



# Data and Compute Platforms for Emerging Applications

Hamid Pirahesh,

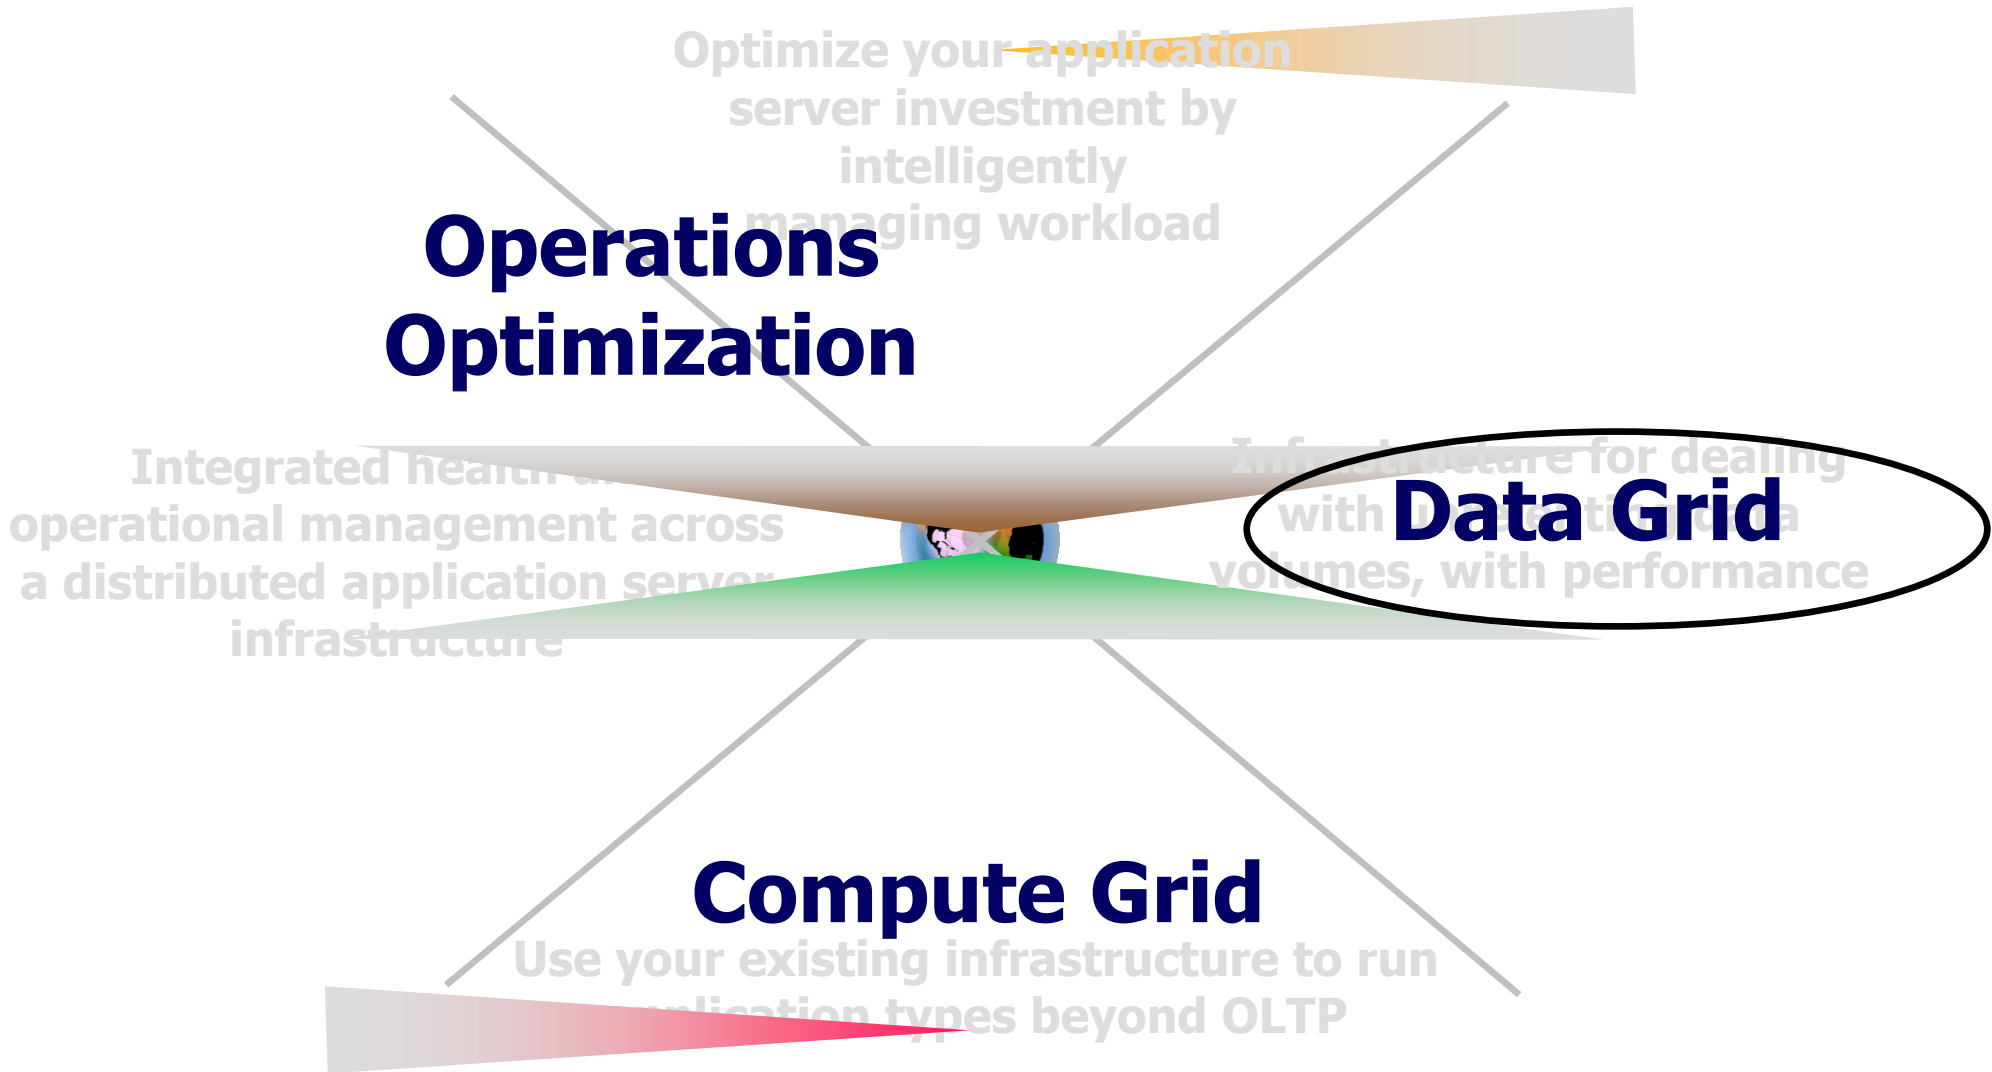
Bo Shekita

IBM

Almaden Research Center

HPTS 2007

# Compute Grid Needs a DBMS: Data Grid



# What is ObjectGrid?

*A flexible framework for realizing high performance, scalable and data-intensive applications*



New York




San Francisco



London



Shanghai



It can be used as a **very powerful cache** that scales from **simple** in-process topologies to **powerful** distributed topologies.

