

RISK YÖNETİMİNDE BAŞARI FAKTÖRÜ “İŞ SÜREKLİLİĞİ YÖNETİMİ” SUCCESS IN RISKMANAGEMENT: “BUSINESS CONTINUITY MANAGEMENT”



19-20.09.2009
İstanbul Teknik
Üniversitesi,
Ayazağa Kampüsü,
Süleyman Demirel
Kültür Merkezi

Maslak – İSTANBUL
TÜRKİYE / TURKEY

Enterprise Data Replication and Security

Yılmaz Çankaya, Uzman Araştırmacı
BSG (Information Systems Security), Tübitak UEKAE

Definition of Enterprise Data Replication



Data Replication Options:

Options you may need to consider during a implementation decision process. They can co-exist in replication scenario.

Host-Based, Appliance-Based, Storage-Based ,
Transaction-Aware, Mirroring or Shadowing

Common Usages:

For which usages could we utilize data replication technologies.

Disaster Recovery, Maintenance, Backup

Data Replication Architectures:

There are a few implementation architectures available along with their pros and cons.

Single Datacenter, Single Master, Multi-Master

Data Replication Options



Definition of Enterprise Data Replication

Host-Based:

Host based replication resides on the application server that needs to have its data replicated.

Appliance-Based:

Unlike the host-based solutions, all intelligence needed to perform the replication is housed in an appliance that resides in the I/O path between the host and the storage.

Storage-Based:

Storage-based replication combines the best aspects of host-based and appliance-based solutions

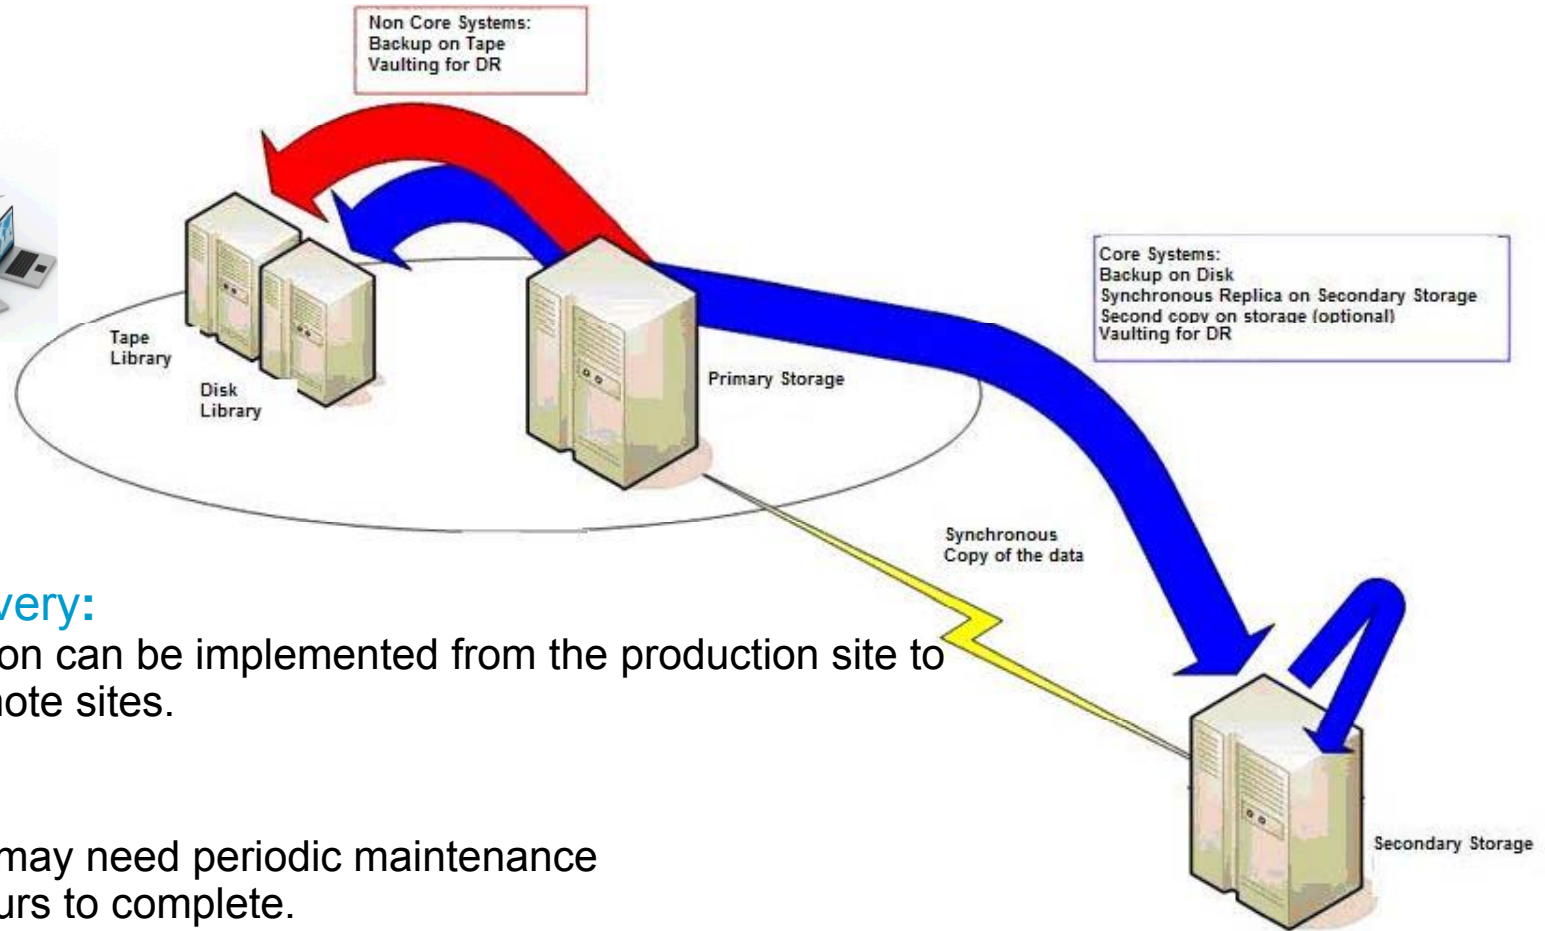
Transaction-Aware:

