

Serverless Computing

DS5110: Big Data Systems

Spring 2025

Lecture 16

Yue Cheng



UNIVERSITY
of VIRGINIA

Some material taken/derived from:

- Berkeley CS 262a (Spring '18) by Ali Ghodsi and Ion Stoica
- Tyler Harter's HotCloud '18 OpenLambda talk

@ 2025 released for use under a [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.

Learning objectives

- Learn the motivation of serverless
- Know the different generations of cloud computing
 - Virtual machines
 - Containers
 - Serverless functions
- Understand current limitations of FaaS

Motivation

When to use the cloud?

- Data
 - Large amounts of data – can't store locally
 - Shared data across users
 - Long-term storage
- Compute
 - Need lots of CPUs for data processing
 - Varying computing demands (resources)
 - No admin (for managing your local hardware)

Based on your inputs, this is the lowest-cost EC2 instance: **t4g.nano**

EC2

Region: Cost: Reserved: Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): Min vCPUs: Min Storage (GiB):

Name	API Name	Memory	vCPUs	Instance Storage	Network Performance	Linux On Demand cost	Linux Reserved cost	Windows On Demand cost	Windows Reserved cost
<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>	<input type="text" value="Search"/>
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NVMe SSD	Up to 25 Gigabit	\$0.272000 hourly	\$0.173000 hourly	\$0.456000 hourly	\$0.357000 hourly
M5A Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	EBS only	Up to 10 Gigabit	\$0.344000 hourly	\$0.219000 hourly	\$0.712000 hourly	\$0.587000 hourly
R5N 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	EBS only	50 Gigabit	\$3.576000 hourly	\$2.253000 hourly	\$5.784000 hourly	\$4.461000 hourly
R5AD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NVMe SSD	10 Gigabit	\$0.262000 hourly	\$0.166000 hourly	\$0.446000 hourly	\$0.350000 hourly
R5N Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	EBS only	Up to 25 Gigabit	\$0.298000 hourly	\$0.188000 hourly	\$0.482000 hourly	\$0.372000 hourly
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NVMe SSD)	50 Gigabit	\$5.424000 hourly	\$3.694000 hourly	\$7.632000 hourly	\$5.902000 hourly
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NVMe SSD)	100 Gigabit	\$10.848000 hourly	\$7.388000 hourly	\$15.264000 hourly	\$11.804000 hourly
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NVMe SSD	Up to 25 Gigabit	\$0.334000 hourly	\$0.211000 hourly	\$0.518000 hourly	\$0.395000 hourly
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800 GiB SSD	Moderate	\$0.853000 hourly	\$0.424000 hourly	\$0.973000 hourly	\$0.565000 hourly
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	EBS only	75 Gigabit	\$3.808000 hourly	\$2.419000 hourly	\$6.752000 hourly	\$5.363000 hourly
T2 Micro	t2.micro	1.0 GiB	1 vCPUs <small>for a 2h 24m burst</small>	EBS only	Low to Moderate	\$0.011600 hourly	\$0.007200 hourly	\$0.016200 hourly	\$0.011800 hourly
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 GiB HDD)	10 Gigabit	\$5.520000 hourly	\$3.216000 hourly	\$6.198000 hourly	\$3.200000 hourly
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NVMe SSD	Up to 25 Gigabit	\$1.356000 hourly	\$0.924000 hourly	\$1.908000 hourly	\$1.476000 hourly
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NVMe SSD	Up to 10 Gigabit	\$1.116000 hourly	\$0.705000 hourly	\$1.668000 hourly	\$1.257000 hourly
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920 GiB SSD	10 Gigabit	\$13.344000 hourly	\$8.223000 hourly	\$16.288000 hourly	\$11.167000 hourly
R5N 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	EBS only	100 Gigabit	\$7.152000 hourly	\$4.506000 hourly	\$11.568000 hourly	\$8.922000 hourly
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800 GiB SSD)	10 Gigabit	\$6.820000 hourly	\$3.392000 hourly	\$7.782000 hourly	\$4.521000 hourly
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	EBS only	Up to 10 Gigabit	\$1.808000 hourly	\$1.141000 hourly	\$3.280000 hourly	\$2.613000 hourly
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	EBS only	Up to 10 Gigabit	\$0.408000 hourly	\$0.257000 hourly	unavailable	unavailable
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800 GiB SSD)	High	\$1.705000 hourly	\$0.848000 hourly	\$1.946000 hourly	\$1.131000 hourly
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NVMe SSD)	Up to 25 Gigabit	\$0.904000 hourly	\$0.616000 hourly	\$1.272000 hourly	\$0.984000 hourly
M5A Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	EBS only	Up to 10 Gigabit	\$0.172000 hourly	\$0.109000 hourly	\$0.356000 hourly	\$0.293000 hourly
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	EBS only	Up to 10 Gigabit	\$3.060000 hourly	\$2.088000 hourly	\$3.428000 hourly	\$2.456000 hourly
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs <small>for a 4h 4.8m burst</small>	EBS only	Moderate	\$0.371200 hourly	\$0.230000 hourly	\$0.433200 hourly	\$0.292000 hourly
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 GiB HDD)	10 Gigabit	\$1.872000 hourly	\$1.272000 hourly	\$3.344000 hourly	\$2.744000 hourly
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NVMe SSD)	25 Gigabit	\$6.912000 hourly	\$4.362000 hourly	\$11.328000 hourly	\$8.778000 hourly
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NVMe SSD)	25 Gigabit	\$2.712000 hourly	\$1.847000 hourly	\$3.816000 hourly	\$2.951000 hourly
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	EBS only	10 Gigabit	\$2.128000 hourly	\$1.344000 hourly	\$3.600000 hourly	\$2.816000 hourly
T2 Large	t2.large	8.0 GiB	2 vCPUs <small>for a 7h 12m burst</small>	EBS only	Low to Moderate	\$0.092800 hourly	\$0.057500 hourly	\$0.120800 hourly	\$0.085500 hourly
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920 GiB SSD	High	\$6.669000 hourly	\$4.110000 hourly	\$9.613000 hourly	\$7.054000 hourly
M5A 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	EBS only	12 Gigabit	\$2.752000 hourly	\$1.751000 hourly	\$5.696000 hourly	\$4.695000 hourly
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	EBS only	25 Gigabit	\$6.048000 hourly	\$3.810000 hourly	\$10.464000 hourly	\$8.226000 hourly
R5A Large	r5a.large	16.0 GiB	2 vCPUs	EBS only	10 Gigabit	\$0.113000 hourly	\$0.071000 hourly	\$0.205000 hourly	\$0.163000 hourly
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16 GiB SSD)	Moderate	\$0.105000 hourly	\$0.073000 hourly	\$0.188000 hourly	\$0.165000 hourly
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	EBS only	20 Gigabit	\$5.424000 hourly	\$3.423000 hourly	\$9.840000 hourly	\$7.839000 hourly
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	EBS only	20 Gigabit	\$4.560000 hourly	\$3.112200 hourly	\$7.504000 hourly	\$6.056200 hourly
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	EBS only	Up to 10 Gigabit	\$0.204000 hourly	\$0.128500 hourly	unavailable	unavailable
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	EBS only	High	\$0.199000 hourly	\$0.126000 hourly	\$0.383000 hourly	\$0.310000 hourly
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480 GiB SSD	Up to 10 Gigabit	\$3.336000 hourly	\$2.056000 hourly	\$4.072000 hourly	\$2.792000 hourly
M5AD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NVMe SSD	Up to 10 Gigabit	\$0.206000 hourly	\$0.132000 hourly	\$0.390000 hourly	\$0.316000 hourly

