ArcGIS Enterprise: Data Storage Strategies

Philip Heede, Hilary Curtis

Agenda

- What is a data strategy and why would I need one?
- Storage options in ArcGIS Enterprise
- Technical architecture of data in ArcGIS Enterprise
- Example data strategies
- Related topics and sessions

Note as the software evolves, so does this information!

What is your role?

- DBA?
- System Architect?
- Analyst?
- Developer?
- Executive?
- All of the above?

ArcGIS Enterprise: Data Storage Strategies

Your organization’s plan for achieving its goals.

What is a data strategy?

A comprehensive plan for how your organization will *store*, *access*, and *manage* your GIS data.

A data strategy is feasible, economical, tailored to your workflows and users, and evolves as necessary.

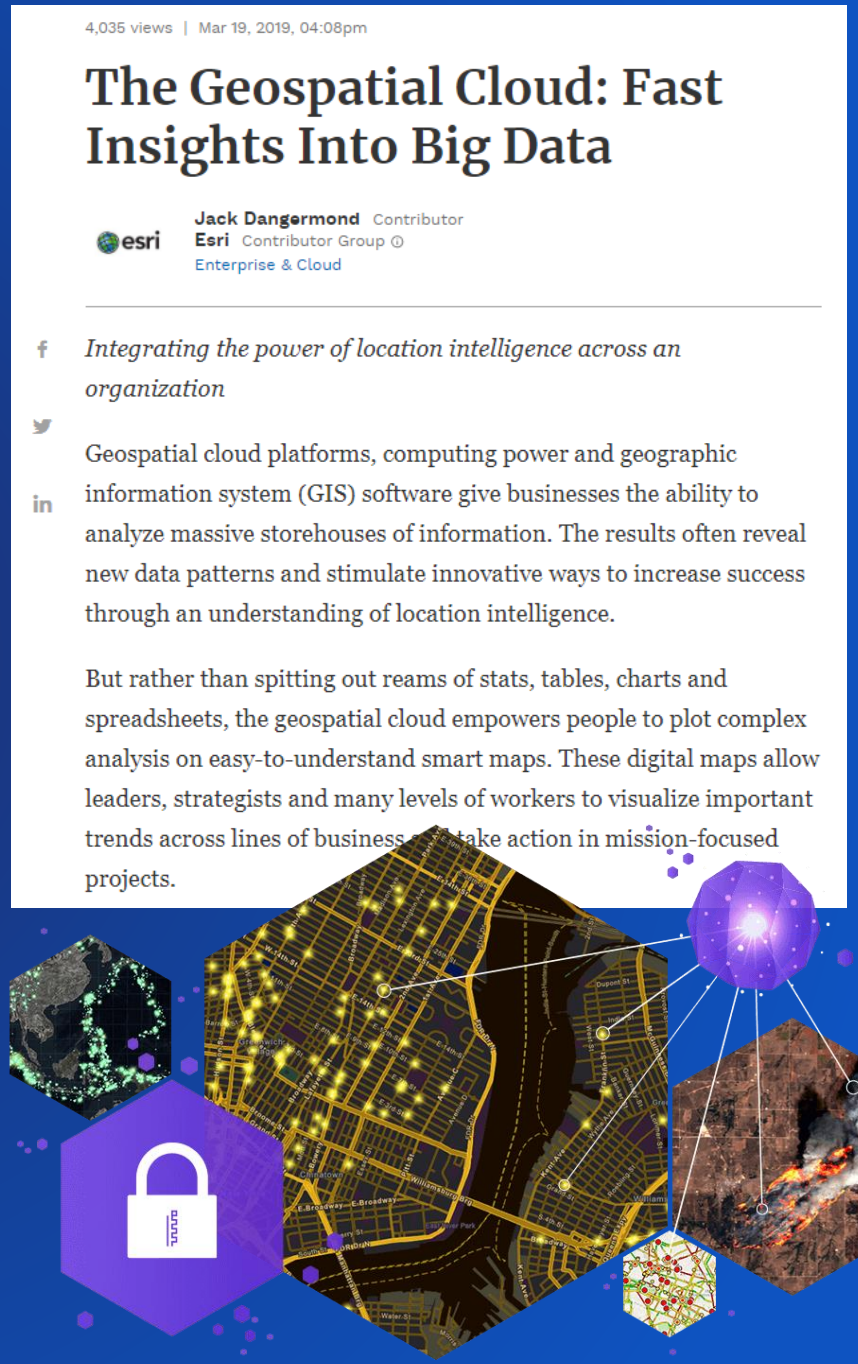

What is a data strategy?

