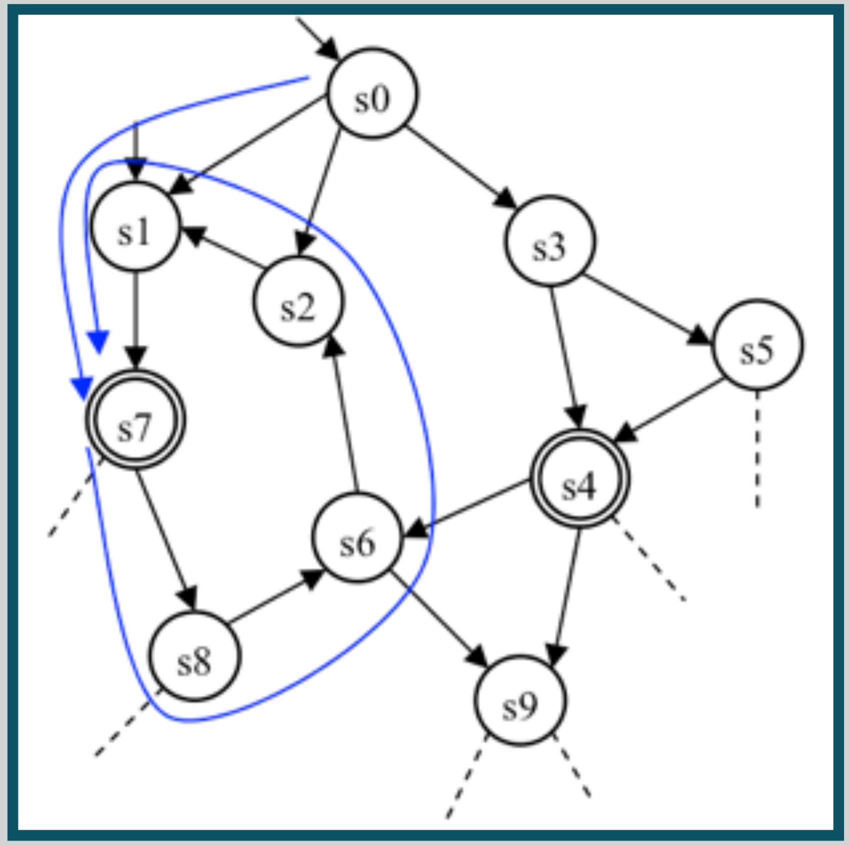

Failure Analysis

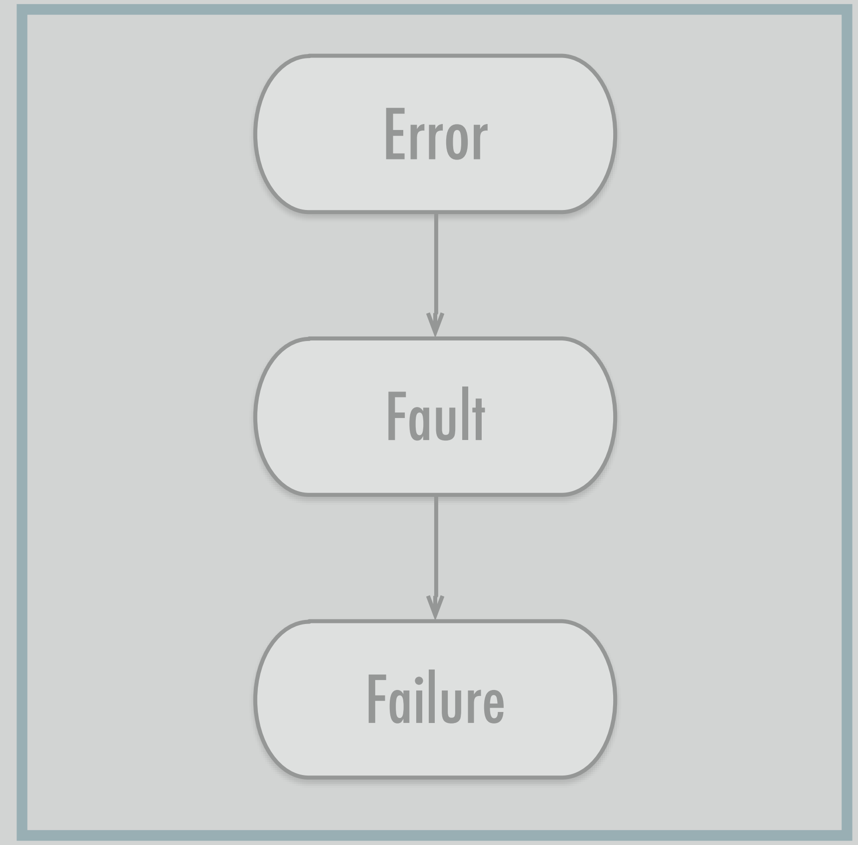

Preliminary Evaluation

- "Elastic Doodle" Service
- Private OpenStack infrastructure
- Input workload:

“Elastic Doodle”

