GenomeHubs: Life Science Informatics & HPC

HubBub 2013 Conference - http://hubzero.org


I m Chris."
"
I m an infrastructure geek."
"
I work for the BioTeam."

BioTeam

Who, What, Why ...

- **Independent consulting shop**
- **Staffed by scientists forced to learn IT, SW & HPC to get our own research done**
- **10+ years bridging the gap between science, IT & high performance computing**
- **We get to see how many groups of smart people tackle similar problems**


ATTENTION
TRY THIS TRAIL ONLY IF YOU ARE IN TOP PHYSICAL CONDITION,
WELL CLOTHED AND CARRYING EXTRA CLOTHING AND FOOD.
MANY HAVE DIED ABOVE TIMBERLINE FROM EXPOSURE.
TURN BACK AT THE FIRST SIGN OF BAD WEATHER.
WHITE MOUNTAIN NATIONAL FOREST

DISCLAIMER.

Why I m here

Not a Hub Expert. Simply a Bio/IT Nerd

- **One of my clients decided to deploy HubZero last year**
  - I was involved in initial eval and source / VM based installs
- **Met HZ team virtually while:**
  - Hacking Submit framework to support the Grid Engine job scheduler
  - Building a chef cookbook to automate Hub installs onto Amazon EC2 node images
- **HubBub 13 folks invited me to talk life science HPC and where Hubs can play a role**
- **May also be a token Corporate/Industry/ Enterprise science type**

Goal: “Talk Fast & Get Out Of The Way”

Quick Sprint 1
Life Science Informatics

Basic Bio-IT Landscape 2
Overview of our apps and requirements

Hub Stories 3
War stories from our eval, testing & deployment

<Time Permitting> 4
Observations, additional details & maybe a promise

QUICK SPRINT: LIFE SCIENCE INFORMATICS

Life Science Informatics

- **Biology and Computer Science are a natural fit**
  - Atoms, molecules & complex structures such as DNA and Protein easily characterized, stored and represented in ways that computers can read & understand
- **We have 3 core problems:**
  - Compute: studying biological data requires sophisticated algorithms, advanced statistical methods and vast amounts of raw computational power
  - Storage: Our field is drowning in petabytes of data. Inexpensive lab instruments now routinely *generate terabyte volumes of data per experiment*
  - Information Management & Triage: Rate at which we are generating new data exceeds rate at which the storage industry is increasing drive capacity. Something has to give.
- **And one big meta problem**
  - More details on that later ...

The Cliche Example

“Next Generation” Genome Sequencing

- **Advances in genome sequencing methods are outpacing Moore s Law**
- **Since the end of the Human Genome Project, cost-per-base of DNA sequencing has dropped 10,000-fold**
- **What took an intense global research effort years and billions of $ can now be done with a $50,000 instrument sitting on a desktop**

Illumina MiSeq

“Personal Sequencer”

Meta Issue:
It s a scary time to be doing Bio-IT

Big Picture / Meta Issue

This is the driver behind many “Bio-IT” design efforts ...

- **HUGE revolution in the rate at which lab platforms are being redesigned, improved & refreshed**
  - Example: CCD sensor upgrade on that confocal microscopy rig just doubled storage requirements
  - Example: The 2D ultrasound imager is now a 3D imager
  - Example: Illumina HiSeq upgrade just doubled the rate at which you can acquire genomes. Massive downstream increase in storage, compute & data movement needs
- **For the above examples, do you think IT was informed in advance?**

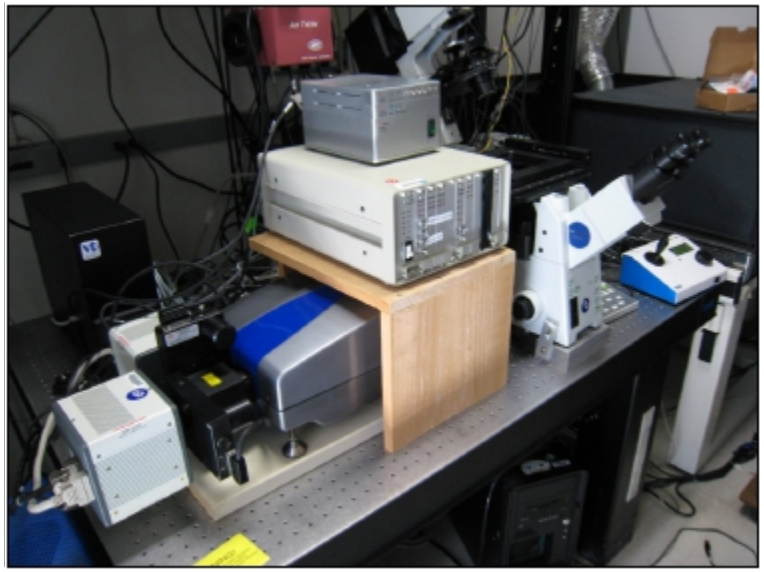

The Central Problem Is ...

Science progressing way faster than IT can refresh/change

- **Instrumentation & protocols are changing FAR FASTER than we can refresh our Research-IT & Scientific Computing infrastructure**
  - Bench science is changing month-to-month ...
  - ... while our IT infrastructure only gets refreshed every 2-7 years
- **We have to design systems TODAY that can support unknown research requirements & workflows over many years (gulp ...)**

The Central Problem Is ...

- The easy period is over
- 5 years ago we could toss inexpensive storage and servers at the problem; even in a nearby closet or under a lab bench if necessary
- That does not work any more; real solutions required

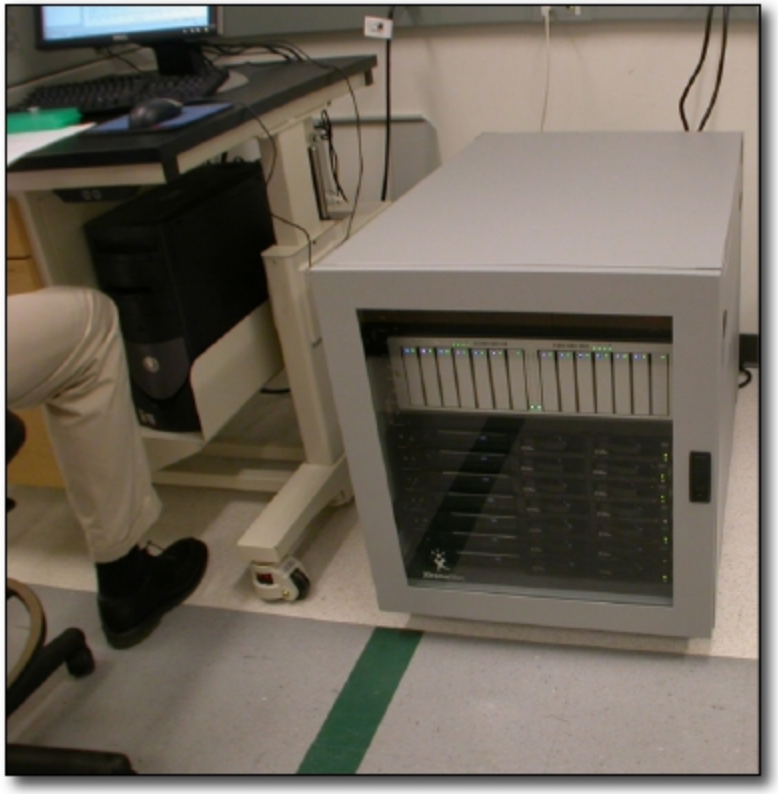

This is our new normal.

