



Informix.  
software

# TimeSeries and Real Time Loader


# Key Points

- Describe two complex DataBlades.
  - TimeSeries
  - Real Time Loader
- Why time series data can be a problem.
- How the TimeSeries Datablade solves these problems.
- Applications.
- Why streaming data can be a problem.
- How the RealTime Loader solves these problems.



# Timeseries

- TimeSeries:
  - What is a time series ?
    - Regular
    - Irregular
  - Where they are used ?
  - What they are used for ?
  - How they are used ?



# Regular Time Series Data

- Timestamps have a regularly repeating pattern of intervals daily, hourly, etc...
- Might be breaks in the pattern:
  - A work week is data captured for 5 days in a row then nothing for 2 days.
- Only one piece of data per interval:
  - If an interval has not been inserted into then it has the value of NULL
  - Intervals not inserted into occupy some space
    - Heuristic: intervals not inserted into at the end of series do not take up space.
- Can be thought of as an array.
- Optimized to return data at offset, not timestamp.

# Irregular Time Series Data

- Data in an irregular time series does not have a regularly repeating pattern of intervals:
  - Any interval may have zero or more pieces of data.
- Missing data takes no space on disk:
  - There really is no concept of missing data.
- Only efficient way to access data is by timestamp:
  - You can retrieve the Nth piece of data, but the code does a linear search.
- Data can be stair stepped:
  - Value persists until next value arrives – for example stock prices.
- Data can be discrete points:
  - Value is valid only at the given time – for example heart beats.

# The Timeseries "Problem"

- Timeseries are usually stored as "tall – thin" tables with a very large number of rows.
- Need fast access to an ordered set of rows.
- May need one index to enforce uniqueness and another for index only read, more space used for index than data.
- Can be difficult to write SQL to work with the data.



# Timeseries in Traditional RDBMS

- Irregular time series of stock "ticks".
- Typical volumes are 50M per day and increasing.
- Could access by "Stock" and "Timestamp" but often need an index on all columns for index only read.

| Stock | Timestamp                 | Price  | Volume |
|-------|---------------------------|--------|--------|
| ABC   | 2006-01-01 09:30:00.00000 | 12.34  | 100    |
| XYZ   | 2006-01-01 09:30:00.01000 | 9.34   | 1000   |
| ABC   | 2006-01-01 09:30:00.19000 | 12.44  | 10     |
| KKK   | 2006-01-01 09:30:01.05000 | 194.00 | 10000  |
| ABC   | 2006-01-01 09:30:01.70000 | 12.54  | 250    |

|                            |       |
|----------------------------|-------|
| Data                       | 35 Mb |
| Index on stock + timestamp | 28 Mb |
| Index on all columns       | 43 Mb |

# Timeseries in Traditional RDBMS

- Can be difficult to write SQL to process tick data:
  - Applying stock splits.
  - Calculating Volume Weighted Average Price (VWAP)

```
select t.tstamp, t.price * s.fac, t.vol/ s.fact
from tick t, split s
where t.stock = s.stock
and t.tstamp::datetime year to day = s.tstamp
and t.stock = "ABC"
and t.tstamp::date between
    ('2006-01-01' and '2006-01-05')
```

```
select stock, ststamp::datetime year to hour,
        sum(price * vol) / sum(vol)
from tick
where stock = "ABC"
and tstamp::date between
    ('2006-01-01' and '2006-01-05')
group by 1, 2
```

# Datablade Solution (1 of 7)

- Store timeseries elements as an ordered set of elements:
  - Uses less space because the "key" is factored out and the time can be represented in 4 bytes (regular).
  - Access is as fast (or faster) than index only read but no additional index.
  - SQL can be made much simpler.



## Datablade Solution (2 of 7)

- Store timeseries data as:
  - Start time.
  - Calendar information.
  - Ordered set of elements.


| Stock | Ticks   |
|-------|---|
| ABC   | 2006-01-01, daily, {(12.34,1234567), (12.56,950000), (12.34,5555567),..}  |
| KKK   | 2006-05-05, daily, {(199.08,678900), (198.55,345000), (198.12,850250),..} |
| XYZ   | 2003-09-01, daily, {(9.34,890567), (9.56,989000), (9.40,1000780),..}      |

## Datablade Solution (3 of 7)

- How a timeseries is stored:
  - Start time and calendar name.
  - Index (btree) of the first element in each page.
  - Each page contains all the elements between 2 timestamps.

|                            |       |
|----------------------------|-------|
| Timeseries                 | 30 Mb |
|                            |       |
| Data                       | 35 Mb |
| Index on stock + timestamp | 28 Mb |
| Index on all columns       | 43 Mb |

- Traditional methods use 2 to 3 times more space or even more depending on the number of NULL values.



## Datablade Solution (4 of 7)

- Performance:
  - Slower for singleton selects.
  - Faster accessing sets of data.
  - Much faster combining timeseries.

|                 | Singleton | 1,000 consec. | Aggregate 80,000 | 1,000 consec.<br>f(t1, t2) |
|-----------------|-----------|---------------|------------------|----------------------------|
| Timeseries      | 1         | 1             | 1                | 1                          |
| Traditional (1) | .5        | 2             | 3                | 5                          |
| Traditional (2) | 2         | 400           | 1000             | 650                        |

- Traditional(2) = index on key columns.
- Traditional(1) = index on key and all columns.

## Datablade Solution (5 of 7)

- Much simpler SQL, consider the "stock split"/calibration problem:

```
select t.tstamp, t.price * s.fac, t.vol/ s.fact
  from tick t, split s
  where t.stock = s.stock
     and t.tstamp::datetime year to day = s.tstamp
     and t.stock = "ABC"
     and t.tstamp::date between ('2006-01-01' and '2006-01-05')
```

```
select func(ticks, splits, '2006-01-01', '2006-01-05')
  from tick
  where stock = "ABC"
```



# Datablade Solution (6 of 7)

- Much simpler SQL, the VWAP problem:

```
select stock, tstamp::datetime year to hour, sum(price * vol) / sum(vol)
  from tick
 where stock = "ABC"
 and tstamp::date between ('2006-01-01' and '2006-01-05')
 group by 1, 2
```

```
select stock, vwap(ticks, '2006-01-01', '2006-01-05', '01:00:00')
  from tick
 where stock = "ABC"
```



# Datablade Solution (7 of 7)

- Other Timeseries Datablade features:
  - Calendar support.
  - VTI to make time series look like traditional tables.
  - C API.
  - Java Class Library.
- Easy to combine with own or 3<sup>rd</sup> party code.



