

Right-sizing Geo-distributed Datacenters for Latency and Availability

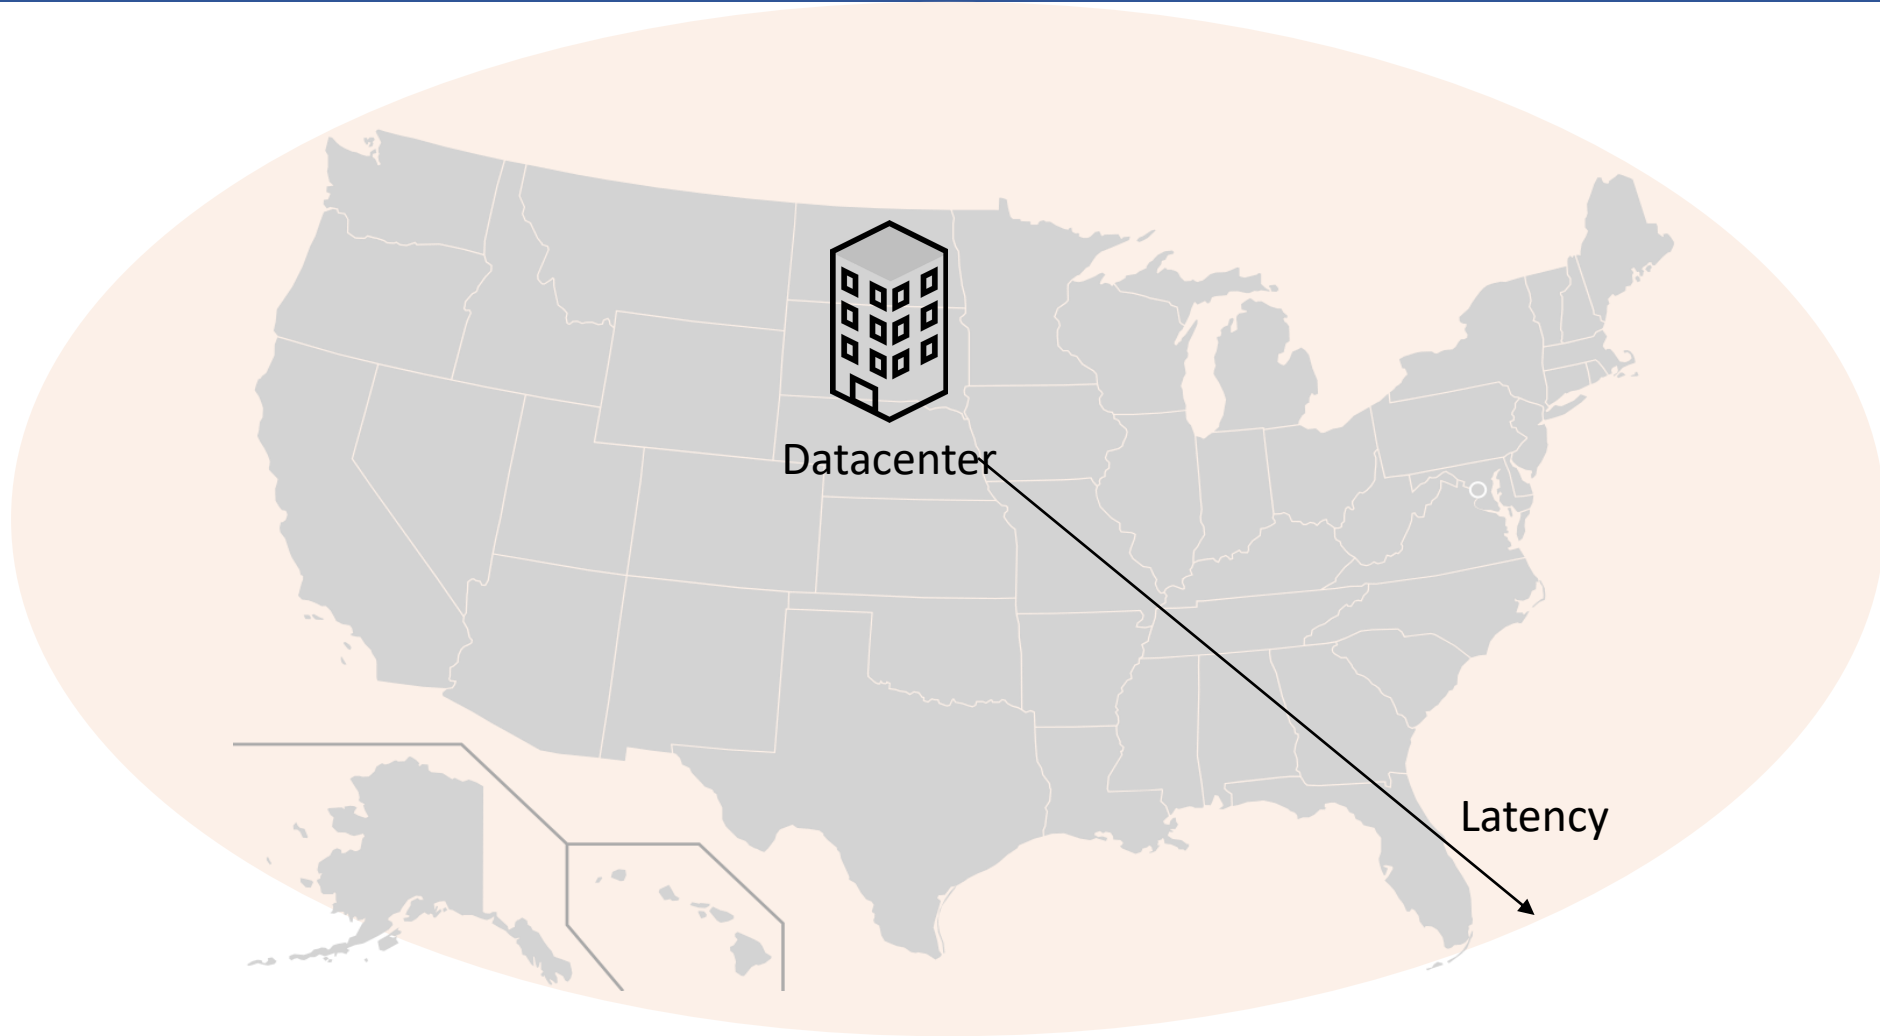
Iyswarya Narayanan[†]

Aman Kansal

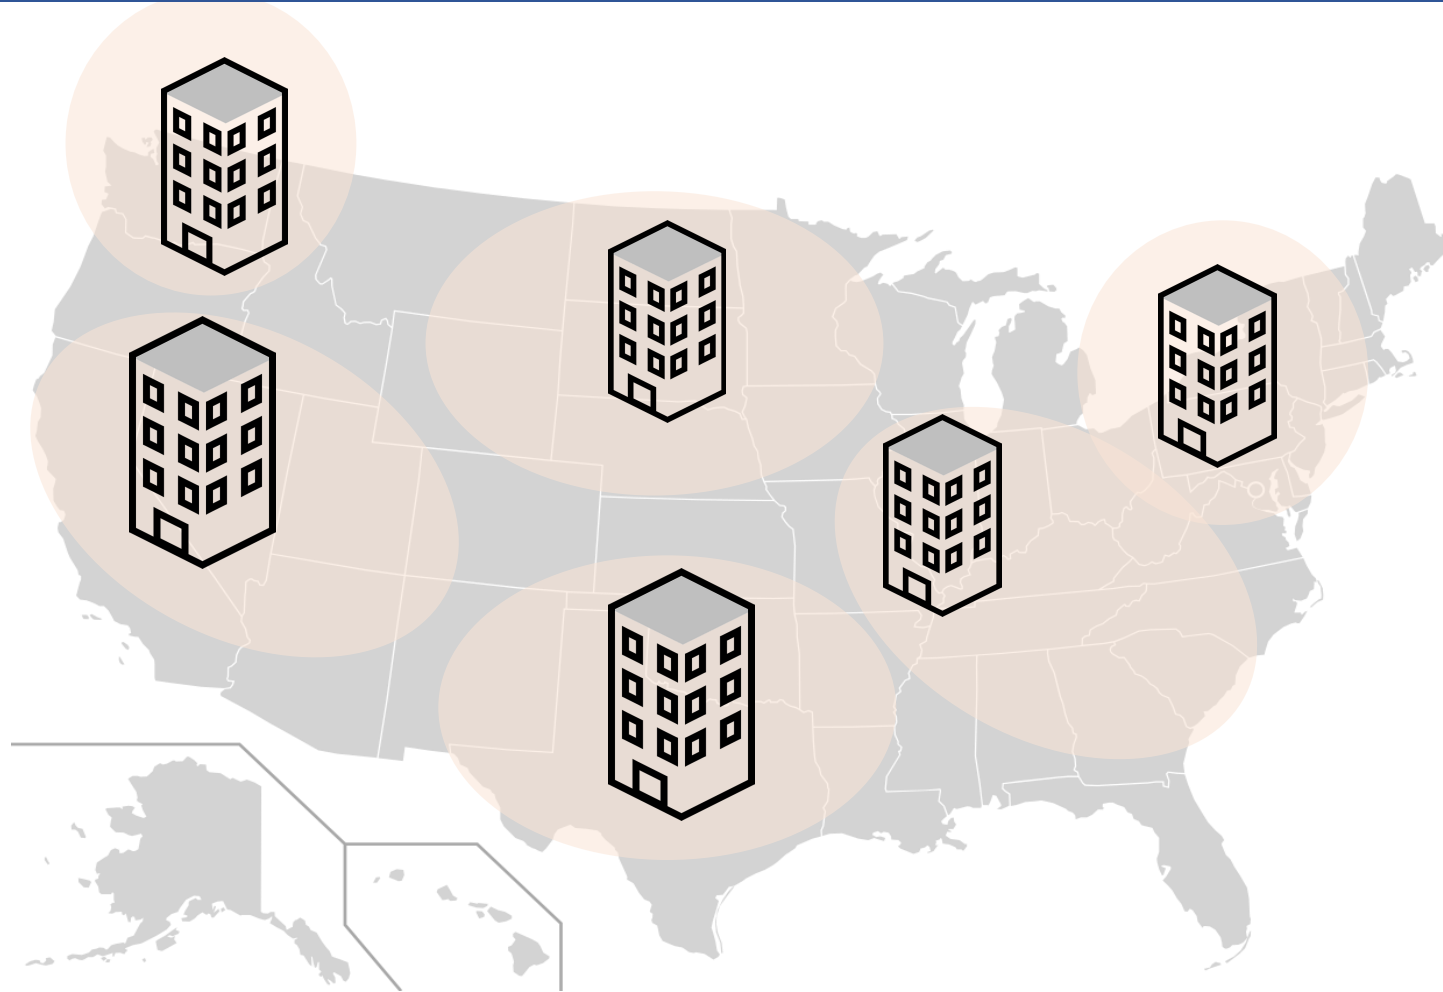
Anand Sivasubramaniam[†]



Latency is important for apps

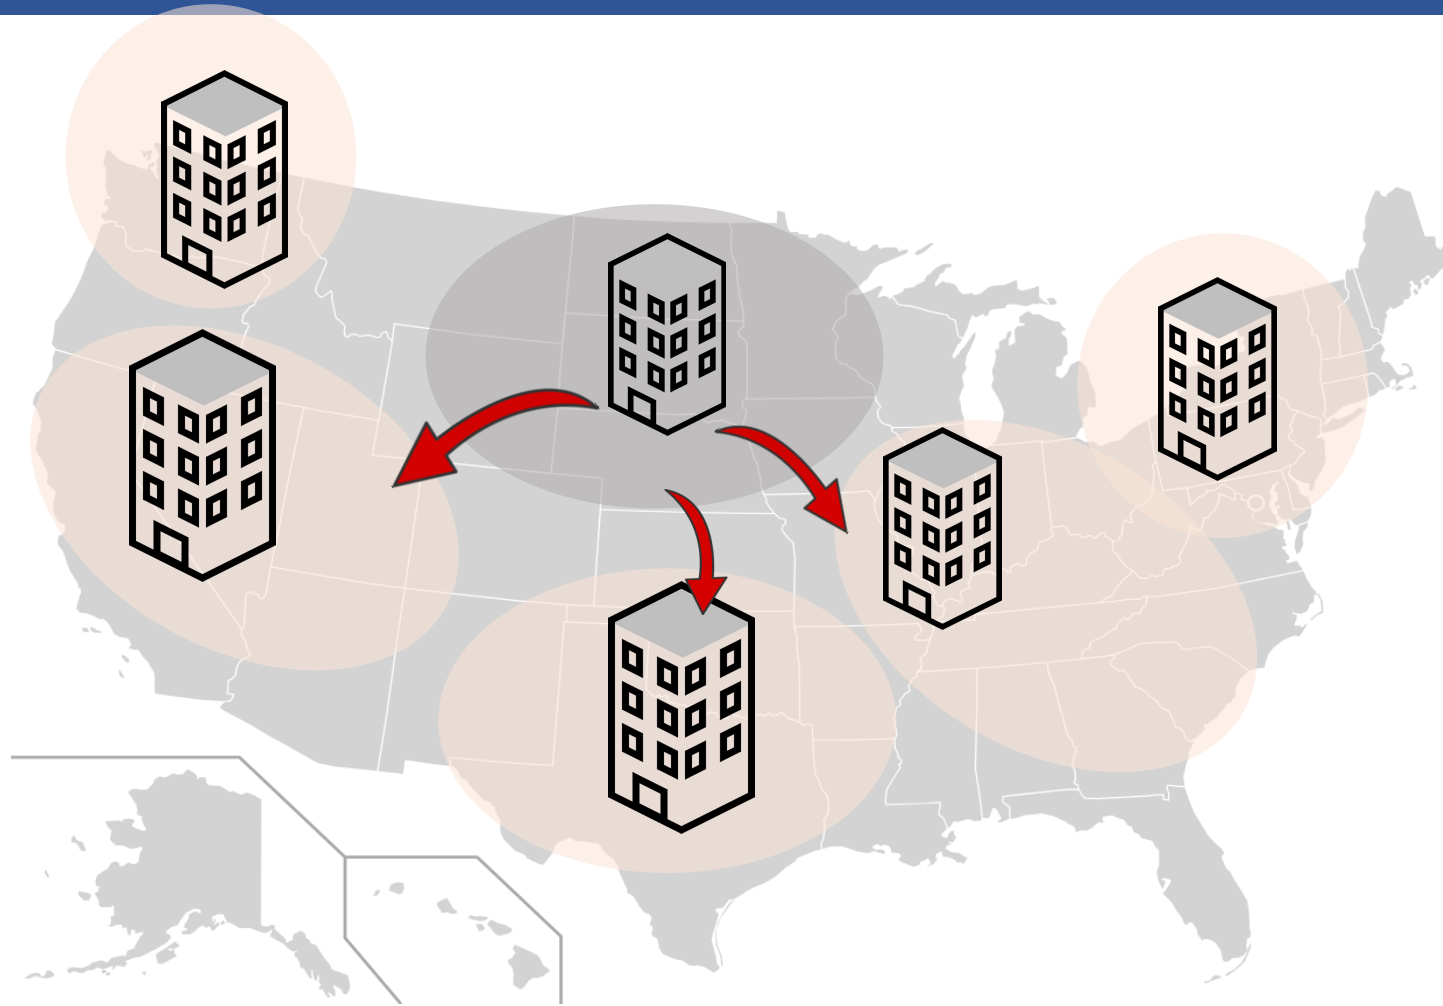


Latency is important for apps



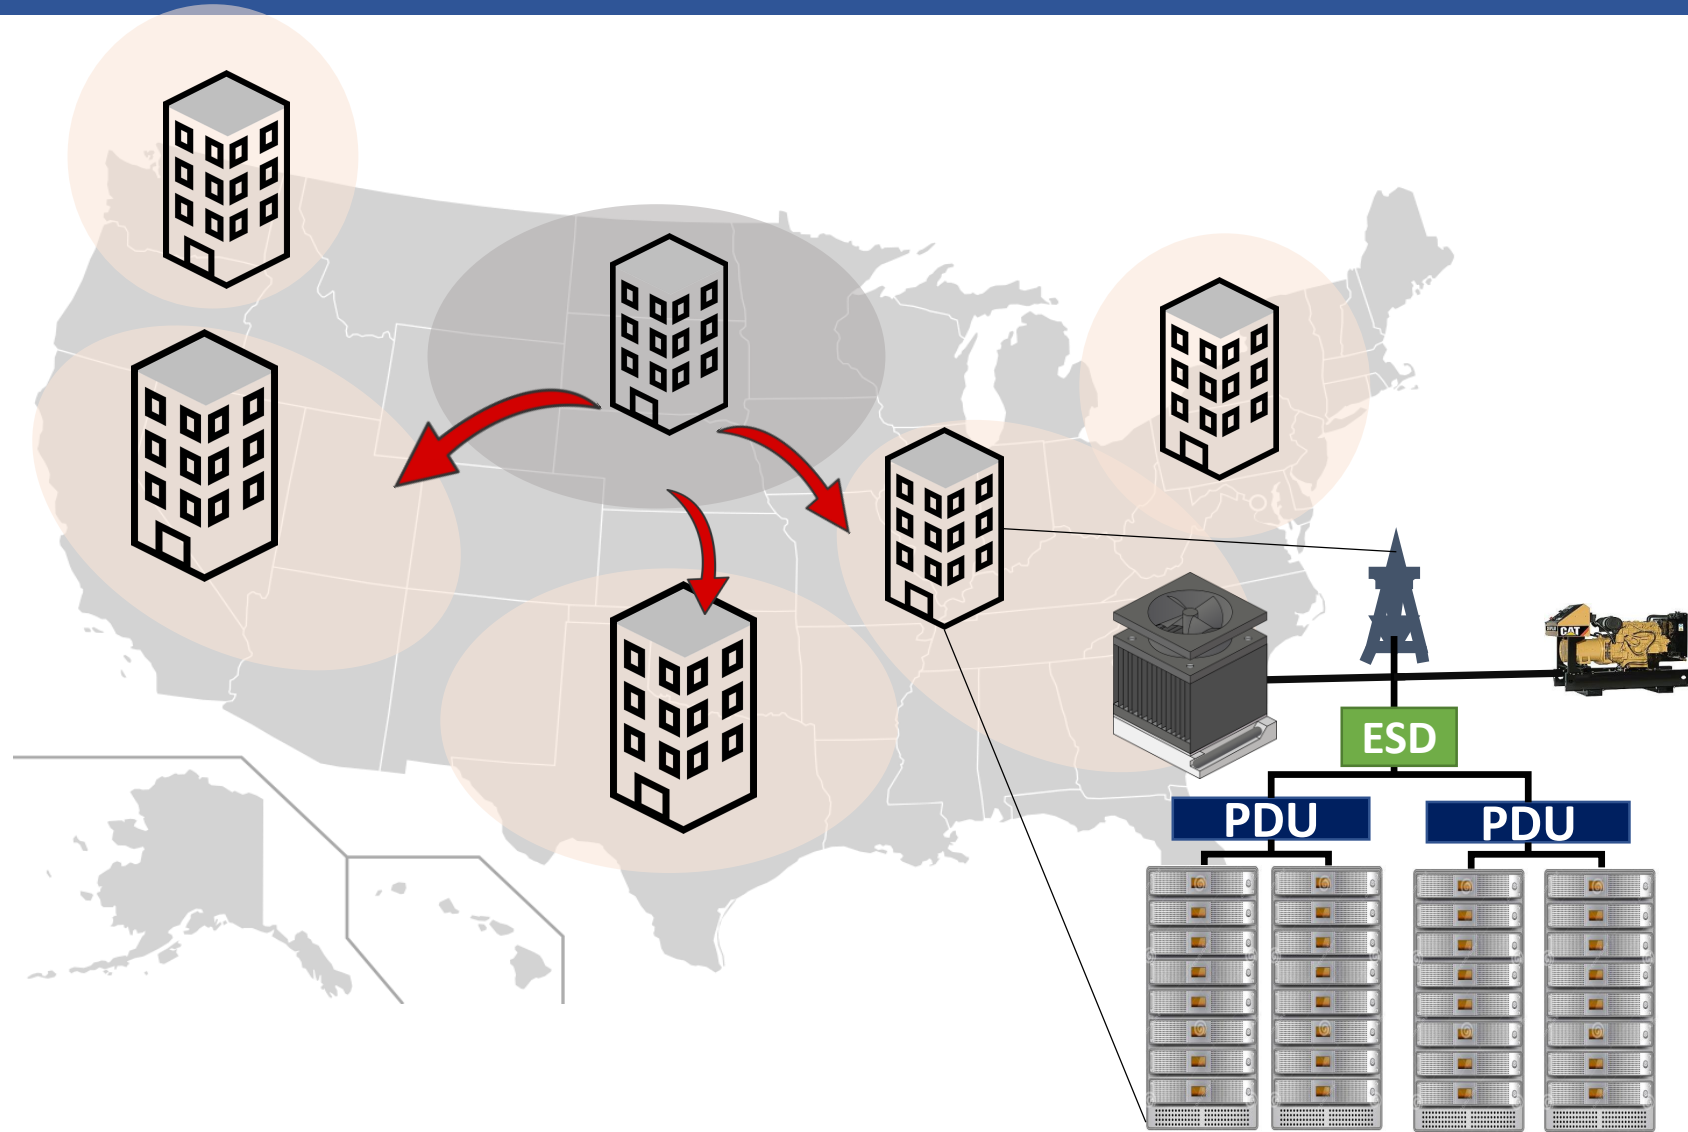
Geo-distribution helps to achieve low latency

