Naples - Workload Management

WM session summary and prospective next steps in R&D and common effort

Torre Wenaus (BNL)

GDB meeting
CERN
April 11 2018

Presentations - Experiments

ATLAS, CMS

- **How HL-LHC challenges impact WM and inform potential R&D and common effort**

ARC, PanDA, DIRAC, OSG

- Asked to give perspective as a WM tool/service provider to “Roadmap towards WM standardization for HL-LHC future commonality in WM”
- Focus of this summary: less of the ‘this is X’ material in favor of ‘this is the future roadmap for X’, highlighting prospects for R&D and common effort

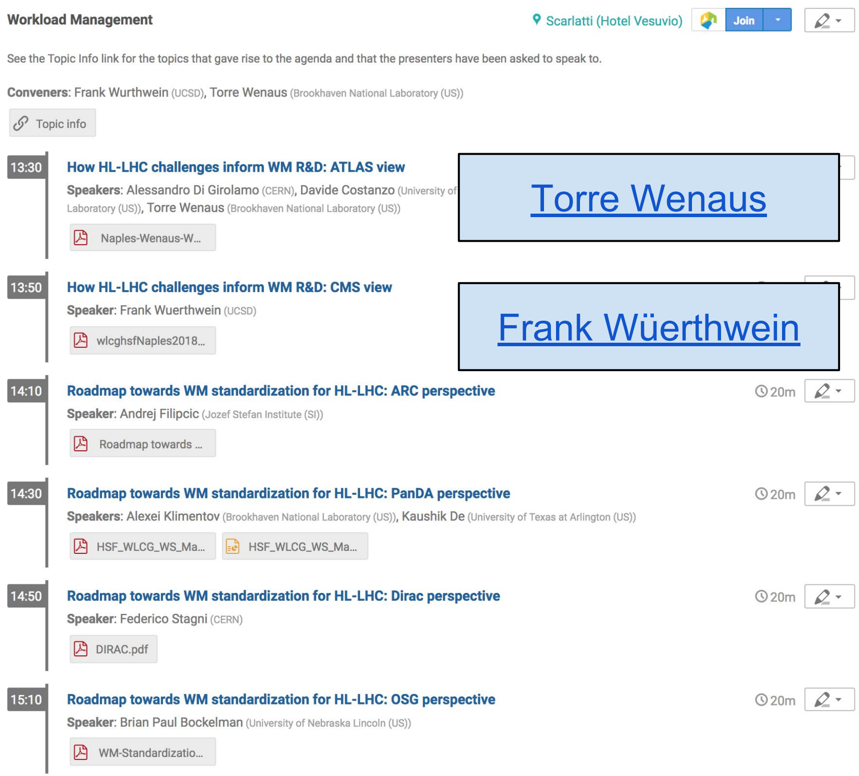
Workload Management — Scarlatti (Hotel Vesuvio) — Join

See the Topic Info link for the topics that gave rise to the agenda and that the presenters have been asked to speak to.

Conveners: Frank Wurthwein (UCSD), Torre Wenaus (Brookhaven National Laboratory (US))

Topic info

13:30 **How HL-LHC challenges inform WM R&D: ATLAS view**
Speakers: Alessandro Di Girolamo (CERN), Davide Costanzo (University of
Laboratory (US)), Torre Wenaus (Brookhaven National Laboratory (US))
Naples-Wenaus-W...

[Torre Wenaus](#)

13:50 **How HL-LHC challenges inform WM R&D: CMS view**
Speaker: Frank Wuerthwein (UCSD)
wlcghsfNaples2018...

[Frank Wüerthwein](#)

14:10 **Roadmap towards WM standardization for HL-LHC: ARC perspective** — 20m
Speaker: Andrej Filipcic (Jozef Stefan Institute (SI))
Roadmap towards ...

14:30 **Roadmap towards WM standardization for HL-LHC: PanDA perspective** — 20m
Speakers: Alexei Klimentov (Brookhaven National Laboratory (US)), Kaushik De (University of Texas at Arlington (US))
HSF_WLCG_WS_Ma... HSF_WLCG_WS_Ma...

14:50 **Roadmap towards WM standardization for HL-LHC: Dirac perspective** — 20m
Speaker: Federico Stagni (CERN)
DIRAC.pdf

15:10 **Roadmap towards WM standardization for HL-LHC: OSG perspective** — 20m
Speaker: Brian Paul Bockelman (University of Nebraska Lincoln (US))
WM-Standardizatio...

