Public Cloud Networking
A Parallel Universe with a Different Geometry

Ivan Pepelnjak (ip@ipSpace.net)
Network Architect

ipSpace.net AG

Who is Ivan Pepelnjak (@ioshints)

Past

- Kernel programmer, network OS and web developer
- Sysadmin, database admin, network engineer, CCIE
- Trainer, course developer, curriculum architect
- Team lead, CTO, business owner

Present

- Network architect, consultant, blogger, webinar and book author

Focus

- SDN and network automation
- Large-scale data centers, clouds and network virtualization
- Scalable application design
- Core IP routing/MPLS, IPv6, VPN

CISCO CCIE EMERITUS

vmware vEXPERT

Should I Stay or Should I Go?

You might find this relevant if…

- You're trying to make sense of how public cloud networking works (and why it behaves that way)
- You have to design (tenant) networking solutions in public clouds

You should have a beer if…

- You've successfully deployed large-scale application architectures across multiple major public clouds
- You're working for a large web property (FANG comes to mind)
- You're running a large private- or public cloud (in which case you probably experienced all I have to say)

What a Weird Land We're Entering

There's no layer-2 in (sane) public cloud

- VMware-based approximations don't count
- We're talking about stuff that scales beyond 1000 hosts

How am I supposed to:

- Move virtual machines
- Implement high-availability clusters
- Deploy firewall clusters
- Migrate workloads from on-premises data center

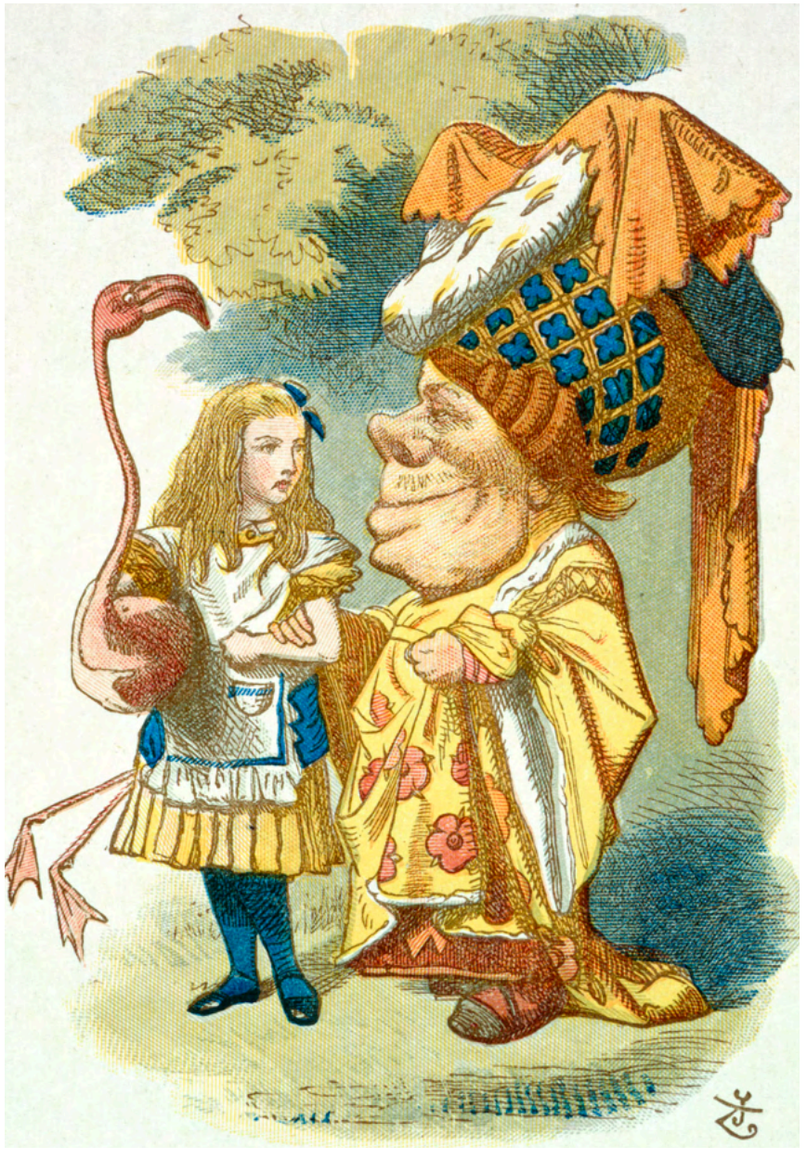

AWS versus Azure: Common Concepts

- Isolated tenant routing domains (VPC, VNet...)
- Multiple subnets within a tenant routing domain
- IP and MAC addresses assigned by the orchestration system
- Strict IP+MAC RPF checks (can be disabled)
- Routing controlled by the orchestration system

Consequences

- You cannot change VM IP or MAC address without an orchestration system API call
- You cannot use FHRP
- Most MAC+IP high availability hacks don’t work
- You cannot run a routing protocol within the cloud to influence forwarding decisions

