Session title: DRS 1.6 and beyond/DRS Access options for compute tools

Work Stream / Subgroup: Cloud Work Stream / DRS

Date (Day/Month)**: Tuesday Sep 17, 2024**
Time (CEST)**: 9:00 - 10:30 am AEST**
Session chair(s): Ian Fore, Michael Lukowski, Brian D O'Connor
Session chair(s) will be present: In-Person
Link to slides: GA4GH Connect 2024 - DRS 1.6 and beyond/DRS Access options for co…
Request for promotion (if desired): **Communications Request Form** — choose "Get help promoting a session at Connect or Plenary"

Aim of meeting:

- Scope out features/improvements for version 1.6 of DRS (to be released for the 2025 Plenary)
- Scoping out improved access options information for compute tools (Ian)

Session description:

- This session will primarily focus on needs gathering from Driver Projects and Implementers in the room
- Our goal will be to scope the DRS 1.6 release and the features that we focus on
- Likely candidates will include items already slated to be improved, such as "defrosting" DRS objects in cold storage
- We will also examine ways we can improve the access options for systems that compute on DRS objects.

Key takeaways and action items for April Connect 2025:

- Implement mechanism to encode geolocation/cloud constraints in DRS metadata
- Explore options for adding defrosting mechanisms to DRS for deep cold storage objects
- Explore adding object creation functionality to the DRS API

- Investigate options for managing trust relationships between DRS servers and clients
- Coordinate integration of Crypt4GH support in DRS
- Encourage driver projects to vote on DRS 1.5 PR #408
- Contribute to the DRS publication draft

Agenda:

	Agenda item	Presented by	Item type	Time
1.0	Welcome	Brian / Ian		5 mins
2.0	DRS 1.5, where have we been?	Brian		10 min
3.0	DRS 1.6 open discussion with attendees	Group		
3.1	Info/session on cold storage impact on access options for compute tools	Ian		20 min
3.2	Conversations w/ Driver Projects, implementers, and interested parties on what's important to them for 1.6			20 min
4.0	Scoping specific priorities	Group		20 min
4.1	Based on the most universally desired features, can we scope those feature requests and identify a champions. Prioritize via EasyRetro or similar			
5.0	DRS Publication			10 min
5.1	Review publication next steps			

Discussion items/ideas

- How to integrate with other Cloud APIs... DRS, WES, TES - Venkat
- Passports AAI working group... goal of better alignment between workstream... trust between a DRS server and Passports (e.g. RAS recognized DRS servers end point) - Surya
- Use case... files generated by a pipeline, care about the output "folder"... how would you use glob with DRS? Maybe DRS can think about a container w/ variable files below. -

Michele

Resources and links:

- See the Cloud WS notes document.

Attendance (Name / Affiliation) : Takudzwa N Musarurwa (University of Cape Town)

Xx
David Steinberg (UC Santa Cruz), Ian Fore (NIH), Nicholas Owen (UCL), Denis Yuen (OICR), Lishuang Shen (Childrens Hospital Los Angeles), Jose Fernandez (BSC | Elixir-es), Bill Clifford,Mu, Rajeev Nedumpally (NIH/CIT), Maxime Hebbar, Nattwet Sriwichai,Oscar, Ayon Meingtjes, Jess Mcgrath, William, Xiao, Jiawen, Matthew McAuliffie (CIT), Mike Warfe, Jilong Liu, Kevin, Alistair Thompson (ARPA-H)

Zoom recording :
https://us02web.zoom.us/rec/share/lX3Yedh-S9zFfHVBaqFKZBRZv5Dest9qbz9XMCwrl-Cj_oyoN0be9_pbn218Ml7v.PTsQ5VqosWhjlk-g
Passcode: R=der2.b

Meeting minutes:

- 2024 goal - continuing to help with adoption and refinement.
- Take a look at the DRS 1.5 and vote your opinion
- Join our publication

  Cold storage - See DRS 1.5 release documentation

  - There is some data that's in cold storage, and without this ability to say, is it available or not, client requests could fail if something's in cold storage. So this allows Velsera or Terra or other compute platforms to understand the state of a DRS object and give more helpful information back to a user saying, this is not available right now.

