Our Experience Supporting Thousands of Community Applications using Containers

Matthew Vaughn
Director, Life Sciences Computing
Texas Advanced Computing Center

July 11, 2017

Contributors

Dr. John Fonner

GitHub & Bioconda integration

Jacquelyn Turcinovic

Continuous integration + on-demand container pipeline

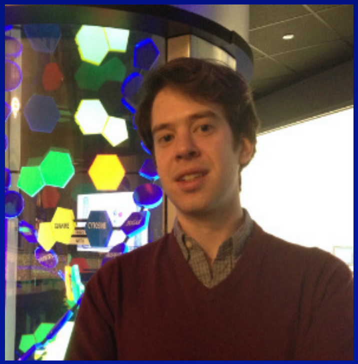

Greg Zynda

Singularity runtime build, configuration, and testing

Dr. Jawon Song

User environment & module system integration

TACC operates a professionally-supported national science cloud

A complex and shifting landscape

- Workflows evolve quickly
  - Half life < 6 months
- Constant exploration around methods, parameterization, and interpretation
- Critical to enhancing *velocity of iteration*
- Collaborative development is a must
  - Reproducible deployment and analysis
  - Facilitate “replaceable” components
- Users ~~want~~ need dozens of packages

Supporting Packages is Hard Work

- Download source/vendor binaries
- Package via RPM build system
  - Affordances for compilers, MPI stacks, heterogeneous environments, modules implementation, and more
- Manually test & request installation
- Admin performs installation

RPM is optimized for simple binary distributions

- Version update - ~60 min effort. Or hours…

Headwinds

- New code written on Debian or BSD laptops
- Bundles of high-level language code (plus modules), binaries written in new languages, shell scripts, plus entire Gnu/Linux user land
- Bootstrapping these is hard for professionals!
  - “Install it yourself” isn’t feasible (or desirable)
    - Suboptimal configuration
    - Privileged dependencies
    - GLIBC (Grrr!)
    - End user support

TACC Systems Landscape

Sandy Bridge, Ivy Bridge, Haswell, Skylake, Knights Landing, Knights Corner

RHEL6, RHEL7, SUSE11, CLE

InfiniBand, Aries, OPA

GPFS, Lustre

Compilers + 1000 packages

How can we scale our support for new packages x architectures?

