



Aster *n*Cluster - Online Backup

Quick Overview

Aster *n*Cluster leverages its "Always Parallel" architecture to deliver high performance backup and restore capabilities required by terabytes to petabytes-scale frontline data warehouses. Using *n*Cluster Online Backup, data can be backed up without requiring a dedicated "backup window". *n*Cluster Online Backup includes powerful capabilities to support different backup policies of an IT department.

Highlights

- Massively parallel backup and recovery for high performance
- Online backup and recovery to eliminate downtime
- Backup storage on dedicated Backup servers or SAN/NAS
- Full and incremental backups
- Cluster-level and table-level backups
- Scheduled backups
- Compressed backups
- *n*Cluster Backup Terminal for managing backup and recovery tasks
- Low-cost disk-based backups using high-density Backup servers

Businesses have witnessed an exponential increase in data volumes in the last few years and continue to see them grow at a rapid pace. At the same time, today's organizations are adopting data-driven decision making strategies that place more stringent availability requirements on frontline data warehouses. Backup technology used in traditional databases is unable to match the multi-terabyte and petabyte scale commonly seen in today's systems. Database backups often don't complete in the dedicated "backup window" and cause significant performance slow-down, if done concurrently with queries. Similarly, recovery processes in traditional database systems have become unacceptably slow, given the large volumes of data that need to be recovered in case of a failure.

To meet the needs of today's frontline data warehousing applications, along with a scalable architecture for the database itself, there is a need for a new backup and recovery architecture that can scale to meet the data protection requirements.

*n*Cluster Online Backup is a part of "Always On" capabilities of the *n*Cluster database. It leverages *n*Cluster's "Always Parallel" architecture to provide powerful disk-based backup and recovery capabilities. Using massive parallelism, these capabilities can effectively handle the large data volumes and strict backup and recovery timelines that today's frontline data warehousing applications need.

Following aspects form the core of *n*Cluster Online Backup architecture:

1. Massive parallelization for providing backup and restore speeds required at the petabyte scale
2. Online operation so that backup and restores don't require system shutdown
3. Incremental backups so that the time taken for every backup does not depend on the amount of historical data
4. Compressed backups to reduce the cost of storage and improve network performance during data transfer

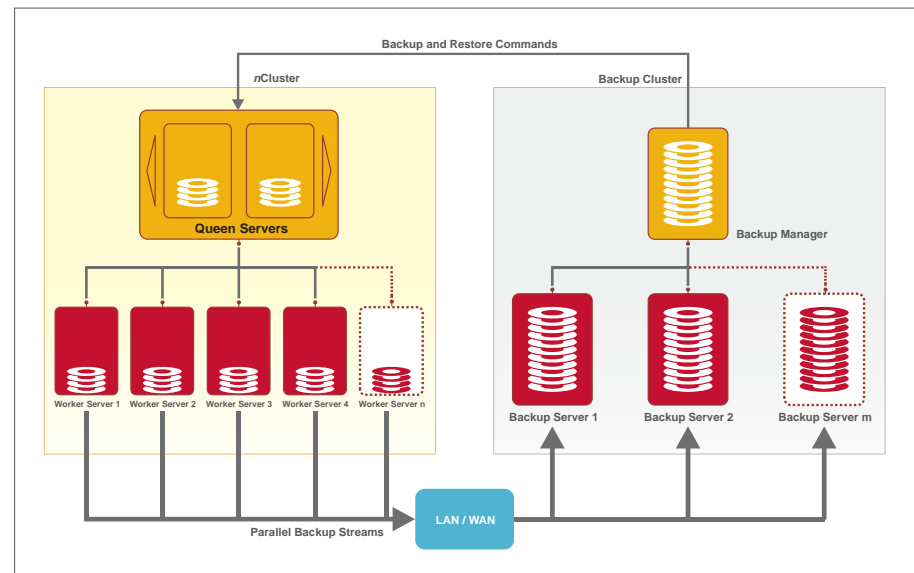


Figure 1: Aster *n*Cluster Online Backup

Backup Architecture

*n*Cluster Online Backup is based on a unique architecture that leverages a cluster of dedicated Backup servers, called Backup Cluster. The Backup Cluster architecture can leverage massive parallelism of Backup servers to provide high-performance backup and recovery. For disaster recovery, Backup Cluster can reside in a location which is geographically separated from *n*Cluster database. Multiple *n*Cluster databases can send data to a single Backup Cluster, providing an opportunity to consolidate backups and reduce costs. Data from *n*Cluster can be backed up to two types of storage targets:

1. **Backup Cluster:** The Backup servers themselves can store backups using their direct attached storage, with disk mirroring for enhanced data protection. Storage-heavy servers can provide high-density storage at a low cost per terabyte, an important consideration for data volumes associated with the frontline data warehousing applications. Backup Cluster is built on incremental scalability principles similar to the *n*Cluster database architecture – more servers can be incrementally provisioned to add backup capacity. Thus, backup storage costs can be managed in a granular manner as data volumes grow. From Backup servers, backup files can be moved to tapes or VTLs, if so required by the IT policies of an organization.
2. **Network Storage:** *n*Cluster Online Backup also provides the flexibility to store backup on network storage (SAN/NAS), if required by an organization's IT policies. If this option is chosen, Backup servers can run backup and restore processes in a massively parallel manner while using the network storage to store backup data.

