Exploiting Big Data in Engineering Adaptive Cloud Services

Patrick Martin
Database Systems Laboratory
School of Computing, Queen's University

Big Data

- 3 v's: ***Volume, Velocity, Variety***

- ***Data analytics*** to support effective on-line decision making

Adaptive Cloud Services

- Software services offered from *private/public/multi/hybrid* clouds
- Exploit ***elasticity*** of clouds
- Exploit ***mobility*** across cloud providers

Claim

Big data has an important role to play in the engineering of adaptive software systems, in general, and adaptive cloud services, in particular.

QuARAM Framework

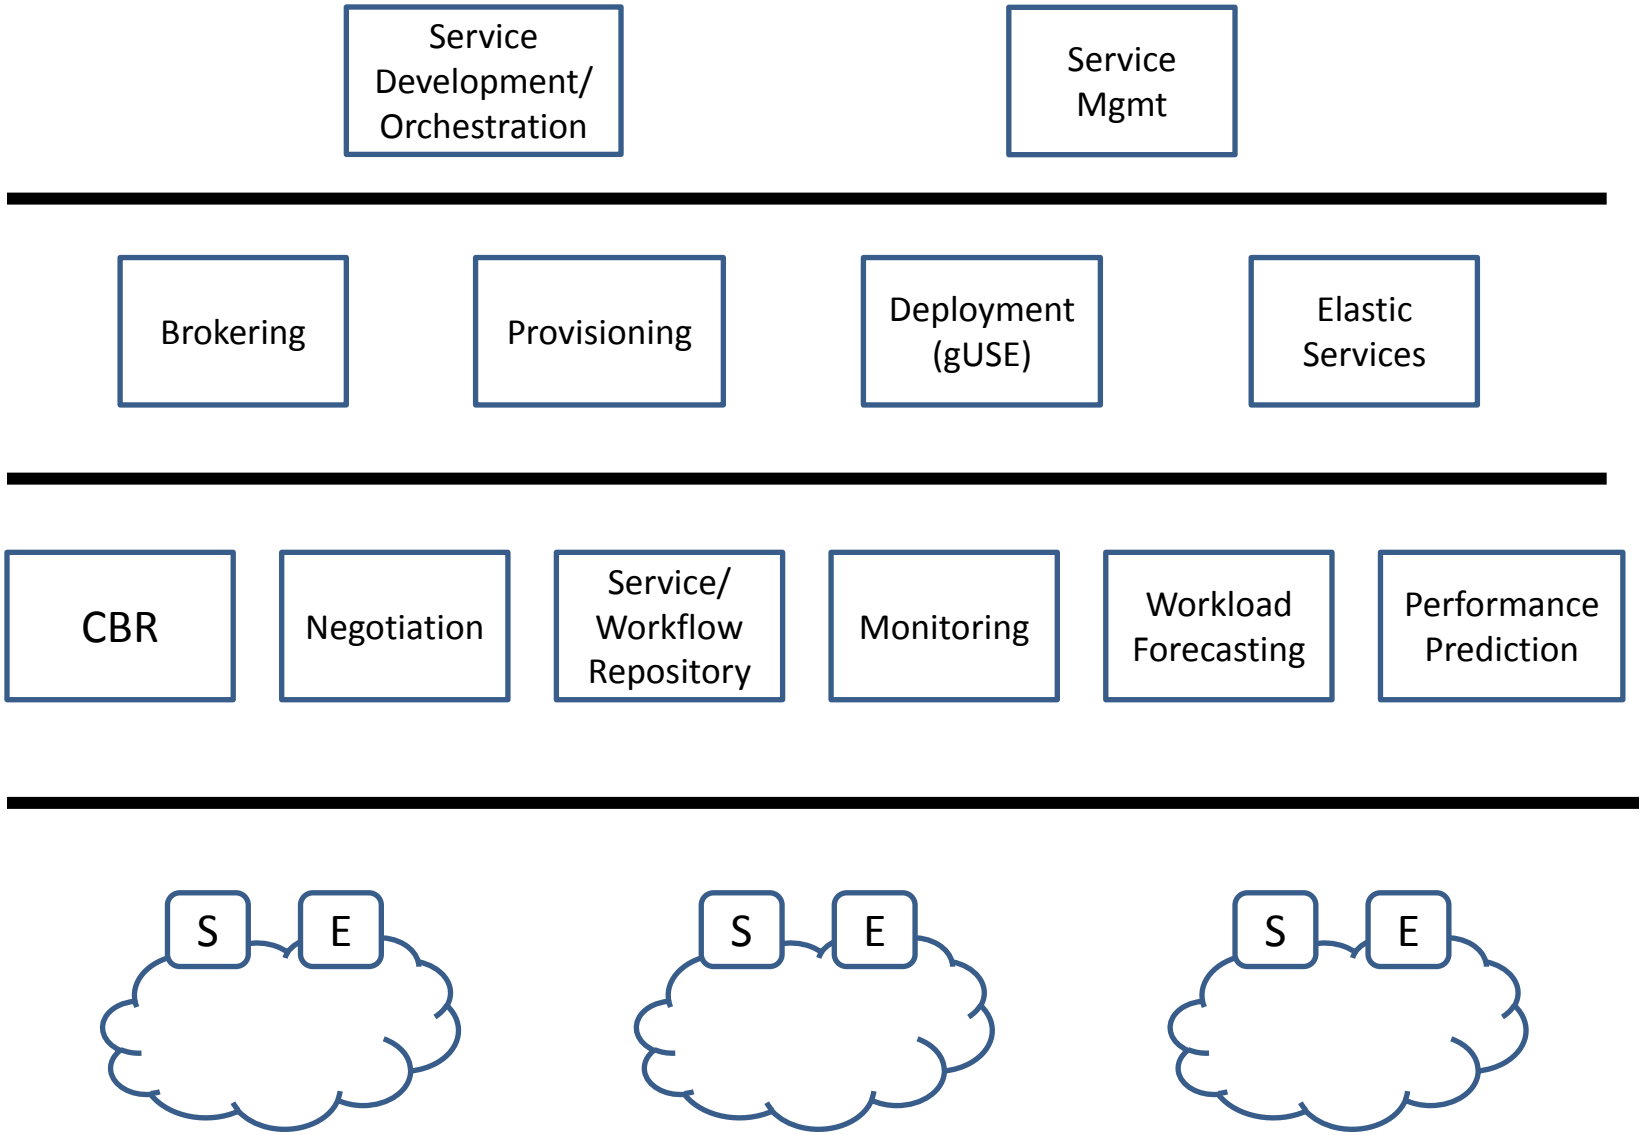

Why QoS-Aware Management?

- Consumer – provider relationship in clouds will rely on SLAs
- Providers will need to support differentiated QoS
- QoS involves application-level metrics that are understood by both parties
- QoS determines resource requirements

Challenges in QoS-Aware Management

- *Complexity* of applications and of cloud environments
- *Varying demands* from an application
- *Heterogeneity* of cloud offerings
- *Varying service quality* from a cloud provider
- *Conflicting goals* of application and cloud provider

