




ETH



Peer-to-Peer Process Execution with OSIRIS

Christoph Schuler*, Roger Weber*, Heiko Schuldt#, Hans-J. Schek*


1st International Conference on Service Oriented Computing,
December 2003, Trento, Italy

***) Database Research Group, Swiss Federal Instit. of Technology ETH, Zurich, Switzerland**

#) University for Health Informatics and Technology UMIT, Innsbruck, Austria


- **Internet-scale Web Service Composition**
 - Composing Web services
 - Services everywhere
- **OSIRIS Architecture**
 - Peer-to-Peer Process Execution
 - Metadata Replication
- **Performance Evaluation**
 - Scalability Characteristics
- **Conclusion and Outlook**

Web Service Composition





- “Mega programming” or “Programming in the large” combine existing service to value added services
 - Recursively reusing services
- Processes allows for visual programming .
- What infrastructure to define and execute processes?
 - BPEL4WS Modeling Tool
 - State-of-the-art process management system

Scale to internet size




- Service composition on **Internet-scale**
- **Number of services** increases dramatically
 - **On-Demand Computing**
 - **GRID Initiatives**
- **Continues fluctuations** on system configuration
- **What system architecture can manage this situation ?**
 - **Classical process management system?**
 - **In 10 years?**

Existing Concepts



- **Process Management**
 - Programming in the Large
 - Visual Programming
 - Reuse of existing services
 - Composite Services
- **GRID Computing**
 - Self adaptation
 - On demand computing
 - Resource Management
 - Publish-and-Subscribe
 - Load Balancing
 - Service Bus

Existing Concepts




- **Peer-to-Peer**
 - Direct communication
 - Scalability
 - Spans large scale communities

A close-up photograph of a document, possibly a ledger or database output, showing several rows of numerical data. The values are positive and include decimal points, such as +2.688, +5.000, +1.500, +1.125, and +1.062.


- **Database Technology**
 - Atomicity
 - Consistency
 - Isolation
 - Durability

OSIRIS Idea

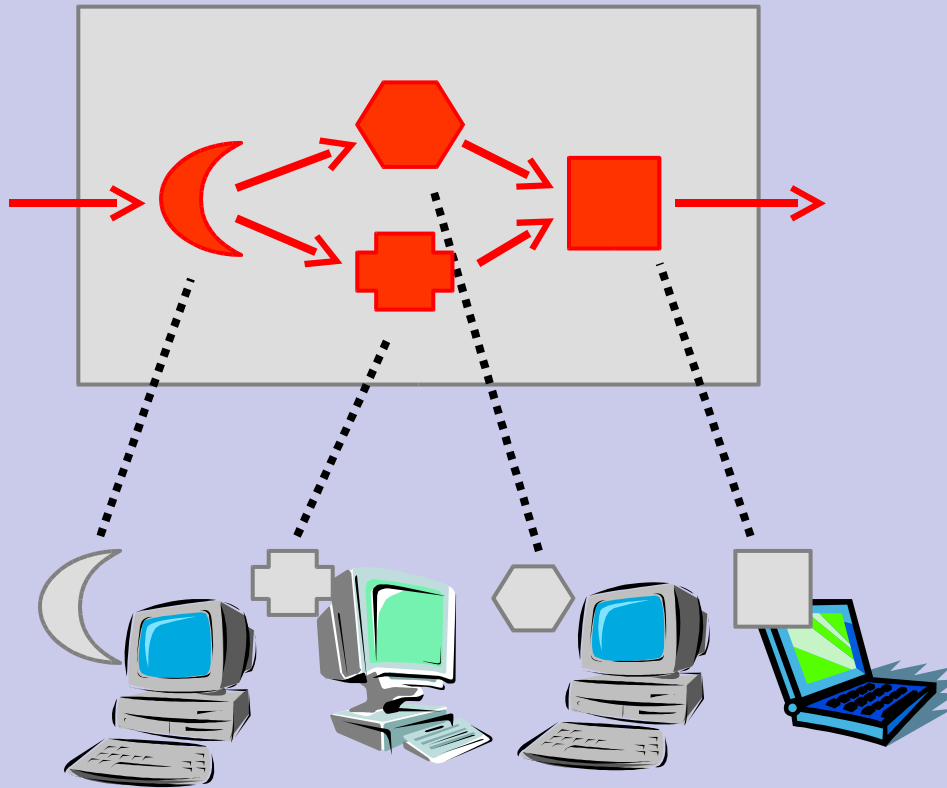


- **Processes** to combine services
 - **Peer-to-peer** execution of processes
 - **Support for dynamic system configuration** (GRID)
 - **Resource and Metadata Management** (GRID)
 - **Providing Execution Guarantees** as known from databases
- **OSIRIS** = **Open Service Infrastructure for Reliable and Integrated Process Support**

