

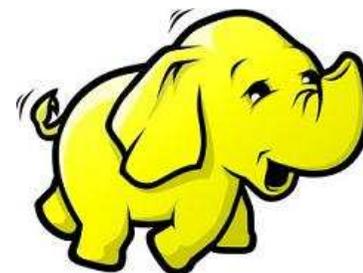
# 第二屆台灣 Hadoop 使用者社群會議

## Hadoop於GIS上之應用 - 以運輸管理為例

辜文元

逢甲大學地理資訊系統研究中心

2010/12/2



Hadoop Taiwan  
User Group

<http://www.hadoop.tw>

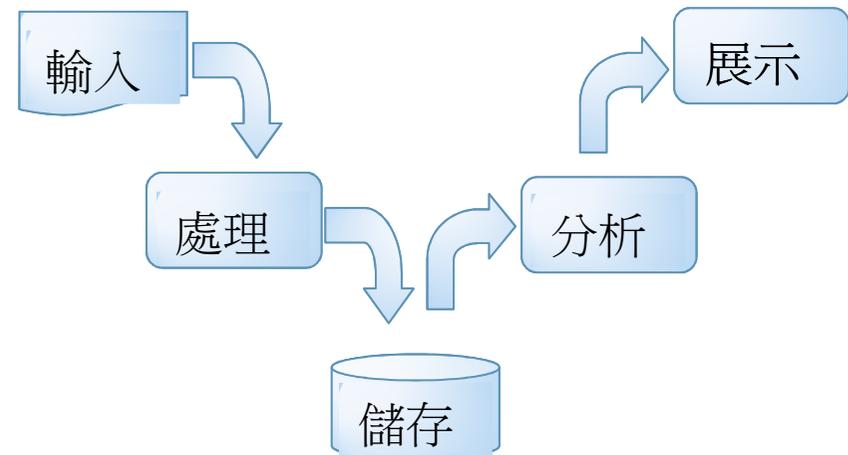
# outline

- 研究背景
- Hadoop - HBase
- 系統分析與設計
- 系統整合
- 結論
- 問題與討論

# 研究背景

# 地理資訊系統-GIS

- Geographic Information Systems (GIS)  
地理資訊系統係為一套整合各項相關地理資訊化作業
- 架構於一完整豐富的地理資料庫之上
- 具有資料擷取、編修、更新、儲存、查詢、處理、分析及展示等功能

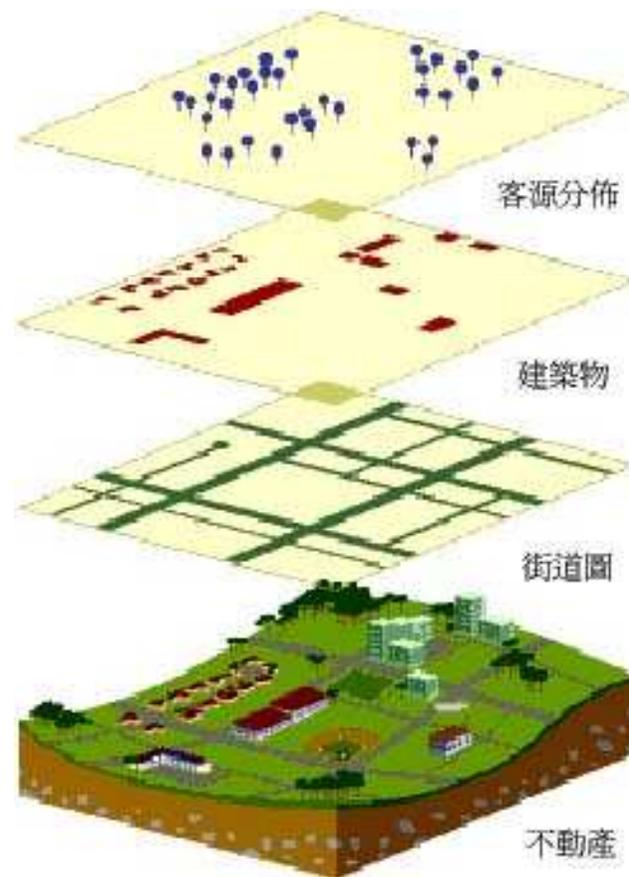


# GIS功能概觀

資料獲取

資料儲存

資料查詢



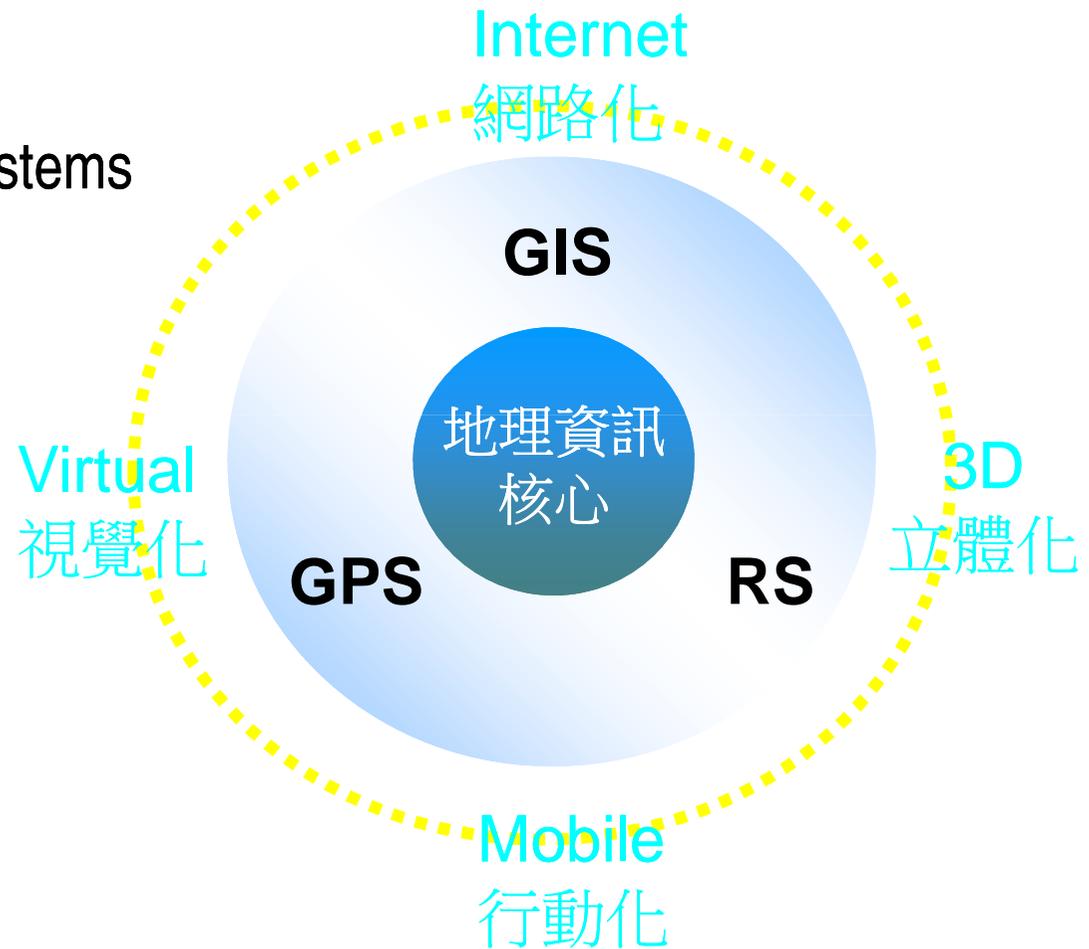
資料分析

成果展示

決策應用

# GIS內涵--技術核心

- Geographic Information Systems  
地理資訊系統
- Remote Sensing  
遙感探測
- Global Positioning System  
全球衛星定位系統



# GIS在運輸資源整合應用

7

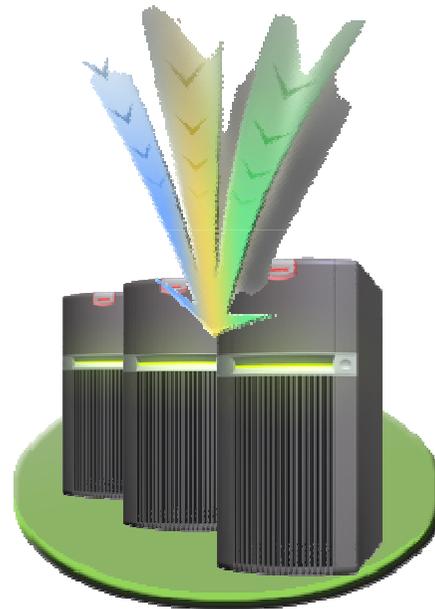
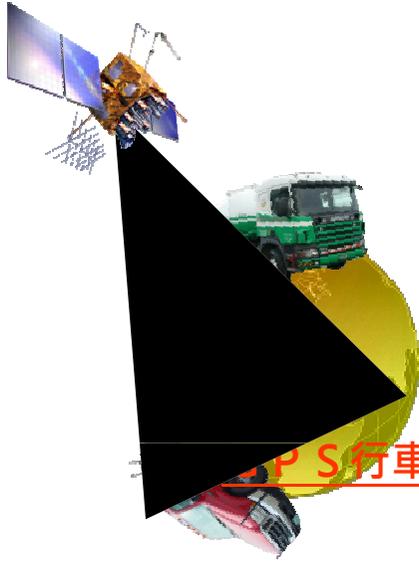
訂單



