



# Technical Challenges with Future HPC Systems

**Rajesh Chhabra**  
**Director- Enterprise Computing, Asia Pacific**  
**Altair**

[www.altair.com](http://www.altair.com)

[www.pbsworks.com](http://www.pbsworks.com)

Email: [info@altair.com.cn](mailto:info@altair.com.cn)

---

## Trends/ Challenges



---

## Cloud computing



---

## Green Computing



---

## Altair




# Leading HPC software provider

Automotive	Aerospace	Heavy Equipment	Government / Defense	Life Sciences/ Earth Sciences	Consumer Goods / Oil & Gas
<p>Drive@earth</p>					

## Upcoming HPC Systems

20 Petaflops, 25 Petaflops, 1 Exaflops  
Million cores  
Petabytes of RAM



Chips with 10s/100s of processors  
Many core boards

Sequoia- DoE to be operational by  
2011 IBM Blue Gene

Petabytes of data generation in weeks/days

Google process over 20Pb of data daily (08)

## HPC Challenges- 5 years ago

- Power
- Cooling
- Floor space
- Data management
  - Data movement is bigger challenge than data hosting
- Memory Bandwidth
- Software Scaleup
  
- Complexity – technical management
  - Skill set




The challenges remain as is, looks more challenging!

Software Scaleup – Most Challenging

# Performance

- Semiconductor industry is obsessed with Moore's law


### CPU Transistor Counts 1971-2008 & Moore's Law



## Performance

---

- We kept pushing the clock speed of a single processor till the Heat & Power became a problem.
  - We could get better performance from a chip by having a second processor instead of a very highly clocked processor
  - **Two cores consume less power than a single core running twice as fast!**
- We moved to Multi-core
  - We quickly realized that we can't just simply add more cores without addressing memory bandwidth issue.
  - Just having extra cores were ineffective because there wasn't enough memory for all these extra cores to contribute effectively.



## Power consumption challenge

---

- Sequoia will consume roughly 6 Megawatt
- Exaflop computer theoretically can take Terawatts of power- not going to happen!
  - Even Gigawatt of power consumption by a single system is not going to be acceptable.
- Worldwide Government HPC market will become around \$18 Billion by 2015

## Data Management

---

- We keep producing data
  - LHC (CERN) can produce terabytes of data per hour.
  - We don't manage data effectively
  - Multiple copies of the same data
- Data is
  - Cheap to store
  - Expensive to move
- Thinking outside the box
  - Partial data save
  - Might be cheaper to reproduce than to save it forever

## Software Scaleup

- Applications scaleup
- 1000s of single core legacy applications
- Just beginning to get 64bit versions of legacy applications
- Most of the applications barely scale up to 64-128cores
- Presently MPI is largely the only tool to write massively parallel applications – not easy!



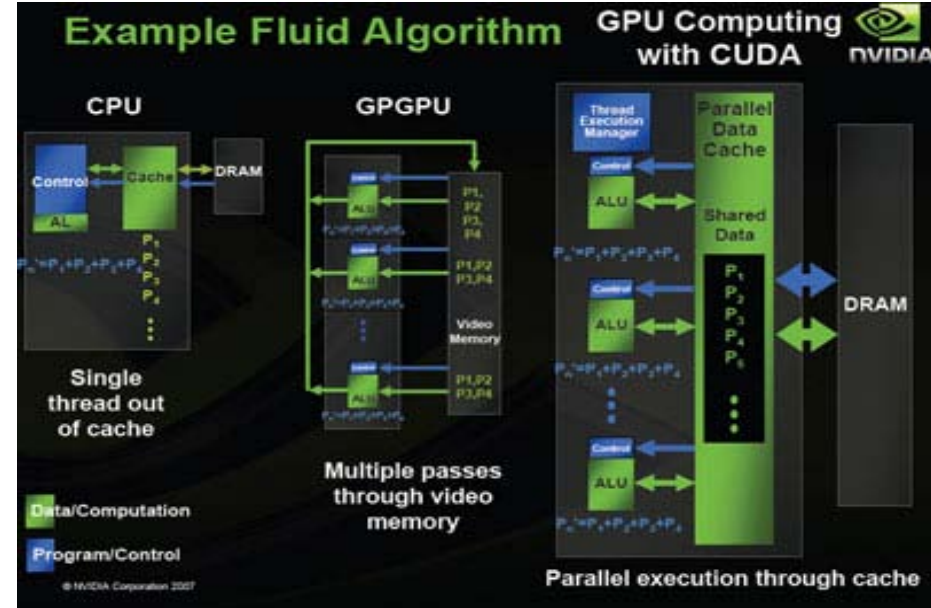
PBS manages 54,000 cores at NASA AMES under a single PBS Server

