





# Outline

- PanDA - Production and Distributed Analysis Workload Management System
- Extending PanDA to Supercomputers
- PanDA for Genome Sequencing studies
- PanDA instance at NRC KI
- Results





# Production and Distributed Analysis System

- PanDA - Production and Distributed Analysis Workload Management System has been developed to address ATLAS experiment at LHC data processing and analysis challenges.
- PanDA was able to cope with increasing LHC luminosity, ATLAS data taking rate, processing and analysis challenges
  - ❖ ATLAS PanDA dedicated talk: [PanDA for ATLAS distributed computing in the next decade](#) (T.Maeno)
- Recently PanDA has been extended to run HEP scientific applications on Supercomputers
- PanDA beyond ATLAS
  - HEP and astro-particle experiments COMPASS and AMS has chosen PanDA as workload management system for data processing and analysis. nEDM with LSST will evaluate PanDA. ALICE is interested in PanDA evaluation for OLCF.
  - JINR (Dubna) is considering PanDA as main WMS for NICA collider
  - Several PanDA instances beyond ATLAS : Dubna, Moscow, Taiwan, Amazon EC2





# Extending PanDA to Supercomputers

- ATLAS is exploring use of supercomputers, HPCs and Clouds via the PanDA system
- Supercomputing centers in USA, Europe, and Asia, in particular the Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF), the National Energy Research Supercomputing Center (NERSC) in USA, Ostrava supercomputing center in Czech Republic, and "Kurchatov Institute" in Russia (NRC KI) are now integrated within the ATLAS workflow via the PanDA WMS.
- Each supercomputer is unique
  - Unique architecture and hardware
  - Specialized Operating System, "weak" worker nodes, limited memory per worker node
  - Unique job submission systems
  - Unique security environment
  - Own usage policy
- PanDA is a pilot based WMS. Using a customized pilot and pilot submission system allows to integrate opportunistic resources
  - Pilot submission to a worker node as it's done for Grid is typically not feasible. The common solution is to run one pilot per supercomputer





# PanDA for Bioinformatics

- The success of the projects to use PanDA beyond HEP (and Grid) has drawn attention from other compute intensive sciences such as bioinformatics.
- Modern biology uses complex algorithms and sophisticated software, which is impossible to run without access to significant computing resources.
- Recent advances of Next Generation Genome Sequencing (NGS) technology led to increasing streams of sequencing data that need to be processed, analyzed and made available for bioinformaticians worldwide.
- Analysis of the ancient genomes sequencing data using popular software pipeline PALEOMIX can take a months even running it on the powerful computer resources.
  - ❖ The PALEOMIX pipeline <https://github.com/MikkelSchubert/paleomix>
- Sophisticated computing software WMS and efficient usage of the supercomputers can greatly enhance this process.



# Ancient DNA studies

- Ancient DNA – is DNA extracted from ancient biological samples, such as paleontological and archaeological finds, the mummified remains, dried plant remains, coprolites.
- Ancient DNA studies allow to establish the phylogenetic relations between species and to test hypotheses about the relationship of changes in the environment and populations evolutionary changes.
- In this study we used the DNA raw reads of the Woolly mammoth lived roughly 38,000 years ago.





# PALEOMIX Pipeline

Special software pipelines are used to analyze an ancient DNA. Pipelines include a sets of software components which are allow fast NGS data processing. PALEOMIX is the one of the most popular pipeline (developed by researchers from Ludovic Orlando's group at the Centre for GeoGenetics, University of Copenhagen, Denmark).

But its use is subject to certain disadvantages:

- It takes a large amount of manual routines, often with the involvement of PALEOMIX experts.
- Specialized computing infrastructure is required
- It is time-consuming.
- For example, the of ancient DNA data analysis, can be considered up to two months, even with the use of powerful computing resources.

