


Virtualization & Grid



Presented by: Bernhard Schott, bschott@platform.com


Date: 27th May 2008



Agenda

- Platform Computing
- Team play: Grid & Virtualization
- Business process speedup
 - Example: “self services” on VMs and complex infrastructures
- Virtualization Technologies & Advanced Use Cases
- Summary

Platform™ Introduction





Platform is a pioneer and the global leader in High Performance Computing infrastructure software, delivering integrated software solutions that enable organizations to improve time-to-results and reduce computing costs.

Over 2200 Customers Worldwide

- Electronics, Financial Services, Manufacturing, Life Sciences, Oil & Gas, Government, Universities & Research, Telco ...

Recognized Leader in HPC, Cluster and Grid Computing

- 15 years global experience
- Worldwide offices, resellers and partners
- 24x7 follow the sun support and services

Growing & Profitable since inception in 1992

- Self-funded
- No debt; money in bank



Office Locations

North America

Toronto (HQ)
San Jose
Washington
Detroit
Los Angeles
Boston
New York

International

China
Japan
Korea
UK
Germany
France

VARs

U.S.
Italy
Israel
Germany
Spain
Korea
Taiwan
Singapore
Japan
U.K.
The Netherlands
Poland
Brunei

South Africa
Sweden
Turkey
India
Malaysia
Thailand
Australia
Austria
France
China
U.A.E.
Portugal



Financial Services

- BNP Paribas
- Citigroup
- Deutsche Bank
- Fortis
- HSBC
- JP Morgan Chase
- Lehman
- Mizuho Financial
- MUFG
- Prudential
- Société Générale
- Sal Oppenheim



Electronics

- AMD
- ARM
- ATI
- Broadcom
- Cadence
- Cisco
- HP
- IBM
- Motorola
- NVIDIA
- Qualcomm
- Samsung
- ST Micro
- Synopsys
- TI
- Toshiba



Industrial Manufacturing

- Airbus
- BMW
- Boeing
- Bombardier
- British Aerospace
- Daimler & Chrysler
- GM
- Lockheed Martin
- Pratt & Whitney
- Toyota
- Volkswagen
- Xi'an Aircraft Design



Life Sciences

- AstraZeneca
- Bristol Myers-Squibb
- Celera
- Dupont
- GSK
- Johnson & Johnson
- Merck
- Novartis
- Novo-Nordisk
- Pfizer
- Wellcome Trust Sanger Institute
- Wyeth



Government & Research

- ASCI
- CERN
- CINECA
- DoD, US
- DoE, US
- ENEA
- ETH
- Fleet Numeric
- GSI
- INFN
- MaxPlanck
- SSC, China
- TACC
- TU Dresden
- Univ Tokyo




Other Business

- Bell Canada
- Cablevision
- Deutsche Telekom
- Ebay
- Starwood Hotels
- Telecom Italia
- Telefonica
- Sprint
- GE
- IRI
- Cadbury Schweppes



Strategic
PartnersPremier
PartnersSelect
Partners

Platform™



Team play:
Grid & Virtualization



Developer



Executive



User



ISV



Executive



I.T. Manager

Platform Enterprise Solution

Platform Accelerate

Platform Symphony DE

Platform LSF APIs

Platform EGO DE

Platform LSF

Platform Symphony

Platform LSF MultiCluster
EnginFrame

Platform Process Manager

Platform VMO

Platform LSF
License Scheduler

Platform Manage

Platform RTM

Platform Manager 5.7

Platform Analytics

Platform EGO



Developer



Executive



User



ISV



Executive



I.T. Manager

Platform Enterprise Solution

Platform Accelerate

Make money
Grow revenue

1. Do more, faster
2. Reduce time to market
3. Simplify use of applications
4. Scalable, grid-enabled apps
5. Improve staff productivity


Platform Manage

Save money
Reduce costs


1. Greater control of infrastructure
2. Reduce TCO
3. Save power & cooling
4. Simplify HPC management

Open & Decoupled Architecture

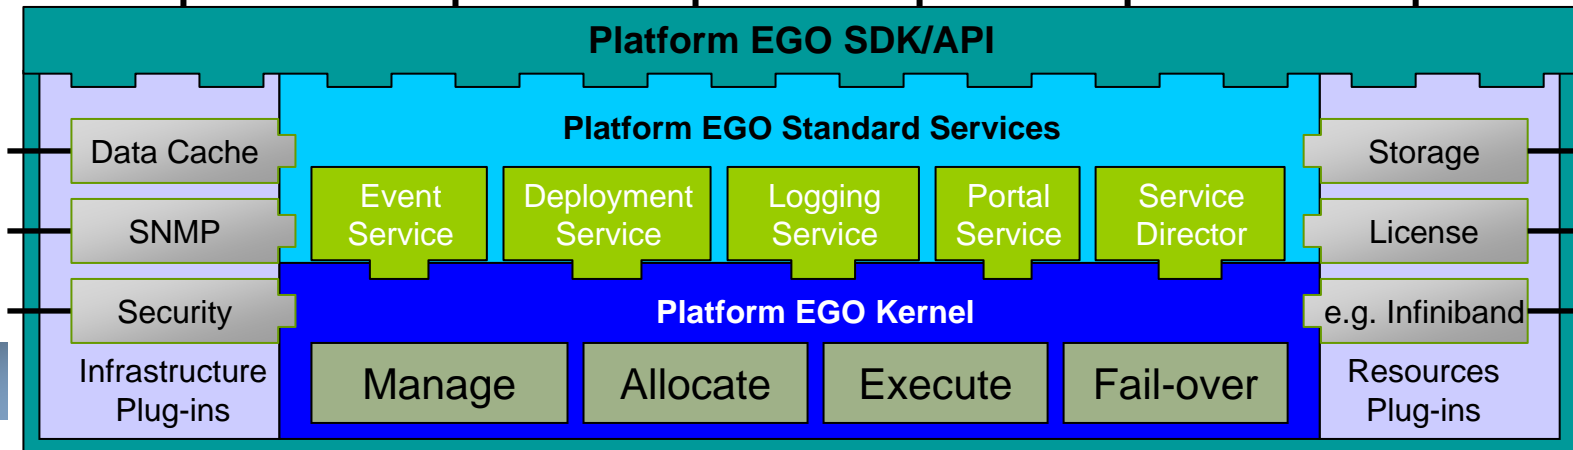
Applications



SOA


Application
Workload
ManagementSystem
Resource
Orchestration

SOI



Grid Devices





- Performance:
 - 10millions jobs per day throughput with >95% job-slot utilization based on EDA job mix, max 5min for failover. (EDA job mix: 5min, 15min, 30min job-runtime)
- Scalability (LSF7.0):
 - n*1000's users, 7500 hosts (up to 10000 hosts dep. on load profile),
 - >100Hz sustained submission rate, (much higher with LSF7.0.3)
 - 0,5 million jobs in single logical LSF cluster at any time
 - 8196's way-parallel jobs with full control down to each single task
- Reliability: self-healing, recovery from incidents, policy driven proactive problem containment, no job loss during operation or in error condition, reconfiguration or failover
- Compatibility: Implemented OGF Standards
 - JSDL / HPC-Basic Profile +extensions / BES
 - DRMAA

