



Esri Best Practices: Implementing an Enterprise Geodatabase

Rasu Muthurakku, Esri Professional Services

Andrew Sakowicz, Esri Professional Services

SEE
WHAT
OTHERS
CAN'T



Agenda

Implementing an Enterprise Geodatabase

- Overview
- Implementation Approach
- Architecture
- Geodatabase Design
- Build
- Workflow Design
- Testing and Tuning
- Maintenance
- Monitor
- Question & Answers

Key Considerations, Best Practices, Recommendations and Lessons Learned!

Audience

Implementing an Enterprise Geodatabase

- **Intermediate**
- **Advanced**



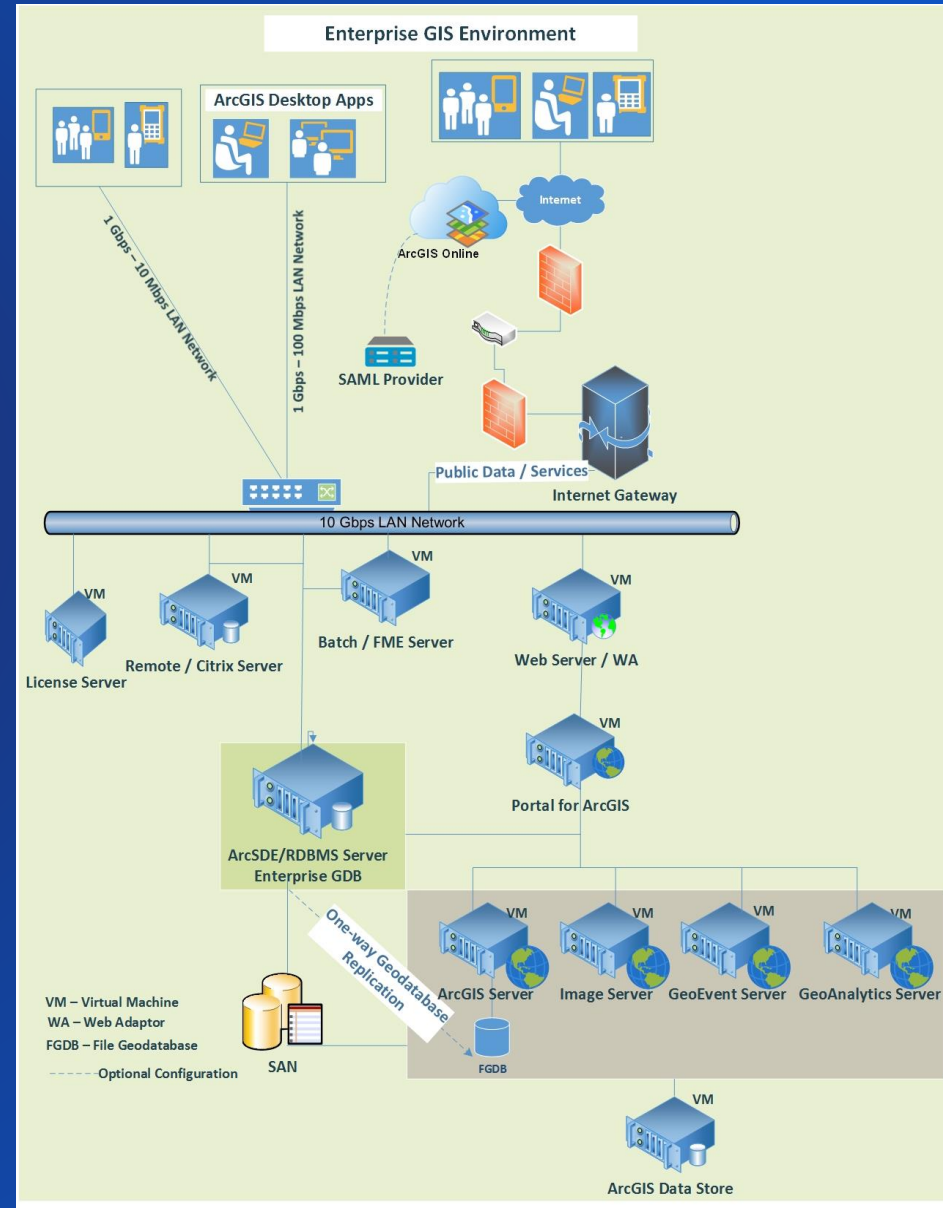
Overview

What is an Enterprise Geodatabase (EGDB)?

Centralized multiuser Geodatabase

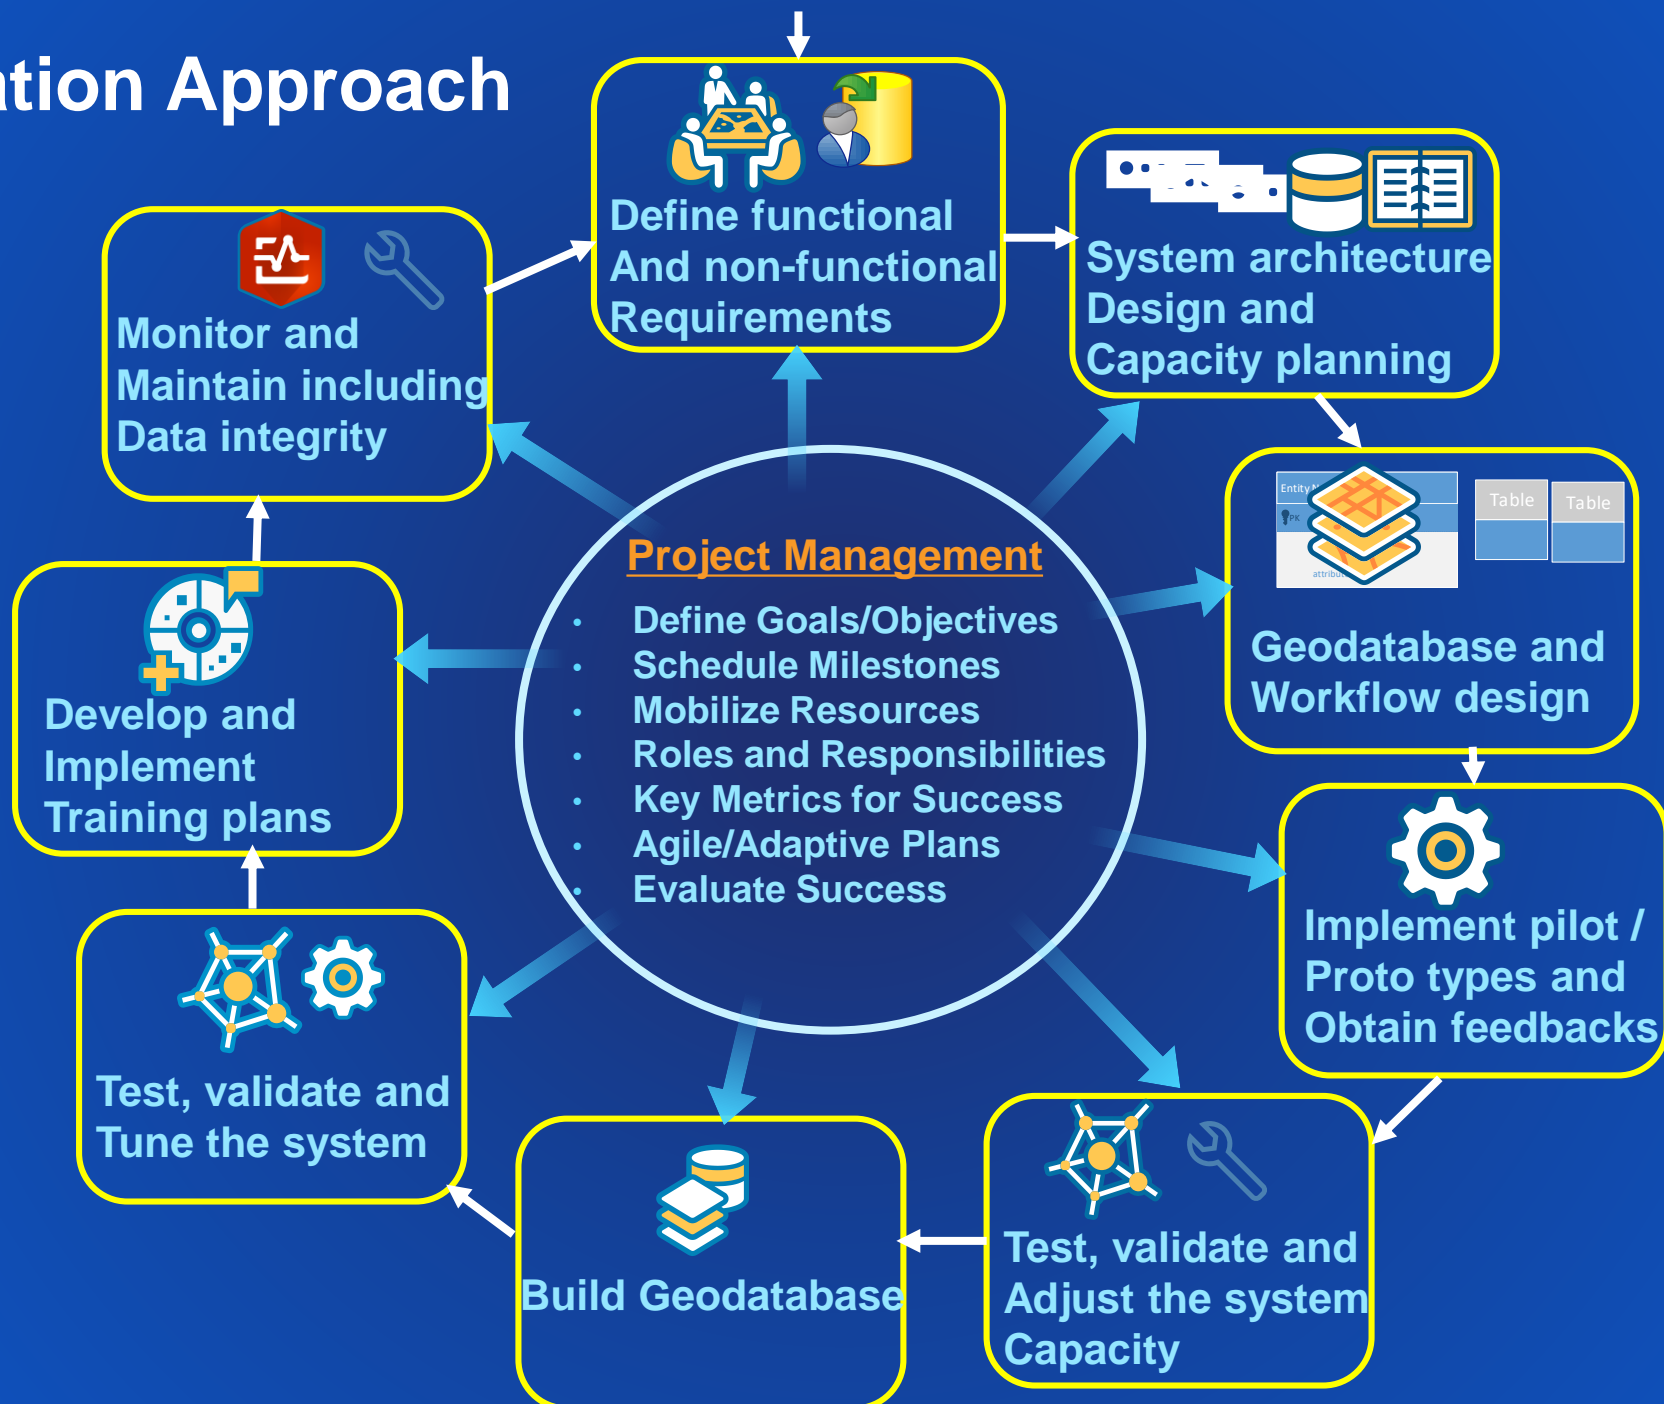
- ArcSDE enables the RDBMS* for GIS data management
 - Oracle
 - Microsoft SQL Server
 - PostgreSQL
 - IBM DB2, Informix
 - SAP Hana
- Extremely large, continuous and centralized GIS database
- Many simultaneous users
- Long transactions and versioned workflows
- SQL types for spatial in all supported RDBMS
- High performance for a very large number of users

*RDBMS – Relational Database Management System



Implementation Approach

Implementation Approach



Architecture

Efficient Implementation

- **System Architecture and Capacity**
 - Separation of read and edit data sources
 - Faster Processors
 - Enough system capacity
 - Better storage and network
 - etc.

- **Software**

- Version selection including the patches
- Utility industry ArcGIS Desktop version 10.2.1 (~~10.2.2~~) and 10.6.1
- Utility Network in ArcGIS Pro >2.1



ArcGIS (Desktop, Engine) Text Performance Patch

Drawing of text when font smoothing is enabled is significantly slower than non-smoothed text. This patch resolves the performance deficit.

