



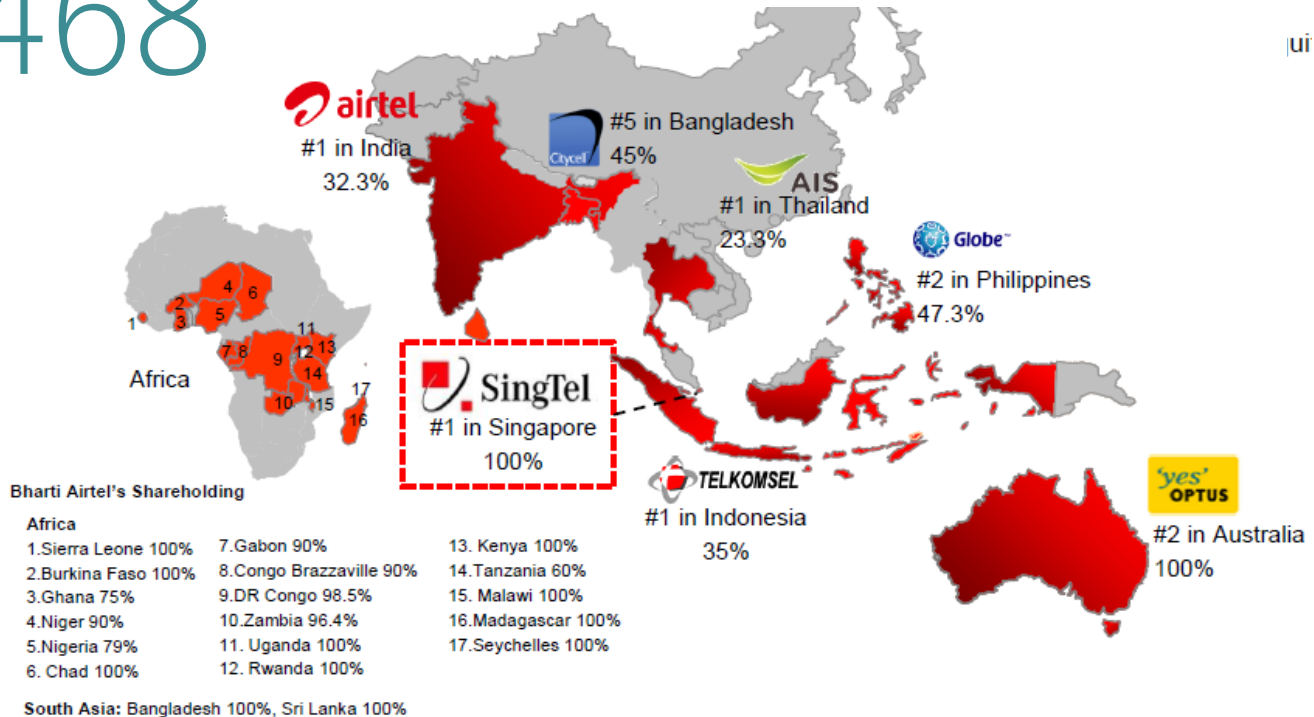
# Software Defined Storage and Private Clouds – Deploying at Scale

