

# Performance Monitoring and Analysis of a Large Online Transaction Processing System

Manoj K. Nambiar, Hemanta Kumar Kalita  
Performance Engineering Research Centre, TCS Mumbai

1 July 2008

Experience certainty. IT Services  
Business Solutions  
Outsourcing

# Agenda

- **Background**
- **ScrutiNet:** A passive performance monitoring tool
- Performance Monitoring & Analysis of an OLTP System:  
**ScrutiNet based approach**
- **Conclusion**
- **References**





# Background

# Scenario

- A large company with 1,00,000 employees
  - Offices across the globe
  - Centralized IT Infrastructure
- Conducts Half yearly employee performance appraisals
  - Process enabled via a centralized web based application
- The entire appraisal process is open for 2 months

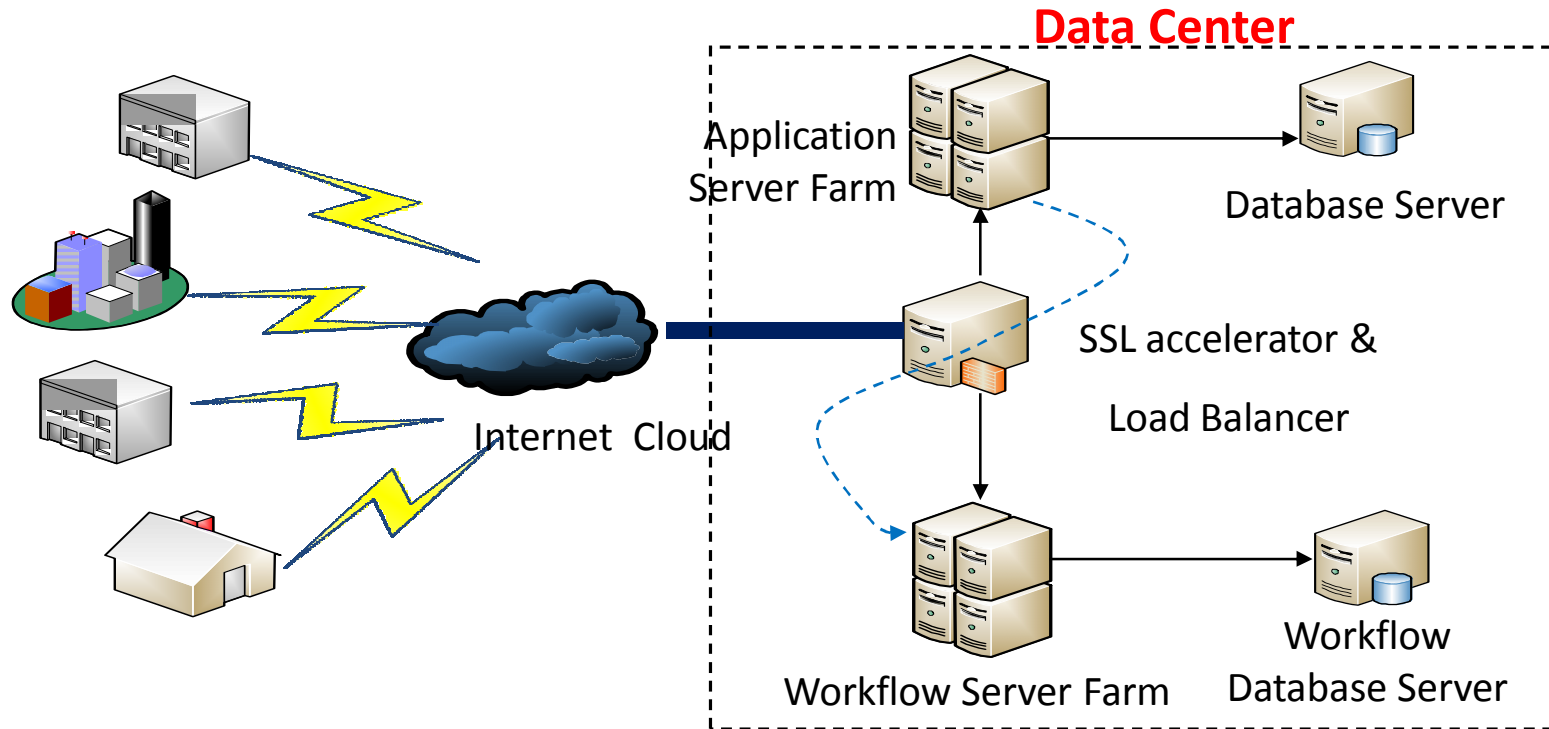


# Typical Process Work-Flow

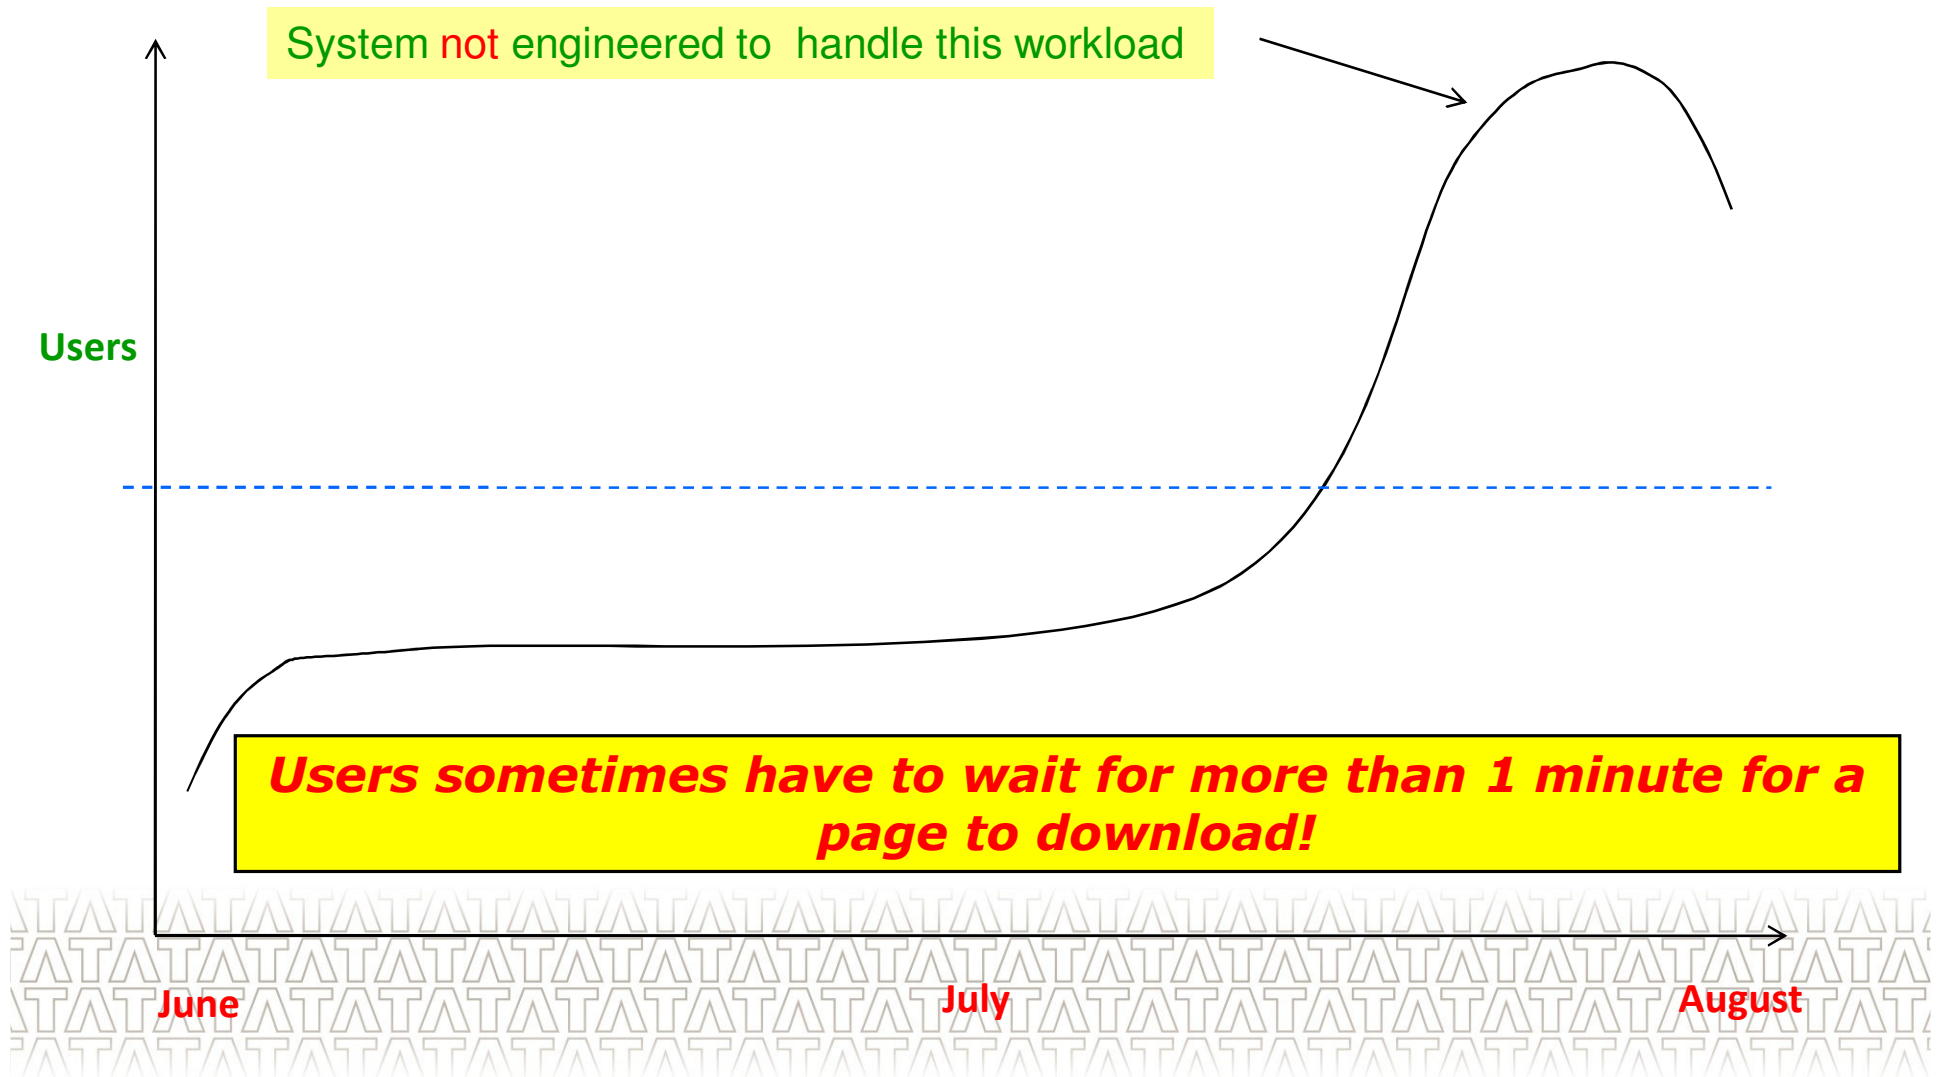
- **Employee** initiates appraisal
- **Immediate supervisor** assigns goals
- **Employee** appraises self
- **Immediate supervisor** rates the employee and puts in comments
- **Unit Head** reviews the ratings
- Ratings normalized by **HR** across the unit
- Evaluation shared with **employee**.

For each transaction listed above the user has to navigate through a series of web pages