Transaction-aware replication offers transaction-level replication.

Mirroring or Shadowing:

Shadowing is an asynchronous process, thus requiring less network bandwidth than synchronous mirroring.

Common Usages



Disaster Recovery:

Remote replication can be implemented from the production site to one or more remote sites.

Maintenance:

Various servers may need periodic maintenance that can take hours to complete.

Backup:

Backing up data is frequently the biggest daily challenge for an IT manager. Backup windows have been shrinking while data has been growing.

Data Replication Architectures



Definition of Enterprise Data Replication

Single Datacenter.

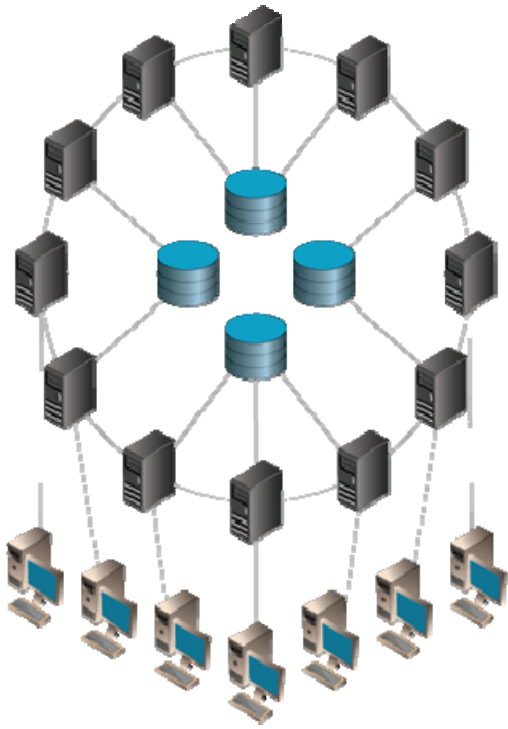
Minimum Availability. No protection in case of natural disasters. Geographical scalability is not an option.

Single Master.

There is a master datacenter that writes go to and other sites replicate to. The replicate sites provide read-only services.

Multi-Master.

All clustered servers are serving reads and writes. All data is consistent. NASDAQ follows this strategy.

