

Big Data for Better Datacenters



Balaji Parimi

@vimAPIGuru

Krishna Raj Raja

@esxtopGuru

Strata 2014, Santa Clara

World Today

Forbes

New Posts

Most Popular

Lists

12 Stocks to BUY for 2014



Kashmir Hill, Forbes Staff

Welcome to The Non-Go Private Game where technology & privacy collide

Follow (1,694)

Follow

243k

TECH | 2/26/2012 @ 11:02AM | 2,333,104 views

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

Facebook knows when you'll break up



By Doug Gross, CNN

November 4, 2010 1:32 p.m. EDT | Filed under: Social Media

David McCandless: The beauty of data visualization

TEDtalksDirector

779 videos

Subscribe



THE VERGE

The first guide for new Transformers 3: Dark of the Moon

LONGFORM . VIDEO . REVIEWS . TECH . SCIENCE . CULTURE . DESIGN . BUSINESS . US & WORLD . FORUM

NEWSLETTER

Bloggers get the same libel protection as traditional journalists, federal...

NOTES

Lightroom for iPad briefly appears on Ad...

Amazon plans to ship your packages before you even buy them

By Erez Shalev on January 26, 2013



Emerging Technology From the arXiv

October 9, 2013

Twitter Datastream Used to Predict Flu Outbreaks

The rate at which people post flu-related tweets could become a powerful tool in the battle to spot epidemics earlier, say computer scientists.



Back in 2006, Google launched its now famous flu trends website. It works on the hypothesis that people make more flu-related search queries when they are suffering from the illness than when they are healthy. So

counting the number of flu-related search queries in a given country gives a good indication of how the virus is spreading.

The Irony

October 25, 2011, 2:45 PM

Article Tools     

Target.com suffers another site outage, this one for more than two hours

The crash is at least the third outage for the newly redesigned e-commerce site.


Thad Rueter
Senior Editor

Topics: [Alertbot](#), [Amazon](#), [AT&T](#), [12 Stocks to BUY for 2014](#), [site crash](#), [site outage](#), [site perform](#), [Eastman](#), [Target](#), [Target.com](#)

The gods of online retail continue to punish Target Corp., as the chain's recently [redesigned e-commerce site](#) crashed again today.

The outage marks at least the third time the new site has crashed since its August debut. The site also crashed less than two weeks ago, after which [Steve Eastman](#), president of Target.com, stepped down to "pursue other opportunities," according to a statement from his former employer.

A Target spokesman confirmed the site went down today at 11:50 a.m. Central time. "We are working diligently to bring the site back online and to ensure that it is operating efficiently," the spokesman says. "Target strives to deliver an exceptional guest experience in our stores and online, and we regret the inconvenience to



The canine mascot in

Forbes

New Posts
+2 posts this hour

Most Popular
Why Companies Lose Talent


Lists
Most Promising Companies

Vic
30 U




Reuven Cohen, Contributor

I focus on disruptive trends in technology and cloud computing.

+ Follow (117)   525

TECH | 7/02/2012 @ 7:37AM | 8,376 views

Cloud Computing Forecast: Cloudy With A Chance of Fail



9 comments, 5 called-out

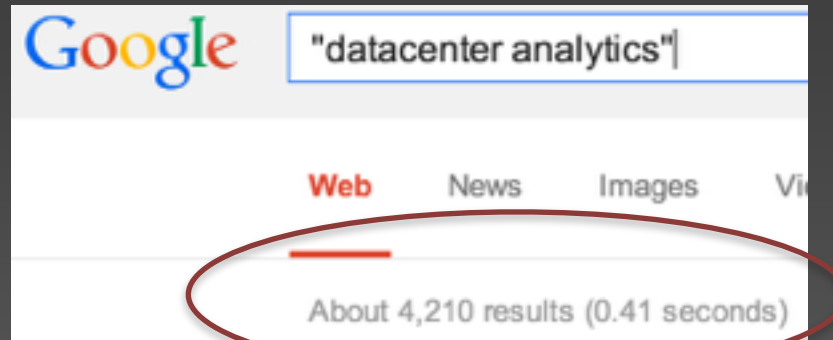
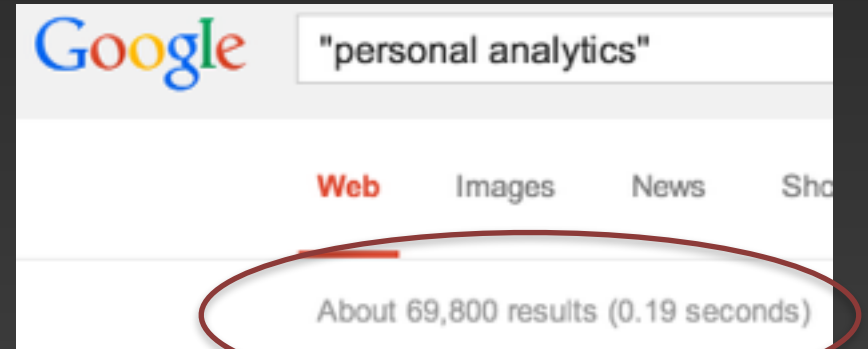
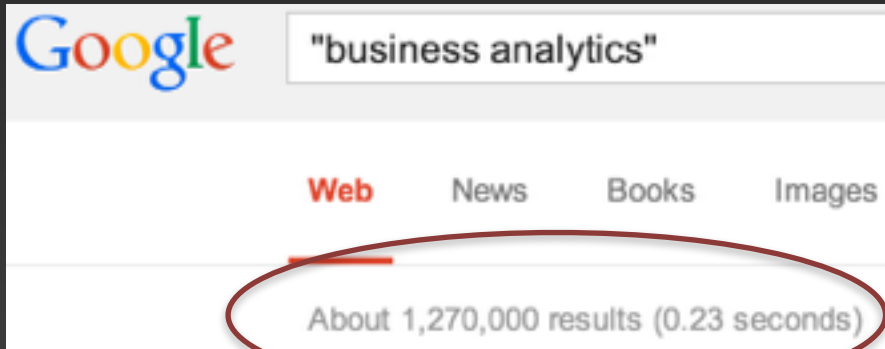
+ Comment Now

+ Follow Comments

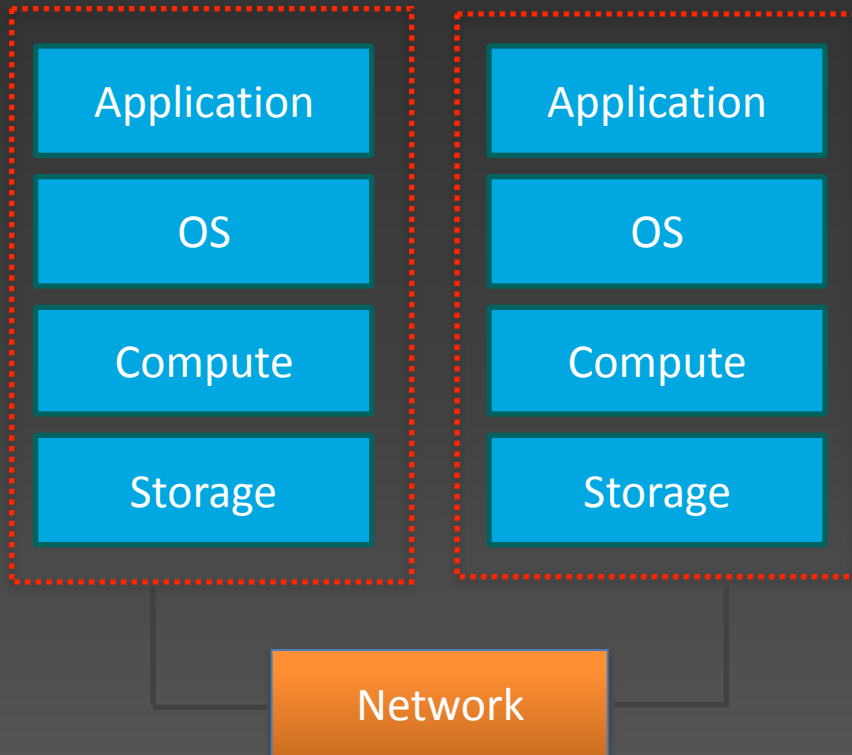
Amazon's Cloud service is having a bad a couple weeks. For the second time in as many weeks Amazon's East Coast cloud crashed during a severe storm that left 1.3 million in the [Washington D.C.](#) area without power. The outage brought down numerous high profile web sites hosted on Amazon including [Netflix](#), [Instagram](#), [Pinterest](#), and [Heroku](#). Making things worse was the fact that other cloud services hosted in the area experienced no downtime.



