Because learning changes everything®

How We Survived (and Thrived) During the Pandemic and Helped Millions of Students Learn Remotely

- **Chinmay Tripathi (He/Him)**
- Sr. Director, Engineering

About McGraw Hill

We are a learning science company

- We deliver educational technology for both K-12 and higher education.
- We partner with 14,000+ authors and educators.
- Our students have answered 12.8 billion questions!

Our Technology Scale

- 200+ million interactions per month
- 80+ development teams; 400+ applications/services
- 150+ accounts/networks over multiple cloud provider
- 50+ Kubernetes and Amazon ECS clusters
- 4,000+ EC2 instances

Mc Graw Hill

Covid-19 Response Timeline

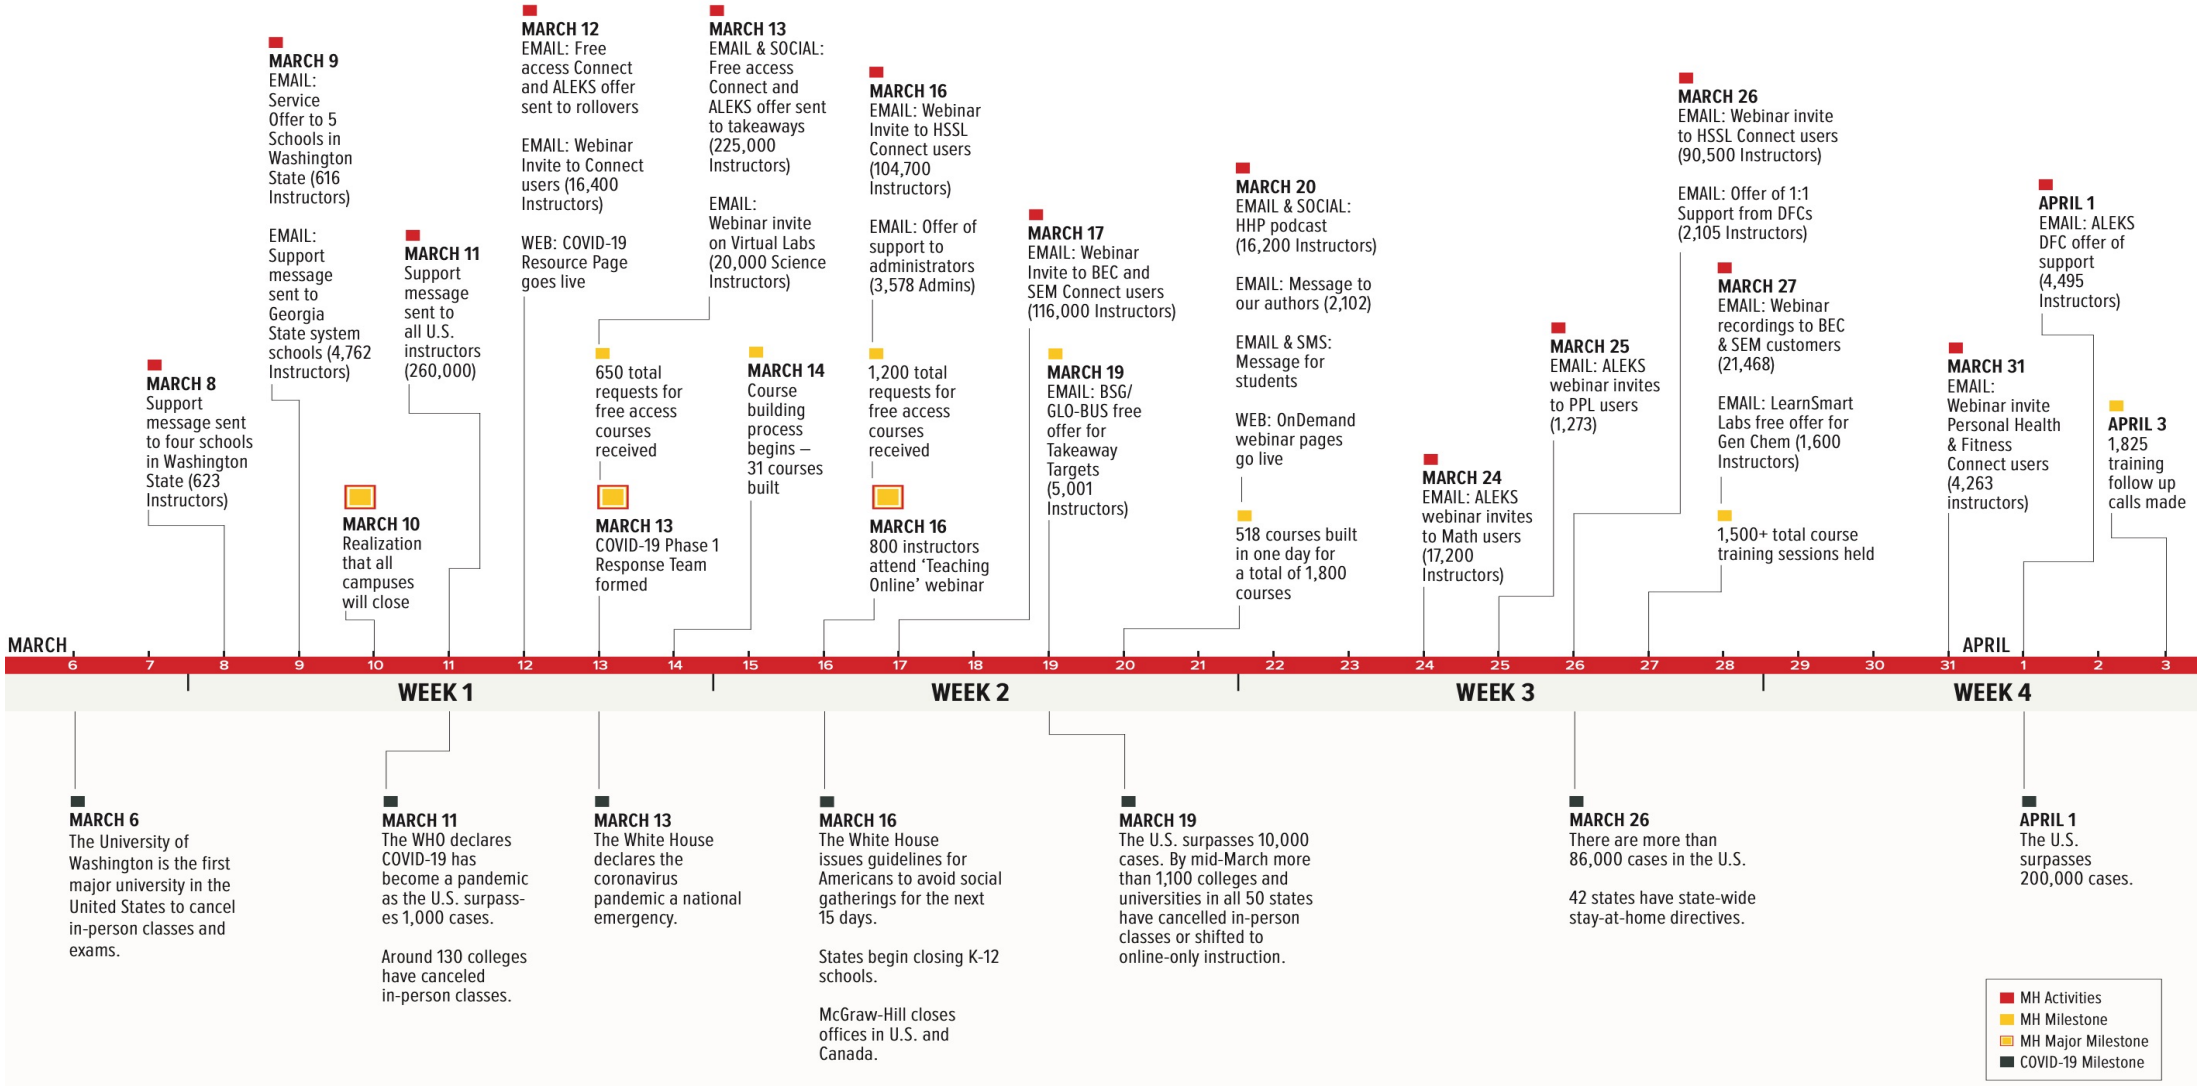

Covid-19 Response by Numbers

- Offer support to schools and Colleges
- 100s of webinars attended by thousands of instructors
- 1000s of training sessions were held
- 1000s of free access courses
- 100,000+ students completed free access courses
- **Zero platform downtime**

