POLITECNICO
MILANO 1863

# About Monitoring
## in a service world

Barbara Pernici
barbara.pernici@polimi.it

SmartGreens, April 25, 2016

## Outline

- Monitoring in the 2010's in the service world
- Monitoring as a design issue
- What to monitor
- What are monitoring data for?
- Challenges and future research work

## Monitoring in the 2010's

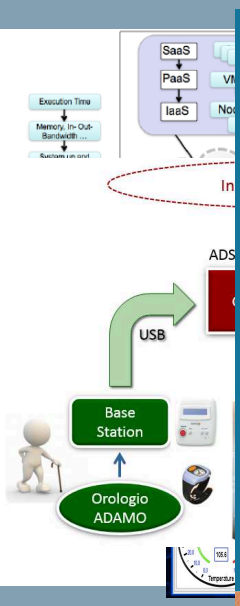





Barbara Pernici

## Adaptive Service-based Information Systems

Key issues:

- Design of adaptive (self-healing and dependable) service-based applications
- Evaluation of the impact of design choices
- Triggering of the most suitable adaption/repair action

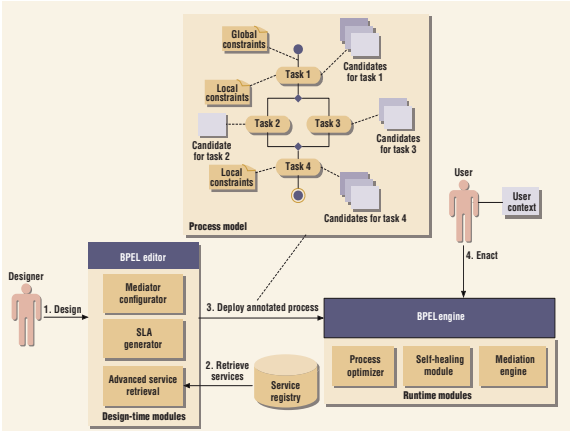


*Italian project: MAIS*
*European projects WS-Diamond, S-Cube*

## Improve the *energy efficiency in data centers and clouds and decrease $CO_2$ emissions*

- Key concepts:
  - Energy efficiency
  - Green related indicators
  - Monitoring
  - Adaptation actions

