ArcGIS Enterprise on Kubernetes: Publishing and Tuning

Gayathri Vijay
Garima Tiwari
Moginraj Mohandas
Shreyas Shinde
Jon Quinn
Laurence Clinton

2021 ESRI USER CONFERENCE

ArcGIS Enterprise on Kubernetes Publishing Services

Gayathri Vijay

Publishing Services in ArcGIS Enterprise on Kubernetes

- Map Services
- Hosted Feature Services
- Bulk Publishing

Map Services in ArcGIS Enterprise on Kubernetes

- Map services can be published by copying all data
- Map services can be published by referencing data
  - Register data store beforehand
- Data store types
  - Folder
  - Database
- Supported database data stores
  - SQLServer
  - PostgreSQL
- Map services can be published using either shared or dedicated instance

Map service referencing a folder data store

- If you plan to use file-based data for your services, you need to look into a pre-requisite workflow that is highly recommended during site creation.
- This workflow will be described during the service management and tuning section later on.

Map Services in ArcGIS Enterprise on Kubernetes

- Overwriting of map services from ArcGIS Pro
- Map service extensions
  - Feature Access
  - WMS
  - WFS
  - Network Analysis
  - KML

Demonstration:

Publish Map Service & Feature Access

Gayathri Vijay

Hosted Feature Services in ArcGIS Enterprise on Kubernetes

- Publish hosted feature services from ArcGIS Pro and from Portal
- Supports overwriting of hosted feature services from Portal and ArcGIS Pro
- Export to different formats (fgdb, shp, csv, excel, etc.) from hosted feature service

Demonstration:

Publish Hosted Feature Service

Gayathri Vijay

Bulk Publishing in ArcGIS Enterprise on Kubernetes

- Register database data store to Portal
- Create feature service in bulk from datasets in a database data store

Demonstration:

Create Feature Services in Bulk from Database Data Stores

Gayathri Vijay

Summary

- Publishing of services in ArcGIS Enterprise on Kubernetes is the same as publishing services in ArcGIS Enterprise on Windows/Linux
- Users can access and use these published services just like they would use a service in ArcGIS Enterprise on Windows/Linux
- In future releases more service types will be supported

ArcGIS Enterprise on Kubernetes Cache Generation & Scene, Tile Layers

Garima Tiwari

esri THE SCIENCE OF WHERE®

Agenda:

Generate / manage map cache in ArcGIS Enterprise on Kubernetes

1. publishing cached map image layers
2. options for cache store
3. migrating existing map caches
4. scaling considerations

Scene, Vector, Tile, Elevation layers

- using packages generated in ArcGIS Pro

Cached services provide fast visualization of your map or scene content.

Cache content comes in various types:

- Raster (png, jpeg)
- Vector (pbf)
- Elevation (LERC)
- Scene (i3s)

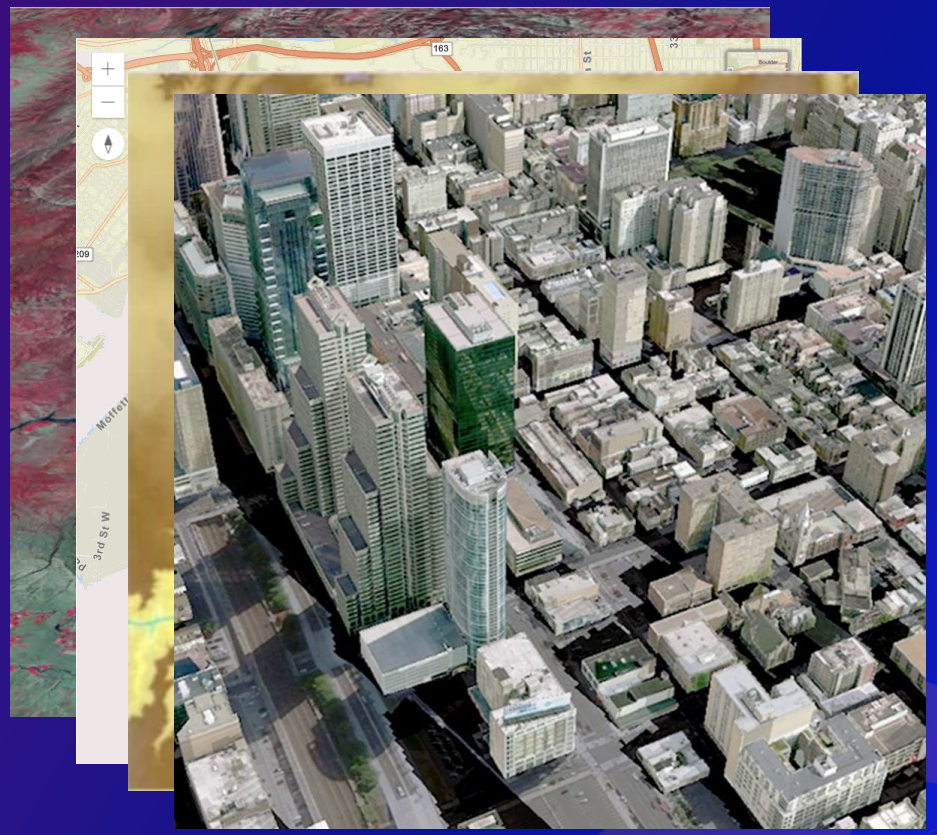

Cache content generation

ArcGIS Enterprise	• Map Image layers • Scene layers
ArcGIS Pro	• Vector tile layers (.vtpk) • Tile layers (.tpkx) • Scene layers (.slpk)
ArcGIS Server	• Map service • Image service