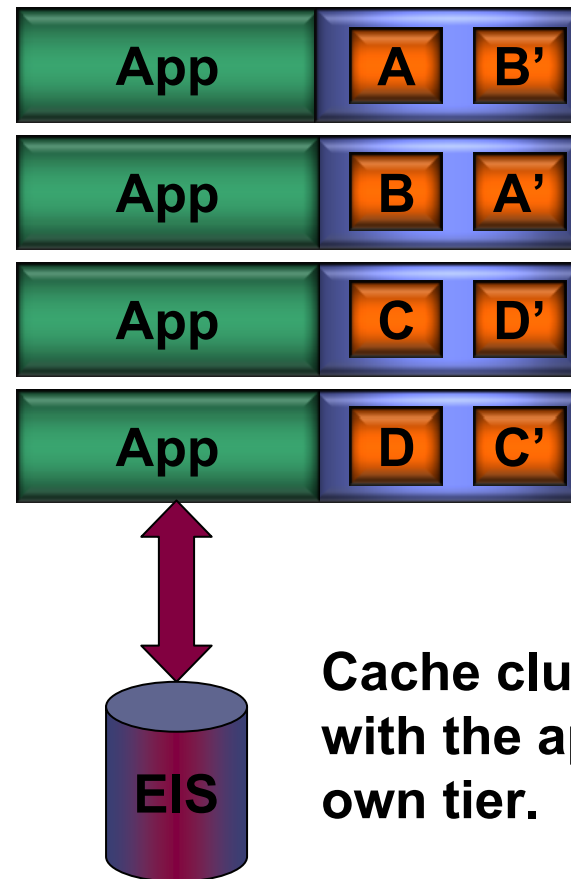
It can be used as a **form of in memory database** to manage application state (and it scales to **1000's of servers**). This is sometimes referred to as **Distributed Application State Management**.

It can be used as a **platform for building powerful Data Grid applications**.



# Distributed ObjectGrid based Cache Operation

- Cluster Coherent cache
- Cache capacity determined by cluster size, not individual JVM Size
- No invalidation chatter
- Cache request handling handled by entire cluster and is linearly scalable
- Load on EIS is lower
- No cold start EIS spikes
- Predictable performance as load increases
- Cached data can be stored redundantly

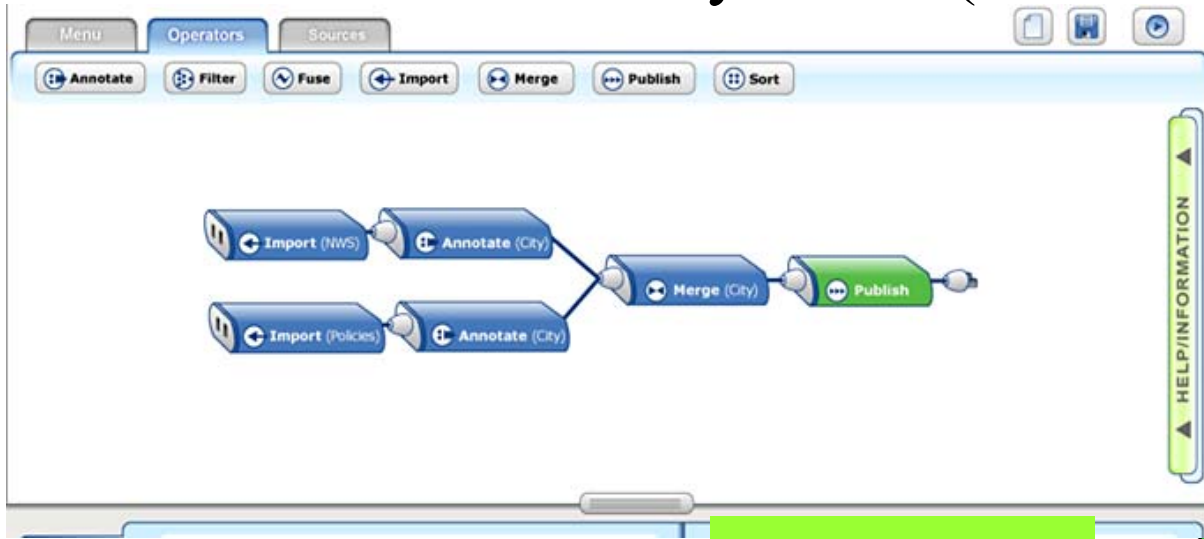


Cache cluster can be co-located with the application or run in it's own tier.

