



When DR Is Not Enough

Maintaining High Availability

**PAUL KOUFALIS
GILLES QUERRET**



Why Are We Here?

- DR and HA are NOT the same thing
- We want to AVOID activating Disaster Recovery plan
- Instead, we want to KEEP the application available on the main production servers

Why Are We Here?

- Disaster Recovery is...
 - Bring it up fast/good after an unexpected **catastrophic** event
 - Typically at another physical location
- MAJOR PITA
 - Call your wife and tell her not to wait up

Why Are We Here?

- High Availability Means...
 - Keep production available
 - AND up-to-date
 - AT the main site
- Databases
- UBrokers (AppServer, WebSpeed...)
- Etc...

Agenda

- Two distinct discussion tracks:
 - DB HA
 - AppServer HA

Introduction

- OpenEdge has no DB HA solution
- We must work to
 - Minimize possibility of DB down time
 - Minimize effect of DB down time
- Can we fix quickly w/out going to DR?

Scenario 1- Running Out of Disk Space

- We've all seen it
 - One FS is full
 - The others have plenty of space

- Often due to uncontrolled BI growth
 - Cannot truncate BI online
 - Cannot truncate BI with OE Replication

Solution

- prostrict addonline
 - Watch out for permissions issues
- Note:
 - I advocate ONE FS for DB/BI/AI
 - This is contrary to many respected DBAs
 - Perf gain is questionable for normal environments

Scenario 2 – Index Corruption

```
SYSTEM ERROR: read wrong dbkey at offset <offset>  
in file <file> found <dbkey>, expected <dbkey>,  
retrying. (9445)
```

- Could be a hardware, O.S. or software issue
- Not as frequent as it used to be

Solution

- Use proutil db -C idxfix online -Bp 64

```
_proutil sports -C idxfix
```

- There will be some impact on prod
 - Depends on the table...

Scenario 3 – Lock Table Overflow

Lock table overflow, increase -L on server (915)

- Of course you need to find and fix the offending process
- BUT...users are screaming NOW
- Let the users continue THEN fix the problem

Solution

- `_proutil db -C increaseto`

```
_proutil sports -C increaseto -L 100000
```

- Can also increase `-B`, `-B2`, `-bibufs`, `-aibufs`, `-Mxs` and `-omsiz`
- Like `prostrct addonline`, there could be permission issues if adding `shmem` segments
- Sorry – cannot decrease online

Scenario 4 – Bad C/S Parameters

- -Mi (min users per server)
- -Ma (max users per server)
- -Mpb (max servers per broker)
- -Mn (max servers)
- -n (max connections)

- $-Mpb * -Ma = \text{max concurrent on } -S \text{ port}$

Scenario 4 – continued...

- I have seen this many times
- Some apps take 2 conx/user
 - When you add 100 users...oops...params only adjusted for 100 connections
- Recent customer issue: SQL application bug was causing large number of connections

Solution

- Segregate ALL the login brokers
- Primary broker: only for shmem connections
- TWO secondary brokers
 - One for ABL
 - One for SQL
- Leave plenty of overhead space in -n and -Mn
 - Licensing and -n are NOT related

Solution – continued...

- You cannot increase `-Mpb` nor `-Ma`
- BUT you can add a 3rd (or 4th) secondary login broker
 - Move group of connections to new broker
 - Ex.: AppServer connections can be moved easily
 - Liberate slots in main broker

Scenario 5 – Imminent Server Death

- Hard drive death rattle (TOC TOC TOC)?
- Blocked fan?
- Weird console error messages?
- Running Windows? (sorry – had to say it)

- There is no RAC-like functionality
 - The DB is going to have to go down
 - Make it quick and almost painless

Solutions

1. ProCluster

- Enable DB as a cluster resource
- Assumes shared SAN containing DB disks

2. Poor-man's clustering

- Mount DB drives on other server
- AIX.: `exportvg + importvg`

3. OE Replication local target

- OER supports two targets
- One to DR site and one local
- Local copy typically used for reporting

- You can use it for quick DR
- A small DNS change
- Restart all clients (ouch! – I know...)

4. Poor man's replication – AKA AI files

- Can keep a local copy up-to-date
- Transfer and apply AI files

- Same as OER solution: two targets
- A small DNS change
- Restart all clients

Scenario 6: Server Maintenance

- Almost same as previous scenario
- Instead of staying down during maintenance, just move PROD DB for a few hours
- Fast, easy, low impact



Part 2 – NameServer & AppServer high availability



AppServer Operating Mode

- Who's using AppServer ?
- We'll only cover Stateless and Statefree appservers
- State-reset and state-aware are deprecated

NameServer Overview

- The NameServer is a Java process managed by AdminService
- A NameServer register AppServer brokers by application service name
- When a client ask for a service name, NameServer returns host and port number of the AppServer broker
- NameServers only use UDP

Stateless Appserver

- Agents are not dedicated to a client
- Client connects to NS
- NS provides address of AppServer(s)
- Client connects to Apsv and invokes request
- Broker passes request to agent
- Agent executes the request, returns result to the broker, which returns result to the client
- Client optionally disconnects from Apsv