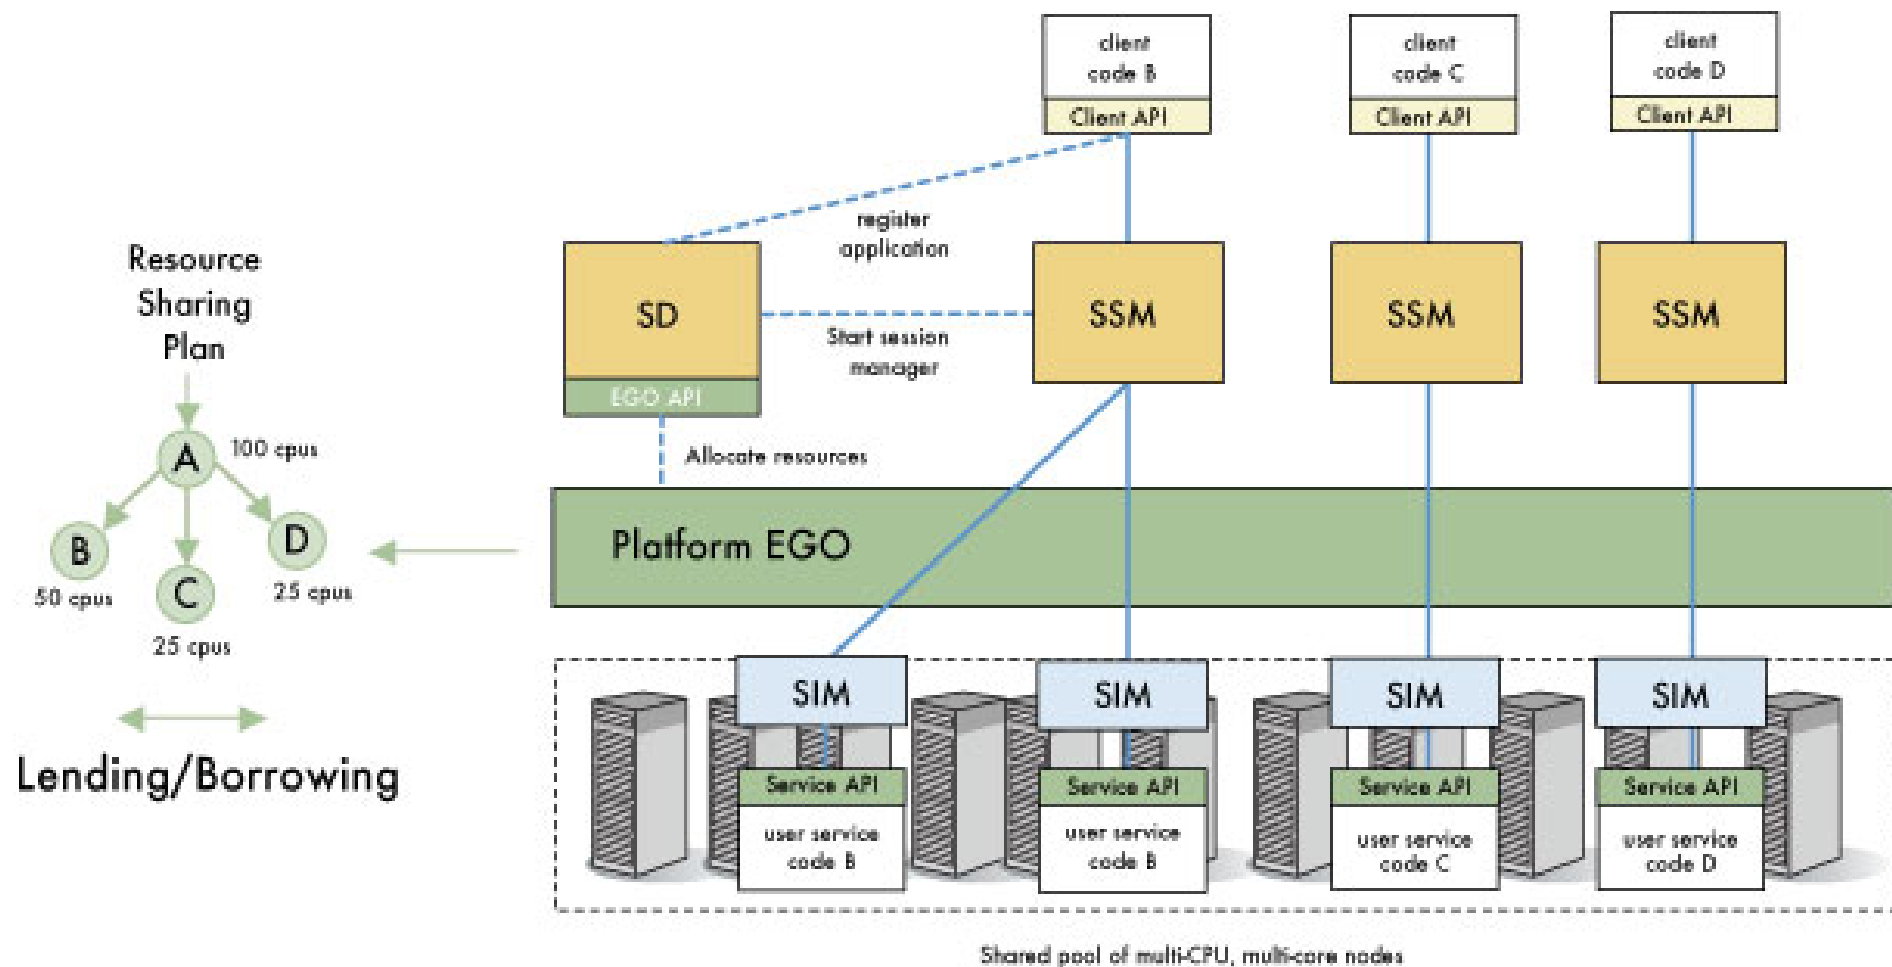
Quality requirements

- Performance and Scalability translates into Reliability
- Reliability can be measured as “MTBF” - Mean Transactions (=Jobs) Between Failure
- Reliability is the prime requirement driving Platform technology

