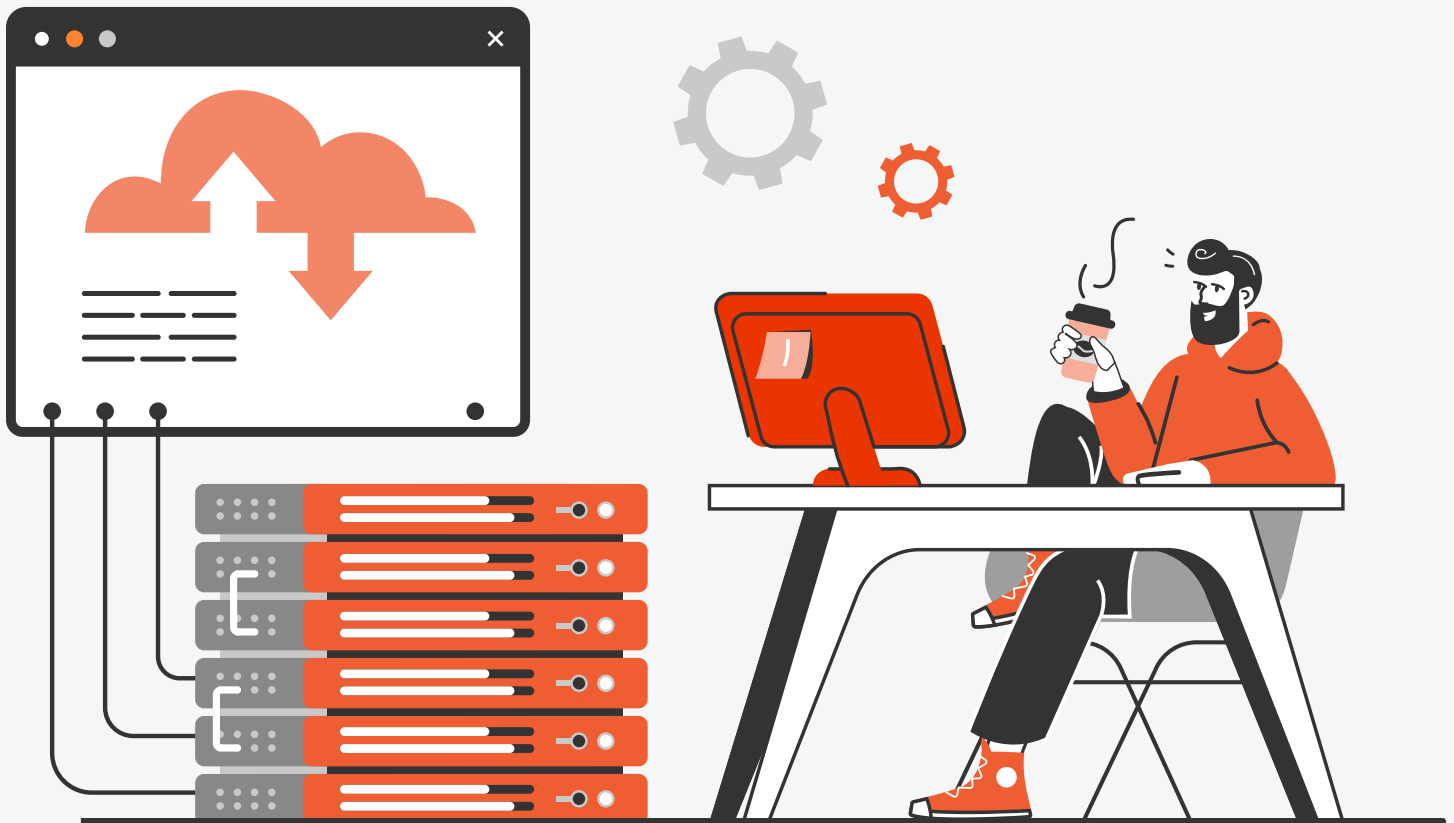


Why SearchStax Cloud is Better than DIY

Get Highly-Available, Fast and Scalable Solr Infrastructure



Introduction

For search-based application development, there has typically been two options for your Solr infrastructure: Build it yourself or buy a completed solution from a third party. Now there is a third option - purchasing a managed software-as-a-service solution that takes the time and pain out of deploying, managing and scaling your Solr infrastructure.