#thecloudistoodamnhard

1. What type of instances?

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NVMe
M5A Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	150 GiB NVMe
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	150 GiB NVMe
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NVMe
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	150 GiB NVMe
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NVMe)
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NVMe)
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NVMe
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800 GiB
M5N 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	150 GiB NVMe
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	0 GiB
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 GiB NVMe)
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NVMe
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NVMe
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920 GiB
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	150 GiB NVMe
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800 GiB)
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	150 GiB NVMe
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	150 GiB NVMe
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800 GiB)
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NVMe)
M5A Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	150 GiB NVMe
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	150 GiB NVMe
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	0 GiB
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 GiB NVMe)
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NVMe)
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NVMe)
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	150 GiB NVMe
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	0 GiB
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920 GiB
M5A 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	150 GiB NVMe
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	150 GiB NVMe
R5A Large	r5a.large	16.0 GiB	2 vCPUs	150 GiB NVMe
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16 GiB)
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	150 GiB NVMe
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	150 GiB NVMe
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	150 GiB NVMe
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	150 GiB NVMe
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480 GiB
M5AD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NVMe

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?
3. What base image?

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?
3. What base image?
4. On-demand or spot?

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?
3. What base image?
4. On-demand or spot?
5. What storage service to use?

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?
3. What base image?
4. On-demand or spot?
5. What storage service to use?
6. And then wait to start...

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?
3. What base image?
4. On-demand or spot?
5. What storage service to use?
6. And then wait to start...
7. Not the end of the horror story:
 1. When to scale out?
 2. When to scale in?
 3. When to switch to different instance types?

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV

#thecloudistoodamnhard

1. What type of instances?
2. How many to spin up?
3. What base image?
4. On-demand or spot?
5. What storage service to use?
6. And then wait to start...
7. Not the end of the horror story:
 1. When to scale out?
 2. When to scale in?
 3. When to switch to different instance types?
8. Go back to Step 1...

EC2Instances.info Easy Amazon EC2 Instance Comparison

EC2 RDS

Region: US East (N. Virginia) - Cost: Hourly - Reserved: 1-year - No Upfront - Columns - Compare Selected Clear Filters CSV

Filter: Min Memory (GiB): 0 Min vCPUs: 0 Min Storage (GiB): 0

Name	API Name	Memory	vCPUs	Instance Storage
MSDN Extra Large	m5dn.xlarge	16.0 GiB	4 vCPUs	150 GiB NV
MSA Double Extra Large	m5a.2xlarge	32.0 GiB	8 vCPUs	E
RSN 12xlarge	r5n.12xlarge	384.0 GiB	48 vCPUs	E
RSAD Extra Large	r5ad.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
RSN Extra Large	r5n.xlarge	32.0 GiB	4 vCPUs	E
I3EN 12xlarge	i3en.12xlarge	384.0 GiB	48 vCPUs	30000 GiB (4 * 7500 GiB NV
I3EN Metal	i3en.metal	768.0 GiB	96 vCPUs	60000 GiB (8 * 7500 GiB NV
R5DN Extra Large	r5dn.xlarge	32.0 GiB	4 vCPUs	150 GiB NV
I2 Extra Large	i2.xlarge	30.5 GiB	4 vCPUs	800
MSN 16xlarge	m5n.16xlarge	256.0 GiB	64 vCPUs	E
T2 Micro	t2.micro	1.0 GiB	1 vCPUs for a 2h 24m burst	E
D2 Eight Extra Large	d2.8xlarge	244.0 GiB	36 vCPUs	48000 GiB (24 * 2000 G
I3EN 3xlarge	i3en.3xlarge	96.0 GiB	12 vCPUs	7500 GiB NV
Z1D 3xlarge	z1d.3xlarge	96.0 GiB	12 vCPUs	450 GiB NV
X1E 16xlarge	x1e.16xlarge	1952.0 GiB	64 vCPUs	1920
RSN 24xlarge	r5n.24xlarge	768.0 GiB	96 vCPUs	E
I2 Eight Extra Large	i2.8xlarge	244.0 GiB	32 vCPUs	6400 GiB (8 * 800 G
R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	E
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	E
I2 Double Extra Large	i2.2xlarge	61.0 GiB	6 vCPUs	1600 GiB (2 * 800 G
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NV
MSA Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	E
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	E
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs for a 4h 4.8m burst	E
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 G
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NV
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NV
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	E
T2 Large	t2.large	8.0 GiB	2 vCPUs for a 7h 12m burst	E
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920
MSA 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	E
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	E
R5A Large	r5a.large	16.0 GiB	2 vCPUs	E
C3 High-CPU Large	c3.large	3.75 GiB	2 vCPUs	32 GiB (2 * 16
R5A 24xlarge	r5a.24xlarge	768.0 GiB	96 vCPUs	E
G3 16xlarge	g3.16xlarge	488.0 GiB	64 vCPUs	E
A1 Double Extra Large	a1.2xlarge	16.0 GiB	8 vCPUs	E
C4 High-CPU Extra Large	c4.xlarge	7.5 GiB	4 vCPUs	E
X1E Quadruple Extra Large	x1e.4xlarge	488.0 GiB	16 vCPUs	480
MSAD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	150 GiB NV


#thecloudistoodamn

1. What type of instances?
2. How many to spin up?
3. What base image?
4. On-demand or spot?
5. What storage service to use?
6. And then wait to start...
7. Not the end of the horror story:
 1. When to scale out?
 2. When to scale in?
 3. When to switch to different instance types?