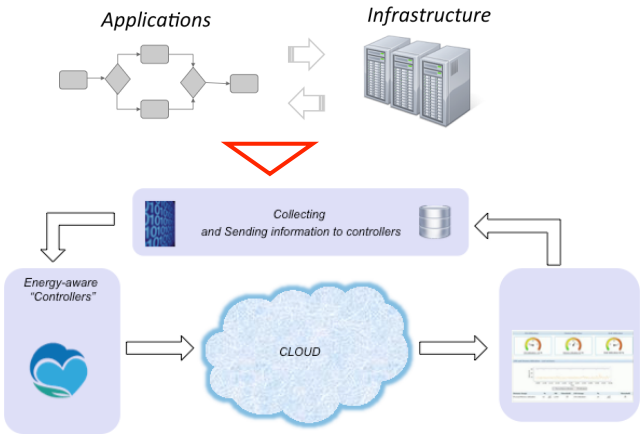


*European projects: GAMES, $ECO_2$Clouds*
*http://eco2clouds.eu*



## The Need for QoS Information in the Service-Oriented Architecture (SOA)

**Many WS with similar functionaly**

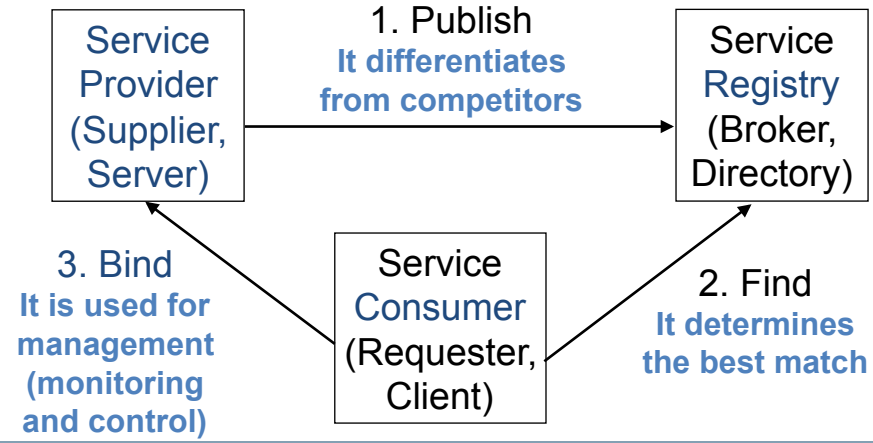

## S-Cube Life-Cycle Model

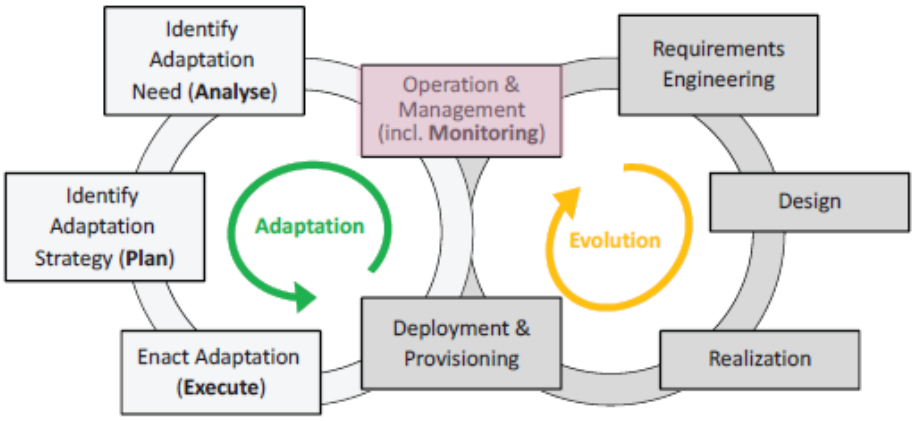




## Monitoring

Needs a systematic approach

To be inserted in the life-cycle
- in the design phase

Main issues:
1. What
2. Usage

## Monitoring

Needs a systematic approach

To be inserted in the life-cycle
- in the design phase

Main issues:
1. **What**
2. Usage



# WHAT TO MONITOR

## Quality parameters/attributes/dimensions

Usually are application dependent

In some case could be only observable

In some other case could be also controllable



## Meta-model for SQM

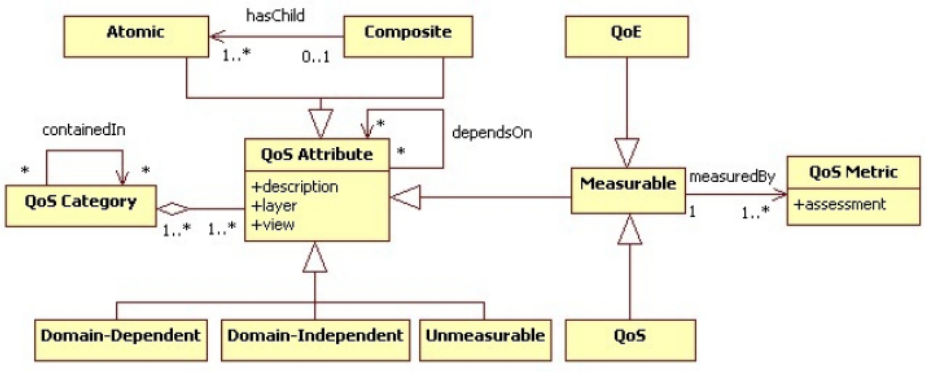


Kritikos et al., CSUR, 2013

## Monitoring characteristics / meta-data

- Sampling frequency
  - Where it is controlled
  - Close to the data
- Quality of monitored data
  - Timeliness
  - Accuracy
  - Completeness
  - Availability
  - ….



## An example for controlling quality

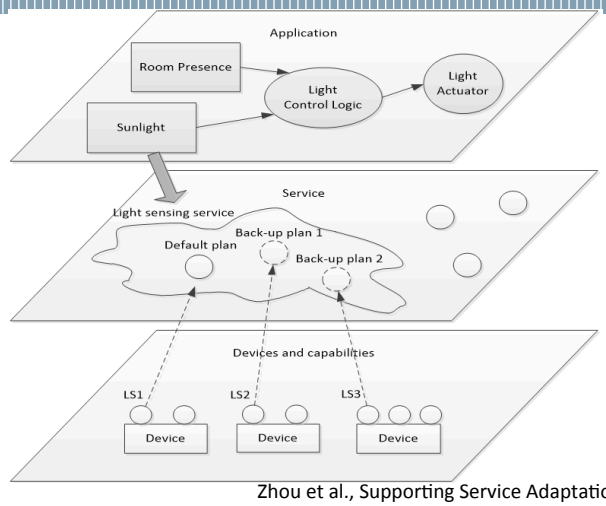


Zhou et al., Supporting Service Adaptation in Fault Tolerant Internet of Things, SOCA 2015

## Issues

Faults are not only concerning availability

- Data quality (accuracy)
  - Occasional errors, intermittent faults
  - Systematic errors
    - stuck-at faults, multiplicative faults, drift faults
- Timeliness



