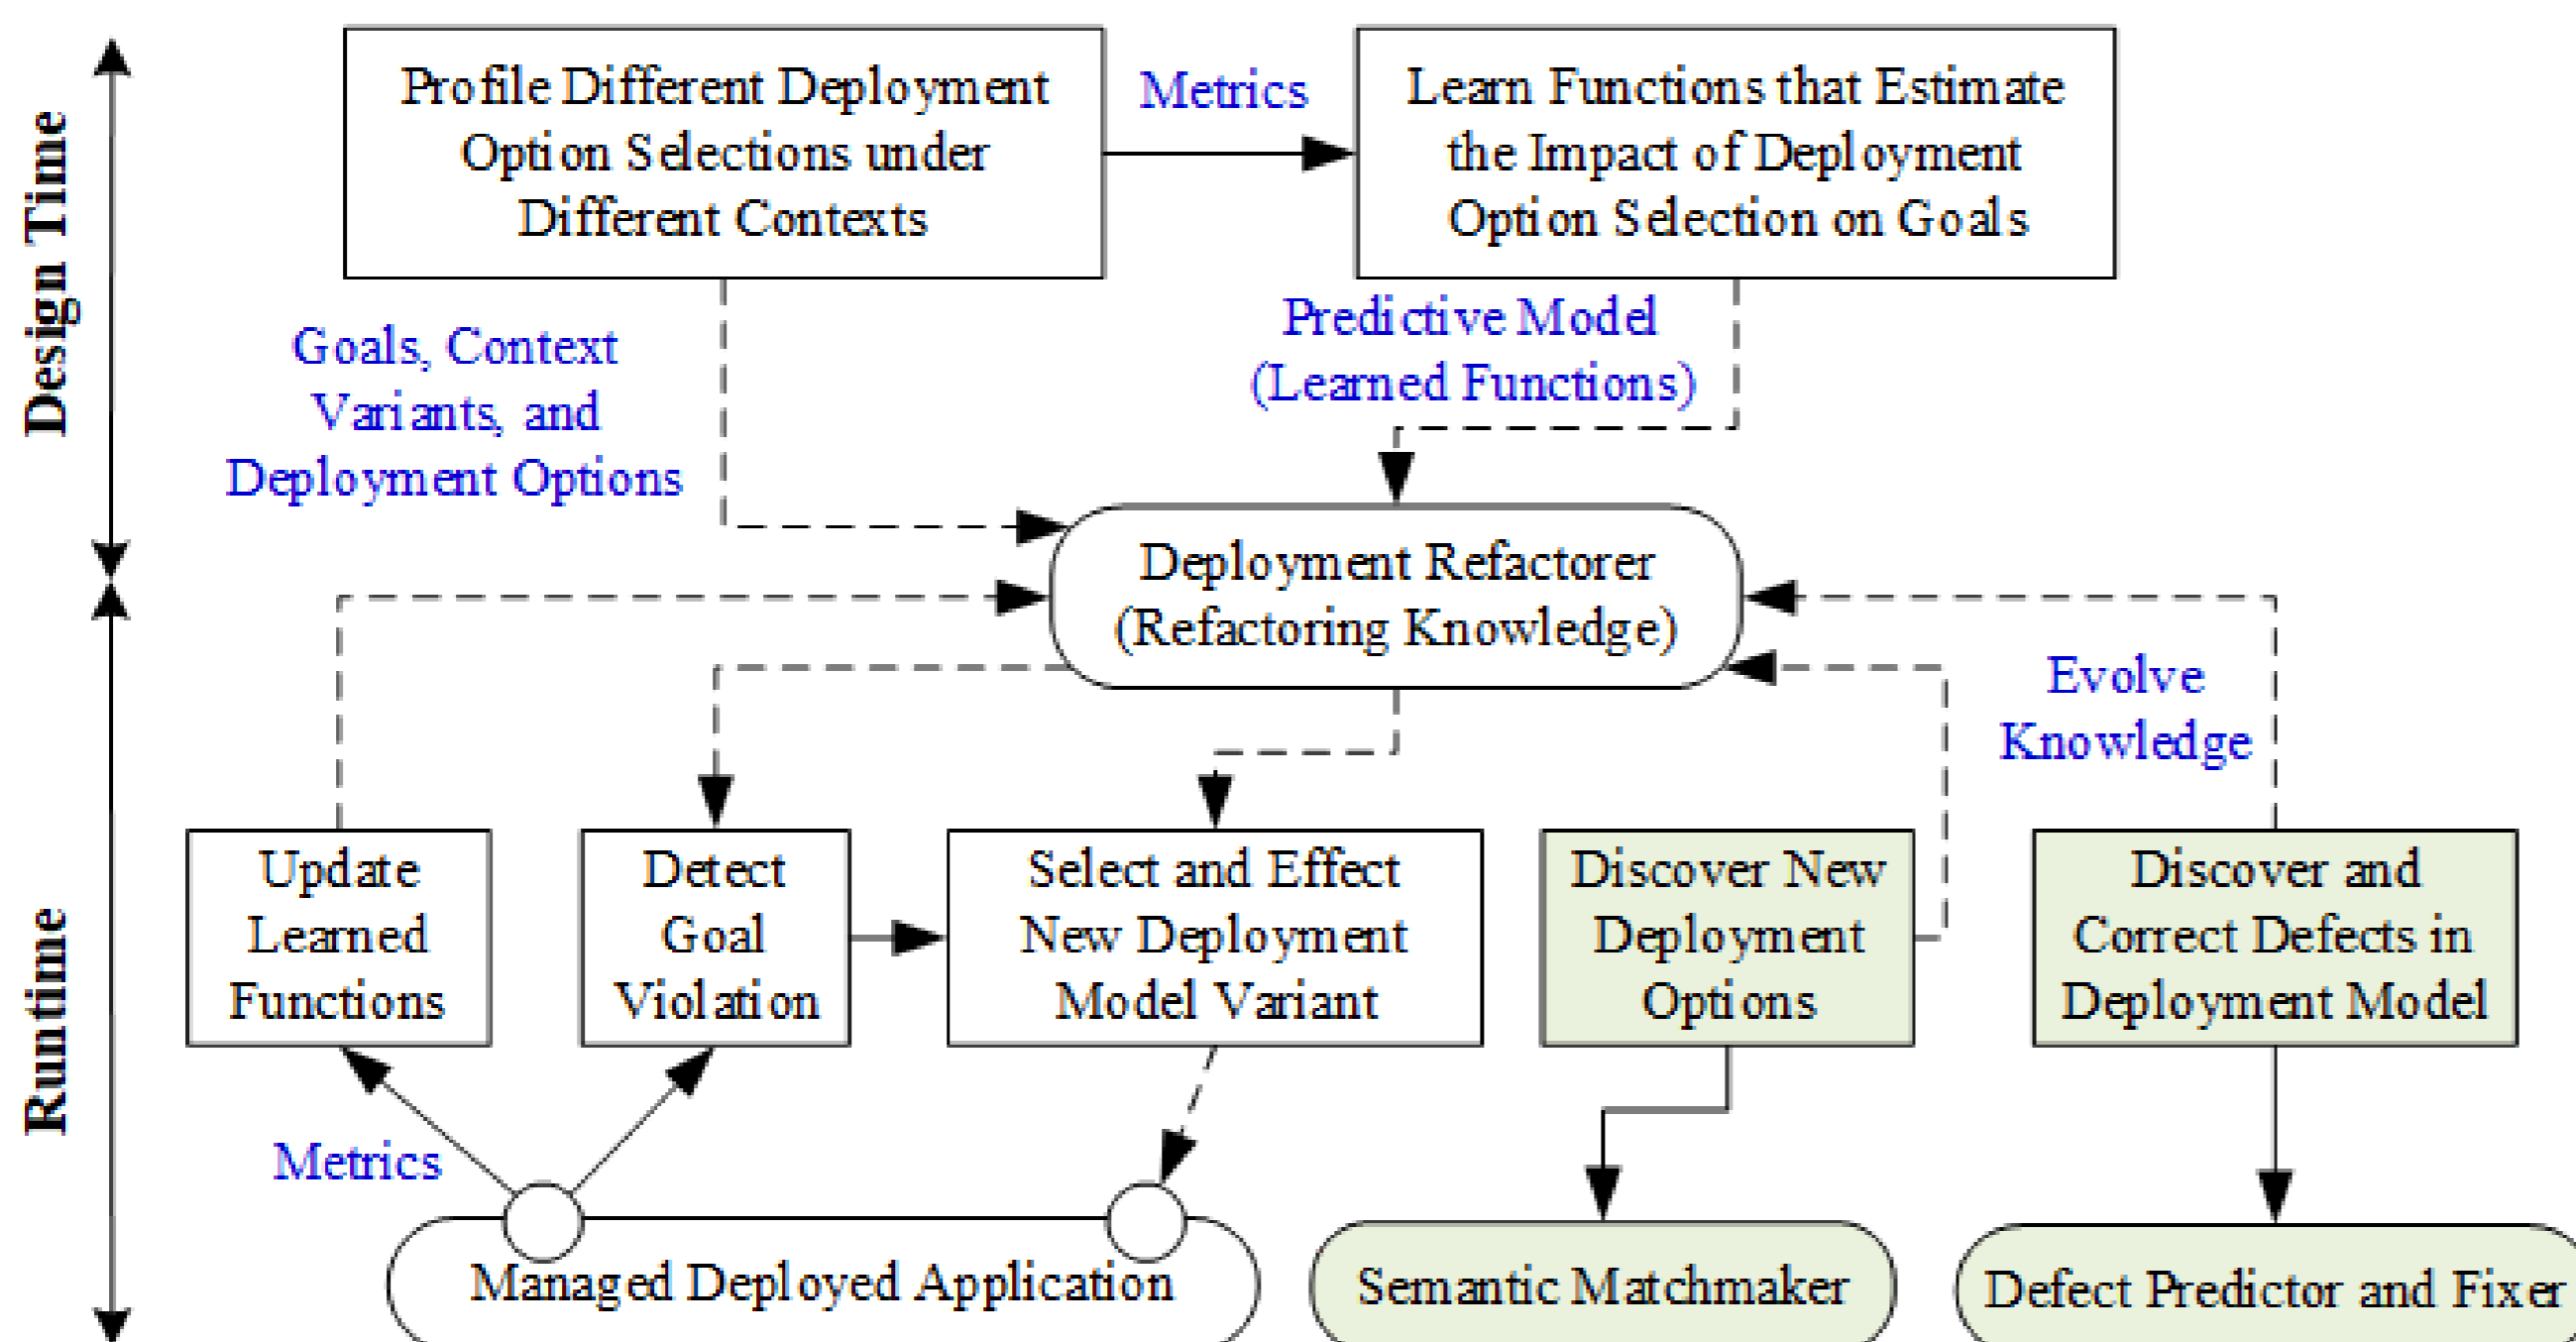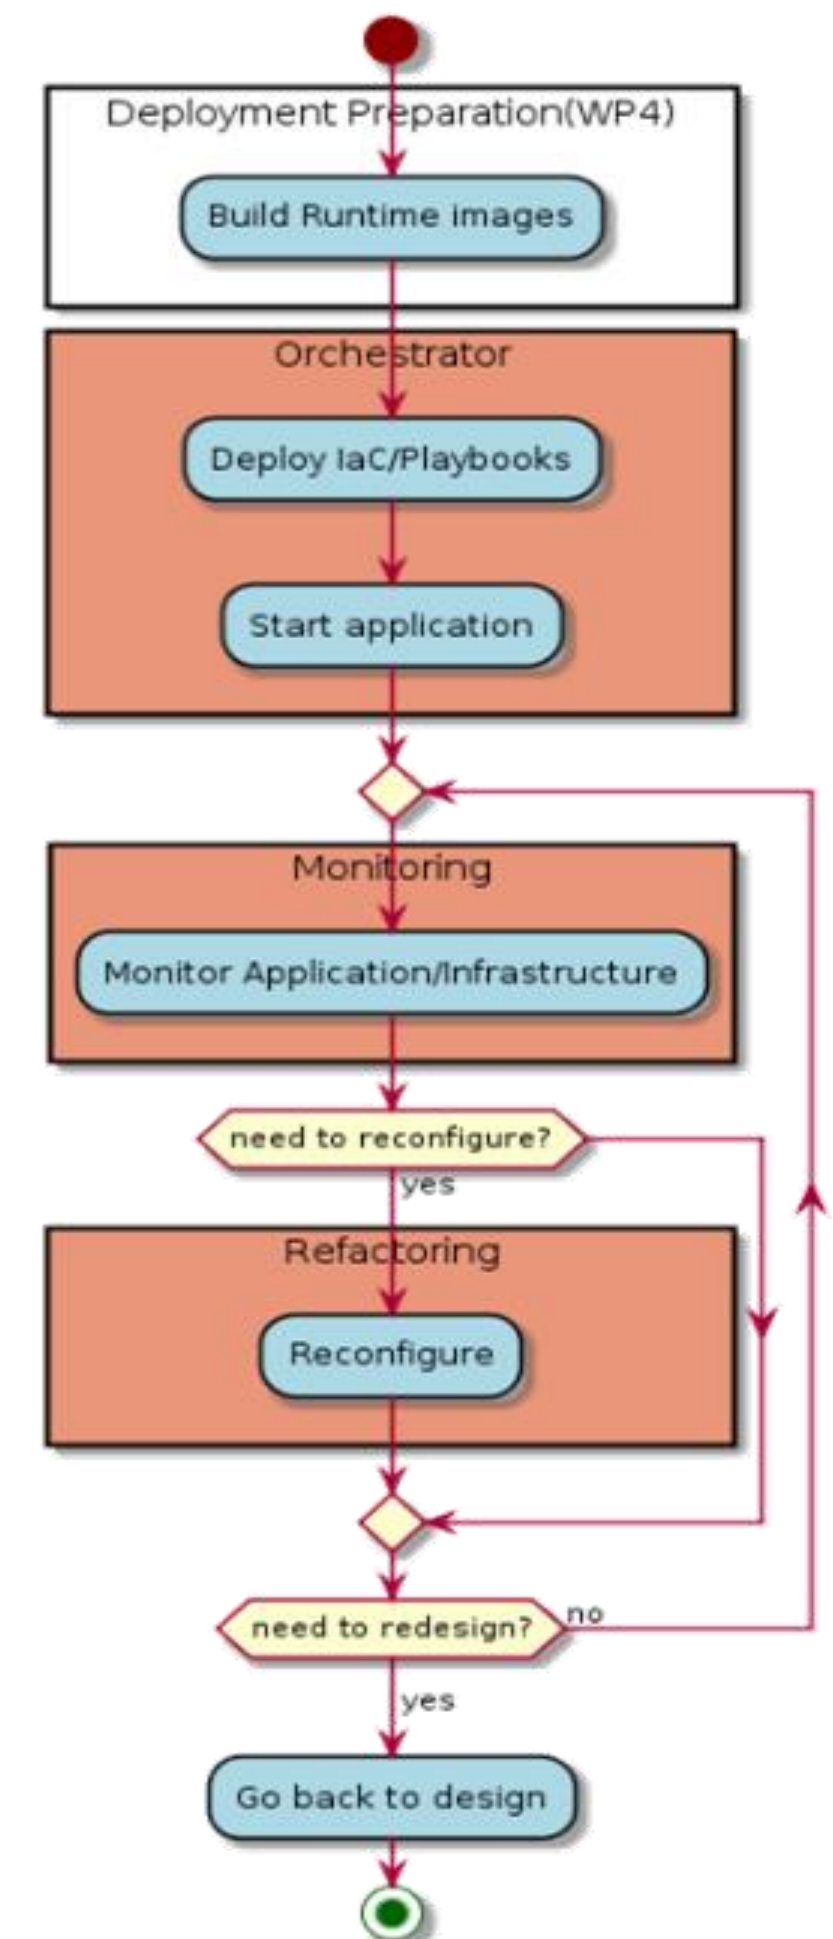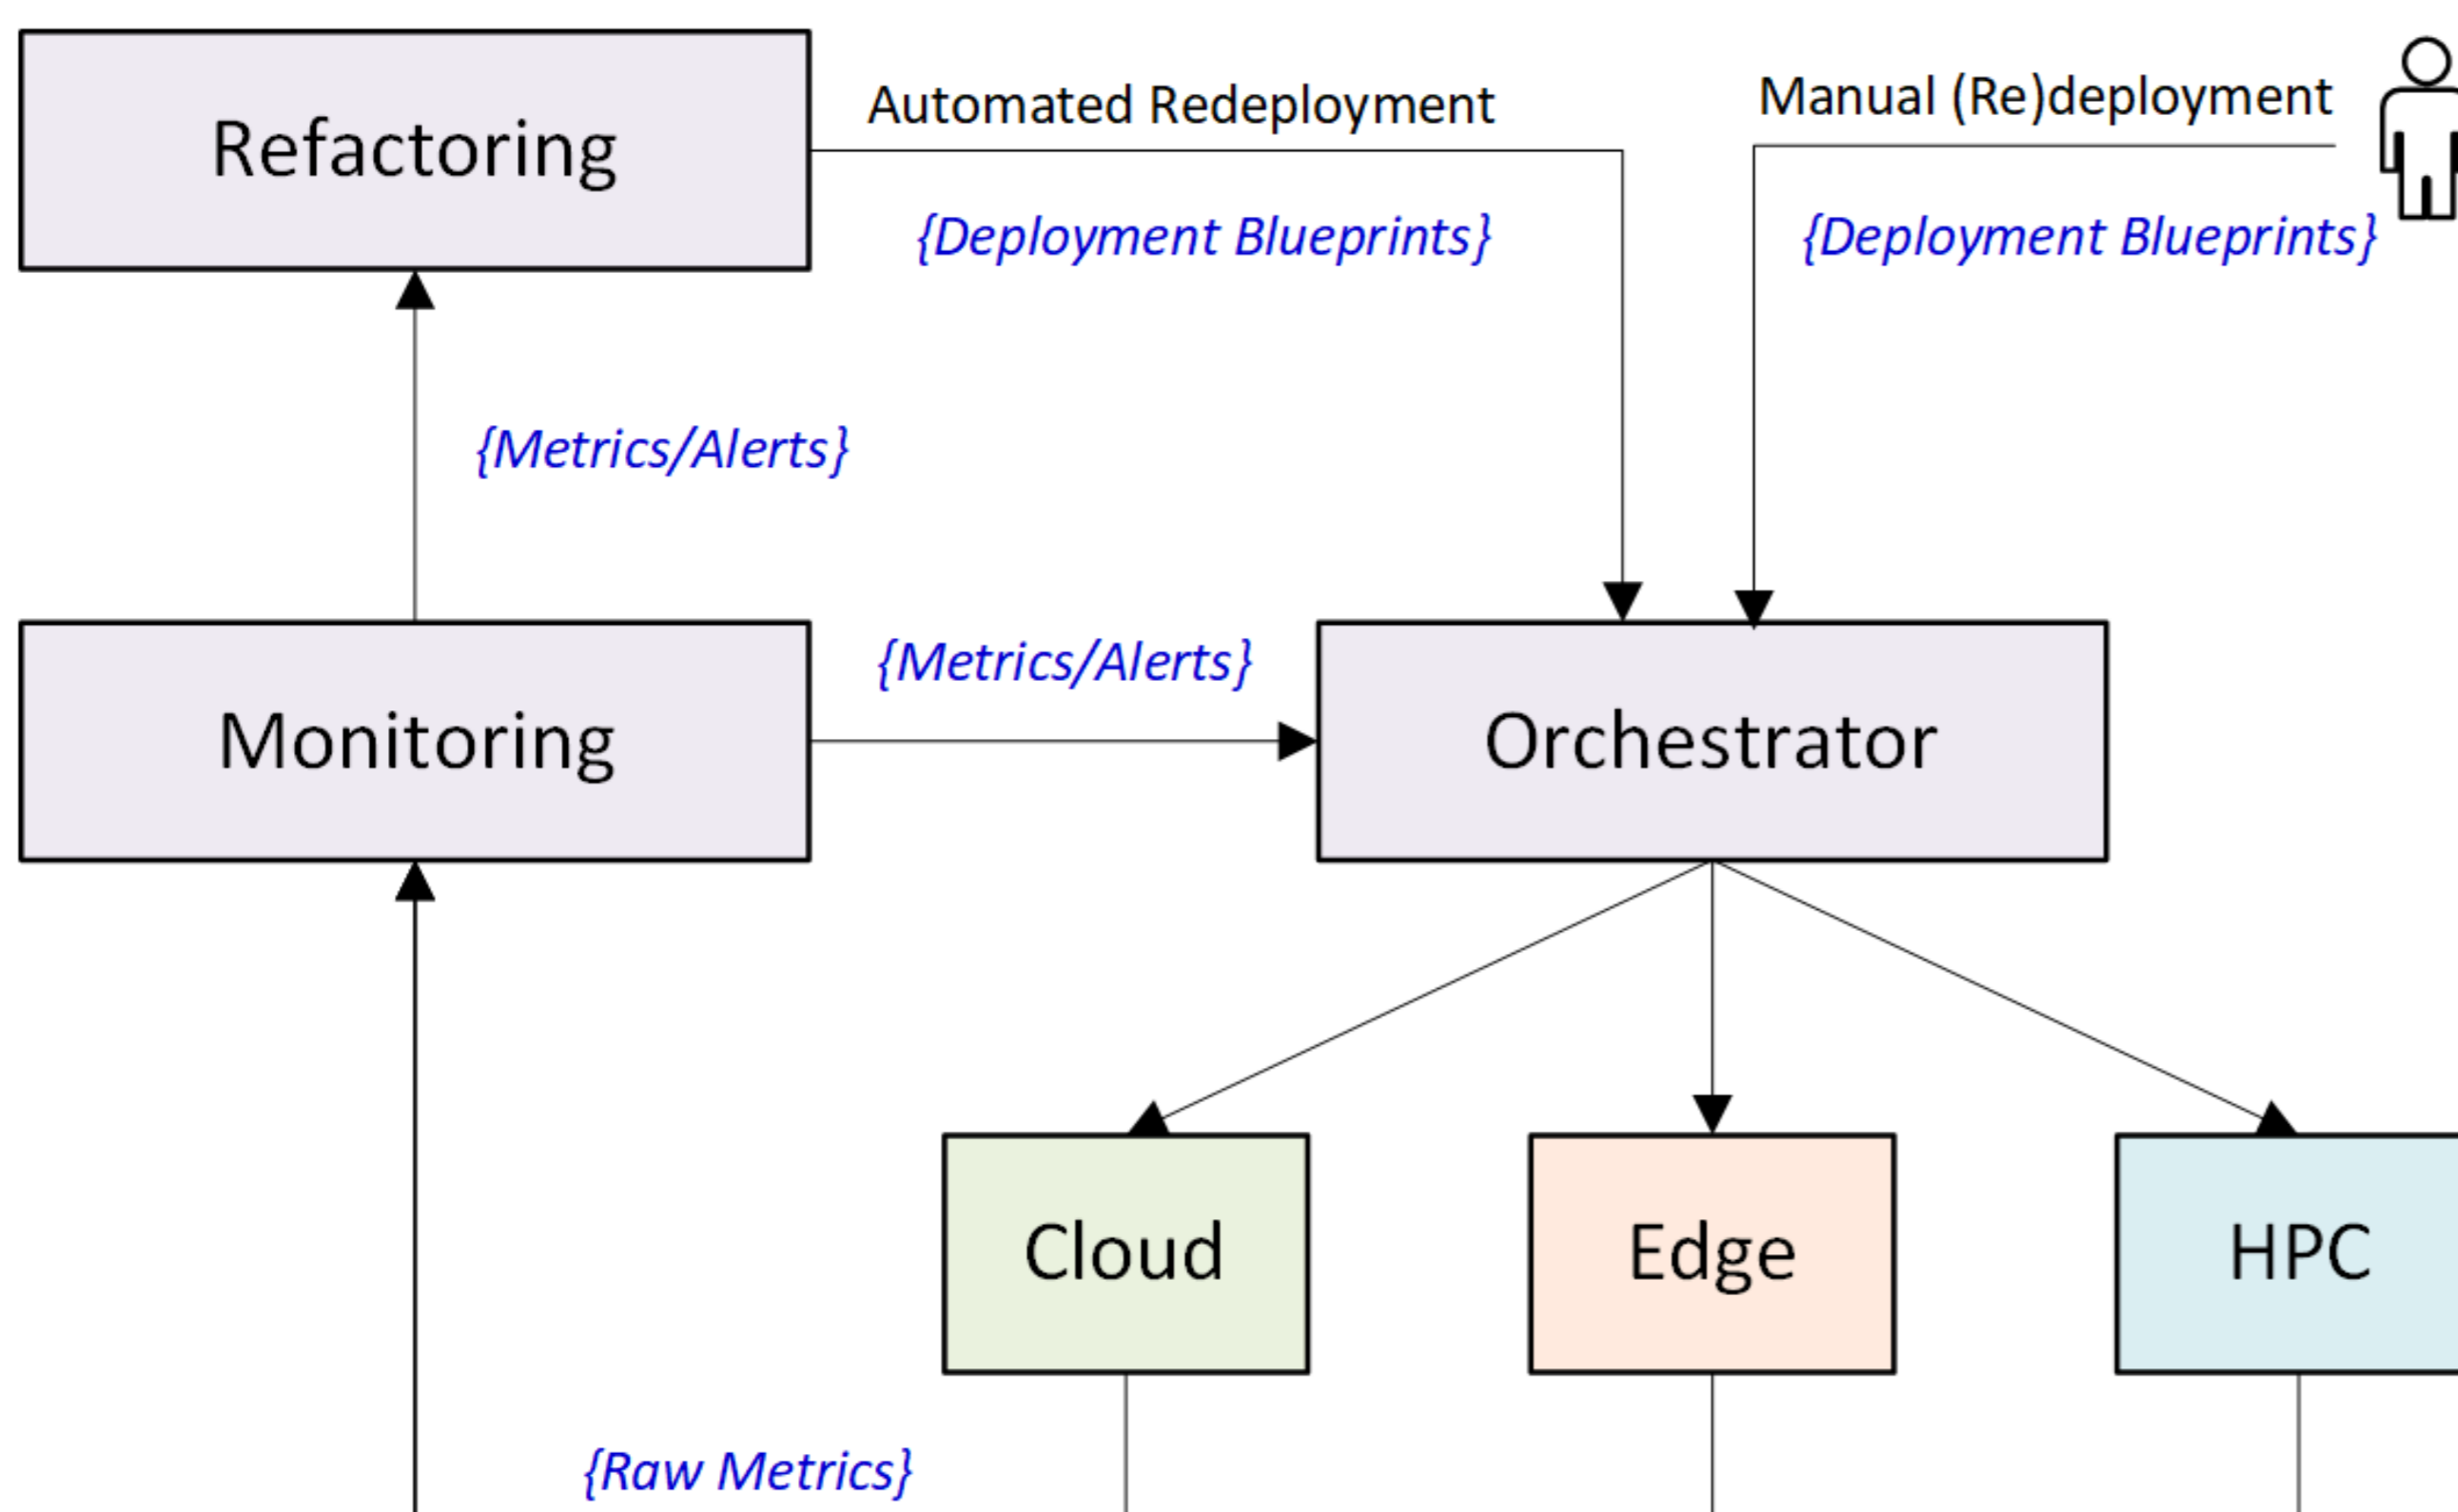# Intelligent Re-deployment Feedback Loop for Hybrid Applications

Statement of Problem

