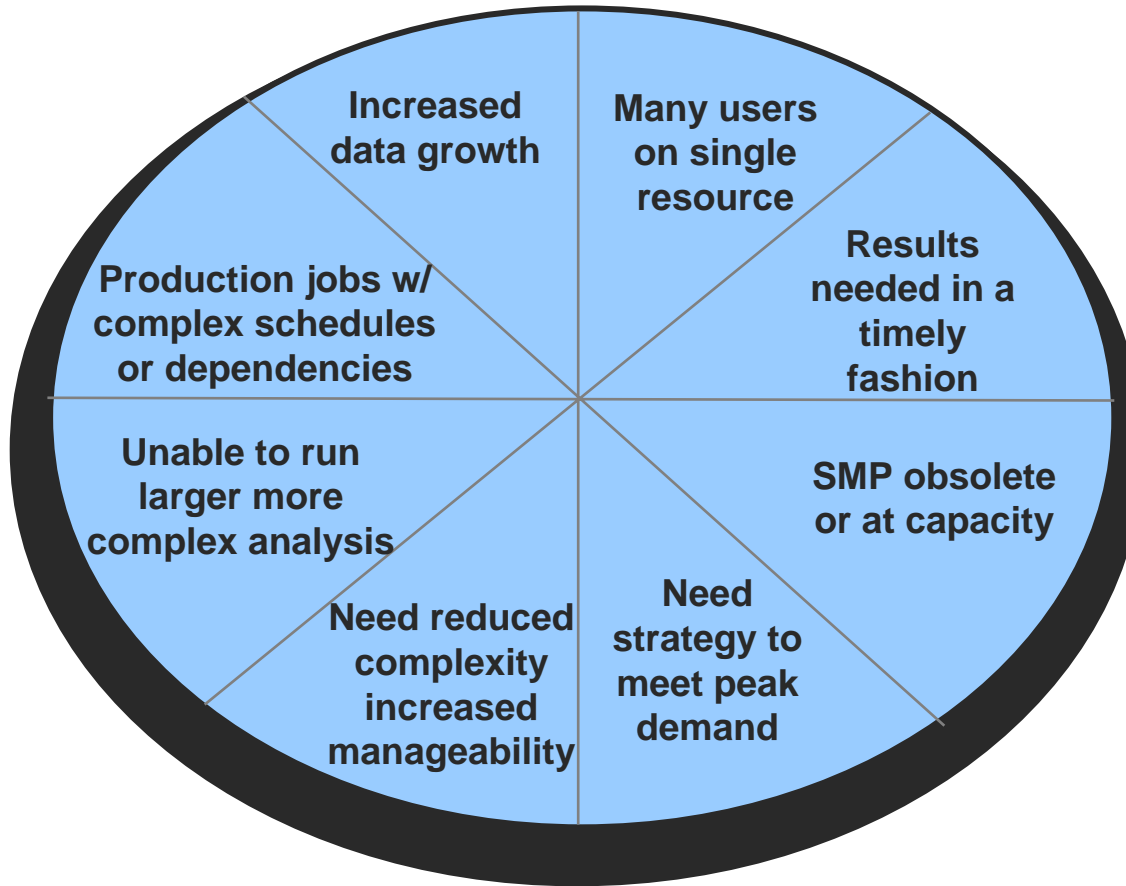


Enterprise: Taking Grid Beyond High Performance Computing (HPC)