Statefree Appserver

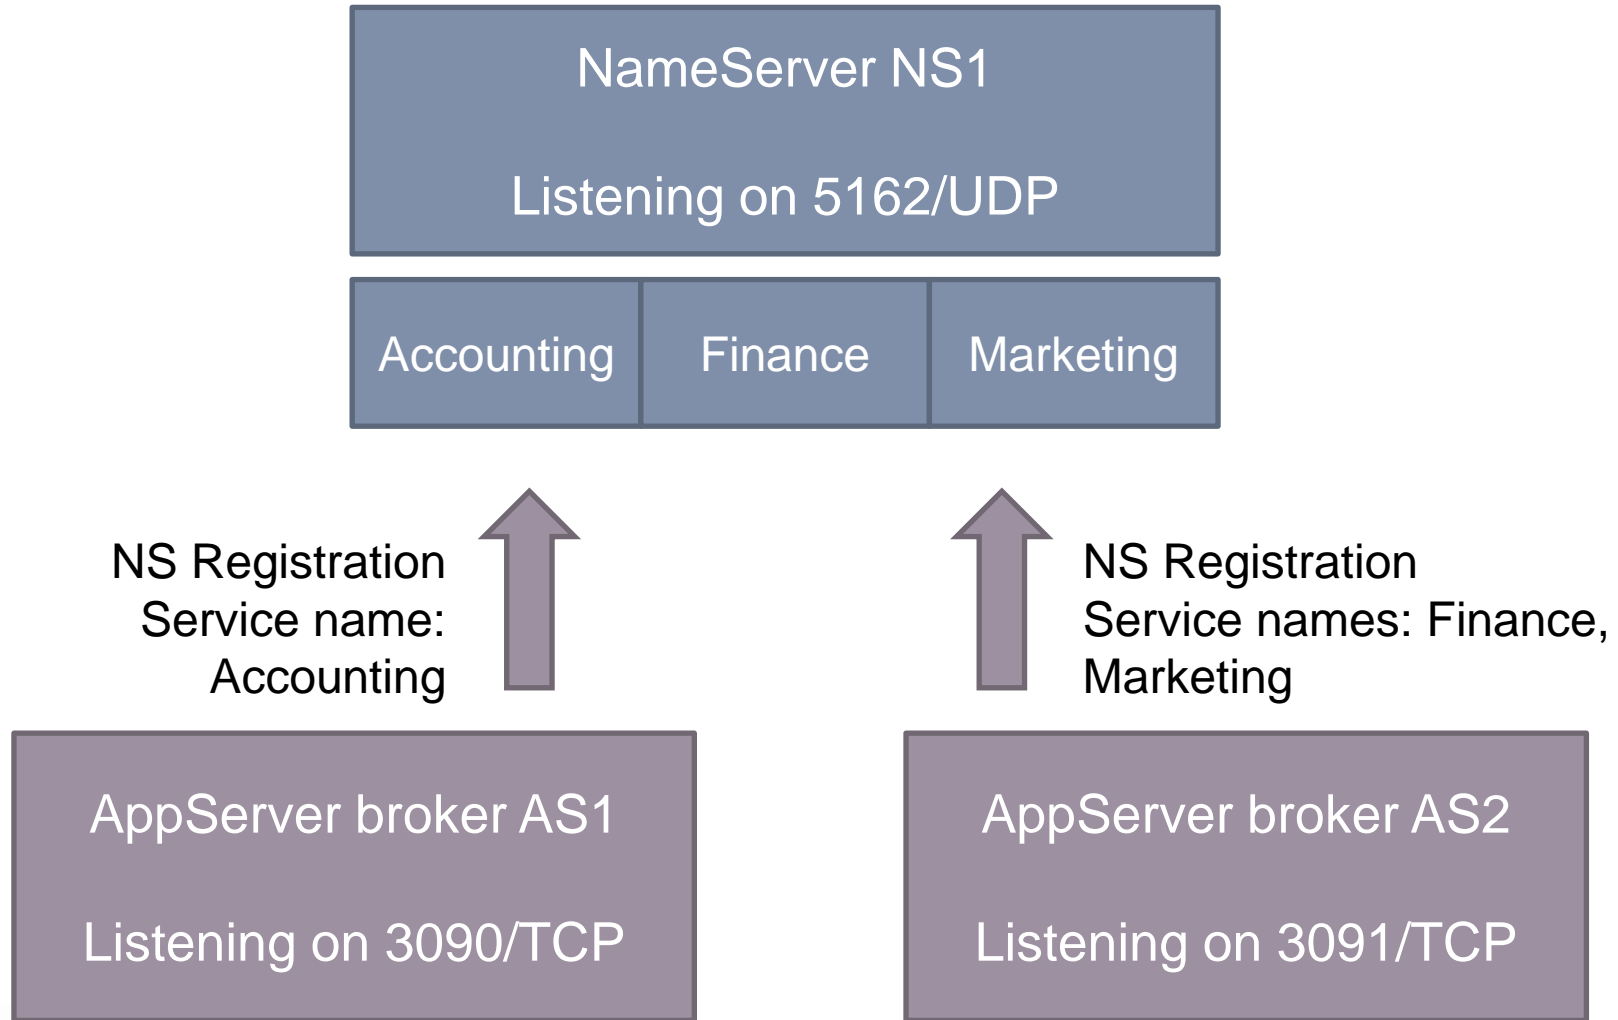
- Same as Stateless EXCEPT:
- NameServer returns a list of AppServers
- Client creates physical connections for each AppServer broker in the list

Stateless or State-free ?

