

Troubleshooting Performance Issues with Enterprise Geodatabases

Ben Lin, Nana Dei, Jim McAbee

blin@esri.com

ndei@esri.com

jmcabee@esri.com

Workshop Agenda

- Performance Troubleshooting (Then & Now)
- Real-World Examples and Troubleshooting
- Tips, Tricks and Guidance
- Tools
- Q&A Open Forum

Few questions to ask first...

- **What is slow?**
 - 30 seconds may be ok for a batch database process, but not an OLTP transaction
- **What is between the end user and the Enterprise Geodatabase?**
 - Understand infrastructure and processes, and at time of day
- **What was the workflow and specific task being done and when?**
 - What was the end user, or system, doing when the issue was noticed

Troubleshooting Enterprise Geodatabase Performance Issues

- **ArcGIS Platform data management includes many solutions**
 - ArcGIS Data Store
 - File Geodatabases
 - Imagery
 - Enterprise Geodatabases
 - Other structures
- **This session focuses on issues with Relational Databases and Enterprise Geodatabases**




Performance Problem Examples

- Example 1: **The Poorly Performing Web Application**
- Example 2: **The Angry DBA**
- Example 3: **The Disappointed GIS Analyst**
- Example 4: **The “It is Slow 50% of the Time” System**
- Example 5: **The Slow Afternoon**
- Example 6: **The Slow Data Maintenance Operations**

