Esri International Developer Summit

Palm Springs, CA

ArcGIS for Server Performance and Scalability: Testing Methodologies

Andrew Sakowicz, asakowicz@esri.com

Frank Pizzi, fpizzi@esri.com

Introductions

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

Agenda

Performance testing

- Definitions
- Process
- Factors
- Tuning
- Monitoring
- Testing

Definitions

Performance

- Speed, e.g. response time (seconds)

Scalability

- **The ability to increase output and maintain acceptable performance**

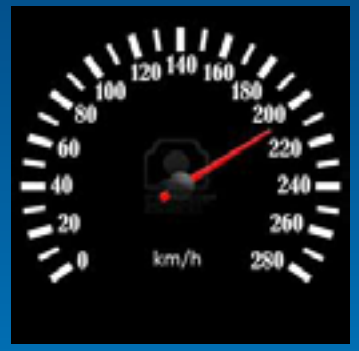

Capacity

- The maximum level of output the system can produce, e.g.
- X cars/sec
- X maps/sec

At capacity

Over capacity

Bottleneck

- Resource(s) limiting the performance or capacity, e.g. CPU, connections

No bottleneck

bottleneck

Think of :
Lanes -as CPU processor
Toll -as ArcGIS Server instances
Cars -as map requests

Step Load and Response Time

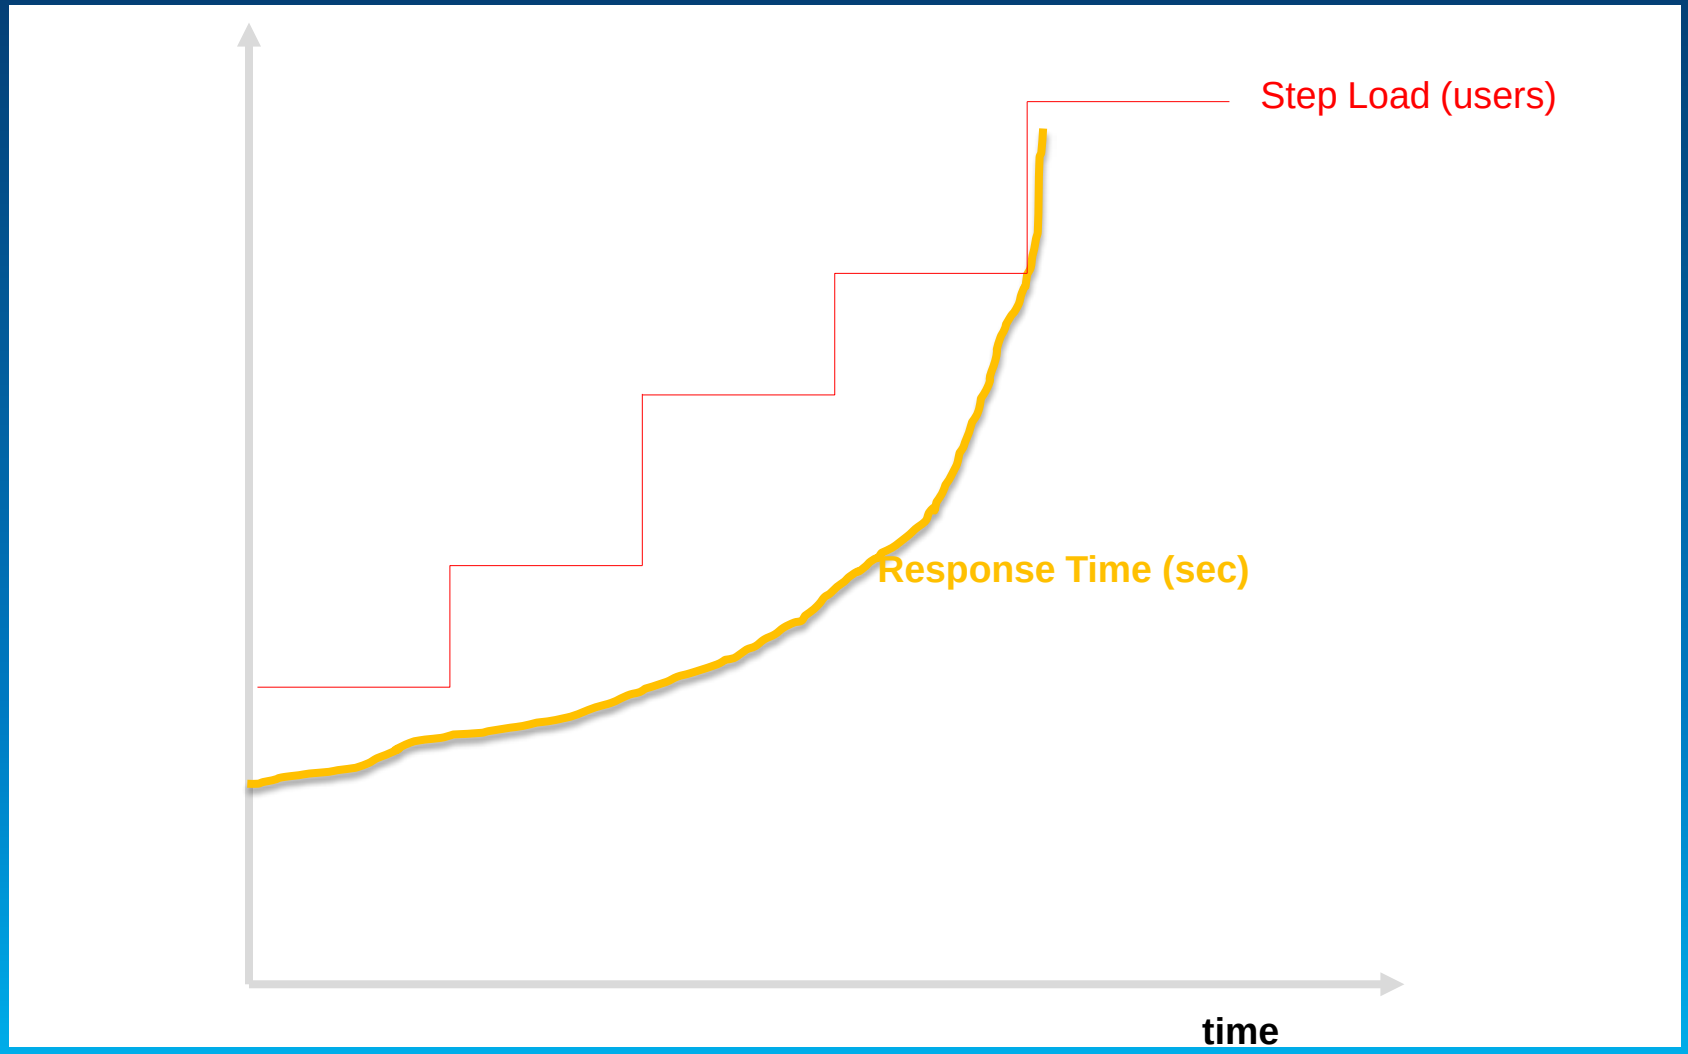

Throughput (request/hr)

