



ArcGIS Enterprise Systems: Performance and Scalability

Amr Wahba

awahba@esri.com

**Esri Developer Summit
Middle East & Africa**

19-21 November 2013
Park Hyatt Dubai



Introductions

- **Target audience**
 - GIS, DB, System administrators
 - Testers
 - Architects
 - Developers
 - Project managers
- **Level**
 - Intermediate


Agenda

- **Definitions**
- **Process**
- **Requirements**
- **Performance Factors – Software**
- **Performance Factors – Hardware**
- **Capacity Planning**
- **Performance Tuning**
- **Performance Testing**
- **Using test results as input to Capacity planning**

Definitions

Performance


- The speed at which a given operation occurs
- E.g. Request response time measured in seconds



Definitions

Scalability

- The ability to increase output and maintain acceptable performance
- Examples
 - Capacity 10 maps/sec and response time 1 second
 - Capacity 1000 cars/hrs and speed 55 mph



Definitions

Capacity

- The maximum level of output the system can produce



Definitions

Bottleneck

- Resource(s) limiting the performance or capacity



Low load

Think of :
lanes - as CPU processor
toll booths – as ArcGIS Server instances
cars - as map requests


Definitions

Step Load and Response Time




Definitions

Throughput (request/hr)



Definitions


Resource utilization: CPU, Memory, Network



ArcGIS Enterprise Systems:

Definitions

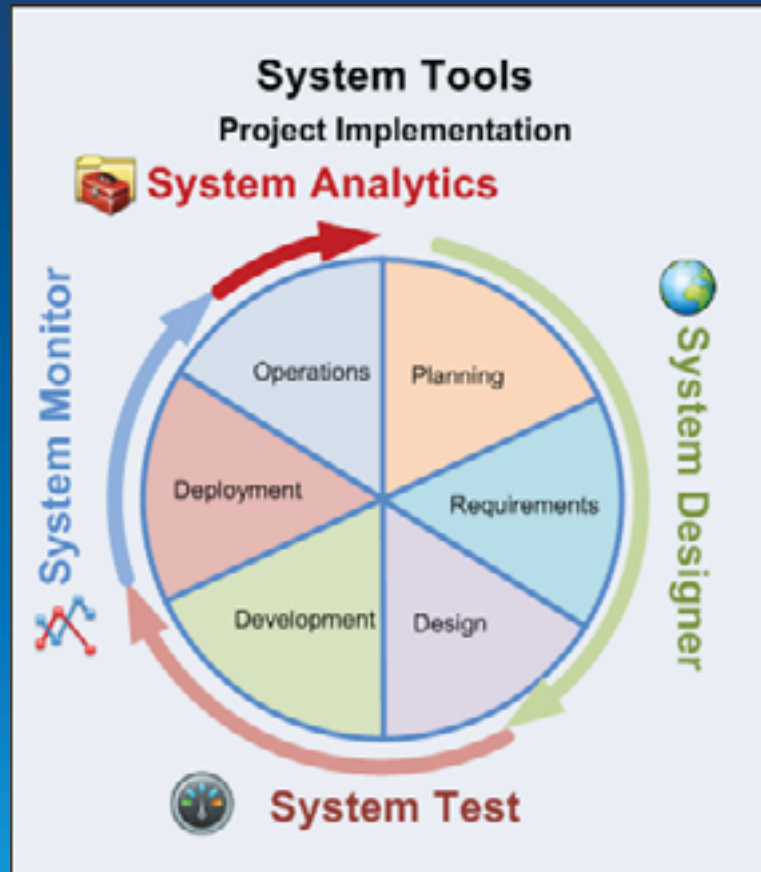
Capacity



Process

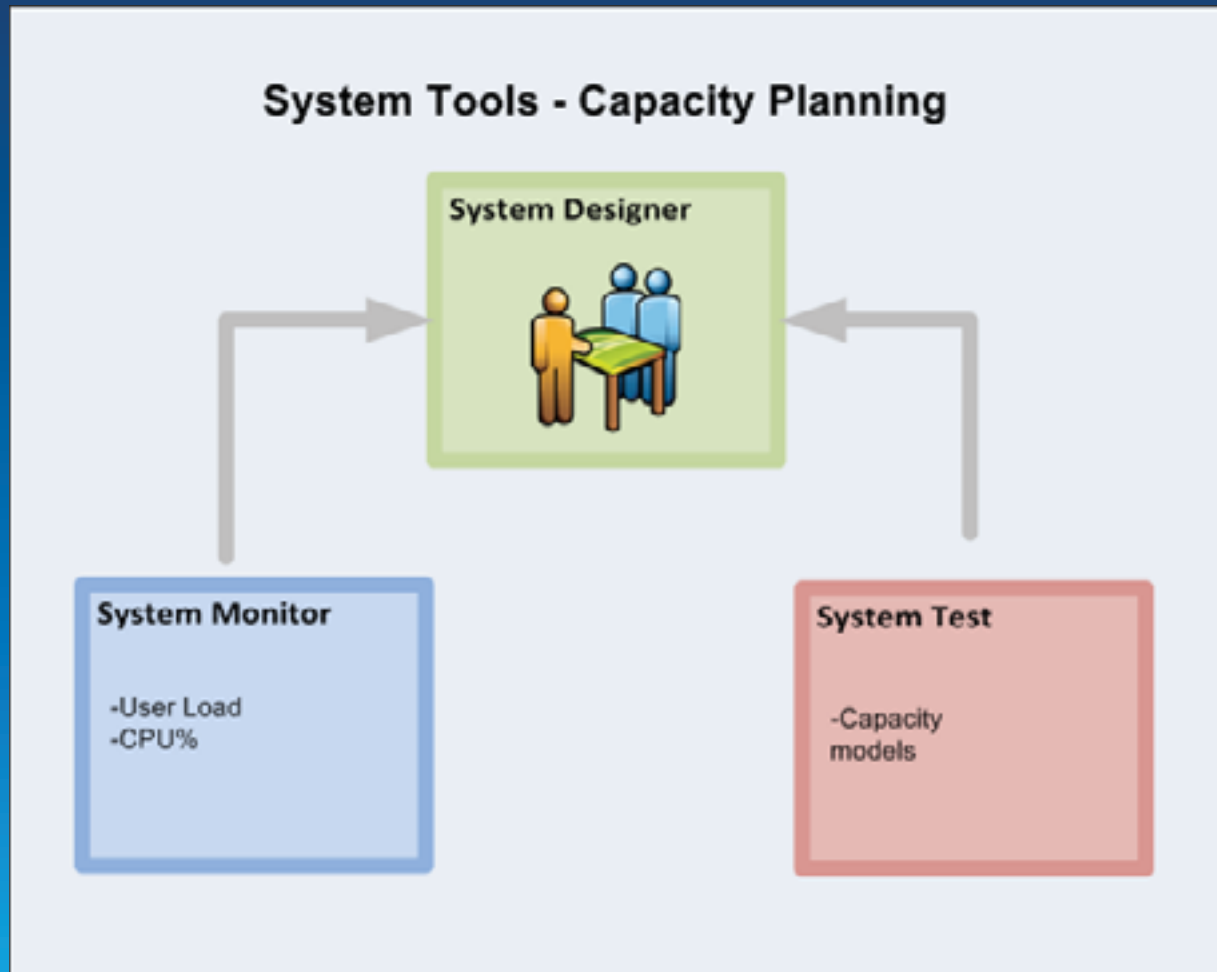
Esri Process and Tools

Holistic approach



Capacity Planning

Tools




Tools

- <http://www.arcgis.com>
- owner:EnterpriseImp
- Show ArcGIS Desktop Content



10 results

Relevance	Title	Owner	Rating	Views	Date
	System Designer				
					
	Open Details				
	<p>A comprehensive tool for planning & designing complete enterprise GIS solutions, including hardware, software, deployment strategy, and capacity forecast.</p> <p> Desktop Application Template by EnterpriseImp Last Modified: July 5, 2013</p> <p>★★★★☆ (2 ratings, 3 comments, 1,186 downloads)</p>				

