



# CloudLightning & HPC: Heterogeneous Computing in the Cloud(\*)



(\*) Funded in part by EU H2020 R&I Project CloudLightning under Grant Agreement No. 643946. (Jan. 2015-Jan. 2018)

Tobias Becker†, Anne C. Elster\*, Geir Amund Hasle\*, Malik M. Khan\*, Perumal Kuppuudaiyar§, Suryanarayanan Natarajan§, Hristina Palikareva†, Muhammad Qasim\*. Norwegian University of Science and Technology, Trondheim, Norway; Maxeler Technologies, United Kingdom†; Intel, Ireland§

## CloudLightning (CL) Overview

### Self-Organization and Self Management (SOSM) of HPC resources

#### Resource Allocation - (Based on Service Requirements) which can be:

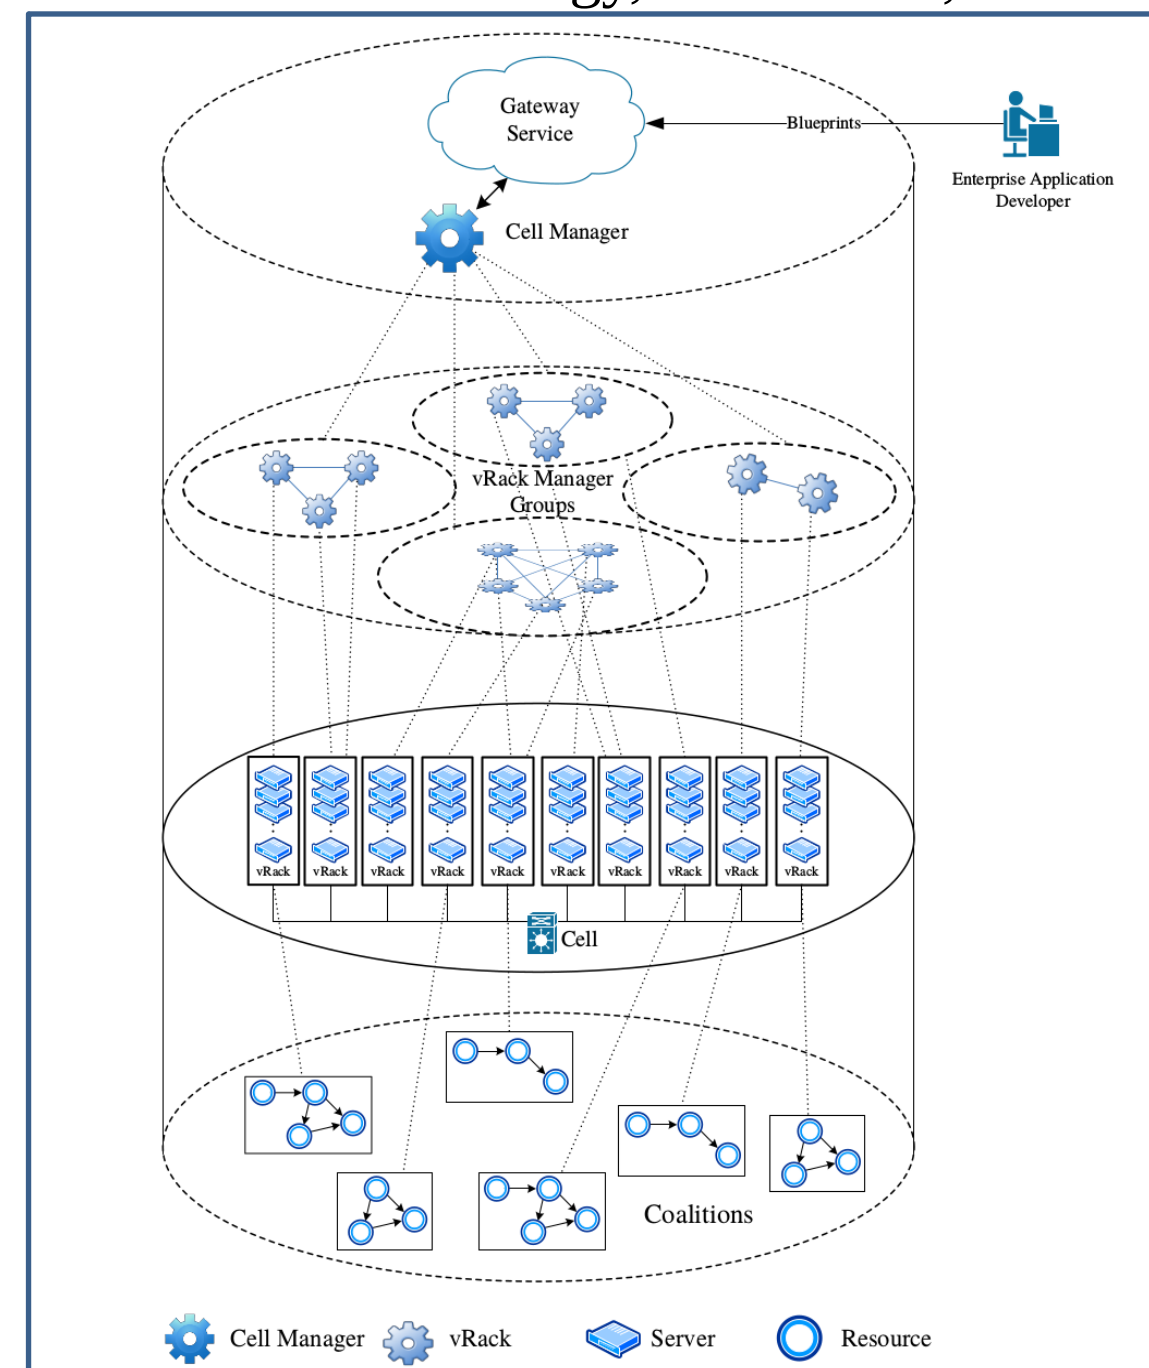
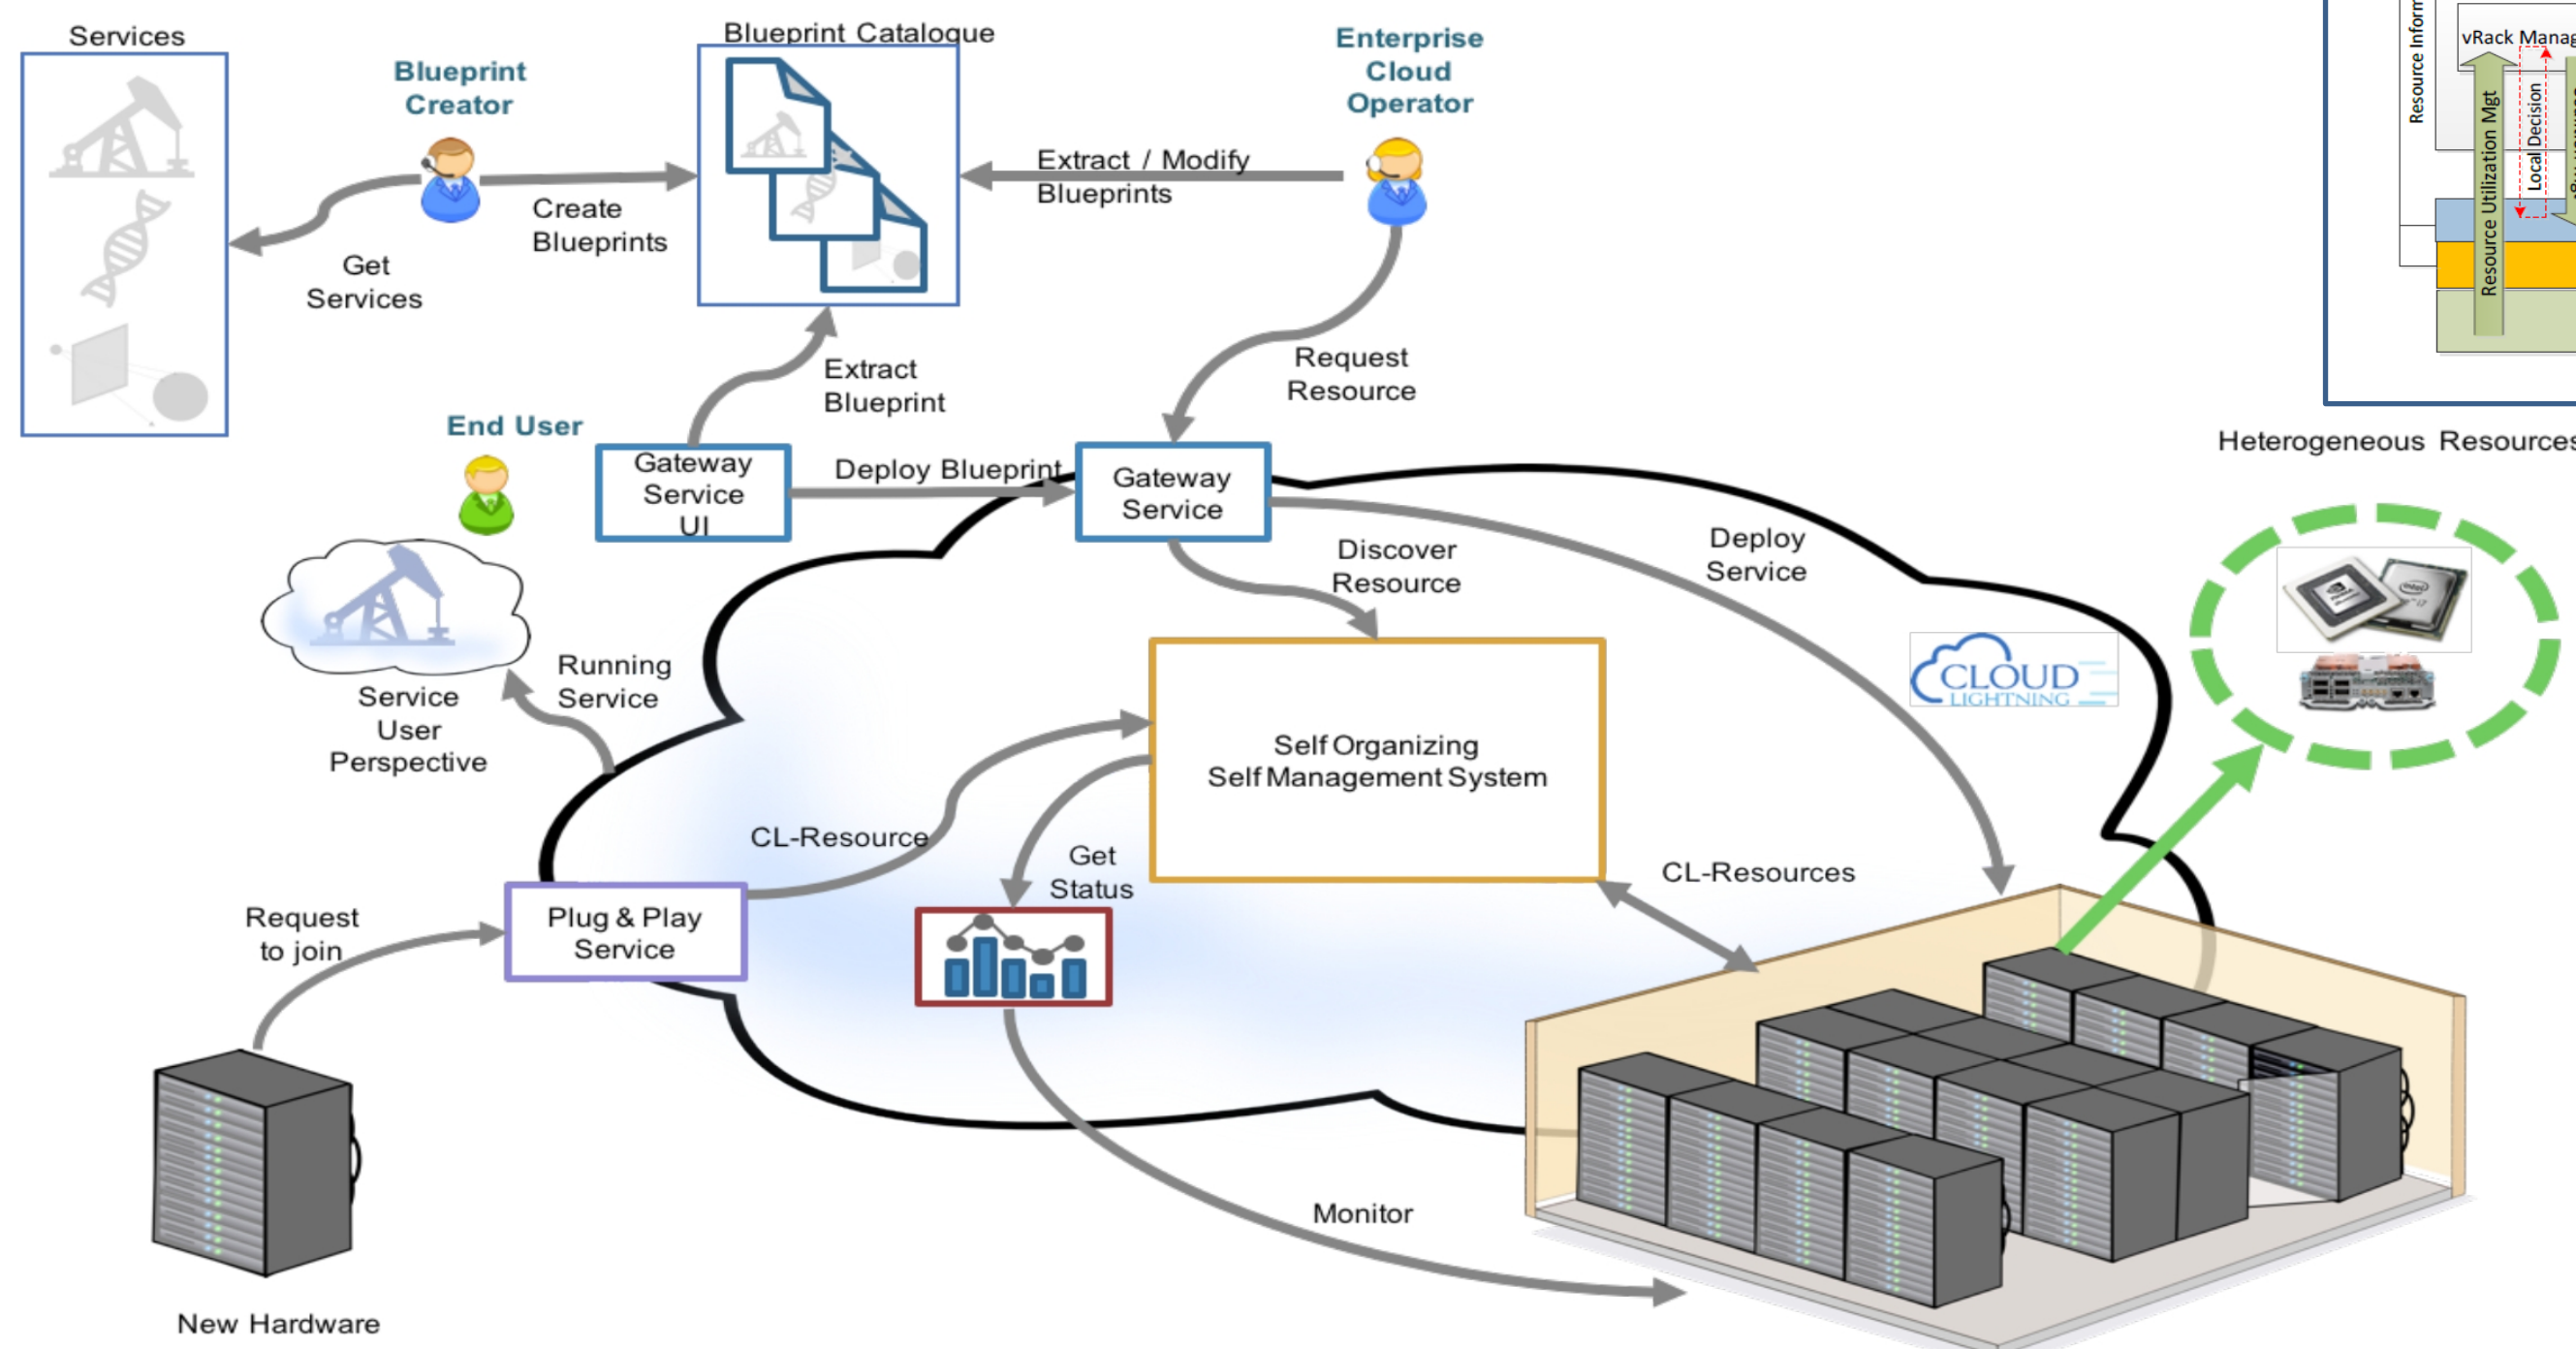
- Bare Metal
- VM (Virtual Machines)
- Containers –the most realistic option for HPC workloads
- Resources are divided into Cells, based on region or location
- Each Cell may have different hardware types, including servers with GPUs, MICs, DFEs
- Each Cell is partition into vRacks which are sets of servers of the same type