# Datablade Practicalities (1 of 7)

- In order of operations:
- How do I add a Datablade to a database ?

On the command line:

```
blademgr  
register TimeSeries.4.10.UC7 in stocks
```


or via dbaccess:

```
execute function  
sysbldprepare('timeseries.4.01.UC7','create');
```



## Datablade Practicalities (2 of 7)

- Calendar Patterns – stored in the `CalendarPatterns` table
  - All time series are associated with these.
  - Used to bound the times at which timeseries data is recorded:
    - Data recorded during 'on' times.
    - Data not recorded during 'off' times.
    - Data recorded for the interval specified within the bounds.
    - Within the calendar definition.
- A Monday to Friday 5 day work week can be represented as 1 day off (Sunday), 5 days on, 1 day off (Saturday):
  - **Insert into `CalendarPatterns` values ('Five\_Day\_Work\_Week', '{1 off, 5 on, 1 off}', day');**
  - Note the brackets '{ }' above.
  - The trailing day is an interval for a single unit of time.



## Datablade Practicalities (3 of 7)

- How do I create a data type that defines each element in a time series ?
  - Must first create a row type whose first column timestamp must be a **datetime year to fraction(5)** value:

```
create row type stockVals (  
    tstamp      datetime year to fraction(5),  
    price       decimal(16,2),  
    high        decimal(16,2),  
    low         decimal(16,2),  
    median      decimal(16,2),  
    volume      integer,  
    trades      integer,  
    ....  
);
```

- Each timestamp must be unique.
- Row types cannot contain types:
  - blob, clob, text, serial

## Datablade Practicalities (4 of 7)

- Once a calendar pattern is established, create the calendar to use it in the table CalendarTable:

```
insert into CalendarTable ( c_name, c_calendar)  
values ('Trading Week','startdate(2010-01-01 00:00:00.00000),  
pattstart(2010-01-07 00:00:00.00000),  
pattname(Five_Day_Work_Week)');
```

- Regular timeseries calendars are used to guide the mapping between the timestamp and the offset
  - Represented by Calendar Data Type, consisting of:
    - Start date of the calendar.
    - Calendar repeating pattern during the duration of the data type for the valid and invalid recording of the data.
    - Interval – calibration of the calendar pattern.
    - Starting date of the calendar pattern
      - Must be  $\geq$  Calendar start date and
      - Less than one calendar pattern length after calendar start.

## Datablade Practicalities (5 of 7)

- Containers are required data structures that hold data for one or more time series.
- You cannot mix data for regular and irregular time series in the same container.
- How do I create a container to store a time series ?
  - **execute procedure**  
**TSContainerCreate("ts\_ctrname","dbspace1","stockVals",40000,40000);**
- A container is made up of index pages and data pages.
- Containers allow time series data to be spread onto many disk partitions.



## Datablade Practicalities (6 of 7)

- How do I create a table that includes a time series ?

```
– create table dailyStocks (  
    stockId          integer,  
    abbrev           char(4),  
    stockData        timeseries(stockVals));
```

- How do I create a time series record ?

```
– insert into dailyStocks values (901, 'IBM',  
    TSCreate('Trading Week', '2010-01-03 00:00:00.00000',  
    20, 0, 0, 'stockCon'));
```



## Datablade Practicalities (7 of 7)

- How do I add data to a time series ?

- update dailyStocks

```
set stockData = PutElem(stockData,  
row(NULL::datetime year to fraction(5), 3.3, 4.4,  
2.2, 3.0, 123456, 789, NULL)::stockData)  
where stock = 'XYZ';
```

- How do I query a row with a time series ?

- select clip(stockData, current - interval(1) month  
to month, NULL) from stockData;

- Get all of IBM since the beginning of the 2001

- select clip(series, '2001-01-01 00:00:00.00000',  
Current) from S\_P\_500 where name = 'IBM';



# Building Applications with Time Series

- Several interfaces are available:
  - SQL
  - VTI
  - SPL
  - Java
  - C-API
- It's a toolkit approach!
- Allow people to build their analytics.



# Time Series SQL Interface

- Time series data is usually accessed through user defined routines (UDR's) from SQL, some of these are:
  - **Clip()** - clip a range of a time series and return it.
  - **LastElem()**, **FirstElem()** - return the last (first) element in the time series.
  - **Apply()** - apply a predicate and expression to a range of elements in a time series.
  - **AggregateBy()** - change the interval of a time series using a aggregate function.
  - **SetContainerName()** - move a time series from one container to another.
  - **BulkLoad()** - load data into a time series from a file.
- See Appendix A for a list of all of SQL Interface UDR's.



# Time Series VTI Interface

- Makes time series data look like standard relational data:
  - Useful for programs that can't handle objects.
  - Useful when application connects with ODBC.
- There is a small (10%) penalty for using VTI.
- Restrictions:
  - A VTI table can only reference one time series column from the base table.
  - No secondary indices are allowed.
- SQL to create a VTI table:
  - *execute procedure tscreatevirtualtab('t\_vti', 't');*

# Stored Procedure (SPL) Example


```
-- count non-null elements in a time series
create function spl_test(arg lvarchar) returning integer
  define var daybar;
  define cnt integer;
  let cnt = 0;
  foreach execute function
    transpose((select series from S_P_500 where name = arg))
  into var
    let cnt = cnt + 1;
  end foreach
  return cnt;
end function;
```

# Time Series C-API Interface

- Client and server versions of the API.
- Treats a time series like a table (sort of).
- Functions to:
  - Open and close a time series.
  - Scan a time series between 2 timestamps.
  - Create a time series.
  - Retrieve, insert, delete, update.
- Plus another 70 functions defined.
- See Appendix C for a list of these.