# Application – Access and architecture



# System Workload

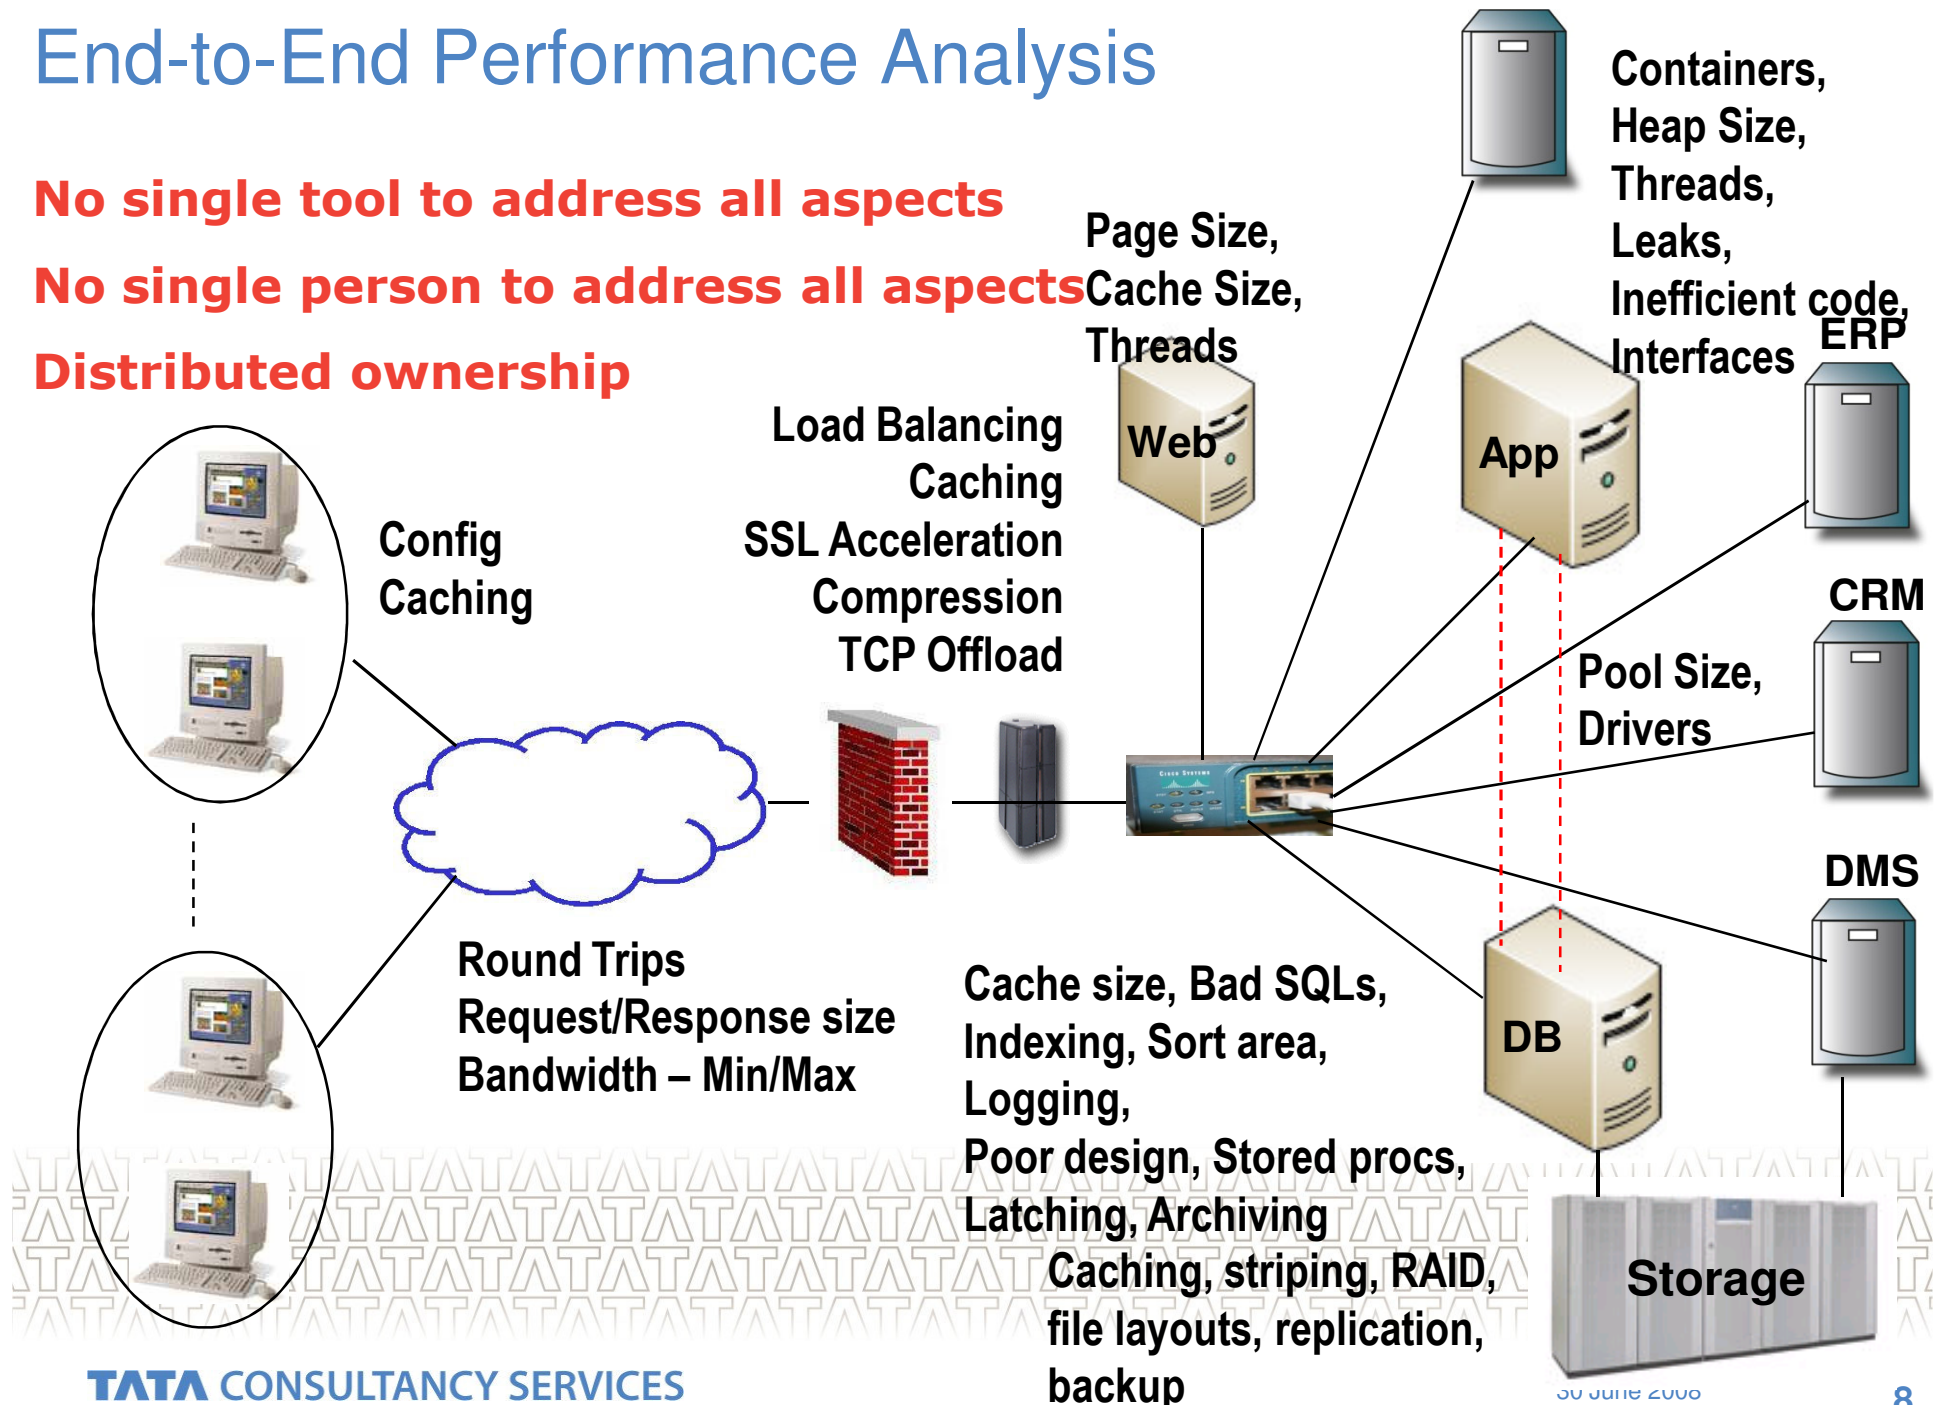


