



# Overcoming the Top Four Challenges to Real-Time Performance in Large-Scale, Data-Centric Applications

Tom Lubinski  
Founder and CEO  
SL Corporation

17 November 2011

- 
- Disclaimers
  - I am not Mike Lee
  - No Mariachi hat
  - No Facial hair
  - A LOT more boring
  - My other computer is a Mac
  - However, we have “shipped” ...



# Select SL RTView Customers



## Financial Services



## E-Commerce/Retail



## Energy



## Telecommunications



## Other



- Software Product company since 1983
- Headquarters in Marin County, CA
- Worldwide presence in Americas, APAC, EMEA
- Over 100,000 licenses sold
- Core expertise in application performance monitoring – special focus on middleware



- Here to talk about Scalability and Performance
- Problem Space:

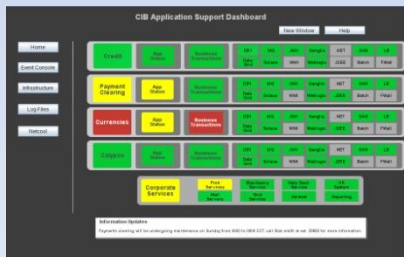
*Collection, Analysis, and Visualization in Real-Time of large volumes of monitoring data from large-scale, complex, distributed applications*

*Keywords: Real-Time, Large Volumes of Data*



## Application Performance Monitoring

- Collect all necessary application-centric and middleware-centric performance data
- Configure data aggregation and persistence, filters, analytics, alerts and displays to deliver information tailored for app support teams




## Oracle Coherence Monitoring

- Understand the behavior of Coherence
- Debug and validate functionality after configuration changes
- Integrate OCM with existing monitoring tools
- Enable quick notification of problems

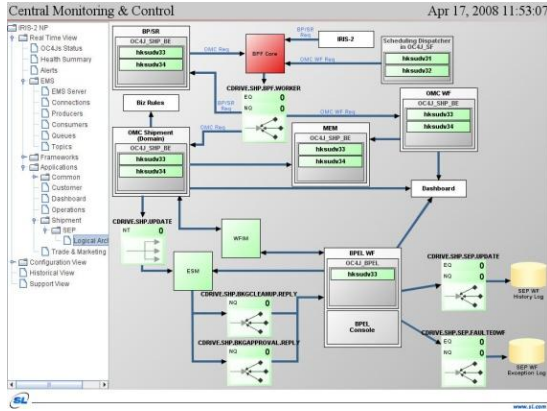


## Middleware Monitoring

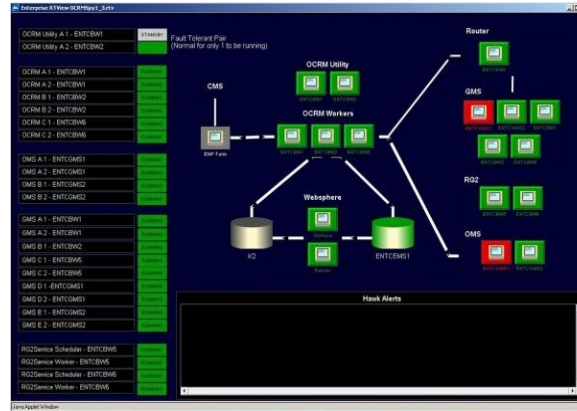
- Determine how applications are interacting with middleware systems
- Assess whether applications are running efficiently and reliably
- Ensure the maximum benefit from an ESB investment