R5A Eight Extra Large	r5a.8xlarge	256.0 GiB	32 vCPUs	
A1 Metal	a1.metal	32.0 GiB	16 vCPUs	
I2 Double Extra Large	i2.2xlarge	61.0 GiB	8 vCPUs	1600 GiB (2 * 800 GiB NVMe)
I3EN Double Extra Large	i3en.2xlarge	64.0 GiB	8 vCPUs	5000 GiB (2 * 2500 GiB NVMe)
M5A Extra Large	m5a.xlarge	16.0 GiB	4 vCPUs	
P3 Double Extra Large	p3.2xlarge	61.0 GiB	8 vCPUs	
T2 Double Extra Large	t2.2xlarge	32.0 GiB	8 vCPUs	for a 4h 4.8m burst
H1 Eight Extra Large	h1.8xlarge	128.0 GiB	32 vCPUs	8000 GiB (4 * 2000 GiB NVMe)
R5D 24xlarge	r5d.24xlarge	768.0 GiB	96 vCPUs	3600 GiB (4 * 900 GiB NVMe)
I3EN 6xlarge	i3en.6xlarge	192.0 GiB	24 vCPUs	15000 GiB (2 * 7500 GiB NVMe)
R4 High-Memory Eight Extra Large	r4.8xlarge	244.0 GiB	32 vCPUs	
T2 Large	t2.large	8.0 GiB	2 vCPUs	for a 7h 12m burst
X1 Extra High-Memory 16xlarge	x1.16xlarge	976.0 GiB	64 vCPUs	1920 GiB (3 * 640 GiB NVMe)
M5A 16xlarge	m5a.16xlarge	256.0 GiB	64 vCPUs	
R5 Metal	r5.metal	768.0 GiB	96 vCPUs	
R5A Large	r5a.large	16.0 GiB	4 vCPUs	
M5AD Extra Large	m5ad.xlarge	16.0 GiB	4 vCPUs	32 GiB (2 * 16 GiB NVMe)
				480 GiB (3 * 160 GiB NVMe)
				150 GiB NVMe

Why is there no “cloud button”?



~~Decision paralysis??~~ Go for Serverless!




Amazon
Lambda



Google Cloud Functions



Microsoft Azure Functions

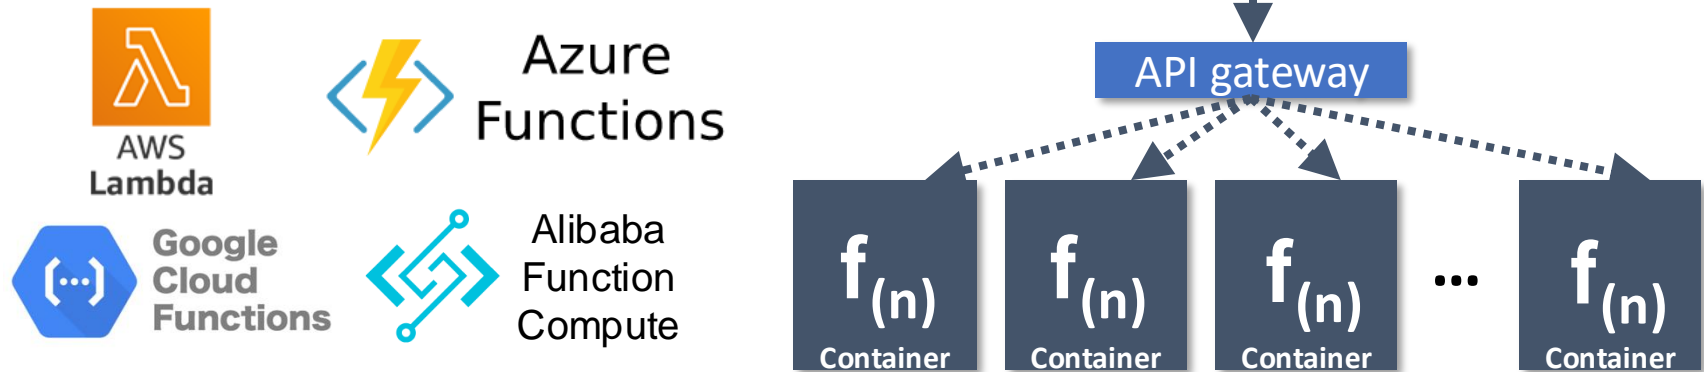


What is serverless computing?

What is serverless computing?

Serverless computing (Function-as-a-Service, or FaaS) is a **programming abstraction** that enables users to upload programs, run them at **(virtually)** any scale, and **pay only for the resources used**

- **Function-as-a-Service (FaaS):** Cloud functions as a basic deployment unit



Function-as-a-Service (FaaS)