Availability is important for apps

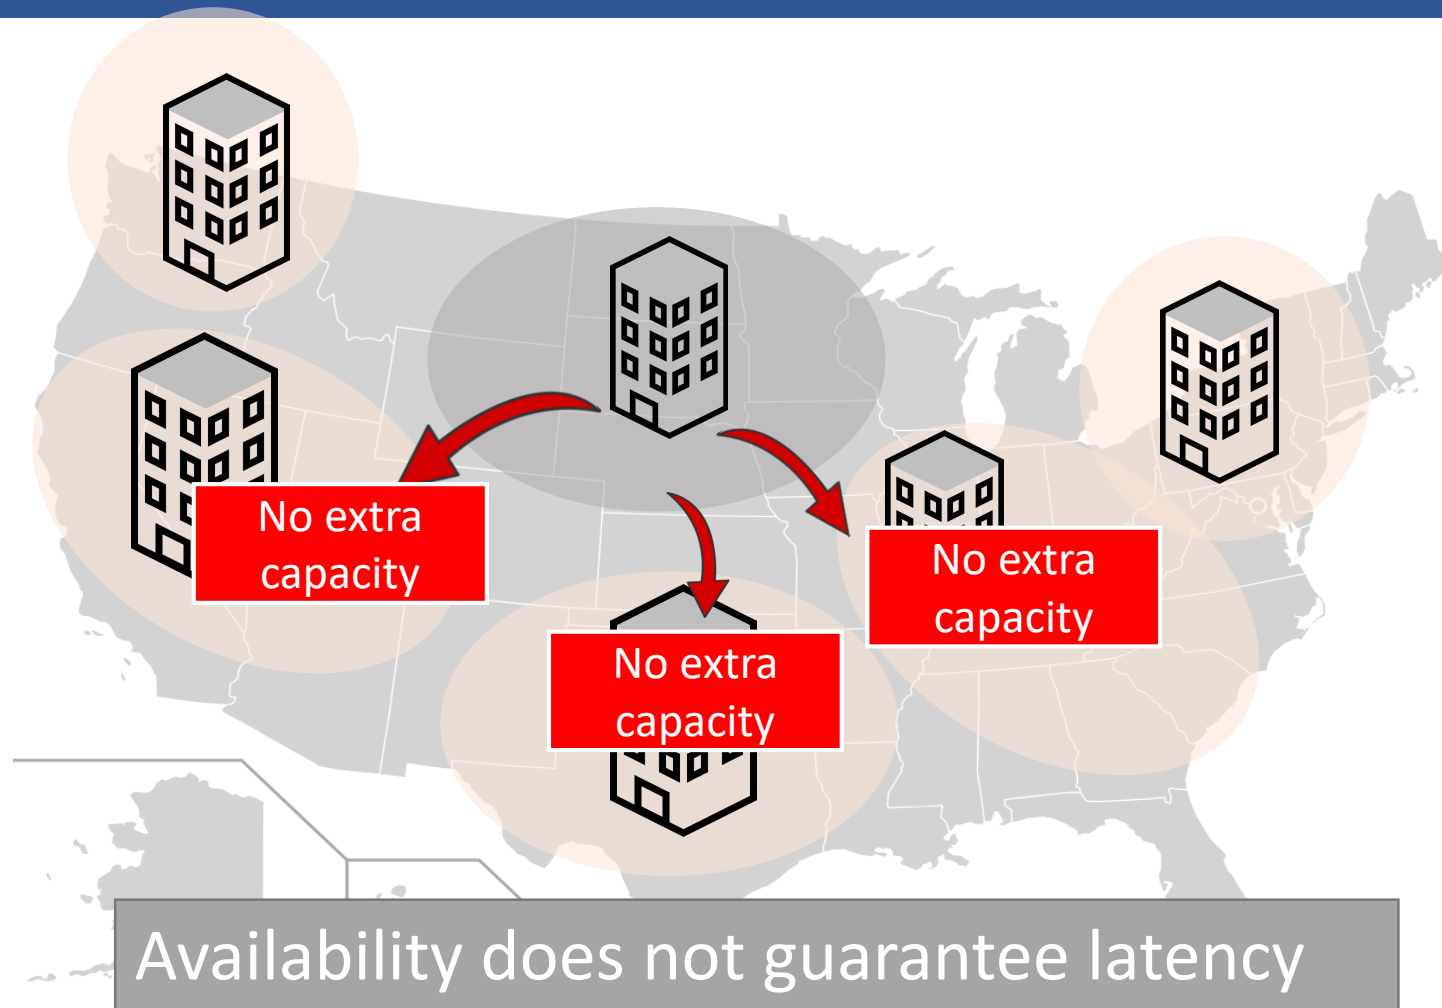


Geo-distribution can be leveraged to achieve high availability

Cloud infrastructure is not elastic

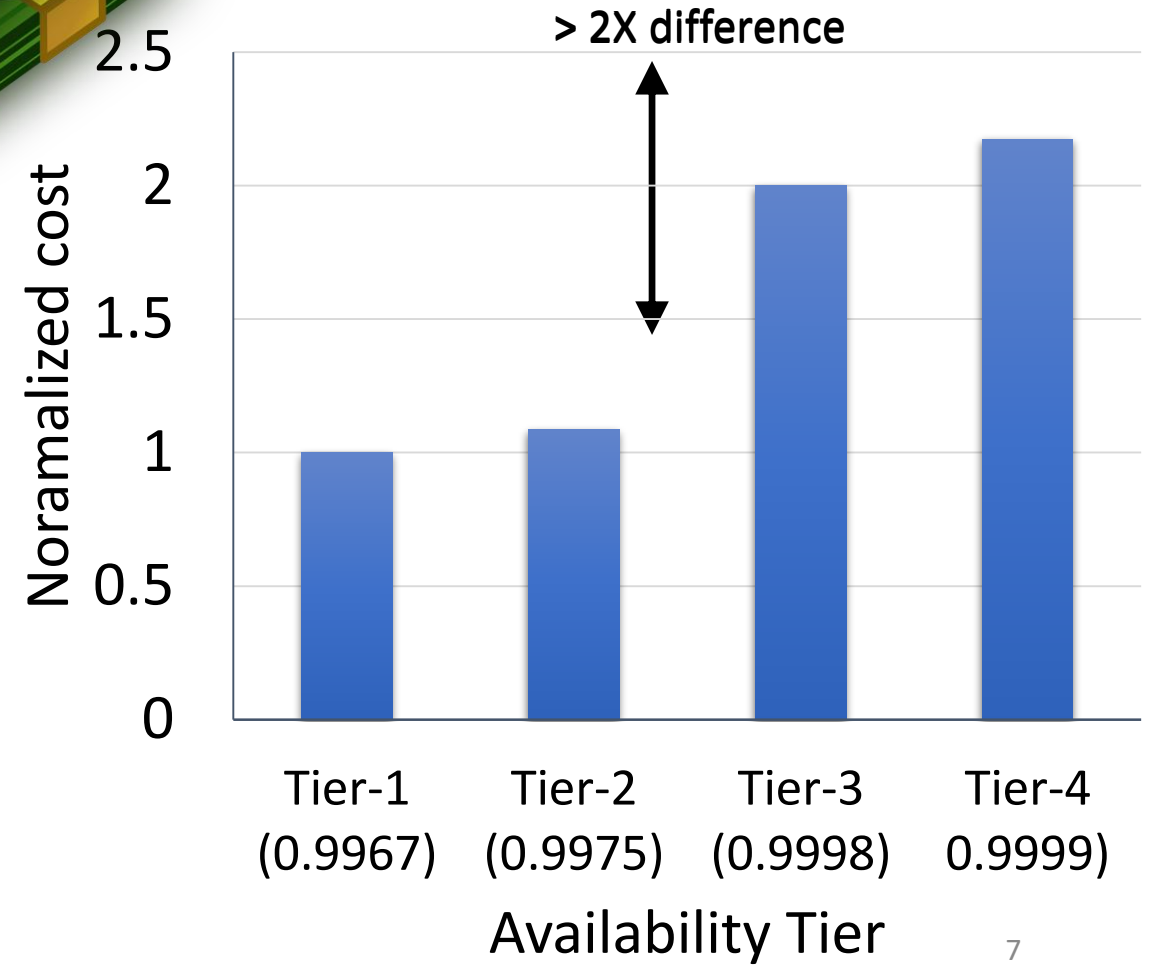
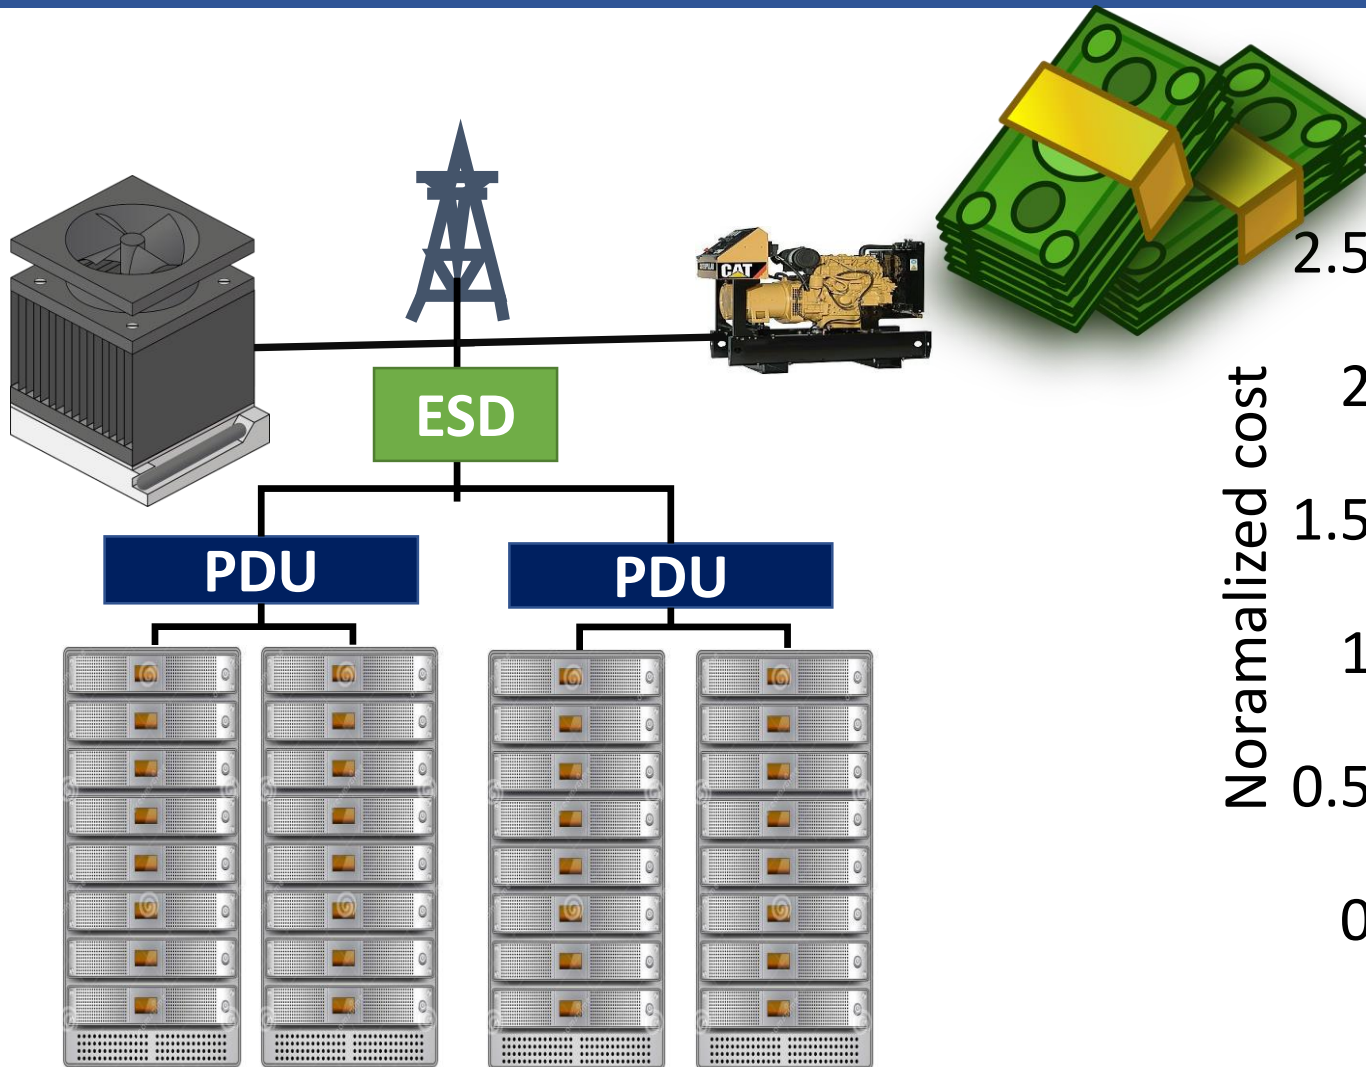


Cloud infrastructure is not elastic

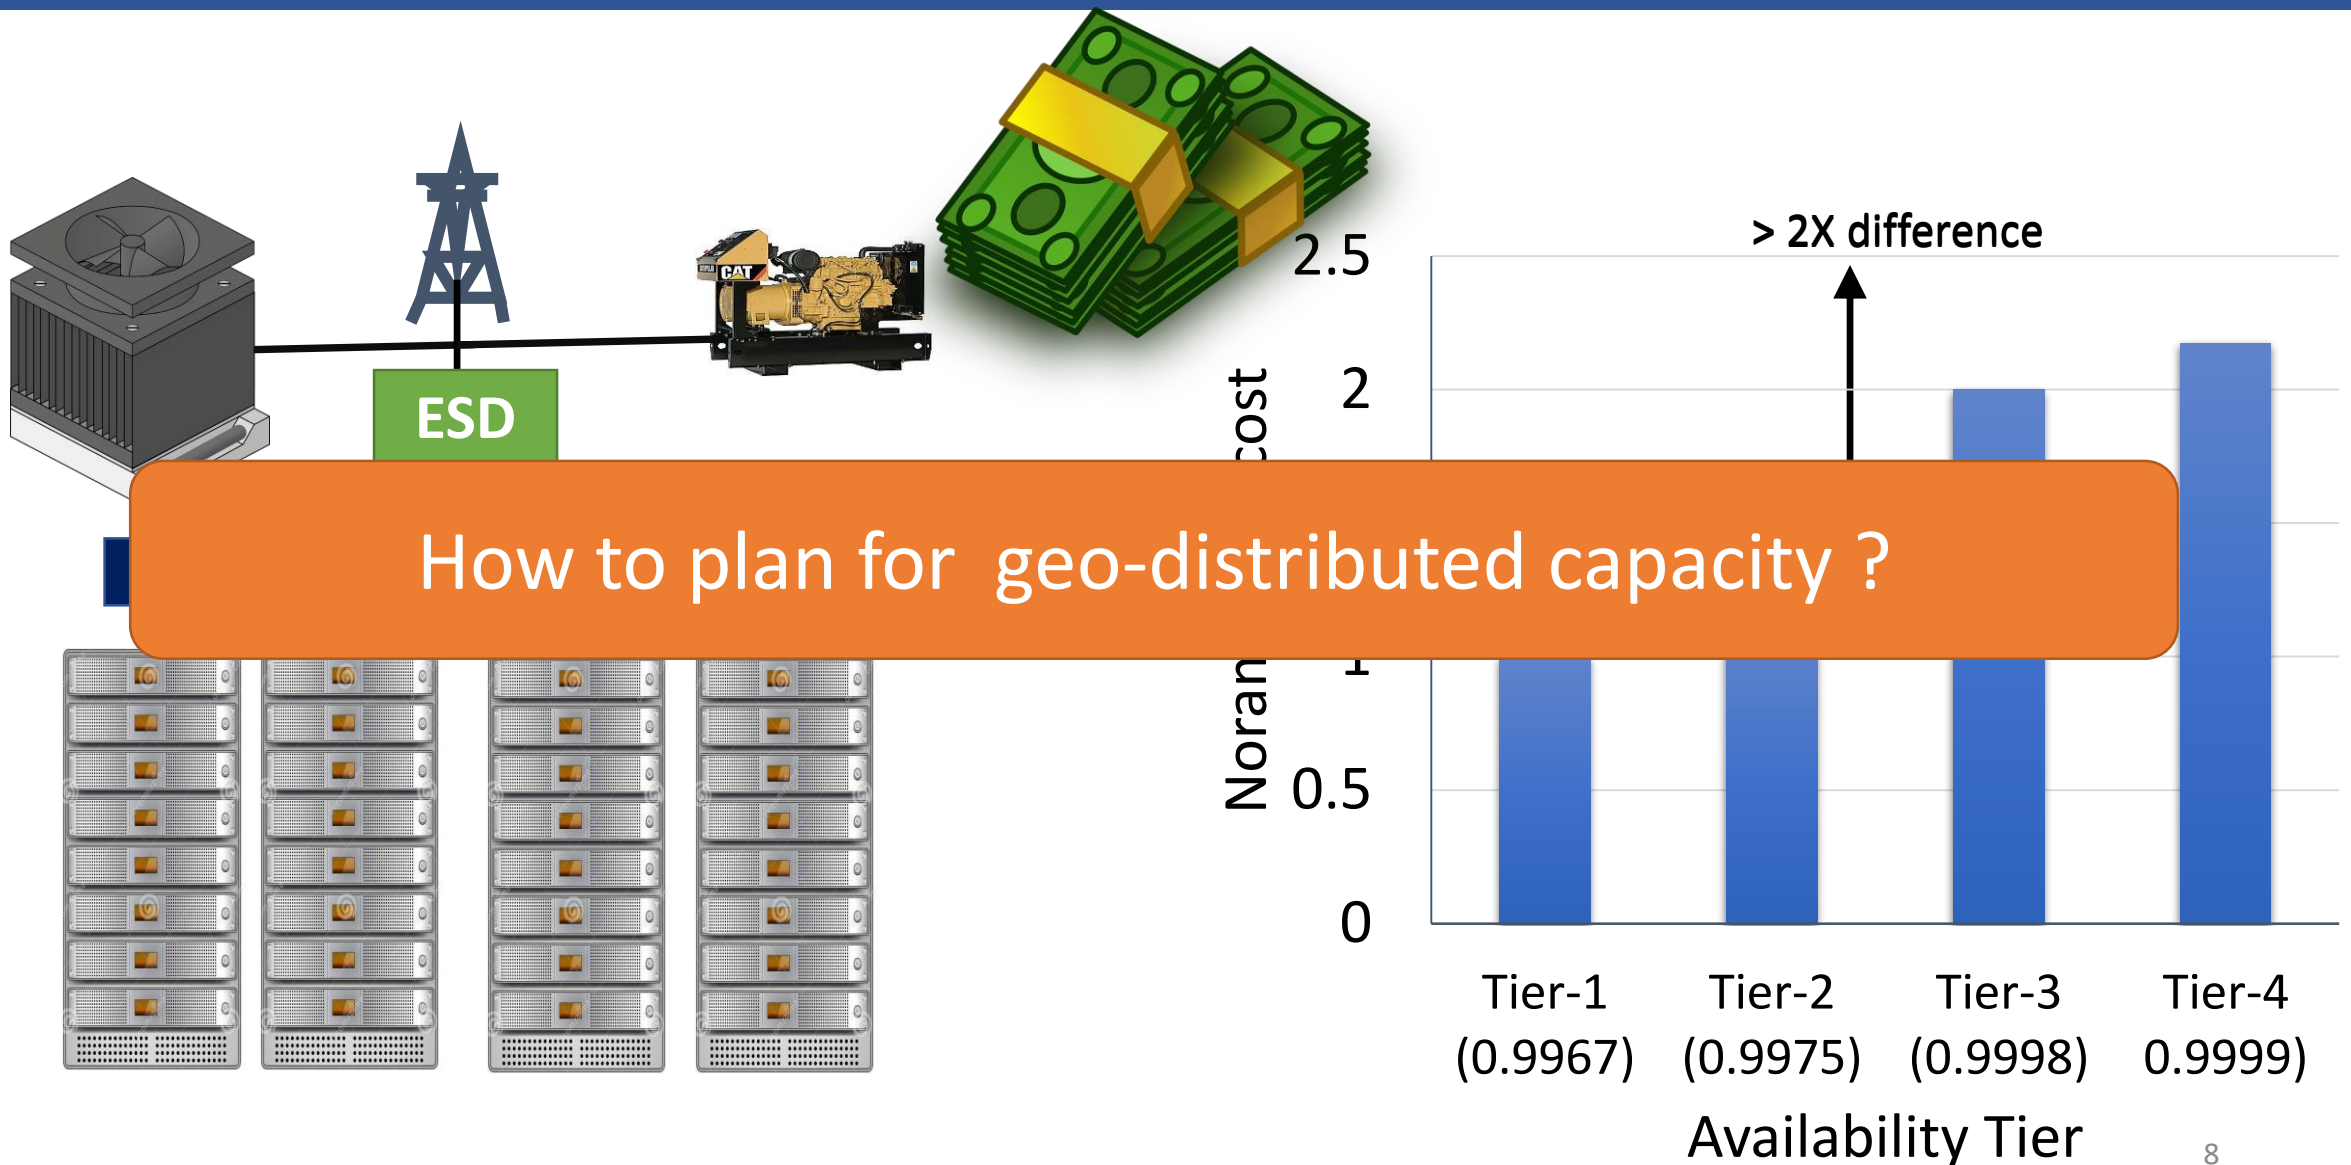


Physical infrastructure capacity has to be planned in advance

Capacity planning is important for operators, too

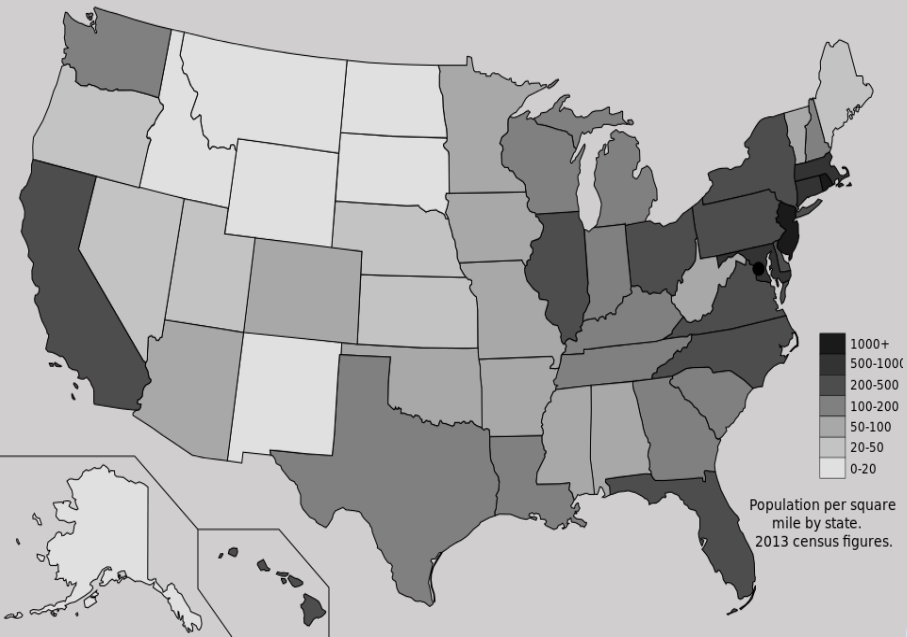


Capacity planning is important for operators, too



Geo-distributed capacity planning

Geography



Datacenter

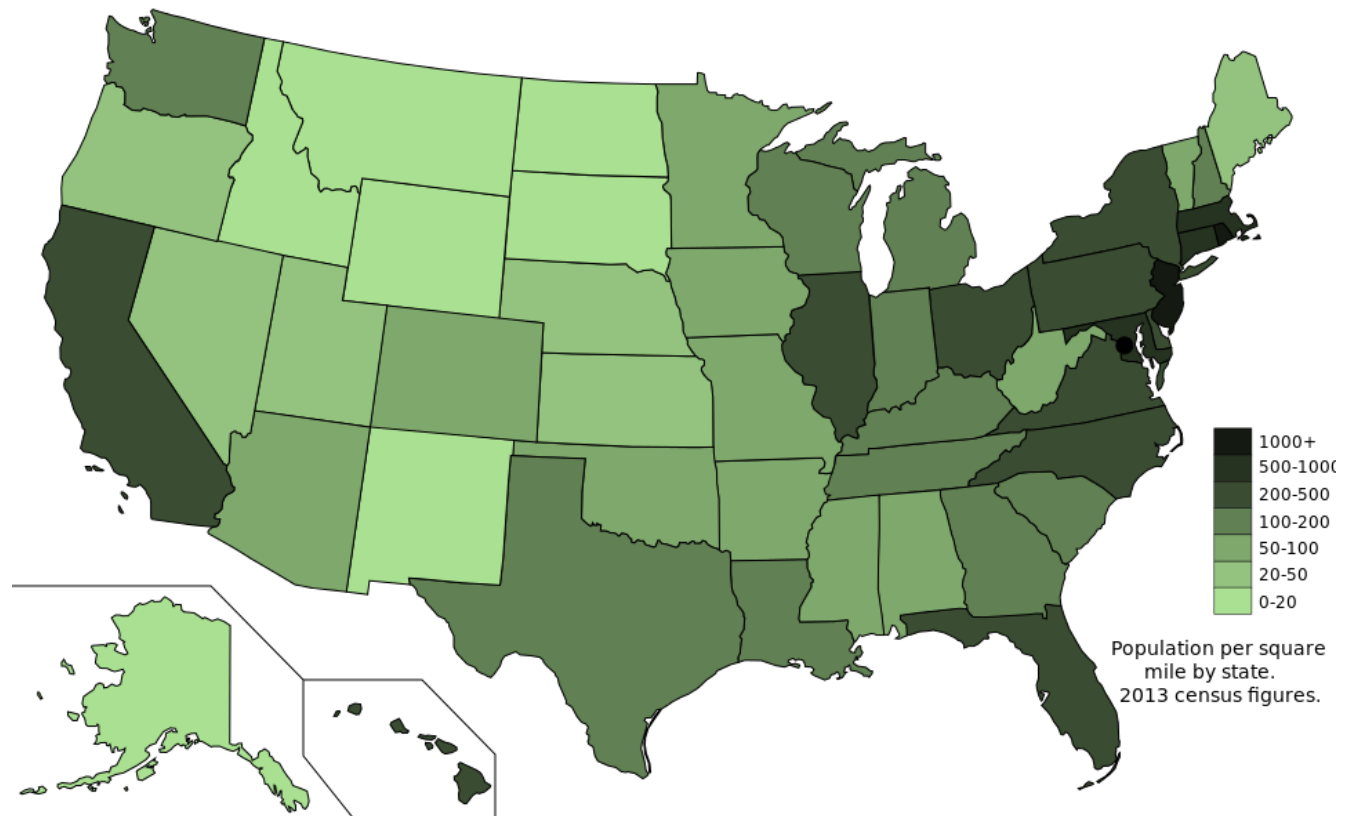


Application



Geographical factors

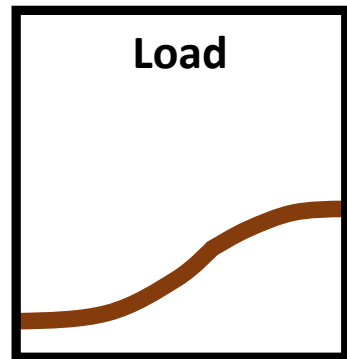
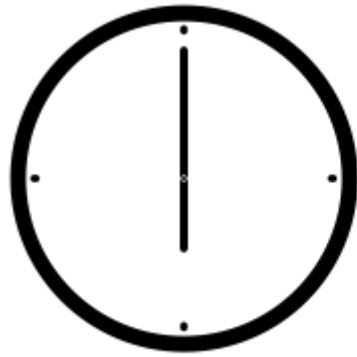
Client distribution may vary with population