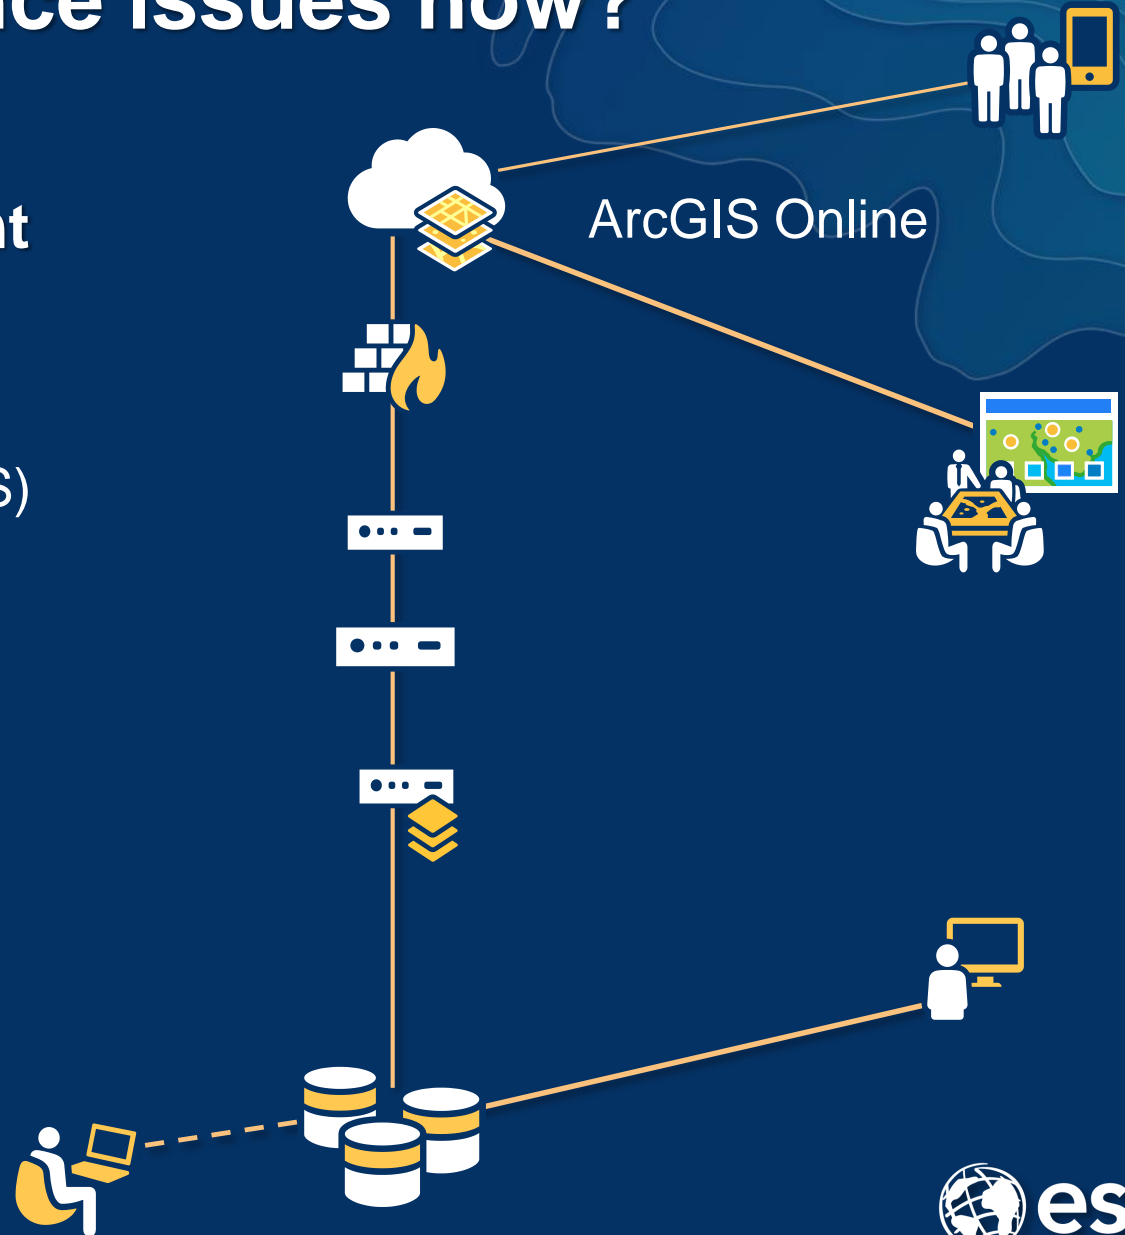

Troubleshooting performance issues **used to be simple**

- Easier to find in the past...
- Geospatial information and use was less pervasive in IT...
- Performance issues were more isolated...
- Fewer users...



Where are the performance issues now?

Now more components are spread out over more deployment environments and locations....



Spatial Data Management is now **Enterprise-Wide**


- Geospatial information & use pervasive in Enterprise IT
- Business System Integration
- Users Enterprise-Wide
- All good, but making performance troubleshooting more difficult