Platform™



Platform Symphony






	Symphony at IBM DCCoD
Scalability 1,000 concurrent clients, 100 applications	20,000+ CPU's simulated on 1,000 physical CPUs in one cluster
CPU Utilization 1-100 clients, 1 sec task, 1KB message, 2,000 CPU	98%
Task Throughput 1KB Message	2,700 messages/sec
Single Task Round Trip	2.4 ms
Single Session Round Trip 100KB common data, 10 second 1KB Task	11.8 ms




Cluster partition usage over the last week

Report period: from 11-19 00:00 to 11-26 00:00

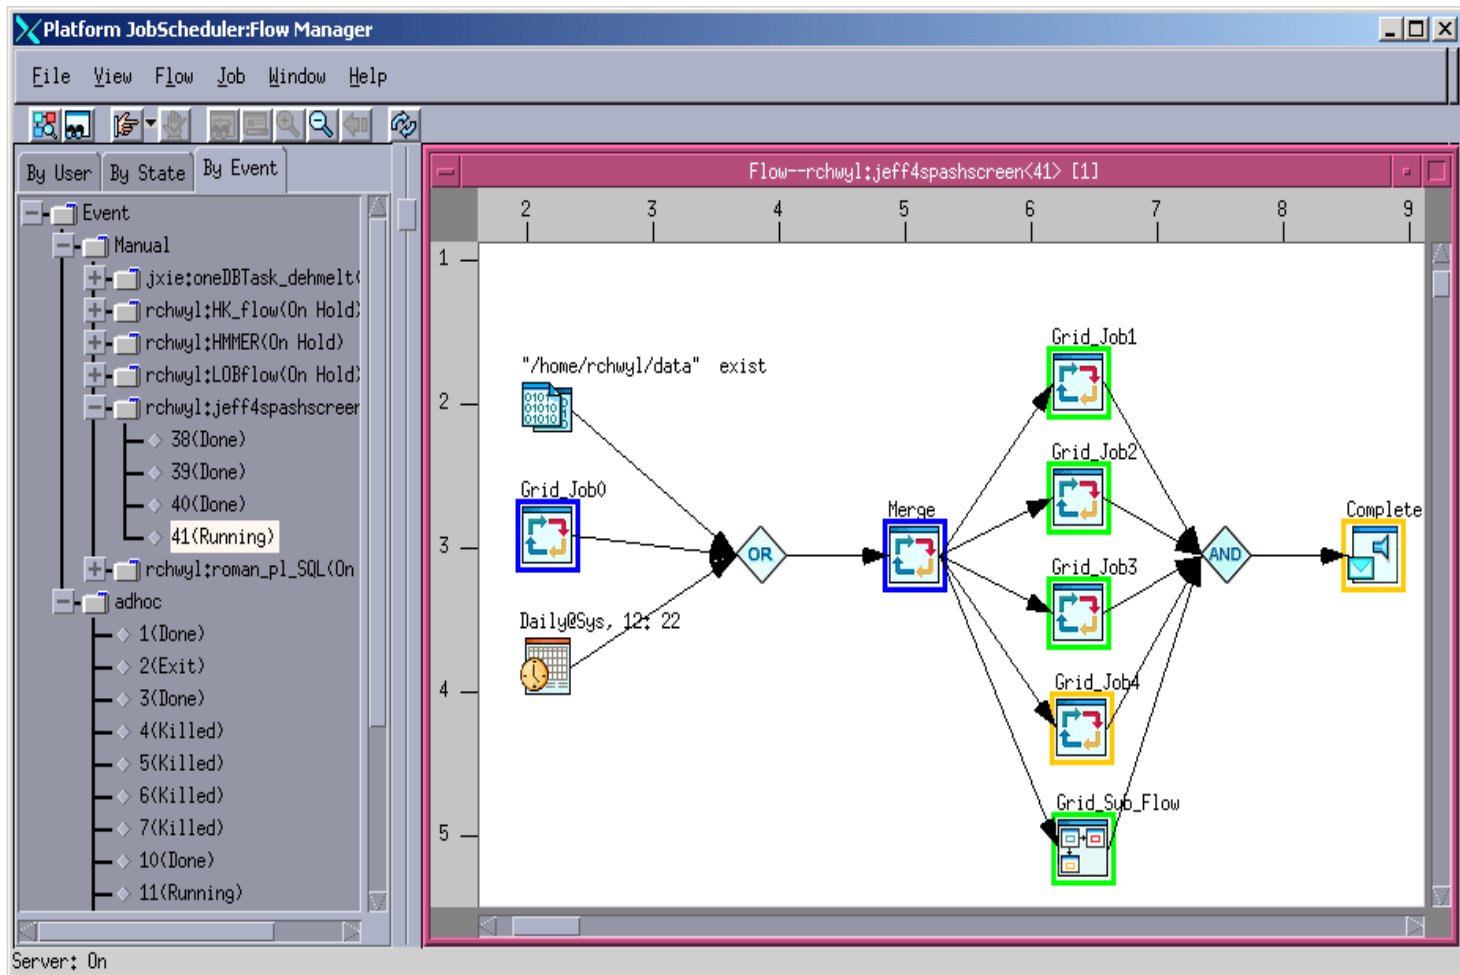


Platform™




Platform Process Manager