## Summary on WHAT

- Definition of monitored variables
- Measuring values
  - How often
  - Where
  - Provider
- Selection of sources
- Design of quality evaluation for monitored data

## Monitoring

Needs a systematic approach

To be inserted in the life-cycle
- in the design phase

Main issues:
1. What
2. **Usage**



# USAGE OF MONITORED DATA

## Assessment and settlement

| | | |
|---|---|---|
| Assessment | Condition Evaluator | The ability to define the party responsible of SLO evaluation |
| | Qualifying Condition | The ability to define conditions that must hold in order to assess an SLO |
| | Obliged | The ability to express the party in charge of delivering what is promised in an SLO |
| | Assessment Schedule | The ability to express the assessment frequency of an SLO |
| | Validity Period | The ability to express the time period in which the SLO is guaranteed |
| | Recovery Actions | The ability to express corrective actions to be carried out when an SLO is violated |
| Settlement | Penalties | The ability to express penalties incurred when one party violates its promises |
| | Rewards | The ability to express rewards incurred when one party overwhelms its promise |
| | Settlement Actions | The ability to express actions concerning the final SLA outcome |



## Usage issues

- Identifying events
- Condition evaluation
- Resources and cost

## Usage issues

- **Identifying events**
- Condition evaluation
- Resources and cost



## Identifying events
## Assessment – thresholds

Thresholds

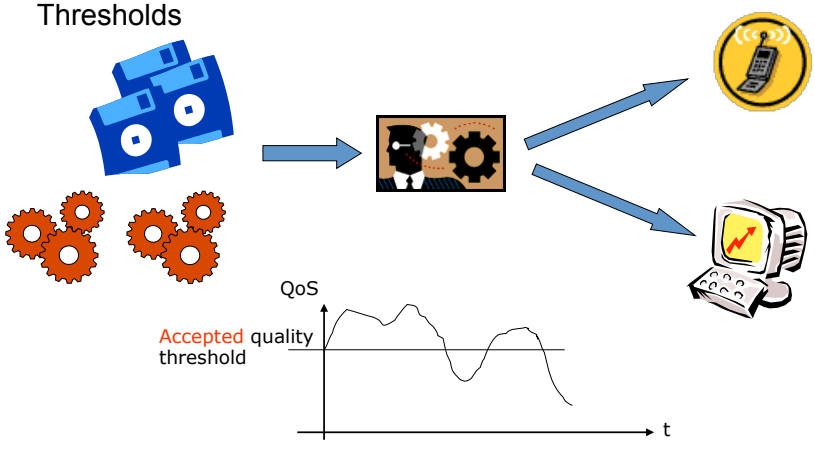

## Definition of thresholds

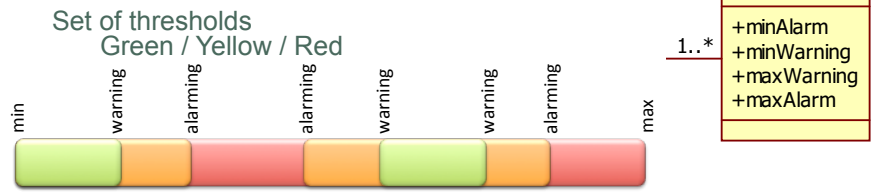


Indicator normalization
Normalization function that allows soft constraints

Indicator aggregation
Composition of monitored data

Ferreira et. al
"Green performance indicators aggregation through composed weighting system"
ICT-GLOW 2012