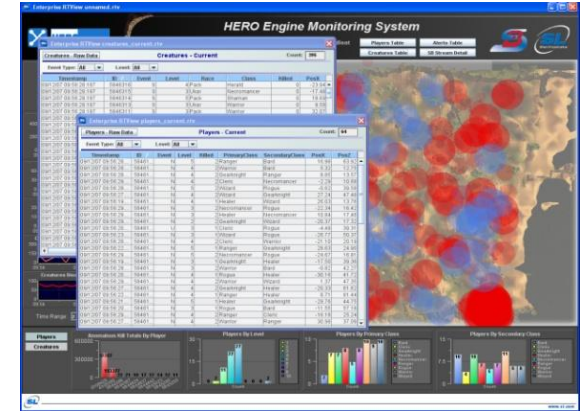
# RTView – Sample Applications



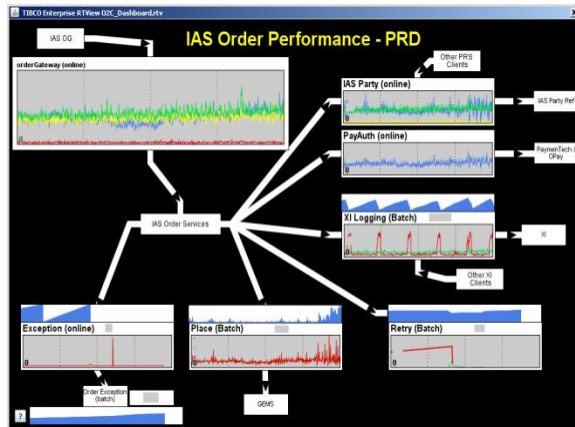
OOCL World Wide Shipment Tracking



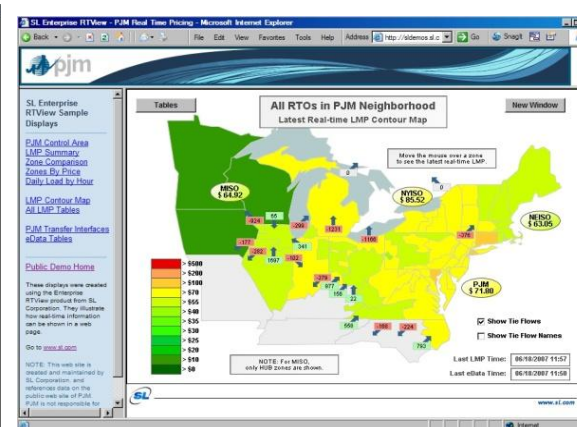
Hospitality Card application at Harrah's casino gaming tables



