

GSAW2004 Grid and Web Service Standards

**Application Virtualization – Grid and
Web Standards producing results**

David Maples – DataSynapse

408 573 6101 dmaples@datasynapse.com

04-Mar-04

Proprietary and Confidential

© 2004, DataSynapse Inc. ALL RIGHTS RESERVED

Application Virtualization

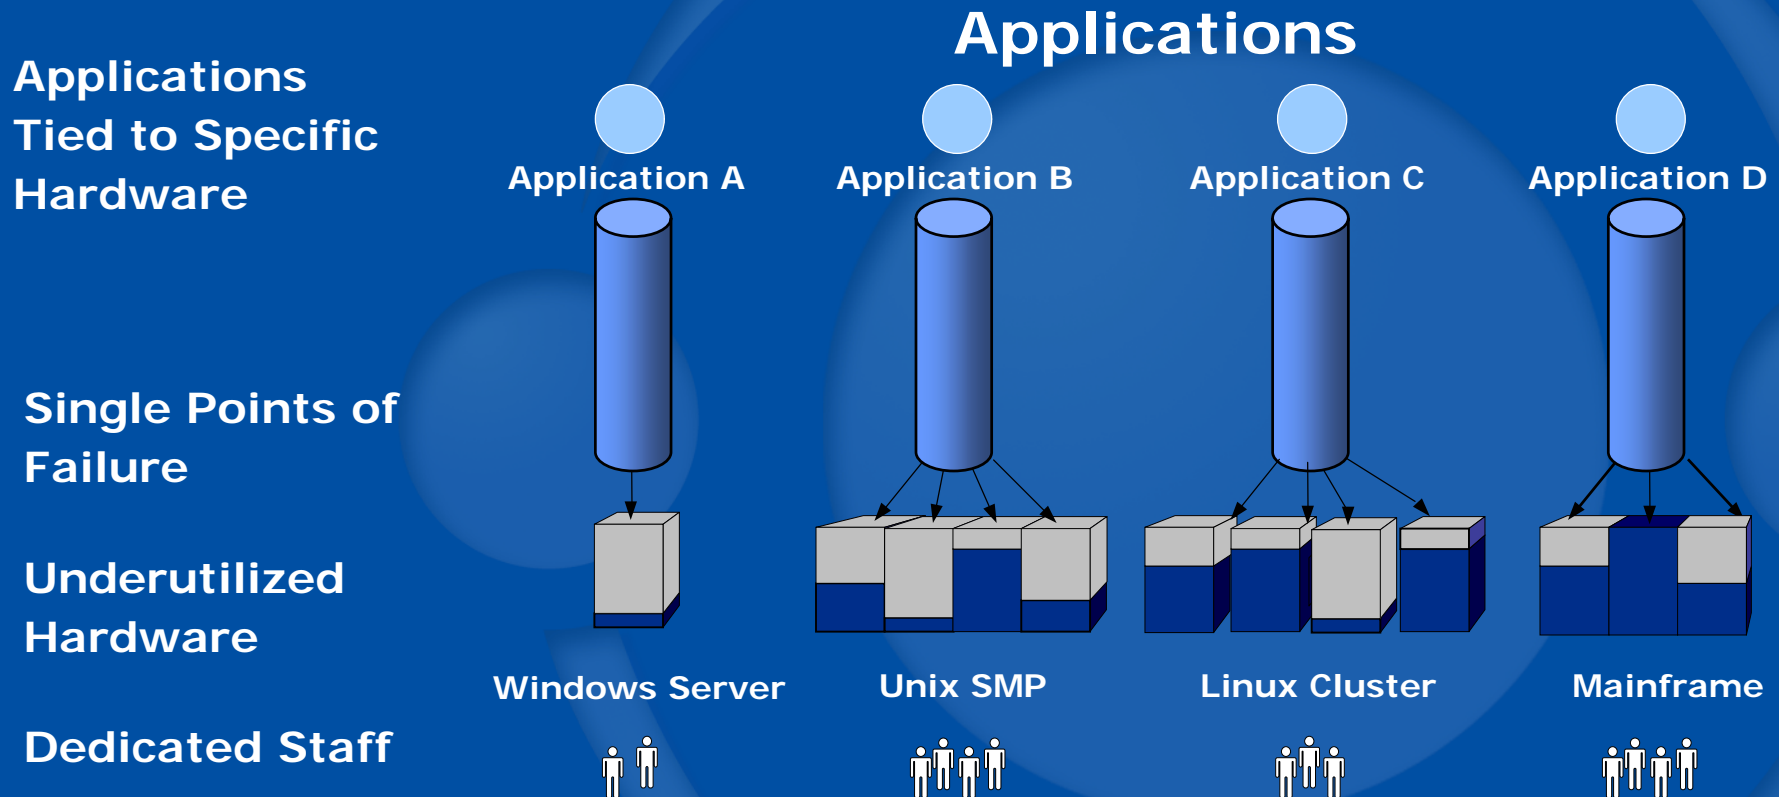
- Commercial Use Cases
- Standards based Architecture
- Reference Architecture with Oracle
- Reference Architecture with IBM

What does GridServer™ do?

GridServer™...

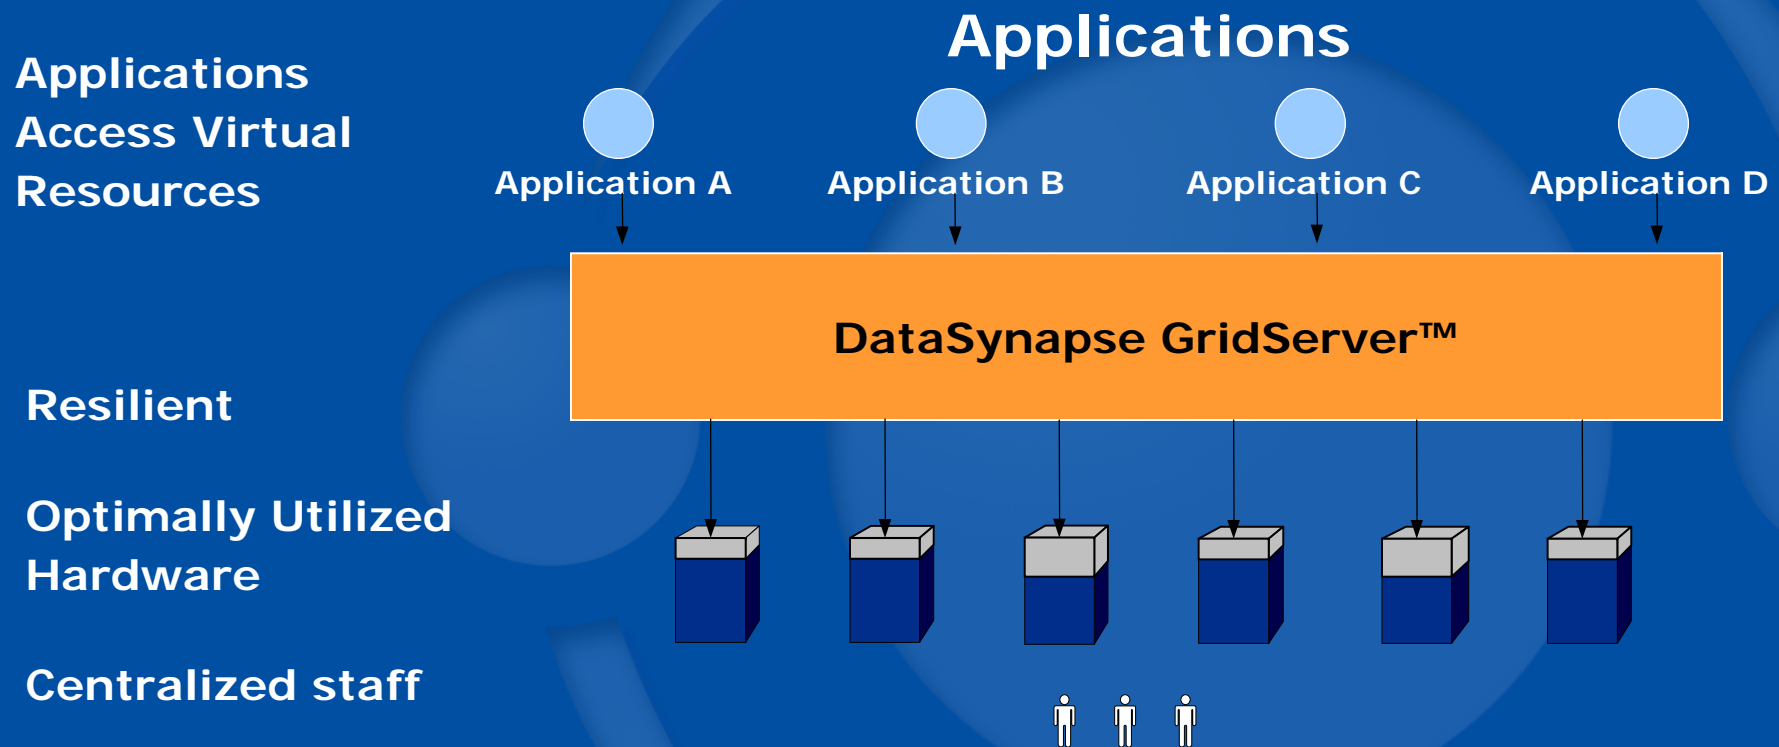
**...reduces the cost and improves performance
through virtualized Service Oriented Architecture
for On Demand computing that complements and
extends existing complex and diverse
infrastructure**

The typical challenge today is dedicated, "Stovepiped" applications tied to underutilized resources



Overall Resource Utilization: 10-30% - No Resource Sharing
IT Staff Focused on Maintenance, Non-Revenue Projects

Application and resource virtualization improves time to results, scale, resilience and time to build and deploy



Resource as optimized resilient utility, Utilization: 90%
IT Staff Focused on revenue-earning business applications

Examples of Virtualization Impact



● GCIB Enterprise Standard

- Real-time risk mgmt
- 14% ↓ IT opex p.a.
- 44% ↓ capex p.a.



● Mainframe/Blades High Volume Transactions

- 5x ↑ Volume capacity
- 50% ↑ SLAs
- 90% ↓ Transaction costs



● Retail, Wealth Mgmt & CIB Enterprise Standard

- \$30-50 MM ↑ Revenue
- 14% ↓ IT budget p.a.
- 30% ↓ capex p.a.

"DataSynapse substantially expands our business capabilities"



Global Grid Forum - www.gridforum.org

- **Formed in 2001 by a Merger of Grid Organizations**
 - European eGrid
 - US Grid Forum
 - Asia Pacific Grid Community
- **Modeled After IETF and IRTF**
 - Meets Three Time Per Year
 - Areas, Working Group and Research Groups
 - Consensus Based
 - Open Membership, Most Work Done on Mailing Lists
 - Creates and Documents the Standards
- **Strong Industry Support**
 - IBM, HP, Sun, Microsoft, SGI, US Govt., many more



DataSynapse is Sponsoring Open Market Grid Standards through the GDC Council

- **Co-sponsors:** DataSynapse & Intel
- **Focus:** Financial Services, non-profit organization
- **Mission:** educating and examining distributed computing standards, emerging and current technologies, and best of breed business practices to accelerate standards



AIG

BANK ONE

BEAR
STEARNS

LEHMAN BROTHERS

Charles Schwab

CREDIT
SUISSE FIRST
BOSTON

Freddie
Mac

UBS

Bank of America

Merrill Lynch

Morgan Stanley

WACHOVIA

Goldman
Sachs

CIBC

JPMorgan

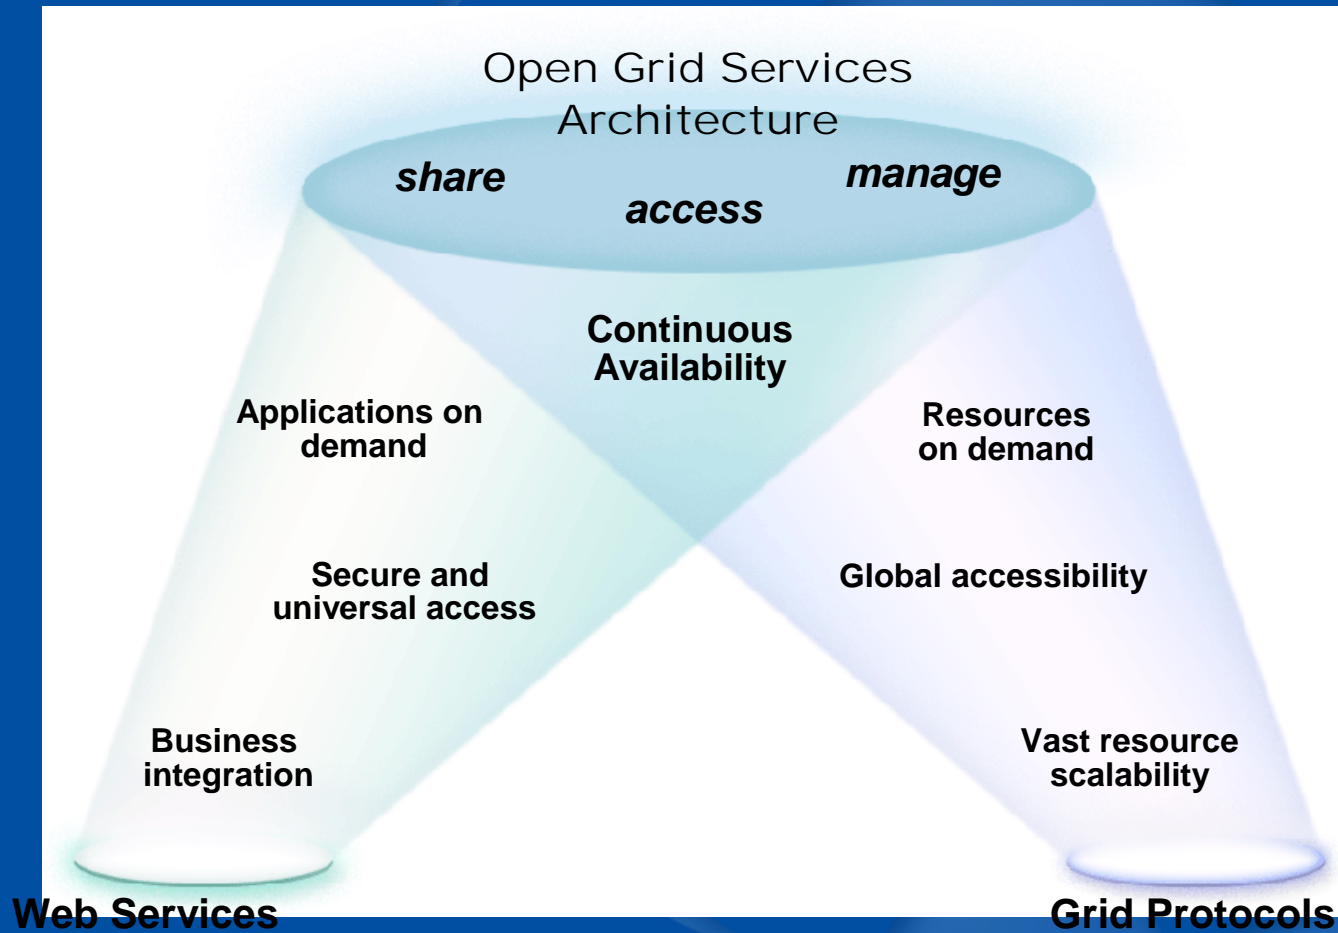
Deutsche Bank



Fidelity Investments

OGSA – The Best of Both Worlds

- Web Services and Grid Protocols



GridServer™ addresses predicable and volatile consumption fulfillment.

GridServer

Predictable Consumption

Decision Intensive

Typical Unit of work > 2 seconds



Enterprise Computing

Analytics

Batch

Compute Intensive Analysis

Data Mining

Decision Support

Reporting

...

Volatile Consumption

Service Intensive

Typical Unit of work > 100 ms



Virtualized Service-Oriented Architecture

Business Workflow

Componentized Applications

Data Warehousing

Portals

Volume-intensive

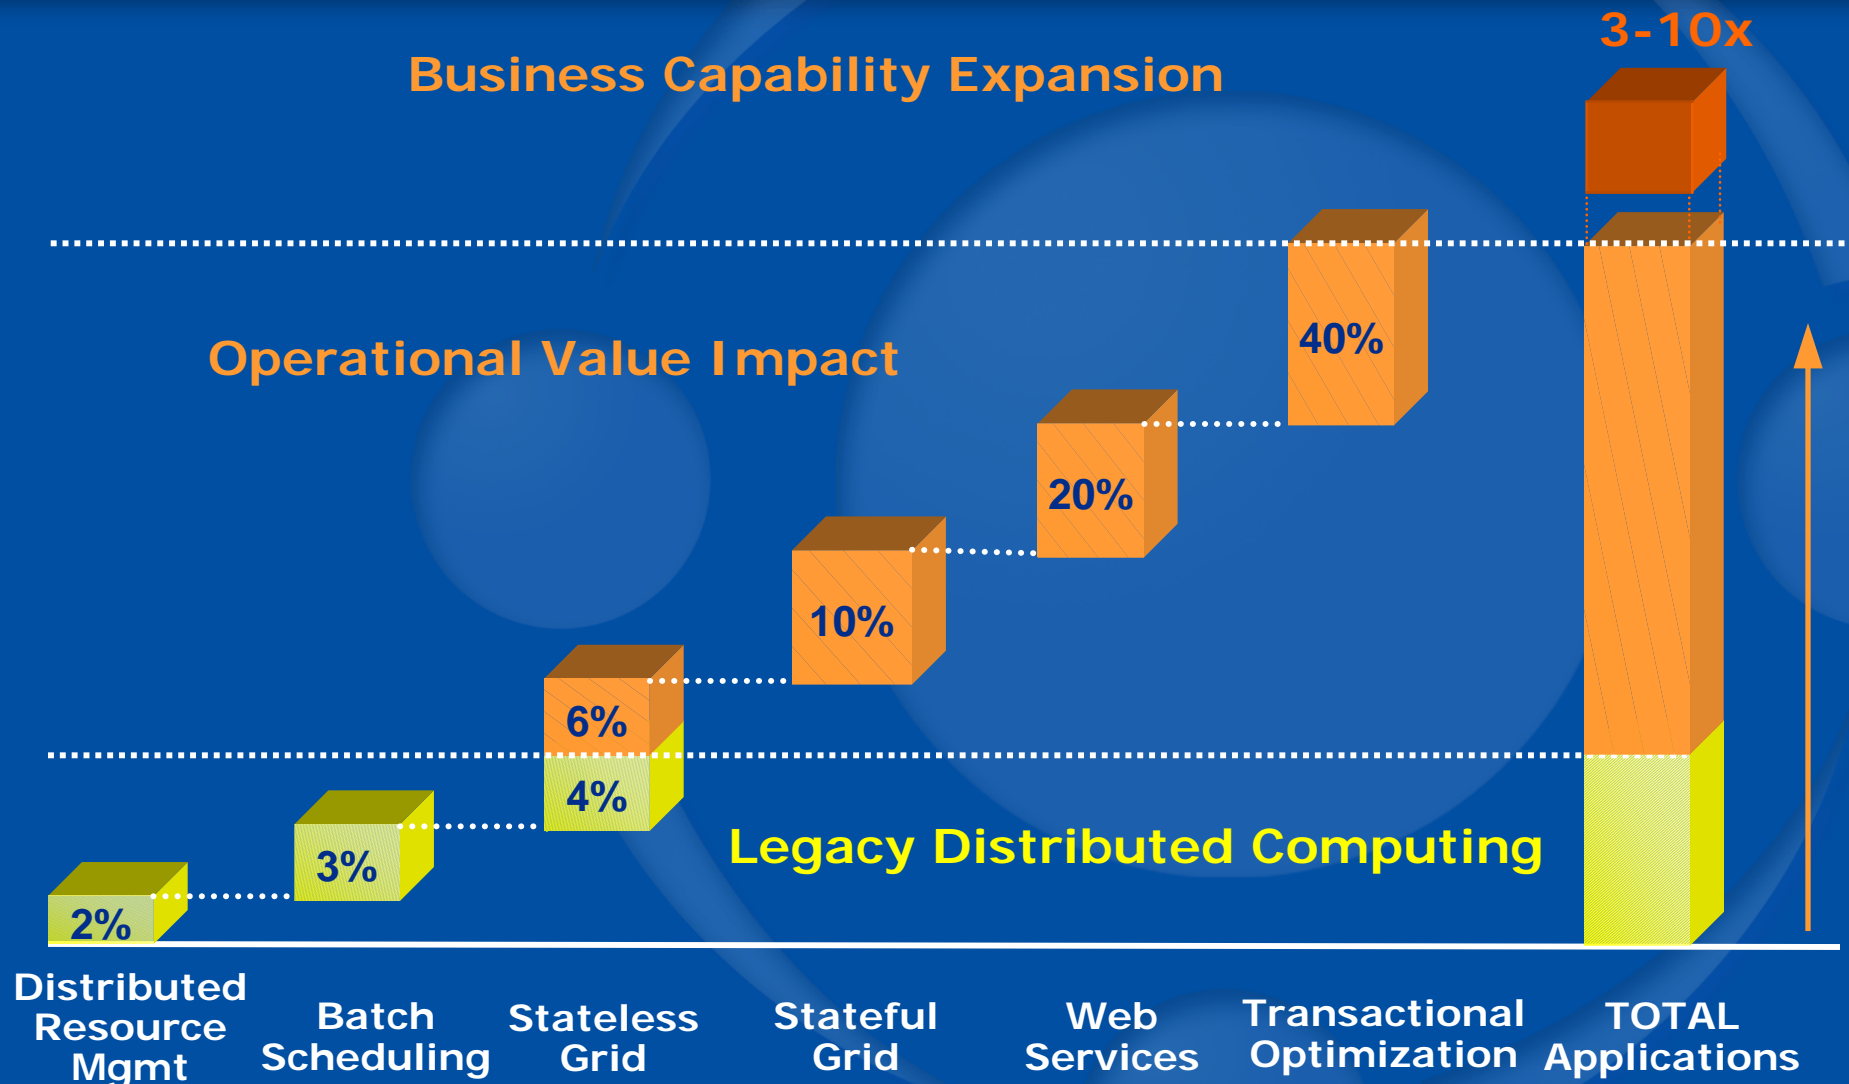
Web Services


...

