





# NASA EOSDIS Cloud Prototype Systems

Chris Lynnes, Mark McInterney, Katie Baynes  
NASA




# Annual distribution is on the same order of magnitude as the total archive volume





# And our archive is scheduled to grow dramatically





If these high volumes are a challenge for robust, experienced data centers to manage...

...what about the science end users?



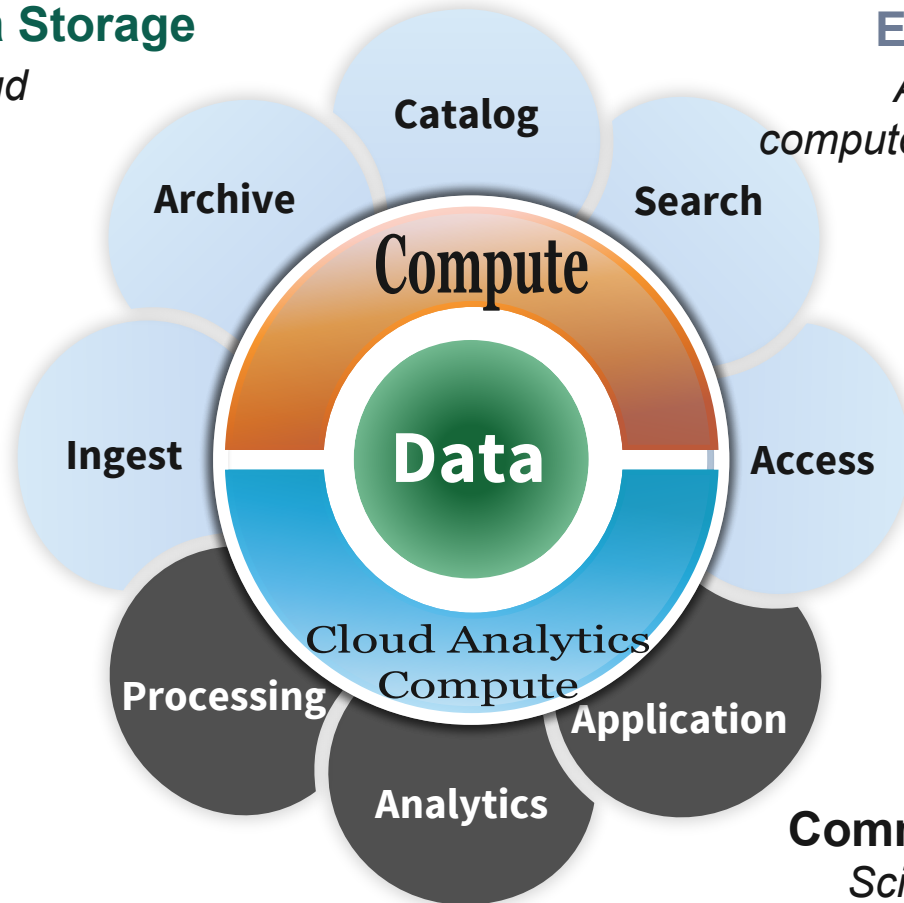
# Conceptual "Data Close to Compute"