Backup and Recovery Process

To manage backup and recovery tasks, one of the servers in Backup Cluster is designated as Backup Manager. Backup Manager runs *n*Cluster Backup Terminal software (and also stores data, like other Backup servers), which provides a functionally rich interface for backup and restore tasks. To start a backup, Backup Manager first contacts the Queen server in *n*Cluster database, which, in turn, sends the backup request to the Worker servers. Next, Worker servers directly connect to the Backup servers and stream data in a massively parallel manner. Backup Manager and the Queen server do not get involved in the actual data transfer so that they do not become performance bottlenecks. It should be noted that the number of Backup servers is not determined by the number of Worker servers, as the data transfer is done in a many-to-many fashion.

Restoration of backup data can also be initiated using *n*Cluster Backup Terminal. During restoration, data is streamed in a massively parallel manner directly from the Backup servers to the Worker servers, providing faster time to recovery. For recovery, both online and offline modes are supported, depending on the parameters selected at the time of backup. Using table-level online recovery, data can be recovered while the system is available for business.

*n*Cluster Online Backup also provides disaster recovery capabilities – full backups can be restored to a secondary site that's running *n*Cluster database software. Disaster recovery in *n*Cluster is fast, as large volumes of data can be quickly transferred by leveraging massive parallelization and backup compression.

Types of Backup

*n*Cluster provides the flexibility to manage backups according to an IT organization's administrative policies. Backups can be made at the level of *n*Cluster database or at

the level of individual tables. *n*Cluster Online Backup supports both full and incremental backups of data. For table-level backup, full backup is currently supported.

A full backup includes all data, metadata, log files, etc. Incremental backup includes only the changes since last backup, minimizing data transferred over the network as well as backup storage requirements.

Scheduled Backups

*n*Cluster Online Backup includes the ability to run scheduled backups according to pre-defined backup policies, minimizing manual intervention for routine backups. Administrators can schedule backups in any of the following three ways:

1. Using the scheduler provided with *n*Cluster backup utility
2. Using an OS scheduler (e.g. CRON jobs)
3. Using third-party schedulers

Backup Compression

*n*Cluster Backup Compression provides powerful compression capabilities. Data can be compressed by the Worker servers before it is sent to the Backup servers. This minimizes backup storage requirements as well as data transfer over the network, speeding up both backup and recovery processes. Compression ratios in the range of 3X to 12X are typical, depending on the nature of data being backed up.

Backup and Recovery Performance

The unique architecture of *n*Cluster Online Backup enables very high levels of performance. As data is streamed from a large number of Workers, tens, hundreds or even thousands of processor cores can send data to Backup Cluster in parallel. On the receiving end, Backup servers can process data in parallel, maximizing throughput. Similarly, the recovery process also benefits from massive parallelism, providing high performance and shorter time to recovery. Compression of backups improves network performance during backup as well as recovery processes, providing further

speed-up. Such high-performance architecture of *n*Cluster Online Backup enables organizations to protect data efficiently even at the petabyte scale.

Backup Management

*n*Cluster Online Backup provides a rich set of backup and restore functionality through *n*Cluster Backup Terminal. Administrators can manage the following tasks using *n*Cluster Backup Terminal:

- Adding and removing Backup servers from Backup Cluster for capacity and performance management
- Starting, pausing, resuming and cancelling backup and restore jobs
- Monitoring the status of ongoing backup and restore jobs
- Scheduling backup jobs
- Deleting backups
- Viewing the history of backups and restores
- Monitoring total and used backup storage capacity (at Backup servers or SAN/NAS, as the case may be)

Cost of Backup

For high-volume data warehouses of today, cost of managing backups can be high. *n*Cluster Online Backup helps organizations reduce the cost of backup in the following manner:

Use of inexpensive commodity servers for backup storage

*n*Cluster Backup Cluster can use inexpensive commodity servers with direct attached storage to minimize the cost of backups. Organizations can select cost-effective server configurations according to their backup requirements. For example, storage-heavy commodity servers with 48 SATA disks (500 GB to 1TB per disk) can provide low cost per terabyte for storing back-ups. In addition, Backup Cluster can be incrementally scaled-out to increase storage capacity at low costs, according to the rate of data growth. This helps organizations manage their storage costs in small incremental steps.

Backup Compression

As mentioned above, *n*Cluster Backup Compression can provide significant reduction in backup storage requirements and also help improve backup and restore performance by minimizing the network traffic.

About Aster Data Systems

Aster Data Systems is a proven innovator in massively parallel processing (MPP) databases for data warehousing and analytics – bringing deep insights on massive data analyzed on clusters of commodity hardware. Co-founded in 2005 by three colleagues in the Stanford Computer Science Ph.D. program, Aster's *n*Cluster database provides patent-pending innovations in scalability, manageability, availability, and analytics. Aster is headquartered in Redwood City, California and is backed by Sequoia Capital, Cambrian Ventures, and First-Round Capital.

