Tintri

MSPWorld
POWERED BY MSPAlliance

Bridging the Gap between Cloud Applications and Storage

#MSPWorld #MSPAlliance @Tintri

Your speaker…

Chuck Dubuque

Senior Director, Product Marketing at Tintri

ChuckD@tintri.com

@ChuckDubuque

Agenda

- What do CSP/MSPs need from storage?
- Why is storage such a challenge?
- Requirements of VM-aware Storage
- Technology solutions
- Conclusion

What do CSP/MSPs need from storage?

Tintri CSP Storage Survey

What is your role in the organization?

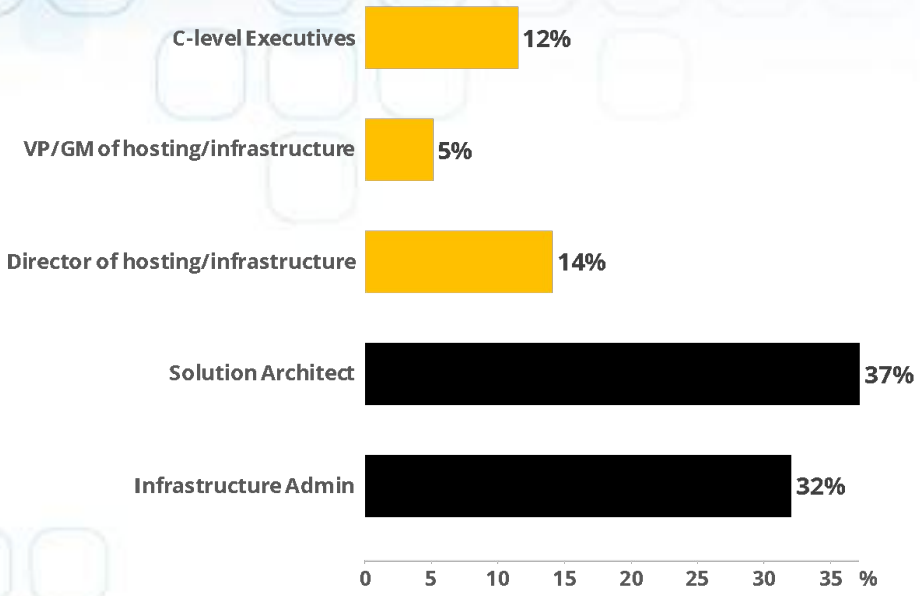

31% are decision makers

Who usually makes the final decision on storage purchases?

