



High Availability and Disaster Recovery

Cherry Lin, Jonathan Quinn

An abstract 3D architectural graphic on the right side of the slide. It features various geometric shapes and planes in shades of blue, orange, and green. Some surfaces are covered with white topographic contour lines. The overall composition is dynamic and modern.

**GIS
INSPIRING
WHAT'S
NEXT**

Managing the Twin Risks to your Operations



Data Loss

Down Time

The Three Approaches



Backups

Snapshot
Ability to go back in time

High Availability

No single point of failure
Machine redundancy

Disaster Recovery

Geographic Redundancy

No single point of failure
Environment redundancy

Choosing Between Them



Complementary

Build On Each Other

Cost and Capability

Backup and Restore

The image features a dark teal background with a white title 'Backup and Restore' centered on the left. The right side is decorated with abstract, colorful geometric shapes in shades of orange, yellow, and light blue, creating a modern, digital aesthetic.

Backups are....



Simple

Highly Effective

Not Disruptive

Under appreciated

ArcGIS Enterprise Backups – WebGIS DR Tool



What the tool backs up

Settings

(Portal, Server, Data Store)

Portal Content

Services

ArcGIS Data Store Data

(relational, scene tiles)

ArcGIS Enterprise Backups – WebGIS DR Tool




What the tool doesn't backup

EGDB or file based data

Traditional cache tiles

How to Backup an ArcGIS Enterprise deployment



Web GIS DR Tool

Property File

- Location
- Portal URL
- Credentials
- Scene Cache?

Automate

WebGIS DR Properties

- **Questions:**
 - **Where is your content, (file system, bucket in S3, container in Azure)**
 - **Where do you want your backup stored, (file system, bucket in s3, container in Azure)**
- **Minimum parameters:**
 - **SHARED_LOCATION** = where each backup will be staged
 - **BACKUP_STORE_PROVIDER** = where to store the backup (file system, or cloud)
 - **PORTAL_ADMIN_URL** = URL to connect to the portal
 - **PORTAL_ADMIN_USERNAME** = administrator's username
 - **PORTAL_ADMIN_PASSWORD** = administrator's password
 - **BACKUP_RESTORE_MODE** = defines if a full or incremental backup will be run

WebGIS DR Tool – Usage

- **Backup**

- Component backups run concurrently
- No downtime while exporting
- **Sample syntax**

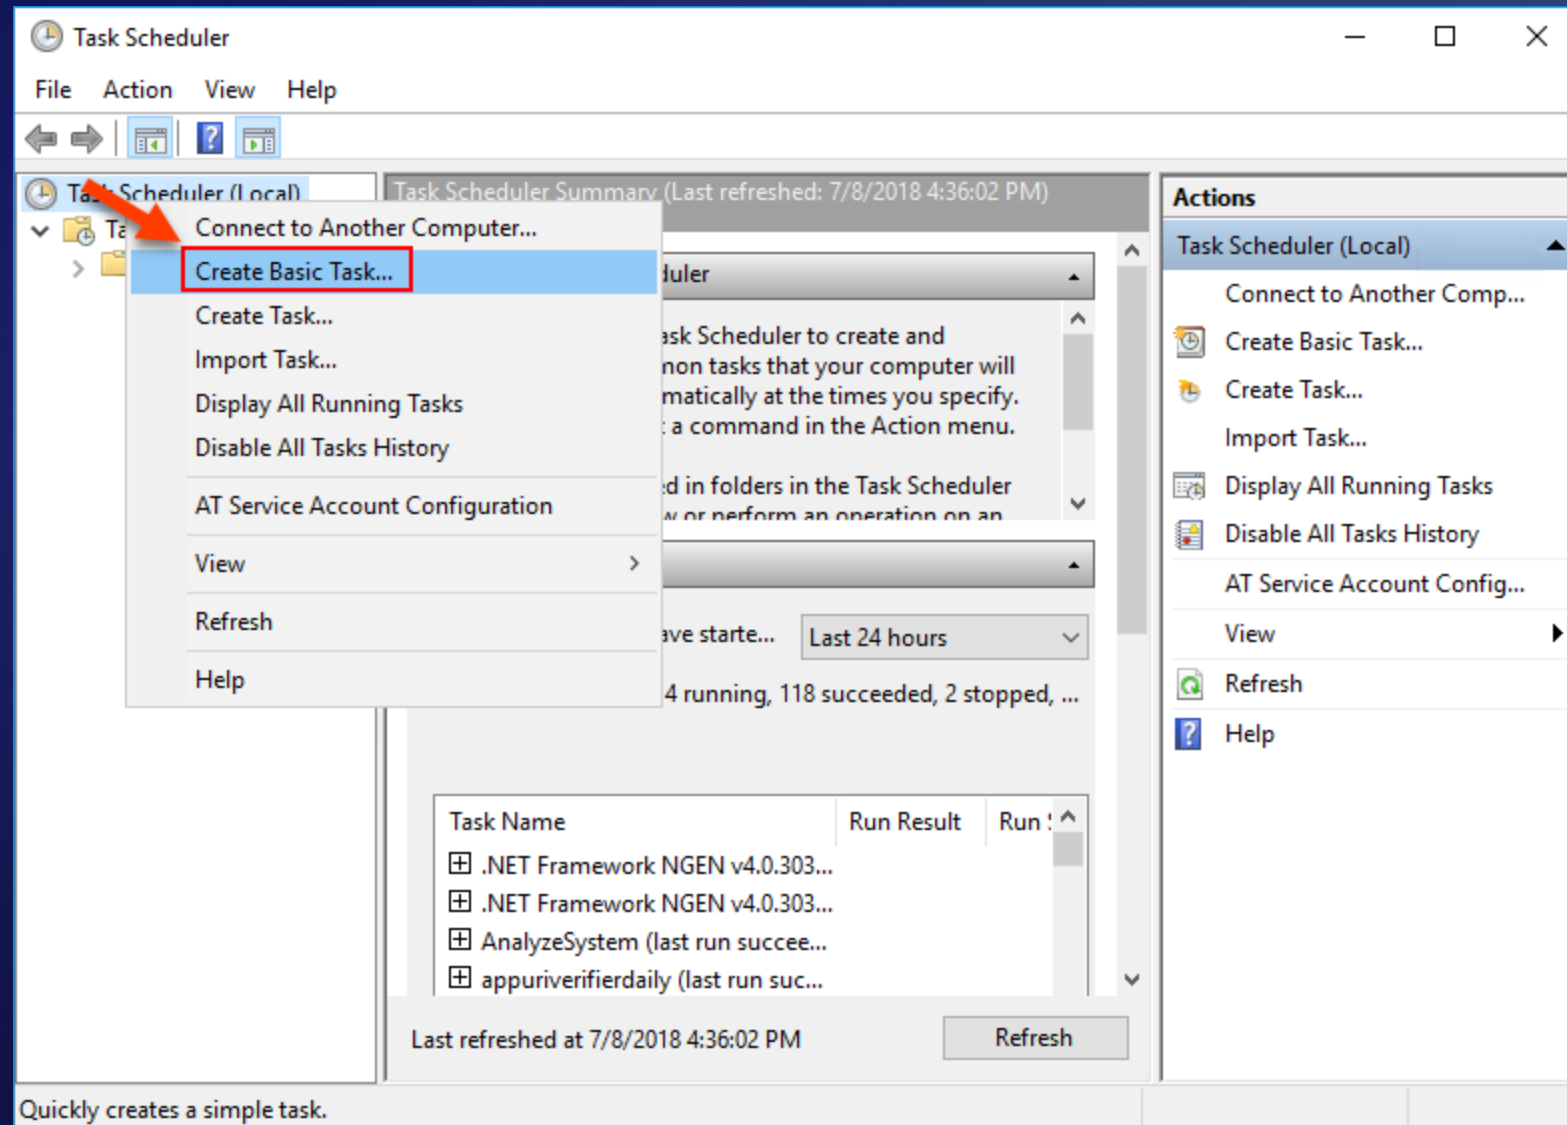
```
C:\Program Files\ArcGIS\Portal\tools\webgisdr>webgisdr.bat -e -f webgisdr.properties
```

- **Restore**

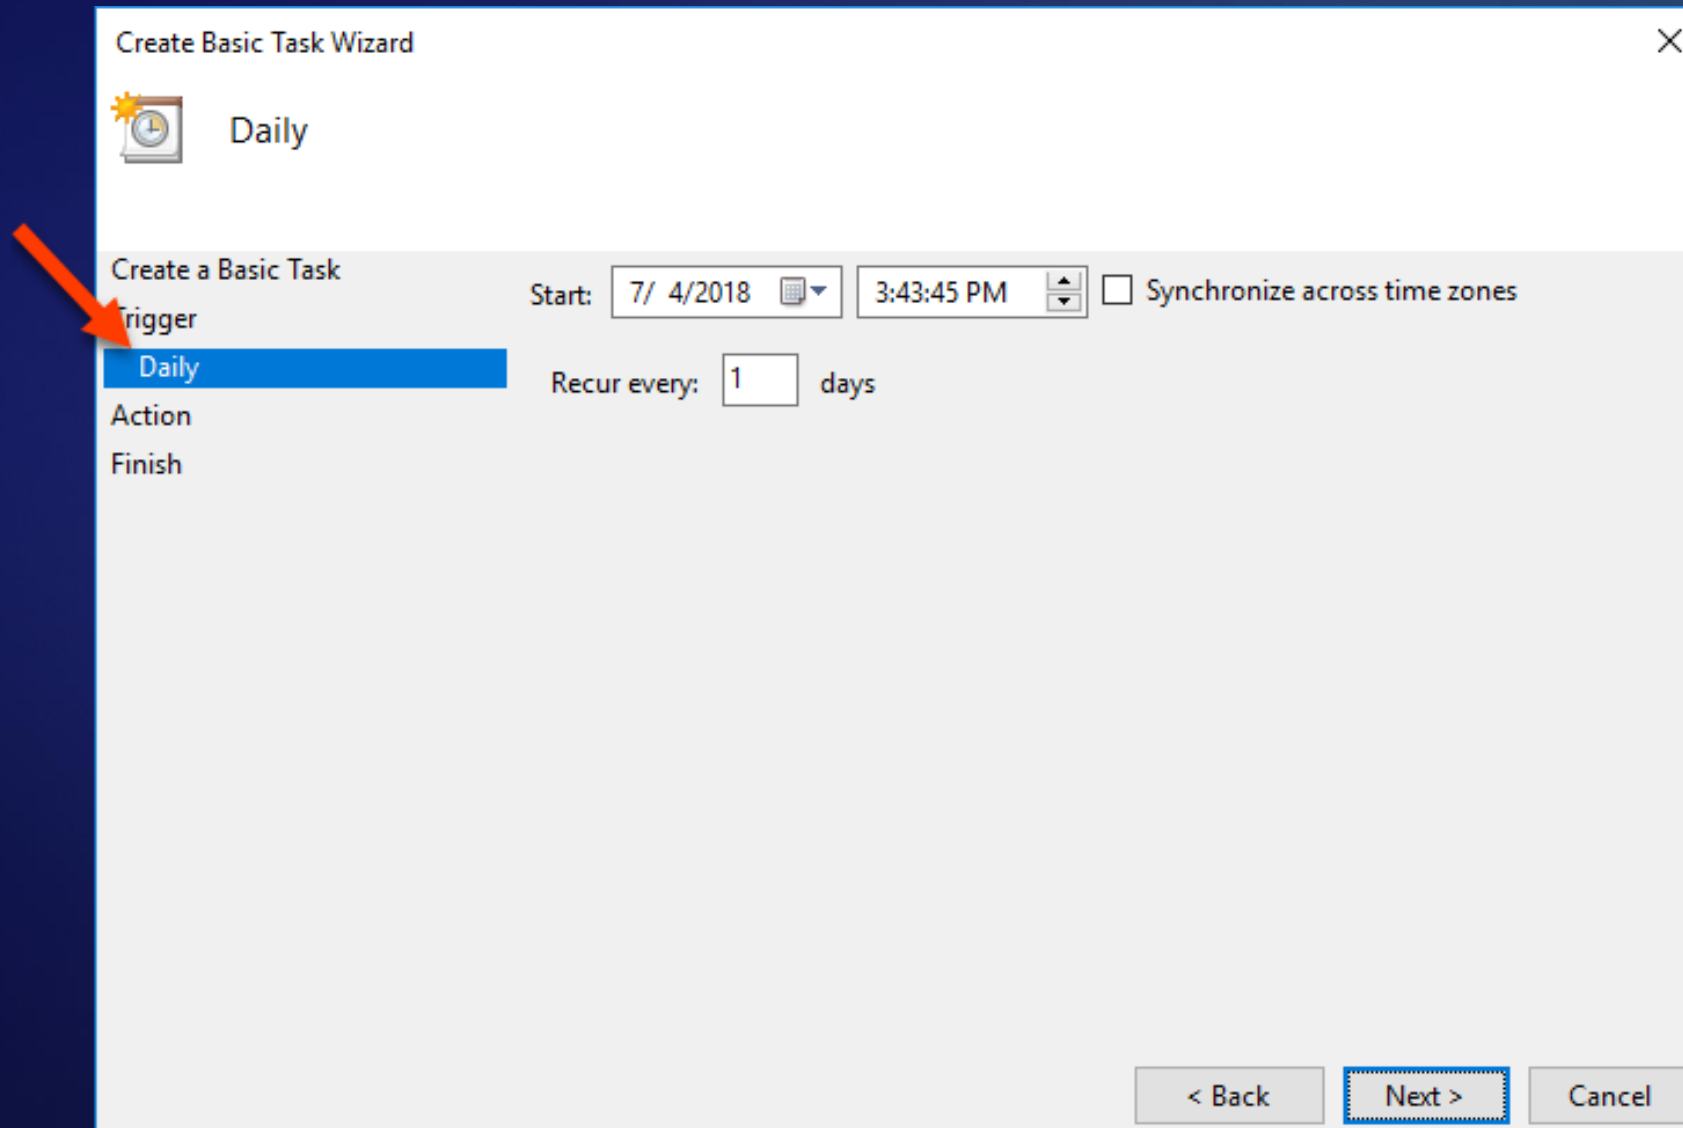
- **Runs sequentially**
 - Data Store → Server → Portal
- **Downtime while restoring**
- **Sample syntax**

```
C:\Program Files\ArcGIS\Portal\tools\webgisdr>webgisdr.bat -i -f webgisdr.properties
```