Capabilities

- Time to results
- Capacity load
- Service Resilience
- Time to build
- Time to deploy

Logical Application Models – Broad Target Applicability

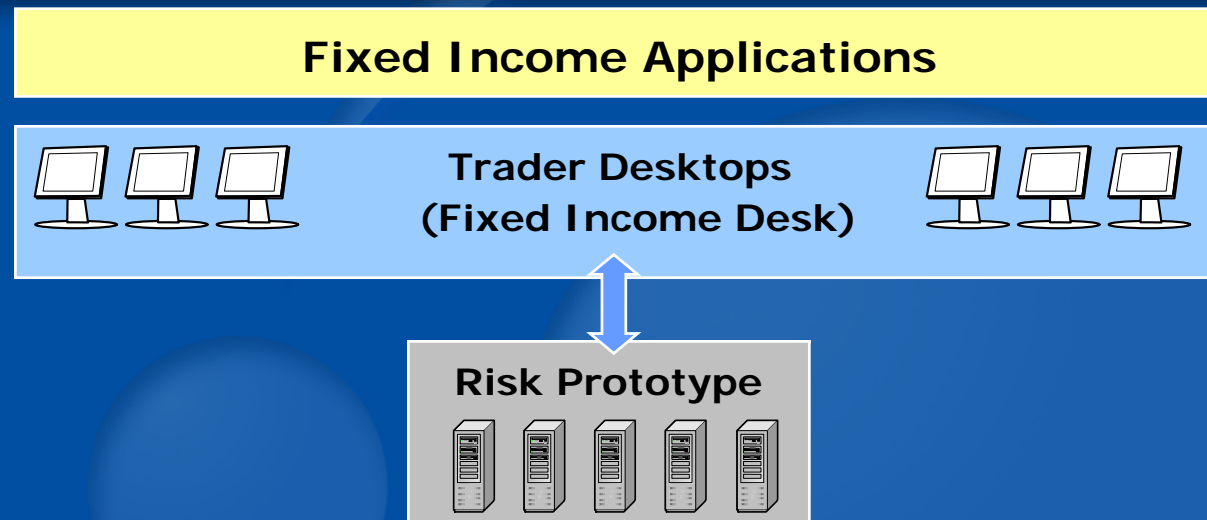




GridServer Success Stories

Predictable and Volatile Consumption

Predictable Consumption – Compute bottlenecks hamper success of STP framework by failing to achieve SLAs



- **Operationally insufficient:**

- Could not scale beyond 5 CPUs
- Lacked Resilience, Fault Tolerance
- Lacked Prioritization, Optimization
- Complex and expensive to manage

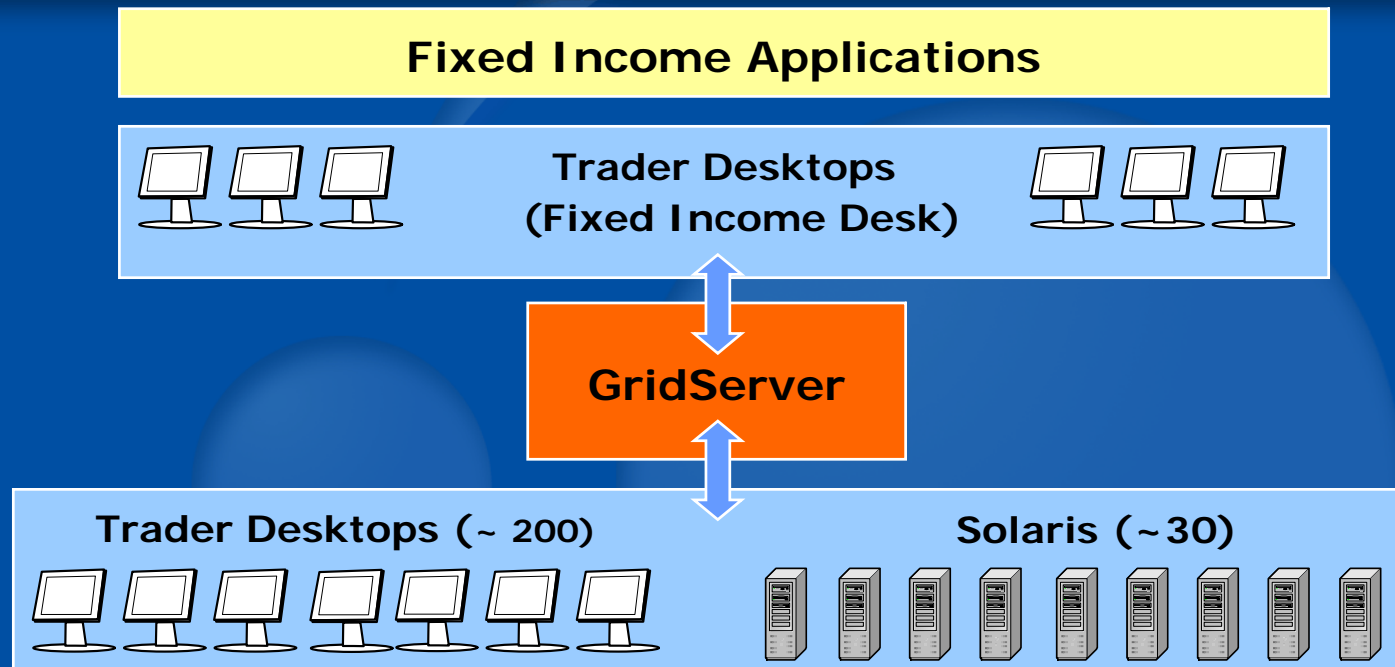
- **Unacceptable Report Turnaround**

- 15 hours, "flying blind"

- **Constrained Trading Volumes:**

- Complex, profitable products
- Limited by time-to-results

Predictable Consumption – GridServer™ provides time-to-results, resilience and scale, better utilizing existing resource



- **Price / Performance Gains:**

- Leverage underutilized hardware

- **Improved Time-to-Results:**

- 15 hours -> 15 minutes
- Batch to -> on demand in-day

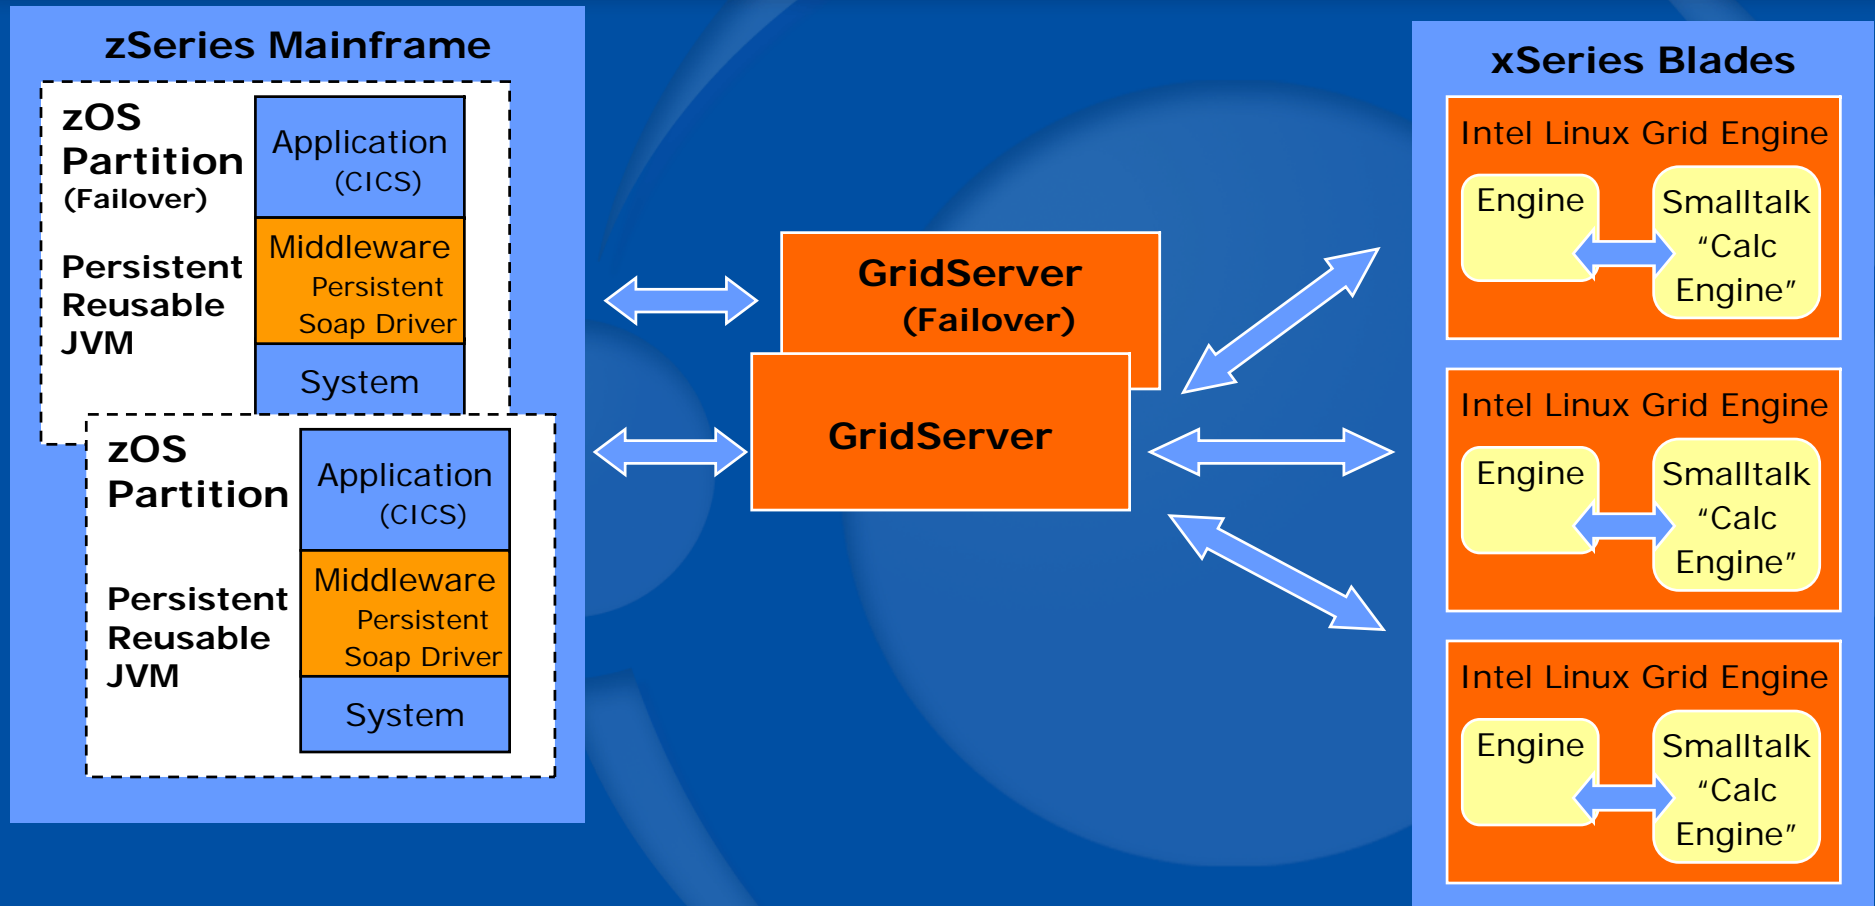
- **Increased Accuracy 25x:**

- 4,000 -> 100,000 simulations

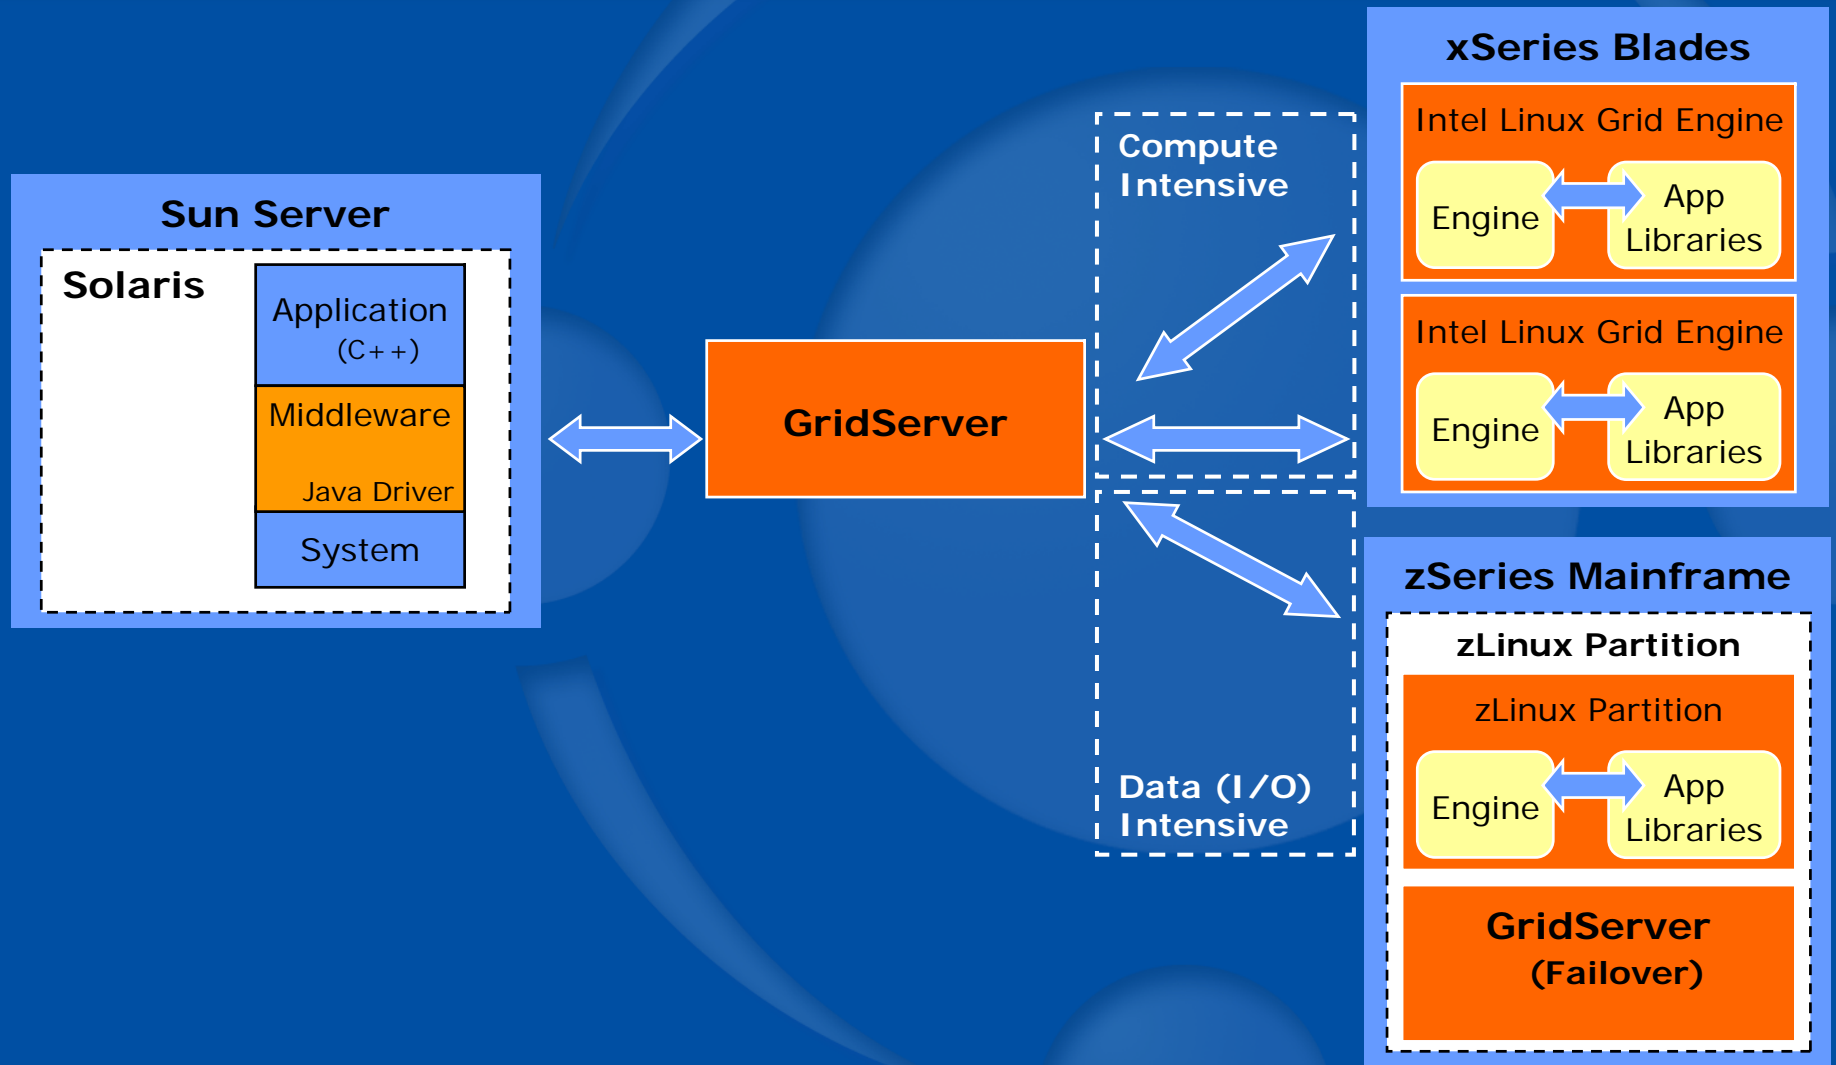
- **Increased Exotics Trading:**

- Volume 3x
- Additional profit per deal \$1MM

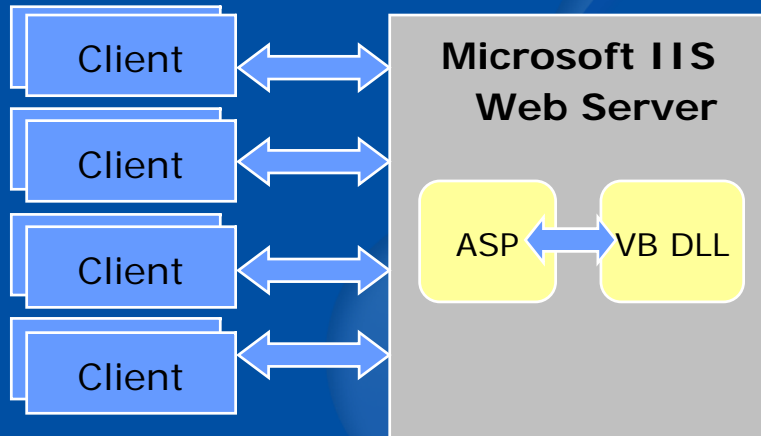
Volatile Consumption – GridServer™ Virtualized Service Oriented Architecture reduces transaction cost by 90%



Volatile Consumption — GridServer™ manages compute-intensive and data-intensive service optimization

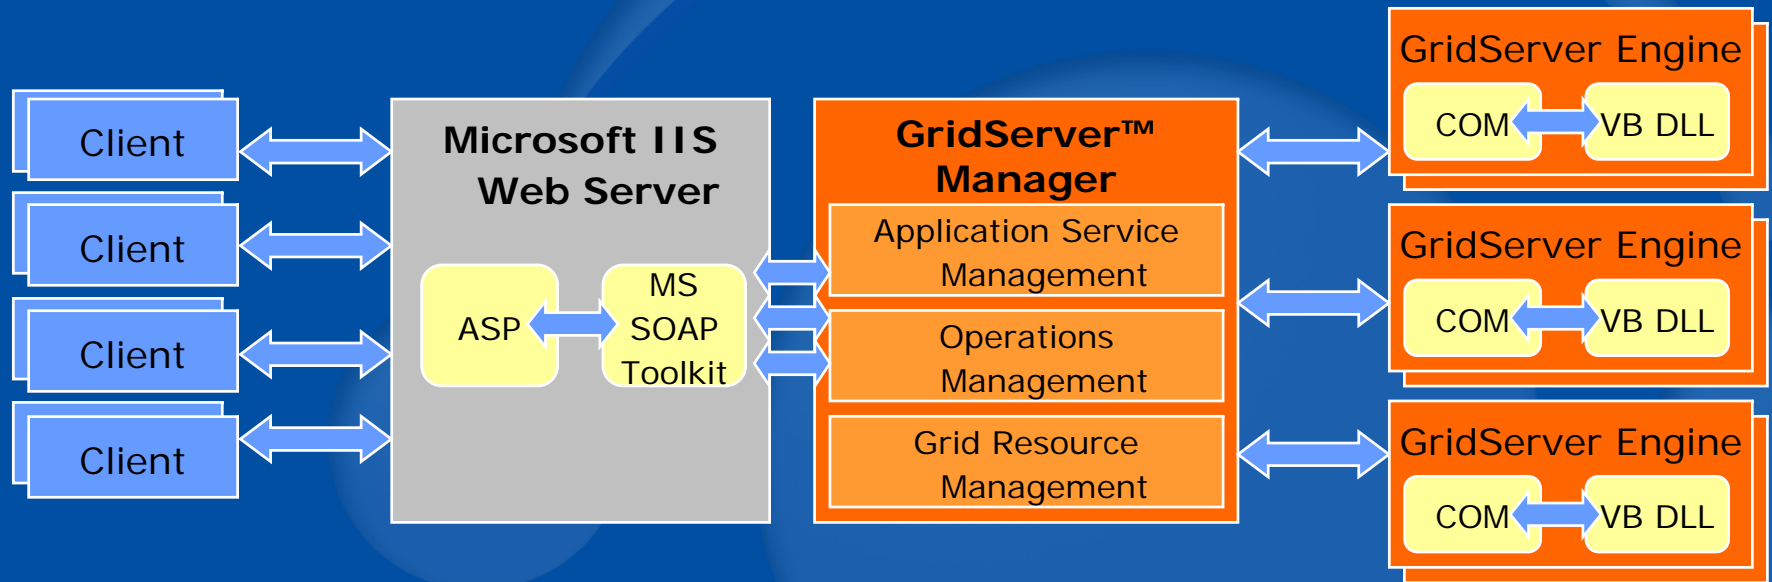


Volatile Consumption – SLA demanded for a resource-intensive web portal could not be met.