The HL-LHC computing landscape

- HL-LHC computing extrapolates to huge costs following today's practices
  - Ergo we can't follow today's practices
- Computing can't be allowed to blow the budget
  - A constrained resource, manage it as such
  - Analysts must understand the costs and tradeoffs in how resources are used: a technical and sociological challenge
  - Expectation management is vital. Understanding of what we can do, by when.
- We need innovation and development to scale
- Leveraging diverse distributed resources knit together by powerful networks
  - The foundation of LHC computing success thus far

Working assumptions

- Flat facilities budget
- ~20%/yr capacity growth
- Development effort is sustained
  - The need for this has been affirmed by ATLAS, LHC management, LHCC, resource board scrutiny, …
  - We rely on this if we're to succeed through innovation and development

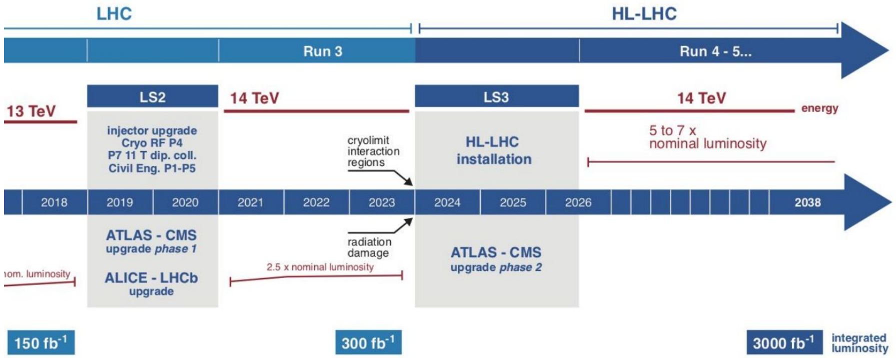

Scaling to HL-LHC: CPU

- ATLAS
  - CPU shortfall has shrunk in last 18mo thanks to tracking speedups; current estimate is ~3x (down from ~6x)
  - Ongoing work will bring further improvements
    - Further reconstruction optimizations
    - Use of truth info in MC track reconstruction
    - Pre-mixing, more fast simulation
  - Substantial WM effort in leveraging opportunistic CPU
- CMS
  - Analysis model revisions should reduce analysis CPU to <10%
  - Pre-mixing would dramatically reduce simu/digi time
  - 25% reco speedup on top of these may bring CPU within budget
- Both
  - Expect substantial use of HPCs, with *data intensive* requirements from our use cases
- Bottom line: **CPU shortfall is serious but we appear to be on track, not the biggest problem**

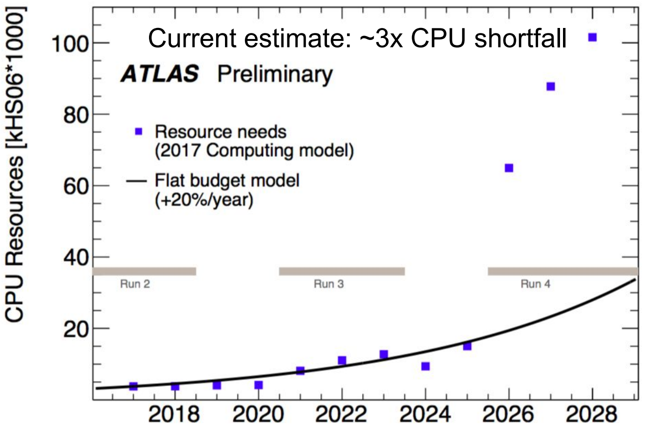

Scaling to HL-LHC: Storage

Size of different event formats covers x500
- RAW ~ 5MB
- AOD ~ 2MB
- MiniAOD ~ 200kB
- NanoAOD ~ 10kB

CMS

- CMS
  - Mini, Nano AODs common across physics analysis groups
  - Degree to which analyses can rely on compact formats, vs. needing subsets of larger formats (which rapidly dominates space), unclear
  - Virtual data not likely practical because of CPU costs
  - Storage budget won’t allow keeping hot/warm data all on disk
- ATLAS
  - ~6x shortfall estimate has held ~steady
  - ‘Opportunistic storage’ basically doesn’t exist
  - Replica counts already squeezed, hard to achieve large gains
  - Working on format size reductions, but hard to achieve large gains; ~100 analysis formats. AODs, DAODs use ⅔ of the disk today.
- Both
  - **Storage shortfall is our biggest problem.**
  - **The best bet for reducing storage budget is to rely more on tape and less on disk.**
  - Technology expectation at this point is that tape will be the cold/cheap storage option that it is today

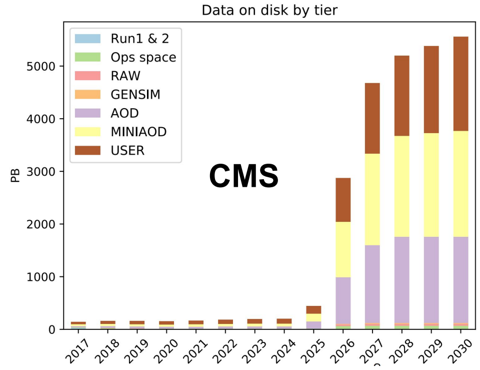

ATLAS disk ~6x short at HL-LHC

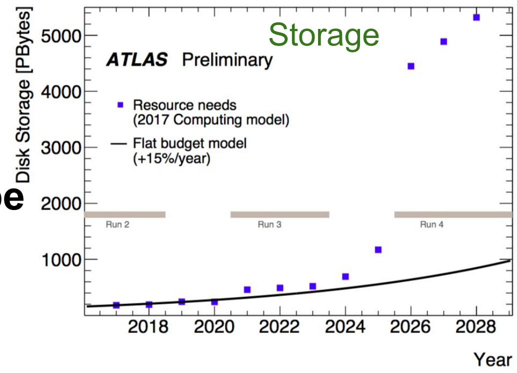

Implications of more tape

Points of ATLAS & CMS agreement:

- There will be managed production stages in the processing chain when we can expect the inputs for a campaign will not all fit on disk.
- Tape usage will need to be orchestrated by workload management and data management working in tandem
  - To reach the needed efficiency and turnaround time to meet analyst expectations
  - Tape-sourced workflows will be delivering analysis data products, so sensitive to turnaround time
  - A full (re)architecting of the workflow + storage stack with WM and DM closely interacting.
- Example: train processing in ‘carousel’ mode: a sliding window of disk-cached data with WM processing it as it appears
  - So disk-resident replica count of the input sample at any one time is <<1
- **We need to calculate, model, analyse just what such workflows imply**
- **Interested to make greater tape usage part of near term R&D**

Granular processing/streaming

- ATLAS couples its interest in tape-sourced WM+DM R&D with interest in granular processing and data streaming
- Relying more on tape facilities at O(10) institutes to serve widely distributed processing increases the importance of efficient, agile, prompt data delivery from source to client
- **Move only the data you need**, to a client ready to consume it
    - **Hide the latencies** involved from the processing
- Can achieve this with **streaming data flows**
    - Don't require big files/datasets to move from A to B before processing starts at B
    - Be **agile, asynchronous, adaptive** to current resource availability
    - Use **knowledge of the task to marshal** and send only needed data
- Enabled by **fine-grained processing** that ATLAS has been developing
    - Flexible partitioning of the work to enable **optimizing the granularity** from full files down to single events
- Granular processing: **Event Service** in early production for simulation
- Granular marshaled inputs: **Event Streaming Service** in early R&D

WM, DM and data lakes

- Tape sourced workflows being one example, WM and DM must work in close tandem for efficient use of storage and full use of CPU resources
  - Preparing data for client consumption…
    - In a data carousel
    - In transforming data optimized for storage into data marshaled for optimal delivery
      - Nice reference: "Organizing, Orchestrating, and Delivering Data From Lakes" white paper by Rob Gardner et al
- If data lakes are the context for distributed DM, they have to foresee close interaction with WM
- Designing these interactions and APIs will be important
  - e.g. in the EOS presentation, the QoS API is along these lines

WM - data lake interactions

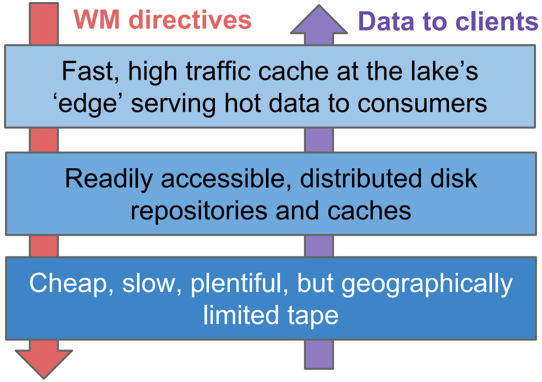

ATLAS & CMS summary: commonality & distinctions

- Agree: computing can't be allowed to blow the budget. A constrained resource. Technical and sociological challenge to expose to analysts the costs and tradeoffs of their processing.
- Agree: expectation management is vital. Understanding of what we can do, by when.
- Agree: addressing CPU shortfall demands much work, but we see paths.
- Agree: the bigger problem and challenge is disk storage.
- Agree: there will be managed production stages in the processing chain when we can expect the inputs for a campaign will not all fit on disk.
- Differ: CMS has one agreed compact analysis data product, ATLAS has ~100; if this persists, ATLAS likely will be hit by this earlier in the chain than CMS.
- Agree: The best bet for reducing storage budget is to rely more on tape and less on disk.
- Agree: tape usage will need to be orchestrated by WM & DM (eg train processing in 'carousel' mode) to reach the needed efficiency and turnaround time to meet analyst expectations. Generally, a full (re)architecting of the workflow + storage stack with WM and DM closely interacting.
- Agree: Interested to make greater tape usage part of near term R&D.
- Agree: we need to calculate, model, analyse just what such workflows imply.
- Differ: ATLAS couples this interest with granular processing/streaming R&D in the data lake context.
- Agree: common interest in data/analysis facility to explore 'declarative programming' for analysis.

Presentations - WM system perspectives

ATLAS, CMS

- How HL-LHC challenges impact WM and inform potential R&D and common effort

ARC, PanDA, DIRAC, OSG

- **Asked to give perspective as a WM tool/service provider to "Roadmap towards WM standardization for HL-LHC future commonality in WM"**
- **Focus of this summary: less of the 'this is X' material in favor of 'this is the future roadmap for X', highlighting prospects for R&D and common effort**

Workload Management — Scarlatti (Hotel Vesuvio) — Join

See the Topic Info link for the topics that gave rise to the agenda and that the presenters have been asked to speak to.