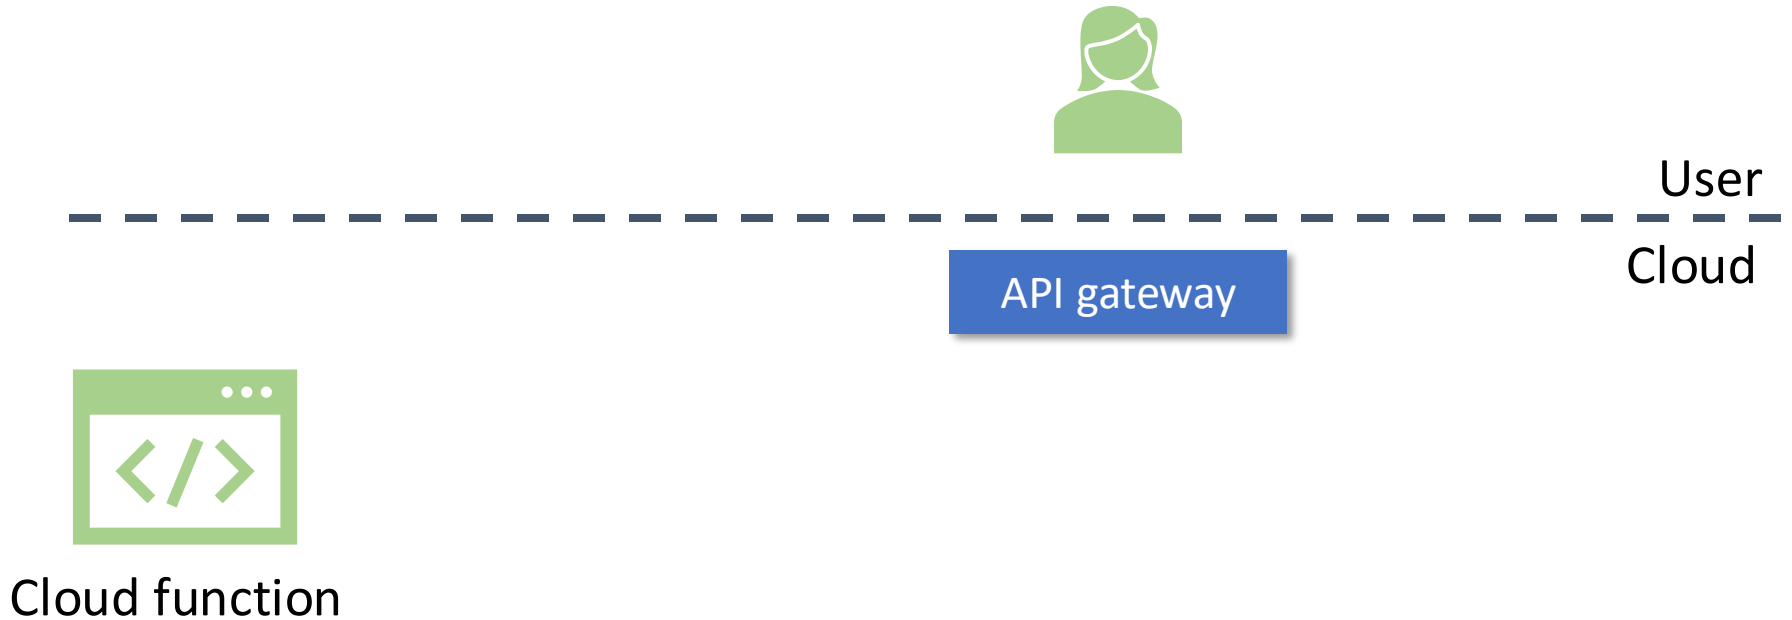


User
Cloud


Function-as-a-Service (FaaS)




Function-as-a-Service (FaaS)



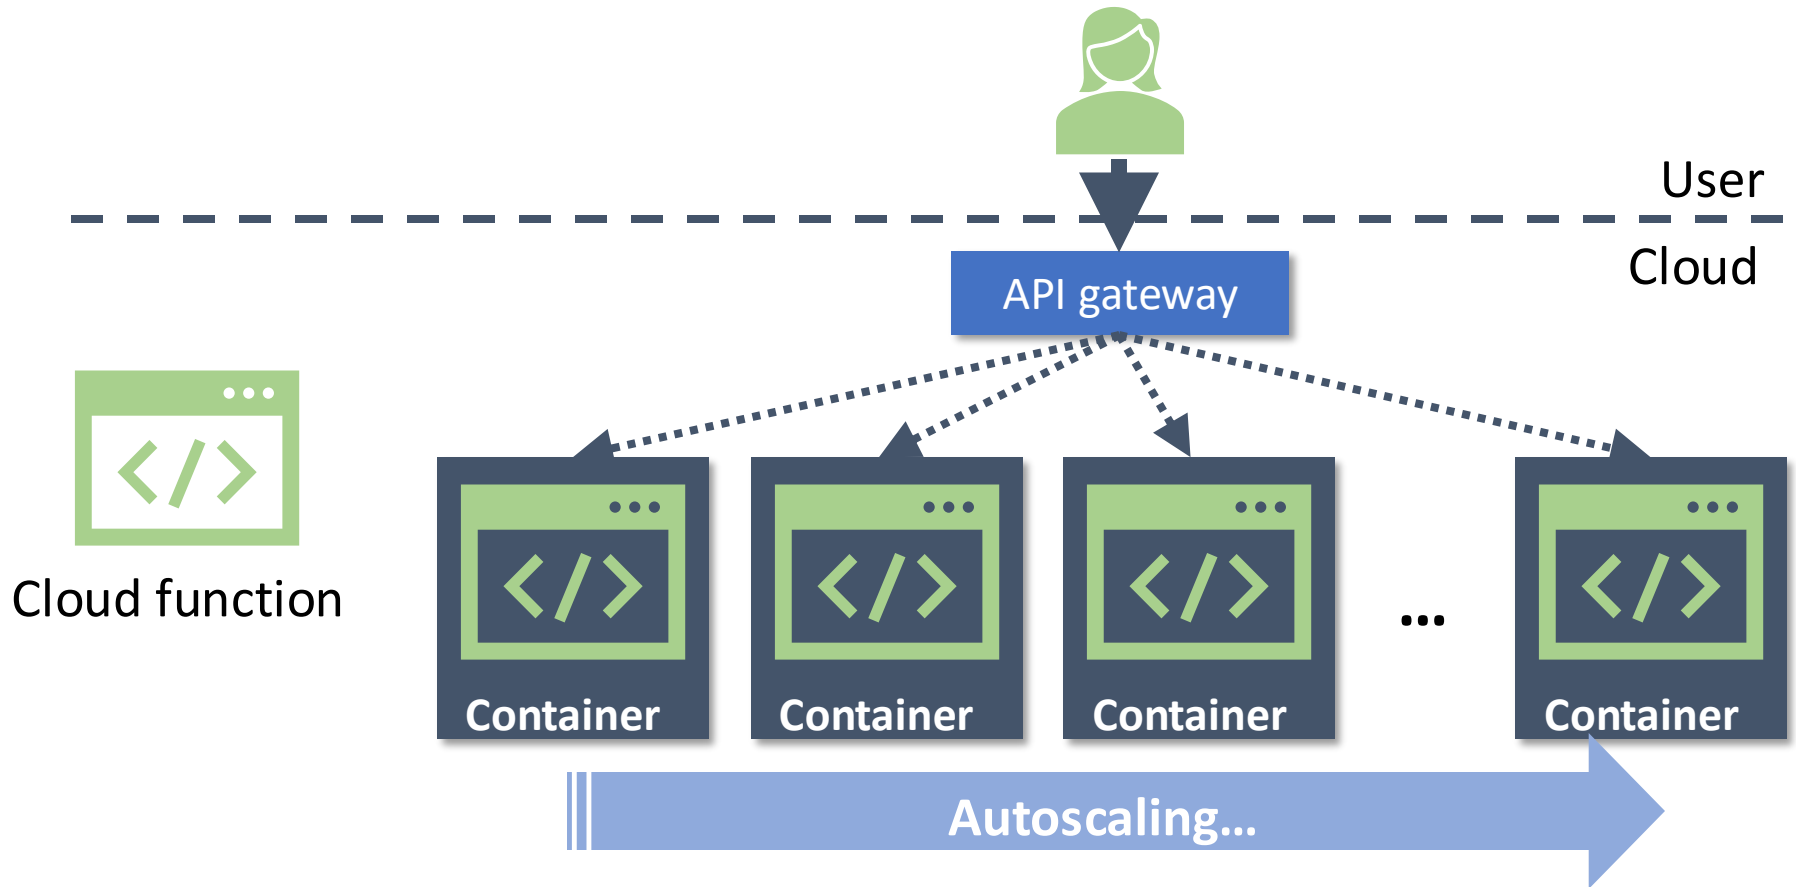
Function-as-a-Service (FaaS)



