



Welcome to VoltDB

A Tutorial

Welcome to VoltDB: A Tutorial

Copyright © 2013, 2014 VoltDB, Inc.



Table of Contents

Preface	iv
How to Use This Tutorial	iv
1. Creating the Database	1
Compiling the Application Catalog	1
Starting the Database	1
Using SQL Queries	2
2. Loading and Managing Data	3
Restarting the Database	3
Loading the Data	4
Querying the Database	4
3. Partitioning	7
Partitioned Tables	7
Replicated Tables	8
4. Schema Updates and Durability	9
Saving and Restoring Data	9
Adding and Removing Tables	10
Updating Existing Tables	11
Restoring Data to a New Catalog	11
Durability and Automated Snapshots	12
5. Stored Procedures	14
Simple Stored Procedures	14
Writing More Powerful Stored Procedures	15
Compiling Java Stored Procedures	17
Putting it All Together	18
6. Client Applications	20
Making the Sample Application Interactive	20
Designing the Solution	20
Designing the Stored Procedures for Data Access	21
Creating the LoadWeather Client Application	22
Running the LoadWeather Application	24
Creating the GetWeather Application	24
VoltDB in User Applications	24
VoltDB in High Performance Applications	25
Running the GetWeather Application	27
In Conclusion	28
7. Next Steps	29

Preface

VoltDB is designed to help you achieve world-class performance in terms of throughput and scalability. However, at its core, VoltDB is still a relational database and can do all the things a traditional relational database can do. So before we go into the fine points of tuning a database application for maximum performance, let's just review the basics.

The following tutorial familiarizes you with the features and capabilities of VoltDB step by step, starting with its roots in SQL and walking through the individual features that help VoltDB excel in both flexibility and performance, including:

- Schema Definition
- SQL Queries
- Partitioning
- Schema Updates
- Stored Procedures

For the purposes of this tutorial, let's assume we want to learn more about the places where we live. How many cities and towns are there? Where are they? How many people live there? What other interesting facts can we find out?

Our study will take advantage of data that is freely available from government web sites. But for now, let's start with the basic structure.

How to Use This Tutorial

Of course, you can just read the tutorial to get a feel for how VoltDB works. But we encourage you to follow along using your own copy of VoltDB if you wish.

The data files used for the tutorial are freely available from public web sites; links are provided in the text. However, the initial data is quite large. So we have created a subset of source files and pre-processed data files that is available from the VoltDB web site at the following address for those who wish to try it themselves:

http://voltldb.com/downloads/technologies/other/tutorial_files.zip

For each section of the tutorial, there is a subfolder containing the necessary source files, plus one subfolder, data, containing the data files. To follow along with the tutorial, do the following:

1. Create a folder to work in.
2. Unpack the tutorial files into the folder.
3. At the beginning of each section of the tutorial:
 - a. Set default to your tutorial folder.
 - b. Copy the sources for that current tutorial into the your working directory, like so:

```
$ cp -r tutorial1/* ./
```

The tutorial also uses the VoltDB command line commands. Be sure to set up your environment so the commands are available to you, as described in the installation chapter of *Using VoltDB*.

Part 1: Creating the Database

In VoltDB you define your database schema using SQL data definition language (DDL) statements just like other SQL databases. So, if we want to create a database table for the places where we live, the DDL schema might look like the following:

```
CREATE TABLE towns (  
    town VARCHAR(64),  
    county VARCHAR(64),  
    state VARCHAR(2)  
);
```

The preceding schema defines a single table with three columns: town, county, and state. We could also set options, such as default values and primary keys. But for now we will keep it as simple as possible.

To create a database using this schema, we take three steps:

1. Save the schema as a text file.
2. Compile the schema into an application catalog.
3. Start the database using the catalog created in step #2.

For now, let's save the preceding schema into a file named `towns.sql`. The following sections explain how to create the application catalog and start the database.

Compiling the Application Catalog

Once you have the schema defined, you can create an application catalog. The application catalog is one way VoltDB achieves the performance it does, by pre-compiling important aspects of the database — such as stored procedures and partitioning information — into a single file.

Although we are not interested in performance right now, we still need to create the application catalog before we can start the database. You create the application catalog by compiling the schema:

```
$ voltdb compile towns.sql
```

In this simplest case, the **voltdb compile** command requires only one argument: the name of the schema DDL file. By default, it creates the catalog as `catalog.jar`. If you want to name your catalog file, you can use the `-o` or `--output` flag, like so:

```
$ voltdb compile -o towns.jar towns.sql
```

Assuming no typos or other errors in the DDL, VoltDB compiles the schema and creates the application catalog `towns.jar`.

Starting the Database

Once you create the application catalog, you are ready to start the database. Again, there are several options available when starting a VoltDB database, which we will discuss later. But for now, we can use the simplest command to start the database using the default options on the current machine with the following command:

```
$ voltdb create towns.jar
```

The **voltodb create** command tells VoltDB to create a new, empty database. The argument tells VoltDB which application catalog to use when creating the database.

Using SQL Queries

Congratulations! You have created your first VoltDB database. Of course, an empty database is not terribly useful. So the first thing you will want to do is create and retrieve a few records to prove to yourself that the database is running as you expect.

VoltDB supports all of the standard SQL query statements, such as INSERT, UPDATE, DELETE, and SELECT. You can invoke queries programmatically, through standard interfaces such as JDBC and JSON, or you can include them in stored procedures that are compiled and included in the application catalog.

But for now, we'll just try some ad hoc queries using the sqlcmd command line interface that VoltDB provides. Create a new terminal window and issue the sqlcmd command from the shell prompt:

```
$ sqlcmd
SQL Command :: localhost:21212
1>
```

The VoltDB interactive SQL command line first reports what database it has connected to and then puts up a numbered prompt. At the prompt, you can enter ad hoc SQL queries, execute stored procedures, or type "exit" to end the program and return to the shell prompt.

Let's start by creating some records using the INSERT statement. The following example creates three records, for the towns of Billerica, Buffalo, and Bay View. Be sure to include the semi-colon after each statement.

```
1> insert into towns values ('Billerica','Middlesex','MA');
2> insert into towns values ('Buffalo','Erie','NY');
3> insert into towns values ('Bay View','Erie','OH');
```

We can also use ad hoc queries to verify that our inserts worked as expected. The following queries use the SELECT statement to retrieve information about the database records.

```
4> select count(*) as total from towns;
TOTAL
-----
      3

(1 row(s) affected)
5> select town, state from towns ORDER BY town;
TOWN          STATE
-----
Bay Village   OH
Billerica     MA
Buffalo       NY

(3 row(s) affected)
```

When you are done working with the database, you can type "exit" to end the sqlcmd session and return to the shell command prompt. Then switch back to the terminal session where you started the database and press CTRL-C to end the database process.

This ends Part One of the tutorial.

Part 2: Loading and Managing Data

As useful as ad hoc queries are, typing in data by hand is not very efficient. Fortunately, VoltDB provides several features to help automate this process.

When you compile your application catalog, VoltDB automatically creates stored procedures to insert records for each table. There is also a command line tool that uses these default stored procedures so you can load data files into your database with a single command. The **csvloader** command reads a data file, such as a comma-separated value (CSV) file, and writes each entry as a record in the specified database table using the default insert procedure.