Published: 04/30/2019

[Read More](#)

Annotation Performance						
S.No	Display Scale	Layer Name	Before Applying Patch in Seconds	After Applying Patch in Seconds	Performance difference in Seconds	Performance improvement in %
1	2000	ServiceAnno	1.02	0.19	0.83	436.84%

Under-utilized / stable IT Infrastructure helps

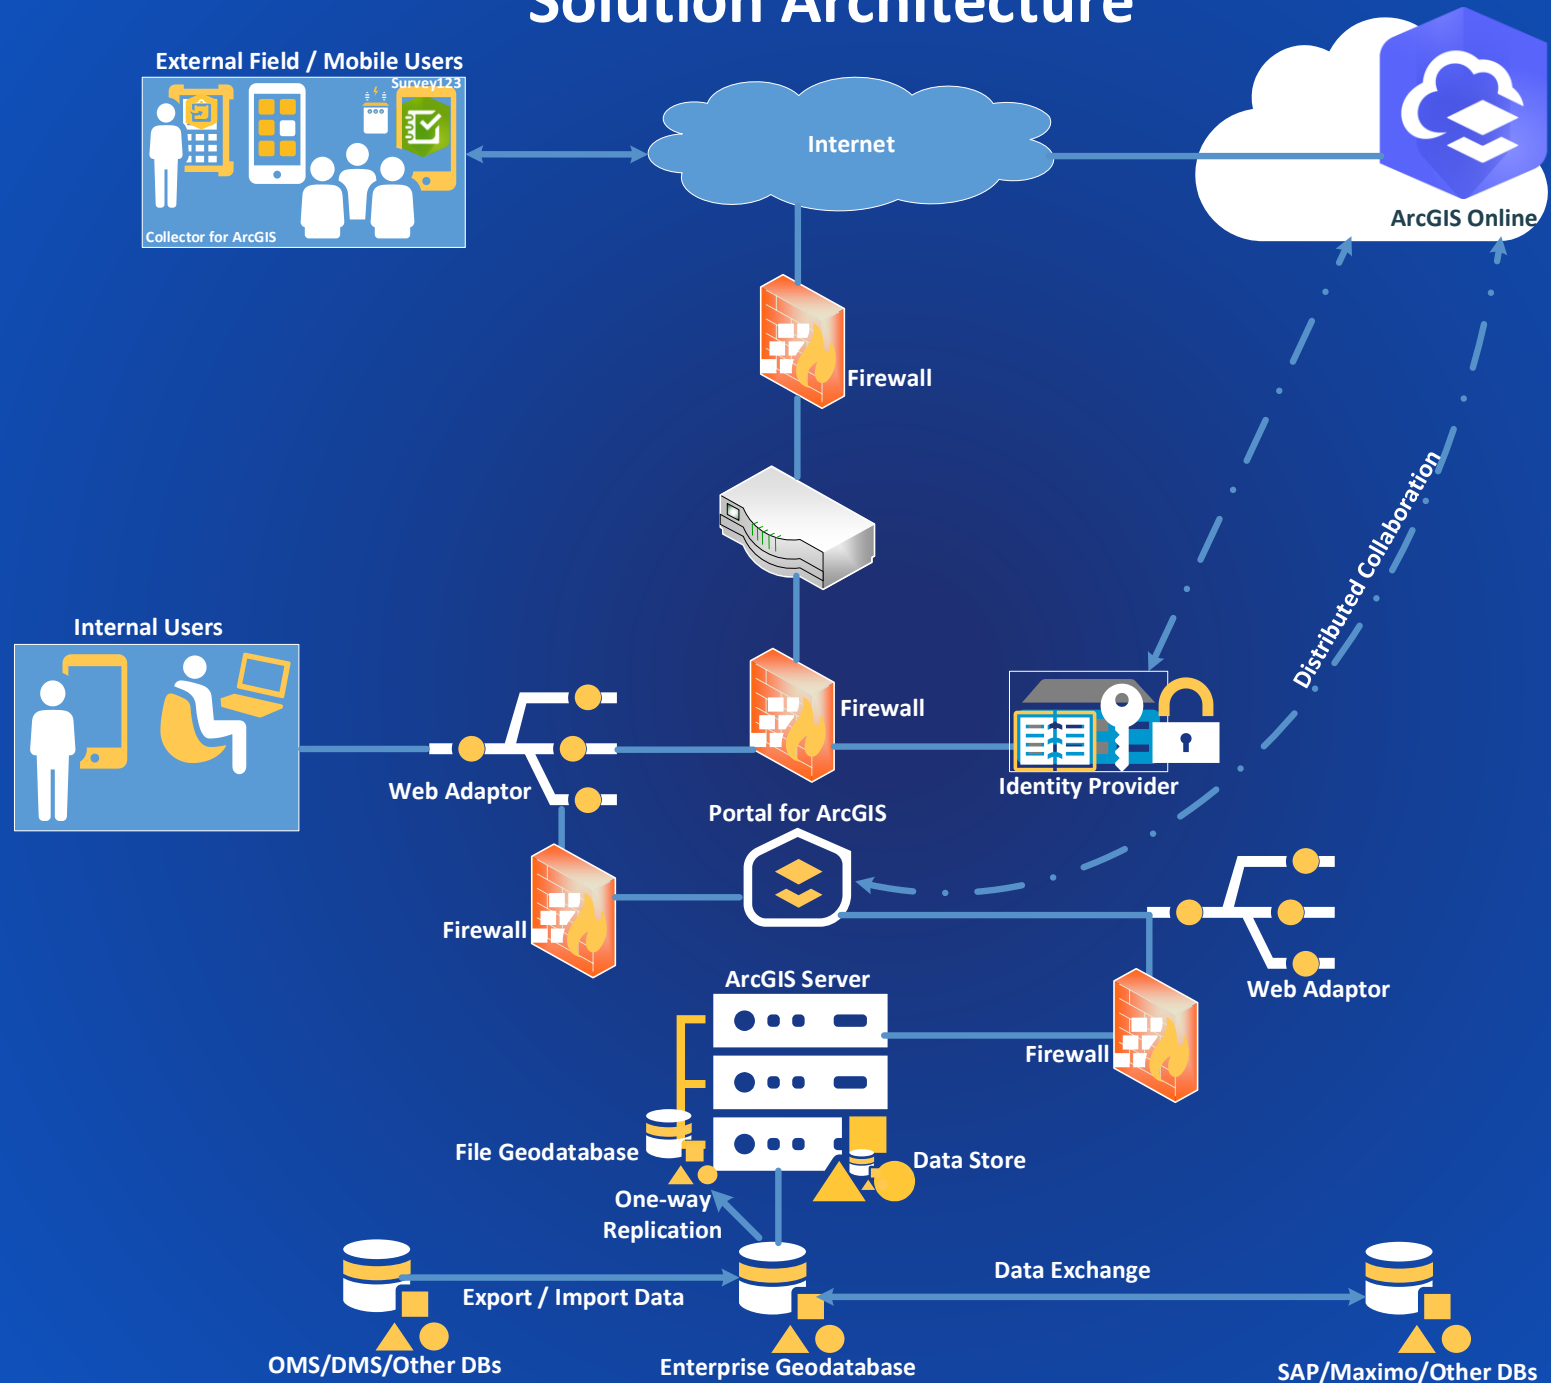
Capacity Planning

- **Define architecture vision / foundation**
 - Describe the System and its relationships
- **Business Architecture**
 - Define the business usage
- **Application Architecture**
 - Plan suitable software solutions / applications
- **Data Architecture**
 - Identify data requirements and management
- **Technology Architecture**
 - Select proper technology & capacity for IT infrastructure



System architecture design reduces cost and improves productivity

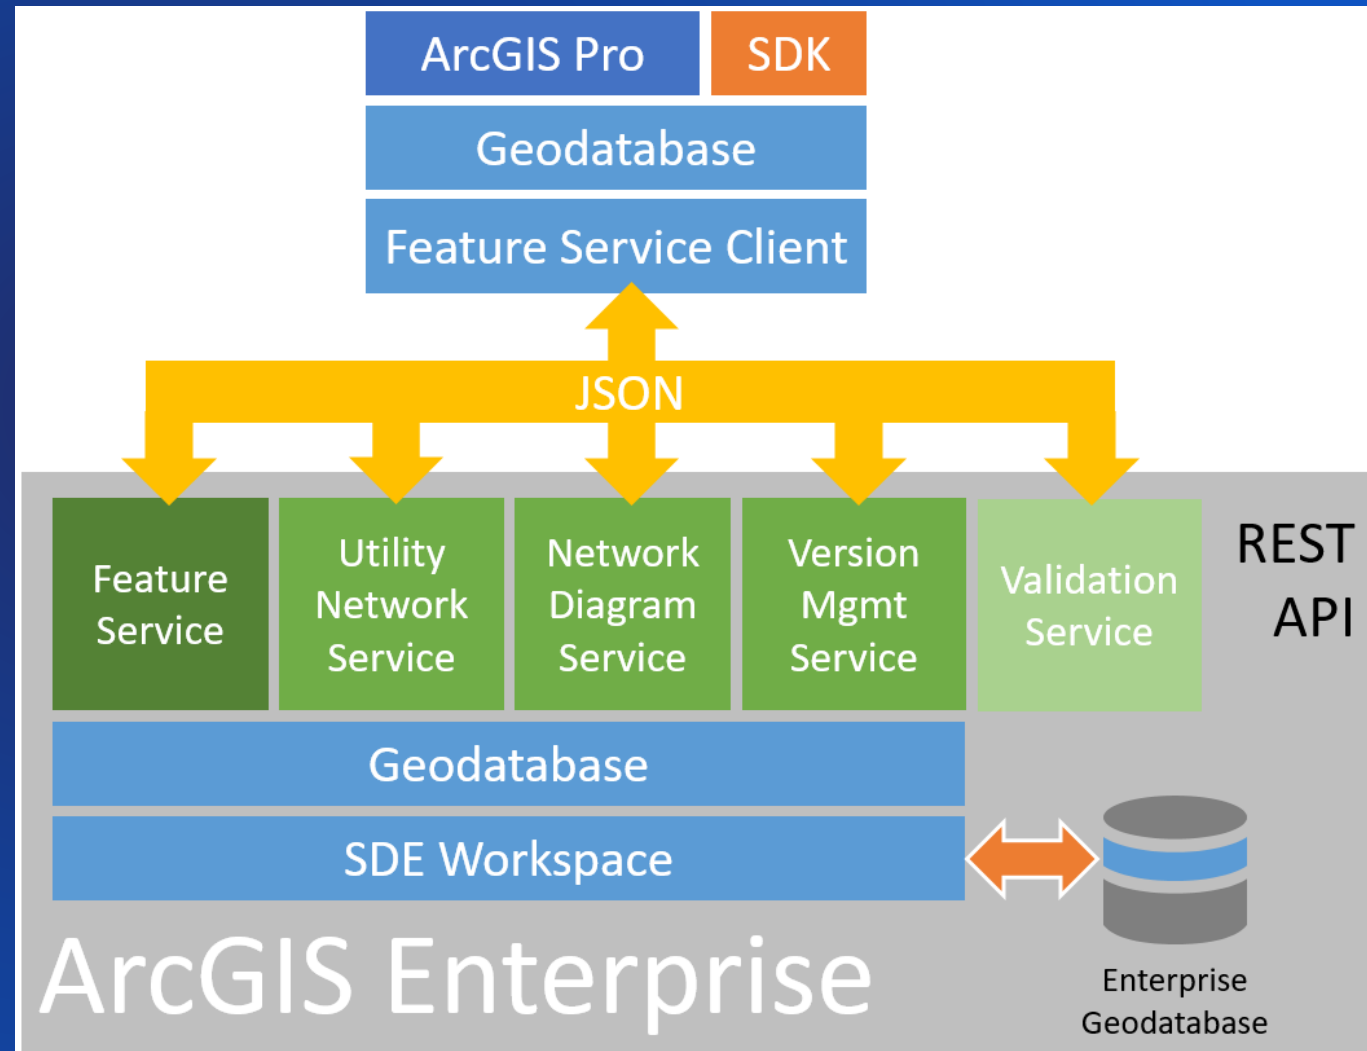
Solution Architecture



New / Additional Solution Architecture Options

Utility Network

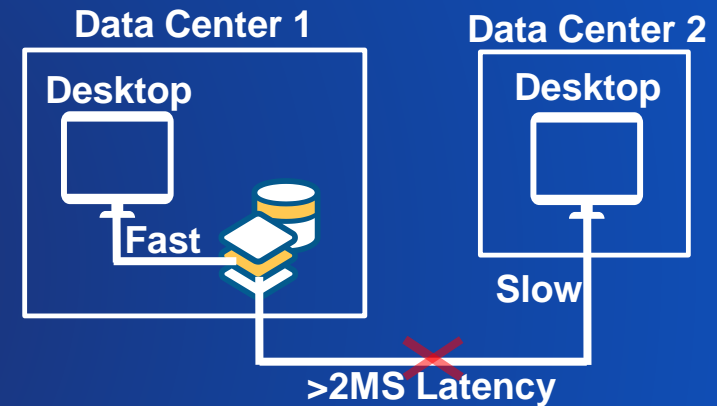
- **New network to manage Utility and Telecom network data**
- **Cross platform support**
 - Any device, anytime, anywhere!
- **Services based architecture**
- **Updated network model**
 - **Connectivity associations**
 - **Containment associations**
 - **Structural attachments**
 - **Multiple terminals**
 - **Expanded tracing framework**
 - **Built in support for network diagrams**



Business Architecture

- Business users
 - Total and concurrent
 - Location
 - Internal and / or external
- Business workflows and user roles
 - Data editors / Managers / Administrators
 - Data readers
 - Power users
 - Web / Mobile users
 - etc.

Data Centers and User Sites





Electric MXD - DC2 Vs DC1 - MXDPerfstat Results					
Scale	Feature Class/Layer	Display in Seconds at Data Center 2 Server	Display in Seconds at Data Center 1 Server	No of Features	Performance Difference in %
25,000	Street Centerline	0.22	0.16	3,449	37.50%
20,000	ROWEdge	0.61	0.23	4,017	165.22%
20,000	Overhead Primary Conductor	0.34	0.28	3,403	21.43%
15,000	Street Centerline	0.23	0.09	2,745	155.56%
15,000	ROWEdge	0.25	0.14	2,461	78.57%
15,000	Overhead Primary Conductor	0.57	0.19	1,944	200.00%

Define Functional and Non-Functional Requirements

Critical to collect NFRs

- **High Availability (HA)**
 - Web application for 24x7 Field / Emergency crew
 - Use clusters with one database for editing
- **Backups**
- **Disaster Recovery (DR)**
- **SLAs (Service Level Agreements)**
 - RPO (Recovery Point Objective)
 - RTO (Recovery Time Objective)



Relevant and justifiable requirements to keep the cost and complexity checked

Application Architecture

- Desktop and web solutions
- Target software versions
 - Operating System
 - Virtualization
 - Partner solutions
- Update service pack levels and patches
 - ArcGIS Desktop TLS Patch
 - ArcGIS (Desktop, Engine) Text Performance Patch
 - Utility and Telecom Update (UTU) Patch 9 for 10.2.1
 - etc.