Analytics



1990s

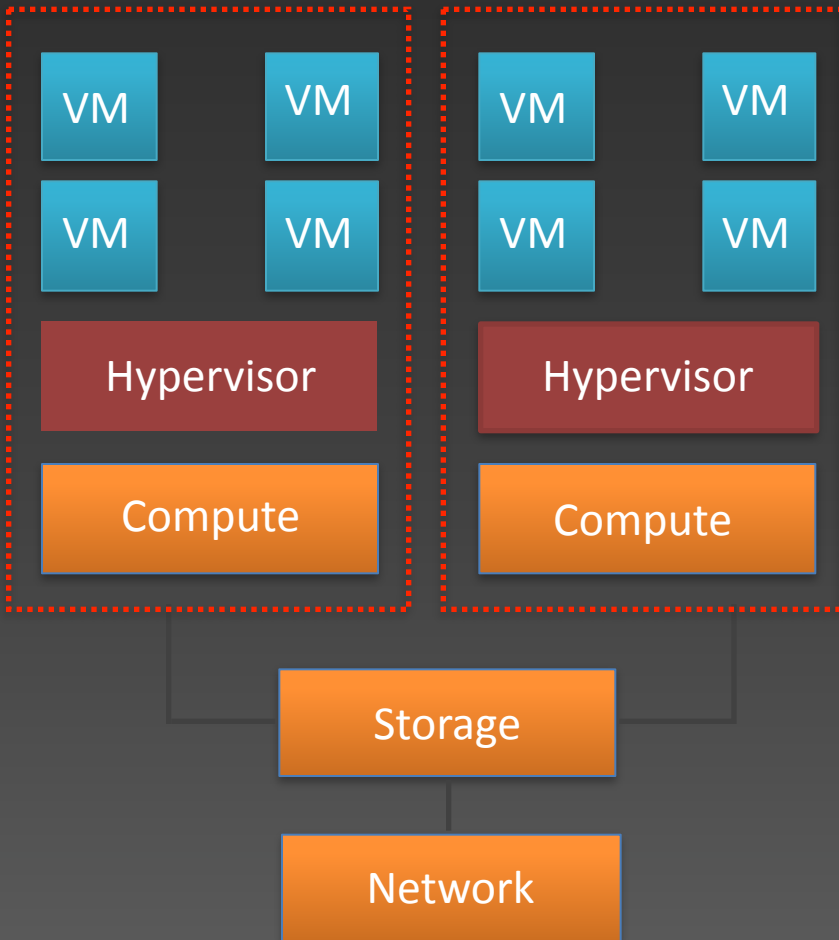


Networking becomes mainstream

No effective means to pool storage and compute capacity

100s of silos to manage

2000s



Virtualization becomes mainstream

Improved resource pooling

Increased utilization


2010s

Cloud has become mainstream

Centralized management

70% of all the x86 OS instances run in virtual machines, 80% by 2016

- Gartner (2012)



Data Centers Have Evolved But
IT Operations Has Not

IT Operations Today

1503Hrs 09/30/02 Data Center

Net IQ - Ops Monitoring

Repository View	Computers Discovered	Running Jobs	Computers w/ Severe Events	Open Severe	Open Events	Event Cross Section			
Total	211	3002	0	1	35	11 Severe	12 Warning	16 Informational	14 Diagnostic
FRANTMON1-QDB	211	3002	0	1	35	1 Severe	11 Warning	11 Informational	12 Diagnostic

Net IQ - Details of Event for Operations

Count	Severity	Event ID	Job ID	Computer	Knowledge Script	Last Occurred	Last Message	Status
1	9	93736	47636	FRAMATCH1	Staples CMThermalStatus	09/29/2002 5:51:41 PM	SNMP Not Functional	Open
1	9	93736	47636	FRAMATCH1	Staples CMThermalStatus	09/29/2002 5:36:41 PM	SNMP Not Functional	Open

Robot Console - Messages

Status	ID	Message	Type	Sev	Time	Date	Queue
WAITING	CPA3394	Load form type "Ship Label" device writer STRP337A (0 B I H R C)	INFO	99	10:25	09/25/02	QSVSOPR
OUTSIDE	CPA3394	Load form type "SH-ELF_TAO" device writer STRP308 (0 B I H R C)	INFO	99	10:26	09/25/02	QSVSOPR
WAITING	CPA3394	Load form type "Del_CAFSGO" device writer STRP307 (0 B I H R C)	INFO	99	10:12	09/25/02	QSVSOPR
OUTSIDE	CPA3394	Load form type "SH-ELF_TAO" device writer STRP308 (0 B I H R C)	INFO	99	10:10	09/25/02	QSVSOPR
OUTSIDE	CPA3394	Load form type "SH-ELF_TAO" device writer STRP308 (0 B I H R C)	INFO	99	10:10	09/25/02	QSVSOPR
OUTSIDE	CPA3394	Load form type "SH-ELF_TAO" device writer STRP308 (0 B I H R C)	INFO	99	10:09	09/25/02	QSVSOPR
OUTSIDE	CPA3394	Load form type "SH-ELF_TAO" device writer STRP308 (0 B I H R C)	INFO	99	10:08	09/25/02	QSVSOPR

OVO Ops Monitoring

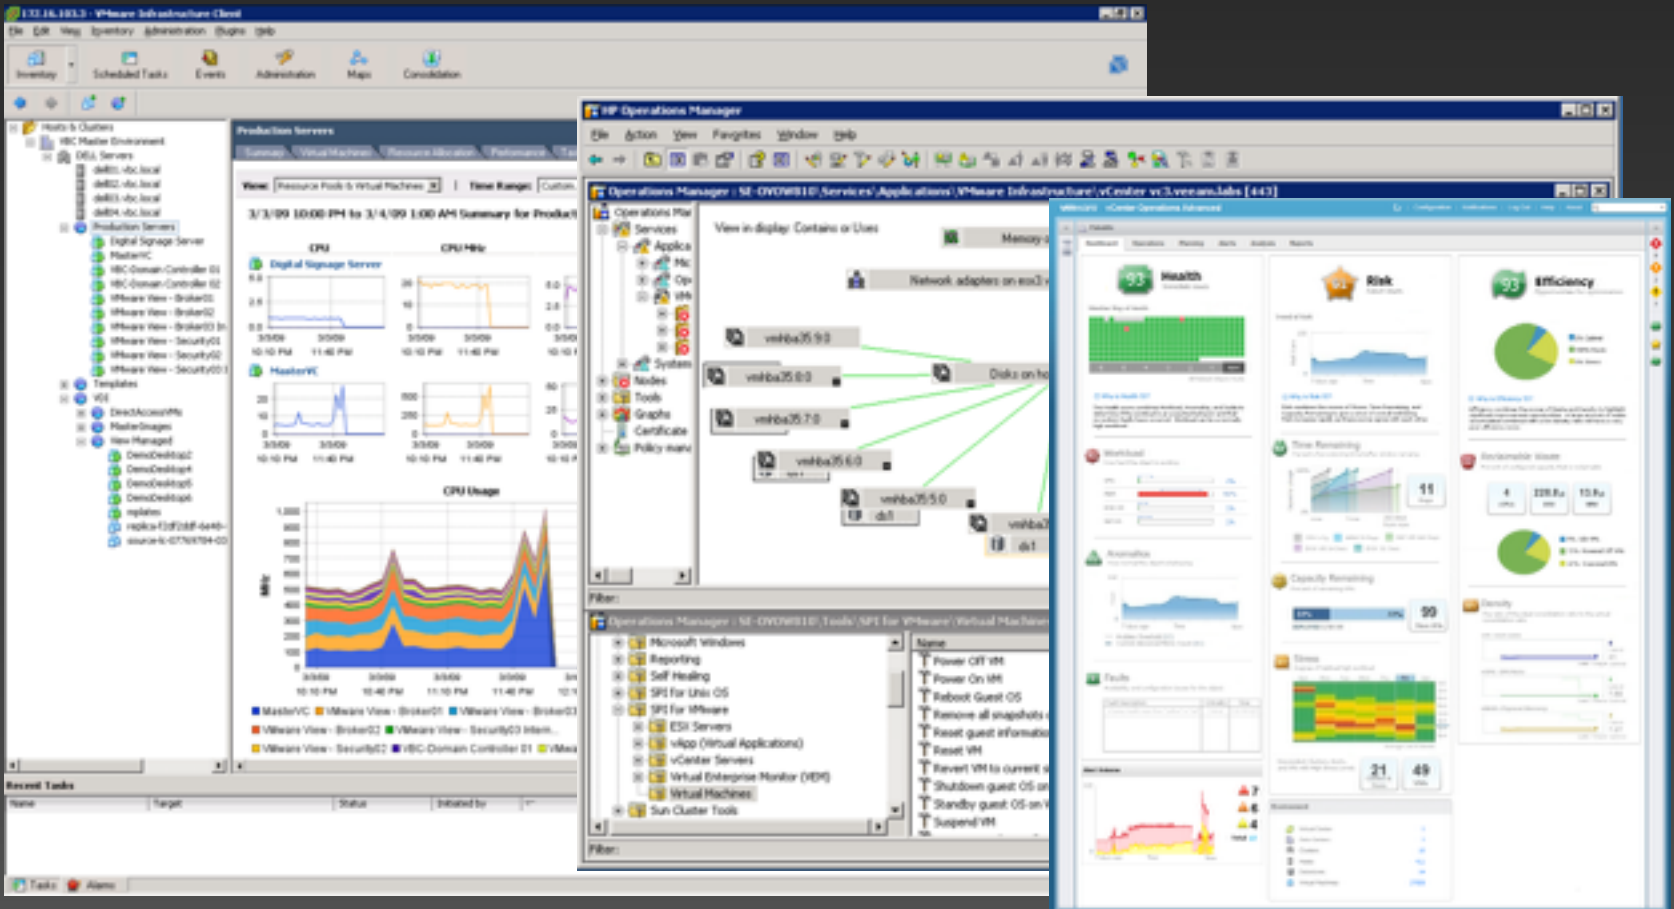
Sev.	SVIRONE	Date	Time	Mode	Application	MsgGroup	Object	Message Text
Warn	--s--	09/25/02	10:58:31	prelude_at	syslogd	Storage	LVM	Sep 25 10:52:17 prelude vmunix: msgcnt 59 vxf
Crit	--c--	09/25/02	11:00:59	viper_stap	alarmgen	Performan	Disk	Disk Bottleneck - Begin Alert -> probability100
Maj	--m--	09/25/02	11:02:23	corvette_s	disk_mon	Staples_U	Root_Slg	/ is at 99.00% (exceeds or equal to threshold in
Warn	---	09/25/03	11:04:57	topaz_stap	OmniBack	Staples_U	omni mail	Backup daily-elderado completed with some em
Norm	--n--	09/25/02	11:06:41	viper_stap	alarmgen	Performan	Disk	Disk Bottleneck - End Alert

Data Center Updates

MQseries upgrade begins this evening 19:00 hours. Questions, concerns - please see Jim Brammer x 38623

Retrofitted

IT Operations Today



Paralyzing

IT Operations Today



Lot of Guesswork!

