Vistronix, Inc®

Ingenuityinnovationintegrity

2013 Esri International User Conference

July 8–12, 2013 | San Diego, California

Vistronix is a leader in Technology Solutions and Information Management Services.

Building Cache Tiles with ArcGIS Server 10.1 Real World Examples from USDA/NRCS/NGCE

Presenter: Art Ullman
Contributor: Dr. Suling Zhao
DATE: 3/20/2013

1851 Alexander Bell Drive Suite 350 Reston, VA 20191
703.734.2270 800.483.2434 www.vistronix.com

Project Overview

- USDA / NRCS has a requirement for high performance Geospatial Web Services to support planning and implementation of conservation practices for farmers and ranchers across all 50 states and territories.
- NRCS users need imagery, elevation data, and vector data delivered over USDA Network to web applications and ArcGIS Desktop.
- Support over 10,000 users in over 3000 counties and over 50TB of geospatial data and imagery.
- National Geospatial Center of Excellence (NGCE) at NRCS is responsible for providing the data and building web services.

Solution –

- Geospatial Web Service based on ArcGIS Server 10.1 SP1
- Use cache tile services to improve performance and scalability for imagery and complex vector layers.
- Vistronix and NGCE staff at the *National Geospatial Center of Excellence* in Fort Worth, Texas will build services using NRCS Geospatial Data Warehouse data and hardware.
- Use ESRI appliance cache tiles for topo maps, transportation and other layers available on the appliance.

Cache Tile Services - Overview

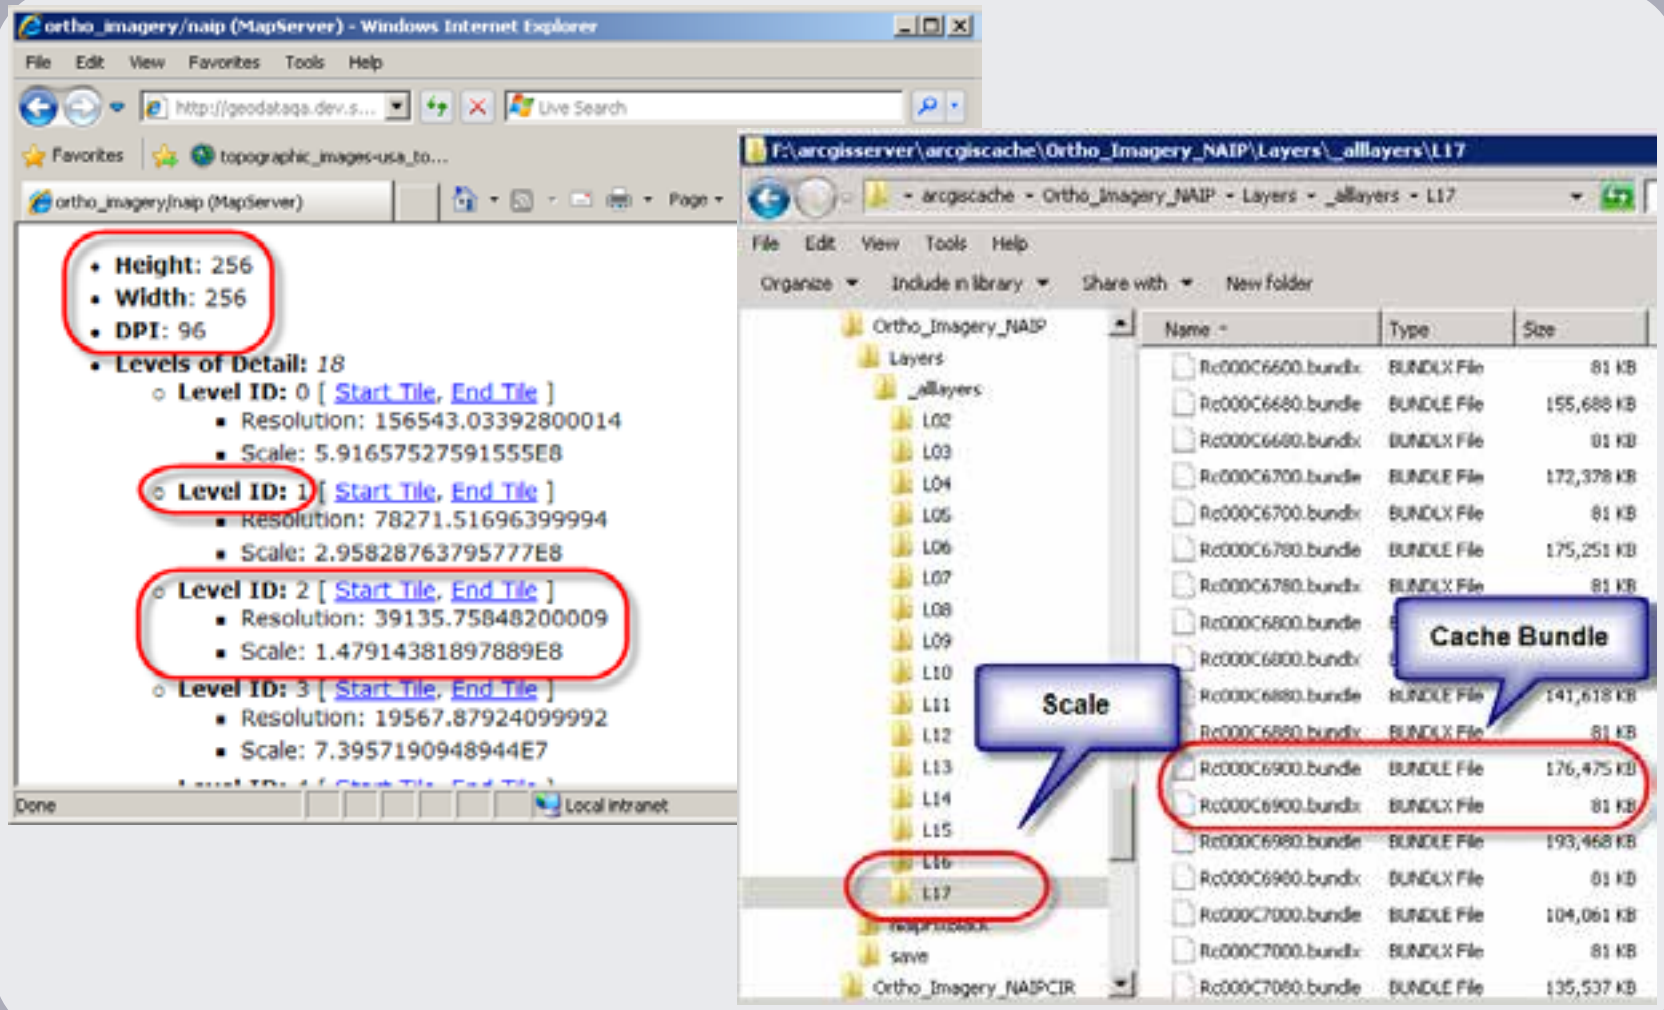

Cache Tile Services - Overview

- Why use cache tile services?
  - Performance
    - Eliminate complex queries and on-the-fly image/mosaic rendering on server.
  - Scalability
    - Server simply reads and sends JPG/PNGs to client. Scalability is determined by how fast the server can read and send JPG/PNGs to client.
  - User Experience
    - Users get continual flow of tiles as they pan and zoom.
    - Application is always responsive and continuously gives visual feedback to the user as tiles are delivered.
    - Users are used to this kind of experience from using Google, ESRI Online and BING maps.

Cache Tile Services – Key Steps

- Design
  - Projection – Web Mercator Auxiliary Sphere in most cases.
  - Cartography - Symbology, labeling, etc.
  - Scales, extent and image format.
  - Tile Scheme – ESRI/Google/BING
- Develop
  - Prepare input datasets
  - Hardware – non-production servers / workstations for building tiles. One large multi-processor machine or a cluster of lower power machines.
  - Storage – each scale requires 3-4 times storage of previous scale.
  - Time – product of spatial extent and scale divided by CPU count.
  - Test small areas and view results!!!
- Deploy
  - Deploy data and cache - or just cache?
  - Copy cache to production server and publish service.