**THE
POWER
TO KNOW[®]**

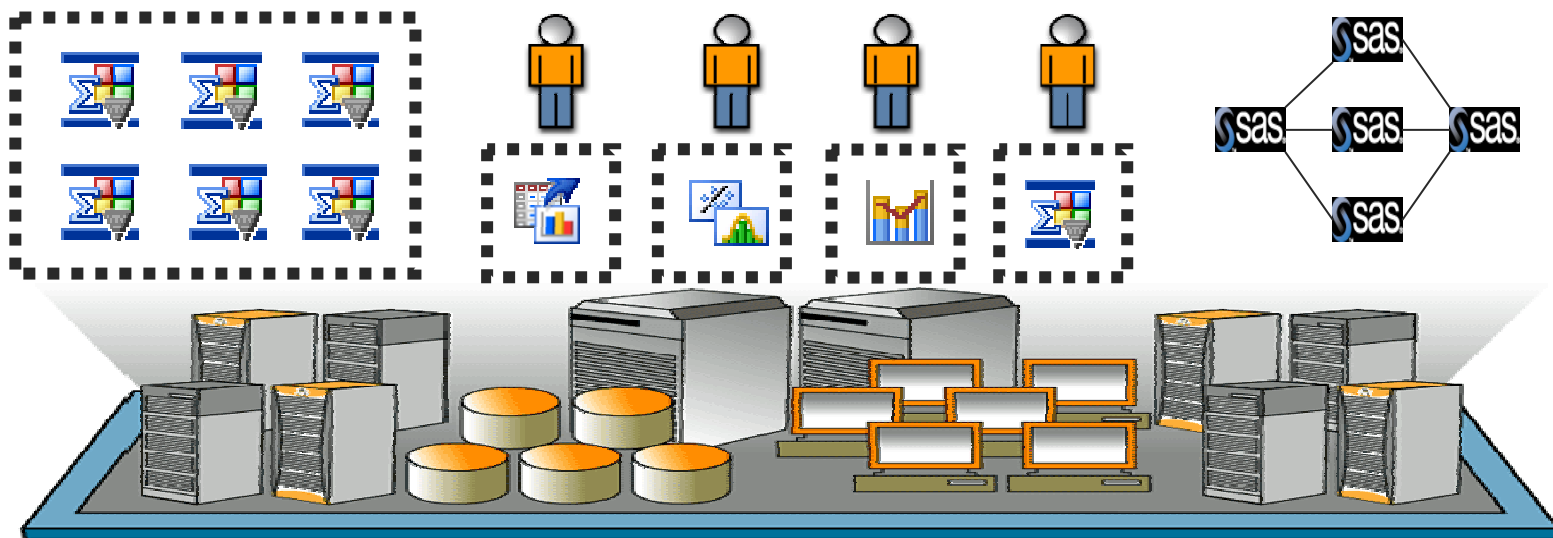
Cheryl Doninger
R&D Director, SAS
OGF Area Director

Enterprise Challenges



Key Grid Capabilities

SAS Grid Manager		
<p>Parallelized Workload Balancing</p>	<p>Workload Balancing</p>	<p>Distributed Enterprise Scheduling</p>
<p>Distribute parallelized workloads to a shared pool of resources.</p>	<p>Distribute workloads to a shared pool of resources.</p>	<p>Distribute jobs within workflows to a shared pool of resources.</p>