A data strategy can take on many forms, from prescriptive manuals, to checklists, to general user guidelines and handbooks.

Whatever form it takes, a data strategy as a integral part of GIS.

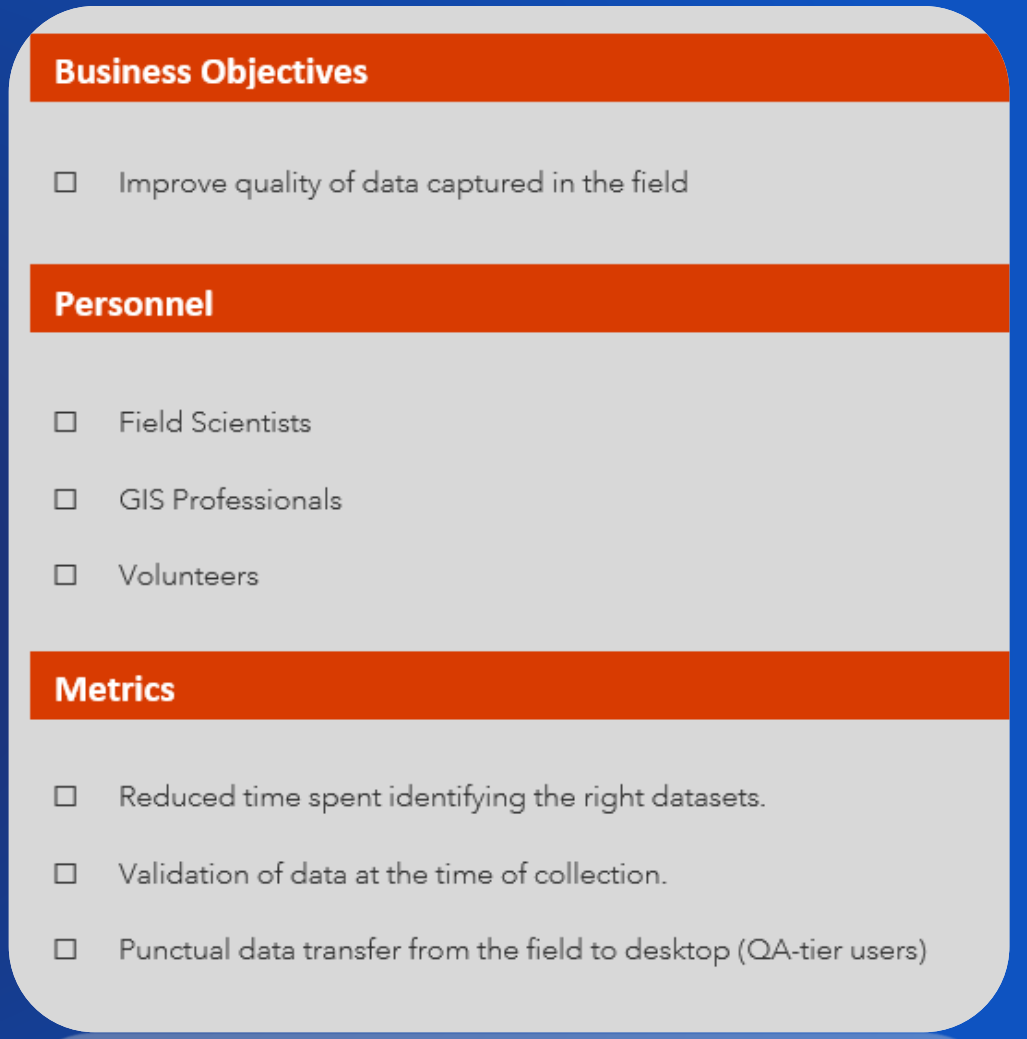

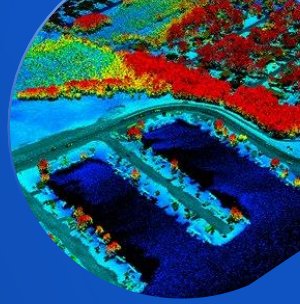

Why is it important to have a strategy for data?