☒ Show ArcGIS Desktop Content

Tools

- **System Designer**
- <http://www.arcgis.com/home/item.html?id=8ff490eef2794f428bde25b561226bda>.
- <http://www.arcgis.com/home/item.html?id=8e655b38f2fc4b778d07dd34f436a978> (MS 2013)
- **System Test (Beta)**
- <http://www.arcgis.com/home/item.html?id=e8bac3559fd64352b799b6adf5721d81>
- **System Monitor (Beta)**
- <http://www.arcgis.com/home/item.html?id=848f48b0f88e4de7a036377197453efe>
- **System CPU**
- <http://www.arcgis.com/home/item.html?id=3e473b63a3254a6ab5f22e6f9608b209>
- **Network Test**
- <http://www.arcgis.com/home/item.html?id=2b8d4ac8b102453bb0437cdd7a07dd27>
- **Mxdperfstat**
- <http://www.arcgis.com/home/item.html?id=a269d03aa1c840638680e2902dadecac>
- **Database Trace Tools**
- <http://www.arcgis.com/home/item.html?id=24c7b251159149848acc9b81cccb8356>
- **ASlog**
- <http://www.arcgis.com/home/item.html?id=5dfe54f1e9fd48068c4ae0c2c4f459c9>

Requirements Phase


- Performance engineering addresses quality attributes.

Functional Requirements



- Visualization
- Analysis
- Workflow Integration

Quality Attribute Requirements



- Availability
- Performance & Scalability
- Security

Requirements Phase

- **Define System Functions**
 - What are the functions that must be provided?
- **Define System Attributes**
 - Nonfunctional requirements should be explicitly defined.
- **Risk Analysis**
 - An assessment of requirements
 - Intervention step designed to prevent project failure
- **Analyze/Profile Similar Systems**
 - Design patterns
 - Performance ranges



Design Phase- Software

Application Types

- <http://resources.arcgis.com/en/communities/enterprise-gis/01n200000023000000.htm>

Rich Client Applications



Desktop applications that operate in stand-alone, connected, and sometimes connected scenarios.

Web Applications



Browser-based applications that operate in connected scenarios and optionally leverage browser plug-ins.

Services



Standards-based service interfaces that support external applications and systems.

Mobile



Mobile applications that operate in stand-alone, connected, and sometimes connected scenarios.

Design Phase - Software

GIS Services—Map Service

Source document (MXD) optimizations

- **Keeping map symbols simple**
 - Setting scale dependency
 - Avoiding multilayer, calculation-dependent symbols
 - Spatial index
 - Avoiding reprojections on the fly
 - Optimizing map text and labels for performance
 - Using annotations
 - Cost for Maplex and antialiasing
 - Using fast joins (no cross database joins)
 - Avoiding wavelet compression-based raster types (MrSid, JPEG 2000)

Design Phase - Software


Geodatabase

- **Keep delta tables and versioning tree small**
 - Reconcile and post
 - Compress
 - Synchronize replicas
- **Rebuild indexes**
- **Update statistics**

Design Phase - Hardware

Hardware Resources

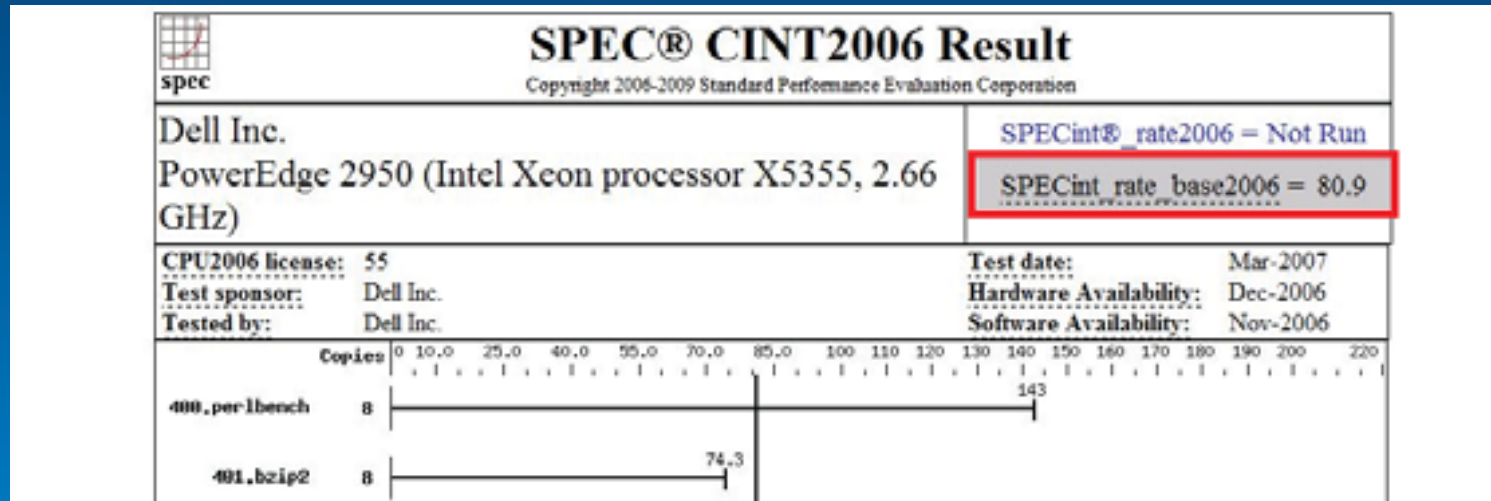
- CPU
- Network bandwidth
 - and latency
- Memory
- Disk



Most well-configured and tuned GIS systems are processor bound.