MXDPERFSTAT – Production Environment

S.No	Display Scale	Layer Name	Before tuning - Display in Seconds	After Tuning (Only SDE Direct Connect and DFQ*) - Display in Seconds	Performance Improvement in %
1	1,000,000	TRANS LINE	0.94	0.37	154.05%
2	500,000	STATION	0.26	0.13	100.00%
3	50,000	TRANS LINE	0.8	0.18	344.44%
4	50,000	STRUCTURE	44.62	0.72	6097.22%
5	25,000	STRUCTURE	46.38	0.49	9365.31%
6	25,000	TRANS LINE	0.83	0.16	418.75%
7	1,000	STRUCTURE	42.26	0.28	14992.86%
8	1,000	TRANS LINE	0.76	0.31	145.16%
9	500	STRUCTURE	41.66	0.27	15329.63%
10	500	TRANS LINE	0.74	0.21	252.38%

*DFQ (Definition Query) = Objectid is not null


Impact of an older RDBMS Version bug

Check system requirements and GDB client compatibility

Technology Architecture

Quality vs Quantity

- IT infrastructure
 - Availability and policies
 - Limitations, preferences and constraints
- Processor selection
 - Key to optimal scalability and performance
 - Save costs by reducing server footprint
 - Check spec rate per core
 - <http://spec.org/cgi-bin/osgresults?conf=cpu2017>
 - E.g. Intel Xeon Gold 6244 16 cores 3.6GHz = 83.125 SPEC rate per core (April 2019 results)



Select higher SPEC rate per core for speed Vs processor density

Server Technology Selection

Do not compensate poor maintenance with top processing power

- Processing power is directly proportional to total Delta table records count

**Intel® Xeon® E5-2637 v4 @
3.40GHz, 8 Cores, 256GB
RAM, ~59.38 SPEC rate/core
Physical machine hosting a
happy GDB!**

```
SQL>
SQL> prompt 2. STATES
2. STATES
SQL>
SQL> select count(*) from sde.states;

Tue Sep 26

COUNT (*)
-----
113,904

SQL>
SQL> prompt 3. STATE_LINEAGES
3. STATE_LINEAGES
SQL> select count(*) from sde.state_lineages;

Tue Sep 26

COUNT (*)
-----
106,274,130

SQL>
SQL> spool off
```

```
---ArcSDE A Table Rows

SQL> select sum(num_rows) as

TOTAL_A_TABLE_ROWS
-----
12,154,746

SQL>

SQL> ---ArcSDE D Table Rows
SQL> select sum(num_rows) as

TOTAL_D_TABLE_ROWS
-----
9,338,738

SQL>
SQL> Spool off
```

Virtualized Database Servers – Key Considerations

Speed is more important than utilization

- Avoid over-commitment
- Ensure less number of vMotions
- Estimated capacity requirements
 - User Load
 - Dedicated operations and transactions
- Application & database complexity



Physical vs. Virtual Server

General	
vSphere DRS:	On
vSphere HA:	On
VMware EVC Mode:	Disabled
Total CPU Resources:	2159 GHz
Total Memory:	12.50 TB
Total Storage:	494.37 TB
Number of Hosts:	25
Total Processors:	800
Number of Datastore Clusters:	1
Total Datastores:	81
Virtual Machines and Templates:	793
Total Migrations using vMotion:	392768

General	
vSphere DRS:	On
vSphere HA:	On
VMware EVC Mode:	Intel® "Sandy Bridge" Generation
Total CPU Resources:	227 GHz
Total Memory:	1.87 TB
Total Storage:	103.92 TB
Number of Hosts:	4
Total Processors:	88
Number of Datastore Clusters:	0
Total Datastores:	42
Virtual Machines and Templates:	314
Total Migrations using vMotion:	3928

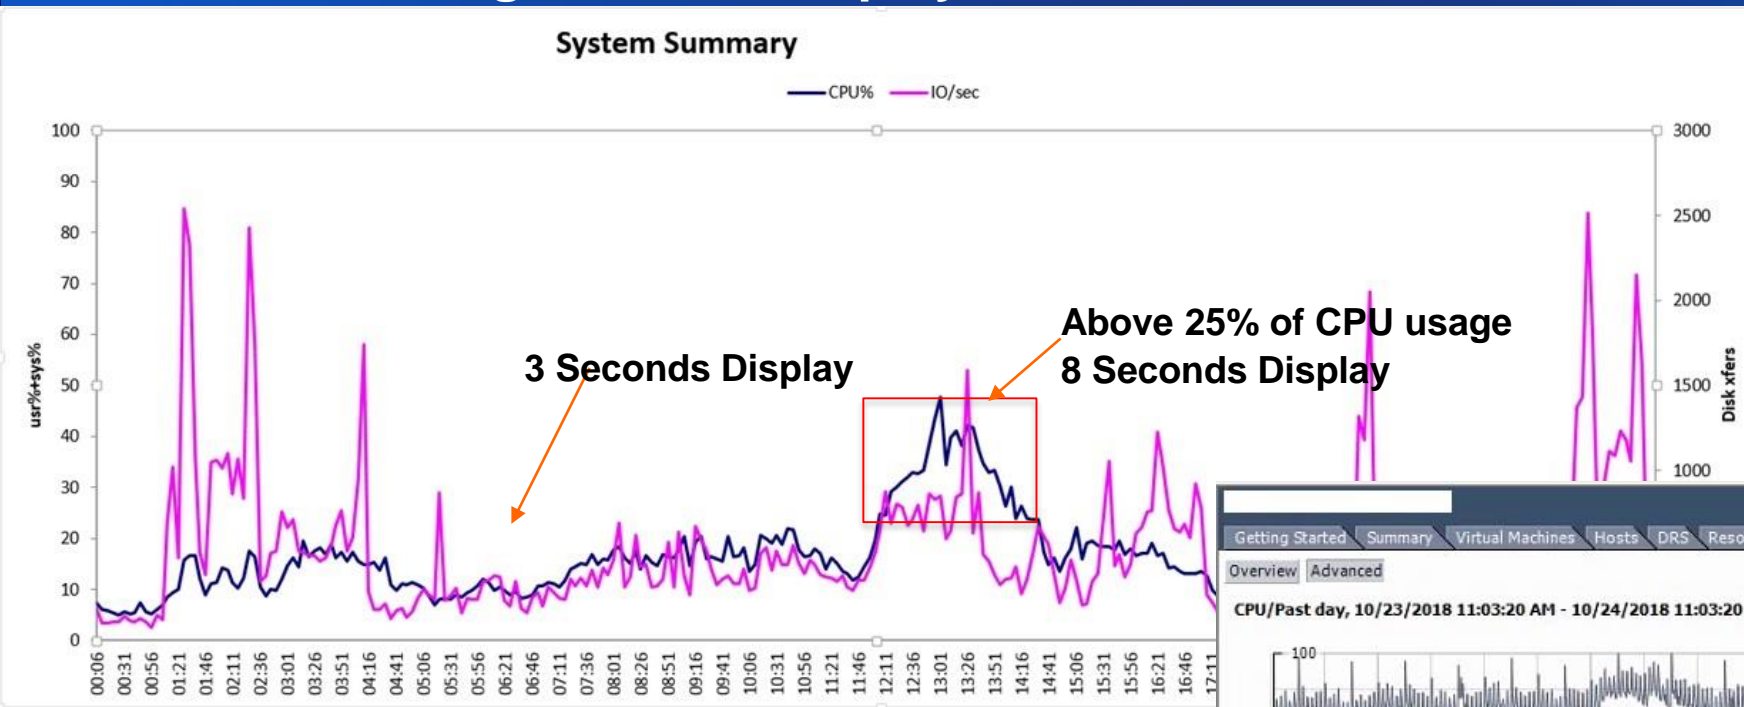
0.3 CPU/VM

General	
vSphere DRS:	On
vSphere HA:	On
VMware EVC Mode:	Intel® "Westmere" Generation
Total CPU Resources:	184 GHz
Total Memory:	2.16 TB
Total Storage:	61.06 TB
Number of Hosts:	5
Total Processors:	92
Number of Datastore Clusters:	0
Total Datastores:	37
Virtual Machines and Templates:	176
Total Migrations using vMotion:	2747

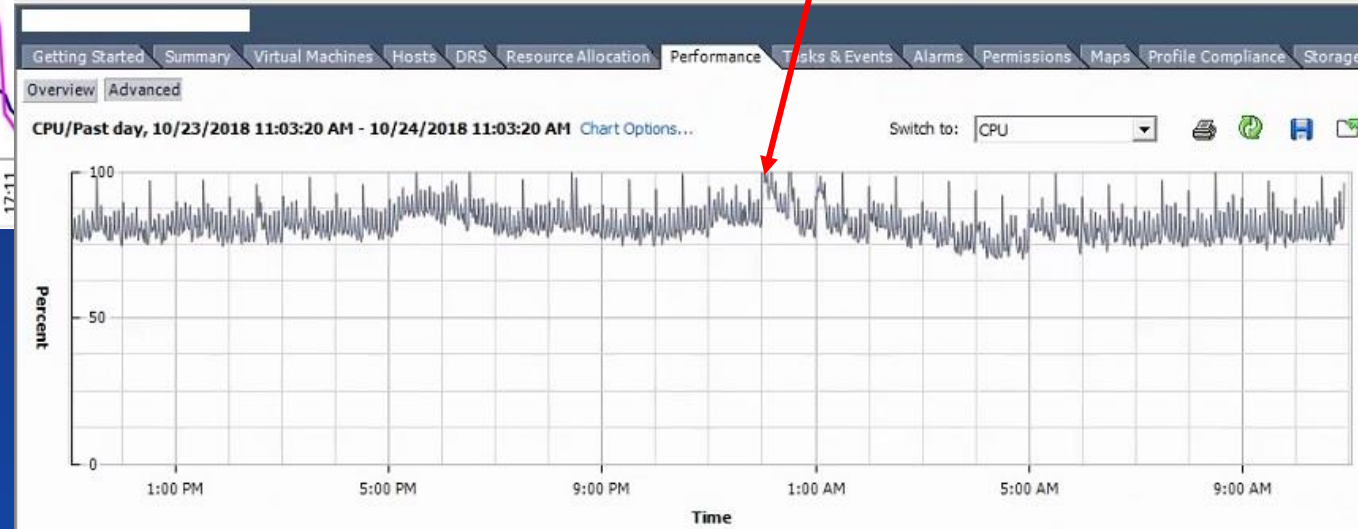
Complex and busy databases need dedicated machines

Symptoms of Over-Committed vCPUs

- vCPU usage Vs MXD display seconds



~100% Virtual cluster utilization!



User load

Estimated vs actual usage

- Number of users, operations and transactions
 - Expected 130 power users (editors) Vs 200+ Actual Users
- Number of outstanding Versions
 - Estimated 600 Vs actual 1250 Versions
- Estimated total delta table records (A# and D# Tables)
 - Designed for 2 Millions Vs 8 Millions records actual



```
HASH_AJ */ SDE_DELETES_ROW_ID, SDE_STATE_ID FROM NIS.D63 WHERE DELETED_AT  
IN (SELECT 1.lineage_id FROM SDE.state_lineages 1 WHERE 1.lineage_name =  
:lineage_name2 AND 1.lineage_id <= :state_id2) AND SDE_STATE_ID >  
:"SYS_B_5") AND a.SDE_STATE_ID = SL.lineage_id AND SL.lineage_name =  
:lineage_name3 AND SL.lineage_id <= :state_id3) V_63
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	169	0.19	0.19	0	0	0	0
Fetch	169	24.66	24.66	0	123445	0	0
total	339	24.86	24.86	0	123445	0	0

Unplanned concurrency overloads Server resources

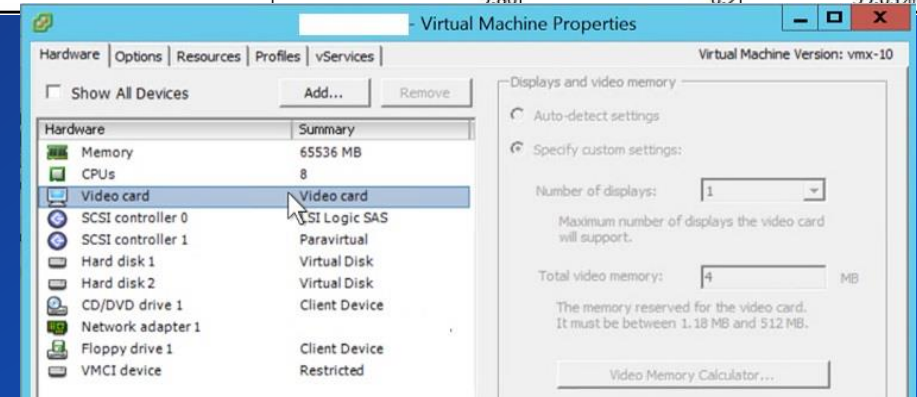
Virtual Environment



- Dedicate the Virtual environment for a large GIS user base
 - Provide a decent provisioning ratio
- Provide GPU* for GDB clients
 - Video RAM >256MB Per Virtual Machine
- Fit Virtual Machine within one CPU NUMA* node
 - # of vCPUs <= number of cores in the CPU socket
- Choose faster CPUs - Spec rate per Core 80+

* GPU – Graphics Processing Unit)
 *NUMA - Non-uniform memory access

Task	Performance - Before CPU Change (in Seconds)	After Change in CPU Entitlement from 1 CPU to 4 CPU (in Seconds)	Performance improvement in %
Opening ArcMap thru Citrix (ArcGIS Desktop Launcher - for setting the licenses)	26.01	21.37	21.71%
ArcFM Locator - Zoom to Distribution	29.45	11.15	164.13%
Zoom to 1000 ft	8.92	7.33	21.69%
Zoom to 500ft	4.36	3.69	18.16%
Zoom out to 100ft	5.02	4.3	16.74%
Zoom out to 50ft	4.19	3.73	12.33%
Place UG Primary Conductor with Ducts (2)	13.3	6.68	99.10%
Distribution Pole	5.46	4.19	30.31%
Split OH Primary Conductor - Custom Split	21.23	15.76	34.71%
Delete Transformer	14.05	7.8	80.13%
Place the above transformer and Jumper	9.85	5.24	87.98%
Place a pole	2.48	0.82	202.44%
Place Primary OH-ABC Phase	5.1	2.89	76.47%
No Save Edits	9.86	6.2	59.03%



Tune Virtual environment before deployment

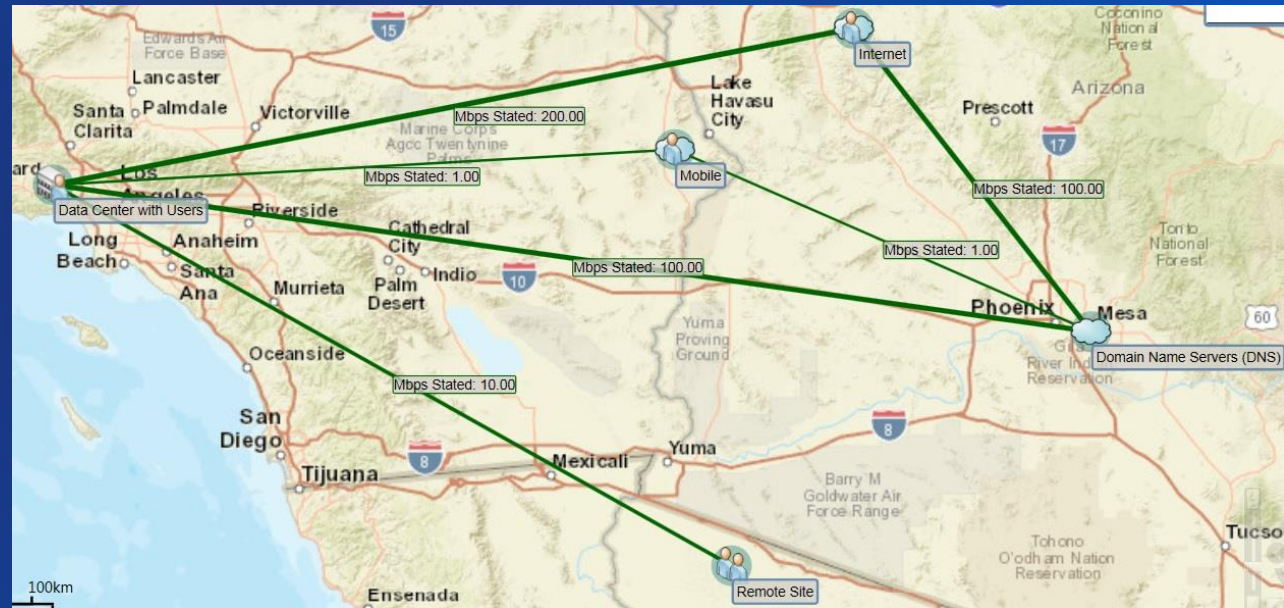


Network Planning

Establish and configure DNS appropriately!

Trace Route: LA Workstation → Phoenix DNS
 LA Database Server ←