## Large Volume Data Storage

*Datasets stored in cloud object storage*

## EOSDIS Applications

*Applications using cloud compute, storage, and services*




## Scalable Compute

*Provision, use, and terminate as needed*

## Community Applications

*Science community brings algorithms to the data.*



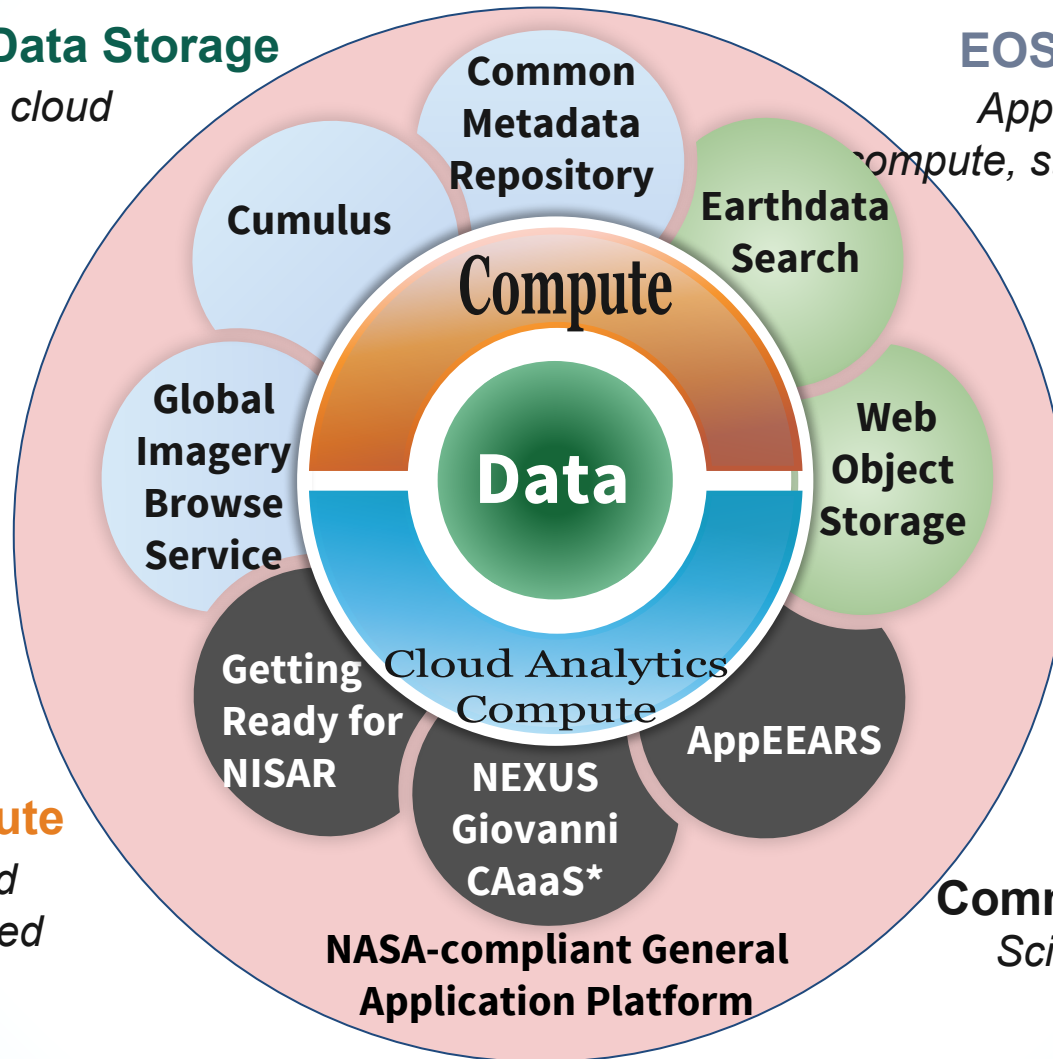
# Conceptual "Data Close to Compute"