Thousands of CPU cores; petabytes of disk

And a related problem ...

- **It has never been easier to acquire vast amounts of data cheaply and easily**
- **Growth rate of data creation/ ingest exceeds rate at which the storage industry is improving disk capacity**
- **Not just a storage lifecycle problem. This data *moves* and often needs to be shared among multiple entities and providers**
  - ... ideally without punching holes in your firewall or consuming all available internet bandwidth

If we get it wrong ...

The stakes are high.

- **Lost opportunity**
- **Missing capability**
- **Frustrated & very vocal scientific staff**
- **Problems in recruiting, retention, publication & product development**
- **And in the clinic: Improper & potentially life-altering medical guidance**

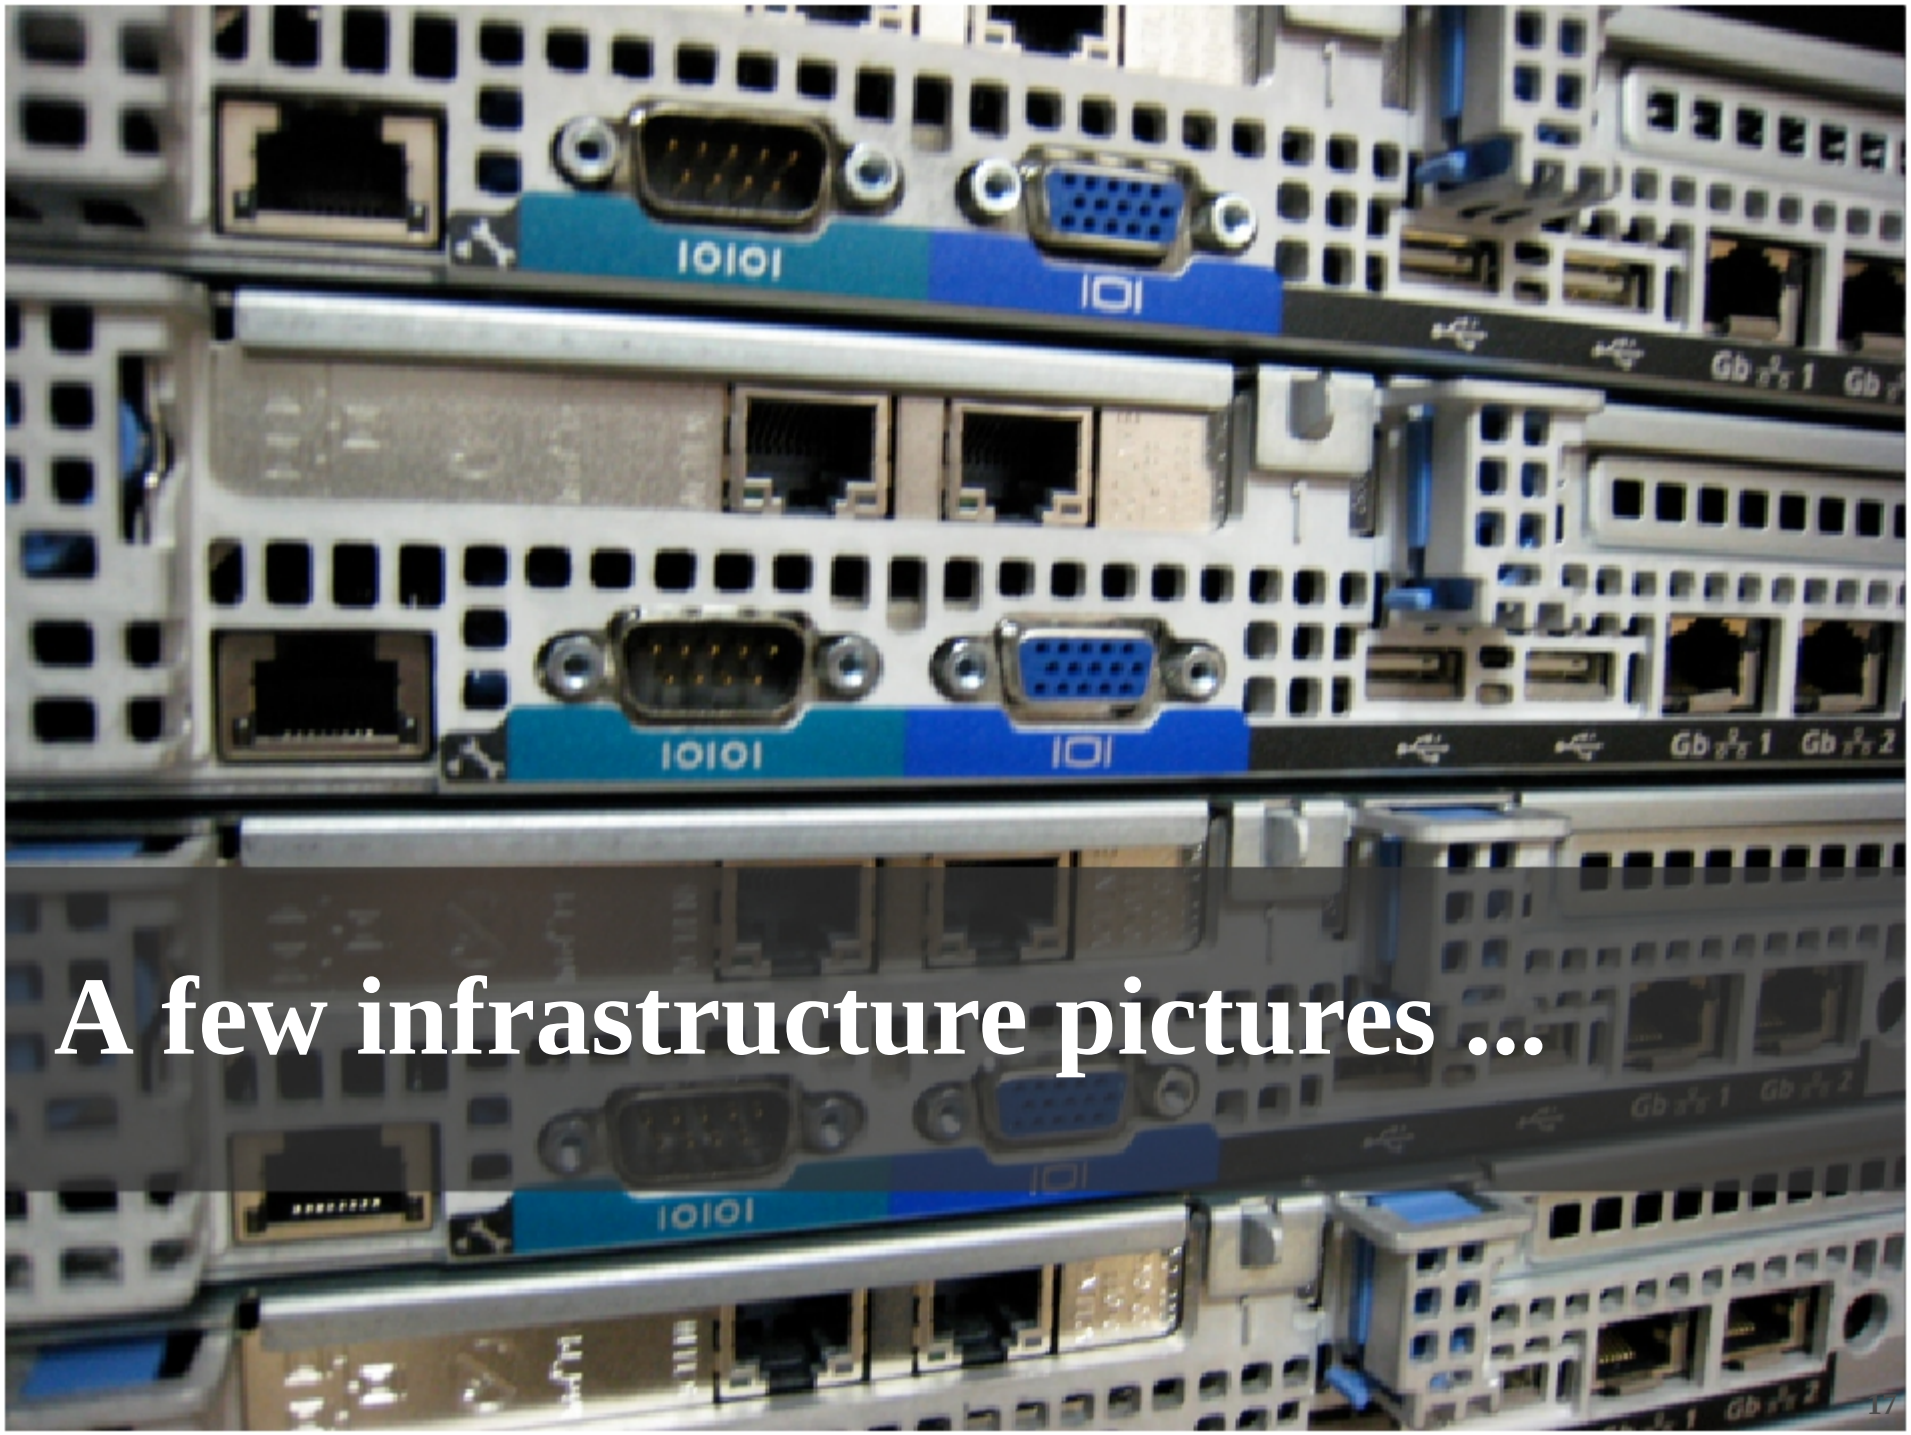
A few infrastructure pictures ...