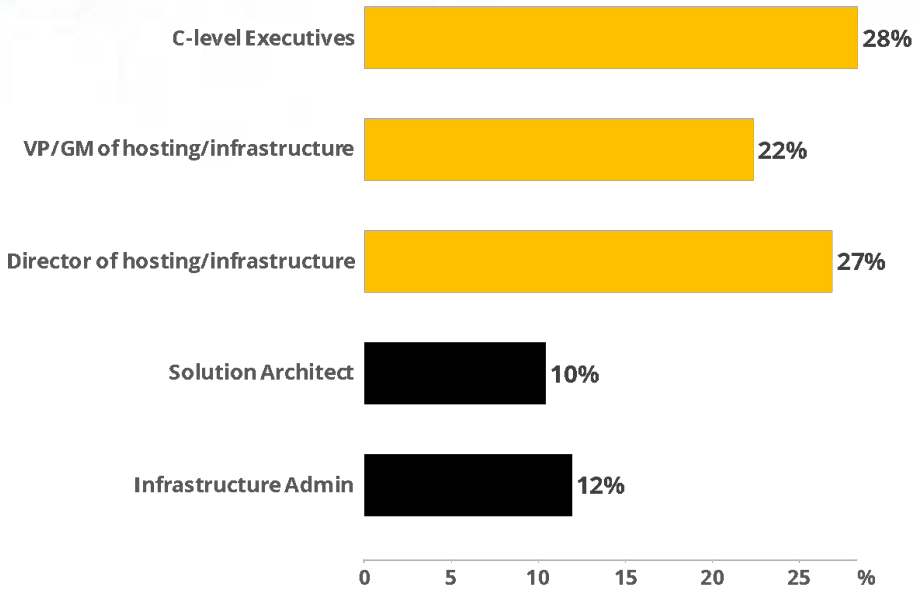

87% of decisions by director and above

Highly Virtualized Environment

75% are over 80% virtualized

25% are 100% virtualized

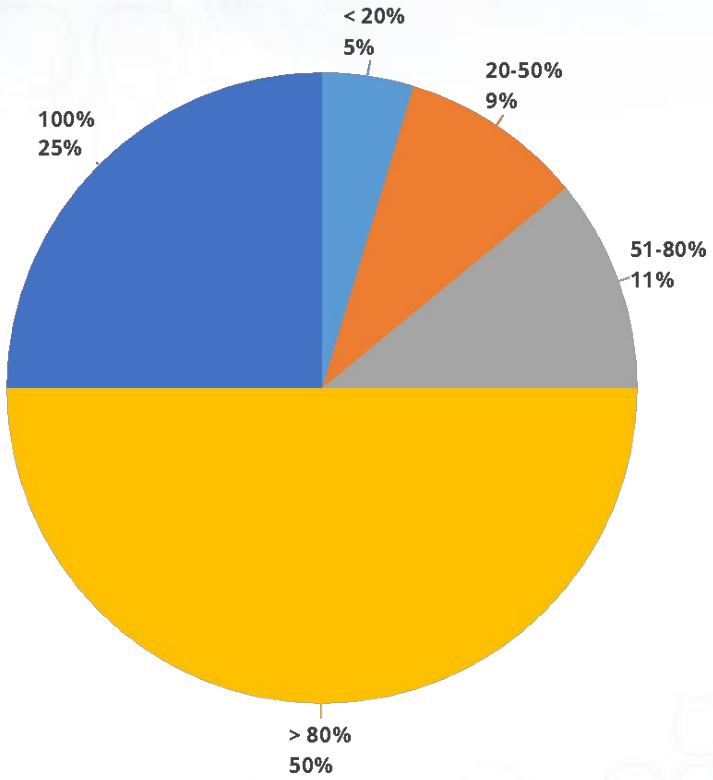

Guaranteed Performance at Scale

84% have over 250 VMs

67% have over 1000 VMs

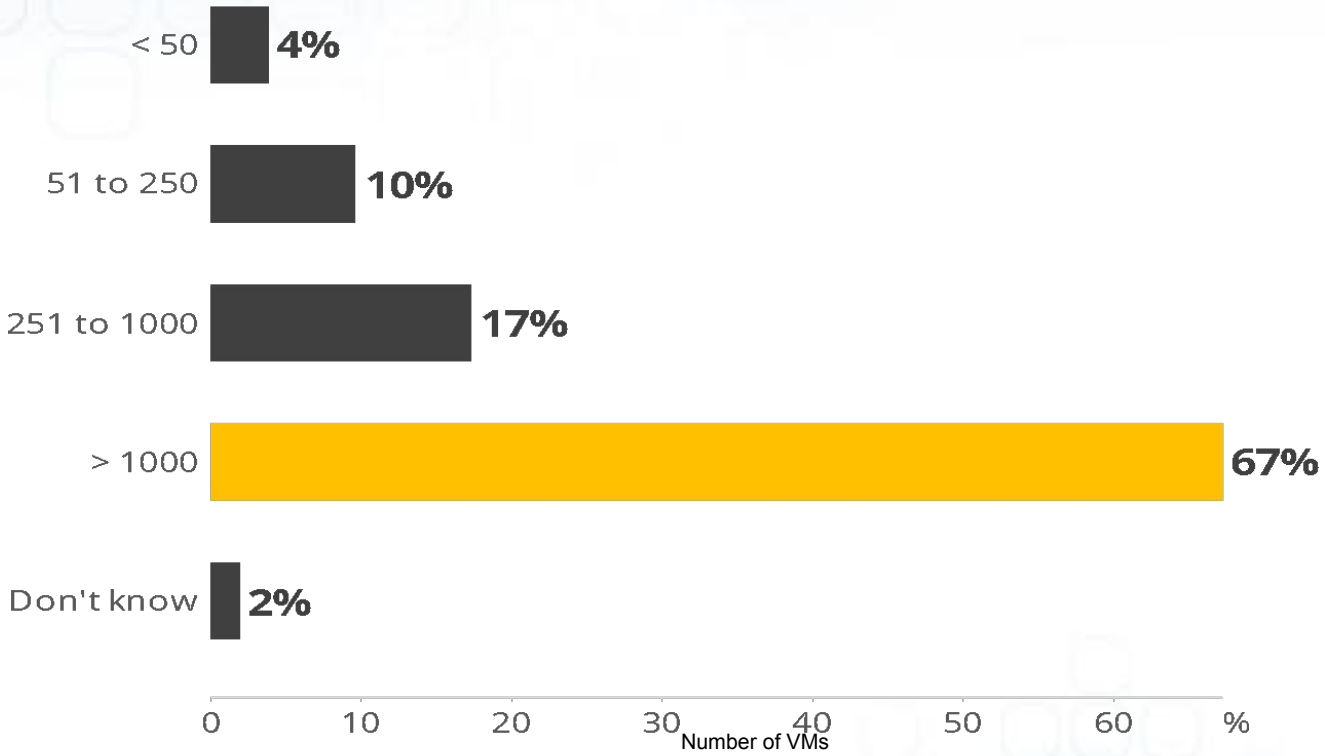

Hypervisor Strategy

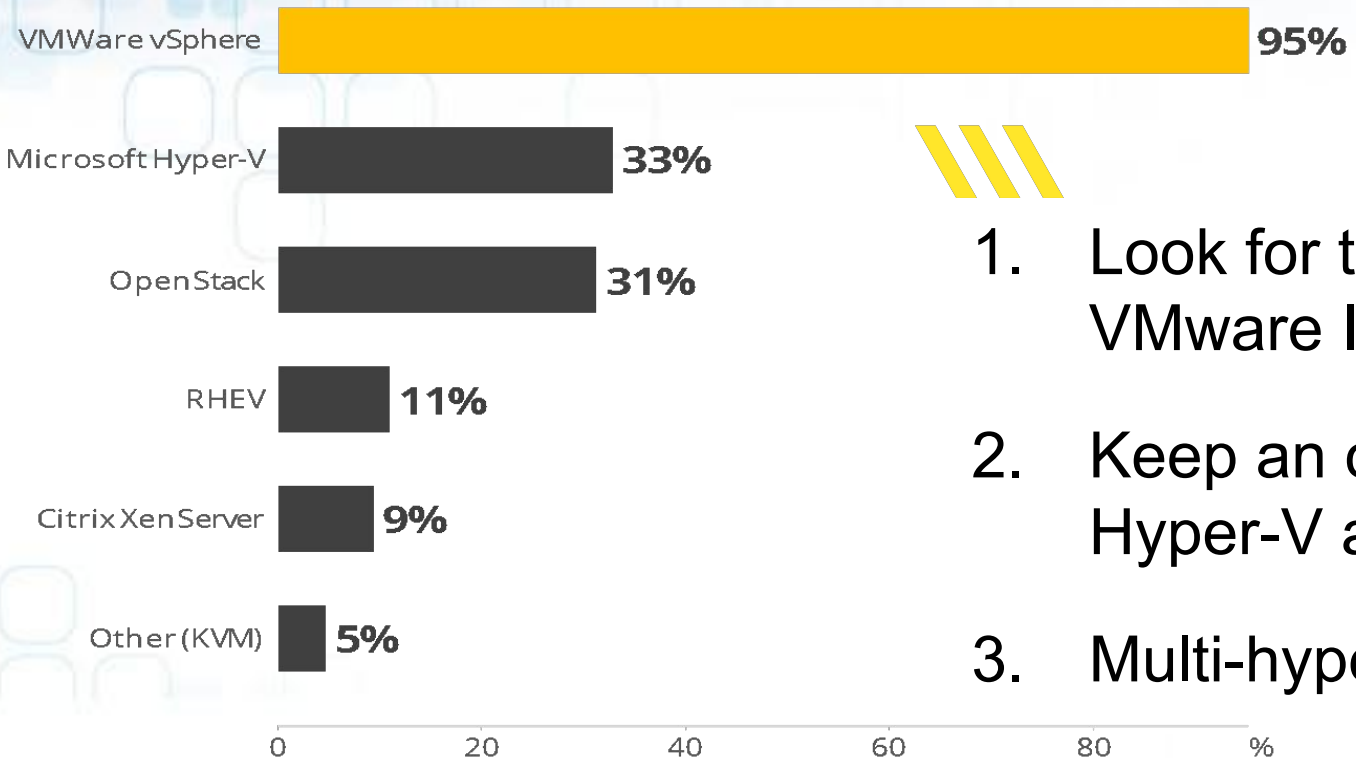

1. Look for the best VMware Integration
2. Keep an open mind for Hyper-V and OpenStack
3. Multi-hypervisor support

Storage Can Be a Make or Break Decision

1/3 spend over 10% of revenue on storage

2-3% is what enterprises spend on storage in contrast

