



Keeping Up with Cloud > Special Edition Dynamic Data: Self-Storage and Dynamic Data: Active Archive

*This Special Edition of “Keeping Up with Cloud” focuses on two Splunk Cloud features that warrant a bit more of the spotlight and a quick refresher. We hope this targeted edition helps you better understand and more effectively position **Dynamic Data – Self-Storage and Active Archive**.*

TL;DR? Here’s Your Executive Summary of Sorts

When selling Splunk Cloud, always understand your customers’ storage requirements, and be aware of upsell opportunities. We continue to deliver active storage as part of the Splunk Cloud service – this is the assurance our customers can easily access and investigate the data that is most important to them.

With Splunk Cloud release 7.1 – 7.2, we introduced additional storage options – the Dynamic Data series of services. **Dynamic Data: Self-Storage (DDSS)** is a feature included with the Splunk Cloud subscription. DDSS empowers customers to move their data from Splunk Cloud to their own Amazon S3 storage environment. This becomes a permanent data storage, required to comply with regulatory standards but no longer accessible by Splunk.

Dynamic Data: Active Archive (DDAA) is an additional service enabling Splunk Cloud customers to cost-effectively retain their less frequently accessed data while maintaining the ability to search on this data using Splunk Cloud. DDAA is a great upsell for all your deals – especially security opportunities. It should become a default component of your sales motion.

Refresher: Let’s Talk Dynamic Data: Self Storage (DDSS)

With the release of Splunk Cloud and Splunk Enterprise 7.1 in April 2018, we announced a new Splunk Cloud data storage service, Dynamic Data: Self-Storage (DDSS). This service empowers customers to move data from Splunk Cloud to their own Amazon S3 environment. Splunk designed this service to help Splunk Cloud customers navigate new compliance regulations and internal mandates that require organizations retain their data for significantly longer periods of time.

“Splunk Cloud offers customers maximum control and real-time access to their data. Splunk Cloud release 7.1 features Dynamic Data: Self-Storage, arming customers with the flexibility to move data from Splunk to their own Amazon S3 storage environment.” – [From the April 2018 Press Release](#)

With DDSS, customers can send data to Splunk Cloud, and control the retention of this data – as usual – on a per-index basis. When a customer’s retention thresholds are met, and data is about to be deleted, indexes configured with DDSS move the data to an Amazon S3 bucket in their organization’s AWS account. Customers are empowered to keep this data for as long as they determine is necessary.

Taking a Deeper Look at DDSS

Last year, our own **Sundeep Gupta**, Director, Cloud Product Management, wrote a blog post entitled, “[Dynamic Data: Self Storage, Compliance, Cloud, and Data Lifecycle](#)” to describe our approach to providing customers with flexible data storage options.

In his blog, Sundeep describes the defining principles of DDSS and takes a deeper look at this service – detailing how data can be moved from a configured index to a defined self-storage location.

[Check out the steps here.](#)

In addition, Sundeep's blog explains how this feature can support customer compliance requirements and search expectations.

“What about when a compliance initiative requires you to search against this data? No problem! You can spin up an instance of Splunk Enterprise and feed it this data. Because Splunk has already indexed this data, it bypasses the licensing metering, meaning you don't need to pay anything to index and search this data,” explains Sundeep, “(DDSS) is designed to move data outside of Splunk Cloud and into your own data storage; an area Splunk does not have access to.”



POSTED BY
Sundeep Gupta



Sundeep leads the product management function for Splunk Cloud. In this role, he is responsible for the definition of Splunk Cloud's evolving cloud-based capabilities. Sundeep is passionate in his belief of applying a structured, data-driven approach to managing professional and personal decisions. As a Splunk customer at two prior companies, he has deep conviction in Splunk's ability to be transformative. Sundeep has degrees in Electrical Engineering and Technology Management.

Andy Hershey Provides a Sales Perspective

When positioning DDSS, **Andy Hershey**, AVP, Global Cloud Sales, offers the following advice:

“Always keep in mind that we have different storage options based on customer needs or their expected needs. For example, customers who need data available to search quickly and easily are well-suited for standard storage – which we've always offered through Splunk Cloud. Perhaps, you have a customer that might be required at some point to export data needed by an auditor. We have the Data Export feature for this – offered at no additional cost with Splunk Cloud. Finally, there are those customers looking for a less expensive way of searching data on occasion. This is where Data Archiving works.”