車輛人員排班



GPS行車記錄



客戶  
運費計算  
貨物即時查詢



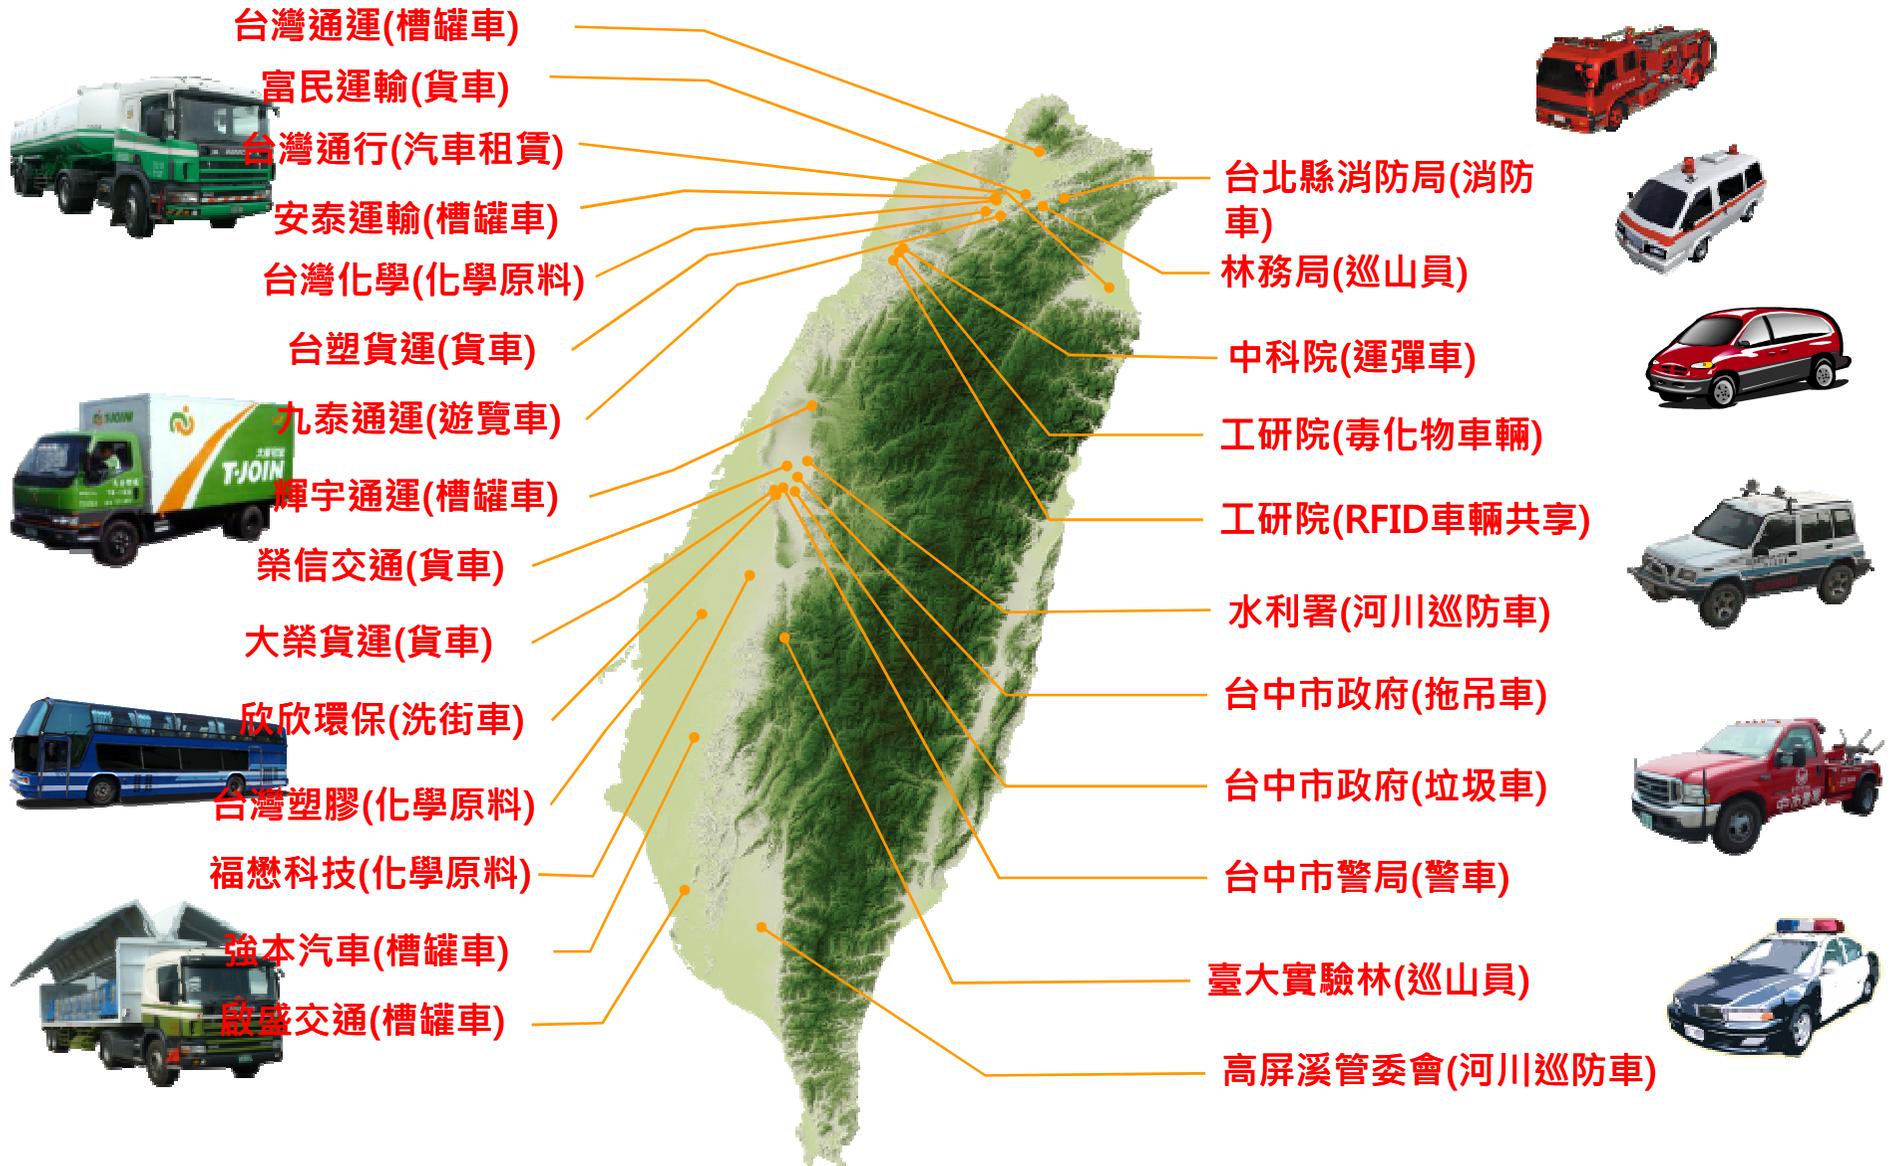
貨運業者  
司機旅費計算  
營運分析



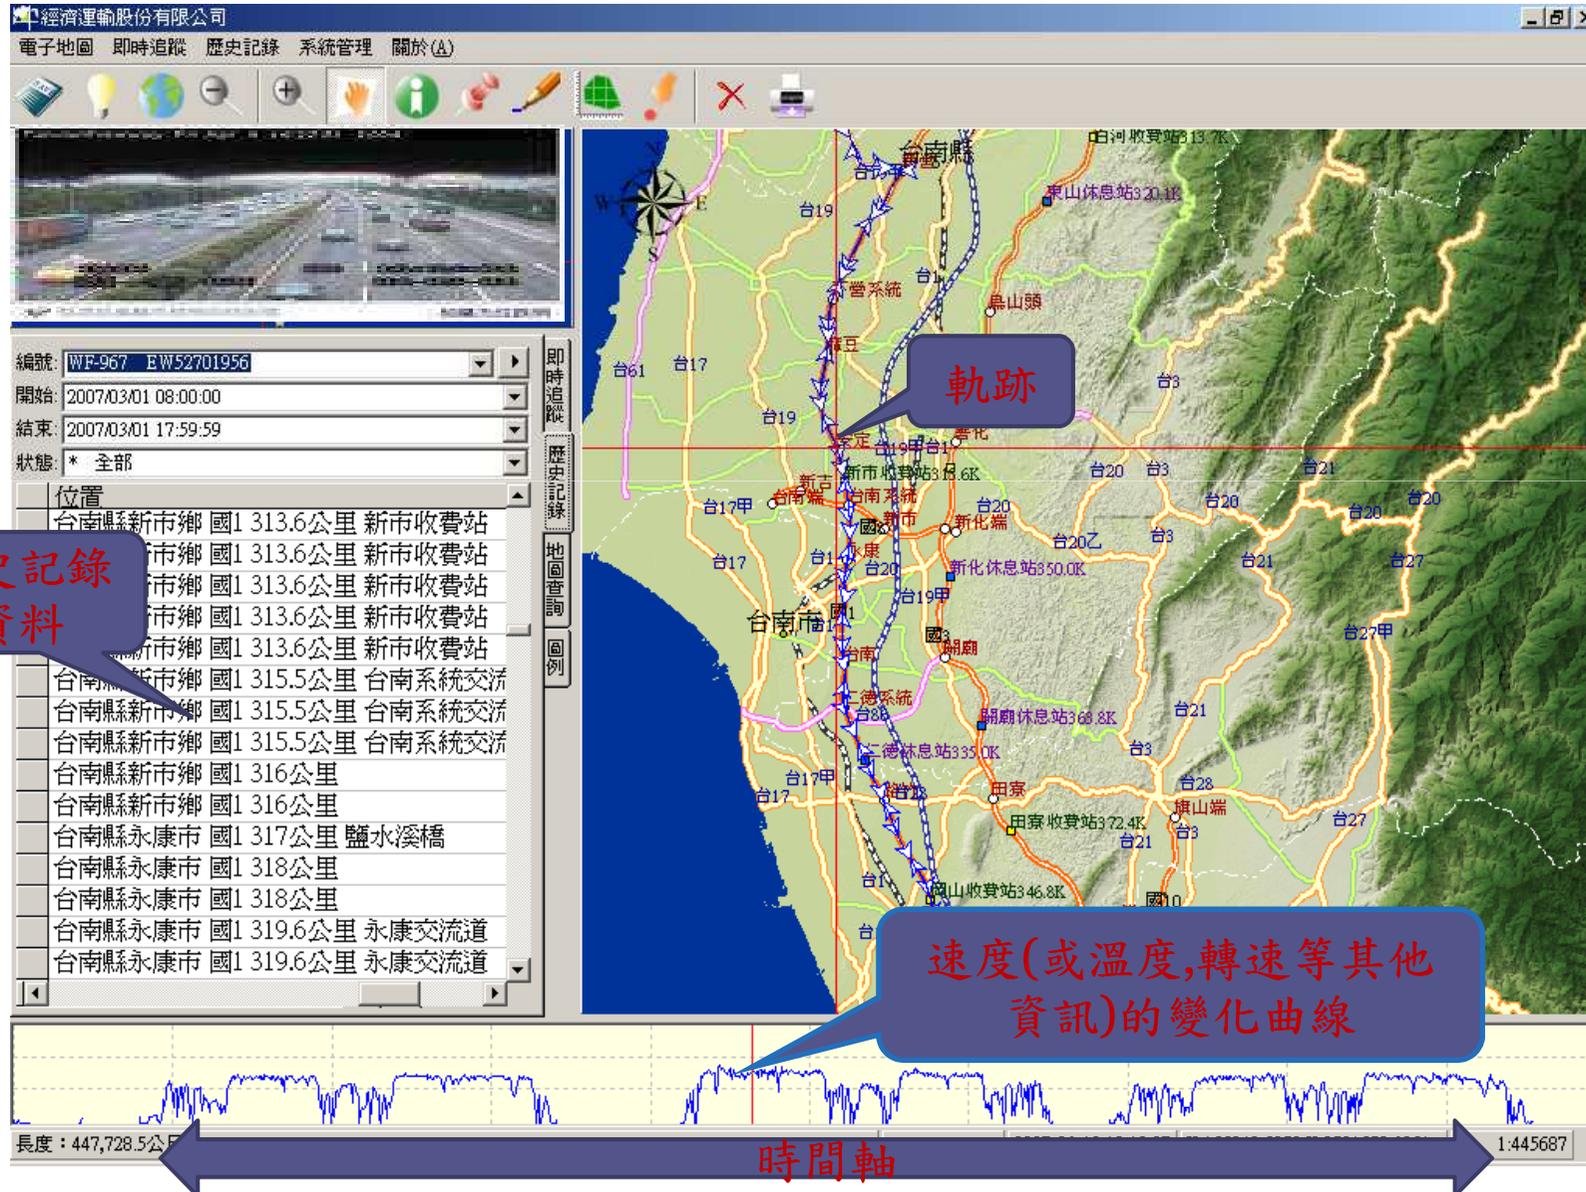
# 運輸管理系統架構



# 實際案例—天眼車隊管理系統



# 提供即時及歷史記錄查詢



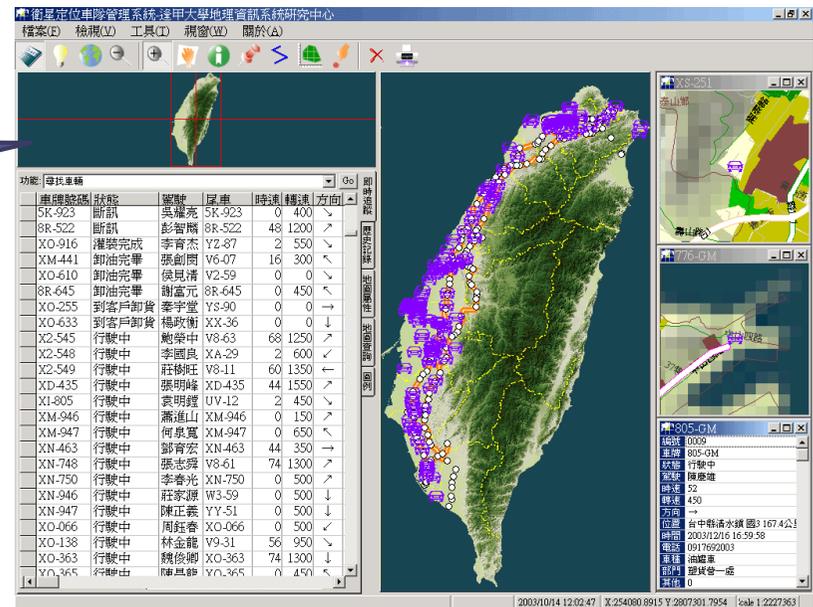
# 營運情形

- 營運車輛數：約3000輛
- 每日資料量：約150萬筆行車記錄 / 150MB儲存空間
- 資料庫：Microsoft SQL Server 2005
- 歷史資料總量：
  - 自91年開始營運
  - 約36億筆行車記錄 / 200GB

# 遭遇問題

- 營運車輛數持續增加，資料庫負荷變重
- 資料不斷累積，影響資料庫存取效率
- 資料庫僅保留3個月行車記錄

尋找雲端 / NoSQL  
Solution



# 尋找適合解決方案

- 雲端技術為基礎
- 分散式儲存架構
- 高可用性
- 延展性



# Hadoop - HBase



# 什麼是HBase

- 基於Hadoop的分散式資料庫解決方案
- OpenSource版的 Google BigTable
- 特性
  - Distributed
  - Column-Oriented
  - Multi-Dimensional
  - High-Availability
  - High-Performance

