

FEATURE: PERFORMANCE ISSUES IN DEVOPS

# Performance Issues? Hey DevOps, Mind the Uncertainty!

Catia Trubiani, Gran Sasso Science Institute

Pooyan Jamshidi, University of South Carolina

Jurgen Cito, Massachusetts Institute of Technology

Weiye Shang, Concordia University

Zhen Ming Jiang, York University

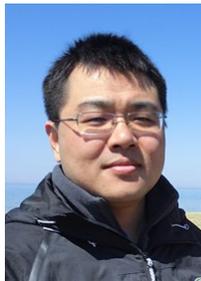
Markus Borg, RISE Research Institutes of Sweden AB

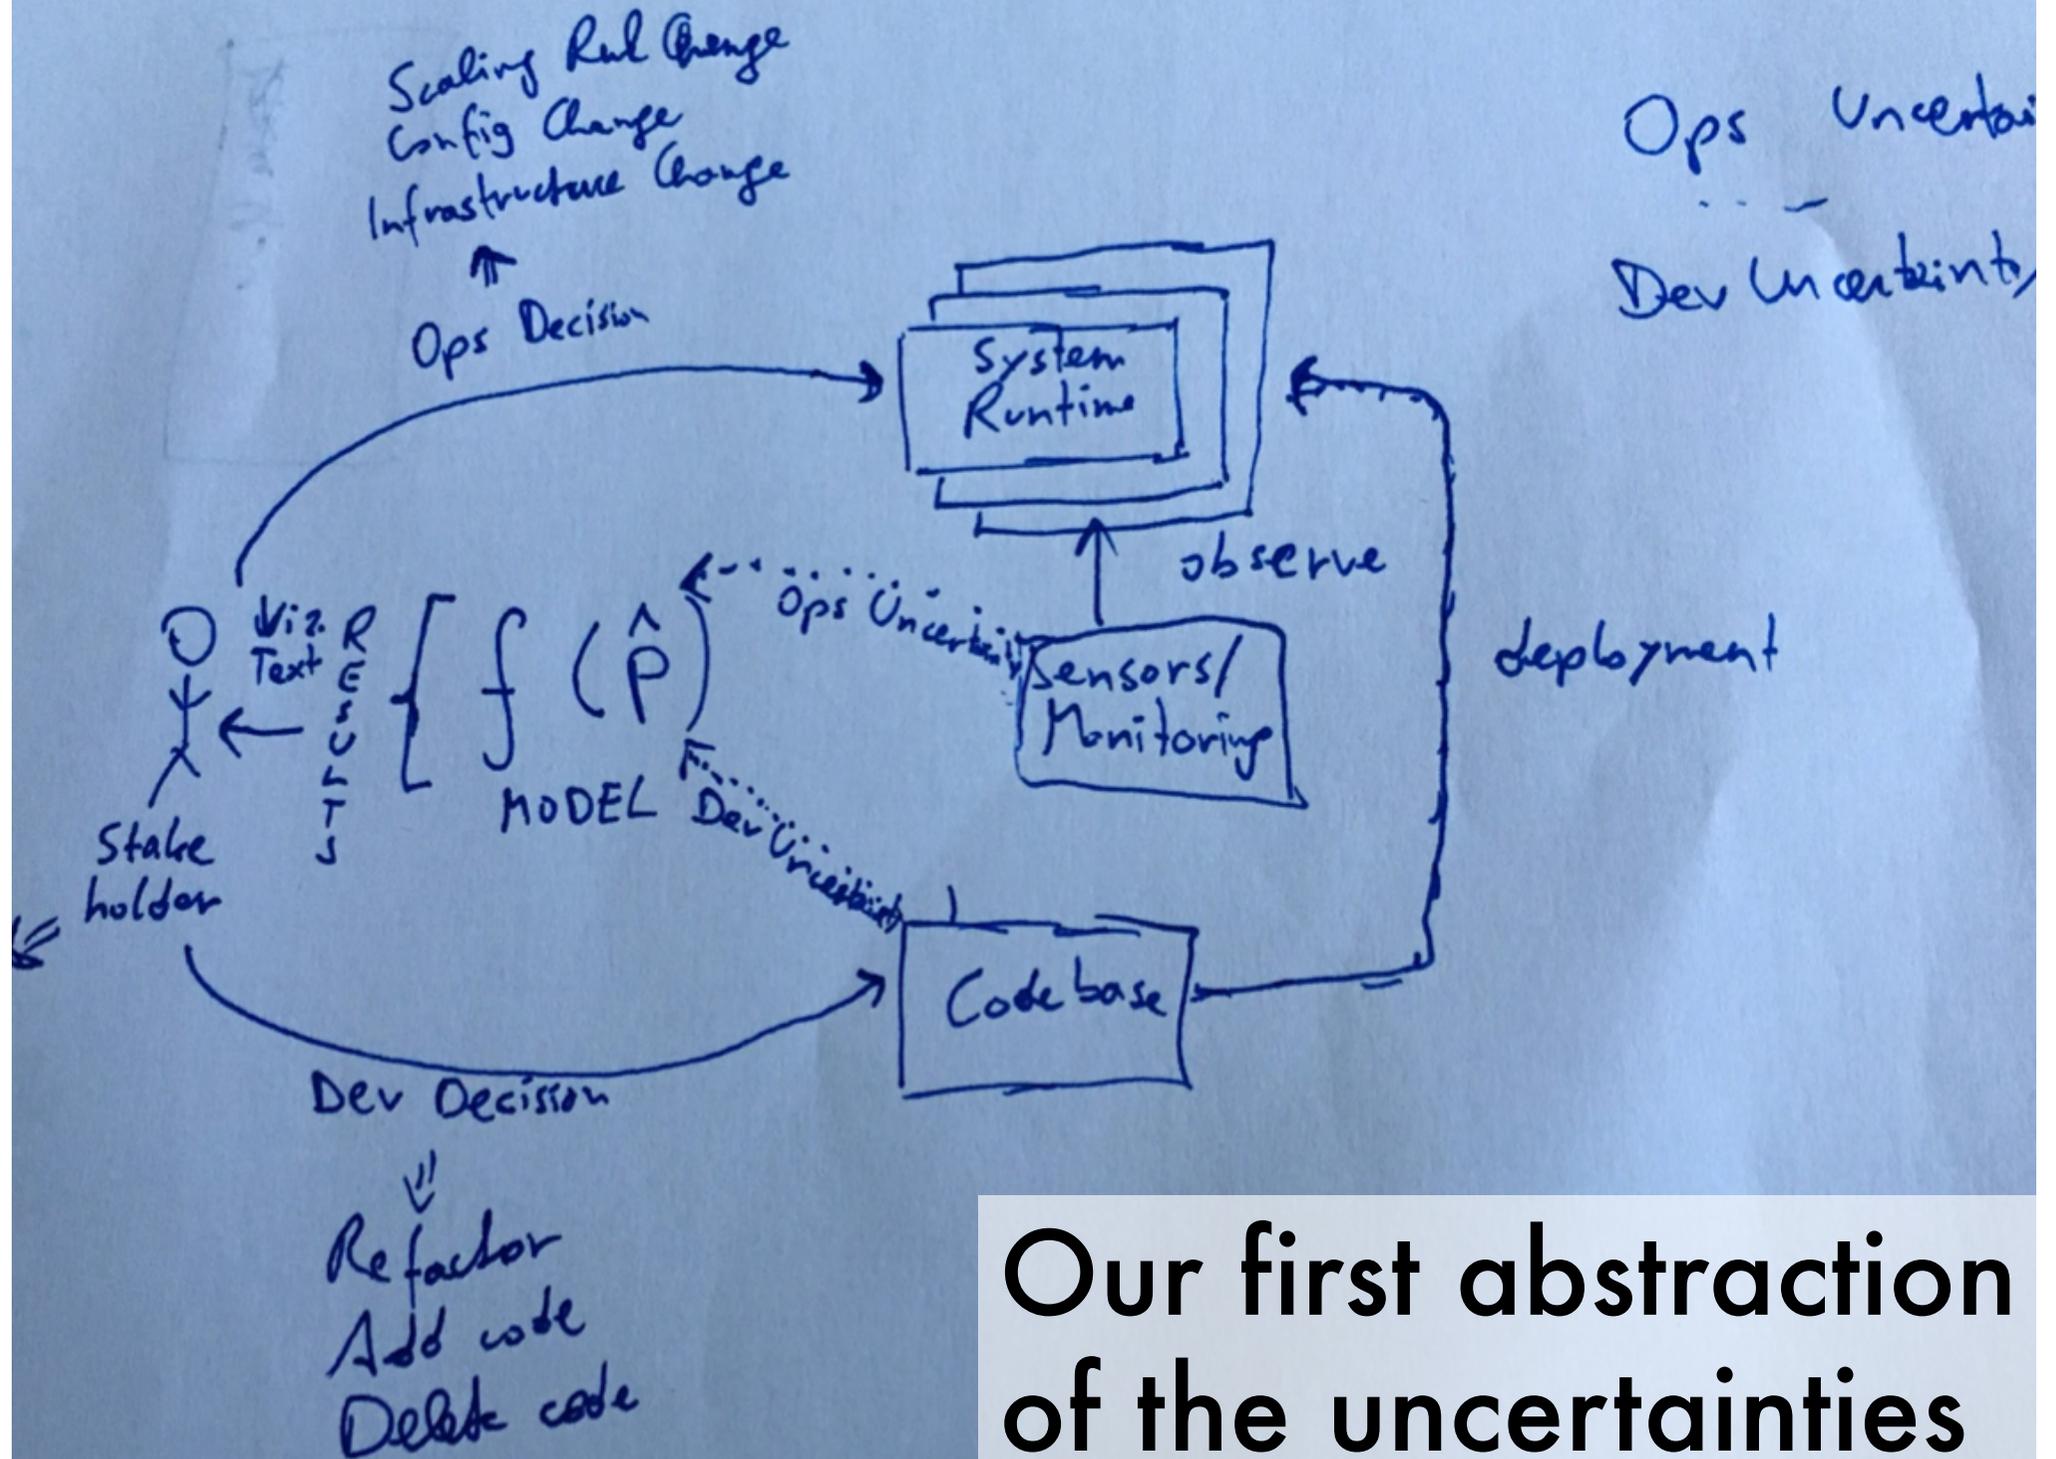


# September 25 – 30 , 2016, GI-Dagstuhl Seminar 16394 Software Performance