Design Phase - Hardware

CPU Processor Speed – Specrate.org



http://www.cpubenchmark.net/cpu_list.php

Design Phase - Performance Factors

Hardware Resources—Memory

Item	Low	High	Delta
XenApp Session	500 MB	1.2 GB	140%
Database Session	10 MB	75 MB	650%
Database Cache	200 MB	200 GB	99,900%
SOC Process (Dynamic Map Service)	50 MB	500 MB	900%
SOC Process (Image Service)	20 MB	1,024 MB	5,020%
SOC Process (Geoprocessing Service)	100 MB	2,000 MB	1,900%
SOM	30 MB	70 MB	133%

Wide ranges of memory consumptions

Design Phase - Hardware

Virtualization


- **Performance depends on configuration and implementation**
 - 5-30%

Overburdened VMs will incur significant performance degradation


Design Phase - Hardware

Network


1. Distance



2. Payload



3. Infrastructure



Design Phase - Hardware

Hardware Resources—Network

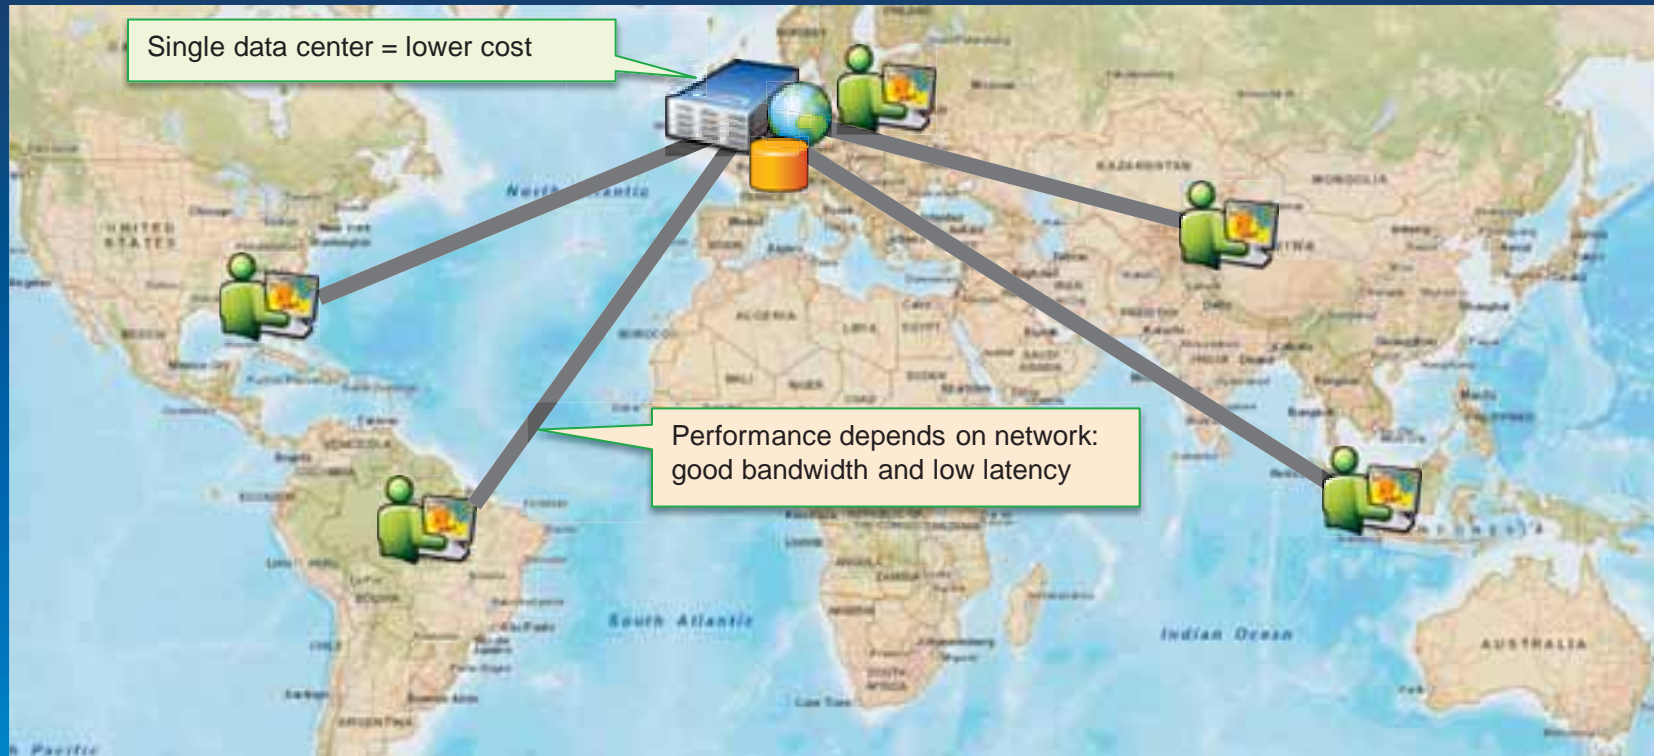
- **Impact of service and return type on network transport time**
 - **Compression**
 - **Content (e.g., Vector vs. Raster)**
 - **Return type (e.g., JPEG vs. PNG)**

					Network Traffic Transport Time (sec)					
					56 kbps	1.54 Mbps	10 Mbps	45 Mbps	100 Mbps	1 Gbps
Application Type	Service/Op	Content	Return Type	Mb/Tr	0.056	1.540	10.000	45.000	100.000	1000.000
ArcGIS Desktop	Map	Vector		10	178.571	6.494	1.000	0.222	0.100	0.010
Citrix/ArcGIS	Map	Vector+Image	ICA Comp	1	17.857	0.649	0.100	0.022	0.010	0.001
Citrix/ArcGIS	Map	Vector	ICA Comp	0.3	5.357	0.195	0.030	0.007	0.003	0.000
ArcGIS Server	Map	Vector	PNG	1.5	26.786	0.974	0.150	0.033	0.015	0.002
ArcGIS Server	Image		JPG	0.3	5.357	0.195	0.030	0.007	0.003	0.000
ArcGIS Server	Map Cache	Vector	PNG	0.1	1.786	0.065	0.010	0.002	0.001	0.000
ArcGIS Server	Map Cache	Vector+Image	JPG	0.3	5.357	0.195	0.030	0.007	0.003	0.000

All built into System Designer

Network Testing

Important for cloud based solutions



Cloud vendors offer multiple location. Which one is the best for your global solution?


Demo

Network Speed Test Tool:





Design Phase—Capacity Planning

- Uncertainty of input information



Design Phase—Capacity Planning

- Uncertainty of input information—Planning hour
- Identify the Peak Planning Hour (most cases)



Design Phase—Capacity Planning

- CPU Factors

1. User load: Concurrent users or throughput
2. Operation CPU service time (model)—performance
3. CPU SpecRate

$$\# CPU_t = \frac{ST_b \times TH_t \times 100}{3600 \times \%CPU_t} \times \frac{SpecRatePerCPU_b}{SpecRatePerCPU_t}$$