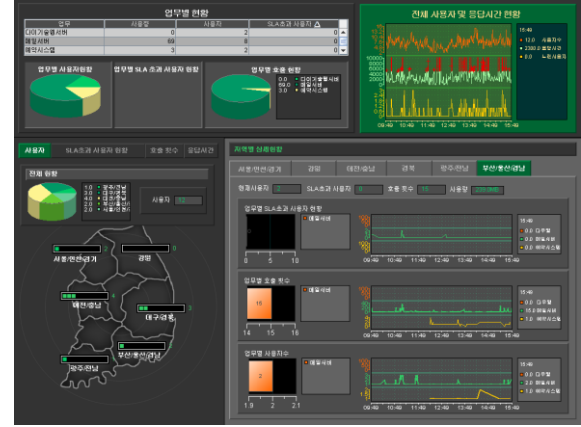
Online Gaming Systems



Tax Season at Intuit



PJM Real-time Energy Pricing

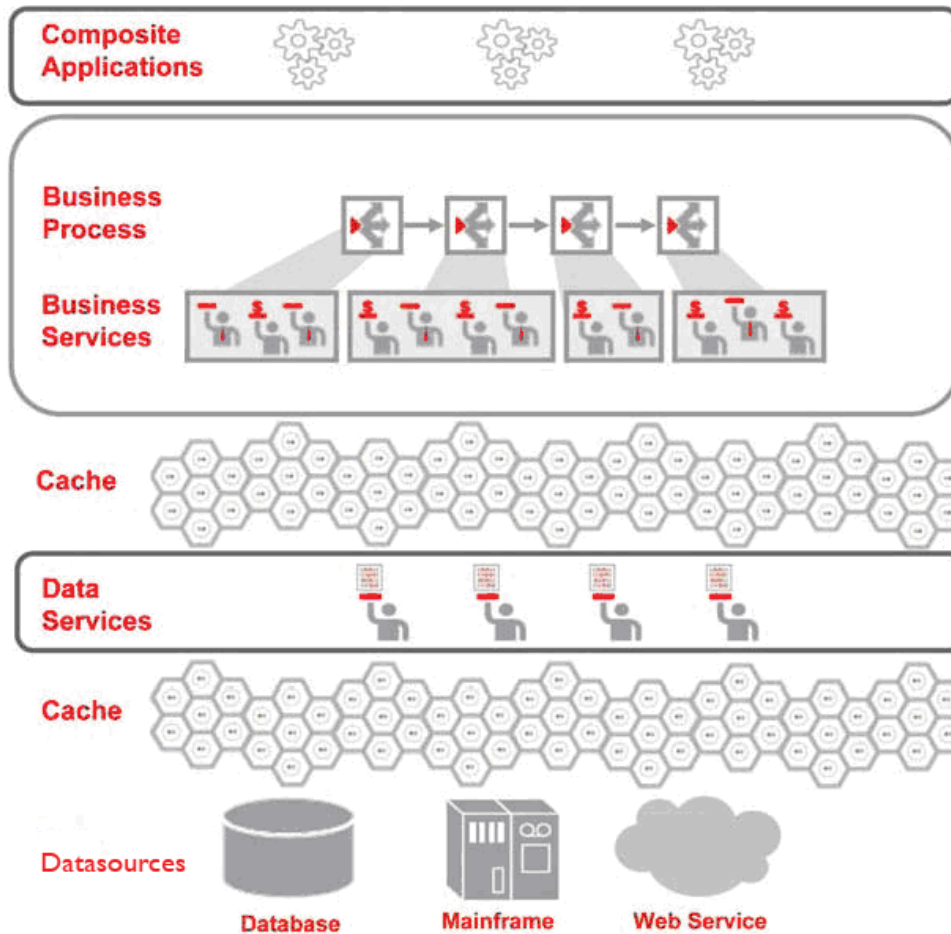


Banking application in Korea

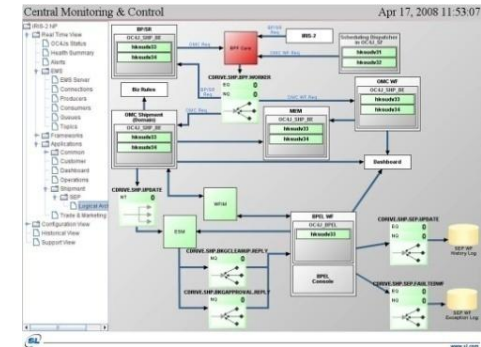
# RTView – Multi-Tier Visibility



## Caching in an SOA Environment



Unified Real-time display of data from all Application tiers



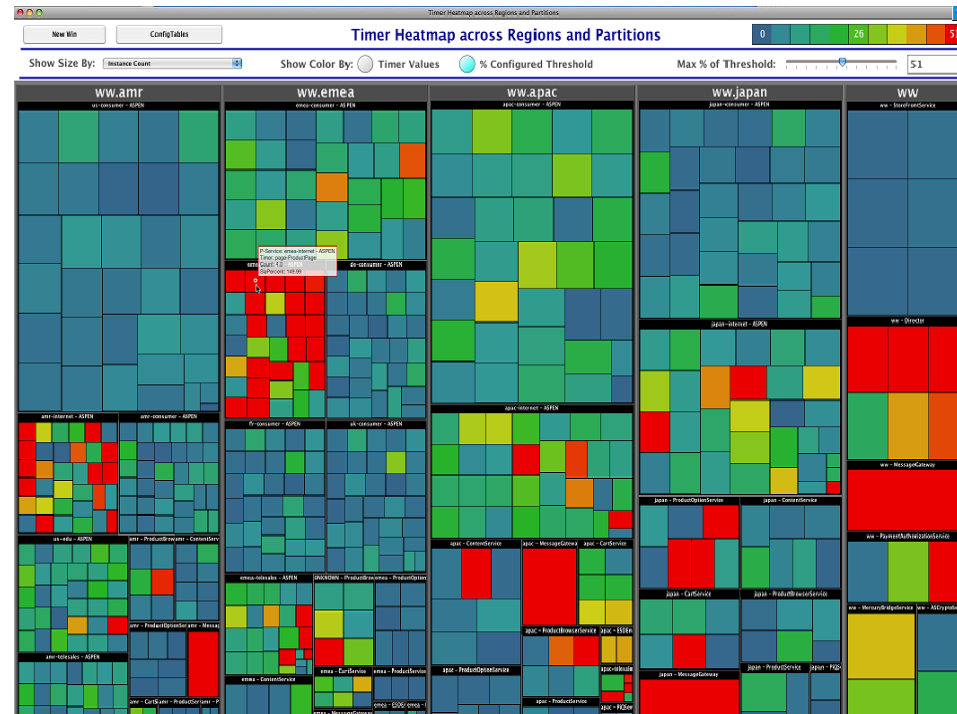
In-depth Monitoring of Middleware Components



# RTView – Large Data Volumes




- Typical large implementation, distributed over several regions with many custom applications
- Heatmap View showing current state of entire system – size represents number of servers for application
- Color represents how close metric is to SLA – large red boxes are worst – drilldown to detail



