

Practical experience with moving applications to PostgreSQL

STREAMLINE
TECHNOLOGIES INC

Radovan Jablonovsky
Database

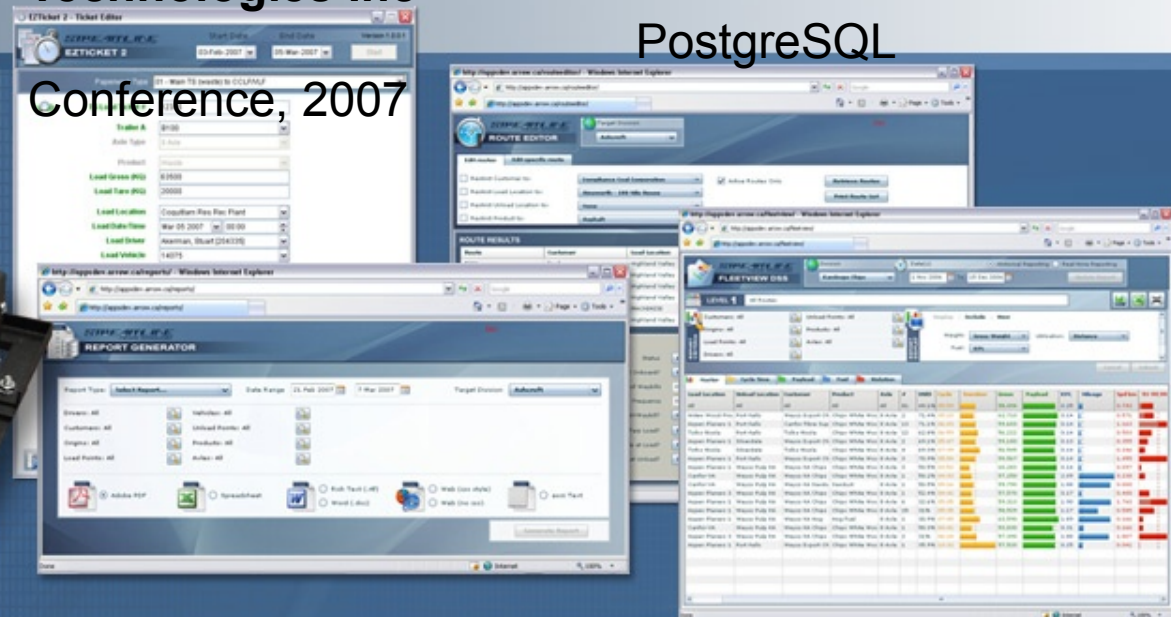
Architect/DBA

Streamline

Technologies Inc

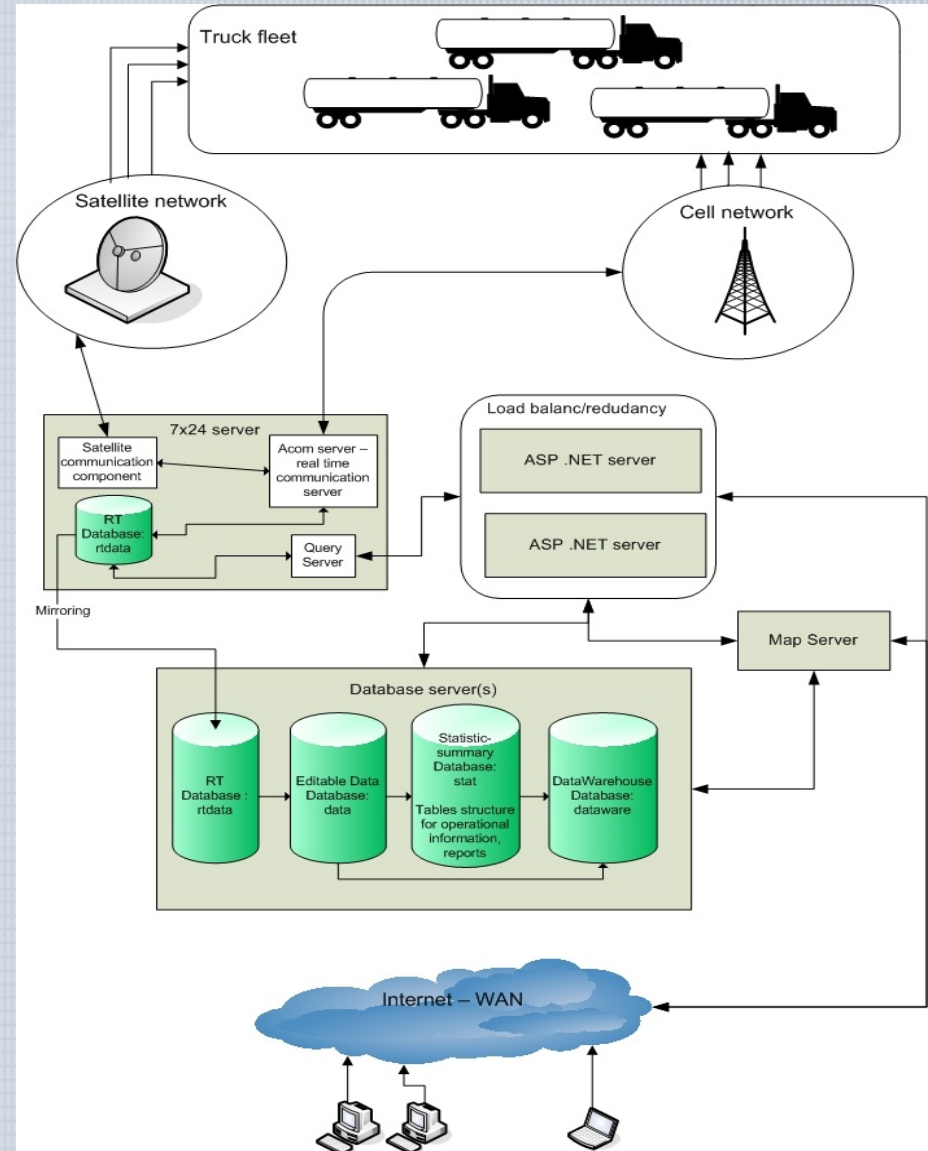
PostgreSQL

Conference, 2007



Overview

- System architecture
- Communication
- Databases
- Servers
- Client web app.
- Custom client/server app.




Applications

- Truck on board Orion data collection and communication system
- Servers applications
 - Map server
 - Data transformation server
 - Web application/services, Web servers: .NET
 - Web reporting server: java
- Client/server applications
- Batch data transformation applications

Orion data collection and communication system



Orion data collection and communication system



Servers applications - Map server

Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://apps.arrow.ca/realtimemapper/

Shaw Webmail (Messenger Exp... PostgreSQL: Documentation: M... http://apps.arrow.ca/ http://apps.a...ltimemapper/

Map / Configure

REAL-TIME TRACKING

Map Tools

Panel View

Map Redraw

Every 5 minutes

Saved Views

01 02 03 04 05 06 07 08 09 10

All Vehicles **Filtered Vehicles** **All Places** **Filtered Places**

| Vehicle # | Driver | Trailers | Communicated D: | Status | Load Place | Unload Place | Product | Customer |
