•rayNET

RayFlow®

Best Practice Application Lifecycle Workflow

Solution Brief – Rayflow Backup und Disaster Recovery

Backup process

In a perfect world, the RayFlow application and database should be backed up at the same time. If SQL is off-box, then this can easily be achieved; however, if SQL is on-box, then the server backup process (if used) could possibly conflict with the database backup process, and vice versa, in which case both processes should be run separately.

Before the application or database is backed up, either the web site that the RayFlow application resides in or the RayFlow application itself should have IIS isolation enabled. This is achieved via the "IP Address and Domain Restrictions" IIS feature on a Windows Server 2012 R2 system:

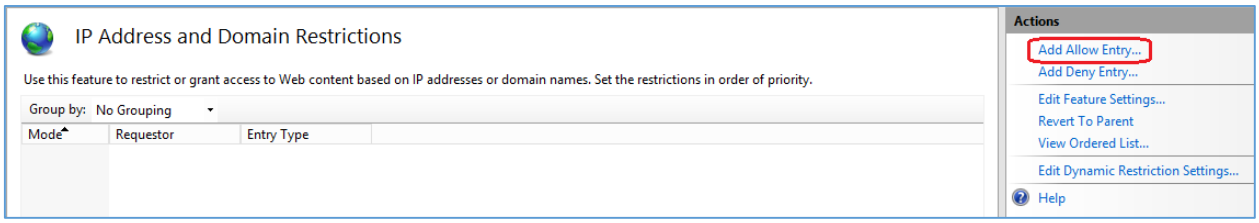

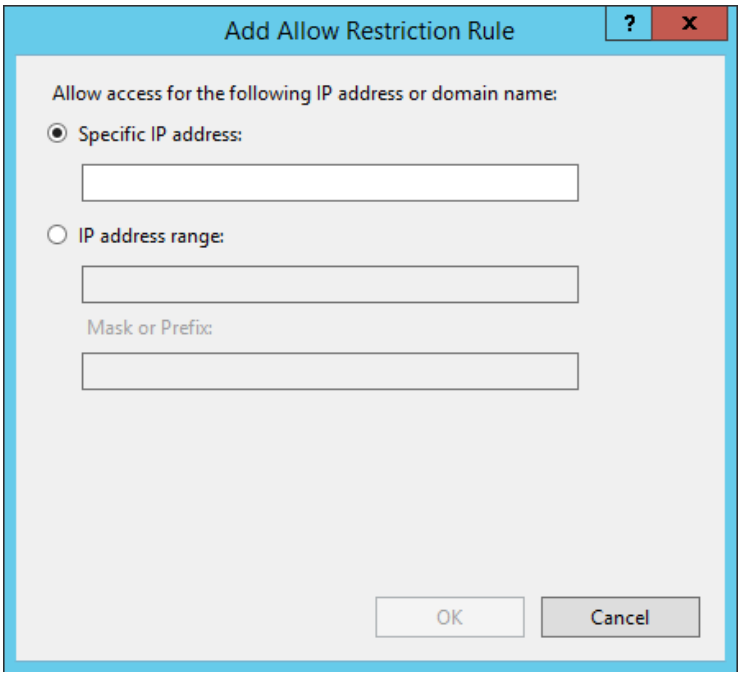

Add the localhost IP Address (127.0.0.1) as well as the servers actual IP Address.