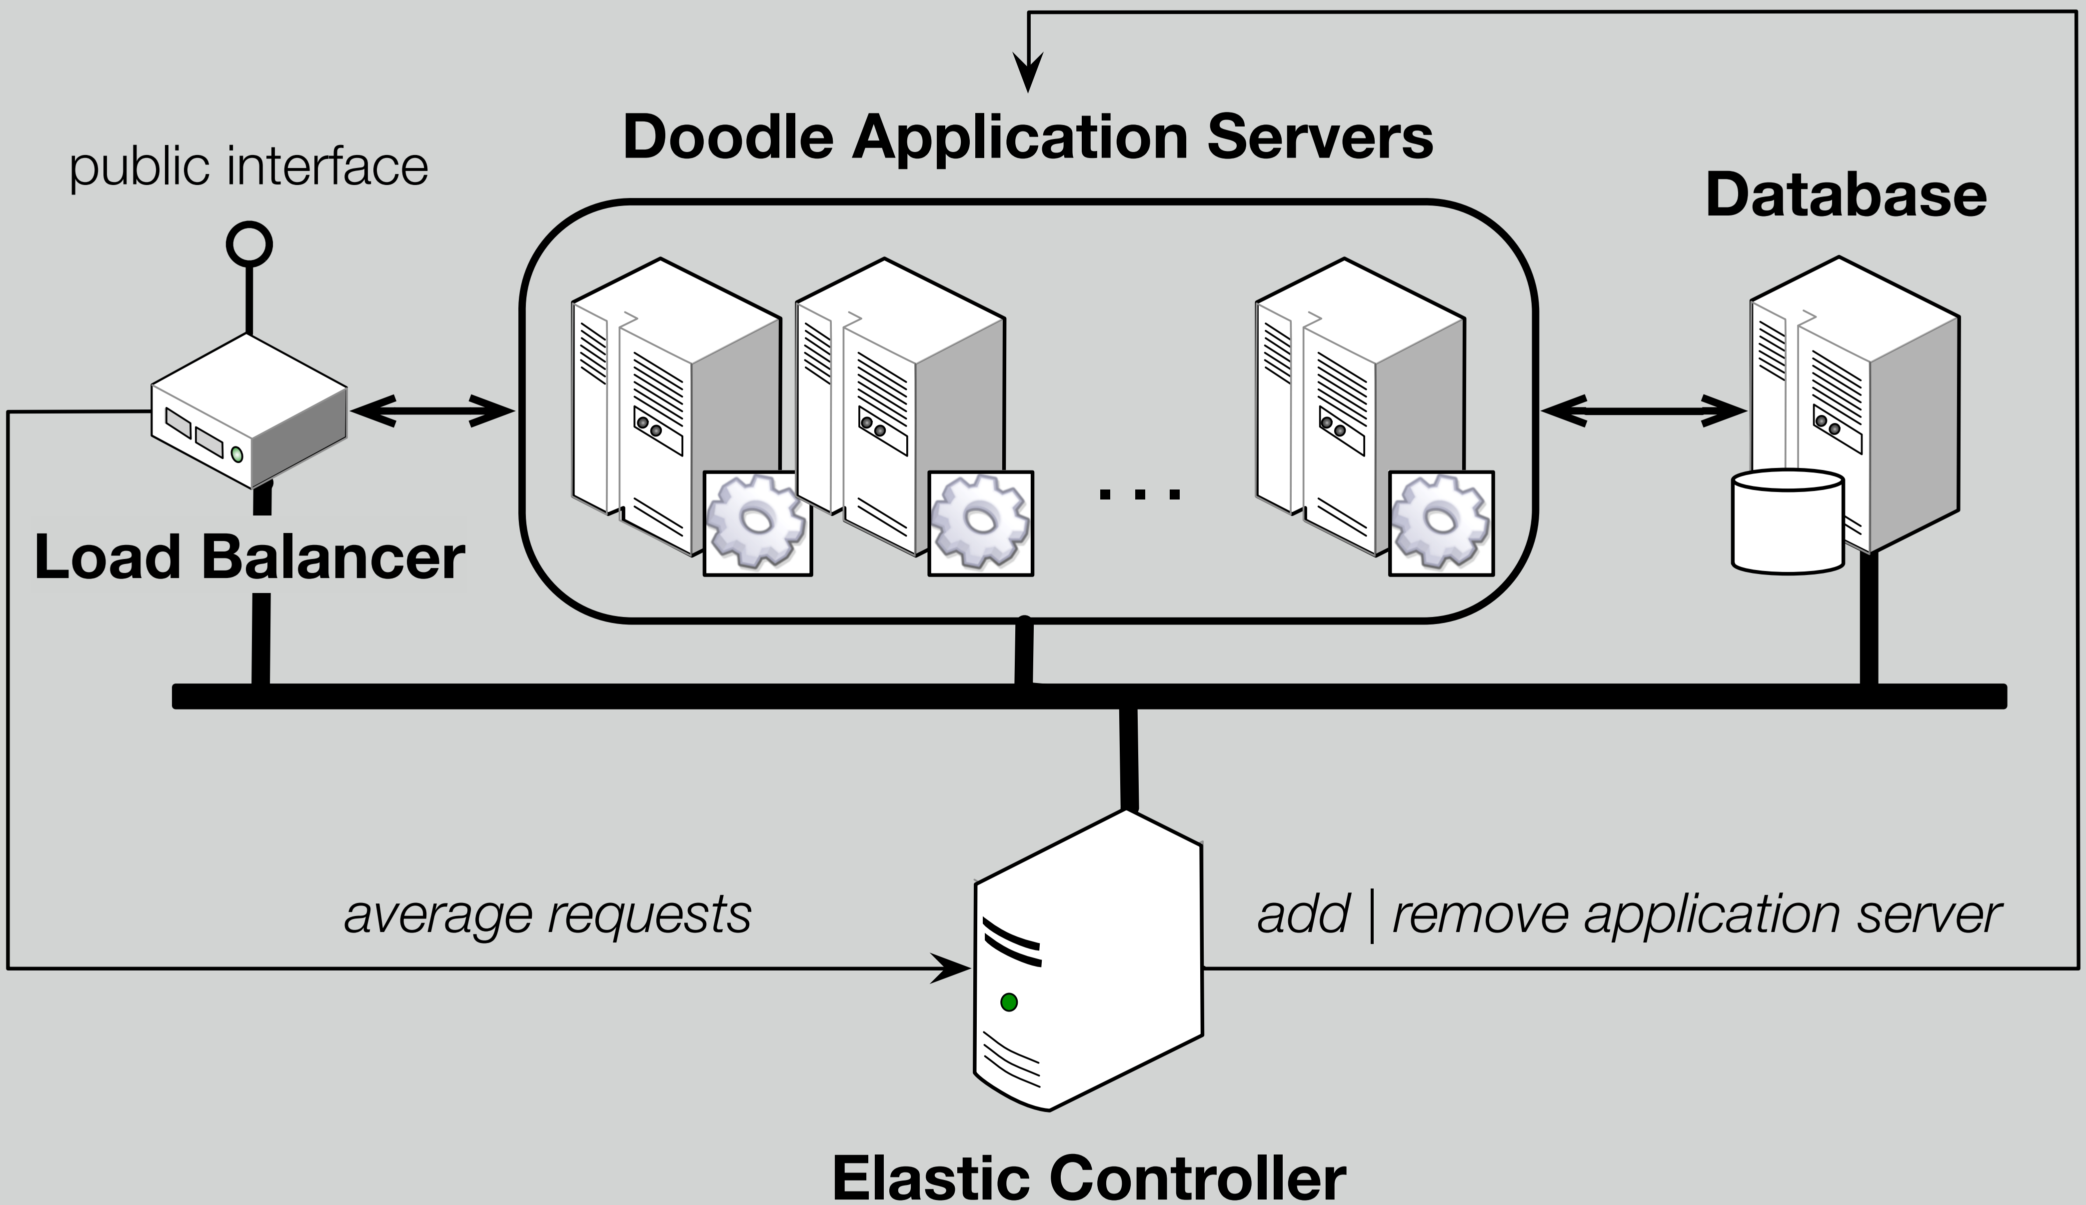

Elastic Doodle Controller

Reads periodically (10 sec) the monitored data applies the following rule-based approach:

- scale-up: if the average number of requests per running application server in the last minute is over a certain maximum threshold, a new instance of application server is allocated; the controller stops its execution for one minute;

- scale-down: if the average number of requests per running application server in the last minute is below a certain minimum threshold, a running instance of application server is deallocated; the controller stops its execution for two minutes.

Properties verified with trace checking

- Resource Thrashing
- Plasticity
- Cool-down Period

Trace checking

Traces				Properties		
ID	Events	Time span (s)	R_{max}	Resource Thrashing	Plasticity	Cool-down Period
T1	15	1102	2	1.44s/120MB	1.20s/117MB	2.29s/126MB
T2	43	635	4	2.83s/135MB	1.47s/122MB	1.42s/121MB
T3	29	641	3	1.77s/131MB	1.21s/118MB	1.62s/126MB
T4	17	499	3	1.20s/117MB	1.27s/116MB	1.38s/116MB
T5	44	644	3	1.94s/135MB	1.45s/122MB	1.45s/122MB

Possible Research Directions

- Refinement of the load model