# Timeseries Application

- Existing applications:
  - Finance (Capital Markets).
  - Some in Manufacturing and Environmental.
  - Smart Electrical Grid Usage
- Possible applications:
  - Telco.
    - Real Time Cell Phone Tower Data
    - Sales/sales analysis.
  - Energy
    - Oil Field Geology data.
  - Water/Sewer System Usage
  - Satellite Telemetry
  - Assembly Line Q/A & Analysis
  - Your application ?

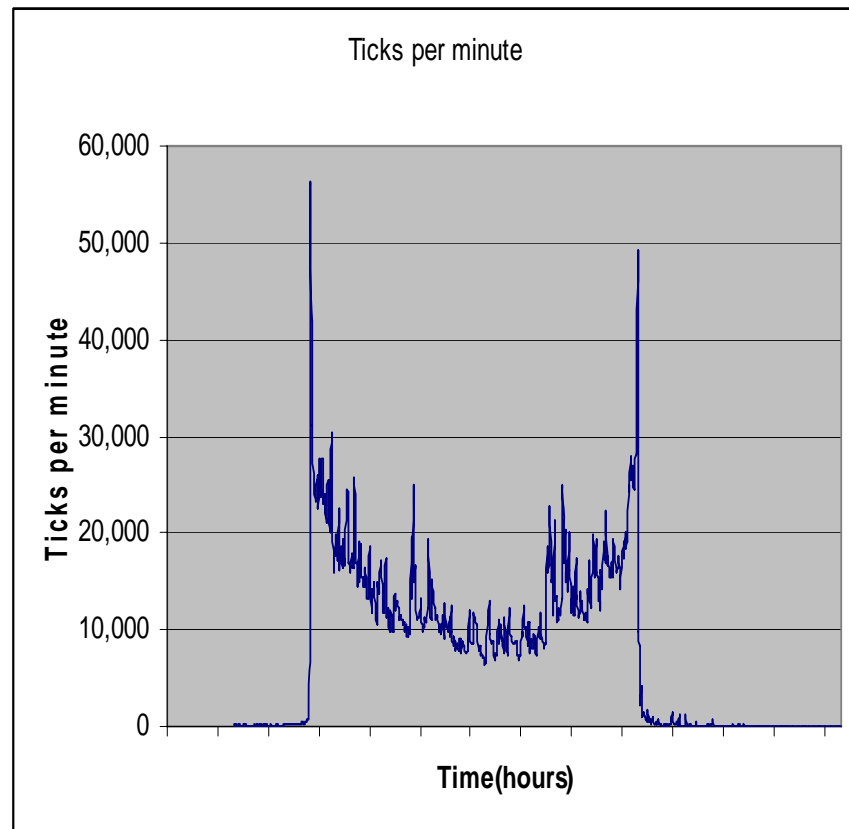


# RealTime Loader (1)

- Next problem after time series "problem" solved.
- Large number of ticks generated at market open/close and in reaction to events.
- Money is made at these times.
- Need to make the data available as soon as possible.