It just so happens that there is data readily available for towns and other landmarks in the United States. The Geographic Names Information Service (GNIS), part of the U.S. Geological Survey, provides data files of officially named locations throughout the United States. In particular, we are interested in the data file for populated places. This data is available as a text file from their web site, http://geonames.usgs.gov/domestic/download_data.htm

The information provided by GNIS not only includes the name, county and state, it includes each location's position (latitude and longitude) and elevation. Since we don't need all of the information provided, we can reduce the number of columns to only the information we want.

For our purposes, let's use information about the name, county, state, and elevation of any populated places. This means we can go back and edit our schema file, `towns.sql`, to add the new columns:

```
CREATE TABLE towns (  
    town VARCHAR(64),  
    state VARCHAR(2),  
    state_num TINYINT NOT NULL,  
    county VARCHAR(64),  
    county_num SMALLINT NOT NUL,  
    elevation INTEGER  
);
```

Note that the GNIS data includes both names and numbers for identifying counties and states. We rearranged the columns to match the order in which they appear in the data file. This makes loading the data easier since the **csvloader** command assumes the data columns are in the same order as specified in the schema.

Finally, we can use some shell magic to trim out the unwanted columns and rows from the data file. The following script selects the desired columns and removes any records with empty fields:

```
$ cut --delimiter="|" --fields=2,4-7,16 POP_PLACES_20120801.txt \  
  | grep -v "|$" \  
  | grep -v "||" > data/towns.txt
```

To save time and space, the resulting file containing only the data we need is included with the tutorial files in a subfolder as `data/towns.txt`.

Restarting the Database

Because we changed the schema we must recompile and restart the database. We use the same commands we used in Part One of the tutorial to do this:

```
$ voltdb compile -o towns.jar towns.sql  
$ voltdb create towns.jar
```


Loading the Data

Once the database is running, we are ready to load our new data. To do this, create a new terminal session, set default to the `/data` subfolder in your tutorial directory, and use the `csvloader` command to load the data file:

```
$ cd data
$ csvloader --separator "|"❶ --skip 1❷ \
           --file towns.txt❸ towns❹
```

In the preceding commands:

- ❶ The `--separator` flag lets you specify the character separating the individual data entries. Since the GNIS data is not a standard CSV, we use `--separator` to identify the correct delimiter.
- ❷ The data file includes a line with column headings. The `--skip 1` flag tells `csvloader` to skip the first line.
- ❸ The `--file` flag tells `csvloader` what file to use as input. If you do not specify a file, `csvloader` uses standard input as the source for the data.
- ❹ The argument, `towns`, tells `csvloader` which database table to load the data into.

The `csvloader` loads all of the records into the database and it generates three log files: one listing any errors that occurred, one listing any records it could not load from the data file, and a summary report including statistics on how long the loading process took and how many records were loaded.

Querying the Database

Now that we have real data, we can perform more interesting queries. For example, which towns are at the highest elevation, or how many locations in the United States have the same name?

```
1> SELECT town,state,elevation from towns order by elevation desc limit 5;
```

TOWN	STATE	ELEVATION
Corona (historical)	CO	3573
Quartzville (historical)	CO	3529
Logtown (historical)	CO	3524
Tomboy (historical)	CO	3508
Rexford (historical)	CO	3484

```
(5 row(s) affected)
```

```
2> select town, count(town) as duplicates from towns
```

```
3> group by town order by duplicates desc limit 5;
```

TOWN	DUPLICATES
Midway	215
Fairview	213
Oak Grove	167
Five Points	150
Riverside	130

```
(5 row(s) affected)
```

As we can see, the five highest towns are all what appear to be abandoned mining towns in the Rocky Mountains. And Springfield, as common as it is, doesn't make it into the top five named places in the United States.

We can make even more interesting discoveries when we combine data. We already have information about locations and elevation. The US Census Bureau can also provide us with information about population density. Population data for individual towns and counties in the United States can be downloaded from their web site, <http://www.census.gov/popest/data/index.html>.

To add the new data, we must add a new table to the database. So let's edit our schema to add a table for population that we will call *people*. While we are at it, we can create indexes for both tables, using the columns that will be used most frequently for searching and sorting, `state_num` and `county_num`.

```
CREATE TABLE towns (
  town VARCHAR(64),
  state VARCHAR(2),
  state_num TINYINT NOT NULL,
  county VARCHAR(64),
  county_num SMALLINT NOT NULL,
  elevation INTEGER
);
CREATE TABLE people (
  state_num TINYINT NOT NULL,
  county_num SMALLINT NOT NULL,
  state VARCHAR(20),
  county VARCHAR(64),
  population INTEGER
);
CREATE INDEX town_idx ON towns (state_num, county_num);
CREATE INDEX people_idx ON people (state_num, county_num);
```

Once again, we put the columns in the same order as they appear in the data file. We also need to trim the data file to remove extraneous columns. The census bureau data includes both measured and estimated values. For the tutorial, we will focus on one population number.

The shell command to trim the data file is the following. (Again, the resulting data file is available as part of the downloadable tutorial package.)

```
$ grep -v "^040," CO-EST2011-Alldata.csv \
  | cut --delimiter="," --fields=4-8 > people.txt
```

Once we have the data and the updated schema, we can stop the database (using CTRL-C), recompile the catalog, restart the database, and load the two tables. So in one terminal session compile and start the database:

```
$ voltdb compile -o towns.jar towns.sql
$ voltdb create towns.jar
```

And in another terminal session load the two data files:

```
$ cd data
$ csvloader --separator "|" --skip 1 --file towns.txt towns
$ csvloader --file people.txt --skip 1 people
```

Now we can join the tables to look for correlations between elevation and population, like so:

```
$ sqlcmd
1> select top 5 min(t.elevation) as height,
2> t.state,t.county, max(p.population)
3> from towns as t, people as p
4> where t.state_num=p.state_num and t.county_num=p.county_num
5> group by t.state, t.county order by height desc;
```

HEIGHT	STATE	COUNTY	C4
2754	CO	Lake	7310
2640	CO	Hinsdale	843
2609	CO	Mineral	712
2523	CO	San Juan	699
2454	CO	Summit	27994

(5 row(s) affected)

It turns out that, even discounting ghost towns that have no population, the five inhabited counties with highest elevation are all in Colorado. In fact, if we reverse the select expression to find the lowest inhabited counties (by changing the sort order from descending to ascending), the lowest is Inyo county in California — the home of Death Valley!

This ends Part Two of the tutorial.

Part 3: Partitioning

Now you have the hang of the basic features of VoltDB as a relational database, it's time to start looking at what makes VoltDB unique. One of the most important features of VoltDB is partitioning.

Partitioning organizes the contents of a database table into separate autonomous units. Similar to sharding, VoltDB partitioning is unique because:

- VoltDB partitions the database tables automatically, based on a partitioning column you specify. You do not have to manually manage the partitions.
- You can have multiple partitions, or sites, on a single server. In other words, partitioning is not just for scaling the data volume, it helps performance as well.
- VoltDB partitions both the data *and* the processing that accesses that data, which is how VoltDB leverages the throughput improvements parallelism provides.