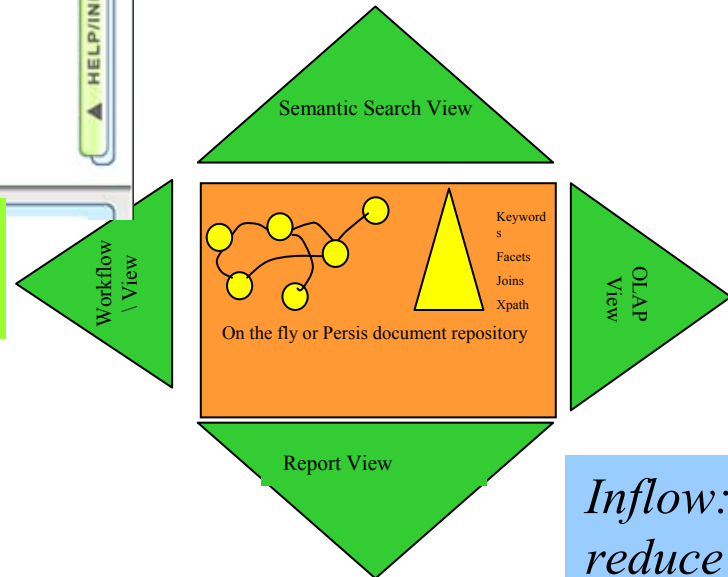
# Data Grid APIs

- These Agents can be invoked by a client using the AgentManager
  - for a single key
  - a group of keys,
  - a partition
  - a whole grid.
- ObjectGrid automatically exploits what ever parallelism is possible for the operation.