# RTView - Drill-Down to Detail Metrics



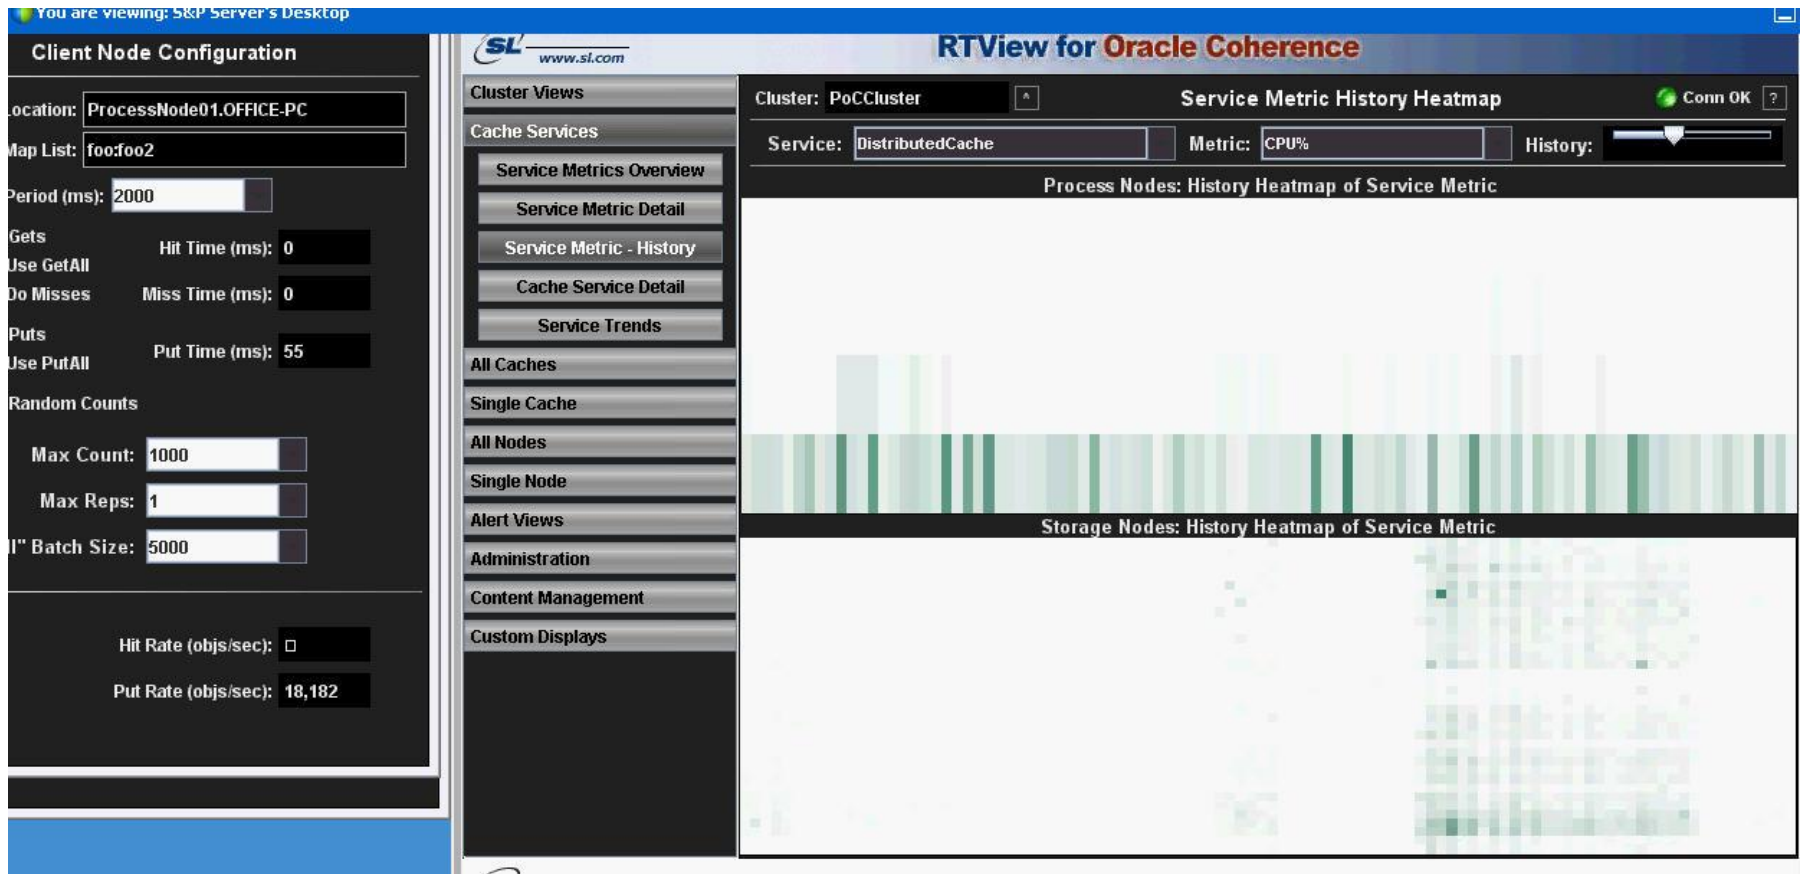
- Drilldown to detail level metrics showing internal metrics from each application
- Sophisticated history and alert view allows fine-tuning of thresholds for each metric