Mc Graw Hill

Covid-19 Impact

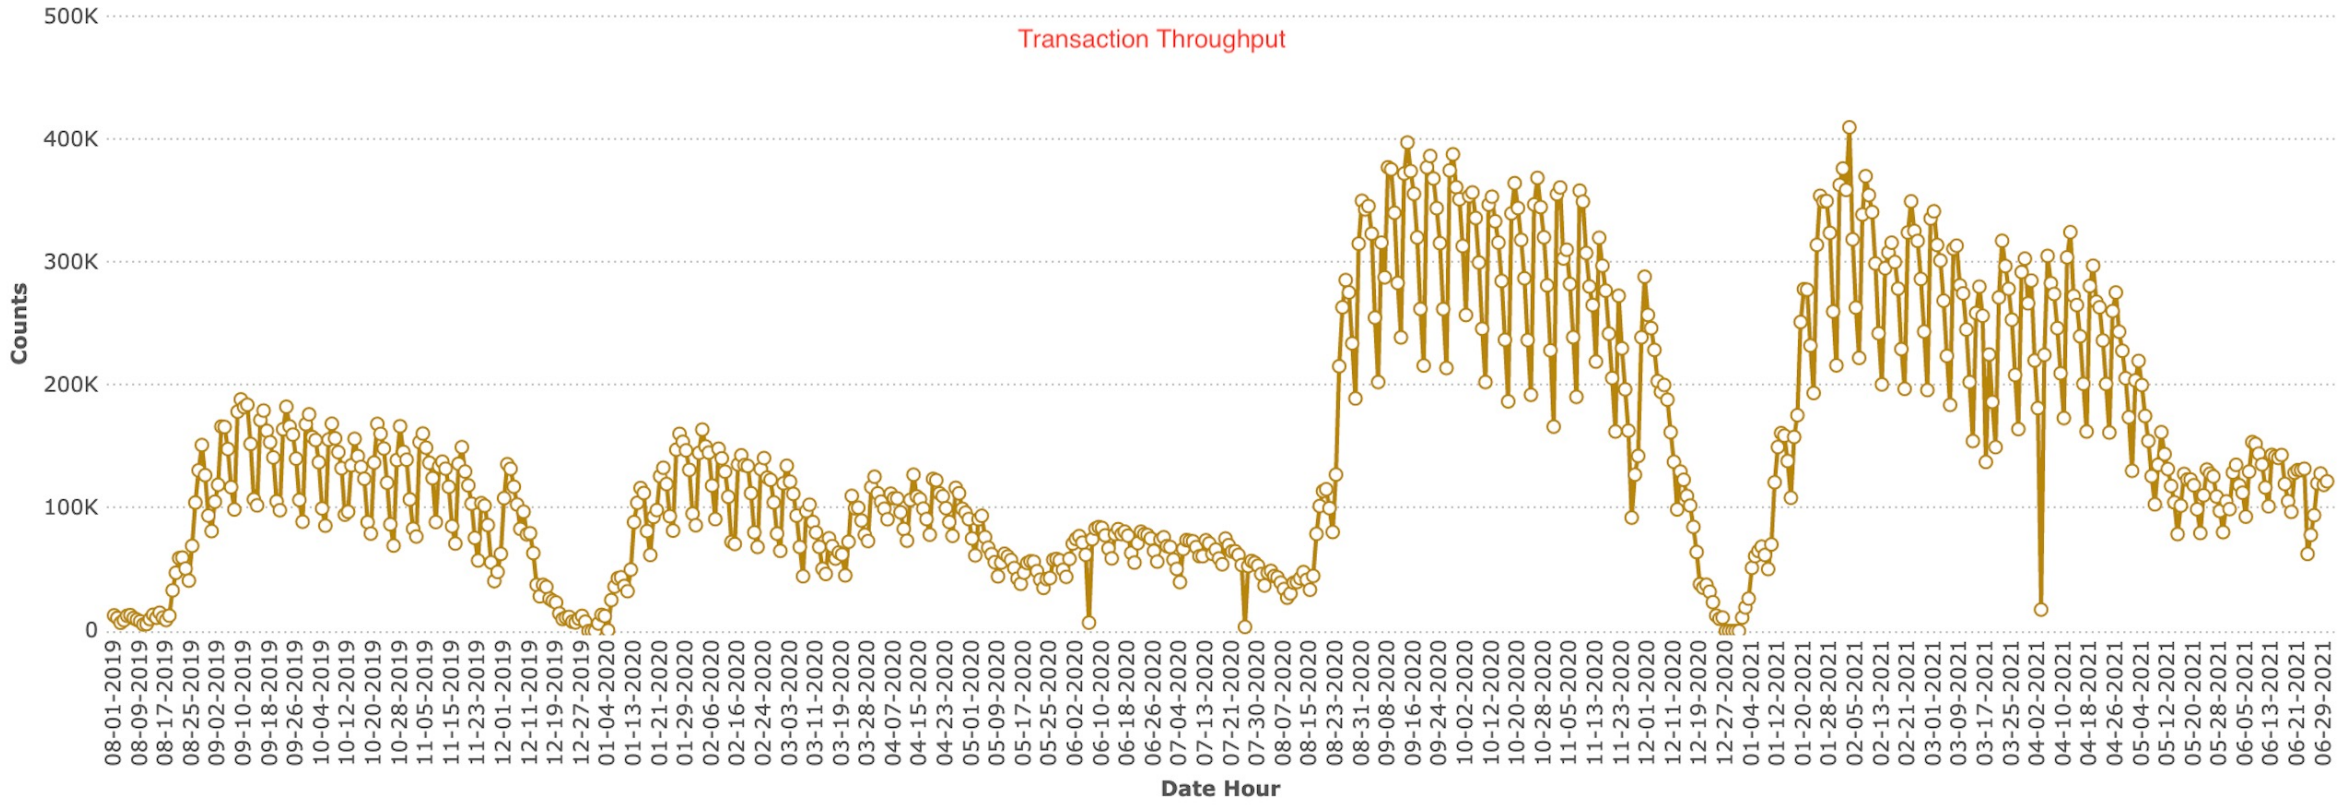

Covid-19 Impact

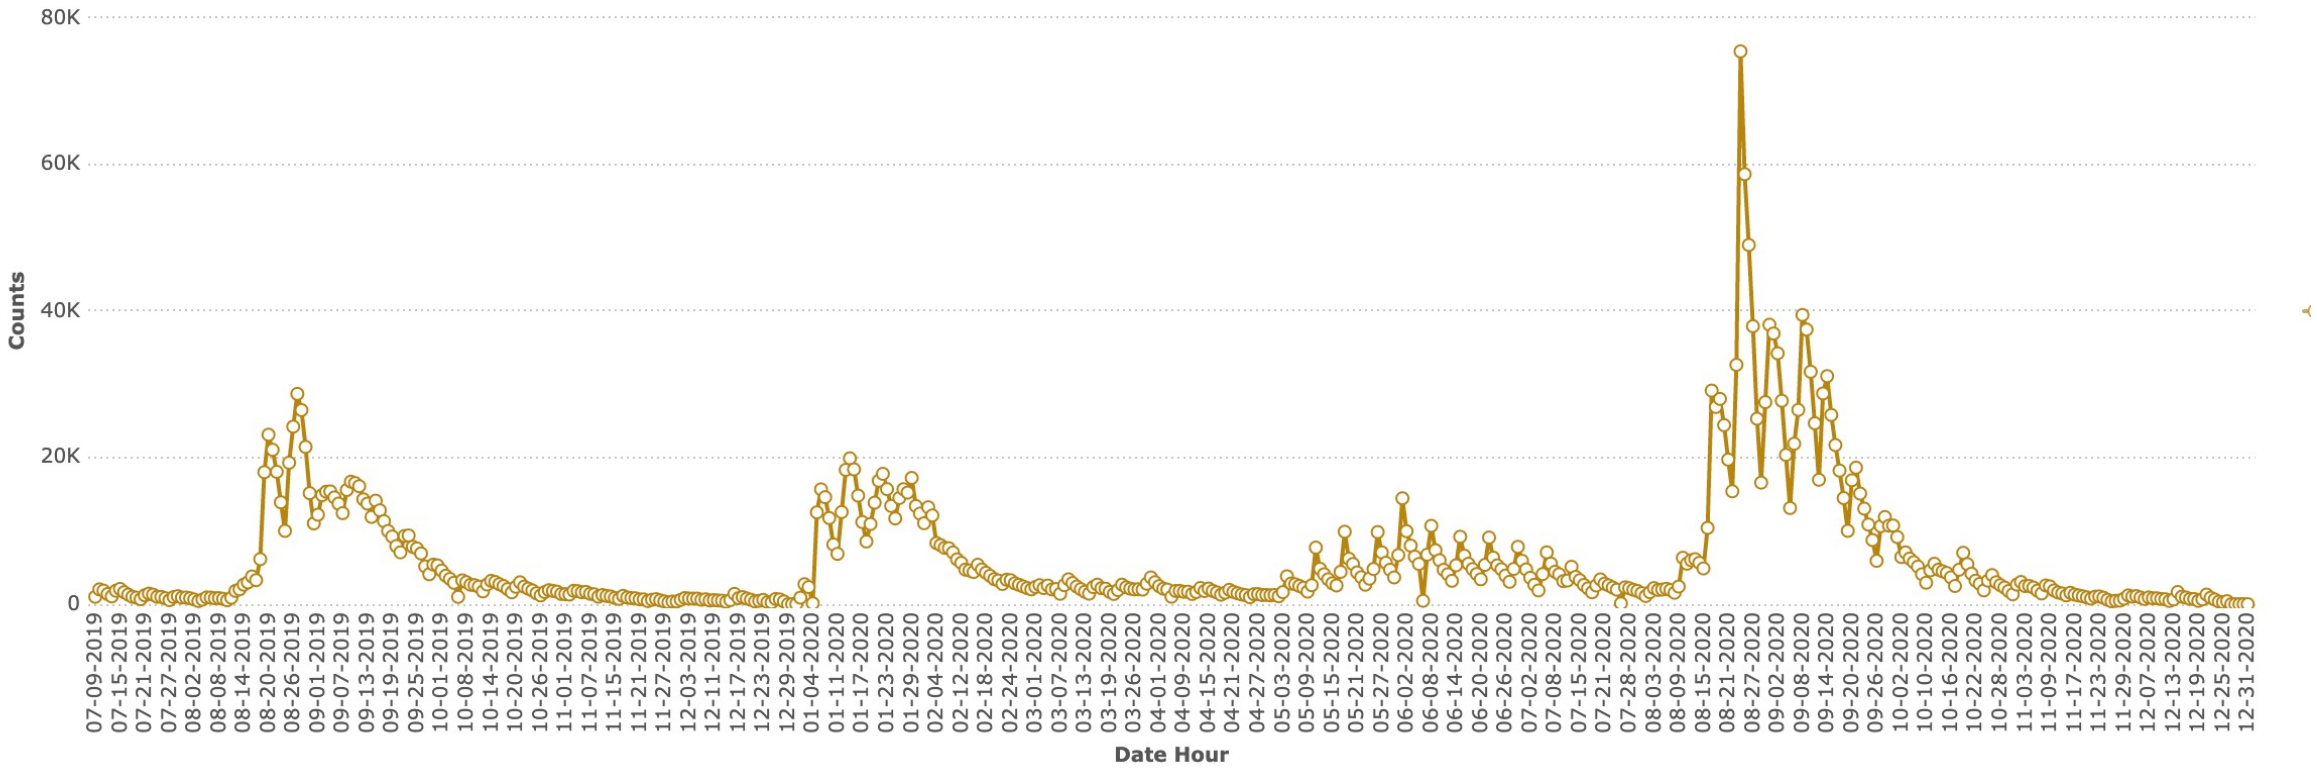

Think Upstream

- Respond vs Prevent
- Scaling
- Preparing for future

Evolution to Multi-Account and Multi Cloud Architecture

Innovator
Small team working on shared goal.

Share House
Multiple teams sharing an account for different projects.

Hosted Services
Handful of centrally managed accounts (dev, prod, etc) are shared by multiple teams.

Multi Account
Projects operate with independence and isolation within agreed rules and services.

Scaling (Non-Technical)

People

- Don’t leave anyone behind!
- Bring cloud learnings into to the whole company
- Form CoEs / Advisory Councils

Processes

- Smart processes vs no processes
- Be agile and not a bottleneck
- Govern via preventative controls, guardrails, and auto-remediation / detective controls

Culture

- “How we get things done”
- Continual Learning
- Empower teams: create ownership, accountability, and collaborative decision making
- Evolve towards homogeneous agile practices and mindset

Scaling (Technical)

- **Governance First**
  - Rules and Regulations
  - Shared Services
  - Preventative and detective controls
- ***Everything* As Code**
  - Infrastructure, cybersecurity, acceptance criteria
  - Accountability, approval, and traceability through Git
  - Re-usable building blocks, supported, and shared
- **Security and Networking**
  - Think forward and plan for growth
  - Create a Cloud Security Policy

Dev+Sec+Ops

- Practice like it's real!
- Bundle security tooling into your DevOps toolchain.
  - Ensure development accounts run under same security rules as production.
  - Code linting and static analysis for everything.
  - Continuous integration isn't just for software devs.
- Develop process to manage, approve and publish AMIs.
  - Automate discovery and enforcement of unapproved AMIs.
  - Be aware of CVEs.
  - Rotate often (monthly?).
  - Fall back by automated patching in place
  - Build "cattle", not "pets".

Centralized Machine Image Publishing

Automate approved & standard AMI Process:

- Build/Test/Patch
- Automate Publishing
- Guardrails to check for non-standard use.
- Exception approval process and testing for marketplace and non-standard use cases.