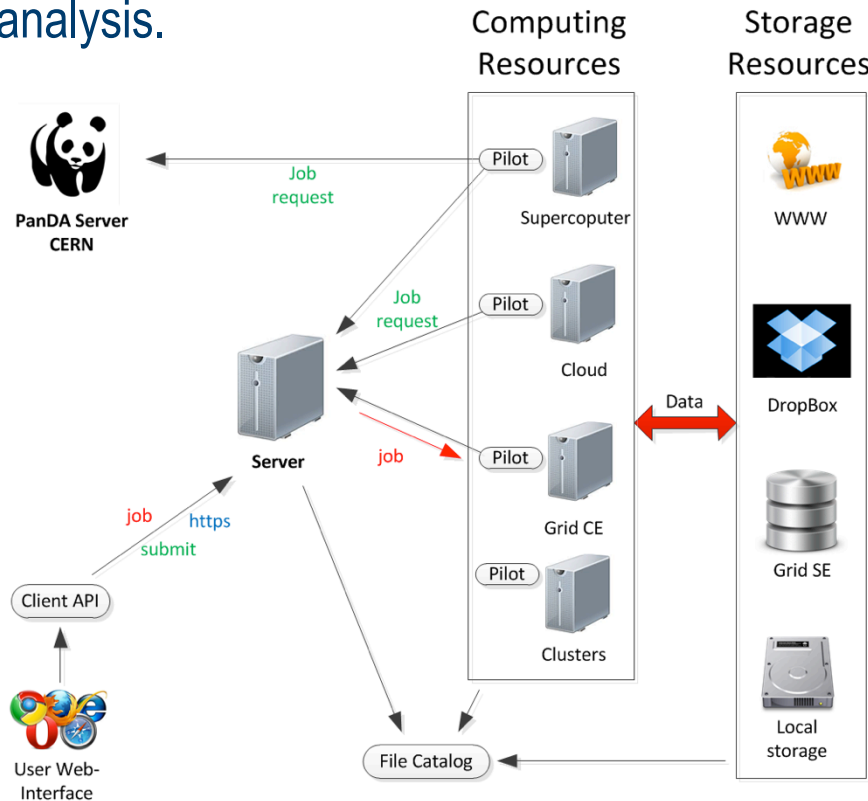
The aim of this study was to research the possibility of using the experience of distributed computing based on PanDA system in HEP to other scientific areas.





# PanDA instance at NRC KI

- The pioneer project of combining the Tier-1, the supercomputer and the cloud platform into a single portal in Kurchatov Institute has begun in 2014 and continues to the present day. This portal is used for biology studies for genome sequencing analysis.