Partitioned Tables

You partition a table by specifying the partitioning column as part of your schema. If a table is partitioned, each time you insert a row into that table, VoltDB decides which partition the row goes into based on the value of the partitioning column. So, for example, if you partition the Towns table on the column Name, The records for all towns with the same name end up in the same partition.

However, although partitioning by name may be reasonable in terms of evenly distributing the records, the goal of partitioning is to distribute both the data *and* the processing. We don't often compare information about towns with the same name. Whereas, comparing towns within a given geographic region is very common. So let's partition the records by state so we can quickly do things like finding the largest or highest town within a given state.

Both the Towns and the People tables have columns for the state name. However, they are slightly different; one uses the state abbreviation and one uses the full name. To be consistent, we can use the State_num column instead, which is common to both tables.

To partition the tables, we simply add a PARTITION TABLE statement to the database schema. Here are the statements we can add to our schema to partition both tables by the State_num column:

```
PARTITION TABLE towns ON COLUMN state_num;  
PARTITION TABLE people ON COLUMN state_num;
```

Having added partitioning information, we can stop the database, recompile our catalog, restart the database and reload the data. The first thing you might notice, without doing any other queries, is that loading the data files is faster.

In fact, when csvloader runs, it creates three log files summarizing the results of the loading process. One of these files, csvloader_TABLE-NAME_insert_report.log, describes how long the process took and the average transactions per second (TPS). Comparing the load times before and after adding partitioning shows that adding partitioning increases the ingestion rate for the Towns table from approximately 5,000 to 16,000 TPS — more than three times as fast! This performance improvement is a result of parallelizing the stored procedure calls across the two sites per host, which is the default setting. Increasing the number of sites per host provides additional improvements, assuming the server has the core processors necessary to manage the additional threads.

Replicated Tables

As mentioned earlier, the two tables `Towns` and `People` both have a `VARCHAR` column for the state name, but its use is not consistent. Instead we use the `State_num` column to do partitioning and joining of the two tables.

The `State_num` column contains the FIPS number. That is, a federal standardized identifier assigned to each state. The FIPS number ensures unique and consistent identification of the state. However, as useful as the FIPS number is for computation, most people think of their location by name, not number. So it would be useful to have a consistent name to go along with the number.

Instead of attempting to modify the fields in the individual tables, we can normalize our schema and create a separate table that provides an authoritative state name for each state number. Again, the federal government makes this information freely available from the U.S. Environmental Protection Agency web site, <http://www.epa.gov/enviro/html/codes/state.html>. Although it is not directly downloadable as a data file, a copy of the FIPS numbers and names for all of the states is included in the tutorial files in the data subfolder as `data/state.txt`.

So let's go and add a table definition for this data to our schema:

```
CREATE TABLE states (  
    abbreviation VARCHAR(20),  
    state_num TINYINT,  
    name VARCHAR(20),  
    PRIMARY KEY (state_num)  
);
```

This sort of lookup table is very common in relational databases. They reduce redundancy and ensure data consistency. Two of the most common attributes of lookup tables are that they are relatively small in size and they are static. That is, they are primarily read-only.

It would be possible to partition the `States` table on the `State_num` column, like we do the `Towns` and `People` tables. However, when a table is relatively small and not updated frequently, it is better to replicate it to all partitions. This way, even if another table is partitioned (such as a customer table partitioned on last name), stored procedures can join the two tables, no matter what partition the procedure executes in.

Tables where all the records appear in all the partitions are called *replicated tables*. Note that tables are replicated by default. So to make the `States` table a replicated table, we simply include the `CREATE TABLE` statement without an accompanying `PARTITION TABLE` statement.

One last caveat concerning replicated tables: the benefits of having the data replicated in all partitions is that it can be read from any individual partition. However, the deficit is that any updates or inserts to a replicated table must be executed in all partitions at once. This sort of multi-partition procedure reduces the benefits of parallel processing and impacts throughput. Which is why you should not replicate tables that are frequently updated.

This ends Part Three of the tutorial.

Part 4: Schema Updates and Durability

Thus far in the tutorial we have restarted the database from scratch and reloaded the data manually each time we changed the schema. This is often the easiest way to make changes when you are first developing your application and making frequent changes. However, as your application — and the data it uses — becomes more complex it is desirable to maintain your database state across sessions.

You may have noticed that in the previous section of the tutorial we modified the schema to add the States table but did not add it to the running database yet. That is because we want to demonstrate ways of modifying the database without having to start from scratch each time.

Note

The following examples make use of functionality available in the VoltDB Enterprise Edition; specifically, the **voltadmin update** command for on-the-fly schema changes. If you are using the VoltDB Community Edition you will need to save, restart, and restore to complete the schema updates manually. See the section called “Restoring Data to a New Catalog” for details.

Saving and Restoring Data

First let's talk about durability. VoltDB is an in-memory database. Each time you start the database with the **create** action, it creates a fresh, empty copy of the database. Obviously, in most real business situations you want the data to persist. VoltDB has several features that preserve the database contents across sessions. We will start by looking at snapshots.

Snapshots are a complete disk-based representation of a VoltDB database, including everything needed to reproduce the database after a shutdown. You can create a snapshot of a running VoltDB database at anytime using the **voltadmin save** command. For example, from our tutorial directory, we can save the data to disk in a snapshot called "townsandpeople".

```
$ HERE=$(pwd)
$ voltadmin save $HERE/voltdbroot/snapshots/ "townsandpeople"
```

The arguments to the **voltadmin save** command are the directory where the snapshot files will be created and the name for the snapshot. Note that the **save** command requires an absolute path for the directory. In the preceding example, we assign the current working directory to a variable so the snapshot can be saved in the subfolder `./voltdbroot/snapshots`. VoltDB creates this folder by default when you start the database. We will learn more about it shortly.

Now that you have a copy of the database contents, we can stop and restart the database. But this time, instead of using the **create** command, we use **recover**:

```
^C
$ voltdb recover
```

When you specify **recover** as the startup action, VoltDB looks for and restores the most recent snapshot. And since the snapshot contains both the catalog and the data, you do not have to specify the catalog on the command line.

We can verify that the database was restored by doing some simple SQL queries in our other terminal session:

```

$ sqlcmd
SQL Command :: localhost:21212
1> select count(*) from towns;
C1
-----
193297

(1 row(s) affected)
2> select count(*) from people;
C1
-----
81691

(1 row(s) affected)

```

Adding and Removing Tables

Now that we know how to save and restore the database, we can add the States table we defined in Part Three. Adding and dropping tables, or changing stored procedures can be done "on the fly", while the database is running. We start by recompiling the application catalog using our new schema definition. Then we use the **voltadmin update** command to update the catalog on the running database.¹

When you update the catalog, you must provide both the catalog and the *deployment file*. VoltDB uses the deployment file to set runtime configuration options, such as the number of servers, how many partitions per host, and so on. Thus far we have used the default settings when starting the database, so have not needed to specify a deployment file. However, VoltDB creates one for you when you do that. So we can use the default deployment file that VoltDB creates to perform the update. So let's recompile the catalog and update the database. We can also check to make sure our new table exists by checking how many records it contains.

```

$ voltdb compile -o towns.jar towns.sql
$ voltadmin update towns.jar voltdbroot/deployment.xml
$ sqlcmd
SQL Command :: localhost:21212
1> select count(*) from states;
C1
---
0

(1 row(s) affected)
3> exit

```

Next we can load the state information from the data file and save a new copy of the database.

```

$ csvloader --skip 1 -f data/states.csv states
$ HERE=$(pwd)
$ voltadmin save $HERE/voltdbroot/snapshots/ "states"

```

¹If you are using the VoltDB Community Edition, you will need to stop and restart the database to change the schema, as described in the section called "Restoring Data to a New Catalog".

