


deep dive into



EC2



Haid-und-Neu-Str. 18, 76131 Karlsruhe
Germany

about me

yes, I caught this myself



David Heidt

DevOps Engineer @msales

lots of aws, lots of ansible

I go fishing

I have two children (less time to go fishing, but still fun)

I play The Legend of Zelda

I'm not a layout professional

david.heidt@msales.com

Twitter: @witsches, @msalestech

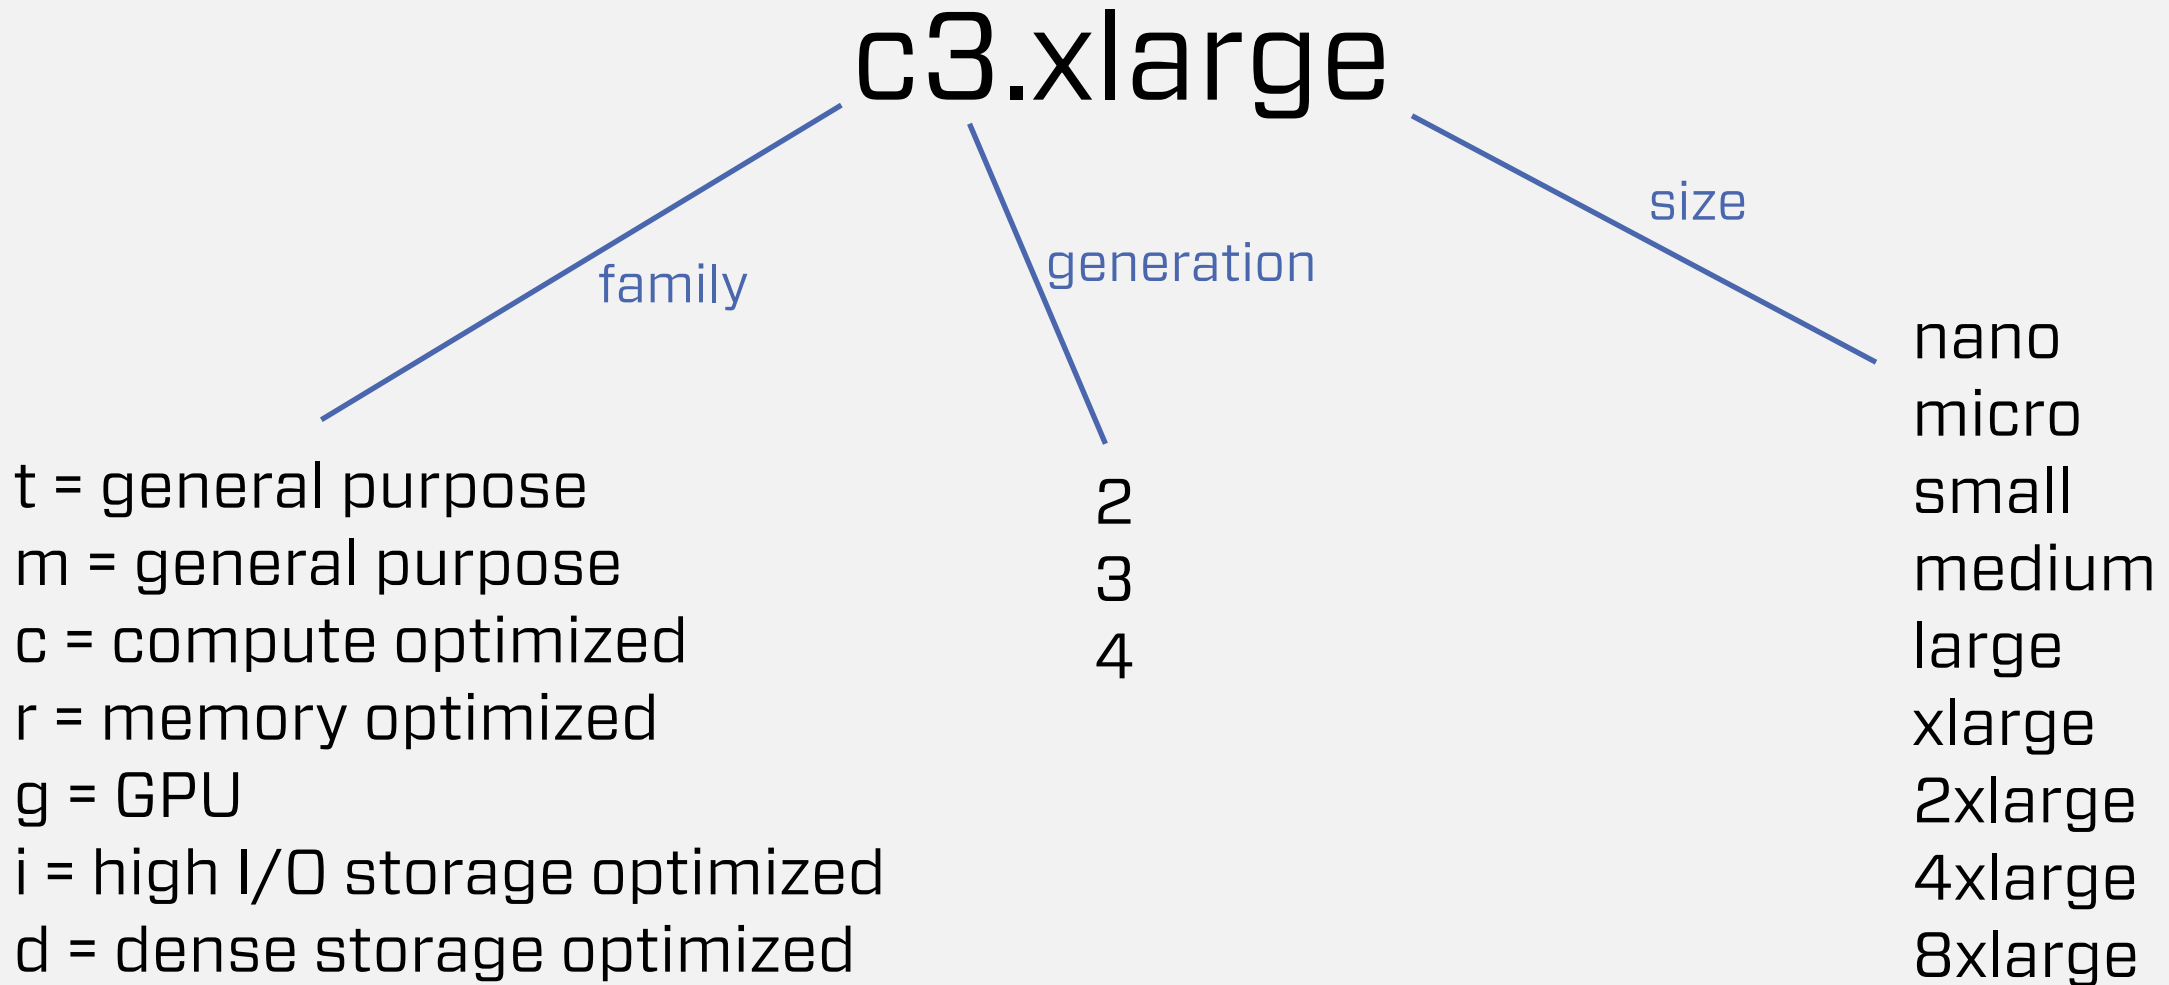
ec2

ec2 = elastic compute cloud

- ▲ launch virtual servers when you need them
- ▲ run them as long as you need them
- ▲ terminate when you don't need them anymore
- ▲ or stop them and pay less



instance types



t != m

general purpose, but different

- ▲ t-instances come with burstable performance & cpu credits
- ▲ IRL t-instances are ideal because they are slim & keen
- ▲ even eligible for free usage tier
- ▲ **t** stops at large, **m** starts at large (so there is no real competition)

instance launch details

- ▲ shutdown behaviour: stop|terminate
- ▲ termination protection
- ▲ Tenancy
- ▲ detailed Monitoring
- ▲ user data

instance storage

- ▲ ebs: magnetic, gp, provisioned iops
- ▲ ebs snapshots
- ▲ ebs performance relates to instance network performance
- ▲ ephemeral
- ▲ instances of the i-family

amis and the marketplace

- ▲ Amazon provides common basic images

for everything else there are:

- ▲ community amis
- ▲ marketplace amis
- ▲ make your own ami

pv vs hvm

- ▲ HVM takes advantage of Intel VT-x technology
- ▲ HVM machines run their own kernel
- ▲ IN PV, Application must go through VMM
- ▲ HVM is therefore usually a bit faster
- ▲ PV may be more secure in some cases

instance networking

- ▲ launch multiple interfaces on an instance
- ▲ ixgbevf driver, enhanced networking
 - higher performance (packets per second), lower latency, and lower jitter
 - not on all OSs available