Function-as-a-Service (FaaS)



Function-as-a-Service (FaaS)



FaaS' core capability

1. (Provider) Manage a set of user-defined functions

FaaS' core capability

1. (Provider) Manage a set of user-defined functions
2. Take an event sent over HTTP or received from an event source

FaaS' core capability

1. (Provider) 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

FaaS' core capability

1. (Provider) 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

FaaS' core capability

1. (Provider) 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

FaaS' core capability

1. (Provider) 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

FaaS' core capability

1. (Provider) 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 the execution terminates

A car analogy

Car analogy



Own a car
(Bare metal servers)



Rent a car
(VPS)



City car-sharing
(Serverless)

Cars are parked **95%** of the time (loige.link/car-parked-95)

How much do you use the car?

<https://www.slideshare.net/loige/building-a-serverless-company-with-nodejs-react-and-the-serverless-framework-jsday-2017-verona>

AWS Lambda: 1st gen → current gen



- Lambda capacity config keeps **evolving**:

1 st gen	300 seconds 900 seconds (15 minutes)	} Current offering
	single-core two-core → up to 6 cores	
	1.5 GB → 10 GB memory	
	512 MB → up to 10GB of /tmp file system	

Python, Java, Node.js, Go, ...

Pricing:

- Fine-grained billing: **1-millisecond** billed duration
- \$0.20 per 1M requests (invocations charge \$)
- \$0.0000166667 for every GB-second (compute time charges \$\$)
- **6,000 1 GB Lambda functions for one second: 10¢**

Concept of serverless is not new

- Google App Engine
 - Fully managed platform as a service (PaaS) for developing and hosting web applications



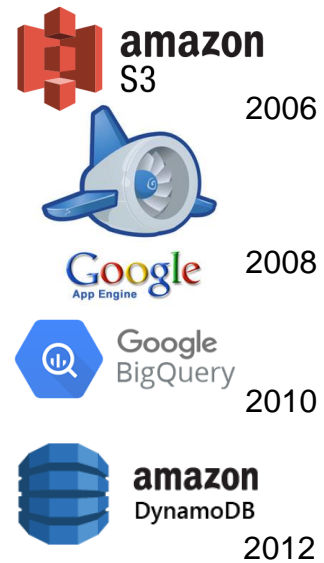
- Google BigQuery
 - Fully managed data warehouse
 - “Arbitrarily” large data and queries
 - Pay per byte being processed
 - No concept of server or cluster



- AWS S3
 - Fully managed object storage service
 - Pay per byte being stored and written
 - No server maintenance or resource scaling





