A field guide for AI-powered data security: How to deliver breakthrough business outcomes

Turbocharge your enterprise with speed, security, scale, simplicity, and smarts from the Cohesity Data Cloud.

COHESITY

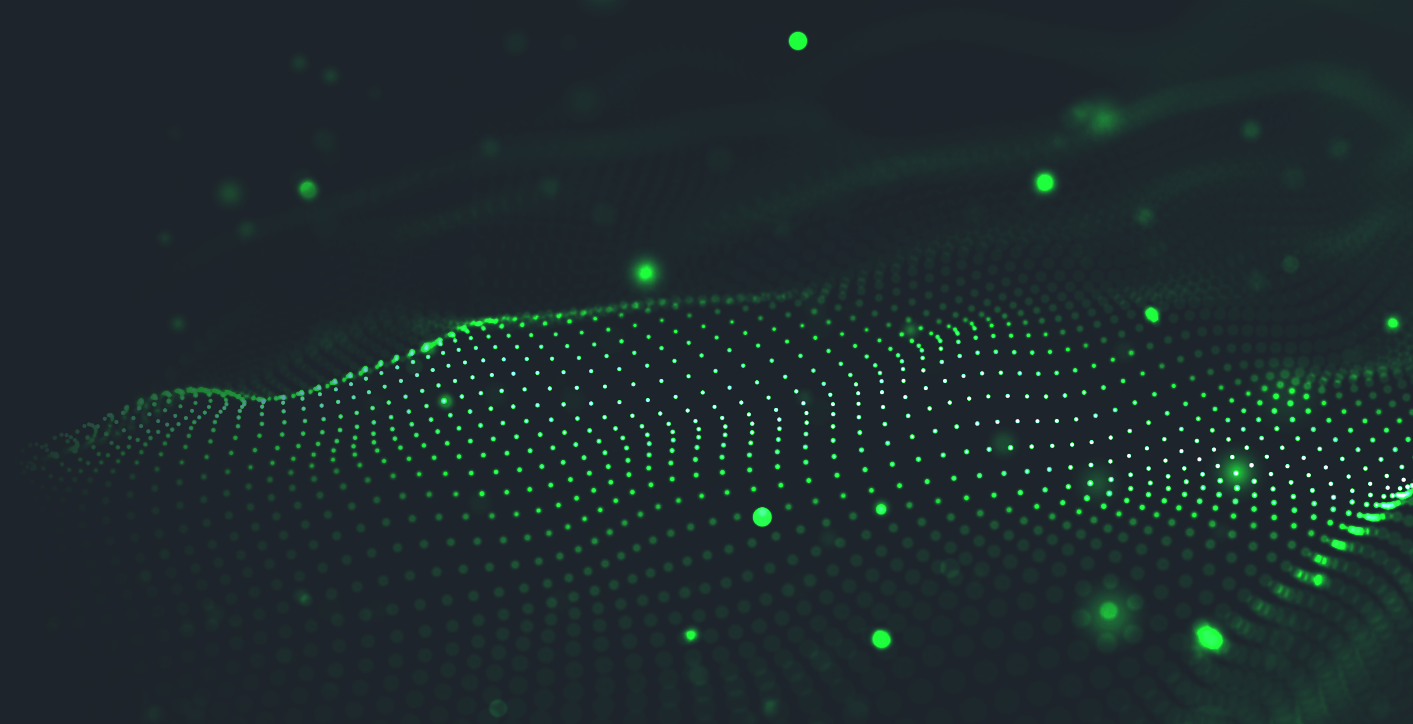

COHESITY

Table of Contents

Executive summary

Three market trends in enterprise data security and management have emerged in recent years:

1. Cyber threats are driving the replacement of traditional backup systems in favor of more resilience-oriented solutions.
2. IT leaders face complexity on a number of fronts: a growing data estate, running atop heterogeneous infrastructure targets; an evolving cyber threat landscape; and numerous regulatory and compliance requirements.
3. Disruption from generative AI technologies brings new opportunities for growth and differentiation, as well as new risk vectors. This is a new twist on the conventional wisdom that data is your most valuable asset.

IT leaders must secure their data estate, so they can "play defense" against attackers. At the same time, leaders must "play offense" and realize an enduring competitive advantage from their data with AI.

The Cohesity Data Cloud is a popular choice for IT leaders who want to address these trends head-on.

Organizations that modernize with the Cohesity Data Cloud often achieve superior outcomes in five key areas:

- **Speed** - They are able to recover from cyberattacks many times faster compared to their previous systems.
- **Security** - They improve their security posture, and detect threats, protect data, and rapidly recover from cyberattacks.
- **Scale** - They can secure and protect their entire data estate on a single platform, even at petabyte scale.
- **Simplicity** - They can run their data estate and perform backup and recovery workflows from a unified control plane and set of APIs.
- **Smarts** - They gain business and operational insights from their data, with advanced AI capabilities.

In this white paper, we summarize the core capabilities of Cohesity Data Cloud that enable these outcomes and provide useful recommendations for how best to achieve them.

Note: The customer outcomes cited in this paper are available from a recently published white paper from IDC and production customer deployments, some of which are public case studies.

Speed: Accelerate cyber recovery

In the event of a cyberattack or other unplanned disruption, it's crucial to return to "business as usual" quickly. The longer you're offline, the greater the financial and reputational impact.

Why do organizations usually struggle to recover quickly from a cyberattack?

The culprit is usually a patchwork portfolio of data management tools built atop older file systems. While these tools have proved useful historically, they often take far too long to rehydrate virtual machines, databases, and NAS environments.