Engineering in the DevOps World

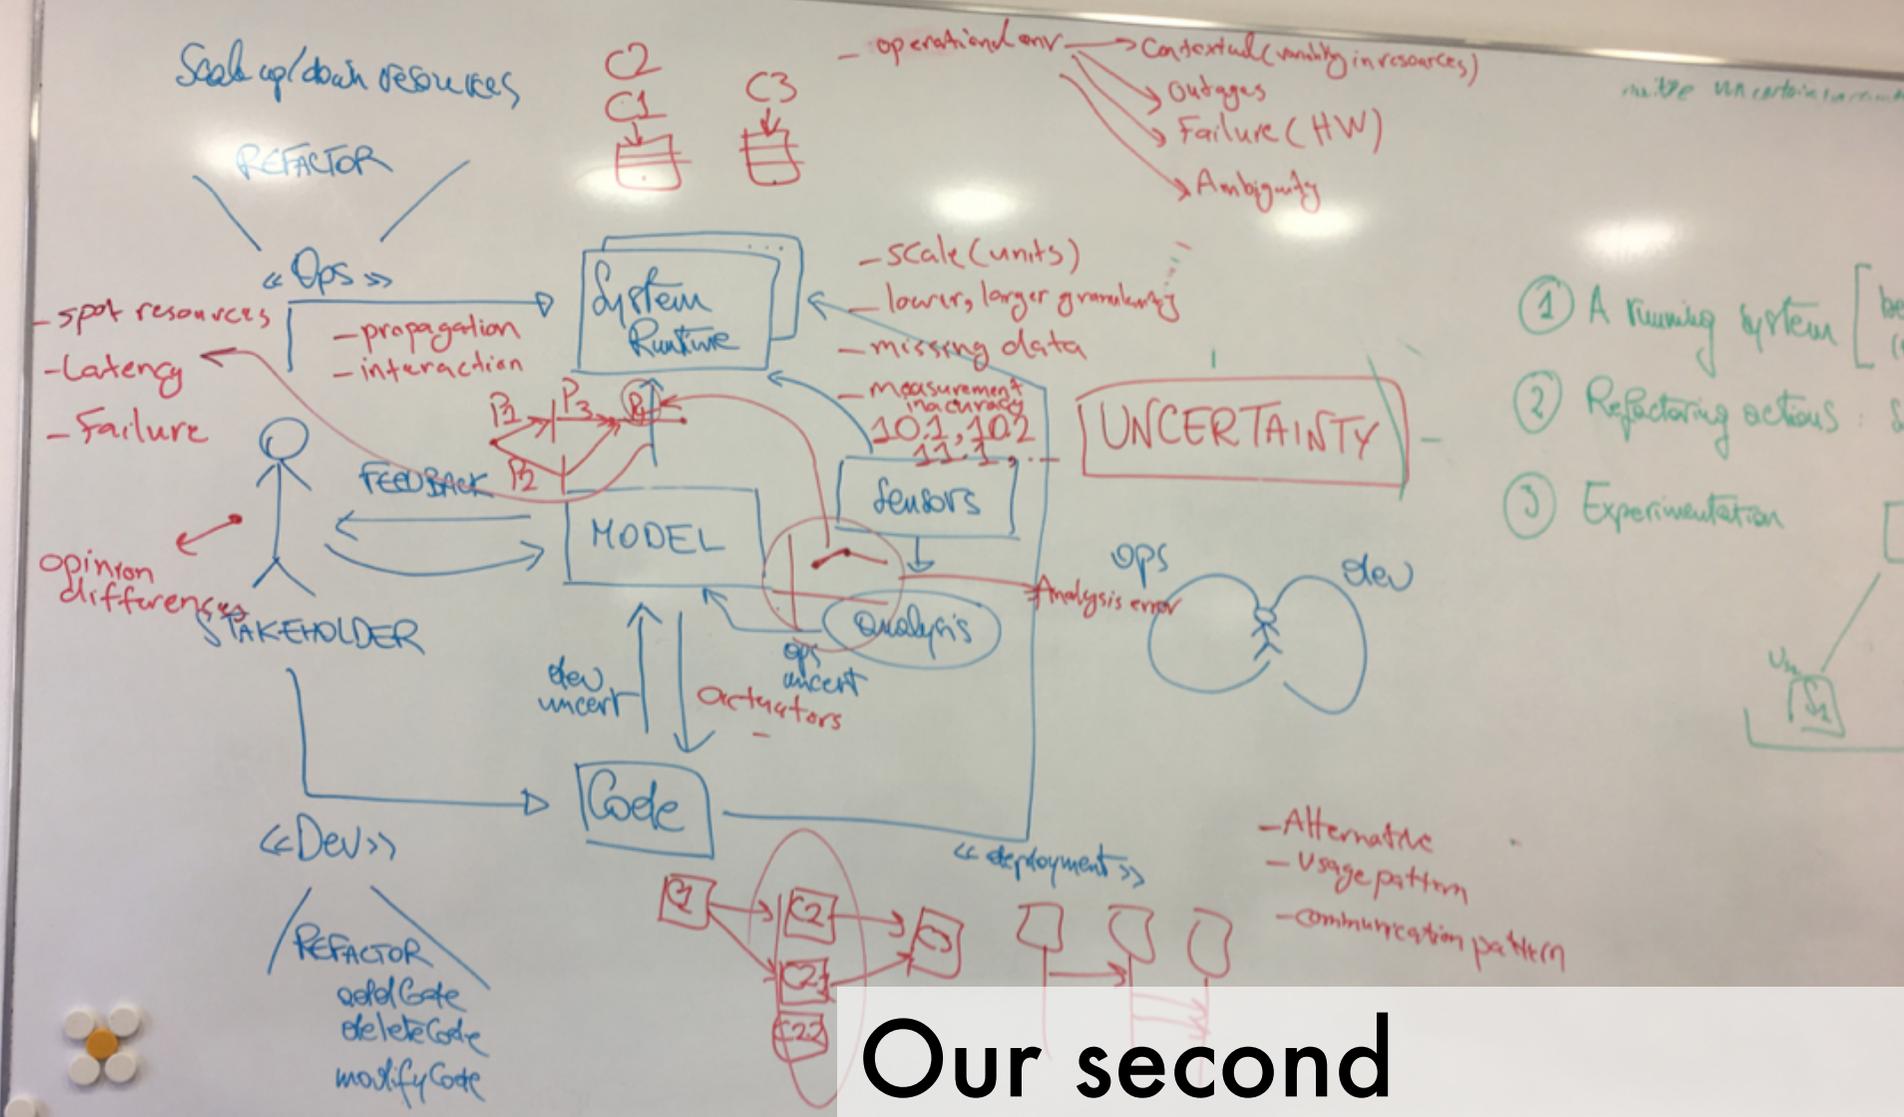


# Sources of Uncertainty in Performance-aware DevOps



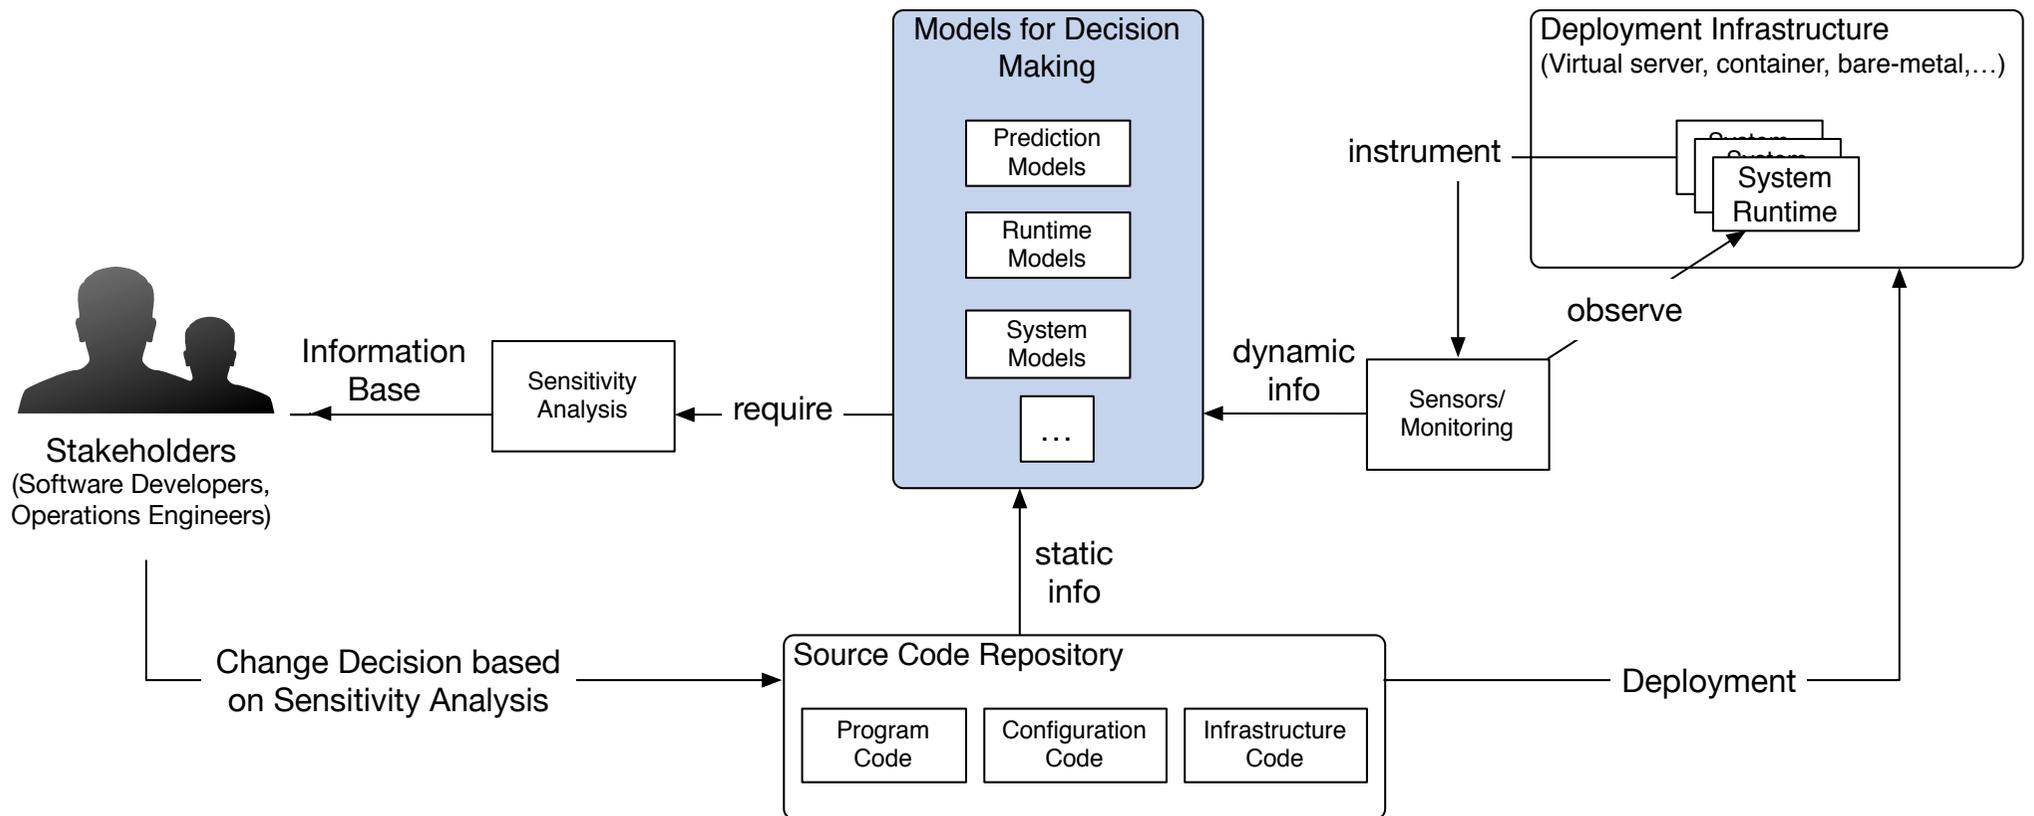


Our first abstraction of the uncertainties

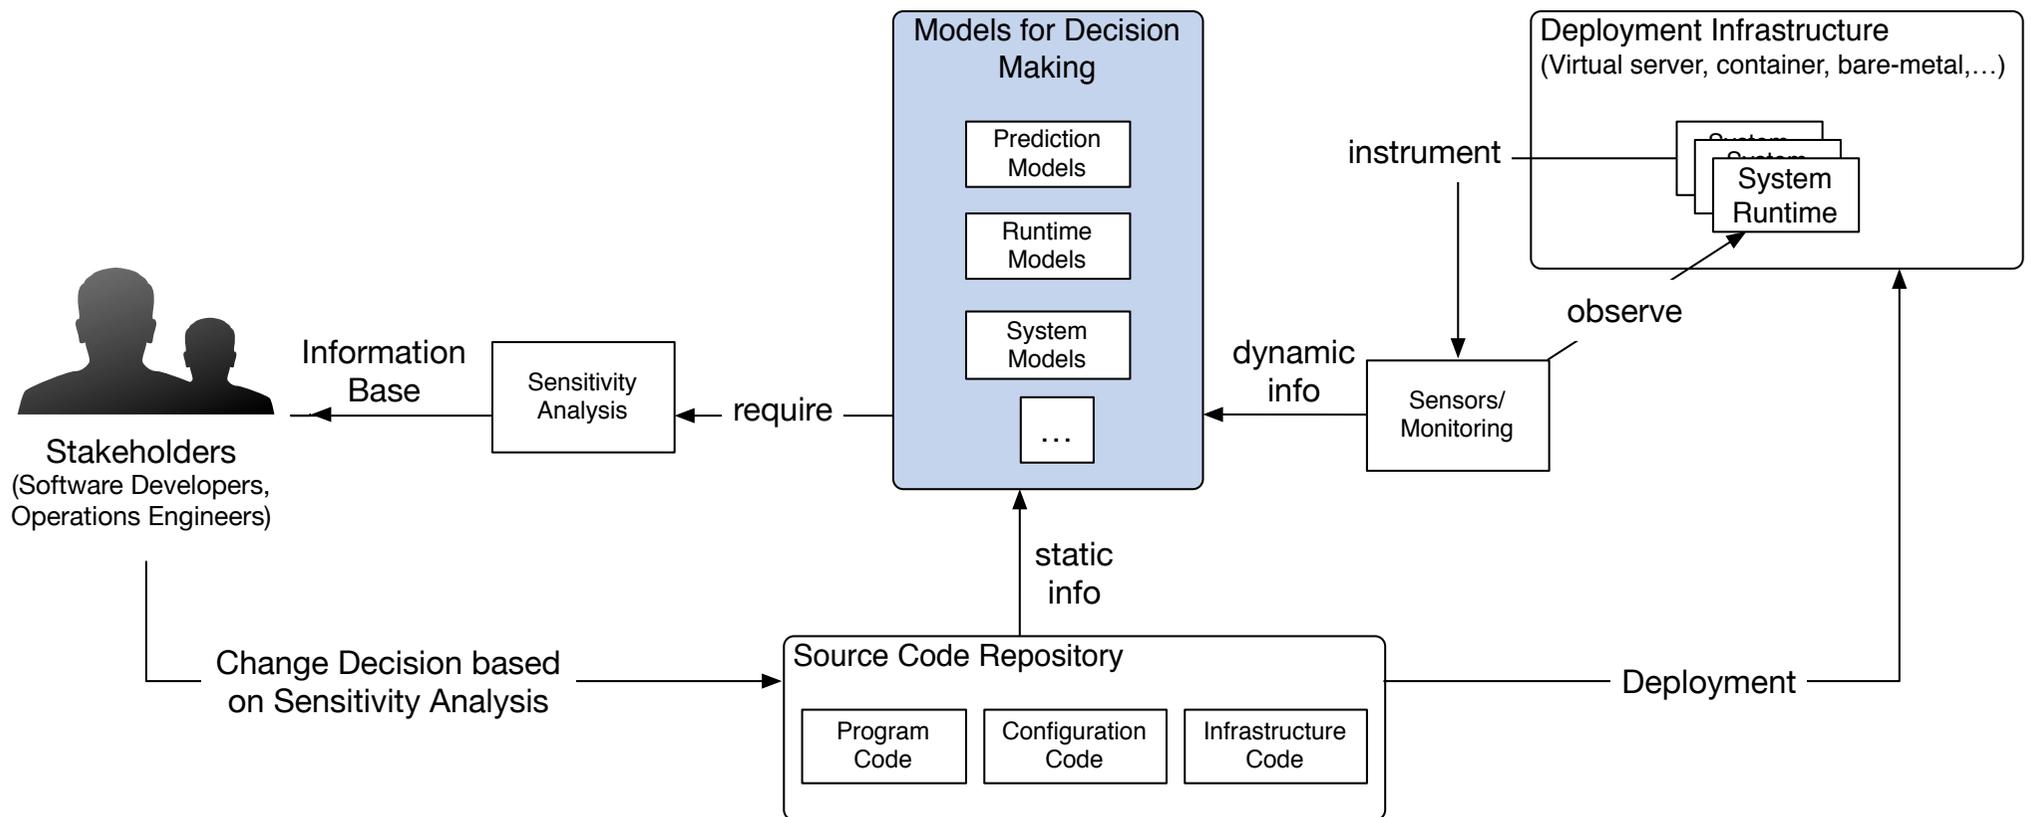