Metadata linking -

DRS publication -

DRS voting - Open now. See PR#408 . Planning to release by 1st October

Ian Fore - DRS cold storage and compute

- Data discovery using Data Connect which Brian has referred to using DRS to access the file and using WES to compute on the data.
- It might be worth checking what some of our assumptions around use of DRS and cloud storage are. The program level motivations are often that we want to use Cloud Storage as this gets us to the point where there's no need for download of data. We can compute the files in place. The benefits of this is, it avoids long download times for gigabyte size files, and it avoids local storage costs which the research funders, one way or another, end up paying for even though they may be incurred locally.

- Using the NCBI DRS service, this is comparing providing a file (a BAM file or a CRAM file) to the SAM tools program running in a TES environment on the Seven Bridges platform, such as the Cancer Research Data Commons Cavatica or biodata catalyst, etc. We were also looking at a couple of scenarios where externally, through WES you can get a signed URL, and that would allow you to download into or provide it to Santos (as Santos will take a URL as input too).
- What was interesting was you never actually saw the download. It streamed directly into SAM tools, and all of the streaming was or the download was taken up in the compute time, compared with a situation where you've downloaded the file from the NCBI bucket, and at this point, we weren't able to import the file directly as part of the compute into The Seven Bridges platform, but which, you can now do and which will be part of the examples of what I'm about to show. But once you have downloaded as a pre step the file into the bucket, and when you run the compute, what actually happens is the file gets transferred out of one bucket into the file storage, and whether I'm correctly calling it another bucket or not, but it's

cloud storage that is available to the compute, to the VM i.e. the Docker container that's running the compute