# Data Model

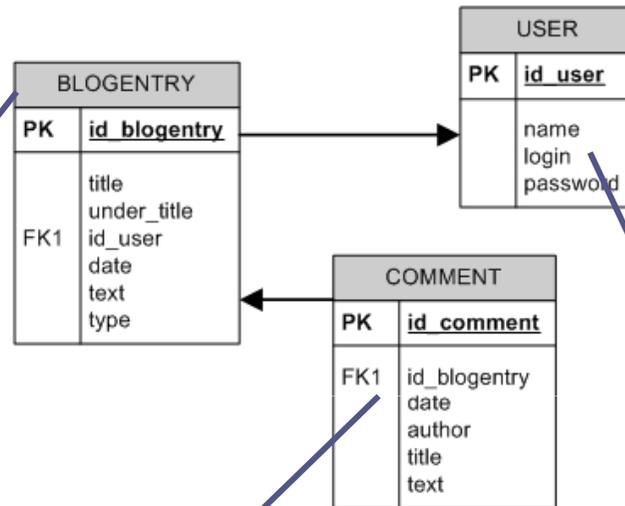
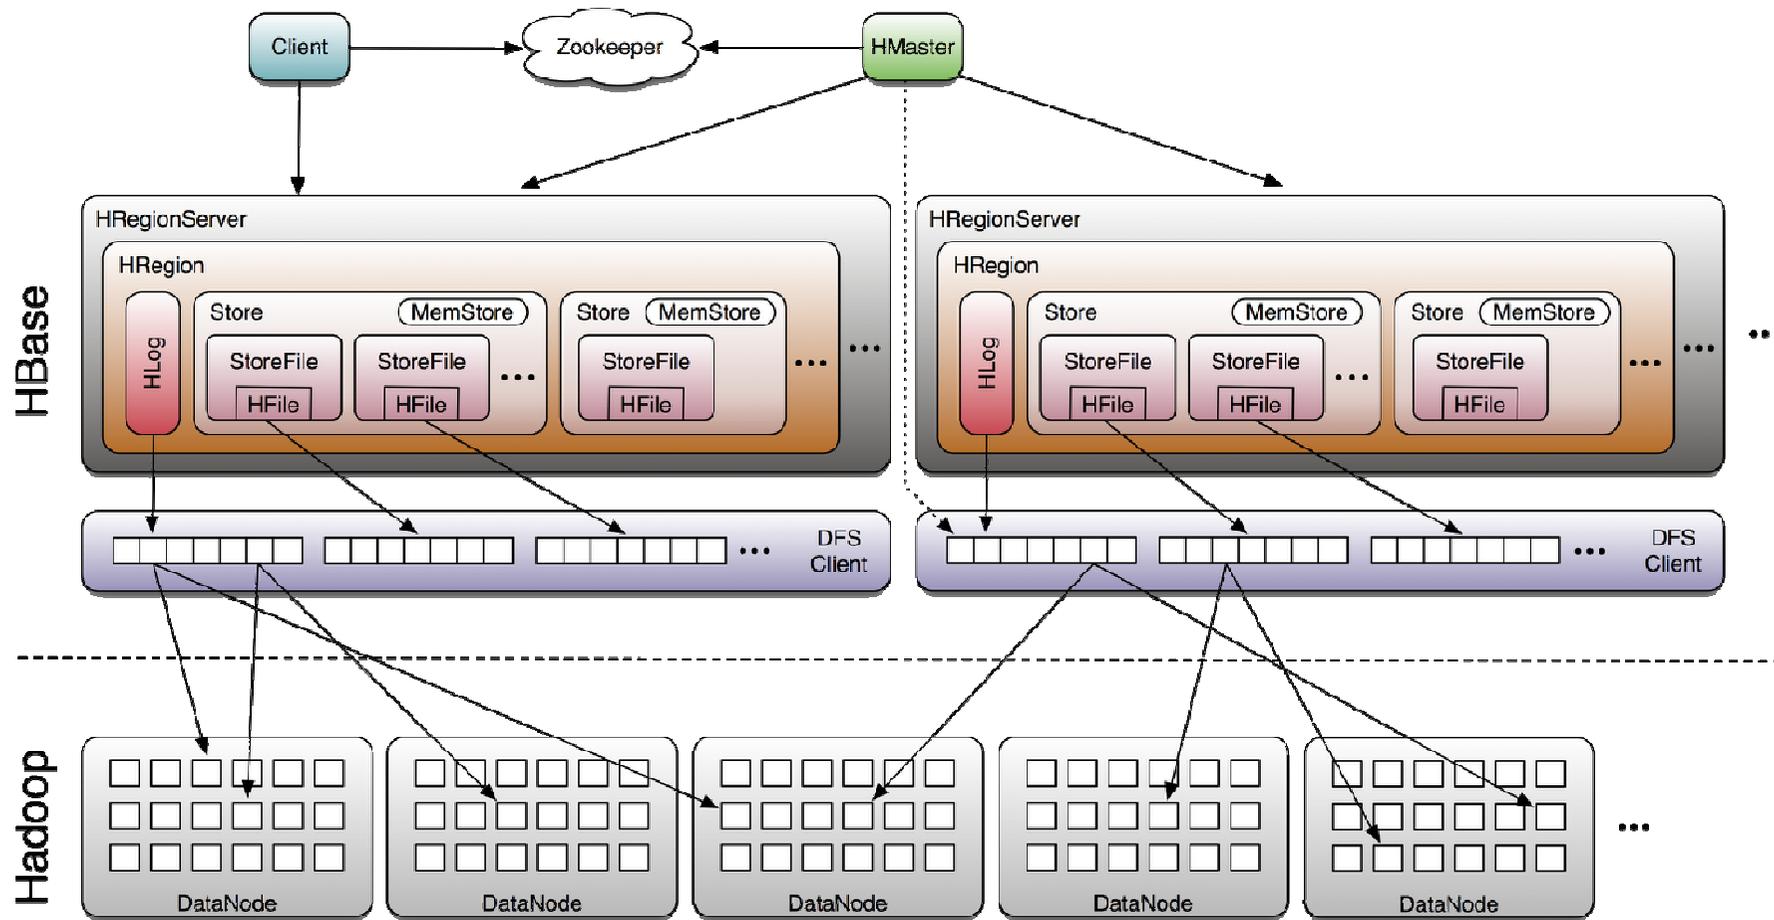


Table	Row Key	Family	Attributs
blogtable	TTYYYMMDDHHmmss	info:	Always contains the column keys author,title,under_title. Should be IN-MEMORY and have a 1 version
		text:	No column key. 3 versions
		comment_title:	Column keys are written like YYMMDDHHmmss. Should be IN-MEMORY and have a 1 version
		comment_author:	Same keys. 1 version
		comment_text:	Same keys. 1 version
usertable	login_name	info:	Always contains the column keys password and name. 1 version

# HBase Architecture

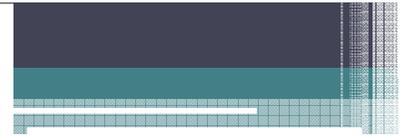


From : <http://www.larsgeorge.com/>



# Table: tbSkyEyeLog

[Master](#), [Local logs](#), [Thread Dump](#), [Log Level](#)



## Table Attributes

Attribute Name	Value	Description
Enabled	true	Is the table enabled
Fragmentation	7%	How fragmented is the table. After a major compaction it is 0%.

## Table Regions

Name	Region Server	Encoded Name	Start Key	End Key
tbSkyEyeLog.1290038593445	<a href="#">cloud-a:60030</a>	318472633		1016753735-2010-03-21_18:48:49
tbSkyEyeLog.1016753735-2010-03-21_18:48:49,1290038593445	<a href="#">cloud-b:60030</a>	1062588561	1016753735-2010-03-21_18:48:49	1026289304-2010-05-17_17:12:28
tbSkyEyeLog.1026289304-2010-05-17_17:12:28,1289980266225	<a href="#">cloud-a:60030</a>	356927351		
tbSkyEyeLog.1026289864-2010-05-03_06:40:56,1289980121793	<a href="#">cloud-b:60030</a>	186378921		
tbSkyEyeLog.1026293045-2010-05-18_06:16:08,1289980121793	<a href="#">cloud-c:60030</a>	210981552		
tbSkyEyeLog.1026293541-2010-05-07_11:59:23,1289982715622	<a href="#">cloud-b:60030</a>	1493815431		
tbSkyEyeLog.1026294468-2010-05-21_22:08:24,1290055565820	<a href="#">cloud-a:60030</a>	1308410366		
tbSkyEyeLog.1026910038-2010-07-15_10:19:37,1290055565820	<a href="#">cloud-b:60030</a>	2117705715		

Contents of directory [/hbase/tbSkyEyeLog](#)

Goto :

Contents of directory [/hbase/tbSkyEyeLog/1062588561](#)

Goto :

[Go to parent directory](#)

Name	Type	Size	Replication	Block Size	Modification Time	Permission	Owner	Group
<a href="#">.regioninfo</a>	file	0.73 KB	3	64 MB	2010-11-18 08:03	rw-r--r--	hadoop	supergroup
<a href="#">Log</a>	dir				2010-11-30 11:28	rw-r--r--	hadoop	supergroup

directory

Type	Size	Replication	Block Size	Modification Time	Permission	Owner	Group
dir				2010-11-30 08:33	rw-r--r--	hadoop	supergroup
dir				2010-11-29 17:38	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:34	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:33	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:33	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:33	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:34	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:20	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:33	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:34	rw-r--r--	hadoop	supergroup
<a href="#">1062588561</a>	dir			2010-11-30 08:33	rw-r--r--	hadoop	supergroup
dir				2010-11-29 17:38	rw-r--r--	hadoop	supergroup
dir				2010-11-30 08:34	rw-r--r--	hadoop	supergroup



# Hadoop 實驗叢集環境



- 於GIS.FCU完成實驗叢集架設
  - 由3部個人電腦組成 Cloud 實驗叢集
    - AMD Phenom X4 2.3GHz, 4G RAM
- 軟體
  - Ubuntu 9.10
  - Hadoop 0.20.1
  - Hbase 0.20.3

## NameNode 'cloud-a:9000'

**Started:** Fri Nov 12 17:03:58 CST 2010  
**Version:** 0.20.2-dev, r  
**Compiled:** Thu Oct 28 16:07:00 CST 2010 by hadoop  
**Upgrades:** There are no upgrades in progress.

[Browse the filesystem](#)  
[Namenode Logs](#)  
[Go back to DFS home](#)

### Live Datanodes : 3

Node	Last Contact	Admin State	Configured Capacity (GB)	Used (GB)	Non DFS Used (GB)	Remaining (GB)	Used (%)	Used (%)	Remaining (%)	Blocks
cloud-a	0	In Service	449.27	203.6	87.51	158.15	45.32		35.2	6934
cloud-b	0	In Service	449.27	203.6	34.14	211.53	45.32		47.08	6929
cloud-c	2	In Service	449.27	203.62	35.66	209.99	45.32		46.74	6933

# 系統分析與設計

Row-key設計

Table設計

既有資料匯入

# 運輸管理系統架構



# 車機回傳格式

- 車輛編號,衛星時間,X,Y(TM2),更改點位標記,未明,狀態,狀態2,速度,方向,轉速,里程,累積里程,差異時間,停留,接收序號,點位編號,車輛標記,方向圖示編號,X,Y(經緯度),接收衛星數,ACC狀態
- **58213063, 2010-01-02 16:29:42**, ,193018, 2488885,False,1,91,,0,0,0,0,1,,184263014,4074 58,0