Serverless trend in industry





Serverless trend in industry


- Major products
- Scope improvement
- Perf/Efficiency improvement

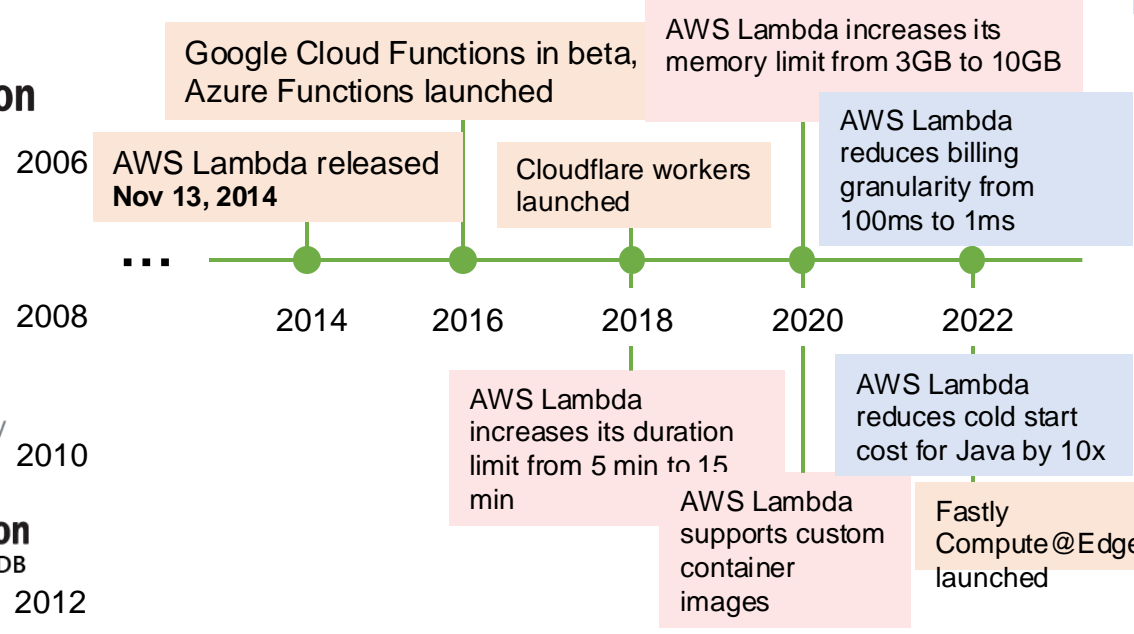
 **amazon**
S3



 **Google**
App Engine


 **Google**
BigQuery


 **amazon**
DynamoDB





Serverless trend in academia


- Major products
- Scope improvement
- Perf/Efficiency improvement

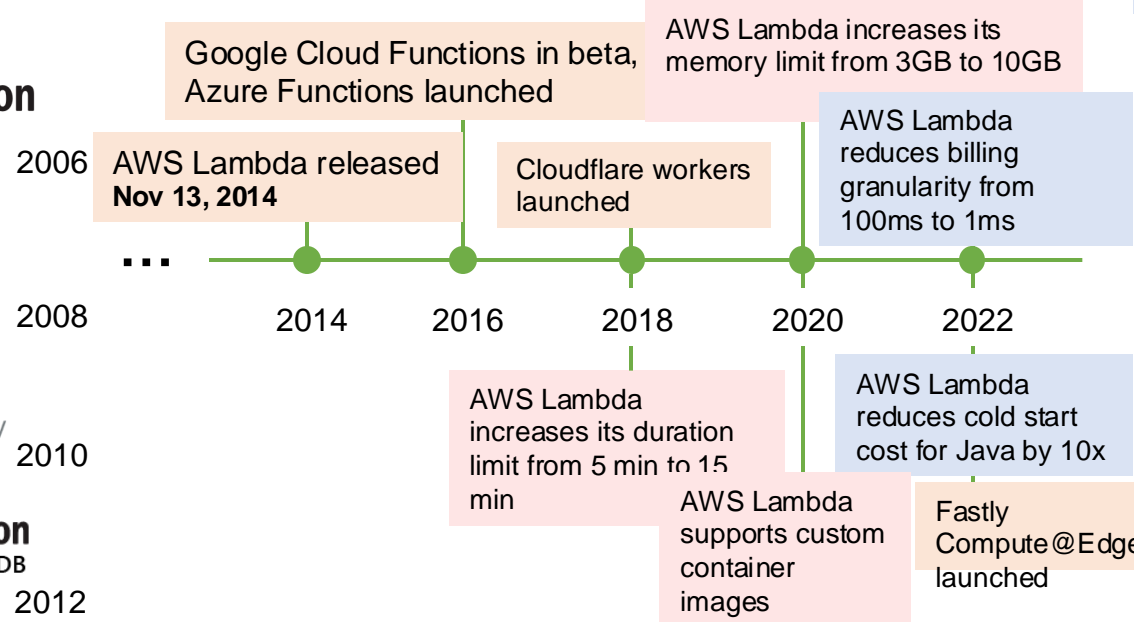
 **amazon**
S3




 **Google**
App Engine

 **Google**
BigQuery

 **amazon**
DynamoDB



'serverless computing' publications [Google Scholar]



Entering Serverless + (Gen)AI

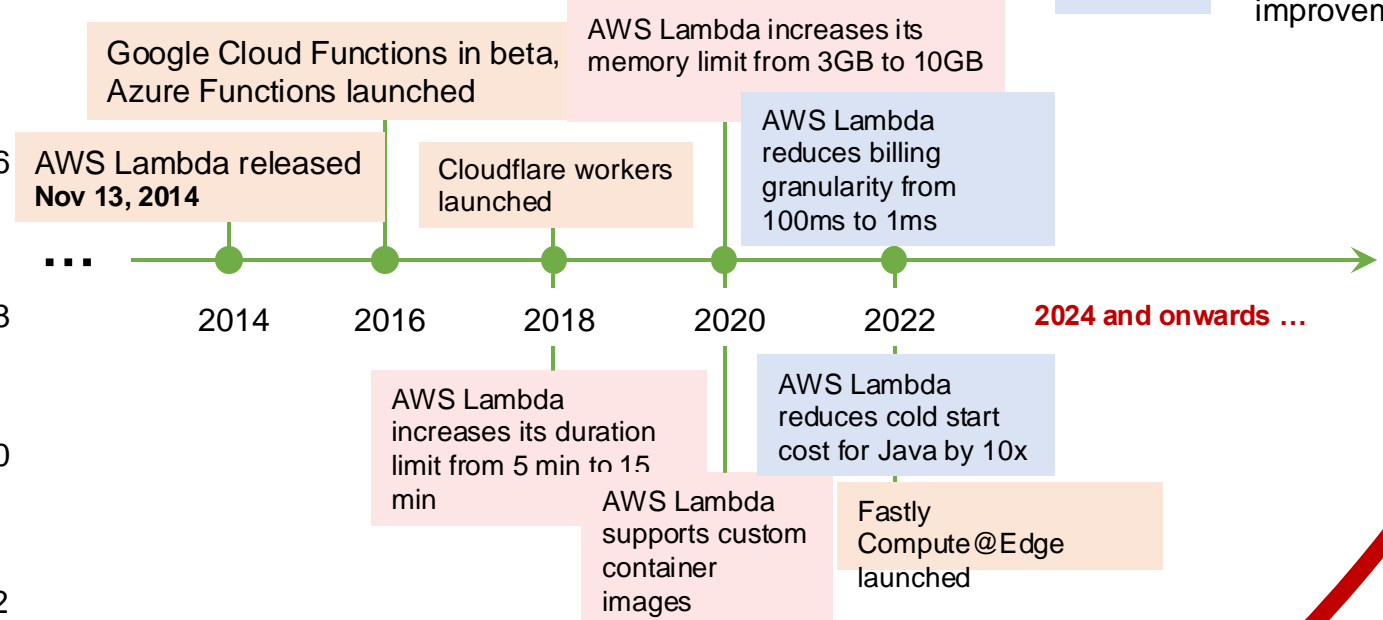
- Major products
- Scope improvement
- Perf/Efficiency improvement

amazon S3
 2006


Google App Engine
 2008

Google BigQuery
 2010

amazon DynamoDB
 2012



'serverless computing' publications [Google Scholar]