subscript t = target

subscript b = benchmark

ST = CPU service time

TH = throughput

%CPU = percent CPU

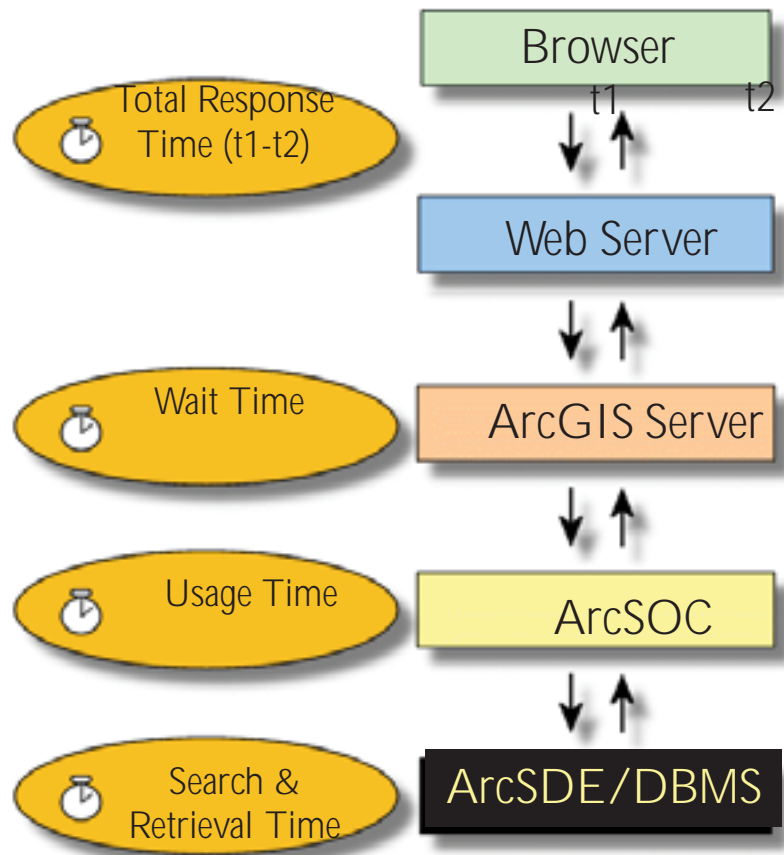
Demo

System Designer

Development phase Tuning: A reproducible test cases


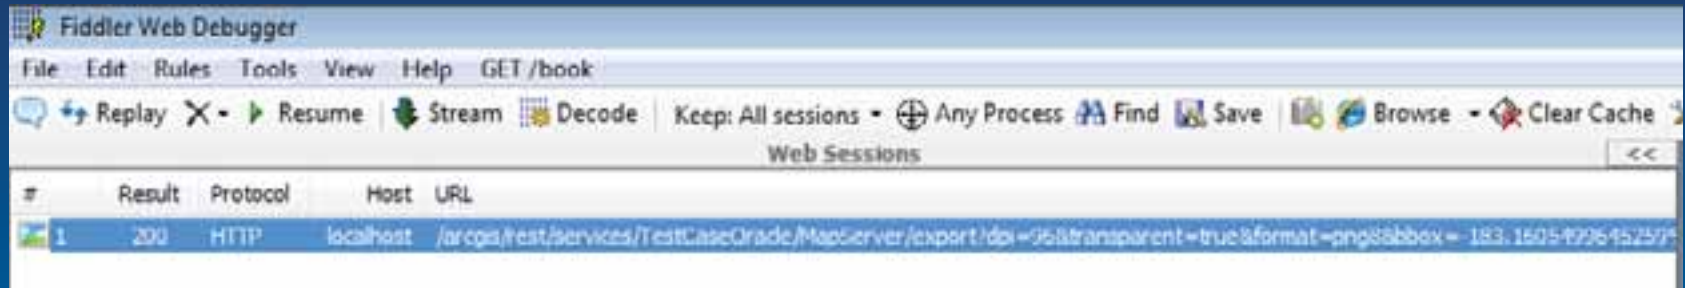
Tuning methodology

Profile each tier starting from the top



Fiddler

Fiddler measurement approximately 5.2 seconds



Server Logs

- *Analyze ArcGIS for Server Statistics & logs*

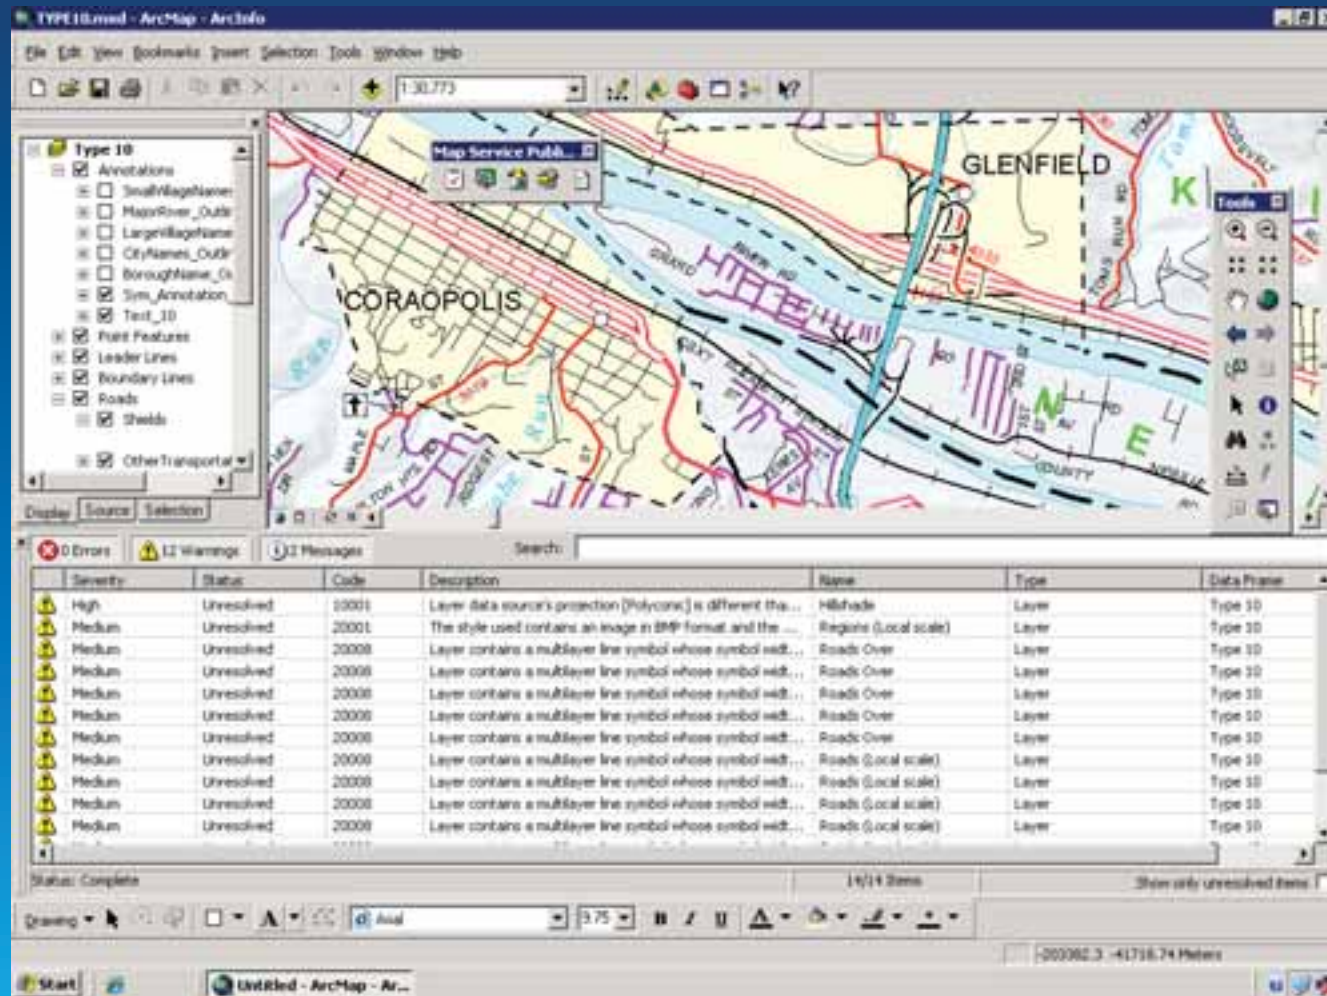
ArcGIS Server Administrator Directory	
Home > services > Riverside_BaseMap.MapServer > statistics	
Service Statistics - Riverside_BaseMap (MapServer)	
Summary	
Busy instances:	0
Free instances:	1
Initializing instances:	0
Not created instances:	1
Maximum instances:	2
Number of transactions:	1
Total busy time (in milliseconds):	38619
Machine: ASAKOWICZ.ESRI.COM	
Busy instances:	0
Free instances:	1
Initializing instances:	0
Not created instances:	1
Maximum instances:	2
Number of transactions:	1
Total busy time (in milliseconds):	38619
Supported Interfaces: REST RSS ATOM	