# End-to-End Performance Analysis

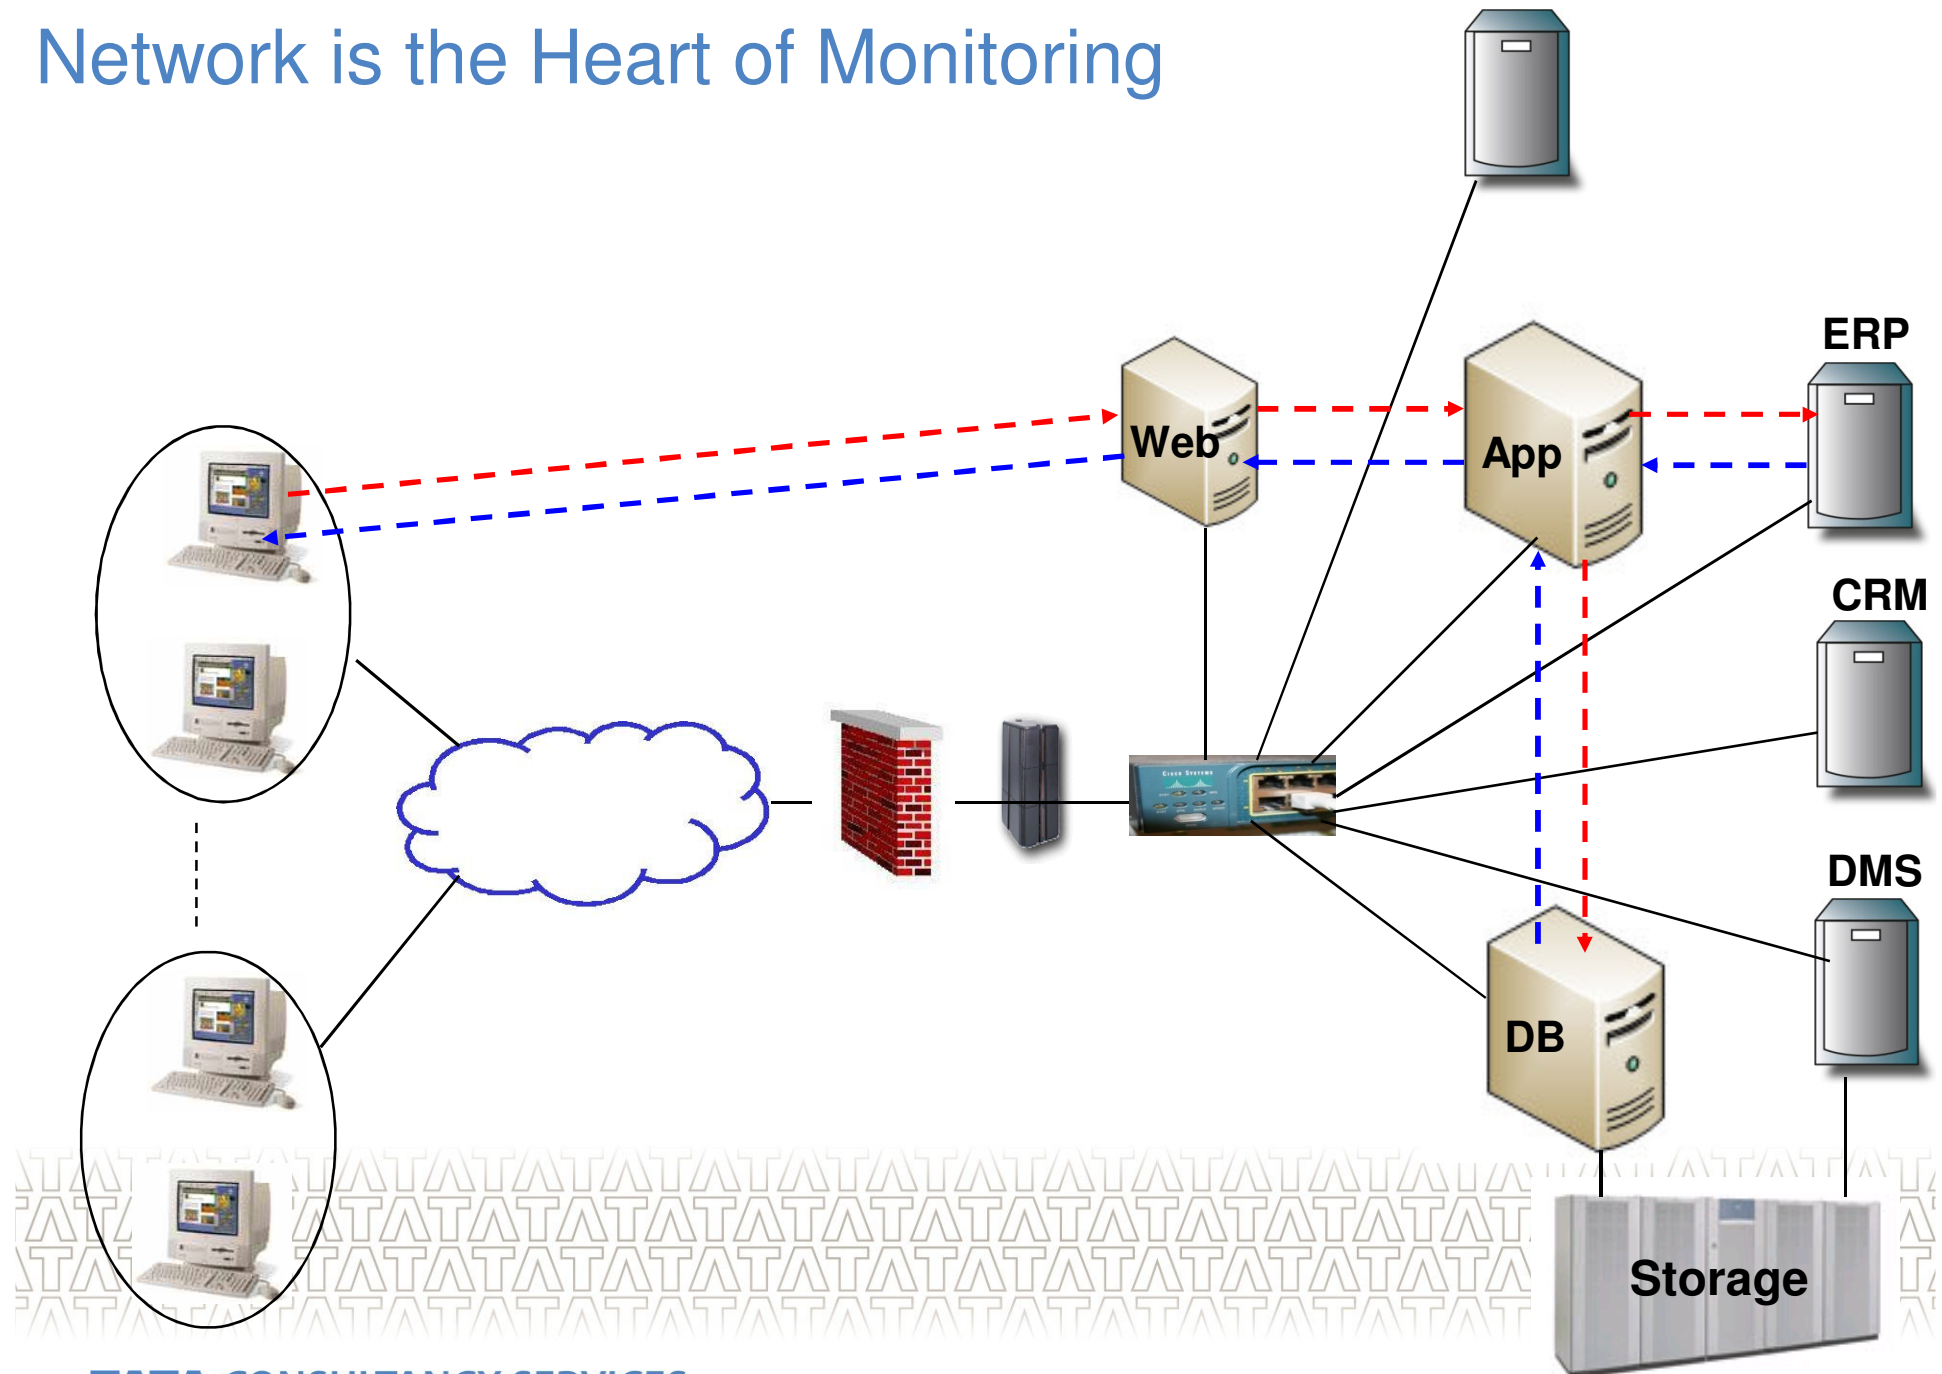
**No single tool to address all aspects**

