



Reflections on Migrating IBM APP Genomic Workflow Acceleration to POWER8

Chandler Wilkerson <chwilk@rice.edu>

HPC Cluster Architect, **Rice University**

#OpenPOWERSummit



 **OpenPOWER™ Summit 2015**

San Jose, CA | March 17-19

About Us

Rice University Research Computing Support Group

Five (and two halves) member team within Rice IT

Seven clusters including IBM POWER7 and POWER8

PowerOmics HW

POWER8 nodes:

- (6) 8247-22L 2 X 12 core P8
- (2) 8247-42L 2 X 12 core P8 + NVIDIA K40 Teslas

Interconnect:

- 36-port FDR IB switch
- Next gen Connect-IB FDR HCAs
- 10-Gigabit Ethernet

Storage:

- IBM GPFS Storage Server (GSS24)

PowerOmics SW

IBM PowerKVM 2.1.0

Ubuntu 14.4 (little-endian)

Red Hat Enterprise Linux 7.0

IBM Platform LSF, PPM, PAC

Mellanox OFED 2.4-1

GPFS 4.1

BioBuilds 2014.11

PowerOmics Layout

One master node with VMs

- Management Node

- PAC/PPM Portal Node

- Red Hat Login Node

- Ubuntu Login Node

Four Red Hat EL 7.0 nodes

Three Ubuntu 14.04.02 nodes

NFS /home and /project space

GPFS /poscratch clustered file system

3D Genome sequencing

Genomics is hard enough in 2D

Discerning the 3D structure of the genome reveals interactions of genes that otherwise seem unrelated

3D Genome sequencing adds a long analysis step to the end of “traditional” sequencing

Genomic Workflow

Utilizes IBM Platform LSF, PPM, PAC

Assorted tools from BioBuilds

Custom software written in Perl, Python, CUDA, C++...

Lots of shell scripts

Genomics Workflow

For each subject's sequencer output (split into two files)

Do

- Split sequenced data into computable chunks

- Run alignment utility on each chunk (parallel)

- Merge aligned chunks

- Output some metrics

- Run large process on merged, aligned chunks

The Problem

This works reasonably for a single sample.

Scaling up slows down the job scheduler prohibitively

Think 20K-40K jobs with dependencies on each other.

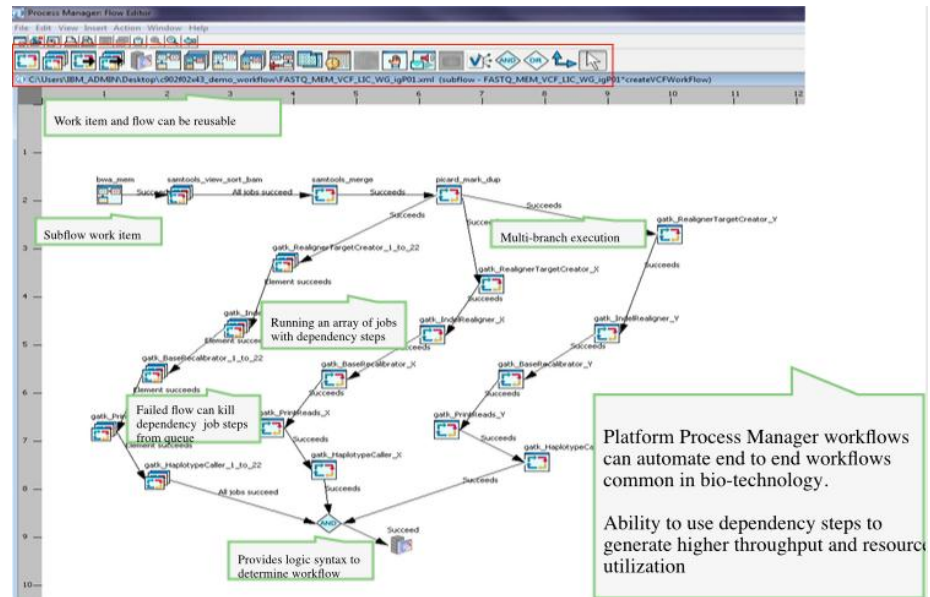
Sometimes things go wrong... Those 10Ks of jobs don't erase themselves.

Managing the Workflow

Platform Process Manager

GUI Interface to create graph based workflow

Coordinates processes in workflow with LSF

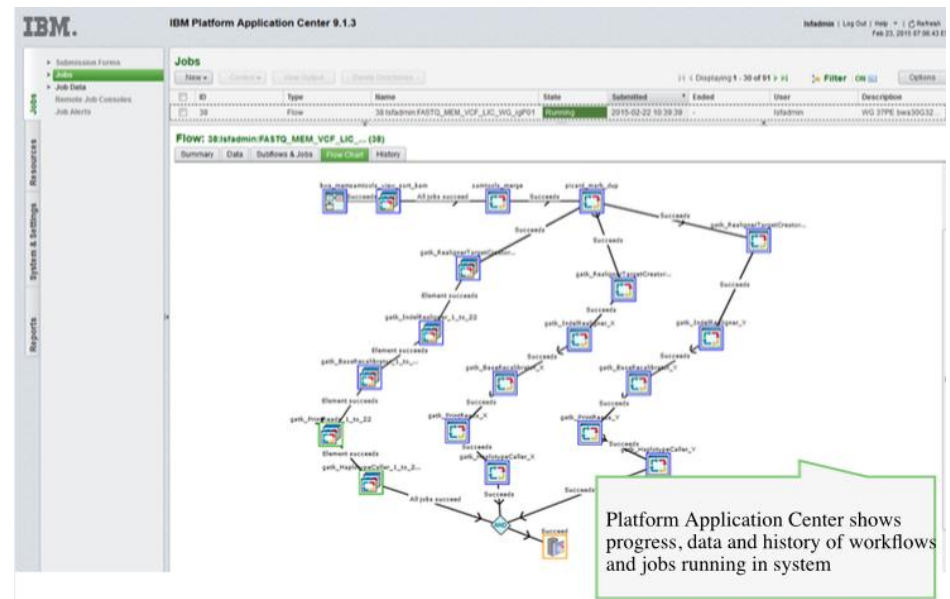


Managing the Workflow

Platform Application Center

Provides a web portal interface to PPM workflows

Allows visual monitoring of workflows



IBM. IBM Platform Application Center 9.1.3

Jobs

ID	Type	Name	State	Submitted	Ended	User	Description
38	Flow	381stadminFASTQ_MEM_VCF_LIC... (38)	Running	2015-02-22 10:39:38		1stadmin	WQ 37PE 1wa30032

Flow: 381stadminFASTQ_MEM_VCF_LIC... (38)

Summary Data Outflows & Jobs Flow Chart History

Platform Application Center shows progress, data and history of workflows and jobs running in system

Admin's Perspective (the good)

OPAL mode (firmware setting) is really cool.

Petitboot saved my OS (twice)

Genomics workflows are rough on clusters

Data transfer delays make for idle clusters

Admin's Perspective (the bad)

OPAL/PowerKVM cannot be managed from HMC

Enumerating ASMs is still a manual process

ASM alone is still not as powerful as other vendors'

PowerKVM hasn't caught up with all our needs yet.

Future Work

FPGA based compression acceleration coming soon

Next generation Teslas

Eagerly awaiting NVlink

Would like to see more work with PowerKVM