IT Operations Should Be Data-Driven,
“Intuition and Guesswork”
Should be Replaced By
“Analytics and Predictive Intelligence”

Datacenter Efficiency

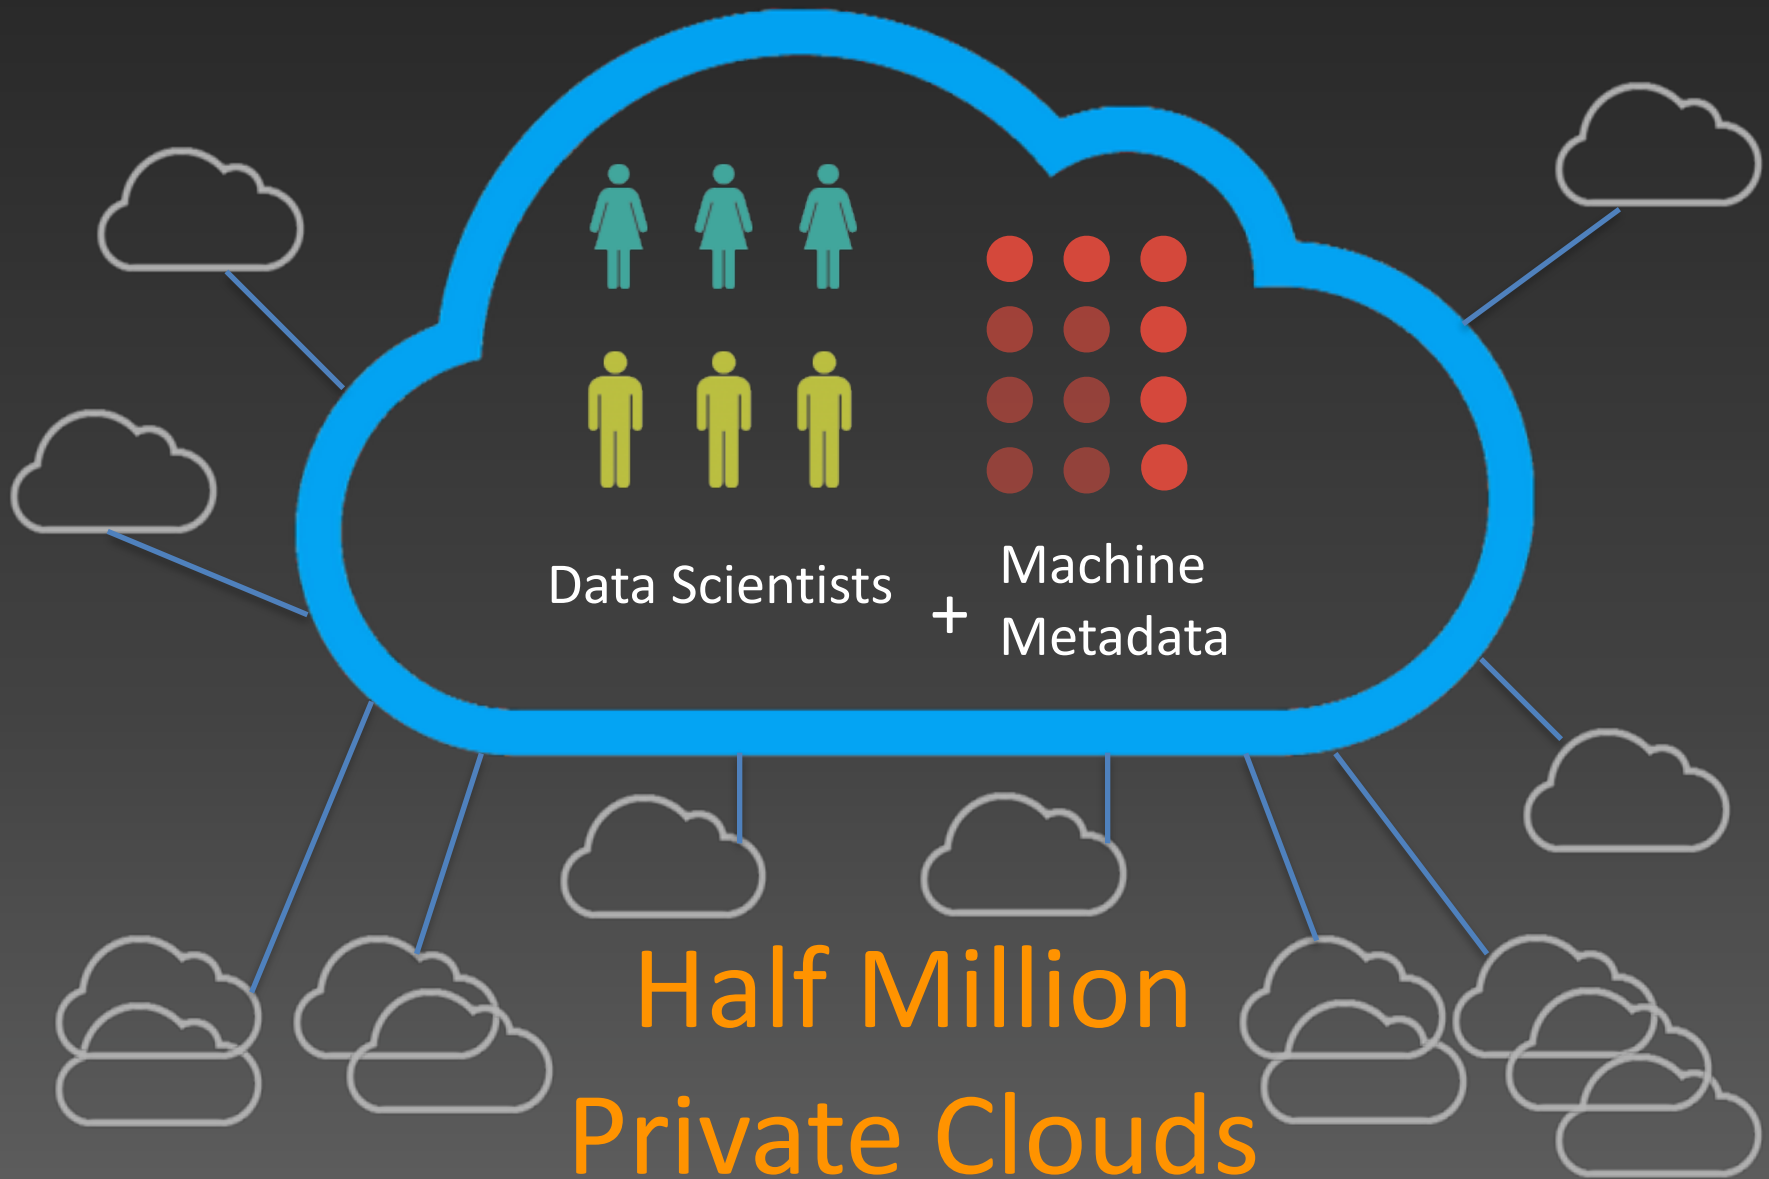
Statistically Significant Machine Metadata Available for Data Scientists



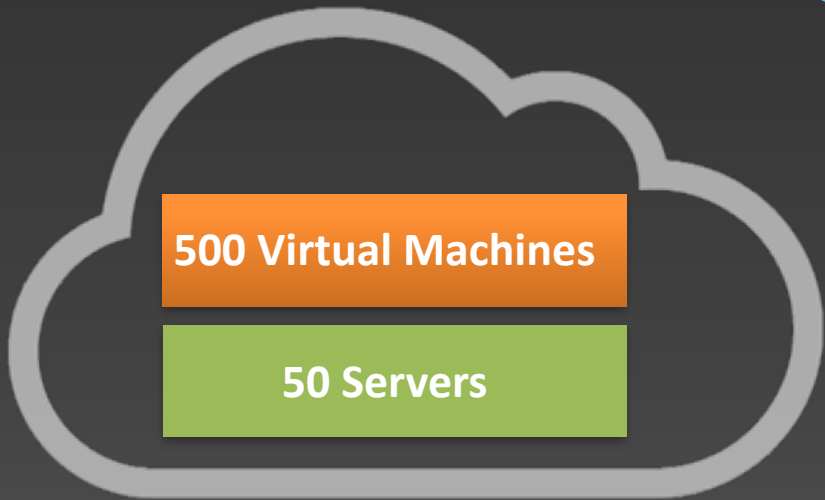
Platform

ice.
pany

Can Private Clouds Match The Efficiency?



