# **Global** Azure BOOTCAMP







#### Designing for scalability and high availability on Microsoft Azure

Vaggelis Kappas

Premier Field Engineer - Microsoft <u>http://autoexec.gr</u> | vaggeliskappas



Session objectives and takeaways Session objective(s):

1. Design highly available applications on the Azure platform

2. Understand best practices and common pitfalls to avoid

3. Learn the inner workings of Azure and how to best work around hidden limitations

Note: This session focuses mostly on the design of Cloud-based apps using IaaS.

# High availability

### What does HA mean to you?

- What are the business requirements?
- Different requirements mean a different architecture







#### **Best For:**

Data deletionProtection for unplanned failuresMission critical appsData corruptionDon't want to re-architect for HANew appsLegal, governance & complianceDon't want the cost of HALocalized failuresLarge scale failuresLarge scale failuresLocalized failures

Mean time between failure (MTBF) Metric used to identify reliability Component failures Cost of delivering cloud infrastructure On-premises vs cloud Hardware failure is inevitable



Designing for the cloud No single point of failure State stays at the edges of your app stack Loosely couple your components Build for scale Process centrally, deliver locally Automate Trust and verify Scale out, not up



# Scaling vertical (up) vs horizontal (out)

- Vertical
- Horizontal
- Cost

### VMSS and autoscale





# Good candidates for Azure

Built versus bought software Software designed for cloud

Applications designed around a single-state datastore







Networking

Storage

All 3



# Compute

## Availability Sets

Governs fault and update domains Tied to a role in your application Required for 99.95% SLA from Microsoft Provides fault domains and upgrade domains



Availability Sets (upgrade domain) Fault Domain is single point of failure Upgrade domain used for patching

ASM (2 FDs)			ARM (3 FDs)		
VM	Update Domain	Fault Domain	VM	Update Domain	Fault Domain
VM1	0	0	VM1	0	0
VM2	1	1	VM2	1	1
VM3	2	0	VM3	2	2
VM4	3	1	VM4	3	0
VM5	4	0	VM5	4	1
				E	2

# Deployment scale



### Azure VMs



**{}** Resource Manager Template

Virtual Network

### Azure VMs



Virtual Network \_\_\_\_\_\_ Availability Set \_\_\_\_\_ Subnet \_\_\_\_\_



Resource Manager Template



# Virtual Machine Scale Sets (VMSS)

- A set of VMs which are identical (cattle)
- Deployed with configuration management techniques
- Can be auto-scaled
- Based off a common base image Portal / ARM template Configuration management



# Deployment scale - VMSS





# Network

# Layer 4 vs layer 7 Load balancers

Туре	Azure Load Balancer	Application Gateway
Protocols	UDP/TCP	HTTP, HTTPS, and WebSockets
Load balancing mode	5-tuple(source IP, source port, destination IP, destination port, protocol type)	Round Robin Routing based on URL
SSL offloading	Not supported	Supported



Availability Set

# Demo Stages of availability

# What do you do if a region fails?

a) Use a Load Balancer to shift to additional resources

b) Use Azure DNS to point to a paired region

c) Use Traffic Manager to change regions

d) Initiate your failover runbook in Azure Site Recovery
 e) Complain to @Azure on Twitter

#### How Traffic Manager handles endpoint loss **Update DNS** 'ou there?) </1 **GFT** No Rezoonské 10s 20s 10s 20s 10s 20s 10s 20s 10s 20s 10s Time-

Traffic Manager

DNS-based load balancing and failover Health probe of endpoints to determine availability Impact of Time To Live (TTL)

Multiple performance policies

- Performance
- Priority
- Weighted









### **Types of storage**

• Locally redundant (LRS) - 3 local copies strongly consistent



• Read access geo-redundant (RA-GRS) – Same as GRS but with read-only access of eventually consistent data

Provides tables, blobs, queues, and files

#### They provide:

- Isolation
- Replication
- Region order
   recovery
- Sequential updates
- Data residency