More exciting research and products in the era of Serverless AI

Entering Serverless + (Gen)AI

Serverless
infrastructure for AI | teams

Announcing Amazon Redshift Serverless with AI-driven scaling and optimization

AWS Lambda increases memory limit from 3GB to 10GB
Posted on: Oct 30, 2024

Serverless Computing In 2024: GenAI Influence, Security, 5G

Serverless enables developers to focus on writing code and building applications, while the cloud provider takes care of the underlying infrastructure.

Jan 4th, 2024 5:00am by Chris J. Preimesberger

Introducing The AWS AI Stack

Serverless Team • Sep 11, 2024

2024 and onwards ...

usenix
THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

ServerlessLLM: Low-Latency Serverless Inference for Large Language Models

Yao Fu, Leyang Xue, Yeqi Huang, and Andong Guo

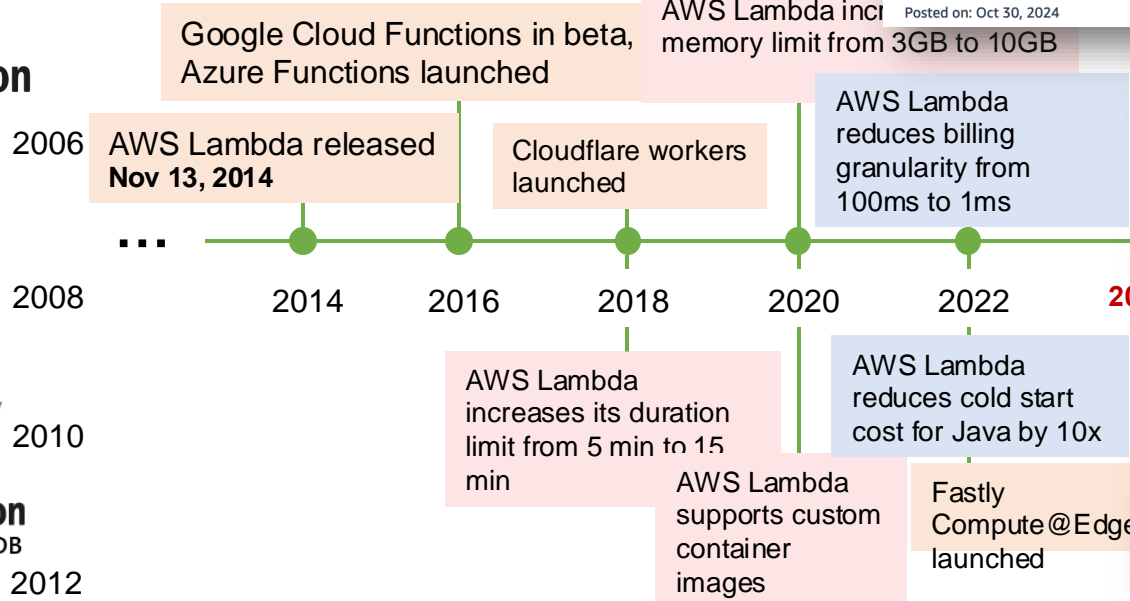
CLOUDFLARE The Cloudflare Blog

Workers AI: serverless GPU-powered inference on Cloudflare's global network


Alibaba Cloud Advances Its Gen AI Capabilities with Serverless Solution

- The cloud computing platform's serverless solution cuts inference costs by half
- Alibaba Cloud hosts its first AI & Big Data Summit in Singapore

and products in the era
of Serverless AI



'serverless computing' publications [Google Scholar]



Desirable properties of FaaS

- Operationally
 - “No-ops” – (almost) no configuration
 - Autoscaling down to 0
 - Closer to pay-per-use (rather than pay-per-allocation)
 - Fine-grained billing



FaaS today

- FaaS is used mostly for simple or coarse-grained tasks
 - Stateless, embarrassingly parallel tasks, simple workflows
 - ETL, software testing, API middleware, image processing, etc.
 - Glue to other serverless backends
- Lots of **problems** are limiting FaaS' **scope**
 - Poor performance (vs. time to run actual code) at scale
 - Mismatch of infrastructure support (e.g., today's OS not designed for FaaS)
 - Rule-breaking research needed to reimagine/broaden the scope
 - Very resource-inefficient and costly for serverless providers
 - Lack of support for accelerators
 - ...
- **Orders of magnitude slower and inefficient for many “killer” apps**
 - ML inference, microservices, ...