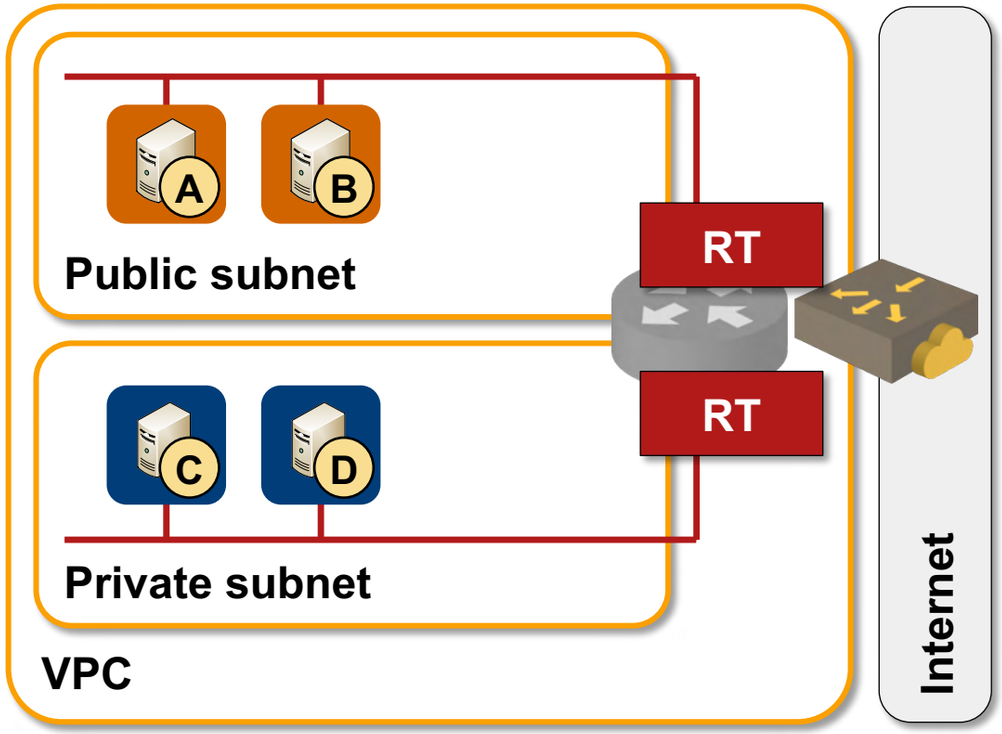

AWS Networking 101

- Each subnet is limited to a single availability zone
- Unicast IPv4 + IPv6 forwarding
- Limited IPv4 multicast support
- Unicast MAC forwarding within the subnet
- No L2 flooding
- Each subnet can have a different route table
- There's no way to influence intra-VPC packet forwarding

Consequences

- Service insertion is interesting
- Intra-VPC service insertion is **really** hard
  (and usually involves a lot of NAT duct tape)
- You could do routing tricks *within a subnet* but not *across subnets*

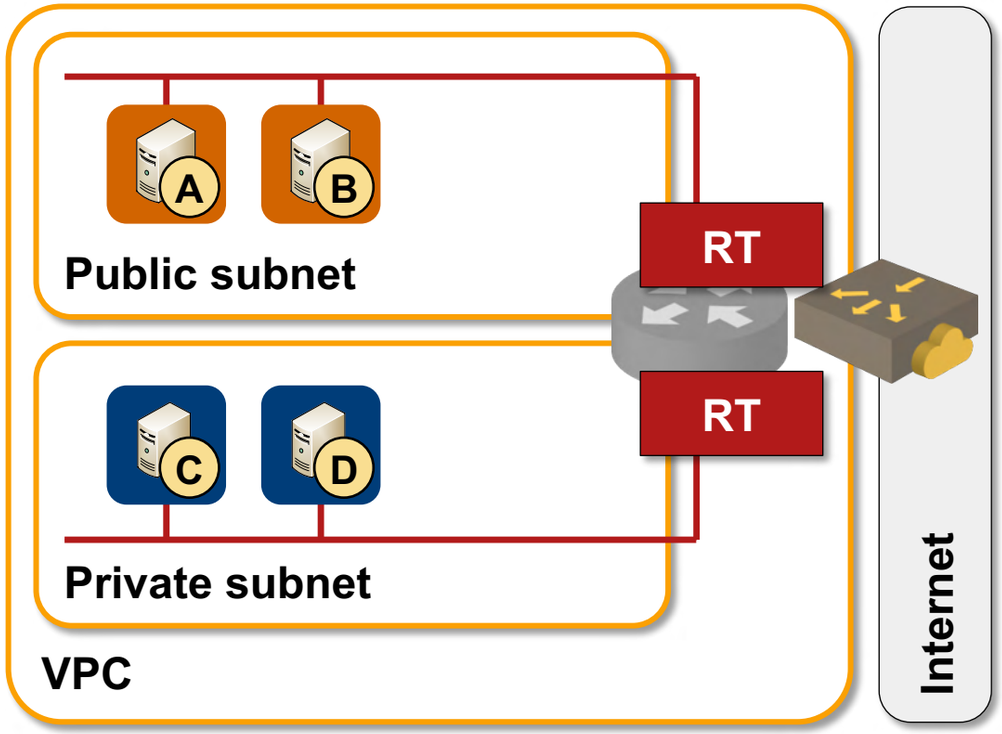

Azure Networking 101

- Subnets span availability zones
- Unicast IPv4 + IPv6 forwarding
- No L2 forwarding (every instance connected directly to a router)
- ARP always returns the MAC address of Azure router
- Each subnet could have a different route table
- Route tables can contain intra-VNet prefixes

Consequences

- Service insertion is relatively easy (but messy)
- Building application swimlanes tied to availability zones is hard(er)

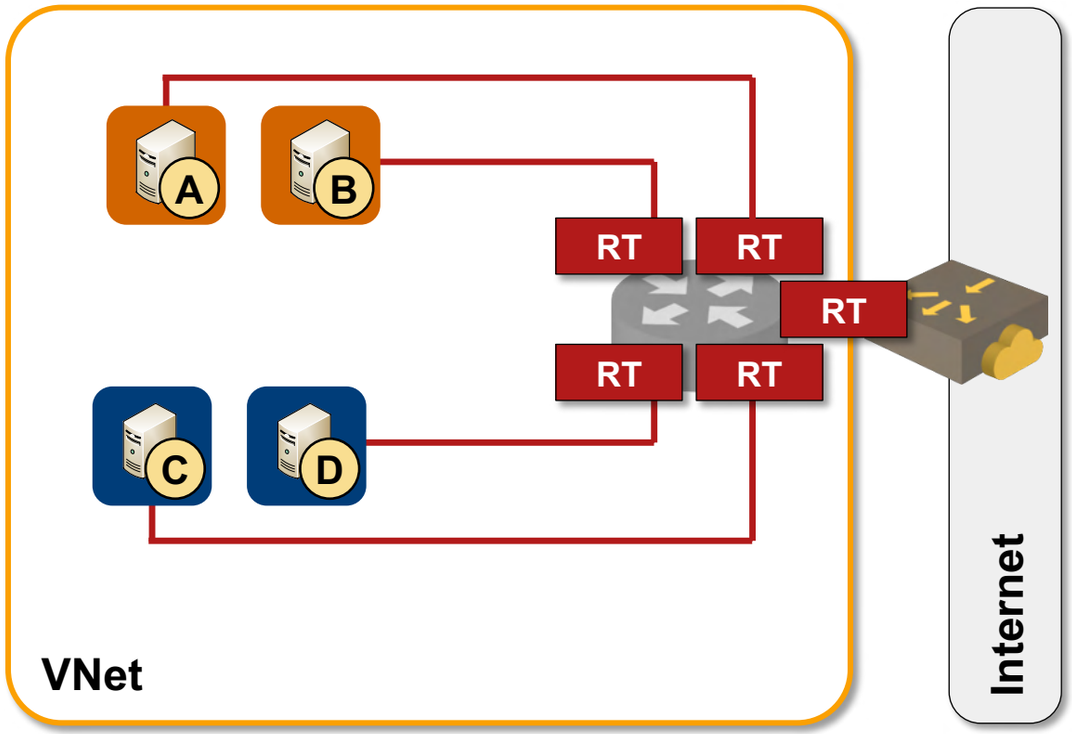

Why Couldn't They Be The Same?