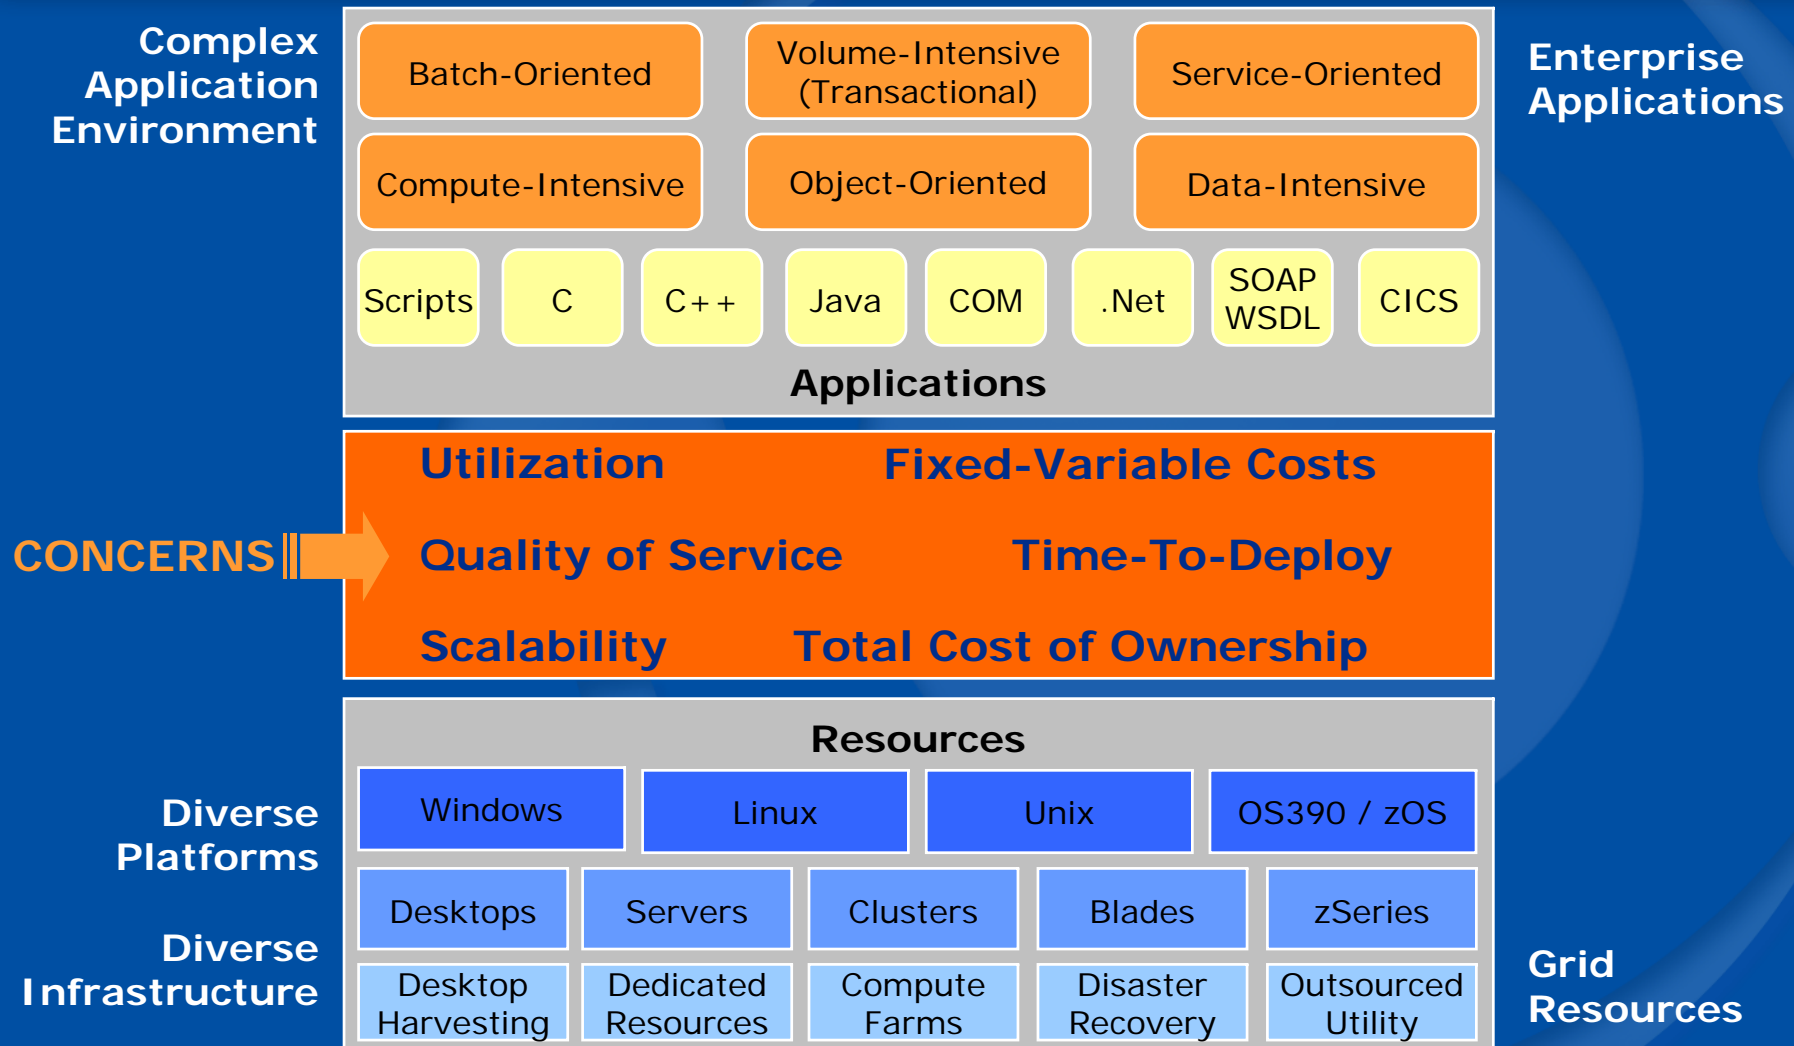
- **Web application performance degraded**
 - as number of clients (users) increases
- **Service levels breached**
 - 'stovepipe' Web Server became performance bottleneck
- **Result – Time-to-results extended beyond acceptable level**

Volatile Consumption – GridServer™ radically improves time-to-results, scale and safeguards service levels for portals.

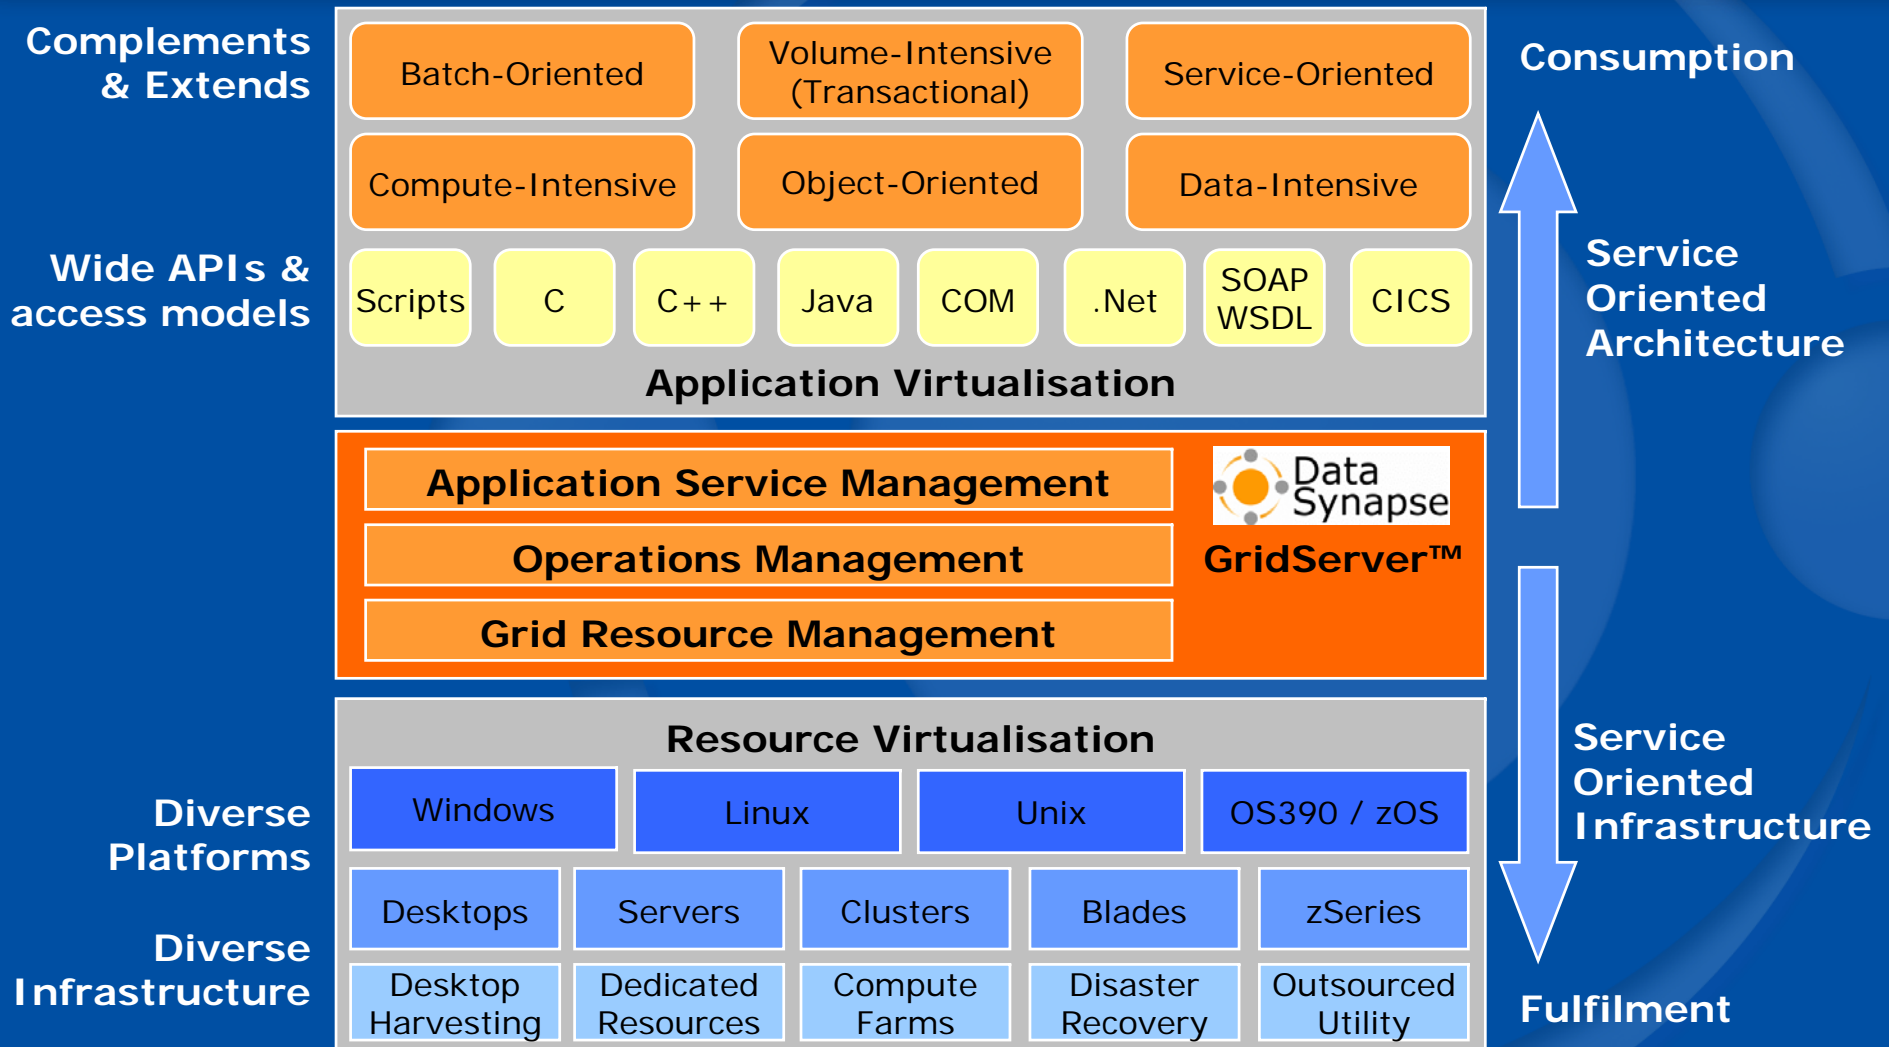


- **Response time improved:**
 - From 500ms to 125ms
- **Improvement in SLA:**
 - 4x
- **GridServer provides basis for future scale**

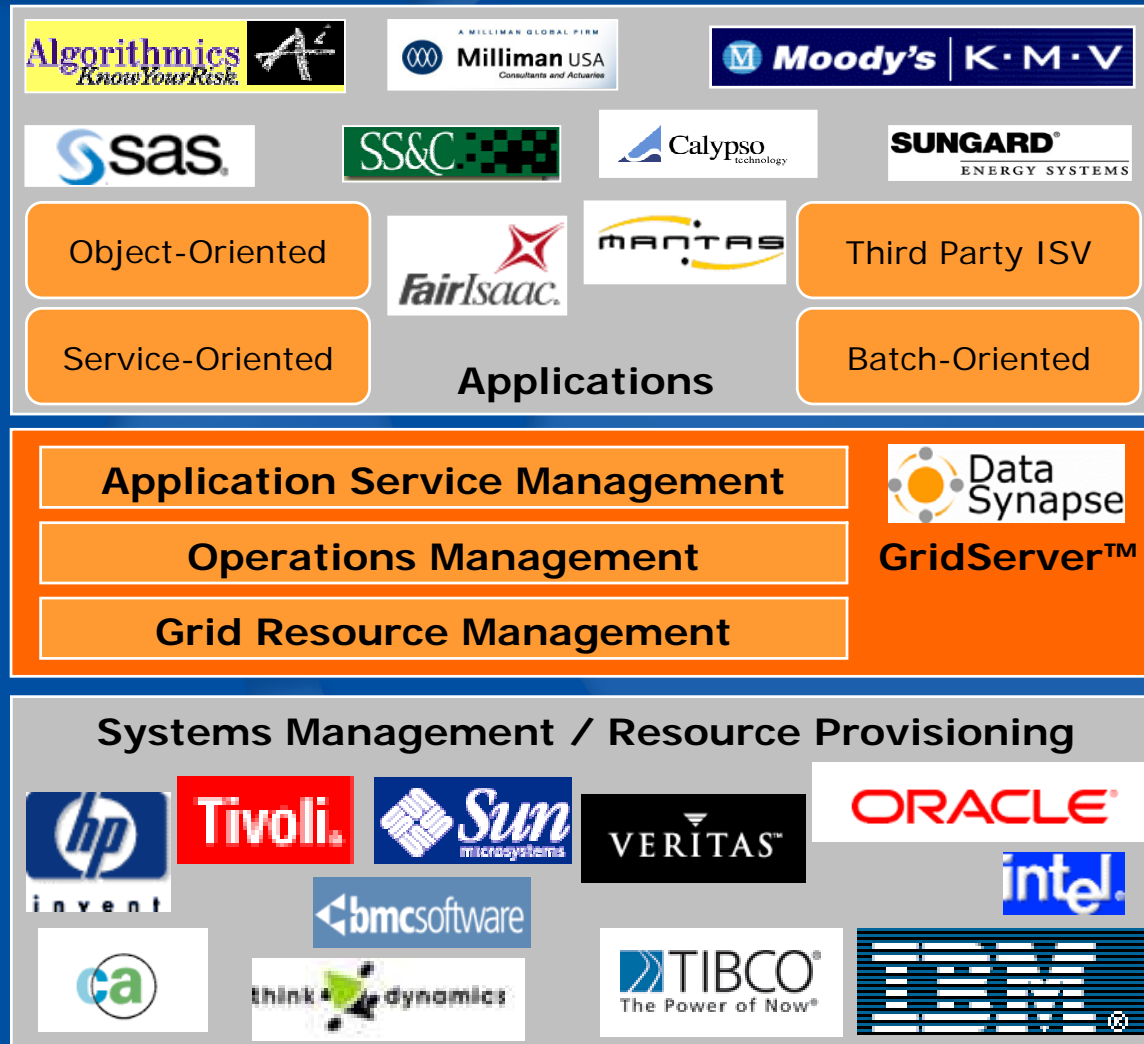
Complexity and diversity represent a roadblock to the coordination of diverse applications and resources



GridServer™ reduces cost and improves performance through a virtualised Service Oriented Architecture that complements and extends the existing complex and diverse infrastructure.



GridServer™ is an “Ecosystem” that is a pluggable, extensible architecture for open interoperability



GridServer™ offers flexible integration strategies

Synchronous & Asynchronous

Clients

Java or
Generated
Proxy

.Net or
Generated
Proxy

C++
API

SOAP

COM
API

C API

Commands
or Scripts

GridServer Manager

Virtualized Invocation

Java
class

.Net
assembly

COM


C++ DLL
or .so

C DLL
or .so

Executables
.exe, .pl, etc.