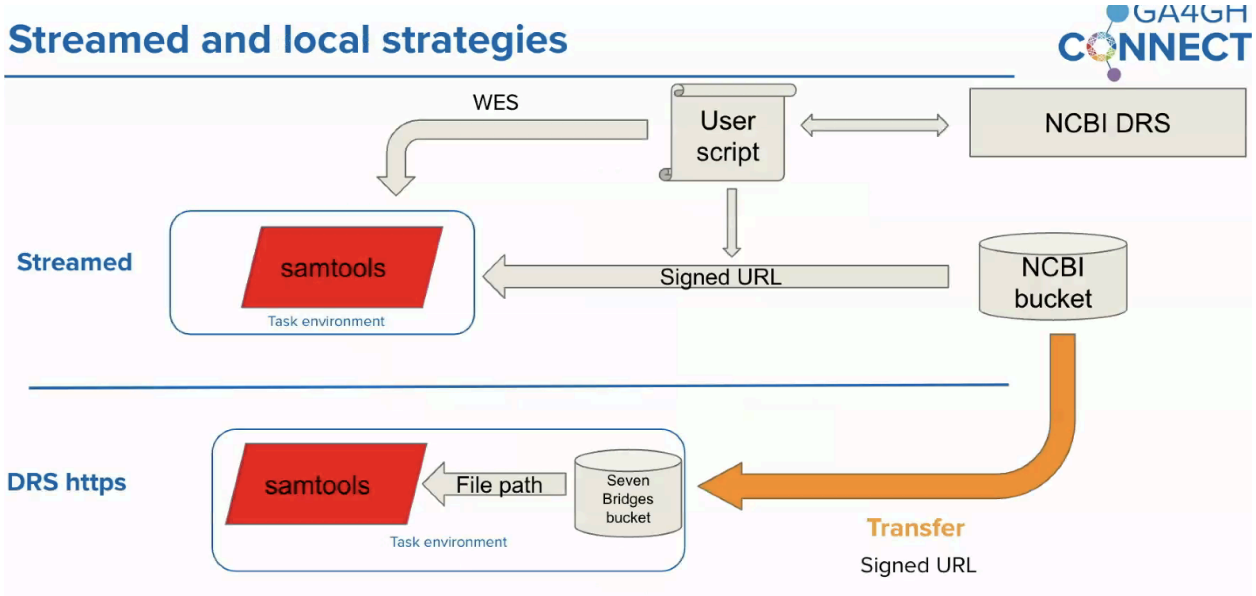

This is describing more the current situation - where the transfer happens, and this is highlighting the transfer, and there's color coordination here with the graphs (the orange trend). What's represented in those graphs in the orange is that transfer that happens when the Seven Bridges task that you create runs and it brings in through assigned URL, the file into the working area.

There are several files involved, but basically they all have this same structure. I'll indicate where that's not the case. But in this case, we've got a file from the Undiagnosed Disease network study, and this file is just available on Amazon (in the US east region). That's the only access method for this file. So this is my smoke and mirrors.

<Add image>

When you set up a project on Cavatica or the other Seven Bridges platforms, you choose whether you want to have that project be on AWS or Google. And we're going to show examples of both. You can see which regions those are using. And in this, I should also say that this is where we have got to have controlled access data, those files on the NCBI DRS are being accessed under RAs and GA4GH Passport control.

Transfer time comparisons for 15.9 UDN bam file

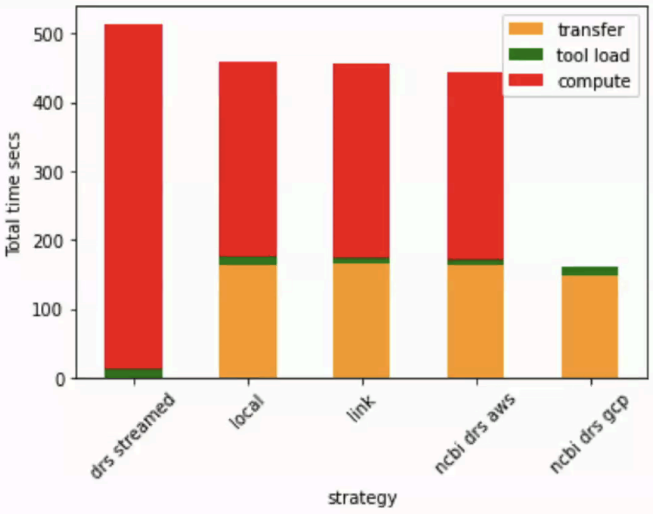

File on AWS in all cases

strategy	Task run on
ncbi drs aws	Amazon Web Services
ncbi drs gcp	Google Cloud Platform

This is in some ways repeating. What we've already seen on the left, where you've got these three phases, which is absent in the first one, and just leaving this for comparison of the history where the file is streamed in. Whether it's pulled in from another Seven Bridges DRS service, but the two interesting columns are those on the right which are about comparing when you set up a project where the compute is done on AWS or on GCP. What was interesting was that the file download transfer step, bringing it into local compute area, and making it available to the file system.