NRCS / NGCE Cache Tile Services

Layer	Type	Size	Max Scale / Resolution
NAIP – Natural Color	Imagery	2.15 TB	1:4513 -1.19m
NAIP – Color Infrared	Imagery	2.10 TB	1:4513 – 1.19m
NHD - Hydrography	Vector	516 GB	1:4513 – 1.19m
Relief - Elevation	Raster	112 GB	1:36111 – 9.55m
WBD – Hydrologic Units	Vector	93.8 GB	1:4513 – 1.19m
State / County	Vector	87.4 GB	1:9026 – 2.38m
Contours	Vector	689 GB	1:4513 – 1.19m
(5) Easement Imagery Layers	Imagery	620 GB each	1:1128 - 0.29m
(7) layers from ESRI Appliance including topo maps and transportation	Img & Vector	27 TB	mixed

NAIP 4 band 1m – Natural Color - bands: 1,2,3

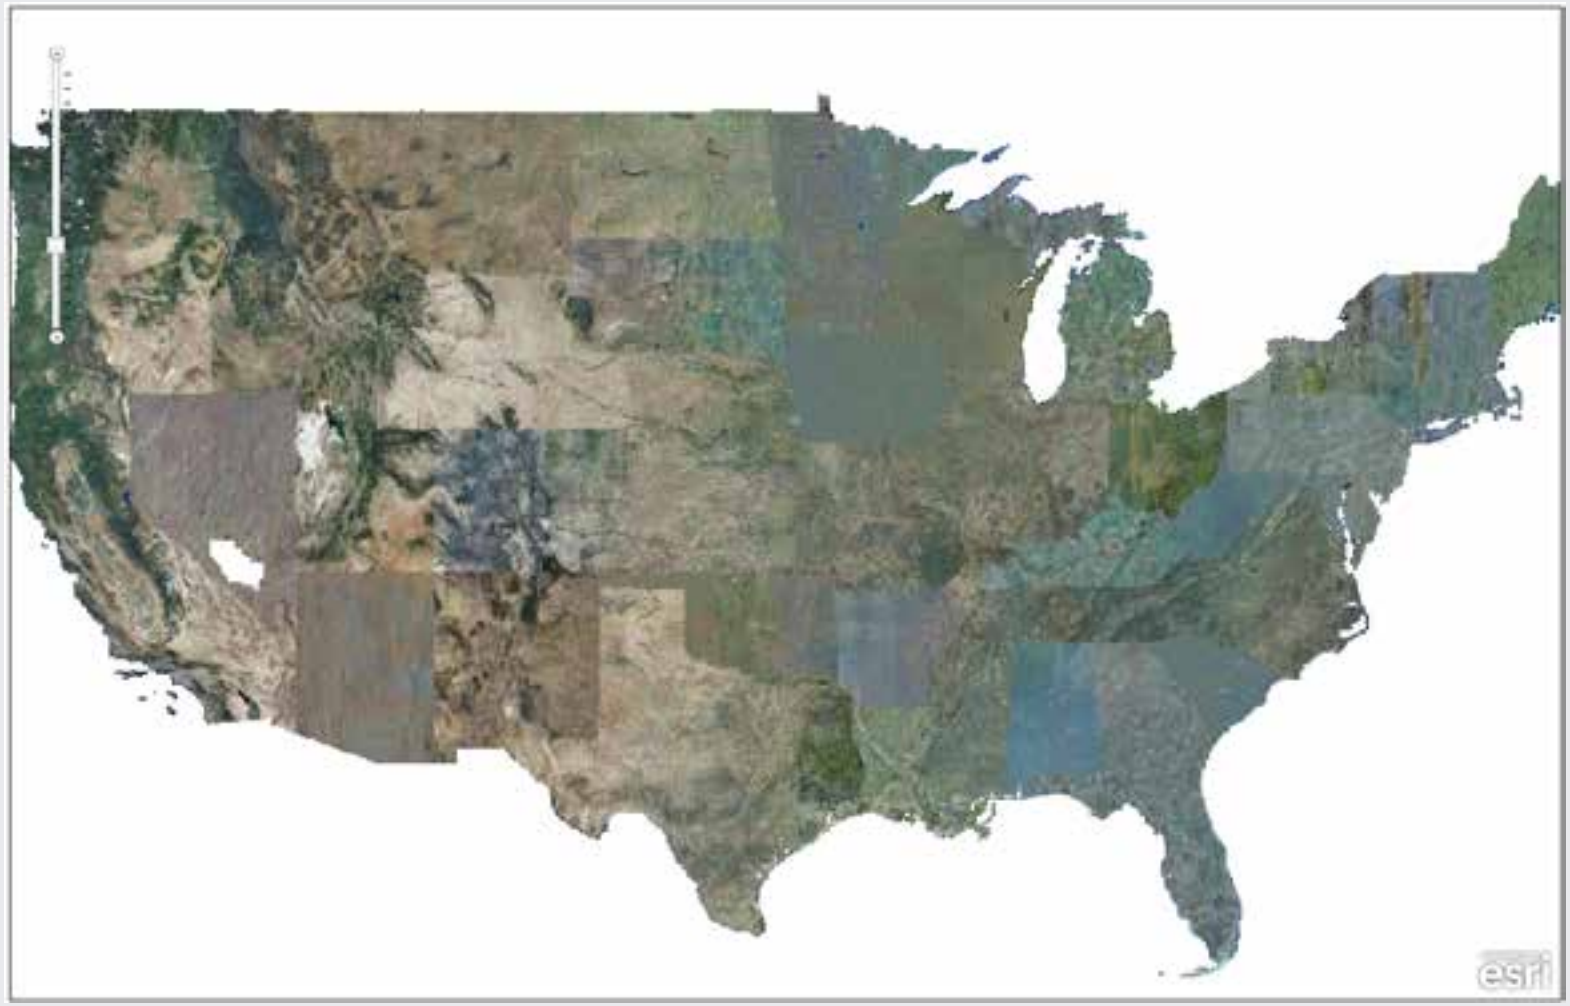

NAIP 4 band 1m – Color Infra-red layer – Bands: 4,1,2

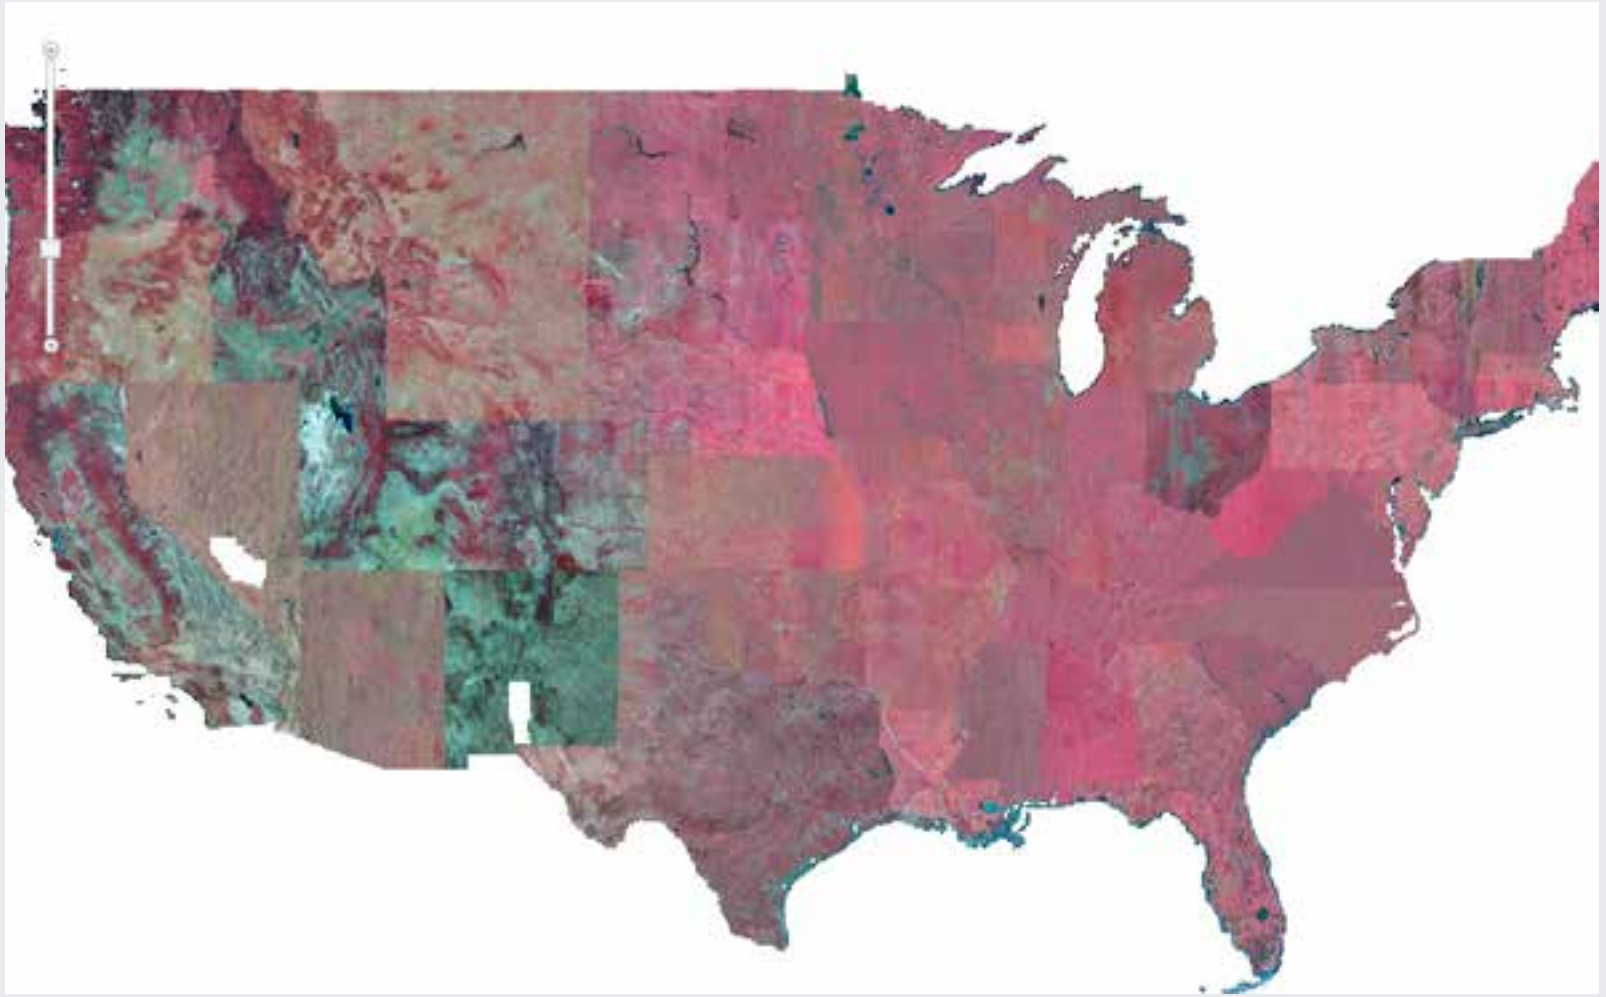

Transparent (PNG) National Hydrography Dataset cache tile – draped over imagery tiles

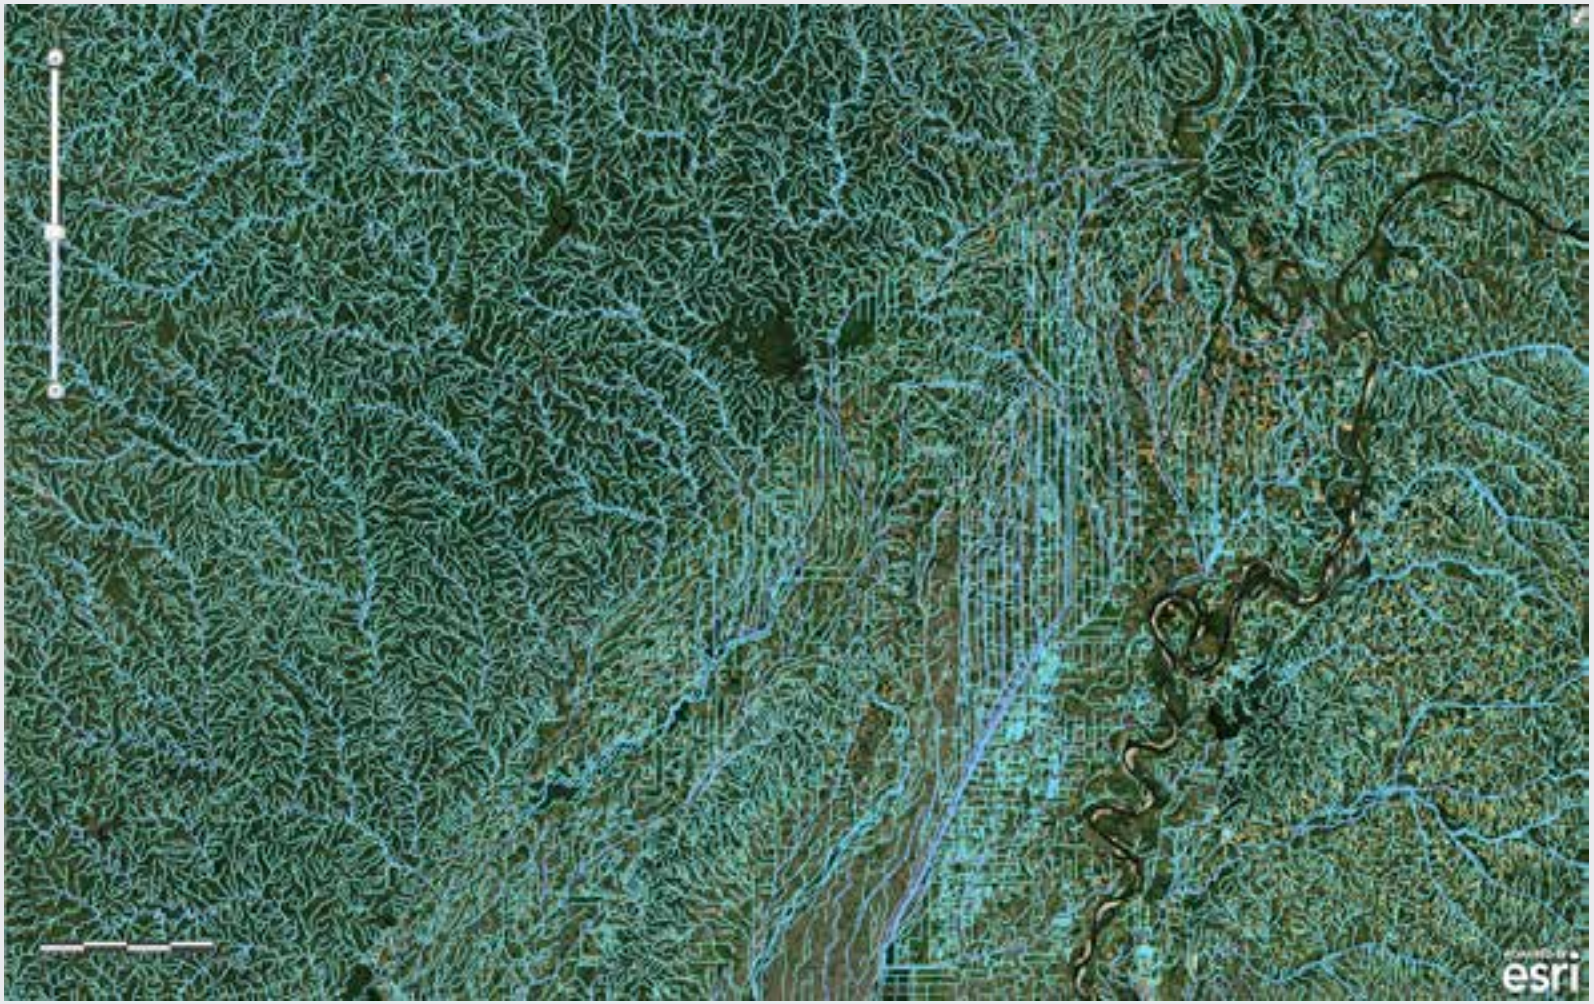