1. Publishing cached Map Services

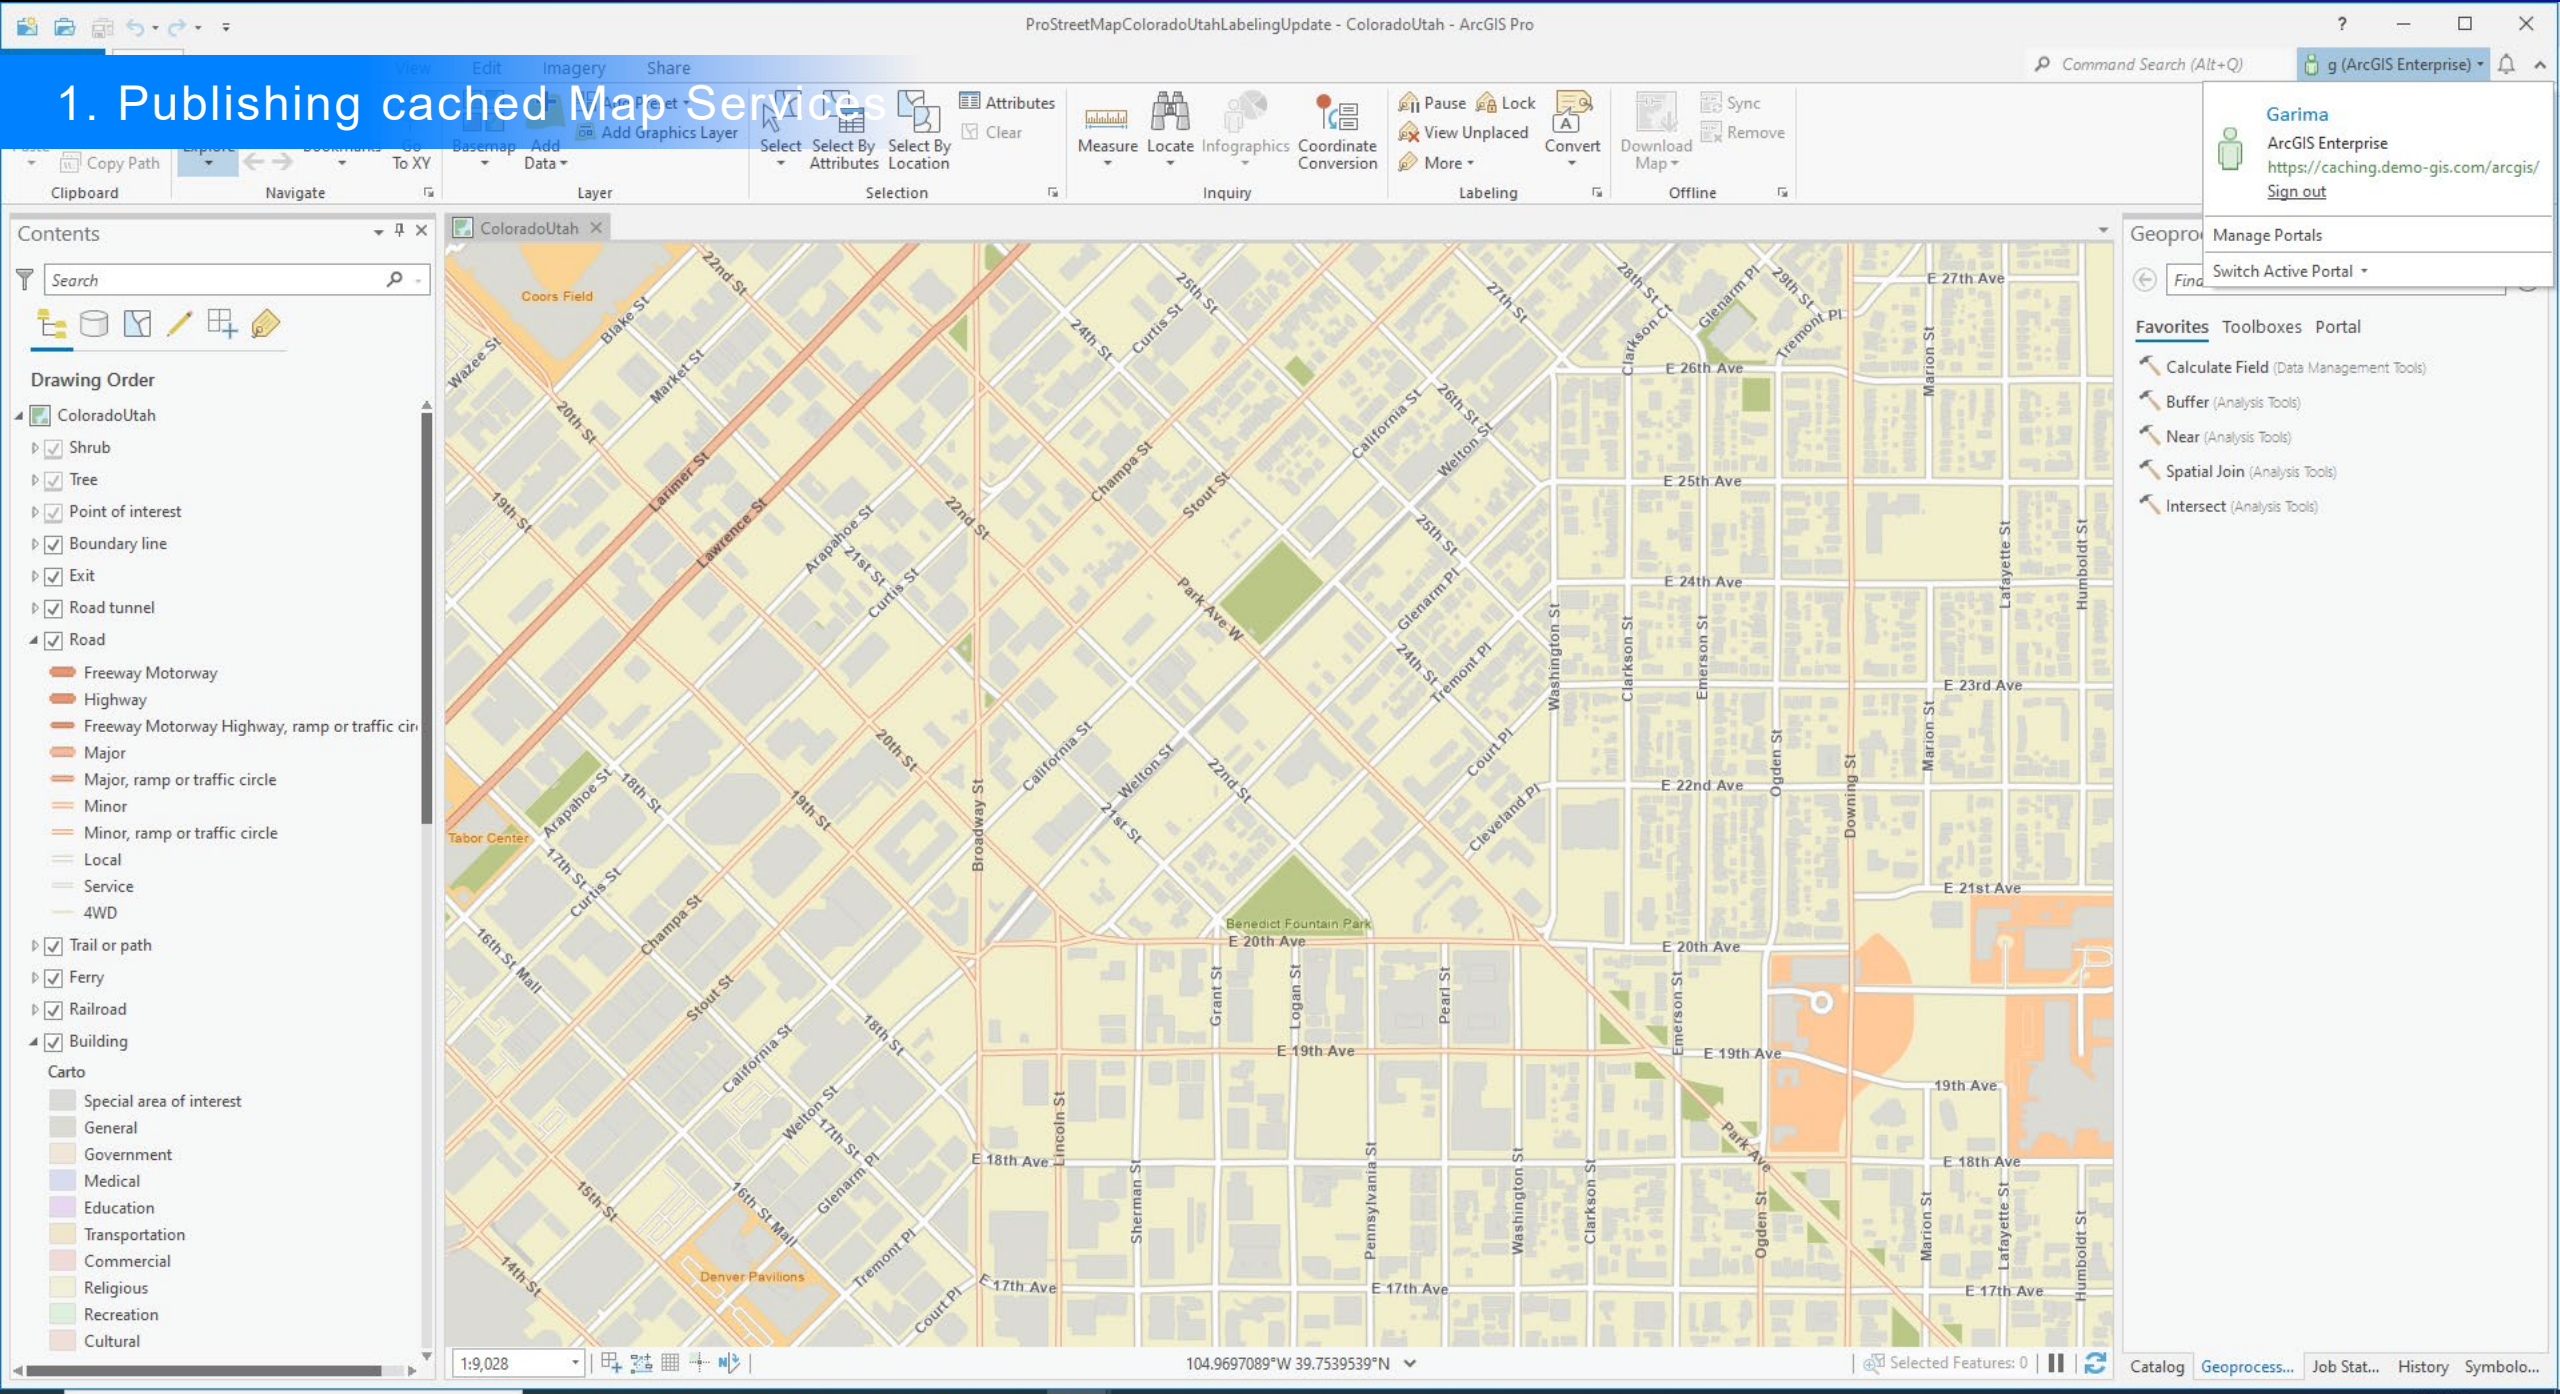

1. Publishing cached Map Services

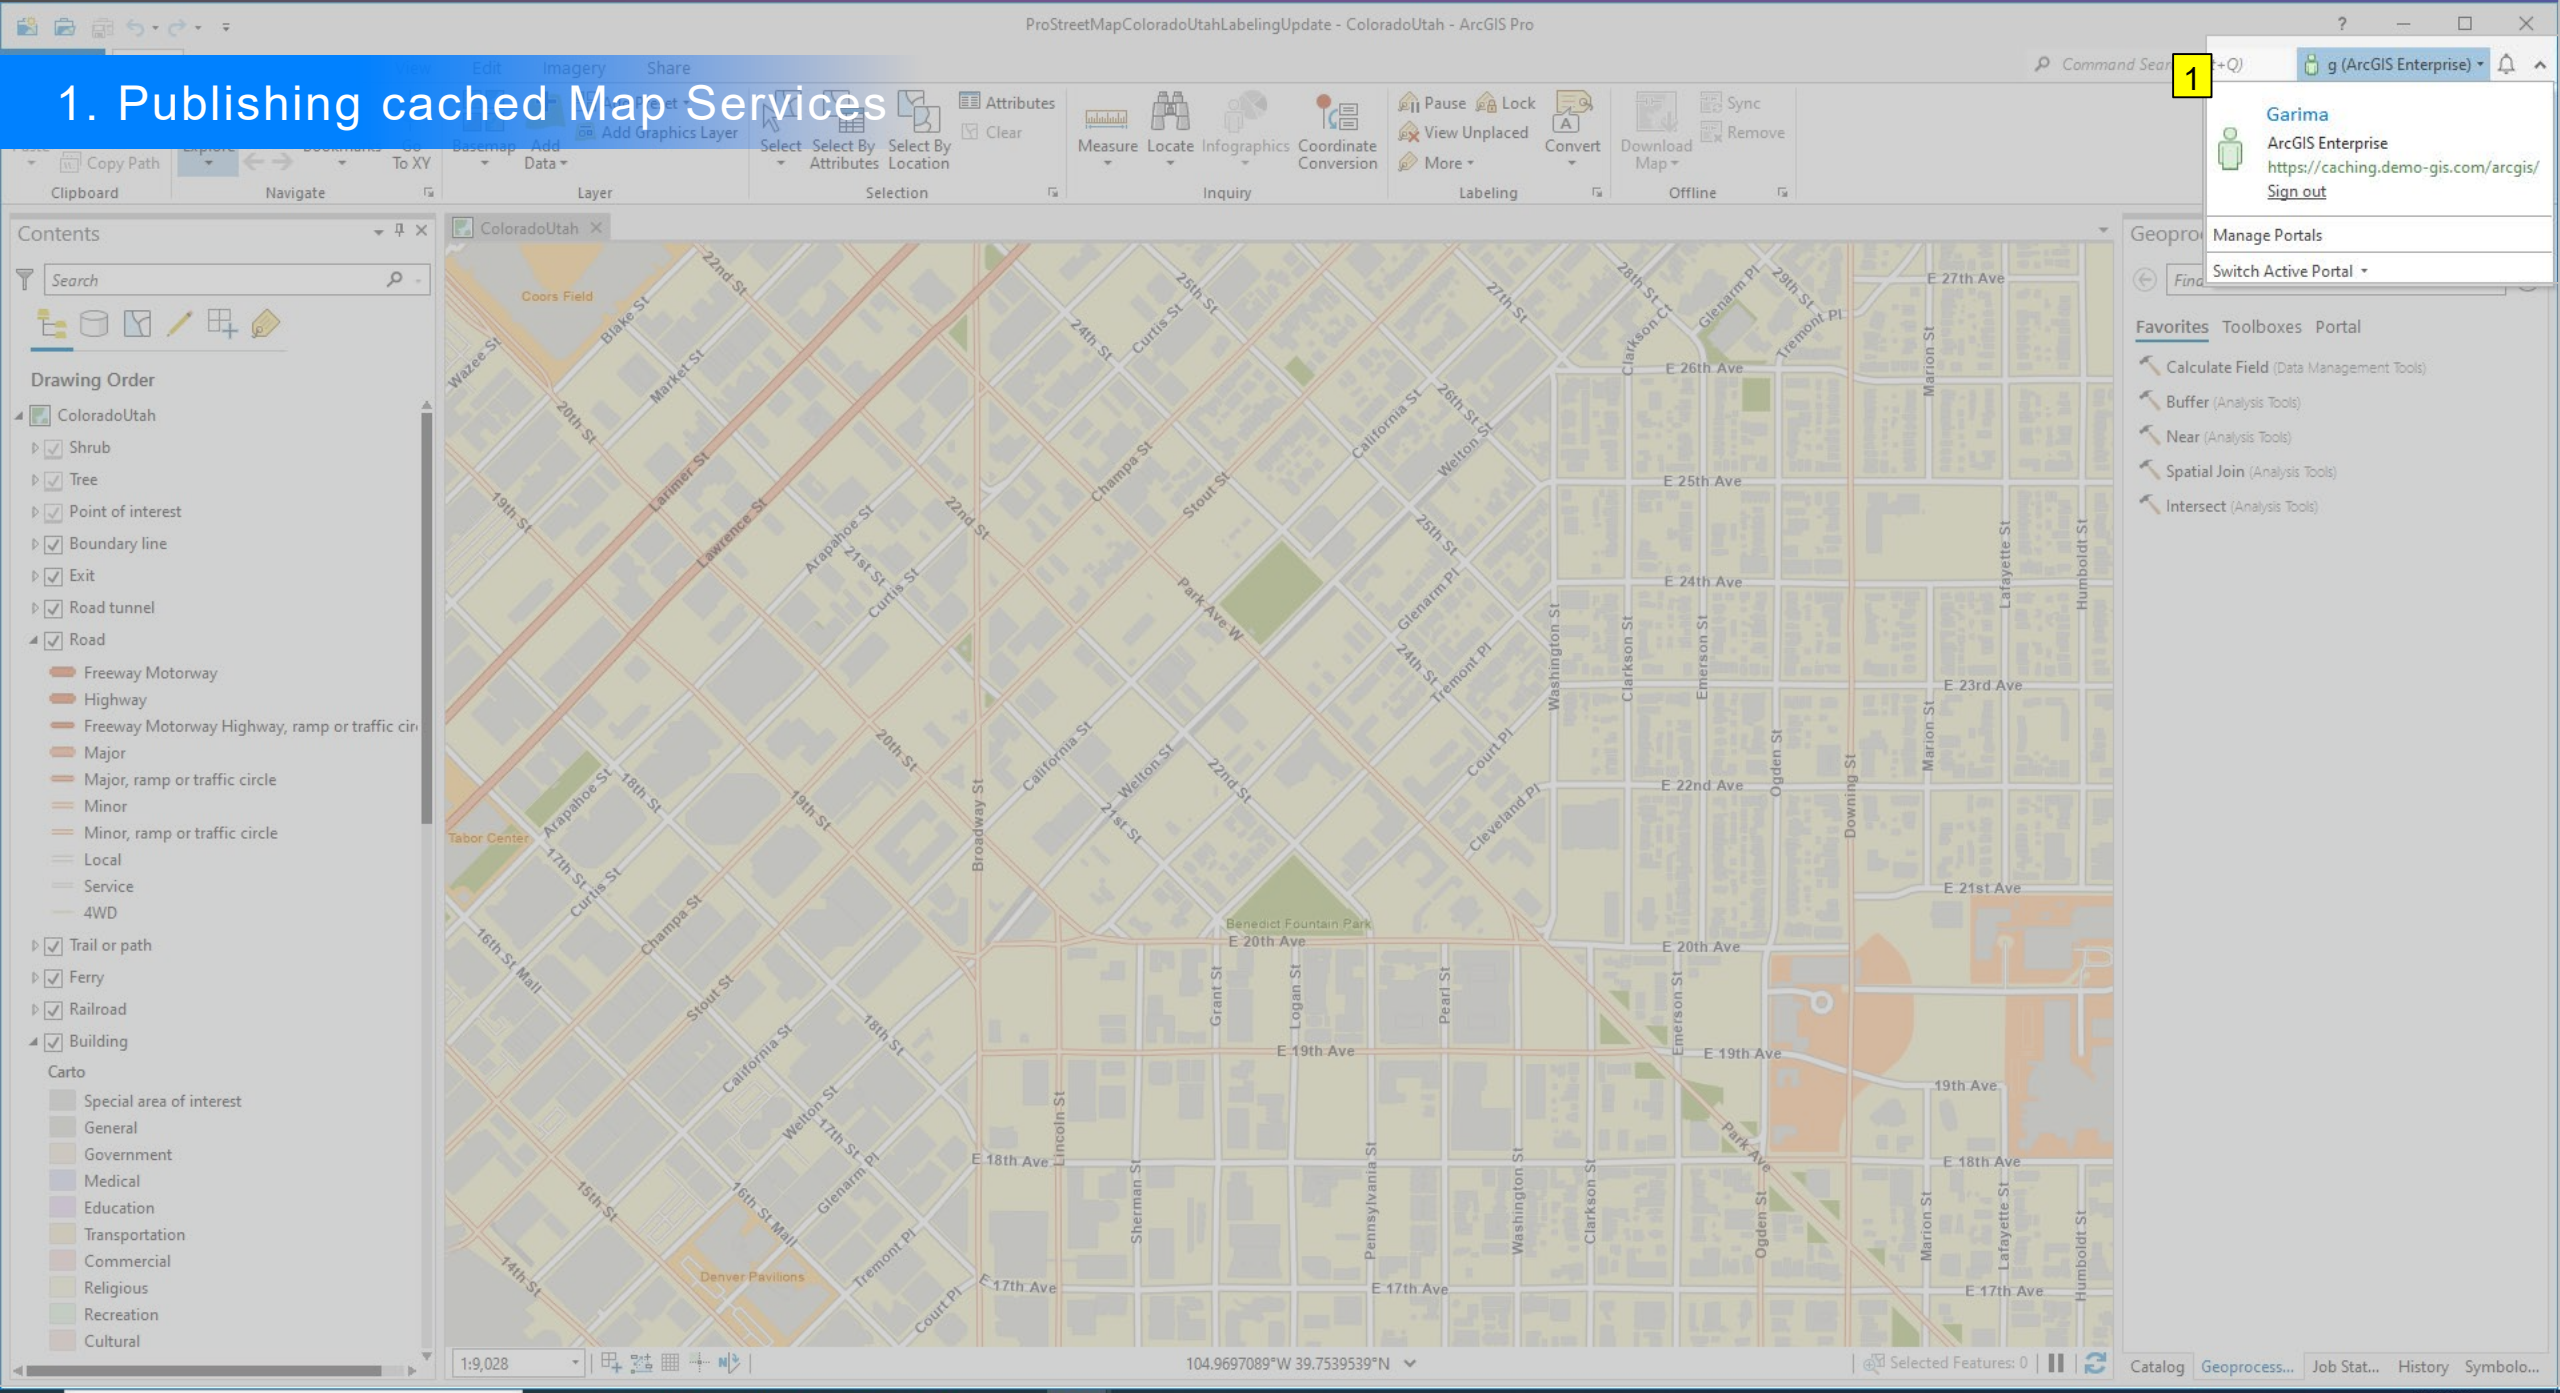

ProStreetMapColoradoUtahLabelingUpdate - ColoradoUtah - ArcGIS Pro

2

Project | Map | Insert | Analysis | Edit | Imagery | Share

Command Search (Alt+Q) | Garima (ArcGIS Enterprise 10.9)