As a result, it's difficult for IT teams to regularly test their cyber recovery plans. When the inevitable attack occurs, even the most thoughtful recovery plans cannot be executed. We all see this play out in news headlines over and over.

One of the inherent advantages of the Cohesity Data Cloud is its speed of recovery. Our speed of cyber recovery is 5 to 10x faster than other alternatives thanks to our platform architecture. We can also help you perform fast restores in a granular or a mass recovery fashion for many workloads, whether virtual machines, databases, or NAS files, and whether on-premises or in the cloud.

Cohesity customer outcomes

- Decreased backup times by 45%
- 10x faster data recovery times compared to legacy infrastructure
- 97% faster file restores

> "Cohesity can help us recover from a large-scale outage affecting thousands of virtual machines within hours—compared to weeks or months with our legacy system. Recovering hundreds of virtual machines...now takes minutes instead of weeks."
>
> - **Guru Vasudeva,** Nationwide

The Cohesity Data Cloud difference

First things first: Meet your predefined backup SLAs

The Cohesity Data Cloud enables you to consistently meet your predefined backup SLAs—the first thing you need to achieve fast, predictable recovery. Full backup operations on traditional systems can take hours, and are often error-prone, driving up operational cost and prolonging recovery times.

Your data estate at your fingertips: Global search

Before you can start a recovery, you have to identify the resources you want to recover. Use the Cohesity Data Cloud's powerful global search to find any virtual machine (VM), file, or object across all your data sources and locations. Results appear in seconds, so you can initiate recovery immediately. Other alternatives require you to know what VMs and files or objects are named, and/or where they are stored, to initiate recovery. That takes valuable time when you have hundreds, even thousands, of siloed systems in your data estate.

Ensure your snapshots are clean with health assurance

Fast cyber recovery requires at-a-glance assurance that your backup snapshots are healthy and ready to be recovered. The Cohesity Data Cloud provides backup verification and snapshot health assurance with automated, policy-based visibility into your snapshot's health within the recovery workflow. This provides peace of mind that your recovery is more predictable and efficient. With other systems, your IT staff may need to identify the best backup copy. Restore the wrong copy, and you make a bad situation worse—recovering from an already compromised snapshot.

Recover with confidence thanks to data and application consistency

The Cohesity Data Cloud features strict consistency: your data is first protected across the cluster before acknowledging the write back to the application. Think of it like an ATM transaction—Cohesity makes sure the transaction processes before it closes. Without these consistency measures in place, you're likely to get a 'data not found error' during an attempted recovery operation. Or worse, you may face a permanent data loss scenario.

Lightning-fast recovery at scale

Once your systems have been verified to be free of infection, you'll want to recover quickly at scale. The Cohesity Data Cloud always maintains an unlimited number of fully-hydrated backup snapshots that can be instantly mounted, making your data readily available (via direct mount) when you need it, while restoring data to production locations in the background. Additionally, our distributed architecture and MegaFile capabilities intelligently "chunk" any data larger than 256 GB, and then streams it across the entire Cohesity cluster for rapid restore.

The freedom of choice: Restore from any point in time

It may seem logical to always look to restore from your latest backup, however, this may not always be a best practice. This is especially true after a ransomware attack when you might need to recover from an older, cleaner copy written before the attack. No matter your situation or business need, Cohesity gives you the flexibility to restore to any previous point in time. Rapid recovery point objectives (RPOs) powered by Cohesity give you flexibility to choose which recovery point is most suitable for your organization to recover predictably. The platform will even suggest a clean point in time for recovery, using our AI technologies.

Recover anywhere, because it's a multicloud world

Your data estate likely spans on-prem environments, public clouds, and edge locations. You should be free to recover your data and applications wherever you'd like, not just the original location. The Cohesity Data Cloud gives you this flexibility. When you have more location options for recovery, you can better meet business SLAs.

Security: Reduce your risk

Bad actors often go after your secondary data in an attack. The reason is twofold. First, backup data is often not protected as robustly as production systems. Second, once the backups are gone, the enterprise has no restoration options.

Because your data is essential to your organization, establishing efficient, reliable, and secure access to your data is crucial. Your data security and management platform should have defenses in place to guard against attacks on your data and platform.

Most of today's security threats can be classified as:

- Ransomware, cyberattacks, and other malware
- Breaches due to unauthorized access and insider threats
- Accidental data leaks and misconfigurations

The Cohesity Data Cloud has strong, built-in defense mechanisms that guard against attacks and breaches that can lead to data loss.

Inspired by web-scale principles, the Cohesity Data Cloud's security-first architecture, combined with secure software development and release practices, helps ensure enterprise-class security.

Cohesity customer outcomes

- $0 ransomware paid in multiple cases
- $2M saved in cyber insurance costs
- 17% more threats detected

> "Little did we know that Cohesity was going to literally rescue Sky Lakes from having to pay a ransom. Cohesity just flat-out worked. We ultimately recovered all of our servers, and Cohesity worked time and time again throughout that process."
>
> - **John Gaede,** Sky Lakes Medical Center

The Cohesity Data Cloud difference

Encryption of data at rest and in motion

Cohesity Data Cloud encrypts all data and data flows within the platform. Encryption prevents unauthorized users from viewing data outside of the platform; data stored in the platform is unintelligible unless accessed and decrypted by an authorized user. Most privacy and industry regulations, notably GDPR, CCPA, PCI, and HIPAA, require organizations to protect sensitive data with encryption.