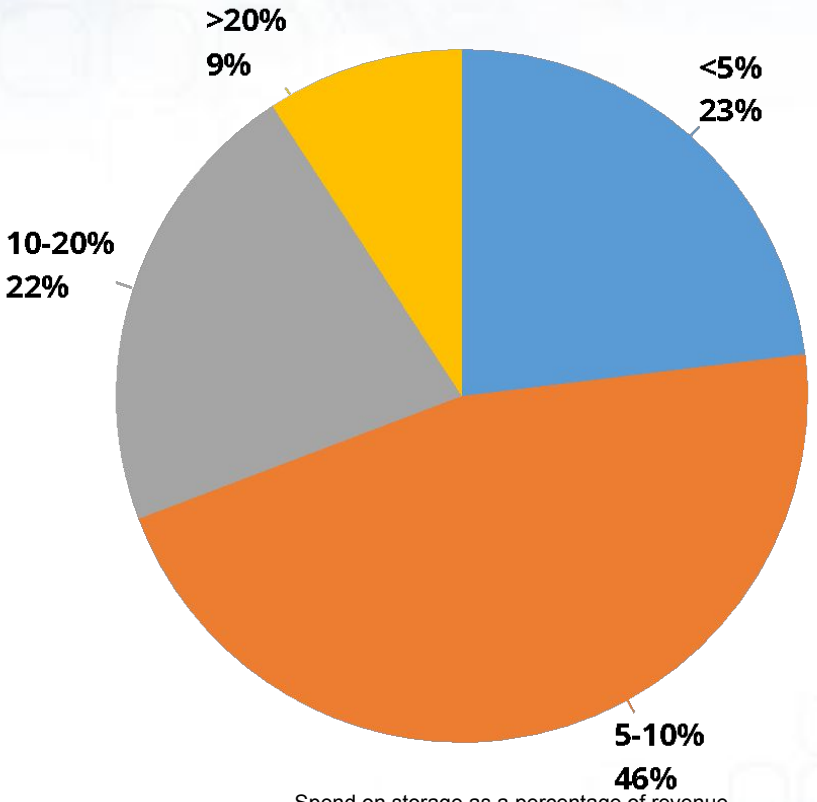

Spend on storage as a percentage of revenue

Top Three Purchase Criteria

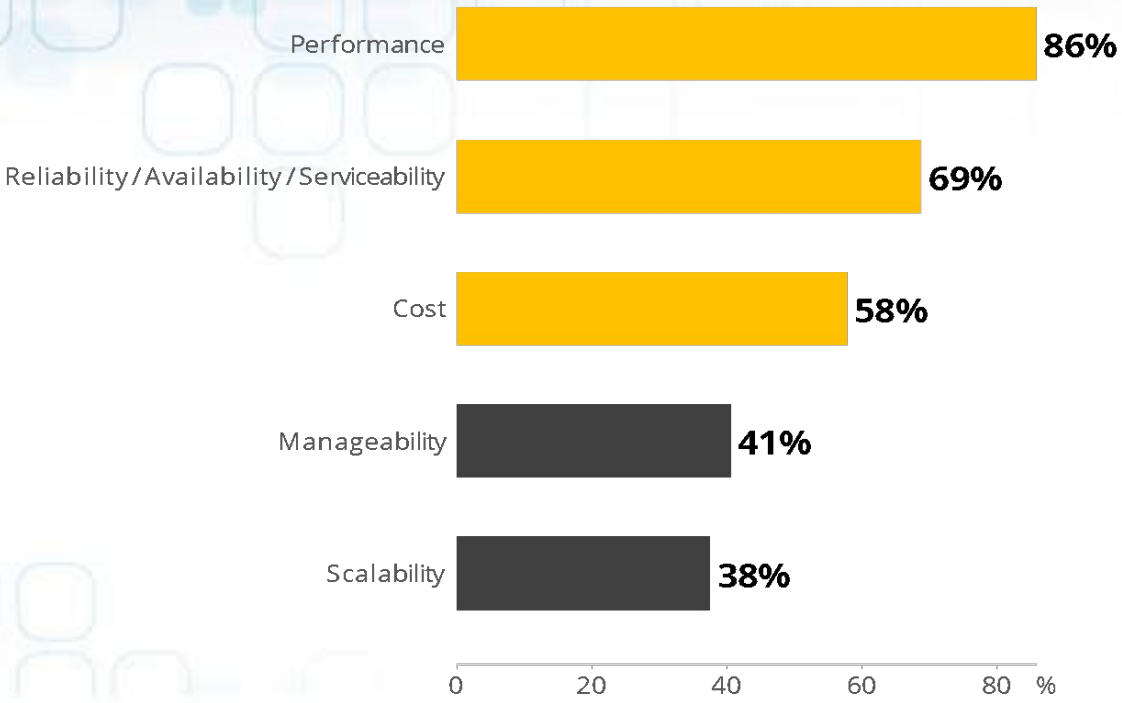

1. Performance
2. Reliability
3. Cost

Top Enabling Features

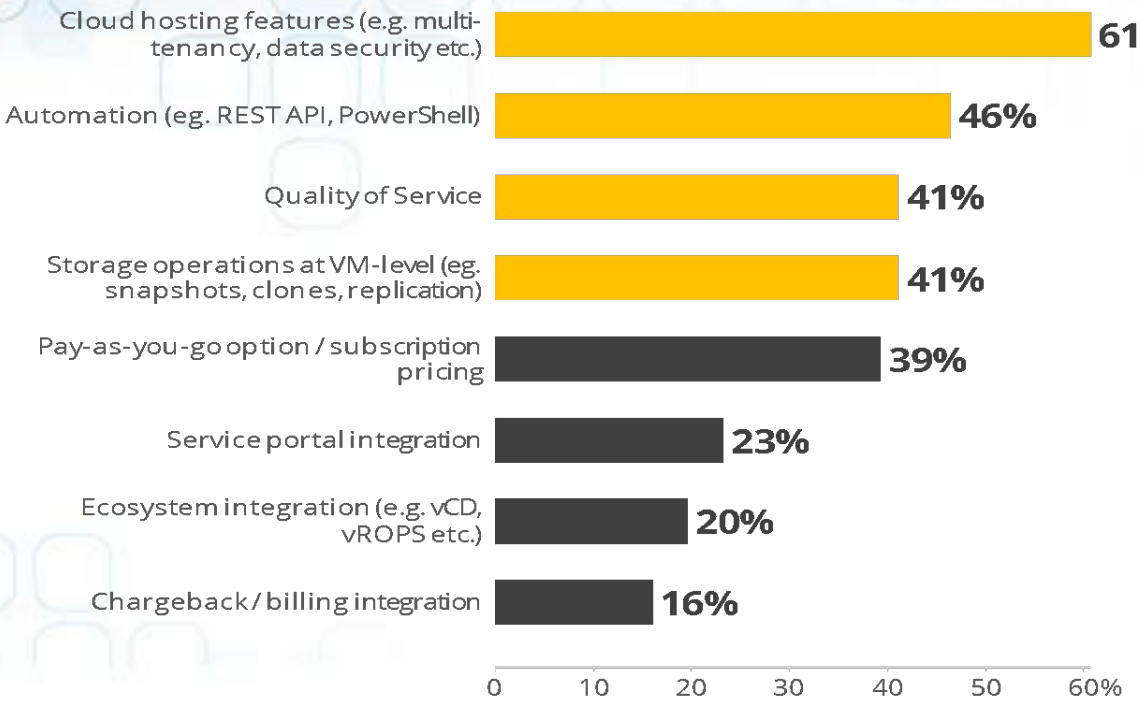

Differentiated Services

1. Cloud hosting
2. Automation
3. QoS and VM-level services

© 2016 Tintri Inc. All Rights Reserved.

Why is storage a challenge for CSPs?

It used to be simple…

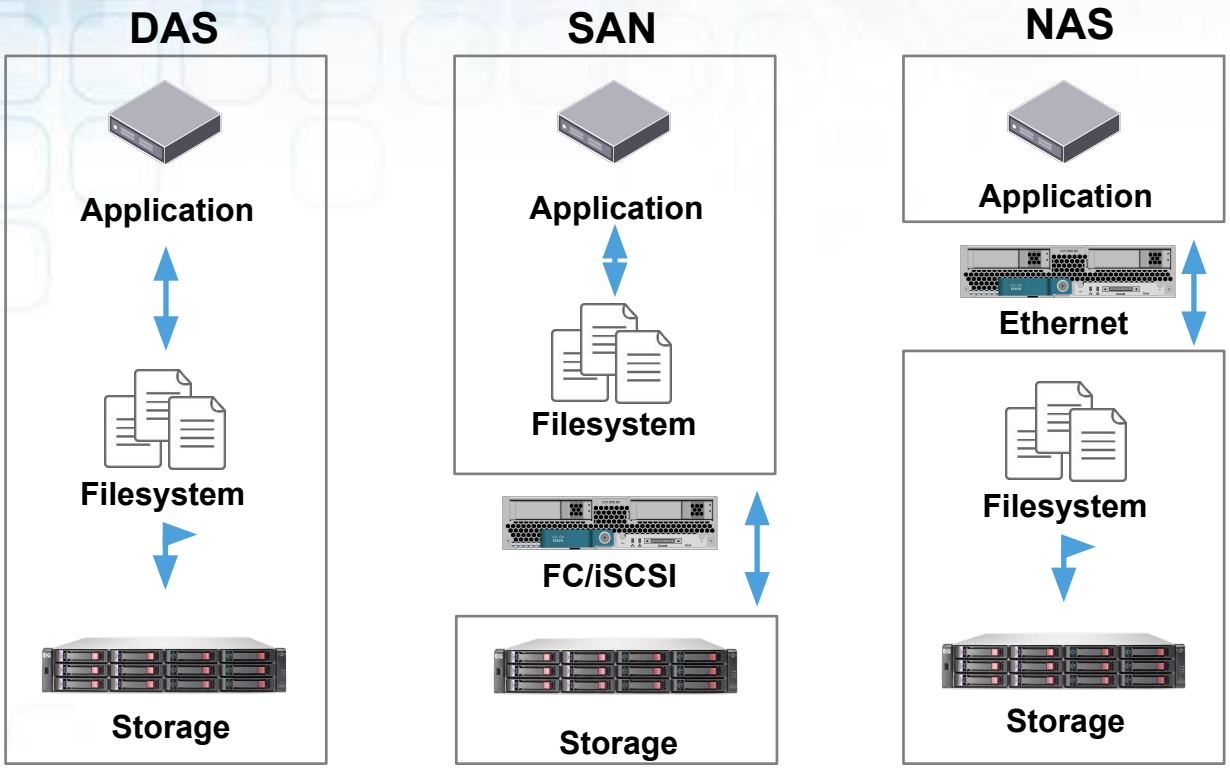

Now it's hard…

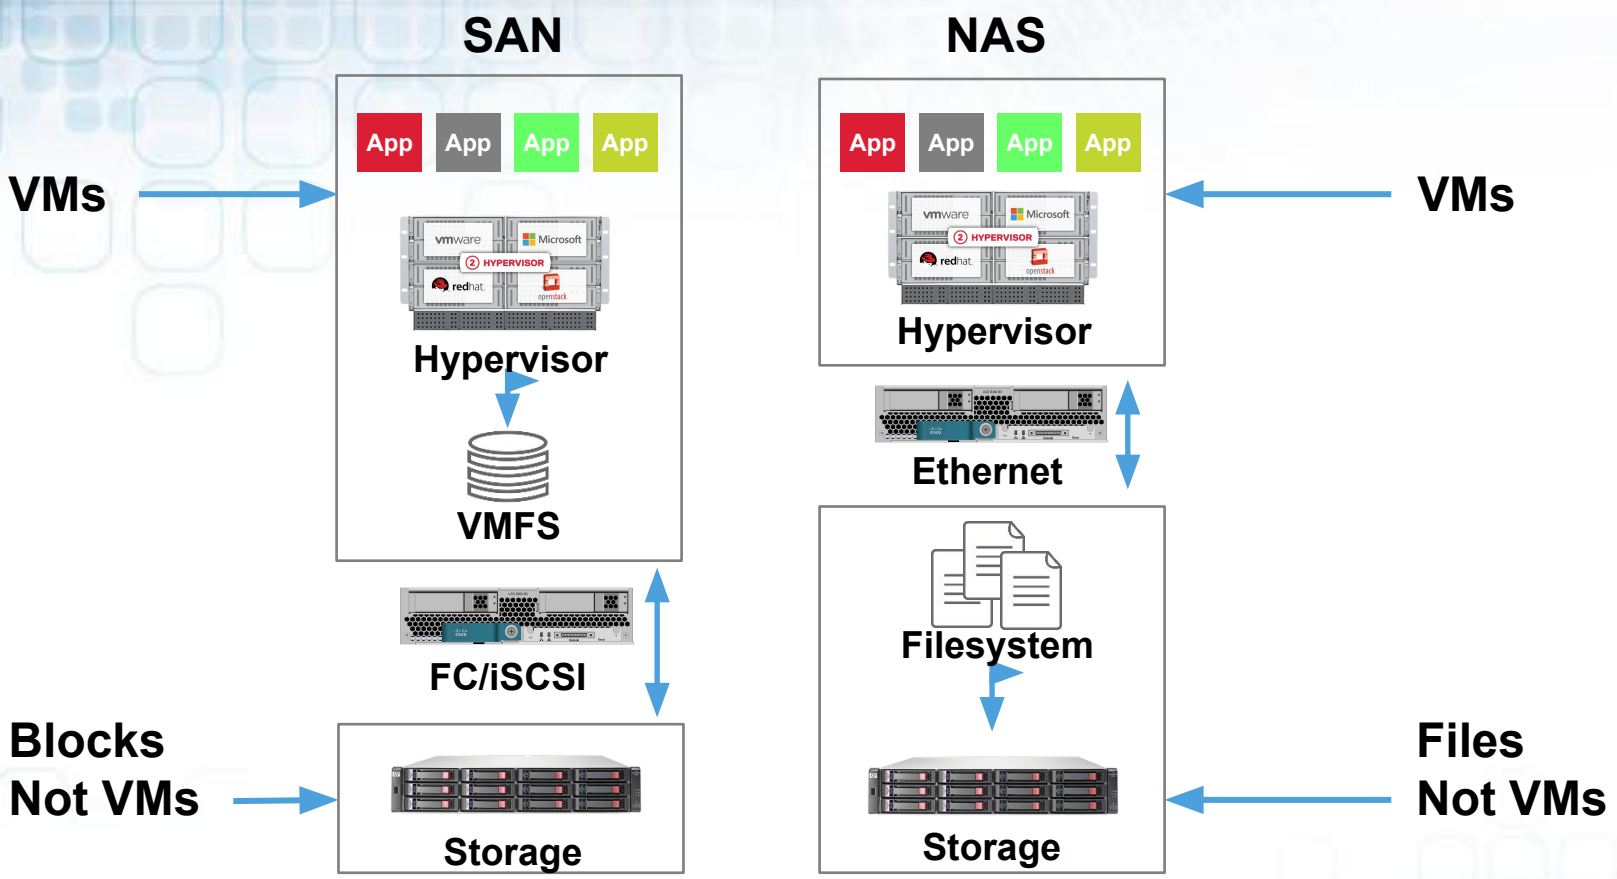
SAN
NAS
VMs
VMs
App
App
App
App
Hypervisor
Hypervisor
VMFS
Ethernet
FC/iSCSI
Filesystem
Blocks
Not VMs
Files
Not VMs
Storage
Storage