Compute comparison

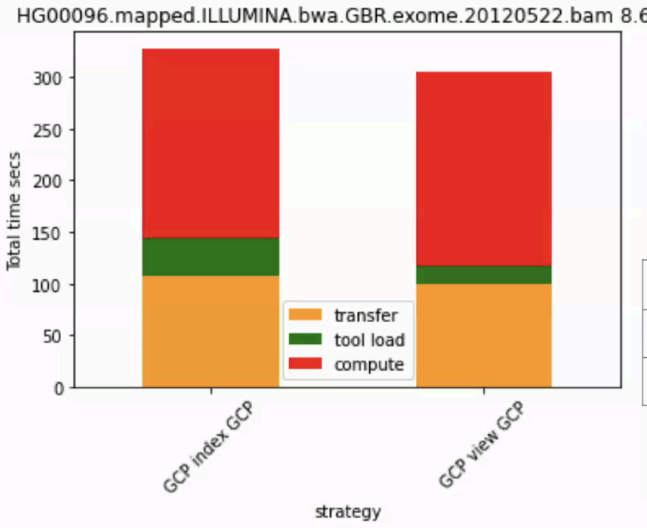

File and compute on GCP in all cases

strategy	Task run
GCP index GCP	Samtools index
GCP view GCP	Samtools view

This is the comparison - before we look at the download times, this is 8.6 gigabyte BAM file, with 1000 genomes. This BAM file is not from the public data. This is a copy of that file that are put in a hot bucket for the purposes of this compute and just to compare it with the other examples, because it's a smaller file, the debt, the transfer time is going to to be less, but with a couple of different SAM tools, programs run on the same file, the the transfer times are very similar.

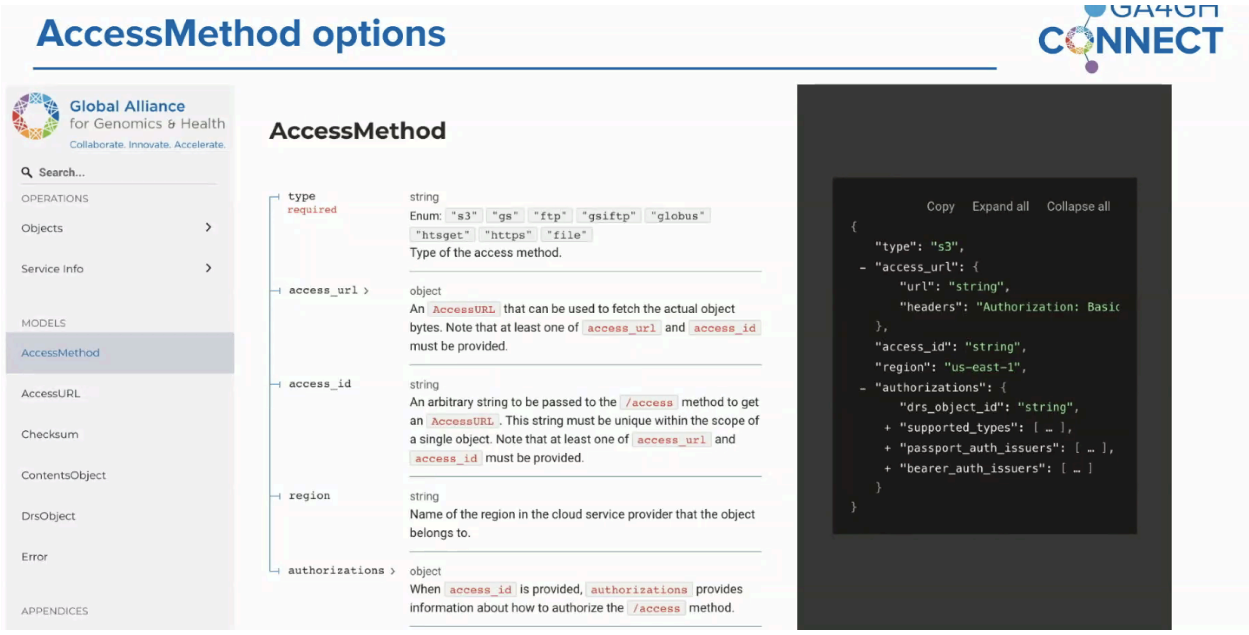

If we look at the access method definition, and Brian's already shown this in part, the type of access method that's being used in all cases. This is essentially HTTPS. When the explanation, I think, for why it is the same time to copy across to Google, to copy across either to Google or to AWS to do the compute is because it is actually doing a download over HTTPS of the file, and it's not using any direct native cloud protocols, which is what s3 and GS might be. But the point is, I think, you know, der says implementations have focused almost entirely on the HTTPS, methods, and so maybe there's something here that we you know might, might look

Implications for data in cold (archival) storage

Direct cloud access may minimize downside of cold storage

Penalty relative to https less than expected

Hypothesis:

If there is a need to copy from the original bucket anyway, is the simplest solution for cold data a direct copy to the storage used for the compute?