This new way to build search applications provides all of the advantages of building your own along with many of the benefits of buying. Your developers' time is one of your most valuable assets and their time should be focused on their core competency of building better applications and focusing on the search experience rather than dealing with provisioning or maintaining the infrastructure.. We couldn't agree more. [SearchStax Cloud](#) is a fully-managed Solr service to help developers get their time back by providing an easy way for them to automate, manage, and scale high-availability Solr in public or private clouds.

Cloud - You're in control. Use yours OR ours

For most modern applications, you'll want to use the cloud for cost savings and disposability of virtual machines. The SearchStax managed Solr offering is available in over 60 regions on AWS, Microsoft Azure and Google Cloud. In addition, our solution is flexible enough to allow usage using your cloud account or ours. If you have a requirement to use your AWS cloud infrastructure due to PII or other security or business reasons, no problem. If you want to use SearchStax's AWS cloud, no problem there either. You're always in control.

Step-by-Step Comparison

Now, let's take a look at a step by step comparison between doing it yourself and using SearchStax Cloud. We'll review the typical projects or activities you'll need to take on as part of setting up and managing a Solr infrastructure. In addition, we'll detail the steps and approximate a level of effort for each one. We'll break the activities into two sections: initial set up and ongoing support. Let's get started!



Initial Set Up

Building a Cluster

Now that you are ready to start setting up your search infrastructure, let's build out the initial development cluster. As you can see from the following comparison, by using SearchStax you have complete control over the specific parameters of your cluster, but at a time and effort savings. Without SearchStax, standing up a cluster can range from a couple of days to weeks depending on the complexity of the application, organization and how many hiccups pop up. With SearchStax, that timeframe is shrunk down to minutes. See the below table for a detailed comparison. The comparison below looks at the time it takes to build one cluster. Typically, teams want to have different environments like development, qa, staging, demo, production and disaster recovery - so the do-it-yourself route can be even longer.

With SearchStax: 10-30 Minutes

1. Select the right plan and [create a new deployment](#) using SearchStax Cloud
2. Go grab a cup of Cappuccino

Do It Yourself: 2-10 Days

1. Create Security groups and do SSH key management
2. Choose and set up proper VMs
3. Set up EBS and partitions appropriately for logs and data directories
4. Upgrade all available base packages to latest versions to avoid known security vulnerabilities
5. Download and Install JDK
6. Download and Set up Zookeeper Ensemble on 3 or more nodes
7. Download Solr on all nodes and configure a Zookeeper Cluster
8. Get SSL Keys, Certs and copy over the the various machines; create java keystores
9. Set up SSL on each node
10. Set up an Elastic Load Balancer along with SSL Certificates
11. Create DNS records for each machine along with DNS records for load balancers
12. Configure Load balancer for the nodes that are running Solr



Monitoring and Alerting

The DIY engineer's or devops job is still not done. Once your cluster is built and running the way you want it, you'll need to implement a monitoring system. You need to ensure that you are aware of the health of not only your servers and cluster as a whole, but also individual metrics that may have an adverse impact on your end-user search performance or stability of your cluster.

With SearchStax: 15 Minutes

1. [Performance Monitoring and Alerting](#) is a core part of the SearchStax Cloud product.
2. Choose the metrics and thresholds that you want the alerts to be invoked on.
3. Enter your email address where you want to receive alerts in the environment dashboard.

