HyperFRAME RESEARCH

Sponsored by: ORACLE

RESEARCH BRIEF

Oracle's Data-Centric AI Security Strategy:

Because You Can't Build an AI Foundation on Quicksand

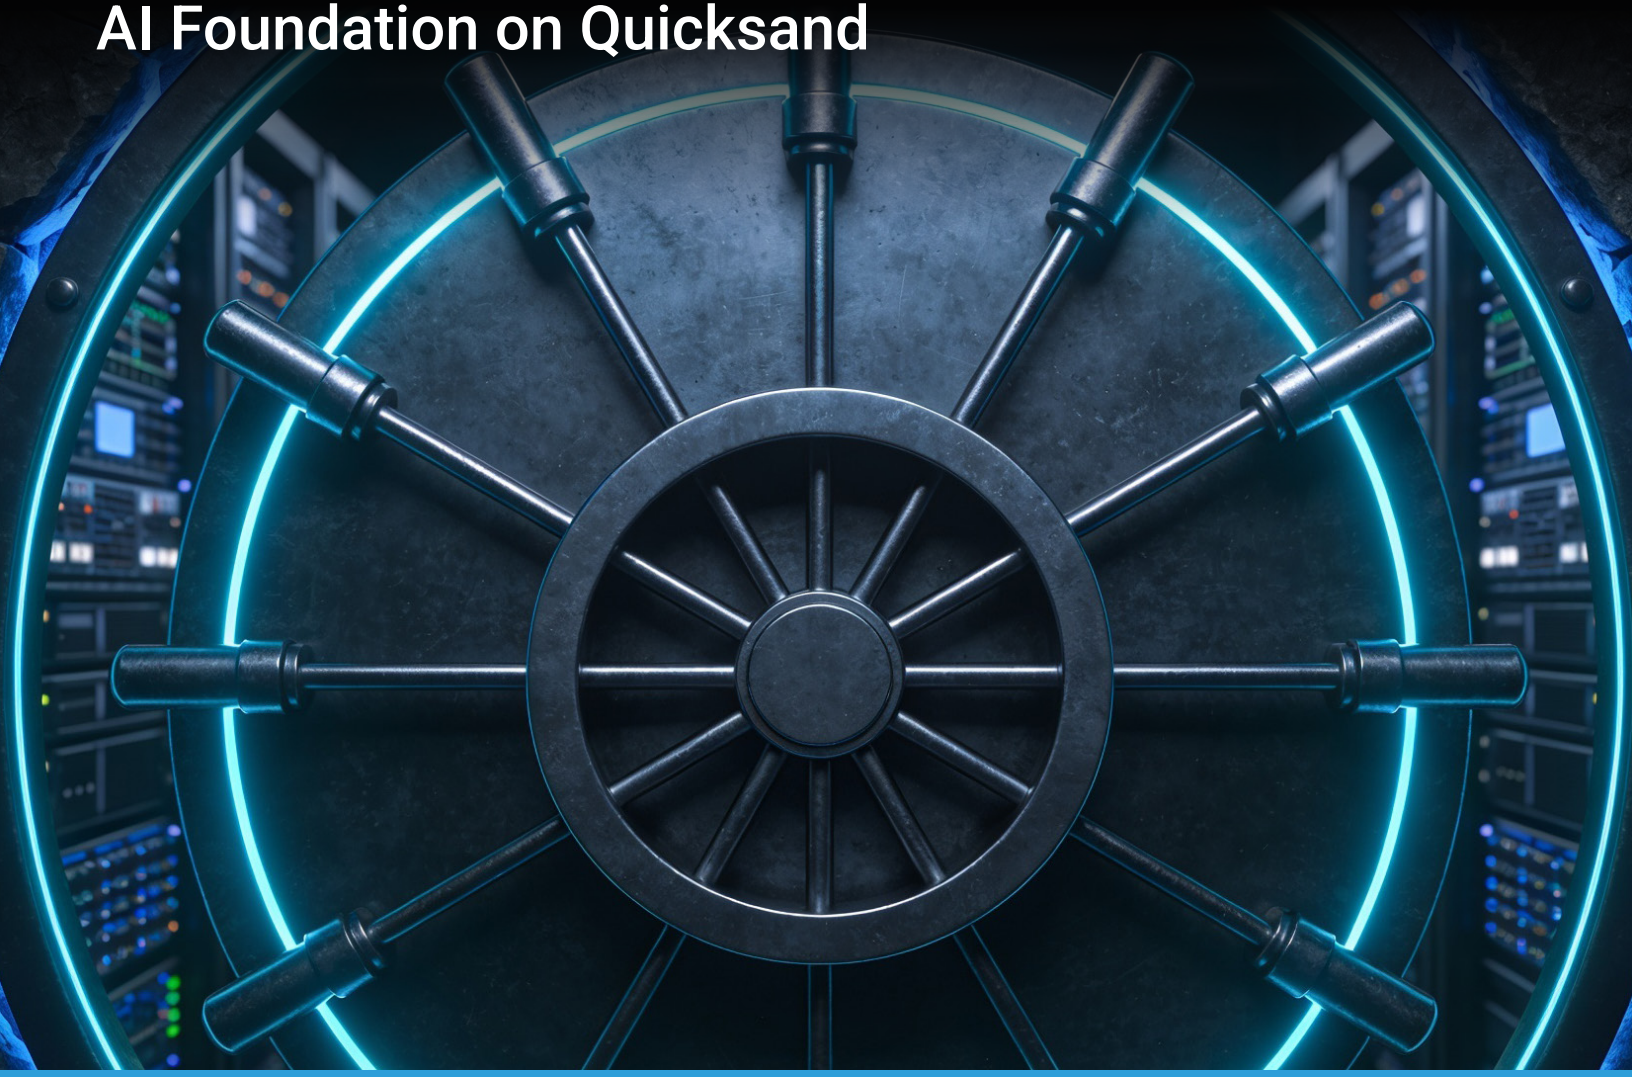

SEPTEMBER 2026

Authors:

Stephanie Walter
Practice Leader AI Stack

Ron Westfall
VP and Practice Leader for Infrastructure and Networking

Steven Dickens
CEO and Principal Analyst

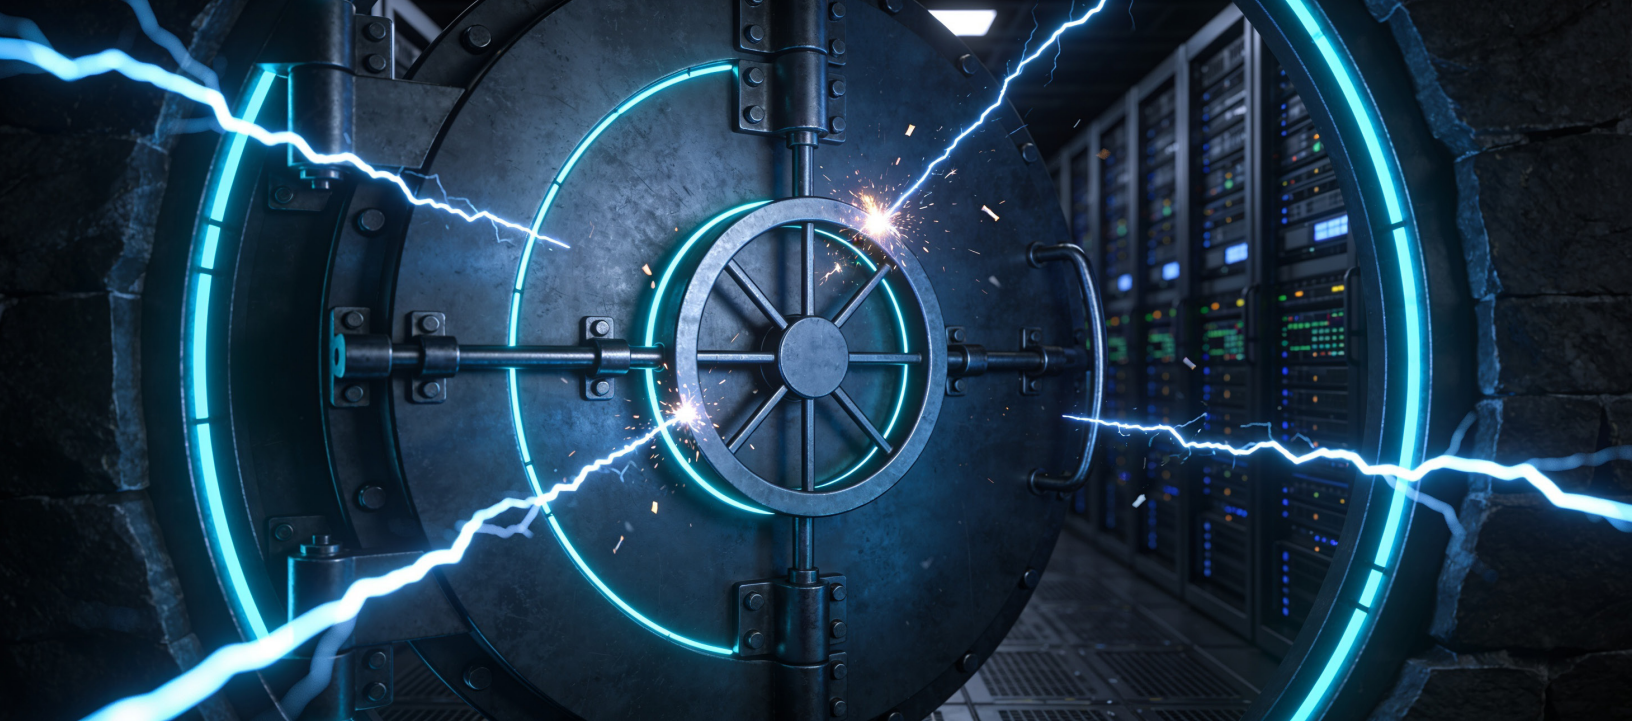

Oracle's Data-Centric AI Security Strategy: Because You Can't Build an AI Foundation on Quicksand

AI is changing the security equation. Threats now move at machine speed, while AI agents can discover and exploit vulnerabilities faster than traditional security models can respond.

Oracle's answer is a data-centric security strategy that anchors governance at the data layer, automates patching and lifecycle management, strengthens operational resilience, and makes critical security capabilities easier to adopt, including widely used tools offered at no cost or discounts of up to 90%.

The premise is simple: AI needs a secure, coherent data foundation. Pipelining data across multiple databases, maintaining consistency across copies, managing fragmented security profiles, and relying on external agent caching creates complexity precisely where enterprises can least afford it. It creates an IT environment akin to quicksand.

And you can't build an AI foundation on quicksand. The more fragmented the data and security architecture becomes, the harder it is to govern, secure, and trust—and AI only accelerates the consequences. The age of AI rewards speed, but if your foundation is built on quicksand, acceleration just means sinking faster.

The Challenge: Securing Data Against Machine-Speed and Scale AI Exploits

Oracle's three-pronged framing of security threats presented by agentic AI highlights a critical paradigm shift where the technology simultaneously amplifies enterprise vulnerability surfaces and supercharges adversarial capabilities through automated exploit generation. To survive this accelerated threat environment, organizations must pivot from perimeter-centric defenses to direct data protection, treating the data itself as the primary perimeter.

Furthermore, the unprecedented speed of AI-driven attacks renders traditional patch-management life cycles obsolete, mandating an evolution toward automated, frequent and rapid patch deployment. By proactively minimizing exposure gaps, enterprises can shrink the window of vulnerability that automated malicious scanners rely on for initial access.

This strategic posture recognizes that countering machine-scale threats requires an equally agile, data-centric security architecture capable of outpacing algorithmic exploits. However, achieving this is often hindered by time-consuming regression testing and strict uptime requirements that force security teams to balance operational stability against cyber threats. In response to these operational bottlenecks, Oracle is expanding access to tools intended to accelerate patching, testing, security assessment, and lifecycle management.