- Convergent evolution
- Different audiences
- Different scalability goals?
- Fixing different problems
- Solving the same problem in different ways (aka *leverage the investment*)

Consequences

- Nobody wants to be limited to the least common denominator
- Real-life tools have cloud-specific plugins or modules (Terraform, Ansible)
- Multi-cloud works best in PowerPoint

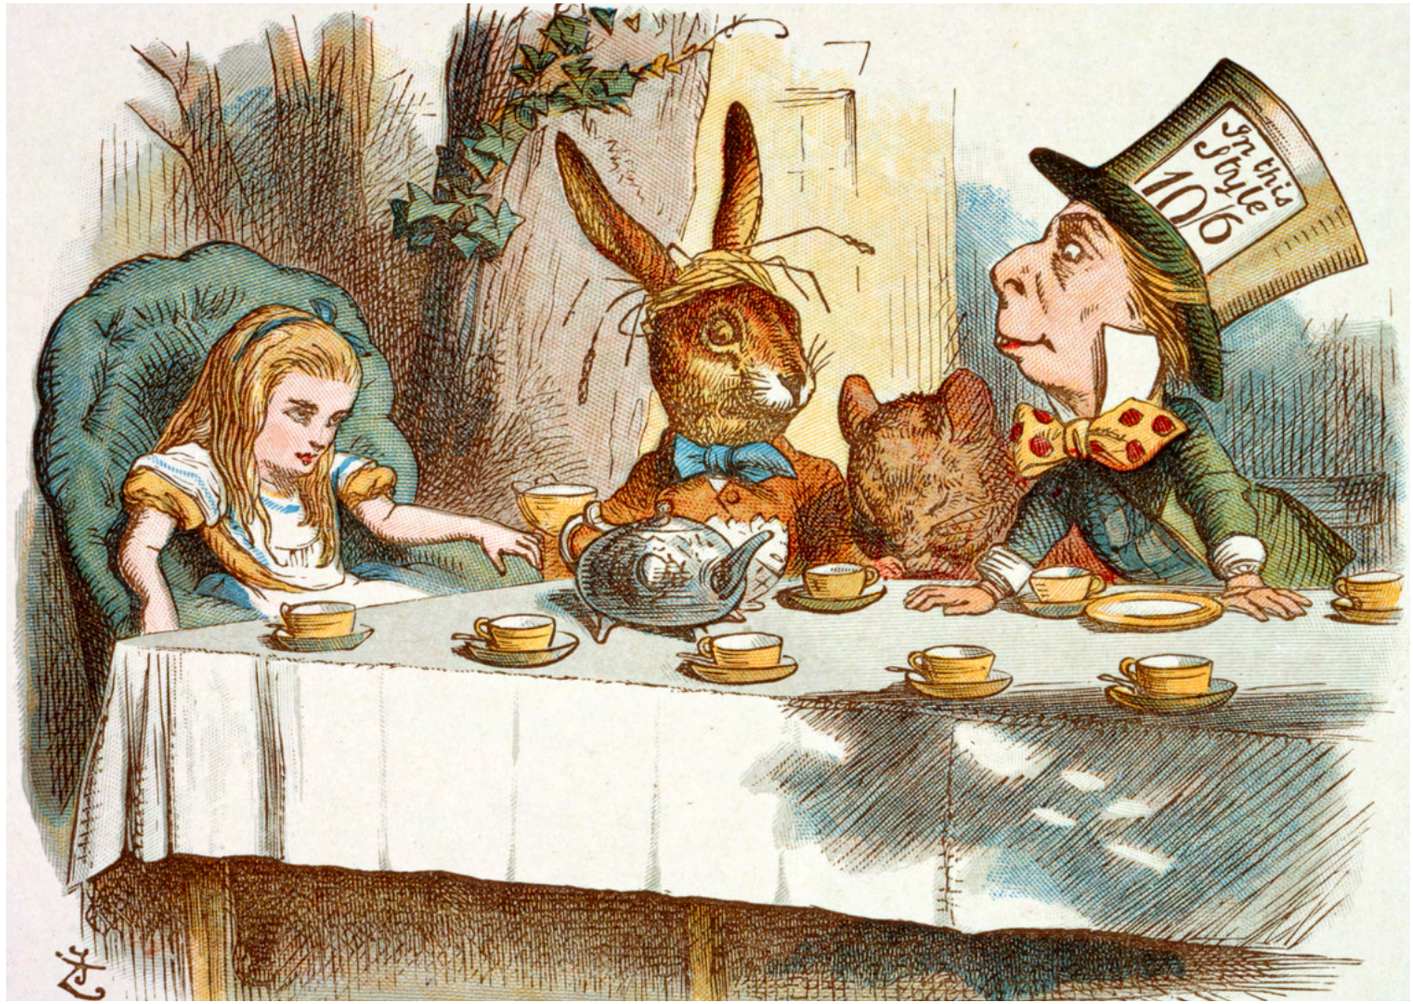

What Can You Do?


