Net Actuate

# Exit Hyperscalers. Keep the Cloud.

## A Framework for Evaluating Global Infrastructure Options

Organizations wanting to repatriate cloud workloads mainly cite financial considerations as a reason for exiting the cloud, but also dependencies on a third-party's proprietary services, price hikes and policy changes, lock-in to the provider's ecosystem, cloud waste, and cost management. These considerations almost exclusively pertain to large public cloud providers like AWS, Azure, GCP—the hyperscalers.

This is especially the case for post 2015 start-ups, that often overrelied on hyperscalers as they lacked initial CAPEX and were taking advantage of free credits, but now need to better manage their OPEX. Some notable examples of cloud repatriation projects are **Ahrefs** and **Basecamp/Hey**, which claim to have saved millions by moving their cloud workloads to on-premises servers. **AWS themselves** quoted cloud exits as a challenge.

In the examples above and many others, the only solution cited to solve cloud-related challenges is running your own hardware. While this is one possible solution, there are reasons why organizations choose the cloud, and those reasons are unlikely to have changed.

The question is whether organizations can have their cake and eat it too—to have their infrastructure outsourced but without hard-to-control costs and overdependence on a hyperscaler's proprietary services.

# Hyperscaler Alternatives

Hyperscaler alternatives that aim at addressing cost-related challenges typically include colocation vendors such as Digital Realty and Summit, and inexpensive cloud providers, such as Digital Ocean and Vultr.

While valid for some use cases, both of these options have drawbacks compared to hyperscaler clouds.

- **COLOCATION** requires a high level of customer responsibility and high upfront CAPEX investment. Colocation vendors only provide space, power, and cooling; so organizations need dedicated infrastructure staff to procure, deploy, and manage the hardware in the colocation data center.
- **INEXPENSIVE CLOUD PROVIDERS** offer a comparable developer experience to the hyperscaler, but have a less mature and comprehensive set of native services. A smaller partner ecosystem and pool of trained and certified engineers entails a degree of investment.

These two options either have organizations opt for a watered down version of a hyperscaler in an attempt to control costs, or take on a huge responsibility for managing both hardware and software.

## Exit Hyperscalers, Enter Open Network Edge

NetActuate, a global provider of managed network and infrastructure services, offers a third path: a hyperscaler alternative that reduces costs without increasing the burden of infrastructure management or requiring customers to learn a new cloud provider's terminology.

NetActuate's Open Network Edge (ONE) concept provides the ideal blend of deployment flexibility, global reach, low latency, and cost-effectiveness. A globally distributed infrastructure-as-a-service solution built using open source technologies, ONE provides out-of-the-box implementation of widely deployed open source software, including operating system, network operating systems, monitoring software, orchestration tools, and everything else needed to scale your platform.

# What does IaaS support that on-prem doesn't?

Organizations opt for cloud-based infrastructure in the first place to have their IT infrastructure respond more easily to business changes and growth. Typically, in instances where your needs include:

- Latency sensitivity
- Global expansion and delivery
- Multi-user scalability and concurrency

As such, if you require any of the above, you will most likely have to opt for an IaaS solution. To help you decide which type of infrastructure model is best suited for your needs, we will review and score hyperscalers, the Open Network Edge, and colocation on the criteria above.

# Hyperscaler Exit Scorecard

We shall evaluate each criterion — capability, implementation, and cost — on a scale from 1 to 5 against Capability and Implementation, while Cost will be a binary evaluation.

- **CAPABILITY:** a score of 5 indicates an extensive native set of features to support the criterion, whereas a score of 1 indicates a lack of native features to support the criterion.
- **IMPLEMENTATION:** a score of 3 indicates no additional work required to implement the feature, while a score of 1 indicates considerable additional work, and a score of 5 indicates considerably less work to deliver (such as a managed service).
- **COST:** This will be evaluated as follows:
    - **NO ADDITIONAL COST** (or not applicable), meaning there are no further charges by the provider to achieve these capabilities. Overhead costs are not considered.
    - **ADDITIONAL COST**, such as purchasing additional services. This will only consider OPEX, such as usage-based pricing for a monitoring service, rather than CAPEX which would entail customers purchasing their own hardware.