#### Resource Utilization - (Self-Optimized)

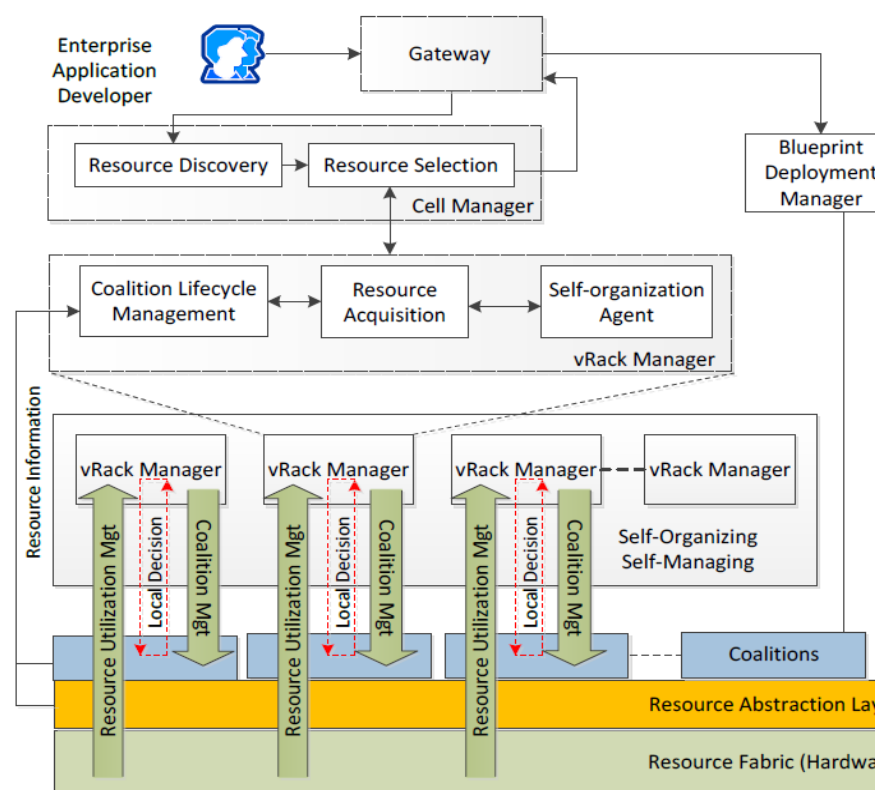
- Alt. 1: Cell manager get request of resources, create them and allocate them directly
- Alt. 2: First discover resources available, then if more than one options, then the Cell Mgr selects the most appropriate option.
- Alt. 3: Same as Alt. 2, except gives back solution rather than option
- vRack managers may create/aggregate the resources in its vRack and is the basic component of self-organization in the CL system. Note this feature makes only sense for larger deployments.

### CURRENT STATUS OF THE OVERALL CLOUDLIGHTNING PROJECT:

- SOSM Architecture defined
- Plugins for resource registration developed
- Tested use cases on individual platforms
- Working on integration, testbed and simulation that includes the OpenStack-based SOSM system



### CloudLightning Architecture



### Heterogeneous Testbed

- MIC System**
  - PC with Xeon Phi card
- GPU System**
  - Dell server blade with NVIDIA Tesla P100
- DFE Cluster**
  - Maxeler MPC-C node with 4x Vectis MAX3 DFEs
- SMP Cluster**
  - Numascale 5-node SMP

### References:


<http://www.cloudlightning.eu> <http://www.opm-project.org> <http://www.mcs.anl.gov/petsc/> <http://www.maxeler.com>

J. Kvalsvik: "Enhancing OPM-based Reservoir Simulation via PETSc integration", NTNU Masters thesis 2015, Advisors: Anne C. Elster(NTNU) and Alf B. Rustad (Statoil) Sven, et al. "Embrex ray tracing kernels for CPUs and the Xeon Phi architecture." ACM SIGGRAPH 2013 Talks. ACM, 2013