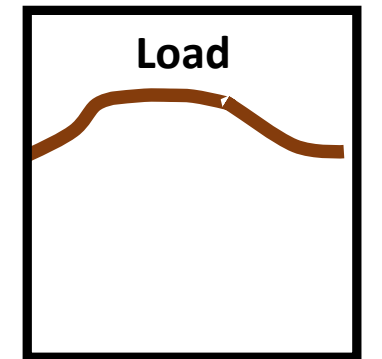
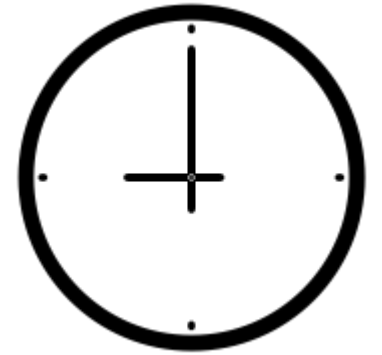
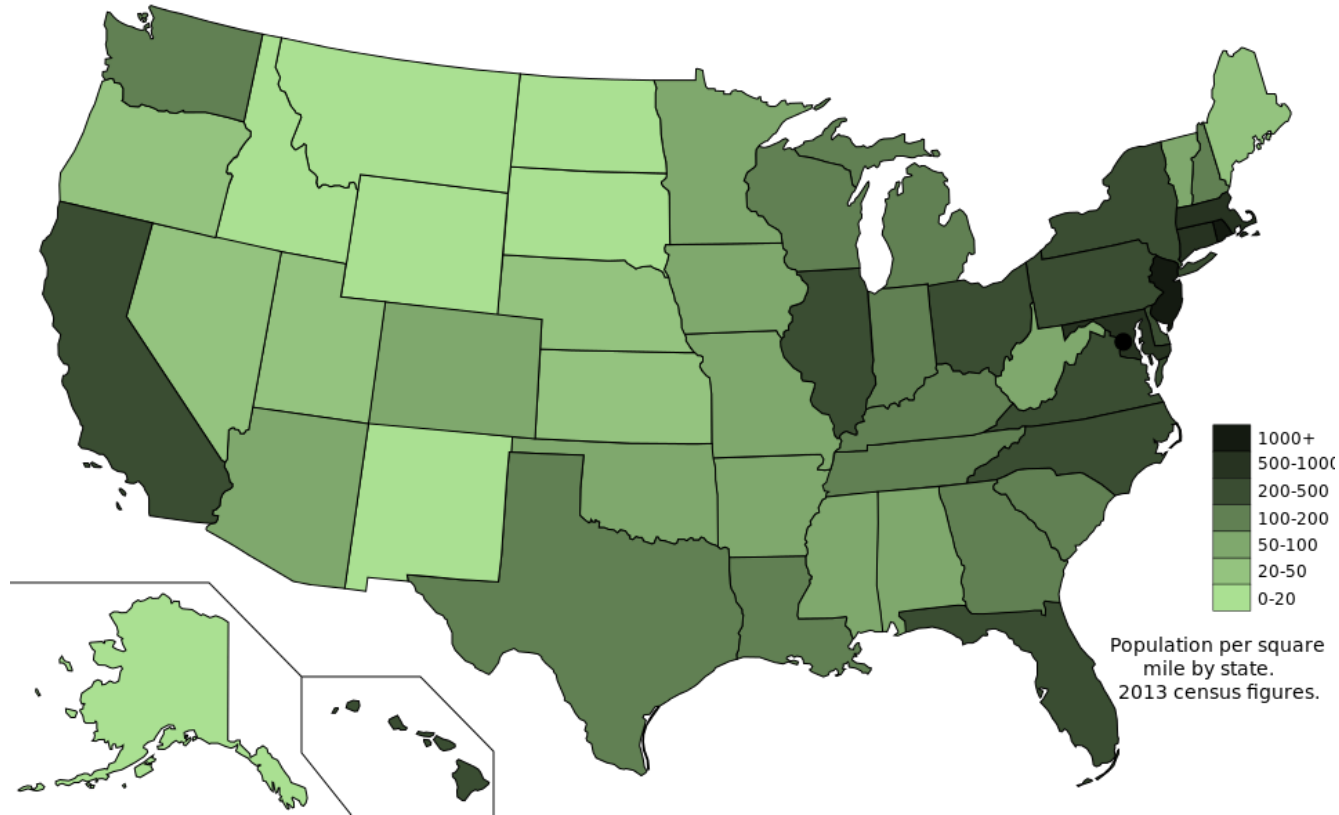
Source: Wikipedia

Geographical factors

Load can vary with time zones



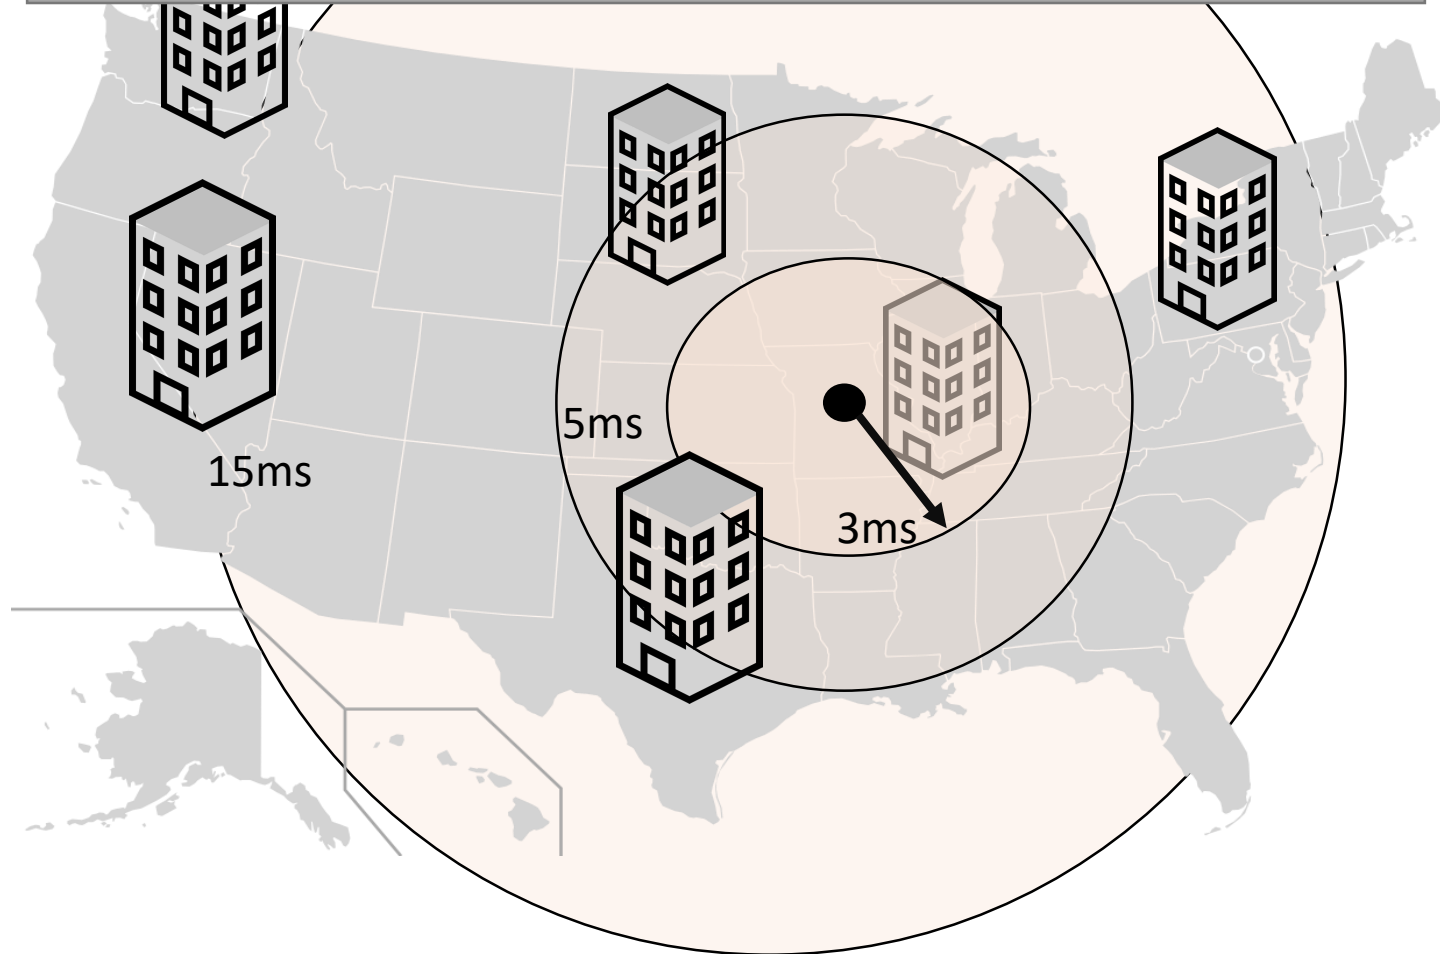
5 am 6 am 7 am



8 am 9 am 10 am

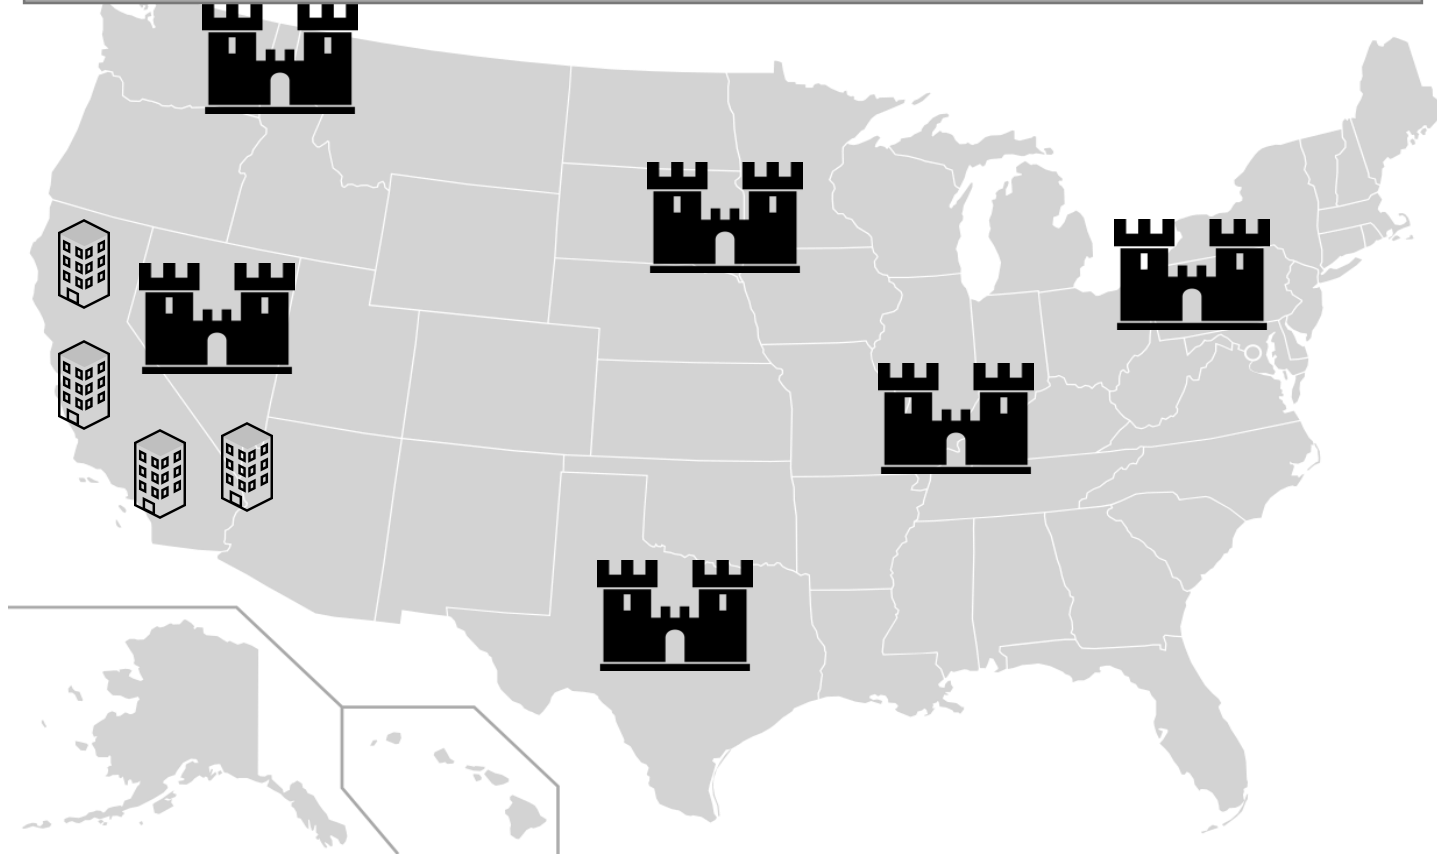
Geographical factors

Latency affects capacity provisioning

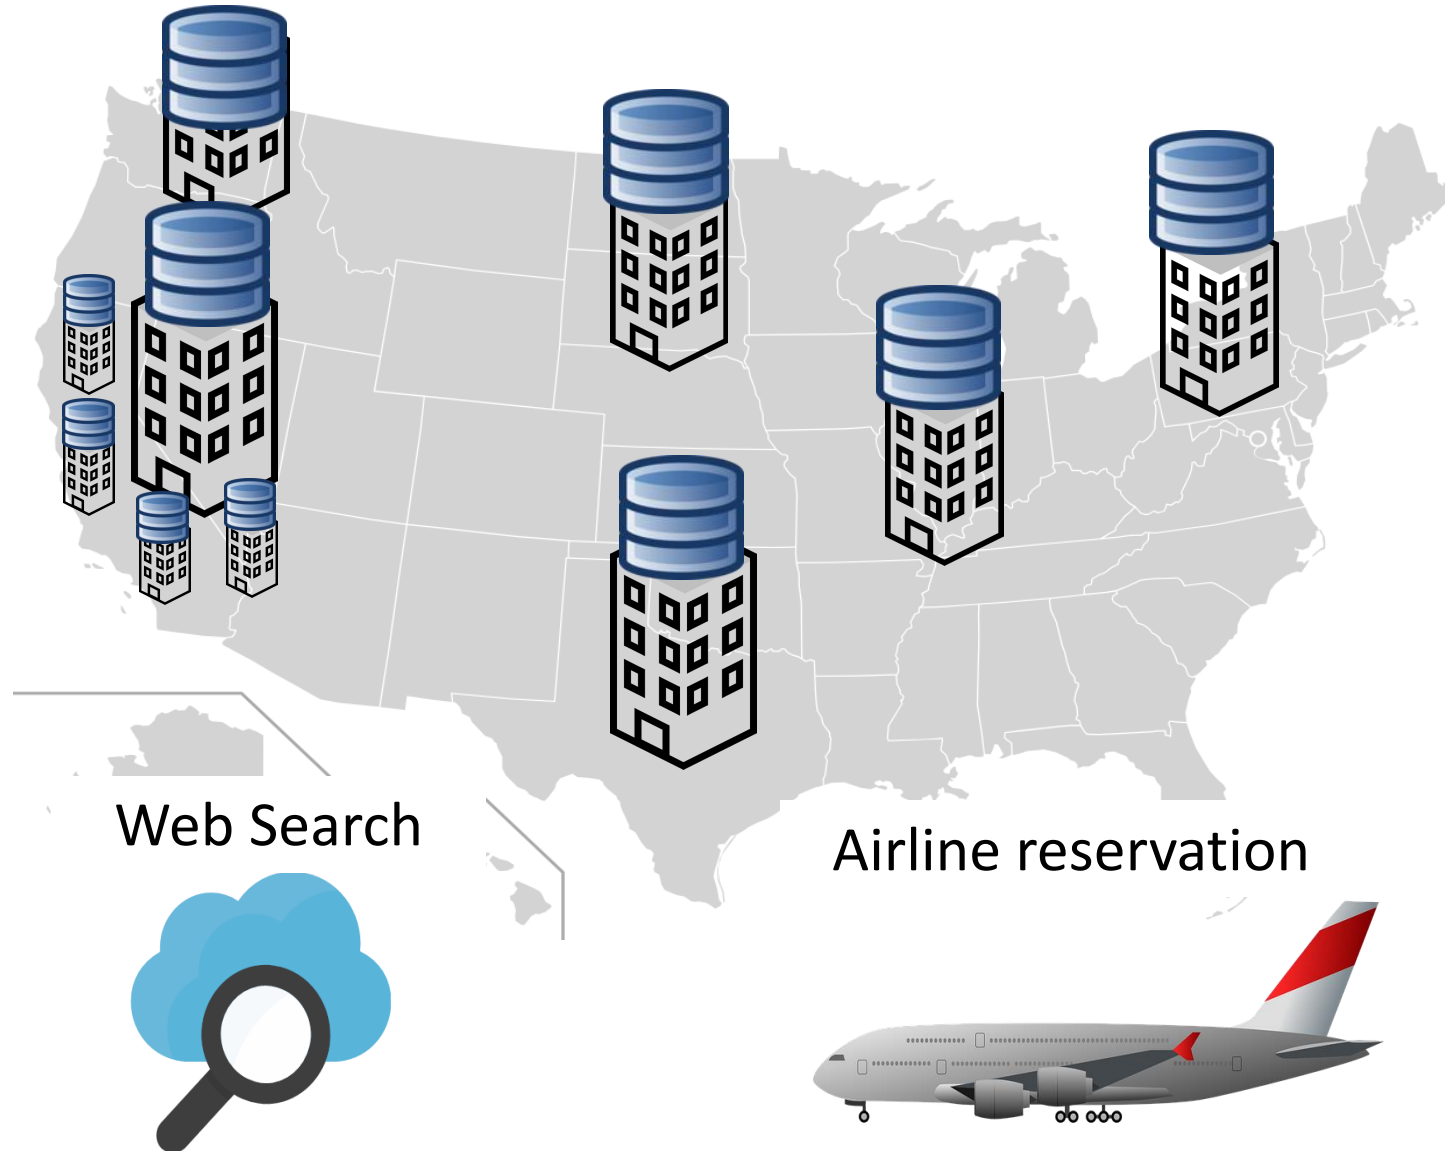


Datacenter factors: Availability

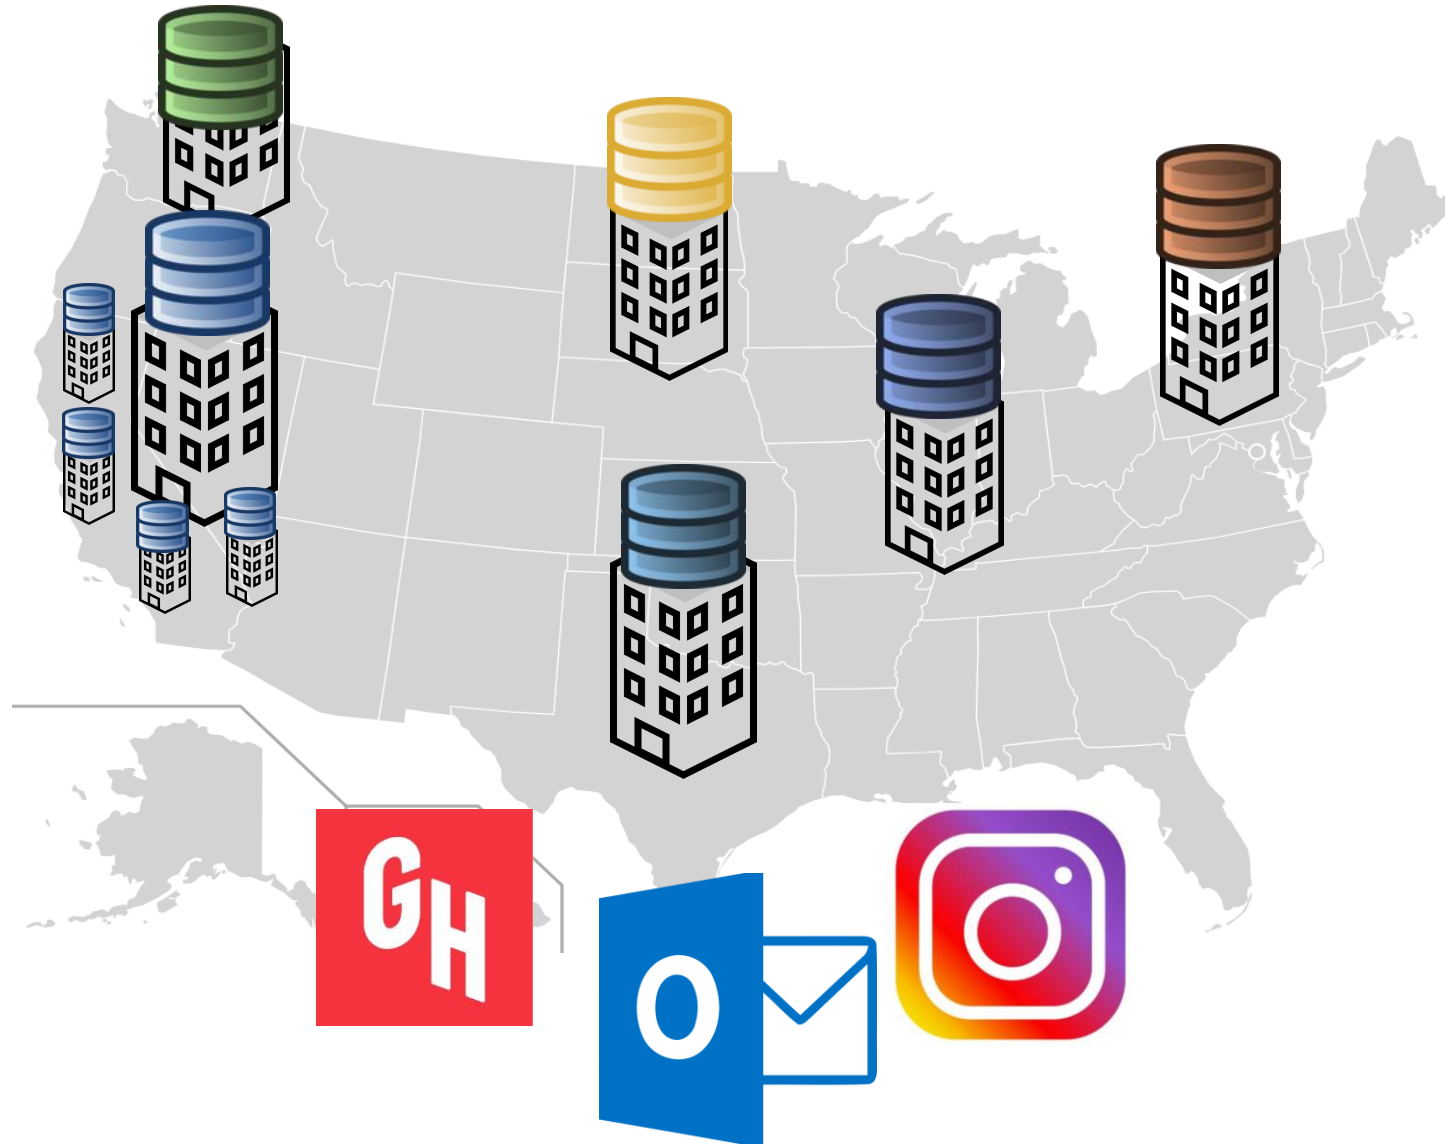
Availability impact geo-distributed spares



