

Characterizing Private Clouds:

A Large-Scale Empirical Analysis of Enterprise Clusters

Ignacio Cano, Srinivas Aiyar, Arvind Krishnamurthy
 icano@cs.washington.edu, sriniva.aiyar@nutanix.com, arvind@cs.washington.edu



Motivation

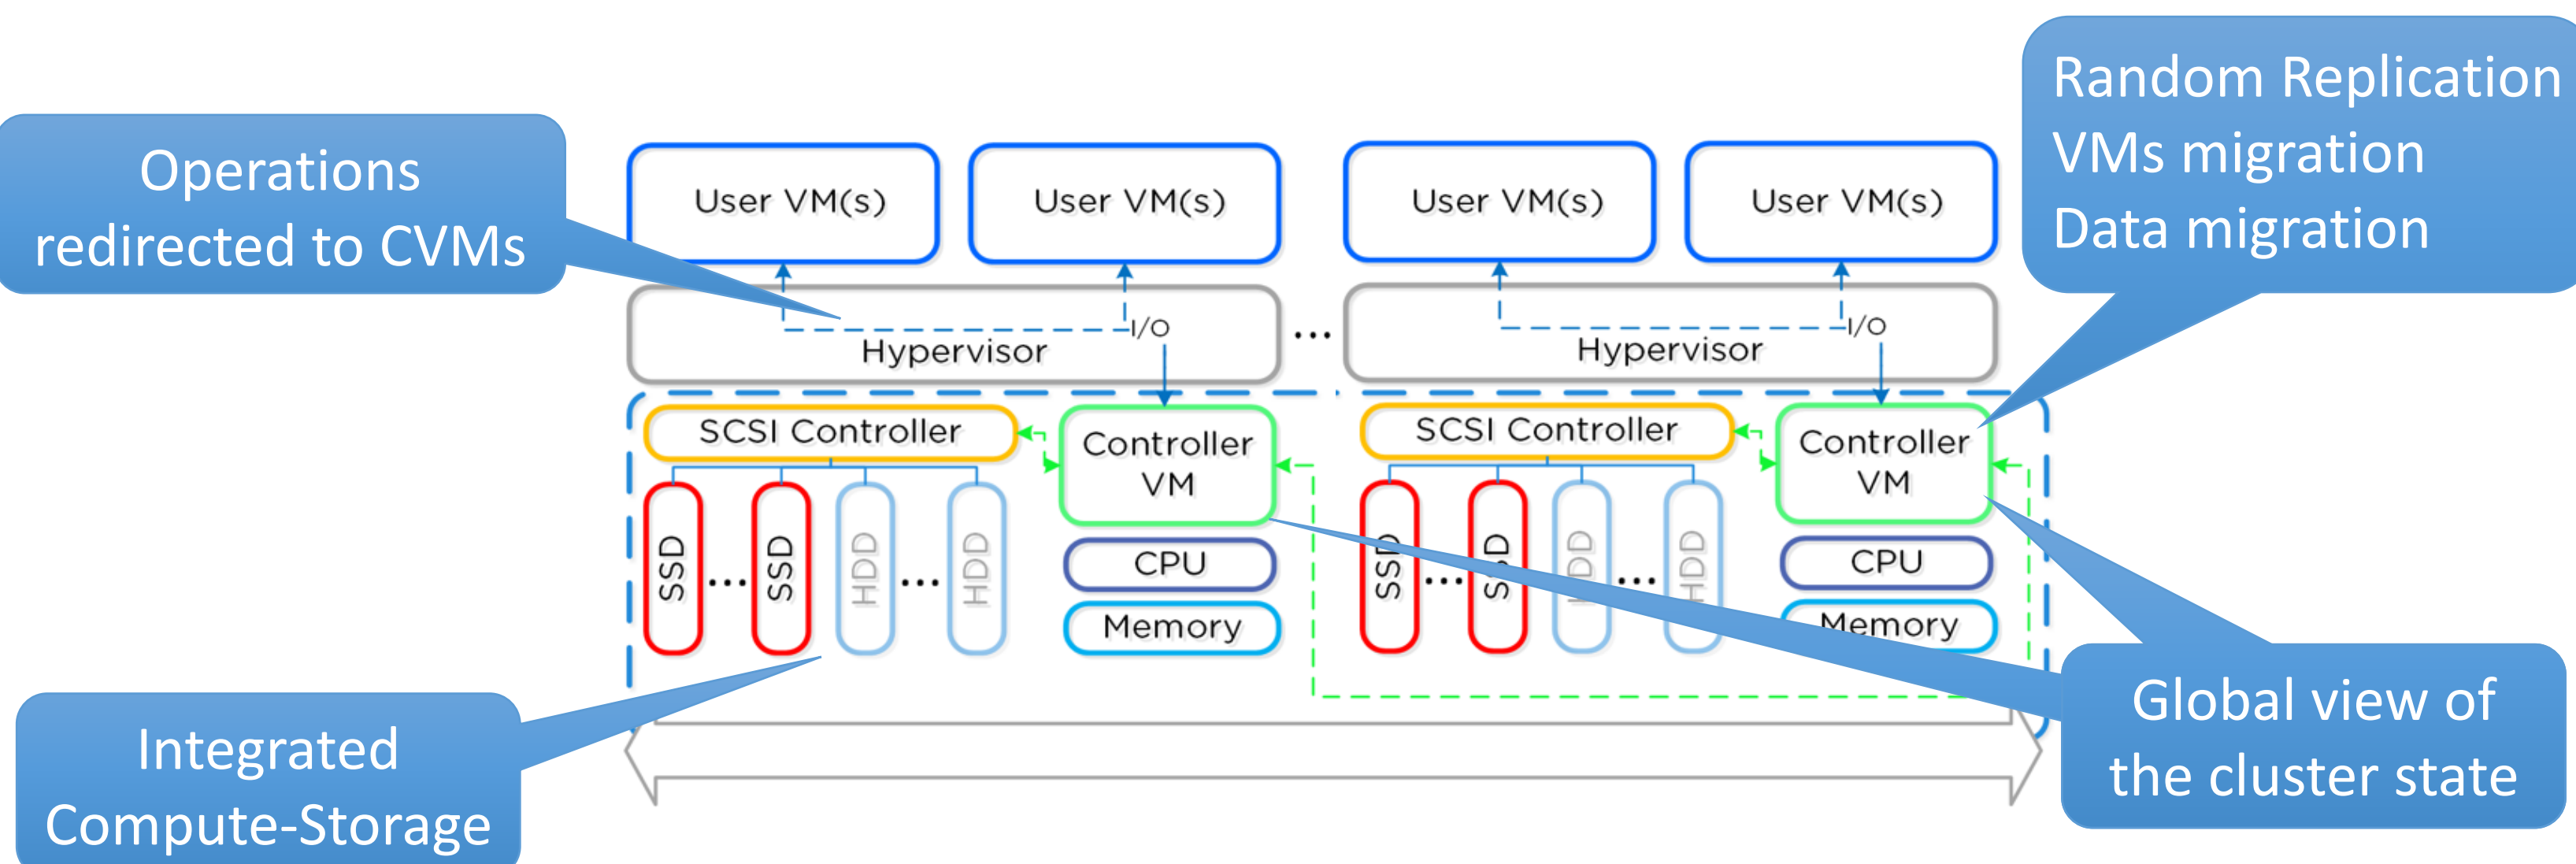


Private Clouds:

- What are the **most common failures**?
- What **type of workloads** are typically run?
- How is the **storage used**? What about **CPU usage**?
- How do additional **replicas** impact **data durability**?
- What **causes** companies to **expand their clusters**?