FORGET POWERPOINT
IT'S TIME FOR THE RED PILL

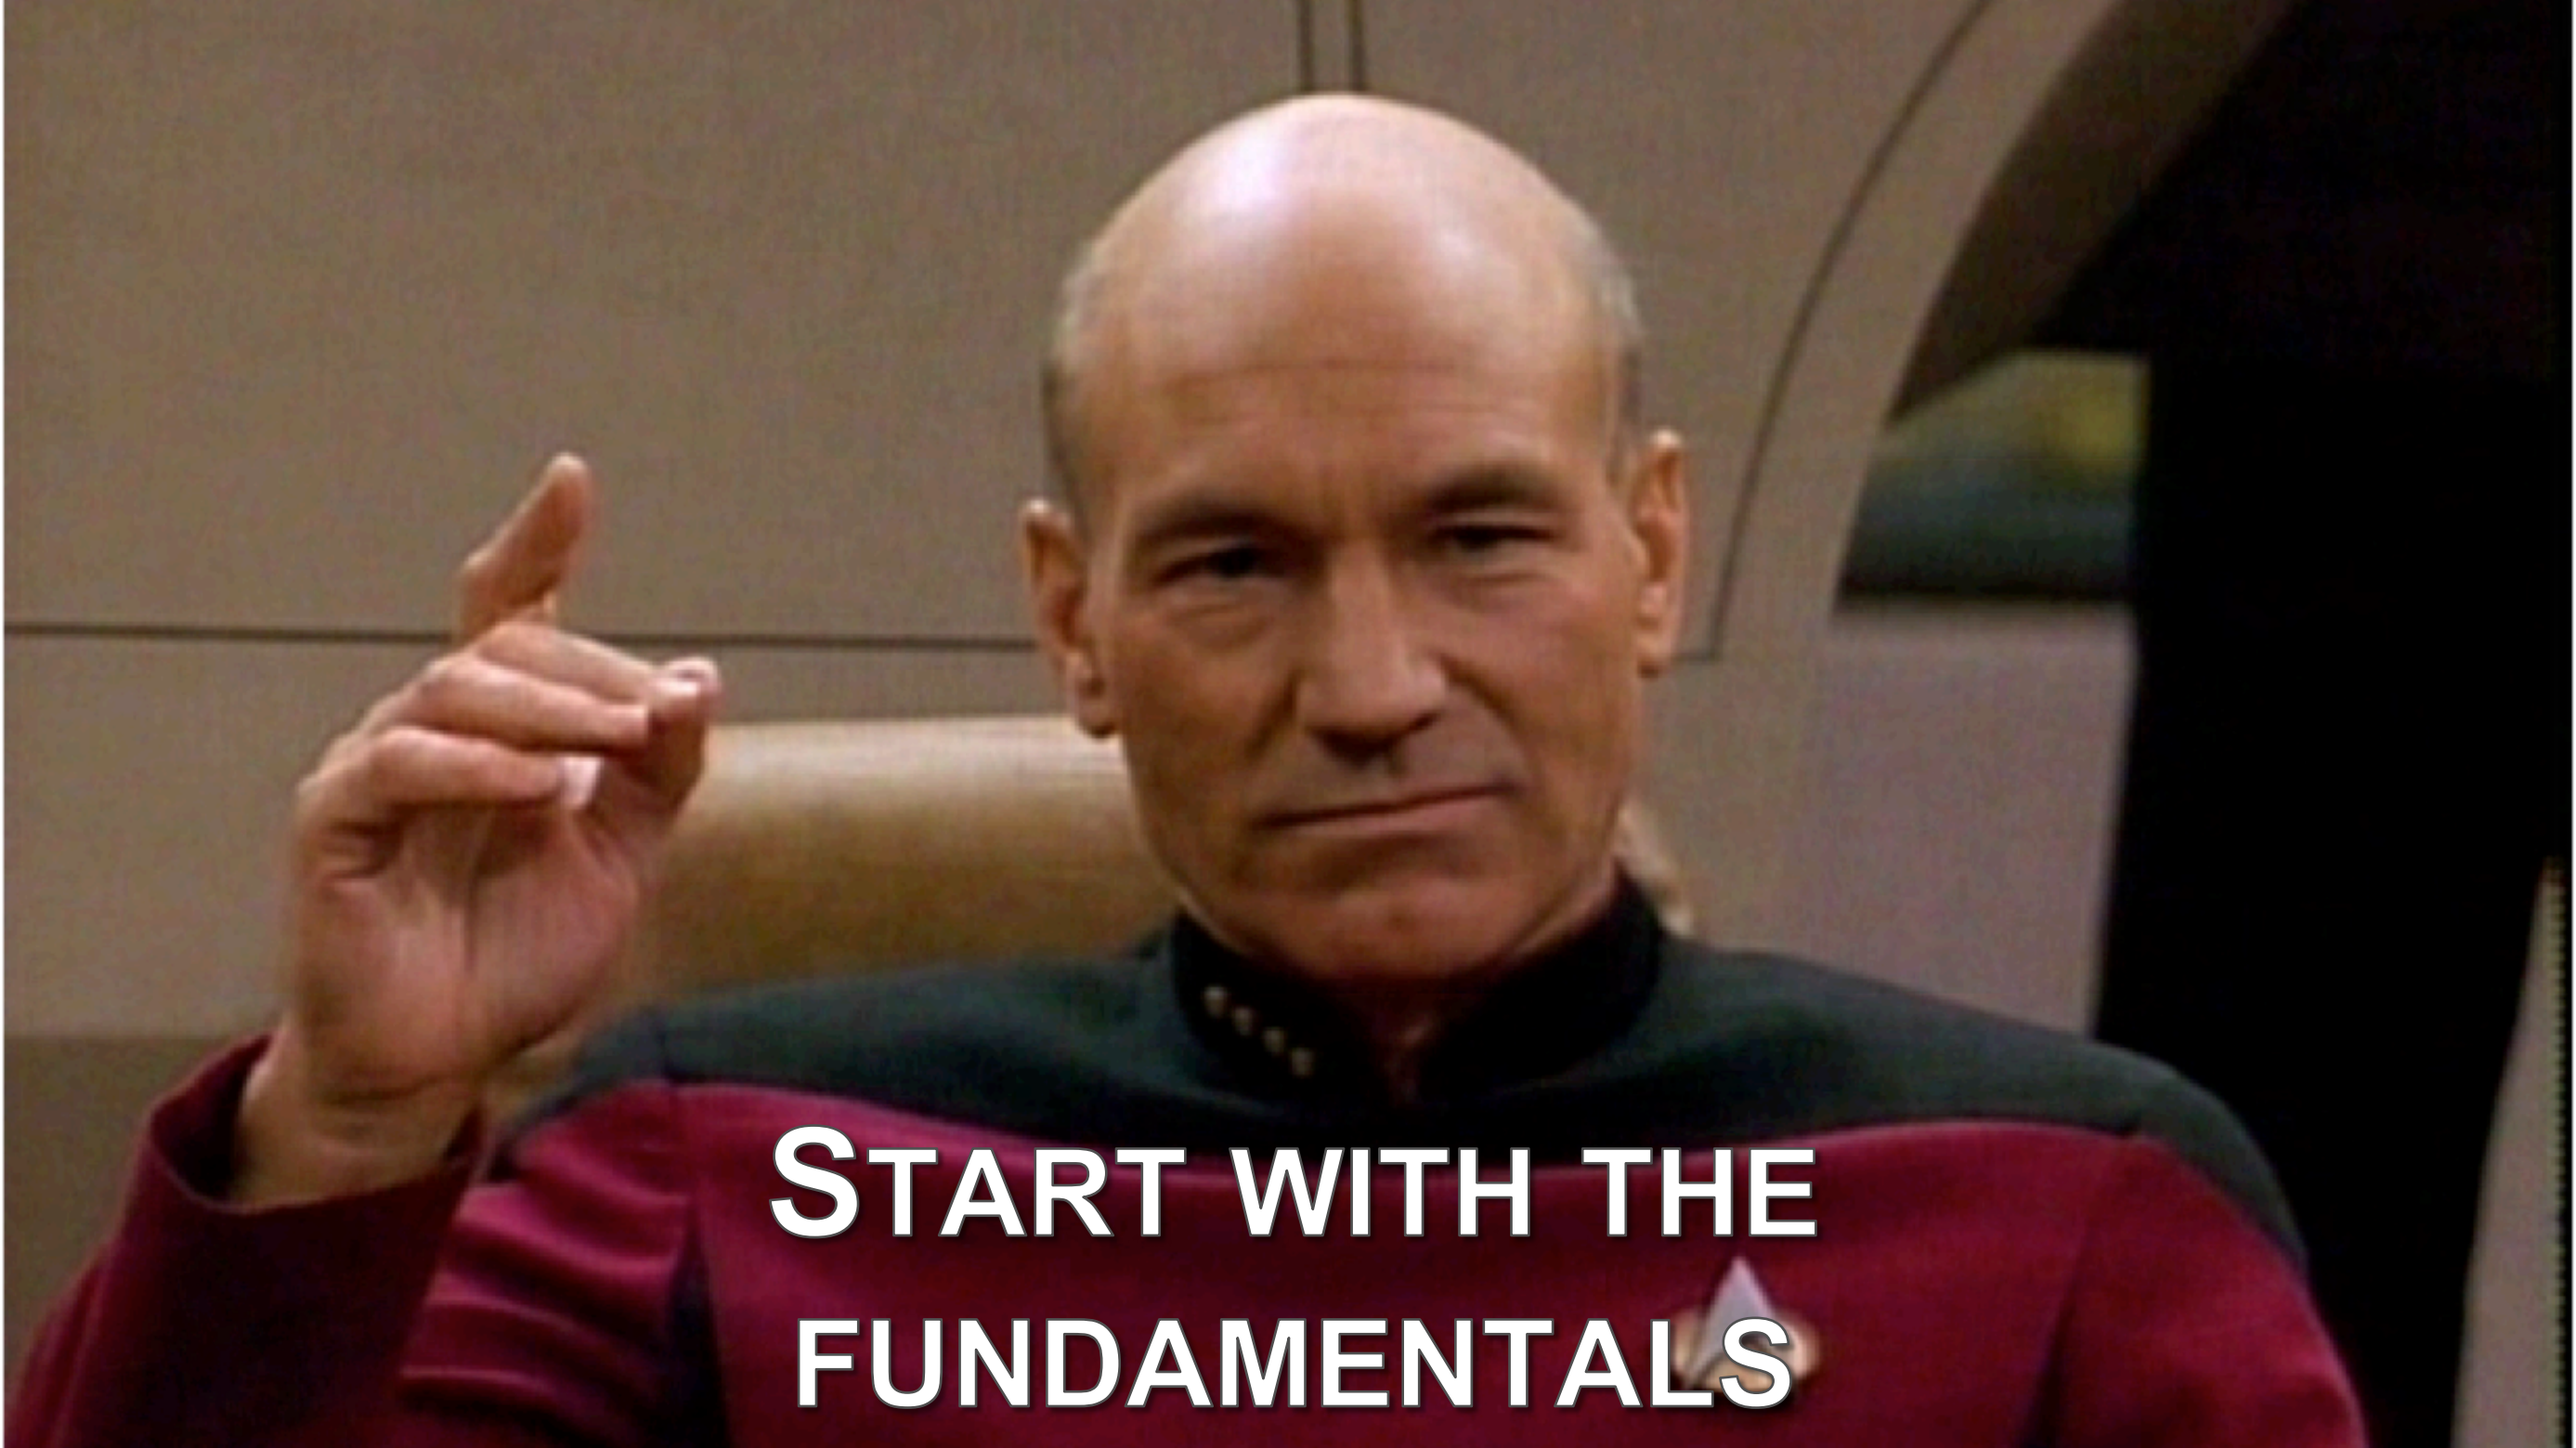