You are working with more data than ever before:

Imagery

3D

Urban

Raster

Real-time

Indoor

Big data

Field

Demographic

Third party

Living Atlas

Unstructured

Vector & tabular

Cloud

Utility networks

Drone

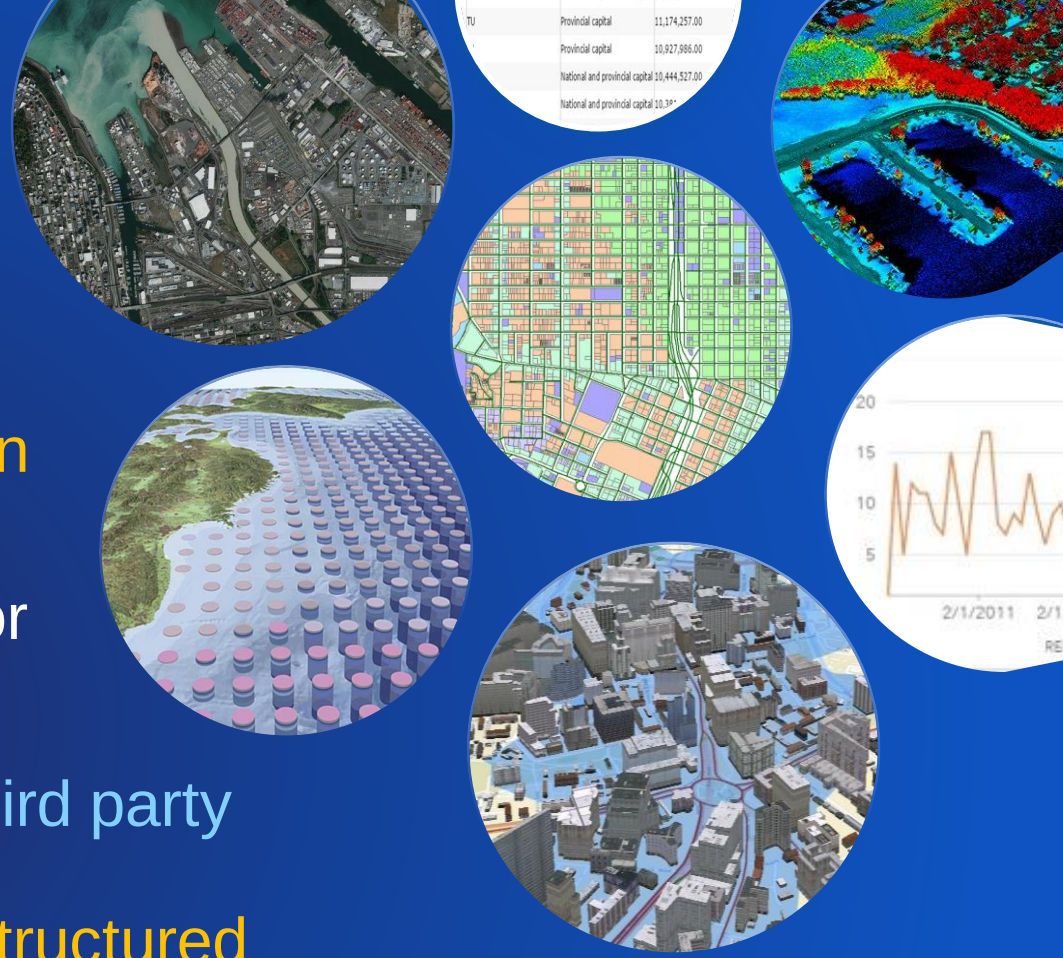

ArcGIS Enterprise supports your data workflows

A data strategy makes the best use out of your data.

