

Serverless Computing

(Lecture 14, cs262a)

Ali Ghodsi and Ion Stoica,
UC Berkeley
March 7, 2018

Today's Papers

Serverless Computing: Current Trends and Open

Problems, Ioana Baldini, Paul Castro, Kerry Chang, Perry Cheng, Stephen Fink, Vatche Ishakian, Nick Mitchell, Vinod Muthusamy, Rodric Rabbah, Aleksander Slominski, Philippe Suter (<https://arxiv.org/abs/1706.03178>)

Serverless Computation with OpenLambda

Authors: Scott Hendrickson, Stephen Sturdevant, and Tyler Harter, Venkateshwaran Venkataramani, Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseauz (<https://www.usenix.org/node/196323>)

Why care?

Rapidly growing in popularity

Change the way we write applications and expose new challenges

“The future of AWS”

– Marvin Theimer,
Distinguished Engineer at AWS



EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 [RDS](#)

Region: [US East \(N. Virginia\) -](#) Cost: [Hourly -](#) Reserved: [1 yr - No upfront -](#) Columns - [Compare Selected](#) [Clear Filters](#)

Filter: Min Memory (GB): Compute Units: Storage (GB):

Name	API Name	Memory	Compute Units (ECU)	vCPUs	Storage	Arch	Network Performance	EBS Optimized: Max Bandwidth	VPC Only	Linux On Demand cost	Linux Reserved cost	Windows On Demand cost	Windows Reserved cost
Cluster Compute Eight Extra Large	cc2.8xlarge	60.5 GB	88 units	32 vCPUs	3360.0 GB (4 * 840.0 GB)	64-bit	10 Gigabit	N/A	No	\$2.000 hourly	\$1.090 hourly	\$2.570 hourly	\$1.336 hourly
Cluster GPU Quadruple Extra Large	cg1.4xlarge	22.5 GB	33.5 units	16 vCPUs	1680.0 GB (2 * 840.0 GB)	64-bit	10 Gigabit	N/A	No	\$2.100 hourly	unavailable	\$2.600 hourly	unavailable
T2 Nano	t2.nano	0.5 GB	Burstable	1 vCPUs	0 GB (EBS only)	64-bit	Low	N/A	Yes	\$0.006 hourly	\$0.005 hourly	\$0.009 hourly	\$0.007 hourly
T2 Micro	t2.micro	1.0 GB	Burstable	1 vCPUs	0 GB (EBS only)	32/64-bit	Low to Moderate	N/A	Yes	\$0.013 hourly	\$0.009 hourly	\$0.018 hourly	\$0.014 hourly
T2 Small	t2.small	2.0 GB	Burstable	1 vCPUs	0 GB (EBS only)	32/64-bit	Low to Moderate	N/A	Yes	\$0.026 hourly	\$0.018 hourly	\$0.036 hourly	\$0.032 hourly
T2 Medium	t2.medium	4.0 GB	Burstable	2 vCPUs	0 GB (EBS only)	64-bit	Low to Moderate	N/A	Yes	\$0.052 hourly	\$0.036 hourly	\$0.072 hourly	\$0.062 hourly
T2 Large	t2.large	8.0 GB	Burstable	2 vCPUs	0 GB (EBS only)	64-bit	Low to Moderate	N/A	Yes	\$0.104 hourly	\$0.072 hourly	\$0.134 hourly	\$0.106 hourly
M4 Large	m4.large	8.0 GB	6.5 units	2 vCPUs	0 GB (EBS only)	64-bit	Moderate	450.0 Mbps	Yes	\$0.120 hourly	\$0.083 hourly	\$0.246 hourly	\$0.184 hourly
M4 Extra Large	m4.xlarge	16.0 GB	13 units	4 vCPUs	0 GB (EBS only)	64-bit	High	750.0 Mbps	Yes	\$0.239 hourly	\$0.164 hourly	\$0.491 hourly	\$0.366 hourly
M4 Double Extra Large	m4.2xlarge	32.0 GB	26 units	8 vCPUs	0 GB (EBS only)	64-bit	High	1000.0 Mbps	Yes	\$0.479 hourly	\$0.329 hourly	\$0.983 hourly	\$0.735 hourly
M4 Quadruple Extra Large	m4.4xlarge	64.0 GB	53.5 units	16 vCPUs	0 GB (EBS only)	64-bit	High	2000.0 Mbps	Yes	\$0.958 hourly	\$0.658 hourly	\$1.966 hourly	\$1.469 hourly
M4 Deca Extra Large	m4.10xlarge	160.0 GB	124.5 units	40 vCPUs	0 GB (EBS only)	64-bit	10 Gigabit	4000.0 Mbps	Yes	\$2.394 hourly	\$1.645 hourly	\$4.914 hourly	\$3.672 hourly
M4 16xlarge	m4.16xlarge	256.0 GB	188 units	64 vCPUs	0 GB (EBS only)	64-bit	20 Gigabit	10000.0 Mbps	Yes	\$3.830 hourly	\$2.632 hourly	\$7.862 hourly	\$5.875 hourly
C4 High-CPU Large	c4.large	3.75 GB	8 units	2 vCPUs	0 GB (EBS only)	64-bit	Moderate	500.0 Mbps	Yes	\$0.105 hourly	\$0.078 hourly	\$0.193 hourly	\$0.170 hourly
C4 High-CPU Extra Large	c4.xlarge	7.5 GB	16 units	4 vCPUs	0 GB (EBS only)	64-bit	High	750.0 Mbps	Yes	\$0.209 hourly	\$0.155 hourly	\$0.386 hourly	\$0.339 hourly
C4 High-CPU Double Extra Large	c4.2xlarge	15.0 GB	31 units	8 vCPUs	0 GB (EBS only)	64-bit	High	1000.0 Mbps	Yes	\$0.419 hourly	\$0.311 hourly	\$0.773 hourly	\$0.679 hourly
C4 High-CPU Quadruple Extra Large	c4.4xlarge	30.0 GB	62 units	16 vCPUs	0 GB (EBS only)	64-bit	High	2000.0 Mbps	Yes	\$0.838 hourly	\$0.621 hourly	\$1.546 hourly	\$1.357 hourly
C4 High-CPU Eight Extra Large	c4.8xlarge	60.0 GB	132 units	36 vCPUs	0 GB (EBS only)	64-bit	10 Gigabit	4000.0 Mbps	Yes	\$1.675 hourly	\$1.242 hourly	\$3.091 hourly	\$2.769 hourly
P2 Extra Large	p2.xlarge	61.0 GB	12 units	4 vCPUs	0 GB (EBS only)	64-bit	High	750.0 Mbps	No	\$0.900 hourly	\$0.684 hourly	\$1.084 hourly	\$0.868 hourly
P2 Eight Extra Large	p2.8xlarge	488.0 GB	94 units	32 vCPUs	0 GB (EBS only)	64-bit	10 Gigabit	5000.0 Mbps	No	\$7.200 hourly	\$5.476 hourly	\$8.672 hourly	\$6.948 hourly
P2 16xlarge	p2.16xlarge	732.0 GB	188 units	64 vCPUs	0 GB (EBS only)	64-bit	20 Gigabit	10000.0 Mbps	No	\$14.400 hourly	\$10.951 hourly	\$17.344 hourly	\$13.895 hourly
G2 Double Extra Large	g2.2xlarge	15.0 GB	26 units	8 vCPUs	60.0 GB SSD	64-bit	High	1000.0 Mbps	No	\$0.650 hourly	\$0.474 hourly	\$0.767 hourly	\$0.611 hourly
G2 Eight Extra Large	g2.8xlarge	60.0 GB	104 units	32 vCPUs	240.0 GB (2 * 120.0 GB SSD)	64-bit	10 Gigabit	N/A	No	\$2.600 hourly	\$1.896 hourly	\$2.878 hourly	\$1.979 hourly
X1 16xlarge	x1.16xlarge	976.0 GB	174.5 units	64 vCPUs	1920.0 GB SSD	64-bit	10 Gigabit	5000.0 Mbps	No	\$6.669 hourly	\$4.579 hourly	\$9.613 hourly	\$7.523 hourly
X1 32xlarge	x1.32xlarge	1952.0 GB	349 units	128 vCPUs	3840.0 GB (2 * 1920.0 GB SSD)	64-bit	20 Gigabit	10000.0 Mbps	No	\$13.338 hourly	\$9.158 hourly	\$19.226 hourly	\$15.046 hourly
R3 High-Memory Large	r3.large	15.25 GB	6.5 units	2 vCPUs	32.0 GB SSD	64-bit	Moderate	N/A	No	\$0.166 hourly	\$0.105 hourly	\$0.291 hourly	\$0.238 hourly
R3 High-Memory Extra Large	r3.xlarge	30.5 GB	13 units	4 vCPUs	80.0 GB SSD	64-bit	Moderate	500.0 Mbps	No	\$0.333 hourly	\$0.209 hourly	\$0.583 hourly	\$0.428 hourly
R3 High-Memory Double Extra Large	r3.2xlarge	61.0 GB	26 units	8 vCPUs	160.0 GB SSD	64-bit	High	1000.0 Mbps	No	\$0.665 hourly	\$0.418 hourly	\$1.045 hourly	\$0.824 hourly
R3 High-Memory Quadruple Extra Large	r3.4xlarge	122.0 GB	52 units	16 vCPUs	320.0 GB SSD	64-bit	High	2000.0 Mbps	No	\$1.330 hourly	\$0.836 hourly	\$1.944 hourly	\$1.490 hourly
R3 High-Memory Eight Extra Large	r3.8xlarge	244.0 GB	104 units	32 vCPUs	640.0 GB (2 * 320.0 GB SSD)	64-bit	10 Gigabit	N/A	No	\$2.660 hourly	\$1.672 hourly	\$3.500 hourly	\$1.989 hourly
I2 Extra Large	i2.xlarge	30.5 GB	14 units	4 vCPUs	800.0 GB SSD	64-bit	Moderate	500.0 Mbps	No	\$0.853 hourly	\$0.424 hourly	\$0.973 hourly	\$0.565 hourly
I2 Double Extra Large	i2.2xlarge	61.0 GB	27 units	8 vCPUs	1600.0 GB (2 * 800.0 GB SSD)	64-bit	High	1000.0 Mbps	No	\$1.705 hourly	\$0.848 hourly	\$1.946 hourly	\$1.131 hourly
I2 Quadruple Extra Large	i2.4xlarge	122.0 GB	53 units	16 vCPUs	3200.0 GB (4 * 800.0 GB SSD)	64-bit	High	2000.0 Mbps	No	\$3.410 hourly	\$1.696 hourly	\$3.891 hourly	\$2.260 hourly

#thecloudistoodamnhard

What type of instances?