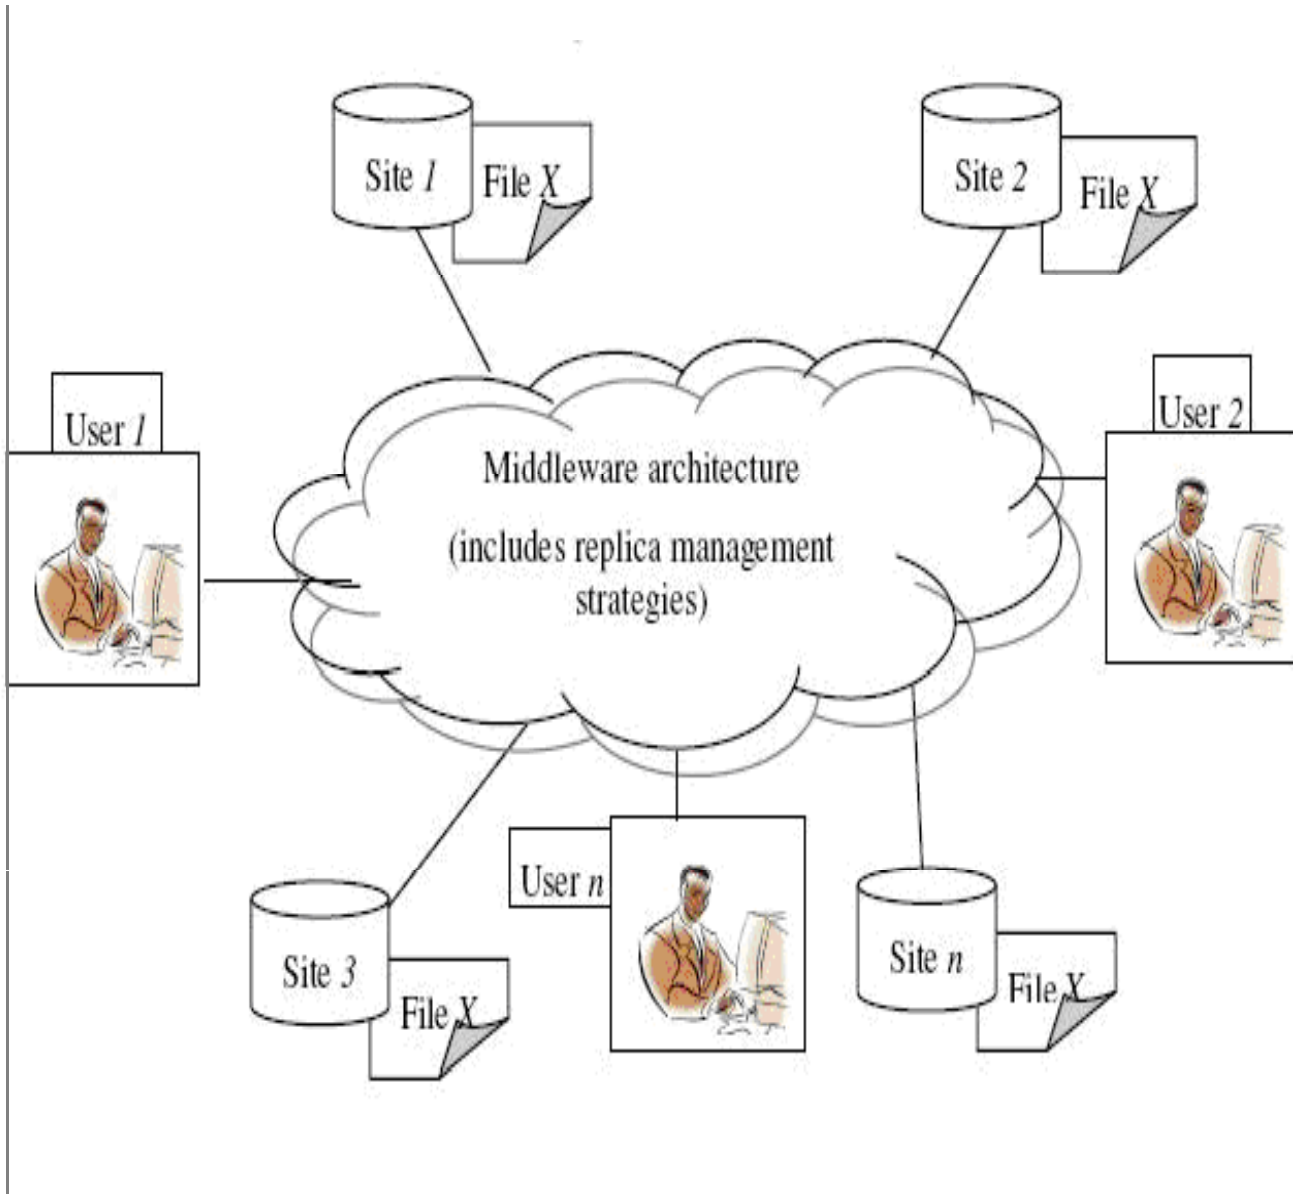


Distributed Systems


Distributed DBMS

Peer-To-Peer (P2P) Systems

Data Grids / Grid Computing



Distributed DBMS



Data Replication
and
Distributed Systems

ROWA and ROWA-Available

The most simple replica control protocol is the Read-One-Write-All (ROWA) protocol. In ROWA protocol, a transaction requests to read an item and the system fetches the value from the most convenient location.