Large-Scale Measurement Study

Nutanix Clusters



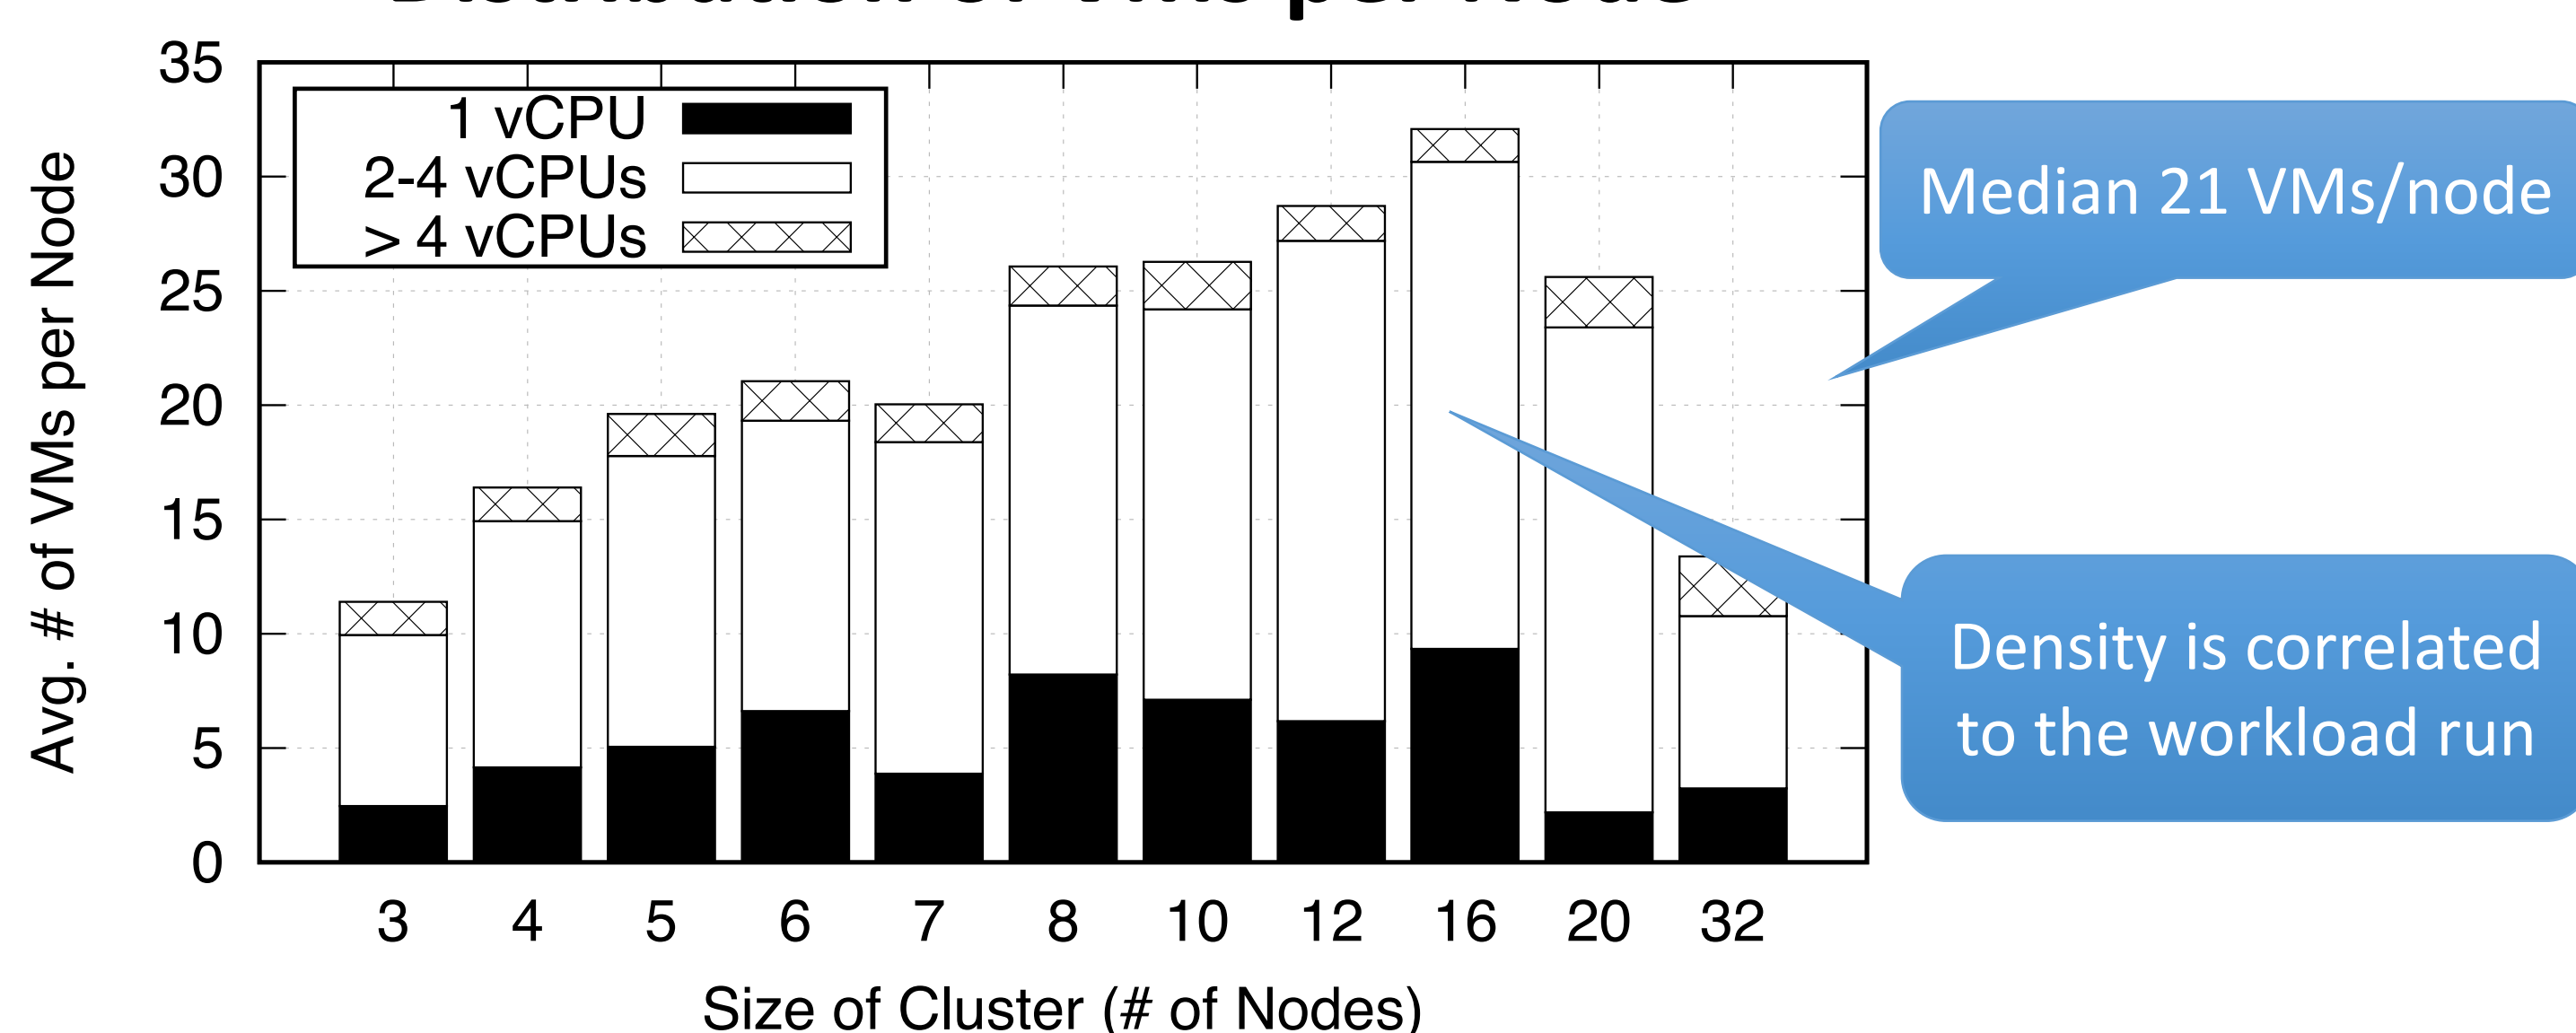
Cluster Profiles

Data Statistics	Value	Data Source	Information	From
# of Clusters	2168	Metrics	Storage, CPU	Aug-2015
# of Nodes	13394	Customers	Trouble Tickets	Dec-2011
Cluster Sizes	3 - 40	Repair and Maintenance	Return Rates	Aug-2013
# of Disks	~ 70K			