Role of Big Data

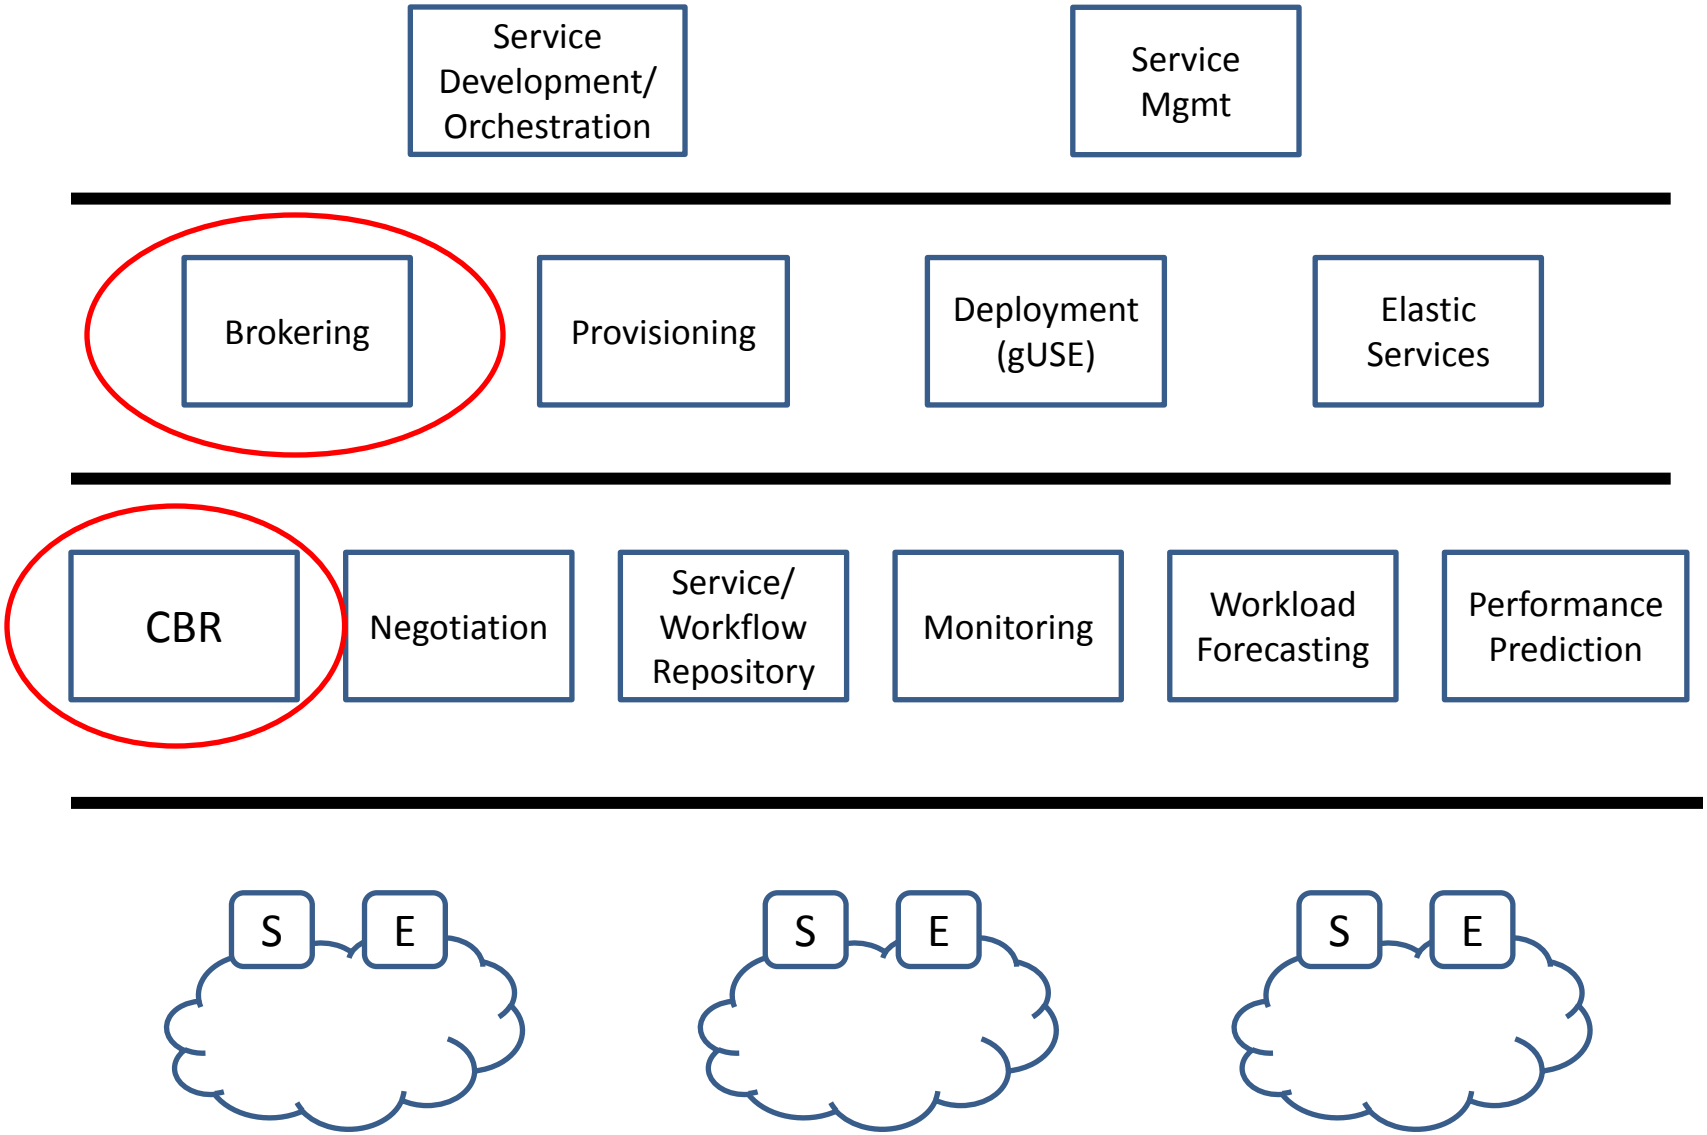

Role of Big Data

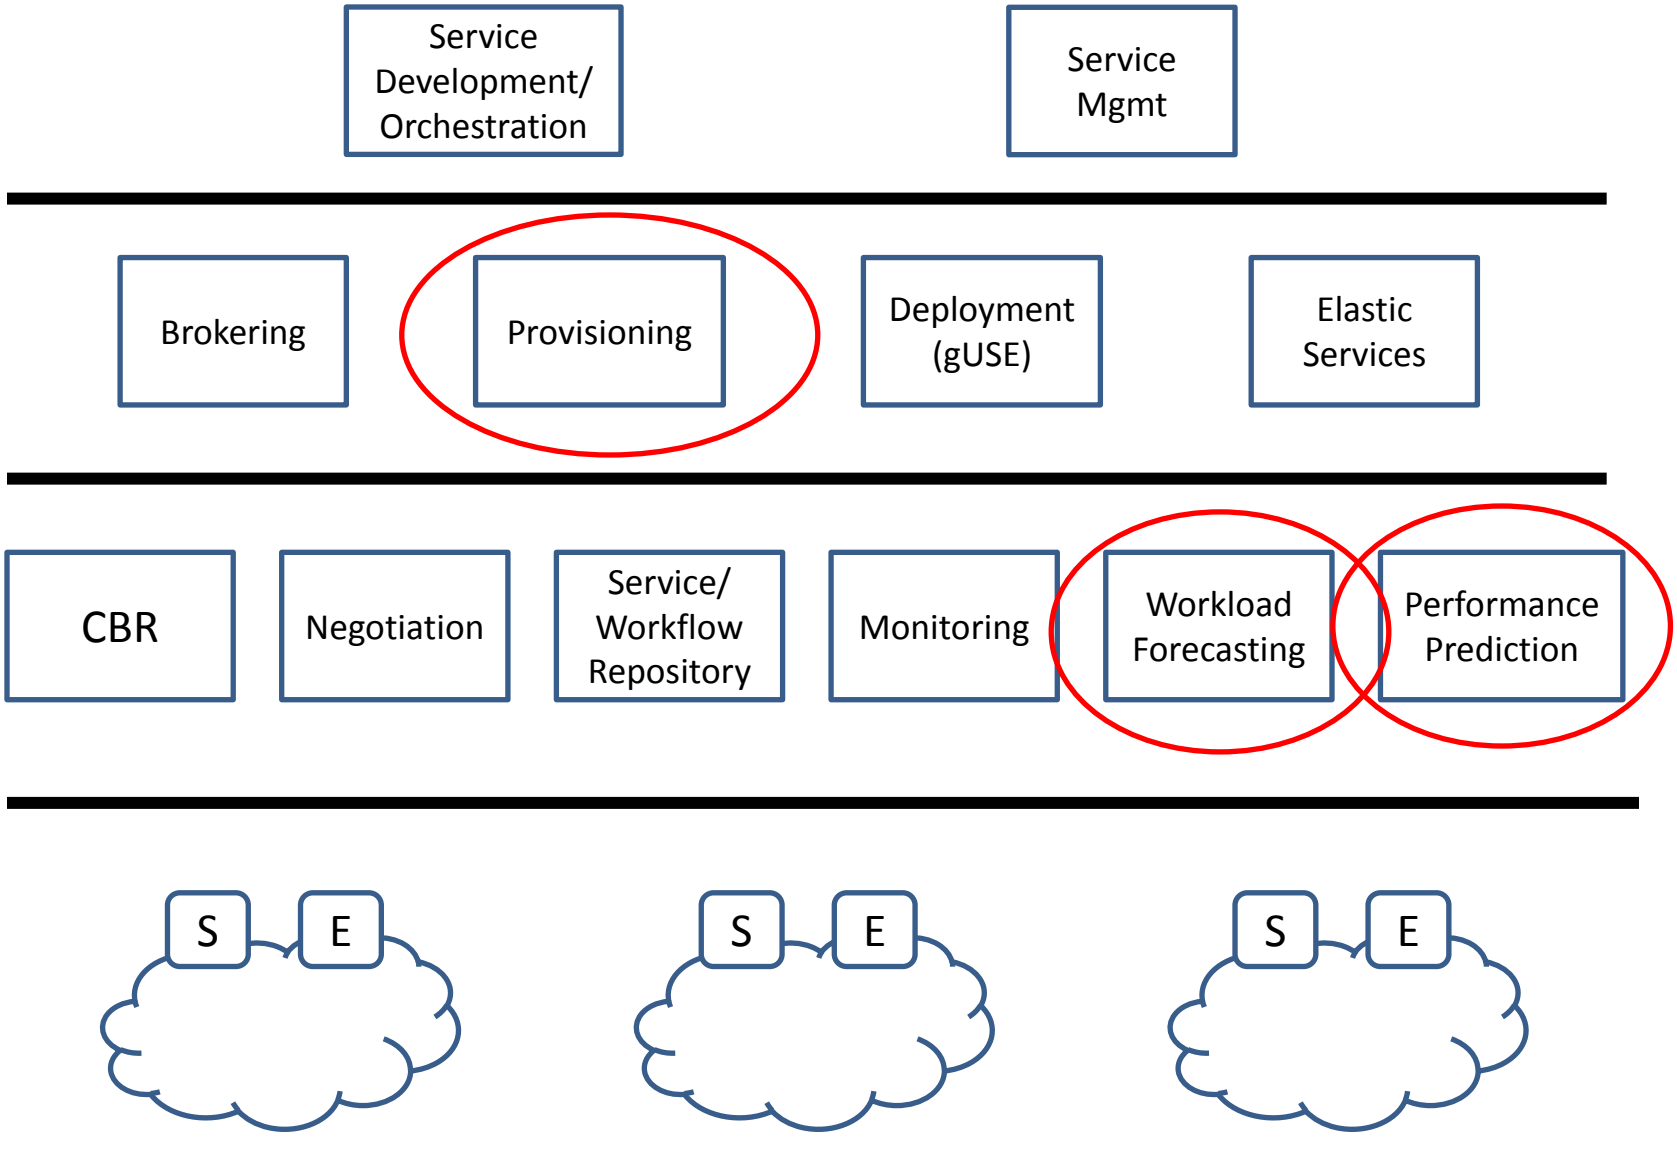

Role of Big Data

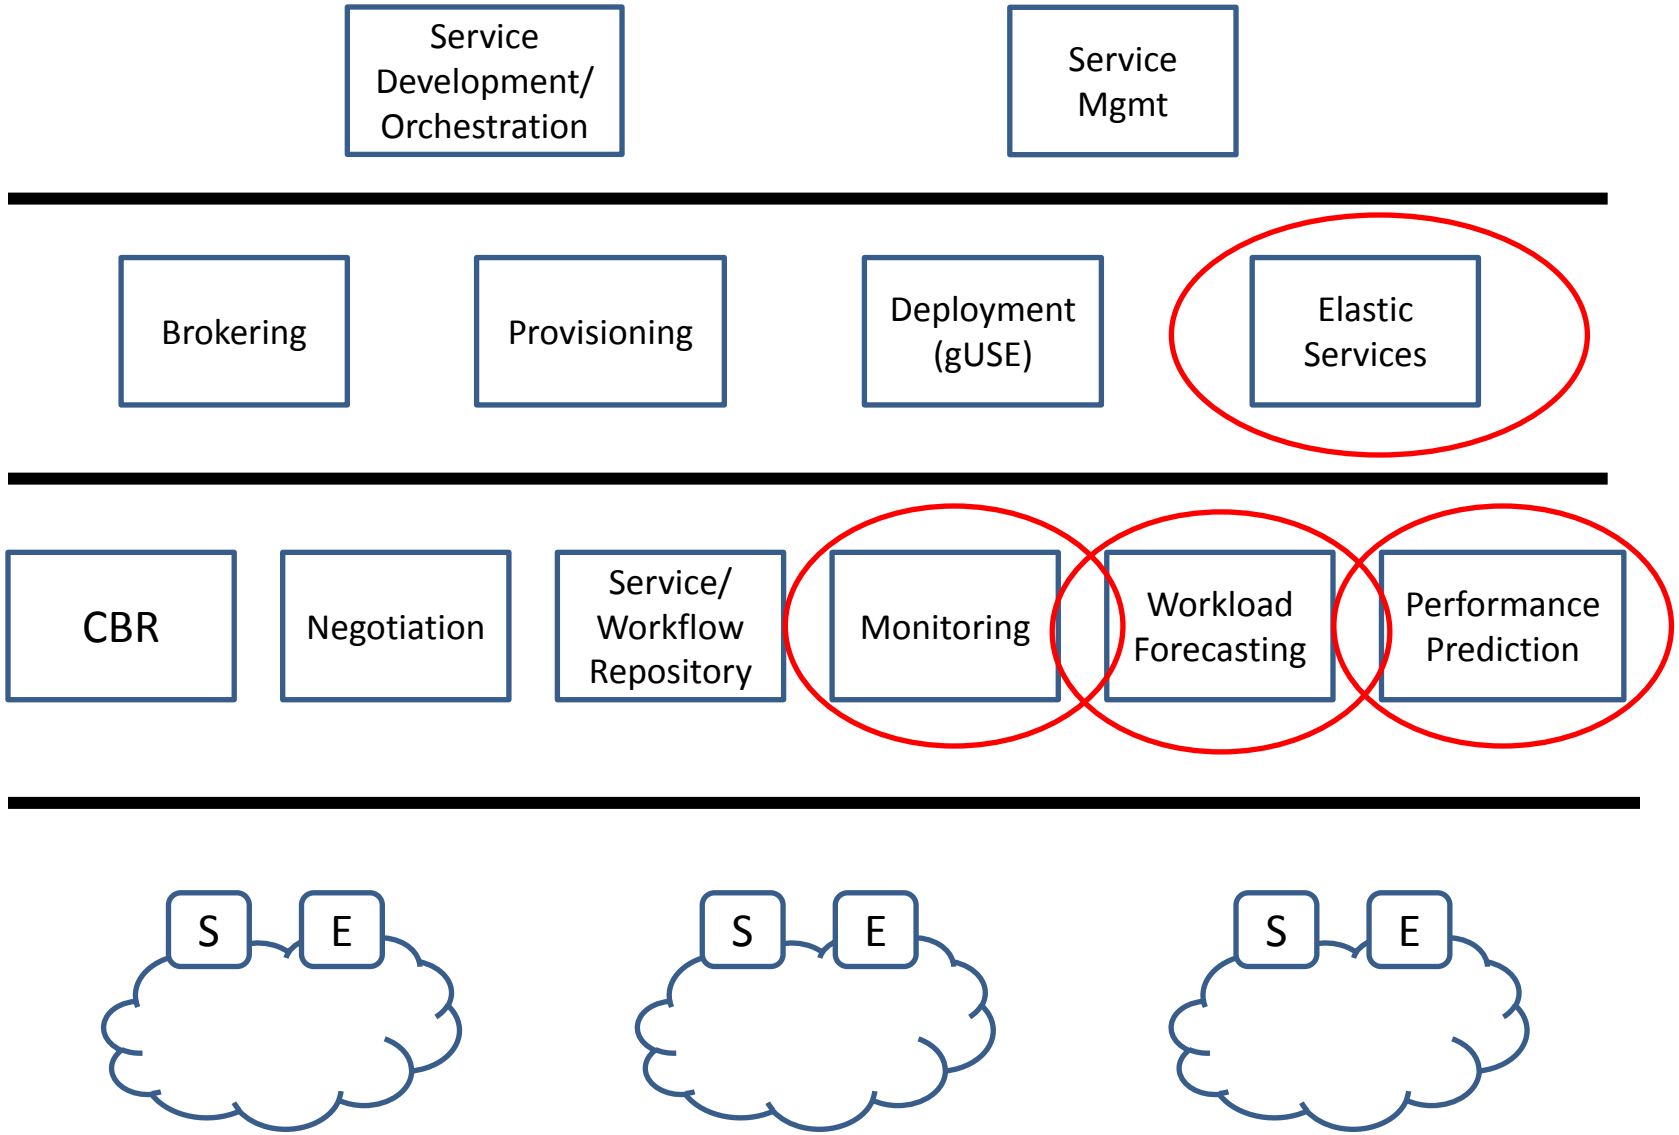

Brokering - CBR

- Big data used to construct case base for cloud provider selection
  - Case features based on resource demands
  - Demands extracted by analyzing performance logs of services in current configurations
- CBR used to find potential matching deployments that can be adapted to new service's requirements

Monitoring

- Big data collected and analyzed during service execution
  - Multiple streams of data
  - ETL, aggregation over time and/or sources
  - Triggers decision-making
  - Used to adapt workload and performance models

Workload Forecasting

- Big data used to
  - Classify requests
  - Characterize resource demands of classes
  - Identify trends in request intensities (time series analysis)
- Models need to adapt to changes in workload

Performance Prediction

- Big data used to build statistical models of performance
  - Clouds increase the number of parameters and the variability in the modeling process
  - Interactions among workload requests must be accounted for in the models
- Models need to adapt to changes in workload and configuration
  - Parameter selection for models
  - Ensemble models

Summary

- Big data plays important role in engineering adaptive cloud services
  - Brokering, provisioning, monitoring, workload forecasting and performance prediction

The big challenge 

Efficient online adaptation of the models!

Runtime Model Management

- Adaptive software systems use a LOT of models
- Adaptation middleware needs to provide model management
  - Store and describe models (metrics, policies, relationships)
  - Efficient maintenance of models at runtime to minimize impact on managed system

Thank you

NSERC
CRSNG