<Msg time="2009-03-16T12:23:22" type="INFO3" code="103021" target="Portland.MapServer" methodName="FeatureLayer.Draw" machine="myWebServer" process="2836" thread="3916" elapsed="0.05221">Executing query.</Msg>

<Msg time="2009-03-16T12:23:23" type="INFO3" code="103019" target="Portland.MapServer" methodName="SimpleRenderer.Draw" machine="myWebServer" process="2836" thread="3916">Feature count: 27590</Msg>

<Msg time="2009-03-16T12:23:23" type="INFO3" code="103001" target="Portland.MapServer" methodName="Map.Draw" machine="myWebServer" process="2836" thread="3916" elapsed="0.67125">End of layer draw: STREETS</Msg>

ArcMap Publish Tool

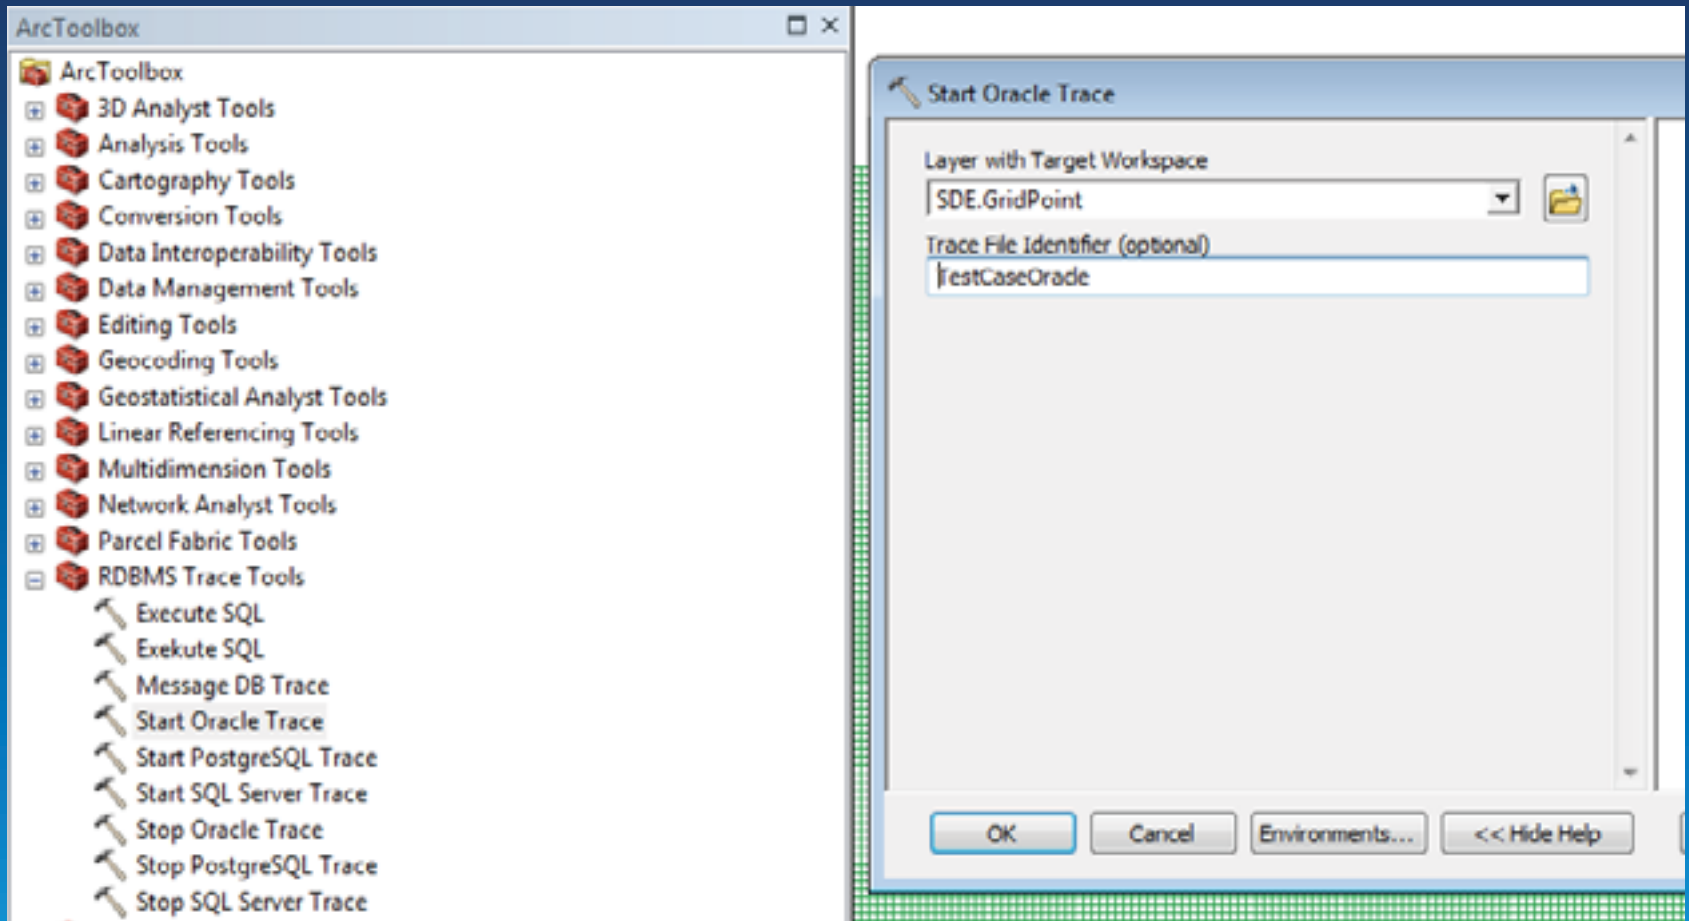


Mxdperfstat

Item	At Scale	Layer Name	Refresh Time (sec)	Recommendations	Features	Vertices	Labeling	Geography Phase (sec)	Graphics Phase (sec)	Cursor Phase (sec)	DBMS CPU	DBMS LIO
1	167,935,665	SDE.GridPoint	4.75	run DBMS trace: oraCPU=4.74; run DBMS trace, check oracle execution plan: oraLIO=130936; check if index exist for query def attributes;	1,998		False	4.74	.00	4.56	4.74	130,936