- Evaluation on industrial-strength case studies
- Run-time monitoring

PESOS 2014

Towards the Formalization of Properties of Cloud-Based Elastic Systems

Marcello M. Bersani
Politecnico di Milano
Milano, Italy
bersani@elet.polimi.it

Domenico Bianculli
University of Luxembourg
Luxembourg, Luxembourg
domenico.bianculli@uni.lu

Schahram Dustdar
TU Wien
Vienna, Austria
dustdar@infosys.tuwien.ac.at

Alessio Gambi
University of Lugano
Lugano, Switzerland
alessio.gambi@usi.ch

Carlo Ghezzi
Politecnico di Milano
Milano, Italy
carlo.ghezzi@polimi.it

Srđan Krstić
Politecnico di Milano
Milano, Italy
srdan.krstic@polimi.it

ABSTRACT

Cloud-based elastic systems run on a cloud infrastructure and have the capability of dynamically adjusting the allocation of their resources in response to changes in the workload, in a way that balances the trade-off between the desired quality-of-service and the operational costs. The actual elastic behavior of these systems is determined by a combination of factors, including the input workload, the logic of the elastic controller determining the type of resource adjustment, and the underlying technological platform implementing the cloud infrastructure. All these factors have to be taken into account to express the desired elastic behavior of a system, as well as to verify whether the system manifests or not such ... first step into these directions, by ... the CLTL^t($\mathcal{D}$) temporal ... to the be-

1. INTRODUCTION

Cloud computing has become a practical solution to manage and leverage IT resources and services. Cloud platforms offer several benefits, among which the ability to access resources or service applications offered as (remote) services, available on-demand and on-the-fly, and billed according to a pay-per-use model.

Cloud providers offer resources and services at three different layers: at the *Software-as-a-Service (SaaS)* layer, users can remotely access full-fledged software applications; at the *Platform-as-a Service (PaaS)* layer, one finds a development platform, a deployment and a run-time execution environment, which is used to run user-provided code in sandboxes hosted on cloud-based premises; at the *Infrastructure-as-a-Service (IaaS)* the user can access computing resources such as virtual machines, block storage, firewalls, load balancers, or networking I/O.

