Information Management

DB2 Cloud Computing: Part #2

Mark Wilding (mwilding@ca.ibm.com)
(DB2 LUW Cloud Computing Architect)

What's Different About Being "in the Cloud"

Traditional	Cloud
Hardware is not always shared	**Shared Hardware** ▪ Can be noisy
Resources are physical	**Virtual Resources** ▪ Is it enough? ▪ More resources on demand
Standards are optional	**Standards**: Standards are highly prevalent ▪ Good for cost reduction ▪ Good and bad for security
Costs are fixed	**PAYG**: Costs are Pay-As-You-Go ▪ Tips and tricks for cost reduction

Non-Cloud

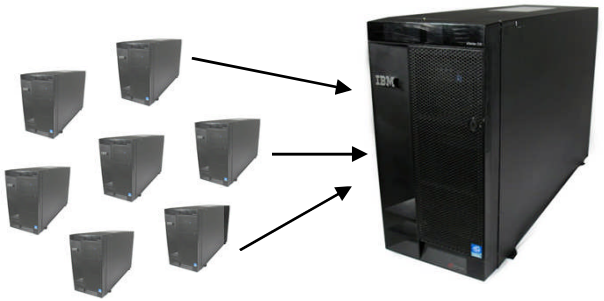

Cloud

DB2 Instance Configuration

- Two main storage types in the cloud
  - **Ephemeral** and/or VM based storage (non-shareable storage)
  - **Persistent** and shareable
- Persistent storage gives the most flexibility in terms of upgrades and movement between VMs.

All on Ephemeral (and boot from VM image)

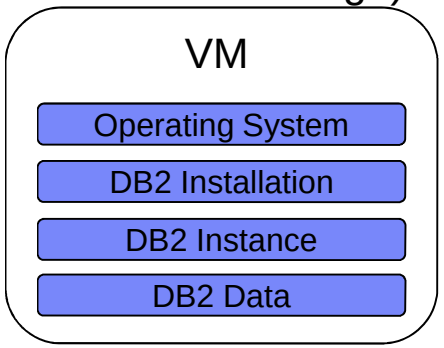

- Running is risky (what if the VM goes down or storage has problem?)
- Movement requires shipping of data (e.g., DB backup/restore)

Boot from VM Image

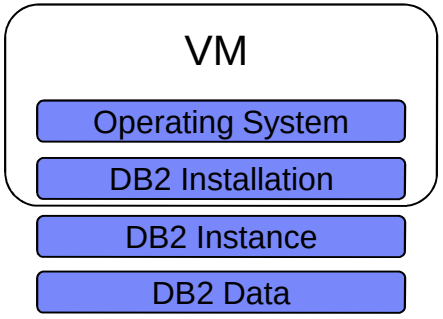

- Can move a DB2 instance and its data to another (e.g., larger) VM
- Upgrades require creating a new VM image

Boot from Persistent Storage

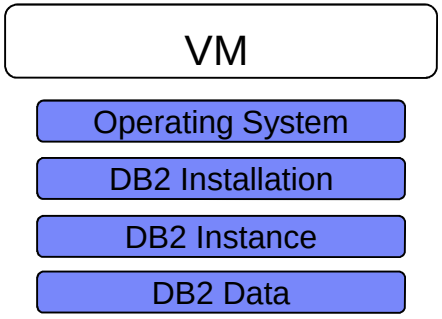

- Compute resources are completely anonymous
- Can move the OS and/or DB2 around the cloud as needed
- Upgrades can be implemented and tested on a clone of the OS/DB2 and then run in production

Cloud: A Potentially Noisy Environment

- For clouds based on server consolidation, you can typically monitor your CPU resources stolen using simple tools such as iostat

```
prompt> iostat
…
avg-cpu: %user %nice %system %iowait %steal %idle
       26.50 0.00  5.50    12.50   54.00  1.50
…
```

%steal is the amount of CPU time that the hypervisor is giving to other VM instances (and away from your VM instance).

