Long Term Archive Frequently Asked Questions

This page is targeted at users who want to learn more on the LOFAR LTA (Long Term Archive). It should answer the most common questions and provide help in case of difficulties in data retrieval. If the information on this page did not succeed to solve any issues you may encounter, please send a mail to Science Support.

Questions

General

- Q: I never worked with the LTA before. Where is a good starting point?
- Q: Why is data retrieval so difficult?
- Q: What is an appropriate amount of data to retrieve?
- Q: What is all this SRM/staging stuff about?
- Q: There are different ways to download. Which one is the best?
- Q: My download speeds are too slow. What can I do?

Troubleshoot

- Q: I did not receive a mail notification that my request was scheduled!
- Q: I did not receive a mail notification that my data is ready for retrieval! Has my request gone lost?
- Q: I got an email that says my staging request has failed! What happened?
- Q: I got an email that says my staging request was only partially successful! What's going on?
- Q: Oops! I made a mistake! How can I stop a request?
- Q: My files only contain some error message instead of data
- Q: My data files are corrupted / I cannot unpack my data
- Q: My downloads fail with error "All Ready slots are taken and Ready Thread Queue is full"
- Q: My downloads don't start / time out
- Q: Http downloads randomly fail with "503 Service Temporarily Unavailable"
- Q: SRM commands fail with error containing "failed with error Java heap space"

Answers

General

I never worked with the LTA before. Where is a good starting point?

New users are advised to read through our tutorial page, the LTA How-To, which guides you through the whole process from getting a user account, over finding your data, to the actual download.

Last update: 2015-02-25 18:11

Why is data retrieval so difficult?

It is important to understand the the data volumes of LOFAR are pretty huge and handling them requires different technologies than what we all know and use in our everyday life. For instance, LTA data is stored on magnetic tape and has to be copied to a hard drive (getting 'staged') before it can be retrieved. To transfer these amounts of data within reasonable time requires careful consideration and special tools. We try to make the LTA as convenient to use as possible, e.g. by providing http downloads for users without Grid certificate and portable Grid tools for those who want or need the extra performance. We are aware of the fact that data retrieval is quite close to the backend technology and we hope to be able to provide solutions with higher abstraction in the future.

What is an appropriate amount of data to retrieve?

This depends. There are two things to consider: The capabilities of your own system and the capabilities of the LTA services. The most important thing to know about LTA capabilities, is that the disk pool that temporarily holds your data and from where it can be downloaded, is of limited capacity. This means that the data you requested is only available for download for a limited time (since the space is needed for new requests at some point). Your data is only guaranteed to stay available for 7 days. It can be re-requested after that, but you should never request more data than you can download within a few days. In most cases, this is limited by the capabilities of your own system, especially your network connection. (And available local storage space, of course.)

As a rule of thumb, we ask you to keep your requests below 5 TB in volume and smaller than 1'000 files. The larger your request, the longer it takes until you can retrieve the first file. Also, please limit the number of requests running in parallel to a few, especially when they contain many files. In principle, we avoid introducing hard limits, but rely on reasonable user behavior. This also means that you can block the system for a long time or, in the worst case, even bring it down. So please act responsibly or we might have to enforce some limits in the future to keep the system available for other users. Be aware, that we may cancel your request(s) in excessive cases to maintain LTA operation.

If you, by accident, staged some 100'000 files or 100 TB of data, please contact Science Support, so that we can stop these requests, thanks!

What is all this SRM / 'staging' stuff about?

These are technical terms that refer to the storage backend of the LTA. Each of the three LTA sites (in Amsterdam, Juelich and Groningen) operates an SRM (Storage Resource Management) system. Each SRM system consists of magnetic tape storage and hard disk storage. Both are addressed by a common file system, where each file has a specific locality: it can be either on disk ('online') or on tape ('nearline') or both. The usual case for LTA data is, that it is on tape only. Since the tape is not directly accessible but placed in a library shelf, the data on it first has to be copied from tape to disk, in order to retrieve it. This process is called 'staging'. Only while the data is (also) on disk, you will be able to download it. (In physics terms, think of it as an excited state.) To save cost, the disk pool is of limited capacity and only meant for temporary caching data that a user wants to access right now. After 7 days, all data is automatically 'released', which means that it may be deleted from the disk storage, as soon as the space is required for other data. It then has to be staged again in order to become accessible again.