Quorum Based

An interesting proposal to update only a subset of replicas and still not compromise with correctness and consistency is based on quorums


Commit Protocol in Distributed DBMS

Two-phase commit (2PC) protocol is the most widely accepted commit protocol in distributed DBMS environment that helps in achieving replica synchronisation

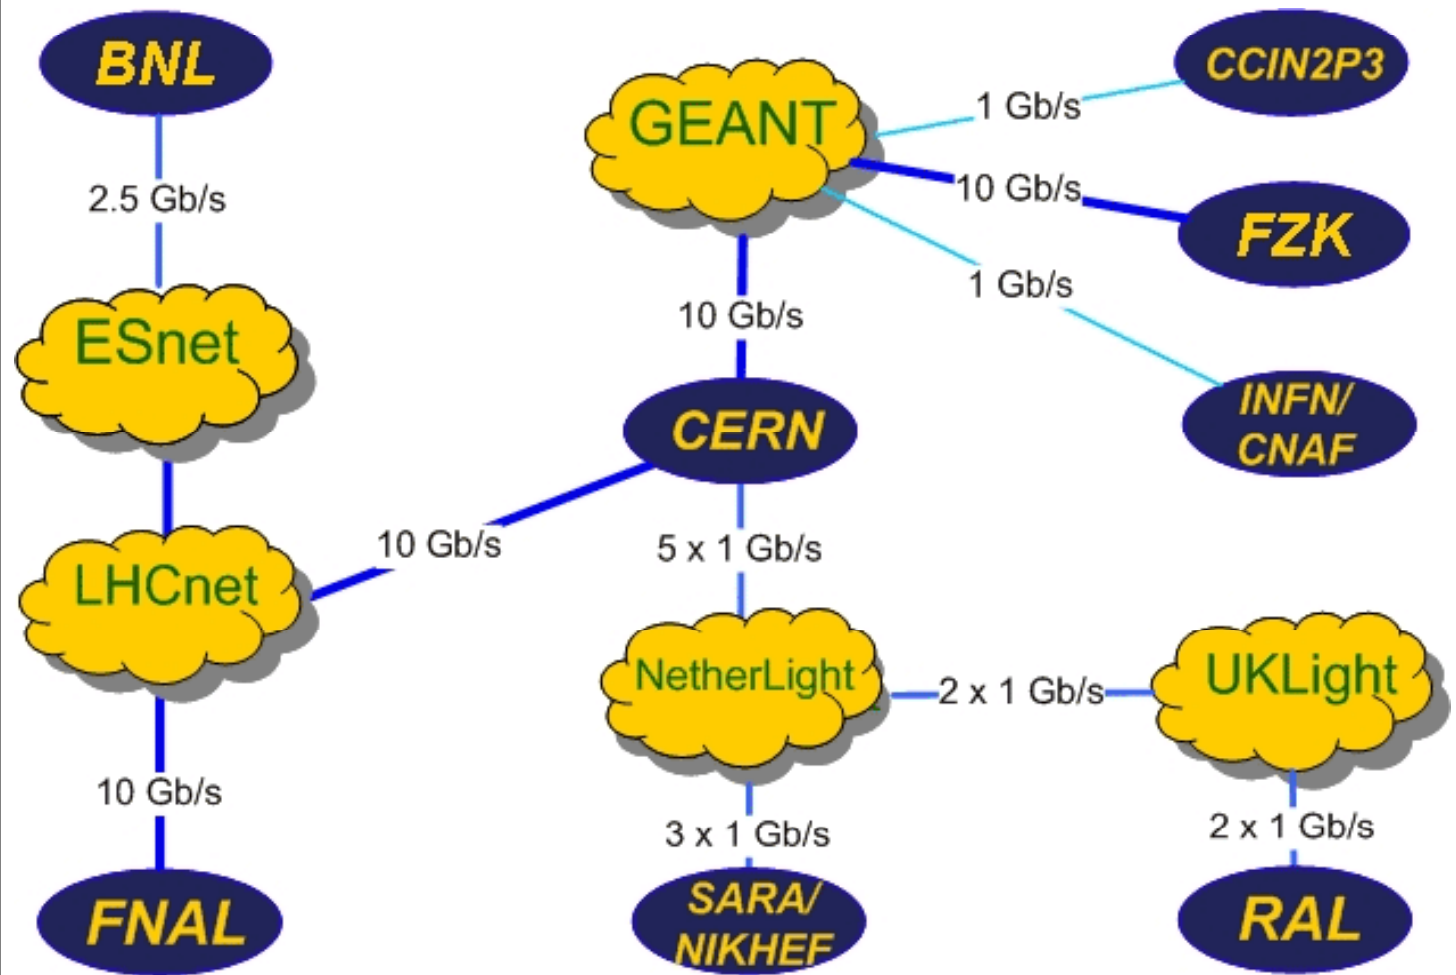
Distributed Data Storage Systems and Replication Strategies

