



# Enabling Cloud Scale Distributed Capabilities


Otto White<sup>1</sup>, Yaoxin Jing<sup>1</sup>, Adrien Ghosn<sup>2</sup>, Anjo Vahldiek-Oberwagner<sup>3</sup>, Mona Vij<sup>3</sup>, Michael Steiner<sup>3</sup>, Lluís Vilanova<sup>1</sup>

HCDS '25

[<otto.white20@imperial.ac.uk>](mailto:otto.white20@imperial.ac.uk)

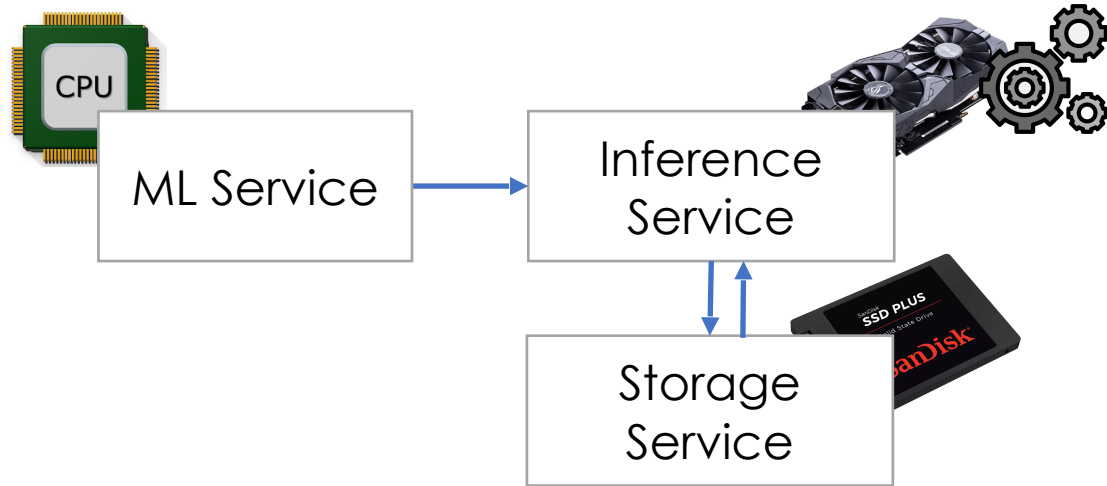
# Security in Disaggregated Cloud

# Security in Disaggregated Cloud



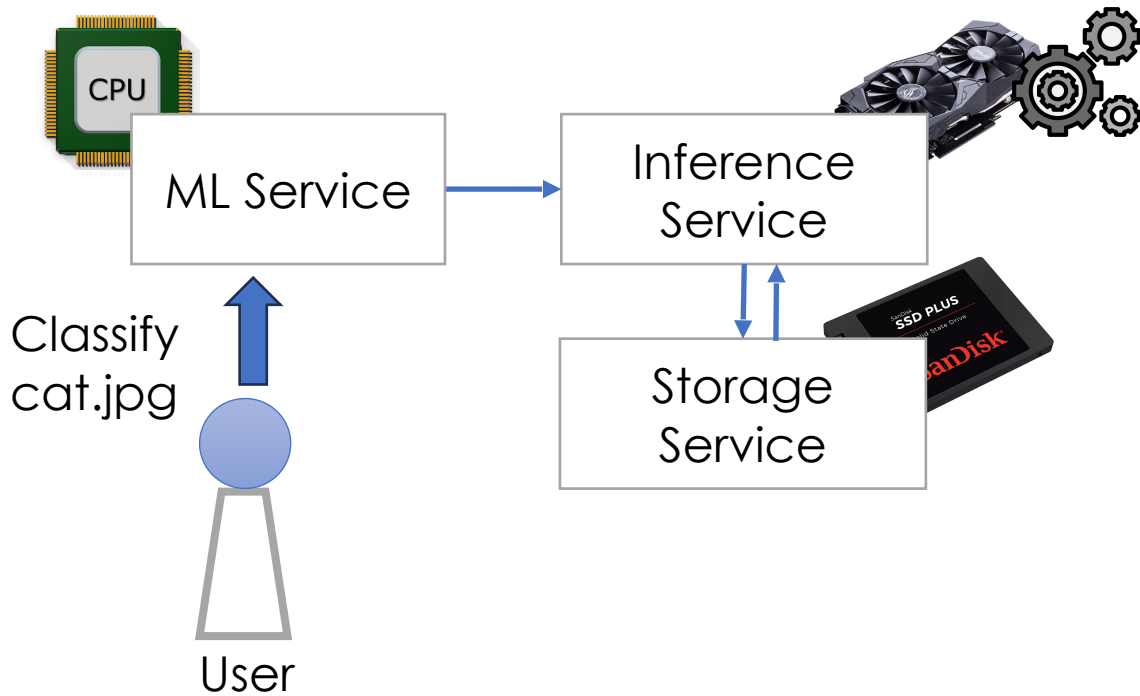
- Disaggregated Service-Oriented Architecture

# Security in Disaggregated Cloud



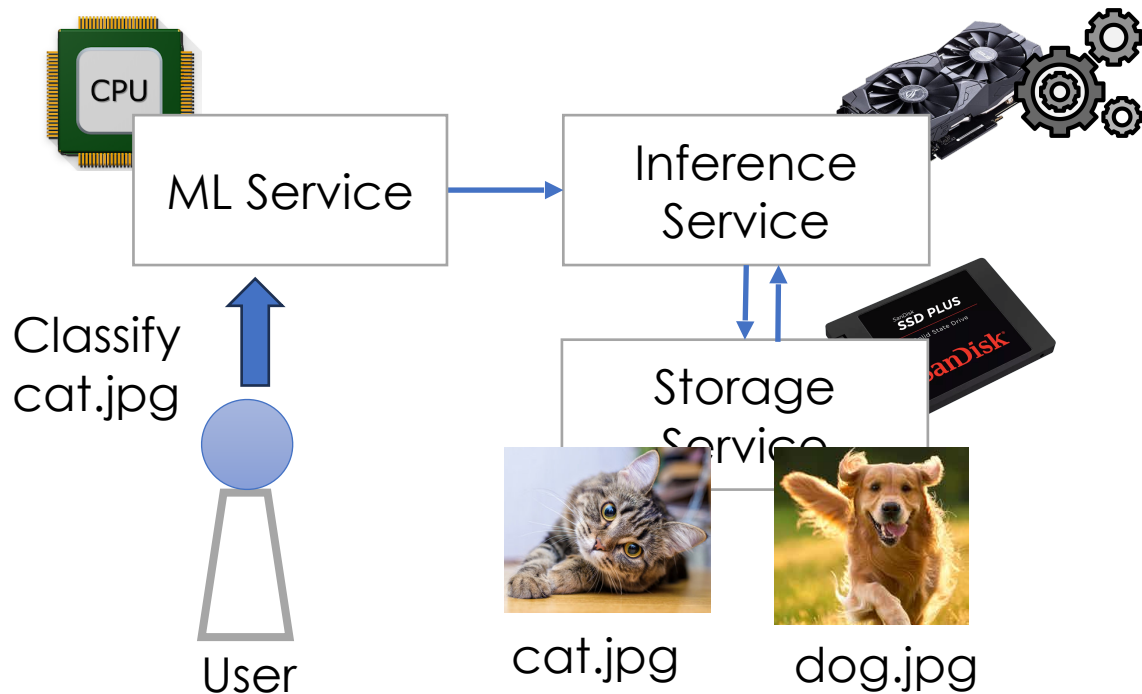
- Disaggregated Service-Oriented Architecture
- Minimizing privilege is essential

# Security in Disaggregated Cloud



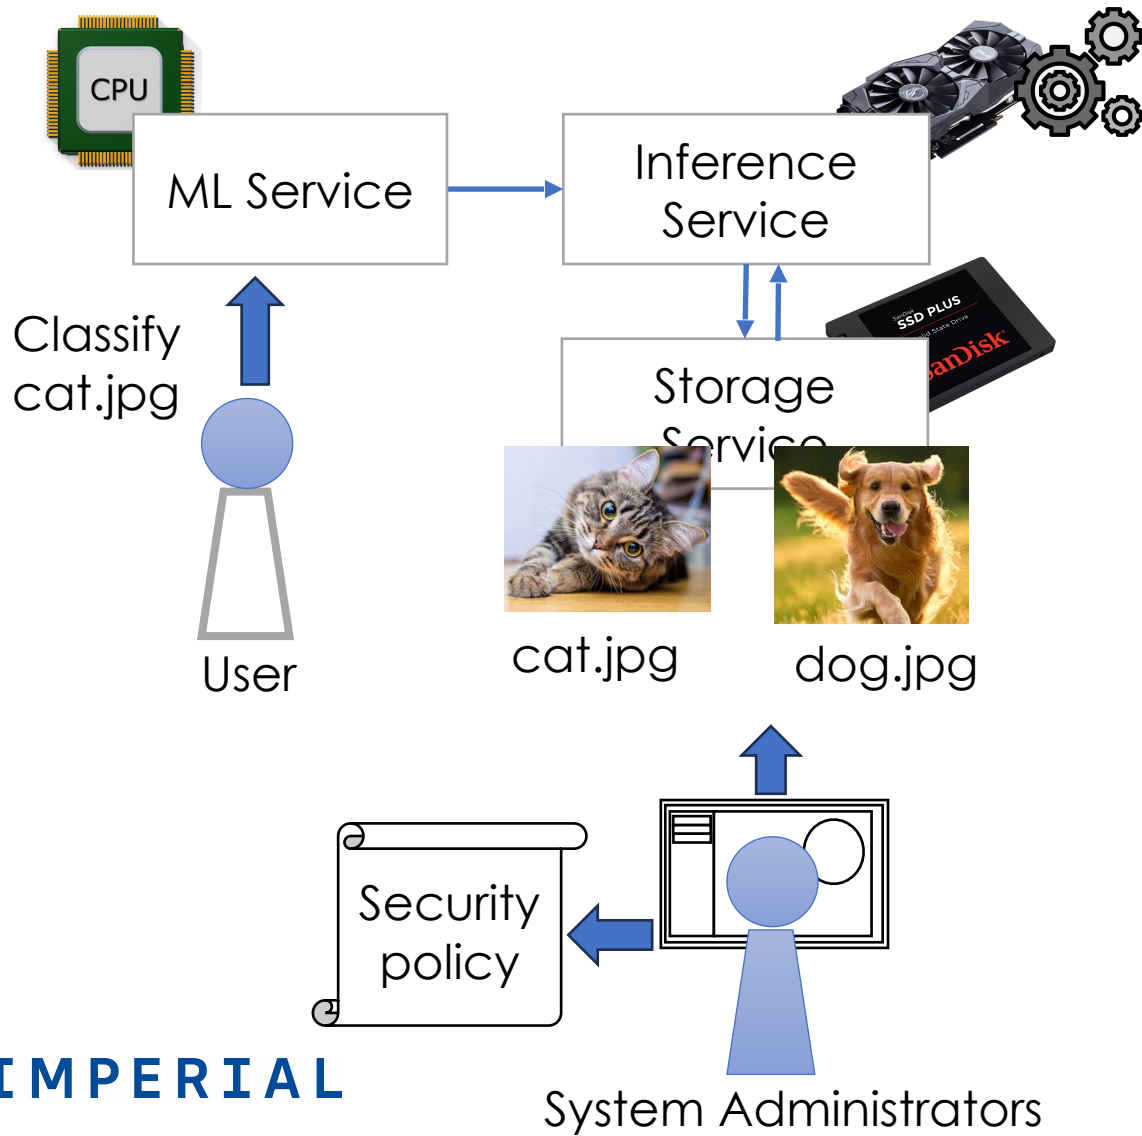
- Disaggregated Service-Oriented Architecture
- Minimizing privilege is essential

# Security in Disaggregated Cloud




- Disaggregated Service-Oriented Architecture
- Minimizing privilege is essential

# Security in Disaggregated Cloud




- Disaggregated Service-Oriented Architecture
- Minimizing privilege is essential

# Security in Disaggregated Cloud



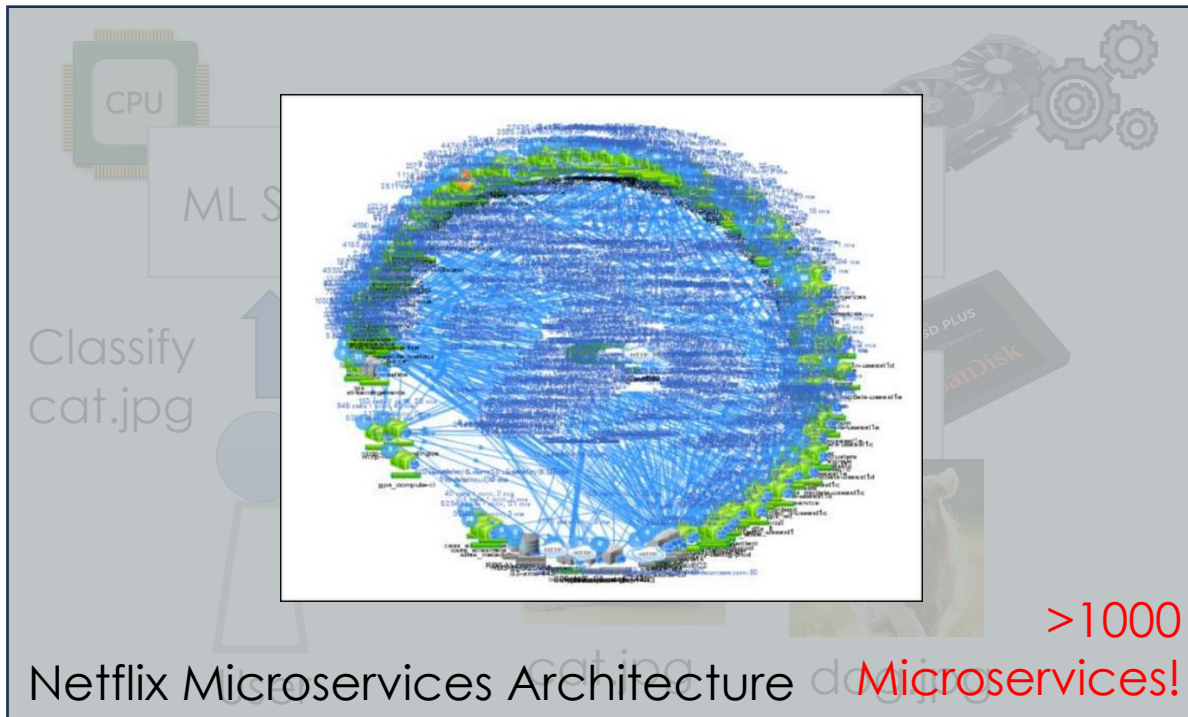
- Disaggregated Service-Oriented Architecture
- Minimizing privilege is essential



System Administrators

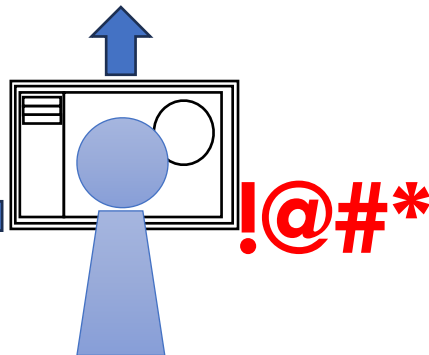



# Security in Disaggregated Cloud




- Disaggregated Service-Oriented Architecture
- Minimizing privilege is essential
- Policies don't scale to complex systems -> over-privilege