Application factors: Replicated data



Application factors: Partitioned data



Factors affecting capacity planning

Geography



Datacenter



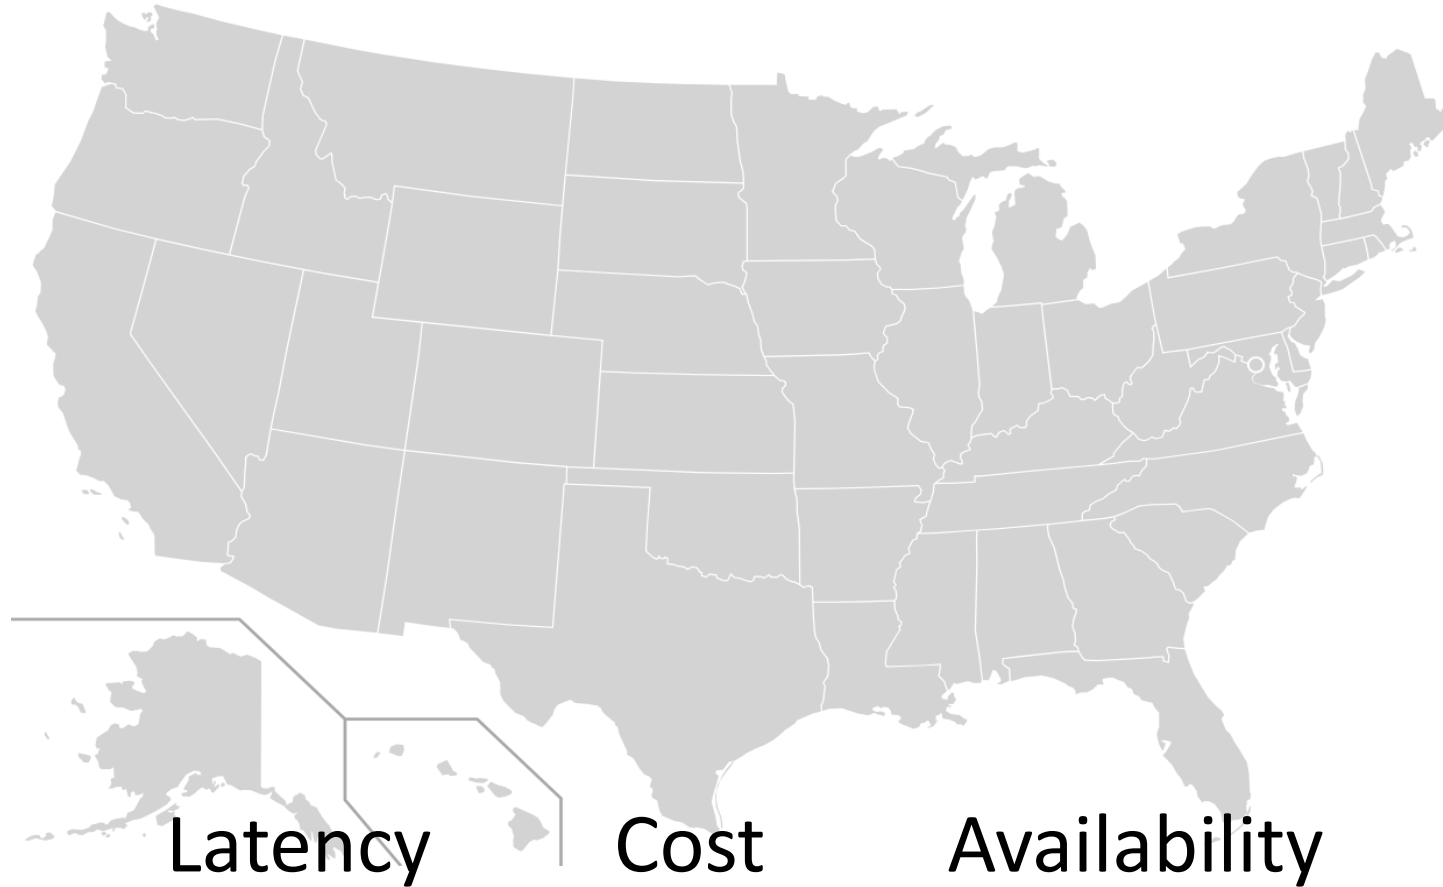
Application



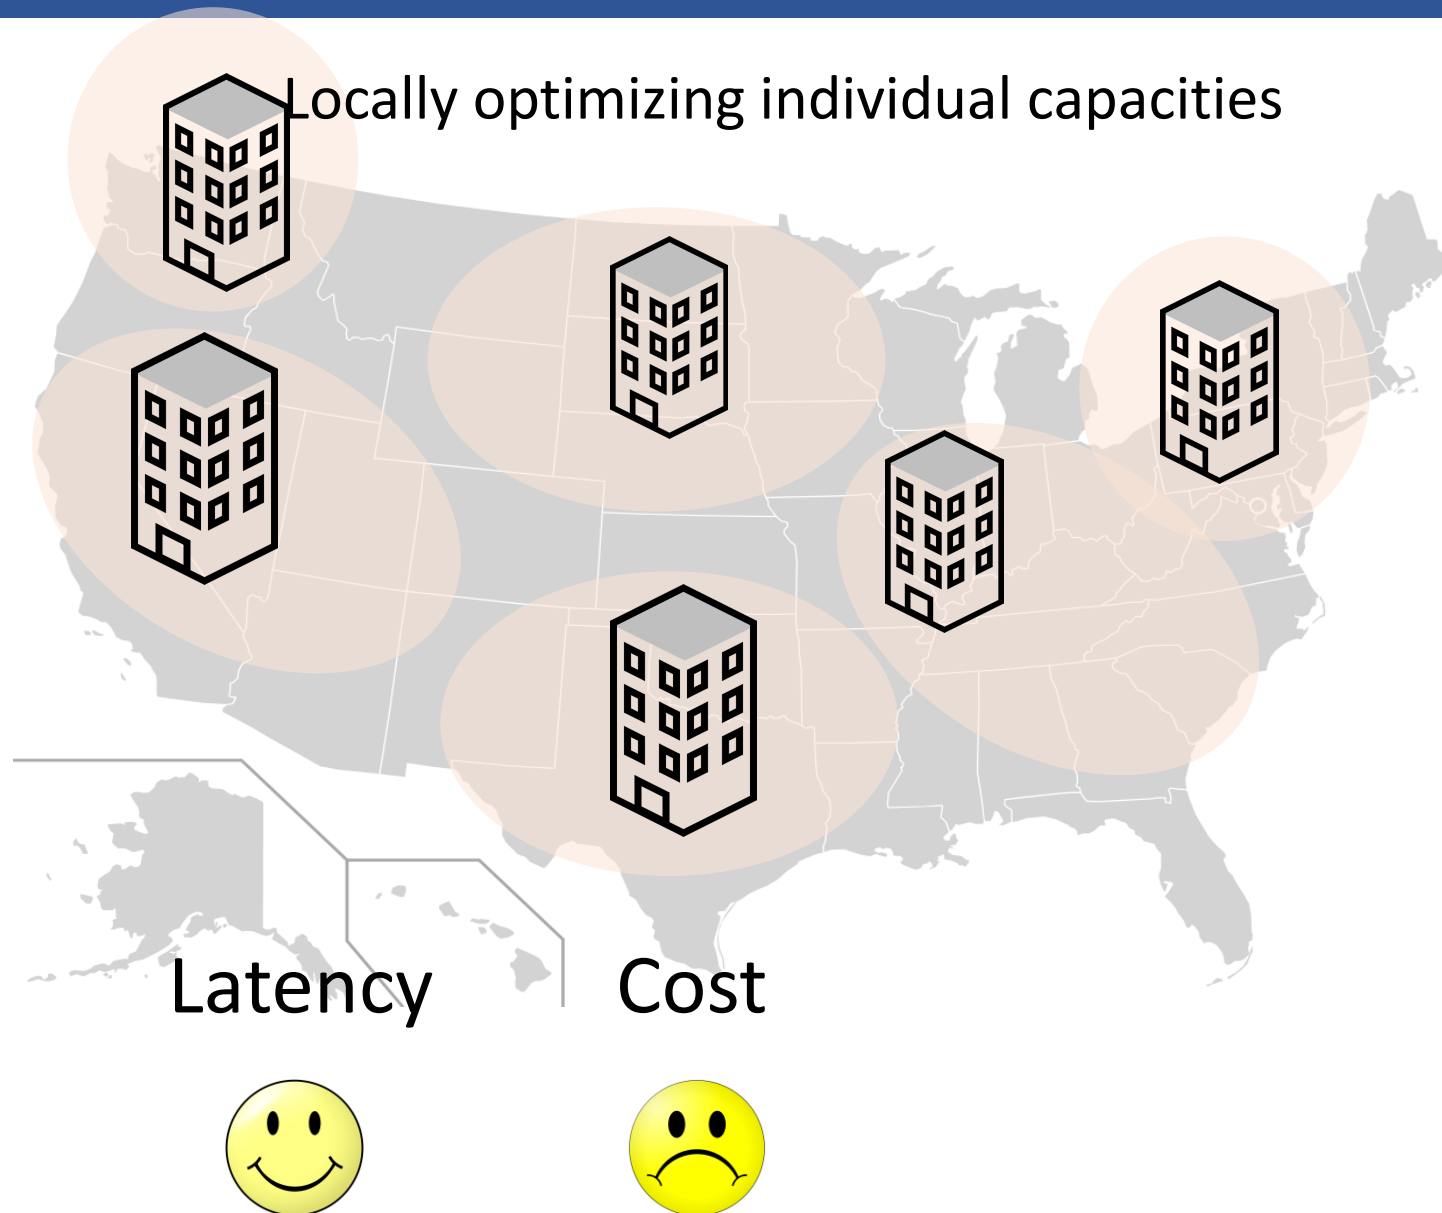
What are the challenges ?