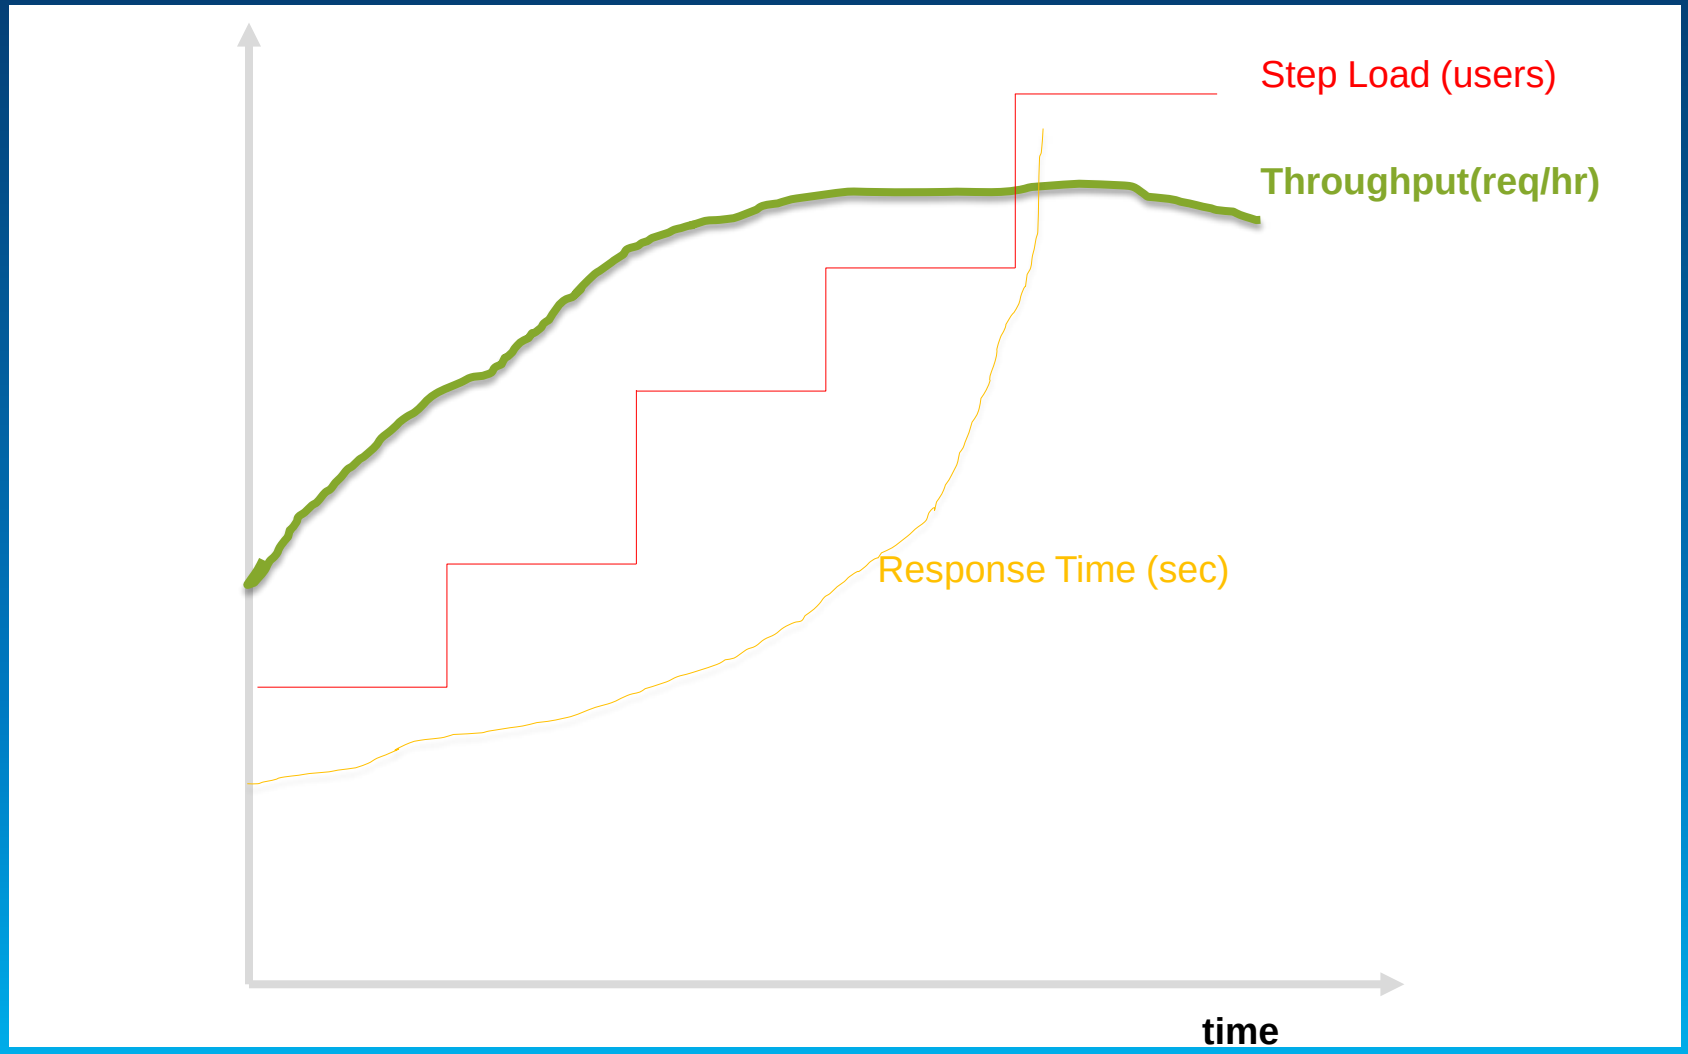

Resource utilization: CPU, Memory, Network

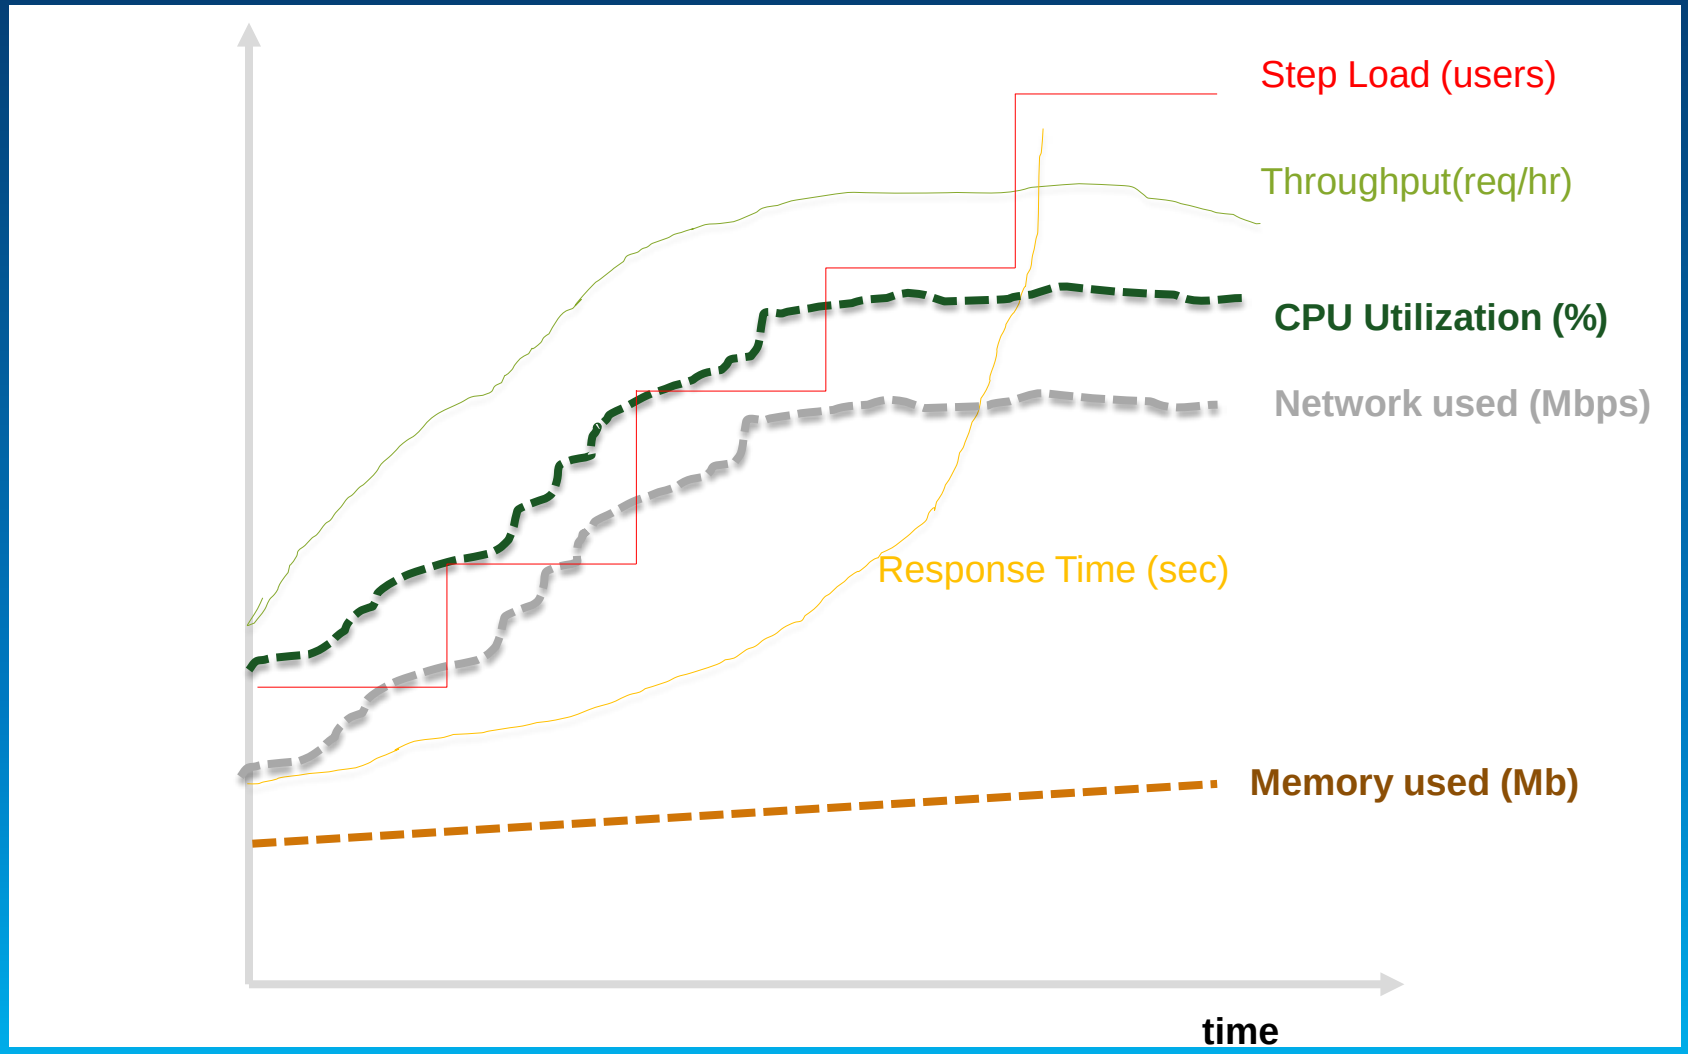

Capacity

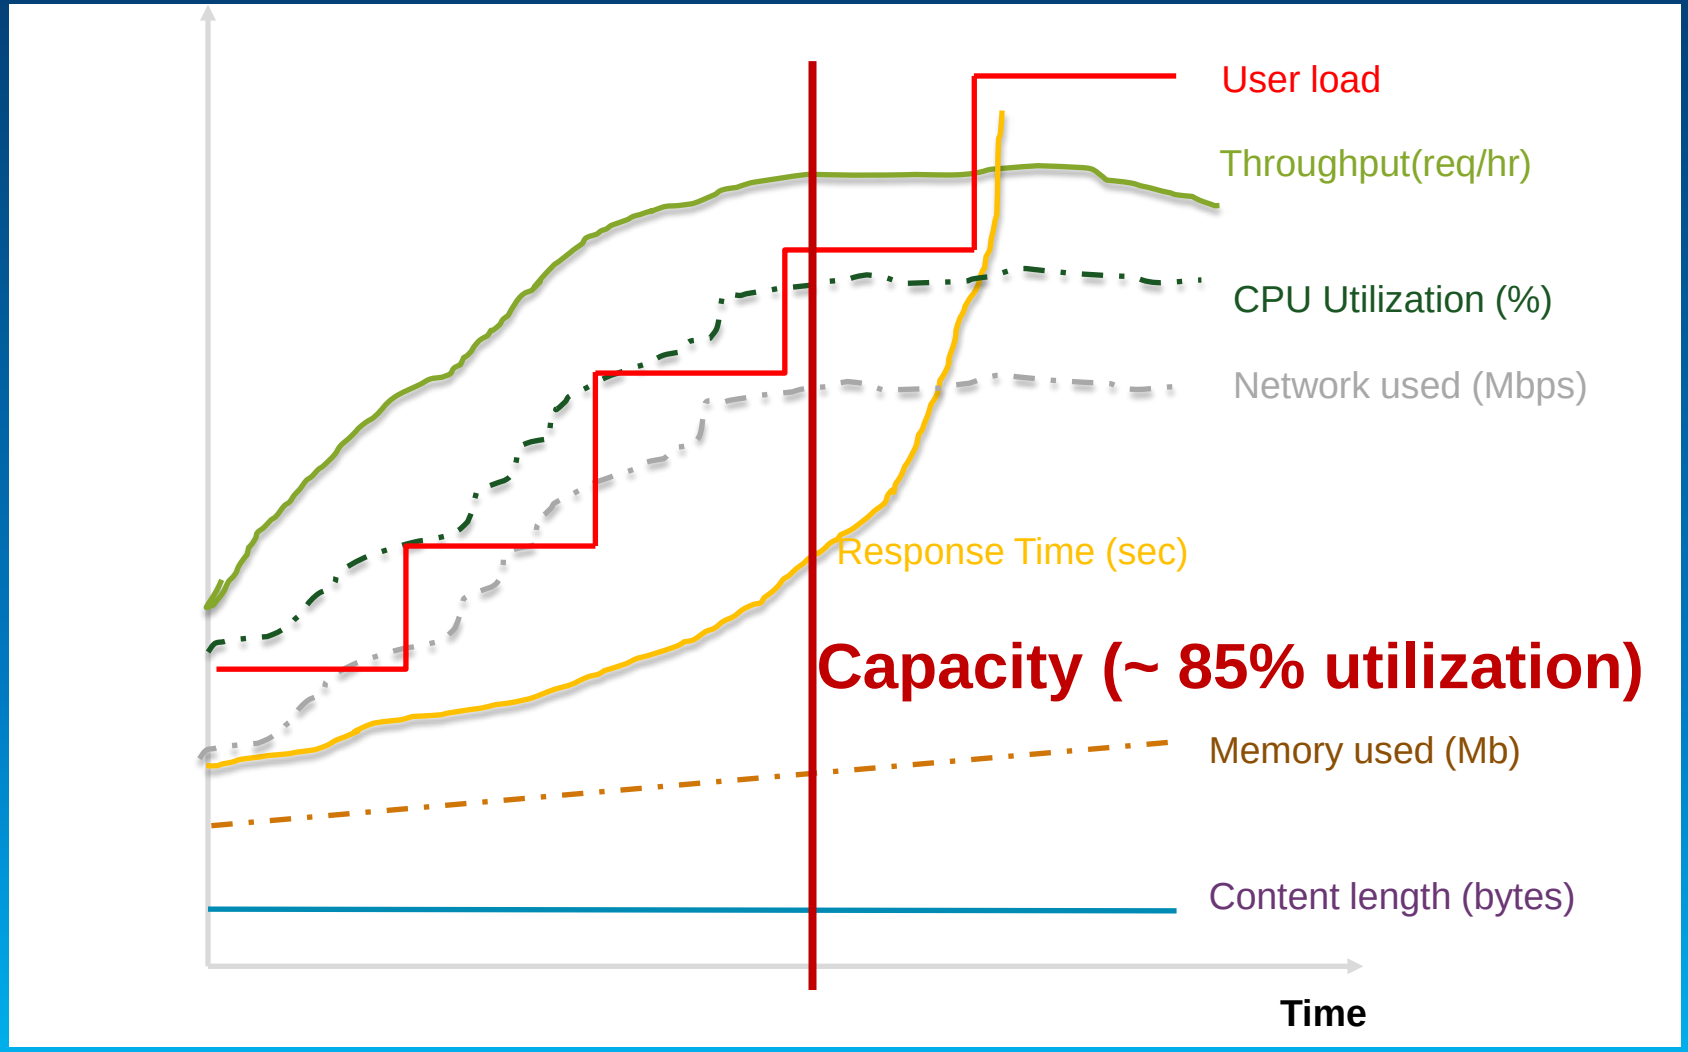
User load
Throughput(req/hr)
CPU Utilization (%)
Network used (Mbps)
Response Time (sec)
Capacity (~ 85% utilization)
Memory used (Mb)
Content length (bytes)
Time

Process

Esri Services Process and Tools

Holistic approach

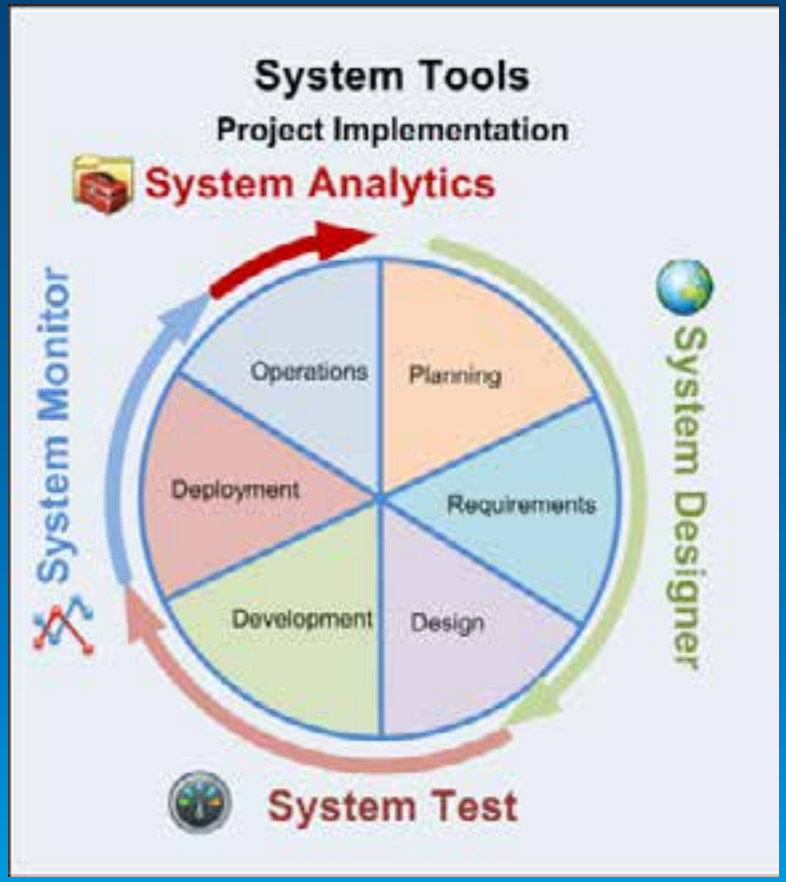

Tools

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

Tools

- System Designer
- http://www.arcgis.com/home/item.html?id=8ff490eef2794f428bde25b561226bda.
- System Test
- http://www.arcgis.com/home/item.html?id=e8bac3559fd64352b799b6adf5721d81
- System Monitor
- http://www.arcgis.com/home/item.html?id=848f48b0f88e4de7a036377197453efe
- System CPU
- http://www.arcgis.com/home/item.html?id=3e473b63a3254a6ab5f22e6f9608b209
- Mxdperfstat
- http://www.arcgis.com/home/item.html?id=a269d03aa1c840638680e2902dadecac
-