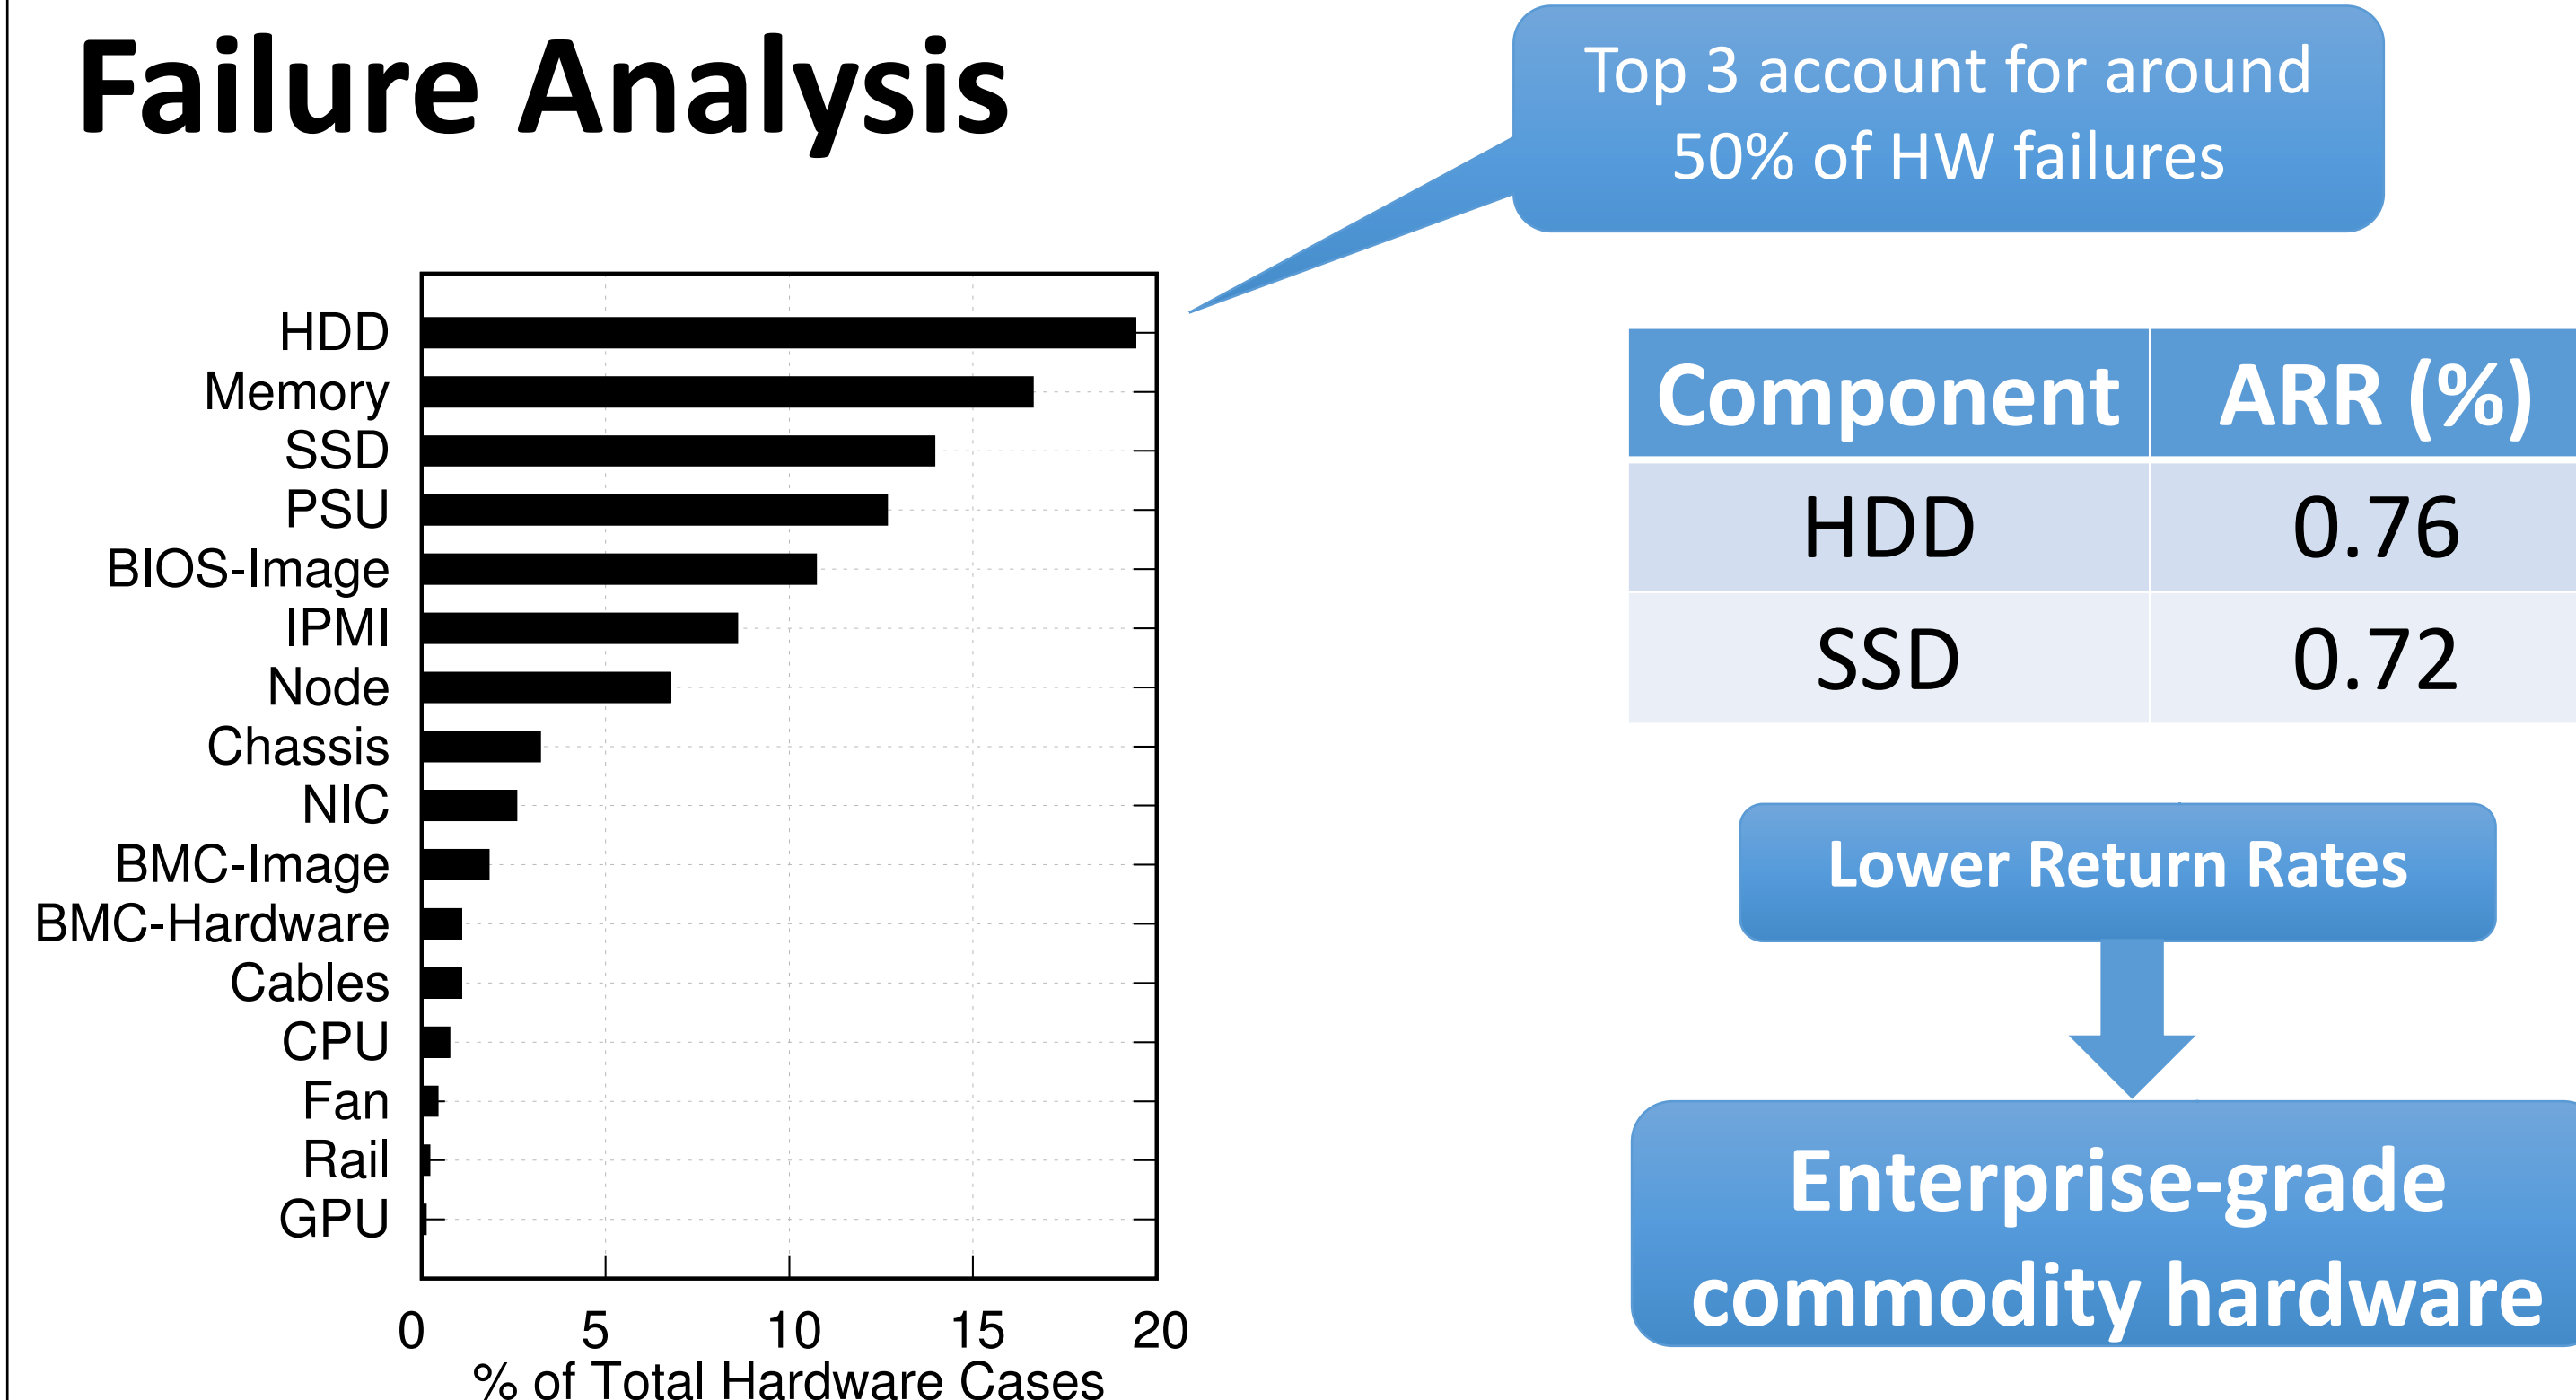
Node Configurations	Storage		Compute		Memory (GB)
	SSD (TB)	HDD (TB)	Cores	Clock Rate (GHz)	
Config-1	1.6	8	24	2.5	384
Config-2	0.8	4	12	2.4	128
Config-3	0.8	30	16	2.4	256