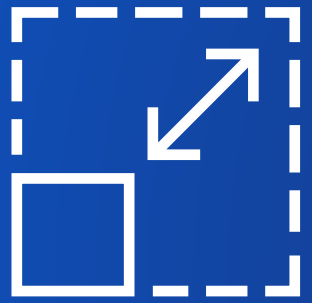

Flexible

Gives you room as data, workflows, and your user base changes and grows.

Accessible

Enables your users to access data when and where they need it (mobile, desktop, web / via direct connections, services).

Functional

Provides the right capabilities and functionality to successfully execute your workflows across your organization.

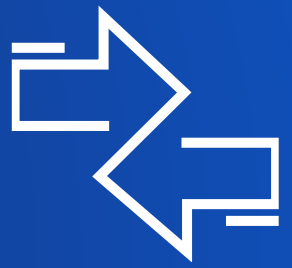

Interoperable

Integrates with other parts of the ArcGIS platform and technology as needed.

TU	Provincial capital	11,174,257.00
	Provincial capital	10,927,986.00
	National and provincial capital	10,444,527.00
	National and provincial capital	10,38*

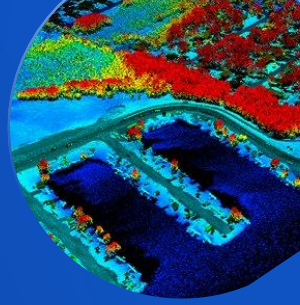

It also gives you a strong foundation to bring on new challenges, workflows, and innovation.

Options for data have evolved throughout the years…

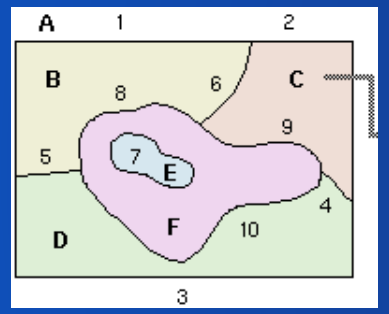

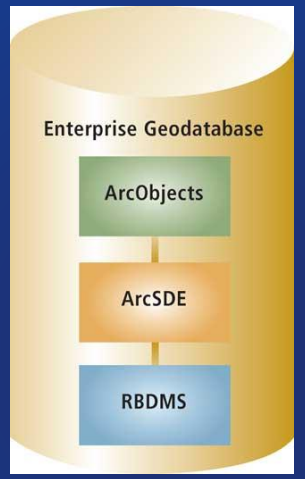

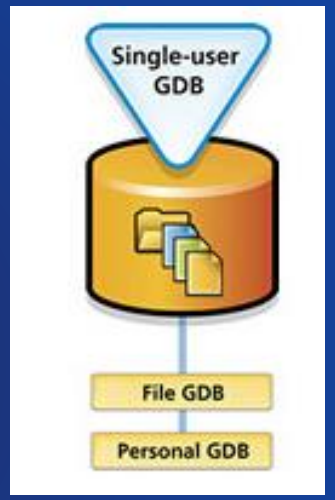

1990s

Coverages

Shapefiles

SDE

(enterprise geodatabase)

Personal geodatabase, file geodatabase

2014

ArcGIS Data Store

Available options in ArcGIS Enterprise today

Geodatabases

- Enable on top of commercial RDBMS
- Large scale, multi-user, authoritative data
- Spatial and attribute integrity across datasets
- Versioning, archiving

Folders & files

- Local or network,
- E.g. file geodatabase
- Storage for different file formats (csv, shp, etc)

ArcGIS Data Store

- Storage included with ArcGIS Enterprise
- 3 different flavors
- Powers hosted data
- Feature data, 3D scenes, high volume real-time and big data

Cloud storage

- Amazon, Azure integrations
- Store map and image caches
- Optional output of raster analytics
- Input and output of big data analysis (vector, tabular)

Big data storage

- Hadoop, Hive
- Input and output of big data analysis (vector, tabular)

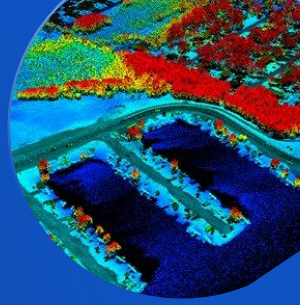

With all of these options, it can seem daunting at first.