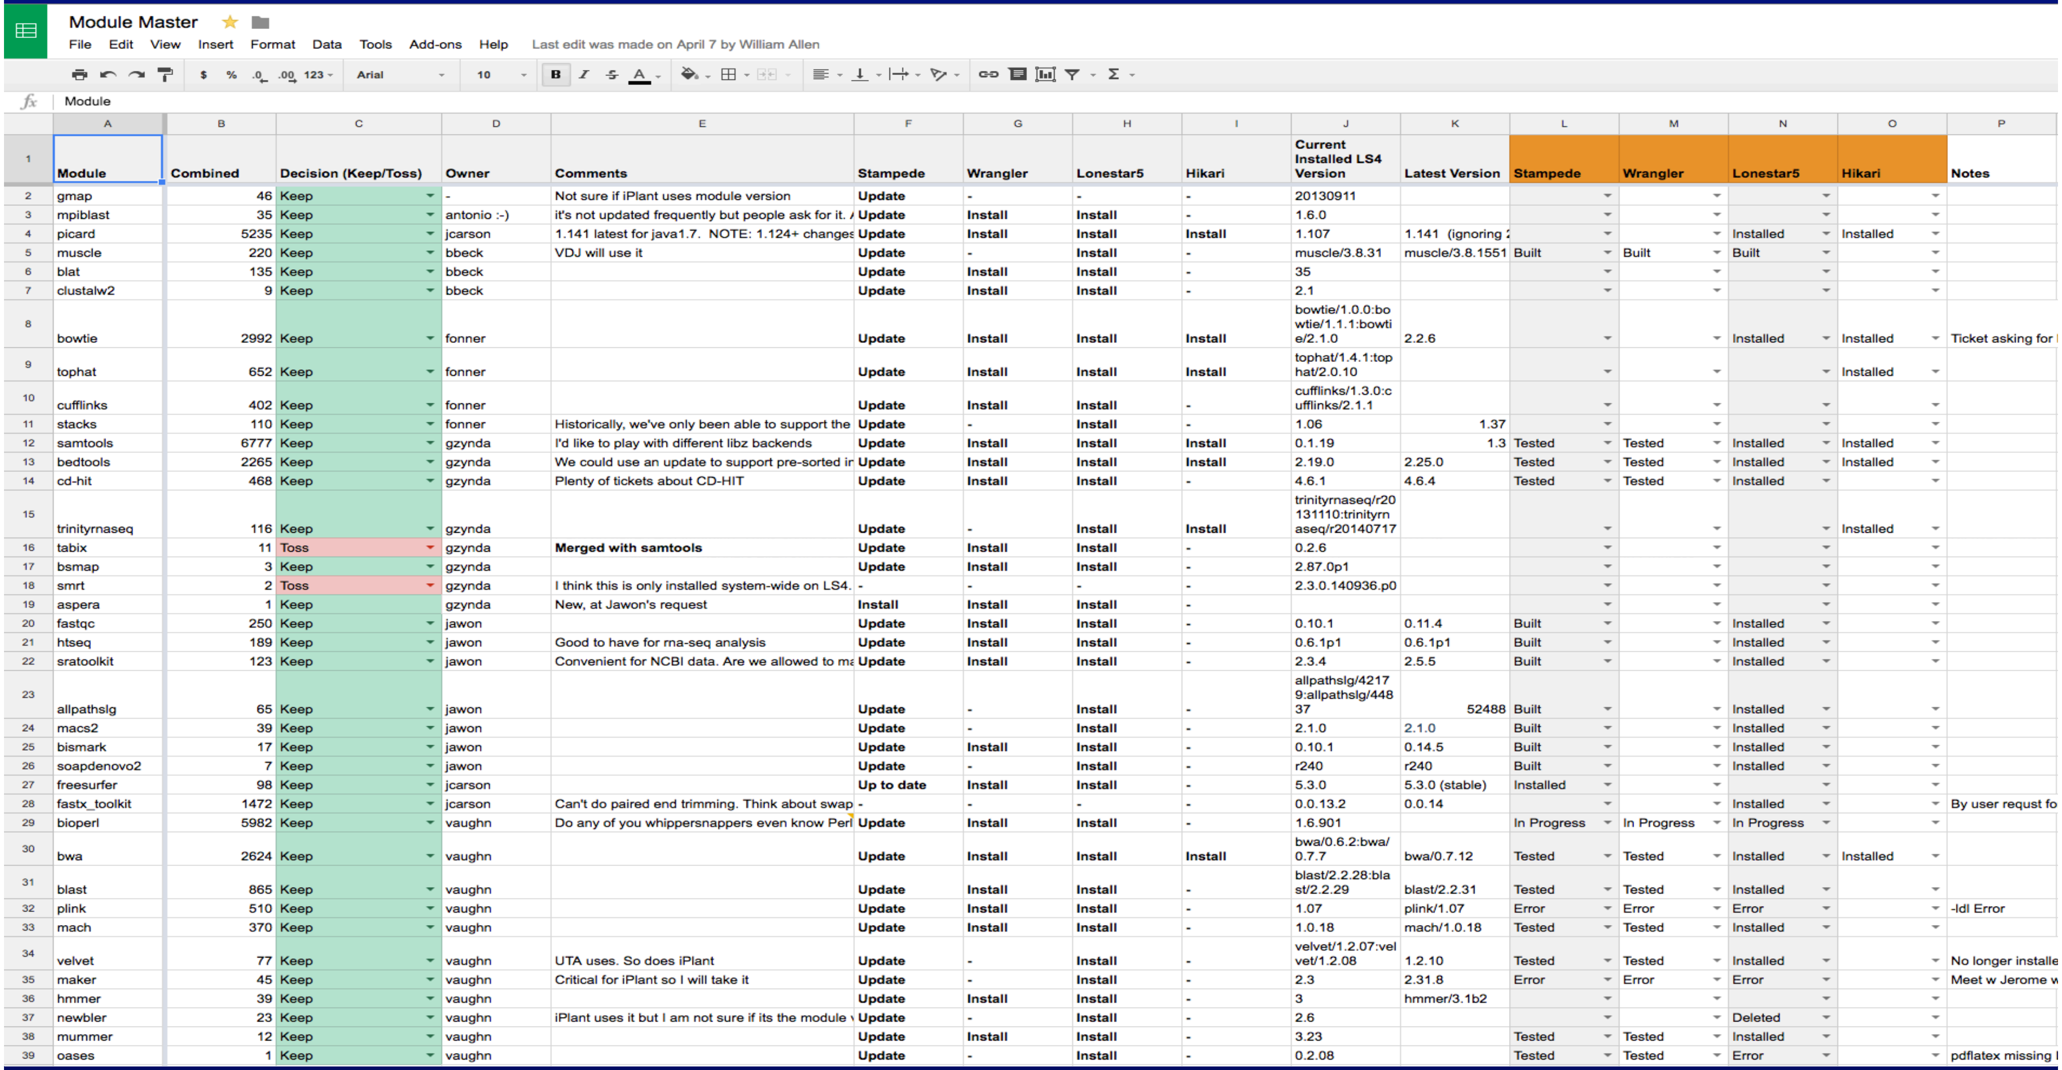

Last edit was made on April 7 by William Allen

fx Module

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1	**Module**	**Combined**	**Decision (Keep/Toss)**	**Owner**	**Comments**	**Stampede**	**Wrangler**	**Lonestar5**	**Hikari**	**Current Installed LS4 Version**	**Latest Version**	**Stampede**	**Wrangler**	**Lonestar5**	**Hikari**	**Notes**
2	gmap	46	Keep	-	Not sure if iPlant uses module version	**Update**	-	-	-	20130911						
3	mpiblast	35	Keep	antonio :-)	it's not updated frequently but people ask for it. /	**Update**	**Install**	**Install**	-	1.6.0						
4	picard	5235	Keep	jcarson	1.141 latest for java1.7. NOTE: 1.124+ changes	**Update**	**Install**	**Install**	**Install**	1.107	1.141 (ignoring			Installed	Installed	
5	muscle	220	Keep	bbeck	VDJ will use it	**Update**	-	**Install**	-	muscle/3.8.31	muscle/3.8.1551	Built	Built	Built		
6	blat	135	Keep	bbeck		**Update**	**Install**	**Install**	-	35						
7	clustalw2	9	Keep	bbeck		**Update**	**Install**	**Install**	-	2.1						
8	bowtie	2992	Keep	fonner		**Update**	**Install**	**Install**	**Install**	bowtie/1.0.0:bowtie/1.1.1:bowtie/2.1.0	2.2.6			Installed	Installed	Ticket asking for
9	tophat	652	Keep	fonner		**Update**	**Install**	**Install**	**Install**	tophat/1.4.1:tophat/2.0.10					Installed	
10	cufflinks	402	Keep	fonner		**Update**	**Install**	**Install**	-	cufflinks/1.3.0:cufflinks/2.1.1						
11	stacks	110	Keep	fonner	Historically, we've only been able to support the	**Update**	-	**Install**	-	1.06	1.37					
12	samtools	6777	Keep	gzynda	I'd like to play with different libz backends	**Update**	**Install**	**Install**	**Install**	0.1.19	1.3	Tested	Tested	Installed	Installed	