Minimizing Privileges Infeasible



System Administrators

# Capabilities for Disaggregated Cloud



ML Service

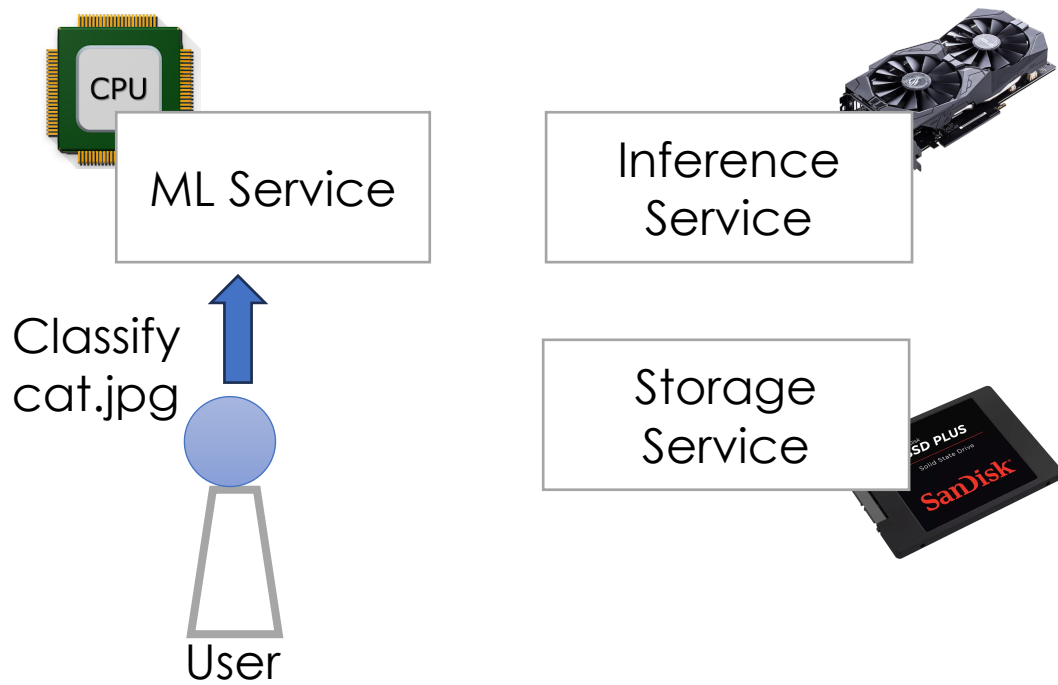


Inference Service

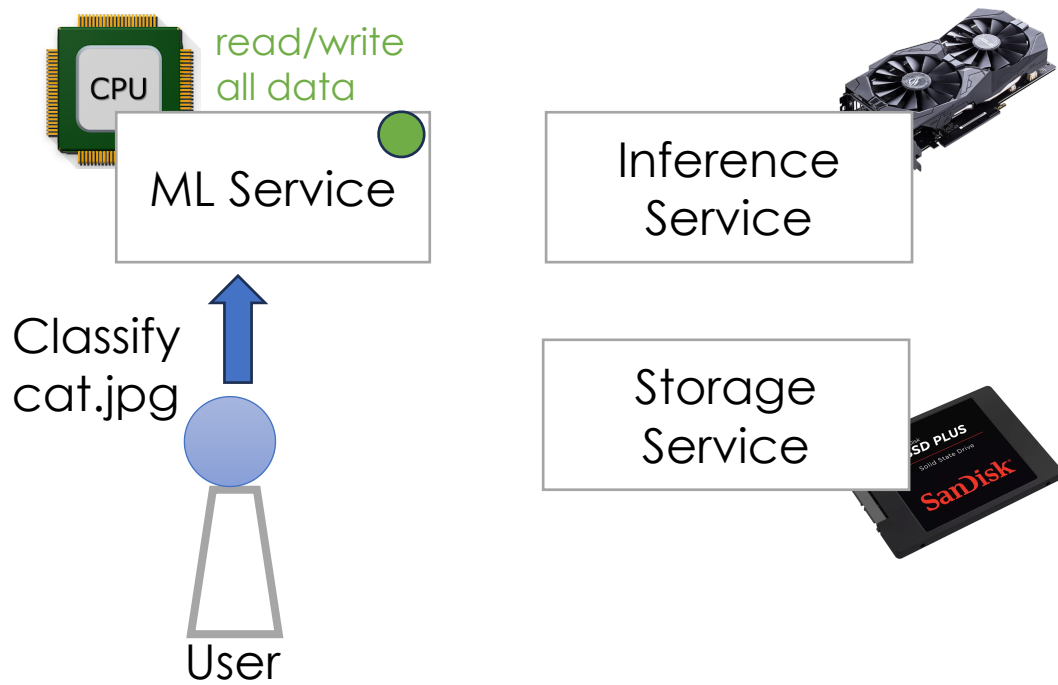
Storage Service



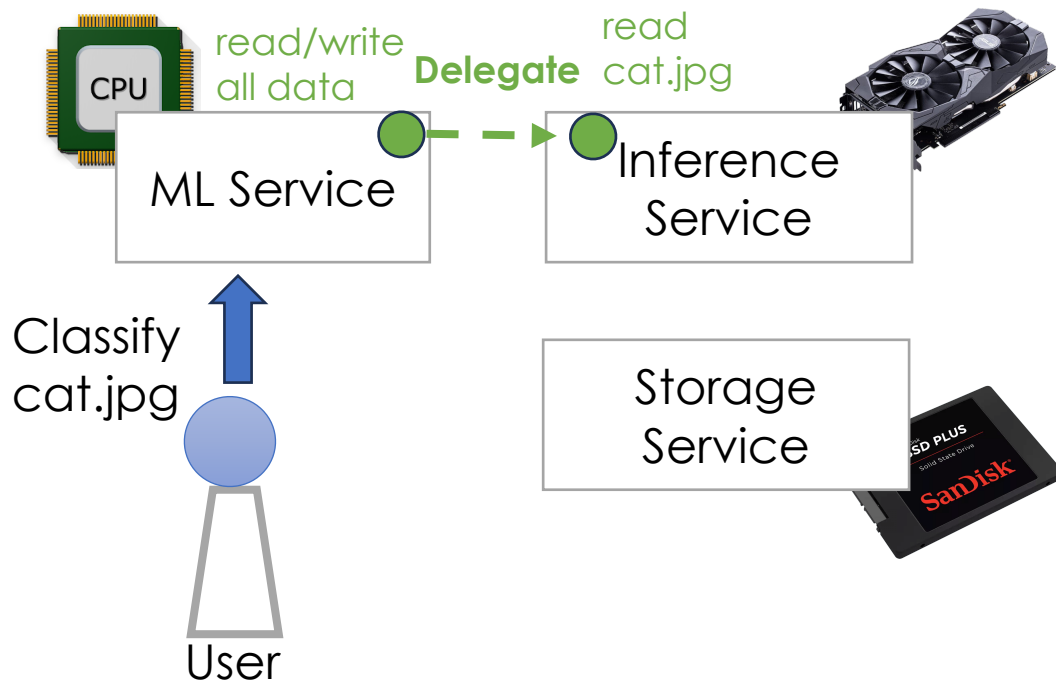
# Capabilities for Disaggregated Cloud



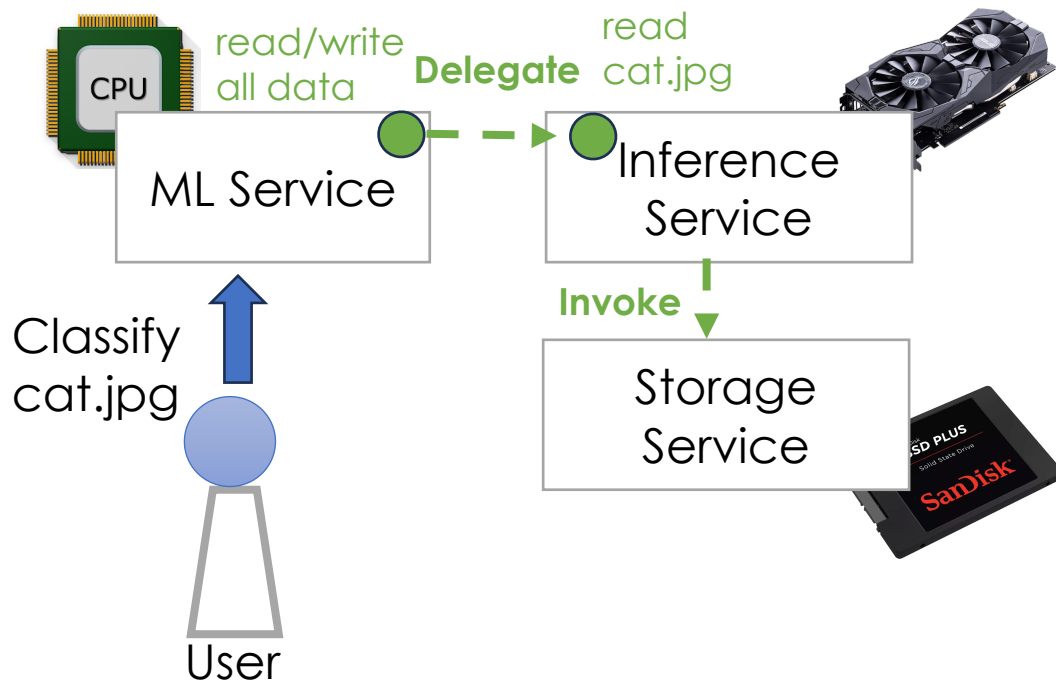
# Capabilities for Disaggregated Cloud



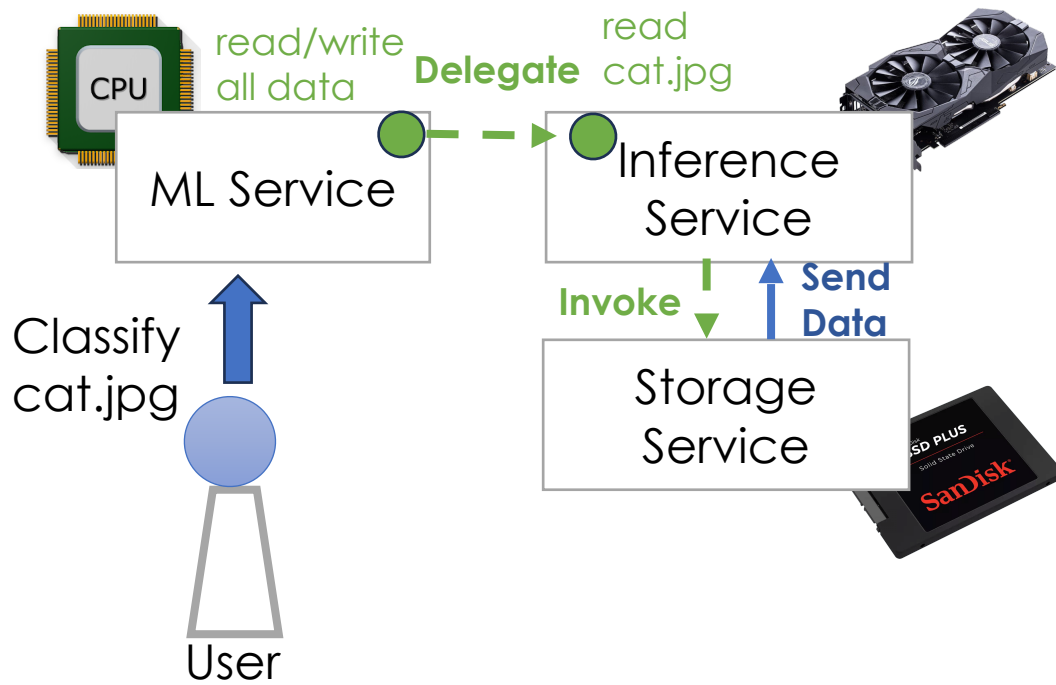
# Capabilities for Disaggregated Cloud



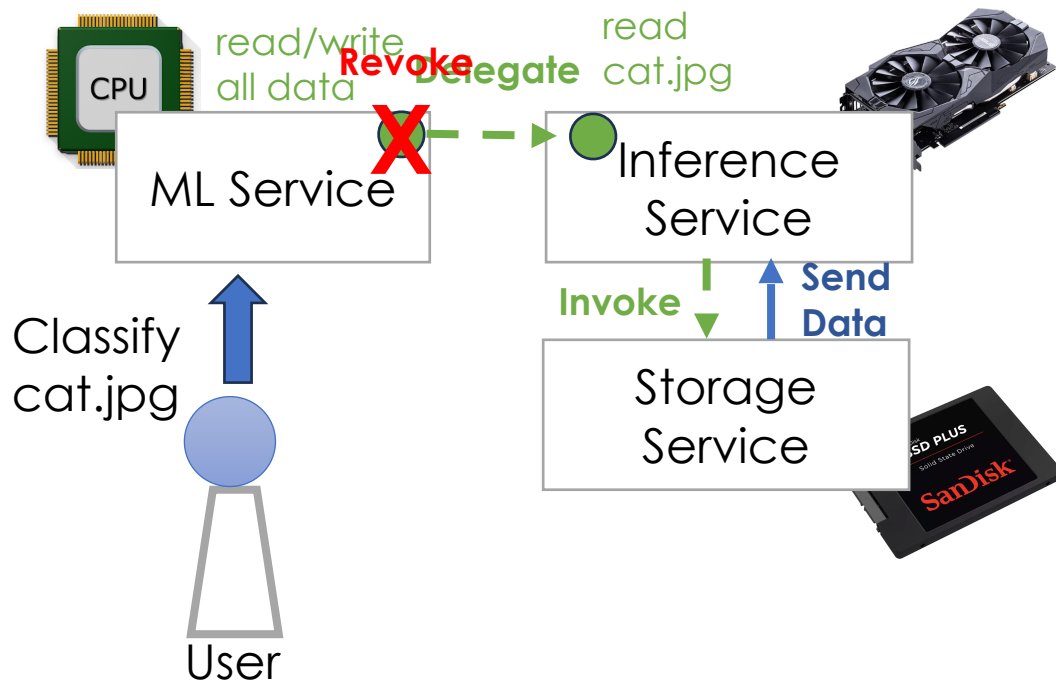
# Capabilities for Disaggregated Cloud



# Capabilities for Disaggregated Cloud

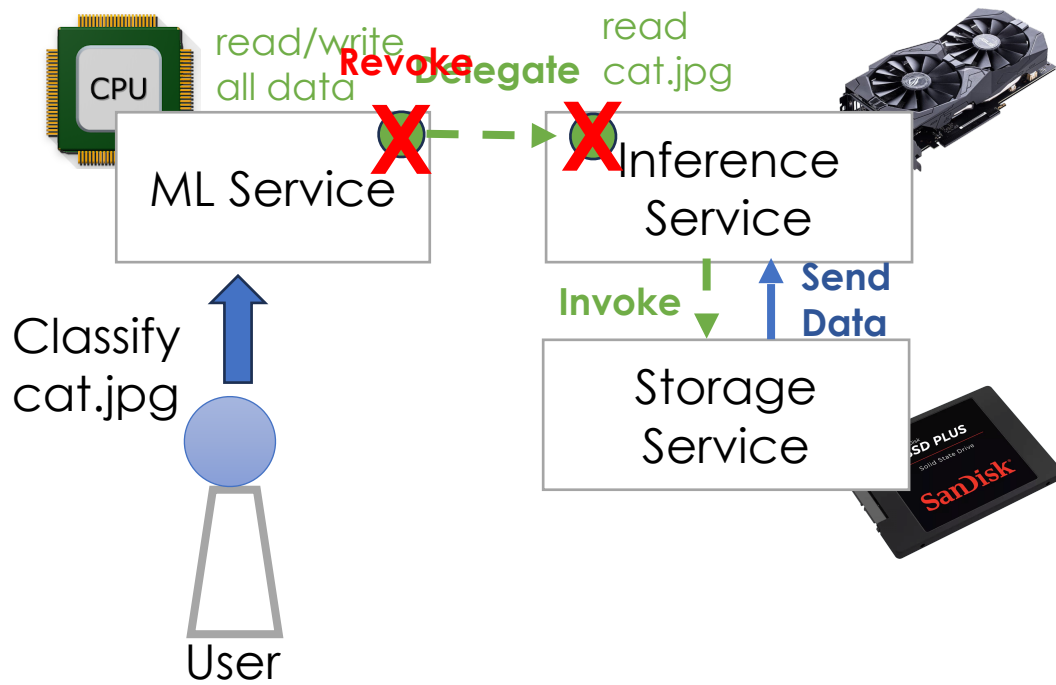


# Capabilities for Disaggregated Cloud



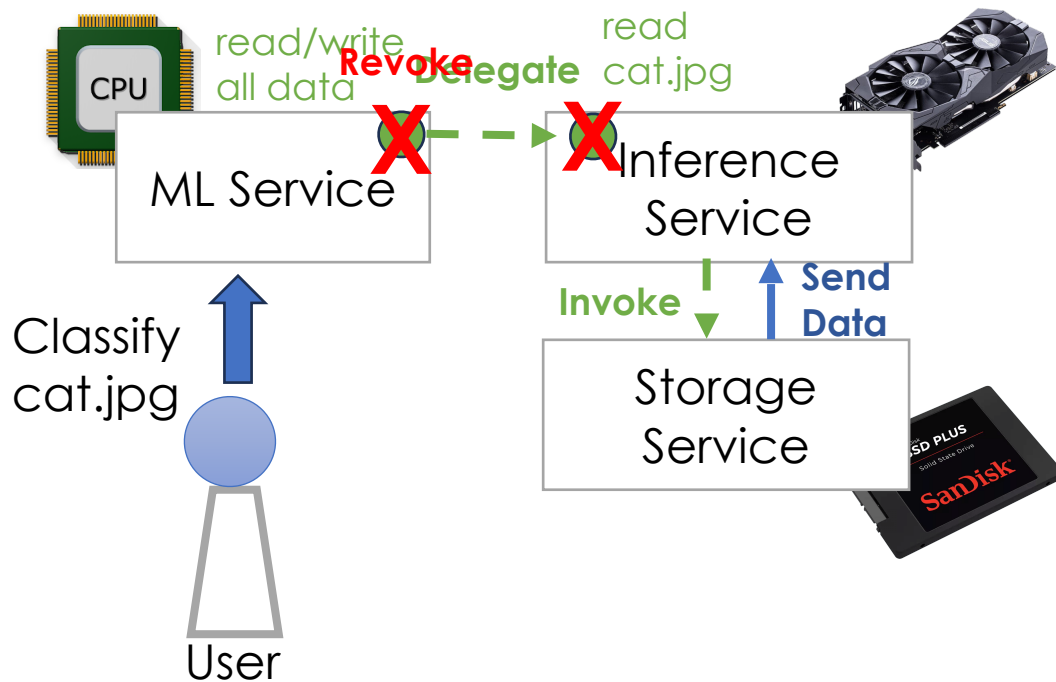


# Capabilities for Disaggregated Cloud



# Capabilities for Disaggregated Cloud

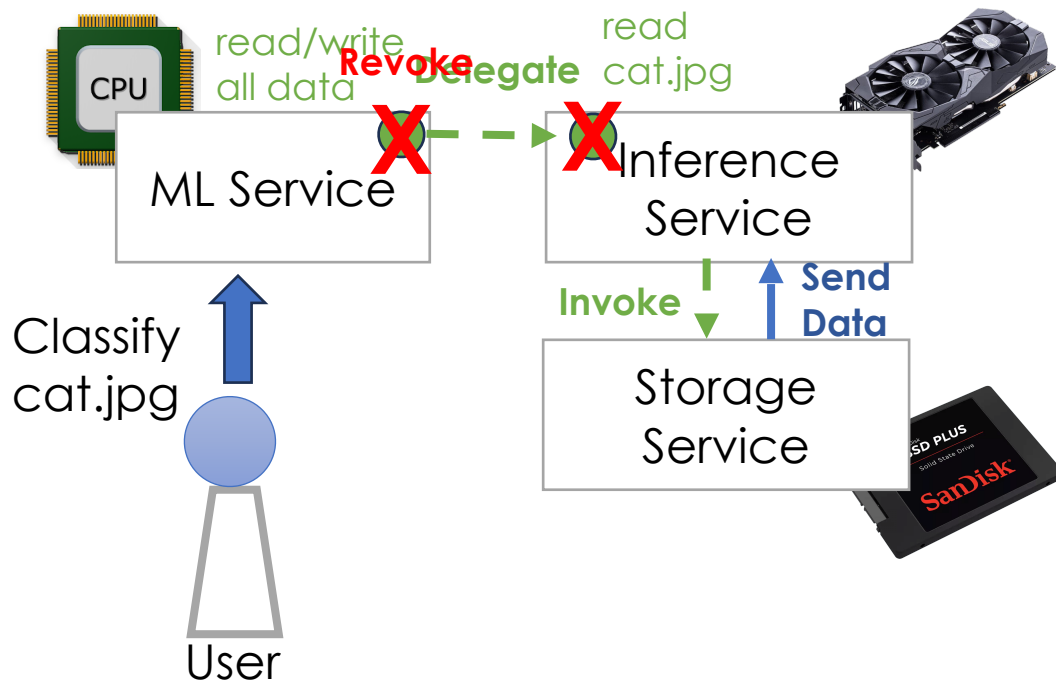
- Application controlled permissions



# Capabilities for Disaggregated Cloud

- Application controlled permissions

– Policy is distributed

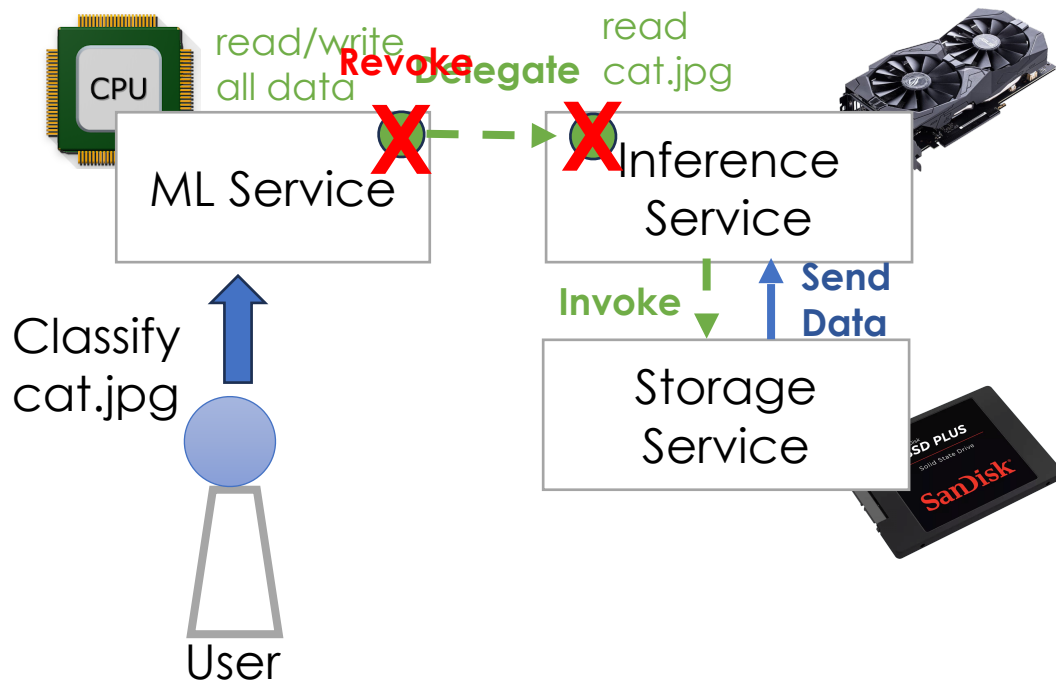


# Capabilities for Disaggregated Cloud


- Application controlled permissions

– Policy is distributed

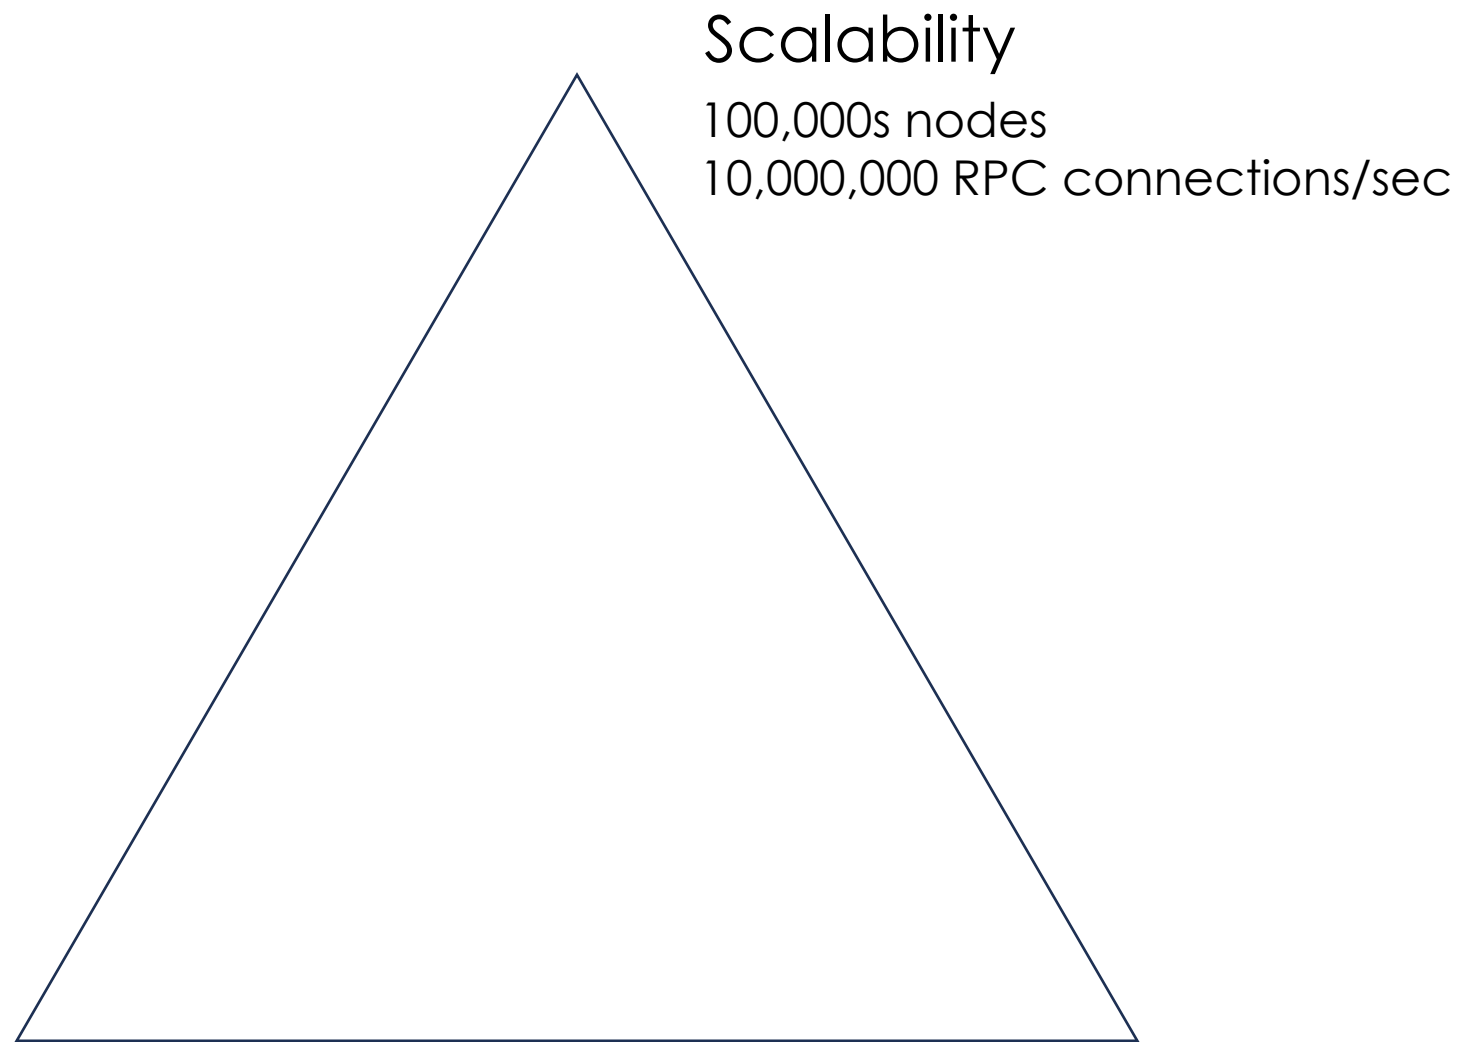
– Privilege minimization scales better to complex systems