**Our second abstraction of the uncertainties**

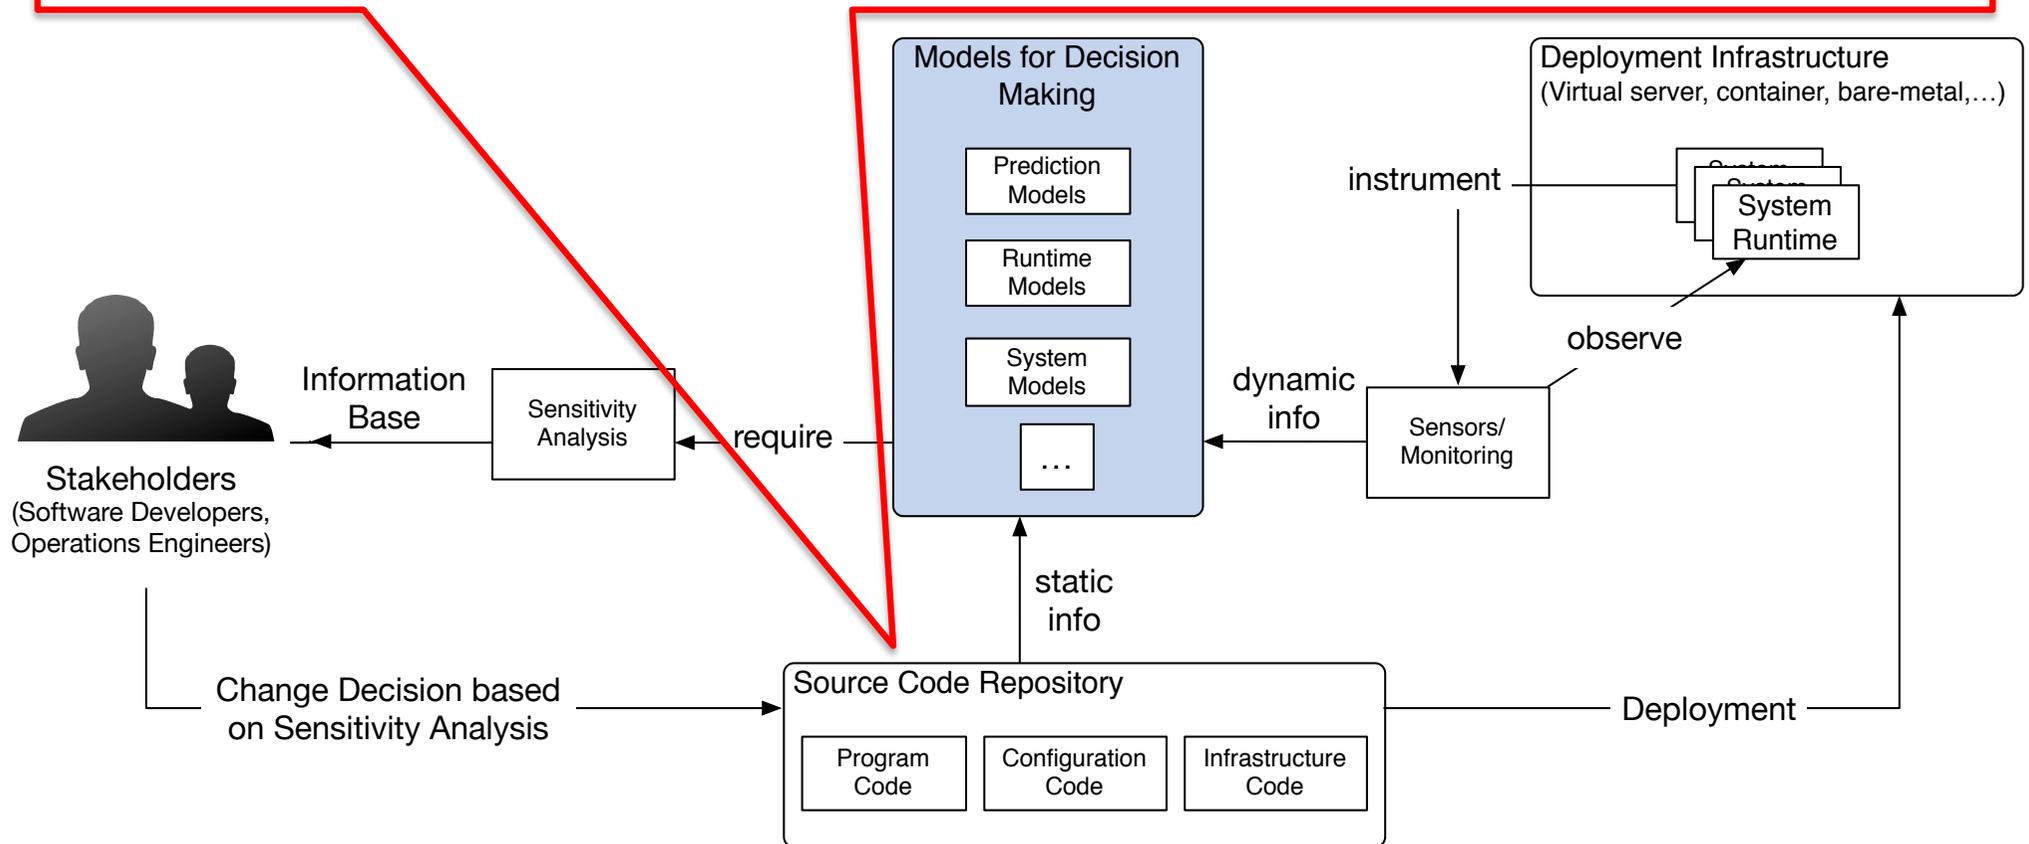
# Our final abstraction of the uncertainties



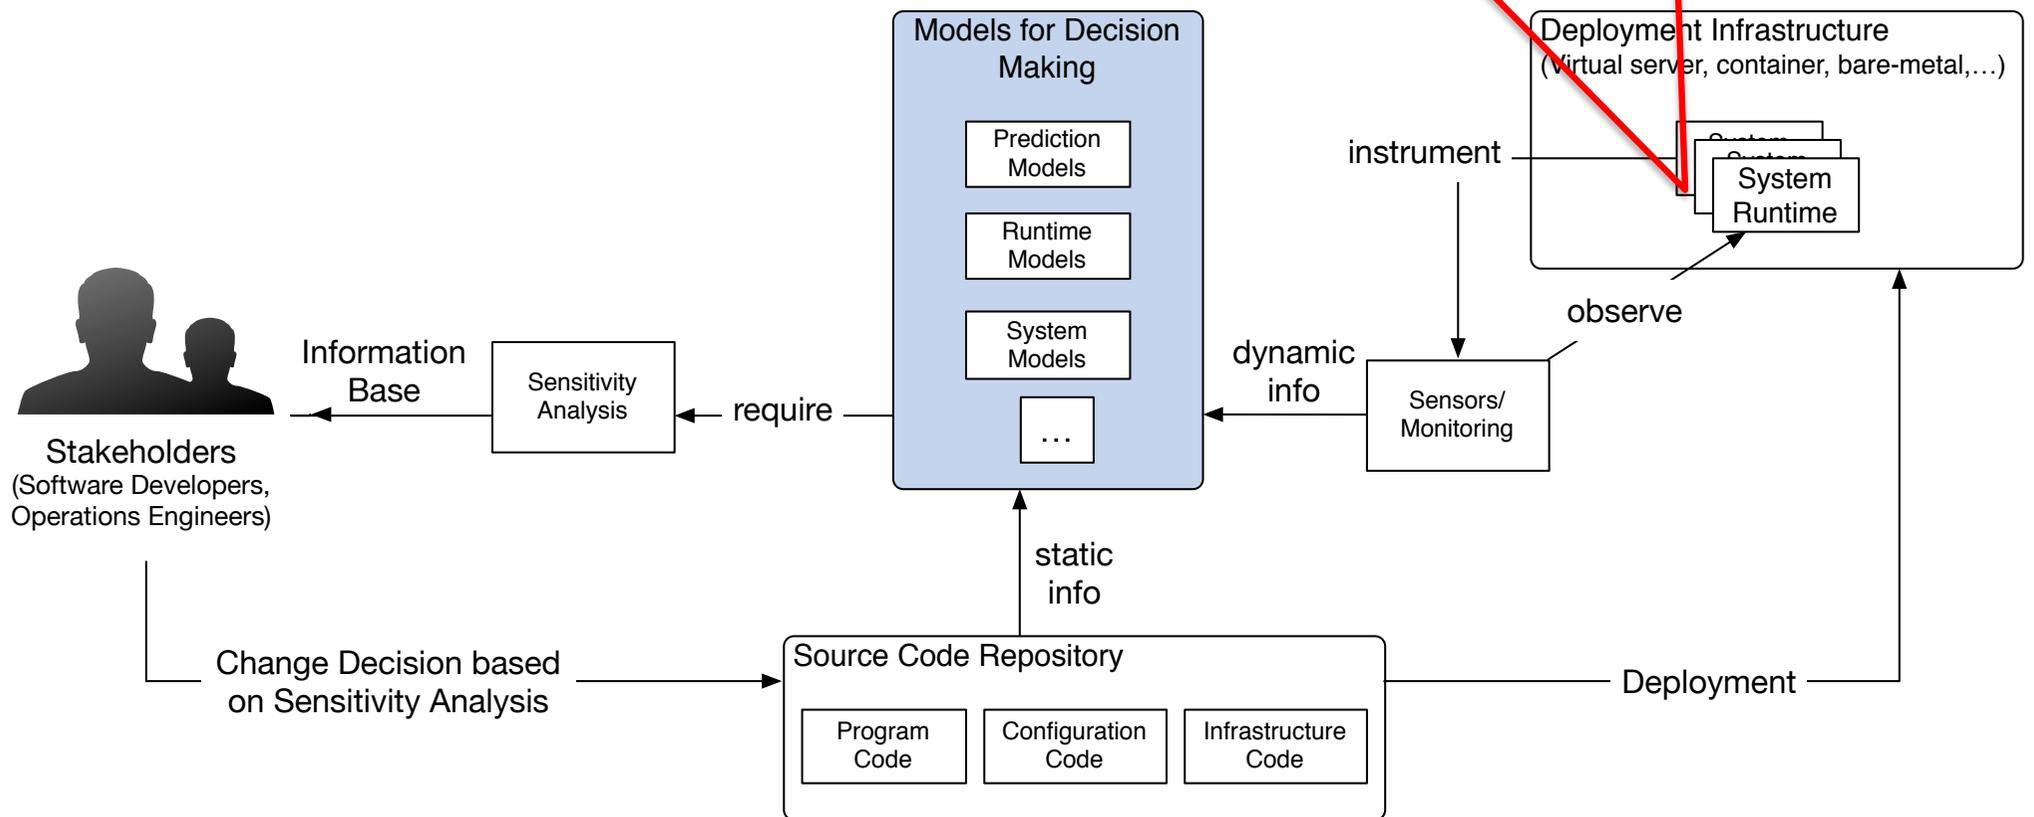
# Deployment infrastructure (DI): Physical or virtual, type of nodes...