Relationship between System Tools

System Monitor and Test inputs into System Designer capacity plan

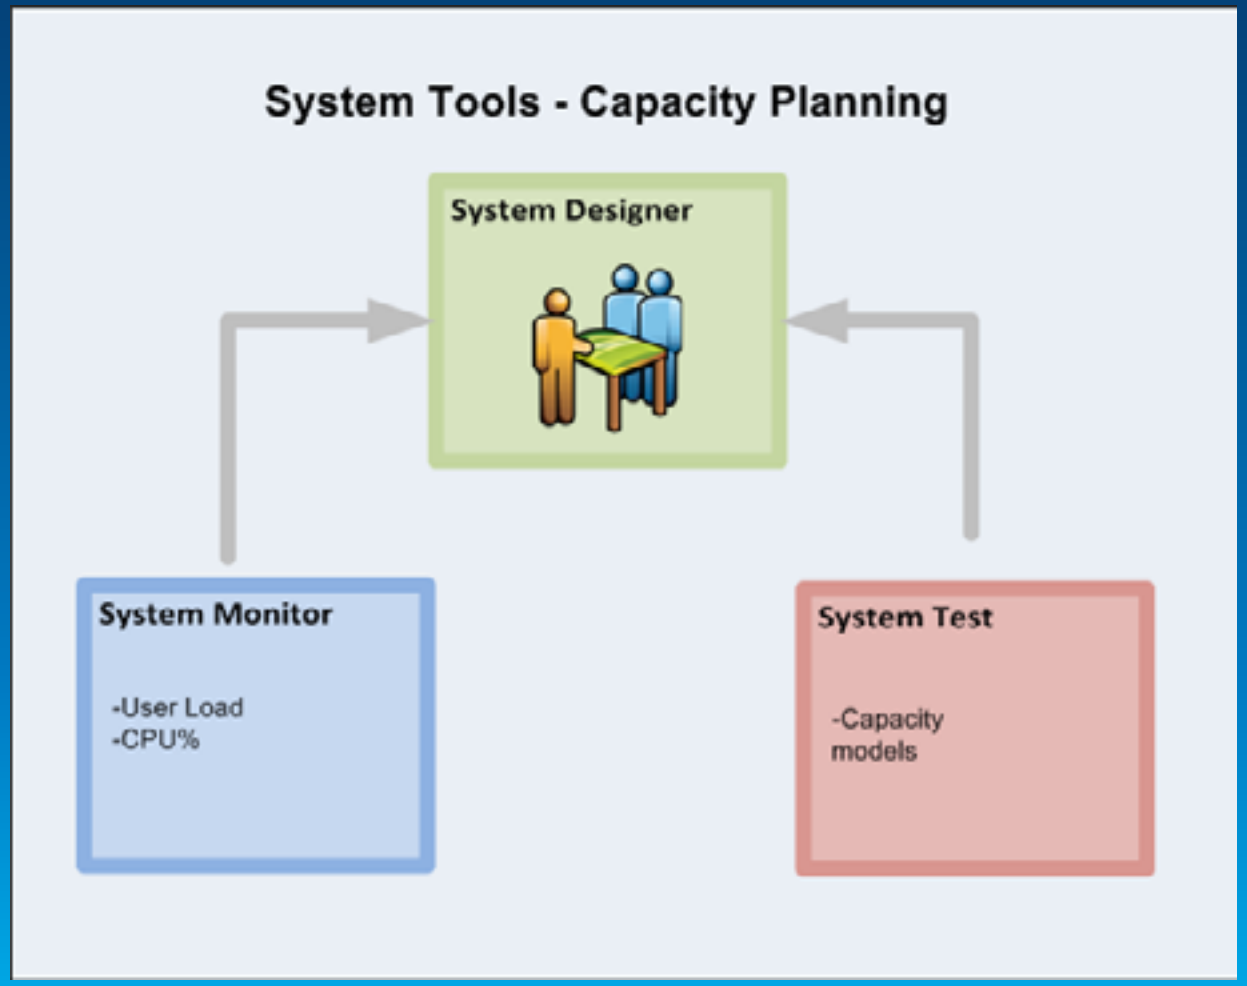

Performance Factors

GIS Services—Map Service

Optimize map document

- Set scale dependency
- Use simple labels and symbols
- Use appropriate attribute and spatial index
- Avoid reprojections on the fly
- Use fast joins (no cross database joins)
- Avoid wavelet compression-based raster types (MrSid, JPEG 2000)

GIS Services—ArcSOC instances

- Start with Max Instances =~ #CPU Cores
- Monitor usage
- Reduce as needed

If max SOC instances are under configured, system will not scale.

Geodatabase

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

Hardware

Most systems are CPU bound

GIS Systems are bound by:

1. **CPU - typical**
2. **Memory – when large number of services**
3. **Disk – Image Service, Synchronization**
4. **Network – low bandwidth deployment**

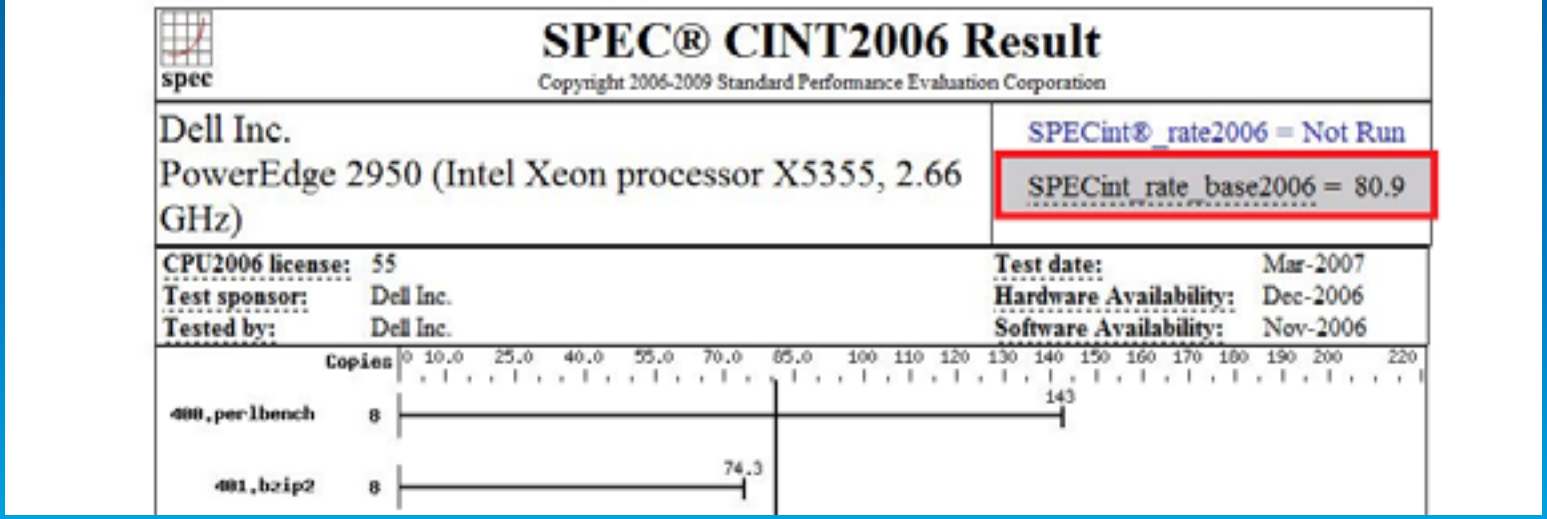

SPEC® CINT2006 Result

Copyright 2006-2009 Standard Performance Evaluation Corporation

Dell Inc.
PowerEdge 2950 (Intel Xeon processor X5355, 2.66 GHz)

SPECint®_rate2006 = Not Run
SPECint_rate_base2006 = 80.9

CPU2006 license: 55
Test sponsor: Dell Inc.
Tested by: Dell Inc.

Test date: Mar-2007
Hardware Availability: Dec-2006
Software Availability: Nov-2006

Copies 0 10.0 25.0 40.0 55.0 70.0 85.0 100 110 120 130 140 150 160 170 180 190 200 220

400.perlbench 8 143

401.bzip2 8 74.3

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

Hardware

Memory critical when publishing very large number of services

Item	Low	High
ArcSOC Map	50 MB	500 MB
ArcSOC Image	20 MB	1,024 MB
ArcSOC GP	100 MB	2,000 MB
XenApp Session	500 MB	1.2 GB
Database Session	10 MB	75 MB
Database Cache	200 MB	200 GB

