



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

Tom Lubinski
Founder and CEO
SL Corporation

7 March 2012

-
- Disclaimers
 - In 30 years, we've learned a lot
(a grizzled veteran)
 - But, we don't know everything ...
 - ... we could be wrong !
 - My other computer is a Mac
 - We have "shipped" ...



Background



Extensive background in real-time process monitoring

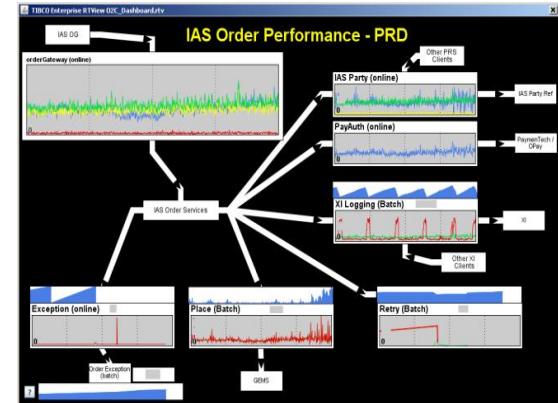
Large volumes of dynamic data

Visualization technologies

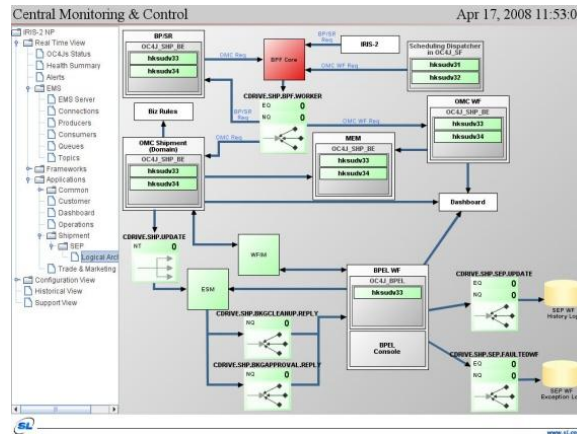
Mission-critical applications



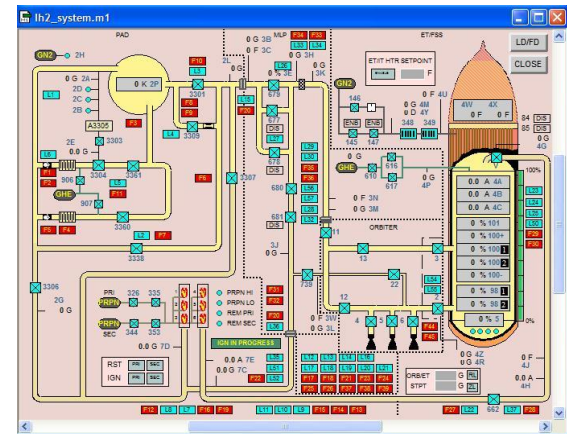
Connecticut Valley Power Grid Management System



Critical Tax Season Applications at Intuit



OOCL World Wide Shipment Tracking



NASA Space Shuttle Launch Control System

- 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

Emphasis: Real-Time, Large Volumes of Data



- 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:

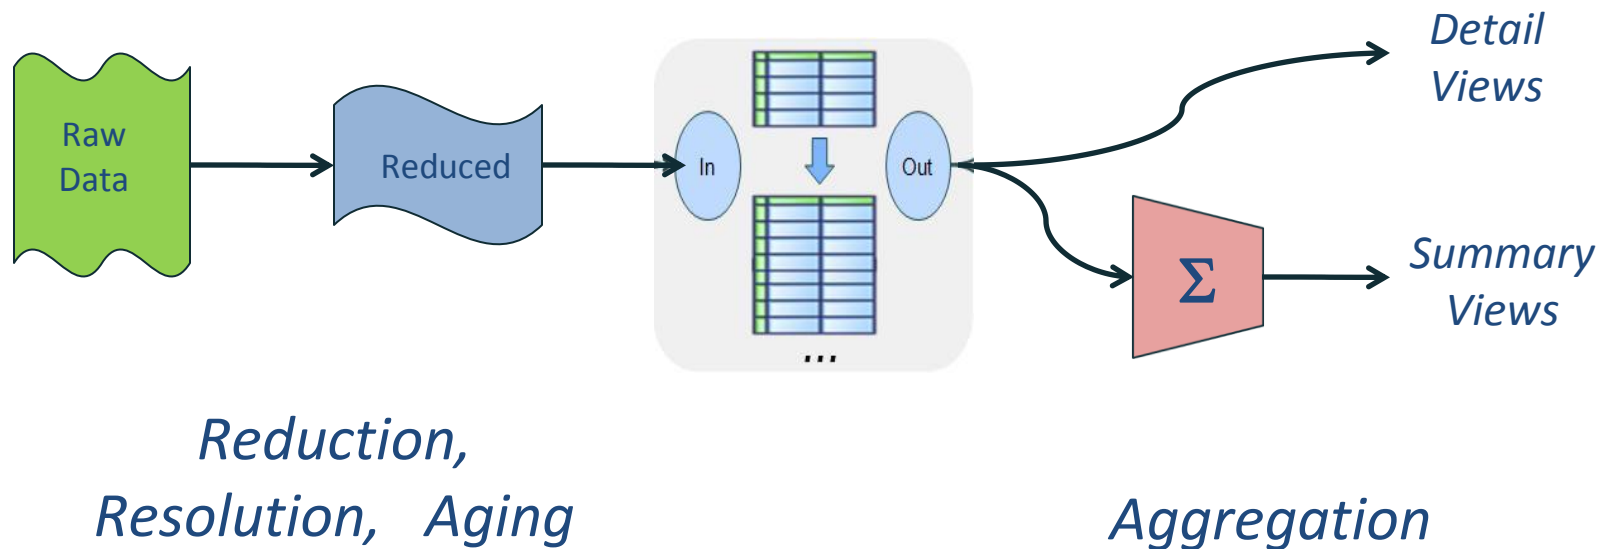


Indexed Insertion - asynchronous real-time data

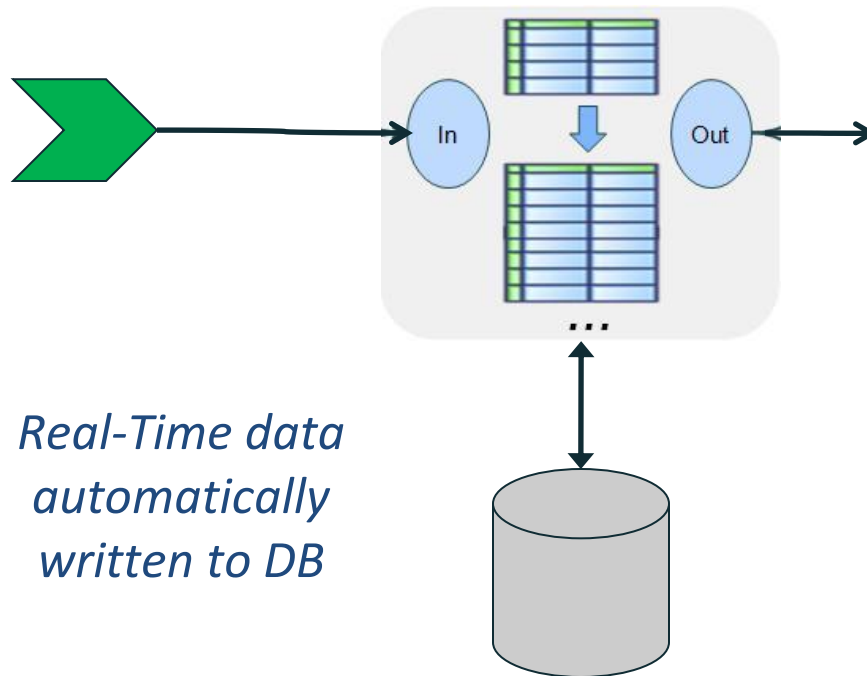
Indexed extraction - optimized transfer to clients



