

Introduction to IBM DB2

IT4GIS
Keith T. Weber, GISP
GIS Director
ISU-GIS Training and Research Center

Pocatello | Idaho Falls | Meridian | Twin Falls



What is it?

- IBM
- DB2 (Database 2)
- An object-relational database:
 - No software limitations (e.g., size of database, number of tables, number of entities per table).
 - Very scalable (it can grow to fit your needs).
 - Efficient for numerous concurrent clients





Pocatello | Idaho Falls | Meridian | Twin Falls



Concurrent Clients

- GIS for the Enterprise
 - Potential concurrent clients



Pocatello | Idaho Falls | Meridian | Twin Falls



Server Installation

- The amount of random-access memory (RAM) required to run any edition of DB2 is 512 MB.
 - 1 GB is recommended for improved performance
- Required disk space depends on type of installation and type of file system

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Professional Tips




- Administration password
 - DBADMIN
 - ADMINISTRATOR

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

DB2 Editions



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Personal Configuration

- **Personal Edition**

- DB2 Personal Edition is a single-user version of the full DB2 product.
 - Has most features available in Express Edition
 - Remote clients cannot connect to this edition
 - Can be remotely administered with DB2 administration tools
 - Net Search Extender
 - Spatial Extender

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Express Configuration

- **Express Edition**

- Built in autonomic manageability features
- *Supports* high availability (HA) architectures
- Net Search and Spatial Extenders
- pureXML
- Homogeneous federation
- Scalable with Features Packs

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Workgroup Configuration

- **Workgroup Server Edition**

- Contains all of Express Edition
- Includes High Availability Feature Pack
 - Online table reorganizations
 - Tivoli System Automation (TSA) high availability services
 - Support for DB2 advanced copy services
 - High Availability Disaster Recovery (HADR)

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Enterprise Configuration

- **Enterprise Server Edition**
 - DB2 Enterprise Edition is designed for large databases with many users. It contains all the functionality of the Workgroup Edition, plus:
 - A license for an unlimited number of client connections
 - Includes services for parallelism, MDCs, MQTs, table partitioning and more
 - Scalable with additional Feature Packs

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Professional Tips



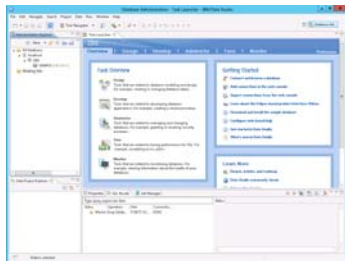
- DB2 licensed per CPU in your server
- Consider DBA overhead
 - Oracle is DBA intensive
 - SQL Server least overhead
 - DB2 to date...

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

DBA

- GUI based database administration
- Alternatively, command prompt can be used.



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

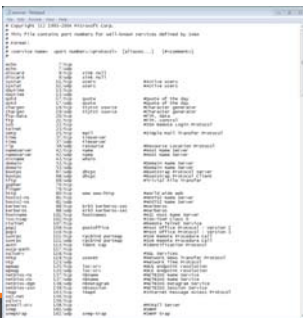
Creating Databases/tables

- Use the Control Center
- A database can be a new **instance** of DB2
- The name you specify must be 1 to 8 characters in length
- To avoid potential problems:
 - do not use the special characters @, #, and \$ in a database name if you intend to have a client remotely connect to a host database.
 - Also, because these characters are not common to all keyboards, do not use them if you plan to use the database in another country.
- On Windows NT-based systems (8, 10, Server 2008, etc.), ensure that no *instance name* is the same as a *service name*.

How do you check this?

Checking Service Names

C:\Windows\System32\drivers\etc




DB2 Specifics

- Pre-fetch
- Buffer pools
- Table data pages



Pre-fetched into the buffer pool!

DB2 Data Types (Numeric)



- FOR BIT DATA (boolean)
- BYTE (0-255)
- SMALLINT (-32,768 to 32,767)
- INTEGER (-2,147,483,648 to 2,147,483,647)
- FLOAT <n>
- DOUBLE PRECISION <n_p,n_s>

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

DB2 Parameters in ArcGIS

- FLOAT < n_p,n_s>
 - n_{precision}(total field length) = 1-6
 - n_{scale} (decimal places) = 1-6
 - n_p,n_s = 5,3 → 26.589 is OK, 256.381 is not
 - Five (5) total characters 2 6 . 5 8 9

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Parameters (cont'd)

- DOUBLE PRECISION <n_p,n_s>
 - n_p = 7 or more
 - n_s = 0 or more

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

DB2 Data Types (Character)

- CHARACTER<n>
- VARCHAR<n>




Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Parameters (cont'd)

- CHARACTER<n>
 - (AKA, String or Text)
 - Example a field named "URL" with $n = 46$
 - <http://giscenter.isu.edu/training/it4gis.htm>




Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

DB2 Data Types (Special)

- DATE
- TIME
- TIMESTAMP



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

DB2 Data Types (Special)

- Stored in special *System managed* tables

- BLOB<n[K|M|G]>
- CLOB<n[K|M|G]>
- DBCLOB<n[K|M|G]>
- GRAPHIC<n>
- VARCHAR<n>





Table Data Pages

- All fields with *standard* data types for each record are contained within a single data page.
- There is a maximum of 255 records stored on each page.
- The ART of efficient data modeling is to minimize wasted space on a page while maximizing the proportion of each page written.

An Instance Example



Each record contains 10 fields
 Each record needs 20 bytes (0.02kb)
 How many records will fit into a 4kb page?

What percentage of a page is written?
 -based on available space?
 -based on available records?

one 8kb page

DB2 for GIS

- **DB2 Spatial Extender** lets you integrate geographic data with your existing business data. It includes:
 - Data types such as points, lines, and polygons
 - Functions such as area, endpoint, and intersect
 - An indexing scheme for spatial data
 - This product is available for all editions of DB2

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Questions?



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY



Key Concepts

- Understand that while data is stored in tables, the tables span TABLE PAGES
- Understand what PRE-FETCH and CACHE are...and how they differ.
- Understand DB2 specific data types

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Your Assignment

- Read IBM DB2 Ref (PDF).
- Read Spatial Data Ext (PDF).
- Complete the exercise
 - Design table pages with the “DB2 Database Administration” exercise