DBMS LIO	DBMS PIO	Source	LayerType	Layer Spatial Reference	LayerQueryDef
130,936		esriDBMS_Oracle,asakowicz,sde:oracle\$asakowicz:1521/gis2,sde	esriGeometryPoint	GCS_WGS_1984	ID<1000

DBMS Trace



Oracle Trace

SQL ID : 71py6481sj3xu

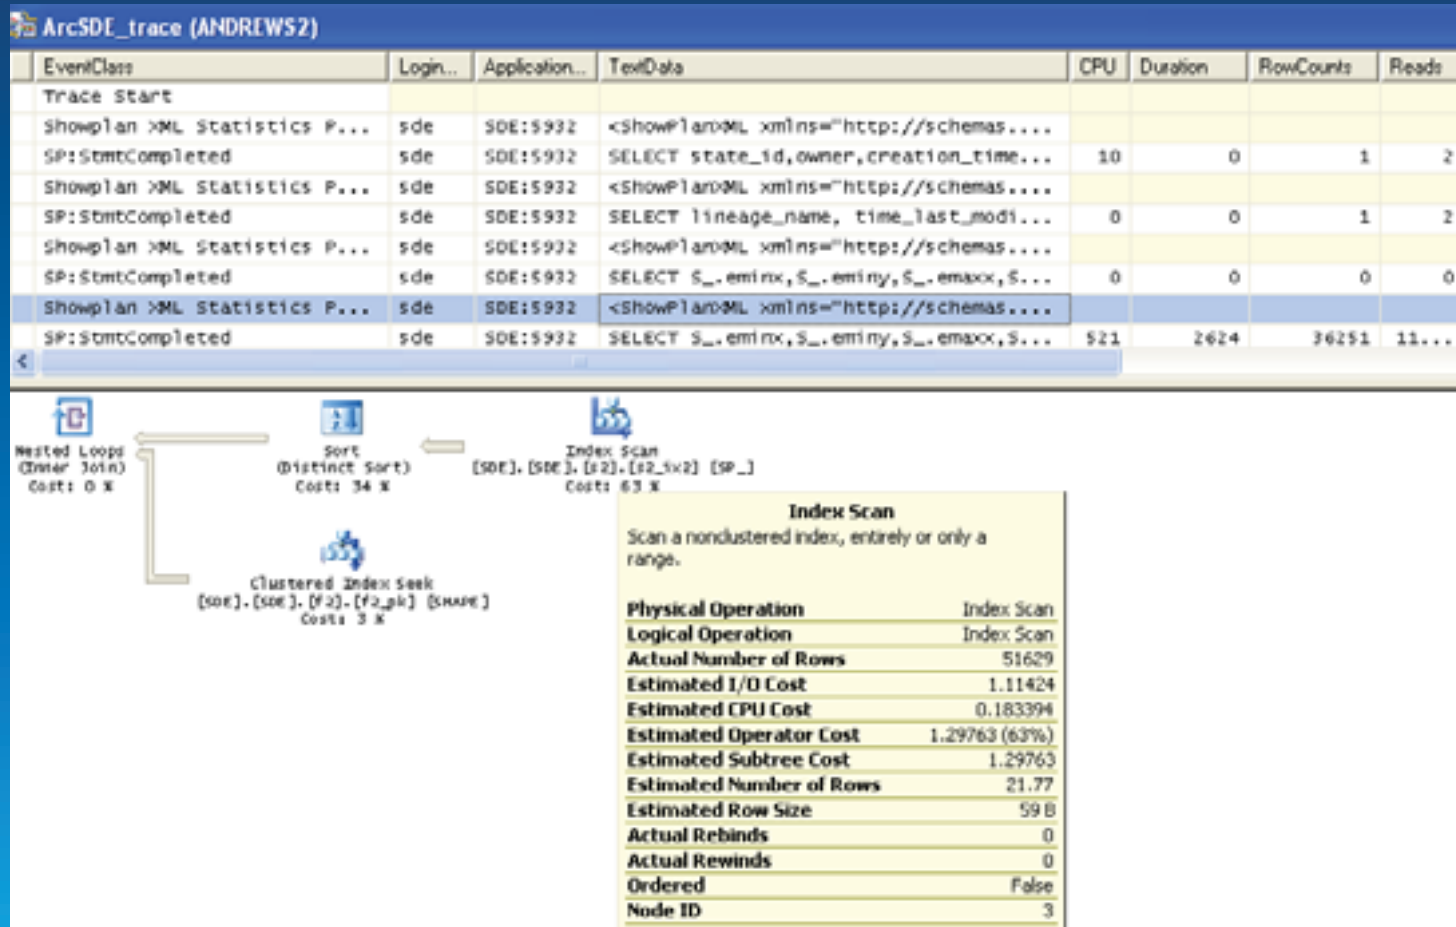
```
SELECT 1 SHAPE, TAXLOTS.OBJECTID, TAXLOTS.SHAPE.points, TAXLOTS.SHAPE.numpts,
TAXLOTS.SHAPE.entity, TAXLOTS.SHAPE.minx, TAXLOTS.SHAPE.miny,
TAXLOTS.SHAPE.maxx, TAXLOTS.SHAPE.maxy, TAXLOTS.rowid
FROM SDE.TAXLOTS TAXLOTS WHERE SDE.ST_EnvIntersects(TAXLOTS.SHAPE, :1, :2, :3, :4) = 1
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	0	0.00	0.00	0	0	0	0
Execute	1	0.07	0.59	115	1734	0	0
Fetch	242	0.78	12.42	2291	26820	0	24175
total	243	0.85	13.02	2406	28554	0	24175

Elapsed times include waiting on following events:

Event waited on	Times Waited	Max. Wait	Total Waited
SQL*Net message to client	242	0.00	0.00
db file sequential read	2291	0.39	11.69
SQL*Net more data to client	355	0.00	0.02
SQL*Net message from client	242	0.03	0.54

SQL Profiler



Performance Testing

Testing Objectives


- **Meet Service-Level Agreement (SLA)**
- **Bottlenecks analysis**
- **Capacity planning**
- **Benchmarking different alternatives**

Testing Tools

Test Tools	Open Source	Pros	Cons
LoadRunner	No	<ul style="list-style-type: none"> • Industry Leader • Automatic negative correlations identified with service-level agreements • HTTP web testing • Click and script • Very good tools for testing SOA • Test results stored in database • Thick client testing • Can be used for bottleneck analysis 	<ul style="list-style-type: none"> • High cost • Test development in C programming language • Test metrics difficult to manage and correlate • Poor user community with few available examples
Silk Performer	No	<ul style="list-style-type: none"> • Good solution for testing Citrix • Wizard-driven interface guides the user • Can be used for bottleneck analysis 	<ul style="list-style-type: none"> • Moderate to high cost • Test metrics are poor. • Test development uses proprietary language. • Test metrics difficult to manage and correlate • Poor user community with few available examples
Visual Studio Test Team	No	<ul style="list-style-type: none"> • Low to moderate cost • Excellent test metric reporting • Test scripting in C# or VB .NET • Unit and web testing available • Blog support with good examples • Very good for bottleneck analysis 	<ul style="list-style-type: none"> • No built-in support for AMF • No thick-client options • Moderate user community