## Identifying events
## Assessment – trends

**Trends**

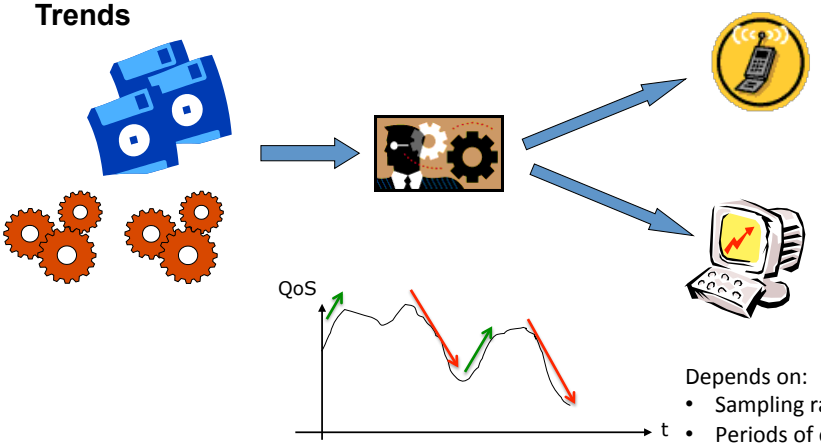


Depends on:
- Sampling rates
- Periods of evaluation

## Assessment
## Events and action feedback analysis – identification of new events

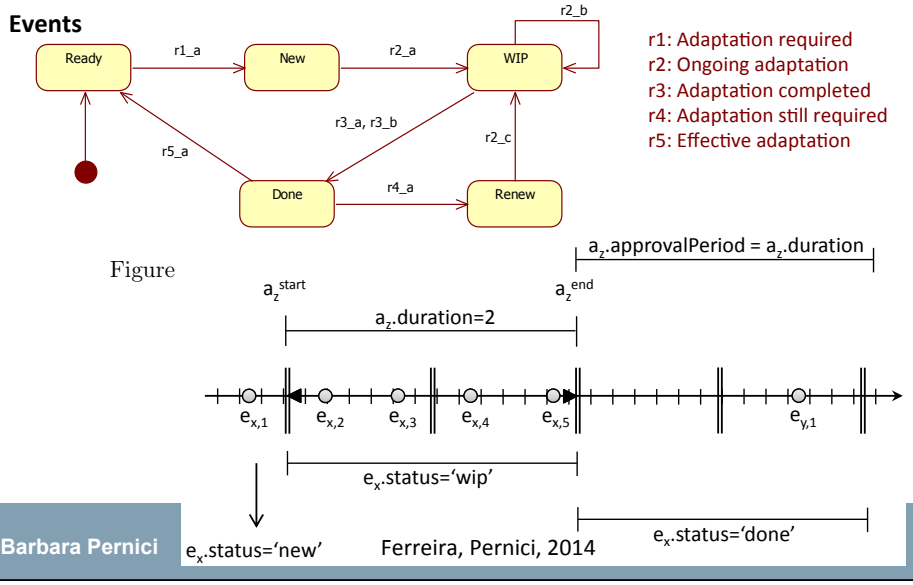


Figure

Ferreira, Pernici, 2014



## Usage issues

- Identifying events
- **Condition evaluation**
- Resources and cost

## Evaluating conditions

Define goals

Some issues

- multiattribute
- user preferences, multiple stakeholders
- context-dependent

Learn from past experience

Heuristics



## Learning goal to goal relations

*GOAL to GOAL* relations are represented through a Bayesian Network. It describes influences existing between the indicators states.

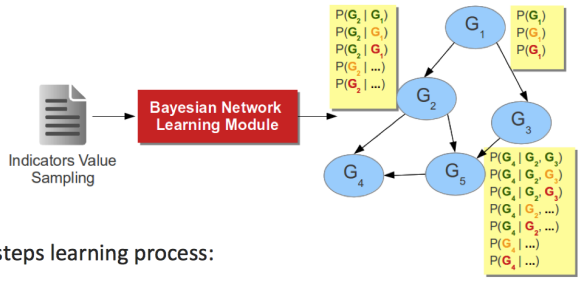


Three steps learning process:

- Structure learning (using correlations)
- Directionality learning (using MMHC algorithm)
- Parameters learning (using Maximum a Priori Estimation)

Vitali et al., 2014

## Selecting actions

- Self-optimization

Criteria

- Stability
- Flexibility
- Elasticity

Focus

- Dependencies between actions



## Learning Action to Goal Relations

ECO$_2$Clouds

*ACTION to GOAL* relations are represented through positive and negative links (positive = action increasing the indicator value).

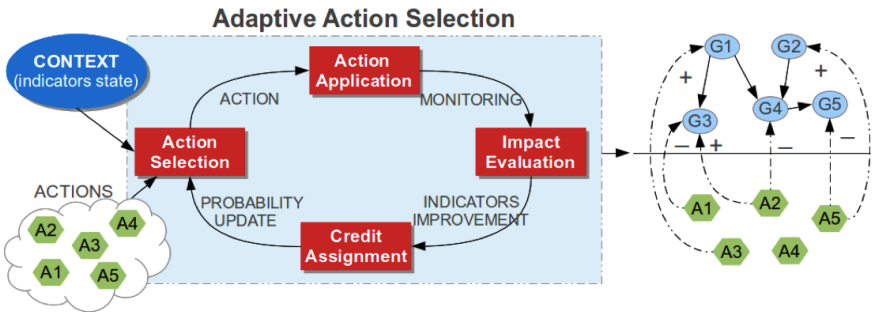


Two steps process:

- *Exploration:* learning the action – goal connections
- *Exploitation:* selecting the best action given a context