Scheduling ArcGIS Enterprise backups - Windows



Scheduling ArcGIS Enterprise backups - Windows



Scheduling ArcGIS Enterprise backups - Windows

Create Basic Task Wizard

Start a Program

Create a Basic Task

Trigger

Daily

Action

Start a Program

Finish

Program/script:

"C:\Program Files\ArcGIS\Portal\tools\webgisdr\webgisdr.bat" Browse...

Add arguments (optional):

-e -f "C:\Program Files\A


Start in (optional):

-e -f "C:\Program Files\ArcGIS\Portal\tools\webgisdr\webgisdr.properties

< Back Next > Cancel

Scheduling ArcGIS Enterprise backups - Windows

Create Basic Task Wizard

 Summary

Create a Basic Task

Trigger
Daily
Action
Start a Program
Finish

Name: Backups

Description:

Trigger: Daily; At 3:43 PM every day

Action: Start a program; "C:\Program Files\ArcGIS\Portal\tools\webgisdr\webgisdr.b

Open the Properties dialog for this task when I click Finish

When you click Finish, the new task will be created and added to your Windows schedule.

< Back Finish Cancel

Scheduling ArcGIS Enterprise backups - Linux

- Creating a cronjob:

```
[ags@wilson ~]$ crontab -e
```

- Cronjob syntax:

```
* * * * * < command >
```

Minute
Hour
Day of the month
Month
Weekday

Examples:

Run the WebGIS DR Tool at 12:00:00 AM every day:

```
0 0 * * * /data/arcgis/portal/tools/webgisdr/webgisdr.sh -e -f /data/arcgis/portal/tools/webgisdr/webgisdr.properties
```

Run the tool every 12 hours every day starting at 12:00:00 AM:

```
0 */12 * * * /data/arcgis/portal/tools/webgisdr/webgisdr.sh -e -f /data/arcgis/portal/tools/webgisdr/webgisdr.properties
```

High Availability

The image features a dark teal background with a subtle grid pattern. The text 'High Availability' is centered in a bold, white, sans-serif font. The corners are decorated with abstract, colorful geometric shapes in shades of orange, yellow, and light blue, creating a modern, digital aesthetic.

Overview

- **What is High Availability**
- **ArcGIS Enterprise High Availability**
 - Components
 - Upgrade
- **Other factors for High Availability**

High Availability (HA)

- **Definition:**
 - A system or component that is continuously operational for a desirably long length of time. Availability can be measured relative to "100% operational" or "never failing." (SLAs)
- **Shorter down time costs more**
- **Elimination of single points of failure.**
- **Availability of a system depends on the availability of all components**

ArcGIS Enterprise



Portal

GIS Services

Hosted Feature and Tile Data



Portal for ArcGIS

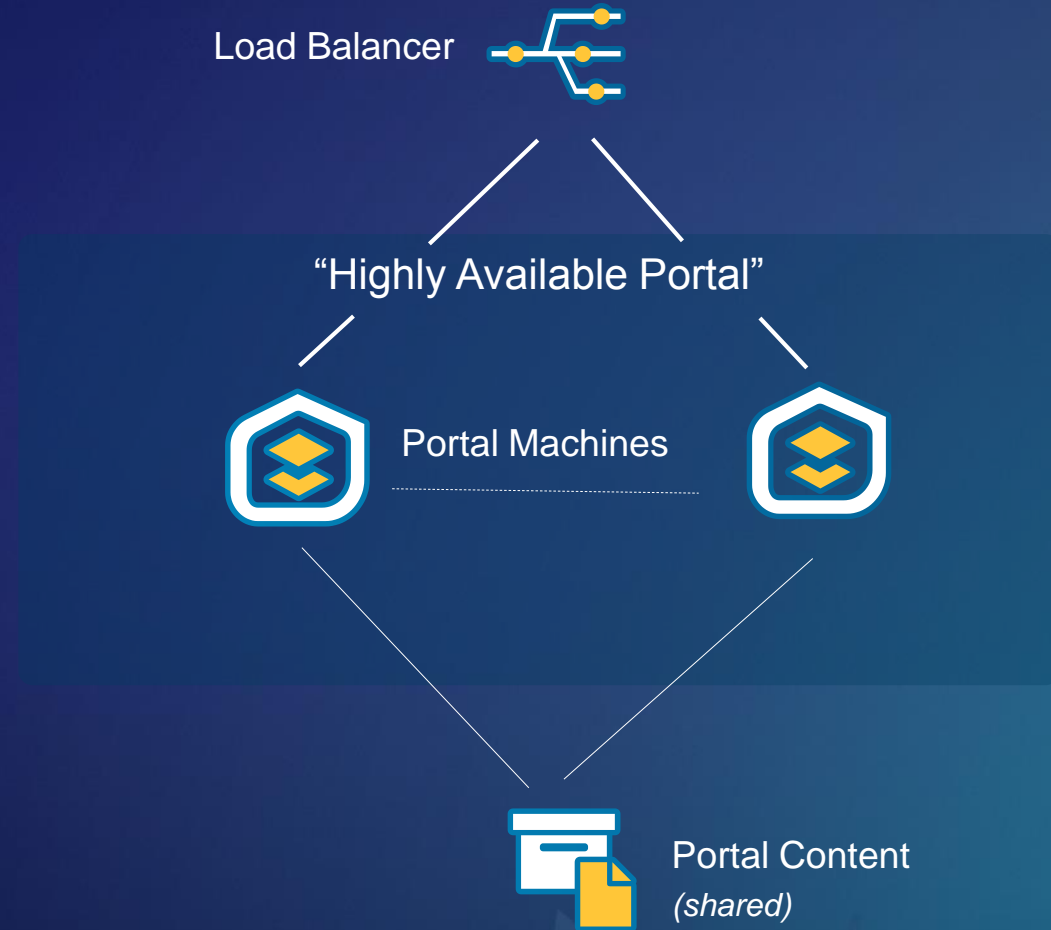


ArcGIS Server



ArcGIS Data Store

Portal for ArcGIS : High Available Deployment



Highly Available Portal

- Two Portal machines
- Both Portal machines take requests
- Internally, there is a difference between the two machines' role:
 - Primary
 - Standby
- Behaves a little bit differently :
 - Standby machine is down (or Portal service stops)
 - **No interruption**
 - Primary is down (or Portal service stops)
 - **Under 30 seconds before standby is promoted**
 - **Improvement at 10.6.1 from typically a few minutes**

Machines:

- SECONDARY.CHERRY.COM **standby** [status](#)
- PRIMARY.CHERRY.COM **primary** [status](#)

Supported Operations: [unregister](#)

Supported Interfaces: [REST](#)


Portal for ArcGIS: Load Balancing Options

ArcGIS Web Adaptor



- Provided by Esri
- Web-Tier Authentication
- Availability dependent on web servers

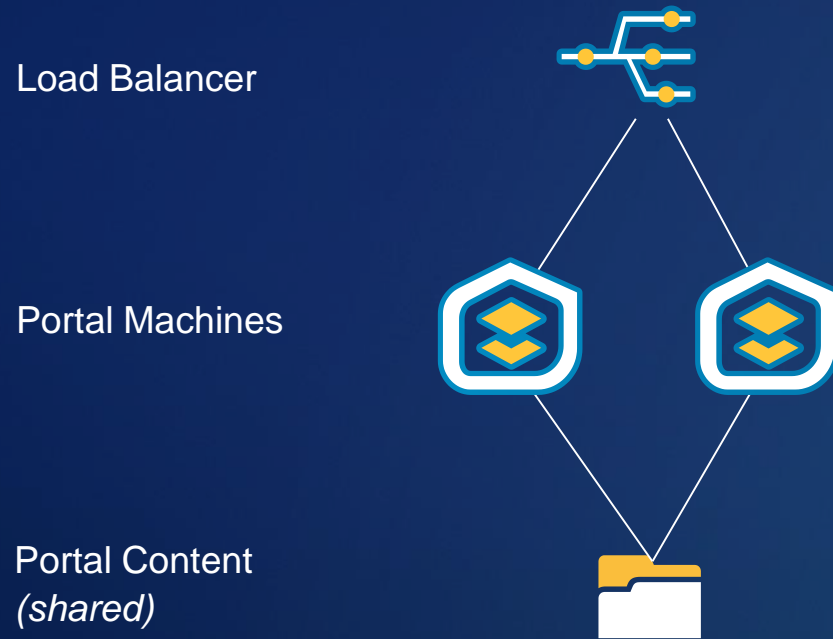
3rd Party Load Balancer



- Not provided by Esri
- Typically already fault tolerant

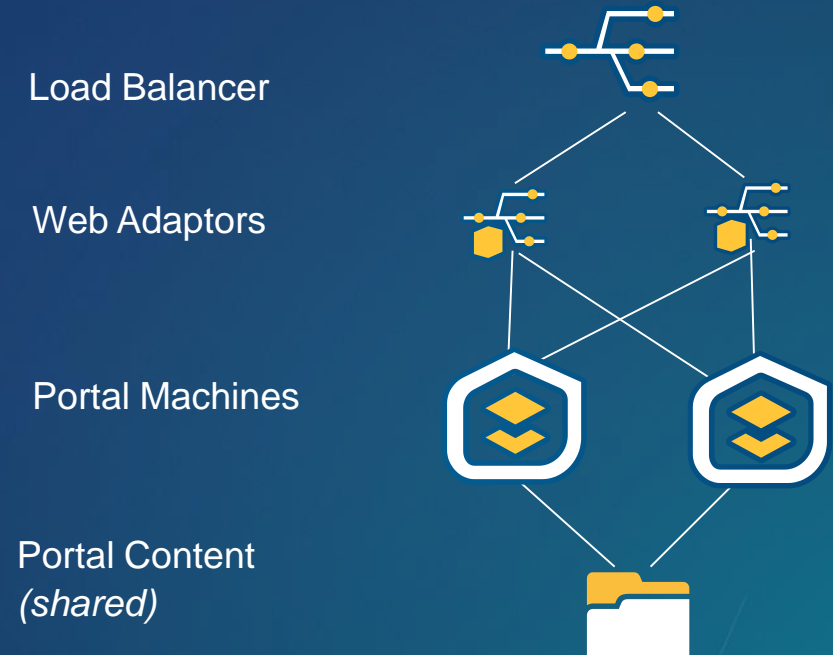
Portal for ArcGIS: High Availability Deployment Patterns

HA Portal with
Load Balancer



- Simpler
- Need certain settings on LB
- Doesn't support Web Tier Authentication