JMeter	Yes	<ul style="list-style-type: none"> • Free • Tool 	<ul style="list-style-type: none"> • Provides only response times • Poor user community with few available examples

Testing

- Test Data



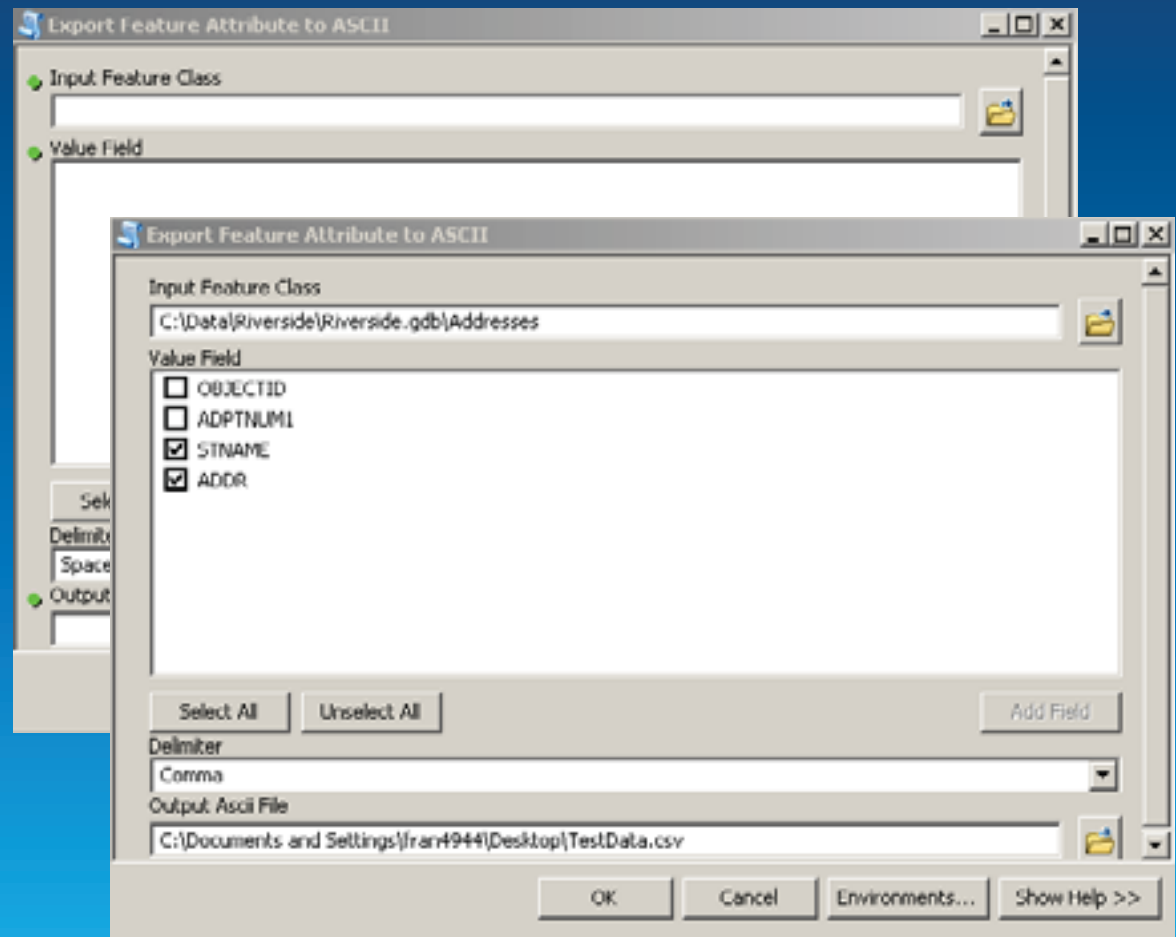
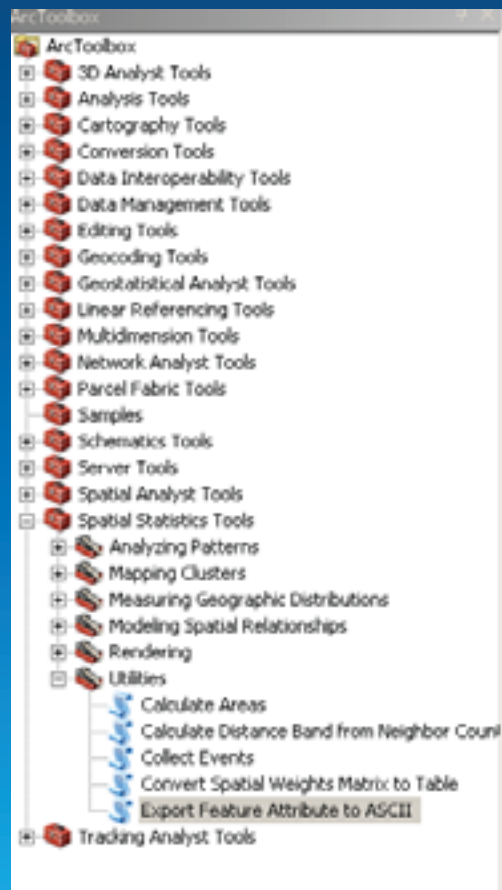
Area of Interest

QueryString	
Name	Value
l	image
dpi	96
transparent	true
format	png8
bbox	("xmin":6219593.737018972,"ymin":2303862.765275147,"xmax":6231478.566248391,"ymax":2311468.277924415,"spatialReference":{"wkid":2230})
bboxSR	2230
imageSR	2230
size	1222,782

Selected Extent
From HTTP
Debugging
Proxy

Testing

Attribute Data



Testing


- Test Scripts
 - Record user workflow based on application user requirements.
 - Create single user web test.
 - Define transactions.
 - Set think time and pacing based on application user requirements.
 - Parameterize transaction inputs.
 - Verify test script with single user.

Development Phase—Testing

- Load Test
 - Create load test.
 - Define user load.
 - Max users
 - Step interval and duration
 - Create machine counters to gather raw data for analysis.
 - Execute.

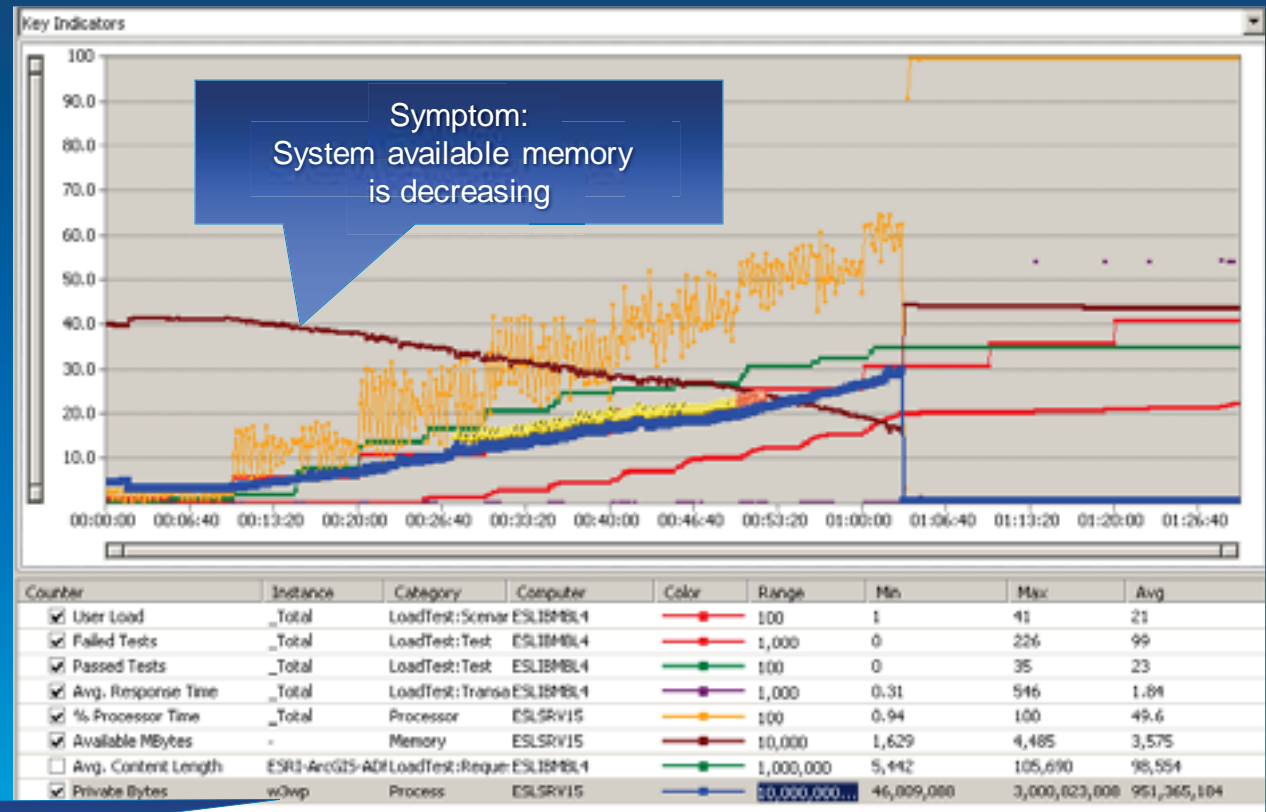
Testing

- Analysis—Compare and correlate key measurements
 - Expected counters correlation: increasing user load, CPU utilization, response time



Testing

- Analysis—Compare and correlate key measurements
- Memory leak example



Root cause:
Web Server process

Demo

System Test –
Enterprise GIS testing tool


Capacity planning: using test results

Test Results as Input into Capacity Planning

- **Load Test Results - Riverside Electric**
 - **Baseline Test with Single Thread**
 - Note* Service Time is Load Independent
 - **Think Time=0**
 - **Evaluate Key Metrics**
 - Throughput
 - Response Time
 - QA Check
 - **Evaluate System Under Test**
 - CPU, Network, Memory, and Disk

Test Results as Input into Capacity Planning

- Load Test Results - Key Indicators




Test Results as Input into Capacity Planning

- Load Test Results - System Metrics



Test Results as Input into Capacity Planning

- Load Test Results – input into capacity models
 - Average throughput over the test duration
 - 3.89 request/sec ~ 14,004 request/hour
 - Average response time over the test duration
 - .25 seconds
 - Average CPU Utilization
 - 20.8%
 - Mb/request = 1.25 Mb



Test Results as Input into Capacity Planning


- Load Test Results – input into CPU capacity model
 - Input from testing
 - #CPUs = 4 cores
 - %CPU = 20.8
 - TH = 14,004 requests/hour
 - SPEC per Core of machine tested = 35