ArcGIS Pro & Desktop


Enterprise Geodatabase

ArcGIS Servers



ArcGIS Pro & Desktop


Enterprise Geodatabase



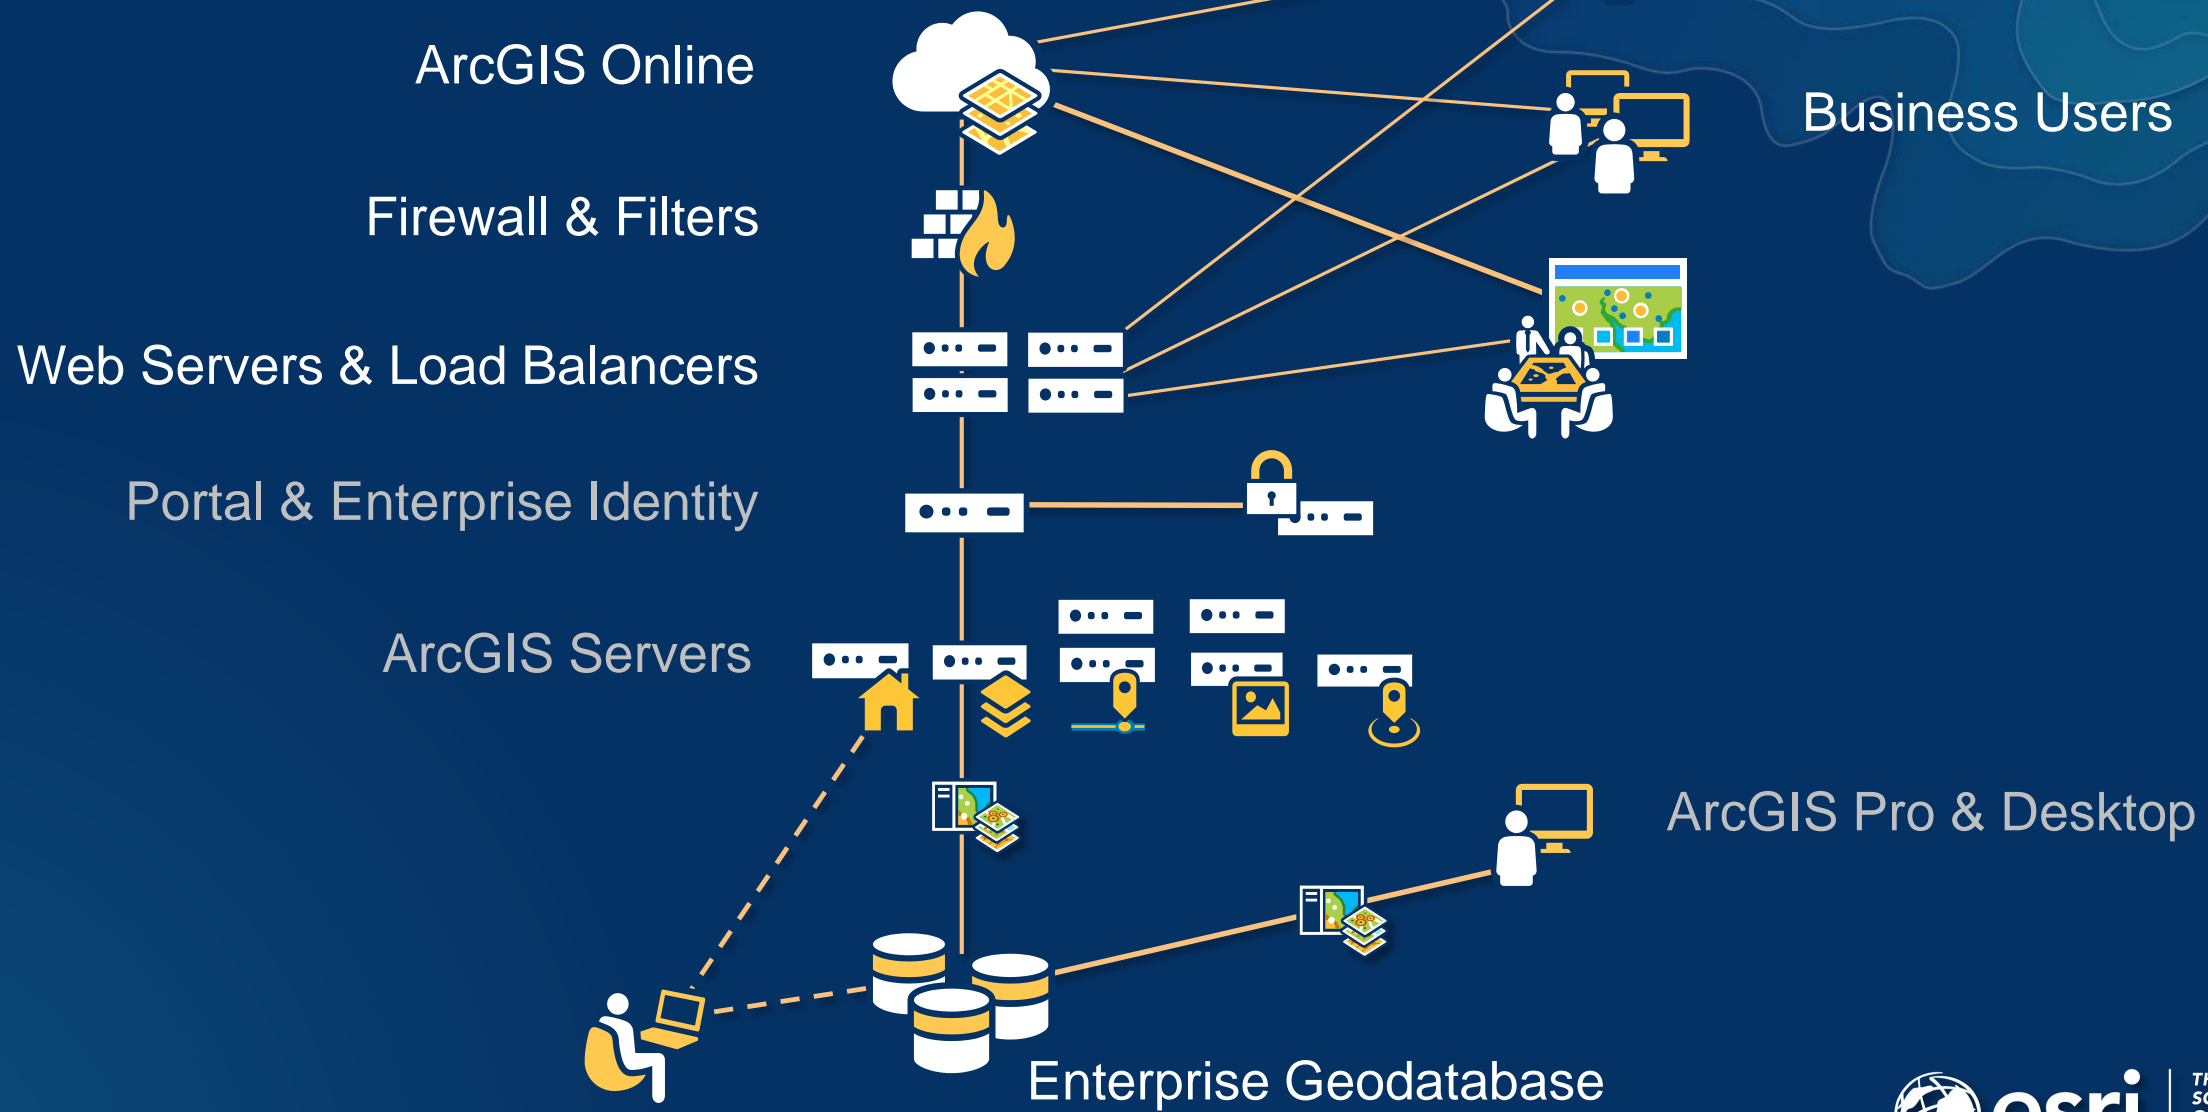
Portal & Enterprise Identity

ArcGIS Servers

ArcGIS Pro & Desktop

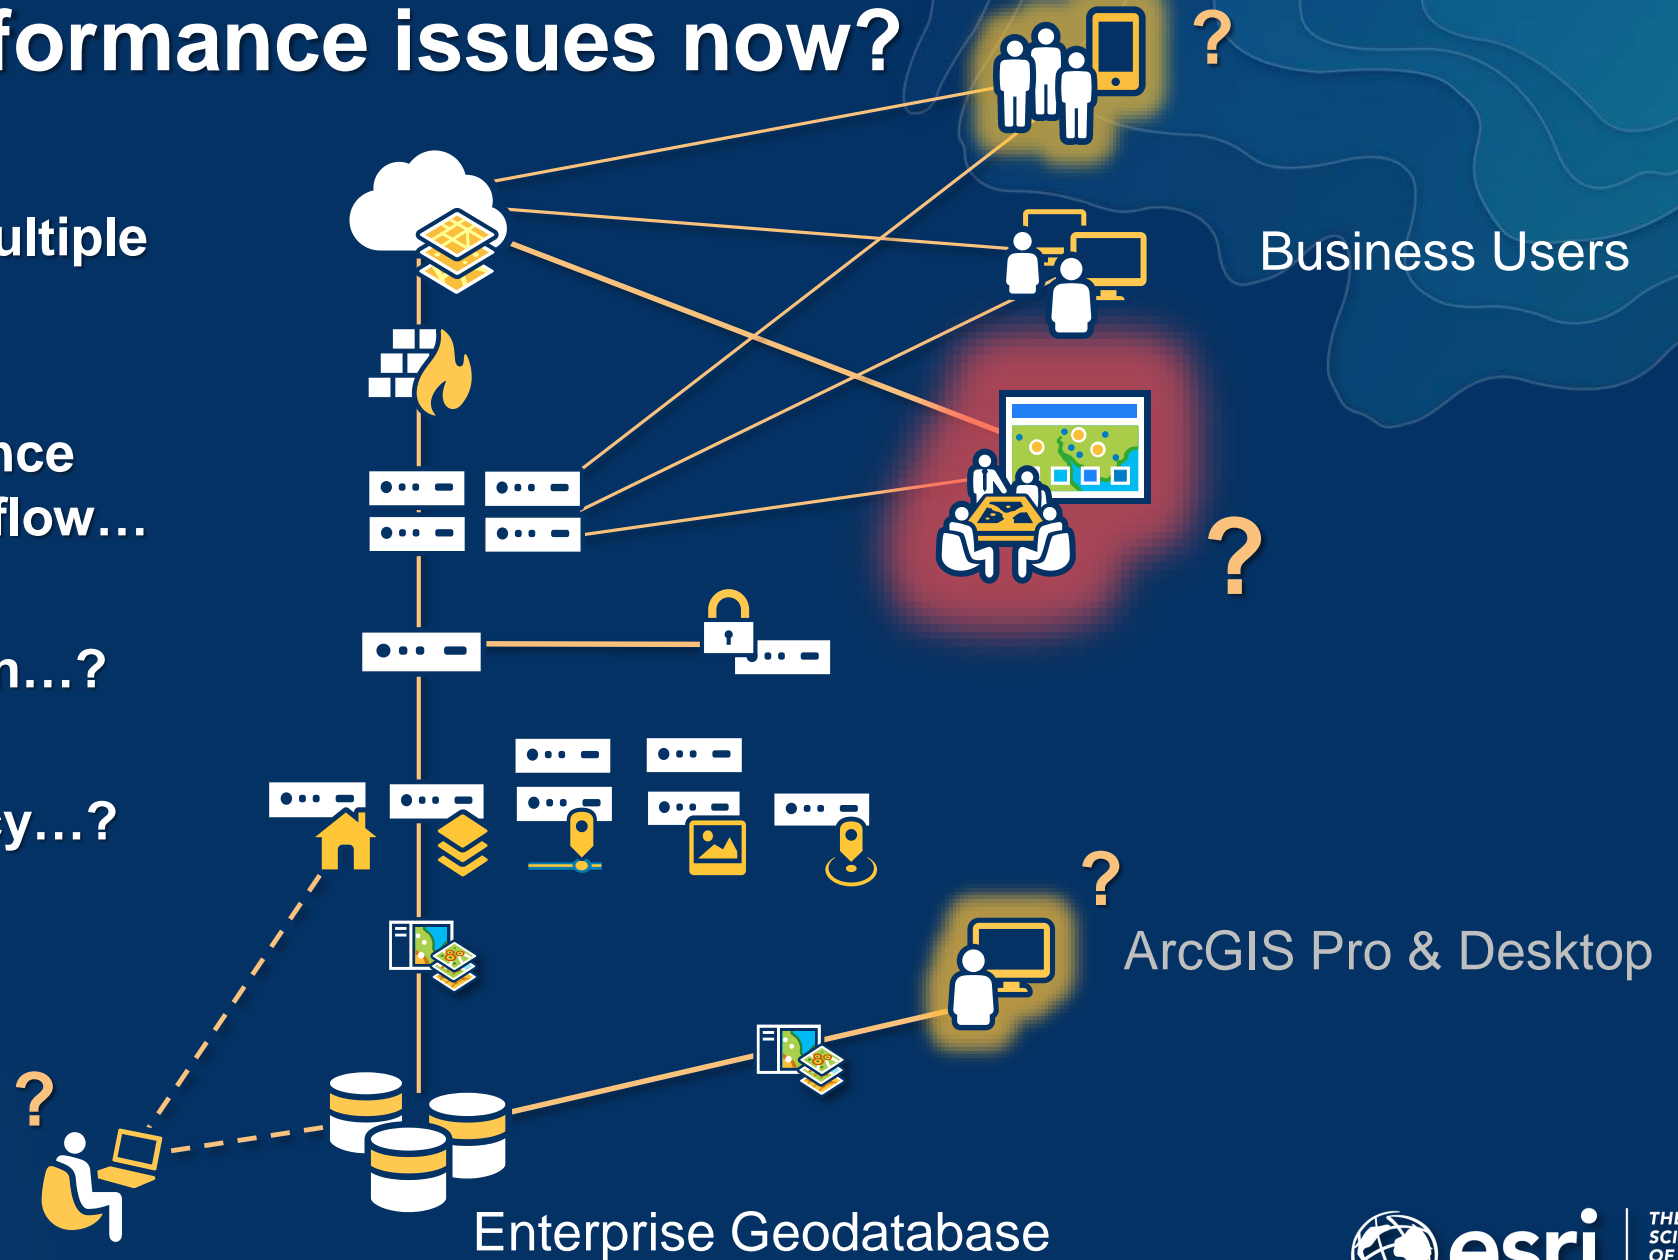


Enterprise Geodatabase




Where the performance issues now?

- May be single or multiple issues...