Hosted Services

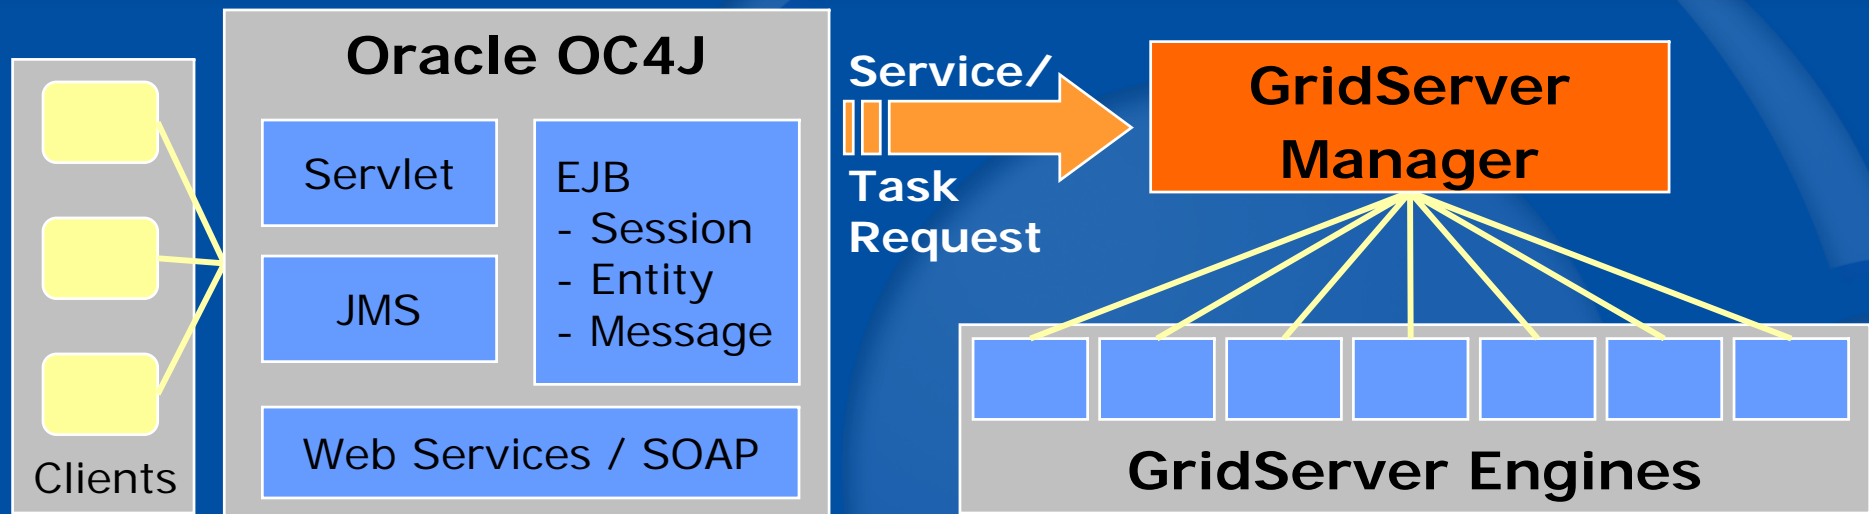
Stateless & Stateful



GridServer™

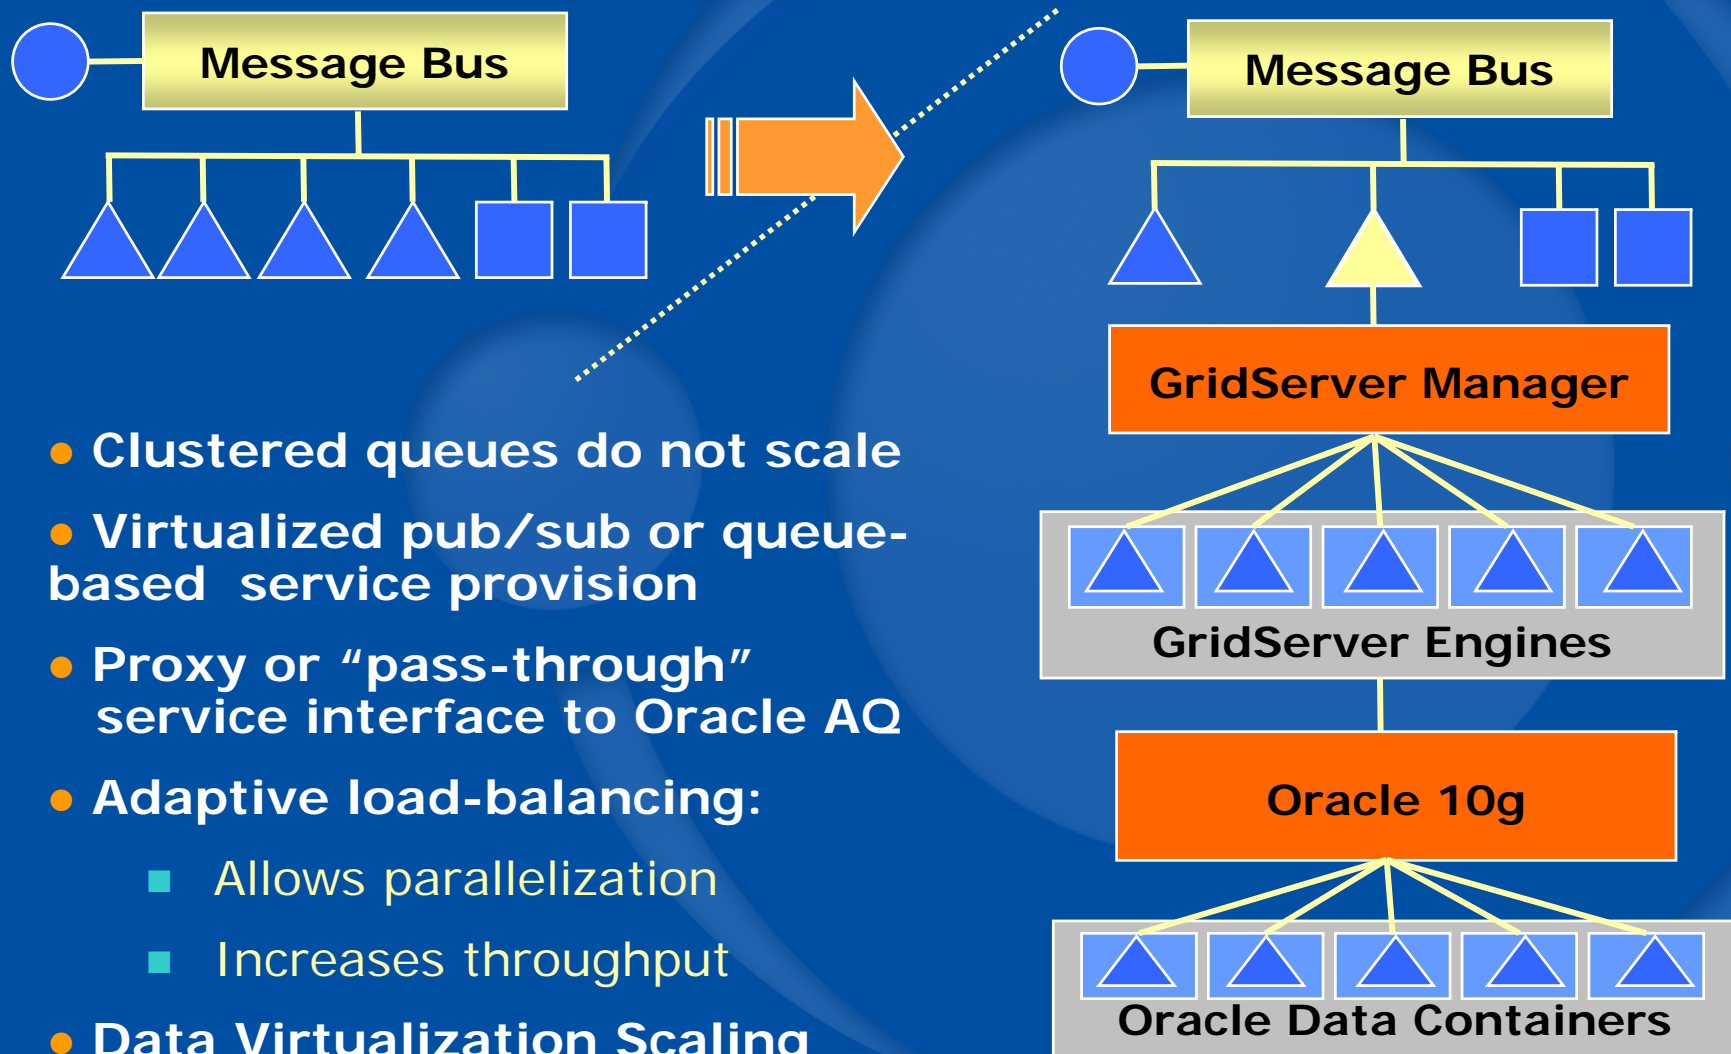
Application Blueprints - Oracle

Blueprint –J2EE Dynamic Scaling Model



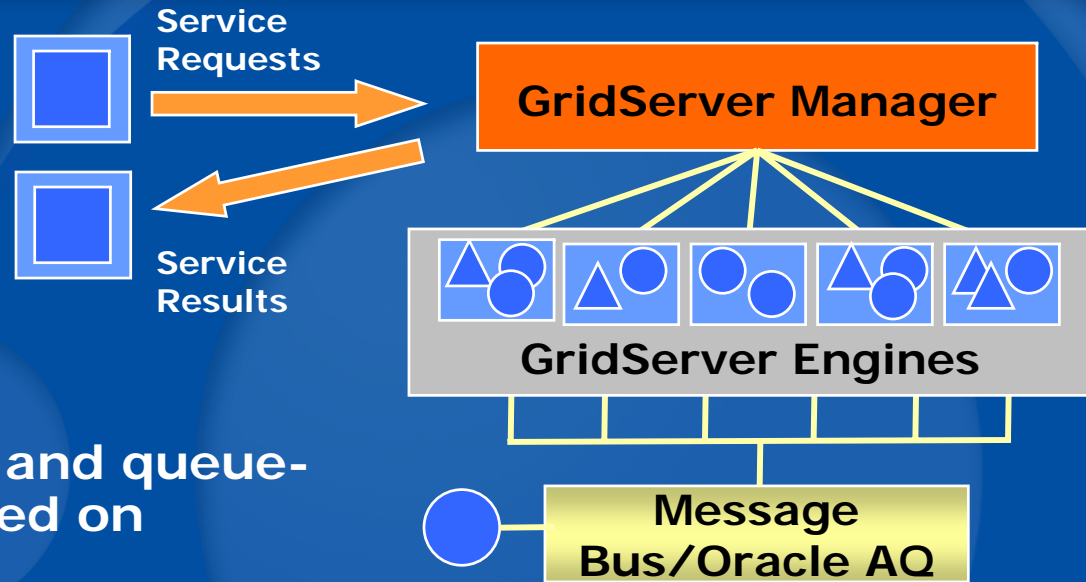
- J2EE components make service/task request to GridServer for stateless and stateful execution
- GridServer provides the virtualized assured resilient execution fabric
- GridServer Engines execute service - J2EE, Java, .Net, binaries, DLLs, shared objects, legacy

Blueprint – EAI Fabric dynamic scaling model



- Clustered queues do not scale
- Virtualized pub/sub or queue-based service provision
- Proxy or “pass-through” service interface to Oracle AQ
- Adaptive load-balancing:
 - Allows parallelization
 - Increases throughput
- Data Virtualization Scaling

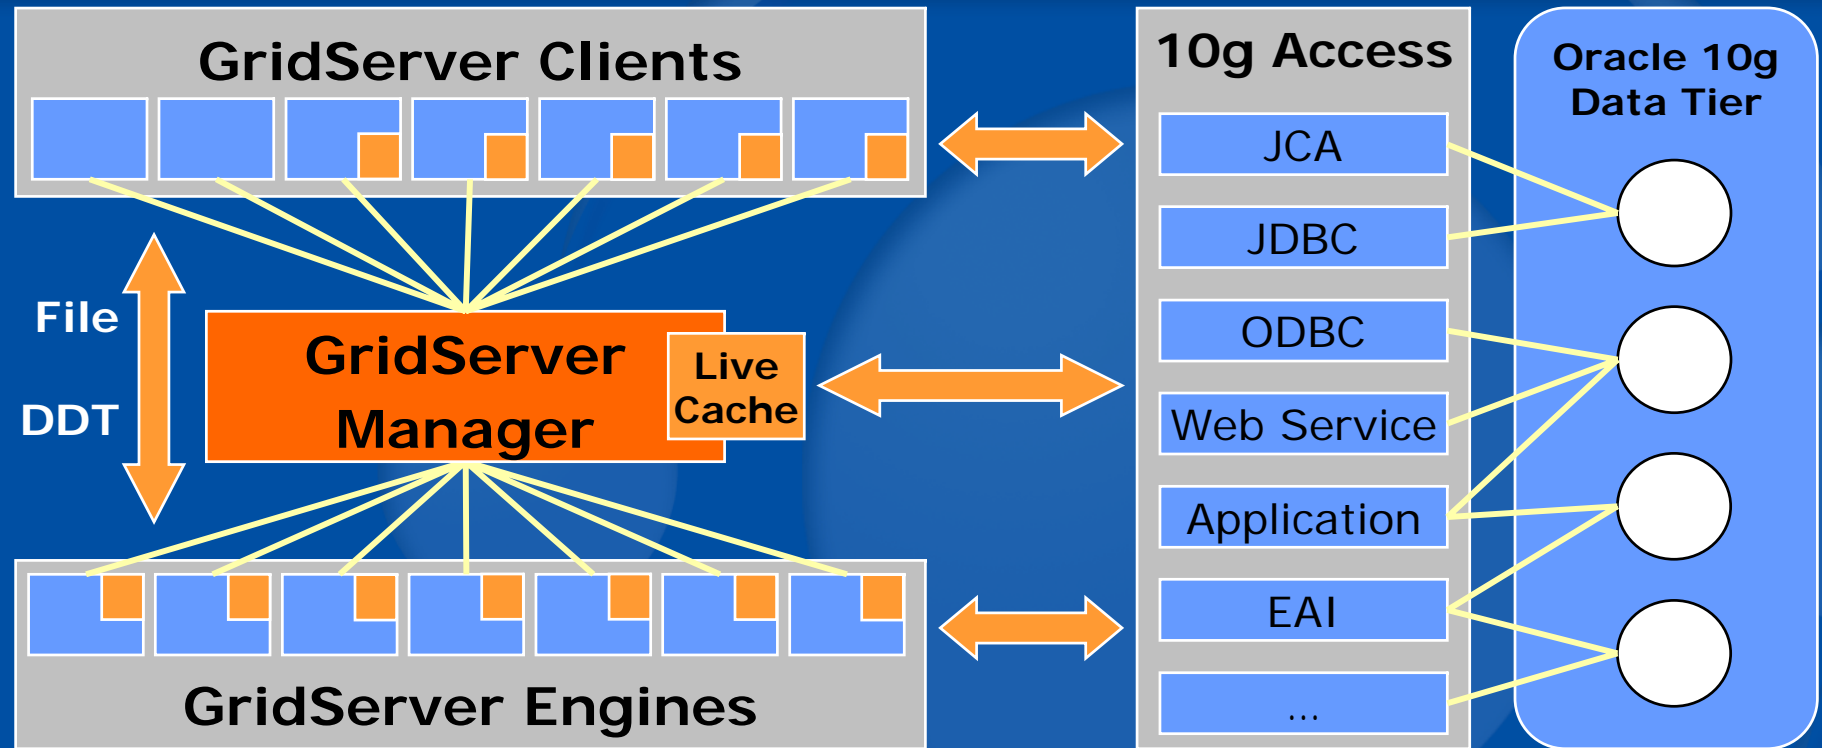
Blueprint – EAI fabric extended consumption model



- Publish / Subscribe and queue-based services hosted on GridServer
- GridServer provides:
 - Application Virtualization
 - Capacity, scale and resilience
- EAI provides:
 - A2A, B2B and B2C Adaptors
 - Transformation , Business Rules

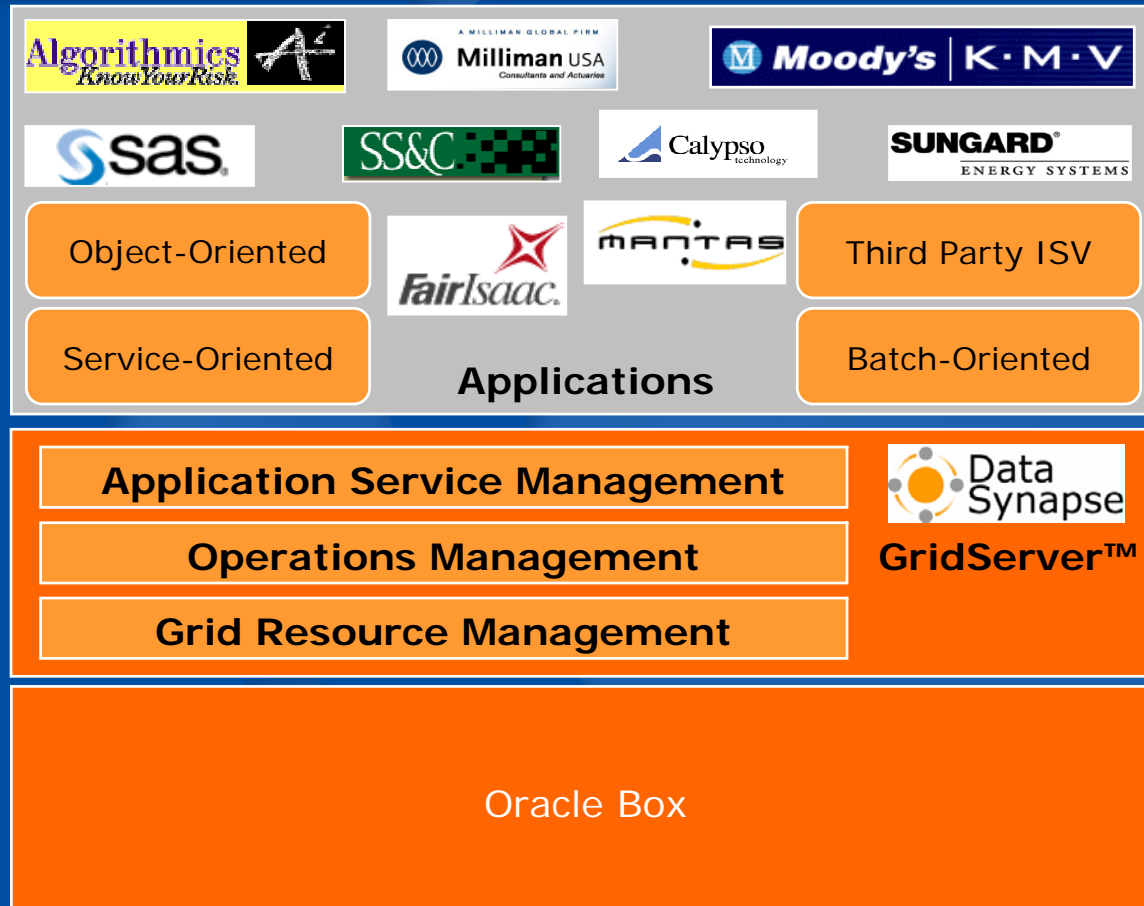



Blueprint – EIS Dynamic Scaling application model



- GridServer Clients, Managers and Engines access EIS through Oracle 10g native access mechanisms
- GridServer provides data distribution mechanisms – file transfer, direct data transfer (DDT) and LiveCache

GridServer™ & Oracle 10g “Ecosystem” that is a pluggable, extensible architecture for open interoperability

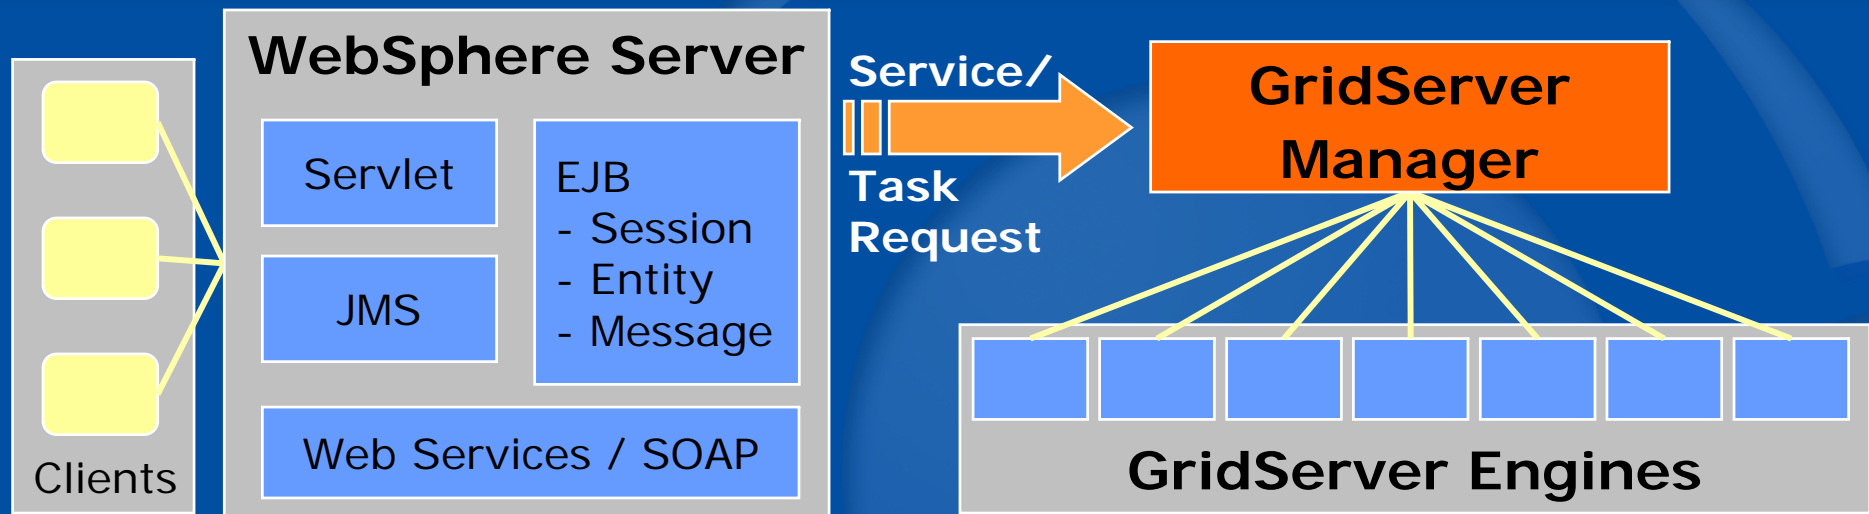




GridServer™

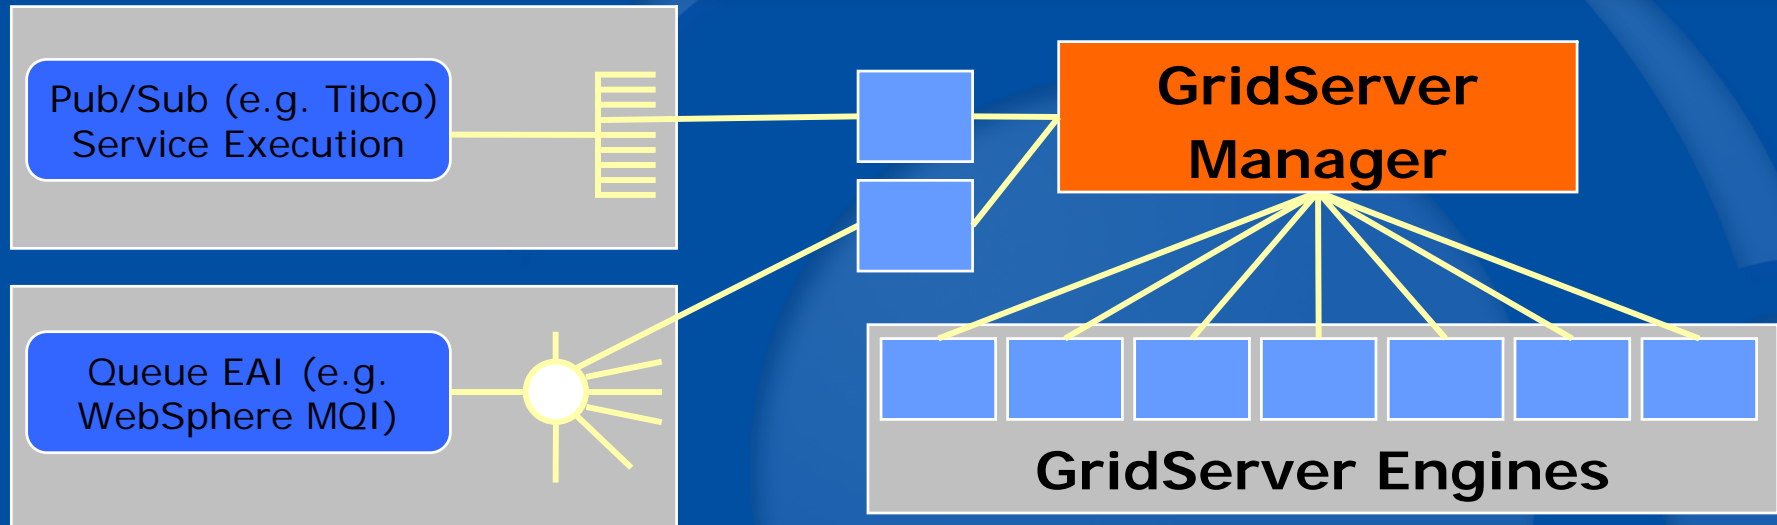
Application Blueprints - IBM

On Demand Scaling WebSphere Application Server



- **WebSphere Application Server (WAS) components make service/task request to GridServer for stateless and stateful execution**
- **GridServer migrates WAS static/deterministic "Clustering" or "spray & pray" to "scale & guarantee" through adaptive load-balancing & parallelization**
- **GridServer Engines execute WAS and non WAS services - J2EE, Java, .Net, binaries, DLLs, shared objects, legacy creating interoperability execution environment for WAS users**