```
C:\Users\ >tracert
Tracing route to
over a maximum of 30 hops:
  1    55 ms    55 ms    55 ms
  2    55 ms    55 ms    55 ms
  3   115 ms    58 ms    62 ms
  4   111 ms   111 ms   112 ms
  5   110 ms   109 ms   110 ms
  6   110 ms   113 ms   110 ms
  7   109 ms   109 ms   109 ms
```



Network Latency's impact on performance

S.No	Task	Performance when <1 ms latency (in Seconds)	Performance during latency fluctuation (in Seconds)	Performance difference in Seconds	Performance Difference in %
1	Open ArcMap	23	90	-67	-291.30%

Network Infrastructure

- Establish higher Network bandwidth (>1Gbps) and less latency (<1 MS)
 - ArcGIS Desktop is sensitive to 1 – 2 MS latency
 - Plan for ~ 1.5 Mbps per concurrent GIS User
- Enable Jumbo Frames between servers
 - All Switches must support – Otherwise don't enable it!
- Validate Network path between GIS user and server locations
- Upgrade lower bandwidth or move GIS user locations

MXDPerfstat Results – 100Mbps Vs 1Gbps port

Scale	Feature Class Name	Response Time in Seconds when the workstation was connected to 100Mbps Phone port	Response Time in Seconds when the workstation was connected to 1Gbps LAN port	Performance improvement in %
50,000	Centerline	8.84	4.18	111.48%
50,000	Control Boundary (WMX)	0.17	0.14	21.43%
2,500	GasMeter	0.78	0.55	41.82%
1,000	GasMeterQualAnno	0.21	0.12	75.00%

Higher bandwidth between EGDB and the clients provide better performance

Storage

Disks

- **Use SSDs (Solid-State Drives)**
- **Plan for ~5000 IOPS (Input/output Operations Per Second)**
- **Avoid LUNs* > 2TB size**
 - Minimum of 4 LUNs that are identical in size
- **Avoid noac mount option**
- **Distributed File System (DFS) is not supported**

***LUN – Logical Unit Number**

SSD delivers better performance for EGDB

ArcSDE Configuration

- **Configure ArcSDE DBTUNE Settings**
 - Use Default Geometry Storage - St_Geometry / Geometry
 - Storage locations
 - etc.
- **ArcSDE initialization parameters**
 - Limit the connections
 - Defaults are good!



Geodatabase Design

Geodatabase Design – Data Modeling

- **Conceptual Design**


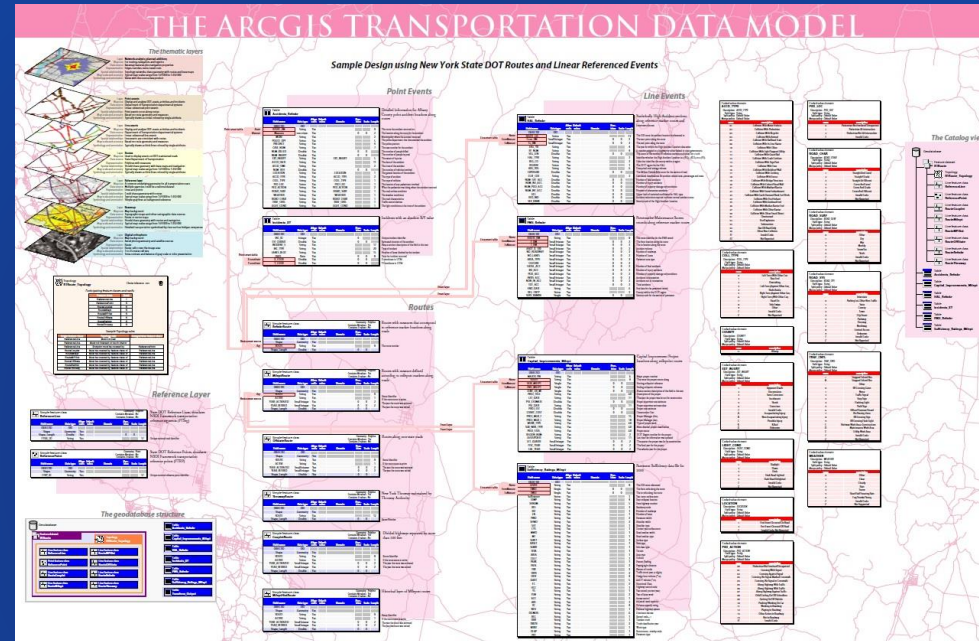
- Identify business requirements
- Identify thematic layers
- Identify required applications
- Document

- **Logical Design**

- Define tabular database structure
- Define relationships
- Determine spatial properties
- Document

- **Physical Design**

- Create and implement model design
- Generate physical schema in the RDBMS / FGDB
- Test and validate
- Document

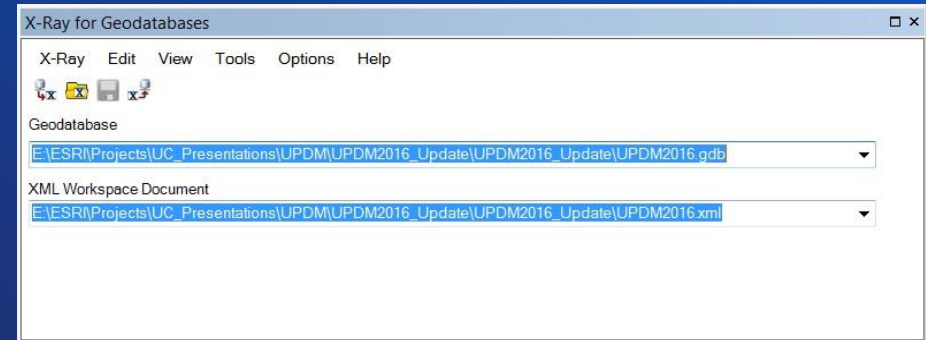


Leverage the existing data models

Geodatabase Design – Tools

- Tools


- X-Ray Add-In
- Geodatabase Diagrammer
- Sparx Systems' Enterprise Architect
- Geometric Network Configuration Manager



Geodatabase Design – Key Considerations

Prevention is better than cure!

- Consider single coordinate system
 - On the Fly Projection is expensive
 - Geometric Network editing does not support “On the Fly Projection”
- Column / Domain names and Field lengths
 - Avoid >10 characters in Field names
 - Apply required text length e.g. Text - 256 Vs NCLOB – 1,073,741,822
 - Select appropriate Field type
 - Apply only the required Precision and Scale
 - Define Not Null Fields



Field Properties	
Alias	CustomerName
Allow NULL values	Yes
Default Value	
Domain	
Length	1073741822

Justify every single Geodatabase element

Geodatabase Design – Key Considerations

Poor Design = Rework, slow performance and bugs



- Avoid XY Tolerance modification
 - Default = 10x Times of XY Resolution
 - Introduces complexity (#NIM090335) for Geometric Network, etc.
 - Impacts performance

A screenshot of the 'New Feature Dataset' dialog box in a GIS application. The dialog has a title bar with a close button. It contains three sections for tolerance settings: 'XY Tolerance' with a text box containing '0.001' and the unit 'Meter'; 'Z Tolerance' with a text box containing '0.001'; and 'M Tolerance' with a text box containing '0.001' and the unit 'Unknown Units'. Below these sections are two buttons: 'Reset To Default' and a blue hyperlink 'About spatial reference properties'. At the bottom, there is a checked checkbox labeled 'Accept default resolution and domain extent (recommended)', which is circled in red. At the very bottom of the dialog are three buttons: '< Back', 'Finish', and 'Cancel'.

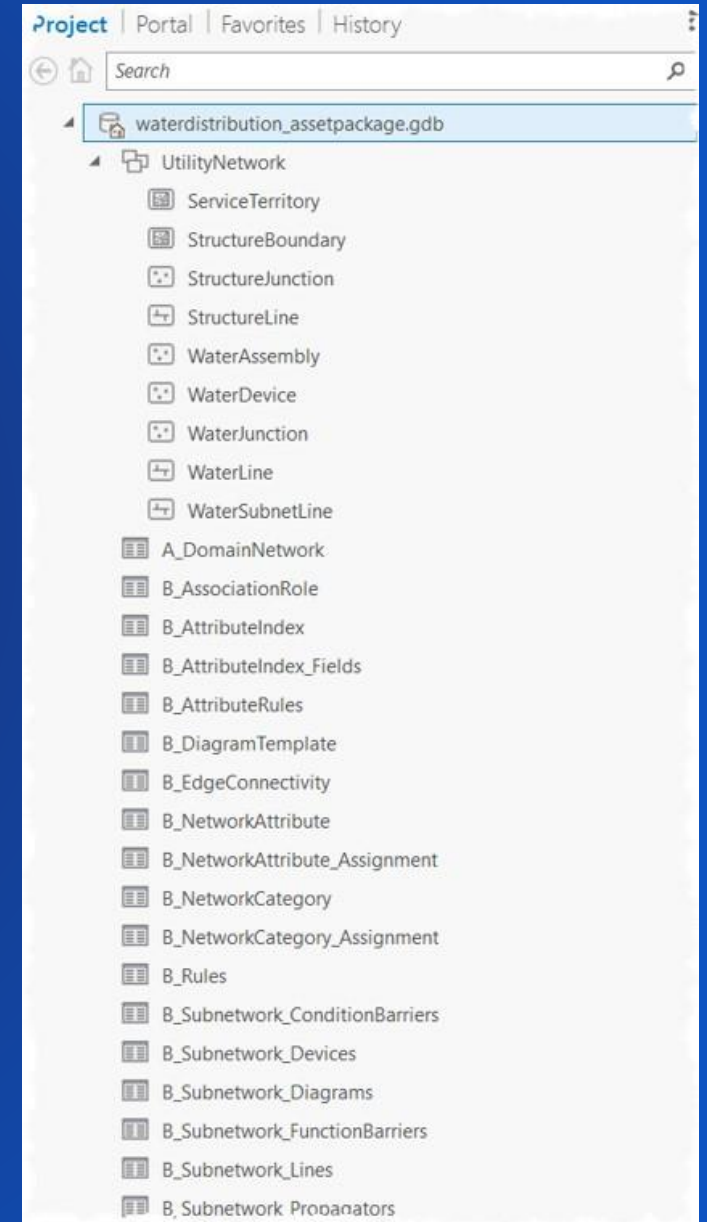
Avoid expensive rework!

Geodatabase Design – Best Practices

Data Model impacts storage and performance


- **Leverage the existing ArcGIS Data Models**
 - Drop redundant Feature Datasets / Classes, Columns, etc.
 - Stand alone Feature Classes are fine!
 - Possibly split the Feature Classes per scale levels
 - Reduce complex and attributed Relationship Classes
 - Test, refine and tune the Data Models
- **Integrate related Feature Classes using topology**
- **Deploy necessary information models**
 - Geometric Network Vs Utility Network

Justify every single Geodatabase element



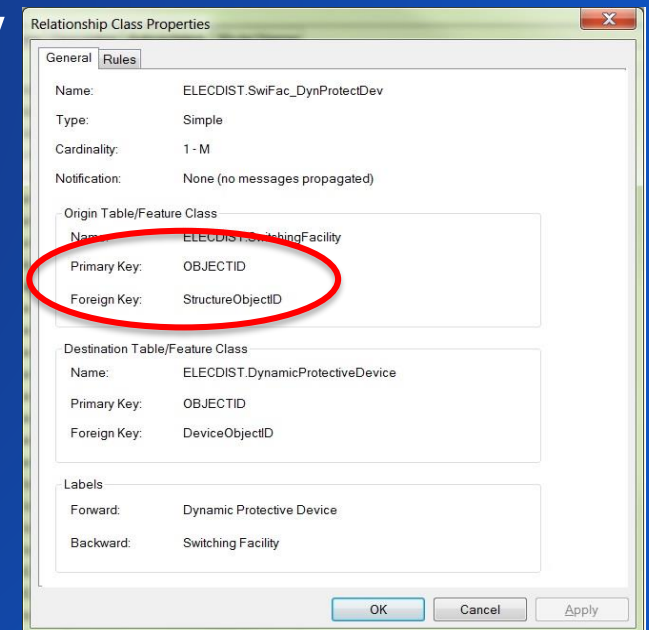
Geodatabase Design – Best Practices

Poor Design = Rework, slow performance and bugs



- **Create Feature Datasets or Databases for each LOBs (Line of Businesses)**
 - Depends on size, access, usage and maintenance
 - E.g. Landbase, Gas, Electric, Water GDBs, etc.
- **No attributed Relationship Classes with no attributes**
- **Use Many to Many Relationship Classes only when necessary**
- **Don't use Objectid as Primary Key for Relationship Classes**
 - Unexpected Replication behavior
 - Additional processing during synchronization

Keep common dataset / database for base map / landbase



Geodatabase Design – Best Practices

Navigate common oversights!

- Review the Labeling requirements ahead of time
 - For Multi-Field complex Labeling
 - Combine them to a new Field and Auto Update
 - Convert Labels to Annotations
- Less Annotation Classes within Annotation Feature Class
- Add Attribute Indexes
 - Label Expression
 - Definition Queries
 - Application Design



Analyze Annotation requirements and choose proper Annotation reference scale

Build

Build Geodatabase

- **Create physical Geodatabase**
 - Separate Data Owner from SDE / DBA user
 - Structure the implementation to Pilot → Phase I → Phase II → Phase III, etc.
 - Enough gap to accommodate the learned lessons
- **Develop Data conversion specification document**
 - Align it with Data Model
- **Team review and demonstration**

Capture, load and maintain data accurately

Geodatabase Access and Management

- Create roles / groups based on the access level
 - Total access will slow connection time!
- Configure client applications to manage data



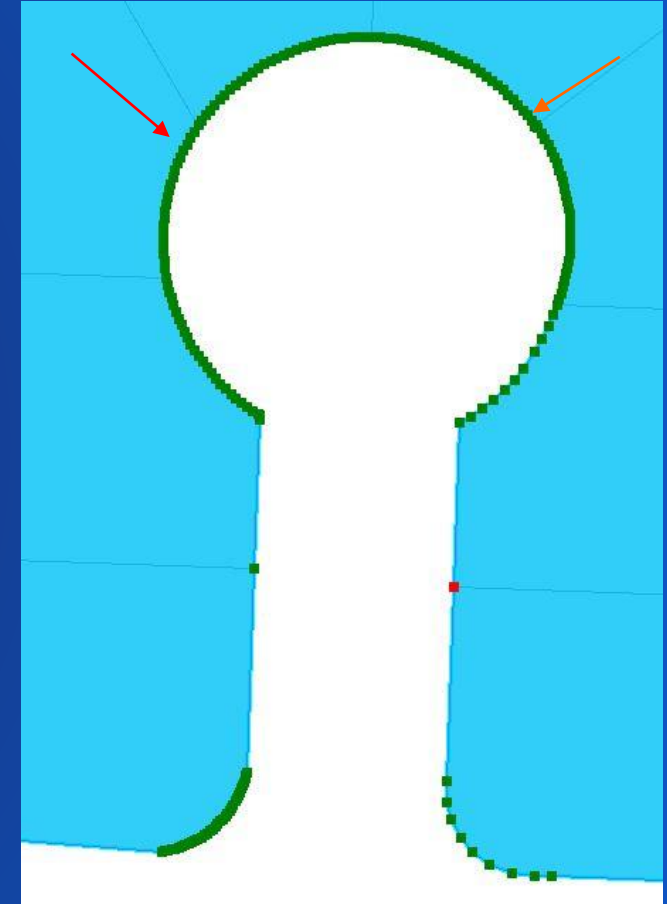
Connection Performance					
S.No	User Name	ArcCatalog Connection in Seconds with all access	ArcCatalog Connection in Seconds with reduced privileges	Performance difference in Seconds	Performance improvement in %
1	Rasu	30	12	18	150.00%
2	Andrew	21	12	9	75.00%

Capture, load and maintain data accurately



Key Data Conversion Considerations

- Extra Vertices introduced by:
 - Conversion process involving CAD systems
 - Geometric Network creation with Snapping ON option
 - Conversion/Update methods
- Develop adequate QA/QC methods and procedures
- Additional Data Reviewer checks
 - Duplicate/invalid geometries
 - Orphan related records
 - Connectivity check
 - etc.



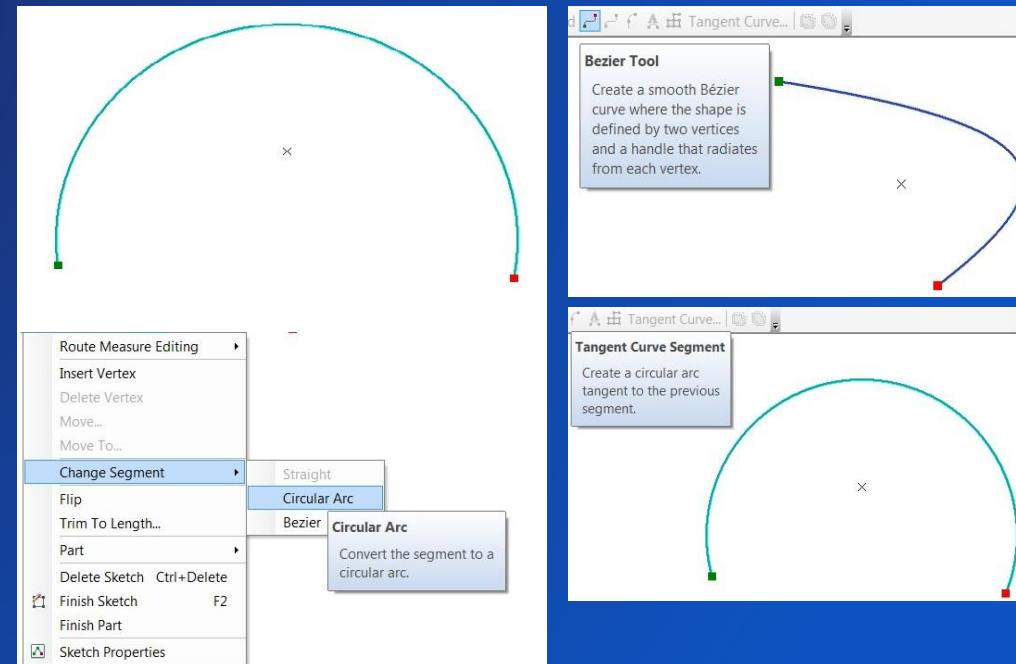
Feature Class/Layer	Total Number of Features	Total Number of Vertices with Geometric Network Snapping	Total Number of Vertices without Geometric Network Snapping	Increase in number of vertices %
Secondary Overhead Conductor	3,712	12,953	7,841	65.20%

Put additional efforts to get the data conversion right

Build Geodatabase – Best Practices

- **Avoid extra vertices**
 - Use curve tools that insert less vertices
- **Turn off snapping during Geometric Network (GN) creation**
 - ArcGIS 10.0 onwards a Vertex is added at every intersection
 - Cannot create GN with >15M edges until ArcGIS 9.3.1
- **Remove additional vertices**
 - Generalize / Simplify
 - Use ArcObjects

Provide tips and tricks for editing crew regularly



Build Geodatabase – Best Practices

Data is the brain of GIS nervous system!

- Aim for 100% Data accuracy
- Either populate or drop empty Fields
- Minimize Data Model / schema changes
- Consolidate GDBs
 - Avoid creating GDB per Geographic locations / regions
- Extra Vertices have performance impact

Scale	Feature Class	Display in Seconds - before Simplify Polygon	Display in Seconds - after Simplify Polygon (without any shape change)	Performance gain in %
50,000,000	Time_Zones	0.59	0.27	118.52%
1,000,000	Time_Zones	0.06	0.02	200.00%
10,000	Time_Zones	0.05	0.02	150.00%
1,000	Time_Zones	0.05	0.02	150.00%

S.No	Display Scale	Layer Name	# of Displayed Features	Number of Vertices - Before Simplify	Number of Vertices - After Simplify	Before Simplification - Display in Seconds	After Simplification - Display in Seconds	Performance Improvement in %
1	50,000	Street_1_inch	35,093	105,695	101,060	2.36	0.5	372.00%
2	50,000	PARCEL_1_Inch	7,922	645,766	188,212	0.37	0.31	19.35%
3	25,000	Street_1_inch	11,192	31,112	29,620	0.69	0.2	245.00%
4	25,000	PARCEL_1_Inch	2,687	168,011	48,540	0.16	0.14	14.29%
5	20,000	Street_1_inch	7,590	20,494	19,574	1.59	0.16	893.75%

Unnecessary vertices add significant performance overhead

Build Geodatabase – Best Practices

- Keep the Data clean and simple
 - Without any topological errors
- **No Coincident Complex Edge Features in Geometric Network**
 - **Most common reason for Geometric Network corruption**
- Unversion the Read Only Feature Classes / Tables
- Use Mosaic Datasets for Raster requirements
 - No massive imagery loading into EGDB

OWNER	TABLE_NAME	NUM_ROWS	BLOCKS	AVG_ROW_LEN	TO_CHAR (LAST_ANALYZED, 'MON/DD/YYHH24:MI:SS')
ARCFM	SDE_BLK_3	1145079	1190598	12	APR/19/16 05:38:10
ARCFM	BK SDE LOGFILE DATA	4231223	8773	10	APR/19/16 05:32:14

Static Raster Data does not need to participate in daily RDBMS backup

Workflow Design

Implementation Dependencies

Understand intrigue challenges

- **Map the project / data dependencies**
 - Business user groups
 - Application designs
 - System integration requirements
 - etc.
- **Organize cross functional skills/support team**
 - Dedicated / Assigned ArcSDE Administrator

Talk to the right people to get the right information



Requirements and Workflows

Selection of GDBs drive efficiency

- Number of users and types of users
- Workflows
 - Multi User editing - **Enterprise / Workgroup GDBs**
 - Single User Editing – **FGDB**
 - Replication - **EGDB → FGDB / EGDB**
 - Read Only / Publication – **FGDB / EGDB**
 - Mobile User offline editing – **EGDB**
 - etc.
- Generally more than one Geodatabase is required

A goal without a plan is just a wish!

Mobile Data Collection - Offline Workflow Planning			
	Data Maintenance Workflow	Short duration Project Work	Ongoing Project Work
Version from which the feature service is published	<u>Default version</u>	<u>Child version</u>	<u>Child version</u>
Offline version is created for each	Downloaded map	User	User
Number of versions created	Many	Few	Few
Latency between offline edits and updates to Default version	Low	High (1 week)	High (Daily)
Maps involved in quality assurance	One map	All maps	All maps
Frequency that offline versions are deleted	Daily	At project completion	Never



Geodatabase Workflows - QA / QC

- **Design and implement QA / QC workflows**
 - **Data requirements for software functions**
 - **Accurate data for business**
 - **Maintain data integrity**

Capture, load and maintain Data accurately



Data Integrity and Validation Strategies

- **Stage 1: Don't allow start editing**
 - Read Only users
 - Without landbase layers
- **Stage 2: No inserts without pre-requisite checks**
 - Out side of editing areas (Pacific Ocean!)
 - Street light without Poles
 - Equipment without structures
 - Required attribute values (WO Number, Number of Phases, etc.) in attribute columns
 - etc.

Capture, load and maintain Data accurately

Continue....

Data Integrity and Validation Strategies

- **Stage 3: Reconcile/Save edits only after rules validation**
 - Domain checks
 - Connectivity rules, etc.
- **Stage 4: Allow to post data with warnings**
 - Run batch processes to perform additional checks
- **Tools**
 - Domains, Subtypes, Topology, etc.
 - Attribute Assistant Add-In
 - ArcGIS Data Reviewer
 - ArcGIS Workflow Manager
 - Business Partner Products
 - Customization

Capture, load and maintain Data accurately

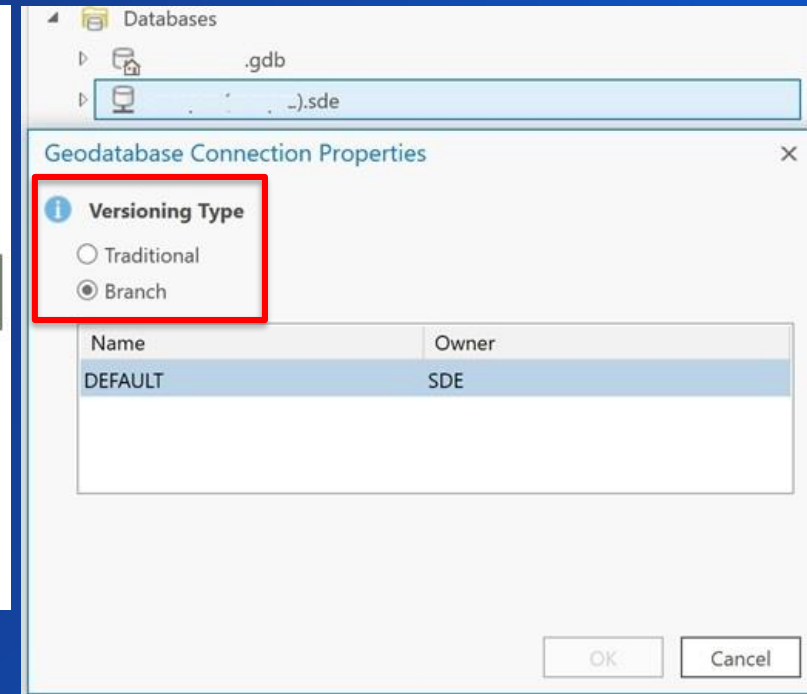
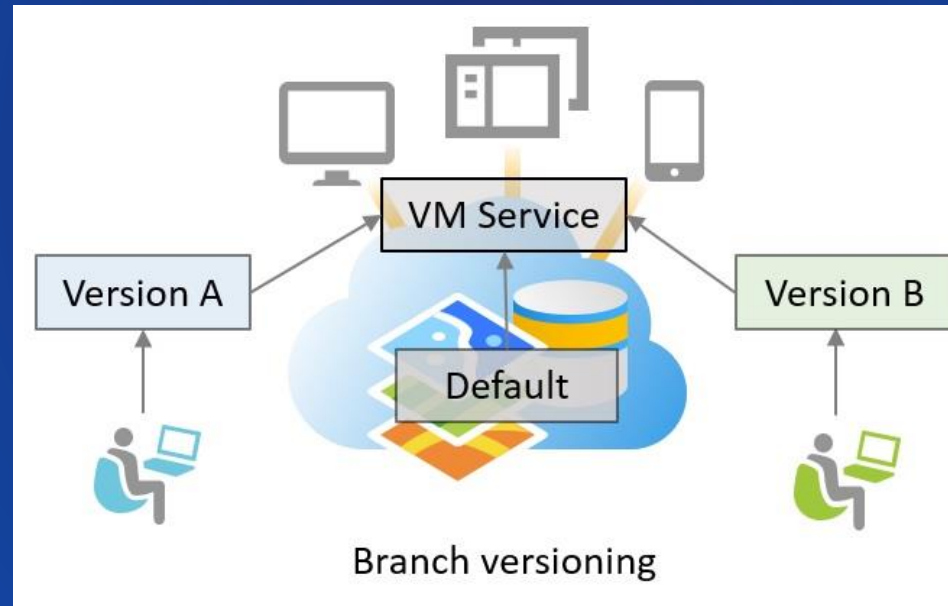
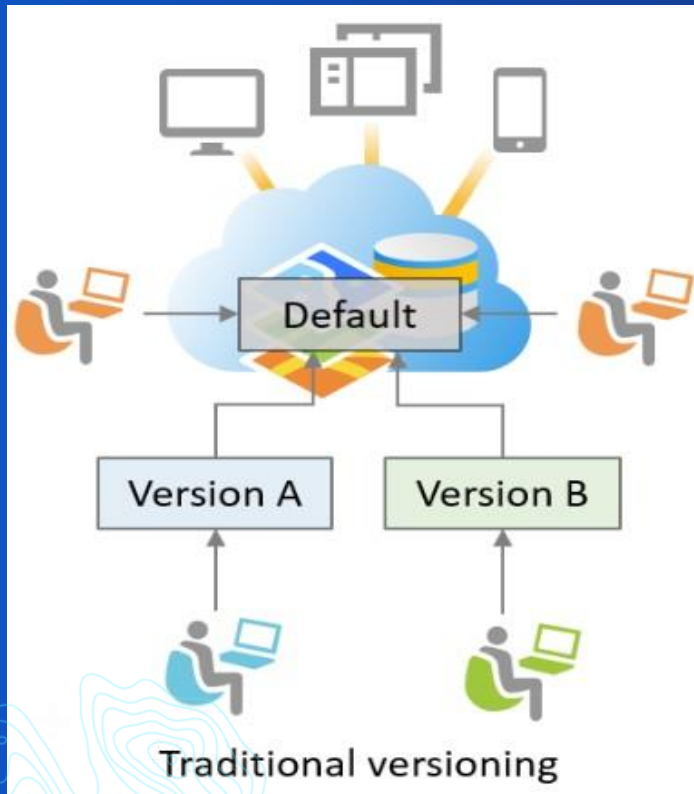
The image displays a comprehensive set of ArcGIS Data Reviewer checks, categorized into 11 groups:

- Database Validation Checks:** Includes checks for geometry generation, attribute validation, and system table validation.
- Table Checks:** Includes checks for table structure, field names, and table relationships.
- Spatial Parameter Evaluation Checks:** Includes checks for spatial parameters, field names, and table relationships.
- Default Checks:** Includes checks for default values, field names, and table relationships.
- Topology Checks:** Includes checks for topology rules, field names, and table relationships.
- Advanced Checks:** Includes checks for advanced topology rules, field names, and table relationships.
- Polygon Checks:** Includes checks for polygon area, field names, and table relationships.
- Z-Value Checks:** Includes checks for z-value ranges, field names, and table relationships.
- Feature on Feature Checks:** Includes checks for feature overlaps, field names, and table relationships.
- Polyline Checks:** Includes checks for polyline length, field names, and table relationships.
- Duplicate Geometry Checks:** Includes checks for duplicate geometries, field names, and table relationships.

Geodatabase Multiuser Workflows

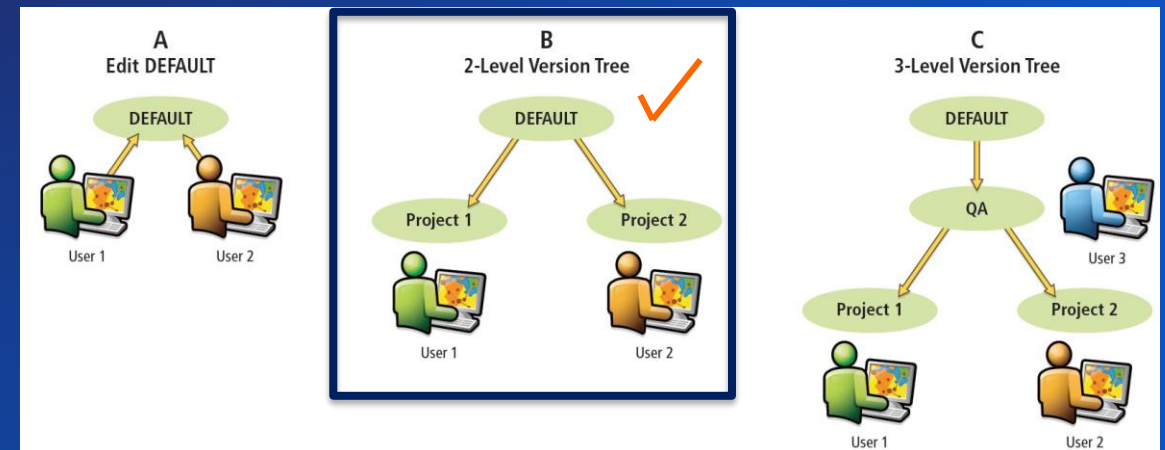
Which Versioning model to choose?

- Traditional Versioning Vs Branch Versioning



Geodatabase Multiuser Workflows – Key Considerations

- Versioning structure
 - Move Edits to Base for simple Feature Classes
 - Recreate the Version after each Post for 3-level Version tree
 - **Problem: Unexpected conflict observed during reconcile**
 - <http://support.esri.com/en/technical-article/000012321>
- Estimate edit volumes and Version durations
- Conflict resolution mechanisms

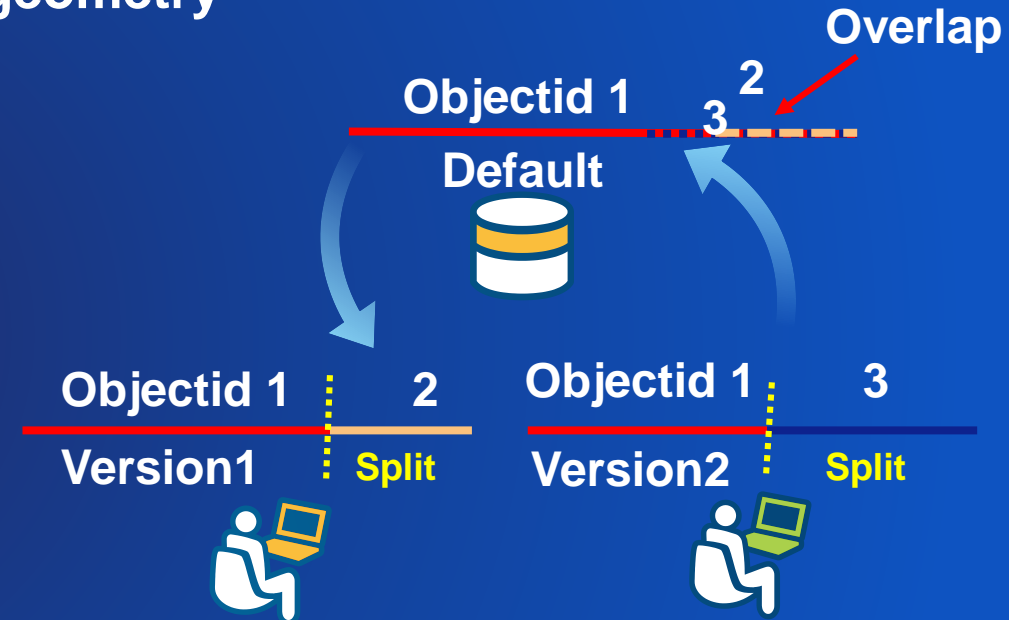


Avoid 3-level version tree

Geodatabase Workflows – Key Considerations





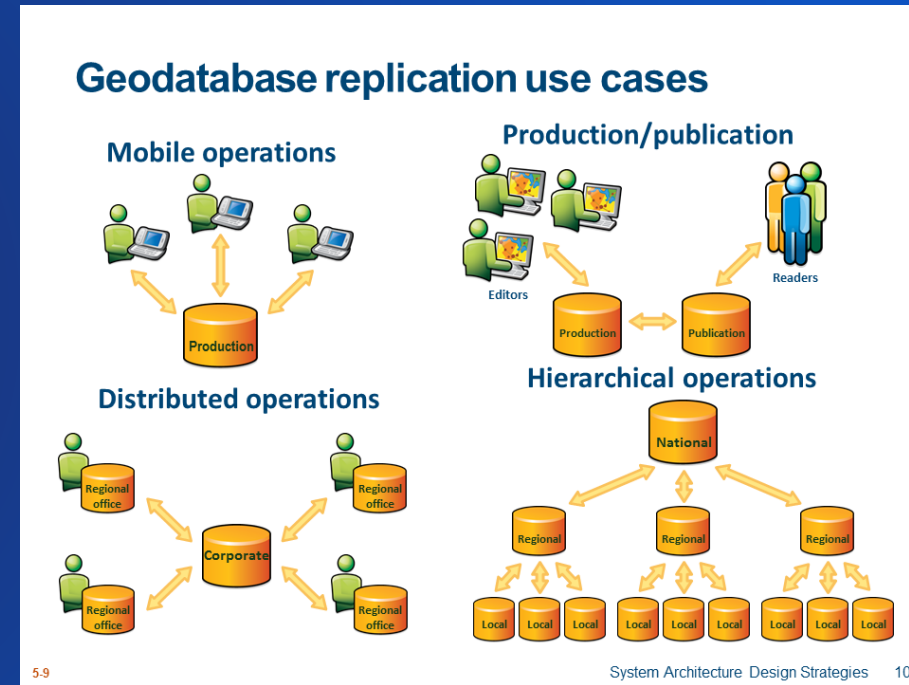
- Conflict for Split operation could introduce duplicate geometry
- Plan bulk loading / mass update
- For regular Data load, consider truncate Vs delete
- Execute batch processes during non-business hours
- Geodatabase Replication - One-Way Vs Two-Way
- Archiving
- Editor tracking



- Conflict on Objectid 1
- Objectid 2 and 3 Overlap

Geodatabase Workflows – Best Practices

- Choose the Versioning workflows appropriately
 - Key to performance and scalability
- Run QA / QC tools regularly
 - Avoid duplicate / invalid geometries, etc.
- Leverage Geodatabase Replication
 - Use One-Way Replication options
 - Parent to Child 
 - Child to Parent 
 - Create Read-Only / Publication Geodatabase
 - Two One-Way Replicas for two separate Datasets/FCs



Geodatabase Workflows – Best Practices

Archiving

- Do not enable Archiving when 100% data update / modification is expected!
- Generally mistaken with the below functionalities / purposes:
 - GDB Editor tracking
 - RDBMS backup and retention
 - Security
 - Data integrity

OWNER	TABLE_NAME	NUM_ROWS	BLOCKS	AVG_ROW_LEN	TO_CHAR(LAST_ANALYZED)
GIS	WSERVICECONNECTION	117,726	12,139	584	APR/10/18 15:32:34
GIS	WSERVICECONNECTION_H	140,242	8,697	528	APR/10/18 15:28:28
GIS	WSERVICECONNECTION_H2	118,057	12,458	605	APR/16/18 08:12:27
GIS	WSERVICECONNECTION_H1	1,653,331	155,838	784	APR/10/18 15:30:13
GIS	S4998_IDX\$	1,653,331		52	APR/24/18 22:07:53

Geodatabase Workflows – Best Practices

- Provide only the required privileges to users
 - Access to large number of tables slows connection performance
- Arrange workflow training for users
 - Conduct tips and tricks session

