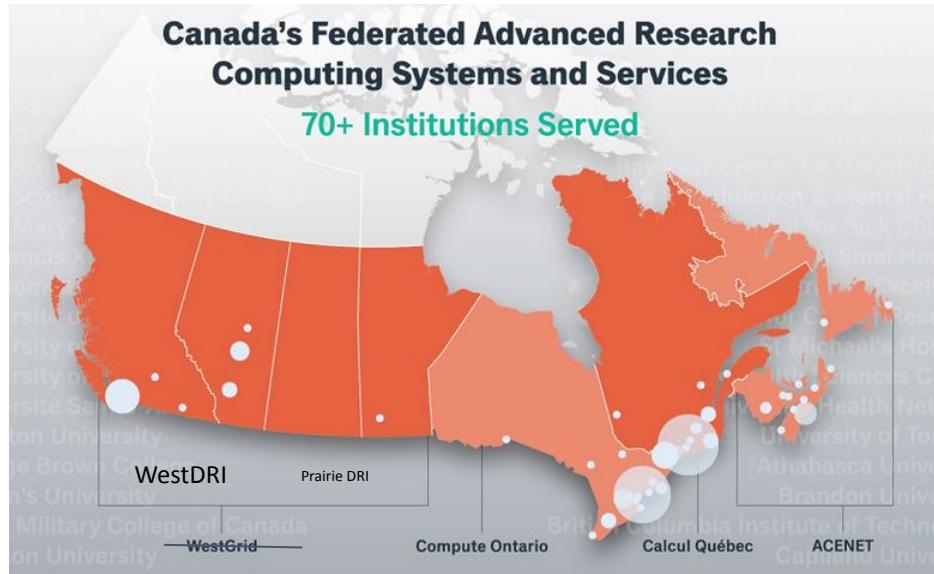


# CVMFS in Canadian Advanced Research Computing

Ryan Taylor  
on behalf of  
CVMFS National Team

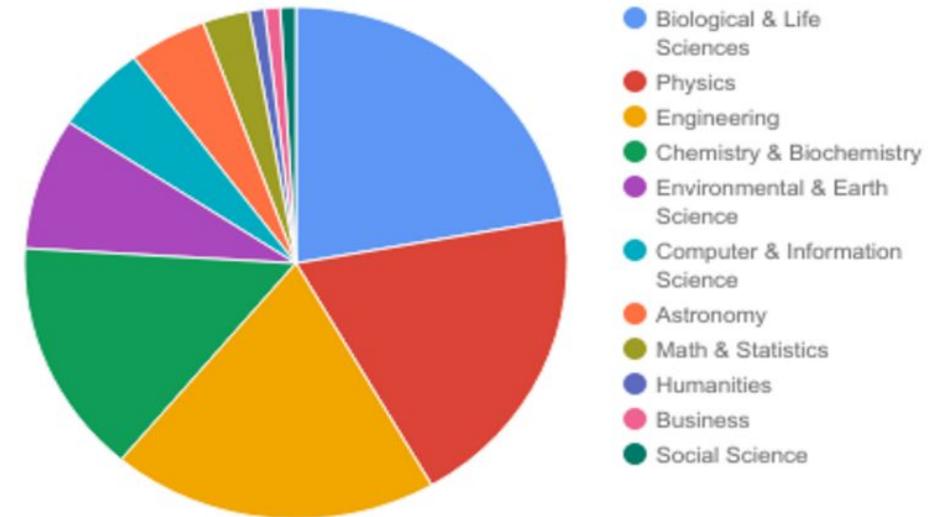