# row-key design #1 / 2

- HBase沒有實體的index
- 資料透過row-key進行搜尋

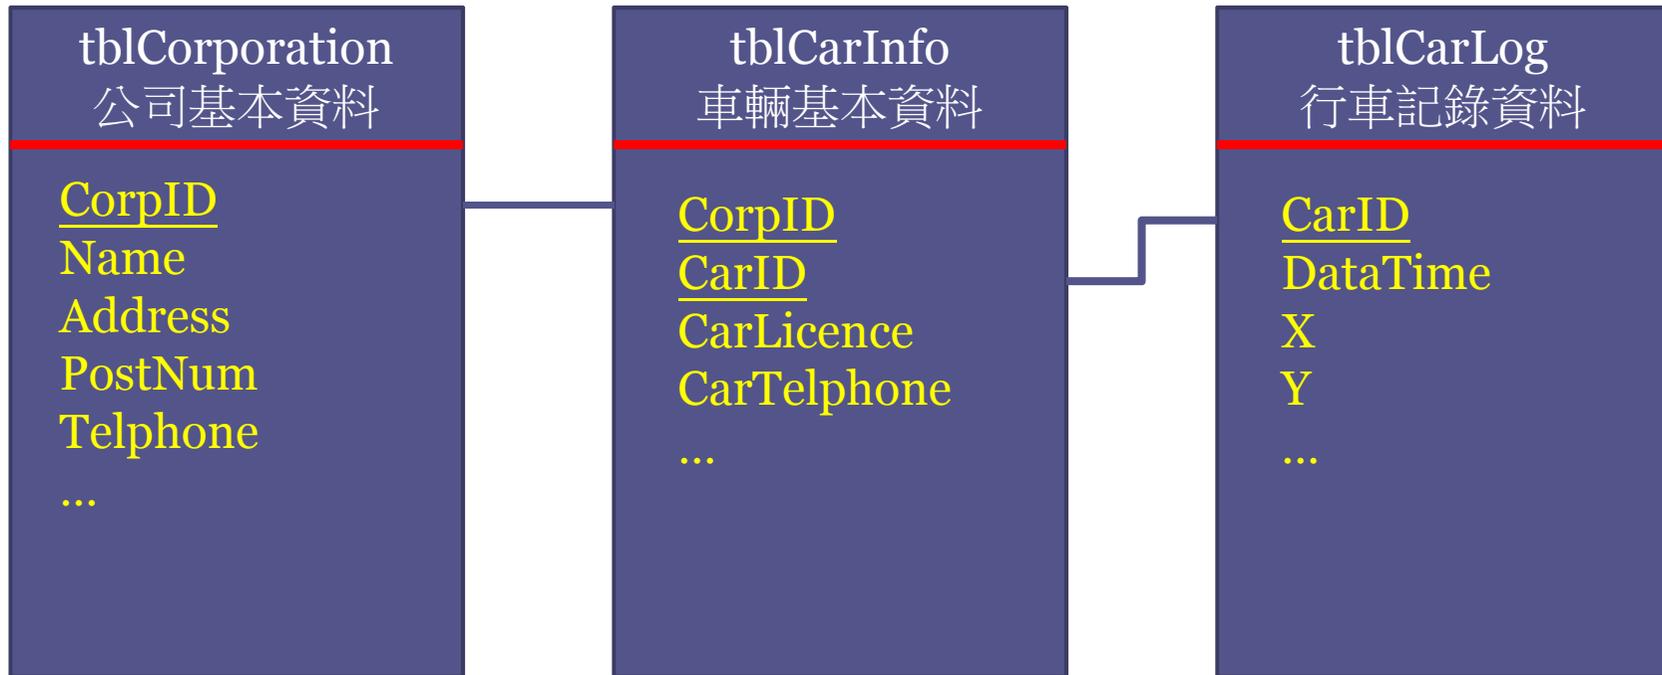
Row Key	Time Stamp	Column (Family) “content:”	Column (Family) “anchor:”	
com.cnn.www	t9	“<html>...”	“anchor:cnnsi.com”	“CNN”
	t8		“anchor:cnnsi.com”	“CNN”
				“anchor:my.lock.ca”
	t6	“<html>...”		

## row-key design #2/2

- 車輛行車記錄具備有時間特性
- 藉由車輛編號和時間組合Rowkey
  - `<carid>_<data date>_<data time>`
  - Ex:
    - Cadid :8S-9115
    - time: 2010-12-02 08:00:00
    - Rowkey: 8S-9115\_2010-12-02\_08:00:00

# HBase table design #1

- 關聯式資料表結構



# HBase table design-#2

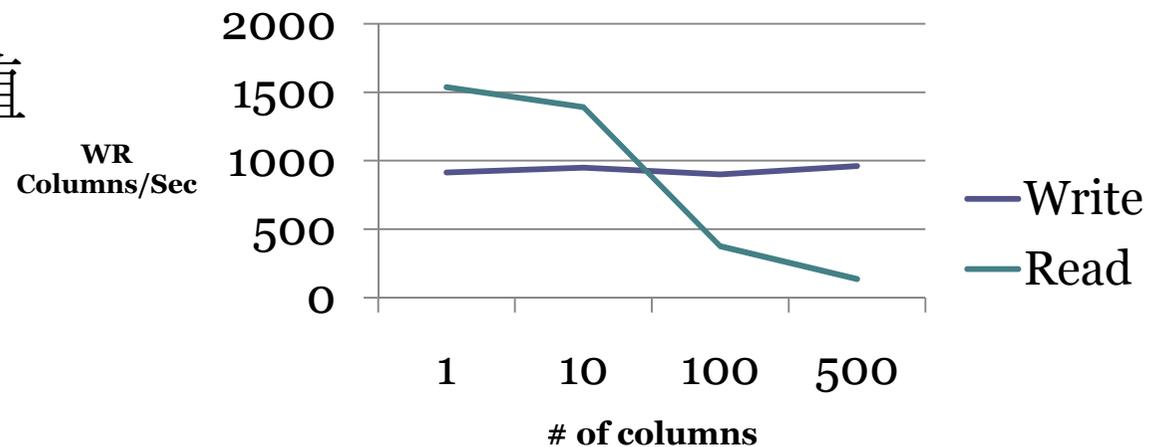
- Family Columns對於“資料列”存取影響

	# of column families (1 Region Server)					# of column families (5 Region Servers)			
Experiment	10	100	500	1000	Experiment	10	100	500	1000
Setup Time (sec)	12.3	19.2	45.9	Timeout	Setup Time (sec)	12.2	18.7	46.4	Timeout
Writes/Sec	164	323	419	Timeout	Writes/Sec	29	153	376	Timeout
Reads/Sec	99	139	122	Timeout	Reads/Sec	119	111	120	Timeout
WriteRows/Sec	16.4	3.23	0.836	X	WriteRows/Sec	2.37	1.53	0.752	X

Tables 1 & 2: Number of 1000-byte values written to and read from a single row with increasing number of column families.

# 實際測試

- 使用Dana測試程式重做實驗
- 3 Region Server
- 測試10次取平均值



	# of columns families			
Experiment	1	10	100	500
Writes/Sec	160.223	400.235	520.136	734.214
Reads/Sec	1169	1142	1114	1116
WriteRows/Sec	160.223	40.023	5.201	1.468

# HBase table design-#3

- 在單一個Column Family中
  - 不同的Column數對於“資料列”存取影響

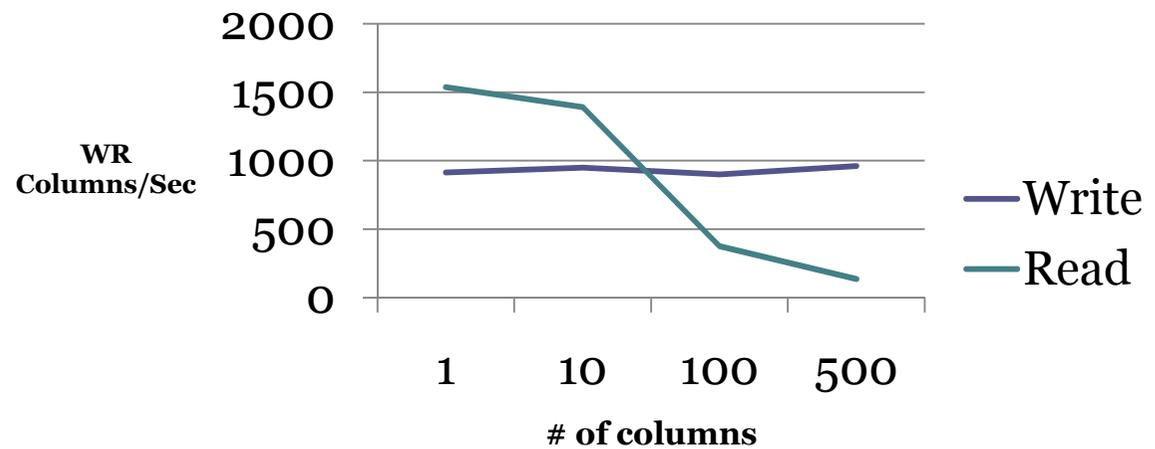
	# of columns (1 Region Server)		
Experiment	10,000	100,000	1,000,000
Writes/Sec	547	403	Crash
Reads/Sec	138	32	Crash

	# of columns (5 Region Servers)		
Experiment	10,000	100,000	1,000,000
Writes/Sec	470	435	Crash
Reads/Sec	96	18	Crash

Tables 3 & 4: Number of 1000-byte values written to and read from a single row with increasing number of columns.

