



State, Context, Adaptability, and Scale for Self-Learning Closed Loop Policies

Liam Fallon
Sven van der Meer
John Keeney

12th December 2017

Quality Attributes for a Policy Framework

The Attributes

- State: Keep track of where we are and what we're doing
- Context: Keep track of what's going on around us
- Adaptability: Able to be changed and to change in response to where we are, what we're doing, and what's going around us

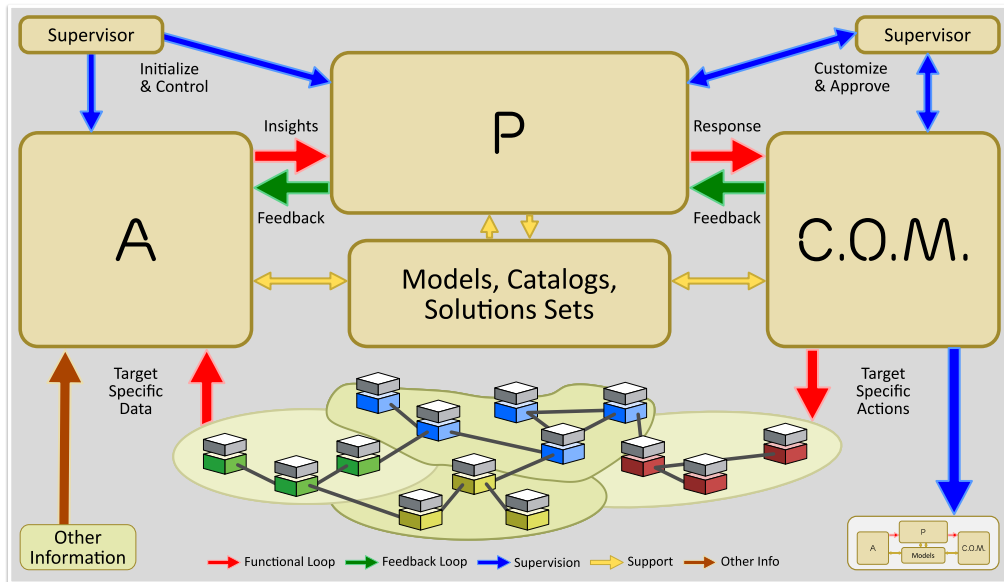
Whilst allowing

- Scale: The ability to do the three bullets above while increasing and decreasing our capacity depending on the load

With the ultimate goal of

- Self-learning: Harnessing and governing machine learning to provide learnt policies

Policy for Closed Loops, Work since 2012



- Analytics, Control and Orchestration becoming complex
- Policies in Control Loops were static, unstructured, and in silos
- Policy Program
 - Established well founded **theory** for policy in Closed Loops
 - Developed a pre-production policy system to apply that theory in **practice**