Vitali et al., 2014

## Cost of monitoring

Providers
- What is monitored
- What is monitorable
- New probes
- Ability to set up alarms
- Aggregation functionalities
- Available storage

Costs
- may depend on sampling rate
- On demand new dimensions



## Dimensions can be many…

E.g. Cloudera

106 categories of metrics

Sampling rate for all metrics one minute
Aggregation functions applied
Storage is limited

## Summary on USAGE

Monitoring design issues

- Assessment
  - Definition of
    - Events
    - Goals
  - Evaluation infrastructure
- Available resources (for storing monitored data for further analysis)
- Cost of monitoring
  - Actual cost
  - Indirect costs (additional load on the system)



## Lessons learnt and future research work About WHAT

1. Which monitored data?
   - Different sources
   - Different granularities
   - Different quality of data
   - ➢ Creating new sources
   - ➢ Combining existing sources

## Lessons learnt and future research work
## About WHAT

1. (bis) Impact of context
    - Unrelated sources
    - How to find the right sources of data?
        - Retrieval
        - Metadata
        - Value driven



## Lessons learnt and future research work
## About WHAT

1. (ter) Understanding monitored variables
    - ➢ Selection of variables
    - ➢ Relationships between variables

## Lessons learnt and future research work
## About USAGE

2. Need to coordinate monitoring activities

   - Variability in monitoring needs may be variable
   - Monitoring activities: to be selected
   - Monitoring can have phases
     - Not always necessary (e.g. setting up vs running system)

   ➢ design a monitoring process, not only a monitoring system



## Lessons learnt and future research work
## About USAGE

3. Monitoring requirements

   a) Specification
   - Definition of monitoring requirements for different goals
   - Requirements related to different types of events and risks
   - Selecting variables, sampling rates, granularity, sources, …
   - Specify accuracy, completeness, timeliness, ...

   b) Control
   - Policies
   - Assessment
   - Satisfying requirements given constraints
     - rules

## Current state and needed research

WHAT

1. Which monitored data?
   - mature research sector
   - selection of monitoring variables is becoming the main issue

USAGE

2. Monitoring processes: Need to coordinate monitoring activities
   Monitoring as a process needs investigation
   Design criteria for monitoring are needed

3. Specification of monitoring requirements
   monitoring requirements need research on informal/semi-formal/formal models and methods



## QUESTIONS?

## References


- Monica Vitali, Barbara Pernici, Una-May O'Reilly: Learning a goal-oriented model for energy efficient adaptive applications in data centers. Inf. Sci. 319: 152-170 (2015)
- Monica Vitali, Barbara Pernici: PiE - Processes in Events: Interconnections in Ambient Assisted Living. OTM Workshops 2015: 157-166
- Cinzia Cappiello, Paco Melia, Barbara Pernici, Pierluigi Plebani, Monica Vitali: Sustainable choices for cloud applications: a focus on CO2 emissions. ICT4S 2014
- Cinzia Cappiello, Barbara Pernici, Laura Villani: Strategies for Data Quality Monitoring in Business Processes. WISE Workshops 2014: 226-238





- Kyriakos Kritikos, Barbara Pernici, Pierluigi Plebani, Cinzia Cappiello, Marco Comuzzi, Salima Benbernou, Ivona Brandic, Attila Kertész, Michael Parkin, Manuel Carro: A survey on service quality description. ACM Comput. Surv. 46(1): 1 (2013)
- Alexandre Mello Ferreira, Barbara Pernici: Using Intelligent Agents to Discover Energy Saving Opportunities within Data Centers. RE4SuSy@RE 2013
- Alexander Kipp, Tao Jiang, Jia Liu, Mariagrazia Fugini, Monica Vitali, Barbara Pernici, Ioan Salomie: Applying green metrics to optimise the energy consumption footprint of IT service centres. IJSSC 2(3): 158-174 (2012)
- Danilo Ardagna, Marco Comuzzi, Enrico Mussi, Barbara Pernici, Pierluigi Plebani: PAWS: A Framework for Executing Adaptive Web-Service Processes. IEEE Software 24(6): 39-46 (2007)

- Giuseppe Aceto, Alessio Botta, Walter de Donato, Antonio Pescapè: Cloud monitoring: A survey. Computer Networks 57(9): 2093-2115 (2013)
- Claudia Foglieni, Mirjana Mazuran, Giovanni Meroni, Pierluigi Plebani: Retrieving Sensors Data in Smart Buildings Through Services: A Similarity Algorithm. ICSOC Workshops 2014: 281-291
- Alexandre Mello Ferreira, Barbara Pernici, Pierluigi Plebani:
- Green Performance Indicators Aggregation through Composed Weighting System. ICT-GLOW 2012: 79-93
- IEEE Computing Edge, April 2016
- Fernando Chirigati, Harish Doraiswamy, Theodoros Damoulas, and Juliana Freire, Data Polygamy: The Many-Many Relationships among Urban Spatio-Temporal Data Sets, ACM SIGMOD, 2016




