

The NGS IT notes

George Magklaras PhD RHCE

Biotechnology Center of Oslo &
The Norwegian Center of Molecular Medicine

University of Oslo, Norway

<http://www.biotek.uio.no>

<http://www.ncmm.uio.no>

<http://www.no.embnet.org>

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

Agenda

- NGS basic notions and history
- Outline of NGS IT challenges
- The storage part
- The networking part
- The computing part
- The big picture





brave
IT
engineer

How not to hit the NGS IT platform brick wall

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

NGS basics: The genome

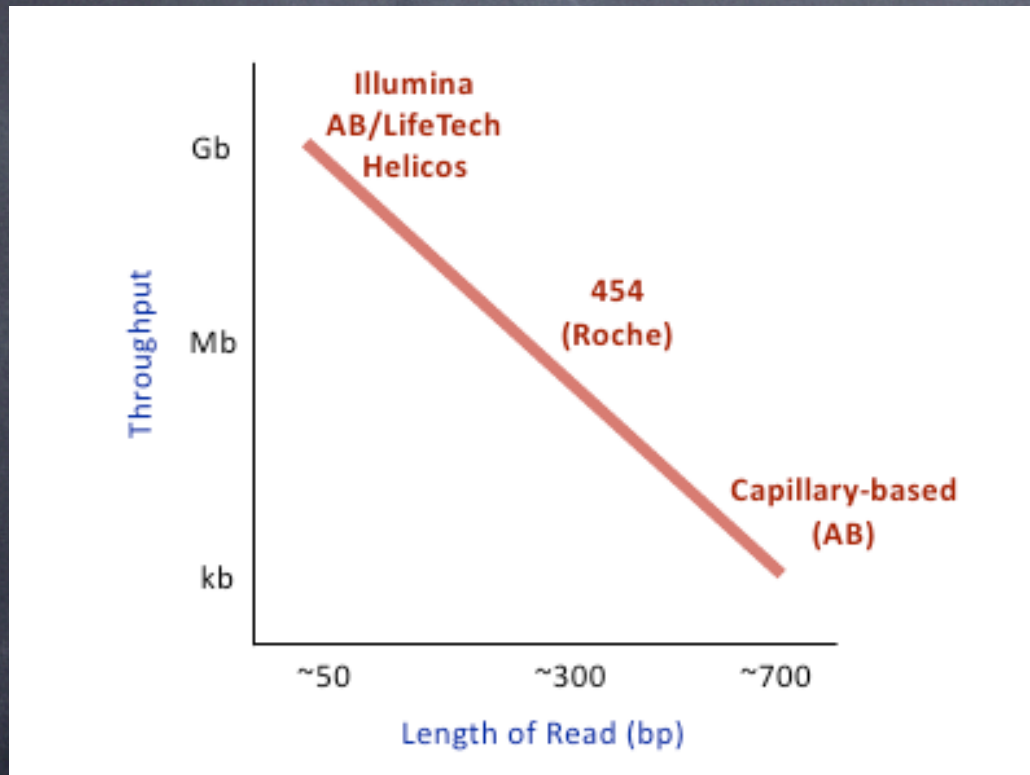
- the genome is the entirety of an organism's hereditary information. It is encoded either in DNA or RNA. The genome includes both the genes and the non-coding sequences of the DNA/RNA.
- The DNA/RNA sequence is the "code". This code is used as the blueprint to construct proteins, the building blocks living things are made of.



NGS (why?)

- Next Generation Sequencing (NGS, also known as High Throughput Sequencing) aims to help us:
 - Record the “code” of life at a reasonable time and cost (sequencing).
 - Understand how living things are made (transcriptome, proteome) and how the “code” works.
 - Understand how the environment affects us (metagenomics) and how we differ from each other (comparative genomics)
 - Assemble this information into a complete genome project

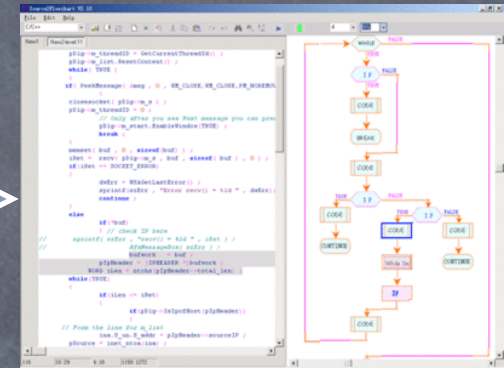
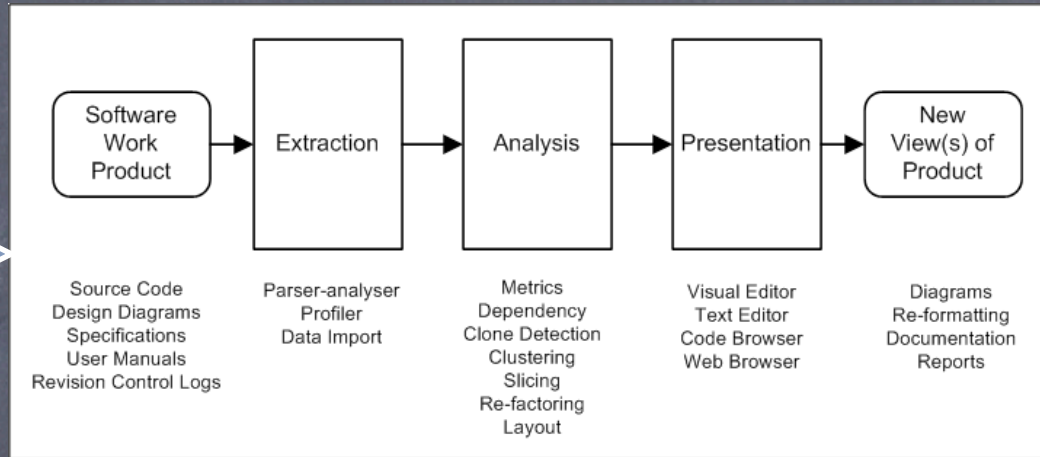
Sequencing throughput (ST)



ST = Sequences generated/
time (OR) cost

Gb - Gigabases Mb - Megabases kb - Kilobases

Reverse engineering workflow (software)



```
or (int j = 0; j < loc; j++) res[j] = buf[j];
return res;

public void checkRes(int[] res) {
    for (int i = 0; i < res.length; i++) {
        if (res[i] != checkRes[i]) {
            // ...
        }
    }
}

decodeMessage(int[] buf) {
    int i = 0;
    while (i < buf.length) {
        // ...
    }
}

public int[] extractMessage(int[] res) {
    int loc = 0;
    while (loc < res.length) {
        // ...
    }
}
```



This program performs functions A and B by using algorithms X and Y.