150,000-300,000 cores system under way


PBS Professional™ working on achieving Million Cores Scalability in the coming years

# GPUs, CPU-GPUs

- GPUs
- 100s of cores
- Targets certain segments of computing
  - Gaming, Animation
- Programmability challenge



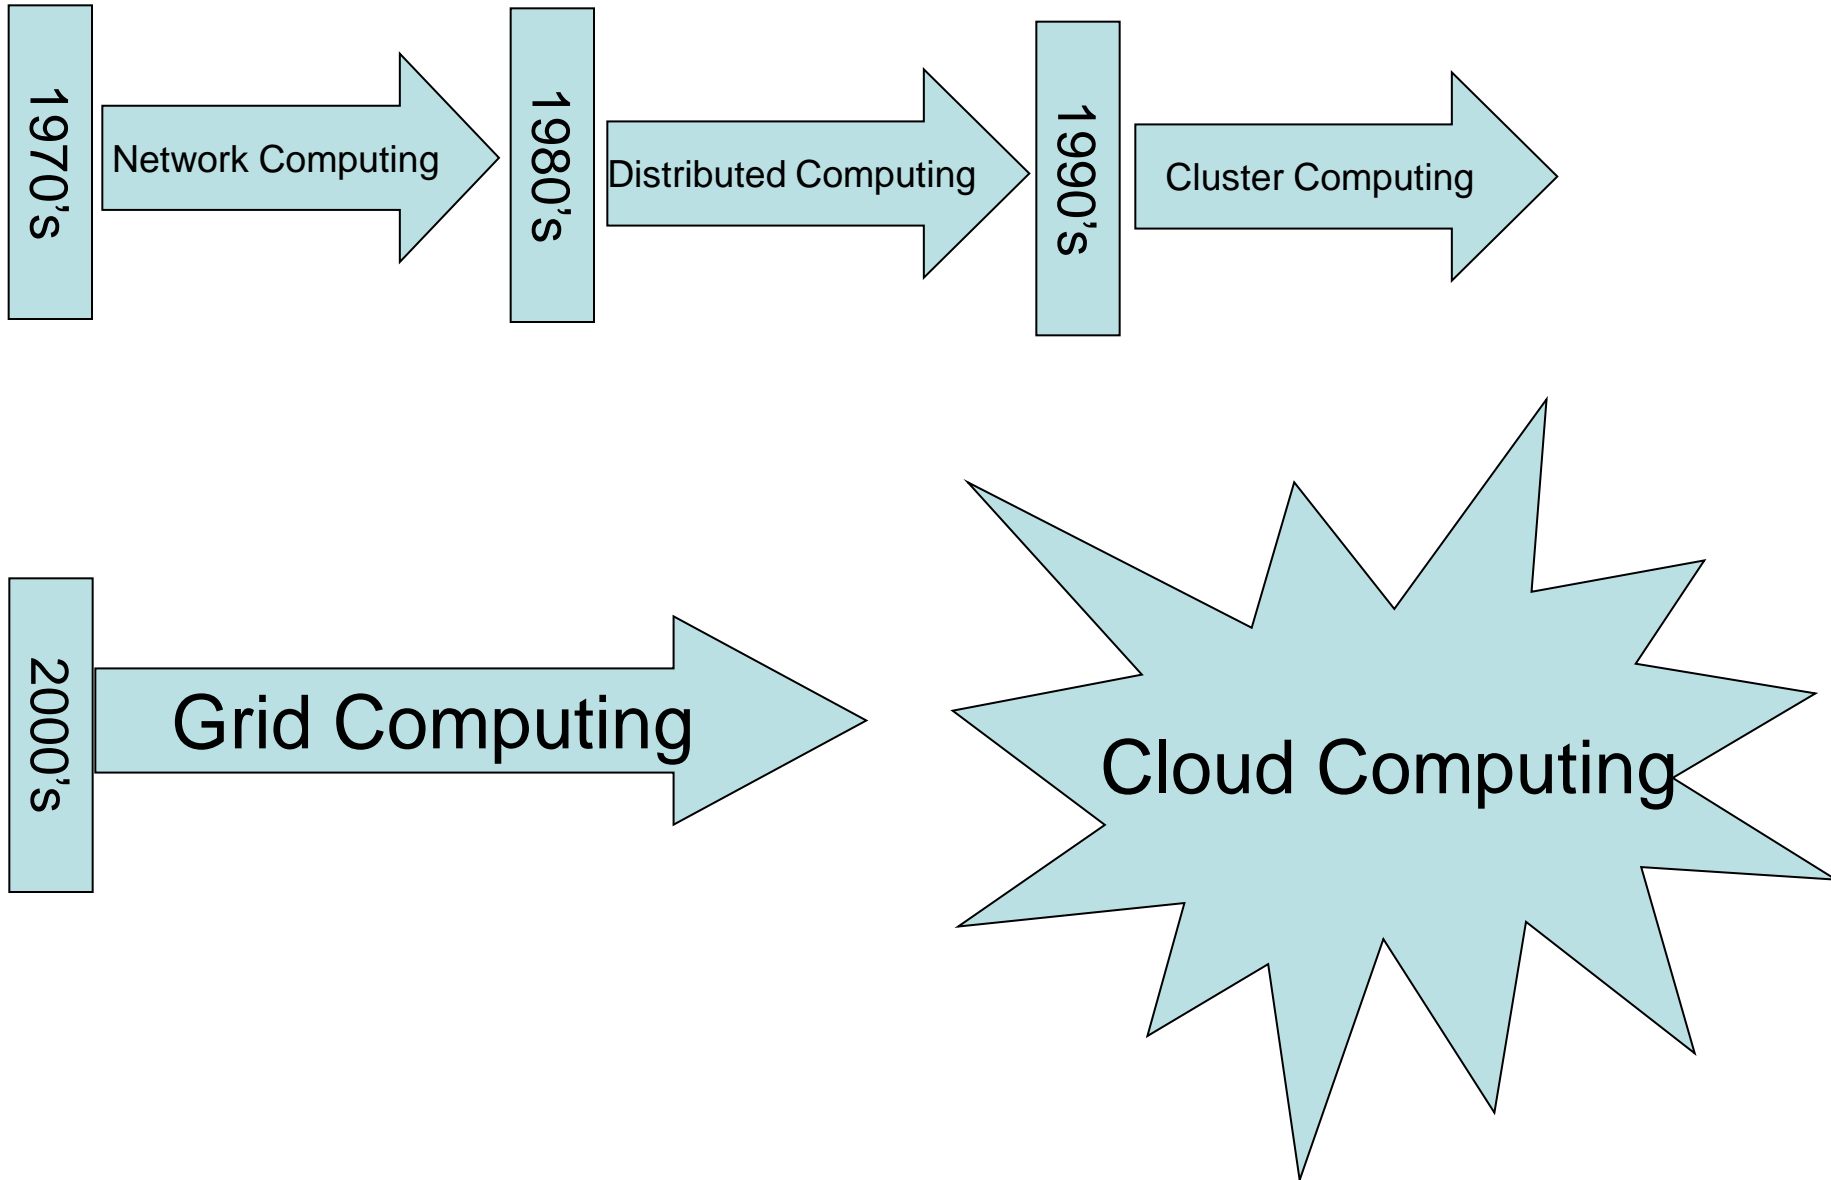
## LARRABEE – TERA-SCALE SOLUTION



- Discrete high end GPU on general purpose platform
- TeraFlops of fully programmable performance
- GPU - >16 cores @ ~2.0GHz, >150W
- JPEG textures, physics acceleration, anti-aliasing, enhanced AI, Ray Tracing etc.