Why Agentic AI Is Forcing an Immediate Shift to Automated, Data-Centric Security

Why Agentic AI Is Forcing an Immediate Shift to Automated, Data-Centric Security

To combat a rapidly escalating threat landscape, Oracle has lowered the near-term cost of selected security, patching, testing, and lifecycle-management tools by making some available at no additional cost for a limited time and offering a 90% discount on one-year term licenses for others. The programs require current paid support and apply to upgrading and patching use cases. While automated patching is standard across all Oracle Database cloud services today, migrating workloads directly to the Oracle Autonomous AI Database provides the strongest long-term security posture; this transition replaces manual administrative overhead with Oracle managed, autonomous patching that minimizes operational costs.

This security modernization is particularly urgent given the rise of agentic AI, which has fundamentally disrupted traditional enterprise threat models. AI agents may be granted broad access to data and tools so they can act independently, increasing the consequences of overprivileged service accounts or inconsistent application-layer controls. Perimeter and application security remain necessary, but they increasingly cannot defend against AI-driven threats and need to be reinforced by granular controls at the data layer.

This research analysis examines how the proliferation of AI agents has introduced an unprecedented attack surface, rendering historic defenses obsolete while shifting the security frontier directly to the data layer. Concurrently, malicious actors are weaponizing adversarial AI to scan for system vulnerabilities and deploy automated exploits at machine speed and scale—effectively stripping security teams of the time required for manual remediation and mandating an immediate shift toward automated, data-centric defense architectures.

Our research indicates that enterprises are actively transitioning to data-centric security frameworks to enforce continuous governance, and direct protection at the source (according to HyperFRAME Research Lens: State of I&O Strategy in AI Era 2H 2026). As such, data security and governance is the leading force shaping design (62%). When AI agents require direct data access and bypass legacy perimeters, the data layer itself becomes the primary security boundary, forcing organizations to implement their storage and data architecture with strict security controls.

Organizations must transition to a data-centric strategy and Oracle is leading the way. With its converged database, Oracle focuses on securing all data types at the source and protecting data wherever AI accesses it with security built-in at the data layer; securing data at speed by staying ahead of AI-driven threats with machine speed automated patching and lifecycle management in the face of shrunken vulnerability windows; and, securing data through resilience, where organizations can recover quickly from disruptions further shrinking vulnerability windows.

Secure at Source: Anchoring Data-Layer Governance to Prevent AI Agent Bypasses and Enable Safe Enterprise Innovation

Application-centric security architectures inherently rely on predictable user journeys, whereas agentic AI introduces non-linear, unpredictable interaction paths that can unintentionally exploit logical gaps within code-level access controls. Relying on developers to consistently hardcode security parameters across a sprawling ecosystem of microservices and applications introduces human error, a flaw that centralized database-level defenses systematically eliminate. By decoupling security logic from the fast-evolving AI application layer and implementing security at the data layer, enterprises can rapidly iterate and deploy new AI tools without needing to audit or rewrite underlying application security code each time.

Deep Data Security is designed to enforce identity- and context-aware row-, column-, and cell-level authorization across agents, applications, and agentic AI workflows accessing governed data. As agentic AI workflows increasingly synthesize unstructured data, moving security boundaries to the source ensures that the underlying data's classification dictates its accessibility, rather than the security posture of the consuming AI application. Additionally, centralizing security within the database reduces the computational and architectural overhead required to maintain identical access control lists across multiple distinct AI training and inference environments. Finally, Deep Data Security preserves end-user and agent identity in database audit records, improving traceability for compliance and forensic analysis.

Oracle SQL Firewall can establish an allowlist of expected SQL activity and block unauthorized statements, reducing the ability of a compromised or manipulated agent to execute unfamiliar query patterns. Its effectiveness still depends on the quality of the allowlist and whether a malicious action resembles previously authorized activity. Oracle Database Vault can separate administrative duties and restrict privileged access to protected data, helping reduce the damage an agent could cause if it obtained or operated through an overprivileged account. Furthermore, these controls can provide an important backstop against prompt injection by preventing the database from returning information that the underlying user or agent is not authorized to access.