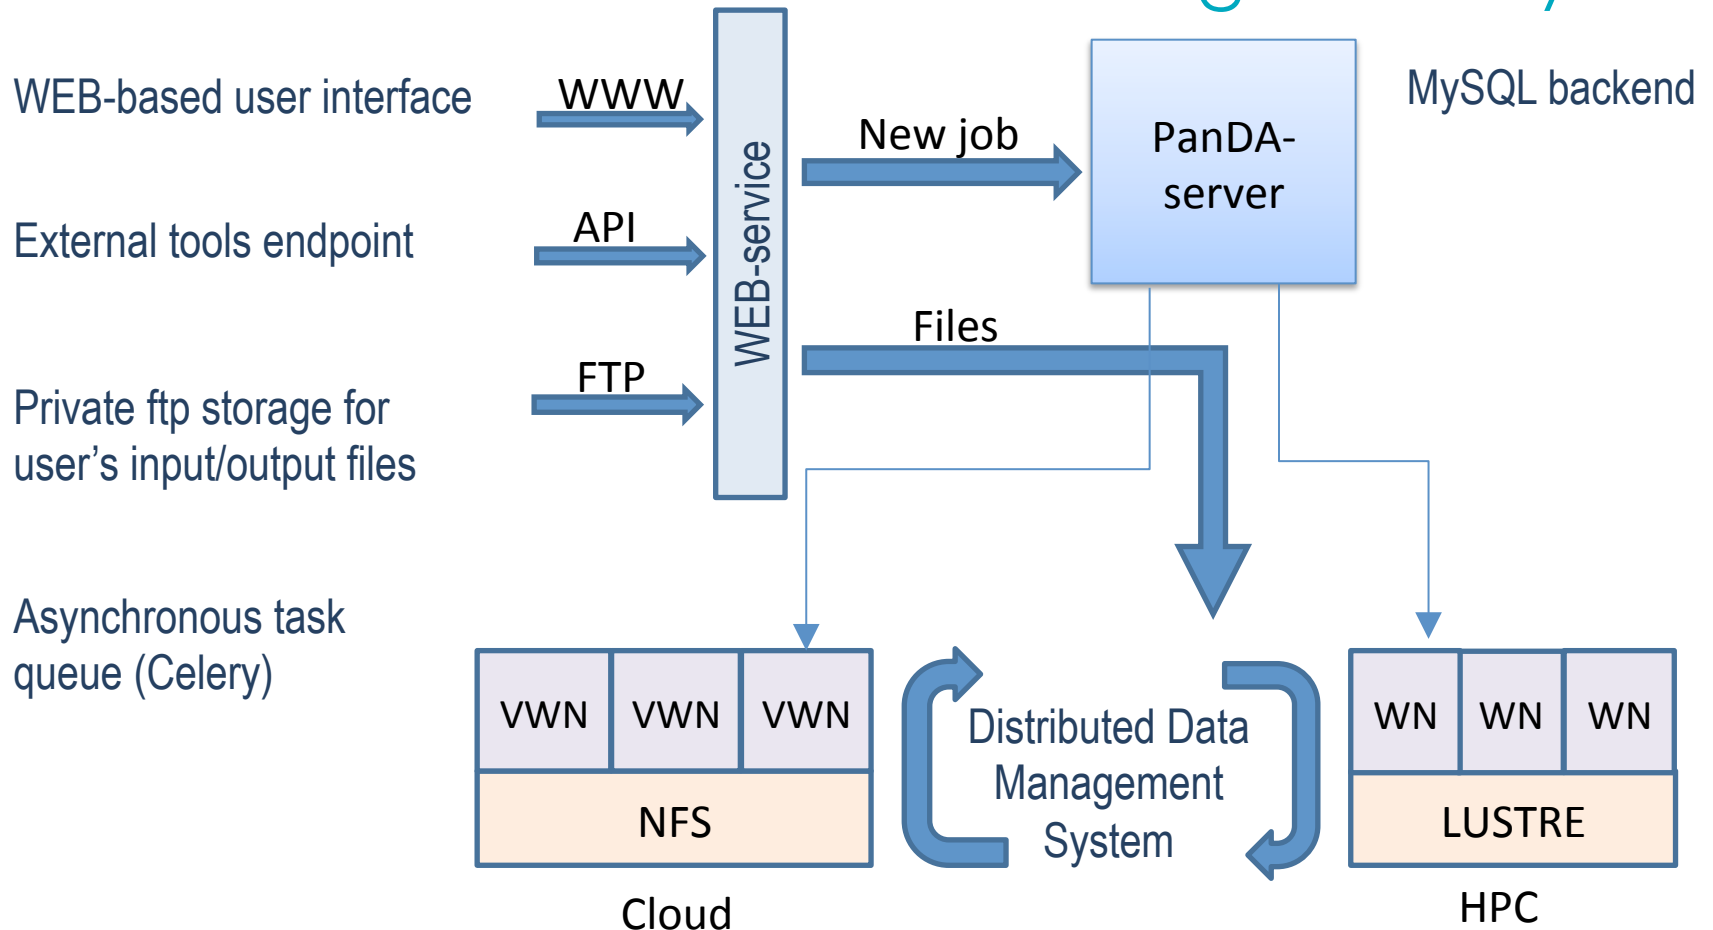


High Performance second generation cluster HPC2 with peak performance 122,9 TFLOPS (commissioned 2011) (10240 cores = 1280 nodes 2x Intel Xeon E5450 3,00ГГц 4 cores 16 GB RAM). #2 in 15-th issue of Russian top50 Supercomputers





# Authorization, API, FTP Storage, Data management system





# Web user interface

**NEW JOB**

SOFTWARE:

---

INPUT FILES: drag & drop    
 1 files ready for upload!