Wide ranges of memory consumptions

Performance Factors - Hardware

Virtualization

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

Overburdened VMs will incur significant performance degradation

Tuning

Tuning methodology

Profile each tier starting from the top

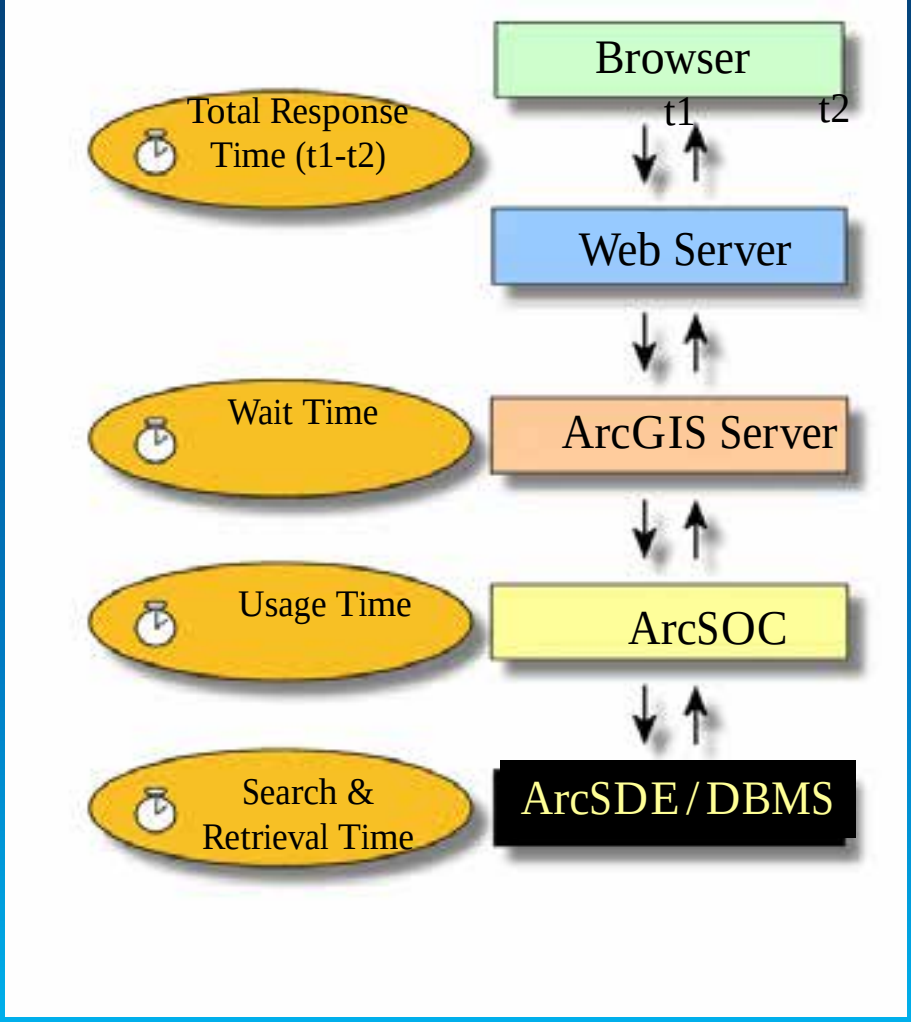

Fiddler

Fiddler measurement approximately 5.2 seconds

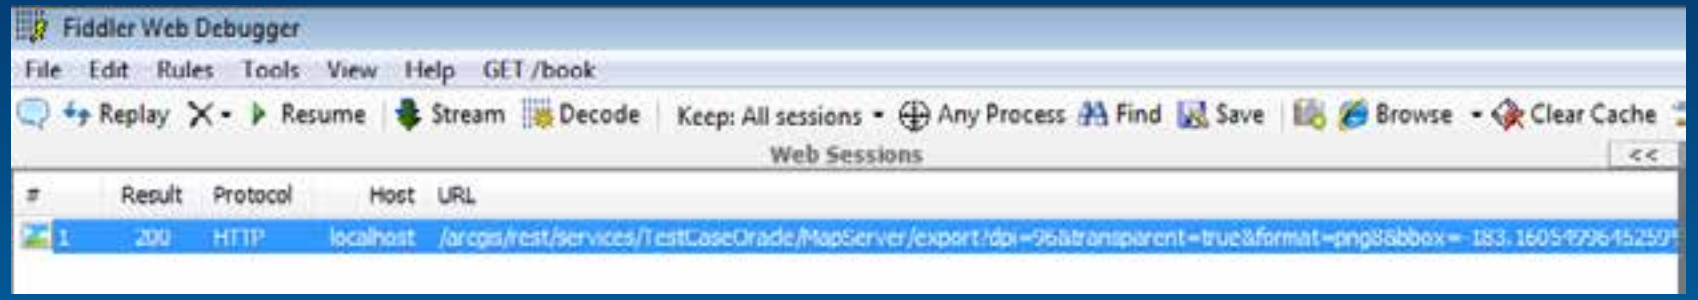

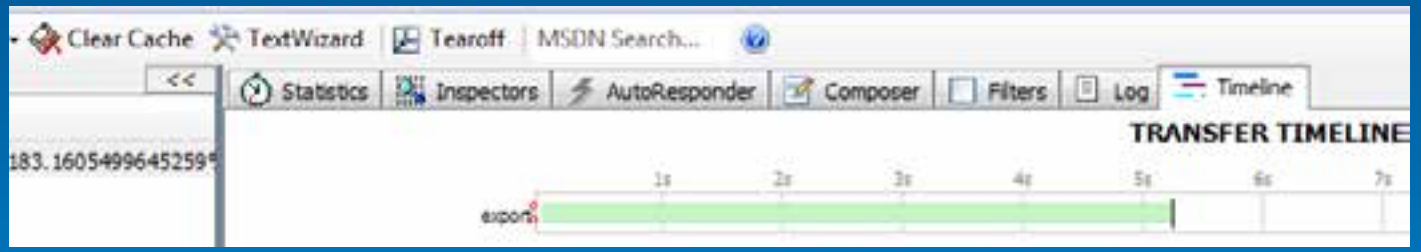

Mxdperfstat

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

Oracle Trace

Compare elapsed time

```
SQL ID: 6p20xry10fv4n Plan Hash: 569628948

SELECT  V__45.st_SHAPE$, V__45.OID,  V__45.st_points,V__45.st_numpts,
  V__45.st_entity,V__45.st_minx,V__45.st_miny,V__45.st_maxx,V__45.st_maxy,
  V__45.st_minz,V__45.st_maxz,V__45.st_minm,V__45.st_maxm,V__45.st_area$,
  V__45.st_len$,V__45.st_rowid
FROM
 (SELECT b.OID,b.GX,b.GY,b.ID,1 st_SHAPE$ ,b.SHAPE.points as st_points,
  b.SHAPE.numpts as st_numpts,b.SHAPE.entity as st_entity,b.SHAPE.minx as
  st_minx,b.SHAPE.miny as st_miny,b.SHAPE.maxx as st_maxx,b.SHAPE.maxy as
  st_maxy,b.SHAPE.minz as st_minz,b.SHAPE.maxz as st_maxz,b.SHAPE.minm as
  st_minm,b.SHAPE.maxm as st_maxm,b.SHAPE.area as st_area$,b.SHAPE.len as
  st_len$,b.rowid as st_rowid FROM SDE.GridPoint b WHERE
  SDE.ST_EnvIntersects(b.SHAPE,:1,:2,:3,:4) = 1  AND b.OID NOT IN (SELECT /*+
  HASH_AJ */ SDE_DELETES_ROW_ID FROM SDE.D45 WHERE DELETED_AT IN (SELECT
  l.lineage_id FROM SDE.state_lineages l WHERE l.lineage_name =
  :lineage_name1 AND l.lineage_id <= :state_id1) AND SDE_STATE_ID = 0) UNION
  ALL SELECT a.OID,a.GX,a.GY,a.ID,2 st_SHAPE$ ,a.SHAPE.points as st_points,
  a.SHAPE.numpts as st_numpts,a.SHAPE.entity as st_entity,a.SHAPE.minx as
  st_minx,a.SHAPE.miny as st_miny,a.SHAPE.maxx as st_maxx,a.SHAPE.maxy as
  st_maxy,a.SHAPE.minz as st_minz,a.SHAPE.maxz as st_maxz,a.SHAPE.minm as
  st_minm,a.SHAPE.maxm as st_maxm,a.SHAPE.area as st_area$,a.SHAPE.len as
  st_len$,a.rowid as st_rowid FROM SDE.A45 a,SDE.state_lineages SL WHERE
  SDE.ST_EnvIntersects(a.SHAPE,:5,:6,:7,:8) = 1  AND (a.OID, a.SDE_STATE_ID)
  NOT IN (SELECT /*+ HASH_AJ */ SDE_DELETES_ROW_ID, SDE_STATE_ID FROM SDE.D45
  WHERE DELETED_AT IN (SELECT  l.lineage_id FROM SDE.state_lineages l WHERE
  l.lineage_name = :lineage_name2 AND l.lineage_id <= :state_id2) AND
  SDE_STATE_ID > 0) AND a.SDE_STATE_ID = SL.lineage_id AND SL.lineage_name =
  :lineage_name3 AND SL.lineage_id <= :state_id3) V__45 WHERE  (ID<1000)


call     count       cpu    elapsed       disk      query    current        rows
------- ------  -------- ---------- ---------- ---------- ----------  ----------
Parse        0      0.00       0.00          0          0          0           0
Execute      1      0.03       0.02          0          0          0           0
Fetch       20      9.67       9.64          0     129581          0        1998
------- ------  -------- ---------- ---------- ---------- ----------  ----------
total       21      9.70       9.66          0     129581          0        1998
```

Elapsed time slightly changed due to different test runs

Oracle Execution plan

```
Misses in library cache during parse: 1
Misses in library cache during execute: 1
Optimizer mode: ALL_ROWS
Parsing user id: 84
Number of plan statistics captured: 1

Rows (1st) Rows (avg) Rows (max)  Row Source Operation
---------- ---------- ----------  ---------------------------------------------------
      1998       1998       1998  VIEW  (cr=131605 pr=0 pw=0 time=512477 us cost=8 size=45906 card=21)
      1998       1998       1998   UNION-ALL  (cr=131605 pr=0 pw=0 time=511602 us)
      1998       1998       1998    FILTER  (cr=131451 pr=0 pw=0 time=508349 us)
      1998       1998       1998     TABLE ACCESS BY INDEX ROWID GRIDPOINT (cr=131451 pr=0 pw=0 time=4
    129600     129600     129600      DOMAIN INDEX (Sel: Default - Undefined) A29_IX1 (cr=2017 pr=0 pw
         0          0          0     NESTED LOOPS  (cr=0 pr=0 pw=0 time=1456 us cost=0 size=44 card=1)
         0          0          0      INDEX RANGE SCAN D45_PK (cr=0 pr=0 pw=0 time=2101 us cost=0 size
         0          0          0      INDEX UNIQUE SCAN LINEAGES_PK (cr=0 pr=0 pw=0 time=0 us cost=0 s
         0          0          0    NESTED LOOPS ANTI (cr=154 pr=0 pw=0 time=2247 us cost=5 size=2367
         0          0          0     NESTED LOOPS  (cr=154 pr=0 pw=0 time=2243 us cost=5 size=2367 car
         0          0          0      TABLE ACCESS BY INDEX ROWID A45 (cr=154 pr=0 pw=0 time=2242 us c
         0          0          0       BITMAP CONVERSION TO ROWIDS (cr=154 pr=0 pw=0 time=2236 us)
         0          0          0        BITMAP AND  (cr=154 pr=0 pw=0 time=2232 us)
         0          0          0         BITMAP CONVERSION FROM ROWIDS (cr=147 pr=0 pw=0 time=455 us)
         0          0          0          SORT ORDER BY (cr=147 pr=0 pw=0 time=454 us)
         0          0          0           INDEX RANGE SCAN A45_STATEID_IX1 (cr=147 pr=0 pw=0 time=439
         0          0          0         BITMAP CONVERSION FROM ROWIDS (cr=7 pr=0 pw=0 time=1768 us)
         0          0          0          SORT ORDER BY (cr=7 pr=0 pw=0 time=1768 us)
         0          0          0           DOMAIN INDEX (Sel: Default - Undefined) A29_IX1_A (cr=7 pr=
         0          0          0      INDEX UNIQUE SCAN LINEAGES_PK (cr=0 pr=0 pw=0 time=0 us cost=0 s
         0          0          0     VIEW PUSHED PREDICATE  VW_NSO_1 (cr=0 pr=0 pw=0 time=0 us cost=0
         0          0          0      FILTER  (cr=0 pr=0 pw=0 time=0 us)
         0          0          0       NESTED LOOPS  (cr=0 pr=0 pw=0 time=0 us cost=0 size=44 card=1)
         0          0          0        INDEX RANGE SCAN D45_PK (cr=0 pr=0 pw=0 time=0 us cost=0 size=
         0          0          0        INDEX UNIQUE SCAN LINEAGES_PK (cr=0 pr=0 pw=0 time=0 us cost=0
```