- Shared resources can fluctuate over time and can cause workload variability

- Monitor shared resources to make sure you understand the impact to your database instance!

Important KPIs to Monitor in the Cloud

- DB2 monitoring works the same way in the cloud as in traditional environments
  - It is the real DB2 running in the cloud (applications and monitoring tools won't know that it is in the cloud)

- However, in a shared environment, the shared resources need to be monitored carefully:
  1. **CPU**: CPU capacity and saturation over time (including %steal)
  2. **I/O:** I/O bottlenecks latency and wait times
  3. **Network:** Network latency and wait times

- Recommendations:
  1. Metrics related to shared resources should be monitored over time
  2. Tuning should be optimized to make the best use of hardware (less reliance on hardware for performance is important)
  3. Leverage WLM (Workload Management) as a buffer against any potential noise

Enterprise: IT Maturity Trends (Intermission)

(100s or 1000s of DBs)

[redacted]

Automation

Automation reduces the costs of mundane tasks (deployment, backups, day-day ops, etc)

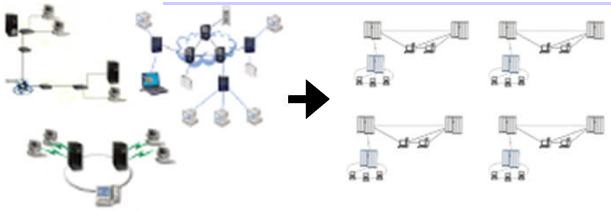

Standardization

Corporate IT standardization reduces complexity. It also reduces required skills and the number of unexpected problems. Building-block HW resources.

Consolidation (often using virtualization)

Costs reduced via virtualization and hardware sharing.

Bare Metal

"90% of the systems are 10% utilized"

Energy/cooling costs

CapEx costs

Operations costs

Security Considerations

- Cloud computing can be very good and very bad for security
- Good: Security best practices can be implemented and automated across the IT environment
- Bad: Any security exposures can be implemented and automated across the IT environment!
- Special considerations are certainly required!