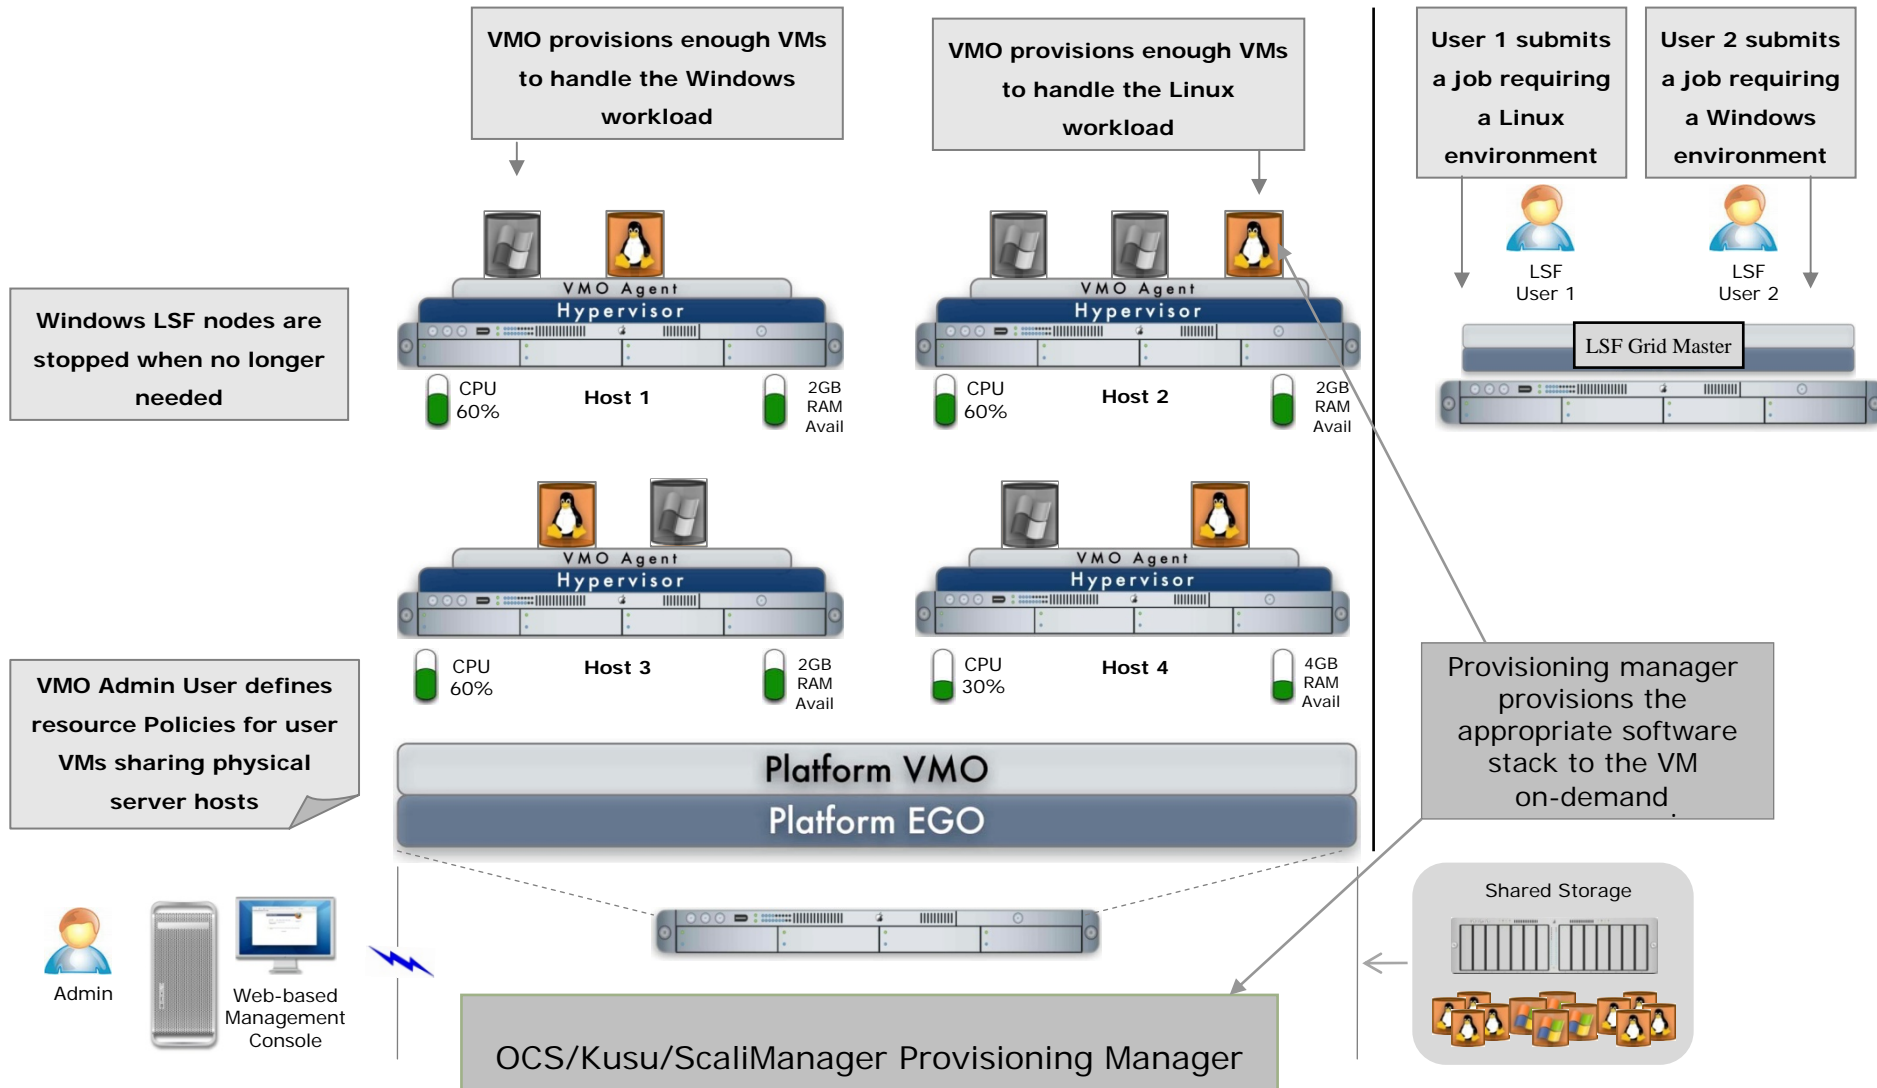
- Workflow Automation

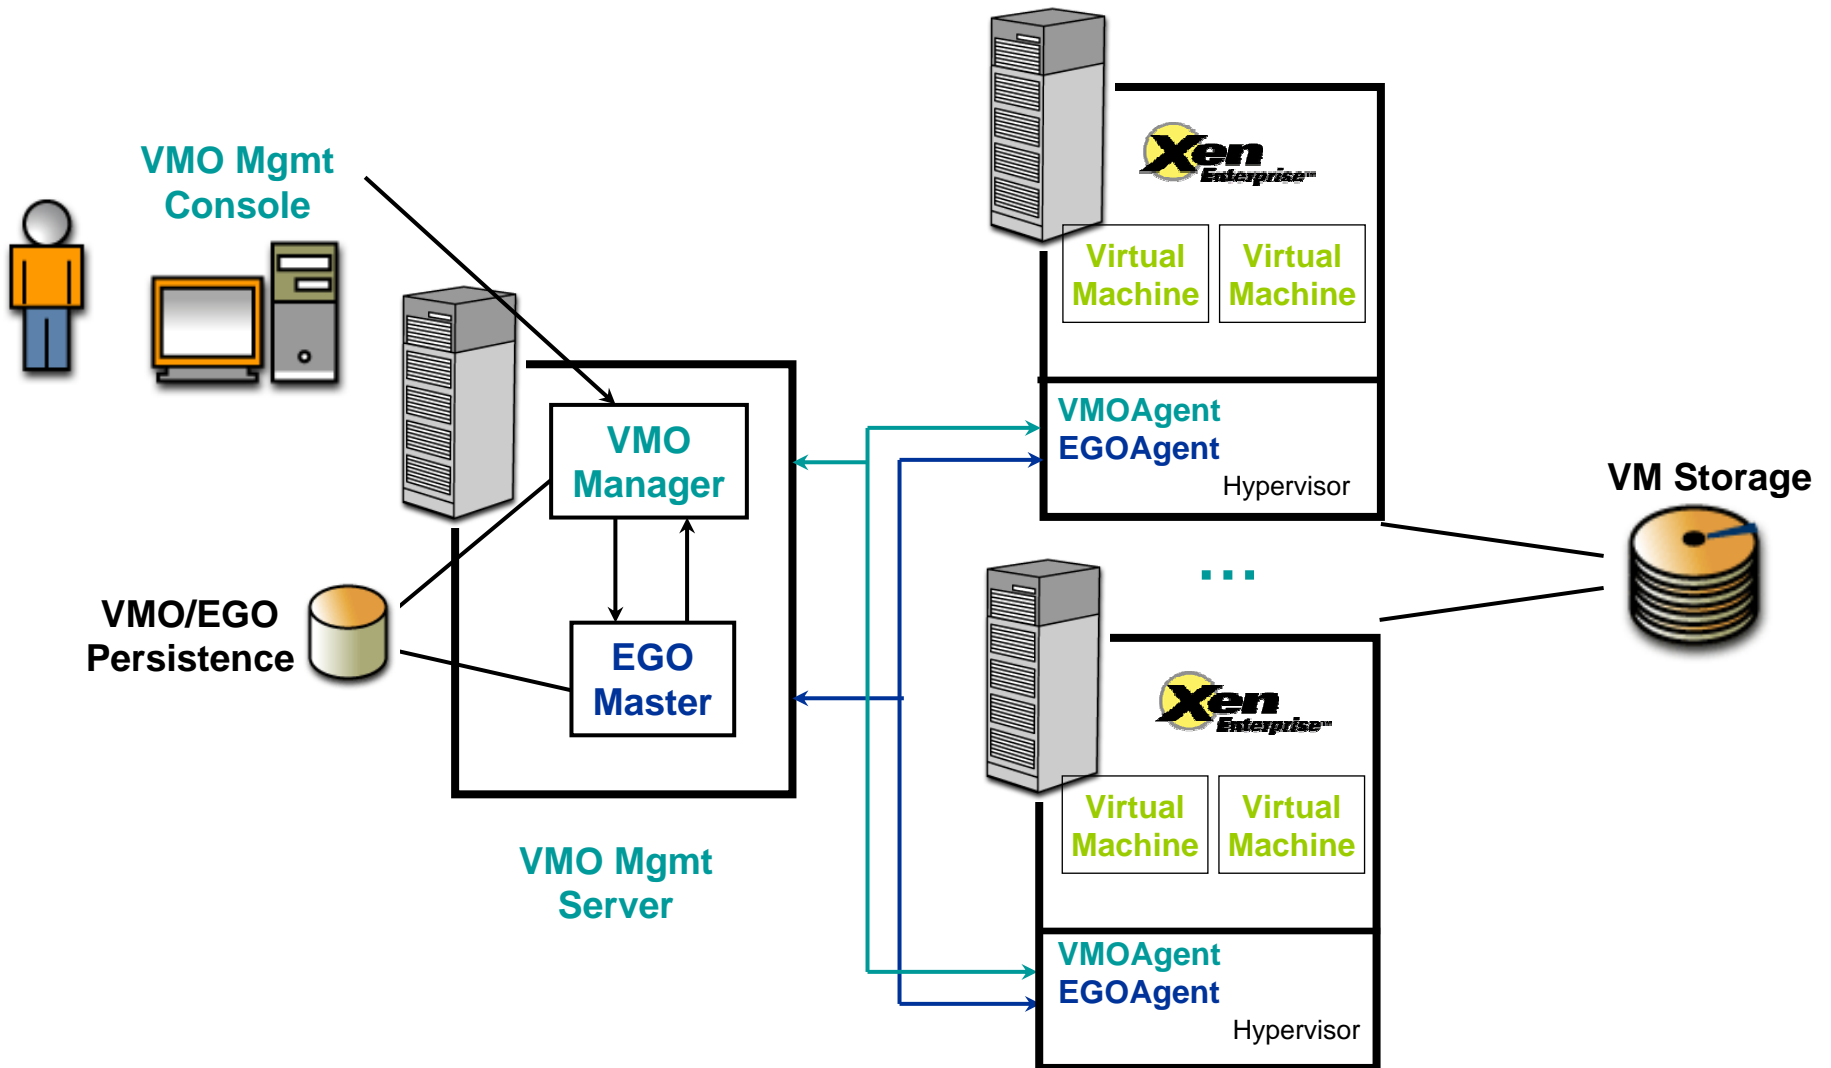


Platform™



Platform VMO






Platform™



Business process
speedup

- Example use case: “versioning”
- Users: chip designers, CAD engineers
- Challenge: customer driven projects
 - Need to use correct application configuration and version of libraries for each project
 - Need for correct application version and patch level – not the latest, but those defined by the customer!
- IT challenge: this morning, give me
 - 3 hosts RHEL4+Patch#47+App-version#7 + component-lib-version#68