Inefficient spatial index

Testing

Testing process

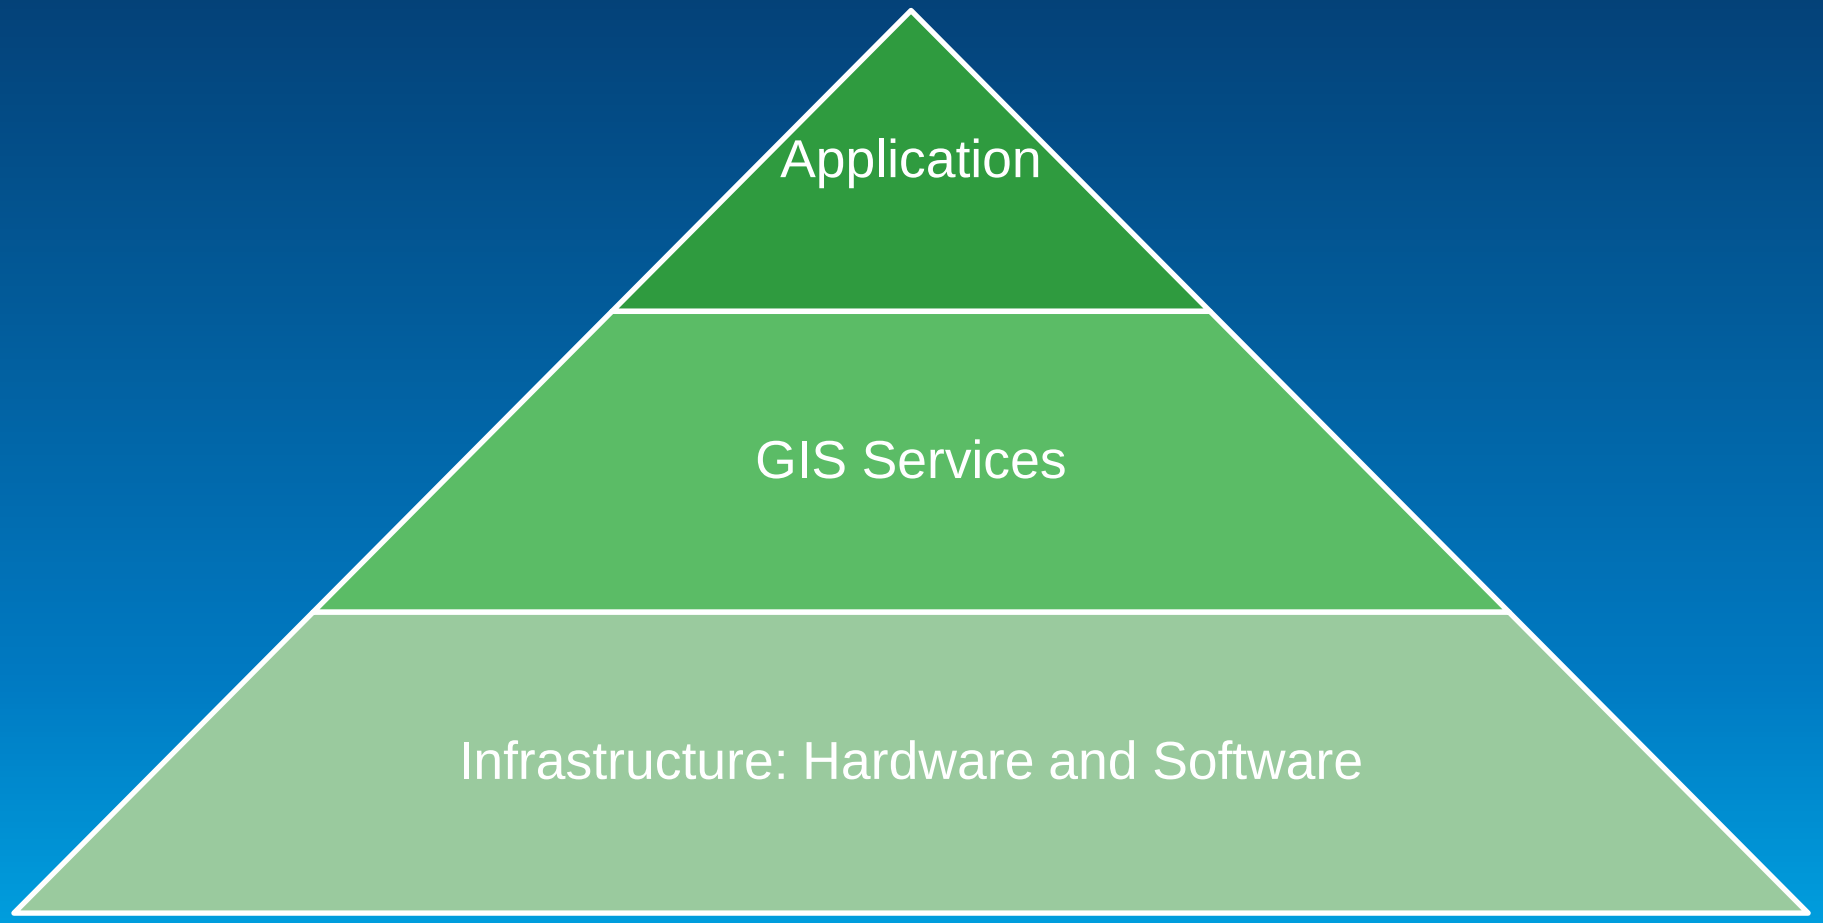

Required skill set

Configuration, Tuning, Testing

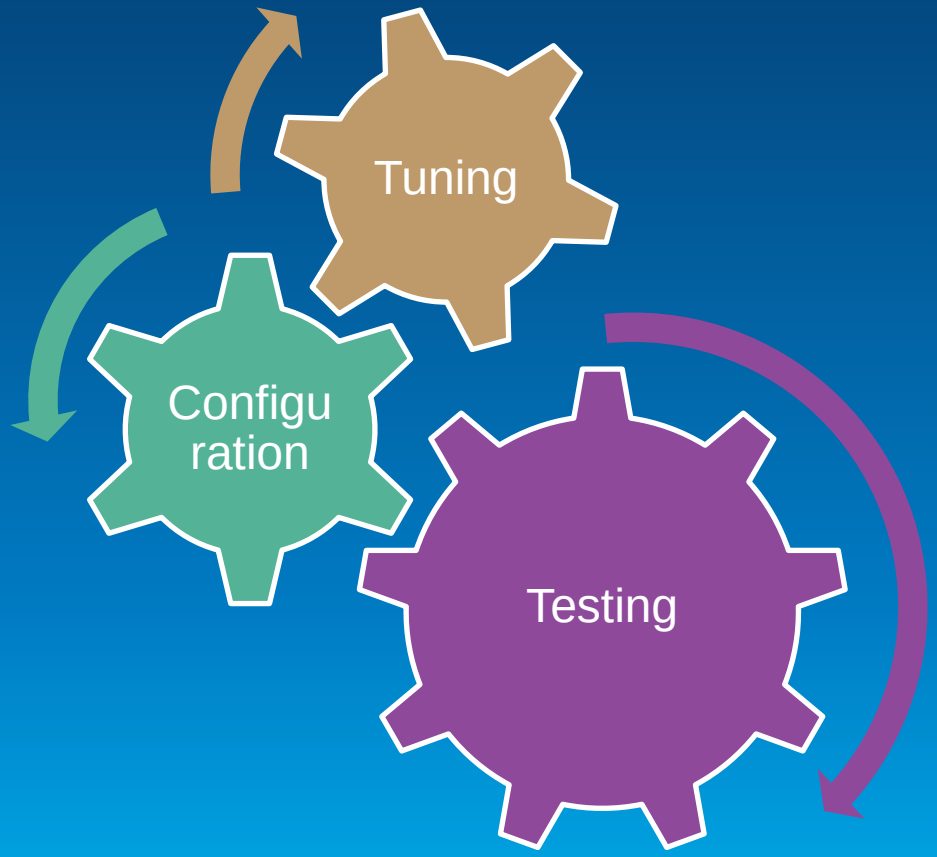

Demo

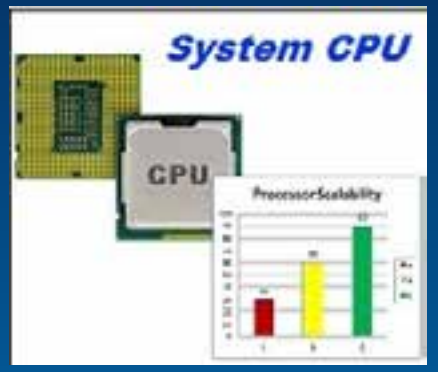

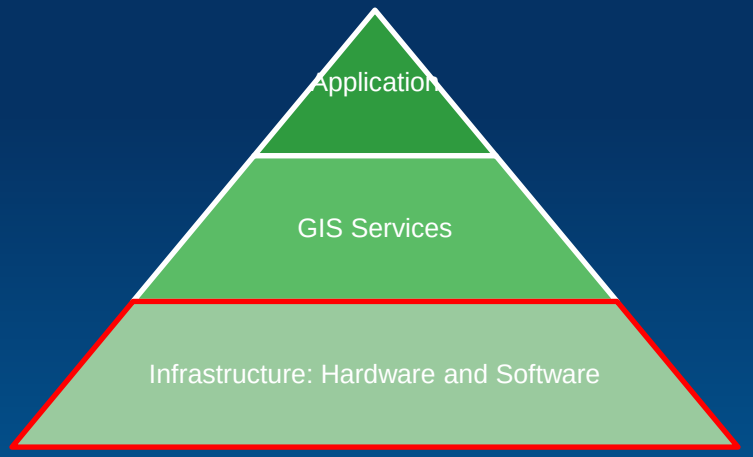

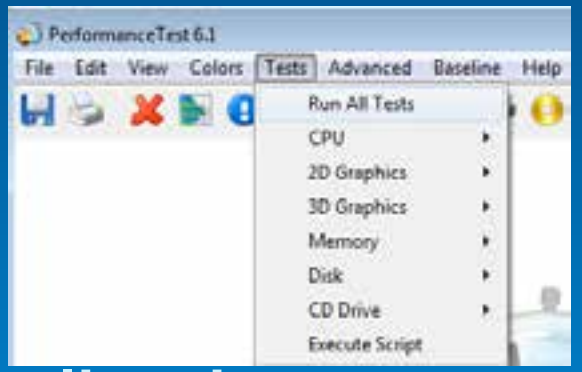

CPU Processor Speed

1. **System CPU**
2. **PassMark**

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

Network test

What is optimal location?

Could server location 1

Could server location 2

Demo

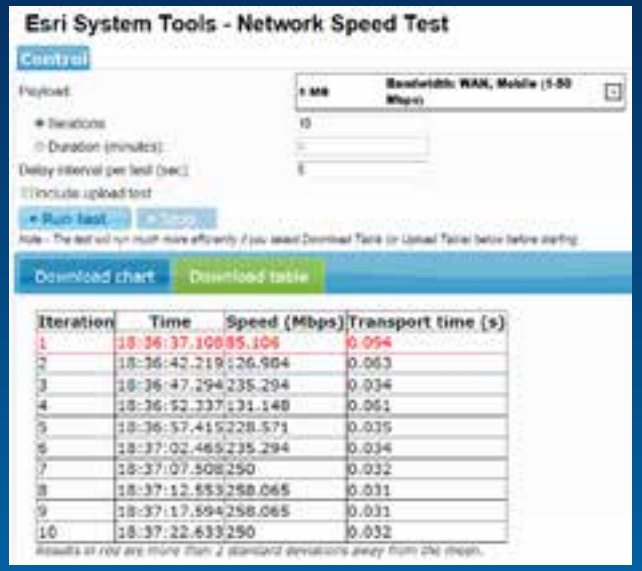

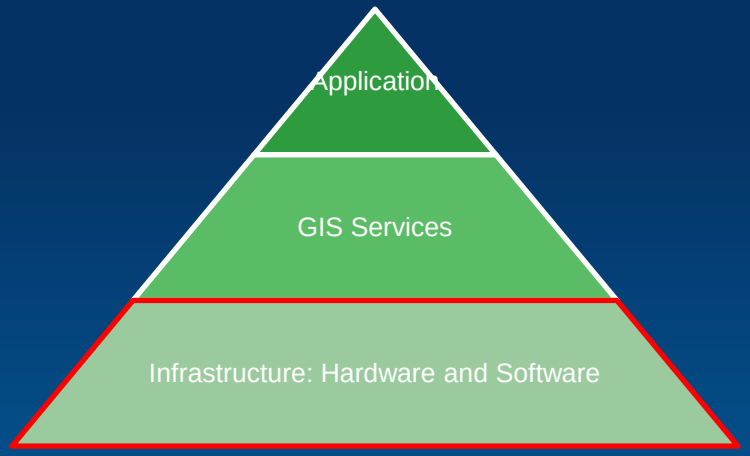

Network Speed Test Tool:
http://localhost/speedtest/

Test tools feature comparison

Tool	Cost	Learning Curve	OS Metrics	GIS Data Generation	GIS Test Automation
Load Runner	High	High	Windows/Linux	No	No
Visual Studio	Medium	High	Windows	No	No
JMeter	Free	High	Requires additional plugin	No	No
System Test	Free	Low	Windows/Linux	Yes	Yes

System Test Tool features

GIS Test Automation

- **ArcGIS Services**
  - Mapping
  - Feature Service
  - OGC
  - Geocoding
  - Image Service
  - Network Analyst
  - Geoprocessing
  - Tile Cache
- **Application Testing**
- **Discipline relevant report**