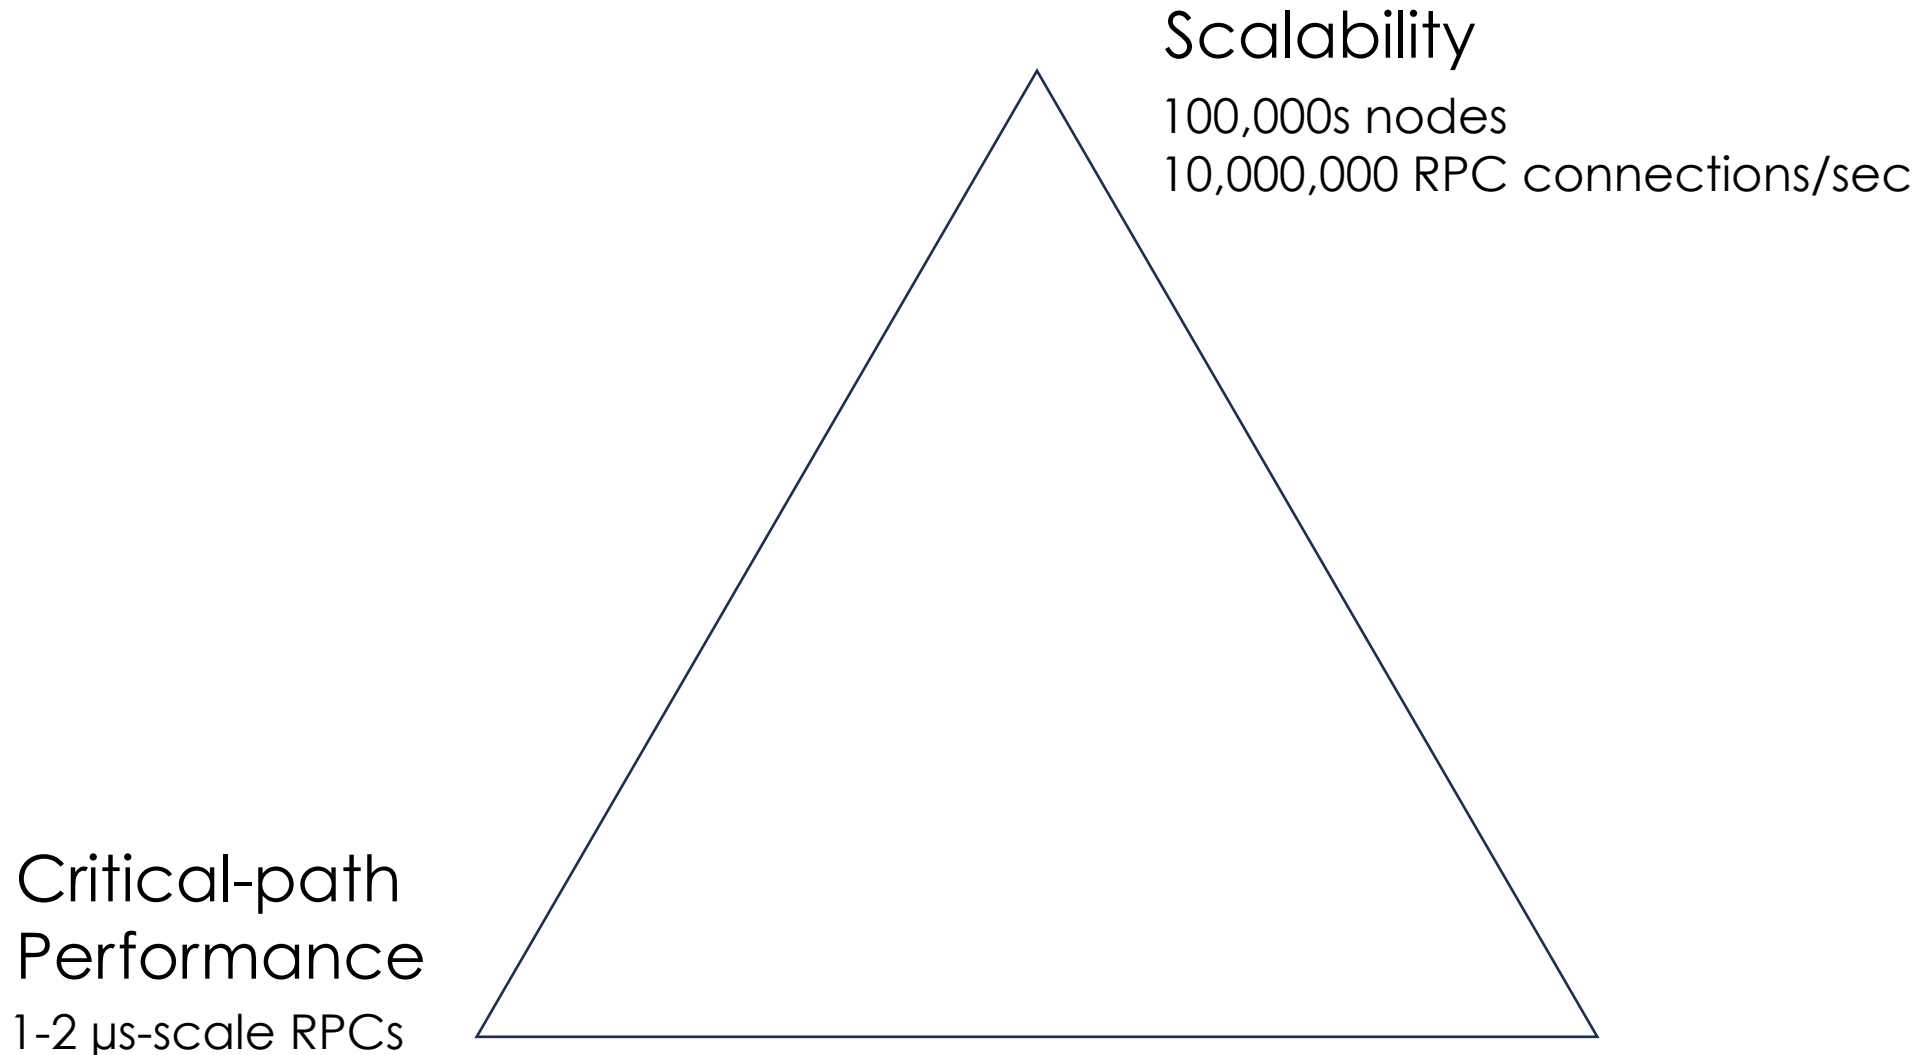
# No Capability Systems Are Cloud Ready



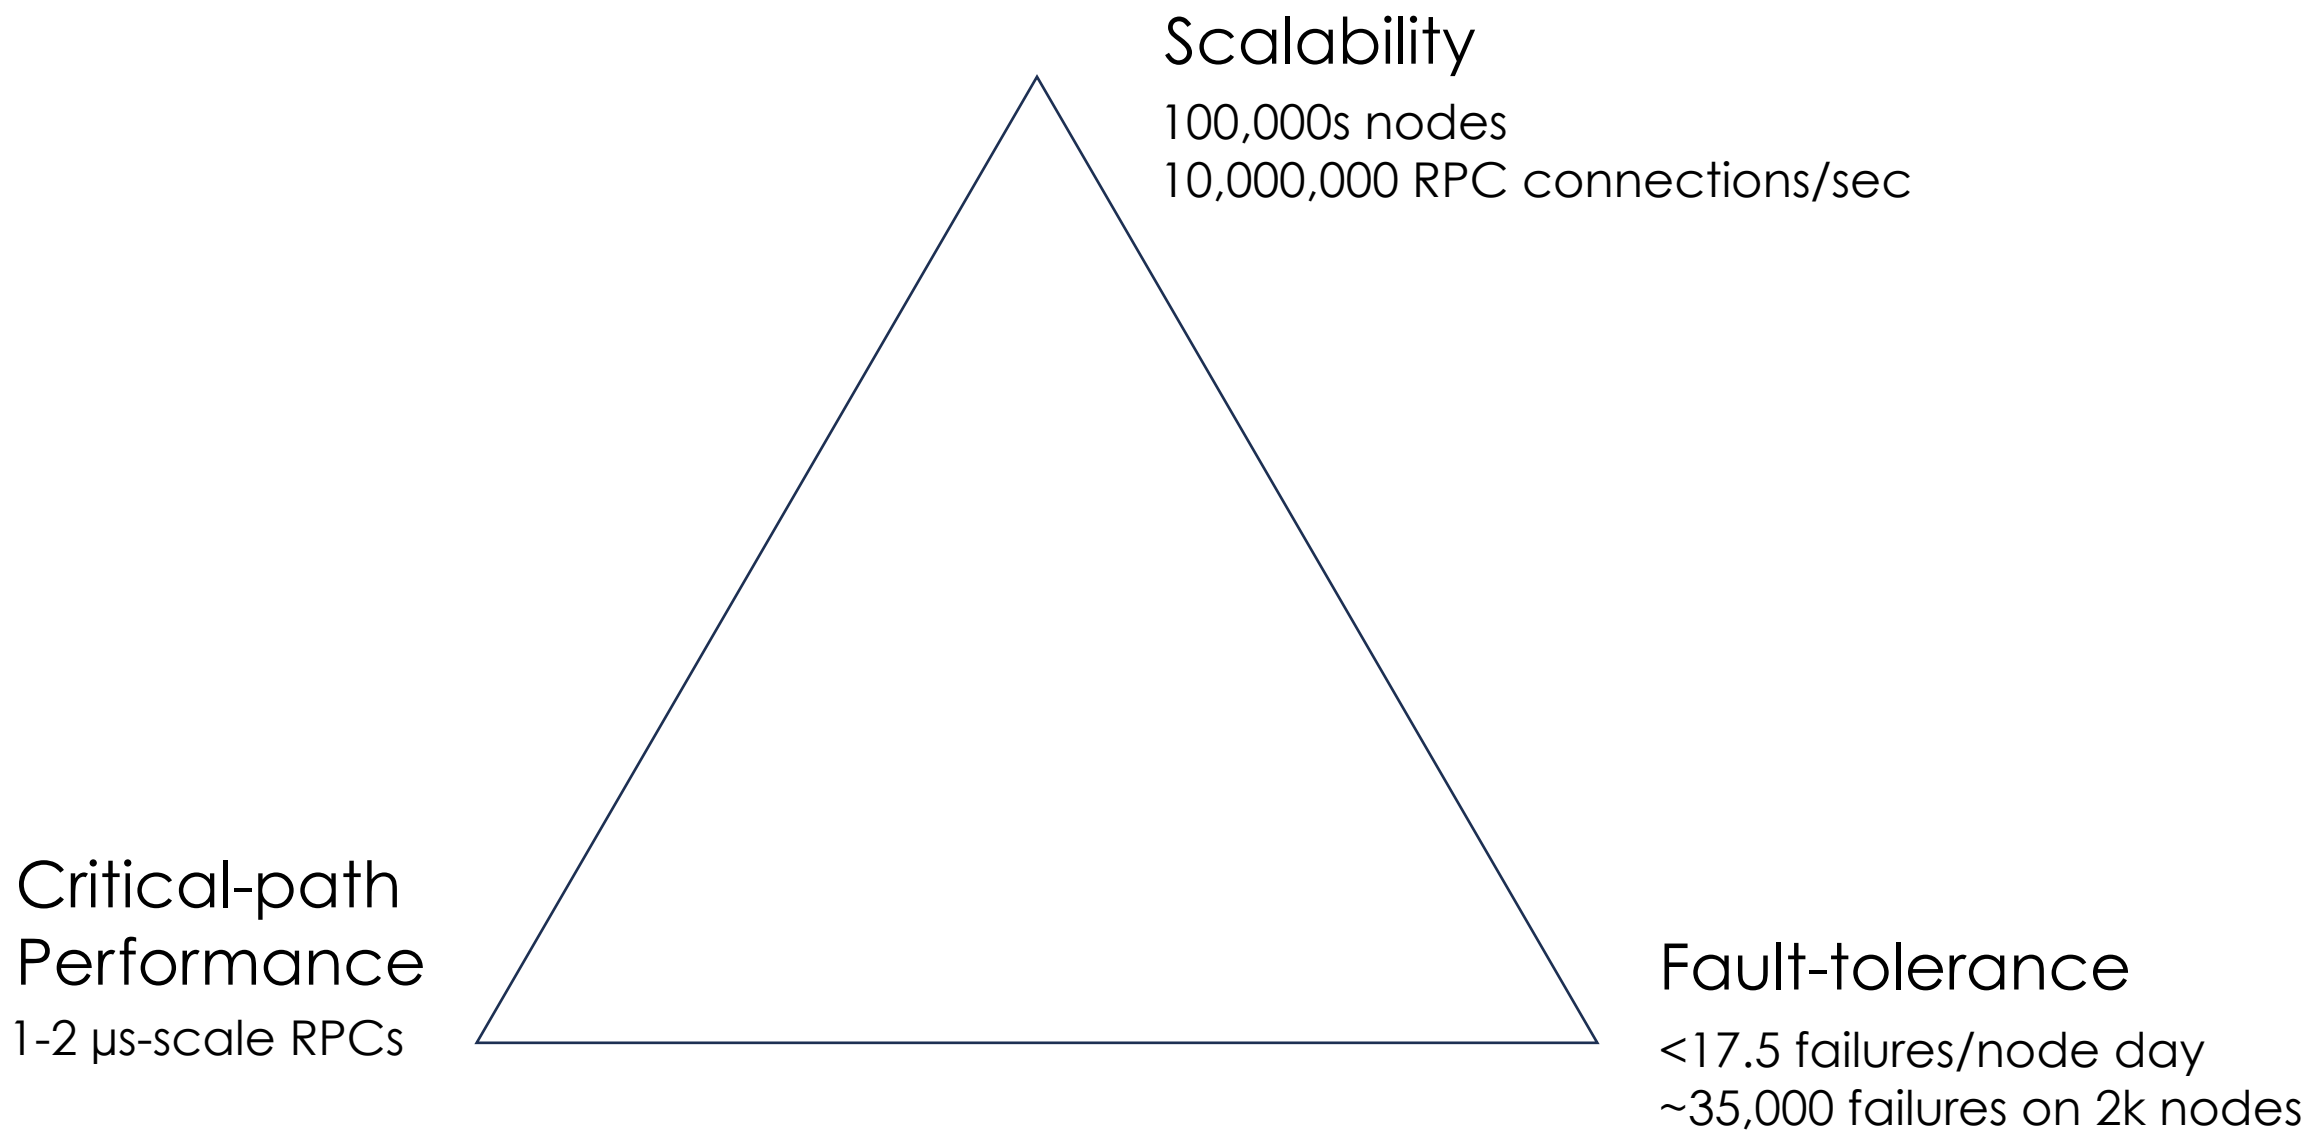
# No Capability Systems Are Cloud Ready



# No Capability Systems Are Cloud Ready

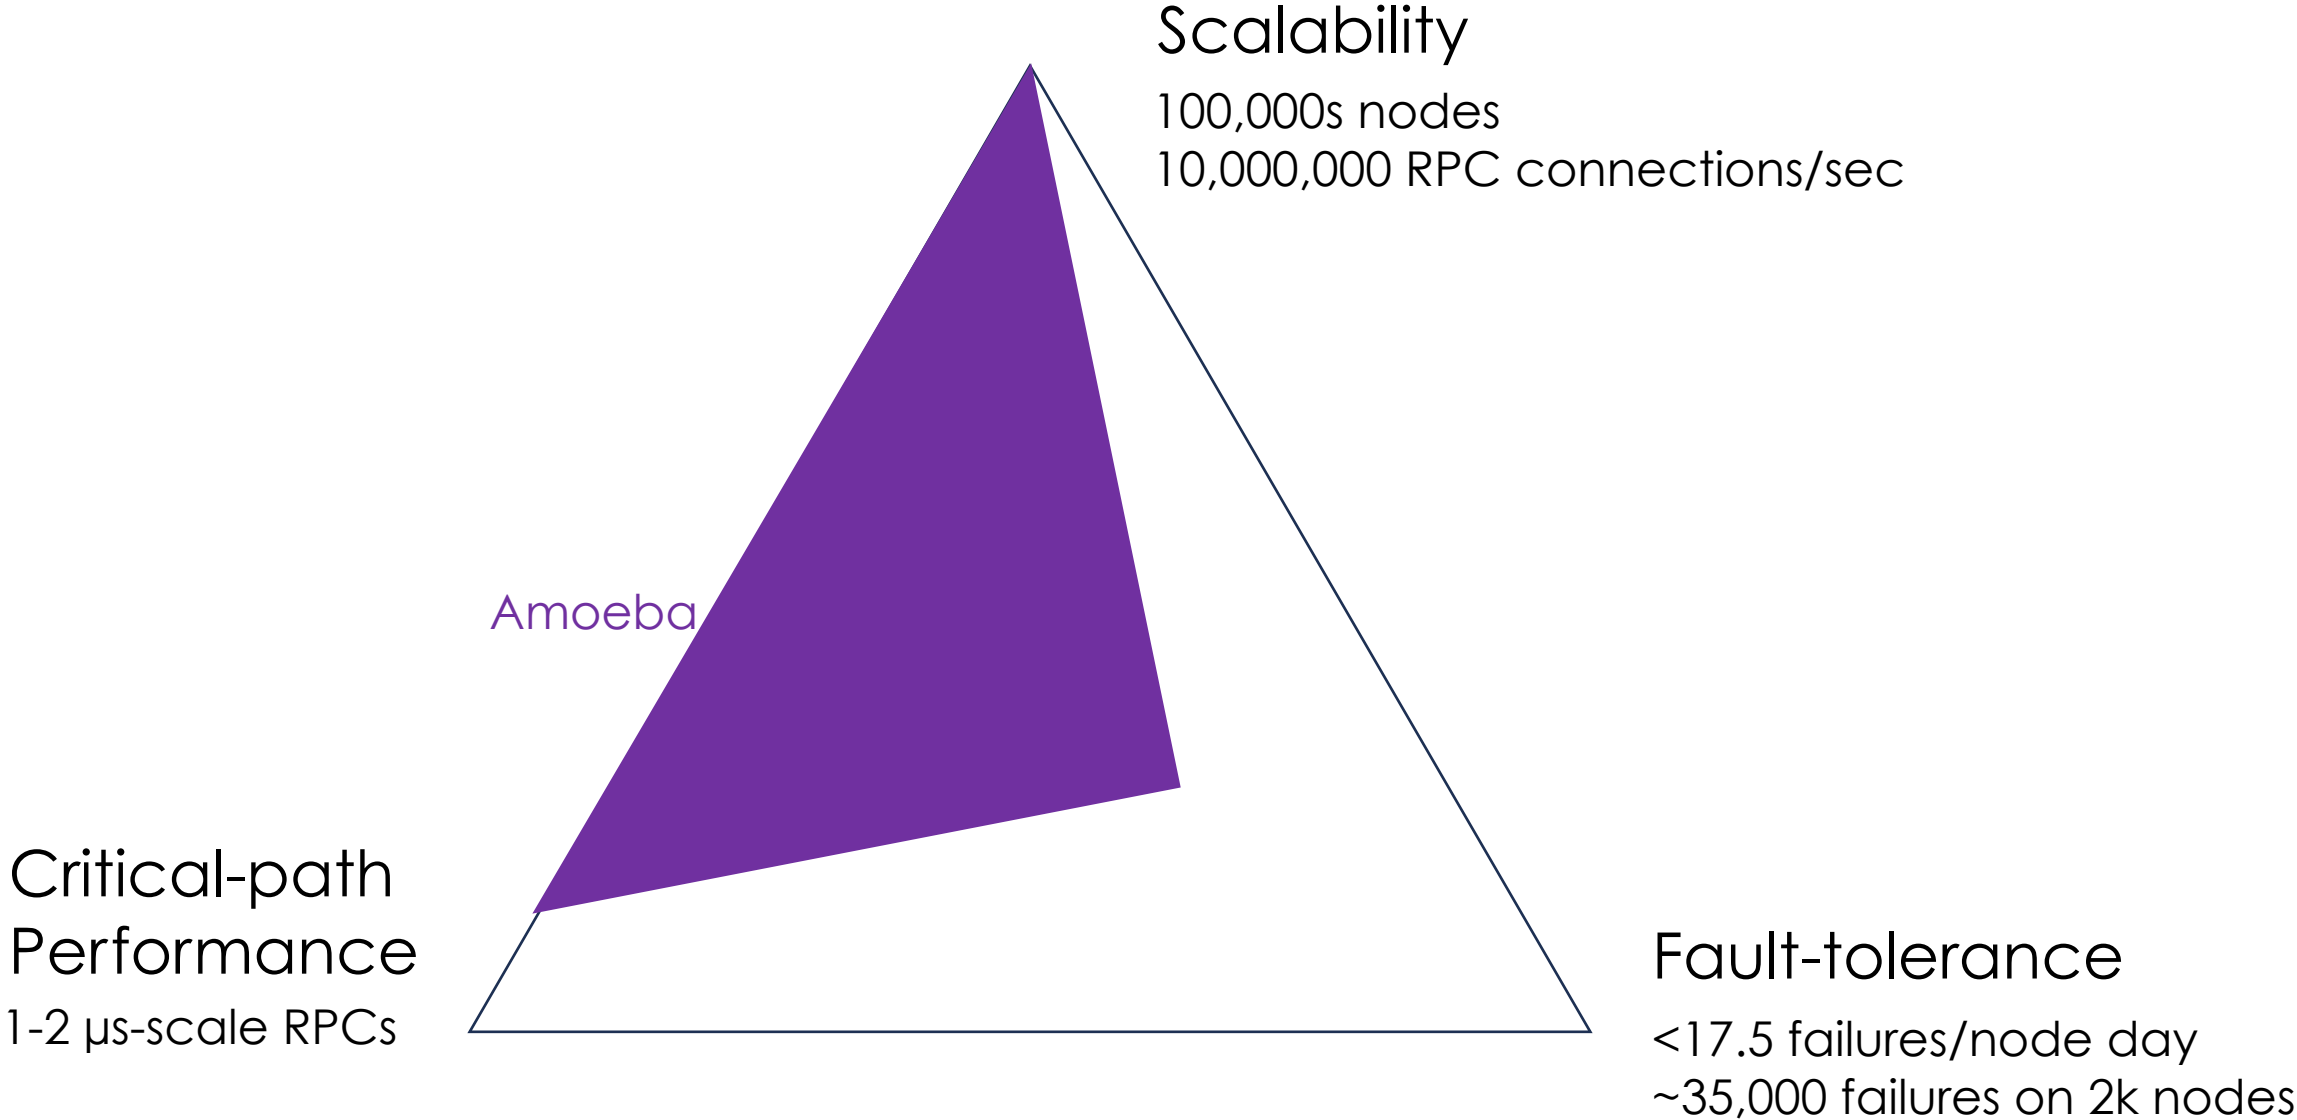


# No Capability Systems Are Cloud Ready

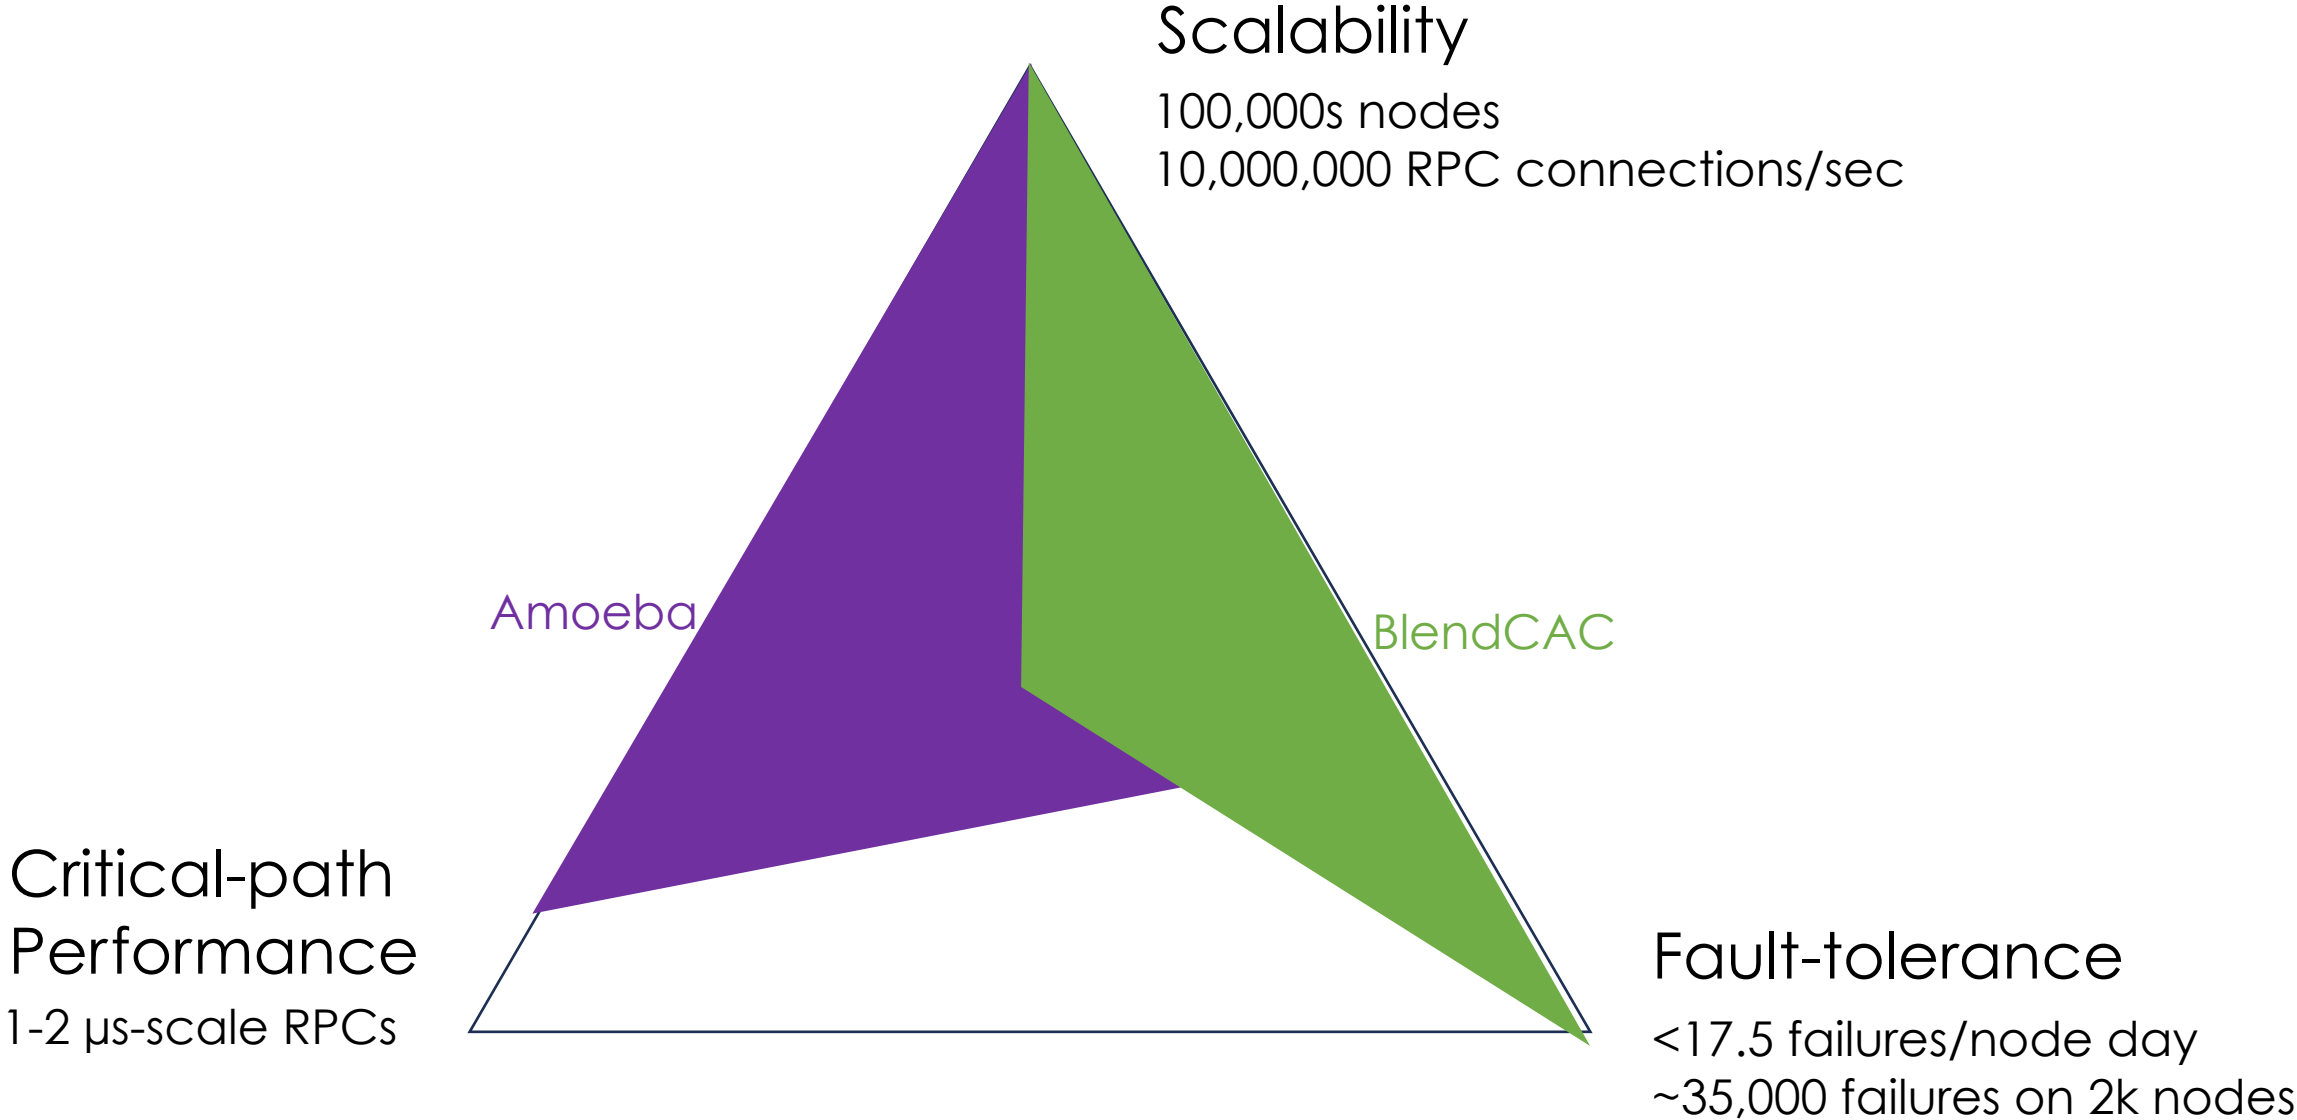




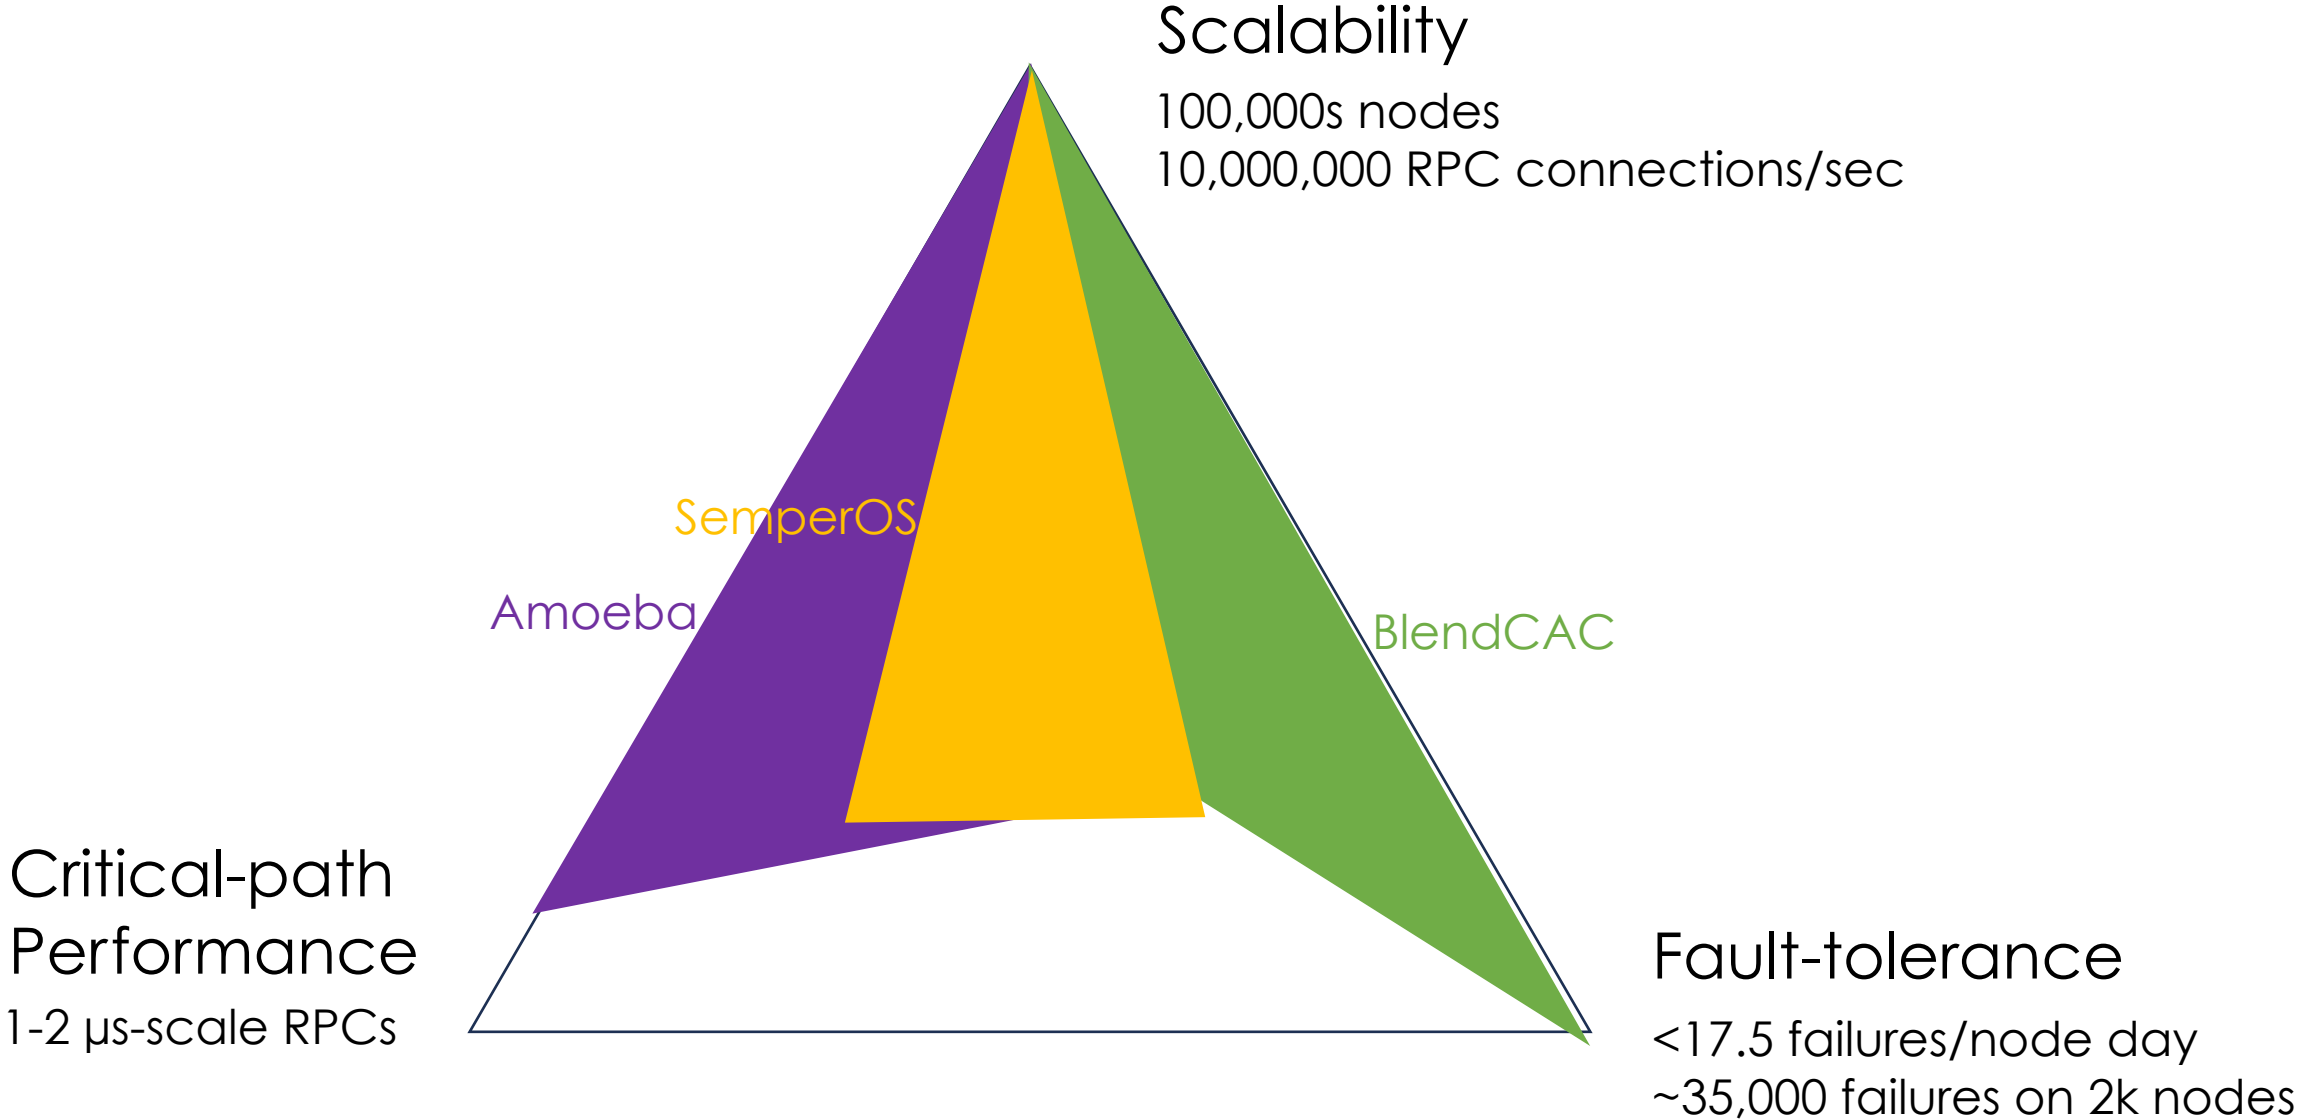
# No Capability Systems Are Cloud Ready



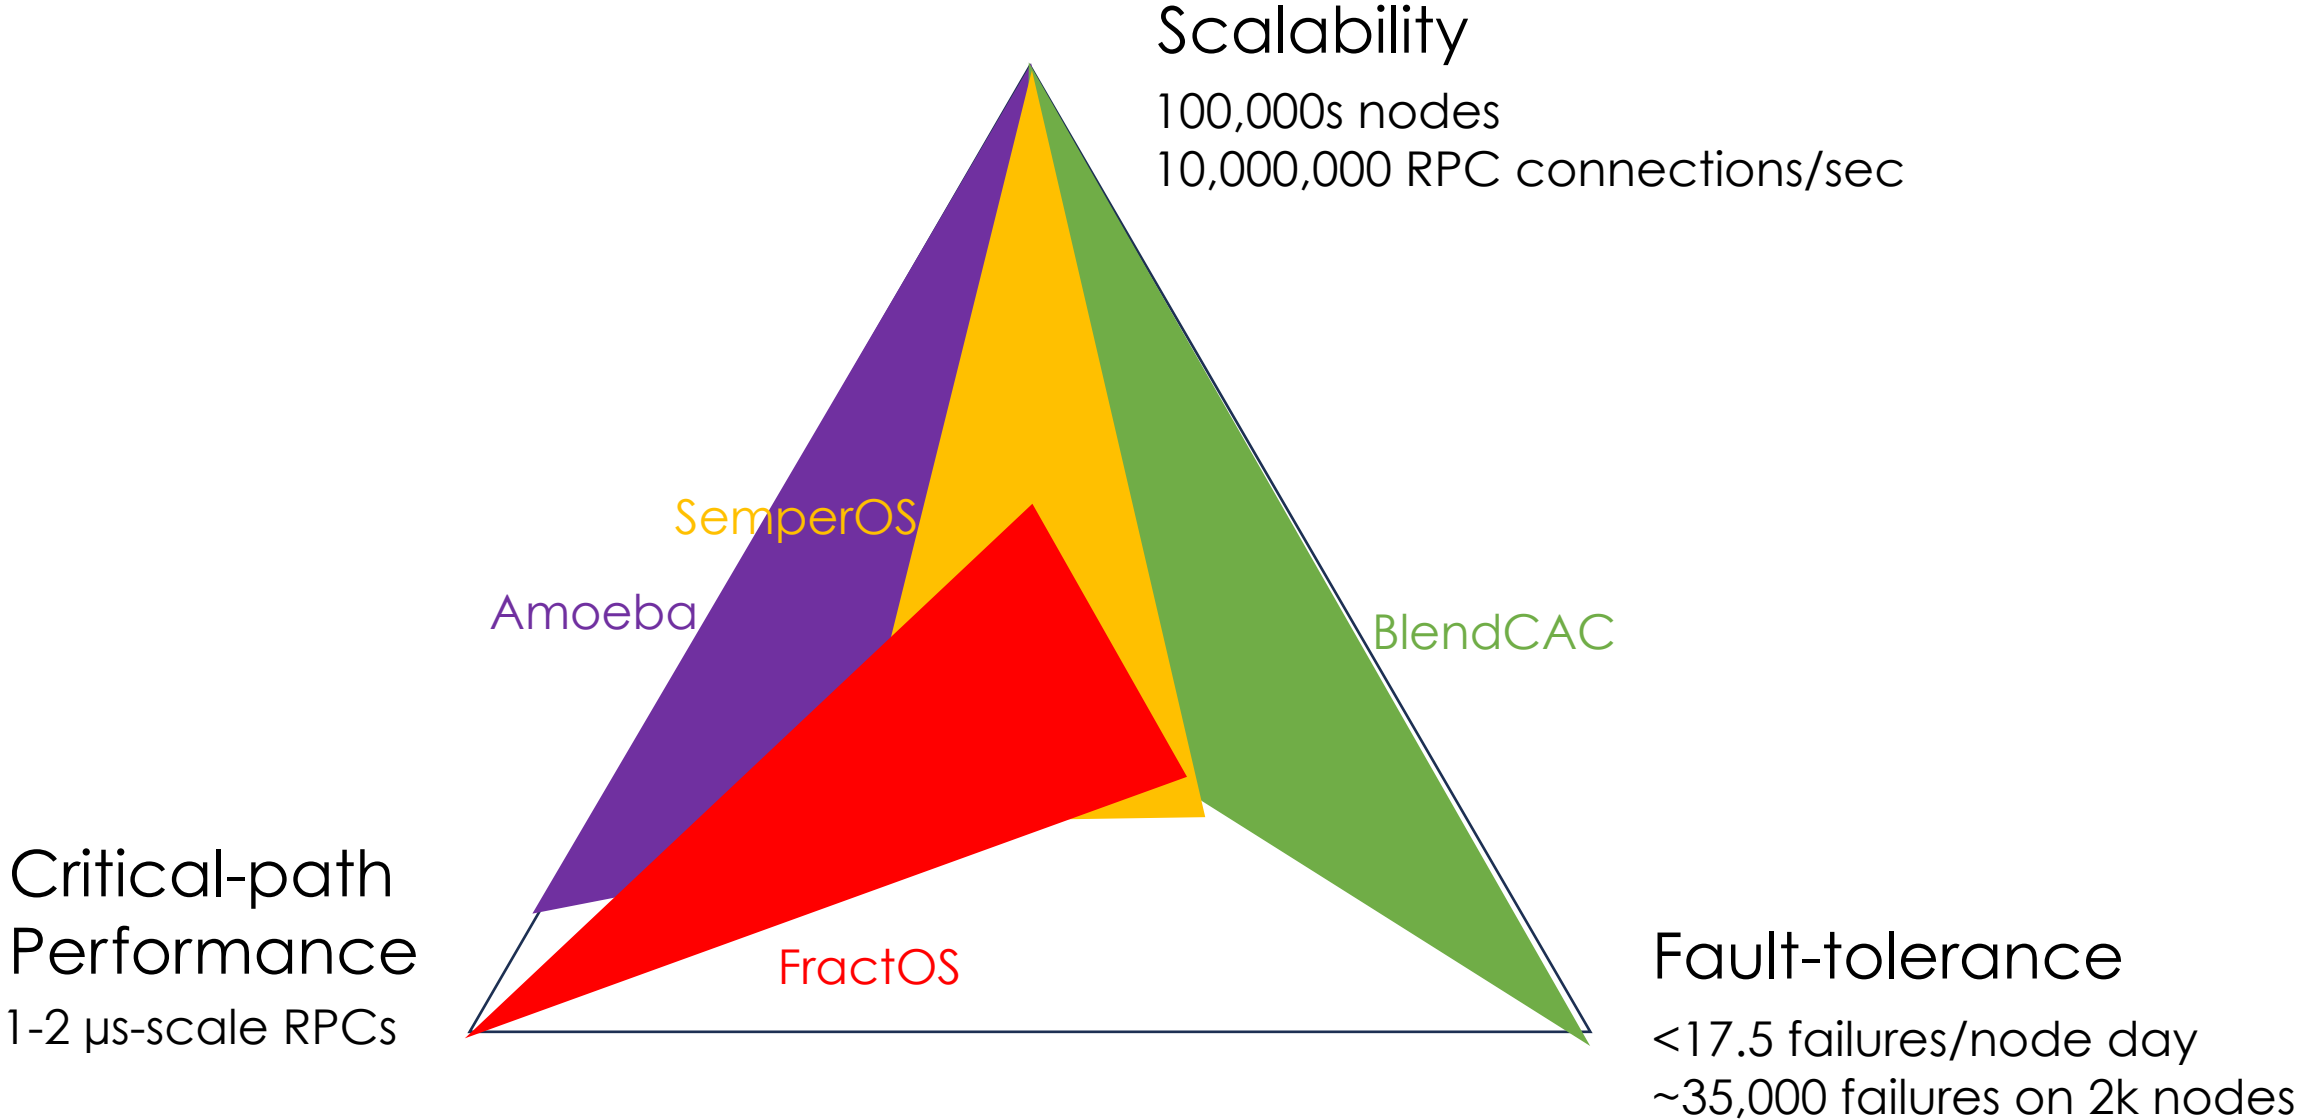
# No Capability Systems Are Cloud Ready



# No Capability Systems Are Cloud Ready



# No Capability Systems Are Cloud Ready



# No Capability Systems Are Cloud Ready

We need all three to achieve cloud scale...

Scalability  
100,000s nodes  
10,000,000 RPC connections/sec