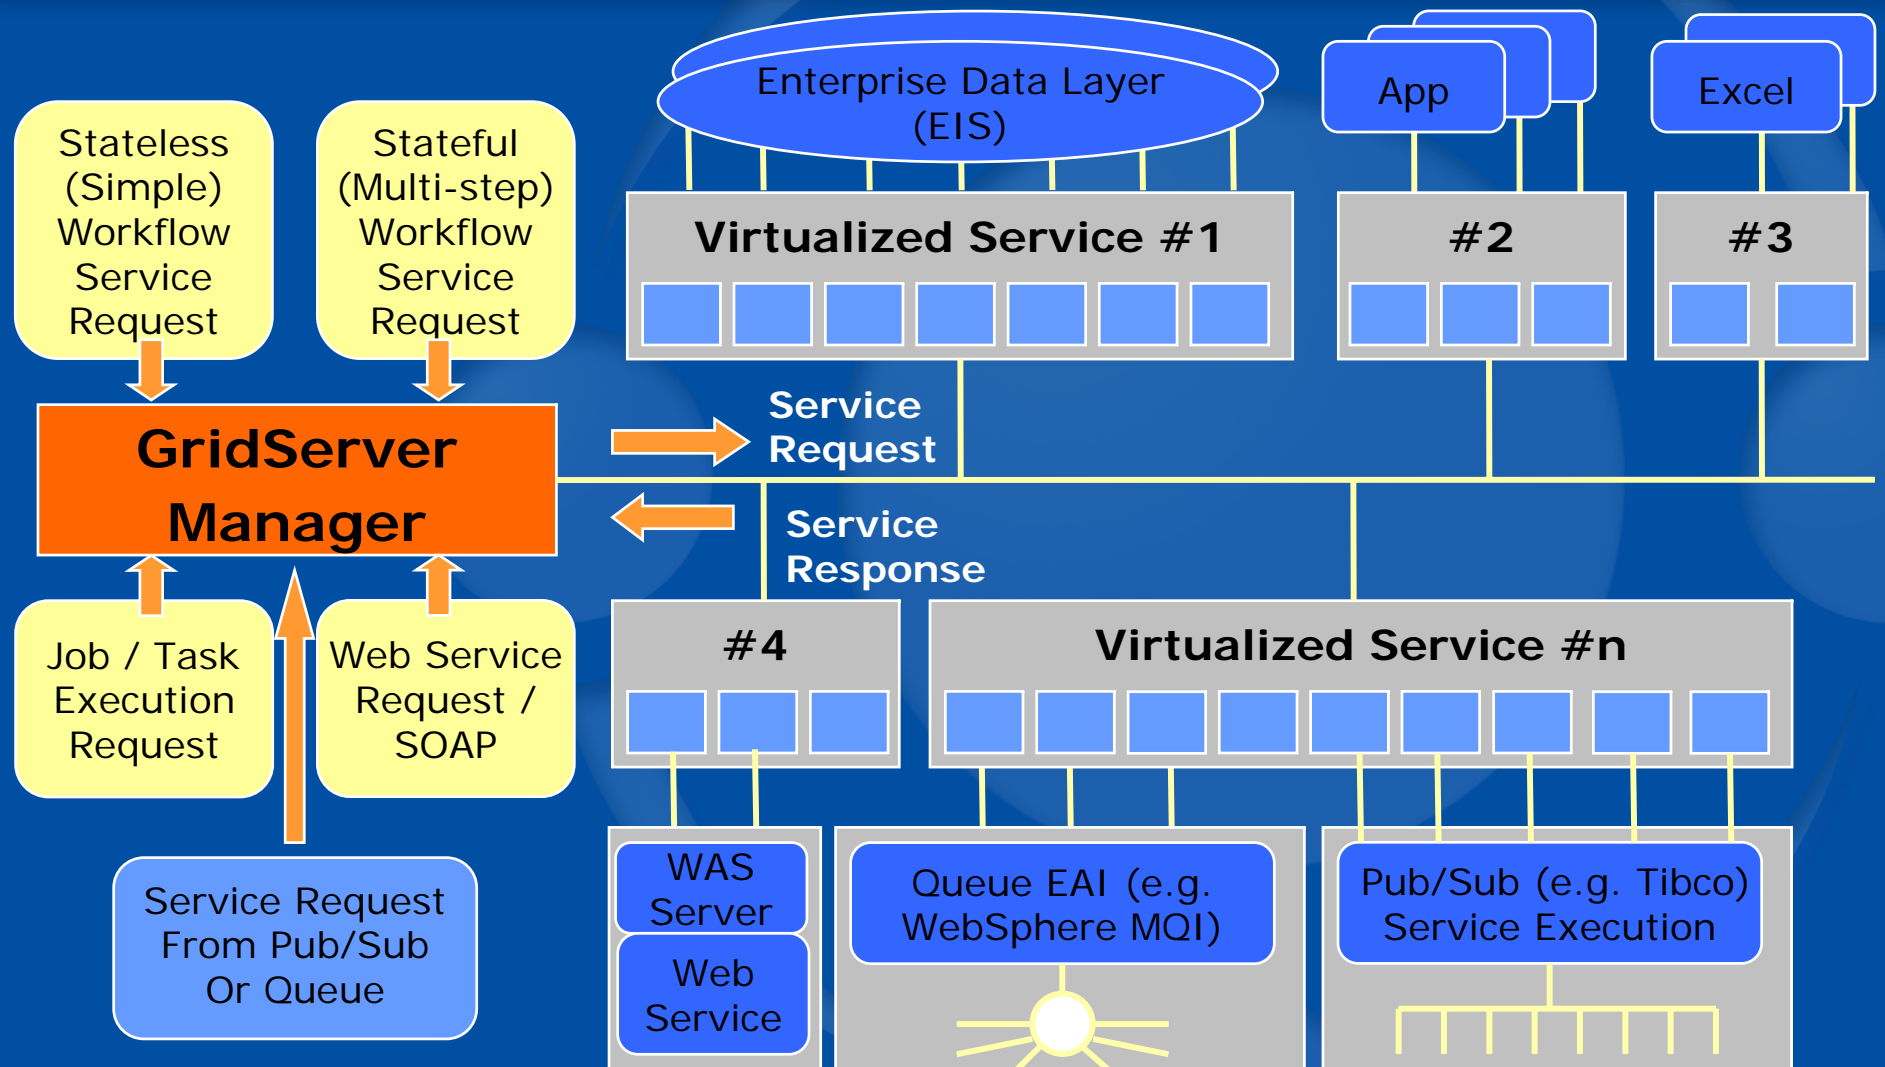
GridServer can scale WAS Enterprise Application Integration (EAI) services



- GridServer acts as the virtualized execution fabric
- Business-level abstraction replaces technical API
- Adaptive load balancing
- Autonomous service configuration
- Autonomous service provisioning

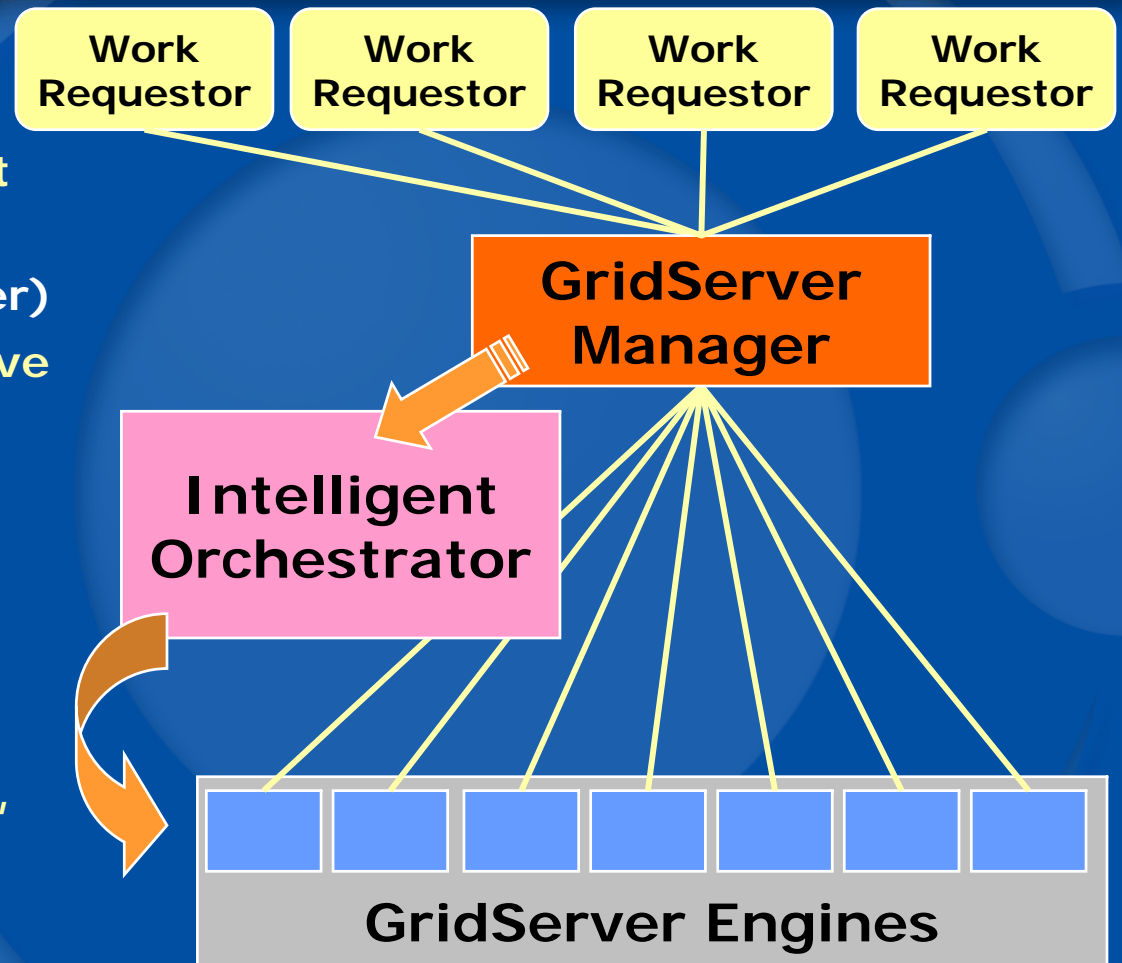
Extended Operating Environment with WAS

Bi-Directional Servicing between WAS & GridServer

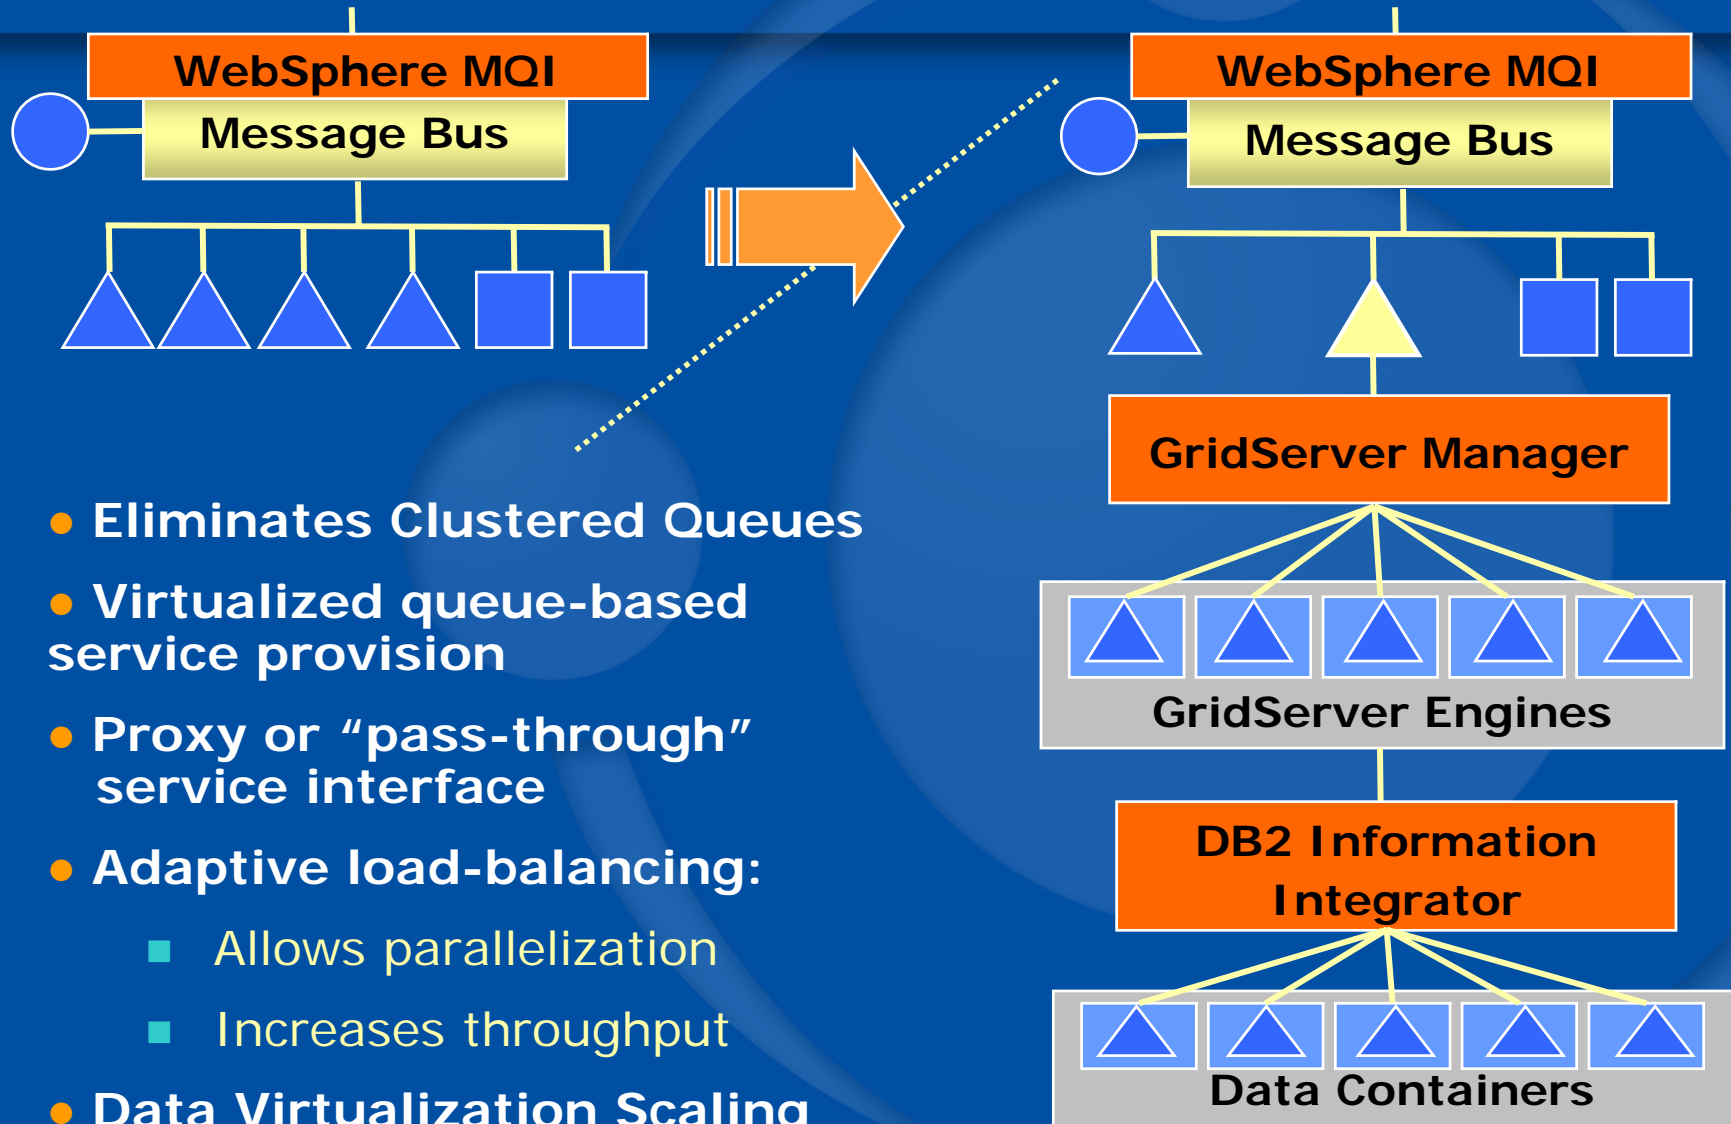


Tivoli TIOP & GridServer Resource 'Sourcing' For The Grid

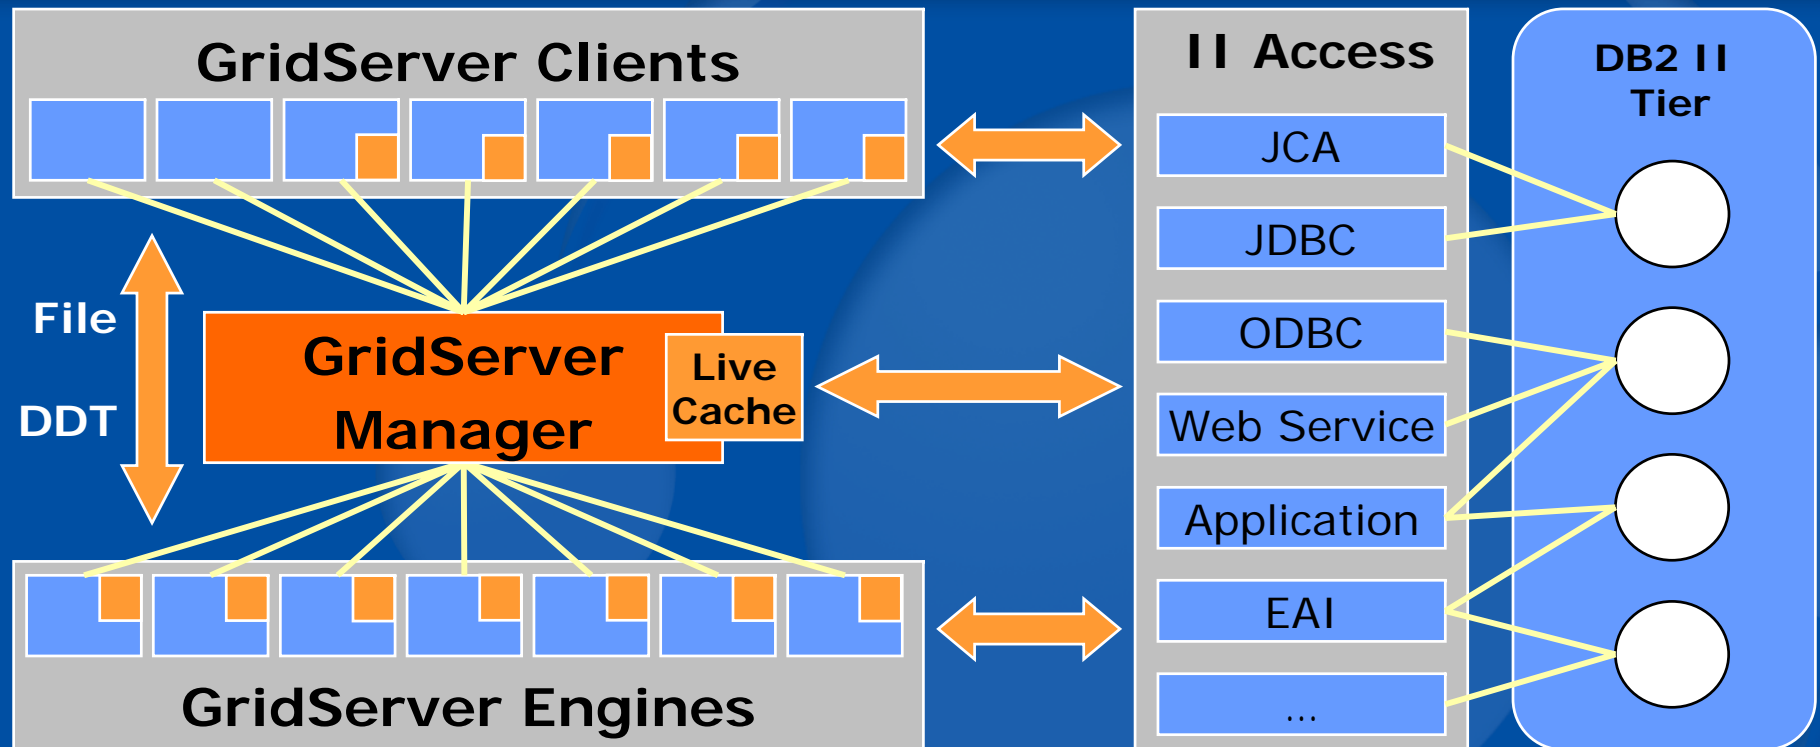
- **Work Requestor**
 - Atomic units of work
 - Service Level Agreement
 - Frequency/rate, Response
- **Resource Manager (GridServer)**
 - Dynamically add / remove resources and services
 - Monitor activity and service level
 - Predict resource requirements
- **Tivoli Intelligent Orchestrator (TIO) Support:**
 - Monitor (CPU Utilization, Arrival Rates)
 - Maintain Service Level Policies / Agreements
 - Add and configure resources on demand