Why VM-aware Storage

Physical Data Center

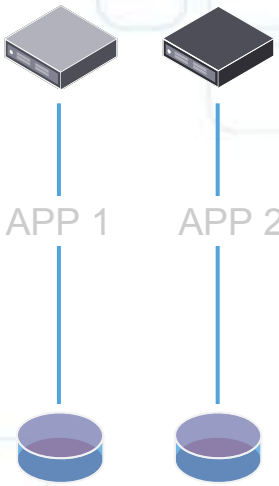

Optimized
Per Application

Virtualized Data Center
Traditional Storage

Optimized for
AGGREGATE LOAD

Virtualized Data Center
VM-aware Storage

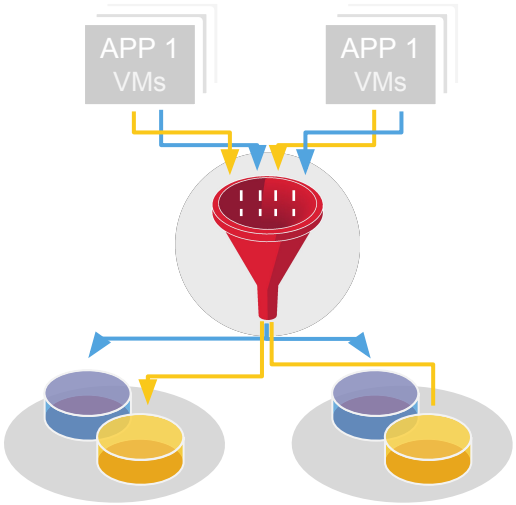

**Optimized
PER APPLICATION!**

VM-Level Analytics

VM-Level Analytics

VM	
tingle pd-t880-01	
tingle pd-t880-02	
tintri-demo-aligned pd-t880-01	
vgentry-demo	
tintri-demo-aligned-not pd-t880-01	
tintri-demo-hostlatency pd-t880-01	
tintri-demo-flass-miss pd-t880-01	
tintri-demo-slownetwork pd-t880-01	
pd-t880-01-RG_SQL_Thick	
pd-scvmm-01	
ExtraHop Discovery Edition	
vc-srm-pri	
vSphere 6GA	
vc-srm-sec	
Jay_SyncVM_demo-dev03	*Oct 07*
Win10-TechPreview--313	*Sep 02*
Ortega-clone-2	
smoot-ws12r2	
Ariix-3	
Merlin Win 10	
jay_test	

A list of VMs is a great start…

VM-Level Analytics

VM	Provisioned GiB	Used GiB
tingle pd-t880-01	53	577.3
tingle pd-t880-02	53	271.7
tintri-demo-aligned pd-t880-01	51	53.2
vgentry-demo	53	18.4
tintri-demo-aligned-not pd-t880-01	51	27.6
tintri-demo-hostlatency pd-t880-01	41	67.7
tintri-demo-flass-miss pd-t880-01	1,065	1,807.7
tintri-demo-slownetwork pd-t880-01	43	4.8
pd-t880-01-RG_SQL_Thick	328	12.3
pd-scvmm-01	**258**	**60.9**
ExtraHop Discovery Edition	20	0.6
vc-srm-pri	208	34.9
vSphere 6GA	108	0.3
vc-srm-sec	208	22.9
Jay_SyncVM_demo-dev03 Oct 07	*0*	*18.2*
Win10-TechPreview--313 Sep 02	*0*	*10.0*
Ortega-clone-2	262	0.01
smoot-ws12r2	44	8.7
Ariix-3	0	0.0
Merlin Win 10	100	10.7
jay_test	102	0.0

Sizes are good too…

VM-Level Analytics

This can become a CSAT issue

VM	Provisioned GiB	Used GiB	IOPS ▼	MBps	Latency ms
tingle pd-t880-01	53	577.3	18,662	190.1	0.4
tingle pd-t880-02	53	271.7	6,376	67.8	2.6
tintri-demo-aligned pd-t880-01	51	53.2	4,357	35.7	0.4
vgentry-demo	53	18.4	2,499	20.5	0.9
tintri-demo-aligned-not pd-t880-01	51	27.6	1,989	16.3	1.6
tintri-demo-hostlatency pd-t880-01	41	67.7	697	11.3	14.1
tintri-demo-flass-miss pd-t880-01	1,065	1,807.7	87	0.4	0.6
tintri-demo-slownetwork pd-t880-01	43	4.8	16	0.3	23.8
pd-t880-01-RG_SQL_Thick	328	12.3	9	0.08	0.6
pd-scvmm-01	**258**	**60.9**	**4**	**0.05**	**0.5**
ExtraHop Discovery Edition	20	0.6	2	0.01	0.0
vc-srm-pri	208	34.9	2	0.01	0.0
vSphere 6GA	108	0.3	1	0.06	0.0
vc-srm-sec	208	22.9	1	0.01	0.0
Jay_SyncVM_demo-dev03 Oct 07	*0*	*18.2*	*0*	*0.0*	*0.0*
Win10-TechPreview--313 Sep 02	*0*	*10.0*	*0*	*0.0*	*0.0*
Ortega-clone-2	262	0.01	0	0.0	0.0
smoot-ws12r2	44	8.7	0	0.0	0.0
Ariix-3	0	0.0	0	0.0	0.0
Merlin Win 10	100	10.7	0	0.0	0.0
jay_test	102	0.0	0	0.0	0.0

