ArchitectNow

ABOUT PRESENTATION

Saving Your Ass(ets)

Azure Resiliency Planning

ArchitectNow Webinar

PRESENTED BY

Sean Whitesell

info@architectnow.net
www.architectnow.net

TRANSFORMATION
THROUGH
TECHNOLOGY

WELCOME TO
ARCHITECT NOW

Whether launching new Cloud or mobile apps or modernizing your legacy platforms we can help you identify the best options and work with you on bringing those ideas to life. To get the ball rolling, reach out and tell us a bit about your needs and we can start identifying solutions. There is no risk, and we can quickly get to the point of providing initial ideas along with rough estimates of the costs and implementation times required with various recommendations.

info@architectnow.net
www.architectnow.net

ArchitectNow

CONTACT INFORMATION

Sean Whitesell
Sr. Cloud Architect
ArchitectNow
swhitesell@architectnow.net

www.ArchitectNow.net
@architectnow
LinkedIn

ArchitectNow

Planning Types

Resiliency Planning vs Disaster Recovery Planning

Planning Types

Resiliency Planning vs Disaster Recovery Planning

Resiliency planning is about preparing for various types of events to mitigate or minimize disaster recovery efforts.

Planning Types

Resiliency Planning vs Disaster Recovery Planning

Resiliency planning is about preparing for various types of events to mitigate or minimize disaster recovery efforts.

Disaster Recovery planning is about procedures (efforts) to be executed after a disastrous event has occurred.
-- restoring from backup

Planning Types

Resiliency Planning vs Disaster Recovery Planning

Resiliency planning is about preparing for various types of events to mitigate or minimize disaster recovery efforts.
-- *making sure backups are being done and actually have data*

Disaster Recovery planning is about procedures to be executed after a disastrous event has occurred.
-- restoring from backup

Planning Types

Don’t let this be your plan.

Horror Stories

What is the cost of downtime?

Cost of missed deadlines?

How much trust can you afford to lose?

Order Fulfillment

Data Backups and Security

Order Processing

Remote Work

Customer Support

eCommerce Sites

Communication Systems

Financial Transactions

Manufacturing and Production

Inventory Management

Regulatory Compliance

HR and Payroll

Horror Stories

Salesforce goes down: On May 9, 2016, the Silicon Valley NA14 instance of Salesforce.com went offline, resulting in an outage that lasted for more than 24 hours. Extensive business damage was inevitable, with customers losing hours and hours of data. Salesforce moved to Amazon Web Services for most of its workloads thereafter as a result.

https://www.spiceworks.com/tech/cloud/guest-article/6-cloud-computing-failures-that-shocked-the-world/

Horror Stories

Salesforce goes down: On May 9, 2016, the Silicon Valley NA14 instance of Salesforce.com went offline, resulting in an outage that lasted for more than 24 hours. Extensive business damage was inevitable, with customers losing hours and hours of data. Salesforce moved to Amazon Web Services for most of its workloads thereafter as a result.

A bad Christmas for Netflix: It was Christmas Eve in 2012, a time of cheer and uninterrupted entertainment that families looked forward to. However, AWS's Elastic Load Balancing service went awry, resulting in Netflix downtime. The aftermath was a whole bunch of disgruntled customers who were depending on the streaming service for a good Christmas. As if this souring relationship between Netflix and AWS was not enough, two years later Netflix rebooted 218 of its production nodes during an AWS update, and 22 failed to reboot—an additional instance of differences between AWS and Netflix.

https://www.spiceworks.com/tech/cloud/guest-article/6-cloud-computing-failures-that-shocked-the-world/

Horror Stories

On Oct 4th, 2021, Meta had an outage that effected Facebook, Instagram, WhatsApp, and several others.

Their routers had an issue involving the Border Gateway Protocol (BGP). This meant no network traffic was able to go in or out of their data center. The main issue was affecting their DNS servers. This outage lasted 6 to 7 hours.

Horror Stories

On Oct 4th, 2021, Meta had an outage that effected Facebook, Instagram, WhatsApp, and several others.

Their routers had an issue involving the Border Gateway Protocol (BGP). This meant no network traffic was able to go in or out of their data center. The main issue was affecting their DNS servers. This outage lasted 6 to 7 hours.

To compound the issue further, Meta’s physical security system of the data center was also on that DNS system. No ID badges worked on the doors. Supposedly someone was able to get a door off the hinges to gain access to reboot the routers.

What constitutes as a disaster?

What constitutes as a disaster?

vecteezy.com

A disaster is a condition in which system(s) and/or business processes are either performing poorly or not at all available due to an event.

A disaster does not have to only be a full-blown regional outage. It’s about the criticality of the systems and related business processes. It’s possible a single business process has such critical importance that if it was not available or not correctly handling data, then millions of dollars could be at risk.

What constitutes as a disaster?

“Resiliency is the ability of a system to gracefully handle and recover from failures, both inadvertent and malicious.” – Azure documentation

https://learn.microsoft.com/en-us/azure/well-architected/resiliency/reliability-patterns#resiliency

How much data can you afford to lose?

DATA LOSS PREVENTION

vecteezy.com

How much data can you afford to lose?

What about the times when data storage is replicating, and it fails? Perhaps there was a network cut (network partition) between two or more systems.

Other chances of data loss include;

- Someone changing a network policy/rule accidently preventing access to a datastore.
- Introduction of bug in code
- Regional outage

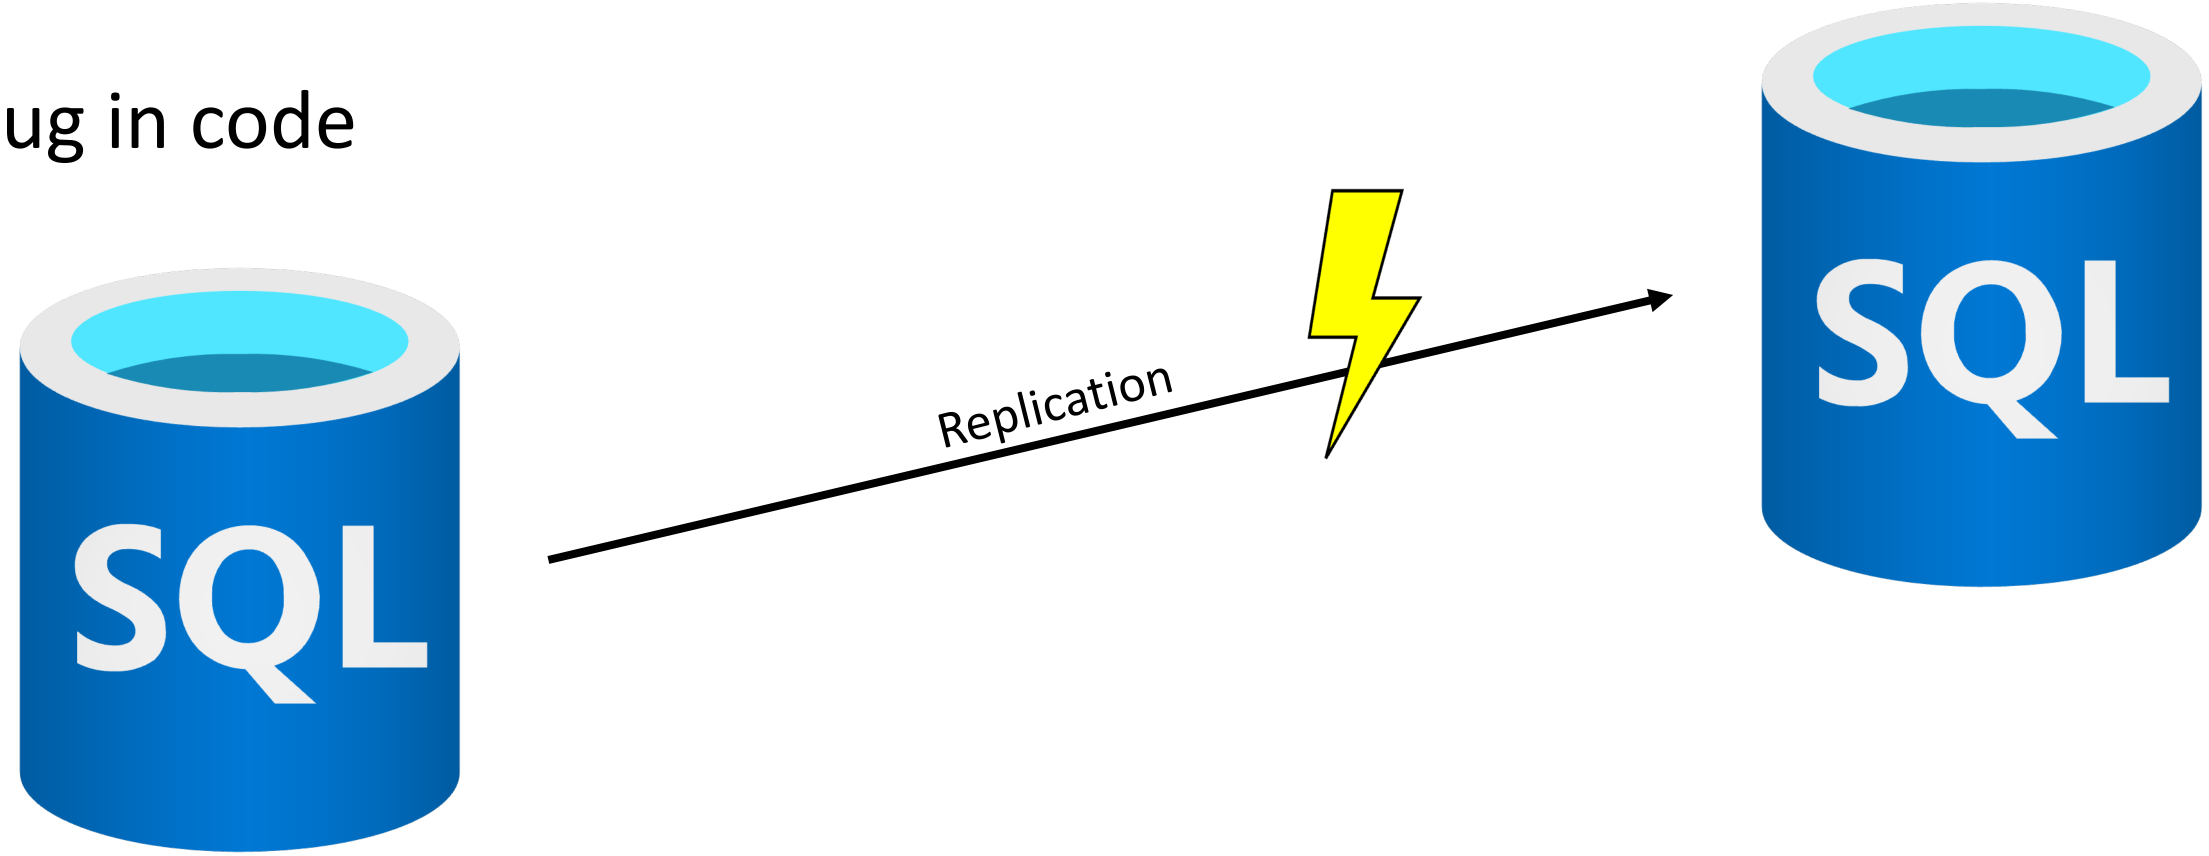

RPO and RTO

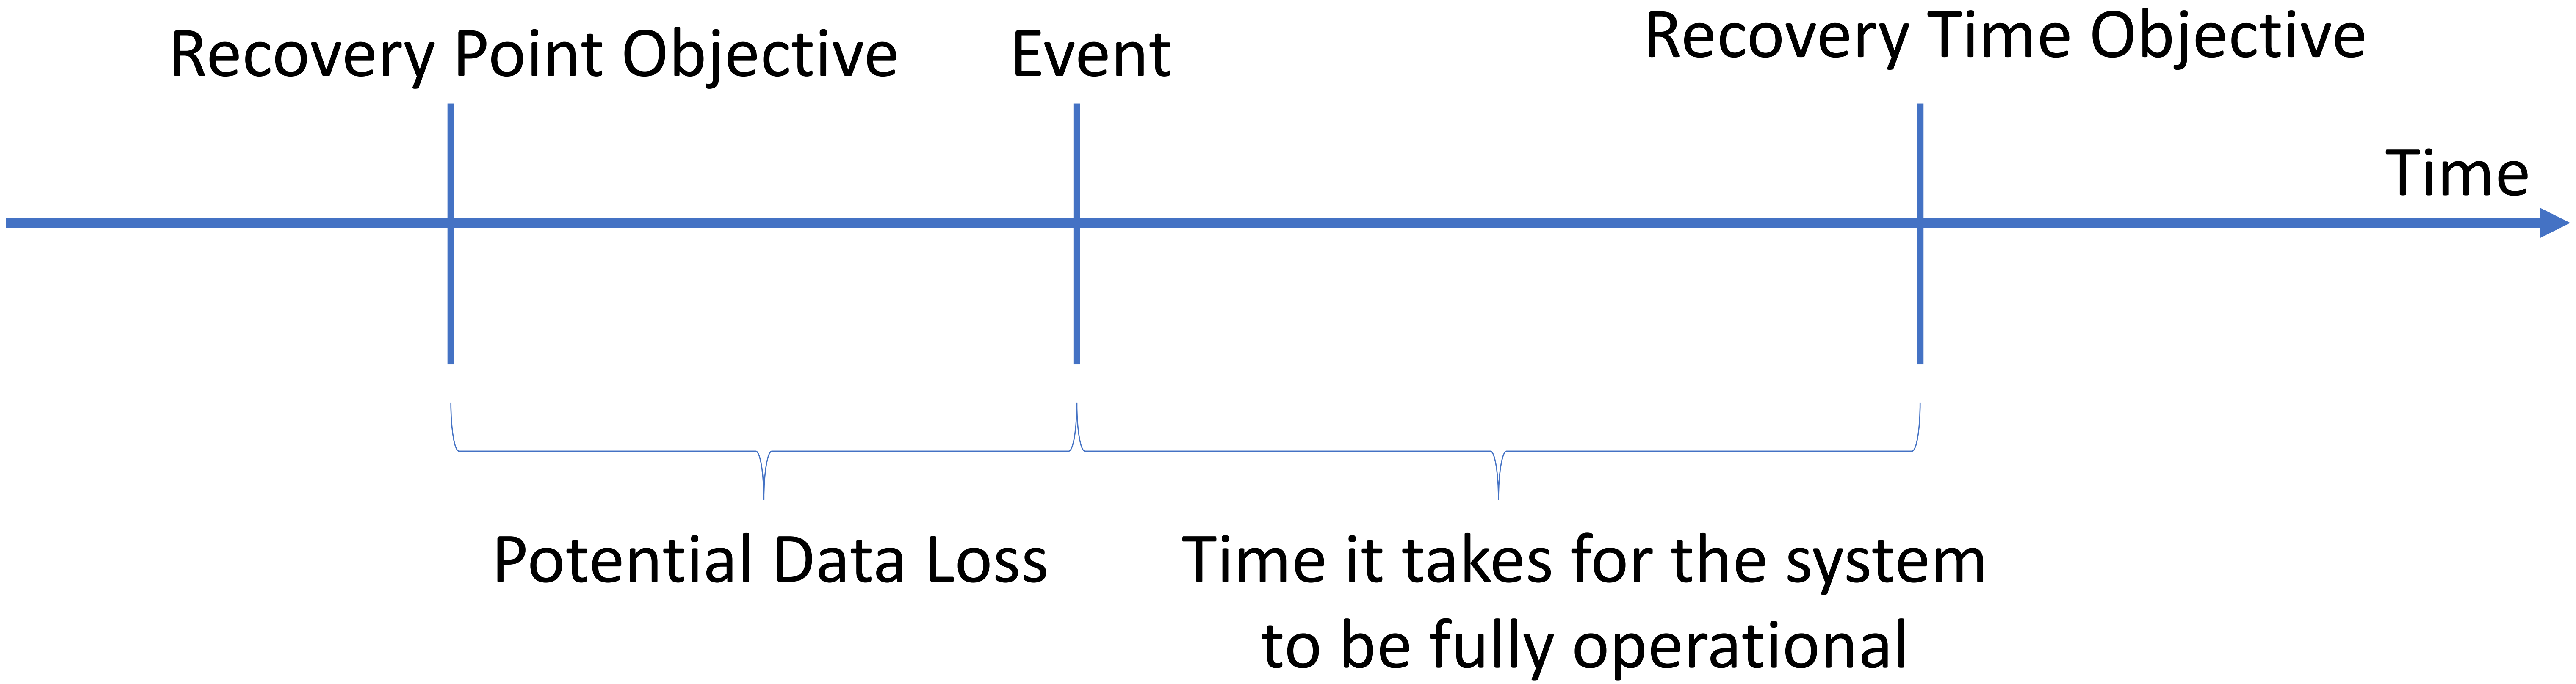
Recovery Point Objective
Event
Recovery Time Objective
Time
Potential Data Loss
Time it takes for the system to be fully operational

Service Level Objectives

also known as - 9's of Availability

31,556,952 seconds in a year

	99%	99.9%	99.99%	99.999%	
31,556,952	31,241,382.48	31,525,395.05	31,553,796.30	31,556,636.43	Seconds of Availability
	315,569.52	31,556.95	3,155.70	315.57	Seconds of Downtime
	5,259.49	525.95	52.59	5.26	Minutes of Downtime
	87.66	8.77	0.88	0.09	Hours of Downtime
	3.65	0.37	-	-	Days of Downtime

SLO applies to each application, server, router, dependent 3rd party systems, etc.

Where's the weakest link(s)?

Consistency and Availability

How should your applications behave when the network fails?