OSIRIS Architecture

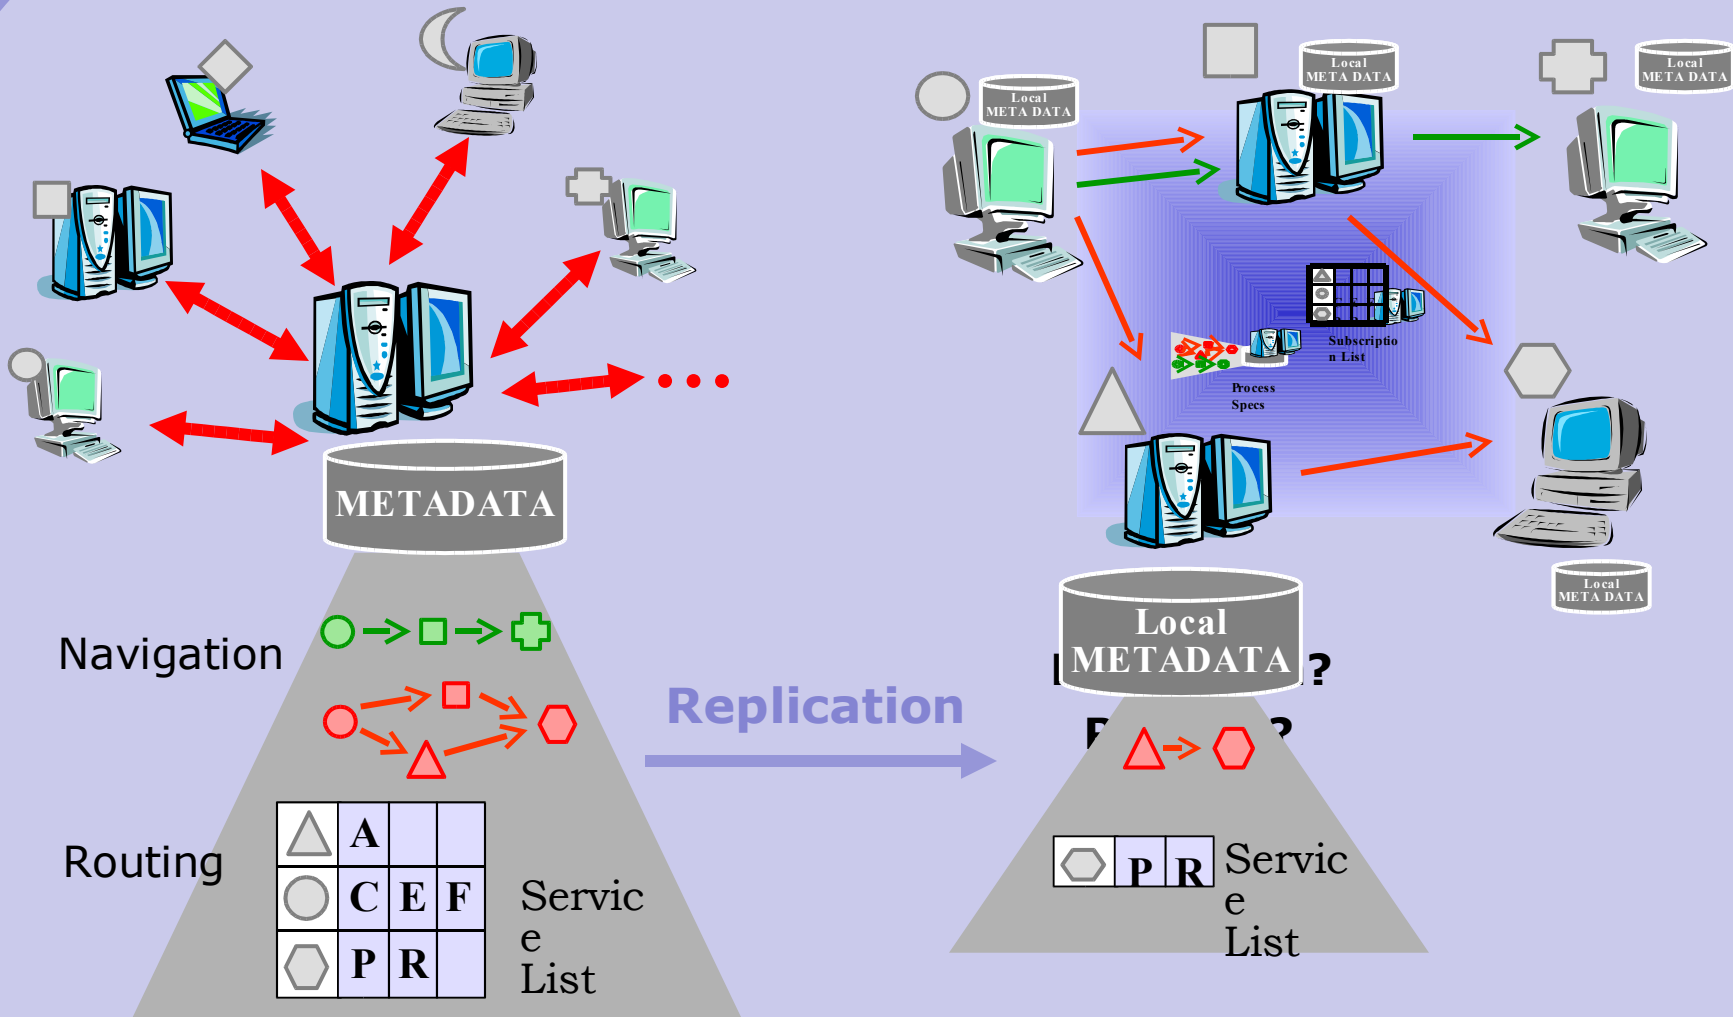


Late Service Binding - Routing

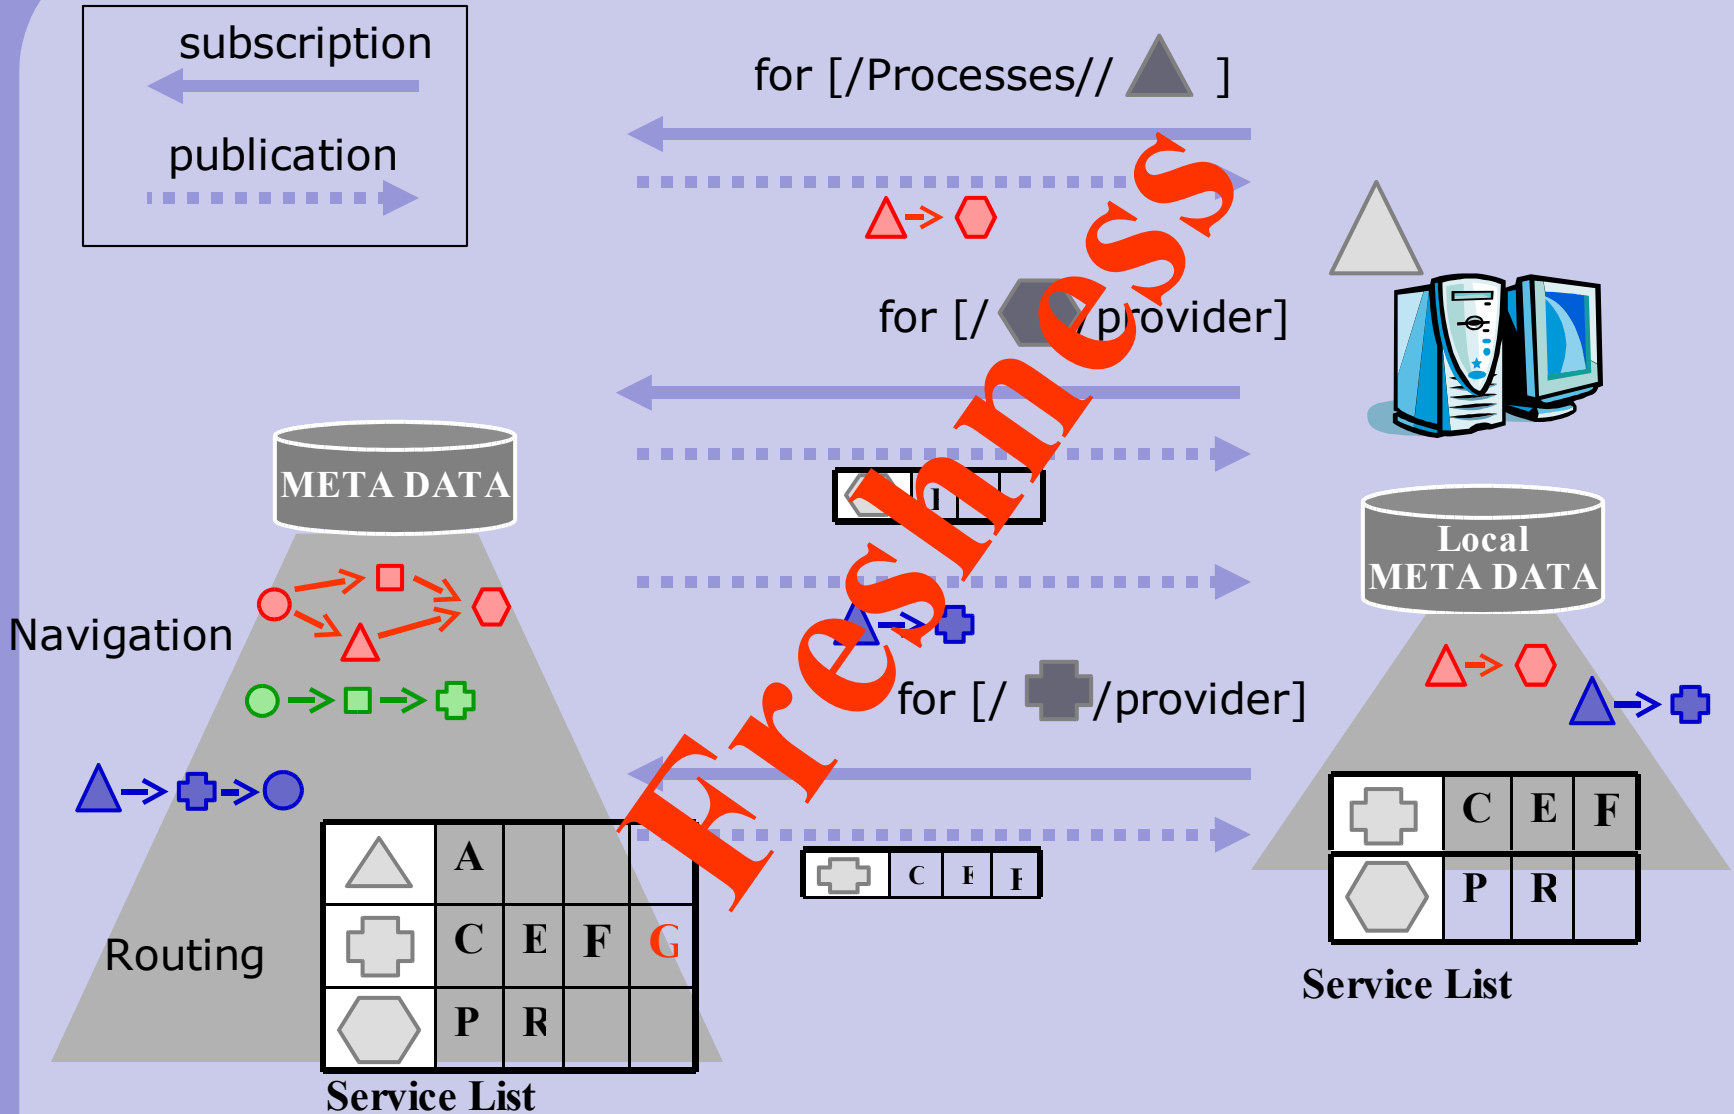


- **More than one** provider of a service type.
- **Process definition**
 - No concrete binding
 - Service Type
 - Semantic Specification
- **Publish-to-ONE** routes process instance to subsequent step
 - Cluster Subscriber
 - “Service Bus”
- **OSIRIS** implements Publish-to-ONE without
 - Central hub
 - Broadcast