Blueprint – EAI Fabric dynamic scaling model

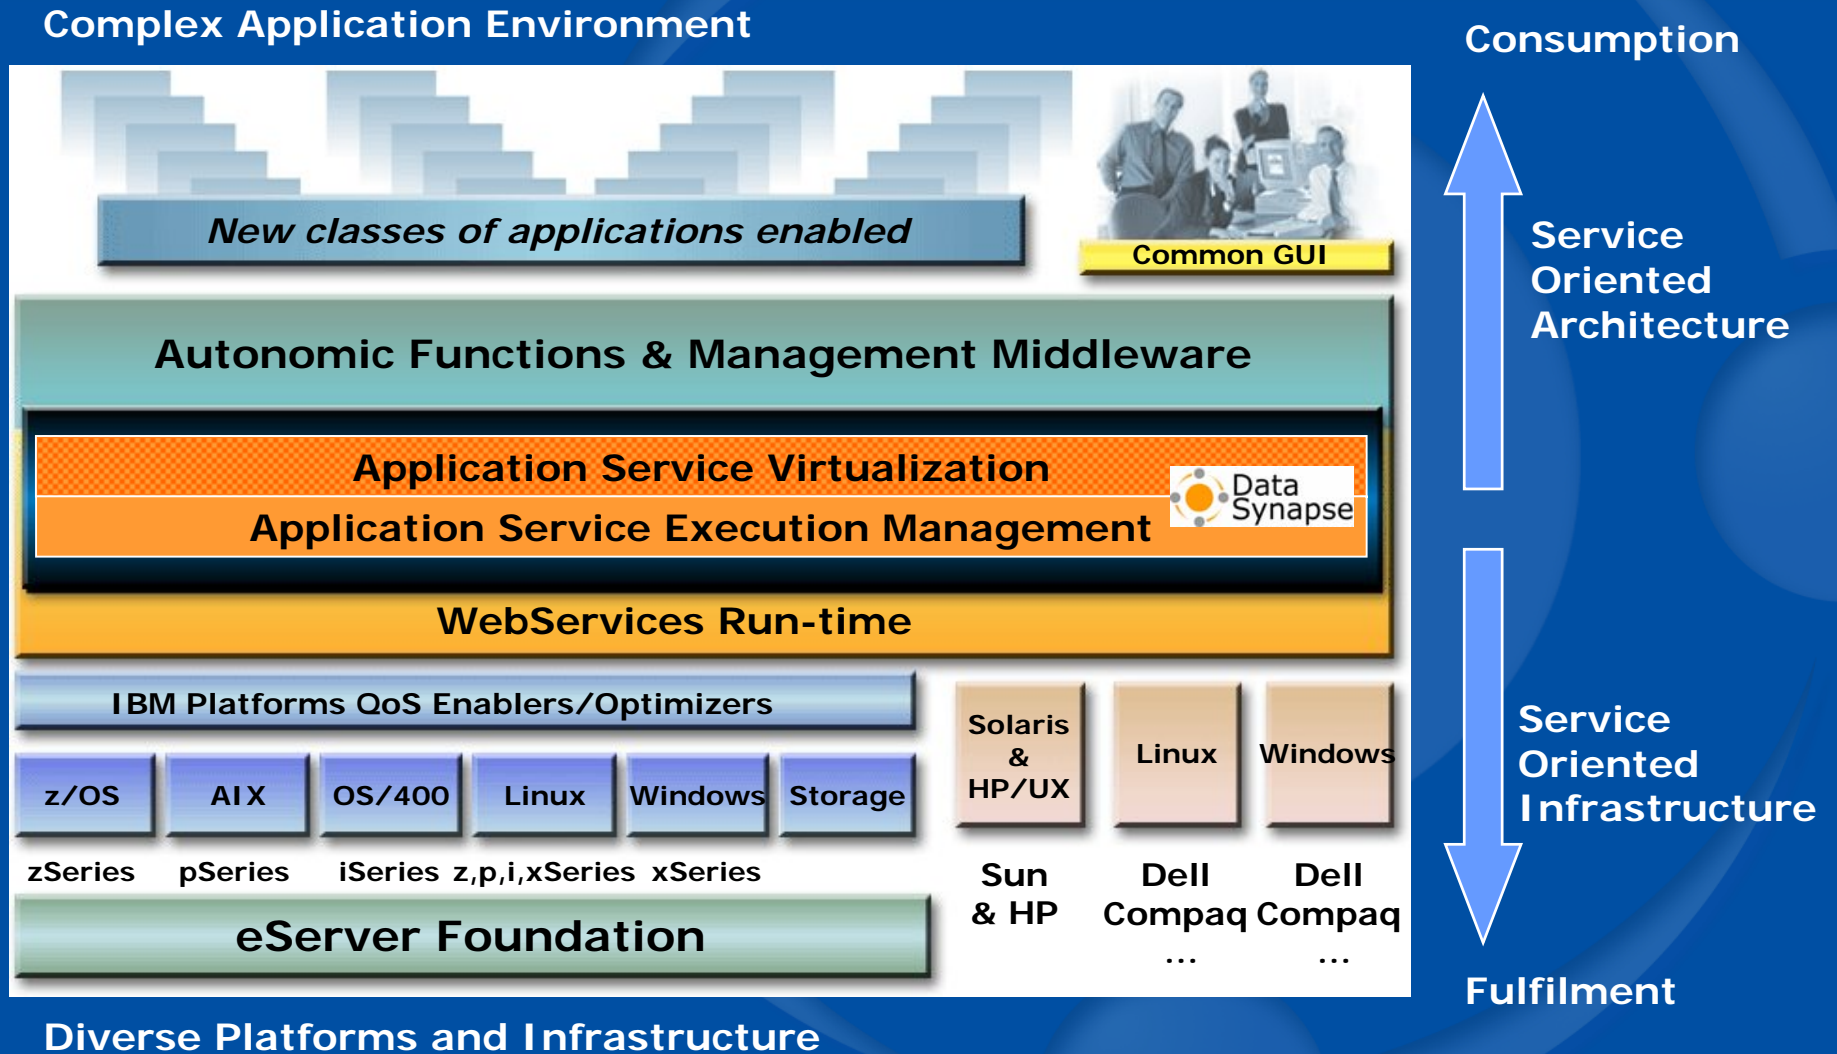



DB2 Information Integrator (II) & GridServer Service & Data Virtualization



- GridServer Clients, Managers and Engines access EIS through DB2 II native access mechanisms
- DB2 II provides common access method and data transaction management

DataSynapse in IBM's Open Architectural Framework






Thank you

GridServer and WebSphere – Use Cases

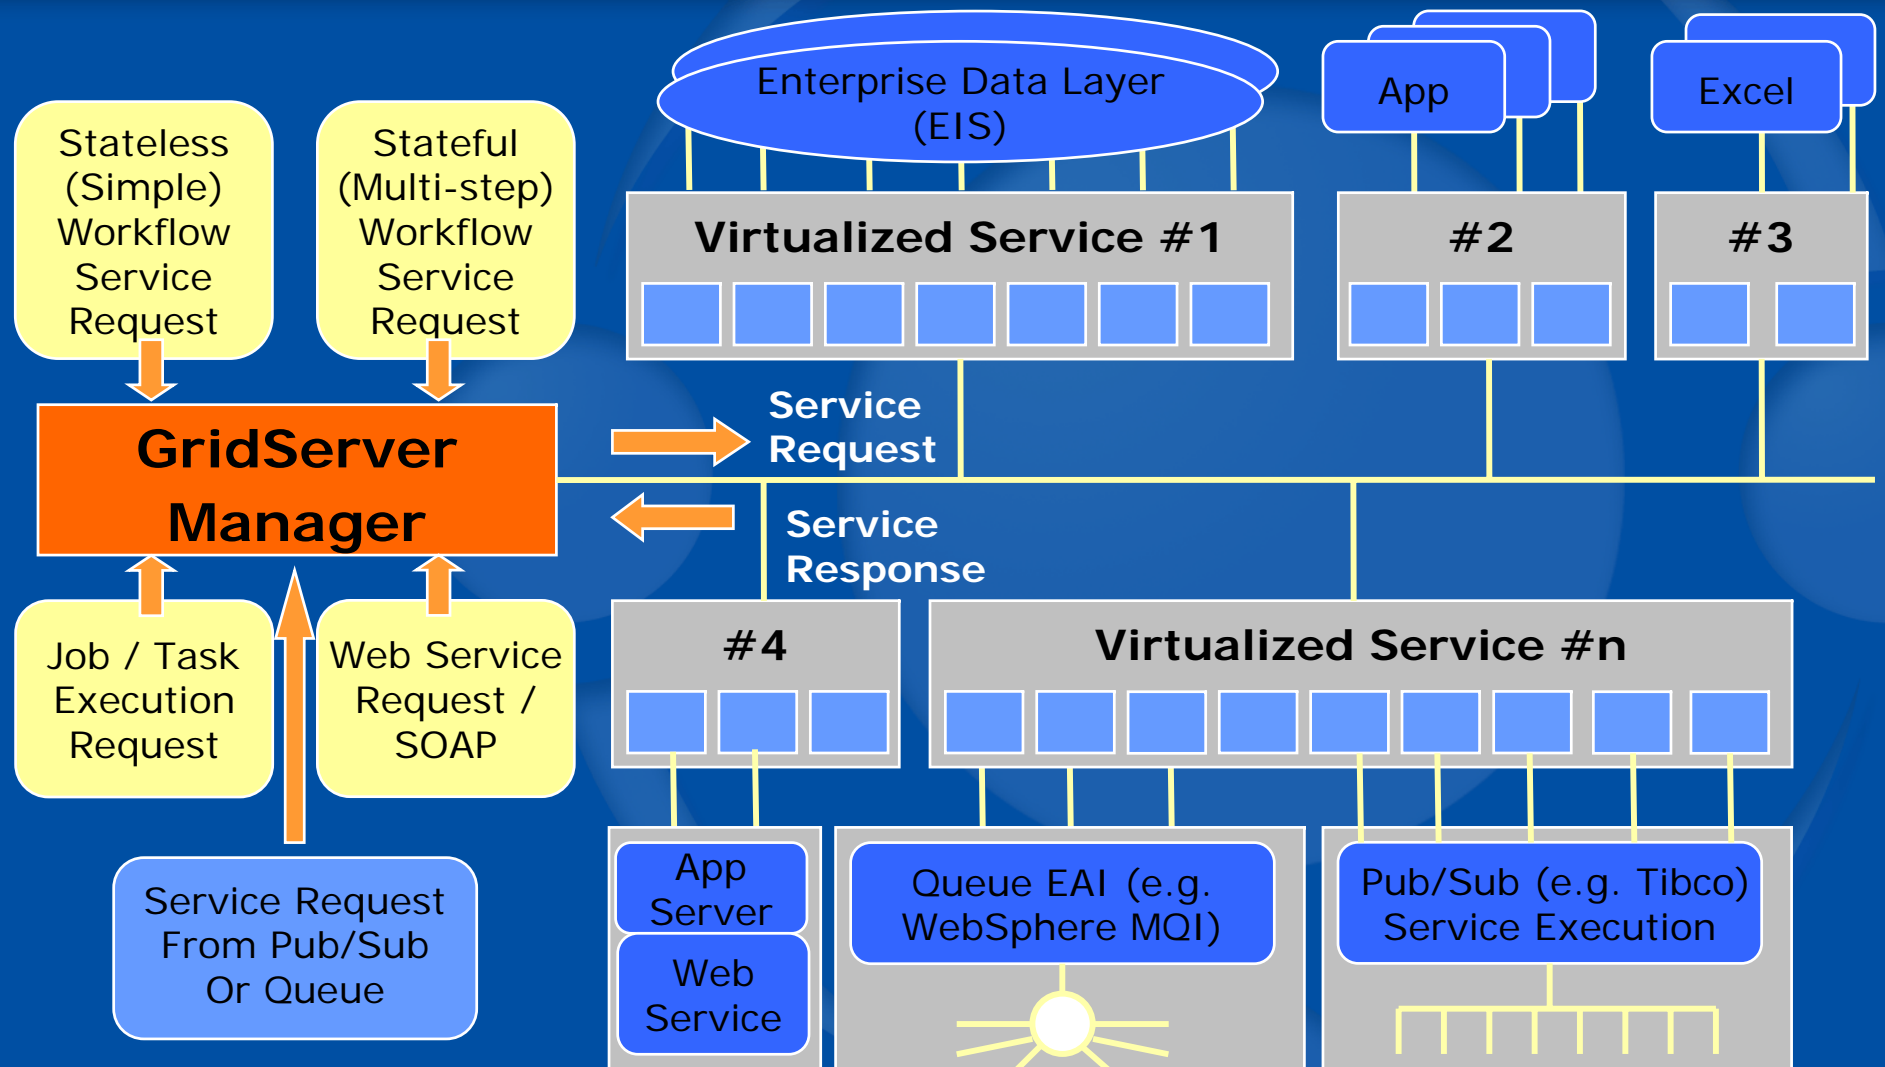
Use GridServer when...	Use WebSphere when...
Volatile consumption – units of work >75ms	Transactions – units of work <75ms
"I/O-light" – unit of work has some CPU bound processing	Purely I/O based
Web Services execution environment for Java <u>and</u> .Net	Web Services execution environment for Java
Runtime dynamic allocation of resources	Cloning & loading JVM latency does not impact throughput
Loosely-coupled service oriented services	
Extend consumption of WAS services from non-WAS environments, for example: <ul style="list-style-type: none">■ .NET or Object Oriented■ Pub/Sub environments	



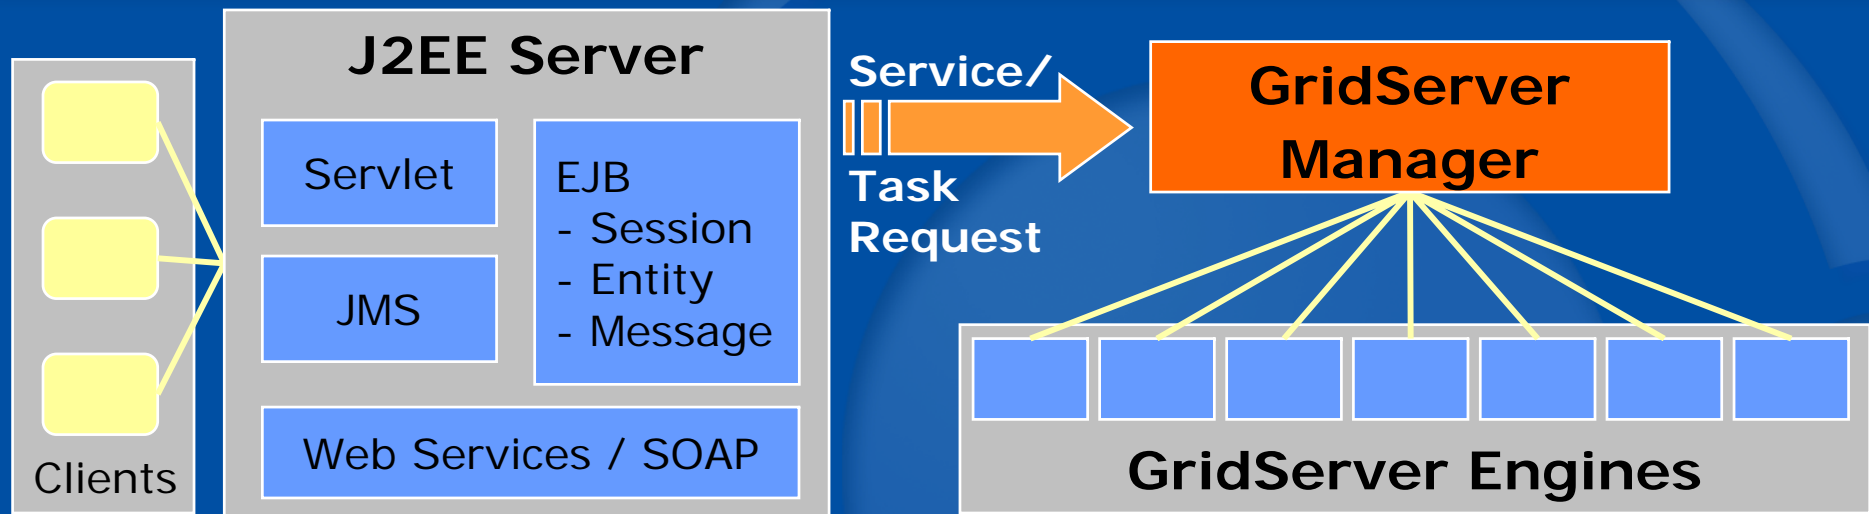
GridServer™

Application Blueprints

Blueprint – GridServer™ enables virtualized service execution that eliminates stovepipe inefficiencies



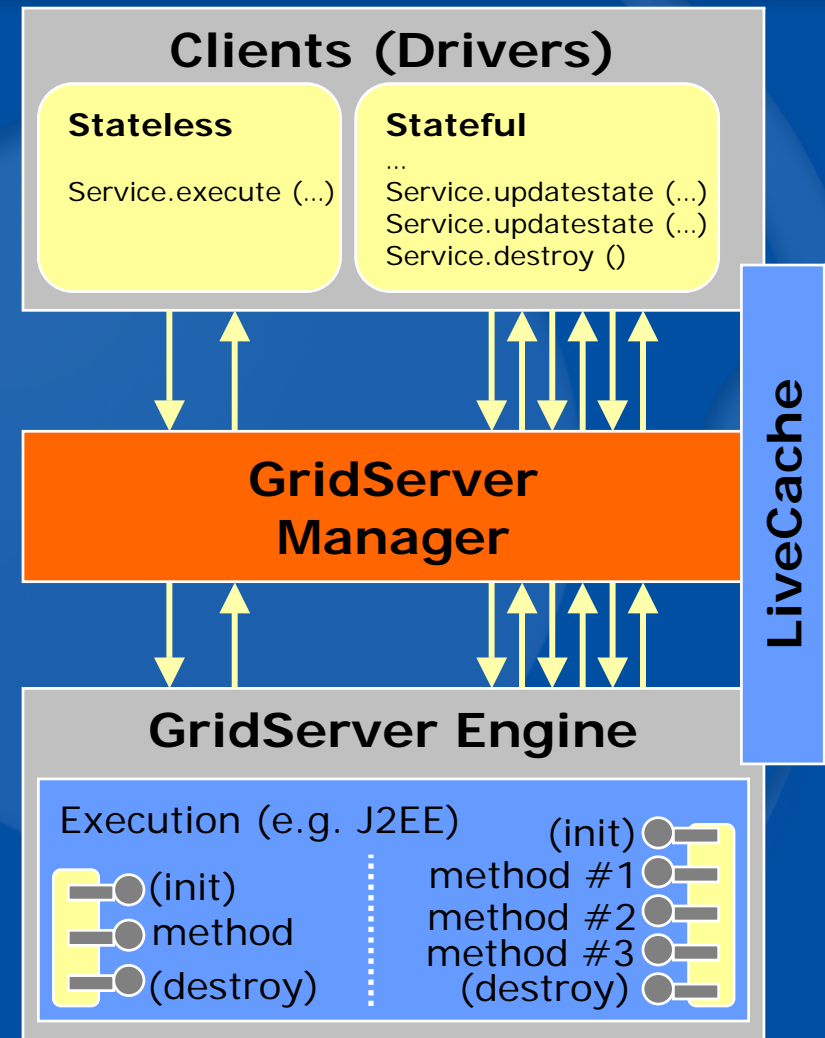
Blueprint – GridServer™ as a virtual service-oriented application execution fabric for J2EE application servers



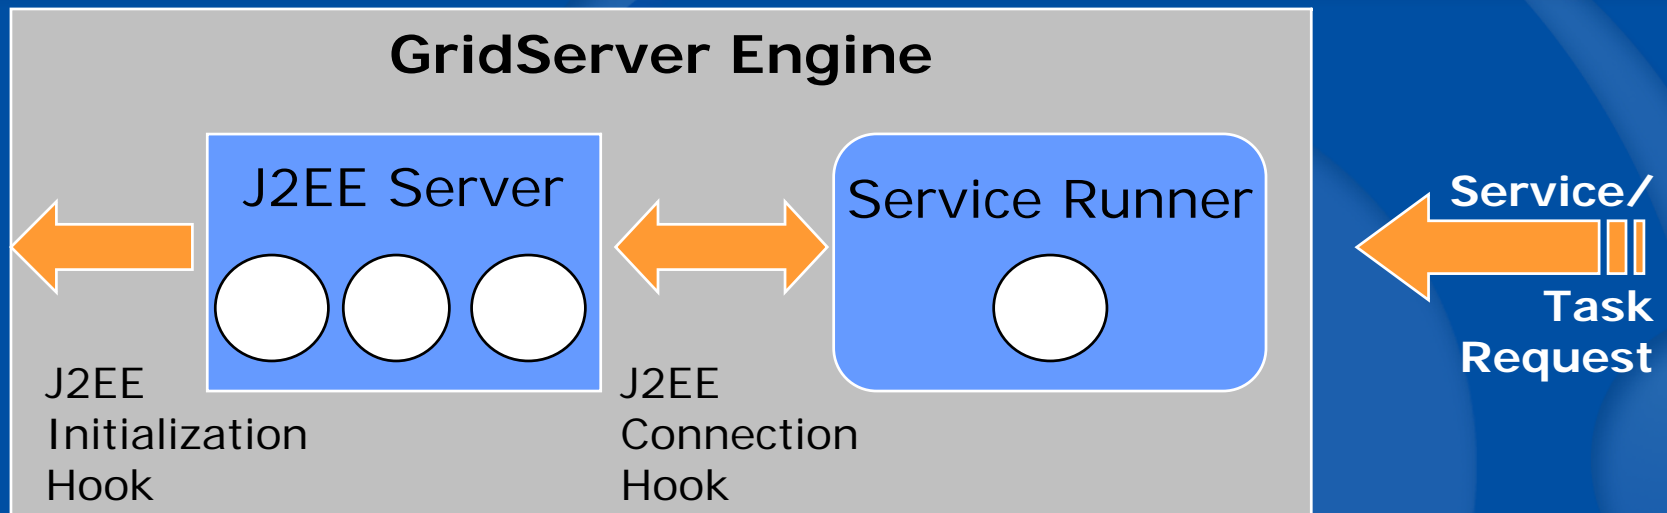
- J2EE components make service/task request to GridServer for stateless and stateful execution
- GridServer provides the virtualized assured resilient execution fabric
- GridServer Engines execute service - J2EE, Java, .Net, binaries, DLLs, shared objects, legacy

Blueprint – GridServer hosts scalable affinity-aware J2EE virtual stateless and stateful services

- **GridServer virtual application fabric supports**
 - Stateless service execution
 - Stateful service execution
 - WS, SOAP, Java, .Net,...
- **Engine is a stateful container:**
 - Provisions services on-demand of Manager to fulfil service consumption requests
- **Durable state maintenance:**
 - LiveCache™ distributed data caching (synchronization)
 - Application checkpointing (recovery)



Blueprint – GridServer™ manages stateless and stateful service and task execution – example: J2EE