- CPU-GPUs
- More programmable
- Application standardization
  - CUDA vs OpenCL
- Yet to see common applications running on GPUs

## Cloud Computing- Evolution of Network Computing



# Cloud– More statements/definitions then Grid ever had!

**PaaS**(Platform as a Service)  
**SaaS** (Software as a Service)  
**HasS** (Hardware as a Service)  
**IaaS** (Infrastructure as a Service)  
**HPCaaS** (HPC as a Service)


-Pay per use  
-Virtualization  
-Dynamic Scalability  
-Public & Private Cloud



Cloud Computing (CC) is the next big thing!



Cloud is different from Grid



Cloud is just the new name for Grid

Amazon's EC2, S3  
Google Apps  
SFDC  
Azure  
Hadoop



Amazon is the CC leader


SFDC is the CC leader

Google will rule CC space

MS Azure is the next big thing in CC space

# What is Cloud?

- Cloud is an increasingly successful commercial term for an IT infrastructure which essentially combine the concepts of
  - Virtualization &
  - On-Demand
- Cloud comprises of “**Virtualization & On Demand**’ technologies and business models
  - **Virtualization**
    - Virtualization Technologies
      - Underlying hardware is hidden from the users
      - Applications are hosted on the web
      - Service Oriented Architecture allow seamless integration of applications
    - Virtualization Business Model
      - Outsourced services
  - **On Demand**
    - On Demand Technologies
      - Infrastructure (Hardware/Software) & Applications on demand
      - Dynamically scalable (or shrinkable) resources
    - On Demand Business Model
      - Pay per use licensing
      - Lower cost outsourcing model



## Types of Cloud

- Public vs Private cloud – As the name goes private cloud refers to the cloud setup in-house, within an enterprise where as public cloud refers to... well a public cloud 😊
  - Private cloud- enterprise clouds
  - Public cloud – Google, Amazon, IBM, HP etc
- Platform Cloud vs Service Cloud
  - Platform Cloud– providing infrastructure, middleware
    - Amazon Elastic Compute Cloud (EC2) – OS, middleware
    - Simple Storage Service(S3) – Raw Disk space
    - IBM, HP, Apple
    - MS Azure (OS for hosting cloud services)
  - Service Cloud- providing applications/services
    - Microsoft online email/exchange service
    - Google apps
    - Salesforce.com CRM
    - Apple MobileMe -



# Major Challenges of Cloud

- **Virtualization SLA**
  - Clouds need to ensure the hosted infrastructure can be made available to their clients delivering higher level of SLA and ROI.
- **Cloud platform standard**
  - While many companies are offering cloud services but there is no standard today to allow “inter-operability” among the cloud providers
- **Cloud interface** – again non-standard interfaces to interact with cloud services.
- **Security**
  - Many organization worry about the data security and aren't convinced about the security models offer by Cloud providers.
  - Govt organizations are also reluctant to host their data out of their country on Cloud
- **Power/cooling Management**– operational costs of running a cloud
- **Cloud monitoring & Billing**
  - A complex infrastructure used by many users require detailed monitoring and complex billing models
- **Application licensing on Cloud**
  - Most of the services offered by Clouds today are free or open source. Simply because Cloud providers can't afford standard commercial licensing model for hosting applications
- **Skill-set**
  - Lack of skilled people to manage and write applications for Cloud.



Customer Ledger Card

Data Field Information Richard & Donna Burke

Select the Charge or Credit Description

Current Ledger Card

>Sponsor

Initials	Day	Date	Description	Comment	Amount	Balance
BOS	Fri	06/30/2006	Pat By Check	Ck 420	279.00	0.00
BOS	Mon	07/03/2006	Tuition-Wobblers	Danny	115.00	115.00
BOS	Mon	07/03/2006	Tuition-Toddlers	Megan	105.00	220.00
BOS	Mon	07/03/2006	Tuition-Before/After	Chris	90.00	310.00
BOS	Mon	07/03/2006	Family Discount	10%	31.00	279.00
BOS	Wed	07/05/2006	Absent Discount	Megan Sick	18.00	261.00
BOS	Fri	07/07/2006	Pat By Check	Ck 429	261.00	0.00
			Agency Adjustment			0.00
			Misc. Credit			
			Bad Debt			
			Deposit Refund			
			Pat By Check			
			Pat By Cash			