Project | Map | Mobile Map | Layer | Geoprocessing | Deep Learning | Package

Web Map | Web Layer | Web Style | Locator | Web Tool | Jobs | Save Web Map | Replace Web Layer | Data Stores | Manage

Project Template | Map File | Layer File | Task Item | Save As

Print Map | Export Map | Capture To Clipboard | Output

- Publish Web Layer
- Overwrite Web Layer
- Save As Offline Service Definition

Contents

Search

Drawing Order

- ColoradoUtah
  - Shrub
  - Tree
  - Point of interest
  - Boundary line
  - Exit
  - Road tunnel
  - Road
    - Freeway Motorway
    - Highway
    - Freeway Motorway Highway, ramp or traffic cir
    - Major
    - Major, ramp or traffic circle
    - Minor
    - Minor, ramp or traffic circle
    - Local
    - Service
    - 4WD
  - Trail or path
  - Ferry
  - Railroad
  - Building
    - Carto
    - Special area of interest
    - General
    - Government
    - Medical
    - Education
    - Transportation
    - Commercial
    - Religious
    - Recreation
    - Cultural

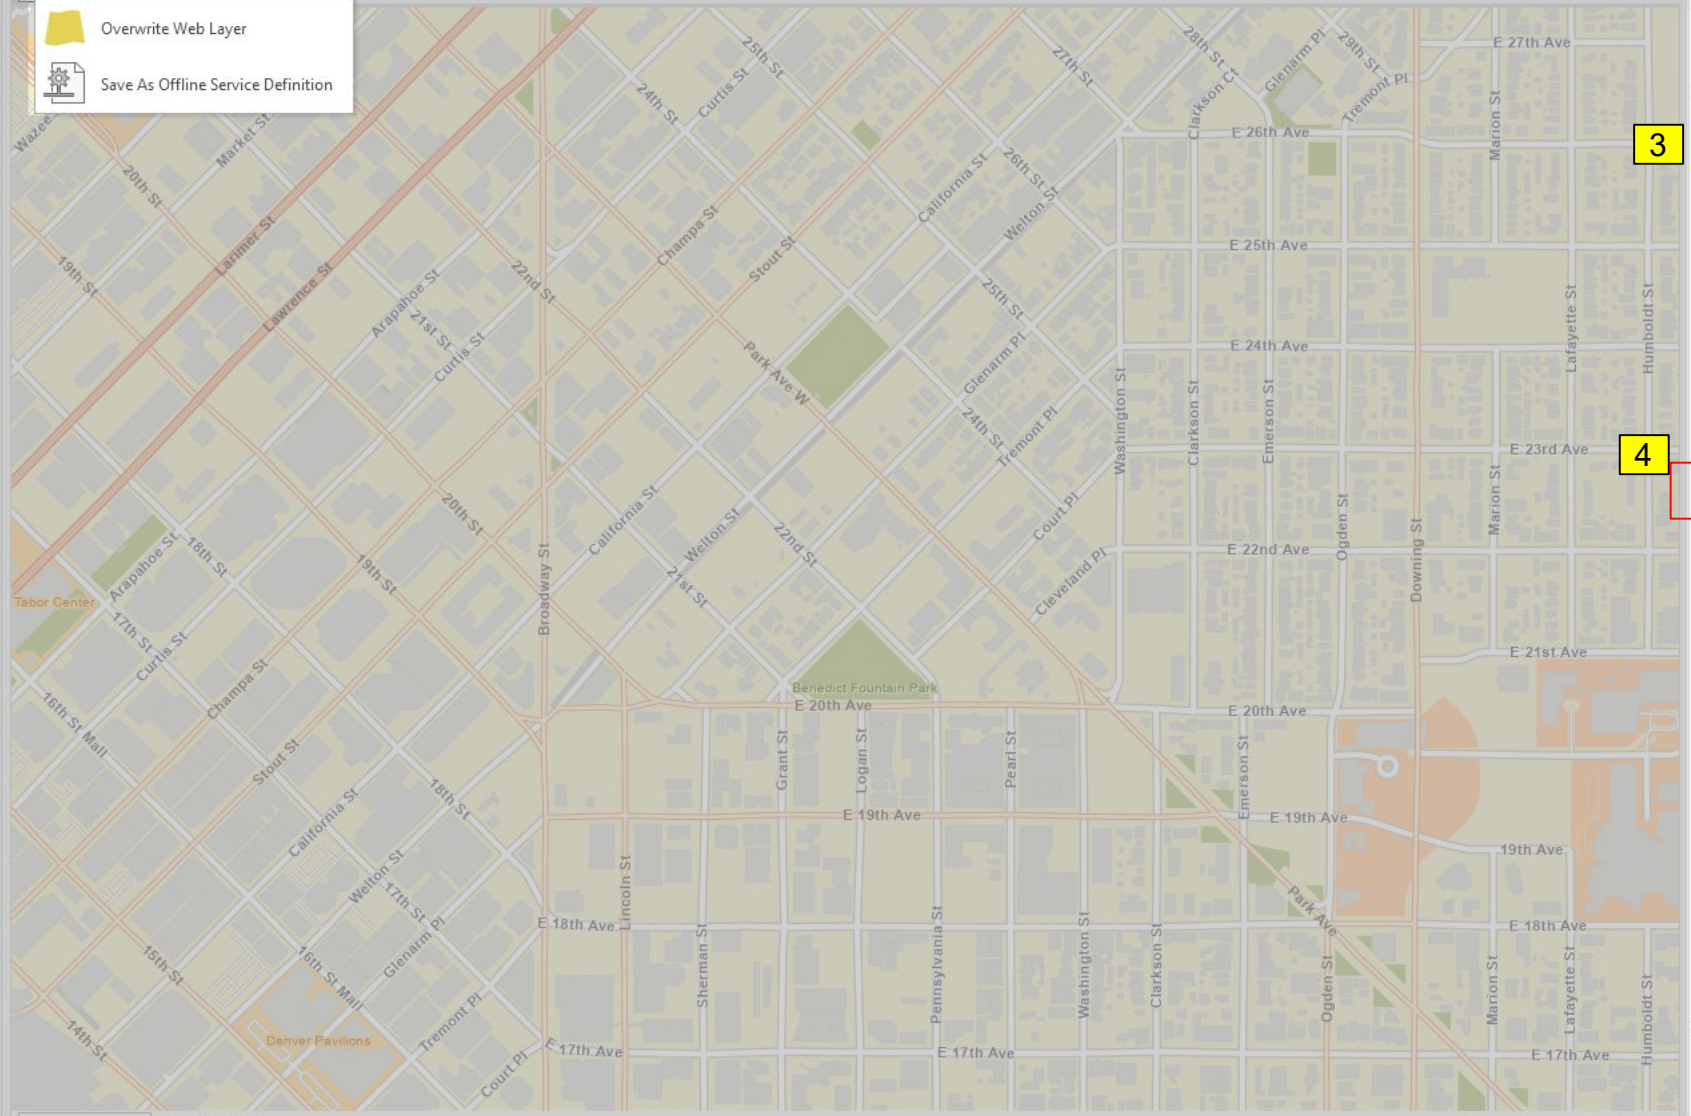

1:9,028 | 104.9941492°W 39.7561314°N | Selected Features: 0

Share As Web Layer

Sharing ColoradoUtah As A Web Layer

General | Configuration | Content

3

Item Details

Name: MyCityParcel

Summary: This is the map data for my City

Tags: MyCity, State, Country, Usage

4

Data and Layer Type

Reference registered data

- Map Image
  - Feature
- Vector Tile
  - Map Image and Feature

Copy all data

- Feature
- Tile
- Map Image
- Vector Tile
  - Feature

Location

Portal Folder: Select or create folder

Server and Folder: https://caching.demo-gis.com/arcgis (Hos

Select or create folder

Finish Sharing

Analyze | Publish | Jobs

Catalog | Share... | Geopr... | Job St... | History | Symb...

Type here to search

8:48 AM 5/26/2021

1. Publishing cached Map Services

ProStreetMapColoradoUtahLabelingUpdate - ColoradoUtah - ArcGIS Pro

Share As Web Layer

Sharing ColoradoUtah As A Web Layer

General | Configuration | Content

Map Image Properties

5

- Operations
- Dynamic Workspaces
- Layer Drawing
  - Dynamically from data
  - Using tiles from a new cache
  - Using an existing cache
- Tiling Scheme
  - ArcGIS Online / Bing Maps / Google Maps
  - Levels of Detail
  - 0 — 23
  - Minimum: Level: 5, Scale: 1:18,489,298
  - Maximum: Level: 19, Scale: 1:1,128
  - Countries
  - Building
  - Image format: PNG
  - Compression:
- Options
  - Cache automatically on the server
  - Cache manually on the server
  - Cache directory: System Managed Data Store
    - System Managed Data Store
    - mys3cache

6

Finish Sharing

Analyze | Publish | Jobs

7

Contents

Drawing Order

- ColoradoUtah
  - Shrub
  - Tree
  - Point of interest
  - Boundary line
  - Exit
  - Road tunnel
  - Road
    - Freeway Motorway
    - Highway
    - Freeway Motorway Highway, ramp or traffic circle
    - Major
    - Major, ramp or traffic circle
    - Minor
    - Minor, ramp or traffic circle
    - Local
    - Service
    - 4WD
  - Trail or path
  - Ferry
  - Railroad
  - Building
  - Carto
    - Special area of interest
    - General
    - Government
    - Medical
    - Education
    - Transportation
    - Commercial
    - Religious
    - Recreation
    - Cultural

ColoradoUtah

1:9,028 | 104.9774981°W 39.7430323°N | Selected Features: 0

Catalog | Share... | Geopr... | Job St... | History | Symb...