```
10568 [W 50:00.764] Long: 1
10569 [W 50:00.764] Long: 1
10570 [R 50:00.764] Long: 0
10571 =====
10572 [W 50:00.764] Command: TableListTables
10573 [W 50:00.764] Long: 4
10574 [R 50:48.765] Long: 0
10575 [R 50:48.765] Long: 7835
10576 [R 50:48.765] Dynamic_Str: "SDE.MV_LPA_BAK"
10577 [R 50:48.765] Long: 1
```

Testing and Tuning

Important step before going live!

Testing

- Test application workflows
 - Functionality
 - Flexibility and consistency
- Conduct single user execution test and measure performance
 - Provides Key Indicators towards scalability

Mxdperfstat Before Tuning in Production	Mxdperfstat After Tuning in Production																																																																		
<p>mxdperfstat</p> <p>1/9/2018 2:58:46 PM Master .mxd layerCount= 6 WGS_1984_Web_Mercator_Auxiliary_Sphere esriMeters X= -9,233,816.99 Y= 4,869,583.83 width= 1200 height= 1000</p> <p>Map Display Performance (sec) for each scale</p> <table border="1"> <thead> <tr> <th>Scale</th> <th>Refresh Time(sec)</th> <th>VisibleLayers</th> </tr> </thead> <tbody> <tr><td>1,000,000</td><td>1.30</td><td>3</td></tr> <tr><td>500,000</td><td>1.56</td><td>3</td></tr> <tr><td>100,000</td><td>1.25</td><td>4</td></tr> <tr><td>50,000</td><td>53.83</td><td>5</td></tr> <tr><td>25,000</td><td>50.40</td><td>5</td></tr> <tr><td>12,000</td><td>38.44</td><td>5</td></tr> <tr><td>5,000</td><td>17.14</td><td>5</td></tr> <tr><td>2,500</td><td>51.52</td><td>5</td></tr> <tr><td>1,000</td><td>44.78</td><td>5</td></tr> <tr><td>500</td><td>28.48</td><td>5</td></tr> </tbody> </table>	Scale	Refresh Time(sec)	VisibleLayers	1,000,000	1.30	3	500,000	1.56	3	100,000	1.25	4	50,000	53.83	5	25,000	50.40	5	12,000	38.44	5	5,000	17.14	5	2,500	51.52	5	1,000	44.78	5	500	28.48	5	<p>mxdperfstat</p> <p>1/12/2018 8:31:01 AM Master .mxd layerCount= 6 WGS_1984_Web_Mercator_Auxiliary_Sphere esriMeters X= -9,233,816.99 Y= 4,869,583.83 width= 1200 height= 1000</p> <p>Map Display Performance (sec) for each scale</p> <table border="1"> <thead> <tr> <th>Scale</th> <th>Refresh Time(sec)</th> <th>VisibleLayers</th> </tr> </thead> <tbody> <tr><td>1,000,000</td><td>1.24</td><td>3</td></tr> <tr><td>500,000</td><td>2.10</td><td>3</td></tr> <tr><td>100,000</td><td>.67</td><td>4</td></tr> <tr><td>50,000</td><td>1.14</td><td>5</td></tr> <tr><td>25,000</td><td>1.37</td><td>5</td></tr> <tr><td>12,000</td><td>1.10</td><td>5</td></tr> <tr><td>5,000</td><td>.81</td><td>5</td></tr> <tr><td>2,500</td><td>.92</td><td>5</td></tr> <tr><td>1,000</td><td>.69</td><td>5</td></tr> <tr><td>500</td><td>.83</td><td>5</td></tr> </tbody> </table>	Scale	Refresh Time(sec)	VisibleLayers	1,000,000	1.24	3	500,000	2.10	3	100,000	.67	4	50,000	1.14	5	25,000	1.37	5	12,000	1.10	5	5,000	.81	5	2,500	.92	5	1,000	.69	5	500	.83	5
Scale	Refresh Time(sec)	VisibleLayers																																																																	
1,000,000	1.30	3																																																																	
500,000	1.56	3																																																																	
100,000	1.25	4																																																																	
50,000	53.83	5																																																																	
25,000	50.40	5																																																																	
12,000	38.44	5																																																																	
5,000	17.14	5																																																																	
2,500	51.52	5																																																																	
1,000	44.78	5																																																																	
500	28.48	5																																																																	
Scale	Refresh Time(sec)	VisibleLayers																																																																	
1,000,000	1.24	3																																																																	
500,000	2.10	3																																																																	
100,000	.67	4																																																																	
50,000	1.14	5																																																																	
25,000	1.37	5																																																																	
12,000	1.10	5																																																																	
5,000	.81	5																																																																	
2,500	.92	5																																																																	
1,000	.69	5																																																																	
500	.83	5																																																																	

Performance Tuning Results for an Editing workflow:

S.No	Task	Performance - Before Tuning in Seconds	Performance - After Tuning in Seconds	Performance Gain in %
1	Opening ArcMap	46.22	15.41	199.94%
2	Open MXD (Master.mxd)	28.79	28.06	2.60%
4	Zoom to selected features	8.46	5.36	57.84%
5	Zoom to 5000	7.25	3.46	109.54%
6	Zoom out to 1000	4.93	2.25	119.11%
7	Zoom out to 500	5.32	0.91	484.62%
8	Start Editing	11.95	7.66	56.01%
9	Place a Structure (1 Pole with Push brace)	9.63	0.93	935.48%
10	Insert a Trans Line	7.85	1.11	607.21%
11	Stop Edits	4.73	1.50	215.33%

Tune - Operating System

Operating System

- Adjust and configure
 - Kernel parameters
 - Settings specific to RDBMS and Network capacity
- Enable Large / Huge Memory Pages for Geodatabases
- Update patches

Operation System's Internal Memory Allocation to RDBMS	
OS Default Memory 4KB/ Page	OS Huge / Large Memory 2MB per Page
RDBMS 32GB / 4KB = 8,388,608 Pages	RDBMS 32GB / 2MB = 16,384 Pages