How many to spin up?

What base image?

What price spot?

And then wait to start.....

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2	RDS												
Region: US East (N. Virginia) Cost: Hourly Reserved: 1 yr - No Uplift - Columns - Compare Selected Clear Filters													
Filter: Min Memory (GB): Compute Units: Storage (GB):													
Name	API Name	Memory	Compute Units (ECU)	vCPUs	Storage	Arch	Network Performance	EBS Optimized: Max Bandwidth	VPC Only	Linux On Demand cost	Linux Reserved cost	Windows On Demand cost	Windows Reserved cost
Cluster Compute Eight Extra Large	c8t.8xlarge	60.5 GB	88 units	32 vCPUs	3380.0 GB (4 * 840.0 GB)	64-bit	10 Gbit/s	N/A	No	\$2,000 hourly	\$1,000 hourly	\$2,370 hourly	\$1,336 hourly
Cluster GPU Quadriple Extra Large	cg1.4xlarge	22.5 GB	33.5 units	16 vCPUs	1680.0 GB (2 * 840.0 GB)	64-bit	10 Gbit/s	N/A	No	\$2,100 hourly	unavailable	\$2,670 hourly	unavailable
T2 Nano	t2.nano	0.5 GB	1 unit	1 vCPU	0 GB (EBS only)	64-bit	Low	N/A	Yes	\$0,006 hourly	\$0,005 hourly	\$0,009 hourly	\$0,007 hourly
T2 Micro	t2.micro	1.0 GB	1 unit	1 vCPU	0 GB (EBS only)	32-bit	Low to Moderate	N/A	Yes	\$0,013 hourly	\$0,005 hourly	\$0,018 hourly	\$0,014 hourly
T2 Small	t2.small	2.0 GB	1 unit	1 vCPU	0 GB (EBS only)	32-bit	Low to Moderate	N/A	Yes	\$0,026 hourly	\$0,010 hourly	\$0,036 hourly	\$0,028 hourly
T2 Medium	t2.medium	4.0 GB	2 units	2 vCPUs	0 GB (EBS only)	64-bit	Low to Moderate	N/A	Yes	\$0,052 hourly	\$0,020 hourly	\$0,072 hourly	\$0,060 hourly
T2 Large	t2.large	8.0 GB	3 units	3 vCPUs	0 GB (EBS only)	64-bit	Low to Moderate	N/A	Yes	\$0,104 hourly	\$0,072 hourly	\$0,134 hourly	\$0,106 hourly
M4 Large	m4.large	8.0 GB	6.5 units	2 vCPUs	0 GB (EBS only)	64-bit	Moderate	450.0 Mbps	Yes	\$0,120 hourly	\$0,083 hourly	\$0,246 hourly	\$0,184 hourly
M4 Extra Large	m4.xlarge	16.0 GB	13 units	4 vCPUs	0 GB (EBS only)	64-bit	High	750.0 Mbps	Yes	\$0,239 hourly	\$0,164 hourly	\$0,491 hourly	\$0,366 hourly
M4 Double Extra Large	m4.2xlarge	32.0 GB	26 units	8 vCPUs	0 GB (EBS only)	64-bit	High	1000.0 Mbps	Yes	\$0,478 hourly	\$0,329 hourly	\$0,983 hourly	\$0,735 hourly
M4 Quadriple Extra Large	m4.4xlarge	64.0 GB	53.5 units	16 vCPUs	0 GB (EBS only)	64-bit	High	2000.0 Mbps	Yes	\$0,958 hourly	\$0,658 hourly	\$1,966 hourly	\$1,499 hourly
M4 Deca Extra Large	m4.10xlarge	160.0 GB	124.5 units	40 vCPUs	0 GB (EBS only)	64-bit	10 Gbit/s	4000.0 Mbps	Yes	\$2,394 hourly	\$1,644 hourly	\$4,914 hourly	\$3,672 hourly
M4 16xlarge	m4.16xlarge	256.0 GB	188 units	64 vCPUs	0 GB (EBS only)	64-bit	20 Gbit/s	10000.0 Mbps	Yes	\$3,830 hourly	\$2,632 hourly	\$7,862 hourly	\$5,875 hourly
C4 High-CPU Large	c4.large	3.75 GB	8 units	2 vCPUs	0 GB (EBS only)	64-bit	Moderate	500.0 Mbps	Yes	\$0,105 hourly	\$0,078 hourly	\$0,193 hourly	\$0,170 hourly
C4 High-CPU Extra Large	c4.xlarge	7.5 GB	16 units	4 vCPUs	0 GB (EBS only)	64-bit	High	750.0 Mbps	Yes	\$0,209 hourly	\$0,150 hourly	\$0,386 hourly	\$0,339 hourly
C4 High-CPU Double Extra Large	c4.2xlarge	15.0 GB	31 units	8 vCPUs	0 GB (EBS only)	64-bit	High	1000.0 Mbps	Yes	\$0,419 hourly	\$0,311 hourly	\$0,773 hourly	\$0,679 hourly
C4 High-CPU Quadriple Extra Large	c4.4xlarge	30.0 GB	62 units	16 vCPUs	0 GB (EBS only)	64-bit	High	2000.0 Mbps	Yes	\$0,839 hourly	\$0,622 hourly	\$1,546 hourly	\$1,357 hourly
C4 High-CPU Eight Extra Large	c4.8xlarge	60.0 GB	124 units	32 vCPUs	0 GB (EBS only)	64-bit	10 Gbit/s	4000.0 Mbps	Yes	\$1,678 hourly	\$1,244 hourly	\$3,091 hourly	\$2,715 hourly
P2 Extra Large	p2.xlarge	61.0 GB	12 units	4 vCPUs	0 GB (EBS only)	64-bit	High	750.0 Mbps	No	\$0,900 hourly	\$0,684 hourly	\$1,094 hourly	\$0,808 hourly
P2 Eight Extra Large	p2.8xlarge	486.0 GB	94 units	32 vCPUs	0 GB (EBS only)	64-bit	10 Gbit/s	5000.0 Mbps	No	\$7,200 hourly	\$0,878 hourly	\$8,872 hourly	\$6,948 hourly
P2 16xlarge	p2.16xlarge	732.0 GB	188 units	64 vCPUs	0 GB (EBS only)	64-bit	20 Gbit/s	10000.0 Mbps	No	\$14,400 hourly	\$10,851 hourly	\$17,344 hourly	\$13,895 hourly
G2 Double Extra Large	g2.2xlarge	15.0 GB	26 units	6 vCPUs	60.0 GB SSD	64-bit	High	1000.0 Mbps	No	\$0,474 hourly	\$0,474 hourly	\$0,787 hourly	\$0,811 hourly
G2 Eight Extra Large	g2.8xlarge	60.0 GB	104 units	32 vCPUs	240.0 GB (2 * 120.0 GB SSD)	64-bit	10 Gbit/s	N/A	No	\$2,600 hourly	\$1,896 hourly	\$2,878 hourly	\$1,979 hourly
X1 16xlarge	x1.16xlarge	976.0 GB	174.5 units	64 vCPUs	1800.0 GB SSD	64-bit	10 Gbit/s	5000.0 Mbps	No	\$6,666 hourly	\$4,579 hourly	\$9,613 hourly	\$7,523 hourly
X1 32xlarge	x1.32xlarge	1952.0 GB	349 units	128 vCPUs	3640.0 GB (2 * 1820.0 GB SSD)	64-bit	20 Gbit/s	10000.0 Mbps	No	\$13,338 hourly	\$9,158 hourly	\$19,228 hourly	\$15,046 hourly
R3 High-Memory Large	r3.large	15.25 GB	6.5 units	2 vCPUs	32.0 GB SSD	64-bit	Moderate	N/A	No	\$0,106 hourly	\$0,100 hourly	\$0,291 hourly	\$0,238 hourly
R3 High-Memory Extra Large	r3.xlarge	30.5 GB	13 units	4 vCPUs	64.0 GB SSD	64-bit	Moderate	500.0 Mbps	No	\$0,333 hourly	\$0,200 hourly	\$0,683 hourly	\$0,428 hourly
R3 High-Memory Double Extra Large	r3.2xlarge	61.0 GB	26 units	8 vCPUs	128.0 GB SSD	64-bit	High	1000.0 Mbps	No	\$0,665 hourly	\$0,419 hourly	\$1,045 hourly	\$0,624 hourly
R3 High-Memory Quadriple Extra Large	r3.4xlarge	122.0 GB	52 units	16 vCPUs	256.0 GB SSD	64-bit	High	2000.0 Mbps	No	\$1,330 hourly	\$0,838 hourly	\$1,844 hourly	\$1,490 hourly
R3 High-Memory Eight Extra Large	r3.8xlarge	244.0 GB	104 units	32 vCPUs	512.0 GB (2 * 256.0 GB SSD)	64-bit	10 Gbit/s	N/A	No	\$2,660 hourly	\$1,677 hourly	\$3,537 hourly	\$2,800 hourly
R2 Extra Large	r2.xlarge	30.5 GB	14 units	4 vCPUs	80.0 GB SSD	64-bit	Moderate	500.0 Mbps	No	\$0,853 hourly	\$0,454 hourly	\$0,973 hourly	\$0,565 hourly
I2 Double Extra Large	i2.2xlarge	61.0 GB	27 units	8 vCPUs	1600.0 GB (2 * 800.0 GB SSD)	64-bit	High	1000.0 Mbps	No	\$1,705 hourly	\$0,848 hourly	\$1,946 hourly	\$1,131 hourly
I2 Quadriple Extra Large	i2.4xlarge	122.0 GB	53 units	16 vCPUs	3200.0 GB (4 * 800.0 GB SSD)	64-bit	High	2000.0 Mbps	No	\$3,410 hourly	\$1,696 hourly	\$3,891 hourly	\$2,260 hourly
I2 Eight Extra Large	i2.8xlarge	244.0 GB	104 units	32 vCPUs	6400.0 GB (8 * 800.0 GB SSD)	64-bit	10 Gbit/s	N/A	No	\$6,820 hourly	\$3,392 hourly	\$7,782 hourly	\$4,821 hourly
D2 Extra Large	d2.xlarge	30.5 GB	14 units	4 vCPUs	6000.0 GB (3 * 2000.0 GB)	64-bit	Moderate	750.0 Mbps	No	\$0,690 hourly	\$0,402 hourly	\$0,821 hourly	\$0,472 hourly
D2 Double Extra Large	d2.2xlarge	61.0 GB	28 units	8 vCPUs	12000.0 GB (6 * 2000.0 GB)	64-bit	High	1000.0 Mbps	No	\$1,380 hourly	\$0,804 hourly	\$1,601 hourly	\$0,885 hourly
D2 Quadriple Extra Large	d2.4xlarge	122.0 GB	56 units	16 vCPUs	24000.0 GB (12 * 2000.0 GB)	64-bit	High	2000.0 Mbps	No	\$2,760 hourly	\$1,608 hourly	\$3,062 hourly	\$1,690 hourly
D2 Eight Extra Large	d2.8xlarge	244.0 GB	118 units	36 vCPUs	48000.0 GB (24 * 2000.0 GB)	64-bit	10 Gbit/s	4000.0 Mbps	No	\$5,520 hourly	\$3,216 hourly	\$6,198 hourly	\$3,300 hourly
H1 High I/O Quadriple Extra Large	h1.4xlarge	60.5 GB	35 units	16 vCPUs	2048.0 GB (2 * 1024.0 GB SSD)	64-bit	10 Gbit/s	N/A	No	\$3,100 hourly	\$1,698 hourly	\$3,580 hourly	\$2,260 hourly
High Storage Eight Extra Large	hs1.8xlarge	117.0 GB	35 units	16 vCPUs	48000.0 GB (24 * 2000.0 GB)	64-bit	10 Gbit/s	N/A	No	\$4,600 hourly	\$2,574 hourly	\$4,601 hourly	\$2,961 hourly
M3 General Purpose Medium	m3.medium	3.75 GB	3 units	1 vCPU	4.0 GB SSD	64-bit	Moderate	N/A	No	\$0,067 hourly	\$0,046 hourly	\$0,150 hourly	\$0,100 hourly
M3 General Purpose Large	m3.large	7.5 GB	6.5 units	2 vCPUs	8.0 GB SSD	64-bit	Moderate	N/A	No	\$0,133 hourly	\$0,092 hourly	\$0,219 hourly	\$0,159 hourly
M3 General Purpose Extra Large	m3.xlarge	15.0 GB	13 units	4 vCPUs	16.0 GB (2 * 8.0 GB SSD)	64-bit	High	500.0 Mbps	No	\$0,266 hourly	\$0,180 hourly	\$0,518 hourly	\$0,397 hourly
M3 General Purpose Double Extra Large	m3.2xlarge	30.0 GB	26 units	8 vCPUs	32.0 GB (2 * 16.0 GB SSD)	64-bit	High	1000.0 Mbps	No	\$0,532 hourly	\$0,360 hourly	\$1,036 hourly	\$0,793 hourly