|--------------|---------------------|------------|-----------------|--------------------|-------------------|--------------------|--------------|------------------|
| 14204 [2649] | Jones, Ernest (201 | [na], [na] | Wed May 16 01:5 | Log Off [LOGOFF] | Coquitlam Res Re | Cache Creek Land | Waste [19] | Wastech Services |
| 14208 [2663] | Fidler, Martin (205 | S15, [na] | Wed May 16 14:4 | Truck Stopped [ST | North Shore Trans | Vancouver Landfill | Waste [19] | Wastech Services |
| 14210 [2666] | [na], [na] (0) [0] | S13, [na] | Wed May 16 13:2 | Truck Stopped [ST | North Shore Trans | Vancouver Landfill | Waste [19] | Wastech Services |
| 14213 [2675] | [na], [na] (0) [0] | S12, [na] | Wed May 16 14:1 | Drive [DRIVELOAD | North Shore Trans | Vancouver Landfill | Waste [19] | Wastech Services |
| 14194 [2680] | Neufeld, Ian (209 | [na], [na] | Wed May 16 15:0 | Disconnect Trailer | CC Reload [7378] | S&R Reload Facilit | Chips [2] | Wastech Services |
| 14214 [2687] | Robinson, Robert | B160, B161 | Wed May 16 14:5 | Drive [DRIVELOAD | Chasm [2369] | S&R Reload Facilit | Sawdust [16] | Wastech Services |
| 14188 [2686] | Desser, Ernest (21 | S15, [na] | Wed May 16 14:4 | Truck Stopped [ST | North Shore Trans | Vancouver Landfill | Waste [19] | Wastech Services |

Find: variab Next Previous Highlight all Match case Phrase not found

Transferring data from webservice.arrow.ca...