Platform data is encrypted at rest using AES-256 encryption. The platform has multiple options for securely managing encryption keys—either using Cohesity's managed Key Management Service (KMS) or via Amazon Web Services KMS or other third-party vendors, such as HashiCorp, Thales, Fortanix, and Entrust.

For data in flight, the Cohesity data management platform uses the TLS standard. TLS encrypts data so eavesdroppers and hackers are unable to see data flowing to and from the platform. This is critical to protect private and sensitive data for security and compliance. Cohesity uses the TLS 1.2/1.3 protocols with mTLS for transport layer security with only FIPS-approved cipher suites with Perfect Forward Secrecy (PFS) protection.

Fault tolerance

The Cohesity Data Cloud is a highly available system, with a fault tolerant architecture that prevents outages. Clusters can continue operating with multiple failures of HDDs and SDDs and nodes, chassis, and racks. The clusters can also sustain faults to power supplies, fans, and networks.

Immutable data storage

Data backed up by the Cohesity Data Cloud will never change from its saved state. Our underlying file system provides immutable backup snapshots to prevent modification or deletion of data. Based on hyperscale architecture, Cohesity stores backed-up data in its secured file system called "Cohesity Views" that is inaccessible from outside a Cohesity cluster. The backup snapshots are stored in a read-only state; no external application or unauthorized user can modify the snapshot.

Any attempts to write to an immutable backup snapshot are written on (zero-cost) clones, which are also marked read-only upon completion of each Protection Run. For any mount-based restores used during Cohesity's instant mass restore process, the internal view is first cloned and then exposed to the external environment, always keeping the internal view inaccessible externally. Writes to internal views during backup are only allowed via trusted internal services and authenticated APIs. For additional security, Cohesity views include DataLock, Cohesity's write once read many (WORM) feature. If DataLock is enabled, the backup snapshot can't be deleted by anyone, including administrators, until the DataLock expires.

Access control: Zero Trust architecture

As defined by the National Institute of Standards and Technology (NIST), Zero Trust is as follows: "… the term for an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources." Zero Trust in the context of the Cohesity Data Cloud focuses on validating the authenticity and authorization of users for any access or changes to the platform.

Multifactor authentication (MFA)

MFA provides strong authentication of users to thwart unauthorized changes to the platform setting or data. MFA improves platform security by requiring users to identify themselves by more than a username and password. Passwords and usernames are susceptible to brute force attacks and can be stolen. MFA requires the user to authenticate login requests with a response only they can provide (such as a mobile phone challenge), or Time-based One-time Password (TOTP). Cohesity supports native MFA or third-party MFA providers such as Ping, Duo, Okta, and more.

Role-based access controls (RBAC)

Granular role-based access control in the Cohesity Data Cloud enables organizations to grant the least privilege required for users to execute their job requirements, minimizing risk and keeping areas outside their responsibilities unreachable. Organizations can restrict Cohesity user roles to specific applications, capabilities, or workflows in the platform, thereby limiting what a user does based on their role and responsibilities. For example, organizations can restrict specific users to only perform backups or data discovery.

Quorum

Cohesity Data Cloud uses quorum features to prevent unilateral changes to the platform within administrative accounts. This crucial control protects against unintentional user error, rogue admins, or compromised accounts. With quorum, user requests to change settings or administrative functions require multiple approvals.

Auditing

The Cohesity Data Cloud maintains a user audit trail for all actions performed on the Cohesity cluster. These records provide proof of compliance and operational integrity. Audit trails can also identify areas of noncompliance by providing information for audit investigations. Audit logs capture user activity for login/logout, changes to data or the data's properties, and job scheduling. The platform organizes logs by categories, such as Active Directory or Cluster, for rapid analysis.

Continuous monitoring

The Cohesity Data Cloud provides environment monitoring to help reduce the risk of human errors and misconfigurations. These capabilities scan the environment, including an array of security configurations, access controls, audit logs, and encryption frameworks, which are critical to protecting the security posture of the data cluster.

The Cohesity Data Cloud capabilities map to the NIST cybersecurity framework

The Cohesity Data Cloud doesn't just protect backup data. It also takes a modern approach to threat detection to enable effective containment of infected systems. These features, coupled with our rapid cyber recovery capabilities, accelerate the cyber incident response process. Our customers enjoy these benefits by having their data under management by Cohesity, lowering effort and cost of implementation.

The NIST cybersecurity framework is a useful way to assess your strategy. It includes six pillars:

- **GOVERN** - The organization's cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored.
- **IDENTIFY** - The organization's current cybersecurity risks are understood.
- **PROTECT** - Safeguards to manage the organization's cybersecurity risks are used.
- **DETECT** - Possible cybersecurity attacks and compromises are found and analyzed.
- **RESPOND** - Actions regarding a detected cybersecurity incident are taken.
- **RECOVER** - Assets and operations affected by a cybersecurity incident are restored.

The first pillar—govern—is developed by the organization, and refined over time.