What it is?



CLOUD FUNCTIONS ^{ALPHA}

A serverless platform for building event-based microservices

Microsoft Azure

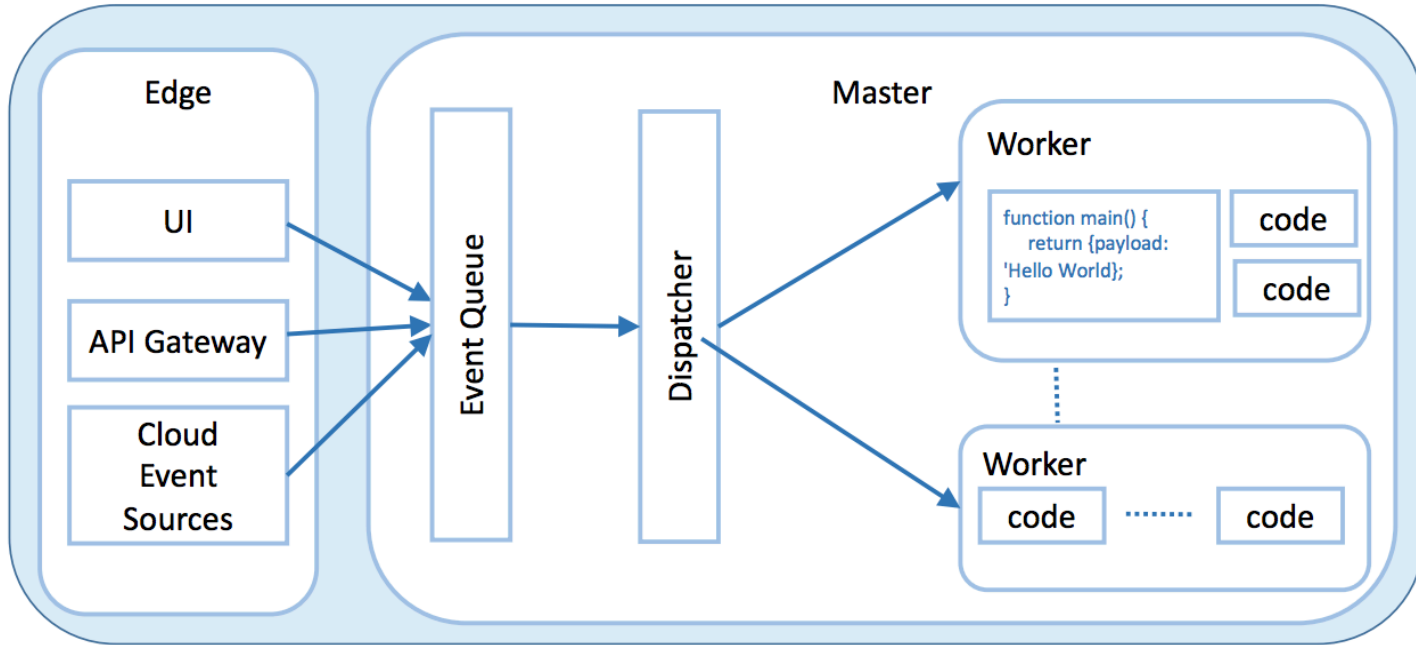
Azure Functions

Process events with a serverless code architecture.

Core capability

1. Manage a set of user defined functions
2. Take an event sent over HTTP or received from an event source
3. Determine function(s) to which to dispatch the event
4. Find an existing instance of function or create a new one
5. Send the event to the function instance
6. Wait for a response
7. Gather execution logs
8. Make the response available to the user
9. Stop the function when it is no longer needed.