Servers applications - Data transformation server

- Receive data from cell network (currently PCS)
- Decode data
- Recalculate and compare information from Trucks to stored on PostgreSQL
- Store data on PostgreSQL
- 24/7 service, windows application

- Servers applications - Web application/services

The screenshot shows a Mozilla Firefox browser window with the following details:

- Title Bar:** Mozilla Firefox
- Menu Bar:** File, Edit, View, History, Bookmarks, Tools, Help
- Address Bar:** http://apps.arrow.ca/
- Search Bar:** Google
- Tab Bar:** Shaw Webmail (...), PostgreSQL: Doc..., http://...ow.ca/ (active), http://a...mapper/, type of telus cell...
- Main Content Area:**
 - STREAMLINE ONBOARD TECHNOLOGIES**
 - APPLICATION LAUNCH PAGE**
 - Navigation Tools:** FleetView (calendar icon), Reports (document icon), RT Tracking (globe icon)
 - Administration Tools:** RouteEditor (compass icon), TripEditor (document icon), AssignUtility (document icon)
- Search Bar (Bottom):** Find: variab, Next, Previous, Highlight all, Match case, Phrase not found
- Status Bar:** Done

STREAMLINE ONBOARD TECHNOLOGIES FLEETVIEW DSS

Division: **Chilliwack Bulk** Date(s): **16 May 2007** Historical Reporting Real-time Reporting

Auto-update **10** min **Update Report**

LEVEL 1 All Routes

REPORT CRITERIA

Customers: All Unload Points: All

Origins: All Products: All

Load Points: All Axles: All

Drivers: All

REPORT DISPLAY

Display: **Include** | **View**

Weight: **Gross Weight**

Fuel: **KPL**

Utilization: **Distance**

Submit

Master Cycle Time Payload Fuel Violation

| Origin | Load Location | Unload | Customer | Product | Axle | # | UtilD | Cycle | Duration | Gross | Payload | KPL | Mileage | Spd km | 93-9' |
|-------------------------|--------------------------|--------|-----------|----------|--------|----|-------|-------|----------|--------|---------|------|---------|--------|-------|
| All | All | All | All | All | All | 30 | 37.8% | 00:25 | | 53,145 | | 1.06 | | 0.755 | |
| Port Kells | Coquitlam Res Rec Plant | Georg | Wastech S | Waste | 8 Axle | 5 | 6.4% | 00:18 | | 57,842 | | 1.26 | | 1.112 | |
| Vancouver Landfill | North Shore Transfer Stn | Vanco | Wastech S | Waste | 7 Axle | 5 | 48.9% | 00:08 | | 53,708 | | 0.75 | | 0.104 | |
| Waste-To-Energy Facilit | North Shore Transfer Stn | Vanco | Wastech S | Waste | 7 Axle | 3 | 42.3% | 00:17 | | 53,377 | | 0.78 | | 0.000 | |
| Vancouver Landfill | Coquitlam Res Rec Plant | Vanco | Wastech S | Waste | 6 Axle | 2 | 11.2% | 00:19 | | 44,675 | | 1.51 | | 0.000 | |
| Port Kells | Matsqui Transfer Stn | Georg | Wastech S | Waste | 8 Axle | 2 | 90.4% | 00:15 | | 56,085 | | 0.83 | | 2.126 | |
| IKO Unid Facility | Coquitlam Res Rec Plant | Georg | Wastech S | Waste | 8 Axle | 1 | 0% | 00:04 | | 58,760 | | 0.02 | | 0.000 | |
| Vancouver Landfill | Coquitlam Res Rec Plant | Waste | Wastech S | Waste | 6 Axle | 1 | 0% | 01:44 | | 44,440 | | 1.58 | | 0.000 | |
| [na] | Coquitlam Res Rec Plant | Vanco | Wastech S | Waste | 6 Axle | 1 | 100% | 00:11 | | 46,430 | | 0.74 | | 0.000 | |
| Waste-To-Energy Facilit | Coquitlam Res Rec Plant | Waste | Wastech S | Waste | 6 Axle | 1 | 47.9% | 01:02 | | 45,470 | | 1.20 | | 0.000 | |
| Waste-To-Energy Facilit | North Shore Transfer Stn | Waste | Wastech S | Waste | 6 Axle | 1 | 20.5% | 00:33 | | 45,000 | | 1.52 | | 0.000 | |
| [na] | IG Machine | [na] | Wastech S | Granules | 8 Axle | 1 | 71% | 01:25 | | 61,400 | | 1.51 | | 0.000 | |
| Port Kells | CC Reload | [na] | Wastech S | Chips | 8 Axle | 1 | 5.4% | 00:39 | | 52,030 | | 1.49 | | 0.827 | |
| [na] | North Shore Transfer Stn | Waste | Wastech S | Waste | 6 Axle | 1 | 62.9% | 01:09 | | 45,170 | | 1.50 | | 0.000 | |
| Port Kells | CC Reload | Georg | Wastech S | Chips | 8 Axle | 1 | 0% | 00:15 | | 58,370 | | 1.56 | | 6.818 | |
| Port Kells | North Shore Transfer Stn | Georg | Wastech S | Waste | 7 Axle | 1 | 21.3% | 00:09 | | 54,150 | | 0.57 | | 3.790 | |
| Port Kells | Surrey Transfer Stn | Georg | Wastech S | Waste | 8 Axle | 1 | 51.4% | 00:44 | | 51,850 | | 1.19 | | 0.098 | |
| Waste-To-Energy Facilit | North Shore Transfer Stn | Waste | Wastech S | Waste | 7 Axle | 1 | 16.8% | 00:21 | | 53,830 | | 1.11 | | 0.000 | |
| Port Kells | North Shore Transfer Stn | Georg | Wastech S | Waste | 8 Axle | 1 | 71.4% | 00:09 | | 58,040 | | 0.67 | | 0.786 | |