Primary	Secondary	Primary	Secondary
North Central US	South Central US	East Asia	South East Asia
South Central US	North Central US	East China	North China
East US	West US	North China	East China
West US	East US	Japan East	Japan West
West US 2	West Central US	Japan West	Japan East
West Central US	West US 2	Brazil South	South Central US
US East 2	Central US	Australia East	Australia Southeast
Central US	US East 2	Australia	Australia East
North Europe	West Europe	Southeast	
West Europe	North Europe	Canada Central	Canada East
South East Asia	East Asia	Canada East	Canada Central
Germany Northeast	Germany Central	UK South	UK West
Germany Central	Germany Northeast	UK West	UK South

https://docs.microsoft.com/en-us/azure/best-practices-availability-paired-regions

Storage – hidden point of failure? Standard vs Premium

Disks live in a storage account A storage account exists on a storage stamp A storage stamp is a single point of failure Storage outage can span fault domains



# Demo: Storage account information tool

# Storage account information <u>http://aka.ms/StorageAccountInfo</u>\*

Х

🔿 🧟 http://storageinfo.w... 🔎 🗕 🖒 😂 Easy Azure Storage Info ... 🗙

This page is designed to test the Azure Storage Info API. Using this tool you can get information about your Azure Storage accounts.

Storage Account Na	me storagetestblob1	ŀ	Public Cloud	$\checkmark$
Get Storage Info				

#### **Results:**

Storage Account UF	RL: storagetestblob1.blob.core.windows.net
Region:	West US
Datacenter ID:	4
Stage:	Production
Туре:	Standard
Stamp:	07
Version:	Primary

\*This tool is not endorsed, supported, or in any way officially endorsed by Microsoft. Use at your own risk. https://storageinfo.azurewebsites.net/api/StorageInfo

**POST to (Yes, it's an Azure Function):** 

#### **Post Content:**

{"storageUrl" : "storagetestblob1.blob.core.windows.net"}

#### **Response (in JSON):**

```
"Region":"West US",
"Datacenter":"4",
"Stage":"Production",
"Type":"Standard",
"Stamp":"07",
"Version":"Primary"
```

Send feedback to aglick@microsoft.com

# Common architectures

Consistency models Strong consistency Session consistency Eventual consistency

Consistency Availability Partition Tolerant == CAP Theorem

# Common single-region deployment- RDBMS



Multi-region deployment Active/Passive Only one region online with DR failover

Active/Active

Round robin (all regions accept traffic)

Active performance Many regions, user is routed to nearest





# Multi-region deployment (active/active)



West EU (Primary)

North EU (Secondary)

# Using a CDN for availability

### Protection from DDoS

Greater availability and scale for your origin

### Other benefits:

Lower latency content delivery Faster data delivery Lower data usage Origin obfuscation



# Demo: Azure CDN

# Chaos Engineering - Fault Injection

- Chaos Monkey
- Chaos Dingo
- Tests for single point of failure
- Automated architecture testing
- Simulates failure



Session objectives and takeaways Session objective(s):

1. Design highly available applications on the Azure platform

2. Understand best practices and common pitfalls to avoid

3. Learn the inner workings of Azure and how to best work around hidden limitations

Note: This session focuses mostly on the design of Cloud-based apps using IaaS.



### Azure High Availability

High Availability Checklist – <u>https://aka.ms/AzureHaChecklist</u>

Designing for High Availability – <u>https://azure.microsoft.com/en-us/documentation/articles/resiliency-high-availability-azure-applications/</u>

Overview of HA and DR - <u>https://azure.microsoft.com/en-us/documentation/articles/resiliency-disaster-recovery-high-availability-azure-applications/</u>

# High availability checklist

## Use Traffic Manager

- Avoid single VMs
- Use load balancers in front of web-facing VMs
- **Put your stateless servers in Availability Sets**
- Use VMSS for your stateless server scaling
- Vse Premium Storage for your production VMs
- Use internal load balancers (or queues) between tiers
- Distribute your database
- Use caches
- Contact support before a high scale event Store static assets in Blob Storage Use a CDN in front of your static assets





# https://aka.ms/cc9cf1

#### Thank You