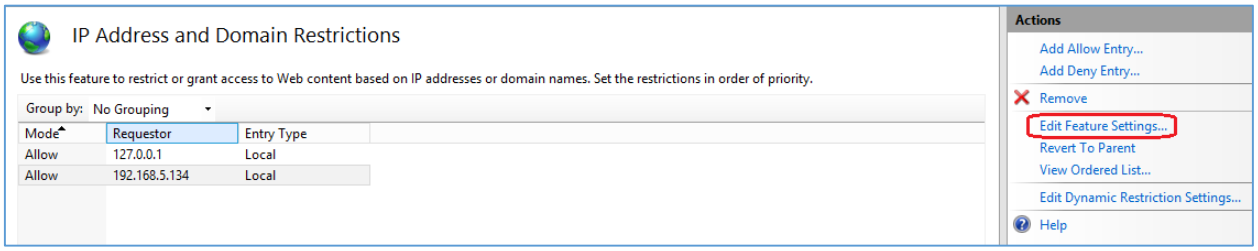

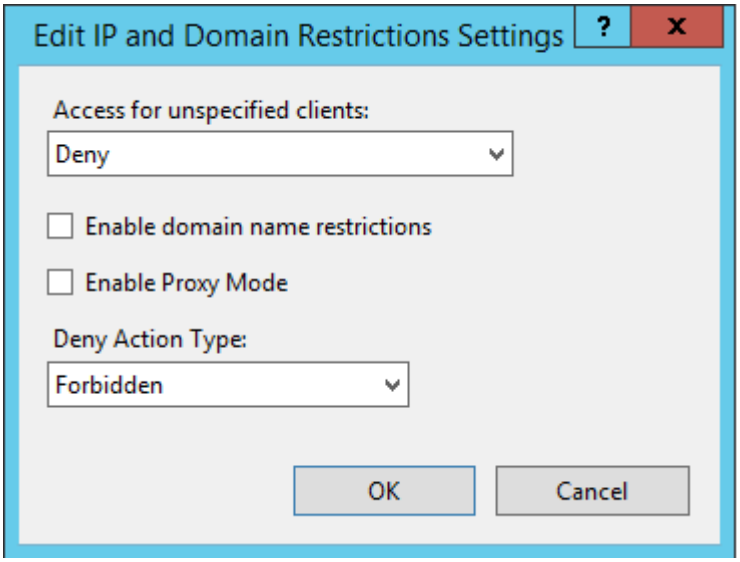

Deny all other connection requests.

Filesystem

The application file system can either be backed up as specific items, or they can be included as part of a complete server backup process. If a company's DR policy utilises an automated server rebuild process, then backing up the application as specific items would be used.

The following is an example of the RayFlow application filesystem location:

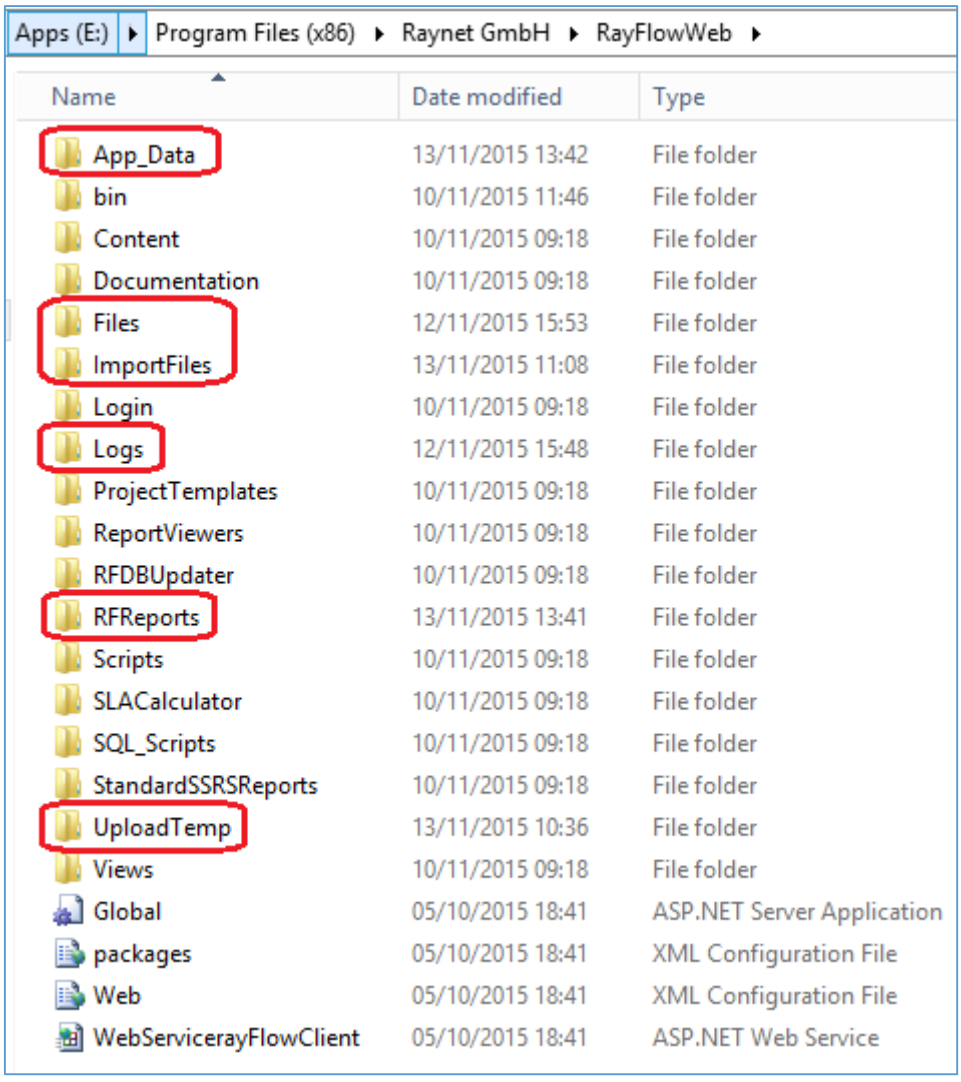

The application folders selected above can be redirected via virtual directories if required. If one or more of those locations have been redirected to a different location on the network (File Server, NAS, SAN, etc.), then that location obviously needs to be backed up too.

Database

The size of the database will determine its backup strategy. Starting off, you may wish to perform daily full backups as RayFlow's initial database size is around 10MB; however, company policy may dictate that once a database has reached a certain size (100s of MB or N GBs), the backup strategy may be changed to "full + incremental" or "full + differential".

Once the application and database backups have been completed, IIS isolation should be disabled:

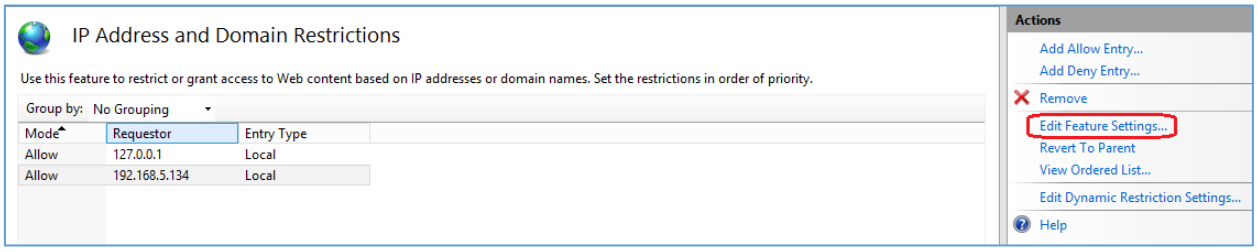

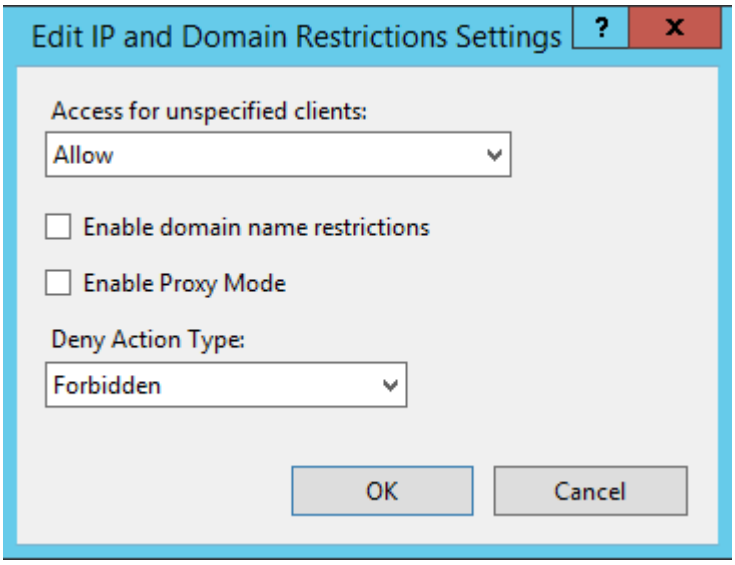

The two specific allow entries can either remain or be removed.

Disaster Recovery

For the database, utilising a SQL Cluster is the best method.

If the application server is a VM, then a snapshot could be taken on a regular basis. If taken daily, then this could also be used as a backup method too instead of a complete server backup process.

Utilising RAID 1 to create a mirrored application drive is an efficient DR method. If the database has to reside on its own SQL server instead of sharing a SQL cluster, then it too could use a RAID 1 solution, unless there are enough HDD's for RAID 5 or RAID 10.

Some server backup products have the ability to create a VM from a backup, which is another DR solution.