


# NIMBIX

*Cloud for the “Missing Middle”*

*Leo Reiter, Chief Technology Officer*

**NIMBIX**

# The Market



(the Status Quo)



(and others... nirvana)


## The Market



(the Status Quo)



(the “Pragmatist in Pain”..  
aka the “Missing Middle”)



**Sandia  
National  
Laboratories**



(and others... nirvana)

## The Missing Middle is (generally)...

- **Underserved**
  - End user services and ISVs focused on availability and productivity, not performance
  - HPC space focused on massive problem areas, often in public sector



## The Missing Middle is (generally)...

- Underserved
- **Under pressure**
- Tight deadlines
- Increased complexity
- Increased competition



## The Missing Middle is (generally)...

- Underserved
- Under pressure
- **Understaffed**
- Scarcity of skills
  - IT, not HPC focus
- Good talent expensive and hard to find
- “Entitled” workforce
- Shrinking budgets
- “More with less”



## Missing Middle Myths

- SMBs who can't afford nor understand HPC
- Individual “inventors” kicking tires
- Just simulation
- Unsophisticated use cases
- Limited upside potential for vendors/service providers
- DIY/open source zealots who want everything free
- Occasional users

## Who really are these people? (hint: everyone)



Illustration by Chris Gash

- Entrepreneurs
- Contractors
- Project-based SMB
- Startups



- Mid-enterprise
- Departmental
- Project-based teams
- Distributed workforce



- Large enterprise
- Fortune 500
- Research Institutes
- Government
- Defense




## Horizontally speaking...

- **Need HPC but don't have it**
  - Skills gap (commodity IT resources)
  - Budget gap (just bought \$\$\$ workstations!)
  - Software license gap (“HPC” versions different and more \$\$\$)
- **Need HPC and already have some in-house**
  - Always out of capacity, jobs queueing
  - Distributed workforce on different networks
  - Aging infrastructure (slow CPU, lacking coprocessors, etc.)
  - Higher priority projects
- **Need HPC but don't know they need it**
  - New use cases such as Machine Learning, HPDA
  - New algorithms applied to existing platforms

**NIMBIX**


# Transforming Digital Transformation

If:



## Transforming Digital Transformation

If:




...Then:




**NIMBIX**

# Transforming Digital Transformation



## How Missing Middle Wins with Cloud

- ✓ Limitless scale
- ✓ Available “right now”
- ✓ Latest HPC class hardware (some clouds), on demand
  - ✓ Bare-metal, InfiniBand, FPGAs, GPUs, “many core”, etc.
- ✓ Turn-key workflows (SaaS clouds)
  - ✓ Simulation, life sciences, rendering, machine learning, HPDA, etc.
  - ✓ **A must for the “pragmatist in pain!”** 
- ✓ Pay-as-you go, contract, or subscription (some clouds)
  - ✓ Low/no setup cost (no CapEx spend)
  - ✓ Options for every budget

## What about those ISVs?

ANSYS Introduces Industry's Most Comprehensive Pay Per Use Licensing Model for Engineering Simulation Software

-- PITTSBURG, March 31, 2016

"Flexibility is critical for a company like ours," says Sukhvinder Kang, chief technology officer for Aavid Thermalloy.

- Multiple teams globally, simultaneous projects
- Needs always changing
- Difficult to predict scale and spend up front

## Cloud Challenges for HPC

- Performance not always as good as dedicated HPC
- “Sticker shock” possible
  - Unlimited scale can mean unlimited spend
- May still require retaining or acquiring HPC skills
  - Setting up and maintaining HPC workflows is not easy

## HPC Clouds rise to Challenges

- **Performance?** Real HPC Architectures
  - InfiniBand
  - High performance storage
  - Coprocessors and accelerators
  - High density and large memory systems
- **“Sticker shock”?** Subscription and contract pricing plans
  - Optimized for customer use cases
  - “Cell phone” plans offer flexibility *and* predictability
- **“Skills gap”?** SaaS for turn-key workflows
  - Sign up, choose workflow+scale, and go



## HPC Cloud Adoption Trends (Use Cases)

(soon)

Energy

- Data sets challenging
- Existing investment

HPDA, Life sciences

- Data sets challenging

Simulation,  
Media,  
Machine  
Learning

- Mature (simulation/media)
- On-demand/project-based
- “Emerging in cloud age” (machine learning)

(Already there)

# Missing Middle Use Cases

How HPC cloud with software-as-a-service helps

## Simulation for Global Team (Enterprise)

- 100+ engineers doing simulation on electronics components:
  - ANSYS IcePak, Fluent, and HFSS/Maxwell primarily
  - Recently acquired small CFD software ISV (required PaaS to onboard)
- Global footprint – teams in California, India, and China
- Project-based simulations
- Challenges:
  - Performance
  - Scale (including large memory machines)
  - Flexibility on licensing/spend
  - Collaboration across teams and geographies

## Bioinformatics – USDA (Government)


- Whitepaper: “Advances in Sequencing Drive Plan and Animal Genomics to the Cloud”
  - Intel + Nimbix + Pacific Biosciences – featuring USDA
- 56x coverage of long-read PacBio data on the lesser grain borer genome
- Challenges:
  - Scale
  - Time
  - Up-front cost
  - Management complexity of infrastructure

“With most short-read data, you can develop in-house sequencing analysis pipelines using software like DNASTar\* and Blast2GO\*, or you can use public services like iPlant\*. However, for larger data sets such as our 50x PacBio data, we need cloud computing with established pipelines, and a vendor with bioinformatics expertise. We would like to have absolute accuracy, but there isn't enough time, money, or manpower. Cloud gives you an alternative. We have pretty good infrastructure here for the next few years, so the cloud supplements our resources. But if someone is new to this, I would advise them to just go to the cloud. Cloud seems to be the future of bioinformatics.”

– Brenda Oppert,  
Research Molecular Biologist,  
Stored Product Insect and  
Engineering Research Unit,  
Agricultural Research Service Center  
for Grain and Animal Health Research,  
US Department of Agriculture

## Data Science-as-a-Service – HUGEdata (Startup)

- “HUGEdata and Nimbix Offer Cloud-Based Big Data Services” – Campus Technology, Big Data News
- Performs HPDA and machine learning on customer data
- Project-based
- Private cloud today, moving to public
- Challenges:
  - Performance
  - Scale
  - Data/compute disaggregation (cost)




## Real-time Photorealistic Rendering (migenius)

- On-demand photorealistic rendering platform (RealityServer):
  - E-commerce (e.g. Boathouse Sports, etc.)
  - Design (Onshape, Bloom Unit for SketchUp, etc.)
- NVIDIA Iray front-end
- Challenges:
  - Supercomputing GPUs, lots of CUDA cores
  - On-demand, pay-as-you-go
  - Invocation performance



## How Cloud HPC Helps Deep Learning

- Training a DNN requires scale, but doesn't run perpetually
- Increasingly training data is cloud-based – e.g.:
  - Images
  - Sounds
  - Texts
  - Other reference material
- Challenges:
  - Large numbers of GPUs, used part time
  - Ability to re-run training on demand
  - Access to large amounts of training data



## Looking Forward...

- Economics drive traditional HPC use cases to cloud
- Increasing network speeds reduce barriers such as data transfer
- Market demand continues to convince traditional ISVs to offer elastic/cloud licensing
  - Or risk disruption from “born in cloud” technologies
- More turn-key workflows (SaaS) for instantaneous domain-specific value (vs. IT “projects”)



Thank you

Questions?