Servers applications - Web reporting server: java

Mozilla Firefox [Window Controls]

File Edit View History Bookmarks Tools Help

Address Bar: <http://apps.arrow.ca/reports/> [Google Search]

Shaw Webmail (Messen... | <http://apps.arrow.ca/> | <http://apps...timemapper/> | <http://apps...a/fleetview/> | <http://app.../reports/>

STREAMLINE
ONBOARD TECHNOLOGIES
REPORT GENERATOR

Report Type: **Trip Report by Trailer** | Date Range: 2 May 2007 - 16 May 2007 | Target Division: **Chilliwack Bulk**

Filters:

| | | | |
|------------------|--------------------------|--------------------|--------------------------|
| Drivers: All | <input type="checkbox"/> | Vehicles: All | <input type="checkbox"/> |
| Customers: All | <input type="checkbox"/> | Unload Points: All | <input type="checkbox"/> |
| Origins: All | <input type="checkbox"/> | Products: All | <input type="checkbox"/> |
| Load Points: All | <input type="checkbox"/> | Axles: All | <input type="checkbox"/> |

Output Formats:

| | | | | |
|--|-----------------------------------|--|---------------------------------------|----------------------------------|
| <input checked="" type="radio"/> Adobe PDF | <input type="radio"/> Spreadsheet | <input type="radio"/> Rich Text (.rtf) | <input type="radio"/> Web (css style) | <input type="radio"/> ascii Text |
| | | <input type="radio"/> Word (.doc) | <input type="radio"/> Web (no css) | |

Generate Report

Find: **variab** [Next] [Previous] [Highlight all] [Match case] [Phrase not found]

Transferring data from webservices.arrow.ca...

Client/server applications

EZTicket 2 - Ticket Editor Start Date: 08-Apr-2007 End Date: 08-May-2007 Version 1.0.0.11

STREAMLINE
EZTICKET 2

Paperwork Type: 09 - CCrk Reload (Chip) to Silverdale

Arrow Probill #

CCRL Ticket #

Siv Mill Ticket #

Trailer A

Axle Type

Product Chips

Load Gross (KG)

Load Tare (KG)

Load Location CC Reload

Load Date/Time May 08 2007 00:00

Unload Location Silverdale

Unload Date/Time May 08 2007 00:00

Unload Driver

Unload Vehicle

Unload Gross (KG)

Unload Tare (KG)

Truck Type Arrow

Haul Type Back Haul - Leg 2

Redirected No

Route

Search Billed:
Search Waybill:
Total (sec):

