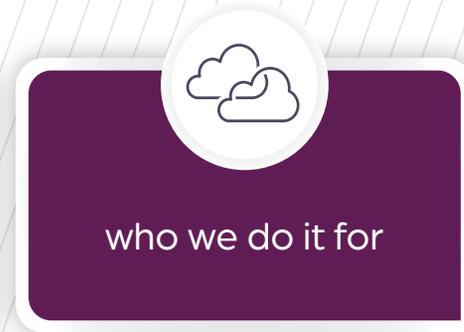
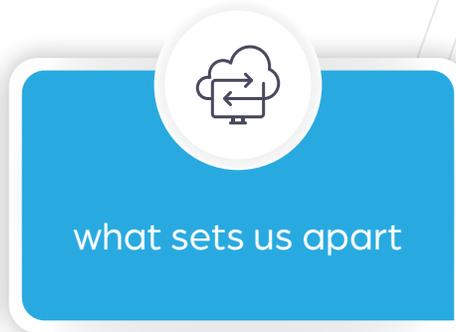
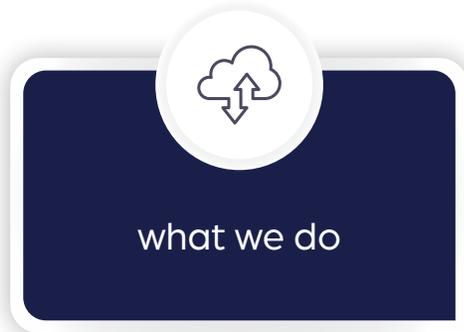




# Moving data and driving innovation

What we do, what sets us apart, and who we do it for.

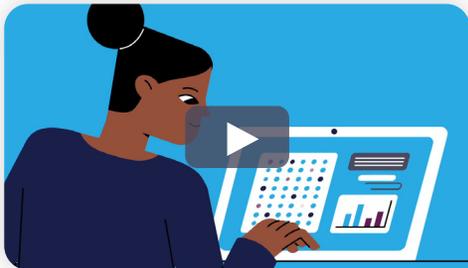
# Table of contents



click any button to go to that section

# What we do

Cirata's advanced data technology offers organizations a highly robust and efficient platform for seamlessly moving, transferring, migrating, and replicating information in data lakes. It does so without disruption to systems or applications, with full automation and at extreme scale and minimal delay. Our cutting-edge technology plays a pivotal role in empowering organizations engaged in advanced analytics, artificial intelligence, and other data-driven operations, enabling them to derive invaluable insights, accelerate innovation, and drive impactful outcomes.



Cirata Explainer video



About Cirata from Stephen Kelly, CEO



# What sets us apart



## The power of scale and immediacy

Organizations can effortlessly transfer and access vast amounts of data at scale, enabling them to leverage the most competitive technologies available. This seamless integration allows for prompt revenue increases, helping businesses maintain a strategic advantage in the market.



## Avoid downtime and business disruption

Seamlessly facilitate data transfer and migration projects, ensuring a smooth transition without any disruption to existing systems. With Cirata, applications can continue to operate flawlessly, allowing businesses to focus on their core operations while leveraging the power of advanced data management solutions.



## Combating data gravity

Empower organizations to overcome the challenges posed by data gravity. By enabling seamless access to and utilization of the latest technologies, our platform empowers businesses to leverage the advantages offered by multiple cloud providers and locations.

# What sets us apart

## Defining use cases



### Cloud migration and data modernization

Our solutions support organizations in their bid to modernize their data, making it more efficient by relocating it to cloud environments that provide better value.



### Continuous data transfer to the cloud

Our technology facilitates a continuous transfer of data to cloud environments, positioning datasets to be readily available for analytics and AI.



### Disaster prevention and recovery

Should a disaster occur, Our technology enables swift data recovery, minimizing potential data loss and ensuring the data's immediate availability.



### Hybrid or Multi-cloud

Our solutions aid in maintaining hybrid-cloud data environments, permitting data to span on-premises, cloud environments, or various locations within different cloud providers.