Consistency and Availability

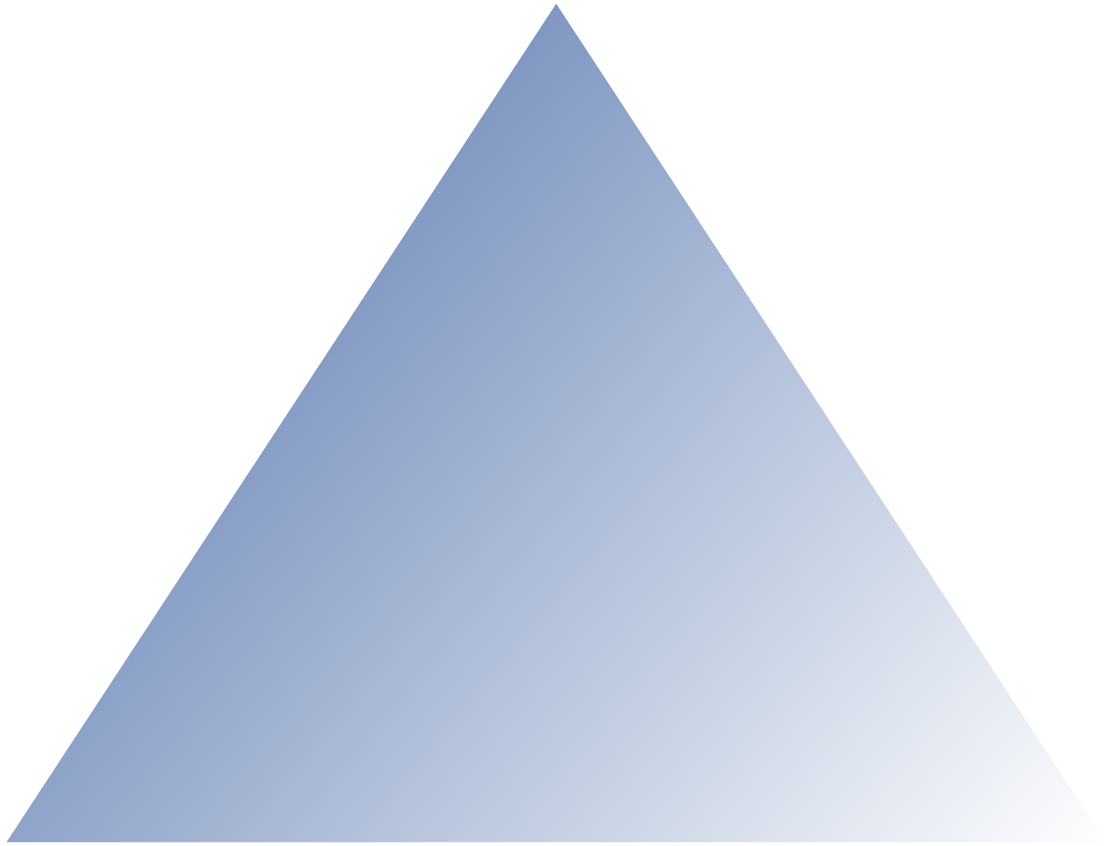

(network) Partition

Wordscapes

Photo credit: https://unsplash.com/photos/WargGLQW_Yk

Consistency and Availability

Availability

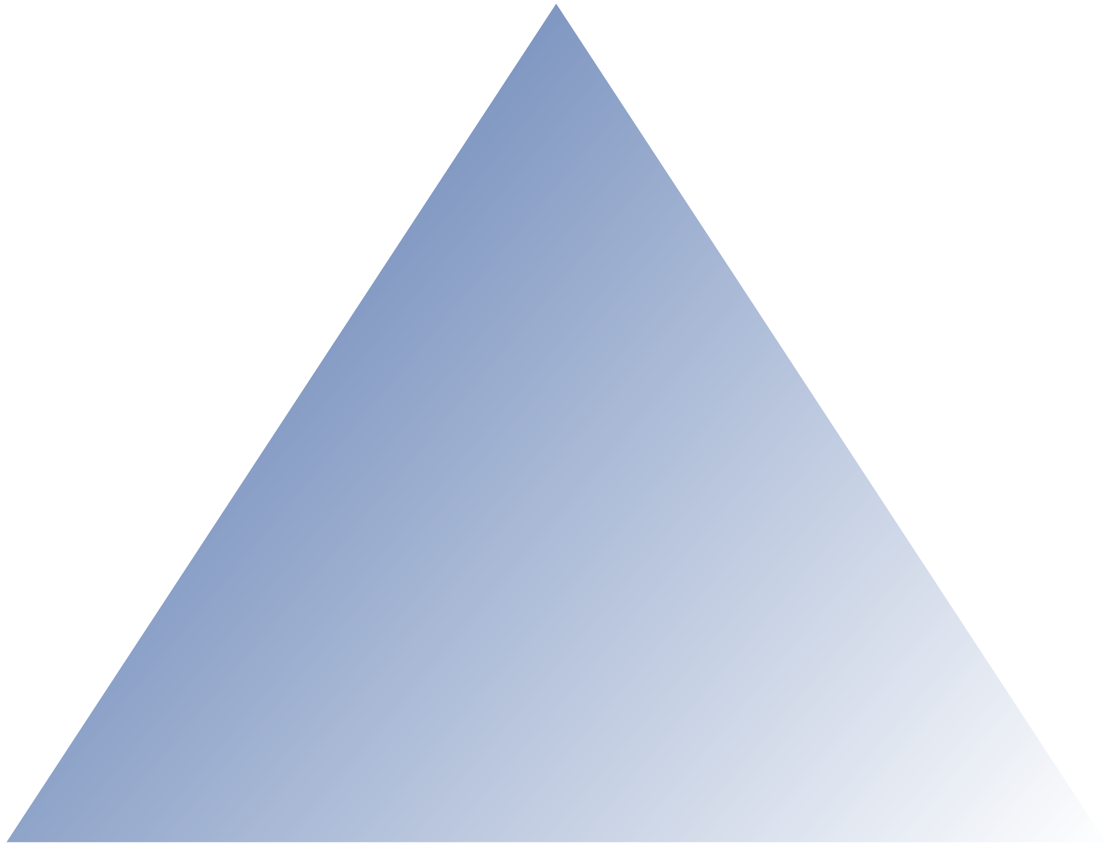

(network) Partition

Wordscapes

Photo credit: https://unsplash.com/photos/WargGLQW_Yk

Consistency and Availability

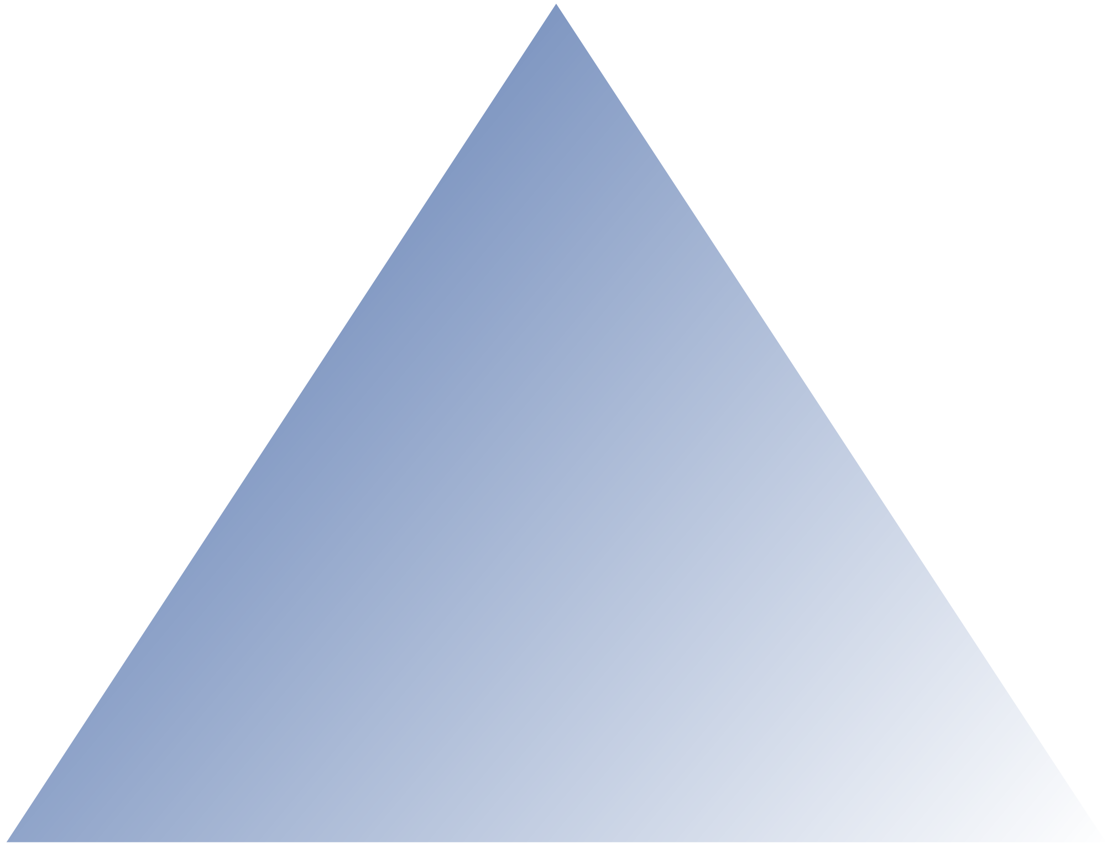

Consistency

(network) Partition

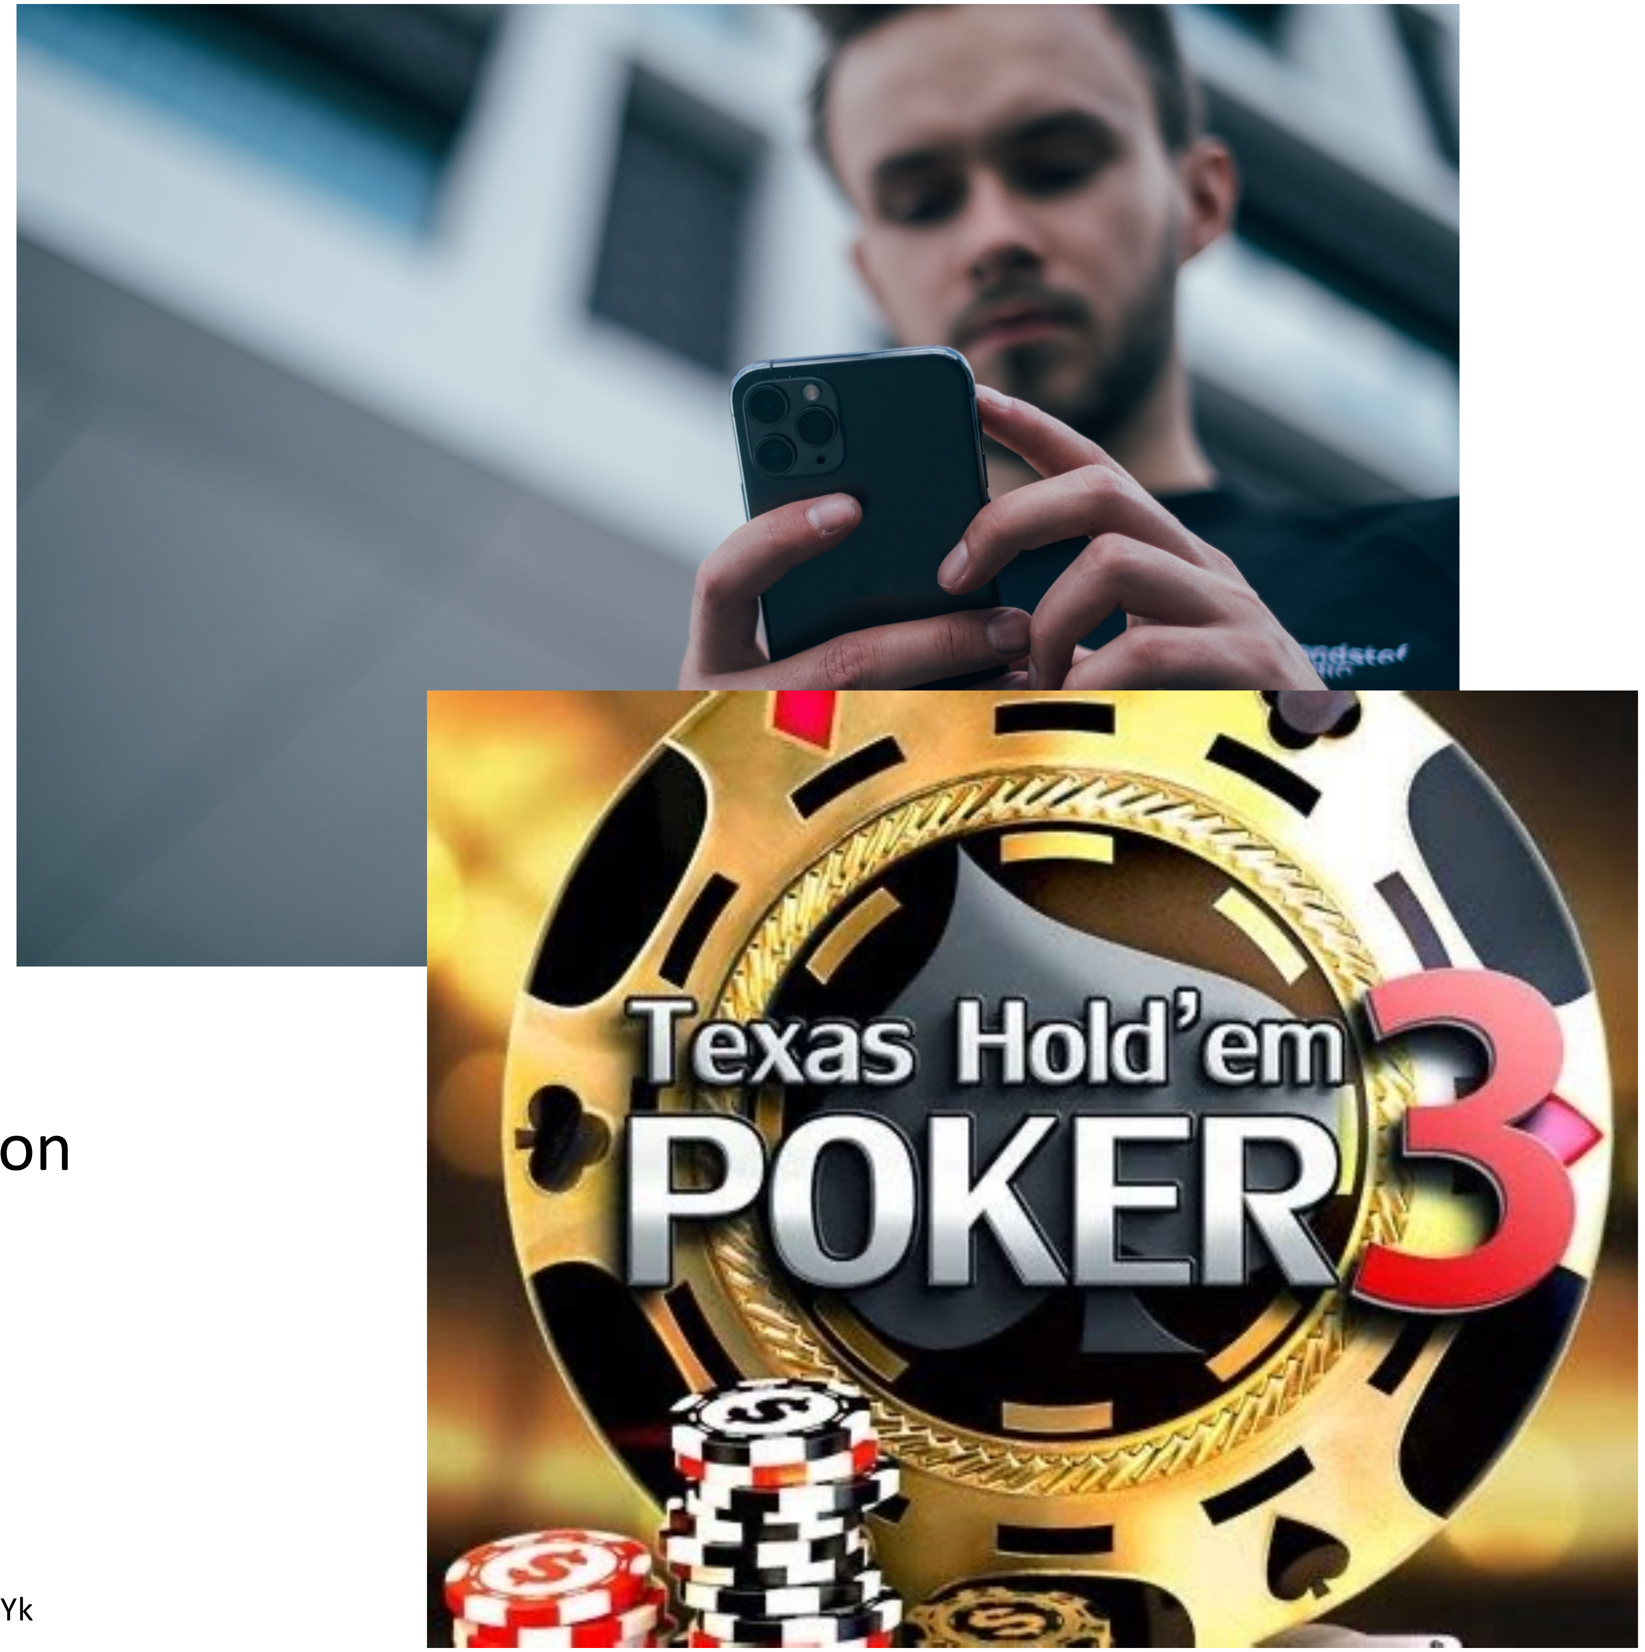