NGS workflow

Prepare Genomic DNA Sample



Determine 5. base

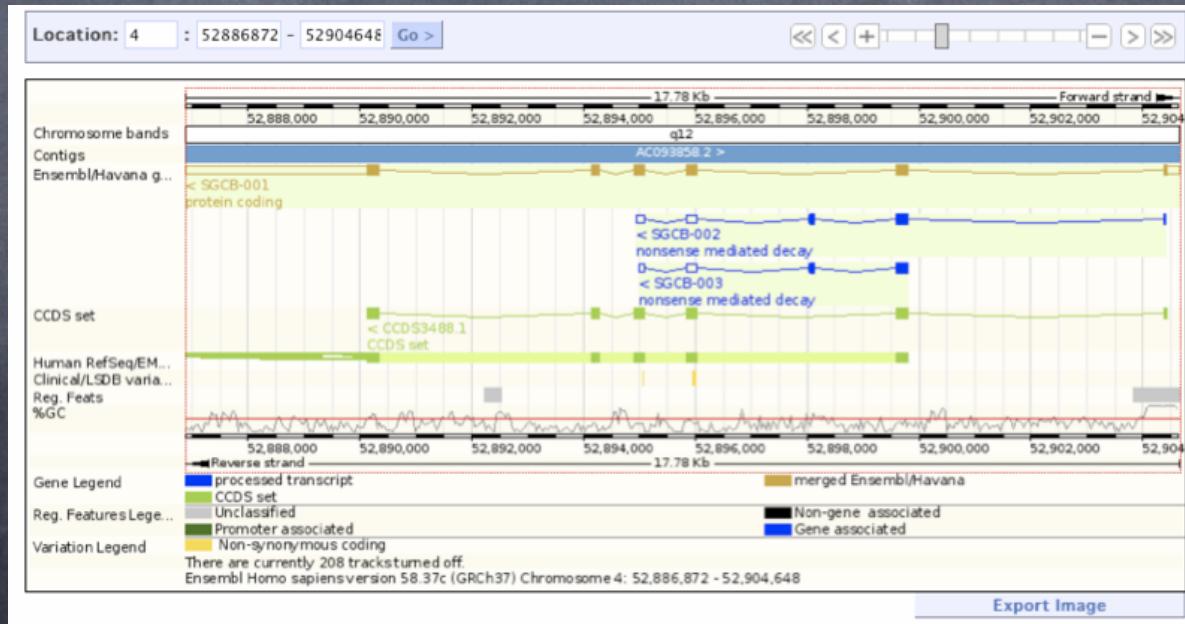


Data Analysis

SNP detection and error masking

```
CACTCCGGTGGTCAC
ATGTGCCCTCCTCACTCCGGTGGT
CTGATGTGCCCTCCTCACTCTGGT
GGTGTGCCCTCCTCACTCCGGTGG
GCTGATGTGCCCTCCTCACTCCGG
...GCTGATGTGCCACCTCACTCCGGTGGTCAC...
```


Sequence Assembly



Annotated Genome Project

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

NGS technologies




454
(Roche)

Illumina
Genome
Analyzer

Applied
Biosystems
SOLiD
Analyzer

Single Molecule
Sequencing
Analyzers

Illumina NGS workflow



Talleres Internacionales de Bioinformática - UNAM- Enero 2012

NGS history (2)



1984

Leroy E. Hood builds the first semi-automated DNA sequencer



1987


ABI 370 Protein Sequencer by Applied Biosciences

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

NGS history




Walter
Gilbert



Frederick
Sanger



Paul
Berg



1980 - Nobel Prize in Chemistry

"for their contributions concerning the determination of base sequences in nucleic acids" (Gilbert, Sanger)

"for his fundamental studies of the biochemistry of nucleic acids, with particular regard to recombinant-DNA" (Berg)


Talleres Internacionales de Bioinformática - UNAM- Enero 2012

NGS history (3)



1996

Pål Nyrén and Mostafa Ronaghi invent pyrosequencing at the Royal Institute of Technology (KTH) Sweden



2004

454 Pyrosequencer

NGS history (4)



David
Walt



Shankar
Balasubramanian



David
Klenerman


"We have an idea that could increase the rate and lower the cost of gene sequencing by [a factor of] 10^4 or 10^5 "

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

The NGS IT challenge

- Life Scientists generate more data than they can analyze. They need the data, but they need IT to manage the NGS workflow.
- You need to know how to:
 - Store: How much and at what speed?
 - Move: NGS data needs to move around.
 - Process the data: CPU and memory requirements.

Enter the Petabyte era



nature International weekly journal of science

[nature news home](#) [news archive](#) [specials](#) [opinion](#) [features](#) [news blog](#) [nature journal](#)

[! comments on this story](#) Published online 3 September 2008 | *Nature* 455, 16-21 (2008) | doi:10.1038/455016a

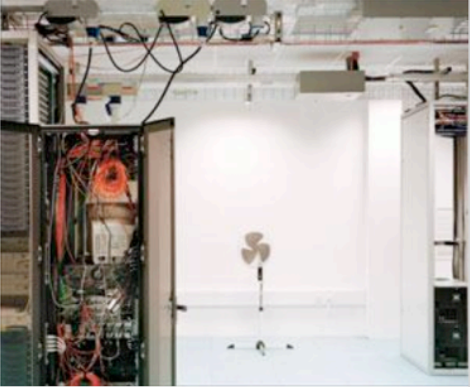
News Feature

Big data: Welcome to the petacentre

What does it take to store bytes by the tens of thousands of trillions? Cory Doctorow meets the people and machines for which it's all in a day's work.

Cory Doctorow

Ten seconds after I stepped into the roar of the data centre at the UK Wellcome Trust Sanger Institute, in rural Cambridgeshire, my video camera croaked: CARD FULL. Impossible. That morning, I'd tossed a handful of thumbnail-sized 32-GB memory cards into my pocket, each one good for a couple of hours' worth of high-definition video. Yet this one had filled up in seconds.