INPUT FILES: ftp dir

INPUT FILES: http url

INPUT FILES: guid

INPUT FILES: container    
 One file one job

---

CORES:

RUN SCRIPT: 

```
mkdir tmp; mkdir out; bam_pipeline run --max-threads=2 --jar:ool=$JAR_ROOT --temp:ool=tmp --destination=out --Mammoth.aaaaaaaaa.yaml
```

TAGS:

Update period: 5 min  
 Show  entries

ID	Owner	PandaID	Distributive	Created	Modified	Attempt	Status
2384	ruslan	2455	paleomix_bam [1.1.0]	21.03.2016 5:53	22.03.2016 3:46	1	finished
2383	ruslan	2454	paleomix_bam [1.1.0]	18.03.2016 8:23	20.03.2016 7:16	0	finished
2382	ruslan	2453	paleomix_bam [1.1.0]	18.03.2016 8:23	20.03.2016 7:16	0	finished
2381	ruslan	2452	paleomix_bam [1.1.0]	18.03.2016 8:23	20.03.2016 6:40	0	finished
2380	ruslan	2451	paleomix_bam [1.1.0]	18.03.2016 8:23	20.03.2016 5:55	1	finished
2379	ruslan	2450	paleomix_bam [1.1.0]	18.03.2016 8:22	20.03.2016 5:31	4	finished
2378	ruslan	2449	paleomix_bam [1.1.0]	18.03.2016 8:22	19.03.2016 21:16	3	failed

GUID	TYPE	LFN	LINK
web_it_78385e6b-144f-4442-b61f-47328c6d32d2	input	loxAfr3.fasta	[http]
web.ruslan_707a6ebf-27c8-4ec9-af7d-1d87f408c6aa	input	Mammoth.aaaaaaaaacs.yaml	[http]
web.it_m_113bde47-7fd7-4442-9ab2-3e9737d001b4	input	Mammoth.1.aaaaaaaaacs.fastq	[http]
web.it_m_4beb6833-24e7-490e-9e9b-a88f8358be1	input	Mammoth.2.aaaaaaaaacs.fastq	[http]
web.ruslan_86f9a9a5-32c2-4f85-8248-4c4902e1dbdc	output	results.tgz	[http] [ftp]
web.ruslan_a9bd8724-1135-43ef-a3c0-a6cb2498afc2	log	job.c3748a94-3add-4af6-b74b-7a6ba9cdc87e.log.tgz	[http] [ftp]

Tasks setup interface:

- Local authentication (OAuth 2.0)
- Distributive selection,
- Input files load,
- Parameters and output file names definition.

FTP (or other) storage for user I/O and flexible data transfer system in backend.

Jobs monitor.

Results from the detailed info page (available only for finished tasks) are links to the I/O files.






# Mammoth DNA mapping on the African elephant reference genome

Previously the woolly mammoth sequenced DNA data analysis was conducted using PALEOMIX pipeline at the 80-cores server with 512 GB RAM. More than 350 gigabytes of data, containing more than 900 million paired reads, was processed. That took about two months.

Overlapping DNA segments (reads)



Sequenced mammoth DNA

...GTTTCCGTTTGGGGGG**CGT**CCAAGGGCCCTTCCC...