Architecture – NAIP Imagery Processing

Tape Library
350TB Archive

GeoTiffs

USDA
FSA/APFO

(Offline)
ArcGIS Server 10.1
For Building Cache

Dell PowerEdge R814
(4) 8-Core Opt 6212
(2) 27TB DAS trays
USB 3 – file transfers

Cache
Tiles

(2) Production
ArcGIS Server 10.1

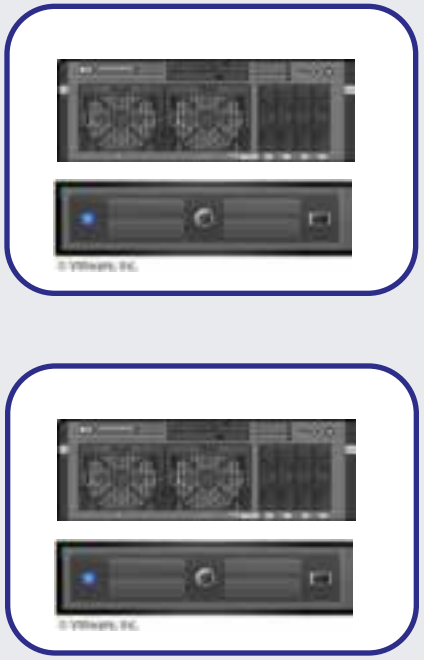

(4) CPU VMs
(1) 27TB DAS tray

Imagery – NAIP

- Input dataset overview
  - 1m NAIP (National Agriculture Imagery Program)
  - 4 band – produce two cache services: Natural Color (RGB) and Color Infra-Red (CIR)
  - Yearly refresh of 1/3 of states
  - Includes full coverage for lower 48 states.
  - 210,000 geotiffs – 23.2 TB
    - Pre-2012 – 20MB JPG compressed geotiffs (3.2TB)*
    - 2012 – 200MB uncompressed geotiffs (20TB)

* Pre-2012 NAIP data was compressed and loaded to ESRI Image Server. Compression was needed to save space.