The data centre at the Wellcome Trust Sanger Institute in Cambridge, UK, under development. S. NORFOLK

Stories by subject

- [Lab life](#)
- [Technology](#)
- [Science in culture](#)

Stories by keywords

- [Petabytes](#)
- [Terabytes](#)
- [Sanger Centre](#)
- [CERN](#)
- [XS5ALL](#)
- [Internet Archive](#)
- [Cory Doctorow](#)
- [energy costs of computing](#)
- [Datacentres](#)
- [bioinformatics](#)
- [LHC](#)
- [data](#)

This article elsewhere

[I fumbled with my camera](#)

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

The Petabyte table

HTS device	runs/ year	Tier 1 Gbytes	Tier 2 Gbytes	Tier 3 Gbytes	Tier 4 Gbytes	TOTAL Tbytes
Illumina	100	9728	100	100	400	990
454	100	200	50	25	75	27
SOLiD	100	6144	100	100	200	80

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

NGS devices data production

Platform	Library/ template preparation	NGS chemistry	Read length (bases)	Run time (days)	Gb per run	Machine cost (US\$)	Pros	Cons	Biological applications
Roche/454's GS FLX Titanium	Frag, MP/ emPCR	PS	330*	0.35	0.45	500,000	Longer reads improve mapping in repetitive regions; fast run times	High reagent cost; high error rates in homo- polymer repeats	Bacterial and insect genome <i>de novo</i> assemblies; medium scale (<3 Mb) exome capture; 16S in metagenomics
Illumina/ Solexa's GA _{II}	Frag, MP/ solid-phase	RTs	75 or 100	4 [‡] , 9 [§]	18 [‡] , 35 [§]	540,000	Currently the most widely used platform in the field	Low multiplexing capability of samples	Variant discovery by whole-genome resequencing or whole-exome capture; gene discovery in metagenomics
Life/APG's SOLiD 3	Frag, MP/ emPCR	Cleavable probe SBL	50	7 [‡] , 14 [§]	30 [‡] , 50 [§]	595,000	Two-base encoding provides inherent error correction	Long run times	Variant discovery by whole-genome resequencing or whole-exome capture; gene discovery in metagenomics

Courtesy Elliott H. Margulies PhD, NIH

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

Tiered storage

- Tier 1: raw unprocessed data as they come out from the instrument (mostly images)
- Tier 2: including base (or colour) calls, intensities and first pass quality scores
- Tier 3: Includes aligned and analyzed data (alignments of all the reads to a reference or de-novo assembly, if required)
- Tier 4: Backup up off-site of Tiers 2 and 3, in order to provide disaster recovery/regulatory requirements.

Tiered storage equations

$$\text{Tier1}_{\text{store}} = \sum (\text{N}_{\text{hts}} \times \text{G}_{\text{bpr}} + (\text{N}_{\text{hts}} \times \text{G}_{\text{bpr}}) / 4) (\times \text{N}_{\text{runs}})$$

N_{hts} =number of per type HTS devices, G_{bpr} =Gigabytes per run

$$\text{Tier2,3}_{\text{store}} = \sum (\text{N}_{\text{runs}} \times \text{G}_{\text{analysis}} + (\text{N}_{\text{runs}} \times \text{G}_{\text{analysis}}) / 3)$$

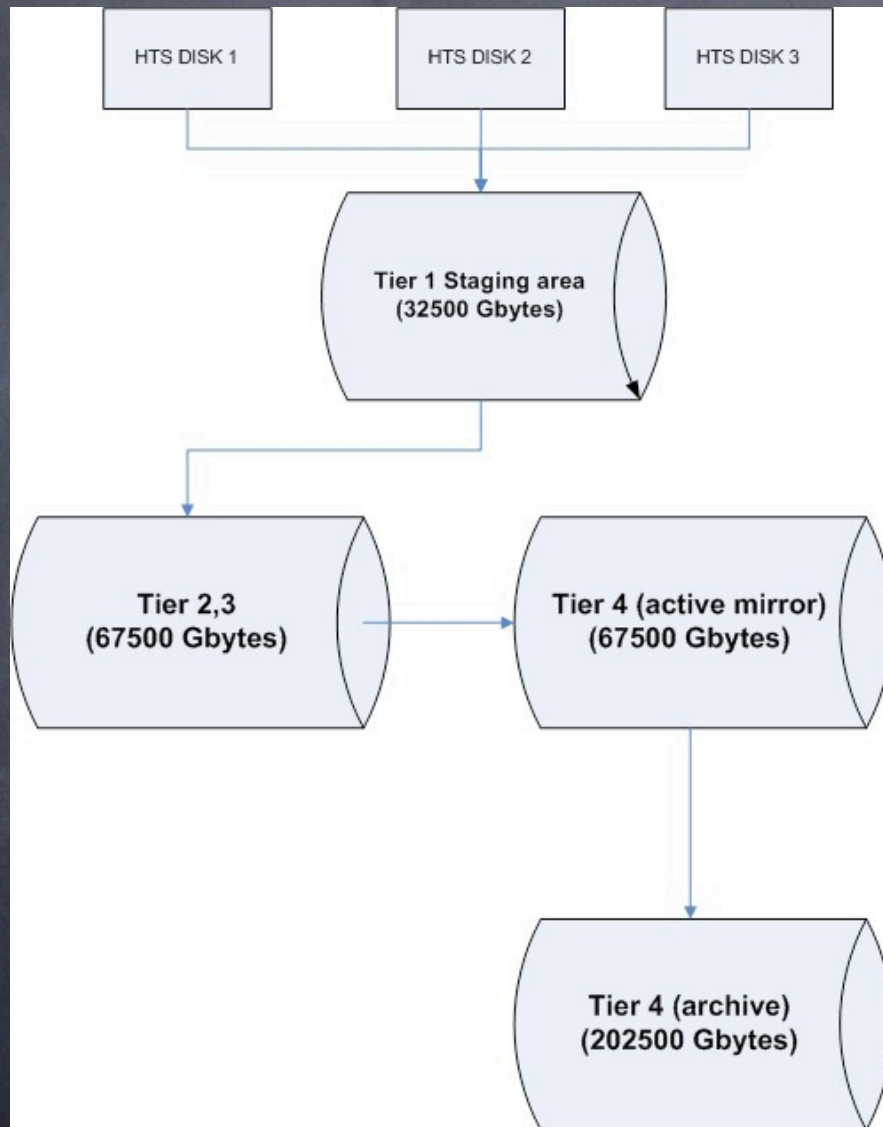
N_{runs} =expected number of runs per year,

$\text{G}_{\text{analysis}}$ =Gigabytes per run for Tiers 2 and 3 (Table 1)

$$\text{Tier4}_{\text{store}} = \text{Tier2,3}_{\text{store}} + \text{R}_{\text{period}} \times \text{Tier2,3}_{\text{store}}$$

R_{period} = number of years to keep the data4_EPamhp

tiered storage example



- 2 x Illumina
- 2 x 454
- 1 x SOLiD
- 3 year data retention period

Filesystems

A filesystem is a key component of the Operating System that dictates how the files are stored and accessed.