 - $ST = (4 \times 3600 \times 20.8) / (14,004 \times 100) = 0.2138 \text{ sec}$
 - Note* very close to Average response time of .25

$$ST = \frac{\#CPU \times 3600 \times \%CPU}{TH \times 100}$$

Test Results as Input into Capacity Planning

- Target values

1. Server SpecRate/core=10.1



2. User load=30,000 req/hr
3. Network=45 Mbps

Test Results as Input into Capacity Planning

- System Designer
 - Input:
 - Throughput=30000
 - ST=0.21
 - Mb/tr=1.25
 - Hardware=80.9 Spec

User Workflow Dialog

Select Site: Client Workflow Name: Web request

Select a Workflow: Web request Total Users: 0 Active Users: 0 Workflow Facing (sec): 0

Workflow Throughput (Workflow/hr): 30000 [Configured]

Workflow Operations: Select Operation: Web request

Operation Name: RTMax (sec): 0 Stated Delay (sec): 0

Avg Think Time (sec): 1 Occurrence: 1 [Configured]

Notes:

Save Discard

Model Review									
Model Assigned									
Selected Model									
Service Type: Map Model Name: ExportMap REST MapService 9.3.1									
Model	Function	Tier	Modified	Service Time(sec)	Queue Time(sec)	CPU Cores Calc	Modified	Mb/Tr	M
✓	Client	Client	○		0.000	0.00	○	1.250	
✓	WebService	Web Services	☑	0.210	0.004	1.01	☑	1.250	

Test Results as Input into Capacity Planning

- System Designer
 - Input
 - Hardware=80.9 Spec

The screenshot shows the 'Hardware Dialog' window with the following configuration:

- Site: Server
- Role: WebServer
- Category: Server
- Switch: Switch0
- Vendor: Dell Inc.
- Processor: [Empty]
- Operating System: Windows Server 2008 R2 64-bit
- Select Hardware Item: WebServer
- [Cores] [SPEC/Core] Hardware: [8] [10.11] PowerEdge 2950 (Intel Xeon processor X5355, 2.66 GHz)
- Processor Name: Intel Xeon X5355
- Processor Speed (MHz): 2666.00
- CPU Cores: 8
- SPEC int rate per Core: 10.11 (circled in red)
- SPEC int rate: 80.90 (circled in red)
- %Max Utilization: 80
- RAM (GB): 0
- Storage (GB): 0
- Platform Virtualization: ☐ Virtual Environment
- Vendor: [Empty]
- CPU Cores Allocated: [Empty]
- Notes: [Empty text area]

Buttons at the bottom: + Add, - Delete, Configured, Save, Discard, Close.

Test Results as Input into Capacity Planning

- System Designer
 - Review results

Model Review Model Assigned										
Selected Model										
Service Type: Map Model Name: ExportMap REST MapService 9.3.1										
Model	Function	Tier	Modified	Service Time(sec)	Queue Time(sec)	CPU Cores Calc	Modified	Mb/Tr	Mbps Calc	Transport(sec)
✓	Client	Client	⊖		0.000	0.00	⊖	1.250	10.42	0.00
✓	WebService	Web Services	✓	0.210	0.046	7.57	✓	1.250	10.42	0.04



Understanding our world.