In this paper, we focus on the IaaS layer, and assume, without loss of generality, that resources offered at this level ... machines. In particular, we consider *cloud-based* ... [7] of computing systems is de- ... Institute of Standards and

Specification of Cloud-Based Elastic Systems

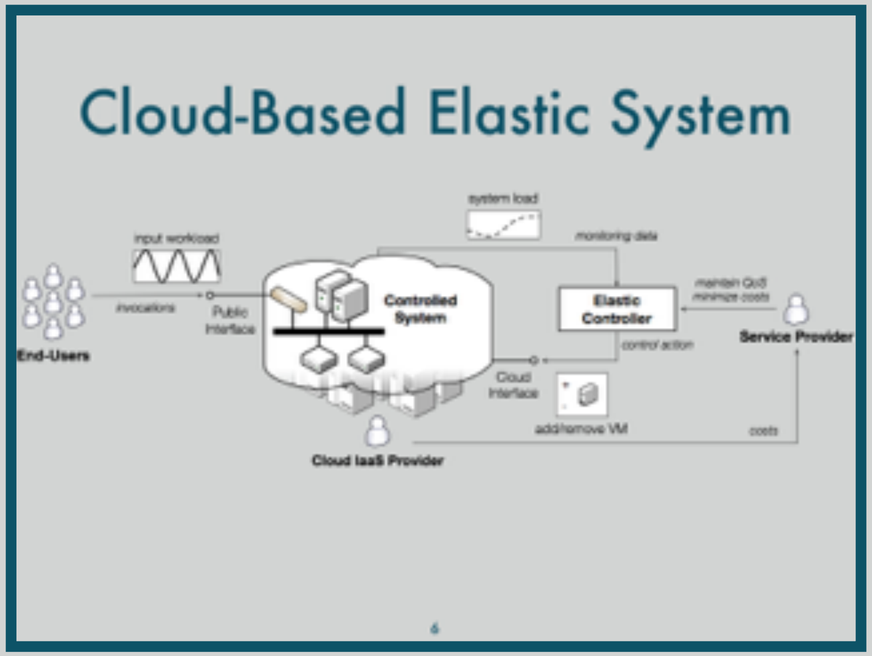

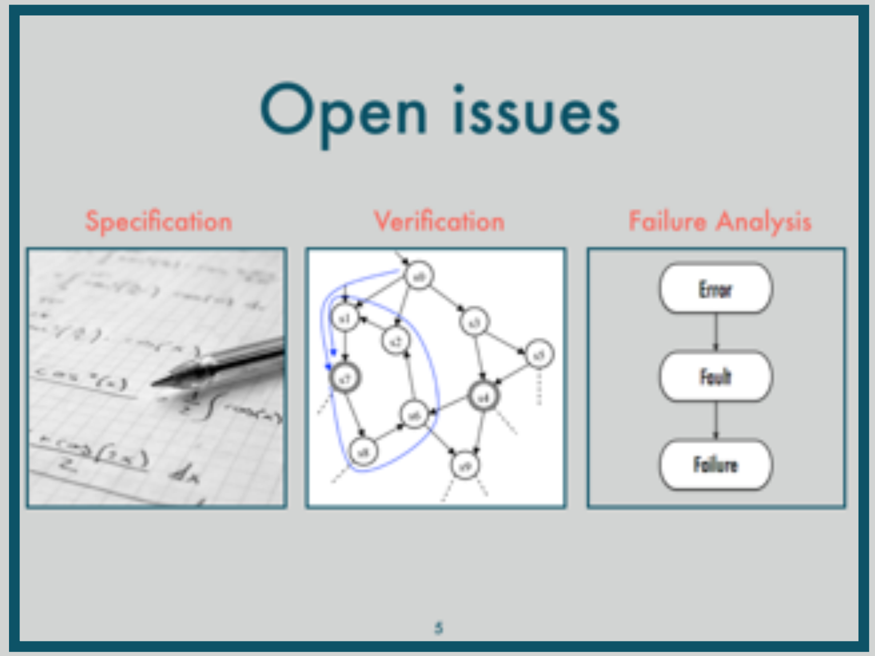

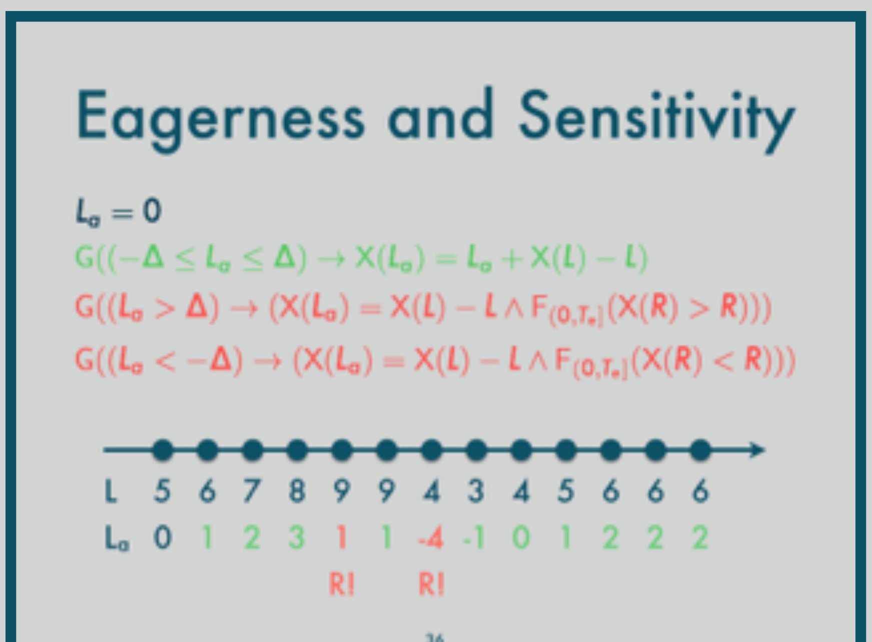

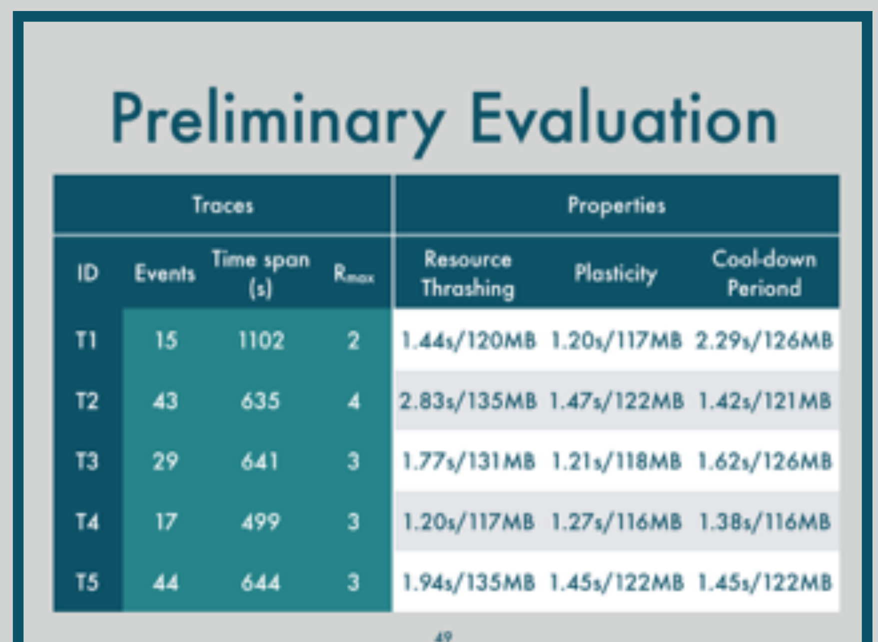

Traces				Properties		
ID	Events	Time span (s)	R_{max}	Resource Thrashing	Plasticity	Cool-down Period
T1	15	1102	2	1.44s/120MB	1.20s/117MB	2.29s/126MB
T2	43	635	4	2.83s/135MB	1.47s/122MB	1.42s/121MB
T3	29	641	3	1.77s/131MB	1.21s/118MB	1.62s/126MB
T4	17	499	3	1.20s/117MB	1.27s/116MB	1.38s/116MB
T5	44	644	3	1.94s/135MB	1.45s/122MB	1.45s/122MB

Specification of Cloud-Based Elastic Systems

Srđan Krstić
Politecnico di Milano

Joint work with Carlo Ghezzi, Domenico Bianculli, Marcello Bersani, Alessio Gambi and Schahram Dustdar