



Rapid Deployment of EPA Enterprise Geospatial Data Warehouse in Katrina's Aftermath

ESRI 2006 User Conference

May 18th, 2006

Torrin Hultgren & Qi Dai

EPA GIS Support Team

Research Triangle Park, North Carolina



Abstract

In the aftermath of Hurricane Katrina, the EPA had an urgent need for current geographic data on everything from known hazardous waste sites, flood conditions, streets and building footprints, and high-resolution satellite imagery. To respond to this need and make the data accessible to EPA personnel in dispersed locations, a high-availability ArcSDE/Oracle data warehouse and ArcIMS server were rapidly deployed in the EPA's National Computing Center in Research Triangle Park, North Carolina. Triage was performed on the terabytes of commercial imagery the NGA made available to federal agencies, and ArcGIS Spatial Analyst and ArcToolbox's Modelbuilder proved equal to the task of rapidly projecting, stretching and mosaicking large volumes of imagery.



EPA Emergency Response Role

- ❖ Field Testing – measuring levels of toxins and contaminants in floodwaters
- ❖ Mitigate risk to critical facilities and sites
- ❖ Contain larger spills when possible
- ❖ Coordinate safe debris removal and disposal
- ❖ Monitor air quality levels at debris burn sites
- ❖ GIS task: rapidly create maps and perform spatial analyses to support field engineers and decision-makers



EPA GIS Community

❖ Diverse Locations:

- ❖ Region 5 – Louisiana
- ❖ Region 4 – Mississippi and Alabama
- ❖ Headquarters EOC – Washington, DC

❖ Mix of EPA Staff and Contractor Support

❖ Mix of GIS Experts and EOC GIS Novices

❖ Most located on the secure, high-speed EPA internal network



EPA Data Needs

❖ EPA Regulated Facilities Data

- ❖ Already in SDE database

❖ Basemap data

- ❖ As feature services
- ❖ From various sources

❖ High-resolution post-Katrina imagery

- ❖ Demand for imagery was intense
- ❖ Response was mandate to NGA to purchase and disseminate as much imagery as possible to responders

Example Feature Layers

Airports	Landfills
Area Codes	Lots
Buildings	Major Institutions
Census Blocks	Major Roads
Cemeteries	Neighborhoods
Curblines	Parcels
Debris Burn Locations	Parish Boundaries
Drainage Lines	Parks
Evacuation Routes	Police Dept Zones
Evacuation Zones	Police Stations
Federal Courts	Populated Places
Federal Lands	Railroads
FEMA Zones	Recreation Areas
Fire Districts	Schools
Fire Stations	Sewer Lines
Fire Zones	Sewer Nodes
Flood Extents	State Boundaries
Golf Courses	Street Centerlines
Hospitals	Topo Map Indices
Hurricane Swath	Tracts
Hydrography	Truck Routes
Interstates	Urban Areas
Address Points	Water Bodies
Lakes	Zip Codes
Land Use	



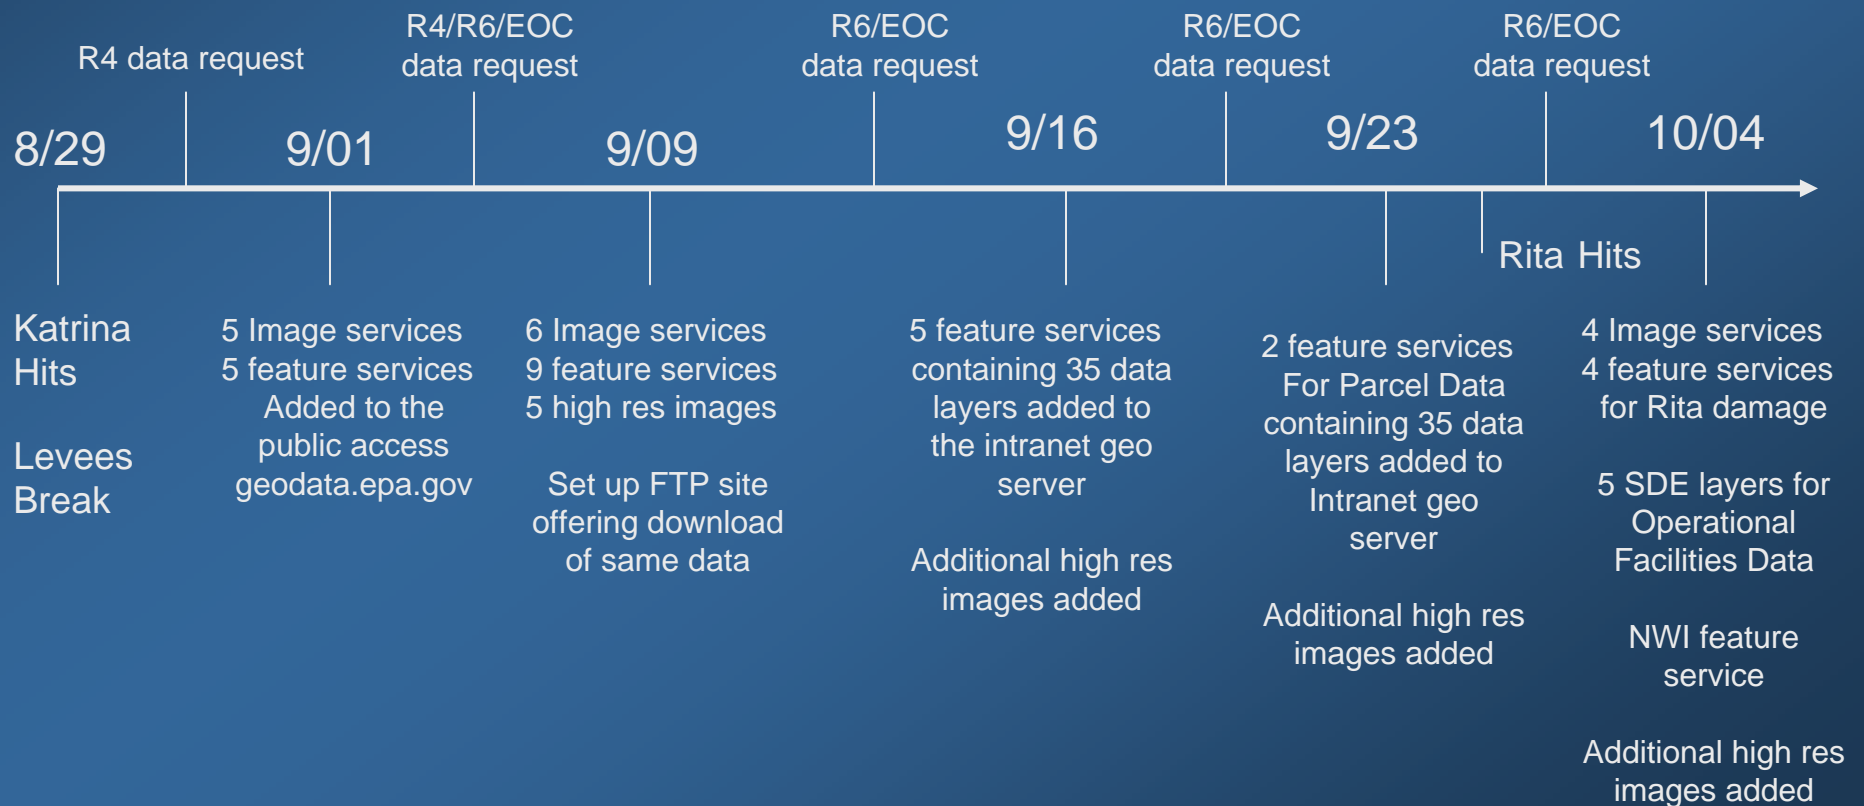
EPA National Computing Center

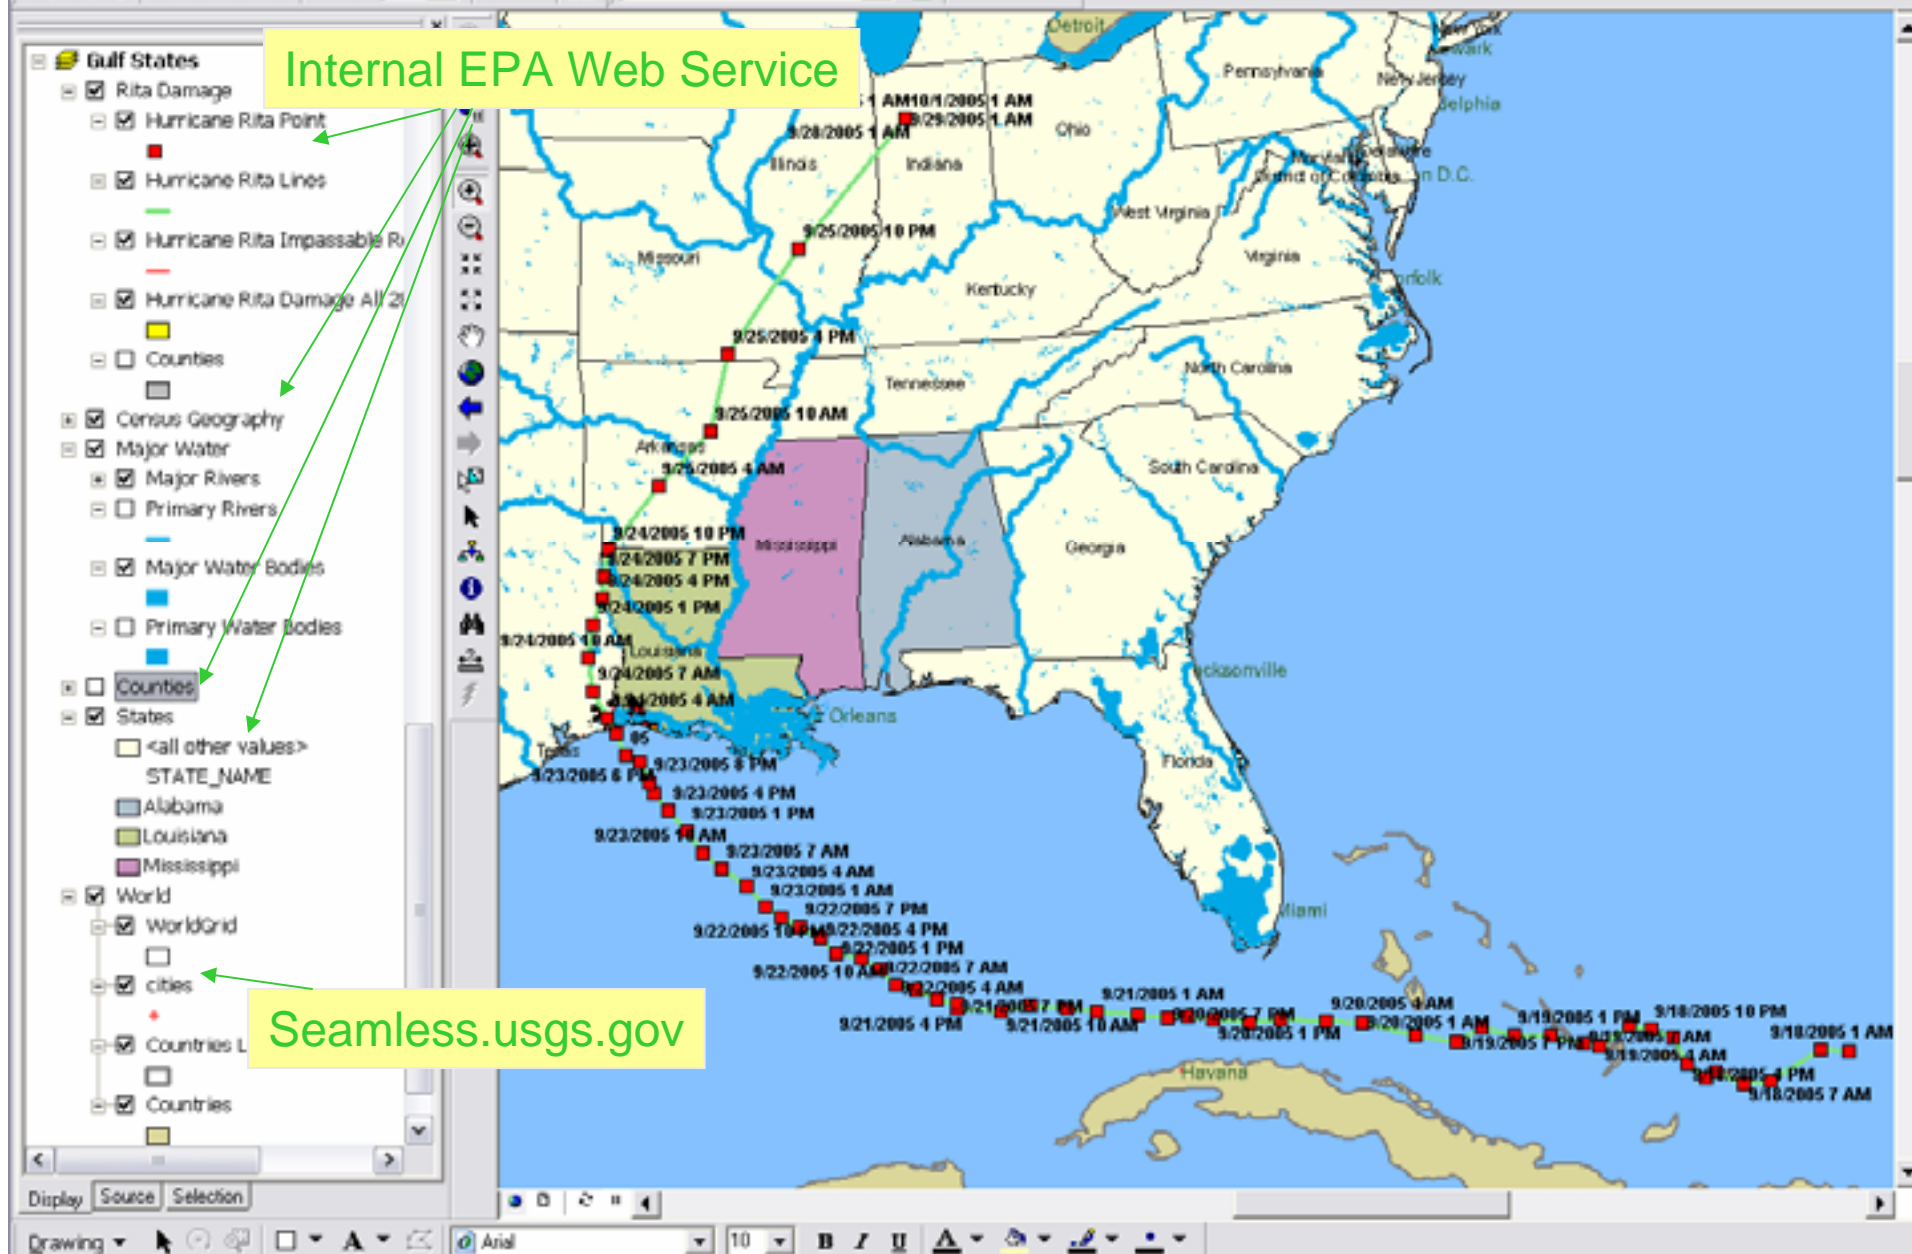
- ❖ Located in Research Triangle Park, NC
- ❖ Industrial strength data center
- ❖ Enterprise computing resources available for all EPA program offices, Research labs, and Regions
- ❖ Official host of all EPA GIS applications and databases for public access
- ❖ Official host of most EPA national GIS applications and databases for intranet access
- ❖ Agency-wide network resource hub and the gateway to the Internet
- ❖ 7x24 operation, backup and restore services, and emergency response
- ❖ EPA GIS group – responsible for public web mapping applications, internal geospatial databases, and general GIS and ESRI desktop support