Capacity planning challenges

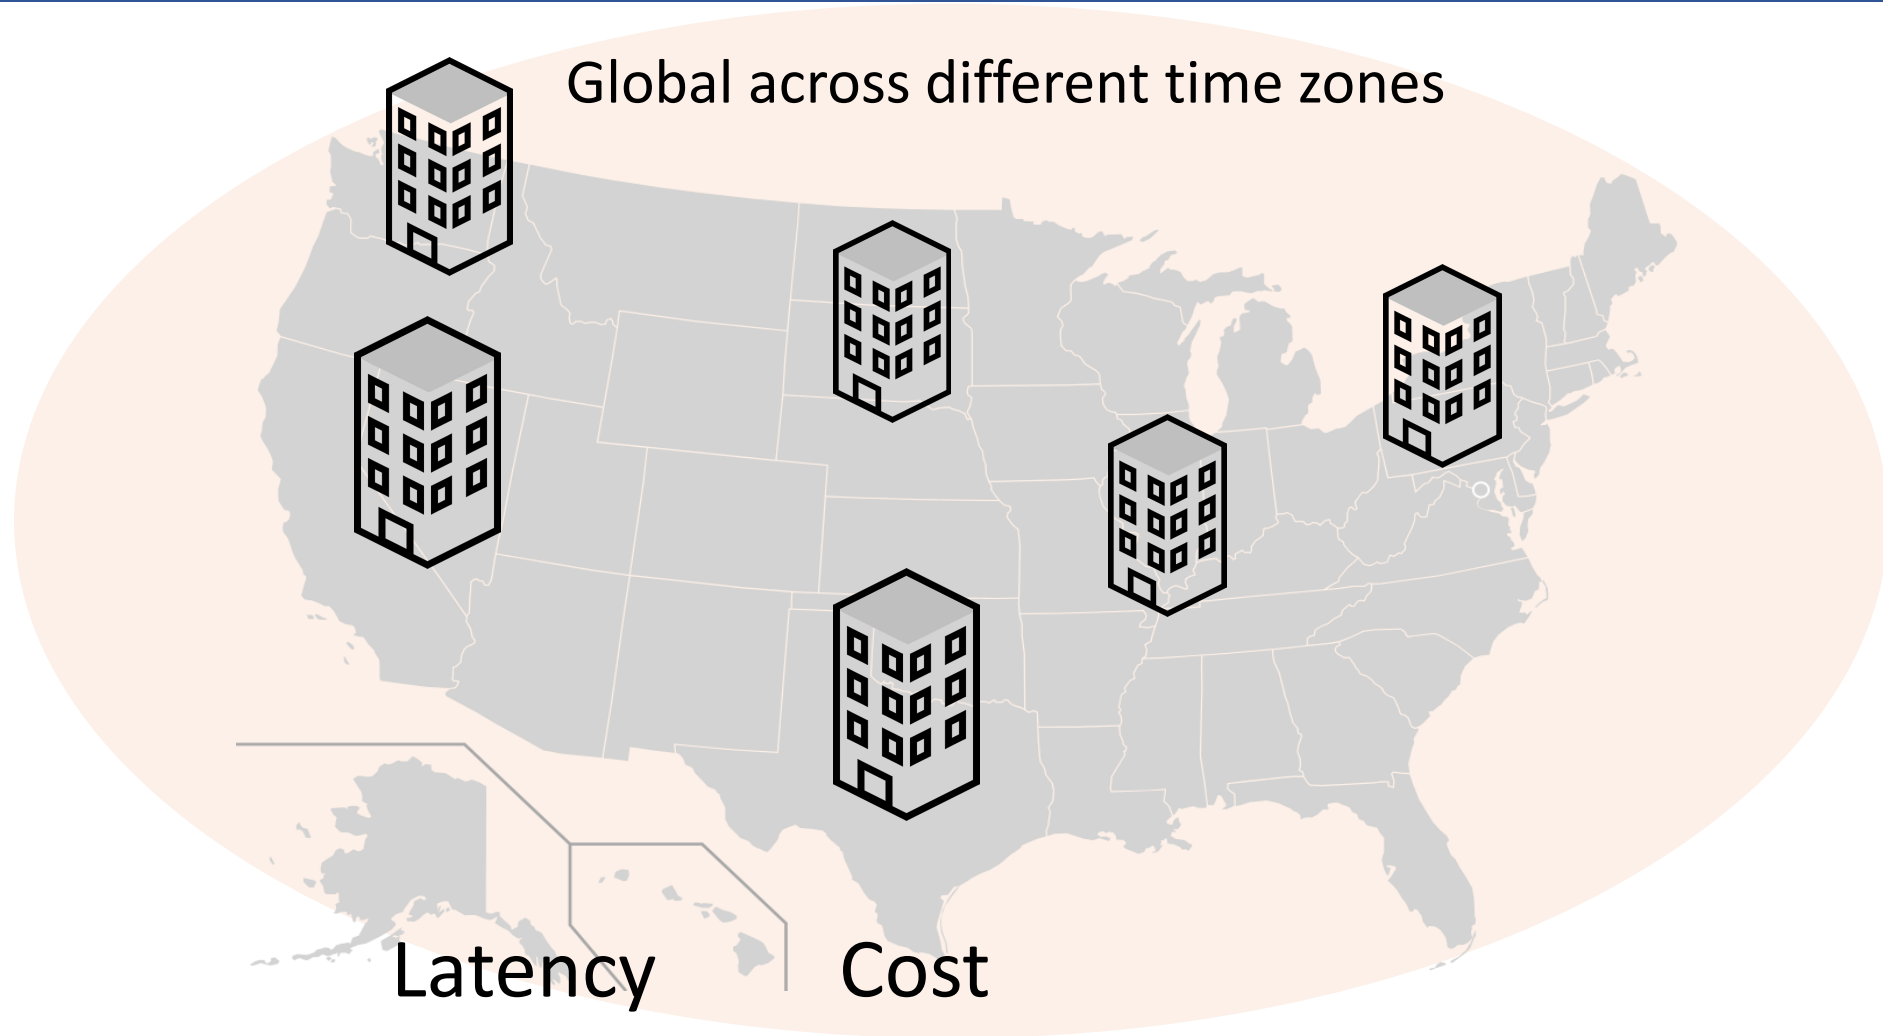
Performance metrics



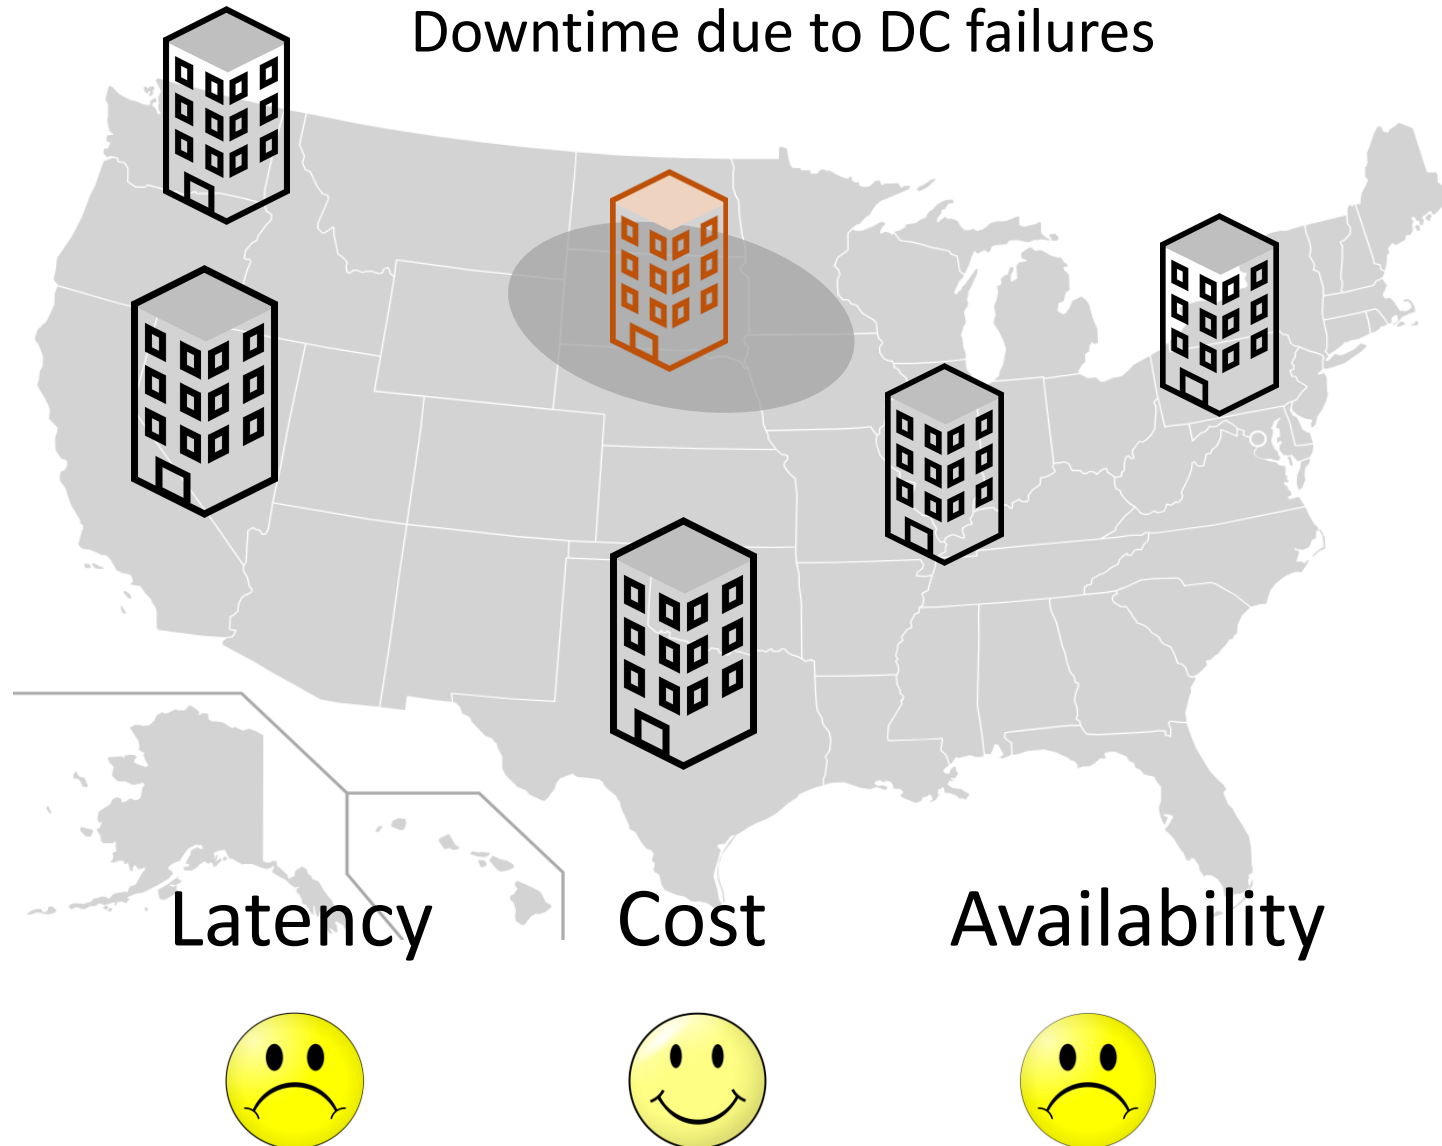
Local planning vs. global planning



Local planning vs. global planning

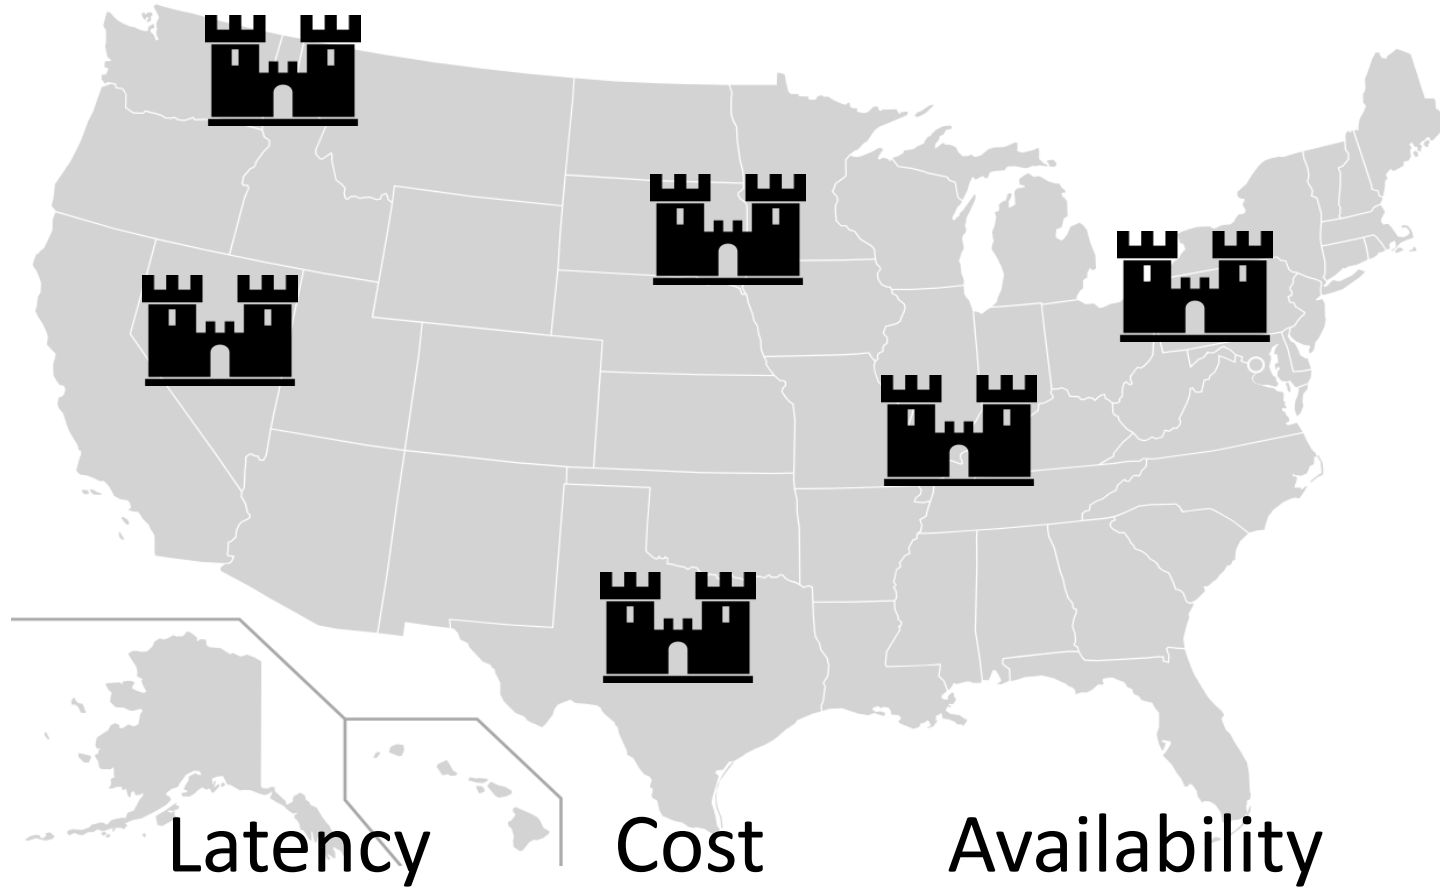


Handling geo-distributed failures

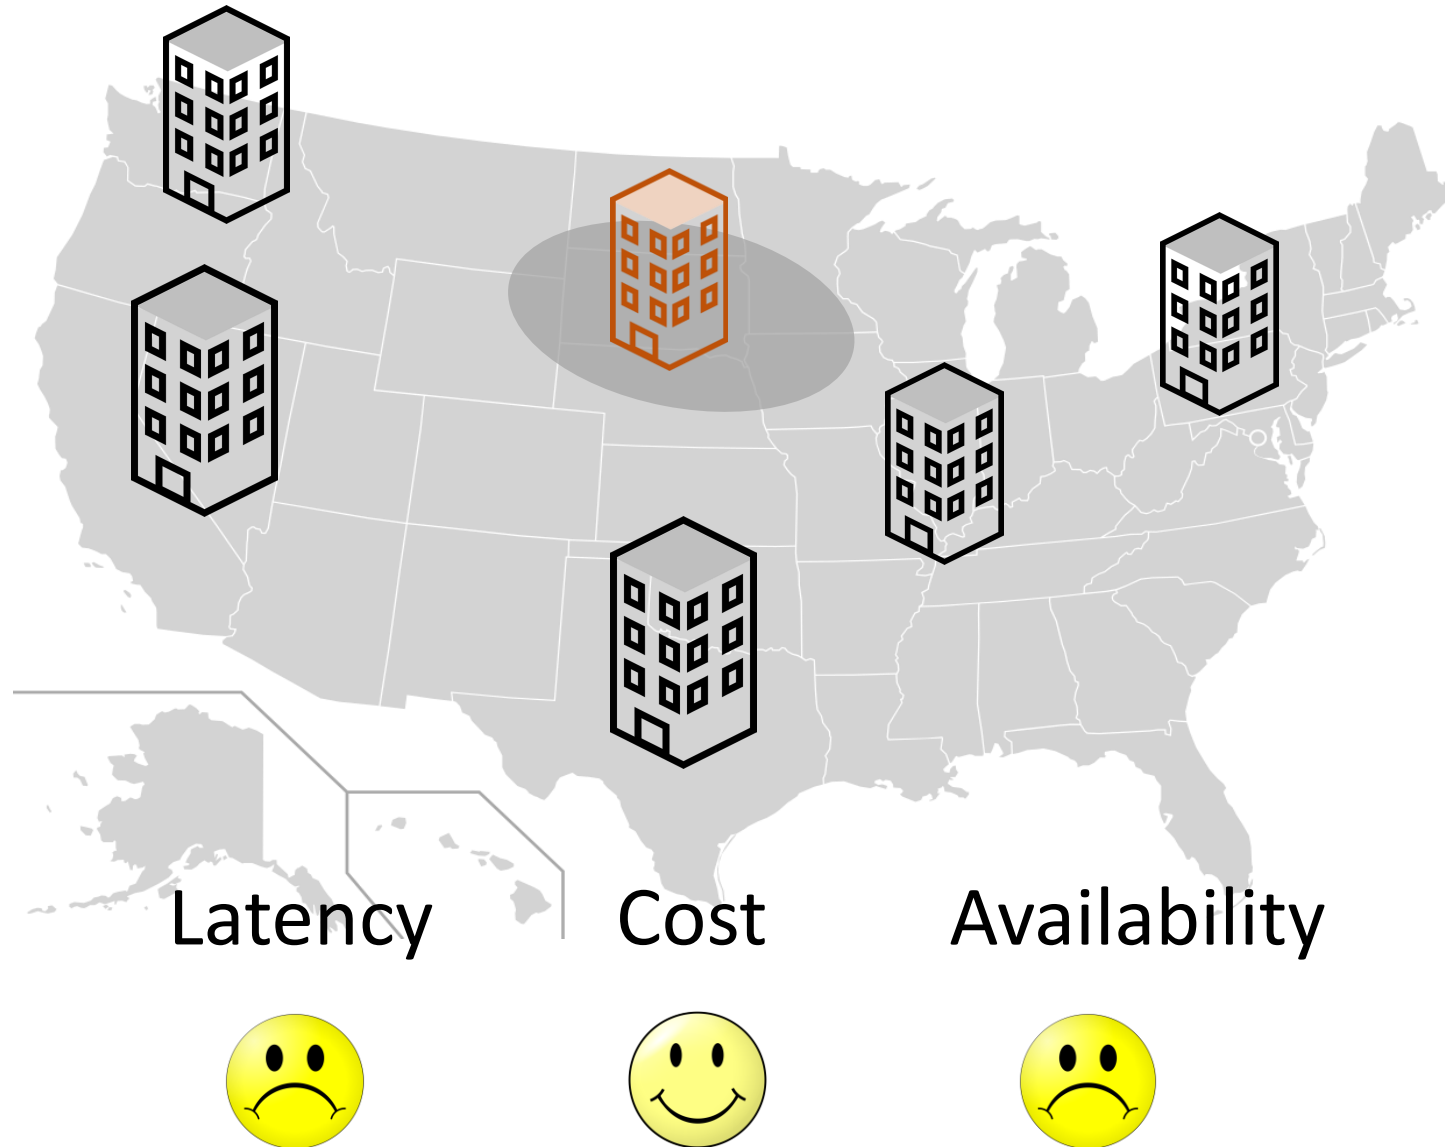


Handling geo-distributed failures

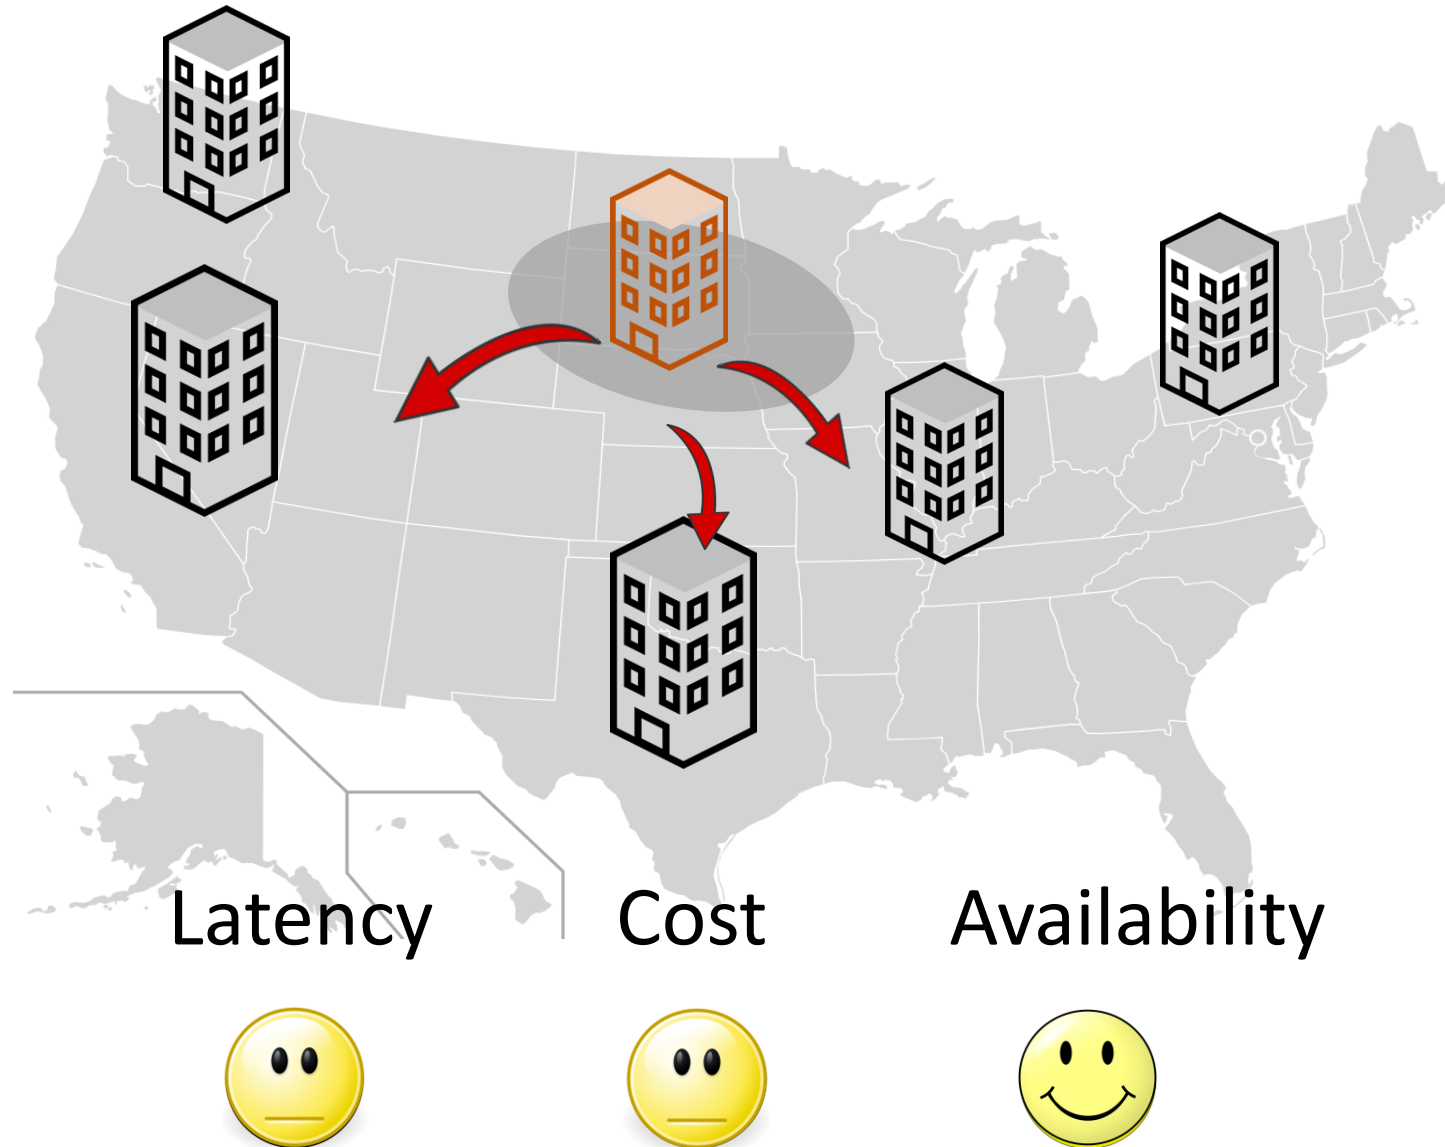
Hardening individual DCs



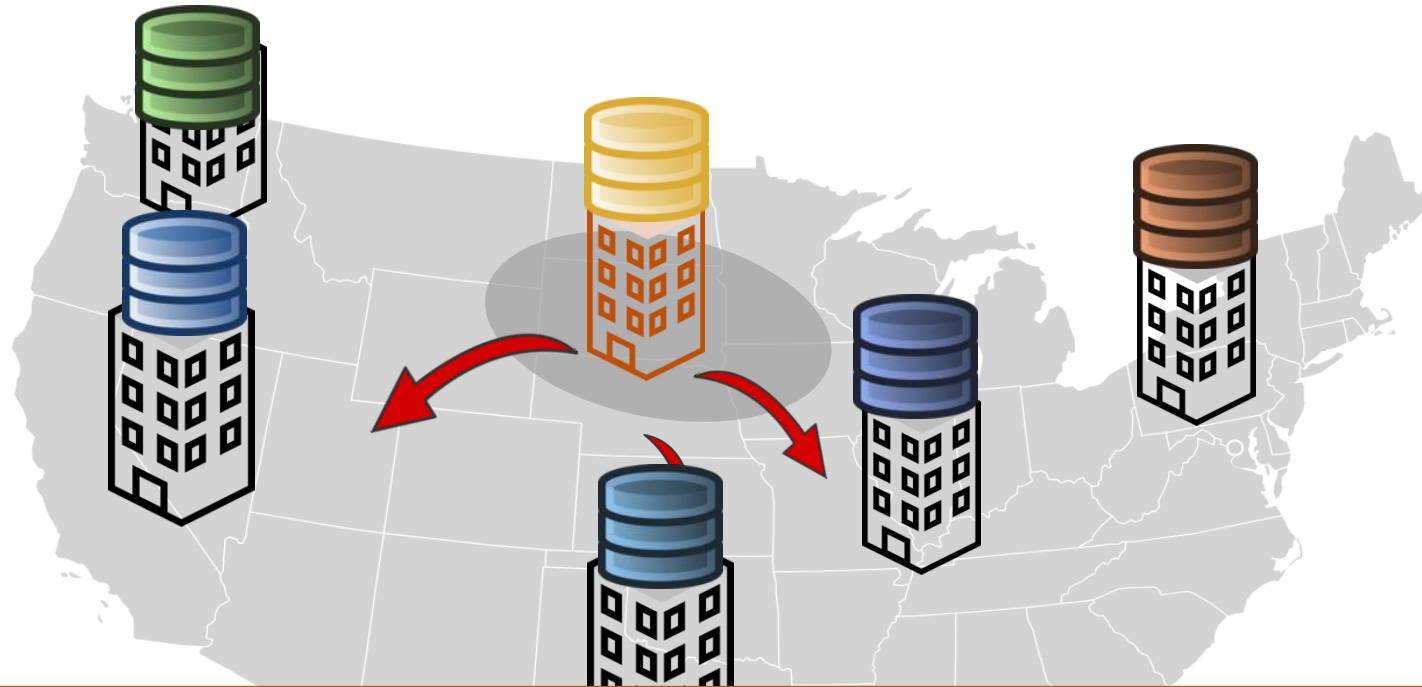
Handling geo-distributed failures



Handling geo-distributed failures



Handling geo-distributed failures



How to guarantee latency and availability for apps
cost effectively ?

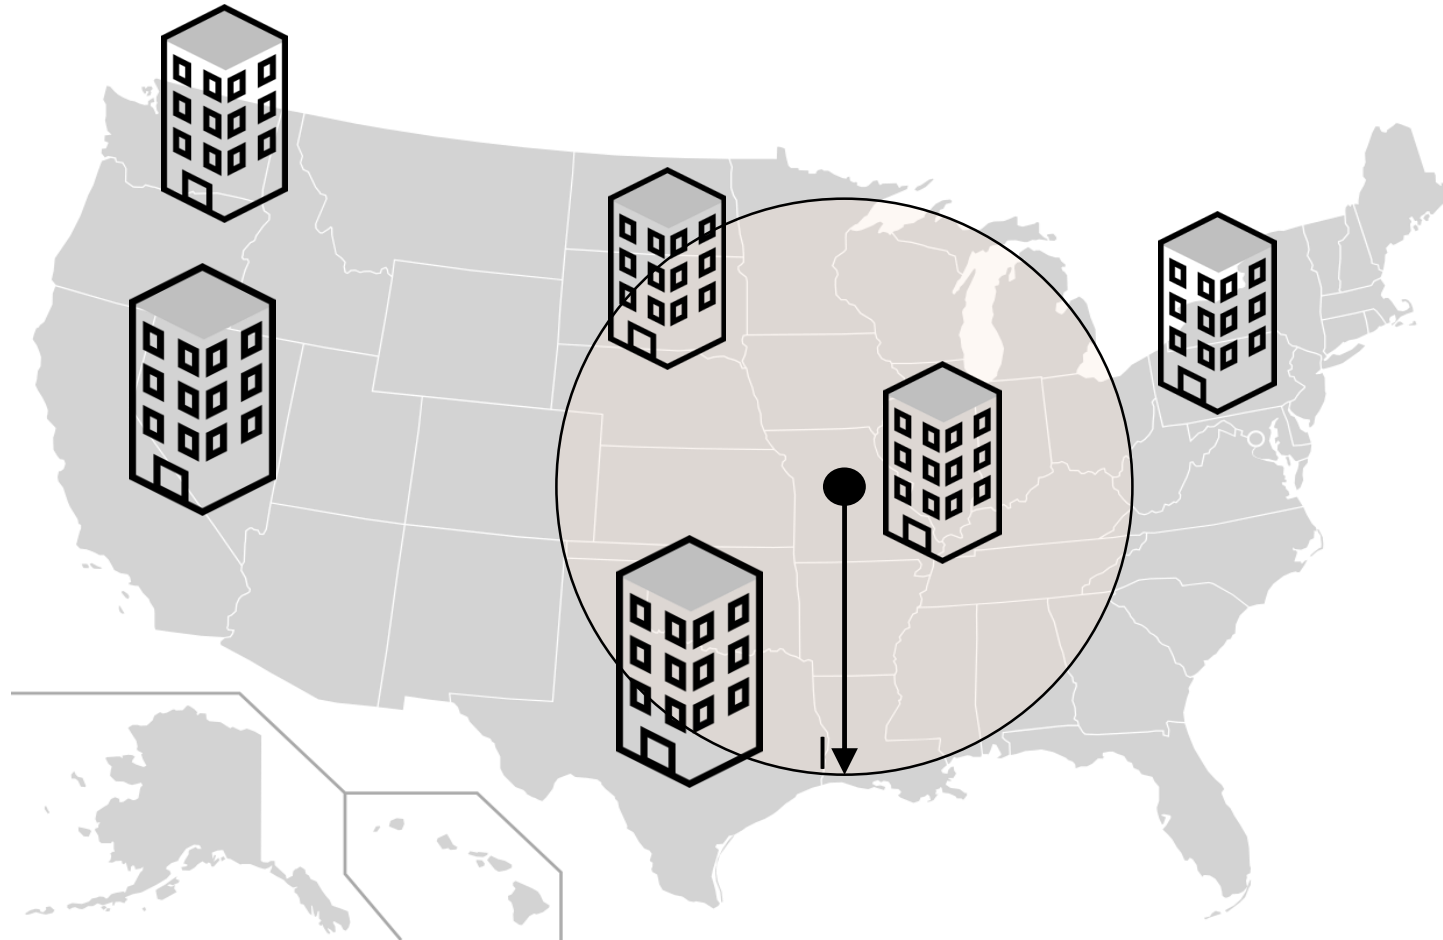
Latency

Cost

Availability

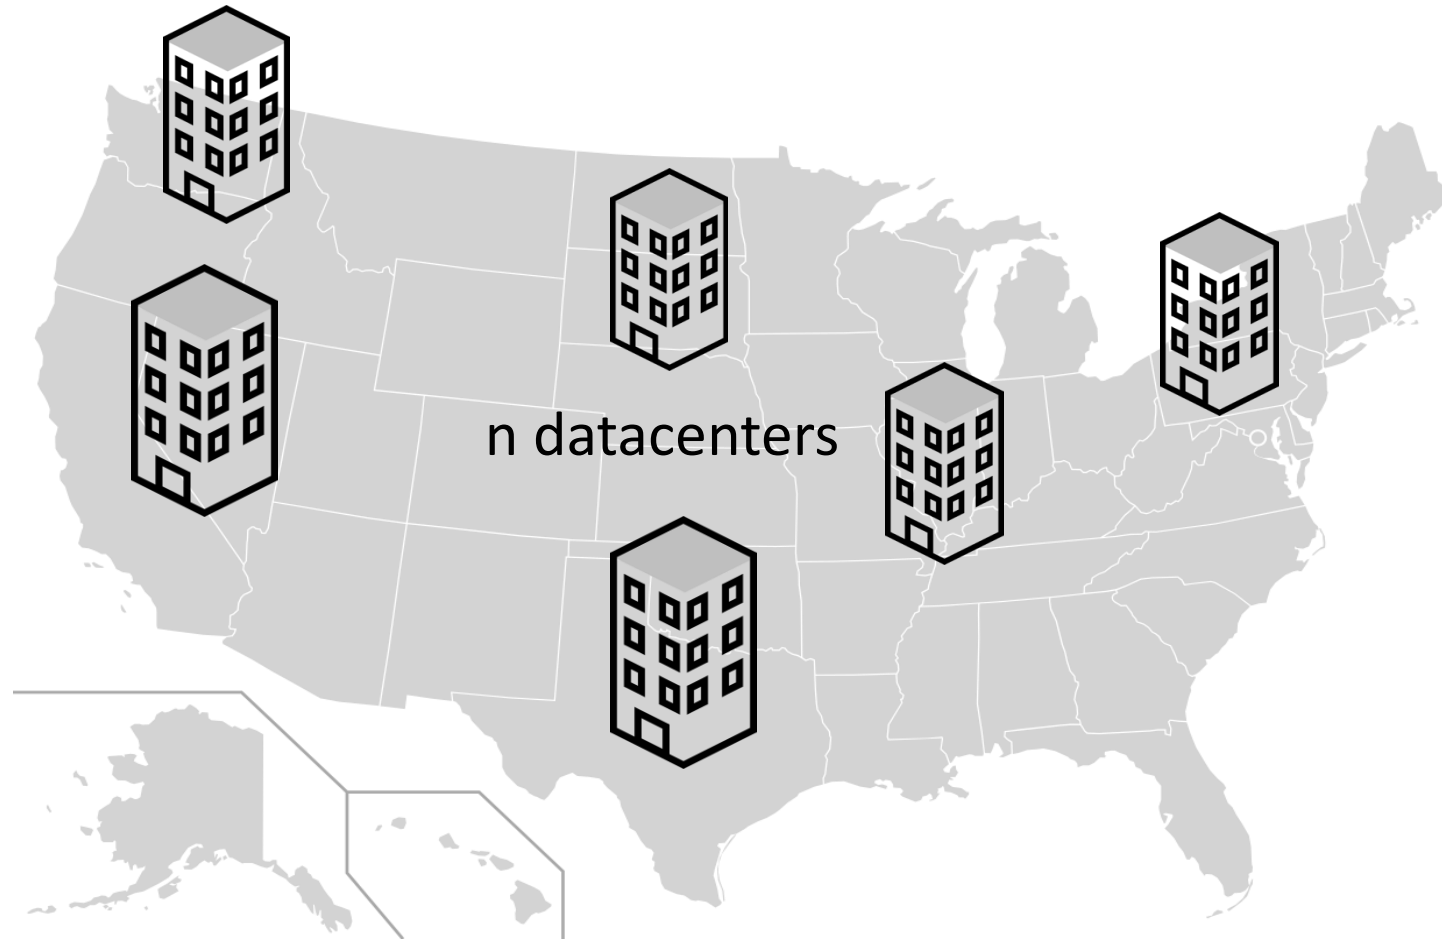


System model



Client desires Tier-4 availability

System model



Optimization objective

$$\min \sum_{j \in \{1, \dots, n\}} x_j C_j$$

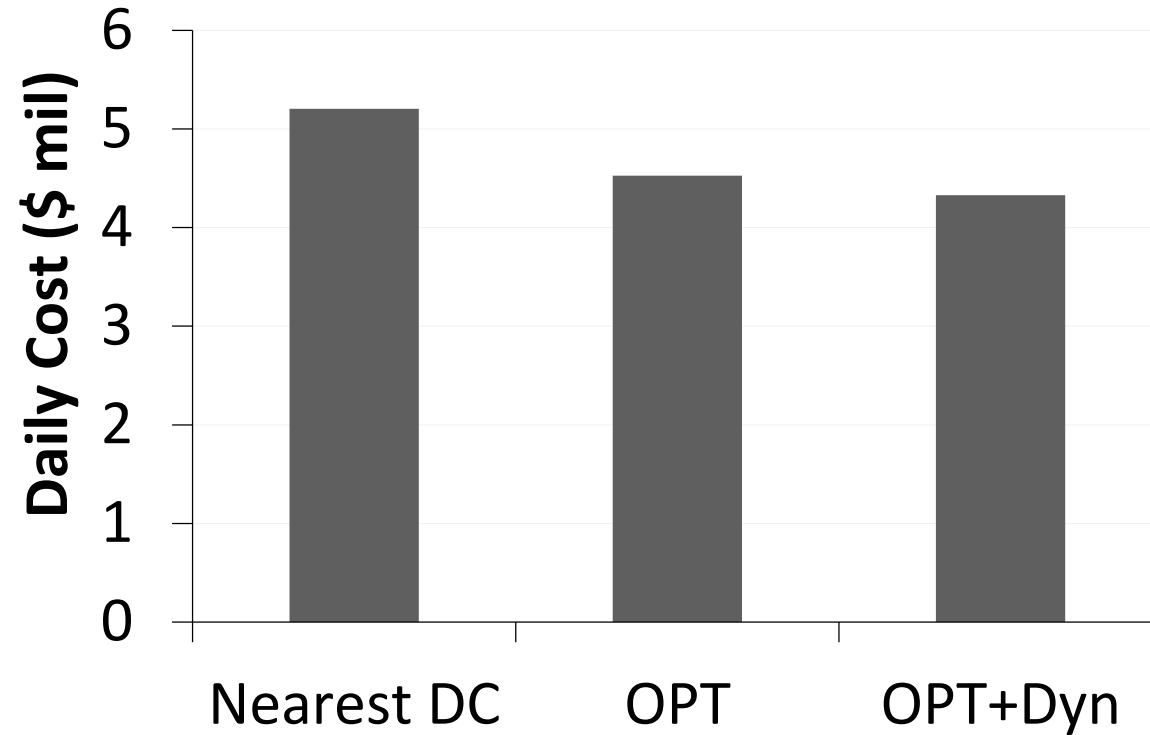
Geo-distribution for cost savings

$$\sum_{\boxed{j:l_{ij} \leq l}} \underbrace{f_{ij0}(t)}_{\text{Distribute across DCs}} \geq \underbrace{d_i(t)}_{\text{Client demand}}, \forall \underbrace{i, t}_{\text{Client index}}.$$