COMPA, a reference architecture for closed loops

Control, Orchestration, Management, Policy, Analytics

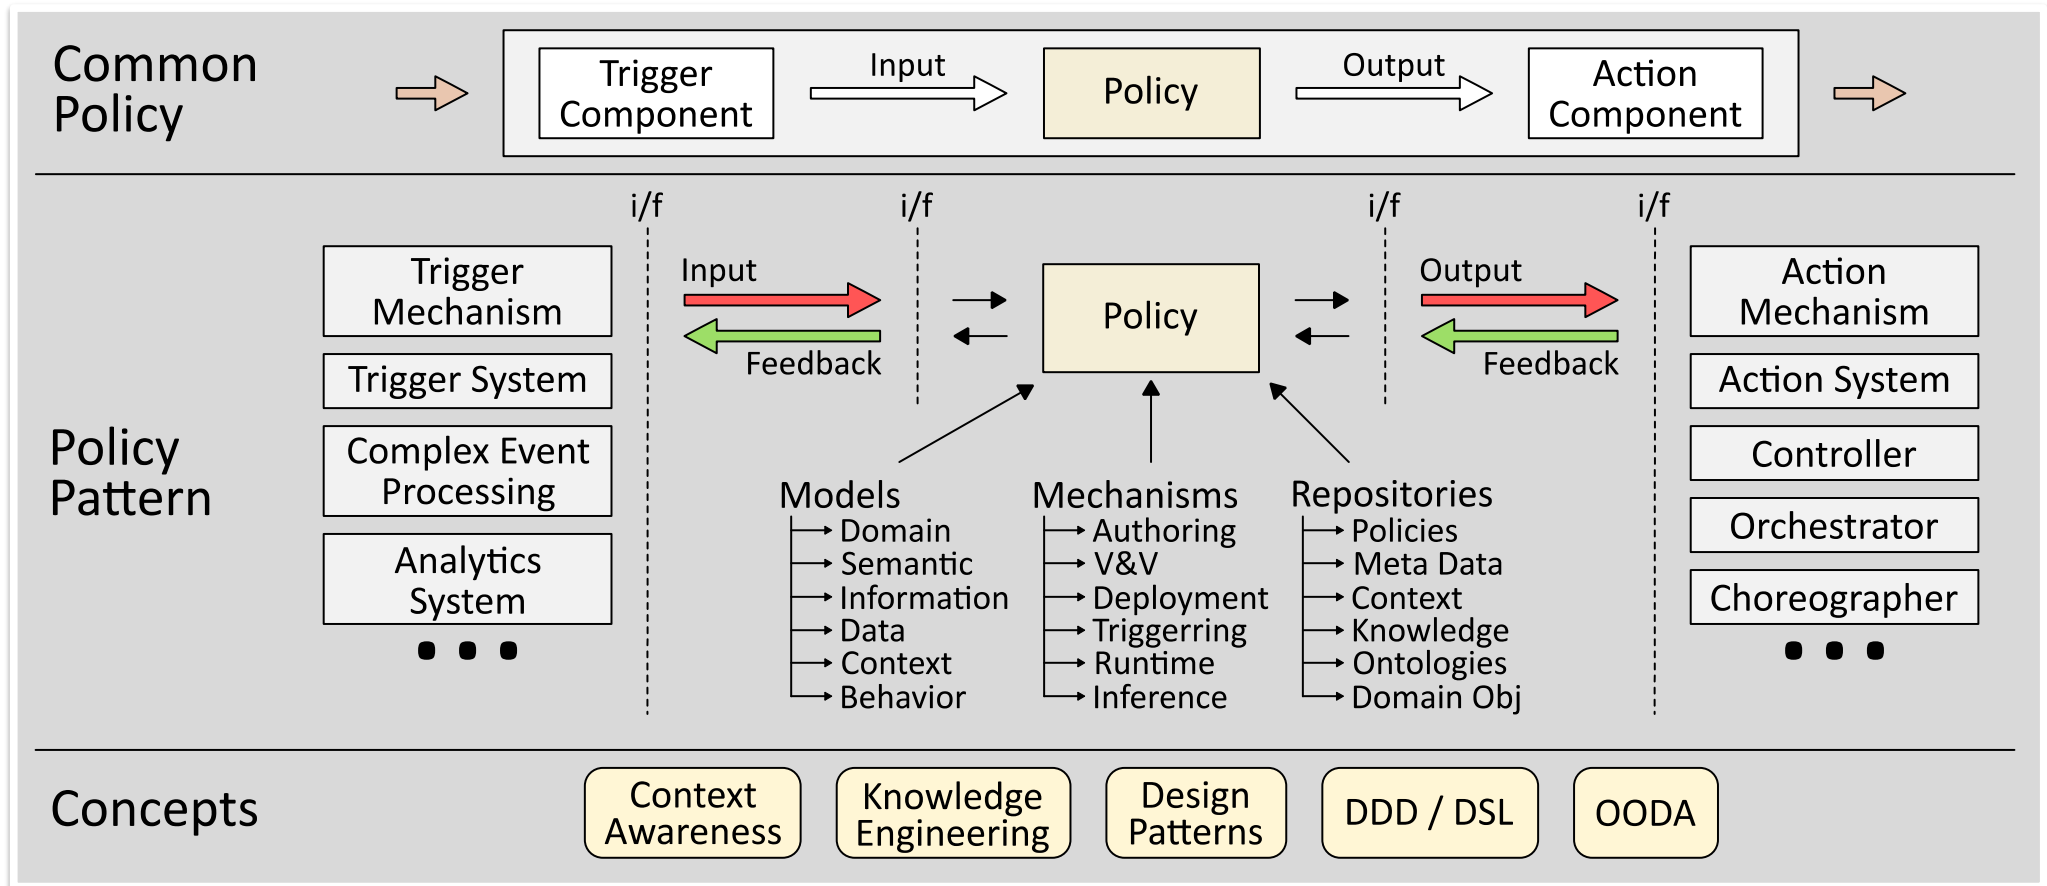
UNIFIED POLICY THEORY (UPT) PRIMER

R^G

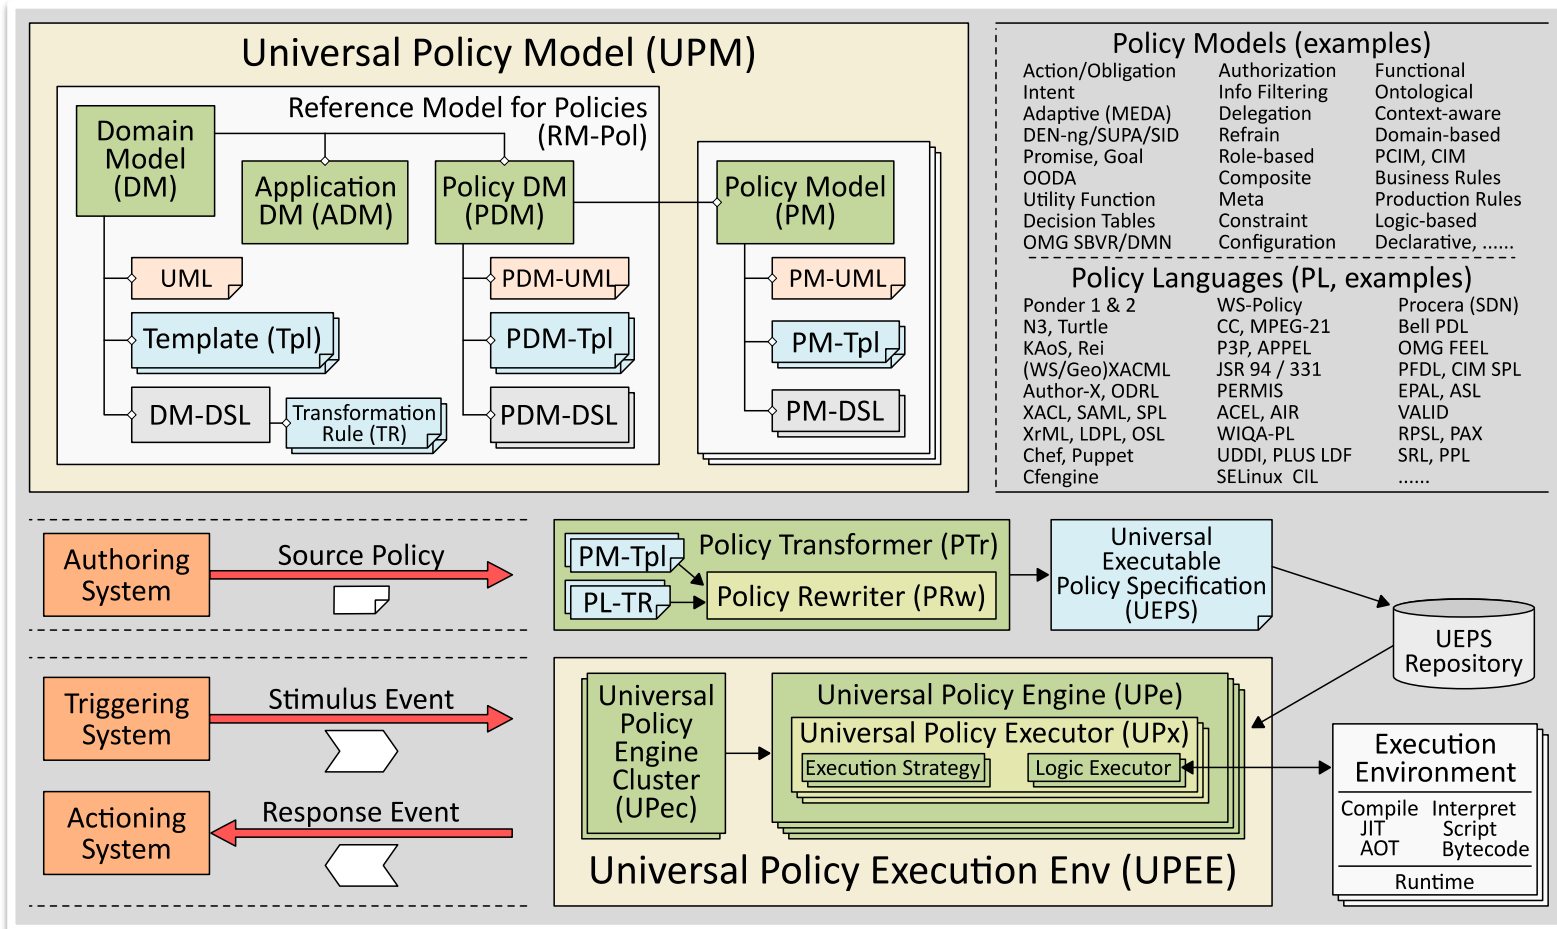
[For full details on Unified Policy Theory, see our research papers in the APEX project in ResearchGate](#)