Type here to search

1:49 PM 5/26/2021

1. Publishing cached Map Services

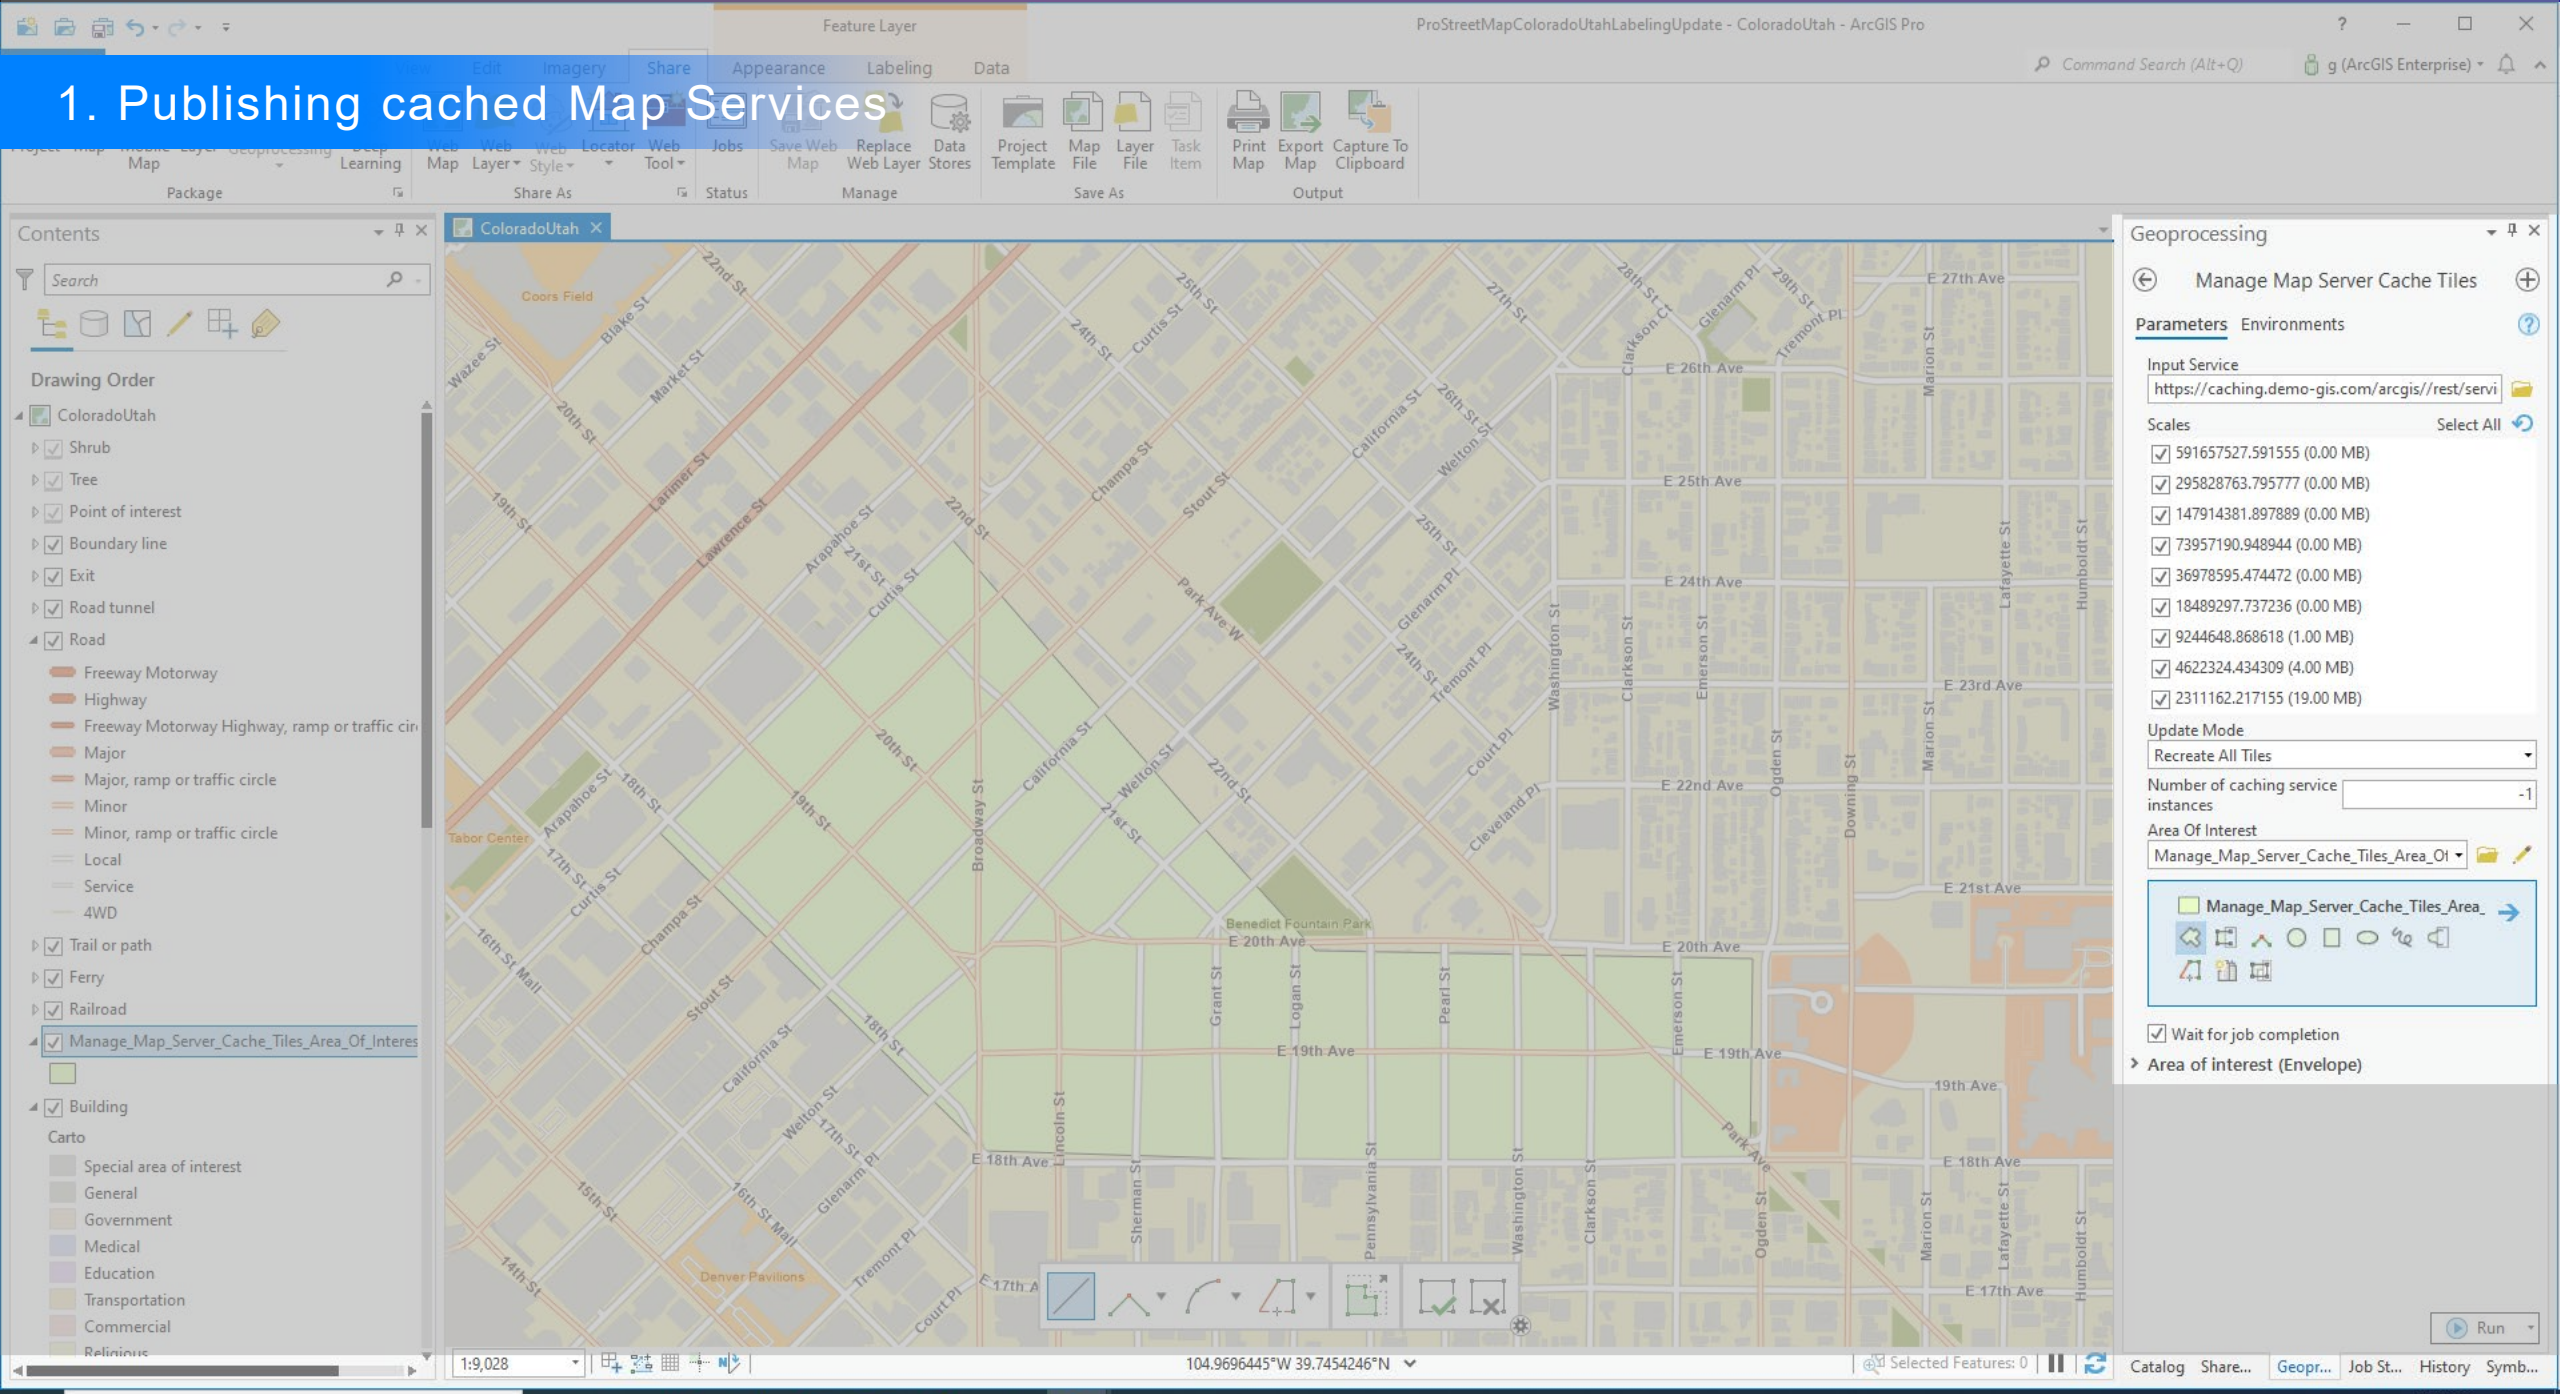

2. Cache Stores

Content - My Content

caching.demo-gis.com/arcgis/home/content.html?view=table&sortOrder=desc&sortField=modified&folder=admin1#content

Home | Gallery | Map | Scene | Groups | Content | Organization

g t
admin1

Content | My Content | My Favorites | My Groups | My Organization | Living Atlas

odified | Filter

Add a data store

STEP 1
Set basic information and type

STEP 2
Configure connection

STEP 3
Configure servers

Select the ArcGIS Server sites to which you want to add your data store. The site must have access to the connection or location you specified in the previous step.