Operational Metadata Per Day



Private Cloud


1 Billion Data Points

25GB

Massive Scale



650 Million Users



60 million tweets




1 Billion Users



2.5 billion pieces of content

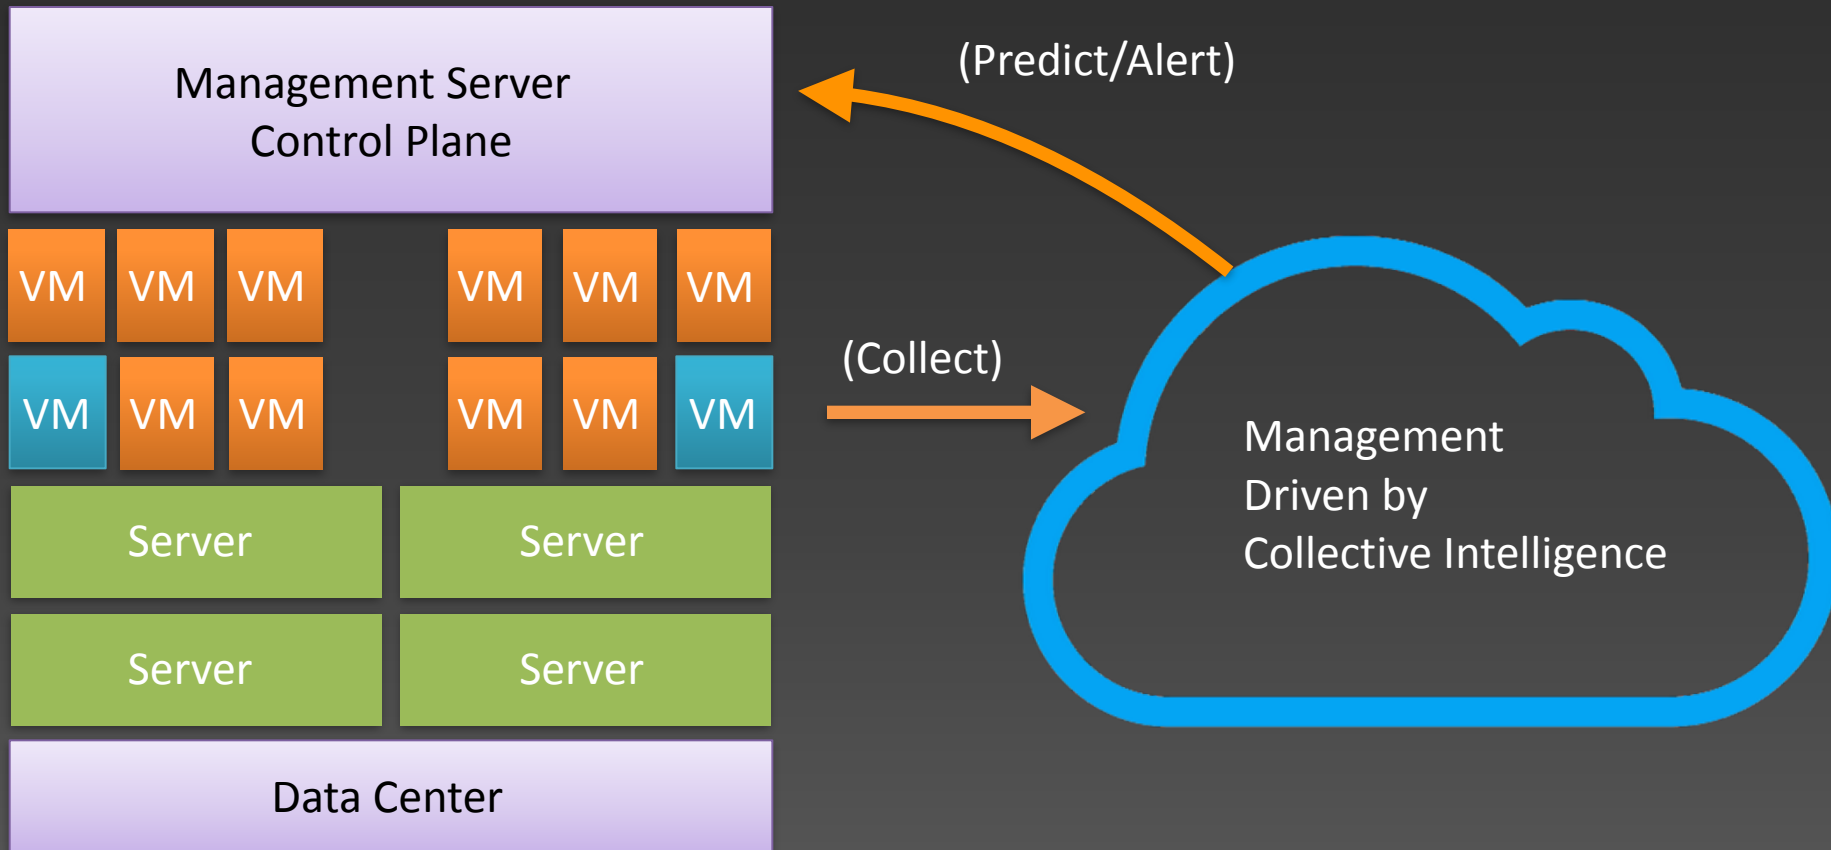
500,000
Datacenters

50 Million VMs

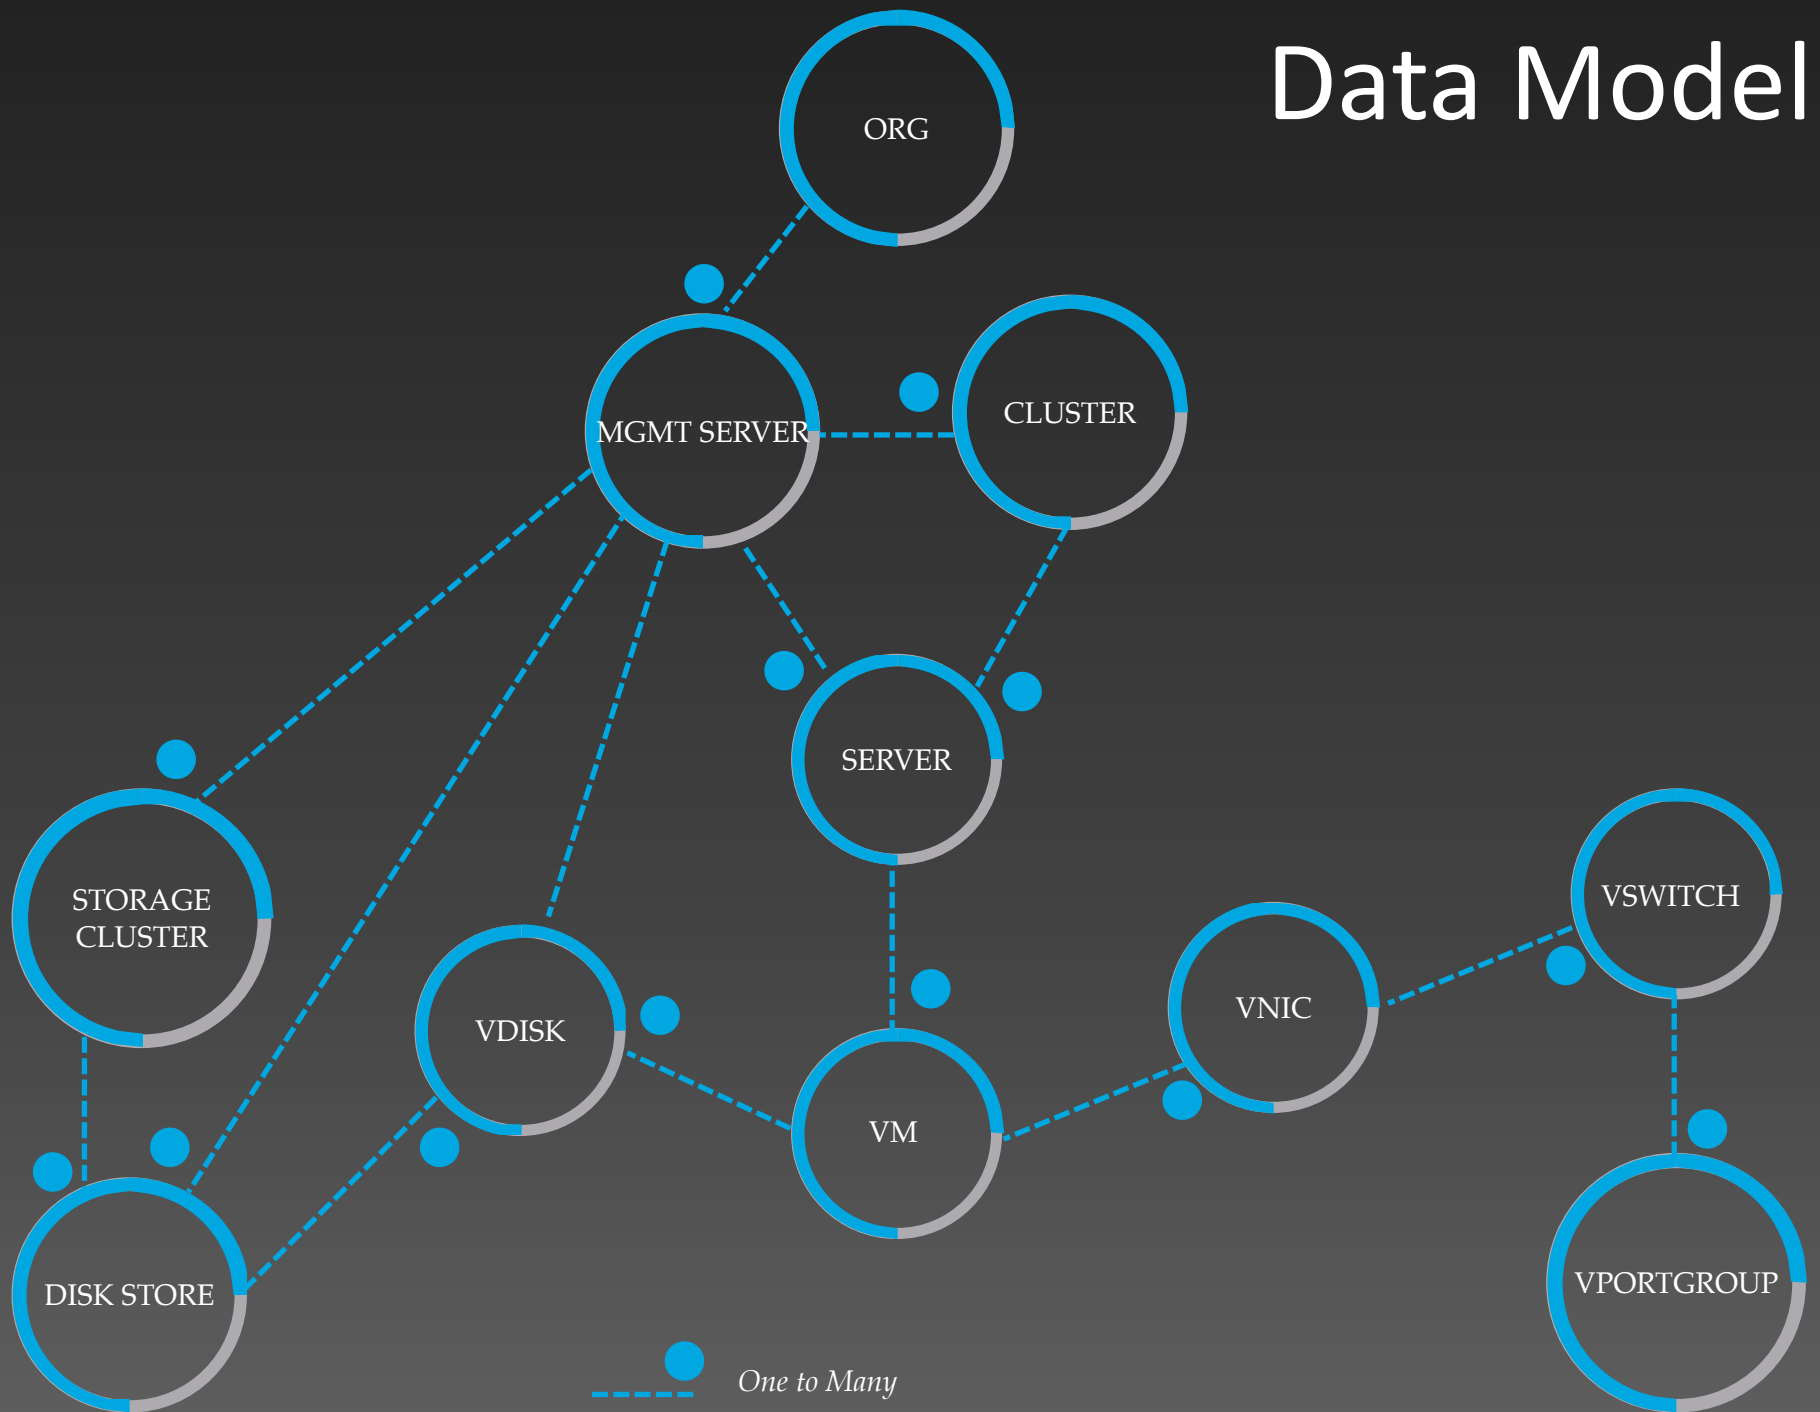



Quadrillion datapoints (10^{15})