## Alert Status



## Observe “internal load balancing” of Data Grid



The screenshot displays the RTView for Oracle Coherence interface. On the left, the 'Client Node Configuration' panel shows settings for 'ProcessNode01.OFFICE-PC' with a map list of 'foo:foo2' and a period of 2000 ms. Performance metrics include Hit Time (0 ms), Miss Time (0 ms), and Put Time (55 ms). The 'Random Counts' section shows a Max Count of 1000, Max Reps of 1, and a Batch Size of 5000. The 'Hit Rate (objs/sec)' is 0, and the 'Put Rate (objs/sec)' is 18,182.

The main interface shows the 'RTView for Oracle Coherence' window with the cluster set to 'PoCCluster'. The 'Service Metric History Heatmap' is displayed for the 'DistributedCache' service, monitoring 'CPU%' over time. The heatmap is divided into two sections: 'Process Nodes: History Heatmap of Service Metric' and 'Storage Nodes: History Heatmap of Service Metric'. The process nodes heatmap shows a dense pattern of green bars, indicating high activity. The storage nodes heatmap shows a sparser pattern of green bars, indicating lower activity.



- Challenge #1:

## Database Performance

Common to see queries taking minutes  
*How can you get real-time that way ?*





- Challenge #2:

## Network Data-Transfer Bandwidth

Bigger pipes, but there's more data to send  
*How do you get the greatest throughput ?*



- Challenge #3:

## Processor Performance

More cores just means more processes !  
*How do you optimize your utilization ?*



- Challenge #4:

## Lack of Real-Time Predictability

Virtualization is the new time-share !

*How can you trust your data ?*

*"time-sharing", "network computer", "cloud", do things ever really change ?*



- Solution – Clues ?
- Facts of Life:

Database – can't live with it, can't live without it

Network – it's a funnel, no way around it

Processor – must limit what you ask it to do

Virtualization - it's erratic, have to compensate





- Solution #1:

## Proper Data Model

Data structures designed for real-time  
In-memory structures to buffer database



*Can your application be ...*



*... like a high-performance racecar ?*



***What is most important part of racecar ?  
(besides the engine)***



**... the Transmission ...**



*For Real-Time performance, it's the **Cache** ...*



**Not a simple  
"current value"  
cache**



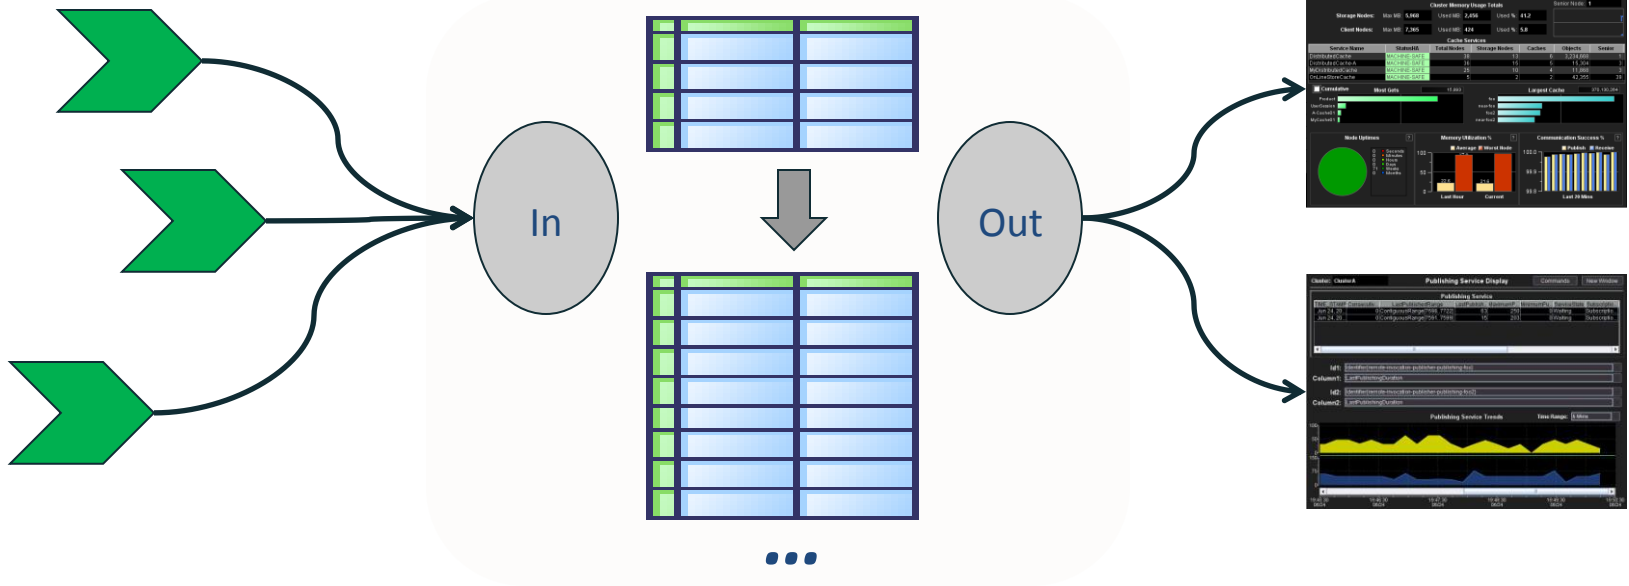
**High-performance  
Real-time Multi-dimensional  
Data Cache**





# Real-Time Cache – optimized for performance !

Current / History Tables:




Indexed Insertion - asynchronous real-time data

Indexed extraction - optimized transfer to clients



## *Real-Time Cache – Data Processing / Aggregation*

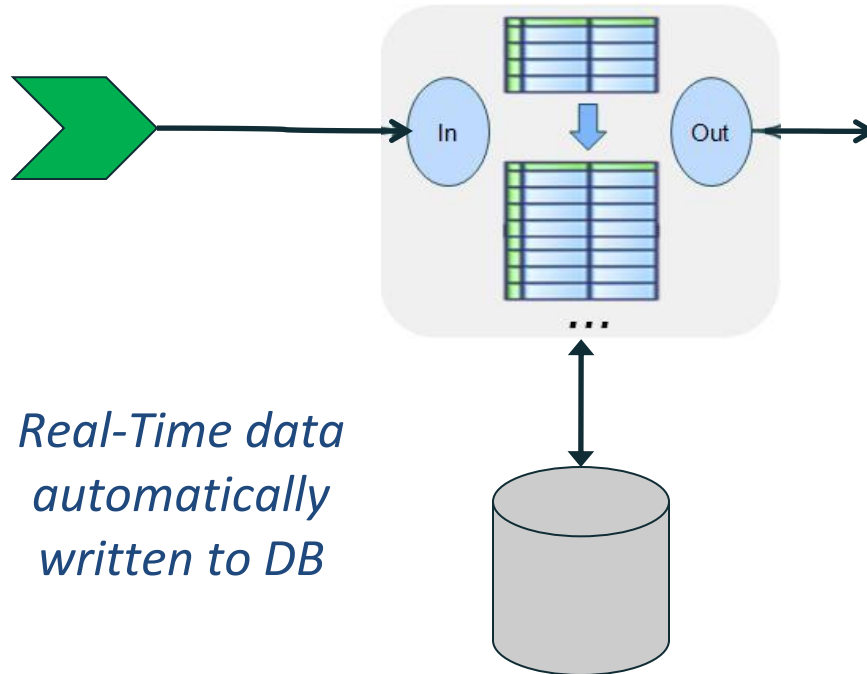


*Reduction,  
Resolution, Aging*


*Aggregation*



## *Real-Time Cache – Database read/write through (optimized for timestamped multi-dimensional data)*



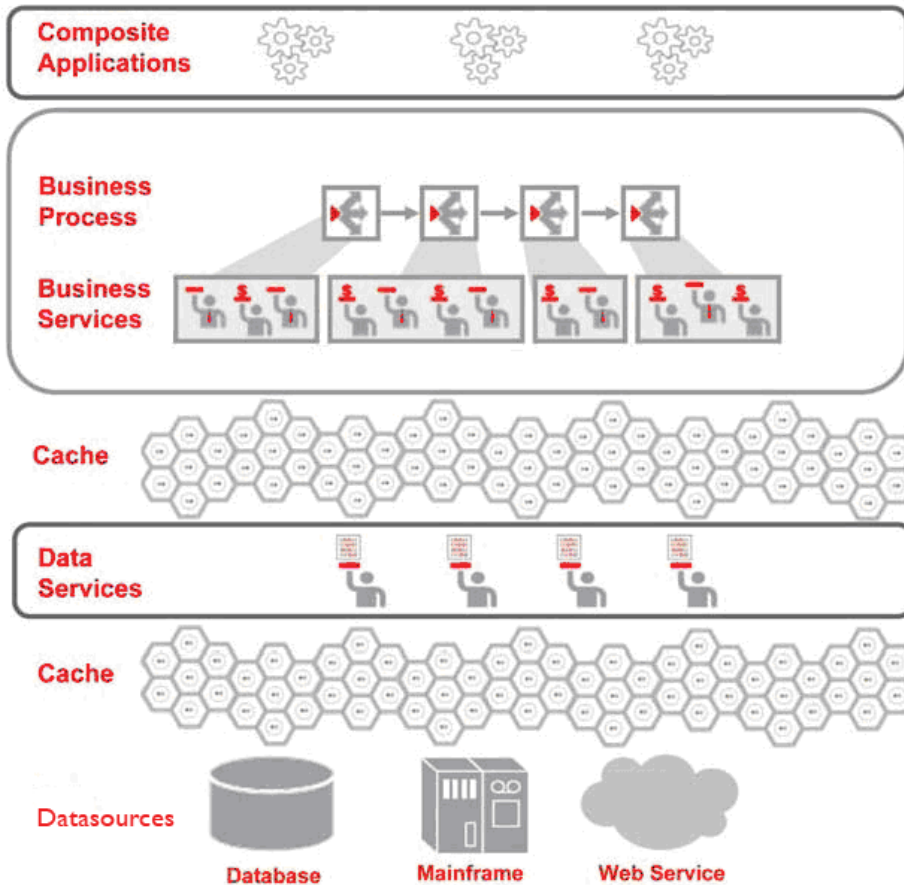
*Real-Time data  
automatically  
written to DB*



*Seamless timeline  
navigation with automatic  
database query*

# This sounds a bit like Oracle Coherence ...

## Caching in an SOA Environment



Buffer database

Read/write through

Listeners

Indexed queries

*What's different ?*

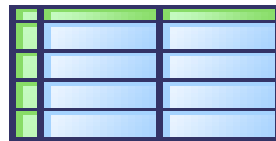


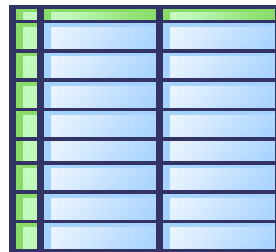


# Different tools for different problems !

## Real-Time Multi-dimensional data:

Current / History Tables:




...

**Multiple rows (time range) of selected columns** returned in one query

Coherence cache distributes objects (rows) = optimized horizontally

Real-Time multi-dimensional cache manages columns and optimizes vertically

---

## Benefits: Indexed Real-Time Caching

Slow SQL queries minimized

Users shielded from database details

Minimize CPU load using effective indexing



- Solution #2

## Server-Side Aggregation

(am I being too obvious with this one ?)

Know the use cases

Joins and GroupBy done on server

SQL does this, but do you need it ?



---

## Problems with SQL Database Queries

Slow

Slower with concurrent queries

If you need it fast, it goes even slowwwwwer !

SQL = Not portable

(Timestamps, especially)



---

Know your problem space !

Real-Time Monitoring:

Join and GroupBy heavily used

We wrote our own!