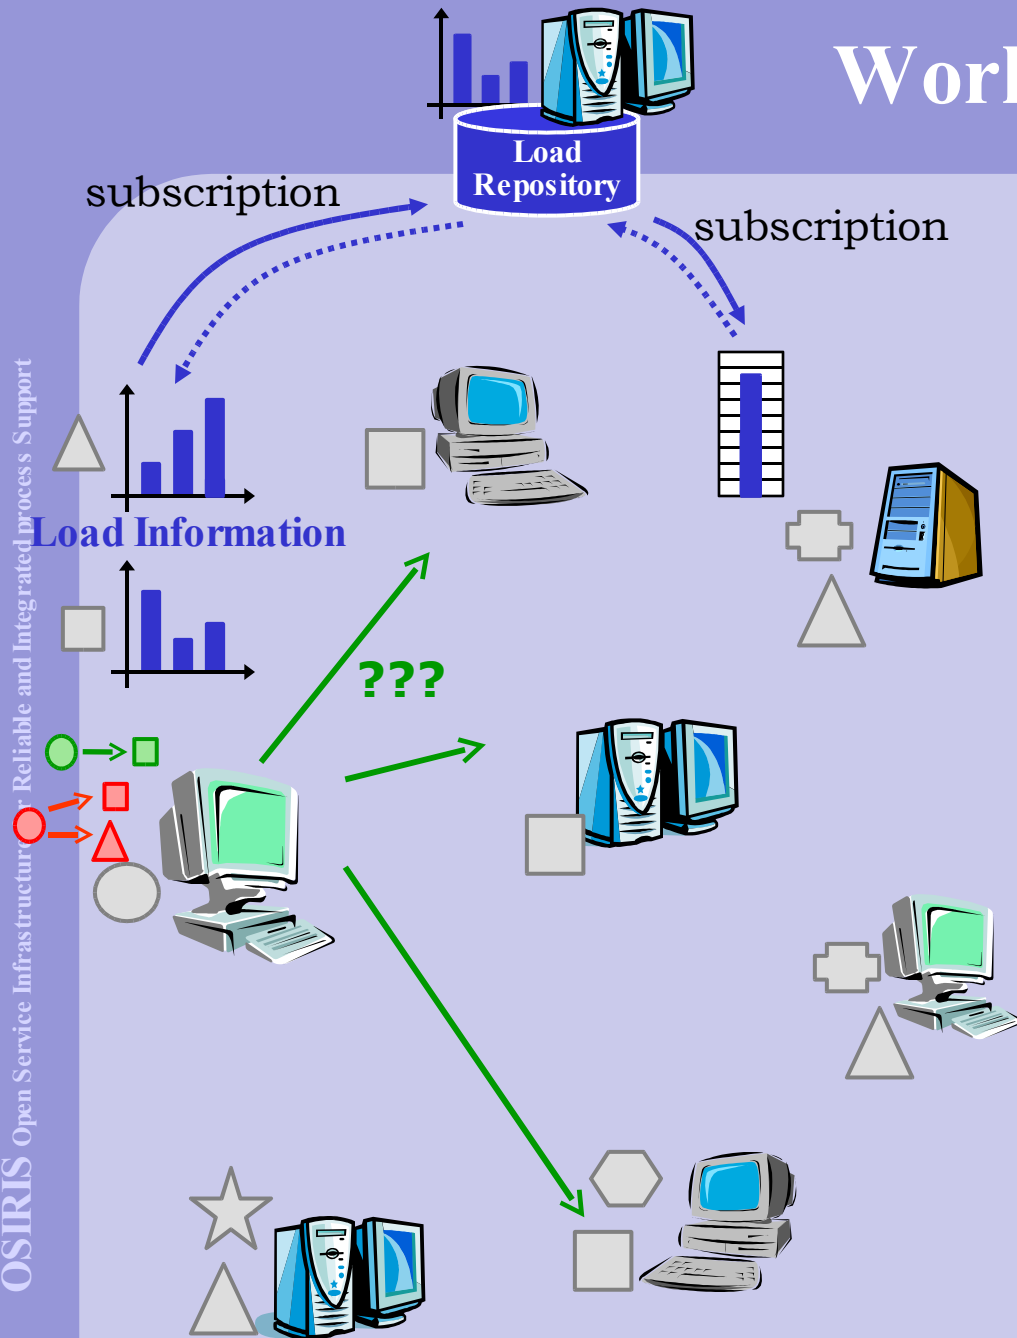
Metadata Replication



Replication Example: Process Data

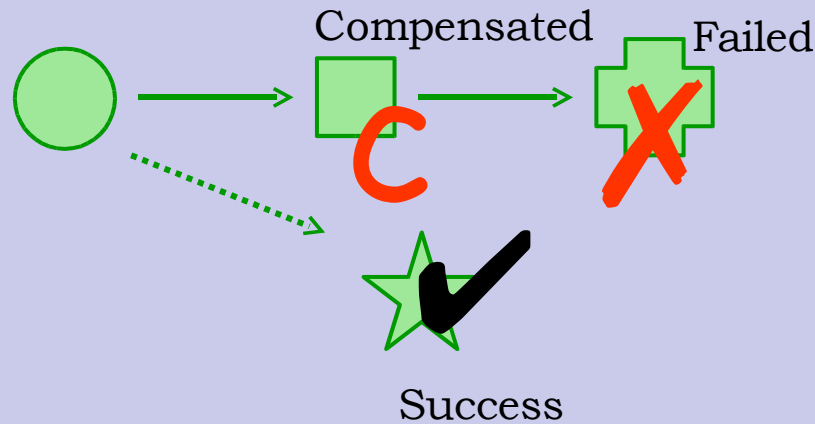


Workload Load Balancing



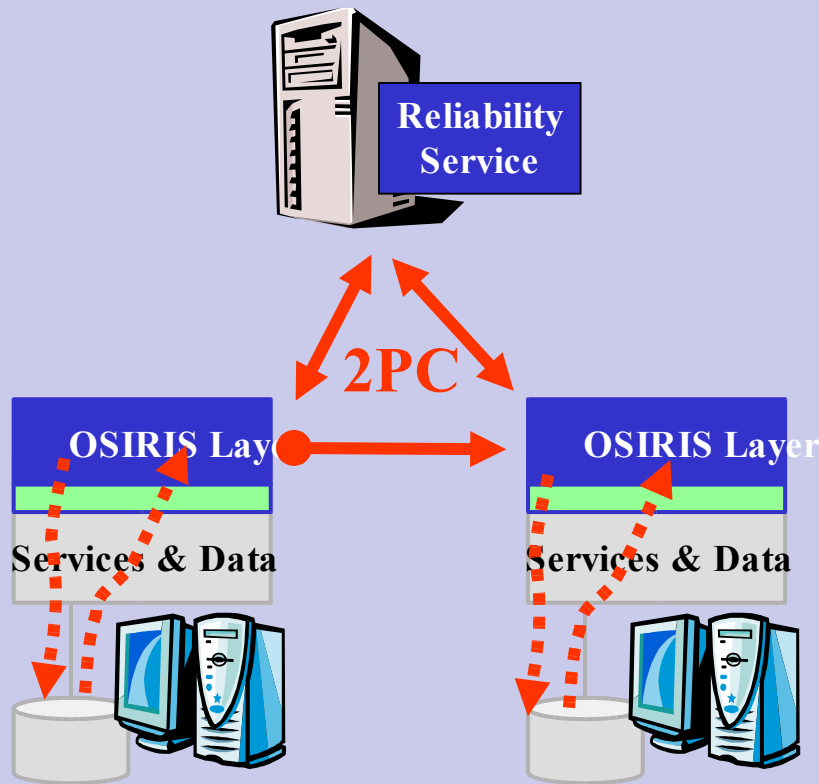
- Workload balancing allows for selecting optimal service provider
- Exploit workload information to route process instance.
- Workload information must be locally available.
- Global **load repository** subscribes for significant changes of provider workload