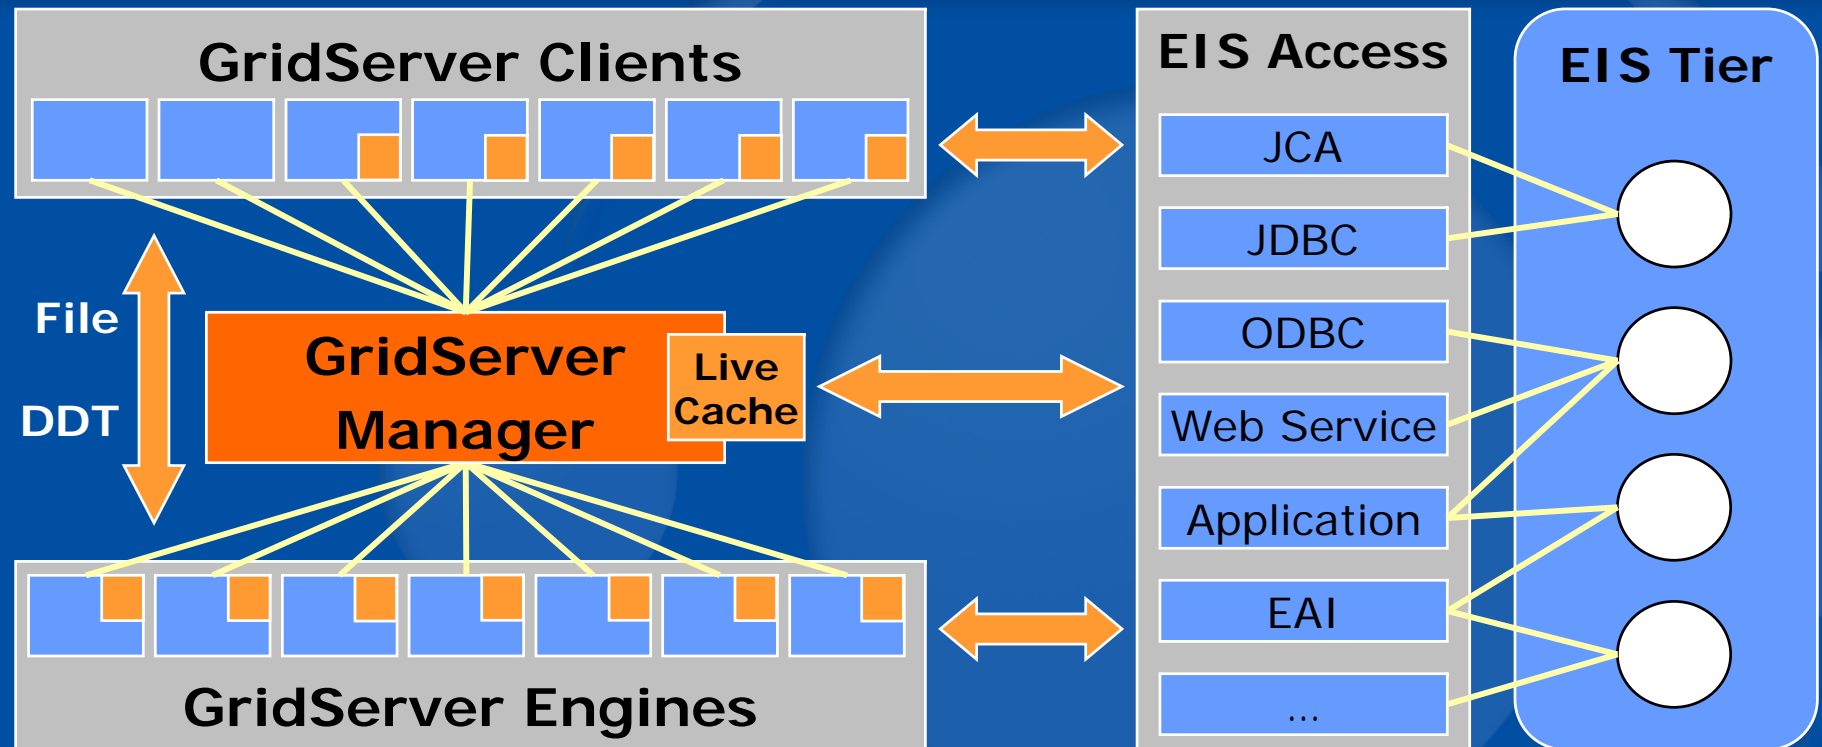


- J2EE can fulfil roles of a stateless and stateful execution end-point on virtualized application fabric
- GridServer Engine can extend to manage lifecycle of long running J2EE Java objects and resources
- Initialize J2EE server
- Access to J2EE services through JNDI

GridServer and Application Server – Use Cases

Use GridServer to compliment AS when...	Use Application Server (AS) when...
Volatile consumption – units of work >75ms	Transactions – units of work <75ms
"I/O-light" – unit of work has some CPU bound processing	Purely I/O based
Web Services execution environment for Java <u>and</u> .Net	Web Services execution environment for Java
Runtime dynamic allocation of resources	Cloning & loading JVM latency does not impact throughput
Loosely-coupled service oriented services	
Extend consumption of AS services from non-AS environments, for example: <ul style="list-style-type: none">■ .NET or Object Oriented■ Pub/Sub environments	

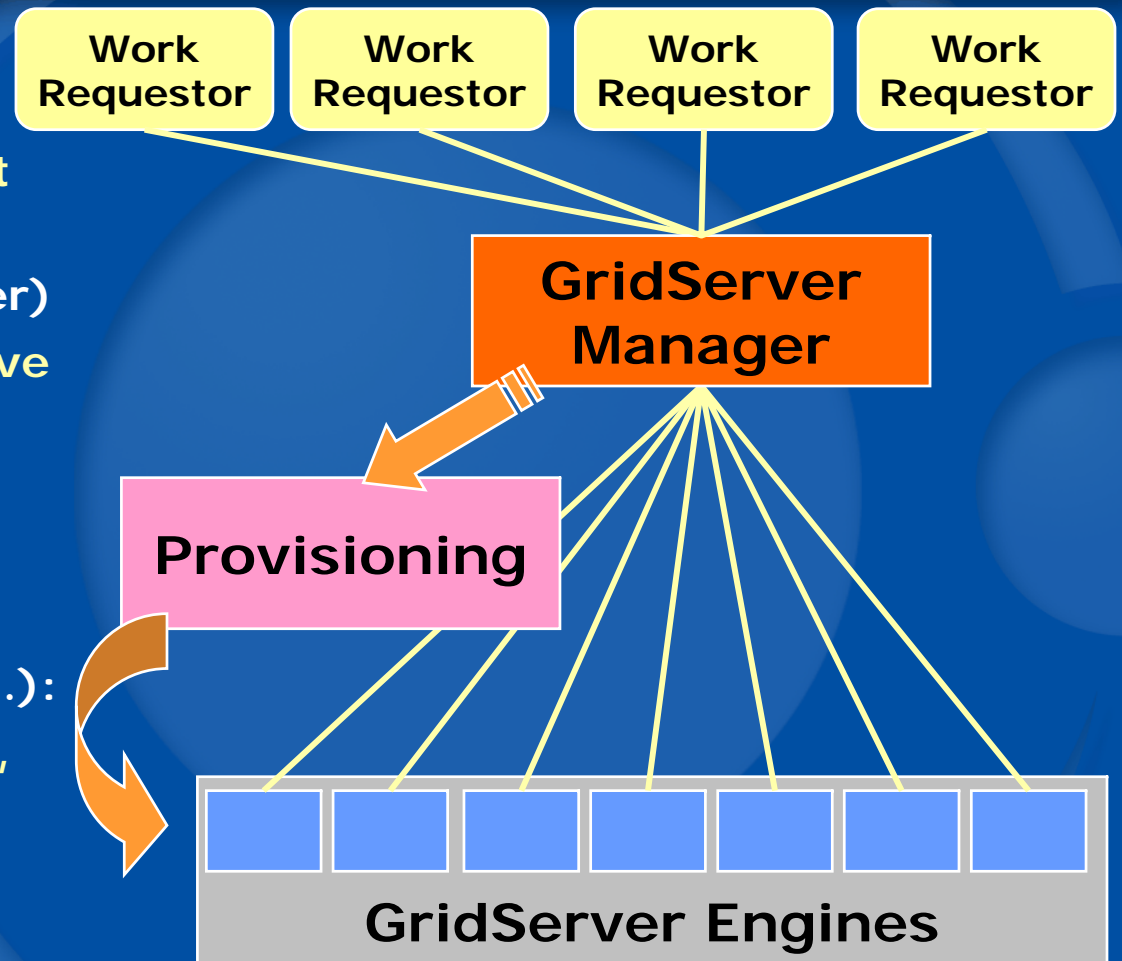
Blueprint – GridServer™ components can access Enterprise Information Sources using standard methods



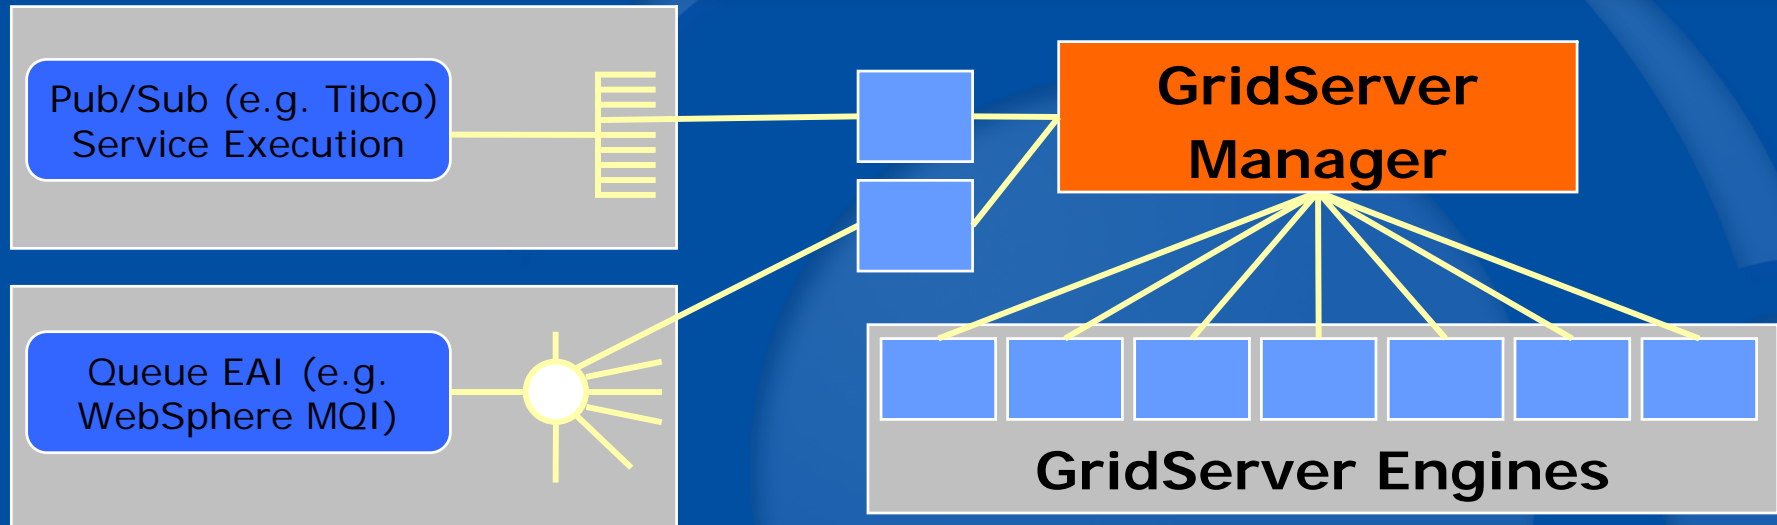
- GridServer Clients, Managers and Engines access EIS through native access mechanisms
- GridServer provides data distribution mechanisms – file transfer, direct data transfer (DDT) and LiveCache

Blueprint –GridServer™ and Provisioning Resource 'Sourcing' For The Grid

- **Work Requestor**
 - Atomic units of work
 - Service Level Agreement
 - Frequency/rate, Response
- **Resource Manager (GridServer)**
 - Dynamically add / remove resources and services
 - Monitor activity and service level
 - Predict resource requirements
- **Provisioning (Tivoli, Veritas,...):**
 - Monitor (CPU Utilization, Arrival Rates)
 - Maintain Service Level Policies / Agreements
 - Add and configure resources on demand

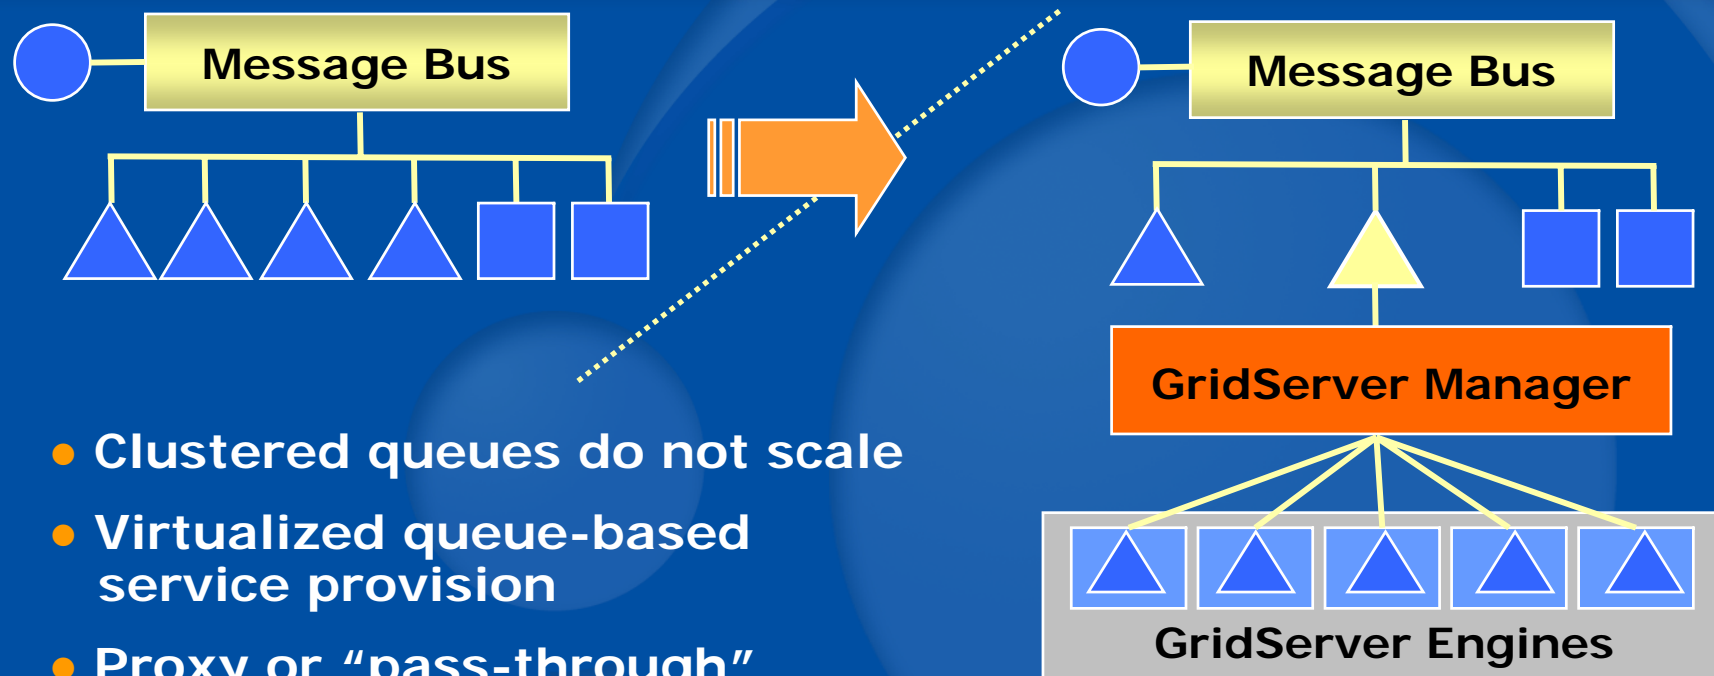


Blueprint – GridServer™ can scale Enterprise Application Integration (EAI) services



- GridServer acts as the virtualized execution fabric
- Business-level abstraction replaces technical API
- Adaptive load balancing
- Autonomous service configuration
- Autonomous service provisioning


Blueprint – GridServer™ scales queue-based services through virtualization



- Clustered queues do not scale
- Virtualized queue-based service provision
- Proxy or “pass-through” service interface
- Adaptive load-balancing:
 - Allows parallelization
 - Increases throughput
- Automated provisioning



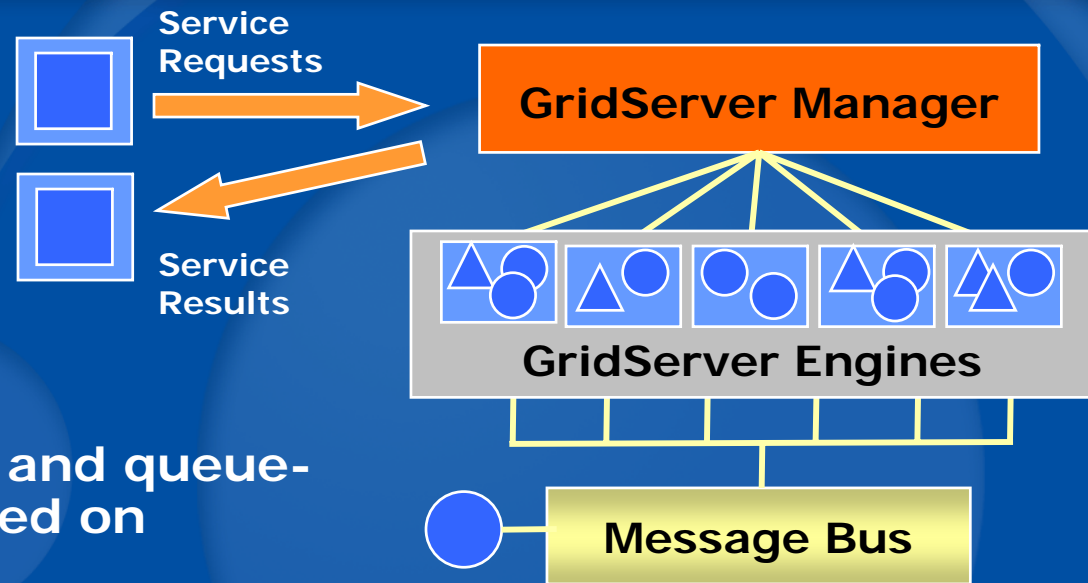
Blueprint – GridServer™ manages the application–service execution fabric



- Application clients serviced through GridServer
- Linearly scaleable virtualized service fabric
- Connect and interoperate:
 - Web Services, SOAP, Java, .Net, C++, COM/DCOM, Excel
- Focus moves to business not technical issues of scale, distribution and resilience




Blueprint – GridServer™ and EAI fabrics unite to provide assured scaleable service execution with wide reach



- Publish / Subscribe and queue-based services hosted on GridServer
- GridServer provides:
 - Application Virtualization
 - Capacity, scale and resilience
- EAI provides:
 - A2A, B2B and B2C Adaptors
 - Transformation , Business Rules

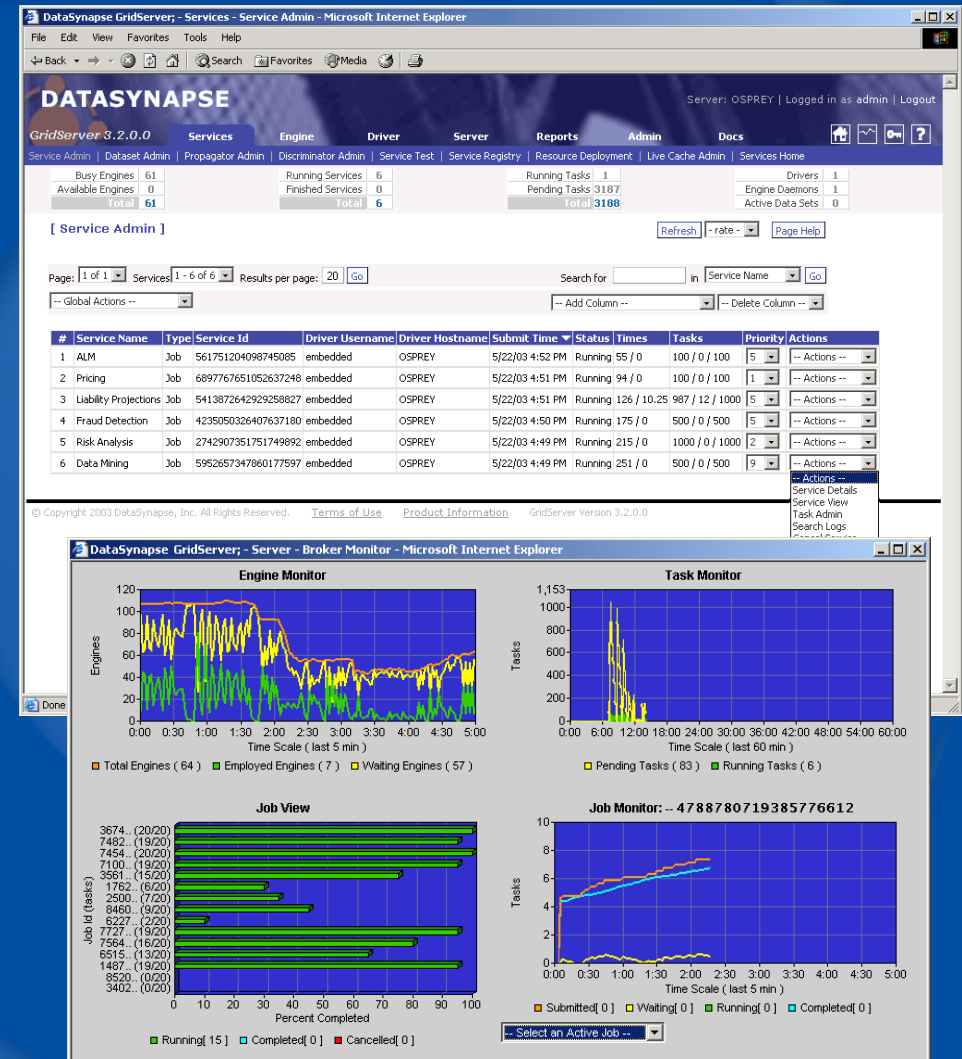




GridServer™ Architecture

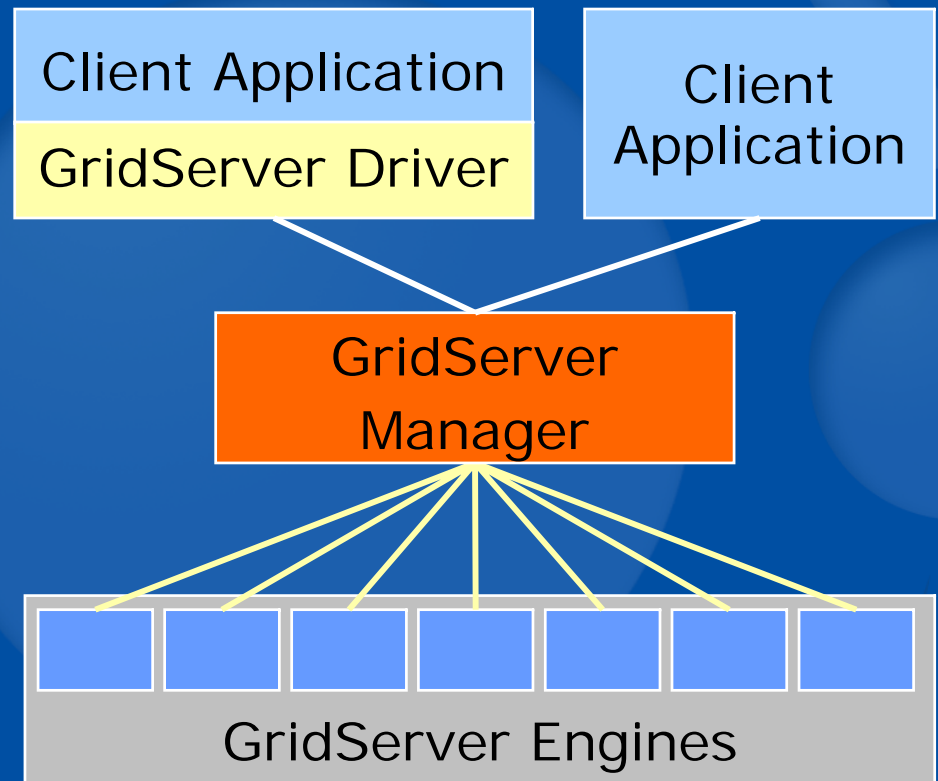
GridServer™ Administration can be accessed from anywhere on the network with web-based interface

- Graphical monitoring
- Engine and Driver installation
- Event-based email notification
- Time-based work or configuration scheduling (batches)
- Security features prevent unauthorized tampering



GridServer™ has three main components

- *Clients interact with GridServer through Drivers...*
 - Batch-Oriented
 - Object-Oriented
 - Service-Oriented
- ...or directly through SOAP
- The *Manager* distributes the service requests to Engines
- *Engines* get the input data and perform the computation
- Platforms:
 - Manager – Java
 - Engines – Windows, Solaris, Linux, pLinux, zLinux
 - Drivers – Java, C, C++, .Net, BatchCOM, SOAP



GridServer™ offers flexible integration strategies

Synchronous & Asynchronous

Clients

Java or
Generated
Proxy

.Net or
Generated
Proxy

C++
API

SOAP

COM
API

C API

Commands
or Scripts

GridServer Manager

Virtualized Invocation

Java
class

.Net
assembly

COM

C++ DLL
or .so

C DLL
or .so

Executables
.exe, .pl, etc.