 - not on all instances available
 - check with `ethtool -i <interface>`

locating the instances

- ▲ global presence
- ▲ multiple AZs/datacenters per region
- ▲ placement groups



so how does payment work?

Use the [calculator](#), Luke!

- ▲ on the fly instances
- ▲ spot instances
- ▲ reserved instances
- ▲ dedicated hosts

other ec2 components

elb (elastic load balancer)

- ▲ http or TCP loadbalancer
- ▲ SSL termination
- ▲ HA
- ▲ Health checking of backends
- ▲ cross -AZ

auto scaling

- ▲ launch or terminate instances based on cloudwatch metrics
- ▲ integrates smoothly with elb
- ▲ react to impacts while being asleep



vpc

VPC = virtual private cloud

is a separate product, but ec2 usage makes no sense without knowing about VPC

- ▲ public ips
- ▲ elastic IPs
- ▲ peering
- ▲ endpoints (currently s3 only)
- ▲ VPN connection to VPCs

security groups

- ▲ filter traffic on IP layer
- ▲ relate to other security groups (within VPC)
- ▲ instances can have multiple security groups
- ▲ once the name of a group is set, it cannot be changed ;)

what else?

cloudwatch

- ▲ metrics
- ▲ custom metrics
- ▲ alarms
- ▲ autoscaling actions

ecs = elastic container service

- ▲ deploy containers to ec2
- ▲ usually docker
- ▲ no extra cost, standard ec2 charges apply

hints for ec2 in production

is a separate product, but ec2 usage makes no sense without knowing about VPC

- ▲ use the aws cli/api or config management!
- ▲ every instance is replaceable
- ▲ always start with elastic IPs if you have static services
- ▲ if you want to scale, keep an eye on your account limits
- ▲ re-use security groups and name them accordingly
- ▲ make your own amis



Thanks.

msales GmbH

Haid-und-Neu-Str. 18, 76131 Karlsruhe
Germany

follow us @msalestech

info@msales.com

www.msales.com

+49 721 91138 0