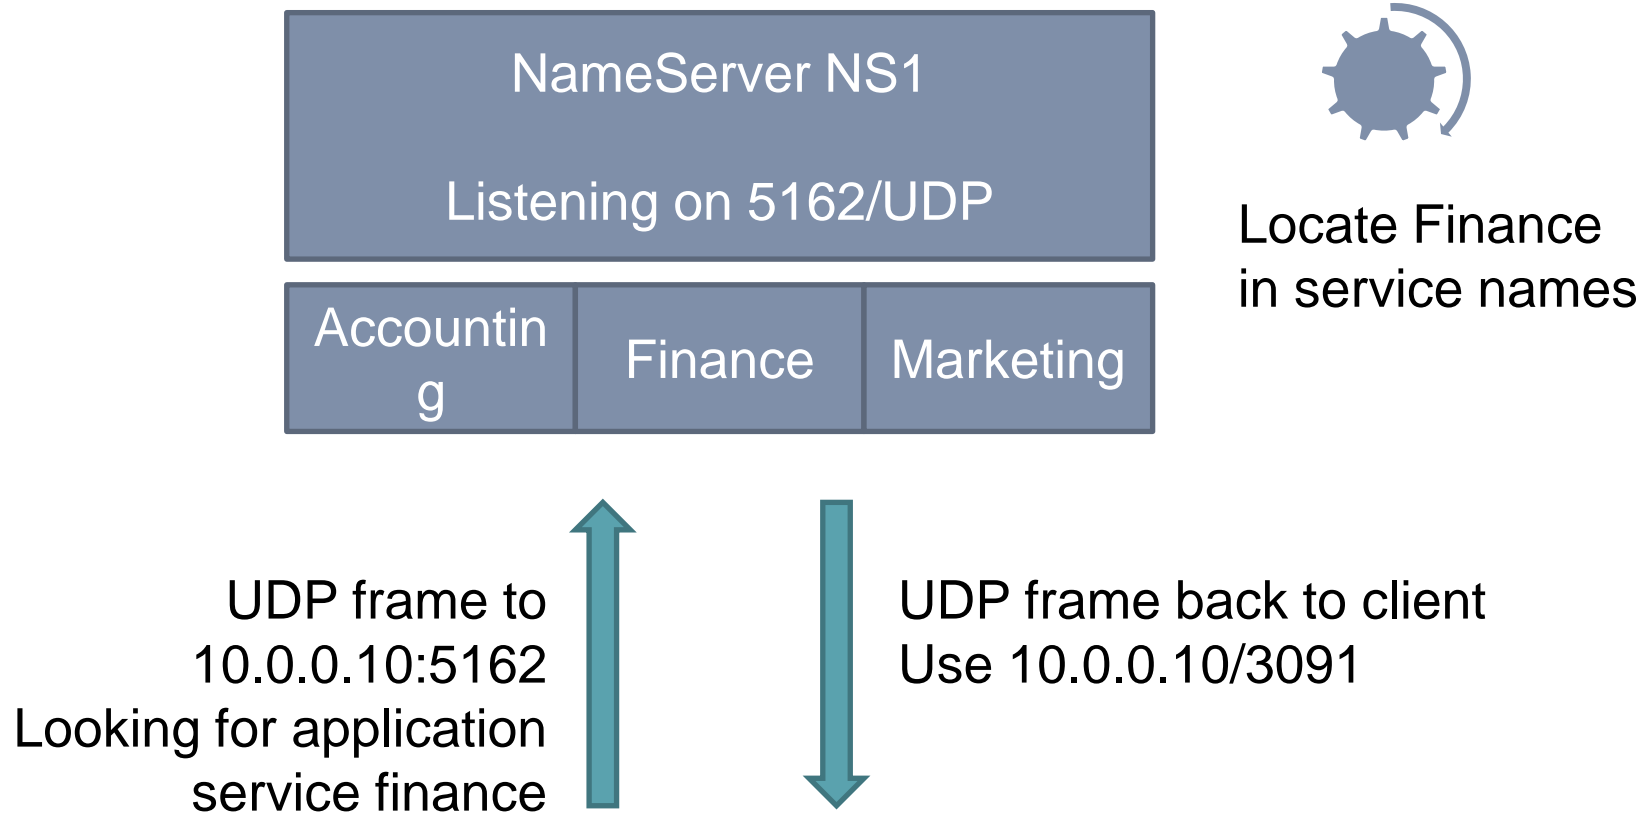
- Not a DBA / SysAdmin decision
- WebServices require State-free
- Context is easier to maintain in Stateless

- Connect / Disconnect required in order to refresh brokers list
 - Responsibility of application
 - Otherwise possible to attempt connection to resource that is no longer available

Simple Setup



NameServer Query



NameServer Neighborhood

- A NameServer can have neighbors
- When NS doesn't recognize an application service name, it forwards the query to its neighbor(s)
- Neighbor is configured as a « Remote NameServer »

Remote NameServer

- Only a place holder pointing to another NS
- Defines host and port of remote NS
- There is no local running process
- Both registrations and queries can be done on remote nameservers

Use Case

- Proxy nameserver on low-end server
- Dispatch requests to real servers
- Move services to different locations without changing the main server URL

NameServer Load Balancing

- Multiple AppServers can register under a single service name
- A priority weight can also be given to a broker and NS will dispatch requests accordingly
- Requires NS Load Balancing license

NS Load Balancing Use Case

- Remote NS required
- Use case during server upgrades
- Either shut down or set priority weight to 0, then upgrade server
- Restart AS or set priority weight to any non-zero value
- Repeat same steps for every server

UDP Broadcast

- NameServers have to listen on same port number on a given subnet
- UDP broadcast packet sent on this subnet
- First packet received will be used, others will be discarded

NS Broadcast Use Case

- High load expected on the server
- Any number of appservers can be added, hosted on different servers
- Any new appserver will register on any Nameserver in the subnet
- Nameservers will automatically dispatch the load on appservers
- Throwing new servers will increase capacity
- Deleting servers can be done when peak expires

When Can You Expect Peaks ?

- All the techniques apply to OpenClient
- Using WebServices Adapter / Appserver Internet Adapter / REST adapter ?
- Releasing version 6 and 6+ of your key product ? Throw new appservers to handle the load, and delete them after having sold 10 million in 3 days !

Remember

- Best to keep PROD available
 - Redundant local OE resources in addition to emergency resources

- Avoid switching to DR site
 - Except for catastrophic emergency

Progresswiz Consulting

- Based in Montréal, Québec, Canada
- Providing technical consulting in Progress[®], UNIX, Windows, MFG/PRO and more
- Specialized in
 - Availability and business continuity planning
 - Performance tuning
 - Security

- www.progresswiz.com

Riverside Software

- Based in Lyon, France
- Technical expertise in OpenEdge and Java environments
- Working on continuous integration process, automated deployment and source code analysis
- contact@riverside-software.fr

Questions?