Workload	Example Applications	Configuration
VDI	Citrix XenDesktop, VMware Horizon/View	Config-1
SERVER	SQL Server, Exchange Mail Server	Config-2 Config-3
BIG DATA	Splunk, Hadoop	Config-3
OTHERS	IT Infrastructure, Custom applications	Mix

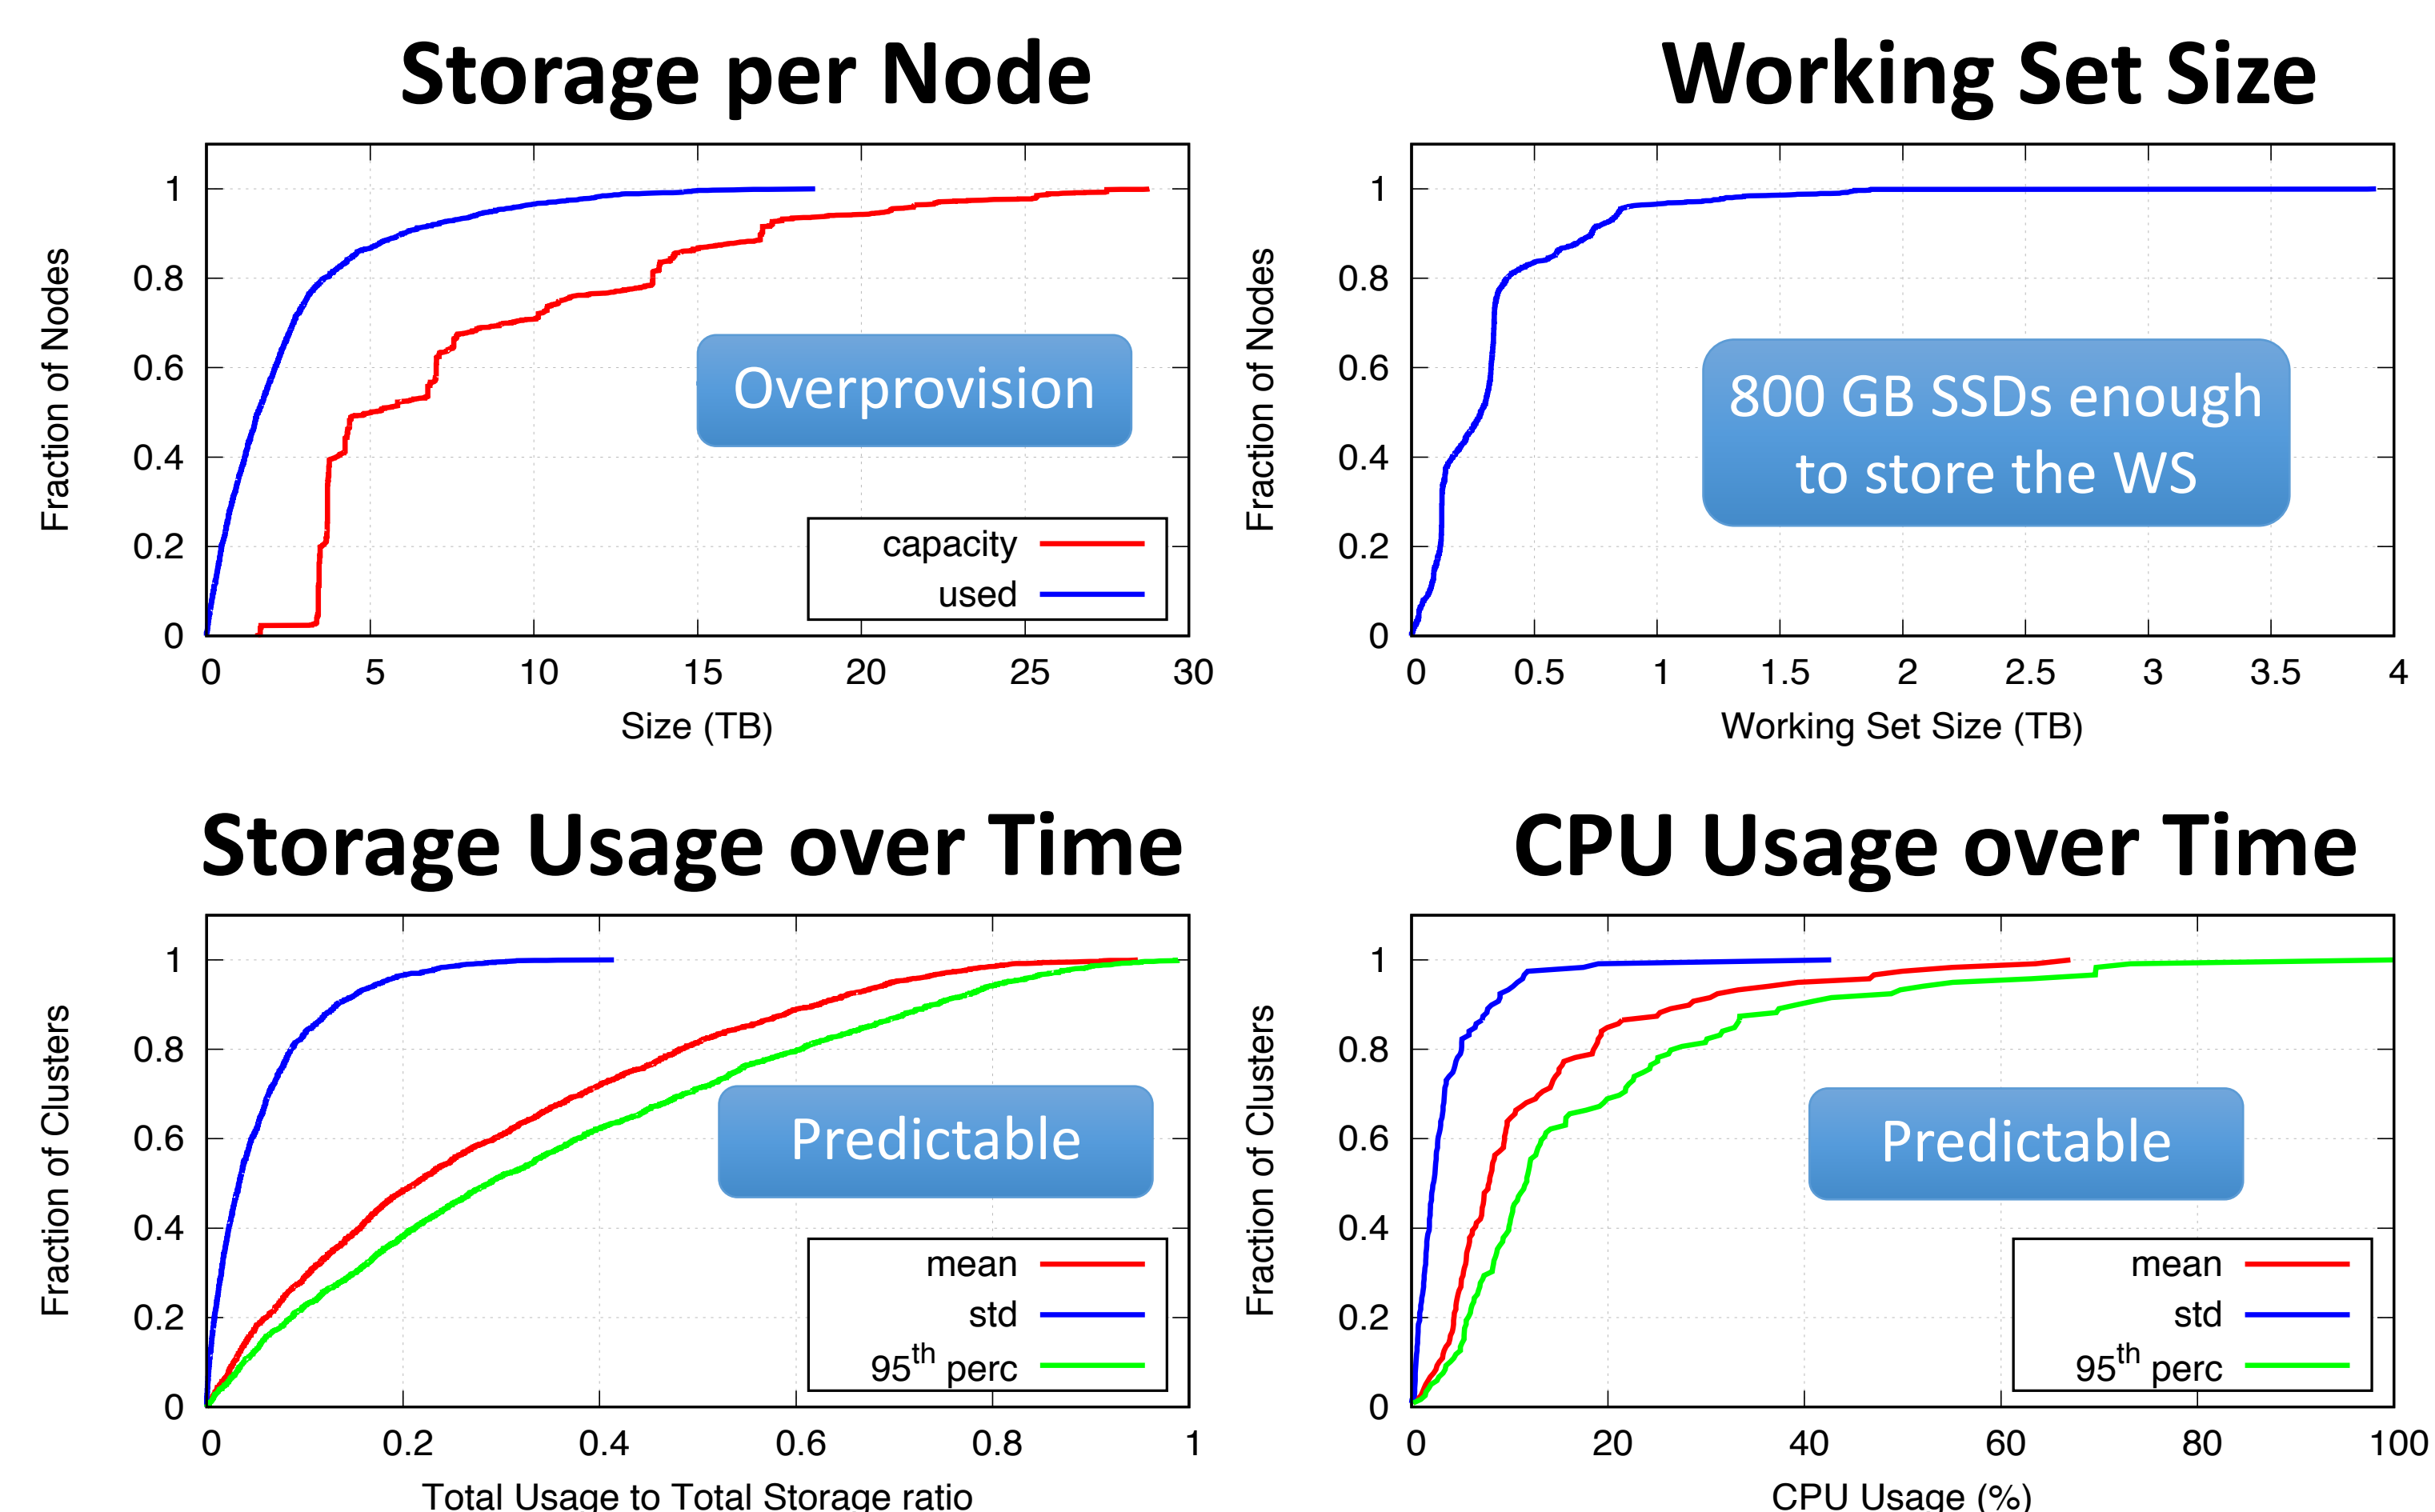
Distribution of VMs per Node



Failure Analysis



Workload Characteristics



Modeling based on Measurements

Durability: estimate probability of data loss

$$\Delta t = \frac{d}{(n-1)r}$$

Time to create new replica

$$P(-Q, \Delta t) \leq (1 - p(\Delta t))^n + np(\Delta t)(1 - p(\Delta t))^{n-1} + (1 - p(\Delta t))^{n-1}$$

Prob of no data loss in Δt (RF2)

$$P_{durability} = P(-Q, \Delta t)^{N(\Delta t)}$$

Prob of no data loss in a year

Cluster Growth: what drives cluster growth?

- **Binary classification** problem, Logistic Regression with L1
- **200 clusters** over a period of **8 months** (15K examples)
- Results:
 1. **Size of cluster**
 2. **Storage needs**
 3. **VM needs**

Conclusions

- Large-Scale Measurement Study of Private Clouds
 - **Lower hardware failure rates**
 - **Nodes overprovisioned**
 - **Stable storage and CPU usage**
- Modeling based on the Measurements
 - **Each extra replica provides substantial durability improvements**
 - **Storage needs drive growth more than compute**