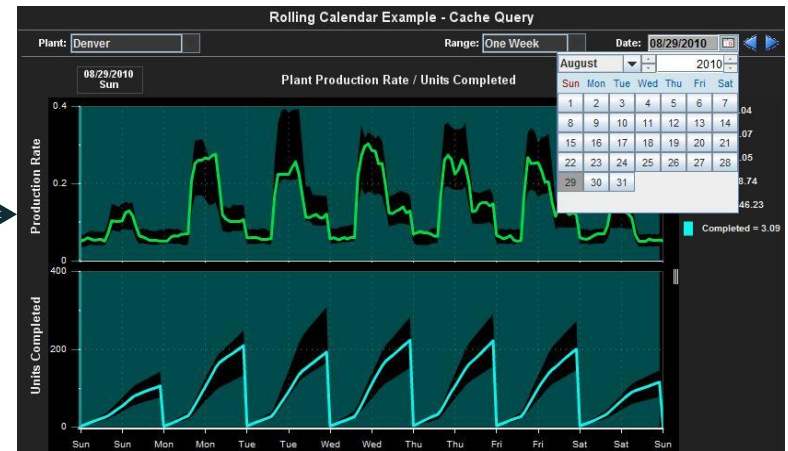
Real-Time Cache – Data Processing / Aggregation



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 ?



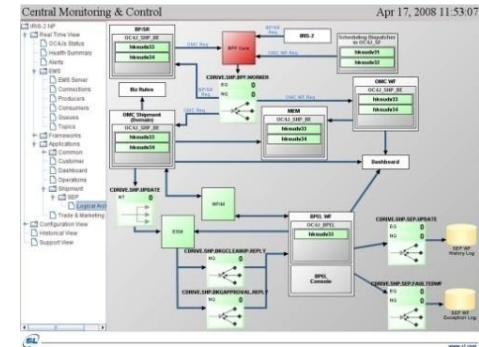
Multi-Tier Visibility into Monitoring Data



Caching in an SOA Environment



Unified Real-time display of data from all Application tiers

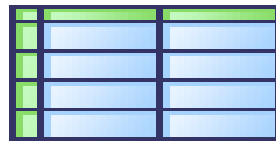


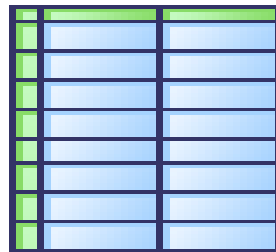
In-depth Monitoring of Middleware Components