IOPS
MBps
Latency

Realtime per VM

CSPs can charge for these

VM-Level Analytics <> Averages

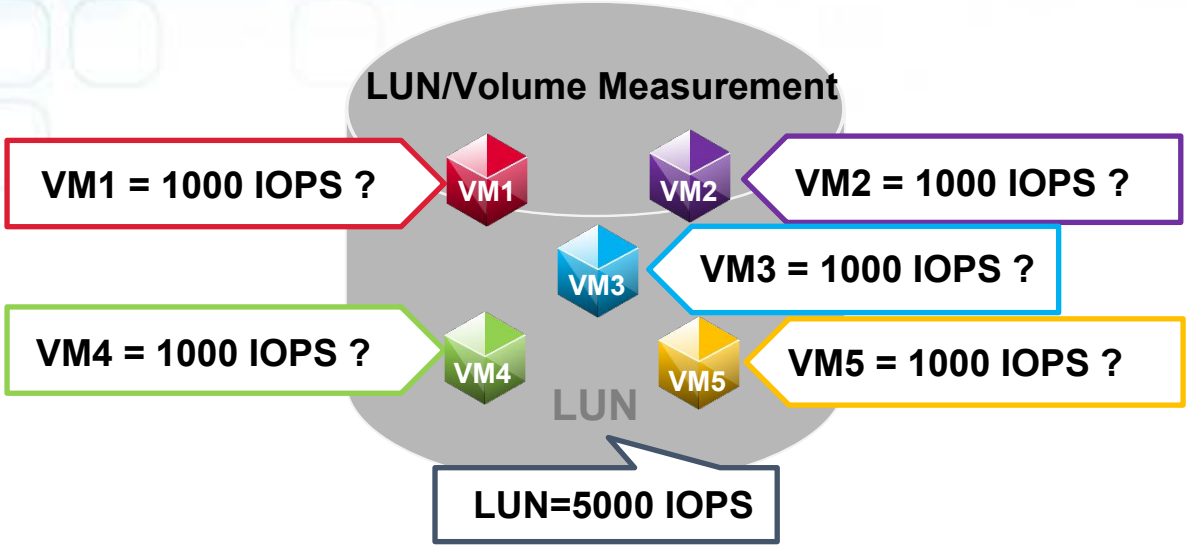

VM-Level Analytics <> Correlations

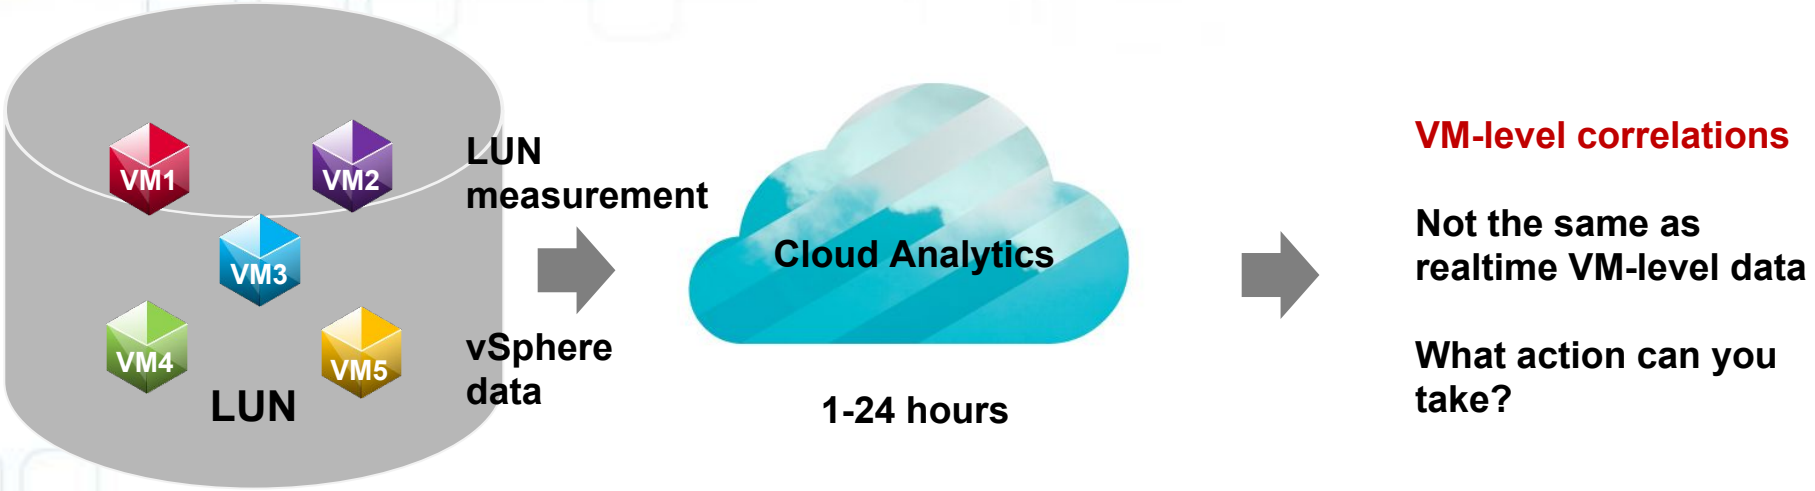

VM-Level Analytics

- Real-time measurement, actionable insights

	IOPS ▼	MBps	Latency ms
VM1	18,662	190.1	0.4
VM2	6,376	67.8	2.6
VM3	4,357	35.7	0.4
VM4	2,499	20.5	0.9
VM5	1,989	16.3	1.6