Photo credit: https://unsplash.com/photos/WargGLQW_Yk

CAP Theorem

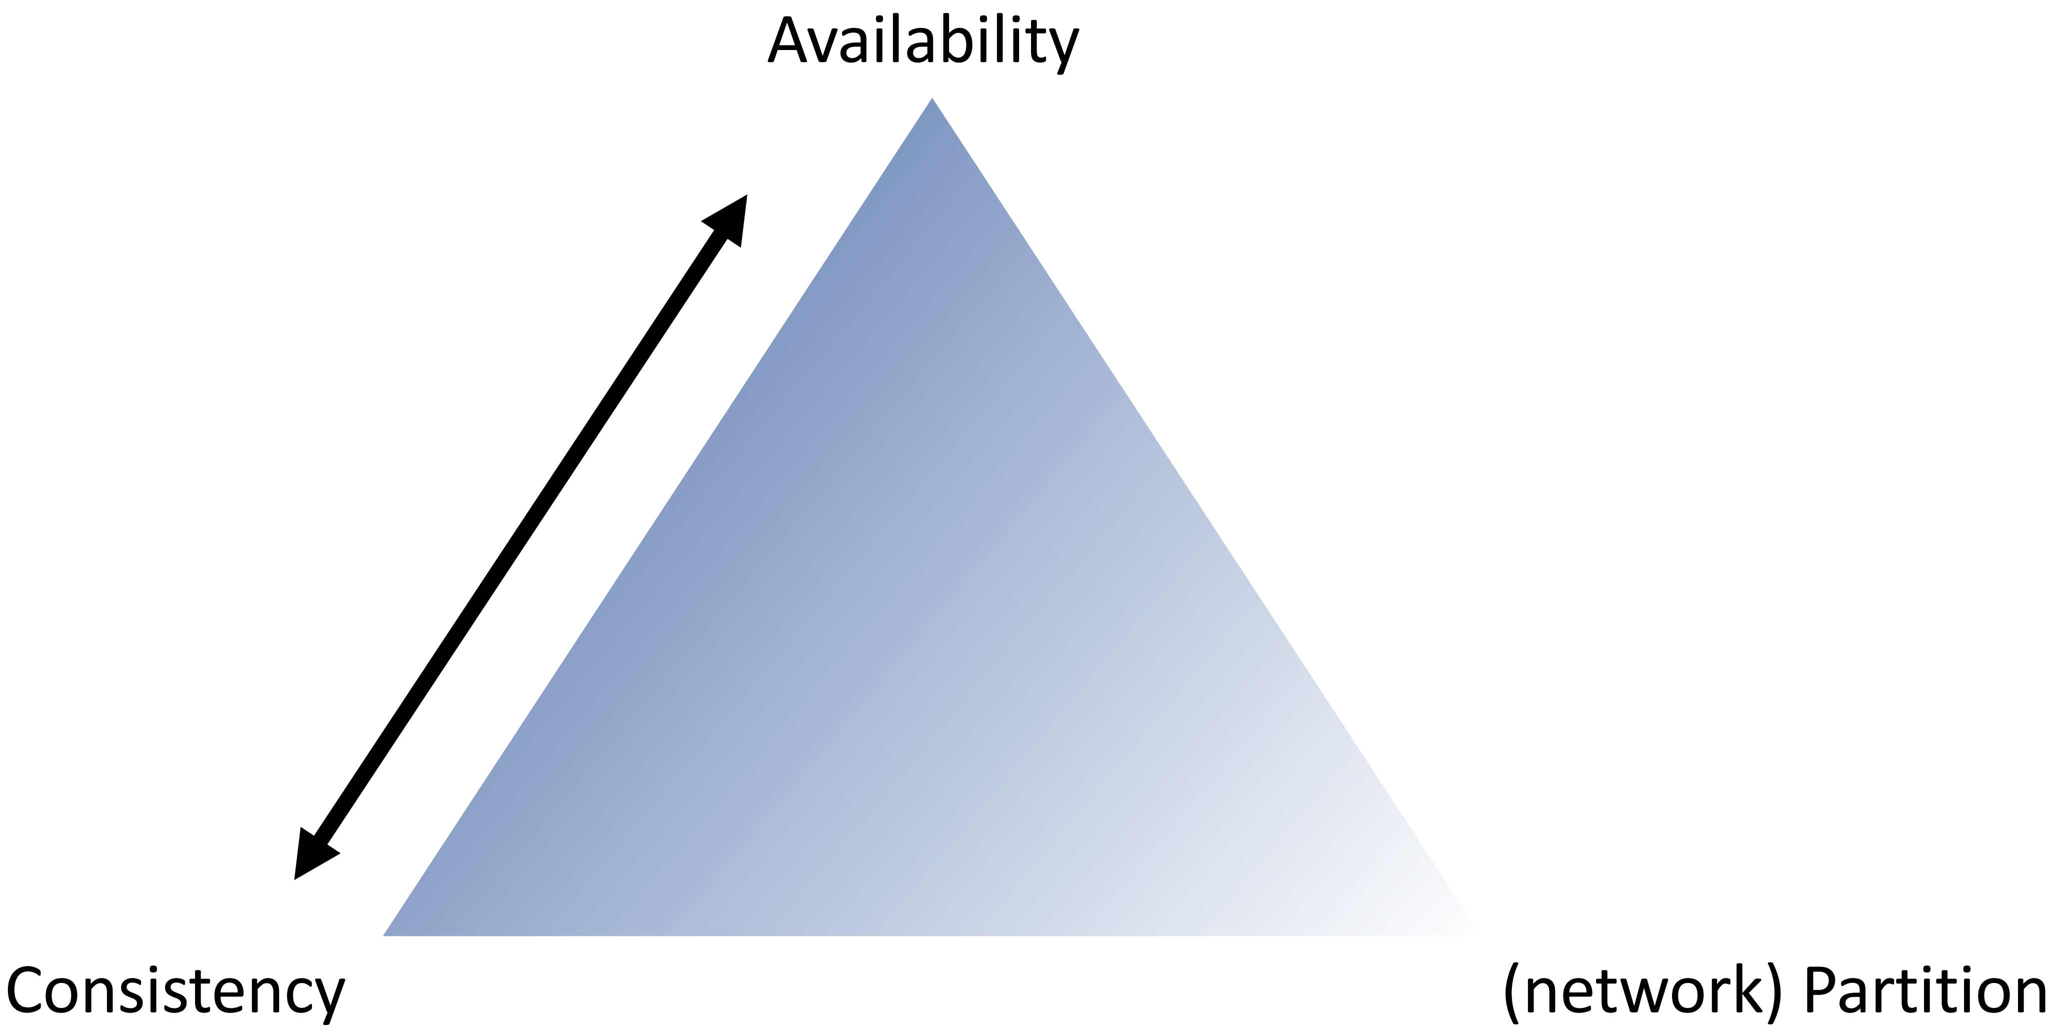

Resilient Hosting

Resilient Hosting – Active/Active

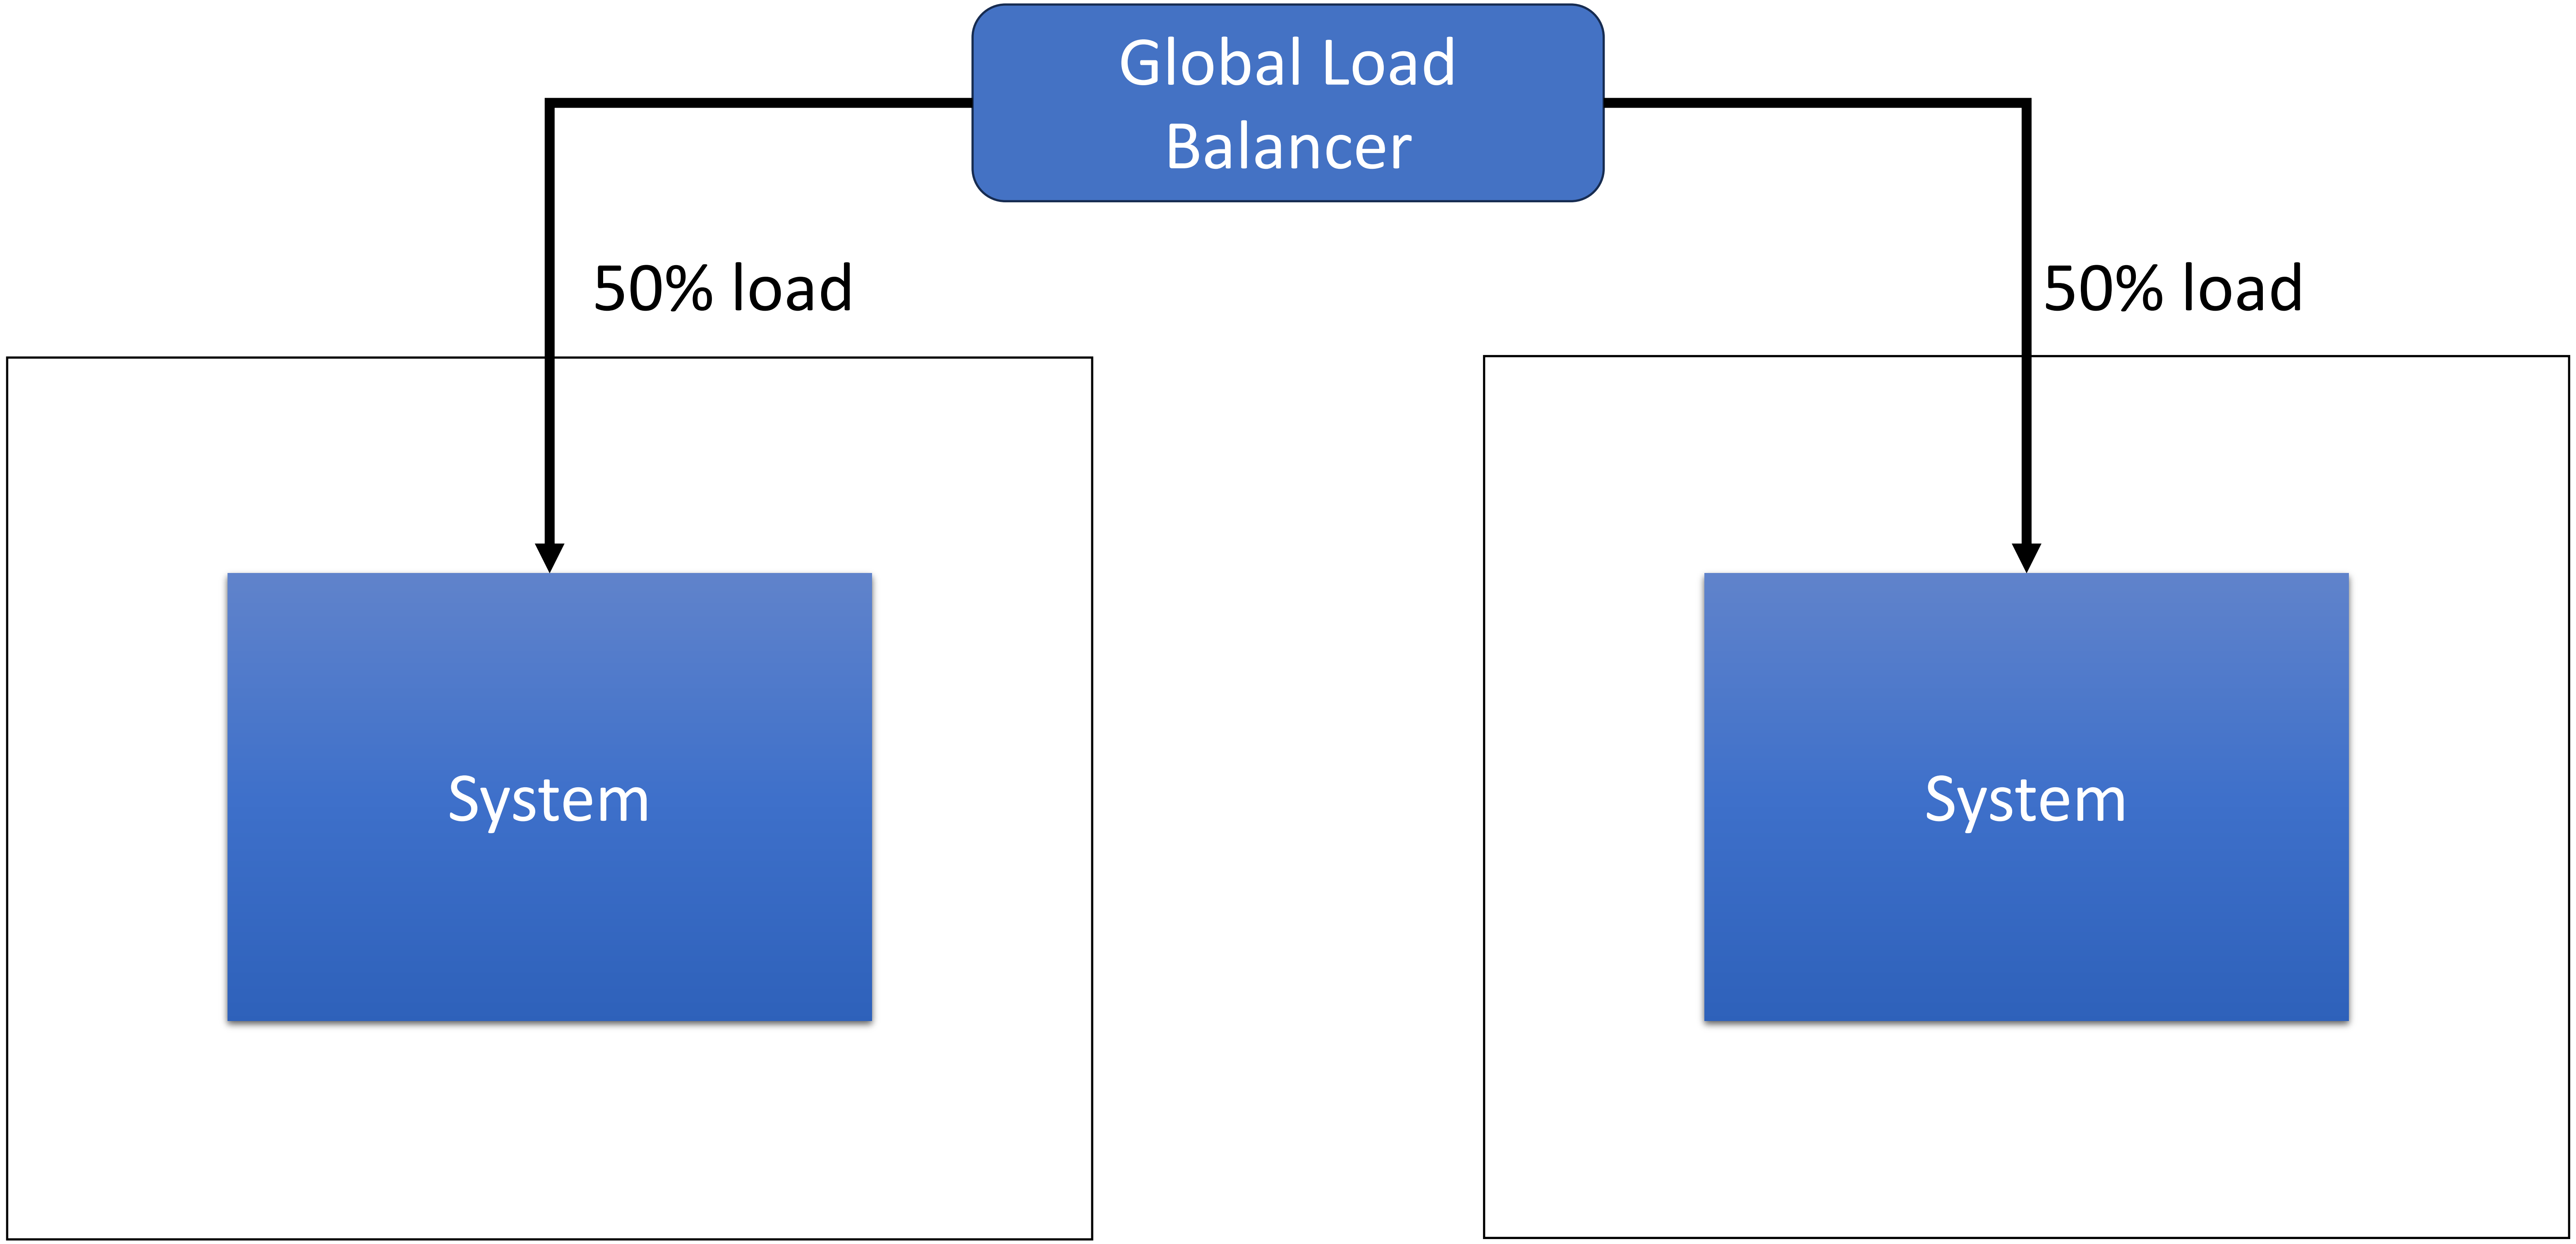

Resilient Hosting – Hot / Hot

also known as
Active / Passive

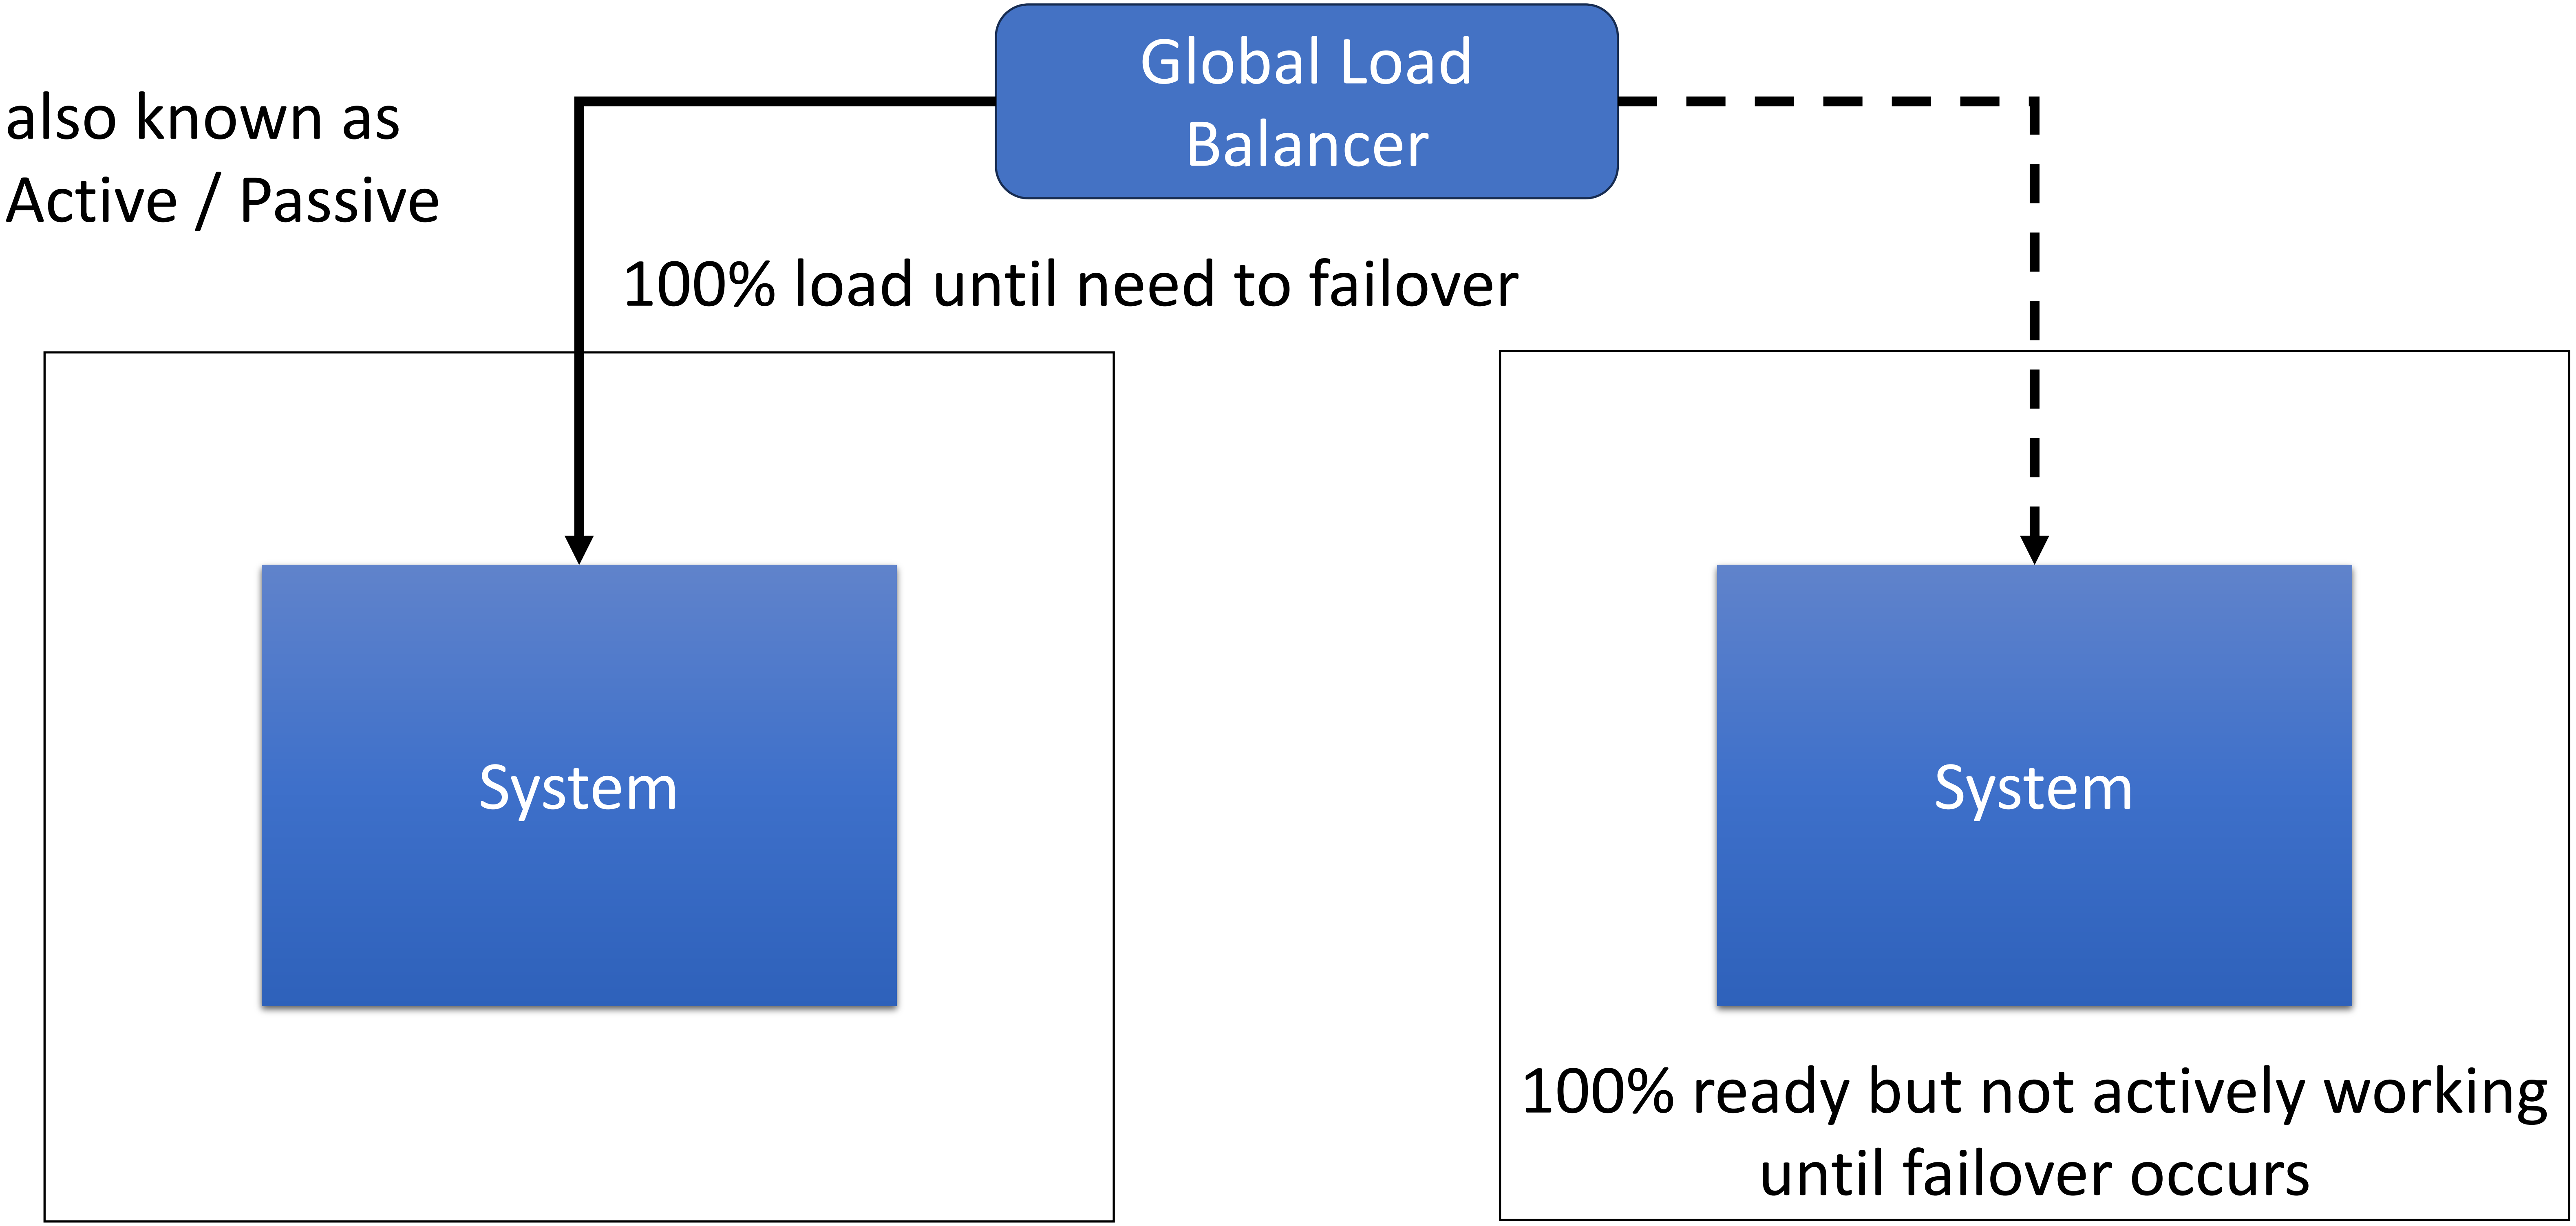

Global Load Balancer

100% load until need to failover

100% ready but not actively working until failover occurs

sometimes called Hot Standby