(That's probably why you are here!)

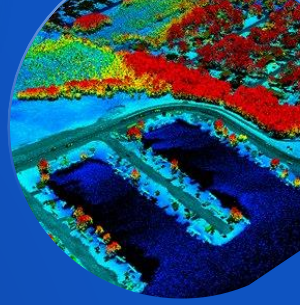

We can break down a lot of these concepts using two terms:

User managed
ArcGIS managed

What does this mean?

User managed:

- Data storage that you manage independently
- You provision, scale, tune, the underlying database
- You make it accessible to ArcGIS Enterprise by registering it

Direct management of the underlying storage

ArcGIS managed:

- Data storage included with ArcGIS Enterprise
- You install the software component as part of your ArcGIS Enterprise deployment

Management through ArcGIS interfaces and APIs

Relationship to data

User managed
(enterprise geodatabase)

The system references the data in place

ArcGIS managed
(hosted, ArcGIS Data Store)

The system hosts (stores) the data for you

This defines the term ‘hosted.’

Storage types

User managed
(enterprise geodatabase)

- Enterprise geodatabase
- Cloud storage
- File shares
- Big data storage

ArcGIS managed
(hosted, ArcGIS Data Store)

ArcGIS Data Store
- Relational
- Tile Cache
- Spatiotemporal

Publishing and access

User managed
(enterprise geodatabase)

- Data doesn’t move: referencing in place
- Accessed through database connections, REST services and items in portal
- Delete the service, data remains

ArcGIS managed
(hosted, ArcGIS Data Store)

- Copy data or publish directly in your portal
- Accessed through REST and items in portal
- The data *is* the service

Use cases

User managed
(enterprise geodatabase)

- Authoritative system of record
- Utility networks and parcels
- Need strict spatial and attribute quality
- Support for multi-user versioning workflows
- Comprehensive, relational database

ArcGIS managed
(hosted, ArcGIS Data Store)

- Often used for self-service portal workflows
- Good alternative for storing file-based data
- Some advanced options (domains, views, etc)
- Relatively isolated, standalone datasets

Whitepaper:
Data in ArcGIS

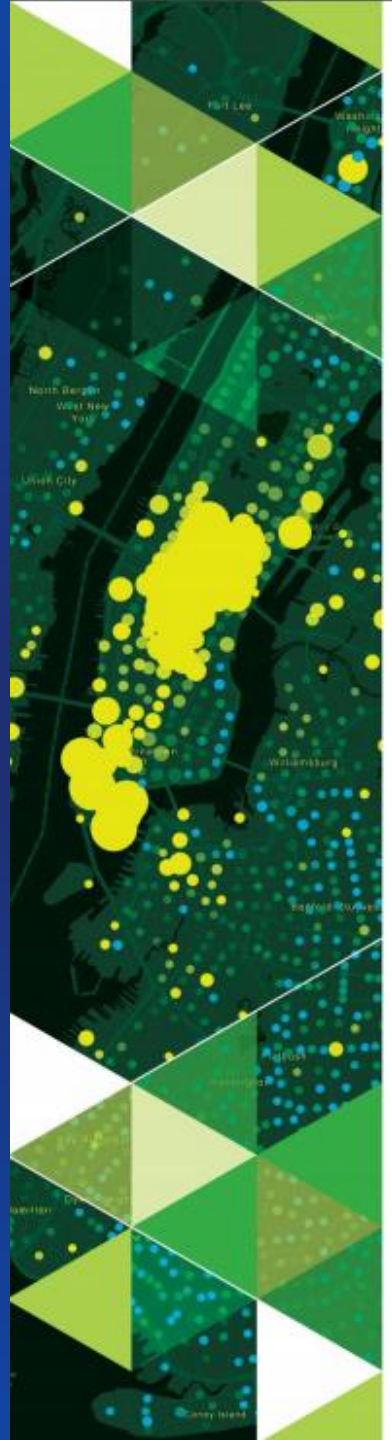

AN ESRI
WHITE PAPER

JUNE 2018

Data in ArcGIS: User Managed and ArcGIS Managed

380 New York Street
Redlands, California 92373-8100 USA
909 793 2853
info@esri.com
esri.com

esri | THE SCIENCE OF WHERE

Whitepaper:
Data in ArcGIS

The below table includes some of the differences between ArcGIS managed and user managed data:

