

# Functional Workload Model


Francois Raab, InfoSizing,  
Yanpei Chen, Cloudera

# Successful Benchmarks

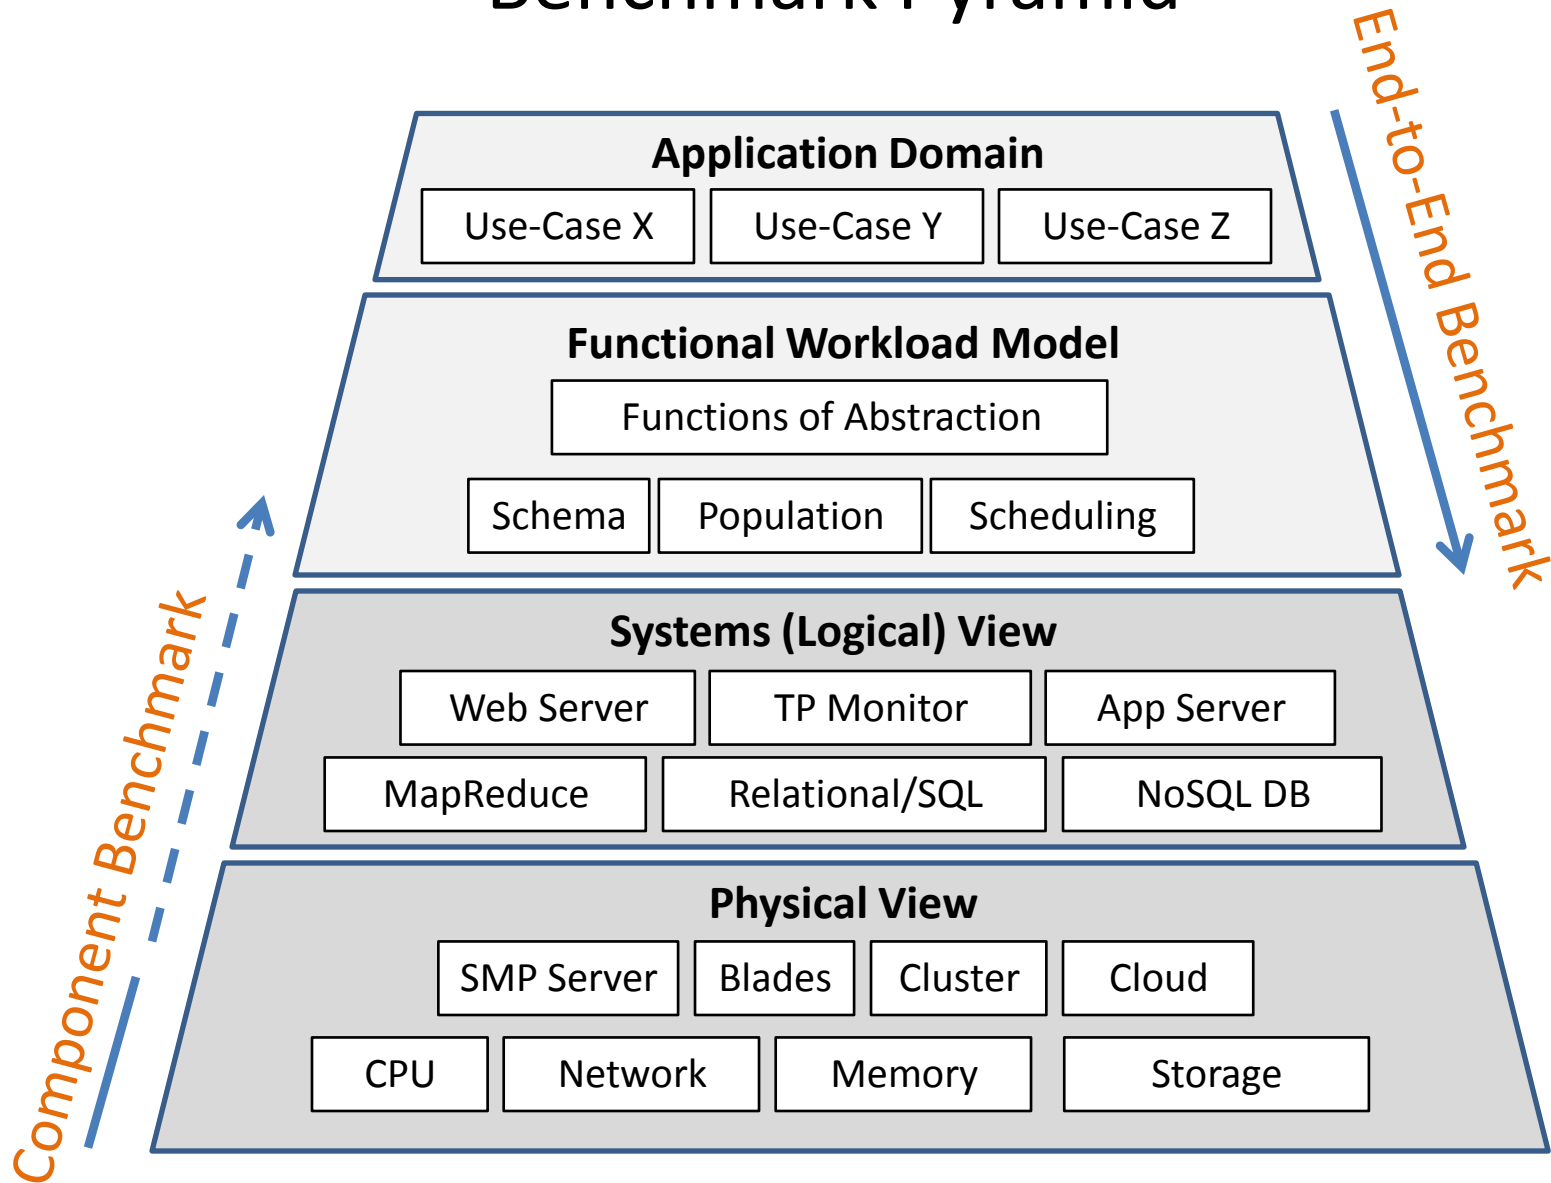
By number of published results (approx.)

| Benchmark                       | Results | Type       |
|---------------------------------|---------|------------|
| SPECcpu (2006 + 2000 + 95)      | 10,800  | Component  |
| SPECjbb (2005 + 2000)           | 1,050   | End-to-End |
| TPC-C                           | 760     | End-to-End |
| SPEC SFS                        | 730     | End-to-End |
| SPECweb (2009 + 2005 + 99 + 96) | 700     | End-to-End |
| TPC-D/H                         | 650     | End-to-End |

# End-to-End Definition Flow

- 
- Identify use cases (real-life applications)
  - Group use cases by ***Application Domain***
  - Select an “interesting” application domain
  - Build ***Functional Workload Model***
    - Refine into ***Functions of Abstraction***
    - Define Schema, Population, Scheduling
  - Complete benchmark definition
    - Execution rules (do’s and don’ts)
    - Constraints (must have’s)
    - Metrics (absolute, composite)

# Benchmark Pyramid



# Proposal: Use Cases Library

- Inventory of specific use cases (deployment)
  - Summary description of context
  - Workload characteristics
    - Data set (4 V's)
    - Basic operations
    - Execution environment
- Classification in Application Domains
  - Commonalities
  - Popularity
  - Adoption rate