## Latency Sensitivity

This refers to use cases where individual users have little tolerance for latency, such as real-time applications, which can include video processing at the edge, autonomous driving, virtual and augmented reality, remote control of IoT devices, trading and retail, and the like.

It includes:

- **NETWORK-INDUCED LATENCY —** takes into consideration the roundtrip latency from a user making a request to the place where it is processed.
- **COMPUTE-INDUCED LATENCY —** takes into consideration the delay associated with processing a request at an infrastructure level, such as cold starts. The actual computation time is not in scope.
- **STORAGE-INDUCED LATENCY —** refers to the processing time associated with retrieving data from or writing data to a storage system, such as using KV stores, in-memory databases, and others.

## NETWORK-INDUCED LATENCY SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●○○<br>Hyperscalers operate their data centers at a global scale, but they are typically used to serve requests from a single location. Edge services in the hyperscalers' portfolio can be used to distribute requests across multiple points of presence. | ●●●●○<br>With an Anycast network and over 40 locations, NetActuate can reach major metropolitan areas in under 30 milliseconds. | ●●○○○<br>Colocation customers get access to global connectivity at a single connection point, while avoiding the need for expensive WAN links, but these setup exercises must be repeated at all locations required by customers. |
| **IMPLEMENTATION** | ●●●○○<br>The customer needs to handle networking functions within the cloud provider, such as configuring VPCs, transit gateways, load balancing, and content delivery. | ●●●●●<br>Fully consumed as-a-service, customers can easily onboard and fully access the global network. | ●○○○○<br>Customers need to procure and deploy their own hardware across their preferred datacenters, which entails significant additional effort compared to virtualized or as-a-service solutions. |
| **COST** | Additional cost | No additional cost | Additional cost |

## COMPUTE-INDUCED LATENCY SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●○○<br>Hyperscalers offer a range of compute services, from persistent VMs to ephemeral containers and serverless functions. Low compute-induced latency can be achieved via persistent virtual machines. Serverless functions can be subject to rather long cold starts. | ●●●●○<br>NetActuate' bare metal and managed hosting services, along-side persistent VMs, are not subject to usage-based pricing. These always-on instances can process requests as soon as received. | ●●●●○<br>Colocation deployments are typically associated with bare metal servers that are continuously operational, meaning that the requests can be served as soon as received. |
| **IMPLEMENTATION** | ●●●●●<br>Compute functions such as VMs and serverless are natural parts of cloud architectures and entirely software-defined. | ●●●●●<br>Provisioning compute instances on the Open Network Edge can be done programmatically using APIs, done via GUI, or even as a managed service. | ●●●○○<br>While there is no additional effort for provisioning compute-sensitive instances in a collocation environment, the customer is responsible for the procurement, deployment, and management of hardware and software. |
| **COST** | No additional cost | No additional cost | No additional cost |

## STORAGE-INDUCED LATENCY SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● | ●●●●○ | ●●●○○ |
| | Hyperscalers offer a comprehensive suite of storage options, including highly optimized in-memory databases, caching services, and low-latency storage offerings. Services like Amazon ElastiCache, Azure Redis Cache, and Google Cloud Memorystore provide sophisticated caching capabilities that can dramatically reduce storage-induced latency. Additionally, hyperscalers have invested heavily in high-performance storage subsystems that offer consistent performance at scale. | NetActuate offers global object storage that is S3 compatible, as well as various custom SAN/NAS/DAS options custom built to customer specifications and exposed via programmatic APIs.<br><br>ONE leverages global NVMe-backed object storage and advanced caching layers reducing latency for data-intensive applications. | In colocation environments, organizations can deploy specialized storage hardware optimized for their specific latency requirements, including all-flash arrays and NVME storage. |
| **IMPLEMENTATION** | ●●●●○ | ●●●●○ | ●○○○○ |
| | WIth a range of storage products and managed services, customers can easily provision and maintain their storage infrastructure | Fully managed provisioning, configuration, and ongoing management are offered, minimizing the customer's operational burden. Storage can be programmatically adjusted to optimize latency requirements dynamically via APIs. | Achieving optimal storage performance requires significant expertise and investment. Additionally, colocation facilities generally don't offer the same level of managed storage services as IaaS proiders, placing more operational burden on the organization to maintain low-latency storage systems. |
| **COST** | Additional cost | Flexibility — Usage based or Predictable, flat-rate pricing model without hidden operational costs ensures controlled storage expenditures, unlike hyperscalers' unpredictable usage-based pricing. | No additional cost |

# Global Expansion and Delivery Scorecard

These are instances where an organization is rapidly entering new geographies and requires a global infrastructure provider to support their expansion.