# What sets us apart

## Partnerships

Cirata has established fruitful partnerships with renowned companies such as *AWS, Microsoft Azure, Databricks, Google Cloud, IBM, Oracle, and Snowflake*.

These collaborations have not only enabled the firm to bring its cutting-edge technology to the market but also provided customers with a wide range of choices and unprecedented flexibility in meeting their unique needs. Through these strategic alliances, Cirata continues to drive innovation and deliver exceptional value to its clients, solidifying its position as a leader in the industry.



[read more ▶](#)



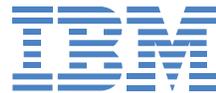
[read more ▶](#)



[read more ▶](#)



[read more ▶](#)



[read more ▶](#)



[read more ▶](#)

# Who we do it for

Cirata's solutions are trusted by hundreds of global brands and industry leaders such as

**Allianz** 



 **GoDaddy**

[read case study](#) ▶

**JUNIPER**  
NETWORKS

[read case study](#) ▶

 **Manulife**

 **NatWest**

 **Sanlam**

[read case study](#) ▶

# Who we do it for

## GoDaddy



### Challenge

GoDaddy utilizes an 800-node Apache Hadoop cluster to hold over 2.5 petabytes of customer-related activity and behavior data. This on-premises data lake is critical for guiding business operations and determining the company's investment strategies. The system is in operation 24x7. It can generate peak loads of more than 100,000 file system events per second, with sustained 12 hour periods processing an average of over 21,000 change operations every second. The challenge for GoDaddy was how to migrate petabytes of actively changing, "live" data when the business depends on the continued operation of applications in the cluster and access to its data.

Any disruption to business operations would be unacceptable and may have prevented a migration from even being attempted.

### Solution

GoDaddy used Cirata's Data Migrator to migrate data from their actively used cluster to AWS S3. Data Migrator performs a single scan of the source datasets and processes the ongoing changes that occur to achieve a complete and continuous data migration. It does not impose any cluster downtime or disruption to production applications and requires no changes to cluster operation or application behavior. Data Migrator enabled GoDaddy to perform their migration without disrupting business operation, and ensured that

datasets were transferred completely even while under active change in a very large and busy Hadoop environment.

### Results

- Using Data Migrator, GoDaddy achieved their initial migration goal—to migrate 500TB (over 8.6 million files) of the 2.5PB to AWS S3.
  - Completed the migration process while maintaining normal business operations at all times.
  - Reduced cost and risk of custom data migration development, enabling engineers to focus on other business critical tasks.
  - Established a new environment using AWS where GoDaddy plans to leverage AWS S3, EMR, Athena and other AWS services to achieve the following:
- Lower risk by moving off current aging hardware.
  - Meet SLAs for critical ETL processing requirements.
  - Create a better experience for their users through faster queries.
  - Greater agility by putting more data and flexible compute in the hands of data consumers.
  - Improved operational efficiency by alleviating the burden of managing the large and complex on-premises hardware and software infrastructure.

# Who we do it for

## GoDaddy testimonial



“At GoDaddy, deep technical knowledge is in our DNA, and we often build applications in-house to support growth. In the use case of a Hadoop to Amazon S3 data migration and replication, we found Cirata’s Data Migrator to be the optimal approach to deliver the best time to value, rather than running a more time-consuming and costly manual migration project internally.”

– Wayne Peacock, Chief Data and Analytics Officer, Godaddy

# Who we do it for

## Juniper



### Challenge

- Business continuity. Juniper wanted to operate on a 24-by-7 basis with no downtime, not even for maintenance.
- Provide remote locations with the same LAN-speed performance, repository access and availability as users at headquarters in Sunnyvale to ensure that remote developers are as productive as possible.
- Have a solution that is transparent to both users and administrators with automated failover and recovery.
- Have the ability to monitor and administer all servers worldwide from a single location.

- Juniper evaluated svnsync and other open source solutions, but couldn't find anything that would meet their requirements.