The Cohesity Data Cloud offers substantial capabilities for the other five pillars:

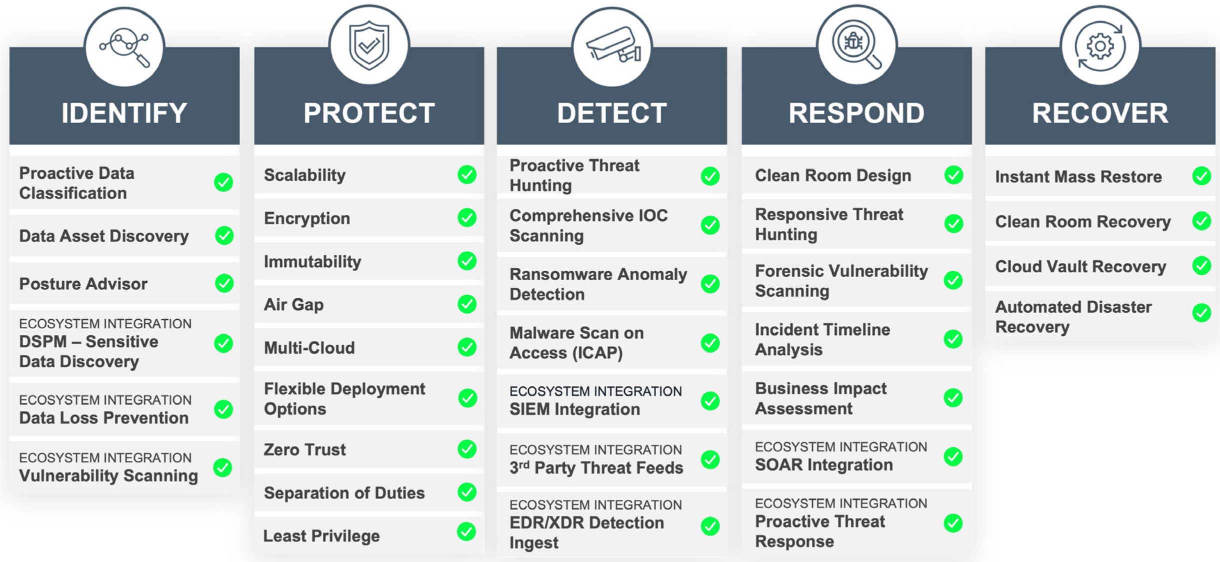

IDENTIFY		PROTECT		DETECT		RESPOND		RECOVER	
Proactive Data Classification	✓	Scalability	✓	Proactive Threat Hunting	✓	Clean Room Design	✓	Instant Mass Restore	✓
Data Asset Discovery	✓	Encryption	✓	Comprehensive IOC Scanning	✓	Responsive Threat Hunting	✓	Clean Room Recovery	✓
Posture Advisor	✓	Immutability	✓	Ransomware Anomaly Detection	✓	Forensic Vulnerability Scanning	✓	Cloud Vault Recovery	✓
ECOSYSTEM INTEGRATION DSPM – Sensitive Data Discovery	✓	Air Gap	✓	Malware Scan on Access (ICAP)	✓	Incident Timeline Analysis	✓	Automated Disaster Recovery	✓
ECOSYSTEM INTEGRATION Data Loss Prevention	✓	Multi-Cloud	✓	ECOSYSTEM INTEGRATION SIEM Integration	✓	Business Impact Assessment	✓		
ECOSYSTEM INTEGRATION Vulnerability Scanning	✓	Flexible Deployment Options	✓	ECOSYSTEM INTEGRATION 3rd Party Threat Feeds	✓	ECOSYSTEM INTEGRATION SOAR Integration	✓		
		Zero Trust	✓	ECOSYSTEM INTEGRATION EDR/XDR Detection Ingest	✓	ECOSYSTEM INTEGRATION Proactive Threat Response	✓		
		Separation of Duties	✓						
		Least Privilege	✓						

When speed meets security: Cohesity clean room design

In the middle of a ransomware attack, it's a mistake to bring your systems online before you've successfully diagnosed the root cause and mitigated the threats and vulnerabilities. InfoSec teams must conduct a forensic analysis of the attack and understand how systems were compromised. Only then can an organization remedy the infected systems and begin the rapid restoration process.

The Cohesity clean room design helps organizations recover clean data into production and then rebuild, or recover systems into a trusted state. This orchestration is done in three phases: initiation, investigation, and mitigation.

Initiation

The first stage of restoring trust in your environment is to create a Minimum Viable Response Capability (MVRC) that allows you to bring critical systems online and begin the investigation. Without defining an MVRC, teams may not have access to communication, such as email or even phones, which further hampers the response.

Investigation

The clean room design allows for an investigation into the attack to create a manifest of IoCs, including files, registry keys, and user accounts for use during recovery—thereby removing threats and patching the vulnerabilities that allowed the attack to be successful. Unfortunately, some people believe that a cursory threat scan is sufficient to remove the threats. It's not. This approach will inevitably cause malware to be reintroduced into the environment and extend the outage.

Mitigation

Once we understand the details of the attack, we can move to the next stage of mitigation, where threats can be removed and clean data recovered. During this stage, we understand what accounts we need to patch and remove. We'll also define the security tools that will be improved to help ensure similar attacks won't get through.

Learn More

Scale: A modern platform built for the cloud era

Before they start the modernization journey, most organizations have a highly fragmented approach to secondary storage. Data is usually managed across a patchwork of point systems, including dedupe appliances, backup servers, cloud gateways, file storage, object storage, and data lakes.

Each silo needs to be provisioned, configured, managed, and updated through its own proprietary UI and processes. This approach is also inefficient: each silo likely sits idle the majority of the time, and data is copied an average of 10 to 15 times across silos. Plus, security capabilities are often "bolted-on" to these systems, rather than "built-in" natively, leading to added cost and complexity.The business rationale for a modern platform is clearly compelling. Retiring dozens of systems into a single, scalable alternative is a superior value.

But not every modern platform is created equal.