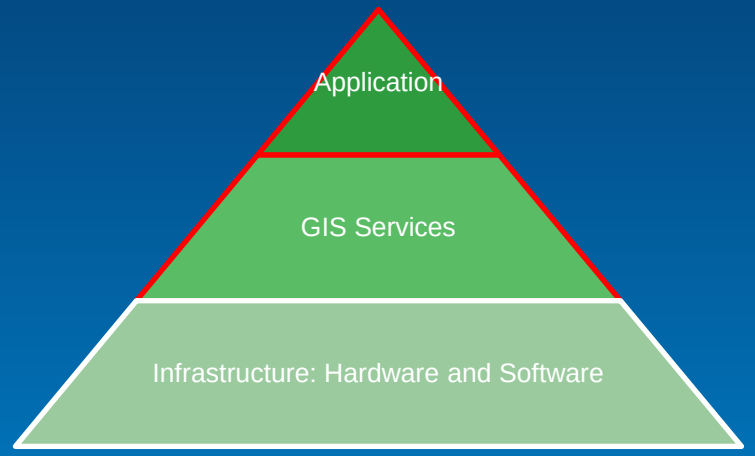

Testing Objectives

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

Designing test

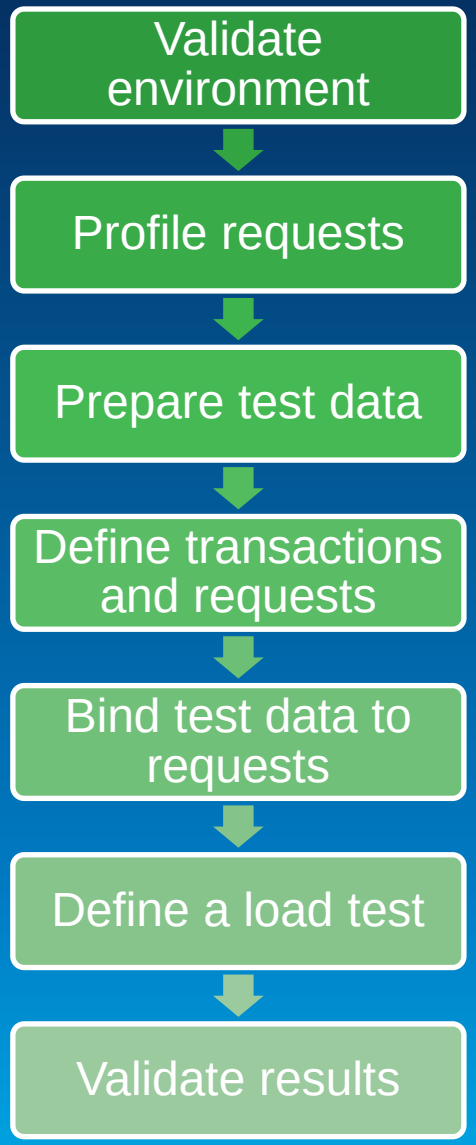

Analyze results

Valid

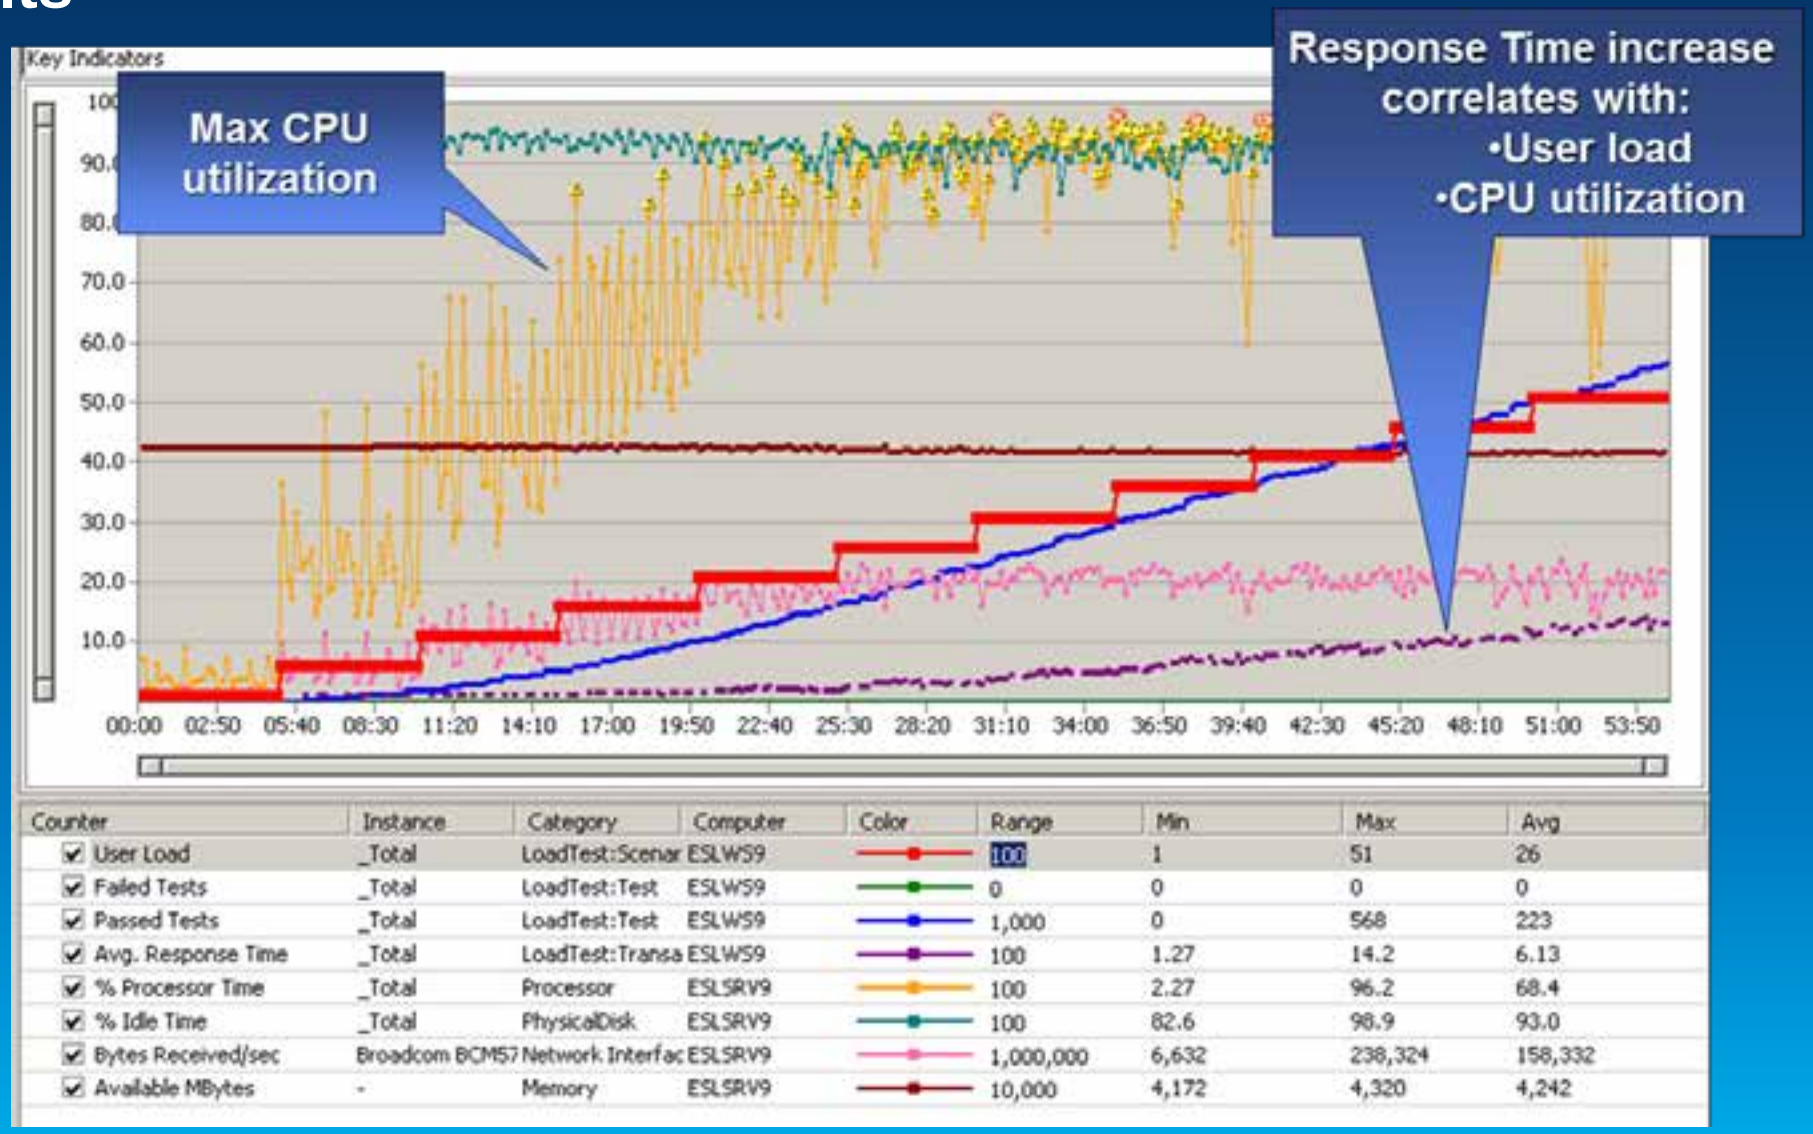

Counter	Instance	Category	Computer	Color	Range	Min	Max	Avg
☑ User Load	_Total	LoadTest:Scenar	ESLWS9		100	1	51	26
☑ Failed Tests	_Total	LoadTest:Test	ESLWS9		0	0	0	0
☑ Passed Tests	_Total	LoadTest:Test	ESLWS9		1,000	0	568	223
☑ Avg. Response Time	_Total	LoadTest:Transa	ESLWS9		100	1.27	14.2	6.13
☑ % Processor Time	_Total	Processor	ESLSRV9		100	2.27	96.2	68.4
☑ % Idle Time	_Total	PhysicalDisk	ESLSRV9		100	82.6	98.9	93.0
☑ Bytes Received/sec	Broadcom BCM57	Network Interfac	ESLSRV9		1,000,000	6,632	238,324	158,332
☑ Available MBytes	-	Memory	ESLSRV9		10,000	4,172	4,320	4,242

Analyze results

Invalid

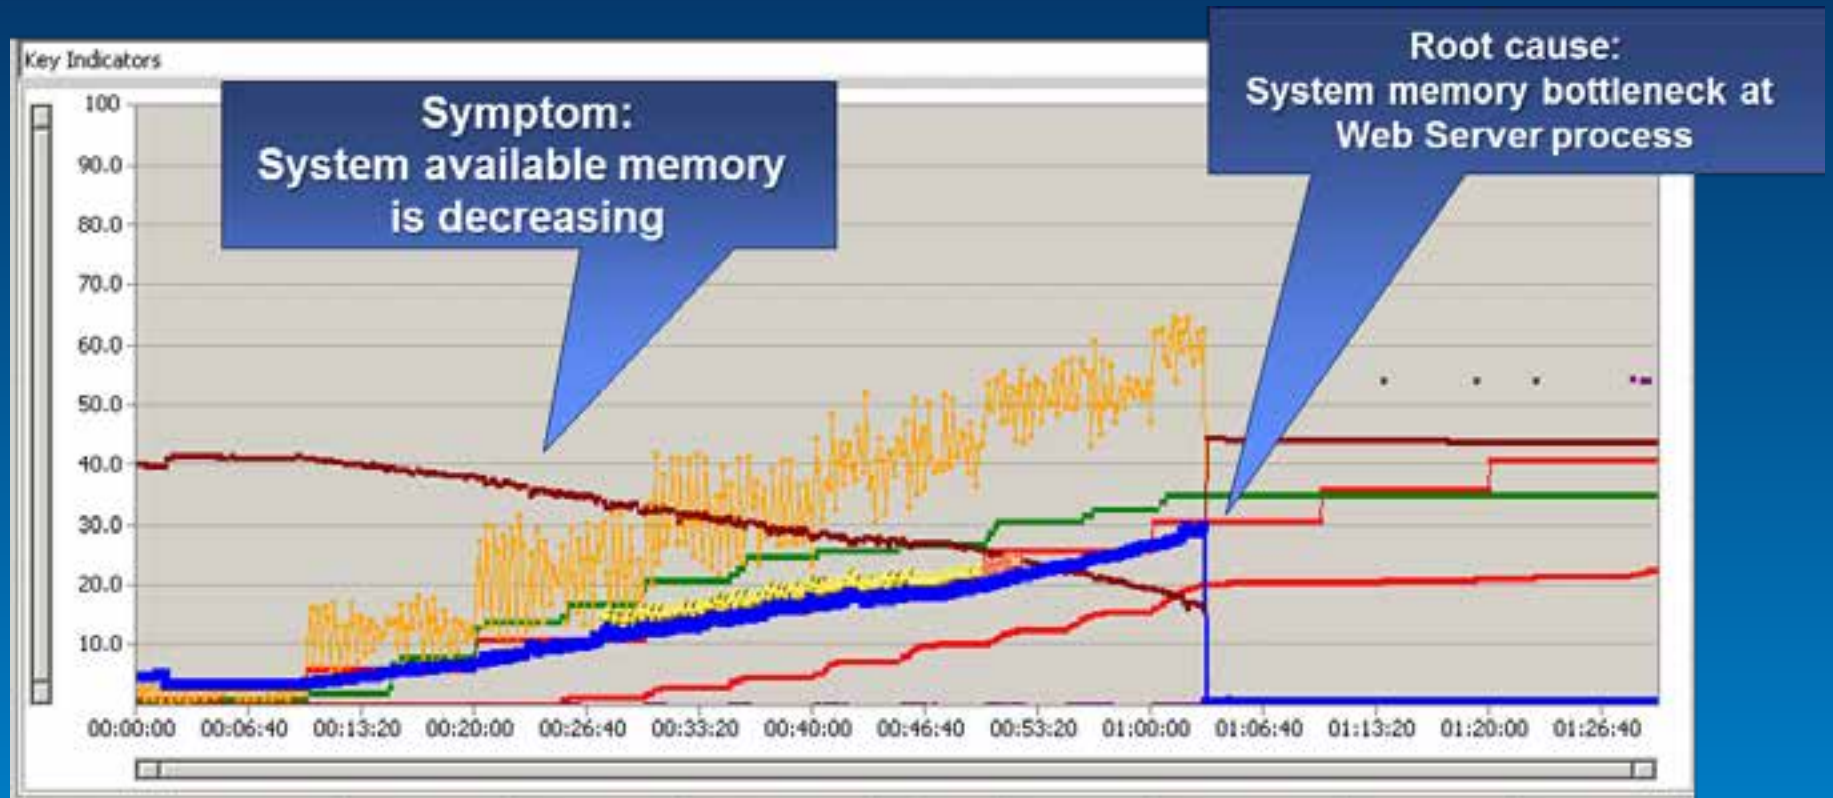

Counter	Instance	Category	Computer	Color	Range	Min	Max	Avg
☑ User Load	_Total	LoadTest:Scenar	ESLIBMBL4		100	1	41	21
☑ Failed Tests	_Total	LoadTest:Test	ESLIBMBL4		1,000	0	226	99
☑ Passed Tests	_Total	LoadTest:Test	ESLIBMBL4		100	0	35	23
☑ Avg. Response Time	_Total	LoadTest:Transa	ESLIBMBL4		1,000	0.31	546	1.84
☑ % Processor Time	_Total	Processor	ESLSRV15		100	0.94	100	49.6
☑ Available MBytes	-	Memory	ESLSRV15		10,000	1,629	4,485	3,575
☐ Avg. Content Length	ESRI-ArcGIS-ADf	LoadTest:Reque	ESLIBMBL4		1,000,000	5,442	105,690	98,554
☑ Private Bytes	w3wp	Process	ESLSRV15		10,000,000...	46,809,088	3,000,823,808	951,365,184