```
HugePages_Total: 16384
HugePages_Free: 5166
HugePages_Rsvd: 1585
HugePages_Surp: 0
Hugepagesize: 2048 kB
DirectMap4k: 67117056 kB
DirectMap2M: 0 kB
```

Tune – RDBMS

RDBMS Configuration and Tuning

- Tune RDBMS
 - Memory allocation from hardware
 - Other initialization parameters
 - Log file settings
 - etc.
- Implement the best practices

RDBMS Tuning - MXDPerfstat Results				
Scale	Feature Class/Layer	Display in Seconds before tuning	Display in Seconds after tuning	Performance gain in %
250,000	Well Top	101.83	0.2	50815.00%
250,000	Last Gas Prod Rate	13.06	0.94	1289.36%
250,000	Cum Gas Production	17.18	0.9	1808.89%
250,000	Last Oil Prod Rate	8.78	0.2	4290.00%
250,000	Cum Oil Production	5.71	0.2	2755.00%
150,000	Well Top	94.04	0.14	67071.43%
150,000	Last Gas Prod Rate	9.87	0.86	1047.67%
150,000	Cum Gas Production	9.05	0.87	940.23%
150,000	Last Oil Prod Rate	4.07	0.11	3600.00%
150,000	Cum Oil Production	3.29	0.12	2641.67%
5,000	Well Top	36.71	0.03	122266.67%
1,000	Well Top	37.19	0.03	123866.67%
1,000	Last Gas Prod Rate	5.13	0.03	17000.00%
1,000	Cum Gas Production	5.1	0.03	16900.00%



Maintenance

Increase Performance and Scalability

Failing to prepare is preparing to fail!



- Workflow estimations
 - Number of outstanding versions
 - Versioning levels
 - Archiving
 - Traditional Vs Branch Versioning
 - etc.
- Maintenance plan
 - Mandatory tasks to keep performance
 - Delta table records Vs CPUs
 - Roles and responsibilities
 - etc.

S.No	Display Scale	Layer Name	Before Maintenance- Display in Seconds	After Maintenance - Display in Seconds	Performance Improvement in %
1	50,000	RoadCL > 10,000	0.34	0.3	13.33%
2	15,000	Water Mains	0.15	0.09	66.67%
3	10,000	MapLink	1.02	0.08	1175.00%
4	10,000	CH2M_Mains	1.15	0.97	18.56%
5	10,000	Leaders	0.58	0.07	728.57%
6	10,000	Annotation	0.69	0.15	360.00%
7	10,000	Water Mains	0.96	0.6	60.00%
8	10,000	CTParcels	1.2	0.42	185.71%
9	6,000	CH2M_Mains	0.68	0.56	21.43%
10	6,000	Water Mains	0.88	0.52	69.23%
11	1,000	SwingTies	7.74	0.05	15380.00%
12	500	SwingTies	6.12	0.74	727.03%
13	500	Water Mains	0.72	0.24	200.00%
14	500	sbDriveways	0.37	0.16	131.25%

**Few private and orphan versions induced the bottlenecks*

Mitigate the risk with proper system capacity & maintenance

Geodatabase Maintenance - Strategies

- Identify a maintenance window and tasks
- Categorize
 - Nightly, Weekly, Monthly and Yearly.
- Classify manual and automated batch processes
 - Design scalable batch processes
- Assign SDE/GIS administrator role
- Monitor




Geodatabase needs maintenance – Plan one

Geodatabase Maintenance – Key Considerations

- Reconcile, Post and Compress schedule
- Underlying RDBMS requires maintenance other than backup
 - Rebuild Index
 - Update Statistics
 - Logs
- Execute automated processes only within maintenance window
- Run repair version tables and metadata
 - (Previously: SDEGDBREPAIR) every ~3 months
 - Fix any inconsistencies
 - Schedule the execution around weekends




Automate the daily maintenance process



Versions maintenance

- Use private Versions for larger versioning environment
 - Keep the total number less
 - Sync frequently
 - Monitor
- Complete the batch processes within maintenance period
 - Improve Hardware and/or Software design
 - Additional CPUs
 - Multi threads
 - Feature/Schema cache




Geodatabase Maintenance – Best Practices

- Identify, Reconcile and Post top 5 blocking Versions every day
 - Blocking Versions cause inefficient Compress
 - Increase in Delta Table records beyond the hardware support level
 - Maintain the lineage length <100