# ARC in Canada



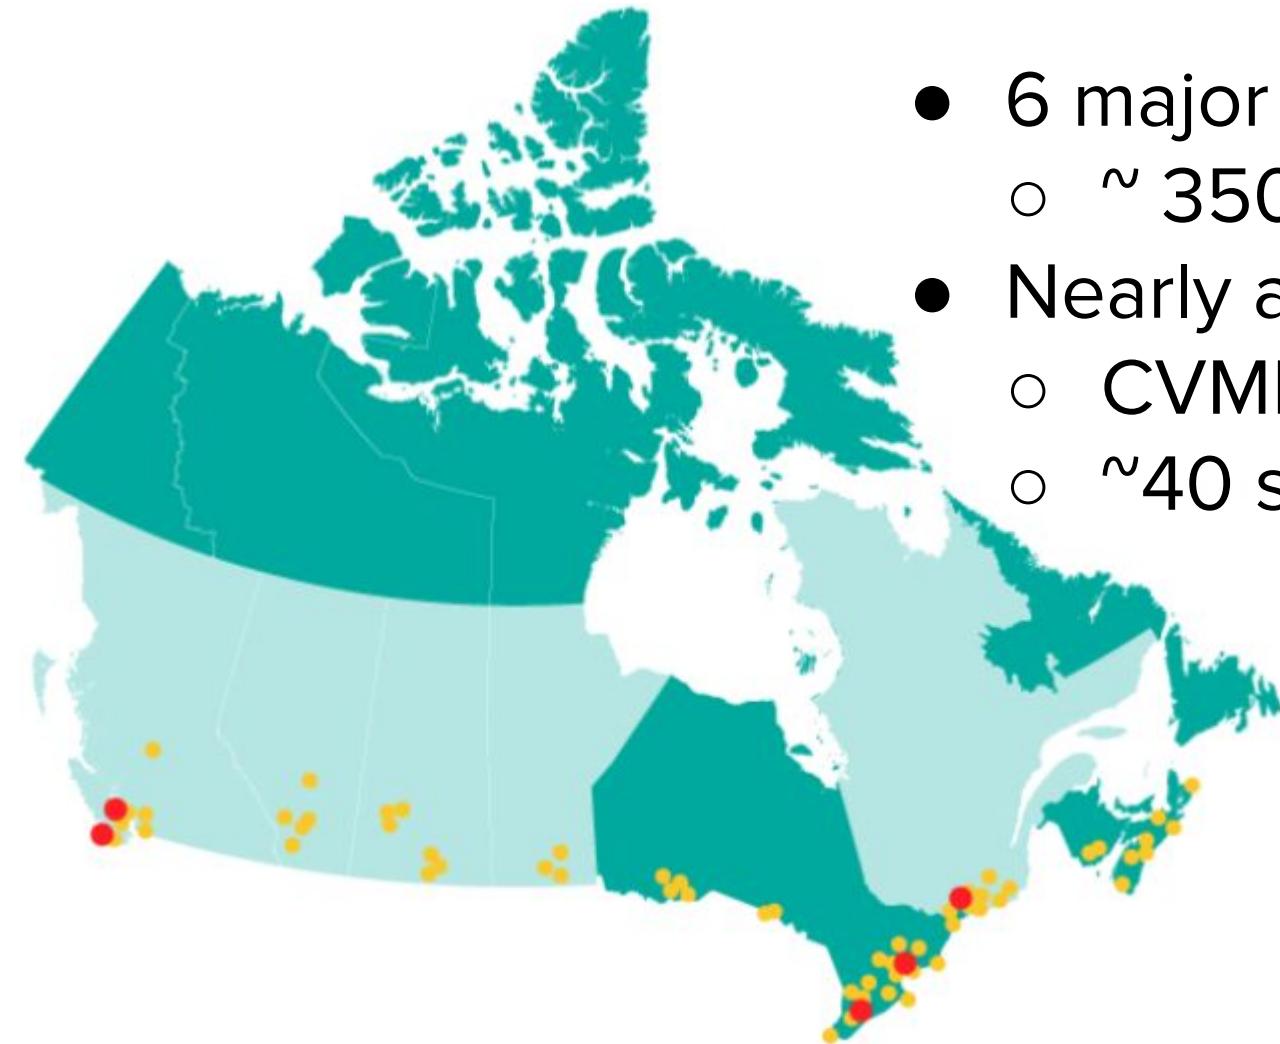
- Compute Canada -> Digital Research Alliance of Canada
- several regional consortia
- ~200 technical staff
- O(10K) user accounts
  - ~ 20% growth per year