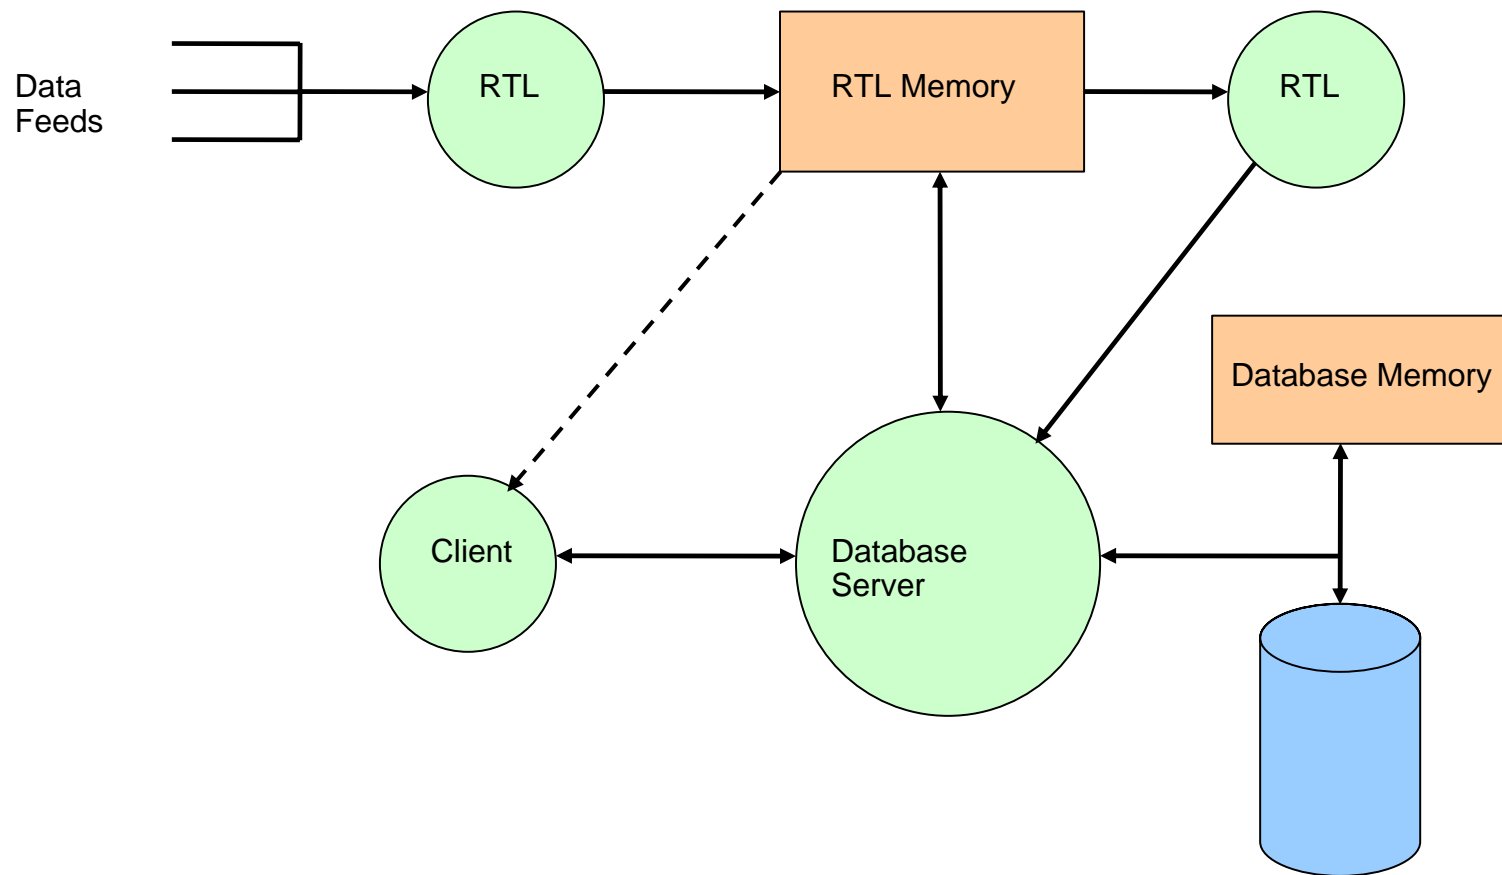
Edison Stock Ticker - 1892



## RealTime Loader (2)

- Traditionally data can be loaded very quickly by loading it in batches. This causes problems:
  - It takes time for the data to become available.
  - The table can be locked while the data is loaded.
- The RealTime Loader solves this problem by holding the data in memory and making it available to the SQL layer as if it was already in the database:
  - Allows real-time analysis of data as it come thru the feeds
    - Uncommitted and committed data reads are possible.

# RealTime Loader Architecture



# RealTime Loader Performance (1 of 3)

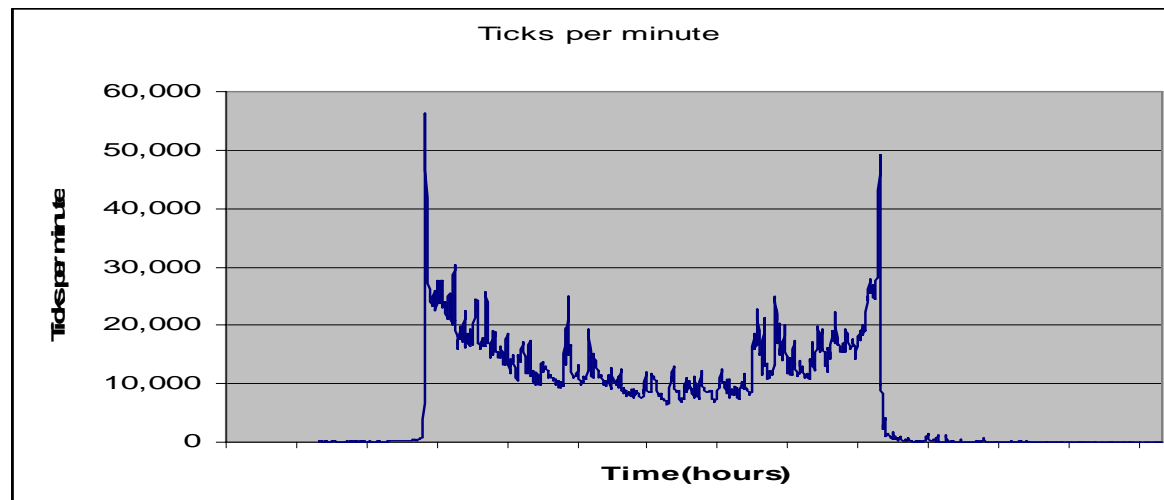
- Load performance of traditional systems depends on the system but rates of 1,000s of ticks per second are reasonable but there is a high latency as the data is batched and loaded.
- With the RTL the data is made accessible within a few ms. There are 2 load rates to consider:
  - How fast the data can be added to shared memory.
  - How fast the data can be loaded into the database.

## RealTime Loader Performance (2 of 3)

- How fast can the data be added to shared memory ?
  - Rates of 10,000s per cpu sec are common.
- How fast can the data can be loaded into the database ?
  - This is roughly the same as the traditional RDBS solution.
- The system can run at the "top" rate until the shared memory allocated to RTL is full.

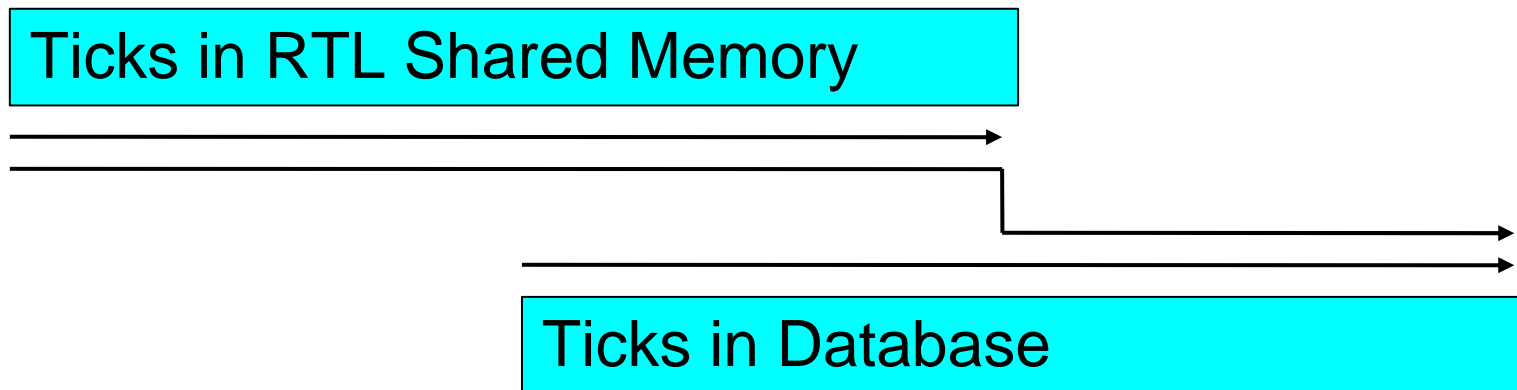
## RealTime Loader Performance (3 of 3)

- The peaks don't last very long, the system can be sized for something just above the average rate.
- Traditional systems have to be sized to handle the peak rate or the user has to accept a high latency during the crucial market open/close periods.