Resilient Hosting – Hot / Warm

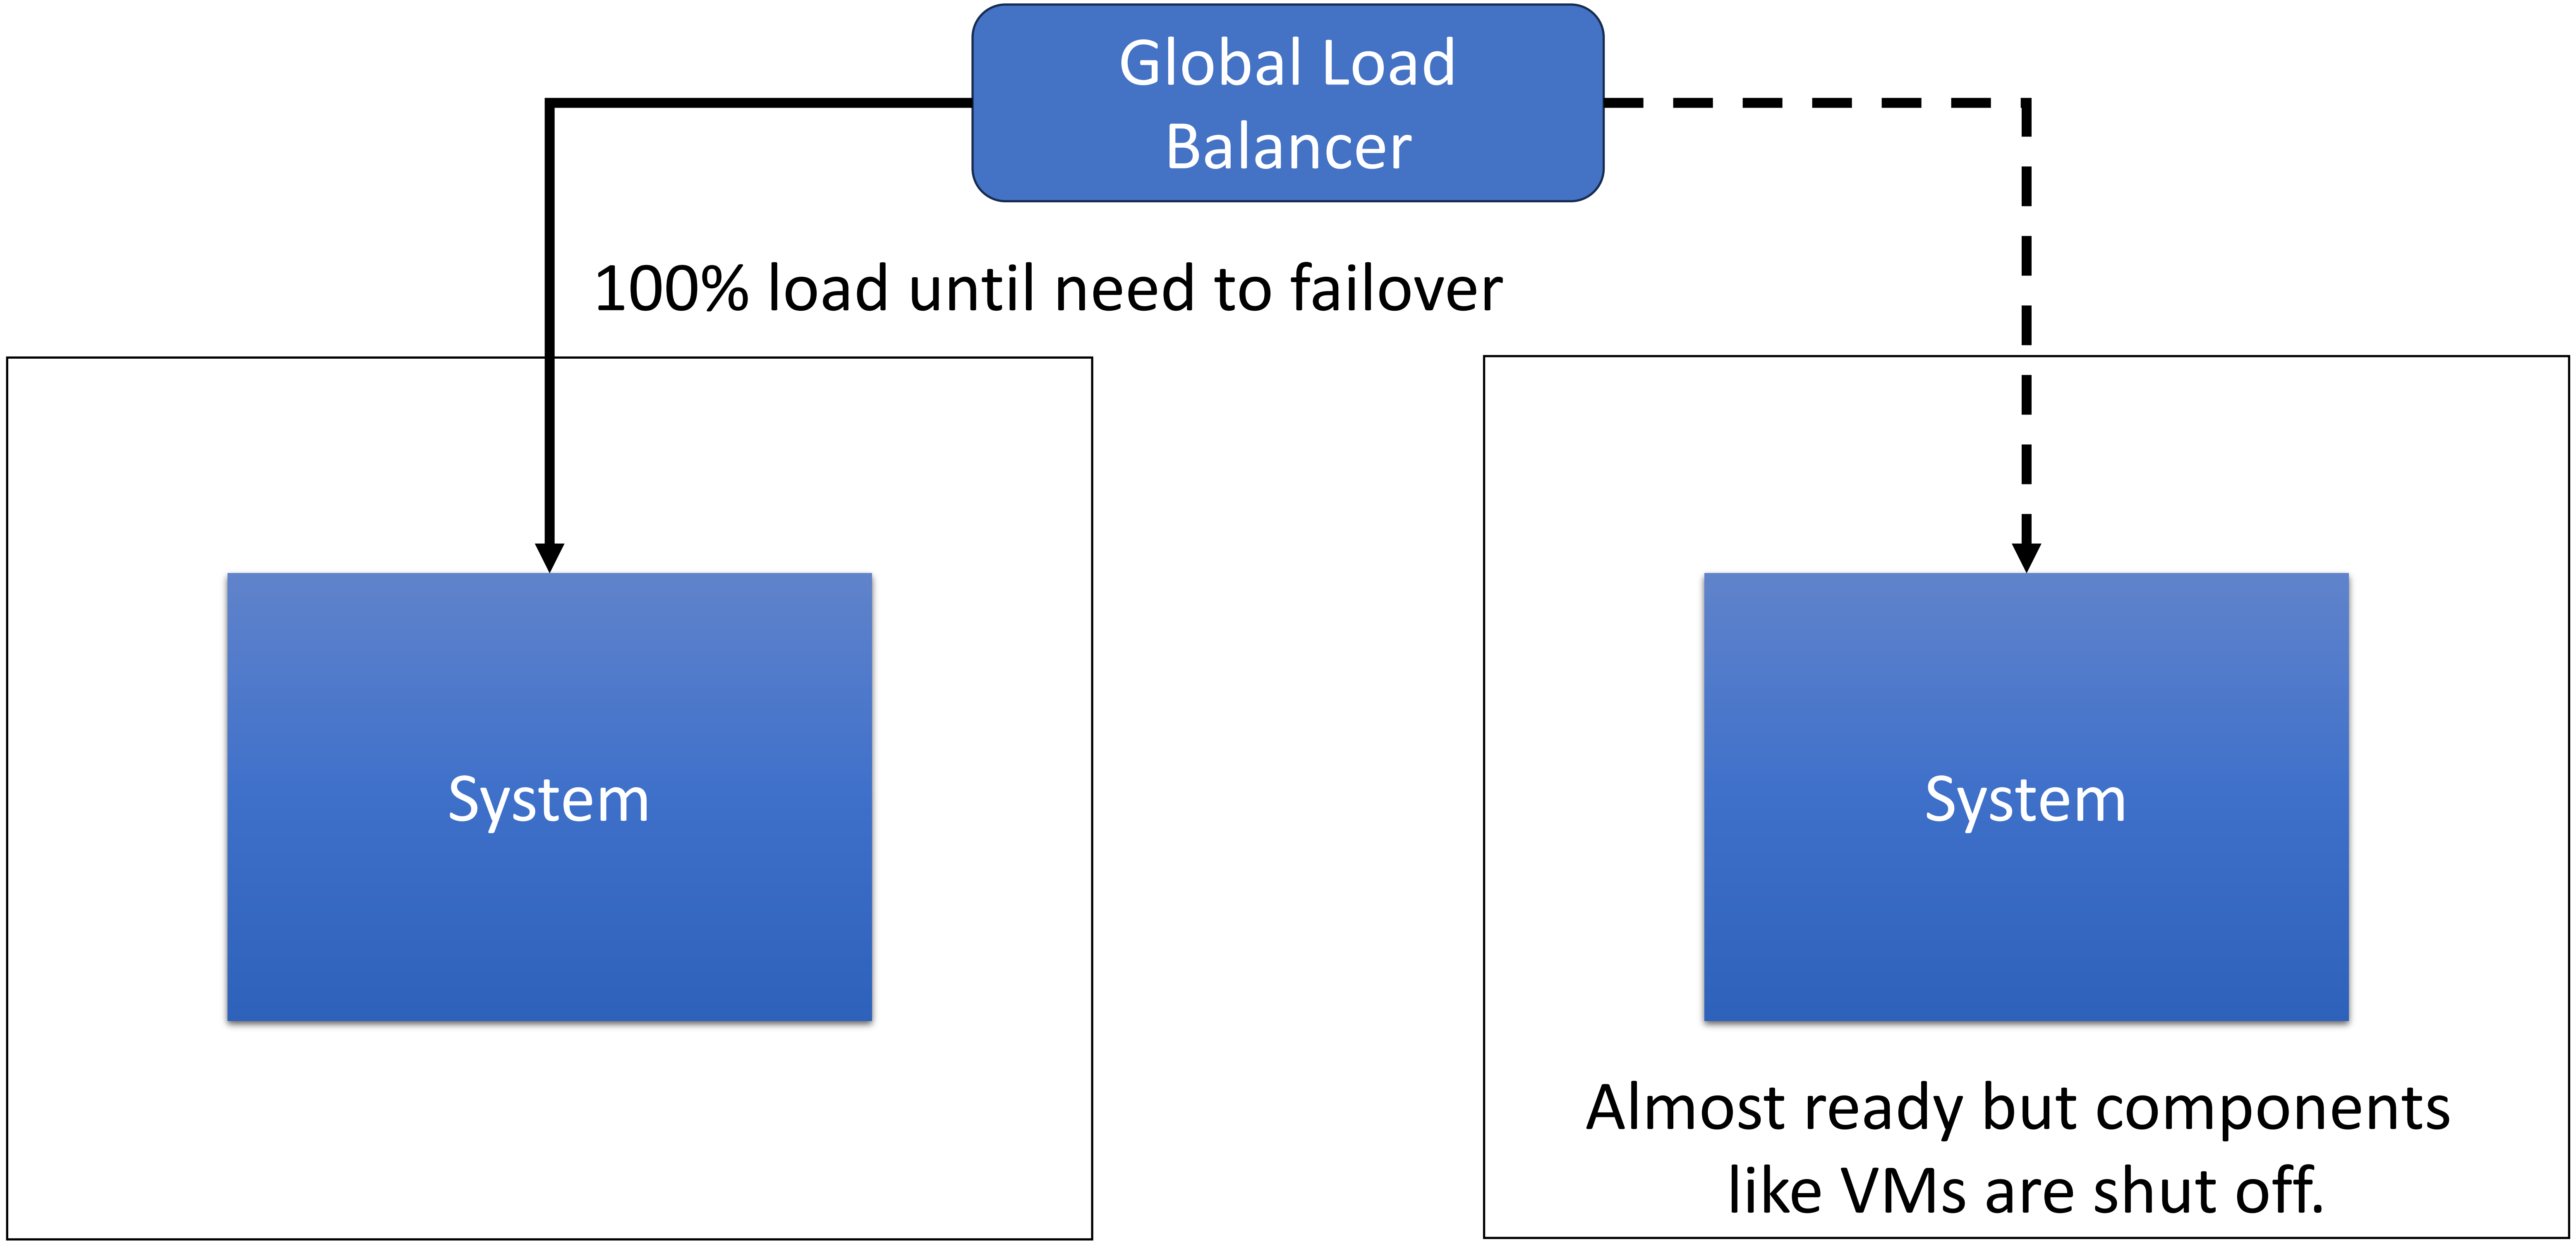

Global Load Balancer

100% load until need to failover

System

System

Almost ready but components like VMs are shut off.

Resilient Hosting – Hot / Cold

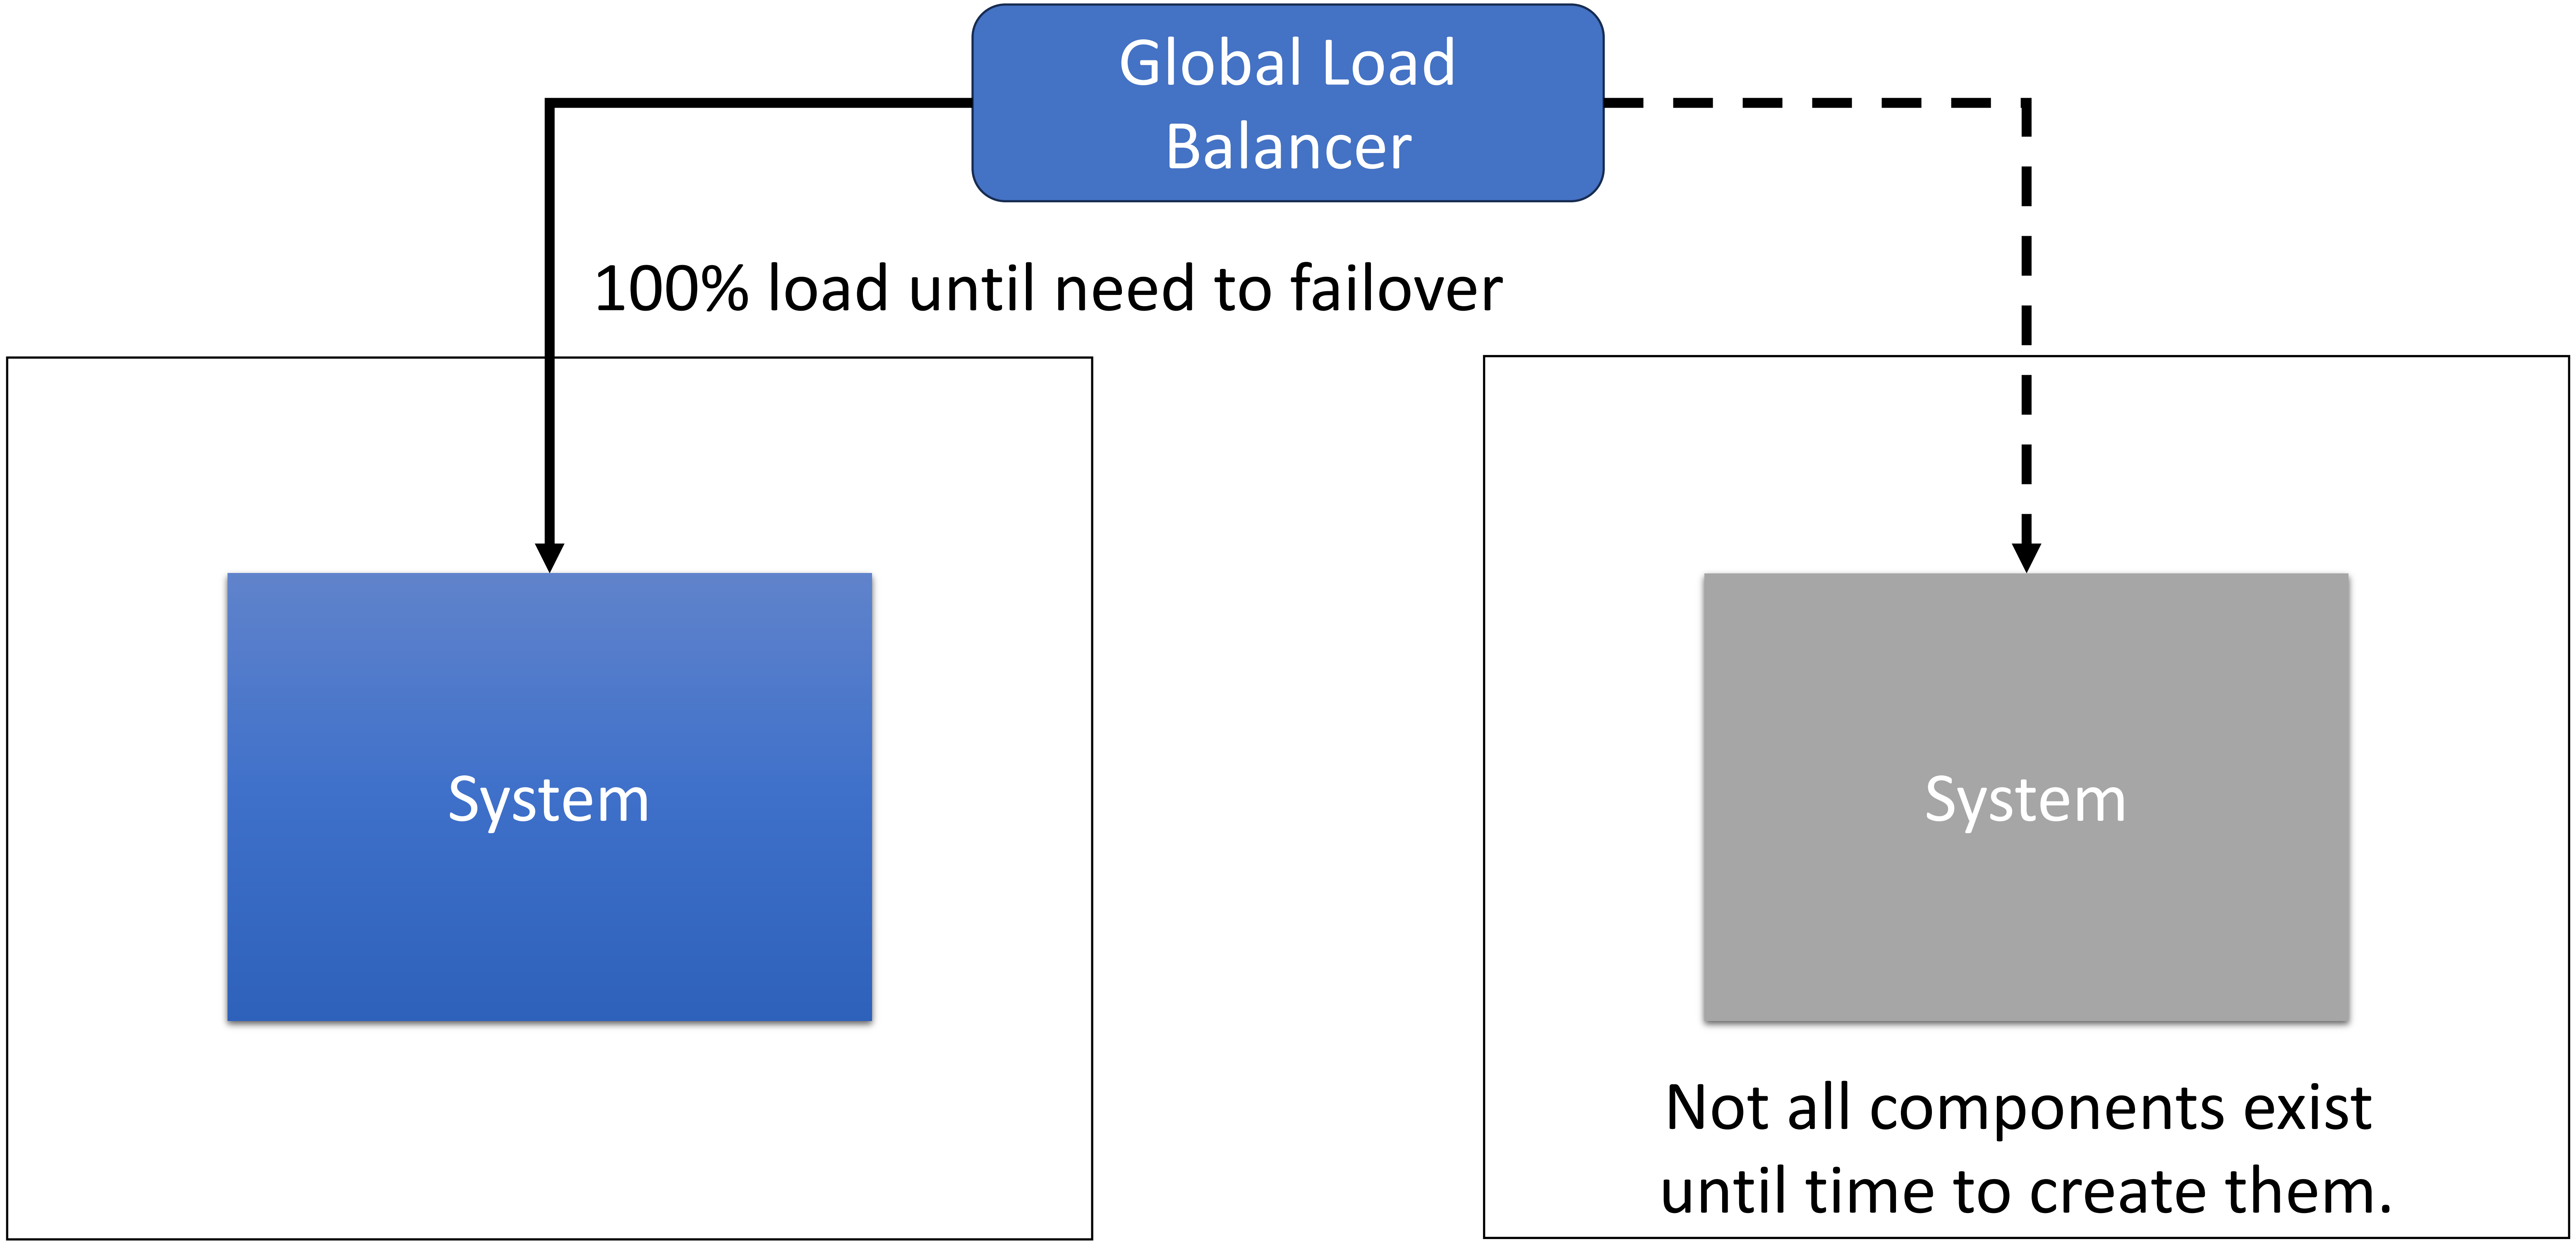

Global Load Balancer

100% load until need to failover

System

System

Not all components exist until time to create them.

Resilient Hosting

Key Point: The type of resilient hosting you choose influences your RTO.

