*Unity
Pierluca Chiodelli
Architecture and Overview - Kaushik Ghosh
New Unity All Flash family (x50F)
Flexible deployment options (purpose-built, converged, virtual)
Virtualisation has helped Unity development tremendously

All the software now included
Block, File or VVol
Snapshots and AppSync Basic
Replication (including RecoverPoint Basic)
Inline compression
D@RE
AV enabler
QoS
Cloud Tiering
Unisphere

Architecture
Active-active fully redundant dual node architecture
unified design: file, block or VVol, sharing the same pool of storage
10 minutes to install, 30 minutes to production
Compact and powerful: cloud integrated 500TB all-flash in 2RU

Architected for all flash
3D TLC NAND flash drive (1wpd > 0.5wpd) for all IO type
Multi-core optimised for best CPU utilisation and low latency
Automatic flash wear balance
Zero impact drive firmware based garbage collection
Per-object in memory log for consistently low response time
Write coalescing with full stripe writes to minimise IO
Inline compression (done in memory)
Mix different flash drives and capacities for lowest cost

Delivering consistent response time and more
1. Write to in memory log (one per volume)
2. Coalesce and fold writes
3. Compress in memory
4. Full stripe write to drive
5. Balance wear and IOPS

Consistent response time
Lower flash consumption
Longer flash endurance
Lower $/GB and $/IOPS

Prioritised IO with Quality of Service
Different service levels by tenant / app

Instantaneous response
Separate limits for each snap and LUN
Flexibility with new burst support
Supports all client types

Flash density

Dynamic Unified Pool - Flexible Drive Configurations
Minimum of 6 drives
Unified pool (block, file, VVols)
Flexibility
Low entry price
Lower cost of ownership
Expand with as few as 1 drive
RAID of the pools represented as R6 or R5, but different under the hood (done in software - this gets decided by the system depending on the configuration)
Spares are now built into the pool - rebuilds are a lot faster

u64 filesystem
256TB scale-up file system
inline compression
file system shrink - reclaim free blocks and shrink file system to reflect new capacity
VMware integrated - full clones and space reserve through VAAI
Multi-tenant NAS servers - independent multi-tenant file stack with ability to set duplicate IP addresses
Cloud archiving and tiering - policy based transparent archival of files to public or private cloud (cloud tiering appliance)

Forever Snapshots: Flexible Low Cost Data Protection
Local -> Remote -> Cloud
Use S3 (cheaper than EC2/EBS)
Support Azure, Amazon, Virtustream, ECS (for file, block support coming)

Thin Clones: For Test, Dev, Analytics and More
Quickly create a fully populated copy of LUN
Shares same data blocks as original LUN (deduplicated)
Independent snapshot and replication topology
Enable separate QoS Limit on each clone and original LUN

Built-in Migration from VNX
FC, iSCSI, NFS and SMB migration from VNX arrays
Built-in solution: self-service without requiring any third-party tool
Migrates LUNs, filesystems, quotas, ACLs, and exports
Transparent to file applications, minimally disruptive for block

CloudIQ
single pane of glass to monitor multiple arrays across the world

access from anywhere - just need laptop / tablet and internet
cloud-based - no need to setup servers and VPN access
Pro-active analysis and support
Support planning and optimisation

HTML-5 Unisphere
VMware VASA and VAAI
REST API, SNMP & SMI-S
Operational Support

[demo - Wei Chen]
256 snapshots per object (volume)
LUNs thin by default (option to allocate if required)

Options for replication?
Support for SRM
Async replication (snap diff) same for file and block (down to 5 minutes)
Mirror capability (ack on both sides)
Looking to do a metro solution for file at the end of this year