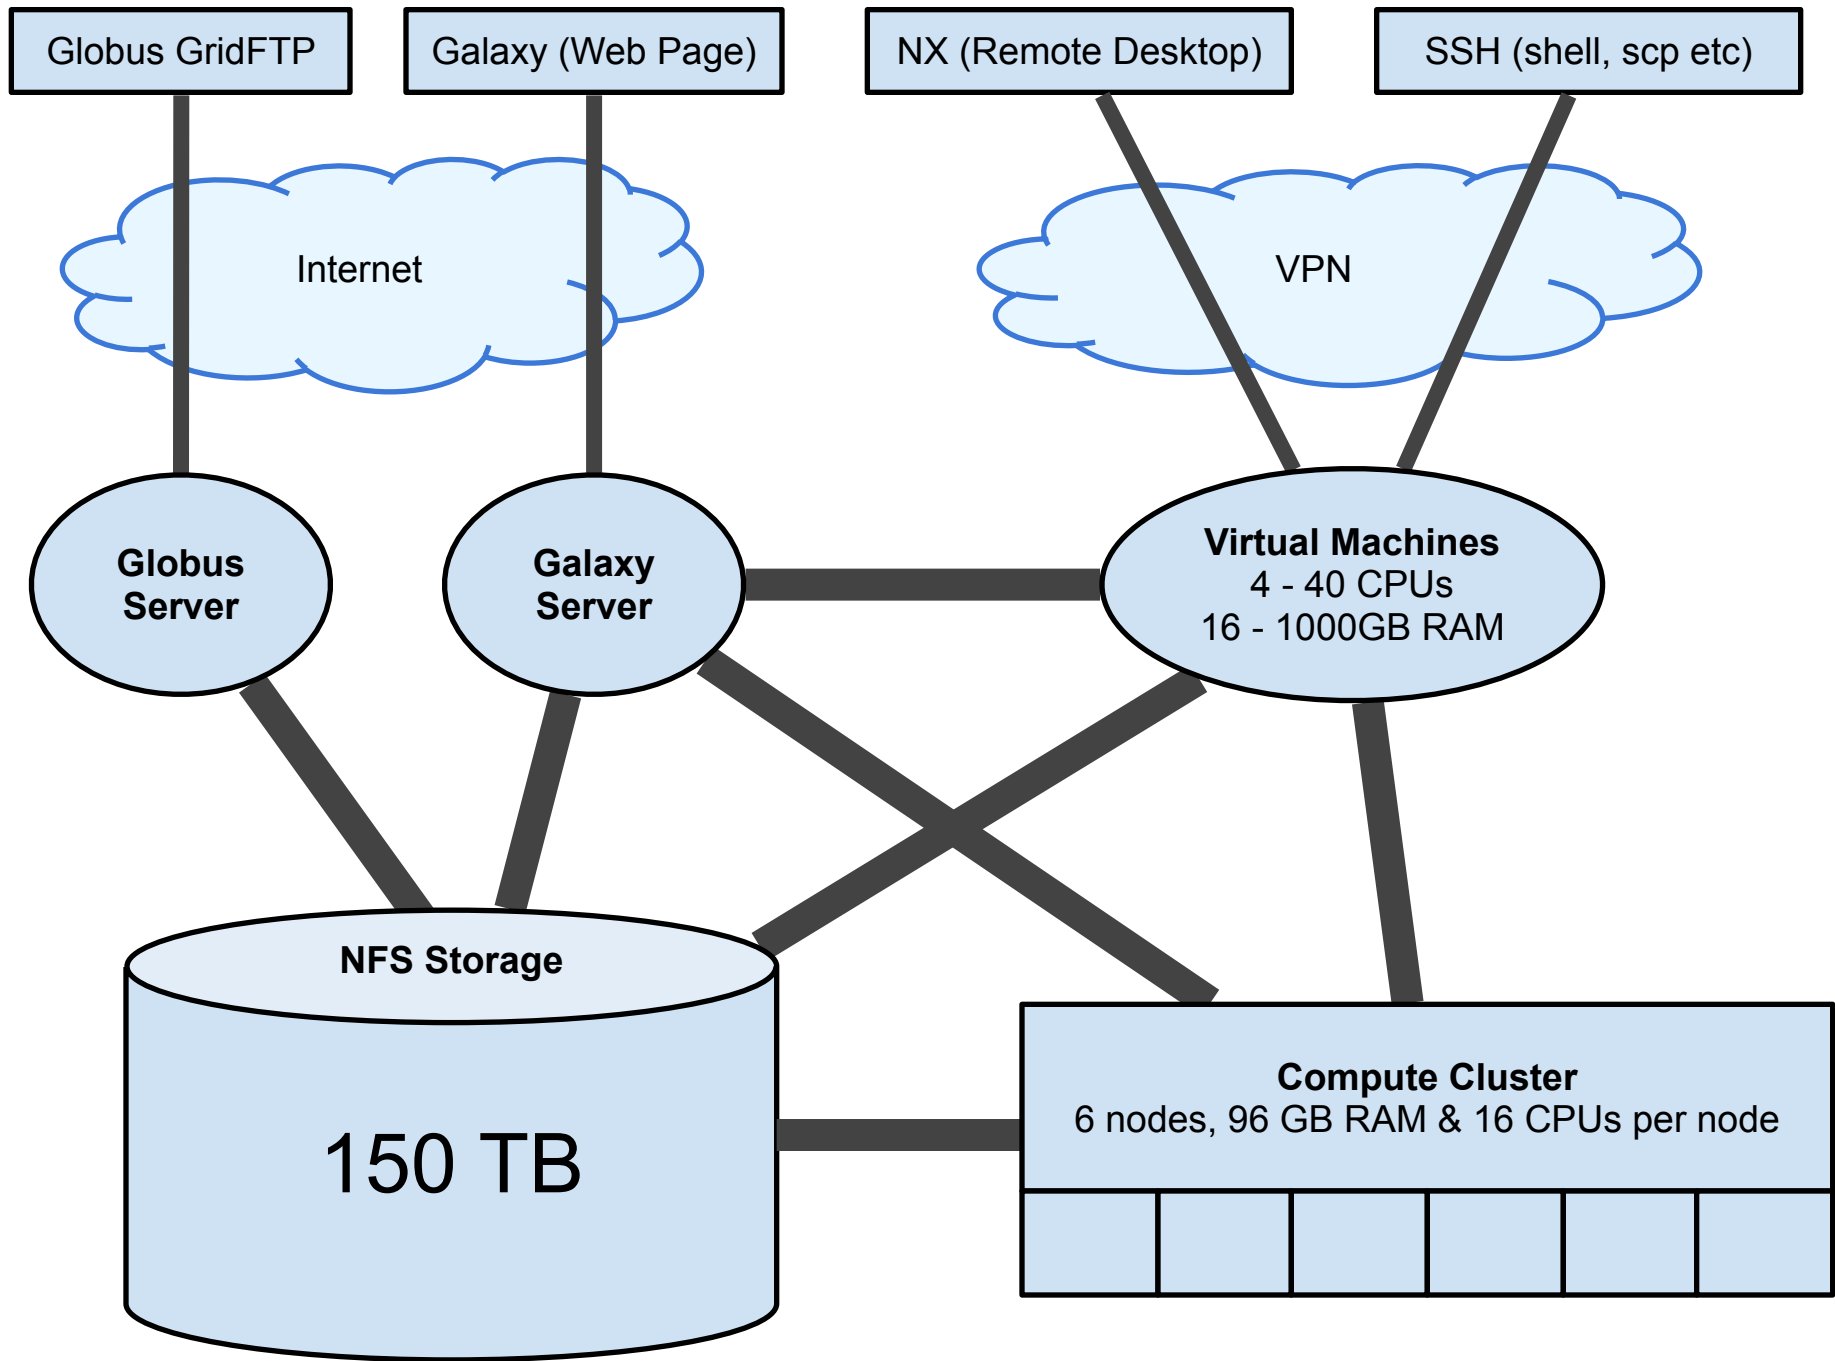


Running and maintaining a reliable production Galaxy server

Shane Sturrock

New Zealand Genomics Ltd





Updating our Galaxy instance

- Many comments this week about how updates can break Galaxy
- Before we do anything, update all the tool_shed tools to their latest versions
- Warn the user base that the update is coming
- We run a second Galaxy server which is a complete mirror of the main production server



Mirror production server

- On the live production server dump the postgres database to a file
- On an admin VM which mounts all file systems
`rsync /home/galaxy with /home/galaxy-test`
- Start the galaxy mirror machine which mounts
`/home/galaxy-test as /home/galaxy`
- Slight modification to universe file to remove cluster support but otherwise the same
- Replace the NZGL logo graphic with one that shows this is the test server to avoid mistakes



Update the test server

- Run the update
- Test it - is it broken? Probably so figure out why
- rsync from live server again and do the update again
- What broke this time? (hopefully less)
- Repeat until I eventually get an update that works (usually takes me a couple of goes)
- Do it once more for luck on the test server to be sure it is repeatable and then shut it down

Update the real server

- One last backup of the database and rsync everything over to the test server's home directory as a backup so this time don't edit anything or run it.
- Send a final warning to any users - since we run jobs on a cluster their jobs in flight should be fine
- Do the update on the real server and check it all works
- If there is something catastrophic there is still a backup of the original configuration in the test server storage so that can be recovered

Updating BLAST databases

- BLAST databases are downloaded in compressed form
- After decompression, a script runs the BLAST database info tool and extracts the name and date for each
- Script checks for new versions and inserts them into the various .loc files in a predefined order while pushing older versions down the list

Reference databases

- We use a script which parallelises the formatting of bowtie, bowtie2, bwa, bismark and samtools indexes
- The script automatically inserts the new databases into the .loc files