Active / Active == lowest RTO

Hot / Cold == highest RTO

Resilient Hosting – Failover & Failback

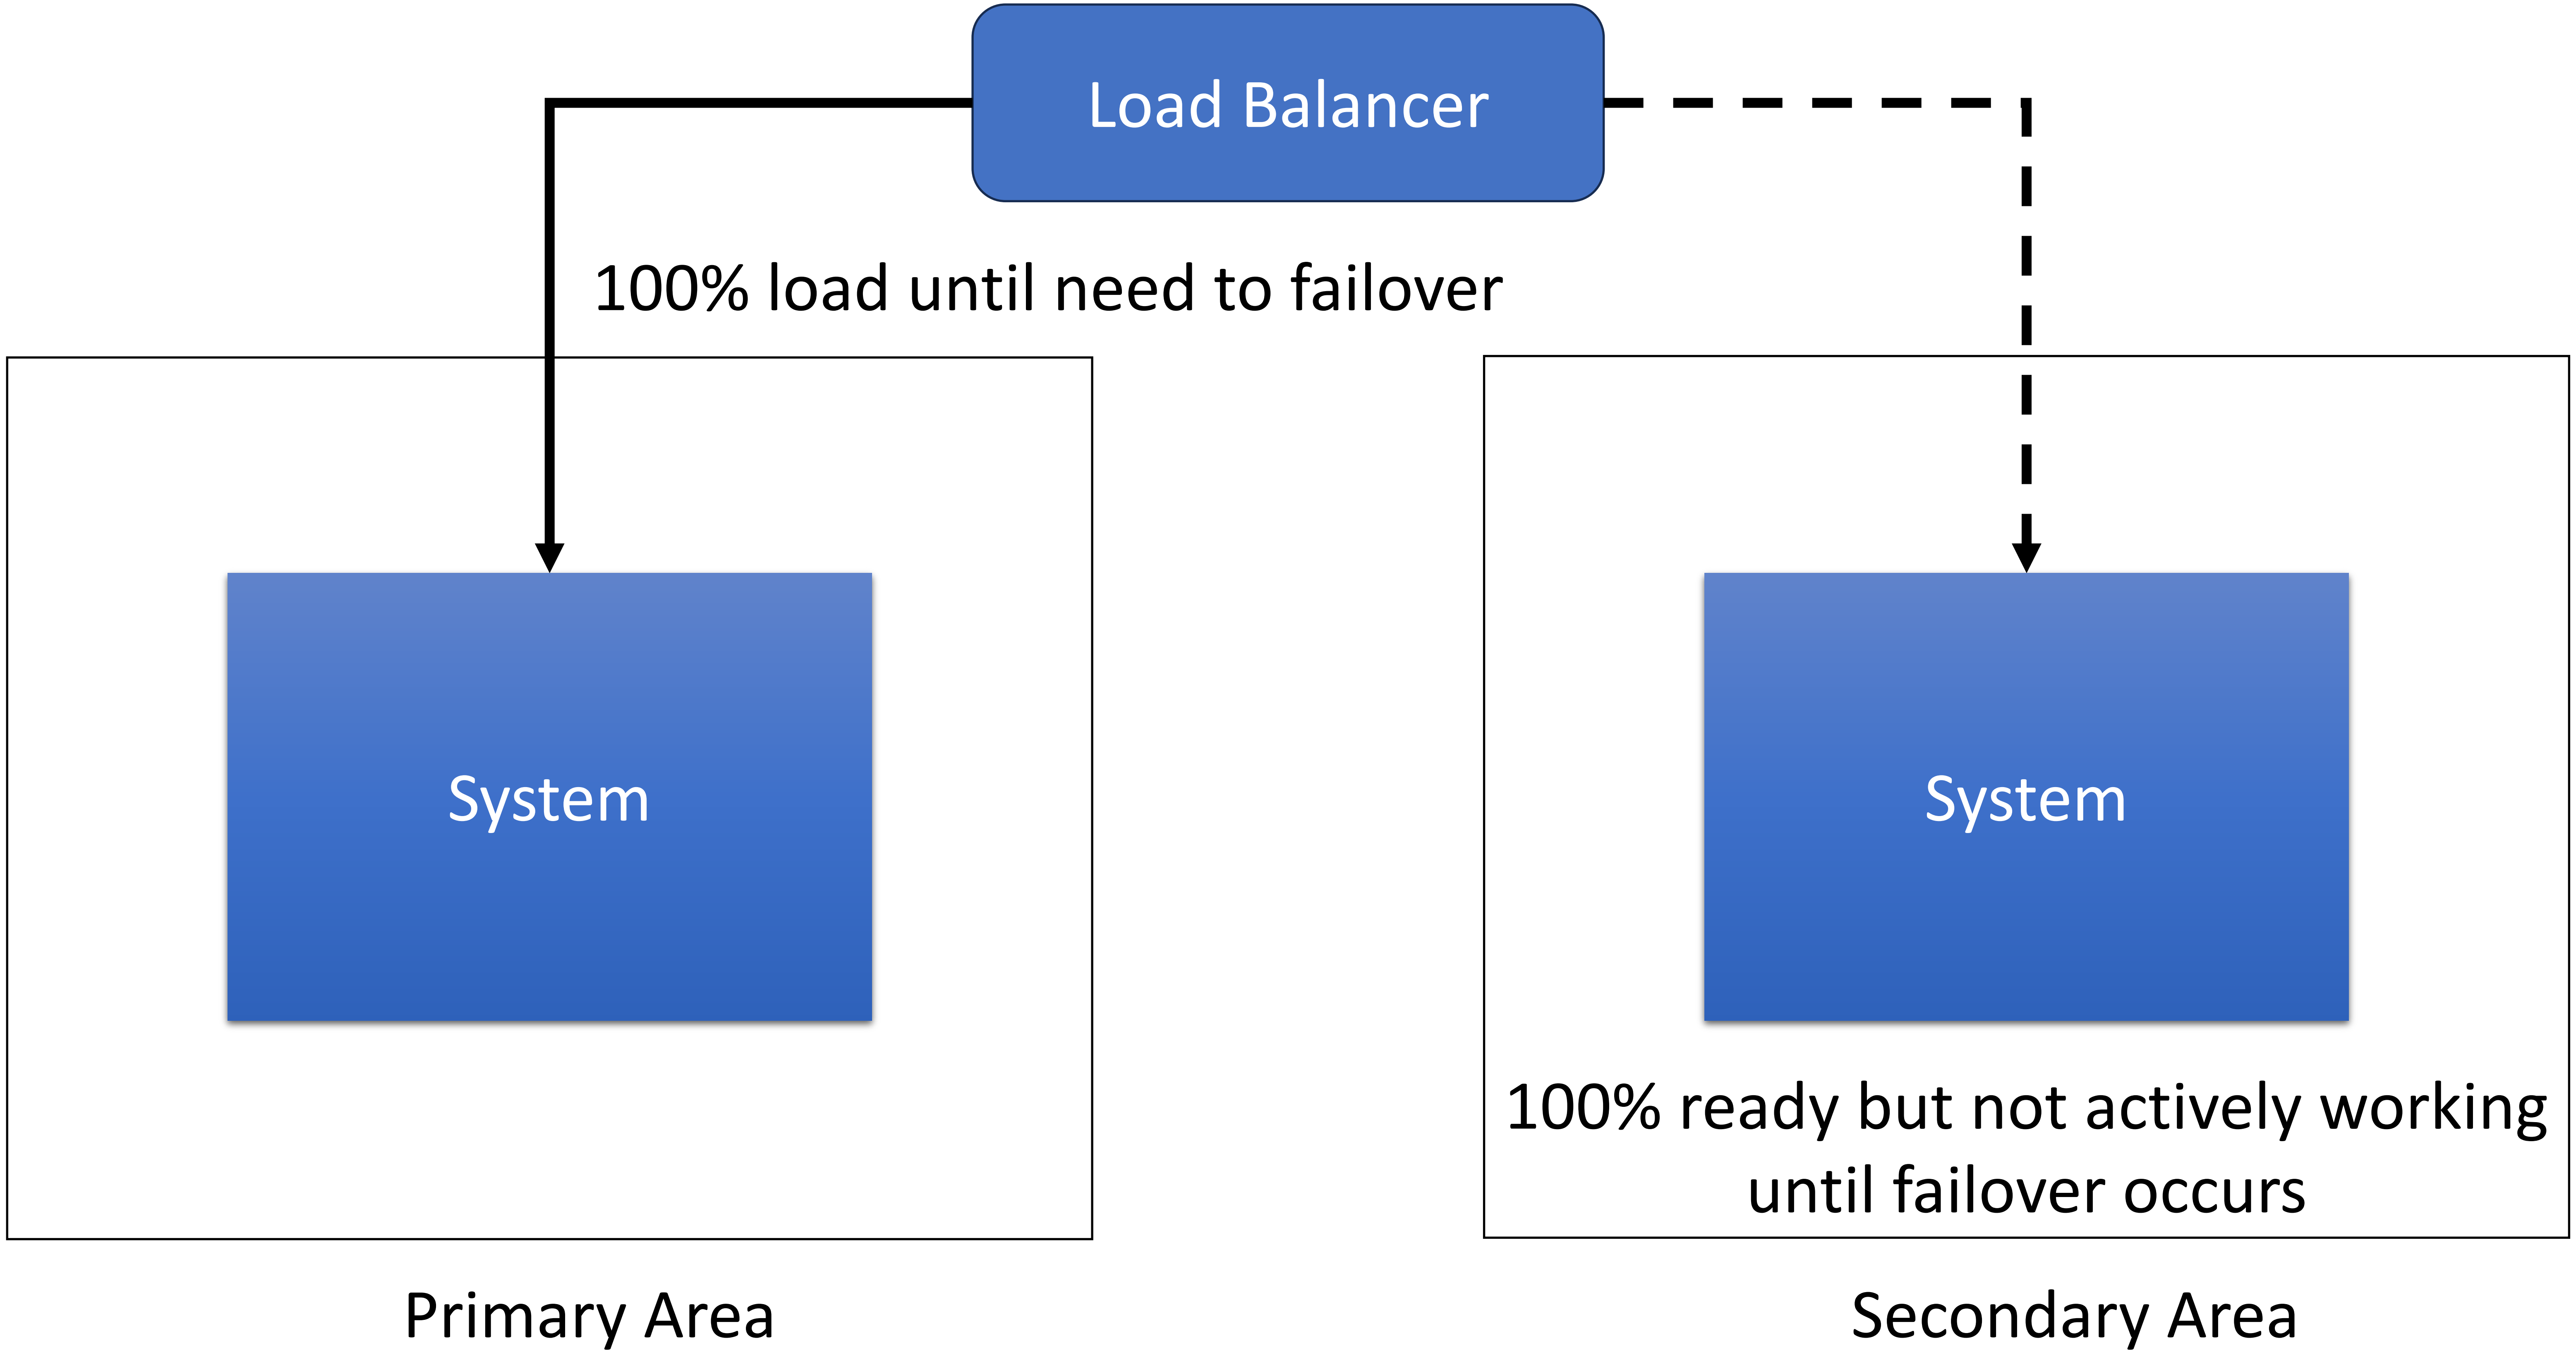

Load Balancer

100% load until need to failover

System

System

100% ready but not actively working until failover occurs

Primary Area

Secondary Area

Resilient Hosting – Failover & Failback

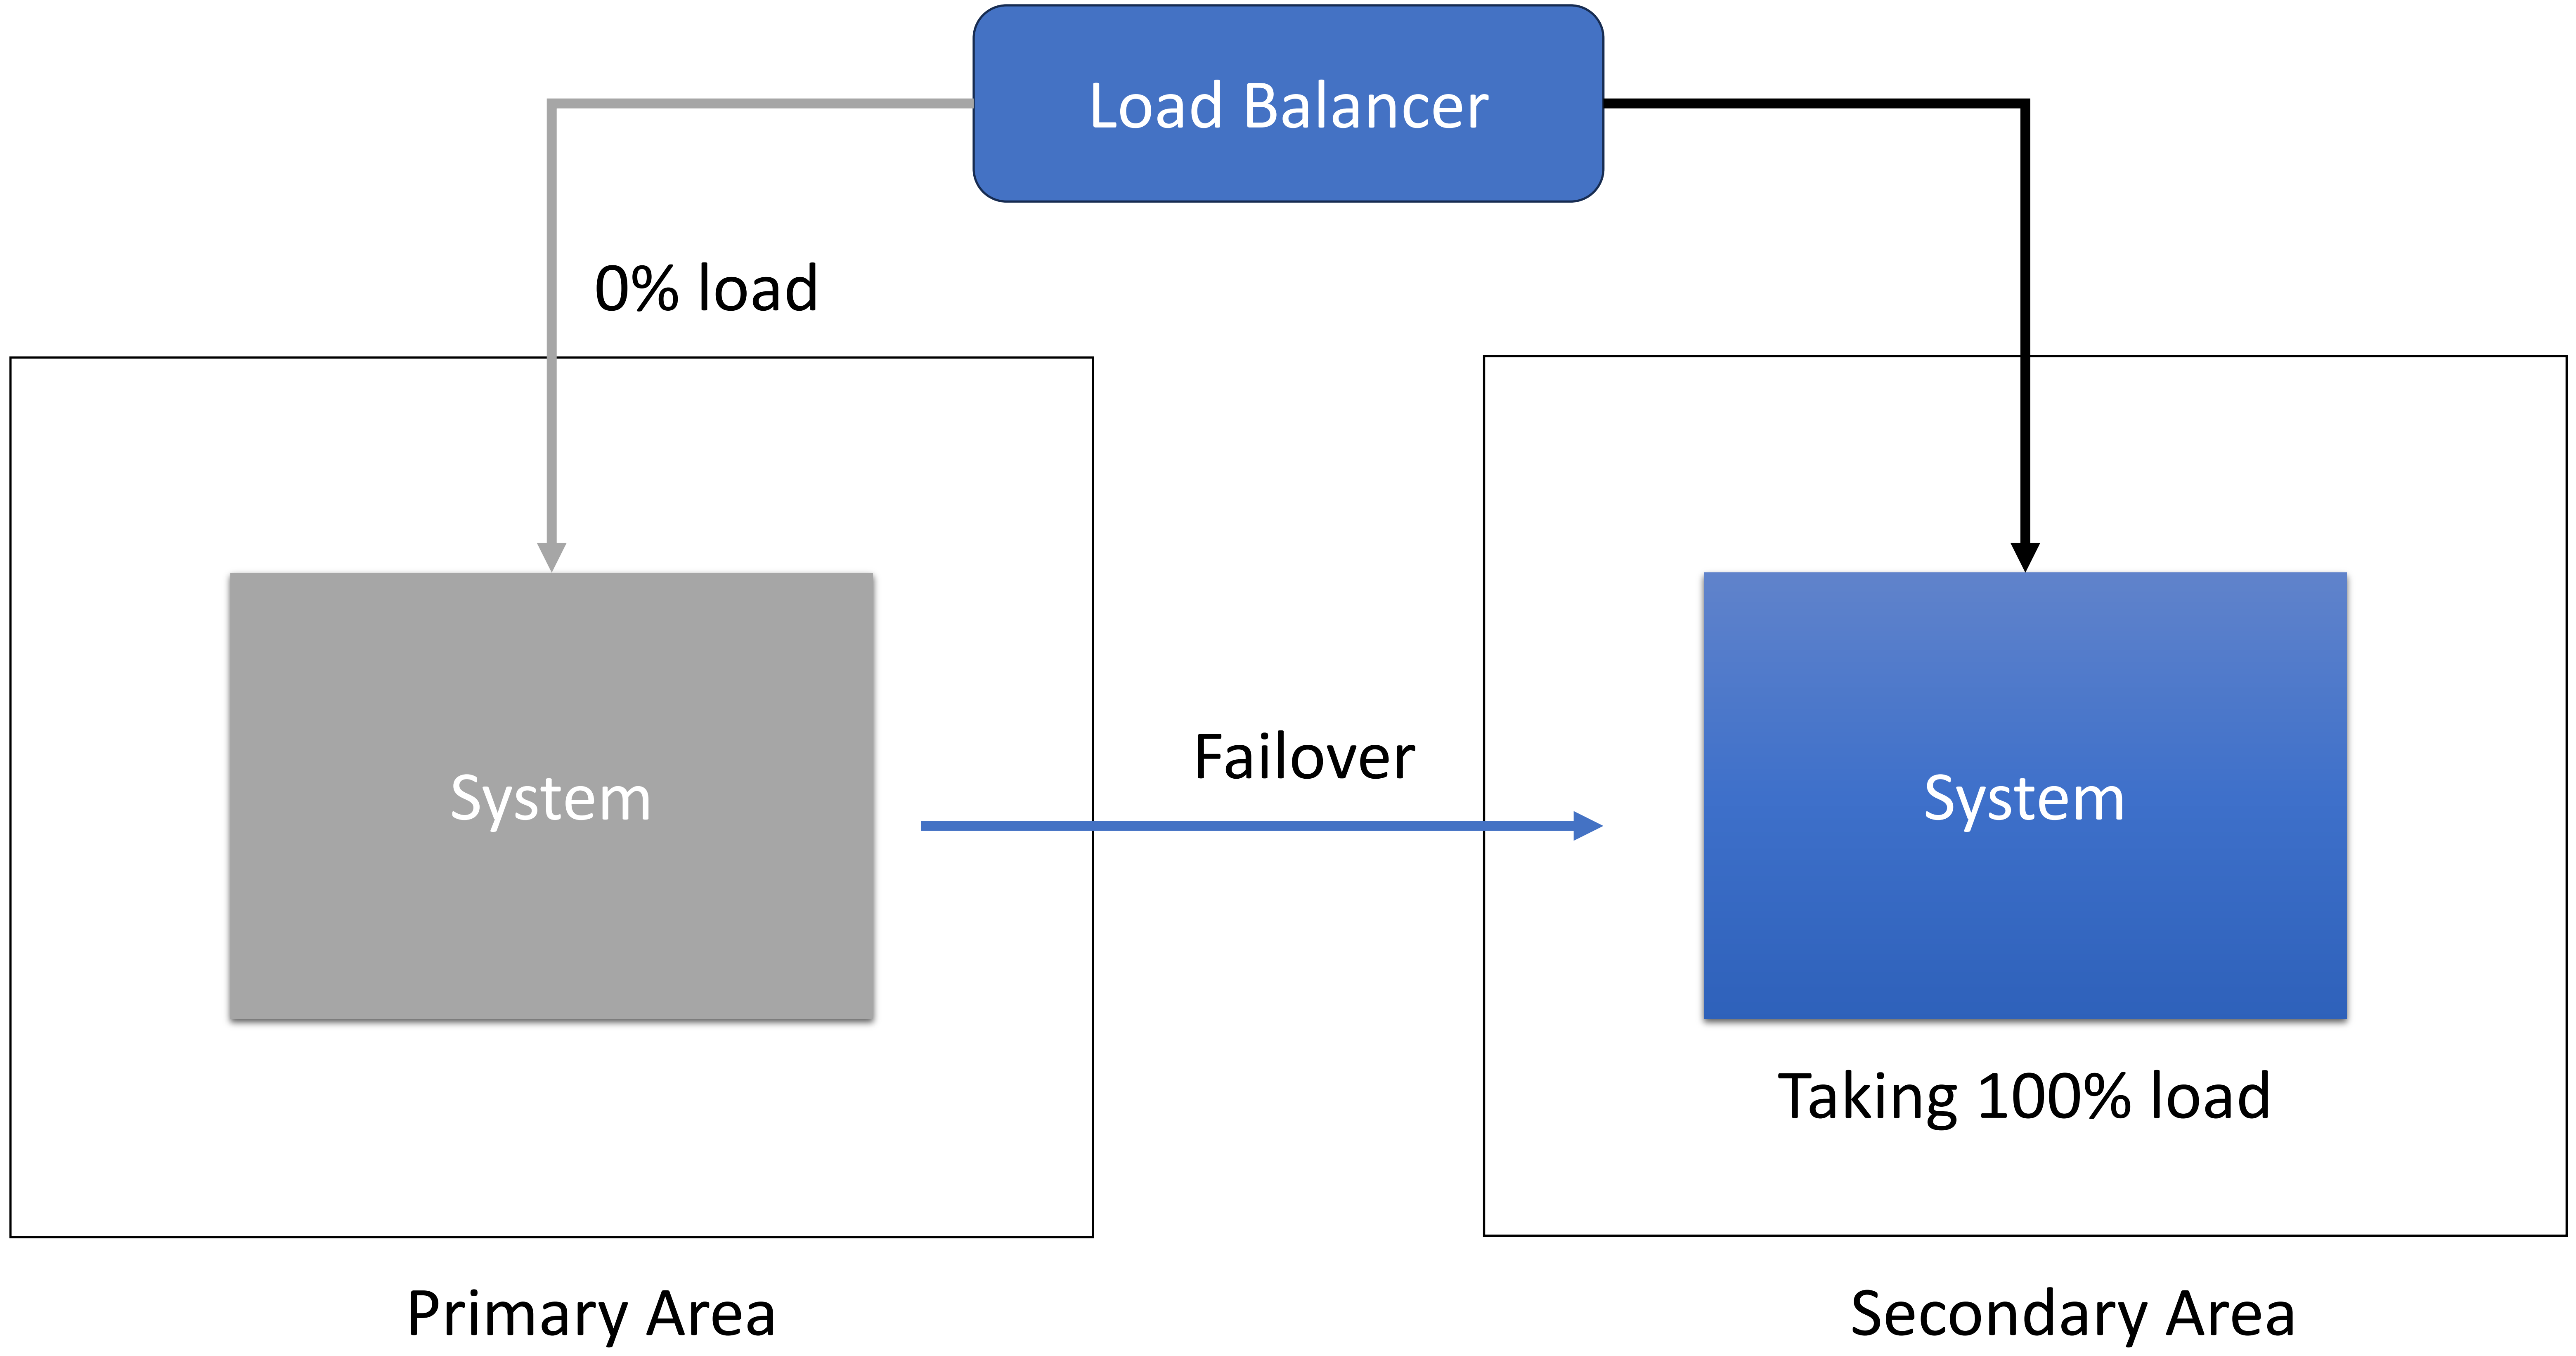

Load Balancer

0% load

System

Failover

System

Taking 100% load

Primary Area

Secondary Area

Resilient Hosting – Failover & Failback

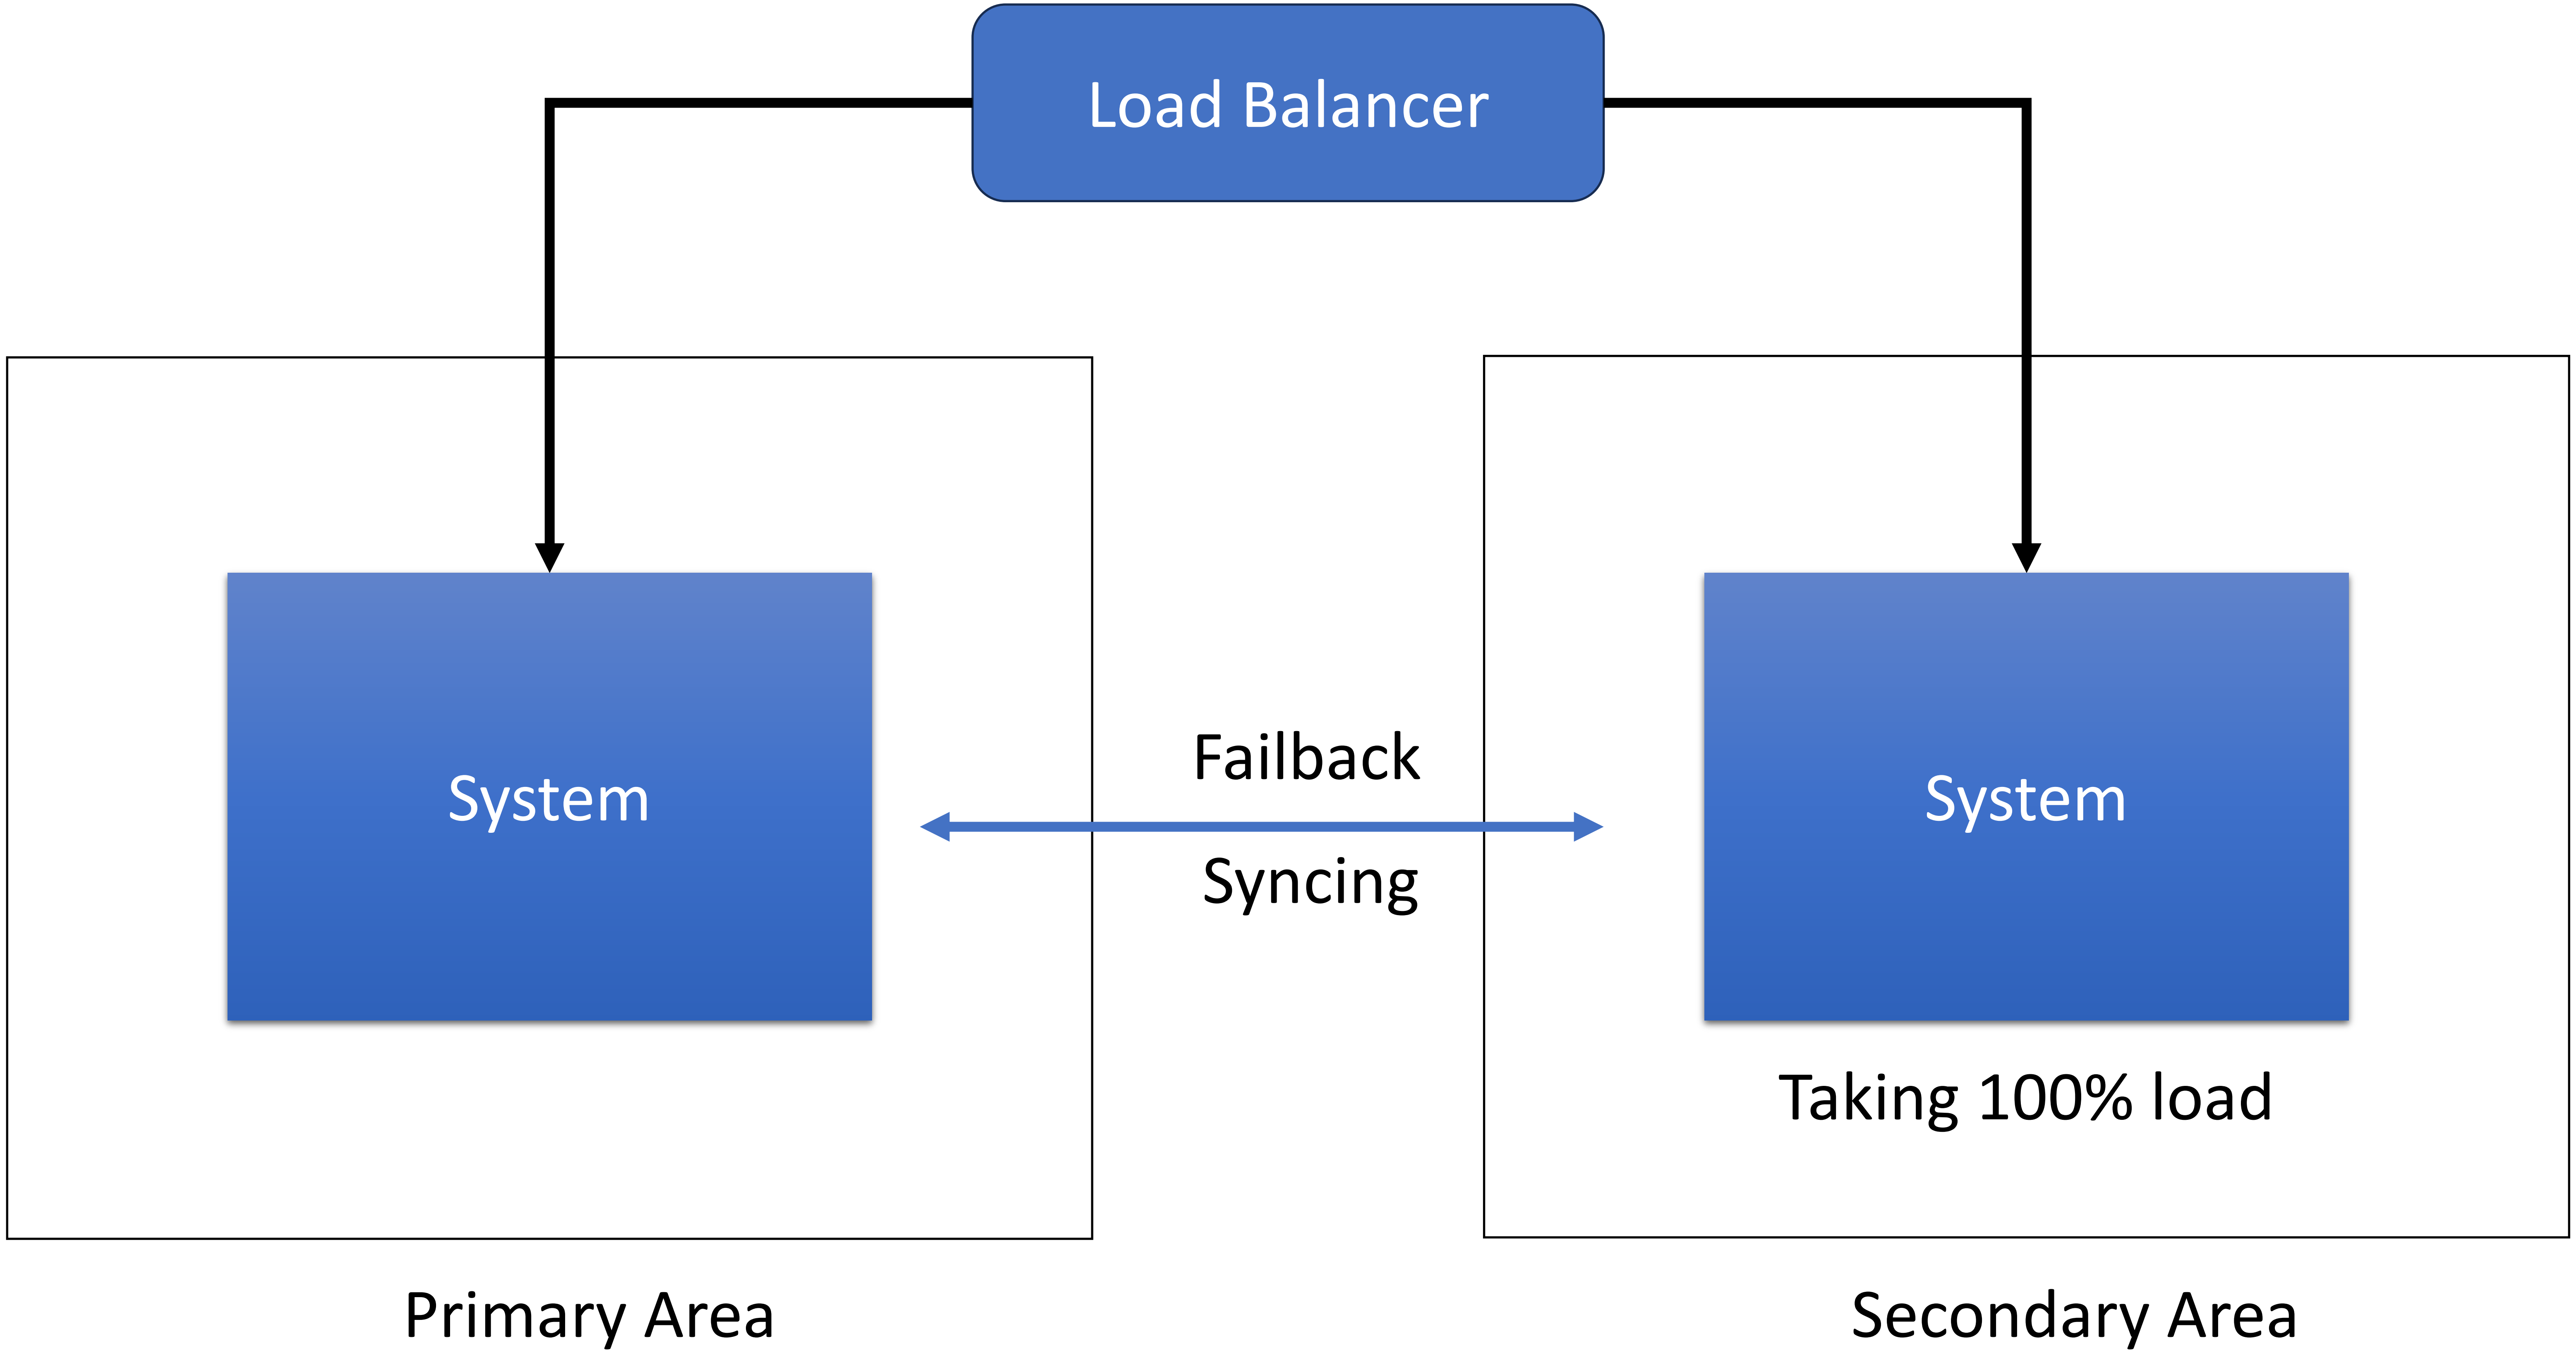

Resilient Hosting – Failover & Failback

Failover can be caused by Azure and/or you.