The cliche image

Lab local HPC & storage

Small core w/ multiple NGS instrument support

Large Core Facility

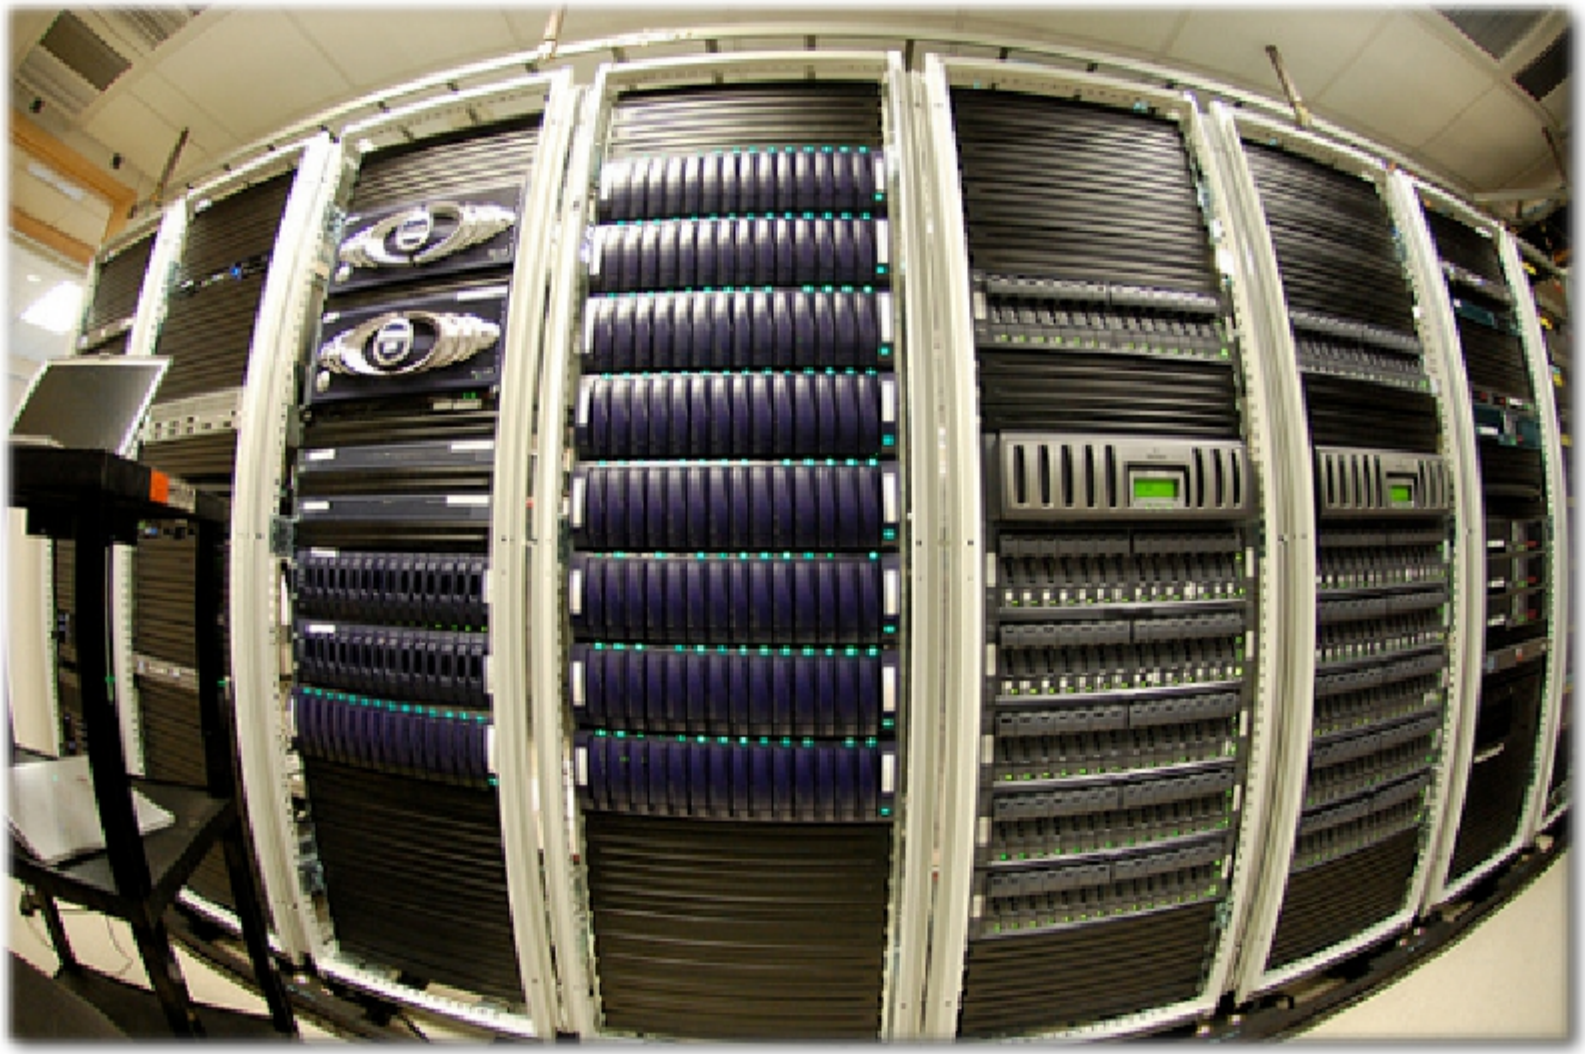

Large Core Facility: Just Storage

Regional Scientific Computing “Hub”

Physical data movement station; Unit= Naked Disk

Physical Ingest Repo & “Naked” Data Archive

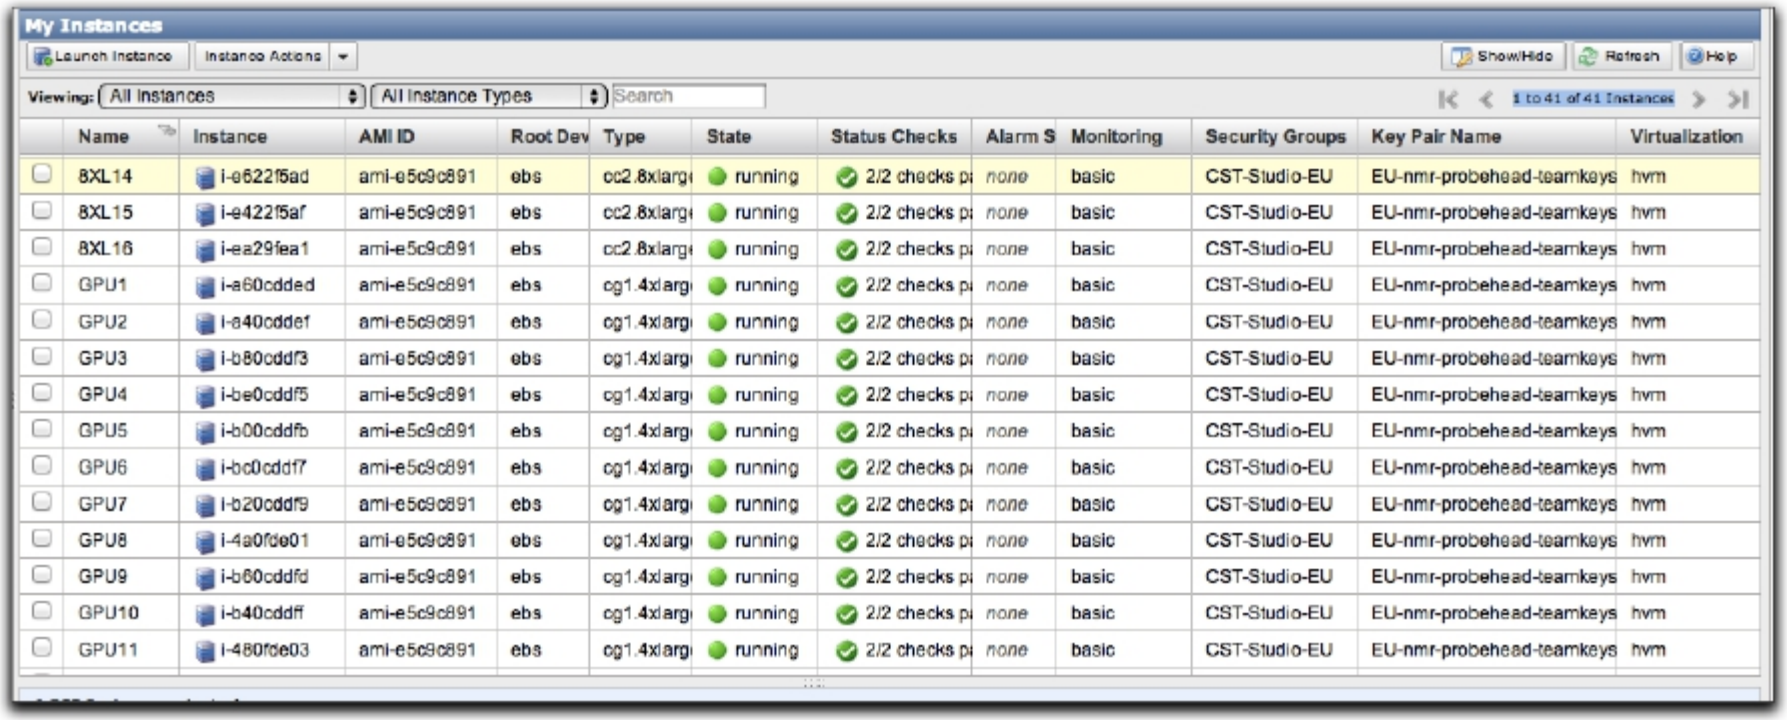

Name	Instance	AMI ID	Root Dev	Type	State	Status Checks	Alarm S	Monitoring	Security Groups	Key Pair Name	Virtualization
8XL14	i-e622f5ad	ami-e5c9c891	ebs	cc2.8xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
8XL15	i-e422f5af	ami-e5c9c891	ebs	cc2.8xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
8XL16	i-ea29fea1	ami-e5c9c891	ebs	cc2.8xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU1	i-a60cdded	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU2	i-a40cdde1	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU3	i-b80cddf3	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU4	i-be0cddf5	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU5	i-b00cddfb	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU6	i-bc0cddf7	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU7	i-b20cddf9	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU8	i-4a0fde01	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU9	i-b60cddfd	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU10	i-b40cddff	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm
GPU11	i-480fde03	ami-e5c9c891	ebs	cg1.4xlarg	running	2/2 checks pa	none	basic	CST-Studio-EU	EU-nmr-probehead-teamkeys	hvm

30 of Amazon's largest nodes + 22 GPU nodes: $30/hr via spot market

Yep. This counts.