- May be a performance issue or poor workflow...
- Where do you begin...?
- And what is urgency...?



Where to Start: **Focus and Isolate**

- Prioritize and isolate
- Check one issue at a time
- Check if it is the database
- Check component resources
- Check workflow, map, service configuration
- Check network design, security, latency
- Check business application/workflow
- Check relationships between issues?



Examples 1 & 2: The Poorly Performing Web App & Angry DBA

- Occurred at same customer site...

Example 1: Poorly performing Web Application – Business Users

- How is application/service configured?

Example 2: Application Generating too much SQL (load) – the Angry DBA

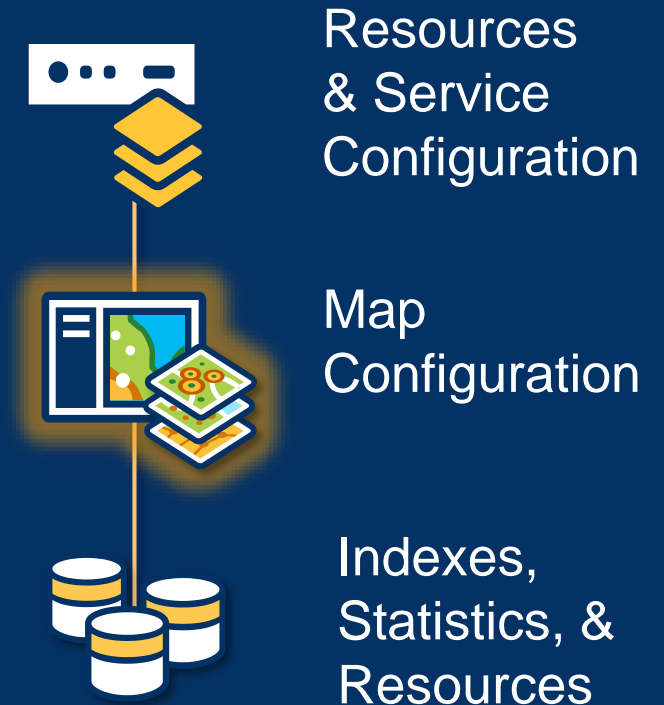
- Which application and how is it configured?