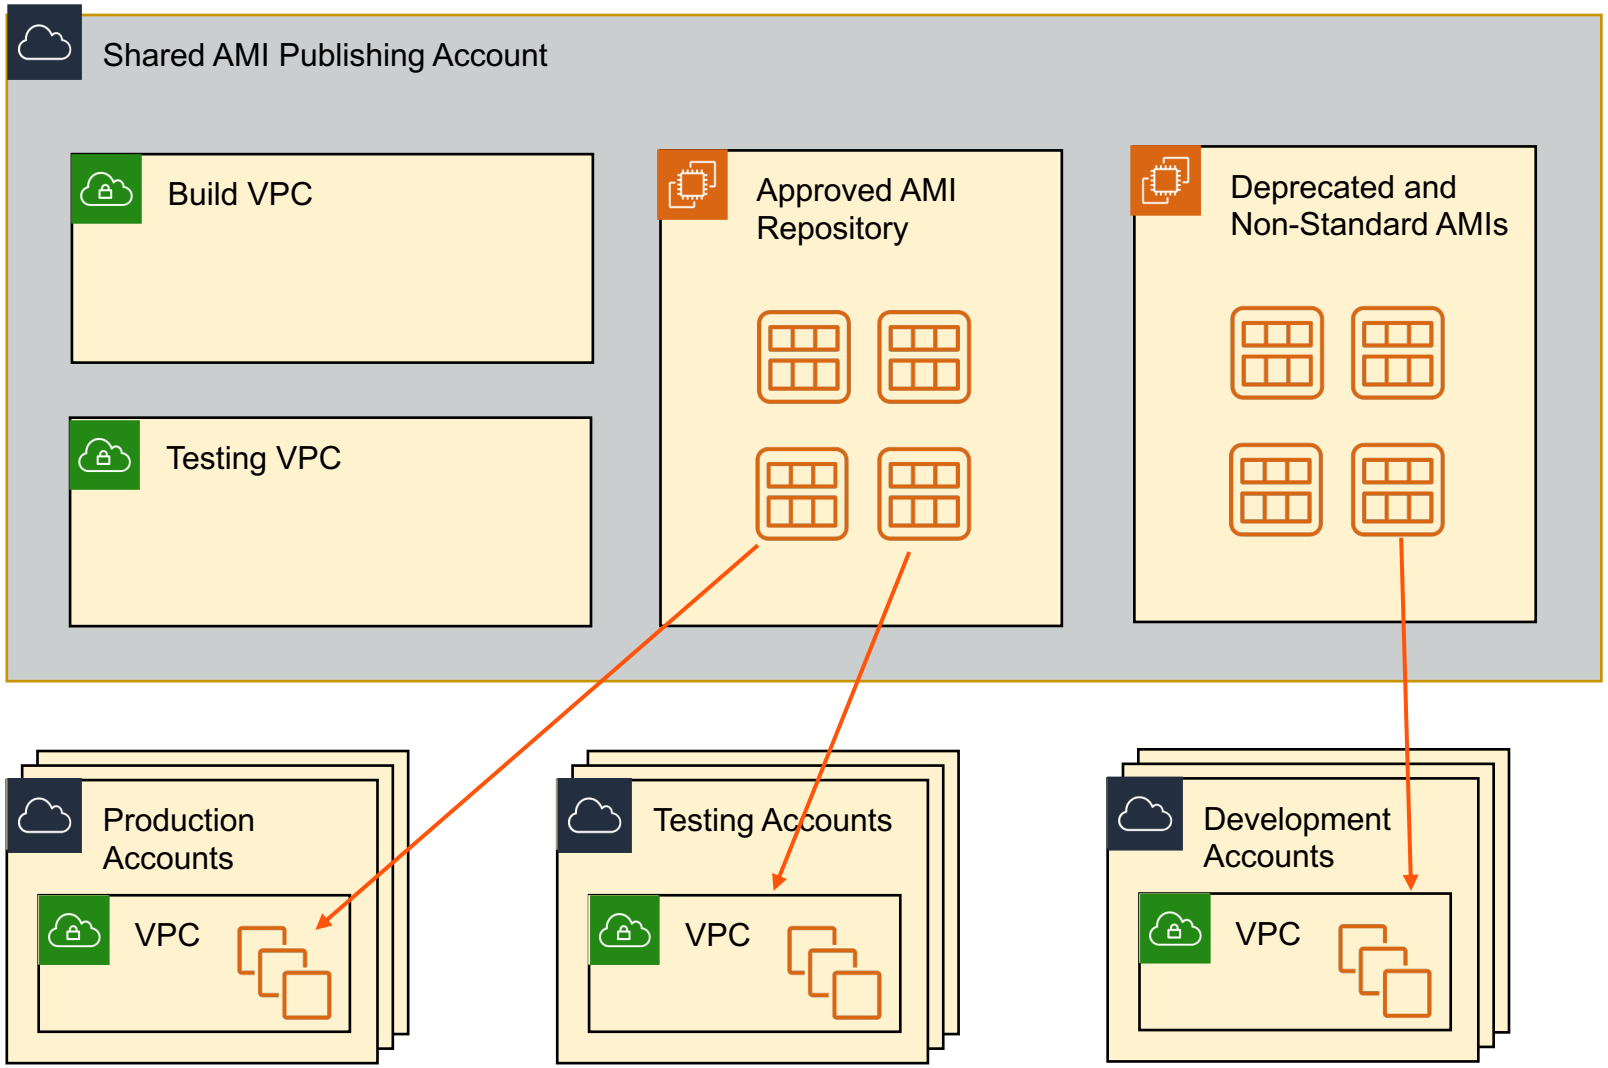

Network Scaling:

Transit Gateway
Cloud Fabric
Cloud Adjacency
DNS Internalization

Unified Network

AWS Transit Gateway enables us to:

- Connect our 150+ accounts via a single gateway.
- Eliminates the need to manage peering between all VPCs.
- Enables central Direct Connect to our Cloud Fabric.
- Dramatically simplifies routing and network management.
- Improves our performance substantially.
- Improves visibility and control.
- Reduces our operational costs.
- Reduces Toil