HA Portal with
Load Balancer & Web Adaptors



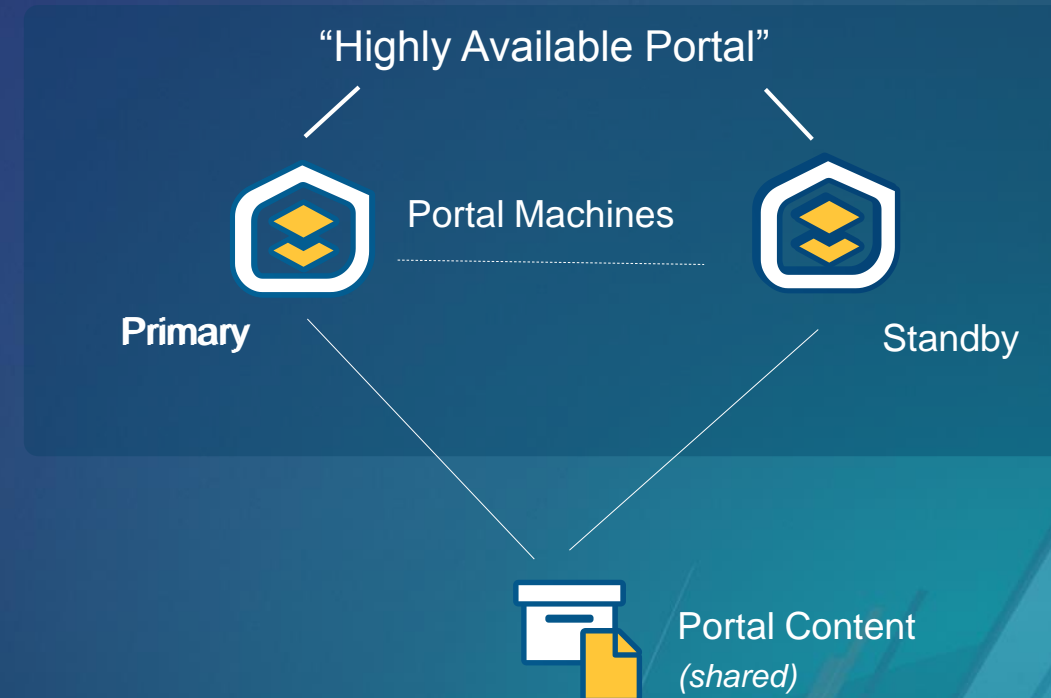
- More complex
- Web Tier Authentication

Portal for ArcGIS: Health Check

- **Provided by Portal for ArcGIS**
 - <https://<webadaptor machine>.domain.com/<context>/portaladmin/healthCheck>
 - <https://<machine>.domain.com:7443/arcgis/portaladmin/healthCheck>
- **Check if Portal is ready to take request. Not individual component, e.g. service, item, etc.**
- **Or your own customized health check**

Upgrade High Availability Portal for ArcGIS

- **There is downtime**
 - Plan
 - Practice
- **Make a backup**
- **Recommended steps**
 - Run setup on STANDBY
 - Run setup on PRIMARY
 - Run upgrade on PRIMARY
- **No need to un-register standby (new at 10.6.1)**
- **Similar other steps as standalone Portal**



Portal for ArcGIS: Key Considerations for HA

- **Two Portal machines**
 - **Primary**
 - **Standby**
- **Highly Available Load Balancer**
 - **Web Tier Authentication**
 - **No single Web Adaptor**
- **Health Check provided for Portal for ArcGIS**
- **Highly Available shared content store**
- **Upgrade: Downtime & steps in order**

ArcGIS Enterprise



Portal

GIS Services



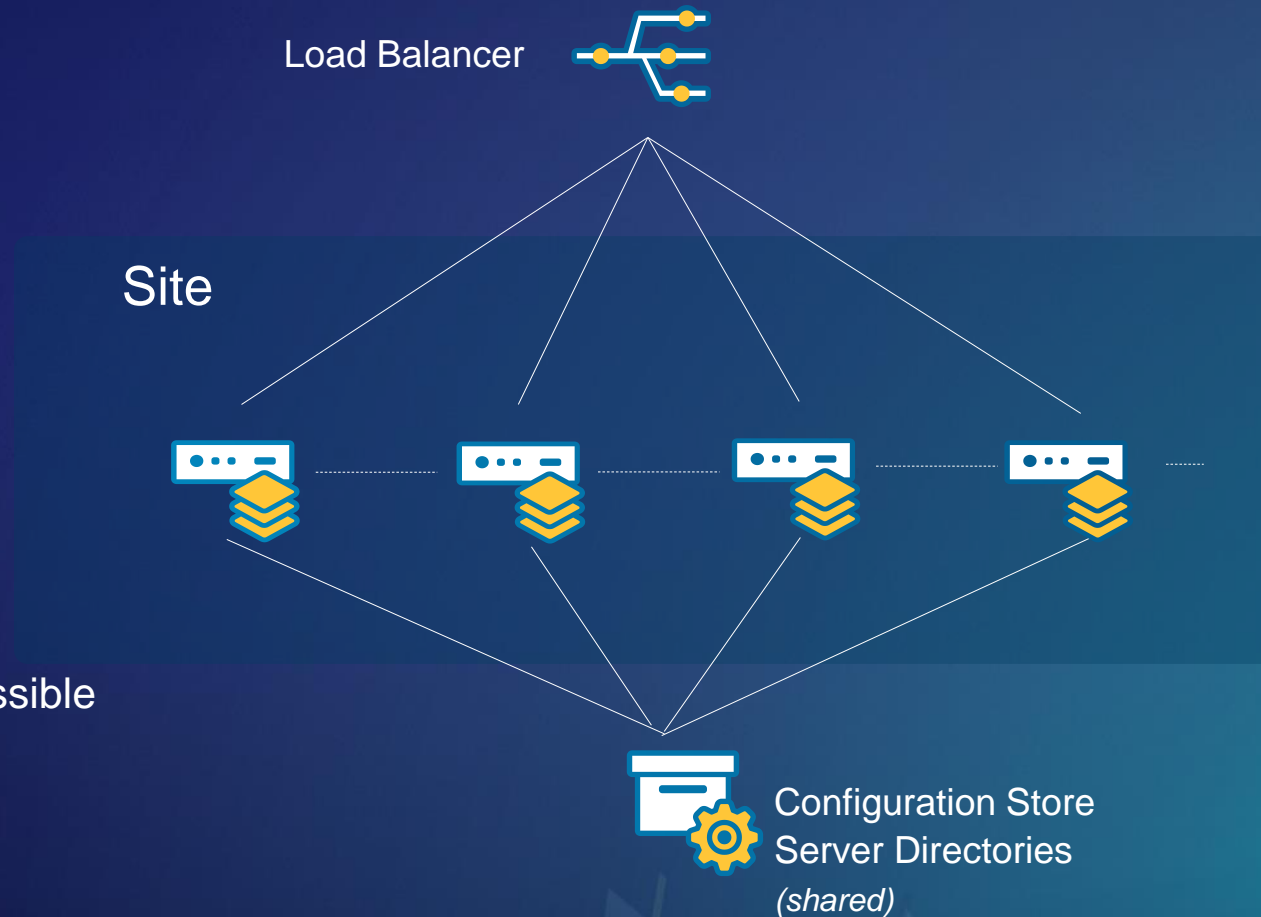
Portal for ArcGIS



ArcGIS Server

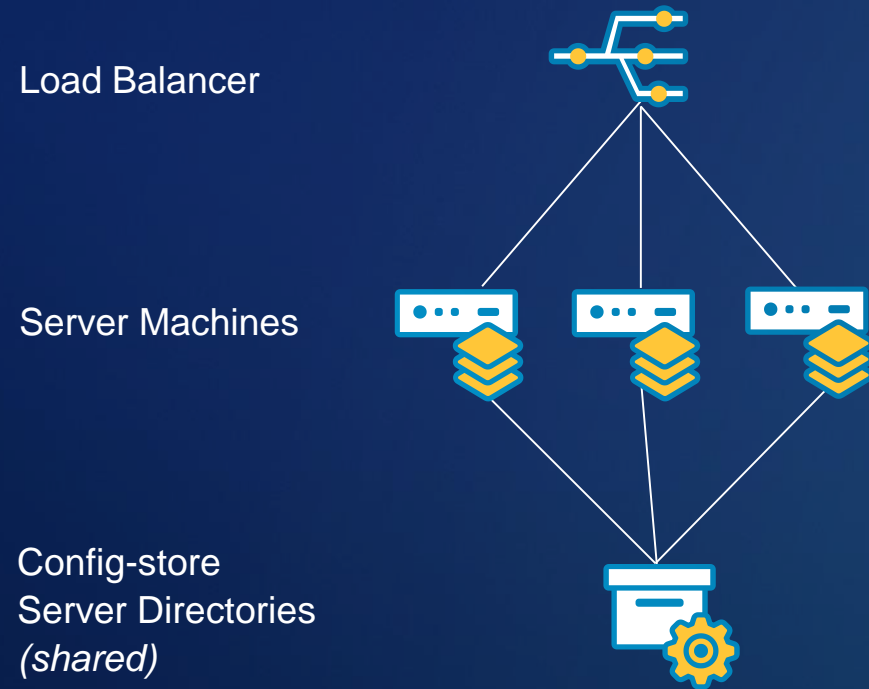
ArcGIS Server: Multiple-Machine Architecture

- Multiple machines
- Identical Roles
- No interruption when any machine is down
- The config-store and server directories need to be accessible to all machines.




ArcGIS Server: High Availability Deployment Patterns

Server Site with
Load Balancer



Server Site with
Load Balancer & Web Adaptors




ArcGIS Server: Health Check

- **Provided by ArcGIS Server**
 - <https://<.....domain.com>/<context>/rest/info/healthcheck>
 - <https://<machine>.domain.com:6443/arcgis/rest/info/healthcheck>
- **Server level health check. Not checking service.**
- **Or your own customized health check**


Upgrade multi-machine ArcGIS Server

- **Install and Upgrade**
- **Same on all machines**
- **Downtime for upgrade one machine**



Portal for ArcGIS and ArcGIS Server: Federation


Portalurl:443



privatePortalurl:7443



Services URL:443



Administrative URL:6443

Add ArcGIS Server

Enter the URLs for accessing and administering your ArcGIS Server site. Also enter credentials for an administrator of the ArcGIS Server site.

Services URL:
Example: https://webadaptor.domain.com/arcgis

Administration URL:
Example: https://gisserver.domain.com:6443/arcgis


Username:

Password:

```
{  
  "portalUrl": "https://webgistesting.net/portal",  
  "privatePortalUrl": "https://webgistesting.net/portal",  
  "portalSecretKey": "29f019ca6ff745aeace5d26bdfc32ca2",  
  "portalMode": "ARCGIS_PORTAL_FEDERATION",  
  "serverId": "7jhSwDZJ6Q6kIurK",  
  "serverUrl": "https://webgistesting.net/server"  
}
```

Portal for ArcGIS and ArcGIS Server: Federation

Portalurl:443



Load Balancer

443

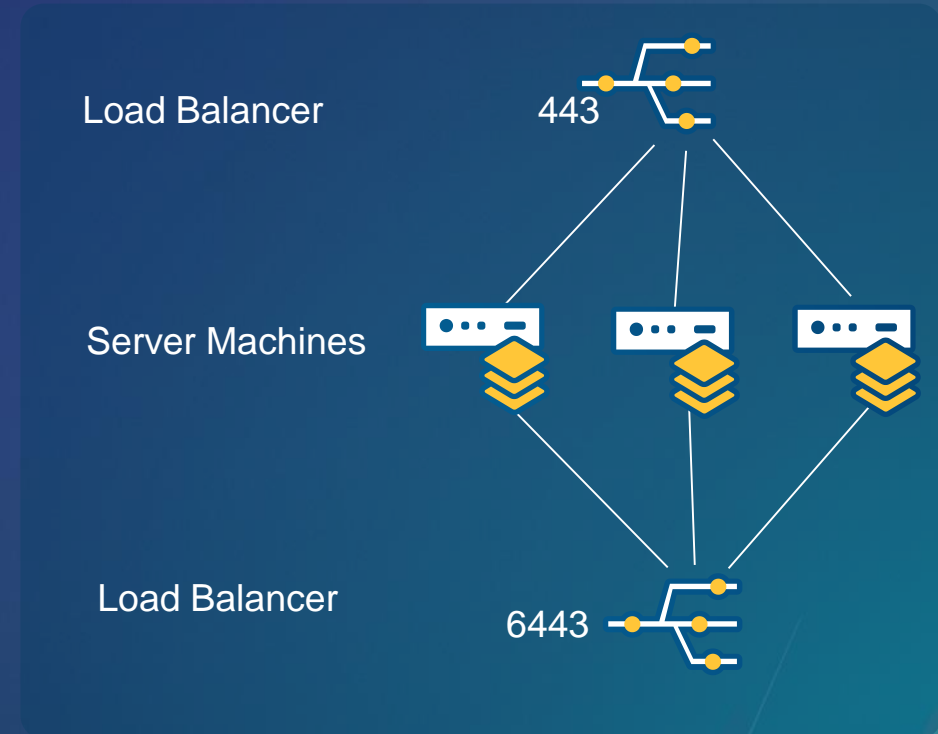
Portal Machines

Load Balancer

7443

privatePortalurl:7443

Services URL:443



Load Balancer

443

Server Machines

Load Balancer

6443

Administrative URL:6443

ArcGIS Server : Key Considerations for HA

- **Highly Available shared config-store and server directories**
- **Health Check provided for ArcGIS Server**
- **Highly Available URLs when communicating with Portal**
 - Portal URL
 - Private Portal URL
 - Services URL
 - Server Administrative URL
- **Install and Upgrade on all machines**