Basic architecture



A virtualization story

Virtualizing the client



Virtualizing the server



Virtualizing the cluster



Virtualizing the cloud

Classic web stack

RPC →

Application

Server


OS

Hardware

Weak virtualization

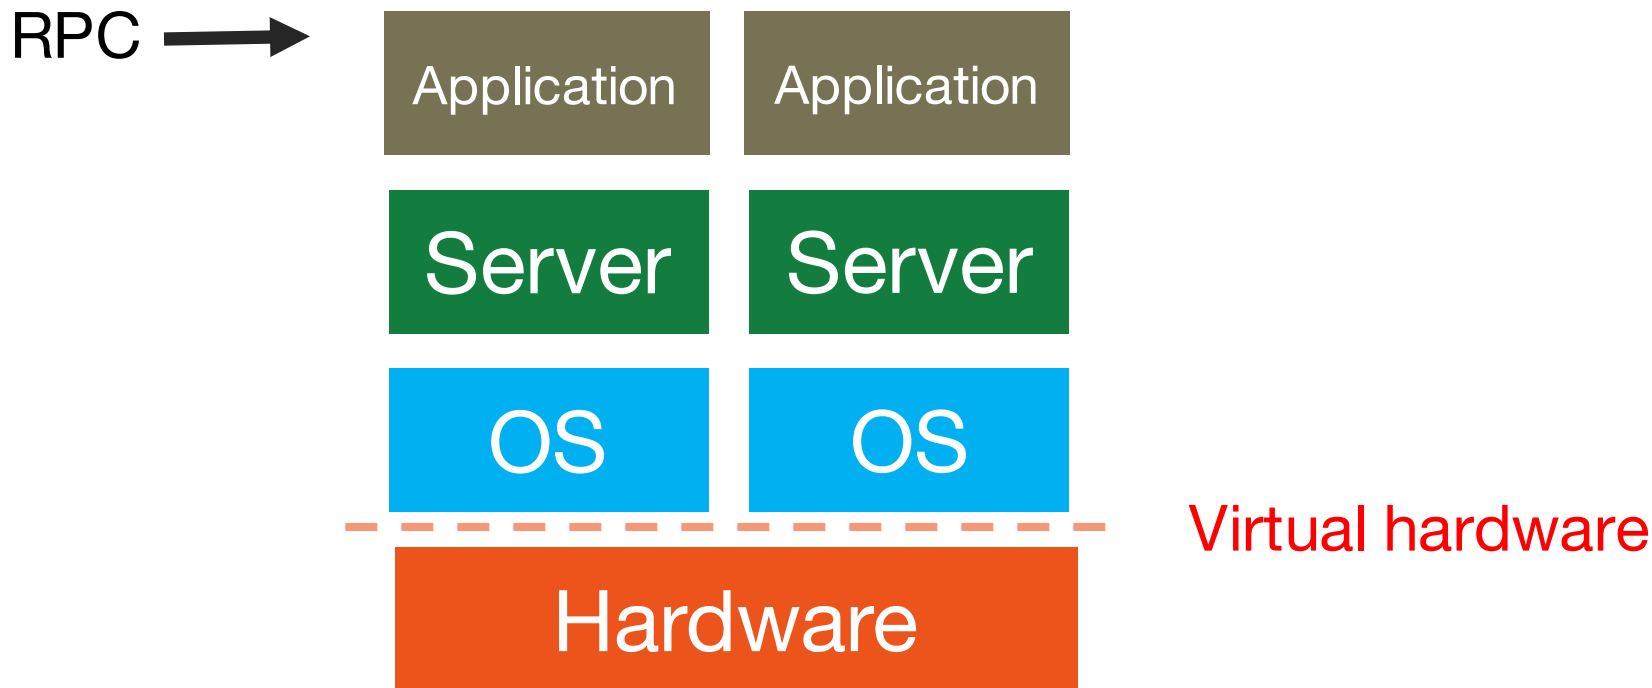
1st Generation: Virtual Machines

RPC →



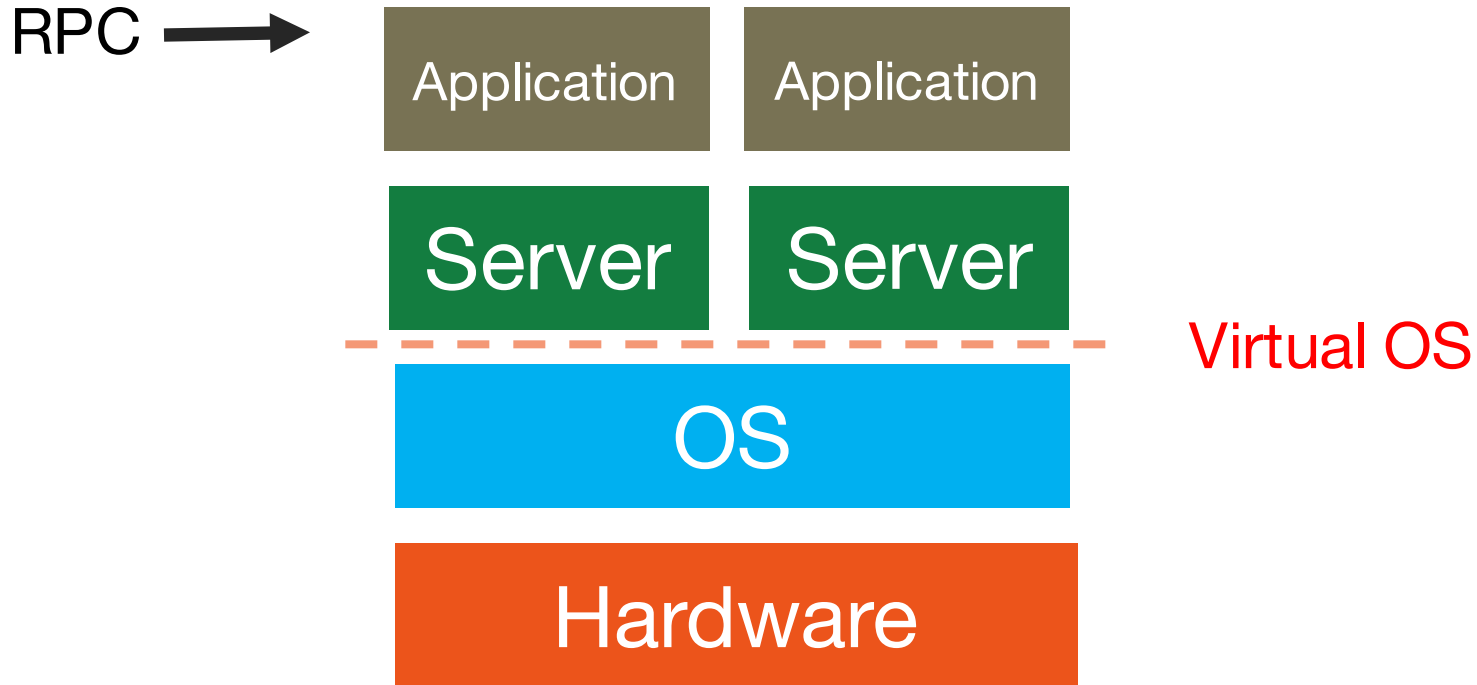
(based on slides at <https://www.usenix.org/node/196323>)

1st Generation: Virtual Machines



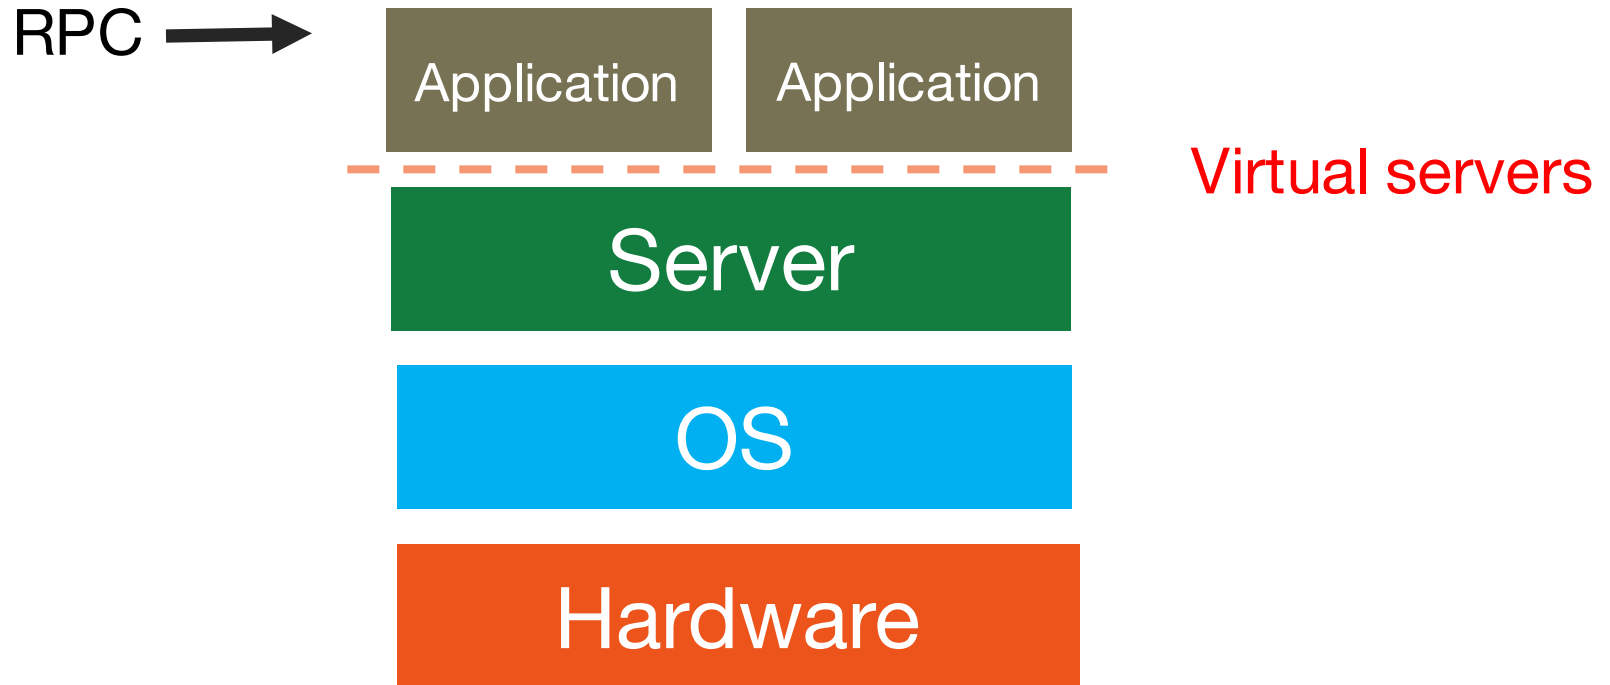
(based on slides at <https://www.usenix.org/node/196323>)

2nd Generation: Containers



(based on slides at <https://www.usenix.org/node/196323>)

3rd Generation: Lambdas



(based on slides at <https://www.usenix.org/node/196323>)

Tradeoffs

Lambdas

Containers

VMs



First serverless app: BigQuery

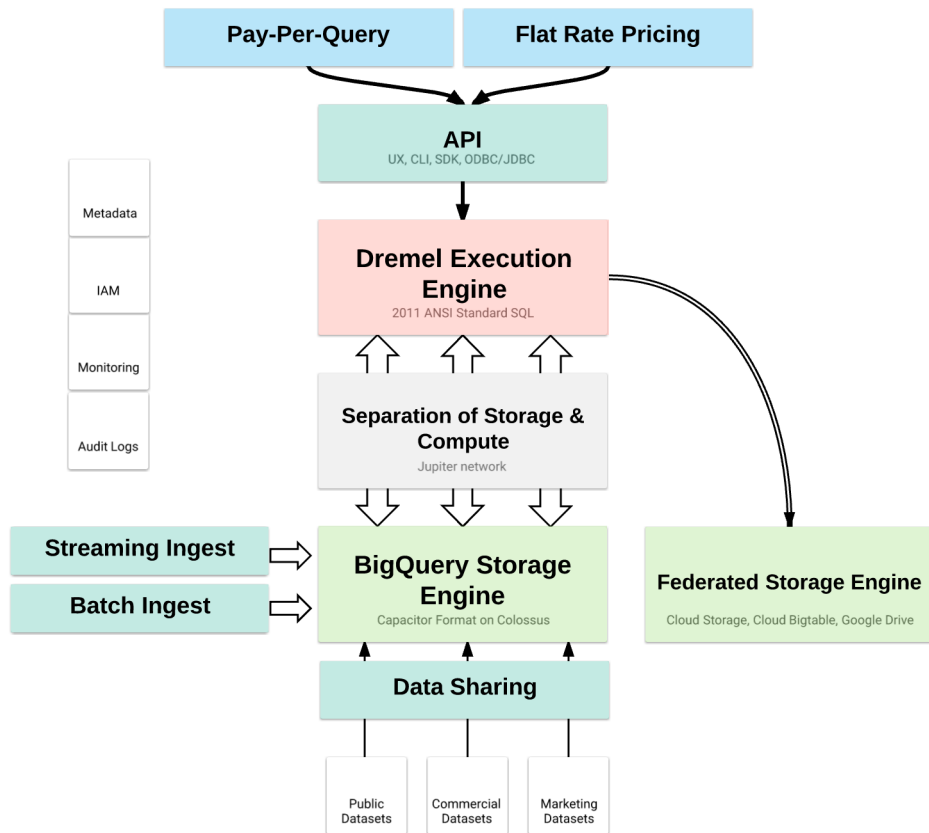


Fully managed Data Warehouse

- “Arbitrarily” large data and queries
- Pay per bytes being processed
- No concept of cluster

Other similar systems

- AWS Athena
- Snowflake
- ...