Jim Pinkerton  
Partner Architect  
Microsoft

# SingTel Group Overview

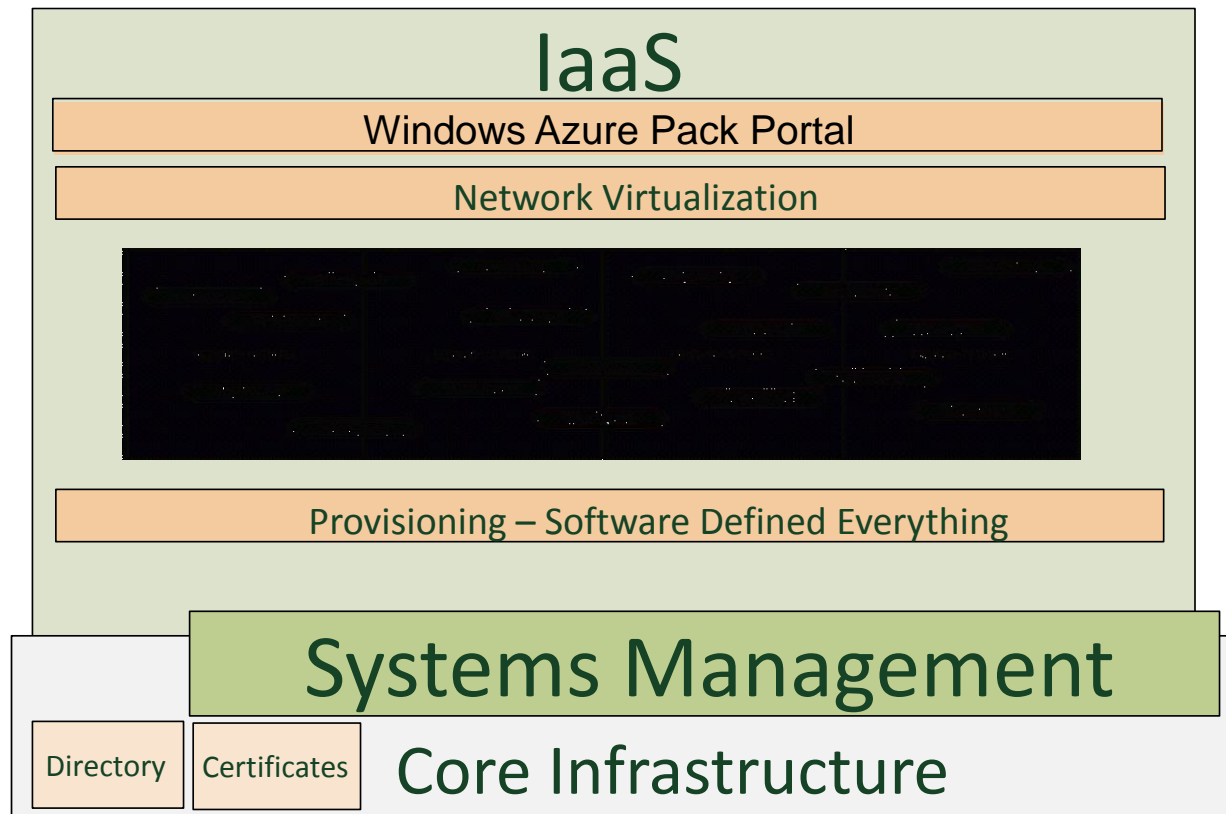
- Asia's leading communications group Providing a wide spectrum of multimedia and ICT solutions.
  - Including voice, data and video services over fixed and wireless platforms. The Group has presence in Asia and Africa

468 million mobile customers in 25 countries.

