

# “Azure? A lap around Cloud-Hosted Services.” Ingo Rammer

May 22, 2009

[MS Techdays 2009] Summary “Azure? A lap around Cloud-Hosted Services.” Ingo Rammer

## What is the cloud about?

Currently we have desktop and server OS's optimized for single machines and a bit of clustering. Once you leave the safe grounds of this one machine, support for transparently consuming resources somewhere in the cloud deteriorates very quickly.

Cloud applications are about:

- Large-scale failover capabilities
- Geographic dispersion (multi-country)
- Deployment support
- Patch Management
- Capacity Management and large-scale clustering

MS Azure tries to answer the move to cloud computing with virtualized computation, scalable storage, transparent management and development support. The concept is that applications will be build on top of a set of services provided by the Azure platform that will support the use of resources in the cloud like they were installed on one machine.



## Development and Deployment

**Azure Application Development** Development for the Azure platform is done on a development fabric simulating the cloud circumstances. It is just way to complex for the moment to debug Azure applications at

run-time in the cloud. Azure provides development storage for the same reasons.

It is a good idea to download the Azure Log Viewer and the Azure Storage Explorer, both third party, when you start developing: Azure Log Viewer can be found at <http://blogs.itmentors.com/bill/2009/02/10/windows-azure-log-viewer/> and Azure Storage Explorer can be located under <http://blogs.itmentors.com/bill/2009/02/10/windows-azure-log-viewer/>

### **Azure Application Deployment**


Applications for Azure are deployed through the Azure Services a portal hosted by MS. You upload your developed application and then the portal allows for zero-downtime upgrades or switch-deployments to bring your updated application into production. This means development and deployment for Azure consist of create-publish cycles.

Under the bonnet virtual machines are created to host your application.

System engineers are provided with status information about the application's health through e-mail, instant messaging or phone calls.


### Scalability and Availability

New to Azure applications is that any request is handled first by a load balancer, build-in in the platform, and then is forwarded to an instance of your application.



Considering scalability and availability we see that Azure really takes things further. We are not talking anymore about just a second server to have a local failover or a few machines in a datacenter to handle the load. No, we're talking about thousands of servers in several datacenters.

The platform allows for turning-the-knob style up- and down-scaling of server capacity. It also supports dynamic capacity schemas by means of follow-the-sun server capacity. You only will have to pay for consumed resources and not for unused capacity required to handle potential peak-loads.



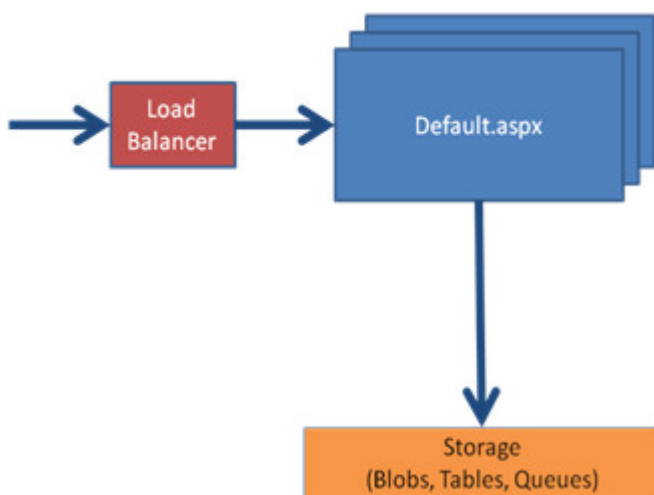
## Storage

Because Azure takes control over your server instances there is no such a thing as sessions state like we are used to in, for example, IIS hosted applications. If you want some data to live longer than one request you need to use Azure Storage. All resources are logical in Azure. The platform maps them transparently to real physical resources.

Azure provides storage accessible to all instances. However there is something like Instance Storage, this should never be used to store session related data. It is just a local cache for transient data.

Azure Storage comes in three tastes ?tables (some TBs) , queues (unlimited number of messages if smaller than 8 kB ) and blobs (up to 50 GB )-and is accessible through a REST or ADO .NET Data Services interface. Data is durably stored and at least three times replicated to guarantee it is never lost.

The logical model known to the application is very simple:



For the moment there are some limitations concerning the Azure Storage. Queues only support polling instead of event based access. Applications require full trust instead off medium trust.

Azure Storage can be accessed using two roles: a **web role** to support synchronous HTTP requests and a **worker role** for asynchronous queued requests. To avoid remarks like vender-lock-in about the storage, MS

made the storage accessible from outside the platform to allow you to copy data locally.

### What is showing at the horizon's

The future will bring following enhancements to Azure:

- More types of roles ("RAW" !)
- Mangement rules, monitoring, dynamic deployment
- Geographic isolation, fault domains, update domains: some homeland security acts do not allow specific data to leave the country so the storage distribution must be done over datacenters within the borders.
- SQL Data Services

---

[www.vanderbist.com](http://www.vanderbist.com)