	Server	Status
☐	https://caching.demo-gis.com/arcgis (Hosting Server)	✓ Normal

Back

Add data store | Cancel

ed
, 2021
, 2021
, 2021
, 2021
, 2021
, 2021
, 2021
, 2021
, 2021
, 2021

> Date modified

> Date created

3. Binding to existing cache content

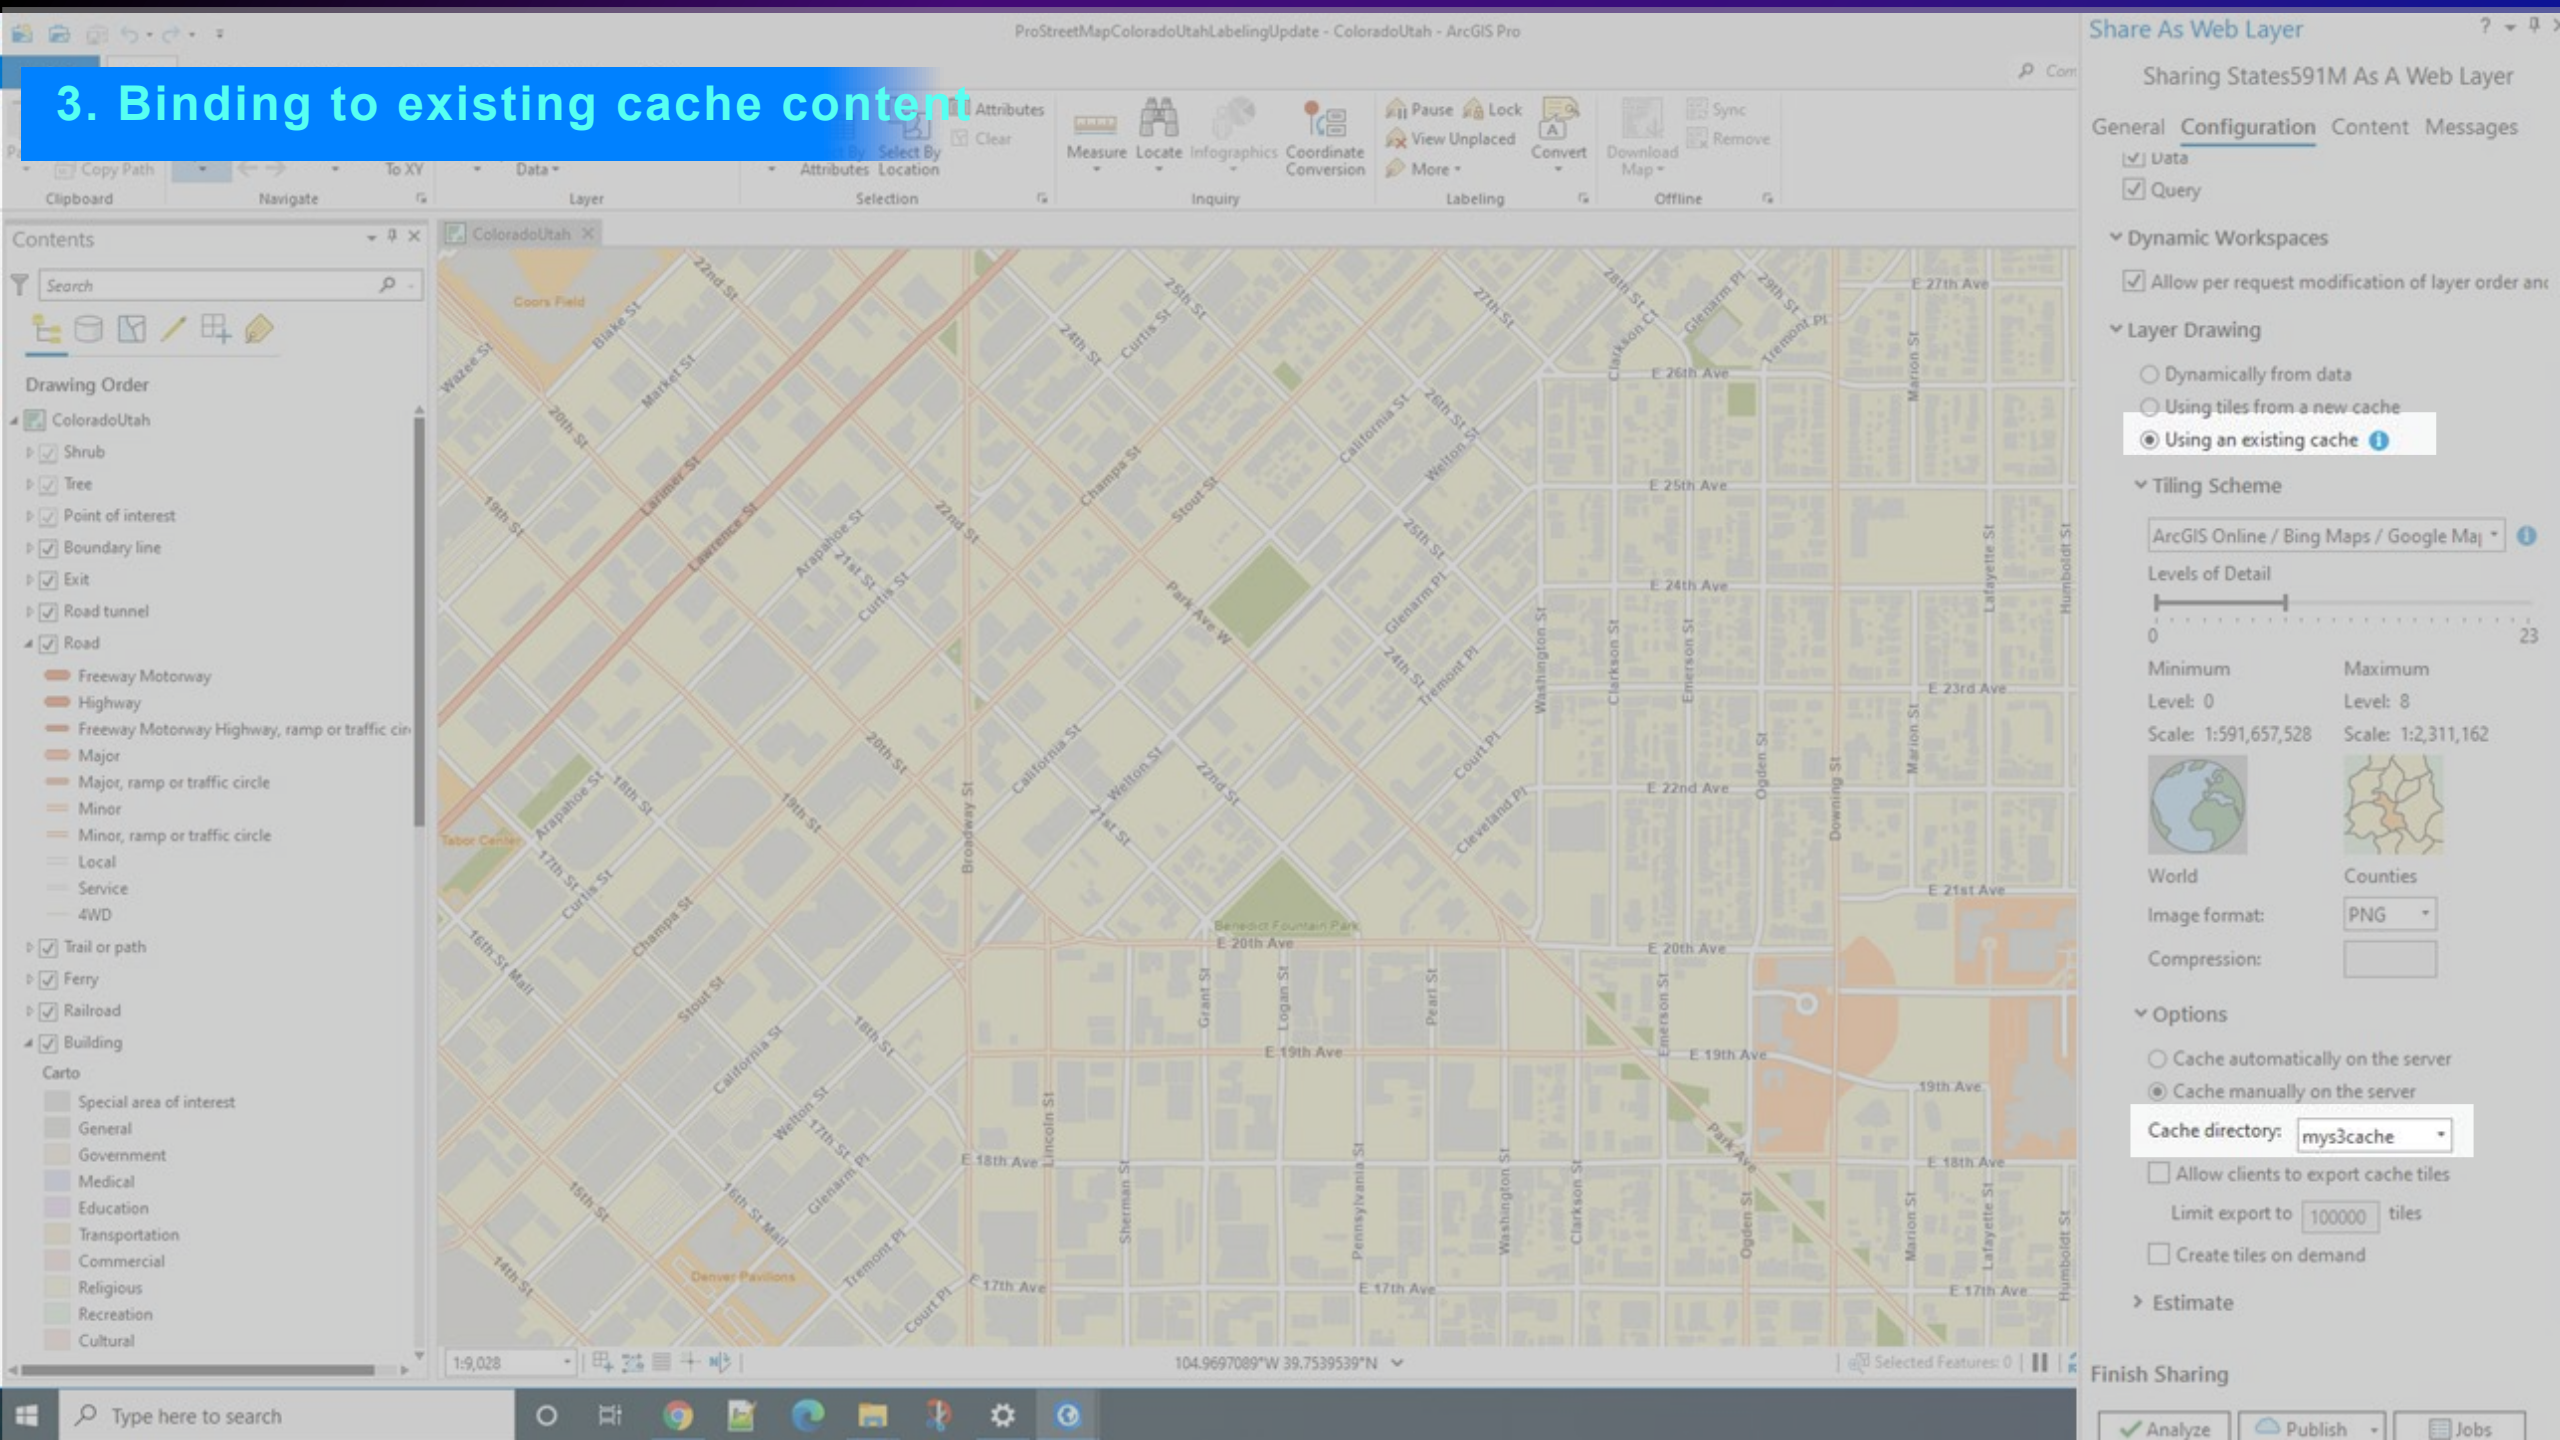

4. Scaling considerations for caching

Cached Map Services

4. Scaling caching instances

caching.demo-gis.com/arcgis/manager/#/overview

ArcGIS Enterprise Manager

Help

g t
admin1

Overview
Logs
Security
Services
System
Settings
Updates

Overview

Last refresh: a few seconds ago

Overview settings

Critical logs

View all

Time	Level	Message	Source
05/27/2021 12:58 AM	Warning	Status check failed for object store.	Enterprise Admin

System managed data stores

Object store	Pods: 2	Healthy	Validate
Relational store		Healthy	Validate
Spatiotemporal and index store	Pods: 3	Healthy	Validate

System services pods

CachingControllers	2 Up, 0 Down	Healthy
CachingTools	5 Up, 0 Down	Healthy
FeatureServiceTools	2 Up, 0 Down	Healthy
PublishingTools	4 Up, 0 Down	Healthy
ReportingTools	2 Up, 0 Down	Healthy

Utility services pods

Geometry	1 Up, 0 Down	Healthy
PrintingTools	2 Up, 0 Down	Healthy
OfflinePackaging	2 Up, 0 Down	Healthy
GeocodingTools	2 Up, 0 Down	Healthy

4. Scaling caching instances

ArcGIS Enterprise Manager

caching.demo-gis.com/arcgis/manager/#/services

ArcGIS Enterprise Manager

Help

g t
admin1

Overview

Logs

Security

1

Services

System

Settings

Services

Allocate resources for GIS, system, and utility services to scale number of pods, set resource limits, and st[illegible]nd start services.

2

GIS services | System services | Utility services

Search services names

1 - 10 of 10

Name	Type
SyncTools	GPServer
CachingControllers	GPServer
[illegible]eportingTools	GPServer
CachingTools	GPServer
SpatialAnalysisTools	GPServer
PublishingTools	GPServer
SharedMapServiceHost	MapServer
FeatureServiceTools	GPServer
SceneCachingTools	GPServer
SceneCachingControllers	GPServer

3

1

Items per page : 10

4. Scaling caching instances

ArcGIS Enterprise Manager

Help

Overview
Logs
Security
Services
System
Settings
Updates

← CachingTools

4

Overview | Settings

Status

Started

Stop service

Mode

Dedicated

Scaling

Number of pods	4

Resource limits

Minimum CPU	0.125
Maximum CPU	8
Minimum memory (GiB)	1.5
Maximum memory (GiB)	16

Service time

Maximum time a client can use a service (seconds)	3000000
Maximum time a client will wait to get a service (seconds)	60
Maximum time a service will take to be started (seconds)	600

Scene, Vector and Tile layers

Scene layers (.slpk)

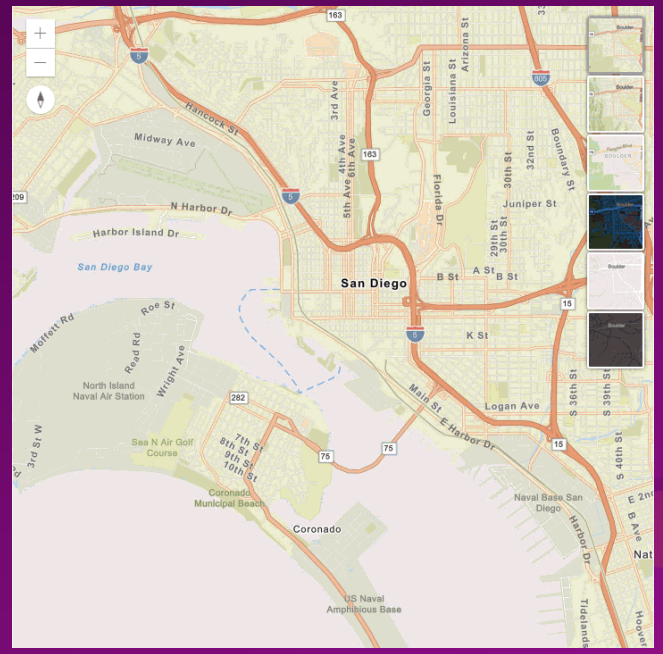

Vector Tile layers (.vtpk)

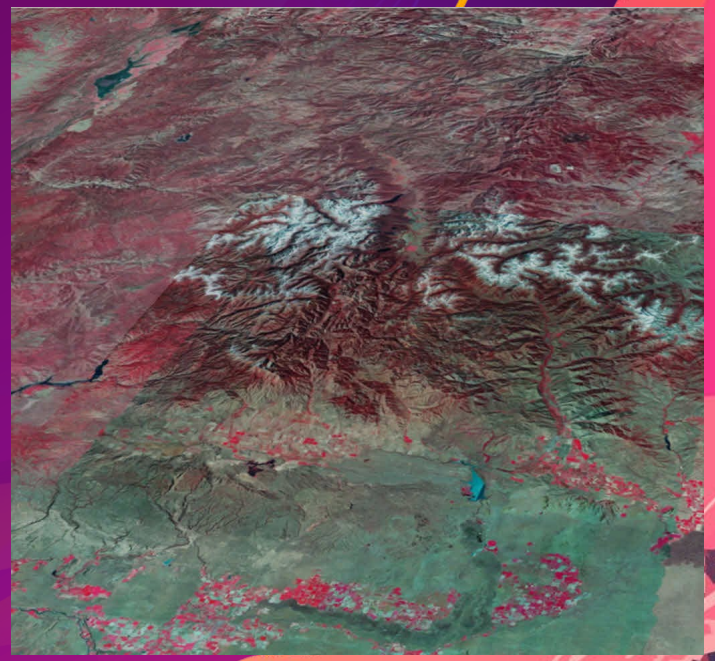

Tile and Elevation Layers (.tpkx)

Hosted Vector/Tile/Scene Layers

Scene, Vector or Tile layers

Extracted i3s or cache content

localhost:8888/notebooks/Documents/Demo_publish_SceneLayer_ref_cloud.ipynb

jupyter Demo_publish_SceneLayer_ref_cloud Last Checkpoint: an hour ago (unsaved changes)

Logout

File Edit View Insert Cell Kernel Widgets Help Trusted Python 3

Run Code

In [1]:
```
import json
from arcgis.gis import GIS
```

In [ ]:
```
# Connect to your GIS
# Create a connection to your portal for publishing
gis=GIS(url="https://dev0016310.esri.com/arcgis", username="[redacted]e)
```