Limitations of FaaS

- **Banned** inbound network
- **No guaranteed** data availability
- Lambdas are **resource-constrained**
- Lambdas have **limited execution time**
- **High cold startup cost** and **invocation cost**


Lambda demo

Cloud evolution history – A virtualization story


Classic server app stack




1st generation: virtual machine (VM)



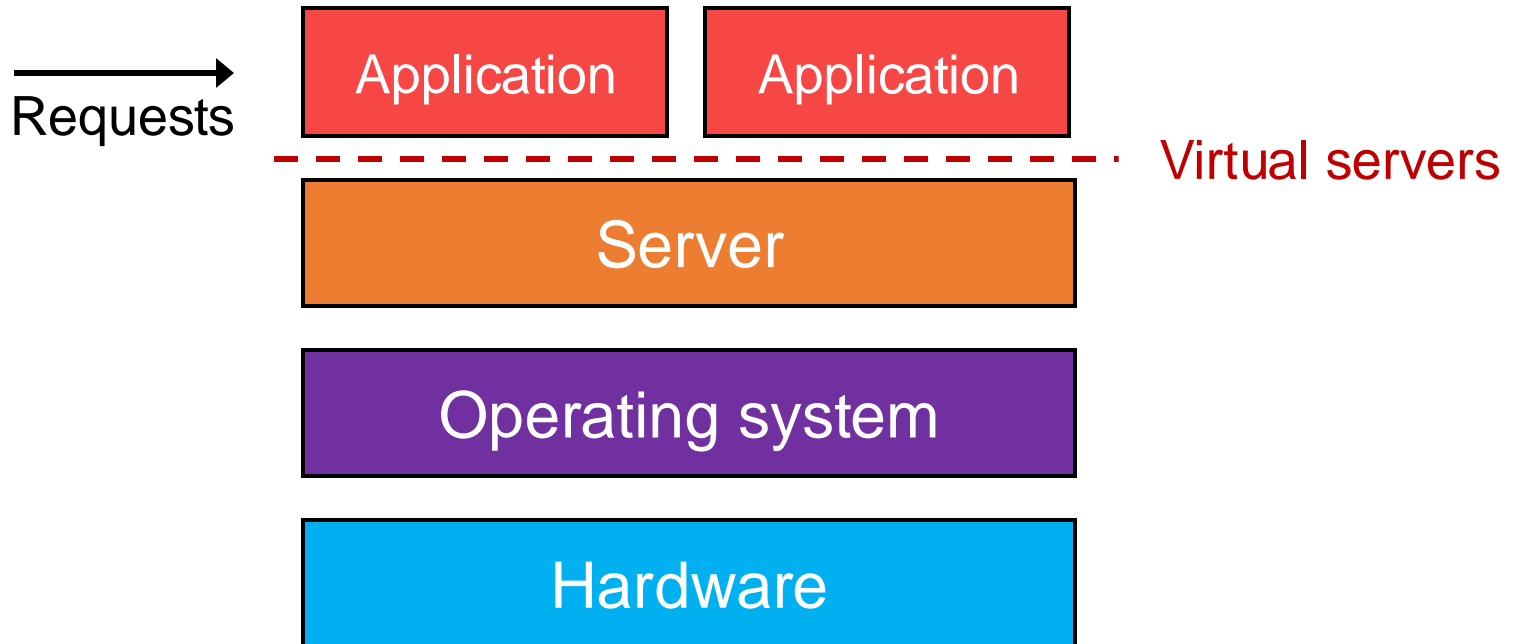
1st generation: virtual machine (VM)



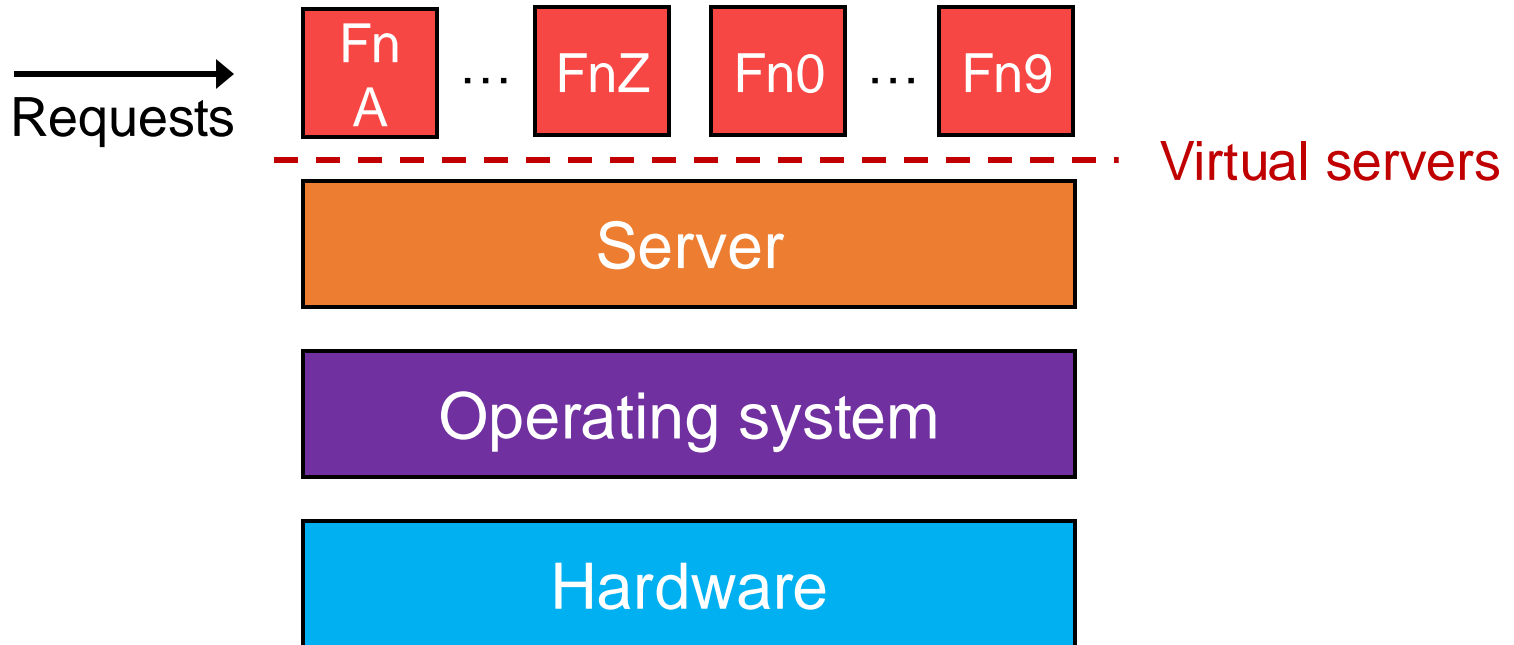
2nd generation: containers



3rd generation: serverless functions



3rd generation: serverless functions



Tradeoff discussion

Serverless functions
(AWS Lambdas)

Containers

VMs

Isolation?

Flexibility?

Overhead?