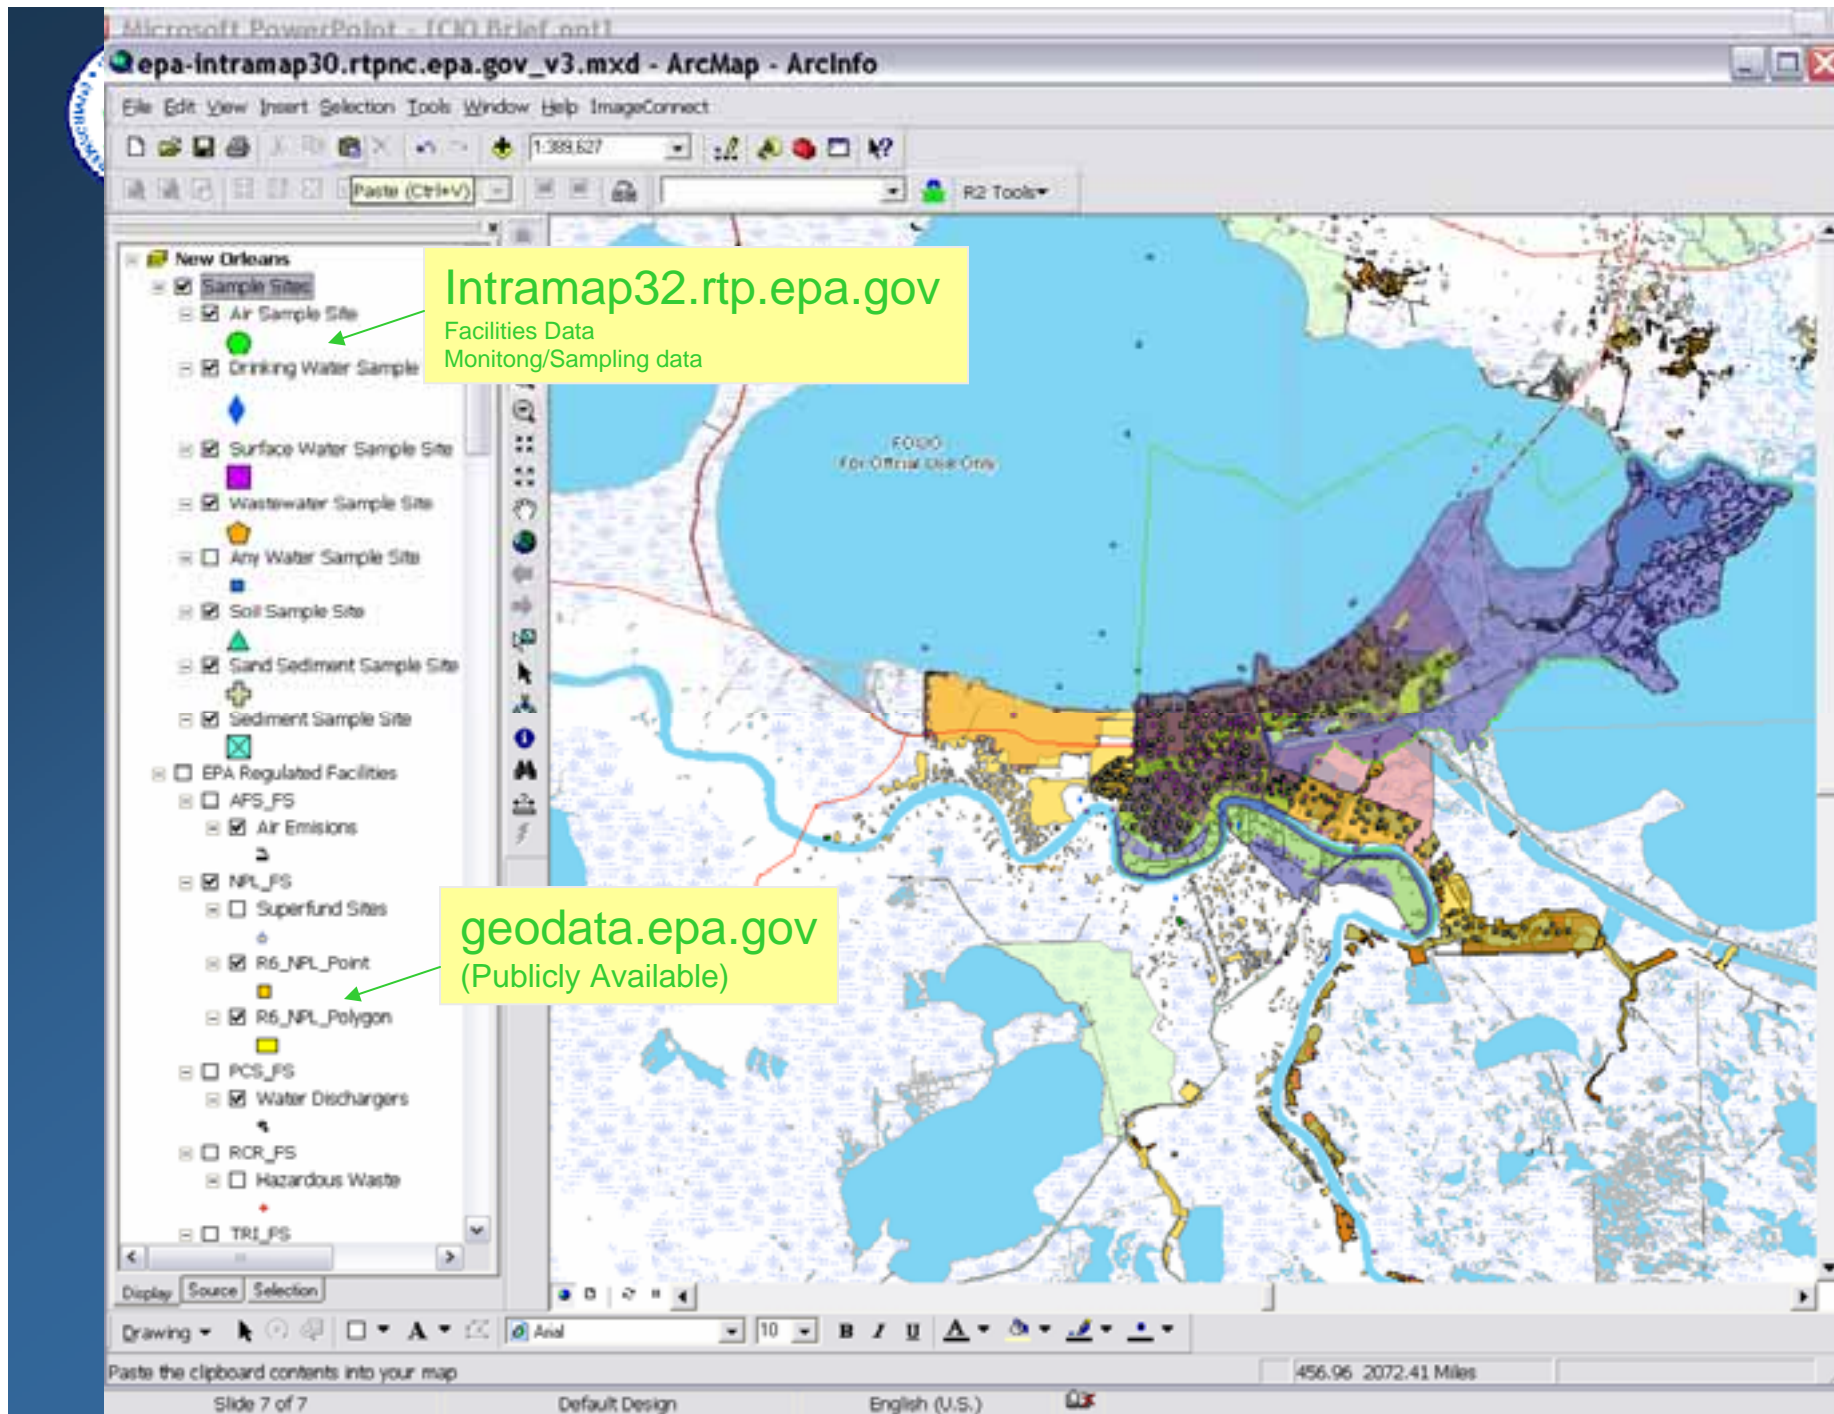



Katrina Response Timeline

- ❖ Created new SDE instance on existing Oracle DB server on day of request
- ❖ Created load-balanced ArcIMS application server configuration the next day









EPA Feature Service

Examples – Local Data

☒ New Orleans City Data

☐ FEDERAL_COURTS



☐ FIRE_STATIONS

☐ HOSPITALS



☐ POLICE_STATIONS



☐ SCHOOLS



☐ SEWER_NODES



☐ WATER_NODES



☐ BUILDINGS

☐ CURBLINES



☐ EVACUATION_ROUTES



☐ INTERSTATES



☐ MAJOR_ROADS



☐ TRUCK_ROUTES



☐ STREET_CENTERLINES



☐ RAILROADS



☐ SEWER_LINES



☐ WATER_LINES



☐ EVACUATION_ZONES

☒ FEMA_FLOOD_ZONES

GSWS.FEMA_ZONES.FLD



100 YR FLOOD

☐ FIRE_ZONES

☐ FIRE_DISTRICTS

☐ NOPD_ZONES

☐ NOPD_SUBZONES



☐ LOTS



☐ PARCELS



☒ NEIGHBORHOODS



☐ VIEUX_CARRE



☒ PARISH_BOUNDARY



☐ LANDUSE

☐ HYDRO



☐ WATER



☐ Jefferson Parrish

Internal Web Services

-Hi-resolution local data

-New Orleans detailed transportation, parcel, administrative information

EPA Feature Service

Examples:

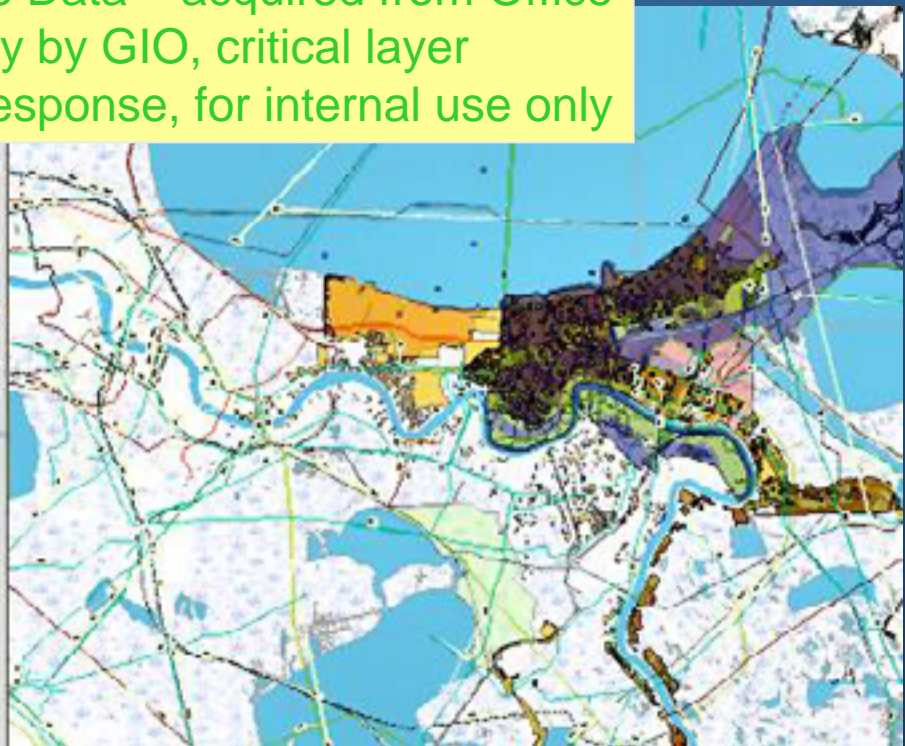
Nationally Consistent Data Services

Internal EPA Web Services

- Broad range of data from GDT national subscription
- National Pipeline Data – acquired from Office of Pipeline Safety by GIO, critical layer for emergency response, for internal use only

- ☒ Transportation
 - ☒ Interstate Highways
 - ☐ Major Highways
 - ☐ Highways
 - ☐ Railroads
 - ☐ LA_GDT_Streets
 - ☐ MS_GDT_Streets
 - ☐ AL_GDT_Streets
 - ☐ Airports
 - ☐ Counties
 - ☐ National Pipeline
 - Other
 - Crude Oil
 - Highly Volatile Liquid (HVL)
 - Liquefied Petroleum Gas (LP)
 - Natural Gas
 - Natural Gas Liquids (NGL)

- ☐ Natural Gas Liquids (NGL)
- ☒ USA Basemap
 - ☐ Schools
 - ☐ Hospitals
 - ☐ Public Buildings
 - ☐ Institutions
 - ☐ Places
 - ☐ Landmarks
 - ☐ Recreation Areas
 - ☐ Golf Courses
 - ☐ Cemeteries
 - ☐ Uninhabited Places
 - ☐ Parks
 - ☐ Federal Lands
 - ☐ River Channels, Floodplains, and Great Lakes