# 實際測試

- 測試10次取平均值
- 寫入100筆資料



	# of columns			
Experiment	1	10	100	500
WriteCols/Sec	914	949	900	961
Reads/Sec	1537	1391	376	135
WriteRows/Sec	914	94.9	9	1.92

# Columns對於效能影響

- HBase為columns-oriented架構
  - 寫入資料必需“逐行”寫入
- 僅量透過Columns family歸納columns
  - 每一個Columns family會儲存到一個檔案
- 資料列結構設計，需視資料特性進行設計
  - 避免設計過多的欄位
  - 僅可能合併資料到一個欄位內

# HBase table design-#4

tbCorporation: 公司基本資料

	CorporationInfo			CarInfo			
row key	Address	PostNum	Telephone	CarId-1	CarId-2	...	CarId-n
逢甲大學	台中市西屯區...	407	04-24516669	<data>	<data>	<data>	<data>
台塑貨運	台北市文山區...	100	02-1234567	<data>	<data>	<data>	<data>
新竹貨運	台北市中山區...	101	02-2345678	<data>	<data>	<data>	<data>
大榮貨運	台北市士林區...	102	02-3456789	<data>	<data>	<data>	<data>

# HBase table design-#5

tbCarInfo: 車輛基本資料

	BaseInfo			LastInfo	
row key	Car_licence	Car_tel	Corp_id	Dtime	Position
OG697373	160-ZS	0970256521	SH	2010/11/29 08:00	<data>
OG697374	FV-531	0926289431	SH	2010/11/29 08:00	<data>
OG697375	FY-579	0926289644	SH	2010/11/29 08:00	<data>

# HBase table design-#6

tbCarLog: 車輛行車資料

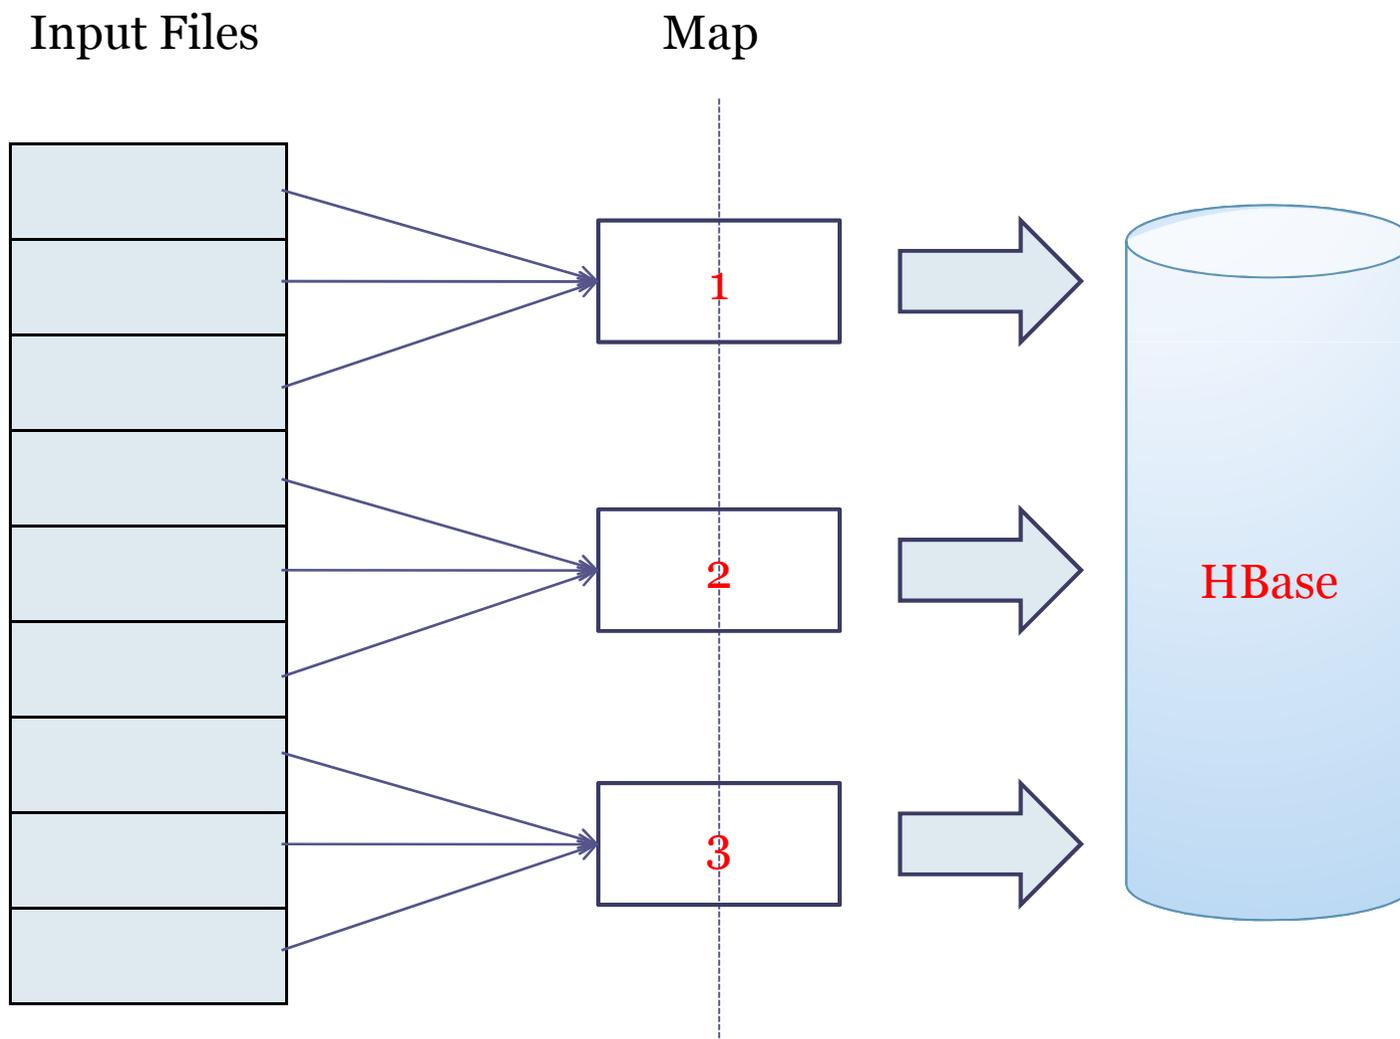
	Log
row key	Data
OG697373_2010-11-29_08:00:00	OG697373, 2010-11-29 08:00:00, ,193018, 2488885,False,1,91,,0,0,0,0,1,,184263014,407458,0
OG697373_2010-11-29_08:01:00	OG697373, 2010-11-29 08:01:00, ,193018, 2488885,False,1,91,,0,0,0,0,1,,184263014,407458,0
OG697373_2010-11-29_08:02:00	OG697373, 2010-11-29 08:02:00, ,193018, 2488885,False,1,91,,0,0,0,0,1,,184263014,407458,0

# 歷史行車記錄匯入

- 如何將長年累積的行車記錄匯入至HBase
  - 目前備份檔採用 **flat text**格式儲存
  - 每個月產生**1**個檔案
  - 每個月平均檔案約 **100~200MB**
  - 總共有約**100**個檔案

# 歷史行車記錄匯入

- 利用 Hadoop MapReduce 進行資料匯入



# 資料匯入情形

- 匯入1億筆行車記錄
- 原始資料約13.6G
- 花費時間
  - 6hr,3mins,5sec
  - 每秒寫入4,712筆

## Hadoop job\_201011121703\_0012 on cloud-a

User: hadoop

Job Name: SkyEye

Job File: [hdfs://cloud-a:9000/opt/hadoop-data/mapred/system/job\\_201011121703\\_0012/job.xml](hdfs://cloud-a:9000/opt/hadoop-data/mapred/system/job_201011121703_0012/job.xml)

Job Setup: [Successful](#)

Status: Succeeded

Started at: Thu Nov 18 10:28:40 CST 2010

Finished at: Thu Nov 18 16:31:46 CST 2010

Finished in: 6hrs, 3mins, 5sec

Job Cleanup: [Successful](#)

Kind	% Complete	Num Tasks	Pending	Running	Complete	Killed	<a href="#">Failed/Killed Task Attempts</a>
<a href="#">map</a>	100.00%	243	0	0	243	0	0 / 3
<a href="#">reduce</a>	100.00%	0	0	0	0	0	0 / 0

	Counter	Map	Reduce	Total
Job Counters	Launched map tasks	0	0	246
	Data-local map tasks	0	0	246
FileSystemCounters	HDFS_BYTES_READ	14,632,152,164	0	14,632,152,164
Map-Reduce Framework	Map input records	102,638,374	0	102,638,374
	Spilled Records	0	0	0
	Map input bytes	14,631,529,419	0	14,631,529,419
	Map output records	0	0	0