13	bedtools	2265	Keep	gzynda	We could use an update to support pre-sorted ir	**Update**	**Install**	**Install**	**Install**	2.19.0	2.25.0	Tested	Tested	Installed	Installed	
14	cd-hit	468	Keep	gzynda	Plenty of tickets about CD-HIT	**Update**	**Install**	**Install**	-	4.6.1	4.6.4	Tested	Tested	Installed		
15	trinityrnaseq	116	Keep	gzynda		**Update**	-	**Install**	**Install**	trinityrnaseq/r20131110:trinityrnaseq/r20140717					Installed	
16	tabix	11	Toss	gzynda	**Merged with samtools**	**Update**	**Install**	**Install**	-	0.2.6						
17	bsmap	3	Keep	gzynda		**Update**	**Install**	**Install**	-	2.87.0p1						
18	smrt	2	Toss	gzynda	I think this is only installed system-wide on LS4.	-	-	-	-	2.3.0.140936.p0						
19	aspera	1	Keep	gzynda	New, at Jawon's request	**Install**	**Install**	**Install**	-							
20	fastqc	250	Keep	jawon		**Update**	**Install**	**Install**	-	0.10.1	0.11.4	Built		Installed		
21	htseq	189	Keep	jawon	Good to have for rna-seq analysis	**Update**	**Install**	**Install**	-	0.6.1p1	0.6.1p1	Built		Installed		
22	sratoolkit	123	Keep	jawon	Convenient for NCBI data. Are we allowed to ma	**Update**	**Install**	**Install**	-	2.3.4	2.5.5	Built		Installed		
23	allpathslg	65	Keep	jawon		**Update**	-	**Install**	-	allpathslg/42179:allpathslg/44837	52488	Built		Installed		
24	macs2	39	Keep	jawon		**Update**	-	**Install**	-	2.1.0	2.1.0	Built		Installed		
25	bismark	17	Keep	jawon		**Update**	**Install**	**Install**	-	0.10.1	0.14.5	Built		Installed		
26	soapdenovo2	7	Keep	jawon		**Update**	-	**Install**	-	r240	r240	Built		Installed		
27	freesurfer	98	Keep	jcarson		**Up to date**	**Install**	**Install**	-	5.3.0	5.3.0 (stable)	Installed				