**No single person to address all aspects**

**Distributed ownership**



# Network is the Heart of Monitoring





# ScrutiNet

A passive performance monitoring tool


# ScrutiNet™



Switch

Non intrusive

Cloned Packets




**Reporting Engine**  
Trace,  
Histogram,  
Averages

ScrutiNet Workstation


GET ./upload/Download.jsp.HTTP/1.0

Resp Time					
Avg	Min		Max		
373.0	323.0		420.0		
Proc Time			NW Time		
Avg	Min	Max	Avg	Min	Max
1.0	-	8.0	371.0	321.0	419.0
Req Send Time					
Avg	Min		Max		
-	-		-		
Client Delay					
Avg	Min		Max		
-	-		-		
Bytes					
Bytes In			Bytes Out		
Avg	Min	Max	Avg	Min	Max
340.0	340.0	340.0	7567.0	7567.0	7567.0



update Login set last\_login='Y' where Empid=X

Resp Time					
Avg	Min		Max		
9.0	-		54.0		
Proc Time			NW Time		
Avg	Min	Max	Avg	Min	Max
9.0	-	54.0	-	-	-
Req Send Time					
Avg	Min		Max		
-	-		-		
Client Delay					
Avg	Min		Max		
-	-		-		
Bytes					
Bytes In			Bytes Out		
Avg	Min	Max	Avg	Min	Max
85.0	85.0	85.0	20.0	20.0	20.0



## What is different in ScrutiNet ?

- Application agnostic.
- Can be used to detect the bottleneck in any request-response based OLTP transaction.
- It does not attempt to decode the application level protocol in the network packets.

## Questions that ScrutiNet answers

- What is the workload on my server?
- What is the Network link utilization of my application?
- Which are the top slow transactions?
- How many user connections are active on my server?
- How much data is transferred in one transaction?
- Where is the transaction spending most of its time?
- When did the transaction start?
- Who requested the transaction?



# ScrutiNet – Transaction Breakdown

## Breaks Delay contribution into

- Server
  - Includes I/O & processing
  - Call to next tiers
- Network (between client & Server)
- Client



# Other transaction specific details

- Transaction Start Time,
- Client IP Address and Port
- Average TCP Round trip time
- Total Bytes Transfer in Request (In) and Response(Out)
- Request content
  - Content of the first packet of the request message in ASCII (Or Hex)





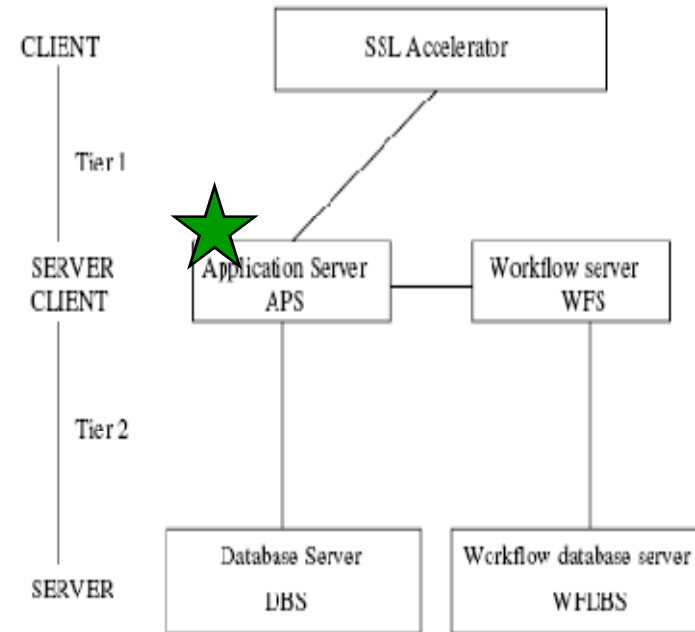


# Performance Monitoring & Analysis of an OLTP System: **ScrutiNet based approach**

1 July 2008

Experience certainty. IT Services  
Business Solutions  
Outsourcing

# ScrutiNet based approach

## Deploy ScrutiNet for Collection of Data.



# ScrutiNet based approach

## Transactions in the Application Server(front tier) sorted on response time

ScrutiNet report shows the top two slow transactions in the Application Server

Request No	Client	Client Port	Percentage client network request delay	Server Reaction Time (seconds)	Server Response Time (seconds)	Average Network Round trip time (seconds)	Network Overhead in the Server Response Time	Percentage Network Overhead in the Server Response Time	Client Window zero Interval	Total Bytes In	Total Bytes Out	Request contents
6785	██████████	3673	0	0.229001	56.12432	0.065364	2.734473	4.872171	49.52164	7864	163373	POST./pages/goalSettingHome.jsf.HTTP/1.1
2817	██████████	3596	95.94055	0.318773	33.51618	0.06017	3.521455	10.506731	26.05994	24049	161844	POST./pages/individualGoalSheetHome.jsf.H
10492	██████████	3701	95.94367	0.271653	29.01595	0.036071	1.179319	4.064381	58.68056	23002	139704	POST./pages/individualGoalSheetHome.jsf.H
13170	██████████	3736	93.31853	0.56624	27.05198	0.029673	0.900574	3.32905	22.44872	23849	174452	POST./pages/individualGoalSheetHome.jsf.H
9907	██████████	3688	98.29304	0.55814	20.58293	0.053934	2.169634	10.540938	14.01995	22547	159694	POST./pages/individualGoalSheetHome.jsf.H
1760	██████████	3810	1.453199	20.27197	20.27475	0.001029	0.002783	0.013726	0	2517	24502	POST./pages/changeAllocationHome.jsf.HT



# ScrutiNet based approach

Request No	Transaction start time	Client	Client Port	Request Send Time (seconds)	Server Reaction Time (seconds)	Server Response Time (seconds)	Average Network Round trip time (seconds)	Network Overhead in the Server Response Time	Client Window zero Interval	Total Bytes In	Total Bytes Out
6785	12/07/2007	[REDACTED]	3673	0	0.229001	56.12432	0.065364	2.734473	49.52164	7864	163373
2817	12/07/2007	[REDACTED]	3596	1.767359	0.318773	33.51618	0.06017	3.521455	26.05994	24049	161844

Request No 6785

POST./pages/goalSettingHome.jsf.HTTP/1.1  
 POST./pages/individualGoalSheetHome.jsf.HTTP/1.1

