GeoCloud Project Report – - LarvaMap tool for modeling trajectories of larval fish

Description of Application

National Oceanic and Atmospheric Administration/National Marine Fisheries Service/ Alaska Fisheries Science Center

The tool is intended to be used by resource managers and fisheries scientists to model the dispersion of fish larvae under varying environmental conditions.

Windows, Apache Tomcat, JRE, Sun Java SMQ, eventually may deploy with ArcServer front end but not done under GeoCloud project

Operational Requirements:

Image type (RAM, local disk) – 4GB RAM

Data storage – 45GB to store model results

Upload monthly – 35GB uploaded once, 1.5GB/day for daily circulation model results to support particle tracking model.

Download monthly – depends on user load, typical user would download kb's to low mb's per model run, estimated maximum of 10GB/month under heavy use.

Deployment in the Cloud

Using existing images – yes. No need to reconfigure.

Loading application with data – simple on-time upload of underlying data files.

Customizing application suite – N/A

Installation scripting – took 2 days in consultation with eGlobalTech

Operations in the Cloud

Monitoring of operations – system not mission critical so relied upon standard Amazon monitoring. Only one situation where there was a technical problem.

Discussion – Basically hassle-free. But our performance requirements are simple.

Operational cost comparison (extrapolate to one year)

Fixed (one-time) costs versus monthly costs

Based on a "server" being a 64bit machine with 8GB RAM, (2) CPU cores, and a “free” operating system. Based on an IT staff rate of $100/hour.

Backup costs
We thought of taking a “snapshot” of all EBS volumes attached to a cloud server as the equivalent to a typical backup of a server in a traditional IT shop.

Cloud backups were quick and could be automated. The only cost associated with the backup was the storage space used by the backup ($0.10 per allocated GB per month). Since EBS snapshots are incremental the storage space used was usually very close to the size of the EBS disk.

Depending on the software, hardware, and offsite service contract, an IT shops backup solution can be a substantial cost. As an example, a BackupExec license for a single Linux host can cost ~$500. There are benefits to a complete backup solution like BackupExec, for example, the ability to restore individual files.

Since LarvaMap can be redeployed very easily, a robust backup solution is not required. Backup costs don’t play a factor when comparing costs. It was worth mentioning here because it could potentially play a huge role in cost comparisons.

Electricity costs
UPS costs
Firewall costs
Physical space costs
AC costs
Network costs
Dedicated IT Staff costs
For purposes of comparison, we lumped the above into one figure. They are nearly impossible to quantify on a server by server basis. The estimate is based on providing all of the above to two servers for a year and an IT staff spending 1 hour/month dedicated to the two servers.

The costs of hardware failure are hidden from cloud users. However, there still needs to be staff that regularly monitor the server and respond to problems. We estimated that they spend 2 hours/month dedicated to the two cloud servers.

Cloud: +$200/month (+$2400/year)

Physical: +$600/month (+$7200/year)

Upgrade costs
Hardware depreciates relatively quickly and needs to be replaced every give years (maybe an overestimate). In the case of LarvaMap, the hardware purchases are relatively low cost and do not play a substantial role in comparing costs.

Hardware failure
The costs of hardware failure are hidden from cloud users. Surely they have hard drives fail all the time, but the cost is built into the hourly rate.

If a hard drive fails in a server, the usual practice is to call the vendor and get a replacement shipped out as quickly as possible (and hopefully under warranty). In the best case there would be a replacement drive on site ready to replace the faulty drive and star the process of the RAID being rebuilt. For comparison purposes, we estimate that it would take an IT staff 6 hours to rectify a hard drive failure in the best situation and that there would be one failure per year.

Cloud: $0
Physical: +$50/month (+$600/year)

Telecommunications

Operations and maintenance support

Security plan development and approval – N/A as GeoCloud is bearing the cost of this.

Issues and Lessons Learned

Security approval process - We would have been unable to start the security approval process without participating I the GeoCloud SandBox. We have no local expertise to support a C and A for the cloud and we do not anticipate having enough cloud users to warrant gaining this expertise.

Recommendations on C&A – Must be done through this type of centralized, multi-agency approach.

Software deployment – The redeployment of the software was remarkably simple. We see very few impediments to this type of redeployment in the future.

Time-to-deploy - Minimal

Failover, redundancy

Project planned future environment – We are currently reengineering parts of the application but anticipate redeploying, at least for testing, in the cloud in the future. We have also been able to share our experiences with other NOAA projects, such as the ERMA development team, that have a much more immediate need to cloud deployments.

Point of contact: Tiffany C. Vance, tiffany.c.vance@noaa.gov, 206-526-6767