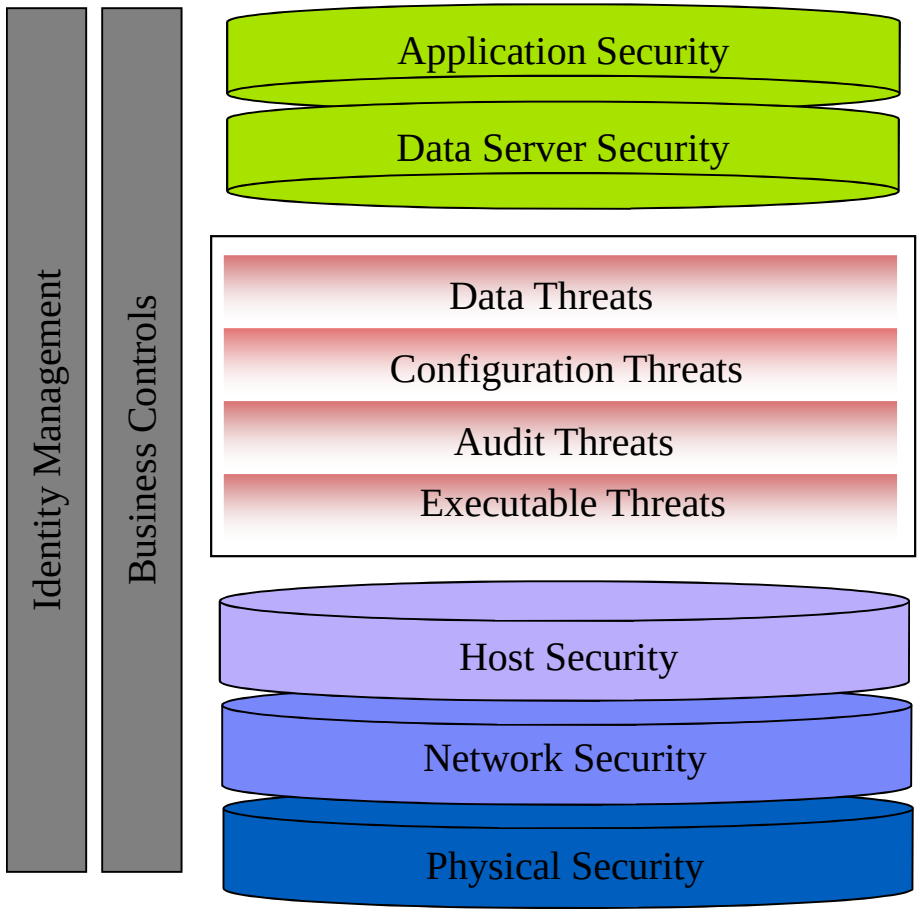

Security: Gartner's 7 Cloud Computing Security Risks
(http://www.networkworld.com/news/2008/070208-cloud.html?page=2)

1. **Privileged user access to data**
   - **From DBA:** Trusted contexts with roles and revoke DATAACCESS from the DBA
   - **From root:** IBM Database Encryption Expert (DEE)
   - **From snoopers:** DB2 has on-wire encryption via SSL
2. **Regulatory compliance**
   - DB2 has Authentication, Authorization, Auditing and Encryption
   - DB2 is EAL4+ certified under Common Criteria
3. **Data location**
   - IBM is a trusted name and has years of experience with hosted systems
4. **Data segregation**
   - DB2 can provide isolation at the instance level, database level, schema level or row level (LBAC)
5. **Recovery**
   - DB2 has HADR for disaster recovery which works both within a cloud and from on-prem to the cloud
6. **Investigative support**
   - DB2 auditing and data segregation methods both capture and limit the scope of investigative efforts
7. **Long-term viability**
   - IBM is a trusted company that has been in business for ~ 100 years

Example: Simple Security Standard

Checklist items	Yes/No
VM does not have any unnecessary software installed or running on it	
There are no active listening ports other than what is required (lsof or netstat on Linux can show this)	
Each VM uses its own firewall to lock down ports to the minimum required set (see previous). No response is used for blocked ports.	
The DB is in a security group that prevents outside access from the raw Internet _or_ the VM firewall only allows connections from the application's IP addresses	
Passwords are randomized and set new for each deployment	
A procedure and policy are in place to roll out a major security patch quickly and to all affected VMs	
Data access is revoked from the DBA if the DBA is not you	
SSL is used for over-the-wire communications, including SQL	
A procedure and policy are in place to audit the VM and DB access and the audit records are sent off of the VM.	

Two Examples with Google Words and Links…

Revoke DBA Access

- **Google key words/examples::**
  - **DB2 REVOKE DATAACCESS ON DATABASE FROM** <user/group/role>
- **URLs**
  - http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.sql.ref.doc/doc/r0000981.html

Configuring Secure Sockets Layer (SSL) support in a DB2 instance

- **Google key words/examples:**
  - **gsk8capicmd -keydb -create -db "mydbserver.kdb" -pw "passw0rd" -stash**
  - db2 update dbm cfg using SSL_SVR_KEYDB D:/temp/gsk/mydbserver.kdb
- URLs
  - http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.admin/db2z_configssl4serv.htm
  - http://www.ibm.com/developerworks/data/library/techarticle/dm-0806sogalad/index.html

Getting Started

- General Questions
  - mwilding@ca.ibm.com
- RightScale
  - http://support.rightscale.com/27-Partners/IBM_DB2
- Amazon
  - http://www.ibm.com/developerworks/downloads/im/udb/ec2.html
- IBM Compute Cloud
  - https://www-949.ibm.com/cloud/developer/dashboard
- WebSphere CloudBurst Appliance
  - http://www-01.ibm.com/software/webservers/cloudburst
- IBM CloudBurst
  - http://www-01.ibm.com/software/tivoli/products/cloudburst
- Optim Data Privacy
  - http://www-01.ibm.com/software/data/optim/core/data-privacy-solution