Conveners: Frank Wurthwein (UCSD), Torre Wenaus (Brookhaven National Laboratory (US))

Topic info

13:30 **How HL-LHC challenges inform WM R&D: ATLAS view** — 20m
Speakers: Alessandro Di Girolamo (CERN), Davide Costanzo (University of Sheffield (GB)), Johannes Elmsheuser (Brookhaven National Laboratory (US)), Torre Wenaus (Brookhaven National Laboratory (US))
Naples-Wenaus-W...

13:50 **How HL-LHC challenges inform WM R&D: CMS view** — 20m
Speaker: Frank Wuerthwein (UCSD)
wlcghsfNaples2018...

14:10 **Roadmap towards WM standardization for HL-LHC: ARC perspective** — 20m
Speaker: Andrej Filipcic (Jozef Stefan Institute (SI))
Roadmap towards ...

14:30 **Roadmap towards WM standardization for HL-LHC: PanDA perspective** — 20m
Speakers: Alexei Klimentov (Brookhaven National Laboratory (US)), Kaushik De (University of Texas at Arlington (US))
HSF_WLCG_WS_Ma... HSF_WLCG_WS_Ma...

14:50 **Roadmap towards WM standardization for HL-LHC: Dirac perspective** — 20m
Speaker: Federico Stagni (CERN)
DIRAC.pdf

15:10 **Roadmap towards WM standardization for HL-LHC: OSG perspective** — 20m
Speaker: Brian Paul Bockelman (University of Nebraska Lincoln (US))
WM-Standardizatio...

ARC - Advanced Resource Collector

Andrej Filipcic

Grid middleware stack developed by Nordugrid collaboration since 2002

- Open Source, mostly volunteer contributors

Key components:

- ARC Compute Element with support for most of the batch system, Boinc, SCEAPI
  - Job control
  - Data transfers and shared cache management
- Client tools:
  - Jobs submission
  - Proxy management
  - File copy
- API: python libraries providing interface to full software stack
  - Custom services
  - Custom clients, eg arcControlTower

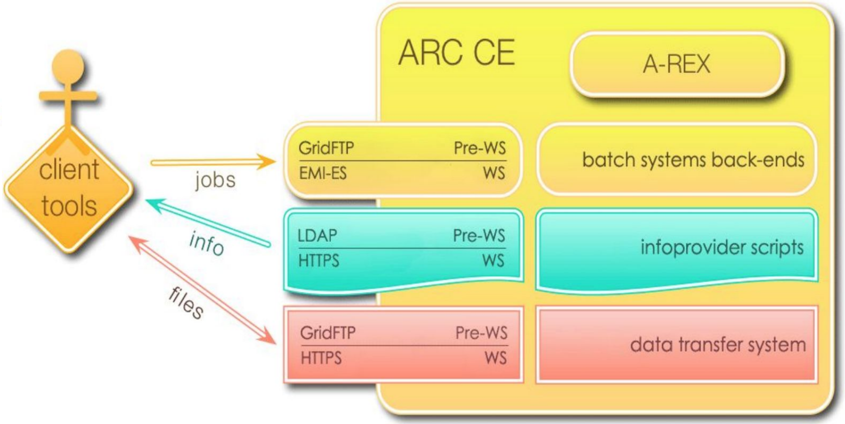

ARC

Resource vs Workload management

- ARC-CE
  - interface to the batch system
  - Used by different tools, eg Dirac, Condor-C, ...
- arcControlTower (aCT)
  - Workload management system interfacing to workflow management system (PanDA)
  - Submission to a (large) set of CEs (similar to Condor-C + schedd)
  - Part of harvester effort

ARC-CE - general purpose

aCT - tuned for ATLAS, but can be adapted to any WFMS

Harvester + aCT submission backend

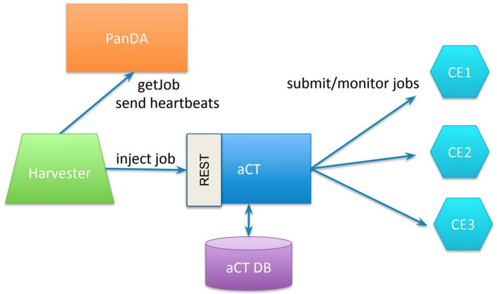

Using the ARC cache

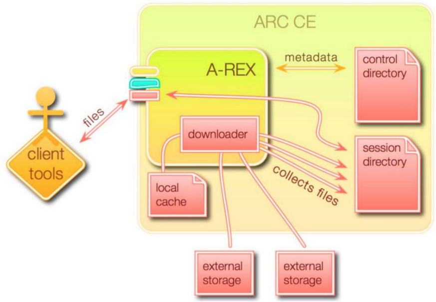

ARC Roadmap

- ARC release 6 coming out in few months - major revision
  - Simplifying the configuration
  - REST submission interface
  - Event driven engine
  - ARCHERY - dns based hierarchical endpoint registry, resource discovery
  - Transition to git
- Future development
  - Extending caching to support other technologies, eg xrootd/xcache
  - Native container handling
  - Alternatives to X509 (scitokens prototype in development)
  - aCT: extending to other user communities

