

WEBVTT

NOTE duration:"00:05:42"

NOTE recognizability:0.800

NOTE language:en-us

NOTE Confidence: 0.785437938

00:00:00.000 --> 00:00:02.776 NeuroStack is an AWS CloudFormation-based

NOTE Confidence: 0.785437938

00:00:02.776 --> 00:00:05.135 tool designed to aid neuroimaging researchers

NOTE Confidence: 0.785437938

00:00:05.135 --> 00:00:07.541 working with AWS cloud computing.

NOTE Confidence: 0.785437938

00:00:07.550 --> 00:00:10.022 It is ideal for neuroimaging researchers

NOTE Confidence: 0.785437938

00:00:10.022 --> 00:00:12.675 transitioning their workflows to AWS or

NOTE Confidence: 0.785437938

00:00:12.675 --> 00:00:14.559 anyone working with neuroimaging at scale.

NOTE Confidence: 0.785437938

00:00:14.560 --> 00:00:16.080 By installing NeuroStack

NOTE Confidence: 0.785437938

00:00:16.080 --> 00:00:17.582 into your AWS account,

NOTE Confidence: 0.785437938

00:00:17.582 --> 00:00:20.179 you construct a pipeline of AWS

NOTE Confidence: 0.785437938

00:00:20.179 --> 00:00:22.578 resources specialized for neuroimaging.

NOTE Confidence: 0.785437938

00:00:22.580 --> 00:00:25.002 NeuroStack is built with the NITRC

NOTE Confidence: 0.785437938

00:00:25.002 --> 00:00:26.518 computational environment, or NITRC-CE,

NOTE Confidence: 0.785437938

00:00:26.518 --> 00:00:29.381 allowing users access to all of the

NOTE Confidence: 0.785437938

00:00:29.381 --> 00:00:32.309 programs pre-installed in the NITRC-CE.
NOTE Confidence: 0.785437938

00:00:32.310 --> 00:00:34.598 In this tutorial, we'll go over how to
NOTE Confidence: 0.785437938

00:00:34.598 --> 00:00:36.827 install NeuroStack into your AWS account,
NOTE Confidence: 0.785437938

00:00:36.830 --> 00:00:39.147 and two examples of using NeuroStack
NOTE Confidence: 0.785437938

00:00:39.147 --> 00:00:41.249 using two different software packages,
NOTE Confidence: 0.785437938

00:00:41.250 --> 00:00:42.591 AFNI and FreeSurfer.
NOTE Confidence: 0.785437938

00:00:42.591 --> 00:00:44.826 To begin using NeuroStack,
NOTE Confidence: 0.785437938

00:00:44.830 --> 00:00:48.448 first sign into your AWS account.
NOTE Confidence: 0.785437938

00:00:48.450 --> 00:00:50.802 Navigate to the NeuroStack page on
NOTE Confidence: 0.785437938

00:00:50.802 --> 00:00:52.990 the NITRC website and go to Downloads.
NOTE Confidence: 0.785437938

00:00:52.990 --> 00:00:55.621 The download is a URL that will bring you
NOTE Confidence: 0.785437938

00:00:55.621 --> 00:00:57.850 to the AWS CloudFormation stack page.
NOTE Confidence: 0.8139093633333333

00:01:01.950 --> 00:01:03.966 Once on the Stack Creation page,
NOTE Confidence: 0.8139093633333333

00:01:03.970 --> 00:01:06.310 everything you need is already preset.
NOTE Confidence: 0.8139093633333333

00:01:06.310 --> 00:01:07.970 Just continue choosing next.
NOTE Confidence: 0.88287323