Key points

Minimal benefit from compute on same platform as storage

Https transfer via DRS urls - is essentially download

Download time counts as compute time

Direct cloud API access methods may have benefit

- Exploration of other access methods?

Direct copy from cold storage may be an efficient option

Securing direct cloud provider access likely more challenging

- But not impossible

Integration testing is valuable

Singapore: Looking at archival cold storage. It is cheaper. The requester has to find ways to restore it. Is there an appetite in DRS to have an option to request access and look at options if the requester is ready to bear the cost to do their own compute. **Ans** - How should cold storage interact? Different tiers of interfacing - cold storage, wait 24 hrs to retrieve but it is cheaper. Different styles of interaction. Async or little bit of delay whatever is acceptable in interaction.

Alaister Thompson - costs can be highly variable. Fast access is expensive. Communication required on the parameters of what DRS will handle and what won't be handled. **Ans**. Access method we have there - In the backlog for long. I wanted to access the date, can we pay for it? "Requester pays" is a possible way forward. Also look at the crowd for their opinion on, who should pay? Requester or someone else? Many in the audience were interested.

Susheel Verma - what exactly does DRS enable? Workflow or client? Any documentation available? **Ans**. Used by Velsera and lot others

BC - what are the priorities for us in DRS 1.6?

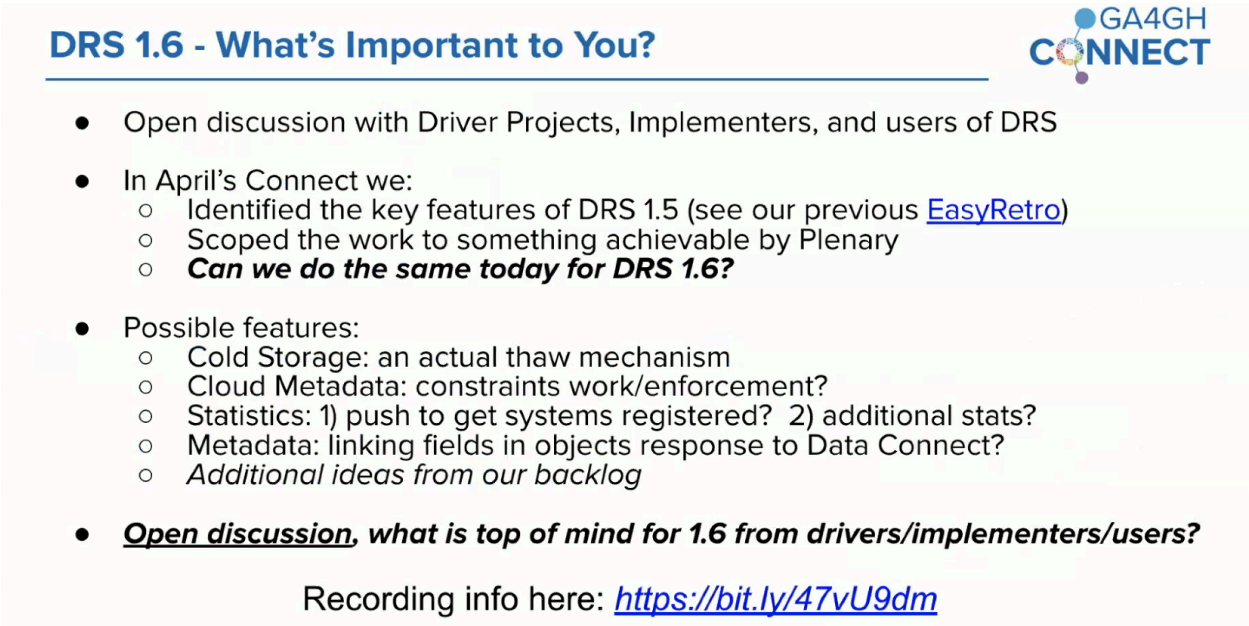

Whiteboard exercise to know more from Driver Projects and Implementers.