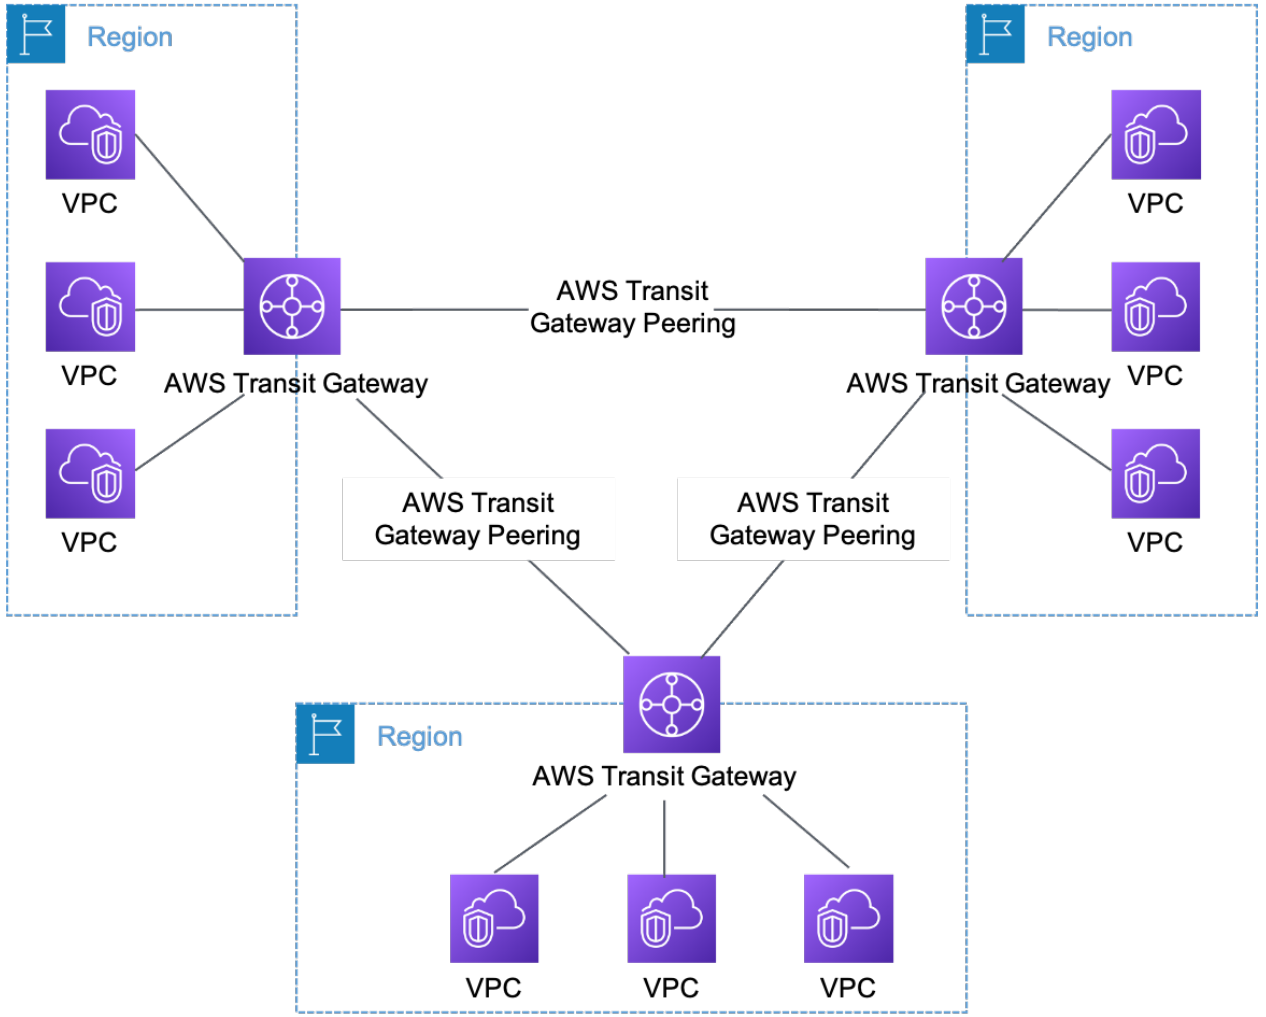

McGraw Hill's Cloud Fabric

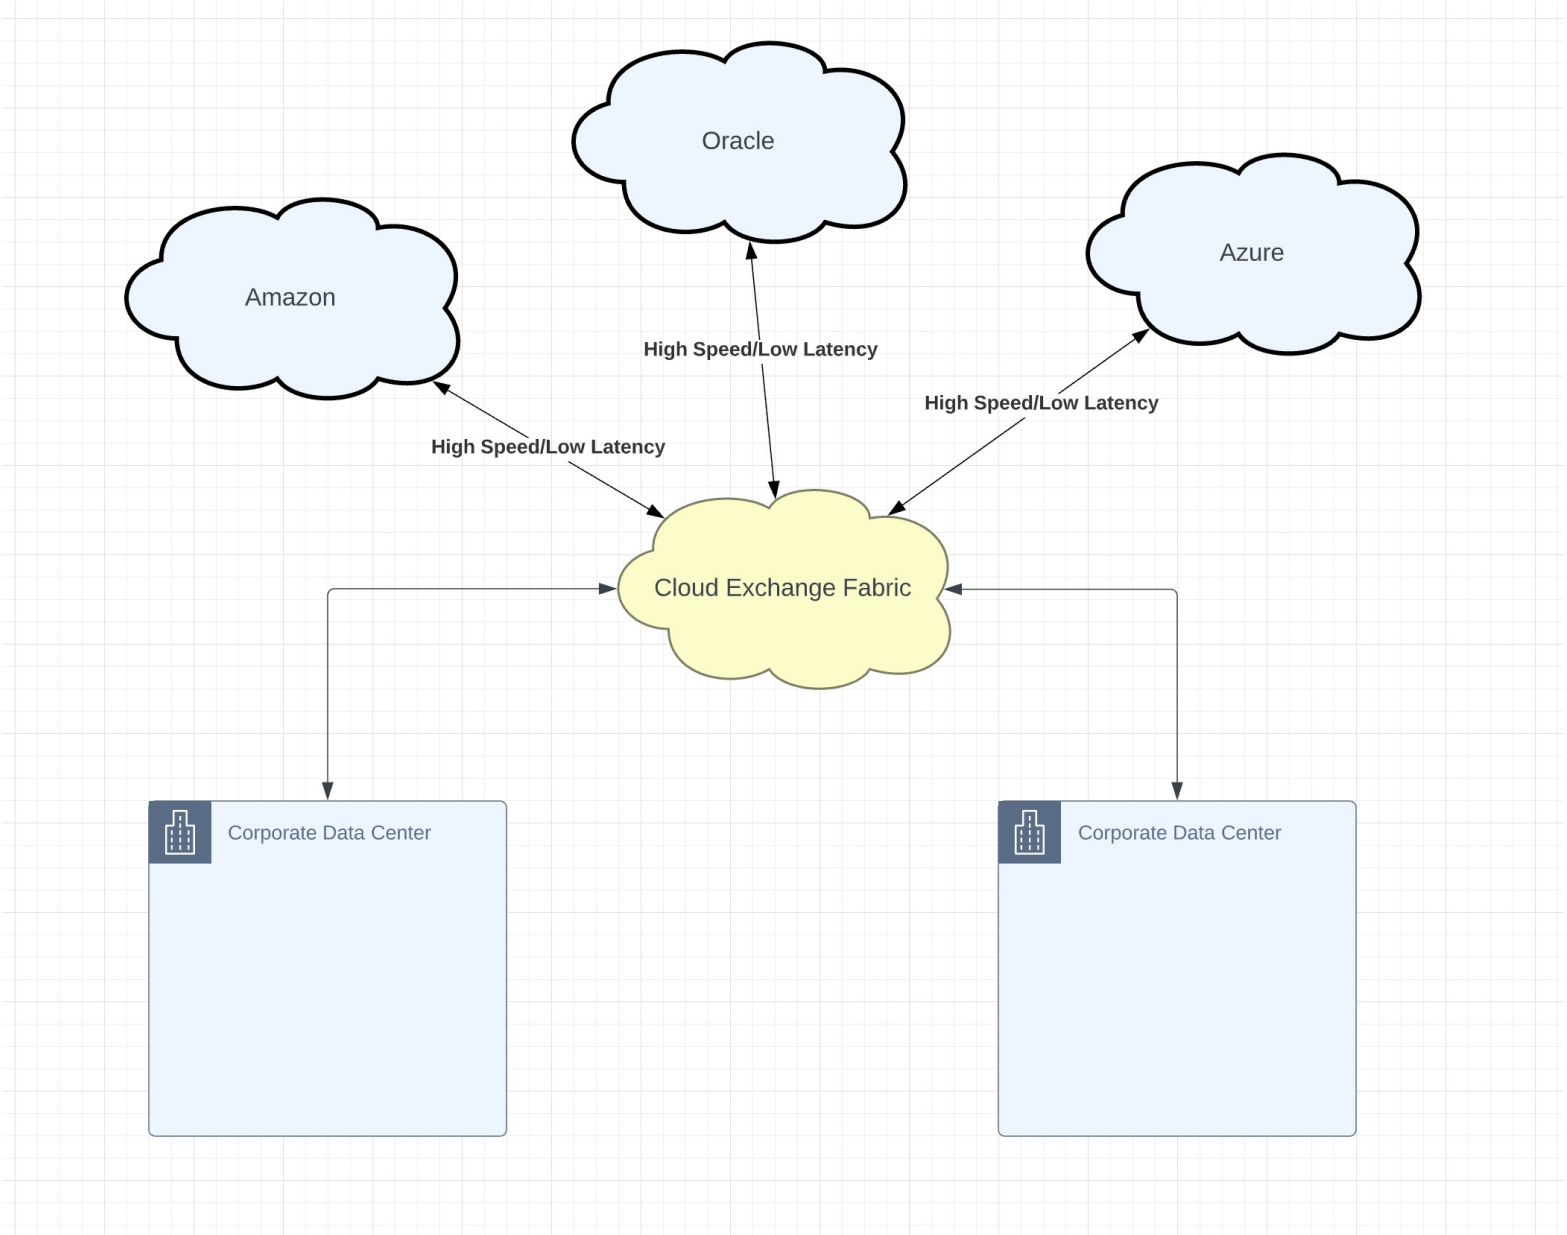

Cloud Adjacency

Leveraging the best features of distinct cloud providers who are **directly** interconnected via a **secure**, **high bandwidth**, and **low latency** network

