

# Adventures in cloud...


## CLIMB

## BEAR Cloud

Simon Thompson

Research Computing Infrastructure

Architect

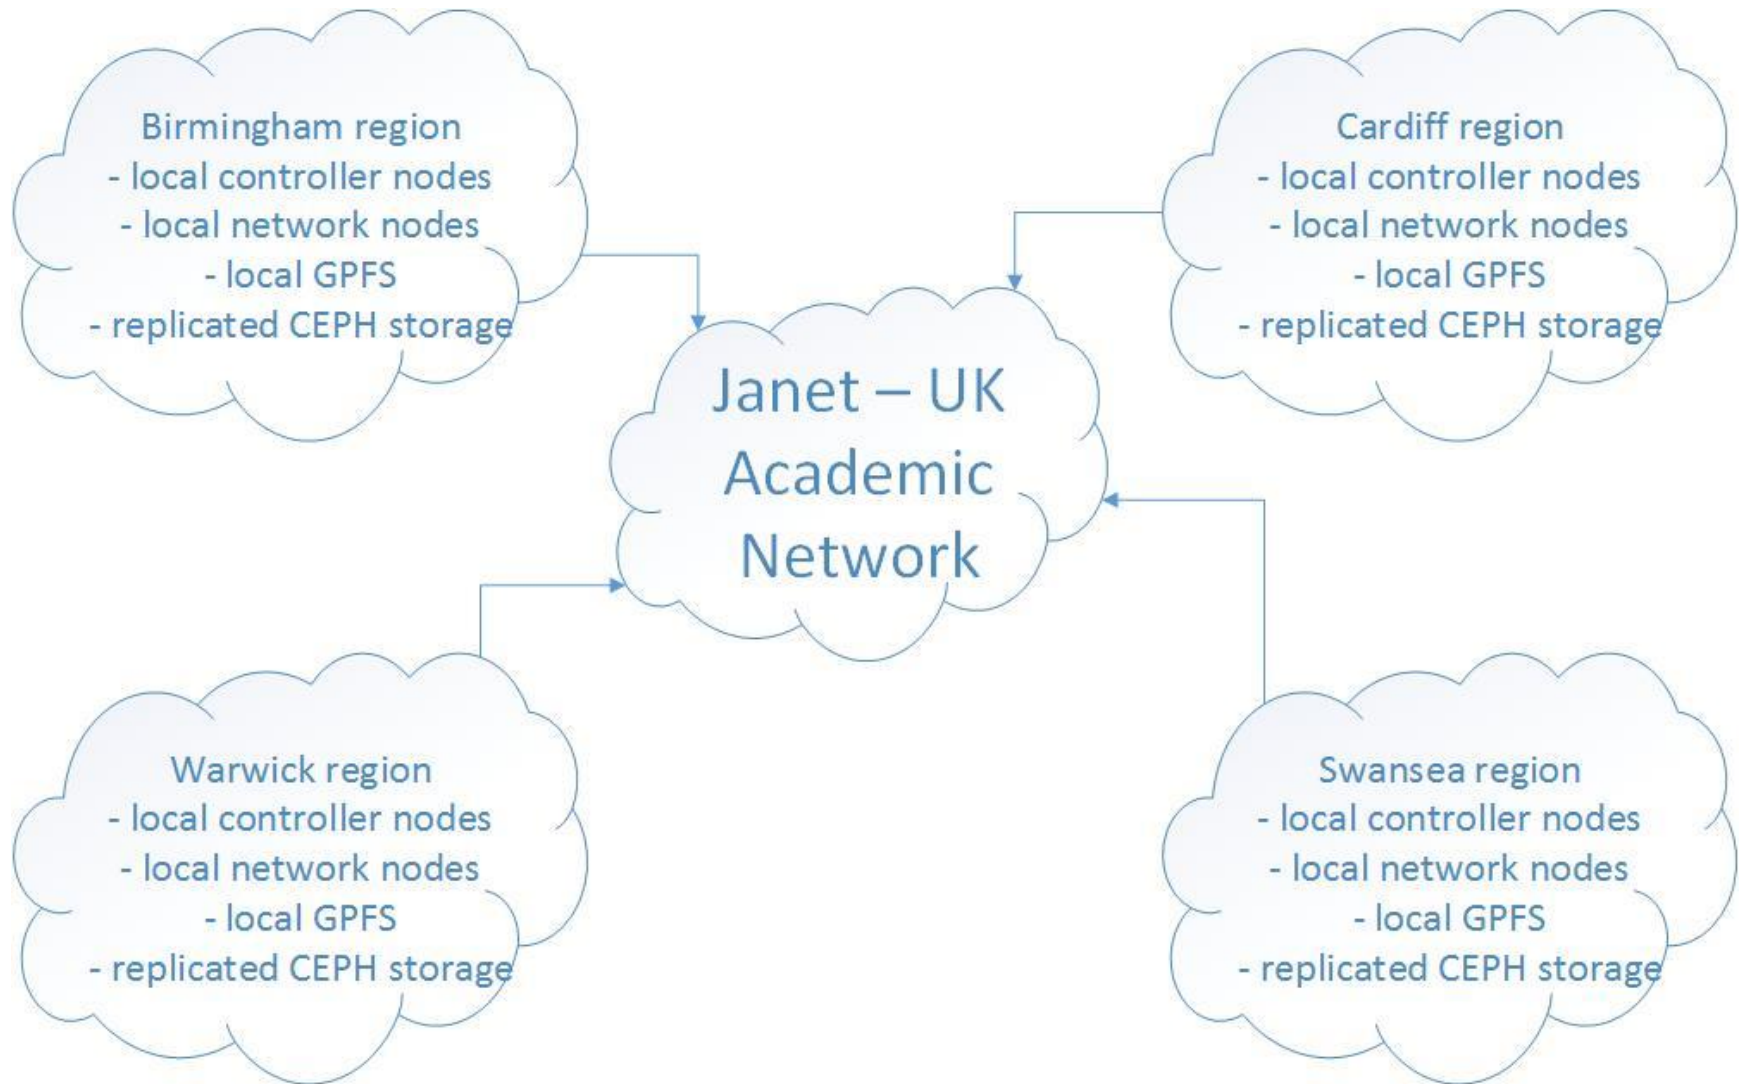


**UNIVERSITY OF  
BIRMINGHAM**

# What is CLIMB?

- MRC funded cloud spanning Bath, Birmingham, Cardiff, Swansea, Warwick
- Led by Mark Pallen (Warwick), Sam Sheppard (Bath), Nick Loman (Birmingham), Tom Connor (Cardiff)
- Supporting microbial bioinformaticians
- Birmingham research computing built the pilot service used to help develop other sites

# CLIMB Overview




vRouter provides encrypted  
ipsec tunnels between  
sites.



Brocade vRouter  
(HA Pair)

Multiple fabric switches, all  
hosts dual attached over 2  
switches



3 controller nodes  
Mariadb cluster


haproxy, keepalived (VRRP) for services  
native HA VRRP L3, multiple DHCP agents  
for neutron



Mellanox Infiniband  
GPFS traffic flows over IB using VERBS,  
will fail back to Ethernet using subnets



GPFS Servers (2-4)  
Multiple v3700  
storage arrays



Compute / Hypervisor nodes  
4 – 16 socket  
512GB – 3TB RAM

# Genome sequencing: 20 years ago



Microbial genome: >\$1,000,000 and 3 years


First human genome \$2,700,000,000 and 13 years

# Genome sequencing: 2015



## Falling fast

In the first few years after the end of the Human Genome Project, the cost of genome sequencing roughly followed Moore's law, which predicts exponential declines in computing costs. After 2007, sequencing costs dropped precipitously.



Microbial genome: \$50 in hours

Human genome \$1,000 in days

# Real-time sequencing in epidemics

Genome sequences are: Universal, portable, digital, comparable, information-rich!

start

Identification

*E. coli O104:H4 in Germany*

Source tracking

*Haiti cholera from Nepal  
Amerithrax*


Control

*Tailored diagnostics  
Influenza infection control and public  
health measures*

Biology