- **GEOGRAPHICAL PRESENCE —** this is the number of data centers or points of presence a provider operates in different geographies. These need to be close to major metropolitan areas across continents.
- **LEGAL, COMPLIANCE AND REGULATIONS —** refers to the legal considerations a customer must work through to operate in that geography. This can include setting up local entities such as limited companies, working with local providers, and even data sovereignty requirements.
- **PROGRAMMABILITY —** is about interacting with the underlying IaaS solution for provisioning and configuring the infrastructure at scale and dynamically without clicking through menus in a graphical user interface.
- **CACHING AND STORAGE —** refers to the ability to distribute and store content closer to end users, reducing retrieval times and origin server load. Effective caching strategies dramatically improve performance while reducing bandwidth costs and origin server load.
- **TRAFFIC ENGINEERING —** this encompasses the technologies and methodologies used to optimize the flow of network traffic across distributed infrastructure. It can include intelligent request routing, congestion management, failover mechanisms, load balancing, using Anycast routing, origin shields, quality of service, etc.

## GEOGRAPHICAL PRESENCE SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●○<br>Hyperscalers operate data centers at a global scale, with multiple locations in North America and Europe, but limited coverage in emerging markets. | ●●●●○<br>NetActuate's >40 global PoPs are distributed across geographies, positioned in major metropolitan areas. | ●●●●○<br>Colocation providers operate data centers in global locations. |
| **IMPLEMENTATION** | ●●●●○<br>Different availability zones can easily be chosen in a cloud environment. Every availability zone needs to have landing zones configured, meaning there is additional effort associated with every new zone. | ●●●●●<br>Customers have access to NetActuate's whole infrastructure footprint without requiring additional per-zone configuration. | ●○○○○<br>Accessing a new colocation datacenter requires a whole new procurement and provisioning process, both for the hardware running on-location, as well as signing a services agreement with the collocation provider for the new location. |
| **COST** | Additional cost<br>Every availability zone needs its supporting infrastructure. Pricing usually varies across availability zones. | No additional cost | Additional cost |

## LEGAL, COMPLIANCE AND REGULATIONS SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● Hyperscalers are compliant with local regulations such that customers have little to no friction from an infrastructure perspective to operate in the local markets. Some, such as AWS even have specific certifications, like the C5 attestation in Germany. | ●●●●● NetActuate maintains extensive global compliance certifications (SOC 2 Type II, ISO 27001, GDPR-ready infrastructure), significantly simplifying entry into regulated markets. ONE includes dedicated compliance consultancy to assist customers in meeting regional regulations seamlessly. | ●●●○○ Colocation providers generally maintain good compliance practices for physical security and environmental controls but offer less assistance with software-level compliance compared to hyperscalers. |
| **IMPLEMENTATION** | ●●●●● Organizations deploying their workloads in hyperscalers have to do little to no additional work to ensure their infrastructure is compliant with regional regulations. | ●●●●● NetActuate provides parity with hyperscalers for facility certifications, so little to no additional work is required for customers. NetActuate operates using a shared responsibility model. | ●●●○○ They typically hold regional certifications for their facilities but place more responsibility on customers to ensure their deployed solutions meet local regulations. |
| **COST** | No additional cost | No additional cost | No additional cost |

## PROGRAMMABILITY SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● Hyperscalers offer exceptional programmability through comprehensive APIs, SDKs for multiple programming languages, and sophisticated Infrastructure as Code (IaC) tools. Their native services like AWS CloudFormation, Azure Resource Manager, and Google Cloud Deployment Manager enable highly automated infrastructure provisioning and management. | ●●●●○ The Open Network Edge is a fully programmable service that can be deeply integrated in CI/CD pipelines. ONE can be programmatically configured using open source tooling such as Terraform and Ansible. | ●●○○○ Some colocation providers have developed APIs and portal-based automation in recent years, and Software-Defined Interconnection that can be programmatically managed, but physical infrastructure changes generally involve longer lead times and more manual intervention. |
| **IMPLEMENTATION** | ●●●●● Cloud providers offer both proprietary and open source tooling for managing the infrastructure as code and integrating in CI/CD pipelines. | ●●●●● This is a core tenet of the Open Network Edge, providing full programmatic access to the solution. Developers can use their preferred tools to interact with the solution. It exposes features via GUI, API, CLI, and SDK. | ●●●○○ Customers have the opportunity of implementing their own tools to manage the hardware, but no services are offered by the collocation provider. |
| **COST** | Additional cost | No additional cost | No additional cost |