```
Versioning statistics
=====
Number of versions: 664
Number of versions blocking DEFAULT: 625
Top 5 blocking versions...
  ANDREW.SN_164022
  ANDREW.SN_162751
  RASU.SN_163090
  RASU.SN_163139
  RASU.SN_164468
Number of states: 4333
Number of state lineages: 86293
DEFAULT versions lineage length: 426
Last compress: JUL-09-2017
```

SQL Server DBs ArcSDE A Table Rows -- ADDS TABLE RECORD COUNT (COUNTER) -- Principle: track record count

Chart resolution: real-time value at collection interval when query less than 12 hrs



Geodatabase Maintenance – Best Practices

- Remove Geoprocessing (GP) history
 - *How To: Automate the process of deleting geoprocessing history*
 - <http://support.esri.com/technical-article/000011751>

- Disable (GP) history for scripts

import arcpy

arcpy.SetLogHistory(False)



Python Tasks	Before Deleting GP History in Seconds	After Deleting GP History in Seconds	Performance Improvement in %
Create Version	400.74	16.3	2358.53%
Delete Version	571.23	14.17	3931.26%

Feature Class Name	Calculating 2 records before GP History Removal in Seconds	Calculating 2 records after GP History Removal in Seconds	Performance Difference in %
Sewer	58.67	8.09	625.22%

Geodatabase Maintenance – Best Practices


- Manage auditing / history tables
 - Reduce the database
 - Backup size
 - Storage
 - Time



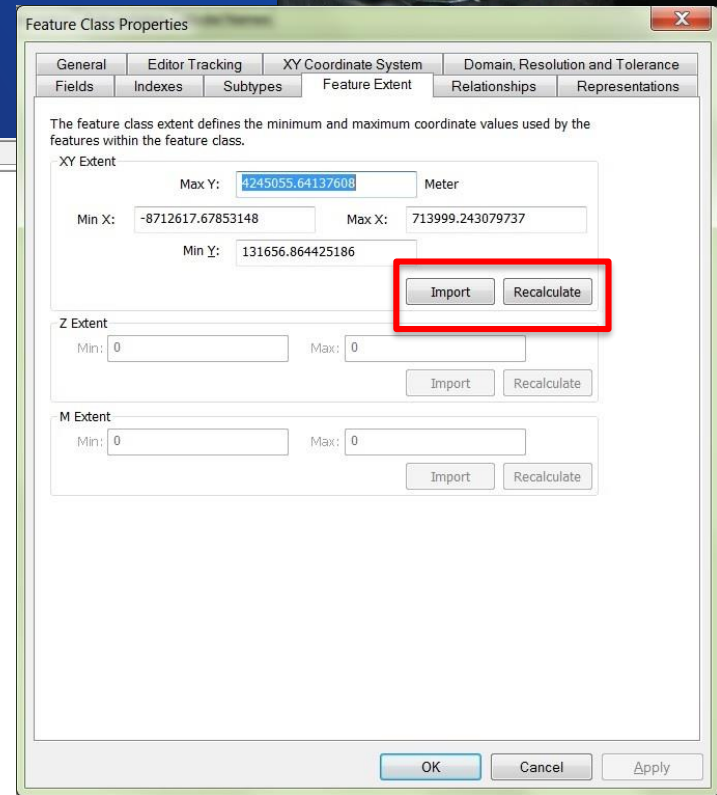
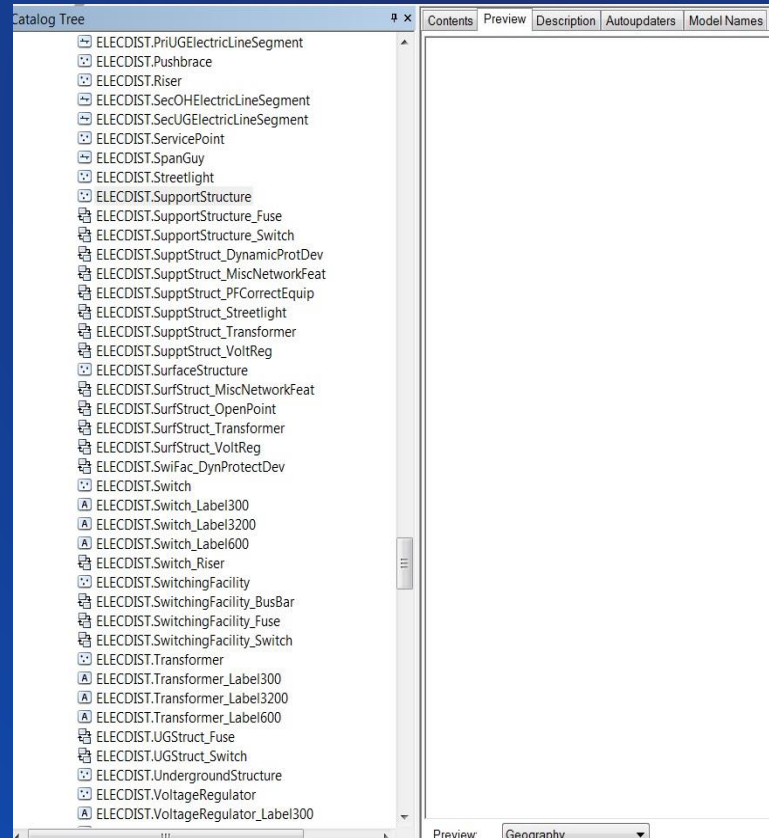
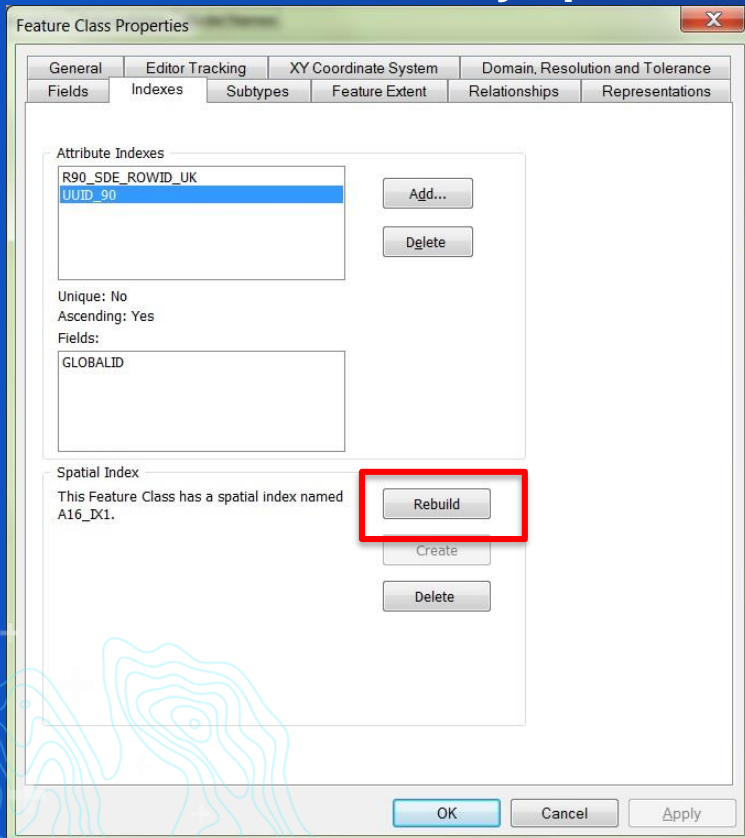
OWNER	TABLE_NAME	NUM_ROWS	BLOCKS	AVG_ROW_LEN	TO_CHAR(LAST_ANALYZED, 'MON/DD/YYHH24:MI:SS')
ARCFM	EDITEDFEATURESTRACKING	12330909	122954	140	APR/23/16 23:44:17
ARCFM	EDITEDGRIDS	12964827	32969	32	APR/23/16 23:47:55
SDE	ARCSDEUSERLOG	56558633	259246	62	APR/24/16 12:28:25

OWNER	TABLE_NAME	NUM_ROWS	BLOCKS	AVG_ROW_LEN	TO_CHAR(LAST_ANALYZED, 'MON/DD/YYHH24:MI:SS')
SDE	GDBM_RECONCILE_HISTORY	4759429	50484	72	AUG/12/17 12:40:34
ARCFM	USAGEINFORMATION	6377236	221057	39	AUG/10/17 15:59:14

Geodatabase Maintenance – Best Practices



- Every 3 - 6 months:
 - Fix the feature class extent first
 - Followed by spatial index rebuild



Geodatabase Maintenance – Best Practices

- Fix the non-empty Feature Classes with no spatial index

```
1 Exec DBMS_STATS.GATHER_SCHEMA_STATS ('ARCFM', estimate_percent=>100,
2   DEGREE=> 7, CASCADE=>TRUE, No_Invalidate=>false);
3 select table_name, GSIZE1, GSIZE2, GSIZE3 FROM sde.layers
4   where gsize1=0 and gsize2=0 and gsize3=0 and table_name in
5   (select table_name from all_tables where num_rows > 0);
```

TABLE_NAME	GSIZE1	GSIZE2	GSIZE3
COATING	0	0	0
CPBOND	0	0	0
CPCABLE	0	0	0
CPGROUNDBED	0	0	0
CPANODE	0	0	0
LINECROSSING	0	0	0
MARKER	0	0	0
PIPEEXPOSURE	0	0	0
TIEINMETHOD	0	0	0
DOCUMENTPOINT	0	0	0
INJECTION	0	0	0
PIGROUTE	0	0	0
PIPEDEPTH	0	0	0
RL_SITE	0	0	0
PIPEJOIN	0	0	0
PIPELENGTH	0	0	0
WELL	0	0	0
LEAK	0	0	0
DOT_CLASS_PREVIOUS_RDETAILS	0	0	0
ALIGNMENT_SHEETS	0	0	0
EXTERNAL_COATING_RDETAILS	0	0	0
MISC_FITTING_DETAILS	0	0	0
ALIG_SHT_CROSS_REF_RDETAILS	0	0	0
CLOSURE_DETAILS	0	0	0
HCA_PREVIOUS_RDETAILS	0	0	0
INJECTOR_DETAILS	0	0	0
ODORANT_RANGE_RDETAILS	0	0	0
OFFLINECOMPSTATIONPOLYGONS	0	0	0
PIG_SIGNAL_DETAILS	0	0	0
PIR_RDETAILS	0	0	0
RIGHT_OF_WAY_RDETAILS	0	0	0
RIVER_WEIGHT_RDETAILS	0	0	0
SHEET_NOTE_RDETAILS	0	0	0
TAP_DETAILS	0	0	0
PIPE_SEG_PIR_BUFF	0	0	0
MAOP_CALC_RDETAILS	0	0	0



Geodatabase Maintenance – Best Practices

- A Sample daily maintenance (batch process) for multi user Geodatabase:
 1. Backup the Database
 2. Synchronize any Replica version.
 3. Delete the orphan / unnecessary versions.
 4. Drop the orphan keyset tables
 5. Reconcile and Post all/eligible Versions – through out the Day
 6. Only Reconcile all versions (>100 versions - parallel reconcile)
 7. Update Database statistics - optional



Continue.....

```
50 END LOOP;
51
52 dbms_output.put_line('Dropped '||cnt||' keyset tables.');
```

53
54 END;
55 /
Dropped 31038 keyset
tables.