ARC

Bold = potential for R&D, common effort

- Roadmap points to
  - Streamlining: simpler config, REST, git
  - **resource discovery, information system**
  - **caching extensions**
  - Native containers
  - **alternatives to X509**
  - **extending aCT to other user communities**
  - **advanced, dynamic workflows**
    - **input, output streaming**
    - **dynamic optimization of task mixture**
    - **Harvester as common (with PanDA) approach**
  - **supporting resource diversity**
    - **metrics and info gathering**
  - **(sub)event level workflows**

PanDA

Alexei Klimentov

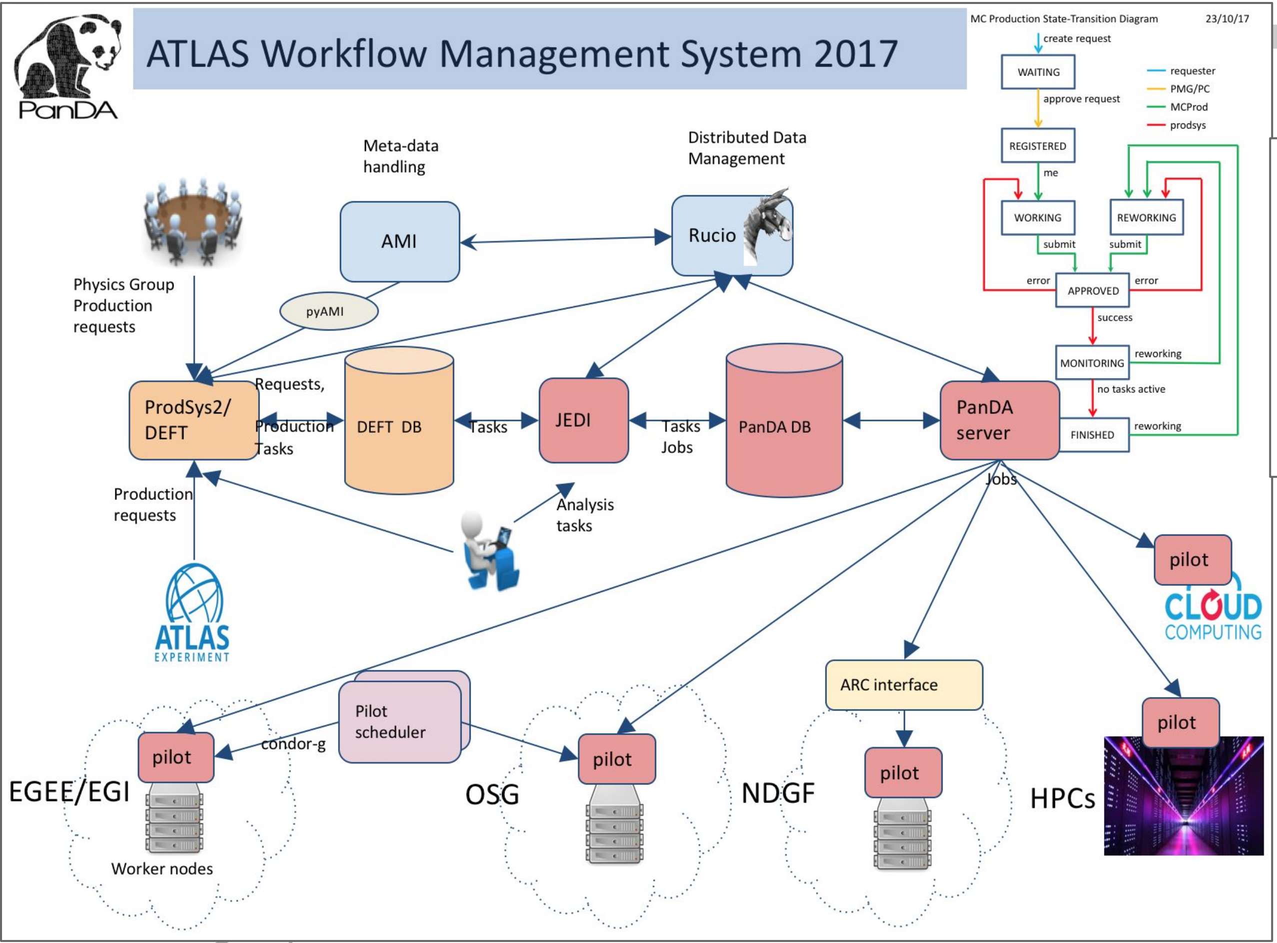

Global ATLAS operations

Up to ~1.2M concurrent jobs

25-30M jobs/month at >250 sites

~1400 ATLAS users

First exascale workload manager in HENP

1.3+ Exabytes processed in 2014 and in 2016

Exascale scientific data processing today

Key features

- Central job queue
  - Unified treatment of distributed resources
  - SQL DB keeps state of all workloads
- Pilot based job execution system
  - Payload is sent only after execution begins on CE
  - Minimize latency, reduce error rates
- Fairshare or policy driven priorities for thousands of users
- Automatic error handling and recovery
- Extensive monitoring
- Modular design