Optimize the Efficiency and Utilization of Computing Resources

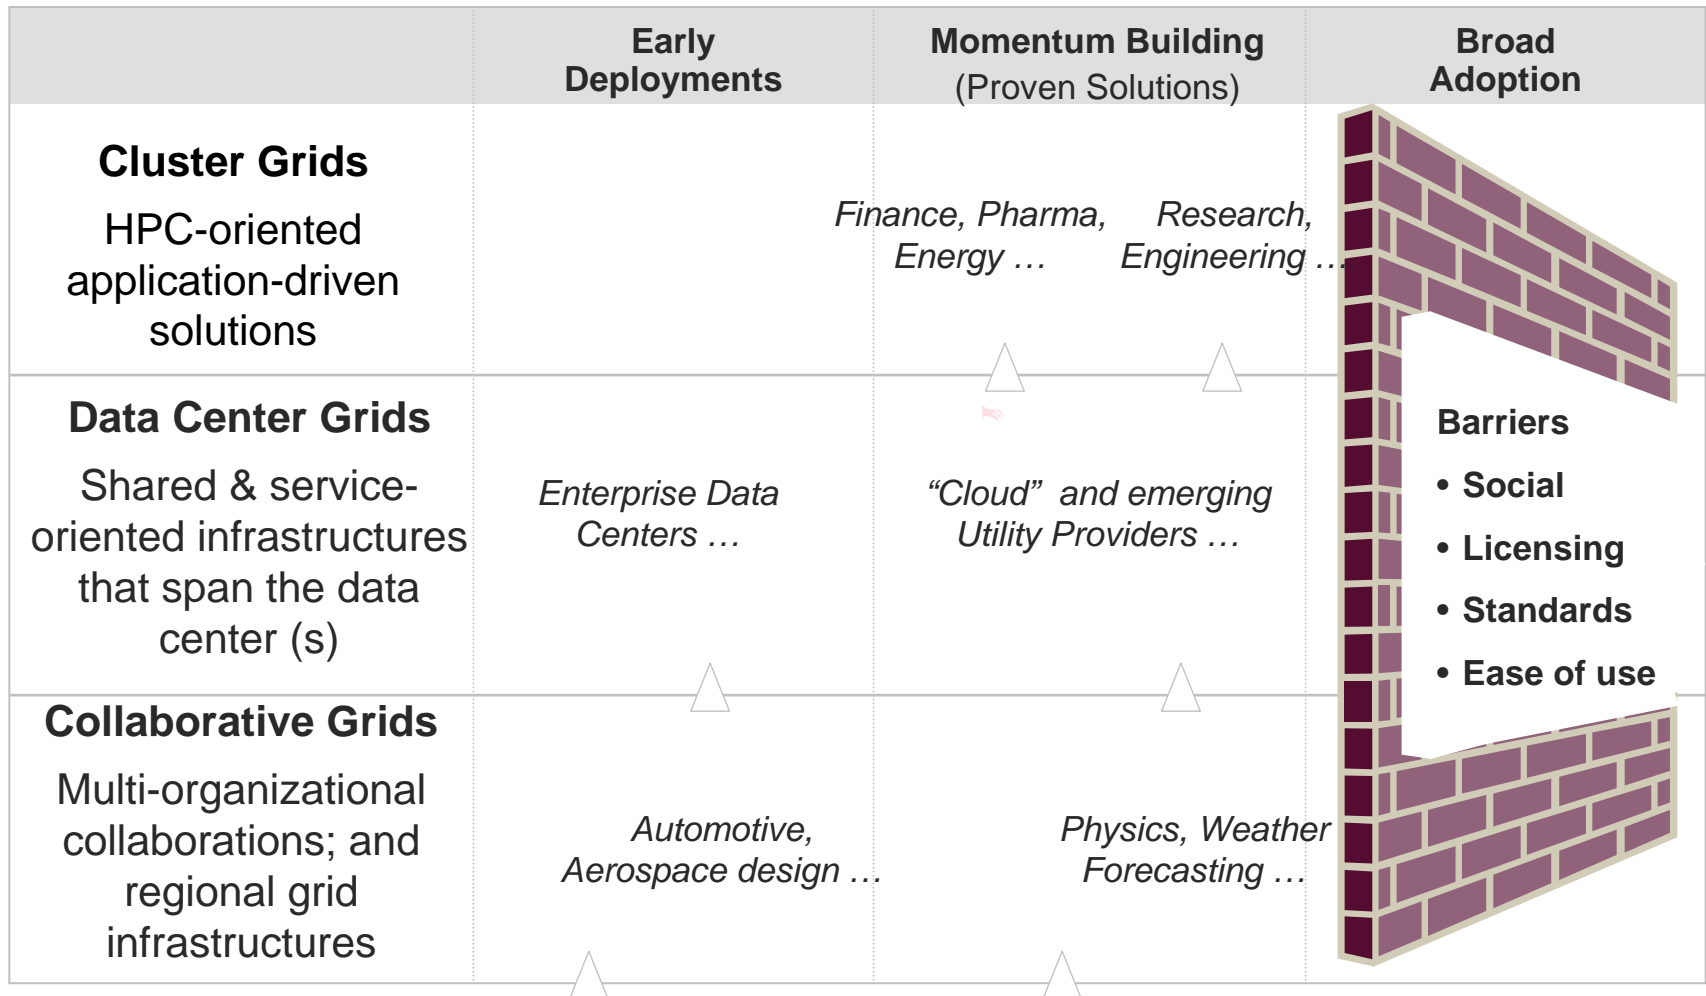
Enterprise Trends

- Industry segments
 - Financial Services, Pharma, Academia, Telco, Energy, Airlines
- Size of environments
 - 8-16 cores for PoC
 - 50, 100, 100+ production
- Operating System
 - Linux, Windows, other flavors of Unix
- Shared file-system is critical
 - Data, data, data
- Primary Goals
 - Replace obsolete hardware
 - Create virtual computing environment with policies to manage multiple applications
 - License flexibility

Grid Aligns With IT

- Provides Enterprise infrastructure
- Consolidation and standardization of hardware
- Enables performance and manageability
- Enforce policies to meet SLA's of entire organization
- Bridge the gap between IT and business users
- Flexibility with management and growth
- Allows ISV's with grid capabilities to align with IT

Popular Grid Usage Patterns Are Emerging ...



This slide courtesy Mark Linesch

Grid Enabling Applications Can Be Expensive

Existing Programs






*Add Grid Computing Program Header
Create Workflow and Schedule It
Automated Tools*

ISV Program with parallel sub-tasks

*Add Grid Computing Program Header
Create Sub-Tasks Using 4GL Syntax
Automated Tools*