Updating Existing Tables

Now that we have a definitive lookup table for information about the states, we no longer need the redundant columns in the Towns and People tables. We want to keep the FIPS column, `State_num`, but can remove the `State` column from each table. Our updated schema for the two tables looks like this:

```
CREATE TABLE towns (  
    town VARCHAR(64),  
    -- state VARCHAR(2),  
    state_num TINYINT NOT NULL,  
    county VARCHAR(64),  
    county_num SMALLINT NOT NULL,  
    elevation INTEGER  
);  
CREATE TABLE people (  
    state_num TINYINT NOT NULL,  
    county_num SMALLINT NOT NULL,  
    -- state VARCHAR(20),  
    town VARCHAR(64),  
    population INTEGER  
);
```

We can recompile the catalog as before.

Many schema changes, including adding and removing columns, can be done on the fly just like adding and dropping tables. However, there are some limitations. For example, you cannot add new uniqueness constraints to an existing index or column. In these cases, the best method is to save the existing data, then restore the data to a fresh copy of the database using the new catalog.

Although the preceding changes are permissible for **voltadmin update**, in the next section we will show you how to change the schema by restoring a snapshot to an updated schema, just so you know how it is done.

Restoring Data to a New Catalog

When you start a database with **recover**, VoltDB restores both the data and the associated catalog. To replace the catalog but keep the data you need to create a new database using the updated catalog and then manually restore the data using the **voltadmin restore** command.

If you haven't already, stop the database, recompile your catalog using the modified schema (which is `townsupdate.sql` in the source files), and create a new database using the updated catalog:

```
^C  
$ cp townsupdate.sql towns.sql  
$ voltdb compile -o towns.jar towns.sql  
$ voltdb create towns.jar
```

Once the database has started, you can use the **voltadmin restore** command from another terminal session to restore the data — but not the catalog — from the previous session. The **restore** command takes the same first two arguments as **save**. That is, the directory where the snapshot resides and the unique identifier.


```
$ HERE=$(pwd)
$ voltadmin restore $HERE/voltdbroot/snapshots/ "states"
HOST_ID  HOSTNAME  SITE_ID  TABLE  PARTITION_ID  RESULT  ERR_MSG
-----
0 pollux  0 STATES  -1 SUCCESS
0 pollux  0 STATES  -1 SUCCESS
0 NULL    0 PEOPLE  1 SUCCESS
0 NULL    0 PEOPLE  0 SUCCESS
0 NULL    0 PEOPLE  1 SUCCESS
0 NULL    0 PEOPLE  0 SUCCESS
0 NULL    0 PEOPLE  1 SUCCESS
0 NULL    0 TOWNS   1 SUCCESS
0 NULL    0 TOWNS   0 SUCCESS
0 NULL    0 TOWNS   0 SUCCESS
0 NULL    0 TOWNS   1 SUCCESS
```

Again, we can use SQL queries to verify that the column has been removed.

```
$ sqlcmd
SQL Command :: localhost:21212
1> select top 1 * from towns order by state_num, county_num;
TOWN          STATE_NUM  COUNTY    COUNTY_NUM  ELEVATION
-----
Autaugaville  1 Autauga   1           49
```

(1 row(s) affected)

Finally, we can save the database one more time. To save space, we can delete the old snapshots from the directory before saving the new one. Like so:

```
$ rm voltdbroot/snapshots/*
$ voltadmin save $HERE/voltdbroot/snapshots/ "tutorial"
```

Durability and Automated Snapshots

Up to now we have manually saved the database when we made changes, which is fine during development or when doing an explicit save and restore to update the catalog. But in most production environments it is important to keep the snapshots up to date in case of accidents or unexpected events, such as system failure.

VoltDB provides several features to ensure the durability of your data. At a minimum, it is a good idea to schedule regular snapshots to maintain a recent copy of the database. Establishing an automated snapshot schedule is one of the configuration options you set when starting the database.

We have been starting VoltDB with the default configuration. However, you can enable specific features using a deployment file. You can see the default deployment file in the voltdbroot subfolder:

```
$ cat voltdbroot/deployment.xml
<?xml version="1.0"?>
<!-- IMPORTANT: This file is an auto-generated ...
... -->

<deployment>
  <cluster hostcount="1" sitesperhost="2" />
  <httpd enabled="true">
    <jsonapi enabled="true" />
  </httpd>
</deployment>
```

The default settings use a single machine, localhost, with two partitions per host and with the http and JSON interfaces enabled. You can copy this file to your working directory and edit it, or create a new deployment.xml file, adding the tags necessary to enable automated snapshots. Your updated deployment file looks like this:

```
<?xml version="1.0"?>
<deployment>
  <cluster hostcount="1" sitesperhost="2" />
  <httpd enabled="true">
    <jsonapi enabled="true" />
  </httpd>
  <snapshot prefix="tutorial"
            frequency="5m"
            retain="3"
  />
</deployment>
```

The attributes of the <snapshot> tag in the preceding example set the name (or prefix) of the snapshots to "tutorial", the frequency of snapshots to every 5 minutes, and tells VoltDB to retain only the three most recent copies. (Older snapshots are deleted to save space.) Now when we start or recover the database we can specify our custom deployment file to have VoltDB automatically snapshot every 5 minutes.

```
$ voltdb recover --deployment=deployment.xml
```

Once automated snapshots are in place, you don't have to manually create a snapshot before stopping the database. You can use the snapshots created automatically by VoltDB, using the command in the preceding example.

Note that the deployment settings are runtime options and are not saved as part of the snapshot. So you should always specify your deployment file on the command line when starting VoltDB unless you want to use the default configuration. VoltDB has several other features that provide additional protection against data loss and increase the resilience of the database, including K-safety, command logging, and database replication. See the *Using VoltDB* manual for more information about availability options.

This ends Part Four of the tutorial.

Part 5: Stored Procedures

We now have a complete database that we can interact with using SQL queries. For example, we can find the least populous county for any given state (California, for example) with the following SQL query:

```
$ sqlcmd
1> SELECT TOP 1 county, abbreviation, population
2>     FROM people, states WHERE people.state_num=6
3>     AND people.state_num=states.state_num
4>     ORDER BY population ASC;
```

However, typing in the same query with a different state number over and over again gets tiring very quickly. The situation gets worse as the queries get more complex.

Simple Stored Procedures

For queries you run frequently only changing the input, you can create a simple stored procedure. *Stored procedures* let you define the query once and change the input values when you execute the procedure. Stored procedures have an additional benefit; because they are pre-compiled as part of the application catalog, the queries do not need to be planned at runtime, reducing the time it takes for each query to execute.