Reference elephant DNA

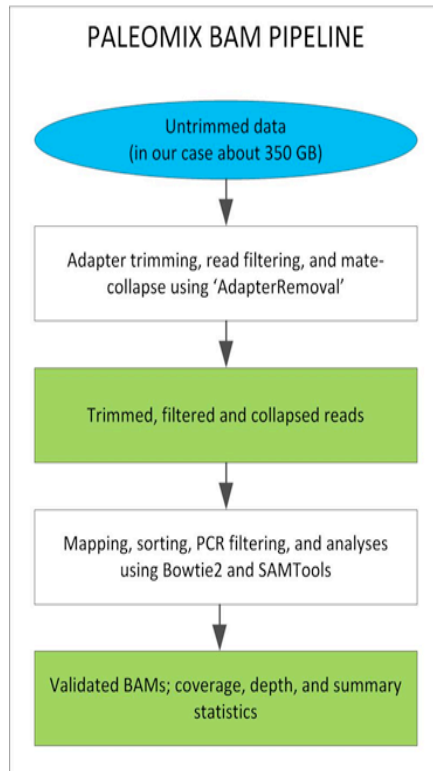
...GTTTCCGTTTGGGGGG**AGT**CCAAGGGCCCTTCCC...



Protein with one different amino acid



# Adaptation of bioinformatics tasks



Paleomix include typical set of biological software used to process NGS data including adapter trimming, read filtering, sequence alignment, genotyping and phylogenetic or metagenomic analysis.

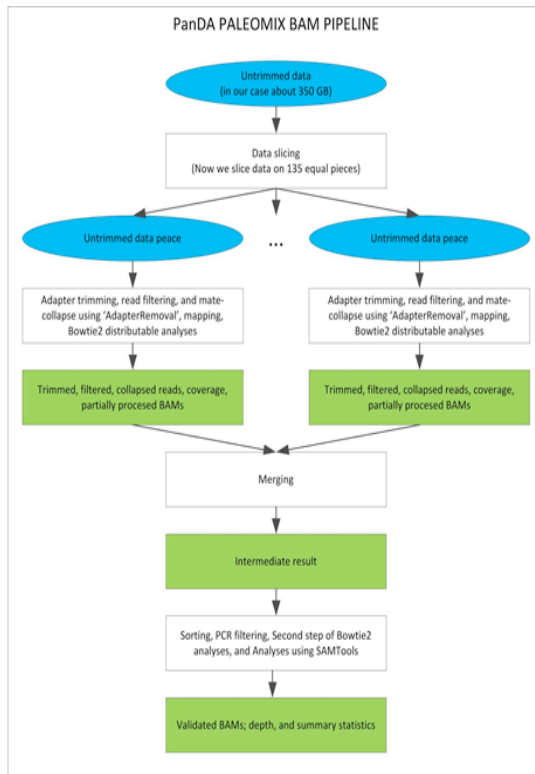
Bioinformatics requires some adaptations:

- RAM >1-2Gb per job.
- Tasks may require multicore or MPI computations.
- FTP storage is preferred over grid catalog.
- Use of Grid X.509 certificates is replaced by local authentication.
- Special software (Bowtie2, Abyss, PALEOMIX) is required