Export For Order Import **Delete Record** **Clear Fields** **Save/Add Load**

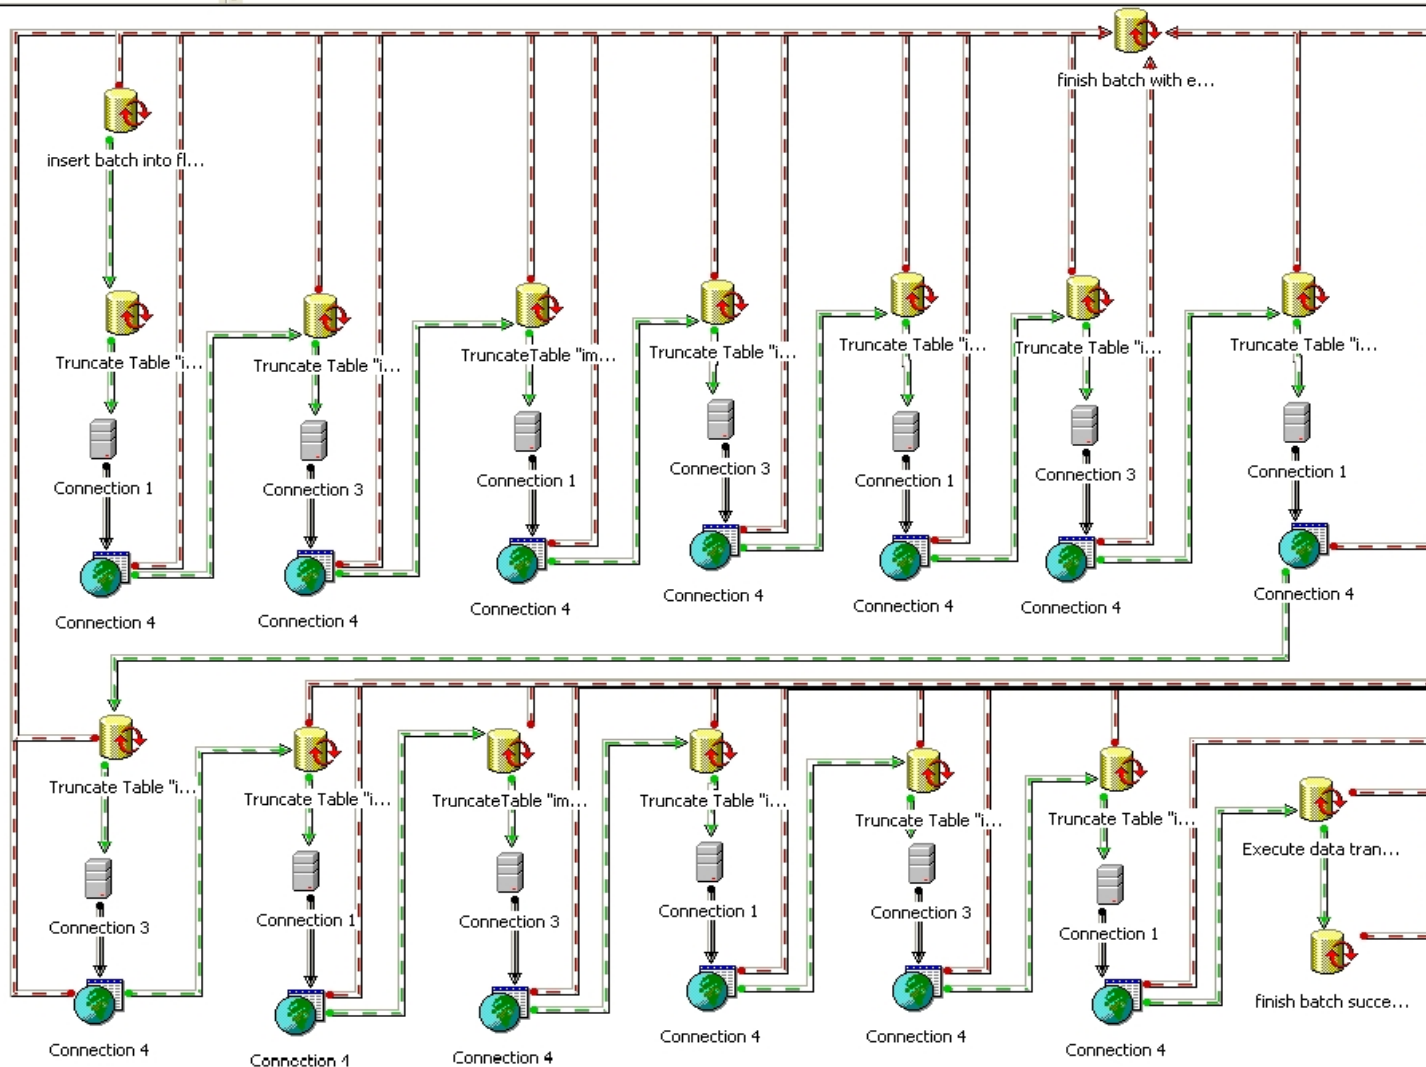
Batch data transformation applications

- Daily synchronization data from MSSQL company resource repository system with PostgreSQL DB – DTS, psqIODBC
- Daily import data from external text file sources (customers, etc) – DTS, psqIODBC
- SQL scripts to transform and calculate data when moving from one schema to other – stat calc., datawarehouse transformation, PHP

Connection



Task



Propagation of PostgreSQL inside company

- Spread informations
- Compare RDBMS features, stability, price, scalability
- Show PostgreSQL support from other companies (Sun, EnterpriseDB, Novell)
- Advantage of following standards (SQL 99)
- Port simple application to PostgreSQL and show performance
- Be prepare for not logical questions and arguments.