Imagery

- Design
  - Source Data - 30TB 4band 1m resolution quarter-quads geotiffs
  - Mosaic Dataset for RGB – bands 1,2,3
  - Reference Mosaic Dataset for CIR – bands 4,1,2
  - 2 Cache Tile Services (RGB/CIR)
  - Mixed Mode
    - JPEG for all internal tiles
    - PNG32 for all tiles that have empty areas - border areas.
  - Google/BING/ESRI Tile Scheme
  - Scales: 1:145,000,000 – 1:4,513 (1.1 m)

Old 9.3 / 10.0 Processing Workflow

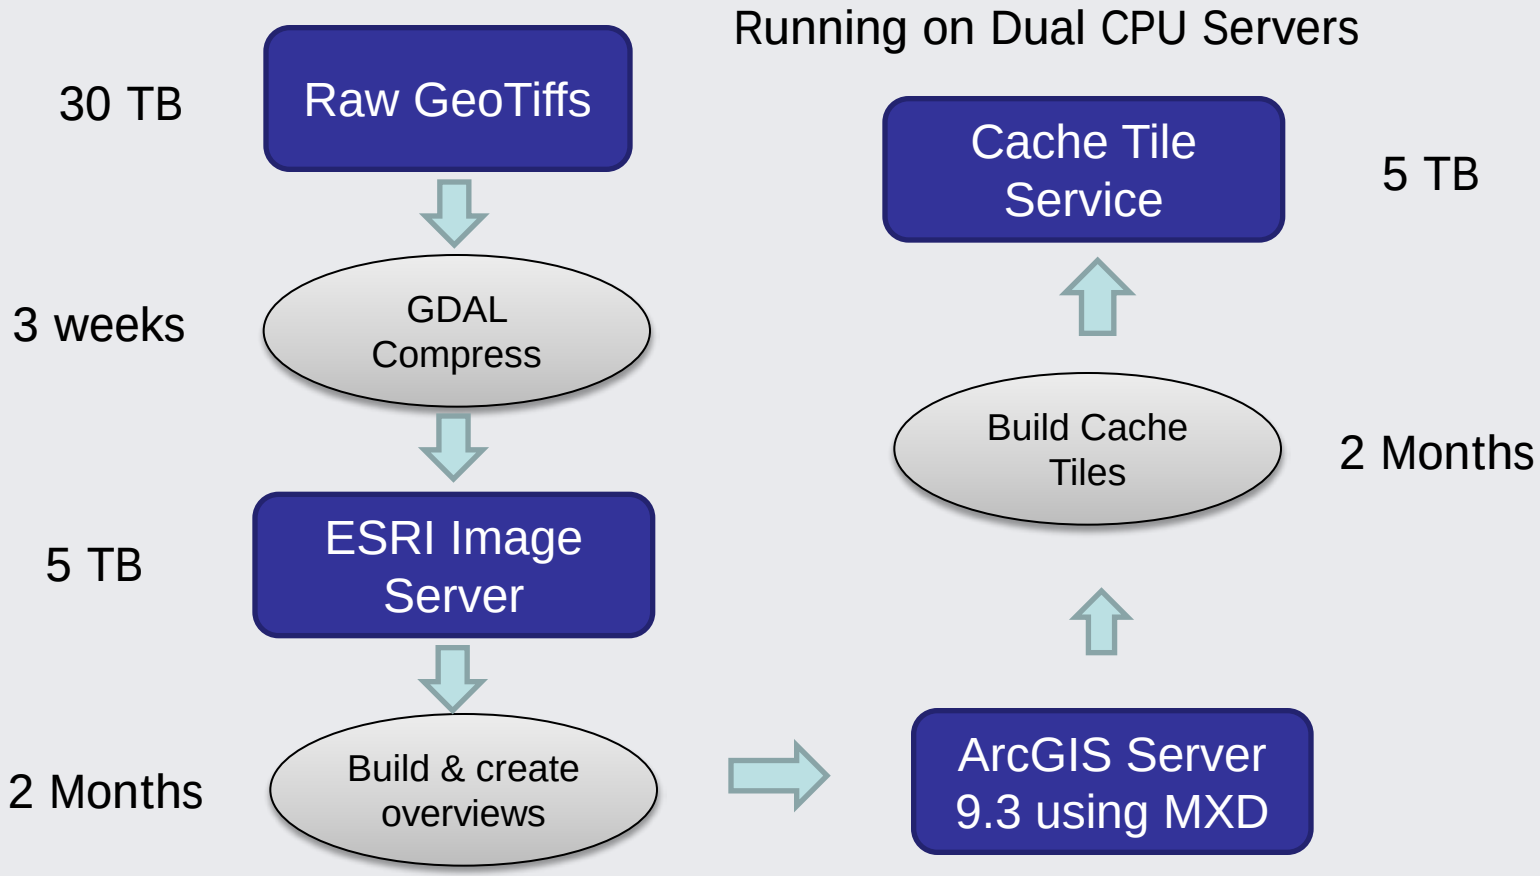

10.1 Processing Workflow

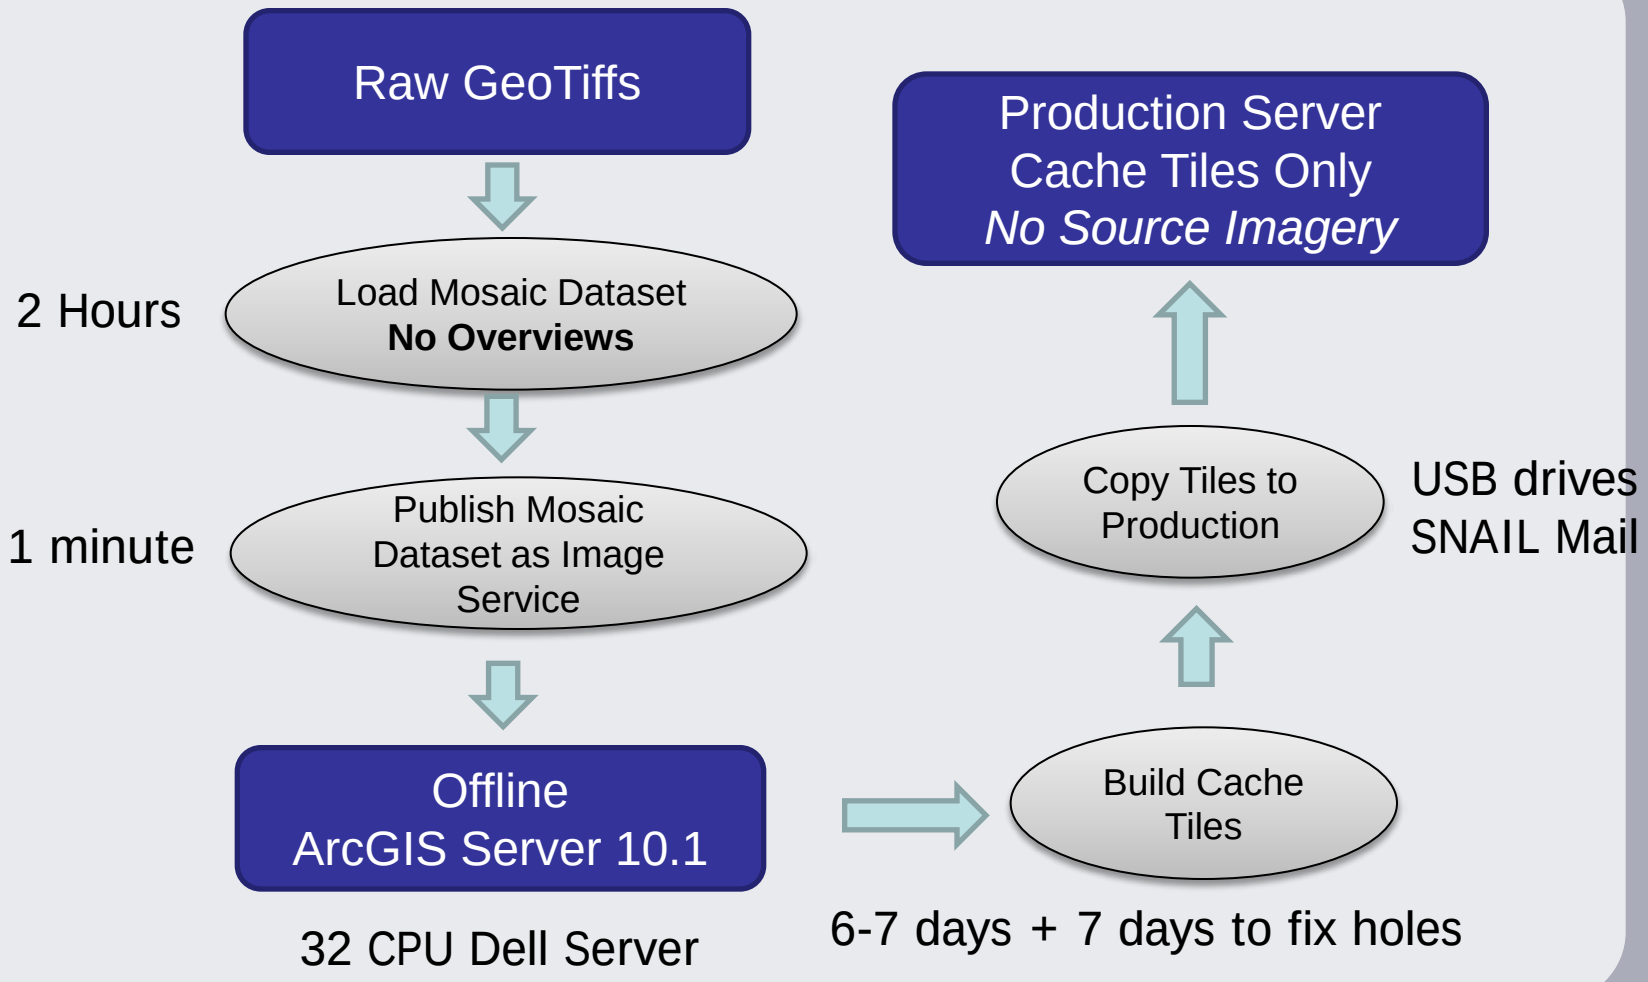

Building Cache for Imagery

- Build Mosaic Dataset (MD)
- Add Rasters to MD
- Fix Boundary if necessary
- Calculate Statistics
- Publish Mosaic Dataset as an Image Service.
Note: requires Image Server Extension for ArcGIS Server
- Set Caching Parameters
- Start process to build cache

Imagery – Building Cache