To create simple stored procedures — that is, procedures consisting of a single SQL query — you can define the entire procedure in your database schema using the `CREATE PROCEDURE AS` statement. So, for example to turn our previous query into a stored procedure, we can add the following statement to our schema:

```
CREATE PROCEDURE leastpopulated❶ AS
  SELECT TOP 1 county, abbreviation, population
     FROM people, states WHERE people.state_num=?❷
     AND people.state_num=states.state_num
     ORDER BY population ASC;
```

In the `CREATE PROCEDURE AS` statement:

- ❶ The label, in this case *leastpopulated*, is the name given to the stored procedure.
- ❷ Question marks are used as placeholders for values that will be input at runtime.

In addition to creating the stored procedure, we can also specify if it is single-partitioned or not. When you partition a stored procedure, you associate it with a specific partition based on the table that it accesses. For example, the preceding query accesses the `People` table and, more importantly, narrows the focus to a specific value of the partitioning column, `State_num`.

Note that you can access more than one table in a single-partition procedure, as we do in the preceding example. However, all of the data you access must be in that partition. In other words, all the tables you access must be partitioned on the same key value or, for read-only `SELECT` statements, you can also include replicated tables.

So we can partition our new procedure using the `People` table by adding a `PARTITION PROCEDURE` statement:

```
PARTITION PROCEDURE leastpopulated ON TABLE people COLUMN state_num;
```

Now when we invoke the stored procedure, it is executed only in the partition where the `State_num` column matches the first argument to the procedure, leaving the other partitions free to process other requests.

Of course, before we can use the procedure we need to compile it into the catalog. Modifying stored procedures can be done on the fly, like adding and removing tables. So we do not need to restart the database, just recompile the catalog and update the running database.¹

```
$ voltdb compile -o towns.jar towns.sql
$ voltadmin update towns.jar deployment.xml
```

Once we update the catalog, the new procedure becomes available. (Although, if you are using interactive commands, as in this tutorial, it is a good idea to exit and reenter `sqlcmd` to make sure it recognizes the new procedure name.) So we can now execute the query multiple times for different states simply by changing the argument to the procedure:

```
$ sqlcmd
SQL Command :: localhost:21212
1> exec leastpopulated 6;
COUNTY          ABBREVIATION  POPULATION
-----
Alpine County    CA              1175

(1 row(s) affected)
2> exec leastpopulated 48;
COUNTY          ABBREVIATION  POPULATION
-----
Loving County    TX              82

(1 row(s) affected)
```

Writing More Powerful Stored Procedures

Simple stored procedures written purely in SQL are very handy as short cuts. However, some procedures are more complex, requiring multiple queries and additional computation based on query results. For more involved procedures, VoltDB supports writing stored procedures in Java.

It isn't necessary to be a Java programming wizard to write VoltDB stored procedures. All VoltDB stored procedures have the same basic structure. For example, the following code reproduces the simple stored procedure *leastpopulated* we wrote in the previous section using Java:

```
import org.voltdb.*;❶

public class LeastPopulated extends VoltProcedure {❷

    public final SQLStmt getLeast = new SQLStmt(
        " SELECT TOP 1 county, abbreviation, population "
        + " FROM people, states WHERE people.state_num=? "
        + " AND people.state_num=states.state_num "
        + " ORDER BY population ASC;" );❸
```

¹If you are using the VoltDB Community Edition see the section called “Restoring Data to a New Catalog”.

```

public VoltTable[] run(integer state_num)❹
    throws VoltAbortException {

    voltQueueSQL( getLeast, state_num );❺
    return voltExecutesSQL();❻
}

```

In this example:

- ❶ We start by importing the necessary VoltDB classes and methods.
- ❷ The procedure itself is defined as a Java class. The Java class name is the name we use at runtime to invoke the procedure. In this case, the procedure name is *LeastPopulated*.
- ❸ At the beginning of the class, you declare the SQL queries that the stored procedure will use. Here we use the same SQL query from the simple stored procedure, including the use of a question mark as a placeholder.
- ❹ The body of the procedure is a single *run* method. The arguments to the run method are the arguments that must be provided when invoking the procedure at runtime.
- ❺ Within the run method, the procedure queues one or more queries, specifying the SQL query name, declared in step 3, and the arguments to be used for the placeholders. (Here we only have the one query with one argument, the state number.)
- ❻ Finally, a call executes all of the queued queries and the results of those queries are returned to the calling application.

Now, writing a Java stored procedure to execute a single SQL query is overkill. But it does illustrate the basic structure of the procedure.

Java stored procedures become important when designing more complex interactions with the database. One of the most important aspects of VoltDB stored procedures is that each stored procedure is executed as a complete unit, a transaction, that either succeeds or fails as a whole. If any errors occur during the transaction, earlier queries in the transaction are rolled back before a response is returned to the calling application, or any further work is done by the partition.

One such transaction might be updating the database. It just so happens that the population data from the U.S. Census Bureau contains both actual census results and estimated population numbers for following years. If we want to update the database to replace the 2010 results with the 2011 estimated statistics (or some future estimates), we would need a procedure to:

1. Check to see if a record already exists for the specified state and county.
2. If so, use the SQL UPDATE statement to update the record.
3. If not, use an INSERT statement to create a new record.

We can do that by extending our original sample Java stored procedure. We can start by giving the Java class a descriptive name, *UpdatePeople*. Next we include the three SQL statements we will use (SELECT, UPDATE, and INSERT). We also need to add more arguments to the procedure to provide data for all of the columns in the People table. Finally, we add the query invocations and conditional logic needed. Note that we queue and execute the SELECT statement first, then evaluate its results (that is, whether there is at least one record or not) before queuing either the UPDATE or INSERT statement.

The following is the completed stored procedure source code.

```

import org.voltdb.*;

public class UpdatePeople extends VoltProcedure {

    public final SQLStmt findCurrent = new SQLStmt(
        " SELECT * FROM people WHERE state_num=? AND county_num=?;");
    public final SQLStmt updateExisting = new SQLStmt(
        " UPDATE people SET population=?"
        + " WHERE state_num=? AND county_num=?;");
    public final SQLStmt addNew = new SQLStmt(
        " INSERT INTO people VALUES (?, ?, ?, ?);");

    public VoltTable[] run(byte state_num,
                          short county_num,
                          String county,
                          long population)
        throws VoltAbortException {

        voltQueueSQL( findCurrent, state_num, county_num );
        VoltTable[] results = voltExecutesQL();

        if (results[0].getRowCount() > 0) { // found a record
            voltQueueSQL( updateExisting, population,
                          state_num,
                          county_num );

        } else { // no existing record
            voltQueueSQL( addNew, state_num,
                          county_num,
                          county,
                          population);

        }
        return voltExecutesQL();
    }
}

```

Compiling Java Stored Procedures

Once we write the Java stored procedure, we need to compile it into the catalog the same way we do with simple stored procedures. But first, the Java class itself needs compiling. We use the Java compiler, **javac**, to compile the procedure the same way we would any other Java program.