# RTL Data Access

- Access can be through the usual SQL layer.
- Even faster access through an API that gets the data directly from the shared memory segment.
- The SQL layer can see all the ticks, or just those in RTL, or just those on the database.



# RealTime Loader Analytics

- The RealTime Loader can do some simple aggregation and statistics on the incoming data making this available to SQL or through the API:
  - VWAP - Volume Weighted Moving Average Price
    - = total amount spent on a security / total number of shares traded.
    - This needs all of the share purchase history stored unbounded by time.
  - EWAP – Exponentially Weighted Moving Average Price
    - A share price moving average where previous moving average data is considered but dampened by the effect of the new, more recent moving average.
    - Only needs 3 pieces of data:
      - the old and new MAP data.
      - plus the damping factor.
  - max, min etc.

# RealTime Loader Applications

- Stock ticker trades data analysis.
- Gaming
- Telco:
  - Cell Phone Tower Call Analysis Apps.
  - Cell Phone Account Apps.
- 24 x 7 Satellite Telemetry Data Applications:
  - Millions of records per minute .....
- Geological Analysis Applications:
  - Oil Exploration Data.
- Land Resource Analysis Applications.
- Water/Sewer Flow Data

# Conclusion

- The TimeSeries DataBlade allows time series data to be:
  - Efficiently stored.
  - Quickly accessed.
  - Conveniently queried.
- The RealTimeLoader allows time series data to be:
  - Quickly stored.
  - Accessed by applications with the minimum of delay.

# Logo





Informix.  
software

# Appendix A - Time Series Built In Routines

## Appendix – Time Series Built In Routines (1 of 9)

| Task Type                                  | Description   | Routine Name            |
|--|---|-------------------------|
| Get info from a timeseries.                | Get the origin.   | <b>GetOrigin</b>        |
|  | Get the interval.   | <b>GetInterval</b>      |
|  | Get the calendar.   | <b>GetCalendar</b>      |
|  | Get the calendar name.  | <b>GetCalendarName</b>  |
|  | Get the container name.   | <b>GetContainerName</b> |
|  | Get the user defined metadata.                                  | <b>GetMetaData</b>      |
|  | Get the metadata type.  | <b>GetMetaTypeName</b>  |
|  | Determine whether a time series is regular.                     | <b>IsRegular</b>        |
|  | Get the instance Id if the timeseries is stored in a container. | <b>InstanceId</b>       |
| Convert between a timestamp and an offset. | Return the offset, given the timestamp.                         | <b>GetIndex</b>         |
|  | Return the timestamp, given the offset.                         | <b>GetStamp</b>         |
| Count the number of elements.              | Return the number of elements.                                  | <b>GetNelems</b>        |
|  | Get the number of elements between two timestamps               | <b>ClipGetCount</b>     |

## Appendix - Time Series Built in Routines (2 of 9)

|                             |  |                         |
|-----------------------------|--|-------------------------|
| Select Individual elements. | Get the element associated with a timestamp. | <b>GetElem</b>          |
|                             | Get the element at or before a timestamp.    | <b>GetLastValid</b>     |
|                             | Get the element after a timestamp.           | <b>GetNextValid</b>     |
|                             | Get the element before a timestamp.          | <b>GetPreviousValid</b> |
|                             | Get the element at a specified position.     | <b>GetNthvalid</b>      |
|                             | Get the first element.                       | <b>GetFirstElem</b>     |
|                             | Get the last element.                        | <b>GetLastElem</b>      |
|                             | Get the last non null element.               | <b>GetLastNonNull</b>   |
|                             | Get the next non null element.               | <b>GetNextNonNull</b>   |
|                             |  |                         |

## Appendix - Time Series Built in Routines (3 of 9)

|                                       |  |                      |
|---------------------------------------|--|----------------------|
| Modify elements or a set of elements. | Add or update a single element.                                | <b>PutElem</b>       |
|                                       | Add or update a single element.                                | <b>PutElemNoDups</b> |
|                                       | Add or update a single element at a given offset.              | <b>PutNthElem</b>    |
|                                       | Add or update an entire set.                                   | <b>PutSet</b>        |
|                                       | Delete an element at a given timepoint.                        | <b>DelElem</b>       |
|                                       | Delete all elements in a specified time range.                 | <b>DelClip</b>       |
|                                       | Delete all elements in a specified time range.                 | <b>DelTrim</b>       |
|                                       | Insert an element.   | <b>InsElem</b>       |
|                                       | Insert a set.  | <b>InsSet</b>        |
|                                       | Update an element.   | <b>UpdElem</b>       |
|                                       | Update a set.  | <b>UpdSet</b>        |
|                                       | Put every element of one time series into another time series. | <b>PutTimeSeries</b> |
| Modify metadata.                      | Update user-defined metadata.                                  | <b>UpdMetadata</b>   |

## Appendix - Time Series Built in Routines (4 of 9)

|   |  |                      |
|---|--|----------------------|
| Make Elements visible or invisible to a scan. | Make an element invisible.   | <b>HideElem</b>      |
|   | Make a range of elements invisible.  | <b>HideRange</b>     |
|   | Make an element invisible.   | <b>RevealElem</b>    |
|   | Make a range of elements visible.  | <b>RevealRange</b>   |
| Check for null or hidden elements.            | Determine if an element is hidden.   | <b>ElemIsHidden</b>  |
|   | Determine if an element is null.   | <b>ElemIsNull</b>    |
| Extract and use part of a time series.        | Extract a period between two timestamps to a set of values and run an expression or function on every entry. | <b>Apply</b>         |
|   | Extract data between two timepoints.   | <b>Clip</b>          |
|   | Clip a certain number of elements.   | <b>ClipCount</b>     |
|   | Extract a period that includes a given time.   | <b>WithinC</b>       |
|   | Extract a period starting or ending at a given time.   | <b>WithinR</b>       |
| Apply a new calendar to a time series.        | Apply a calendar.  | <b>ApplyCalendar</b> |