- And **publishes** global changes to peers

Providing “Database like” Guarantees



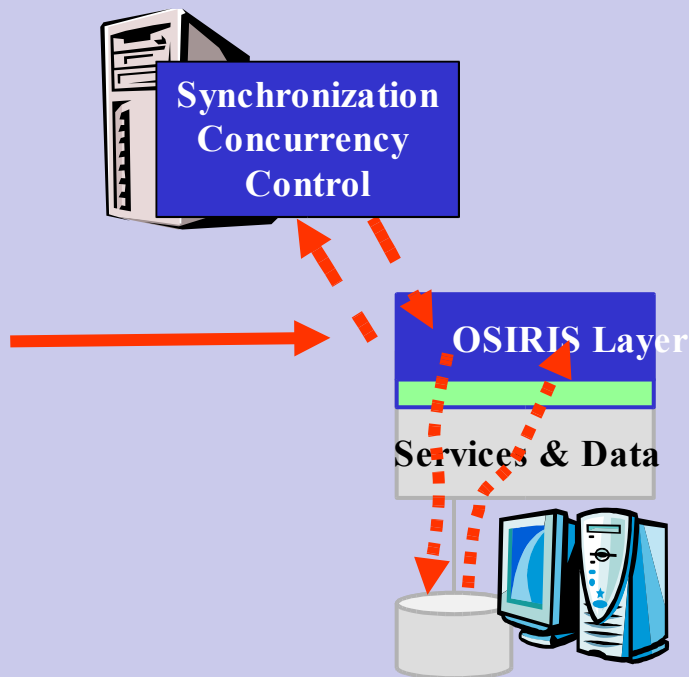
- **Transactional Processes**
 - **Guaranteed Termination**
 - **compensation and alternative execution paths**
- Persistent process routing
 - 2PC Protocol from Peer to Peer
 - Reliability Service
- Concurrency Control on Service level
 - Intercepting Service Call
 - Global Concurrency Control Service