Usually, you don't have to worry about the details. But be aware, that data retrieval is a two-step procedure: 1) preparation for download ('staging') and 2) the download itself. Also, take care not to request too much data at the same time.

There are different ways to download. Which one is the best?

That depends. In short terms: Http downloads are the easiest (e.g. via wget), downloads via SRM tools can be faster and are encouraged for large amounts.

The SRM systems which the LTA sites operate are integrated in the Grid. To work with them directly, you need a Grid certificate. To allow users without a Grid certificate to download LTA data, we operate webservers as a frontend to the SRM backend. These webservers provide the requested data via http downloads. The webservers are not excessively capable machines and meant for occasional users. If you retrieve huge amounts of data on a regularly timescale, please work with SRM directly, especially if you own a Grid certificate. We provide a portable Grid toolkit to make it as easy to set up as possible.

You may want to read this FAQ Answer as well to make a decision: My downloads are too slow. What can I do?

For user instructions, refer to the LTA How-To page.

My downloads are too slow. What can I do?

First of all, you have to check how slow your download really is. If wget shows an estimated time of arrival of several hours, this does not necessarily mean that the download is 'slow': some files in the LTA are also just really huge. In most cases, your local network connection will be the bottleneck. For instance, a standard 'Fast Ethernet' network connection allows download speeds of around 12 MB/s at a maximum. Our systems are able to handle that. In case you can rule out your network connection as the bottleneck: there are different ways to download your data and not all provide the same performance. By our experience, this is the order of performance:

- Http downloads are the slowest option. The speed is limited by the server's network connection (~120 MB/s), which is shared by all users, and an upper limit per download (around ~30 MB/s) for technical reasons. If your download maxes out at the per-download limit, you may try to start up to four downloads in parallel. Note: There is no performance benefit to expect from more than four parallel transfers! However, there is a connection limit, which you may trigger if you start too many parallel downloads.
- SRMCP is the faster option in most cases, since you work with the SRM backend directly.
- SRMCP + copy-script seems to be the fastest solution available. It uses globus-url-copy, which is reported to have superior performance over the default srmcp transfer.

You may also want to read this FAQ Answer for further explanation: There are different ways to download. Which one is the best?

Troubleshoot

Last update: 2015-02-25 18:11

I did not receive a mail notification that my request was scheduled!

If the LTA catalog did not show any error when you submitted your request, then it is safe to assume that your request was registered in our staging system. Usually, you should get a notification mail that this has happened within a few minutes. If you did not receive the notification within an hour, then our staging service may be down. Note that your request is not lost in this case and will be picked up after the service is back online. In urgent cases or if you are not sure that something went wrong while submitting your request, please contact Science Support.

I did not receive a mail notification that my data is ready for retrieval! Has my request gone lost?

After you got a notification that your requests was scheduled, it is in our database and there's hardly a possibility that it got lost. Staging requests can take up to a day or two, but will finish a lot sooner in most cases. This depends on your request's size but also on how busy the storage systems are by other user's requests at the moment. Sometimes, the LTA storage systems are down for maintenance and this can delay the whole procedure. You can check for downtimes here.

It is not alarming when your request did not finish in 24 hours, even when your last request finished within 10 minutes. In urgent cases or if you did not receive a notification after 48 hours, please contact Science Support.

I got an email that says my staging request has failed! What happened?