ArcGIS Enterprise



Portal

GIS Services

Hosted Feature and Tile Data



Portal for ArcGIS



ArcGIS Server



ArcGIS Data Store

Spatiotemporal Big Data Store

Title: Data Store Management Best Practices

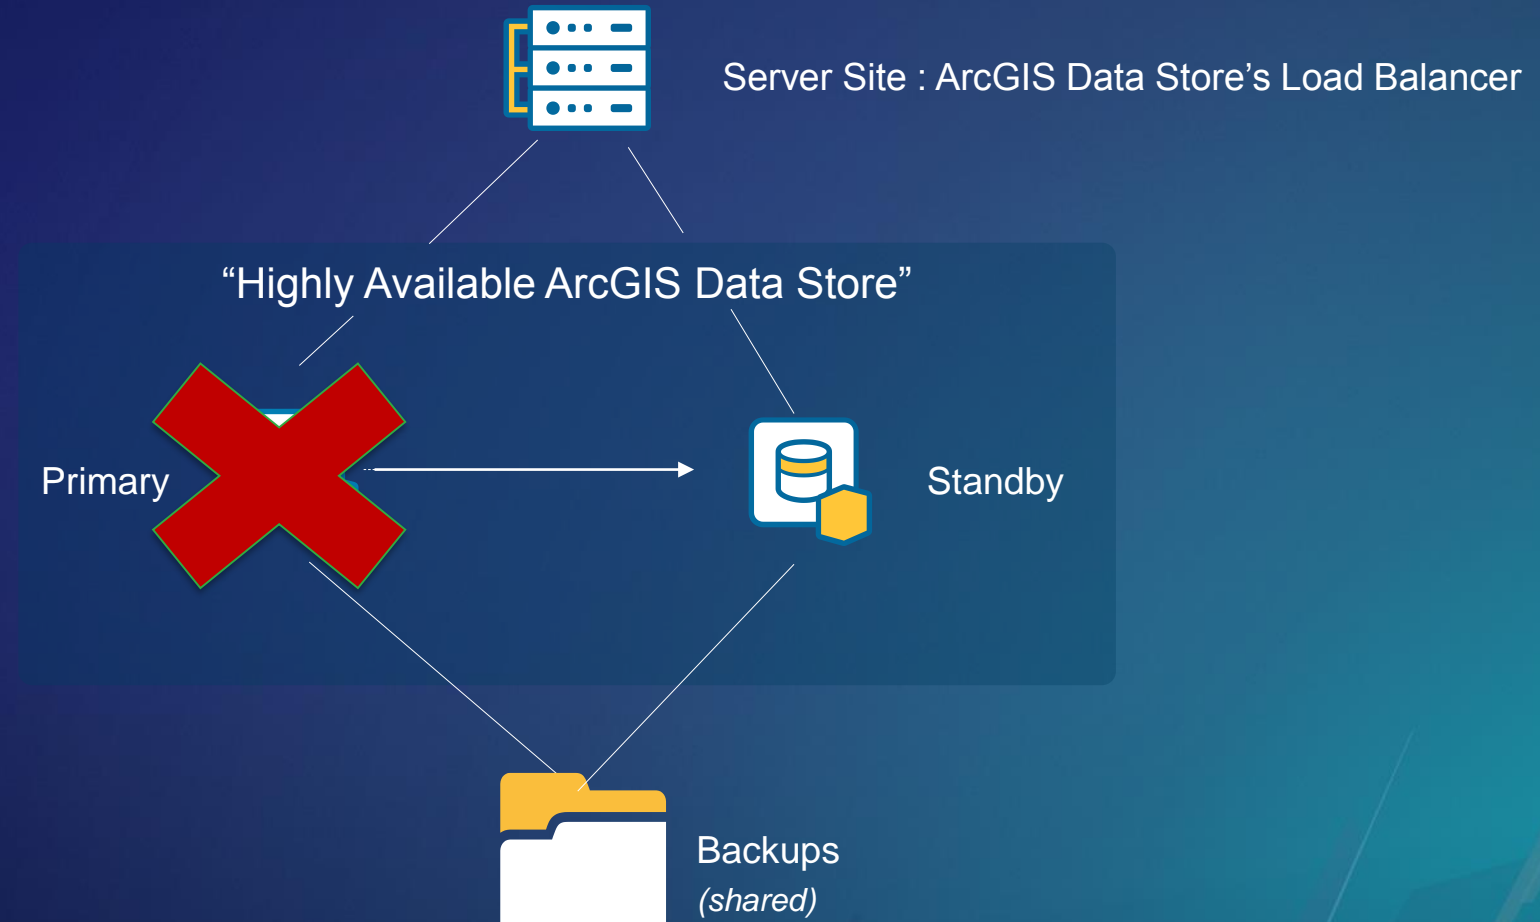
Date: 07/12/2017

Time: 1pm - 2pm

Location: SDCC

Room: SDCC - Room 10

ArcGIS Data Store: High Availability Architecture




ArcGIS Data Store: Failover Scenarios

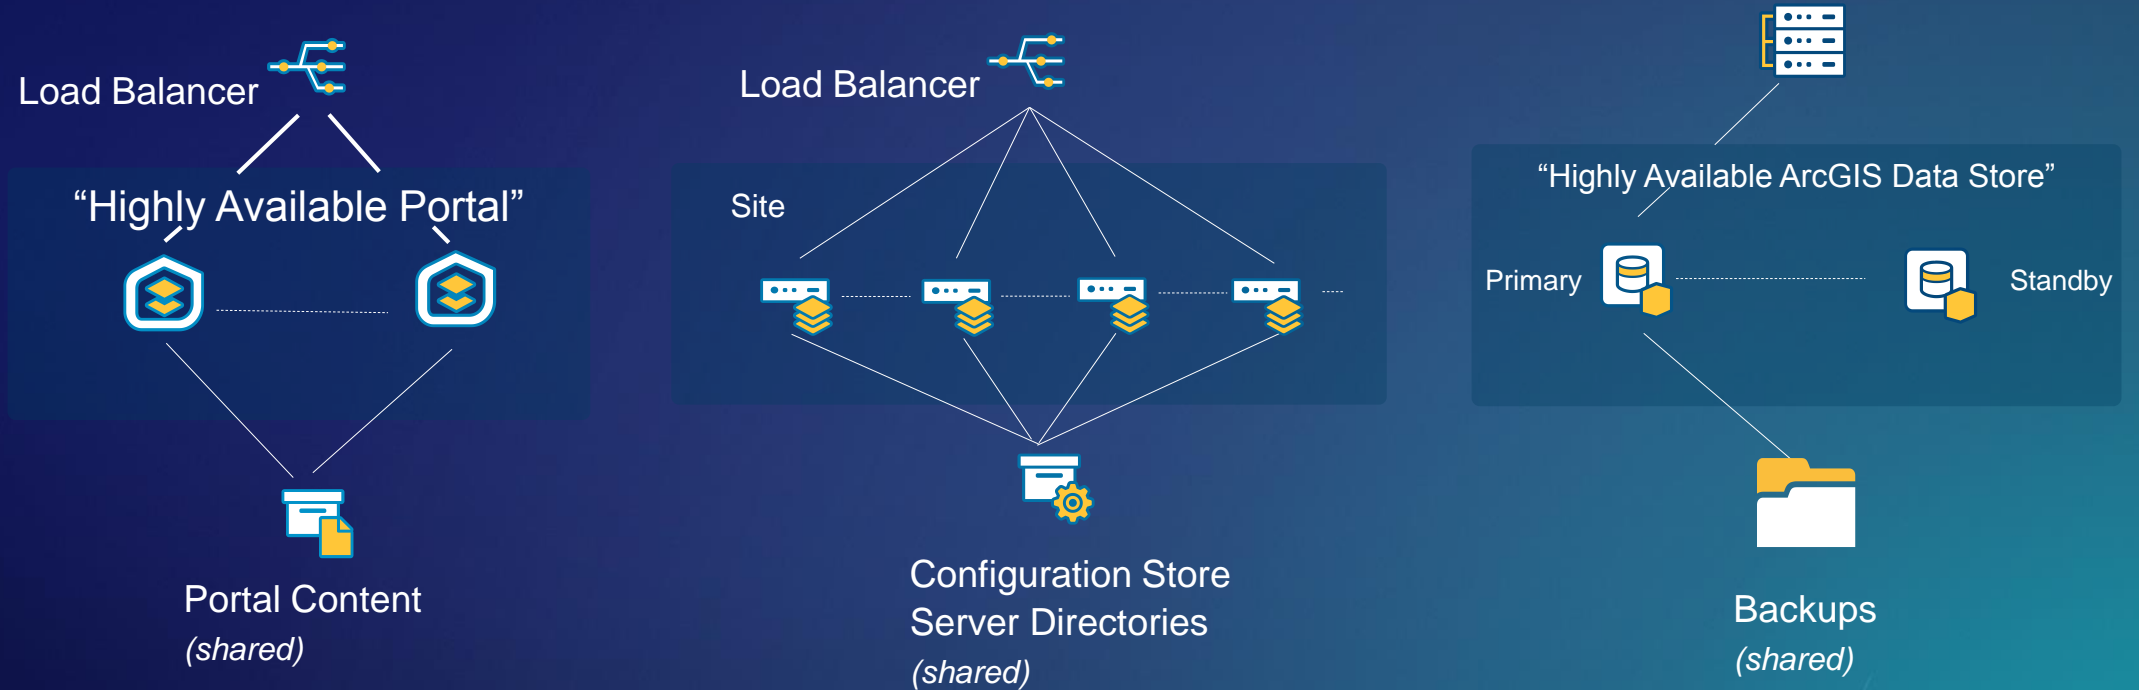
- **Primary ArcGIS Data Store stops working: Define Failure**
 - Computer crashes
 - Gets unplugged
 - Lose network connectivity
 - etc
- **Not “gracefully” shutdown**
 - Data Store service stops
- **<http://server.arcgis.com/en/documentation/> → Search “Fail over scenarios”**