Commonly used disk filesystems: ext3/4 (Linux), NTFS (Windows), HFS+ (MACOSX), ZFS (Sun), XFS (SGI)

Shared/clustered/SAN filesystems: GFS (RedHat), XSAN (Apple)

Distributed File Systems (Network File Systems): NFS(9), CIFS/SMB

Distributed parallel fault-tolerant file systems: GPFS (IBM), XtremeFS, OneFS (Isilon), PanFS (Panasas), Lustre

Talleres Internacionales de Bioinformática - UNAM- Enero 2012

Filesystem Requirements

- Next Gen Sequencing (NGS) filesystems need to:
 - Be scalable in size
 - Be Scalable in the number of IOPS for read/writes/nested directory access
 - Allow concurrent access
 - Have file redundancy/replication features

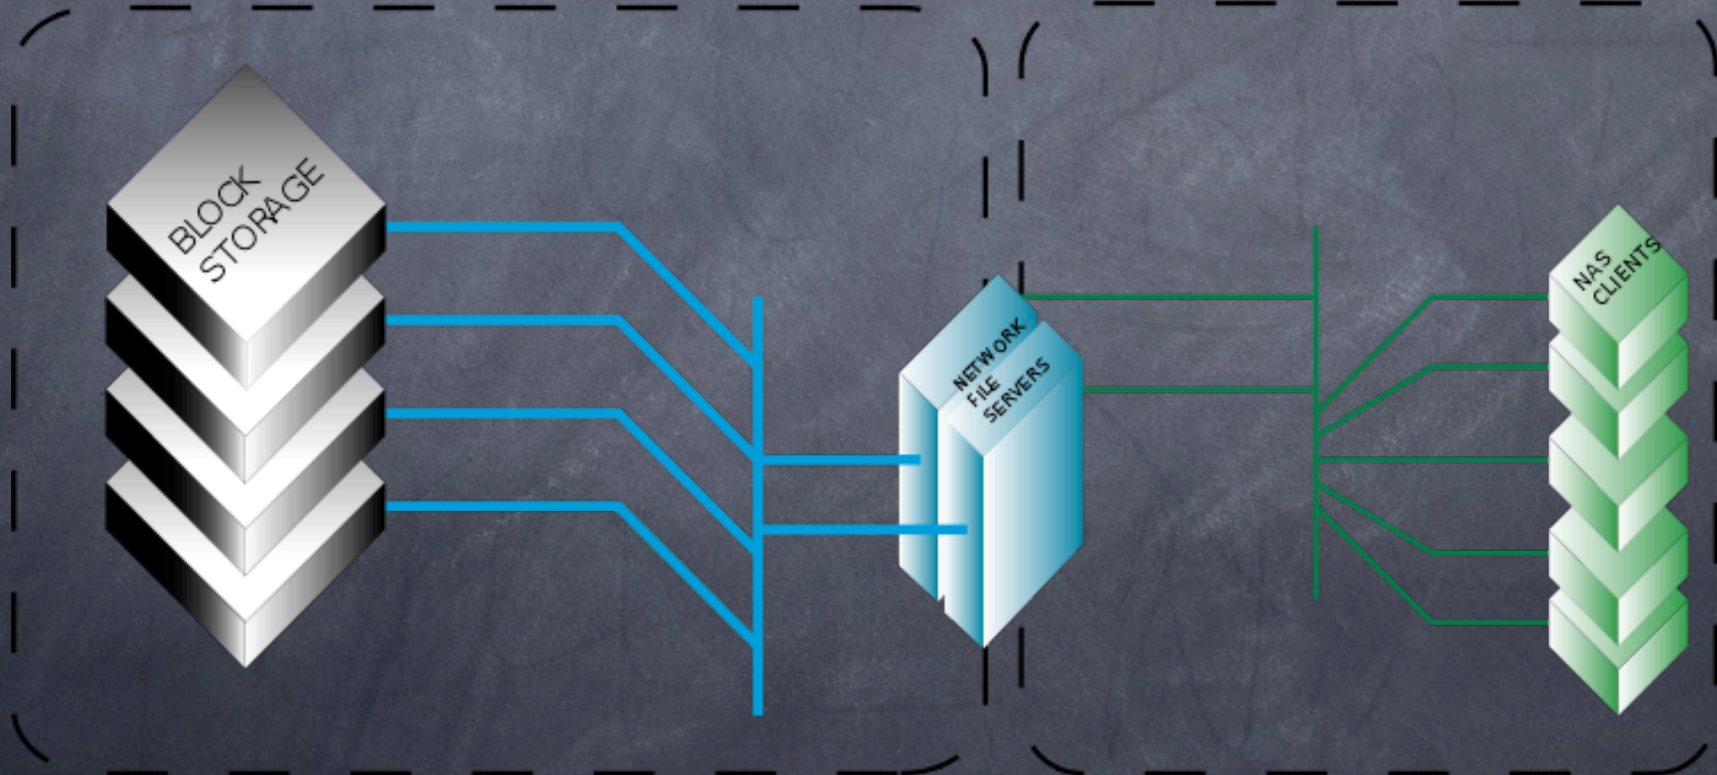
The data network and storage (2)