- Few Questions...

- Are they the same problem?
- Where is the problem?
- Hint - database traces and sde intercepts may not be best first step...is it the database?

Examples 1 & 2: Resolution: The Poorly Performing Web App & Angry DBA

- **Poorly performing Web Application – Business Users**
- **Application Generating too much SQL (load) – DBA**
- **Process**
 - Review of Web Application, found (75 layers using same feature class with definition queries)
 - Isolated issue to an application issue NOT database
 - (overly complex map document)
 - Check indexing on definition query fields or symbology
 - Check scale-dependency on layers
 - Read-only Publication Web App (consider publication File Geodatabase)



Example 3: The Disappointed GIS Analyst

- **Slow performance editing data – GIS Analyst**
- **Is it the same issues that were causing examples 1 & 2 issues? No**
 - Starting edit, creating feature, saving edit, snapping?
 - ArcGIS Desktop / Pro (and Server) = “Thick” Client (heavy traffic)
 - What is the total number of feature classes in feature dataset / map?
 - How many rows in the state_lineages table?
 - When was the last compress?
 - Where is editor located, where is database located...

Where to Start?

- Distance = Latency
- Services can communicate more efficiently across distances (WAN and Internet)
- Database communication is much more “chatty”
- Connections directly to database need to be close

ArcGIS Server
Services



Enterprise
Geodatabase

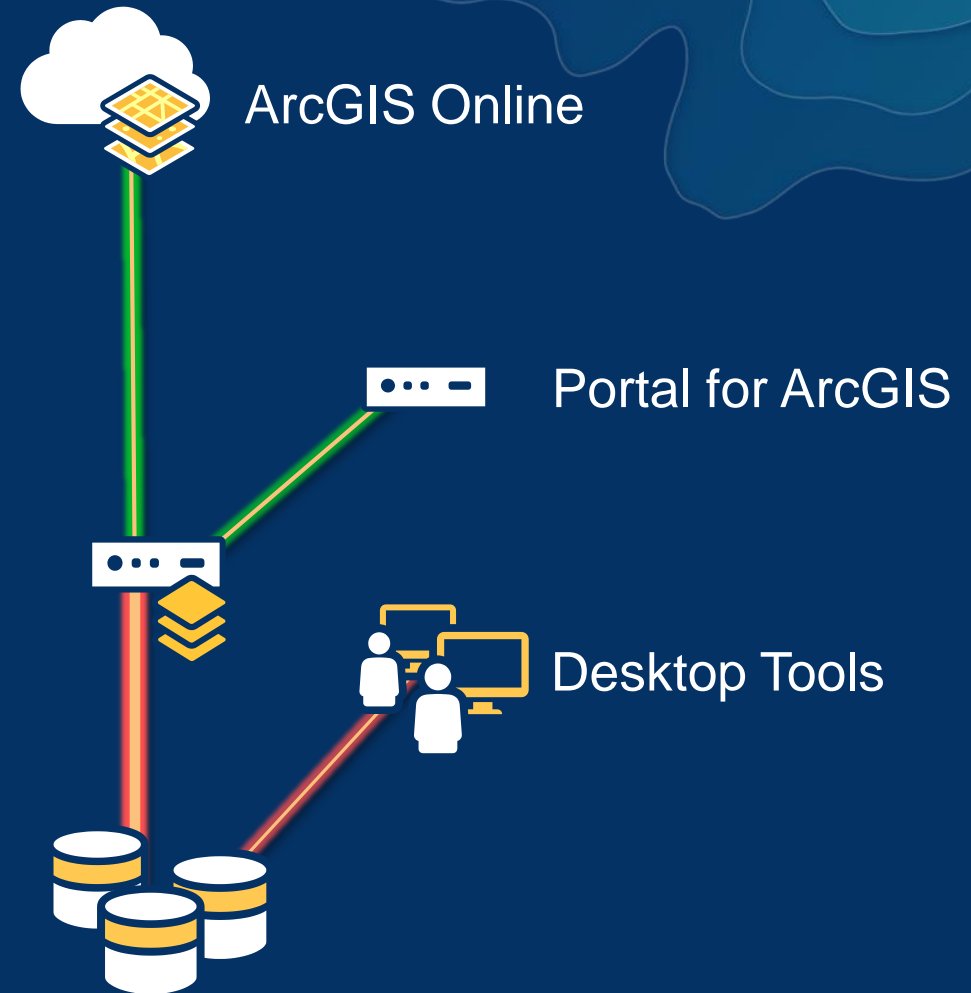
Component Location is Critical

Services

Services-based communication best for configuration prone to latency


SQL and File access

Components communicating via SQL or file-based access must be as close as possible to reduce effects of latency.