Cohesity customer outcomes

- 40K+ M365 users
- Billions of objects
- Multi-PB of data stored in a cloud cyber vault
- VMware, Oracle, Physical, and SAP HANA at PB scale

The Cohesity Data Cloud difference

To effectively consolidate secondary storage silos, enterprises need a file system that can handle the requirements of multiple use cases simultaneously. It must provide strong performance for both sequential and random read/write operations, scalable snapshots, encryption, cloud tiering, and so on. Further, the underlying platform must provide native integration with the public cloud to support multicloud topologies. All of this must be done on a web-scale architecture to manage ever-increasing volumes of data effectively.

Cohesity invented the hyperconverged architecture many years ago. It meets all these requirements and more. Our customers run in the tens of petabytes, with some of our largest customers securing and protecting petabytes in the hundreds.

A new file system for a new era

The Cohesity Data Cloud is designed to effectively consolidate and manage all secondary data, including backups, files, objects, test/dev, and analytics data, on a web-scale platform that spans from the edge to the cloud.

This file system, called SpanFS, is designed to span everything:

- **Scale**: SpanFS provides unlimited scale across many hyperconverged nodes. SpanFS is completely distributed and doesn't have a master node. It scales linearly, and dynamically rebalances data as nodes are added or removed. It provides always-on availability, nondisruptive upgrades, and a pay-as-you-grow consumption model.
- **Private and public cloud**: SpanFS manages data across private data centers and public cloud sites. The public cloud can be used for archival, tiering, or replication. For replication, SpanFS is deployed in the public cloud to manage data in support of multiple use cases.
- **Secondary storage**: SpanFS supports data protection, files, objects, test/dev copies, and analytics data. It supports all the key capabilities required by these use cases, including globally distributed NFS, SMB and S3 storage, unlimited snapshots, global dedupe, encryption, replication, global indexing and search, and good performance for both sequential and random operations.
- **Tenants**: SpanFS supports multiple tenants with strong QoS capabilities, data isolation between tenants, separate encryption keys, and role-based access control.
- **Media tiers**: SpanFS spans across SSD and HDD media tiers and uses the most appropriate tier based on IO profiles.

Unlimited scalability in testing and in the real world

The Cohesity Data Cloud is designed to maximize scalability. Everything in the file system is completely distributed across all the nodes in a cluster. There is no master node and no single point of bottleneck. The data and the IO are dynamically balanced across the nodes, and individual nodes can be added or removed to adjust capacity or performance with no downtime. The system provides always-on availability and the data remains available even in the event of multiple node failures. Software updates are completely nondisruptive and done with rotating node updates.

Over 4000 of our customers depend on the Cohesity Data Cloud to secure and protect their data estate. We also regularly benchmark the performance of our platform.

Cohesity tested the scalability of our platform atop Microsoft Azure. The cluster was scaled from 8 to 256 nodes. As shown below, IO throughput scaled almost linearly for both sequential and random IO with inline dedupe turned on.

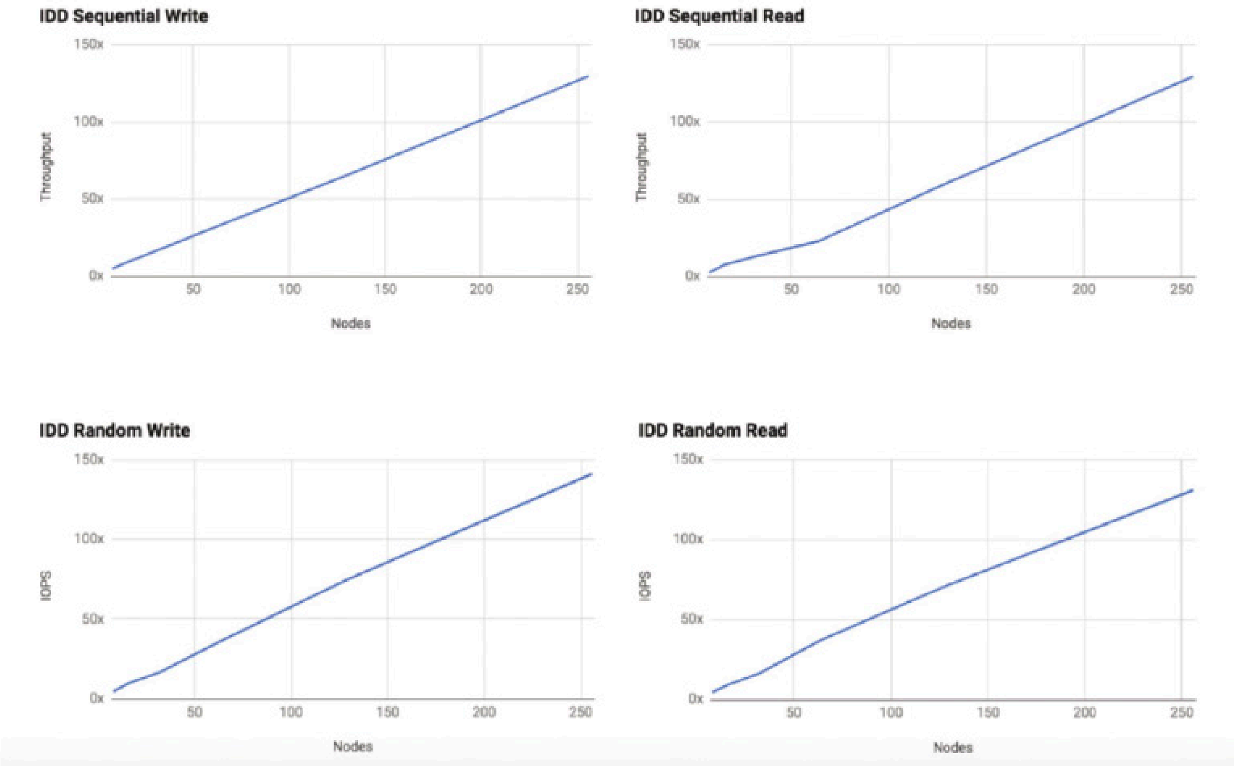

Random and sequential read/write performance of the Cohesity Data Cloud with inline dedupe turned on, scaling from 8 to 256 nodes in Microsoft Azure

Support for the data sources that matter

Your enterprise uses more data sources than ever. The Cohesity Data Cloud offers deep and wide support for the most important systems so you can enjoy economies of scale.