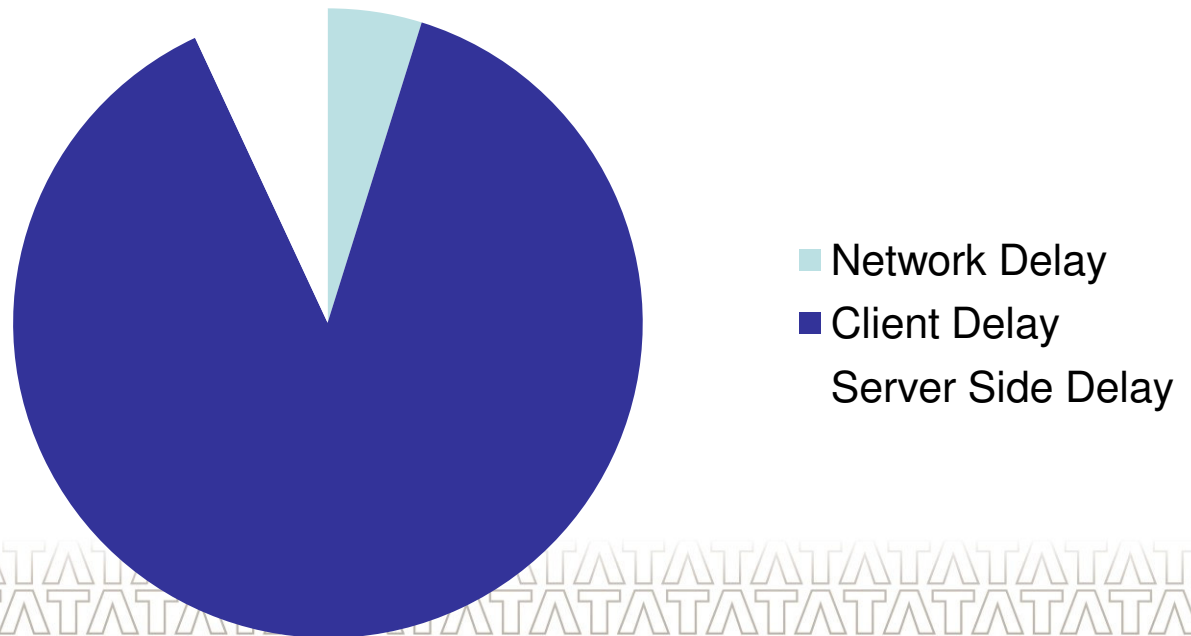
Request No 2817

Request no 6785 passes from the SSL Accelerator (port 3673) to the Application Server (APS) and the Application Server takes **56.12** seconds to respond.



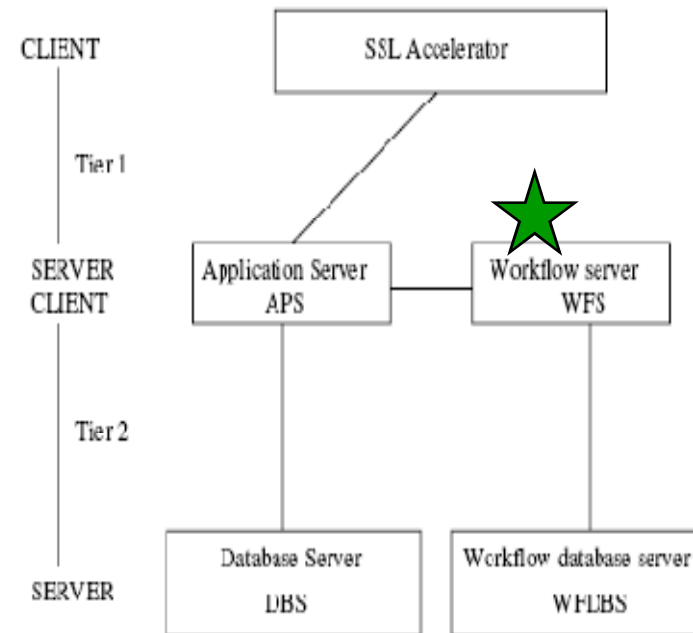


# Application Server Transaction 6785 - Example

- Client – Load Balancer
- Server – Application Server
- Overall Response Time = 56.12 seconds
- Network Delay = 2.7 seconds
- Client Delay = 49.52 seconds
- Request size = 7KB
- Response Size = 160 KB



# Another Example

## Monitoring of Workflow server



# ScrutiNet based approach

The slowest transaction in the Workflow Server is an XML based query called QUERY PROPERTY (Request no 582) which is sent by the Application Server.

Request No	Transaction start time	Client	Client Port	Request Send Time (seconds)	Server Reaction Time (seconds)	Server Response Time (seconds)	Round Trip Time (seconds)	Average Network Round trip time (seconds)	Network Overhead in the Server Response Time	Client Window zero Interval	Total Bytes In	Total Bytes Out
582	12/07/2007	[REDACTED]	53307	0	19.37399	19.37402	0.062015	0.03107	0.000034	0	1300	3765

Request No	Request contents
582	<pre>&lt;env:Envelope xmlns:env="http://schemas.xmlsoap.org/soap/envelope/" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xmlns:soapenc="http://schemas.xmlsoap.org/soap/encoding/" xmlns:xsd="http://www.w3.org/2001/XMLSchema"&gt;&lt;env:Header&gt;&lt;/env:Header&gt;&lt;env:Body&gt;&lt;n1:getTasks xmlns:n1="http://InitiateWorkflow/worklistComponent1INF"&gt;&lt;whereClause&gt;TASK.ORIGINATOR.= '133111'.and.TASK.KIND.IN(105).and.TASK.STATE.IN(2comma8comma3).and.QUERY_PROPERTY1.NAME.like.'%originator%'.and.upper(QUERY_PROPERTY1.STRING_VALUE).like.upper('%133751%')..and.QUERY_PROPERTY2.NAME.like.'%ProcessType%'.and.QUERY_PROPERTY3.NAME.like.'%requestType%'.and.QUERY_PROPERTY4.NAME.like.'%desc%'.and.QUERY_PROPERTY5.NAME.like.'%dueDate%'.and.QUERY_PROPERTY6.NAME.like.'%applicationName%'.and.upper(QUERY_PROPERTY6.STRING_VALUE).like.upper('%Speed.Workflow%')..and.QUERY_PROPERTY7.NAME.like.'%NotificationType%'.and.QUERY_PROPERTY8.NAME.like.'%ProcessName%'.and.QUERY_PROPERTY9.NAME.like.'%workflowName%'&lt;/whereClause&gt;&lt;orderBy&gt;&lt;/orderBy&gt;&lt;skipTuples&gt;0&lt;/skipTuples&gt;&lt;threshold&gt;0&lt;/threshold&gt;&lt;num_query_props&gt;9&lt;/num_query_props&gt;&lt;/n1:getTasks&gt;&lt;/env:Body&gt;&lt;/env:Envelope&gt;</pre>

