



Critical Role of Flash in Guaranteeing Cloud Storage Quality of Service

Dave Wright

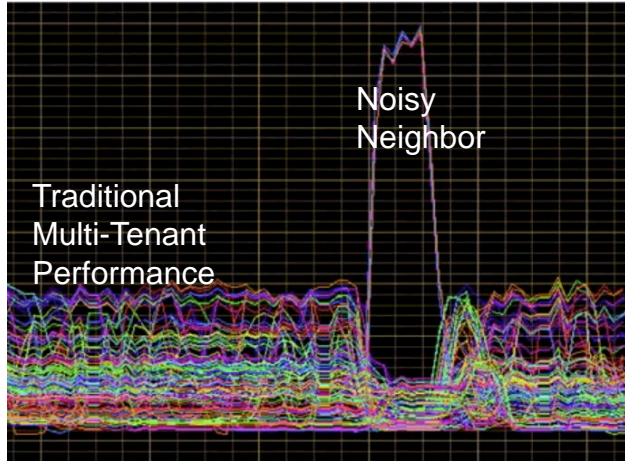
Founder & CEO

What is Guaranteed Quality of Service (QoS)?

- The ability to deliver consistent application performance regardless of system condition or application activity.

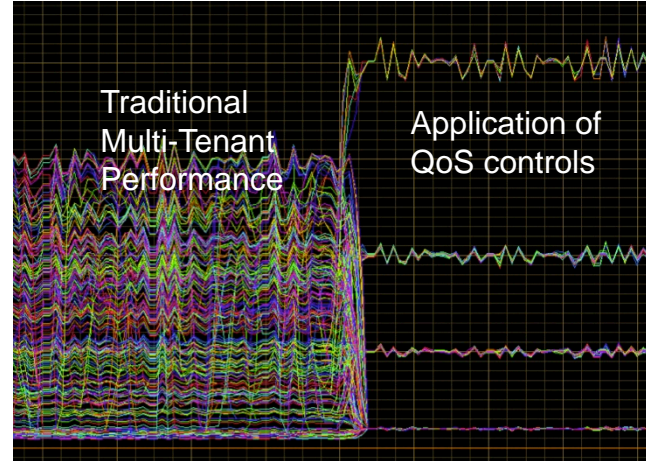


QoS Essential in Multi-Tenant Clouds



The Noisy Neighbor Effect

- Individual tenant impacts other applications
- Unsuitable for performance sensitive apps



Multi-tenant QoS in Practice

- Create fine-grained tiers of performance
- Application performance is isolated
- Performance SLAs enforced

Barriers to Guaranteed QoS in the Cloud

- Multi-application, multi-tenant IO contention: “The IO Blender”
- Dynamic infrastructure, rapid growth
- Hidden infrastructure, no user visibility of hardware & software failures
- Lack of global application visibility and control: “The Noisy Neighbor”
- Wide range of application performance requirements



Designing an architecture to guarantee QoS

Six key elements identified to address shortcomings in traditional storage architectures.

- All-SSD architecture
- True scale-out architecture
- RAID-less data protection
- Balanced load distribution
- Fine-grain QoS control
- Performance virtualization

Requirement #1: All-SSD architecture


Deliver consistent latency for every I/O



- Spinning disks can serve only a single IO at a time
- Deep queue depths can lead to highly variable latency
- Flash allows for multiple, parallel IOs
- Consistent, low latency for every IO


Requirement #2: Scale-out architecture

Linear, predictable performance gains as system scales




Requirement #3: RAID-less data protection

Predictable performance in any failure condition



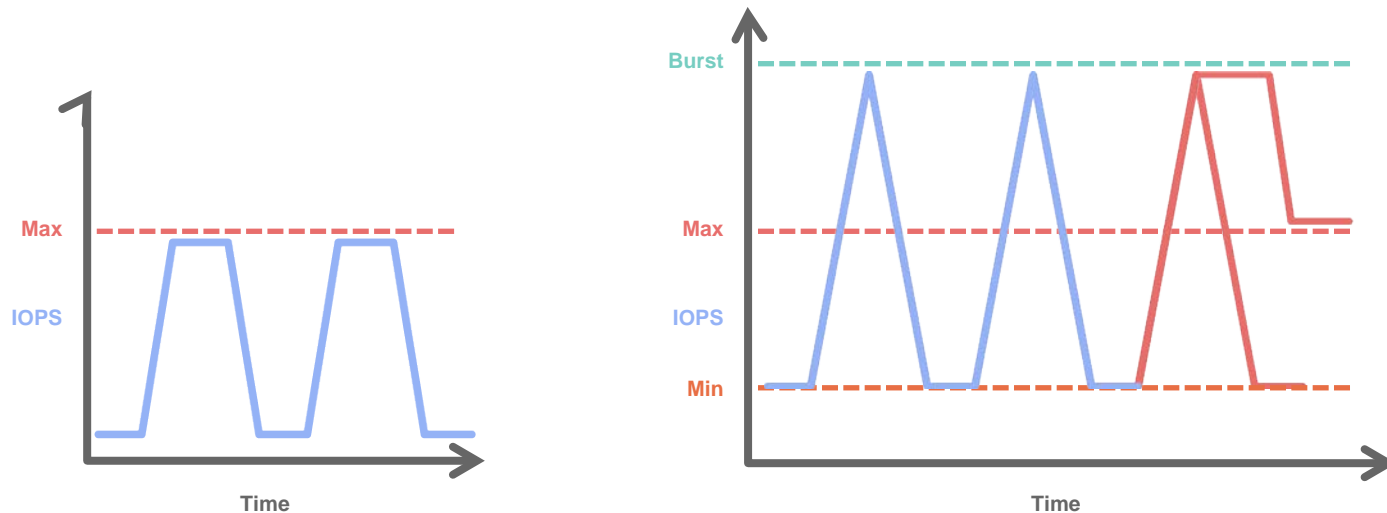
Requirement #4: Balanced load distribution

Eliminate hot spots that create unpredictable I/O latency




Requirement #5: Fine-grain QoS control

Completely eliminate noisy neighbors, and guarantee volume performance



Requirement #6: Performance virtualization

Deliver performance resources independent of capacity and on demand



Summary

- Flash is a key component in an architecture to deliver QoS to the cloud
- Disk and hybrid systems will never be able to deliver consistent performance
- However, delivering Guaranteed QoS requires a complete architecture, not just flash