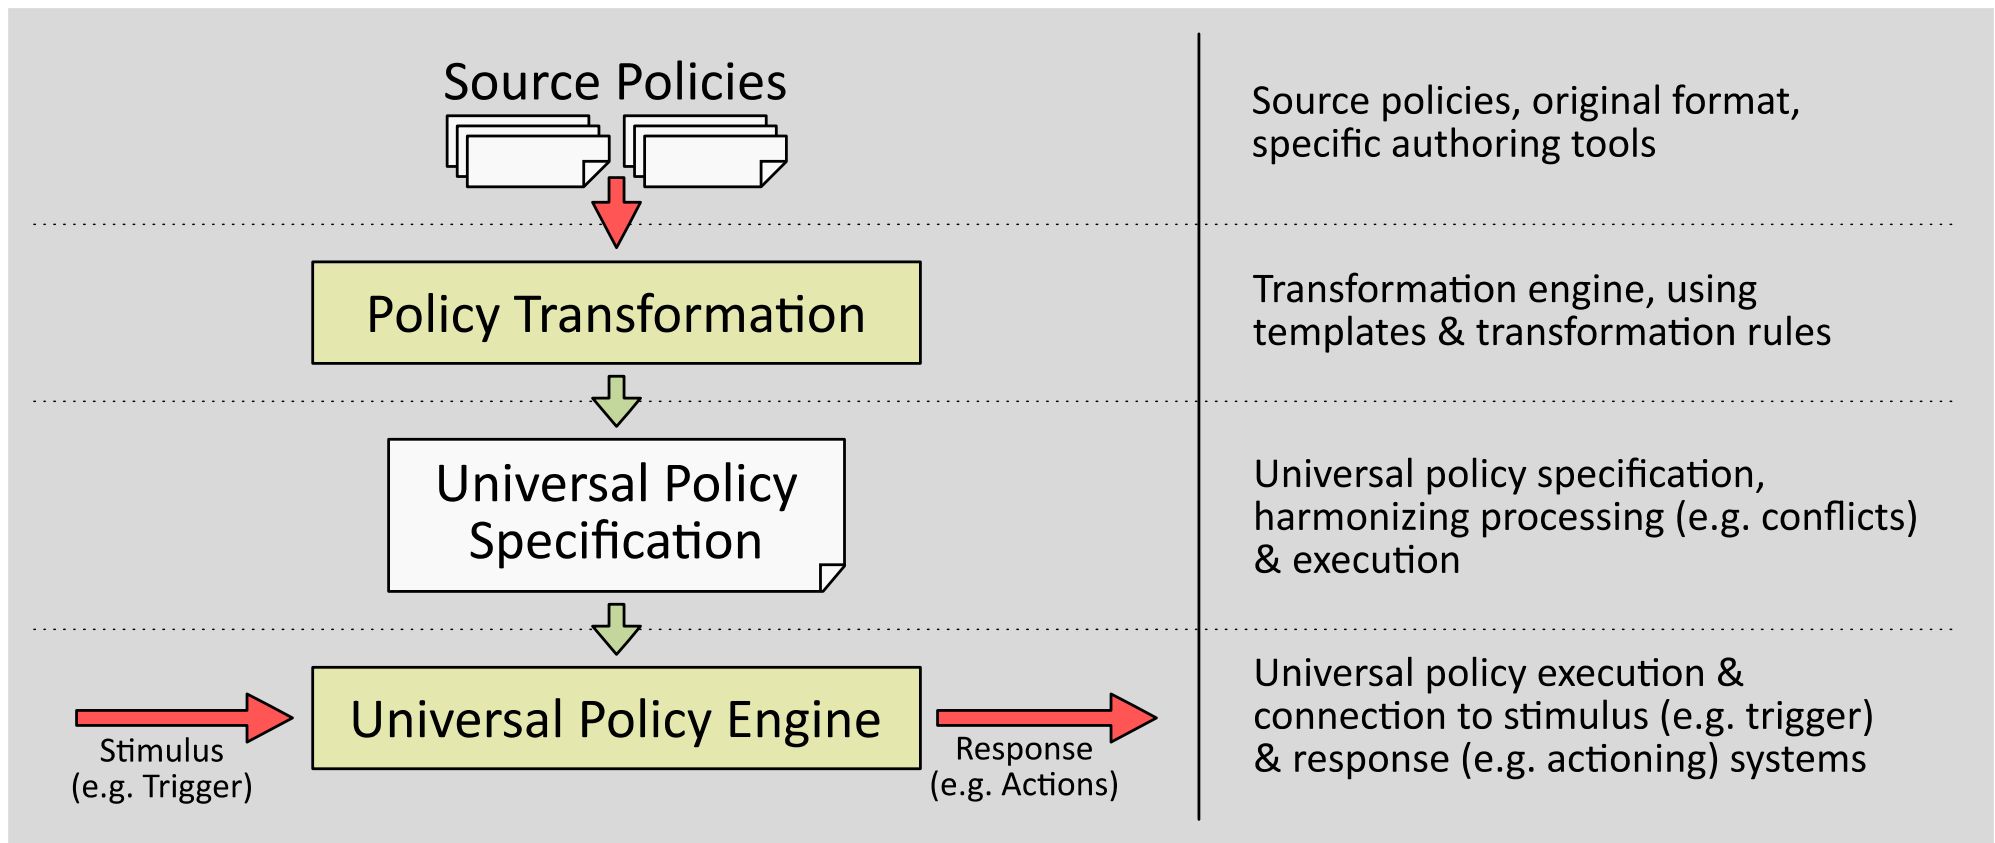
Policy Patterns and Roles



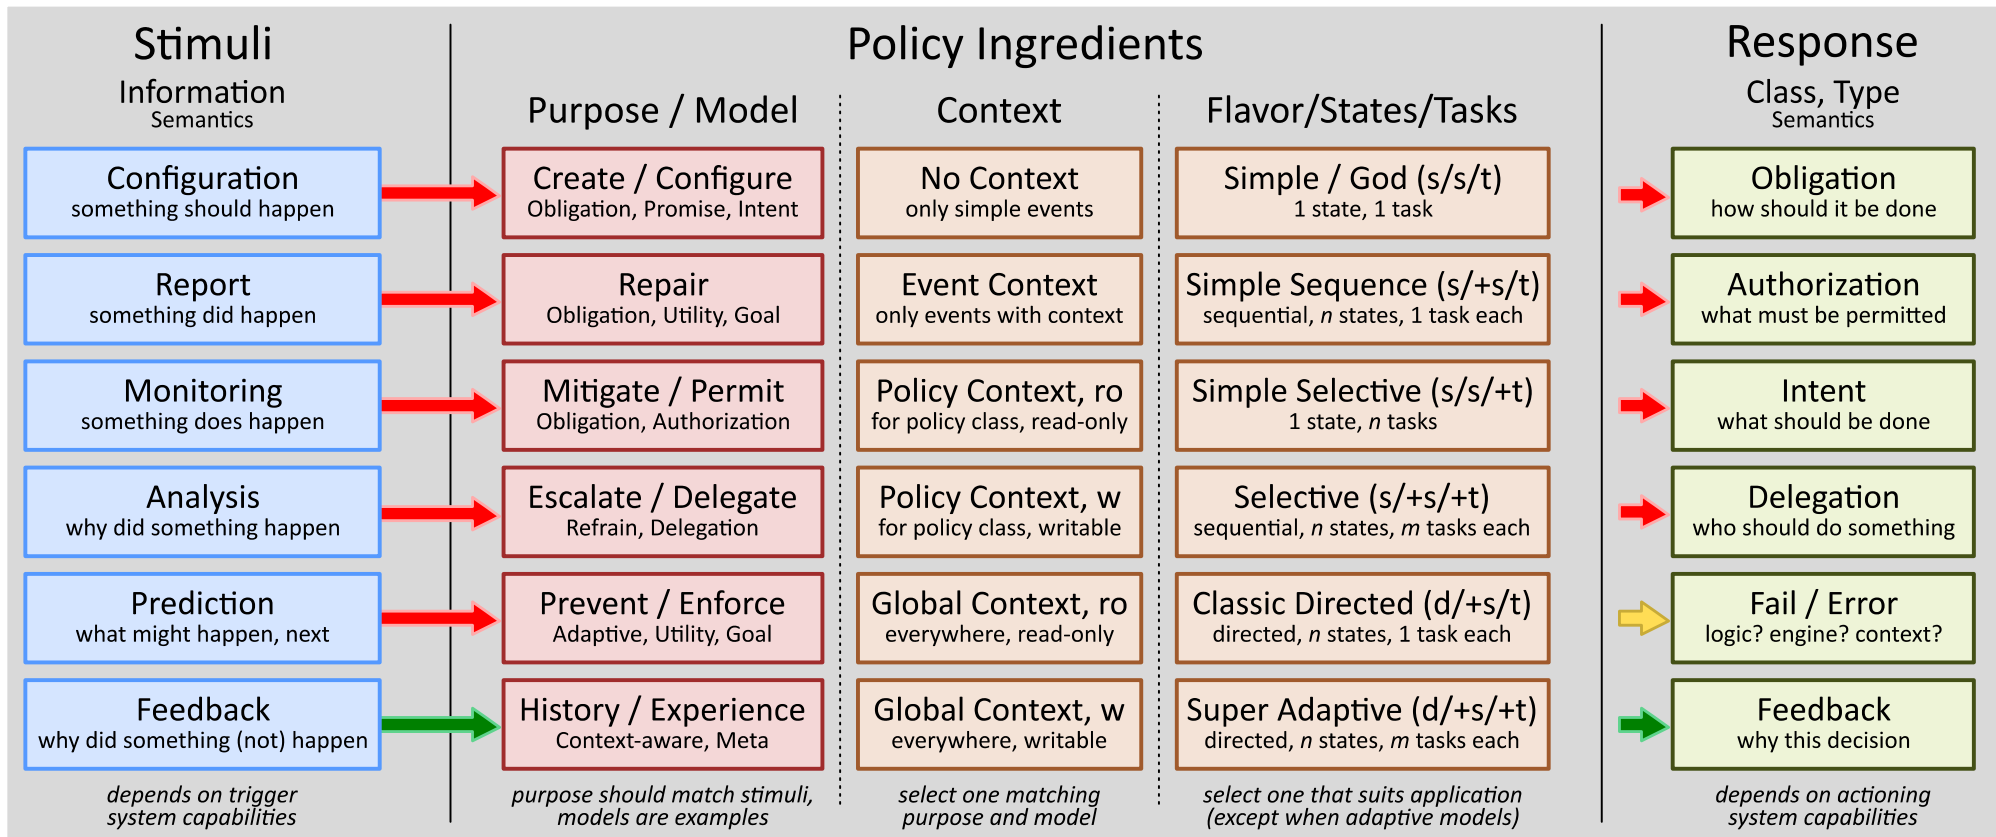
Unified Policy Model



Policy Transformations



Policy Matrix



APPLYING UNIFIED POLICY THEORY IN A SYSTEM



Concept & Implementation

Theory

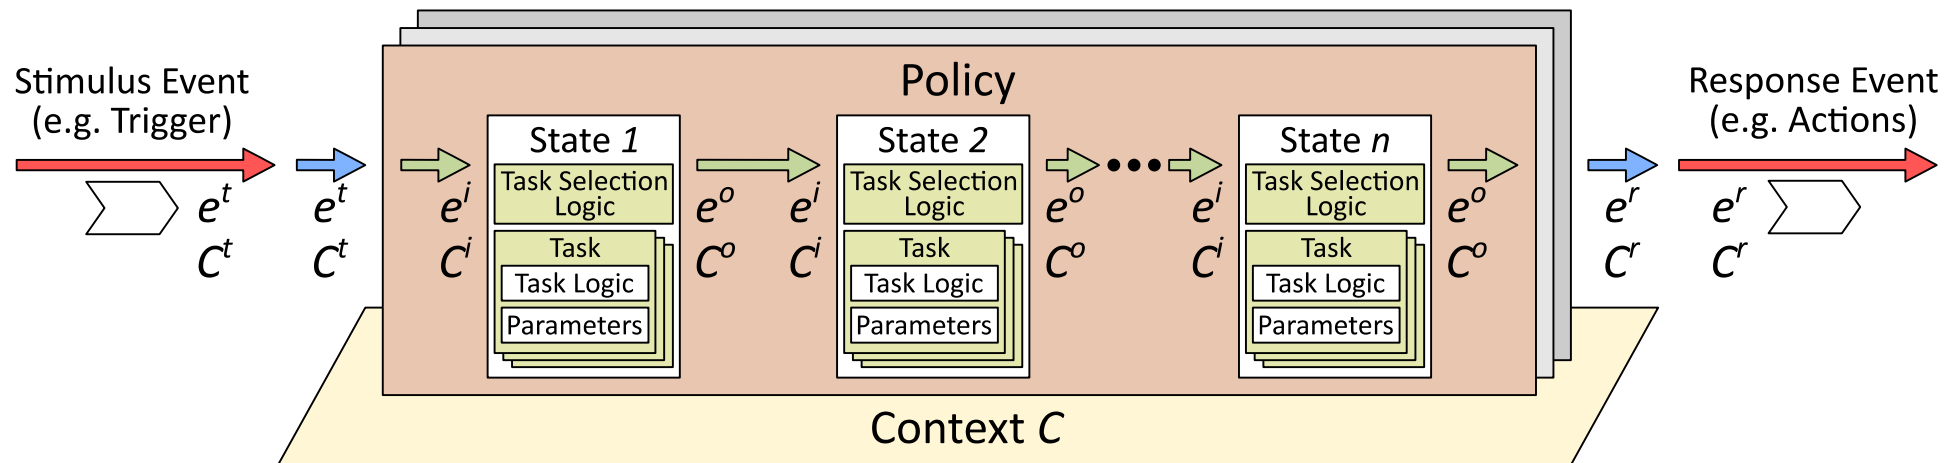
- Harmonize policy models
- Provide single execution environment
- Facilitate conflict processing
- Features
 - Context aware
 - Adaptive logic selection
 - Flexible clustering options
 - Flexible deployment options
 - Flexible policy deployment