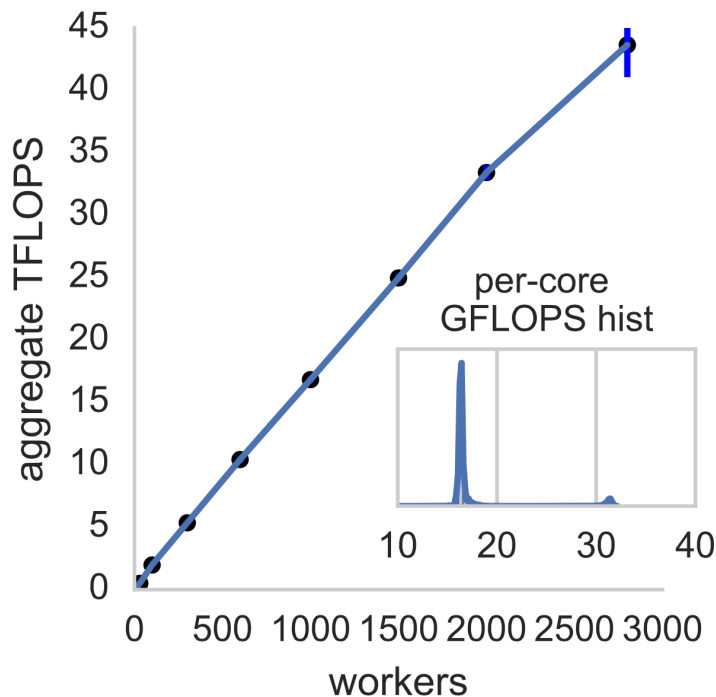
AWS lambdas

Serverless functions: typically read/write data from/to S3

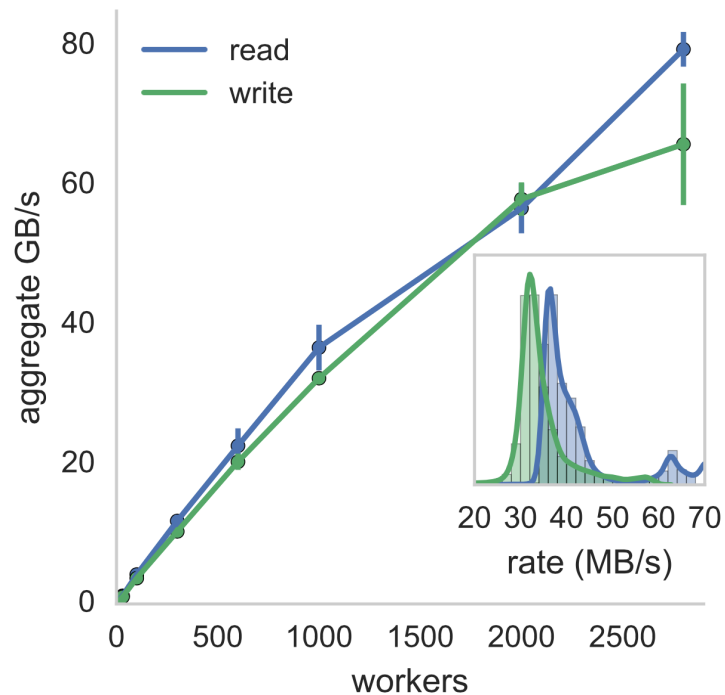
- 300 seconds
- single-core (AVX2)
- 512 MB in /tmp
- 1.5GB RAM
- Python, Java, Node
- Sub-second billing: 3,600 threads for one second → 10¢

Lambda Scalability

Compute



Data



A case study: PyWren API

```
import pywren
import numpy as np
```

```
def addone(x):
    return x + 1
```

```
wrenexec = pywren.default_executor()
xlist = np.arange(10)
futures = wrenexec.map(addone, xlist)
```

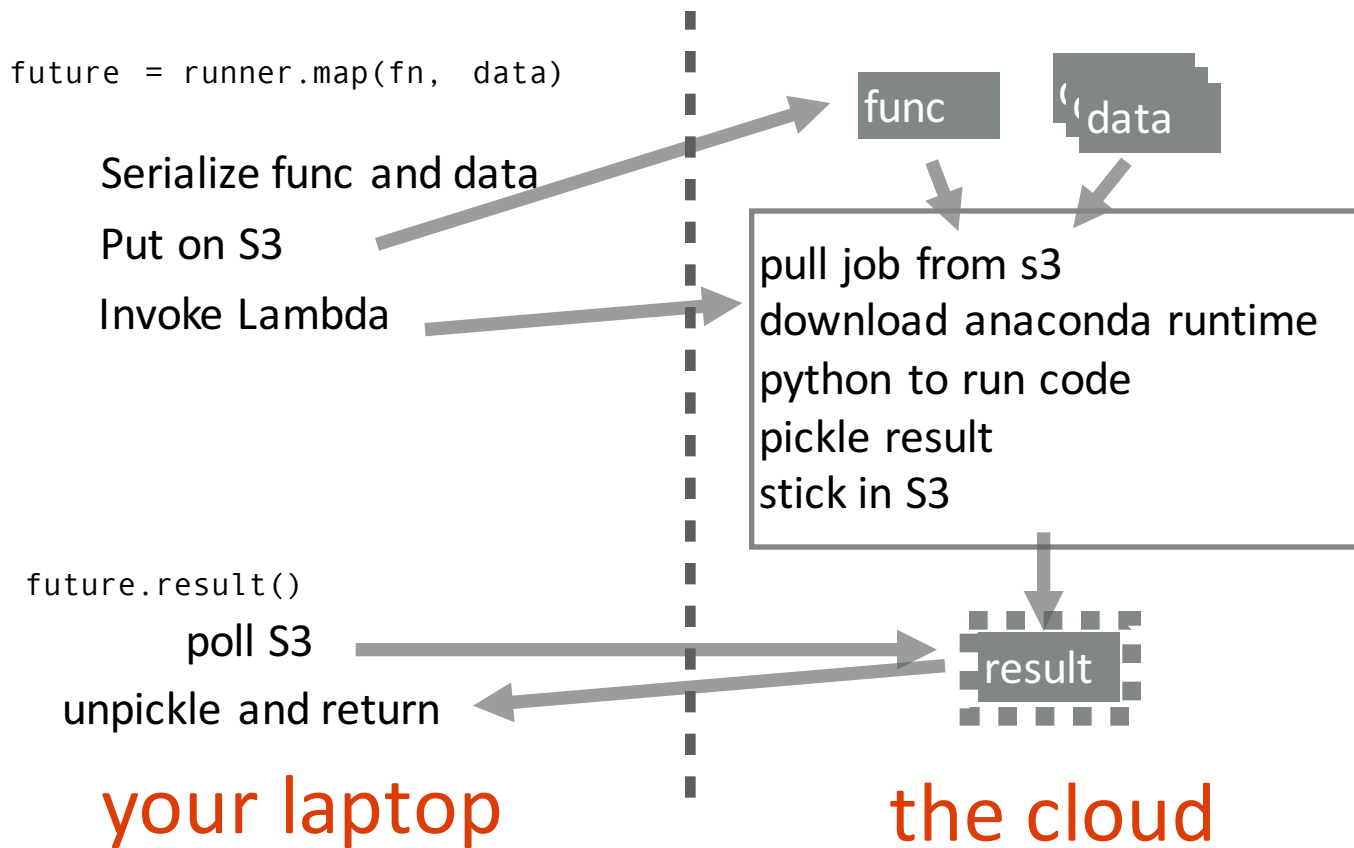
```
print [f.result() for f in futures]
```



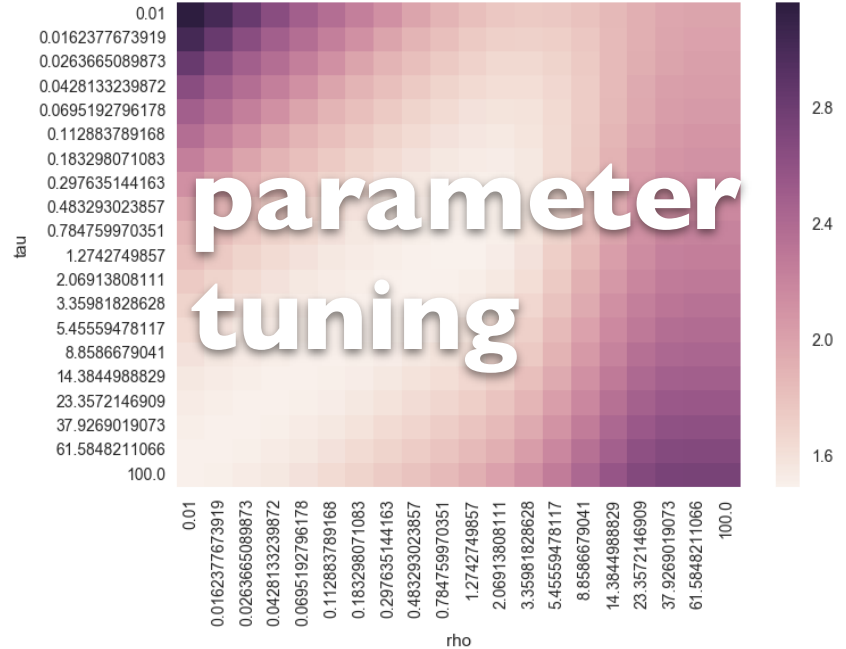
The output is as expected:

```
[1, 2, 3, 4, 5, 6, 7, 8, 9, 10]
```

How it works?



You can do A LOT OF work with map!