In [ ]:
```
## Step 4a. Create a dictionary for your service configuration
ds = gis.datastore
service_conf= {
  "type": "SceneServer",
  "serviceName": "LAS_S3_point_cloud",
  "properties": {
    "pathInCachedStore" :"/v17_i3s/M_R_LAS.i3srest",
    "cacheStoreId": "16469fabfa274a39af059303cd307354",
  }
}
service_conf
```

In [ ]:
```
## Step 4b. Publish the scene layer content
##          located in data store managed by you as a scene layer

items = ds.publish(config=service_conf, server_id="OkzFw2JpsPgwetm7", folder="/",
                   description="BSL by python API",
                   tags="pythonAPI test").result()
items

items[0]
```

Summary

- Publishing, caching and scaling workflows are consistent across Windows, Linux & Kubernetes
- For faster cache generation use replicated content in folder data store
- System managed data store:
  - Map Image layer cache content
  - Hosted Scene, Vector and Tile layers
- Additional cache stores: AWS S3 and Azure blob stores
  - Bind to pre-existing cache
  - Publish Scene, Vector and Tile layers referencing pre-extracted cache or i3s content

Caching workflows can be precisely scaled, on per service basis, with fine grained resource management available on ArcGIS Enterprise on Kubernetes

ArcGIS Enterprise on Kubernetes Service Management and Tuning

Moginraj Mohandas

esri THE SCIENCE OF WHERE®

2021 ESRI USER CONFERENCE

ArcGIS Enterprise Manager – service management

- **Services**
  - Services power your geospatial capabilities
  - Different types – Map, Feature service, GeoProcessing…
  - Running inside pods – quick, scalable, isolated
    - Dedicated services run on dedicated pods
    - Shared services run on shared pods

Services

Clients
INGRESS
Shared Map
Catalog
Kubernetes
Docker
Spatial Analysis
Dedicated Map
Kubernetes
Docker
Dedicated Map
Dedicated Map
Kubernetes
Docker
Ignite
in-memory cache
Kubernetes
Docker
METERING
Prometheus
Grafana
Kubernetes
Docker
Configurations
Apps
JS API
Kubernetes
Docker
Tiles & Scenes
Shared Map
Kubernetes
Docker
Geocode
Publishing Tools
Kubernetes
Docker
Shared FS
Caching Tools
Kubernetes
Docker
MinIO
Blob Store (2x)
Kubernetes
Docker
Enterprise & Cloud
Data Stores
Portal App
Admin API
Kubernetes
Docker
Portal API
Help Doc
Kubernetes
Docker
Queue (2x)
Manager
Kubernetes
Docker
Postgres (2x)
Kubernetes
Docker
Elastic Search (3x)
Kubernetes
Docker
PERSISTENT VOLUME CONTROLLER
VOLUMES

ArcGIS Enterprise Manager – service management

- **Service Tuning/Scaling**
  - Depending on usage and demand, tune services
    - Use service usage statistics provided out of the box – exposes the Metrics Viewer to obtain rate of requests, response times, succeeded/failed requests etc.
  - Scale horizontally by upping the number of pods for a service
    - Dedicated -> Scale out the specific pods for a service
    - Shared -> Scale out the shared pods, think about moving to dedicated
  - Update resources for pods – Memory/CPU
  - Add more resources to the cluster as necessary(nodes)

Demonstration:

Experience for Service Management and Tuning

Moginraj Mohandas

Demonstration:

Usage Metrics

Shreyas Shinde

ArcGIS Enterprise Manager – service management

- **Additional considerations**

  - If using data from folder shares(for service publishing, caching workflows etc)
    - Admins are highly recommended to register root folder paths during org configuration.
      - Avoids disruption to the deployment when folder paths are registered after org is configured.
      - Publishers will not be interrupted when registering folder paths that contain their data, before publishing from Pro.
    - Registering brand new data paths after org is configured will restart many service pods. Can cause disruption.
    - Only administrators will be allowed to register any new data paths. Publishers will be blocked unless a root path has not already been registered.
    - If registering new data paths after configuring org, admins are highly recommended to do this during maintenance windows, to minimize disruption to services.

ArcGIS Enterprise on Kubernetes High Availability & Disaster Recovery

Laurence Clinton and Jon Quinn

2021 ESRI USER CONFERENCE

Overview

- Backup and Restore
  - What is included
  - Where is it stored
  - How are they created and restored
- High Availability
  - Architecture profiles

Backup and Restore

ArcGIS Enterprise Backups

- Backups are integral to minimizing data loss after a disaster or outage
- Integrated within ArcGIS Enterprise as opposed to an external, client-type tool

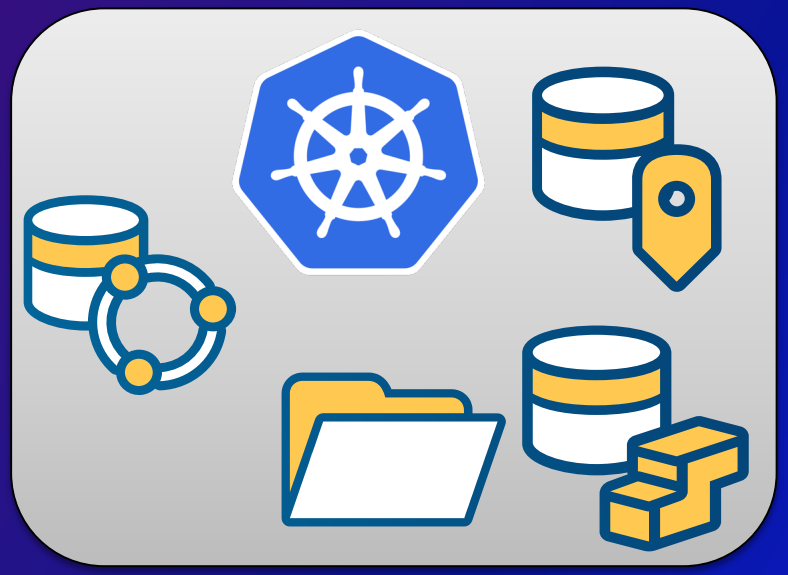

What is included in a backup?

- User data
  - Content and services
  - Groups, users, and other metadata
  - Map and tile cache
- Configuration information
  - Organization settings
  - Server type configuration
- Kubernetes infrastructure objects

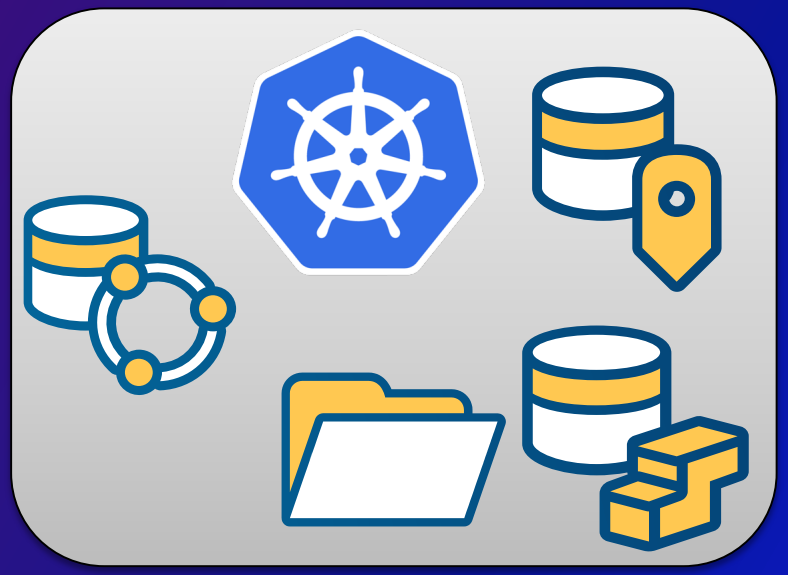

Where is it stored?

- Supported by either a pre-created or dynamically provisioned persistent volume
  - Stored in a object store, similar to where user data is stored within the deployment

*Support for cloud native storage options coming post 10.9

How do I create backups?

- All operations are completed using the Administrator API
  - Registering a staging location and backup store
  - Creating the backup
- Can be scripted using the language of your choice

What happens during a restore?

- Kubernetes infrastructure objects restored
- Configuration data restored
- User data restored

How do you restore a backup?

- Restoring a backup is done via the Administrator API

- Original deployment is accessible
  - Restore using the Administrator API

- Original deployment is not accessible
  1. Deploy ArcGIS Enterprise
  2. Use existing persistent volume to register new backup store
  3. Restore backup

High Availability

Architecture Profiles

- One size does not fit all
- Support for different profiles that match an organizations needs

Architecture Profiles

- Three default profiles:
  - Development
  - Standard Availability
  - Enhanced Availability

Development

- Non-production
- Easy to deploy
- No mandate for high availability or achieving a specific SLA
- Moderate tolerance for downtime

Development

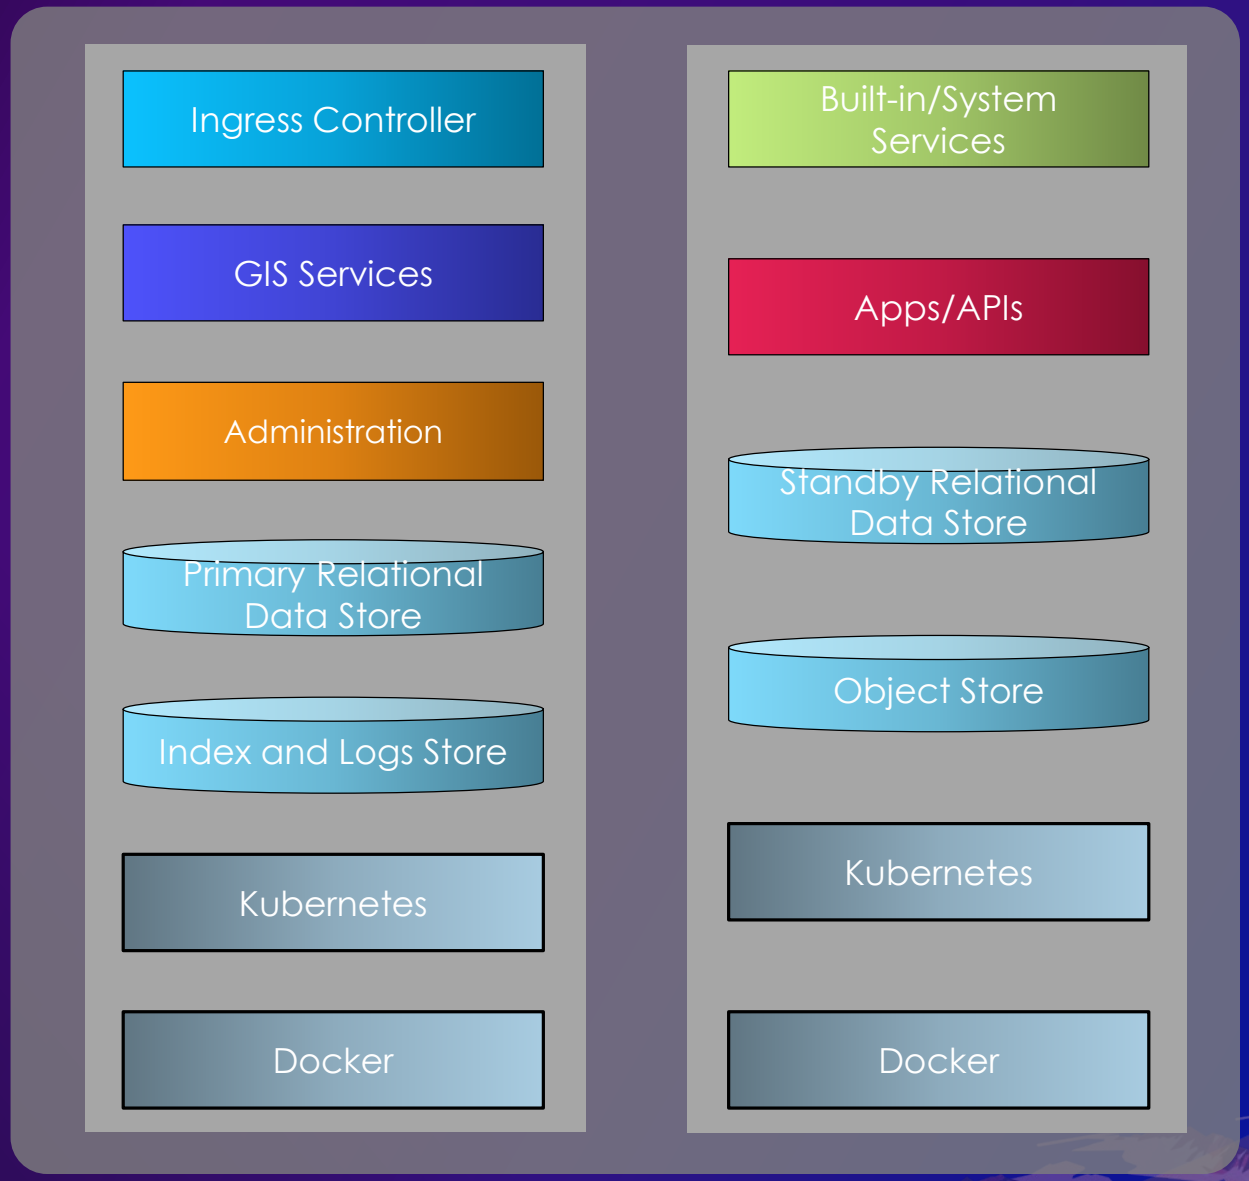

High Availability Profiles

- HA is a spectrum, not a switch
- Inverse relationship between amount of downtime acceptable (RTO) and cost
  - Lower downtime, higher cost
- Designed for production environments and for organizations that have an HA mandate
- No single points of failure
  - Redundancy
  - Scalability