# PALEOMIX Pipeline optimization

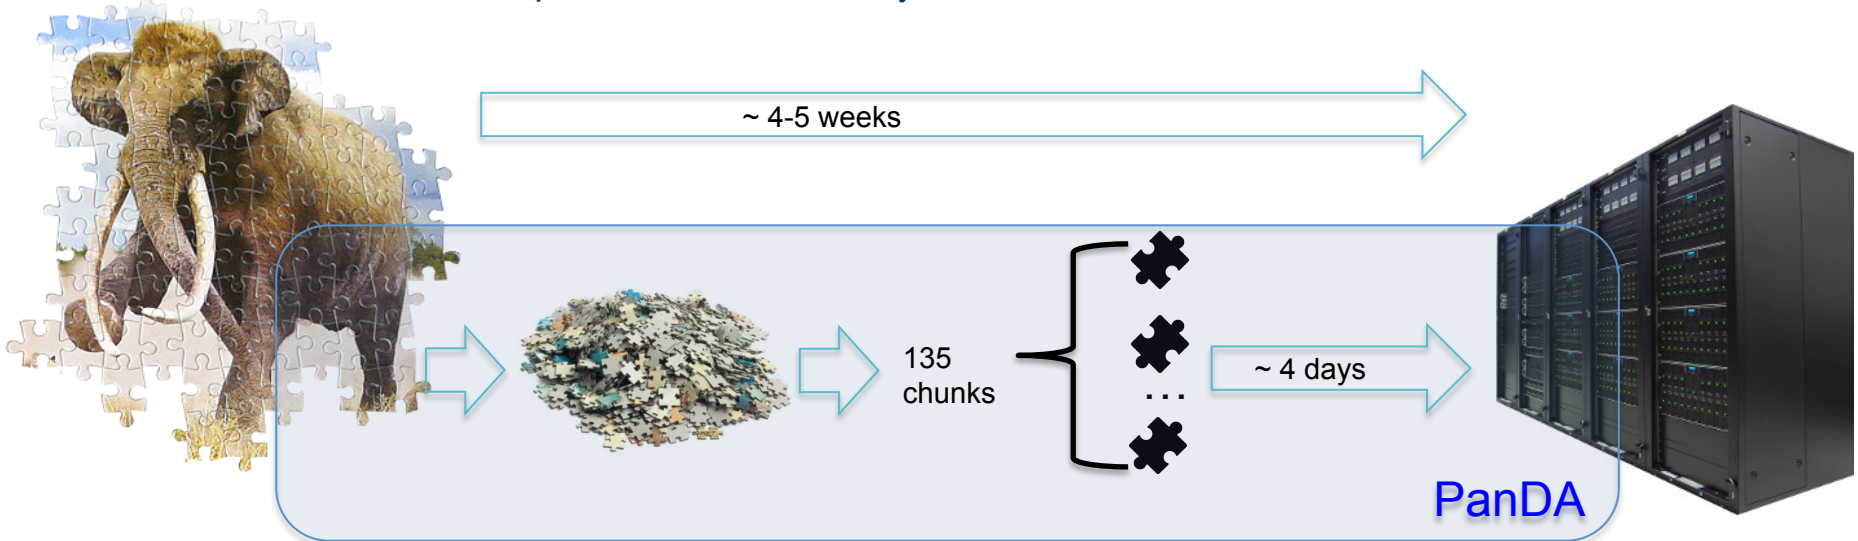


- To hide execution complexity and manual routines from end-users we introduced original pipelines control system and seamlessly integrated it into the portal.
- This system automatically (and without user prompt) split input data, prepare and run sub-tasks as ordinary PanDA jobs and merge results.
- So we assume that every pipeline contains several steps, each consists of some preparation and one or many standard PanDA jobs.



# HEP computing methods and tools for NGS studies

- The portal fits standard PanDA computational scheme and shows most efficiency for compute-intensive tasks that could be split into many jobs to be computed in parallel.
- To run pipeline we split input files into chunks which are run in parallel on different nodes as separate inputs for PALEOMIX and finally merge output file, it is very similar to what it done by ATLAS to process and to simulate data.
- Using software tools developed initially for HEP and Grid reduced payload execution time for Mammoths DNA sample from weeks to days.





# Conclusion

- We adapted the PALEOMIX pipeline to run it on a distributed computing environment powered by PanDA.
- We used PanDA to manage computational tasks on a multi-node parallel supercomputer.
- To run pipeline we split input files into chunks which are run in parallel on different nodes as separate inputs for PALEOMIX and finally merge output file, it is very similar to what it done by ATLAS to process and to simulate data.
- We dramatically decreased the total walltime because of jobs (re)submission automation and brokering within PanDA, what was earlier demonstrated for the ATLAS applications on the Grid.
- Use of the software tools developed initially for HEP and Grid reduced the total time of the Mammoths DNA analysis task from weeks to days.





# Acknowledgements

NRC "Kurchatov Institute" laboratory team work is conducted with support from Ministry of Education and Science Russian Federation, contract № 14.Z50.31.0024, and with partial support from RFBR (Russian Foundation for Basic Research) project № 16-37-00249 mol\_a.

Many thanks to PanDA SW team.