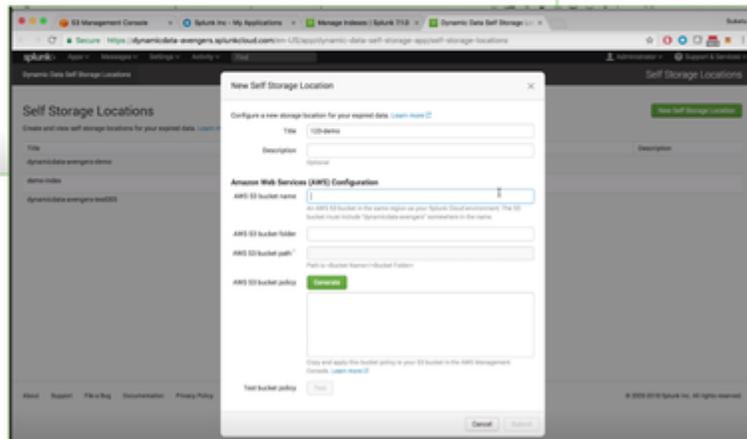
Andy also encourages sellers to stay away from selling “storage” and instead provide options around Dynamic Data – since it differentiates Splunk from pure storage to a way to flexibly manage their data based on need.

Get Technical with a Video Walk-Through

If you want a more technical perspective of DDSS, check out the [Splunk Enterprise 7.1 - Technical Enablement Feature Video - Dynamic Data Self Storage video](#). It starts with an overview and goes into a demo walk-through on how the features work.

Dynamic Data Self Storage Overview

- ▶ Cost effective long term data retention for regulatory compliance and advanced analytics use cases
- ▶ Export data from indexes in Splunk Cloud to AWS S3 buckets
- ▶ AWS S3 bucket needs to be in the same region like customers Splunk Cloud environment
- ▶ Available for managed Splunk Cloud instances
- ▶ Customers can restore data to a Splunk Enterprise instance



Refresher II: Now Let's Talk Dynamic Data: Active Archive

With the release of Splunk Cloud and Splunk Enterprise 7.2 in October 2018, we announced the latest release in the Dynamic Data service series called Active Archive. Dynamic Data: Active Archive (DDAA) enables Splunk Cloud customers to meet regulatory and compliance requirements by cost-effectively retaining their less-frequently accessed data. Customers can specify which data to send to their Active Archive, and then they retain the ability to search on this data with Splunk Cloud.

“In a heightened regulatory and compliance environment, the breakthroughs in the new Splunk Cloud enables [Cox Automotive] to cost-effectively meet new business requirements. The ability to quickly recall data for audits and security investigations saves a huge amount of time gathering data, so we can spend more time investigating and taking action on it.” – [From the October 2018 Press Release](#)

DDAA is priced at \$240 per storage unit. This represents a 60% cost reduction compared to actively searchable data (“warm data”) in Splunk Cloud which is priced at \$600 per storage unit.

It's different than DDSS in two key ways:

- DDAA enables customers to restore data from archive and then search on this restored data with Splunk Cloud (i.e., the restored data becomes “warm”)
- Splunk manages data storage with DDAA, whereas with DDSS customers set up and manage data storage in their own Amazon S3 account. Also, once moved into storage, this data is no longer accessible by Splunk Cloud (i.e., the data becomes “cold.”)

Dynamic Data: Data Retention Options in Splunk Cloud

Customers can learn even more about data management options with Splunk Cloud in this blog post. The author, who is no longer with Splunk, writes, "It's all about giving our customers flexibility and choice. In talking with customers, it became very clear that they wanted choices in how their data is managed in Splunk Cloud. DDSS is designed to provide a path for customers to self-manage data as it ages out. Customers are then responsible for restoring this data should the need arise to search against it. With DDAA, on the other hand, Splunk manages the archival and restoration of data."

Below are the two key differences between DDAA and DDSS:

- 1. Data Management:** With **DDAA**, Splunk provides complete lifecycle management of the archive on your behalf and remains the custodian of your data. Just like your Active Searchable data, Splunk manages all aspects of archive availability, durability, security and privacy requirements on your behalf. With **DDSS**, you the customer are responsible for data once it ages out. You define an Amazon S3 self-storage location and decide data from which indexes lands there. Once the data lands in self-storage, you're in complete control.
- 2. Data Restore:** **DDAA** enables you to request a slice of your data to be restored back into your Splunk Cloud instance. The entire workflow is fully integrated into Splunk Web so your archived data is available at your fingertips with predictable time between retrieval to search. With **DDSS**, if you want to search against the data stored in a self-storage location, you will need to restore it to a separate Splunk instance.

Read the [full blog post](#) to get a deeper dive on how DDAA works.