SAN

FiberChannel, iSCSI, or AoE

NAS

SMB, NFS, AFS



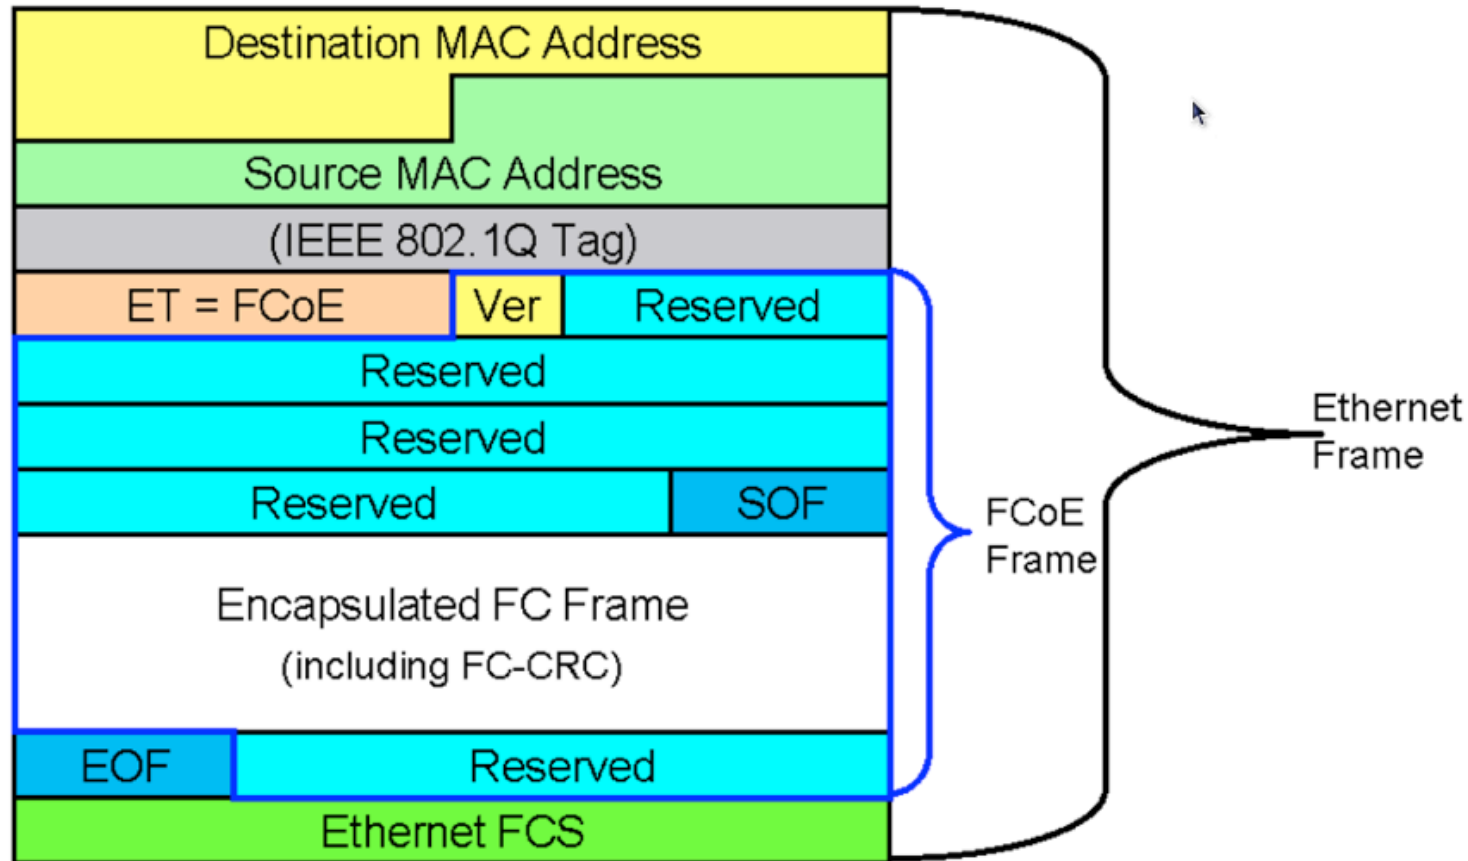
Talleres Internacionales de Bioinformática - UNAM- Enero 2012

The data network and storage (3)

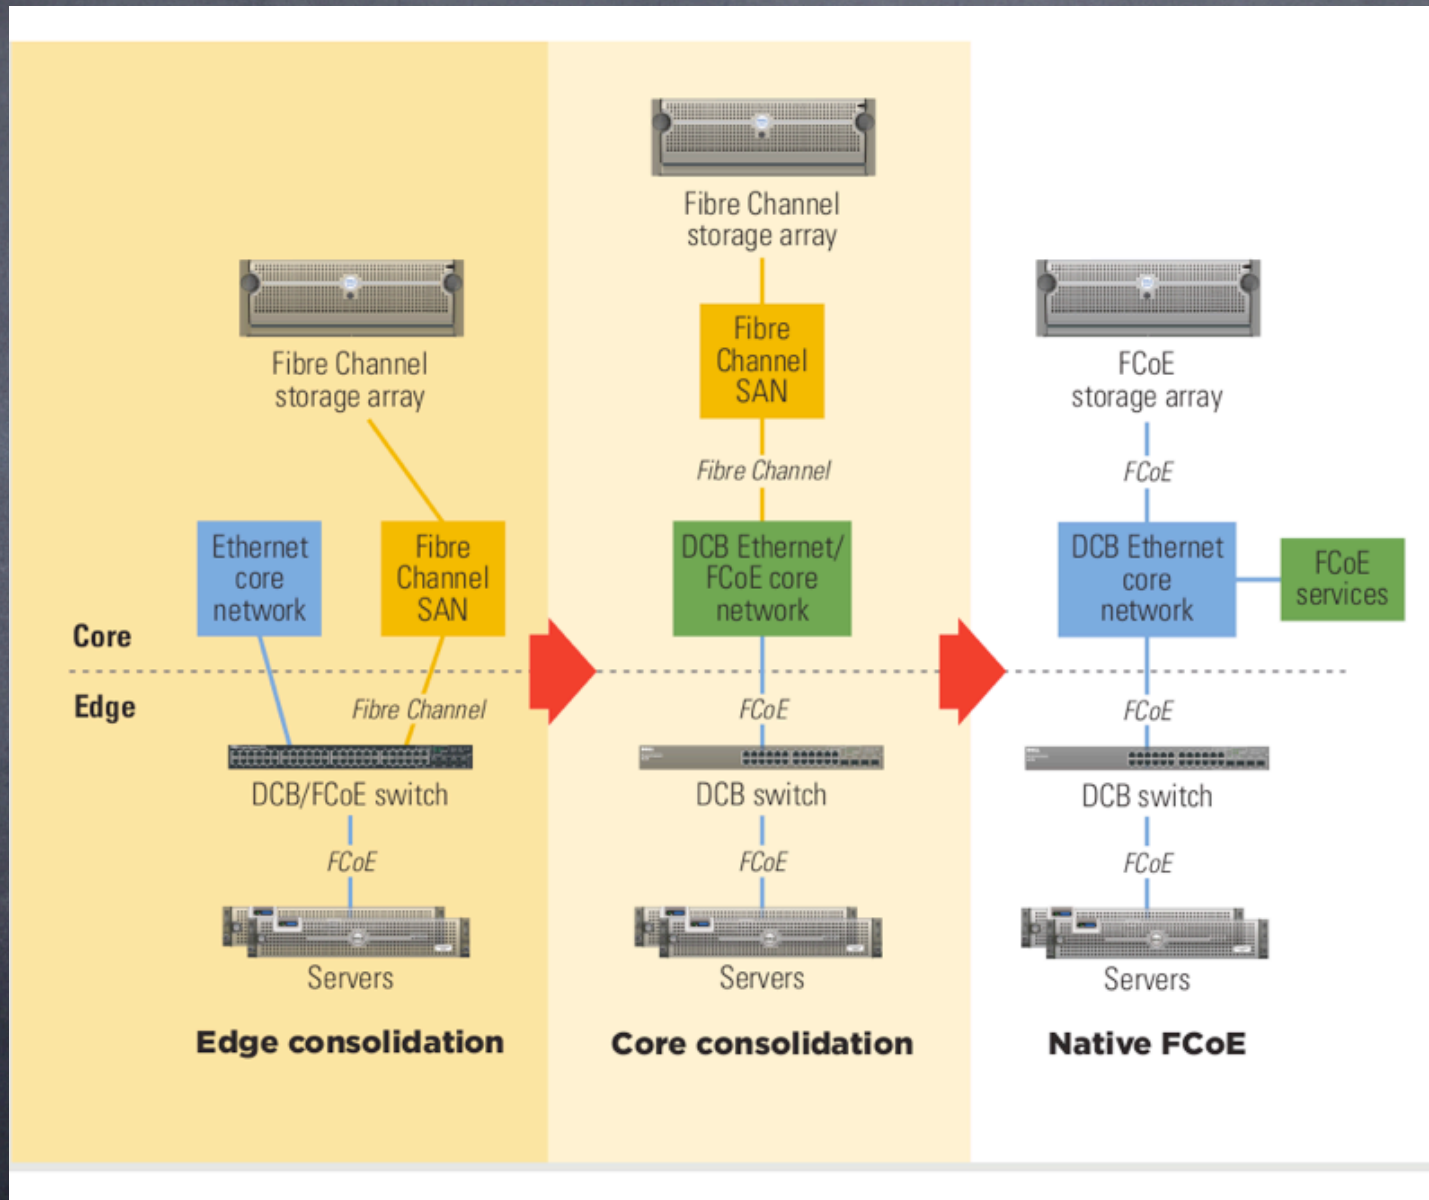
- Questions for your IT architect/system administrator(s):
 - Can you afford the pure Fiber Channel solutions today?
 - How many storage interconnects you have (GigE, FC, Infiniband).
 - Would it not be nice to have a smaller number of storage interconnects (consolidation)?