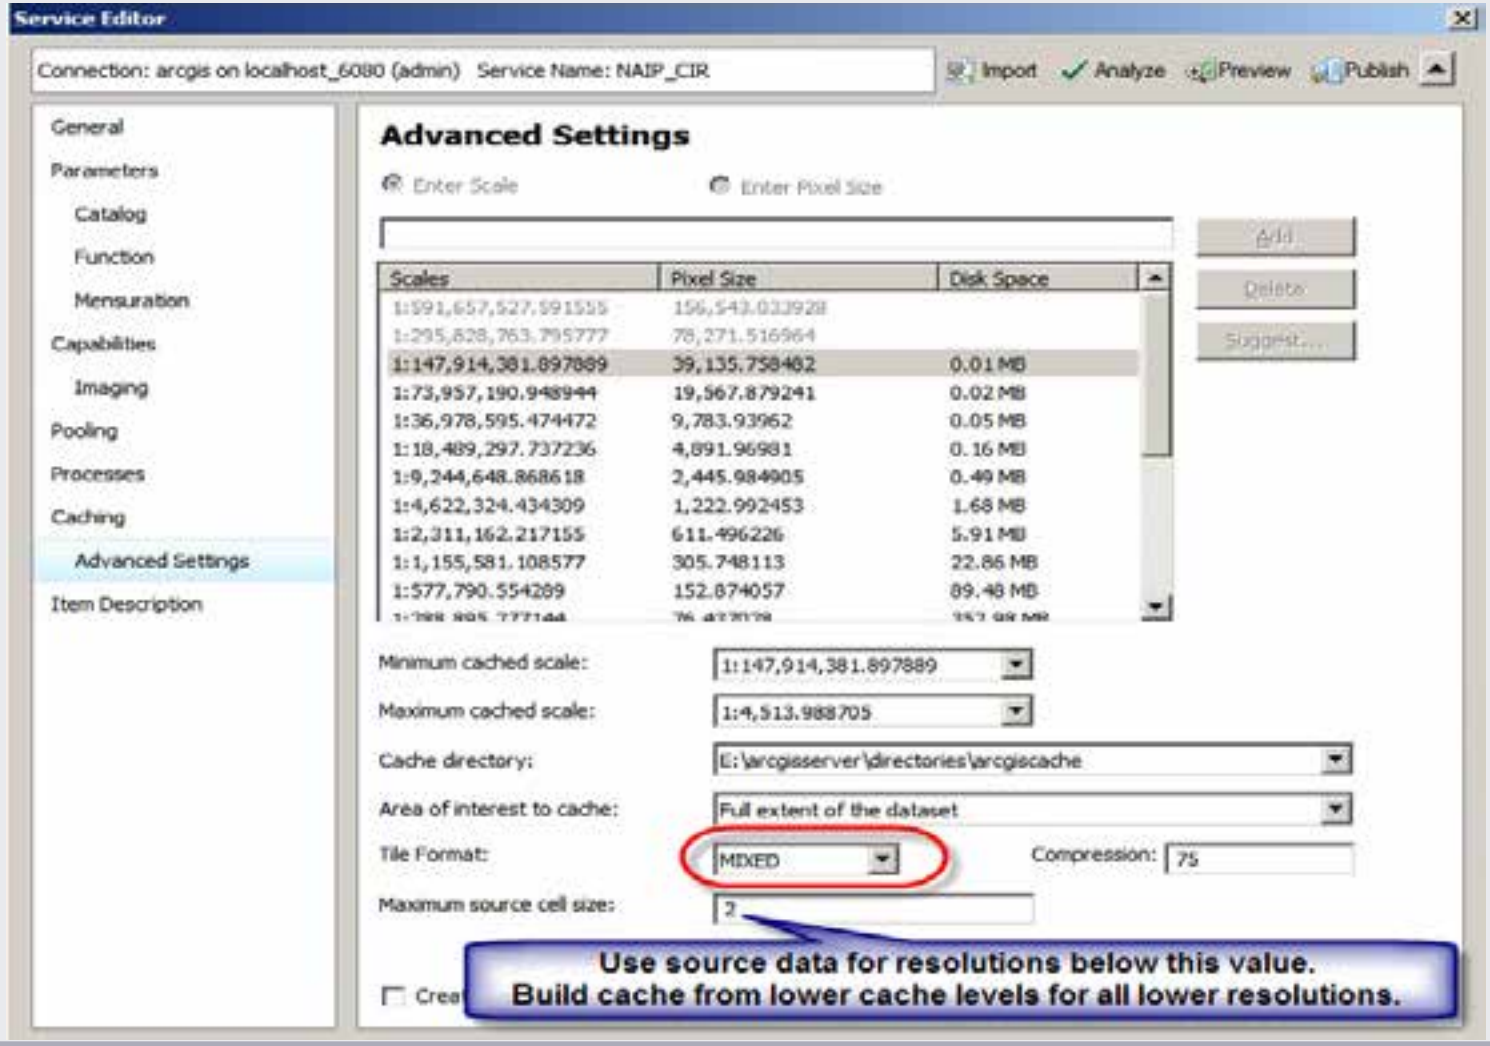

Building Cache from Image Services

- Publish Mosaic Datasets as Image Services.
- In ArcGIS10.1 you can build Cache Tiles directly from a Mosaic Dataset published as an Image Service (no MXD needed).
- Cache for Image Services are built from bottom up (high resolution first).
  - ArcGIS uses the *Max Cell Size* parameter to determine whether to build cache from source data or from cache tiles at higher resolution. Set the Max Cell Size to slightly higher than source resolution to build all lower res. caches from higher res. cache.
  - It is faster to build cache from higher resolution cache than from source data. Source data files can be large – requiring more IO.
  - You do not need to build pyramids on source data if you are *not* going to serve out the source data as Image Service.