## Large Volume Data Storage

*Datasets stored in cloud object storage*

## EOSDIS Applications

*Applications using cloud compute, storage, and services*




## Scalable Compute

*Provision, use, and terminate as needed*

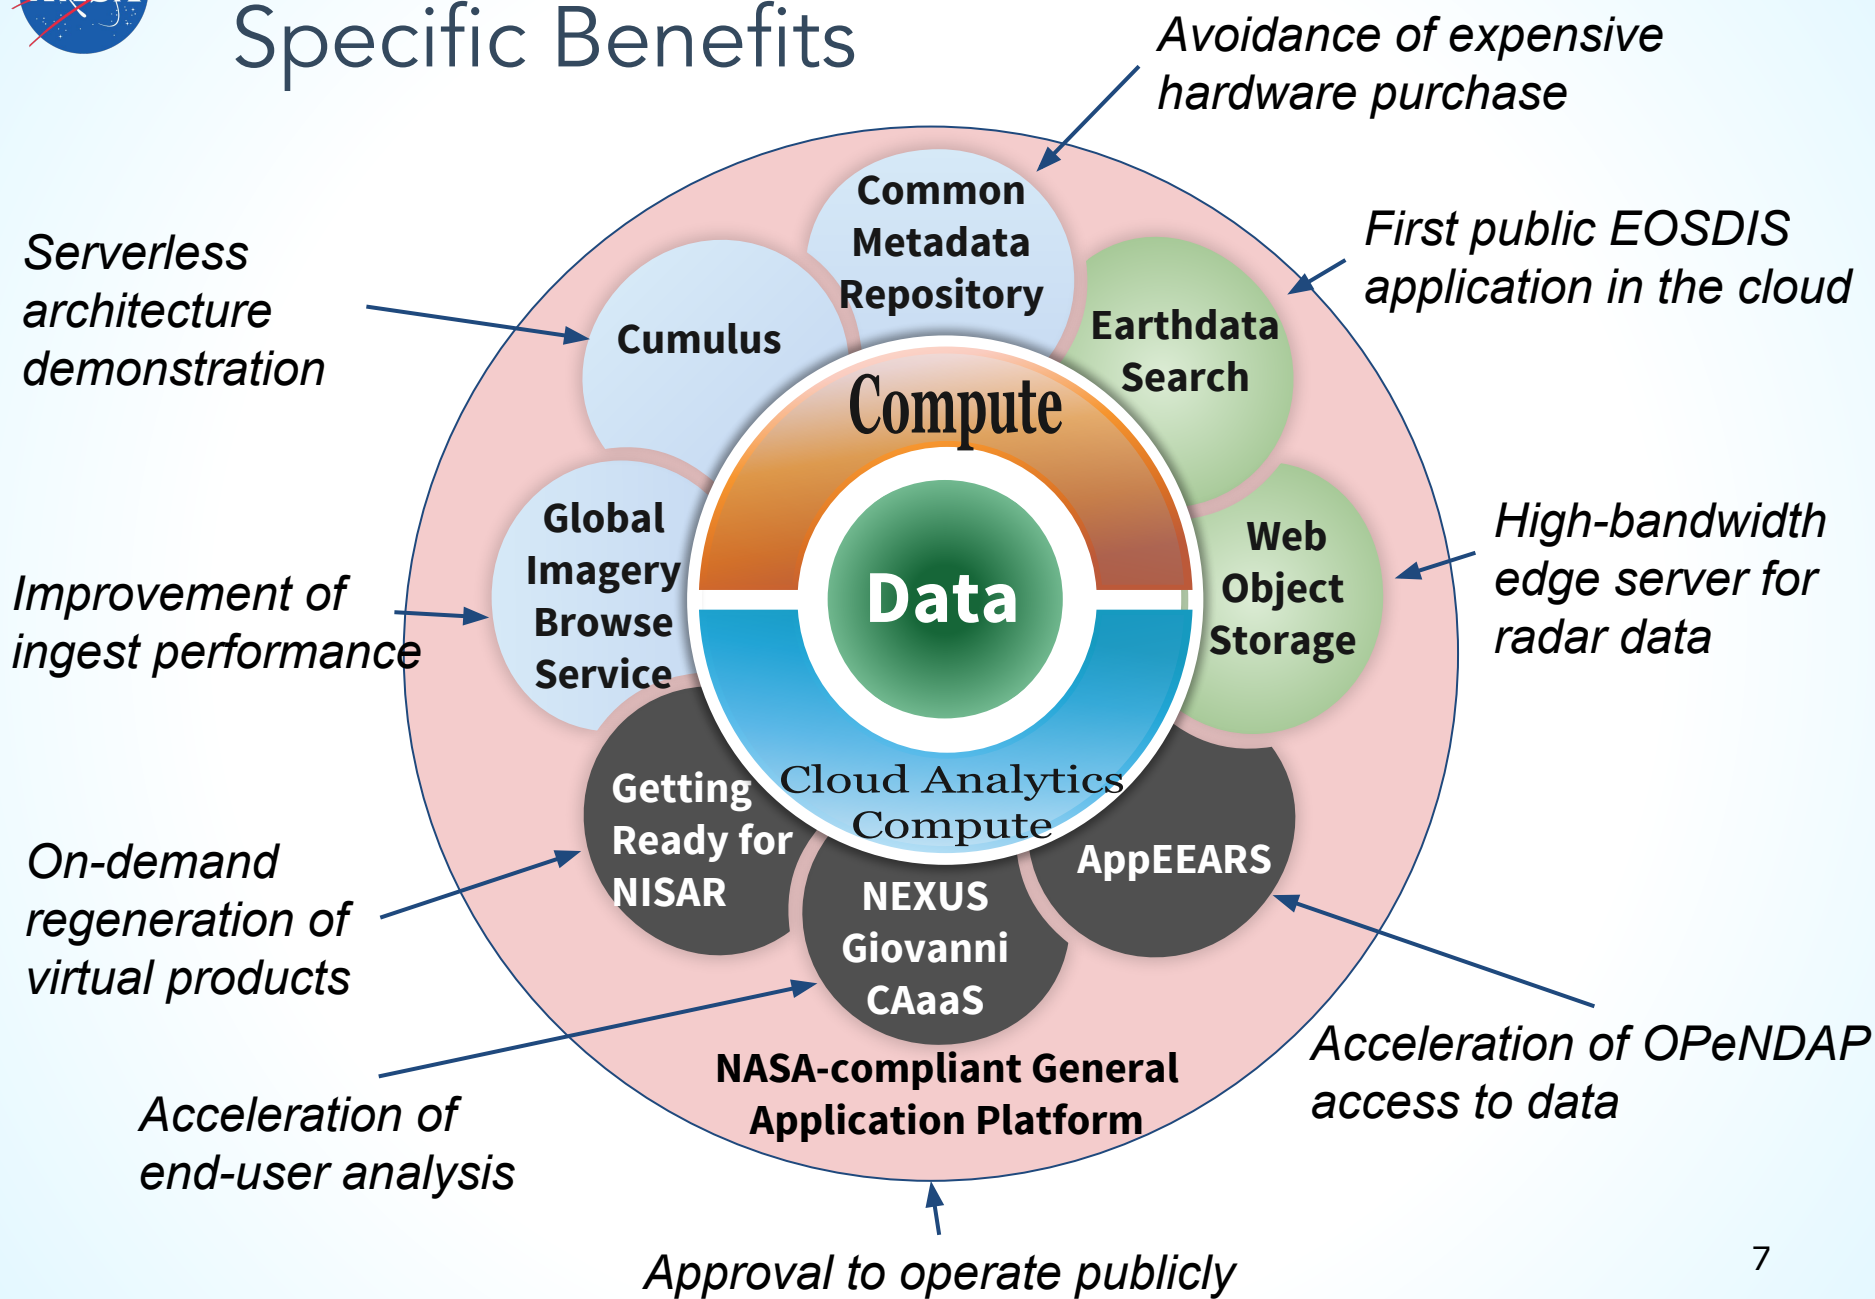