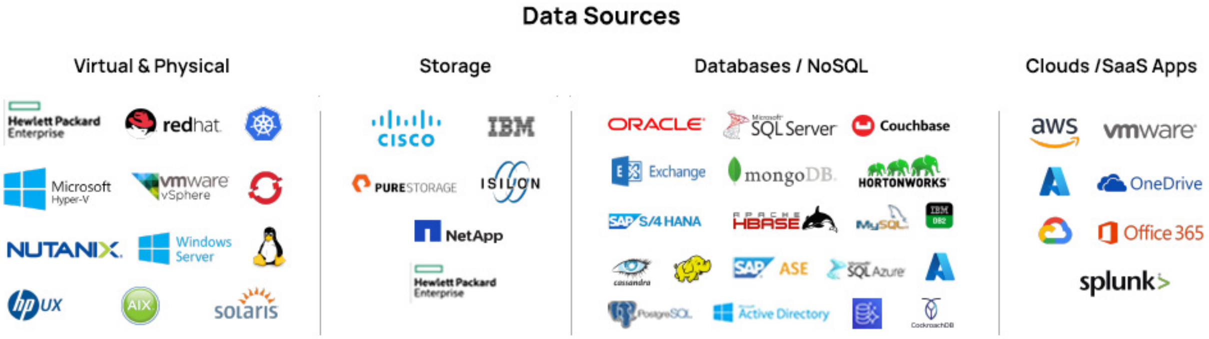

Simplicity: Secure and protect your data with ease

Traditional data management systems are full of silos. They largely consist of backup software, target storage, media and master servers, and bolt-on cloud gateways—all from different vendors. IT teams must monitor, keep track of, and think about multiple things while negotiating with multiple vendors. This fragmented, slow, on-premises, and expensive architecture drives the desire to modernize with a single platform that can support an enterprise-wide data estate in a much simpler way.

The Cohesity Data Cloud was born in the modern cloud era. We believe everything should be radically simplified, from our hyperconverged architecture to the user experience. We call it consumer-simple, yet enterprise-secure—inspired by companies like Google and Apple, where many of our founding team came from.

Our management control plane gets applause from users, analysts, and even design studios as one of the most beautiful user experiences in tech. We've dramatically simplified how you protect all of your workloads; you can do it all from a single pane of glass. Advanced capabilities in security, with the Security Console, and in AI, with Cohesity Gaia, are at your fingertips. (Learn more about Gaia starting on page 17.)

Cohesity customer outcomes

- 18 systems retired in favor of one modern platform
- 63% more VMs managed with each FTE
- 36% more efficient security and backup teams

> "We choose Cohesity because of its simplicity and superior automation capabilities."
>
> - **Rocco Amico,** Desjardins

The Cohesity Data Cloud difference

Manage all your Cohesity clusters from a single user interface

A single dashboard aggregates operational data, including the health of your global Cohesity infrastructure, the status of protected objects and SLAs, performance (throughput and IOPS), and more. Administrators can also:

- Orchestrate cluster upgrades (one or more) across locations simultaneously or scheduled to meet your business requirements.
- Easily access multiple pre-canned or tailored global reports. Generate global reports including but not limited to protection status, SLA violations, storage utilization, and reduction. Each of these reports can be exported and shared in multiple formats, including PDF and CSV.
- Set your desired thresholds and receive aggregated alerts of potential issues when exceeding these thresholds.

Integrated marketplace

Discover diverse apps and integrations in Cohesity's marketplace for streamlined data management. Optimize your strategy in one convenient location.

Comprehensive API ecosystem

- **Seamless Integrations**. Our APIs are crafted for easy integration, providing a seamless connection with our customer's existing workflows and applications.
- **API-First Architecture**. Every use case covered by Cohesity Data Cloud has an associated API. All our APIs are well documented and can be curtailed as per our customer's use case.
- **Comprehensive Documentation**. Navigate through our extensive documentation effortlessly. Clear, concise, and regularly updated documentation ensures a smooth integration process, reducing development time and potential roadblocks.

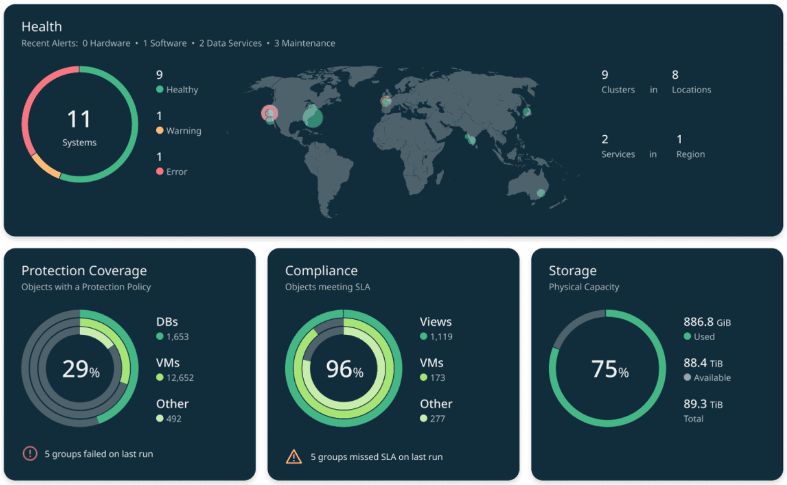

Smarts: AI-powered operational and business insights

Enterprises now expect their modern systems to use embedded AI to continually improve performance and efficiency. We call this capability “AI-powered operational insights.” Your infrastructure and InfoSec teams can use these features to run your data estate more intelligently, and use AI to investigate and remediate potential threats faster.