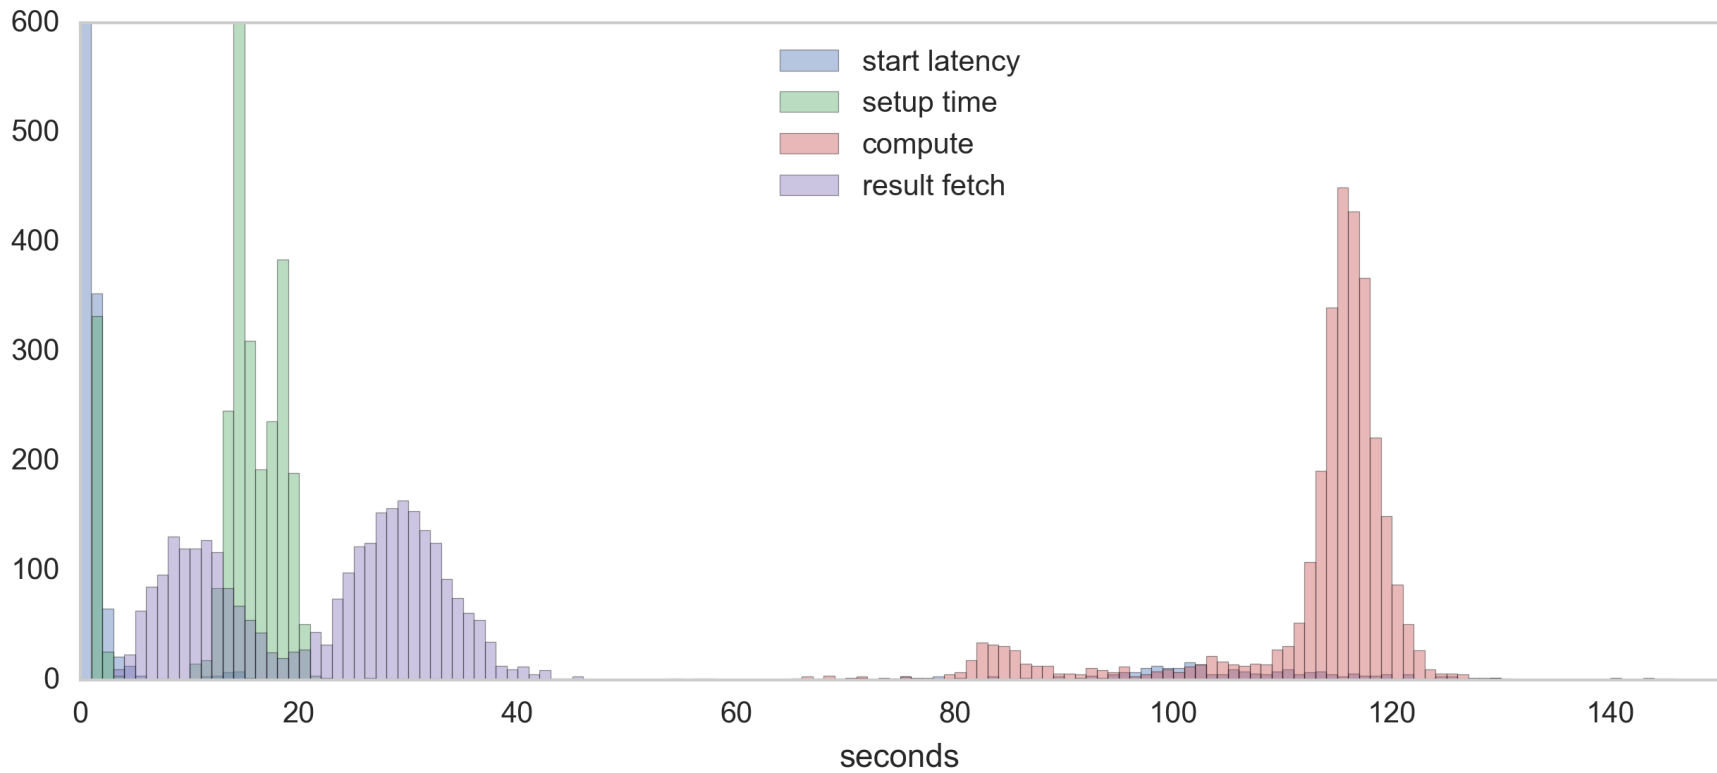


Preprocess 1.4M images from

IMAGENET

Compute GIST
image descriptor
(some random
python code off
the internet)





Want our runtime to include



Start

1205MB

conda clean

977 MB

eliminate pkg

946 MB

Delete non-AVX2 MKL

670 MB

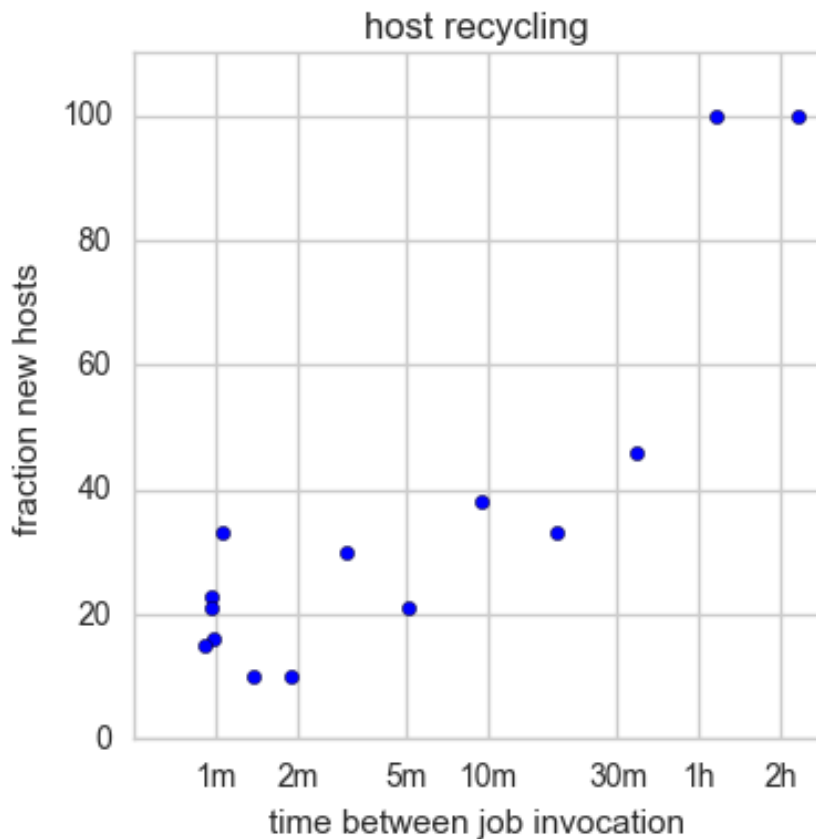
strip shared libs

510MB

delete pyc

441MB

Understanding host allocation



Beyond Maps...

Numpywren: **linear algebra library** on top of Pywren/Lambdas

Use Amazon S3 to store sharded intermediate data, stream shards to CPUs for parallel computation

How well does it work?



N	D	Cores	Dgemm runtime	Aggregate FLOPS	Peak FLOPS
1.2 Million	4096	3000	1320s	8.9 TFLOPs	40 TFLOPs
1.2 Million	18432	3000	2520s	21 TFLOPs	40 TFLOPs

How well does it work?




N	D	Cores	Dgemm runtime	Aggregate FLOPs	Peak FLOPs
1.2 Million	4096	3000	1320s	8.9 TFLOPs	40 TFLOPs
1.2 Million	18432	3000	2520s	21 TFLOPs	40 TFLOPs

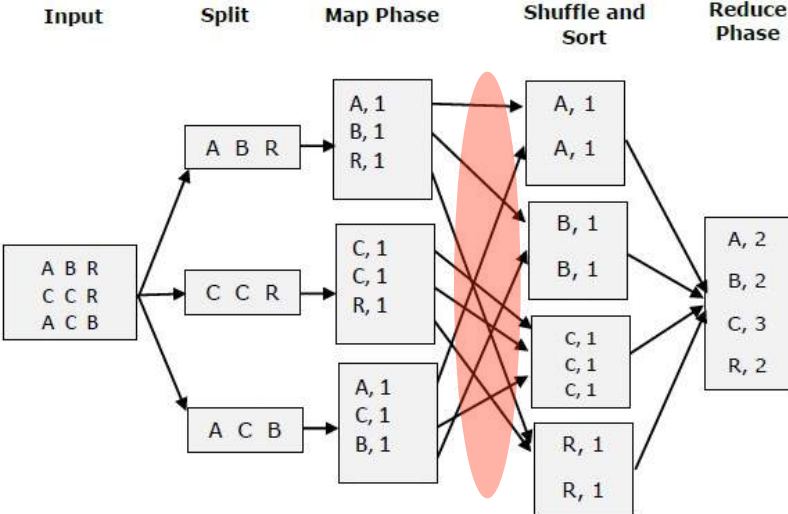
0.5 * Peak FLOPs!



From Map to MapReduce

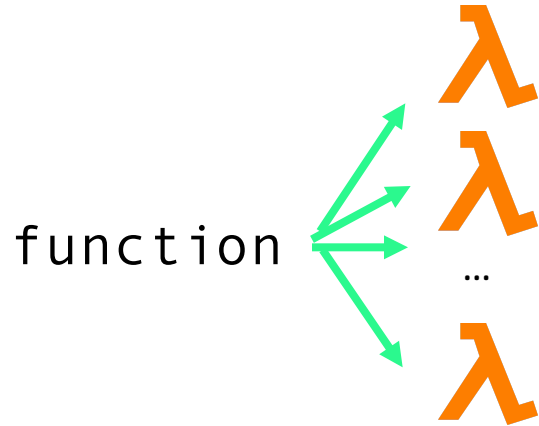


Map()




A baby analytics job: word count

The missing piece: Shuffle

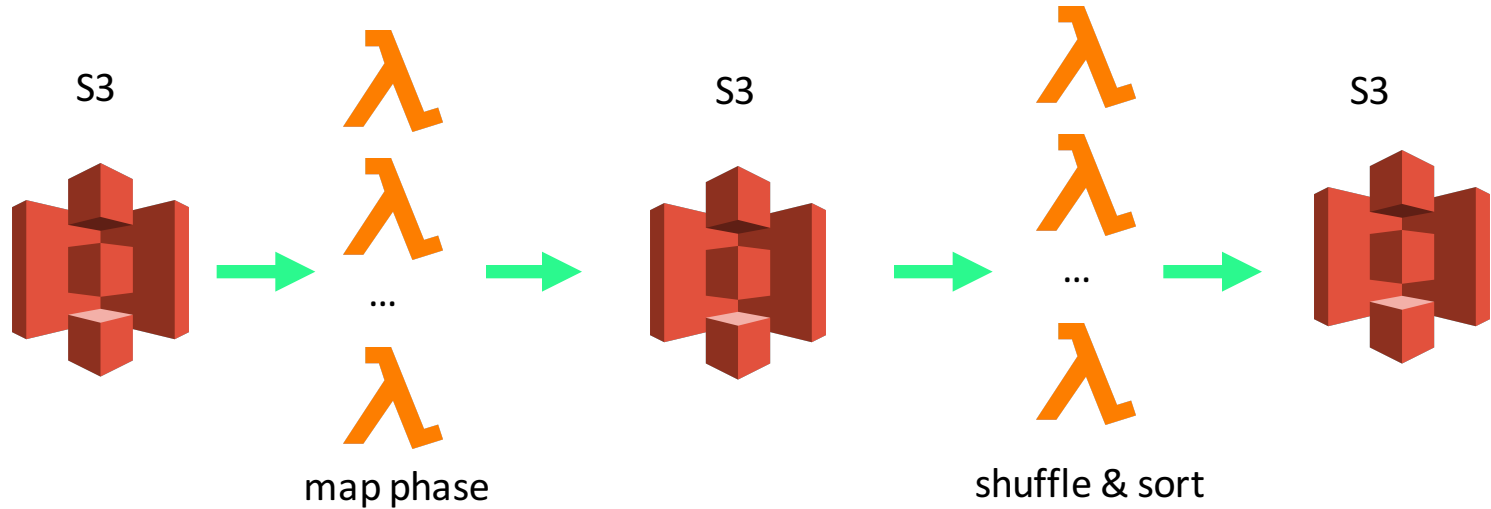


Map()