Practice

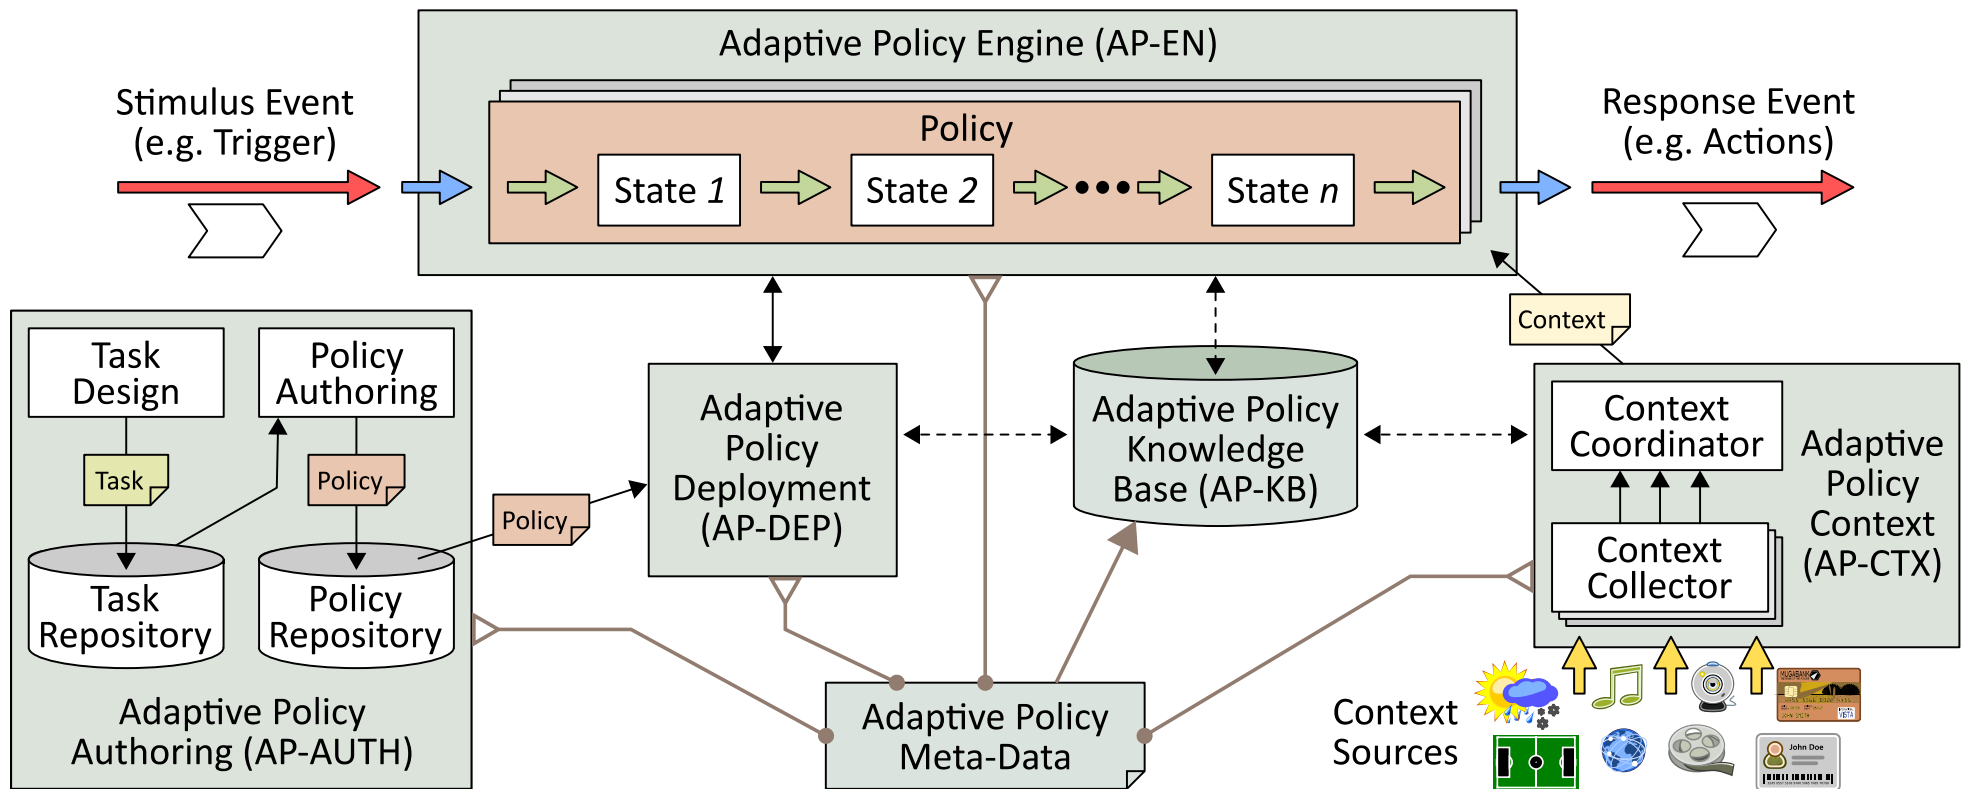
- An editor to create policy models
- An engine to run policy models
- Control of state and context
- Features
 - Context defined at run time using metadata
 - Logic loaded at run time
 - Policy Deployed as metadata
 - Policies/context distributed for scale