 - Afternoon: ... another project ... but keep my morning machines until customer calls back

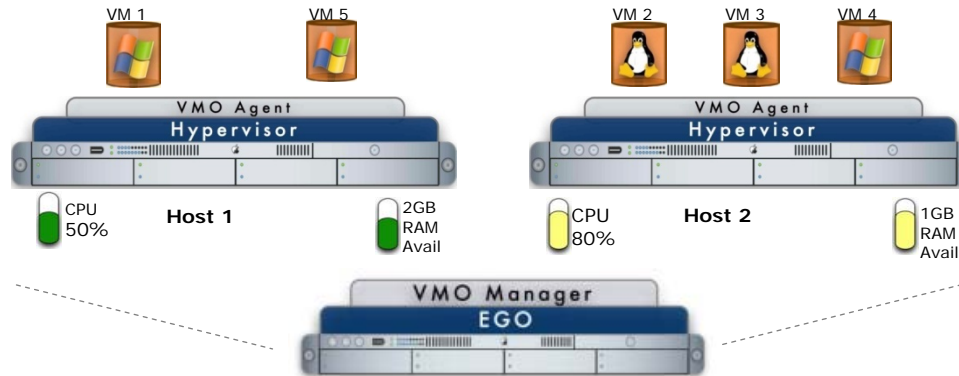
- Example use case: “versioning”
- Solution: “self service”
- Users request VMs with project specific properties by GUI or command line (CLI)
- CLI allows script integration → automation of project specific setups
 - Instead of scrolling through GUI-Listings for correct choices, have an icon on your desktop – 1 click – here you are!
- Tremendous reduction of operational costs
- Better reproducibility – less errors