Modern data platforms offer even greater upside: to help leaders accelerate their AI strategy by bringing large language models (LLMs) to their secondary data. We refer to this concept as “AI-powered business insights.”

A modern data security and management platform should use AI to optimize all the aforementioned outcomes, and use retrieval augmented generation (RAG) to provide business insights from high-quality backup data.

Cohesity customer outcomes

- 48% less time to remediate threats
- 45% less staff time to detect threats
- 26% more productive compliance teams
- 46% reduction in TCO
- 233% three year ROI

> “Cohesity Gaia allows us to query our rich store of research data and quickly find relevant work. It will also allow our researchers to use their native language to query the system. This could prove incredibly valuable in accelerating the rate of our research and discovery.”
>
> - **Ryan Reed,** head of IT Products and Services, JSR Corporation

The Cohesity Data Cloud difference

AI-powered operational insights

The Cohesity Data Cloud ingests telemetry from your deployment to continually optimize efficiency of data storage, accelerate detection of potential anomalies, and assess the impact of potential data breaches.

Ransomware detection

The Cohesity Data Cloud offers early detection to help limit the damage caused by ransomware. Cohesity's ransomware detection techniques can analyze changes in:

- The volume of data read or written over time—noticing whether the volume of data protected from one protection job run to the next changed dramatically.
- Analyzing changes in data entropy—very good for detecting when the contents of files may have changed significantly due to encryption or compression.
- Changes in the nature of the file system—identifying when a large number of files has been added, deleted, or modified.

Anomaly detection

The Cohesity Data Cloud immediately analyzes data ingested from production environments on every backup for telltale signs of unusual activity or data changes. This activity may indicate a ransomware attack. A central dashboard displays alerts for anomalies based on how the timing and frequency of data reads and writes, and how the randomness of data and files change, including files added, deleted and modified. Using the Cohesity Helios anomaly detection feature, organizations can set alerts for conditions that could indicate ransomware or other malicious activity.

Data classification

Data proliferation defines the growing locations, volume, and diversity of data across organizations. With proliferation, organizations need automation to track the growing sources of critical and sensitive data to help ensure that data protection, security, analytics, governance, and privacy don't have gaps in coverage.

The Cohesity Data Cloud provides automated data classification so organizations can discover, classify, and tag sensitive information and ensure that protections and safeguards meet compliance and SLA requirements. AI algorithms create data-privilege maps detailing the location and classification of data.

These predefined policies help organizations meet their global and regional requirements for GDPR, CCPA, HIPAA, and other regulations. And organizations can use 100+ predefined patterns to create policies tuned for their specific challenges and needs.

Capacity prediction and planning

The Cohesity Data Cloud uses your capacity history and usage patterns to predict your future capacity needs. Forecasts are displayed in an intuitive graphical format, so you can take action to help ensure you have enough capacity in the coming months.

Threat intelligence

Cyberattacks use deceptive tactics to spread across your network. The Cohesity Data Cloud uncovers elusive threats using AI-driven threat detection with curated threat feeds featuring simple point-and-click threat hunting. You can even create custom YARA rules to identify advanced threats targeting your environment.

Data entropy

The Cohesity Data Cloud analyzes rates of changes in your data, a concept called "data entropy." When the rate of change in data is widely different from past trends, perhaps due to encryption or compression, that's a sign of a potential ransomware attack. The Cohesity Data Cloud will alert administrators accordingly.

AI-powered business insights

The Cohesity Data Cloud is a modern platform that's ready-made for the AI era. Your high-quality backups are indexed and stored in a powerful file system purpose-built for the AI services that are taking the world by storm. Cohesity Gaia, our AI product built atop our platform, is a RAG service that lets you have a conversation with your data.

A RAG engine for your enterprise data

Connect your data to Cohesity Gaia in any format from Cohesity-managed backups. Cohesity Gaia vectorizes the data, creating a baseline for answering questions on your enterprise data.

AI-powered conversational search UI

Start having a conversation with your data. Using common language, ask questions about your data and dig deeper into datasets.

Context-rich answers

Generate answers and insights based on questions about your enterprise data. Cohesity Gaia uses RAG AI to give more accurate answers, streamlining the compliance investigation process.

Fine-grained RBAC policies

Cohesity Gaia incorporates advanced, fine-grained RBAC policies that precisely govern access to the Gaia APIs. These specialized policies help ensure that only authorized users can access and manage the stored data, mitigating the risk of data exposure or unauthorized access. The RBAC system is designed to accommodate various levels of access and privileges, offering granular control over user permissions and actions within the Cohesity Gaia environment.

Conclusion and next steps

IT leaders must secure and protect more data, in more places, as attackers grow more sophisticated over time. Meanwhile, industry regulations grow more complex. And AI is evolving at a dizzying pace.