Do It Yourself: 3-20 Days

1. Make a choice for system monitoring. There are many tools out there: commercial, open source, brand new or long-standing and battle-tested. Your choice.
2. Customize the tool to capture search throughput, latency and other metrics
3. Customize the tool to capture indexing throughput, latency, pending documents, commit frequency, etc
4. Enable or Customize the tool to capture garbage collection level metrics
5. Customize the tool to capture caching metrics warm-up time, cache hit ratios, cache evictions, etc.
6. Customize the tool to capture zookeeper related metrics - queue sizes, connections, etc.
7. Configure specific widgets to display graphs for specific / new metrics
8. Deploy & configure the tool on each Solr Server
9. Deploy and configure the tool on each zookeeper instance
10. Plan out and list the thresholds and alerts that you need to set up
11. Integrate into an alerting tool for sending email alerts



Log Management

Most search applications will create a large log file pretty quickly. Without proper log management in place, this can quickly fill your drive. In addition, to get easy access to those logs, you either have to put something in place for your developers to request access to the logs or develop a user access management system so that they can log in to the servers and get the logs.

With SearchStax: 0

1. [Log Management](#) is included with SearchStax Cloud. Moreover, you have access to your logs from any of your instances right from the dashboard. In addition, these logs are archived and stored in scalable storage on your cloud provider.

Do It Yourself: 12–48 Hours

1. Identify each log on the system that needs to be rotated by logrotate.d
2. Build an agent that pushes these logs to a centralized repository (e.g. S3)
3. Create configuration files in/etc/logrotate.d for each of them
4. Restart logrotate.d



Backup Set up and Test

Regularly backing up your application is a best practice in our industry. In case of data corruption, unexpected hardware malfunction or accidental situations, you can go back to your previously stored backups and restore your entire application quickly. For search systems with large amount of data, re-indexing of your entire index can take many hours if not up to days. And sometimes, reindexing is not an option. When putting a backup strategy in place, you need to think about frequency of backups, develop scripts for performing backups, identifying and planning for redundant storage (where backups can be safely stored), and a process for restoring and validating backups. Learn more about [SearchStax Cloud backups](#).

With SearchStax: 5-10 Minutes

1. Daily Backups are included in our Gold, Platinum and Platinum Plus plans. You can create one-time backups or create a custom schedule on an hourly, daily or weekly basis.

Do It Yourself: 2-5 Days

1. Identify storage location for keeping backups
2. Develop a script that that copies the backup of all the index directories along with Zookeeper configuration files
3. Test the backup script to ensure backups that are created are valid
4. Restore the backup to ensure no file has been missed
5. Set up a cron or scheduled process to run periodically
6. Change the cron file to send an email if backups fail
7. Create a script that deletes old backups beyond a certain date / day (depending upon your backup retention policy)
8. Test all of the above



Building a Disaster Recovery (DR) Plan and Environment

We all hope disasters don't happen, but smart companies are better prepared for it. Whether that's because of ensuring zero downtime for your customers or SLA guarantees or compliance, a disaster recovery site in a separate data center and/or location is advised.

Setting up a disaster recovery site involves setting up a similar cluster, and the ability to test and ensure the DR plan is in place in case of a catastrophic event. Let's look at what it takes to set up a DR site, develop scripts for performing backups, identifying and planning for redundant storage (where backups can be safely stored), and a process for restoring and validating backups.

With SearchStax: 0

Gold and Platinum Support Plans

Cold, Warm or Hot Disaster Recovery is available as an add-on for our Gold and Platinum Plans. Learn more about [SearchStax Cloud Disaster Recovery options](#).

Platinum Plus Support Plans

Disaster Recovery is automatically included with our Platinum Plus plans. Let our Customer Success onboarding experts identify your needs and set up disaster recovery that meets your business needs.

Do It Yourself: 2-10 Days

1. All the steps needed to build a cluster previously
2. If using Solr 6.x, set up cross datacenter replication
3. Set up failover mechanisms to ensure either manually or automatically traffic is moved to the DR cluster in case of a need
4. Set up monitoring and alerting for the DR cluster
5. Set up Log management and access for the DR cluster



Ongoing Support

Now that you have your initial set up and deployment completed, let's take a look at ongoing management and support activities.

Scaling

Any application that's built to be successful is eventually going to need to scale. Whether that's because of more data that you want to bring into your Solr deployments or if it's more query traffic or use cases that start to put more query load on your clusters.