## Use Cases

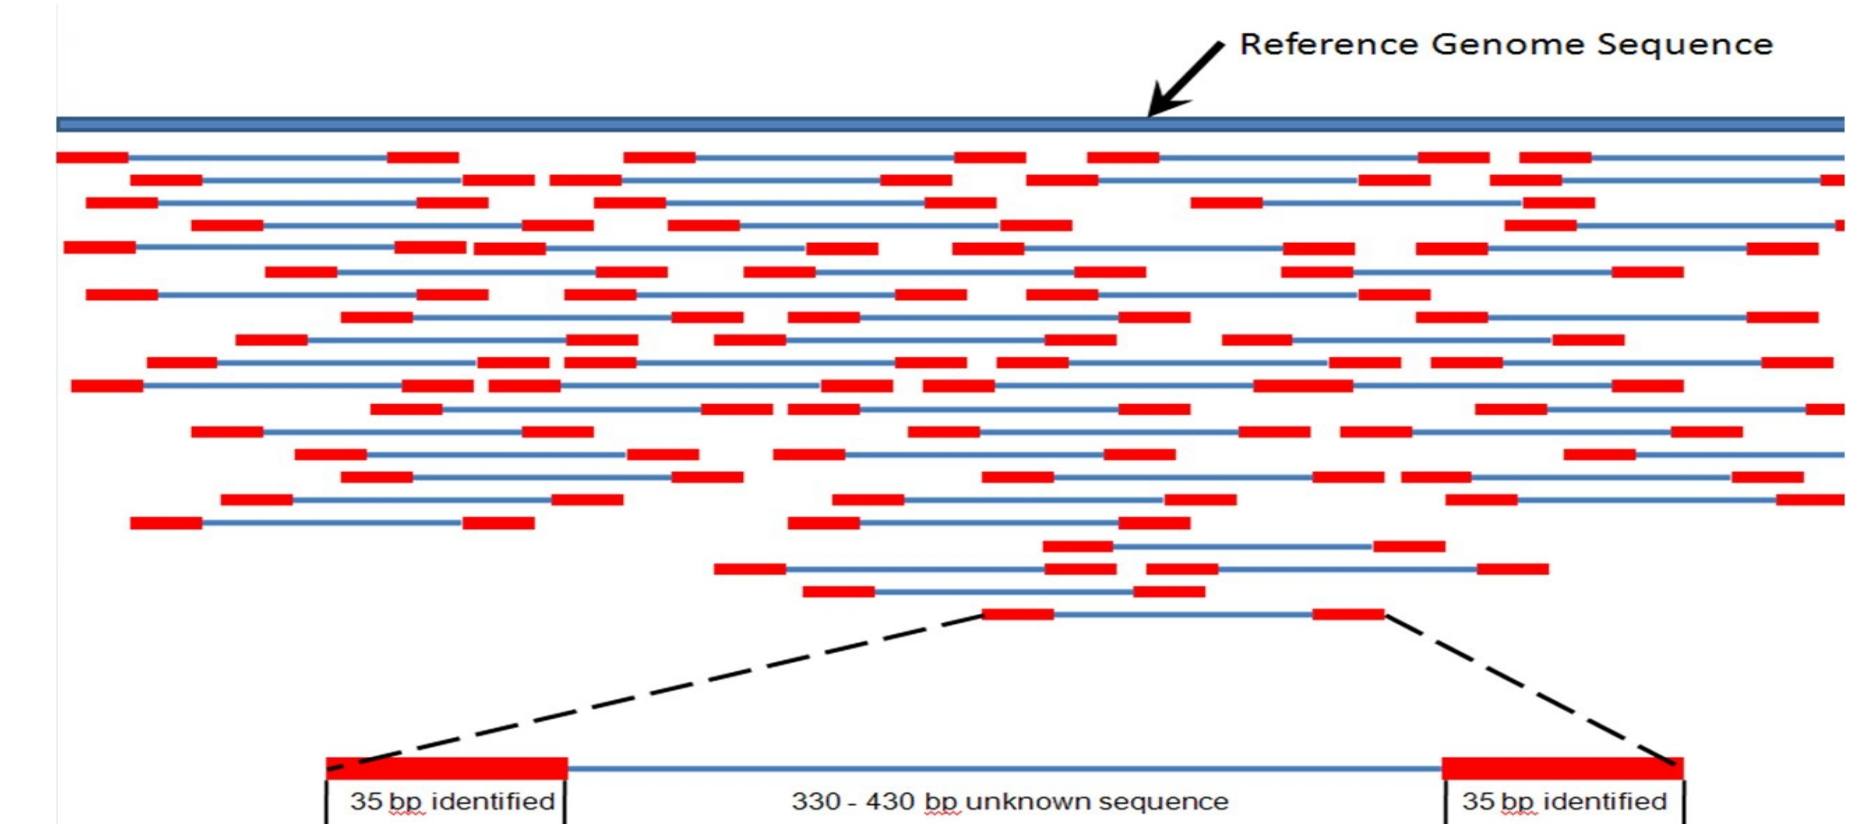
### Oil & Gas:

Reservoir simulation case using OPM upscaling of relative permeability



### Genomics:

Sequence alignment for whole genome analysis



### Self-Optimized Libraries

- ATLAS BLAS/MKL/cuBLAS/FFTW/cuFFT

### Ray Tracing:

Generating image by tracing path of light through pixels