Imagery – Building Cache

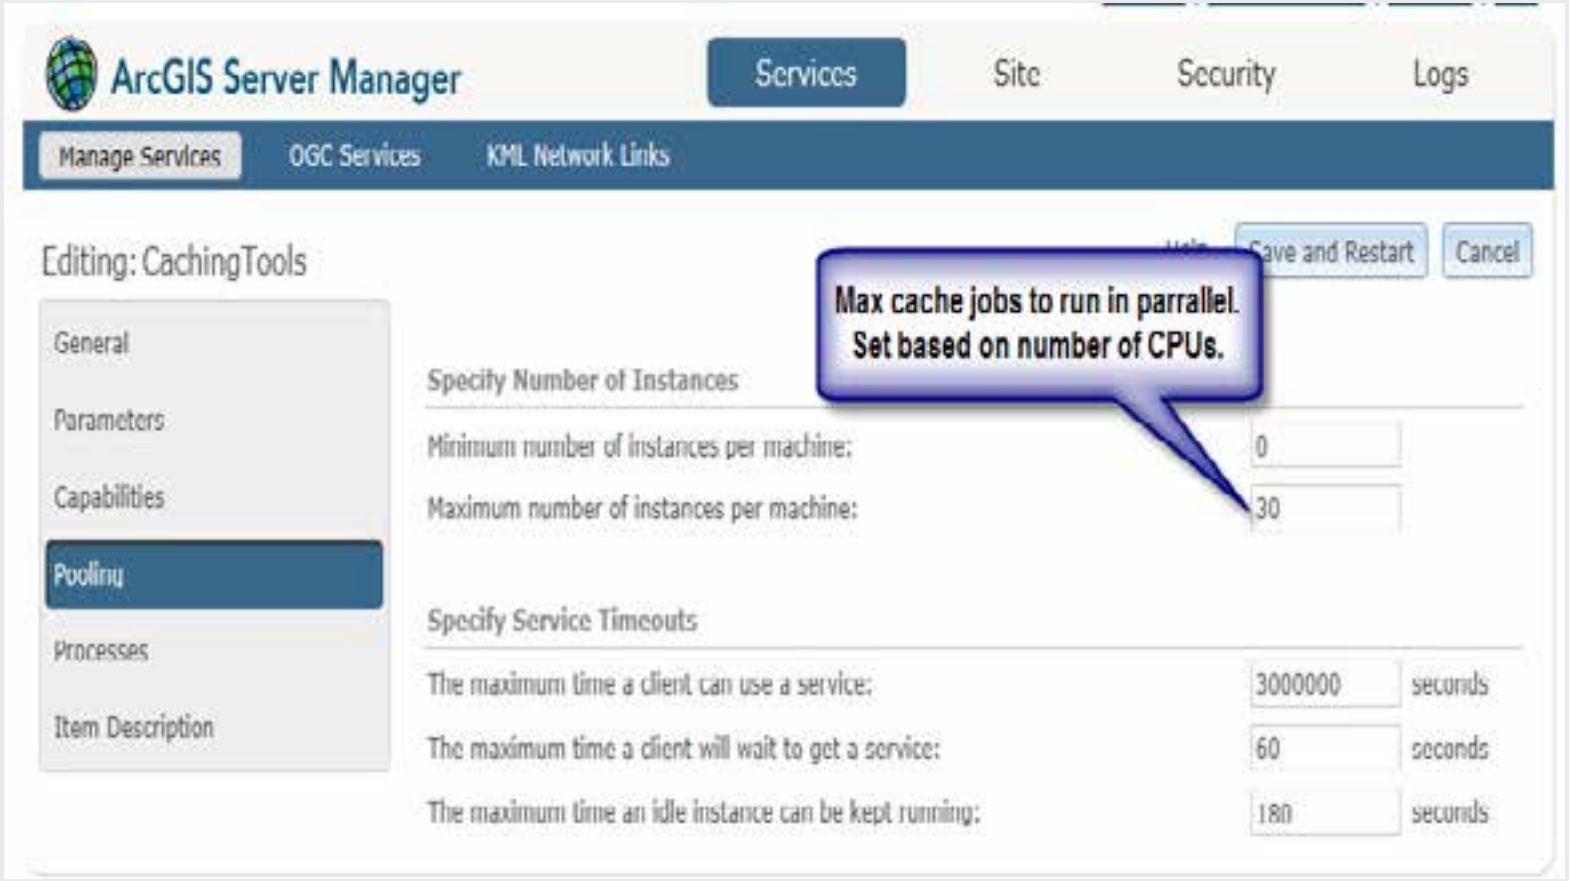

Imagery – Building Cache

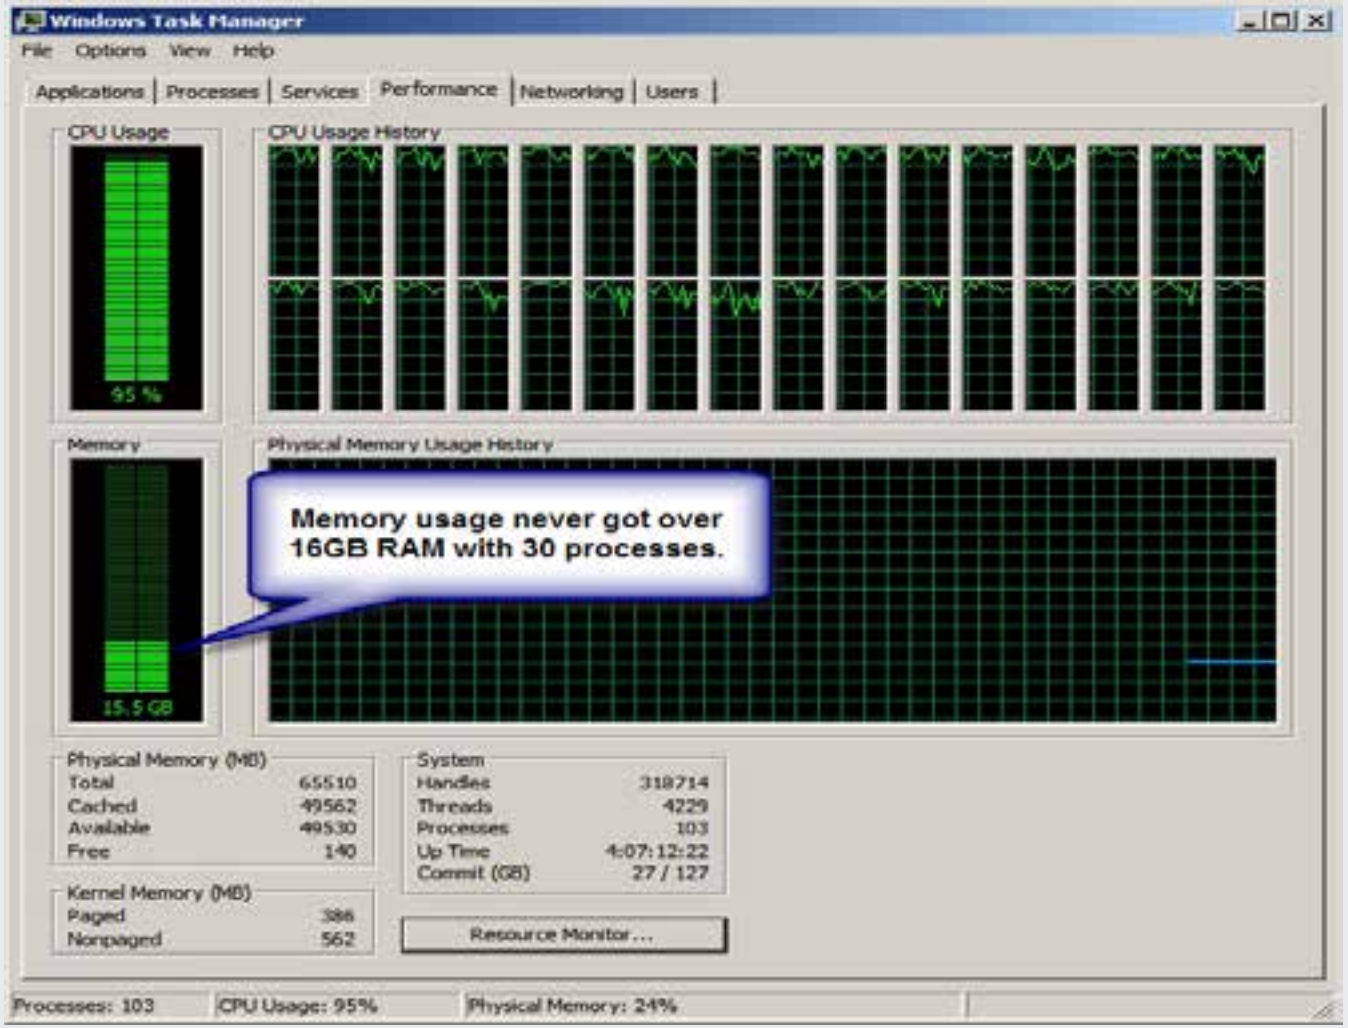

Imagery – Building Cache

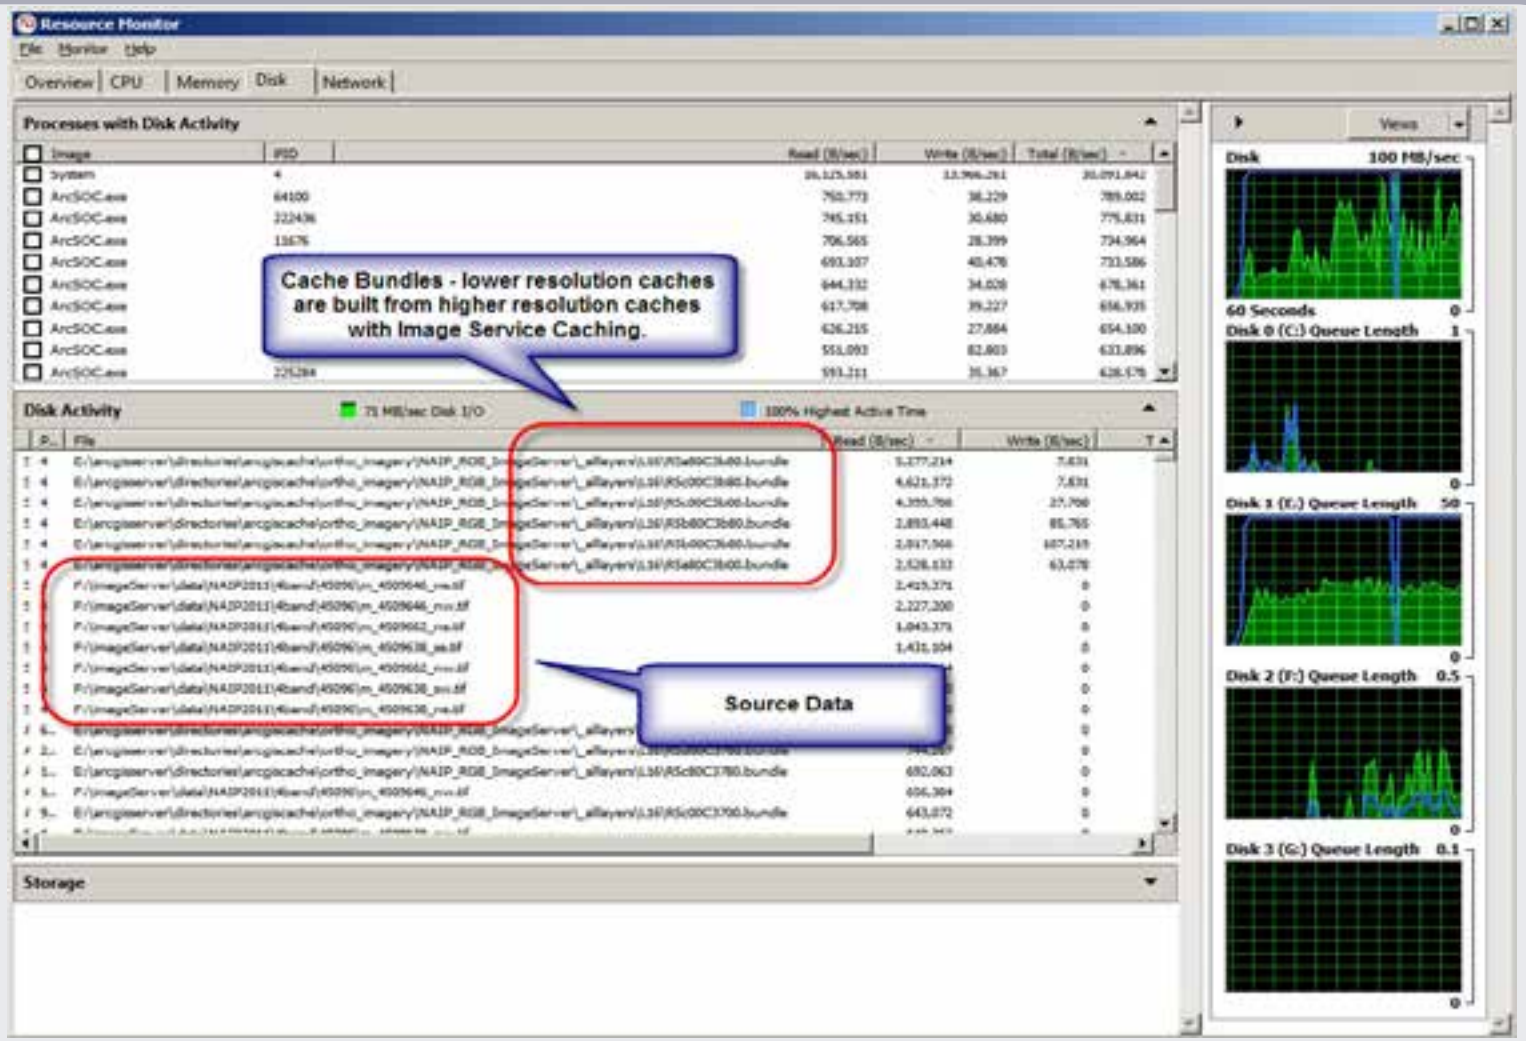

Imagery – Building Cache

- 6 days later

- Holes?
- Because lower level caches are built from higher level cache, it is easier to spot holes. Missing tiles percolate up to the top of the pyramid for Image Services.
- What caused the holes?
  - 108 bad TIF files – not caught by Mosaic Dataset because headers were good, but files were somehow truncated.

Imagery – Building Cache

- Fix holes
  - Ran tiffinfo/python scripts to check all images.
  - Replaced bad imagery – 108 quarter-quads.
  - Re-ran cache for areas with holes.
  - Still had problems in some areas.
  - Used status.gdb to determine bundles with problems. Deleted “suspect” bundles, and re-ran.
  - 1 week to fix all problems.
  - Lesson learned -
    - Check source imagery before building cache.
    - It would have been faster to delete the cache layer and rebuild from scratch.

Imagery – Deploying Cache

- Copy tiles to server.
- Rename directories.