## Appendix - Time Series Built in Routines (5 of 9)

|  |  |                        |
|--|--|------------------------|
| Create and load a time series.                 | Load data from a client file.  | <b>BulkLoad</b>        |
|  | Create a regular empty time, a regular populated time series, or a regular time series with metadata.          | <b>TsCreate</b>        |
|  | Create an irregular empty time, an irregular populated time series, or an irregular time series with metadata. | <b>TsCreateIrr</b>     |
| Find the intersection or union of time series. | Build the intersection of multiple time series and optionally clip the result.                                 | <b>Intersect</b>       |
|  | Build the union of multiple time series and optionally clip the result.  | <b>Union</b>           |
| Iterator functions.                            | Convert time series data to tabular form.  | <b>Transpose</b>       |
| Aggregate Functions.                           | Return a list (collection of rows) containing all elements in a time series.                                   | <b>TSSetToList</b>     |
|  | Return a list of columns in a timeseries.  | <b>TSColNameToList</b> |
|  | Return a list of columns in a timeseries.  | <b>TSColNumToList</b>  |
|  | Return a list containing the columns of a time series plus non time series columns.                            | <b>TSRowToList</b>     |
|  | Return a list containing the columns of a time series plus non time series columns.                            | <b>TsRowNameToList</b> |
|  | Return a list containing the columns of a time series plus non time series columns.                            | <b>TsRowNumToList</b>  |

## Appendix - Time Series Built in Routines (6 of 9)

|  |  |                      |
|--|--|----------------------|
| Used within the apply function to perform statistical calculations on a time series. | Perform a sum over a time series type.   | <b>Sum</b>           |
|  | Sum <b>SMALLFLOAT</b> or <b>DOUBLE PRECISION</b> values.                                 | <b>TSAddPrevious</b> |
|  | Compute the decay function.  | <b>TSDecay</b>       |
|  | Compute a running average over a specified number of values.                             | <b>TSRunningArg</b>  |
|  | Compute a running correlation between two time series over a specified number of values. | <b>TsRunningCor</b>  |
|  | Compute a running median over a specified number of values.                              | <b>TsRunningMed</b>  |
|  | Compute a running sum over a specified number of values.                                 | <b>TSRunningSum</b>  |
|  | Compute a running variance over a specified number of values.                            | <b>TSRunningVar</b>  |
|  | Compare <b>SMALLFLOAT</b> or <b>DOUBLE PRECISION</b> values.                             | <b>TSCmp</b>         |
|  | Return a previously saved value.   | <b>TSPrevious</b>    |

## Appendix - Time Series Built in Routines (7 of 9)

|  |   |                 |
|--|---|-----------------|
| Perform an arithmetic operation on one or two time series. | Add two time series together.   | <b>Plus</b>     |
|  | Subtract one time series from another.                                    | <b>Minus</b>    |
|  | Multiply one time series by another.                                      | <b>Times</b>    |
|  | Divide one time series by another.  | <b>Divide</b>   |
|  | Raise the first argument to the power of the second.                      | <b>Pow</b>      |
|  | Get the absolute value.   | <b>Abs</b>      |
|  | Exponentiate the time series.   | <b>Exp</b>      |
|  | Get the natural logarithm of a time series.                               | <b>Logn</b>     |
|  | Get the modulus or remainder of a division of one time series by another. | <b>Mod</b>      |
|  | Return the argument, is bound to the unary + operator.                    | <b>Negate</b>   |
|  | Round the time series to the nearest whole number.                        | <b>Positive</b> |
|  |   |                 |
|  |   |                 |

## Appendix - Time Series Built in Routines (8 of 9)

|   |  |                        |
|---|--|------------------------|
| Perform an arithmetic operation on one or two time series (cont'd). | Round the time series to the nearest whole number.   | <b>Round</b>           |
|   | Get the square root of the time series.  | <b>Sqrt</b>            |
|   | Get the cosine of the time series.   | <b>Cos</b>             |
|   | Get the sine of the time series.   | <b>Sin</b>             |
|   | Get the tangent of the time series.  | <b>Tan</b>             |
|   | Get the arc cosine of the time series.   | <b>Acos</b>            |
|   | Get the arc sine of the time series.   | <b>Asin</b>            |
|   | Get the arc tangent of the time series.  | <b>Atan</b>            |
|   | Get the arc tangent for two time series.   | <b>Atan2</b>           |
|   | Apply a binary function to a pair of time series, or to a time series and a compatible row type or number. | <b>ApplyBinaryTsOp</b> |
|   | Apply a unary function to a time series.   | <b>ApplyUnaryTsOp</b>  |
|   | Apply another function to a set of time series.  | <b>ApplyOpTsSet</b>    |
| Reset the origin.   | Reset the origin.  | <b>SetOrigin</b>       |

## Appendix - Time Series Built in Routines (9 of 9)

|                                    |  |                           |
|------------------------------------|--|---------------------------|
| Aggregate values in a time series. | Aggregate values in a time series.                                       | <b>AggregateBy</b>        |
|                                    | Aggregate values in a time series over a specified time range.           | <b>AggregateRange</b>     |
| Create a time series that lags.    | Create a time series that lags the source time series by a given offset. | <b>Lag (regular only)</b> |
| Manage containers.                 | Create a container.  | <b>TsContainerCreate</b>  |
|                                    | Destroy a container.   | <b>TsContainerDestroy</b> |
|                                    | Set the container name.  | <b>SetContainerName</b>   |
|                                    |  |                           |
|                                    |  |                           |
|                                    |  |                           |

- Following routines are used only with regular time series:
  - Lag
  - PutNthElem
  - TSCreate
- The **TSCreateIrr** function is only used with regular time series.