Demo

System Test - Map

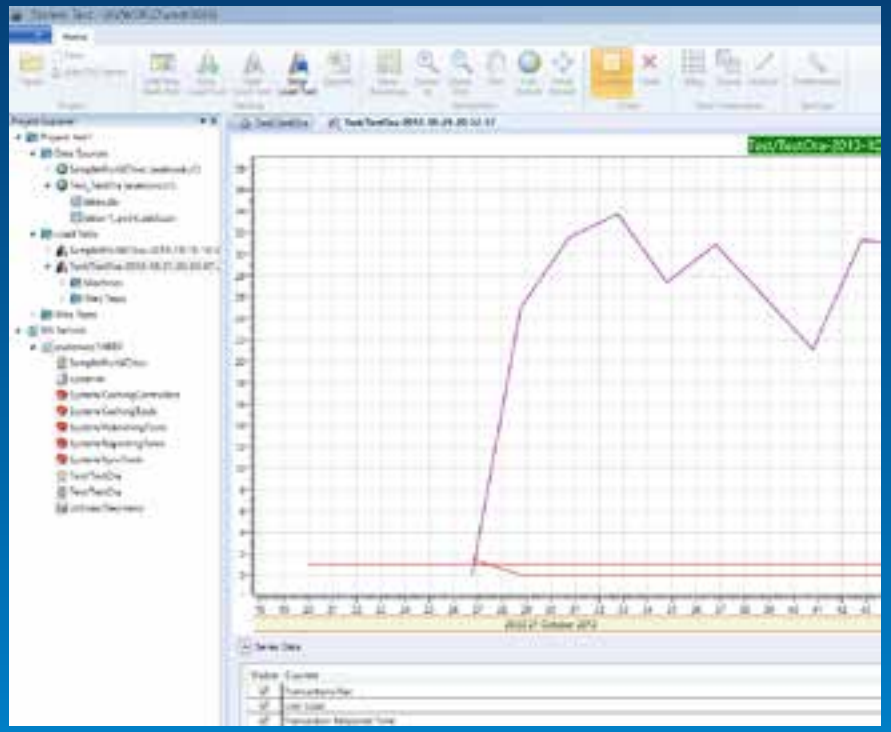

Impact of ArcGIS Server max instances

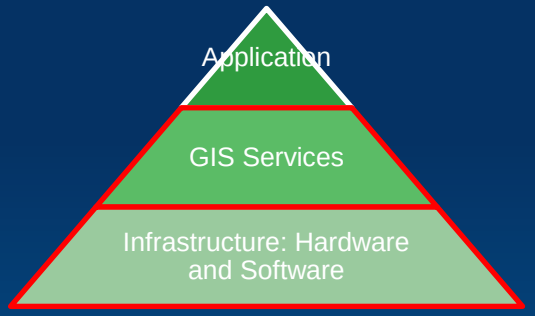

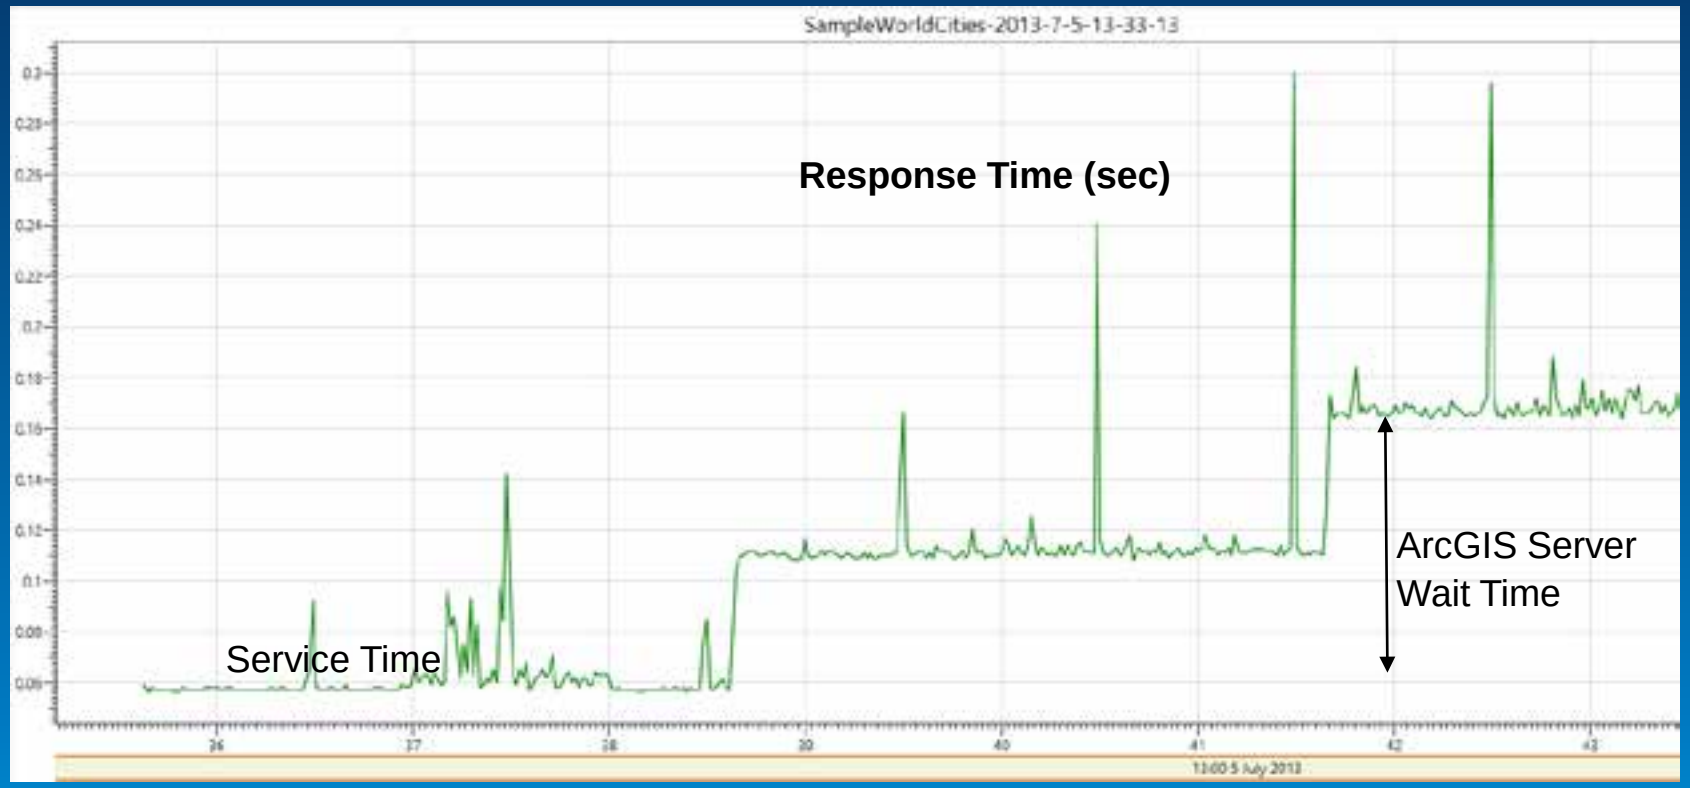

CPU ST/Tr @ ASAKOWICZ1		
Step Load	Avg. Value	Std. Deviation
1	0.054	0.018
2	0.048	0.009
3	0.045	0.009

Transaction Response Time @ ASAKOWICZ1		
Step Load	Avg. Value	Std. Deviation
1	0.060	0.023
2	0.111	0.023
3	0.168	0.026

Impact of CPU speed

CPU, Throughput and Response time ratios are similar

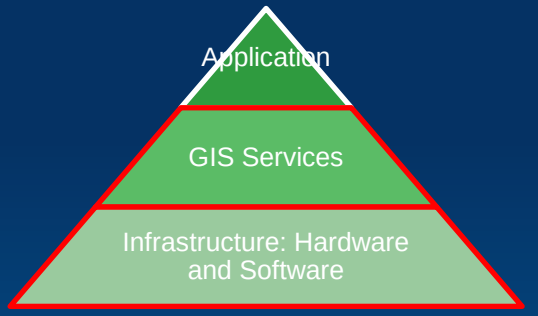

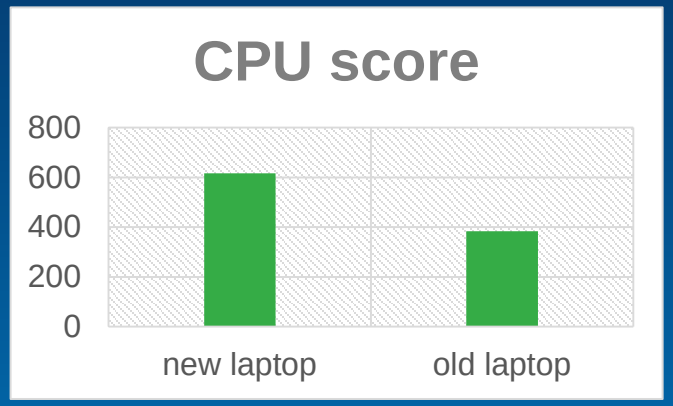

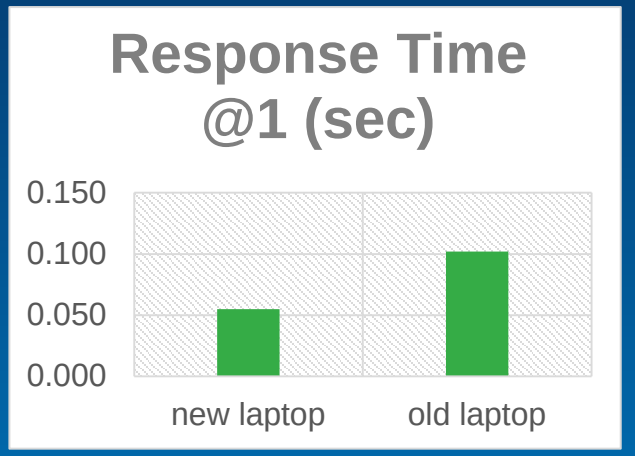

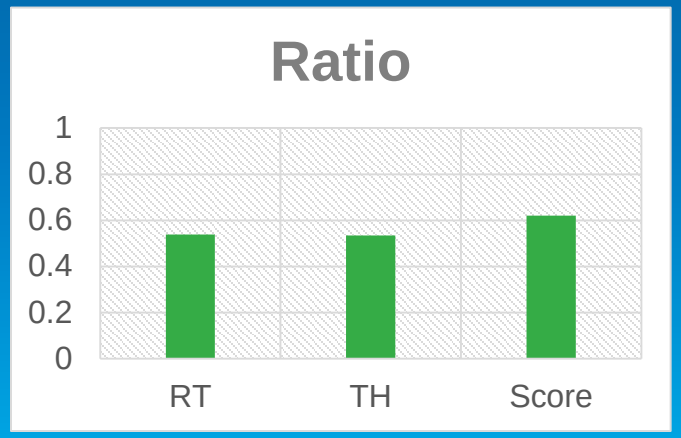

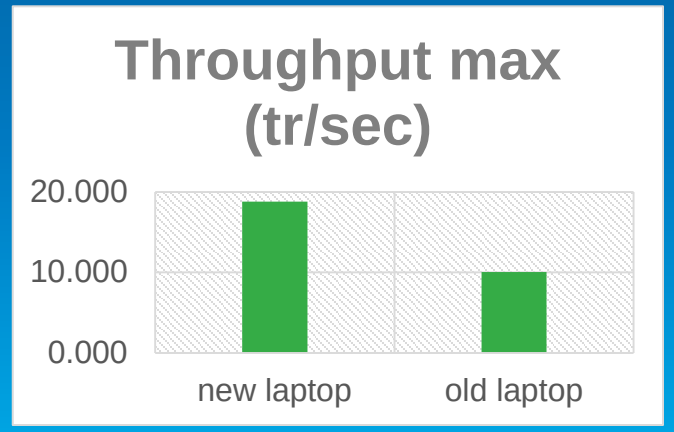

Demo

System Test - Network Analysis

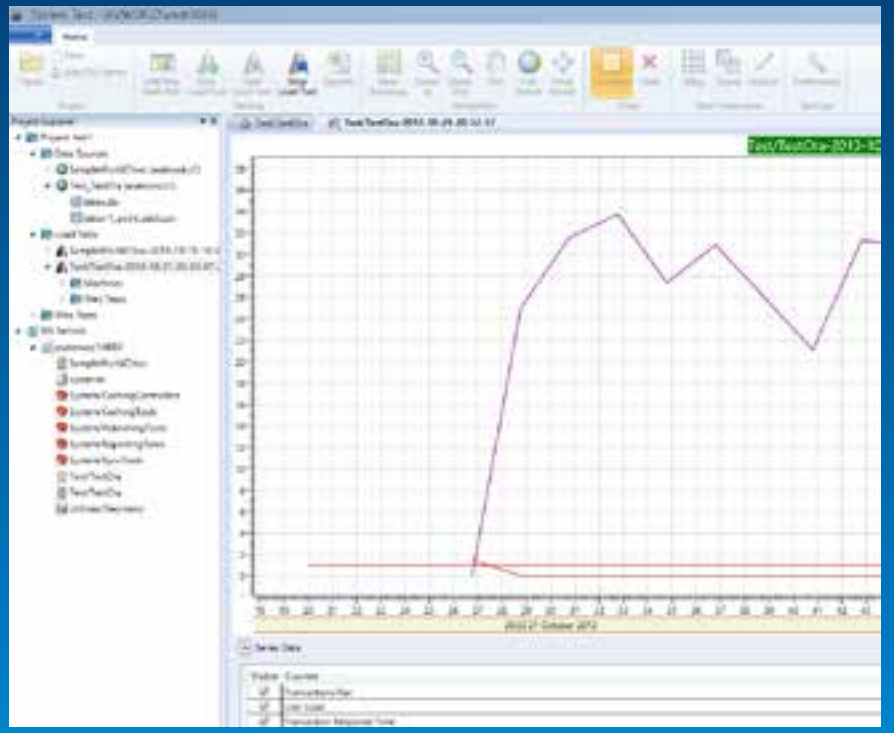

Demo

System Test - GP

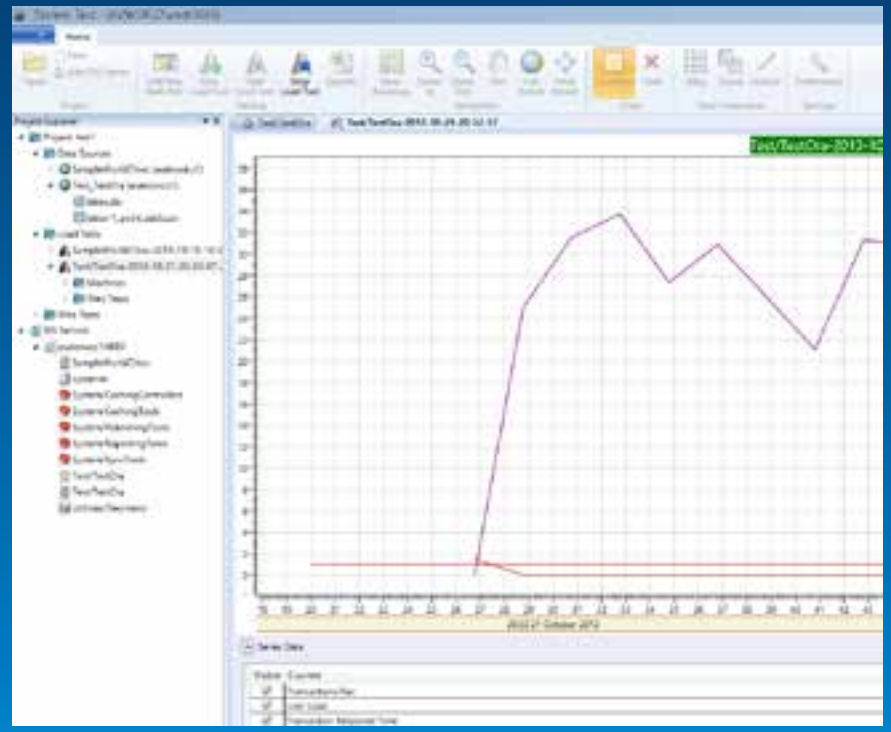

Monitoring

Monitoring Enterprise GIS

Challenges

- Multiple administrators
- Multiple disparate monitoring/diagnostic tools
- Data collected in a reactive fashion: on demand and for limited time
- Correlation of data with different timestamp is difficult
- ArcGIS administrators do not have access to all tools, data and reports
- Challenging to quickly identify the root cause and take appropriate measures