Resilient Hosting – Failover & Failback

Failover can be caused by Azure and/or you.

Failback may not be under your control.

Anatomy of an Azure Region

Anatomy of an Azure Region

54+ Regions around the world

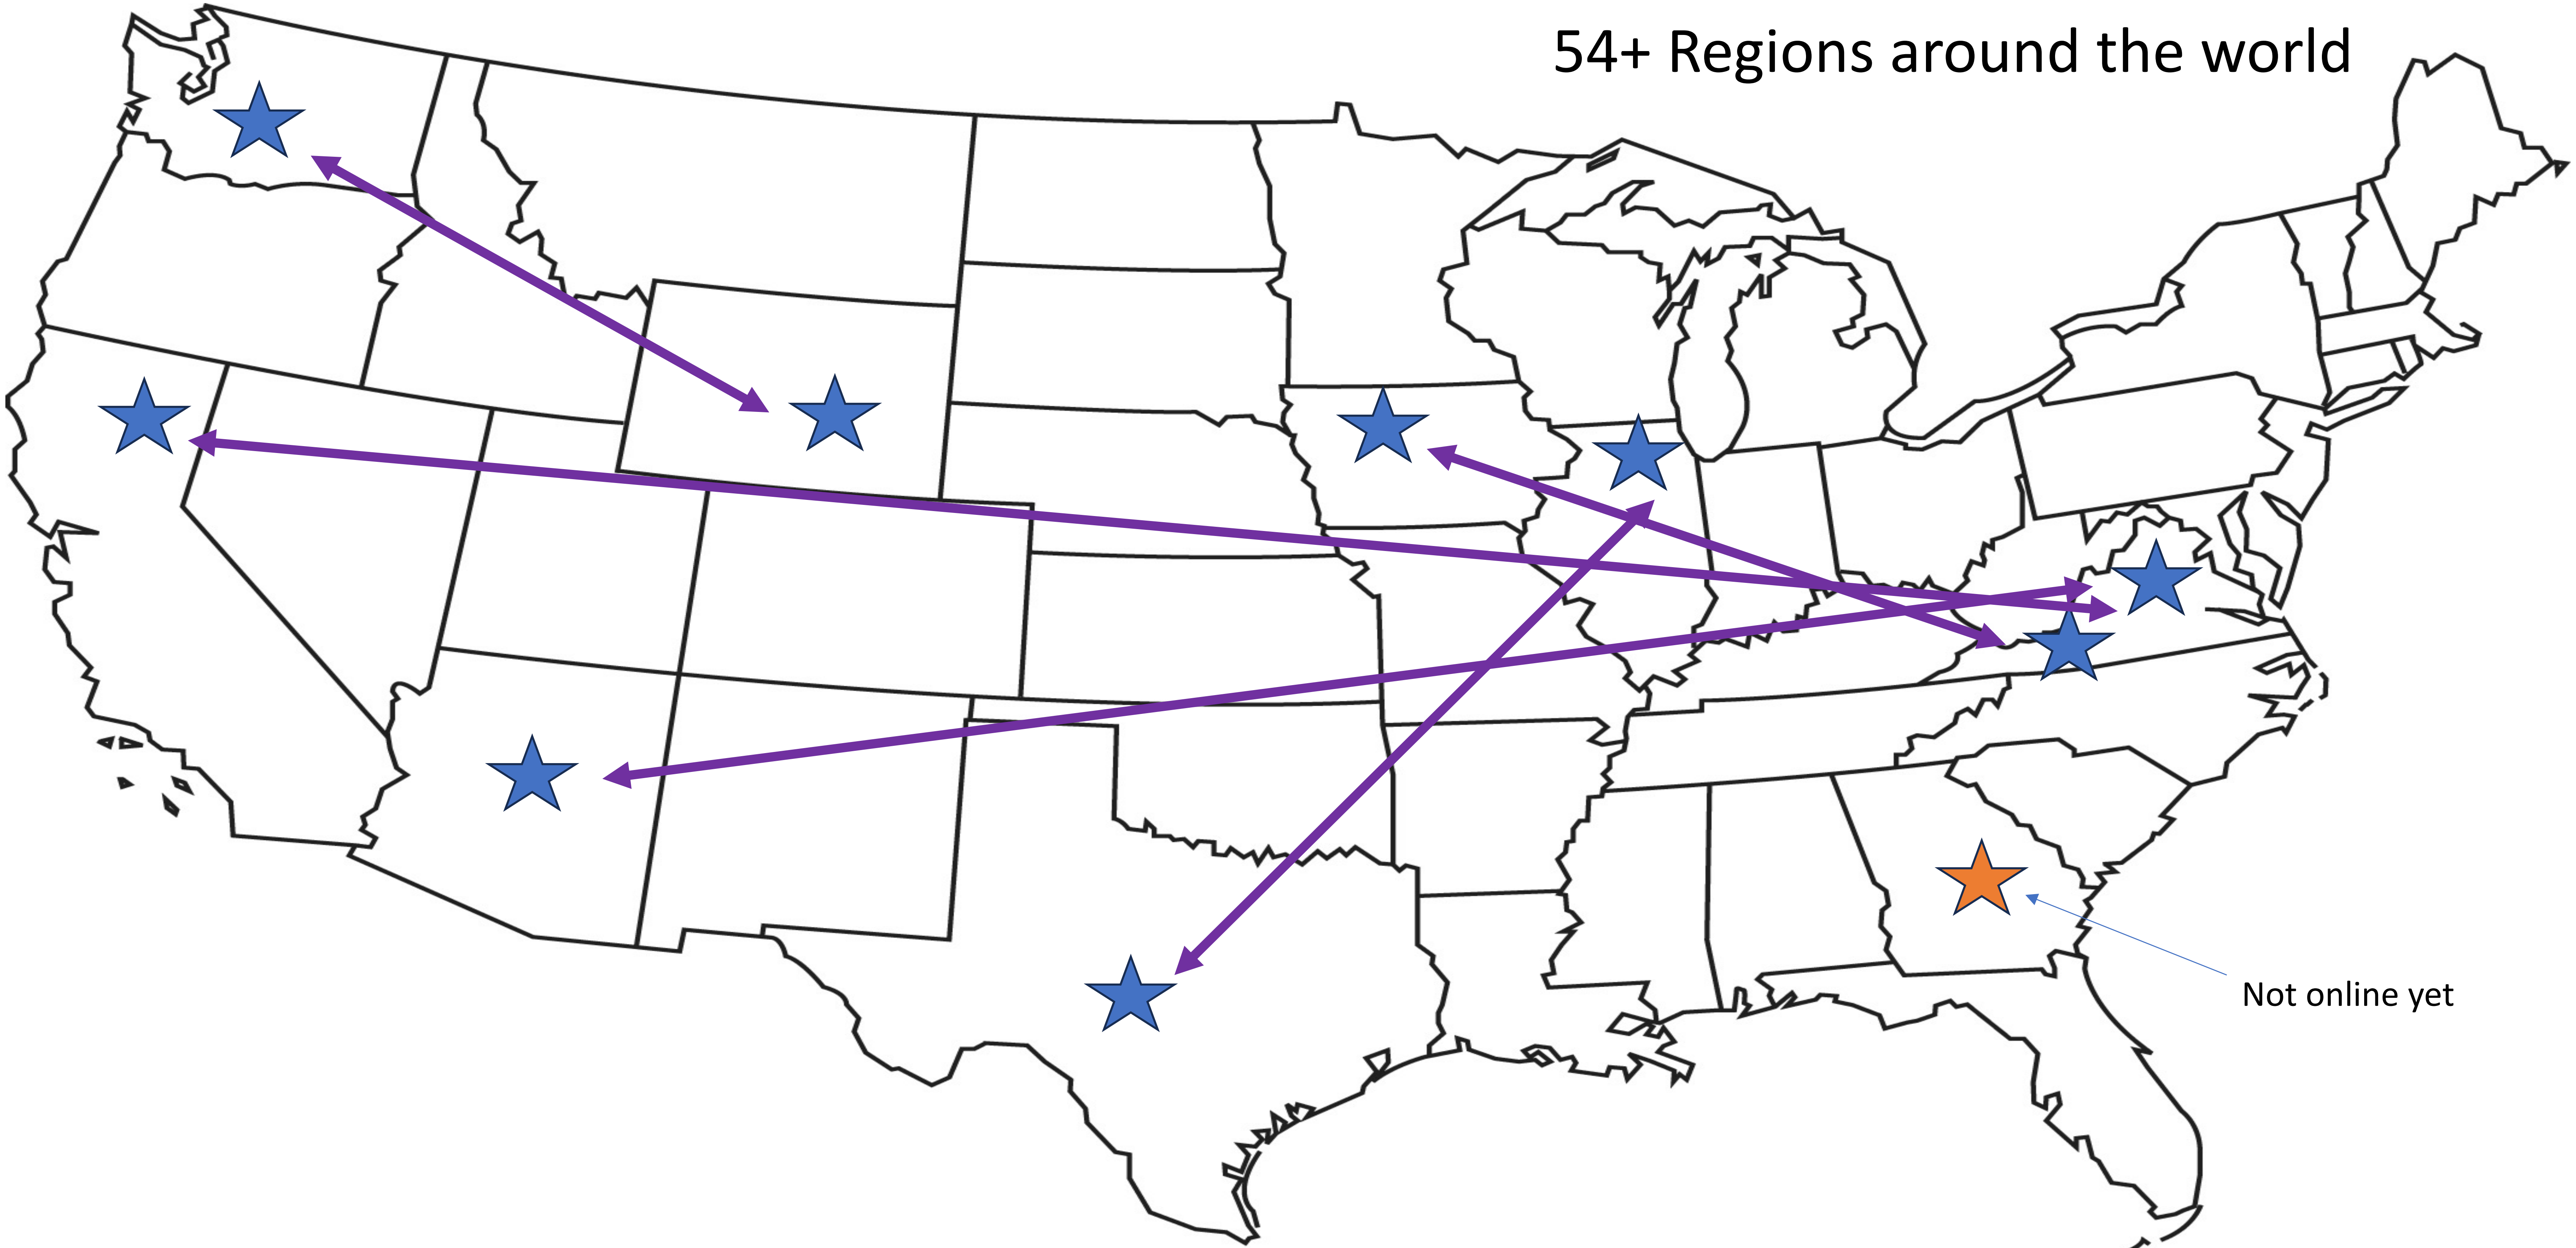

Anatomy of an Azure Region

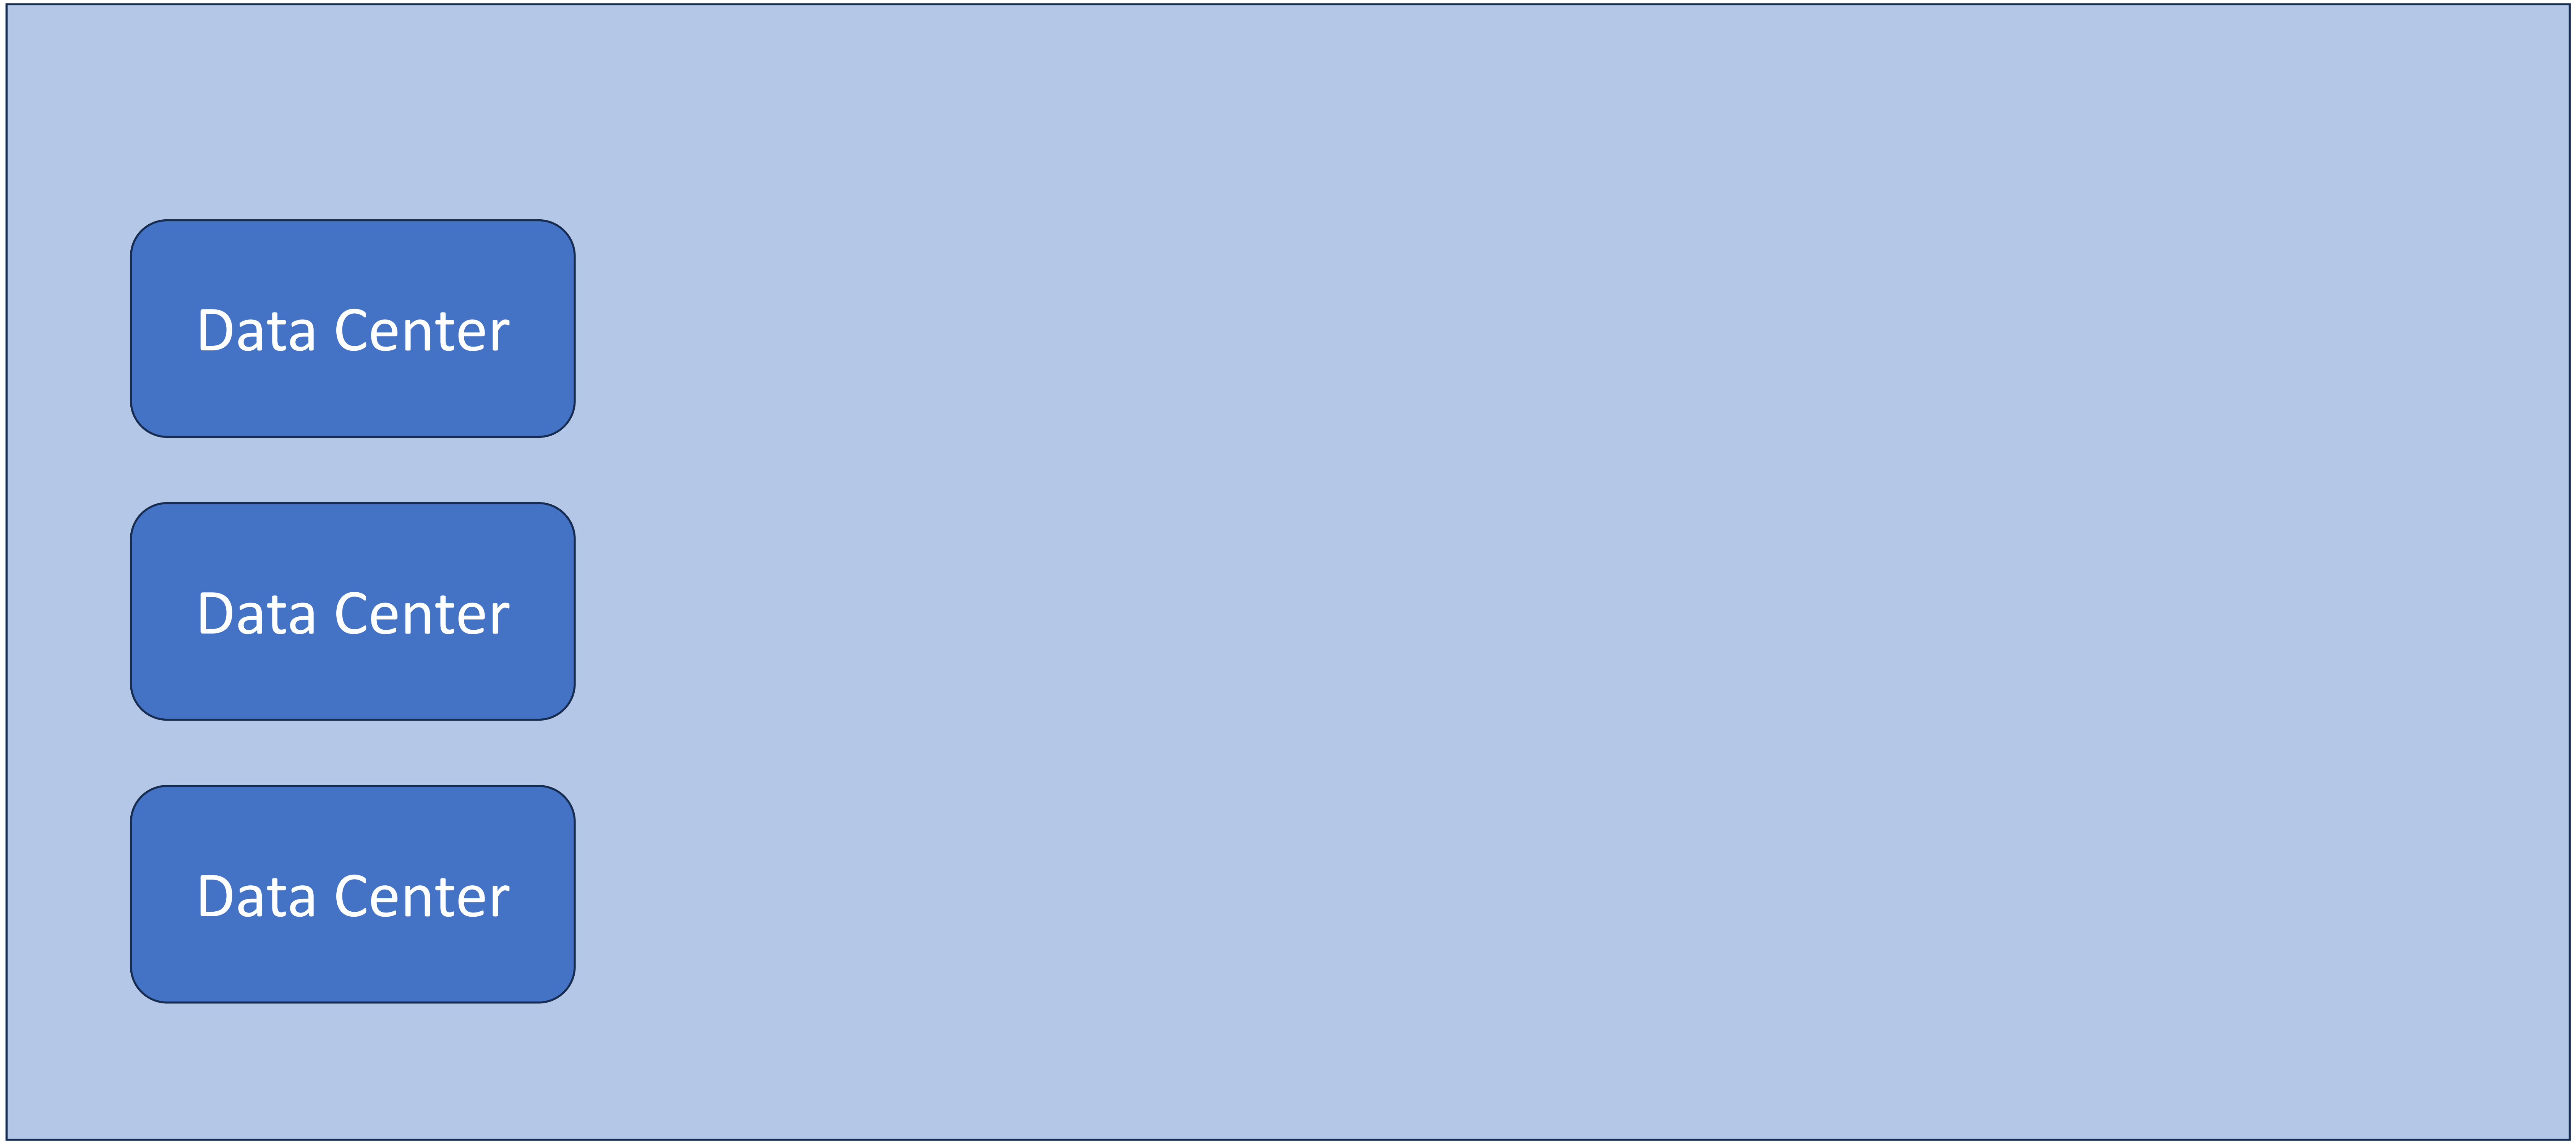

Data Center

Data Center

Data Center

Anatomy of an Azure Region

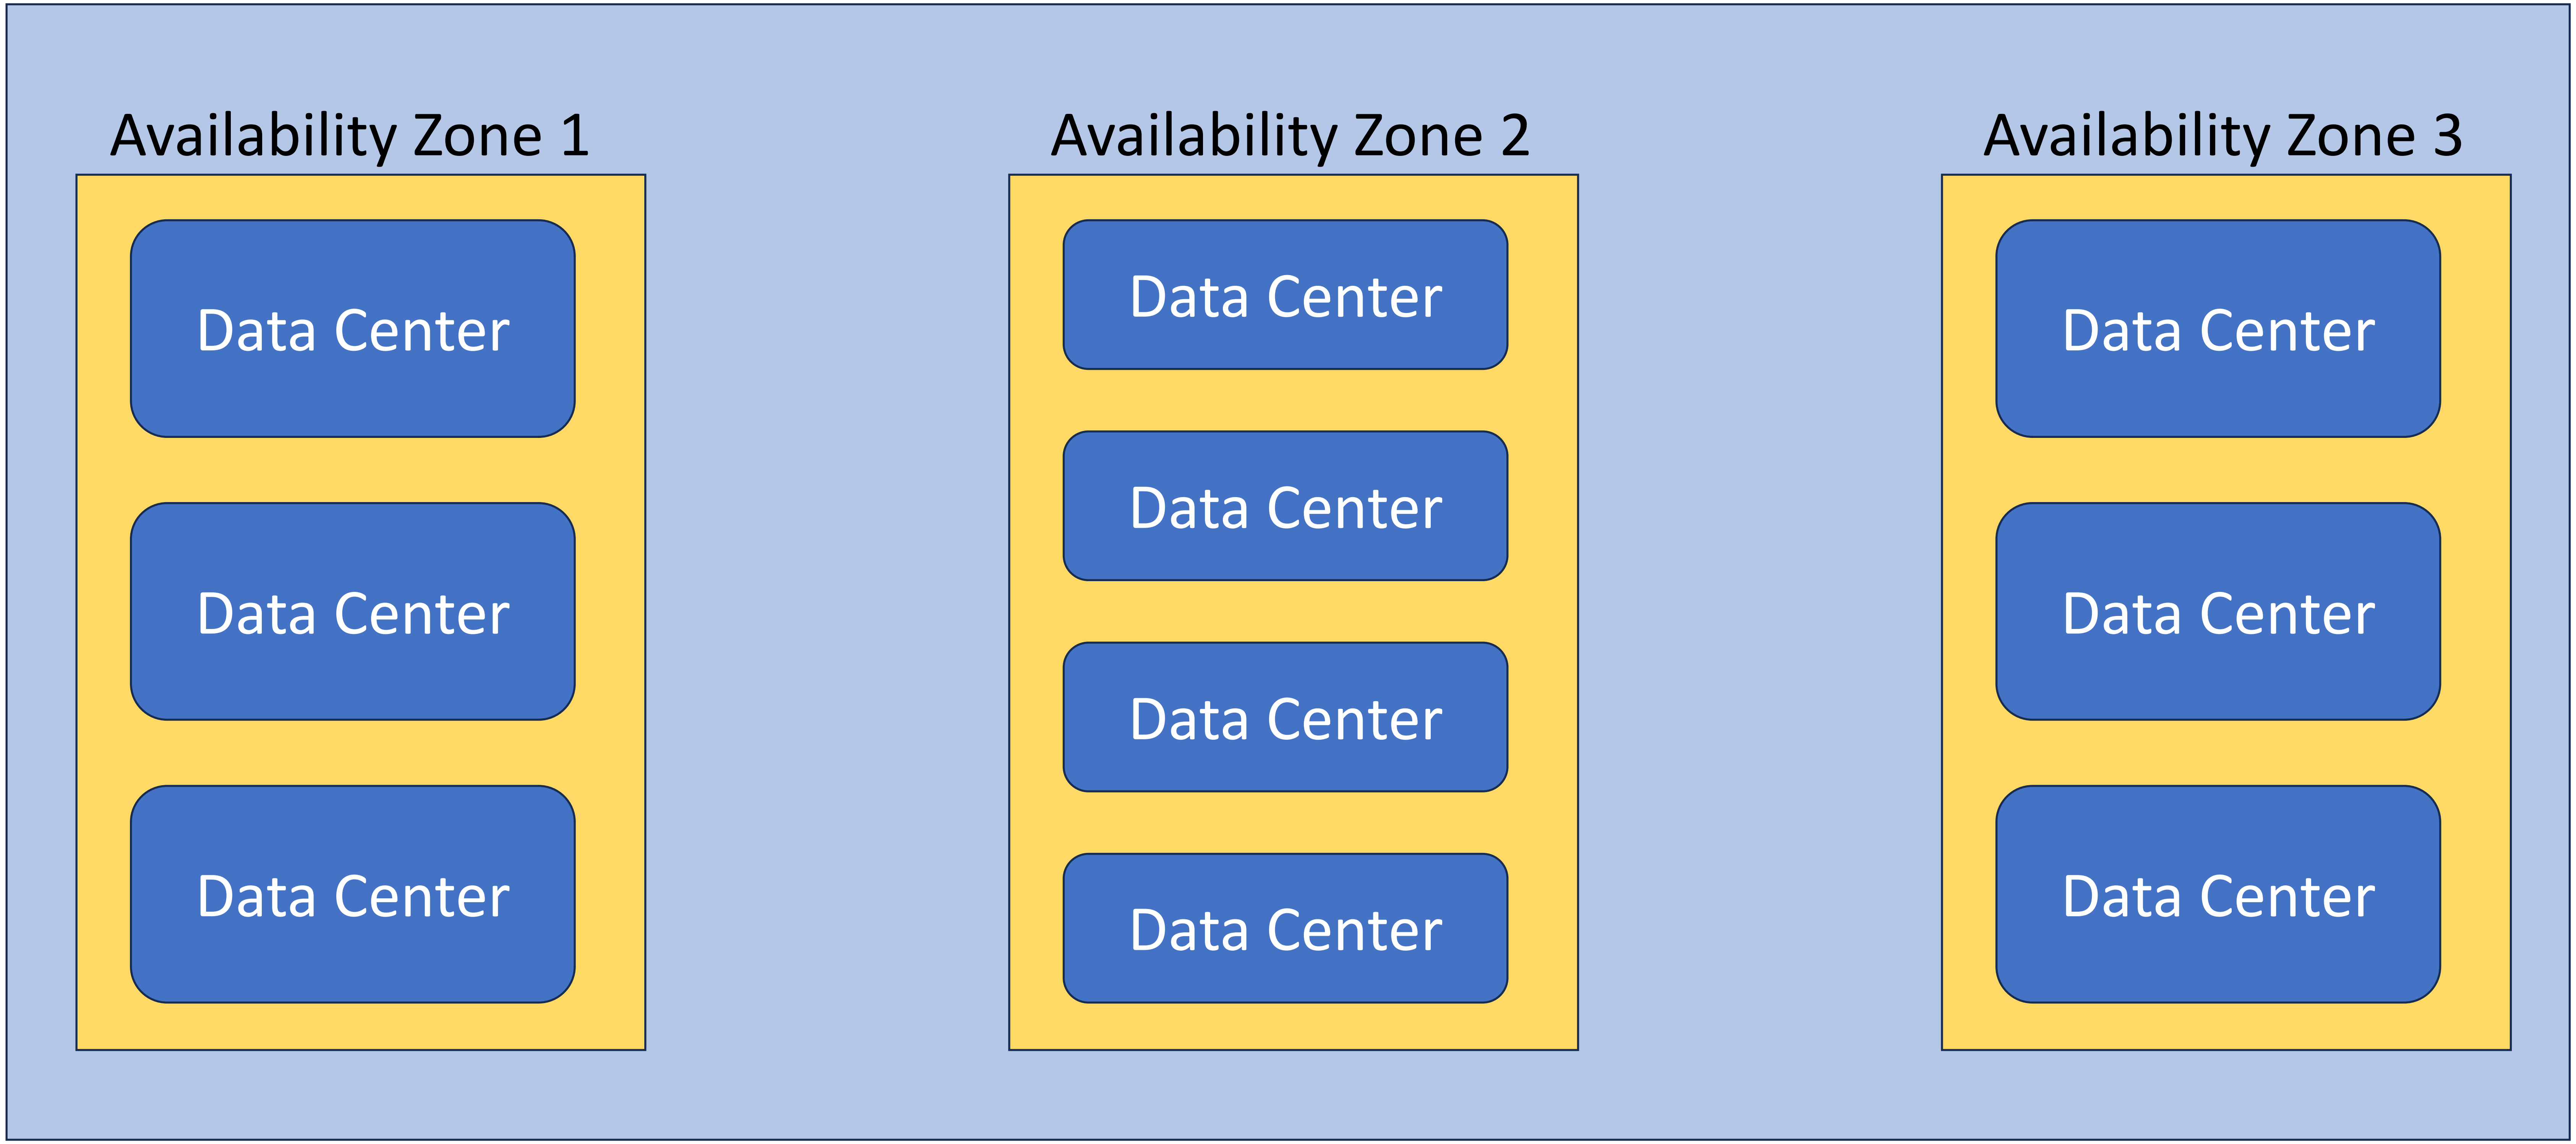

Availability Zone 1

Data Center

Data Center

Data Center

Availability Zone 2

Data Center

Data Center

Data Center

Data Center

Availability Zone 3

Data Center

Data Center

Data Center

Anatomy of an Azure Region

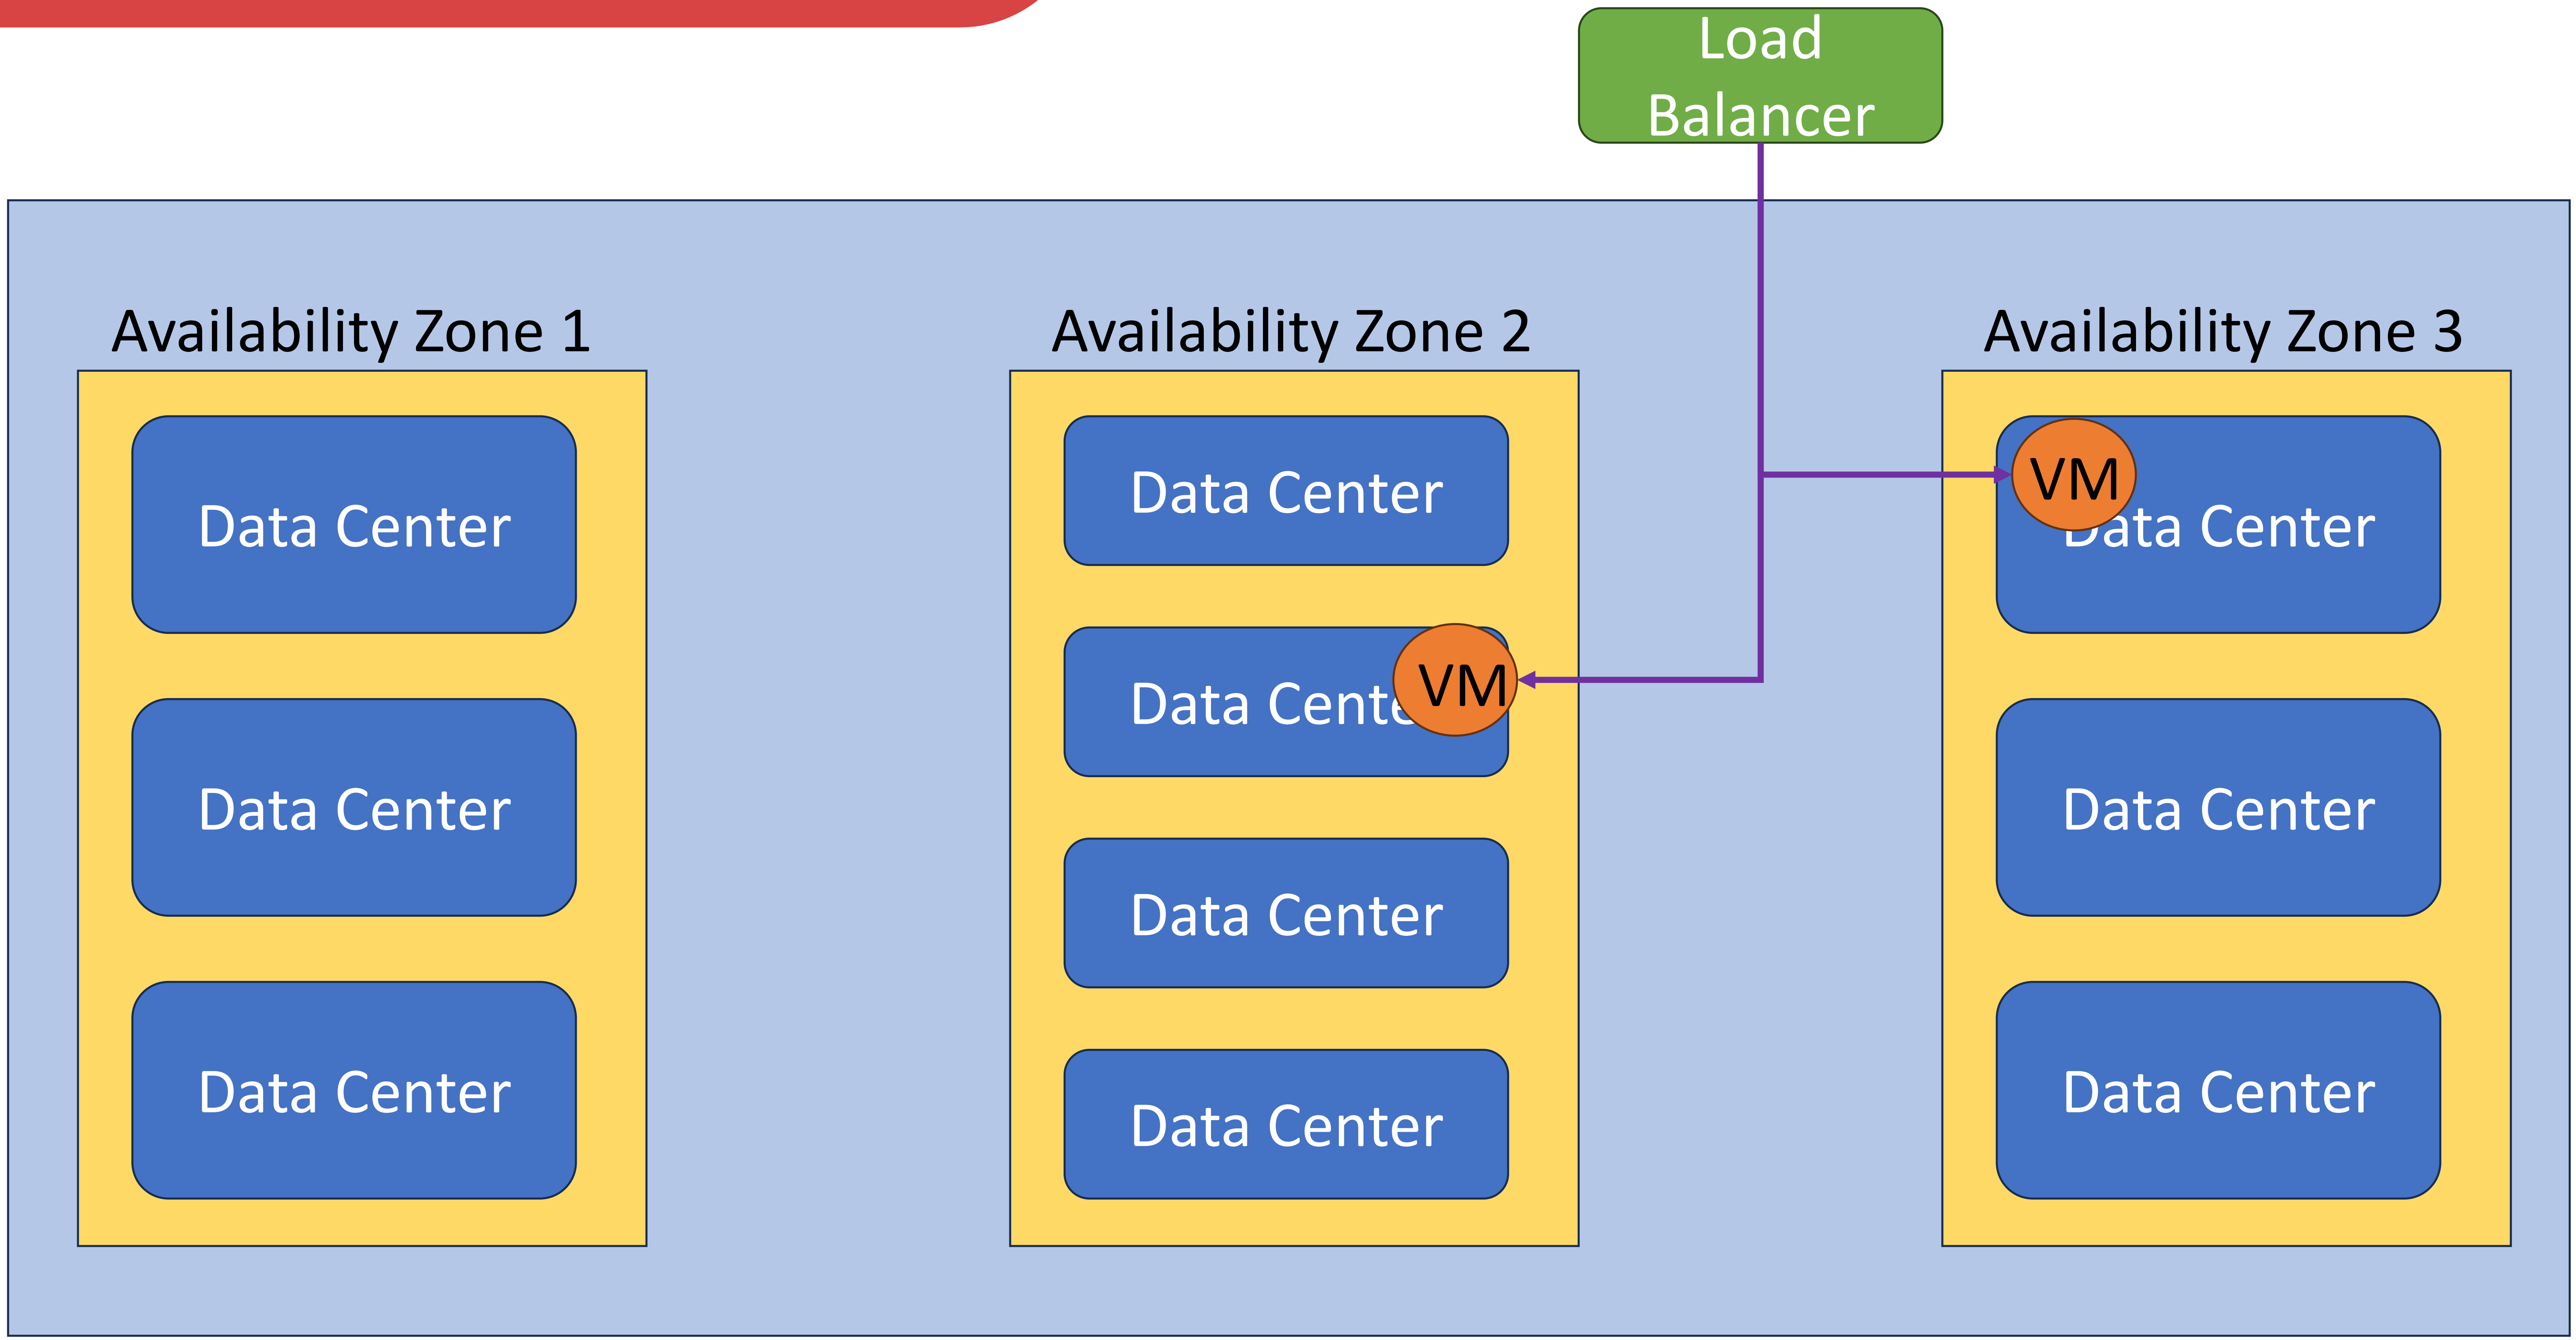

Considerations

Considerations

Security

- Accounts created in other subscriptions / tenant
- Network rules
  - Possibly different IP addresses in DR area (vnet, region, etc)

Service Level Agreements

- Those of the products/services you rely on
- Those of the products/services you offer

RPO / RTO

- Per DB option
- Per custom service
- Cost of being down / unavailable

Data Residency Laws

- DR area must abide by same data residency laws as primary area

Implementing Resiliency

Implementing Resiliency

Storage replication

- Paired region
- RPO due to replication latency
- Cosmos DB geo-replication
- Archiving
    - Geo-Replication
    - Healthcare, Financial Data, etc.