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



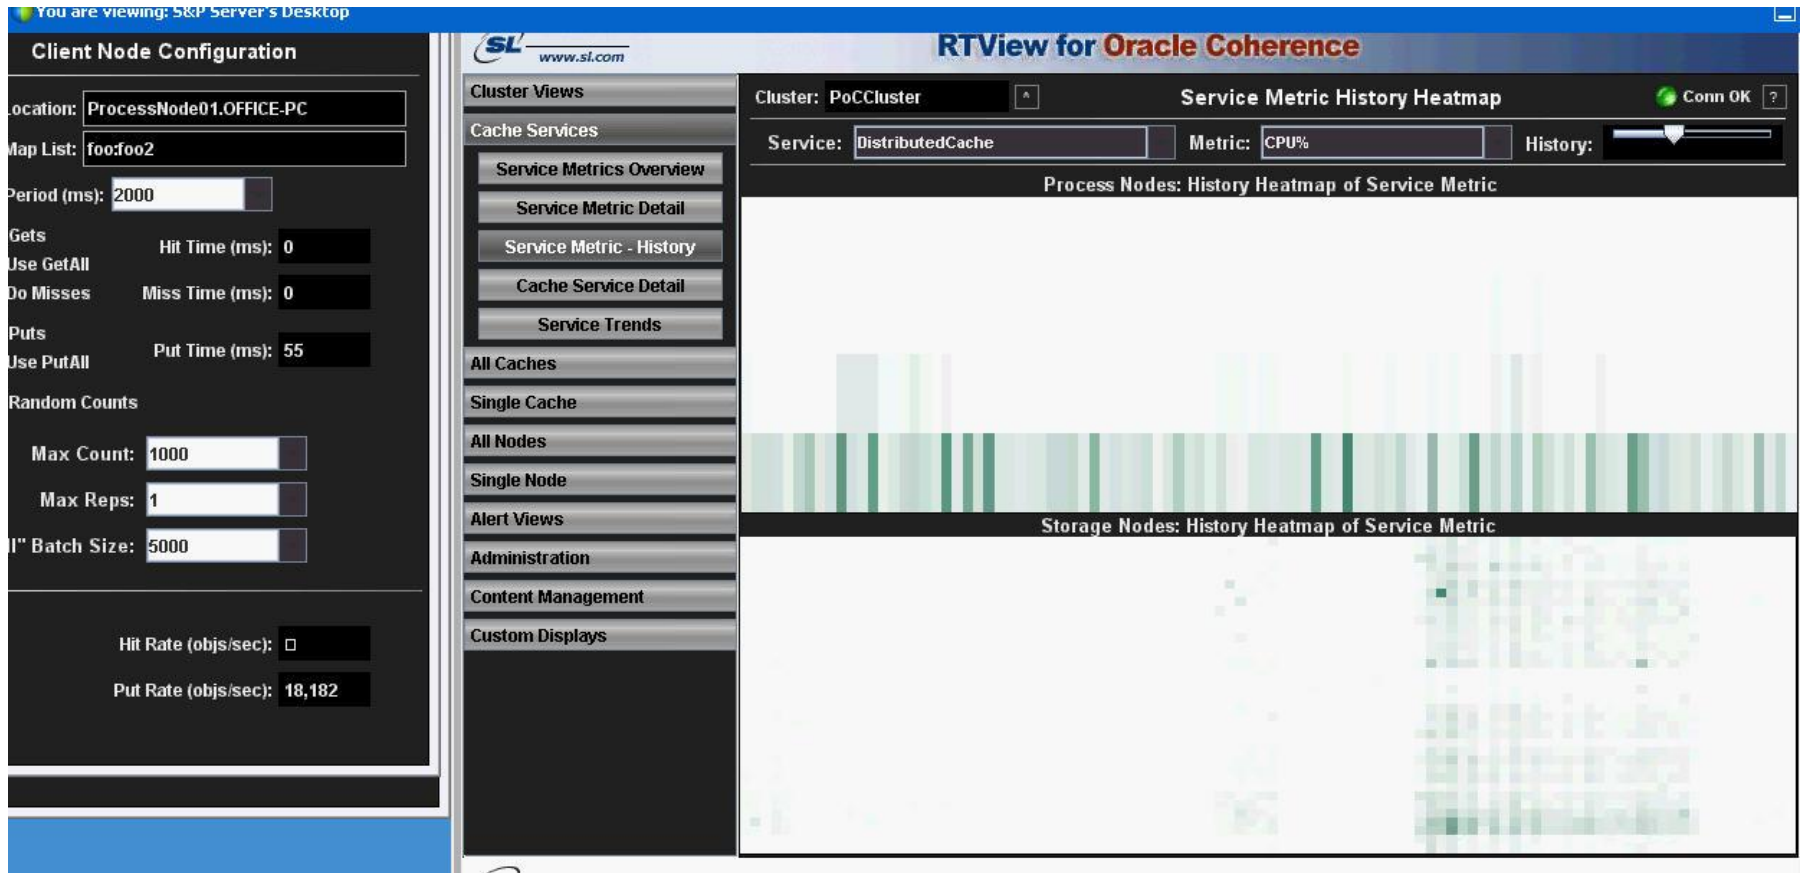
Display of Large Data Volumes



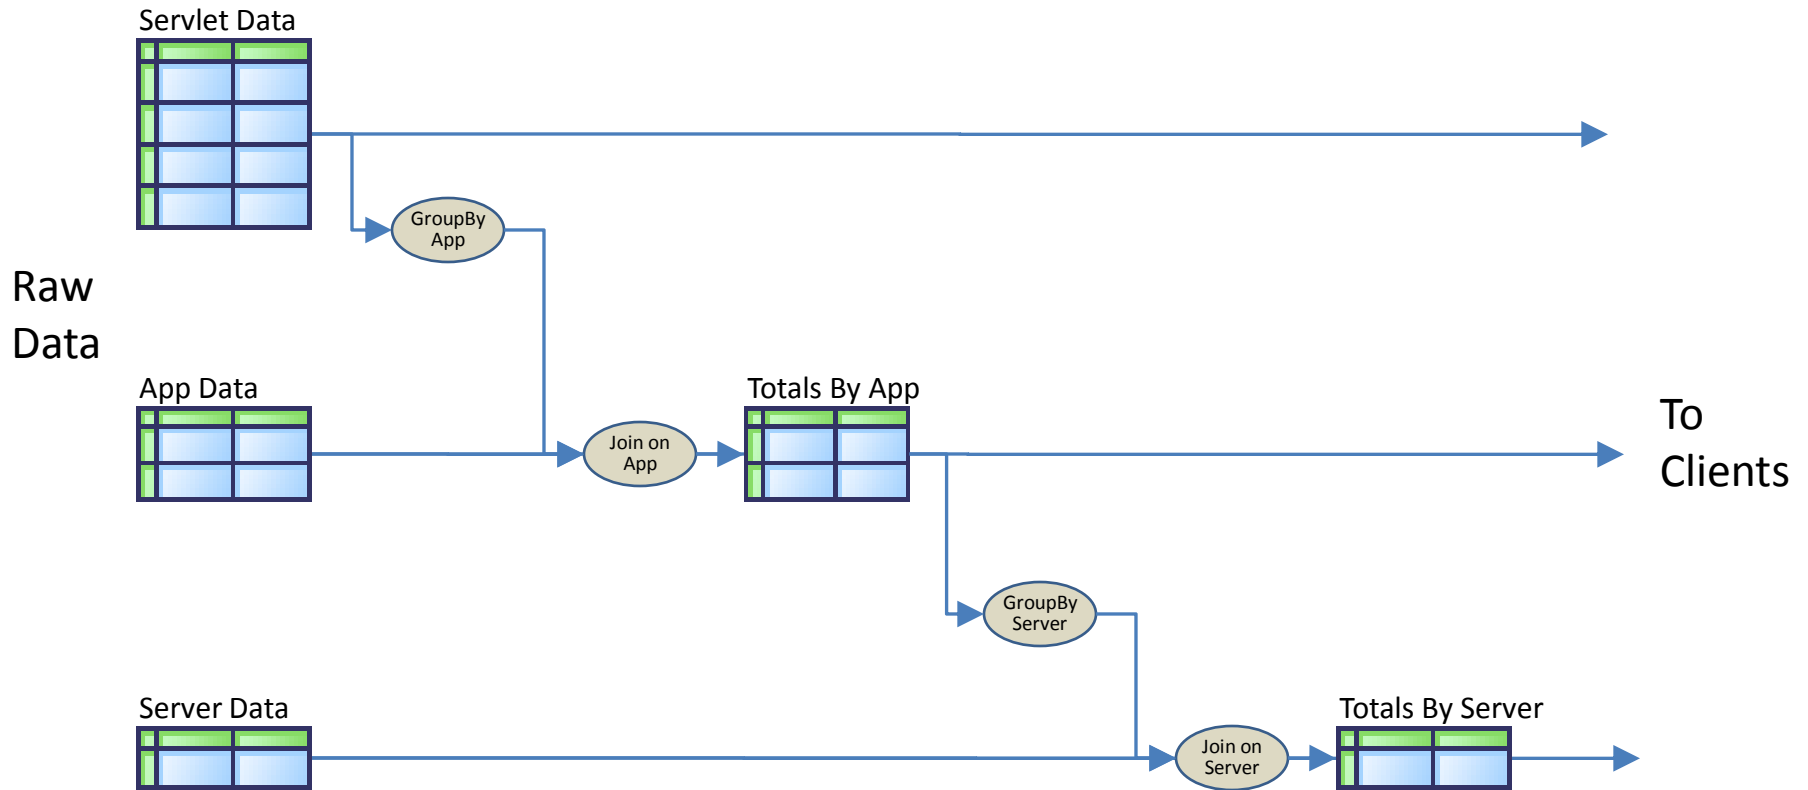
- Typical large implementation, distributed over several regions with many custom applications
- Heatmap View showing current state of entire system – size represents number of servers for application
- Color represents how close metric is to SLA – large red boxes are worst – drilldown to detail