# What about analytics on (Business) Documents?

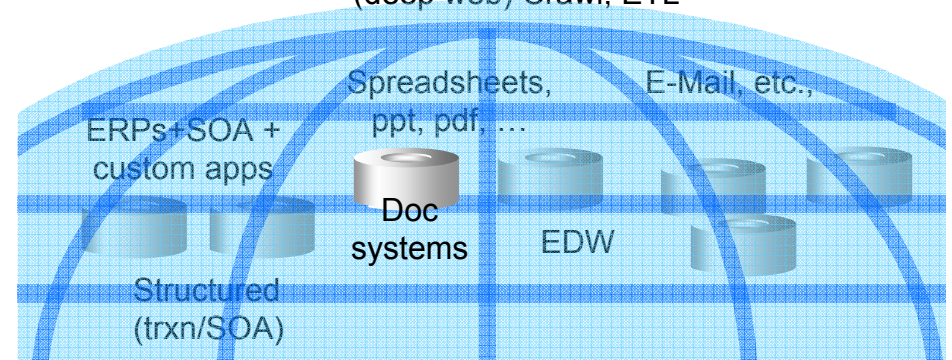


*OutFlow: map reduce platform*

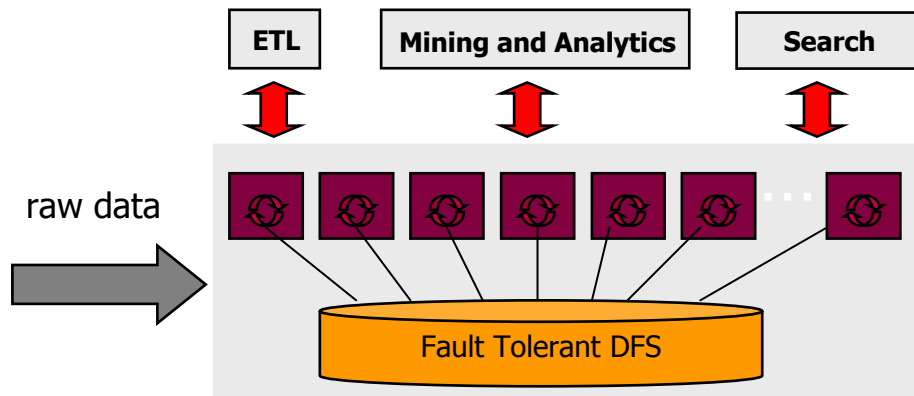