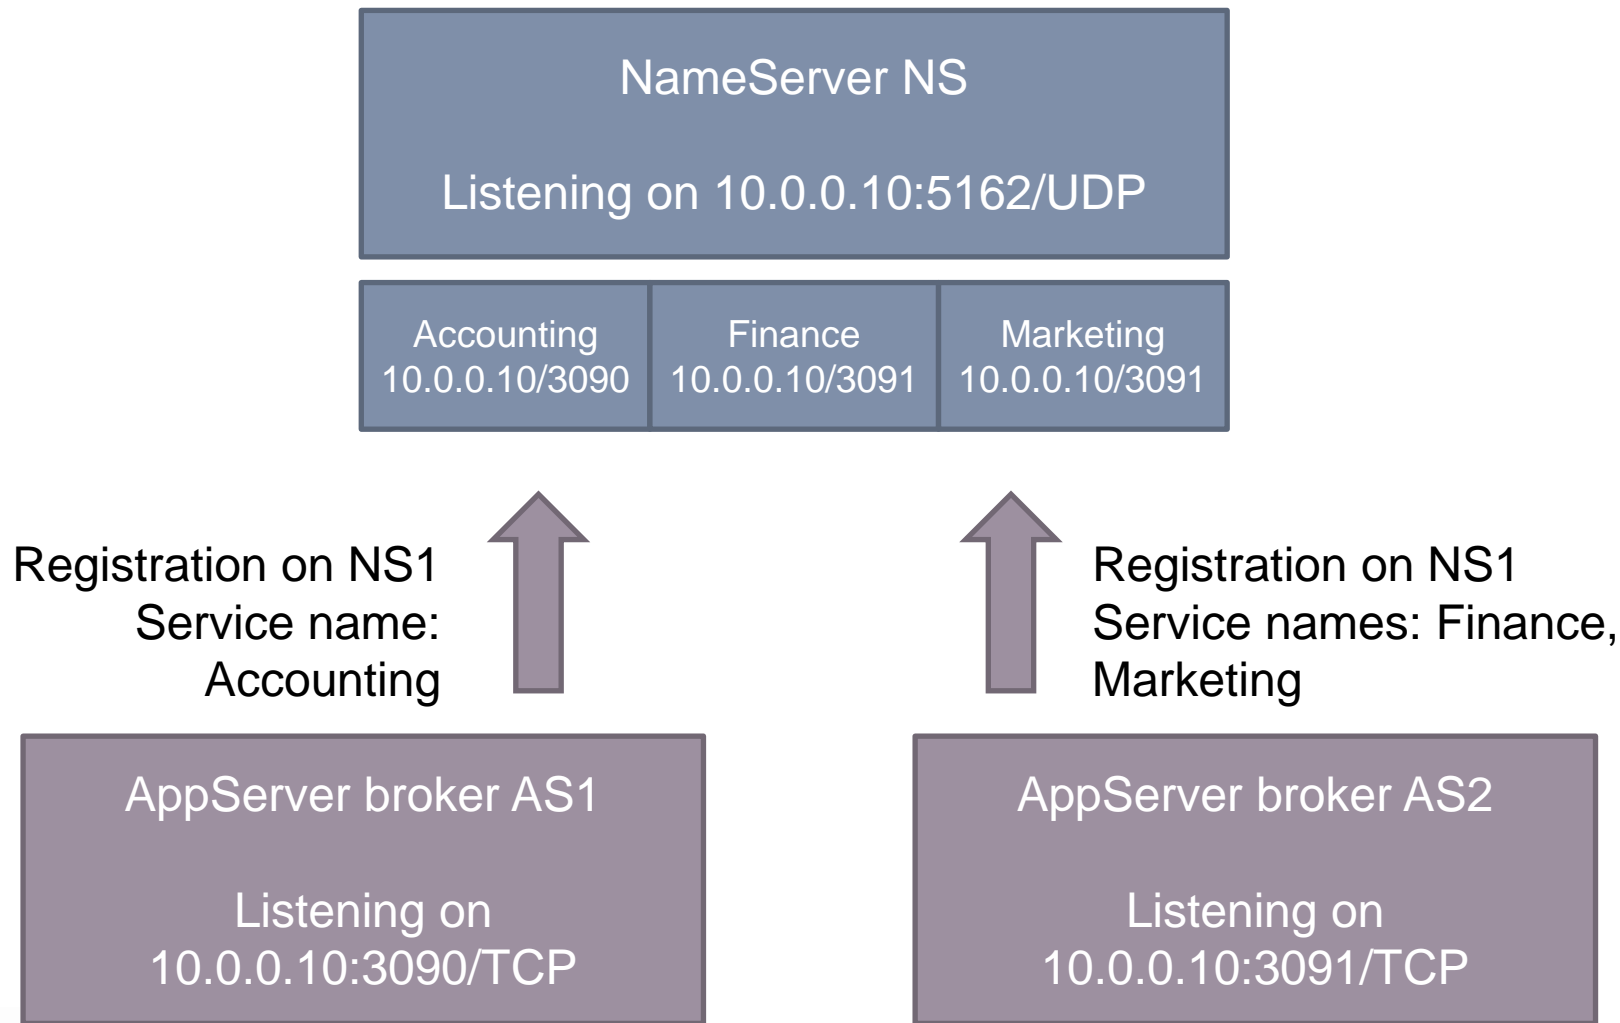


Progress Exchange Resource Portal

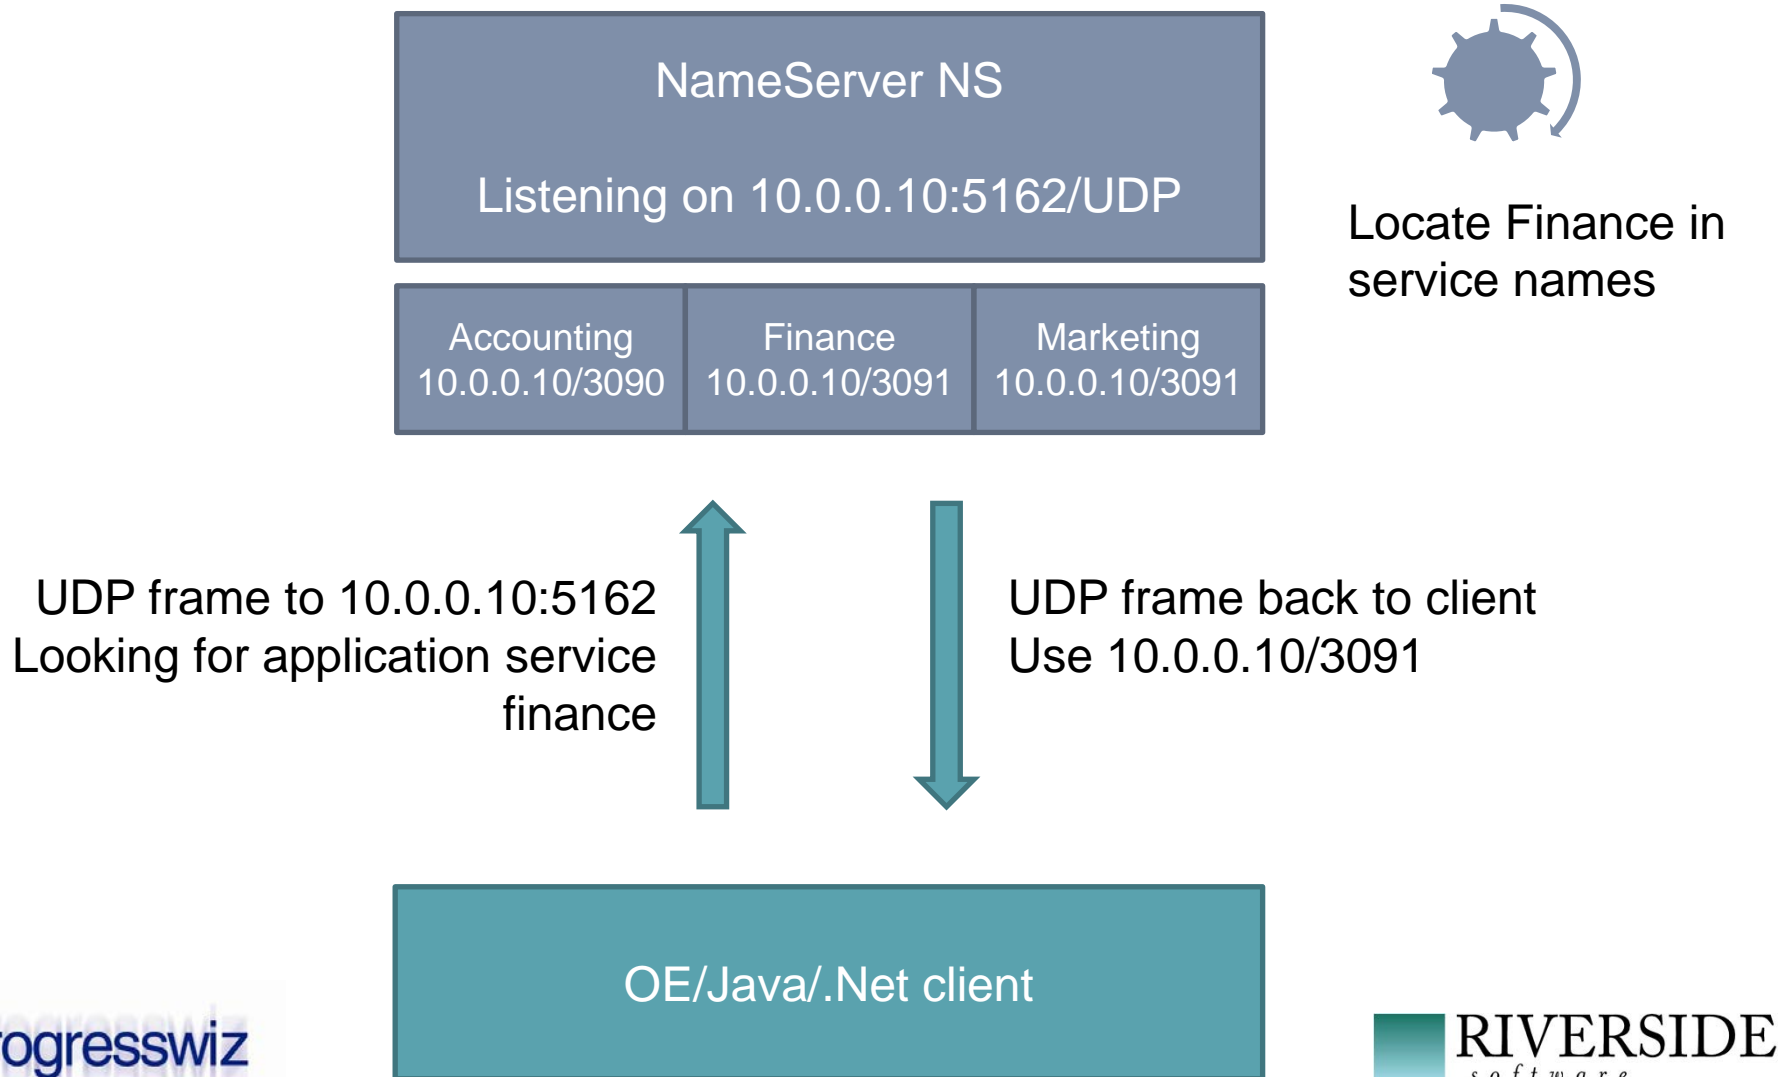
- Get session details & presentation downloads
- Complete a survey
- Access the latest Progress product literature