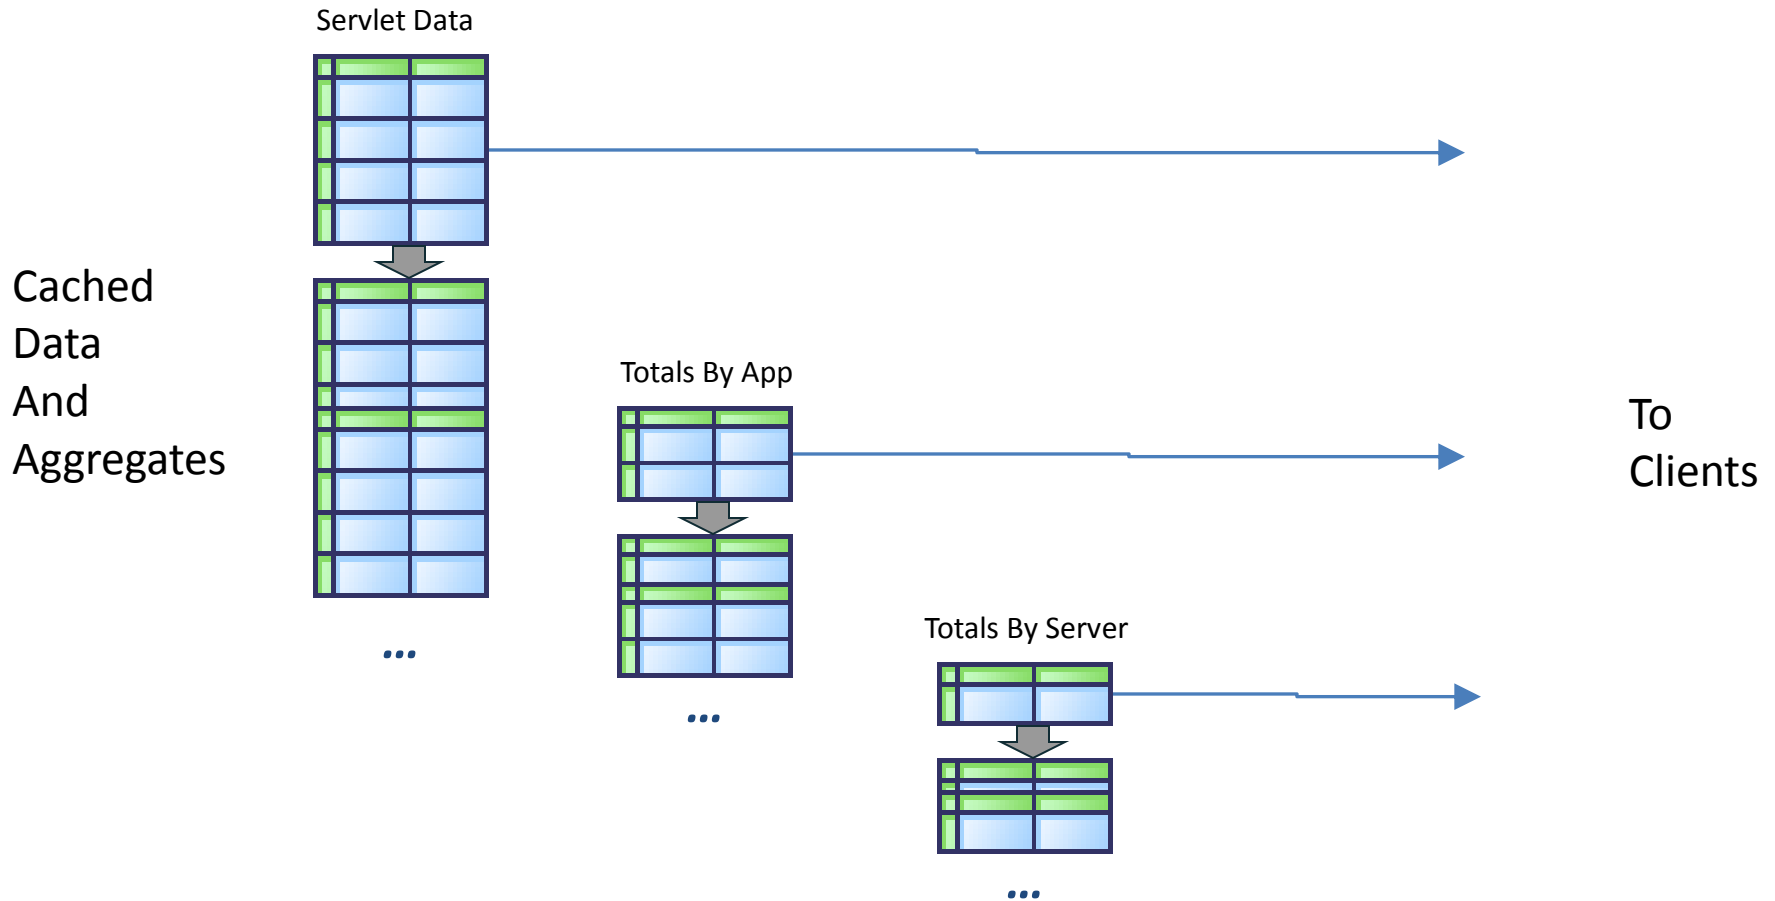
Observe "internal load balancing" of Data Grid