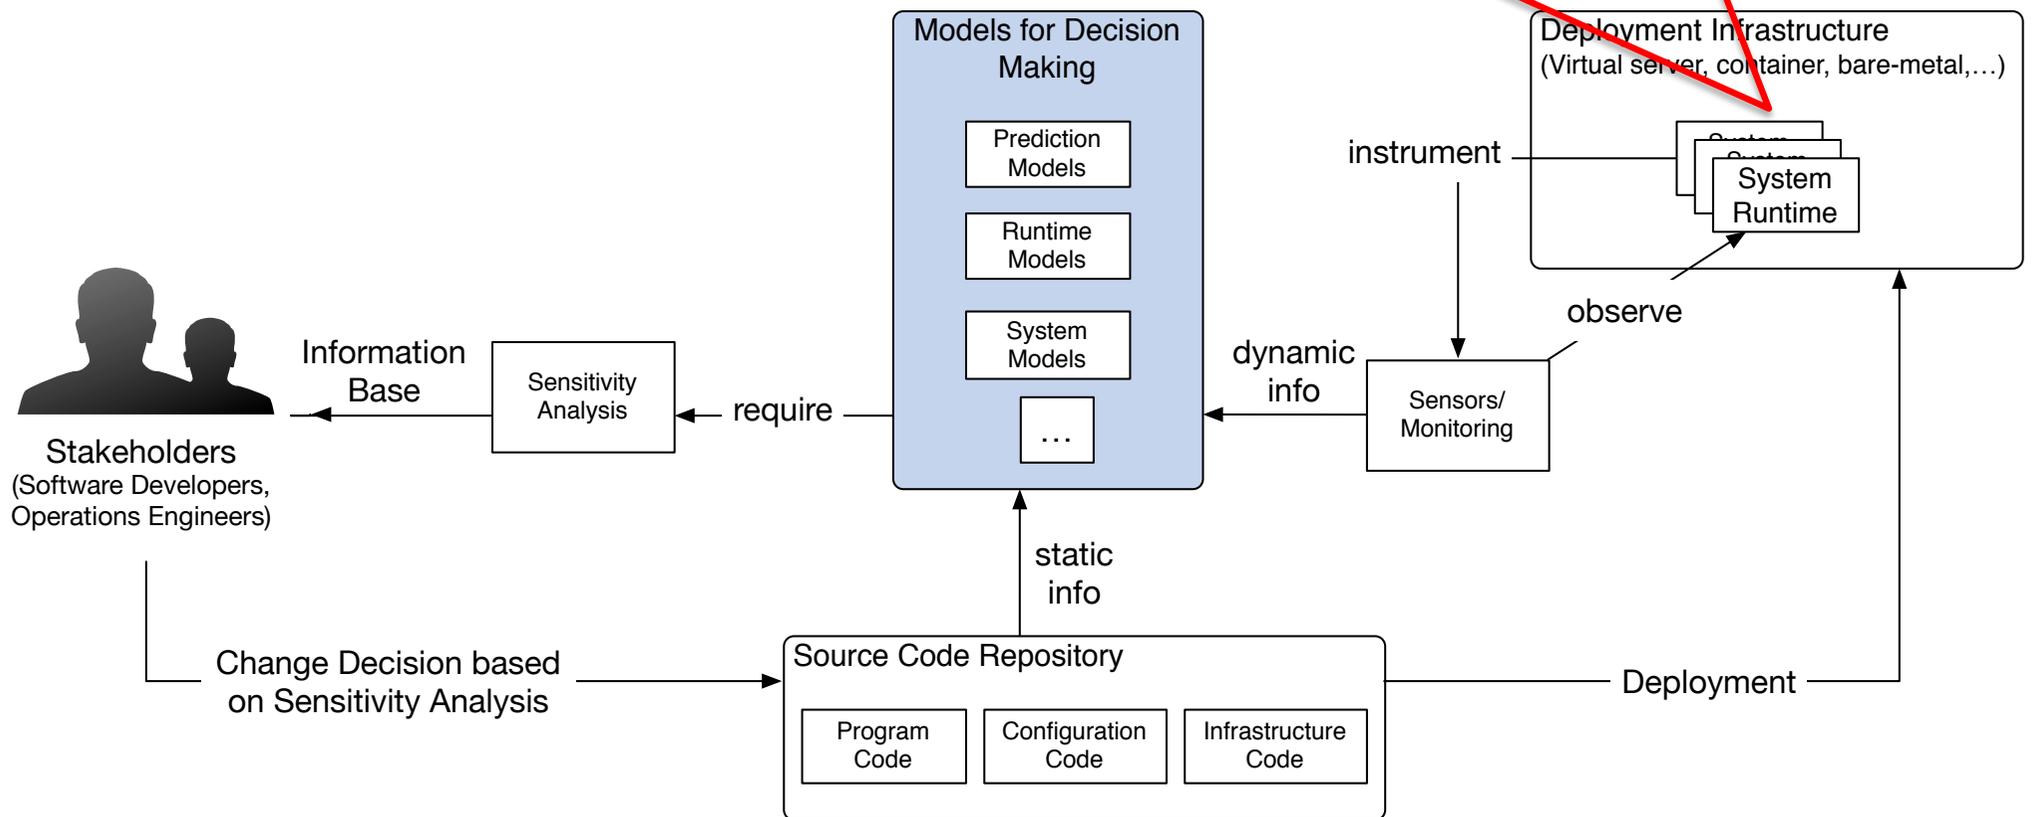
# Software versions and code changes (SCs): Code versioning, upgrade, patch



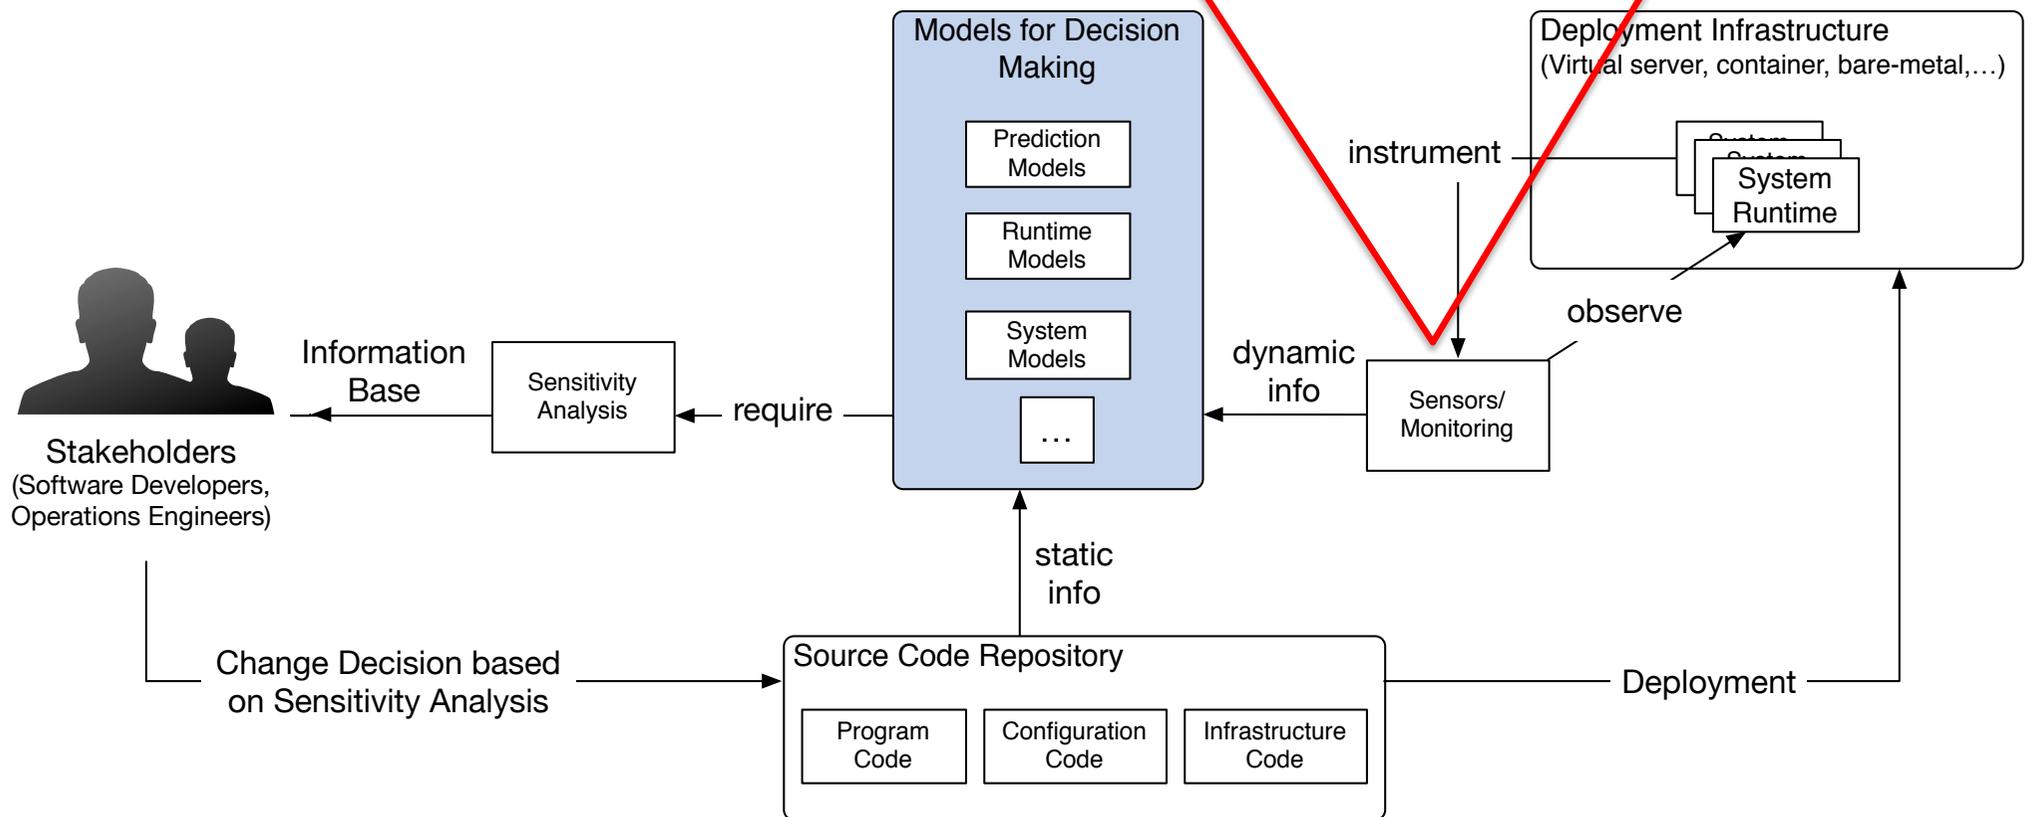
# Configuration parameters (CPs):



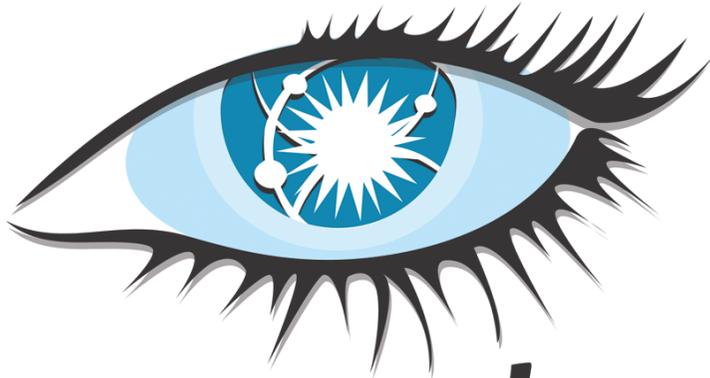
# Workload fluctuation(WF): User behavior, benchmark



# Monitoring and sensor accuracy (MS): active monitoring, instrumentation, and sensors



# We conduct a case study



***cassandra***



Yahoo Cloud Service Benchmark (YCSB)