## Community Applications

*Science community brings algorithms to the data.*

\*CAaaS = Climate Analytics as a Service

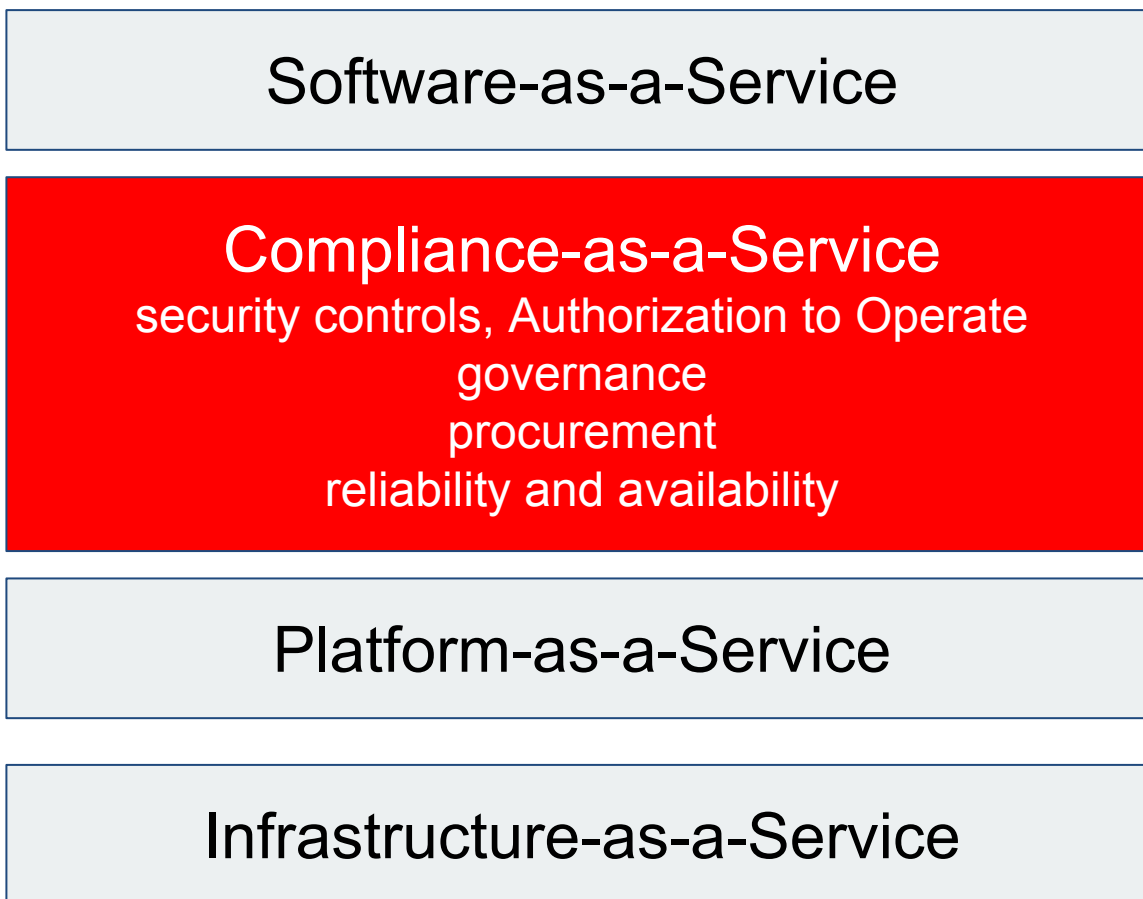


# Specific Benefits





# NASA-compliant General Application Platform: “Compliance-as-a-Service”







# Key Cloud Benefits for Science Users

Immediate direct access to data  
Immediate availability of compute

Lower Barrier to Entry

Dynamic scaling for bursting  
Variety of CPU and GPU

Scalable Performance

Purchase compute as needed  
Spot market for cost control  
No need for local copy of data