Imagery: Need for Triage

- ❖ The volume of imagery was in terabytes
- ❖ Imagery was provided by many sources – USGS, FEMA, NOAA, USACE, NGA, commercial providers, various universities
- ❖ The same images were posted online in many places, and most sites were updated daily
- ❖ Many websites had no ability to preview images prior to download
- ❖ Field GIS personnel responsible for making maps “five minutes ago”
 - ❖ Lacked bandwidth to handle “firehose” of downloadable imagery
 - ❖ Requested triage assistance from EPA RTP GIS Group



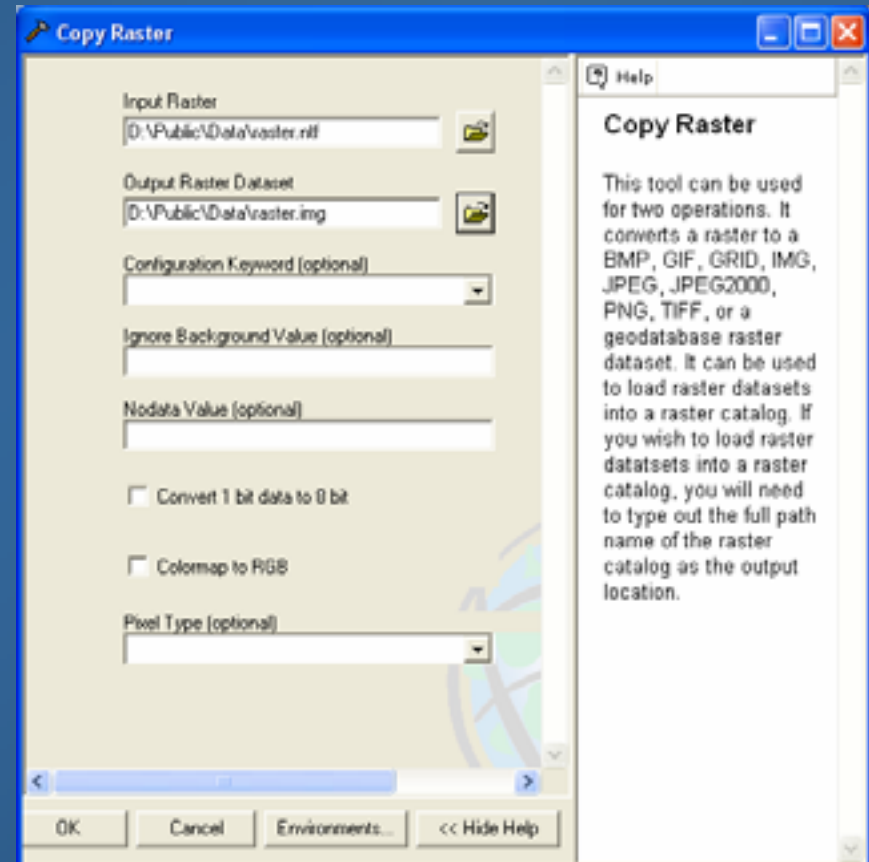
Response to Need for Triage

- ❖ RTP Staff monitored all imagery sites and maintained a list of potentially useful cloud-free images
- ❖ RTP Staff processed and republished restricted NGA imagery as ArcIMS Image Services
- ❖ Kept responders informed of updates to imagery via daily conference calls



Geoprocessing: *Format Conversion*

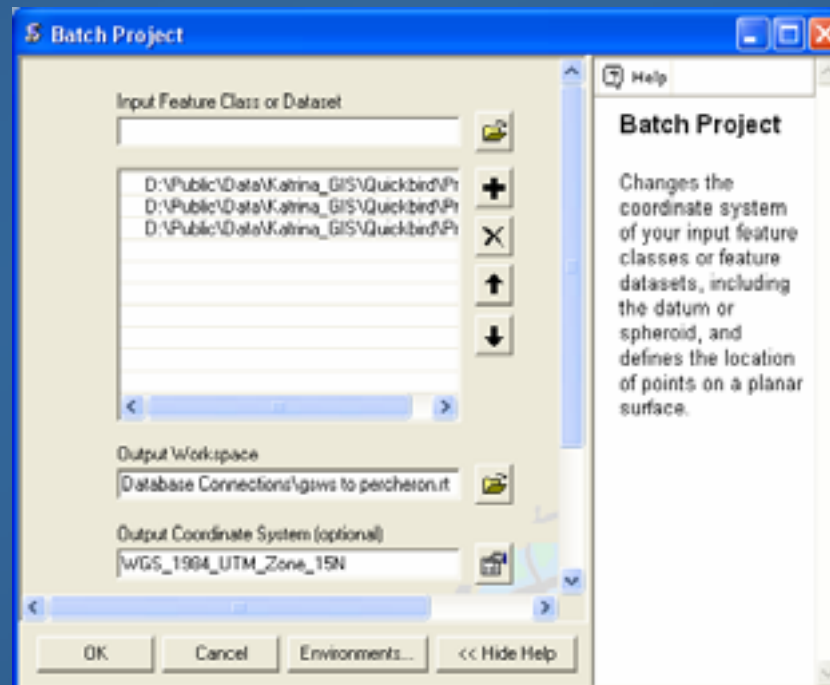
- ❖ Raw, unrectified imagery posted as National Imagery Transmission Format (NITF) files
- ❖ Support in ENVI 4 and GeoXpress 4.0 was poor to nonexistent, but ArcGIS can recognize and process NITF reliably
- ❖ Copy Raster tool used for loading into SDE and conversion to MrSID compressible format





Geoprocessing: Reprojection

- ❖ New Orleans is divided by two UTM zones, imagery arrived in both projections
- ❖ ArcGIS can project images on-the-fly but all images in a SDE mosaic or Raster Catalog* must have the same projection





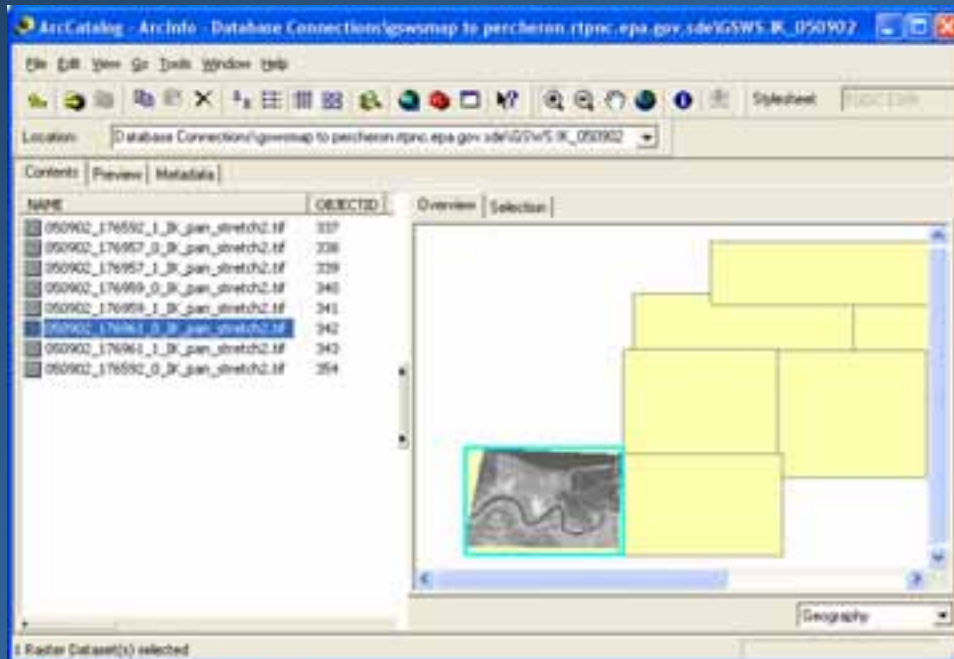
Geoprocessing: Stretching

- ❖ When viewing a standalone or SDE Raster with ArcCatalog or ArcMap, a 2 standard deviation stretch is applied by default
- ❖ When viewing an ArcIMS image service, no stretch can be applied
 - ❖ Raster must be saved with stretched values
- ❖ Stretch image in ArcMap, then use Reclassify from Spatial Analyst Toolbar (not ArcToolbox)
 - ❖ Current display lookup table loads by default
 - ❖ Can save table for manual editing or repeat use
 - ❖ Also useful for converting background values to nodata
 - ❖ Reduces file size by decreasing bit-depth



Geoprocessing: Mosaicking

❖ ArcSDE has two options for raster storage:



❖ Single mosaicked raster

- One set of pyramids, leading to faster display at smallest scales

❖ Raster catalog

- Eases management as individual rasters can be added and removed without affecting others

❖ Raster Catalogs are improving in 9.2!



Geoprocessing: Pyramids

- ❖ Can be built automatically during copy/upload process
 - ❖ Specify in the geoprocessing environment settings
 - ❖ Pay attention to pyramid levels, not just resampling technique
- ❖ Pyramids will occasionally be corrupted – stripes or missing levels
 - ❖ Deleting and rebuilding won't always fix
 - ❖ Must choose different resampling technique