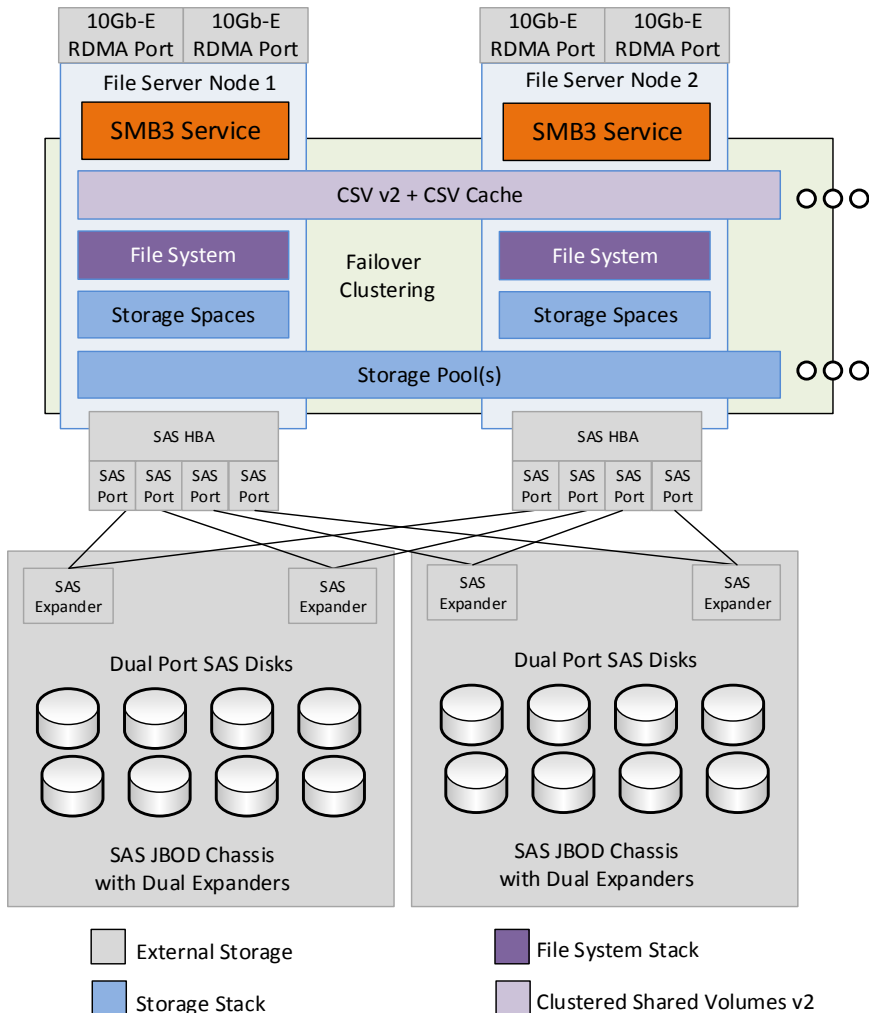


# SingTel Service Provider Cloud OS

- SingTel Service Provider Cloud OS Platform is the start of a journey
  - Initially Infrastructure as a Service (IaaS)



# Software Defined Storage Architecture



- All Storage in the solution is from a file server using SMB3 with RDMA
  - Dual 10 GbE per server
  - Each server in the file server is active-active
- Software Defined Storage Stack
  - SMB3 RDMA access protocol
  - Clustered Shared Volumes
  - **Storage Spaces with Tiering**
  - JBOD Enclosures
    - Mix of SSD and HDD
- Fault Tolerance is critical
  - Network
  - Server
  - HBA/SAS links
  - JBOD
  - Disk
  - Power
- Based on Windows Server 2012 R2

# VM Boot Storm Workload

## - The value of flash -

- **Goal: Reliably boot ~1800 VMs**
  - Understand number of VMs that can be deployed at one time
  - Cold/first boot performance analysis
  - Post Tiering optimization
- **Design Goal: Size the storage subsystem**
  - **Pre-Tiering and Post-Tiering**
    - 200 GB SSD
    - 800 GB SSD

# Flash Tiering Results - VM Boot Storm

Test Configuration	Test Run	# of VMs created	# of VMs booted	Per-VM Boot Time in seconds			
				Min	Average	Max (90 <sup>th</sup> Perc)	Max
800GB SSD	Baseline (Pre-tier) run	1792	1792	17	173.4	462	811
	Post-tier run	1792	1792	17	21.6	24	24
200GB SSD	Baseline (Pre-tier) run	1792	1468	23	583	996	1629
	Post-tier run	1792	1480	83	499	672	827

**After balancing the tiers:**

- 200GB SSD Setup: ~320 VMs failed to boot
- 800GB SSD Setup: All VMs booted


**Takeaway:** Full rack **can simultaneously boot ~1800 VMs** in cold-boot power recycle scenario

**After balancing the tiers:**

- 800GB SSDs **reduce average boot times by 23x**
- Takeaway:** Tiering **accelerates** subsequent VM boots

# The Money Slide...

Pre and Post Tier VM Boot Times Comparison (800GB )



**After balancing the tiers:**

- 800GB SSDs **reduce average boot times by 23x**
- Average of 173 seconds reduced to 22 seconds
- Maximum of 811 seconds reduced to 24 seconds

- Including SSD in Software Defined Storage makes a huge difference
- SSD for IOPs balanced with HDD for capacity is critical for a cost effective SDN Solution
- Correctly sizing the SSD tier using real world workload analysis is critical