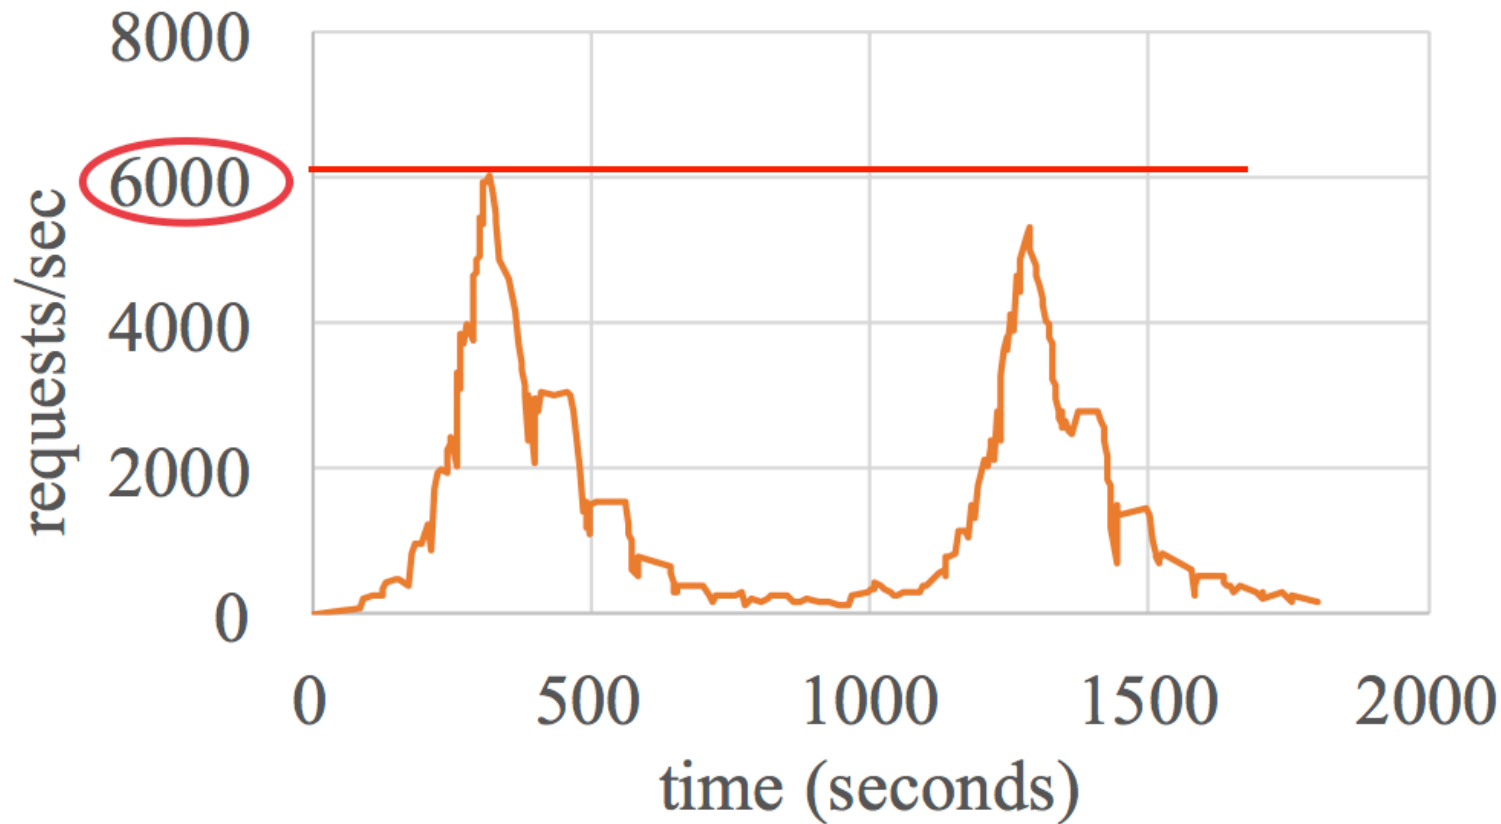
A baby analytics job: word count

The missing piece: Shuffle



Now we can do everything! Or is it?

Rate-limiting on S3 (SlowDown error)



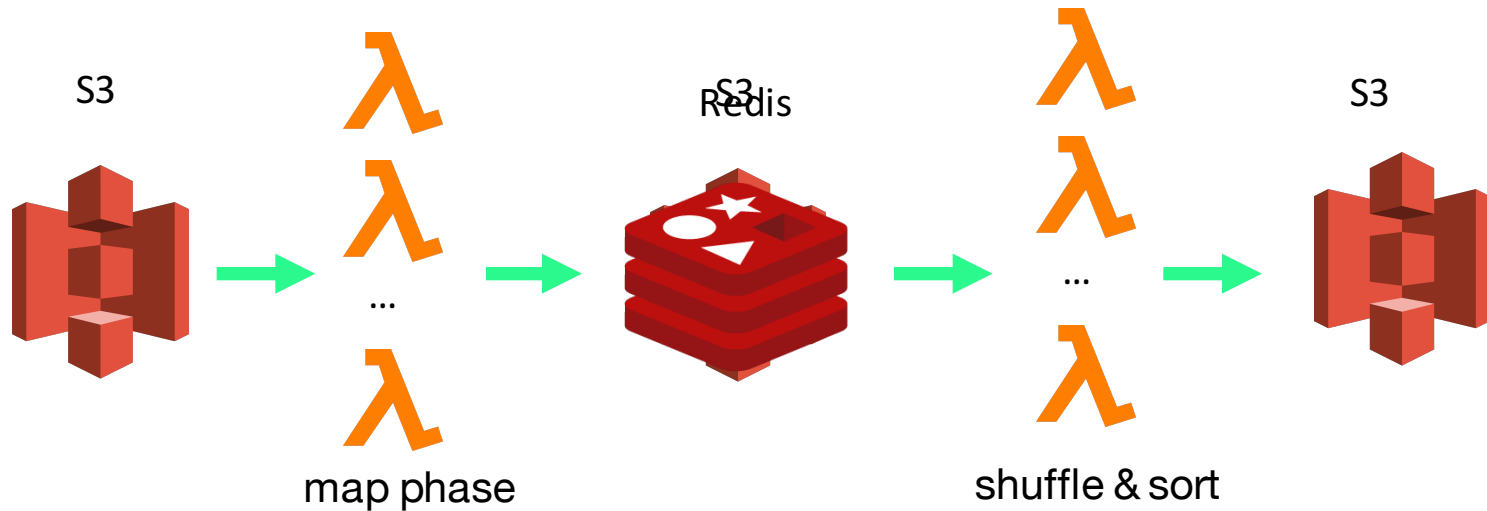
Doing 100TB Sort

CloudSort record: Spark completes in **2983s** with **\$144**

Back of envelope calculation for serverless sort:

- 1GB per container: 100000 mappers, 100000 reducers
- 10,000,000,000 = 10^{10} files to shuffle
- $10^{10} / 6000 = \mathbf{19 \text{ days!}}$ and \$\$\$\$\$\$