Whether that's vertical scaling or horizontal scaling, you'll have to ensure that the process of scaling is designed in a way that you can easily scale up or down your clusters. More so, you'd like not to have any downtime or performance impact to your end-users while this activity is being performed.

In certain use cases where demand is seasonal or unpredictable, it becomes even more critical for you to optimize the utilization of your resources and not have too much extra capacity lying around which will increase your costs. In typical installations, we see enterprises scaling up or down 2-3 times a year. That can be attributed to seasonal changes, user adoption / growth or the ever increasing data. With today's applications, not all that information is available upfront, so the system needs to be able to adapt to changing application needs quickly. Typically, teams want to have different environments like development, qa, staging, demo, production and disaster recovery - so the do-it-yourself route can be even longer.

With SearchStax: 0-30 Minutes
depending on your plan

Gold Support Plans

To Scale Up:

1. Go to the SearchStax Cloud dashboard and click on "Add Node" and specify the number of additional nodes you want to add
2. Execute Solr APIs to create new replicas or split shards, etc.

To Scale Down:

1. Go to the SearchStax Cloud dashboard and click on "Delete Node" Platinum Support Plans

Platinum Support Plans

Have margaritas on the beach and let a SearchStax Expert Support Engineer scale your clusters up or down depending upon the need

Do It Yourself: 2-4 Days per occurrence
3x/year=6-12 days/year

To Scale Up:

1. As per your needs, if you were to double the capacity of your cluster, you'll have to add another set up of servers and perform all the activities that you performed in Cluster Set up except the load balancer activities
2. Set up monitoring & alerting for these set of servers following the same steps that is within the Monitoring section
3. Execute Solr APIs to create new replicas or split shards, etc. This may take you a 12-24 hours, or more if you run into surprises or hurdles

To Scale Down:

1. Take the node out of the load balancer pool
2. Tear down the node that is not an issue
3. Remove all monitoring and alerting for these nodes that were put in the monitoring tool



Upgrades

Below are some common reasons why software upgrades in general are important for your application:

- Upgrades introduce improvements and new features that continually increase your efficiency and time to market
- Upgrades address/correct bugs encountered in prior versions
- Support is geared to address the most common, most recent issues faced by the majority of customers
- Avoid being stuck with obsolete versions

Apache Solr version upgrades consist of major and minor version upgrades. Major Solr version upgrades can contain changes that are not backward-compatible with existing applications. However, minor version upgrades can be generally applied without having an impact on your application.

Let's look at what's involved in performing an upgrade of your Solr deployment:

With SearchStax: 0

1. SearchStax will automatically upgrade or patch your Solr version as needed
2. With our Platinum Plus plan, you will receive priority patching for Solr
3. Tell your boss it's taken care of and you rock!

Do It Yourself: 2-4 Days

2x/year=4-8 days/year

1. Identify the version you plan to upgrade to
2. Take a backup of your existing Solr Index to ensure you can restore in case of any issues
3. Set up a Solr instance with the version you want to upgrade to
4. Take the index from one of the machines within the cluster and deploy to the new instance
5. Perform an in-place upgrade and verify if there are any changes that need to be made
6. Once verified, set up a new cluster take one machine out of the cluster pool and perform an upgrade
7. Add that machine back to the pool and ensure no errors are happening
8. In certain cases, you might have to start up a new cluster, migrate your data to that cluster and replace this cluster with the existing one



Backup and Restore Updates

As your cluster configuration changes over time, you need to evolve your backup and recovery procedures and/or scripts that you've put in place. These changes can be either due to node changes or due to the fact that you have a unique collection / shard / replication set up that needs to be backed up appropriately.

They may also be used by performance testing teams to ensure their golden set of data is replaced after a period of time so that they have more control and predictability on their tests.