This means that the SRM server could not fulfill the request at all. This might mean that the system itself is fine, but none of the files from your request could be staged (e.g. missing files). Check the error message from your mail notification for details. The notification can also indicate that there is a general problem with the SRM system or with the staging service itself, i.e. something is broken or down for maintenance. We try to detect all temporary issues and only inform users in case that something is wrong with their request itself, but we cannot foresee all eventualities. If you cannot make sense out of the error message, or don't know how to deal with it, please contact Science Support.

If you used the xmlrpc interface to submit your request, please first check whether you made a mistake and e.g. entered the wrong SURLs.

Note: We get notified of these issues as well and will usually re-schedule failed requests due to server issues after the problem was solved. So please first check whether you got a 'Data ready for retrieval' notification for the same request id after the error notification. If you did, the problem was already resolved.

I got an email that says my staging request was only partially successful! What's going on?

In general, this means that the SRM system works fine, but there was a problem processing your request. As a result, some of your files could be staged, some could not. Your mail notification should include a list of which files could not be prepared for download successfully and also include an error

message to indicate the cause. If the error message says 'Incorrect URL: host does not match', this means that you combined files in a requests that are stored on two different SRM locations (e.g. one file at surfSARA and one file at Target). When one SRM location gets the request, it can only stage the local files. You have to request the files from different locations independently, to prevent this. Other messages should be self-explanatory, e.g. if a file is missing. If you cannot make sense out of the error message, or don't know how to deal with it, please contact Science Support.

If you used the xmlrpc interface to submit your request, please first check whether you made a mistake and e.g. entered the wrong SURLs.

Note: We get notified of these issues as well and will usually re-schedule failed requests due to server issues after the problem was solved. So please first check whether you got a 'Data ready for retrieval' notification for the same request id after the error notification. If you did, the problem was already resolved.

Oops! I made a mistake! How can I stop a request?

Unfortunately, this is currently not possible for you as a user. Stay calm and ask Science Support to stop the request for you.

My files only contain some error message instead of data

Please read the error message carefully. In many cases, it should give you some indication of what went wrong. If this does not help you, please contact Science Support or retry after a few hours.

Important: If you use wget with option '-c', please note the following: wget does not check the contents of an existing file, so when restarting wget with option '-c' (continue) to retrieve the failed files, it will append the later data chunk to the existing file that contains the error message (and not the first section of you data). Make sure to delete the existing error files (should be obvious by the small file size) before calling 'wget -ci' again, to avoid corrupted data. If you already ended up with a corrupted file, you have to delete that and re-retrieve the whole file.

My data files are corrupted

Check if the files are much smaller than you expect. Something might have gone wrong with the transfer. Please check the beginning of your files, e.g. with 'less'. If there is an error message, please refer to this answer. Otherwise, please try to re-retrieve an affected file. If this does not help, please contact Science Support.

My Downloads fail with error "All Ready slots are taken and Ready Thread Queue is full"

This usually means the SRM server system is overloaded and you should try again in a few hours.

Last update: 2015-02-25 18:11

My downloads don't start / time out

Maybe the SRM system is down for maintenance, please check http://web.grid.sara.nl/cgi-bin/lofar.py. If there is nothing going on, there is probably something wrong with the download service. Please try again a bit later and notify Science Support, if the issue persists.

Http downloads randomly fail with "503 Service Temporarily Unavailable"

This can indicate too many users downloading at the same time. Please try again a bit later. There is also a limit of simultaneous downloads you are allowed to start yourself. Please limit yourself to four simultaneous downloads, the overall download rate will not improve with a larger number of connections.

SRM commands fail with error containing "failed with error Java heap space"

The SRM tools ignore the system's default Java heap space settings. Either reduce the amount of files in that request or increase the SRM-specific heap space by setting an environment variable 'SRM_JAVA_OPTIONS' with a higher value (e.g. '-Xms256m -Xmx256m'; default is '-Xms64m - Xmx64m').

From:

https://www.astron.nl/lofarwiki/ - LOFAR Wiki

Permanent link:

https://www.astron.nl/lofarwiki/doku.php?id=public:lta_faq&rev=1424887900

Last update: 2015-02-25 18:11