## References


S. Zhou, K.-J. Lin, J. Na, C.-C. Chuang, and C.-S. Shih, “Supporting service adaptation in fault tolerant internet of things,” in Service-Oriented Computing and Applications (SOCA), 2015 IEEE 8th International Conference on, Oct 2015, pp. 65–7

## Two monitoring examples

- Cloud monitoring and energy efficiency / environmental impact
- Business processes and data quality



## Cloud monitoring

ECO2Clouds

Energy efficiency (IT)

Environmental impact

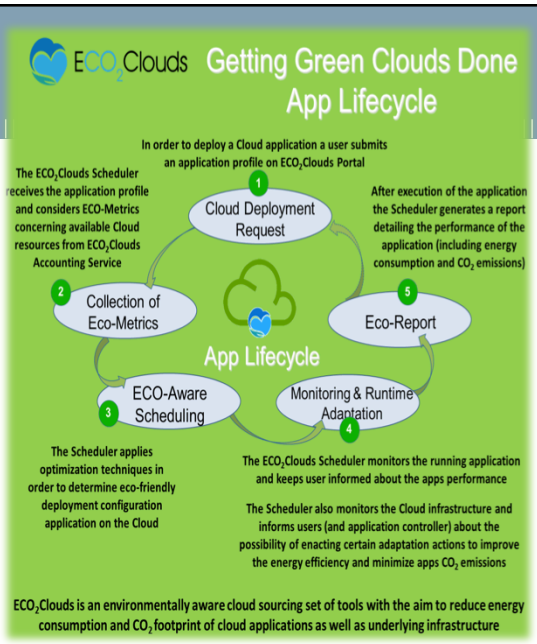
ECO2Clouds Getting Green Clouds Done
App Lifecycle
In order to deploy a Cloud application a user submits an application profile on ECO2Clouds Portal
The ECO2Clouds Scheduler receives the application profile and considers ECO-Metrics concerning available Cloud resources from ECO2Clouds Accounting Service
1
Cloud Deployment Request
After execution of the application the Scheduler generates a report detailing the performance of the application (including energy consumption and CO2 emissions)
2
Collection of Eco-Metrics
5
Eco-Report
App Lifecycle
3
ECO-Aware Scheduling
Monitoring & Runtime Adaptation
4
The Scheduler applies optimization techniques in order to determine eco-friendly deployment configuration application on the Cloud
The ECO2Clouds Scheduler monitors the running application and keeps user informed about the apps performance
The Scheduler also monitors the Cloud infrastructure and informs users (and application controller) about the possibility of enacting certain adaptation actions to improve the energy efficiency and minimize apps CO2 emissions
ECO2Clouds is an environmentally aware cloud sourcing set of tools with the aim to reduce energy consumption and CO2 footprint of cloud applications as well as underlying infrastructure



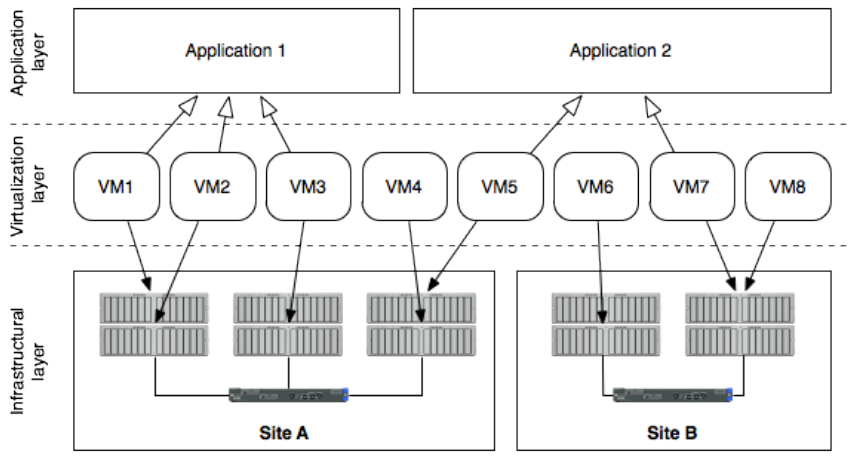
ECO2Clouds
Application layer
Application 1
Application 2
Virtualization layer
VM1
VM2
VM3
VM4
VM5
VM6
VM7
VM8
Infrastructural layer
Site A
Site B