Providing “Database like” Guarantees



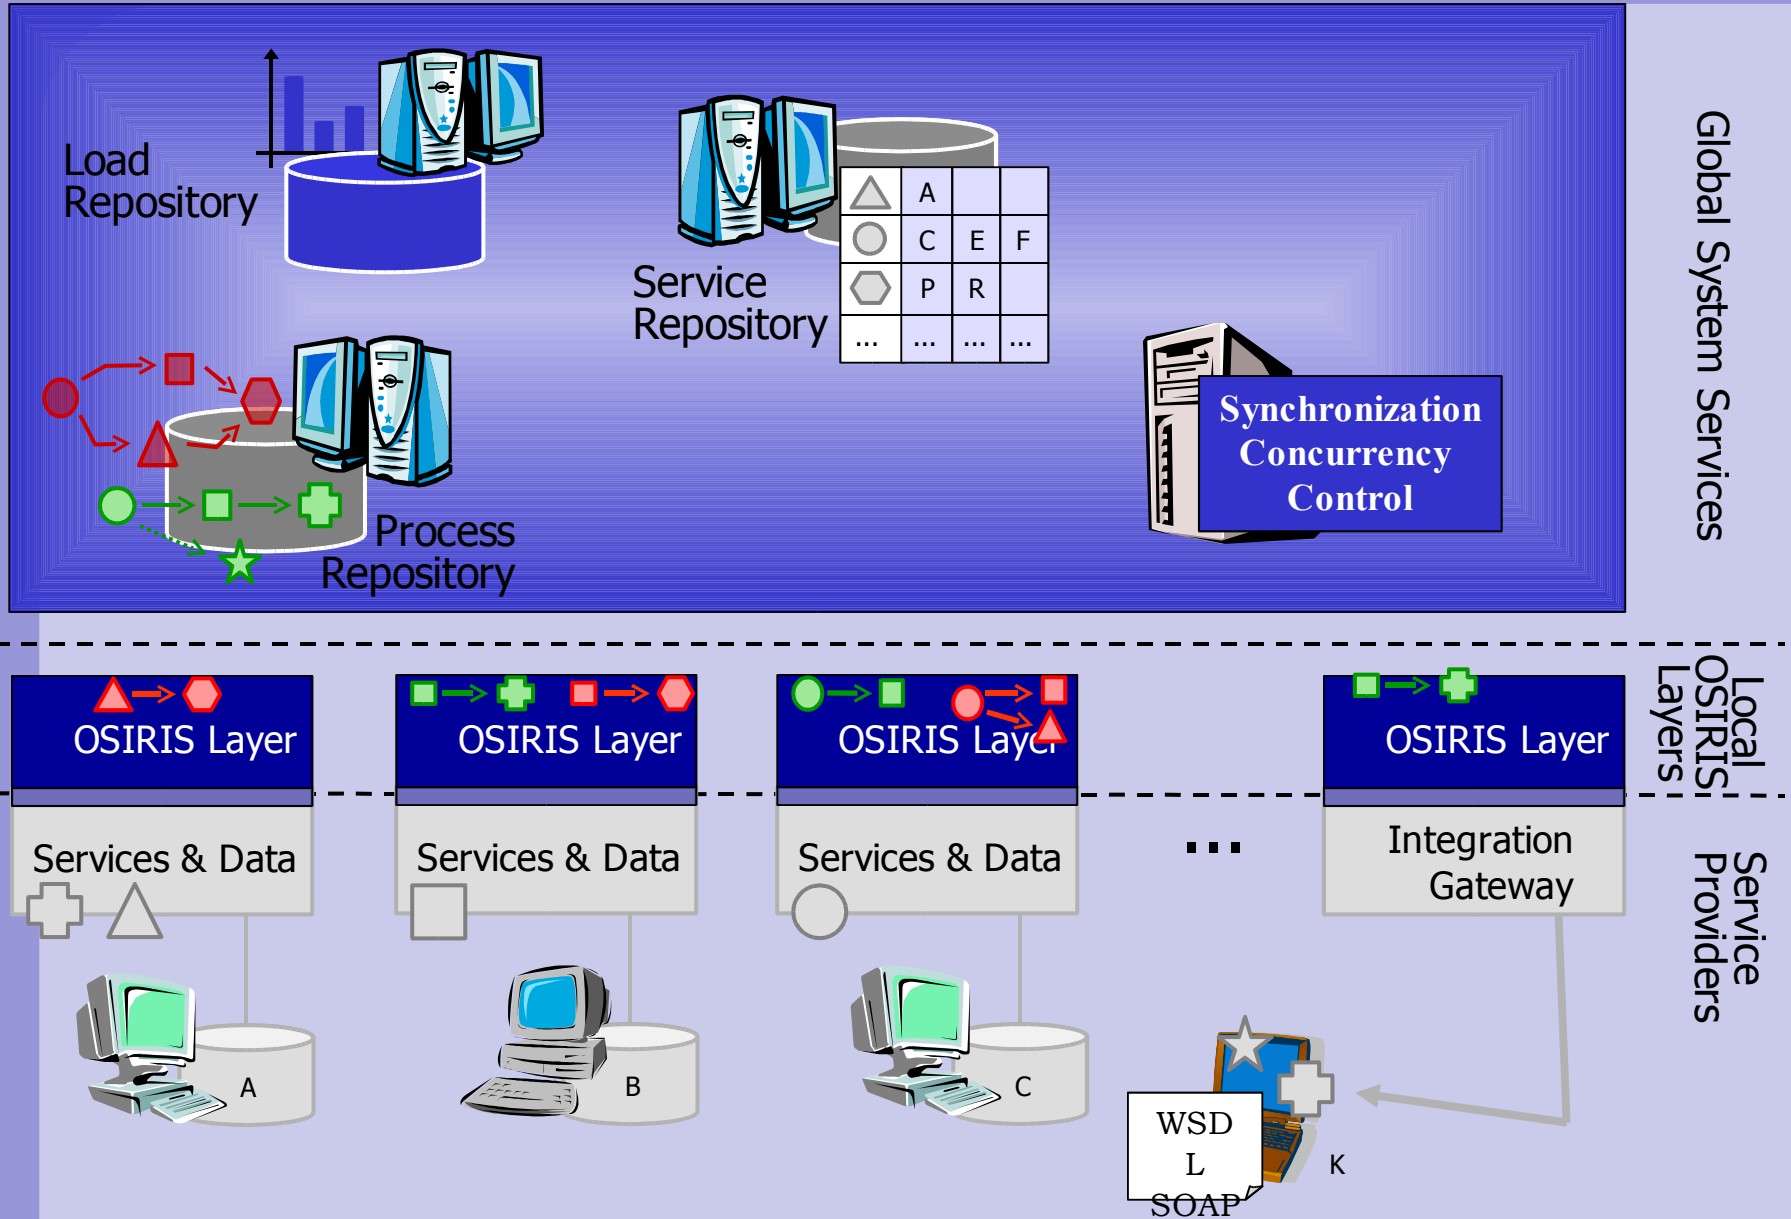
- Transactional Processes
 - Guaranteed Termination
 - compensation and alternative execution paths
- **Persistent process routing**
 - **2PC Protocol from Peer to Peer**
 - **Reliability Service**
- Concurrency Control on Service level
 - Intercepting Service Call
 - Global Concurrency Control Service