Basic Bio/IT Landscape

Core Compute

Compute related design patterns largely static

- **Linux compute clusters are still the baseline compute platform**
- **Even our lab instruments know how to submit jobs to common HPC cluster schedulers**
- **Compute is not hard. It s a commodity that is easy to acquire & deploy in 2013**

File & Data Types

We have them all

- **Massive text files**
- **Massive binary files**
- **Flatfile databases**
- **Spreadsheets everywhere**
- **Directories w/ 6 million files**
- **Large files: 600GB+**
- **Small files: 30kb or smaller**

Application characteristics

- **Mostly SMP/threaded apps performance bound by IO and/or RAM**
- **Hundreds of apps, codes & toolkits**
- **1TB - 2TB RAM High Memory nodes becoming essential**
- **Lots of Perl/Python/R**
- **MPI is rare**
  - Well written MPI is even rarer
- **Few MPI apps actually benefit from expensive low-latency interconnects***
  - *Chemistry, modeling and structure work is the exception

Storage & Data Management

- **LifeSci core requirement:**
  - Shared, simultaneous read/write access across many instruments, desktops & HPC silos
- **NAS = easiest option**
- **Scale Out NAS products are the default standard**
- **Parallel & Distributed storage for edge cases and large organizations**

Storage & Data Management

- **Storage & data mgmt. is the #1 infrastructure headache in life science environments**
- **Most labs need *peta capable* storage due to unpredictable future**
  - Only a small % will actually hit 1PB
  - Often forced to trade away performance in order to obtain capacity
- **Object stores and commodity Nexentastor-style methods are making significant inroads**

Data Movement & Data Sharing

- **Peta-scale data movement needs**
  - Within an organization
  - To/from collaborators
  - To/from suppliers
  - To/from public data repos
- **Peta-scale data sharing needs**
  - Collaborators and partners may be all over the world

We Have Both Ingest Problems

Physical & Network

- **Significant physical ingest occurring in Life Science**
  - Standard media: naked SATA drives shipped via Fedex
- **Cliche example:**
  - 30 genomes outsourced means 30 drives will soon be sitting in your mail pile
- **Organizations often use similar methods to freight data between buildings and among geographic sites**

Physical Ingest Just Plain Nasty

- **Easy to talk about in theory**
- **Seems easy to scientists and even IT at first glance**
- **Really really nasty in practice**
  - Incredibly time consuming
  - Significant operational burden
  - Easy to do badly / lose data

Huge Need For Network Ingest

1. Public data repositories have petabytes of useful data
2. Collaborators still need to swap data in serious ways
3. Amazon becoming an important repo of public and private sources
4. Many vendors now deliver to the cloud

Cloud(s)

Mainstream in life science for quite some time

- **IaaS clouds offer excellent pressure release valve when rapidly changing scientific requirements can t be satisfied by on-premise infrastructure**
- **Economics can t be ignored**
- **Popular meeting ground for data swapping and collaboration**
- **Data providers pushing cloud delivery over physical media**
- **Interesting AWS use cases for archive and downloader pays methods**

```
#!/bin/sh

knife ec2 server create \
-d chef-full \
-N hubzeroTestNode \
-f t1.micro \
--image ami-4d20a724 \
--security-group dag-HZ \
-i ~/bioteam/bioteam-IAM-admins-v1.pem \
--ssh-key bioteam-IAM-admins-v1 \
-x admin
```

Cloud Hubs & Portals

The ‘neutral’ meeting ground ..

- **Many types of entities need to meet, collaborate and exchange life science data**
- **Data sharing hubs and portals becoming popular on public IaaS clouds like AWS**
- **Why?**
  - Far easier than punching holes in your firewall and issuing VPN credentials to outsiders

Cloud Data Repositories

Compelling economics

- **IaaS clouds becoming center of gravity for large scale scientific data hosting**
- **Why?**
  - Very compelling pricing
  - You don t pay for the bandwidth used by consumers of your data
  - AWS has some very interesting downloader pays models that seem to be a good fit for grant-funded science with mandated multi-year data accessibility requirements

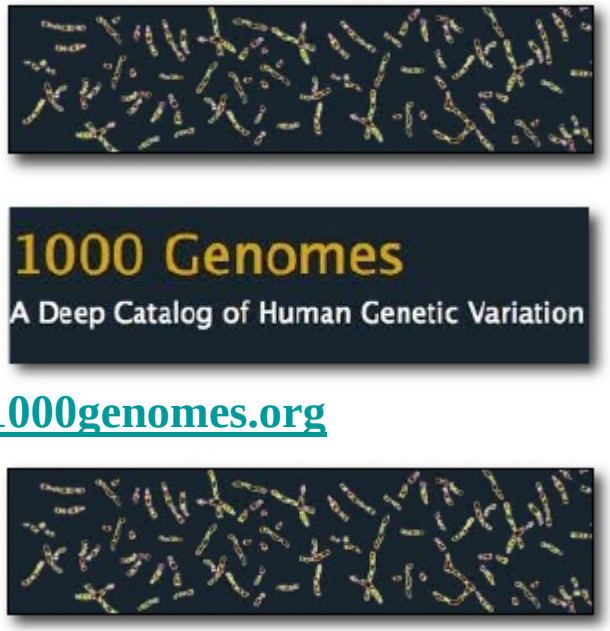

www.1000genomes.org

It all boils down to this ...

Life Science In One Slide:

- **Huge compute needs but not intractable and generally solved via Linux HPC farms. Most of our workloads are serial/batch in nature**
- **Ludicrous rate of innovation in lab drives a similar rate of change for our software and tool environment**
- **With science changing faster than IT, emphasis is on agility and flexibility - we ll trade performance for some measure of future proofing**
- **Buried in data. Getting worse. Individual scientists can generate petascale data streams.**
- **We have all of the Information Lifecycle problems: Storing, Curating, Managing, Sharing, Ingesting and Moving**