www.progress.com/exchange2014

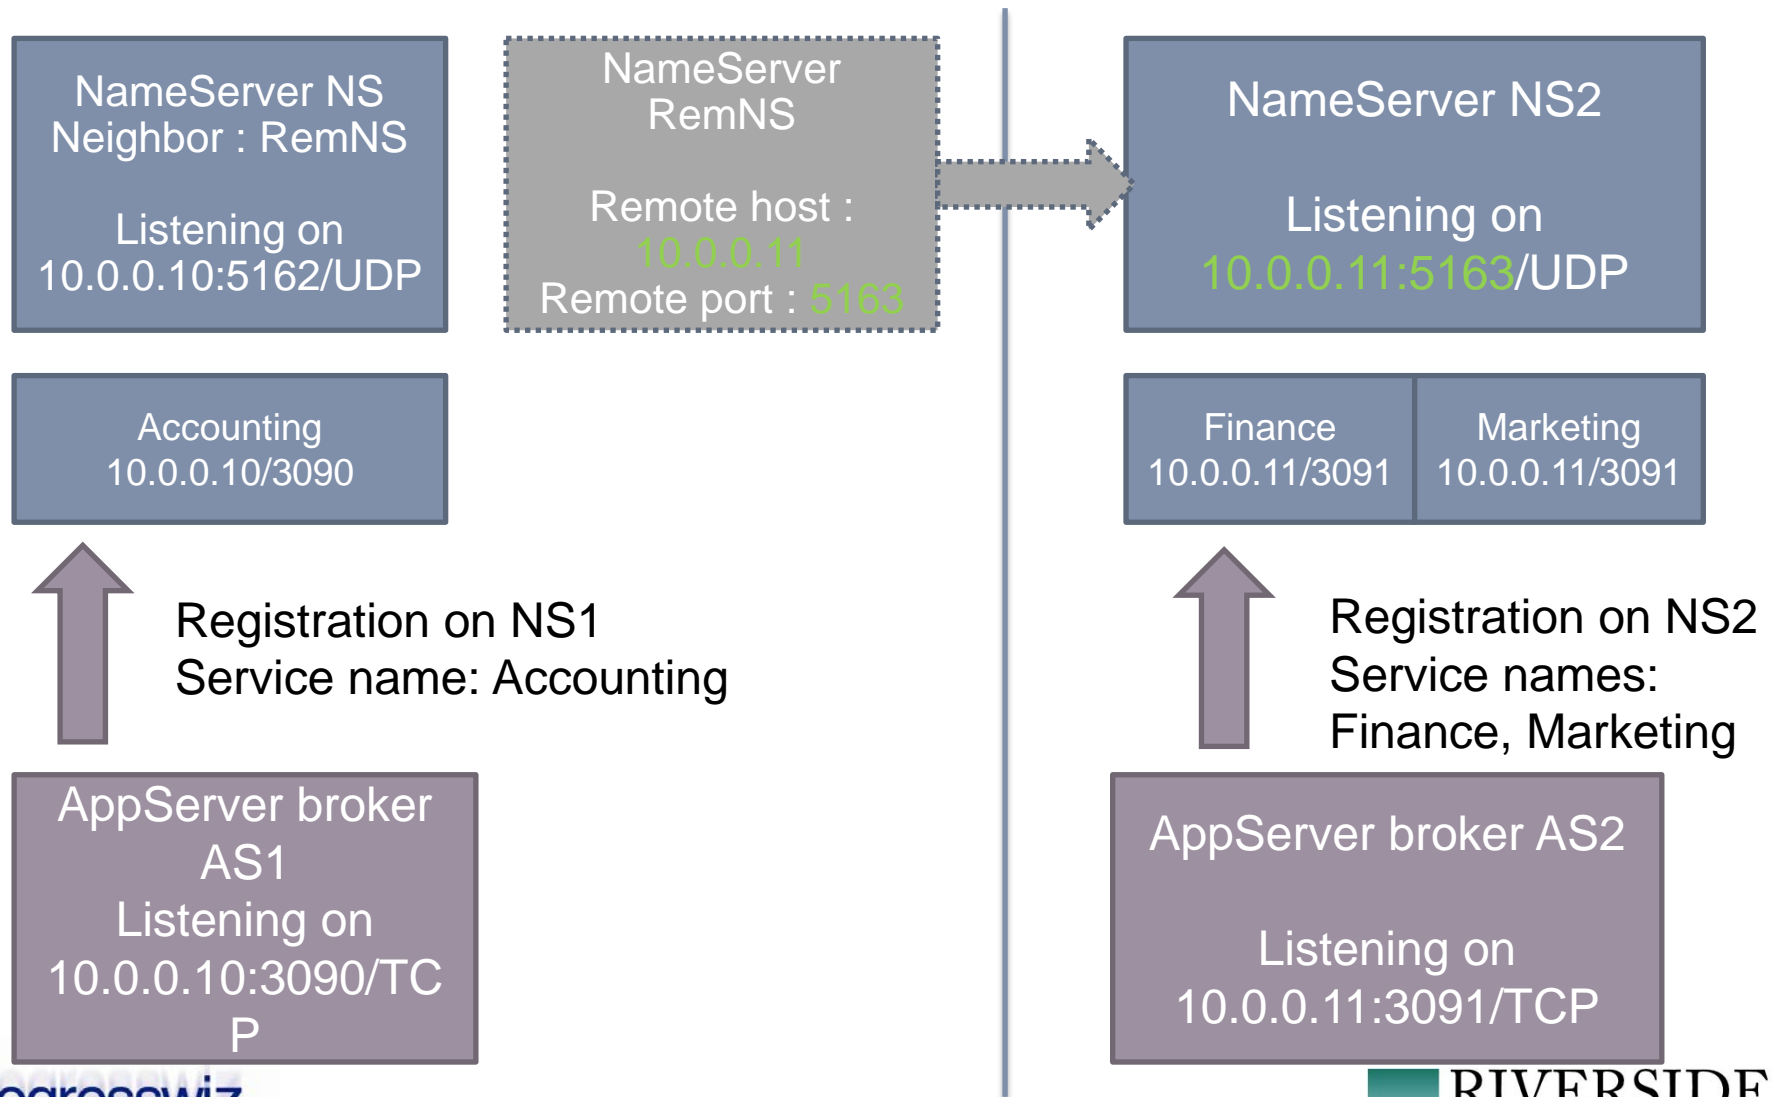
Reference Diagram – Simple Setup



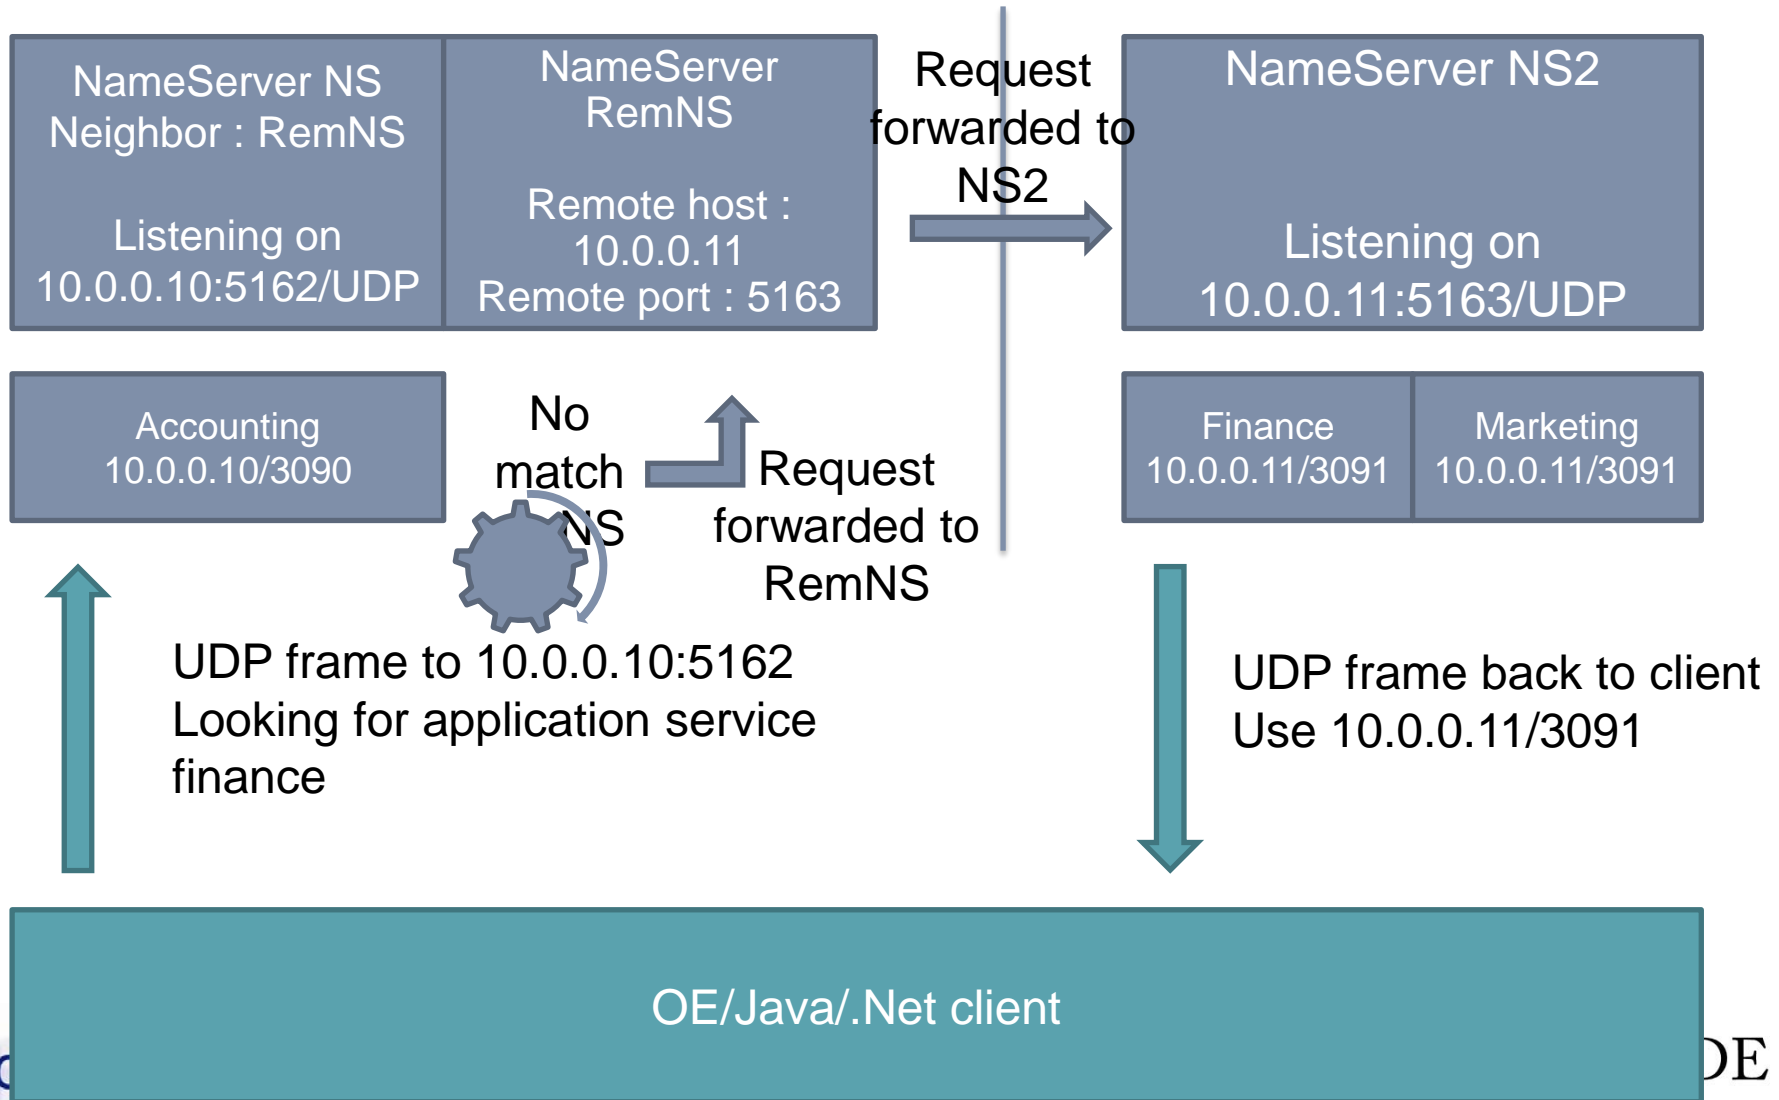
Reference Diagram - NameServer Query