VM-Level Analytics

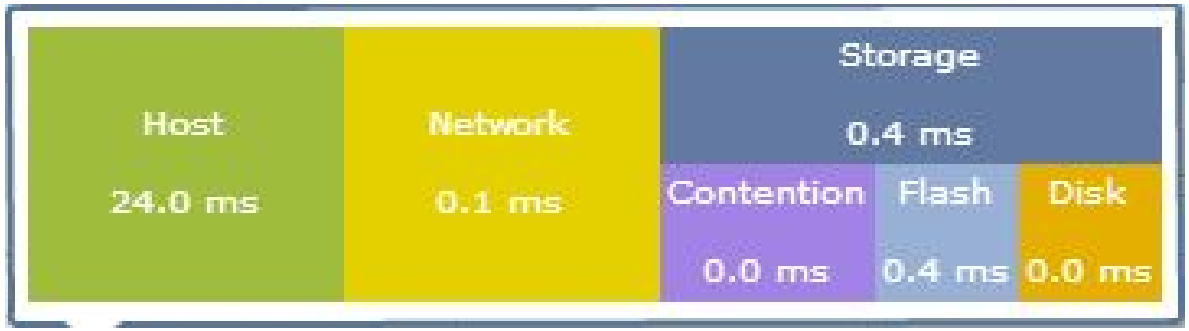

End-to-End Latency

VM-Level Storage Operations

Array-based storage operations

✓ Fast

✓ Storage efficient

✓ No host or network impact

✗ LUN/Volume level

Array-based replication

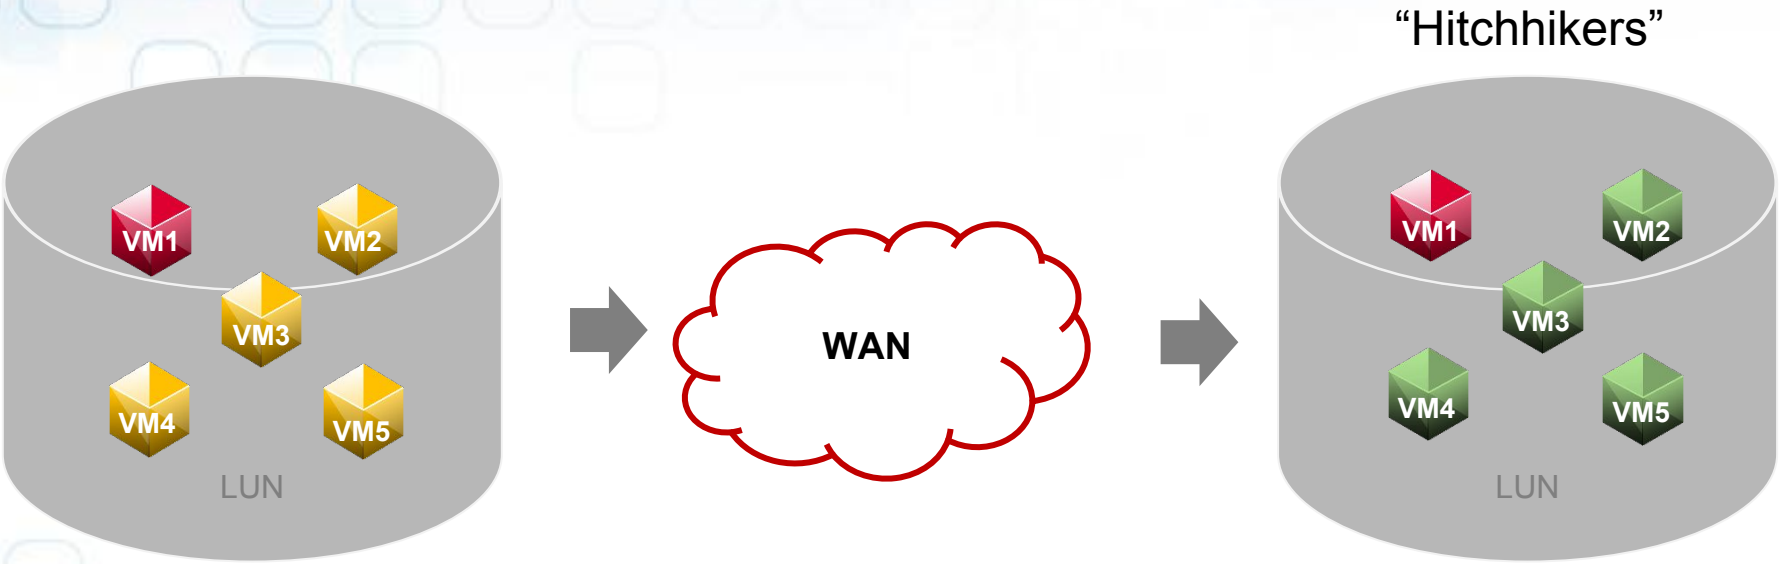

Array-based snapshotting

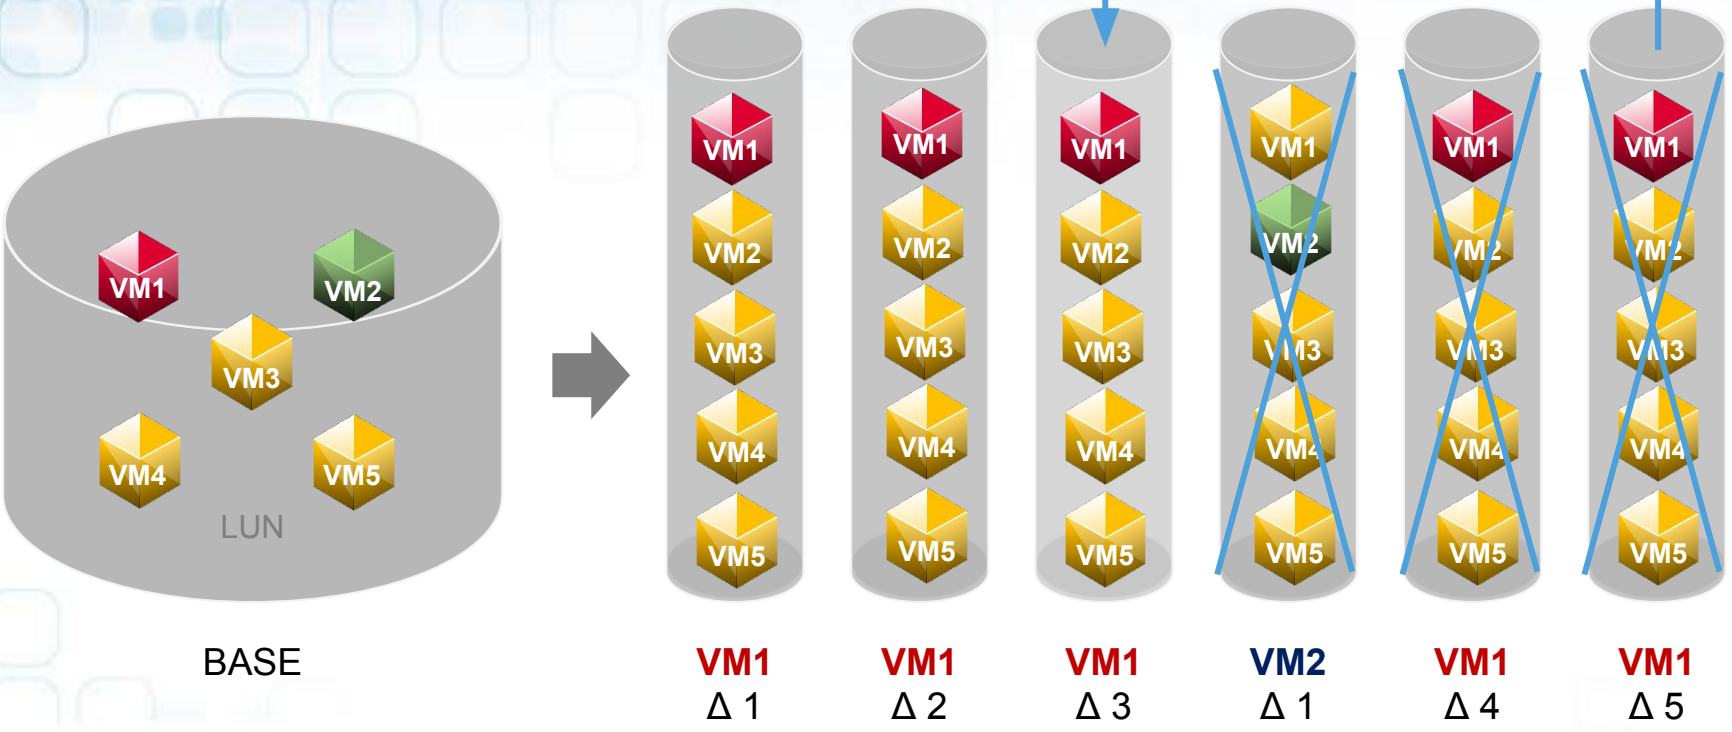

Software-based storage operations

✗ Slow

✗ Storage inefficient

✗ Host CPU and network impact

✓ VM-level flexibility

VM-aware storage operations

- ✓ Fast
- ✓ Storage efficient
- ✓ No host or network impact
- ✓ VM-level flexibility

VM-Level Quality of Service (QoS)