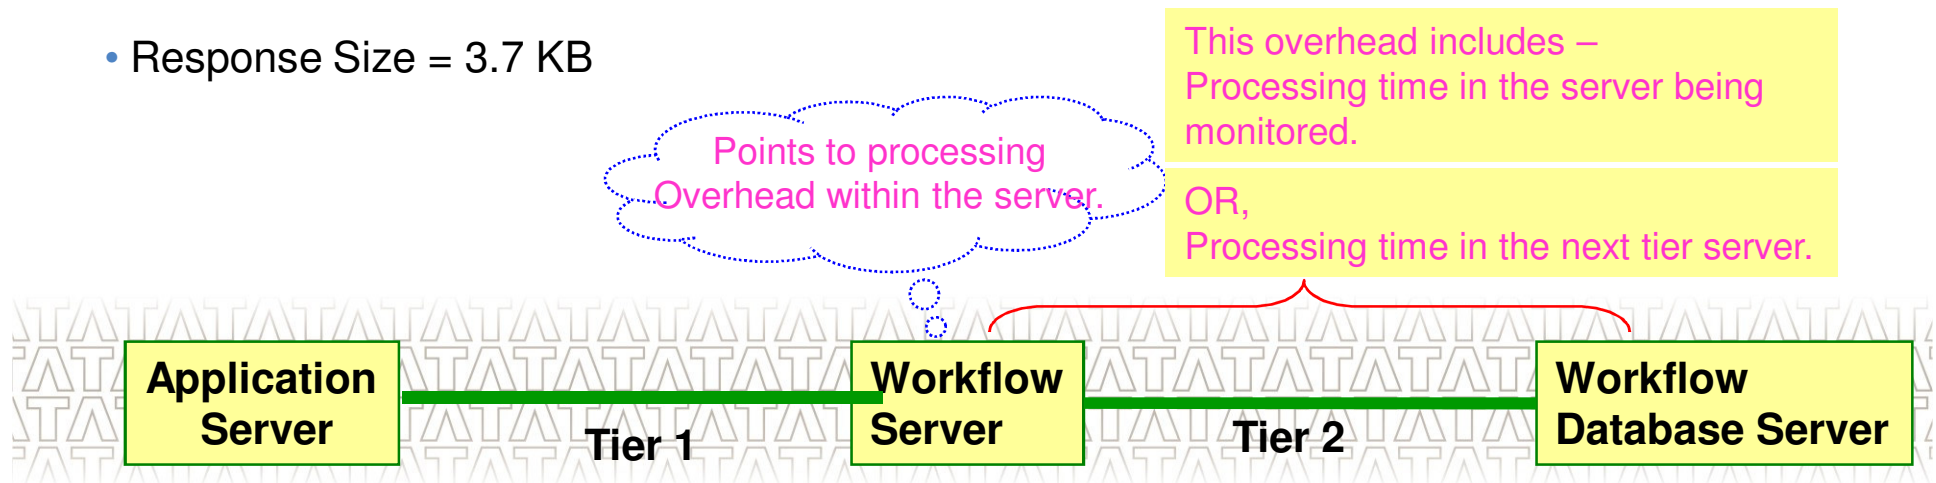


# Workflow Server Transaction 582- Example

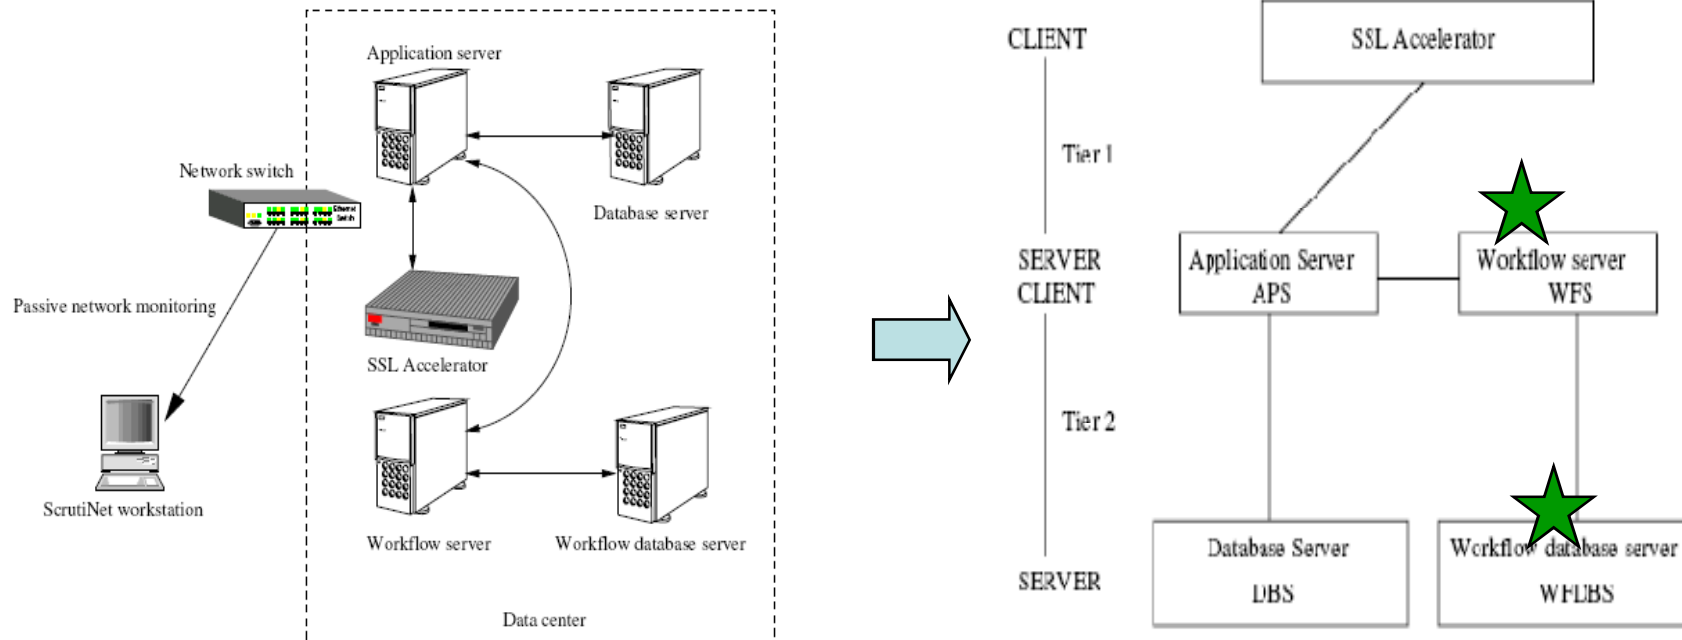
- Overall Response Time = **19.37** seconds
- Network Delay = 0 seconds
- Client Delay = 0 seconds
- Server Delay = **19.37** seconds
- Request size = 1.3 KB
- Response Size = 3.7 KB

**Client – Application Server**


**Server – Workflow Server**



# Correlation of transactions in the next tier




# Transaction Nesting



From the development team, it is understood that only one connection to the Workflow Database Server is made for each workflow transaction.

# ScrutiNet Correlation Report

Contribution of next tier to the response time



We can see that the connection on port **62277** contributes significantly to the response time of the workflow transaction **QUERY PROPERTY**.

Response Time of selected txn = 19.374 seconds			
Final txn resp time = 19.374			
Network overhead = 3.4e-005 (0.000175493 %)			
Client Delay = 0.000000 (0%)			
Next Tier overheads per connection .....			
Response Time in next tier for connection on port 55695 = 0.178429 (0.92097%)			
Response Time in next tier for connection on port 55697 = 0.206702 (1.0669%)			
Response Time in next tier for connection on port 49427 = 0.00307 (0.015846%)			
Response Time in next tier for connection on port 61982 = 1.31589 (6.79203%)			
Response Time in next tier for connection on port 62232 = 0.895182 (4.62053%)			
Response Time in next tier for connection on port 65439 = 0.000432 (0.00222979%)			
Response Time in next tier for connection on port 55685 = 0.132921 (0.686079%)			
Response Time in next tier for connection on port 65440 = 0.00067 (0.00345824%)			
Response Time in next tier for connection on port 65441 = 0.000555 (0.00286466%)			
Response Time in next tier for connection on port 62276 = 0.363481 (1.87613%)			
<b>Response Time in next tier for connection on port 62277 = 19.3564 (99.9089%)</b>			
Response Time in next tier for connection on port 62278 = 1.53336 (7.91451%)			
Response Time in next tier for connection on port 55694 = 0.02255 (0.116393%)			
Response Time in next tier for connection on port 58042 = 0.004071 (0.0210127%)			
Total overhead in next tier = 24.0137 (123.948%)			