Hosted Services

Stateless & Stateful

GridServer™ offers flexible integration strategies

- **Service-Oriented**

- Standards-based model
- Thin client
- Promotes re-use
- Language Interoperability

- **Object-Oriented**

- Rich (empowered) client
- Access to cache
- Simple programming models
- Advanced programming models

- **Batch-Oriented**

- Script based
- Parametric parallel
- No re-compiling
- Least integration

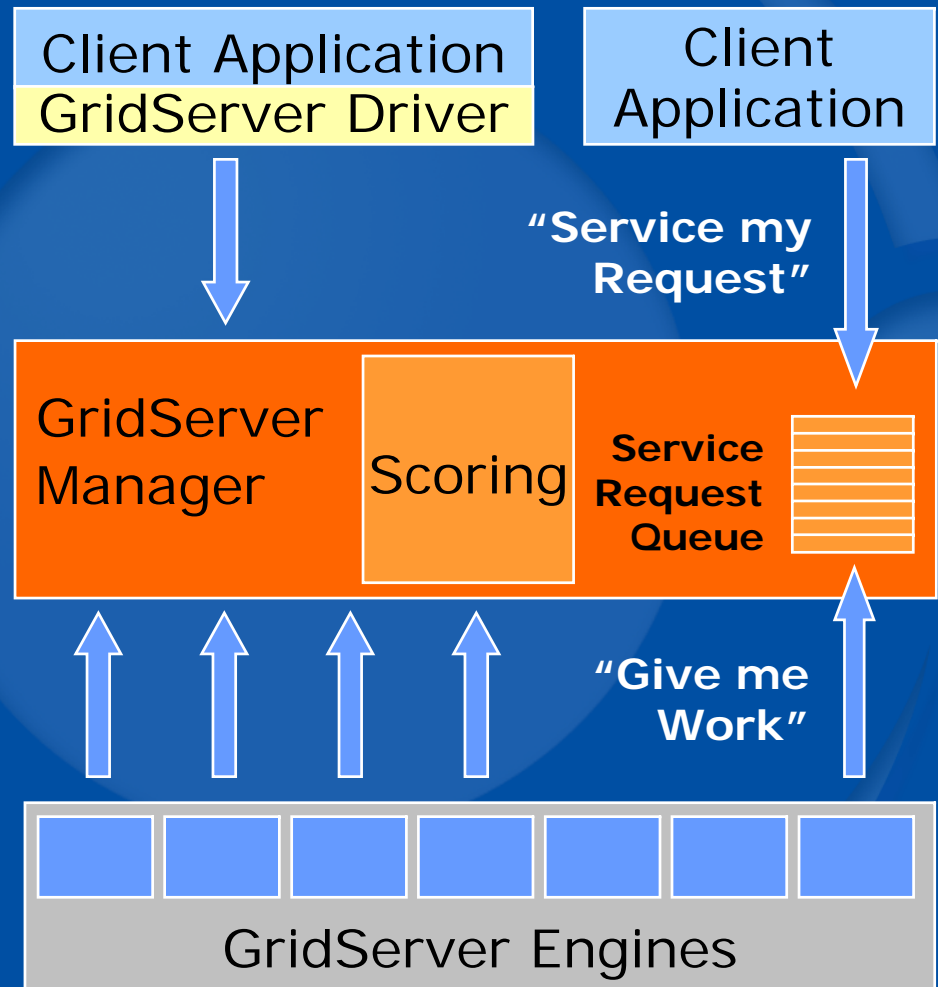
GridServer has sophisticated scheduling capabilities based on an extensible scoring algorithm

- **Scheduling**

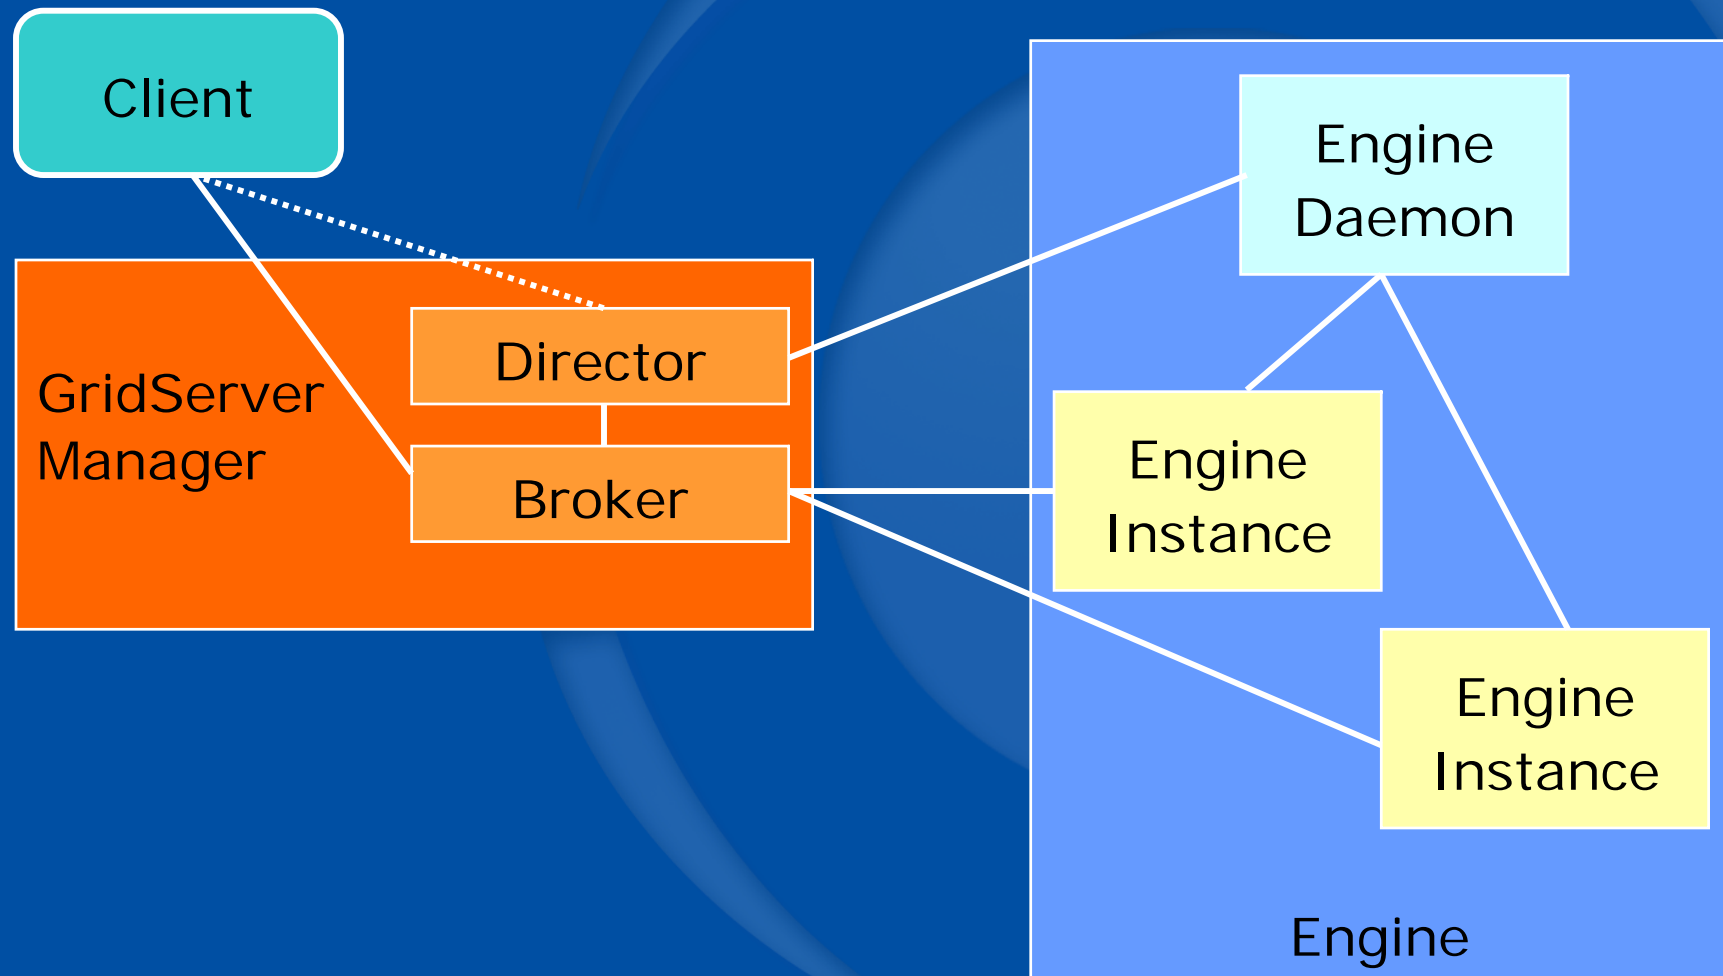
- Scoring
- Affinity

- **Sophisticated Control**


- Priorities
- Urgent Priority Services
- Pre-emption
- Engine Blacklisting
- Task Discriminators
- Redundant Instances (long running tasks)



GridServer™ - Components




As GridServer™ scales, components of GridServer™ Manager are distributed across machines



- Director has little work to do—effectively unlimited scalability
- Brokers can be added to extend capacity
- Lightweight, self-throttling communication enables hundreds of Engines per Broker


GridServer™ can operate in fault tolerant configurations



- **Primary and Secondary Director on different machines**
- **At least two Brokers on different machines -more Brokers as needed**
- **Multiple Engines -if an Engine fails or is interrupted, its current service request will be given to another Engine**

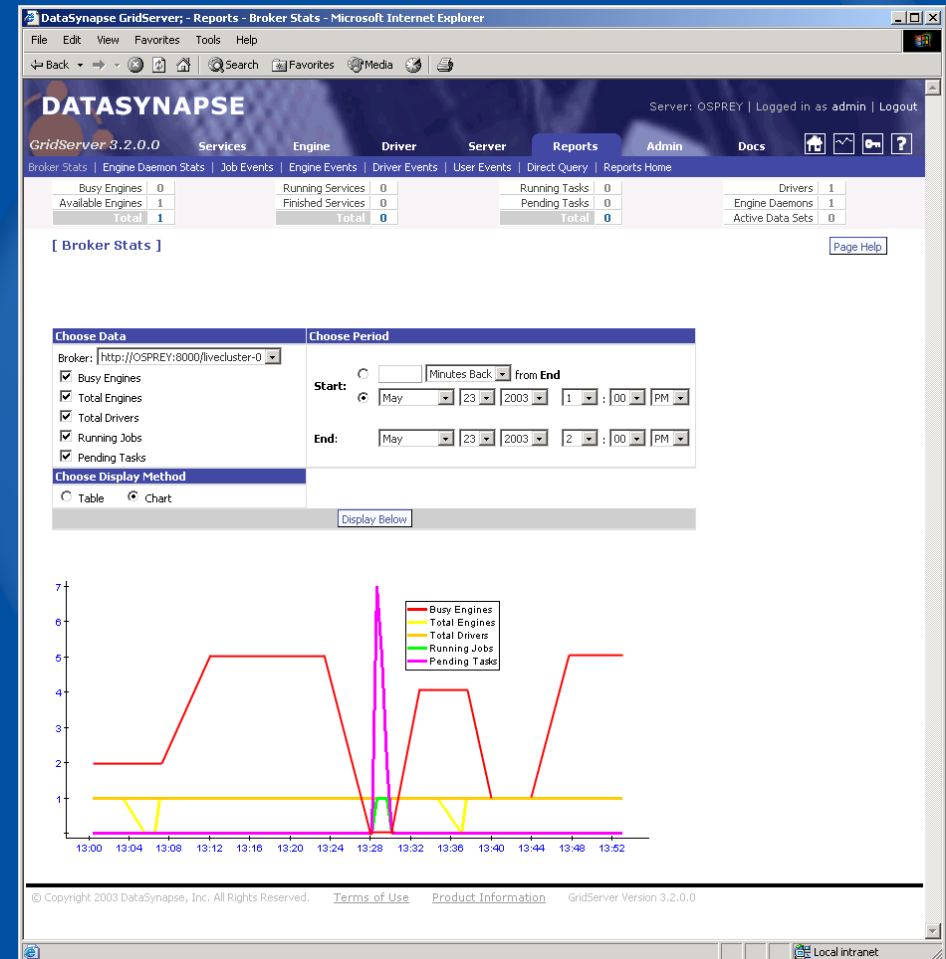
GridServer™ provides multiple mechanisms for data distribution

- File Replication
- Direct Data Transfer
- LiveCache™
 - Aggressively cached by application or engines
 - Simple dictionary interface
 - APIs in all Supported Languages
 - Updates implemented by deletion of cache item and lazy initialization
 - 100's of updates per second
 - Non-transactional




GridServer™ offers ease of administration

- Sophisticated reporting facility for accounting, auditing, performance monitoring, etc.
- Most features also accessible programmatically, via Web Services




Adaptive Service Level Agreements (SLAs) are the basis for achieving optimal resource usage across applications

- establish baseline SLAs by application for:
 - resource consumptive applications
 - business critical applications
- Initially establish simple SLAs then refine and optimize
- Realize SLAs that vary over time matching in-day, ad-hoc and batch characteristics of each application
- Monitor and manage SLA fulfilment across a common virtualized application infrastructure
 - Optimizes underlying resource usage
 - Supports autonomous SLA fulfilment and balancing
 - Enables automated provisioning

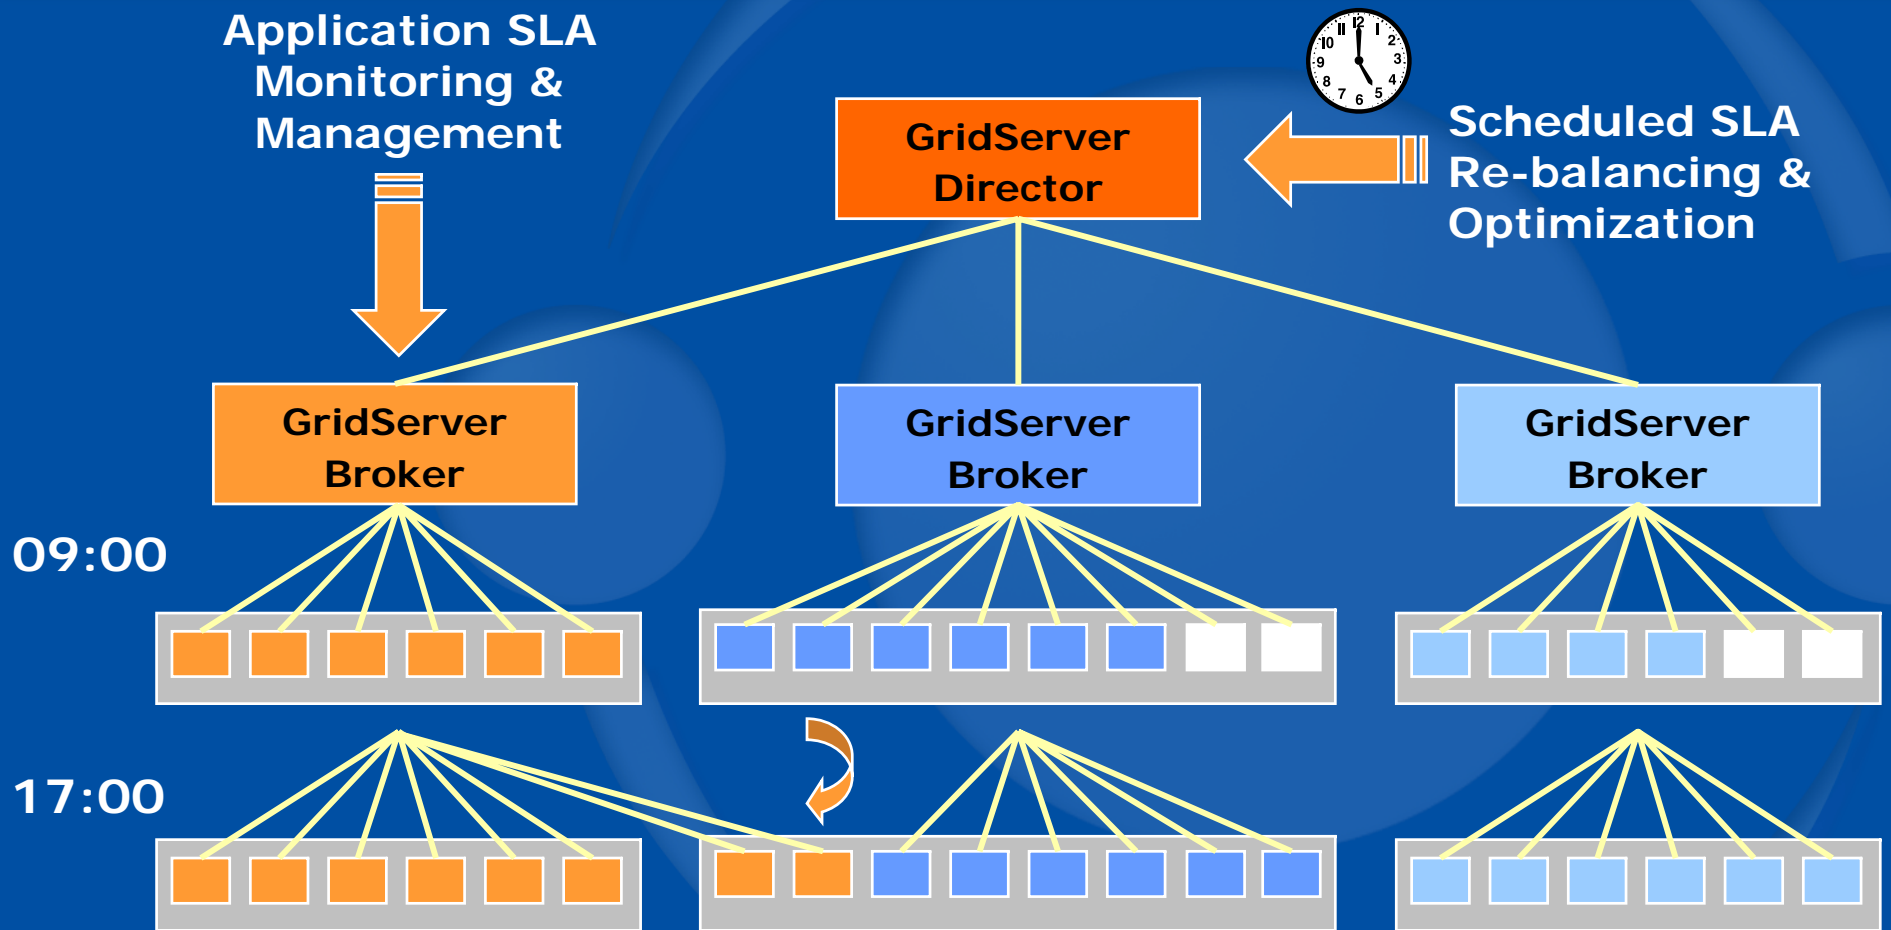


GridServer™ supports Adaptive SLAs that vary across time according to the characteristics of each application

Compute Resource



GridServer™ schedules and re-balances SLAs to optimize resource usage



Enterprise Virtualization Stack – Application & Data Services