ISV Products and Solutions

Automatic Parallel/Grid Code Generation



Business Analytics on Grid

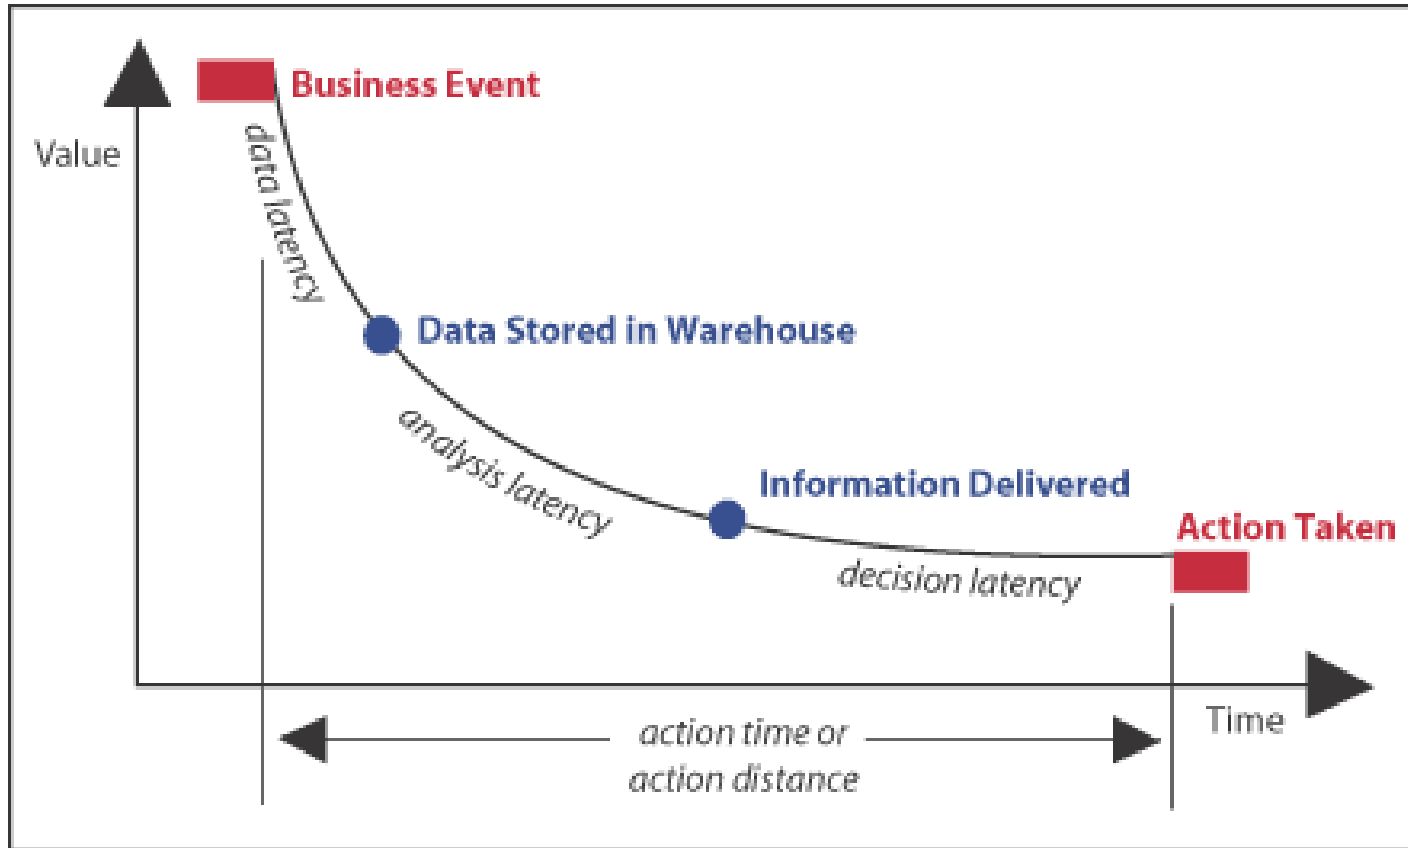
The screenshot displays the SAS Enterprise Miner interface for a project named 'test1'. The main workspace shows a workflow diagram titled 'Home Equity'. The workflow starts with a 'Home Equity' data source, which is processed by 'Data Partition'. From 'Data Partition', the data is split into several parallel paths:

- One path goes to 'StatExplore'.
- Another path goes to 'Cluster [emdlinclus2]', which then leads to 'Drop' and 'Regression'.
- A third path goes to 'Rule Induction [emdlinclus3]', which leads to 'Decisions'.
- A fourth path goes to 'Neural Network [emdlinclus4]', which leads to 'Ensemble'.
- A fifth path goes to 'Principal Components [emdlinclus1]', which leads to 'Ensemble'.
- A sixth path goes to 'Transform Variables', which leads to 'DMNeural'.

All these parallel paths eventually converge into a final 'Model Comparison' node. A red dotted line highlights a path from 'Data Partition' through 'Cluster', 'Drop', 'Regression', 'Rule Induction', 'Decisions', 'Ensemble', and 'DMNeural'. The status bar at the bottom indicates 'Running 4 nodes' and 'Connected to EMDLIN - Logical Workspace Server'.

Parallel Processing Reduces Time to Results

Why Consider Grid Computing?



(Courtesy: Bolder Technology)

How To Increase Enterprise Adoption? Make It Easier!!!

- More grid enabled solutions
 - Standards
- More information
 - Reference architectures
 - Case studies/success stories
 - Best practices
- Pricing/licensing

SAS – Primary Goal...

"When we first started in 1976, we let our customers know that our goal was to make SAS useful for them. Over thirty years, that's still the number one goal: make sure that we develop software that our customers need, and want, and are happy with. And if you have a happy customer, that's one of the great keys to success in business."

Dr. Jim Goodnight

Founder and CEO, SAS