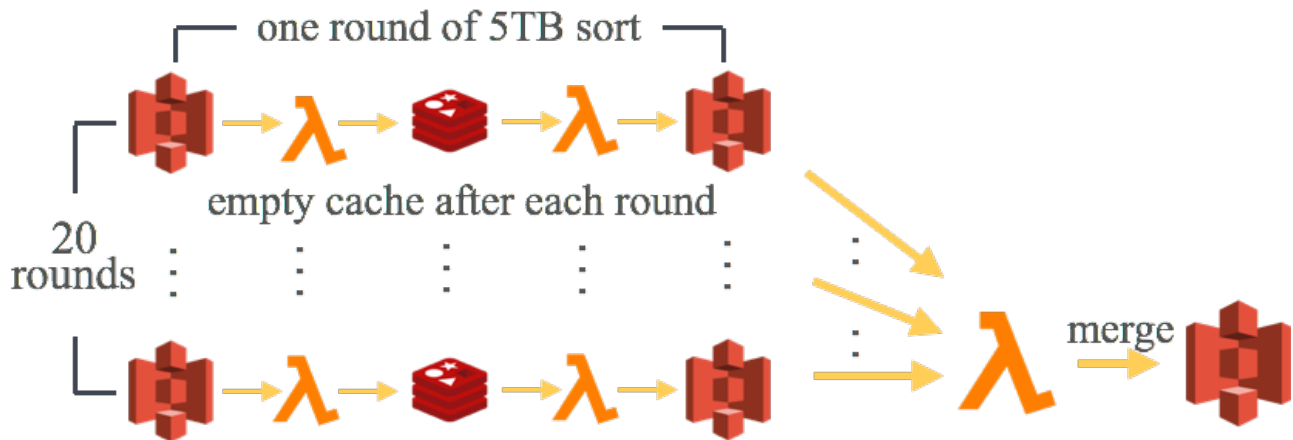
Elastic memory in cloud



S3: cheap for storage/access
low throughput

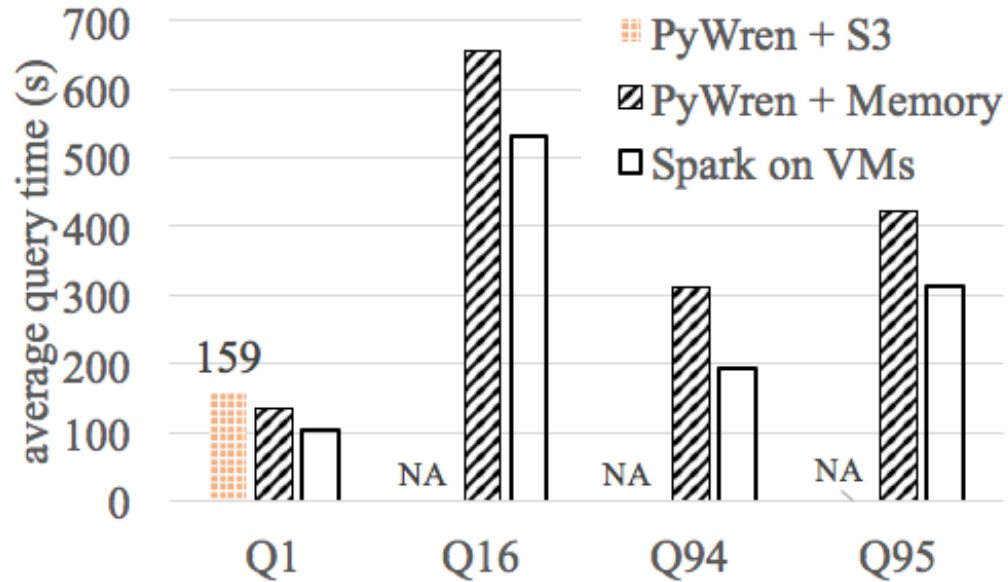
elastic memory: high throughput
expensive?

Doing 100TB Sort



Spark on VMs:	2983 secs	\$144
PyWren:	2945 secs	\$163 (\$143 Lambda + \$32 Redis)

TPC-DS Queries



Comparable performance for shuffle-intensive queries.

ExCamera

People can make changes to a word-processing document

- The changes are instantly visible for the others

Goal: people can interactively edit and transform a video

- The changes are instantly visible for the others

Encoding, Fast and Slow: Low-Latency Video Processing Using Thousands of Tiny Threads,

Sadjad Fouladi, Riad S. Wahby, and Brennan Shacklett, *Stanford University*;

Karthikeyan Vasuki Balasubramaniam, *University of California, San Diego*; William Zeng, *Stanford University*;

Rahul Bhalerao, *University of California, San Diego*; Anirudh Sivaraman, *Massachusetts Institute of Technology*;

George Porter, *University of California, San Diego*; Keith Winstein, *Stanford University*,. NSDI 2017.


<https://ex.camera>

(slide from <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/fouladi>)



"Apply this awesome filter to my video."

(slide from <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/fouladi>)



"Look everywhere for this face in this movie."




"Remake Star Wars Episode I without Jar Jar."

Video Encoding/Compression


Exploit the temporal redundancy in adjacent images.

- Store the first image on its entirety: a key frame.
- For other images, only store a "diff" with nearby images: an interframe.



4K video @15Mbps: key frame ~ 1 MB; an interframe is ~ 25 KB

Video Encoding/Compression



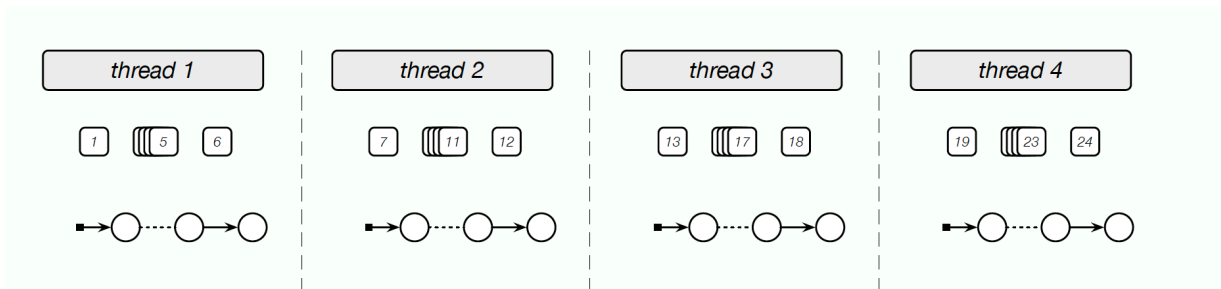
Key idea:

- split video in very small chunks
- encode chunks in parallel
- stitch the encoded chunks

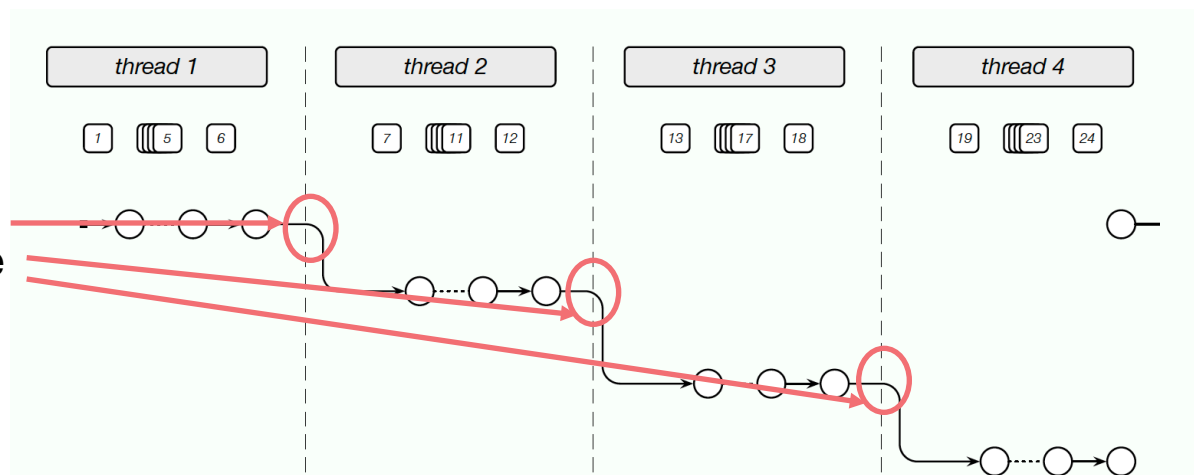
Challenge: Inter-frames are **stateful**, i.e., need state from other frames to reconstruct/encode

Parallel encoder

Split into tiny
chunks & encode
each chunk
(VP8 encoder)



Stitch chunks;
Need to exchange
state!



mu, supercomputing as a service

Library for designing and deploying general-purpose parallel computations on a commercial “cloud function” service.

Starts thousands of threads in seconds and manages **interthread communication** (<https://github.com/excamera/mu>)

- Computation no longer stateless!
- Needed because S3 has high latency

14.8-minute 4K Video @20dB

vpxenc Single-Threaded	453 mins
vpxenc Multi-Threaded	149 mins
YouTube (H.264)	37 mins
ExCamera[6, 16]	2.6 mins

Are Lambdas limitations fundamental?

Recall:

- 300 seconds
- single-core (AVX2)
- 512 MB in /tmp
- 1.5GB RAM

Challenges

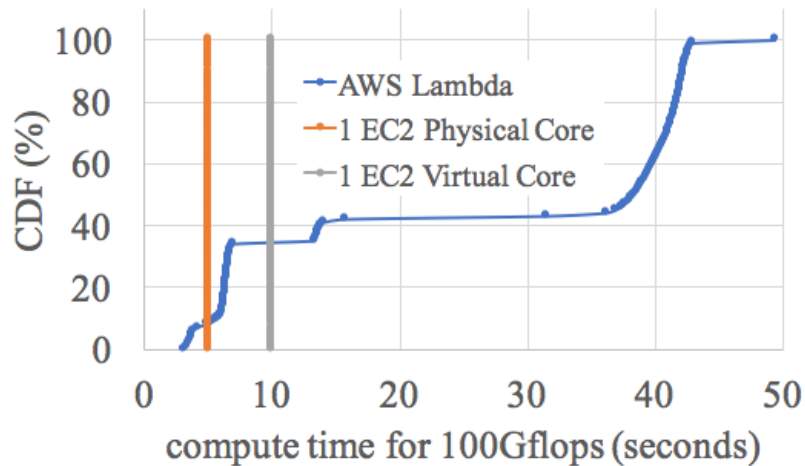
Fast, elastic storage

- Needed for shuffle
- Would enable ExCamera like apps without interthread communication
 - We implemented ExCamera with Redis, simpler, more robust, and as fast

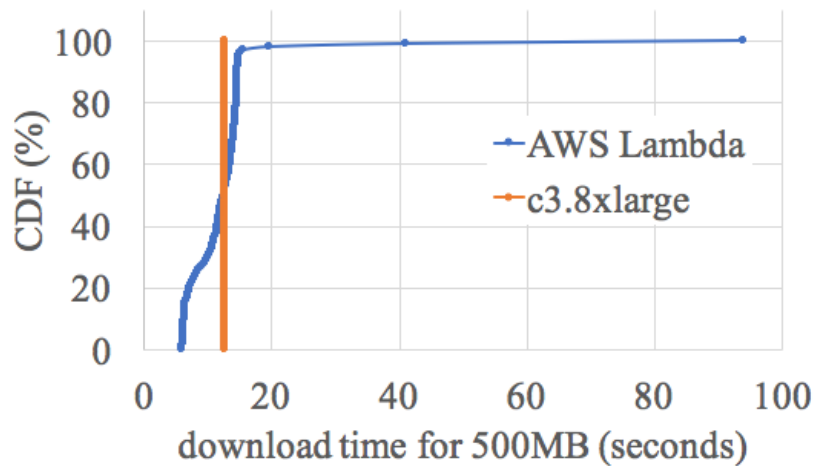
What do we need from such fast storage:

- Elasticity
- Key-value or similar API
- Notification based mechanism (think Chuby)

Performance Isolation



Compute performance has higher variance on Lambda



Network performance has higher variance on Lambda

Challenges

Fast, elastic storage

Performance isolation

Others:

- Security
- Start up time

Summary

Serverless, the next level of virtualizing and multiplexing cloud

Initially restricted to maps, but only a matter of time before extending to other programming models

- Video editing in real time
- Linear algebra
- MpaReduce