User logs into VM Management Console

User sees only VMs and VM templates that they have been assigned

User starts up VM with VMO selecting the appropriate host




Templates & VMs are stored on shared storage



Admin creates VMO users

Admin creates VMs and VM templates

Admin assigns user to VMs and/or VM templates

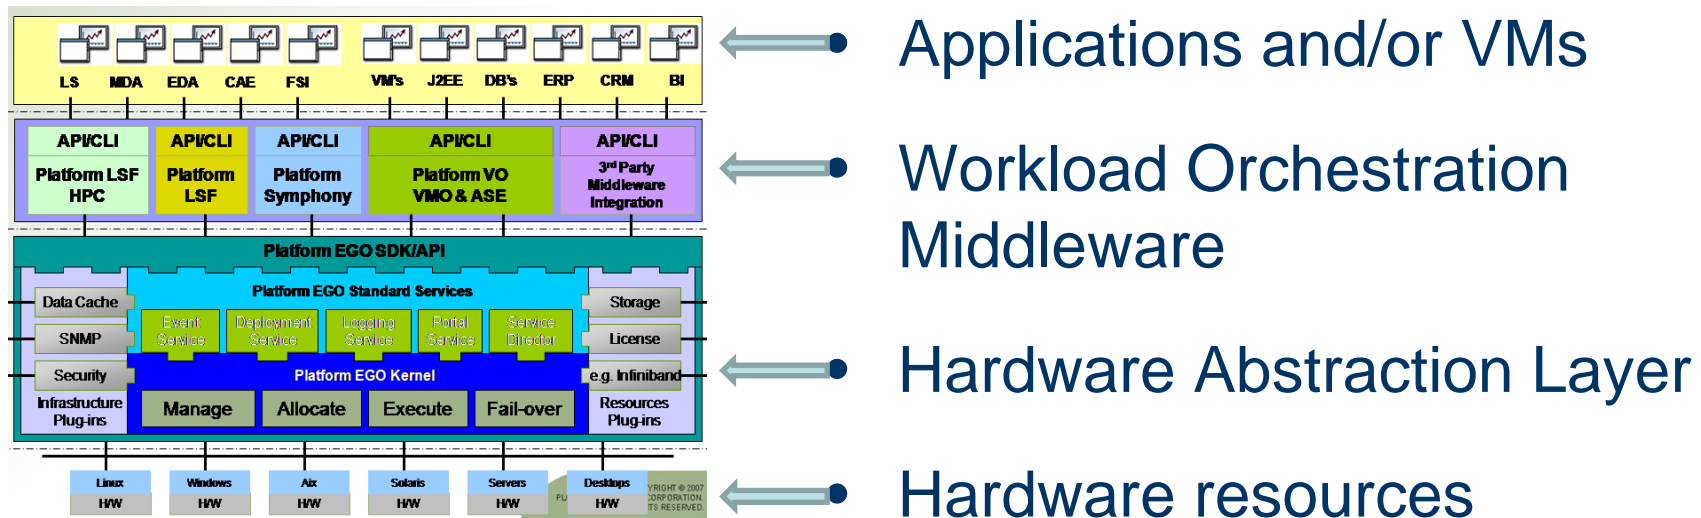


Platform™

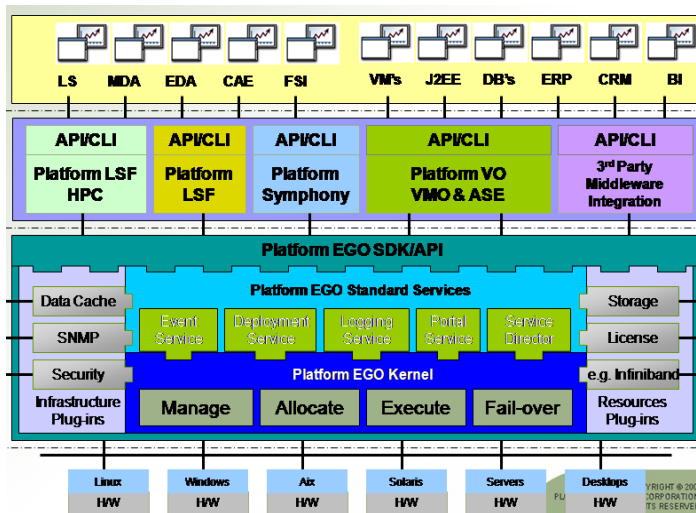


Virtualization
Technologies &
Advanced Use Cases

- Virtualization is taking place on several layers using different technologies:
 - Virtual Machines (VM)
 - Grid technology



- Virtual Machines can
 - be launched by the Grid
 - themselves participate in the Grid
 - become dynamic components in the Grid



VMs launched by the Grid become resources in the same Grid



- Automated deployment of compute nodes
 - dynamic workload driven integration (growth) and release (shrink) of resources
- Per request provisioning of multi-tier systems
- Multi-Tenant infrastructure:
 - separate Grids per consumer group
 - Fully independent configuration
- On demand multi-site extension
 - Transparent workload transfer to remote sites
 - Integration of resource providers like Amazon, IBM, HP



LSF 7.0.2 & VMO 3 are configured and running

VMO RRB maps resource strings in LSF to VMs managed by VMO

Job is submitted

If the job persists in the queue for a defined period of time a specific VM is started