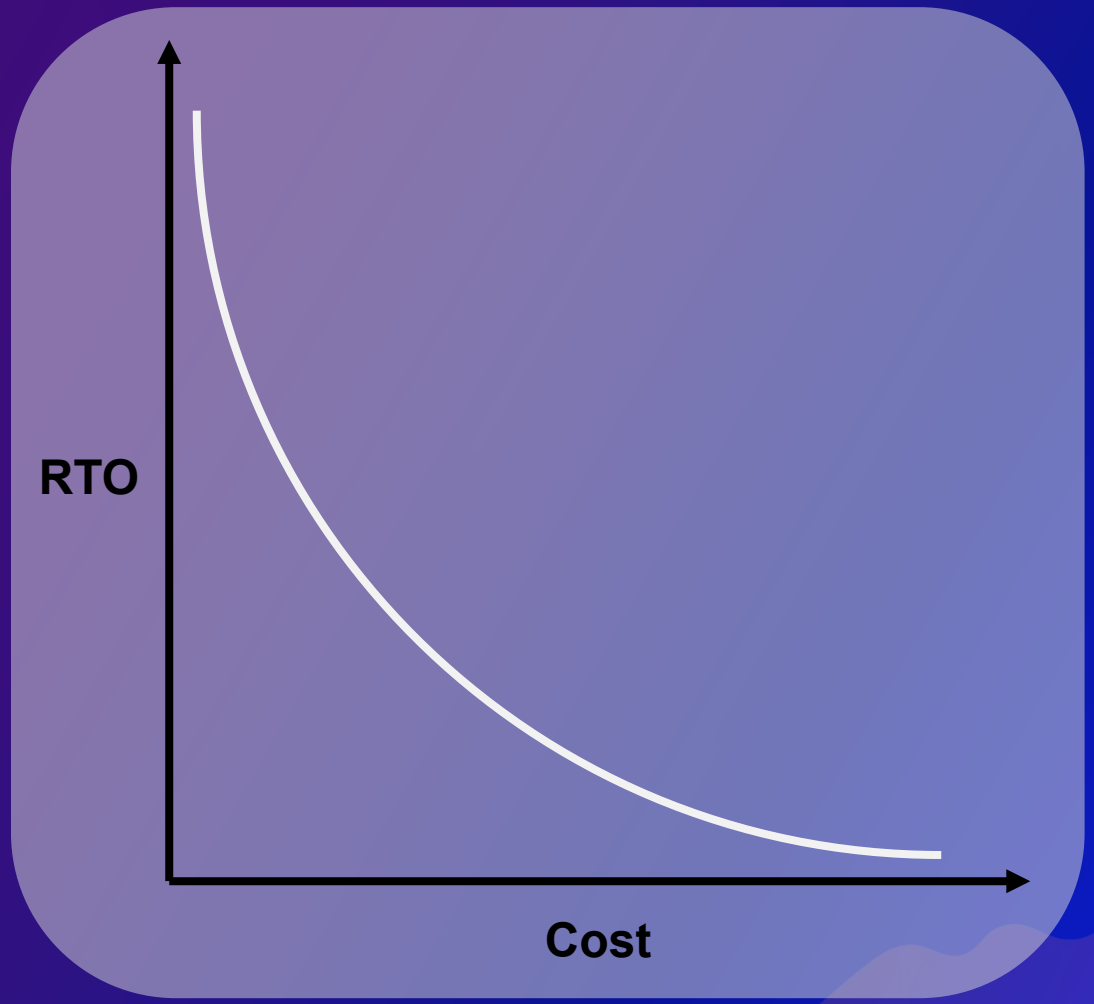

Standard Availability

- Minimum to moderate downtime
- Potential loss of functionality in the event of a failure
- Appropriate for medium sized customers

Standard Availability

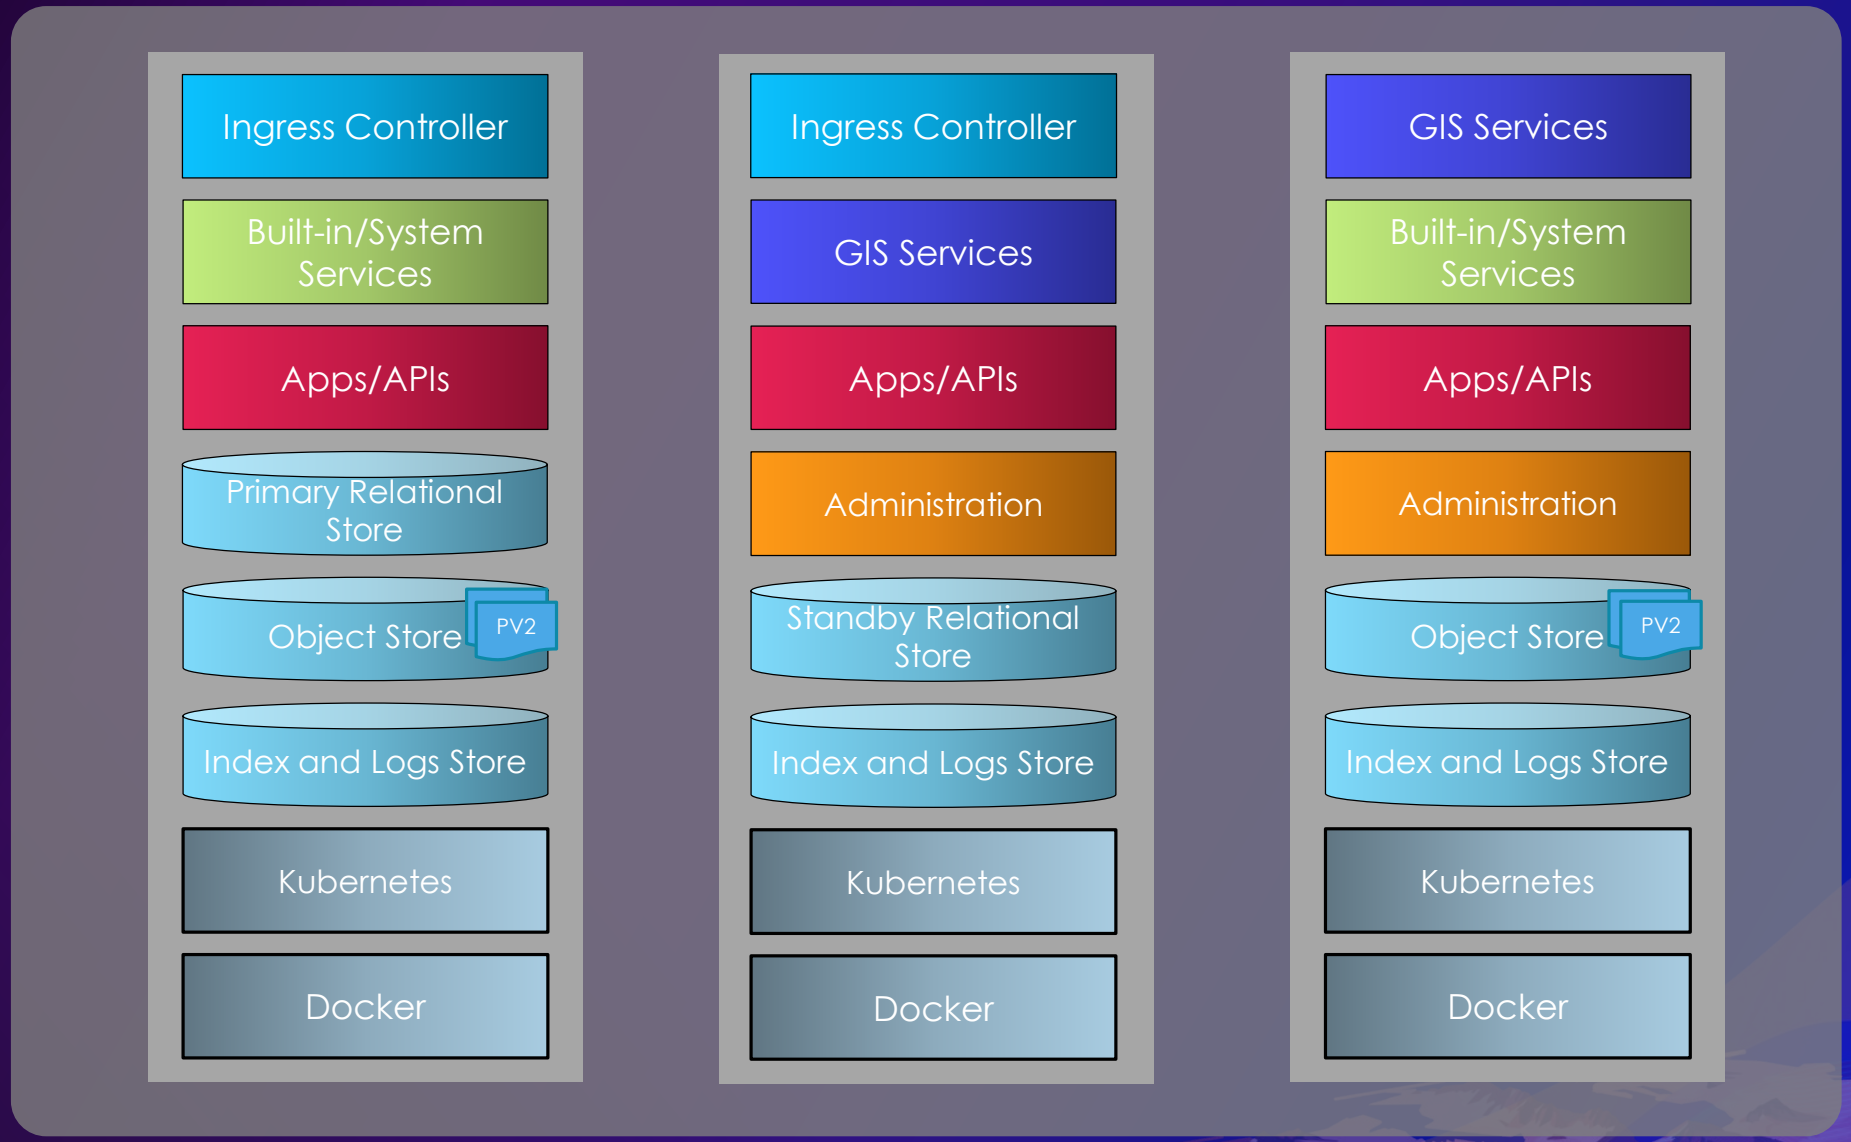

Enhanced Availability

- Minimum downtime
- No potential loss of functionality in the event of a failure
- Appropriate for large sized customers