Critical-path Performance  
1-2  $\mu$ s-scale RPCs

Fault-tolerance  
<17.5 failures/node day  
~35,000 failures on 2k nodes

# Problem: Centralized Metadata in Existing Capability Systems

GPU Service

Storage  
Service

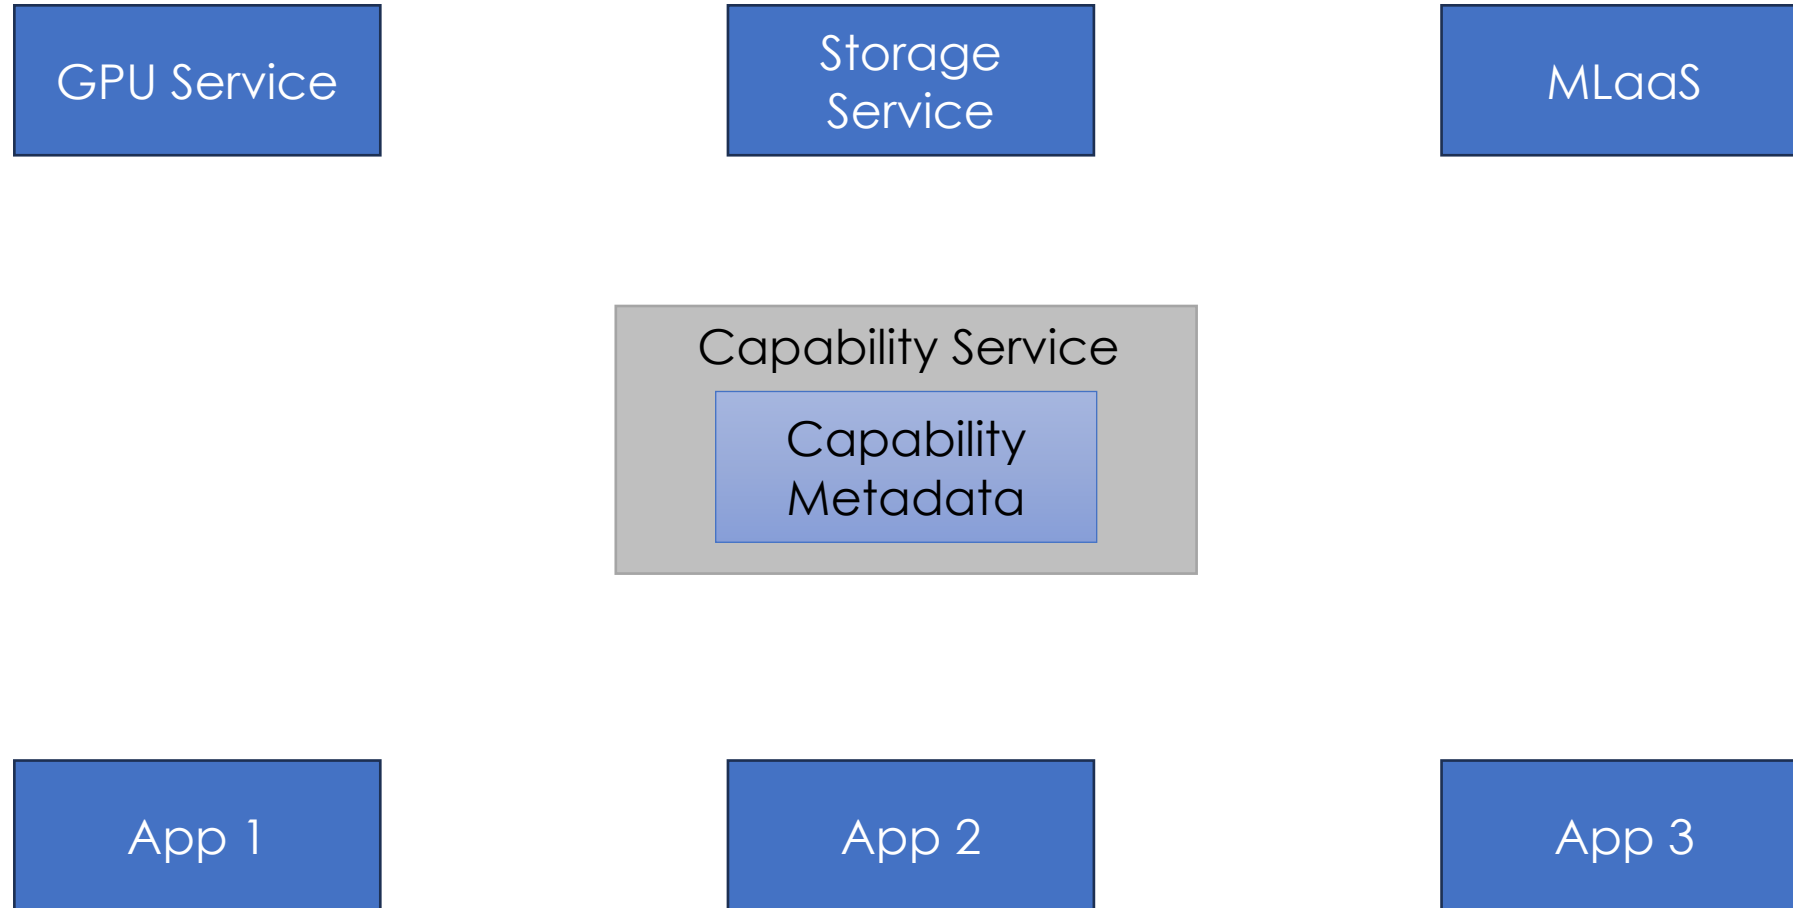
MLaaS

App 1

App 2

App 3

# Problem: Centralized Metadata in Existing Capability Systems



# Problem: Centralized Metadata in Existing Capability Systems

GPU Service

Storage Service

MLaaS



- Ownership, permissions, ...
- Capability operation processing

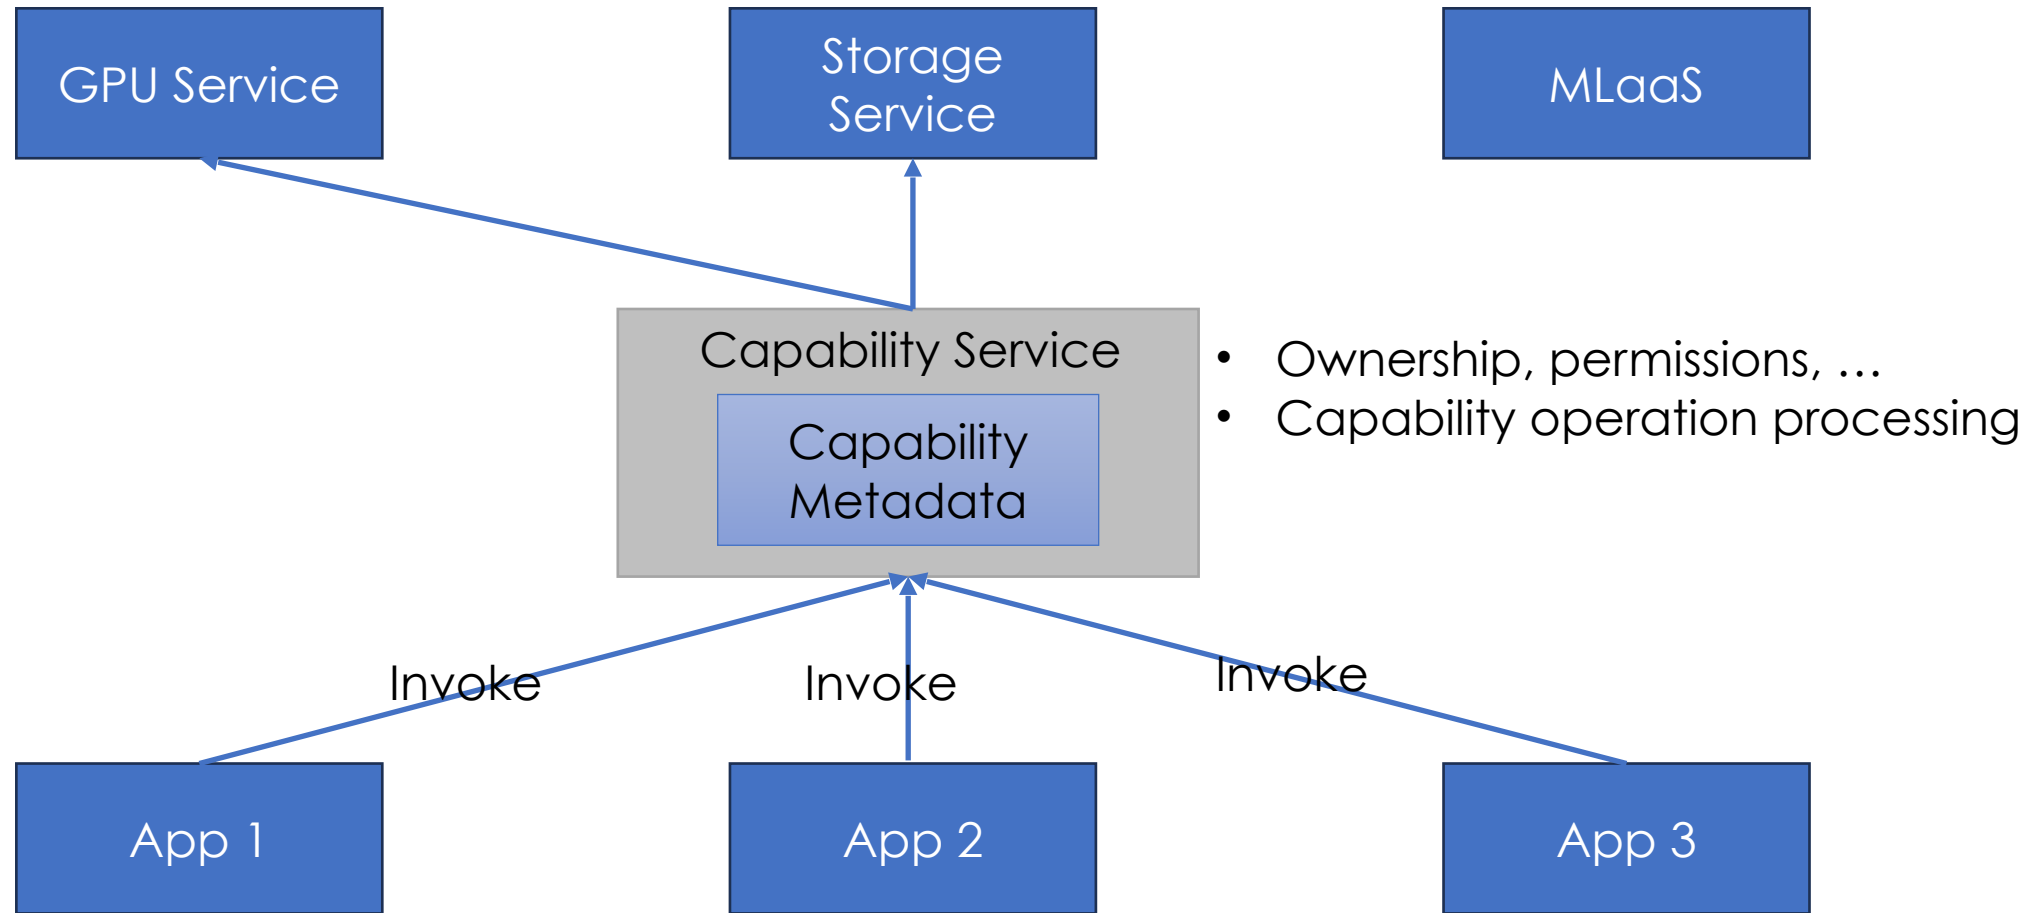
App 1

App 2

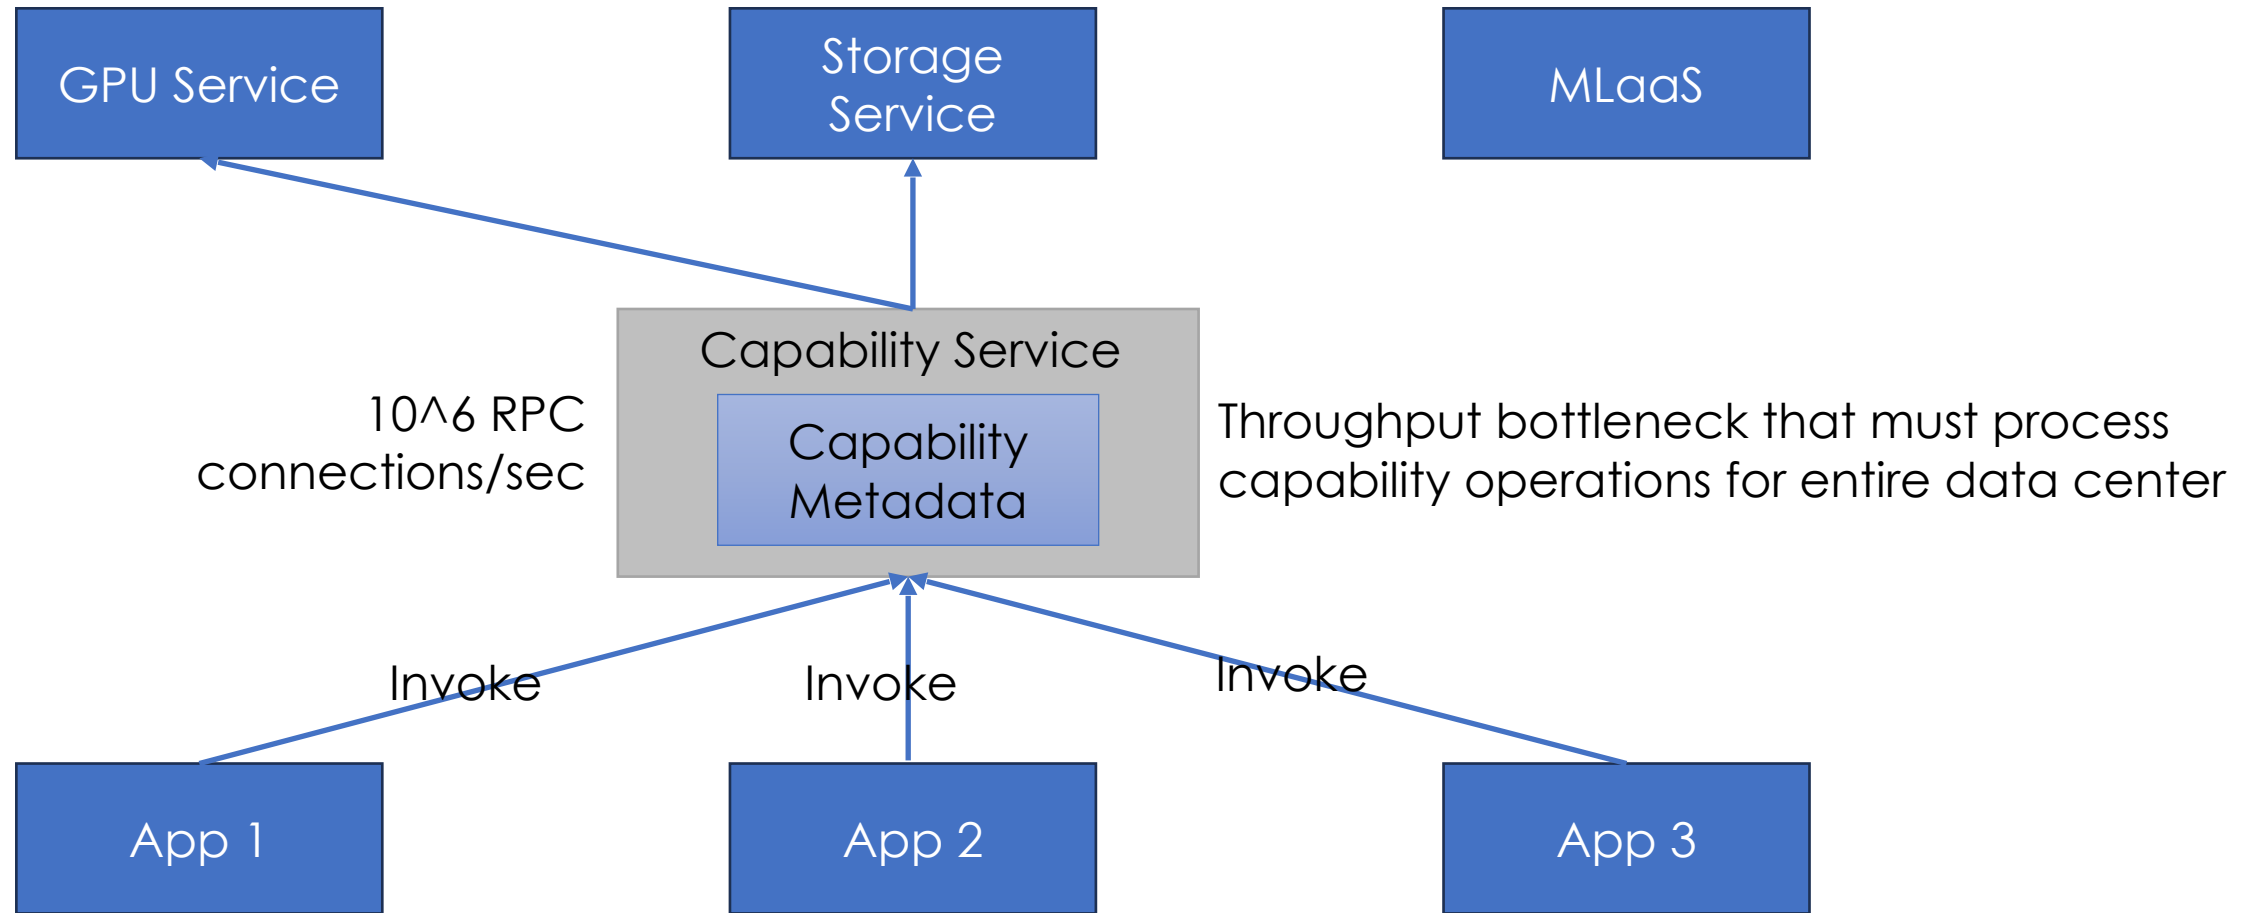
App 3



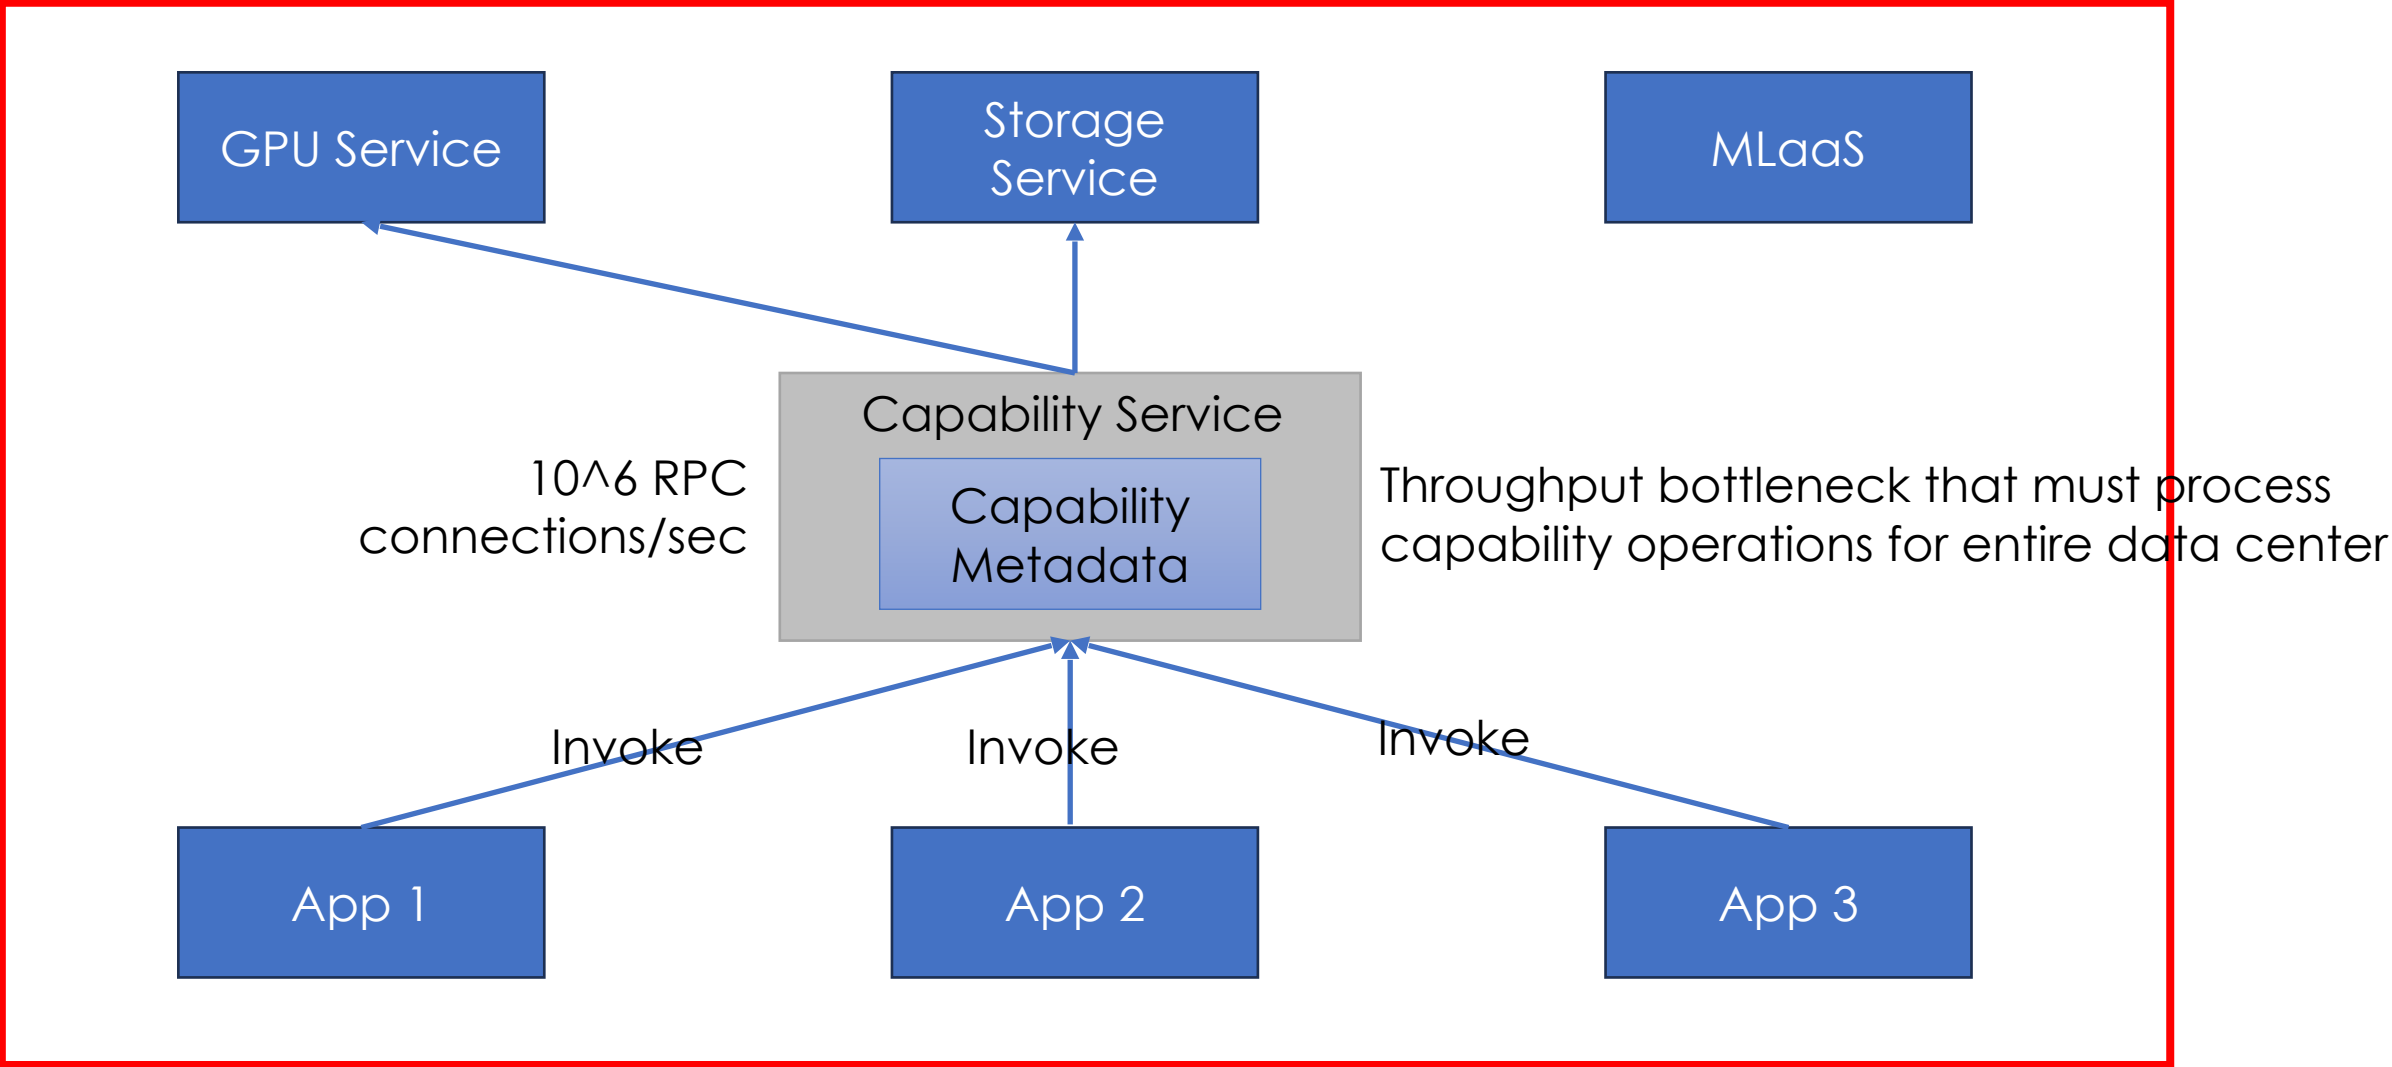
# Problem: Centralized Metadata in Existing Capability Systems