The data network and storage (4)

FCoE



The data network and storage - FCOE(5)



NGS Computing

- Multi-core CPU intensive jobs : Image processing, Quality Score
- Multi-core memory intensive jobs: Sequence Assembly
- Sequence assembly is the most challenging computational task in NGS, besides data storage


Sequence assembly graphs

```

AACCGTTAAGACCAAGTCTTTCCGACTCTCGA x 4
ACCGTTAAGACCAAGTCTTTCCGACTCTCGAC x 2
ACCGTTAAGACCAAGTCTTTCCGACTCTCGGC x 2
CCGTTAAGACCAAGTCTTTCCGACTCTCCACT x 1
CGTTAAGACCAAGTCTTTCCGACTCTCCGCTG x 2
GTTAAGACCAAGTCTTTCCGACTCTCGACTCG x 1
GTTAAGACCAAGTCTTTCCGACTCTCGACTCG x 1
TTAAGACCAAGTCTTTCCGACTCTCCACTCGA x 2
TTAAGACCAAGTCTTTCCGACTCTCCGCTCGA x 1
TAAGACCAAGTCTTTCCGACTCTCCACTCGAA x 2
TAAGACCAAGTCTTTCCGACTCTCCGCTCGAA x 2
TAAGACCAAGTCTTTCCGACTCTCCACTCGAA x 1
TAAGACCAAGTCTTTCCGACTCTCCACTCGAA x 1
TAAGACCAAGTCTTTCCGACTCTCCACTCGAA x 1
GACCAAGTCTTTCCGACTCTCCGCTCGAACCT x 1
GACCAAGTCTTTCCGACTCTCCACTCGAACCT x 1
ACCAAGTCTTTCCGACTCTCCACTCGAACCTT x 1
CCAAGTCTTTCCGACTCTCCACTCGAACCTTT x 1
CAAGTCTTTCCGACTCTCCACTCGAACCTTTA x 1
CAAGTCTTTCCGACTCTCCACTCGAACCTTTA x 1
AAGTCTTTCCGACTCTCCGCTCGAACCTTTAG x 1
AAGTCTTTCCGACTCTCCACTCGAACCTTTAG x 1
AGTCTTTCCGACTCTCCGCTCGAACCTTTAGG x 1
GTCCTTTCCGACTCTCCACTCGAACCTTTAGGT x 1
GTCCTTTCCGACTCTCCGCTCGAACCTTTAGGT x 1
TCCTTTCCGACTCTCCGCTCGAACCTTTAGGTG x 2
TCCTTTCCGACTCTCCACTCGAACCTTTAGGTG x 1
CFTTCCGACTCTCCACTCGAACCTTTAGGTGT x 1
CFTTCCGCTCTCTCCGCTCGAACCTTTAGGTGT x 1
TTTCCGACTCTCCACTCGAACCTTTAGGTGTA x 2
TTTCCGACTCTCCGCTCGAACCTTTAGGTGTA x 1
TTCCGACTCTCCACTCGAACCTTTAGGTGTA x 2
TCGGACTCTCCACTCGAACCTTTAGGTGTA x 3
CCGACTCTCCGCTCGAACCTTTAGGTGTA x 1
CCGACTCTCCGCTCGAACCTTTAGGTGTA x 1
GGAACCTCCGCTCGAACCTTTAGGTGTA x 1
GACTCTCCGCTCGAACCTTTAGGTGTA x 1
ACTCTCCACTCGAACCTTTAGGTGTA x 1
CTCTCCGCTCGAACCTTTAGGTGTA x 1
CTCTCCACTCGAACCTTTAGGTGTA x 1
CTCCACTCGAACCTTTAGGTGTA x 1
TCCACTCGAACCTTTAGGTGTA x 2
TCCGCTCGAACCTTTAGGTGTA x 1
CCGCTCGAACCTTTAGGTGTA x 1
TTGGCAATTCGGTTCAGAAAGCCTGAGAGCCGAGCTTGAAATCCACATTTCTCGGCTGC
    
```