**We measure system performance  
when altering the system based on  
different source of uncertainty**

**Deployment infrastructure: 2**

**Software versions: 3**

**Configuration parameters: 6 (1024  
possible configurations)**

**Workload: 6**

**We first alter configurations and  
keep others unchanged**

**Deployment infrastructure: 2** 

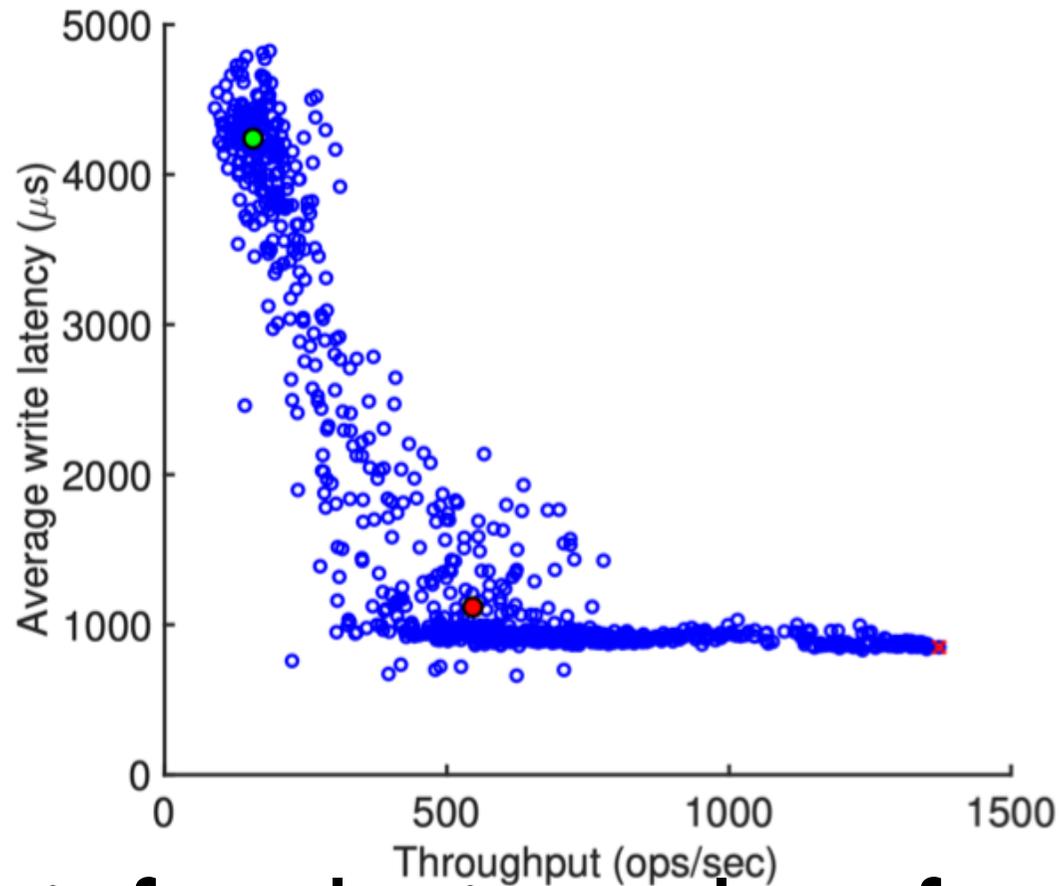
**Software versions: 3** 

**Configuration parameters: 6 (1024  
possible configurations)**



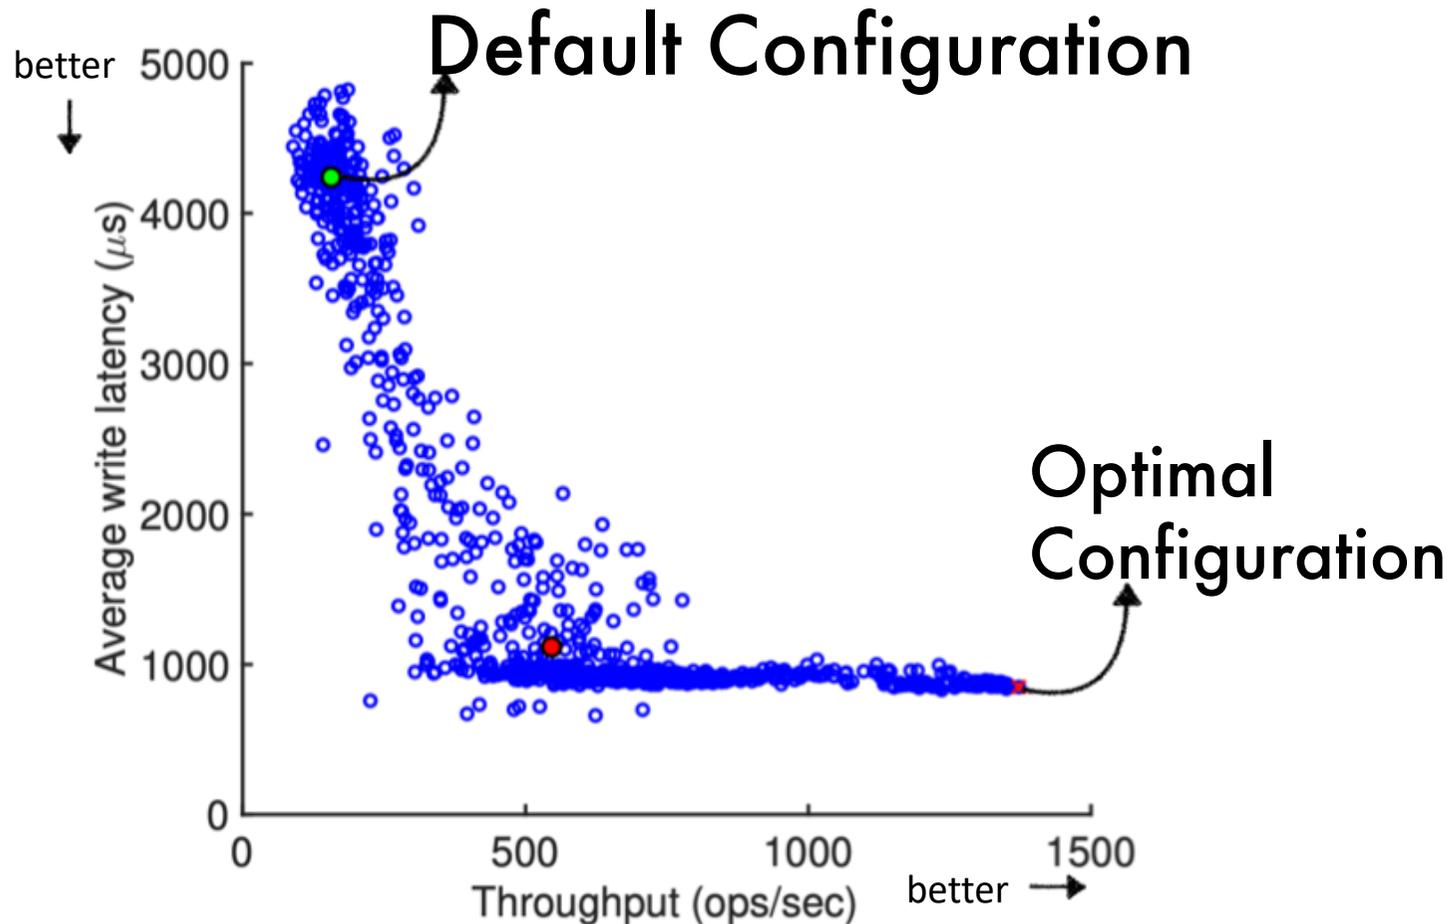
**Workload: 6** 

# There exist large uncertainty of performance when varying configurations



The plot is for altering value of configuration parameters when fixing all other aspects

The default configuration is typically bad and the optimal configuration is noticeably better than median



**We start to alter other aspects  
unchanged**

**Deployment infrastructure: 2** 

**Software versions: 3** 

**Configuration parameters: 6 (1024  
possible configurations)** 

**Workload: 6** 

# We measure the top/bottom configurations that are common between two settings

Decision	ID	Source	Target	Top	Bottom	Top/bottom	Correlation	Correlation (10%)
DI	ec <sub>1</sub>	h2-A-V3	h1-A-V3	0.0980	0.1569	0.0589	0.0364	-0.0078
SC	ec <sub>2</sub>	h1-A-V3	h1-A-V2	0.0490	0.0588	0.0098	-0.1266	-0.0527
	ec <sub>3</sub>	h1-A-V3	h1-A-V1	0.1176	0.0376	0.08	0.1424	0.0696
WF	ec <sub>4</sub>	h2-A-V3	h2-B-V3	0.0392	0.0686	0.0294	-0.1732	0.0139
	ec <sub>5</sub>	h2-A-V3	h2-C-V3	0.1373	0.1275	0.0098	0.0318	0.0381
	ec <sub>6</sub>	h2-A-V3	h2-D-V3	0.1471	0.1176	0.0295	0.0088	0.0172
	ec <sub>7</sub>	h2-A-V3	h2-E-V3	0.0490	0.0686	0.0196	-0.0704	0.0127
	ec <sub>8</sub>	h2-A-V3	h2-F-V3	0.0686	0.1373	0.0687	0.0217	0.0078
SC-WF	ec <sub>9</sub>	h1-A-V3	h1-B-V1	0.1078	0.1765	0.0687	0.1001	-0.0302
DI-SC-WF	ec <sub>10</sub>	h2-A-V3	h1-B-V1	0.1078	0.1176	0.0098	-0.0327	0.0192

Before altering

After altering

Potion of common top configuration before/after altering

What is altered

Decision	ID	Source	Target	Top	B
DI	ec <sub>1</sub>	h2-A-V3	h1-A-V3	0.0980	0
SC	ec <sub>2</sub>	h1-A-V3	h1-A-V2	0.0490	0
	ec <sub>3</sub>	h1-A-V3	h1-A-V1	0.1176	0
WF	ec <sub>4</sub>	h2-A-V3	h2-B-V3	0.0392	0
	ec <sub>5</sub>	h2-A-V3	h2-C-V3	0.1373	0

**Correlation of each configuration's performance before/after altering**

Top	Bottom	Top/bottom	Correlation	Correlation (10%)
0.0980	0.1569	0.0589	0.0364	-0.0078
0.0490	0.0588	0.0098	-0.1266	-0.0527
0.1176	0.0376	0.08	0.1424	0.0696
0.0392	0.0686	0.0294	-0.1732	0.0139
0.1373	0.1275	0.0098	0.0318	0.0381
0.1471	0.1176	0.0295	0.0088	0.0172

# The percentage of top/bottom common configurations between two settings are low

	Top	Bottom	Top			(10%)
	0.0980	0.1569	0.0			
	0.0490	0.0588	0.0			
	0.1176	0.0376	0.0			
	0.0392	0.0686	0.0			
	0.1373	0.1275	0.0			
	0.1471	0.1176	0.0			
	0.0490	0.0686	0.0196		-0.0704	0.0127

The best/worst configuration of one setting typically do not apply for another setting.