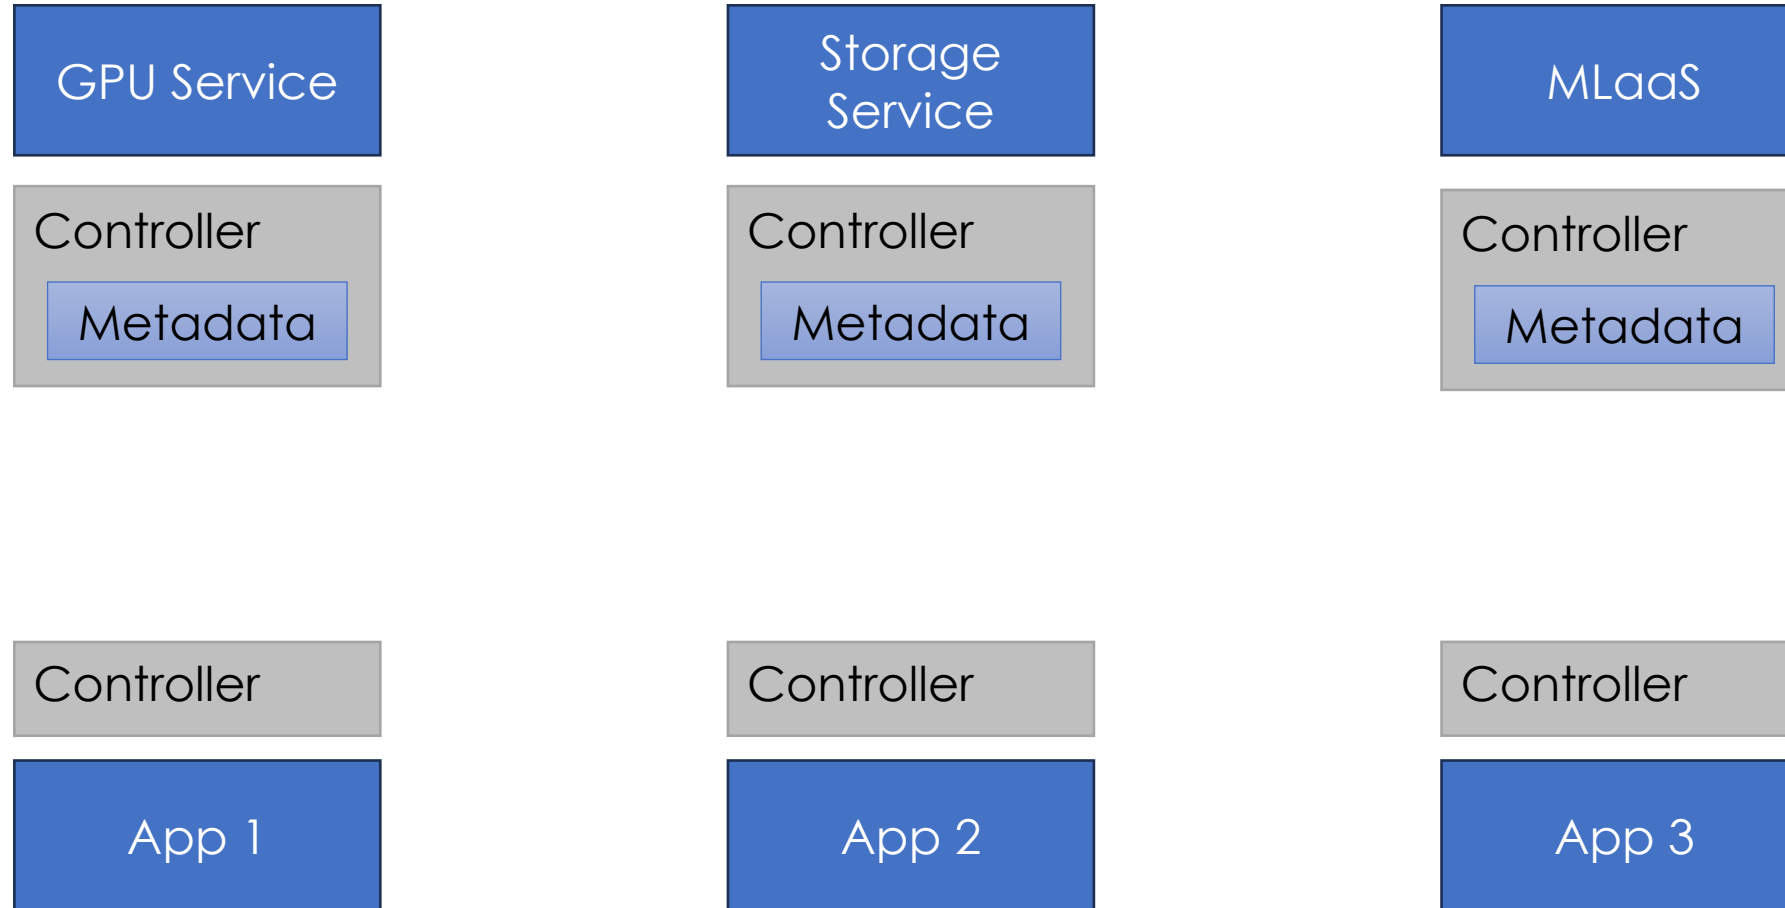
# Problem: Centralized Metadata in Existing Capability Systems



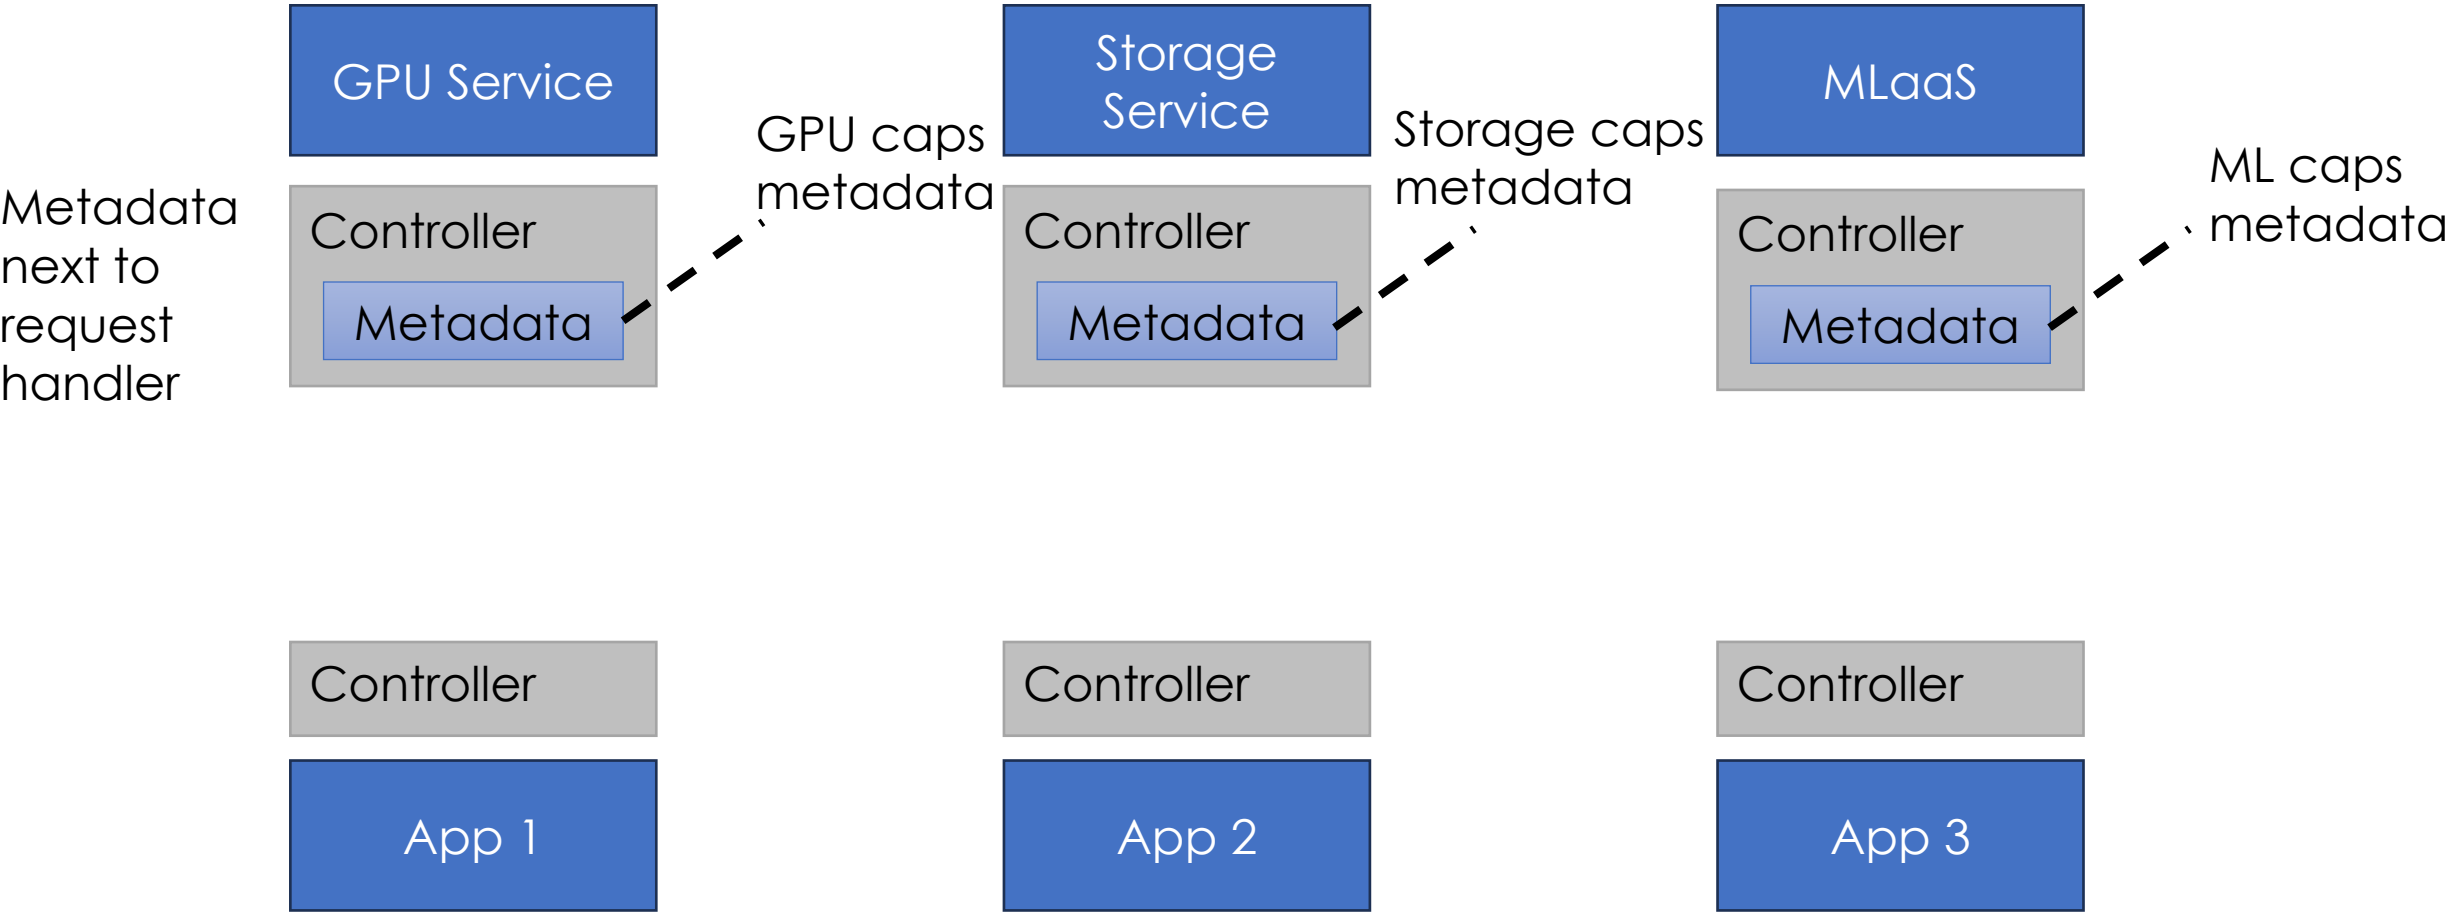
# Problem: Centralized Metadata in Existing Capability Systems