Upgrade High Availability ArcGIS DataStore

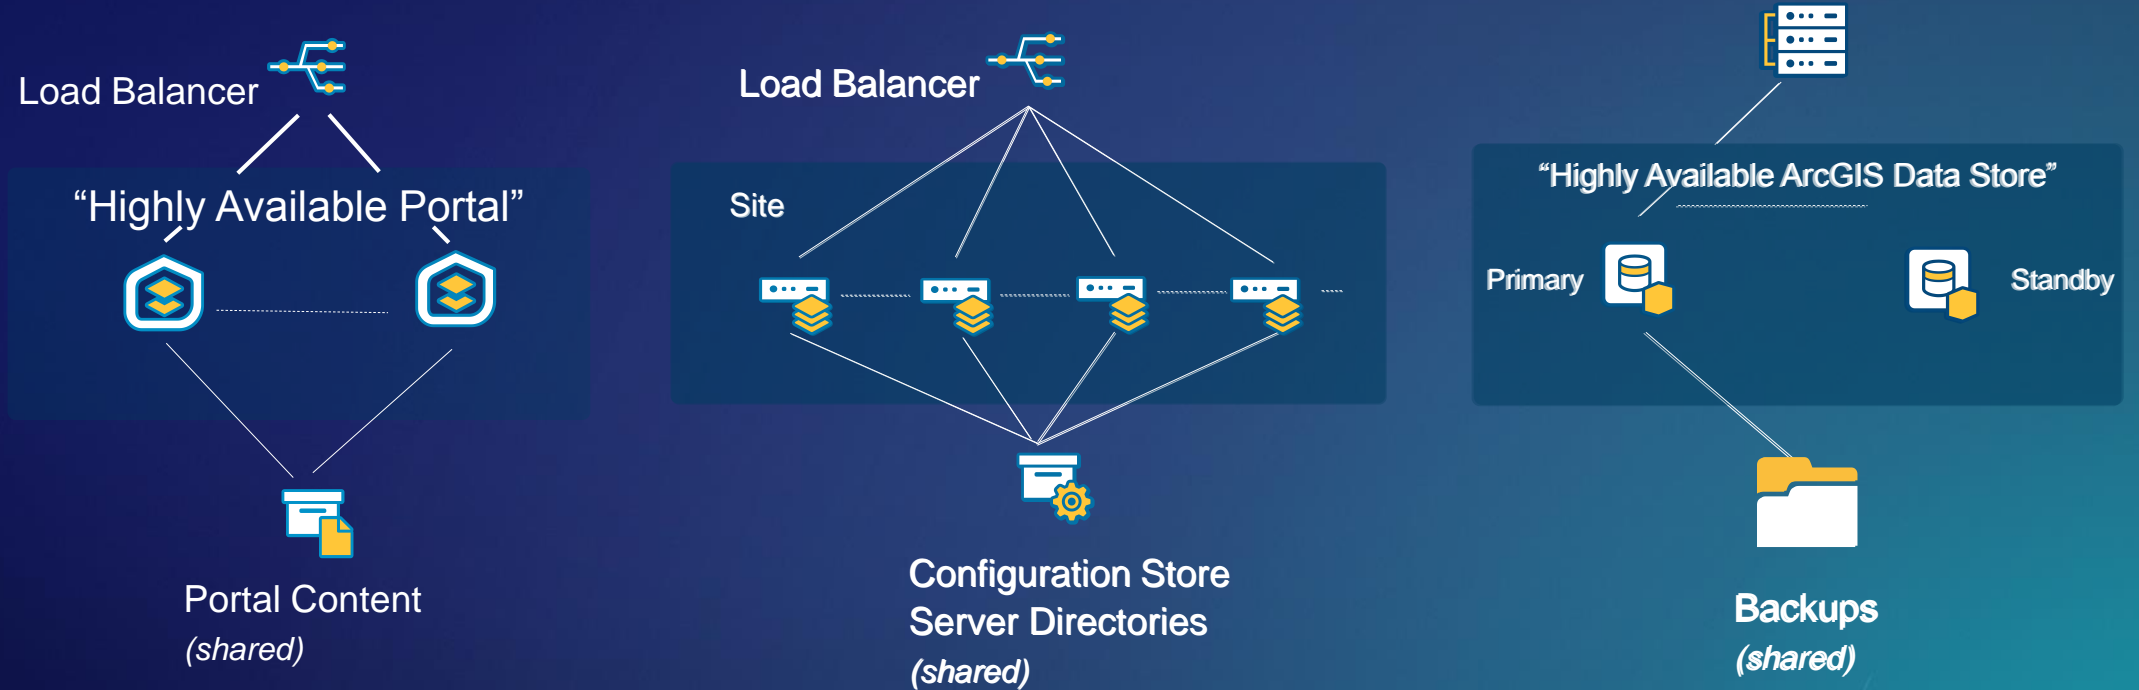
- Run setups on both machines
- On PRIMARY, run configure to upgrade
 - Relational
 - TileCache



ArcGIS Enterprise High Availability Deployment



Upgrade ArcGIS Enterprise High Availability Deployment




• Upgrade Order



Native Cloud Implementations

- Cloud Store
- Caching Directory
- Data Input Directory
- Backup/Restore



Register a cloud store on your ArcGIS Server

Type:	Amazon S3
Cloud Store Name:	Amazon S3
Credential Type:	Microsoft Azure Storage
	Microsoft Azure Data Lake Store
	Alibaba Cloud Storage
Access Key Id:	Huawei Cloud Storage

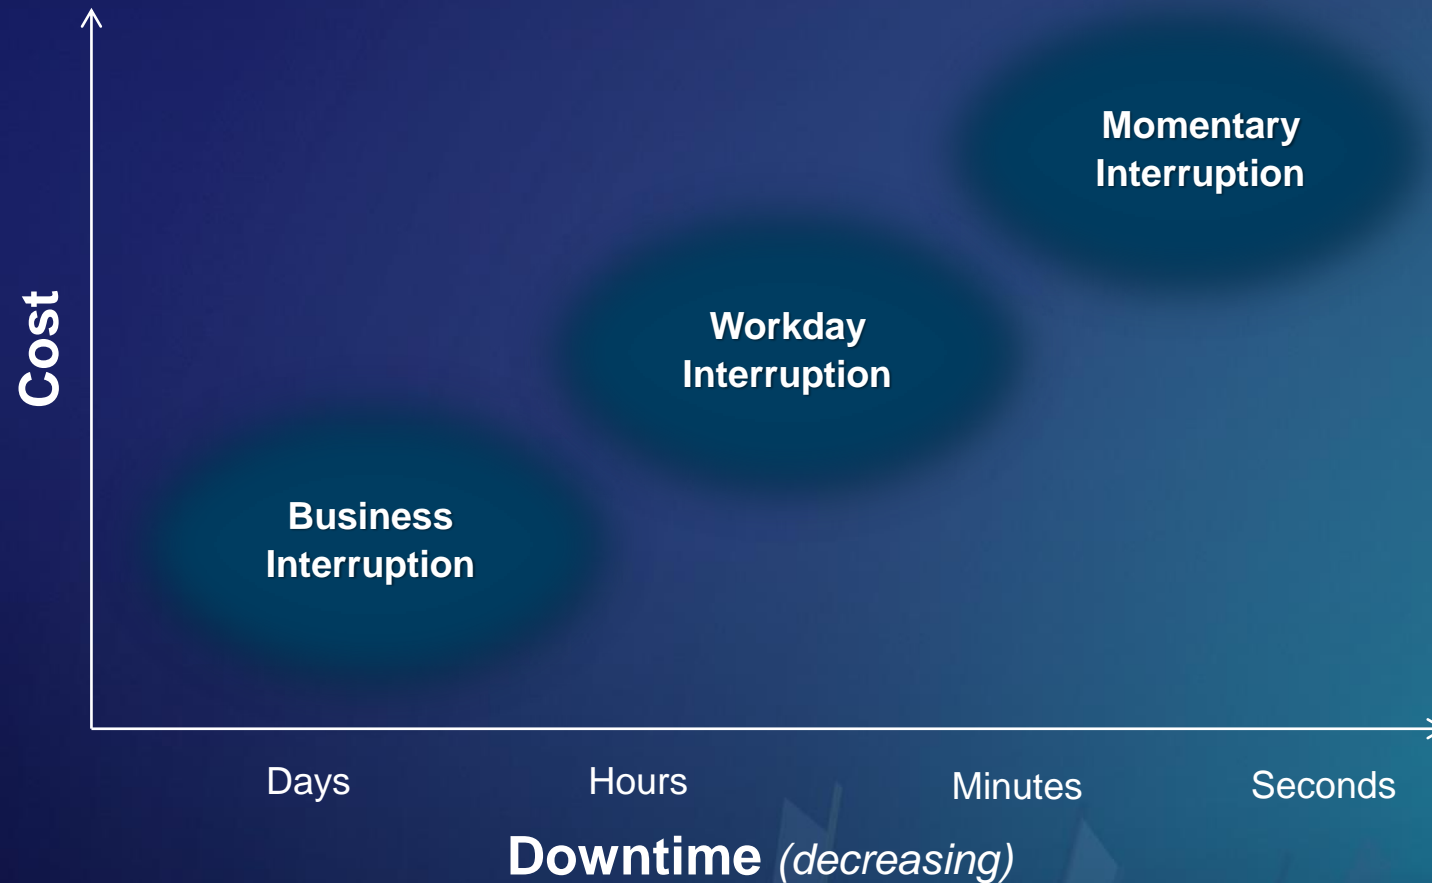
ArcGIS Enterprise HA: Part of Your HA Architecture

- **Your Data**
 - Enterprise GeoDatabase
 - File based Data
- **Software**
 - Web Server
 - Software Load Balancer
- **Hardware**
 - File Server
 - Network
- **People**
 - HA?
 - IT strong?

ArcGIS Enterprise HA: IT Governance

- **Ensure the effective and efficient use of IT**
- **Policies and procedures highly disciplined**
 - **Planned and updated in a timely manner**
 - **Documented clearly**
 - **Tested properly**
 - **Exercised with staff**

ArcGIS Enterprise HA: Spectrum, Not a Switch



Disaster Recovery

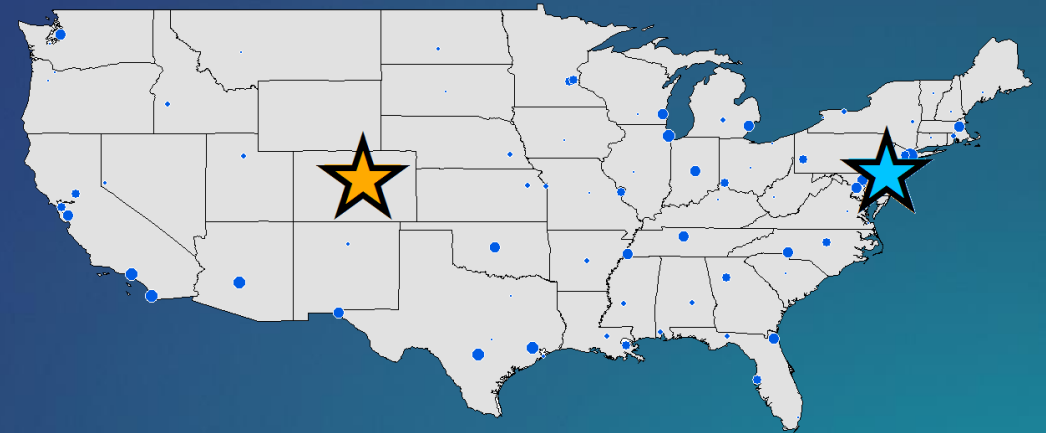
Geographic Redundancy

Agenda

- **What is geographic redundancy**
- **Using the Web GIS DR tool**
- **Roadmap to being geographically redundant**

Overview

- Geographically separate data centers
- Duplicated configurations and data between the two data centers
- Components within data centers are typically highly available
- WebGIS DR Tool is used to move snapshots of data from primary to standby
- Complex disaster recovery option

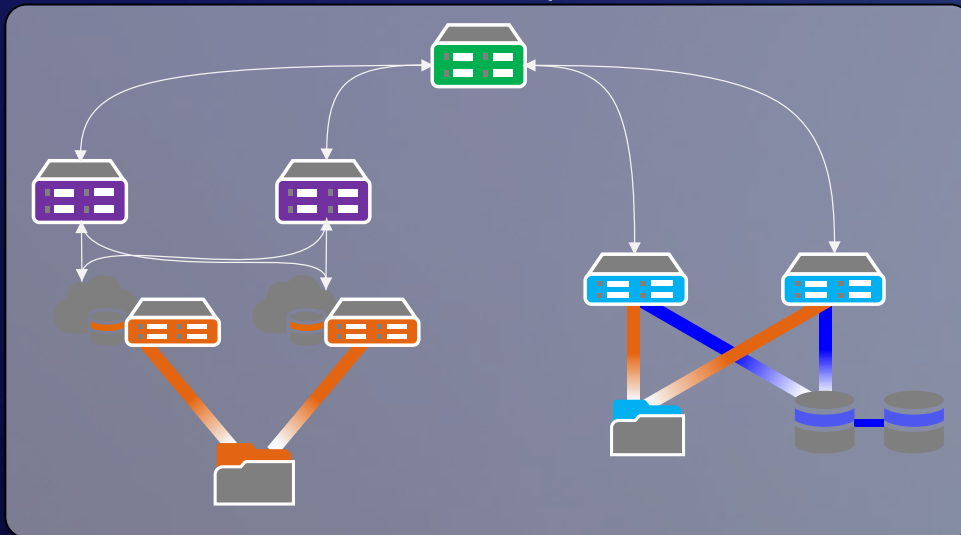


Geographic Redundancy

Public Portal URL - <https://mysite.esri.com/portal>
Services URL – <https://mysite.esri.com/server>