- Create and publish “dummy” MXD.

  At NRCS we chose to only publish the tiles – not the raw imagery - in order to save space. This requires that you create a dummy MXD with a simple vector with dimensions that match the cache tile envelope (CONUS outline).

  The cache tile directory needs to be renamed to match the folder naming structure of the MXD based service.

  Example: ortho_imagery_naip\layers

  ```
  folder-name_service-name\layers
  ```

NAIP – Tile Directory Sizes

Level	Resolution	Scale	Size on Disk
10	1:577,790	152 m	.197 GB
11	1:288,895	76 m	.654 GB
12	1:144,447	38 m	3 GB
13	1:72,223	19 m	9 GB
14	1:36,111	9.5 m	34 GB
15	1:18,055	4.7 m	132 GB
16	1:9,027	2.38 m	473 GB
17	1:4,513	1.19 m	1,510 GB

NAIP Imagery – Color Infra-Red

- Create a Reference Mosaic Dataset
- Add a function to specify band 4, 1, 2

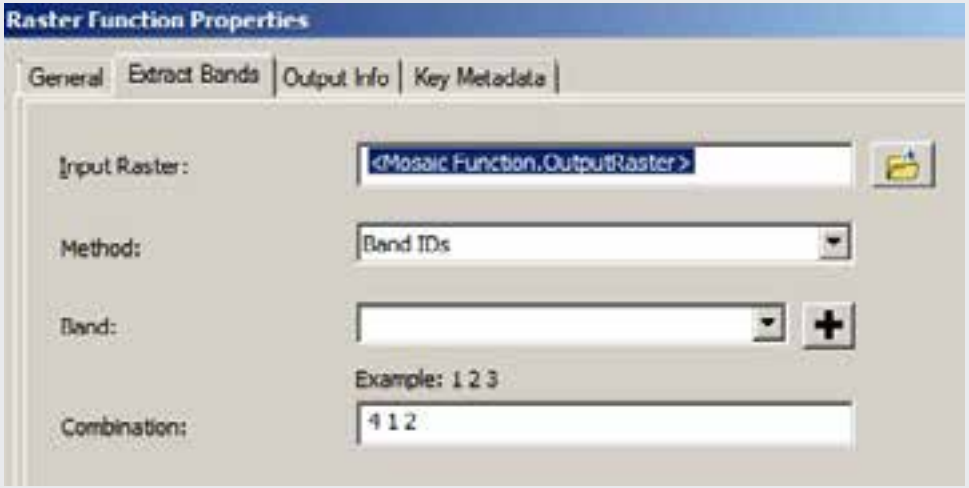

- Publish the Reference Mosaic Dataset
- Build Cache

Imagery – Stewardship Land Imagery – SLI

- Sparsely populated high resolution imagery.
- Remove the NODATA value in the header and use Radiometry tool to fix footprints and adjust boundaries.