Example 3: Resolution - The Disappointed GIS Analyst

- Ensure ArcGIS Server, Desktop and other direct connections to database are close to database server.
- Provide Virtual Desktop (e.g. VMware Horizon, Citrix, MS Terminal Services) - VDI



Example 4: The “It is Slow 50% of the Time” System

- Editors said performance was ok half the time
- What about time of day?
- Are maintenance practices being followed?
 - Not as much in certain scenarios
 - Or a cold backup done at night without shutting down ArcGIS Server services or managing connections




Citrix ArcGIS
Desktop Editing

Enterprise
Geodatabase


Example 4: Resolution - It is Slow 50% of the Time

- There were two data centers, but only one had the Enterprise Geodatabases
- Citrix based ArcGIS Desktops were deployed at both.
- 50% of time editor might get desktop at remote data center



Example 5: The Slow Afternoon

- Systems got slow for users as it got late in the day
- What are the workflows?
- What are the maintenance practices for those?
 - Any conflicting?
 - What time of day?
- Large numbers of editors (100+) doing versioned editing concurrently on complex data
- Also field editing taking place
- Time of day did matter here...



Example 5: Resolution - The Slow Afternoon

- Review of versioned editing management tables State_Lineages table growing very rapidly due to various editing workflows (tens of thousands of rows per hour).
- Geodatabase compress operations were being done nightly
- Editors and field crews did take a lunch break mid-day, allowing for a mid-day compress to take place along with database statistics update.

Example 6: The Slow Data Maintenance Operations

- Geodatabase Administrator stated the Reconcile / Post / Compress geoprocessing operations gets slower over time
- About 50,000 edits performed daily
- Isolate and review one operation – Reconcile
- What is the work performed?
- Set up batch script to run with no edits made
- Capture SDEIntercept
- Review log (1st run VS Nth run)



Enterprise Geodatabase

Example 6: Resolution – The Slow Data Maintenance Operations

- Geoprocessing logging history in metadata
- Viewing the logging history

ArcCatalog 10.x: Customize > ArcMap Options window > Metadata tab

Select FGDC style

- ArcGIS 10.x (Default - ON) : ArcGIS PRO (Default – OFF)
- Scripting - arcpy.SetLogHistory(False)
- Resources : Delete history (as desired)

Clear GP History Add-in

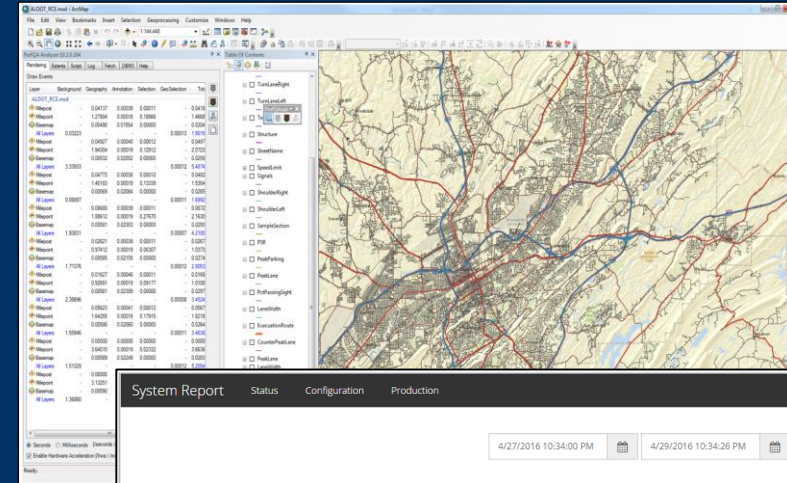
How To: Automate the process of deleting geoprocessing history

Troubleshooting – Some Typical Scenarios

- **Inefficient application configuration**
 - Overly complex map document, too many layers, complex symbology, etc..
- **Inefficient database maintenance**
 - Geodatabase compress, statistics, index maintenance
 - Compress_log – does it exist?
 - State_lineages – 10's of millions of rows?
- **Inefficient architecture**
 - Large physical distance between Database and “thick” clients – Desktop and Server
- **Resource Configuration Issue**
- **Security Configuration**
 - Is web sense or other web filtering present?