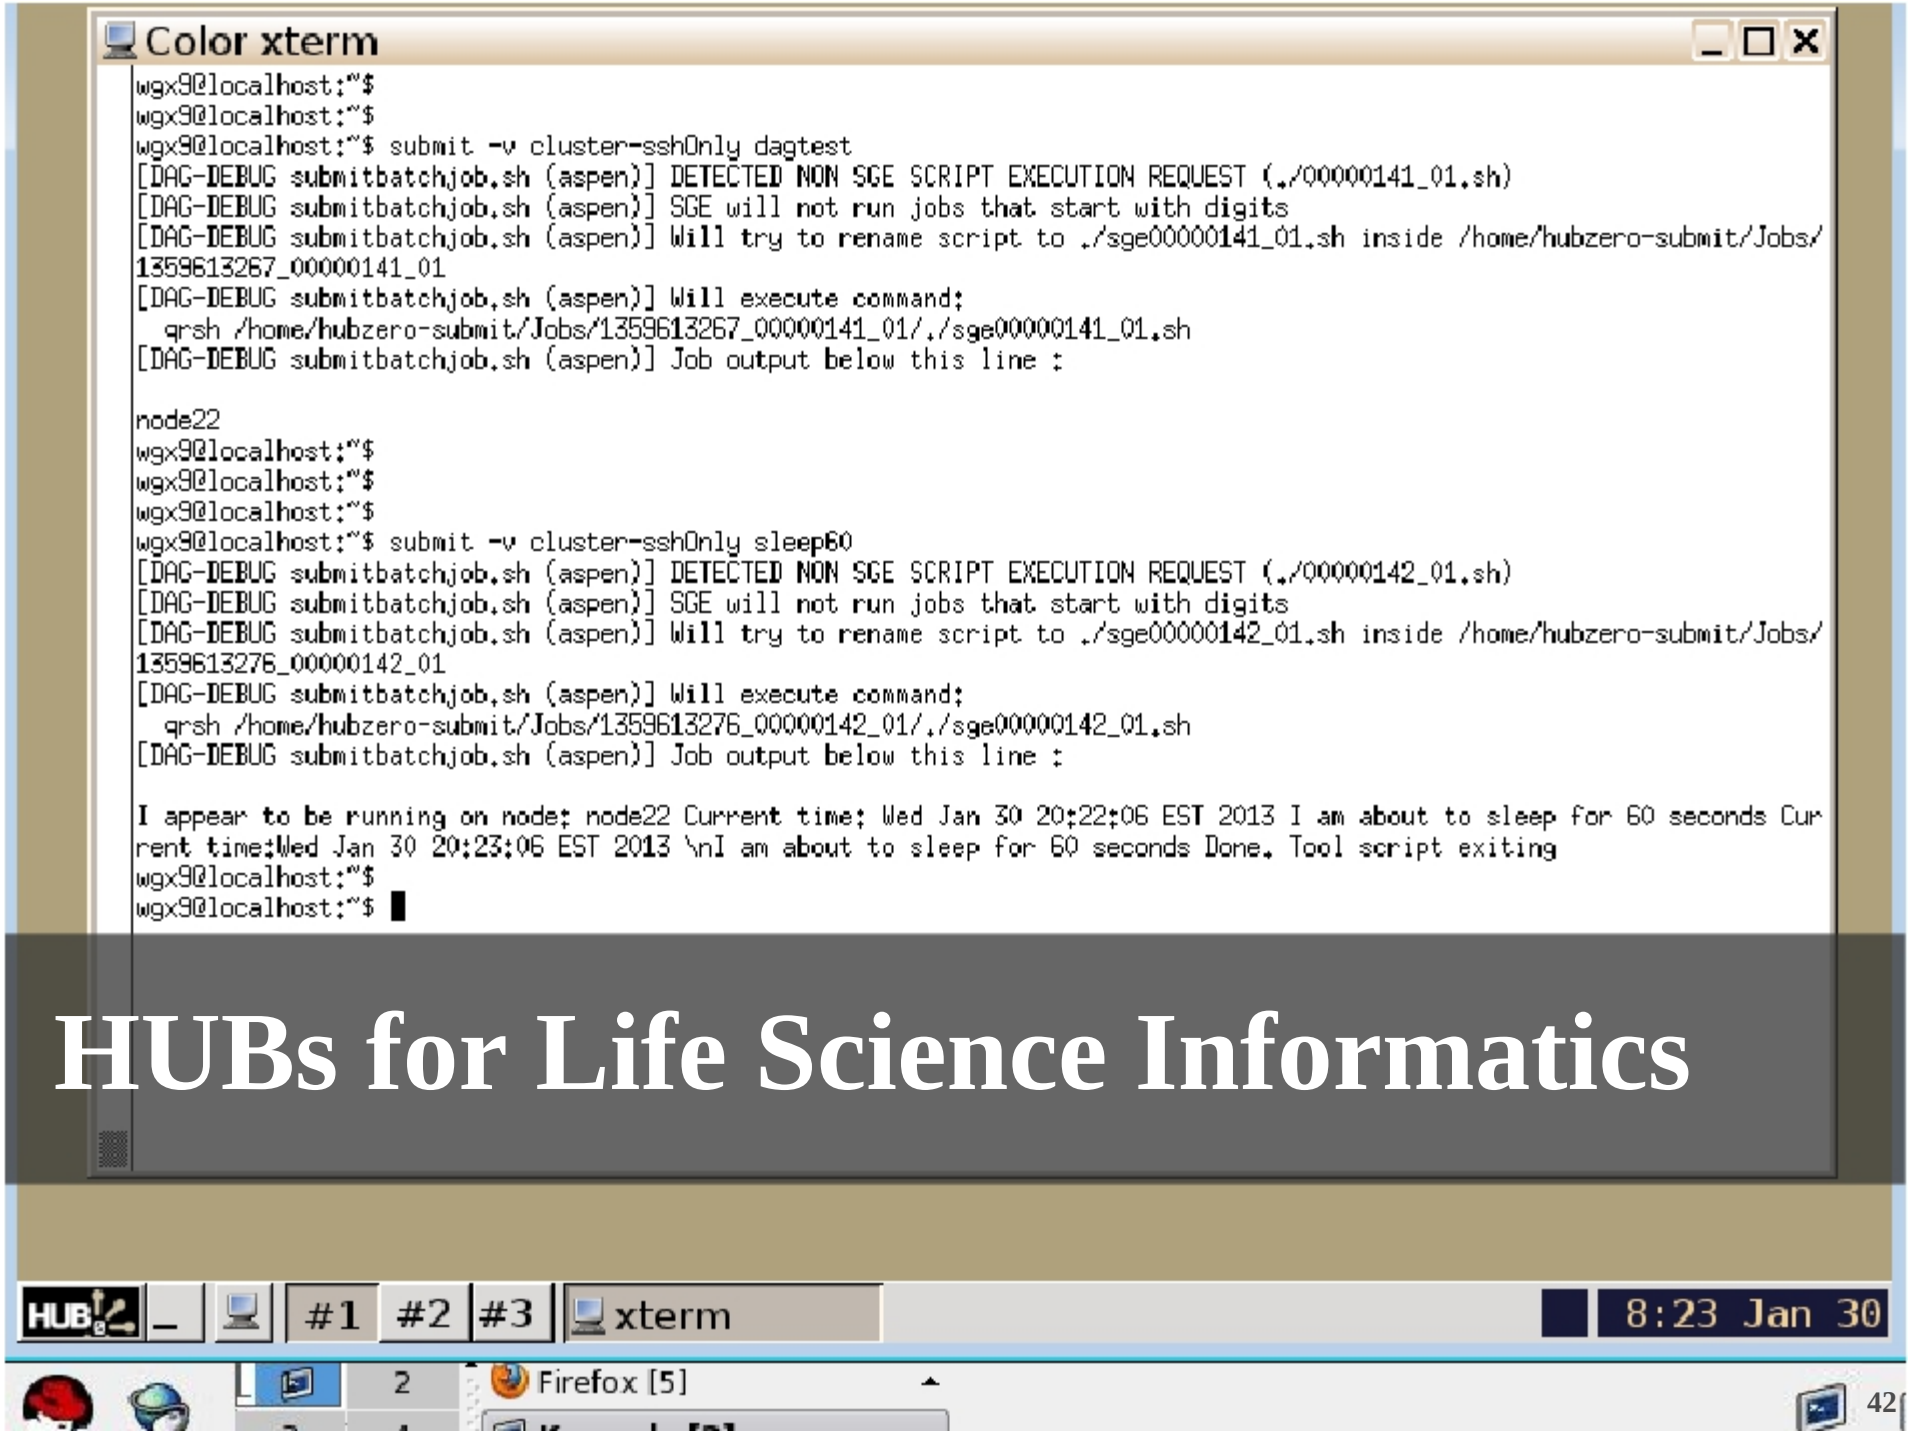

HUBs for Life Science Informatics

hubzero and one .gov

My Hub Story

- **Sub-contractor to much larger company that won HPC and scientific support contract for division of a very large US .gov disease and public health organization**
- **Our team has very a specific technical/HPC & scientific computing support mission**
  - ... and a general mission to enhance collaboration, data and knowledge sharing very diverse group of dedicated professionals.
- **HZ recommended by a team member and supported by senior internal sponsor**

hubzero and one .gov, cont.