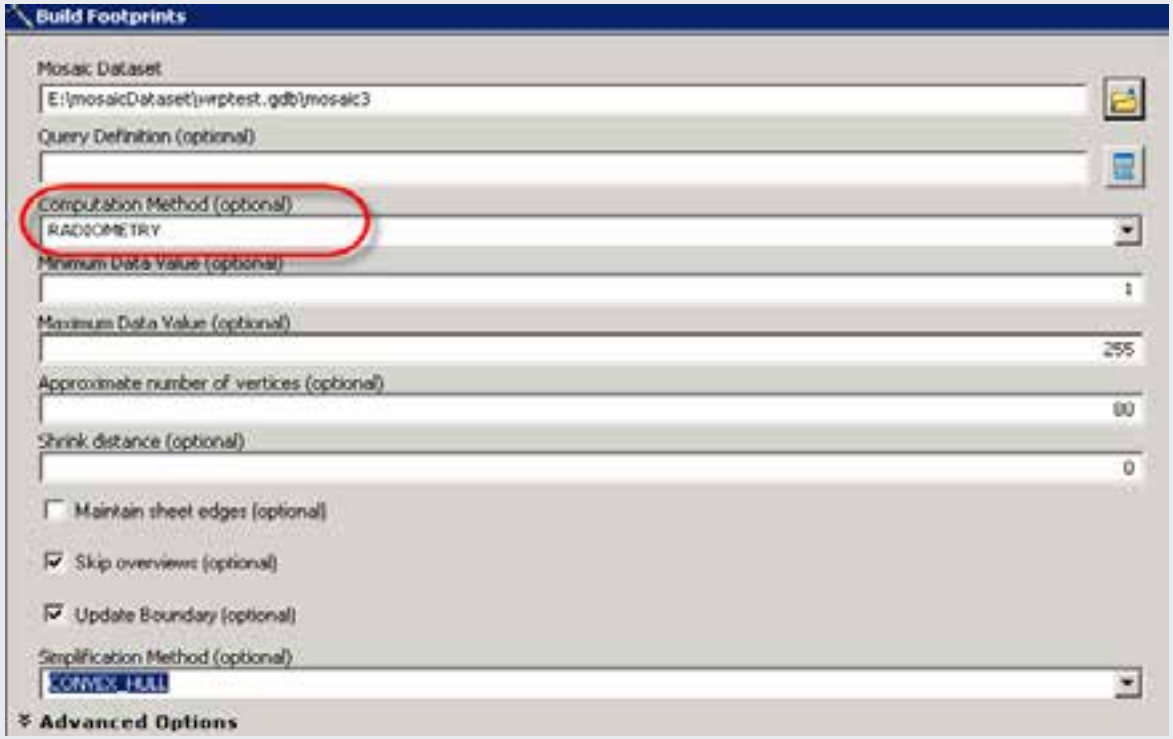

Cache for Vector Layers

- Cache complex vector layers. The more complex, the greater the need for cache tiles. Labeling also has a severe impact on Web Map Services for vector layers.
  - National Hydrography Dataset – very large and very detailed vector layer.
  - Roads – multiple layers with complex labeling.
  - State / County – high usage with complex polygons around the coast and water ways.

Imagery – Deploying Cache

- Why cache State / County boundaries? It is only 30MB.
- High activity (needed by all applications and users).
- More complex than you would think. Generalized layers can be used at low resolution, but detailed features are needed at high resolution. Louisiana has over 500,000 vertices. Server would need to query and render Map Image every time a user panned around Louisiana coast.

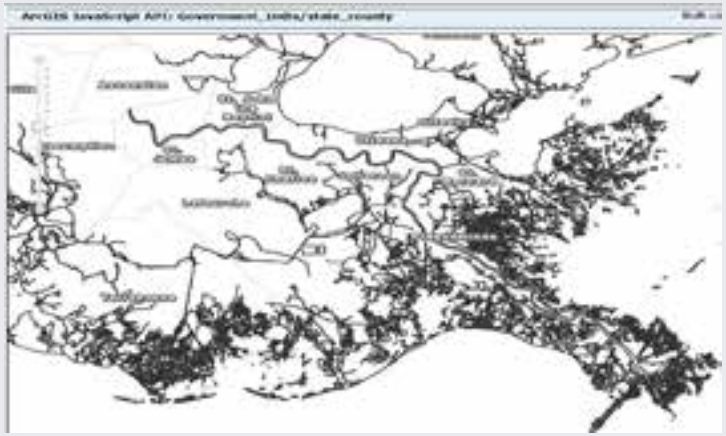

Cache for Vector Layers

- NGCE decided to create transparent vector layers that could be draped over other layers (i.e. imagery).
- Used PNG32. PNG (which is a mix of PNG8 and PNG32) did not look very clean. Note: cache never looks exactly like non-cached services. Test and check.
- Symbology is difficult to design since the layers can be draped over imagery with very different coloration.
- Anti-aliasing and masking helps greatly with labeling.
- Source data was loaded to File Geodatabases.
- Generalized layers were used for lower resolution scales to improve cartography and speed up the caching process.

Cache for Vector Layers

- Symbology and labeling are challenging with vector layers. Issues: color, size, placement, halos. There are also issues with duplicate labels (on tile boundaries).
- Run sample areas at all resolutions.
- View tiles on top of sample imagery – make sure to check against different geographies, i.e. dark green vegetation, desert areas, urban areas and water ways.
- Base layers can be created by building mashups of vector layers over background imagery.

Tools

- ArcGIS Server 10.1 has quite a few tools for managing cache tiles

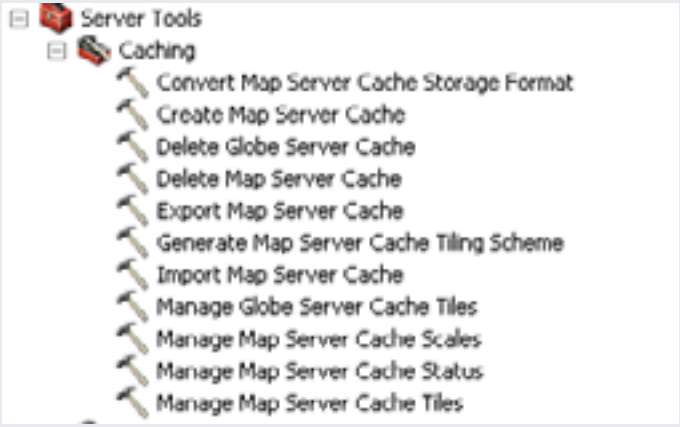

- Caches can be directly viewed in ArcMap and ArcCatalog on the server by opening the cache folder with ArcMap/ArcCatalog.
- There is a geo-processing tool for creating Tile Packages (TPK) for offline access to cache tiles for specific areas and scales.
- The import tool can be used to replace tiles for areas and scales of a cache. Note: bug in export tool in 10.1 – fixed in 10.2.

Summary

- Space requirements are directly proportional to the extent of the dataset and the number of scales. Multiply current size by 4 to determine total size for building an extra scale.
- Only build cache for the area you need and scales to support application and input dataset. (It does not make sense to create 1 meter resolution cache from 2 meter source data).
- Time requirements are directly proportional to the extent of the dataset and the number of scales divided by the number of CPUs. If you do not have a large SMP workstations then use multiple machines in a cluster.
- Use low-end DAS or large internal drives to save on cost of storage.
- Design, test and view at different scales over different backgrounds before committing 2-3 weeks of run time.
- Make sure to clear local cache when checking a cache layer so you can be sure you are looking at the new version of your tiles.
- Allocate time in your schedule to fix problems and to rebuild parts of cache.
- Cache Tiles are well worth the time and disk space!!!!

Questions