Video: Data Retention in Splunk Cloud

As you've already learned, in Splunk Cloud, customer data has a dynamic lifecycle that goes through different phases: Searchable, Archive and Self-Storage. As such, every Splunk Cloud subscription includes searchable storage that runs at the speed of the business and is immediately available for analytics workloads.

Watch this customer-facing video to learn how to enable retention rules with Splunk Dynamic Data: Active Archive. This process ensures your customers' data lifecycle matches their data retention needs, while ensuring they remain compliant.



Calculating Dynamic Data: Active Archive Capacity

Is Dynamic Data Active Archive (DDAA) a fit for your customer's long-term data retention requirements? Then, you'll need to calculate how much DDAA capacity they will need. Here is a quick example:

If the customer has a one-year data retention requirement and the quote is either one year or multi-year, the math is relatively straightforward.

- The customer has XGB of data ingest per day, then the retention requirement is 365 days x (365x GB). 90X GB is accounted for using the searchable storage included in the Splunk Cloud base subscription, so your DDAA quote needs to account for 275X GB.

- Since the DDAA SKU is sold in 500GB increments, your quote should include 275X / 500 units of DDAA for every annual subscription period.

Also note there are different DDAA SKUs for encryption at rest (EAR) and GovCloud. Therefore, if your customer's Splunk Cloud base subscription includes EAR, then choose the DDAA EAR SKU for the quote. If your customer's Splunk Cloud is within the GovCloud region, then choose the DDAA GovCloud SKU.

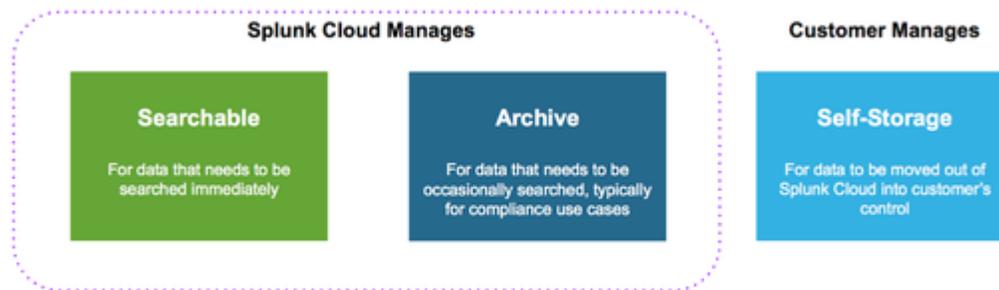
If your customer has a more complex data retention requirement, go ahead and post your questions to Cloud Selling Chatter.

Here's an Overview of the Splunk Cloud Data Storage Options

Sundeep Gupta, in his internal blog about Splunk Cloud Data Storage Options, says "Our goal is to increase the overall support for managing customer's data lifecycle in Splunk Cloud."

"When customers ask you questions on storage options, be sure to ask your customer appropriate questions around the descriptions below to ensure we understand exactly what type of storage is appropriate."

[Click here to get a side-by-side comparison](#) of the Dynamic Data Storage Options offer through Splunk Cloud.



Splunk Cloud Customer Presentation with Dynamic Data Slide

This presentation introduces Splunk Cloud. The main deck (slides 1-13) covers the Splunk Cloud value proposition. The appendix includes additional slides which are intended to be used to build custom presentations. Appendix slides cover the following topics: total cost of ownership, technical overview (including latest Splunk Cloud features), customer evidence, and readiness.

This content is designed for prospect customers, including new logos and also Splunk Enterprise customers who are considering migrating to Splunk Cloud. Target personas include Splunk administrators, as well as IT and security professionals.

Slide 38 focuses on Dynamic Data. Feel free to use this slide in your other customer presentations.

© 2019 SPLUNK INC.

splunk > cloud

The benefits of Splunk deployed and managed as a secure, reliable and scalable service



Service Excellence Fast and Nimble Saves Time and Money

splunk > listen to your data

Sales Resources - Your One-Stop-Shop



Field Enablement

- [Cloud Selling Chatter](#) (Find and solicit answers to all Cloud Q's)
- [LAER for Cloud](#) (Cloud Sales Methodology)
- [Field Enablement Cloud Homepage](#) (Home for all Cloud Resources)
- [Cloud Deal Structuring Guide](#) (Structuring & Pricing Cloud Deals)
- [Splunk Cloud Terms of Service](#)
- [Splunk Cloud Service Definition](#)
- [Splunk Cloud FAQ](#)

Thanks for reading! If you have questions or need more support, please contact any one of us on the Splunk Cloud Sales Leadership team via email.

 Like (0)