Adaptive Policy Engine: Event Flow & Context

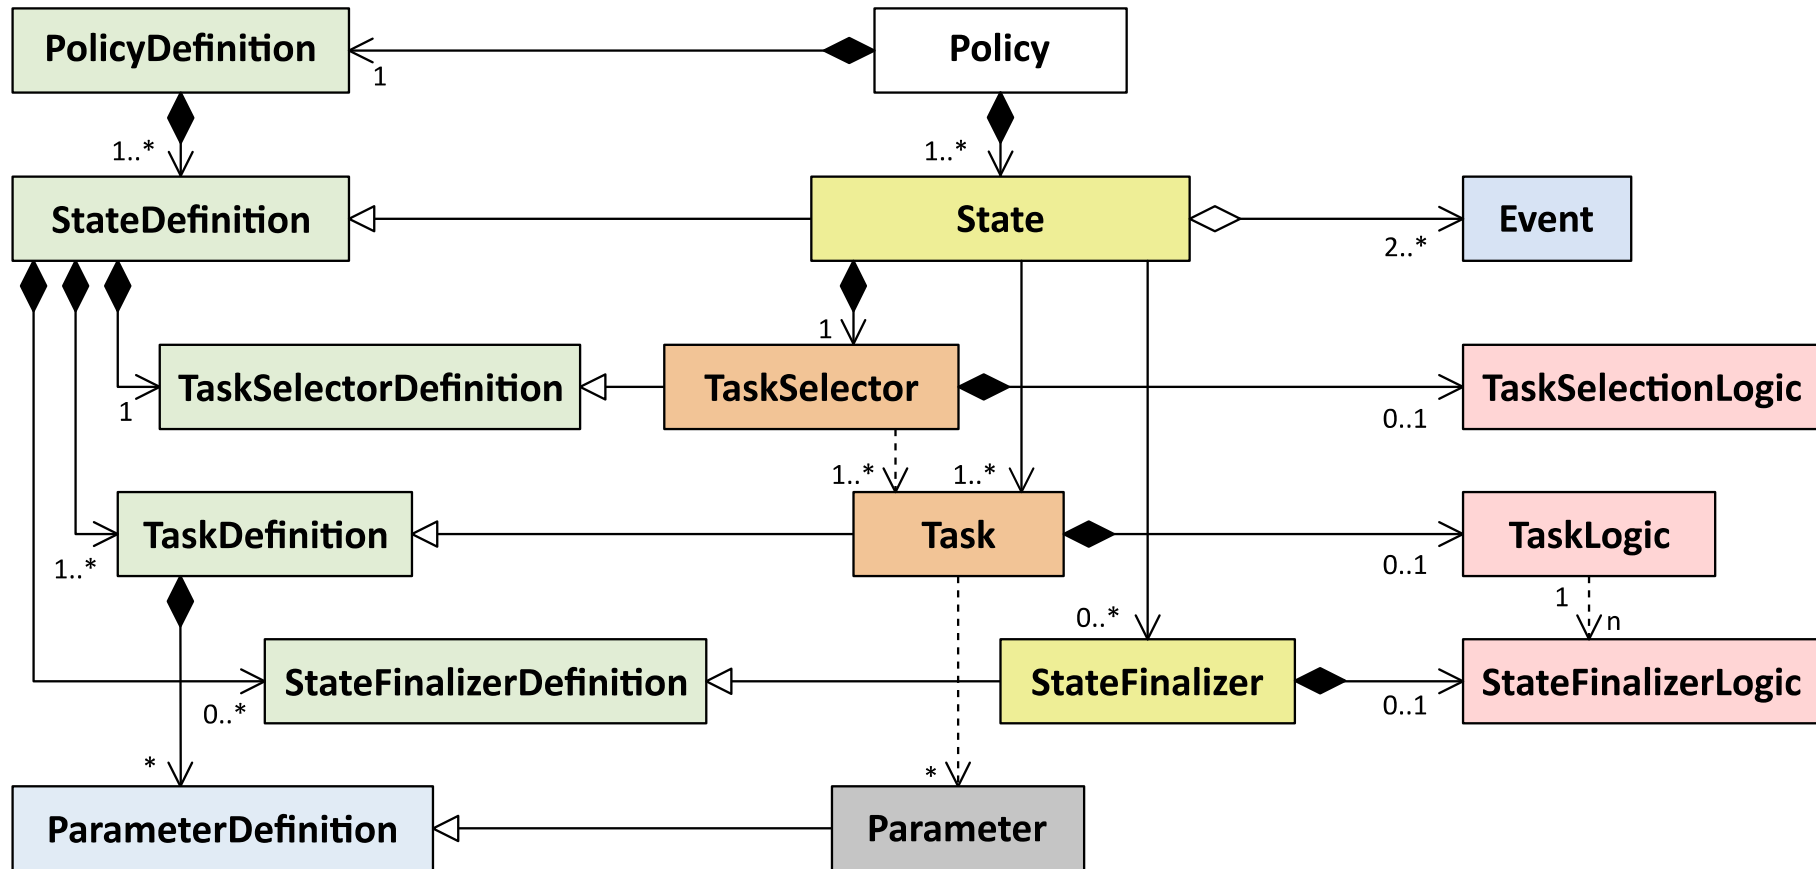
- Flow: **Trigger** → Engine → Policy → Engine → **Actioning System**
- Context: in all events, per policy type, global (r/w), external (r)