Andrew (NCBI) - Understand the ecosystem. Dashboard expanding on the metrics such as endpoints, data through the endpoints. This helps in seeing NCBI is evolving as an ecosystem. **Ans**. 2025 plans

Michael (Gen3) - Metadata is the big thing. Recently PR Data Connect as best practice for metadata. How to metadata ? Separate metadata service in Gen3 and not. **Ans**. Not wanting to connect DRS with metadata.

Surya - point on trust between DRS server and client. More trusted servers are coming into the ecosystem. Onerous to maintain over an extended period of time without compromising. Who is the right authority to maintain valid DRS servers? **Ans**. Seeing in platforms different combos with many DRS servers, trust is an issue. Needs to be in conjunction with the Passport group and not in silo. 40 PB of data. More feedback requested on the trust issue.

Michael - potential interest in Crypt4GH. There are mods to file for Crypt4GH standard and it becomes a problem in DRS. If you change bytes, it becomes a new object. How to manage or handle DRS? **Ans**. DRS ID is with MD5 checksum. Things not matching as data changed. Crypt4GH isn't simple and more about what DRS ID and sharing the original and new checksum for the same file.

Susheel - expanding on the above. We are a Data repo. We access data and many times it's being transferred. Mechanism required (DRS is download) to write. Sharing if the same data exists on different networks. How do DRS endpoints accept DRS objects and can be accessed later? Secondary point - if there's a copy of file, rather than transfer point it to the copy rather than share the copy. No opportunity for discovery of copies. **Ans**. Redirection based on relationships and awareness of the existence of the copy (MD5 checksum checks). Overcomes single point of data access as closer copy is available. Least expensive options explore more as it works internally in the DRS server.

Minimize the load.

Maxime - producer has a file but doesn't want it copied multiple times. No copies as it increases storage cost. Giving read access to the data. And also have a

central registry. **Ans**. DRS servers in different geo locations and there is a closer copy existing in DRS server. Ian showed access is faster on the cloud and it becomes the cheapest option. Are more standards or API required to expand on the registry? Mesh card helps in expanding the network. Can the registry help with this?

Don't need users creating multiple copies which leads to increased cost

Andrew - should I download? Or spread my compute across the cloud? **Ans**. The DRS registry is useless to DRS servers. Unawareness of the compute and copy functions.

Surya - locations of the data provider and associated restrictions. Freedom and open access are challenging. **Ans**. This would be answered in the Cloud session on geo location and sending compute to those nodes.

Susheel - DUO codes are available but what is it assigned to? Doesn't seem to be data. Can this be assigned to DRS objects? **Ans**. Interfacing with other Work Streams will help in building synergy.

Ian - encryption goes all the way. DRS will end up being the pass through. What's happening with compute tools and if they are ready with encrypted data? Room for discussions on it. **Ans**. It isn't done in isolation and more support is required. Cloud Work Stream session would discuss more and avoid isolation.

David - DRS has evolved to be fit for purpose as it is focused. What am I excited to enhance DRS? Whole ecosystem vision needs more work. Ways to find working compute with distributed data - DRS, WES, TES. DRS 1.x will have room for improvements.

Mirroring has potential value, no spec change required. DRS provides more access methods due to data availability across zones. DRS servers can call other DRS servers. As a client this should be hidden. DRS server can add value and not DRS API. **Ans**. Combination of Data Connect can help in this. DRS has been lively and that's why there aren't too many requests due to the maturity of the DRS.