Distributed applications may consist of numerous components that work together.
Each component may have defined some list of resources it requires or requests in order to run.
Based on available resources in the computing environment (e.g. cluster), a deployment plan is developed to run the application.
A component may be deployed on various physical and/or virtual machines which satisfy the prerequisite requirements of that component, based on what is available.
At some point, the application may not be satisfying its requirements, and may need to be redeployed.
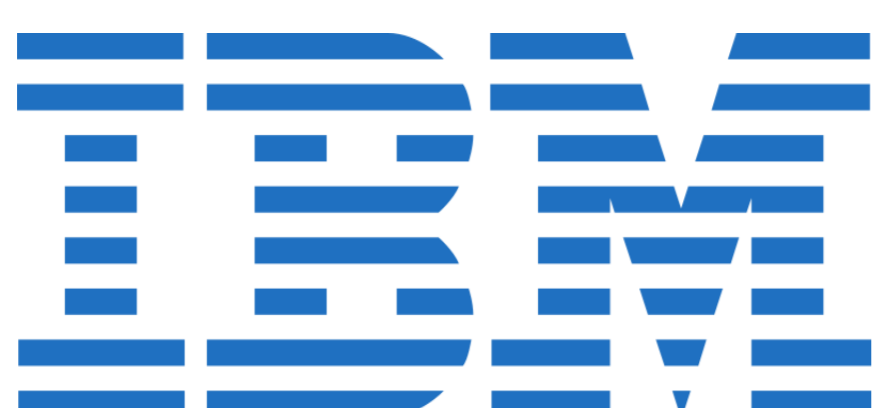
Proposed Solution

Monitoring information is continuously collected at the infrastructure level.
If available, performance information is also collected from the application.
If the performance of the application is not satisfactory with its current deployment, it may be possible to improve its performance by deploying it differently on the existing or on other hardware.
The monitoring information is fed to the Redeployment component, which recommends a new deployment of the distributed application.

Enablers

- Containers
  - Fast to boot
  - Fast to scale
  - Scale at fine granularity
  - Isolation
  - Kubernetes
- Control Theory
  - Automatic
  - Lightweight
  - Scalable
  - Design-time guarantees
  - PID controller

Contacts: Indika Kumara (i.p.k.weerasingha.dewage@tue.nl), Kalman Meth (meth@il.ibm.com), Giovanni Quattrocchi (giovanni.quattrocchi@polimi.it)