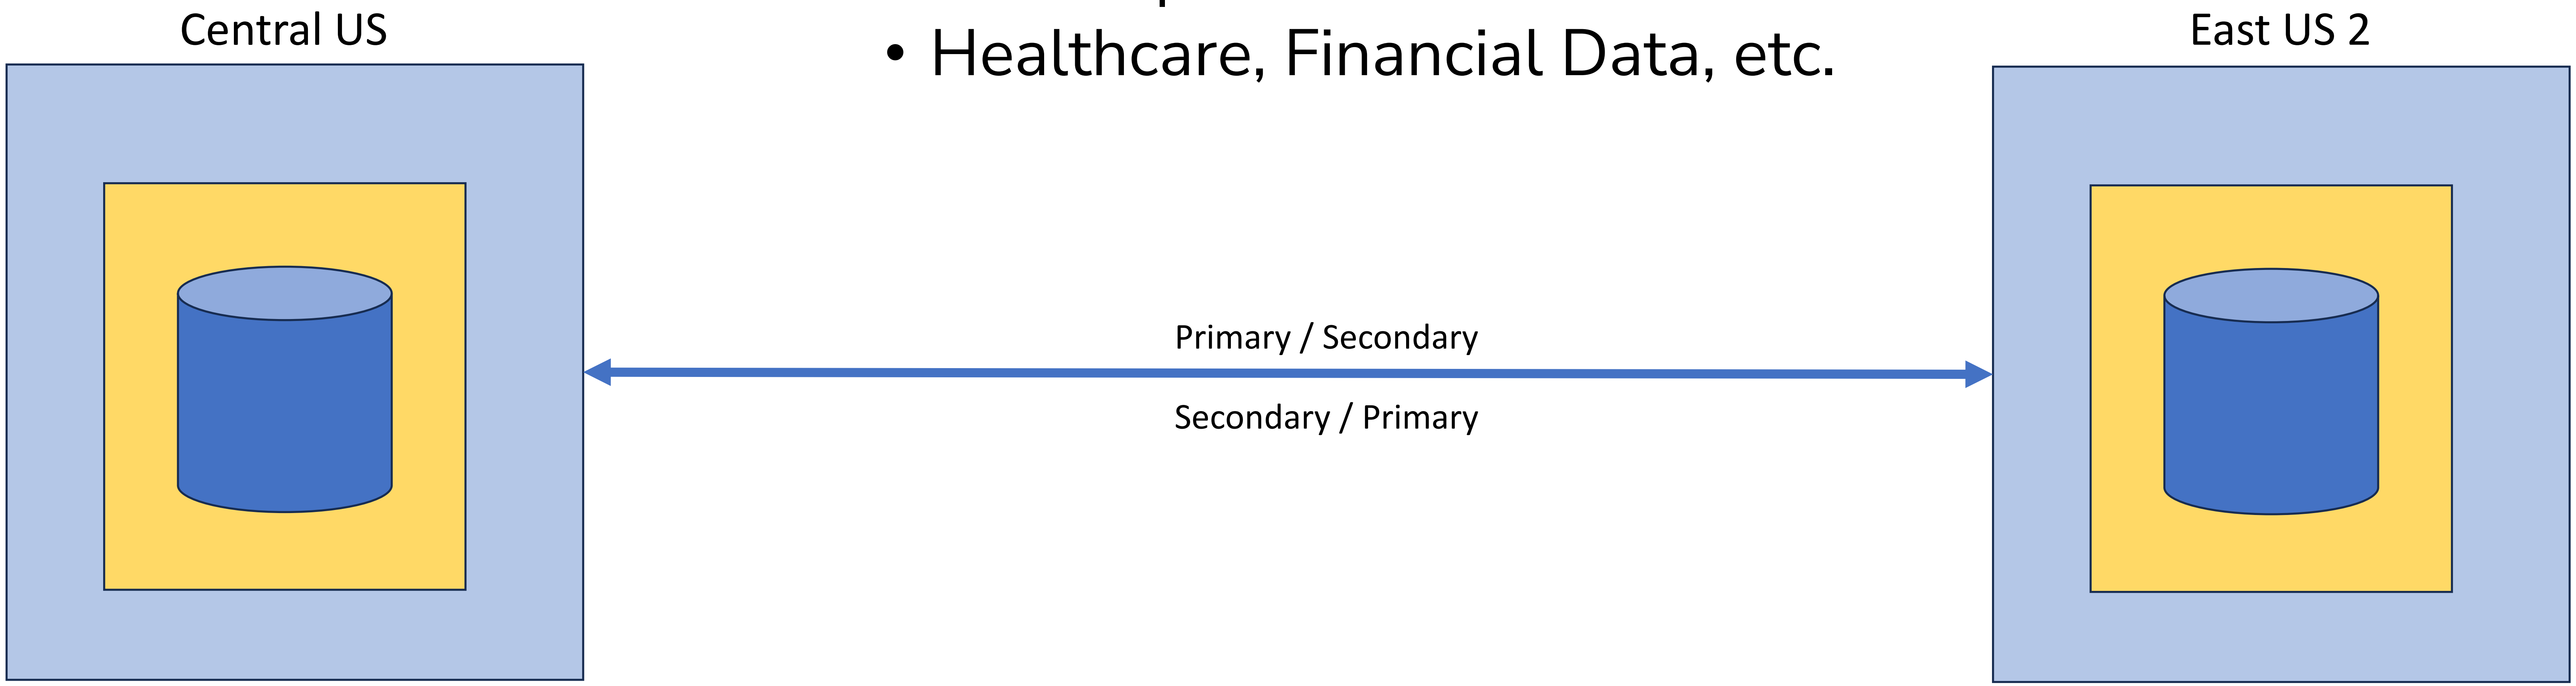

Implementing Resiliency

Network failover / latency

- Azure Front Door
- Azure Traffic Manager
- Azure Application Gateway
- Azure Cross-Region Load Balancer
- Azure Load Balancer
- Azure DNS Zone
- Network Security Group
  - Application Security Groups instead of specific IP addresses
  - Service Tags to allow replication e.g., Storage
- Network Virtual Appliance

Implementing Resiliency

Azure Site Recovery

- Plans
- Orchestration
  - Priority of resources to create (DB before service)
  - On-prem to cloud

Playbook

- Compliance / Regulatory requirements
- Contains runbooks

Runbook

- How to setup a server or configure a load balancer
- Automated as much as possible
- Azure Automation

Implementing Resiliency

(Re)creating resources

- Infrastructure as Code
- Automation / Azure DevOps
- Regional resource capacity

Security

- Access to container registry/images
- Access to vaults
- Accounts on VMs

Implementing Resiliency

Configuration Information

- Azure App Config
  - geo-replicated
  - sensitive info stored in Azure Key Vault
- Azure Key Vault
  - backups

Implementing Resiliency

- Containers
  - Azure Container Registry geo-replication
  - Azure Kubernetes Service
  - Azure Container Apps

Implementing Resiliency

VM

- Images
  - backups
- Configuration (does it know it's a backup? Affects logging)

Compliance / Regulations

- Auditing
- Policies
- Azure Policy

Data

- Azure SQL DB - Active geo-replication
- RPO
- Cosmos DB geo-replication

Implementing Resiliency

Azure Monitoring

- Logging
- Metrics
- Alerting
- Action Groups
- Baseline
  - Know what nominal is
  - Ability to determine issues

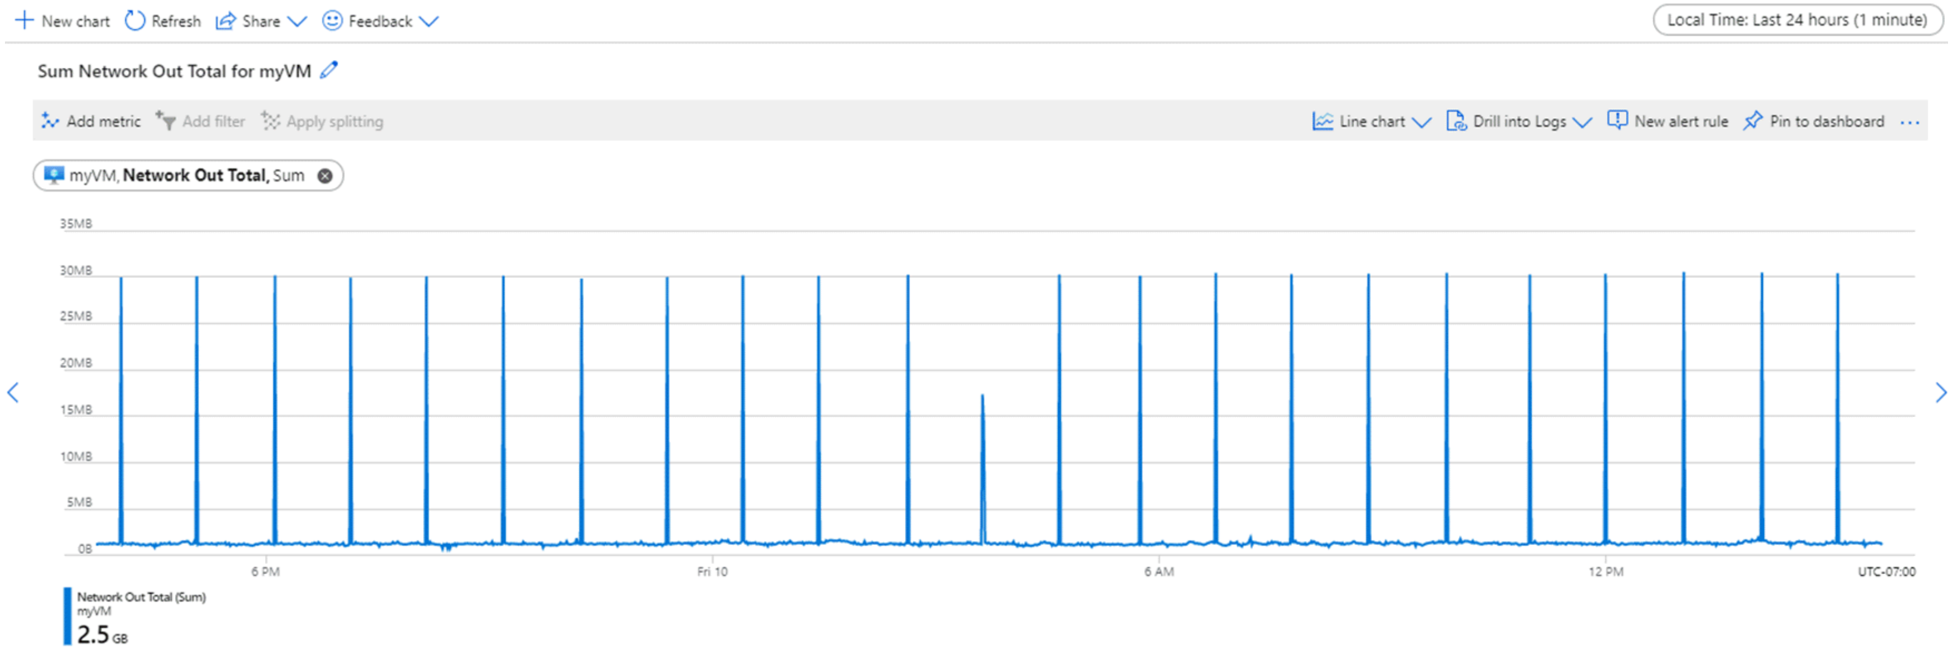

Sum Network Out Total for myVM

myVM, **Network Out Total**, Sum

Testing Your Plans

Testing Your Plans

In 2016, Delta Airlines had a fire in a data center. It was quickly extinguished. But it was discovered that 300 of the 7,000 servers were not wired to backup power source. The impact cost Delta nearly $100M in lost revenue.

Testing Your Plans

A hospital was going to test their power generators they have had for years. They notified many there may be a disruption during the day.

Hours later it was noticed that there was no disruption. When inquired, it was noted that they could not continue test after realizing ...

Testing Your Plans

A hospital was going to test their power generators they have had for years. They notified many there may be a disruption during the day.

Hours later it was noticed that there was no disruption. When inquired, it was noted that they could not continue test after realizing there was no fuel in their state-of-the-art emergency power generators.

Testing Your Plans

Planning

- Multiple teams engaged
- All members know the tools required
- Action Groups are accurate
- Alerts are created
- Consider security breach
- Consider network outage
- Consider not being able to access Azure Portal
- Phased approach
  - Not everything at once

This Photo by Unknown Author is licensed under CC BY-SA-NC

Testing Your Plans

Azure Chaos Studio

- Does real changes to infrastructure, not simulated
- Not available in all Azure regions yet
- Azure Load Testing to generate traffic

Testing Your Plans

Azure products evolve too

- Azure Site Recovery
- Azure Backup
- Azure Migrate
- Azure Resource Mover

Testing Your Plans

Code deployment

- When deploying to primary location
  - Also send to App Services, VMs, etc in secondary location

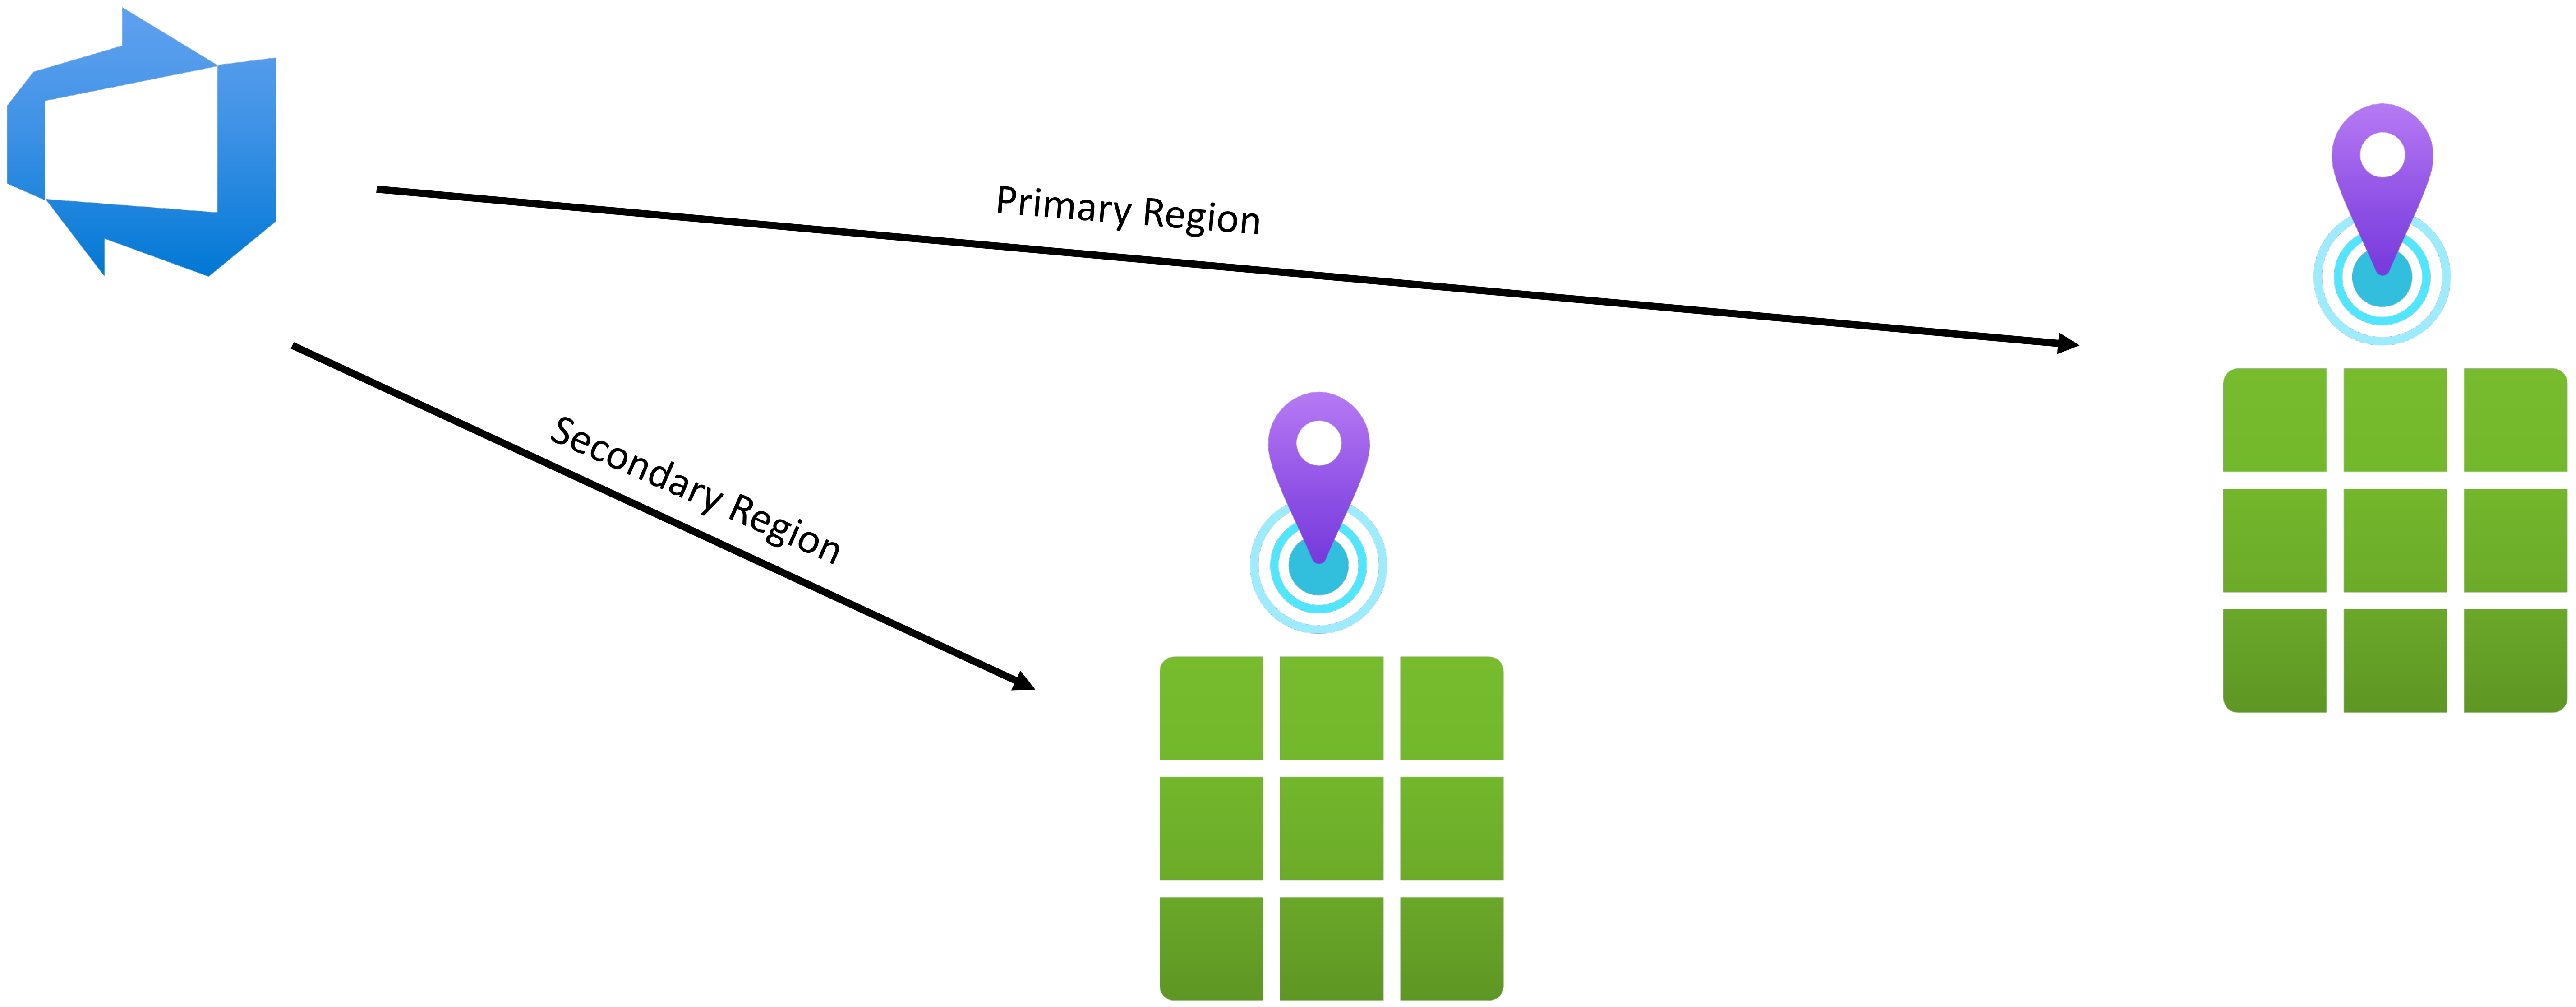

Testing Your Plans

Challenge All Assumptions!

This Photo by Unknown Author is licensed under CC BY

Challenge All Assumptions!

Knight Capital, a trader with U.S. equities, etc., ultimate went under due to a software issue. They reused a feature flag when deploying a new large feature. The issue was that the new code was only deployed to 7 of the 8 servers. The 8th server executed old code when the feature flag was enabled.

The code on the 8th server purchased around $7 billion in stocks the first hour of trading the day of Aug 1st, 2012. The following summer, Knight Capital was acquired by a rival company at a loss of over $400 million.

Evolution

Biz changes

Architecture changes

Testing needs to change

Resiliency and DR plans needs to change

This Photo by Unknown Author is licensed under CC BY-SA

Know Your Assumptions

Examples

- All the code is in a repo and apps can be recreated at any time
- The repo is globally accessible all the time
- Our DNS is safe and unable to be hacked
- The scripts that built our infrastructure will always work again.
- It's cheaper to apologize to customers than to build and test DR plans
- Our DR plans don't need to be tested or can't be tested
- All Azure services are always available
- Those who need to know, know
- You told your boss, so they told their boss
- Infinite scaling
- The backups are good (story of empty backup)

META problem, DNS issue; peeps couldn’t get into server room to reboot router.

The ArchitectNow Planning Approach

high level overview

The ArchitectNow Planning Approach

Analyze and understand what you currently have

The ArchitectNow Planning Approach

Analyze and understand what you currently have

Evaluate the criticality of each item

- (can be gone, can be down…..needs as much availability as possible)

The ArchitectNow Planning Approach

Analyze and understand what you currently have

Evaluate the criticality of each item

- (can be gone, can be down…..needs as much availability as possible)

Evaluate order of importance

- Grouping of resources per business need
- Order of importance
  - e.g., Processing payments before Scheduling

The ArchitectNow Planning Approach

Analyze and understand what you currently have

Evaluate the criticality of each item

- (can be gone, can be down…..needs as much availability as possible)

Evaluate order of importance

- Grouping of resources per business need
- Order of importance
  - e.g., Processing payments before Scheduling

Evaluate ability to replicate

- Some architectures require changes before DR is possible

The ArchitectNow Planning Approach

Analyze and understand what you currently have

Evaluate the criticality of each item

- (can be gone, can be down…..needs as much availability as possible)

Evaluate order of importance

- Grouping of resources per business need
- Order of importance
  - e.g., Processing payments before Scheduling

Evaluate ability to replicate

- Some architectures require changes before DR is possible

Are Azure resources available in other regions?

- Capacity issues affect all cloud providers

The ArchitectNow Planning Approach

Analyze and understand what you currently have

Evaluate the criticality of each item

- (can be gone, can be down…..needs as much availability as possible)

Evaluate order of importance

- Grouping of resources per business need
- Order of importance
  - e.g., Processing payments before Scheduling

Evaluate ability to replicate

- Some architectures require changes before DR is possible

Are Azure resources available in other regions?

- Capacity issues affect all cloud providers

Evaluate product availability

- Older version may not be available
- Product upgrade may be required (MySQL on VM, etc)

The ArchitectNow Planning Approach

Analyze and understand what you currently have

Evaluate the criticality of each item

- (can be gone, can be down…..needs as much availability as possible)

Evaluate order of importance

- Grouping of resources per business need
- Order of importance
  - e.g., Processing payments before Scheduling

Evaluate ability to replicate

- Some architectures require changes before DR is possible

Are Azure resources available in other regions?

- Capacity issues affect all cloud providers

Evaluate product availability

- Older version may not be available
- Product upgrade may be required (MySQL on VM, etc)

Target region network

- IP address cannot overlap if peered/VPN

The ArchitectNow Planning Approach

How can ArchitectNow help you?

Need an Azure audit? We can help.

The ArchitectNow Planning Approach

Whether launching new Cloud or mobile apps or modernizing your legacy platforms we can help you identify the best options and work with you on bringing those ideas to life. To get the ball rolling, reach out and tell us a bit about your needs and we can start identifying solutions.

There is no risk, and we can quickly get to the point of providing initial ideas along with rough estimates of the costs and implementation times required with various recommendations.

ArchitectNow

Thank you!

info@architectnow.net
www.architectnow.net