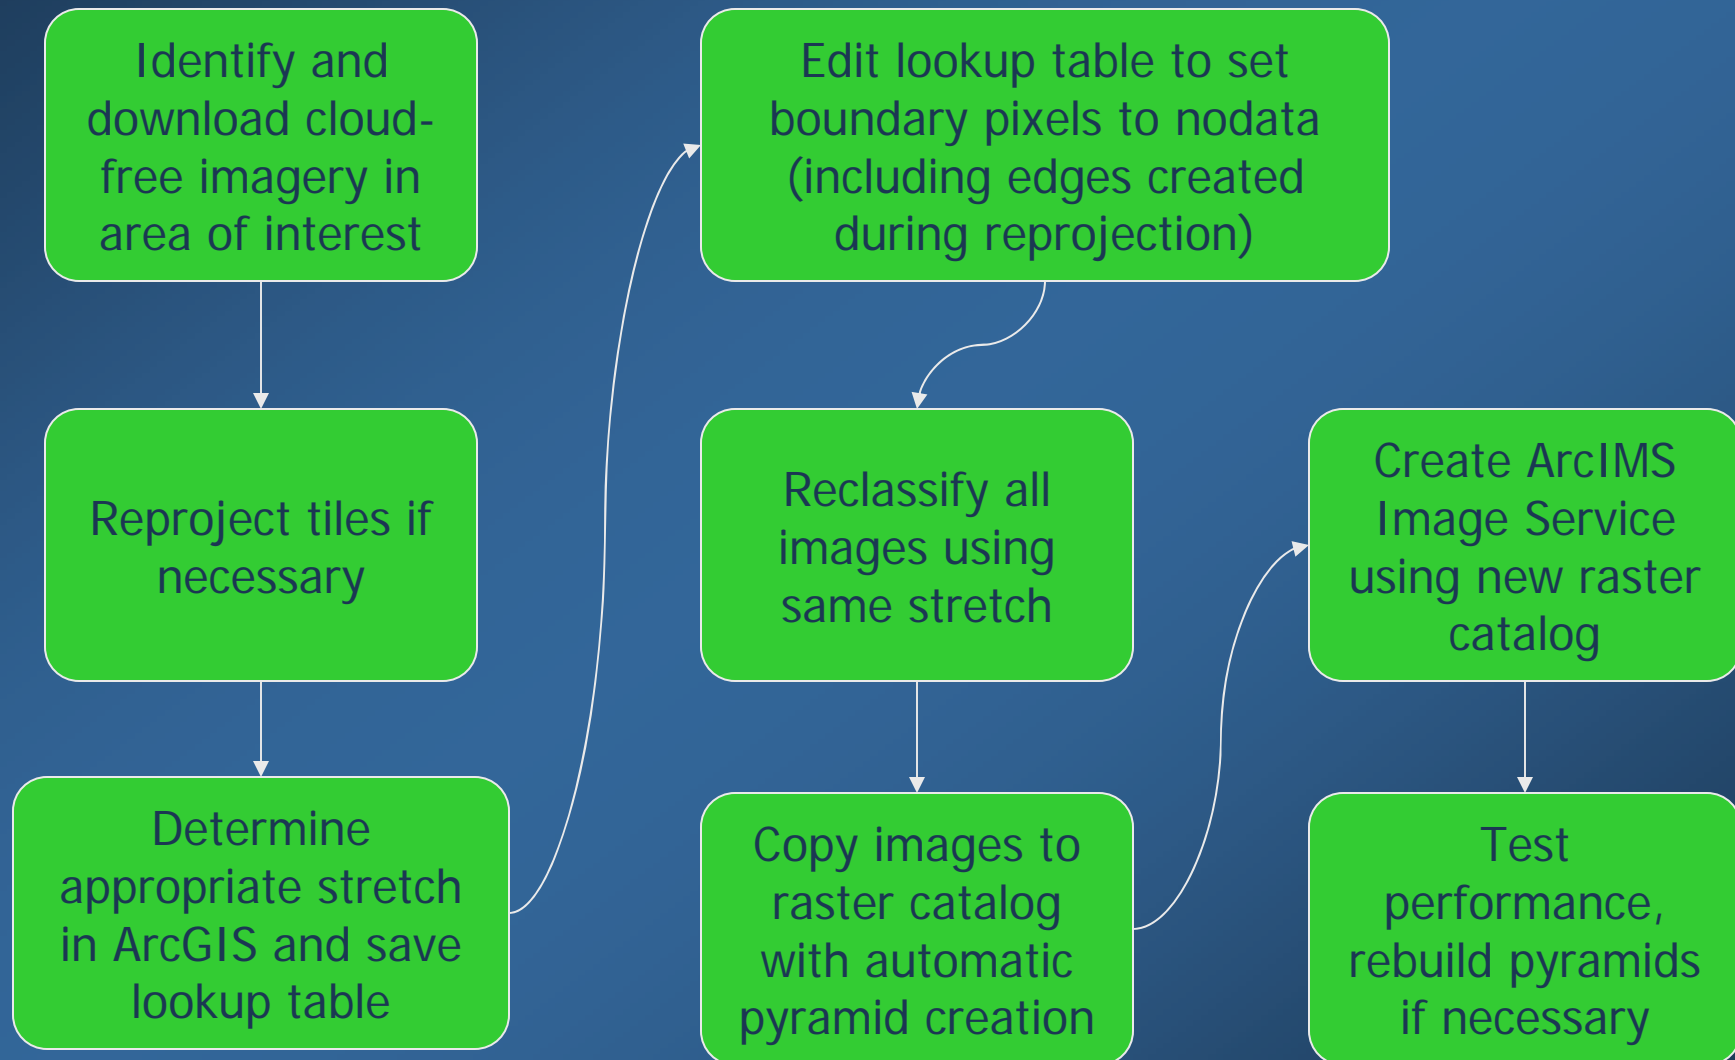


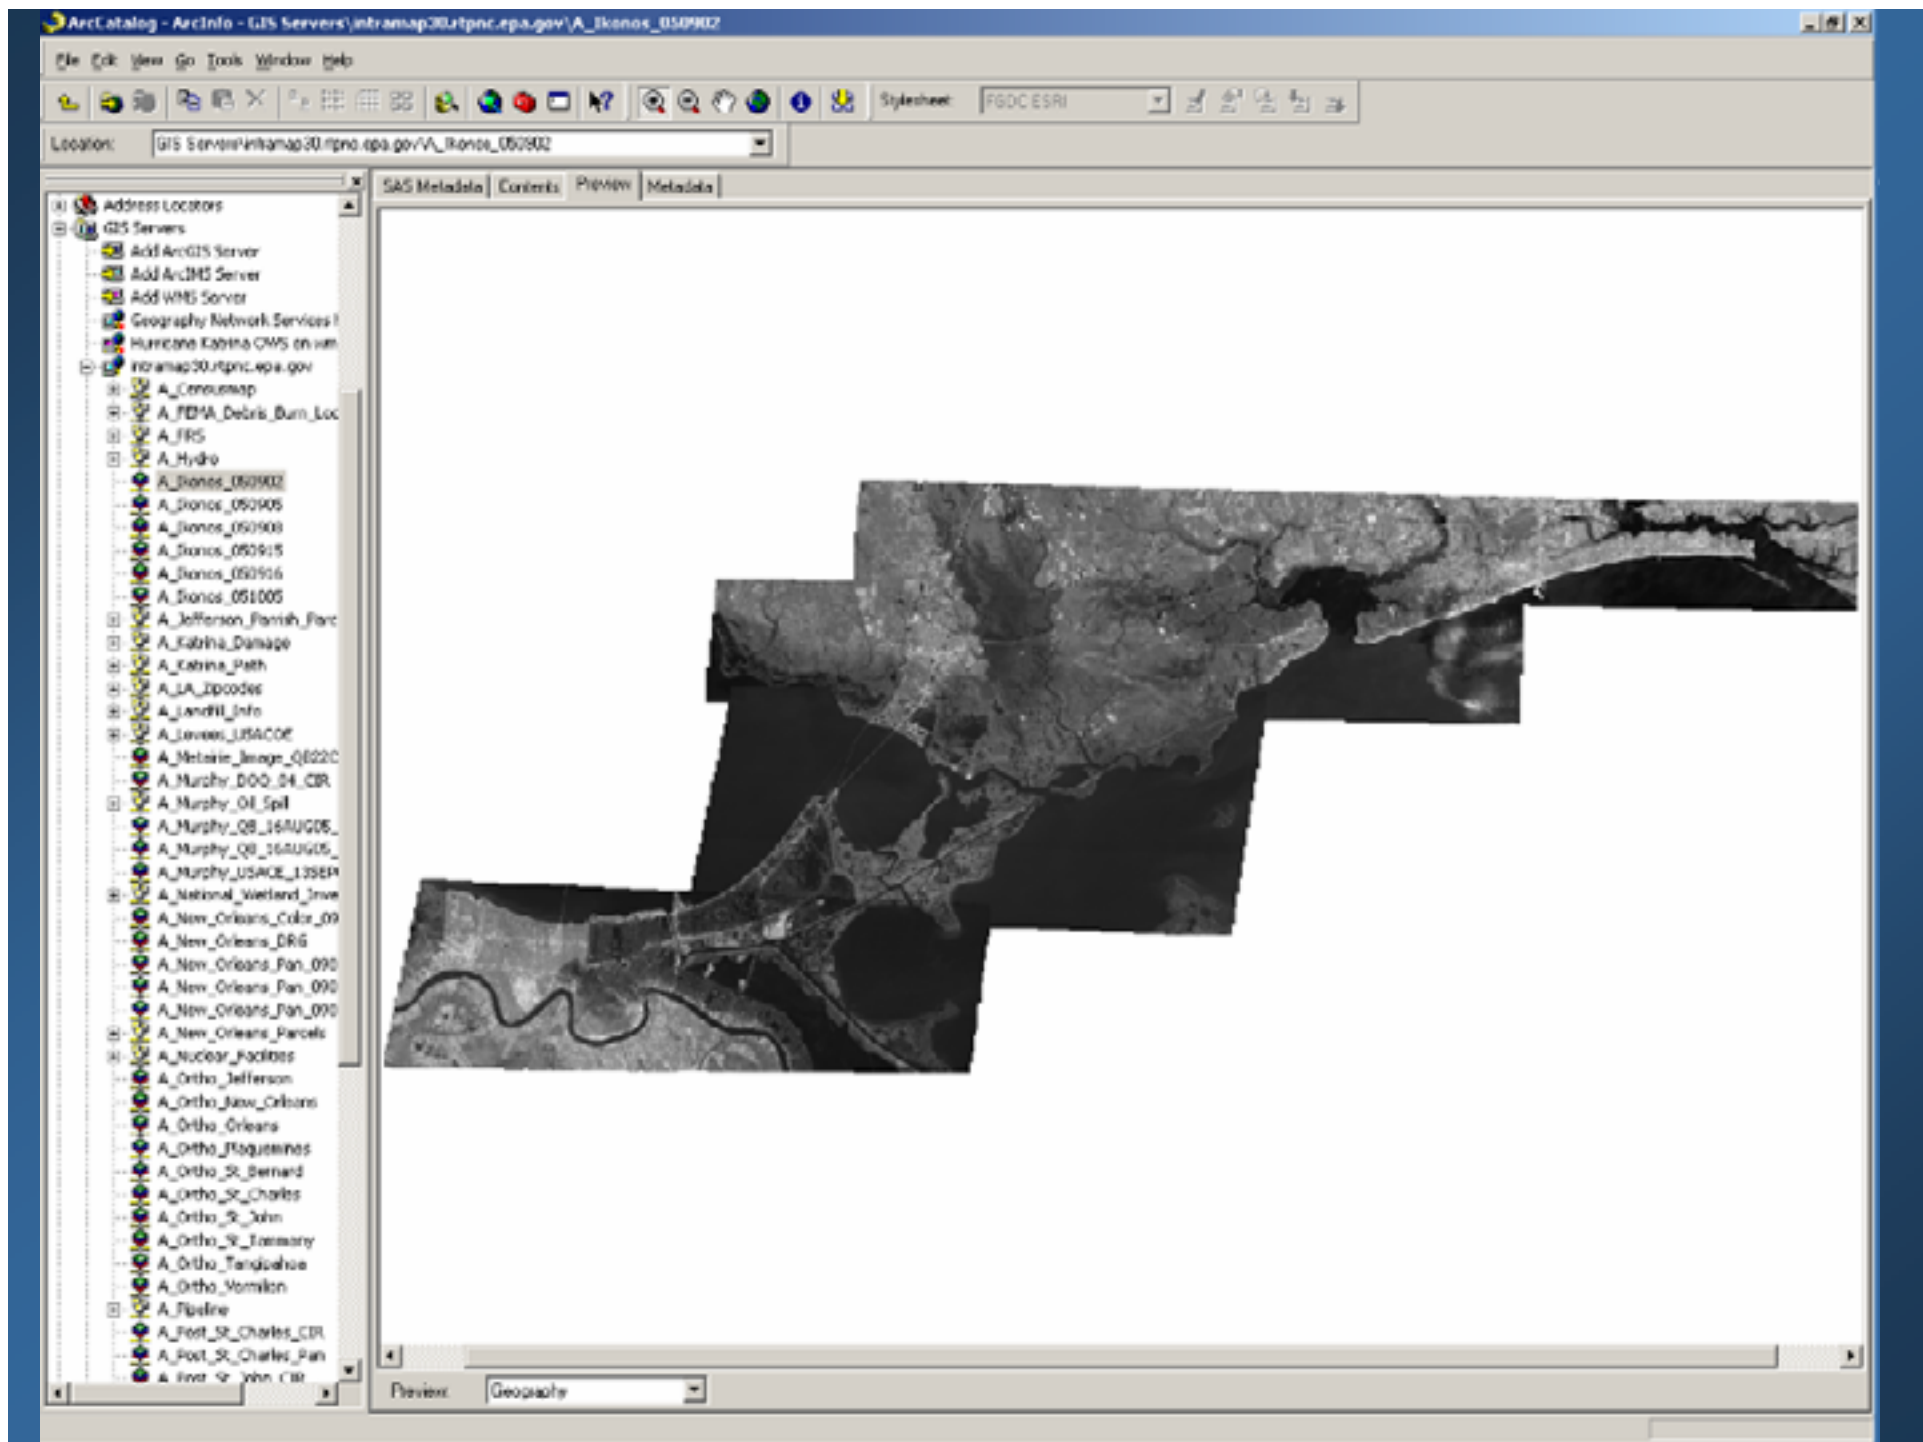
Geoprocessing: Batch Processes

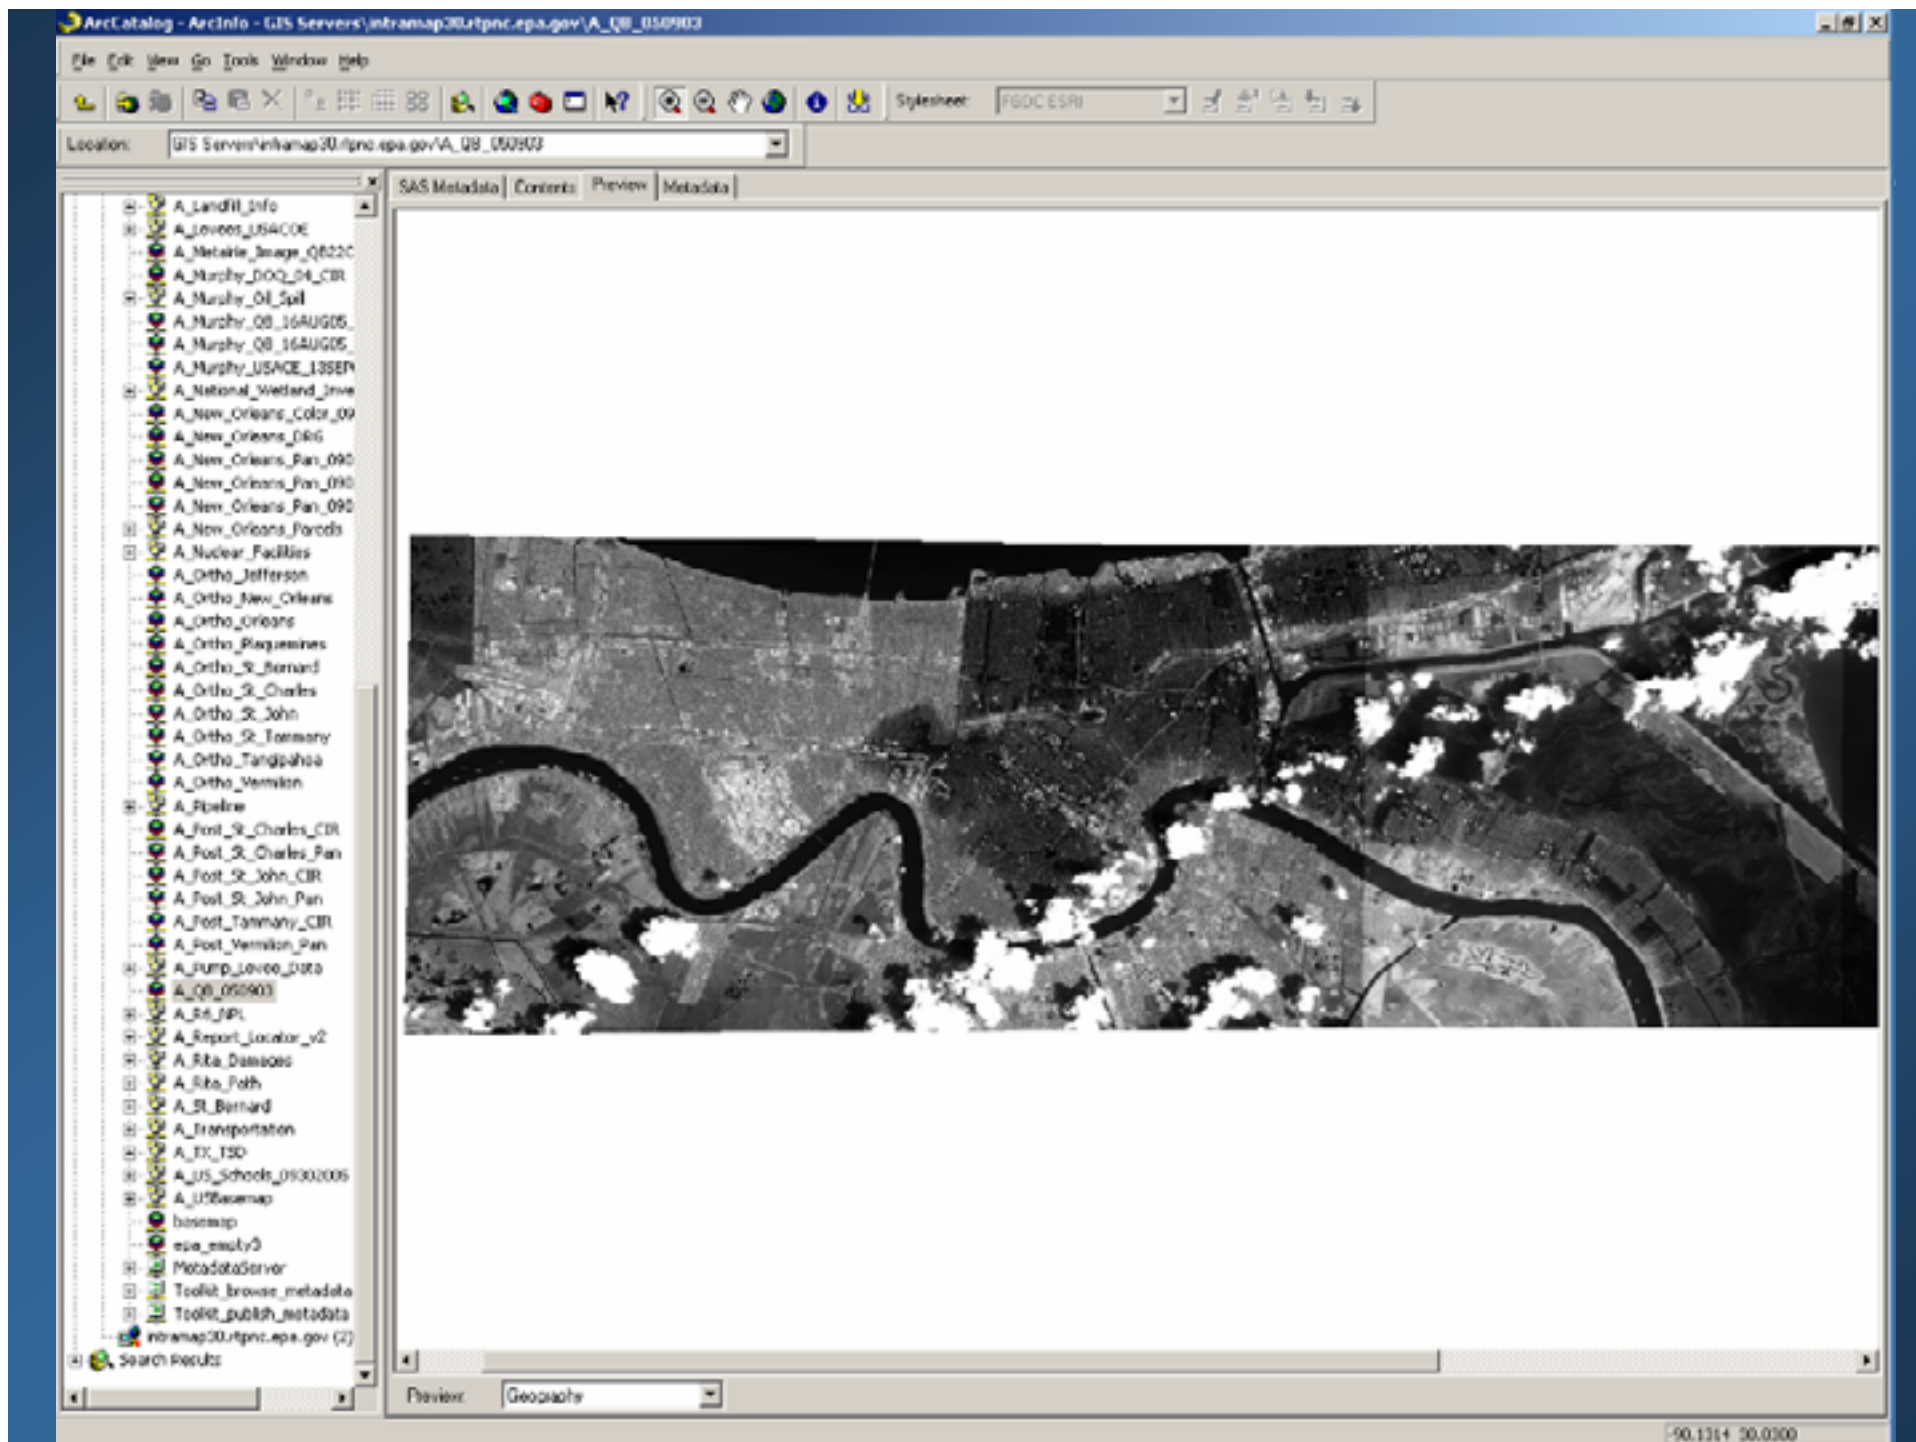
- ❖ Quick and dirty batch processing is easiest using the ArcCatalog Command Line
 - ❖ Copy and paste commands from text file
- ❖ ModelBuilder better for polished processes
 - ❖ Allows a nice GUI dialog box for setting input parameters
- ❖ When dealing with very large images in crisis response mode, quick and dirty is fine