ECO2Clouds

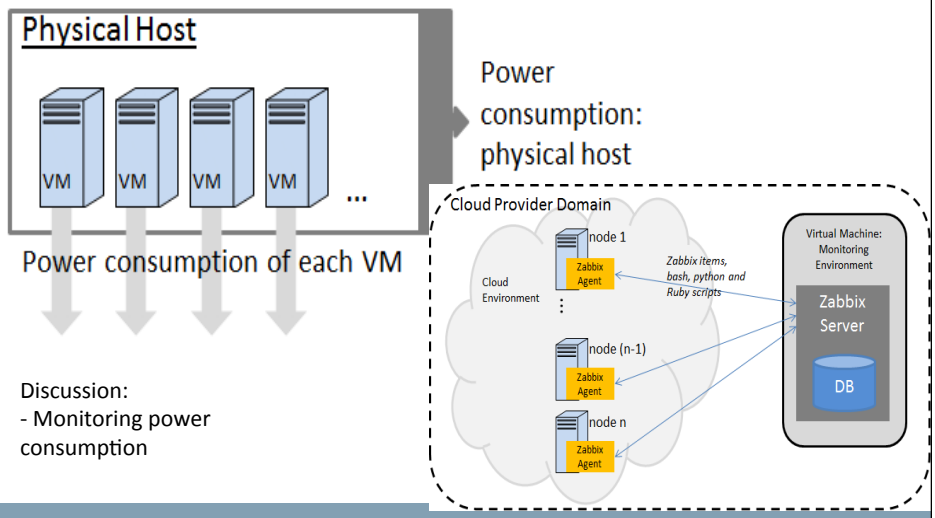


Discussion:
- Monitoring power consumption



## Monitoring or modeling?

ECO2Clouds

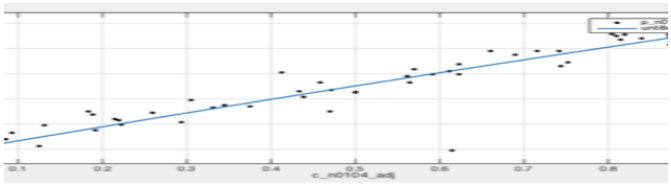

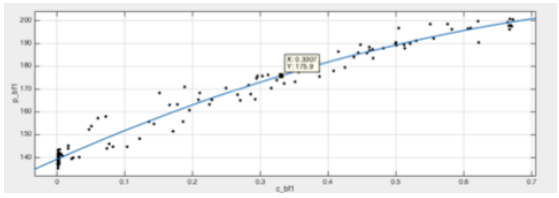

## Consumption models for applications

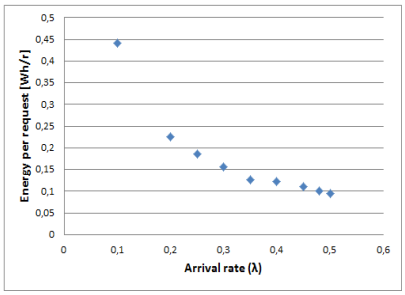


Figure 20: Energy per job

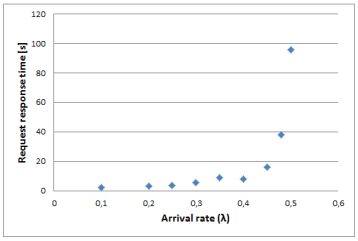


Figure 21: Response time per job



## Monitoring business processes

Monitored data

- Mining processes
- Process events
- KPIs, some examples
  - Time
    - Starting, ending time, time of events
  - GPI (environmental impact)
  - Process specific
    - Achievement of **goals**

## Monitoring business processes

Poor quality negatively affects the efficiency and effectiveness of business processes

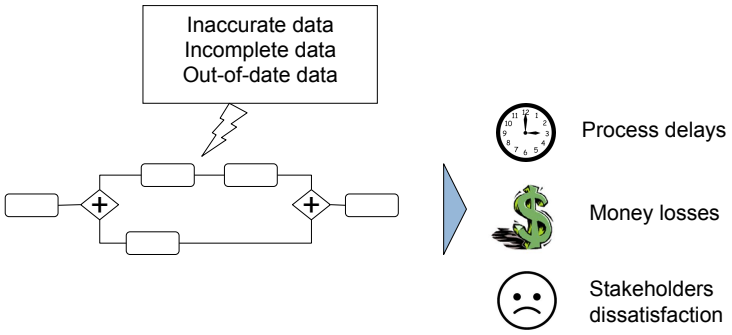




## Quality-aware redesign of the business process

Insert Data Quality blocks inside the process to improve its data quality level.