WestDRI    

- All research disciplines supported
- Free access for any researcher at a Canadian institution



# ARC in Canada



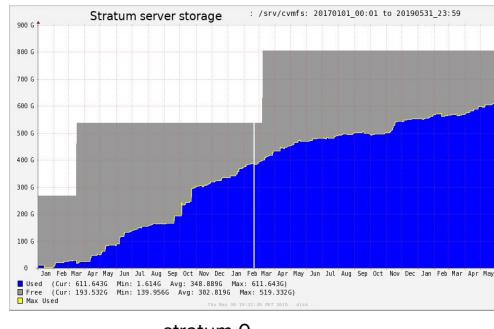
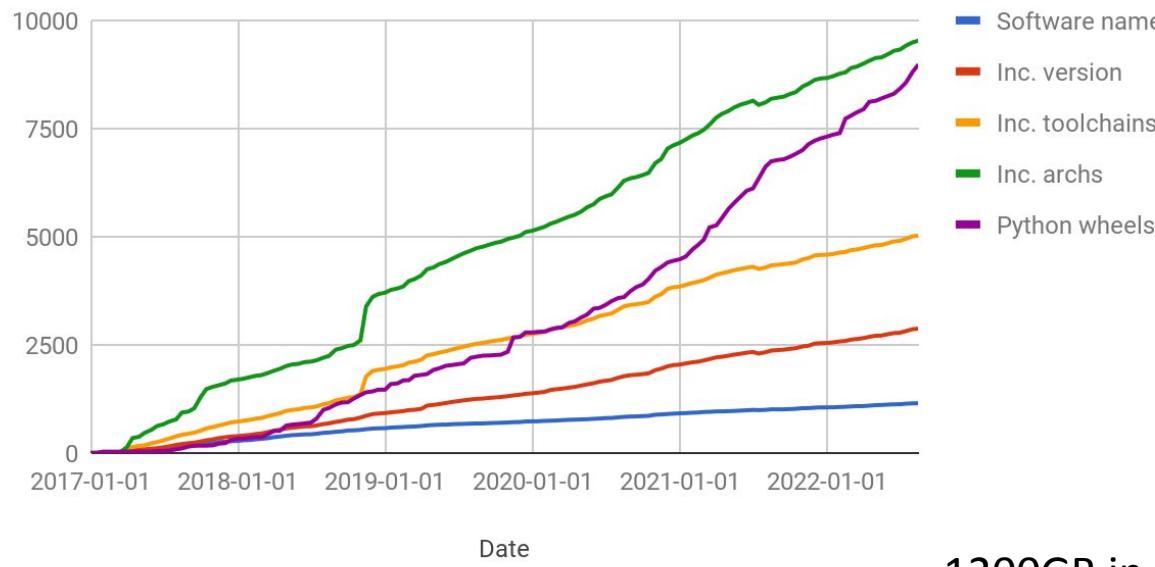
- 6 major national systems
  - ~ 350K cores, 50 PF, 200 PB disk
- Nearly all research software on CVMFS
  - CVMFS National Team: 8 members
  - ~40 staff members can publish software

Site	System	Type	Production
UVic	<b>Arbutus</b>	Cloud	2016H2
SFU	<b>Cedar</b>	General	2017H1
Waterloo	<b>Graham</b>	General	2017H1
U of T	<b>Niagara</b>	Large MPI	2018H1
ETS	<b>Beluga</b>	General	2018H2
ETS	<b>Narval</b>	General	2021H2

New EL8 cluster →

# Compute Canada CVMFS Software Stack

Number of software packages available through modules and python wheels



## Available software

- 1,150 scientific applications
- 9,500 modules (version/arch/toolchain)
- 9,000 python wheels
- no OS dependence, works on any Linux

