


# Secure Network Analytics Data Store

## How well are you managing your network data?

Many organizations face significant challenges when it comes to effectively managing the collection and storage of their network telemetry in an efficient and scalable manner. This is especially pronounced for large enterprises and service providers with massive network footprints and exceptionally high flow per second (FPS) volumes, as they are faced with problems related to ingestion bandwidth, query performance, long-term data retention, and data resiliency. As a result, practitioners at these firms are often forced to implement unconventional workarounds that come with undesirable tradeoffs to satisfy their network telemetry and data storage needs.

Large organizations need a solution that provides scalable network telemetry collection and storage, highly responsive query times, and reliable data resiliency as core capabilities. The Stealthwatch Data Store provides an improved database architecture for solving these problems by decoupling ingest and data storage enabling new ways of efficiently managing data.



### Common network telemetry collection, storage, and analysis challenges

- **Ingestion:** Organizations with large or expanding network footprints face scalability challenges and increased expenses as they must continuously purchase additional sensors or Flow Collectors to handle continuously growing ingest volumes
- **Query performance:** For large enterprises, the task of running queries on large data sets is incredibly computationally expensive and can take upwards of 24 hours – this leads to operational inefficiencies by slowing down remediation efforts and draining finite computational bandwidth
- **Data retention:** Many organizations are unable to retain the amount of network telemetry that they need to fulfill compliance requirements, forcing them to either purchase expensive third-party storage solutions or free up room in their proprietary databases to avoid legal risks should they be audited
- **Data resiliency:** Organizations that lack sufficient backup storage capacity are at risk of losing valuable data if one of their critical backup data storage systems fails

## Required solution components

### Data Store (DS 6200):

- Secure Network Analytics Manager 2210
- Flow Collector 4210
- Data Store 6200

### Virtual Data Store:


- Virtual Secure Network Analytics Manager
- Virtual Flow Collector
- Virtual Data Store (L-ST-DS-VE-K9)

### Additional notes:

- The Data Store will be supported by Secure Network Analytics versions 7.3 and above
- Both the Data Store 6200 and the Virtual Data Store consist of a set of three Data Node appliances
- Data Node appliances are not sold separately

## How it works

The Data Store cluster sits between the Stealthwatch Management Console and Flow Collectors. One or more Flow Collectors ingest and de-duplicate flow data, perform analyses, and then send the flow data and its results directly to the Data Store. This flow data is then distributed equally across a Data Store, which is comprised of a minimum of three Data Node appliances. The Data Store facilitates flow data storage and keeps all your network telemetry in one centralized location as opposed to having it spread across multiple Flow Collectors in a distributed model. This new centralized model offers greater storage capacity, flow rate ingestion, and increased resiliency versus the distributed model.



The illustration above depicts the components and architecture of the Data Store solution. Similarly, to the current deployment model, Flow Collectors leverage enterprise telemetry to ingest NetFlow. However, unlike the standard model, the Data Store offers a new design in which ingest and data storage functions are performed independently from one another.

Organizations can purchase either a single hardware Data Store 6200 or a Virtual Data Store, or can scale in terms of FPS and data retention by adding multiple Data Stores together as a single database cluster. The table below shows typical ingestion rates and retention numbers for the DS 6200:

Flow Rate (FPS)	Flow Collector 4210	Data Store 6200 for 90 days storage	Data Store 6200 for 180 days storage	Data Store 6200 for 360 days storage
250K	1	1	1	2
500K	1	1	2	4
1M	2	2	4	8
2M	4	4	8	-
3M	6	6	12	-

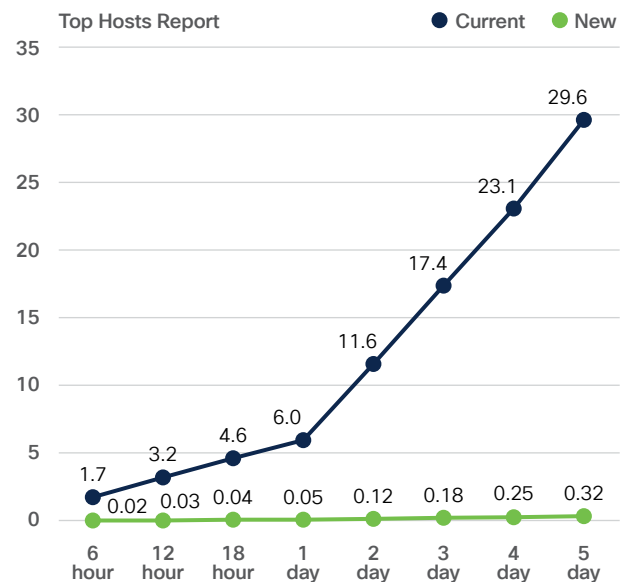
The number of Data Stores required can be customized based on an organization's traffic profile (density of traffic, number of hosts, etc.) to satisfy their specific ingestion and retention requirements.

## Key Benefits

**Increased data ingest capacity:** Data Stores can be combined to create a single cluster of data nodes capable of monitoring over 3 million flows per second (FPS) to aid in relieving ingestion bandwidth challenges for organizations with high flow volumes.

**Enterprise-class data resiliency:** Telemetry data is stored redundantly across nodes to allow for seamless data availability during single node failures helping to ensure against loss of telemetry data. Deployments with 2 Data Stores or more can support up to 50% of data node loss and continue to operate\*. The Data Store also supports redundant inter-connection switches to remain fully operational during network upgrades and unplanned outages.

**Query and reporting response times improved by a significant magnitude:** The Data Store provides drastically improved query performance and reporting response times of at least 10x faster than those offered by the distributed deployment model. It can also perform an increased number of concurrent queries whether through APIs or the Stealthwatch Management Console web UI. These query improvements stand to deliver substantial operational efficiency gains. Through the ability to run reports and get answers more quickly, the Data Store enables practitioners to pinpoint and respond to threats more quickly to expedite triage, investigation, and remediation workflows.



Internal testing in both lab and production environments demonstrated a dramatic improvement in reporting response times with top host reports that prior took 29 hours being reduced to just 19 minutes.

\*Depending on your hardware configuration and installation.

## Key Benefits

**Storage scalability:** The Data Store offers organizations with growing networks enhanced flexibility around data storage scalability through the ability to add additional DS 6200s to grow their database cluster.

**Long-term data retention:** Scalable and long-term telemetry storage capabilities enable long-term flow retention of up to 1-2 years' worth of data with no need to add additional Flow Collectors offers the additional benefits:

- **Ability to satisfy regulatory requirements:** The ability to store everything on your network for longer periods of time, aids the fulfillment of regulatory and compliance requirements related to data retention.
- **Reduced complexity and cost savings:** The Data Store's scalable model reduces complexity by eliminating the need to cobble non-integrated storage solutions together to retain data. This also results in cost savings by removing the need to purchase expensive third-party storage solutions or extra Flow Collectors.
- **Larger data sets available to perform investigations on:** With 1-2 years' worth of telemetry stored, all of your NetFlow data is at your fingertips and available for you to perform queries on from the Stealthwatch Management Console.



---

Learn more about the Stealthwatch Data Store [here](#).