*E. coli O104:H4 in Germany  
Monitor response to human  
interventions e.g. vaccines,  
treatments*





**Bill Gates** @BillGates · 22h

From Ebola to Zika, this “lab in a suitcase” provides crucial data for outbreaks: [b-gat.es/1XIKQkZ](https://b-gat.es/1XIKQkZ) via @verge




↻ 899


❤ 1.7K







# CLIMB: International bioinformatics collaborations facilitated by elastic storage, compute



Training performed on high memory CLIMB servers (1Tb RAM with 60 vCPUs)

<http://simpsonlab.github.io/>



Loman, Quick, Simpson Nature Methods 2015


# Real-time analysis of Ebola virus evolution

2016 Jan 27

Apr Jul Oct 2015 Apr Jul Oct 2016

Region










Epi week: 2014-Feb-10  
Decimal time: 2014.121




# Genome sequencing: now



Zibraproject.org, real-time analysis pipeline on CLIMB, [nextstrain.org/zika](https://nextstrain.org/zika)





- Genome sequencing can now be done everywhere and anywhere
- Genome datasets doubling faster than Moore's law
- Tension between democratisation and centralisation
- I/O a bottleneck: bring compute to datasets
- For value of biology to be unlocked we need tools for large-scale data sharing and integration
- Cutting-edge academic bioinformatics software remains stubbornly hard to use

# Training the next generation of bioinformaticians

MRC Cloud Infrastructure for Microbial Bioinformatics

all categories ▶ Latest New Unread (5) Top Categories

Topic Category

Salmonella outbreak	Wellcome Trust DTP Tutorials
Resolving ssh login problems	Tutorials
NGS for beginners	Wellcome Trust DTP Tutorials
Data formats and databases	Wellcome Trust DTP Tutorials
Beginner's Unix	Wellcome Trust DTP Tutorials
Accessing CLIMB and Setup	Wellcome Trust DTP Tutorials
Documentation for intro to CLIMB	Documentation
Correct directory for files I'm ftp'ing for galaxy?	
Cannot copy/paste to and from the ubuntu desktop from Windows	Service