Alastair - optimization through region, networking, etc. taking costs into consideration. TLS use and terms of trust - I want object to be used with some

region or compute env. TLS can be used to provide optimized access. Supporting the ability for DRS servers to expand on access is great. **Ans**. Box added on the whiteboard.

Crypt4GH support - we can achieve it in generic, without new stuff. Three needed - metadata for DRS object, expanding service info (not difficult), and if data is encrypted is new object and headers have metadata. Crypt4GH is good or is it narrowly defined and maybe expanded later. **Ans**. It was good.

David - Curious on content availability and moderation. Same file, different IDs. Use Hash to find DRS objects? **Ans**. Path of decouple. Endpoints are required to check MD5 and DPs are asked on their views.

Susheel - Hashes would be ideal. Hash lookup - it comes in through network effect. More DRS, more IDs shared. General IDs which can be understood by other DRS servers. Hashing mechanisms are not important right now. Nothing on spec side. It is about minting IDs and current spec supports it. **Ans**. Not practical to move over to checksums. Giving alternate endpoints isn't necessarily required right now esp. If DP doesn't need it, fine.

Andrew - Terra DRS hub. Intelligent handling of knowing about all the mirrors.

Surya - in network effect, user authentication. Different DRS servers have their own authentication. **Ans**. Bulk access serves this.

From the chat - @Susheel - The boolean 'available' is really only a stop gap - to handle what would otherwise raise an error. Sending people off to dig into specific APIs once they find out the data is in cold storage doesn't seem a helpful direction. Per DRS being an abstraction layer above specific cloud implementations it seems reasonable to address that most cloud providers have some tiered (3,4, more) storage options

The abstraction to ask for delivery to a certain cloud location seems workable. SRA has an implementation to do this. There are security issues to address, but something to work on.
Also - caveat that there are many variables involved in this that I haven't set out to explore e.g. the thaw costs that @Alastair mentioned. I looked for the Storage Transfer examples I showed - and none were visible. Am sure that's not the whole story! Integration testing is key.

Whiteboard results

Enforcement might be tricky. Jared has made a pull request. Registering a new DRS object.

Takeaway - these need champions. As follow up in Cloud WS, have champions move this forward through discussions, pull requests, etc. If champions aren't there, they usually don't move forward. Helps on focusing for the next year. DRS has hit the mark on its functionality and what's added is useful. Add self to the authors list. Consider joining the Cloud WS.

From the chat:

02:03:21 Venkat's iPhone: That depends on your tier and support level for cloud storage

02:03:36 Venkat's iPhone: Also were these experiments in the same region or across regions?

02:30:10 Ian Fore: @Susheel - The boolean 'available' is really only a stop gap - to handle what would otherwise raise an error. Sending people off to dig into specific APIs once they find out the data is in cold storage doesn't seem a helpful direction. Per DRS being an abstraction layer above specific cloud implementations it seems reasonable to address that most cloud providers have some tiered (3,4, more) storage options

The abstraction to ask for delivery to a certain cloud location seems workable. SRA has an implementation to do this. There are security issues to address, but something to work on.

Also - caveat that there are many variables involved in this that I haven't set out to explore e.g. the thaw costs that @Alastair mentioned. I looked for the Storage Transfer examples I showed - and none were visible. Am sure that's not the whole story! Integration testing is key.

02:30:50 David Steinberg: Content-addressability - get object by hash

02:33:11 Venkat's iPhone: Agreed

02:34:08 Venkat's iPhone: There is no example of that I. The real world

02:34:19 Venkat's iPhone: We need to. Hold the example and show what is possible

02:50:51 Ian Fore: Is there a link to the MIro? Board?

—-

Summary:

The meeting focused on the progress and future directions of DRS, a data access and storage system. Key points included the addition of a Boolean flag to indicate data availability in cold storage, which simplifies client requests. Ian Fore highlighted the benefits and challenges of cold storage and data access, noting that direct cloud copy methods can be efficient. The discussion emphasized the need for DRS 1.6 to address geolocation constraints, object creation, and passport trust mechanisms. Participants also stressed the importance of metadata management and the potential for a central registry to optimize data access and storage.