A history of strategies used to achieve replication in Distributed Data Storage Systems.

Data Grids / Grid Computing



Data Replication
and
Distributed Systems



Data Replication and Security



Distributed Systems and Security

Distributed systems make security assurance particularly difficult, as these systems naturally cross administrative and trust boundaries

Encryption

As the modest precaution, encryption of the transfer wire should be provided.
(i.e. SSL, IPSEC...)



Security of Backup Data



Encryption Key Management

But while encryption is proliferating, the means for tracking and managing the keys that make encryption schemes workable has not kept up.

Backups

What if someone were to break into your office and not only steal your computers but also your external drives

Host-Based vs. Appliance-Based Tape Encryption

Remote And Branch Offices

Any company with remote offices or branch offices requires solutions that can support corporate data protection

Stay Up-to-date

Challenges in Data Replication



- Read-only queries
- Update transactions
- Managing mobile clients

Data Consistency:

High precision applications may require strict consistency (i.e 1SR) of the updates made by transactions.

Downtime during new replica creation:

If strict data consistency is to be maintained, performance is severely affected if a new replica is to be created.

Maintenance overhead:

If the files are replicated at more than one sites, it occupies storage space and it has to be administered.

Lower write performance:

Performance of write operations can be dramatically lower in applications requiring high updates in replicated environment, because the transaction may need to update multiple copies.

Today's Technology in Data Replication



Distributed Cache Products:

DCPs, so called *in-memory data grids*, *data fabrics*.

Virtualization:

Virtualization is quickly becoming one of the most highly utilized technologies in business today.

Mobile Environments:

Mobile environments are characterized by a frequent change in their resources.

Definition of Data Replication

Data Replication Options

Common Usages

Data Replication Architectures

Data Replication and Distributed Systems

Distributed DBMS

ROWA and ROWA-Available

Quorum Based

Commit Protocol in Distributed DBMS

Distributed Data Storage Systems and Replication Strategies

Peer-To-Peer (P2P) Systems

Data Grids / Grid Computing

Data Replication and Security

Security of Backup Data

Encryption Key Management

Backups

Host-Based vs. Appliance-Based Tape Encryption

Remote And Branch Offices

Stay Up-to-date

Today's Technology in Data Replication