# The correlation of configurations performance between two settings decreases with noise

	Top	Bottom	Top/bottom	Correlation	Correlation (10%)
<b>The same configuration typically have different performance for different settings.</b>				0.0364	-0.0078
				-0.1266	-0.0527
				0.1424	0.0696
				-0.1732	0.0139
				0.0318	0.0381
				0.0088	0.0172
				-0.0704	0.0127

Correlation of with injected white noise as Monitoring and sensor accuracy (MS) uncertainty

Top			Correlation	Correlation (10%)
0.09			0.0364	-0.0078
0.04			-0.1266	-0.0527
0.1176	0.0576	0.06	0.1424	0.0696
0.0392	0.0686	0.0294	-0.1732	0.0139
0.1373	0.1275	0.0098	0.0318	0.0381
0.1471	0.1176	0.0295	0.0088	0.0172

Motoring noise worsen the uncertainty

# What should practitioners do?



Conduct **additional experiments** that further reduce the uncertainty



Identify and handle the **root cause** of the uncertainty



If the uncertainty cannot be easily reduced or handled, **uncertainty quantification approaches** should be considered.

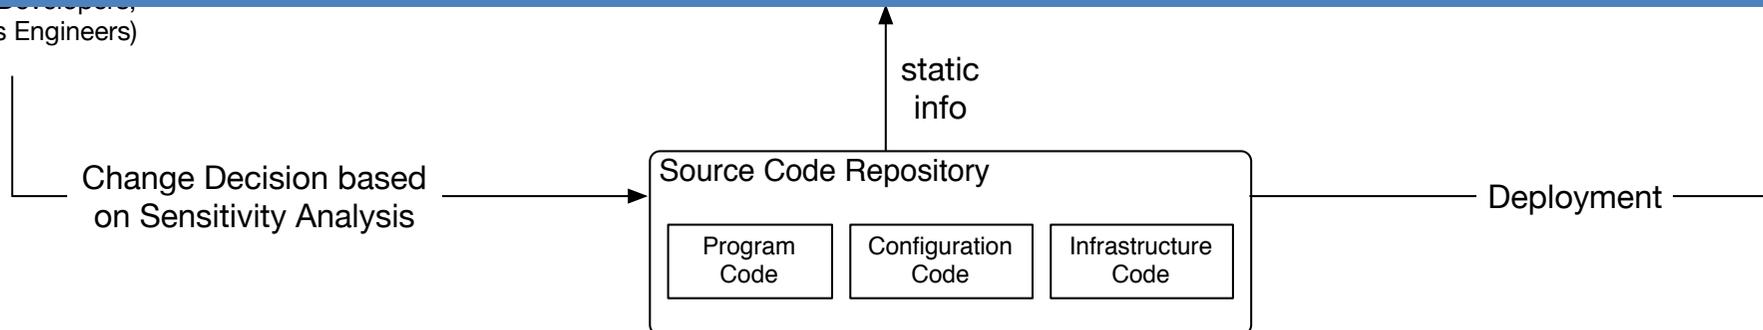
**Deployment infrastructure (DI):  
Physical or virtual, type of nodes...**

Models for Decision Making

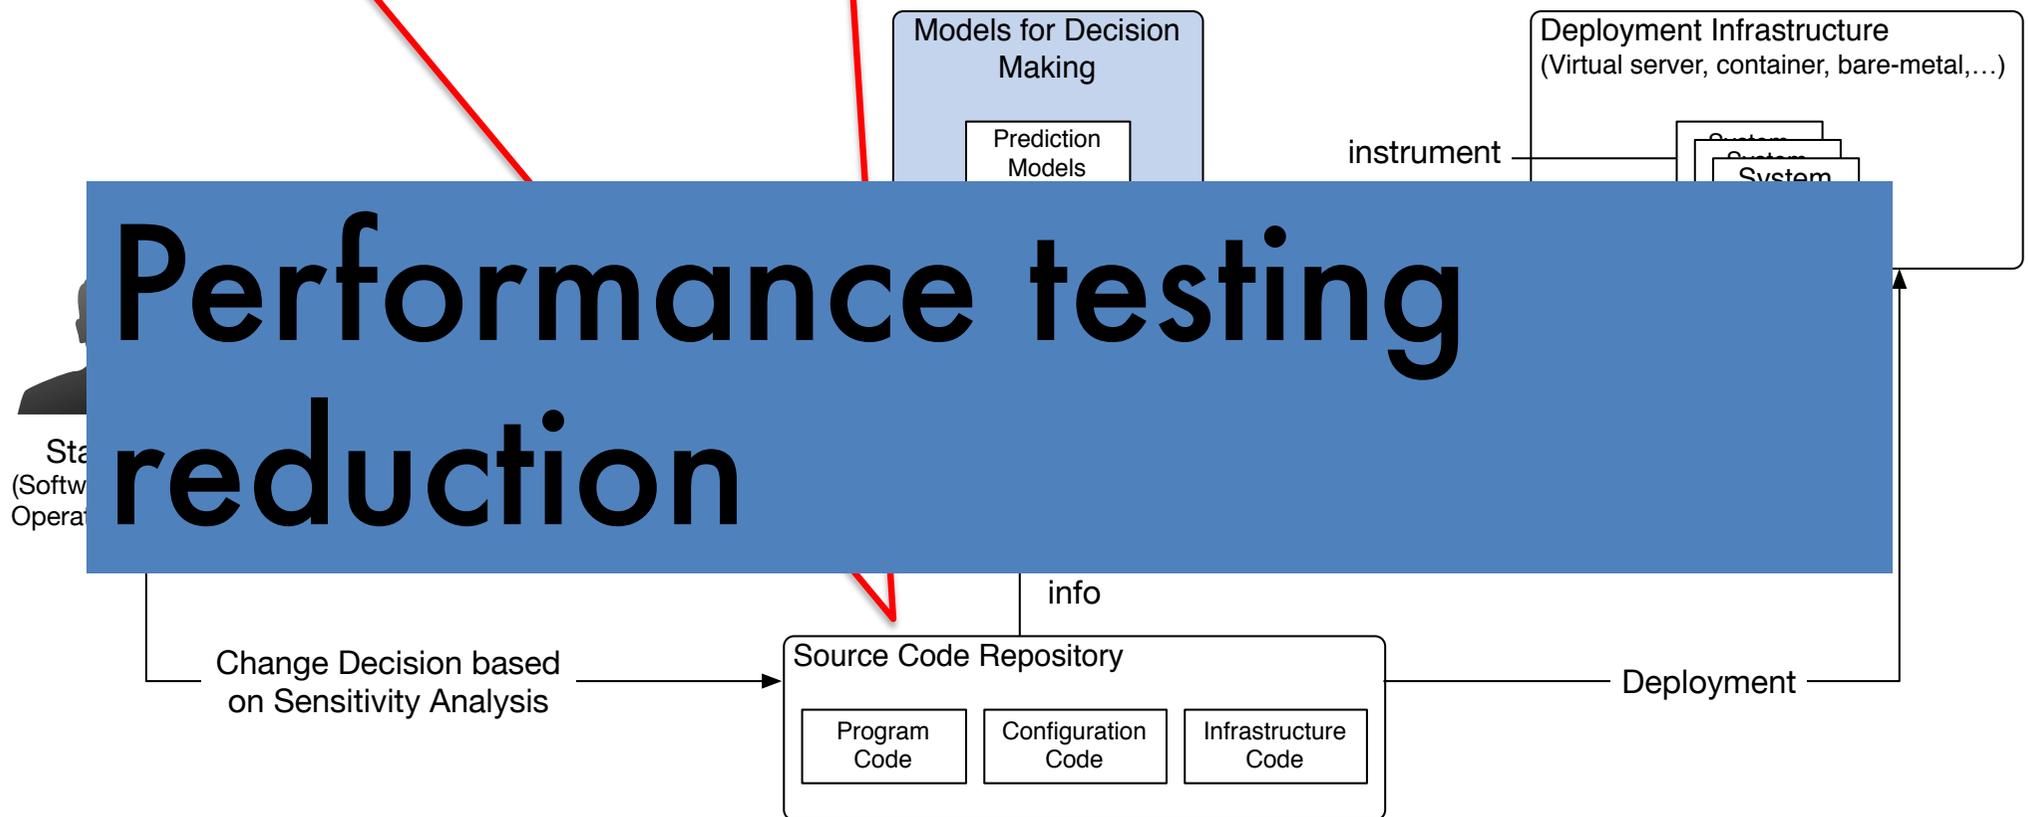
Deployment Infrastructure  
(Virtual server, container, bare-metal,...)

# User acceptance testing and canary deployment

St  
(Software Developers,  
Operations Engineers)



# Software versions and code changes (SCs): Code versioning, upgrade, patch



# Configuration parameters (CPs):

Models for Decision Making

Deployment Infrastructure  
(Virtual server, container, bare-metal,...)