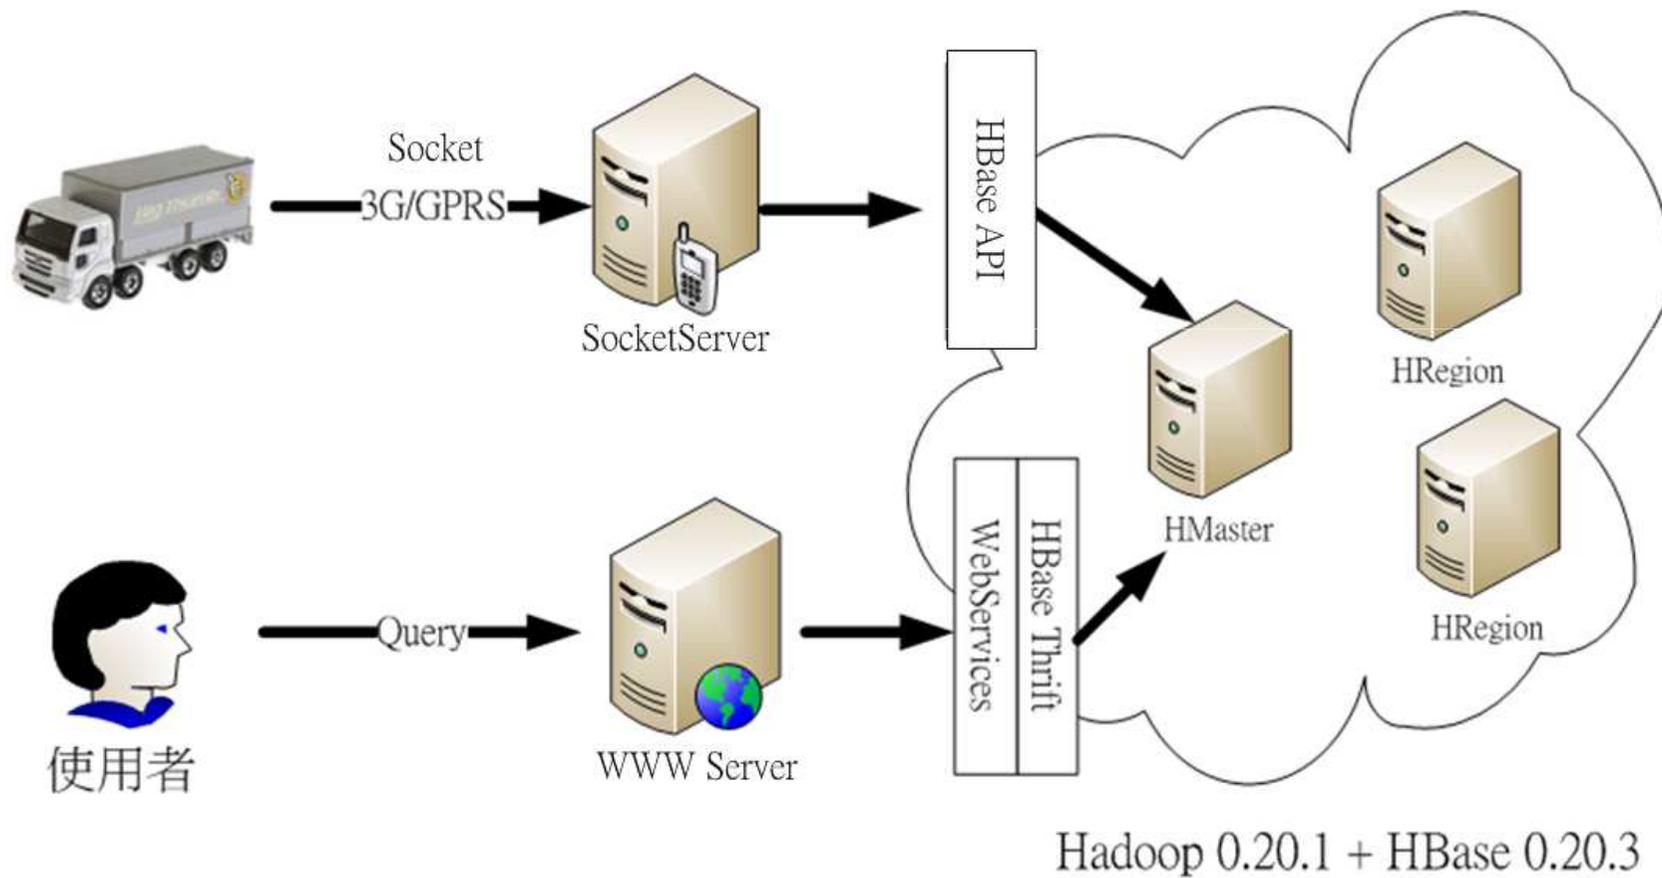
# 系統整合

資料讀取

資料儲存

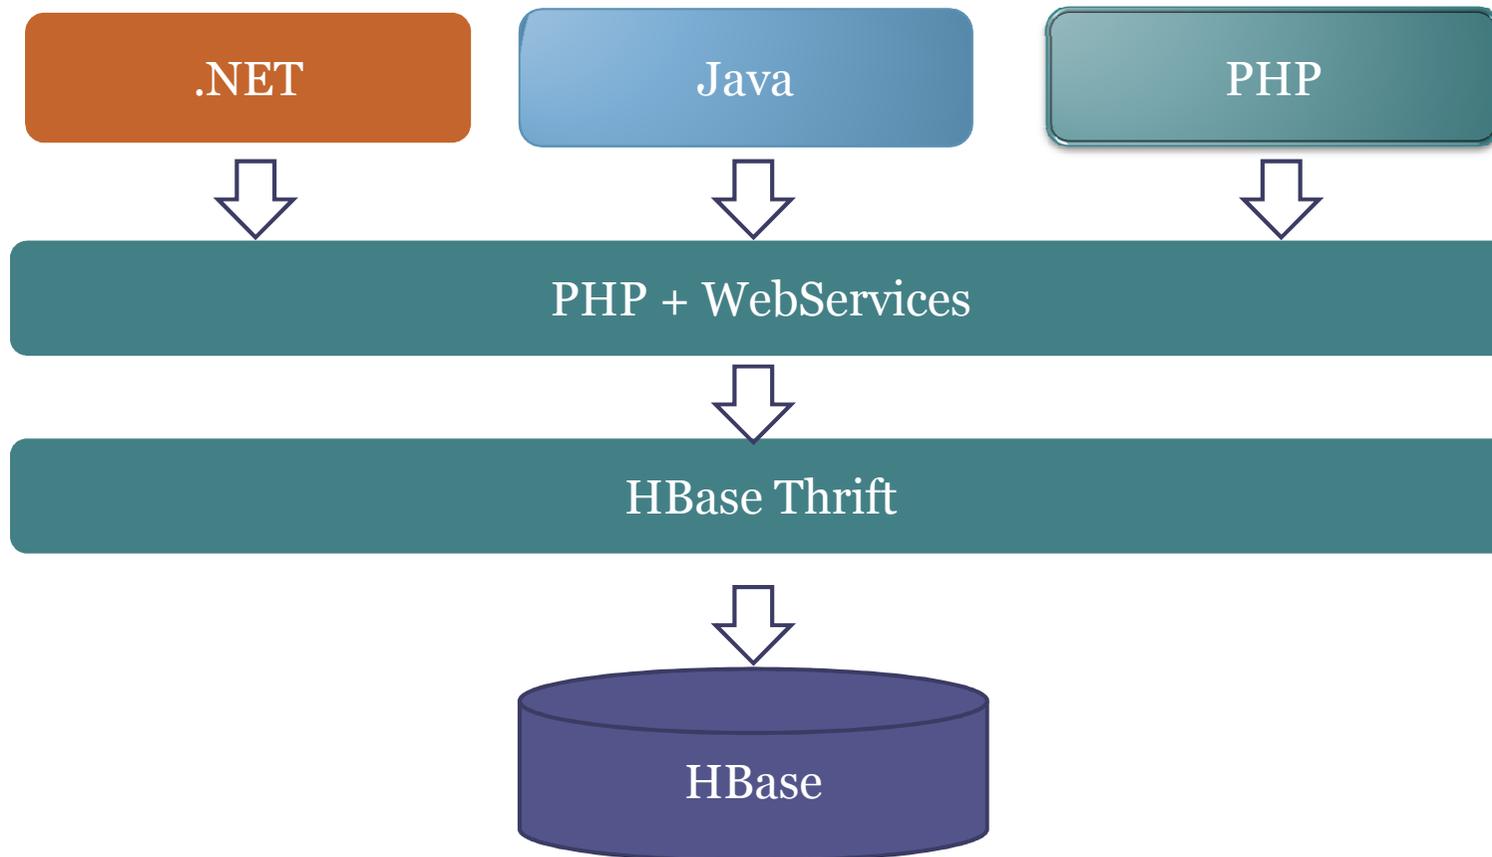
展示系統整合

# 系統整合架構



# 如何提供異質系統讀取資料

- 既有系統為 Windows 平台
- 利用WebServices提供資料服務



# HBase WebServices

- 提供歷史資料查詢

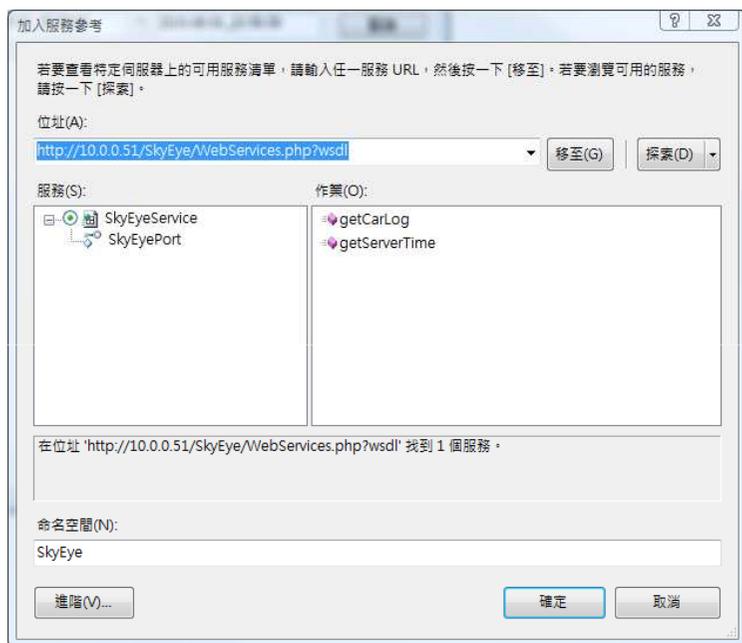


```
<?xml version="1.0" ?>
- <definitions name="SkyEye" targetNamespace="urn:SkyEye" xmlns:wSDL="http://schemas.xmlsoap.org/wSDL/"
  xmlns:soap="http://schemas.xmlsoap.org/wSDL/soap/" xmlns:tns="urn:SkyEye"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:SOAP-
  ENC="http://schemas.xmlsoap.org/soap/encoding/" xmlns="http://schemas.xmlsoap.org/wSDL/">
  <types xmlns="http://schemas.xmlsoap.org/wSDL/" />
  - <message name="getCarLogRequest">
    <part name="carid" type="xsd:string" />
    <part name="startDate" type="xsd:string" />
    <part name="endDate" type="xsd:string" />
  </message>
  - <message name="getCarLogResponse">
    <part name="output" type="xsd:string" />
  </message>
  <message name="getServerTimeRequest" />
+ <message name="getServerTimeResponse">
+ <portType name="SkyEyePort">
+ <binding name="SkyEyeBinding" type="tns:SkyEyePort">
+ <service name="SkyEyeService">
</definitions>
```