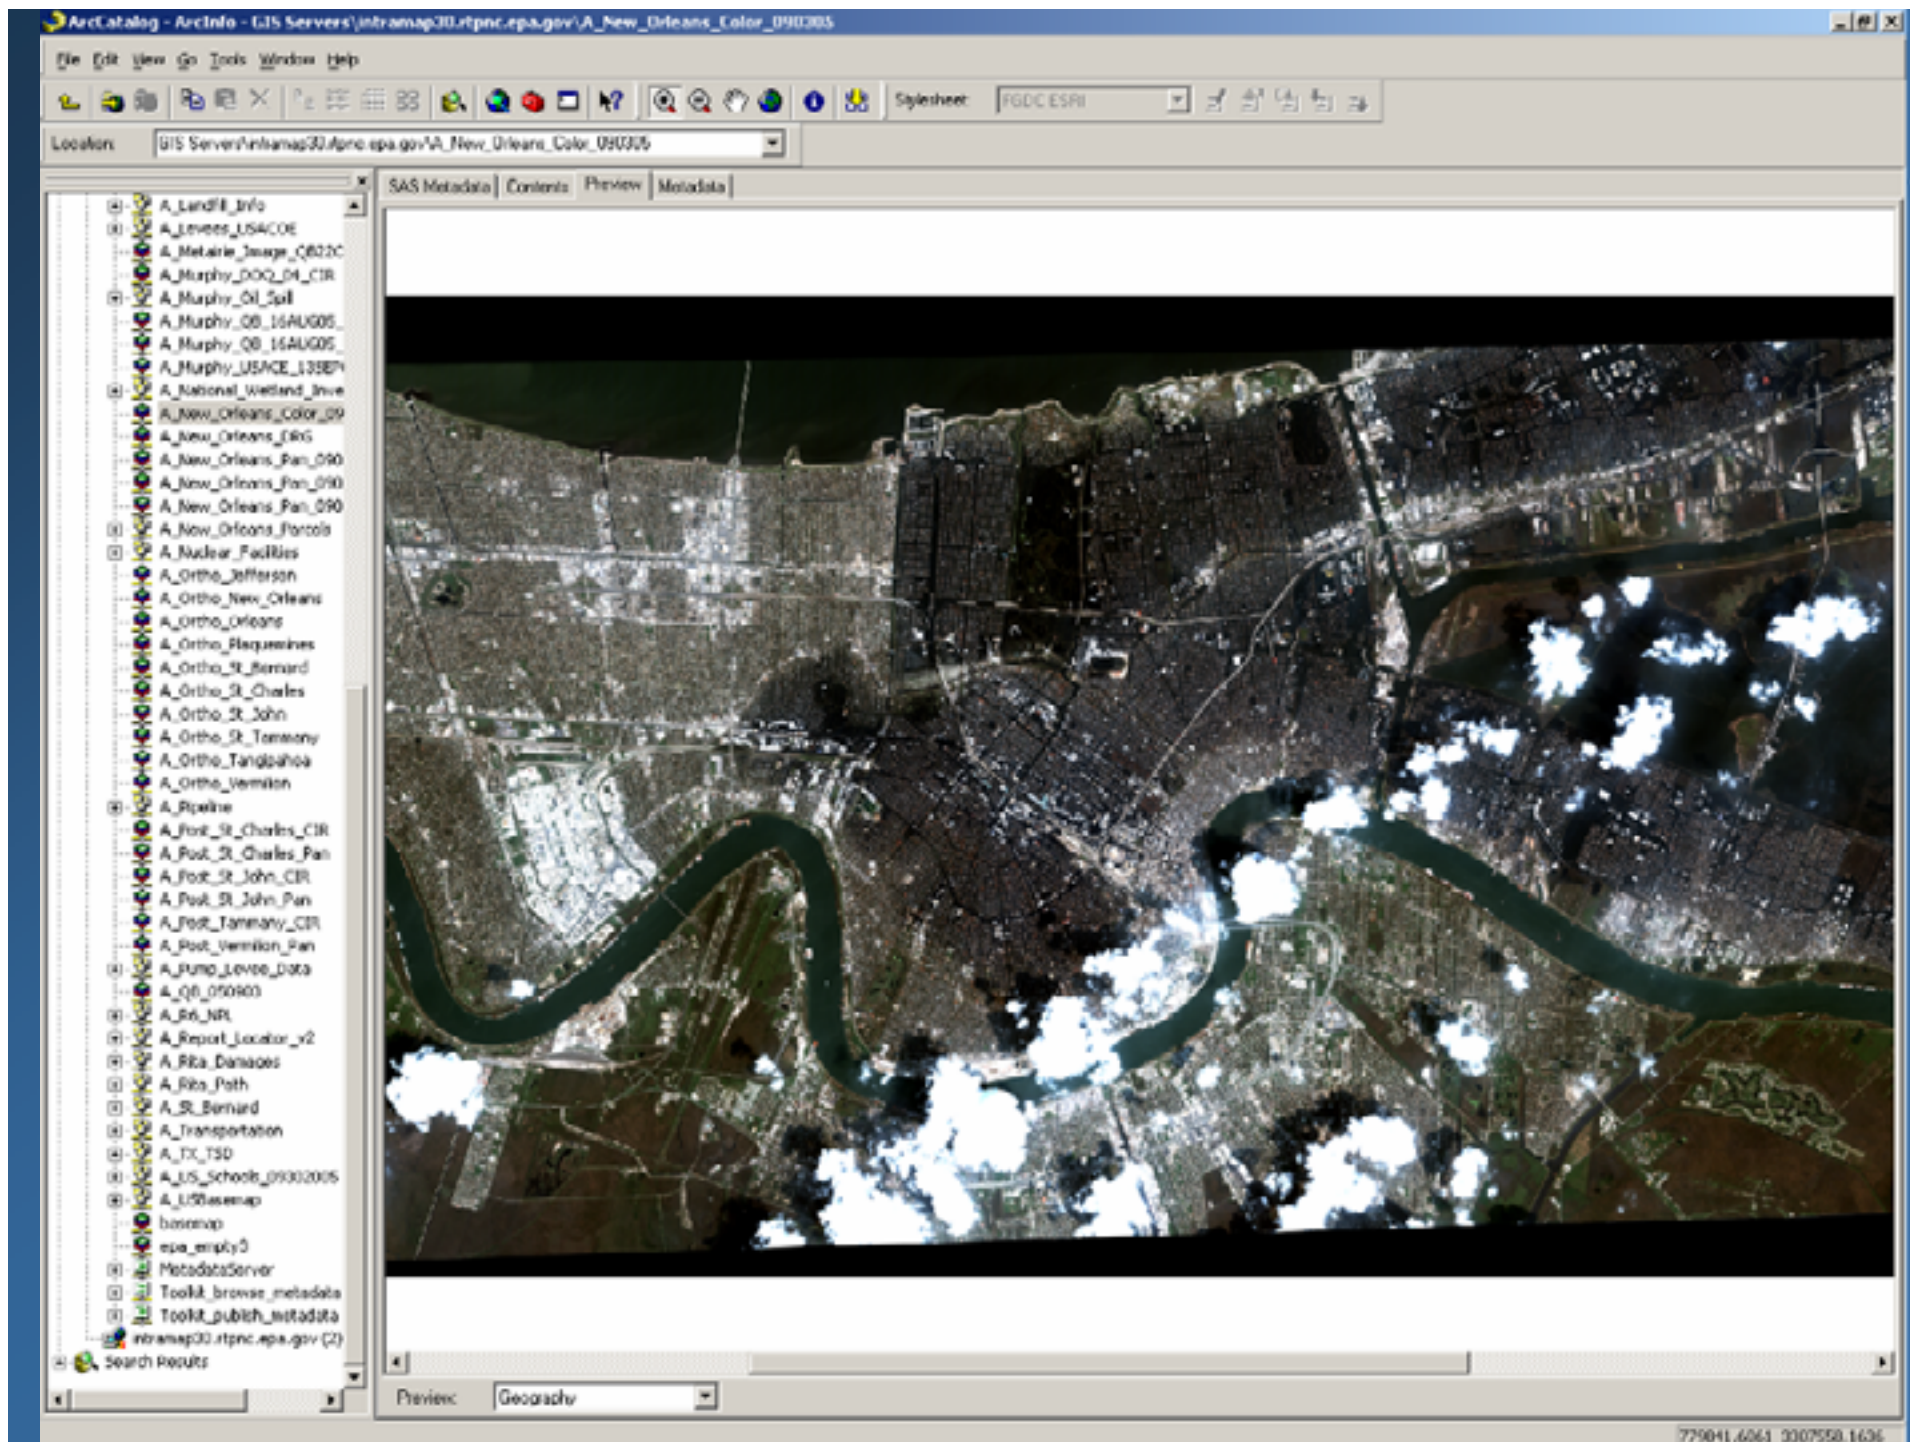


Workflow: Download to Service






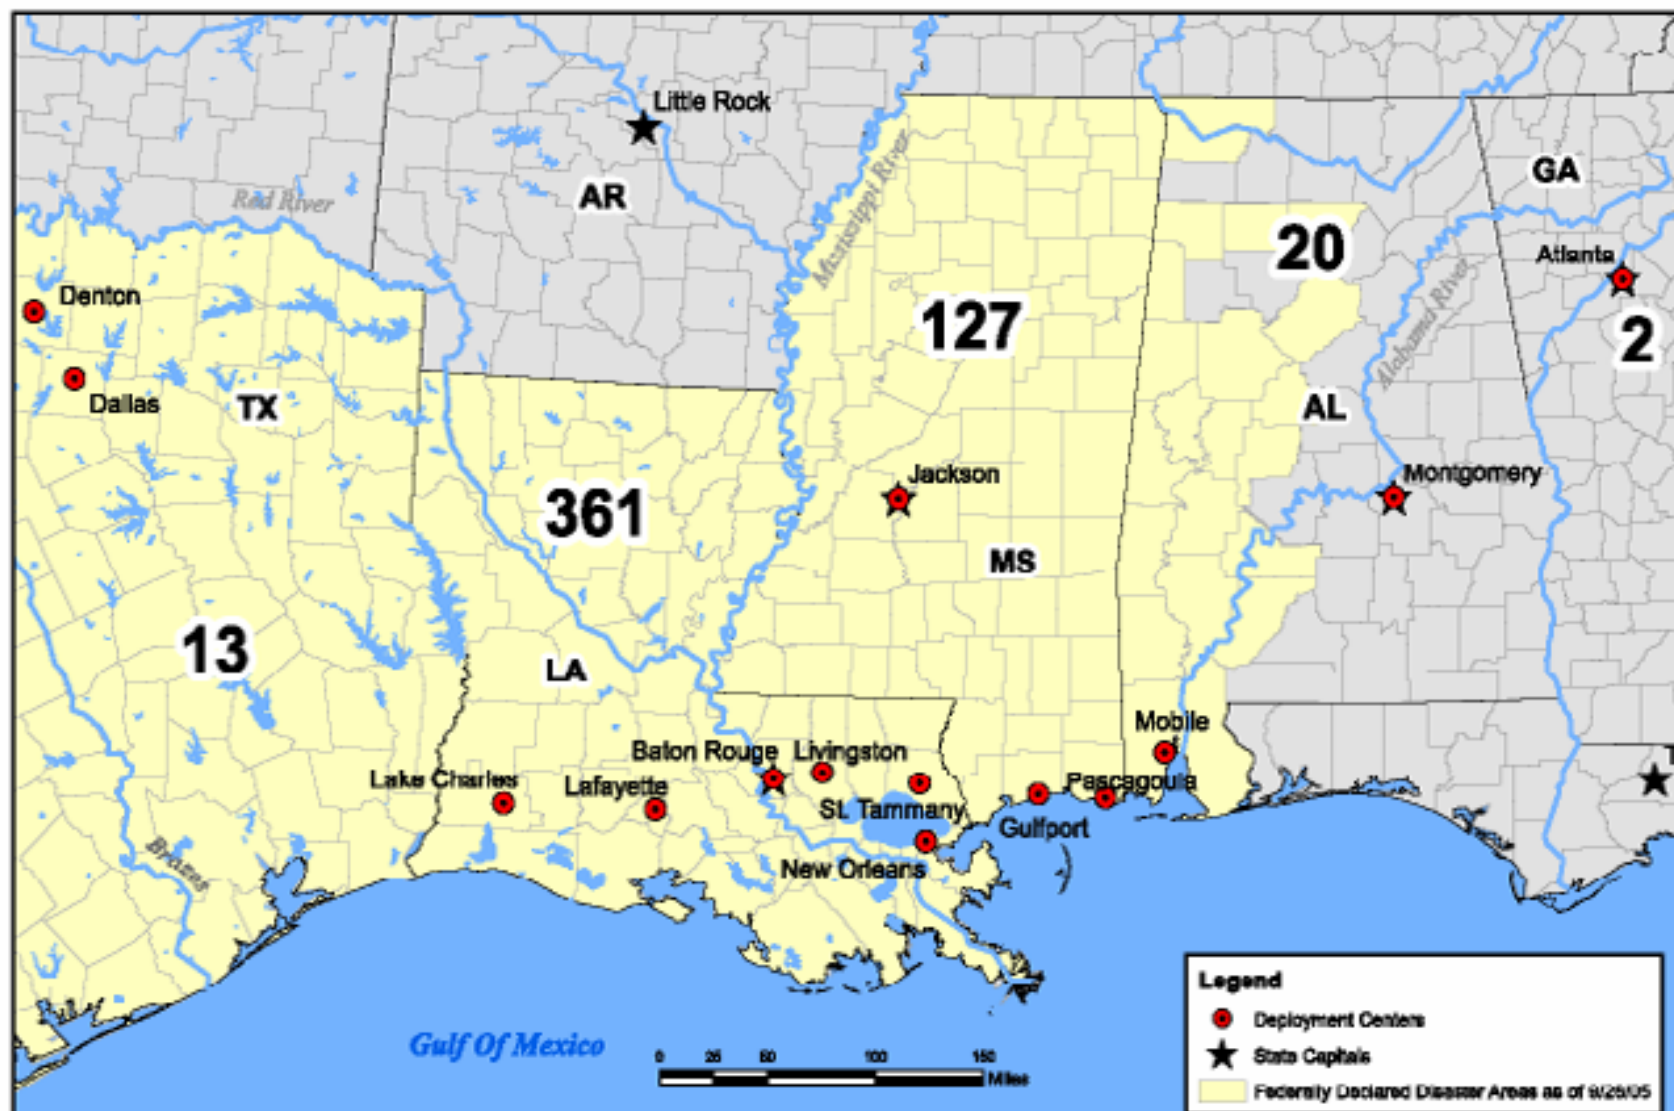




- New Orleans
 - ☒ DPA Regulated Facilities
 - ☒ APS_FS
 - ☒ Air Emissions
 - ☒ HPL_FS
 - ☒ Superfund Sites
 - ☒ HPL_FS
 - ☒ HPL_FS
 - ☒ HPL_FS
 - ☒ HPL_FS
 - ☒ Water Discharges
 - ☒ HPL_FS
 - ☒ Hazardous Waste
 - ☒ HPL_FS
 - ☒ Toxic Releases
 - ☒ Land Use Information
 - ☒ New Orleans City Data
 - ☒ Jefferson Parish
 - ☒ Transportation
 - ☒ National Pipeline
 - ☒ USA Streets
 - ☒ New Orleans Flooding and Drain
 - ☒ Zip Codes from Louisiana GIS
 - ☒ Census Geography
 - ☒ Water - Water
 - ☒ Detailed Water
 - ☒ Airports
 - ☒ Courtes
 - ☒ States
 - ☒ QuickBird New Orleans
 - ☒ Inframap30 New Orleans



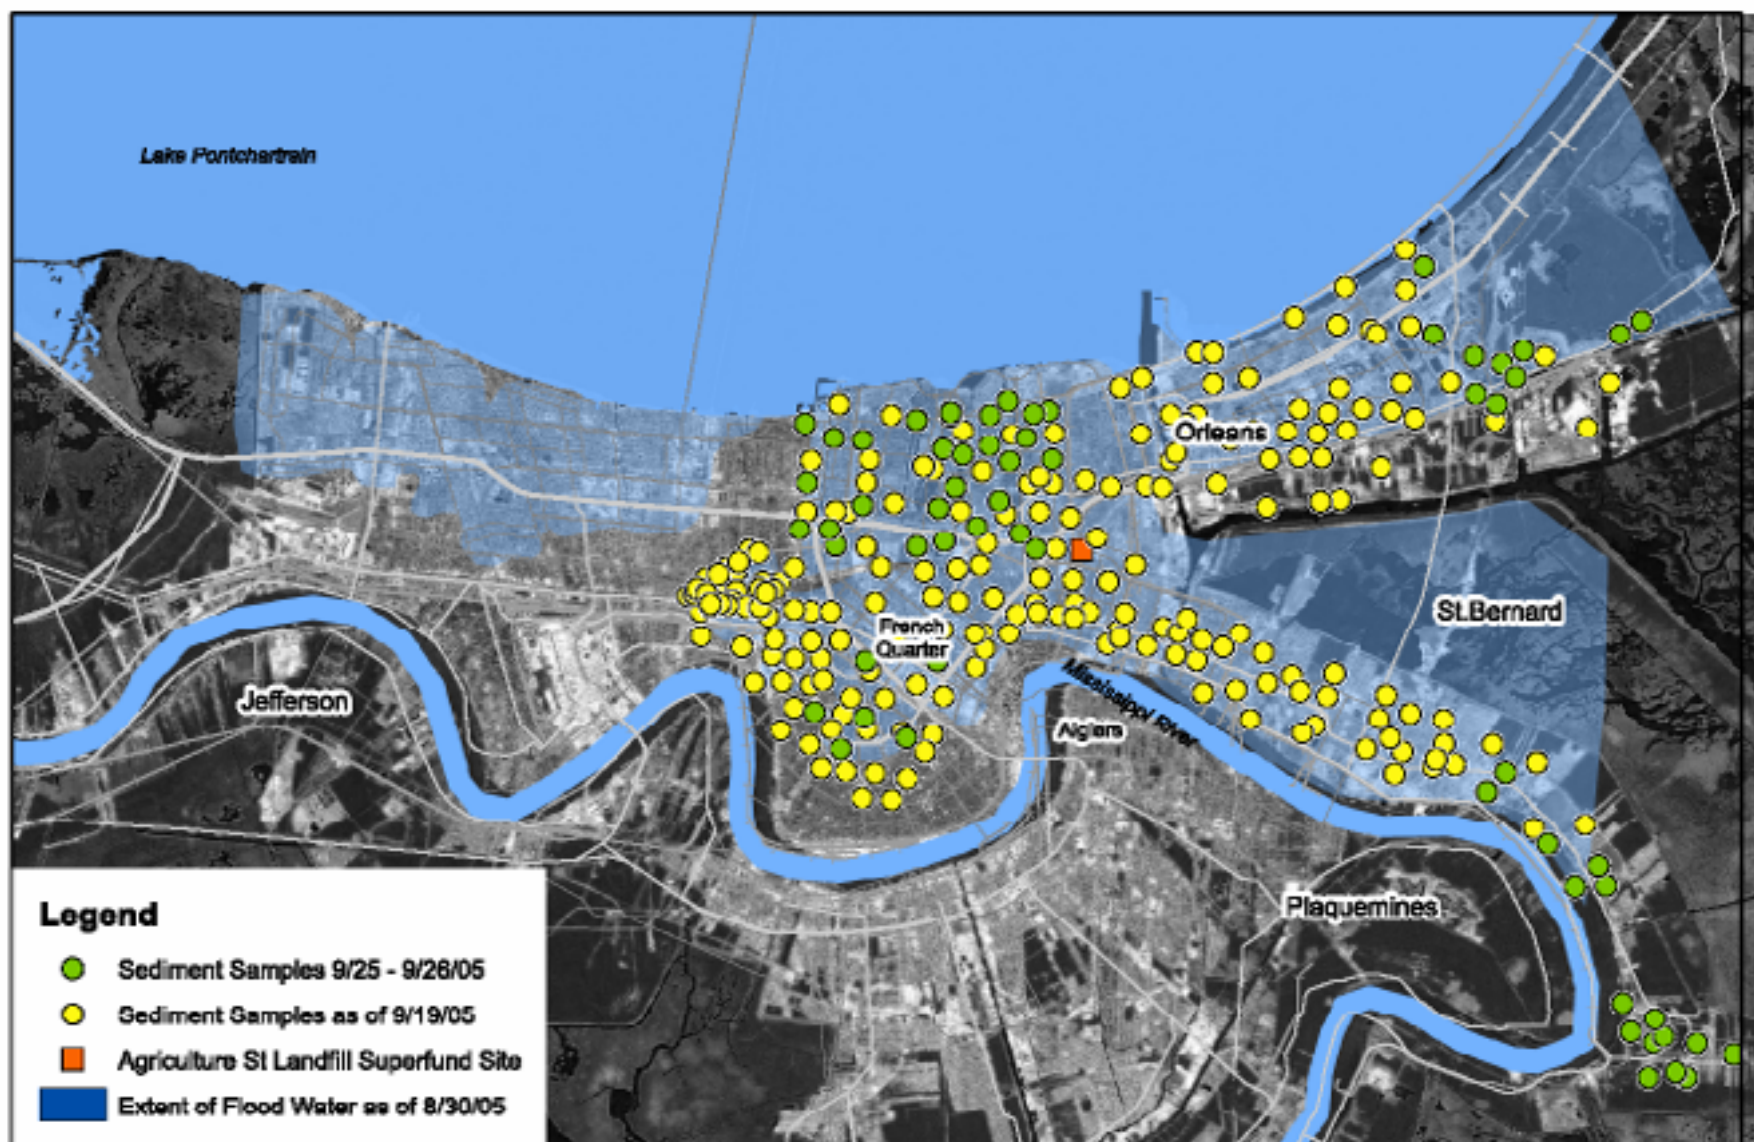
MAP CREATED FROM
 EPA INFRAMAP30 DATA
 FOR PUBLICATION TO THE PUBLIC
 UNDER THE PROJECT V.I. - 10/1/2000
 See the Inframap30 user manual for background
 information and project details.