Void Sort Help Exit



## Challenges for an average HPC user/admin

---

- Users
  - Easier ways to do data transfer between PC and backend HPC
  - Better ways to monitor jobs and take actions
  - Performance
    - Speed -Fast cluster interconnect but often a slow router/switch in between the user & HPC system
- Single interface for multiple HPC Systems
  - Institutions with multiple HPC systems
  - Job management among multiple HPC systems
- Advanced Policies
  - Reservations, Backfilling,
- License management
  - Monitoring, management, policies

# Products



**PBS Catalyst**



**e-Compute**



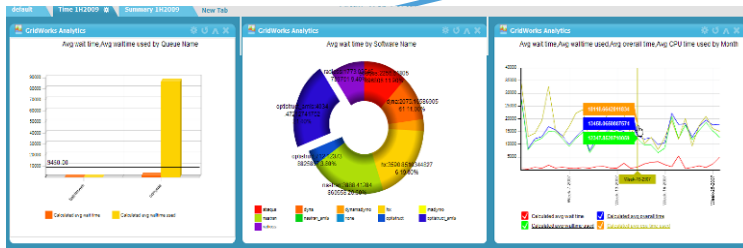
**e-BioChem**




**PBS Professional**



**e-Render**



**PBS Analytics**



**Web Services API**

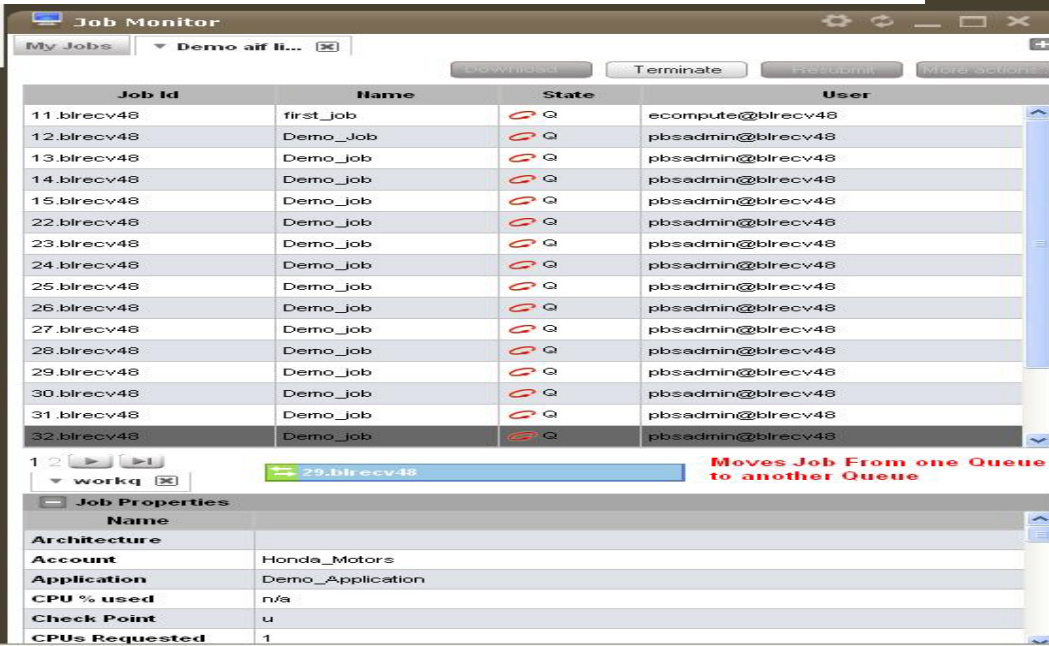


**Green Provisioning**

# PBS Catalyst Web

## • Features

- Job Submission
- Job Monitoring
- File Management
- Application Administration
- Site Administration
- User Preferences