28	fastx_toolkit	1472	Keep	jcarson	Can't do paired end trimming. Think about swap	-	-	-	-	0.0.13.2	0.0.14			Installed		By user requst fo
29	bioperl	5982	Keep	vaughn	Do any of you whippersnappers even know Perl	**Update**	**Install**	**Install**	-	1.6.901		In Progress	In Progress	In Progress		
30	bwa	2624	Keep	vaughn		**Update**	**Install**	**Install**	**Install**	bwa/0.6.2:bwa/0.7.7	bwa/0.7.12	Tested	Tested	Installed	Installed	
31	blast	865	Keep	vaughn		**Update**	**Install**	**Install**	-	blast/2.2.28:blast/2.2.29	blast/2.2.31	Tested	Tested	Installed		
32	plink	510	Keep	vaughn		**Update**	**Install**	**Install**	-	1.07	plink/1.07	Error	Error	Error		-ldl Error
33	mach	370	Keep	vaughn		**Update**	**Install**	**Install**	-	1.0.18	mach/1.0.18	Tested	Tested	Installed		
34	velvet	77	Keep	vaughn	UTA uses. So does iPlant	**Update**	-	**Install**	-	velvet/1.2.07:velvet/1.2.08	1.2.10	Tested	Tested	Installed		No longer installe
35	maker	45	Keep	vaughn	Critical for iPlant so I will take it	**Update**	-	**Install**	-	2.3	2.31.8	Error	Error	Error		Meet w Jerome w
36	hmmer	39	Keep	vaughn		**Update**	**Install**	**Install**	-	3	hmmer/3.1b2					
37	newbler	23	Keep	vaughn	iPlant uses it but I am not sure if its the module v	**Update**	-	**Install**	-	2.6				Deleted		
38	mummer	12	Keep	vaughn		**Update**	**Install**	**Install**	-	3.23		Tested	Tested	Installed		
39	oases	1	Keep	vaughn		**Update**	-	**Install**	-	0.2.08		Tested	Tested	Error		pdflatex missing

Don’t reinvent the wheel

- Similar problems faced by engineers in the Internet-scale economy. Their solutions often leverage a cloud mindset.
  - If it fails, swap it out with an identical copy
  - Modify capacity via horizontal scaling
  - Provision *on demand* rather than *on schedule*