Business

- Leverage vendor relationships
- Avoid “cloud lock-in”
- Reduce multi-cloud overhead
- Compensating controls

Developers / Engineers

- Additional flexibility and more choices
- Cross-cloud team cooperation
- Use best features from best providers without compromise

Better solutions for end users!

DNS Internalization

Provide McGraw Hill with a **highly-resilient**, **performant**, and **multi-region-capable** internal DNS infrastructure that addresses the following goals:

- Provide a unified internal DNS fabric across all AWS VPCs.
- Ensure all internal traffic remains restricted to the internal network.
- Minimize or eliminate public exposure of internal services.
- Allow for automated management of internal DNS zones and resource records.
- Minimize the number of network hops required to resolve internal records.
- Allow all VPCs to resolve any other internal domain or resource in other VPCs.
- Allow on-premise workloads to resolve internal domains in the AWS environment and allow VPCs to resolve on-premise internal domains and resources.

Mc Graw Hill

DNS Internalization (Solution)

Leverages Route 53 Resolver, Route 53 Private Hosted Zones, and forwarding rules to allow resolution of domains and resources.

Works across multiple accounts, between workloads running on AWS, in adjacent clouds, and on-premise environments.

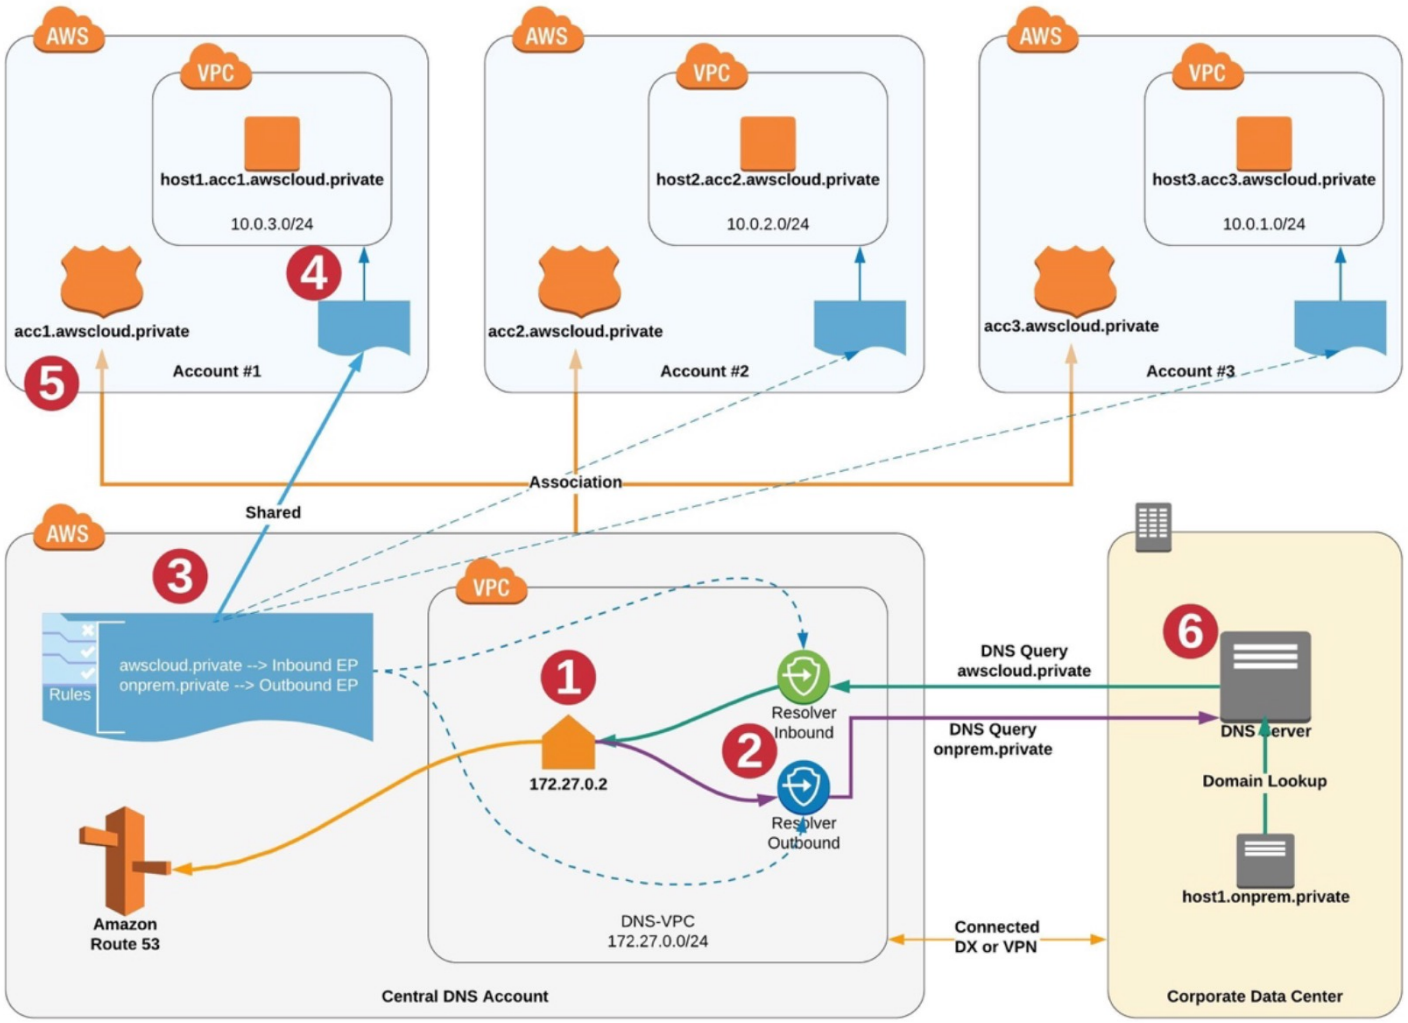

Guardrails and Security Controls

- Disallow public access to your databases
- Detect Disk, Buckets, database storage encryption
- Disallow public access to S3 buckets
- Disallow deletion of critical network and security constructs
- Disallow unapproved regions
- Disallow actions as a root user
- Disallow changes to automation IAM roles
- Disallow port 22 from Internet
- Disallow unapproved machine images

Cost Management

(In-progress) Leveraging Cloud Nuke (and similar tools) in:

- **Account Factory:** Ensure new accounts are resilient and have guardrails in-place.
- **Sandbox Accounts:** Ensure (a) automated clean-up, (b) practice automation-by-default.