- 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)

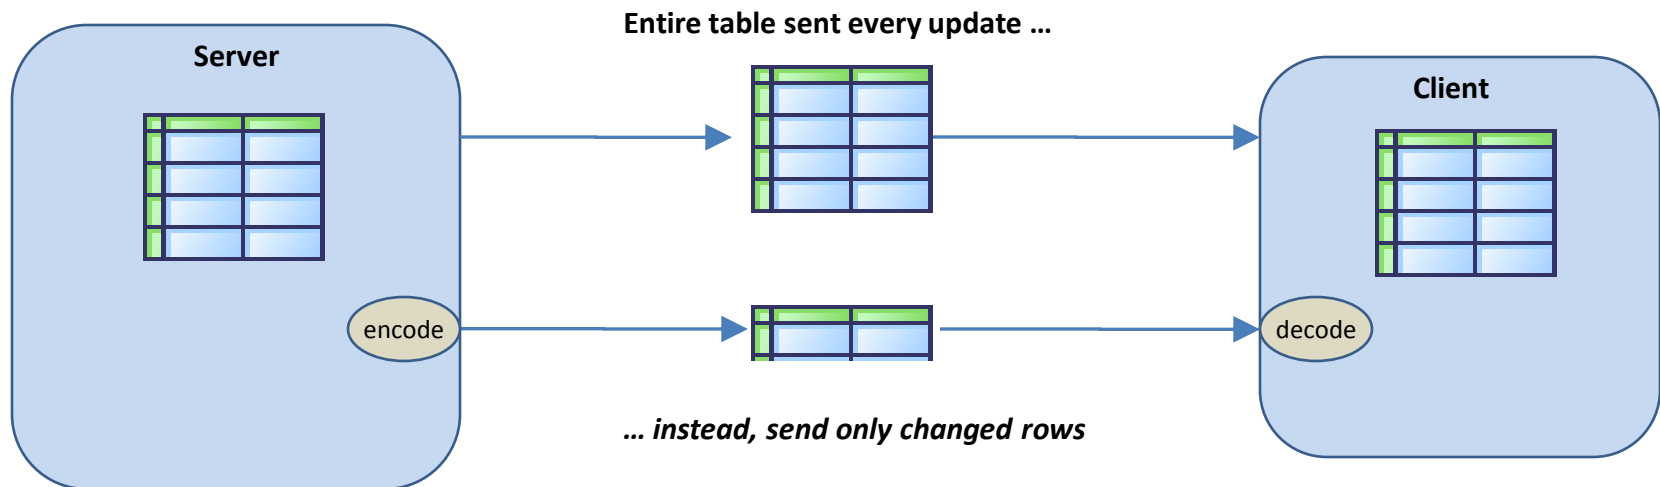


... even simple, non-proprietary algorithms can make big difference

- Pattern #2:

Data Current / Changed

(large data tables with sparse real-time updates)