EXTRAORDINARY DEGREE OF AUTOMATABLE REPRODUCIBILITY

Automating Reproducibility

- Packaging systems
  - RPM, Deb, APK
- Configuration management
  - Chef, Puppet, Ansible
- Continuous integration systems
  - Jenkins, Travis, CircleCI
- Virtualization
  - VMWare, OpenStack, Commercial IaaS
- Linux Containers

Linux Containers

- Linux supports several obscure features
  - virtual networks
  - process control groups (cgroups)
  - logical volume mapping
  - chroots, overlay, and loopback filesystems
  - security profiles
  - iptables
- Individually useful. Packaged as *Linux Containers* (LXC). Extended by Docker, Inc.

Docker for Research Software

- Delivery organized around registries
  - DockerHub
  - Quay.io
  - Private registries
- Integrated with source management and CI
- Name-spaced conventions for discovery & sharing
- Build instructions via plaintext template
- Orchestration via Swarm, K8S, etc.

Docker Drawbacks for HPC

- Requires “modern” 3.10+ kernel
- Some features need additional specialized capability
- Docker assumes root privileges
- HPC can support multitenant Docker
  - User namespaces + TLS + secrets escrow
- But… if the goal is abstracting software why bother until it’s easier (if at all)?

Singularity (Singularity.lbl.gov)

- No root owned daemon processes
- No user contextual changes or root escalation allowed (ever)
- Encapsulation of the environment
- Containers are image-based
- Facile integration with Docker ecosystem

Singularity 2.1 has been in production use on Stampede 1 since 10/2016 but only for one-off cases…

Singularity^2

- Every* supported package is a container
- Integrate with the modules system for discoverability and interoperability with native installed codes
- Automate as much as possible of the build and test process
- Leverage author- or community-built containers for codes. *Over-ride when needed*.

Crowdsourcing Research Software

- Conda
  - Package, dependency and environment management for any language
- BioConda
  - Conda package manager specializing in Bioinformatics. 2000+ Github recipes + build system that turns them into Conda packages.
- BioContainers
  - Community framework that maintains BioConda packages as Docker containers