EPA Deployment For Hurricanes Katrina and Rita as of 9/28/05

For Official Use Only
Source: EPA, ES&D






New Orleans Sediment Sampling as of 9/27/05


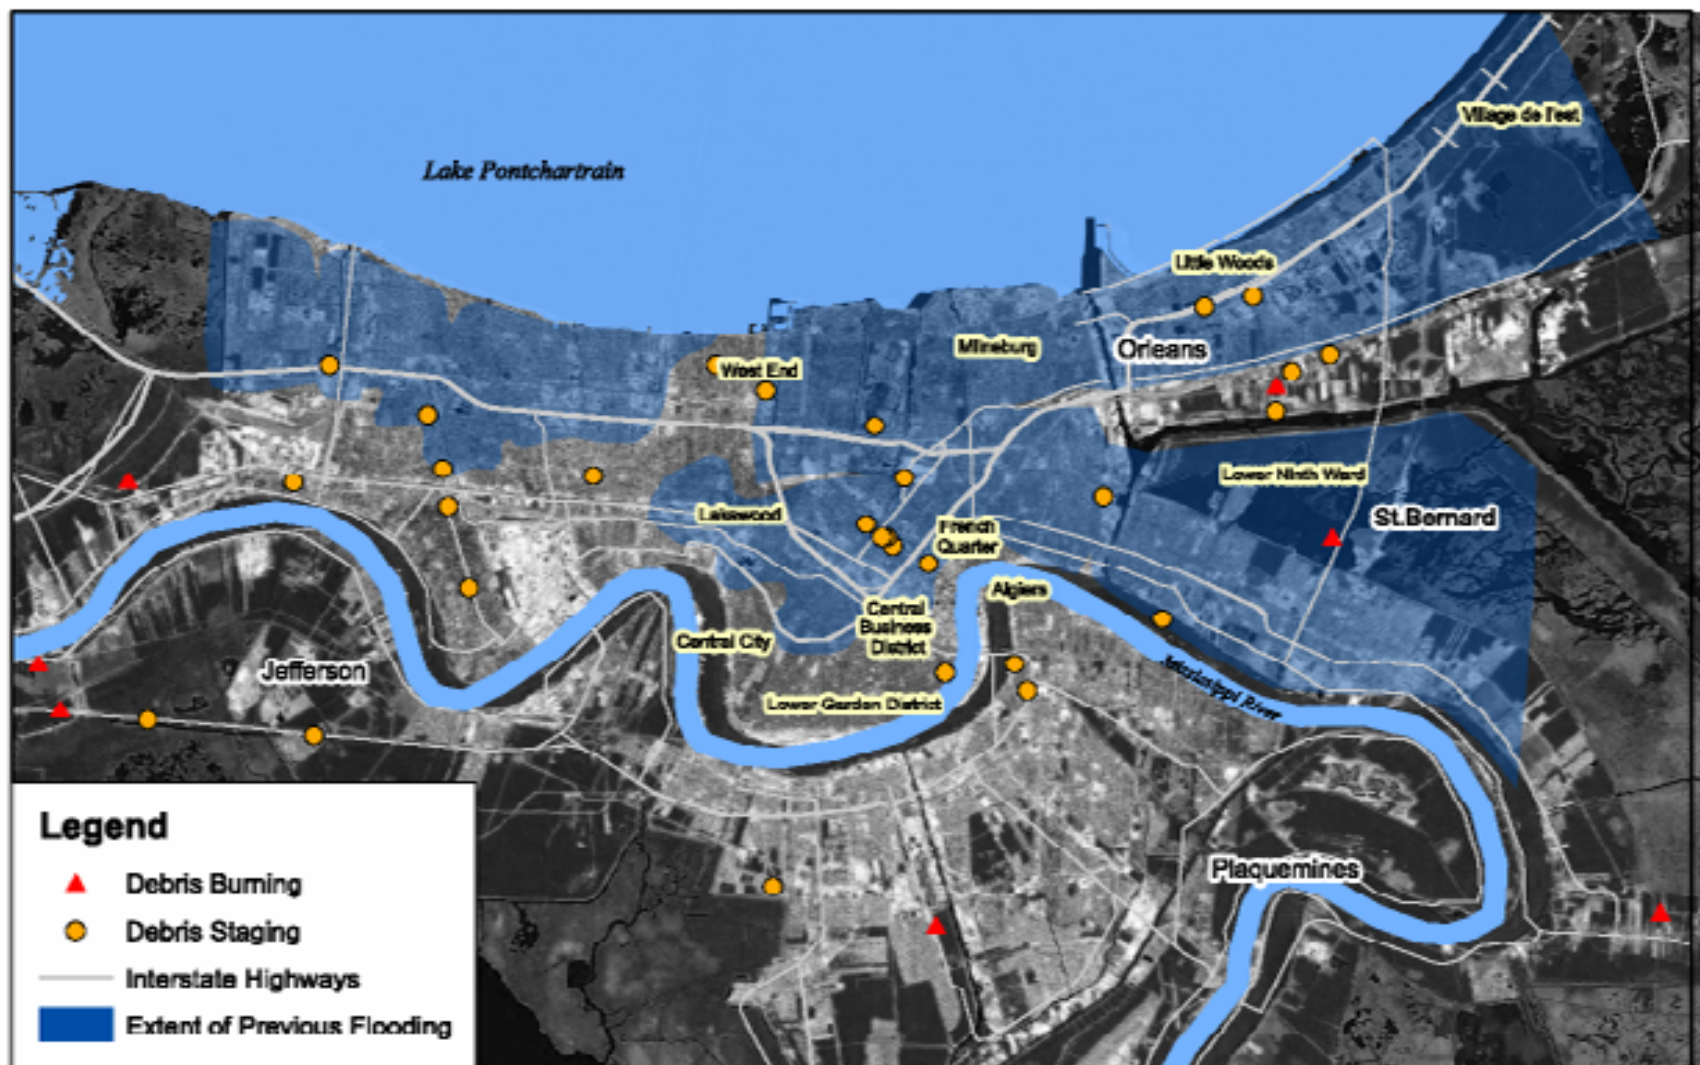
0 0.8 1.6 2.4 3.2 4 Miles

For Official Use Only
Source: EPA, NOAA, USGS and EBRU Data





Source: EPA, ESRI
For Official Use Only

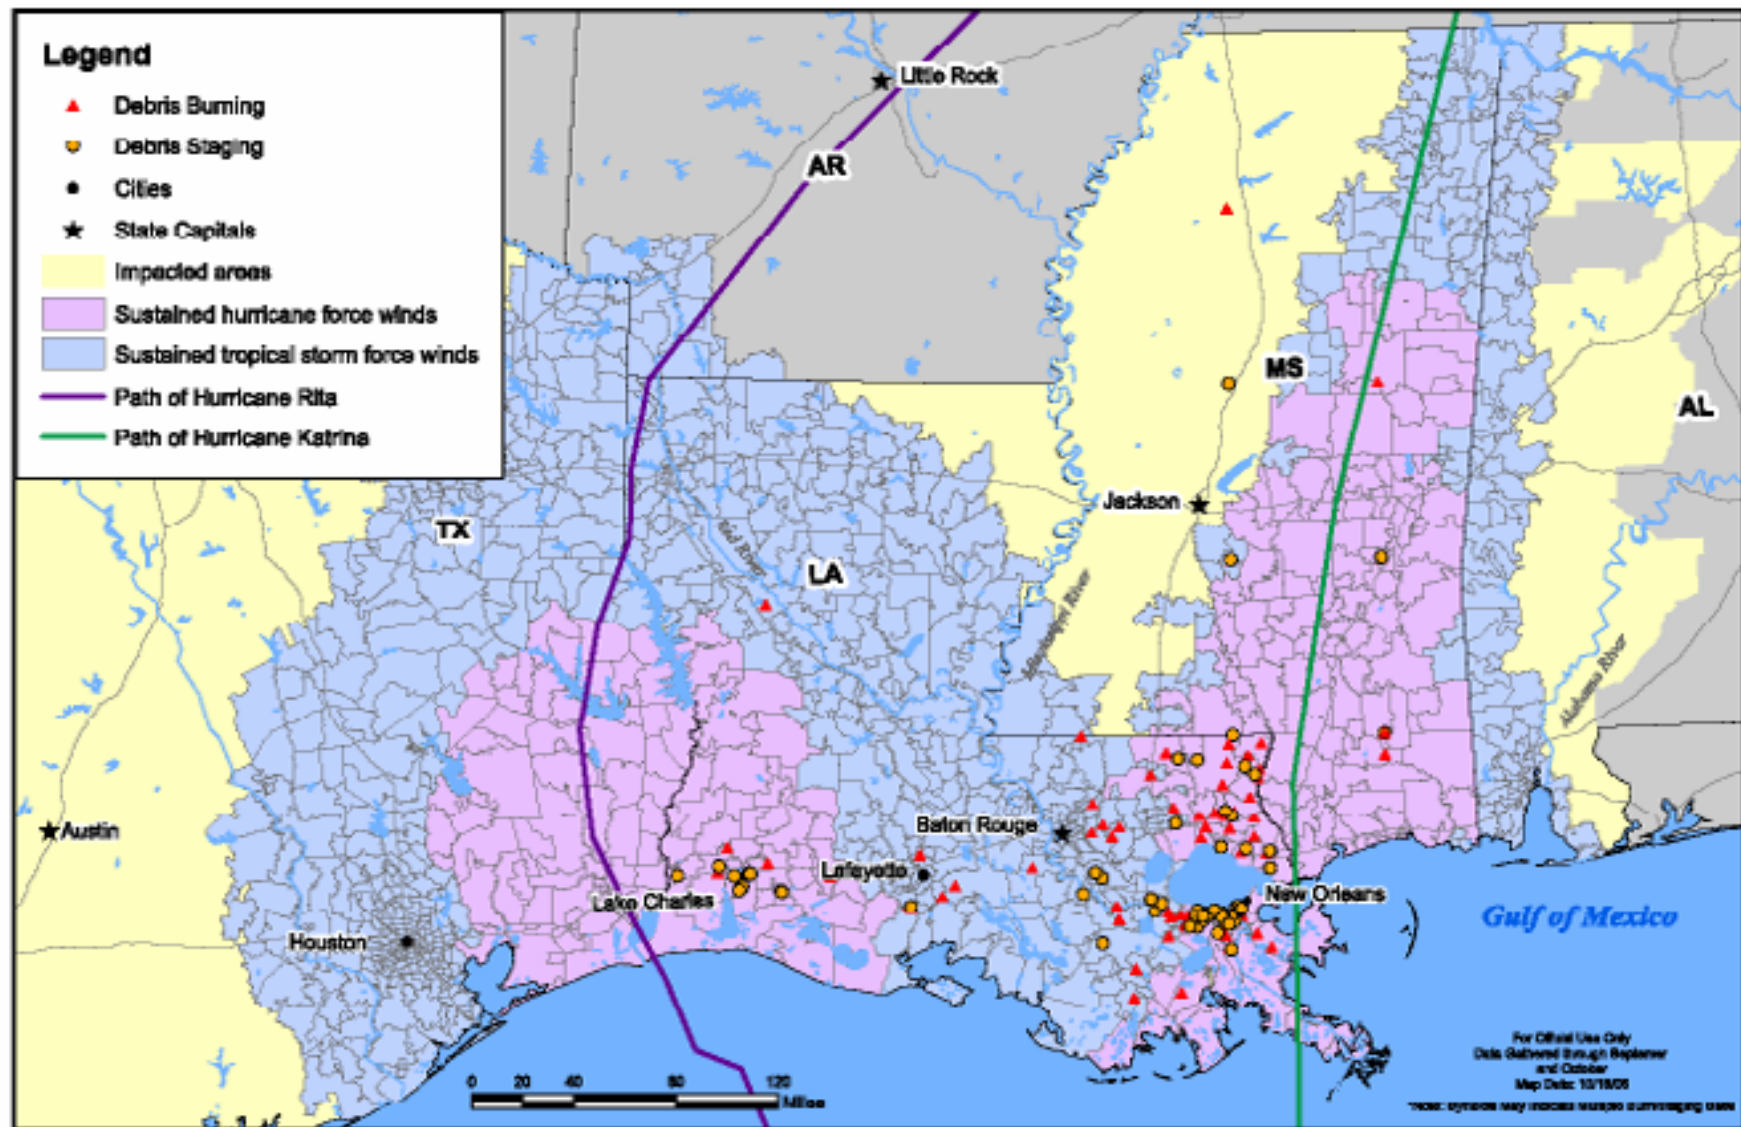


New Orleans Debris Burning and Staging Sites as of 10/14/05

0 0.8 1.6 2.4 3.2 4
Miles

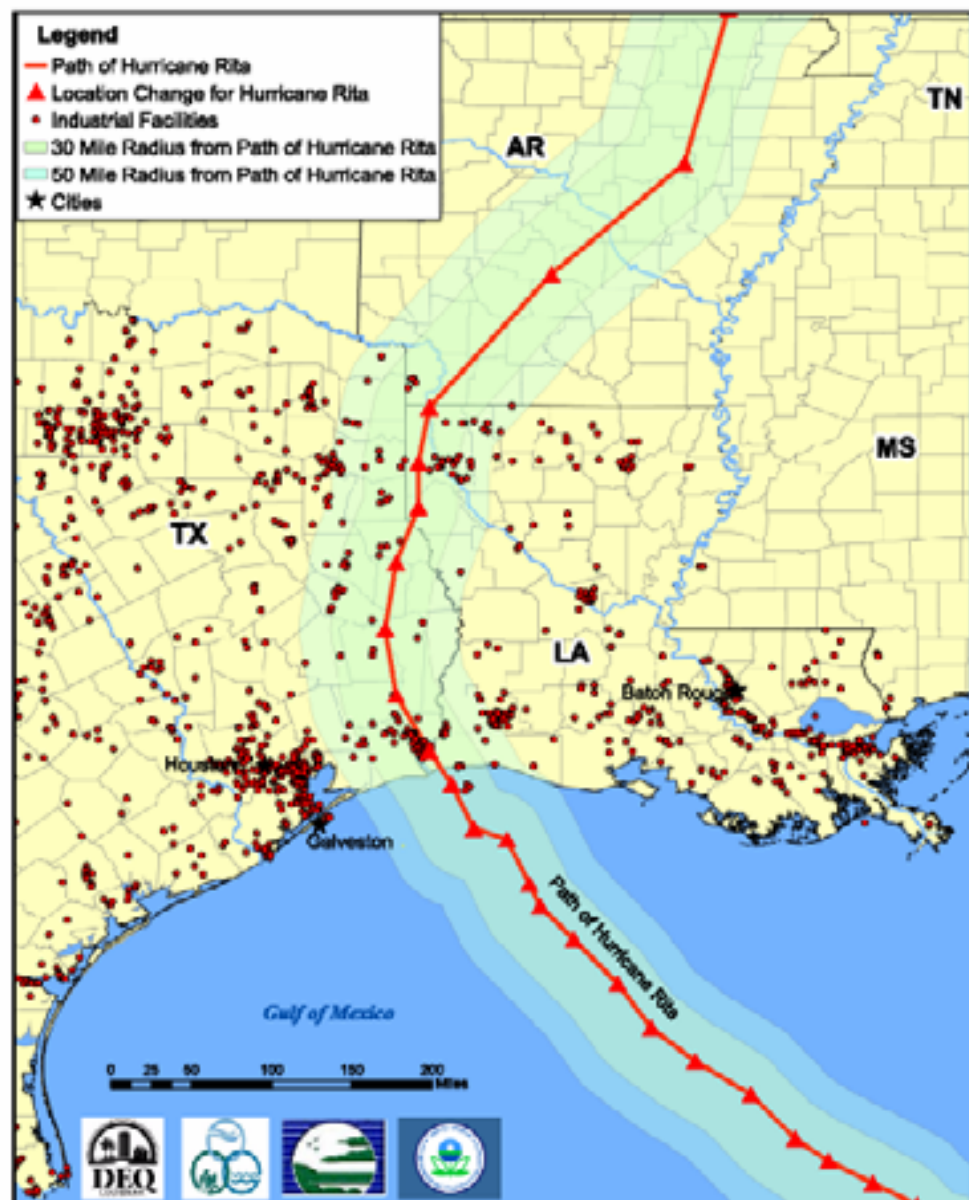
For Official Use Only
map date: 10/14/05





Wind Speed by Zip Code and Debris Burning/Staging Sites in Federally Declared Disaster Areas for Hurricanes Katrina & Rita





Industrial Sites In Path of Hurricane Rita as of 9/27/05

For Official Use Only



ArcIMS Web Service Reality Check

- ❖ Network connectivity is unreliable in the field
- ❖ Responders prefer local copies of all data to live web services
- ❖ Feature services can be exported/clipped but image services cannot
- ❖ EPA HQ EOC utilized image services heavily
- ❖ Gulf coast responders preferred to download compressed imagery via ftp



Lessons Learned/Implemented

- ❖ Formalized process for EPA central GIS support to perform imagery triage and post compressed images on ftp site for download
- ❖ Added many intranet feature services, starting with GDT Streets, then census boundaries and EPA facilities
- ❖ Took the time to tune and optimize the performance of the services
- ❖ Set up ESRI Data Delivery Extension for large extracts



Performance Optimization

- ❖ Optimize the OS system of the database server
- ❖ Optimize the RDMS (Oracle or SQL server)
- ❖ Optimize the ArcSDE server environment
 - ❖ The dbtune table
 - ❖ The giomgr.defs
- ❖ Optimize the individual SDE data layers
 - ❖ The Spatial Index
 - Oracle Spatial and IBM DB2 Spatial extender
 - ArcSDE Ge0database
 - ❖ The Attribute Index
 - ❖ Minimize on-the-fly operation with big dataset
- ❖ Optimize the SDE clients



Optimize the ArcSDE Server Environment

❖ The "dbtune" table

- ❖ Control how ArcSDE clients create objects.
- ❖ Determine which tablespace a table or index is created in.
- ❖ Define the extent size of the data objects
- ❖ Specify 1 of 4 available storage formats
- ❖ sdedbtune command to export / import dbtune table
 - Sdedbtune -o export – export the current dbtune table into a text file format
 - Edit the text file
 - Sdedbtune -o import – re-load the update dbtune file into the database table

❖ The "giomgr.defs" file

- ❖ Maximum SDE connection
- ❖ TCPKEEPALIVE




Optimize the Spatial Index

- ❖ What is the spatial index?
- ❖ How do you know if the spatial index need optimized?
- ❖ How to change the spatial index?
 - ❖ ArcSDE command line
 - ❖ ArcGIS desktop GUI



What Is The Spatial Index?