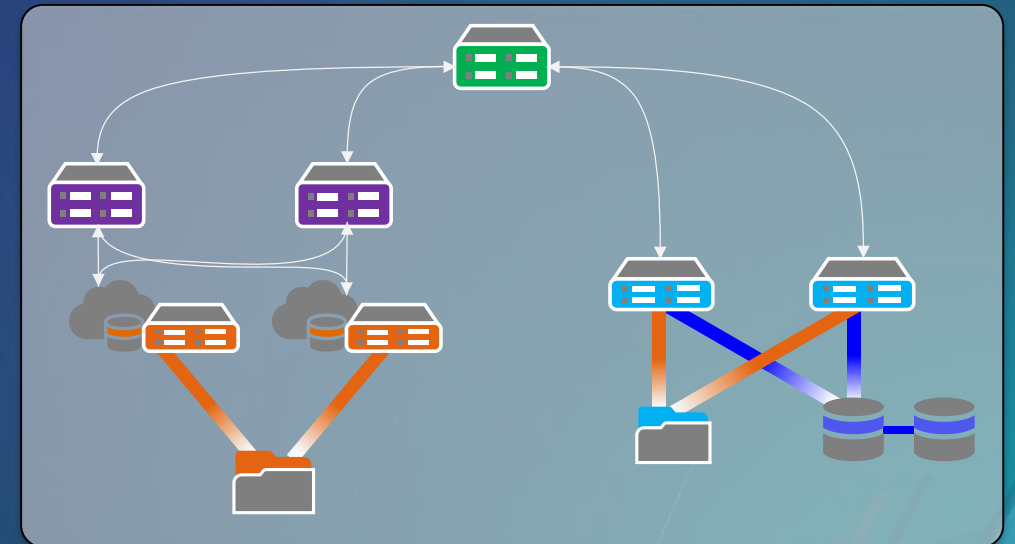


198.0.0.1



East coast data center (primary)

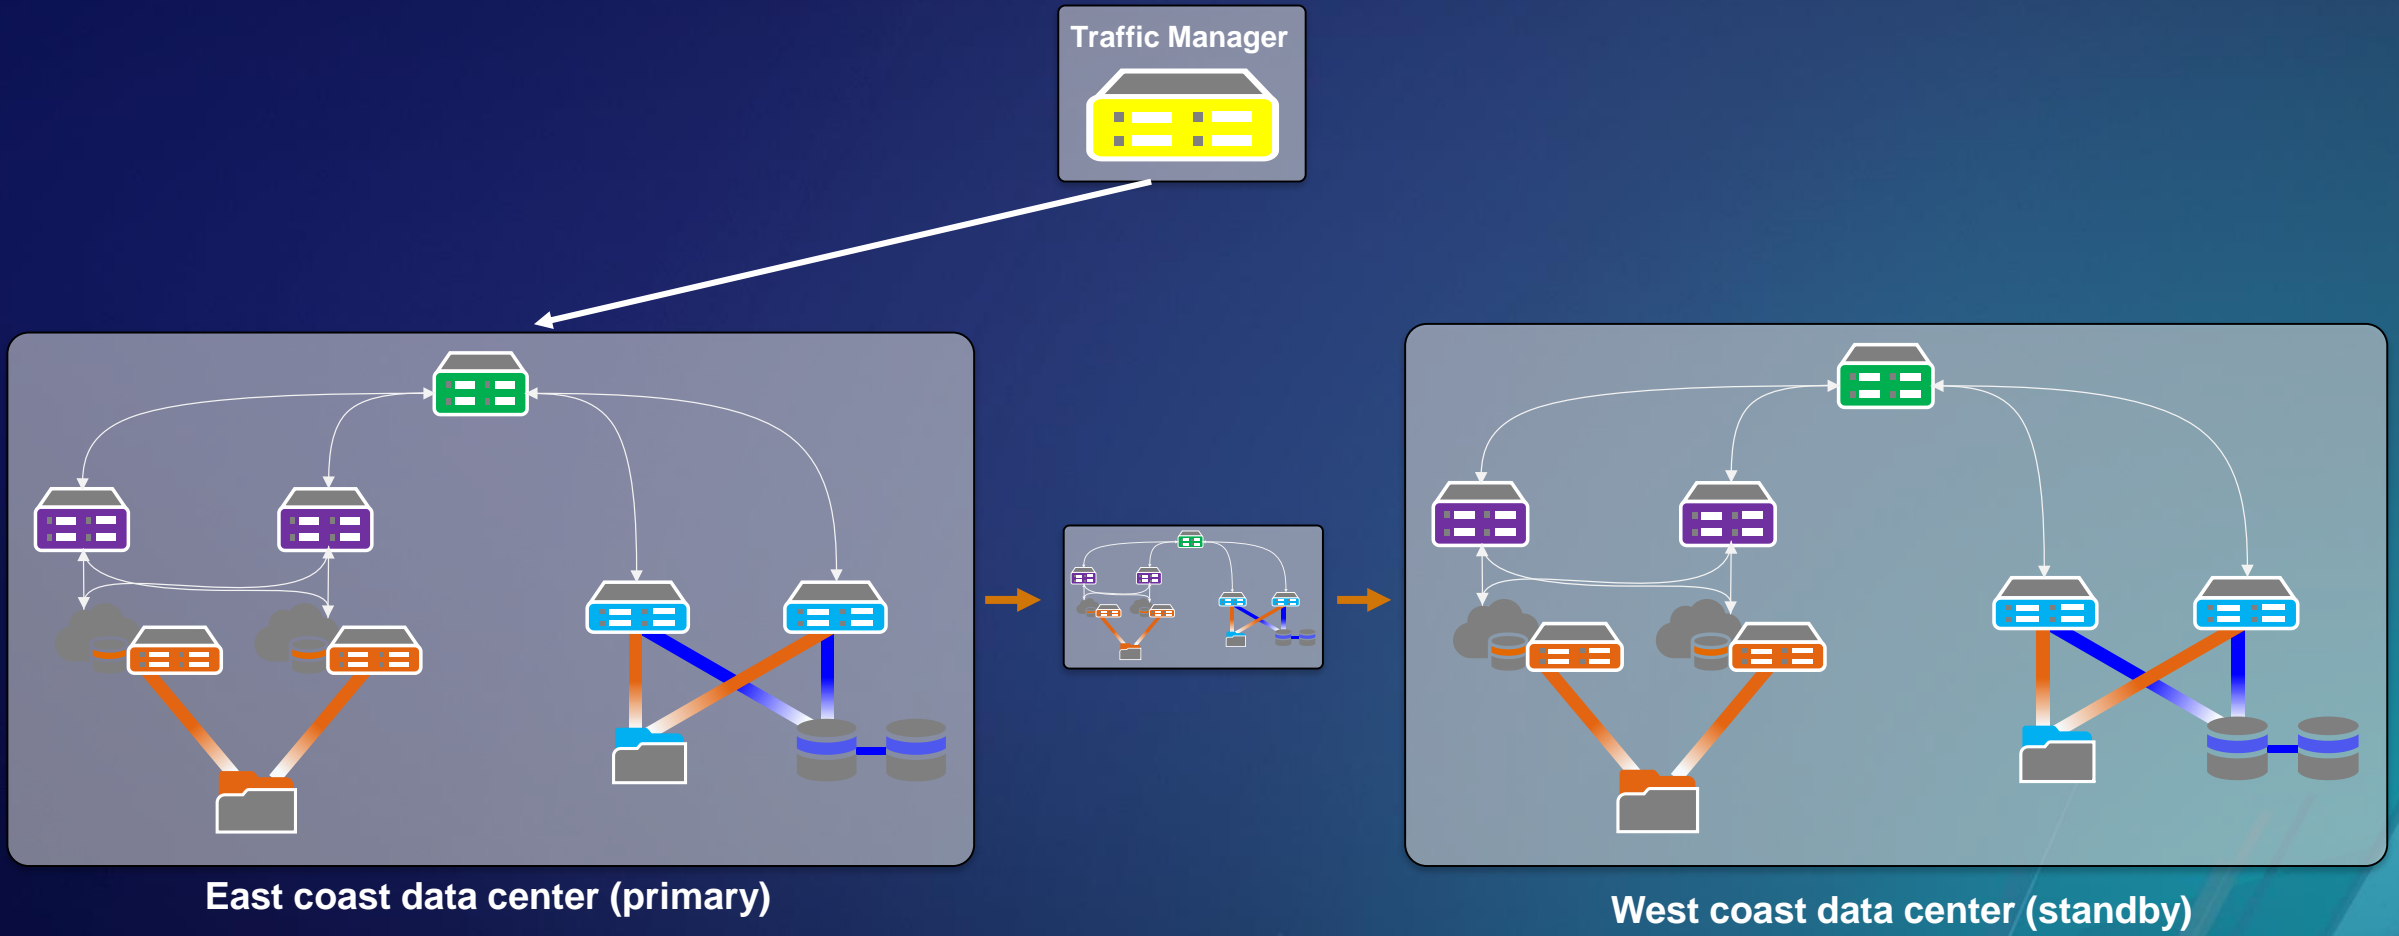
198.0.0.2



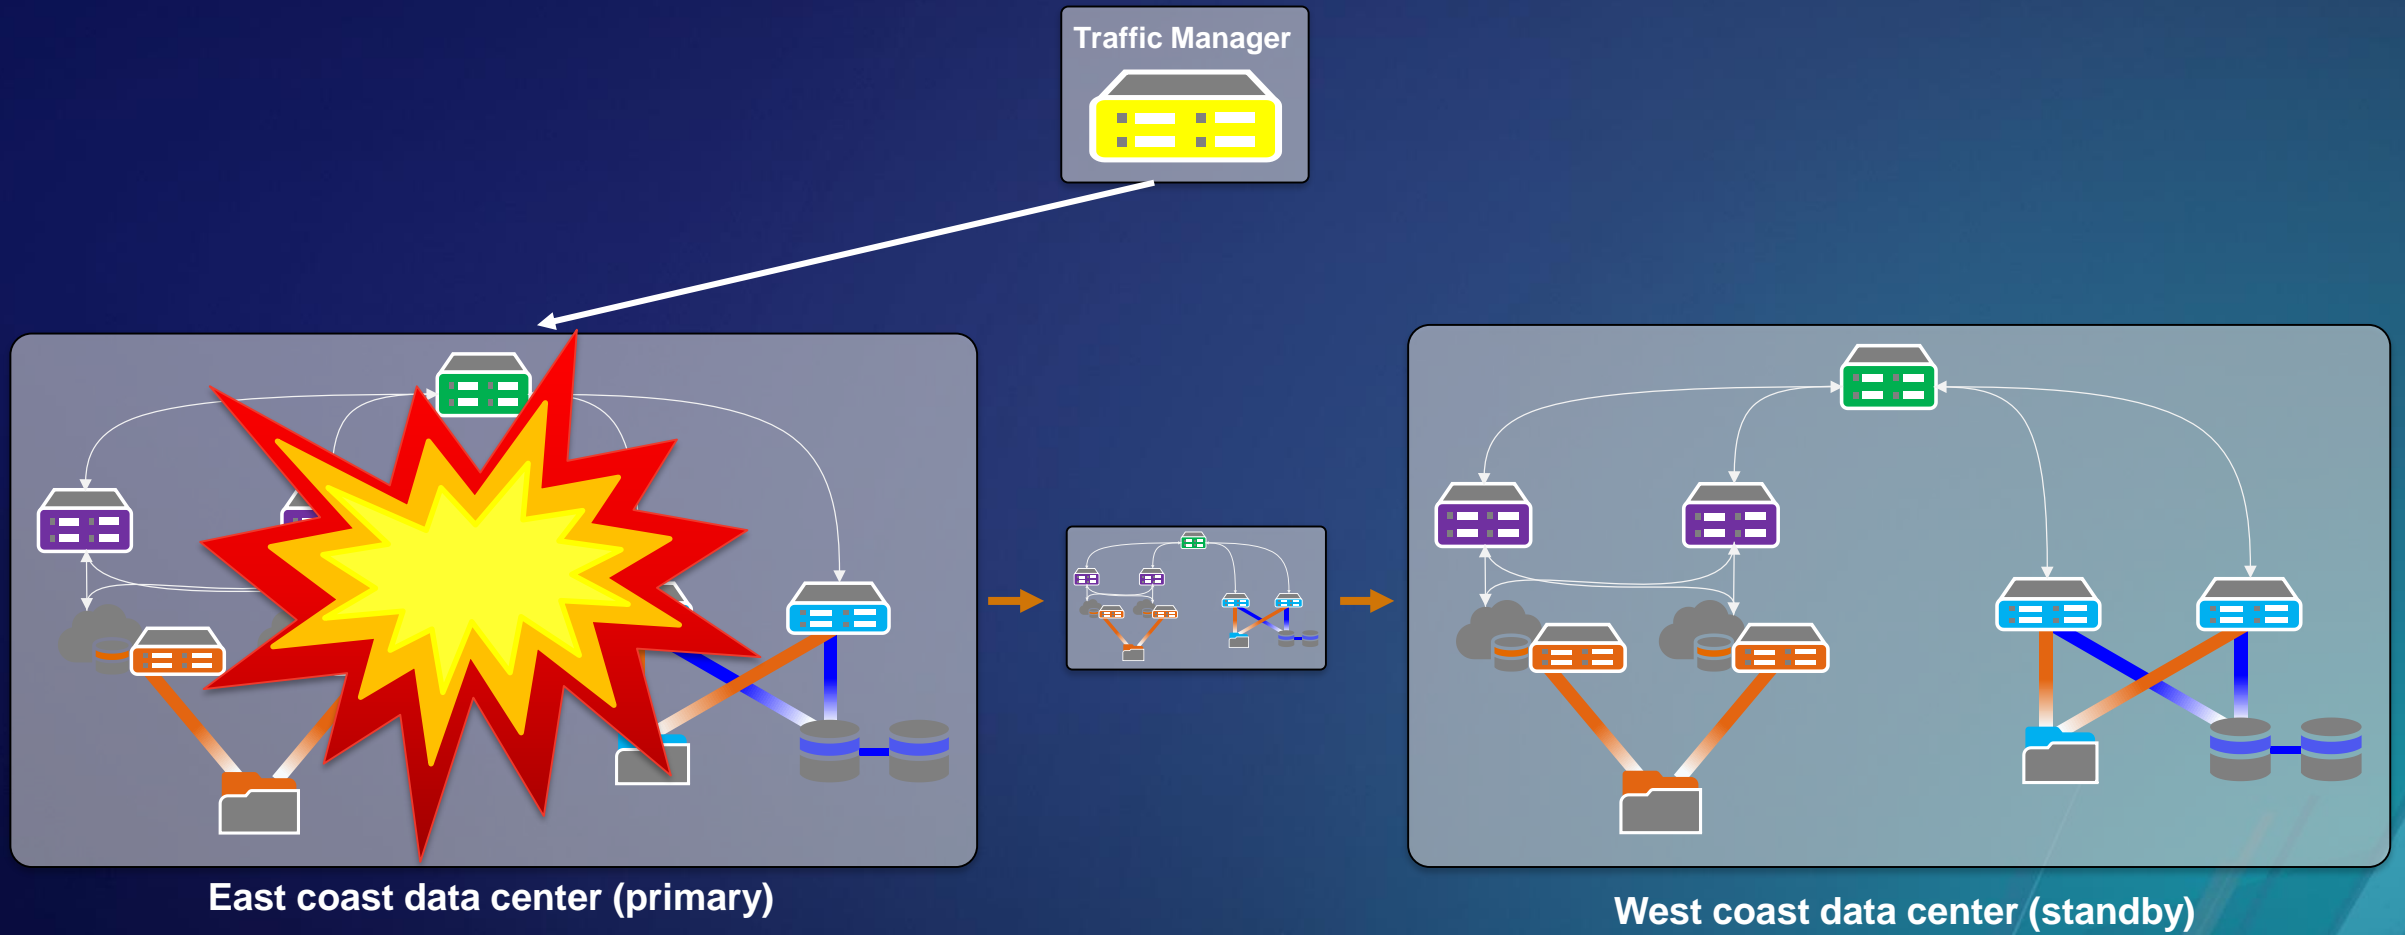
West coast data center (standby)