Informix.  
software

# Appendix B – Calendar Pattern Routines

## Appendix - Calendar Pattern Routines

|  |   |                         |
|--|---|-------------------------|
| Return the intersection of two calendar patterns.                                  | Calendar pattern that has every interval on that was on in both calendar patterns; the rest are off. If the given patterns do not have the same interval unit, the pattern with the larger interval unit is expanded to match the other.                      | <b>AndOp</b>            |
| Obtain the start date of the pattern for a given calendar.                         | Returns the start dates of the calendar patterns for each calendar.   | <b>CalPattStartDate</b> |
| Convert a calendar from a smaller calendar pattern into a larger calendar pattern. | Collapses the given calendar pattern into destination units, which must have a larger interval unit than that of the given calendar pattern, i.e. Converting hours to days.   | <b>Collapse</b>         |
| Convert a calendar from a larger calendar pattern into a smaller calendar pattern. | Converts the given calendar pattern into the destination units, which must have a smaller interval unit than that of the given calendar pattern.  | <b>Expand</b>           |
| Return an inverted calendar intervals pattern.                                     | Turns all on intervals off and all off intervals on in the given calendar pattern.  | <b>NotOp</b>            |
| Returns the union of the two calendar patterns.                                    | This function returns a calendar pattern that has every interval on that was on in either of the calendar patterns; the rest are off. If the two patterns have different sizes of interval units, the resultant pattern has the smaller of the two intervals. | <b>OrOp</b>             |

## Appendix – Calendar Routines

|  |  |                     |
|--|--|---------------------|
| Return the intersection of two calendars.  | Returns a calendar that has every interval on that was on in both calendars; the rest are off. The resultant calendar takes the later of the two start dates and the later of the two pattern start dates. If the two calendars have different size interval units, the resultant calendar has the smaller of the two intervals.         | <b>AndOp</b>        |
| Return the number of valid intervals in the named calendar between the two timestamps. |  | <b>CalIndex</b>     |
| Return a set of valid timestamps within a range.                                       | There are two syntaxes: the first syntax specifies the range as between two given timestamps. The second syntax specifies the number of valid timestamps to return after a given timestamp.  | <b>CalRange</b>     |
| Return the timestamp at a given number of calendar intervals after a given timestamp.  |  | <b>CalStamp</b>     |
| Returns a <b>DATETIME</b> value containing the start date of a given calendar name.    |  | <b>CalStartDate</b> |
| Return a calendar that is the result of two others combined with the OR operator       | This function returns a calendar that has every interval on that was on in either calendar; the rest are off. The resultant calendar takes the earlier of the two start dates and the two pattern start dates. If the two calendars have different sizes of interval units, the resultant calendar has the smaller of the two intervals. | <b>OrOp</b>         |



Informix.  
software

# Appendix C - Time Series API Information

# Appendix - Time Series API's

## Application Programmer Interfaces for a Time Series:

- Two types -
  - Client version – *tsfeapi*
    - Sample Location: `$INFORMIXDIR/extend/TimeSeries.4.07.UC7/lib/tsfeapi.a` (Unix/linux)
    - Similar directory for Windows.
    - Operates on a private copy of timeseries datum.
    - Changed data must be written back into the database.
      - Function *ts\_copy* is used prior to the execution of a prepared statement.
  - Server version – *tsbeapi*
    - Sample Location: `$INFORMIXDIR/extend/TimeSeries.4.07.UC7/lib/tsbeapi.a` (Unix/Linux)
    - Similar directory for Windows.
    - Operates directly on timeseries datum.
      - Direct execution of prepared/unprepared statements.
- Header file – *tseries.h* -
  - Must be included in header whenever using *bsfeapi* or *tsfeapi*

## Appendix - Time Series API – Data Structures

- `ts_timeseries` – the header for a time series data.
- `ts_tscan` – allows you to look at no more than two time series elements at a time.
- `ts_tsdesc` – A structure containing a time series and the data structures it works with.
- `ts_tselem` – A pointer to one element of a time series.

# Appendix – Time Series API Routines

| Task type                      | Description   | API Routine                      | SQL Routine                  |
|--------------------------------|---|----------------------------------|------------------------------|
| Open and close a time series.  | Open a time series.   | <b>ts_open()</b>                 |                              |
|                                | Close a time series.  | <b>ts_close()</b>                |                              |
|                                | Return a pointer to the time series associated with the given time series descriptor.               | <b>ts_get_ts()</b>               |                              |
| Create and copy a time series. | Create a time series.   | <b>ts_create()</b>               | <b>TSCreate, TSCreateIrr</b> |
|                                | Create a time series with metadata.   | <b>ts_create_with_metadata()</b> | <b>TSCreate, TSCreateIrr</b> |
|                                | Copy a time series.   | <b>ts_copy()</b>                 |                              |
|                                | Free all memory associated with a time series created with <b>ts_copy()</b> or <b>ts_create()</b> . | <b>ts_free()</b>                 |                              |
|                                | Copy all elements of one time series into another.  | <b>ts_put_ts()</b>               | <b>PutTimeSeries</b>         |

## Appendix – Time Series API Routines (cont'd)

| Task type                                     | Description   | API Routine                   | SQL Routine       |
|---|---|-------------------------------|-------------------|
| Scan a time series.                           | Scan a time series.   | <b>ts_begin_scan()</b>        |                   |
|   | Retrieve the next element from a scan.                                    | <b>ts_next()</b>              |                   |
|   | End a scan.   | <b>ts_end_scan()</b>          |                   |
|   | Find the timestamp of the last element returned by the <b>ts_next()</b> . | <b>ts_current_timestamp()</b> |                   |
|   | Return the offset for the last element returned by <b>ts_next()</b> .     | <b>ts_current_offset()</b>    |                   |
| Make elements visible or invisible to a scan. | Make an element invisible.  | <b>ts_hide_elem()</b>         | <b>HideElem</b>   |
|   | Make an element visible.  | <b>ts_reveal_elem()</b>       | <b>RevealElem</b> |
|   |   |                               |                   |

## Appendix – Time Series API Routines (cont'd)

| Task type                                      | Description   | API Routine                | SQL Routine             |
|--|---|----------------------------|-------------------------|
| Select individual elements from a time series. | Get the element closest to a given timestamp.         | <b>ts_closest_elem()</b>   |                         |
|  | Get the element associated with a given timestamp.    | <b>ts_elem()</b>           | <b>GetElem</b>          |
|  | Get the element at a specified position.              | <b>ts_nth_elem()</b>       | <b>GetNthElem</b>       |
|  | Get the first element.                                | <b>ts_first_elem()</b>     | <b>GetFirstElem</b>     |
|  | Get the last element.                                 | <b>ts_last_elem()</b>      | <b>GetLastElem</b>      |
|  | Find the next element after a given timestamp.        | <b>ts_next_valid()</b>     | <b>GetNextValid</b>     |
|  | Find the last element before a given timestamp.       | <b>ts_previous_valid()</b> | <b>GetPreviousValid</b> |
|  | Find the last element at or before a given timestamp. | <b>ts_last_valid()</b>     |                         |