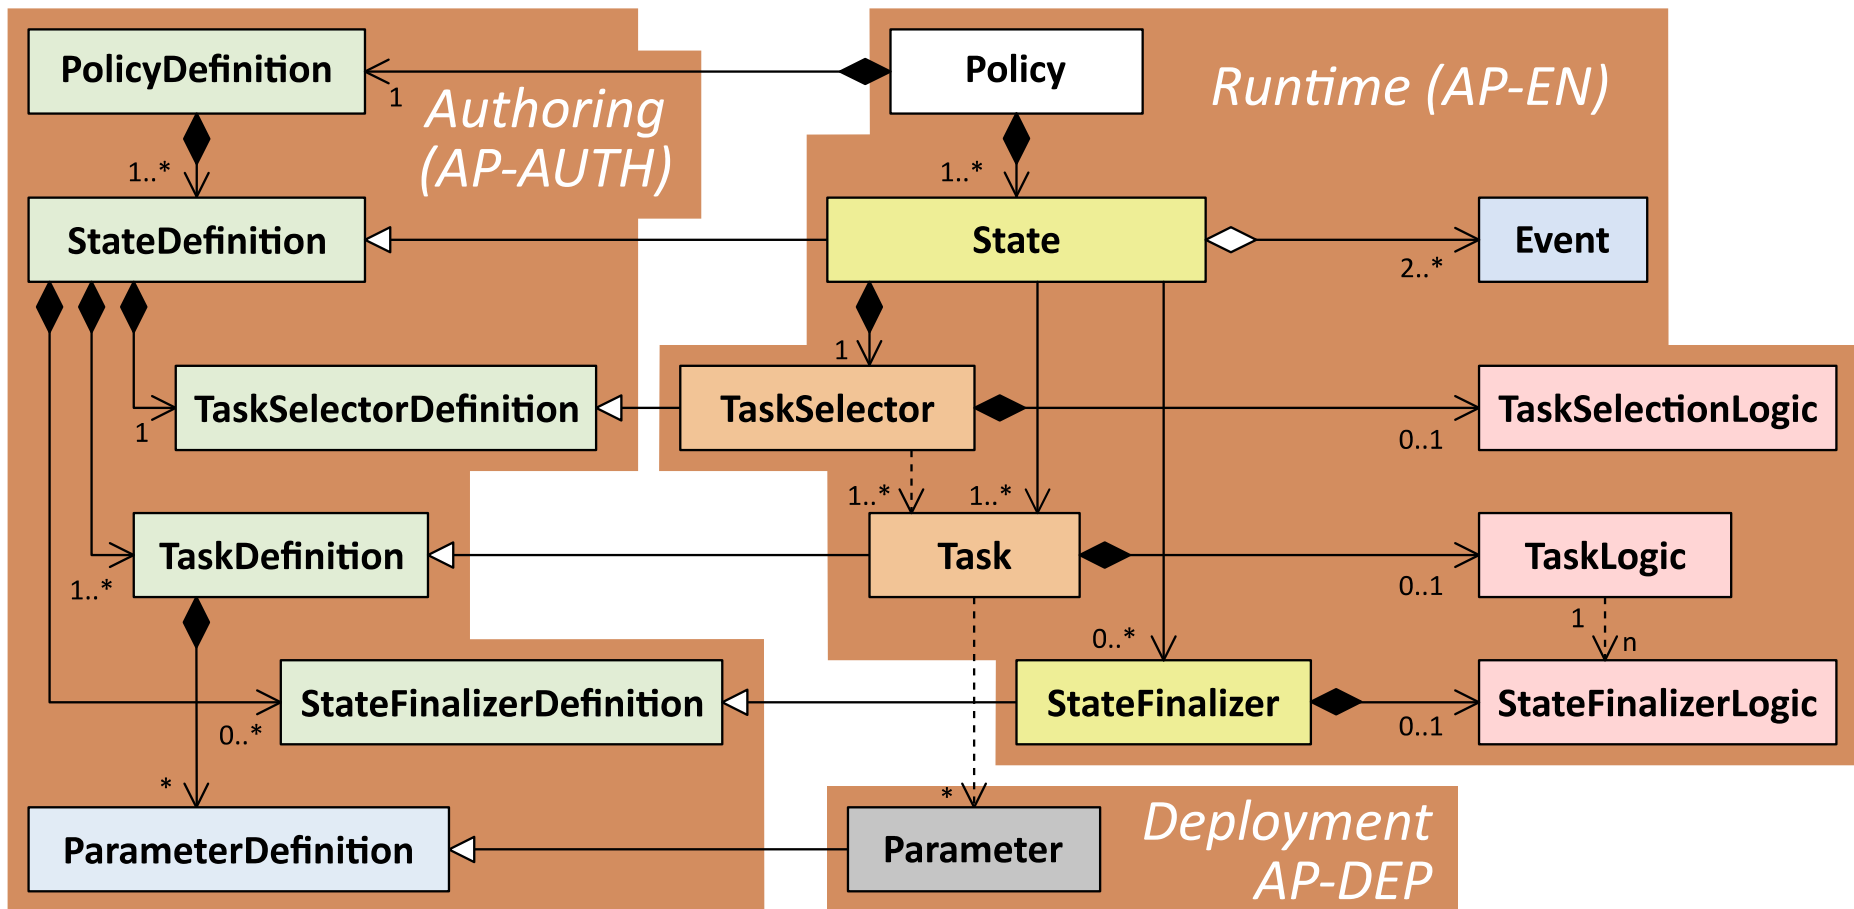
Policy Environment Components & Flow



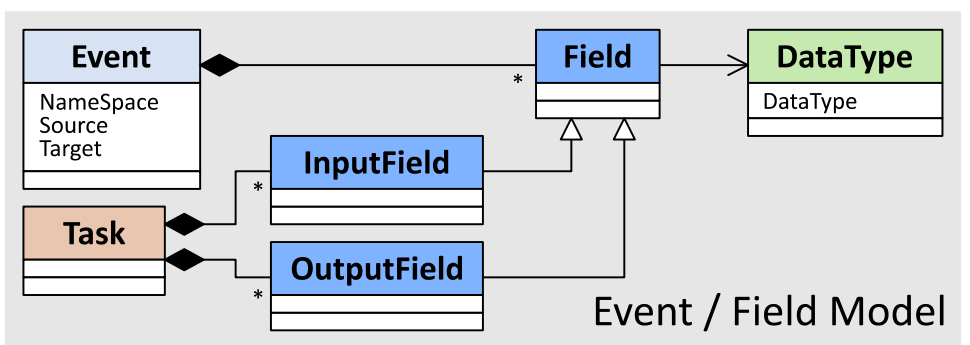
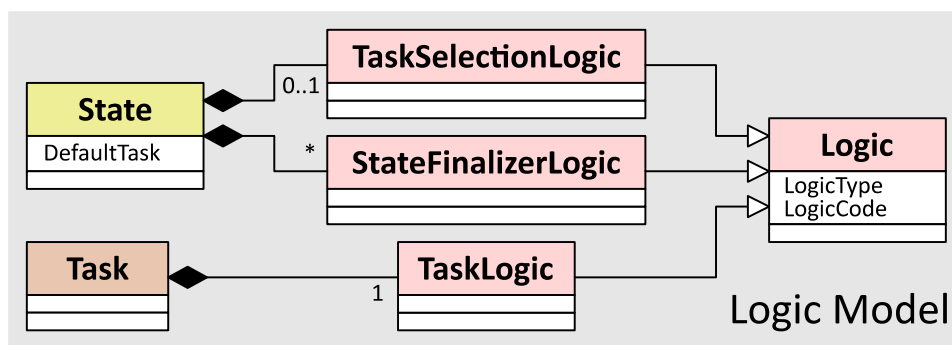
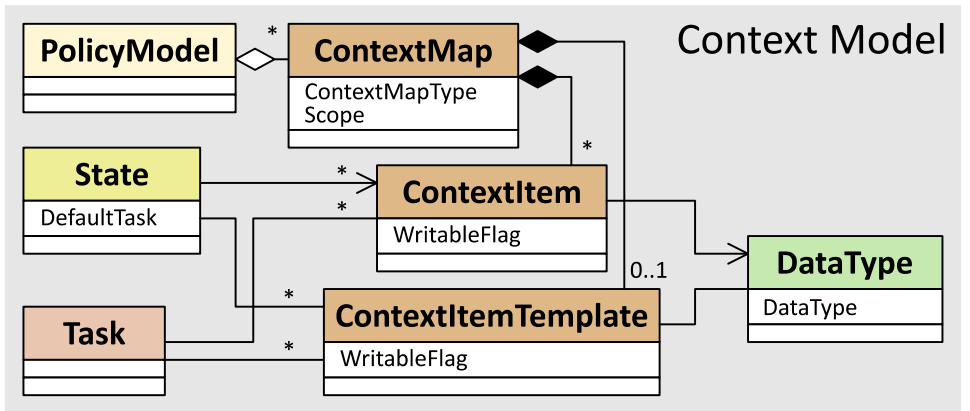
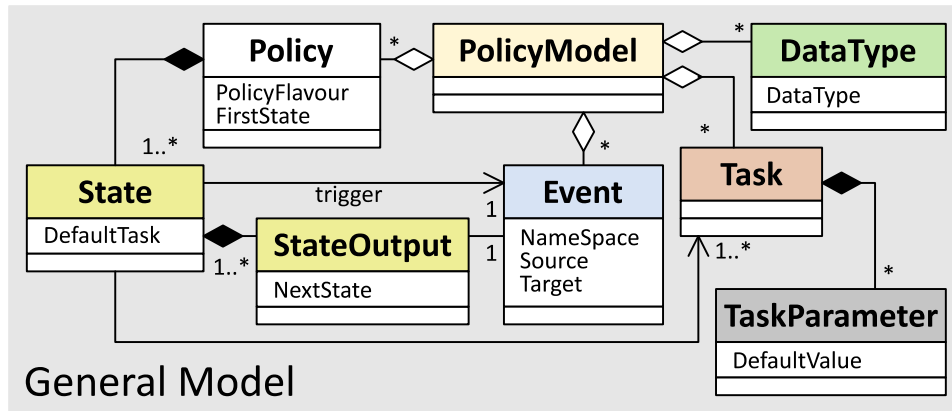
The Universal Policy Theory (UPT) Policy Model as UML



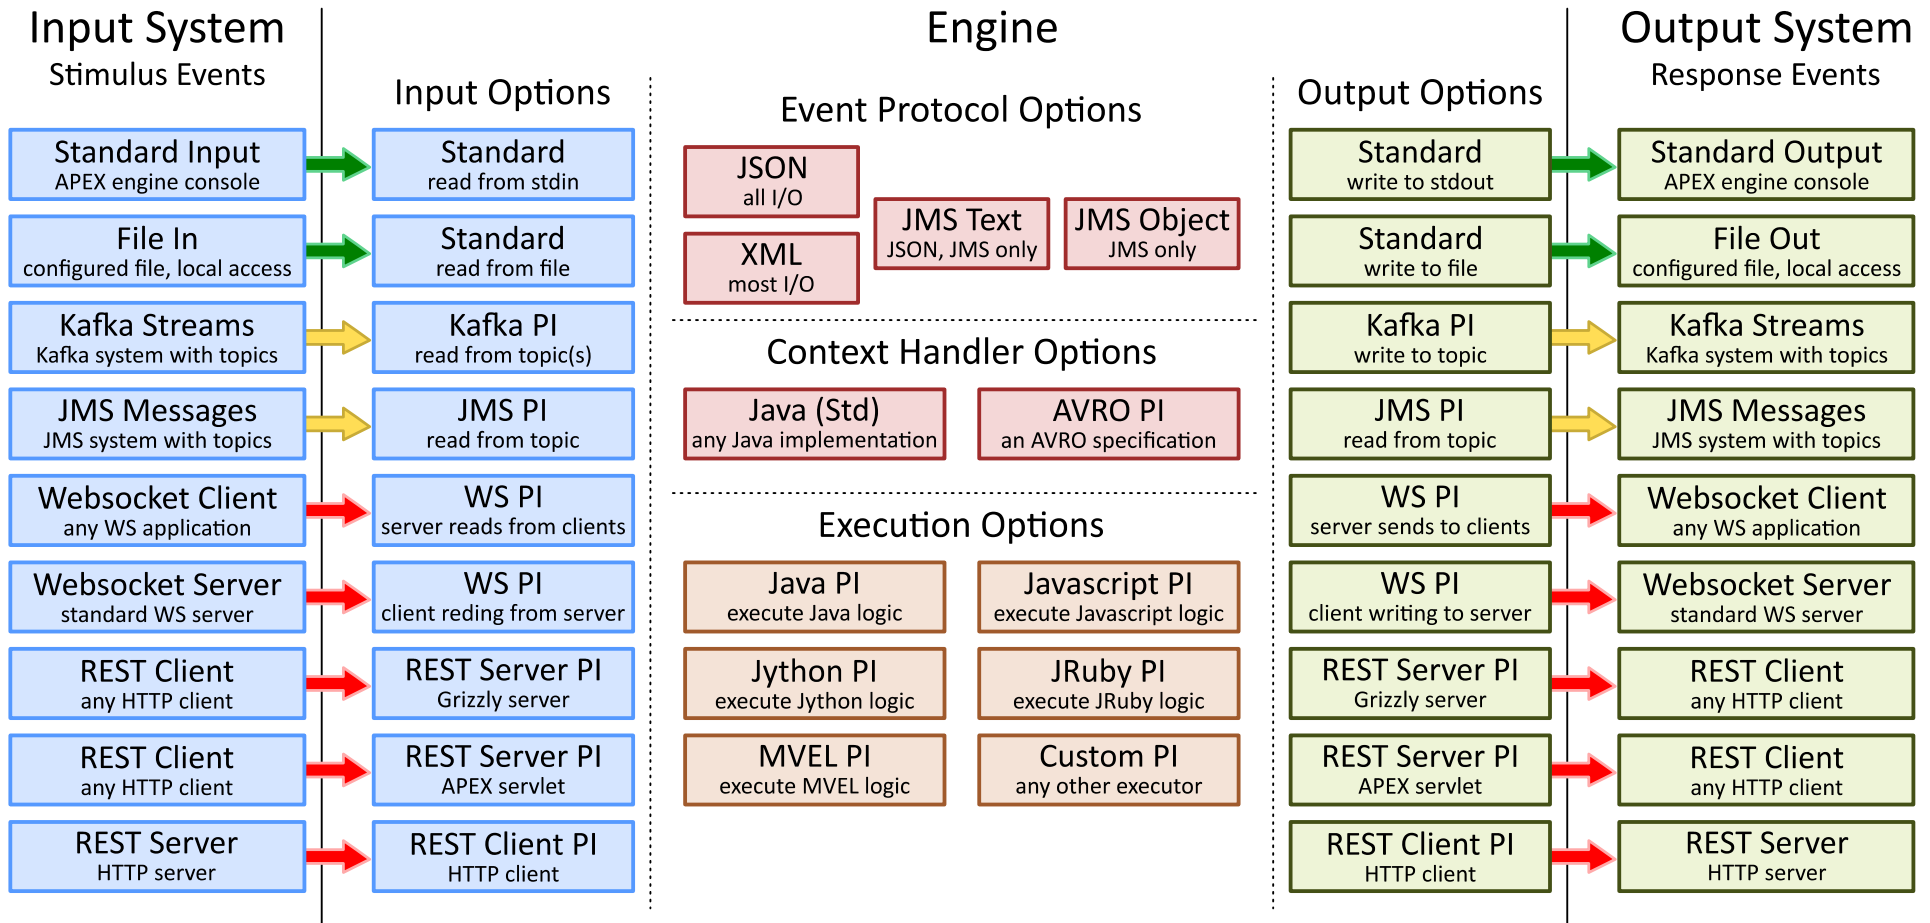
The UPT Policy Model mapped to a Policy System