### Solution

- Juniper implemented Subversion WAN Clustering (SVN MultiSite Plus) with Subversion LAN Clustering at each site. Juniper now has a solution that delivers LAN-speed reads and writes at every location. All of the servers stay continuously in sync so that users at every site have access to the latest changes regardless of where they originated.
- When a server is taken offline, users automatically and transparently failover to another server and keep

on working. As soon as the server comes back online it resynchronizes automatically without administrators having to do anything.

- When an entire site is taken offline, users transparently failover to Subversion servers at another location and continue working. When the site comes back online, its servers resynchronize automatically with Subversion servers at other locations and the users are transparently switched back.
- The solution is transparent to users and administrators.

### Results

- Remote users have the same LAN-speed performance and read/write access as users at headquarters in Sunnyvale.
- 24-by-7 operation achieved with no downtime, not even for maintenance.
- Subversion servers stay continuously in sync so that users at every location access the latest changes.
- Provided continuous hot backup with automated fail over and recovery.
- Juniper can monitor and administer all of the servers from one location.
- Deployed to all servers worldwide within a couple of hours.

# Who we do it for

## Juniper testimonial

JUNIPER  
NETWORKS

”

“We really needed some way to ensure that all of our sites were up and running all of the time. We needed to be 24-by-7 globally with the same LANspeed performance and access at all three locations. Cirata’s (formerly WANdisco) active/active WAN clustering for Subversion was the only solution we found.”

— Angela Thomas, Development Tools Manager, Juniper Networks

# Who we do it for

Sanlam



## Challenge

Meeting the high availability objectives that Sanlam established for their CDP implementation meant that data in the DR environment needed to be kept consistent with that in production. Data is continuously being changed in production — either new data ingested, or existing data updated — so Sanlam required a solution that could replicate the data changes as they occurred in production to the DR environment.

Sanlam investigated several tools, such as Cloudera Replication Manager and DistCp (Distributed Copy) — an open-source tool provided with their Hadoop distribution. However, these tools do not support replication of changes as they occur. Instead, users must create schedules to periodically replicate data

incrementally. In the event of a failure, any changes made since the prior replication run will be lost, impacting both the RPO and RTO. Furthermore, these DistCp-based solutions run as standard MapReduce jobs competing for resources with other processes, which can impact production system performance.

Sanlam needed a solution that could replicate the actively-changing production data in near-real time without impacting performance so they could meet their high availability objectives.

## Solution

Following a thorough technical evaluation, Sanlam selected Cirata Data Migrator to automate the migration and ongoing replication of data between

their primary and secondary (DR) environments. Data Migrator is a safe and reliable cloud migration solution that automates the migration and replication of Hadoop data and Hive metadata to the cloud or between data centers. Data Migrator deployment is performed in minutes and requires no changes to applications or business operations. Migrations and replication of any scale can begin immediately and be performed while the source data is under active change, without requiring any production system downtime or business disruption. Data Migrator is the ideal solution for Sanlam and enables them to support near-zero recovery time objective (RTO) and recovery point objective (RPO), which were critical objectives for their business intelligence platforms.

## Results

- Original data migration of 70TB of data performed with no business disruption
- Ongoing replication of changes to production data as they occur (20TB per week)
- No disruption or impact to existing production environment
- Established DR environment that supports near-zero RTO and RPO
- DR environment provides additional capacity for more analytics, AI, and ML processing
- Data Migrator enables Sanlam to easily move data to other sources as well, such as for their future cloud requirements

# Who we do it for

## Sanlam testimonial



”

“We were impressed with Cirata’s unique capabilities for real-time replication without disrupting our production environments. It was critical to us that data in our disaster recovery cluster is kept in sync with our primary cluster as closely as possible to enable a near-zero recovery point objective and recovery time objective.”

— Jacques Joubert, Big Data Manager and Architect at Sanlam

Ready to modernize your data to fuel AI and analytics with Cirata?

talk to an expert

visit [www.cirata.com](http://www.cirata.com)