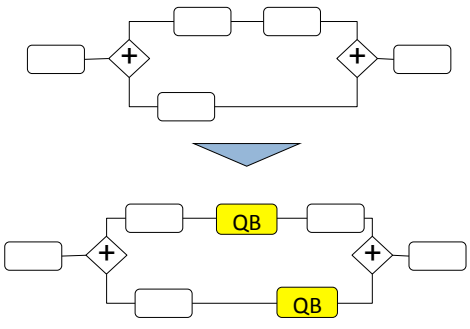

## Where to monitor? Local checks

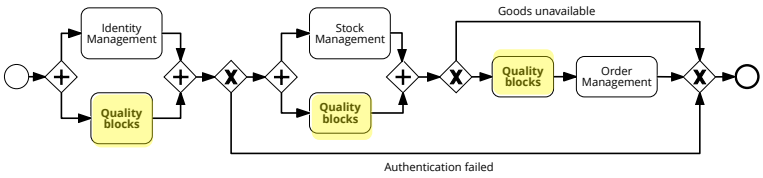




## Wher to monitor? Preliminary checks

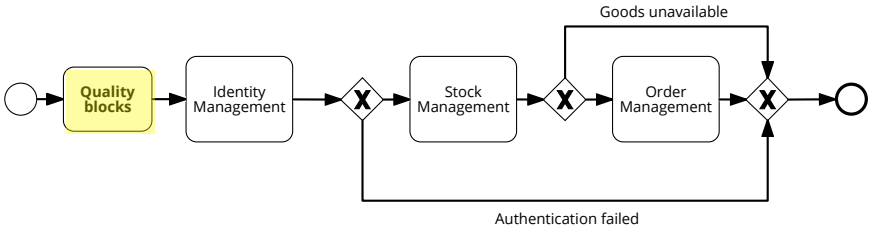

## Design of the Data Quality Blocks

We defined four types of Data Quality blocks:

•*Accuracy Data Quality blocks*: they reveal the correctness of the data values used in the process. We defined accuracy quality blocks for both textual and numeric data types

•*Completeness Data Quality blocks*: they search for the "null" values inside the data sources that are used by the tasks of the analyzed process.

•*Consistency Data Quality blocks:* check the consistency among values (e.g., cities and postal codes)

•*Timeliness Data Quality blocks:* check the validity of data in terms of their age and updateness



## Local check

Using the local check, each important activity is checked:

- If data are updated the block is inserted before the activity
- If data are read the block is inserted in parallel

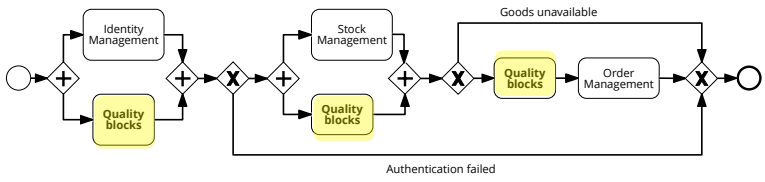


This first configuration is suitable for simple processes, since such an accurate monitoring increases the complexity of the process flow and the execution time.

## Preliminary Check

*Preliminary check* takes into account the impact of quality blocks on execution time, but has also the goal of preventing to insert low quality values in databases. Therefore this configuration envisions that a quality block is inserted *before the execution of the process*

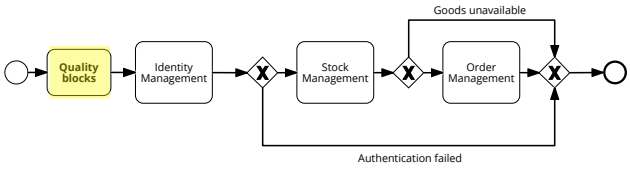


This solution is particularly useful for processes that have often to deal with low quality data since quality is monitored and problems are revealed before the process execution



## Where to monitor? Parallel Check

Using *Parallel check* the quality check is performed in parallel with the process execution, in order to decrease the impact on the execution time

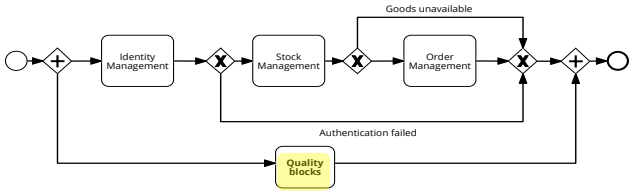


This solution is useful for processes in which quality constraints are most of the time satisfied. It is also useful if the different activities are performed by different organizations