```
SDE_VERSION_NAME
-----
HEALTHCOMMENT
-----
RASU.SYNC_SEND_96817_3
Critical: Orphaned replica version (http://support.esri.com/technical-article/000010858)
RASU.SYNC_SEND_90803_2
Critical: Orphaned replica version (http://support.esri.com/technical-article/000010858)
```

Geodatabase Maintenance – Best Practices



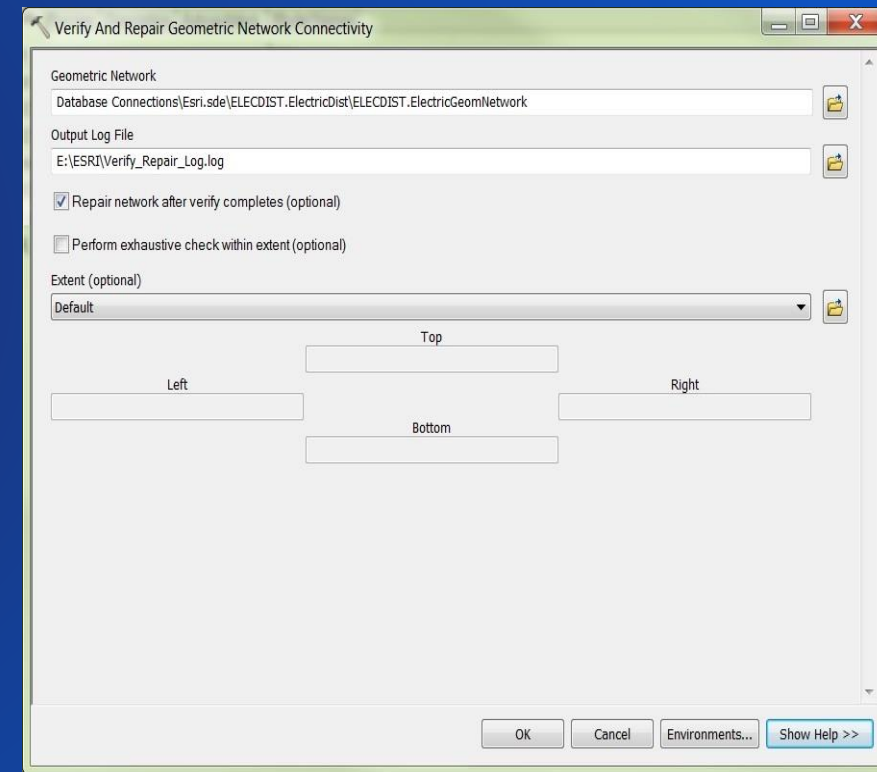
- Continue.....

8. Pause the SDE connections
9. Kill the existing or orphaned user connections
10. Truncate dynamic tables
 - A. state_locks; table_locks; object_locks; layer_locks; process_information; <user>.SDE_LOGFILE_DATA;
11. Start the Compress process
12. Un-pause the SDE connection
13. Rebuild indexes in RDBMS for all the Schema Owners and SDE
14. Update RDBMS statistics for all Schema Users and SDE.

OWNER	TABLE_NAME	NUM_ROWS	BLOCKS	AVG_ROW_LEN	TO_CHAR(LAST_ANALY
RASU	SDE_LOGFILES	7,488	244	93	JUN/17/19 22:03:13
RASU	SDE_LOGFILE DATA	86,832,159	171,379	10	JUN/03/19 22:02:56
ANDR	SDE_LOGFILES	783	13	91	MAY/14/19 22:02:05
ANDR	SDE_LOGFILE DATA	28,887,671	57,017	10	AUG/26/18 06:04:45

Geodatabase Maintenance – Geometric Network

- Every 3 – 6 Months run Esri’s Verify And Repair Geometric Network connectivity tool
- Only GIS Administrator should follow the below steps:
 - Create a new Version under SDE.Default.
 - Create a SDE connection document with the new Version.
 - Run the Verify and Repair tool
 - With “Repair network after verify completes” option
 - Reconcile, Post and Delete the newly created version.
 - Compress the Database.



Caution: Individual users should not run it. Generates larger number of delta table records and it can produce more conflicts when run under SDE.Default version directly.

Monitor

ArcGIS Monitor

- Add RDBMS Queries
- Monitor key performance indicators
- Keep 15 minutes sampling interval

ArcGIS Monitor Categories > Database

Home Availability Alerts Categories Site Reports Help

Prod

- ArcGIS Datastore
 - ArcGIS Datastore Activity
 - BlockRead(msec)
 - BlockWrite(msec)
 - BufferCacheHit(Ratio)
 - Collection Time(sec)
 - Connections
 - ConnectionsPortal
 - ConnectionsPortal(%)
 - IndexToTableScan(Ratio)
 - ReadPhysicalToMemory(Ratio)
 - TableLocksOther
 - TableLocksShared
 - TempFiles
 - Trf(PerSec)

Collection Status

Alerting	Monitor	Name
✓	✓	DataStore Activity
✓	✓	Oracle Activity
✓	✓	Oracle Performance
✓	✓	DataStore Performance
✓	✓	SQL Server Activity
✓	✓	SQL Server Performance
✓	✓	DataStore Status
✓	✓	SQL Server Status
		Oracle Status

Production Database

- ArcMap Active Sessions
- ArcSDE A Table Rows
- ArcSDE D Table Rows
- ArcSDE Editors
- ArcSDE State Lineage Depth
- ArcSDE State Lineages
- ArcSDE States
- ArcSDE Versions
- ArcSDE Versions >07 Days Old
- ArcSDE Viewers
- Block Read Latency ms
- Buffer_Cache_Hit_Ratio
- Collection Time(sec)
- Database_CPU_Time_Ratio
- Executions_Per_Sec
- Executions_Per_Txn
- Hard Parse count/Sec
- Hard Parse count/Trnsxn
- Host Avg CPU Last Min
- Host Process
- IO_Megabytes_per_Second
- IO_Requests_per_Second
- Logical_Reads_Per_Sec
- Logical_Reads_Per_Txn
- Physical Reads Per Sec
- Response Time /Trnsxn
- Total User Memory (Mb)
- Transaction/Sec

1/21/2018 12:00:00 AM 1/27/2018 11:59:59 PM OK 8 Hours 12 Hours Day Week

ArcSDE State Lineages

Chart Resolution: hourly average when query greater than 12 hrs

Category Database, type DB query, source DB query, time Timespan, stat name DB query

ID	Collection	Name	Instance	Min	Avg	p5	p50	p75	p95	p99	Max	LastUpdated	Comments
1	Prod	Add Tables Oracle Stats	2197136	2225871.16	2198031	2217611	2251894	2264732	2265925	2265954	2018-01-11 11:24:00-06:00		
2	Prod	Delete Tables Oracle Stats	1481956	1514668.22	1483060	1514105	1540245	1544028	1549550	1549550	2018-01-11 11:24:00-06:00		
3	Prod	Add Tables Oracle Stats	620209	650699.38	620450	645498	666022	685179	696552	696552	2018-01-11 11:24:00-06:00		
4	Prod	Delete Tables Oracle Stats	600846	636172.18	601045	631832	651928	671062	682420	682420	2018-01-11 11:24:00-06:00		
5	Prod	State Lineages	13220	73214.32	13690	52102	58491	244017	331950	343105	2018-01-11 11:24:00-06:00		
8	Prod	Add Tables Oracle Stats	37584	38138	37584	37728	38086	41146	44291	44291	2018-01-11 11:24:00-06:00		
9	Prod	Delete Tables Oracle Stats	35637	35927.69	35637	35697	35868	37566	39350	39350	2018-01-11 11:24:00-06:00		
10	Prod	State Lineages	395	1475.53	414	875	1438	4736	14615	14615	2018-01-11 11:24:00-06:00		
12	Prod	States	311	1128.08	316	930	1479	2134	2528	2672	2018-01-11 11:24:00-06:00		
13	Prod	Lineage Depth Default	67	563.7	72	380	768	1466	1718	1852	2018-01-11 11:24:00-06:00		
15	Prod	Versions	634	640.36	637	637	641	651	651	651	2018-01-11 11:24:00-06:00		
16	Prod	States	42	166.55	42	122	244	448	595	648	2018-01-11 11:24:00-06:00		
18	Prod	Lineage Depth Default	20	61.76	20	40	91	182	220	243	2018-01-11 11:24:00-06:00		
21	Prod	Compress Sec	39	65.84	39	66	75	121	121	121	2018-01-11 11:24:00-06:00		
23	Prod	Versions	50	64.92	50	66	68	78	78	78	2018-01-11 11:24:00-06:00		
27	Prod	State Lineages	4	8.49	4	4	15	23	34	34	2018-01-11 11:24:00-06:00		
33	Prod	Compress Sec	0	1.82	0	1	1	18	18	81	2018-01-11 11:24:00-06:00		
34	Prod	States	3	6.27	3	3	11	15	25	25	2018-01-11 11:24:00-06:00		
35	Prod	Lineage Depth Default	3	5.88	3	3	10	15	23	23	2018-01-11 11:24:00-06:00		
41	Prod	Versions	4	4	4	4	4	4	4	4	2018-01-11 11:24:00-06:00		
61	Prod	Compress Sec	0	0.13	0	0	0	0	0	50	2018-01-11 11:24:00-06:00		

Statistics: only table reflects true statistics like min, max, percentile

Item	Category	Counter	Instance	Min	Avg	P5
GISP	Db	ArcSDE State Lineages	ArcSDE State Lineages	67,000	1,535,921	90,000
GWDP	Db	ArcSDE State Lineages	ArcSDE State Lineages	9,877,000	102,939,334	15,224,000

ArcGIS Monitor – EGDB Add-on

- Configure EGDB Add-ons per RDBMS
- Use EGDB health tool

database_name	owner_name	delta_table_name	base_table_name	comment
GIS_Common	dbo	a146	ADDITIONAL_PROPERTY	Critical: Delta tables exist and there is no compress history
GIS_Common	dbo	D146	ADDITIONAL_PROPERTY	Critical: Delta tables exist and there is no compress history
GIS_Common	dbo	a147	ER_	Critical: Delta tables exist and there is no compress history
GIS_Common	dbo	D147	ER_	Critical: Delta tables exist and there is no compress history
GIS_Common	dbo	a148	FACILITY	Critical: Delta tables exist and there is no compress history
GIS_Common	dbo	D148	FACILITY	Critical: Delta tables exist and there is no compress history

1	stats_name	table_name	type_desc	owner_name	modify_date	statistics_recompute_setting	statistics_user_created	comment
2	d160_pk	D160	USER_TABLE	dbo	42710.54988	Recompute	Not user created	Warning: No statistics update information available. Review statistics updating procedures.
3	d160_idx2	D160	USER_TABLE	dbo	42710.54988	Recompute	Not user created	Warning: No statistics update information available. Review statistics updating procedures.
4	a160_rowid_idx1	a160	USER_TABLE	dbo	42710.55096	Recompute	Not user created	Warning: No statistics update information available. Review statistics updating procedures.
5	a160_state_idx2	a160	USER_TABLE	dbo	42710.55096	Recompute	Not user created	Warning: No statistics update information available. Review statistics updating procedures.
6	d189_pk	D189	USER_TABLE	dbo	42710.55745	Recompute	Not user created	Warning: No statistics update information available. Review statistics updating procedures.
7	d189_idx2	D189	USER_TABLE	dbo	42710.55745	Recompute	Not user created	Warning: No statistics update information available. Review statistics updating procedures.

- <https://community.esri.com/community/implementing-arcgis/blog/2019/05/10/using-egdbhealth-to-evaluate-a-geodatabase?et=watches.email.blog>

Monitor, interpret and respond

Tools for Implementation Assistance

ArcGIS Monitor

- **ArcGIS Monitor**
 - ArcGIS Monitor is a tool for monitoring and analyzing your enterprise GIS system
 - <https://www.esri.com/en-us/arcgis/products/arcgis-monitor/overview>
- **MXDPerfstat**
 - An ArcGIS Engine command line tool to diagnose typical mxd performance problems
 - <https://www.arcgis.com/home/item.html?id=a269d03aa1c840638680e2902dadecac>
- **System Designer**
 - A comprehensive tool for designing and capacity planning of GIS solutions.
 - <https://www.arcgis.com/home/item.html?id=8ff490eef2794f428bde25b561226bda>
- **System Log Parser**
 - A reporting tool specifically designed for analyzing ArcGIS server and service logs
 - <http://www.arcgis.com/home/item.html?id=a29649a3d87d4cae84374e5d711dc3aa>

Questions and Answers

Contact Info:

Rasu Muthurakku

rasu@esri.com

Andrew Sakowicz


asakowicz@esri.com

See us here

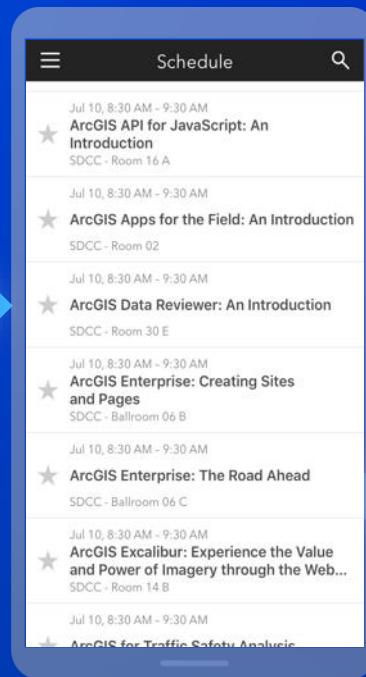
WORKSHOP	LOCATION	TIME FRAME
<ul style="list-style-type: none">• Enterprise Geodatabase: Automating Administration Tasks Using Python	<ul style="list-style-type: none">• SDCC - Expo Demo Theater 04	<ul style="list-style-type: none">• Thursday 7/11/2019 10:00 AM - 10:45 AM
<ul style="list-style-type: none">• Geodatabase: Ensuring Data Quality with Attribute Rules and Contingent Values	<ul style="list-style-type: none">• SDCC - Ballroom 06 E	<ul style="list-style-type: none">• Thursday 7/11/2019 4:00 PM - 5:00 PM

Please Share Your Feedback in the App

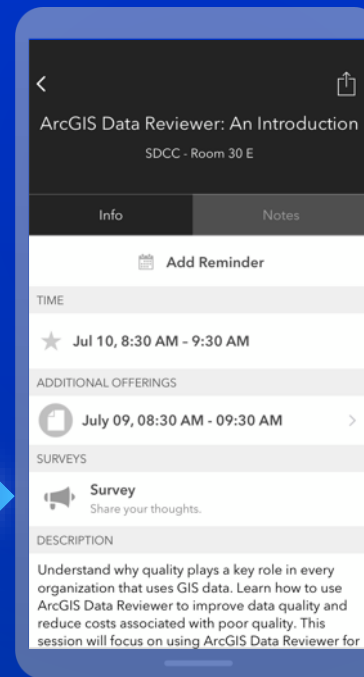
Download the Esri Events app and find your event




Select the session you attended



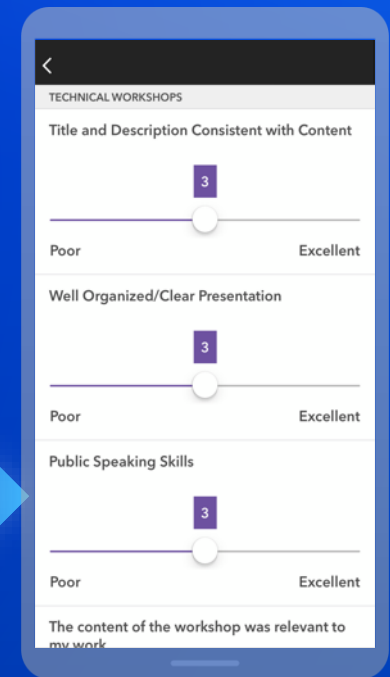
Scroll down to "Survey"



Log in to access the survey



Complete the survey and select "Submit"



Thanks!



esri

THE
SCIENCE
OF
WHERE