Action Items

- Implement mechanism to encode geolocation/cloud constraints in DRS metadata
- Explore adding object creation functionality to the DRS API
- Investigate options for managing trust relationships between DRS servers and clients
- Coordinate integration of Crypt4GH support in DRS
- Encourage driver projects to vote on DRS 1.5 PR #408Contribute to the DRS publication draft

Outline

DRS 1.5 Documentation and Best Practices

- Speaker 1 discusses the best practices for leveraging GA4 GH in the discovery workstream, emphasizing the importance of metadata and DRS IDs.
- The documentation includes recommendations for using Data Connect and DRS, focusing on metadata and DRS IDs.
- Speaker 1 mentions the availability of a flag to indicate whether data is available in cold storage, which is crucial for client requests.
- The cloud metadata includes region and cloud string, with plans to add more variables in DRS 2.0.

DRS 1.5 Documentation Changes and Publication

- Speaker 1 highlights the documentation changes, including best practices for using Data Connect to store metadata and keeping the DRS API simple.
- The DRS publication is nearing completion, with a draft expected by the end of 2024, and encourages participation from the community.
- Speaker 1 emphasizes the importance of voting on DRS 1.5 and the goal of achieving neutral or thumbs up from all voters by October 1.
- The session transitions to Ian Fore's presentation on cold storage and data access for compute purposes.

Ian Fore's Presentation on Cold Storage and Data Access

- Ian Fore discusses the evolution of data discovery, access, and compute using DRS, Data Connect, and Wes.

- The comparison of compute and storage permutations reveals that streaming data into compute tools is not always faster than downloading and storing data locally.
- Ian Fore explains the current situation with the NCBI DRS service, highlighting the transfer of data from cold storage to hot storage for compute purposes.
- The session includes examples and statistics on the transfer times of data between different cloud providers and storage tiers.

Discussion on Cold Storage and Data Access Mechanisms

- Ian Fore's presentation raises questions about the benefits of compute on the same platform as storage and the implications for cold storage.
- The discussion includes the potential for direct cloud API access methods to minimize the downside of cold storage.
- Speaker 1 and other participants emphasize the importance of experimentation to guide the design of DRS and understand the performance of different storage and access methods.

Connectivity between DRS objects and metadata services.

- Sushil from Seven Bridges disc The conversation touches on the need for requester pays mechanisms and the challenges of securing direct cloud provider access.

Audience Input and Prioritization for DRS 1.6

- Speaker 1 opens the floor for audience input on what they want to see in DRS 1.6, encouraging participants to use the Zoom whiteboard to vote on features.
- Andrew Gerhardt from NCPI emphasizes the need for a robust dashboard to map out the ecosystem and track data flow.
- Michael Lewski from Gen3 highlights the importance of metadata and the need for better uses, the need for trust between DRS servers and clients and the potential for a central registry of valid DRS servers.

Further Discussion on DRS 1.6 Features and Prioritization

- Maxim Hebbard from Singapore raises the issue of data producers wanting to minimize the number of copies of their data.
- Sushil extends the discussion on trust and the need for a central registry to minimize storage and transfer costs.
- David Glazer from Verily questions the need for new DRS features, suggesting that the focus should be on optimizing the existing ecosystem.
- Alistair Thompson from Gen3 proposes using attested TLS to provide cost factors and trust information for DRS servers.
- Final Voting and Prioritization of DRS 1.6 Features
- Speaker 1 gives participants a few minutes to vote on the most important features for DRS 1.6 using the Zoom whiteboard.
- The top-voted features include encoding geolocation constraints, object creation through the API, and passport and trust mechanisms.
- Speaker 1 emphasizes the need for champions to move these features forward and encourages participants to join the cloud sessions to discuss further.