Providing “Database like” Guarantees



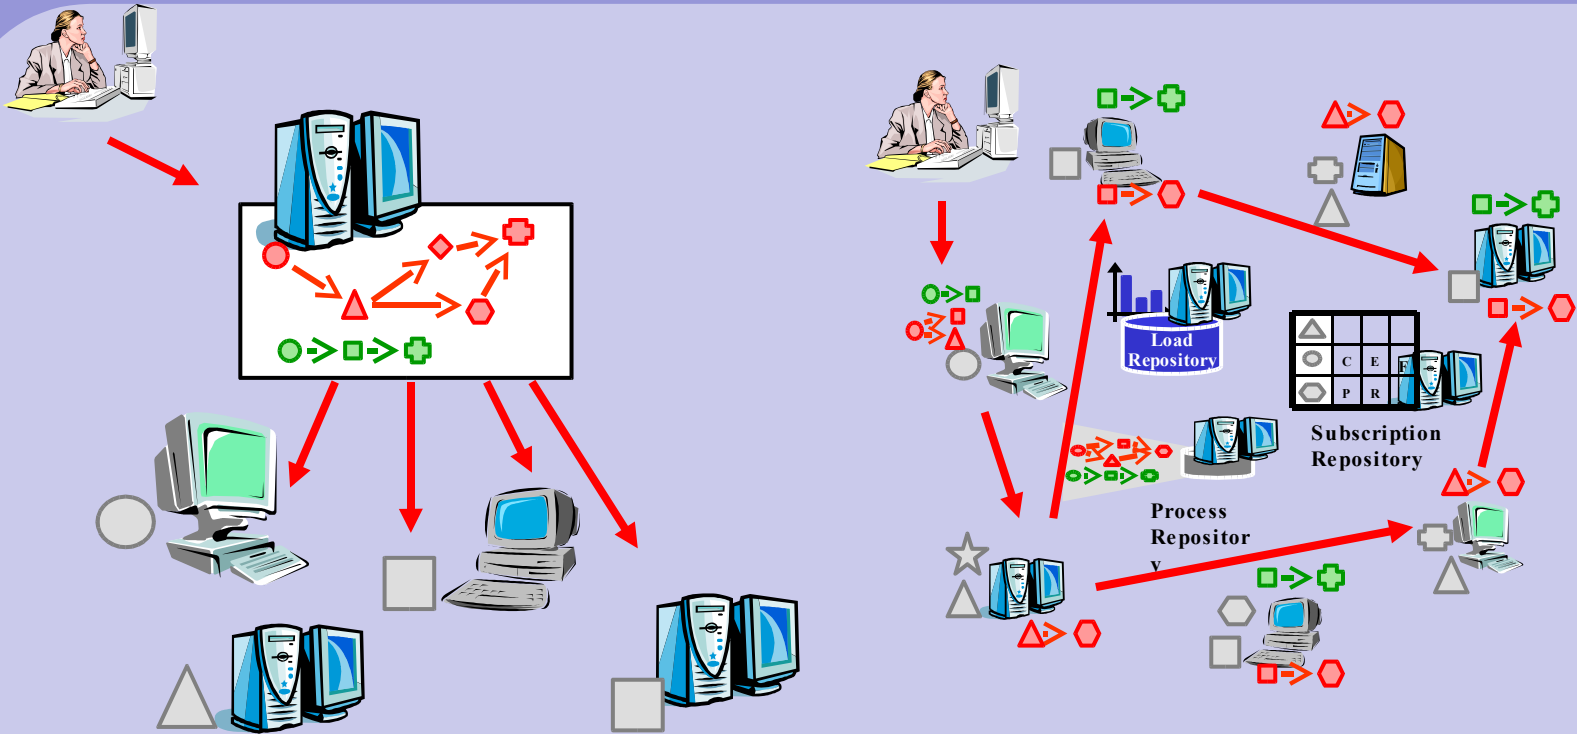
- Transactional Processes
 - Guaranteed Termination
 - compensation and alternative execution paths
- Persistent process routing
 - 2PC Protocol
 - Reliability Service
- **Concurrency Control on Service level**
 - **Intercepting Service Call**
 - **Global Concurrency Control Service**