## CACHING AND STORAGE SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● Hyperscalers offer a CDN service on top of their infrastructure that provides caching for global delivery.<br><br>Hyperscalers have developed advanced features like origin failover, automatic compression, and media optimization. Their caching systems typically include comprehensive monitoring and analytics that help organizations optimize their content delivery strategy. | ●●●●● Built-in advanced caching with globally distributed Varnish CDN, optimized for high-IOPS workloads. Integrated CDN, combined with Anycast IP delivery, significantly improves global performance and reduces bandwidth costs. | ●○○○○ Colocation providers do not offer any caching or storage services.<br><br>Colocation facilities provide the physical infrastructure where organizations can deploy their own caching and storage solutions, but generally don't offer managed caching services. |
| **IMPLEMENTATION** | ●●●●○ These CDN services are available as add-ons to existing hyperscaler footprints and tightly integrate with the rest of the environment, but entail additional setup. | ●●●●● Turnkey CDN and caching solution requires minimal configuration from users, supporting instant-on deployments via GUI, APIs, or Terraform. | ●○○○○ This approach requires organizations to design, implement, and maintain their own caching architecture, often involving third-party CDN services or self-managed caching software like Varnish or NGINX. |
| **COST** | Additional cost | No additional cost | No additional cost |

## TRAFFIC ENGINEERING SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● Hyperscalers offer comprehensive load balancing solutions at multiple levels, including global load balancing, regional load balancing, and application-level load balancing.<br><br>Their traffic management capabilities typically include DDoS protection, rate limiting, and traffic prioritization features. Advanced capabilities like AWS Global Accelerator provide optimized routing over the hyperscaler's private backbone network. | ●●●●● ONE includes comprehensive Anycast routing, ECMP (Equal-Cost Multi-Path) load balancing, DDoS protection, and advanced traffic prioritization. Real-time BGP optimization through APIs or a user-friendly dashboard enhances network performance and reliability. | ●●●●○ The multi-carrier nature of most colocation facilities can help organizations define network designs with route optimization, multi-homing, and carrier redundancy. |
| **IMPLEMENTATION** | ●●●●○ The traffic management capabilities are powerful but may require integration of multiple services to achieve a comprehensive solution. | ●●●●● Zero-configuration network optimization is included. Advanced capabilities can be tailored via APIs and managed network services without additional complexity. | ●●○○○ Implementing these capabilities requires significant networking expertise and ongoing operational management. Organizations must design, deploy, and maintain their own traffic engineering infrastructure, which can be complex and resource-intensive. |
| **COST** | Additional cost | No additional cost | No additional cost |

# Multi-user Scalability and Concurrency Scorecard

This refers to instances where the volume of request and time spent processing each of them will have an impact on the overall experience and the scalability of the solution.

- **FAILOVER AND RESILIENCE —** this ensures continuous service availability even during component failures, traffic spikes, or regional outages. It includes redundancy mechanisms, high availability configurations, active-active configuration, Anycast-based networks, automated failover processes, and disaster recovery capabilities across distributed infrastructure. Facilities and physical Infrastructure reliability capabilities such as redundant power and disaster recovery is out-of-scope.
- **POP RESOURCES —** the solution's ability for dealing with noisy neighbors, providing dedicated hardware, and scaling CPU/Memory/Storage/Port speeds per single location.
- **ANALYTICS AND TELEMETRY —** refers to the capabilities for monitoring, measuring, and analyzing system performance and user interactions across distributed infrastructure. This includes real-time metrics collection, distributed tracing, anomaly detection, capacity planning tools, and actionable insights that enable organizations to optimize their applications for multi-user scalability.