**One database transaction takes 99% of the time!**

# Offending SQL Transaction

Request No	Request contents
73592	.....g.....WASD.\$6.....<.....O....."3.jv.l.....a: M.....O....."3.jv.l.....a: M.....
73593	.....^.)...W.....SELECT.DISTINCT.TA.TKIID.comma.TA.ACTIVATED.comma.TA.APPLIC_NAME.comma.TA.COMPLETED.comma.TA.LAST_MODIFIED.comma.TA.NAME.comma.TA.ORIGINATOR.comma.TA.OWNER.comma.TA.TKTID.comma.TA.STATE.comma.QP1.NAME..QP1NAME.comma.QP1.STRING_VALUE..QP1SVALUEcomma.QP2.NAME..QP2NAME.comma.QP2.STRING_VALUE..QP2SVALUEcomma.QP3.NAME..QP3NAME.comma.QP3.STRING_VALUE..QP3SVALUEcomma.QP4.NAME..QP4NAME.comma.QP4.STRING_VALUE..QP4SVALUEcomma.QP5.NAME..QP5NAME.comma.QP5.STRING_VALUE..QP5SVALUEcomma.QP6.NAME..QP6NAME.comma.QP6.STRING_VALUE..QP6SVALUEcomma.QP7.NAME..QP7NAME.comma.QP7.STRING_VALUE..QP7SVALUEcomma.QP8.NAME..QP8NAME.comma.QP8.STRING_VALUE..QP8SVALUEcomma.QP9.NAME..QP9NAMEcomma.QP9.STRING_VALUE..QP9SVALUE.FROM.QUERY_PROPERTY.QP3comma.QUERY_PROPERTY.QP8comma.QUERY_PROPERTY.QP5comma.QUERY_PROPERTY.QP7comma.QUERY_PROPERTY.QP1comma.QUERY_PROPERTY.QP2comma.QUERY_PROPERTY.QP6comma.QUERY_PROPERTY.QP4comma.WORK_ITEM.WIcomma.TASK.TAcomma.QUERY_PROPERTY.QP9.WHERE.(WI.OBJECT_ID.=.TA.TKIID.AND.TA.CONTAINMENT_CTX_ID.=.QP5.PIID.AND.TA.CONTAINMENT_CTX_ID.=.QP7.PIID.AND.TA.CONTAINMENT_CTX_ID.=.QP1.PIID.AND.TA.CONTAINMENT_CTX_ID.=.QP2.PIID.AND.TA.CONTAINMENT_CTX_ID.=.QP6.PIID.AND.TA.CONTAINMENT_CTX_ID.=.QP4.PIID.AND.TA
73594	.....m..
76104	.M.....g.....WASD.\$6.....<.....WASD.....\$.....6.....O....."3.jv.l.....a: M.....O....."3.jv.l.....a: M.....
76105	.....w.....O....."3.jv.l.....a: M.....O....."3.jv.l.....a: M.....

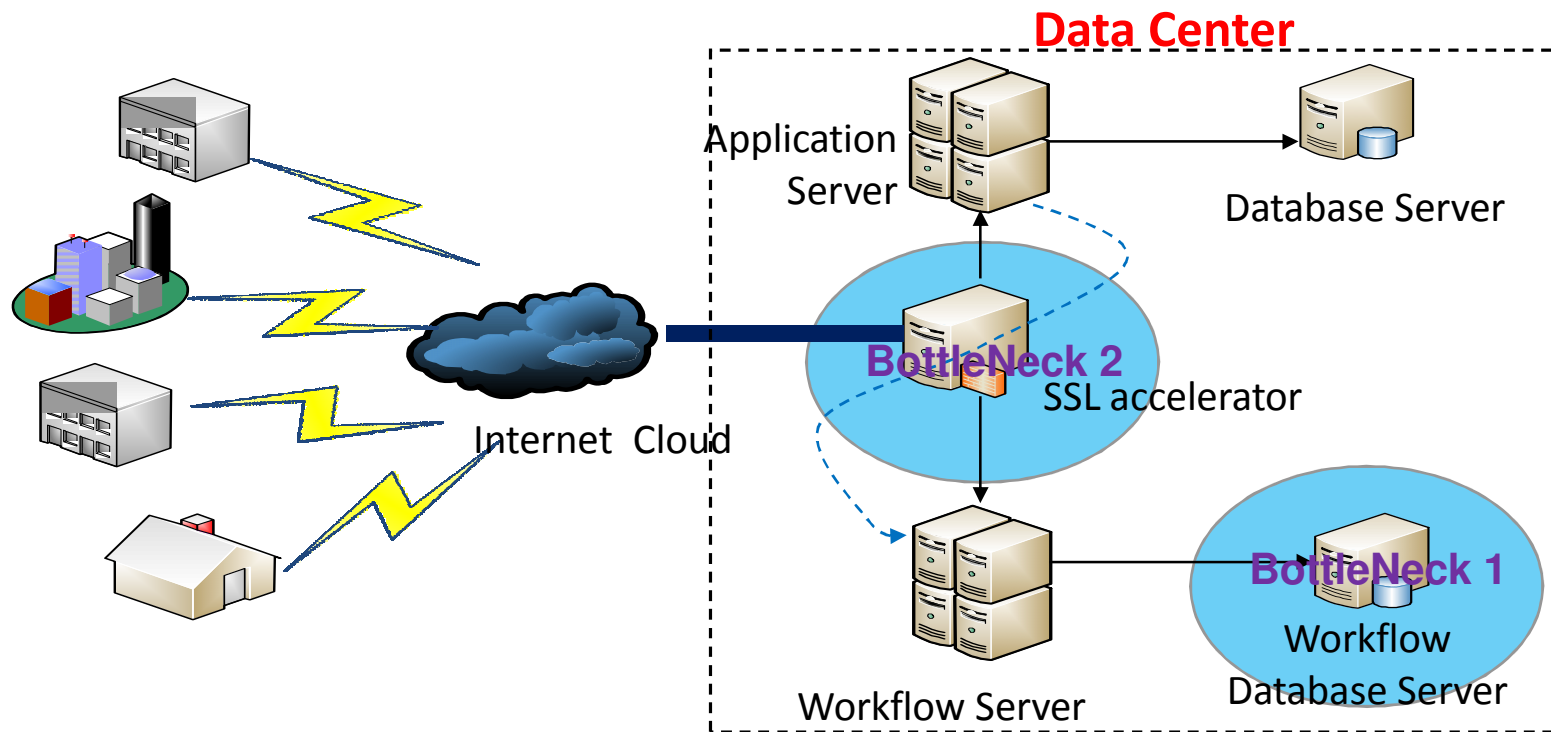
Transactions on connection port 62277 on the Workflow Database Server for the execution of QUERY PROPERTY transaction on the Workflow Server

# Useful aid to correlation analysis

- Obtain a ScrutiNet Log with no concurrency for all transactions in a test environment.
  - i.e. on a single user using the system
- The output is a signature for every end user transaction on every server
- This information is very useful for correlation analysis of transactions across tiers



# An OLTP System – Bottleneck Identified



# Limitations

- Pipelining
  - Transactions cannot be identified accurately
  - However, It can be still useful as a bottleneck detection tool
- Difficult to determine transaction type in encrypted packets
  - Response time delay and other info is still available
- Presence of NAT devices
  - Difficult to identify clients or their access locations when the end users IP addresses are NATed





# Conclusion

# Summary

- In this paper we presented our approach of Performance Monitoring and Analysis of an OLTP System with an example.
- Passive Monitoring approach is **most beneficial for benchmarking** exercises where the system being measured is not affected.
- ScrutiNet is a passive performance monitoring tool developed by the authors at Tata Consultancy Services
  - A black box approach to begin with
  - Availability of application level information speeds eases analysis effort.
- Authors can be contacted for more information on ScrutiNet.





Thank You – Q&A  
[m.nambiar@tcs.com](mailto:m.nambiar@tcs.com)