OSIRIS: The big Picture




- **Internet-scale Web Service Composition**
 - Composing Web services
 - Services everywhere
- **OSIRIS Architecture**
 - Peer-to-Peer Process Execution
 - Metadata Replication
- **Performance Evaluation**
 - Scalability Characteristics
- **Conclusion and Outlook**

Evaluation Setting



- **Centralized vs. Peer-to-Peer process execution**


- One basic process type
- Every activity takes 2s



- Increasing number of providers and processes


Process Throughput – Centralized vs. Peer-to-Peer

Central




- Increasing overall load in the system
- **Centralized Approach:** with 8 nodes, the limit of scalability is reached
- **OSIRIS (P2P):** scales well, already 16 peers sufficient to cover the full load

OSIRIS




Process Response Times – Centralized vs. Peer-to-Peer

Central

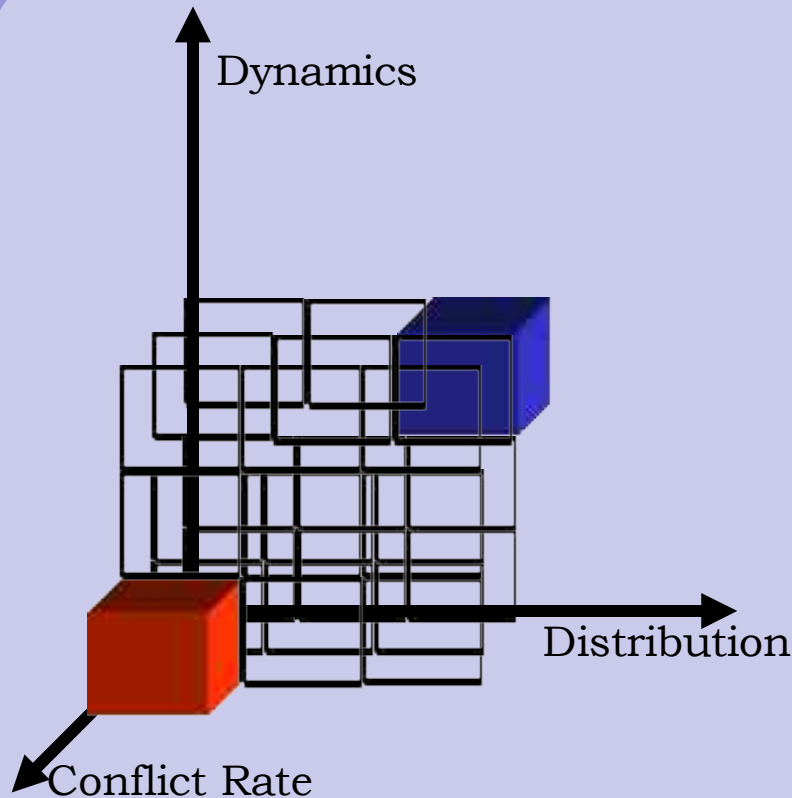


- Same scenario than before
- **Centralized Approach:** coordinator as bottleneck – adding more than 8 nodes will not further decrease avg. response time
- **OSIRIS (P2P):** no bottleneck, almost perfect scalability

OSIRIS



- Dramatically increasing number of services and processes in the internet need a new architecture to high scalability.
- Combining ideas from different areas together with a sophisticated replication mechanism allows for realizing a truly distributed peer-to-peer process management system.
- The OSIRIS infrastructure has a high potential to **scale** with respect to increasing **number of providers** as well as to increasing **number of processes**.
- Applied to Information Space Management of ETHWorld (Virtual Campus of ETH Zurich)
- First evaluation results presented here support the scalability predictions.



- A more complex benchmark setting will show, for what class of application OSIRIS architecture is best suited.
- What concepts can be applied to partially improve existing approaches.
 - Service GRID infrastructure
 - Central process management
- Acknowledgment: Work supported by the IBM Labs Germany, Böblingen with a IBM Hardware SUR Grant

Thank you ...

Questions?