From an AI stack perspective, Oracle's strongest argument is not that application and perimeter security disappear, but that agent authorization cannot depend entirely on the application making the correct decision every time. Agents introduce more dynamic access paths, more generated queries at machine speed, and a greater chance that user intent will be misinterpreted. Enforcing identity and access policy at the database is the last bastion of defense creating a final control boundary when the agent bypasses the application layer.

Secure at Speed: Accelerating Enterprise Patching and Proactive Risk Mitigation to Counter Machine-Velocity AI Threats

As attackers leverage machine-speed automation to exploit system vulnerabilities, modern organizations can no longer afford prolonged patching cycles and must prioritize closing exposure gaps immediately. To accelerate these security operations, Oracle is streamlining patch management, centralizing risk monitoring, and minimizing maintenance downtime through automated tools such as Database Lifecycle Management. By now offering these critical patching, testing, and compliance capabilities for free or at a 90% discount, Oracle removes procurement barriers to help enterprises rapidly mitigate operational risk and counter escalating AI-driven threats.

The significance goes beyond patching speed. As agents gain access to databases, tools, and production systems, every unpatched vulnerability can expand the potential blast radius of an agent credential or compromised workflow. Automated lifecycle management reduces that exposure, but enterprises still need regression testing, workload visibility, and clear rollback paths to ensure that faster patching does not introduce new operational risks. Oracle has thought this through and offers Real Application Testing to assess application impact before deploying patches, lowering operational risk. The objective is not simply to patch at machine speed, but to shorten vulnerability windows without sacrificing production stability.

Secure through Resilience: How Oracle's Advanced Recovery Frameworks Neutralize AI-Era Disruption and Help Minimize Downtim

As enterprise AI transitions from peripheral experimentation to core operational infrastructure, security incidents threaten more than just data exfiltration—they risk total business paralysis. Because absolute threat prevention cannot be guaranteed, modern security relies on an organization's capacity to minimize downtime and rapidly restore trusted data without manual intervention during a crisis. Oracle addresses this challenge by anchoring operational resilience in robust framework technologies—such as Zero Data Loss Recovery solutions, Globally Distributed AI Database, and Maximum Availability Architecture—utilizing immutable backups, automated failovers, and Raft-based consensus algorithms. By delivering transaction-level protection and failover measured in seconds, Oracle provides the continuous uptime, data consistency, and near-zero recovery objectives required for stateful, transaction-intensive AI workloads.

Furthermore, incorporating immutable backups into a Zero Data Loss air-gapped framework provides a crucial defense against modern, highly sophisticated ransomware strains that actively target secondary storage and recovery catalogs to force ransom compliance. This native integration of high availability and continuous data protection eliminates the latency and fragmentation inherently created when piecing together disparate third-party backup vendors. Ultimately, Oracle transforms resilience from a reactive disaster recovery exercise into a competitive operational advantage, enabling enterprises

to deploy mission-critical agentic AI automation with stronger and more predictable business-continuity characteristics.

Resilience in the agentic era extends beyond restoring the database. Enterprises must be able to re-establish trust in the context, state, permissions, and actions surrounding an interrupted agent workflow. Oracle's transaction-level recovery capabilities provide an important foundation, but customers will still need application-level controls to determine which agent actions should resume, be replayed, or require human review after recovery.

Oracle's Architectural Advantage in the AI Database Market

With respect to the competitive environment, several database vendors have pieces of what comprises Oracle Deep Data Security, but there is not a single major competitor currently offering the same precise combination as a single, database-native authorization framework explicitly designed for agentic AI.

Oracle Deep Data Security in Oracle AI Database 26ai is more than conventional row-level security. Its differentiator is that the database can establish the actual end-user security context behind an application/AI agent and then enforce declarative authorization at row, column, and individual-cell level during SQL execution, regardless of whether the access came through an application, analytics tool, RAG workflow, AI agent, or direct SQL path.

That distinction matters because if an agent is compromised through prompt injection or simply generates inappropriate SQL, the agent shouldn't inherit a broad database service account's access. The database independently limits the result to what the human principal is entitled to see. Oracle explicitly positions this as moving authorization out of application/agent code and enforcing it at the data source.