With SearchStax: 10-15 Minutes

***Part of product offering**

1. Use the SearchStax Cloud dashboard to initiate one-time backups or create a custom backup schedule (daily, weekly)
2. Utilize cloning feature as part of SearchStax Cloud to start a new deployment based on the existing backup that you've create

Do It Yourself: 12-48 Hours

2x/year=1-4 days/year

1. Change script to adapt to the new needs.
2. Develop scripts to perform restores or recovery
3. Test the integrity of the backup and scripts



Managing & Supporting the Production Infrastructure

Great now that you have monitoring set up, you'll need to have a process in place to respond to the things that your monitoring uncovers as you support your production application. Working with numerous applications and systems, we've seen tons of issues. To list out just a few: issues related to performance degradation, SSL certificate updates, GC pauses, authentication issues, JVM heap changes, OS level security patches, node replacements, disk drive failures, network failures, zookeeper queue issues, replica recovery, security audits and more. They may also be used by performance testing teams to ensure their golden set of data is replaced after a period of time so that they have more control and predictability on their tests.

With SearchStax: 0-30 Minutes
depending on your plan

Standard and Gold Customers

1. Receive alert
2. Determine severity of the issue and file ticket with our support team
3. Let us handle it

Platinum Support Customers

1. Receive alert
2. Roll over and go back to sleep because the SearchStax support team will automatically handle it for you and contact you if they need your approval on anything

Do It Yourself: 1-4 Days
12x/year=12-48 days/year

1. Get an alert around an anomaly at an inconvenient time
2. Drop what you're doing; whether that's going home or to your office or get up in the middle of the night and rush to your computer or office
3. Assess the situation: What Severity level is the situation? Is the service running? Slow? Memory leaks? Errors?
4. Diagnose the problem
5. Call in whatever resources you need to address the problem; other developers or team members if needed.
6. Begin work to resolve the issue, possibly restarting servers, changing configurations, starting new load balancers, new servers, patching code, etc.



Summary

Building your own search application can be a very time consuming process. Even when you are “done”, there is still work to do in the form of responding to issues and events, applying security patches and upgrades, and keeping software versions up to date. A Solr-as-a-service offering like SearchStax Cloud automates a huge part of this work without losing flexibility or control enabling your developers to focus on the search experience. The table below summarizes the potential savings of a basic search application.

To learn more about SearchStax Cloud, go to searchstax.com/managed-solr

INITIAL SET UP	SEARCHSTAX	DIY
Build Cluster	10-30 minutes	2-10 days
Monitoring & Alerting	15 minutes	3-20 days
Log Management	0: Built into product	12-48 hours
Backups Set up & Test	5-10 minutes	2-5 days
Build a Disaster Recovery Cluster	0-60 minutes	2-10 days
Total	30-115 minutes	9.5-47 days
ONGOING SUPPORT	SEARCHSTAX	DIY
Scaling	0-30 minutes	2-10 days
Upgrades	0: Built into product	3-20 days
Baackup/Recovery	10-15 minutes	12-48 hours
Production Support	5-15 minutes: Built into product	2-5 days
Total	15-60 minutes	23-72 days
Total Time	45-175 minutes	32.5-119 days

The traditional DIY build or buy options are now surpassed by using SearchStax Cloud. Let us deploy, manage, and scale your search infrastructure for faster and more agile development, streamlined operations and lower total cost of ownership or TCO.



SEARCHSTAX MAKES POWERFUL SEARCH EASY

We offer end-to-end SaaS search solutions - with fully-managed Solr Service on the backend and powerful website search on the frontend.

SearchStax Cloud is a fully-managed SaaS solution that automates, manages and scales Solr infrastructure in public or private clouds on AWS, Azure and Google Cloud. Developers use SearchStax Cloud to:

- Implement faster with less risk
- Focus resources on building better search experiences
- Spend less time managing and supporting Solr

SearchStax Studio delivers advanced, modern and personalized search for your website or custom application. Marketers can deliver best-in-class site search experiences with:

- Actionable search insights
- Self-service, easy-to-use tools for the Marketing team
- Quick implementation delivering rapid ROI for Developers

SearchStax Sales

Headquartered in California

sales@searchstax.com

+1.844.973.2724



SEARCHSTAX