Cost


Experiment replication enabled  
Sharing of complex analysis workflows

Repeatability & Sharing

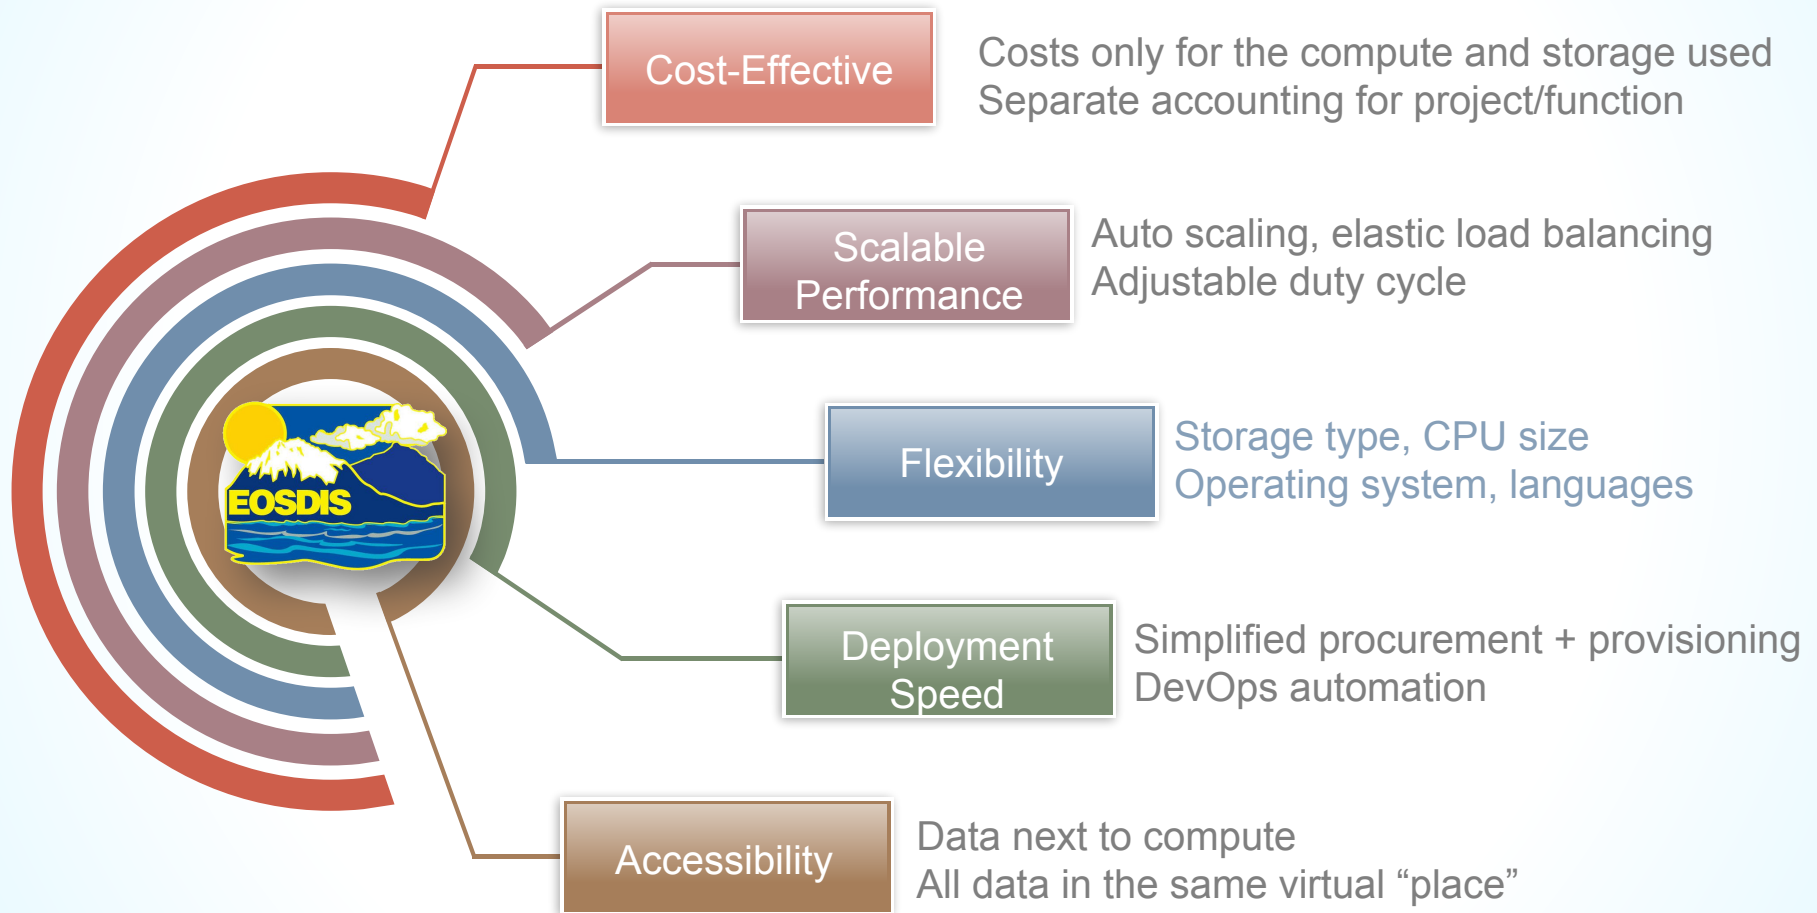
Decrease in startup time for projects  
Analysis at scale for anyone  
Rapid iteration enabled


Promoting Science Discovery





# Key Cloud Benefits for Data Systems (EOSDIS)





# Decision Considerations

## High level decision consideration for individual project prototypes and capabilities to operationalize into AWS (commercial cloud)



Is AWS (commercial cloud) affordable?




Is NASA IT Security compliance and tactical operations achievable in AWS (commercial cloud)?



Is performance equal to or better than current on-premises solutions?



Can we operate “Operationally” in AWS (commercial cloud), technical and business?



# Key Lessons - So Far

- Cloud-native re-architecture is preferable to forklift/lift-and-shift
- Savings from Dev / Test environments
  - Full-workload testing:  $N_{ops} = N_{perftest}$
  - On-Premise:  
$$N_{ops} + N_{perftest} = 2N \times 24hr \times 7d = 336 N\text{-hr}$$
  - Cloud:  
$$(N_{ops} \times 24hr \times 7d) + (N_{perftest} \times 10hr \times 5d) = 218 N\text{-hr} = 35\% \text{ savings}^*$$
- Serverless architectures look promising
  - Reduce needed code down to business logic only