Loman Labz Wellcome Trust AAMR DTP Student Group

## Wellcome Trust AAMR DTP Student Group

### Running servers at Bham

Name	Created	Flavor	Status	IP address	Refresh
tom-gvl	2016-10-16T21:15:26Z	climb.user	Active	147.188.173.143	Options
emily-gvl	2016-10-16T21:14:29Z	climb.user	Active	147.188.173.142	Options
tutor-gvl	2016-10-15T12:32:04Z	climb.user	Active	147.188.173.140	Options
fernanda-gvl	2016-10-14T11:12:55Z	climb.user	Active	147.188.173.132	Options
gloria-gvl	2016-10-14T11:11:59Z	climb.user	Active	147.188.173.131	Options
liying-gvl	2016-10-14T11:10:26Z	climb.user	Active	147.188.173.128	Options
farhana-gvl	2016-10-14T11:09:24Z	climb.user	Active	147.188.173.127	Options
sam-gvl	2016-10-14T11:07:35Z	climb.user	Active	147.188.173.123	Options
jessica-gvl	2016-10-14T11:06:44Z	climb.user	Active	147.188.173.122	Options
alice-gvl	2016-10-14T10:58:25Z	climb.user	Active	147.188.173.121	Options
alastair-gvl	2016-10-14T10:56:44Z	climb.user	Active	147.188.173.120	Options


New PhD students get their own personal server for research

Online course documentation




# System Outline

- 4 sites
- Connected over Janet
- Different sizes of VM available; personal, standard, large memory, huge memory
- Able to support >1,000 VMs simultaneously (1:1 vCPUs/vRAM : CPUs/RAM)
- ~7PB of object storage across 4 sites (~3PB usable, replicated)
- 500TB of local high performance storage per site
- A single system, with common log in, and between site data replication\*
- System has been designed to enable the addition of extra nodes / Universities



# And now BEAR Cloud



# Why Cloud?

- HPC is not ideal for some types of workload
  - Activation energy to use
  - Software is difficult to force-fit
- Many emerging community tools/portals for research
- University strategic investment for Life Sciences (CaStLeS) – as an urgent driver
- Supporting Golden Science
- Encouraging good research citizenship



# Why Cloud?

- Large memory or intensive applications that aren't suited to HPC
- Easing transition from desktop to highly powered resources

# What BEAR Cloud is **NOT**?

- An HPC facility
  - Keen to work with groups who need isolated HPC
- A way of just getting “Admin” rights on a PC
- For running web servers
  - Though web fronted scale out services are welcome

# What it provides


- Large numbers of cores per system/VM (up to 20 cores)
- Large amounts of memory (up to 128GB)
- Fast connectivity to storage and research networks
- Ability to use Infiniband
- Low latency interconnect

# Why not public cloud?

- Data movement and storage costs
- Potential IP questions
- Capital expenditure on “as a service” difficult to manage, control
- Business models based on over-allocation of CPUs and memory
- Support for other technologies is (currently) poor (Infiniband, RDMA, MPI, GPU)
- Not optimised and purpose built for research workloads

# Mitigating VM “overheads”

- Mellanox ConnectX-4 VPI
  - Infiniband (100Gb EDR)
  - 10Gb Ethernet to hypervisors
  - Hardware offload for network functions
  - SR-IOV Infiniband & Ethernet



# Mitigating VM “overheads”

- High performance storage from DDN
  - Fast access to storage for data
  - Spectrum Scale plugins for OpenStack
  - Liberal use of SSD to tweak “bottlenecks”



# Mitigating VM “overheads”


- Water cooled technology from Lenovo
  - Turbo more sustained = extra “free” Ghz
  - (and its really efficient to cool!)
- Intel Broadwell
  - Per core clock scaling




User portals



GitLab



Research Data  
Network



Research group HPC  
instances

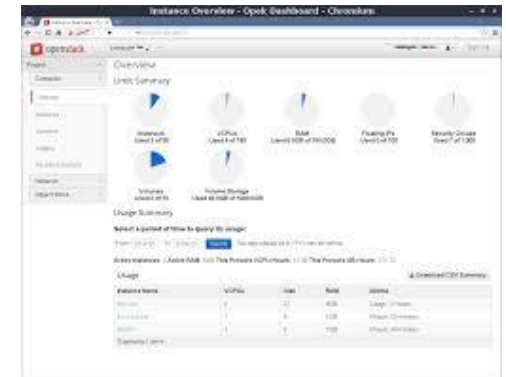
Self VM provisioning  
(Horizon)



Research Data Storage



BlueBEAR HPC



# And next?

- More community portals enabled
  - Working with research groups
- Tighter integration with research data services
- Automated workflows?
- Data management services?
- Containers, Docker, Kubernetes
- Science reproducibility?
- Sharing of resource and research?
- GPU enablement?
- Integration work with public or research cloud?