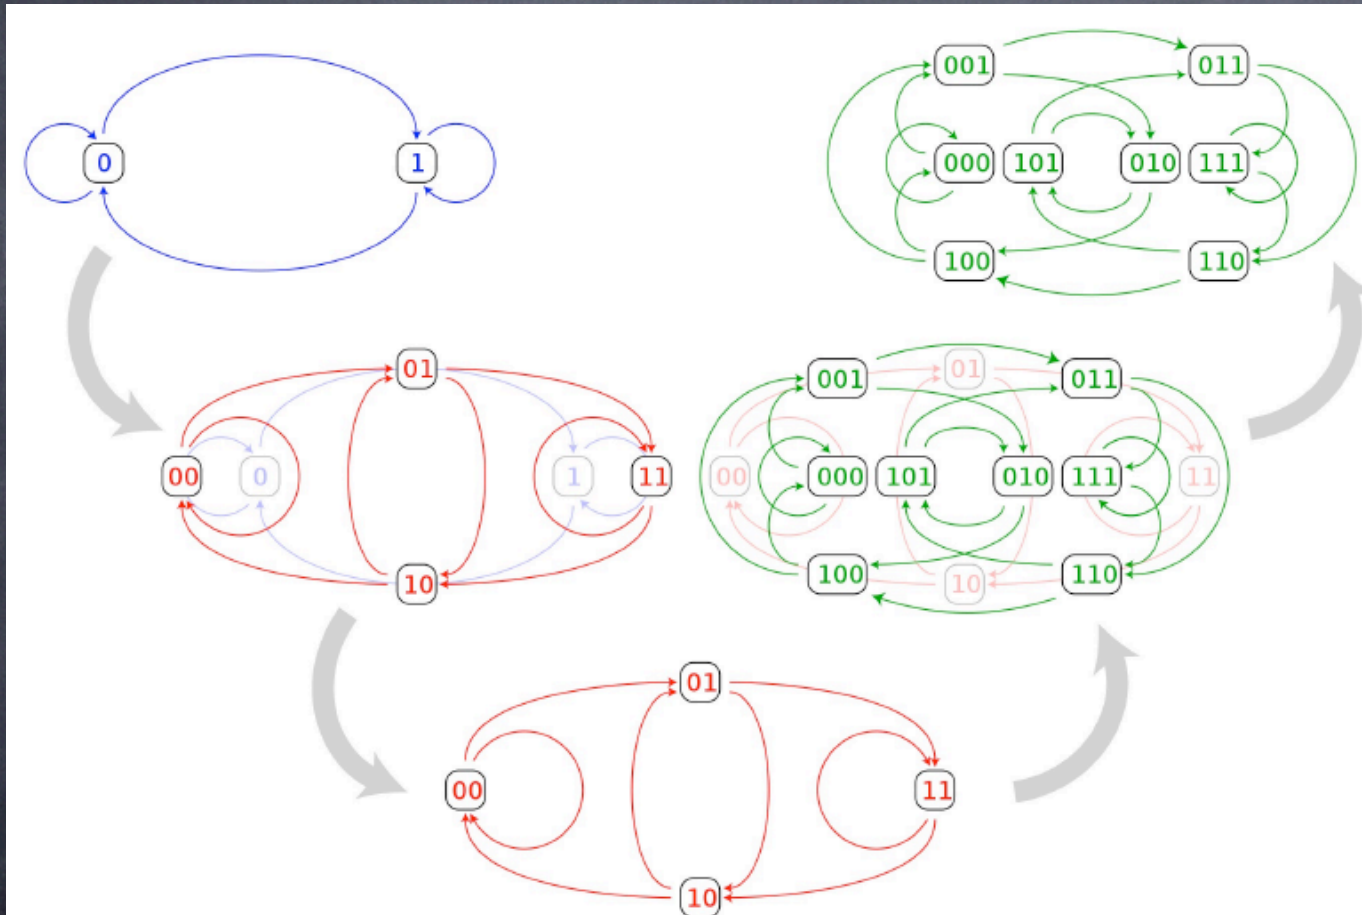
Sequencing Reads

R_1 ACGATACCGTAGA
 R_2 TACCGTAGATGCA
 R_3 GTAGATGCAGTAC



Graph Data Structure

De Bruijn Graph



<http://www.nature.com/nbt/journal/v29/n11/full/nbt.2023.html>

How to get the complete sequence from short read overlapping sequences

Common Sequence Assemblers

- Velvet: <http://www.ebi.ac.uk/~zerbino/velvet/>
- ALLPATHS-LG: <ftp://ftp.broadinstitute.org/pub/crd/ALLPATHS/Release-LG/>
- SOAPdenovo: <http://soap.genomics.org.cn/soapdenovo.html>
- Cortex: <http://cortexassembler.sourceforge.net/>
- ABySS: <http://www.bcgsc.ca/platform/bioinfo/software/abyss>
- Curtain: <http://code.google.com/p/curtain/>

Sequence assembly requirements

- Installation: OpenMP, other middleware and libraries
- Large RAM requirement:
 - Get at least one memory fat node: RAM \geq 256 Gbytes (512 Gb to 1Tb of RAM are common)
 - Standalone or part of a queue/batch system
 - Multi/core \leftrightarrow Multi-thread