	User managed (enterprise geodatabase)	ArcGIS managed (hosted data in the ArcGIS Data Store)
Database Technology	Option of RDBMS (Oracle, SQL Server, PostgreSQL, Informix, Db2, SAP HANA, etc.)	The database is chosen by ArcGIS; not a bring-your-own-database nor a general-purpose database.
User Access	Multiuser, with both direct connect from desktop applications and via REST-based map and feature services.	Multiuser, solely via REST-based map and feature services.
Rendering	Map services can render the output server-side using map image layers. Both map and feature services can render client-side via feature layers. Map image layers support advanced renderers and cartographic options not available in feature layers.	Hosted feature layers only support client-side rendering with out-of-the-box symbol set and cartographic options.
Versioning and archiving	Supports both traditional versioning and branch versioning. Archiving historical snapshots is supported.	Versioned editing is not supported; the last edit submitted is stored for the feature layer. No archiving.
Topology	Topology rules can be created and enforced.	Topology rules not supported.
Scalability	The underlying RDBMS can be scaled to support a large number of users, editors and data using native RDBMS features and scaling of the server hosting the database. Individual map and feature services that reference user managed data sources can have a substantial memory footprint on the server. Typical configurations can support in the range of hundreds of services.	The relational data store type of ArcGIS Data Store can be scaled vertically; adding more capacity & resources to a single machine. A secondary machine can be added as passive backup and a failover host. Individual hosted feature layers have very low memory footprint on the server. Hundreds to thousands of hosted feature layers can be supported on a standard configuration.
Use as managed database for hosting server	Not supported, however the geodatabase can be registered as an additional data store, or the managed database on a non-hosting server.	Supported – only the relational ArcGIS Data Store can be registered as the managed database for the hosting server site.

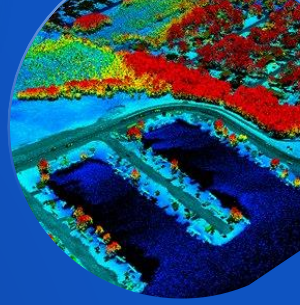

The ArcGIS Data Store is not intended to replace your enterprise geodatabase.

It is a complement to your existing storage options and can be used in conjunction with them.

As part of your data strategy, consider what workflows your organization needs and where best to store and access your data.