within the latency limit

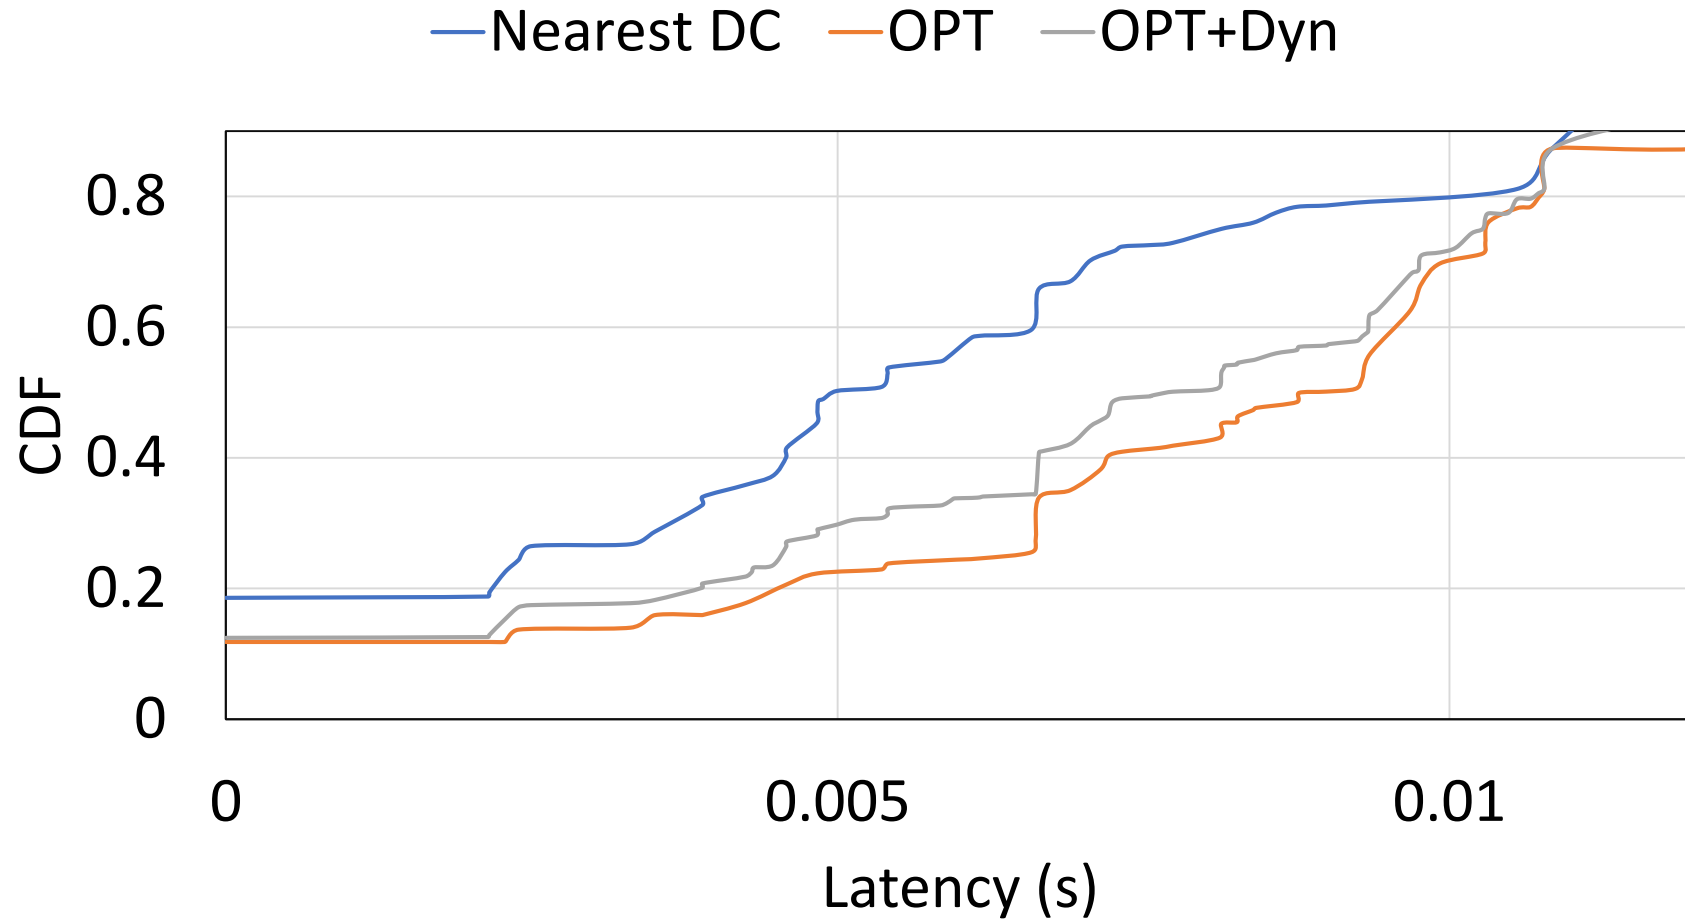
$$\underbrace{C_j}_{\text{DC capacity}} - \sum_i \underbrace{f_{ij0}(t)}_{\text{Assigned load from all clients}} \geq 0, \forall \underbrace{j, t}_{\text{DC index}}$$

Geo-distribution for cost savings



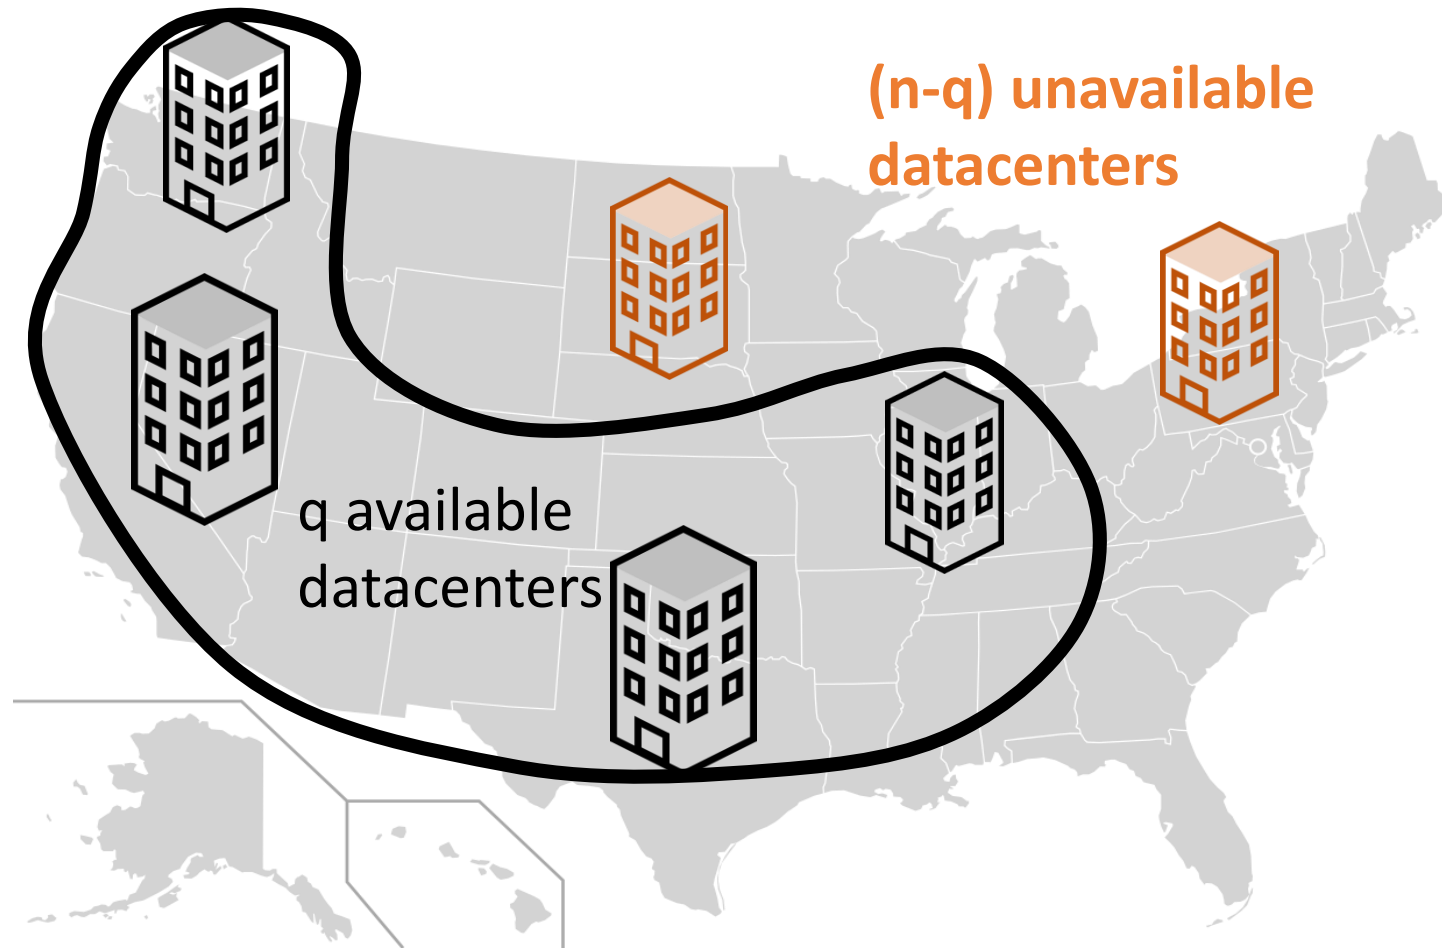
Geographical factors help reduce provisioned capacity by 16%

Geo-distribution for high availability

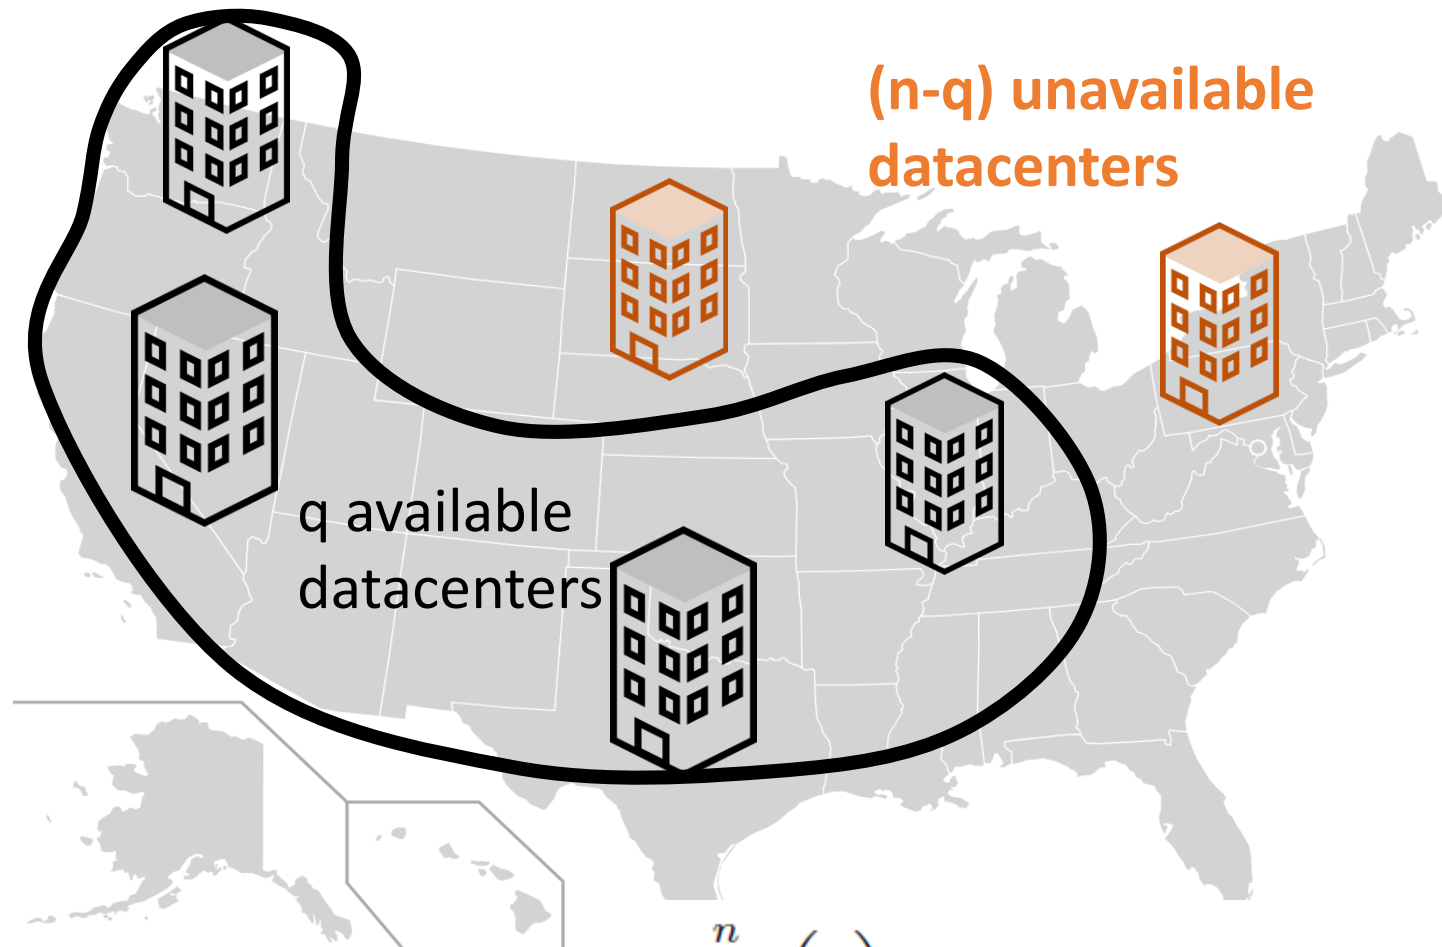


Capacity reduction by optimizing latency for a large fraction of clients

System model



System model



$$\text{System Availability} = \sum_{i=q}^n \binom{n}{i} a_d^i (1 - a_d)^{n-i}$$

Geo-distribution for high availability

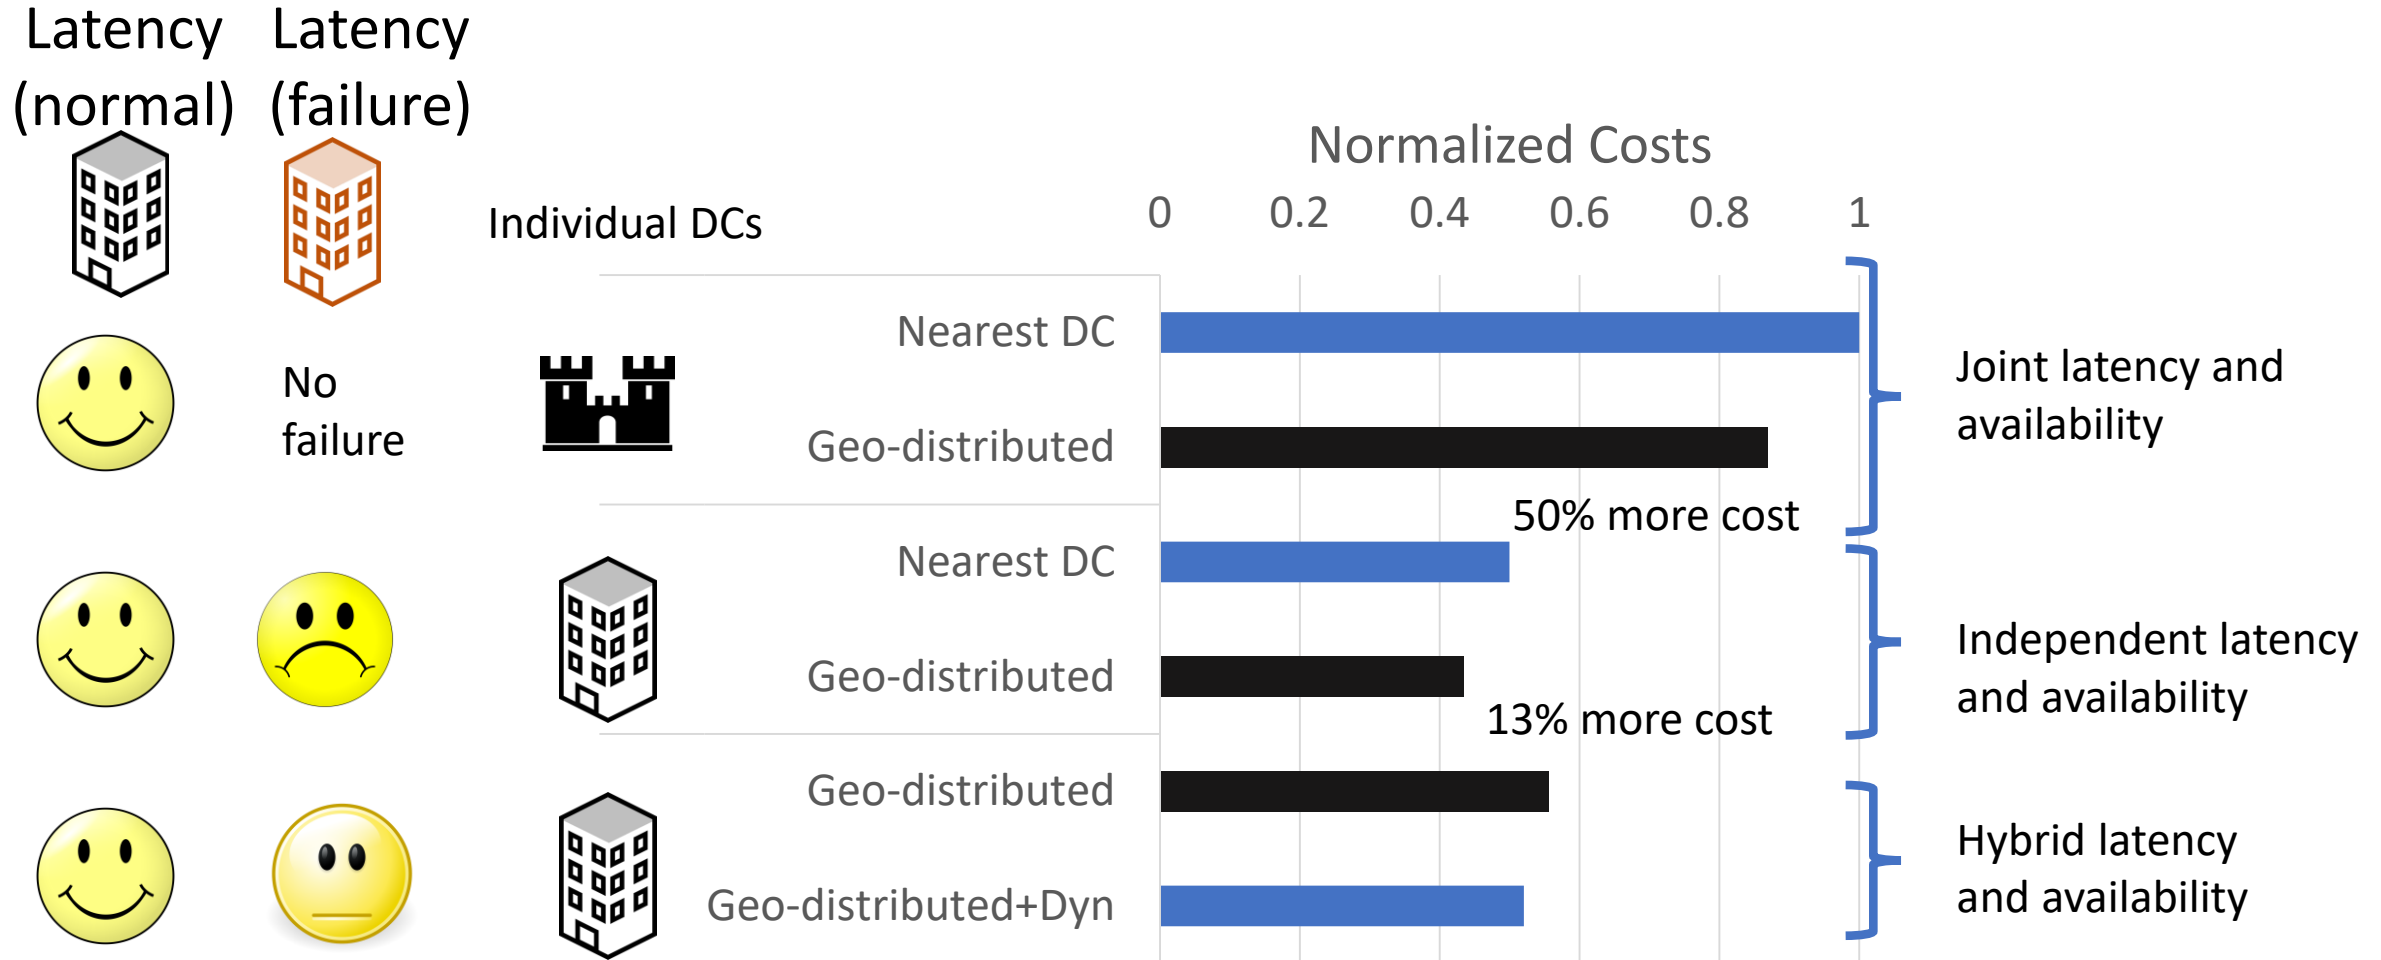
Latency guarantee even if a single DC has failed

$$\sum_{j: j \neq k, l_{ij} \leq l'} f_{ijk}(t) \geq \underbrace{d_i(t)}_{\text{Client demand}}, \forall k > 0, \underbrace{i}_{\text{Client index}}, t.$$