When compiling stored procedures, the Java compiler must be able to find the VoltDB classes and methods imported at the beginning of the procedure. To do that, we must include the VoltDB libraries in the Java *classpath*. The libraries are in the subfolder `/voltdb` where you installed VoltDB. For example, if you installed VoltDB in the directory `/opt/voltdb`, the command to compile the `UpdatePeople` procedure is the following:

```
$ javac -cp "$CLASSPATH:/opt/voltdb/voltdb/*" UpdatePeople.java
```

Once we compile the source code into a Java class, we can then declare the stored procedure in our schema, in much the same way simple stored procedures are declared. But this time we use the `CREATE PRO-`

PROCEDURE FROM CLASS statement, specifying the class name rather than the SQL query. We can also partition the procedure on the People table, since all of the queries are constrained to a specific value of State_num, the partitioning column. Here are the statements we add to the schema.

```
CREATE PROCEDURE FROM CLASS UpdatePeople;
PARTITION PROCEDURE UpdatePeople ON TABLE people COLUMN state_num;
```

The updated schema definition is included in the tutorial files as townsupdate.sql. So, if you are following along, rename the file to towns.sql to bring your sample files up to date.

```
$ cp townsupdate.sql towns.sql
```

Putting it All Together

OK. Now we have a Java class file for the new stored procedure and an updated schema. We are ready to recompile the catalog.

Just as Java needs to know where the VoltDB classes are when it compiles your stored procedure, VoltDB needs to know where your stored procedure classes are when it compiles the catalog. You tell VoltDB where to look by adding the location of the stored procedure class files to the classpath. You can add them to the classpath using the `-c` or `--classpath` flag to the **voltdb compile** command.

If you compile your Java classes into a subdirectory, which is common for larger projects with multiple classes, that subdirectory would be the path you should specify. Since our example does not specify an alternate location for the Java classes, the classes are in the current working directory and we can use `./` to identify their location.

So we can use the same commands we have used throughout (with the addition of the `--classpath` flag) to compile our catalog and update the database:

```
$ voltdb compile --classpath=./ -o towns.jar towns.sql
$ voltadmin update towns.jar deployment.xml
```

Obviously, we don't want to invoke our new procedure manually for each record in the People table. We could write a program to do it for us. Fortunately, there is a program already available that we can use.

The **csvloader** command normally uses the default INSERT procedures to load data into a table. However, you can specify a different procedure if you wish. So we can use **csvloader** to invoke our new procedure to update the database with every record in the data file.

First we must filter the data to the columns we need. We use the same shell commands we used to create the initial input file, except we switch to selecting the column with data for the 2011 estimate rather than the actual census results. We can save this file as `data/people2011.txt` (which is included with the source files):

```
$ grep -v "^040," data/CO-EST2011-Alldata.csv \
| cut --delimiter="," --fields=4,5,7,11 > data/people2011.txt
```

Before we update the database, let's just check to see which are the two counties with the smallest population:

```
$ sqlcmd
SQL Command :: localhost:21212
1> SELECT TOP 2 county, abbreviation, population
2> FROM people,states WHERE people.state_num=states.state_num
3> ORDER BY population ASC;
COUNTY          ABBREVIATION  POPULATION
-----
Loving County    TX             82
Kalawao County   HI             90
```

(2 row(s) affected)

Now we can run **csvloader** to update the database, using the **-p** flag indicating that we are specifying a stored procedure name rather than a table name:

```
$ csvloader --skip 1 --file data/people2011.txt \
-p UpdatePeople
```

And finally, we can check to see the results of the update by repeating our earlier query:

```
$ sqlcmd
SQL Command :: localhost:21212
1> SELECT TOP 2 county, abbreviation, population
2> FROM people,states WHERE people.state_num=states.state_num
3> ORDER BY population ASC;
COUNTY          ABBREVIATION  POPULATION
-----
Kalawao County   HI             90
Loving County    TX             94
```

(2 row(s) affected)

Aha! In fact, the estimates show that Loving County, Texas is growing and is no longer the smallest!

Part 6: Client Applications

We now have a working sample database with data. We even wrote a stored procedure demonstrating how to update the data. To run the stored procedure we used the pre-existing csvloader utility. However, most applications require more logic than a single stored procedure. Understanding how to integrate calls to the database into your client applications is key to producing a complete business solution. In this lesson, we explain how to interact with VoltDB from client applications.

VoltDB provides client libraries in a number of different programming languages, each with their own unique syntax, supported datatypes, and capabilities. However, the general process for calling VoltDB from client applications is the same no matter what programming language you use:

1. Create a client connection to the database.
2. Make one or more calls to stored procedures and interpret their results.
3. Close the connection when you are done.

This lesson will show you how to perform these steps in several different languages.

Making the Sample Application Interactive

As interesting as the population and location information is, it isn't terribly dynamic. Population does not change that quickly and locations even less so. Creating an interactive application around this data alone is difficult. However, if we add just one more layer of data things get interesting.

The United States National Weather Service (part of the Department of Commerce) issues notices describing dangerous weather conditions. These alerts are available online in XML format and include the state and county FIPS numbers of the areas affected by each weather advisory. This means it is possible to load weather advisories correlated to the same locations for which we have population and elevation data. Not only is it possible to list the weather alerts for a given state and county, we could also determine which events have the highest impact, in terms of population affected.

Designing the Solution

To make use of this new data, we can build a solution composed of two separate applications:

- One to load the weather advisory data
- Another to fetch the alerts for a specific location

This matches the natural break down of activities, since loading the data can be repeated periodically — every five or ten minutes say — to ensure the database has the latest information. Whereas fetching the alerts would normally be triggered by a user request.

At any given time, there are only a few hundred weather alerts and the alerts are updated only every 5-10 minutes on the NWS web site. Because it is a small data set updated infrequently, the alerts would normally be a good candidate for a replicated table. However, in this case, there can be — and usually are — multiple state/county pairs associated with each alert. Also, performance of user requests to look up alerts for a specific state and county could be critically important depending on the volume and use of that function within the business solution.

So we can normalize the data into two separate tables: `nws_alert` for storing general information about the alerts and `local_event` which correlates each alert (identified by a unique ID) to the state and county it applies to. This second table can be partitioned on the same column, `state_num`, as the towns and people tables. The new tables and associated indexes look like this:

```
CREATE TABLE nws_event (  
  id VARCHAR(256) NOT NULL,  
  type VARCHAR(128),  
  severity VARCHAR(128),  
  SUMMARY VARCHAR(1024),  
  starttime TIMESTAMP,  
  endtime TIMESTAMP,  
  updated TIMESTAMP,  
  PRIMARY KEY (id)  
);  
  
CREATE TABLE local_event (  
  state_num TINYINT NOT NULL,  
  county_num SMALLINT NOT NULL,  
  id VARCHAR(256) NOT NULL  
);  
  
CREATE INDEX local_event_idx ON local_event (state_num, county_num);  
CREATE INDEX nws_event_idx ON nws_event (id);  
  
PARTITION TABLE local_event ON COLUMN state_num;
```

It is possible to add the new table declarations to the existing schema file. However, VoltDB allows you to compile multiple schema into a single catalog. So to help organize our source files, we can create the new table declarations in a separate schema file, `weather.sql`. We will also need some new stored procedures, so we won't recompile the catalog right now. But you can view the `weather.sql` file in your tutorial directory to see the new table declarations.