Reference Diagram - Remote NameServer



Reference Diagram - NameServer Query



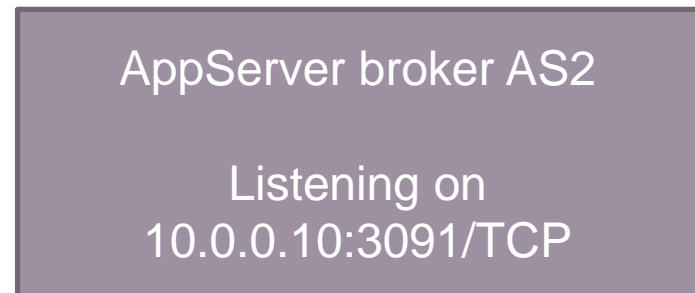
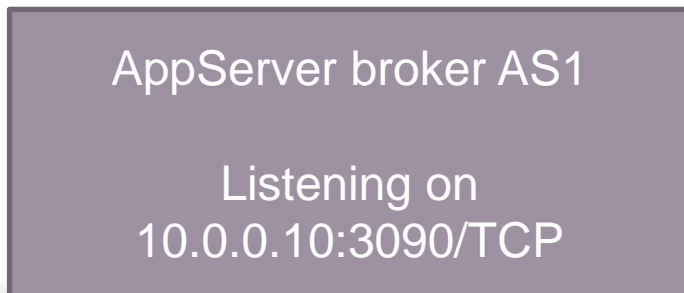
Reference Diagram - NS Load Balancing