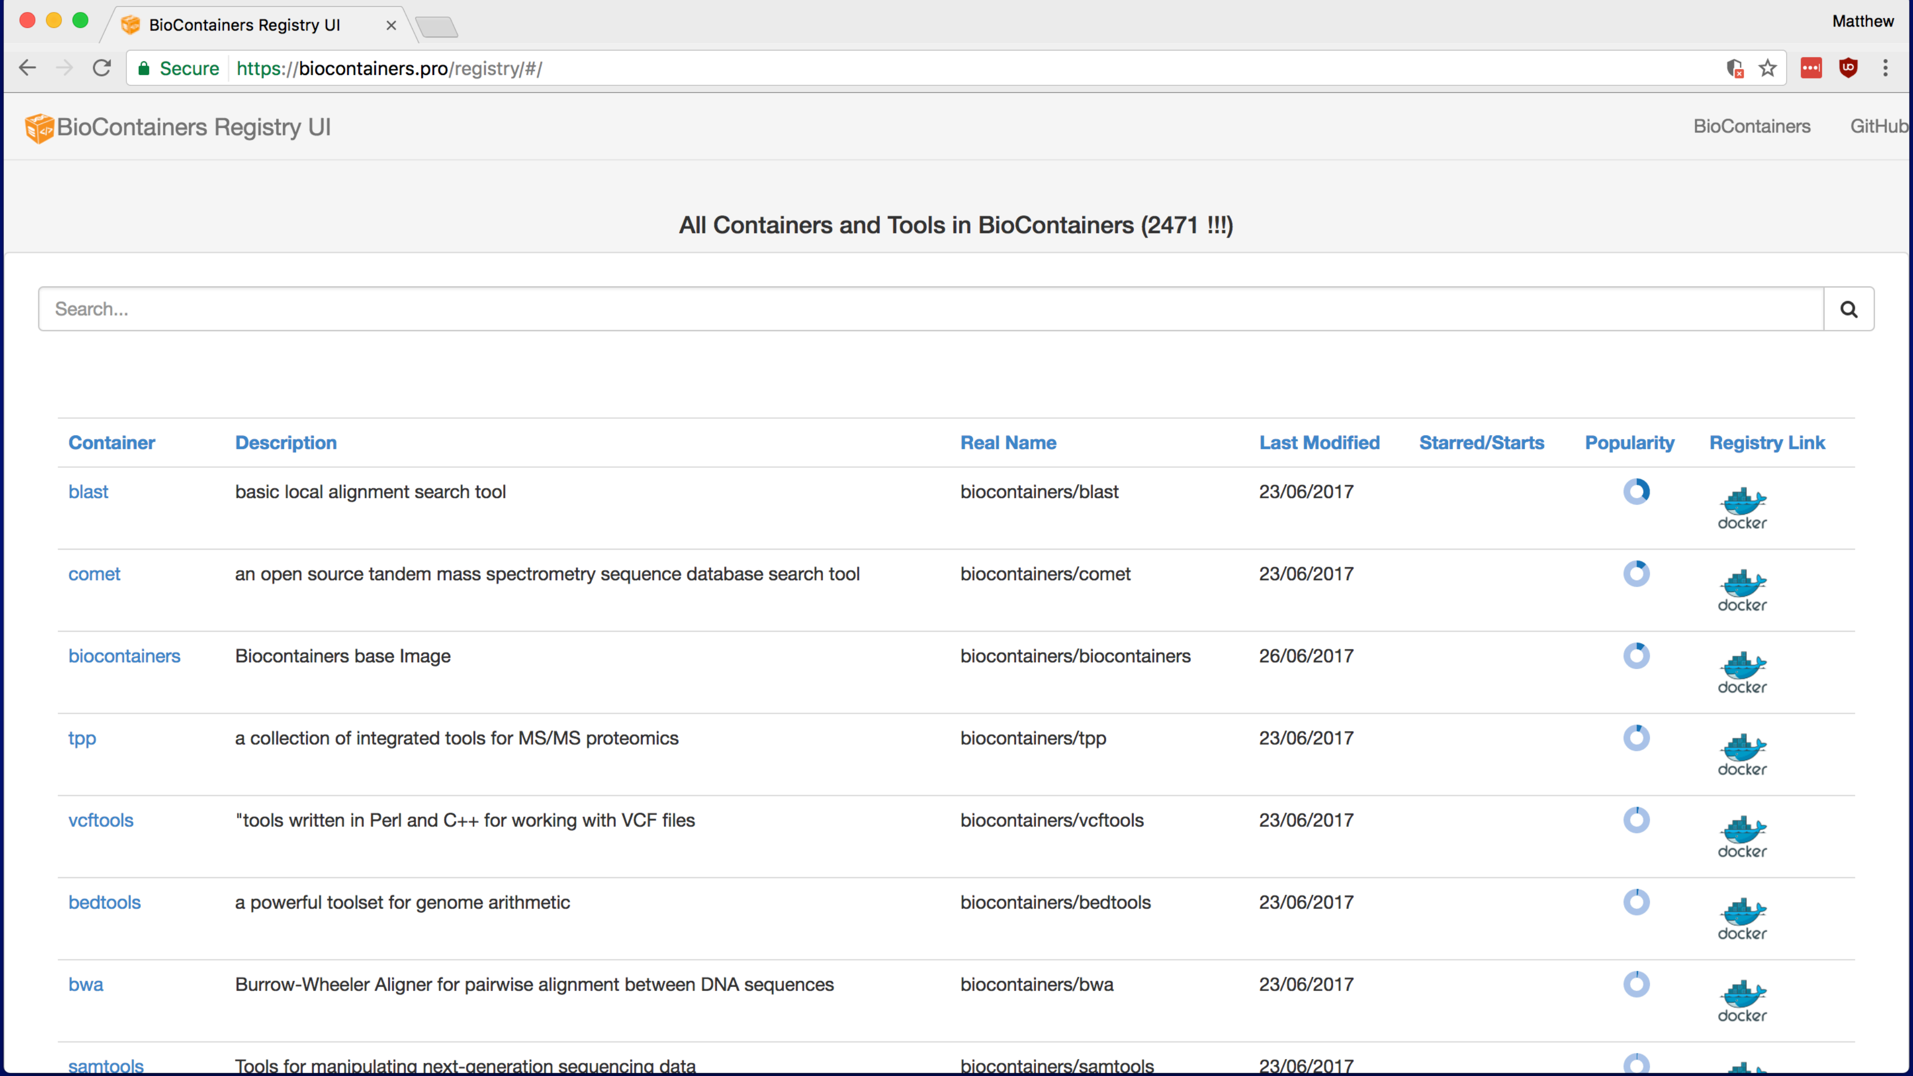

All Containers and Tools in BioContainers (2471 !!!)

Search...

Container	Description	Real Name	Last Modified	Starred/Starts	Popularity	Registry Link
blast	basic local alignment search tool	biocontainers/blast	23/06/2017			docker
comet	an open source tandem mass spectrometry sequence database search tool	biocontainers/comet	23/06/2017			docker
biocontainers	Biocontainers base Image	biocontainers/biocontainers	26/06/2017			docker
tpp	a collection of integrated tools for MS/MS proteomics	biocontainers/tpp	23/06/2017			docker
vcftools	"tools written in Perl and C++ for working with VCF files	biocontainers/vcftools	23/06/2017			docker
bedtools	a powerful toolset for genome arithmetic	biocontainers/bedtools	23/06/2017			docker
bwa	Burrow-Wheeler Aligner for pairwise alignment between DNA sequences	biocontainers/bwa	23/06/2017			docker
samtools	Tools for manipulating next-generation sequencing data	biocontainers/samtools	23/06/2017			

Target User Experience

```
ml singularity
ml biocontainers/bwa/0.7.5

singularity run $TACC_BWA

Program: bwa (alignment via Burrows-
Wheeler transformation)
Version: 0.7.15-r1140
Contact: Heng Li lh3@sanger.ac.uk

Usage:   bwa <command> [options]


…
```