# Relevant configuration isolation

Stakeholder  
(Software Developer, System Operator)

Change Decision based on Sensitivity Analysis

Source Code Repository

Program Code

Configuration Code

Infrastructure Code

static info

Deployment

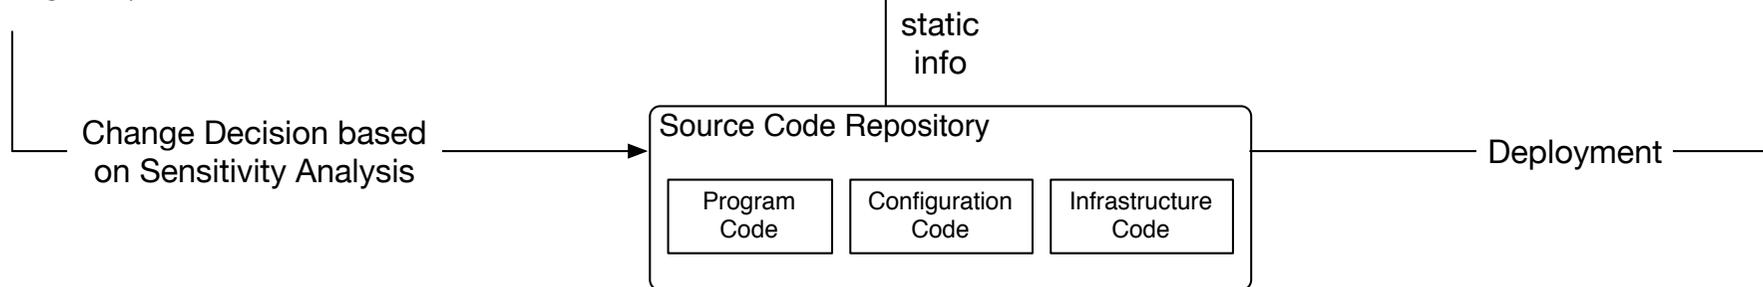
# Workload fluctuation(WF): User behavior, benchmark

## Workload recovery and verification

Models for Decision Making

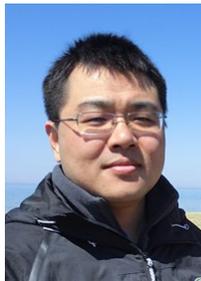
Deployment Infrastructure  
(Virtual server, container, bare-metal,...)

(S  
Op

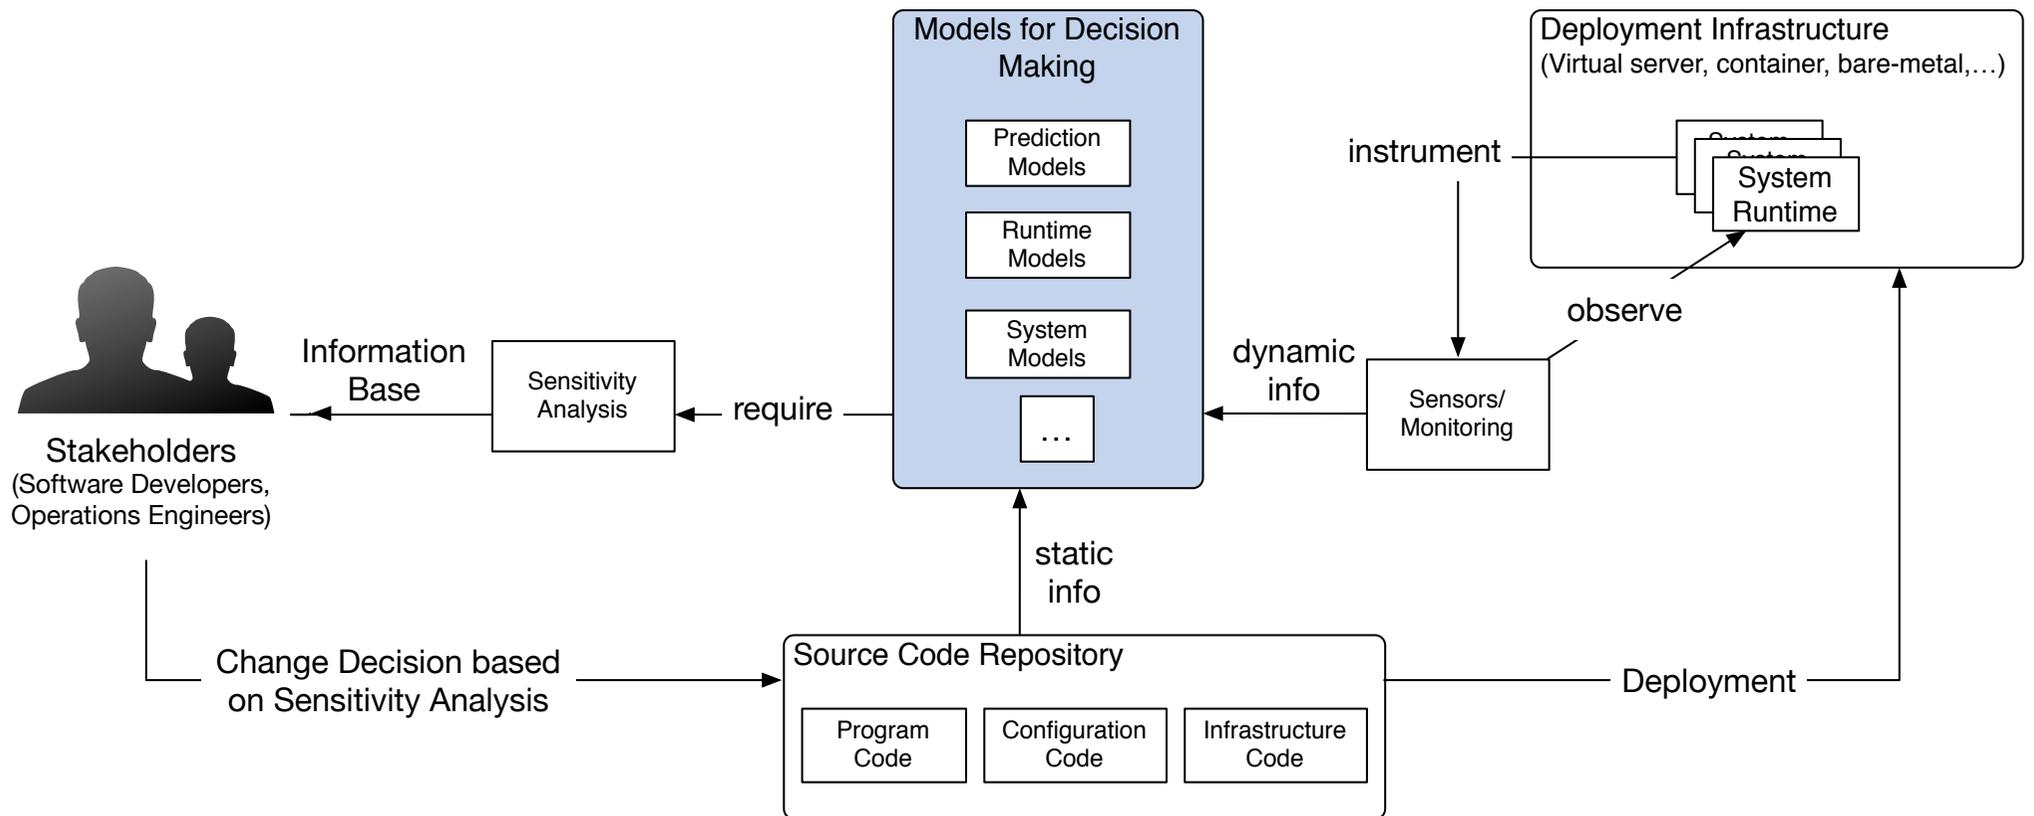




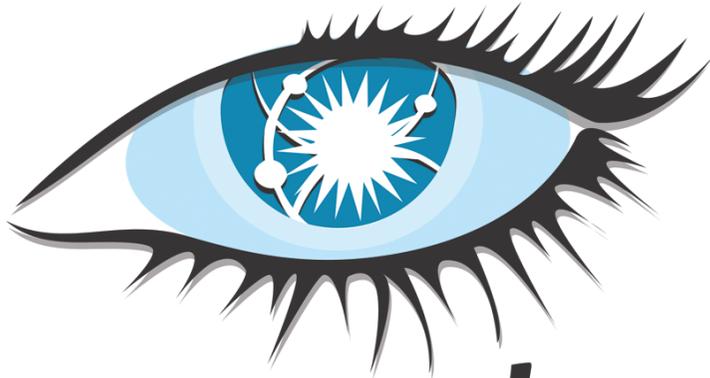
# Sources of Uncertainty in Performance-aware DevOps



# Our final abstraction of the uncertainties



# We conduct a case study



***cassandra***



Yahoo Cloud Service Benchmark (YCSB)

# What should practitioners do?



Conduct **additional experiments** that further reduce the uncertainty



Identify and handle the **root cause** of the uncertainty

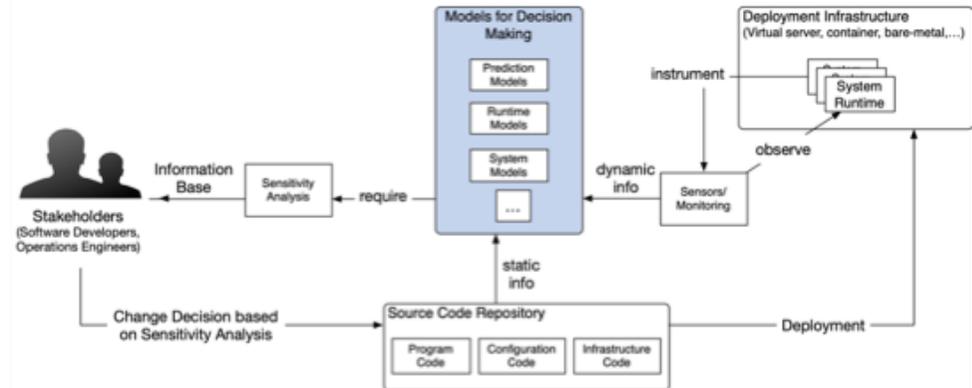


If the uncertainty cannot be easily reduced or handled, **uncertainty quantification approaches** should be considered.

# Sources of Uncertainty in Performance-aware DevOps



## Our final abstraction of the uncertainties



## We conduct a case study



Yahoo Cloud Service Benchmark (YCSB)

## What should practitioners do?



Conduct **additional experiments** that further reduce the uncertainty



Identify and handle the **root cause** of the uncertainty



If the uncertainty cannot be easily reduced or handled, **uncertainty quantification approaches** should be considered.