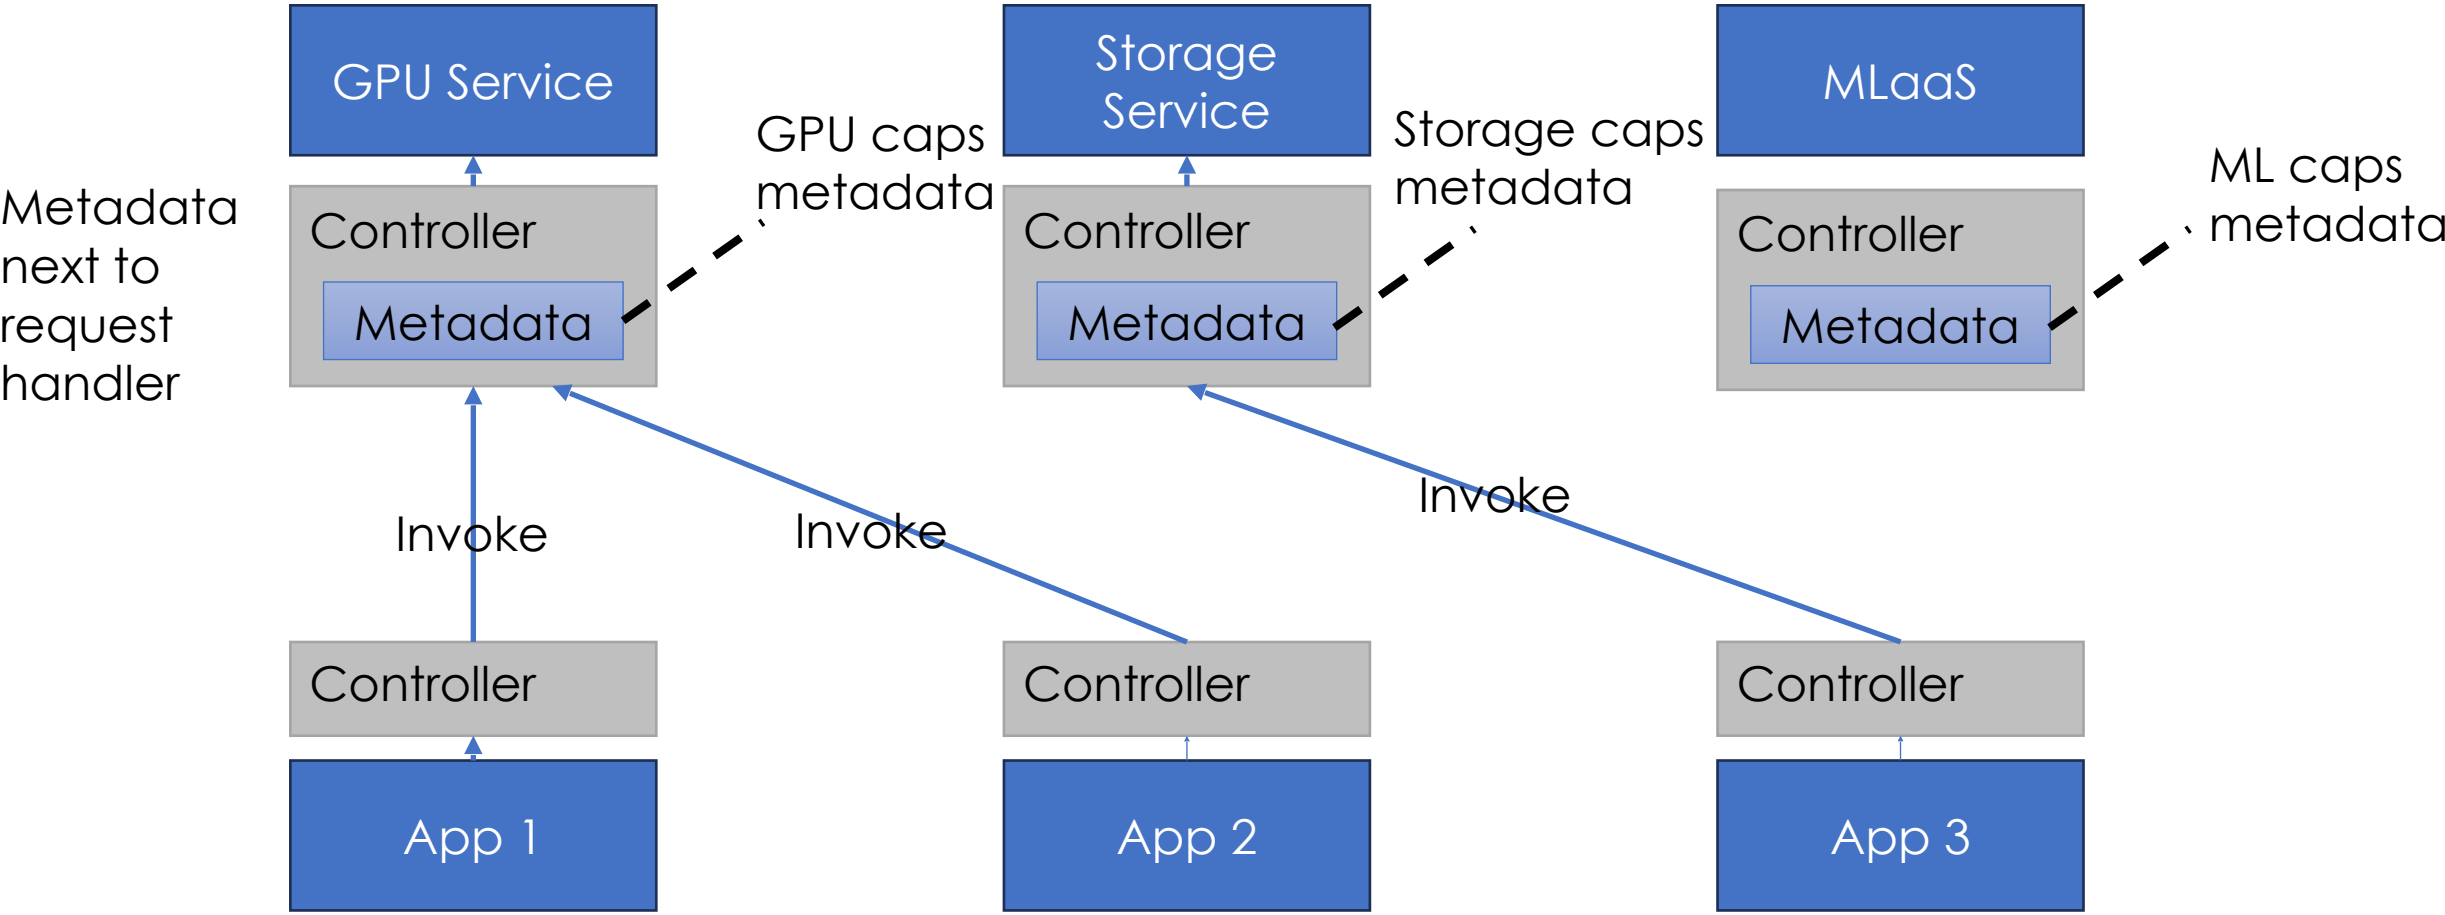
# Owner-based Metadata Sharding



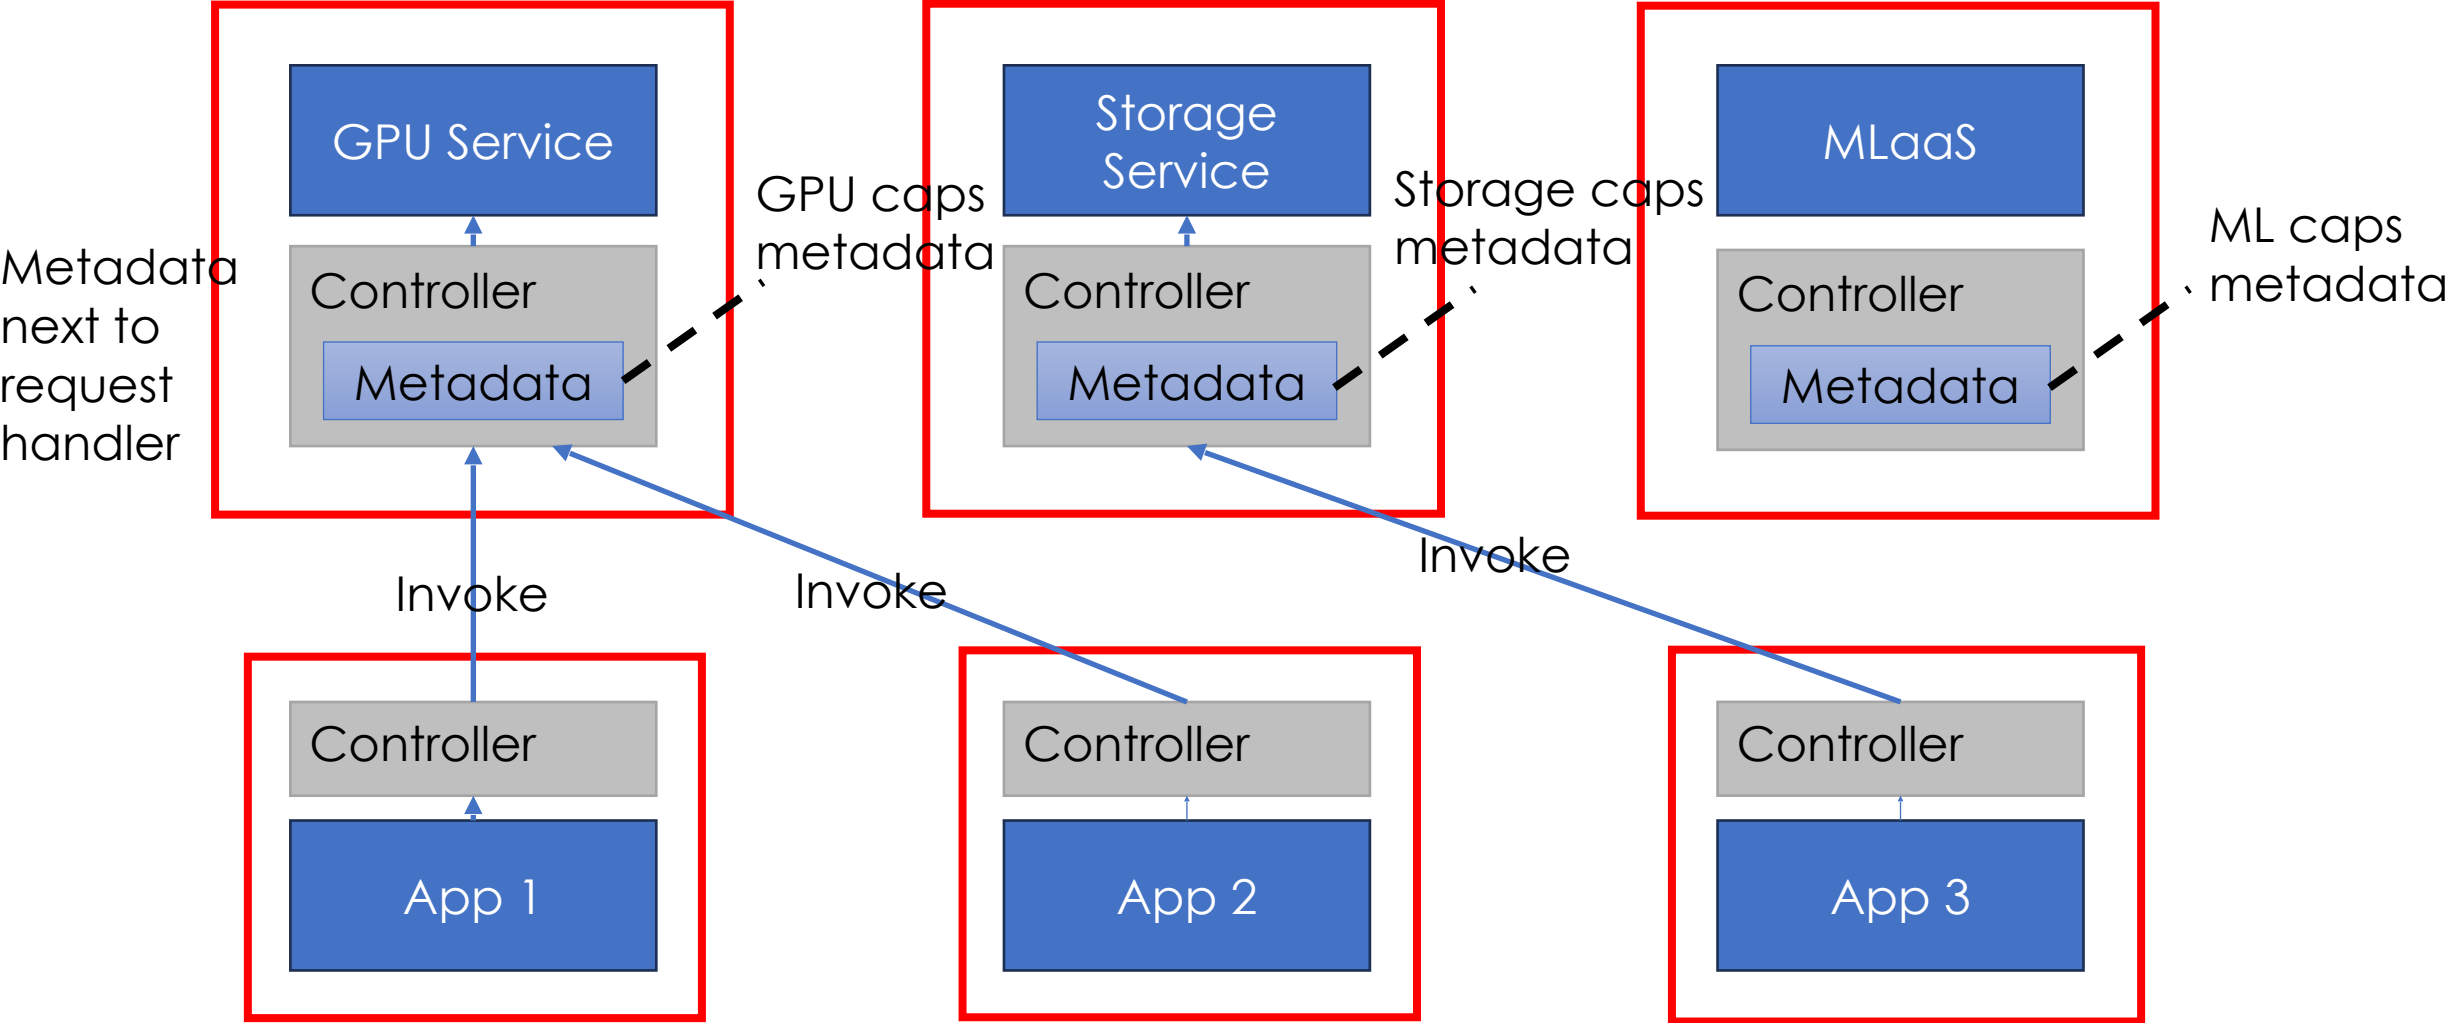
# Owner-based Metadata Sharding



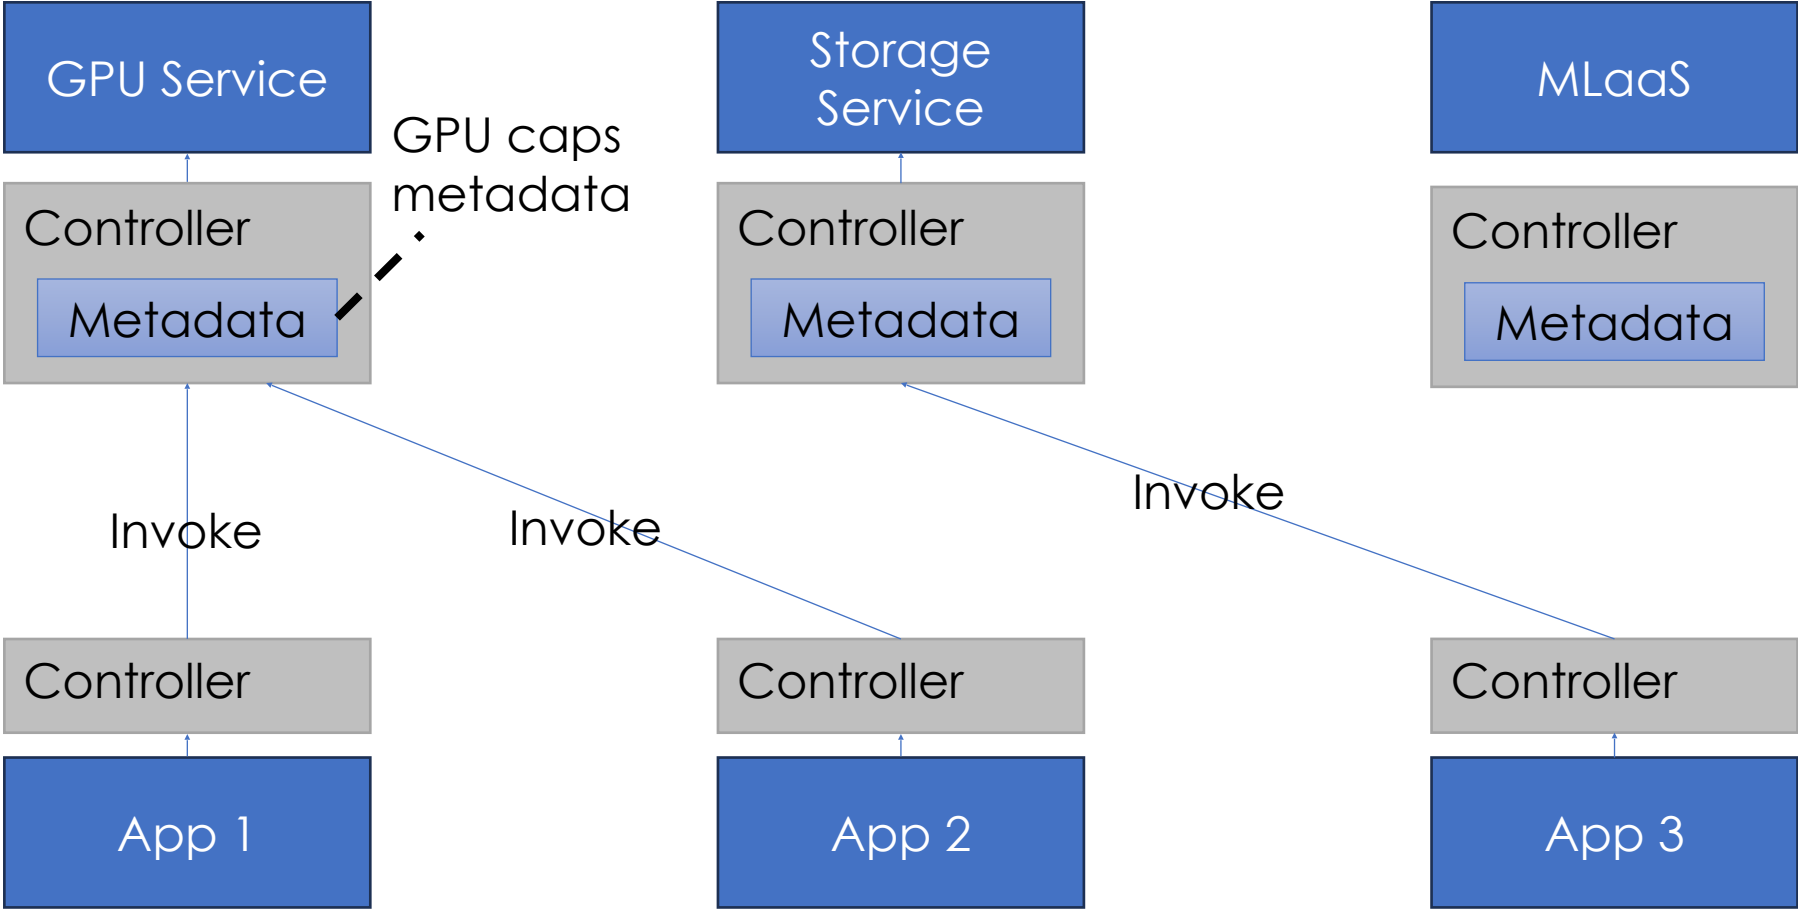
# Owner-based Metadata Sharding



# Owner-based Metadata Sharding

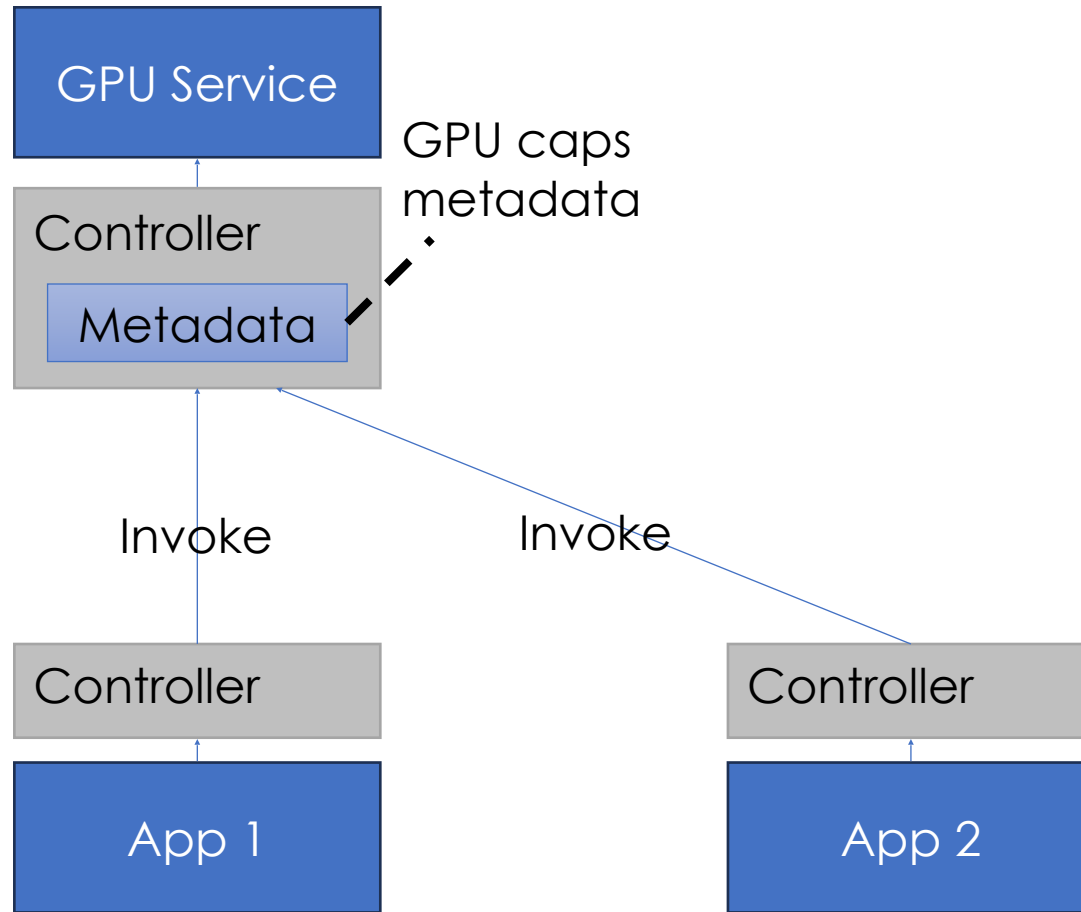


# Problem: Inefficient Revocation in Existing Capability Systems

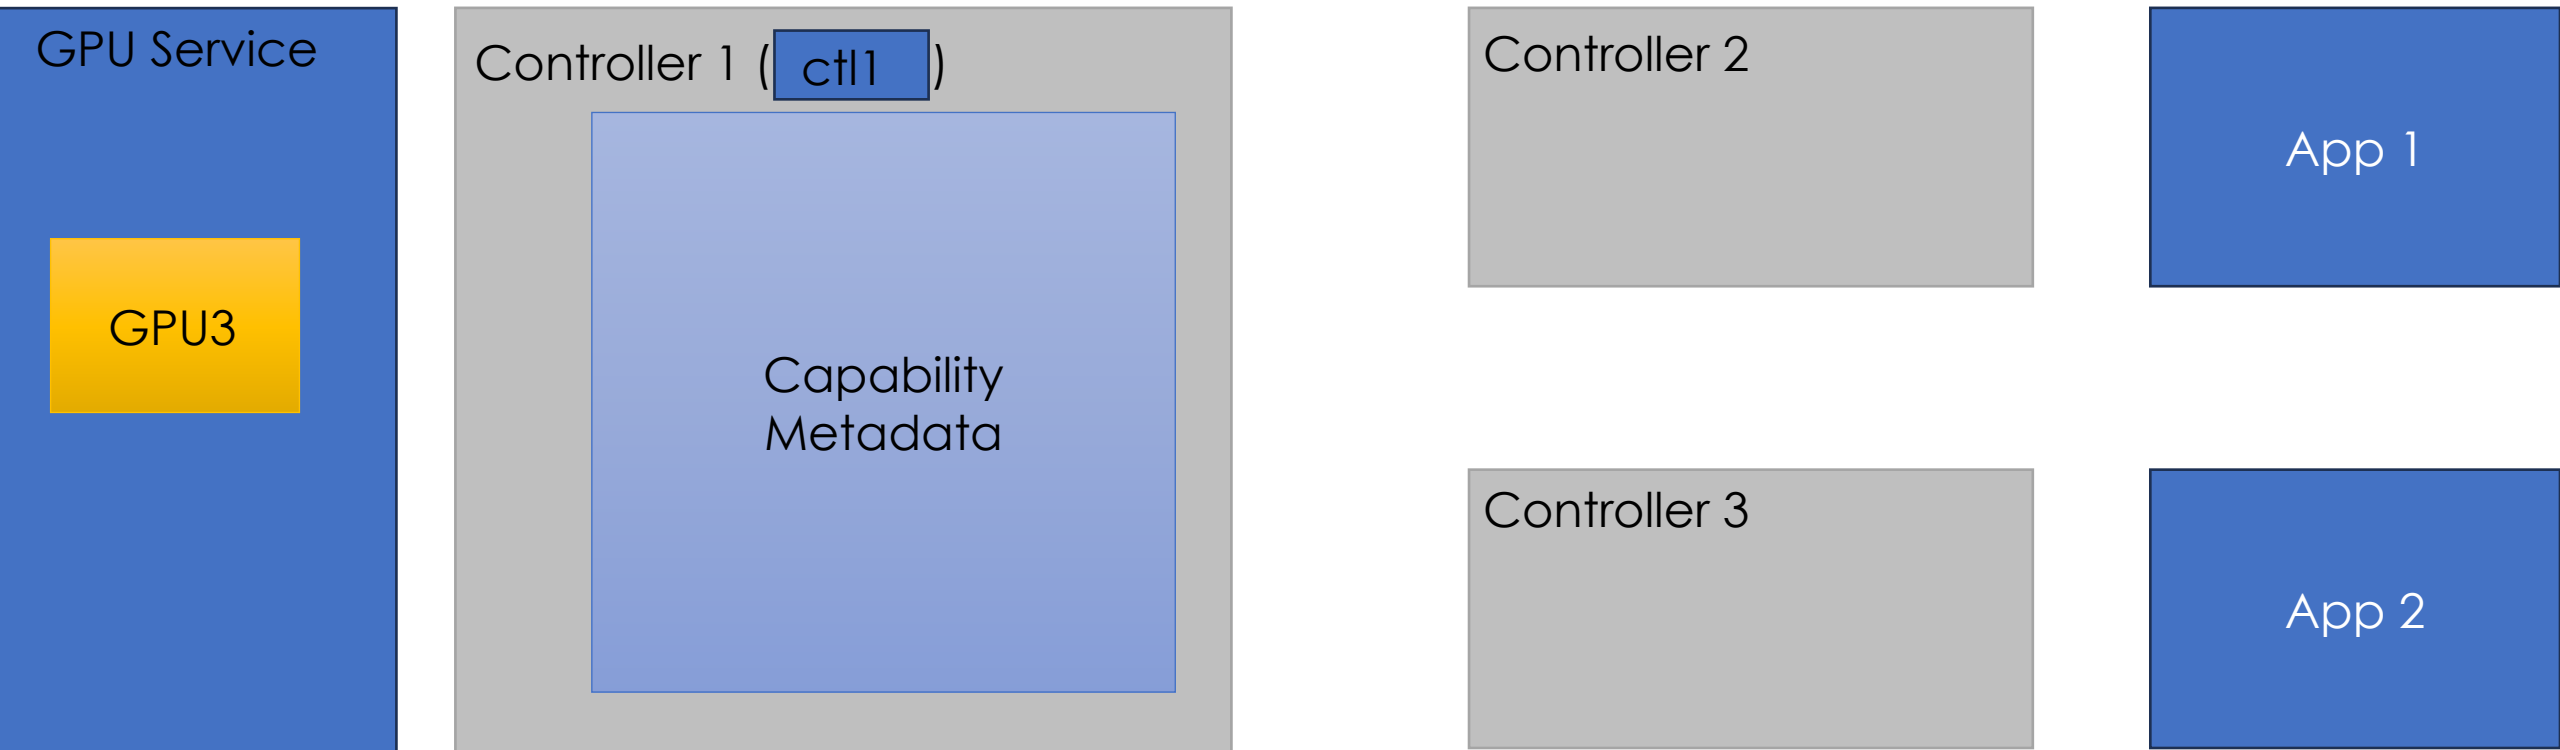




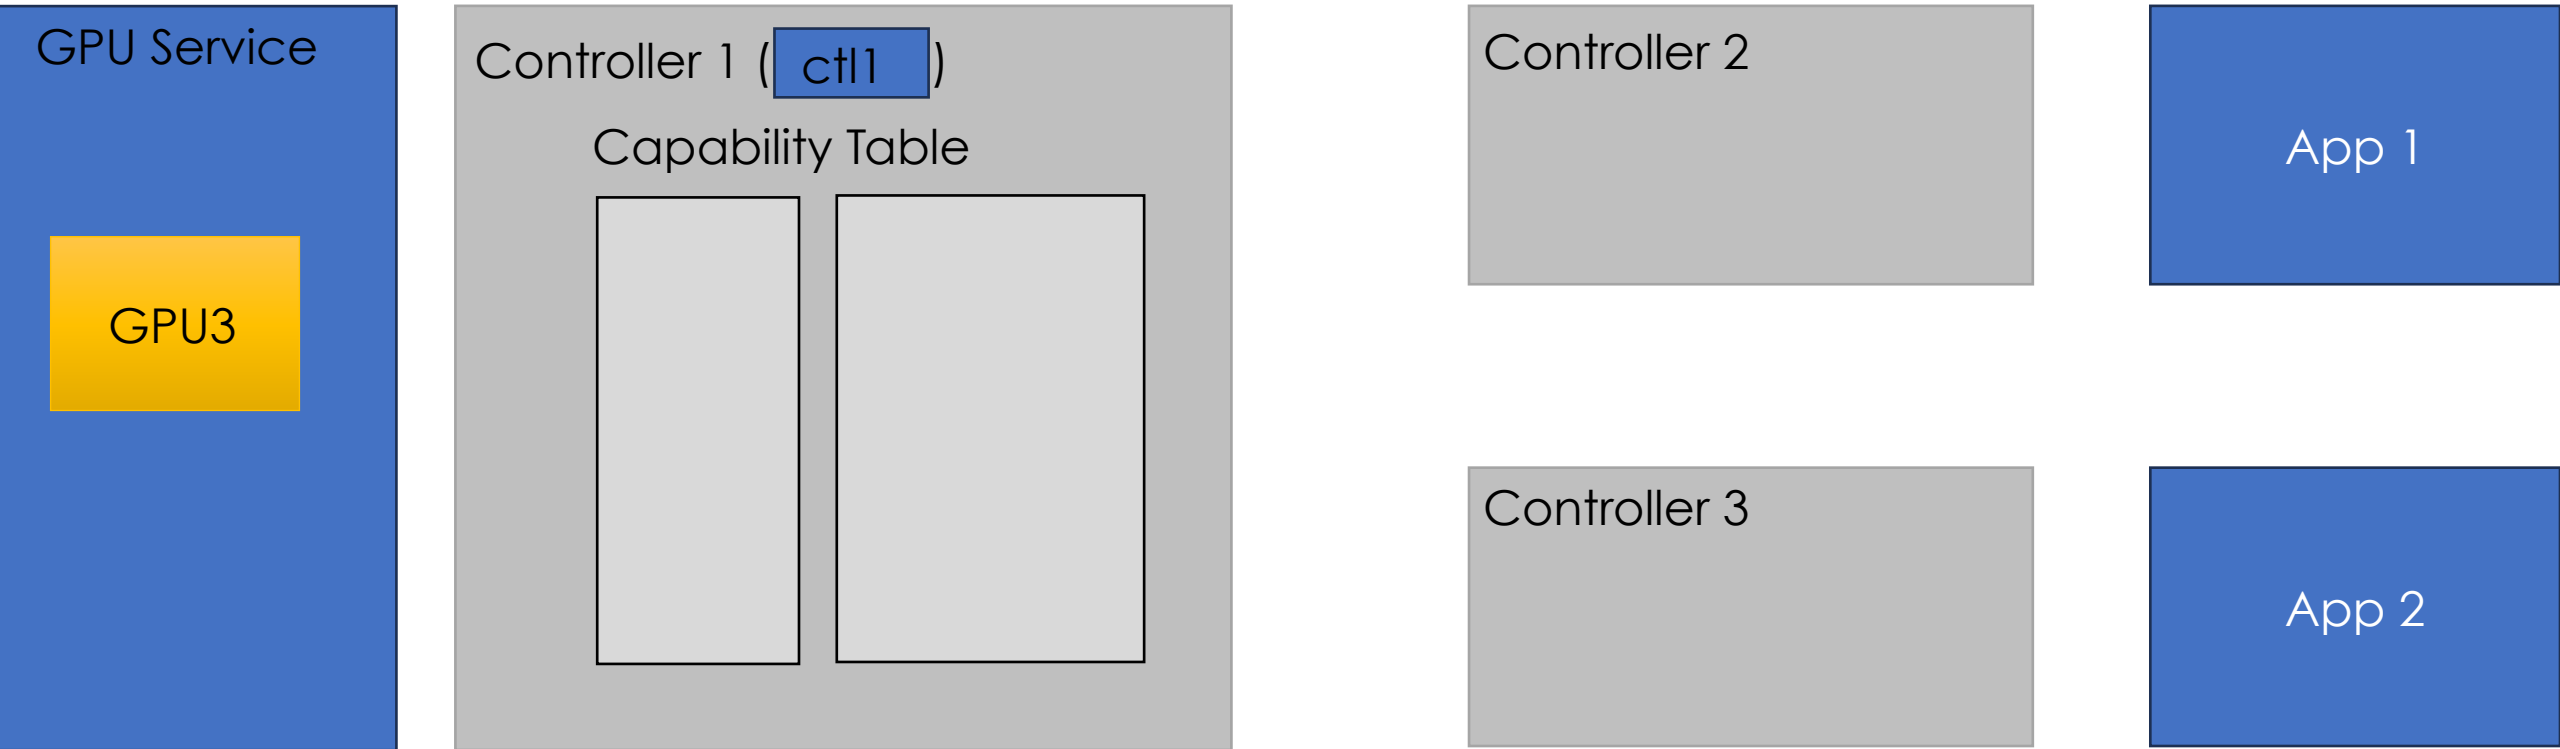
# Problem: Inefficient Revocation in Existing Capability Systems



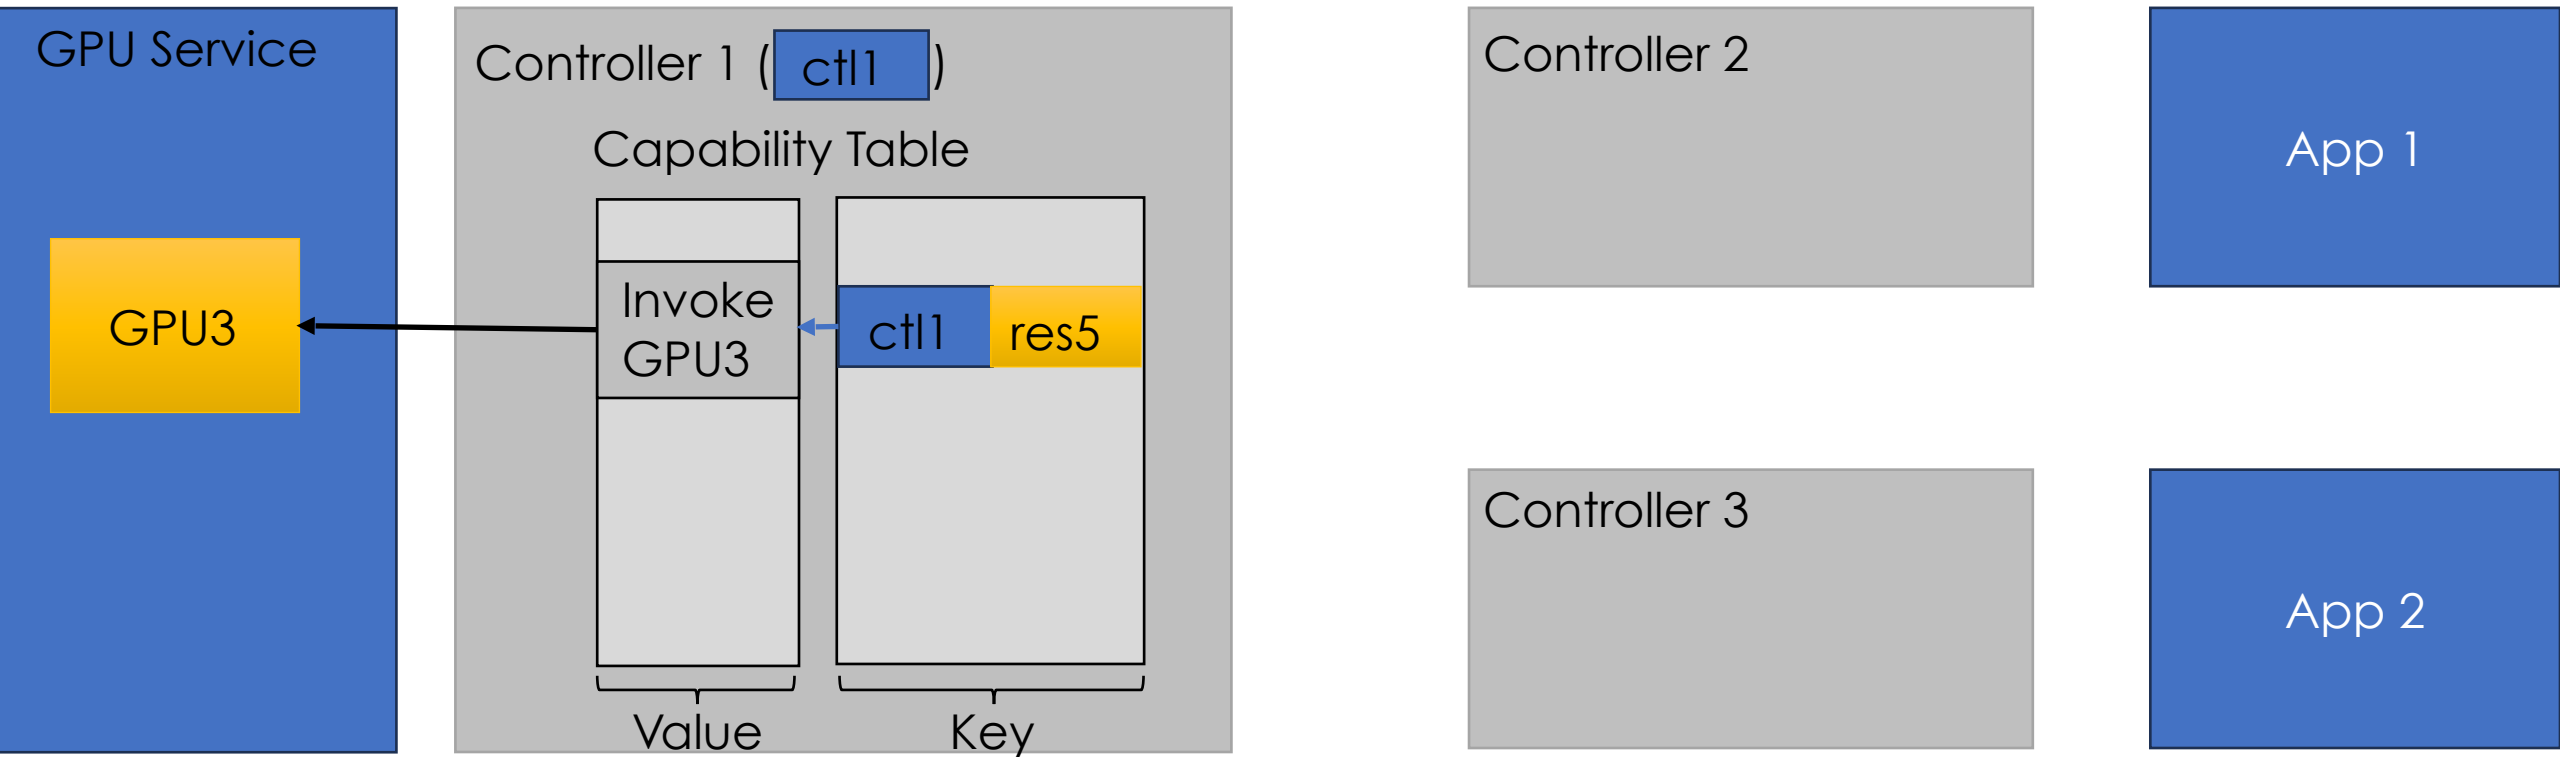
# Problem: Inefficient Revocation in Existing Capability Systems