## FAILOVER AND RESILIENCE SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● Hyperscalers offer services such as global load balancers, managed database replication, cross-region data synchronization, and automated failover mechanisms. Services like AWS Route 53, Azure Traffic Manager, and Google Cloud Load Balancing provide sophisticated health-checking and routing capabilities that auto-matically direct traffic away from impaired resources. | ●●●●● Robust global Anycast network offers built-in HA, active-active configurations, and automated failover capabilities across >40 global PoPs, enhancing uptime and performance. | ●●●○○ Some colocation providers typically design their facilities to Tier III or Tier IV data center standards, with multiple independent distribution paths and concurrent maintainability. Others offer direct connections between their facilities, simplifying the implementation of active-active architectures. |
| **IMPLEMENTATION** | ●●●●○ There is some configuration work required to set up the high availability services, but these follow documented architectural patterns and guided setups. | ●●●●● Out-of-the-box resilience provided without manual intervention. Advanced configurations and automated disaster recovery scenarios are facilitated via easy-to-use APIs or managed options based on BGP Anycast.Custom BCP/ DR strategies also exist that can be managed and implemented by NetActuate without customer intervention. | ●●○○○ Customers are responsible for designing, implementing, and managing their own resilience solutions at the network, server, and application levels. Synchronizing and managing distributed systems across colocation facilities requires more operational effort compared to hyperscalers' integrated services. |
| **COST** | Additional cost | Additional cost | No additional cost |

## PoP RESOURCES SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● | ●●●●● | ●●●●○ |
| | Hyperscalers operate full data centers, which can provide as much scalability as an organization would need at a single location. While Edge locations may have fewer resources compared to full-sized data centers, these can also scale enough for most use cases.<br><br>Hyperscalers implement resource isolation mechanisms to prevent "noisy neighbor" problems in their multi-tenant environments. Technologies like hypervisors with dedicated CPU cores, memory allocation controls, and storage QoS mechanisms help maintain consistent performance. Some hyperscalers offer dedicated instance options that provide exclusive access to underlying hardware for applications with particularly stringent performance requirements. | NetActuate's edge PoPs support dedicated resource allocation, offering customized CPU, memory, storage, and port-speed scalability, reducing noisy neighbor issues through dedicated hardware or isolated VM allocations.<br><br>Going beyond Hypervisors, NetActuate also allows customers to provision bespoke hardware, dedicated to them as well as colocate their own hardware and mix/match in a combination of environments. | By owning and controlling their hardware, organizations can eliminate multi-tenant resource contention entirely, ensuring consistent performance even under heavy load. This approach is particularly valuable for applications with specific resource requirements or performance characteristics that don't align well with standardized cloud offerings. |
| **IMPLEMENTATION** | ●●●●● | ●●●●○ | ●○○○○ |
| | All appliances software-defined and provisioning is done programmatically. | Simple scaling through customer portal/API; managed service offerings available may require service requests to spin up additional servers. | Managing distributed hardware across multiple colocation facilities introduces significant operational complexity, particularly for organizations without established processes for global hardware deployment and management. |
| **COST** | Additional cost | Additional cost | No additional cost |

## ANALYTICS AND TELEMETRY SCORECARD

| | Hyperscaler | Open Network Edge | Colocation |
|---|---|---|---|
| **CAPABILITY** | ●●●●● Hyperscalers offer comprehensive analytics and telemetry capabilities that span their entire service portfolio. These include real-time metrics collection, log aggregation, distributed tracing, and sophisticated alerting capabilities. Services like AWS CloudWatch, Azure Monitor, and Google Cloud Monitoring provide detailed insights into application performance, resource utilization, and user experience across global deployments. | ●●●●● Advanced telemetry based on server resources and network flow data offers deep visibility into infrastructure performance, real-time metrics, logging, tracing, and actionable analytics exposed via the NetActuate portal. | ●●○○○ Traditional colocation providers offer limited native monitoring capabilities, typically focused on facility metrics like power utilization, and associated costs. |
| **IMPLEMENTATION** | ●●●●○ There is additional work required to deploy monitoring services, but these are tightly integrated in the wider hyperscaler portfolio, so deployment and data ingestion setups are relatively straightforward. | ●●●●● Customers can use default Dashboards under metrics and observability, as well as configure additional alerts or services. | ●○○○○ Organizations must implement their own monitoring solutions to gain visibility into application performance and user experience, which entails a whole tech stack deployment, from ingesting hardware and kernel-level logs to deploying their data pipelines solution and analytics platforms. |
| **COST** | Additional cost | No additional cost | No additional cost |

# User Scorecard

Now it's your turn. We've developed this scorecard for you to rank how important each one of these areas is to your business. Rank the needs of your company to get a sense for which areas you'll need to focus on when choosing your optimal global infrastructure. If you'd like to schedule a pressure-free 30-minute call with one of our engineers to learn how NetActuate can help you build and deploy a solution, click the "Schedule a Call" button.