Engine Execution Model



Engine Configuration and Plugins



DISTRIBUTED CONTEXT FOR POLICY



Distributed Context for Policy

- Policy work required context sharing across policy engines
- We wanted structured context just like what's available in management models
 - Think MIBs, Yang objects, UML classes, Java Beans, XML entities, JSON objects, ...
- We went looking for a model distribution system that
 - Provides distributed context (somehow classified information)
 - Supports locking
 - Supports monitoring
 - Supports persisting
- No such distributed context framework existed

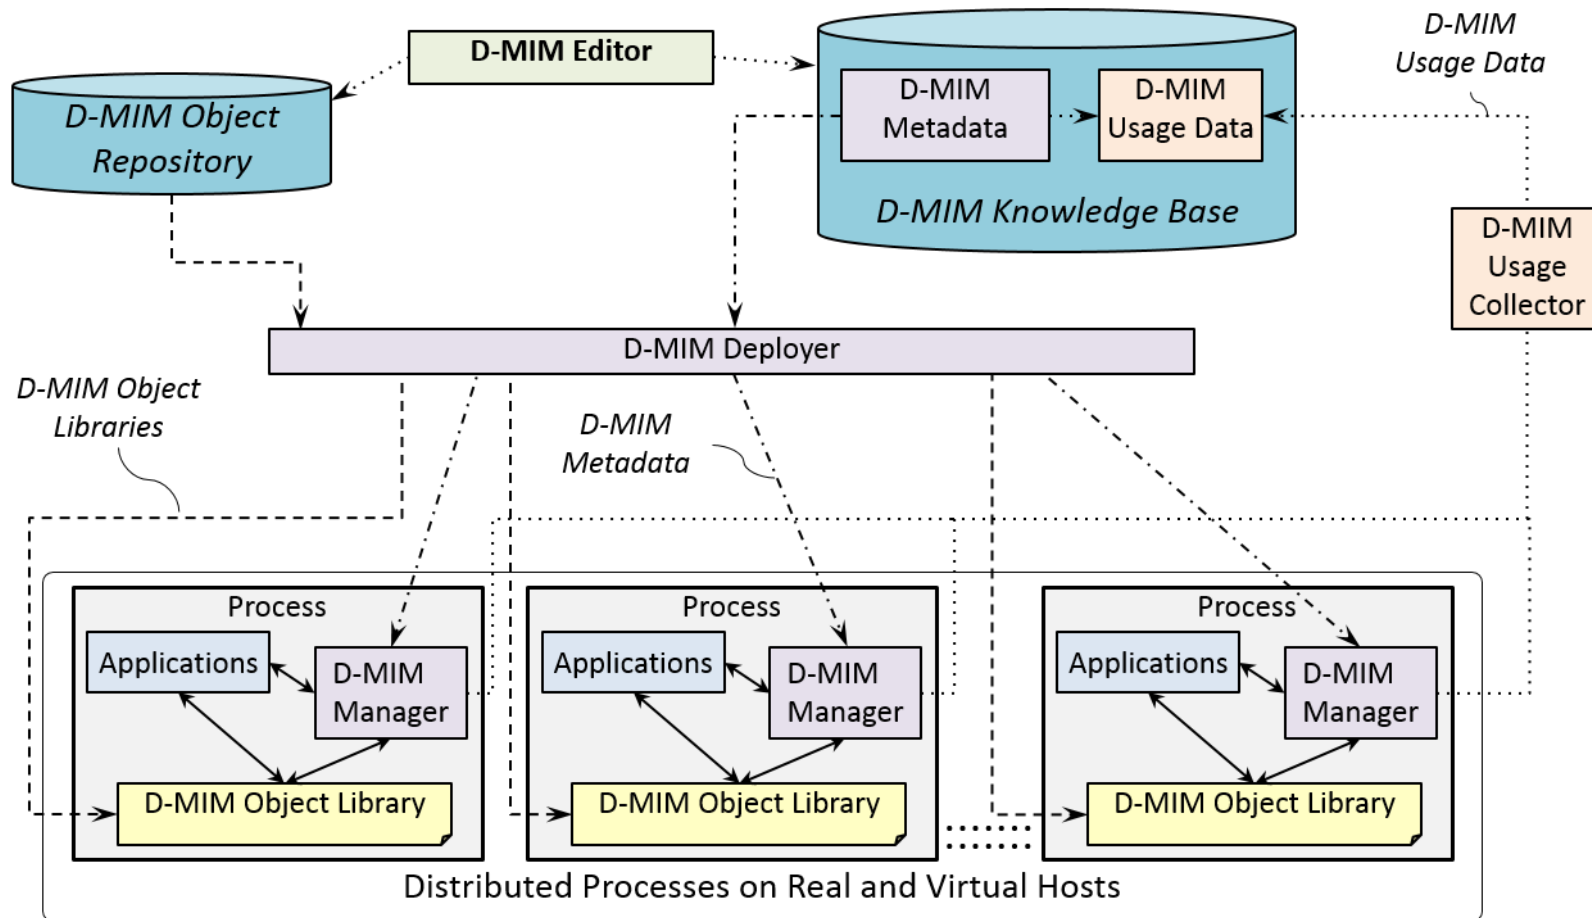
Distribution Frameworks

- Numerous frameworks for distributing unstructured hash maps
 - Distribution of maps of objects keyed by objects
 - E.g.: Hazelcast and Infinispan
- Some frameworks for locking
 - Transactional frameworks such as Narayana (Jboss JTA implementation) are slow
 - and very expensive in resource usage
 - No locking support on specific entries on distributed maps
 - No integrated persistence support
 - No integrated monitoring of CRUD operations on map entries