Enhanced Availability

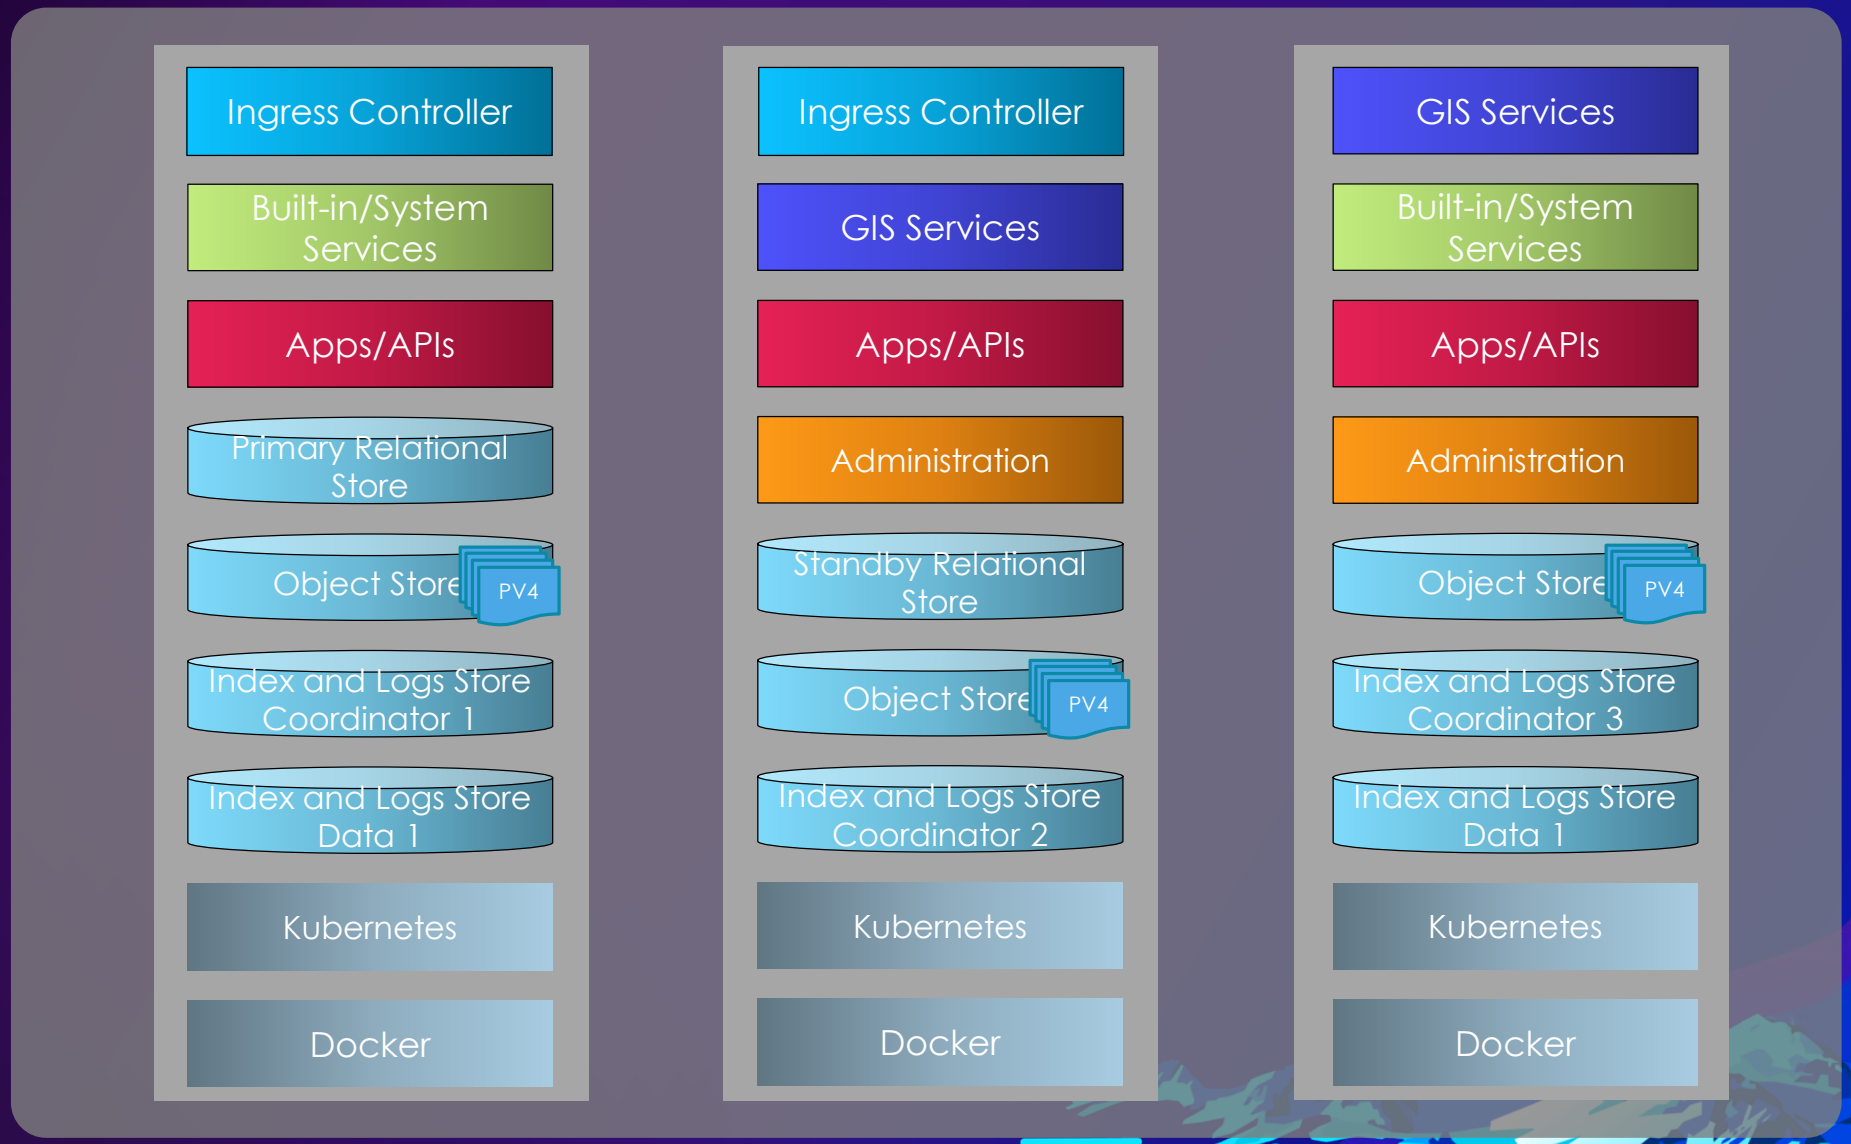

High Availability – Standard Availability vs Enhanced Availability

Category	Standard	Enhanced
Number of logs and index store pods	3	5
Logs and index store pods mode	Mixed mode	Dedicated mode
One object store pod failure	Read-only	Read-Write

High Availability and Data Stores

- Role and purpose of the data stores
  - Relational
  - Object
  - Spatiotemporal and Index
- How high availability works with each store type
- Demo

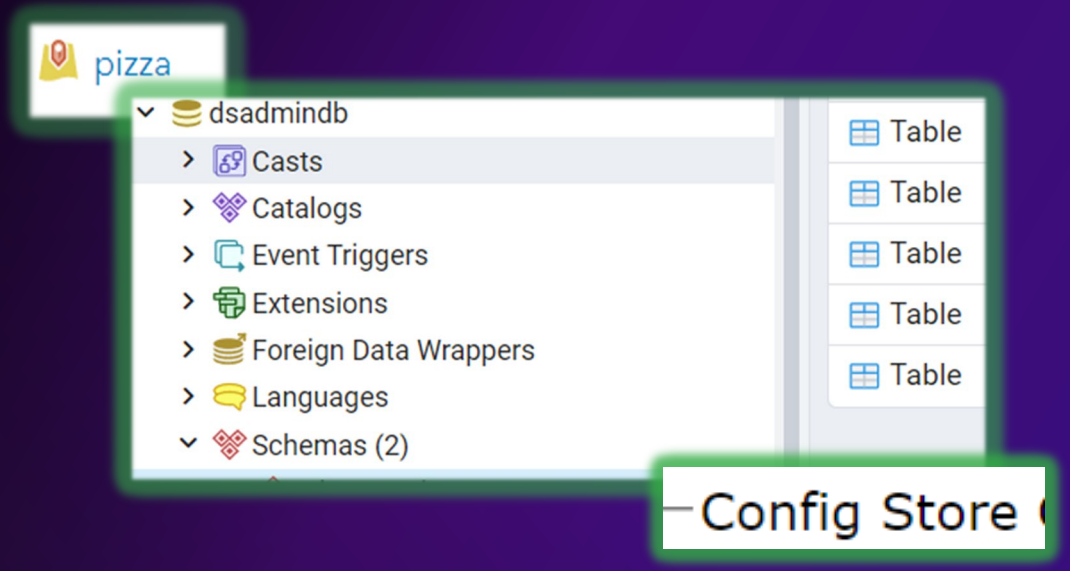

Relational role

- Storage for
  - **Hosted feature data**
  - **Administration (geowarehouse)**
  - **Configuration settings**

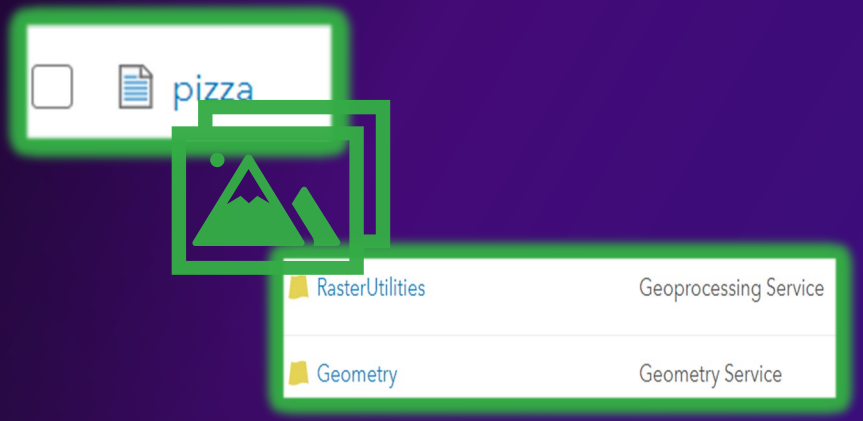

Object store role

- Storage for
  - **Uploaded items and saved content**
  - **Hosted tile and image layer caches**
  - **Geoprocessing output**

INFO	Backup the pgsql directory
INFO	Starting watchdog thread run.
INFO	

Type	Count in store	Count in index
users	47	47
groups	449	449
search	20236	20236

Spatiotemporal and Index role

- Storage for
  - **logs**
  - **Portal indexes**
  - **High velocity data (future)**

Relational data store high availability

- Primary Standby architecture
- Kubernetes operators determine failover
  - Readiness (ready for traffic?)
  - Liveness (restart pod?)
- Decisions
  - Recreate pod
  - Failover to standby

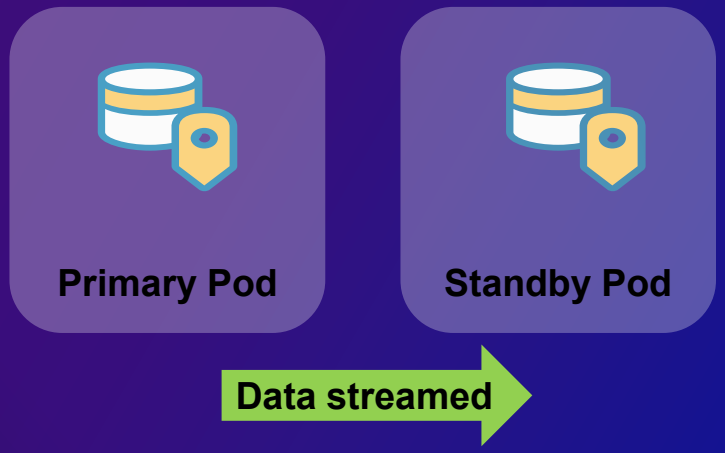

Object store high availability

- Liveness and readiness checks
- Configured as standalone or cluster
  - 1 PV as standalone or multiples of 4 PVs for cluster
- If greater than half nodes\disks available
  - > Read write

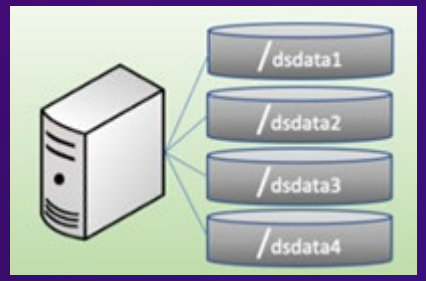

Object store pod

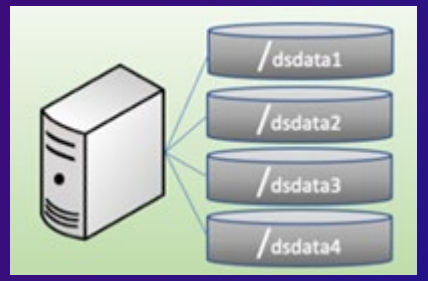

Object store pod

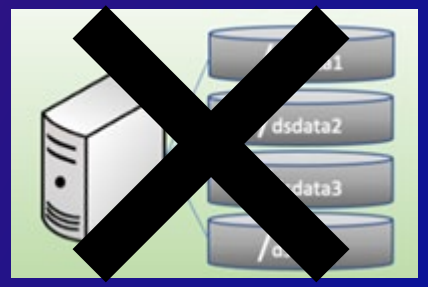

Not available

Spatiotemporal and index store high availability

- Liveness and readiness checks
- Sharding replication for HA

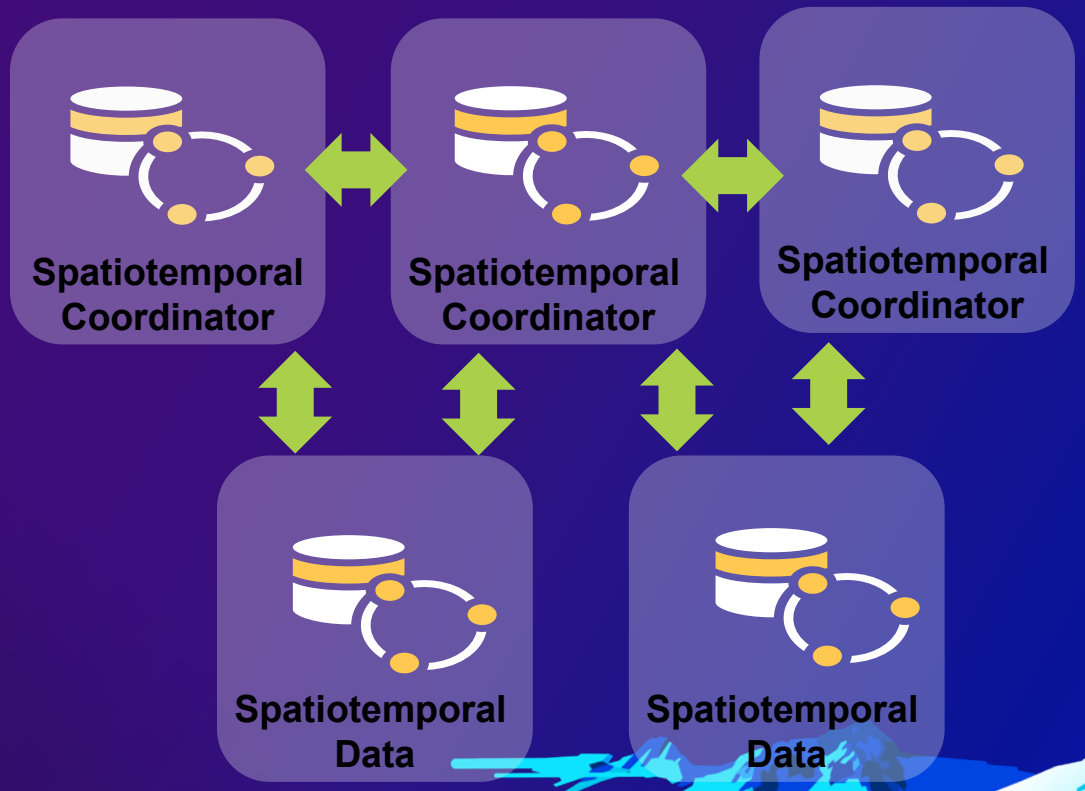

Demonstrations:

High Availability Liveness and Failover

Laurence Clinton

esri®

THE SCIENCE OF WHERE®

Copyright © 2021 Esri. All rights reserved.