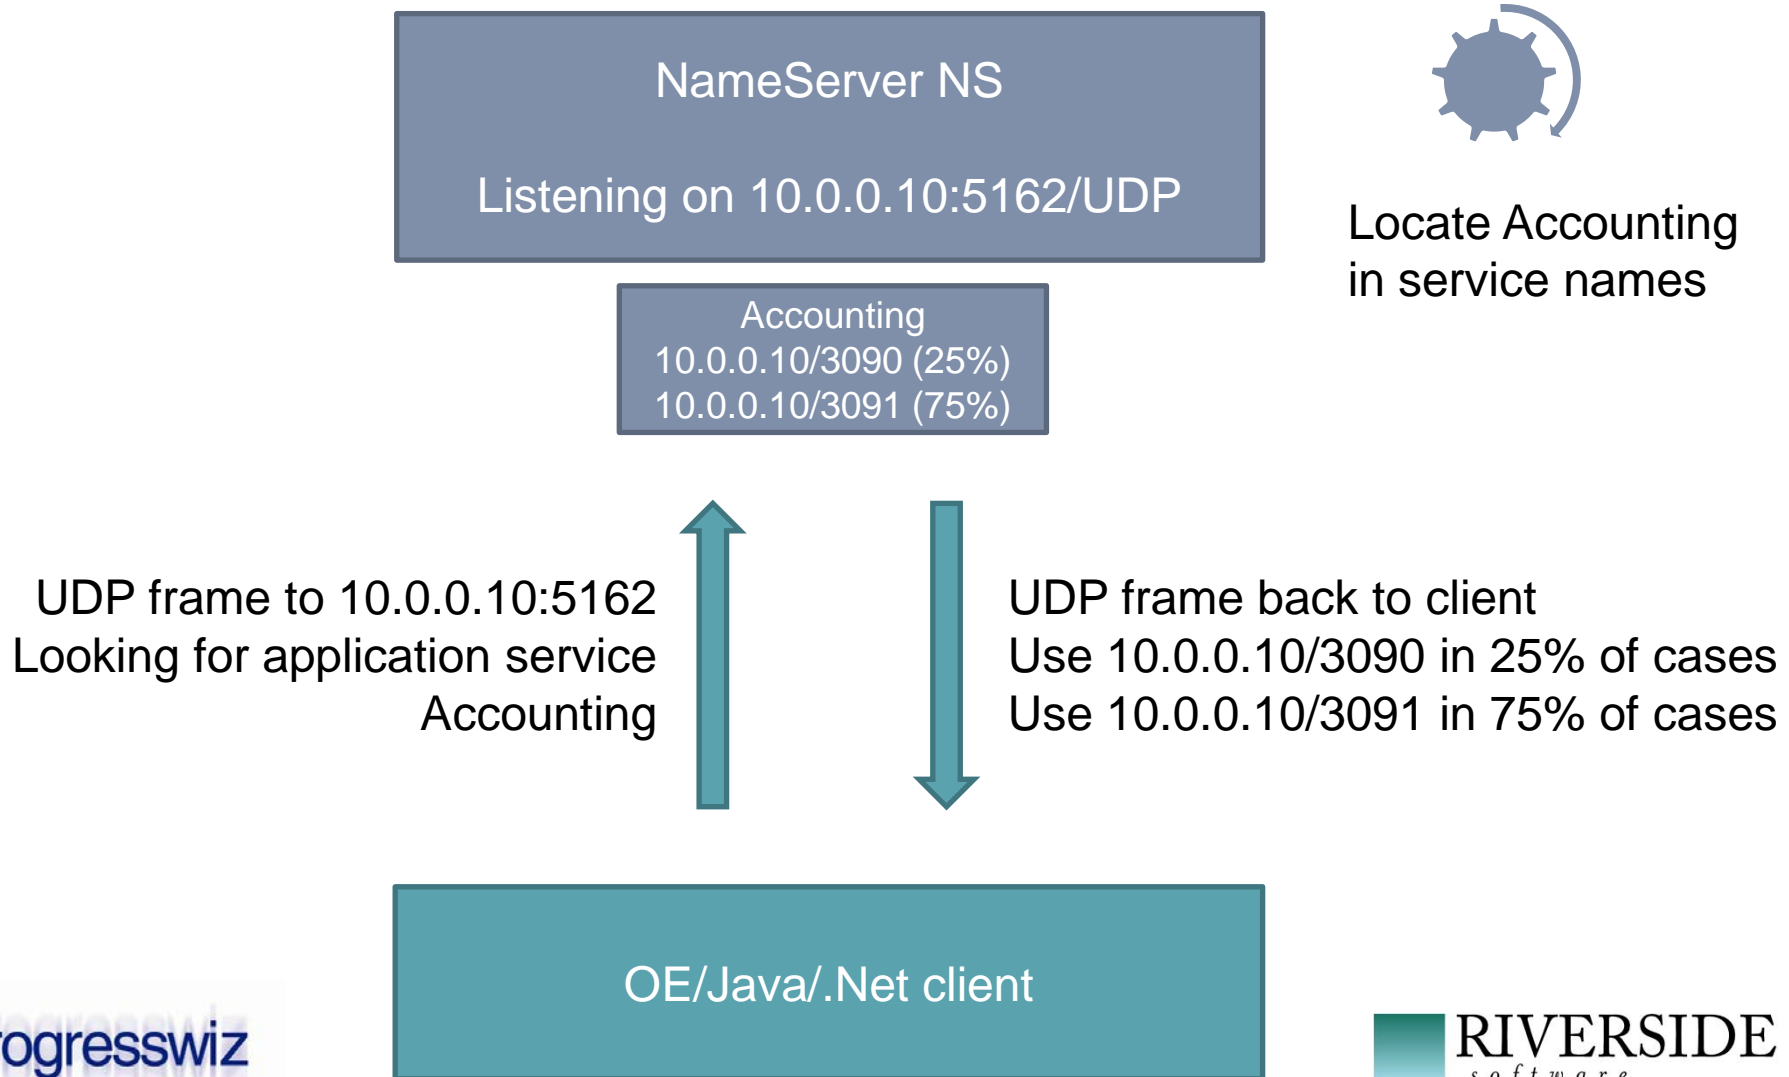
Registration on NS1
Service name: Accounting
Weight: 10



Registration on NS1
Service name: Accounting
Weight: 30



Reference Diagram - NameServer Query



Reference Diagram - NS Broadcast