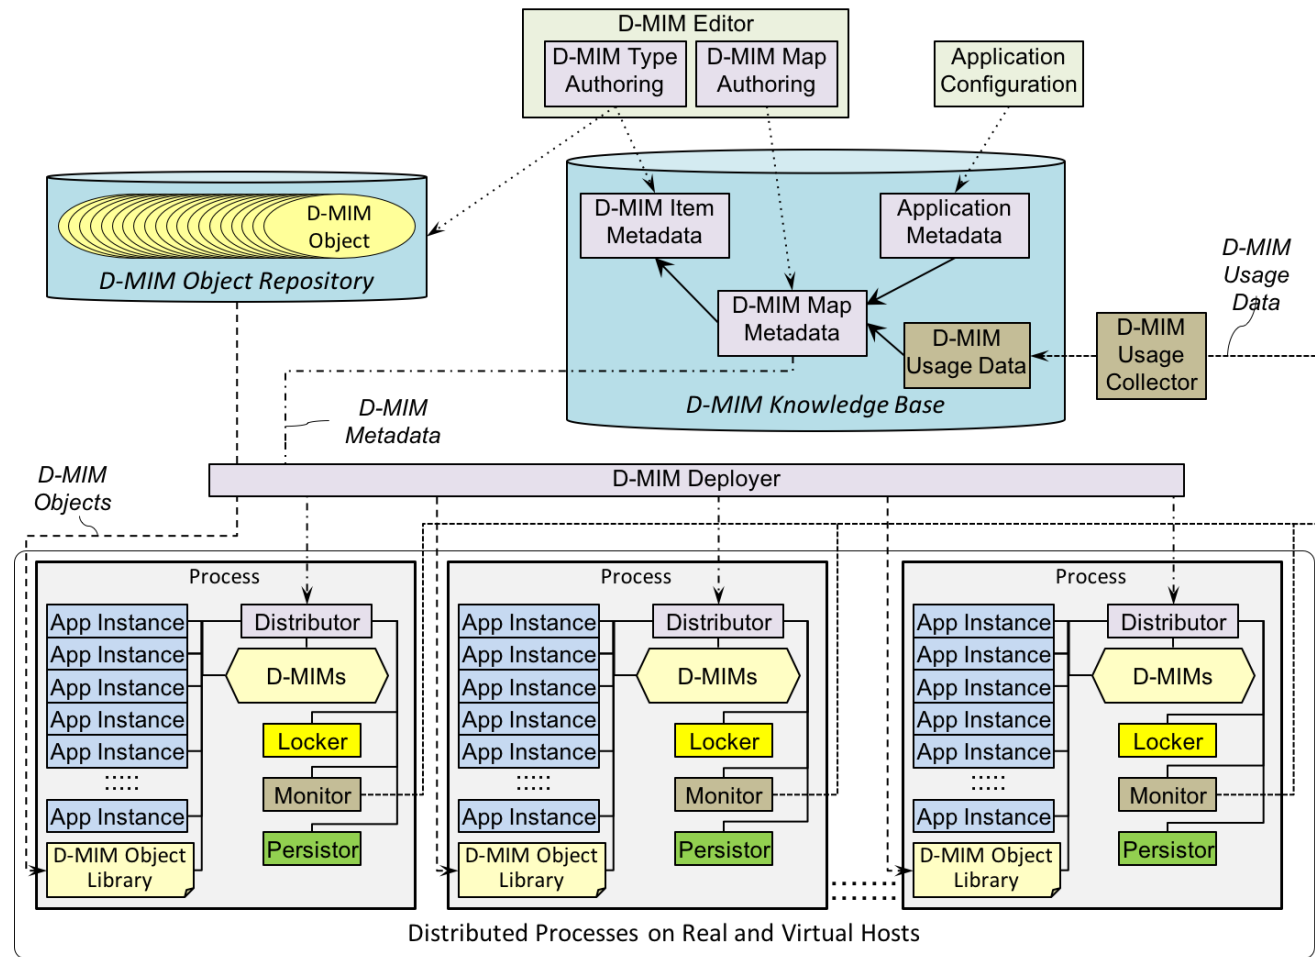
Distributed MIMs

- We decided to build a system for context that had strong structure support (just like MIMs)
 - Provides an interface to users to define highly structured context
 - Provides an interface to users to read, write, create, delete, persist, lock, and monitor context
 - Provides a plug-in architecture that allows existing distribution, locking, persistence and monitoring frameworks to be used
- D-MIM users can use distributed MIM maps
 - transparently on multiple processes, hosts, and geographic locations
 - Changes to one D-MIM copy are propagated to all others
 - Unified monitoring is supported
 - Unified locking is supported

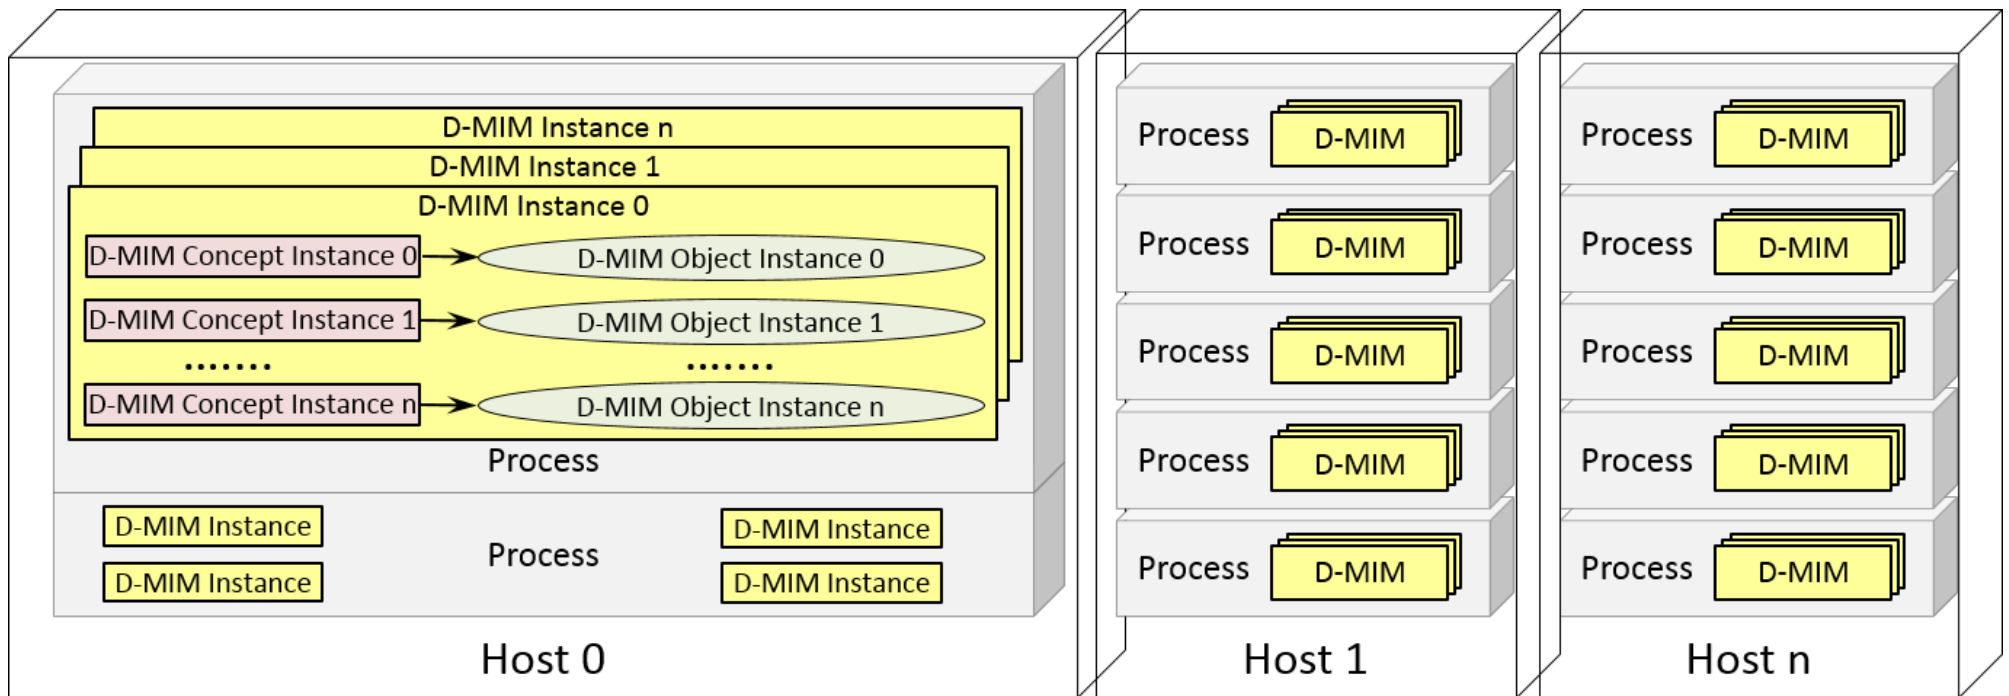
Distributed MIMs (D-MIMs) in Policy Engines



D-MIMs in a Distributed System



Distributed MIM instances on Hosts



APPLYING TO ONAP



UPT and UPEE in ONAP

- A model for policies and policy engines
 - Drools and beyond Drools, state handling
 - XACML engine, stateless
- UPM model distribution using ONAP Policy Framework
- D-MIMs and Context in Drools/plugin for Drools?
- Editor integration for policy authoring
- Context and conflict
 - Design time
 - Deployment time
 - Runtime identification
 - Runtime mitigation



ERICSSON

ADDENDUM: SOME POLICY TERMINOLOGY



Policy & Engine

- *Policy* is an artefact that governs the choices in behaviour of a system
 - Separation of mechanism from policy
 - Has capabilities, defined in a specification, explicit / implicit trigger
 - Has multiple, partially relative, dimensions
 - Example for choices: in Network Management choices are operations on managed objects
- A *Policy Engine* is responsible for executing policies
 - Receiving triggers
 - Execute relevant policies
 - Receive actions from executed policies
 - Return / forward them
 - With all “-ities: scalability, performance, security, ...

System & Application

- A *Policy System* controls and manages life cycle of policies
 - Functional and non-functional capabilities
 - Life cycle
 - Authoring
 - Deployment
 - Execution (using the engine)
- A *Policy Application* realizes a policy system
 - Builds / implements functional and non-functional capabilities

Policy Variants

Context-aware, Adaptive & Adaptable

Context-aware Policy

- makes different decisions based on context information
- static decision-making behavior
- *Fewer policies:* same trigger, different context

Policies are more flexible

Adaptable Policy

- Can change its decision making behavior
- Based on an external activity
 - outside the policy
- *Fewer policies:* same policy, multiple behaviors

Adaptive Policy

- Can change its decision making behavior
- Based on an internal activity
 - inside the policy
- *Fewer policies:* same policy, multiple behaviors
- *Policy can adapt* to target shifts

Policy Variants

Context-aware, Adaptive & Adaptable

Context-aware Policy

- Different trigger context results in different situations
- External context as part of situation or decision making can change decision
- Not understood context might signal shift in automation target

Adaptable Policy

- Non-context-aware
 - Set policy parameter
 - Set policy state logic
- Context-aware (trigger/external)
 - Based on context, set policy parameters and/or use different policy state logic

Adaptive Policy

- Change on automation target resulting in new/altered context
- Set policy parameter/state logic due to policy internal context
- Change state logic



ERICSSON