START WITH THE
FUNDAMENTALS

It’s Just Another Case of Alternate Geometries

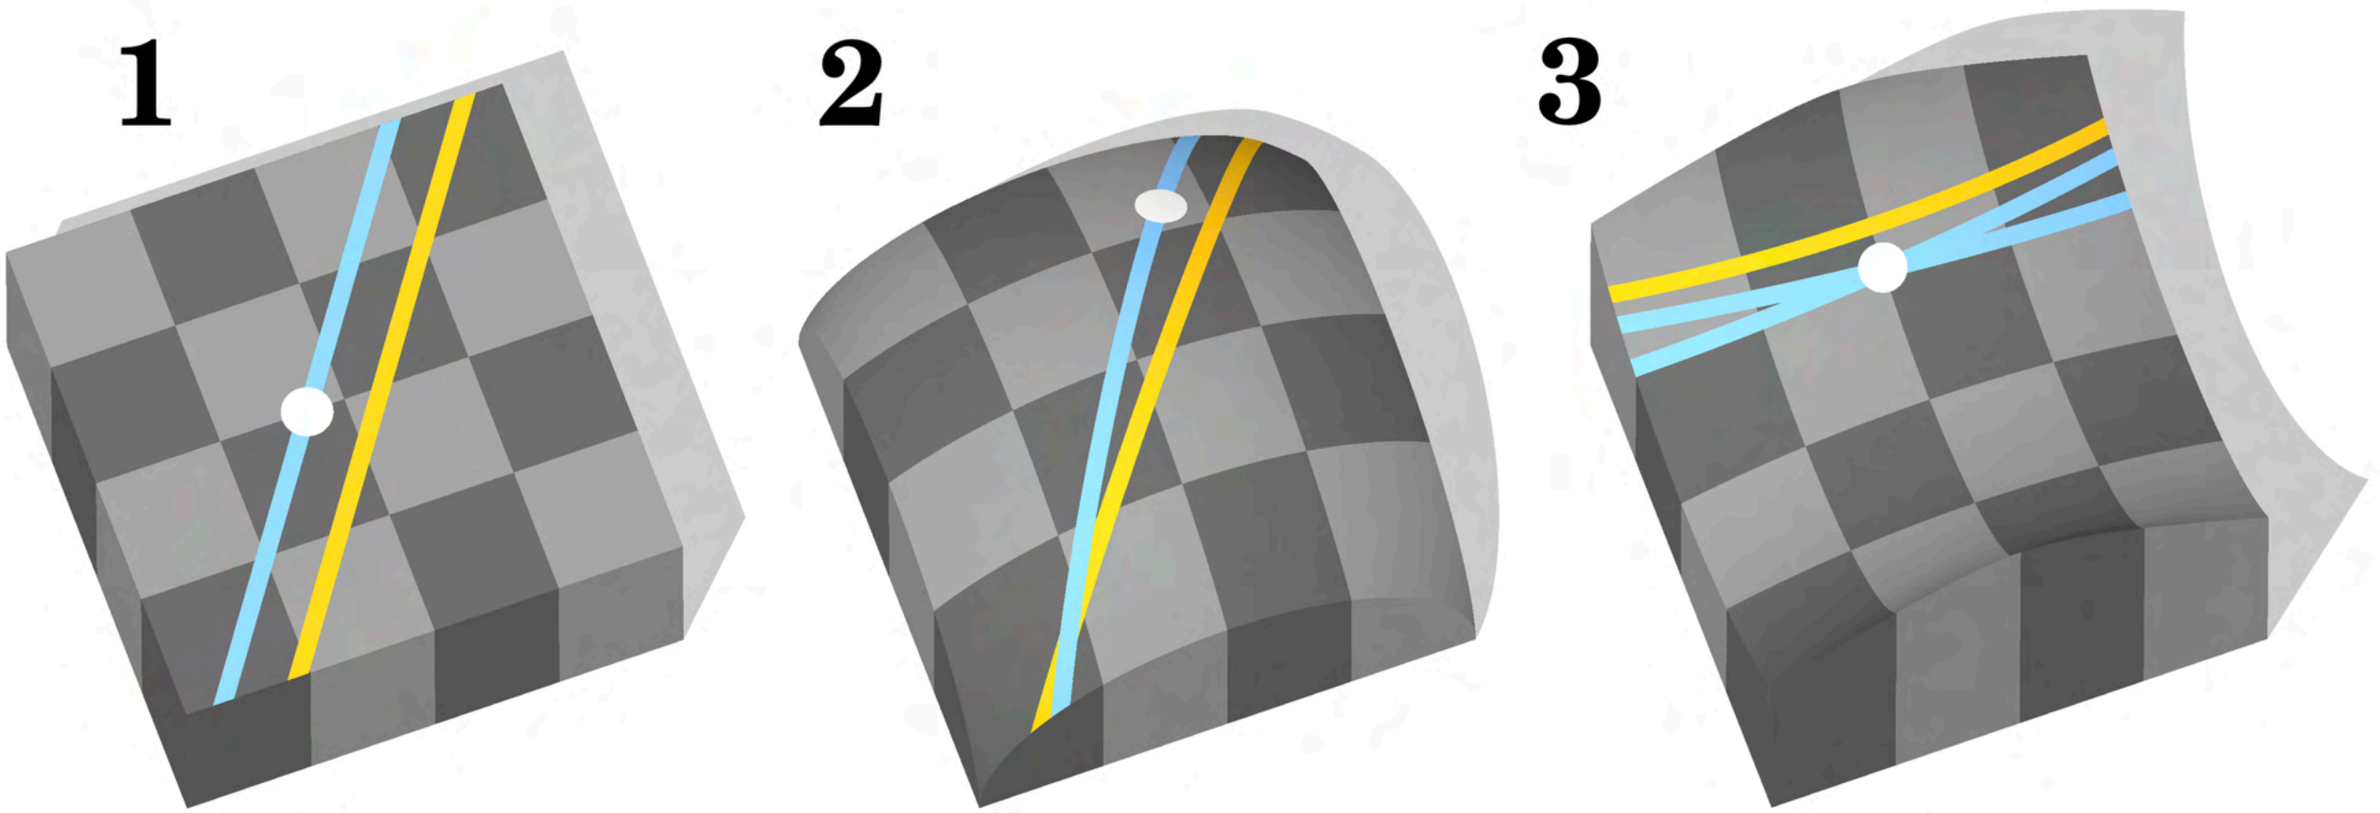


Always Wonder…
And Ask “WHY?”