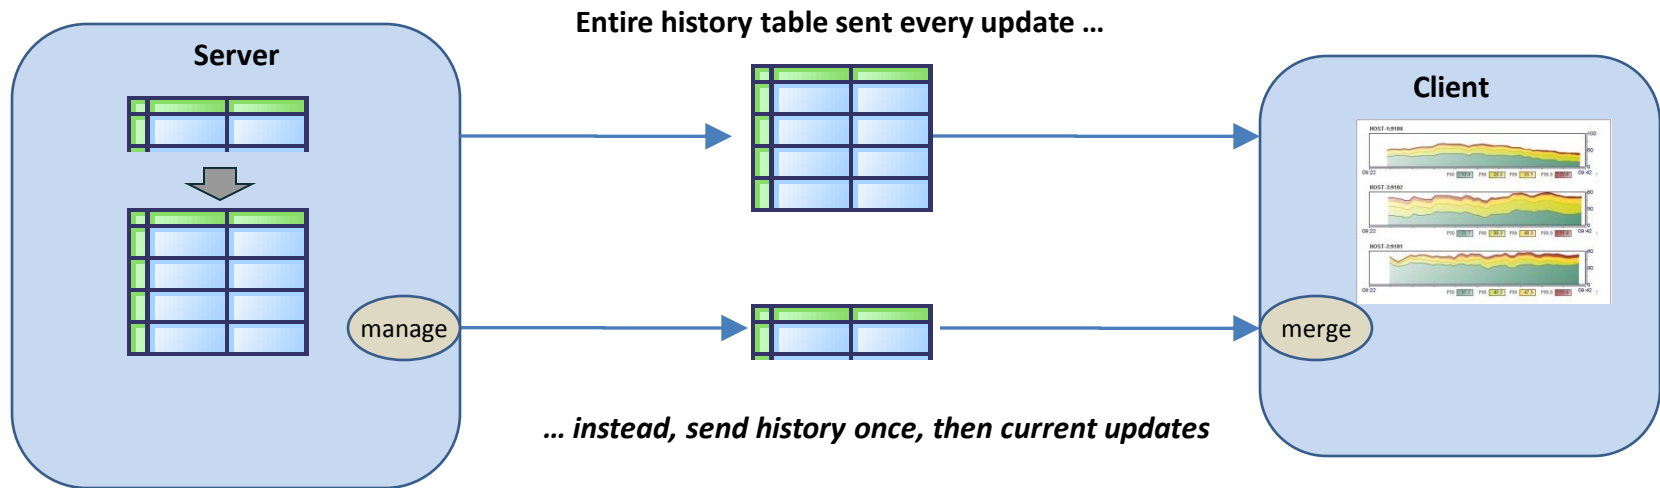


... little more complex, requires indexing

- 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

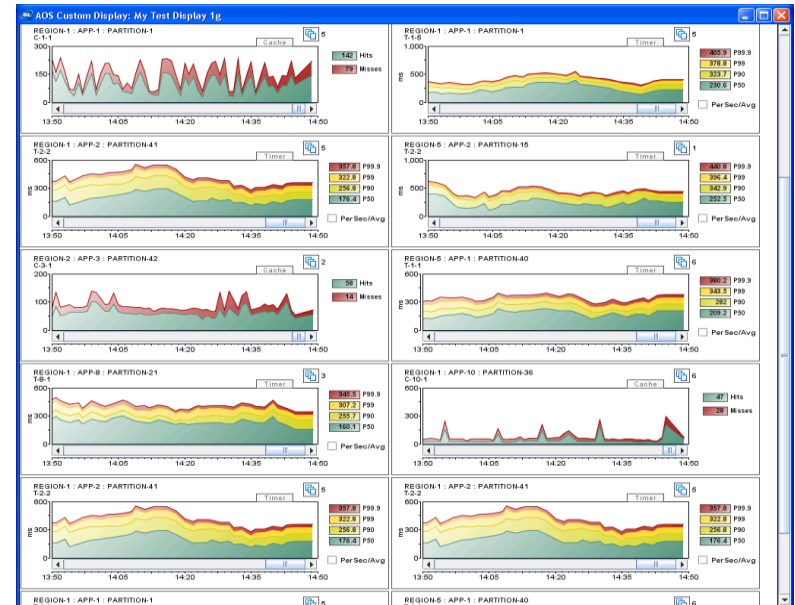
... requires registration logic coupled with cache



Drill-Down to Detail Metrics



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



Alert Status

Site: Service: Partition:

Timer	P50th	P90th	P95th	P99.9th	Status	Minutes of Violation	Severity	Group
Traffic ASPEN	NA	NA	2000	4000	FAILING	2	ERROR	APP-DEV
Traffic ASPEN Cached	NA	NA	1200	1400	OK	0	ERROR	APP-DEV
Traffic ASPEN Errors	NA	NA	10000	NA	OK	0	ERROR	APP-DEV
page-ASCConfigurableProductAjaxResponse	NA	NA	100	200	OK	0	ERROR	APP-DEV
page-ASCC//MessageValidation	NA	NA	100	200	ALERTING	11	ERROR	APP-DEV



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
for more into about SL and RTView

Don't miss SL Booth on Exhibit Floor !