File Edit View Search Terminal Help

```

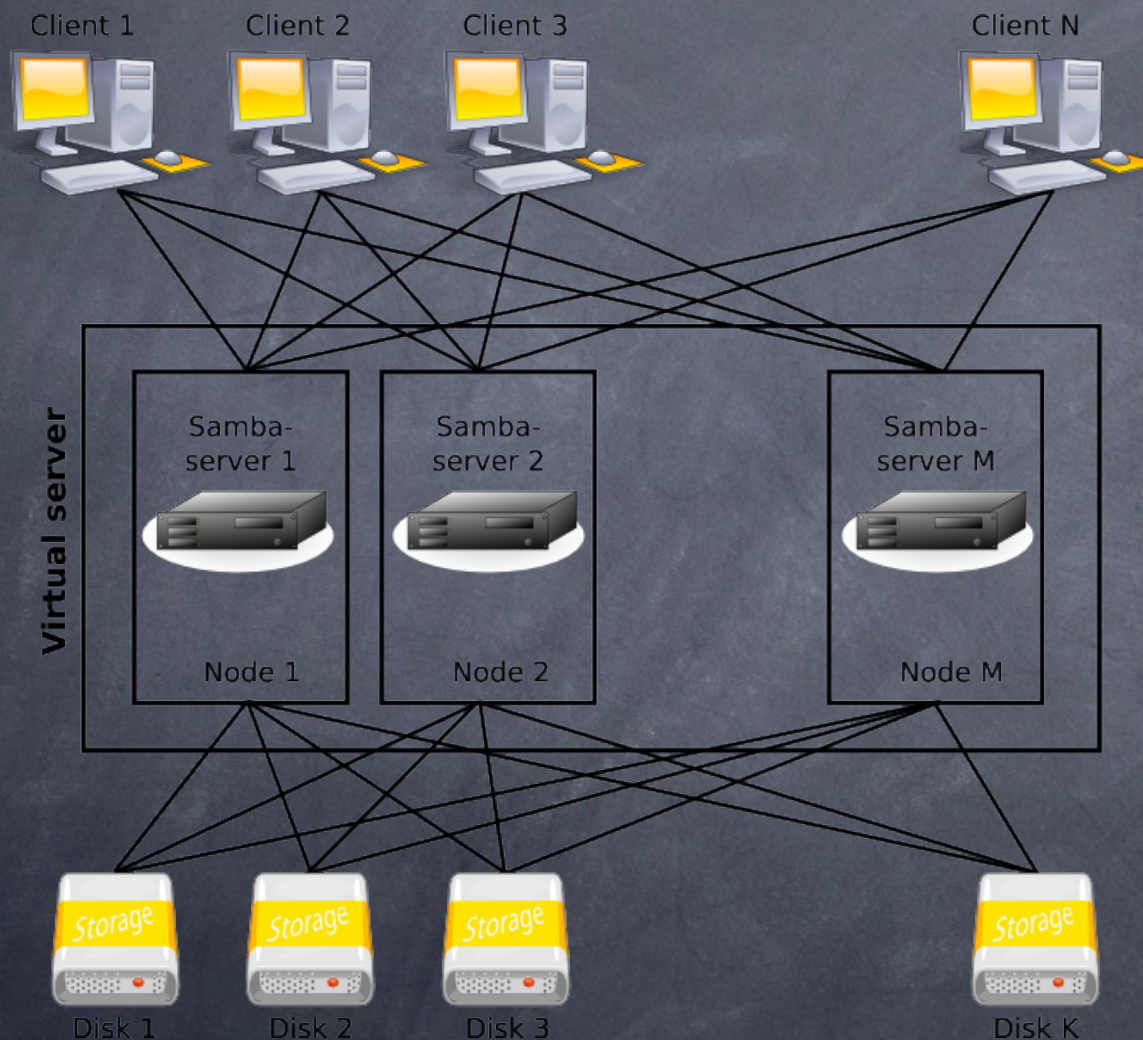
top - 23:24:56 up 1 day, 15:36, 1 user, load average: 32.18, 32.22, 32.20
Tasks: 996 total, 33 running, 963 sleeping, 0 stopped, 0 zombie
Cpu(s): 73.2%us, 26.8%sy, 0.0%ni, 0.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 264532284k total, 236729744k used, 27802540k free, 80780k buffers
Swap: 10485752k total, 0k used, 10485752k free, 159672436k cached

```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
12101	biotekab	20	0	2372m	2.0g	18m	R	100.0	0.8	301:25.17	Ray
12108	biotekab	20	0	2504m	2.1g	18m	R	100.0	0.8	303:39.61	Ray
12112	biotekab	20	0	2504m	2.1g	18m	R	100.0	0.8	303:43.25	Ray
12117	biotekab	20	0	2502m	2.1g	18m	R	100.0	0.8	302:35.79	Ray
12121	biotekab	20	0	2504m	2.1g	18m	R	100.0	0.8	303:56.15	Ray
12123	biotekab	20	0	2504m	2.1g	18m	R	100.0	0.8	303:39.20	Ray
12127	biotekab	20	0	2504m	2.1g	18m	R	100.0	0.8	303:27.20	Ray
12129	biotekab	20	0	2504m	2.1g	18m	R	100.0	0.8	303:42.23	Ray
12100	biotekab	20	0	2372m	2.0g	18m	R	99.6	0.8	300:27.78	Ray
12102	biotekab	20	0	2372m	2.0g	18m	R	99.6	0.8	298:05.20	Ray
12104	biotekab	20	0	2371m	2.0g	18m	R	99.6	0.8	299:11.82	Ray
12105	biotekab	20	0	2371m	2.0g	18m	R	99.6	0.8	299:06.95	Ray
12106	biotekab	20	0	2452m	2.1g	18m	R	99.6	0.8	303:27.31	Ray
12109	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:10.03	Ray
12110	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:08.66	Ray
12113	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:23.37	Ray
12114	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:33.50	Ray
12115	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:40.59	Ray
12118	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	302:43.28	Ray
12119	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:01.71	Ray
12120	biotekab	20	0	2502m	2.1g	18m	R	99.6	0.8	302:42.04	Ray
12122	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:25.47	Ray
12124	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:21.28	Ray
12125	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	302:21.90	Ray
12126	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	302:54.00	Ray
12128	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	303:19.78	Ray
12130	biotekab	20	0	2372m	2.0g	18m	R	99.6	0.8	300:30.61	Ray
12131	biotekab	20	0	2504m	2.1g	18m	R	99.6	0.8	301:48.06	Ray
12103	biotekab	20	0	2372m	2.0g	18m	R	99.3	0.8	298:35.48	Ray
12107	biotekab	20	0	2504m	2.1g	18m	R	99.3	0.8	303:26.74	Ray
12111	biotekab	20	0	2504m	2.1g	18m	R	99.3	0.8	303:52.98	Ray
12116	biotekab	20	0	2504m	2.1g	18m	R	99.0	0.8	303:37.71	Ray
57419	root	20	0	19884	1988	944	R	1.0	0.0	0:02.69	top
3106	root	20	0	192m	5916	1756	S	0.3	0.0	3:09.93	snmpd



Tier 1 Architecture



samba
ctdb

<http://ctdb.samba.org/>

Bill of materials (data network)



Cisco Nexus Switch 5000 series



QLE8152
Dual Port 10GbE Ethernet to
PCIe Converged Network
Adapter (CNA).
www.qlogic.com

Bill of materials (storage and computing)




Dell | EMC CX4-960
(8Gbit and 4Gbit FC/
10 with FCoE
support modules)



Dell R815 Memory fat
node - 32 cores- 512
Gbytes of RAM



Dell 1950, 64 Gbytes
of RAM/Qlogic CN
cards (as access/
front end nodes),
8 cores.



Questions?

admin@embnet.uio.no

and

<http://www.embnet.org/join/ContactRegistration>

Talleres Internacionales de Bioinformática - UNAM- Enero 2012