Future Challenges

- New physics workflows
  - also new ways how Monte-Carlo campaigns are organized
- New strategies
  - "provisioning for peak"
- Integration with networks (via DDM, via IS and directly)
- Data popularity -> event popularity
- Address new computing model
- Address future complexities in workflow handling
  - Machine learning and Task Time To Complete prediction
  - Monitoring, analytics, accounting and visualization

WMS – DDM coupled optimizations

- WMS will evolve to enable new data models
- Data lakes, data ocean, caching services, SDN, DDN,…
- Another level of granularity (from datasets to events)

PanDA

Future development. Harvester

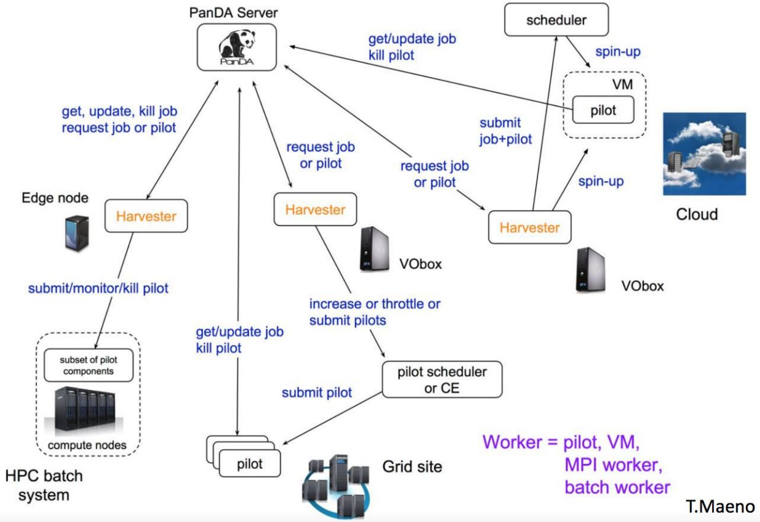

- To address wide spectrum of computing resources/facilities available to ATLAS and experiments in general
- New model : PanDA server- harvester-pilot
- The project was launched in Dec 2016

Primary objectives :

- To have a common machinery for diverse computing resources
- To provide a common layer in bringing coherence to different HPC implementations
- To optimize workflow executions for diverse site capabilities

Harvester Status

- Architecture designed and implemented
- Harvester for cloud
  - In production : CERN+Leibniz+Edinburgh resources (1.2k CPU cores
  - Work in progress : HLT farm @ LHC Point1, Google Cloud Platform
- Harvester for HPC
  - In production :
    - Theta/ALCF, Titan (OLCF)
    - ASGC (non-ATLAS Vos)
    - Cori+Edison / NERSC
    - KNL@BNL
- Harvester for Grid
  - Core SW is ready
  - Many scalability test are planned in 2018 before commissioning
    - harvester is currently running at BNL (~800 jobs). Migration to full scale production is ongoing at BNL

PanDA

Bold = potential for R&D, common effort

- Roadmap points to
  - **close integration with key (distinct) systems: info system, DDM system(s), metadata system(s)**
  - **reliance on key middleware (HTCondor)**
  - **Harvester: common, coherent, optimized provisioning for diverse resources**
  - **integration with networks**
  - new physics workflows, MC organization
  - **address new computing model**
  - **ML, analytics**
  - **granularity, data streaming**
  - **WM support for co-processor capabilities and workflows that utilize them**
  - **tighter WMS-DDM coupling: data lakes, fine grained processing (datasets => events)**
  - **data lakes R&D (Google 'data ocean')**

DIRAC: The Interware

[Federico Stagni]

- A software framework for distributed computing
- A **complete** solution to one (or more) user community
- Builds a layer between users and resources

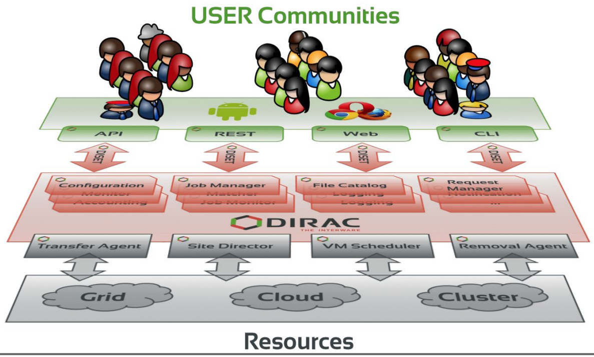

- A framework shared by multiple experiments/projects, both inside HEP and in astronomy, and life science
- Experiment agnostic
- DIRAC Consortium
- Installations:
  - Single-VO
  - Multi-VO (DIRAC as-a-service)
    with limited functionalities