Public portal URL and services URL need to be the same
Referenced data paths need to be the same


Geographic Redundancy




Geographic Redundancy



Geographic Redundancy

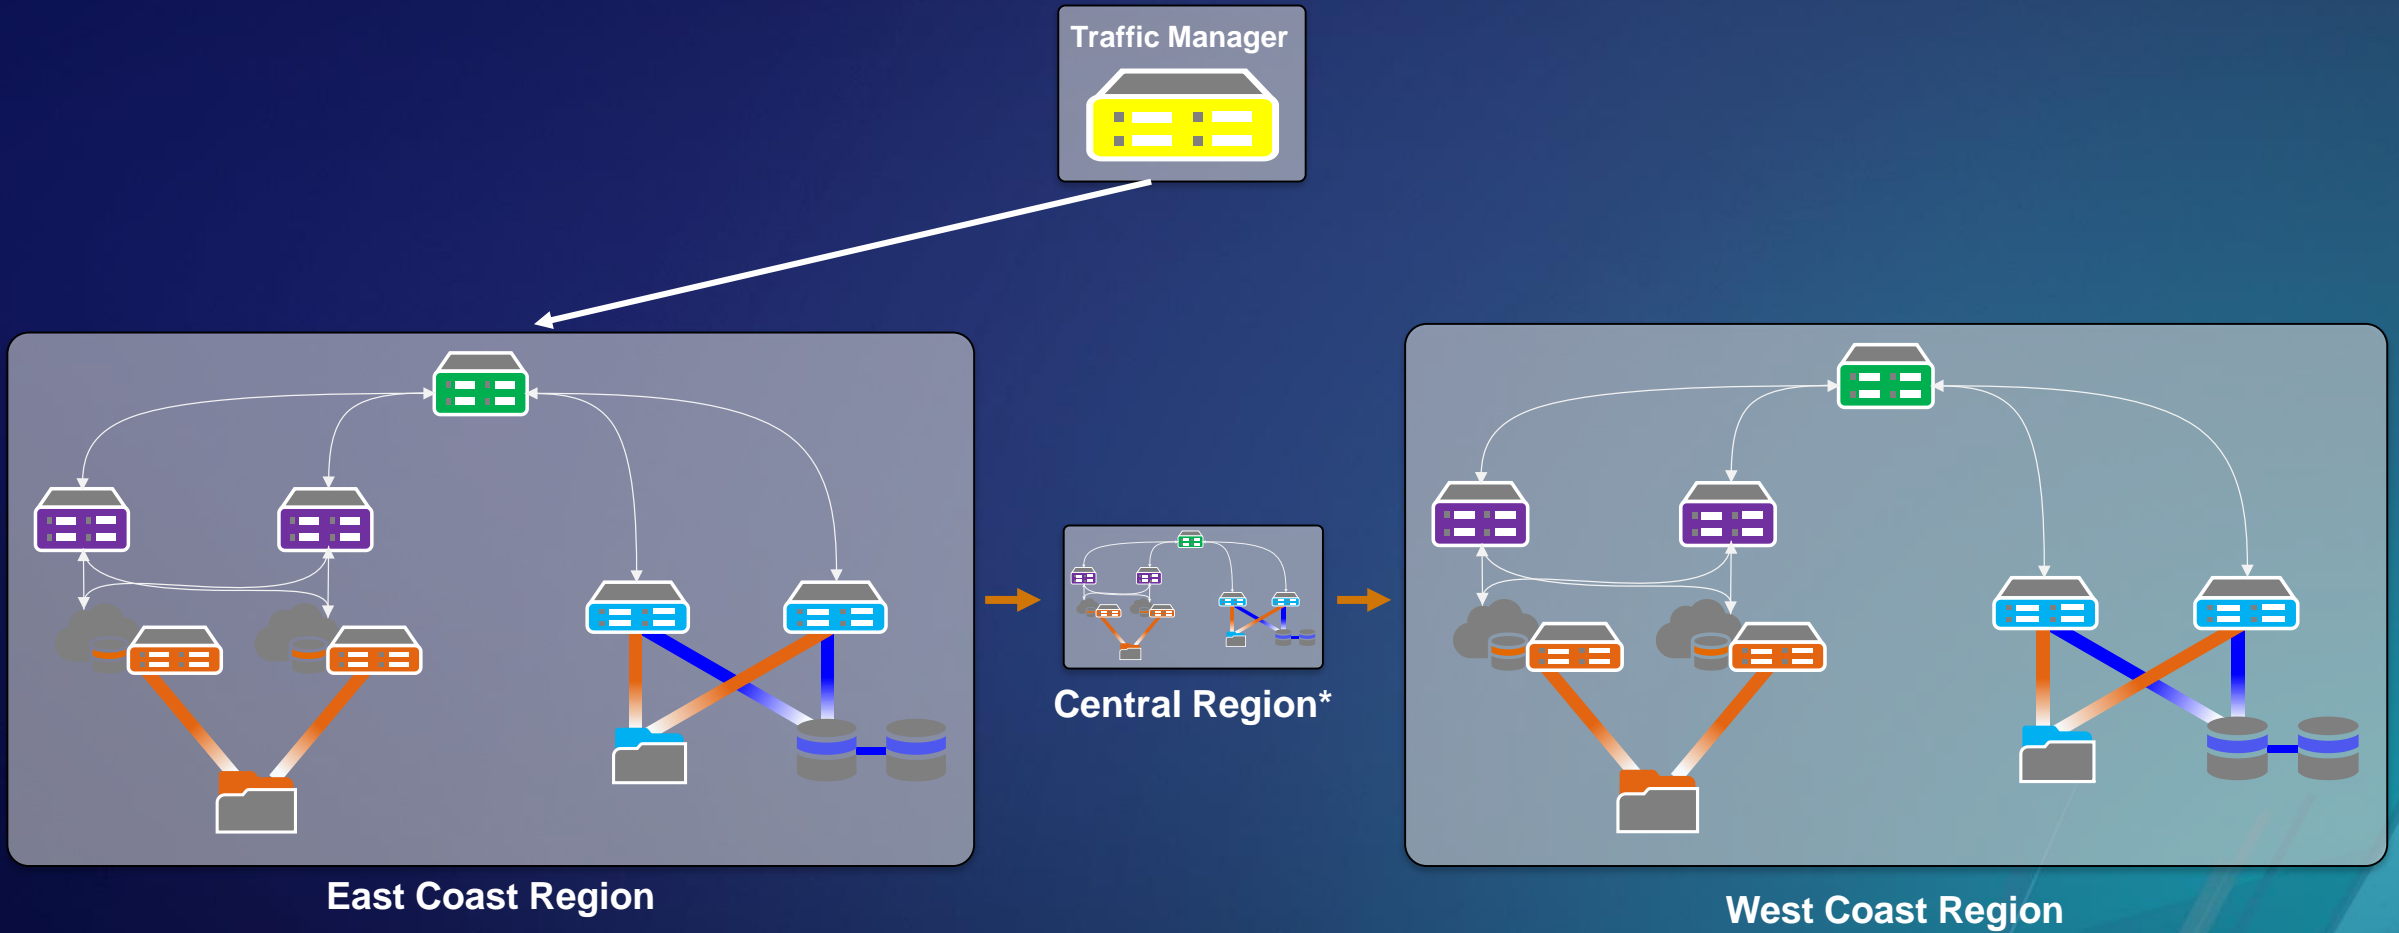


East coast data center (primary)



West coast data center (standby)