## Appendix – Time Series API Routines (cont'd)

| Task type                                 | Description   | API Routine  | SQL Routine                            |
|---|---|--|--|
| Update a time series.                     | Insert an element.                                  | <b>ts_ins_elem()</b>                                 | <b>InsElem</b>                         |
|   | Update an element.                                  | <b>ts_upd_elem()</b>                                 | <b>UpdElem</b>                         |
|   | Delete an element.                                  | <b>ts_del_elem()</b>                                 | <b>DelElem</b>                         |
|   | Put an element in a place specified by a timestamp. | <b>ts_put_elem()</b><br><b>ts_put_elem_no_dups()</b> | <b>PutElem</b><br><b>PutElemNoDups</b> |
|   | Append an element.                                  | <b>ts_put_last_elem()</b><br>(regular only)          |  |
|   | Put an element in a place specified by an offset.   | <b>ts_put_nth_elem()</b><br>(regular only)           |  |
| Modify metadata.                          | Update metadata.                                    | <b>ts_update_metadata()</b>                          | <b>UpdMetaData</b>                     |
| Convert between an index and a timestamp. | Convert timestamp to index.                         | <b>ts_index()</b>                                    | <b>GetIndex</b>                        |
|   | Put an element in a place specified by an offset.   | <b>ts_time()</b>                                     | <b>GetStamp</b>                        |

## Appendix – Time Series API Routines (cont'd)

| Task type  | Description  | API Routine   | SQL Routine |
|--|--|---|-------------|
| Transform an element.                            | Create an element from an array of values and nulls.   | <code>ts_make_elem()</code><br><code>ts_make_elem_rowdesc()</code><br><code>ts_make_elem_no_dups()</code> |             |
|  | Convert an <code>MI_ROW</code> value to an element.  | <code>ts_row_to_elem()</code>   |             |
|  | Convert an element to a <code>MI_ROW</code> value.   | <code>ts_elem_to_row()</code>   |             |
| Extract column data from an element.             | Free memory from a time series element created by <code>ts_make_elem()</code> or <code>ts_row_to_elem()</code> | <code>ts_free_elem()</code>   |             |
|  | Get a column from an element by name.  | <code>ts_get_col_by_name()</code>   |             |
|  | Get a column from an element by number.  | <code>ts_get_col_by_number()</code>   |             |
|  | Pull columns from an element into values and nulls arrays.   | <code>ts_get_all_cols()</code>  |             |
| Create and perform calculations with timestamps. | Compare two timestamps.  | <code>ts_datetime_cmp()</code>  |             |
|  | Get fields from a timestamp.   | <code>ts_get_stamp_fields()</code>  |             |
|  | Create a timestamp.  | <code>ts_make_stamp()</code>  |             |

## Appendix – Time Series API Routines (cont'd)

| Task type   | Description   | API Routine                      | SQL Routine |
|---|---|----------------------------------|-------------|
| Create and perform calculations with timestamps (cont'd). | Calculate the number of intervals between two timestamps. | <b>ts_timestamp_difference()</b> |             |
|   | Subtract N intervals from a timestamp.                    | <b>ts_row_to_elem()</b>          |             |
|   | Add N intervals to a timestamp.                           | <b>ts_elem_to_row()</b>          |             |
| Get information about element data.                       | Find the number of a column.                              | <b>ts_col_id()</b>               |             |
|   | Return the number of columns contained in each element.   | <b>ts_get_col_cnt()</b>          |             |
|   | Get type information for a column specified by number.    | <b>ts_get_colinfo_number()</b>   |             |
|   | Get type information for a column specified by name.      | <b>ts_get_colinfo_name()</b>     |             |
|   | Determine if an element is hidden.                        | <b>ts_ELEM_HIDDEN</b>            |             |
|   | Determine if an element is NULL.                          | <b>ts_ELEM_NULL()</b>            |             |

## Appendix – Time Series API Routines (cont'd)

| Task type                           | Description  | API Routine                         | SQL Routine             |
|-------------------------------------|--|-------------------------------------|-------------------------|
| Get information about a timeseries. | Get the name of calendar associated with a timeseries. | <code>ts_get_calname()</code>       | <b>GetCalendar Name</b> |
|                                     | Return the number of elements in a timeseries.         | <code>ts_nelems()</code>            | <b>GNelems</b>          |
|                                     | Return the flags associated with the timeseries.       | <code>ts_get_flags()</code>         |                         |
|                                     | Get the name of the container.                         | <code>ts_get_containername()</code> | <b>GetContainerName</b> |
|                                     | Determine if the timeseries is in a container.         | <code>TS_IS_INCONTAINER()</code>    |                         |
|                                     | Get the origin of the timeseries.                      | <code>ts_get_origin()</code>        | <b>GetOrigin</b>        |
|                                     | Get the metadata associated with the timeseries.       | <code>ts_get_metadata()</code>      | <b>GetMetaData</b>      |
|                                     | Determine if the timeseries is irregular.              | <code>TS_IS_IRREGULAR()</code>      |                         |

## Appendix – Time Series API Routines (cont'd)

| Task type                         | Description  | API Routine                       | SQL Routine     |
|-----------------------------------|--|-----------------------------------|-----------------|
| Get information about a calendar. | Return the number of valid intervals between two timestamps.                 | <code>ts_cal_index()</code>       | <b>CalIndex</b> |
|                                   | Return all valid timepoints between two timestamps.                          | <code>ts_cal_range()</code>       | <b>CalRange</b> |
|                                   | Return a specified number of timestamps starting at a given timestamp.       | <code>ts_cal_range_Index()</code> |                 |
|                                   | Return the timestamp at a given number of intervals after a given timestamp. | <code>ts_cal_stamp()</code>       | <b>CalStamp</b> |

# Logo