Designing the Stored Procedures for Data Access

Having defined the schema, we can now define the stored procedures that the client applications need. The first application, which loads the weather alerts, needs two stored procedures:

- `FindAlert` — to determine if a given alert already exists in the database
- `LoadAlert` — to insert the information into both the `nws_alert` and `local_alert` table

The first stored procedure is a simple SQL query based on the `id` column and can be defined in the schema. The second procedure needs to create a record in the replicated table `nws_alert` and then as many records in `local_alert` as needed. Additionally, the input file lists the state and county FIPS numbers as a string of six digit values separated by spaces rather than as separate fields. As a result, the second procedure must be written in Java so it can queue multiple queries and decipher the input values before using them as query arguments. You can find the code for this stored procedure in the file `LoadAlert.java` in the tutorial directory.

These procedures are not partitioned because they access the replicated table `nws_alert` and — in the case of the second procedure — must insert records into the partitioned table `local_alert` using multiple different partitioning column values.

Finally, we also need a stored procedure to retrieve the alerts associated with a specific state and county. In this case, we can partition the procedure based on the `state_num` field. This last procedure is called `GetAlertsByLocation`.

The following procedure declarations complete the `weather.sql` schema file:

```
CREATE PROCEDURE FindAlert AS
  SELECT id, updated FROM nws_event
  WHERE id = ?;

CREATE PROCEDURE FROM CLASS LoadAlert;

CREATE PROCEDURE GetAlertsByLocation AS
  SELECT w.id, w.summary, w.type, w.severity,
         w.starttime, w.endtime
  FROM nws_event as w, local_event as l
  WHERE l.id=w.id and
         l.state_num=? and l.county_num = ? and
         w.endtime > TO_TIMESTAMP(MILLISECOND,?)
  order by w.endtime;

PARTITION PROCEDURE GetAlertsByLocation
  ON TABLE local_event COLUMN state_num;
```

Now the stored procedures are written and the additional schema file created, we can compile the Java stored procedure and create the new application catalog. Remember, we wrote the additional schema information into a separate file, so we need to specify both schema files on the command line when we invoke the VoltDB compiler. We can also give the resulting catalog a new name to distinguish it from earlier versions and then update the database:

```
$ javac -cp "$CLASSPATH:/opt/voltdb/voltdb/*" LoadAlert.java
$ voltdb compile --classpath=./ -o weather.jar towns.sql weather.sql
$ voltadmin update weather.jar deployment.xml
```

Creating the LoadWeather Client Application

The goal of the first client application, `LoadWeather`, is to read the weathers alerts from the National Weather Service and load them into the database. The basic program logic is:

1. Read and parse the NWS alerts feed.
2. For each alert, first check if it already exists in the database using the `FindAlert` procedure.
 - If yes, move on.
 - If no, insert the alert using the `LoadAlert` procedure.

Since this application will be run periodically, we should write it in a programming language that allows for easy parsing of XML and can be run from the command line. Python meets these requirements so we will use it for the example application.

The first task for the client application is to include all the libraries we need. In this case we need the VoltDB client library and standard Python libraries for input/output and parsing XML. The start of our Python program looks like this:

```
import sys
from xml.dom.minidom import parseString
from voltdbclient import *
```

The beginning of the program also contains code to read and parse the XML from standard input and define some useful functions. You can find this in the program `LoadWeather.py` in the tutorial directory.

More importantly, we must, as mentioned before, create a client connection. In Python this is done by creating an instance of the `FastSerializer`:

```
client = FastSerializer("localhost", 21212)
```

In Python, we must also declare any stored procedures we intend to use. In this case, we must declare `FindAlert` and `LoadAlert`:

```
finder = VoltProcedure( client, "FindAlert", [
FastSerializer.VOLTTYPE_STRING,
] )
```

```
loader = VoltProcedure( client, "LoadAlert", [
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING,
FastSerializer.VOLTTYPE_STRING
] )
```

The bulk of the work of the application is a set of loops that walk through each alert in the XML structure checking if it already exists in the database and, if not, adding it. Again, the code for parsing the XML can be found in the tutorial directory if you are interested. But the code for calling the VoltDB stored procedures is the following:

```
# Check to see if the alert is already in the database.
response = finder.call([ id ])
if (response.tables):
    if (response.tables[0].tuples):
        # Existing alert
        cOld += 1
    else:
        # New alert
        response = loader.call([ id, wtype, severity, summary,
                               starttime, endtime, updated, fips])
        if response.status == 1:
            cLoaded += 1
```

Note how the application uses the response from the procedures in two different ways:

- The response from `FindAlert` (`finder`) is used to check if any records were returned. If so, the alert already exists in the database.
- The response from `LoadAlert` (`loader`) is used to verify the status of the call. If the return status is one, or success, then we know the alert was successfully added to the database.

There is additional information in the procedure response besides just a status code and the data returned by the queries. But LoadWeather shows two of the most commonly used components.

The last step, once all the alerts are processed, is to close the connection to the database:

```
client.close()
```

Running the LoadWeather Application

Because Python is a scripting language, you do not need to compile your code before running it. However, you do need to tell Python where to find any custom libraries, such as the VoltDB client. Simply add the location of the VoltDB client library to the environment variable PYTHONPATH. For example, if VoltDB is installed in your home directory as the folder `~/voldb`, the command to use is:

```
$ export PYTHONPATH="$HOME/voldb/lib/python/"
```

Once you define PYTHONPATH, you are ready to run LoadWeather. Of course, you will also need weather alerts data to load. A sample file of weather data is included in the tutorial files data directory:

```
$ python LoadWeather.py < data/alerts.xml
```

Or you can pipe the most recent alerts directly from the NWS web site:

```
$ curl http://alerts.weather.gov/cap/us.php?x=0 | python LoadWeather.py
```

Creating the GetWeather Application

Now the database contains weather data, we can write the second half of the solution — an application to retrieve all of the alerts associated with a specific location. In a real world example, the GetWeather application is relatively simple, consisting of a single call to the GetAlertsByLocation stored procedure for the user's current location. Run manually, one query at a time, this is not much of a test of VoltDB — or any database. But in practice, where there can be hundreds or thousands of users running the application simultaneously, VoltDB's performance profile excels.

To demonstrate both aspects of our hypothetical solution, we can write two versions of the GetWeather application:

- A user interface showing what it looks like to the user and how easy it is to integrate VoltDB into such applications.
- A high-performance application to emulate real world loads on the database.

VoltDB in User Applications

The first example adds VoltDB access to a user application. In this case, a web interface implemented in HTML and Javascript. You can find the complete application in the `/MyWeather` folder in the tutorial directory. Run the application by opening the file `GetWeather.html` in a web browser. If you use the sample alert data included in the tutorial directory, you can look up alerts for Adams County in Colorado to see what the warnings look like to the user.

Most of the application revolves around the user interface, including HTML, CSS, and Javascript code to display the initial form and format the results. Only a very small part of the code is related to VoltDB access.

In fact, for applications like this VoltDB has simplified programming interfaces that do not require the explicit setup and tear down of a normal database application. In this case, we can use the JSON interface,

which does not require you to open and close an explicit connection. Instead, you simply call the database with your query and it returns the results in standard JavaScript Object Notation (JSON). VoltDB takes care of managing the connection, pooling queries, and so on.