VM Quality of Service

Conventional FIFO = I/O Traffic Jam

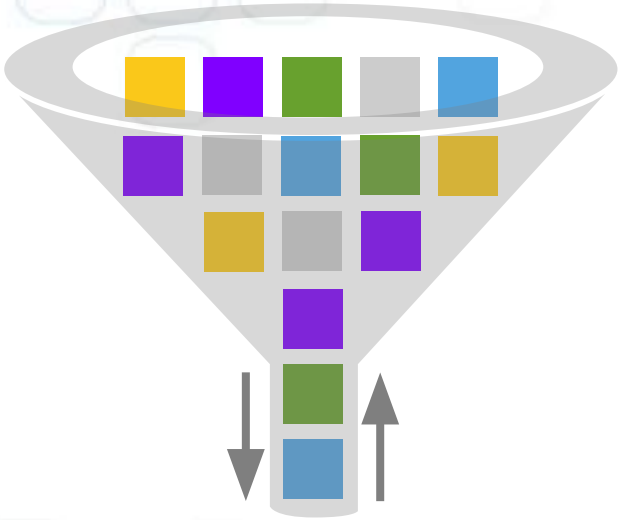

Per-VM and Per-vDisk Queueing

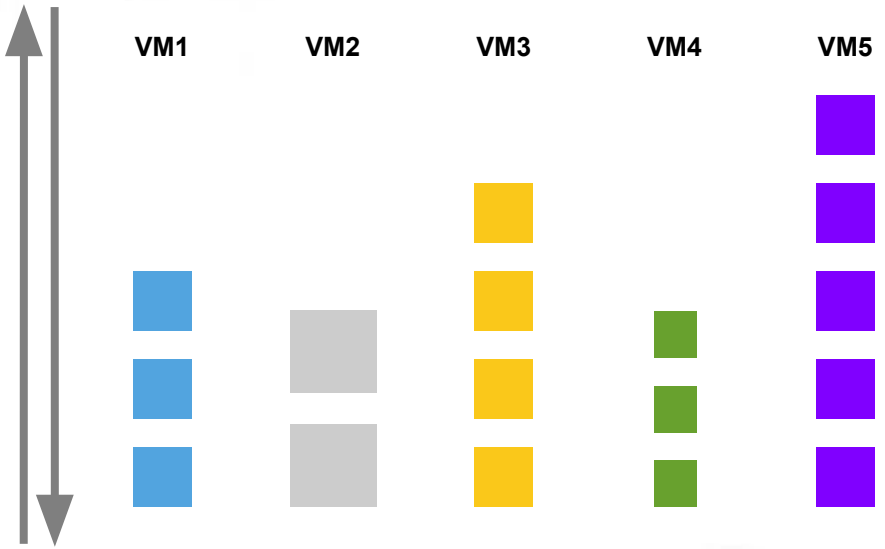

Automatic VM-level Isolation
Between tenants
Between VMs

VM-level QoS is real-time

- Visual guidance on how / when to throttle IOPS for each individual VM

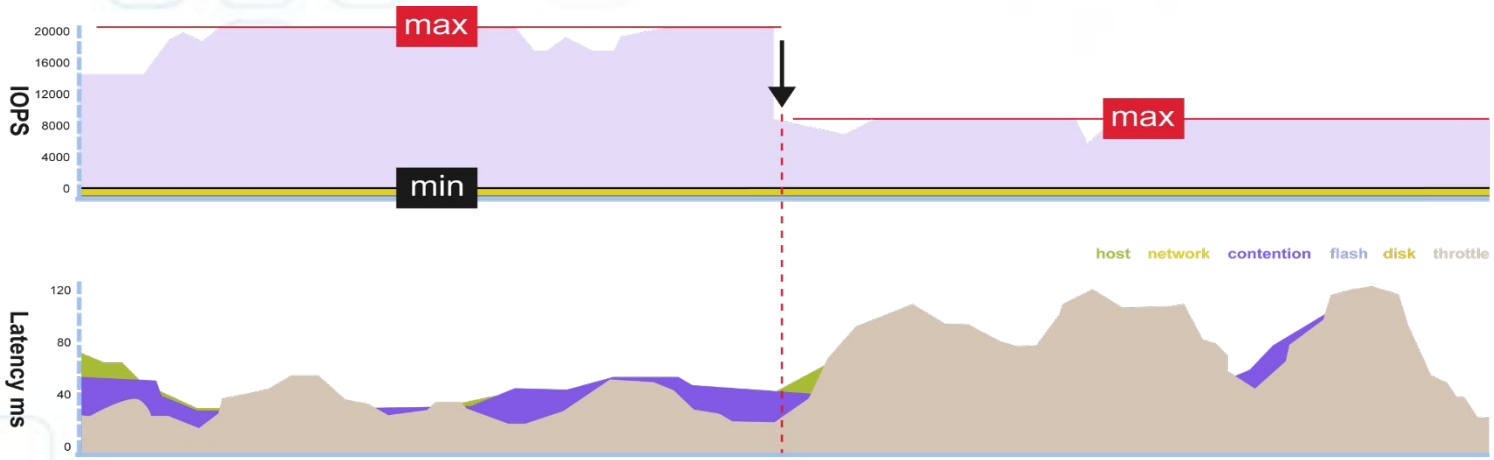

Immediate feedback on the latency / contention impact

Policy-based VM-level QoS

- Establish service level tiers and guarantee performance to your customers

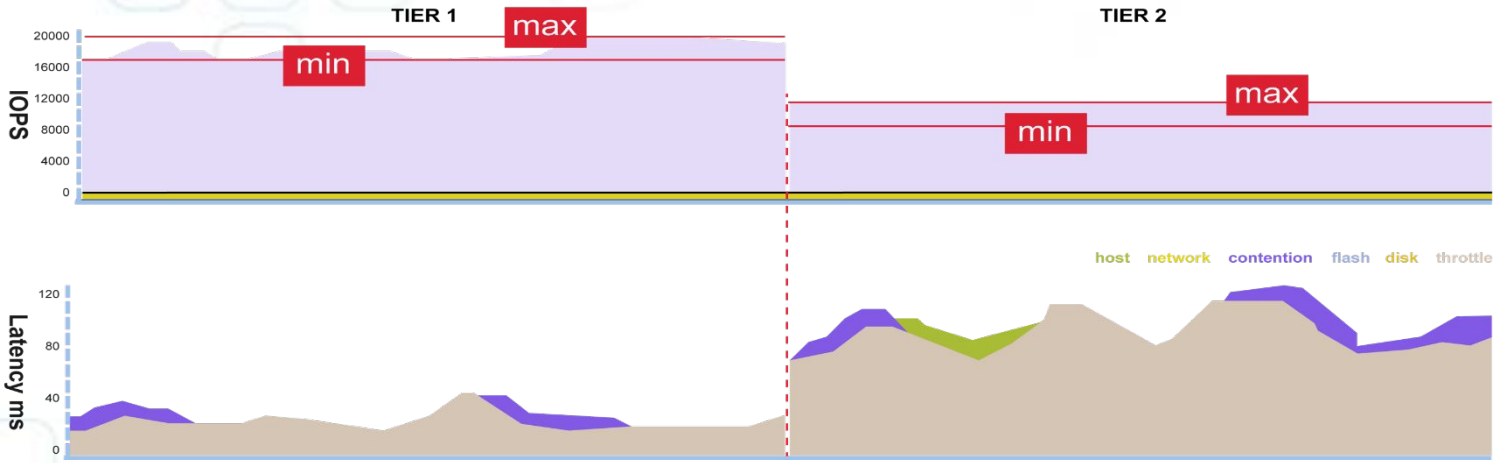

Provide visual chargeback and identify upsell opportunities

VM-Level Automation

Set VM1 to replicate to Site2 every 1 hour

Conventional storage **vs.** **VM-aware storage**

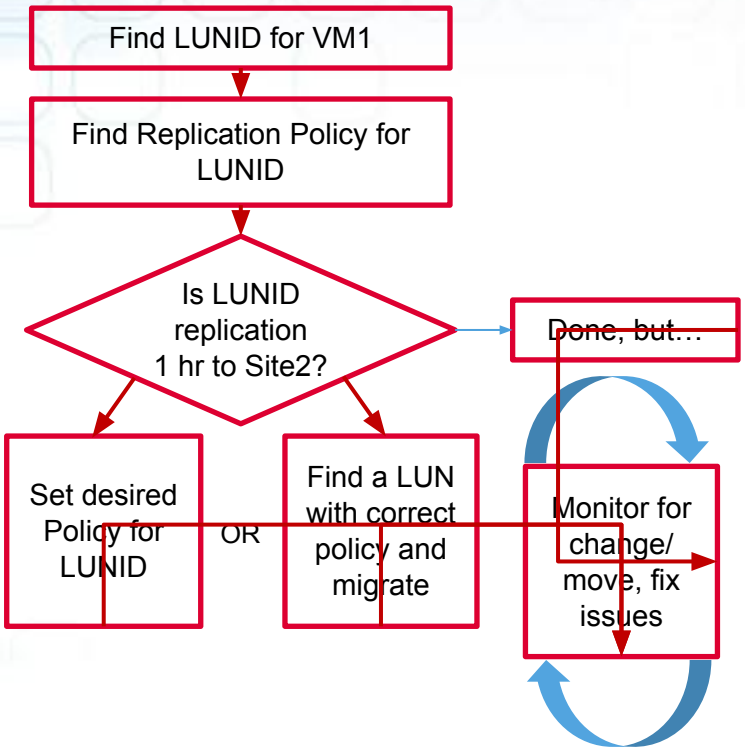

Set VM1 to replicate to Site2 every 1 hour.

Technology Solutions

Array-Offloading APIs

VAAI

ODX

Without API

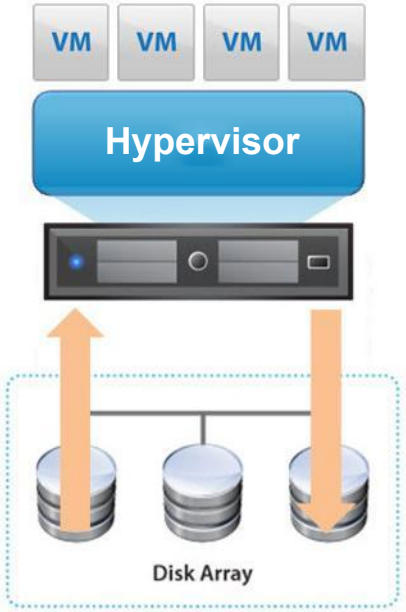

All data copied first to host, then back to array

With API

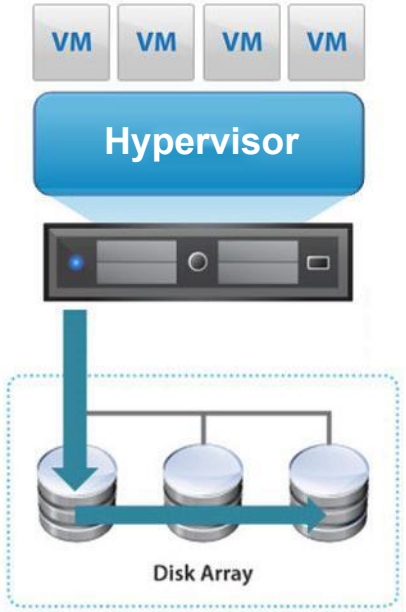

Copy command issued to array, array performs operation

Still LUN/ Volume operations

VMware Storage APIs

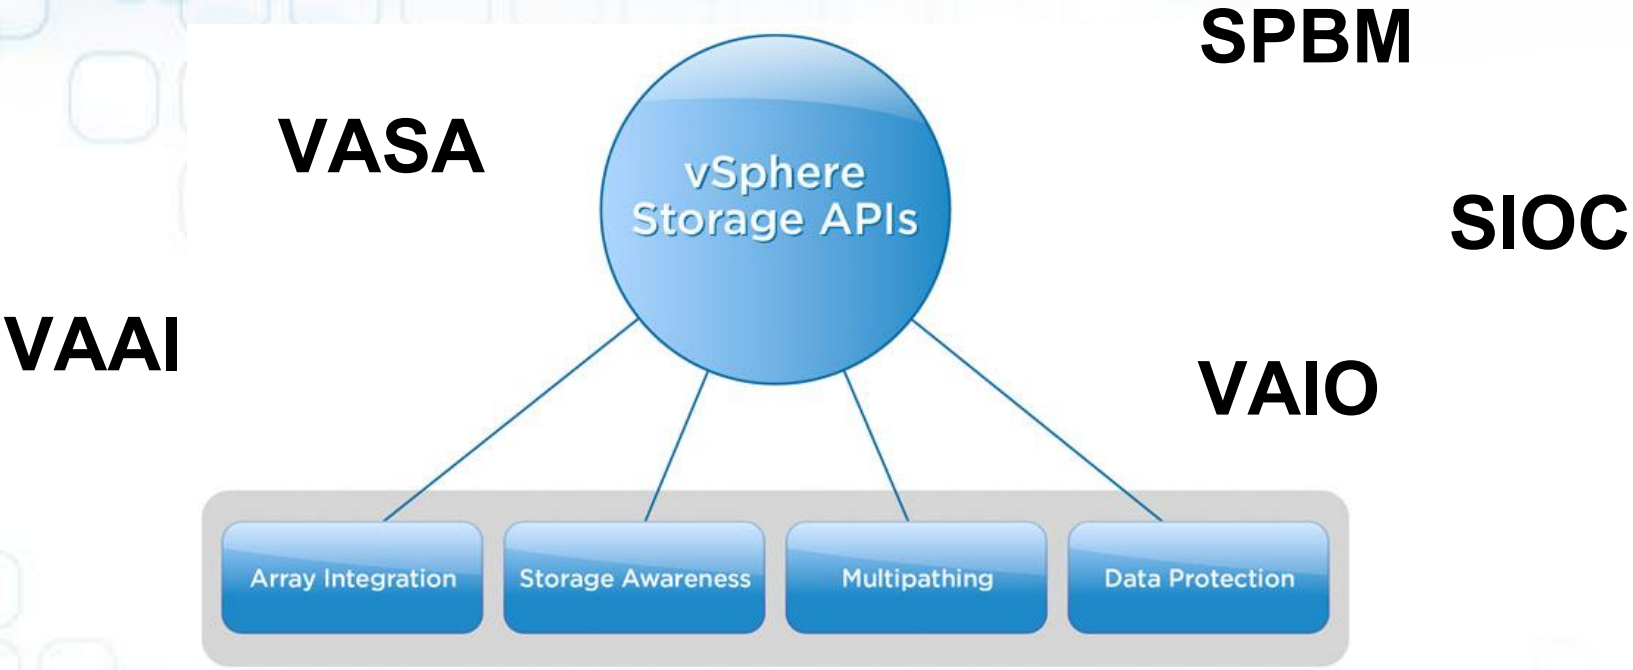

vSphere Virtual Volumes (VVols)