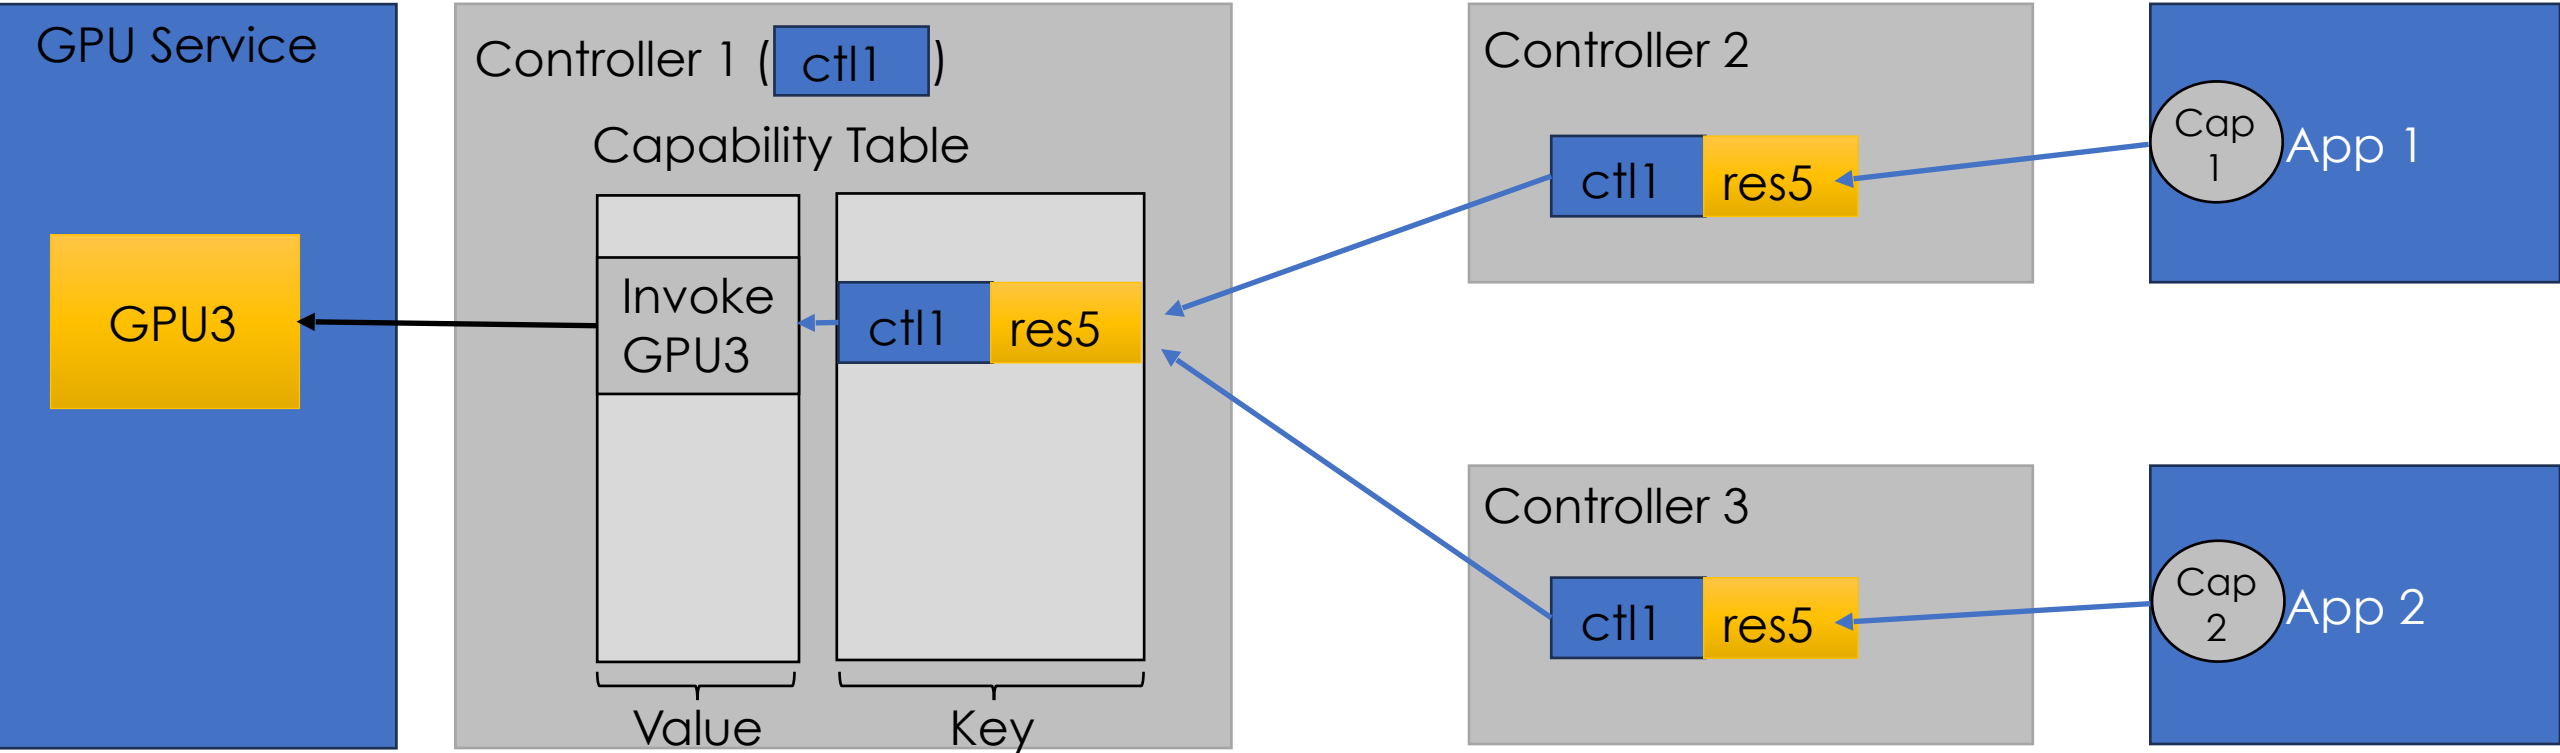
# Problem: Inefficient Revocation in Existing Capability Systems



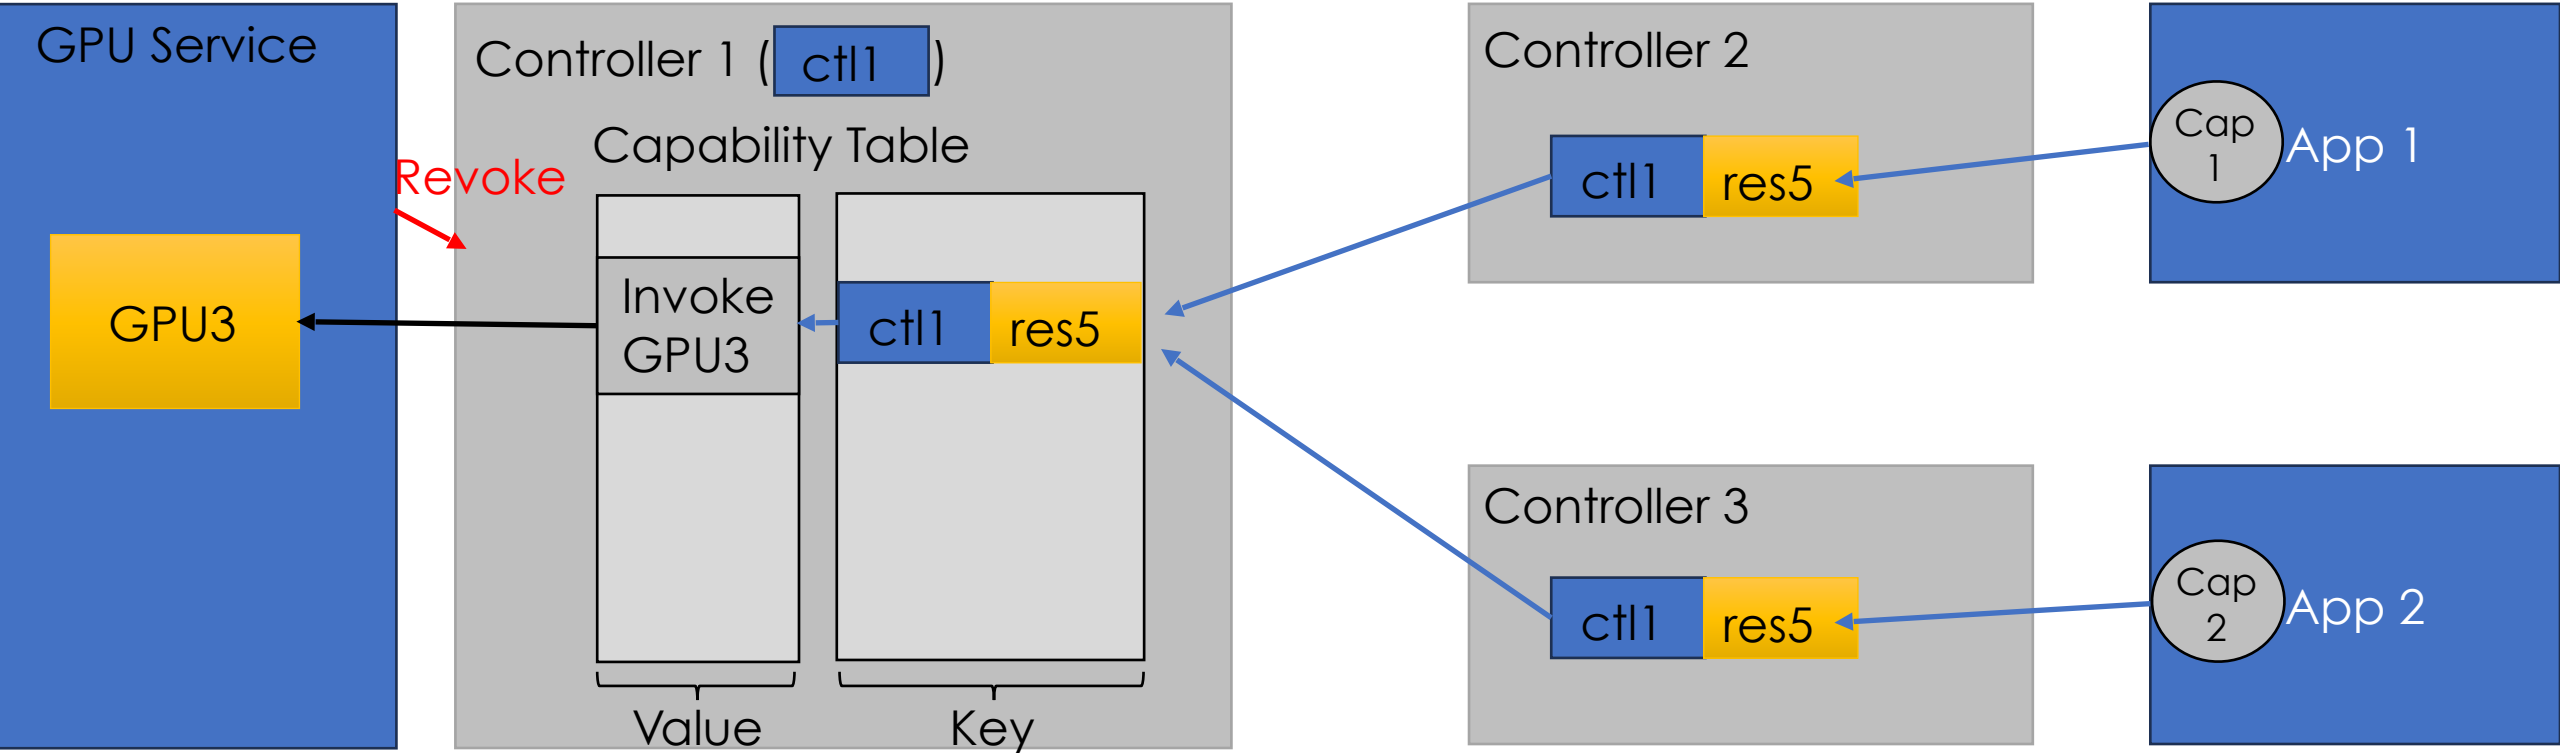
# Problem: Inefficient Revocation in Existing Capability Systems



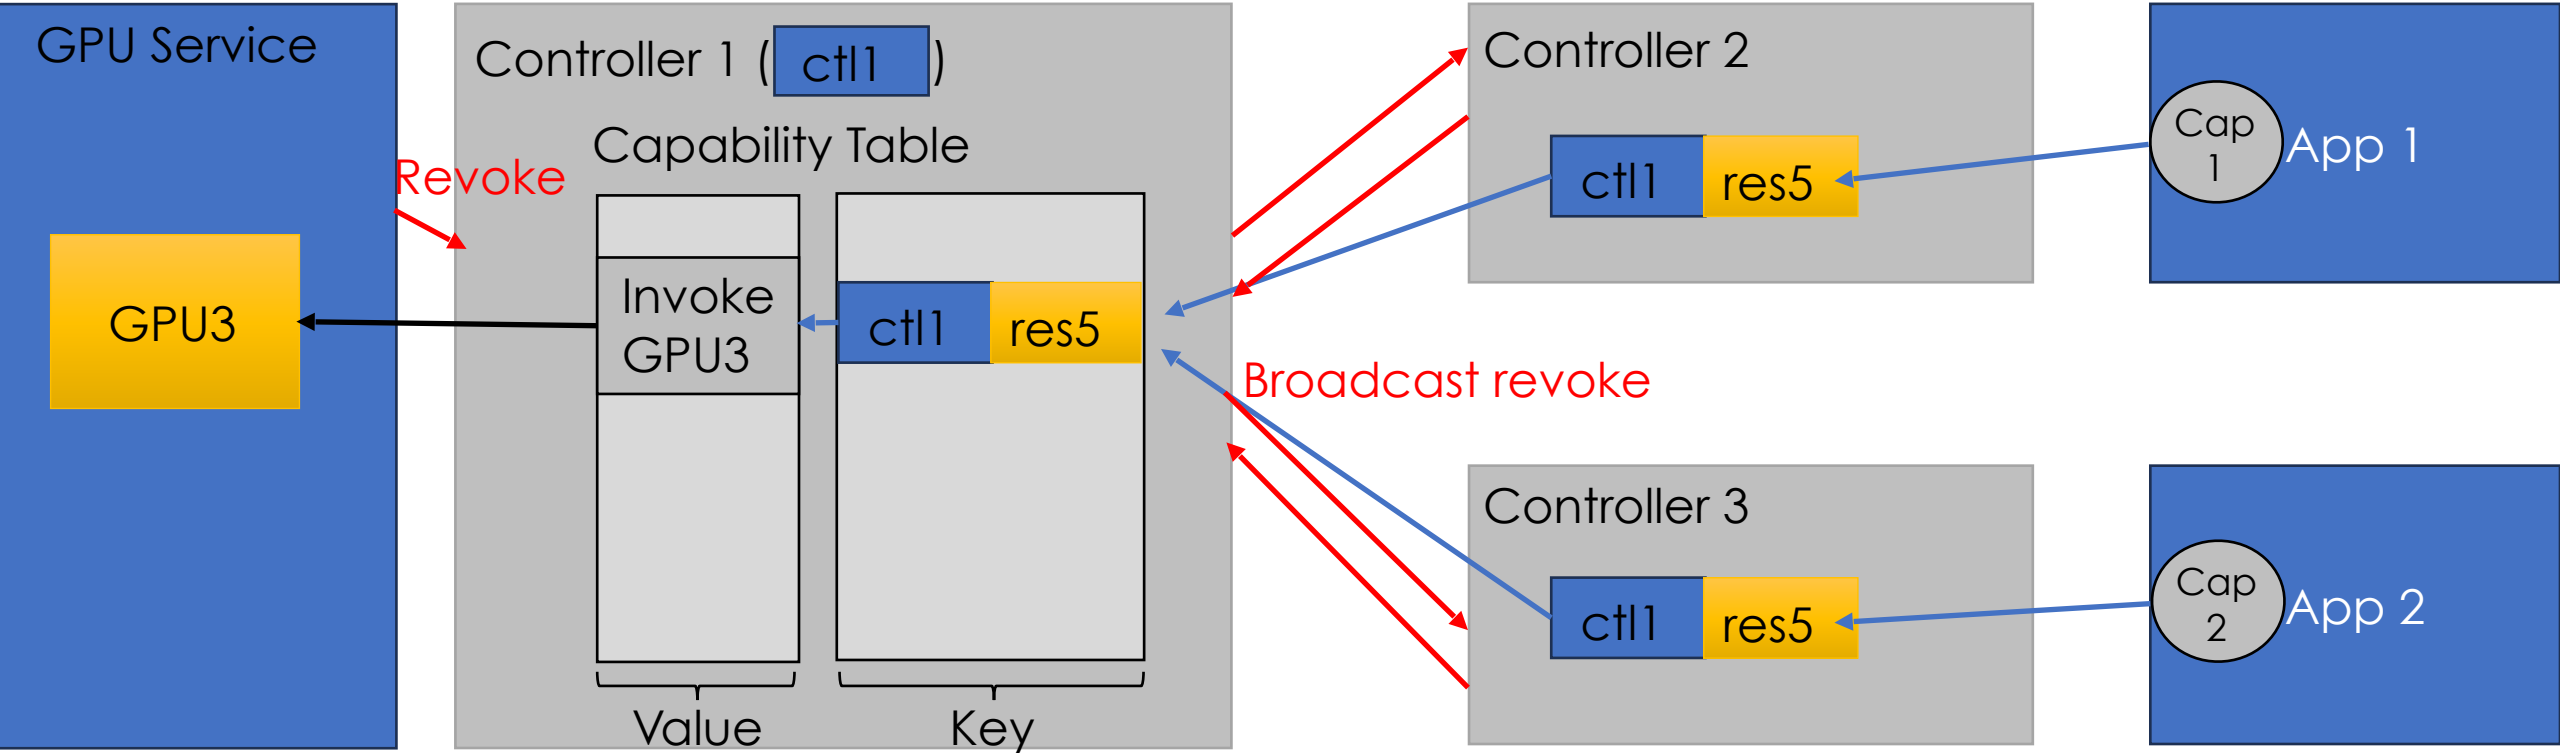
# Problem: Inefficient Revocation in Existing Capability Systems



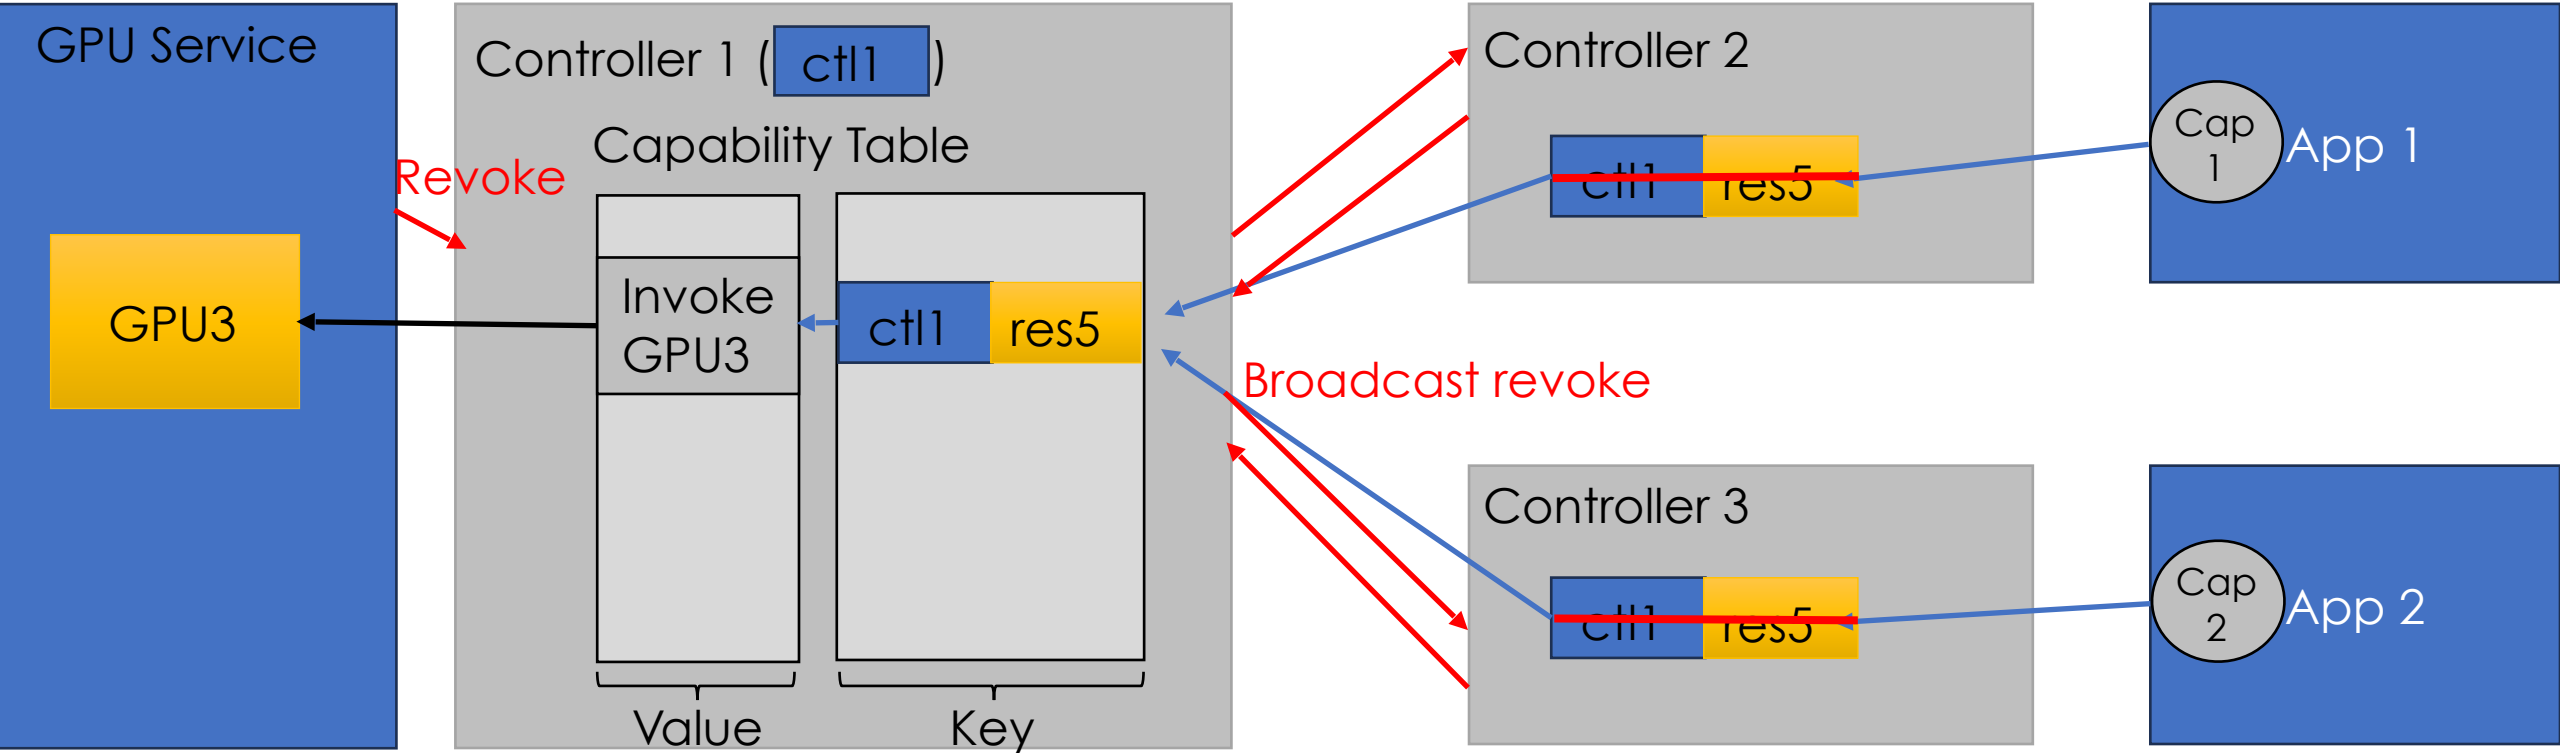
# Problem: Inefficient Revocation in Existing Capability Systems