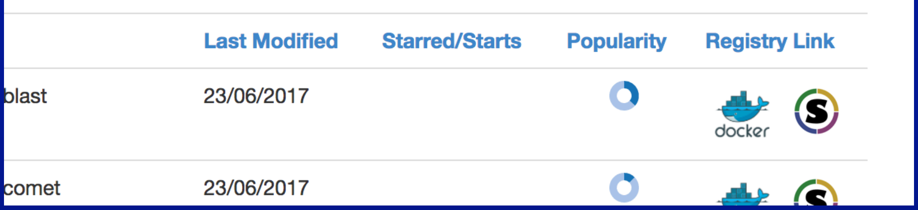

	Last Modified	Starred/Starts	Popularity	Registry Link
blast	23/06/2017			docker
comet	23/06/2017			

```
singularity pull
shub://biocontainers/bwa:0.7.5
```

- Selected container images from multiple registries cached world-readable on Stockyard
- Images are continually built and maintained
  - TACC-side CI pipeline now
  - Collaborating with Biocontainers to make this automatic
- Select package/versions marked by TACC for support via modules system
  - Explicitly tested and supported
  - Others via discovery CLI
- Modules automatically built for all TACC systems
- Support (via Jetstream) for on-demand container image builds via TACC web services

Speed Bumps

- Variable kernel capabilities
- Lack of key kernel modules (Cray)
- Mounted filesystems sans OverlayFS
- Bulky image size (and lack of local build support)
- Shifting security vulnerability landscape
- Socializing users to new workflow(s)

Outcomes

- Use nearly any public Docker image
  - Bioinformatics & computational biology
  - Machine & deep learning
  - Data analytics applications
  - Bespoke applications!
- Develop locally. Package with Docker. Share & use nearly anywhere
  - Increase collaborativity
  - Increase velocity of exploration

Thanks!

Discussion…