Geographic Redundancy – Cloud deployments

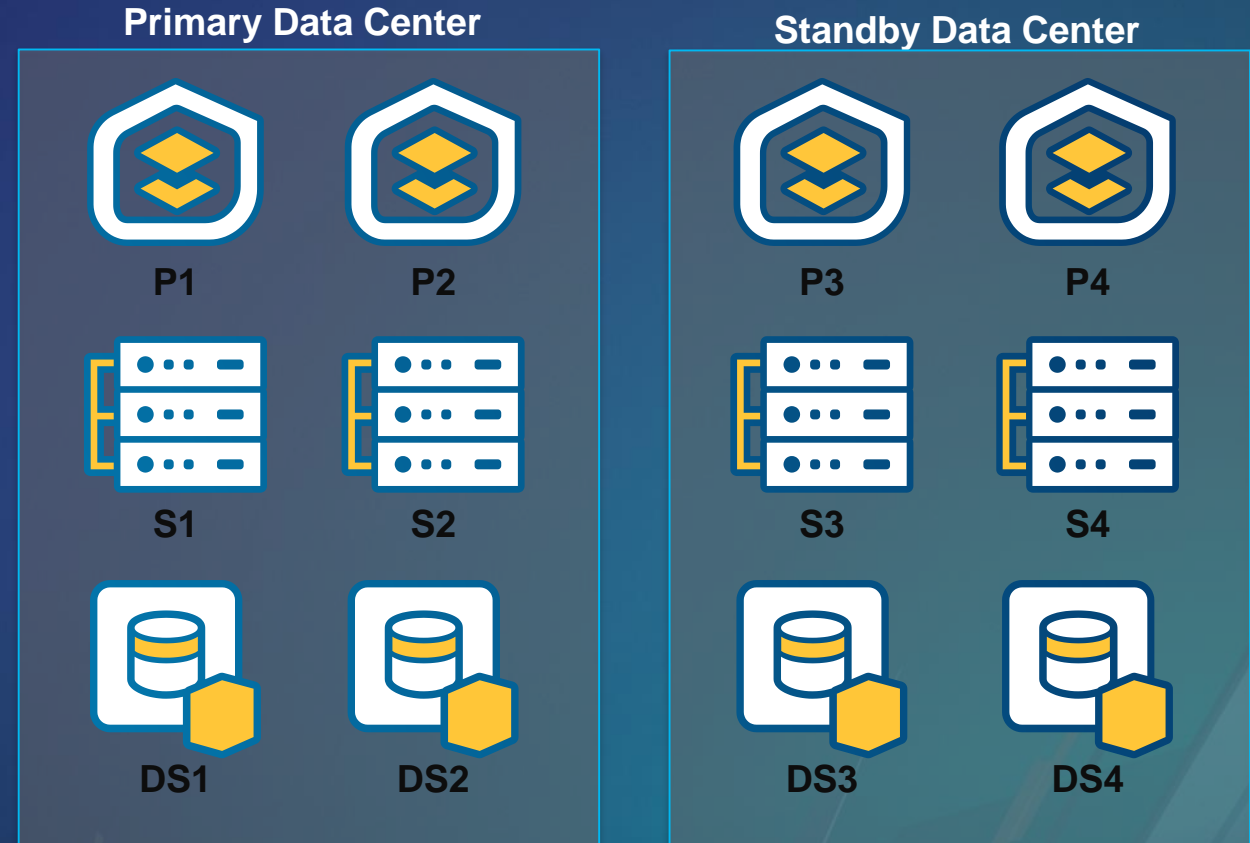


Roadmap for geographic redundancy

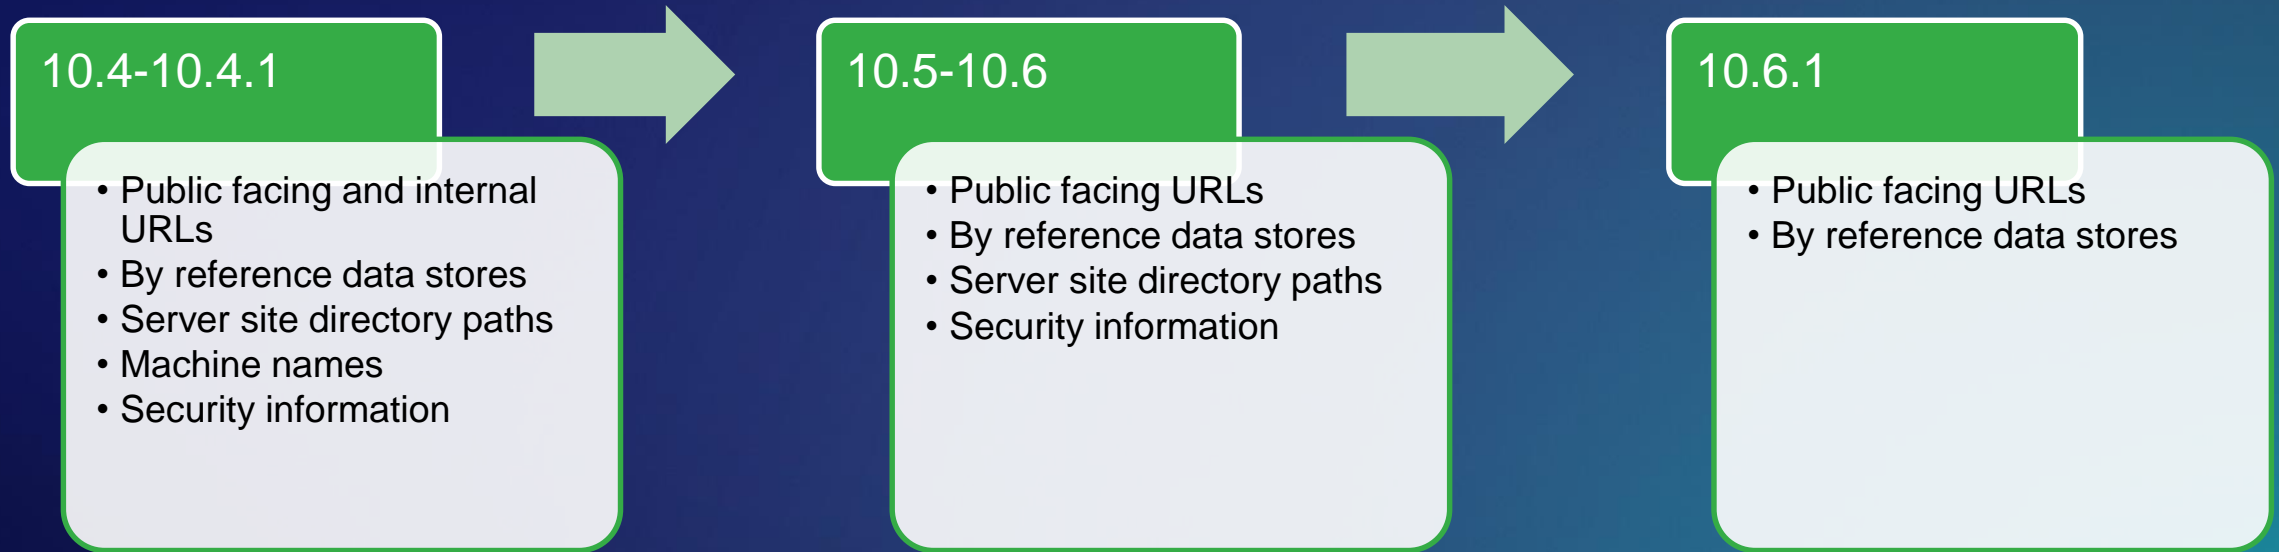
1. **Duplicate the deployment between primary and standby data centers**
2. **Create snapshots of the primary data center**
3. **Apply snapshots to the standby data center**
4. **Monitor your standby data center**

Duplication

- Number of machines should be the same
- Identical URLs between data centers
 - Public Portal URL
 - Services URL
- Identical paths to data and connections to databases or enterprise geodatabases



Duplication – What needs to be the same?

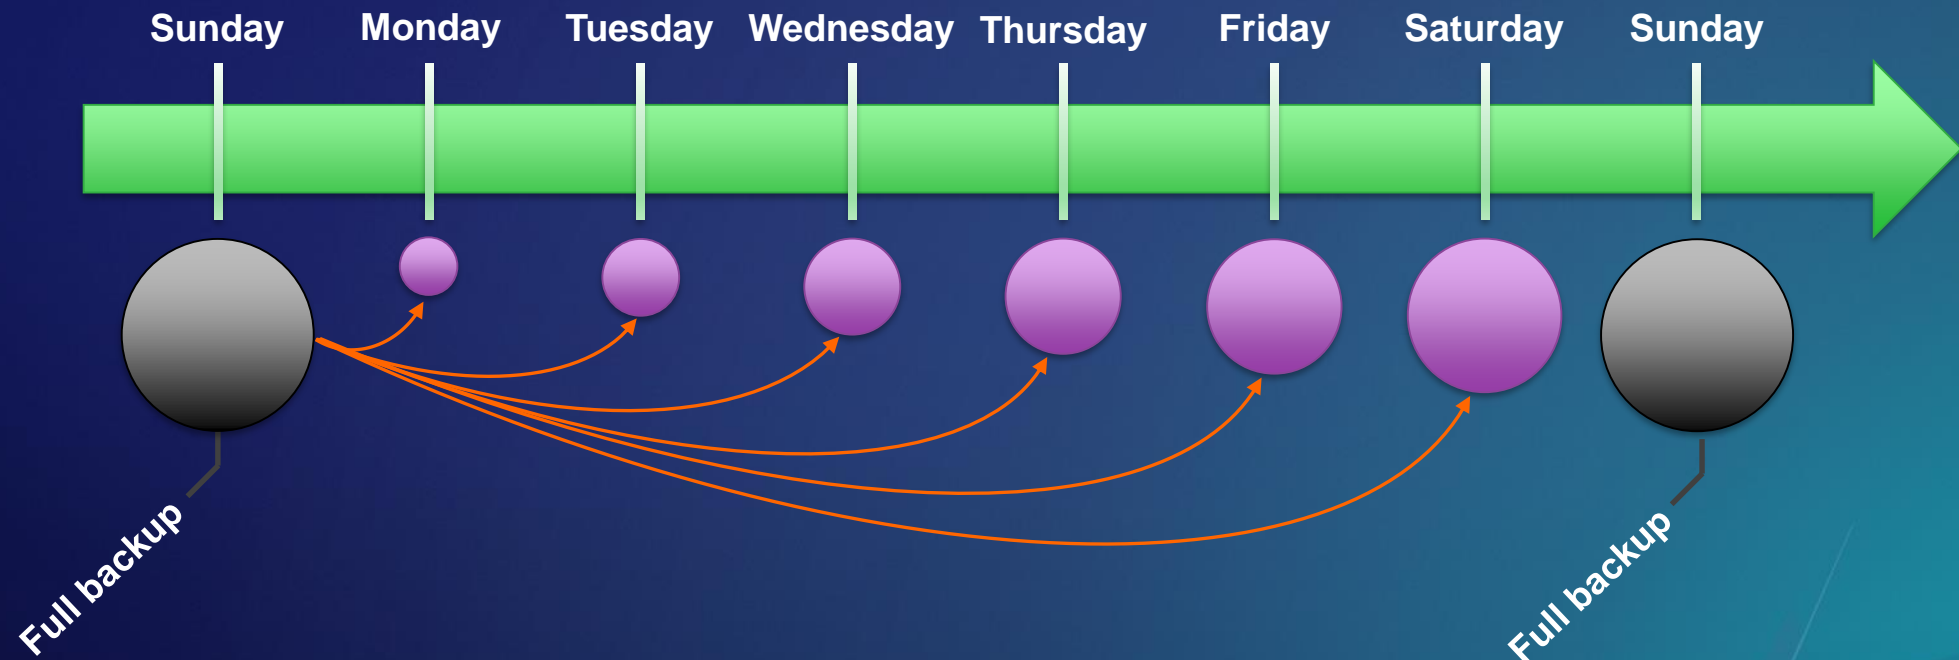


Creating snapshots

- **Full snapshot**
 - Create an initial snapshot of all of the data within the ArcGIS Enterprise
 - Internally defines a base time that will be used for an incremental snapshot
- **Incremental snapshot**
 - Creates a snapshot of all of the data that has been created or modified since the last full backup
 - Decreases the time it takes to synchronize content, services, and data between primary and standby

Creating incremental snapshots

- Creates a snapshot of all data added or modified since the last full snapshot



Monitoring and Failover

- **QC process on standby ArcGIS Enterprise**
 - Checking the index within Portal
 - Validating federated Servers
 - Validating data stores using Server Admin
 - Checking important services or applications
- **Detecting when components fail within a data center**
 - Monitoring the healthCheck URLs of Portal and Server
- **Failing over data centers should be a manual, deliberate decision**

ArcGIS REST Services Directory

[Home](#) > [healthCheck](#)

[JSON](#)

Server Health Check

Health Check successful, the site is ready

Takeaway points

- **Important to understand the requirements of geographic redundancy as a disaster recovery option**
- **Take advantage of the Web GIS DR tool to move snapshots of the deployment from primary to standby**
- **Geographic redundancy is a complex disaster recovery option**

Success Stories with HA or DR




- Let us know if you have a success story to share

Please Take Our Survey on the App

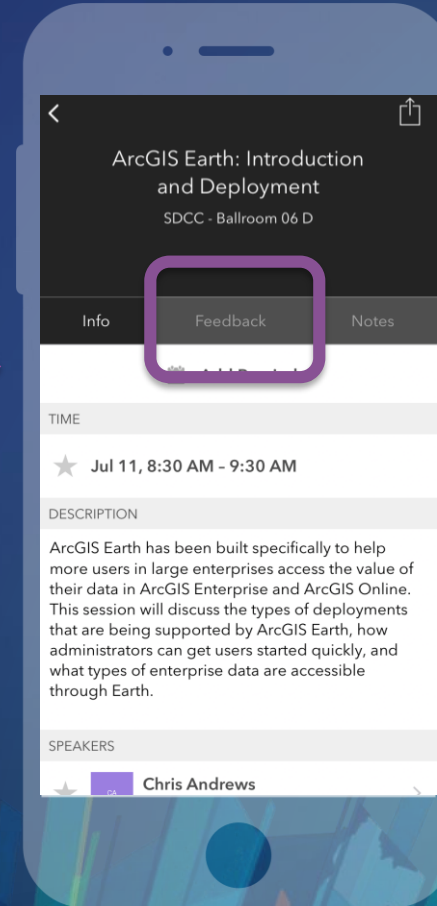
Download the Esri Events app and find your event




Select the session you attended



Scroll down to find the feedback section



Complete answers and select "Submit"





esri

**THE
SCIENCE
OF
WHERE**