Troubleshooting - Tools

- ArcGIS PerfQAnalyzer and MxdPerfstat
- ArcGIS System Log Parser
- Esri System Monitor
- OS performance metrics and logs
- Database Specific
 - SDEIntercept – communication between ArcGIS and Database
 - RDBMS Database Traces – SQL and execution plans within database



```
SQL ID: cfsranagv70vh  
Plan Hash: 2103630758  
SELECT MJRRD.OBJECTID  
FROM  
GDB.MJRRD, GDB.CNTY  
1 and CNTY.NAME = 'M'
```

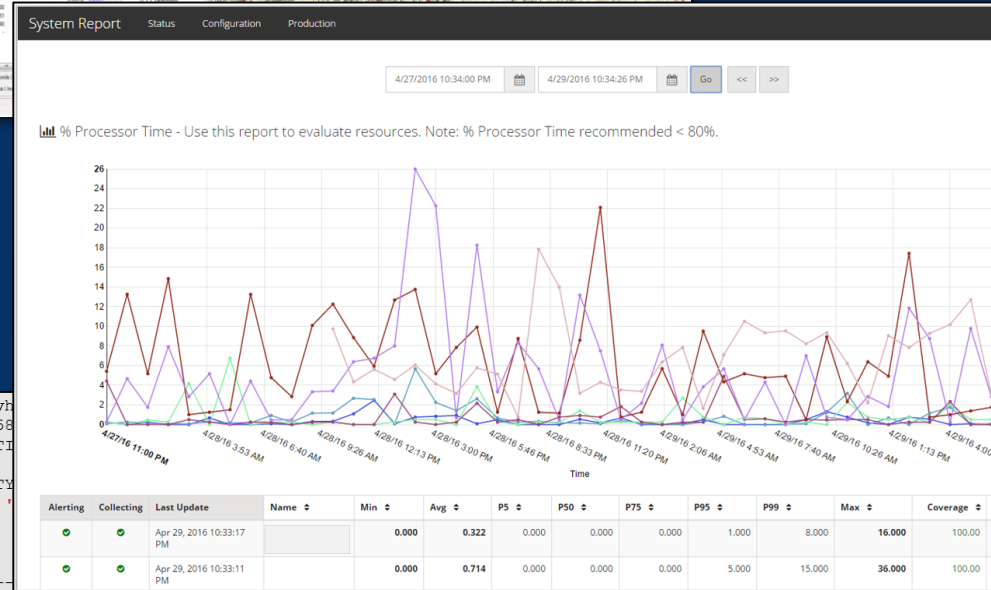
call	count
------	-------

Parse	1	0.01	0.02	0	222	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	22.76	23.07	34322	34953	0	48
total	3	22.77	23.09	34322	35175	0	48

Misses in library cache during parse: 1
Optimizer mode: ALL_ROWS
Parsing user id: 218

Rows	Row Source Operation
------	----------------------

48	NESTED LOOPS (cr=34953 pr=34322 pw=0 time=1598 us cost=18864 size=1546 card=1)
1	TABLE ACCESS FULL CNTY (cr=626 pr=0 pw=0 time=0 us cost=170 size=2302 card=2)
48	TABLE ACCESS FULL MJRRD (cr=34327 pr=34322 pw=0 time=1457 us cost=9347 size=395 card=1)



Troubleshooting – Esri Support

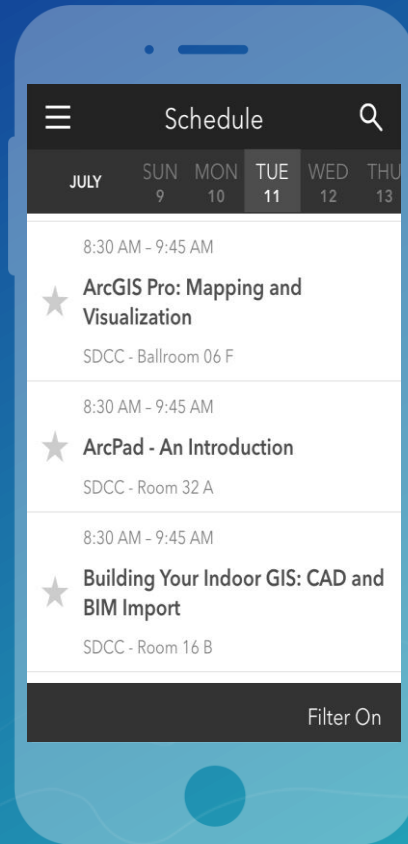
- Review behavior on a test instance (if applicable)
- Isolate behavior
- Generate SDEIntercept / DBMS trace (single action)
- Open a Tech Support case
- A database backup might be requested
- More...

Please Take Our Survey on the Esri Events App!


Download the Esri Events app and find your event




Select the session you attended



Scroll down to find the survey



Complete Answers and Select "Submit"



Q & A

Note: We will take as many questions as possible during this time, but note that some questions may be better suited to post workshop session discussion and follow up at Technical Support Island.

We will be available after the session for this.



esri

THE
SCIENCE
OF
WHERE