Challenges





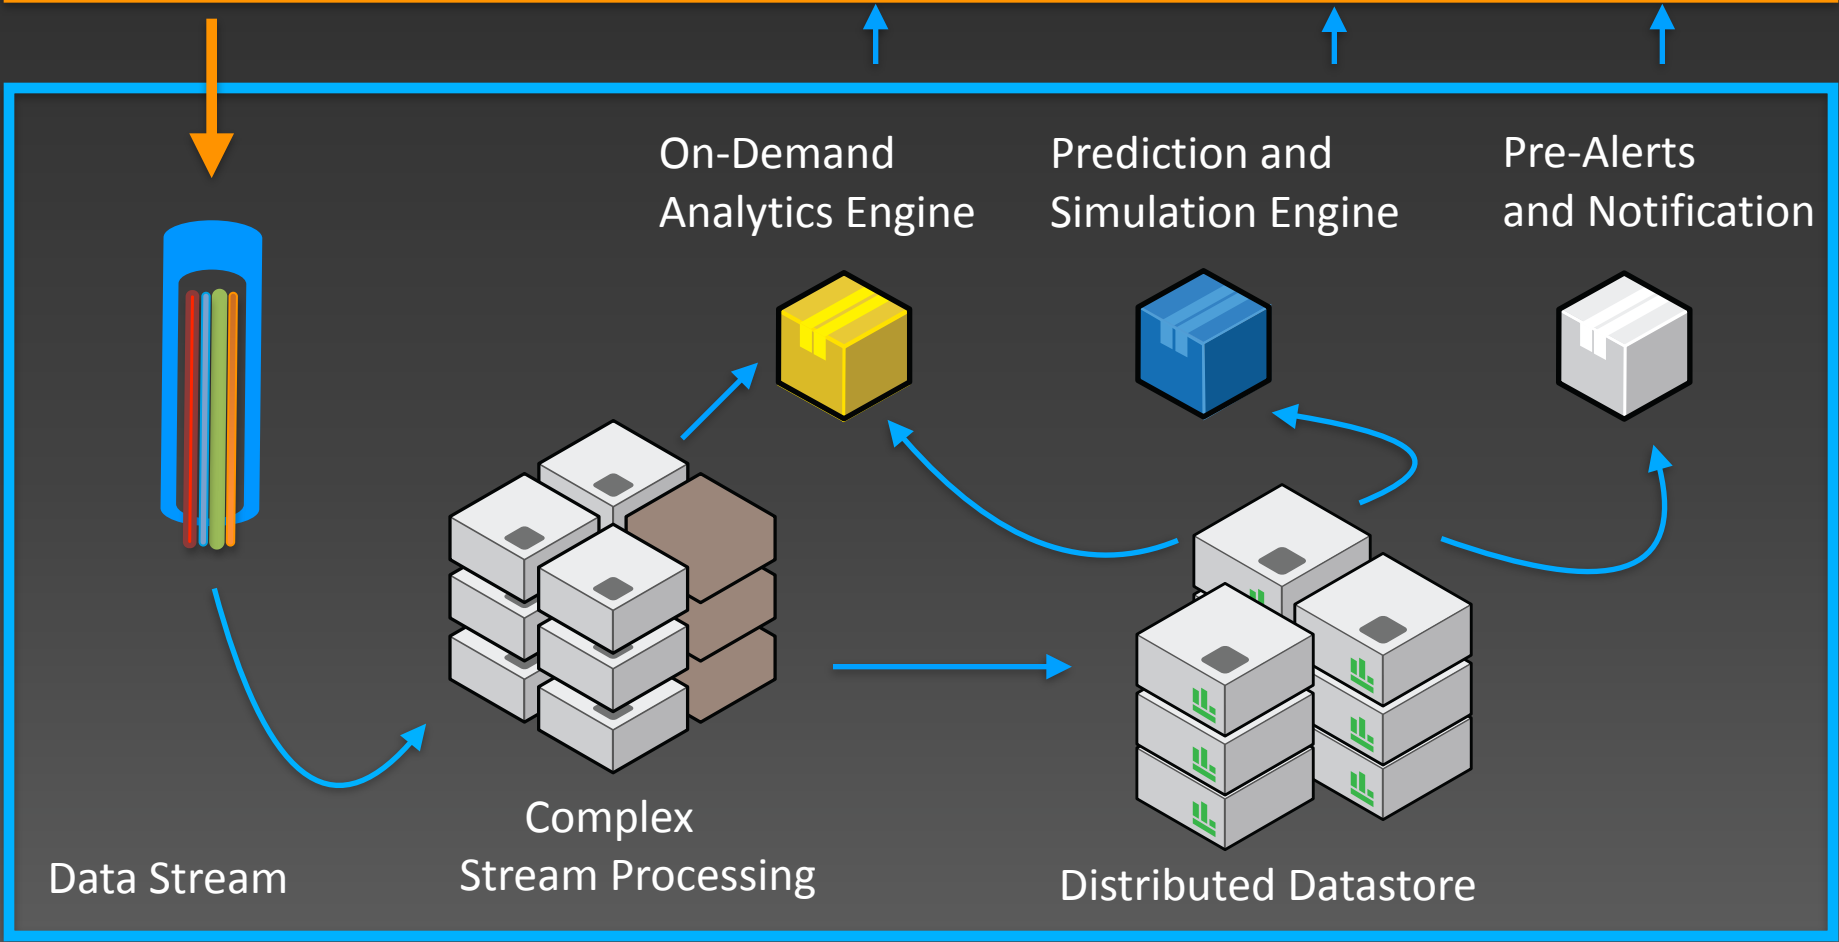

Data Model



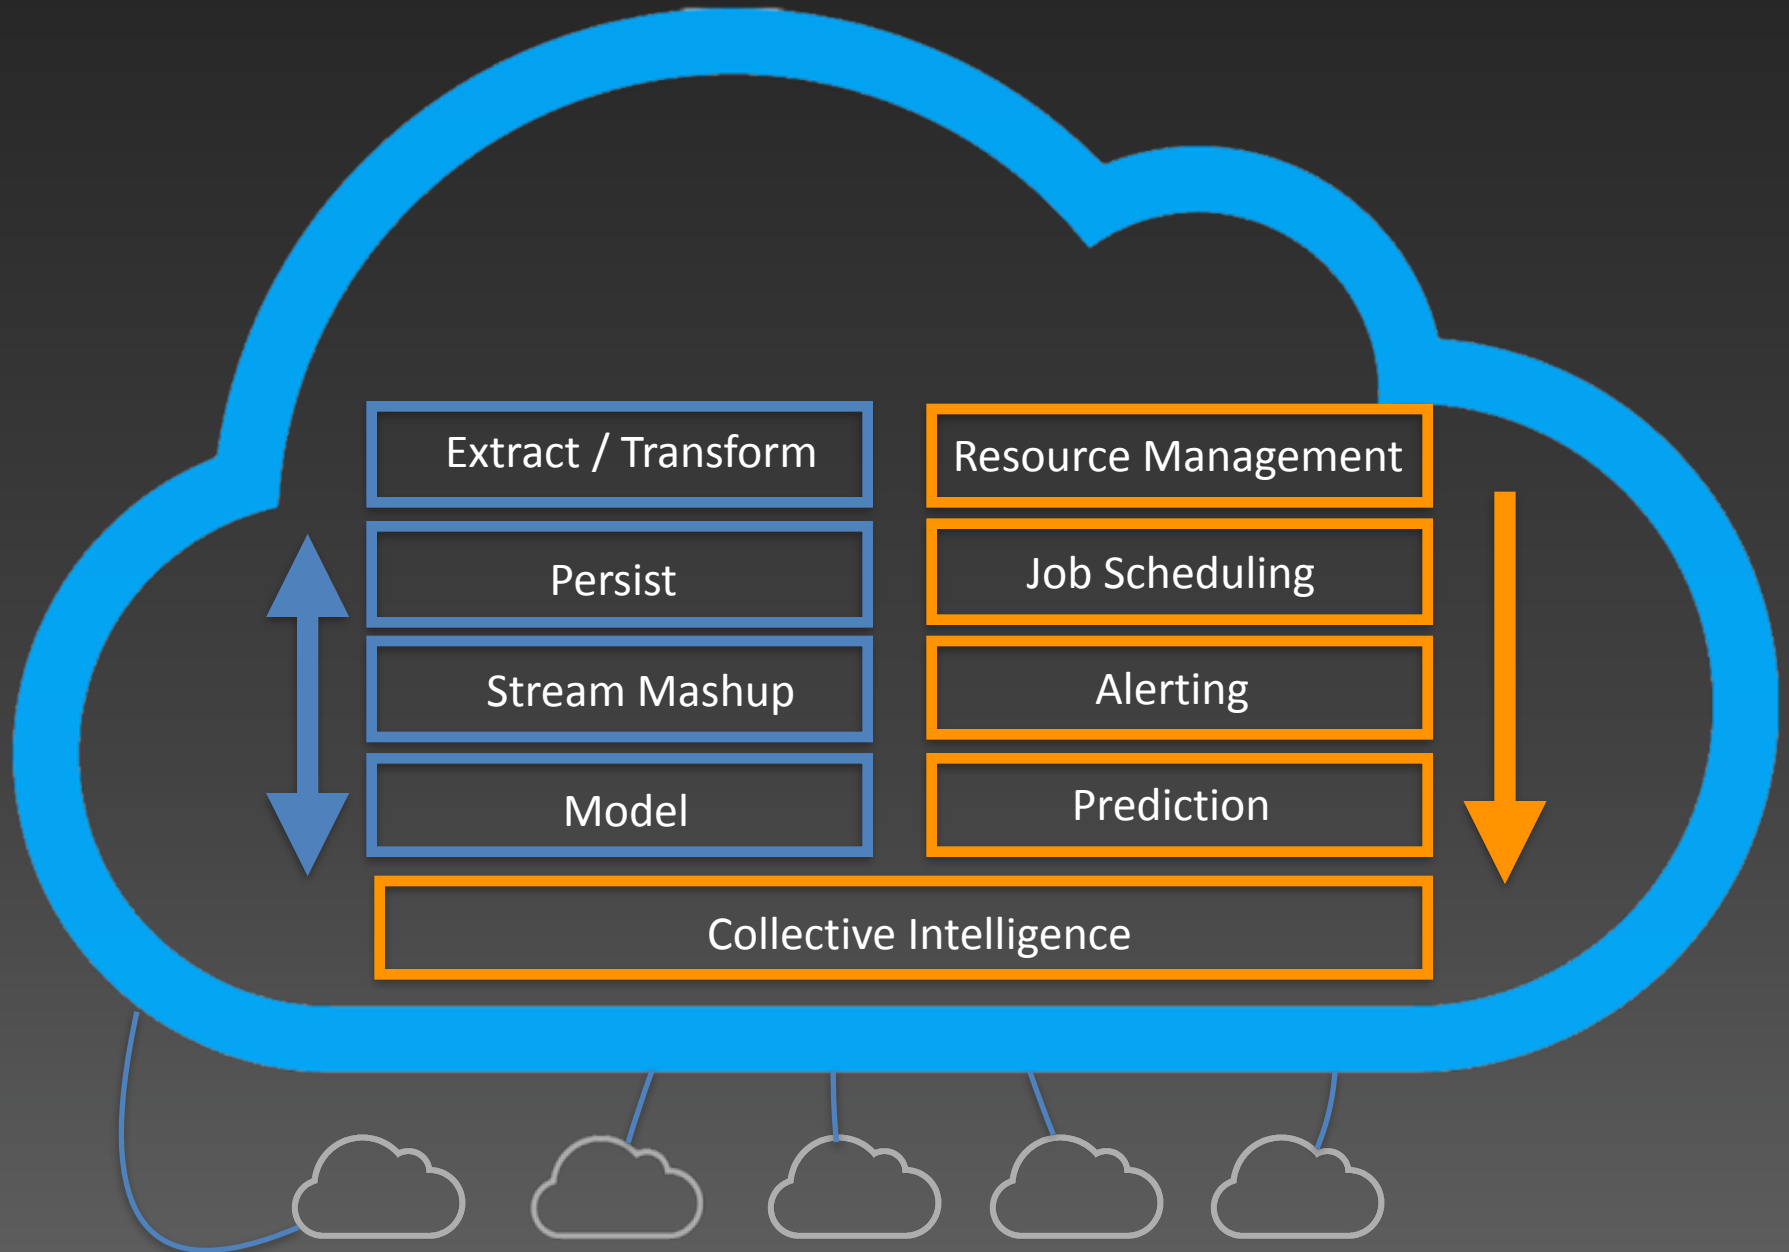


Types of Data Streams



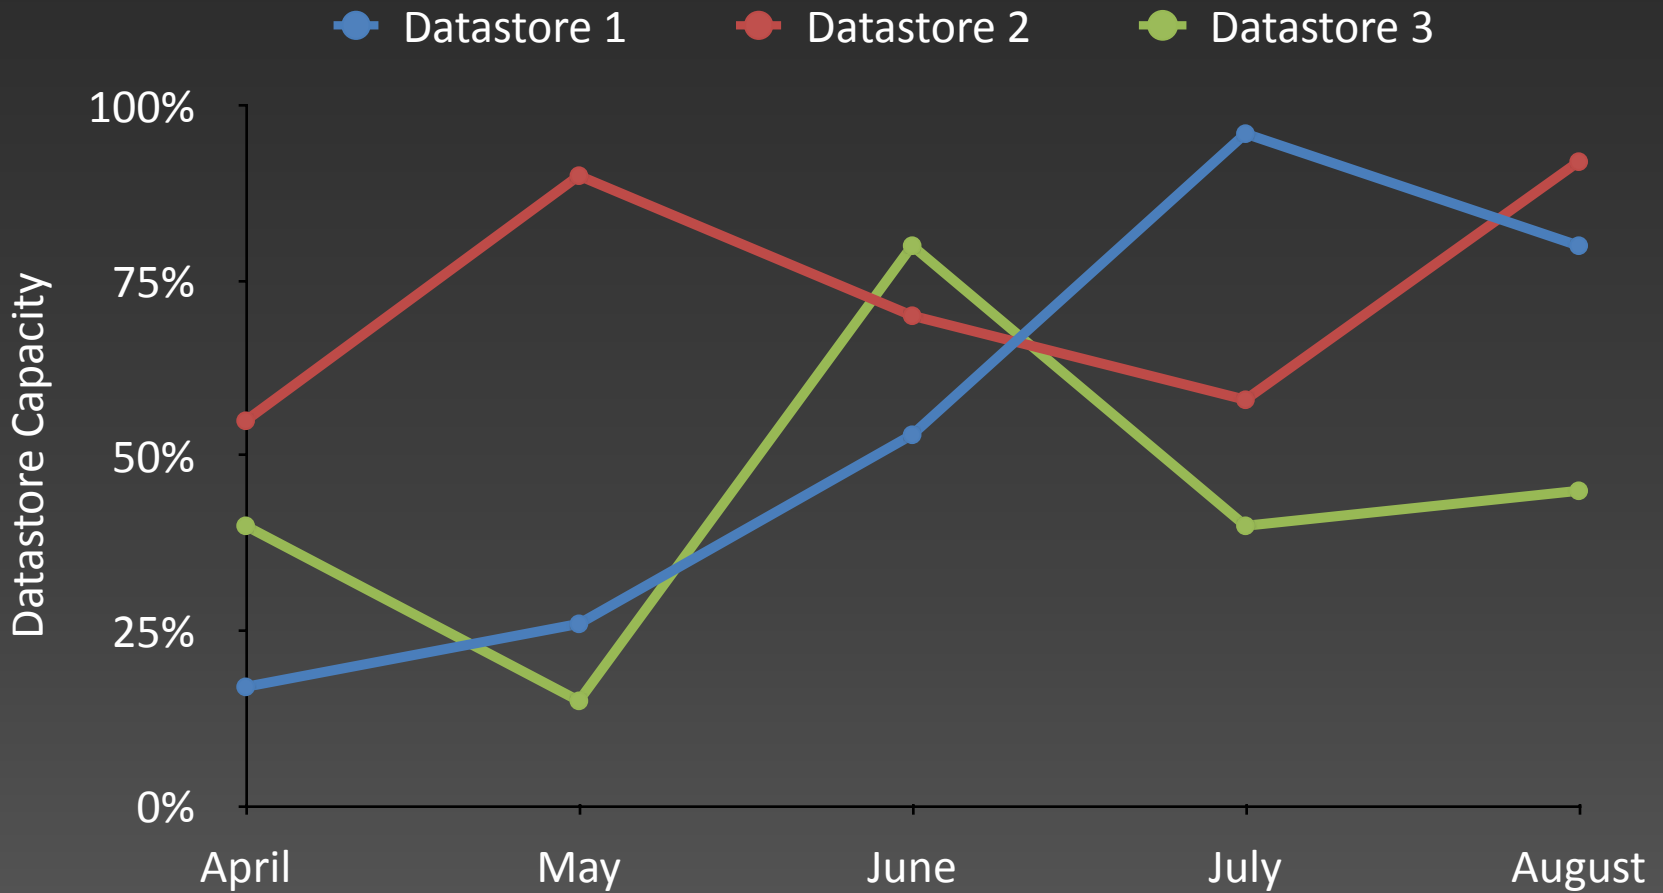


Enabling Private Clouds



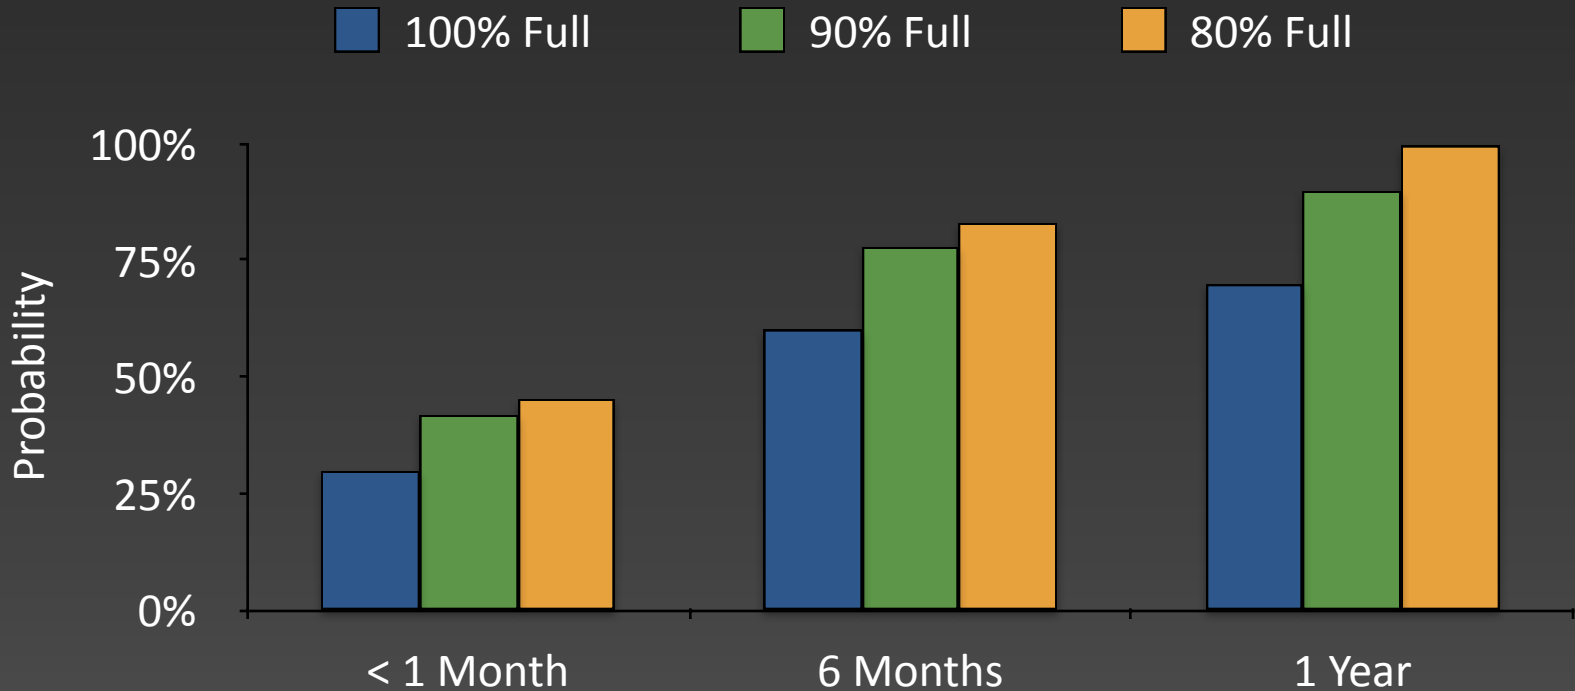
Use Cases

Predicting Out of Disk Space Event



Disk Space Usage is Hard to Predict

Predicting Out of Disk Space



Probability Distribution Using Monte Carlo Simulation

Discovering Causality



Predicting Potential Outage

The screenshot shows the Knowledge Base Advisor interface. At the top, there is a search bar with the text "FILTER BY KEYWORD: Type a keyword(s), VCenter, cluster, host or VM name." Below the search bar, there are several filter buttons: "Symptom: Purple Screen of Death/PSOD", "Custom: tardis.cphy.local", "Custom: psod", and "Custom: tard", along with a "RESET" button. Below the filters, it displays "31 OF 1411 RESULTS" and a severity breakdown: "80 CRITICAL", "341 HIGH", and "598 MEDIUM SEVERITY". A search result is highlighted with a red box, showing the title "Restarting NEC Baseboard Management Controller (BMC) fails with a purple diagnostic screen or ESXi/ESX host becomes unresponsive", source "VMware", updated date "02/04/2014", and applies to "15 Hosts". Below the search result, the category "Troubleshooting" is shown. The "APPLIES TO: 15 HOSTS" text is circled in red. Below this, a table lists the affected hosts.

Knowledge Base Advisor

FILTER BY KEYWORD: Type a keyword(s), VCenter, cluster, host or VM name.