*Inflow: map reduce platform*

Analyze, Integrate  
(deep web) Crawl, ETL



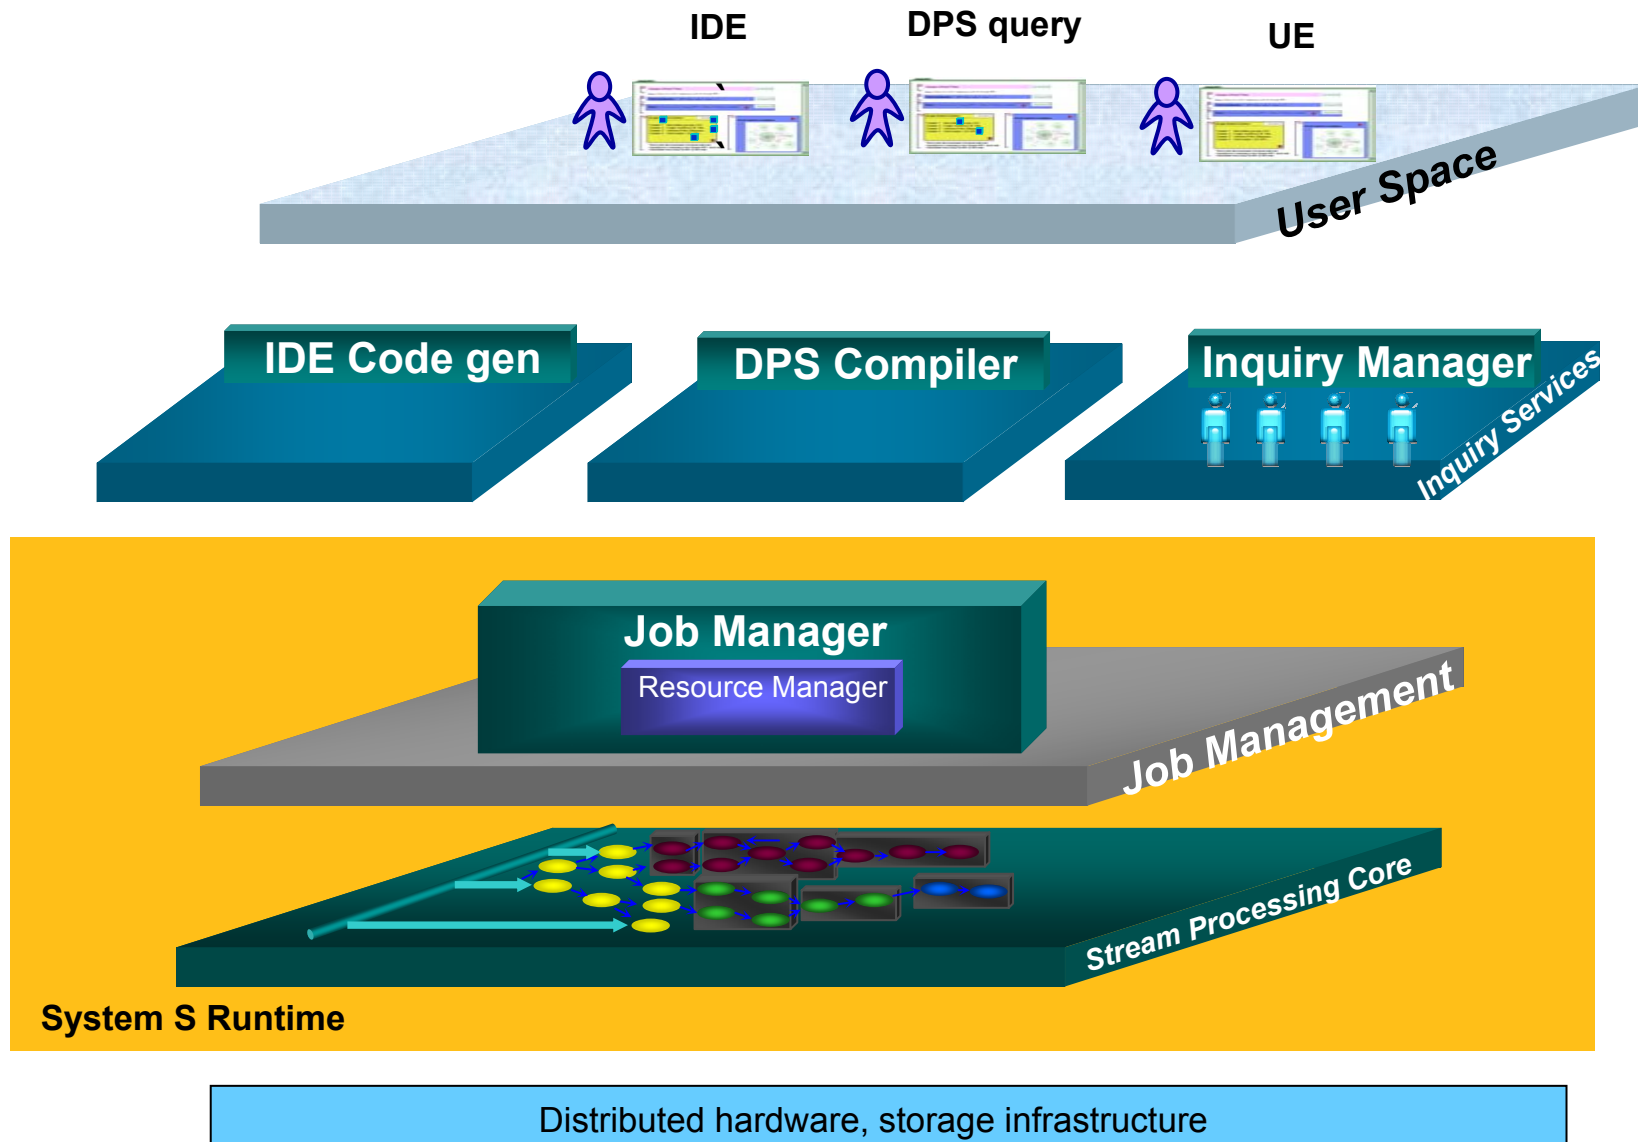
# impliance-- Exploratory Project in Cloud Computing at IBM Research



