

5X
POWER SAVING

1/2
AWS COSTS

100
PERCENT SCALABLE



SEAMLESS TRANSITION



SELF-HEALING



QUANTUM ENCRYPTED

If You Run Ai, Where Will You Go?

Are you stuck with antiquated data center infrastructure that takes months to configure and is costly to manage? The increase in SIGNIFICANT energy costs, the scalability and seamless transition point to HyperCloud as the smart path forward. **We promise.**

Promise 1

Scale to massive traffic. with performance and resilience that scale with it.

Promise 2

No lifecycle management. No version incompatibilities or product integration.

Promise 3

Highly available. Use pure peer resources, no bottlenecks or protocol interfaces.

Promise 4

Enable multiple true administrative tenancy leveraging 1 cloud underpinning.

Promise 5

Zero people, therefore no human error, inconsistency and remove fragility.

Promise 6

Self healing, organizing and assembling. Requiring no people (like public cloud).

Promise 7

Always high performance. Unthrottled use and performance shall increase linearly with scale or better.

Promise 8

High auditable and transparent security without exception.

Promise 9

Have a consumptive abstraction. Focus on the workload and value creation, not infrastructure.

Promise 10

Delivered completely at a fixed price. We the CIOs own it complete. TCO is lower and different.

Contact us for more Info@Indexr.ai **408.123.4567**



Indexr.ai

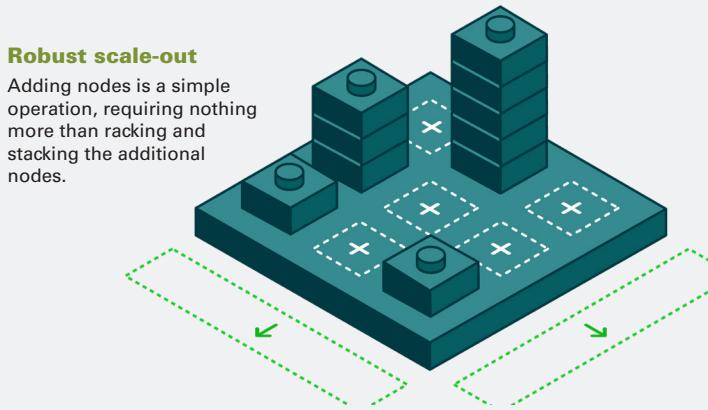
WORLD'S FASTEST GROWING INDEX

MORE ▾

More Deliverables. More Growth. More Profits.

- Stability under heavy load
- True high availability (no single supervisor choke point)
- Near-infinite scale-out capability
- Independent scaling of compute, storage, and networking
- Stateless, ephemeral node architecture
- Auto-configuring, hot-swappable nodes
- Self-healing fleet intelligence
- Rapid rack-to-ready deployment (hours, not weeks)
- Single-binary, non-disruptive stack upgrades
- Low administrative overhead / low-touch operations
- Predictable, fixed-cost ownership model
- Lower total cost of ownership (TCO) vs hyperscale OpEx
- Runs on purpose-built, task-specific hardware
- Optimized resource utilization (no wasted CPU/RAM)
- Efficient GPU management and passthrough for AI/ML
- Support for multiple GPU per VM and live GPU migration
- Built-in support for VMs and containerized workloads
- Integrated S3-compatible object storage
- Mix HDD / SSD / NVMe in a single cluster
- Crash-consistent snapshots and remote replication
- Native replication to remote HyperClouds or any S3 endpoint
- Certified data immutability for ransomware protection
- Automated, cluster-native disaster recovery
- Secure provenance (auditable hardware, firmware, software lineage)
- US-designed/manufactured hardware and supply-chain controls
- CVE Numbering Authority responsibilities and proactive CVE checks
- FIPS 140-2 support and FIPS 203 / post-quantum crypto support
- FedRAMP/FISMA/DoD compliance posture (deployed for classified use)
- Air-gap capable operation and air-gap upgrades (sneakernet support)
- Post-quantum full disk encryption with hardware token
- Hardware root of trust and tamper-resistant firmware practices
- Fine-grained RBAC and tenant isolation (secure multi-tenancy)

- Fine-grained RBAC and tenant isolation (secure multi-tenancy)
- Policy-driven resource allocation and QoS controls
- Reduced vendor lock-in through full-stack ownership and APIs
- Unified management plane (single pane of glass)
- Rich API + IaC friendliness (programmable automation)
- Native observability (OpenMetrics/Prometheus exporter Grafana)
- Integrated marketplace and one-click application deployment
- Deterministic performance scaling as nodes are added
- Optimized power and thermal profile (low power/heat design)
- Smaller physical footprint per delivered capacity
- Energy efficiency lowers operating expenses and cooling needs
- Rapid fault detection and automatic remediation
- Simplified firmware/driver/OS provenance for audits
- Predictable capacity planning and chargeback models
- Simplified compliance evidence and audit trails
- Reduced human error through automation and abstraction
- Better support for edge, tactical & disconnected environments
- Consistent behavior across sites (same hardware + software stack)
- Ability to mix CPU architectures while maintaining feature parity
- Marketplace templates and reproducible, operator-provided images
- Fast, deterministic upgrades from any baseline to current release
- Designed for AI/ML lifecycle (marketplace, GPU, orchestration)



Indexr.ai

WORLD'S FASTEST GROWING INDEX

Contact us for more: Info@Indexr.ai **408.123.4567**