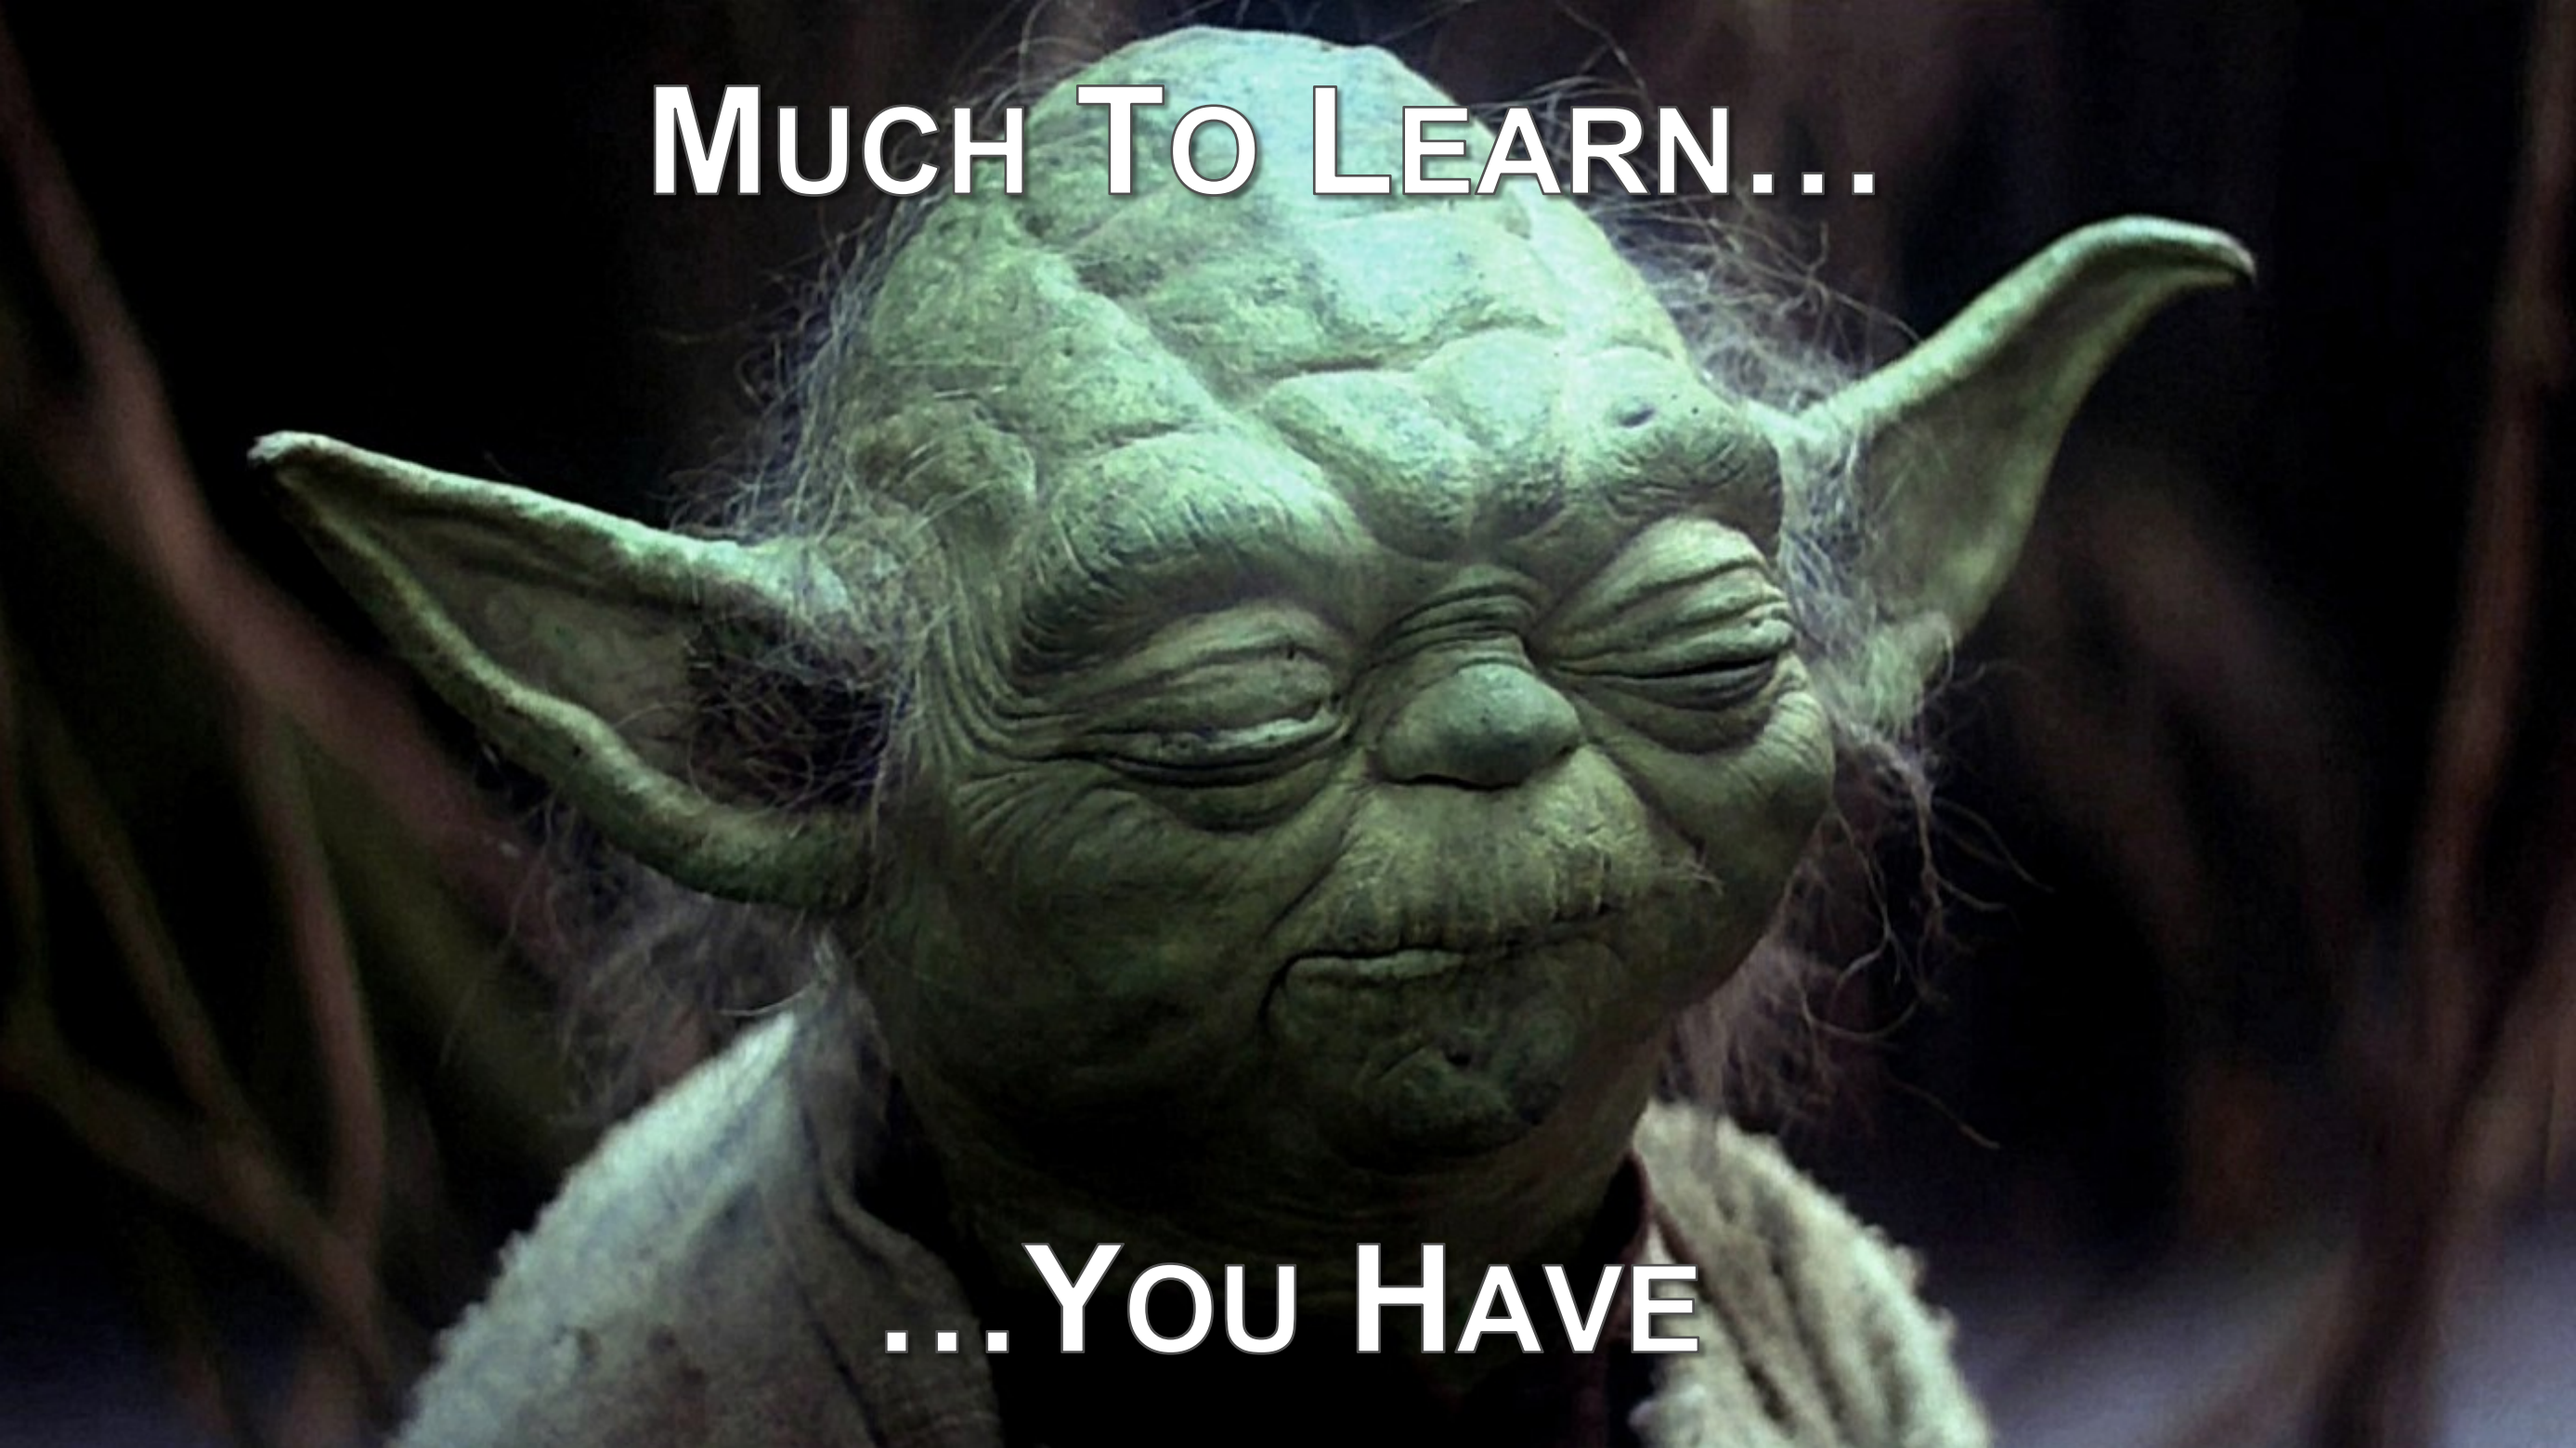
Much To Learn...
...You Have


MASTER THE CLOUD…
…YOU WILL

ipSpace.net Networking in Public Cloud Deployments Course

Core modules

- Public cloud principles
- Automation and infrastructure-as-code
- Compute and storage
- Networking in public clouds
- Network security
- Network services

Interesting concepts

- IPv6 in public cloud deployments
- Generic public cloud security (authentication, logging, monitoring, testing)
- High-availability architectures
- Hybrid clouds and multi-cloud






Hands-On Exercises

Practice public cloud networking while learning about it:

- Plan your own public cloud deployment
- Select your automation tools
- Deploy a simple compute/storage solution
- Create a full-blown virtual network
- Convert your virtual network into a dual-stack deployment
- Add security components
- Add application firewall and load balancers
- Deploy your workload across multiple availability zones and regions
- Add VPN connectivity

Choose your own environment

- Public cloud of your choice (we can help you if you choose AWS or Azure)
- Git is recommended, use any public repository you feel comfortable with
- Use an automation tool of your choice (but use it)
- Use tasks from the hands-on exercises or something equivalent from your $work

Working on the hands-on exercises

- Whenever you have time
- All submitted exercises are reviewed
- Work alone or in a team

Questions?

Web:	ipSpace.net
Blog:	blog.ipSpace.net
Email:	ip@ipSpace.net
Twitter:	@ioshints

Data center:	ipSpace.net/NextGenDC
Automation:	ipSpace.net/NetAutSol
Webinars:	ipSpace.net/Webinars
Consulting:	ipSpace.net/Consulting