VVols helps VM admins deploy storage features at VM-level

Protocol Endpoints

Storage Containers

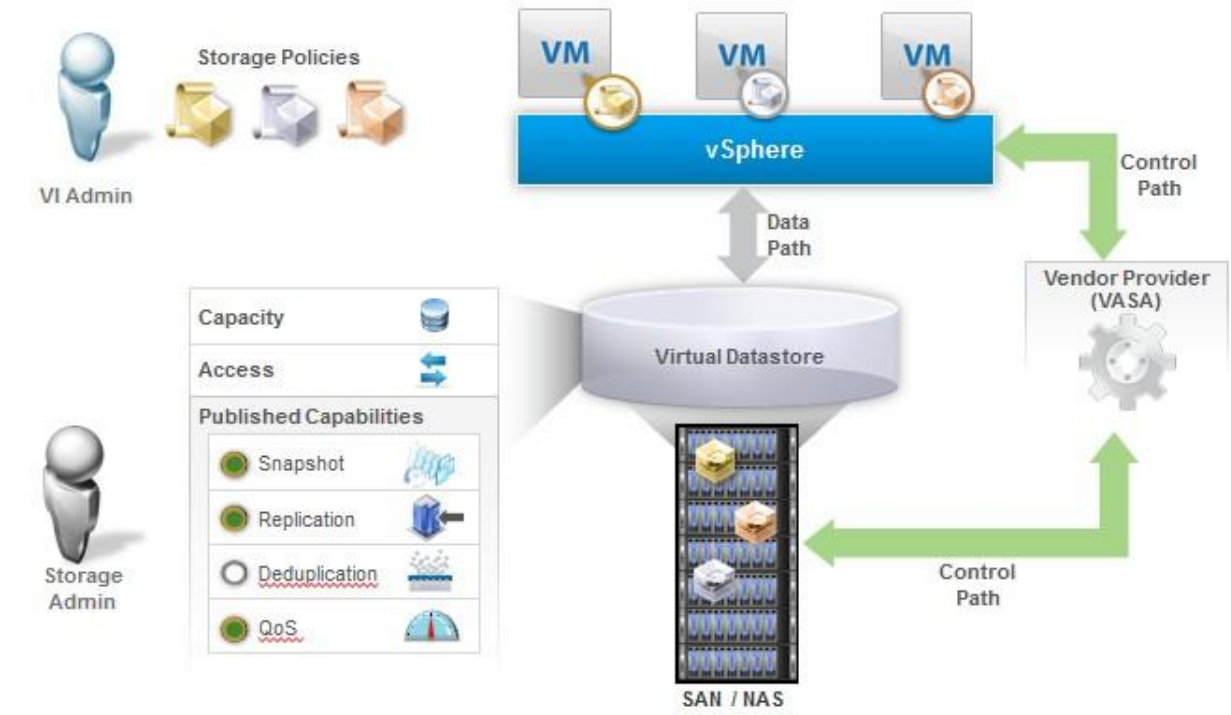

Hyperconverged

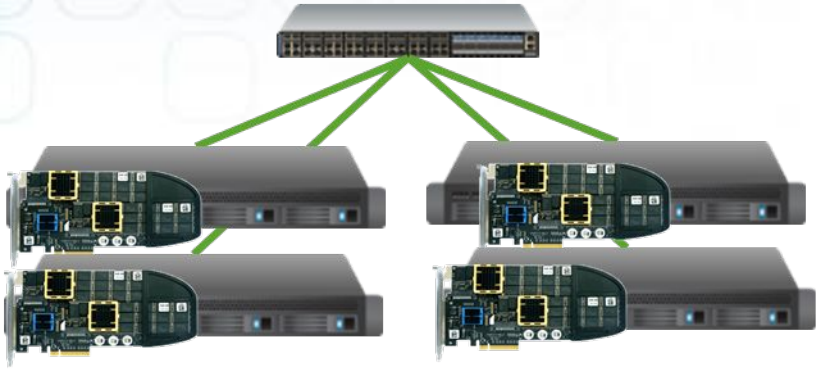

- “Invisible” storage
- Closer to VMs
- What is the filesystem underneath?
- Scaling issues beyond 3-5 nodes

VM-Level Filesystem

VM-aware Filesystem

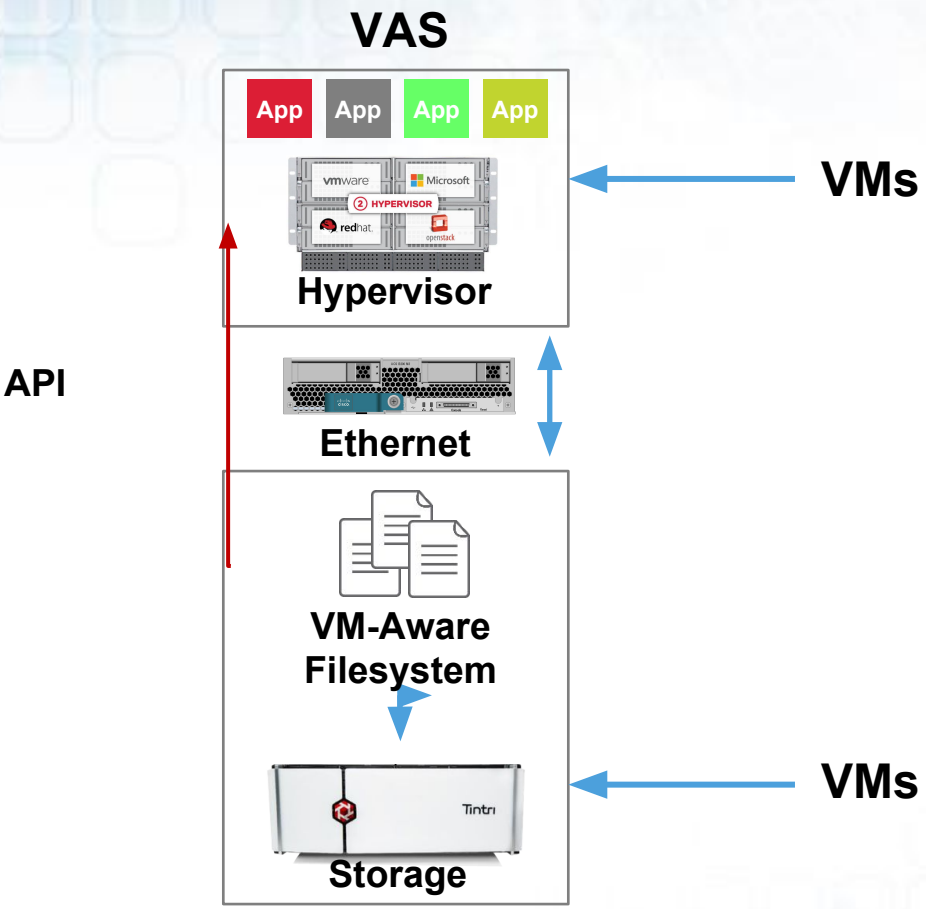

VM-aware Filesystem

Visibility

Performance

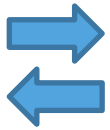

IO Scheduling

	Visibility	Performance	IO Scheduling
Convention	• Can "list" VMs and VMDKs • Filesystem is either block or file	• Optimize at the LUN/Volume Level • Unpredictable VM performance	• IO scheduled at LUN/Volume level • Noisy neighbors and contention
VM-aware	• Sees VMs and VMDKs • Optimized for VMs • Uses VM as unit of mgmt	• Optimized per-VM • Predictable VM performance	• IO scheduled at a VM level • VM-level isolation

Why VM-aware Storage

Physical Data Center

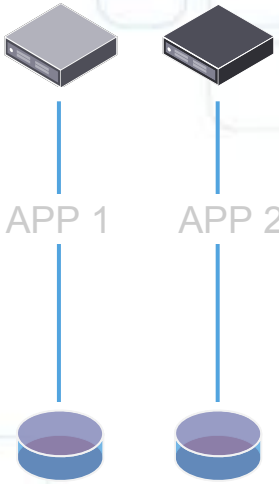

Optimized
Per Application

Virtualized Data Center
Traditional Storage

Optimized for
AGGREGATE LOAD

Virtualized Data Center
VM-aware Storage

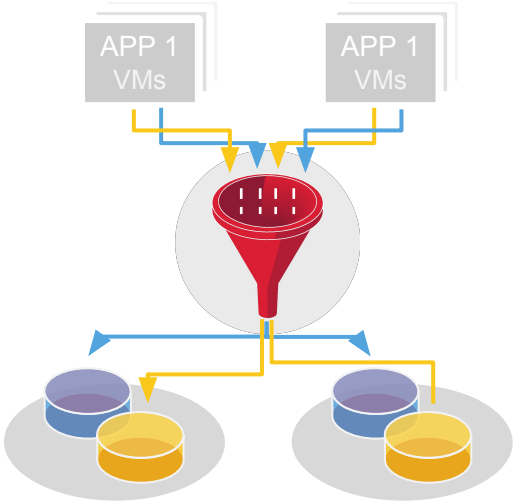

**Optimized
PER APPLICATION!**

Why VM-aware Storage with Cloud

Physical Data Center

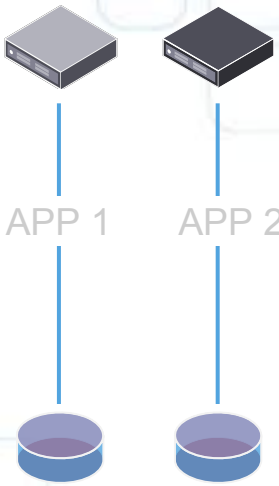

Optimized
Per Application

Virtualized Data Center
Traditional Storage

Optimized for
AGGREGATE LOAD

Could your traditional storage handle your growth?

2X?
10X?
Number of VMs in your DC today

2X?
10X?
Rate of change (create/ destroy) of VMs

Conclusion

Conclusion

- Cloud Service Providers REQUIRE VM-aware storage
- Five requirements for VM-aware storage
  - VM-level analytics: Realtime and actionable
  - VM-level storage operations
  - VM-level quality of service (QoS)
  - VM-level automation
  - VM-level filesystem

Questions?

ChuckD@tintri.com

@ChuckDubuque

www.tintri.com