# Problem: Inefficient Revocation in Existing Capability Systems

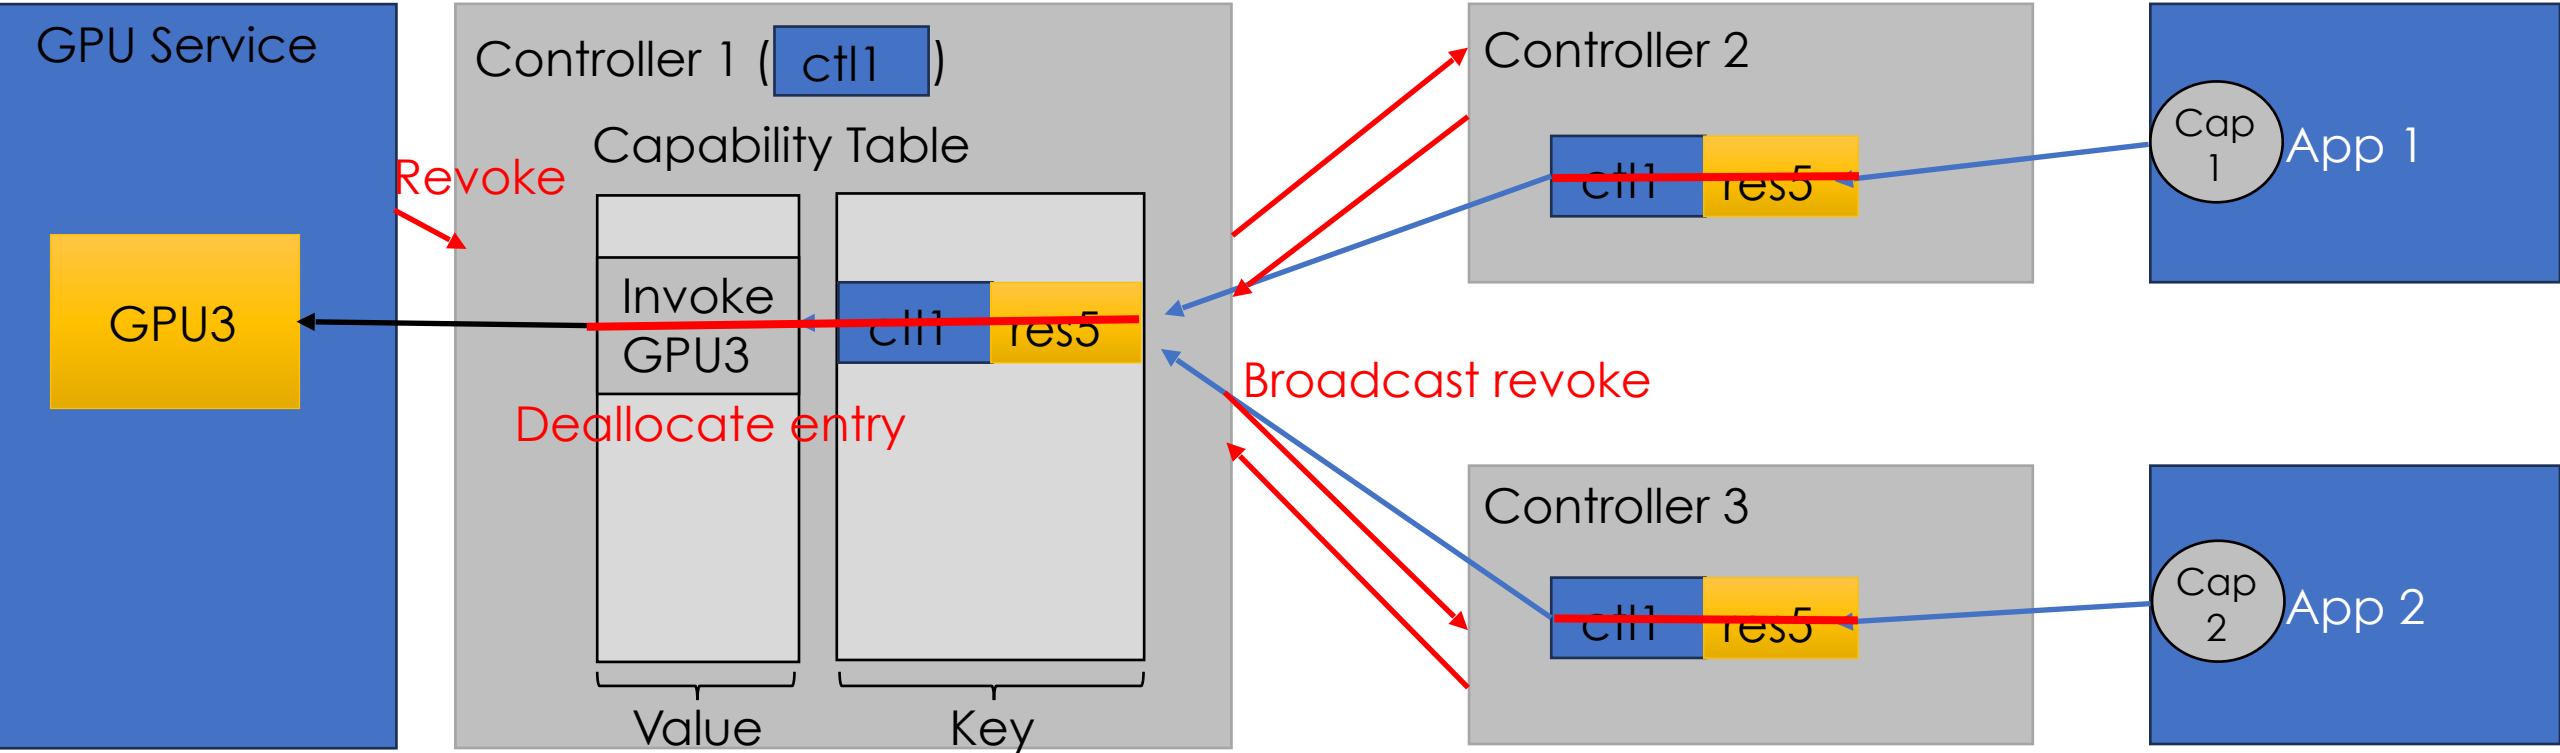


# Problem: Inefficient Revocation in Existing Capability Systems

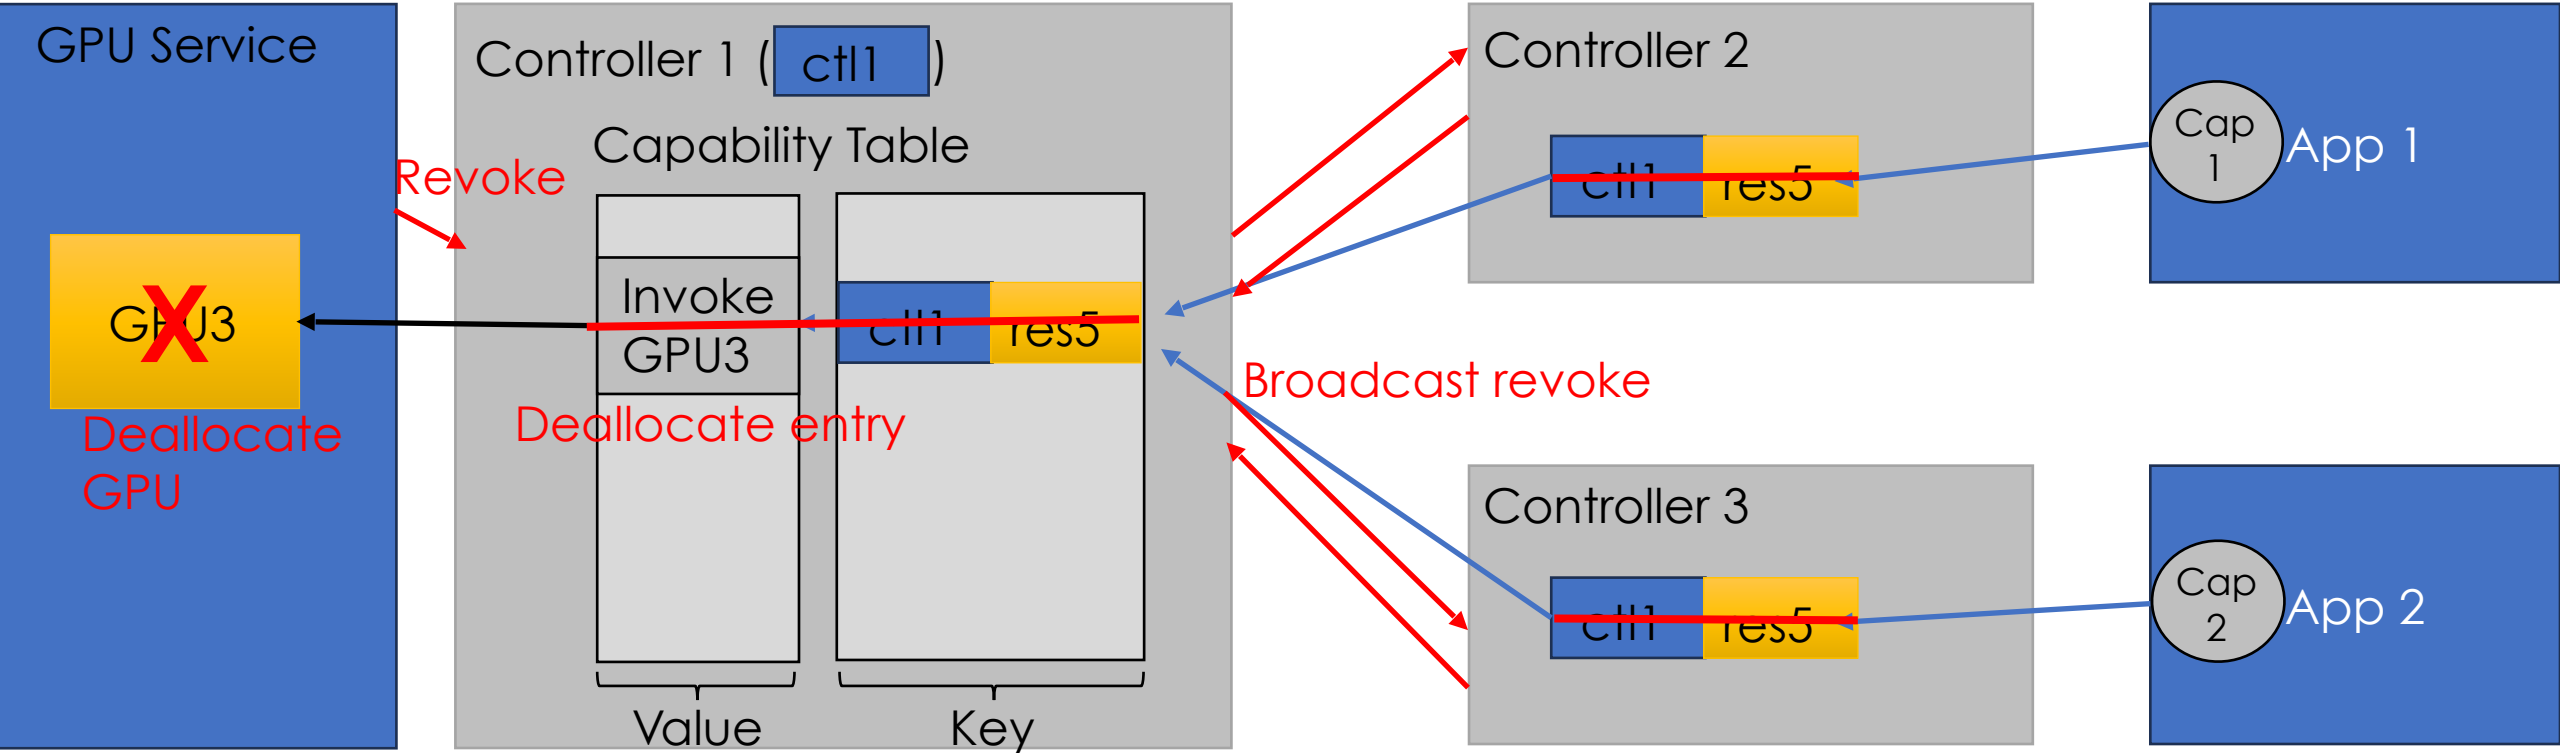





# Problem: Inefficient Revocation in Existing Capability Systems



# Problem: Inefficient Revocation in Existing Capability Systems




# Problem: Inefficient Revocation in Existing Capability Systems




Revocation requires  $O(N)$  network messages (and CPU)


# Versioned Capabilities




# Versioned Capabilities



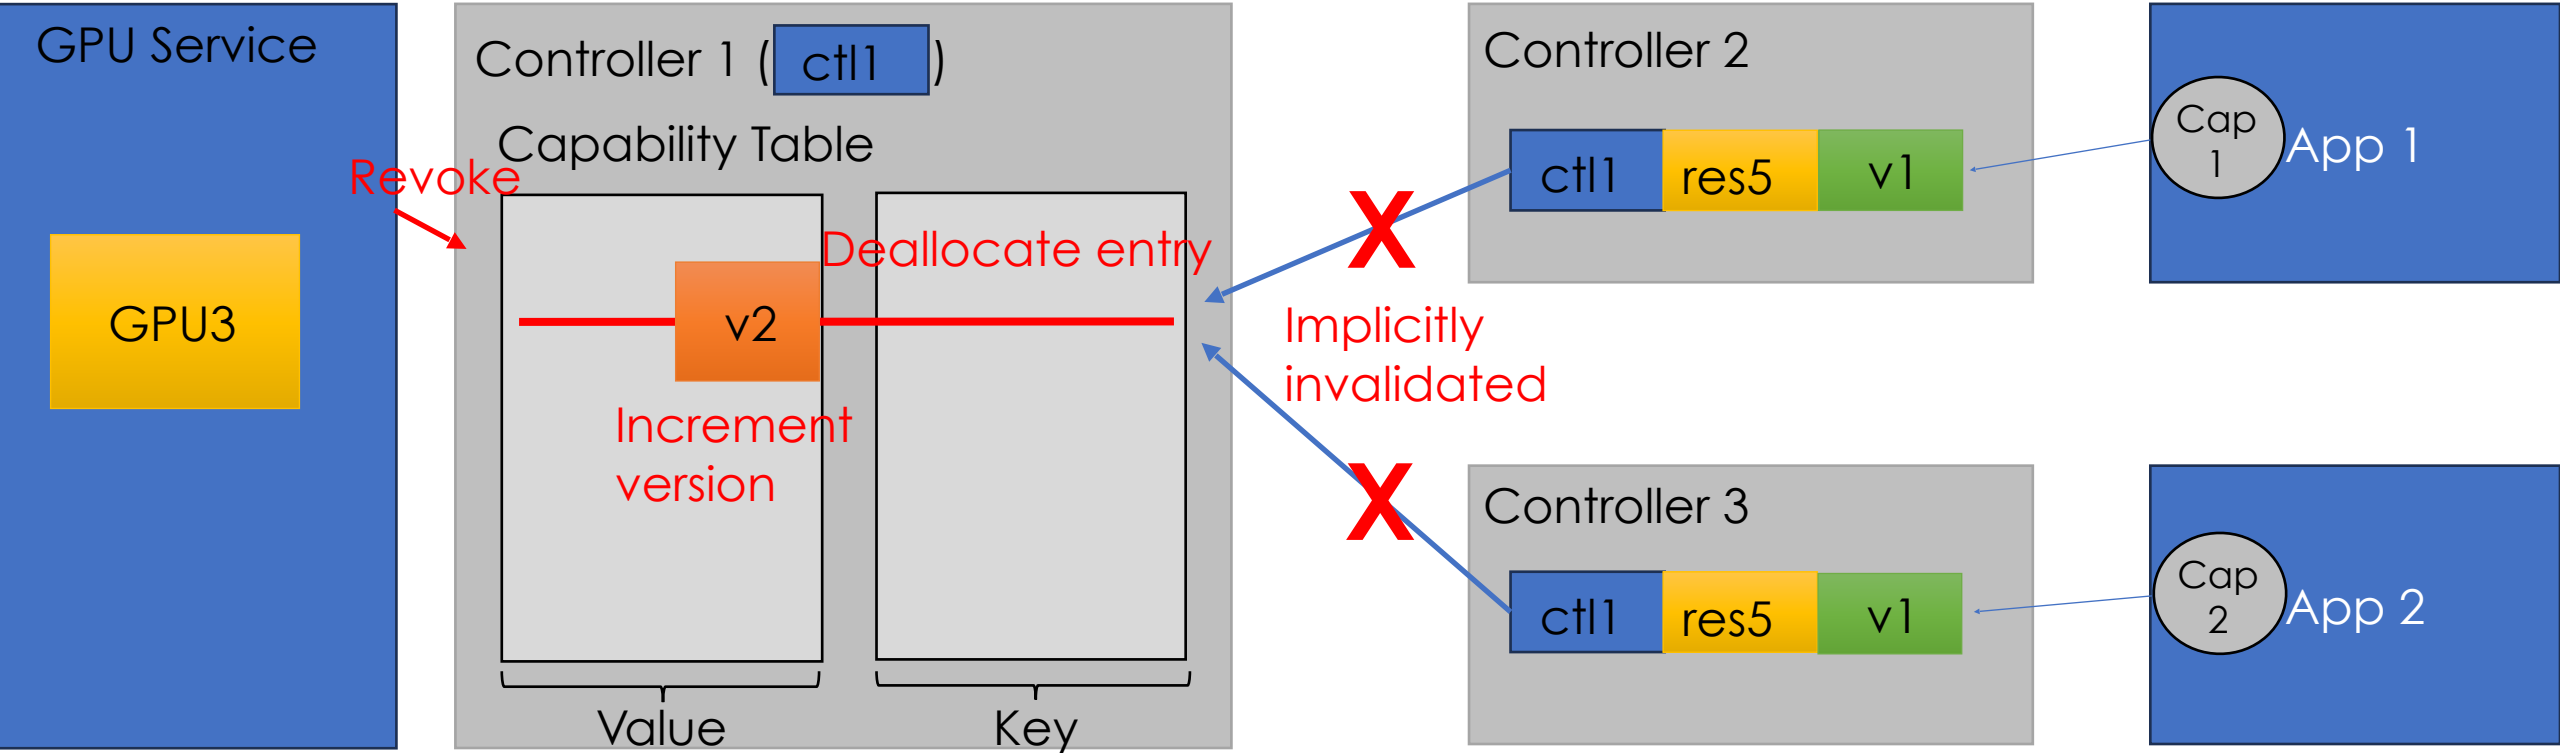
# Versioned Capabilities



# Versioned Capabilities

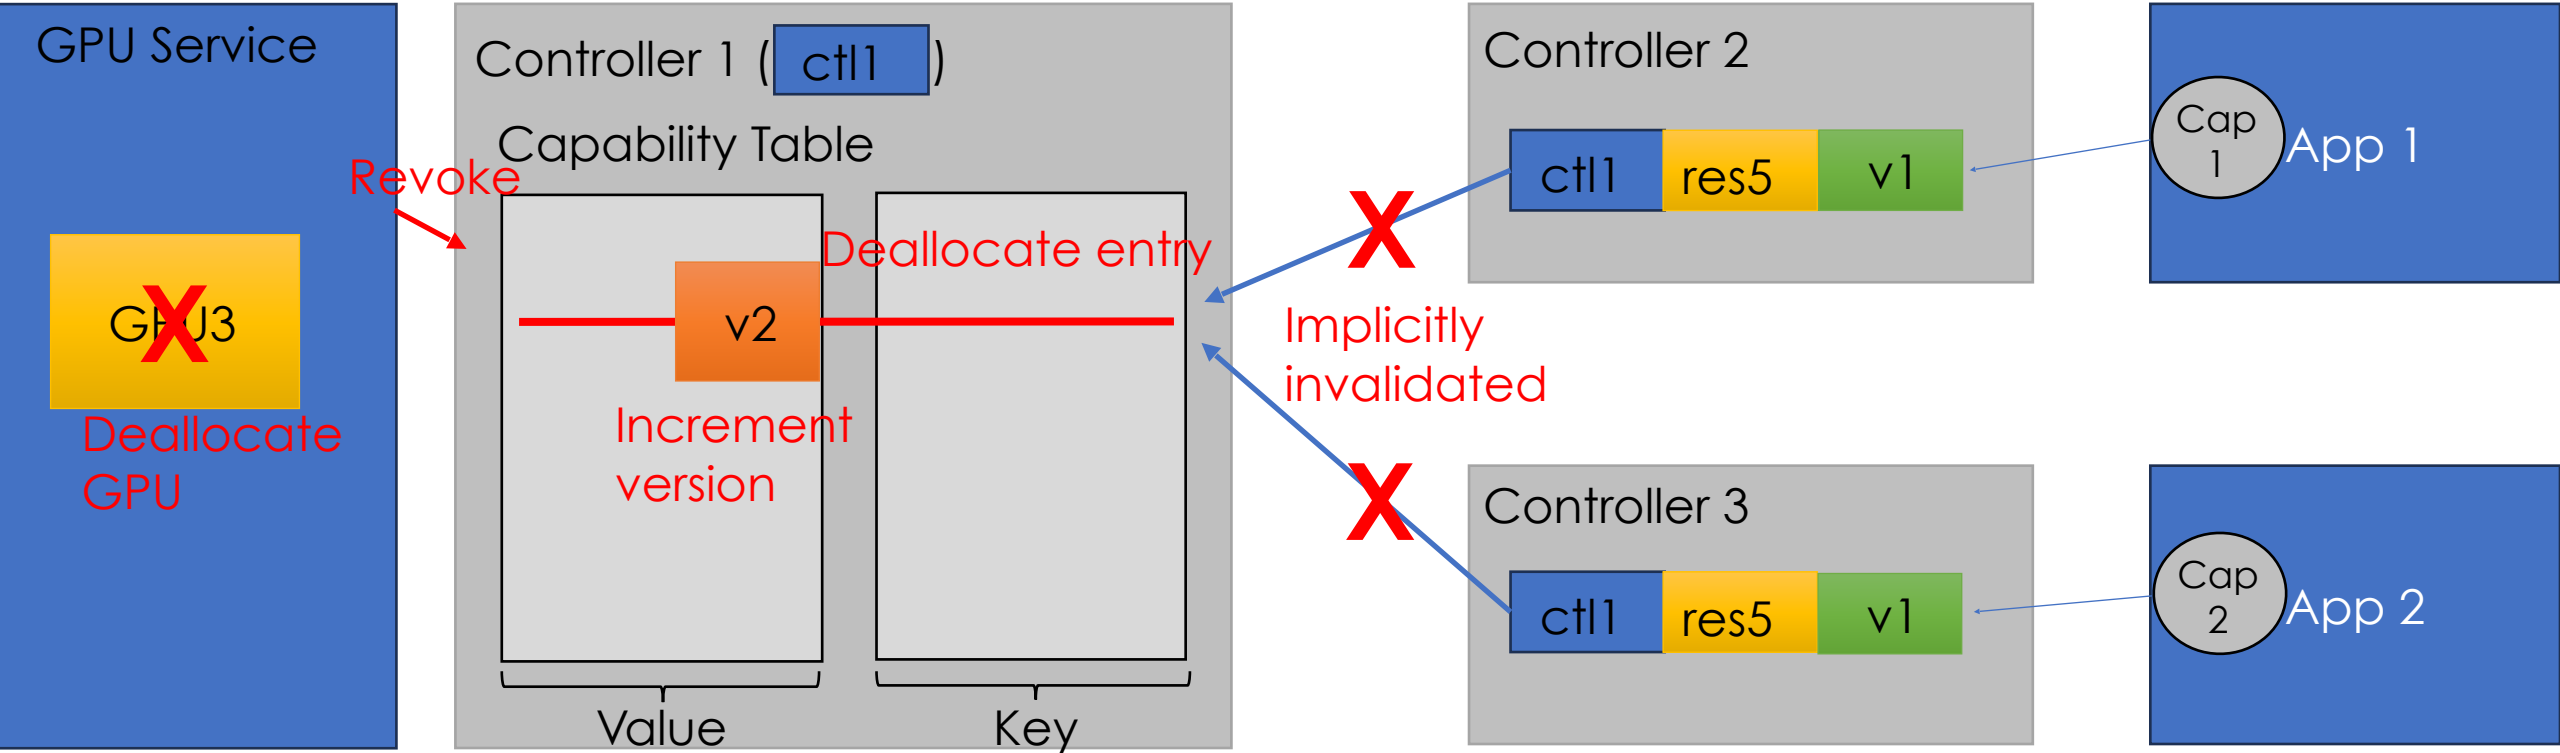


# Versioned Capabilities






# Versioned Capabilities

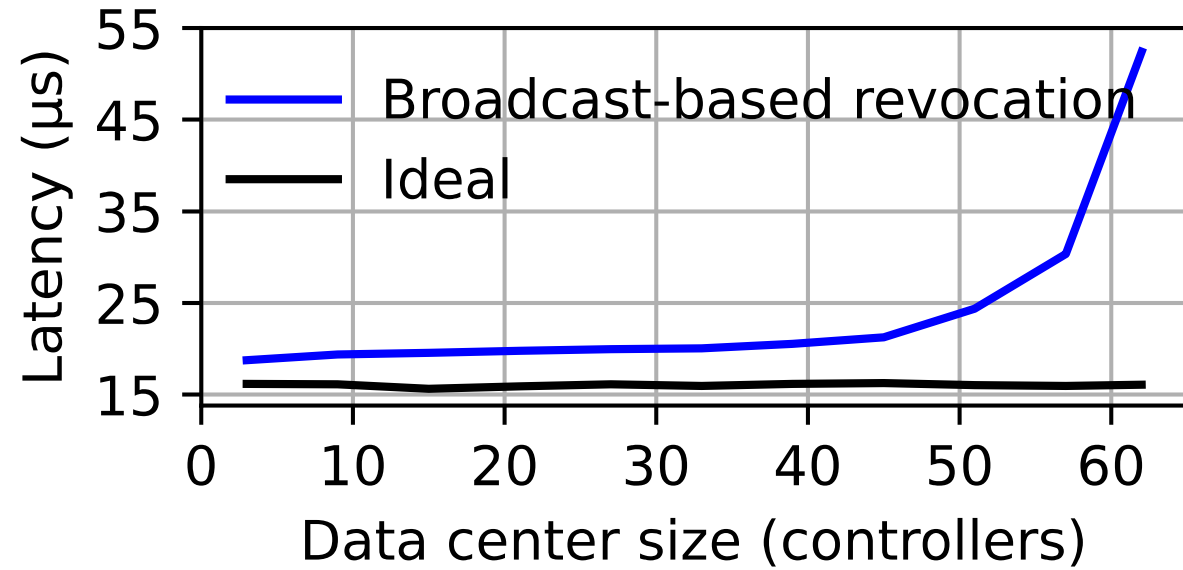


# Versioned Capabilities




Revocation requires zero network transmission (or CPU)


# Evaluation





# Evaluation





# Evaluation



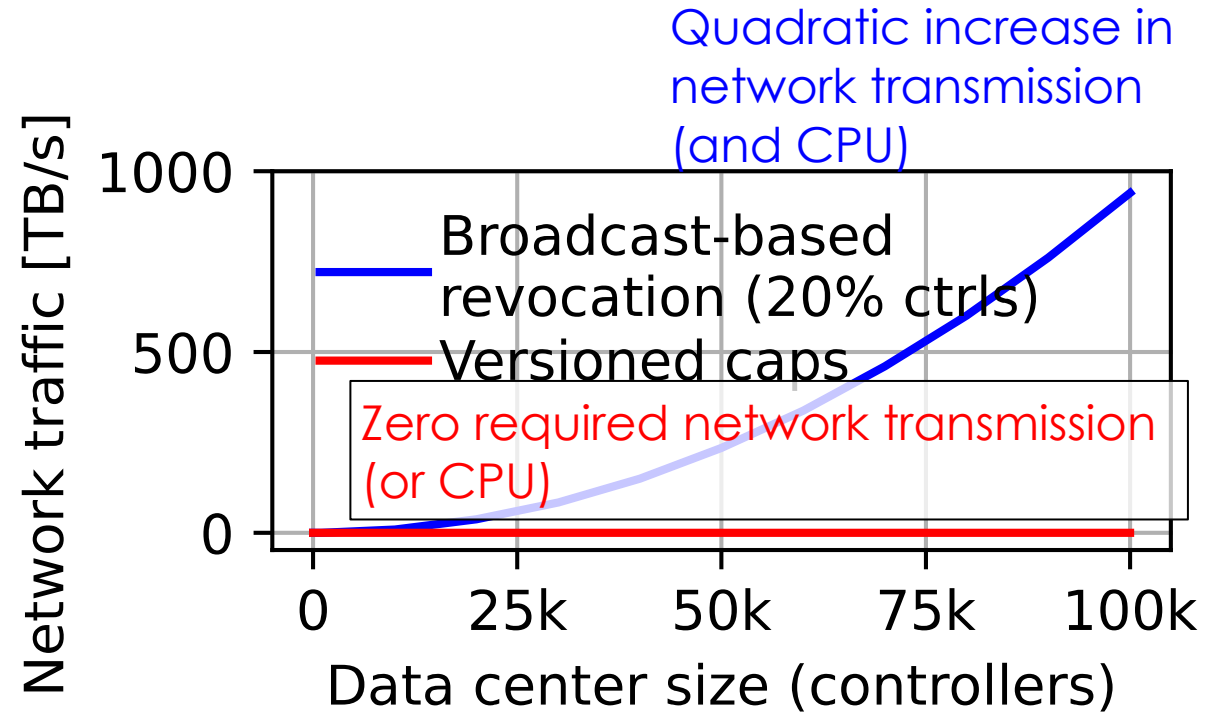

# Evaluation



# Evaluation

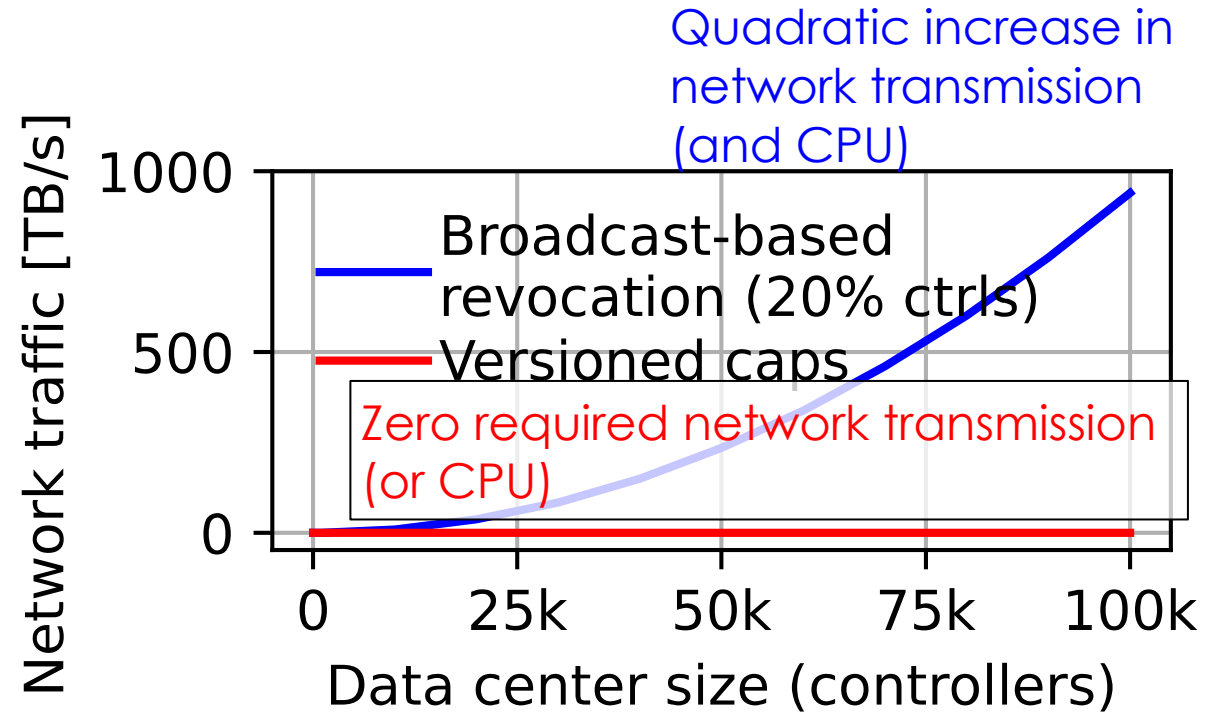



# Evaluation





# Evaluation



- Constant-cost, low  $\mu\text{s}$ -scale cap operations (eval in paper)

# Conclusion

# Conclusion

- Three key innovations
  - Owner-based metadata sharding, versioned capabilities and guards (detail in paper)

# Conclusion

- Three key innovations
  - Owner-based metadata sharding, versioned capabilities and guards (detail in paper)
- To achieve all three cloud-scale challenges

# Conclusion

- Three key innovations
  - Owner-based metadata sharding, versioned capabilities and guards (detail in paper)
- To achieve all three cloud-scale challenges
  - Failure isolation




# Conclusion

- Three key innovations
  - Owner-based metadata sharding, versioned capabilities and guards (detail in paper)
- To achieve all three cloud-scale challenges
  - Failure isolation
  - Constant-cost,  $\mu$ s-scale cap operations (eval in paper)




# Conclusion

- Three key innovations
  - Owner-based metadata sharding, versioned capabilities and guards (detail in paper)
- To achieve all three cloud-scale challenges
  - Failure isolation
  - Constant-cost,  $\mu$ s-scale cap operations (eval in paper)
  - Constant system overheads (network transmission, CPU usage, latency overheads) regardless of DC size



# Conclusion

- Three key innovations
  - Owner-based metadata sharding, versioned capabilities and guards (detail in paper)
- To achieve all three cloud-scale challenges
  - Failure isolation
  - Constant-cost,  $\mu$ s-scale cap operations (eval in paper)
  - Constant system overheads (network transmission, CPU usage, latency overheads) regardless of DC size



Otto White 🤗 🙌  
otto.white20@imperial.ac.uk