So the actual database call only takes up two statements in the file `GetAlerts.js`;

- One to construct the URL that will be invoked, identifying the database server, the stored procedure (`GetAlertsByLocation`), and the parameters.
- Another to do the actual invocation and specify a callback routine.

```
var url = "http://localhost:8080/api/1.0/" +
        "?Procedure=GetAlertsByLocation&Parameters=" +
        "[" + statenum + "," + countynum + "," + currenttime + "]";
callJSON(url, "loadAlertsCallback");
```

Once the stored procedure completes, the callback routine is invoked. The callback routine uses the procedure response, this time in JSON format, much the same way the `LoadWeather` application does. First it checks the status to make sure the procedure succeeded and then it parses the results and formats them for display to the user.

```
function loadAlertsCallback(data) {
    if (data.status == 1) {
        var output = "";
        var results = data.results[0].data;
        if (results.length > 0) {
            var datarow = null;
            for (var i = 0; i < results.length; i++) {
                datarow = results[i];
                var link = datarow[0];
                var descr = datarow[1];
                var type = datarow[2];
                var severity = datarow[3];
                var starttime = datarow[4]/1000;
                var endtime = datarow[5]/1000;
                output += '<p><a href="' + link + '">' + type + '</a> '
                    + severity + '<br/>' + descr + '</p>';
            }
        } else {
            output = "<p>No current weather alerts for this location.</p>";
        }
        var panel = document.getElementById('alerts');
        panel.innerHTML = "<h3>Current Alerts</h3>" + output;
    } else {
        alert("Failure: " + data.statusstring);
    }
}
```

VoltDB in High Performance Applications

The second example of `GetWeather` emulates many users accessing the database at the same time. It is very similar to the voter sample application that comes with the VoltDB software.

In this case we can write the application in Java. As we did before with LoadWeather, we need to import the VoltDB libraries and open a connection to the database. The code to do that in Java looks like this:

```
import org.voltdb.*;
import org.voltdb.client.*;

[ . . . ]

    /*
     * Instantiate a client and connect to the database.
     */
    org.voltdb.client.Client client;
    client = ClientFactory.createClient();
    client.createConnection("localhost");
```

The program then does one ad hoc query. You can add ad hoc queries to your applications by calling the `@AdHoc` system procedure with the SQL statement you want to execute as the only argument to the call. Normally, it is best for performance to always use stored procedures since they are precompiled and can be partitioned. However, in this case, where the query is run only once at the beginning of the program to get a list of valid county numbers per state, there is little or no negative impact.

You use system procedures such as `@AdHoc` just as you would your own stored procedures, identifying the procedure and any arguments in the `callProcedure` method. Again, we use the status in the procedure response to verify that the procedure completed successfully.

```
ClientResponse response = client.callProcedure("@AdHoc",
    "Select state_num, max(county_num) from people " +
    "group by state_num order by state_num;");
if (response.getStatus() != ClientResponse.SUCCESS) {
    System.err.println(response.getStatusString());
    System.exit(-1);
}
```

The bulk of the application is a program loop that randomly assigns a state and county number and looks up weather alerts for that location using the `GetAlertsByLocation` stored procedure. The major difference here is that rather than calling the procedure synchronously and waiting for the results, we call the procedure asynchronously and move immediately on to the next call.

```
while ( currenttime - starttime < timelimit) {

    // Pick a state and county
    int s = 1 + (int)(Math.random() * maxstate);
    int c = 1 + (int)(Math.random() * states[s]);

    // Get the alerts
    client.callProcedure(new AlertCallback(),
        "GetAlertsByLocation",
        s, c, new Date().getTime());

    currenttime = System.currentTimeMillis();
    if (currenttime > lastreport + reportdelta) {
        DisplayInfo(currenttime-lastreport);
        lastreport = currenttime;
    }
}
```

Asynchronous procedure calls are very useful for high velocity applications because they ensure that the database always has queries to process. If you call stored procedures synchronously, one at a time, the database can only process queries as quickly as your application can send them. Between each stored procedure call, while your application is processing the results and setting up the next procedure call, the database is idle. Essentially, all of the parallelism and partitioning of the database are wasted while your application does other work.

By calling stored procedures asynchronously, the database can queue the queries and process multiple single-partitioned in parallel, while your application sets up the next procedure invocation. In other words, both your application and the database can run at top speed. This is also a good way to emulate multiple synchronous clients accessing the database simultaneously.

Once an asynchronous procedure call completes, the application is notified by invoking the callback procedure identified in the first argument to the `callProcedure` method. In this case, the `AlertCallback` procedure. The callback procedure then processes the procedure response, which it receives as an argument, just as your application would after a synchronous call.

```
static class AlertCallback implements ProcedureCallback {
    @Override
    public void clientCallback(ClientResponse response) throws Exception {
        if (response.getStatus() == ClientResponse.SUCCESS) {
            VoltTable tuples = response.getResults()[0];
            // Could do something with the results.
            // For now we throw them away since we are
            // demonstrating load on the database
            tuples.resetRowPosition();
            while (tuples.advanceRow()) {
                String id = tuples.getString(0);
                String summary = tuples.getString(1);
                String type = tuples.getString(2);
                String severity = tuples.getString(3);
                long starttime = tuples.getTimestampAsLong(4);
                long endtime = tuples.getTimestampAsLong(5);
            }
            txns++;
            if ( (txns % 50000) == 0) System.out.print(".");
        }
    }
}
```

Finally, once the application has run for the predefined time period (by default, five minutes) it prints out one final report and closes the connection.

```
    // one final report
    if (txns > 0 && currenttime > lastreport)
        DisplayInfo(currenttime - lastreport);

    client.close();
```

Running the GetWeather Application

How you compile and run your client applications depends on the programming language they are written in. For Java programs, like the sample GetWeather application, you need to include the VoltDB JAR file in the class path. If you installed VoltDB as `~/voltdb`, a subdirectory of your home directory, you can add the VoltDB and associated JAR files and your current working directory to the Java classpath like so:


```
$ export CLASSPATH="$CLASSPATH:$HOME/voltdb/voltdb/*:$HOME/voltdb/lib/*:./"
```

You can then compile and run the application using standard Java commands:

```
$ javac GetWeather.java
$ java GetWeather
Emulating read queries of weather alerts by location...
.....
1403551 Transactions in 30 seconds (46783 TPS)
.....
1550652 Transactions in 30 seconds (51674 TPS)
```

As the application runs, it periodically shows metrics on the number of transactions processed. These values will vary based on the type of server, the configuration of the VoltDB cluster (sites per host, nodes in the cluster, etc) and other environmental factors. But the numbers give you a rough feel for the performance of your database under load.

In Conclusion

How your database schema is structured, how tables and procedures are partitioned, and how your client application is designed, will all impact performance. When writing client applications, although the specifics vary for each programming language, the basics are:

- Create a connection to the database.
- Call stored procedures and interpret the results. Use asynchronous calls where possible to maximize throughput.
- Close the connection when done.

Much more information about how to design your client application and how to tune it and the database for maximum performance can be found online in the *Using VoltDB* manual and *VoltDB Performance Guide*.

Part 7: Next Steps

This tutorial has introduced you to the basic features of VoltDB. There are many more capabilities within VoltDB to optimize and enhance your database applications. To see more examples of VoltDB in action, see the sample applications that are included in the VoltDB kit in the `/examples` subfolder. To learn more about individual features and capabilities, see the *Using VoltDB* manual for more information.