3

"Horizontal" extensibility

-

For specific requirements

WebAppDIRAC | RESTDIRAC | VMDIRAC | COMDIRAC

DIRAC — Core project

Pilot

Externals

Each project is independently versioned

A DIRAC release is composed by all the projects (strong dependency)

"Vertical" extensibility

-

Community driven

VO DIRAC

VO WebAppDIRAC

WebAppDIRAC

DIRAC

DIRAC

- Computing
  - **CEs**: ARC, CREAM, HTCondor, "SSH" for standalone BS, ...
  - **Batch**: LSF, BQS, SGE, PBS/Torque, SLURM, Condor,...
  - Clouds, BOINC, HPC, "desktops"
- Storage
  - SRM2, GSIFTP, XRoot, http, DIPs, …
    - ECHO, CTA, ...
- Catalog
  - DIRAC FC, LFC, LHCb Bookkeeping, ….
- Information services
  - BDII, GOCDB, ...
- DBs, MQs, LogBackends
  - MySQL, Oracle, ElasticSearch
  - stomp
  - file, MQ (stomp -- ActiveMQ, RabbitMQ), ES

- DIRAC supports multiple grid middlewares and infrastructures (EGI, WLCG, OSG, NorduGRID, etc)
  - Other types of grids can be supported
    - As long we have customers needing that
- Standalone clusters behind a BS
  - access through SSH/GSISSH tunnel
    - a really thin layer that we call "SSH CE"

- CLOUD resources
- VAC resources
- BOINC Volunteer resources
- HLT farm (LHCb)

DIRAC THE INTERWARE

Summary

- DIRAC Pilots are the real federator of "any" computing resource
- Used by all the DIRAC communities in every DIRAC installation
  - Single or multi-VO
- Flexible: highly configurable, easy to extend
- Actively developed

Vision, and roadmap, won't deviate while approaching Run3 and HL-LHC

8th DIRAC Users' Workshop: 22-25 May, Lyon

DIRAC

Bold = potential for R&D, common effort

- Roadmap points to
  - **Ongoing broad multi-experiment support through Consortium**
  - **Support other types of grids/fabrics as they appear and are interesting to DIRAC users**
  - **Dirac pilots as the federator of “any” computing resource**
    - Support for CE free resources via ssh
    - Support for CE & BS free resources via “Make a machine a pilot machine”, e.g. pilot launch in VM contextualization
  - **Clean, agnostic, extensible API**
  - **“Vision, and roadmap, won’t deviate while approaching Run3 and HL-LHC”**

OSG

Brian Bockelman

- Workflow management in OSG follows the overlay model:
  - User defines a workflow in a VO-specific workflow scheduling component.
  - “Factory” requests resource allocations from a distributed set of computing resources to meet workflow needs. When started, these new resources are unassigned and go into a central pool of free resources.
  - Workflow scheduling component examines available resources and assigns workflows according to some policy.

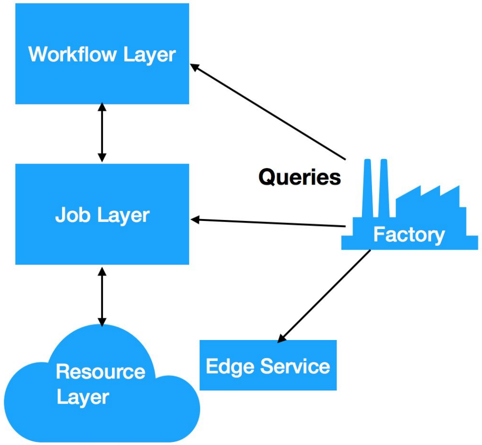

This is a bit of the classic view, which is beginning to evolve as the “resource layer” starts to include more non-grid resources.

- Within the OSG, there are two common setups for the overlay model:
  - PanDA serves as the workflow manager. AutoPyFactory serves as the resource factory.
  - HTCondor serves as the job manager (workflow layer can also be in HTCondor or elsewhere). GlideinWMS serves as the resource factory.
- The resources involve tend to be allocated from batch clusters in chunks of 1 core (to N cores, but all on one node).

Technology-Free Talk

OSG

Baseline Capacity

- Through some combination of Brownian motion and detailed planning, we've arrived at some concept of a compute service:
  - Run as a site-level service.
  - Remotely accessible.
  - Some mechanism for remote authorization / authentication.
  - Bulk submit of jobs (job descriptions), query, cancel.
  - Movement of input and output sandboxes.
  - Transfer of credentials.
- Fits well with our "classic grid" model. Shouldn't this be a starting point for discussing new resources?