- Massive clusters of commodity servers connected by high-bandwidth Ethernet
  - **Flexibility:** semi-structured data with chaotic schemas
  - **Modern Hardware:** large memories, multi-core CPUs, commodity clusters
  - **Scale** to 1000s of nodes
  - **Super-low TCO and Simple Design**
  - **Leverage Hadoop's DFS, MapReduce, and BigTable equivalent (Hbase)**
  - **Focus on building new DBMS layers above Hbase core**
  - **SW is built with failures in mind (no gold plated HW)**

# System S at IBM Research

## High-level Architecture



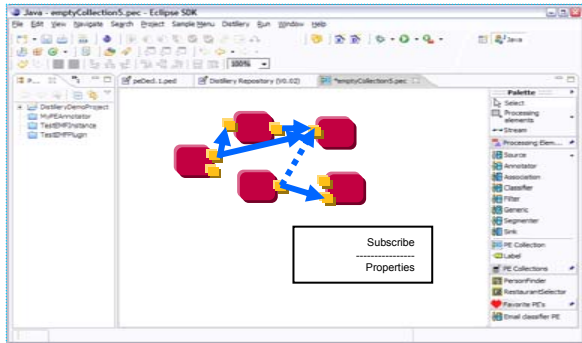
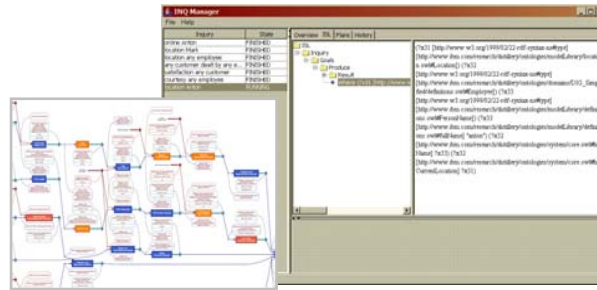