Job ID	Name	State	User
11.blrecv48	first_job	Q	ecompute@blrecv48
12.blrecv48	Demo_job	Q	pbsadmin@blrecv48
13.blrecv48	Demo_job	Q	pbsadmin@blrecv48
14.blrecv48	Demo_job	Q	pbsadmin@blrecv48
15.blrecv48	Demo_job	Q	pbsadmin@blrecv48
22.blrecv48	Demo_job	Q	pbsadmin@blrecv48
23.blrecv48	Demo_job	Q	pbsadmin@blrecv48
24.blrecv48	Demo_job	Q	pbsadmin@blrecv48
25.blrecv48	Demo_job	Q	pbsadmin@blrecv48
26.blrecv48	Demo_job	Q	pbsadmin@blrecv48
27.blrecv48	Demo_job	Q	pbsadmin@blrecv48
28.blrecv48	Demo_job	Q	pbsadmin@blrecv48
29.blrecv48	Demo_job	Q	pbsadmin@blrecv48
30.blrecv48	Demo_job	Q	pbsadmin@blrecv48
31.blrecv48	Demo_job	Q	pbsadmin@blrecv48
32.blrecv48	Demo_job	Q	pbsadmin@blrecv48

Job Properties	
Name	
Architecture	
Account	Honda_Motors
Application	Demo_Application
CPU % used	n/s
Check Point	u
CPUs Requested	1

## • Highlights

- Access to multiple PBS Clusters via simple one tab per server interface
- Simple and Intuitive Graphical User Interface
- Multitasking
- Persistent dashboards (Quickly start from where you stopped)
- Get started with “any” widget.
- Move between queues.
- Reorder jobs in Queue.

## e-Render (Cloud Gateway) used for Malaysia's first 3D Animation movie

- MIMOS sets up a Cloud Computing Infrastructure in Malaysia (2007)
- e-Render delivered to MIMOS
- First Malaysian 3D animation movie in 2009
  - Grossed RM6 million in 1 month
  - Top movie for the first 2 weeks
  - In progress to release the movie in other countries.
  - [www.gengthemovie.com](http://www.gengthemovie.com)




# Manage your HPC/Cloud using Analytics/HiQube



# License utilization by projects

day: 2008 – top 10 projects

cost duration



years

2009	<u>2008</u>	2007	2006
------	-------------	------	------

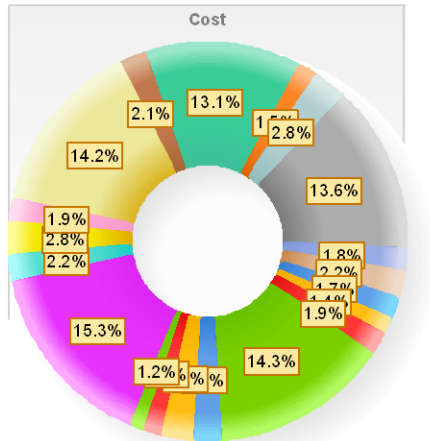
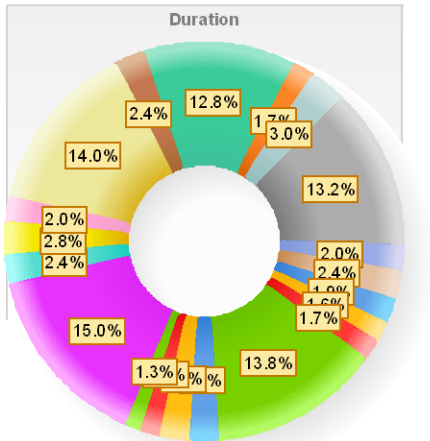
months

				Total
Jan	Apr	Jul	Oct	
Feb	May	Aug	Nov	
Mar	Jun	Sep	Dec	

days

								Total
01	02	03	04	05	06	07	08	
09	10	11	12	13	14	15	16	
17	18	19	20	21	22	23	24	
25	26	27	28	29	30	31		

day: 2008 – G1



- Project 2
- Project 7
- Project 4
- Project 9
- Project 1
- Project 6
- Project 3
- Project 8
- Project 5
- Project 10
- Project 12
- Project 17
- Project 14
- Project 19
- Project 15
- Project 20
- Project 11
- Project 16
- Project 13
- Project 18

filter by group:

user group:

G1

dollars...

delta analysis...

compare...


altair specific...