Integrated continuous end-to-end monitoring

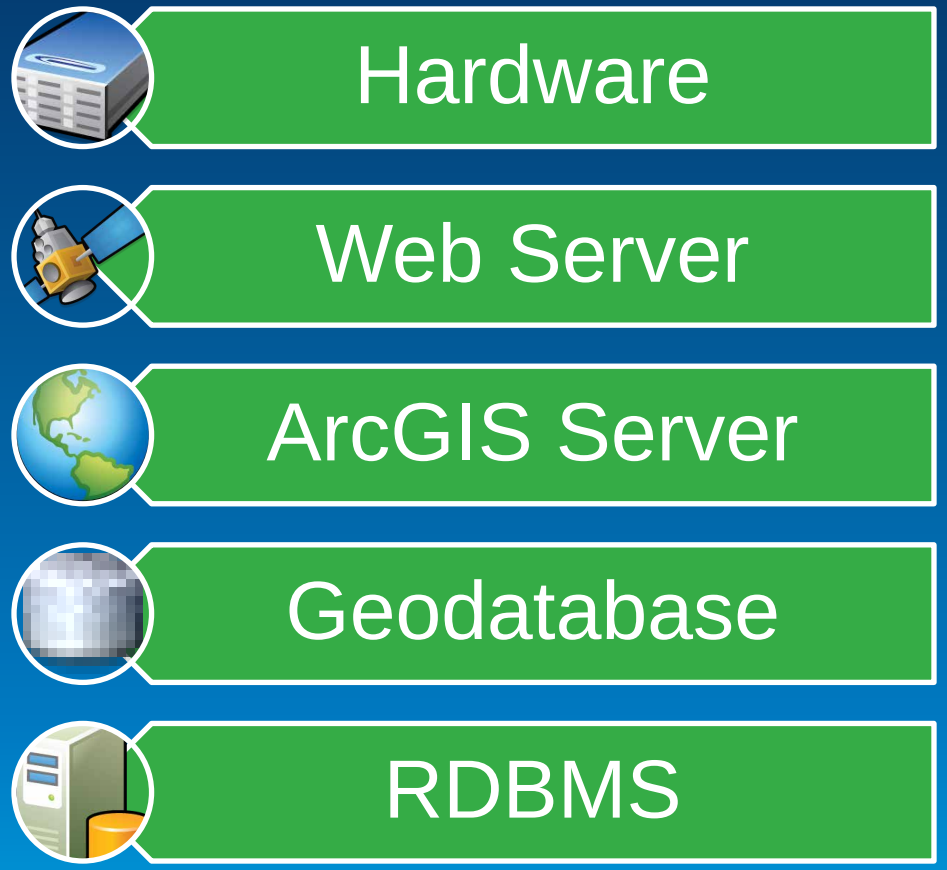

Key Performance Indicators (KPI)

Management KPI

Usage

Performance

Availability

Demo

System Monitor

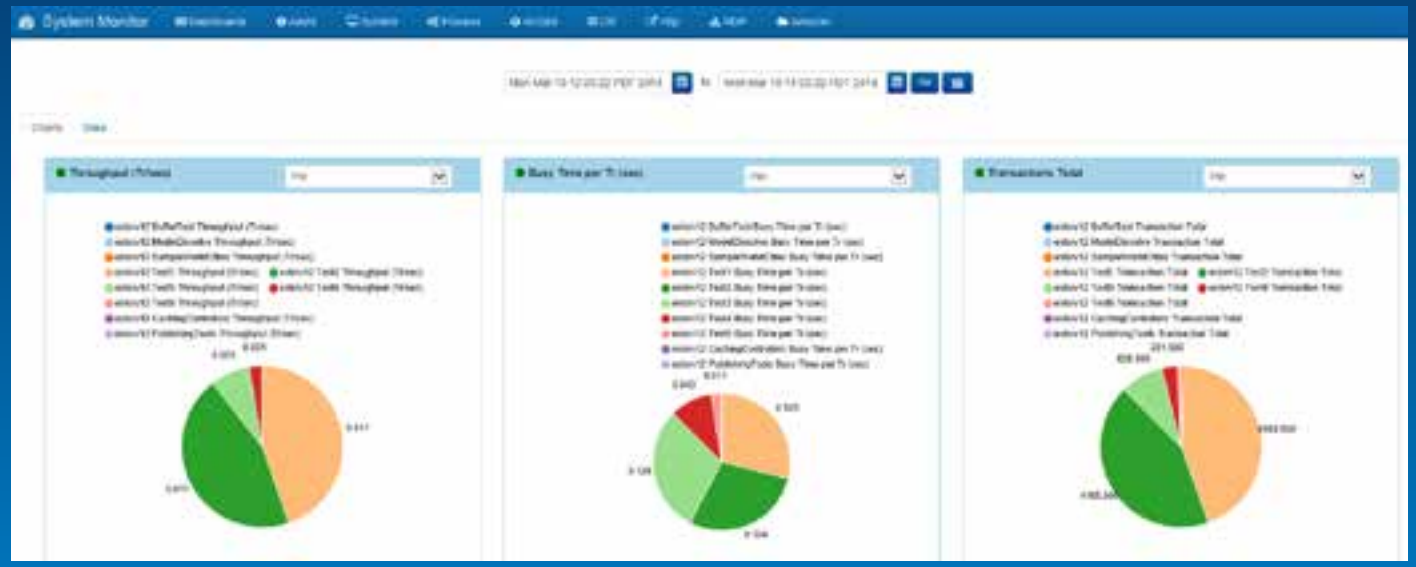

Summary

- Process
- Skills
- Tools

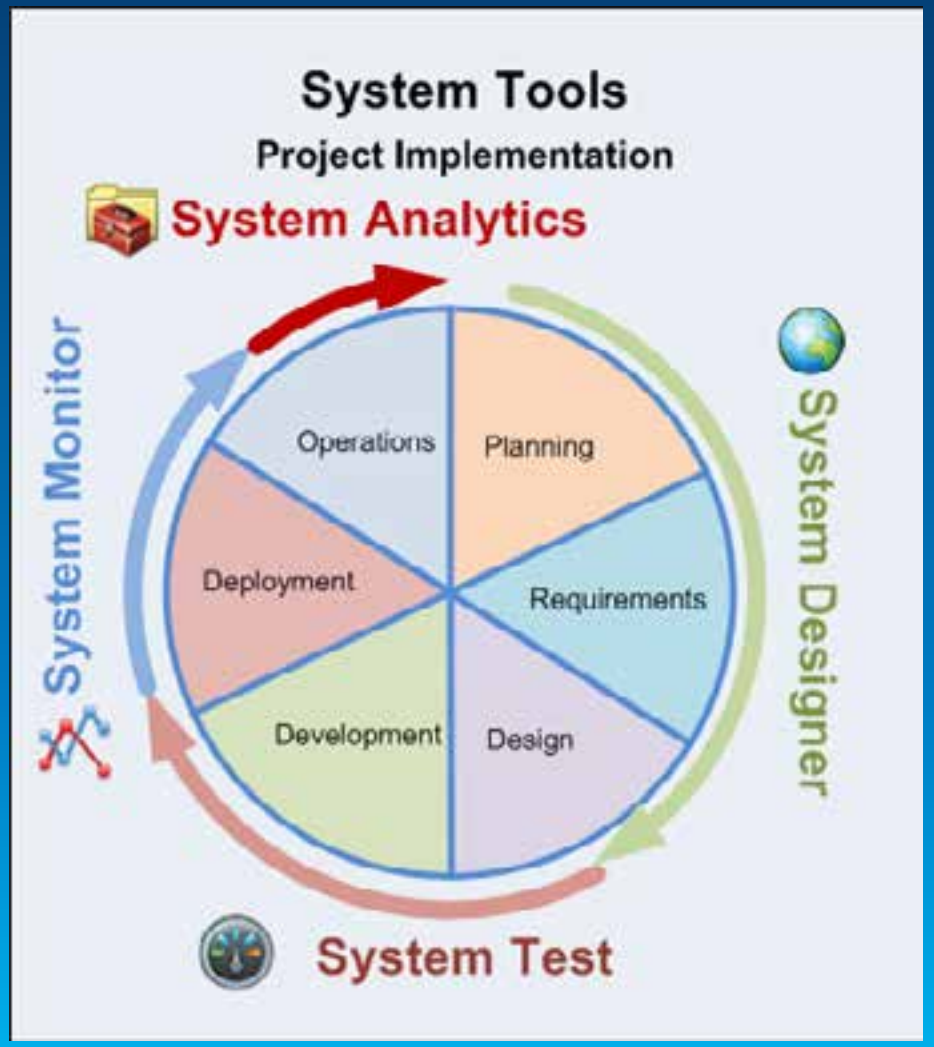

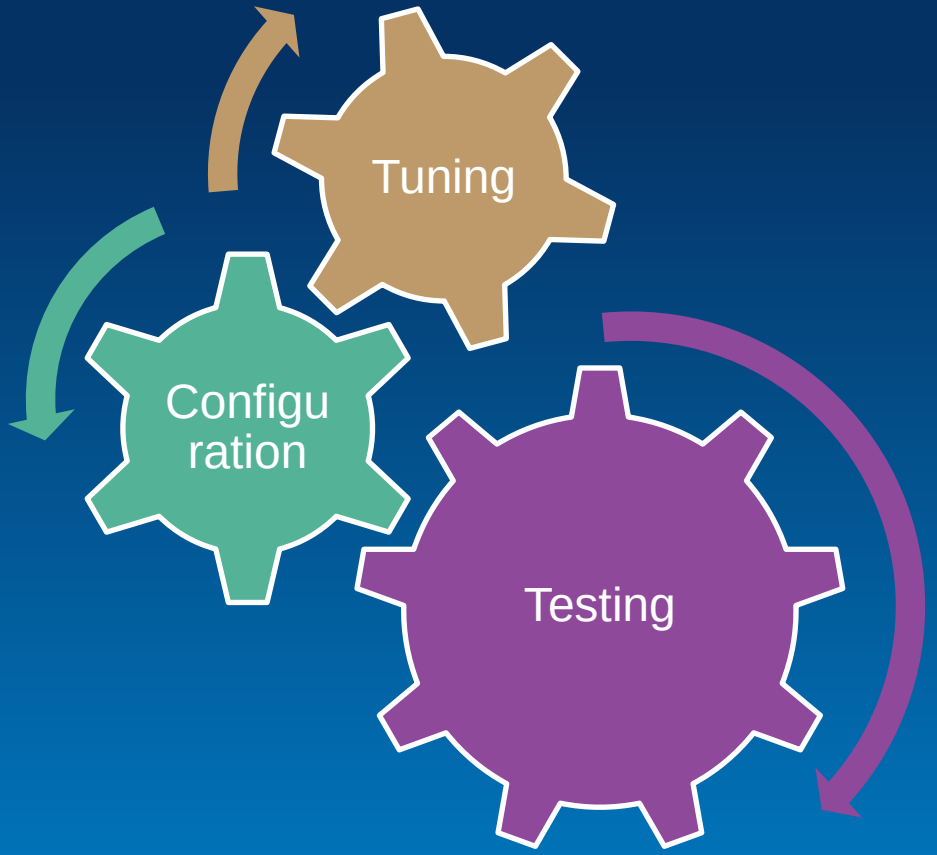

Survey

- http://www.esri.com/events/devsummit/session-rater

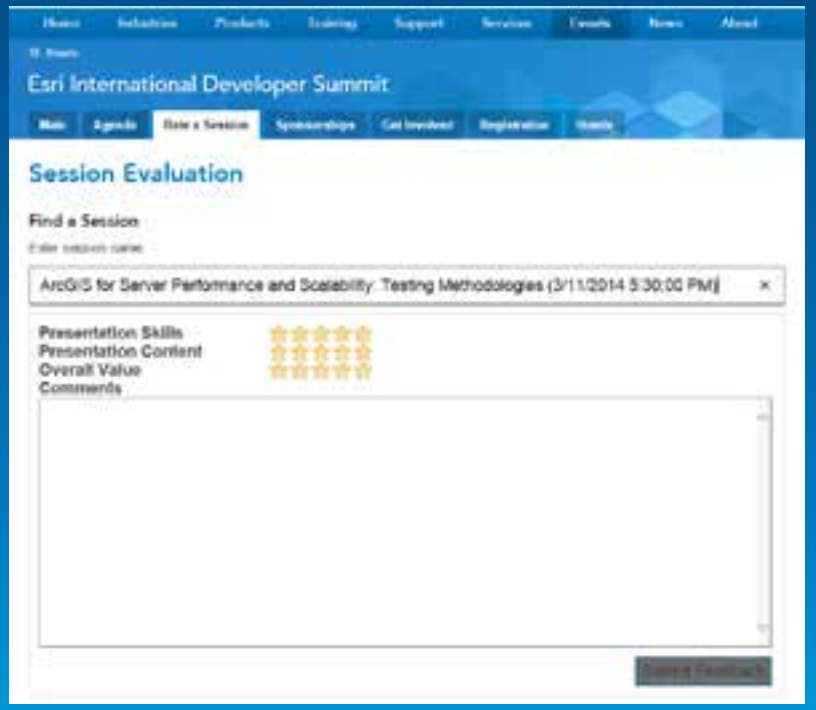

Questions

asakowicz@esri.com
fpizzi@esri.com

Understanding our world.