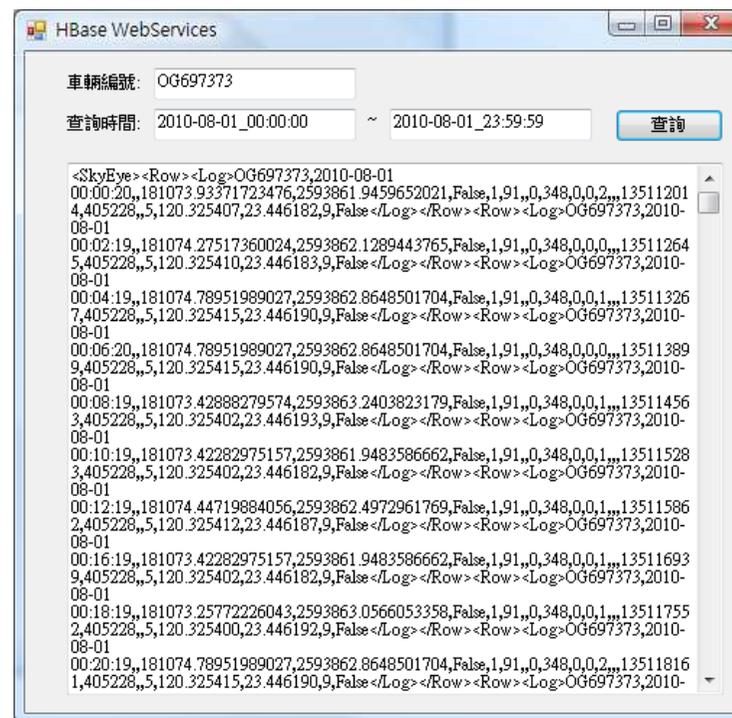
# .NET 整合範例

## Step2 :程式中呼叫WebServices取得歷史資料

```
private void btnQuery_Click(object sender, EventArgs e)
{
    SkyEyeService pService = new SkyEyeService();
    textBox1.Text =
        pService.getCarLog(
            txtCarID.Text, txtStart.Text, txtStop.Text);
}
```



**Step 1:**  
在 VS 2008 專案中加入服務參考



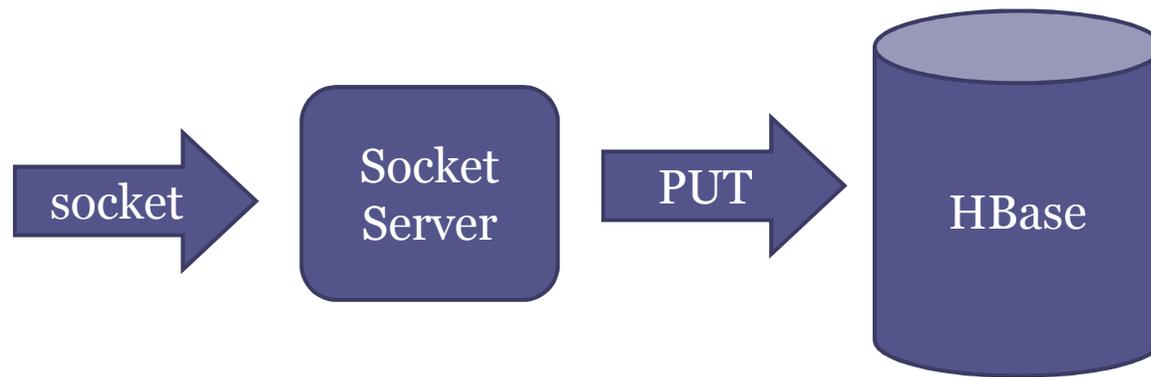
**Step3 :取得歷史資料**

# 車機資料如何儲存至HBase

- 透過HBase Java API
- 撰寫Socket Server接收車機回傳資料

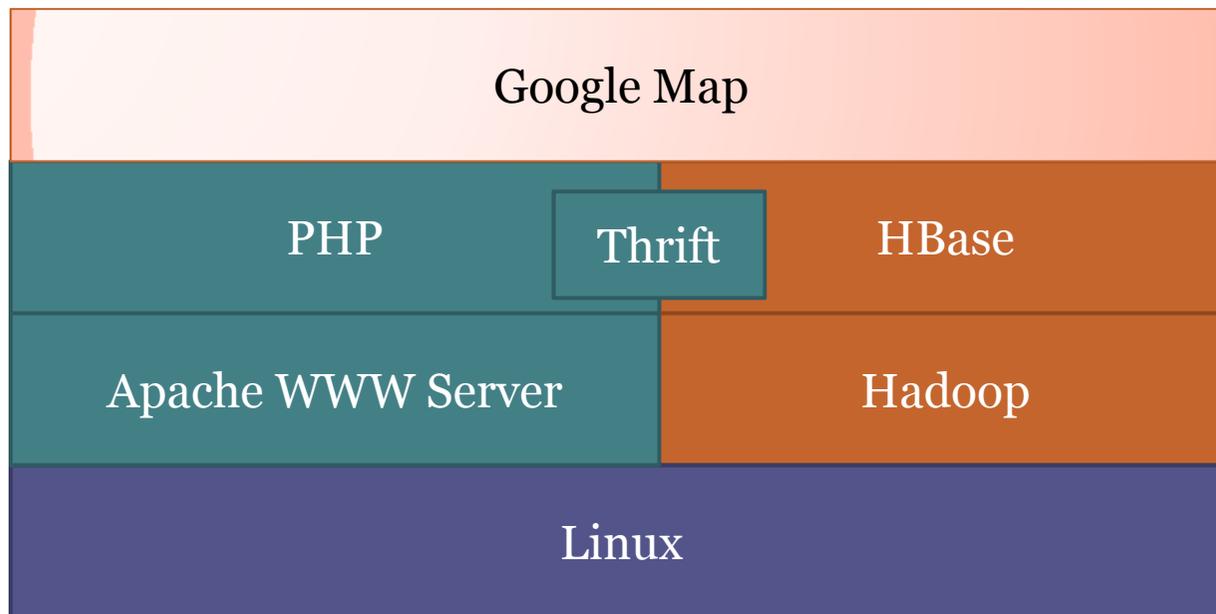


Socket client



# 前端展示程式整合

- Google Map
- PHP + HBase thrift





# 系統整合畫面

## Corporation List

Corporation	PostNum	Address	Function
台塑貨運	116	臺北市文山區木柵路三段220號	<a href="#">Edit</a> <a href="#">ShowCars</a>
大榮貨運	106	臺北市大安區新生南路2段86號	<a href="#">Edit</a> <a href="#">ShowCars</a>

[return to mainpage](#)

[close or Esc Key](#)

CarInfo:OG698001  
CarInfo:OG698002  
CarInfo:OG698003  
CarInfo:OG698004  
CarInfo:OG698005  
CarInfo:OG698006  
CarInfo:OG698007  
CarInfo:OG698008

[close or Esc Key](#)

台塑貨運

PostNum	116
Address	臺北市文山區木柵路三段220號
<input type="button" value="Save"/>	

## 結論 - 整體效益

- 解決巨量資料儲存問題
- **Hadoop**及**HBase**建置費用低
- 具延展性，擴充性高
- 分散式架構，降低單一伺服器負荷
- 支援 **MapReduce**，有助於資料分析

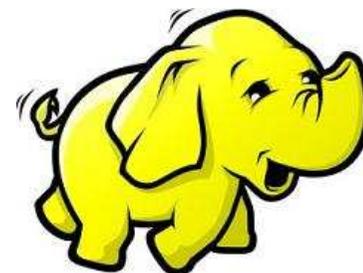
# 結論 - HBase的挑戰

- HBase 穩定性
  - 資料反覆進行大量寫入測試時，偶會造成HBase crash
- HBase安全性
  - 如何控管使用者權限
  - 資料備份機制
- HBase 管理性
  - 視覺化的管理工具

# 參考網站

- Hadoop官網
  - <http://hadoop.apache.org>
- Taiwan Hadoop Forum
  - <http://hadoop.nchc.org.tw/phpbb/>
- Hadoop技術論壇
  - <http://bbs.hadoopor.com/>
- 中國雲計算論壇
  - <http://bbs.chinacloud.cn/showforum-16.aspx>

簡報結束  
敬請指教



Hadoop Taiwan  
User Group

<http://www.hadoop.tw>