ArcSDE builds a spatial index by applying a grid to the feature class. It records features that fall within each grid cell in an index table (the S table of the feature class schema).





How Do Spatial Index Grids Work?

- ❖ Client submits a spatially-constrained query
“retrieve all land parcels that overlap a ruptured storage tank’s contamination plume”
- ❖ Server determines which parcels share an index grid with the plume
- ❖ Server finds shapes with overlapping envelopes
- ❖ Feature shapes compared to plume shape
- ❖ The key: only look at the data that meet your query criteria and ignore all other data





How To Check The Spatial Index Grid?

❖ ArcSDE command: `sdelayer -o si_stats`

Layer 5 Spatial Index Statistics:
Level 1, Grid Size 0.1

Grid Records:	212781							
Feature Records:	199991							
Grids/Feature Ratio:	1.06							
Avg. Features per Grid:	196.47							
Max. Features per Grid:	3887							
% of Features Wholly Inside 1 Grid:	93.96							
Spatial Index Record Count By Group								
Grids:	<=4	>4	>10	>25	>50	>100	>250	>500
Features:	199991	0						
% Total:	100%	0%						

This value should be < 2. If it exceeds 4, the spatial index need to be modified

This value should be > 100 & < 200



This value should be < 4000

This value should be > 80

•The majority of the features should have fewer than four records.



Examples of Improving The Spatial Index



Layer 1 Spatial Index Statistics:
Level 1, Grid Size 200

Grid Records: 107061
Feature Records: 41236
Grids/Feature Ratio: 2.60
Avg. Features per Grid: 2.93
Max. Features per Grid: 33
% of Features Wholly Inside 1 Grid: 38.19

Spatial Index Record Count By Group

Grids:	<=4	>4	>10	>25	>50	>100	>250	>500
Features:	38309	2928	621	163	60	24	9	3
% Total:	92%	7%	1%	0%	0%	0%	0%	0%

Too Small!



Layer 1 Spatial Index Statistics:
Level 1, Grid Size 4000

Grid Records: 43441
Feature Records: 41236
Grids/Feature Ratio: 1.05
Avg. Features per Grid: 283.93
Max. Features per Grid: 728
% of Features Wholly Inside 1 Grid: 94.94

Spatial Index Record Count By Group

Grids:	<=4	>4	>10	>25	>50	>100	>250	>500
Features:	41235	1	0	0	0	0	0	0
% Total:	99%	0%	0%	0%	0%	0%	0%	0%

Much better!

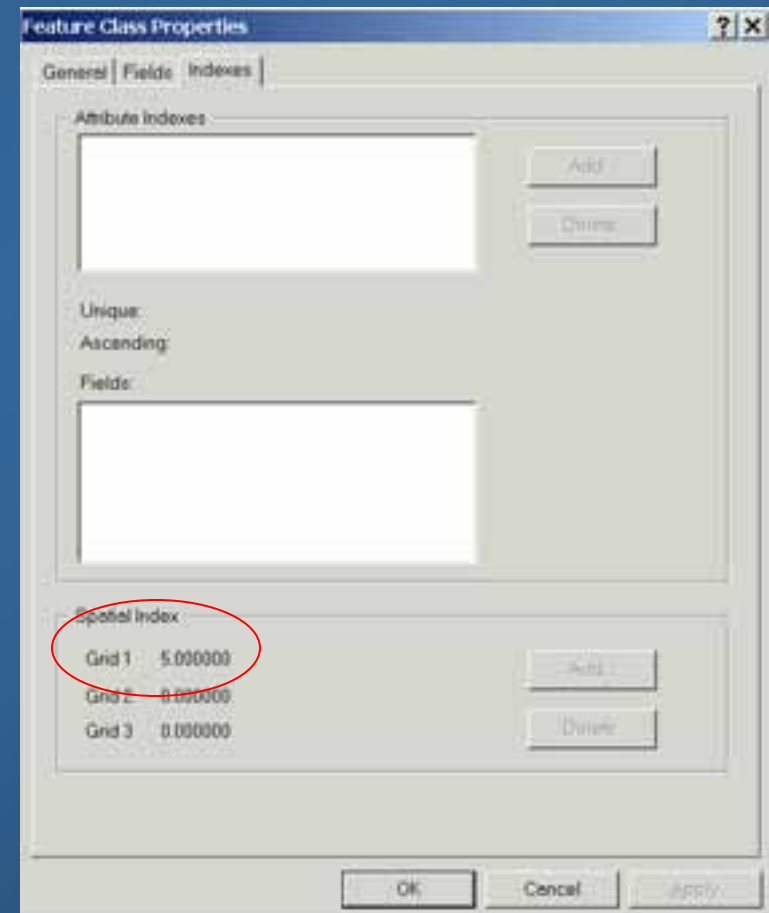


How To Change The Spatial Index Grid?

ArcSDE Command Line:

- ❖ Use "sdelayer -o load_only_io" to drop the spatial index. No spatial queries or data loading is allowed on this layer while it is in load-only mode.
- ❖ Use "sdelayer -o alter -g n,n,n" to specify new grid sizes. (Specify zero for the second or third grid size if not used.)
- ❖ Use "sdelayer -o normal_io" to rebuild the spatial index and make the layer accessible again.


ArcCatalog GUI





Spatial Index Recommendations

- ❖ Experiment with grid settings. Can change grid size any time.
- ❖ Start with 3X to the LARGER of:
 - ❖ Your layers average feature size or
 - ❖ Your average querying area
- ❖ Never make grids smaller than your features
- ❖ Most of the time **one grid level** is enough
- ❖ Keep average number of features per grid between 100 and 300
- ❖ Try not to let the max number of features per grid exceed 2000 to 4000
- ❖ Maintain a high % of features wholly in 1 grid
 - ❖ >90%=great; >80%=good; <80%=change
- ❖ Tuning grids to a “general” setting works well for most queries
- ❖ Automation of the optimization routine



1x of average
feature extent.



3x of average
feature extent.



What is the Attribute Index?

- ❖ Two columns: indexed field and key
- ❖ Index table has the same number of records as the attribute table.
- ❖ The indexed field will be sorted with different indexing technology
- ❖ The index table could be very big



How to Build Attribute Indexes

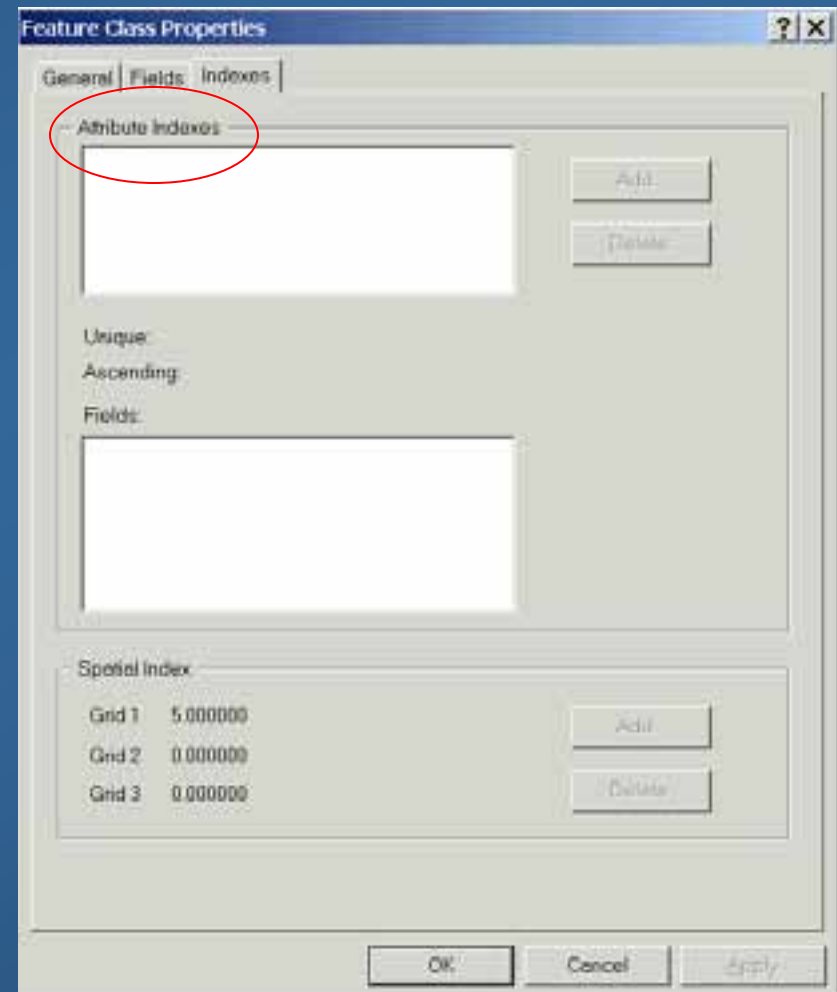
❖ ArcSDE Command Line:

- ❖ Use "sdetable" SDE command to drop the attribute index.
- ❖ Use "sdetable" SDE command to build new attribute index

❖ ArcGIS Toolbox Wizard Interface

❖ ArcCatalog GUI

❖ Oracle Tools





Attribute Index Recommendations

- ❖ Index attribute fields you query often on
- ❖ Indexing both the “where” and “select” columns helps business table queries in the same way as it helps the S table queries.
- ❖ Different indexing technologies
 - ❖ unique; bitmap; unsorted; reverse
- ❖ Check application logic
 - ❖ Avoid unnecessary multiple criteria in “where” clause



Minimize On-The-Fly Operations

- ❖ The name system of the GDT street data contains 3 separated columns:
 - ❖ Prefix: e.g. South; Name: e.g. Miami; Street type: e.g. Drive
- ❖ To use the full name (South Miami Drive) for labeling, you have to join these 3 columns on-the-fly.
- ❖ To avoid on-the-fly operations, you can add a new attribute column to the database that contains all three parts of the name.
- ❖ If the dataset is small, the improvement may not be obvious. If the dataset is big, the improvement could be quite significant.



Performance Comparison Tests

- ❖ The same database server (SDE/Oracle)
- ❖ Same dataset: GDT street data
- ❖ Same SDE client: ArcGIS Desktop with connection to the SDE/Oracle database directly
- ❖ The before optimization test:
 - ❖ Default ESRI spatial index
 - ❖ No attribute index
 - ❖ No new attribute columns with all 3 name components
- ❖ The after optimization test:
 - ❖ Rebuilt the spatial index 8 times
 - ❖ Build attribute indexes for the FCC code and name label columns
 - ❖ Add a new attribute columns with all 3 name components



Results: Before vs After Optimization

- ❖ On average, the “after optimization” test is **30% faster** than the “before optimization” test
- ❖ In several cases, the performance improvement is over 50%
- ❖ For desktop operation, the improvement is quite obvious with big dataset
- ❖ For web application, the improvement is quite obvious even with small and mid size dataset



Web Service Considerations

- ❖ ArcIMS Image/Feature Services
 - ❖ Setting up proper number of spatial servers
 - ❖ Setting up proper number of instances
- ❖ ArcIMS OGC WMS/WFS connectors
- ❖ XML/SOAP based Geo Web Services
- ❖ Other Considerations



Conclusion/Summary

- ❖ EPA GIS Emergency Response is a team effort supporting the field personnel and decision-makers
- ❖ Imagery triage is a perfect task to offload onto centralized support
- ❖ ESRI geoprocessing and server tools met the challenge
- ❖ Taking stock of lessons learned is a fruitful interagency exercise that should help us be better prepared for the next crisis



Acknowledgements

RTP GIS Support Team

GIS/Remote Sensing Project Manager:

Dick Dulaney

GIS Analysts:

Mark S. Murphy

Catherine Harness

Binfei Chen

Senior System Administrator:

Janet Braxton

Customer Support Representative:

Claudia Benesch

EPA GIS Personnel

Geospatial Information Officer:

Brenda Smith

National Hosting Operations Coordinator:

Tim Richards

Region 2 GIS Coordinator:

Harvey Simon

Region 10 GIS Coordinator:

Scott Augustine



Contact Information

Qi Dai

GIS Technical Lead

dai.qi@epa.gov

919-767-7285

Torrin Hultgren

GIS Analyst

hultgren.torrin@epa.gov

919-767-7312

79 T.W. Alexander Drive, Building 4201

Research Triangle Park, NC 27709