# System S-- Big Picture


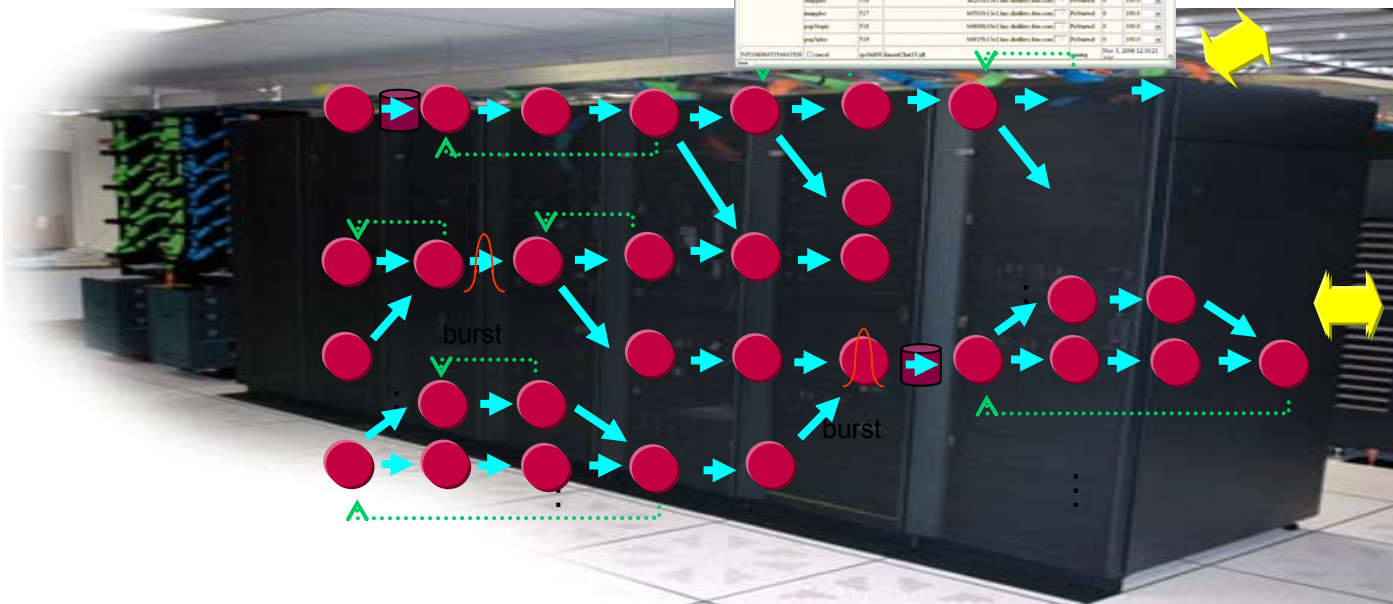
## Inquiry Services: Automated, goal-based application composition

IDE: High productivity integrated tool set  
Eclipse, OSGi, XML, Java, C++,



System orchestration and resource management

Job ID	Artifact	Event ID	JOB Name	State	Submit Date
2495894854627814	Classical	isp48495	ServiceJobTypeJob	Waiting	Thu, 9, 2006 12:18:12 AM
	PE_M	PE_LM	Node_M	None	None
	AppC	AppL	W001911411 Job-Definition-Process	Predefined	30
	Personnel	Person	W001911411 Job-Definition-Process	Predefined	30
77427579111790149	Classical	isp48495	ServiceJobTypeJob	Waiting	Thu, 9, 2006 12:18:26 AM
	PE_M	PE_LM	Node_M	None	None
	AppC	AppL	W001911411 Job-Definition-Process	Predefined	30
	Personnel	Person	W001911411 Job-Definition-Process	Predefined	30
	AppC	AppL	W001911411 Job-Definition-Process	Predefined	30
	Personnel	Person	W001911411 Job-Definition-Process	Predefined	30
	AppC	AppL	W001911411 Job-Definition-Process	Predefined	30
	Personnel	Person	W001911411 Job-Definition-Process	Predefined	30
744128844714642108	Classical	isp48495	ServiceJobTypeJob	Waiting	Thu, 9, 2006 12:18:12 AM



⋮