00:01:11.020 --> 00:01:13.652 On the final page, check the box to

NOTE Confidence: 0.88287323
00:01:13.652 --> 00:01:16.288 accept that CloudFormation will create
NOTE Confidence: 0.88287323
00:01:16.288 --> 00:01:19.144 IAM resources before creating the stack.
NOTE Confidence: 0.88287323
00:01:19.150 --> 00:01:21.064 NeuroStack will take several minutes
NOTE Confidence: 0.88287323
00:01:21.064 --> 00:01:23.259 to set up your AWS resources.
NOTE Confidence: 0.88287323
00:01:23.260 --> 00:01:27.350 When it says "Create Complete", it's finished.
NOTE Confidence: 0.88287323
00:01:27.350 --> 00:01:28.700 Let's run a simple job through
NOTE Confidence: 0.88287323
00:01:28.700 --> 00:01:30.230 NeuroStack to see how it works.
NOTE Confidence: 0.88287323
00:01:30.230 --> 00:01:31.861 We're going to convert an MRI file
NOTE Confidence: 0.88287323
00:01:31.861 --> 00:01:33.389 from one file type to another.
NOTE Confidence: 0.88287323
00:01:33.390 --> 00:01:36.000 We're going to use an AFNI program to convert
NOTE Confidence: 0.88287323
00:01:36.000 --> 00:01:38.506 from a NIFTI filetype to an AFNI filetype.
NOTE Confidence: 0.88287323
00:01:38.510 --> 00:01:39.812 The job should only take a few
NOTE Confidence: 0.88287323
00:01:39.812 --> 00:01:41.800 minutes to run, so you can see the
NOTE Confidence: 0.88287323
00:01:41.800 --> 00:01:43.230 process from beginning to end.
NOTE Confidence: 0.88287323
00:01:43.230 --> 00:01:44.550 I've downloaded the template
NOTE Confidence: 0.88287323

00:01:44.550 --> 00:01:46.880 script and a test imaging file from
NOTE Confidence: 0.88287323

00:01:46.880 --> 00:01:48.350 the NeuroStack GitHub Repo.
NOTE Confidence: 0.88287323

00:01:48.350 --> 00:01:50.555 I'll navigate to the AWS S3 console where I
NOTE Confidence: 0.88287323

00:01:50.555 --> 00:01:52.985 can see that NeuroStack has configured 3 S3
NOTE Confidence: 0.88287323

00:01:52.990 --> 00:01:55.025 storage buckets and input,
NOTE Confidence: 0.88287323

00:01:55.025 --> 00:01:56.653 output, and script buckets.
NOTE Confidence: 0.810976631

00:02:00.040 --> 00:02:01.625 I'll upload the script to
NOTE Confidence: 0.810976631

00:02:01.625 --> 00:02:03.210 the NeuroStack script S3 bucket.
NOTE Confidence: 0.935507962857143

00:02:08.670 --> 00:02:10.644 Let's take a look at the script.
NOTE Confidence: 0.935507962857143

00:02:10.650 --> 00:02:12.540 The top section contains commands that
NOTE Confidence: 0.935507962857143

00:02:12.540 --> 00:02:14.910 aren't intended to be altered by the user.
NOTE Confidence: 0.935507962857143

00:02:14.910 --> 00:02:16.395 The bottom section is meant
NOTE Confidence: 0.935507962857143

00:02:16.395 --> 00:02:18.250 to be modified by the user.
NOTE Confidence: 0.935507962857143

00:02:18.250 --> 00:02:20.586 There are three basic parts to the section:
NOTE Confidence: 0.935507962857143

00:02:20.590 --> 00:02:23.628 copying data from our input S3 bucket,
NOTE Confidence: 0.935507962857143

00:02:23.630 --> 00:02:25.670 performing an operation on the data,

NOTE Confidence: 0.935507962857143

00:02:25.670 --> 00:02:27.854 in this case using an AFNI command

NOTE Confidence: 0.935507962857143

00:02:27.854 --> 00:02:30.347 to convert file types, and copying our

NOTE Confidence: 0.935507962857143

00:02:30.347 --> 00:02:33.140 processed data to our output S3 bucket.

NOTE Confidence: 0.935507962857143

00:02:33.140 --> 00:02:34.869 An AFNI file consists of two

NOTE Confidence: 0.935507962857143

00:02:34.869 --> 00:02:36.358 files a .BRICK and a .HEAD.

NOTE Confidence: 0.935507962857143

00:02:36.360 --> 00:02:39.804 So we're actually copying 2 files here.

NOTE Confidence: 0.935507962857143

00:02:39.810 --> 00:02:42.474 The variable aSub is defined as the

NOTE Confidence: 0.935507962857143

00:02:42.474 --> 00:02:44.299 filename before the first period.