Compare RDBMS features, stability, price, scalability

| Purpose | PostgreSQL 8.1 | MySQL 5.0 | FireBird | Oracle 9.i | MSSQL2000 | Sybase |
|--|---|--|---|--|--|---|
| Use with mainstream business application | 4 - Breadth of features support mainstream workloads. | 3 – very new major version – it still include many bugs. | 3 - Support mainstream workloads. Different internal and external SQL | 5 – All standard and extended features support mainstream workloads. Excelent scalability and tuning capability | 4 - All standard and extended features support mainstream workloads. | 4 - All standard and extended features support mainstream workloads. |
| Target for migration from Oracle, DB2, MS SQL, or other server class RDBMS | 5 – ver 8.1 SQL 99 compliance and database features ease migration | 5 - SQL 99 compliance and database features ease migration | 4 - SQL 92 entry level compliance | 5 - SQL 99 compliance and database features ease migration | 5 – SQL 99 compliance and database features ease migration | 5 - SQL 99 compliance and database features ease migration |
| Distribute with a closed source application | 5 - BSD license *, PostGIS GPL lincese ** | 3 - Requires commercial license. | 4 - and *** | 1 – Commercial license required | 1 – Commercial license required | 1 – Commercial license required |
| Use for production quality J2EE/ASP/PHP web site | 4 - Leverages scalability and database features. | 3 - new major version has few bugs and few missing features, missing GIS | 4 - Leverages scalability and database features. | 5 - Leverages scalability and database features. | 4 - Leverages scalability and database features. | 4 - Leverages scalability and database features. |
| Use for datawarehouse, data mining, GIS app. | 5 – All features for complex queries, GIS module, extended data types | 4 – Extremely fast load process, very fast response to simple queries. Complex queries have to be manually tuned, lack of data types | 3 – Queries have to be manually tuned, lack of data types | 5 – All features for complex queries, GIS module, extended data types, standard interface to many 3-th party tools | 4 – All features for complex queries, missing internal GIS module-, lack of provided data types –user-defined data types | 4 – All features for complex queries, missing internal GIS module – external module SQS, lack of provided data types –user-defined data types |
| Cost for version to utilize 2 dualcore CPUs and 8GB RAM | Total: 0 USD | 595 USD/server on Linux Total: 595 USD | Total: 0 USD | Enterprise version 40,000 USD/CPU(*0.75/core) on Linux Total.: 120000 USD | Enterprise version 19,999 USD/CPU + Ent. Win Server - 3,999 USD Total: 43997 USD | 31,679 CAN/CPU on Linux Total: 63358 CAN |

* BSD license allows free use within closed source application

** GPL license allows free use but custom changes to source code have to be open too.

*** [IPL v 1.0](http://www.gnu.org/licenses/licenses.html) and [IDPL](http://www.gnu.org/licenses/licenses.html) licenses are similar to GPL. Their definitions are on

<http://firebird.sourceforge.net/index.php?op=doc&id=ipl> and <http://firebird.sourceforge.net/index.php?op=doc&id=idpl>

Port simple application to PostgreSQL, performance

Test between MySQL 4.1 and PostgreSQL 8.1. Lower value is the better. Indexes are created on the same fields

| Query type/task type | MySQL 4.1 result in sec | PostgreSQL 8.1 result in sec |
|---|----------------------------|-------------------------------------|
| FleatView level 3 – 10users x 20 queries; Used table testvwFulCycle in db test, no indices | 1176.47 | 344.83 |
| FleatView level 3 – 2users x 5 queries; Used table testvwFulCycle in db test, no indices | 38.46 | 8.93 |
| FleatView level 3 – 10users x 20 queries; Used table testvwFulCycle in db test, indices | 747.74 | |
| FleatView level 3 – 2users x 5 queries; Used table testvwFulCycle in db test, indices | 27.03 | 17.86 |
| Load 393635 rows of vwFullCycle table data into db – MySQL MyISAM engine | 7.76 | 27.1 – use only 1 storage engine |
| Load 393635 rows of data from txt file into db – MySQL InnoDB engine | 28.28 | 28.5 – use only 1 storage engine |