DCs within the latency limit (l') Distribute across all DCs except k

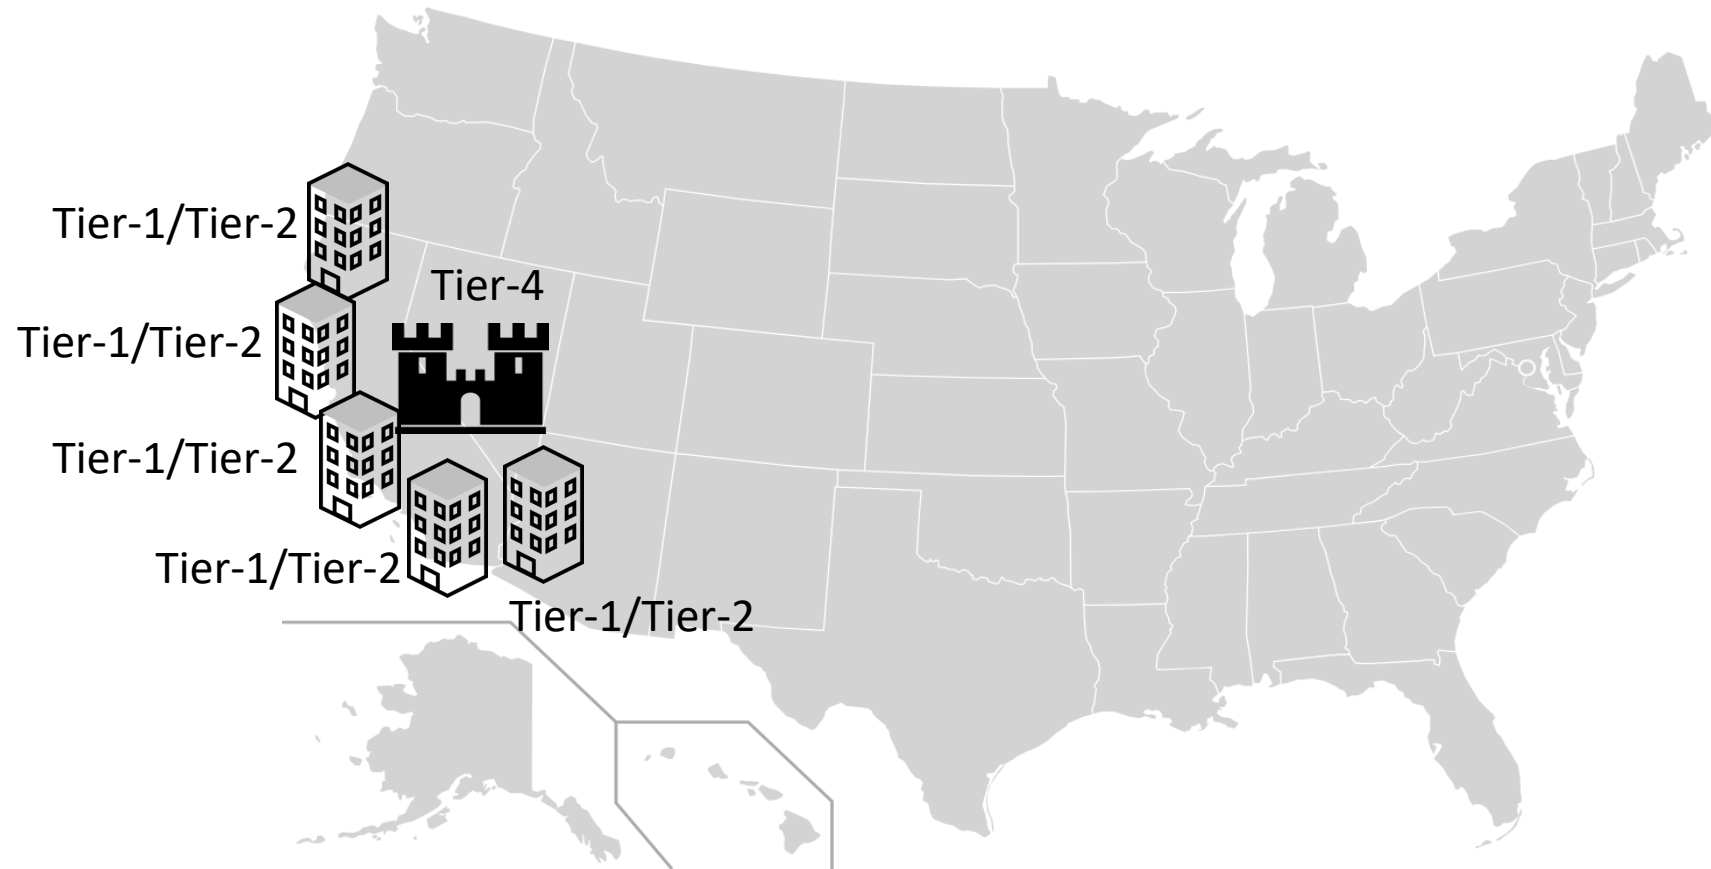
$$C_j - \sum_i f_{ijk}(t) \geq 0, \forall k (k > 0, k \neq j), j, t$$

Geo-distribution for high availability



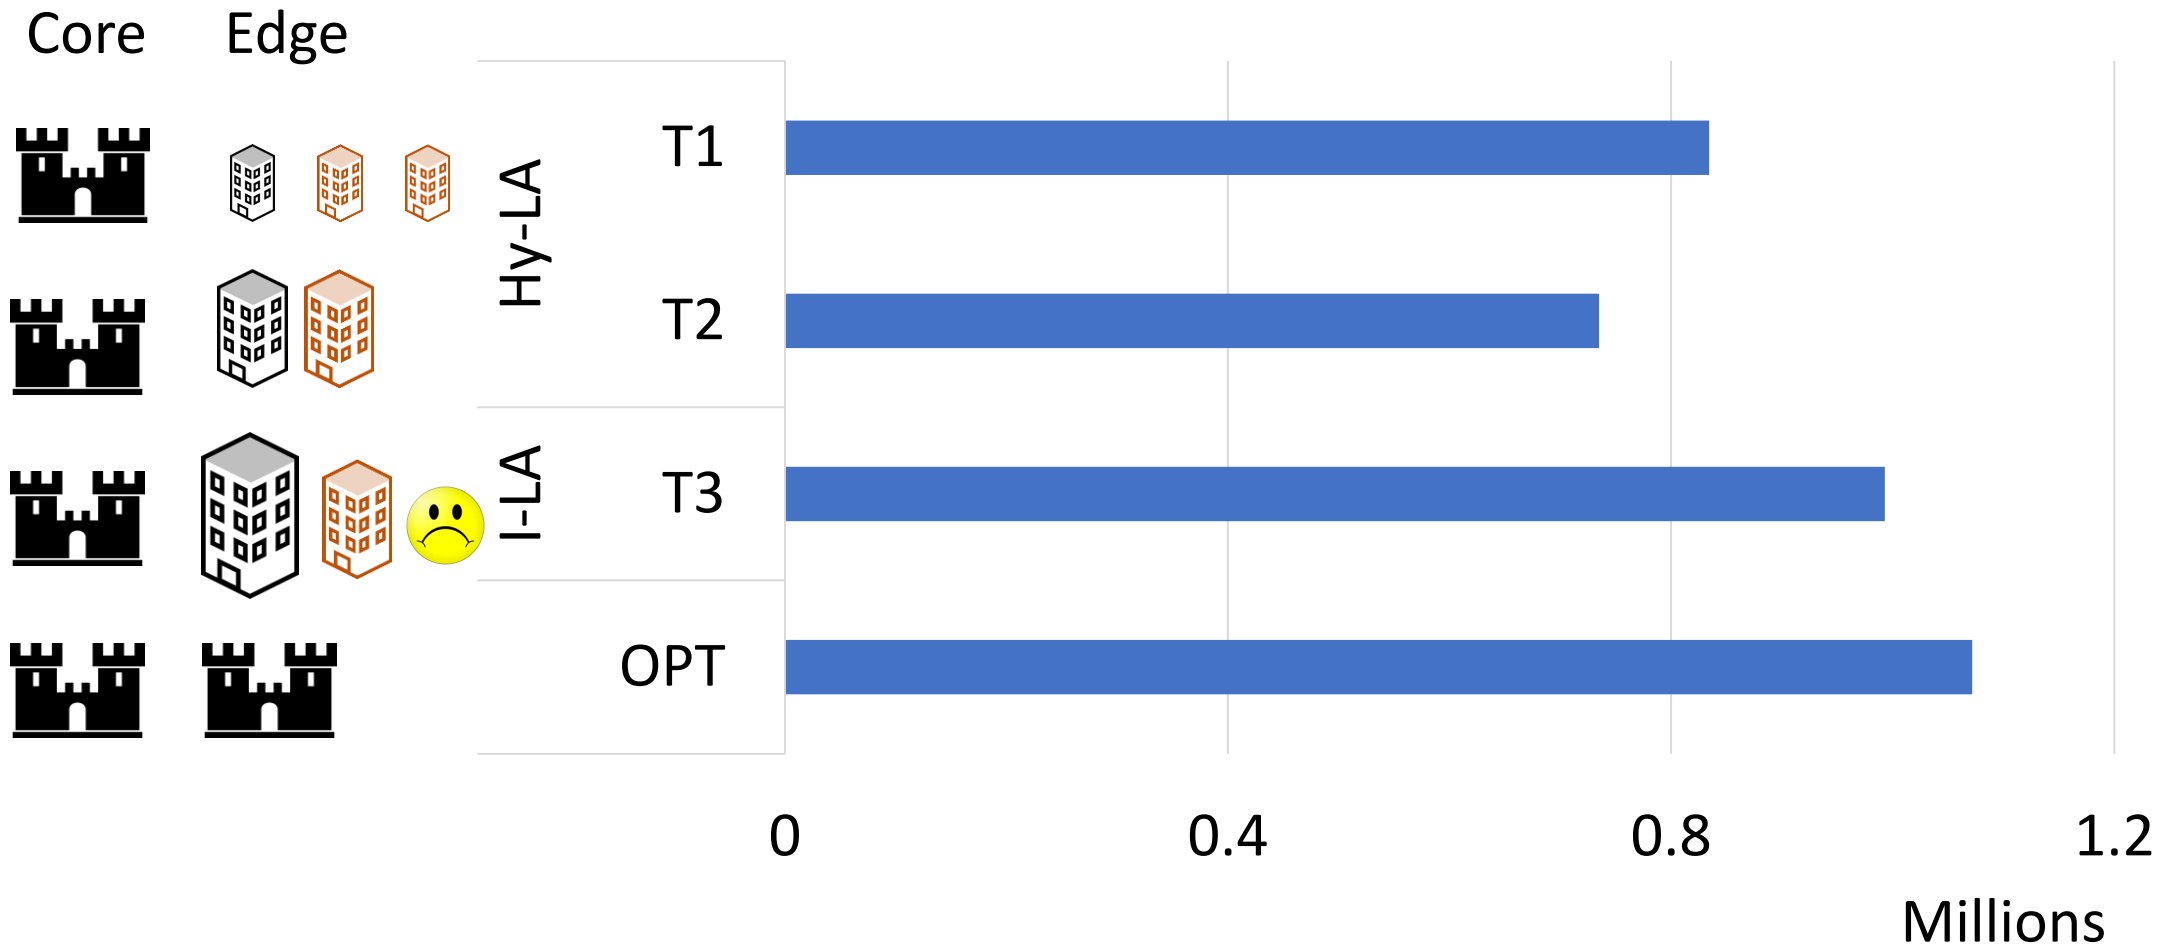
Geo-distribution helps to achieve latency and availability at low costs

Geo-distributed core and edge datacenters



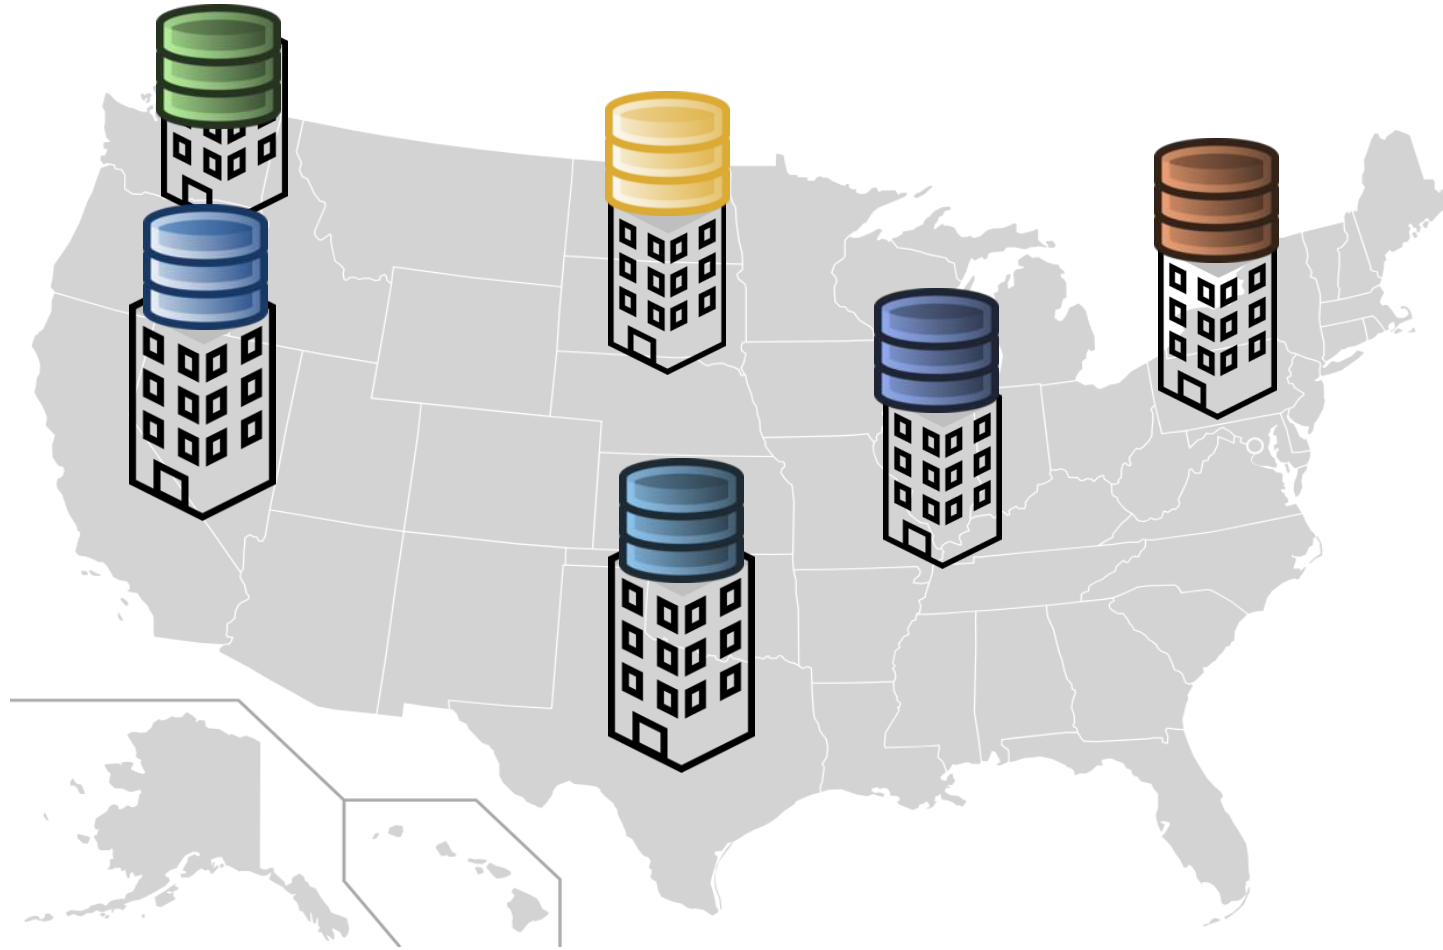
Extend the previous equation to handle multiple failures

Geo-distributed core and edge datacenters



Cost benefit from using low availability DC reversers

Application factors: Replicated data



Geo-distribution for availability

Replicate all
data before
failure

$$\left\{ \sum_{d: d \neq j, l_{ij} \leq l} r_{ijd}(t) \geq f_{ij0}(t), \forall i, j, t \right.$$

Synchronous
or
Asynchronous

Make sure that
enough replicas
are available
upon redirection
due to failures

$$\left\{ \sum_{j: j \neq d, l_{id} < l'} r_{ijd}(t) \geq \delta f_{idk}(t), \forall k > 0, k \neq j, i, d, t \right.$$

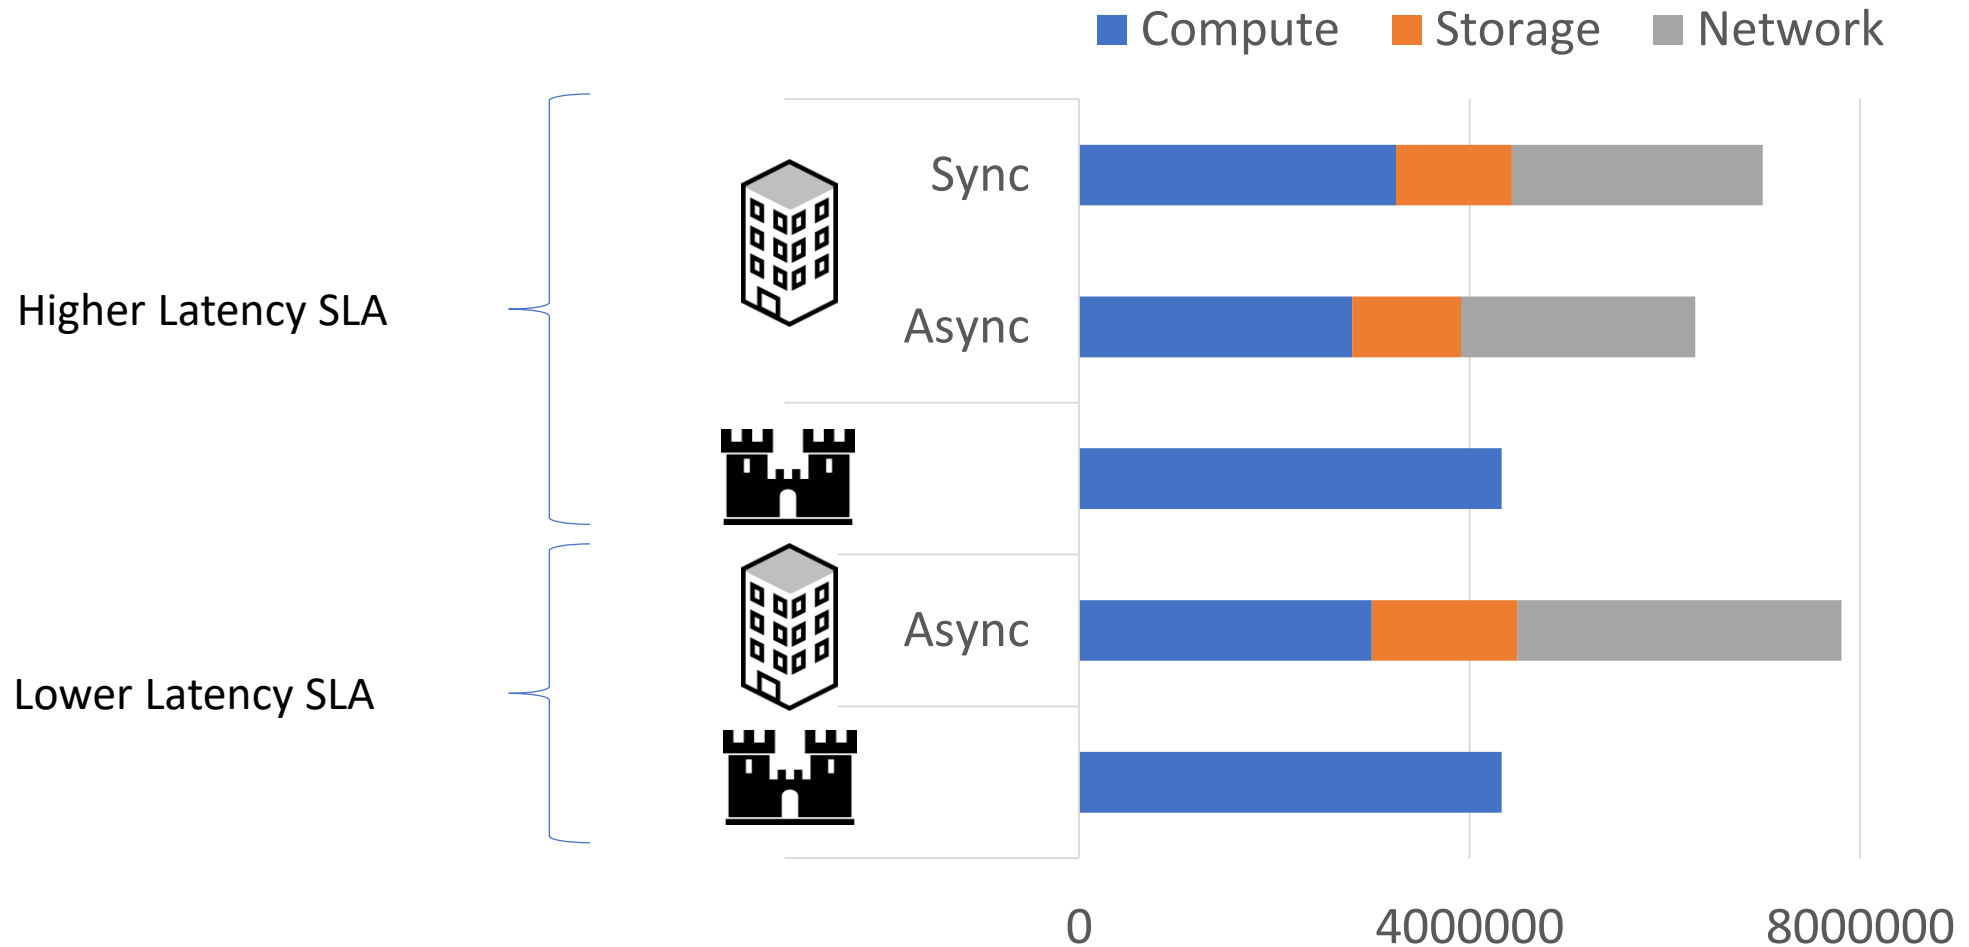
Geo-distribution for availability

$$\min \left[\mathcal{C} \times C_j + \overbrace{(\mathcal{S} \times \mathcal{SC}_{\mathcal{R}}) \sum_{i,j,k,t} r_{ikj}(t)}^{\text{Storage Costs}} + \overbrace{\mathcal{N} \sum_{j,k,t} N_{jk}(t)}^{\text{Network Costs}} \right]$$

$$N_{jk}(t) = \mathcal{NC}_{\mathcal{R}} \sum_i \underbrace{(r_{ijk}(t) + r_{ikj}(t))}_{\text{Replication between i and j}}, \forall t, j, k (j \neq k)$$