NOTE Confidence: 0.935507962857143

00:02:44.300 --> 00:02:46.550 So in this example we have

NOTE Confidence: 0.935507962857143

00:02:46.550 --> 00:02:48.560 a test file called test,

NOTE Confidence: 0.935507962857143

00:02:48.560 --> 00:02:49.905 which means that aSub

NOTE Confidence: 0.935507962857143

00:02:49.905 --> 00:02:51.250 will take the value "test".

NOTE Confidence: 0.4369919

00:02:53.910 --> 00:02:55.940 I'll upload the test imaging file to

NOTE Confidence: 0.4369919

00:02:55.940 --> 00:02:57.876 the NeuroStack input bucket. When I

NOTE Confidence: 0.4369919

00:02:57.876 --> 00:02:59.507 upload a file into the input bucket,

NOTE Confidence: 0.4369919

00:02:59.510 --> 00:03:02.205 it will immediately begin processing the file
NOTE Confidence: 0.4369919

00:03:02.205 --> 00:03:05.188 according to my script in the script bucket.
NOTE Confidence: 0.4369919

00:03:05.190 --> 00:03:06.540 It will take several minutes
NOTE Confidence: 0.4369919

00:03:06.540 --> 00:03:07.890 for the job to complete.
NOTE Confidence: 0.4369919

00:03:07.890 --> 00:03:09.370 When it has finished successfully,
NOTE Confidence: 0.4369919

00:03:09.370 --> 00:03:11.491 we will see our processed data in
NOTE Confidence: 0.4369919

00:03:11.491 --> 00:03:13.559 the NeuroStack output S3 bucket.
NOTE Confidence: 0.4369919

00:03:13.560 --> 00:03:14.550 In the meantime,
NOTE Confidence: 0.4369919

00:03:14.550 --> 00:03:16.200 let's watch our job status.
NOTE Confidence: 0.4369919

00:03:16.200 --> 00:03:17.999 We can do that by navigating to
NOTE Confidence: 0.4369919

00:03:17.999 --> 00:03:19.346 the AWS batch console.
NOTE Confidence: 0.4369919

00:03:19.346 --> 00:03:21.880 The job will move from "Submitted" to
NOTE Confidence: 0.4369919

00:03:21.950 --> 00:03:23.860 "Running". When it has finished running,
NOTE Confidence: 0.4369919

00:03:23.860 --> 00:03:26.004 it will either move to "Succeeded" or "Failed".
NOTE Confidence: 0.7631748

00:03:35.010 --> 00:03:37.604 Great, we can see that our job succeeded.
NOTE Confidence: 0.7631748

00:03:37.610 --> 00:03:39.314 We can see more details about

NOTE Confidence: 0.7631748

00:03:39.314 --> 00:03:40.860 our job by clicking on it.

NOTE Confidence: 0.7631748

00:03:40.860 --> 00:03:42.684 And we can see that the job's total

NOTE Confidence: 0.7631748

00:03:42.684 --> 00:03:43.920 running time was 12 seconds,

NOTE Confidence: 0.7631748

00:03:43.920 --> 00:03:45.332 which is the computation

NOTE Confidence: 0.7631748

00:03:45.332 --> 00:03:47.097 time that we're charged for.

NOTE Confidence: 0.7631748

00:03:47.100 --> 00:03:48.676 We can also look at the log of

NOTE Confidence: 0.7631748

00:03:48.676 --> 00:03:50.267 our job under Log Stream Name.

NOTE Confidence: 0.858164316666667

00:03:56.980 --> 00:03:58.800 And when we navigate to our output

NOTE Confidence: 0.858164316666667

00:03:58.800 --> 00:04:00.465 S3 bucket, we can see that our

NOTE Confidence: 0.858164316666667

00:04:00.465 --> 00:04:01.610 two processed files are there.