My Hub Story

- **I was part of a group that did**
  - Initial evaluation
  - Initial trial deployments
  - Technical assessment report for the go|no-go decision
  - resources required to operate report
- **Along the way ...**
  - Made ugly hacks to Submit() framework to enable integration with Grid Engine managed HPC Clusters
  - ... met a few HZ team members virtually

hubzero and one .gov, cont.

My Hub Story

- **And in my free time (non-work) ...**
  - Started work on Opscode Chef Cookbook for automated installation of Hubzero onto AWS Cloud and other automated DevOps environments
  - http://bioteam.net/2012/12/hubzero-on-the-cloud/

```
#!/bin/sh

knife ec2 server create \
 --verbose \
 --template-file /Users/dag/opscode-platform-projects/bioteam/chef-repo/.chef/bootstrap/debian6-gems.erb \
 --availability-zone us-east-1b \
 --security-group dag-HZ \
 --node-name hubzeroTestNode \
 --flavor t1.micro \
 --image ami-4d20a724 \
 --ssh-user admin \
 --identity-file /Users/dag/opscode-platform-projects/bioteam/bioteam-IAM-admins-v1.pem \
 --ssh-key  bioteam-IAM-admins-v1
```

Current Status

Hubzero and .gov story

- **In production & formally deployed**
  - ... but in small-scale pilot mode
- **Successful integration with Grid Engine and our HPC cluster & storage environment**
- **Primary current use is documentation**
- **Secondary focus will be Collaboration/Projects**
- **Beginning to recruit users and encourage the more social collaboration and self-publishing features**

Future Status

Our sweet spots and interest areas ...

- **Very significant interest in courseware & learning**
- **Very significant interest in the DataStore features**
- **Interest in tags, especially automated tagging**
- **Interest in the fine-grained security model & role-based access to content**
- **Interest in automated/API-driven content creation. Possibly by lab instruments that auto publish experimental result summaries w/ links to data repository location**

Time permitting ...

image: shanelin via flickr

Time?

A few more details

1. My take on our Hubzero assessment
2. Enterprise integration war stories
3. Potential Personal Promise

My $.02: Internal Hubzero assessment

Speaking for myself, not my employer or .gov team !

- **Initial Impressions**
  - Very positive
  - The public hubs are doing a great job at evangelizing
  - Particularly good job at delivering a CMS that understands the org structure and working behaviors of real scientists
  - Great foundation for what we were looking to do
- **Observed, however:**
  - Not many commercial/industry users publicly visible
  - HZ not easily deployed in enterprise environments (OS/vm/etc)
  - The most sophisticated Hubs appear to be hosted at Purdue where they depend on internal expert support and may also leverage knowledge that otherwise may not be well documented

My $.02: Internal Hubzero assessment

Speaking for myself, not my employer or .gov team !

- **Assessment summary outcome:**
  - Worth pursuing; great features and capabilities
  - IT Infrastructure requirements for running a hub are trivial relative to the stuff we have already supporting peta-scale genomics
- **However:**
  - Human resource requirements were non trivial and if senior management was serious about deploying a hub in any meaningful way they would have to commit real dollars and real staff to the effort.

My $.02: Internal Hubzero assessment

Speaking for myself, not my employer or .gov team !

- **Human resources required for a successful Hubzero deployment:**
  - Salaried FTE or part-time internal employee (NOT AN OUTSIDE CONTRACTOR) required to build community, configure the Hub and serve as Community Manager
  - Need a Web Developer w/ creative skills and UI experience (the best hubs have great UI and visuals)
  - Need a Joomla CMS expert on contract or on-call
  - Existing IT staff would need C++/Rapture training

War Stories

Episodes that drove the assessment results

1. Outsider vs Insider issues
2. Enterprise deployment issues
3. Authentication & identity management
4. Featured Article adventures
5. Submit() and Rapture adventures

One last thing ...

One last thing

- **In 2013 when I can ...**
  - ... Launch a completely integrated 500 CPU/GPU cluster on AWS with 1 command typed from my macbook
  - ... Launch a full Wordpress stack with global CDN, SSL-aware load balancers, geographically replicated SQL datastores and an elastic fleet of webservers via a single CloudFormation .json template
- **It s hard to believe that Hubzero is not available for everyone as a single button click on Amazon or Google Compute**

One last thing

- **I ve benefited personally and professionally from open source over the years; have always tried to contribute back in kind**
- **If there is interest in making Hubzero available in forms other than VM files and package-based installs I might be able to contribute**
- **Personal interests include:**
  - Packaging Hubzero for AWS Marketlpace ($0.00 product)
  - Creating Opscode Chef Cookbook that can auto build and deploy Hubzero from packages on any Debian/Ubuntu OS