Introduction PostgreSQL to developers

- Adjust to new or different datatypes
 - Timestamp – MSSQL datetime
 - Serial – MSSQL bigint identity
 - Bigserial – MSSQL integer identity
 - Timestamp with time zone – MSSQL has no such data type
- Functions – MSSQL stored procedures
- Triggers – new,old row approach versus spec.
table inserted, deleted

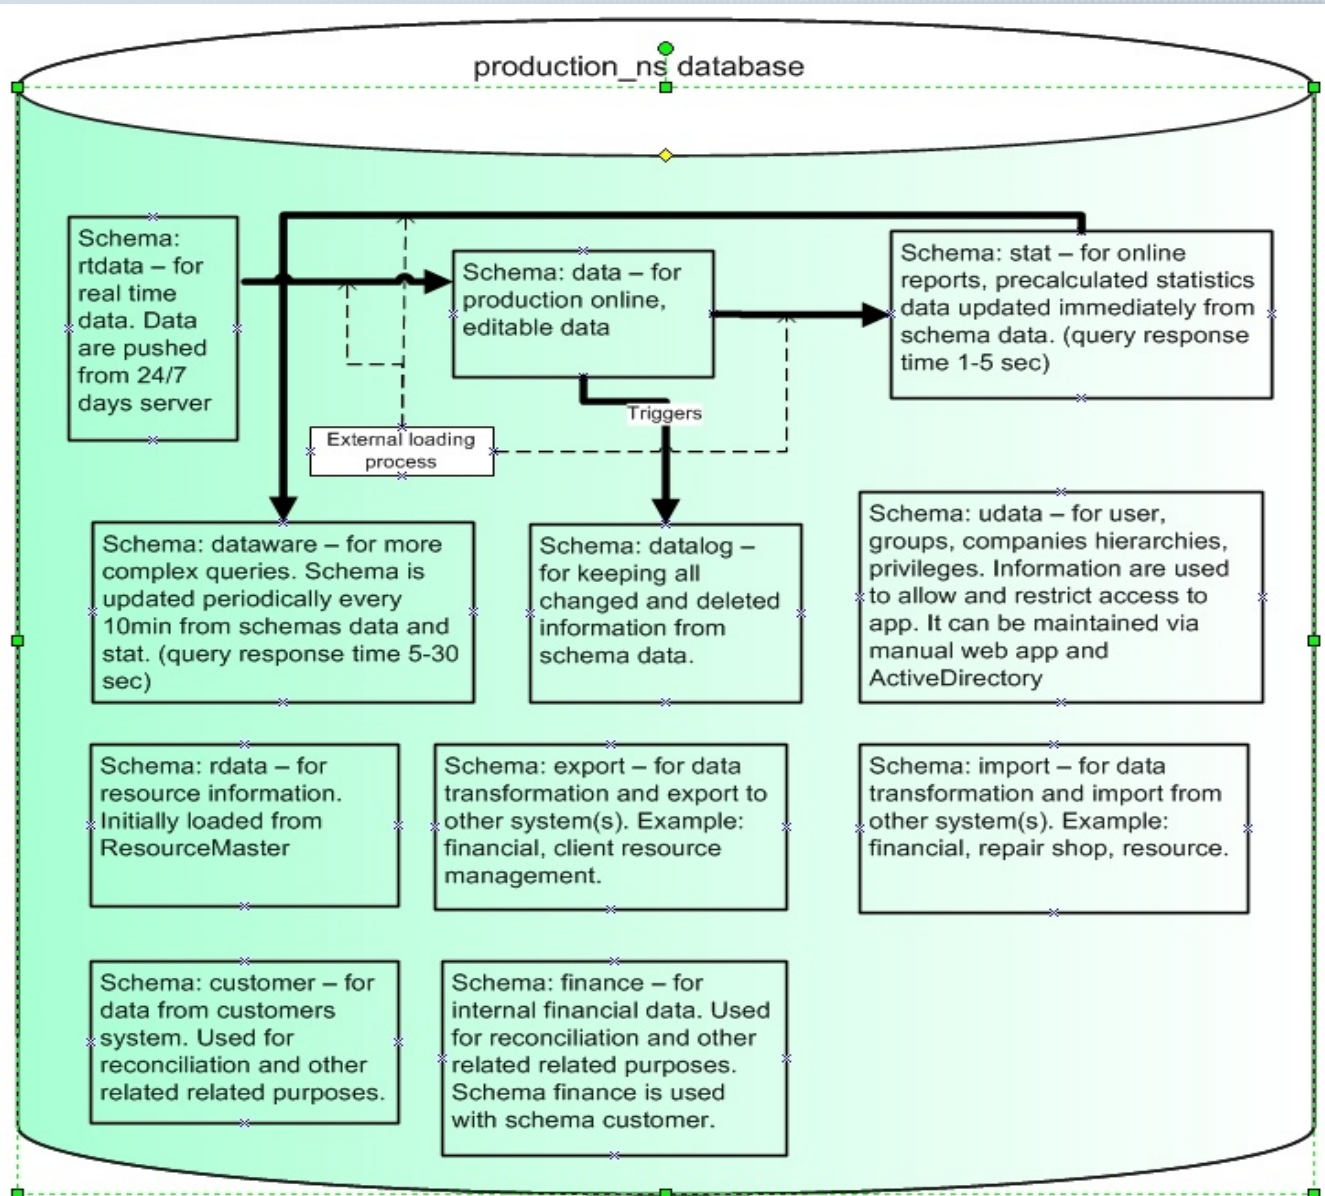
Introduction PostgreSQL to developers

- Problems with data connection components
 - Windows OS
 - ODBC - psqlODBC
 - .NET – slower performance
 - Borland C++ Builder – slow performance, big overhead
 - Native PostgreSQL data components
 - .NET
 - <http://www.crlab.com/pgsqlnet/> - problems with timestamp
 - Borland C++ Builder 5-6
 - microOLAP - PostgresqlDAC problems with bigint, timestamp
 - pgExpress – more for Delphi then for C++ Builder

Introduction PostgreSQL to developers

- Problems with data connection components
 - *NIX OS
 - Native libraries in C, C++ - very good and stable
 - PHP – use native connection to PostgreSQL
 - PEAR
 - DB – old but STABLE, almost no bugs, db abstraction layer
 - MDB2 – current, still some bugs, db abstraction layer
 - ADOdb – MS ADODB like interface, db abstraction layer
 - PECL – in development, bugs, sister to PEAR, compiled packages, db abstraction layer PDO

Database Design. What is correct approach?



Database Design

- Real time schema – eliminate overhead
 - No custom triggers
 - Constraints
 - Foreign Key
 - use indexes wisely. Less is more
- Datawarehouse schema
 - Star and snowflake schemata
 - Transformations data
 - Functions versus SQL batch
 - Indexes

Data migration from MSSQL 2000 to PostgreSQL

- MSSQL DTS
 - Advantage
 - simple development
 - wide variety of drivers – connections
 - flexibility
 - Disadvantage:
 - Performance