Best practices and cost improvements rolled into:

- Foundational building blocks (Terraform modules):
  - Amazon EBS gp3 volumes
  - Auto-scaling with ECS Capacity Providers + right-sized instances
- Custom automation to catch misconfigured AWS accounts or leftover cloud resources.
- Cost Anomaly alerting

Demo – Monitoring as Code

- Monitoring as Code

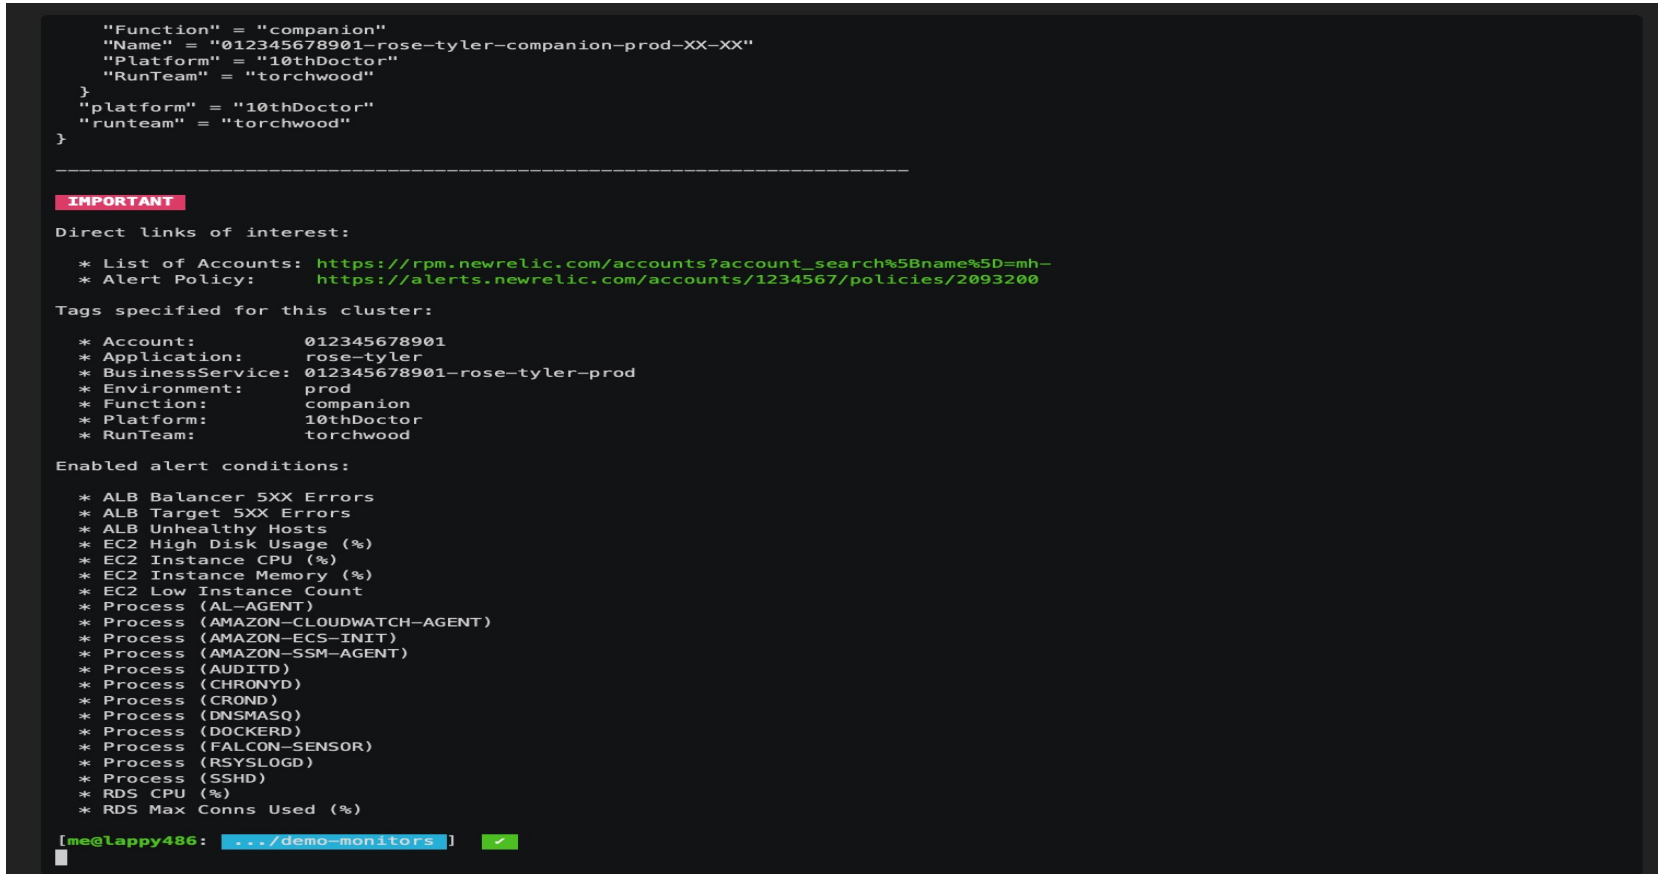

Horizon


Serverless ◦ Network Segmentation ◦ Zero Trust Architecture ◦ Deprecation of SSH ◦ Intent-Based Networking ◦ AI Ops ◦ Advanced Automation

Because learning changes everything.®

Thank You!

Chinmay Tripathi

- **Sr. Director, Engineering**

chinmay.tripathi@mheducation.com

linkedin.com/in/chinmaytripathi

https://careers.mheducation.com/

"Want to empower educators and learners across the world?"
We are hiring!"