The Cohesity Data Cloud is a popular choice for enterprises like yours to survive and thrive in an increasingly complex environment. The platform can deliver remarkable business outcomes in five key areas:

Outcome	Before	With Cohesity Data Cloud
Speed	Recovery in days or weeks	Recovery in minutes or hours
Security	Bolted-on; limited intrinsic capabilities	Zero Trust principles; deep support for NIST cybersecurity framework best practices
Scale	Performance degradation as data volumes increase; siloed environments offer minimal automation	Consistent performance with scale-out hyperconverged architecture; support for a large ecosystem of data sources and cloud providers
Simplicity	Fragmented systems resulting in significant operational cost and complexity	Single UI and API to manage enterprise data estates at petabyte-scale (and beyond)
Smarts	Often requires traditional ETL procedures, or DIY development, to combine AI models with enterprise data	Native integration with AI capabilities for operational and business insights

From here, we recommend these next steps:

1. Establish your desired outcomes for cyber resilience, and benchmark them against your current capabilities.
2. Develop a rough prioritization or a roadmap of desired workloads and use cases for cyber resilience.
3. Assess the ROI and TCO of a modern data platform compared with your current solution.
4. Choose your solution based on product demonstrations, proven ROI and TCO calculations, and support for roadmap priorities.
5. Deploy your chosen solution, and execute your roadmap.
6. Experiment with generative AI capabilities and refine your overall AI strategy.

Once your modern platform is in place, conduct regular testing and "red team" attacks to assess your recovery capabilities against cyberattacks.

Sign up for a Cohesity Ransomware Resilience Workshop

Ransomware is on the rise. How prepared is your organization? Experience firsthand what it's like to have your company data held to ransom, discover potential vulnerabilities in your current setup, and learn from our team of industry experts—all in just 2 hours.

Sign Up

COHESITY

About Cohesity

Cohesity is a leader in AI-powered data security and management. We make it easy to secure, protect, manage, and get value from data—across the data center, edge, and cloud. Cohesity helps over 4,000 organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring malicious behavior, and rapid recovery at scale. Cohesity solutions can be delivered as a service, self-managed, or provided by a Cohesity-powered partner.

Learn how Cohesity can accelerate your journey to modern data security and management at www.cohesity.com.

Appendix A: Security is a team sport

Use existing tools with the Cohesity Data Cloud to help you identify, protect, detect, respond, and recover from cyberattacks

Most organizations have security tools for detecting malware, viruses, and vulnerabilities. Many also have established practices and policies for monitoring the health and safety of IT resources. The Cohesity Data Cloud embraces this entire ecosystem: our platform is designed to integrate into an organization's existing security operations.

This approach helps ensure consistent implementation of user authentication and data encryption, complements existing solutions and processes for incident response and management, and uses the threat and vulnerability detection solutions already in place.

Integrations support several use cases as described below.

Enhance threat hunting and response

Use curated threat feeds in the Cohesity Data Cloud to identify threats and known indicators of compromise in your backups—either proactively or during incident response.

Partner: CROWDSTRIKE

Understand risk exposure using backup data

Scan backup data with the Cohesity CyberScan app powered by Tenable—any time, including during incident recovery—to avoid reinjecting known vulnerabilities back into production.

Partner: tenable

Secure access to Cohesity clusters

Secure and control access to your Cohesity clusters through integrations with identity and access management (IAM) solutions to minimize risk.

Partners: CYBERARK, CISCO, okta, PingIdentity, Microsoft

Prioritize sensitive data protection

Identify gaps in protecting sensitive workloads and accelerate their protection to reduce cyber recovery and compliance risks.

Partners: WIZ, CYERA, Normalyze, DASERA

Prevent exfiltration of sensitive data

Identify sensitive data in your backups and prevent data loss.

Partner: zscaler™

Respond to and recover from a ransomware attack

Automatically protect data in response to a ransomware attack detected by Cisco XDR. Ingest data protection telemetry and execute automated recovery playbooks from your SOC tool.

Partners: cisco paloalto NETWORKS CROWDSTRIKE Microsoft servicenow

These integrations are part of our Data Security Alliance.

Learn from your peers

Check out these resources to better understand how we've helped customers like you.

Videos

- JSR Corporation turns to Cohesity for cyber resiliency
- AutoNation scales up and boosts data protection with Cohesity and AWS
- Why Broadcom chose Cohesity
- Nationwide Increases Resilience—and Saves $2M Annually—with Cohesity
- Why Delta chose Cohesity
- Why Salesforce chose Cohesity

Case studies

- The French Red Cross trusts Cohesity for migration, management, backup, and data protection
- Beckman Coulter secures data management with Cohesity backup as a service
- Nasdaq gets security, support, and SaaS with Cohesity
- Desjardins Chooses Cohesity-Cisco Solution for Faster Backup and Recovery, Simplicity, and Automated Data Management

Recommended reading

We think you'll find the following white papers, guides, and blogs helpful.

- Modern data security and management topologies: A guide for IT leaders
- An executive's guide to modern data security and management
- The business value of HPE solutions with Cohesity
- Cohesity SpanFS and SnapTree
- Cohesity Gaia: An executive's guide to accelerating enterprise AI with retrieval augmented generation and secondary data
- Evaluating cyber and data resilience solutions
- Introducing the Cohesity clean room design
- Cohesity Data Cloud Security White Paper (customer only, login required)
- Cohesity Platform Security Hardening Checklist (customer only, login required)
- Cohesity Platform Security Pocket Guide (customer only, login required)

COHESITY

www.cohesity.com

2000054-002-EN 9-2024