 - Resources

Database GUI tools

- CASE – Toad data modeler formally know as CASE Studio 2, ErWin
- SQL Editor – pgAdmin III 1.6.3, Navicat PostgreSQL, Maestro PostgreSQL: comparison
- Graphical Query Editor – important for windows developer
- Administration tools – pgAdmin III for daily maintenance, command line for setup security, performance tuning, server restart

Database GUI tools

- SQL Editor – pgAdmin III 1.6.3, Navicat PostgreSQL, Maestro PostgreSQL: comparison
 - Similar features: database object browsing
 - Difference: export, import data
 - Stability
- Graphical Query Editor – important for windows developer
 - Help with building/starting building queries
 - Creating complex queries with joins
 - Have visual grasp of problem/query
 - Improve performance of programmer
- Administration tools – pgAdmin III for daily maintenance, command line for setup security, performance tuning, server restart
 - Excellent for remote administration tasks
 - Problems with using on Novell linux10 (formerly SLES10)
 - command line - most powerful

Difference between MSSQL and PostgreSQL

- Case sensitivity - difference between MSSQL, MySQL and PostgreSQL
- Index usage – Which index use in PostgreSQL versus MSSQL. Clustered indexes
- Scalability. Advantage of *NIX over Windows, 64bit OS
- TSQL – programmatic approach to solve performance issues
- PL/PGSQL – problem with performance in some situations

Difference between MSSQL and PostgreSQL

- TSQL – programmatic approach to solve performance issues
 - Example: Calculating Running Totals

| Solution | No Index | Index |
|-----------------------------|-----------------|--------------|
| Temp Table/Cursor | 2 sec | 2 sec |
| The "Celko" Solution | 20 sec | 20 sec |
| The "Guru's Guide" Solution | 38 sec | 17 sec |