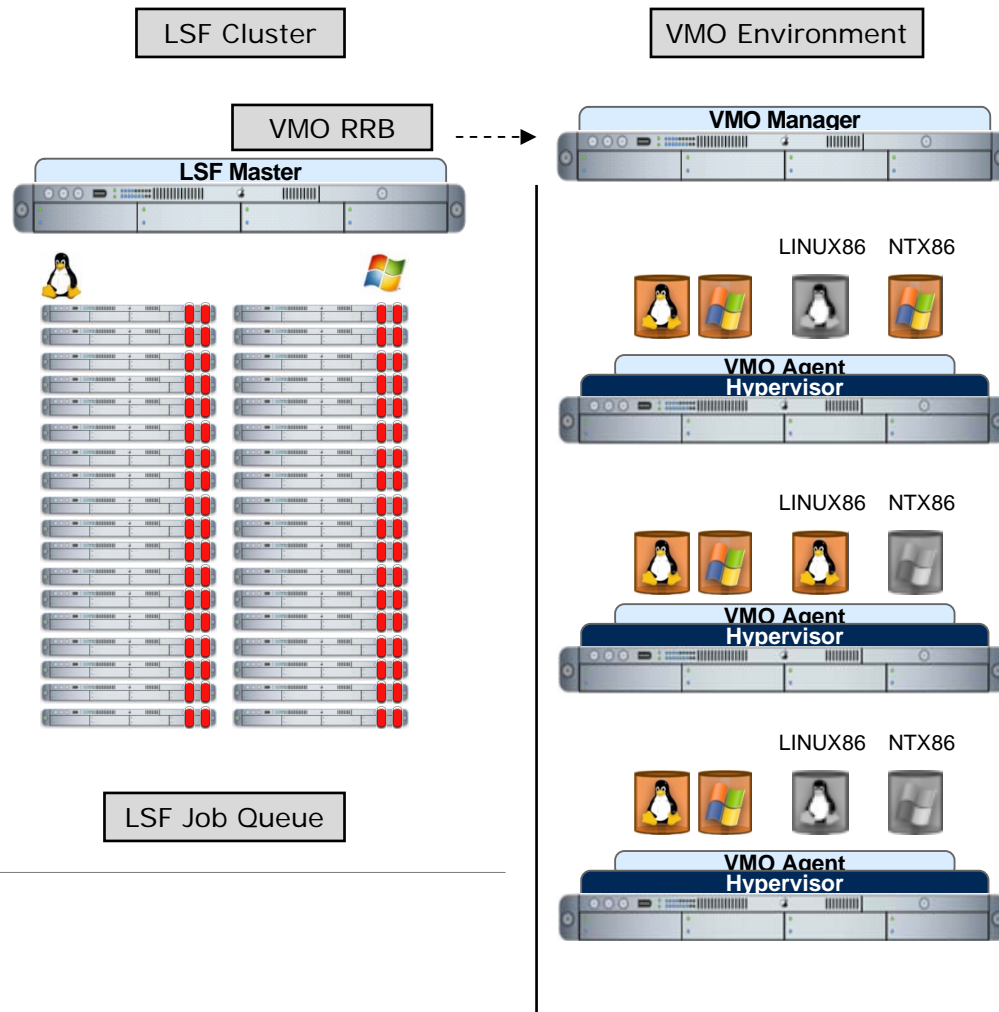
There are now resources available, the job is(are) executed

Another job is submitted

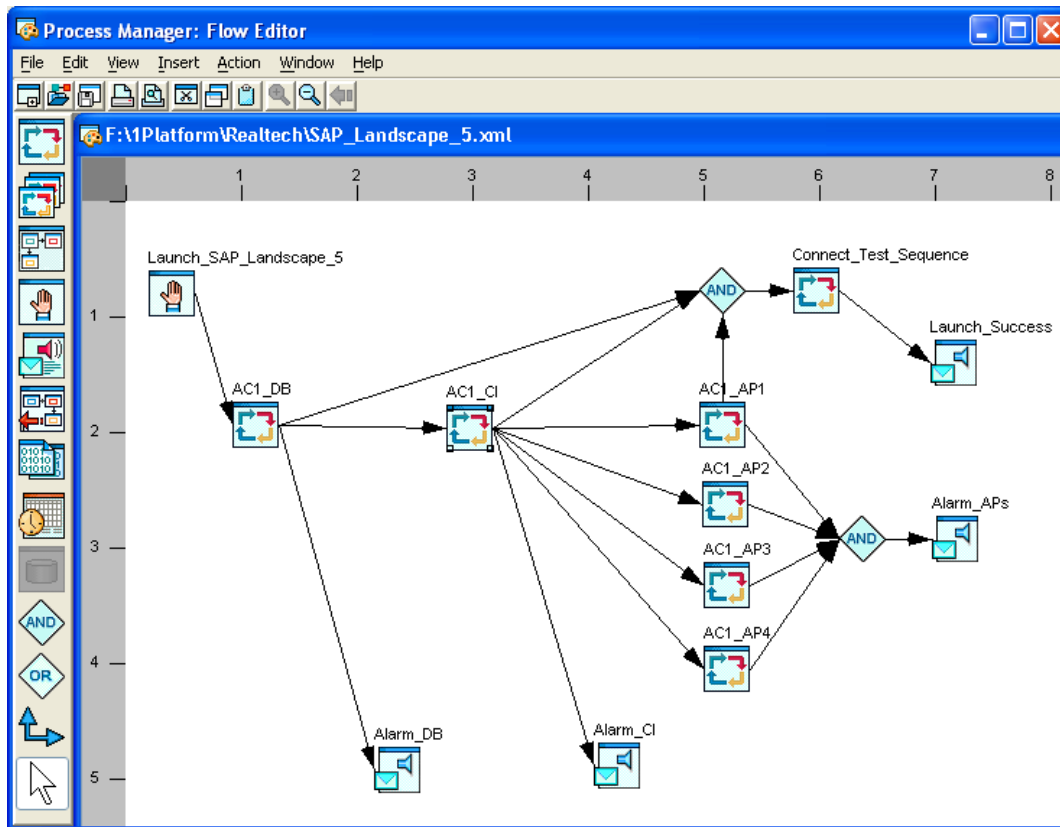
If the job persists in the queue for a defined period of time a different, specific VM is started