Performed in real-time on server-side data

Optimized for real-time requirements






- Example: Server-Side Aggregation/Caching



- Each cache can maintain its own history



- Result: trend chart of Totals by History has all data available immediately
- Using SQL would require:

Query 3 tables

2 GroupBys, 2 Joins, + Join on Timestamp (not portable)



---

## Benefits: Server-Side Aggregation

Client requests and gets exactly what is needed

Client processing = zero

Server processing = done ahead of time

Current/History for aggregates readily available (No SQL)

Response time = fast



- Solution #3

Use Appropriate Design Patterns

Server-Side vs. Client-Side Processing  
Efficient Data Transfer Patterns






- Pattern #1:

## Data Compaction

(obvious, initial approach for any data transfers)




*... even simple, non-proprietary algorithms can make big difference*

- Pattern #2:

## Data Current / Changed

(large data tables with sparse real-time updates)




*... little more complex, requires indexing*

- Pattern #3:

## Data History / Current

(trend chart invoke with real-time updates)

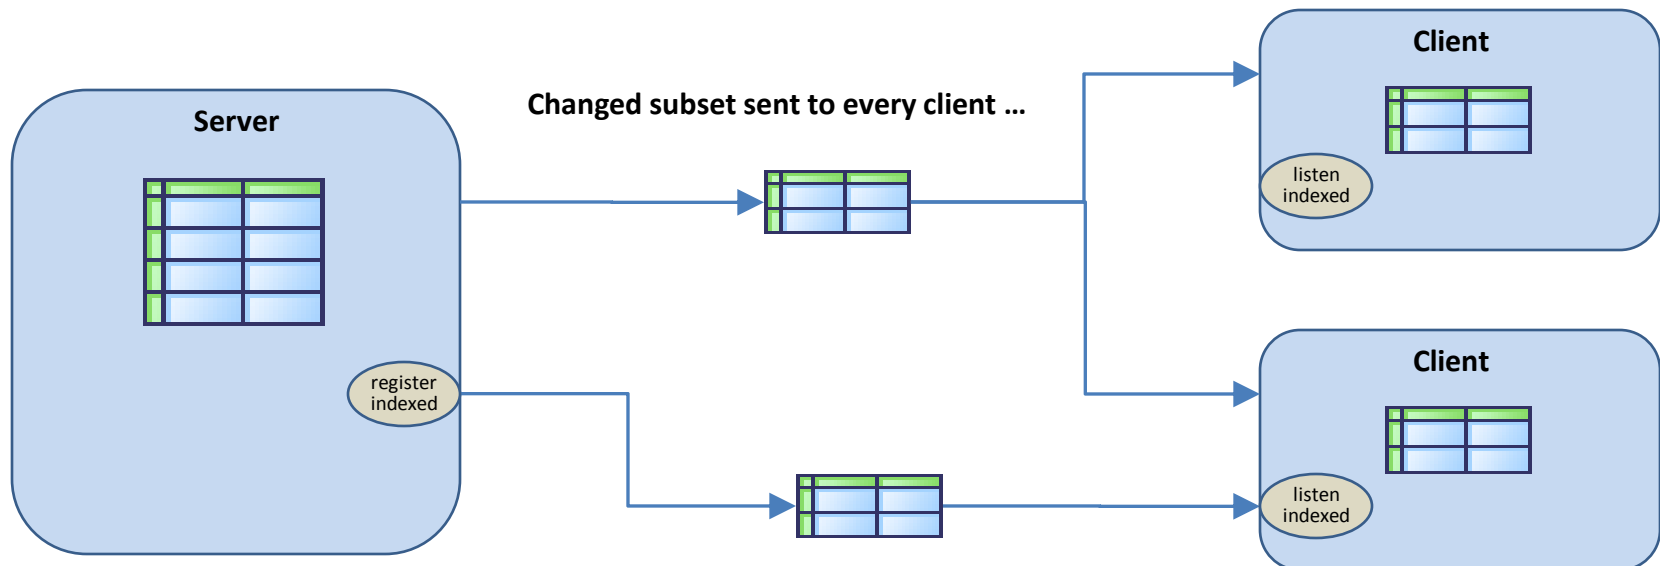


... similar to current / changed pattern, but specific to history

- Pattern #4:

## Data Current / Subset

(optimizing transfer of data subsets to multiple clients)



*... instead, send subset only to registered client*

*... requires registration logic coupled with cache*



---

## Benefits: Design Patterns for Data Transfer

Same problem over and over again solved similar way

Reduce load on network

Optimize response time – no unnecessary data





- Conclusion #1:

Know your data !

Data Model designed for real-time  
In-memory structures to buffer database  
Server-side aggregations



- Conclusion #2

Respect Design Patterns !

Server-Side vs. Client-Side Processing

Efficient Data Transfer Patterns

Don't over-generalize – solve the problem





# Questions?

See [www.sl.com](http://www.sl.com)  
for more into about SL and RTView