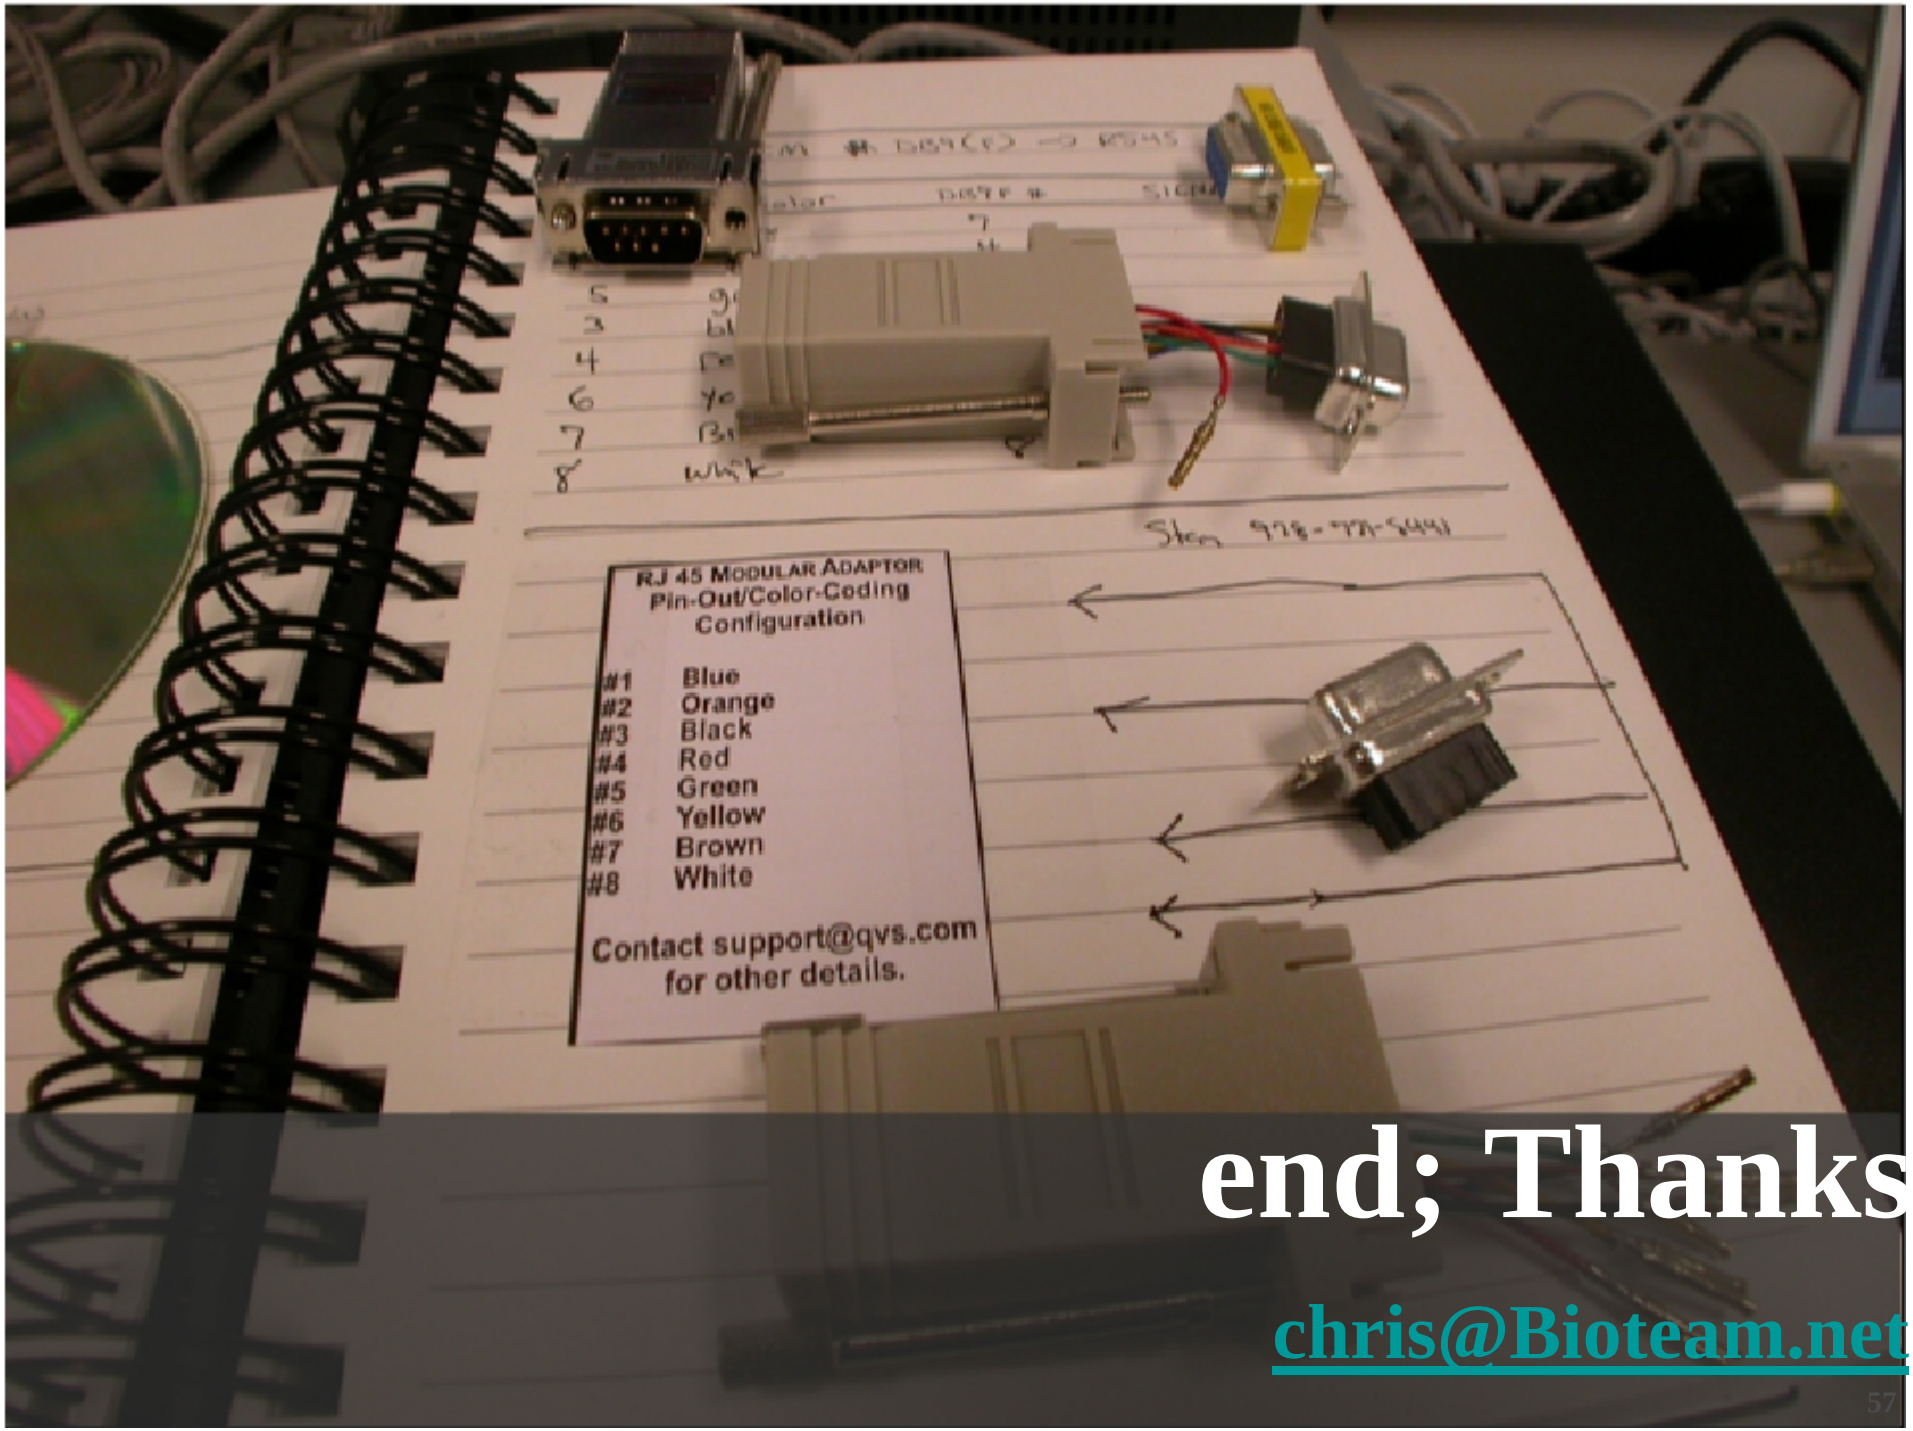

end; Thanks

chris@Bioteam.net