Type	Modules
AI/data/tools	122
Bioinformatics	223
Chemistry	65
Geo/Earth	20
Mathematics	54
MPI libraries	4
Physics	27
Visualisation	32
Misc	416



# Migration to gateway publishing system

- Jan 2022: deployed new EL8 stratum servers, gateway service
- Repository migration via mkfs, demote to stratum 1, replicate, promote to stratum 0
  - With eliminate-hardlinks operation for > EL6, regenerates all catalogs
- Switchover from old to new stratum 0 via DNS update
  - Procedure fine for dev repos/servers
  - For prod, some churn and congestion at some peripheral sites
  - Incident at a (diskless) site using preloaded alien cache
- GW generally works smoothly, easier to operate/maintain than ad-hoc publishing scripts
- But we no longer have control of publisher settings

*Thanks to Enrico for notes*

# Open access to Compute Canada repos

- Available out of the box thanks to global config repo!
  - `source /cvmfs/soft.computeCanada.ca/config/profile/bash.sh`
- [Documentation](#)
- Using Cloudflare for resilience against DDoS
  - [Load balancer](#) of 4 stratum 1 origins
    - Two pools, East and West
  - [Dynamic steering](#) determines closest origin
  - Automated health checks and email alerts
  - [Zero-downtime](#) failover avoids returning HTTP errors
- BCNet peers research network with Cloudflare
  - Also benefits e.g. users pulling from Docker Hub

# Current projects and plans

- ~ 20 TB genomics data repository being launched
- Use S3 storage to support larger (~PB?) data repos, e.g. Biobank
  - with ACLs (network and POSIX) to protect restricted content
- Container image repository, for both auto and manual publishing
  - Use gateway keys to limit manually published images to separate path
  - Container tools not yet compatible with gateway

# CVMFS Ansible roles

- Client: <https://github.com/cvmfs-contrib/ansible-cvmfs-client>
- Stratum 0,1 server: (internal but potentially shareable)
  - Sets up and manages pubkeys, GC, log rotation, server meta info, replication, storage, firewalls, geodb key, cvmfs-servermon, etc.
  - Supports repository migration procedure (stratum 0-1 conversion)
  - Supports arbitrary repository configuration
  - Supports gateway setup, including API keys and both basic and advanced repo.json configuration

# Ansible inventory example

```
# define repo config for stratum servers
cvmfs_repo_list:
  - name: cvmfs-config.dev.computeCanada.ca
  - name: test.dev.computeCanada.ca
  - name: soft-dev.computeCanada.ca
    gateway: true
    server_conf:
      CVMFS_FILE_MBYTE_LIMIT: "2048"
      CVMFS_VIRTUAL_DIR: "true"
  - name: data.dev.computeCanada.ca
    gateway: true
    server_conf:
      CVMFS_FILE_MBYTE_LIMIT: "30720"
      CVMFS_MIN_CHUNK_SIZE: "16777216"
      CVMFS_AVG_CHUNK_SIZE: "33554432"
      CVMFS_MAX_CHUNK_SIZE: "67108864"
  - name: images.dev.computeCanada.ca
    gateway: true
    gateway_keys:
      # For users to manually publish images in the store directory
      - id: store_images.dev.computeCanada.ca
        path: "/store"
      # For auto publishing by container tools
      - id: auto_images.dev.computeCanada.ca
        path: "/"
```



```
$ cat /etc/cvmfs/gateway/repo.json
{
  "repos": [
    "soft-dev.computeCanada.ca",
    "data.dev.computeCanada.ca",
    {
      "domain": "images.dev.computeCanada.ca",
      "keys": [
        {
          "id": "store_images.dev.computeCanada.ca",
          "path": "/store"
        },
        {
          "id": "auto_images.dev.computeCanada.ca",
          "path": "/"
        }
      ]
    }
  ],
  "version": 2
}
```



# Discussion