Symptom: Purple Screen of Death/PSOD Custom: tardis.cphy.local Custom: psod Custom: tard RESET

31 OF 1411 RESULTS 80 CRITICAL 341 HIGH 598 MEDIUM SEVERITY

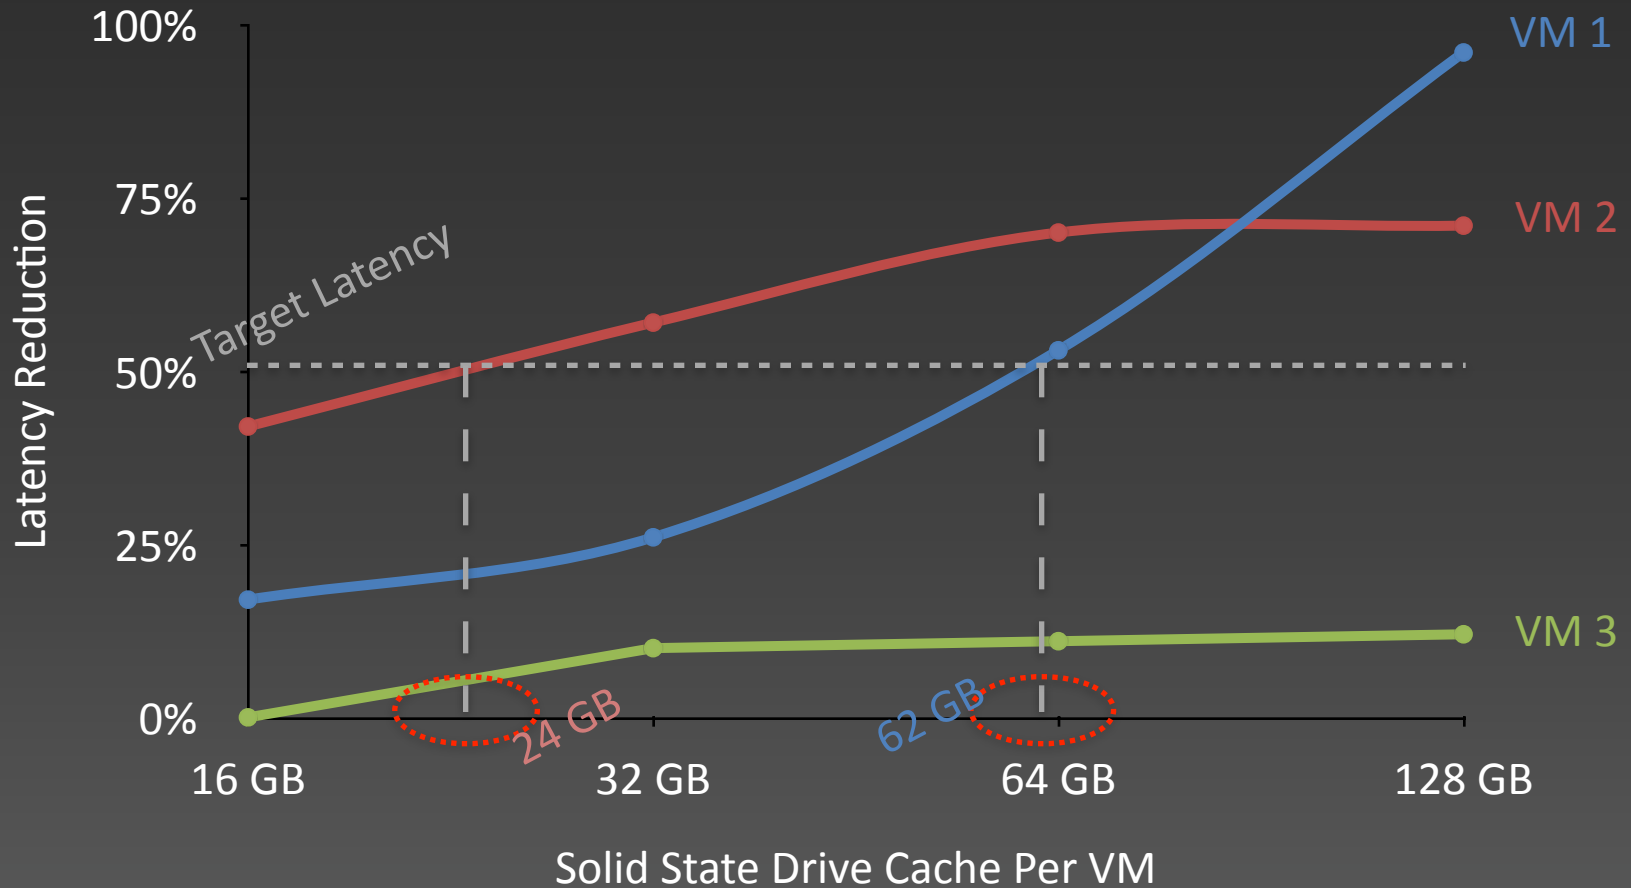
Restarting NEC Baseboard Management Controller (BMC) fails with a purple diagnostic screen or ESXi/ESX host becomes unresponsive
SOURCE: VMware UPDATED: 02/04/2014 APPLIES TO: 15 Hosts SYMPTOM: Purple Screen of Death/PSOD, Compatibility Error

CATEGORY: Troubleshooting

APPLIES TO: 15 HOSTS

VCENTER	CLUSTER	HOST
vc.cphy.local	(Standalone Hosts)	tardis.cphy.local

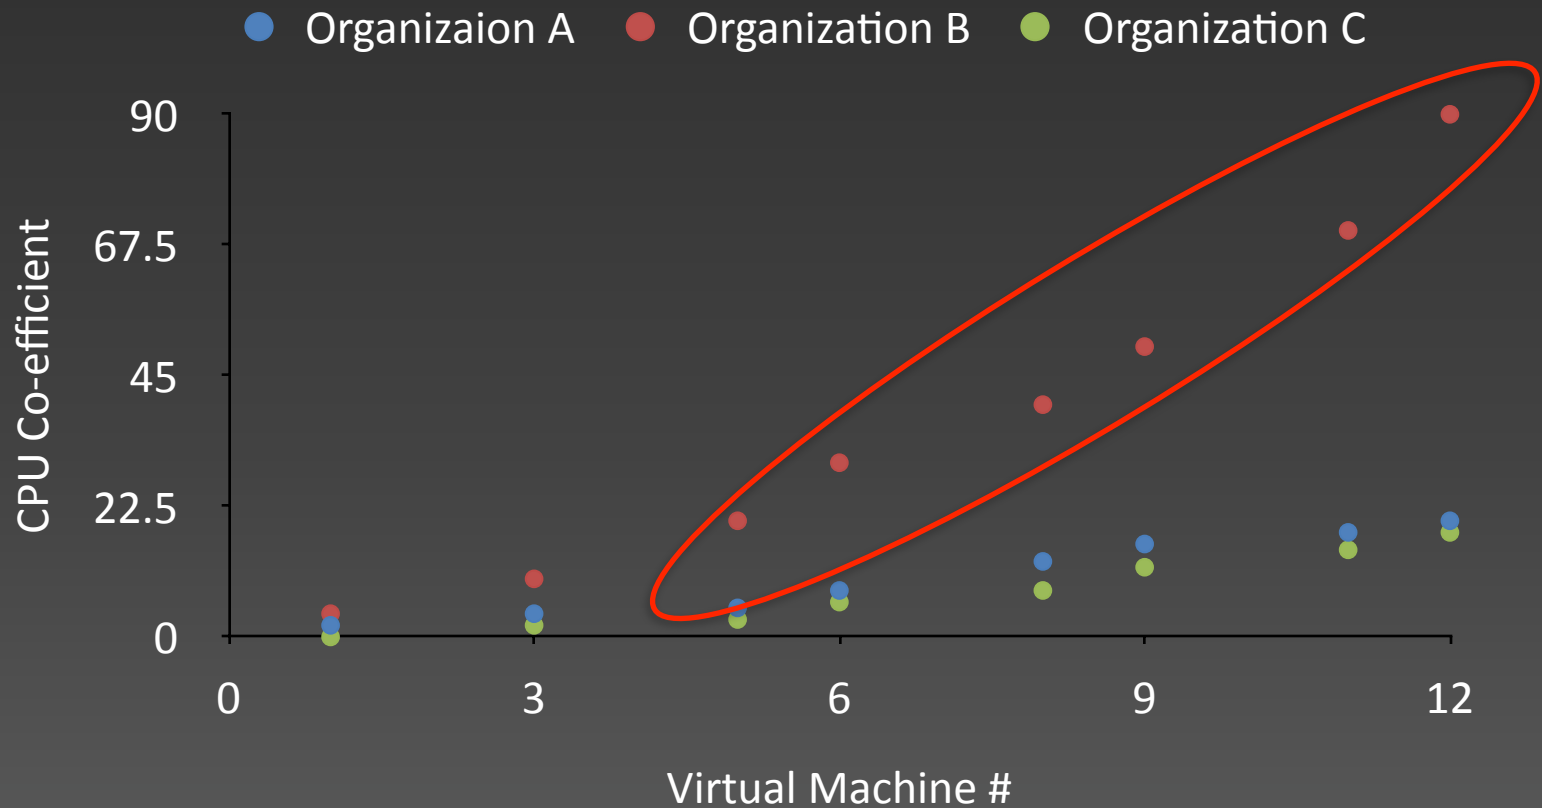
Predicting SSD Performance Impact



Predicting Configuration Outliers

ATTENTION	<ul style="list-style-type: none">● VM AND APPLICATION MONITORING ARE DISABLED <p>IMPORTANCE</p> <p>HIGH</p> <p>RECOMMENDATION</p> <p>For additional VM protection, turning on "VM and Application Monitoring" is recommended.</p>	<p>DESCRIPTION</p> <p>In the case of a Guest OS level failure, HA can restart the failed Guest OS. This feature is called VM Monitoring, but is sometimes also referred to as VM-HA.</p>
NOTEWORTHY	<ul style="list-style-type: none">● MULTIPLE SHARED STORAGE IS NOT PRESENT <p>IMPORTANCE</p> <p>MODERATE</p> <p>RECOMMENDATION</p> <p>The use of multiple shared datastores are recommended.</p>	<p>DESCRIPTION</p> <p>Datastore heartbeating requires a minimum of two datastores to function.</p>

Predicting Performance Outliers



Some Closing Thoughts



Towards a New World



Thank You



Balaji Parimi

@vimAPIGuru

Krishna Raj Raja

@esxtopGuru