NOTE Confidence: 0.881239327777778

00:04:04.050 --> 00:04:05.076 In this example,

NOTE Confidence: 0.881239327777778

00:04:05.076 --> 00:04:07.128 we've uploaded a single subject file,

NOTE Confidence: 0.881239327777778

00:04:07.130 --> 00:04:08.908 but we can upload multiple files at

NOTE Confidence: 0.881239327777778

00:04:08.908 --> 00:04:10.536 a time for simultaneous processing

NOTE Confidence: 0.881239327777778

00:04:10.536 --> 00:04:12.551 by uploading the files either

NOTE Confidence: 0.881239327777778

00:04:12.551 --> 00:04:14.652 manually or through the AWS CLI.
NOTE Confidence: 0.881239327777778

00:04:14.652 --> 00:04:16.969 If you have many files to process,
NOTE Confidence: 0.881239327777778

00:04:16.970 --> 00:04:19.940 the CLI will be more efficient.
NOTE Confidence: 0.881239327777778

00:04:19.940 --> 00:04:22.040 Let's use NeuroStack to process
NOTE Confidence: 0.881239327777778

00:04:22.040 --> 00:04:23.440 the subject through FreeSurfer.
NOTE Confidence: 0.881239327777778

00:04:23.440 --> 00:04:25.784 To do this, we'll need to change our
NOTE Confidence: 0.881239327777778

00:04:25.784 --> 00:04:27.820 script in the user modified section.
NOTE Confidence: 0.881239327777778

00:04:27.820 --> 00:04:30.375 I have the FreeSurfer setup commands.
NOTE Confidence: 0.881239327777778

00:04:30.380 --> 00:04:33.296 I copied the data from our input S3 bucket.
NOTE Confidence: 0.881239327777778

00:04:33.300 --> 00:04:35.916 Freesurfer requires a license file to run,
NOTE Confidence: 0.881239327777778

00:04:35.920 --> 00:04:38.688 so I'll copy my license file as well.
NOTE Confidence: 0.881239327777778

00:04:38.690 --> 00:04:40.895 I'll then do recon-all, the command
NOTE Confidence: 0.881239327777778

00:04:40.895 --> 00:04:43.254 to process through FreeSurfer, and once
NOTE Confidence: 0.881239327777778

00:04:43.254 --> 00:04:45.788 that process is done I'll copy my
NOTE Confidence: 0.881239327777778

00:04:45.788 --> 00:04:48.450 output files to my output S3 storage.
NOTE Confidence: 0.881239327777778

00:04:48.450 --> 00:04:50.268 Now that we've modified our script,

NOTE Confidence: 0.881239327777778
00:04:50.270 --> 00:04:52.377 we can upload our script and the
NOTE Confidence: 0.881239327777778
00:04:52.377 --> 00:04:54.105 FreeSurfer license file to the
NOTE Confidence: 0.881239327777778
00:04:54.105 --> 00:04:55.897 script bucket and our test data to
NOTE Confidence: 0.881239327777778
00:04:55.958 --> 00:04:57.903 our input bucket. Again, by uploading
NOTE Confidence: 0.881239327777778
00:04:57.903 --> 00:04:59.307 to our input bucket,
NOTE Confidence: 0.881239327777778
00:04:59.310 --> 00:05:01.576 we started our process running.
NOTE Confidence: 0.881239327777778
00:05:01.576 --> 00:05:03.298 FreeSurfer will take several hours to run.
NOTE Confidence: 0.752192962352941
00:05:11.470 --> 00:05:13.530 It's the next day and I can see our
NOTE Confidence: 0.752192962352941
00:05:13.587 --> 00:05:15.547 FreeSurfer output in the output S3 bucket.
NOTE Confidence: 0.752192962352941
00:05:15.550 --> 00:05:18.205 I can go to the Batch dashboard and see
NOTE Confidence: 0.752192962352941
00:05:18.205 --> 00:05:20.775 that FreeSurfer took six hours to run.
NOTE Confidence: 0.752192962352941
00:05:20.780 --> 00:05:23.028 We hope you will find NeuroStack useful
NOTE Confidence: 0.752192962352941
00:05:23.028 --> 00:05:25.138 for your neuroimaging work on AWS.
NOTE Confidence: 0.752192962352941
00:05:25.140 --> 00:05:26.270 If you have questions about
NOTE Confidence: 0.752192962352941
00:05:26.270 --> 00:05:27.400 how to use NeuroStack,
NOTE Confidence: 0.752192962352941

00:05:27.400 --> 00:05:29.032 please visit the NeuroStack Forum

NOTE Confidence: 0.752192962352941

00:05:29.032 --> 00:05:31.120 on the NITRC website. Thank you.