Schedule a Call

| Name | Email | How important is this to you? | Which is your priority? |
|---|---|---|---|
| **LATENCY-SENSITIVE** | Network-induced Latency | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Compute-induced latency | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Storage-induced latency | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| **GLOBAL EXPANSION** | Geographical presence | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Legal, compliance and regulations | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Programmability | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Caching and storage | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Traffic Engineering | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| **MULTI-USER SCALABILITY AND CONCURRENCY** | Failover and resilience | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | PoP resources | not very ○○○○○ very | Low OPEX ○ Convenience ○ |
| | Analytics and telemetry | not very ○○○○○ very | Low OPEX ○ Convenience ○ |

# Go global today.
# Scale without fear.

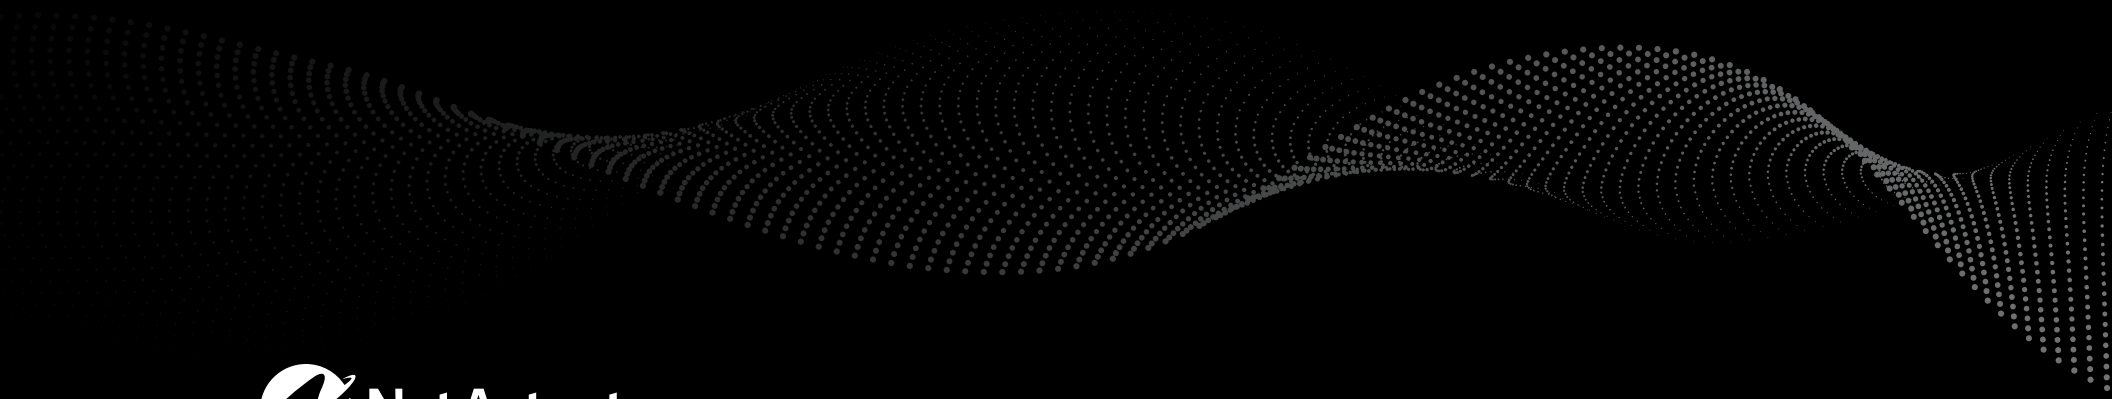

NetActuate

linkedin.com/company/netactuate-inc

github.com/netactuate

youtube.com/@NetActuate

**netactuate.com**