Microsoft SQL Server and Azure SQL have largely separate mechanisms that customers assemble. Microsoft does not currently offer a unified database-native authorization model specifically designed around an AI agent acting on behalf of an end-user, comparable to Deep Data Security.

PostgreSQL, including AWS RDS and Aurora PostgreSQL, provide many of the primitives needed to construct a secure agentic application, but developers generally have to build the identity propagation, context management, and policy architecture themselves. There is no direct PostgreSQL or AWS equivalent to Deep Data Security as an integrated framework.

Google AlloyDB and Cloud SQL for PostgreSQL are similar. Customers can use PostgreSQL Row-Level Security together with Google Cloud IAM and Google's broader security infrastructure. That provides useful building blocks for protecting data accessed by applications and AI systems, but it remains an assembly of capabilities rather than a unified database authorization framework specifically intended to preserve and enforce an end user's security context through an AI agent.

Snowflake's model is primarily centered on cloud data and analytics governance rather than Oracle's combination of transactional database workloads, enterprise applications, and AI agents operating under an individual user's identity.

MongoDB approaches part of the security problem differently. Its role-based access controls are complemented by capabilities such as Queryable Encryption, which can protect sensitive fields while still allowing certain operations on encrypted data. But encryption addresses a different problem from Deep Data Security's contextual authorization model: determining which individual rows, columns, or cells a particular end user should be allowed to access when an application or AI agent acts on that user's behalf.

Databricks governs what data a user or AI application can see within the lakehouse. Oracle Deep Data Security goes even further by making the database the authorization authority for what an end-user, application, or AI agent is entitled to do with individual business data—including at the cell level.

Lastly, IBM does not package RCAC as an AI-agent security framework that combines end-user identity propagation, contextual authorization, and cell-level controls in the way Oracle Deep Data Security does.

The combination of end-user identity propagation, database-native contextual authorization, row/column/cell-level enforcement, and consistent enforcement across application, analytics, and AI-agent access paths makes Oracle AI Database the optimal platform for enterprise agentic AI.

Key Takeaways and Conclusion

In an evolving threat landscape where high-speed AI attacks and autonomous agents target sensitive assets, Oracle urges organizations to shift their security focus directly to the data layer. To achieve this defense, Oracle recommends transitioning to its Autonomous AI Database, upgrading to long-term support releases like Oracle AI Database 26ai, and leveraging zero-cost or discounted lifecycle management and testing tools. This strategic roadmap enables modern enterprises to close AI-driven vulnerability windows, streamline and test patch deployments, and establish permanent operational resilience where risks are highest.

Oracle's strategic emphasis on data-layer security validates the necessary market shift to prioritizing data-layer security, as traditional perimeter defenses prove increasingly porous against autonomous, AI-driven threats capable of discovering and exploiting zero-day vulnerabilities in real time. By embedding autonomous self-patching and risk auditing directly into Autonomous AI Database, the company reduces the human latency associated with database patching, security assessment, and recovery, particularly when customers use this service.

Oracle's strategy also addresses a major economic and operational bottleneck by offering discounted lifecycle management and testing tools, lowering the friction for risk-averse organizations hesitant to undertake more frequent database patching. Moreover, unifying relational, vector, and unstructured data under a single autonomous defense perimeter eliminates the configuration errors typical of multi-vendor, piecemeal security stacks. Oracle's approach moves enterprise cybersecurity from a reactive, perimeter-focused posture to a holistic model where security and resilience controls remain enforceable closer to the data itself. After all, the cost of not doing anything at all to prevent a cybersecurity breach is more expensive than actually implementing the right strategy in the first place.

ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

CONTACT HYPERFRAME RESEARCH:

Steven Dickens
CEO & Principal Analyst | HyperFRAME Research

Email Address:

steven.dickens@hyperframeresearch.com

Telephone Number:

+1 845 505 1678

X: @StevenDickens3

LinkedIn: Steven Dickens

BlueSky: Steven Dickens

CONTRIBUTORS

Ron Westfall
VP and Practice Leader for
Infrastructure and Networking

Stephanie Walter
Practice Leader, AI Stack

Steven Dickens
CEO & Principal Analyst

INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

LICENSING

DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.


HyperFRAME
RESEARCH