Architecture

ArcGIS Enterprise architecture

ArcGIS Enterprise portal

ArcGIS Server

ArcGIS Data Store

ArcGIS Managed

Enterprise geodatabases, folders, cloud storage

User managed

Workflow: Publish by *reference* from ArcGIS Pro

ArcGIS Enterprise portal
Feature layer (item) created

ArcGIS Server
Feature service created

ArcGIS Data Store

Enterprise geodatabases, folders
Data remains here

Workflow: Publish as a *copy* from ArcGIS Pro

ArcGIS Enterprise portal
Feature layer (item) created

ArcGIS Server
Feature service created

ArcGIS Data Store
Copy of data stored here

Enterprise geodatabases, folders
Data copied from here

Workflow: Directly upload a CSV to your portal

ArcGIS Enterprise portal
Feature layer (item) created

ArcGIS Server
Feature service created

ArcGIS Data Store
Data stored here

Example strategy

A data strategy can take on many forms

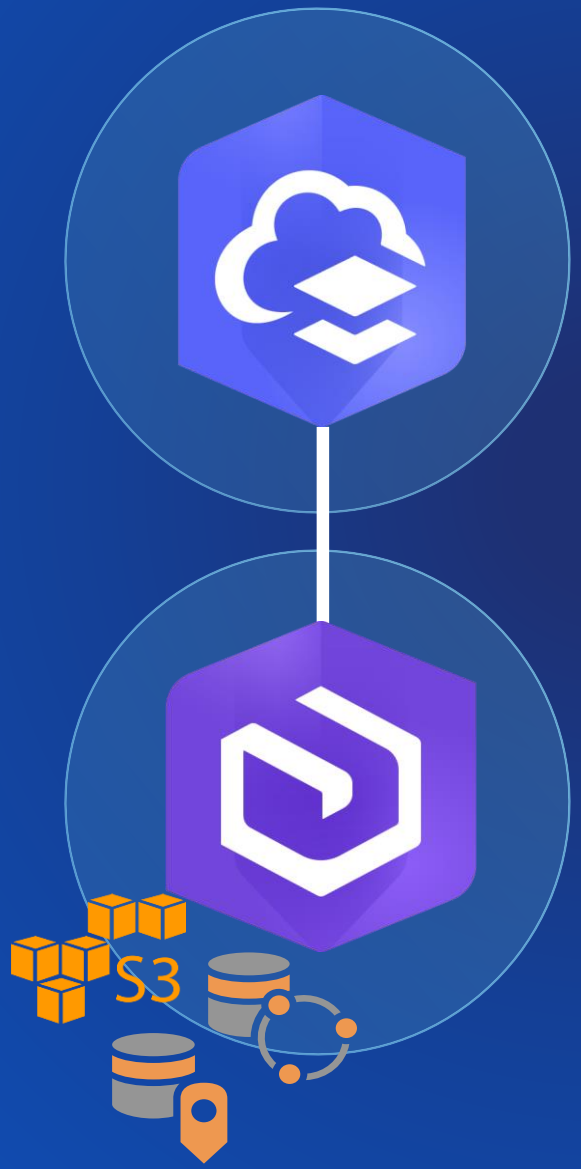

ArcGIS Online

- Public content
- Open data
- Non-employees (volunteers, contractors)
- Collaborated data from ArcGIS Enterprise for field operations

ArcGIS Enterprise

Enterprise geodatabase:

- Continuous, multi-user datasets

Hosted data:

- Innovation, projects, PoC, learning
- Replacement for file geodatabases

A data strategy can take on many forms

Business Objectives	Personnel	Metrics	Applications and Tools
Improve quality of data captured in the field	Field scientists, Geologists, GIS Professionals	• Reduce space for taking maps offline • Validate data at the time of collection • Ensure that collected data is punctually provided to QA-tier users	• ArcGIS Enterprise • ArcGIS Pro • Collector • Offline map areas • Domains

Where do I start?

Think about your end goal first

And then work backwards

Start here:
"I want to make a cake."

Then put the pieces in place to get there:
Butter, flour, a mixing bowl, an oven, ...

Working from your end goal backwards

What do you want to do?

I want to maintain an accurate inventory of parcels in my city.

The cake

I need to be able to have many editors working at once and to track changes.

The ingredients

My users have ArcGIS Pro licenses and we're using ArcGIS Enterprise.

The supplies

I'll use an enterprise geodatabase and use branch versioning off of web services.

The recipe

Example questions to jumpstart your strategy:

Collected

How will we capture your data?

Edited

Who will need to make changes?

Kept accurate

What type of quality assurance is needed?

Accessed

Who needs to be able to find and use it?

Scaled

Will our data grow? Will our userbase grow?

Used

What is the function of our data? Where will it be used?

Related topics

Related topics

- **Distributed collaboration**
  - *Sharing data between ArcGIS Enterprise environments and with ArcGIS Online*
- **ArcGIS Enterprise sites**
  - *Tailored landing pages for your users to discover and interact with your GIS*
- **Bulk publishing**
  - *A new option for publishing all of your enterprise geodatabase data as web services*

Related sessions (catch the recordings and slides!)

- ArcGIS Enterprise: Publishing Content and Services
- ArcGIS Enterprise: Best Practices for Layers and Service Types
- Spatial Data in ArcGIS: The Big Picture

Thank you!

Questions? Comments?