There are now resources available, the job is(are) executed

When the VMs remain inactive for a defined period of time they are shutdown

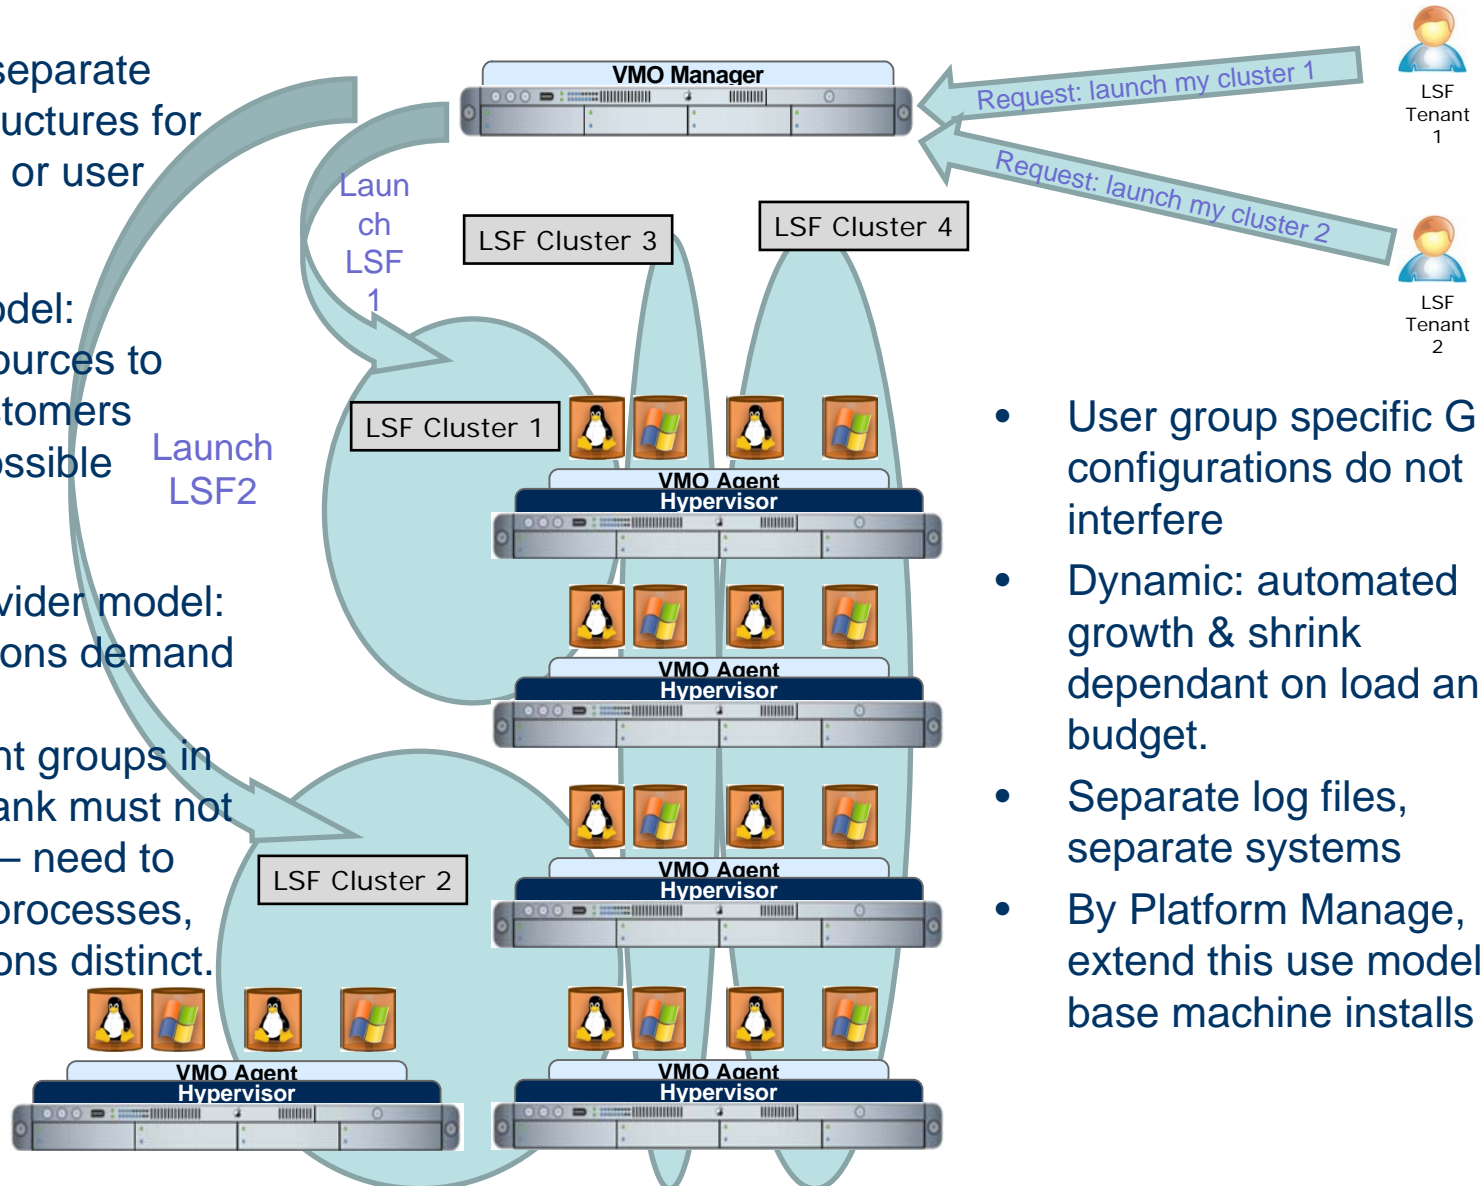


- Provisioning Workflow Automation
 - Multi-tier infrastructure initiation




- Instantiate separate System for Audit
- Temporarily run new (/older) version for compatibility
- Testing of new functionality
- Test upgrade impact
- Training instance / delete after use

- Instantiate separate Grid infrastructures for each tenant or user group
- Provider model: provide resources to multiple customers with best possible separation.
- Internal provider model: e.g. regulations demand certain risk management groups in the same bank must not collaborate – need to keep data, processes, and operations distinct.




- User group specific Grid configurations do not interfere
- Dynamic: automated growth & shrink dependant on load and budget.
- Separate log files, separate systems
- By Platform Manage, extend this use model to base machine installs

- On demand multi-site extension
 - Transparent workload transfer to remote sites
 - Integration of resource providers like Amazon, IBM, HP, or my own company sites



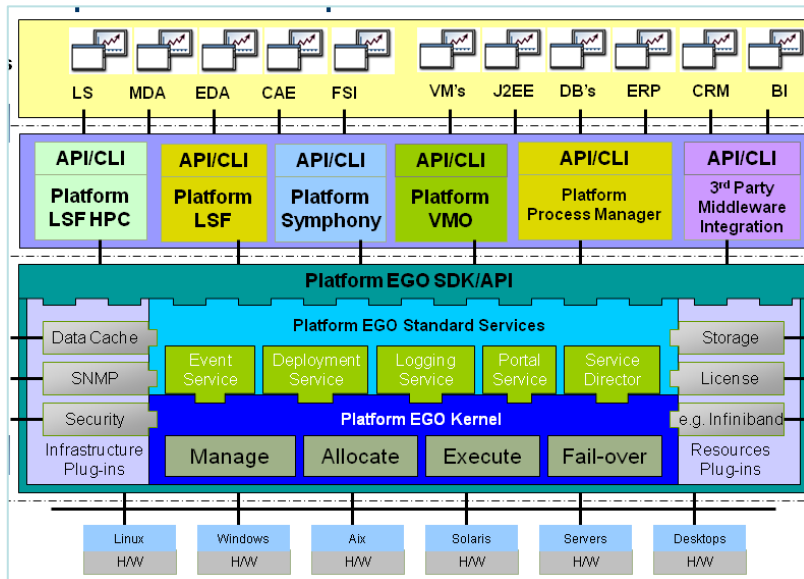
- Virtualized access to geographically distributed resources
- In use since on global scale since 1996

Platform™



Summary

- VM & Grid together enable significant advantages, cost reductions, and business process speedup
- All components and technologies have proven their value and reliability since many years in industry
- Lets reconnect and detail your solution!



Bernhard Schott
Dipl. Phys.
Business Development Program Manager

Platform Computing GmbH

Frankfurt Office

Direct +49 (0) 69 348 123 35

Mobile +49 (0) 171 6915 405

Email: bschott@platform.com

Skype: [bernhard_schott](https://www.skype.com/user/bernhard_schott)

Web: <http://www.platform.com/>



Thank you

www.platform.com