Replication between i and j

Geo-distribution for high availability

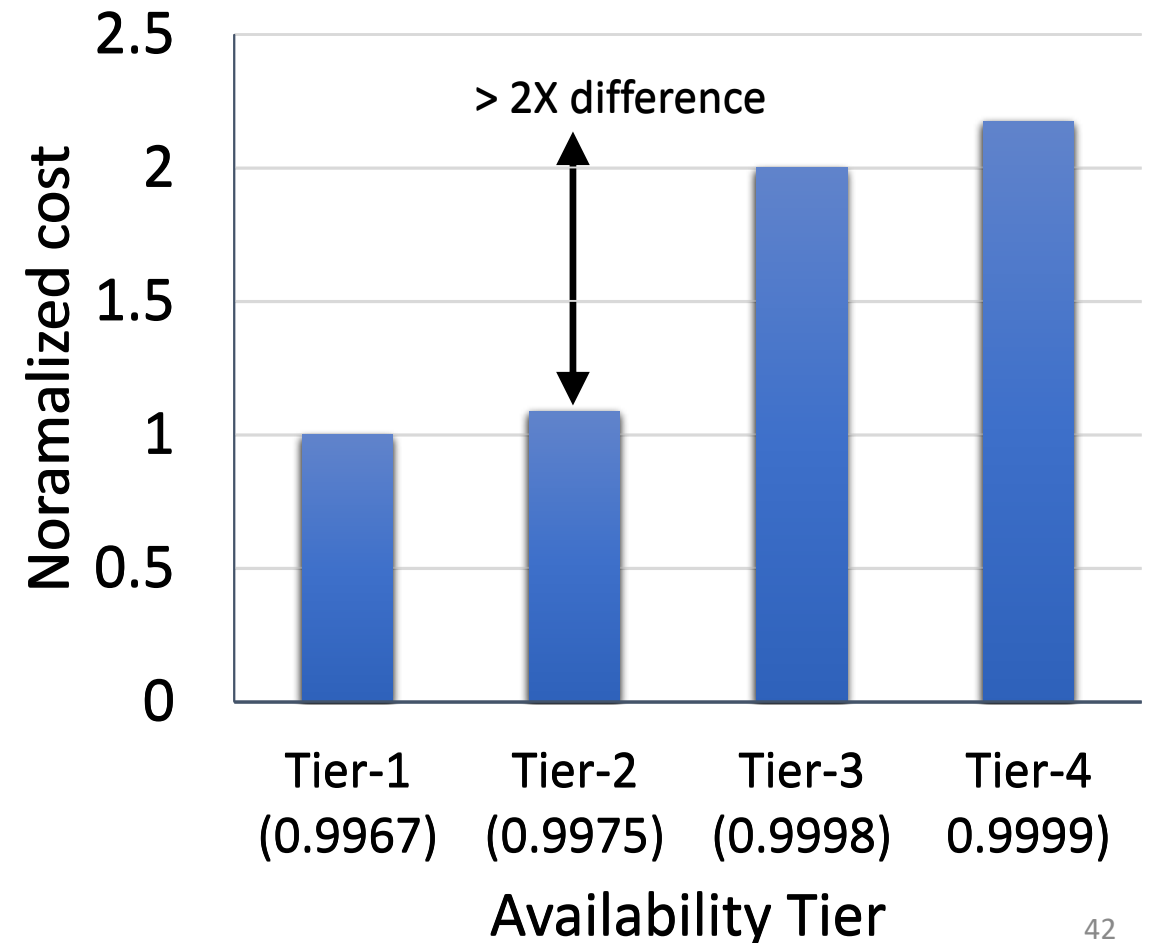
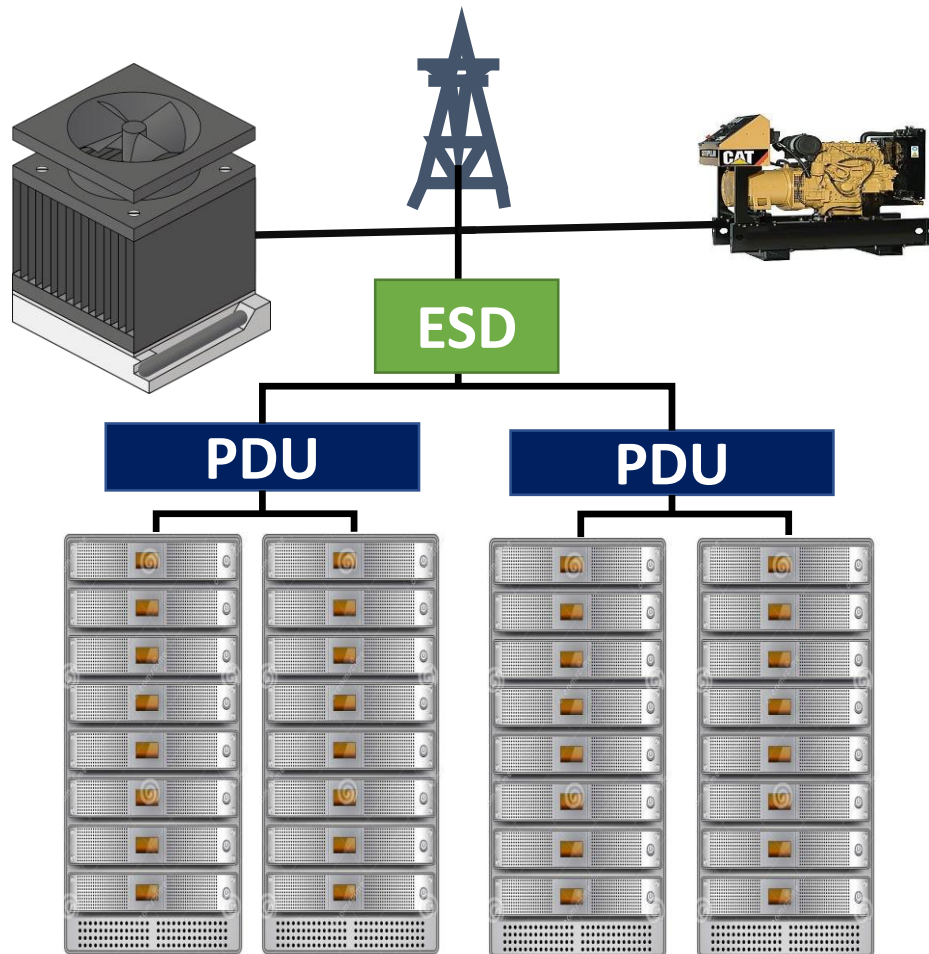


Replication costs outweigh the benefits

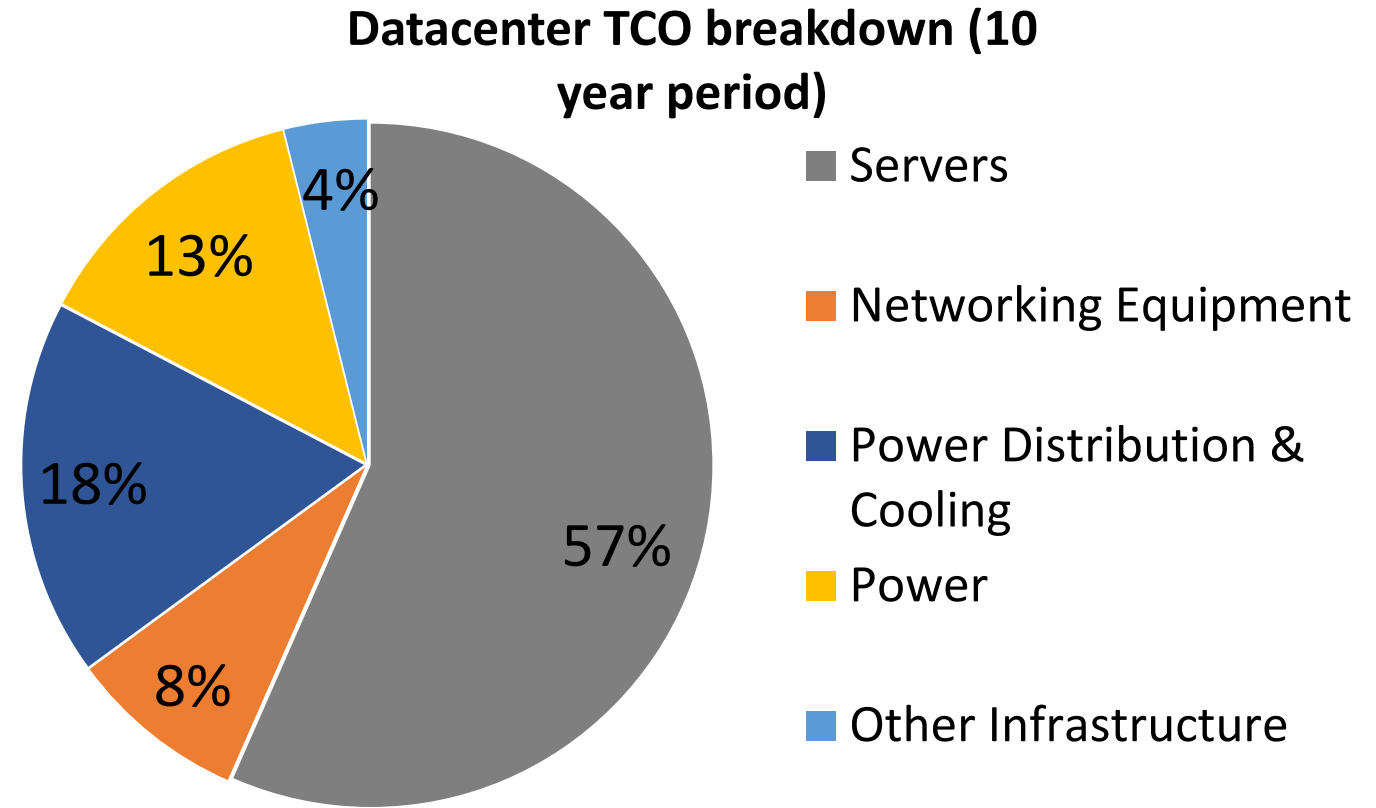
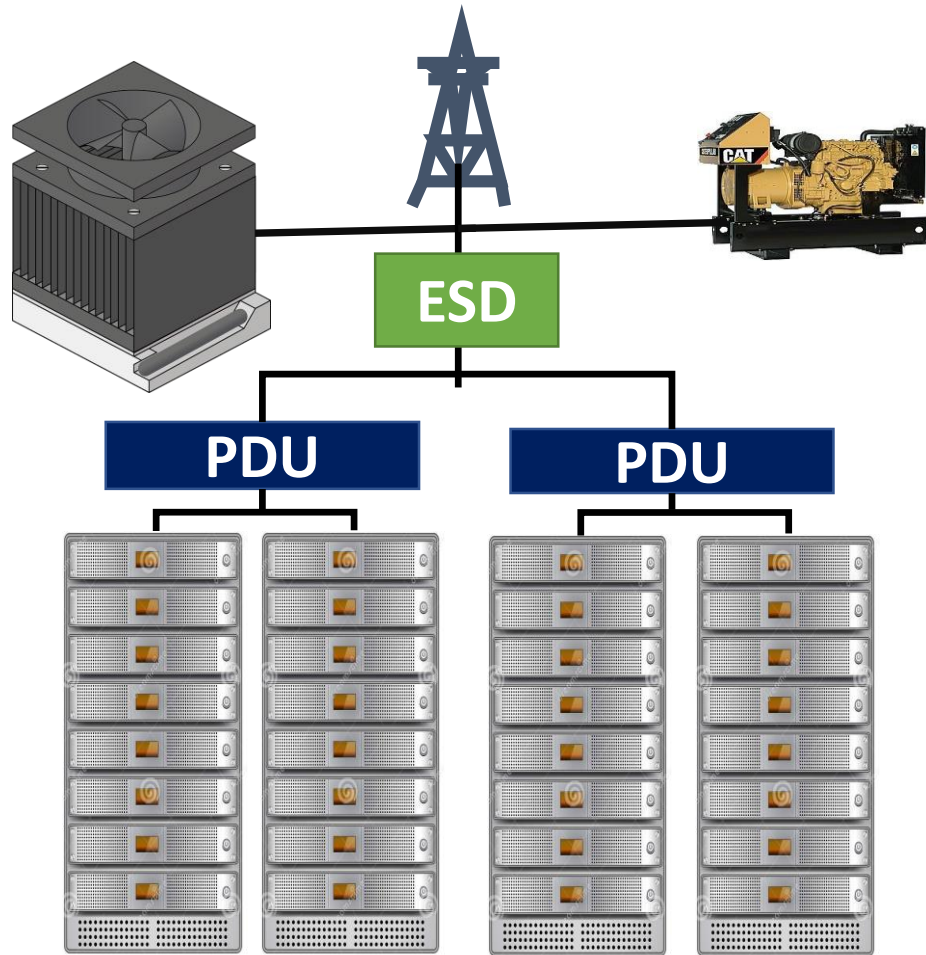
Summary

- Latency and availability are important for cloud apps
- Geo-distributed capacity planning
 - Geographical factors, Datacenter factors and Application factors
- Capacity optimization framework
 - Geo-distribution for lower costs
 - Geo-distribution for joint latency and availability
 - Geo-distributed replication aware capacity planning
 - More results in paper

Capacity planning – Important for operators



Capacity planning – Important for operators



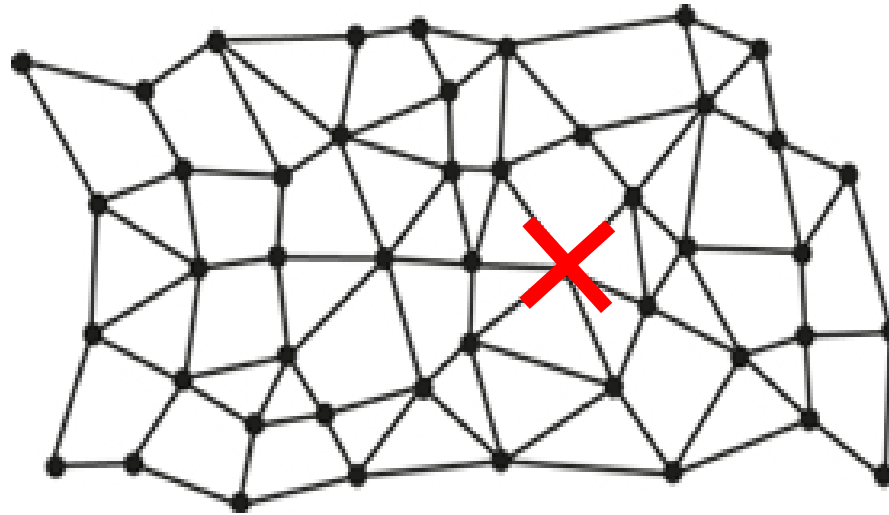
Source: Perspectives Datacenter Cost And Power

Availability is important for applications

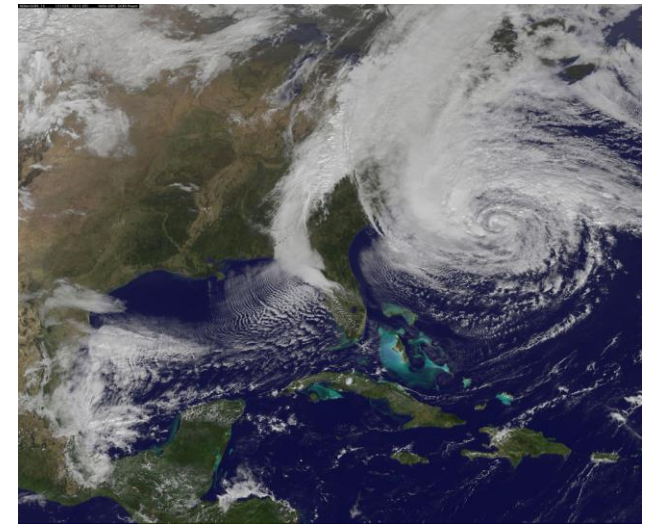
DCs are prone to failures



Power failures

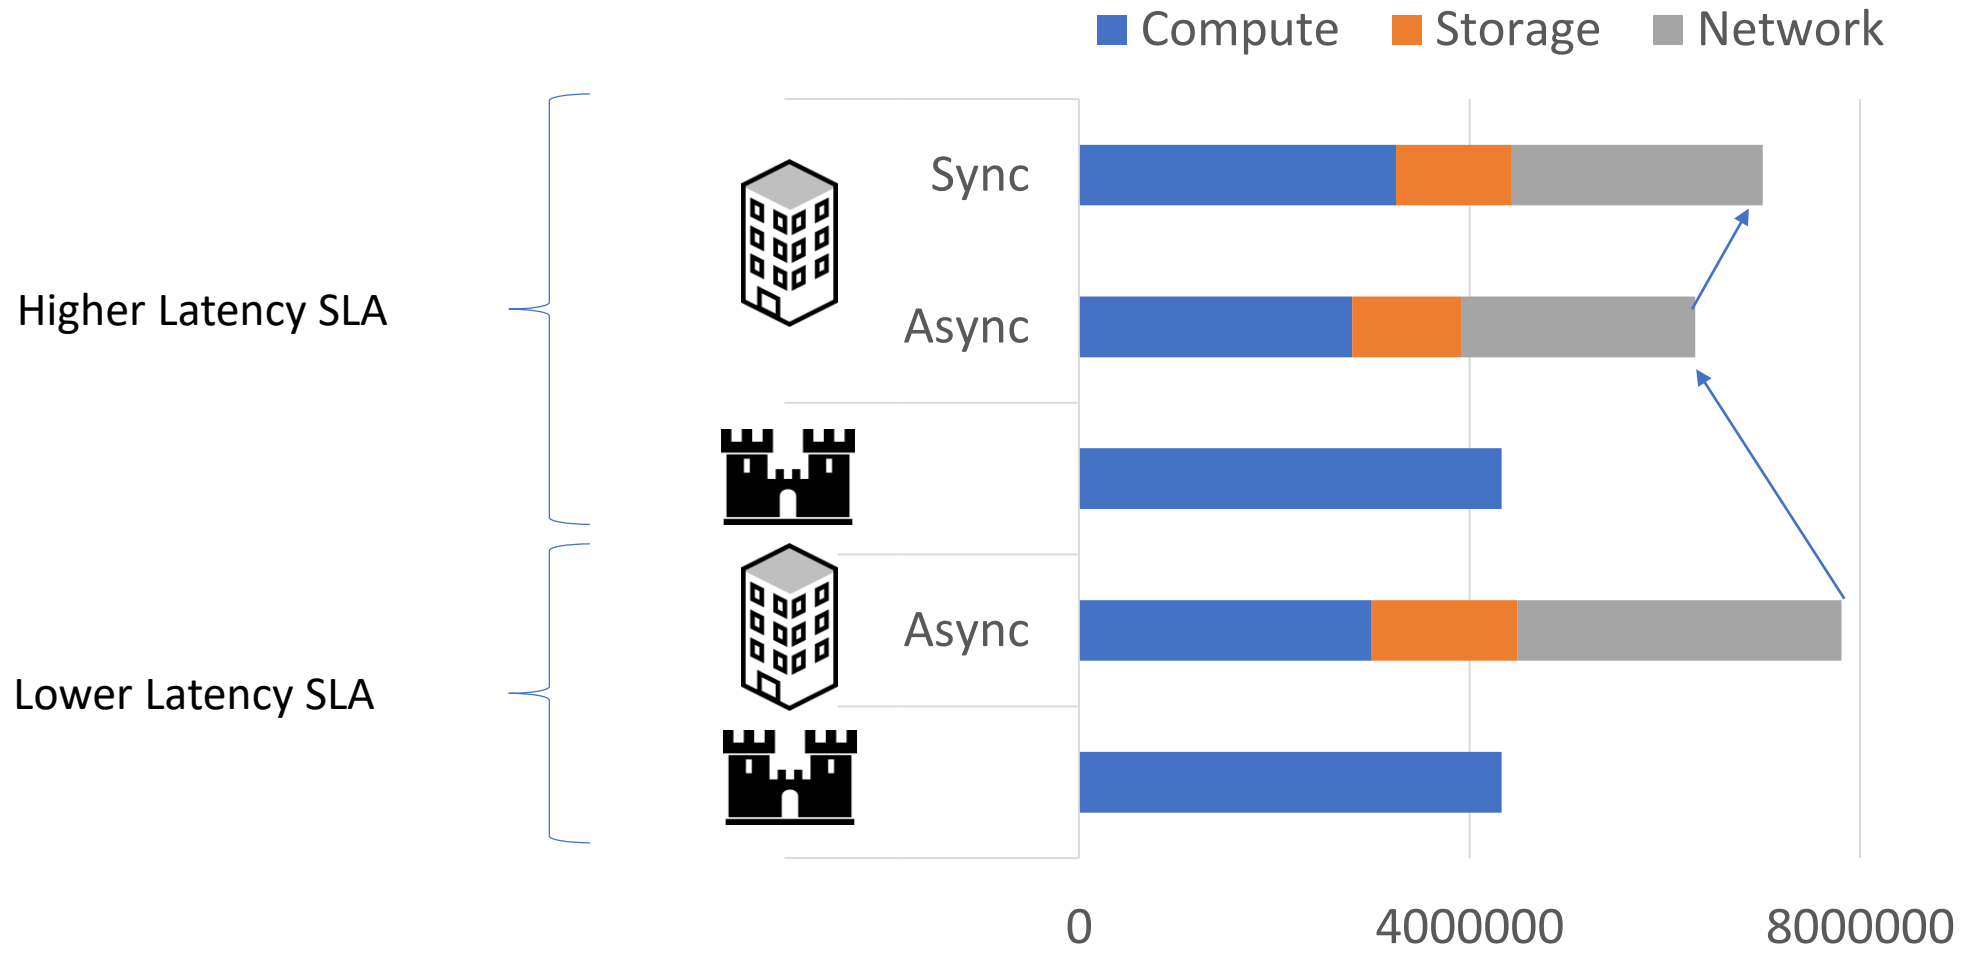


Network failures



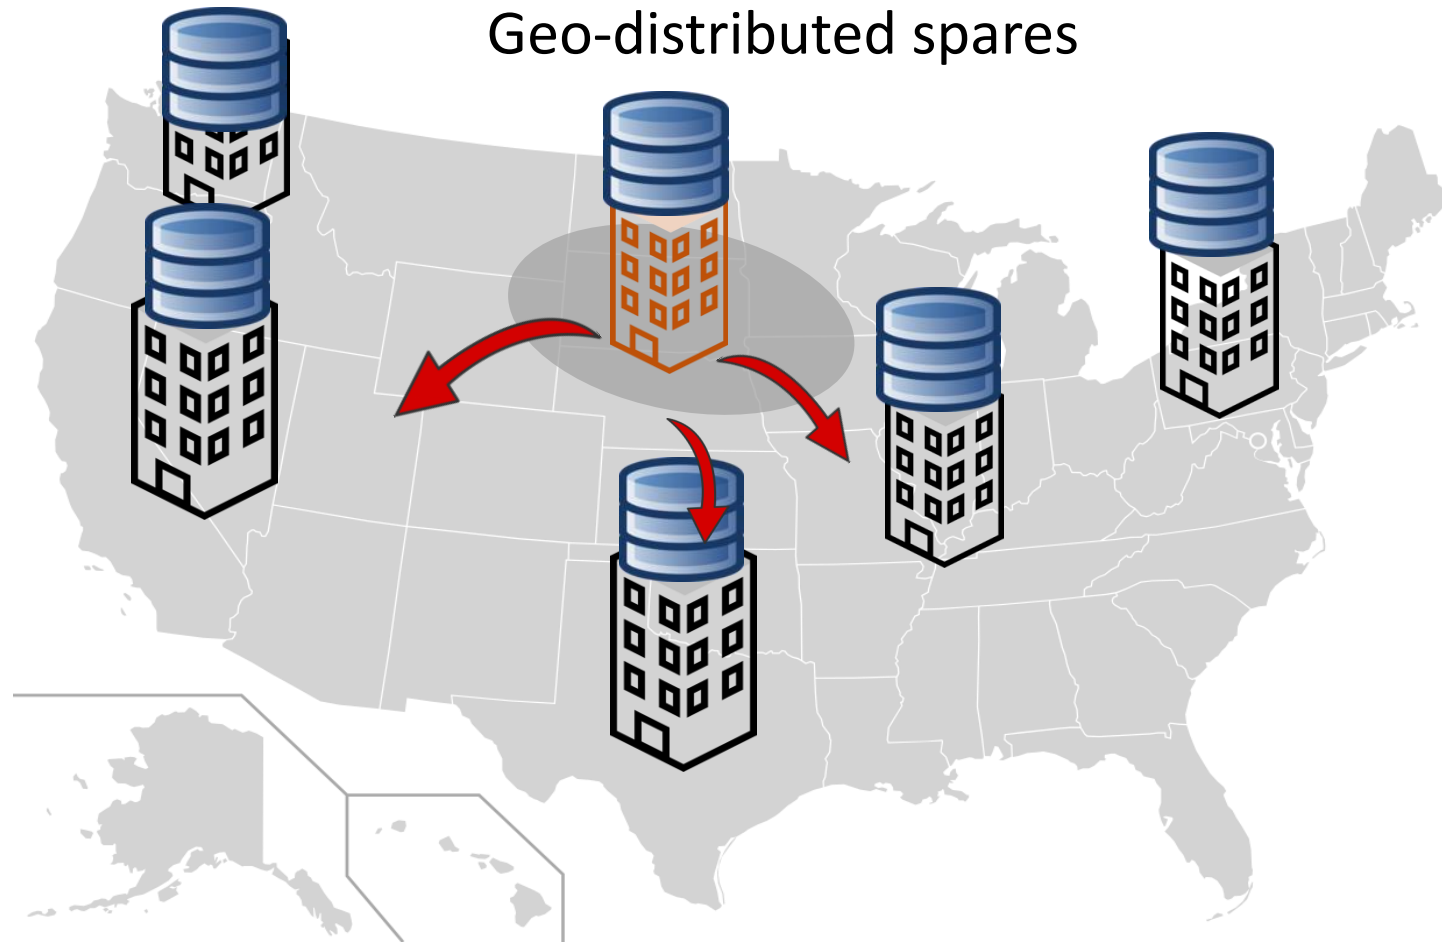
Natural calamities

Geo-distribution for high availability



Replication costs outweigh the benefits

Replication to use geo-distributed spares



Replication to use geo-distributed spares