# PBS Analytics Charts

## Q-What applications are used most, by whom?


A-Abaqus, Dyna, and Optistruct are the top apps, used by Engr

Top 10 Queues for Total number of jobs by Group

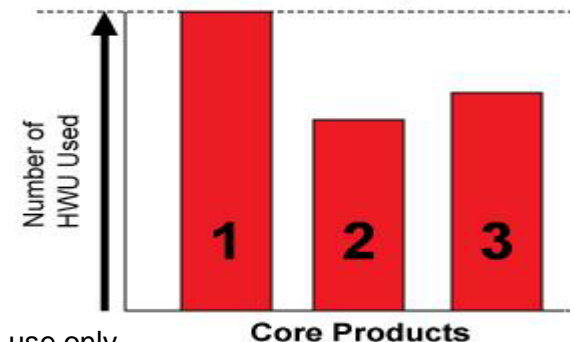


## Unique Licensing

- Floating Licenses
- Tokens Based
- Leveling – Altair's Patented licensing mechanism

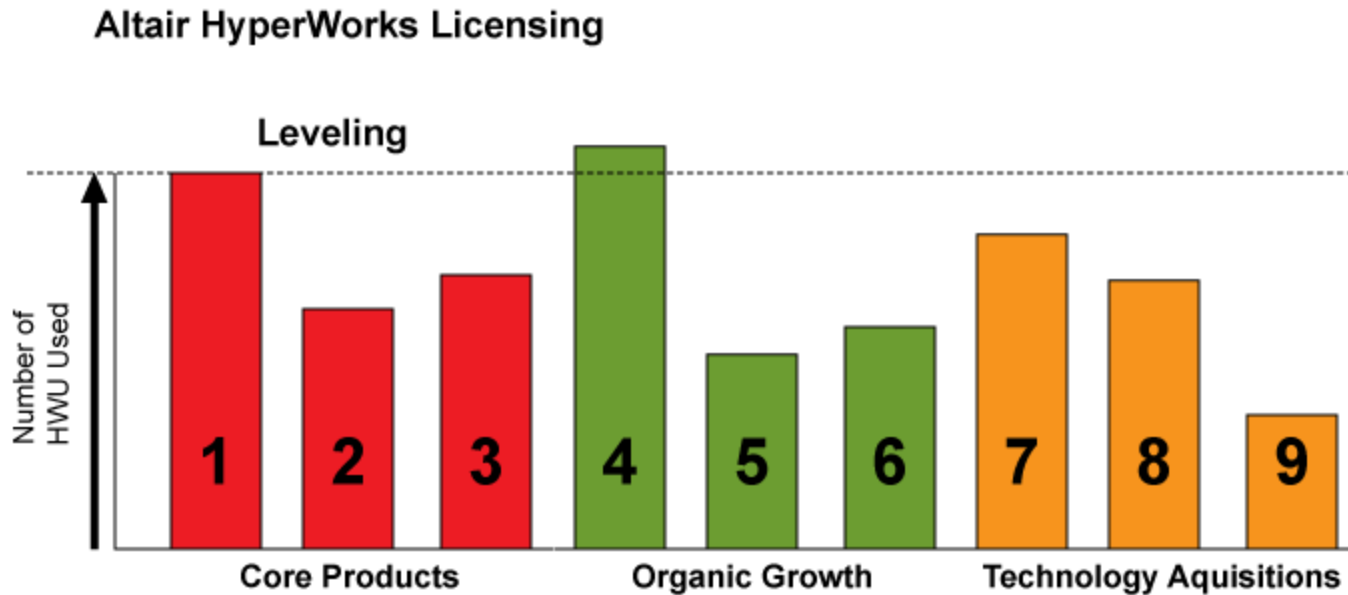


Altair HyperWorks Licensing  
Leveling



# Application licensing

*Unit leveling offers a tremendous value to organizations*



## References

---

- <http://dvice.com/pics/simbbbluegenel.jpg>
- [http://www.laptops-drivers.com/wp-content/uploads/2009/10/nvidia\\_fermi.jpg](http://www.laptops-drivers.com/wp-content/uploads/2009/10/nvidia_fermi.jpg)
- [http://www.scientific-computing.com/images/exhibitors/HPC\\_ISC\\_Finisar\\_web.JPG](http://www.scientific-computing.com/images/exhibitors/HPC_ISC_Finisar_web.JPG)
- [http://en.wikipedia.org/wiki/File:Transistor\\_Count\\_and\\_Moore%27s\\_Law\\_-\\_2008.svg](http://en.wikipedia.org/wiki/File:Transistor_Count_and_Moore%27s_Law_-_2008.svg)