Evolving beyond baseline

- We see hints of the baseline breaking down at the large HPC sites:
  - Sites are disinterested in running services for an experiment (maybe not even a community?)
  - No ability for pilots to be "pushed" to the edge service.
  - No ability for pilots to "pull" payloads due to network access restrictions.
  - Need to pre-stage data or pre-setup storage environment (configure burst buffers prior to job launch).
- In some cases, a complete breakdown: non-job-based models for backfilling HPCs.
- The community sees utilizing HPC as a core part of the HL-LHC challenge! So, how do we tackle this?

OSG

Words of the day:
standardization
common APIs
common approaches
common projects

- Brian took a simultaneously lower level and higher level perspective on what the realistic prospects are for commonality, and where those reside in the stack
- The realistic prospects are at the lower level, and we should realize them
- Our community should approach new sites with one voice, particularly in terms of providing an experiment-independent baseline interface to a computing resource
- Shouldn't a common compute service be a baseline starting point for discussing new resources?
    - Large HPCs have to this point been the breaking point for such a baseline
    - How do we, with our HPC friends, tackle this?
- We should tackle greenfield projects together, especially as we look at new site functionality
- At higher levels, realistically we should look for common approaches, common capabilities; maybe not standardization
- Let the community pick best-of-breed

What's the forum for tackling this?

WM system talks: summary impressions

- All have strong, ongoing communities; none of them is going away
- No visible prospects for merging systems, at the macro level at least
  - But close joint efforts do exist
    - e.g. ARC - PanDA, now extending into Harvester commonality
- No sign that standardizing on one WM system for the HL-LHC era is a likely prospect
- Best prospect: common approaches, common capabilities
  - Maybe common elements in achieving those
  - Particularly new/emerging elements
  - Particularly at the lower level resource interface
    - Addressing HPC complexity and heterogeneity is a common issue, should address it in common

Prospective R&D and common efforts

From the ATLAS, CMS talks

- Tape sourced workflows
  - Analyzing requirements, implications, attainable capability / performance
  - WLCG has a WG addressing tape
- Close coupling between DM and WM in the data lakes / DOMA R&D
  - Make WM an explicit part of the R&D
  - API / CDN requirements from WM
- Data marshalling at the storage / in the lake
  - Proposed in Rob Gardner et al white paper
- Data streaming, granular processing, WAN access, latency hiding
- New analysis models based on declarative programming interfaces to remote cloud/data center analysis resources
- How much do we do via ROOT
  - In latency hiding, data marshaling & packaging, compression management, caching, new analysis models …
- Across it all: good metrics and deep analytics
  - And leveraging ML in analytics

Prospective R&D and common efforts

More from the WM system talks

- Uniform, experiment-independent baseline interface to a computing resource
  - Including large HPCs
  - Especially greenfield sites, and/or adding new site functionality
  - Integrating new resources, grids, … as they appear
- Scalable, uniform resource provisioning, the federator of “any” computing resource
  - Harvester interest extends to at least two of them (PanDA, ARC)
  - At a minimum… common means of submitting pilot wrappers to sites
- Metrics and analytics
  - A ‘resource value’ metric? e.g. ATLAS toying with an ‘ATLAS Coin’ notion
    - One day one CPU (2GB of RAM) is 1 ATLAS Coin
    - Metric to meaningfully convey real costs and tradeoffs of how resources are used; make analysts aware
    - An aspect of managing expectations
- Gathering and serving information and metadata
  - eg. site capabilities (memory, CPU power, LAN and WAN throughput, I/O power, …)

More prospective R&D and common efforts

- Integrating with and leveraging networks
  - Cost matrix based on throughput performance, factored into brokerage decisions
  - Active network controls driven dynamically by WM, DM (the SDN dream)
- Support for co-processor capabilities and workflows that utilize them
  - Utilization has been low so attention has been low
- Expanding commonality in foundation middleware
  - HTCondor…
- Alternatives to X509
  - Eager for the work on this to bear fruit

Conclusion

- As in all things HL-LHC, it's too early to have a good understanding of the WM problem and the technology toolkit we'll have to address it
- But not too early to give attention to the likely problem areas
- Commissioning prior to Run 4 in ~2025 is only 7 years away!
- CPU is probably on track as long as we keep smart and keep working
- Storage looks to be the bigger problem
- If tape remains the cold cheap storage option it is today (as we expect), then using more tape in more sophisticated ways is part of the solution
- Close coupling between DM and WM will be required
- WM roadmap likely doesn't include standardization on one system
    - Commonality possible at the new component/capability level
- Plenty of work to do, plenty of R&D prospects, some in concert with DM
    - Some with clear near/mid term benefits
    - Next step, get some new R&D concretely defined and launched

Thank you

- Thank you to co-chair Frank, and all the speakers and participants in the Naples presentations and discussions
- Thank you to a great many in ATLAS Distributed Computing (ADC) for inputs and discussions, and to Simone Campana, Ian Bird
- Draws on parts of the emerging “WLCG Strategy towards HL-LHC” by Ian Bird, Simone Campana et al
- For interesting reading very much aligned with ideas presented here see the “Organizing, Orchestrating, and Delivering Data From Lakes” white paper by Rob Gardner et al